

Using Liferay DXP 7.2

A Complete Guide

THE LIFERAY DOCUMENTATION TEAM

Richard Sezov, Jr.

Jim Hinkey

Stephen Kostas

Jesse Rao

Cody Hoag

Nicholas Gaskill

Michael Williams

Liferay Press

Using Liferay DXP 7.2
by The Liferay Documentation Team
Copyright ©2020 by Liferay, Inc.

This work is offered under the following license:

Creative Commons Attribution-Share Alike Unported



You are free:

1. to share—to copy, distribute, and transmit the work
2. to remix—to adapt the work

Under the following conditions:

1. Attribution. You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).
2. Share Alike. If you alter, transform, or build upon this work, you may distribute the resulting work only under the same, similar or a compatible license.

The full version of this license is here:

<http://creativecommons.org/licenses/by-sa/3.0>

This book was created out of material from the Liferay Docs repository. Where the content of this book and the repository differ, the site is more up to date.

CONTENTS

| | |
|--|--------------|
| Contents | i |
| Preface | xliii |
| Conventions | xliii |
| Publisher Notes | xliv |
| | |
| I Using Liferay DXP | |
| 1 The Liferay Distinction | 1 |
| 2 What's New in Liferay 7.2! | 3 |
| 2.1 Key New Features | 3 |
| 2.2 Experience Creation | 3 |
| 2.3 Personalization | 6 |
| 2.4 Bulk Management | 8 |
| 2.5 Business Operations | 8 |
| 2.6 Headless Capabilities | 9 |
| 2.7 Developer Experience | 11 |
| 2.8 What's Next | 11 |
| 3 Setting Up | 13 |
| 4 Where Configuration Happens | 15 |
| 5 Configuration Scope | 17 |
| 6 System Wide Settings | 19 |
| 7 System Settings | 21 |
| 7.1 Editing System Configurations | 21 |
| 7.2 Configuration Scope | 23 |
| 7.3 Exporting and Importing Configurations | 25 |
| 8 Understanding System Configuration Files | 27 |

| | | |
|-----------|---|-----------|
| 9 | Creating Configuration Files | 29 |
| 9.1 | Key/Value Syntax | 29 |
| 9.2 | Multi-Value Settings | 30 |
| 9.3 | Escaping Characters | 31 |
| 9.4 | Typed Values | 31 |
| 9.5 | Deploying a Configuration File | 31 |
| 9.6 | Configuration Files and Clustering | 32 |
| 10 | Factory Configurations | 33 |
| 11 | Server Administration | 37 |
| 12 | Server Administration: Resources | 39 |
| 13 | Server Administration: External Services | 41 |
| 13.1 | ImageMagick Configuration | 42 |
| 13.2 | Xuggler Configuration | 42 |
| 14 | Setting Up a Virtual Instance | 45 |
| 15 | Virtual Instances | 47 |
| 16 | Configuring Virtual Instances | 49 |
| 17 | Email Instance Settings | 51 |
| 17.1 | Account Created Notification | 51 |
| 17.2 | Email Sender | 52 |
| 17.3 | Email Verification Notification | 52 |
| 17.4 | Mail Host Names | 52 |
| 17.5 | Password Changed Notification | 53 |
| 17.6 | Password Reset Notification | 53 |
| 18 | Instance Configuration Instance Settings | 55 |
| 18.1 | Appearance | 55 |
| 18.2 | Contact Information | 56 |
| 18.3 | General | 56 |
| 18.4 | Terms of Use | 57 |
| 19 | User Authentication Instance Settings | 59 |
| 19.1 | General | 59 |
| 19.2 | Reserved Credentials | 61 |
| 20 | Users Instance Settings | 63 |
| 20.1 | Personal Menu | 63 |
| 20.2 | Default User Associations | 63 |
| 20.3 | Fields | 64 |
| 21 | More Platform Section Instance Settings | 65 |
| 21.1 | Infrastructure | 65 |
| 21.2 | Localization | 65 |

| | |
|---|------------|
| 21.3 Analytics | 66 |
| 21.4 Third Party | 66 |
| 22 Using Liferay’s Script Engine | 67 |
| 23 Invoking Liferay Services From Scripts | 69 |
| 23.1 Related Topics | 70 |
| 24 Running Scripts From the Script Console | 71 |
| 24.1 Running the Sample Script | 71 |
| 24.2 Predefined Variables | 72 |
| 24.3 Tips | 73 |
| 25 Leveraging the Script Engine in Workflow | 75 |
| 25.1 Injected Variables | 75 |
| 25.2 Variables that are Always Available | 75 |
| 25.3 Variables Injected into Task Nodes | 76 |
| 25.4 Scripting Examples | 76 |
| 25.5 Calling OSGi Services | 78 |
| 25.6 Related Topics | 79 |
| 26 Script Examples | 81 |
| 26.1 Example 1: Presenting New Terms of Use to Users | 81 |
| 26.2 Example 2: Embedding HTML Markup in Script Outputs | 82 |
| 26.3 Example 3: Show Exceptions in the Script Console | 83 |
| 26.4 Example 4: Implement a Preview Mode | 84 |
| 26.5 Example 5: Plan a File Output for Long-Running Scripts | 84 |
| 26.6 Related Topics | 85 |
| 27 Custom Fields | 87 |
| 27.1 Adding Custom Fields | 88 |
| 27.2 Editing a Custom Field | 89 |
| 28 Managing Users | 93 |
| 29 Users and Organizations | 95 |
| 29.1 What are Users? | 95 |
| 30 Adding, Editing, and Deleting Users | 97 |
| 30.1 Adding Users | 97 |
| 30.2 Editing Users | 98 |
| 30.3 Deleting Users | 100 |
| 31 User Management: Additional Topics | 103 |
| 31.1 Password Resets | 103 |
| 31.2 Adding an Administrative User | 103 |
| 31.3 Gender | 104 |
| 31.4 User Profile Pictures | 104 |
| 31.5 Numeric Screen Names | 106 |

| | |
|--|------------|
| 32 Organizations | 109 |
| 32.1 When to Use Organizations | 109 |
| 32.2 What can Organization Administrators Do? | 110 |
| 32.3 Organization Roles and Permissions | 112 |
| 32.4 Organization Sites | 112 |
| 33 Managing Organizations | 113 |
| 33.1 Adding Organizations | 113 |
| 33.2 Editing Organizations | 114 |
| 33.3 Organization Types | 114 |
| 34 Roles and Permissions | 117 |
| 34.1 Deleting Asset Containers | 118 |
| 34.2 Default Liferay Roles | 118 |
| 35 Managing Roles | 121 |
| 35.1 Creating Roles | 122 |
| 35.2 Assigning Users to a Role | 122 |
| 36 Defining Role Permissions | 123 |
| 36.1 Delegating Social Activities Configuration | 125 |
| 37 Managing User Data | 127 |
| 37.1 Anonymizing Data | 128 |
| 37.2 Manual Anonymization | 129 |
| 38 Sanitizing User Data | 133 |
| 38.1 The Personal Data Erasure Screen | 134 |
| 38.2 Delete the User | 136 |
| 39 Exporting User Data | 137 |
| 39.1 Exporting and Downloading | 137 |
| 39.2 Examining Exported Data | 139 |
| 40 User Groups | 143 |
| 41 Creating a User Group | 145 |
| 41.1 Assigning Members to a User Group | 145 |
| 42 User Groups and Site Membership | 149 |
| 43 User Group Sites | 151 |
| 43.1 Creating User Group Sites From Site Templates | 151 |
| 43.2 Creating User Group Sites Manually | 152 |
| 43.3 Legacy User Group Sites Behavior | 153 |
| 44 Configuring User Group Permissions | 155 |
| 45 Editing User Groups | 157 |

| | |
|---|------------|
| 46 Password Policies | 159 |
| 46.1 Adding and Configuring Password Policies | 159 |
| 46.2 Assigning Users to a Password Policy | 160 |
| 46.3 Default Policy Properties | 161 |
| 47 Auditing Users | 163 |
| 48 Viewing Audit Events | 165 |
| 48.1 Finding Audit Events | 166 |
| 49 Configuring Audits | 171 |
| 49.1 Reporting Audit Events in Liferay’s Logs and Console | 171 |
| 49.2 Configuring Audit Events for Scheduled Liferay Jobs | 172 |
| 49.3 Enabling or Disabling Audit Events Entirely | 173 |
| 49.4 The End of the Story | 173 |
| 50 Web Experience Management | 175 |
| 51 Authoring Content: Structured and Inline Content | 177 |
| 51.1 Structured Web Content | 177 |
| 51.2 Inline Content | 178 |
| 51.3 What’s Best for Your Use Case? | 178 |
| 52 Building a Site | 181 |
| 53 Site Management | 183 |
| 54 Understanding Site Management | 185 |
| 54.1 Site Scope | 185 |
| 54.2 Site Hierarchies | 186 |
| 54.3 Site Members | 188 |
| 54.4 Page Sets | 188 |
| 55 Adding Sites | 189 |
| 56 Adding Pages to Sites | 193 |
| 57 Creating Pages | 195 |
| 58 Creating Content Pages | 199 |
| 58.1 Creating Page Fragments | 199 |
| 58.2 Building Content Pages | 200 |
| 58.3 The Content Page Interface | 200 |
| 58.4 Personalizing Content Pages | 200 |
| 59 Content Page Elements | 201 |
| 59.1 Editable Elements | 203 |
| 59.2 Editable Text | 203 |
| 59.3 Editable Images | 203 |
| 59.4 Editable Links | 204 |

| | |
|--|------------|
| 60 Content Page Management Interface | 207 |
| 60.1 Sections | 208 |
| 60.2 Section Builder | 210 |
| 60.3 Widgets | 210 |
| 60.4 Page Structure | 212 |
| 60.5 Look and Feel | 212 |
| 60.6 Comments | 213 |
| 60.7 The Title Bar | 213 |
| 61 Building Content Pages | 217 |
| 61.1 Creating a Content Page | 218 |
| 62 Propagation of Changes | 223 |
| 63 Creating Page Fragments | 225 |
| 63.1 Creating and Managing Fragments | 225 |
| 64 Exporting and Importing Fragments | 229 |
| 64.1 Exporting Fragments | 229 |
| 64.2 Importing | 231 |
| 64.3 Importing Collections | 231 |
| 64.4 Importing Individual Page Fragments | 232 |
| 65 Using Widget Pages | 233 |
| 66 Creating Widget Pages | 235 |
| 66.1 Adding a Widget Page | 235 |
| 66.2 Adding Widgets to a Page | 237 |
| 67 Creating Widget Pages from Templates | 239 |
| 67.1 Sharing Widget Page Templates | 241 |
| 68 Building a Responsive Site | 243 |
| 69 Built-in Mobile Support | 245 |
| 69.1 Using the Device Simulator | 246 |
| 69.2 Designing Mobile Friendly Pages | 246 |
| 70 Mobile Device Rules | 249 |
| 71 Creating Mobile Device Rules | 251 |
| 72 Mobile Device Actions | 255 |
| 72.1 Mobile Device Rules Example | 257 |
| 73 Using Full Page Applications | 261 |
| 73.1 Configuring the Page | 261 |
| 74 Managing Site Navigation | 265 |

| | |
|--|------------|
| 75 Page Hierarchy | 267 |
| 75.1 Creating a Page | 267 |
| 75.2 Organizing Pages | 268 |
| 75.3 Public and Private Pages | 270 |
| 75.4 Page Options | 270 |
| 76 Creating and Managing Navigation Menus | 271 |
| 76.1 Creating a Navigation Menu | 271 |
| 76.2 Managing Menus | 272 |
| 76.3 Modifying Menus | 273 |
| 77 Displaying Navigation Menus | 275 |
| 77.1 Navigation Menu Widget | 275 |
| 77.2 Choosing a Navigation Menu | 275 |
| 77.3 Display Template | 276 |
| 77.4 Menu Items to Show | 277 |
| 78 Building Sites from Templates | 279 |
| 79 Creating a Site Template | 281 |
| 80 Managing Site Templates | 285 |
| 81 Propagating Changes from Site Templates to Sites | 287 |
| 81.1 Site Template Page Behavior | 287 |
| 81.2 Merging and Resetting Changes | 288 |
| 82 Sharing Site Templates | 291 |
| 83 Configuring Sites | 293 |
| 84 Configuring Site Settings | 295 |
| 84.1 Details | 295 |
| 84.2 Membership Options | 295 |
| 84.3 Site Hierarchies | 295 |
| 84.4 Pages | 296 |
| 84.5 Categorization | 296 |
| 84.6 Site URL | 296 |
| 84.7 Documents and Media | 297 |
| 84.8 Site Template | 298 |
| 84.9 Asset Auto Tagging | 298 |
| 84.10 Custom Fields | 298 |
| 85 Social Settings and Languages | 299 |
| 85.1 Ratings | 299 |
| 85.2 Mentions | 299 |
| 85.3 Languages | 299 |
| 86 Advanced Site Settings | 301 |
| 86.1 Default User Associations | 301 |

| | |
|---|------------|
| 86.2 Analytics | 301 |
| 86.3 Maps | 303 |
| 86.4 Recycle Bin | 303 |
| 86.5 Content Sharing | 303 |
| 87 Customizing Personal Sites | 305 |
| 88 Importing/Exporting Sites and Content | 307 |
| 88.1 Backing Up and Restoring Pages and Their Content | 307 |
| 88.2 Page Export Example | 308 |
| 88.3 Export Templates | 309 |
| 89 Styling Apps and Assets | 311 |
| 90 Styling Widgets with Widget Templates | 313 |
| 90.1 Creating a Widget Template | 313 |
| 90.2 The Template Editor | 314 |
| 90.3 Configuring Widget Templates | 316 |
| 91 Widget Template Example | 319 |
| 92 Setting a Default Widget Template | 321 |
| 92.1 Default Widget Template Example | 322 |
| 93 Customizing Page Options | 325 |
| 94 Configuring Page Sets | 327 |
| 95 Configuring Page Sets | 329 |
| 95.1 Themes | 329 |
| 95.2 Using a Custom Logo for a Site | 331 |
| 96 Advanced Page Set Options | 333 |
| 96.1 Executing JavaScript in Site Pages | 333 |
| 96.2 Merge Public Pages | 333 |
| 96.3 Rendering Pages for Mobile Devices | 334 |
| 96.4 Robots | 334 |
| 96.5 Notifying Search Engines of Site Pages | 334 |
| 97 Configuring Individual Pages | 335 |
| 98 Individual Page Settings | 337 |
| 98.1 General | 337 |
| 98.2 Name and Friendly URL | 337 |
| 98.3 Page Layout | 338 |
| 98.4 Categorization and SEO | 338 |
| 98.5 Categorization | 338 |
| 98.6 SEO | 339 |
| 98.7 Look and Feel | 339 |
| 98.8 Advanced Settings | 340 |

| | |
|--|------------|
| 98.9 Query String | 341 |
| 98.10 Custom Fields | 341 |
| 98.11 Embedded Widgets | 341 |
| 98.12 Customization Settings | 341 |
| 98.13 JavaScript | 341 |
| 98.14 Mobile Device Rules | 342 |
| 99 Personalizing Pages | 343 |
| 99.1 Enabling Page Customizations | 343 |
| 99.2 Customization Permissions | 344 |
| 99.3 Customizing Pages | 344 |
| 99.4 Viewing Customized Pages | 345 |
| 99.5 Customization Example | 346 |
| 100 Changing Page Permissions | 347 |
| 101 Configuring Widgets | 349 |
| 102 Look and Feel Configuration | 351 |
| 102.1 General Settings | 351 |
| 102.2 Text Styles | 352 |
| 102.3 Background Styles | 352 |
| 102.4 Border Styles | 352 |
| 102.5 Margin and Padding | 354 |
| 102.6 Advanced Styling | 355 |
| 103 Exporting/Importing Widget Data | 357 |
| 103.1 Exporting Widget Data | 359 |
| 103.2 Importing Widget Data | 359 |
| 104 Communication Between Portlet Widgets | 361 |
| 105 Sharing Widgets with Other Sites | 363 |
| 105.1 Any Web Site | 363 |
| 105.2 Facebook | 363 |
| 105.3 OpenSocial Gadget | 364 |
| 105.4 Netvibes | 365 |
| 106 Widget Permissions | 367 |
| 107 Widget Scope | 369 |
| 108 Configuration Templates | 373 |
| 108.1 Summary | 374 |
| 109 Managing Members in Your Site | 375 |
| 110 Adding Members to Sites | 377 |
| 110.1 Administrating Site Membership | 377 |
| 110.2 Adding Members to a Site | 377 |

| | | |
|------------|--|------------|
| 110.3 | Removing User Membership from a Site | 378 |
| 110.4 | Assigning Site Roles | 379 |
| 111 | Creating Teams to Empower Site Members | 381 |
| 112 | Managing Content | 385 |
| 113 | Managing Web Content | 387 |
| 114 | Publishing Basic Web Content | 389 |
| 115 | Creating Web Content | 391 |
| 116 | Using the Web Content Editor | 393 |
| 116.1 | Basic Editor Functions | 393 |
| 116.2 | Editing the Article Source | 394 |
| 116.3 | Web Content Options | 396 |
| 117 | Publishing Web Content | 399 |
| 117.1 | Editing Published Content | 401 |
| 118 | Other Content Options | 403 |
| 118.1 | Localizing Content | 403 |
| 118.2 | Xuggler for Embedding Video | 404 |
| 118.3 | XML Format Downloads | 405 |
| 118.4 | Subscribing to Content | 405 |
| 118.5 | Organizing Structure Names | 406 |
| 119 | Designing Structured Content | 409 |
| 120 | Creating Structured Web Content | 411 |
| 121 | Editing Structures | 413 |
| 121.1 | Structure Fields | 413 |
| 122 | Configuring Structure Fields | 417 |
| 122.1 | Structure Default Values | 418 |
| 123 | Structure Settings | 421 |
| 123.1 | Assigning Permissions | 421 |
| 123.2 | WebDAV URL | 421 |
| 124 | Designing Web Content with Templates | 423 |
| 125 | Adding Templates with Structures | 425 |
| 126 | Embedding Widgets in Templates | 427 |
| 127 | Using Taglibs in Templates | 429 |
| 128 | Assigning Template Permissions | 431 |

| | |
|--|------------|
| 128.1Assigning Permissions for Individual Templates | 431 |
| 129Display Page Templates for Web Content | 433 |
| 130Creating Display Page Templates | 435 |
| 131Display Page Template Example | 437 |
| 131.1Publishing with Display Page Templates | 438 |
| 132Managing Content Sets | 443 |
| 132.1Creating and Displaying Content Sets | 443 |
| 132.2Content Set Personalization | 444 |
| 132.3Converting Asset Publisher Configurations to Content Sets | 444 |
| 133Creating Content Sets | 445 |
| 133.1Creating a Manual Content Set | 445 |
| 133.2Creating a Dynamic Content Set | 447 |
| 134Displaying Content Sets | 449 |
| 134.1Configuring the Asset Publisher for Content Sets | 449 |
| 134.2Adding Items to an existing Content Set | 451 |
| 135Converting Asset Publisher Configurations to Content Sets | 455 |
| 136Organizing Content with Tags and Categories | 457 |
| 137Tagging Content | 459 |
| 138Defining Categories for Content | 461 |
| 139Targeted Vocabularies | 465 |
| 139.1Single and Multi-valued Vocabularies | 466 |
| 139.2Separated Entries | 466 |
| 140Geolocating Assets | 469 |
| 140.1Geolocating Web Content | 469 |
| 140.2Geolocating Data Lists | 470 |
| 140.3Geolocating Documents and Media | 471 |
| 141Publishing Content Dynamically | 475 |
| 142Defining Content Relationships | 477 |
| 142.1Related Assets Widget | 477 |
| 143Publishing Assets | 481 |
| 144Querying for Content | 483 |
| 145Selecting Assets | 485 |
| 145.1Selecting Assets Manually | 485 |
| 145.2Selecting Assets Dynamically | 487 |

| | | |
|------------|---|------------|
| 145.3 | Selecting a Content Set | 489 |
| 146 | Configuring Display Settings | 491 |
| 147 | Configuring Asset Publisher Subscriptions | 495 |
| 148 | Publishing RSS Feeds | 497 |
| 149 | Configuring RSS Feeds | 499 |
| 150 | The RSS Publisher Widget | 503 |
| 150.1 | Using the RSS Publisher Widget | 504 |
| 151 | Restoring Deleted Assets | 507 |
| 152 | Configuring the Recycle Bin | 509 |
| 153 | Using the Recycle Bin | 511 |
| 153.1 | Drag and Drop | 513 |
| 154 | Recycle Bin Intelligence and Support | 515 |
| 155 | Collaboration | 517 |
| 156 | Managing Documents and Media | 521 |
| 157 | Publishing Files | 523 |
| 158 | Adding Files to a Document Library | 527 |
| 158.1 | Granting File Permissions and Roles | 527 |
| 158.2 | Using the Add Menu | 528 |
| 159 | Creating Folders | 531 |
| 159.1 | Adding a Folder | 531 |
| 159.2 | Document Type Restrictions and Workflow | 532 |
| 159.3 | Setting Folder Permissions | 533 |
| 160 | Using the Documents and Media Management Bar | 535 |
| 160.1 | View Types | 535 |
| 160.2 | The Info Panel | 536 |
| 160.3 | Finding and Arranging Items | 537 |
| 160.4 | Selecting Items | 537 |
| 161 | Viewing File Previews | 539 |
| 161.1 | File Preview Apps | 539 |
| 161.2 | Managing Files | 541 |
| 161.3 | The Info Panel | 542 |
| 162 | Editing Images | 543 |
| 163 | Publishing Files | 547 |

| | |
|---|------------|
| 163.1Using the Media Gallery | 547 |
| 164Checking Out and Editing Files | 551 |
| 165Sharing Files | 555 |
| 165.1Sharing Files in Documents and Media | 556 |
| 165.2Working with Shared Files | 556 |
| 165.3Managing Shared Files | 558 |
| 166Configuring Sharing | 563 |
| 166.1Global Configuration | 563 |
| 166.2Instance Configuration | 563 |
| 166.3Site Configuration | 564 |
| 167Desktop Access to Documents and Media | 567 |
| 168Linking to Google Drive™ | 569 |
| 168.1Install the App | 569 |
| 168.2Configure Your Google Project | 570 |
| 168.3Configure Your Portal | 570 |
| 168.4Creating Linked Files | 571 |
| 169Metadata Sets | 575 |
| 169.1Managing Metadata Sets | 575 |
| 169.2Creating Metadata Sets | 576 |
| 170Document Types | 581 |
| 170.1Managing Document Types | 581 |
| 170.2Creating Document Types | 581 |
| 171Online File Creation and Editing with Google Docs™ | 585 |
| 172Configuring Google Docs™ Integration | 589 |
| 172.1Configure Your Google Project | 589 |
| 172.2Configuring the Portal | 589 |
| 173Creating and Editing Files with Google Docs™ | 593 |
| 173.1Authentication | 593 |
| 173.2Creating Files | 593 |
| 173.3Editing Files | 595 |
| 173.4Multiple Editing Sessions | 596 |
| 174Integration with Microsoft Office 365™ | 599 |
| 175Configuring Office 365™ Integration | 601 |
| 175.1Register an Application with the Microsoft Identity Platform | 601 |
| 175.2Configuring Liferay DXP | 601 |
| 176Creating and Editing Documents and Media Files with Office 365™ | 605 |
| 176.1Authentication | 605 |

| | | |
|------------|---|------------|
| 176.2 | Creating Files | 605 |
| 176.3 | Editing Files | 608 |
| 177 | Store Types | 611 |
| 177.1 | Simple File System Store | 611 |
| 177.2 | Using the Advanced File System Store | 612 |
| 177.3 | Using Amazon Simple Storage Service | 613 |
| 178 | Liferay Sync | 615 |
| 179 | Administering Liferay Sync | 617 |
| 180 | Installing Liferay Sync’s Prerequisites | 619 |
| 181 | Configuring Liferay Sync | 621 |
| 182 | Preventing Accidental File Deletion in Liferay Sync | 625 |
| 183 | Ensuring Liferay Sync Security | 627 |
| 183.1 | Liferay Sync Permissions Demonstration | 627 |
| 184 | Using Liferay Sync on Your Desktop | 629 |
| 185 | Installing and Configuring the Desktop Liferay Sync Client | 631 |
| 185.1 | Installing the Liferay Sync Desktop Client | 631 |
| 185.2 | Configuring the Liferay Sync Desktop Client | 632 |
| 186 | Using the Liferay Sync Desktop Client | 637 |
| 186.1 | Using Sync Preferences | 638 |
| 187 | Using Your Local Liferay Sync Folder | 643 |
| 188 | Using Liferay Sync on Your Mobile Device | 645 |
| 189 | Connecting Liferay Sync Mobile | 647 |
| 190 | Managing Files and Folders in Liferay Sync Mobile | 651 |
| 191 | Adapting Your Media Across Multiple Devices | 655 |
| 192 | Installing Adaptive Media | 657 |
| 192.1 | Adaptive Media’s Modules | 657 |
| 192.2 | Processing Animated GIFs | 658 |
| 193 | Adding Image Resolutions | 659 |
| 193.1 | Adding a New Image Resolution | 660 |
| 194 | Managing Image Resolutions | 663 |
| 194.1 | Disabling Image Resolutions | 663 |
| 194.2 | Enabling Image Resolutions | 663 |
| 194.3 | Editing Image Resolutions | 664 |

| | | |
|------------|---|------------|
| 194.4 | Deleting Image Resolutions | 664 |
| 194.5 | Generating Missing Adapted Images | 664 |
| 194.6 | The Recycle Bin and Adapted Images | 665 |
| 195 | Creating Content with Adapted Images | 667 |
| 195.1 | Including Adapted Images in Content | 667 |
| 195.2 | Using Adapted Images in Structured Web Content | 668 |
| 195.3 | Staging Adapted Images | 669 |
| 196 | Migrating Documents and Media Thumbnails to Adaptive Media | 671 |
| 196.1 | Adding the Replacement Image Resolutions | 671 |
| 196.2 | Creating the Adaptive Media Images | 672 |
| 196.3 | Running the Migration Process | 672 |
| 197 | Advanced Configuration Options | 675 |
| 198 | Collaborating | 677 |
| 199 | Publishing Blogs | 681 |
| 200 | Adding Blog Entries | 683 |
| 201 | Using the Blog Entry Editor | 689 |
| 201.1 | Using the Editor's Text View | 690 |
| 201.2 | Using the Editor's Code View | 690 |
| 202 | Managing Blog Entries | 693 |
| 202.1 | View Types | 694 |
| 202.2 | Finding and Arranging Blog Entries | 694 |
| 202.3 | Selecting Blog Entries | 694 |
| 202.4 | Sharing Blog Entries | 695 |
| 203 | Configuring the Blogs App | 697 |
| 204 | Displaying Blogs | 699 |
| 205 | Aggregating Blogs | 703 |
| 206 | Highlighting Recent Bloggers | 705 |
| 207 | Creating Forums with Message Boards | 707 |
| 208 | Creating Message Boards | 709 |
| 208.1 | Site-scoped Message Boards | 709 |
| 208.2 | Page-scoped Message Boards | 709 |
| 208.3 | Globally-scoped Message Boards | 710 |
| 209 | Configuring Message Boards | 713 |
| 209.1 | General Setup | 713 |
| 209.2 | Email Setup | 714 |

| | |
|---|------------|
| 209.3Thread Priorities | 714 |
| 209.4User Ranks | 715 |
| 209.5RSS | 715 |
| 210Message Board Permissions | 717 |
| 211Message Board Categories | 719 |
| 211.1Adding Categories | 719 |
| 211.2Adding Subcategories | 721 |
| 211.3Moving and Merging Categories | 722 |
| 212User Subscriptions and Mailing Lists | 723 |
| 212.1User Subscriptions | 723 |
| 212.2Mailing Lists | 724 |
| 213Using the Message Boards | 727 |
| 213.1Posting New Threads | 728 |
| 213.2Participating in Message Board Threads | 729 |
| 214Managing Message Boards | 731 |
| 214.1Locking Threads | 731 |
| 214.2Moving Threads | 733 |
| 214.3Deleting Threads | 733 |
| 214.4Banning Users | 733 |
| 214.5Splitting Threads | 733 |
| 214.6Editing Posts | 734 |
| 214.7Post Permissions | 734 |
| 215Mentioning Users | 735 |
| 216Working Together with the Wiki | 739 |
| 217Getting Started with Wikis | 741 |
| 217.1Configuring Wikis | 741 |
| 217.2Adding Wikis | 742 |
| 217.3Wiki Node Options | 742 |
| 218Adding and Editing Wiki Pages | 745 |
| 219Using the Wiki on Site Pages | 749 |
| 220Wiki Page Details | 753 |
| 220.1Details | 754 |
| 220.2History | 754 |
| 220.3Incoming/Outgoing Links | 755 |
| 220.4Attachments | 755 |
| 221Other Wiki Widgets | 757 |
| 221.1Page Menu | 757 |
| 221.2Tree Menu | 758 |

| | |
|---|------------|
| 221.3Wiki Display | 759 |
| 222Sending Alerts and Announcements | 761 |
| 222.1Creating Alerts and Announcements | 761 |
| 222.2User Configuration | 763 |
| 222.3Alert and Announcement Roles | 765 |
| 223Managing Notifications and Requests | 767 |
| 223.1Managing Notifications | 767 |
| 223.2Managing Multiple Notifications | 768 |
| 223.3Managing Requests | 768 |
| 224Using the Knowledge Base | 769 |
| 225Creating Knowledge Base Articles | 771 |
| 225.1Authoring Articles in the Editor | 771 |
| 225.2Importing Knowledge Base Articles | 773 |
| 226Managing the Knowledge Base | 775 |
| 226.1Setting the Knowledge Base’s Options | 775 |
| 226.2Managing Knowledge Base Articles | 777 |
| 227Knowledge Base Templates | 779 |
| 227.1Creating Templates | 780 |
| 227.2Managing Templates | 780 |
| 228Responding to Knowledge Base Feedback | 781 |
| 229Knowledge Base Display | 783 |
| 229.1Displaying Different Article Sets | 785 |
| 230Other Knowledge Base Widgets | 787 |
| 230.1Knowledge Base Article | 787 |
| 230.2Knowledge Base Section | 788 |
| 230.3Knowledge Base Search | 789 |
| 231Importing Knowledge Base Articles | 791 |
| 232Knowledge Base ZIP File Requirements | 793 |
| 233Knowledge Base Importer FAQs | 797 |
| 234Knowledge Base System Settings | 799 |
| 234.1Source URL Settings | 799 |
| 234.2Importer File Convention Settings | 800 |
| 234.3Section Names Setting | 801 |
| 235Inviting Members to Your Site | 803 |
| 236Assets | 805 |

| | | |
|------------|---|------------|
| 237 | Configuring Asset Auto Tagging | 807 |
| 237.1 | Configuration Levels | 807 |
| 237.2 | Global Configuration | 808 |
| 237.3 | Instance-level Configuration | 808 |
| 237.4 | Site-level Configuration | 809 |
| 238 | Auto Tagging Images | 811 |
| 238.1 | Configuring TensorFlow Image Auto Tagging | 811 |
| 238.2 | Configuring Google Cloud Vision | 813 |
| 238.3 | Configuring Microsoft Cognitive Services | 814 |
| 239 | Auto Tagging Text | 817 |
| 239.1 | Configuring Google Cloud Natural Language Text Auto Tagging | 818 |
| 239.2 | Configuring OpenNLP Text Auto Tagging | 818 |
| 240 | Creating A Social Network | 821 |
| 241 | Using the Activities Widget | 823 |
| 242 | Connecting Users | 825 |
| 243 | Exporting Widgets To Other Websites | 827 |
| 243.1 | Sharing OpenSocial Gadgets | 827 |
| 244 | Integrating with Facebook | 829 |
| 244.1 | Facebook Sign On | 829 |
| 244.2 | Using Your Widget as Facebook Applications | 829 |
| 245 | Using Social Bookmarks | 831 |
| 246 | Search | 833 |
| 246.1 | Elasticsearch | 833 |
| 246.2 | Search Features | 834 |
| 246.3 | Search Bar | 835 |
| 246.4 | Search Results and Relevance | 835 |
| 246.5 | Search Facets | 835 |
| 247 | What's New in Search | 837 |
| 247.1 | New and Improved Widgets | 837 |
| 247.2 | Custom Filter | 837 |
| 247.3 | Sort | 837 |
| 247.4 | Search Insights | 838 |
| 247.5 | New Search Admin Functionality | 838 |
| 247.6 | Search Engine Info | 838 |
| 247.7 | Field Mappings | 838 |
| 247.8 | Indexing Progress | 839 |
| 247.9 | New System Settings | 839 |
| 247.10 | New Infrastructure | 839 |
| 247.11 | Elasticsearch Support | 839 |

| | | |
|------------|---|------------|
| 247.1 | Application-Specific Indexes | 839 |
| 247.1A | API Enhancements | 840 |
| 247.1M | Multi-Language Search | 840 |
| 248 | Configuring Search Pages | 841 |
| 248.1 | Search Page Templates | 841 |
| 248.2 | Default Search Pages | 842 |
| 248.3 | Manual Search Page Configuration | 845 |
| 248.4 | Legacy Search Experience | 846 |
| 249 | Searching for Assets | 847 |
| 249.1 | Search Bar | 847 |
| 249.2 | Entering Search Terms | 848 |
| 249.3 | Matching Exact Phrases: Quoted Search | 848 |
| 249.4 | Prefix Searching | 848 |
| 249.5 | Configuring the Search Bar | 849 |
| 249.6 | Search Suggestions | 851 |
| 250 | Facets | 853 |
| 250.1 | Using Facets | 854 |
| 250.2 | Multiple Facet Selection | 855 |
| 250.3 | Facets and Friendly URLs | 856 |
| 251 | Site Facet | 859 |
| 252 | Type Facet | 861 |
| 253 | Tag and Category Facets | 863 |
| 254 | Folder Facet | 865 |
| 255 | User Facet | 867 |
| 256 | Modified Facet | 869 |
| 257 | Custom Facet | 873 |
| 257.1 | Finding Indexed Fields | 874 |
| 258 | Search Results | 877 |
| 259 | Display Settings | 879 |
| 260 | Filtering Search Results with the Custom Filter Widget | 883 |
| 260.1 | Adding and Configuring Custom Filters | 883 |
| 260.2 | Custom Filter Configuration Options | 884 |
| 261 | Sorting Search Results with the Sort Widget | 887 |
| 261.1 | Adding and Configuring the Sort Widget | 887 |
| 261.2 | Configuring the Sort Widget | 888 |
| 261.3 | Controlling the Sort Order | 890 |

| | | |
|------------|--|------------|
| 262 | Search Results Behavior | 891 |
| 262.1 | Filtering Results with Facets | 891 |
| 262.2 | Search Results Relevance | 891 |
| 262.3 | Permissions and Search Results | 892 |
| 262.4 | Initial Permissions Checking | 892 |
| 262.5 | Final Permissions Checking | 892 |
| 262.6 | Search and Staging | 893 |
| 262.7 | Result Summaries | 893 |
| 262.8 | Highlighting | 894 |
| 263 | Search Insights | 897 |
| 263.1 | Inspecting The Search Query String | 897 |
| 263.2 | Explaining Search Results | 898 |
| 264 | Searching for Localized Content | 901 |
| 264.1 | What is Localized Search? | 902 |
| 264.2 | Fully Localized | 902 |
| 264.3 | Fully Localized for Search | 902 |
| 264.4 | Site-Localized Search | 902 |
| 264.5 | Assets Supporting Localized Search | 903 |
| 264.6 | Web Content Articles: | 903 |
| 264.7 | Categories: | 903 |
| 264.8 | Documents and Media File Entries: | 904 |
| 264.9 | Dynamic Data Mapping Fields: | 904 |
| 264.10 | Examples | 904 |
| 264.11 | Fully Localized Search for Web Content Articles | 904 |
| 264.12 | Site-Localized Search for Documents and Media | 905 |
| 265 | Configuring Search | 907 |
| 265.1 | System Scoped Search Configuration | 907 |
| 265.2 | Default Keyword Query | 907 |
| 265.3 | Default Search Result Permission Filter | 909 |
| 265.4 | Index Status Manager | 909 |
| 265.5 | Indexer Writer Helper | 909 |
| 265.6 | Index Registry | 910 |
| 265.7 | Index Query Preprocessor | 910 |
| 265.8 | Reindex | 910 |
| 265.9 | Engine Helper | 911 |
| 265.10 | Permission Checker | 911 |
| 265.11 | Title Field Query Builder | 911 |
| 265.12 | Elasticsearch 6 | 911 |
| 265.13 | Search Web | 912 |
| 265.14 | Search Administration | 912 |
| 265.15 | Portal Properties | 913 |
| 265.16 | Site Scoped Search Configuration | 913 |
| 265.17 | Widget Scoped Search Configuration | 914 |
| 266 | Low Level Search Options: Searching Additional or Alternate Indexes | 915 |

| | | |
|------------|---|------------|
| 266.1 | Configuring Low Level Search | 915 |
| 266.2 | Example: Searching an Alternate Index | 917 |
| 267 | Search Tuning: Synonym Sets | 919 |
| 267.1 | Requirements and Limitations | 919 |
| 267.2 | Creating and Managing Synonym Sets | 920 |
| 267.3 | Using Synonym Sets | 921 |
| 267.4 | Known Issues | 921 |
| 267.5 | Related Resources | 921 |
| 268 | Search Tuning: Customizing Search Results | 923 |
| 268.1 | Availability | 924 |
| 268.2 | Requirements and Limitations | 924 |
| 268.3 | Creating and Managing Result Rankings | 925 |
| 268.4 | Adding Aliases | 925 |
| 268.5 | Activating and Deactivating Aliases | 926 |
| 268.6 | Pinning and Hiding Results | 926 |
| 268.7 | Adding Results | 927 |
| 268.8 | Re-Ordering Pinned Results | 928 |
| 268.9 | Result Rankings Scope and Permissions | 928 |
| 268.10 | Result Rankings Aliases versus Synonyms | 929 |
| 268.11 | Known Issues | 929 |
| 269 | Forms | 931 |
| 269.1 | Forms and Lists | 932 |
| 269.2 | Which Form Builder Should I Use? | 932 |
| 270 | Creating and Managing Forms | 933 |
| 270.1 | Viewing Forms | 933 |
| 270.2 | Building a Form | 935 |
| 270.3 | Accessing Forms | 938 |
| 271 | Managing Form Entries | 941 |
| 271.1 | Viewing Form Entries | 941 |
| 271.2 | Exporting Form Entries | 942 |
| 271.3 | Deleting Form Entries | 943 |
| 272 | Form Field Types | 945 |
| 273 | Form Rules | 951 |
| 273.1 | The Anatomy of a Form Rule | 951 |
| 273.2 | Creating Form Rules: Rule Builder | 952 |
| 273.3 | Conditions | 953 |
| 274 | Action: Show and Hide | 955 |
| 275 | Action: Require | 957 |
| 276 | Action: Enable and Disable | 959 |

| | |
|--|-------------|
| 277Action: Jump to Page | 961 |
| 278Action: Autofill | 963 |
| 278.1Using Inputs with Autofill | 964 |
| 279Action: Calculate | 969 |
| 280Form Element Sets | 971 |
| 280.1Creating Element Sets | 971 |
| 280.2Using Element Sets | 971 |
| 281Data Providers | 975 |
| 281.1Adding a Basic Data Provider | 976 |
| 281.2Using a Data Provider in a Select Field | 976 |
| 281.3Granting Data Provider Permissions | 978 |
| 281.4Data Provider Configuration | 978 |
| 281.5Troubleshooting Data Provider Errors | 979 |
| 282Auto-Save | 981 |
| 283Translating Forms | 983 |
| 284Autocompleting Text Fields | 987 |
| 284.1Configuring Autocomplete | 987 |
| 285Form Success Pages | 991 |
| 286Workflow and Forms | 993 |
| 286.1Enabling Workflow in a Form | 994 |
| 286.2Testing the Workflow | 995 |
| 287Duplicating Forms and Form Fields | 997 |
| 288Form Pages | 999 |
| 289Help Text, Placeholder Text, and Predefined Values | 1003 |
| 290Validating Text and Numeric Fields | 1005 |
| 290.1Validating Text Fields | 1006 |
| 290.2If Input Contains/Does Not Contain | 1006 |
| 290.3If Input Is not URL/Email | 1007 |
| 290.4If Input Does Not Match | 1007 |
| 290.5Validating Numeric Fields | 1007 |
| 291Enabling CAPTCHA on Form Submissions | 1011 |
| 292Form Notifications | 1013 |
| 293Redirecting Users | 1015 |
| 294Form Permissions | 1017 |

| | | |
|------------|---|-------------|
| 295 | Styling Form Pages | 1019 |
| 296 | Dynamic Data Lists | 1023 |
| 296.1 | System Configuration | 1024 |
| 297 | Creating Data Definitions | 1025 |
| 298 | Managing Data Definitions | 1031 |
| 299 | Creating Data Lists | 1033 |
| 299.1 | Creating List Records | 1033 |
| 299.2 | Configuring the Dynamic Data Lists Display Widget | 1035 |
| 300 | Using Templates to Display Forms and Lists | 1037 |
| 300.1 | Managing Display and Form Templates | 1037 |
| 301 | Creating Form Templates | 1039 |
| 302 | Creating Display Templates | 1041 |
| 302.1 | Display Template Editor | 1042 |
| 303 | Workflow | 1045 |
| 303.1 | What's New with Workflow? | 1045 |
| 303.2 | DXP Feature: Workflow Metrics | 1046 |
| 303.3 | Control Panel Reorganization | 1046 |
| 303.4 | Workflow Definition Permissions: System Settings | 1048 |
| 303.5 | Embedded Workflows | 1048 |
| 304 | Activating Workflow | 1049 |
| 304.1 | Workflow Assets | 1049 |
| 304.2 | Activating Workflow in Applications | 1050 |
| 304.3 | Workflow Behavior | 1050 |
| 305 | Managing Workflows | 1055 |
| 305.1 | Workflow Definition Publication Permissions | 1055 |
| 305.2 | Adding, Editing, and Deleting | 1056 |
| 305.3 | Uploading Workflow Definitions | 1056 |
| 305.4 | Published Versus Unpublished | 1057 |
| 305.5 | Workflow Versions | 1057 |
| 306 | Reviewing Assets | 1059 |
| 306.1 | Asset Submission to Workflow | 1060 |
| 306.2 | Assigning the Task | 1060 |
| 306.3 | Completing the Task | 1062 |
| 307 | Workflow Metrics: The Service Level Agreement (SLA) | 1065 |
| 307.1 | Adding SLAs | 1065 |
| 308 | Workflow Metrics: Reports | 1069 |
| 308.1 | Understanding Reports | 1069 |

| | | |
|------------|--|-------------|
| 308.2 | Pending Items | 1070 |
| 308.3 | Workload by Step | 1070 |
| 308.4 | Completed Items | 1071 |
| 308.5 | Completion Velocity | 1071 |
| 308.6 | Items View | 1072 |
| 308.7 | Filtering by SLA Status | 1072 |
| 308.8 | Filtering by Process Status and Completion Period | 1072 |
| 308.9 | Filtering by Process Step | 1073 |
| 308.10 | Combining Filters | 1074 |
| 308.11 | Item Details | 1074 |
| 309 | Workflow Designer | 1077 |
| 310 | Managing Workflows with the Workflow Designer | 1079 |
| 310.1 | Adding New Workflow Definitions with the Workflow Designer | 1079 |
| 310.2 | Saving and Publishing Workflow Definitions | 1079 |
| 310.3 | Adding Nodes | 1080 |
| 310.4 | Node Settings | 1081 |
| 310.5 | Related Topics | 1082 |
| 311 | Workflow Definition Nodes | 1083 |
| 311.1 | Node Actions and Notifications | 1083 |
| 311.2 | Actions | 1084 |
| 311.3 | Notifications | 1085 |
| 311.4 | Start and End Nodes | 1087 |
| 311.5 | State Nodes | 1087 |
| 311.6 | Related Topics | 1087 |
| 312 | Affecting the Processing of Workflow Definitions | 1089 |
| 312.1 | Transitions | 1089 |
| 312.2 | Forks and Joins | 1091 |
| 312.3 | Conditions | 1092 |
| 312.4 | Related Topics | 1093 |
| 313 | Creating Tasks in the Workflow Designer | 1095 |
| 313.1 | Assignments | 1095 |
| 313.2 | Resource Action Assignments | 1096 |
| 313.3 | Scripted Assignments | 1097 |
| 313.4 | Related Topics | 1098 |
| 314 | Kaleo Forms | 1099 |
| 314.1 | Creating Kaleo Forms Process | 1099 |
| 314.2 | Adding Records to a Process | 1102 |
| 315 | Segmentation and Personalization | 1103 |
| 315.1 | Defining Segments | 1103 |
| 315.2 | Integrating Segments with Analytics Cloud | 1104 |
| 315.3 | Personalizing Experiences | 1104 |
| 315.4 | Content Page Personalization | 1104 |

| | | |
|------------|--|-------------|
| 315.5 | Content Set Personalization | 1104 |
| 316 | The Segment Editor | 1105 |
| 316.1 | User Properties | 1106 |
| 316.2 | Organization Properties | 1106 |
| 316.3 | Session Properties | 1107 |
| 316.4 | Operations and Conjunctions | 1107 |
| 317 | Creating User Segments | 1109 |
| 317.1 | Creating a Custom Segment | 1109 |
| 317.2 | Managing Segments | 1110 |
| 318 | Creating Segments with Custom Fields and Session Data | 1113 |
| 318.1 | Creating a Custom Field | 1113 |
| 318.2 | Defining a Segment with a Custom Field | 1114 |
| 318.3 | Extending a Segment With Session Data | 1115 |
| 319 | Using Analytics Cloud With User Segments | 1117 |
| 319.1 | Creating a New Segment | 1117 |
| 319.2 | Getting Segment Analytics | 1117 |
| 320 | Personalization Experience Management | 1119 |
| 320.1 | Managing Content Page Personalization | 1119 |
| 320.2 | Managing Content Set Personalization | 1121 |
| 320.3 | Previewing User Experiences | 1121 |
| 321 | Content Page Personalization | 1123 |
| 321.1 | Creating the Default Page | 1123 |
| 321.2 | Defining Custom Experiences | 1124 |
| 322 | Content Set Personalization | 1127 |
| 322.1 | Creating and Setting the Default Content Set | 1127 |
| 322.2 | Personalizing the Content Set | 1128 |
| 323 | Recommending Content Based on User Behavior | 1129 |
| 323.1 | Adding Tags to Track User Behavior | 1130 |
| 323.2 | Displaying Content Based on User Behavior | 1131 |
| 324 | A/B Testing | 1133 |
| 325 | Enabling A/B Testing | 1135 |
| 325.1 | Setting A/B Testing Permissions | 1135 |
| 326 | Creating A/B Tests | 1137 |
| 327 | Running A/B Tests | 1141 |
| 328 | Monitoring A/B Test Results | 1143 |
| 329 | Publishing A/B Test Variants | 1145 |

| | | |
|------------|--|-------------|
| 330 | Content Publication Management | 1149 |
| 331 | Staging | 1151 |
| 332 | Enabling Staging | 1153 |
| 333 | Enabling Local Live Staging | 1155 |
| 334 | Enabling Remote Live Staging | 1157 |
| 335 | Configuring Servers for Remote Live Staging | 1161 |
| 335.1 | Applying Patches When Using Remote Staging | 1163 |
| 335.2 | Configuring Remote Staging's Buffer Size | 1163 |
| 336 | Enabling Page Versioning and Staged Content | 1165 |
| 337 | Publishing Staged Content Efficiently | 1169 |
| 337.1 | Understanding the Publication Process | 1169 |
| 337.2 | Planning Ahead for Staging | 1170 |
| 338 | Using the Staging Environment | 1171 |
| 339 | Staging Content | 1173 |
| 340 | Advanced Publication with Staging | 1177 |
| 340.1 | Date | 1177 |
| 340.2 | Pages | 1177 |
| 340.3 | Content | 1178 |
| 340.4 | Deletions | 1179 |
| 340.5 | Permissions | 1180 |
| 341 | Managing Content Types in Staging | 1181 |
| 341.1 | Referenced Content | 1182 |
| 341.2 | Version History | 1182 |
| 341.3 | Previews and Thumbnails | 1182 |
| 341.4 | Vocabularies | 1183 |
| 341.5 | Deletions | 1183 |
| 342 | Staging Processes and Templates | 1185 |
| 343 | Disabling Staging | 1187 |
| 343.1 | Disabling Local Live Staging | 1187 |
| 343.2 | Disabling Remote Live Staging | 1187 |
| 343.3 | Steps to Disable Staging | 1188 |
| 344 | Publishing Single Assets From a Staged Site | 1189 |
| 345 | Organizing Pages for Staging | 1193 |
| 346 | Using Multi and Single Page Variations | 1195 |
| 346.1 | Enabling Page Versioning | 1196 |

| | |
|--|-------------|
| 346.2Using Site Pages Variations | 1196 |
| 346.3Using Page Variations | 1197 |
| 346.4Managing Variation Permissions | 1198 |
| 347Creating Multi and Single Page Variations | 1199 |
| 348Merging Site Pages Variations | 1201 |
| 349Managing Permissions | 1203 |
| 350Scheduling Web Content Publication | 1205 |
| 351Managing Apps | 1207 |
| 352Managing and Configuring Apps | 1209 |
| 352.1Managing Apps in Production | 1209 |
| 352.2Using the App Manager | 1210 |
| 352.3Using the Components Listing | 1211 |
| 353Using the Liferay Marketplace | 1213 |
| 353.1Finding and Purchasing Apps | 1213 |
| 353.2Managing Purchased Apps | 1215 |
| 353.3Renewing a Purchased App | 1217 |
| 354Installing Apps Manually | 1219 |
| 354.1Using the Control Panel to Install Apps | 1219 |
| 354.2Using Your File System to Install Apps | 1220 |
| 354.3Manually Deploying an LPKG App | 1220 |
| 355App Types | 1221 |
| 355.1OSGi Modules | 1221 |
| 355.2Portlets | 1222 |
| 355.3Web Plugins | 1222 |
| 355.4Templates and Themes | 1222 |
| 355.5Liferay Marketplace App Packages | 1223 |
| 356Blacklisting OSGi Bundles and Components | 1225 |
| 356.1Blacklisting Bundles | 1225 |
| 356.2Reinstalling Blacklisted Bundles | 1226 |
| 356.3Blacklisting Components | 1227 |
| 356.4Re-enabling Blacklisted Components | 1228 |
| 357Polls | 1229 |
| 357.1Creating a Poll | 1229 |
| 357.2Adding a Poll to a Page | 1230 |
| 357.3Viewing Poll Results | 1231 |
| 358Using the Calendar | 1235 |
| 358.1Configuring the Calendar | 1235 |

| | | |
|------------|---|-------------|
| 359 | Using the Calendar Widget | 1239 |
| 359.1 | Adding New Calendars | 1240 |
| 359.2 | Adding Events to a Calendar | 1240 |
| 359.3 | Additional Event Functions | 1240 |
| 359.4 | Customizing Email Notifications | 1245 |
| 360 | Calendar Resources and Porting | 1247 |
| 360.1 | Calendar Resources | 1247 |
| 360.2 | Exporting and Importing Calendar Data | 1248 |
| II | Deploying Liferay DXP | 1251 |
| 361 | Deploying Liferay DXP | 1253 |
| 362 | Obtaining Liferay DXP | 1255 |
| 362.1 | Liferay Tomcat Bundle | 1255 |
| 362.2 | Downloading the Liferay WAR and Dependency JARs | 1256 |
| 363 | Preparing for Install | 1257 |
| 363.1 | JDK Requirements | 1257 |
| 363.2 | JVM Requirements | 1258 |
| 363.3 | Preparing a Database | 1259 |
| 363.4 | Using the Built-in Data Source | 1259 |
| 363.5 | Using a Data Source on Your Application Server | 1260 |
| 363.6 | Limiting Database Access | 1260 |
| 364 | Installing Liferay DXP | 1263 |
| 364.1 | Extracting a Liferay DXP Bundle | 1263 |
| 364.2 | Installing the JDBC Driver | 1263 |
| 364.3 | Running Liferay DXP for the First Time | 1264 |
| 364.4 | Using the Setup Wizard | 1264 |
| 364.5 | Portal | 1264 |
| 364.6 | Administrator User | 1264 |
| 364.7 | Database | 1265 |
| 365 | Installing Liferay DXP on Tomcat | 1267 |
| 365.1 | Installing Dependencies | 1267 |
| 365.2 | Configuring Tomcat | 1268 |
| 365.3 | Database Configuration | 1272 |
| 365.4 | Mail Configuration | 1273 |
| 365.5 | Deploying Liferay DXP | 1273 |
| 366 | Installing Liferay DXP on Wildfly | 1275 |
| 366.1 | Installing Dependencies | 1275 |
| 366.2 | Running Liferay DXP on Wildfly in Standalone Mode vs. Domain Mode | 1276 |
| 366.3 | Configuring Wildfly | 1277 |
| 366.4 | Database Configuration | 1280 |
| 366.5 | Mail Configuration | 1281 |

| | | |
|------------|---|-------------|
| 366.6 | Deploying Liferay DXP | 1282 |
| 367 | Installing Liferay DXP on JBoss EAP | 1283 |
| 367.1 | Installing Dependencies | 1283 |
| 367.2 | Running Liferay DXP on JBoss EAP in Standalone Mode vs. Domain Mode | 1285 |
| 367.3 | Configuring JBoss | 1285 |
| 367.4 | Database Configuration | 1288 |
| 367.5 | Mail Configuration | 1289 |
| 367.6 | Deploy Liferay | 1290 |
| 368 | Installing Liferay DXP on WebLogic 12c R2 | 1291 |
| 368.1 | Configuring WebLogic's Node Manager | 1291 |
| 368.2 | Configuring WebLogic | 1292 |
| 368.3 | Setting Liferay DXP Properties | 1293 |
| 368.4 | Installing Liferay DXP Dependencies | 1293 |
| 368.5 | Database Configuration | 1294 |
| 368.6 | Mail Configuration | 1295 |
| 368.7 | Deploying Liferay DXP | 1295 |
| 369 | Installing Liferay DXP on WebSphere | 1297 |
| 369.1 | Preparing WebSphere for Liferay DXP | 1297 |
| 369.2 | Configuring the WebSphere Application Server | 1299 |
| 369.3 | Setting up JVM Parameters for Liferay DXP | 1299 |
| 369.4 | Removing the secureSessionCookie Tag | 1301 |
| 369.5 | Installing Liferay DXP's Dependencies | 1301 |
| 369.6 | Ensuring that Liferay DXP's portlet.jar is loaded first | 1302 |
| 369.7 | Database Configuration | 1302 |
| 369.8 | Mail Configuration | 1305 |
| 369.9 | Creating a WebSphere-Managed Mail Session (Optional) | 1305 |
| 369.10 | Verifying WebSphere Mail Provider | 1306 |
| 369.11 | Enable Cookies for HTTP Sessions | 1306 |
| 369.12 | Enable UTF-8 | 1306 |
| 369.13 | Deploy Liferay DXP | 1307 |
| 369.14 | Setting the JDK Version for Compiling JSPs | 1307 |
| 369.15 | Start Liferay DXP | 1309 |
| 370 | Activating Liferay DXP | 1311 |
| 371 | Setting Up Marketplace | 1313 |
| 371.1 | Server is Firewalled without Access to the Internet | 1313 |
| 371.2 | Limited Database Access | 1314 |
| 372 | Trial Plugin Installation | 1315 |
| 372.1 | Installation Process | 1315 |
| 372.2 | FAQ | 1316 |
| 373 | Document Repository Configuration | 1319 |
| 374 | Using the Simple File System Store | 1321 |

| | |
|---|-------------|
| 375Using the Advanced File System Store | 1323 |
| 376Using Amazon Simple Storage Service | 1325 |
| 377Using the DBStore | 1327 |
| 378Configuring Liferay DXP | 1329 |
| 379Configuring Mail | 1331 |
| 379.1Configuring Liferay DXP's Built-in Mail Session | 1331 |
| 379.2Built-in Mail Session in the Control Panel | 1332 |
| 379.3Built-in Mail Session Portal Properties | 1332 |
| 379.4Configuring a Mail Session on the Application Server | 1333 |
| 379.5Configuring default email senders | 1333 |
| 380Locales and Encoding Configuration | 1335 |
| 380.1Time Zones | 1335 |
| 380.2Set the JVM Time Zone to GMT | 1336 |
| 380.3Friendly URLs and Locales | 1337 |
| 380.4Modifying Language Keys | 1337 |
| 380.5Right to Left | 1337 |
| 380.6Localizing User Names | 1338 |
| 380.7Related Topics | 1338 |
| 381Liferay DXP Clustering | 1339 |
| 382Point all Nodes to the Same Liferay DXP Database | 1341 |
| 382.1Read-Writer Database Configuration | 1341 |
| 382.2JDBC | 1341 |
| 382.3JNDI | 1343 |
| 383Configure Documents and Media the Same for all Nodes | 1345 |
| 384Clustering Search | 1347 |
| 385Enabling Cluster Link | 1349 |
| 385.1Enabling Cluster Link | 1349 |
| 385.2Multicast Over UDP | 1351 |
| 385.3Checkpoint: | 1351 |
| 385.4Unicast over TCP | 1352 |
| 385.5JDBC Ping | 1353 |
| 385.6S3 Ping | 1354 |
| 385.7Other Pings | 1354 |
| 385.8Using Different Control and Transport Channel Ports | 1354 |
| 385.9Modifying the Cache Configuration with a Module | 1356 |
| 385.10Conclusion | 1356 |
| 386Auto Deploy to All Nodes | 1357 |
| 387Updating a Cluster | 1359 |

| | | |
|------------|---|-------------|
| 388 | Rolling Restarts | 1363 |
| 388.1 | New Modules and Plugins | 1363 |
| 388.2 | Updating Existing Modules and Plugins | 1364 |
| 388.3 | Applying Fix Packs (DXP only) | 1364 |
| 388.4 | Reverting Fix Packs (DXP only) | 1364 |
| 388.5 | Portal Properties controlled by portal-ext.properties | 1364 |
| 388.6 | System Settings controlled by Configuration Admin Files | 1364 |
| 388.7 | Application Server or JVM setting modifications | 1364 |
| 388.8 | Java Version Updates | 1365 |
| 388.9 | Related Topics | 1365 |
| 389 | Blue-Green Deployment | 1367 |
| 389.1 | Related Topics | 1367 |
| 390 | Configuring Remote Staging in a Clustered Environment | 1369 |
| 391 | Content Delivery Network | 1373 |
| 391.1 | Using CDN for Performance Enhancements | 1373 |
| 391.2 | Liferay CDN Requirements | 1374 |
| 391.3 | Configuring Liferay DXP to Use a CDN | 1375 |
| 392 | Tuning Guidelines | 1377 |
| 392.1 | Database Connection Pool | 1377 |
| 392.2 | Deactivating Development Settings in the JSP Engine | 1378 |
| 392.3 | Thread Pool | 1378 |
| 393 | Java Virtual Machine Tuning | 1381 |
| 393.1 | Garbage Collector | 1381 |
| 393.2 | Code Cache | 1382 |
| 393.3 | Java Heap | 1382 |
| 393.4 | JVM Advanced Options | 1383 |
| 394 | Installing a Search Engine | 1385 |
| 394.1 | Choosing a Search Engine | 1385 |
| 394.2 | End User Feature Limitations of Liferay's Solr Integration | 1385 |
| 394.3 | Developer Feature Limitations of Liferay's Solr Integration | 1386 |
| 394.4 | Elasticsearch Java Distribution Compatibility | 1387 |
| 395 | Elasticsearch | 1389 |
| 395.1 | Embedded vs. Remote Operation Mode | 1390 |
| 395.2 | Troubleshooting Elasticsearch Integration | 1390 |
| 396 | Preparing to Install Elasticsearch | 1393 |
| 396.1 | Sizing Your Deployment | 1393 |
| 396.2 | CPU | 1393 |
| 396.3 | Memory | 1394 |
| 396.4 | Disk | 1394 |
| 396.5 | Cluster Size | 1394 |
| 396.6 | Networking | 1394 |

| | | |
|------------|---|-------------|
| 397 | Installing Elasticsearch | 1395 |
| 397.1 | Step One: Download a Supported Version of Elasticsearch | 1396 |
| 397.2 | Step Two: Install Elasticsearch | 1397 |
| 397.3 | Step Three: Install Elasticsearch Plugins | 1397 |
| 397.4 | Step Four: Name Your Elasticsearch Cluster | 1398 |
| 397.5 | Step Five: Configure Liferay DXP to Connect to your Elasticsearch Cluster | 1398 |
| 397.6 | Step Six: Restart Liferay DXP and Reindex | 1399 |
| 398 | Configuring the Liferay Elasticsearch Connector | 1401 |
| 398.1 | Starting Elasticsearch | 1402 |
| 398.2 | Configuring the Liferay Elasticsearch Connector | 1402 |
| 398.3 | Configuring the Connector in the Control Panel | 1402 |
| 398.4 | Configuring the Connector with an OSGi .config File | 1403 |
| 398.5 | Configuring a Remote Elasticsearch Host | 1405 |
| 398.6 | Clustering Elasticsearch in Remote Operation Mode | 1406 |
| 398.7 | Elasticsearch Connector System Settings, By Operation Mode | 1406 |
| 399 | Advanced Configuration of the Liferay Elasticsearch Connector | 1407 |
| 399.1 | Adding Settings and Mappings to the Liferay Elasticsearch Connector | 1407 |
| 399.2 | Additional Configurations | 1407 |
| 399.3 | Adding Index Configurations | 1408 |
| 399.4 | Adding Type Mappings | 1408 |
| 399.5 | Overriding Type Mappings | 1409 |
| 399.6 | Multi-line YAML Configurations | 1410 |
| 400 | Elasticsearch Connector Settings: Reference | 1411 |
| 400.1 | Configurations only Affecting the Embedded Elasticsearch Server | 1413 |
| 401 | Installing Liferay Enterprise Search Security | 1415 |
| 401.1 | Enabling X-Pack Security | 1415 |
| 401.2 | Setting Up X-Pack Users | 1415 |
| 401.3 | Enabling Transport Layer Security | 1416 |
| 401.4 | Generate Node Certificates | 1416 |
| 401.5 | Enable TLS for Elasticsearch 7 | 1417 |
| 401.6 | Elasticsearch 6 TLS | 1418 |
| 401.7 | Example Elasticsearch Security Configuration | 1418 |
| 401.8 | Install and Configure the Liferay Enterprise Search Security app | 1419 |
| 401.9 | Disabling Elasticsearch Deprecation Logging | 1420 |
| 402 | Backing Up Elasticsearch | 1421 |
| 402.1 | Creating a Repository | 1421 |
| 402.2 | Taking Snapshots of the Cluster | 1422 |
| 402.3 | Restoring from a Snapshot | 1423 |
| 403 | Upgrading to Elasticsearch 6.5 | 1425 |
| 403.1 | Re-index | 1425 |
| 403.2 | Reverting to Elasticsearch 6.1 | 1426 |
| 404 | Upgrading to Elasticsearch 6.8 | 1427 |

| | | |
|------------|---|-------------|
| 404.1 | Re-index | 1427 |
| 404.2 | Reverting to Elasticsearch 6.1 | 1428 |
| 405 | Upgrading to Elasticsearch 7 | 1429 |
| 405.1 | Backing up Application-Specific Indexes | 1430 |
| 405.2 | Blacklisting Elasticsearch 6 | 1430 |
| 405.3 | Re-index | 1430 |
| 405.4 | Reverting to Elasticsearch 6 | 1430 |
| 406 | Installing Liferay Enterprise Search | 1433 |
| 407 | Installing Liferay Enterprise Search Monitoring | 1435 |
| 407.1 | Enable Encrypting Communication (TLS/SSL) in Elasticsearch and in Liferay DXP | 1436 |
| 407.2 | Enable Data Collection | 1436 |
| 407.3 | Install Kibana | 1436 |
| 407.4 | Configure Kibana with Authentication | 1437 |
| 407.5 | Configuring Kibana with Encryption (TLS/SSL) | 1437 |
| 407.6 | Configuring the Liferay Enterprise Search Monitoring app | 1438 |
| 407.7 | Monitoring in Liferay DXP | 1439 |
| 407.8 | Example Kibana Configuration | 1439 |
| 408 | Liferay Enterprise Search: Learning to Rank | 1441 |
| 408.1 | Disabling Learning to Rank on a Search Page | 1441 |
| 408.2 | Prerequisites | 1442 |
| 408.3 | Technical Overview | 1442 |
| 408.4 | Model Training | 1443 |
| 409 | Configuring Learning to Rank | 1445 |
| 409.1 | Step 1: Install the Learning to Rank Plugin on Elasticsearch | 1445 |
| 409.2 | Step 2: Training and Uploading a Model | 1445 |
| 409.3 | Upload the Model to the Learning to Rank Plugin | 1446 |
| 409.4 | Step 3: Enable Learning to Rank | 1447 |
| 410 | Installing Solr | 1449 |
| 410.1 | Blacklisting Elasticsearch-Only Features | 1449 |
| 411 | Installing Solr: Basic Installation | 1451 |
| 411.1 | Installing and Configuring Solr 7 | 1451 |
| 411.2 | Installing and Configuring the Liferay Solr Adapter | 1453 |
| 411.3 | Stopping the Elasticsearch Connector | 1453 |
| 411.4 | Install and Configure the Solr Connector | 1454 |
| 412 | Installing Solr: High Availability with SolrCloud | 1457 |
| 412.1 | Configure the Solr Adapter for SolrCloud | 1459 |
| 413 | Solr Connector Settings | 1461 |
| 413.1 | Solr 7 | 1461 |
| 414 | Securing Liferay DXP | 1463 |

| | | |
|------------|--|-------------|
| 414.1 | Authentication Overview | 1463 |
| 414.2 | Authorization and Permission Checking | 1464 |
| 414.3 | Additional Security Features | 1464 |
| 414.4 | Secure Configuration and Run Recommendations | 1465 |
| 415 | Logging into Liferay DXP | 1467 |
| 415.1 | Authentication Types | 1467 |
| 415.2 | The Sign In Portlet | 1468 |
| 416 | Service Access Policies | 1471 |
| 416.1 | Managing Service Access Policies | 1472 |
| 416.2 | Service Access Policy Module | 1474 |
| 416.3 | Summary | 1475 |
| 416.4 | Related Topics | 1475 |
| 417 | Authentication Verifiers | 1477 |
| 417.1 | Authentication Verification Process Overview | 1477 |
| 417.2 | Basic Auth Header | 1478 |
| 417.3 | Digest Auth Header | 1479 |
| 417.4 | HTTP Tunnel Extender | 1479 |
| 417.5 | Image Request Authentication Verifier | 1479 |
| 417.6 | Portal Sessions Auth Verifiers | 1479 |
| 417.7 | Request Parameter Auth Verifiers | 1480 |
| 417.8 | Tunnel Authentication Verifiers | 1480 |
| 417.9 | Related Topics | 1480 |
| 418 | LDAP | 1481 |
| 418.1 | Connection | 1481 |
| 418.2 | Checkpoint | 1482 |
| 418.3 | Instance Settings vs. System Settings | 1482 |
| 418.4 | Security | 1483 |
| 418.5 | Configuring LDAP Import/Export | 1483 |
| 418.6 | Users | 1483 |
| 418.7 | Groups | 1484 |
| 418.8 | Export | 1485 |
| 418.9 | Related Topics | 1486 |
| 419 | Configuring LDAP | 1487 |
| 419.1 | Export | 1487 |
| 419.2 | General | 1487 |
| 419.3 | Import | 1488 |
| 419.4 | Servers | 1489 |
| 419.5 | LDAP Options Available in System Settings | 1489 |
| 420 | Token-based Single Sign On Authentication | 1491 |
| 420.1 | Required SiteMinder Configuration | 1492 |
| 420.2 | Summary | 1492 |
| 421 | Authenticating with OpenID Connect | 1493 |

| | | |
|------------|--|-------------|
| 421.1 | Creating a Client in OpenID Connect Provider | 1493 |
| 421.2 | Configuring an OpenID Connect Provider Connection | 1494 |
| 421.3 | Enabling OpenID Connect Authentication | 1495 |
| 421.4 | Signing In With OpenID Connect | 1495 |
| 422 | OpenAM Single Sign On Authentication | 1497 |
| 423 | CAS (Central Authentication Service) Single Sign On Authentication | 1501 |
| 424 | Authenticating Using SAML | 1503 |
| 424.1 | What's new in Liferay Connector to SAML 2.0 | 1503 |
| 424.2 | Important SAML Paths | 1504 |
| 424.3 | Single Sign On | 1504 |
| 424.4 | Identity Provider Initiated SSO | 1504 |
| 424.5 | Service Provider Initiated SSO | 1506 |
| 424.6 | Single Log Off | 1507 |
| 424.7 | Identity Provider Initiated SLO | 1508 |
| 424.8 | Service Provider Initiated SLO | 1509 |
| 424.9 | Related Topics | 1510 |
| 425 | Setting up Liferay DXP as a SAML Identity Provider | 1511 |
| 425.1 | Storing Your Keystore | 1511 |
| 425.2 | Configuring Liferay DXP as a SAML Identity Provider | 1512 |
| 425.3 | Changing the Identity Provider Settings | 1513 |
| 425.4 | Checkpoint | 1515 |
| 426 | Registering a SAML Service Provider | 1517 |
| 426.1 | Checkpoint | 1518 |
| 427 | Setting up Liferay DXP as a SAML Service Provider | 1521 |
| 427.1 | Checkpoint | 1523 |
| 427.2 | Setting Up Liferay DXP as a SAML Service Provider in a Clustered Environment | 1523 |
| 428 | Changing the Settings for Service Provider and Identity Provider Connections | 1525 |
| 428.1 | Changing the SAML Identity Provider Connection Settings | 1526 |
| 428.2 | Configuring Negotiation Via metadata.xml | 1528 |
| 429 | NTLM Single Sign On Authentication | 1531 |
| 429.1 | Summary | 1532 |
| 430 | OpenID Single Sign On Authentication | 1533 |
| 430.1 | OpenID at the System Scope | 1533 |
| 430.2 | OpenID at the Instance Scope | 1534 |
| 431 | Authenticating with Kerberos | 1535 |
| 431.1 | How Kerberos Authentication Works | 1536 |
| 431.2 | Configuring Kerberos Authentication | 1537 |
| 431.3 | Creating the User Keytab | 1537 |
| 431.4 | Configuring Your Web Server | 1537 |

| | | |
|------------|--|-------------|
| 431.5 | Connecting Liferay DXP to Active Directory over LDAP | 1538 |
| 431.6 | Configuring your Clients | 1539 |
| 432 | Configuring CORS | 1541 |
| 432.1 | Enabling CORS for Liferay DXP Services | 1541 |
| 432.2 | Enabling CORS for JAX-RS Applications | 1543 |
| 433 | AntiSamy | 1545 |
| 433.1 | Configuring AntiSamy | 1545 |
| 433.2 | Using Wildcards | 1546 |
| 434 | OAuth 2.0 | 1547 |
| 434.1 | Flow of OAuth 2.0 | 1547 |
| 434.2 | OAuth 2.0 Terminology | 1549 |
| 434.3 | Creating an Application | 1550 |
| 435 | OAuth2 Scopes | 1553 |
| 435.1 | Creating a Scope for a JSONWS Service | 1553 |
| 435.2 | Creating the Authorization Page | 1554 |
| 436 | Authorizing Account Access with OAuth2 | 1557 |
| 436.1 | Authorization Code Flow | 1557 |
| 436.2 | PKCE Extended Authorization Code Flow | 1558 |
| 436.3 | Client Credentials and Resource Owner Flows | 1559 |
| 436.4 | Token Use | 1559 |
| 436.5 | Revoking Access | 1559 |
| 436.6 | Summary | 1561 |
| 437 | Upgrading to 7.0 | 1563 |
| 438 | Planning for Deprecated Applications | 1567 |
| 439 | Test Upgrading a Liferay DXP Backup Copy | 1569 |
| 439.1 | Preparing a Test Server and Database | 1569 |
| 439.2 | Copy the Production Installation to a Test Server | 1570 |
| 439.3 | Copy the Production Backup to the Test Database | 1570 |
| 440 | Pruning the Database | 1571 |
| 440.1 | Remove Duplicate Web Content Structure Field Names | 1571 |
| 440.2 | Find and Remove Unused Objects | 1571 |
| 440.3 | Objects From the Large/Populated Tables | 1572 |
| 440.4 | Common Object Types Worth Checking | 1572 |
| 440.5 | Test with the Pruned Database Copy | 1574 |
| 441 | Example: Removing Intermediate Journal Article Versions | 1575 |
| 442 | Upgrading Your Test Server and Database | 1579 |
| 442.1 | Install Liferay on a Test Server and Configure It to Use the Pruned Database | 1579 |
| 442.2 | Tune Your Database for the Upgrade | 1579 |
| 442.3 | Upgrade the Database | 1579 |

| | | |
|------------|---|-------------|
| 442.4 | Test the Upgraded Portal and Resolve Any Issues | 1580 |
| 442.5 | Checkpoint: You've Pruned and Upgraded a Production Database Copy | 1580 |
| 443 | Preparing to Upgrade the Liferay DXP Database | 1581 |
| 443.1 | Remove All Unused Objects You Identified Earlier | 1581 |
| 443.2 | Test Using the Pruned Database | 1581 |
| 443.3 | Upgrade Your Marketplace Apps | 1582 |
| 443.4 | Publish all Staged Changes to Production | 1582 |
| 443.5 | Synchronize a Complete Backup | 1582 |
| 444 | Preparing a New Liferay DXP Server for Data Upgrade | 1583 |
| 444.1 | Request an Upgrade Patch from Liferay Support (Liferay DXP only) | 1583 |
| 444.2 | Install Liferay | 1583 |
| 444.3 | Install the Latest Upgrade Patch or Fix Pack (Liferay DXP only) | 1584 |
| 444.4 | Migrate Your OSGi Configurations (7.0+) | 1584 |
| 444.5 | Migrate Your Portal Properties | 1584 |
| 444.6 | Update Your Portal Properties | 1584 |
| 444.7 | Convert Applicable Properties to OSGi Configurations | 1585 |
| 444.8 | Update Your Database Driver | 1586 |
| 444.9 | Configure Your Documents and Media File Store | 1586 |
| 444.10 | Configure Kerberos in place of NTLM | 1587 |
| 444.11 | Disable Indexing | 1587 |
| 445 | Upgrading the Liferay DXP Data | 1589 |
| 446 | Tuning for the Data Upgrade | 1591 |
| 446.1 | Tuning the Database Upgrade Java Process | 1592 |
| 446.2 | Tuning the Database Transaction Engine for Executing Updates | 1593 |
| 446.3 | IBM DB2 | 1593 |
| 446.4 | MariaDB | 1593 |
| 446.5 | Microsoft SQL Server | 1593 |
| 446.6 | MySQL | 1593 |
| 446.7 | Oracle Database | 1594 |
| 446.8 | PostgreSQL | 1594 |
| 446.9 | Tuning the Database Transaction Log | 1594 |
| 446.10 | IBM DB2 | 1594 |
| 446.11 | MariaDB | 1594 |
| 446.12 | Microsoft SQL Server | 1594 |
| 446.13 | MySQL | 1595 |
| 446.14 | Oracle Database | 1595 |
| 446.15 | PostgreSQL | 1595 |
| 447 | Configuring the Data Upgrade | 1597 |
| 447.1 | Configuring the Core Upgrade | 1597 |
| 447.2 | Configuring app-server.properties | 1598 |
| 447.3 | Configuring portal-upgrade-database.properties | 1598 |
| 447.4 | Configuring portal-upgrade-ext.properties | 1599 |
| 447.5 | Example Upgrade Configuration | 1599 |

| | |
|---|-------------|
| 447.6Configuring Non-Core Module Data Upgrades | 1600 |
| 448Upgrading the Core Using the Upgrade Tool | 1601 |
| 448.1Upgrade Tool Usage | 1601 |
| 448.2Running and Managing the Core Upgrade | 1602 |
| 449Upgrading Modules Using Gogo Shell | 1603 |
| 449.1Command Usage | 1603 |
| 449.2Listing module upgrade processes | 1604 |
| 449.3Executing module upgrades | 1605 |
| 449.4Checking upgrade status | 1605 |
| 449.5Executing verify processes | 1606 |
| 450Executing Post-Upgrade Tasks | 1607 |
| 450.1Tuning Your Database for Production | 1607 |
| 450.2Re-enabling Search Indexing and Re-indexing Search Indexes | 1607 |
| 450.3Enabling Web Content View Permissions | 1608 |
| 451Upgrading a Sharded Environment | 1609 |
| 451.1Add Configurations Before the Data Upgrade | 1609 |
| 451.2Upgrade and Update Properties | 1610 |
| 452Migrating From Audience Targeting to Segmentation and Personalization | 1613 |
| 453Migrating User Segments | 1615 |
| 453.1Upgrade Process | 1615 |
| 454Manually Migrating from Audience Targeting | 1619 |
| 454.1User Attribute Rules | 1619 |
| 454.2Session Rules | 1619 |
| 454.3Behavior Rules | 1619 |
| 454.4Migrating Custom Rules | 1620 |
| 454.5Migrating Display Portlets | 1620 |
| 454.6User Segment Content Display | 1620 |
| 454.7Asset Publisher Personalization | 1620 |
| 455Deprecated Apps in 7.2: What to Do | 1621 |
| 455.1Foundation | 1621 |
| 455.2Personalization | 1622 |
| 455.3Web Experience | 1623 |
| 455.4Forms | 1623 |
| 455.5Security | 1623 |
| 455.6User and System Management | 1624 |
| 455.7Related Topics | 1624 |
| 456Apps in Maintenance Mode | 1625 |
| 456.1Related Topics | 1625 |
| 457Maintaining Liferay DXP | 1627 |

| | | |
|------------|---|-------------|
| 458 | Patching Liferay DXP | 1629 |
| 459 | Patching Basics | 1631 |
| 459.1 | Fix Packs | 1631 |
| 459.2 | Hotfixes | 1631 |
| 459.3 | Service Packs | 1632 |
| 459.4 | How Patches are Tested | 1632 |
| 460 | Using the Patching Tool | 1633 |
| 460.1 | Installing the Patching Tool | 1633 |
| 460.2 | Executables | 1634 |
| 461 | Installing Patches | 1635 |
| 461.1 | Handling Hotfixes and Patches | 1636 |
| 461.2 | Fix Pack Dependencies | 1637 |
| 461.3 | Updating the Patching Tool | 1637 |
| 461.4 | Cleaning Up | 1637 |
| 462 | Working with Patches | 1639 |
| 462.1 | Including support-info in Support Tickets | 1639 |
| 462.2 | Uninstalling Patches | 1639 |
| 462.3 | Showing collisions between patches and deployed plugins | 1640 |
| 462.4 | Separating Patches from the Installation | 1640 |
| 462.5 | Restoring the Separated Patch Files | 1641 |
| 463 | Configuring the Patching Tool | 1643 |
| 464 | Patching Tool Basic configuration | 1645 |
| 465 | Patching Tool Advanced Configuration | 1647 |
| 465.1 | Using Profiles with the Patching Tool | 1647 |
| 466 | Installing patches on the 7.0 WAR | 1649 |
| 466.1 | Prerequisites | 1649 |
| 466.2 | Install the patch on the Liferay DXP WAR and artifacts | 1649 |
| 466.3 | Related Topics | 1650 |
| 467 | Keeping up with Fix packs and Service Packs | 1651 |
| 468 | Backing up a Liferay DXP Installation | 1653 |
| 468.1 | Backing up Source Code | 1653 |
| 468.2 | Backing up Liferay DXP's File System | 1654 |
| 468.3 | Backing up Liferay DXP's Database | 1654 |
| 469 | Monitoring Liferay DXP | 1655 |
| 470 | Monitoring Garbage Collection and the JVM | 1657 |
| 470.1 | VisualVM | 1657 |
| 470.2 | JMX Console | 1657 |
| 470.3 | Garbage Collector Verbose Logging | 1658 |

| | | |
|------------|---|-------------|
| 471 | Managing Liferay DXP with Liferay Connected Services | 1661 |
| 472 | Getting Started with LCS | 1663 |
| 473 | LCS Preconfiguration | 1671 |
| 473.1 | Downloading the LCS Client App | 1671 |
| 473.2 | Preconfiguring LCS to Connect Through a Proxy | 1673 |
| 473.3 | JVM App Server Arguments | 1674 |
| 473.4 | LCS Client App Properties | 1674 |
| 473.5 | Ensuring Access to LCS | 1675 |
| 473.6 | NTP Server Synchronization | 1675 |
| 473.7 | Configuring WebSphere | 1675 |
| 473.8 | Installing the LCS Client App | 1676 |
| 473.9 | Upgrading the LCS Client App | 1677 |
| 474 | Registering Your Liferay DXP Server with LCS | 1679 |
| 474.1 | Determining Your Server's LCS Connection Status | 1682 |
| 475 | Using LCS | 1685 |
| 476 | What LCS Stores About Your Liferay DXP Servers | 1687 |
| 477 | Managing LCS Users in Your Project | 1689 |
| 477.1 | Managing LCS Roles | 1690 |
| 478 | Using the Dashboard | 1693 |
| 478.1 | Using the Project View | 1693 |
| 479 | Managing LCS Environments | 1695 |
| 479.1 | Creating Environments | 1695 |
| 479.2 | Working with Environments | 1697 |
| 480 | Managing LCS Servers | 1699 |
| 480.1 | Details | 1700 |
| 480.2 | Server Settings | 1700 |
| 480.3 | Page Analytics | 1701 |
| 480.4 | Snapshot Metrics | 1701 |
| 480.5 | Fix Packs | 1703 |
| 480.6 | Portal Properties | 1703 |
| 481 | Managing Your LCS Account | 1707 |
| 481.1 | Using Web Notifications | 1709 |
| 482 | Managing Liferay DXP Subscriptions | 1711 |
| 482.1 | Decommissioning Servers | 1713 |
| 482.2 | Elastic Subscriptions | 1713 |
| 483 | Understanding Environment Tokens | 1715 |
| 484 | Troubleshooting Your LCS Connection | 1717 |

| | |
|--|-------------|
| 484.1LCS Grace Periods | 1718 |
| 484.2Connection Grace Period | 1718 |
| 484.3Subscription Grace Period | 1719 |
| 484.4Troubleshooting | 1720 |
| 484.5Server Can't Reach LCS | 1720 |
| 484.6Subscription Issues | 1720 |
| 484.7Invalid Token | 1721 |
| 484.8Increasing Log Levels | 1721 |
| 484.9Control Panel | 1722 |
| 484.10log4j | 1722 |
| 485Deployment Reference | 1725 |
| 486Liferay Home | 1727 |
| 487Portal Properties | 1729 |
| 488System Properties | 1731 |
| 489Database Templates | 1733 |
| 489.1MariaDB | 1733 |
| 489.2MySQL | 1733 |
| 489.3PostgreSQL | 1734 |
| 490Comparing Patch Levels | 1735 |
| 491Patching Tool Configuration Properties | 1737 |
| 492Troubleshooting Deployments | 1739 |
| 493Liferay DXP Failed to Initialize Because the Database Wasn't Ready | 1741 |
| 493.1Related Topics | 1741 |
| 494Sort Order Changed with a Different Database | 1743 |
| 495Using Files to Configure Module Components | 1745 |
| 495.1Configuration File Formats | 1745 |
| 495.2Naming Configuration Files | 1746 |
| 495.3Resolving Configuration File Deployment Failures | 1747 |
| 495.4Related Articles | 1747 |

PREFACE

Welcome to the world of Liferay DXP! This book was written for anyone who has any part in setting up, using, or maintaining a web site built on Liferay DXP. For the end user, it contains everything you need to know about using the applications included with Liferay. For the administrator, you'll learn all you need to know about setting up your site with users, sites, organizations, and user groups, as well as how to manage your site's security with roles. For server admins, it guides you step-by-step through the installation, configuration, and optimization of Liferay DXP, including setting it up in a clustered, enterprise-ready environment. Use this book as a handbook for everything you need to do to get your Liferay DXP installation running smoothly, and then keep it by your side as you configure and maintain your Liferay-powered web site.

Conventions

The information contained herein has been organized in a way that makes it easy to locate information. The book has two parts. The first part, *Using Liferay Portal*, describes how to configure and use a freshly installed Liferay Portal. The second part, *Deploying Liferay Portal*, is for administrators who want to install Liferay Portal and optimize its performance.

Sections are broken up into multiple levels of headings, and these are designed to make it easy to find information.

Source code and configuration file directives are presented monospaced, as below.

Source code appears in a non-proportional font.

Italics represent links or buttons to be clicked on in a user interface.

Monospaced type denotes Java classes, code, or properties within the text.

Bold describes field labels and portlets.

Page headers denote the chapters and the section within the chapter.

Publisher Notes

It is our hope that this book is valuable to you, and that it becomes an indispensable resource as you work with Liferay DXP. If you need assistance beyond what is covered in this book, Liferay offers training¹, consulting², and support³ services to fill any need that you might have.

For up-to-date documentation on the latest versions of Liferay, please see the documentation pages on Liferay Learn.⁴

As always, we welcome feedback. If there is any way you think we could make this book better, please feel free to mention it on our forums or in the feedback on Liferay Learn. You can also use any of the email addresses on our Contact Us page.⁵ We are here to serve you, our users and customers, and to help make your experience using Liferay DXP the best it can be.

¹<https://learn.liferay.com>

²<https://www.liferay.com/consulting>

³<https://help.liferay.com>

⁴<https://learn.liferay.com>

⁵<https://www.liferay.com/contact-us>

Part I

Using Liferay DXP

THE LIFERAY DISTINCTION

Your web presence is a big deal, and the software platform that runs it must be up to the task. You must be able to update your site easily, to provide your partners with the services they need, and to enable your users and customers to interact with you. It shouldn't be complicated: in today's cloud-based environment, you should be able to make changes with a click of a button and dynamically allocate nodes whenever you need them.

Liferay DXP offers all of this and more. It is a mature, stable, open source platform with a strong heritage of serving some of the biggest sites in the world—and some of the smallest, too.

The documentation here shows you how to use it. This is the **User** section, which describes the features of an installed Liferay DXP, how to configure its applications, and how to build your website.

The **Developer** section is divided into five subsections:

Application Development shows you how you can build applications using industry standard tools and frameworks on Liferay DXP.

Tutorials lead you step-by-step through specific tasks, such as developing web applications, upgrading old applications, creating themes, and more.

Frameworks show you all the Liferay APIs and development frameworks you can use in your applications to streamline development.

Customization explains the myriad of ways Liferay DXP can be customized to your exact specifications.

Reference is a collection of material showing developers the options available for various APIs.

The **Installation and Upgrades** section shows administrators how to obtain, install, and configure Liferay DXP on bare metal, virtualized environments, and the cloud.

If you're coming from an older version, read on to learn what's new in 7.2.

WHAT'S NEW IN LIFERAY 7.2!

The latest version of Liferay DXP delivers powerful tools to help businesses create and personalize any experience across their solutions, leverage the flexibility of a decoupled CMS architecture, and streamline business operations.

Liferay DXP meets the needs of today's digital-first business teams to create experiences rapidly across channels. It equips enterprises with a wide variety of easy-to-use applications and tools to build tailored solutions and experiences on a flexible platform.

For a full feature list, please read the *Liferay DXP 7.2 Features Overview* or contact sales@liferay.com.

2.1 Key New Features

Ensure your business evolves and stays relevant with new features designed to support great customer experiences across all stages of the customer journey.

2.2 Experience Creation

Great digital experiences are imperative in today's competition for new business. Companies also need to make sure they stay relevant after the sale, investing in long-term customer relationships that cultivate loyalty and repeat business. Here are some new features to help you ensure you're creating and delivering excellent experiences all the time.

Content Authoring

7.0 evolves our content authoring and management capabilities significantly, making it even easier and more intuitive for the non-technical content creator to create and manage engaging content.

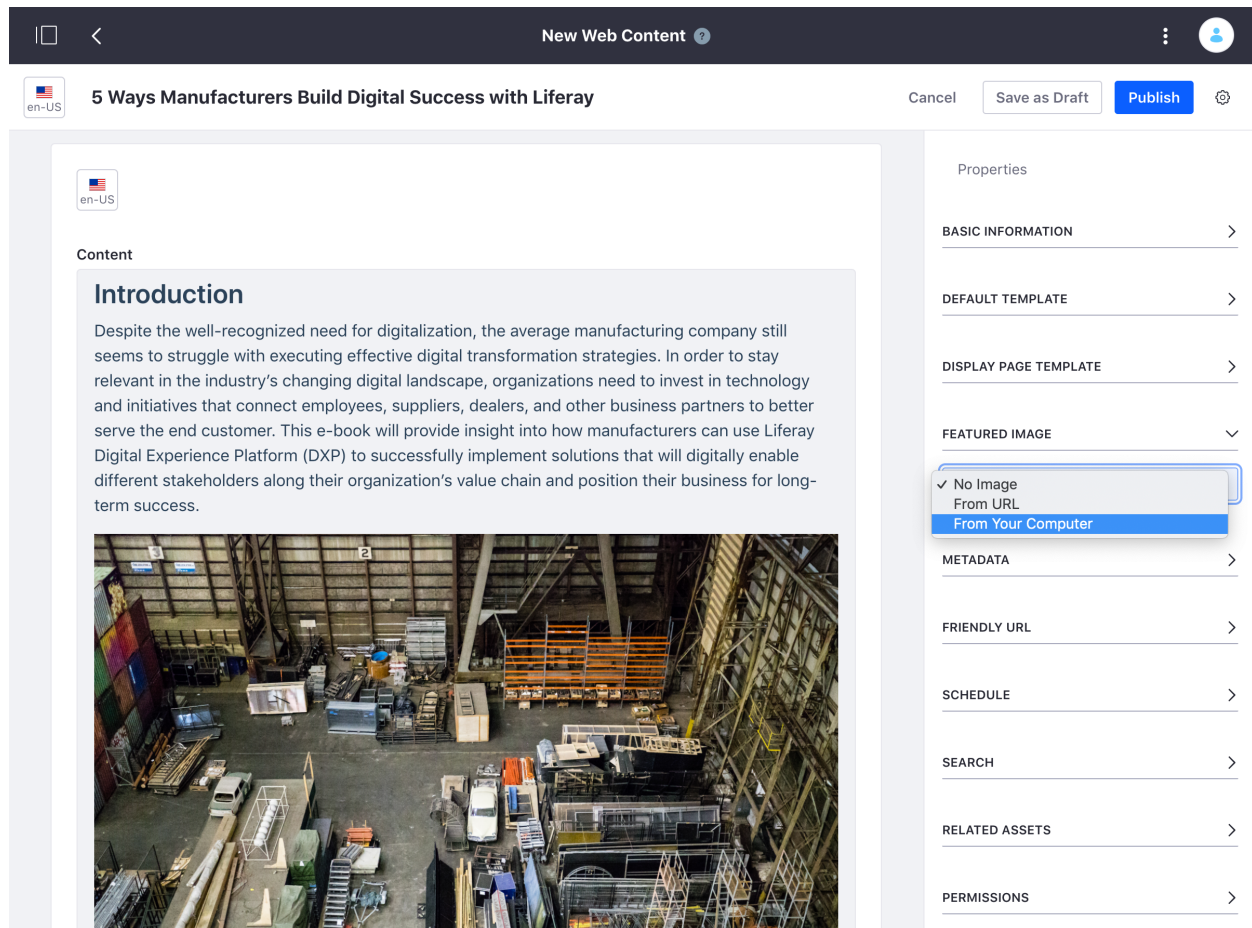


Figure 2.1: Content authors can build pages out of building blocks called *Fragments*.

Site Building

Content pages now allow users to create experiences that have both applications and content on the same page, giving businesses increased flexibility to create experiences tailored to fit the needs of different end users. More tools have been provided to allow for richer layouts and functionality for the non-technical user to build a visually appealing site experience easily.

Fragments

This release expands on the functionality available to create fragments, which previously required developer involvement. For marketers looking to create a fragment for their use case, we provide new functionality that allows a non-technical user to create simple fragments easily through the page editor itself by dragging and dropping out-of-the-box components into a container.

Fragments Toolkit/CLI

Improvements to the web developer's experience has also been made for 7.0. A set of front-end toolkits—including a CLI tool—is provided so web developers can write fragments in their own code editors and import/export them without needing to redeploy.




Figure 2.2: Fragments make it easy to build pages.




Figure 2.3: Use the Fragments Editor or download the toolkit and use your own tools.

In-Context Editing and Content Previews

An improved site building interface shows users how different content would appear to different visitors. Liferay DXP allows content creators to preview content in context of how it would look like on a live site.

Content Usages

For structured content that can be mapped and delivered across multiple locations, a new feature has been introduced to allow the content creator to view where specific pieces of web content are being used and reused across their channels.

A/B Testing (DXP only)

Content creators can use A/B testing to create and customize tests to evaluate which elements on Content Pages perform better and edit content accordingly. Liferay DXP's A/B tests leverage Bayesian statistics to identify the probable values of lift for a variant, allowing your business to make more informed decisions. Native integration with Liferay Analytics Cloud allows for data collected for the running test to determine a winning variant.

2.3 Personalization

In 7.0, segmentation and personalization capabilities have been moved into the core product. This allows for a more seamless integration of content creation and personalization functionality and helps streamline the process of creating segments while creating personalized experiences.

Session Rules

There are many different pieces of information available to track a web visitor. Session rules help further identify their distinct digital fingerprint per visit and is something that marketers can leverage to identify who is a repeat visitor and who is a new visitor.

Rule Builder

Being able to sort intelligently through vast amounts of data and associated individuals can be tedious; this process should be easy and not have to be curated manually. Marketers should instead be able to build lists of individuals based on a specific set of criteria using rules, leading to time savings and efficiency. The rule builder helps enable more effective targeting by reducing the amount of manual work needed to create a specific list of people who fall under certain pieces of criteria.

Content Sets

Content sets help address the use case of personalizing a feed of content for different groups of people. Users can easily target and personalize multiple pieces of content in a fragment or widget and have it display accordingly for the visitor.

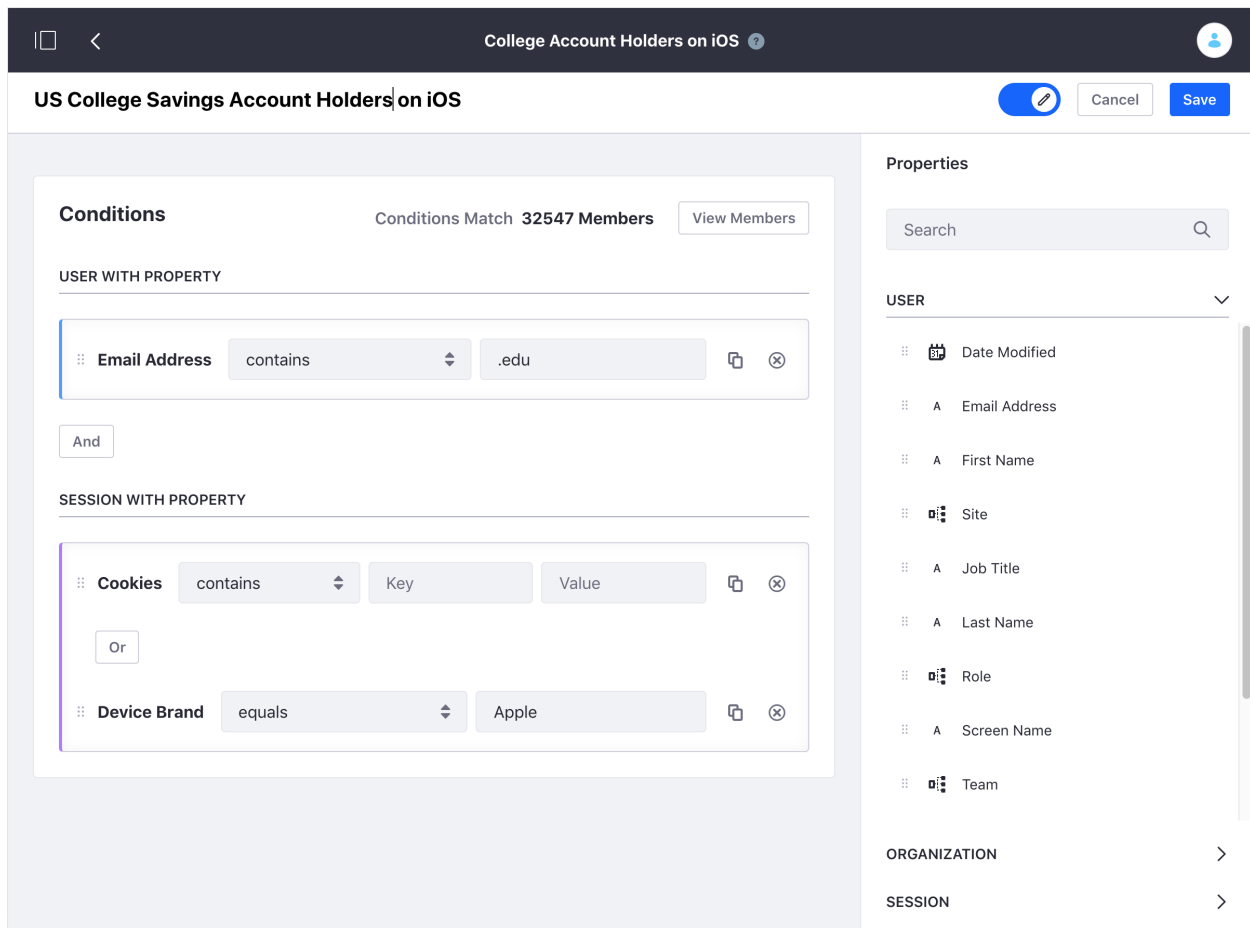


Figure 2.4: The Rule Builder provides a drag-and-drop interface that helps you build exactly the criteria you need to target the right information to the right users.

Experiences

The new Experiences feature allows a user to create different variations of a page. All the tools available for creating a page in the layout editor are also available for making these page variations. Non-technical users can easily tailor the messaging, images and even widgets on a page.

Content Recommendations (DXP only)

Content recommendations leverages interest models generated in Liferay Analytics Cloud to recommend content on Liferay DXP. The option to do so is available through Content Sets and automatically filters content based on interest keywords. Liferay Analytics Cloud uses AI to cluster topics and model long-term interest for known and unknown visitors. Keywords are taken from the categories, tags, and keywords of content.

2.4 Bulk Management

Enterprises regularly manage large volumes of content and work that gets generated on their digital properties. The need to manage and accurately categorize these assets can be overwhelming, if not impossible, to maintain and use. Bulk management features provided in 7.0 help reduce the amount of manual labor through automation and tools that perform bulk actions.

Auto-Tagging

Users can automatically assign the correct metadata for images, documents and web content. Having rich metadata assigned to the proper content helps to not only reduce the amount of manual, tedious work, but also improves the searchability of assets. The resulting metadata that accumulates helps establish the foundation needed for content personalization efforts and content automation. An Auto-Tagging API lets developers extend the functionality to tag any asset with any selected auto-tagging service.

Bulk Operations

Improvements have been made to make the experience of managing tags, categories, and file operations much easier at scale.

Automatic Document Versioning Policies

Refinements to how documents can be versioned are provided in 7.2. The functionality for document versioning has been expanded to let users define their own versioning policies and apply them to Liferay DXP. This is in addition to the existing default versioning system Liferay DXP provides out-of-the-box.

2.5 Business Operations

Digital enterprises today must possess efficient back office organization to support the delivery of great customer experiences. It is crucial to fulfill these customer interactions and ensure those experiences are delivered and supported in a timely manner. Achieving this kind of efficiency, for a great customer experience, requires involvement from different departments and stakeholders and affects a lot of disparate content and processes. 7.0 helps enterprises navigate these problems with features to help support streamlining and optimizing those processes and operations.

Forms API

Reduce the time it takes for IT to deliver custom applications for the business. Automate data collection through applications.

Workflow Metrics

Workflow metrics helps users gain insight into how long certain workflow events take to complete. Users can set deadlines on a workflow process's events; these configurations are called Service Level Agreements (SLAs). Once defined, use Workflow Reports to measure compliance with the SLAs. The analytics and metrics generated can help you understand the throughput performance of your processes in a given timeframe, allowing users to better optimize their processes.

Workflow Reports

Cut down the time it takes to serve your customers. New reports for workflow processes help users identify operational bottlenecks and work efficiently to increase revenue goals.

Online Document Editing

Integration with Office 365 and Google Docs allows the creation and editing of documents stored in Liferay DXP. Users can manage documents, presentations, and spreadsheets with the power that the suite provides and store them in Liferay DXP's document repository for future access. This feature takes advantage of the existing permissions system, versioning, and sharing capabilities already included with 7.0.

P2P Asset Sharing

In addition to being able to create and edit documents easily, users can now comment and share documents with other registered users. This allows for faster coordination without the need of assigning custom permissions.

Improved Identity Management Tools

New improvements have also been made to help tackle tedious tasks of handling personal data erasure requests. Logical hierarchies for assets are now in place to help simplify bulk deletion and anonymization during the review process. Filtering and scoping capabilities across Liferay DXP help provide greater context when reviewing personal data, and administrators can be informed when specific applications will delete or anonymize data during the auto-anonymization process. All of these improvements help ensure compliance with GDPR and other data protection laws.

2.6 Headless Capabilities

Give your developers the freedom to create any presentation layer and content creators the ease of traditional content management tools. 7.0 provides a decoupled architecture to allow for a range of approaches towards content management and delivery. Enterprises can choose a traditional CMS or experience management approach, a full headless approach or a hybrid approach.

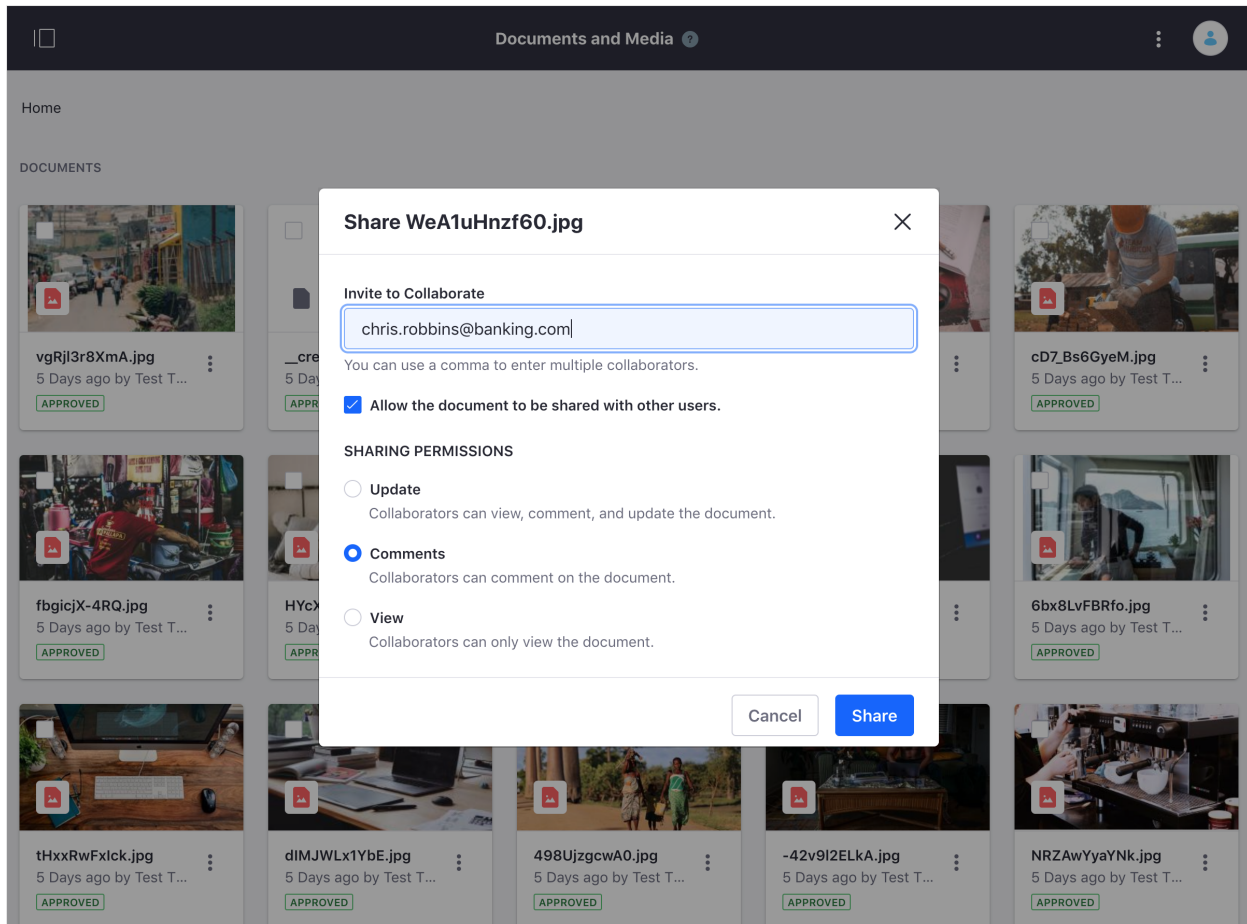


Figure 2.5: You can share assets with other users easily.

OpenAPI

Liferay's API layer supports the OpenAPI specification, the most popular open source framework for REST APIs. This allows for greater flexibility, security and ease of integration as the OpenAPI specification is language-agnostic, allows for security audits, and clients can understand and consume services without knowledge of server implementation or access to the server code.

Headless CMS

Enable developers to manage web content headlessly. Headless capabilities in this version are also provided to help enable omnichannel experiences from a single source of data. Content delivery APIs allow for developers to deliver richer, faster and more responsive user experiences, no matter the device. Front-end developers can leverage their native tools and frameworks they're already familiar with to build sites. These APIs provide access to a variety of assets and content in Liferay DXP.

Content participation APIs are also provided to support user interactions with comments and ratings for web content and assets.

2.7 Developer Experience

In addition to the new headless capabilities that are provided on 7.0, improvements to developer tools and the developer experience has been made. These tools help developers accelerate time to market and tailor the platform for your business needs.

Upgrade Tool

A revamped Upgrade Tool simplifies the upgrade process from previous versions of Liferay DXP. This all-in-one tool allows system administrators to reduce the time spent troubleshooting by helping them plan their upgrade process for database migrations, checking properties, and restarting failed upgrades.

Front-end Toolkits

Front-end toolkits allow developers to create an application using their favorite JavaScript framework. The toolkit resolves any conflicting JavaScript packages and bundles the application properly to deploy it into Liferay DXP seamlessly.

Struts Removed

Though Liferay DXP never used a version of Struts with security vulnerabilities, the Struts library has been removed, and it is no longer a part of the product.

2.8 What's Next

Learn more about how Liferay can help your business take the next step in your digital strategy. Request a demo from one of our team members at liferay.com/request-a-demo.

SETTING UP

If Liferay DXP is anything, it's configurable. As the core is shrinking due to its increased modularity, it's important that all the applications in Liferay are also configurable.

Breaking it down, three types of applications must be configurable:

1. The platform itself
2. Liferay's out-of-the-box applications
3. Custom applications

To this end, Liferay's engineers have made the platform and its applications configurable, and created a mechanism for developers to make their applications configurable.

WHERE CONFIGURATION HAPPENS

Liferay's configuration takes place in the following places:

User Interface: configuration through Liferay's UI is stored in a database. The values set in the UI always override configurations set in properties files.

Properties files: properties files that set default behavior may be included in the platform or the modules. Keep in mind that these settings can always be overridden by a system administrator in the UI. To find what properties are configurable this way, visit docs.liferay.com.

The UI location where these configuration options appear depends on the scope you want to affect with the settings you choose.

CONFIGURATION SCOPE

Depending on the configuration scope of a setting you change, you'll impact the platform and its applications with more or less granularity. At one end of the spectrum, you can affect the scope of the whole system. Configurations made at the system scope affect all virtual instances, Sites, and widgets. At the opposite end of the spectrum, configurations made at the widget level provide configuration settings only for that instance of the widget.

Take Language settings, for example. You can set the default language used by the virtual instance. You can also set the default language of a Site. Some applications even let you set the default language used by a specific piece of content.

Here's an overview of the available configuration scopes:

System: Configuring Liferay and its applications through System Settings provides system scoped configurations and sets default values for all virtual instances, sites, or widgets.

Virtual Instance: Configuring in Instance Settings provides settings that act on the specific virtual instance for which they are made, including Sites and widgets in the virtual instance.

Site: Configurations made at the Site scope, where you select the Site to configure in the Site selector, provide settings that take place only in that Site. Alternate configurations can be made in different Sites.

Widget Scope: Configuring a specific widget only provides a configuration for that particular widget.

Scopes in Liferay are hierarchical so that one scope can set the default values for the underlying sub-scopes. For example, making a system-wide configuration sets the default values for all virtual instances, sites or widgets of the system. If a different configuration is set at a level with more granularity (for example, the widget scope), it takes precedence over the settings made at less granular scopes (for example, the virtual instance scope).

This section contains articles on configuring Liferay at the System and Instance scopes:

System wide configuration:

- System Settings is the primary location for system configuration.
- Server Administration contains some lower-level server configuration options, such as logging.

Setting up a virtual instance:

- Virtual Instances is where virtual instances are added and edited.
- Instance Settings is the primary location for a virtual instance's configuration.
- Custom Fields is where additional database fields are added to existing virtual instance entities.

All of these are accessed through the Control Panel. Start by learning to configure modules system-wide in the System Settings Control Panel app.

SYSTEM WIDE SETTINGS

It can be hard to keep track of all the configuration interfaces. The Control Panel's Configuration section houses a lot of the higher level (for example, system and instance scoped) configuration options. This section considers the configuration options dealing with the *System* scope. Configuration at the system level affects all the *Virtual Instances* of Liferay in the system.

- System Settings is the primary location for system configuration.
- Server Administration contains some lower-level server configuration options, such as logging.

Get started by learning about System Settings.

SYSTEM SETTINGS

Liferay DXP is modular, meaning it's composed of many applications divided into even smaller “chunks” of functionality. The applications, and sometimes even code chunks, are configurable at several scopes, as discussed in the introductory article for this section.

In System Settings, administrators make system scoped configuration changes and set system-wide default configurations for other scopes. System Settings is located in Control Panel → Configuration → System Settings.

7.1 Editing System Configurations

System Settings is organized into sections (for example, Content) and categories (for example, Workflow) based on the functionality being configured. There's also a Search bar to make finding configuration entries easier. Search for the name of a specific configuration entry, or even a specific field within an entry.

Changing a configuration isn't difficult:

1. Find the configuration entry you need, either by searching or browsing the sections and categories.
2. Open the configuration form for the entry.
3. Make any changes you'd like, then click *Save*. Your configuration changes are saved and applied throughout the system.

Important: Content generated using templates (e.g., FreeMarker templates and Application Display Templates) is cached. Cached content might not reflect configuration changes until the cache is invalidated (cleared). The Server Administration → Resources tab provides cache clearing options.


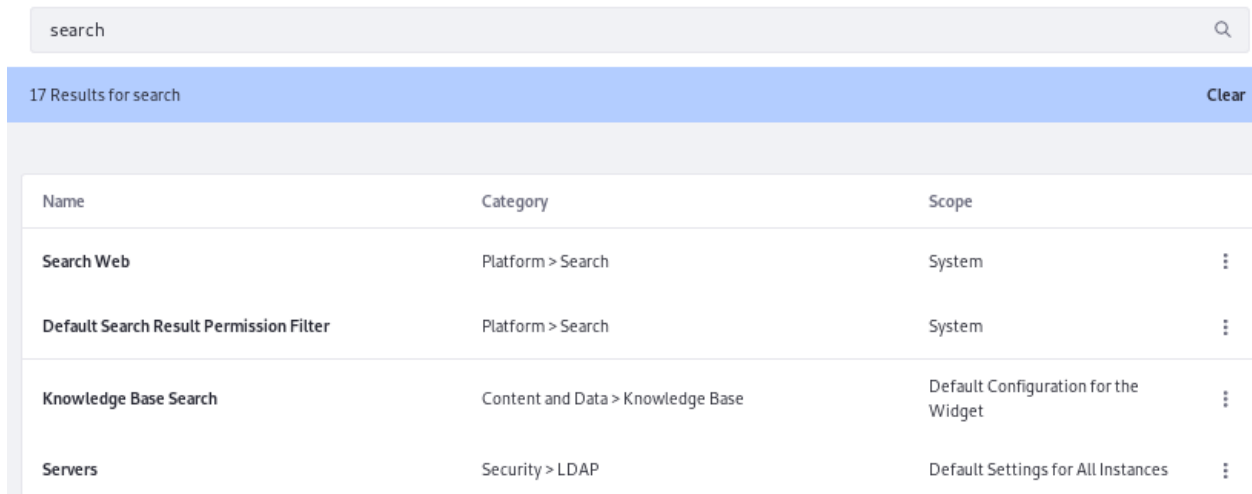


Figure 7.1: System Settings are accessed through the Control Panel.



| Name | Category | Scope |
|---|-----------------------------------|--------------------------------------|
| Search Web | Platform > Search | System |
| Default Search Result Permission Filter | Platform > Search | System |
| Knowledge Base Search | Content and Data > Knowledge Base | Default Configuration for the Widget |
| Servers | Security > LDAP | Default Settings for All Instances |

Figure 7.2: System Settings are organized by section and category.





Figure 7.3: After saving changes to a configuration, the actions *Reset Default Values* and *Export* are available for it.

If you make some regrettable configuration decisions and can't recall exactly what you did, start over by clicking the actions button (), then clicking *Reset Default Values*.

7.2 Configuration Scope

While browsing the categories of System Settings, you'll notice that clicking into a topic (for example, Blogs) reveals entries at different scopes. All the settings here act at the system scope. For scopes labeled other than System, these configurations act as defaults. In other words, they identify where the system-wide configuration is overridden. True system-scoped configurations (those under a

category's *System Scope* heading) are not overridden anywhere. There are four values that you'll see under Scope:

- *System Scope*: Any System Settings configuration made for system scoped entries becomes the final value for the application in a system-wide fashion. It affects the whole system and isn't overridden anywhere else.




Figure 7.4: Some System Settings entries are system scoped.

- *Virtual Instance Scope*: Configuration at the Virtual Instance level is overridden in Instance Settings.




Figure 7.5: Some System Settings are virtual instance scoped.

- *Site Scope*: Configuration at this scope is overridden in each site.




Figure 7.6: Some System Settings are site scoped.

- *Widget Scope*: Configuration at this scope is overridden in each Widget Instance (like the Blogs example below).

If a configuration changed in System Settings is also configurable at a different scope, the System Settings value acts as a default that can be overridden. Once a configuration change is made at a more granular scope, making a change at the system level doesn't do anything.

For example, allowing comments is configurable for each Blogs Entry. Set the default behavior at Control Panel → Configuration → System Settings → Blogs. In the Blogs Entry under Widget Scope, disable the *Enable Comments* checkbox.

Now add a Blog Entry to a Site's Content & Data → Blogs application. Then go to a public page and add the Blogs Widget to the page. Click the Options button (⋮) for the widget and select *Configuration*. You'll see the same Enable Comments checkbox, and its default is now false

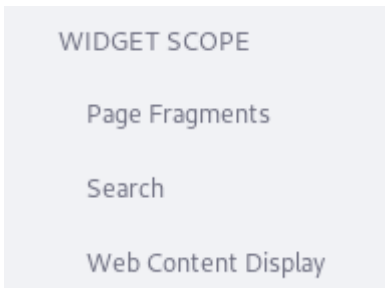



Figure 7.7: Some System Settings entries are widget scoped.

(unchecked). Checking the box in the Widget Configuration screen breaks its link with the System Settings entry. Changing the System Settings configuration has no effect on this widget anymore.

7.3 Exporting and Importing Configurations

What if you change many default configurations in System Settings, and then need to make the same changes in another installation? Don't worry, you don't need to remember and duplicate every choice you made. The System Settings application lets you export a single entry's configurations, or all the settings you made in the System Settings interface. The exported files are deployable to any other installation of the same version.

To export a single entry's configurations, click the actions button (), then click *Export*. A .config file containing your configuration downloads to your system.

To export all the configuration changes you've made in System Settings, click the System Settings options button (), then click *Export All Settings*. The .config files for all the entries you edited then download in a ZIP file.

To make these configurations active in the destination system, unzip and place the .config files in the [Liferay_Home]/osgi/configs folder.

Now you know what System Settings is and how to use it. All that's left is to explore the entries to see what configuration options are available. If you aren't sure what something does, check the documentation for the feature you're interested in, as specific configurations are covered there.

UNDERSTANDING SYSTEM CONFIGURATION FILES

The System Settings application is convenient for making system-scoped configuration changes and setting default configurations for other scopes. But there's another supported configuration approach: configuration files. You can use configuration files to transfer configurations from pre-production systems to production systems, or between any other Liferay DXP systems. Sometimes developers choose to distribute the default configuration for their applications via configuration file. Whatever the reason, configuration files offer another configuration approach.

Configuration files use the `.config` property value format defined by the Apache Felix Configuration Admin framework.

Important: Content generated using templates (e.g., FreeMarker templates and Application Display Templates) is cached. Cached content might not reflect configuration changes until the cache is invalidated (cleared). The Server Administration → Resources tab provides cache clearing options.

Note: The `.cfg` file format is common in OSGi environments, and it's a supported format, but `.config` files are preferable since they allow specifying a property value's type, and allow multi-valued properties. The syntax described in these articles is for `.config` files.

CREATING CONFIGURATION FILES

System Settings provides an *Export* option that becomes available once you modify a configuration entry. Exporting is the recommended way to create `.config` files: you download a `.config` file containing the entry's settings in a key-value format. Liferay DXP exports an entry's total available configuration keys and values, even if only one value was changed. You can export a single configuration entry or the entire set of modified configurations.

To avoid a file name conflict, name configuration files using a unique identifier. For example, the Journal Service entry, which backs Web Content functionality, has this file name:

```
com.liferay.journal.configuration.JournalServiceConfiguration.config
```




Figure 9.1: The Web Content System Settings entry has the back-end ID `com.liferay.journal.configuration.JournalServiceConfiguration`.

9.1 Key/Value Syntax

The syntax for all keys and values in a `.config` file is the same:

```
configurationName="value"
```

For single value configurations without special characters, that's all there is to know. Settings with multiple values and certain characters require slight modifications.

9.2 Multi-Value Settings

Configuration entries can have properties that accept multiple values. For example, a configuration property for specifying supported file extensions needs more than one value. Here's how to write a multi-value setting in a `.config` file:

```
multiValueSetting=["Value 1","Value 2", ...]
```

Do not use a space character between values (after the comma). The property won't be loaded.

Open the Web Content category in System Settings (under the Content section), and select *Web Content* for the virtual instance scope. You'll see what looks like multiple single value entries for *Characters Blacklist*:

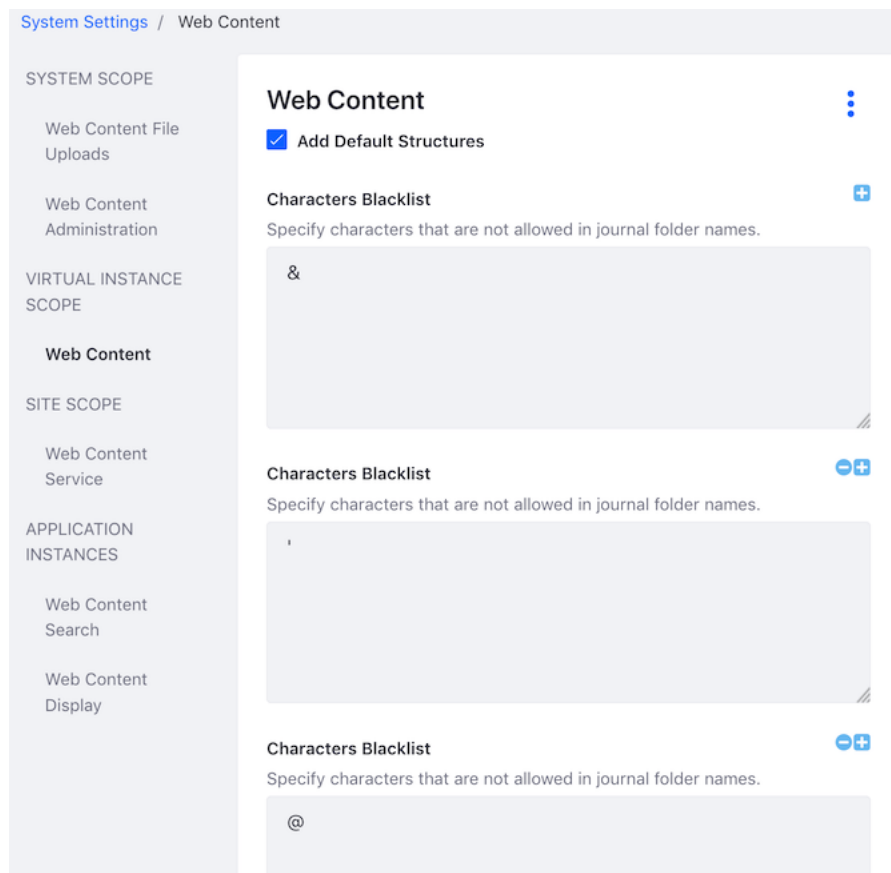


Figure 9.2: The Web Content System Settings entry has many *Characters Blacklist* fields.

In the configuration file, this is really a single key with an array of comma-separated values:

```
charactersblacklist=["&",".","@","\\"","}",":","\=",">","/","<","[","{","%","+","#","\","?","\",";","*","~"]
```

9.3 Escaping Characters

Double quotes (") and equals signs (=) must be *escaped* in .config files. Escaping is using another character to denote that a character shouldn't be used in its normal way. Since double quotes and equals signs are already used in .config files, escaping them tells the framework not to read them the normal way, but to pass them through as part of the value. Use a backslash to escape characters in the .config file:

```
charactersblacklist=["&","\"","\\="]
```

This setting illustrates a multi-value setting with a regular, unescaped character (&), and two escaped ones (\" and \=).

Along with the mandatory escaping of double quotes and equals characters, it's beneficial to escape spaces inside values to avoid problems. In this example, a backslash is used before each space character to ensure it's read and processed properly:

```
blacklistBundleSymbolicNames=["Liferay\ Marketplace","Liferay\ Sharepoint\ Connector"]
```

If you don't escape spaces yourself, the framework adds the backslash for you after deployment.

9.4 Typed Values

The .config file format supports specifying the type of a configuration value by inserting a special type marker character. Because Liferay DXP already knows the correct type for each configuration property, the type characters are only useful for informational purposes. For example, a configuration with a boolean type has *B* just before the value to mark it as a boolean type:

```
addDefaultStructures=B"true"
```

If you see type markers in .config files, you can safely ignore them. The example included above functions identically without the type marker:

```
addDefaultStructures="true"
```

9.5 Deploying a Configuration File

Once you have a configuration file, deploy it. It's registered and the targeted configuration values are updated automatically.

To deploy the .config file, place it in your Liferay Home's `osgi/configs` folder. To change the configuration further, you can edit the .config file directly or use System Settings.

9.6 Configuration Files and Clustering

In a clustered environment, each node needs the same configuration values for each entry. For example, all nodes should use the same *Blogs* configuration settings. To accomplish this, deploy copies of the same `.config` file to each node's `[Liferay Home]/osgi/configs` folder.

—

```
header-id: factory-configurations
```

```
Factory Configurations
```

```
[TOC levels=1-4]
```

Configurations supporting multiple entries are called factory configurations.

Factory Configuration Example: JAX-WS and JAX-RS web services are supported. These services must use a CXF Endpoint, which is a context path where the web services are deployed and accessed. Endpoints can be created via factory configuration by navigating to the CXF Endpoints entry in System Settings (System Settings → Platform → Web API → CXF Endpoints). Click ADD, enter the desired configuration values, then repeat the process to add as many CXF Endpoint configurations as needed. Creating CXF Endpoint configurations also creates CXF Endpoints themselves. This is how factory configurations work.

If your service supports factory configurations, use the convention of calling the configuration's first instance `-default.config`:

The next instance contains a unique subname (something other than default). It's good practice to use a descriptive name:

To follow the CXF Endpoints example described above, if Liferay's developers had shipped an initial CXF Endpoint `.config` file with Liferay DXP, it would have been named this:

If this `-default.config` configuration specifies a context path for REST web services, and you create another endpoint with a different context path for SOAP web services, your second configuration file could be named:

Note: Some System Settings entries (like the CXF Endpoints entry) don't ship with a configuration file, so anything you create is the first occurrence. However, if you configure one and export it to obtain the `.config` file, it doesn't use the `-default.config` naming convention. Instead it's given a guaranteed unique identifier for its subname, like this: `com.liferay.portal.remote.cxf.common.configuration.CXFEndpointPublisherConfiguration-a6f67e48-6dca-49c6-bf6b-8fd5e6016b2d.config` | This guarantees that the file has a unique name. If you're exporting the configuration file for deployment in a separate system, you can rename the exported file to use a more descriptive subname. If you rename the file and deploy it to the same system it was exported from, the new subname marks it as an entirely new configuration. You'll end up with an additional configuration instance in this case, not just a renamed one.

FACTORY CONFIGURATIONS

Configurations supporting multiple entries are called *factory configurations*.

Factory Configuration Example: Adding Organization types is supported, and is useful if you need to model real-life hierarchies or enforce hierarchical rules. In Liferay DXP, each Organization type is created via a factory configuration entry in System Settings.

If a service is meant to support factory configurations, its System Settings entry has an ADD button.



Figure 10.1: If a System Settings entry has an ADD button, it's suitable for factory configurations.

As with single-instance configurations, you can set factory configurations in the System Settings interface (as described in the example above) or via configuration files. Name a standard single-instance configuration file like this:

```
my.service.ServiceConfiguration.config
```

If your service supports factory configurations, use the convention of calling the configuration's first instance `-default.config`:

```
my.service.ServiceConfiguration-default.config
```

The next instance contains a unique *subtype* (something other than *default*). It's good practice to use a descriptive name:

```
my.service.ServiceConfiguration-port9080.config
```

In the Organization type example, the default Organization type (aptly named *organization*) is created by a `-default.config` file named

```
com.liferay.organizations.internal.configuration.OrganizationTypeConfiguration-default.config
```

Following the example from the Adding a New Organization Type article, you could add the *League* type with a configuration file named

```
com.liferay.organizations.internal.configuration.OrganizationTypeConfiguration-league.config
```

Some System Settings entries that support factory configuration don't ship with a configuration file for the default instance (e.g., the Anonymous User entry). If you export a factory configuration file to obtain the `.config` file, it doesn't use the `-default.config` naming convention. Instead, whether it's the first occurrence or an additional one, it's given a guaranteed unique identifier for its subtype:

```
com.liferay.user.associated.data.web.internal.configuration.AnonymousUserConfiguration-6befcd73-7c8b-4597-b396-a18f64f8c308.config
```

This guarantees that the file has a unique name. If you're exporting the configuration file for deployment in a separate system, you can rename the exported file to use a more descriptive subtype. If you rename the file and deploy it to the same system it was exported from, the new subtype marks it as an entirely new configuration. You'll end up with an additional configuration instance in this case, not just a renamed one.

Warning:: For configuration entries supporting factory configurations, omitting the subtype from a `.config` file's name causes System Settings to disallow adding new entries for the configuration (only the configuration entry targeted by this `.config` file). This is caused by a known bug. See LPS-76352 for more information. Once an improperly named configuration file is deployed, you can't add any entries for the configuration in question from its System Settings entry. For example, if you deploy the following file to configure an Organization Type, not only does this not add an Organization Type, it also prevents you from adding any via System Settings:

```
com.liferay.organizations.internal.configuration.OrganizationTypeConfiguration.config
```

Deploying an erroneous (lacking a subtype) `.config` file doesn't disable anything permanently. Just rename the file using the proper convention described above or remove it entirely and start over.

In many cases, configuration files can be used to force a factory configuration scenario, but not all configurations can be used this way. It's best to stick to the intended use cases. Use System Settings as described above to determine if using factory configurations is a good idea. If not, stick to the single occurrence mode of configuration (specifying only one configuration file for the service).

—

header-id: server-administration
Server Administration

[TOC levels=1-4]

Server Administration lets you manage and monitor your Liferay DXP server. Access the application by clicking Control Panel → Configuration → Server Administration.

Figure 1: The Resources tab of Server Administration shows a graph of your server's memory usage.

Server Administration's functionality is segmented into these tabs:

Resources: View memory usage and perform management tasks like running the garbage collector, clearing the database cache, and more. For more information, see Resources.

Log Levels: View and set logging levels. You can make dynamic modifications of log levels for any class hierarchy in Liferay DXP. Custom objects not on the list can be added with the Add Category tab. Changes to the log level near the top of the class hierarchy (such as at com.liferay) also change log levels for all the classes under that hierarchy. Modifications unnecessarily high in the hierarchy generate too many messages to be useful.

Properties: View JVM and portal properties. This tab has two sub-tabs: System Properties and Portal Properties. The System Properties tab shows an exhaustive list of system properties for the JVM, as well as many

Create a backup copy of the Document Library repository and Liferay DXP database.

Configure the new file store in System Settings → Platform: File Storage.

In this tab (Server Administration → Data Migration), select the repository hook for the file store you configured and click Execute.

Make sure the data migrated correctly.

Configure the new repository as the default.

If you used a portal-ext.properties file to configure the repository, restart the server.

Mail: Instead of using a portal-ext.properties file to configure a mail server, you can configure a mail server from this tab. If your message boards receive mail, you can connect a POP mail server. If Liferay DXP sends mail (useful for sending notifications to users), you can connect to an SMTP server. Note that if you configure mail server settings here in System Settings, these settings override any mail server settings in your portal-ext.properties file.

External Services: Configure external services for generating file previews. For more information, see the article on External Services.

Script: A scripting console for executing migration or management code. The Groovy scripting language is supported out of the box. .

Shutdown: Schedule a shutdown that notifies logged-in users of the impending shutdown. You can define the number of minutes until the shutdown and a message to display. Liferay DXP displays the message at the top of users' pages for the duration of time you specified. When the time expires, all pages display a message saying the portal has been shut down. The server must then be restarted to restore access.

SERVER ADMINISTRATION

Server Administration lets you manage and monitor your Liferay DXP server. Access the application by clicking *Control Panel* → *Configuration* → *Server Administration*.




Figure 11.1: The Resources tab of Server Administration shows a graph of your server's memory usage.

Server Administration's functionality is segmented into these tabs:

Resources: View memory usage and perform management tasks like running the garbage collector, clearing the database cache, and more. For more information, see Resources.

Log Levels: View and set logging levels. You can make dynamic modifications of log levels for any class hierarchy in Liferay DXP. Custom objects not on the list can be added with the *Add Category* tab. Changes to the log level near the top of the class hierarchy (such as at `com.liferay`) also change log levels for all the classes under that hierarchy. Modifications unnecessarily high in the hierarchy generate too many messages to be useful.

Properties: View JVM and portal properties. This tab has two sub-tabs: System Properties and Portal Properties. The System Properties tab shows an exhaustive list of system properties for the JVM, as well as many Liferay DXP system properties. You can use this information for debugging purposes or to check the currently running configuration. The Portal Properties tab shows an exhaustive list of the current portal property values. For explanations of these properties, see the portal properties reference documentation.

Data Migration: Migrate documents from one repository to another. For example, you can migrate your documents to a new repository on a different disk or in a new format. Here are the

steps:

1. Create a backup copy of the Document Library repository and Liferay DXP database.
2. Configure the new file store in *System Settings* → *Platform: File Storage*.
3. In this tab (*Server Administration* → *Data Migration*), select the repository hook for the file store you configured and click *Execute*.
4. Make sure the data migrated correctly.
5. Configure the new repository as the default.
6. If you used a `portal-ext.properties` file to configure the repository, restart the server.

Mail: Instead of using a `portal-ext.properties` file to configure a mail server, you can configure a mail server from this tab. If your message boards receive mail, you can connect a POP mail server. If Liferay DXP sends mail (useful for sending notifications to users), you can connect to an SMTP server. Note that if you configure mail server settings here in System Settings, these settings override any mail server settings in your `portal-ext.properties` file.

External Services: Configure external services for generating file previews. For more information, see the article on External Services.

Script: A scripting console for executing migration or management code. The Groovy scripting language is supported out of the box.

Shutdown: Schedule a shutdown that notifies logged-in users of the impending shutdown. You can define the number of minutes until the shutdown and a message to display. Liferay DXP displays the message at the top of users' pages for the duration of time you specified. When the time expires, all pages display a message saying the portal has been shut down. The server must then be restarted to restore access.

SERVER ADMINISTRATION: RESOURCES

The Server Administration app's Resources tab contains several server wide actions that an administrator can execute. These include the following items:

Run the garbage collector: Tells the JVM to free memory by running the garbage collector.

Generate a thread dump: Generates a thread dump for later scrutiny to determine the presence and location of any deadlocks. Useful during performance testing, but you must add a logger category for `com.liferay.server.admin.web.internal.portlet.action.EditServerMVCActionCommand` and set it to `INFO` before executing.

Clear content cached by this VM: Clears content stored in the local cache. Only local JVM scope Ehcache content is cleared, not clustered Ehcache. 1

Clear content cached across the cluster: Clears the content of the entire clustered Ehcache. 1

Clear the database cache: Clears the database cache. Does not clear any Ehcache content except database results at the persistence layer. 1

Clear the direct servlet cache: Clears the direct servlet cache. In case emergency fixes must be applied, this action allows an administrator to clear out the cache manually to force JSPs to reload.

The direct servlet context optimizes JSP serving performance by caching and accessing the generated servlets directly instead of accessing them over the application server's dispatcher chain. This function is only suitable for cases where no filter is required for the JSPs; it should be enabled for production mode to improve performance, but disabled for development mode to allow JSP servlets to be reloaded on the fly. See the Direct Servlet Context section of the `portal.properties` file for details. 1

Verify database tables of all plugins: Checks all tables against their indexes for data retrieval accuracy.

Verify membership policies: Checks that existing Site membership policies were correctly applied and automatically makes updates. If the Liferay DXP database is changed manually or is hacked—resulting in a user assigned to a Site in violation of a site membership policy—this action triggers the verification methods of all implemented Site membership policies. Changes are automatically made to any memberships in violation.

Reset preview and thumbnail files for Documents and Media: Regenerates previews of each item in your Documents and Media libraries.

Clean up permissions: Removes permissions on the Guest, User, and Power User Roles to simplify the management of User Customizable Pages. The Add To Page permission is removed

from the Guest and User Roles for all portlets, and is reduced in scope for Power Users from portal-wide to User Personal Site.

Clean up portlet preferences: This action cleans up database entries if portlet preferences become orphaned in the Liferay DXP database.

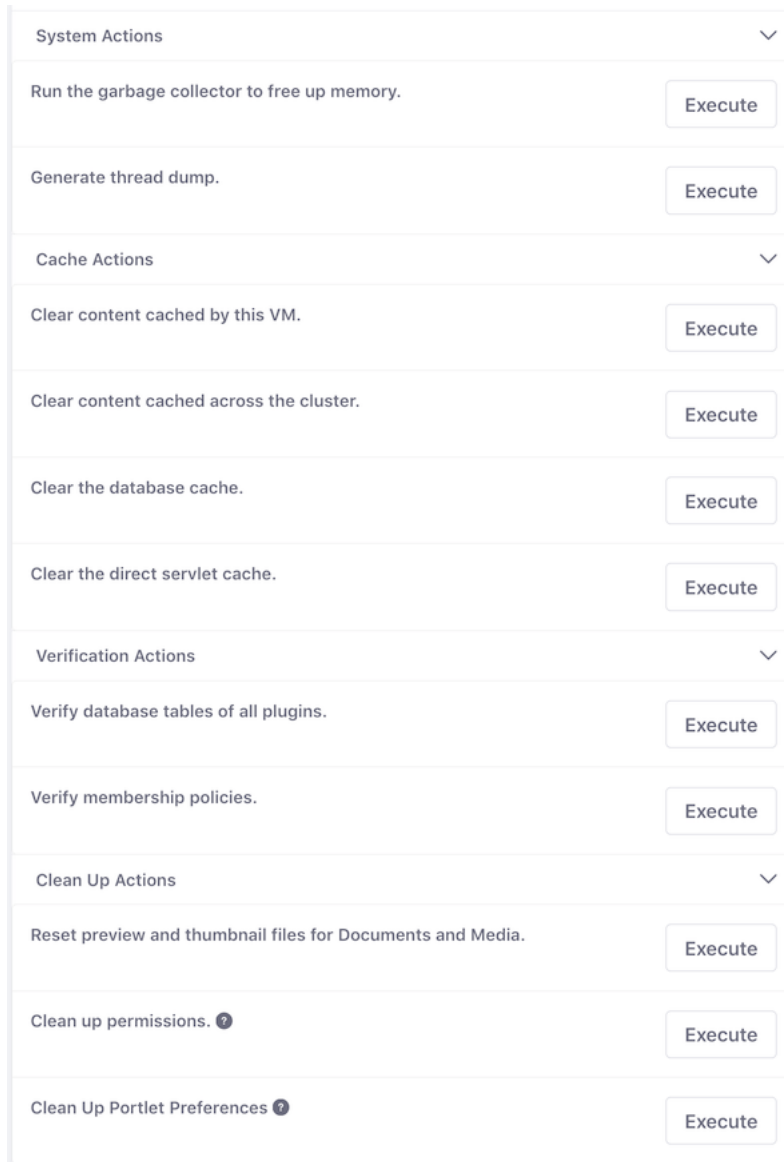


Figure 12.1: The Resources tab of Server Administration lets you execute several server maintenance tasks.

[1] Caching occurs at multiple levels. Some higher caching layers aren't aware of lower caching layers. Always clear the cache at the lowest (most granular) layer possible, even if you've already cleared a higher level cache.

SERVER ADMINISTRATION: EXTERNAL SERVICES

Users can upload and share any type of file via the Documents and Media library, a customizable and permissions-enabled online repository for files (see publishing files for more information). PDFBox is included with Liferay DXP and generates automatic previews for certain file types (mostly PDFs). You can install two additional tools to generate previews for other file types:

- **OpenOffice:** or **LibreOffice:** Convert and generate previews for many types of documents.
- **ImageMagick:** Generate higher-quality image previews for many types of images.
- **Xuggler:** Convert and generate previews for audio and video files.

As of Liferay 7.1, OpenOffice/LibreOffice is configured in OSGi Configuration Admin instead of portal properties. To adjust these settings, go to *Control Panel* → *Configuration* → *System Settings* → *Connectors* → *OpenOffice Integration*. You can also adjust these settings by deploying a `com.liferay.document.library.document.conversion.internal.configuration.OpenOfficeConfiguration.config` file to your `[Liferay Home]/osgi/configs` folder. See the Breaking Changes document for details.

Once you've installed ImageMagick and Xuggler, you can use the Server Administration app's External Services tab to configure Liferay DXP to use them. Make sure to choose the correct versions of these tools for your operating system. Install the latest stable versions, as older versions may not run properly with Liferay DXP. ImageMagick must be installed manually, but you can install Xuggler from the Control Panel.

Tip: If you're running Liferay DXP on a Linux server and experience a problem enabling Xuggler, check your server's glibc version. For Xuggler to work, you may need to update glibc to version 2.6 or later.

13.1 ImageMagick Configuration

Before configuring ImageMagick to generate image and PDF previews, install it and its dependency, Ghostscript. This differs by operating system: on Linux, both are likely already installed. They are not likely to be installed on Windows, but may be on macOS.

1. Download and install ImageMagick.
2. Download and install Ghostscript.

Once installed, you must enable ImageMagick in the Server Administration app's External Services tab, or in a `portal-ext.properties` file. If using `portal-ext.properties`, add the following lines and make sure the search path points to the directories containing the ImageMagick and Ghostscript executables. You may also need to configure the path for fonts used by Ghostscript when in macOS or Unix environments.

```
imagemagick.enabled=true
imagemagick.global.search.path[apple]=/opt/local/bin:/opt/local/share/ghostscript/fonts:/opt/local/share/fonts/urw-fonts
imagemagick.global.search.path[unix]=/usr/local/bin:/usr/local/share/ghostscript/fonts:/usr/local/share/fonts/urw-fonts
imagemagick.global.search.path[windows]=C:\\Program Files\\ImageMagick
```

Follow these steps to instead enable ImageMagick from the Server Administration app's External Services tab:

1. In the *Control Panel*, navigate to *Configuration* → *Server Administration*, then click the *External Services* tab.
2. Expand the ImageMagick and Ghostscript section and select *Enabled*.
3. Verify that the paths to the ImageMagick and Ghostscript executables are correct.

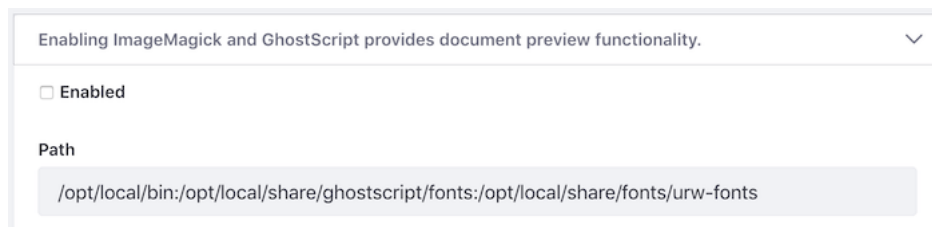


Figure 13.1: Enable ImageMagick and Ghostscript, and verify that the paths are correct.

4. Set the Resource limits to use.

13.2 Xuggler Configuration

Follow these steps to install and configure Xuggler:

1. In the *Control Panel*, navigate to *Configuration* → *Server Administration*, then click the *External Services* tab.
2. In the Xuggler section, select the Xuggler .jar file that matches your operating system. Then click *Install*.
3. Restart your application server.
4. Enable Xuggler for your portal. There are two ways to do this:
 - In the Control Panel, navigate to the *Server Administration* → *External Services* tab, select *Enabled*, then click *Save*.
 - Add the following line to your `portal-ext.properties` file and restart your application server:

```
xuggler.enabled=true
```

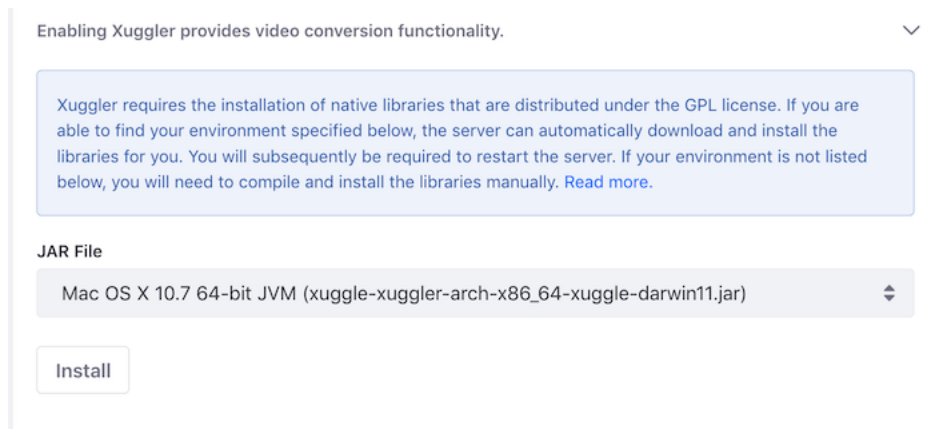


Figure 13.2: Install Xuggler.

SETTING UP A VIRTUAL INSTANCE

Once Liferay DXP is installed, the configuration begins. Recall that configuration happens at different scopes. Here we're covering configuration at the virtual instance scope. There's an important difference between the system scope and the instance scope. The system scope is the highest level scope you can make configurations at. All virtual instances are impacted by configuration done at this scope. The instance scope applies only to one particular virtual instance.

Virtual instances have unique domain names but share a server and database. Each virtual instance can have independent data and configurations.

The articles in this section cover these topics:

- Adding a virtual instance
- Configuring a virtual instance

Get started by learning how to add a virtual instance.

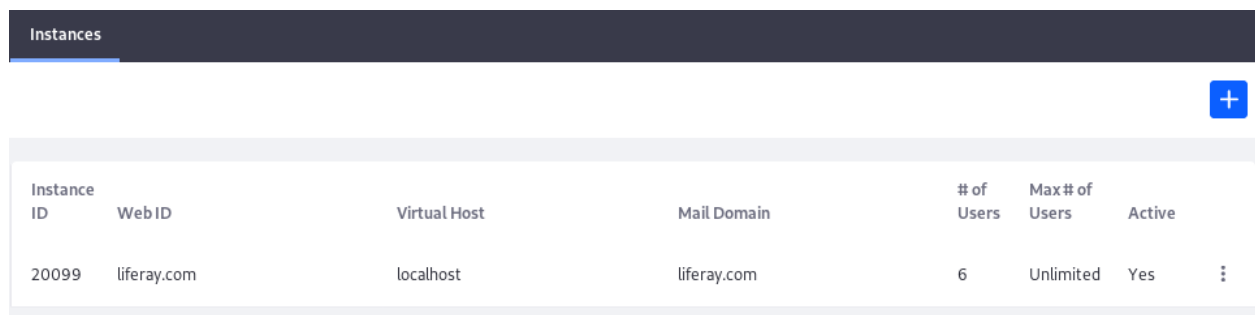
VIRTUAL INSTANCES

Here's a quick scenario: you already have a server hosting a Liferay DXP installation and a database. It has many Users, Sites, and specific instance settings. If you require a second similar installation, then adding a *Virtual Instance* might be right for you.

You can run more than one Virtual Instance on a single server with a shared database, but separate data and configurations. Users are directed to the correct Virtual Instance via its unique domain name. And because Virtual Instances share an application server and OSGi container, they also share these customizations:

- Custom code deployed by developers and administrators.
- System-scoped configurations (e.g., .config files, changes made in *Control Panel* → *Configuration* → *System Settings*).
- Application server configuration.

Administrators can manage Virtual Instances in *Control Panel* → *Configuration* → *System Settings*.




| Instance ID | Web ID | Virtual Host | Mail Domain | # of Users | Max # of Users | Active |
|-------------|-------------|--------------|-------------|------------|----------------|--------|
| 20099 | liferay.com | localhost | liferay.com | 6 | Unlimited | Yes |

Figure 15.1: Add and manage virtual instances of Liferay in the Control Panel's *Configuration* → *Virtual Instances* section.

Follow these steps to create a Virtual Instance:

1. Before you create a Virtual Instance, configure its domain name in your network.
2. Go to *Control Panel* → *Configuration* → *Virtual Instances*.

3. Click the *Add* button () . This opens the *New Instance* form.
4. Complete the *New Instance* form as follows:

Web ID: The instance's ID. Using the domain name is a common convention.

Virtual Host: The domain name you configured in your network. When users are directed to your server via this domain name, they'll be sent to the Virtual Instance that contains their data.

Mail Domain: The mail host's domain name for the Virtual Instance. Email notifications are sent from the instance using this domain.

Max Users: The maximum number of user accounts the Virtual Instance supports. Enter *0* to support unlimited users.

Active: Whether the Virtual Instance is active. Note that inactive Virtual Instances aren't accessible to anyone, even the administrator.

Click *Save* when you're done filling out the form.

Now you can navigate to the instance using its domain name. You're brought to what looks like a clean install of Liferay DXP. This is your new Virtual Instance! You can configure it any way you like. The remaining articles in this section show you how to configure an instance's settings.

CONFIGURING VIRTUAL INSTANCES

To access instance settings, open the Control Panel and navigate to *Configuration* → *Instance Settings*. The Instance Settings are organized into three sections: PLATFORM, SECURITY, and CONTENT AND DATA. Here we focus on the instance settings available under the PLATFORM section. The CONTENT AND DATA settings are specific to each application, and the SECURITY instance settings are covered in their respective articles, listed below for reference:

- LDAP
- OAuth 2
- SSO:
 - CAS Server
 - Token based SSO
 - Open ID Connect
 - OpenSSO

The PLATFORM instance settings are covered in the articles shown below:

- Email
- Instance Configuration
- User Authentication
- Users
- More Platform Settings

EMAIL INSTANCE SETTINGS

The Email configuration is your one-stop-shop for all your email notification needs. To access the Email configuration settings for your instance, open the Control Panel and navigate to *Configuration* → *Instance Settings* and select the *Email* category under the PLATFORM section. The Email configuration contains six entries:

- Account Created Notification
- Email Sender
- Email Verification Notification
- Mail Host Names
- Password Changed Notification
- Password Reset Notification

Each configuration entry is described in the corresponding section below.

17.1 Account Created Notification

The Account Created Notification entry defines the email templates, with and without the new User's password included in the body of the text, for the message sent to notify Users when they create a new account. You can specify whether this email is sent by checking/unchecking the *Enabled* checkbox at the top. Use the template variables listed at the bottom of the configuration under the "Definition of Terms" heading to help build your email template.

17.2 Email Sender

The Email Sender entry specifies the virtual instance’s administrative Name and Address for email notifications, declared as the `[$FROM_NAME$]` and `[$FROM_ADDRESS$]` variables respectively in the email templates. By default, they are from the `admin.email.from.name` and `admin.email.from.address` portal properties). This name and email address appear in the *From* field in all email messages sent by the virtual instance.

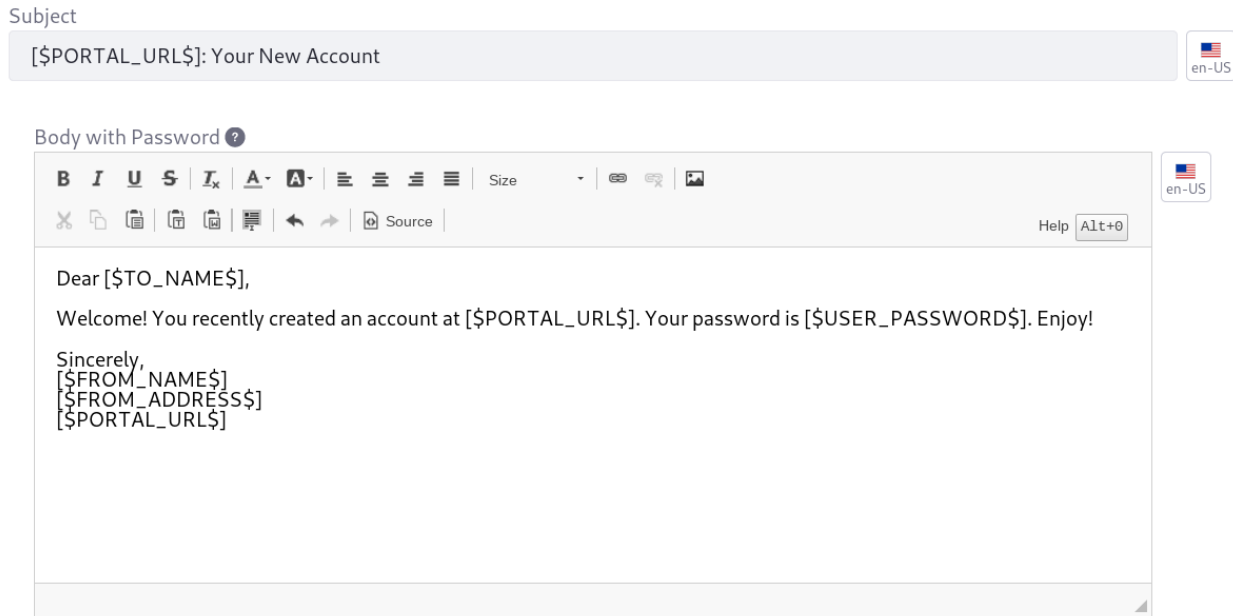


Figure 17.1: Customize the email template for the email messages sent to new Users.

17.3 Email Verification Notification

The Email Verification Notification entry defines the email template for the message sent to Users when asked to verify their Email Address. Use the template variables listed at the bottom of the configuration under the “Definition of Terms” heading to help build your email template.

17.4 Mail Host Names

The Mail Host Names entry specifies which mail host names are owned by your organization for the virtual instance. Enter one mail host name per line, besides the one specified in the General

tab.

17.5 Password Changed Notification

The Password Changed Notification entry defines the email template for the message sent to notify Users when their password has changed. Use the template variables listed at the bottom of the configuration under the “Definition of Terms” heading to help build your email template.

17.6 Password Reset Notification

The Password Reset Notification entry defines the email template for the message sent to notify Users when a request has been made to reset their password. Use the template variables listed at the bottom of the configuration under the “Definition of Terms” heading to help build your email template.

Definition of Terms

[\$FROM_ADDRESS\$]test@liferay.com

[\$FROM_NAME\$]Test Test

[\$PORTAL_URL\$]http://localhost:8080

[\$TO_ADDRESS\$]The address of the email recipient

[\$TO_NAME\$]The name of the email recipient

[\$USER_ID\$]The user ID

[\$USER_PASSWORD\$]The user password

[\$USER_SCREENNAME\$]The user screen name

Figure 17.2: There are some handy variables available for use in email templates.

INSTANCE CONFIGURATION INSTANCE SETTINGS

The Instance Configuration settings define the basic configuration information for the virtual instance, from the appearance to the Terms of Use for your Users to agree to. Access the Instance Configuration settings from the Control Panel's *Configuration* → *Instance Settings* section, and select the *Instance Configuration* category under the *PLATFORM* section. The Instance Configuration contains four entries:

- Appearance
- Contact Information
- General
- Terms of Use

Each configuration entry is described in the corresponding section below.

18.1 Appearance

The Appearance configuration entry defines the default logo and overall look and feel for the virtual instance. It's organized into LOGO and LOOK AND FEEL sections:

LOGO: Change the default logo and check/uncheck the *Allow site administrators to use their own logo?* checkbox, enabled by default, to specify whether site administrators can upload a logo when they configure a site. When configuring a new logo, be careful to choose an image file that fits the space. Large images might overlap with the navigation.

LOOK AND FEEL: Set the default theme(s) for the instance and Control Panel.

18.2 Contact Information

The Contact Information configuration entry specifies how to contact the organization that owns the virtual instance. It's divided into several sections:

ADDRESSES: Specify the primary, mailing, shipping, P.O. Box, etc. address of the organization.

PHONE NUMBERS: Provide the fax, local, etc. phone numbers for the organization.

ADDITIONAL EMAIL ADDRESSES: Specify any additional email addresses associated with the organization.

WEBSITES: Specify the public and/or intranet websites for the organization.

Developers can query for this contact information in their applications. Note that you can add and remove additional entries in a section with the plus and minus buttons respectively.

18.3 General

The General entry specifies the virtual instance's configuration information, landing pages, and the associated organization's basic legal information. It has two sections:

Main Configuration: Configure this information for the virtual instance:

- Set the name of the entity responsible for running the virtual instance.
- Set the mail domain.
- Set the virtual host.
- Set the URLs to a CDN configured to serve static resources.

Navigation: Set a home page for your virtual instance as well as default landing and logout pages. To set these pages, use the page's relative URL that follows the domain. For example, to set the default landing page to `http://localhost:8080/web/guest/login`, use `/web/guest/login`. You can use the variable `${liferay:screenName}` as part of the address to redirect Users to their personal pages upon login. For example, the User `marvin` has this default URL to his personal page:

```
http://localhost:8080/user/marvin
```

To make sure he's directed there on login, place `/user/${liferay:screenName}` in the Default Landing Page field. These URLs can also be set at the system scope, in a `portal-ext.properties` file:

```
default.landing.page.path=
default.logout.page.path=
company.default.home.url=
```

All virtual instances share the values specified in the properties file. Changes made in Instance Settings override the values set in the properties file. For more information, see the Portal Properties documentation.

Additional Information: Specify a Legal name, ID, company type, SIC code, ticker symbol, industry and industry type for the owner of the virtual instance.

18.4 Terms of Use

The Terms of Use entry contains everything you need to provide a custom Terms of Use agreement for your Users. Terms of Use are important when you need them, but not all Sites do. They're not listed first, but they're one of the first things to configure for your virtual instance, whether you require them or not.

Since the Terms of Use page is enabled by default, one of your first actions should be to disable or replace the default, placeholder terms. You can disable the requirement for all Users to read the Terms of Use or set the Group ID and Article ID for the Web Content Article that contains your Terms of Use. "How do I find a web content article's Group ID and Article ID," you ask? No problem. The Group ID is the ID of the Site the Web Content is associated with. To find the Group/Site ID, follow these steps:

1. Go to *Site Menu* → *Configuration* → *Settings*.
2. Find the Site ID field in the General tab. Enter it into the Group ID field.

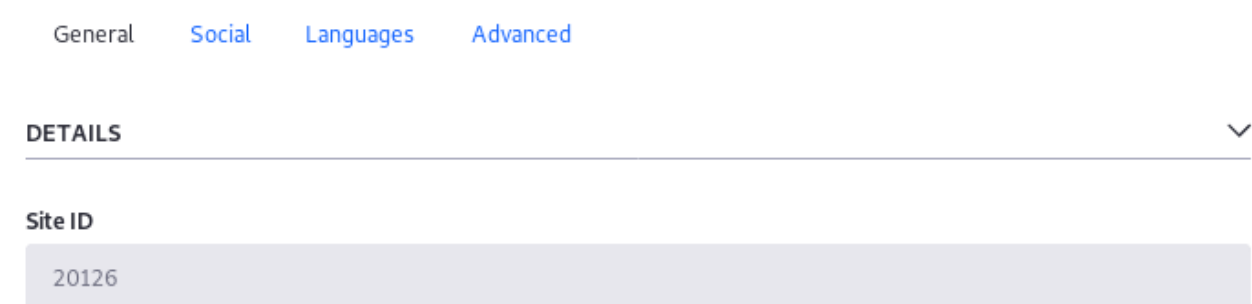


Figure 18.1: The Site ID in Site Settings is the Group ID in the terms of Use configuration.

To find the Web Content Article's ID, follow these steps:

1. Go to the *Site Menu* → *Content & Data* → *Web Content*.
2. Click on your Terms of Use article. Its ID appears at the top of the screen, with the Version and Workflow Status.




Figure 18.2: The Web Content Article ID is displayed in the edit screen.

Save the configuration. All your users signing in for the first time see your Terms of Use article. Users must agree to the Terms of Use to register their User Accounts.

Terms of Use



Welcome to the best site in the history of the internet.

There's are some things you should know:

1. Trolls are not tolerated.
2. Especially cave trolls.
3. Orcs are only slightly more tolerable.
4. Goblins are repugnant skulkers, and will be unceremoniously kicked out upon discovery.

...

I Agree

I Disagree

Figure 18.3: Turn a Web Content Article into the Site's Terms of Use page.

USER AUTHENTICATION INSTANCE SETTINGS

The User Authentication settings define how Users can authenticate, the various authentication methods that are required for them, and the screen names and email addresses that are reserved and can't be registered by Users.

Access the User Authentication settings in the Control Panel's *Configuration* → *Instance Settings* section, and select the *User Authentication* category under the *PLATFORM* section.

User Authentication contains two entries:

- General
- Reserved Credentials

Each configuration entry is described in the corresponding section below.

19.1 General

The General configuration entry contains several general authentication settings:

- Authenticate by email address (default), screen name, or User ID (a numerical ID auto-generated in the database—not recommended).
- Enable/Disable automatic log in. If enabled, a User can check a box which will cause the Site to “remember” the login information by placing a cookie on the browser. If disabled, Users must always log in manually.
- Enable/Disable forgotten password functionality.
- Enable/Disable request password reset links.
- Enable/Disable account creation by strangers. If running an Internet site, leave this enabled so visitors can create accounts on your Site.




Figure 19.1: Configure general authentication behavior and settings for external authentication systems.

- Enable/Disable account creation by those using an email address in the domain of the company running the Site (which is set on the General page of Instance Settings). This is handy if you're using Liferay to host both internal and external web sites. Make sure all internal IDs are created by administrators but external Users can register for IDs themselves.
- Enable/Disable email address verification. If enabled, Users receive a verification email with a link back to the virtual instance, verifying that the email address they entered is valid.

By default, all settings except for the last are enabled. User authentication by email address is an important default for the following reasons:

1. An email address is unique to the User who owns it.
2. People remember their email addresses. A Users who hasn't logged in for a while could forget their screen name.
3. If email address isn't used to authenticate, a User might change her email address then forget to update the email address in her profile. If this occurs, no notifications sent by the virtual instance will reach the User. Keeping the email address at the forefront of a User's mind when she logs in helps ensure the User keeps it current.

19.2 Reserved Credentials

The Reserved Credentials configuration entry specifies the screen names and email addresses Users aren't allowed to use. This prevents Users from creating IDs that look like administrative IDs or that have reserved words in their names.

Learn to configure a third party authentication service or set up Single Sign On (SSO) in the security documentation.

USERS INSTANCE SETTINGS

The Users configuration defines the look and feel of the Personal Menu, the default user associations for Users, and defines whether specific fields are available when a new User creates an account.

To access the Users settings, Open the Control Panel and navigate to *Configuration* → *Instance Settings*, and select the *Users* category under the *PLATFORM* section.

The Users Instance configuration contains three entries:

- Personal Menu
- Default User Associations
- Fields

Each configuration entry is described in the corresponding section below.

20.1 Personal Menu

The Personal Menu configuration entry specifies whether personal applications use the same look and feel as the current site or if they should use the look and feel of the My Dashboard pages instead. You can also specify whether the personal menu appears in the Control Menu (as it did in past versions of Liferay DXP) by toggling the *Show in Control Menu* option on and off.

20.2 Default User Associations

The Default User Associations configuration entry defines the default Sites, Organization Sites, Roles, and User Groups you want all new Users assigned to automatically. By default, new Users are only assigned to the Users Role. User groups are handy tools for pre-populating your Users' private Sites with pages, assigning Roles and permissions, and managing site membership. If you

update this configuration after your Users have created their accounts, don't worry. You can apply the updates to existing Users by checking the *Apply to Existing Users* checkbox. Changes take effect the next time the User signs in.

20.3 Fields

The Fields configuration entry contains settings for enabling/disabling the fields listed below on the Add/Edit User Form:

- Autogeneration of screen names
- Birthday field
- Gender field

MORE PLATFORM SECTION INSTANCE SETTINGS

The Instance Settings also contain settings for Infrastructure and Localization, as well as integrations for Analytics and Third Party map APIs. Each configuration entry is described in the corresponding section below.

21.1 Infrastructure

The Infrastructure category contains settings to specify which content types are validated during the Import/Export process and whether to delete temporary LAR files during the Import/Export process.

To access these settings, open the Control Panel and navigate to *Configuration* → *Instance Settings* and select *Infrastructure* under the PLATFORM heading.

21.2 Localization

The Localization category contains the configuration for setting the default instance language and time zone for the virtual instance. You can also configure the Language Selector's look and feel from the Widget Scope.

To access the Localization settings, Open the Control Panel and navigate to *Configuration* → *Instance Settings*, and select the *Localization* category under the *PLATFORM* section.

21.3 Analytics

The Analytics category defines the available analytics systems for the virtual instance. Enter an analytics system or remove one of the two pre-configured options (google and piwik). Activate these systems here, and configure them at the site level.

To access Analytics configuration settings, open the Control Panel and navigate to *Configuration* → *Instance Settings*, and select the *Analytics* category under the *PLATFORM* section.

21.4 Third Party

The Third Party category specifies the maps API provider for geolocalized assets. Choose OpenStreetMap or Google Maps as the maps API provider.

To access Third Party configuration settings, open the Control Panel and navigate to *Configuration* → *Instance Settings*, and select the *Third Party* category under the *PLATFORM* section.

USING LIFERAY'S SCRIPT ENGINE

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay DXP provides a robust script engine for executing Groovy scripts. You can execute scripts to perform maintenance tasks like data cleanup, user maintenance operations, bulk invocations of Liferay's API, or even system level operations in the scripting console.

These tutorials cover the following scripting topics:

- Invoking Liferay services
- Running scripts from the script console
- Using the script engine with workflow
- Script examples

The most common use of the scripting console is to invoke Liferay's services, so that's covered first. Familiarity with Liferay's API makes the scripting console a useful tool.

INVOKING LIFERAY SERVICES FROM SCRIPTS

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Many scripting scenarios require invoking Liferay services. Liferay `*ServiceUtil` classes are the fastest and most convenient way to invoke Liferay services in the script console. You can use Groovy to invoke Liferay services the same way you would use Java. Groovy's syntax facilitates writing concise, elegant scripts.

This first example illustrates correct syntax for interacting with Liferay services. It uses `UserLocalServiceUtil` to retrieve a list of users and print their names to Liferay's log. To do this, you could deploy a module with Java code like this:

```
import com.liferay.portal.kernel.model.User;
import com.liferay.portal.kernel.service.UserLocalServiceUtil;
import java.util.List;

...

int userCount = UserLocalServiceUtil getUsersCount();
List<User> users = UserLocalServiceUtil getUsers(0, userCount);

for (User user:users) {
    System.out.println("User Name: " + user.getFullName());
}

...
```

Or you could use Groovy—based on Java—and do the whole thing right from the script console with the same code:

```
import com.liferay.portal.kernel.model.User;
import com.liferay.portal.kernel.service.UserLocalServiceUtil;
import java.util.List;

int userCount = UserLocalServiceUtil getUsersCount();
List<User> users = UserLocalServiceUtil getUsers(0, userCount);

for (User user:users) {
    System.out.println("User Name: " + user.getFullName());
}
```

You can even make the code somewhat Groovier:

```
import com.liferay.portal.kernel.service.UserLocalServiceUtil
```

```
userCount = UserLocalServiceUtil.getUsersCount()
users = UserLocalServiceUtil.getUsers(0, userCount)
for (user in users){
    System.out.println("User Name: " + user.getFullName())
}
```

Groovy scripts that invoke Liferay services are easy to write and execute in the script console. How to run scripts is next.

23.1 Related Topics

Running Scripts From the Script Console
Leveraging the Script Engine in Workflow
Script Examples

RUNNING SCRIPTS FROM THE SCRIPT CONSOLE

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The script console provides a single view for executing Groovy scripts and printing their output. It has predefined variables that facilitate printing output and working with portlets and users. Here you'll learn these things:

- How to execute a script in the script console
- The predefined variables available in the script console
- Tips for running scripts in the script console

Important: The script console is for system operations and maintenance and not for end users. Limit script console access to portal administrators.

Start with running the script console's sample script.

24.1 Running the Sample Script

Here's how to run the sample script in the script console:

1. Sign in as an administrator.
2. In the Product Menu, navigate to *Control Panel* → *Configuration* → *Server Administration*.
3. Click on *Script*. This is the script console. The default sample script prints the User count to the console output.

```
// ## Groovy Sample ##  
  
number = com.liferay.portal.kernel.service.UserLocalServiceUtil.getUsersCount();  
  
out.println(number);
```

4. Click *Execute* and check the script console *Output* for the User count.

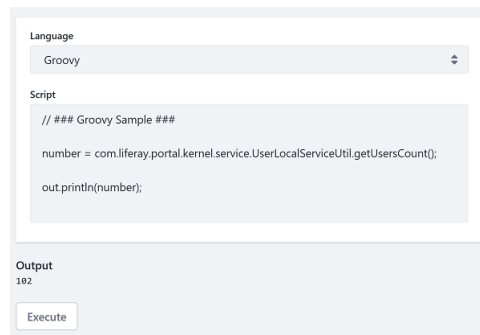


Figure 24.1: The script console’s sample Groovy script prints the User count to the console’s *Output* section.

The Groovy sample invokes the Liferay service utility `UserLocalServiceUtil` to get the user count. Then it uses `out` (a built-in `PrintWriter`) to write the count to the script console. Note that if you use `System.out.println` instead of `out.println`, your output is printed to Liferay’s log file rather than to the script console.

24.2 Predefined Variables

Here are the predefined variables available to scripts executed in the script console:

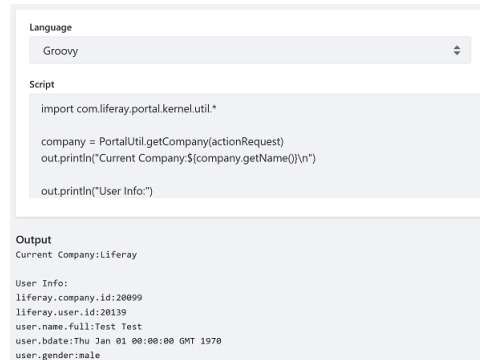
- `out` (`java.io.PrintWriter`)
- `actionRequest` (`javax.portlet.ActionRequest`)
- `actionResponse` (`javax.portlet.ActionResponse`)
- `portletConfig` (`javax.portlet.PortletConfig`)
- `portletContext` (`javax.portlet.PortletContext`)
- `preferences` (`javax.portlet.PortletPreferences`)
- `userInfo` (`java.util.Map<String, String>`)

This script demonstrates using the `actionRequest` variable to get the portal instance’s `Company`:

```
import com.liferay.portal.kernel.util.*

company = PortalUtil.getCompany(actionRequest)
out.println("Current Company:${company.getName()}\n")

out.println("User Info:")
userInfo.each {
    k,v -> out.println("${k}:${v}")
}
```



```
Language
Groovy

Script
import com.liferay.portal.kernel.util.*

company = PortalUtil.getCompany(actionRequest)
out.println("Current Company:${company.getName()}\n")

out.println("User Info:")

Output
Current Company:Liferay

User Info:
liferay.company.id:20099
liferay.user.id:20139
user.name.full:Test Test
user.date:Thu Jan 01 00:00:00 GMT 1970
user.gender:male
```

Figure 24.2: Here's an example of invoking a Groovy script that uses the predefined `out`, `actionRequest`, and `userInfo` variables to print information about the company and current user.

24.3 Tips

Keep these things in mind when using the script console:

- There is no undo.
- There is no preview.
- Permissions checking is not enforced for local services.
- Scripts are executed synchronously. Avoid executing scripts that might take a long time.

For these reasons, use the script console cautiously. Test your scripts on non-production systems before running them on production.

Of course, Liferay's script engine can be used outside of the script console. Next, you'll learn how workflows leverage Liferay's script engine.

LEVERAGING THE SCRIPT ENGINE IN WORKFLOW

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay's Kaleo workflow engine provides a robust system for reviewing and approving content in an enterprise environment. Even if you don't leverage scripts, it's a powerful and robust workflow solution. Adding scripts takes it to the next level. These scripts aren't run from the script console, but are embedded in XML workflow definitions and run during the execution of the workflow.

25.1 Injected Variables

Usually when you're scripting in Groovy, you must define your variables.

```
KaleoInstanceToken kaleoInstanceToken=new KaleoInstanceToken();
```

In workflow scripts, there are several pre-defined variables injected into your script context, to be called without defining them first.

25.2 Variables that are Always Available

These variables are available from anywhere that you can run a workflow script:

kaleoInstanceToken (KaleoInstanceToken) A workflow instance and corresponding instance token (the `KaleoInstanceToken`) are created each time a User clicks *Submit for Publication*. Use the injected token to retrieve its ID, by calling `kaleoInstanceToken.getKaleoInstanceId()`. This is often passed as a method parameter in a script.

userId The `userId` returned is context dependent. Technically, the logic works like this: if the `KaleoTaskInstanceToken.getCompletionUserId()` is null, check `KaleoTaskInstanceToken.getUserId()`.

If that's null too, call `KaleoInstanceToken.getUserId()`. It's the ID of the last User to intervene in the workflow at the time the script is run. In the created node, this would be the User that clicked *Submit for Publication*, whereas it's the ID of the reviewer upon exit of the review node of the Single Approver definition.

workflowContext (Map<String, Serializable>) The workflow context is full of useful information you can use in your scripts. Usually you'll pass this as a parameter to a method that requires a `WorkflowContext` object, but all of the `WorkflowContext`'s attributes are available in the script as well. The workflow context in the script is context dependent. If a call to `ExecutionContext.getWorkflowContext()` comes back null, then the workflow context is obtained by `KaleoInstanceModel.getWorkflowContext()`.

25.3 Variables Injected into Task Nodes

If a `kaleoTaskInstanceToken` has been created:

kaleoTaskInstanceToken (KaleoTaskInstanceToken) The task's token itself is available in the workflow script. Use it to get its ID, to use in other useful programmatic workflow activities, like programmatic assignment.

taskName (String) The task's own name is accessible (returns the same as `KaleoTak.getName()`).

workflowTaskAssignees (List<WorkflowTaskAssignee>) If the script is inside a task node, get a List of its assignees.

kaleoTimerInstanceToken (KaleoTimerInstanceToken) If a task timer exists, use the `kaleoTimerInstanceToken` to get its ID, by calling `kaleoTimerInstanceToken.getKaleoTimerInstanceId()`.

25.4 Scripting Examples

The final step in a workflow runs a script that makes content available for use. The snippet below accesses the Java class associated with the workflow to set content's status to *approved*.

```
<script>
  <![CDATA[
    import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
    import com.liferay.portal.kernel.workflow.WorkflowConstants;

    WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("approved"), workflowContext);
  ]]>
</script>
<script-language>groovy</script-language>
```

At virtually any point in a workflow, you can use Liferay's script engine to access workflow APIs or other Liferay APIs. There are a lot of different ways you could use this. Here are a few practical examples:

- Getting a list of users with a specific workflow-related role

- Sending an email to the designated content approver with a list of people to contact if he is unable to review the content
- Creating an alert to be displayed in the Alerts portlet for any user assigned to approve content

Of course, before you try any of this, you need to know the appropriate syntax for inserting a script into a workflow. In an XML workflow definition, a script can be used in any XML type that can contain an actions tag: those types are <state>, <task>, <fork> and <join>. Inside of one of those types, format your script like this:

```
<actions>
  <action>
    <script>
      <![CDATA[*the contents of your script*]]>
    </script>
    <script-language>*your scripting language of choice*</script-language>
  </action>
  ...
</actions>
```

Here's an example of a workflow script created in Groovy. This one is used with a Condition statement in Kaleo. It accesses Liferay's asset framework to determine the category of an asset. The script uses the category to determine the correct approval process automatically. If the category legal has been applied to the asset, the asset is sent to the Legal Review task upon submission. Otherwise, the asset is sent to the Default Review task.

```
<script>
  <![CDATA[
    import com.liferay.portal.kernel.util.GetterUtil;
    import com.liferay.portal.kernel.workflow.WorkflowConstants;
    import com.liferay.portal.kernel.workflow.WorkflowHandler;
    import com.liferay.portal.kernel.workflow.WorkflowHandlerRegistryUtil;
    import com.liferay.asset.kernel.model.AssetCategory;
    import com.liferay.asset.kernel.model.AssetEntry;
    import com.liferay.asset.kernel.model.AssetRenderer;
    import com.liferay.asset.kernel.model.AssetRendererFactory;
    import com.liferay.asset.kernel.service.AssetEntryLocalServiceUtil;

    import java.util.List;

    String className = (String)workflowContext.get(
        WorkflowConstants.CONTEXT_ENTRY_CLASS_NAME);

    WorkflowHandler workflowHandler =
        WorkflowHandlerRegistryUtil.getWorkflowHandler(className);

    AssetRendererFactory assetRendererFactory =
        workflowHandler.getAssetRendererFactory();

    long classPK =
        GetterUtil.getLong((String)workflowContext.get(
            WorkflowConstants.CONTEXT_ENTRY_CLASS_PK));

    AssetRenderer assetRenderer =
        workflowHandler.getAssetRenderer(classPK);

    AssetEntry assetEntry = assetRendererFactory.getAssetEntry(
        assetRendererFactory.getClassName(), assetRenderer.getClassPK());

    List<AssetCategory> assetCategories = assetEntry.getCategories();

    returnValue = "Default Review";

    for (AssetCategory assetCategory : assetCategories) {
```

```

        String categoryName = assetCategory.getName();

        if (categoryName.equals("legal")) {
            returnValue = "Legal Review";

            return;
        }
    }
]]>
</script>
<script-language>groovy</script-language>

```

Within a workflow, the next task or state is chosen based on the return value. See some examples of workflow scripts by accessing the embedded workflows and inspecting the XML.

25.5 Calling OSGi Services

How do you call OSGi services from a workflow script, accounting for the dynamic environment of the OSGi runtime, where services your script depends on can disappear without notice? Use a service tracker. That way you can make sure your code has access to the service it needs, and if not, do something appropriate in response. Here's a little example code to show you how this might look in Groovy:

```

import com.liferay.journal.model.JournalArticle;
import com.liferay.journal.service.JournalArticleLocalService;
import com.liferay.portal.scripting.groovy.internal.GroovyExecutor;

import org.osgi.framework.Bundle;
import org.osgi.framework.FrameworkUtil;
import org.osgi.util.tracker.ServiceTracker;

ServiceTracker<JournalArticleLocalService, JournalArticleLocalService> st;

try {
    Bundle bundle = FrameworkUtil.getBundle(GroovyExecutor.class);

    st = new ServiceTracker(bundle.getBundleContext(), JournalArticleLocalService.class, null);
    st.open();

    JournalArticleLocalService jaService = st.waitForService(500);

    if (jaService == null) {
        _log.warn("The required service 'JournalArticleLocalService' is not available.");
    }
    else {
        java.util.List<JournalArticle> articles = jaService.getArticles();
        if (articles != null) {
            _log.info("Article count: " + articles.size());
        } else {
            _log.info("no articles");
        }
    }
}
catch(Exception e) {
    //Handle error appropriately
}
finally {
    if (st != null) {

```

```
        st.close();
    }
}
```

If you read the article on service trackers, the only odd looking piece of the above code is the `getBundle` call: why is `GroovyExecutor.class` passed as a parameter? The parameter passed to the `FrameworkUtil.getBundle` call must be a class from the OSGi bundle executing the workflow script. This is different from the context of a plugin project, where you'd want to get the bundle hosting the class where you're making the call (using `this.getClass()`, for example). Note that for another scripting engine, you must pass in a concrete class from the particular bundle executing your script.

The combination of Liferay's script and workflow engines is incredibly powerful. Since, however, it enables users to execute code, it can be dangerous. When configuring your permissions, be aware of the potential consequences of poorly or maliciously written scripts inside a workflow definition. For more information on creating workflow definitions, see the workflow documentation.

25.6 Related Topics

Running Scripts From the Script Console
Script Examples

SCRIPT EXAMPLES

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Here are some examples to help you use Liferay’s script console. Note: Most of these originated from a Liferay blog post.

The following scripts are Groovy scripts but they can be adapted to other languages.

- Example 1: Presenting New Terms of Use to Users
- Example 2: Embedding HTML Markup in Script Outputs
- Example 3: Show Exceptions in the Script Console
- Example 4: Implement a Preview Mode
- Example 5: Plan a File Output for Long-Running Scripts

26.1 Example 1: Presenting New Terms of Use to Users

This example retrieves user information from the database, makes changes, and then saves the changes in the database. Suppose that your company has updated the terms of use and wants present users with the updated terms of use whenever they sign in next. When they agree to the terms of use, a boolean attribute called `agreedToTermsOfUse` is set in their user records. As long as the value of this variable is true, users aren’t presented with the terms of use when they sign in. If you set this flag to false for each user, each user must agree to the terms of use again before they can sign in.

1. Enter and execute the following code in the script console:

```
import com.liferay.portal.kernel.service.UserLocalServiceUtil

userCount = UserLocalServiceUtil.getUsersCount()
users = UserLocalServiceUtil.getUsers(0, userCount)
```

```
for (user in users) { println("User Name: " + user.getFullName() + " -- " +
user.getAgreedToTermsOfUse()) }
```

This code prints each user's `agreedToTermsOfUse` attribute value.

2. Replace that with this script:

```
import com.liferay.portal.kernel.service.UserLocalServiceUtil

userCount = UserLocalServiceUtil getUsersCount()
users = UserLocalServiceUtil getUsers(0, userCount)

long currentUserId = Long.parseLong(userInfo.get("liferay.user.id"))

for (user in users){

    if(!user.isDefaultUser() && (user.getUserId() ≠ currentUserId)) {

        user.setAgreedToTermsOfUse(false)
        UserLocalServiceUtil.updateUser(user)

    }

}
```

This sets each user's `agreedToTermsOfUse` attribute to false. It skips the default user as well as the default admin user that's currently signed in and running the script.

3. Click *Execute*.

4. Verify the script updated the records by running the first script again.

All users (except the default user and your user) have been updated.

You've enabled the new terms of use agreement for all users to accept.

26.2 Example 2: Embedding HTML Markup in Script Outputs

The output of the script console is rendered as HTML content. Thus, you can embed HTML markup in your output to change its look and feel. Here's an example:

```
import com.liferay.portal.kernel.service.*

number = com.liferay.portal.kernel.service.UserLocalServiceUtil.getUsersCount();
out.println(
    """
        <div style="background-color:black; text-align: center">
            <h1 style="color: #37A9CC; font-size:xx-large">${number}</h1>
        </div>
    """);
```

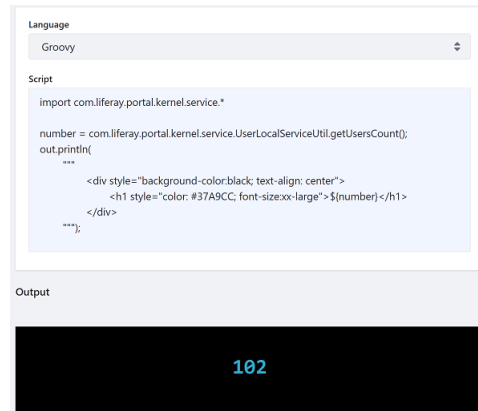



Figure 26.1: Here's an example of invoking a Groovy script that embeds HTML markup in the output of the script.

26.3 Example 3: Show Exceptions in the Script Console

When any exception occurs during script execution, the error message is always the same:

Your request failed to complete.

This message gives no detail about the error. To find information about the error and what caused it, you must usually examine the server logs.

You can, however, use the following technique to make exception details appear in the script console. Wrap your code with a try / catch block and print the stacktrace to the console output from the catch clause. Note that even this technique does not catch script syntax errors. Here's an example:

```
try {
    nullVar = null
    out.println(nullVar.length())
} catch(e) {
    out.println("""<div class="portlet-msg-error">${e}</div>""")
    e.printStackTrace(out)
}
```



Figure 26.2: Here's an example of a Groovy script that catches exceptions and prints exception information to the script console.

26.4 Example 4: Implement a Preview Mode

Since Liferay's script console does not provide an undo feature, it can be convenient to set up a kind of preview mode. The purpose of a preview mode is to determine any permanent effects of a script before any information is actually saved to the Liferay database. The preview mode consists in using a `previewMode` flag which determines whether the operations with permanent effects should be executed or not. If `previewMode` is true, all the data that would be permanently affected by the script is printed instead. Then you can see an outline of the data impacted by the script. If everything is okay, switch the flag so the script can make permanent updates to the database.

Here's an example Groovy script that sets users to inactive. Clearly, you'd want to test this with preview mode before running it:

```
import java.util.Calendar
import com.liferay.portal.kernel.service.*
import com.liferay.portal.kernel.model.*
import com.liferay.portal.kernel.dao.orm.*
import static com.liferay.portal.kernel.workflow.WorkflowConstants.*

//
// Deactivate users never logged and created since more than 2 years
//

previewMode = true // Update this flag to false to really make changes

Calendar twoYearsAgo = Calendar.getInstance()
twoYearsAgo.setTime(new Date())
twoYearsAgo.add(Calendar.YEAR, -2)

DynamicQuery query = DynamicQueryFactoryUtil.forClass(User.class)
    .add(PropertyFactoryUtil.forName("lastLoginDate").isNull())
    .add(PropertyFactoryUtil.forName("createDate").lt(twoYearsAgo.getTime()))

users = UserLocalServiceUtil.dynamicQuery(query)

users.each { u ->
    if(!u.isDefaultUser() && u.getStatus() != STATUS_INACTIVE) {
        out.println(u.getEmailAddress())
        if(!previewMode) {
            UserLocalServiceUtil.updateStatus(u.getUserId(), STATUS_INACTIVE)
        }
    }
}

if(previewMode) {
    out.println('Preview mode is on: switch off the flag and execute '
        + 'again this script to make changes to the database')
}
```

26.5 Example 5: Plan a File Output for Long-Running Scripts

Important: The script console is for system operations and maintenance and not for end users. Limit script console access to portal administrators.

When a script has been running for a long time, the console could return an error even though the script can continue running and potentially conclude successfully. But it's impossible to know the outcome without the corresponding output!

To bypass this limitation, you can send the output of the script console to a file instead of to the console itself or to the Liferay log. For example, consider this script:

```
import com.liferay.portal.kernel.service.*
import com.liferay.portal.kernel.dao.orm.*

// Output management

final def SCRIPT_ID = "MYSCRIPT"
outputFile = new File("${System.getProperty("liferay.home")}/scripting/out-${SCRIPT_ID}.txt")
outputFile.getParentFile().mkdirs()

def trace(message) {
    out.println(message)
    outputFile << "${message}\n"
}

// Main code

users = UserLocalServiceUtil.getUsers(QueryUtil.ALL_POS, QueryUtil.ALL_POS)
users.each { u ->
    trace(u.getFullName())
}
```

The script above creates a subfolder of Liferay Home called `scripting` and saves the script output to a file in this folder. After running the script above, you can read the generated file without direct access to the file system. Here's a second script that demonstrates this:

```
final def SCRIPT_ID = "MYSCRIPT"
outputFile = new File("${System.getProperty("liferay.home")}/scripting/out-${SCRIPT_ID}.txt")
out.println(outputFile.text)
```

One advantage of using a dedicated output file instead of using a classic logger is that it's easier to get the script output data back. Getting the script output data would be more difficult to obtain from the portal log, for example, because of all the other information there.

26.6 Related Topics

Running Scripts From the Script Console
 Leveraging the Script Engine in Workflow
 Using Liferay's Script Engine

CUSTOM FIELDS

Have you ever wondered why there's no *Head Circumference* field in the form for adding users to Liferay DXP? Probably because most sites based on Liferay DXP don't need it. If you're an administrator at the Lunar Resort, however, you certainly need to know your guests' head circumference so you can provide them with a properly fitting helmet.

Many of Liferay DXP's assets and resources let you add new fields to their edit forms. Here's the complete list:

- Blogs Entry
- Calendar Event
- Document
- Documents Folder
- Knowledge Base Article
- Knowledge Base Folder
- Message Boards Category
- Message Boards Message
- Organization
- Page
- Role
- Site
- User
- User Group
- Web Content Article
- Web Content Folder
- Wiki Page

Developer Use Case: Adding custom fields to Liferay DXP resources affords flexibility to developers. Suppose you must limit the number of users that can be assigned to a particular Role. First an administrator creates a custom field called *max-users* for the Role. A developer then creates a module that inserts logic before a user is added to that Role. If the logic detects that the maximum number of Role users would be exceeded by completing the action, an exception is thrown and the action does not complete.

27.1 Adding Custom Fields

To add custom fields, find the Custom Fields entry beneath the Control Panel's Configuration heading.

To add a custom field to one of the listed entities,

1. Choose a resource by clicking on it.
2. Click the add (+) button.
3. Choose a field type: Text Area, Input Field, Dropdown, Checkbox, Radio, Geolocation, Date, True/False.
4. Add a name that's used as a key for accessing the field programmatically.

Note: The Key you enter here is the name of the new field. It's stored in the database and used by developers to access the custom field with the `<liferay-ui:custom-attribute />` tag. It is also used to label the field in the UI.

4. Choose the Data Type of field and set any advanced properties.




Figure 27.1: At The Lunar Resort, a Head Circumference field is necessary for all users.

5. Click Save.

That's it.

Once you have a custom field for a resource, go find it in the UI of the resource. First find the UI location for the resource, and all its custom fields are displayed in a Custom Fields panel. For example, consider the Users UI:

Navigate to *Control Panel* » *Users and Organizations*. Click on a User to open the Edit User form and scroll down to find your custom field.

You can also leverage your custom field in Content Pages and Display Page Templates. See Editable Elements for more information.




Figure 27.2: The Custom Fields panel is found at the bottom of the Edit User form.

27.2 Editing a Custom Field

You can't change the key or field type of a custom field, but you can delete it and create a new one if necessary.

Edit an individual custom field's permissions by clicking the field's kebab menu (☰), then *Permissions*. Permission can be granted or removed for these actions:

- Delete
- Permissions
- Update
- View

Custom fields make many of Liferay DXP's entities extensible directly from the administrative user interface. Use them as is or combine them with some back-end code, and you have yet another powerful, flexible feature at your disposal. As they're fond of saying at The Lunar Resort, "The sky is certainly not the limit."

Edit Input Field

Field Name

Head-Circumference

The custom field key is used to access the field programmatically through the `<liferay-ui:custom-attribute />` tag.

Data Type

Decimal

Starting Value

0.0

ADVANCED PROPERTIES

Precision

32-bit

Precision sets the length of the number.

Visible with Update

NO

Setting a custom field to visible with update permission means that a user with update permission can view this field. This setting overrides the value of hidden in this case.

Hidden

NO

Setting a custom field to hidden means that the field's value is never shown in any user interface besides this one. This lets the field be used for more obscure and advanced purposes such as acting as a placeholder for custom permissions.

Secret

NO

Setting a custom field to secret means that typing is hidden on the screen. Use this for passwords.

Searchable as Keyword

YES

Setting a custom field to searchable means that the value of the field is indexed when the entity (such as User) is modified. Only `java.lang.String` fields can be made searchable. Note that when an field is newly made searchable, the indexes must be updated before the data is available to search.

Save

Cancel

Figure 27.3: The exact Custom Fields configuration options you use depend on the field type you choose.




Figure 27.4: You can delete a custom field, edit it, or configure its permissions.

MANAGING USERS

Ever heard a retailer advertise as a “one stop shop” for anything you want? The idea is they have so much stuff that whatever you want is probably there. Liferay’s Control Panel is like this. Where do you create Users, Organizations, or Sites? Where do you configure permissions and plugins and pretty much anything else? You do it from the Control Panel.

The Control Panel is divided into six main areas: Users, Sites, Apps, Configuration, and Workflow. The Users section lets you create and manage Users, Organizations, User Groups, Roles, and Password Policies. If monitoring has been enabled, you can also view all the live sessions of your Users.

Anonymous User: *Anonymous Anonymous* is used for the new Managing User Data functionality. Created the first time an administrator clicks *Delete Personal Data* for a User, *Anonymous Anonymous* is a deactivated User assigned anonymized assets. The Anonymous User is configurable, so the name and configuration details might be different in your virtual instance.

Begin exploring Liferay’s User Management functionality by reading about adding and editing users.

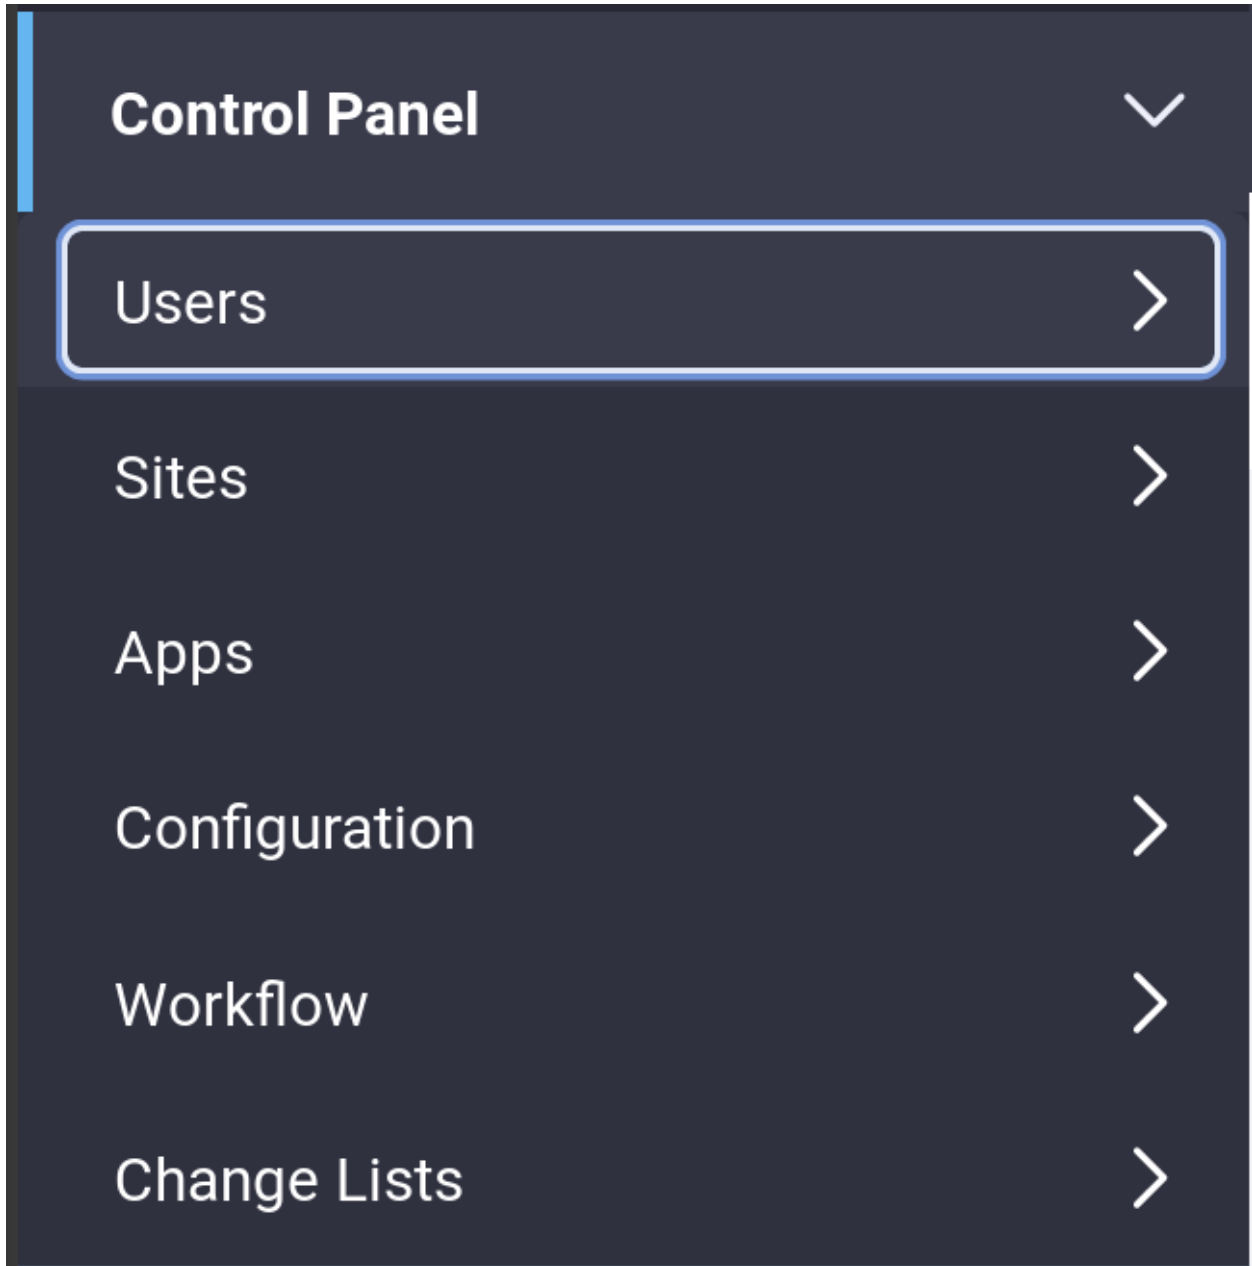


Figure 28.1: Administrators can access the Control Panel from the Product Menu.

USERS AND ORGANIZATIONS

Users and *Organizations* are fundamental entities. If your site requires people (even just a set of site administrators) to have accounts to do anything, you need to know about users. If your users are at all divided hierarchically, like into departments, you'll find that organizations are helpful.

You're probably not surprised to hear that Users and Organizations are managed in the Control Panel's *Users and Organizations* section. If it were any different, it'd be weird.

Consider the Lunar Resort site. Consider what you'd do if

- An employee leaves the company to join that pesky competitor, Martian Resort and Luxury Spa.
- An employee joins the resort as a new Mechanical Crew member.
- An employee is promoted from Crew Supervisor to Department Head and needs the requisite permissions.
- You need to organize the users by department.
- A new department is added to the Lunar Resort and the employees need their own internal website.
- An employee gets married, and their name changes.

The user tasks listed above are all resolved in the Users and Organizations section of the Control Panel.

29.1 What are Users?

In case there's any confusion over the term, a User is an entity that can sign into the portal and do something. Generally a User has more privileges, called Permissions, than a Guest of your site, who does not sign in. Users are assigned Roles, and Roles define the User's privileges.

Understanding Users is pretty straightforward. Organizations are a bit trickier, but a smart administrator like you is undoubtedly up to the challenge. Read more about Organizations [here](#).

The remaining articles in this section give you guidance on managing (creating, deleting, editing, and more) Users and Organizations.

ADDING, EDITING, AND DELETING USERS

At the root of managing Users is adding, editing, and deleting them. As long as you're the Administrative user, you can do all these things and more.

30.1 Adding Users

Here's how to edit Users:

1. From the Product Menu, click *Control Panel* → *Users* → *Users and Organizations*.
2. In the Users tab, click the *Add* button .
3. Fill out the Add User form and click *Save*. At a minimum, provide a Screen Name, First Name, Last Name, and Email Address for the User.

Note: Screen names and email addresses are not interchangeable. A screen name cannot contain an '@' symbol because it is used in the URL to a User's private page.

The Add User functionality is split over several independent forms. Saving the first form creates the User, and then you'll see a success message saying

Success. Your request completed successfully.

After submission of the first form, you see a larger form with many sections. The one you're on is the Information section. To the left is a navigation pane where you can continue configuring the user you're adding by clicking through the available sections. The options in the left menu change as you click through the tabs at the top. Peruse the sections for the three tabs (General, Contact, Preferences) and fill in all the applicable information.


| Name | Screen Name | Job Title | Organizations | User Groups |
|----------------------|-------------|-------------------|---------------------------|-------------|
| ■ Melvin Dooitrong | melvin | Intern | Physical Plant Department | ⋮ |
| ■ James Jeffries | james | Lunar Associate | Physical Plant Department | ⋮ |
| ■ Rex Nihilo | rex | Upsale Manager | Physical Plant Department | ⋮ |
| ■ Ziltoid Omniscient | ziltoid | Coffee Reviewer | Physical Plant Department | ⋮ |
| ■ Jane Slaughter | jane | Chief of Security | Physical Plant Department | ⋮ |
| ■ Marvin Smart | marvin | Lunar Associate | Physical Plant Department | ⋮ |

Figure 30.1: Add Users from the Users and Organizations section of the Control Panel.

You don't have to fill anything else out right now. Just note that when the user account was created, a password was automatically generated. If Liferay was correctly installed and a mail server was set up, an email message with the User's new password was sent to the User's email address.

If you haven't set up a mail server, click the *Password* item from the General menu and manually set a password for your new user. Enter the new password twice.

30.2 Editing Users

If you click on *Users and Organizations* in the Control Panel, you'll see your own user's account in the list of Users, along with any others. To change something about a particular user, click the *Actions* button () next to that user.

Choosing *Edit* takes you back to the Edit User page where you can modify any aspect of the User account including the screen name, email address, first name, last name, Site and Organization memberships, Roles, etc.

Choosing *Permissions* allows you to define which Roles have permissions to edit the User.

Choosing *Manage Pages* allows you to configure the personal pages of a User.

Choosing *Impersonate User* opens another browser window that loads the site as if you were the User so you can test your User management on a User to make sure you're achieving the desired

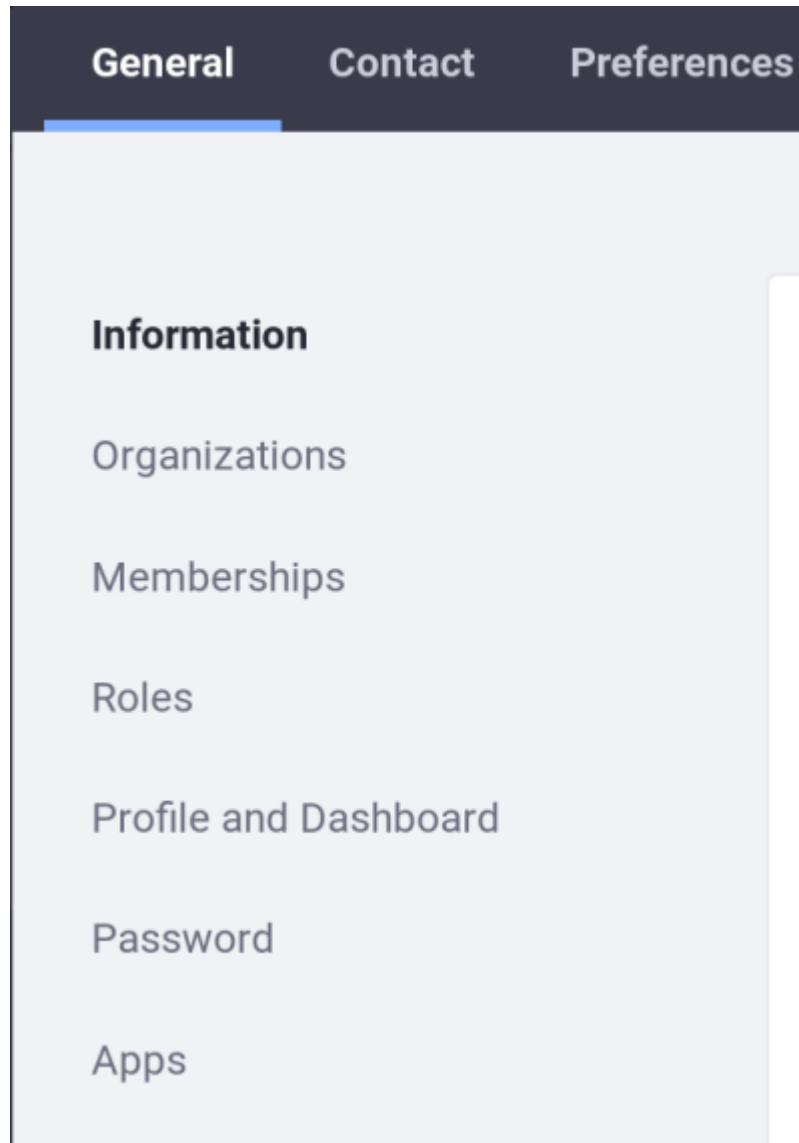
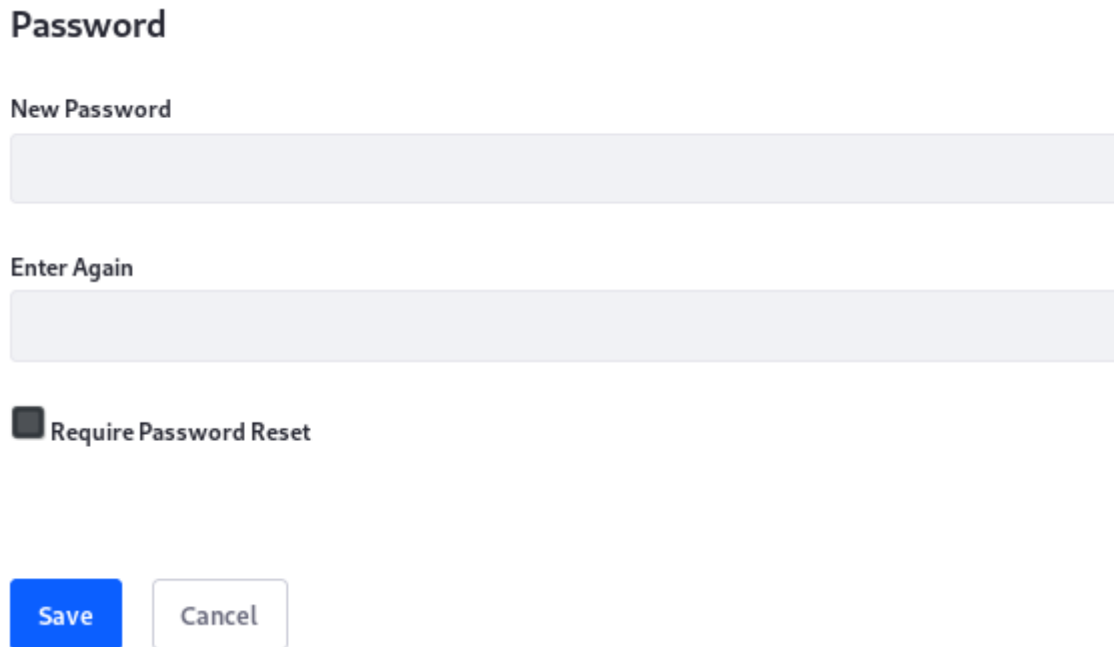


Figure 30.2: At a minimum, enter a screen name, email address, and first name to create a new user account. Then you'll be taken to the Information form and can continue configuring the user.



The screenshot shows a form titled "Password" with two input fields: "New Password" and "Enter Again". Below the fields is a checkbox labeled "Require Password Reset" which is currently checked. At the bottom of the form are two buttons: a blue "Save" button and a white "Cancel" button with a grey border.

Figure 30.3: Enter the password twice to manually set the password for a user. If the Password Policy you're using is configured to allow it, select whether to require the user to reset their password the first time they sign in to the portal.

behavior, without having to repeatedly log out of your administrator account and into the User's account.

Choosing *Deactivate* deactivates the user's account. The User is still in your database along with all the rest of your Users, but the account is deactivated, so the User cannot sign in to the portal. You can toggle between active and inactive Users in the Users view. If all the Users are active, this filtering option doesn't appear.

Choosing *Erase Personal Data* deletes the User's personal data.

Choosing *Export Personal Data* lets you download the User's personal data.

Most Users can't perform any of the above actions. In fact, most Users won't have access to the Control Panel at all. You can perform all of the above functions because you have administrative access.

30.3 Deleting Users

You must be careful when deleting Users. To guard against accidental deletion of Users, a two-step process must be followed: deactivate first, then delete.

1. Find the User to delete in the Users tab of *Control Panel* → *Users* → *Users and Organizations*. If you have a lot of Users, save time by searching for the User.
2. Click the *Actions* menu for the User and select *Deactivate*. You're asked to confirm that you want to deactivate the User. Click *OK*.

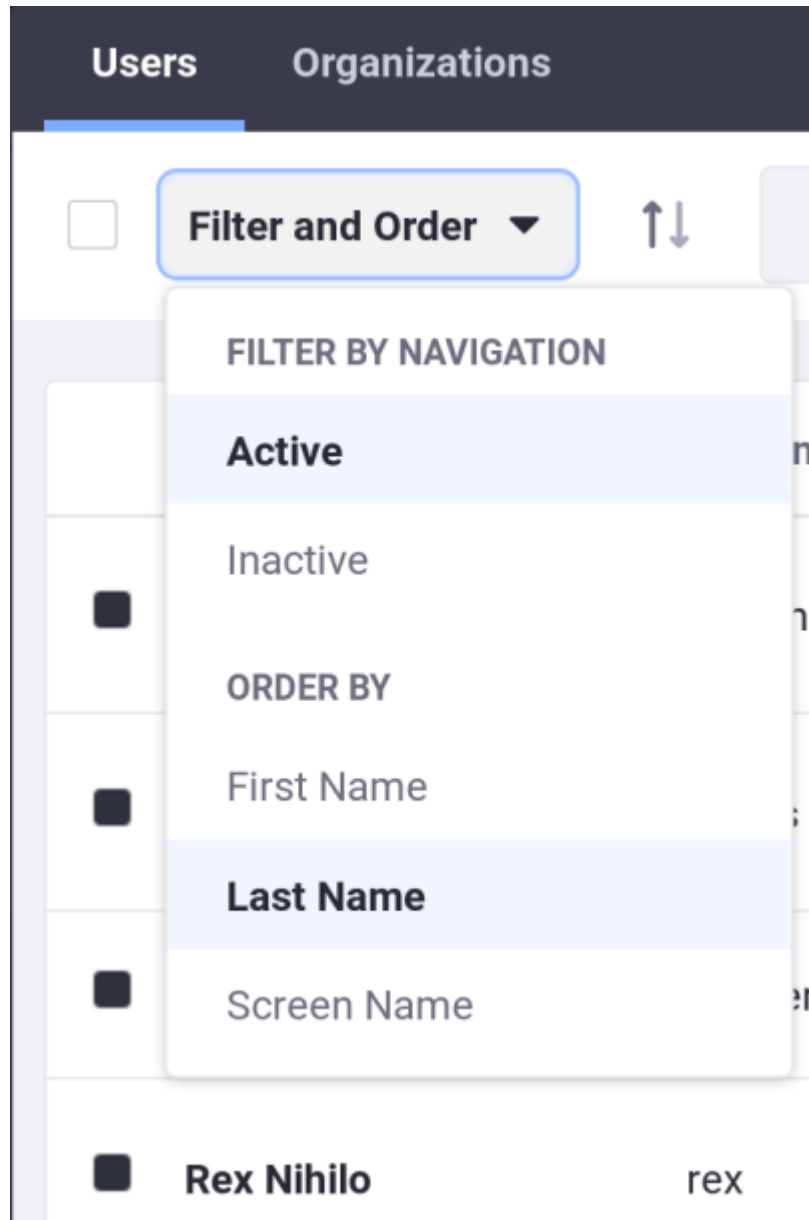


Figure 30.4: You can choose whether to view active or inactive (deactivated) portal users in the users list found at *Product Menu* → *Control Panel* → *Users* → *Users and Organizations*.

You'll see a success message and the User disappears, but isn't gone yet.

3. By default the Users table displays only Active users. Click on *Filter and order* in the top of the table and a dropdown menu appears. Click *Inactive*, and you can see the User you just deactivated.
4. Click the Actions menu again, and click *Delete* if you really mean to delete the User. Confirm that you want to delete the User, and now the User is gone. This time, it's for real.

Deactivated Users: Deactivating a User means the User can't log in to the portal. He/she has no more permissions in the Sites and pages of the portal than a guest, although the account still exists in the system.

Users are reactivated when an administrator finds them in the Users table (be sure you're filtering the table results by Deactivated users), clicks the Actions menu, and selects Activate. There's no confirmation window for activation: they're automatically restored to their former status once Activate is clicked.

Now you understand the basic principles of User administration. There are important additional topics in the next article that you should consider mandatory information for all portal administrators, so do continue reading.

USER MANAGEMENT: ADDITIONAL TOPICS

You've learned the basics on adding and editing Users, but there are additional important topics that go beyond the most basic tasks an administrator must perform. Read on to learn about these.

31.1 Password Resets

The Add User functionality includes a *Require Password Reset* checkbox at the bottom of the Password form. The default password policy does not even allow administrators to deselect this option. As the administrator, however, you can modify the default password policy so that this box becomes usable.

1. Navigate to *Password Policies* in Control Panel → Users.
2. Click on the *Default Password Policy*.
3. Deselect the *Change Required* switcher in the Password Changes section. Now you can decide whether users you add must reset their passwords.

See Password Policies for more information on editing the default policy or creating your own.

31.2 Adding an Administrative User

If you're setting things up for the first time, you're likely to be using the default administrator account, the account of one of those famous Liferay Administrators, *Test Test* or her cousin, *Joe Bloggs*. Because these are default accounts, hackers know about them, so it's better to set up your own administrator account. Add a user with your information, then give your user account the same administrative rights as the default administrator's account:

1. Click the *Roles* link in the left navigation pane (in the *Edit User* page's *General* tab). This page of the form shows the Roles to which your account is currently assigned. No Roles appear by default (the User role does not appear since it can't be removed).
2. Click *Select* under Regular Roles and assign the Administrator Role to your user account. A dialog box pops up with a list of all the regular (portal-scoped) Roles in the portal. Select the Administrator Role from the list (click *Choose*). The dialog box disappears and the Role is added to the list of Roles associated with your account. You are now a portal administrator. Log out and then log back in with your own user account.

Power Users: Users are not assigned the Power User Role by default. The Power User Role grants more permissions than the User Role. If the User Role is sufficient for you, ignore the Power User Role. Alternatively, use it to provide a second level of User permissions and assign it to those Users. If there are certain custom permissions that you'd like all of your portal Users to have, you can grant these permissions to the User Role. You can also customize the default Roles a new User receives via *Default User Associations*. This is covered in *Setting Up a Virtual Instance*.

In production, you should always delete or disable the default administrator account to secure your portal.

31.3 Gender

To collect data on users' genders, enable the binary gender field in the *Add User* form or create a custom field that meets your needs.

Enable the binary field by including the following in `portal-ext.properties`:

```
`field.enable.com.liferay.portal.kernel.model.Contact.male=true`
```

31.4 User Profile Pictures

Users have profile pictures. Administrative Users can upload images in the *Edit User* form, and any User can update her own account information, including image, from her personal site (*My account* → *Account Settings*).

If no image is explicitly uploaded for a User's profile picture, a default User icon is assigned as the User avatar. By default the User's initials are displayed (First Name then Last Name) over a random color.

If the initials-based approach for generating User profile pictures isn't suitable for your portal, disable the inclusion of Users' initials in the default icons:

1. Navigate to *Control Panel* → *Configuration* → *System Settings*.

Information

PERSONAL INFORMATION

Screen Name *

ray

Email Address *

ray@liferay.com

Language

English (United States)

Prefix



Change

Delete

Figure 31.1: Upload images for user avatars in the Edit User form.



Figure 31.2: The default user profile picture is an icon with the user initials over a randomly colored bubble.

2. In the Platform section, click *Users* → *User Images*.
3. Deselect *Use Initials for Default User Portrait*.

Now, instead of the default icon, which is a colorful circle containing the user's initials, the icon is a gray circle containing the approximate shape of a human being.




Figure 31.3: If you disable the default initials-based profile picture, this icon is used instead.

This is just the default. To override it with your own default image:

1. Create at least one image that is a 100x100 px square. Place it somewhere on the application server's classpath. For example, in Tomcat you could place it in the `tomcat/webapps/ROOT/WEB-INF/classes` folder.

2. Set the following property in a `portal-ext.properties` file:

```
image.default.user.portrait=image-filename-here.png
```

This overrides the value of this portal property:

```
image.default.user.portrait=com/liferay/portal/dependencies/user_portrait.png
```

****NOTE:**** If you are using the binary field to collect information on users' genders (see above), then you'll have two default images to override. Set these properties instead:

```
image.default.user.female.portrait=image-filename.png
image.default.user.male.portrait=image-filename.png
```

3. Restart the application server.

Note: There's a way to adjust which initials are displayed and in what order, so you can make the default user icon (with the user initials) work for your locale. These settings are configured in a Language Settings module, so kidnap a friendly developer, give him a cup of coffee, and tell him the settings you want to change:

`lang.user.default.portrait=initials` sets the type of icon to use for avatars. The default value is *initials*. If set to *initials*, the next property configures which initials to display, and in what order. Alternatively, specify *image*, which gives you the same non-initials default image shown above.

`lang.user.initials.fields=first-name,last-name` determines which initials appear in the user portrait and in what order. The setting here only matters if `lang.user.default.portrait` is set to *initials*. Valid values are first name, middle name, last name, with first and last name as the defaults.

31.5 Numeric Screen Names

In prior versions, numeric user screen names were disabled out of the box via the default portal property

```
users.screen.name.allow.numeric=false
```

Other user management systems (LDAP, for example) did not have the same restriction, which made importing users more difficult. Administrators first had to set the above property to true before importing and hope that no screen names conflicted with site IDs. In 7.0, this property defaults to true and there's no danger of numeric screen names conflicting with site IDs:

```
users.screen.name.allow.numeric=true
```


This means you're free to set a user screen name to *24601*, or whatever other number you can think of, and imports from systems that allow numeric screen names go more smoothly. That's everything you need to know to take advantage of this feature. Keep reading to understand what enabled the change.

Because users have personal sites, the URL to user *24601*'s personal site is

```
http://localhost:8080/web/24601
```

Meanwhile, a default site URL to cleverly named *Test Site* is

```
http://localhost:8080/web/test-site
```

There's no conflict here, but two conditions could easily lead to one:

1. *Test Site*'s group ID matches the number chosen for the user's screen name. Each site has a unique numeric identifier in the database, called group ID. There's nothing stopping it from matching the user's numeric screen name, so it could easily be *24601* just like the hypothetical user above.
2. A site administrator comes along and changes the site's friendly URL to match its `groupId`. Hello, URL conflict! Now the site's URL matches the user's URL:

```
http://localhost:8080/web/24601
```

This conflict is no longer possible. In 7.0, a site's friendly URL is not allowed to be numeric. See for yourself:

1. Navigate to the site's *Configuration* → *Site Settings* → *Site URL* section.
2. In the Friendly URL section, enter *24601* and save the form. A failure message appears if you don't have a user with the matching screen name:

The friendly URL may conflict with another page.

You'll see this failure message if there's an existing conflict with a user screen name:

Please enter a unique friendly URL. Site [user-first-name user-last-name] has the same friendly URL.

Next, learn about collecting users in organizations.

ORGANIZATIONS

An *Organization* groups *Users* hierarchically. For example, you can model a company's departments (i.e., Human Resources and Customer Support) with Organizations. Organizations often have their own Sites. The *how-to* portion of managing Organizations is in the next article, *Managing Organizations*. This article contains important conceptual information on what Organizations are and when they're needed.

Many simple portal designs don't use Organizations at all; they only use sites. The main purpose of Organizations is to enable distributed User management. Portal administrators can delegate some user management responsibilities to Organization administrators. If you don't anticipate needing to delegate User management responsibilities, your portal design probably doesn't need to include Organizations.

User Groups and Organizations: It's easy to confuse User Groups (covered in a separate article) with Organizations since they both group Users. User Groups are an ad hoc collection of Users, organized for a specific function. In the Lunar Resort, if you wanted a group of bloggers, for example, it wouldn't make sense to assign the Sales Department the role of blogging (see the article on Roles if you're not sure what they are). The Sales Department users could blog whenever a new T-shirt design became available in the Lunar Resort store, but they probably wouldn't be as diligent about announcing the new Rover Racing schedule. Instead, creating a User Group containing one individual from each department who is responsible for blogging would make more sense. Read the article on User Groups to learn more about how to use them in your portal.

32.1 When to Use Organizations

To decide whether your portal design should include Organizations, think about its function. A photo-sharing web site could be powered by Sites only. On the other hand, Organizations are useful for corporations or educational institutions since their users can be placed into a hierarchical structure. Don't think that Organizations are only for large enterprises, though. Any group hi-

erarchy, from large government agencies all the way down to small clubs, can be modeled with Organizations. Also, don't think that you must decide between an Organization-based structure or a Site-based structure for assembling your portal's Users. Users can belong both to Organizations and to independent Sites. For example, a corporation or educational institution could create a social networking site open to all Users, even ones from separate Organizations.

To illustrate what an Organization is, consider a potential Organization of the Lunar Resort's Intranet. The company hierarchy has three tiers: The Lunar Resort, its departments, and divisions within each department.

- Lunar Resort—The top-level Organization.
 - Physical Plant Department—Department of users that keep the place running.
 - * Grounds Crew—Users that maintain the grounds.
 - * Janitorial Crew—Users who keep the resort clean.
 - * Mechanical Crew—Users who fix stuff, like lunar rovers.
 - Recreation Department—A department that makes sure much fun is had by guests of the Lunar Resort.
 - * Golf Instructors—Teach guests how to golf on the moon.
 - * Rover Race Instructors—Teach guests how to drive the lunar rovers.
 - * Lunar Sherpas—Lead guests on moon hikes.
 - Sales Department—A department of users who sell things to Lunar Resort guests.
 - * Up-sale Group—Make sure guests know how easy it is to improve their stay by spending more money.
 - * Souvenir and Memorabilia Group—Peddle souvenirs to Lunar Resort guests.
 - * Retail Group—Maintain the Lunar Resort store, which contains basic necessities, since guests are coming all the way from Earth.
 - Sentient Organism Resources Department—Department of Users that hire, fire and regulate intra-company relationships. We'd call it Human Resources, but what's stopping Martians from applying? Nothing!

Each department is a sub-Organization of the resort, and each division is a sub-Organization of the department.

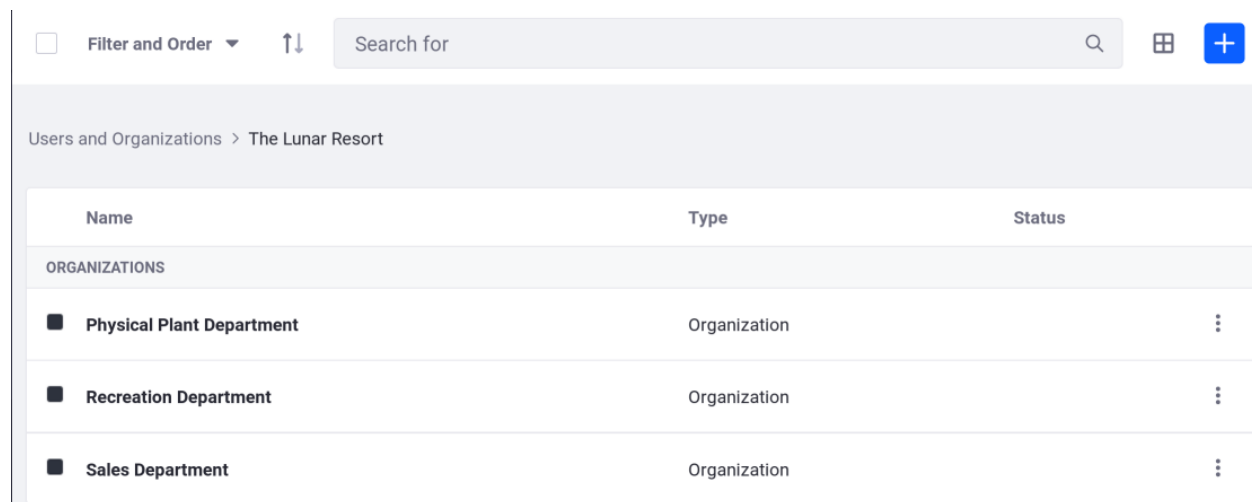
32.2 What can Organization Administrators Do?

Whenever you have a collection of Users that fits into a hierarchical structure, you can use Organizations to model those Users. Organization administrators can manage all the Users in their Organization *and* in any sub-Organization. Referring to the hierarchy above, for example, an Organization administrator of the Lunar Resort could manage any Users belonging to the resort itself, to any of the departments, or to any of a department's subdivisions. An Organization Administrator of the Physical Plant Department can manage any Users belonging to the Physical Plant Department itself, or to the Grounds Crew, the Janitorial Crew, or the Mechanical Crew. However, an administrator of the Physical Plant Department can't manage Users belonging to the Recreation Department or the Retail Group organization.

Organizations and sub-Organization hierarchies can nest to unlimited levels. Users can be members of one or many Organizations. The rights of an Organization administrator apply both to his/her Organization and to any child Organizations. Members of child Organizations are implicit members of their parent Organizations. This means, for example, that members of child Organizations can access the private pages of their parent Organizations. This behavior can be customized in the Organizations section of the portal-ext.properties file where the properties specific to organizations are listed.

Since Organizations are designed for distributed user administration, Organization Administrators have an entirely different set of privileges than Site Administrators. Site Administrators are responsible for the pages, portlets, content, and membership of their Sites. To this end, they can set the membership type to Open, Restricted, or Private. They can also add Users to or remove Users from their Sites but cannot manage the Users themselves. If an Organization has a Site attached to it, the Organization Administrator has the same rights as a Site Administrator for managing the Site's content, but an Organization Site's members are the members of the Organization. Thus Organization administrators have more user management permissions than Site administrators: they can edit users belonging to their Organization or any sub-Organization. They cannot add existing portal Users to their organization, but they can create new Users within their Organization. Only portal administrators can add existing users to an Organization.

Organization Administrators can't access the Control Panel by default, but it's not necessary. In their personal Sites, Organization administrators can click the *My Organizations* link to gain access to any Organizations they manage.



The screenshot shows a web interface for managing organizations. At the top, there is a navigation bar with a 'Filter and Order' dropdown, a search bar, and a '+ Add' button. Below the navigation bar, the breadcrumb 'Users and Organizations > The Lunar Resort' is visible. The main content area is a table with three columns: 'Name', 'Type', and 'Status'. The table lists three organizations under the heading 'ORGANIZATIONS': 'Physical Plant Department', 'Recreation Department', and 'Sales Department', all of which are of type 'Organization' and have a vertical ellipsis icon in the 'Status' column.

| Name | Type | Status |
|---------------------------|--------------|--------|
| ORGANIZATIONS | | |
| Physical Plant Department | Organization | ⋮ |
| Recreation Department | Organization | ⋮ |
| Sales Department | Organization | ⋮ |

Figure 32.1: The My Organizations application lets Organization Administrators manage their organizations in their personal site.

32.3 Organization Roles and Permissions

A huge time-saving benefit of including Organizations into your portal design is that Organization administrators can assign Organization-scoped Roles to members of the entire Organization. For example, consider an IT Security group in a corporate setting. You could have a sub-Organization of your IT organization that handles security for all applications company-wide. If you grant the IT Security Organization the portal administrator Role, all members of the Organization get administrative access to the entire system. Suppose further that a User in this Organization was later hired by the Human Resources department. The simple act of removing the User from the IT Security Organization also removes the User's administrative privileges, since the privilege came from the IT Security Organization's Role. By adding the User to the HR Organization, any roles the HR Organization has (such as access to a benefits system in the portal) are transferred to the User. In this manner, you can design your portal to correspond with your existing organization chart and Users' permissions are granted according to their positions in the chart.

Of course, this is only one way to set things up. If you have more complex requirements for permissions within an Organization, you can create custom Organization-scoped Roles to assemble the permissions you wish to grant to particular Users. Alternatively, attach a Site to your Organization and use Site Teams to assemble the sets of permissions (see below). See the Roles and Permissions article for more detail.

32.4 Organization Sites

Does your Organization need to have its own Site? If an organization has an attached Site, the Organization's administrators are treated as the Site administrators. This means that they can manage the pages, portlets, and content of the Site as well as the Users of the Organization. Members of an Organization with an attached Site are treated as members of the Organization's Site. This means that they can access the private pages of the Organization's Site, along with any portlets or content there. Attaching Sites to Organizations allows portal administrators to use Organizations to facilitate distributed portal administration, not just distributed User administration.

That's a lot of information on Organizations. Next, learn how to create and manage Users and Organizations.

MANAGING ORGANIZATIONS

If you're not entirely sure what Organizations are or whether you need them, start here. This article gets right to the practical stuff: how to manage Organizations.

33.1 Adding Organizations

Add an Organization (perhaps start by adding the *Physical Plant Department* organization to the Lunar Resort):

1. Click *Users and Organizations* from Control Panel → Users.
2. Go to the *Organizations* tab and click the *Add* button. Fill out the Name field at a minimum.
3. If you're creating a child Organization, use the Parent Organization *Select* button to select an Organization in the system to be the direct parent. Click the *Remove* button to remove the currently configured parent.
4. Click *Save* when finished filling out the Add Organization form.

As when creating a new user, once you submit the form a success message appears and you have access to a new form which lets you enter additional information about the Organization. Organizations can have associated multiple email addresses, postal addresses, web sites, and phone numbers. The Services link can be used to indicate the operating hours of the Organization, if any.

Tip: After creating an Organization, assign the desired user to the Organization Owner Role. The Organization Owner can do everything that an organization Administrator can. In addition to their full administrative rights within the Organization, they can do these things:

- Appoint other Users to be Organization Administrators
- Appoint other Users to be Organization Owners
- Remove the memberships of other Organization Administrators or Owners

Organization Administrators can't make these Role assignments and can't manage the memberships of other Organization Administrators or Owners.

33.2 Editing Organizations

To edit an Organization, go to the Users and Organizations section of the Control Panel and click the *Organizations* tab. All active Organizations are listed. Click the *Actions* button next to an Organization. This shows a list of actions you can perform on this Organization.

- *Edit* lets you specify details about the Organization, including addresses, phone numbers, and email addresses. You can also create a Site for the Organization.
- *Manage Site* lets you create and manage the public and private pages of the Organization's Site. This only appears after a Site has been created for the Organization.
- *Assign Organization Roles* lets you assign Organization-scoped Roles to Users. By default, Organizations are created with three Roles: Organization Administrator, Organization User and Organization Owner. You can assign one or more of these Roles to Users in the Organization. All members of the Organization automatically get the Organization User Role so this Role is hidden when you click Assign Organization Roles.
- *Assign Users* lets you search and select Users to be assigned to this Organization as members.
- *Add User* adds a new User and assigns the User as a member of this Organization.
- *Add Organization* lets you add a child Organization to this Organization. This is how you create hierarchies of Organizations with parent-child relationships.
- *Delete* removes this Organization. Make sure the Organization has no Users in it first. You'll be prompted for confirmation that you want to delete the Organization. If there are Users in the Organization or if there are sub-Organizations, you must remove the Users and delete the sub-Organizations before deleting the parent Organization.

If you click the Organization name you can view both a list of Users who are members of this Organization and a list of all the sub-Organizations of this Organization.

33.3 Organization Types

By default, Liferay DXP only includes the *Organization* type. Configure the existing type or add additional types using the aptly named Organization Type entry in System Settings. There are two main reasons to configure Organization types:

1. Organizations usually correlate to real-life hierarchical structures. Calling them by their real names is helpful for administrators and Users. In the Major League Baseball (MLB) example, *League*, *Division*, and *Team* Organization types are useful.
2. Enforce control over which Organizations can be top level Organizations and the type of sub-Organization allowed for each parent Organization type. For example, MLB would not allow Division Organization types to be sub-Organizations of Team Organizations.

Organization Type

This configuration was not saved yet. The values shown are the default.

Name

League

Country Enabled

Country Required

Rootable

Children Types +

Division

Save Cancel

Figure 33.1: Create new organization types through the System Settings entry called Organization Types.

Check out the configuration options that configure the default *Organization* type and then configure an additional type.

To add another Organization type called *League*, enter these options into the configuration form:

Name: *League* Adds League to the list of Organization types that appear in the Add Organization menu.

Country Enabled: *True* Enables the Country selection list field on the form for adding and editing League types.

Country Required: *False* Specifies that the *Country* field is not required when adding a League.

Rootable: *True* Enables Leagues as a top level Organization. Limit League to sub-Organization status by excluding this property.

Children Types: *Division* Specifies Division as the only allowable sub-Organization type for the League parent type.

Once you configure additional Organization types and click Save, you'll find your new type(s) available for selection in the Add Organization form.

The screenshot shows a mobile application interface for adding a division. The title bar at the top is dark with a back arrow and the text "Add Division". The form is white and contains the following elements:

- Name:** A text input field containing "AL East".
- Type:** A dropdown menu currently showing "Division".
- Parent Organization:** A table with two columns: "Name" and "Type". It contains one row: "American League" (Name) and "League" (Type). A "Remove" button with an "x" icon is to the right of the row.
- Select:** A button below the parent organization table.
- Organization Site:** A section with a "Create Site" checkbox and a downward arrow.
- Categorization:** A section with a rightward arrow.
- Buttons:** "Save" (blue) and "Cancel" (white) buttons at the bottom.

Figure 33.2: Custom configuration types are available in the Add Organization form.

Users can join or be assigned to Sites when they share a common interest. Users can be assigned to Organizations when they fit into a hierarchical structure. User groups provide a more ad hoc way to group users than sites and Organizations. You'll look at them next.

ROLES AND PERMISSIONS

If a *Role* were to win a Grammy or an Oscar or some other ego-feeding popularity contest, it better remember to thank all its *permissions* groupies during the acceptance speech, because they're the ones doing the real work. The Role is just the pretty face, so to speak.

Roles collect permissions that define a particular function, according to a particular scope. Roles collect permissions, and Users are assigned to Roles.

Note: Roles are assigned to Users, but it's tedious to assign each User to a Role intended for lots of Users. Recall that Users are grouped in Sites, Organizations, and User Groups. Implicitly assign regular scoped permissions to Users by assigning a Role directly to one of these User groupings.




Figure 34.1: Assign Users to a role, directly or by their association with a Site, Organization, or User Group.

Take a Message Board Administrator Role, for example. A Role with that name should have permissions relevant to the specific Message Board portlets delegated to it. Users with this Role inherit the permissions collected underneath the umbrella of the Role.

In addition to regular Roles, Site Roles, and Organization Roles, there are also Teams. Teams can be created by site administrators within a specific Site. The permissions granted to a Team are defined and applied only within the Team's site. The permissions defined by regular, Site, and Organization Roles, by contrast, are defined at the global level, although they are applied to

different scopes.

Regular role Permissions are defined at the global level and are applied at the global scope.

Site role Permissions are defined at the global level and are applied to one specific Site.

Organization role Permissions are defined at the global level and are applied to one specific Organization.

Team Permissions are defined within a specific Site and are assigned within that specific Site.

Note: Some permissions cannot be handled from the control panel. Asset-level permissions (for instance, permission to edit an individual blog post, or view a folder in the Documents and Media library) are managed from the individual asset. See Widget Permissions for details.

34.1 Deleting Asset Containers

A Web Content Folder contains Web Content articles. The Web Content Folder is an asset container, and the Web Content Article is an asset. It's possible to give a Role permission to delete an asset container without giving the Role permission to delete individual assets. In that case, beware: if a Role assignee deletes an asset container with individual assets in it, the individual assets themselves are deleted as well.

Besides Web Content Folders, examples of asset containers include Bookmarks Folders, Message Boards Categories, Wiki Nodes, and Documents and Media Folders.

You might not need to create a Role for a certain functionality. Liferay provides many pre-configured Roles for your convenience.

34.2 Default Liferay Roles

In the Roles Application appears a list of all the Roles in Liferay, by scope.

These are some of the pre-configured regular Roles:

- **Guest:** The Guest role is assigned to unauthenticated users and grants the lowest-level permissions.
- **User:** The User role is assigned to authenticated Users and grants basic permissions (mostly *Add to Page* permissions for their own Sites).
- **Power User:** The Power User Role grants more permissions than the User Role. It's an extension point for distinguishing regular Users from more privileged Users. For example, you can set things up so that only Power Users have personal sites.
- **Administrator:** The administrator Role grants permission manage the entire portal, including global portal settings and individual Sites, Organizations, and Users.

These are some of the pre-configured site roles:

- **Site Member:** The Site Member Role grants basic privileges within a Site, such as permission to visit the Site's private pages.
- **Site Administrator:** The Site Administrator Role grants permission to manage *almost* all aspects of a Site including site content, site memberships, and site settings. Site Administrators cannot delete the membership of or remove roles from other Site Administrators or Site Owners. They also *cannot* assign other Users as Site Administrators or Site Owners.
- **Site Owner:** The Site Owner Role is the same as the Site Administrator Role except that it grants permission to manage *all* aspects of a Site, including permission to delete the membership of or remove Roles from Site Administrators or other Site Owners. They *can* assign other users as Site Administrators or Site Owners.

These are some of the pre-configured organization roles:

- **Organization User:** The Organization User role grants basic privileges within an Organization. If the Organization has an attached Site, the Organization User Role implicitly grants the Site member Role within the attached Site.
- **Organization Administrator:** The Organization Administrator Role grants permission to manage *almost* all aspects of an Organization including the Organization's Users and the Organization's Site (if it exists). Organization Administrators cannot delete the membership of or remove Roles from other Organization Administrators or Organization Owners. They also *cannot* assign other Users as Organization Administrators or Organization Owners.
- **Organization Owner:** The Organization Owner Role is the same as the Organization Administrator Role except that it grants permission to manage *all* aspects of an Organization, including permission to delete the membership of or remove Roles from Organization Administrators or other Organization Owners. They *can* assign other Users as Organization Administrators or Organization Owners.

Tip: It's easy to overlook the differences between owner type roles and administrator type roles for Sites and Organizations. Site and Organization administrators *cannot* remove the administrator or owner Role from any other administrator or owner, and they *cannot* appoint other Users as site or organization administrators or owners.

In contrast, site and organization owners *can* do those things.

Roles, and the permissions granted with their assignment, are foundational components in Liferay. Understanding their uses and configuration enhances your ability to configure Liferay DXP to suit your organizational needs.

MANAGING ROLES

You manage Roles and Permissions in the Control Panel (*Control Panel* → *Users* → *Roles*). There you'll find an application for creating Roles, granting them permissions, and assigning Users to them. Roles can be scoped by portal, Site, or Organization.

Defining a Role's permissions is a topic deserving its own article. Read here about defining a role's permissions.

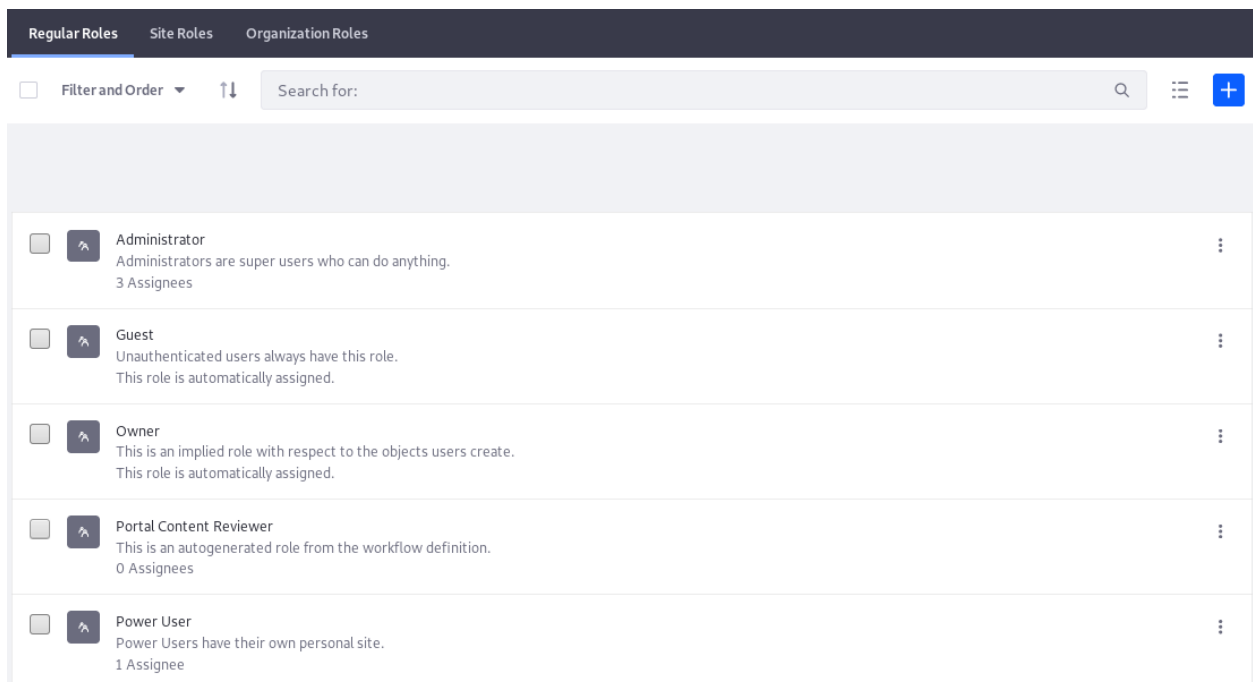


Figure 35.1: The Roles application lets you add and manage roles for the global (Regular), Site, or Organization scope.

35.1 Creating Roles

Determine the scope of the Role you must create. Roles can be scoped globally (Regular Roles), to a specific Site (Site Roles), or to an Organization (Organization Roles).

To create a Role:

1. Click the tab for the proper Role scope, then click the *Add* (+) button.
2. Enter a title and description. The title field is required but the description is optional.
3. Enter a Key, if desired. This required field provides a key that can be used to refer to the Role programmatically. It's auto-populated with the title text, but you can override it if desired.
4. Click *Save*.

Now the Role is present in the database and ready for further configuration.

35.2 Assigning Users to a Role

Assign users to a Role in the Assignees tab of the Add/Edit Role form. Roles are assigned to Users, Sites, Organizations, or User Groups. Here's how to assign the User Group Manager Role created in the last section to Users:

1. In the Assignees tab of the Add/Edit Role form, click the second level tab for *Users*.
2. Click the Add button (+).
3. Select the Users you want to add to the Role and click *Add*.

If assigning a group, note that all Users assigned to that group inherit the Role as well.

That's a good start, but your Role isn't worth the database row it occupies without defining permissions for the Role. Read the next article to learn how.

DEFINING ROLE PERMISSIONS

Roles collect permissions, so when Users are given a Role, they receive all the permissions defined by the Role.

If you create a Role with permission to access something in the Control Panel, keep in mind that the *View Control Panel Menu* permission is automatically granted.

Consider a Role called User Group Manager. Define the permissions for the User Group Manager Role so that assigned Users can add Users to or remove Users from any User Group:

1. Go to the Control Panel and then click on *Users* → *Roles*.
2. On the Regular Roles screen, click *Add* (+).
3. After naming your Role, click *Save*.
4. Click on the *Define Permissions* tab.
5. Drill down in the menu on the left to *Control Panel* → *Users* → *User Groups*.
6. Under the *General Permissions* heading, flag *Access in Control Panel* and *View*. This lets user group managers access the User Groups Control Panel portlet and view existing User Groups.
7. Since you want User Group managers to be able to view User Groups and assign members to them, also check the *Assign Members* and *View* permissions under the *Resource Permissions* → *User Group* heading.
8. There's one last necessary permission you might not think of in association with this Role. In *Control Panel* → *Users* → *Users and Organizations*, User Group managers need *View* permission on the User resource. Grant this permission.
9. Click *Save*.

Now the User Group Manager Role has all the permissions necessary for adding Users to User Groups. After all, User Group managers can view User Groups, assign members, and access User Groups in the Control Panel. The permission to view Users in the Control Panel was necessary because you must view Users to assign them as members of a Role. Without this permission, User Group managers see an empty list if they try to add Users to a Role.




Figure 36.1: When defining permissions on a Role, the Summary view provides a list of permissions that have already been defined for the role. The area on the left side of the screen lets you drill down through various categories of permissions.

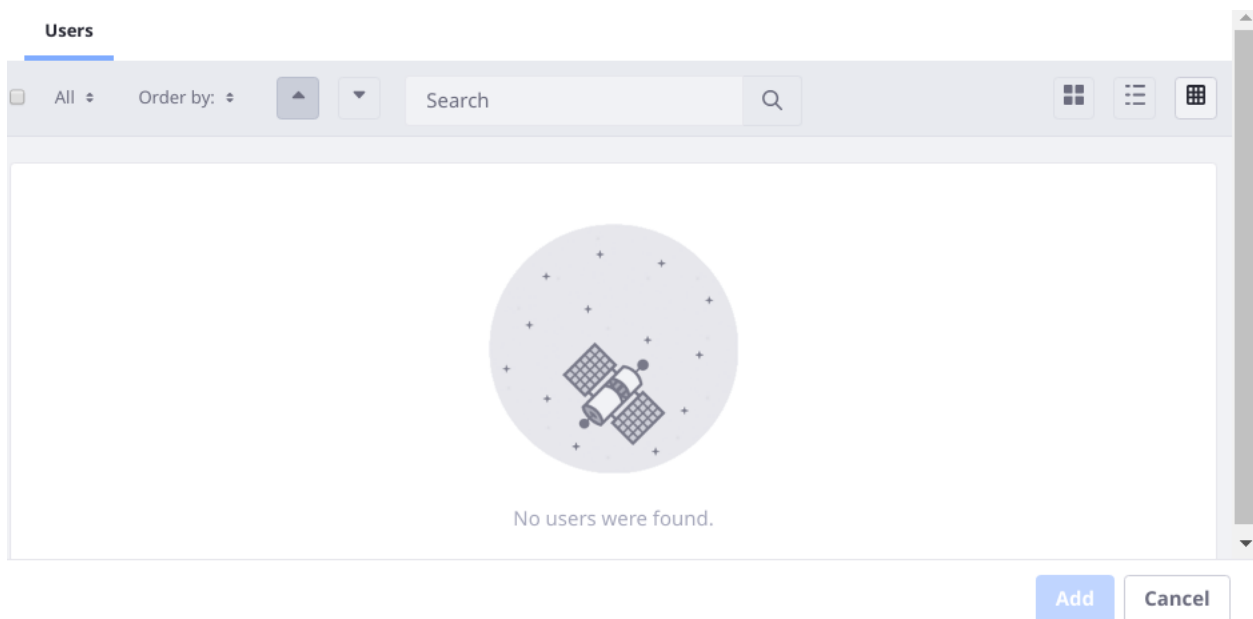


Figure 36.2: Users assigned to the User Group Manager Role can't find any users to add unless they have view permissions on the User resource.

Note: The Roles application in the Control Panel is not the only place where permissions are configured. You can configure a Role's permissions on a resource at a more granular level. For example, from a particular application instance, click its *Options* (⚙️) menu and select *Permissions*. There you can configure permissions for the resource that overlap with those configured in the Control Panel's Roles application. However, permissions granted or removed in the Control Panel override those made at the more granular level.

There are three categories of permissions: *Control Panel*, *Site Administration*, and *User*. By default, Users can manage their User accounts via the permissions belonging to the User category. Site Administrators can access the site administration tools belonging to the Site Administration category. Portal Administrators can access the entire Control Panel. For custom Roles, you can mix and match permissions from as many categories as you like.

The permissions in the Site Administration → Applications categories govern the content that can be created by portlets such as the Wiki and Message Boards. If you pick one of the portlets from this list, you see options for defining permissions on its content. For example, if you pick Message Boards, you see permissions for creating categories and threads or deleting and moving topics.

Site application permissions affect the application as a whole. Using the Message Boards as an example, an application permission might define who can add the Message Boards portlet to a page.

The Control Panel permissions affect how the Control Panel appears to the User in the Control Panel. The Control Panel appears differently for different Users, depending on their permissions. Some Control Panel portlets have a Configuration button, and you can define who gets to see that. You can also fine-tune who gets to see various applications in the Control Panel.

If you want to change the scope of a permission, click the *Change* link next to the gear icon next to the permission and then choose a new scope. After you click *Save*, you'll see a list of all permissions currently granted to the Role. From the Summary view, you can add more permissions or go back to the Role Application default view by clicking on the *Back* (⏪) icon.

Sometimes you might find that a certain permission grants more or less access than what you expected—always test your permissions configurations!

36.1 Delegating Social Activities Configuration

There's a permission that allows Site administrators to delegate responsibility for configuring social activities to other Users. To add this permission to a Role, click *Actions* next to the desired Role and select *Define Permissions*. Find the *Site Administration* → *Configuration* → *Social Activity* permissions category. Flag all of the permissions and then click *Save*:

- Access in Site Administration
- Configuration
- Permissions
- Preferences
- View

The screenshot displays the 'Message Boards' configuration page. On the left is a sidebar with a search bar and a list of navigation items: Summary, Control Panel, Site Administration, Build, Content, Web Content, Blogs, Documents and Media, Dynamic Data Lists, Forms, Knowledge Base, Dynamic Data Mapping Data P, Message Boards, Polls, Wiki, Kaleo Forms Admin, and Categorization. The main content area is titled 'Message Boards' and is divided into three sections:

- GENERAL PERMISSIONS**: A table with 6 rows, each with a checkbox, an action name, and site access options.

| Action | Sites |
|-------------------------------|-----------|
| Action | Sites |
| Access in Site Administration | All Sites |
| Configuration | All Sites |
| Permissions | All Sites |
| Preferences | All Sites |
| View | All Sites |
- RESOURCE PERMISSIONS**: A section header with a gear icon.
- MESSAGES**: A table with 5 rows, each with a checkbox, an action name, and site access options.

| Action | Sites |
|--------------|-----------|
| Add Category | All Sites |
| Add File | All Sites |
| Add Message | All Sites |
| Ban User | All Sites |

Figure 36.3: You can fine-tune which actions are defined for a role within a specific application like the Message Boards.

Once these permissions are granted, Role assignees can manage the site's Social Activities.

Roles allow portal administrators to define various permissions in whatever combinations they like. This gives you as much flexibility as possible to build the Site you have designed.

MANAGING USER DATA

Internet users are increasingly and justifiably concerned about how their personal data is processed by the systems they use. The enforcement of GDPR is a crystallization of these concerns into legislative action. Companies processing the personal data of EU residents must adopt appropriate measures to protect User data. Of course, legal requirements like those in GDPR only explain one reason for companies to develop policies for ensuring their users' right to privacy. The market demands site owners show higher levels of responsiveness to User inquiries into how their data is stored and processed.

Liferay is aware of the need for functionality to address User data management, and added two important features toward this end:

Erase and/or anonymize data associated with a User Administrative Users go through a step by step process, choosing to erase certain pieces of data and anonymize others.

Export a User's personal data Export ZIP files containing the data associated with a User.

These features are tools that get you closer to meeting two of GDPR's technically challenging requirements, the *right to data portability* and the *right to be forgotten*.

Note: It is Liferay's sincerest hope that through the User Management functionality of Liferay DXP, companies processing the personal data of their website's users can satisfy the requirements of GDPR. However, the tools discussed here and anywhere else in the documentation, including those directly aimed at addressing GDPR requirements (as in this article) do not guarantee compliance with the legal requirements of GDPR. Each company or individual whose website processes user personal data and is under the jurisdiction of GDPR must carefully determine the precise steps necessary to ensure they are fully compliant with GDPR.

37.1 Anonymizing Data

Deleting User data is the safest way to honor *right to be forgotten* requests. When User data must be preserved, automatic anonymization of the data is in order. Users being anonymized must have their identifiers (for example, User ID and User Name) removed from content they've interacted with. However, portal content usually requires this information for its applications to work properly. Therefore, the User's identifiers must be replaced by something, or someone. Meet the new User, *Anonymous Anonymous*, identity swapper *extraordinaire*. This deactivated User is dedicated to be the User whose identifiers are assigned to anonymized content. This identity swap is an important step in the anonymization process, but additional manual intervention may be necessary to truly achieve anonymization.


as11-44-6667.jpg



✦ Uploaded by Anonymous Anonymous, 6/14/18 6:38 PM Average (0 Votes)
☆☆☆☆☆

Figure 37.1: Anonymized content is presented with the User Anonymous Anonymous's identifying information.

Here at Liferay, we've grown fond of *Anonymous Anonymous*. If you'd rather start from scratch or assign an existing User to be the Anonymous User, get rid of *Anonymous Anonymous* and configure your own Anonymous User.

The anonymous user is programmatically created for each instance the first time an Administrator clicks *Delete Personal Data* from a User's Actions menu (). If you haven't yet done that, no Anonymous User exists.

The easiest way to set up a new User as the Anonymous User is to edit an existing Anonymous User configuration, passing in the new Anonymous User's User ID.

To edit an existing configuration:


1. Go to Control Panel → Configuration → System Settings → Users → Anonymous User.
2. Edit the existing configuration, providing a different User ID.

Get the User ID from Control Panel → Users → Users and Organizations. Click on the User and find the User ID in the Information screen of the Edit User application.

3. Click *Update*.

To create a new Anonymous User:

1. Create a User use for data anonymization. Alternatively, you can use an existing User.
2. If there's already an Anonymous User configured for the instance, there are two ways to remove it:

Delete the User entirely. Deleting the User simultaneously deletes its configuration as the Anonymous User. Go to Control Panel → Users → Users and Organizations. If it's an active User, first deactivate, then delete the User. The default Anonymous Anonymous User is deactivated by default. Simply delete the User in this case. Click the Actions button () and select *Delete*.

If you don't want to delete the User, just delete the User's configuration as the Anonymous User. Go to Control Panel → Configuration → System Settings → Users → Anonymous Users.

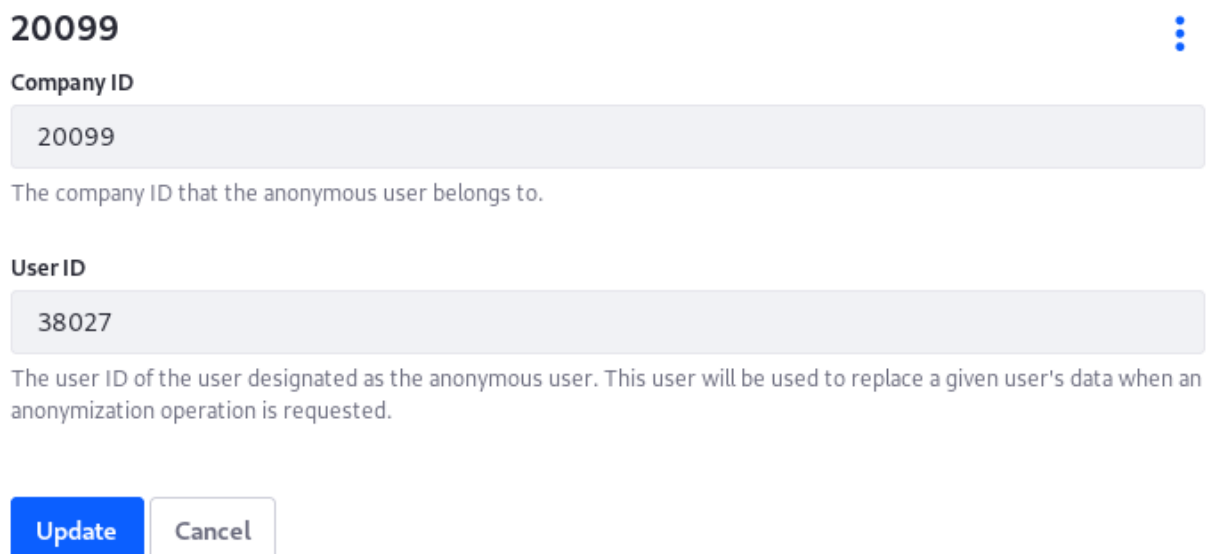
3. Add a new Anonymous User configuration. Click the *Add* button.


4. Fill out the two fields, Company ID and User ID.

Get the Company ID from Control Panel → Configuration → Virtual Instances. The Instance ID and Company ID are the same.

Get the User ID from Control Panel → Users → Users and Organizations. Click on the User and find the User ID in the Information screen of the Edit User application.

There can only be one Anonymous User configured for each instance.



20099 

Company ID

20099

The company ID that the anonymous user belongs to.

User ID

38027

The user ID of the user designated as the anonymous user. This user will be used to replace a given user's data when an anonymization operation is requested.

Update **Cancel**

Figure 37.2: Assign your own Anonymous User from Control Panel → Configuration → System Settings → Users → Anonymous User.

37.2 Manual Anonymization

Anonymizing just the User's identification fields is often not enough. If a User named Ziltoid Omniscient complains about The Lunar Resort's coffee in a Message Boards Message and in it

signs the post with *Supreme Leader of Ziltoidea 9*, anonymizing this post would remove the User's name (Ziltoid Omniscient) and replace it with Anonymous Anonymous, but searching the Internet for *Ziltoidea 9* quickly reveals that the post was written by Ziltoid the Omniscient. There can be user-entered personal data within the content of an application. You must manually edit such content to remove identifying details.

Announcing Fresh Roasted, Lunar-Grown Coffee

We're tireless when it comes to improving things at the resort, but caffeine helps.



As the head of Culinary Engineering at the resort, I'm delighted to announce that our hydroponically grown coffee is now available in The Lunar Resort Café.

Comment (1) 0 0 P



1 Comment

[Subscribe to Comments](#)



Type your comment here.

Reply



Anonymous Anonymous

Reply 0 0

42 Seconds Ago

Been there, tried it. Fetid and Foul! The moon environment is not suitable for such an endeavor. Wake up and smell the coffee!

Sincerely, The Supreme Leader of Ziltoidea 9



0 0

Reply

Figure 37.3: Even though this Message Boards Message (a comment on a blog post in this case) is anonymized, it should be edited to remove User Associated Data from the content of the message.

SANITIZING USER DATA

One of the technically challenging requirements of the General Data Protection Regulation (GDPR) is *the right to be forgotten*. The purpose of this article is not to go into the details of this requirement, but to show you how the personal data erasure functionality can assist you in satisfying this requirement.

A simple way to think of what it means to be *forgotten* by software is to consider a scenario where a new portal administrator is hired immediately after a User's right to be forgotten request has been honored. The new portal administrator has access to all of the Site's content and administration capabilities. Despite this, the administrator must not be able to glean information that could lead her to knowing the identity of the User whose personal data was erased.

Conceptually, forgetting a User means two things, at a minimum:


- Erasing the User's identifying information from the system. In Liferay DXP, this entails removing the User from database tables and search indexes.
- Erasing or anonymizing content the User has interacted with so it cannot be tracked to a real person.

Users can already be deactivated and then deleted, so why add new functionality? Deleting removes the User from the table of Users in the database. The User's information is preserved in other locations, however. In a standard User deletion scenario, all of a User's personally created content is still assigned to the User and her identifiers (User ID and User Name) still appear in the UI next to content associated with her. This unintentional preservation of user-identifying data is inadequate for satisfying some of the GDPR requirements and is the primary reason why the data erasure functionality was added in 7.0.

Note: Personal data erasure can help companies in their attempts to satisfy the requirements of GDPR. Using the data erasure tool described here provides no guarantee of compliance with the legal requirements of GDPR. Each company or individual whose website processes user personal data and is under the jurisdiction of GDPR must carefully determine the precise steps necessary to ensure they are fully compliant with GDPR.

To begin sanitizing a user's data,

1. Go to Control Panel → Users → Users and Organizations.

2. Click the Actions button for a User () and select *Delete Personal Data*. If you have not deactivated the user, you will be asked to do so.

The User's Personal Data Erasure screen appears.

38.1 The Personal Data Erasure Screen

You can browse all data the user has posted on the system. Click *Personal Site* to browse data from that site.




Figure 38.1: From here, you can browse all data the user posted on his or her personal Site.

Click *Regular Sites* to browse any data posted in regular Liferay sites.

To review the user's data, click the item. For example, Pepper seems to have posted a blog entry on her personal Site. Clicking that entry reveals the title of that blog entry.

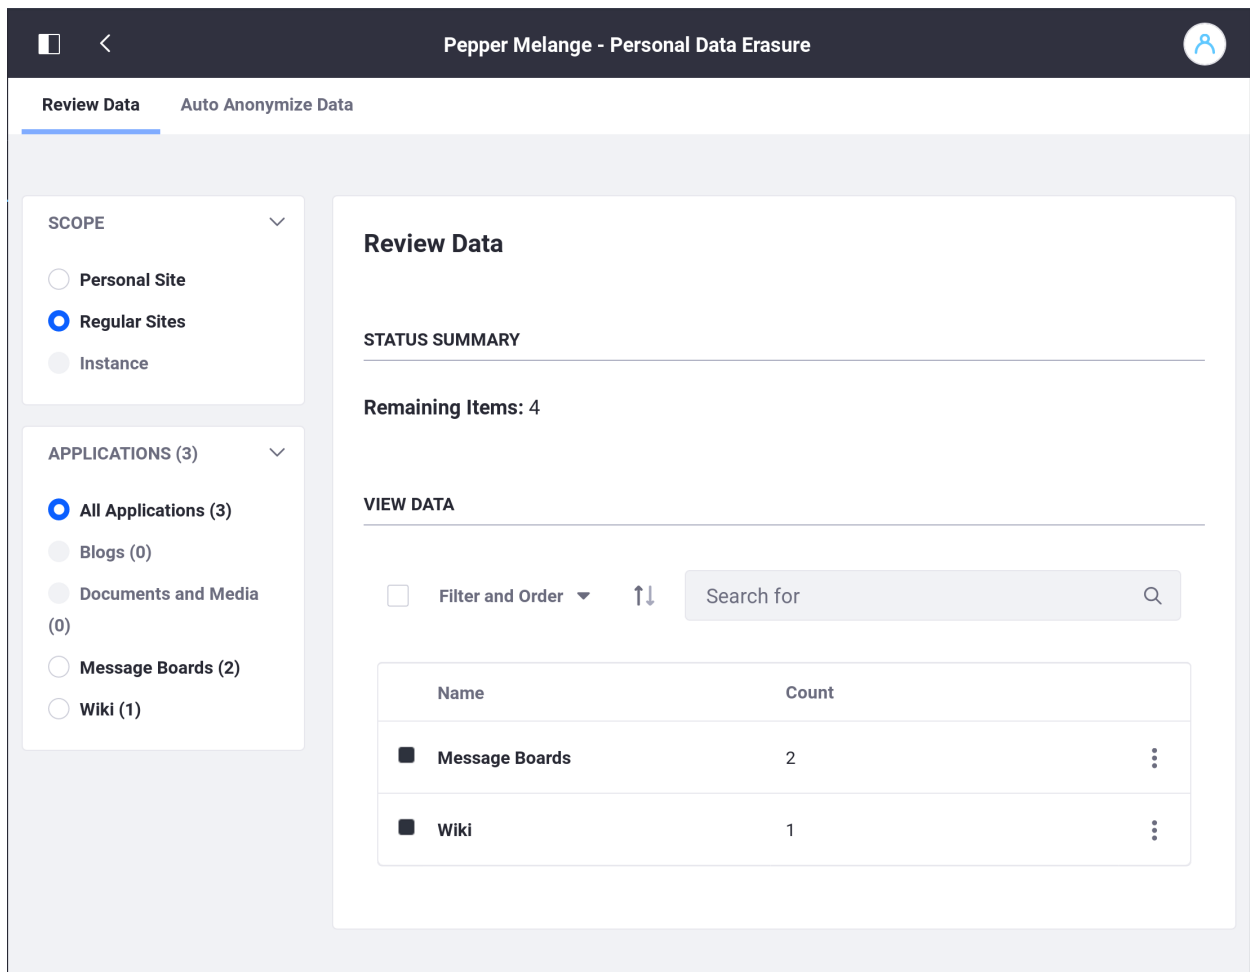


Figure 38.2: Choose Regular Sites to browse all data posted by the user on administratively-created Sites.

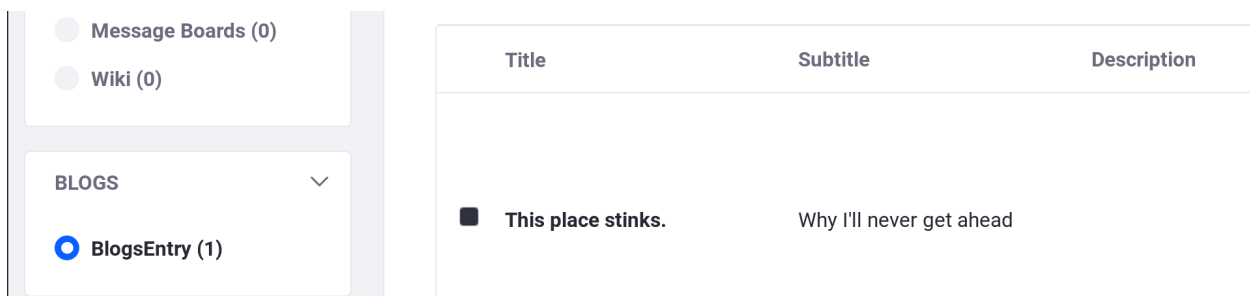



Figure 38.3: Pepper's blog entry might need review.

To review any entry, click it. You're brought to the edit mode of the application (in this case, Blogs), where you can make any changes to the content that are necessary.

To manage (anonymize or delete) all the items for an application at once:

1. Click the Actions button () for the application.
2. If you're sure all items for an application can be safely deleted, choose *Delete*.
3. If you're sure simple anonymization is good enough for all of an application's items, choose *Anonymize*.

Use the interface to browse through the Sites, applications, and data.

38.2 Delete the User

Once all data is reviewed, deleted, edited, and/or anonymized as appropriate, delete the User. A dialog box pops up automatically when you're finished. This step is simple: Click *OK*.

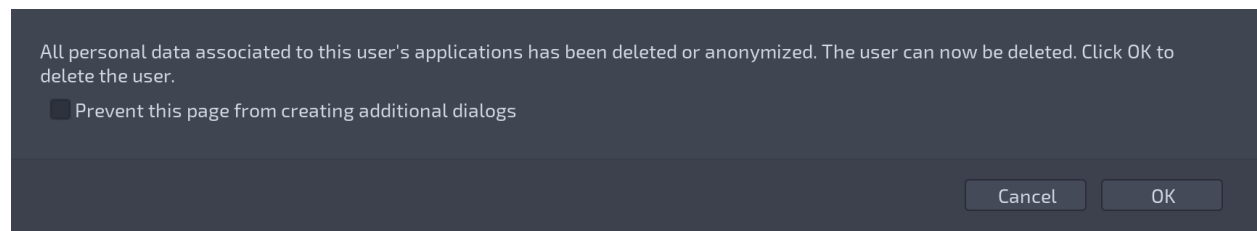


Figure 38.4: To finish the data erasure process, delete the User.

Now the User's data is anonymized or deleted, and the User is also deleted.

EXPORTING USER DATA

User Management practices must account for the EU's General Data Protection Regulation. One of its tenets is that Users have a right to *data portability*.



Data portability means that a User has the right to receive their personal data in a machine-readable format.

Note: Personal data export can help companies in their attempts to satisfy the requirements of GDPR. Using the export tool described here provides no guarantee of compliance with any GDPR requirement. Each company or individual whose website processes user personal data and is under the jurisdiction of GDPR must carefully determine the precise steps necessary to ensure they are fully compliant with GDPR.

The Control Panel's User Management system now natively supports the export of a User's personal data to a ZIP file for download. The data format for the files containing the data is XML.

39.1 Exporting and Downloading

To export user data,

1. Go to *Control Panel* → *Users* → *Users and Organizations*.
2. Find the User and click the Actions button (), then click *Export Personal Data*.
This opens the User's Export Personal Data screen.
3. If there are no existing export processes shown, there's only one thing to do: click the *Add* button (). The tool for exporting the User's data appears.
4. Most of the time you want to export all the available data. Click *Select Items*, and all applications containing User data are selected in the UI.

Export Personal Data

Please select the applications for which you want to start an export process.

Select Items

| Application | Items | Last Available Export |
|---|-------|-----------------------|
| <input checked="" type="checkbox"/> Announcements | 0 | - |
| <input checked="" type="checkbox"/> Blogs | 0 | - |
| <input checked="" type="checkbox"/> Bookmarks | 0 | - |
| <input checked="" type="checkbox"/> Contacts Center | 0 | - |
| <input checked="" type="checkbox"/> Documents and Media | 0 | - |
| <input checked="" type="checkbox"/> Message Boards | 2 | - |
| <input checked="" type="checkbox"/> Wiki | 0 | - |

20 Entries ▾ Showing 1 to 7 of 7 entries. ◀ 1 ▶

Export

Figure 39.1: The Export Personal Data tool lets you export all or some of the User's data.

5. Click *Export*. You're taken back to the User's Export Personal Data screen, but now there's an export process in the list.

| Export Processes | | | |
|------------------|------------------------------------|--|---|
| Filter and Order | | | |
| SUCCESSFUL | | | |
| Message Boards | Create Date: 2018.06.22 - 05:54 PM | Completion Date: 2018.06.22 - 05:54 PM | ⋮ |
| Message Boards | Create Date: 2018.06.22 - 05:47 PM | Completion Date: 2018.06.22 - 05:47 PM | ⋮ |

Figure 39.2: Once User data is successfully exported, the export process is displayed in the User's Export Personal Data list.

6. Download the data. Click the Actions button (⋮) for the process and select *Download*.

39.2 Examining Exported Data

So what does the exported data look like?

```
<?xml version="1.0"?>
<model>
  <model-name>com.liferay.message.boards.model.MBMessage</model-name>
  <column>
    <column-name>messageId</column-name>
    <column-value><![CDATA[38099]]></column-value>
  </column>
  <column>
    <column-name>statusByUserId</column-name>
    <column-value><![CDATA[38045]]></column-value>
  </column>
  <column>
    <column-name>statusByUserName</column-name>
    <column-value><![CDATA[Jane Slaughter]]></column-value>
  </column>
  <column>
    <column-name>userId</column-name>
    <column-value><![CDATA[38045]]></column-value>
  </column>
  <column>
    <column-name>userName</column-name>
    <column-value><![CDATA[Jane Slaughter]]></column-value>
  </column>
  <column>
    <column-name>subject</column-name>
    <column-value><![CDATA[Great list. I was thinking of bringing the family,...]]></column-value>
  </column>
  <column>
    <column-name>body</column-name>
    <column-value><![CDATA[<p>Great list. I was thinking of bringing the family, but I don't]]></column-value>
  </column>
</model>
```

```
    actually believe humans have ever been to the moon, so I guess it
    would be silly to book a trip! LOL!</p>]]></column-value>
  </column>
</model>
```

In this example, User Jane Slaughter made a Message Boards Message post, and her User information was recorded in the MBMessage model's database table.

This actually corresponds with a comment on a Blogs Entry:

Exporting User data informs Site owners and Users about how much personal data the sight may have.

Is a Lunar Resort Vacation Right for You?



Test Test
27 Minutes Ago

The Moon is beautiful this time of year, and we're offering all kinds of online specials that are enticing more people than ever to visit us! This makes us excited, but we understand that not everyone is cut out for a moon vacation. Here are the top ten signs that you might not be ready for a trip to The Lunar Resort:

10. You don't believe in the moon landing or The Lunar Resort's existence.

9.

Comment (1) 1 0



1 Comment

[Subscribe to Comments](#)



Type your comment here.

Reply



Jane Slaughter
25 Minutes Ago

Reply 1 0

Great list. I was thinking of bringing the family, but I don't actually believe humans have ever been to the moon, so I guess it would be silly to book a trip! LOL!



Figure 39.3: A Comment on a blog post is User Associated Data.

USER GROUPS

A User Group is a list of Users created for a specific purpose. User Groups can be created across the hierarchical boundaries of Organizations. For example, an administrator could create a Teachers User Group for adding all members to multiple Sites, assign them all to a Regular Role, and create a common set of profile pages for all teachers in the User Group. User Groups are integrated with Roles, Sites, Site Templates, and permissions. This flexibility means that there are many different use cases for User Groups. The articles in this section show you how to work with User Groups to serve the most common use cases.



User Groups are most often used in these scenarios:

- **Manage Site membership:** Grant Site membership to all Users in a User Group. Using the previous example, the Teachers user group could be added as a member of the Sites *University Employees* and *Students and Teachers Collaboration*. All users in that User Group would become members of those Sites.
- **Manage user personal pages:** Provide predefined public or private pages to the users in the user group. For example, the Teachers user group could be created to ensure the home page on all teachers' personal Sites has the same layout and applications.
- **Collect permissions:** Assign Roles and permissions to a group of Users that don't share an organization. For example, in a university's portal, a user group could be created to group all teachers independently of their departments (organization). This would make it easier to assign one or several Roles at once to all the teachers.

Read on to learn how to manage User Groups.

CREATING A USER GROUP




Follow these steps to create a user group:

1. Open the Menu () and select *Control Panel* → *Users* → *Users Groups*. Any existing user groups appear here in a table.
2. Click the *Add* button (). This opens the *New User Group* form.
3. Give your user group a name and description.
4. If you want to create *My Profile* and/or *My Dashboard* pages for the user group's members, select a Site Template to use from the *My Profile* and *My Dashboard* selector menus.
5. Click *Save*. The new user group then appears in the table.

Note that new User Groups don't have any Users. The next section shows you how to add members to a user group.

41.1 Assigning Members to a User Group

Follow these steps to add existing users to a user group:

1. If you're not already there, open the Menu () and select *Control Panel* → *Users* → *Users Groups*. The available user groups appear here in a table.
2. Click *Actions* () → *Assign Members* for the user group you want to add users to. The group's existing users appear in a table.
3. Click the *Add* button (). This opens a list of the users you can select.

Name *

Teachers

Description

The user group for all teachers.

The pages of a user group cannot be accessed directly by end users. Rather, each member of the user group will see the pages automatically as part of the pages of their personal dashboard and profile. To allow users to make changes, enable the customization options of each page.

My Profile

None

My Dashboard

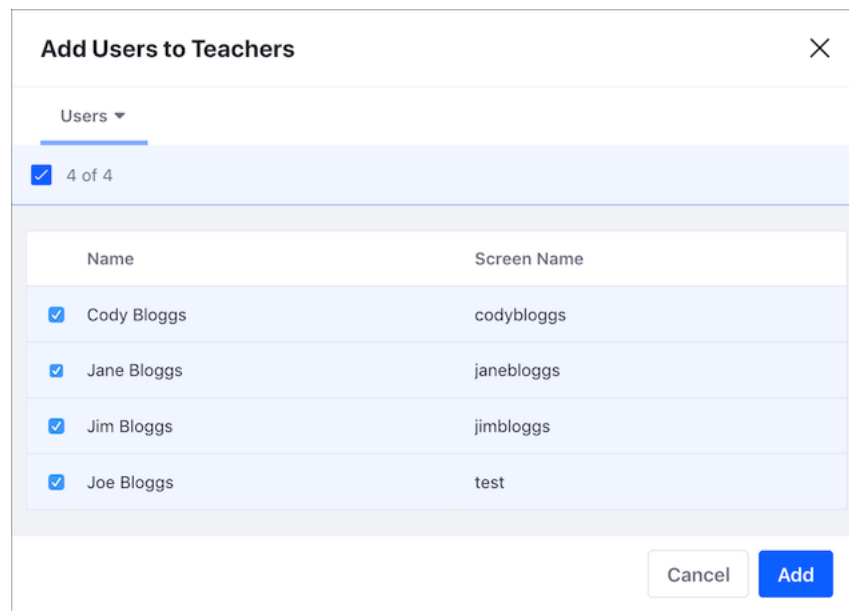
None

Figure 41.1: The New User Group form.

| User Groups | | | |
|-----------------------------------|----------------------------------|--------|---|
| Name | Description | Users | |
| <input type="checkbox"/> Teachers | The user group for all teachers. | 0 User | ⋮ |

Figure 41.2: The user group you just created now appears in the table.

4. Select one or more users from the list, then click *Add*. This adds the selected users to the group, and returns you to the table containing the group's users. The users you added now appear in the table.



Add Users to Teachers ✕

Users ▾

4 of 4



| Name | Screen Name |
|---|-------------|
| <input checked="" type="checkbox"/> Cody Bloggs | codybloggs |
| <input checked="" type="checkbox"/> Jane Bloggs | janebloggs |
| <input checked="" type="checkbox"/> Jim Bloggs | jimbloggs |
| <input checked="" type="checkbox"/> Joe Bloggs | test |

Figure 41.3: Select the users to add to the user group.

USER GROUPS AND SITE MEMBERSHIP

User Groups are used to manage Site membership. When you assign a User Group to a Site, the group's Users become members of that Site. This is one of the main use cases for User Groups.

Follow these steps to assign a user group to a Site:

1. Open the Menu () , select the Site you want to work in, then open its Site Administration menu.
2. In the Site Administration menu, select *People → Memberships*. This opens the Site Memberships screen.
3. In Memberships, select the *User Groups* tab. This tab displays any User Groups currently assigned to the Site.
4. Click the *Add* button () , select any User Groups you want to assign to the Site, then click *Done*. The user groups you selected now appear in the User Groups tab.

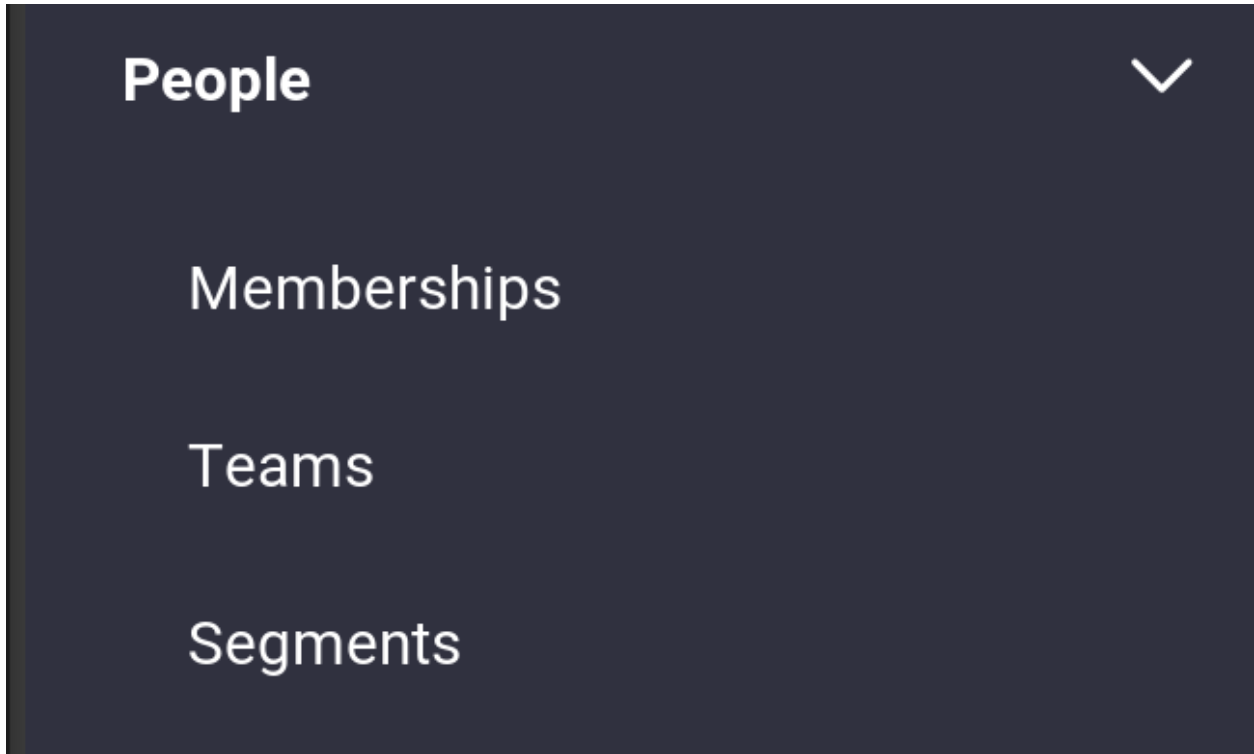


Figure 42.1: Select *Memberships* from the People menu.

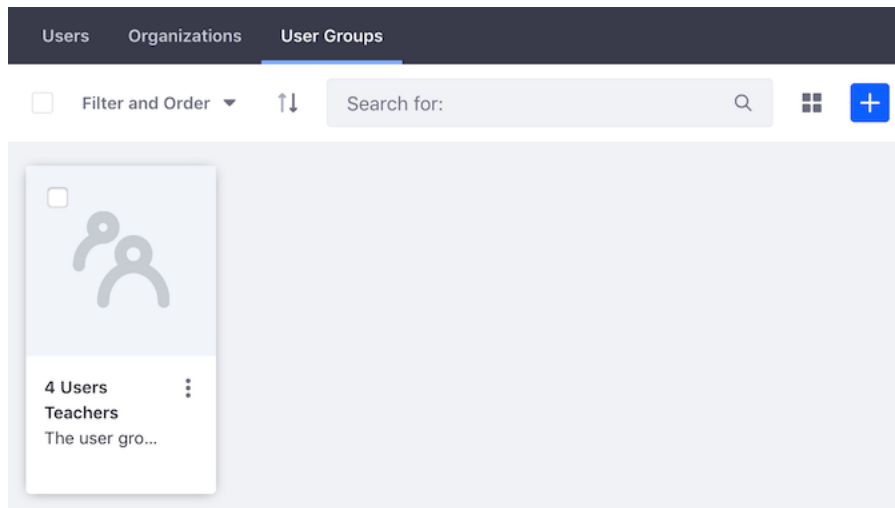


Figure 42.2: The User Groups tab in Memberships shows the User Groups currently assigned to the Site.



USER GROUP SITES

Each User has a personal Site that consists of public (Profile) and private (Dashboard) pages. A *User Group Site* determines the base pages of the User Group members' personal Sites. If the User Group Site pages are added to a User's Profile pages, then the User Group Site is a public Site, accessible to anyone with the URL (`http://www.[sitename].com/web/[username]`). If the User Group Site pages are added to the user's Dashboard pages, then the User Group Site is a private Site. A mixed approach can also be used, where both private and public pages are added for the User Group Site. If Users belong to multiple User Groups, all the pages from those User Group Sites are made part of their personal Sites.

When creating a user group, you can create the user group Site via the Site Templates available for selection in the *My Profile* and *My Dashboard* selector menus. You can also create a User Group Site later, either manually or via a Site Template.

43.1 Creating User Group Sites From Site Templates

Follow these steps to create a User Group Site from a Site Template, for a User Group that already exists:

1. Open the Menu () and select *Control Panel* → *Users* → *User Groups*. User groups appear in a table.
2. Click *Actions* () → *Edit* for the User Group you want to create a Site for. This opens a form that you can use to edit the User Group. Note that this is the same form that appears when creating a user group.
3. To use a Site Template to create a public profile for the Users on their *My Profile* Site, select that Site Template from the *My Profile* menu. To use a Site Template to create private pages for the Users on their *My Dashboard* Site, select that Site Template from the *My Dashboard* menu. Note that you can also do both.

4. Click *Save*.

Now, when one of the group's Users navigates to their *My Profile* or *My Dashboard* Sites, the content of those Sites reflect the Site Template(s) you selected.



User Group Site pages function similarly to regular Site Template pages, with an important exception: User Group Site pages aren't copied for each user. They're shown dynamically along with any custom pages that Users may have on their personal Site. For this reason, Users can't modify pages inherited from the User Group. If needed, the User Group administrator can define certain areas of a page as customizable, like with regular Sites. This lets Users add and configure widgets in the specified area of the page.

This flexibility lets you achieve almost any desired configuration for a User's personal Site without having to modify it directly. When Users are assigned to a User Group, they'll immediately have access to the User Group's Site pages from their personal Sites.

Note: Site Templates have an option that propagates changes made to the Site Template. If you use a Site Template with this option enabled, the User Group Sites update automatically when that template changes. If you disable this option but enable it again later, the template's pages are copied to the Users' Sites, overwriting any changes they may have made. For more information on the automatically propagating Site Template changes, see Site Templates.

43.2 Creating User Group Sites Manually

You can create a User Group's Site manually, instead of basing it on a Site Template. Follow these steps:

1. Open the Menu () and select *Control Panel* → *Users* → *Users Groups*. User groups appear here.
2. Click *Actions* () → *Manage Pages* for the user group you want to create a Site for. This opens the *Pages* window. Note that this is the same window you use for creating pages.
3. Create the public and/or private pages that you want to use for the Users' *My Profile* and/or *My Dashboard* Sites. Public pages you create here become pages on users' *My Profile* Site, while private pages become pages on users' *My Dashboard* Site.

When you return to User Groups in the Control Panel, you can access a User Group's public and/or private pages via these links in the User Group's *Actions* button ():

- **Go to Profile Pages:** Opens the User Group's public *My Profile* page(s) in a new browser window.
- **Go to Dashboard Pages:** Opens the User Group's private *My Dashboard* page(s) in a new browser window.

In the new window, you can add more pages and portlets and configure Site settings.

43.3 Legacy User Group Sites Behavior

Since the inheritance of User Group Site pages is now dynamic, even if there are hundreds of thousands of Users, even millions, there isn't an impact in performance. Versions of Liferay Portal and Liferay DXP prior to 7.0 required User Group pages be copied to each User's personal Site. If you long for the old days, or if you're upgrading from an older version and must keep that behavior, enable it by adding the following line to your `portal-ext.properties` file:

```
user.groups.copy.layouts.to.user.personal.site=true
```

When this property is set to true, the template pages are copied to a User's personal Site once, and then may be modified by the User. This means that if changes are made to the template pages later, they only affect Users added to the User Group after the change is made. Users with administrative privileges over their personal Sites can modify the pages and their content if the *Allow Site Administrators to Modify the Pages Associated with This Site Template* box has been checked for the template. When a User is removed from a User Group, the associated pages are removed from the User's personal Site. If a User is removed from a group and is subsequently added back, the group's template pages are copied to the User's Site a second time. Note that if a User Group's Site is based on a Site Template and an administrator modifies the User Group's Site Template after users have already been added to the group, those changes only take effect if the *Enable propagation of changes from the Site Template* box for the User Group was checked.



CONFIGURING USER GROUP PERMISSIONS

Administrators commonly create User Groups so the group's Users can take some specific action in a Site. This is done by assigning the permissions for that action to a Role, and then assigning that Role to the User Group. This grants the Role's permissions to the User Group's Users.

Follow these steps to grant permissions to Users in a User Group:

1. Create the User Group.
2. Assign the User Group to a Site.
3. Create the Site Role and define its permissions.
4. Assign the Role to the User Group.

For instructions on the first three steps, click those links above. This article shows you how to assign a Role to a User Group:

1. Open the Menu () , select the Site to work in, then open its Site Administration menu.
2. In the Site Administration menu, select *People* → *Memberships*. This opens the Memberships screen.
3. In Memberships, select the *User Groups* tab. This tab displays User Groups currently assigned to the Site.
4. Click the *Actions* button () for the User Group you want to assign to a Role, and select *Assign Site Roles*. This opens the Assign Site Roles dialog.
5. In the Assign Site Roles dialog, select the Role from the list and click *Done*.

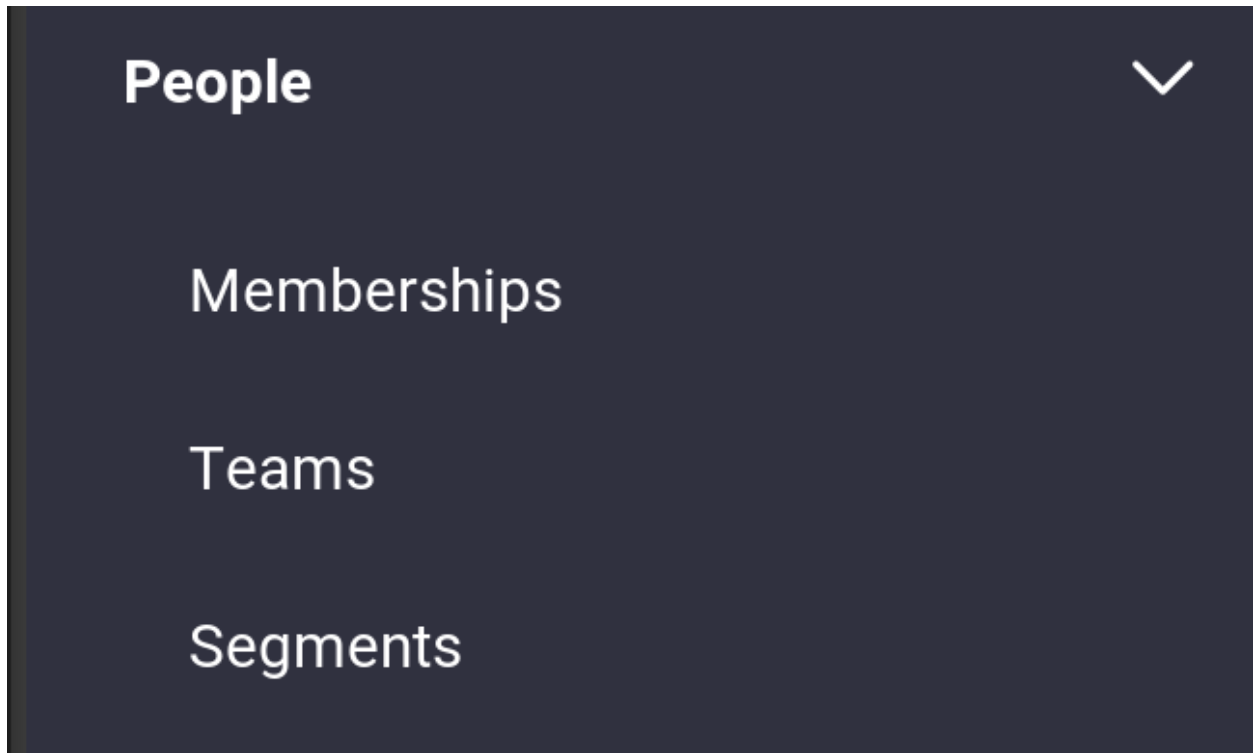


Figure 44.1: Select *Memberships* from the Site Administration menu.

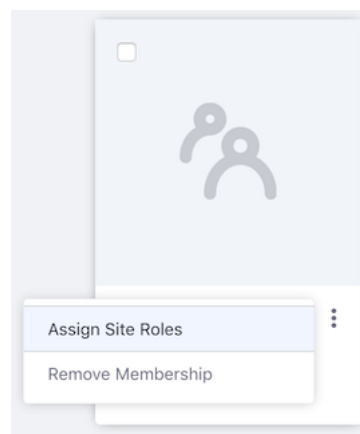


Figure 44.2: Select *Assign Site Roles* for the user group.




EDITING USER GROUPS


You can access User Groups from *Control Panel* → *Users* → *User Groups*. Managing membership is the most common action you'll likely perform on a User Group.

| Name | Description | Users |
|-----------------------------------|----------------------------------|--------|
| <input type="checkbox"/> Teachers | The user group for all teachers. | 0 User |

Figure 45.1: The user groups appear in a table.

Follow these steps to add/remove users to/from a User Group:

1. Click the User Group's name or description. Alternatively, you can click the User Group's *Actions* button () and select *Assign Members*. This presents a list of the User Group's users.
2. To remove a User from the User Group, click the X button next to that User. To remove multiple Users at once, check each User's checkbox and then click the trash icon () that appears in the Management Bar above the User list.
3. To add Users to the User Group, click the *Add* button (). In the dialog that appears, select the Users and click *Add*.

Other options are available in each User Group's *Actions* button ():

Edit: Modify the User Group's name or description, or choose Site templates to use for the User Group's Sites.

Permissions: Assign permissions for viewing and managing the User Group.

User Group Pages Permissions: Assign permissions for managing the User Group's Site pages.

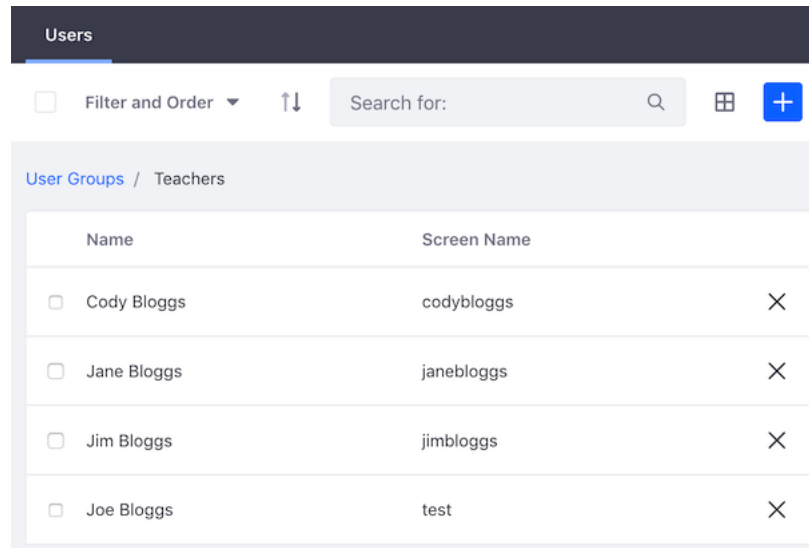


Figure 45.2: The list of Users lets you manage the User Group’s membership.

Manage Pages: Manually manage the User Group’s Site pages. See the documentation on user group Sites for details.

Assign Members: Add/remove Users to/from the User Group. This is described in detail above.

Delete: Remove the User Group. Note that you can’t delete a User Group that contains Users. You must first remove the Users from the group.

If your User Group has public and private Site pages, the options *Go to Profile Pages* and *Go to Dashboard Pages* also appear in your User Group’s Actions menu. Clicking one of these links opens that Site in a new browser window. See the documentation on user group Sites for details.

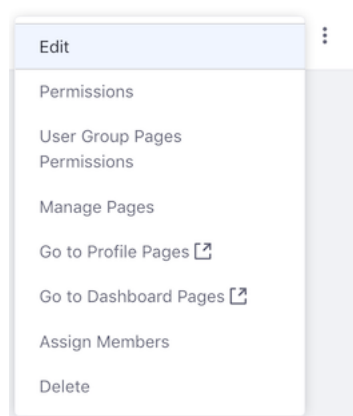


Figure 45.3: The Actions menu for a user group.

PASSWORD POLICIES


Password policies enforce password rules to help users specify secure passwords. Use the default policy that ships with Liferay (modified or as is), or create your own policies. Password policies can be assigned to Users or Organizations, or set as the default policy used throughout a virtual instance.


46.1 Adding and Configuring Password Policies

To add or edit password policies,

1. Navigate to *Control Panel* → *Users* → *Password Policies*.

There's already a default password policy in the system.

2. Edit this the same way you edit other resources: click *Actions* () and then click *Edit*.

3. To add a new policy, click the *Add* () button.

Provide the *Name* (required), *Description*, and specific configuration options for your new password policy.

There are several configuration categories.

Password Changes Allow or disallow users to change their passwords, and set a time limit on the validity of password reset links.

Password Syntax Checking If enabled, require users to use a certain syntax when choosing a password. You can disallow dictionary words, set a minimum length, and more in this section.

Password History If enabled, decide how many passwords to keep in the history, preventing users from reusing an old password.

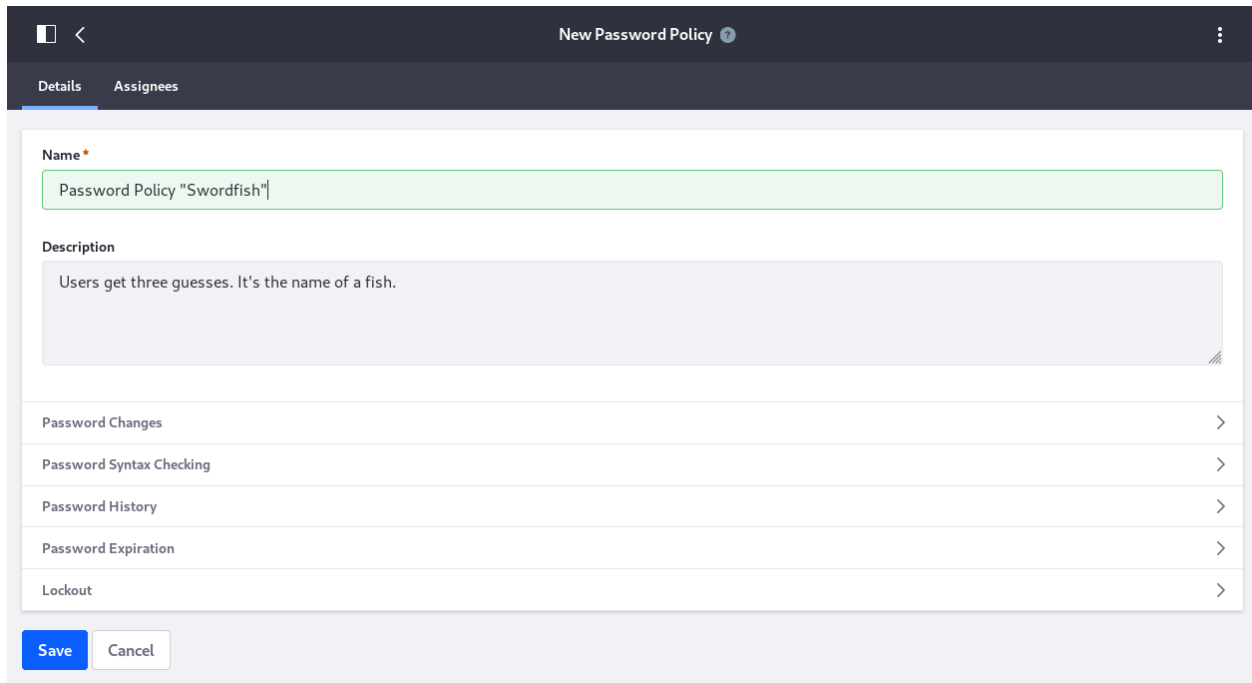


Figure 46.1: You can create new password policies to suit your needs.

Password Expiration Decide whether to expire passwords after a specified time. If enabled, specify how long passwords are valid, when and whether to send a warning, and how many times the user can log in after the password is expired before needing to set a new password (called a *Grace Limit*).


Lockout If enabled, set a maximum number of failed authentication attempts before the account is locked, how long the number of attempts is stored, and the lockout duration.

Self Destruct If enabled, set the time after lockout before Liferay self destructs catastrophically, sending the world into apocalyptic chaos, out of which self-aware robots arise and recolonize the world, enslaving the surviving remnant of humanity for their own nefarious purposes. We recommend you keep this disabled.

Just making sure you were paying attention; that last one doesn't actually exist.

Once you configure the policy, click *Save* to add it to the list of ready-to-use password policies.

46.2 Assigning Users to a Password Policy

To use the default password policy, you don't have to do anything: like its name suggests, it's the default. If you create a new password policy, however, you must assign users to it. To do this click *Actions* () → *Assign Members*.




Figure 46.2: Assign members to new password policies to make them take effect.

Choose whether to assign users directly or to assign organizations to the password policy, then click *Add* (+).

Once assignments are saved, the password policy is in effect. Did you know you can change the default password policy and configure it using Liferay's `portal.properties` file?

46.3 Default Policy Properties

The Default Password Policy is set as the default and configured in Liferay's `portal.properties` file. Find the properties that start with `passwords.default.policy`. To make changes, including changing the default policy, add whichever properties and values you choose to modify in your `portal-ext.properties` file, as usual. Restart the application server and your changes take effect.

```
#
# Set the properties of the default password policy.
#
...
passwords.default.policy.name=Default Password Policy
...
```

As you can see, Password Policies give you a simple yet powerful way to set password rules.

AUDITING USERS

You've just finished lunch and are ready to get back to work. You have a Site you use to manage your project and before you left, you were about to create a folder in your Documents and Media library for sharing some requirements documentation. Sitting down at your desk, you navigate to the repository and attempt to create the folder.

You do not have permission to perform this action, Liferay DXP helpfully tells you.

"What?" you blurt in surprise. "This is *my* project!"

"Ah, you too?" asks a co-worker from over the cube wall. "I lost access to a wiki I was updating just a few minutes ago. I was about to enter a support ticket for it."

"Forget the ticket. Let's go see the admin now," you say.

And off you go, two floors down and to the far end of the building where, as you approach, you can already hear stress in the admin's voice as he tries to reassure someone on the phone.

"Yes, Mr. Jones. Yes, I'll fix it." (*Jones? The president of the company?* goes through your mind.) "I'll get on it right away, Mr. Jones. It was just a mistake; I'll fix it. Thank you, Mr. Jones," and he hangs up the phone.

"Problems?" you ask the admin, whose name is Harry. He does look rather harried.

"Yeah, Tom," he says. "Somebody changed a bunch of permissions—it wasn't me. I'm assuming you and Dick are here because of the same problem?"

"Yup," you say. "I lost access to a document repository folder."

"And I lost access to a wiki," Dick says.

"It was probably due to some Site membership change. Let's take a look at the Audit app in the Control Panel and see what happened."

Sometimes you need to know what Users are doing and exactly who is doing it. If you're a DXP subscriber, you can find this out with the Audit app. In combination with some settings in `portal-ext.properties`, the Audit app shows you all the activity that occurs on your server. You can quickly find out what changes were made and by whom. If you've delegated permission granting to any group of people, this is an essential feature you're likely to use.

VIEWING AUDIT EVENTS

The Audit app shows activities in your Liferay DXP installation. Access it by navigating to *Control Panel* → *Configuration* → *Audit*. The app displays a searchable list of captured events. You can browse the list, but searching it is typically faster.

This figure shows that John Watson logged in and performed some actions on the site. Click an entry to view details about any of these events.

The screenshot shows the 'Audit' app interface. At the top, there is a search bar with the placeholder text 'Keywords' and a 'Search' button. Below the search bar, there are navigation controls: 'Page 1 of 3', '20 Items per Page', and 'Showing 1 - 20 of 43 results.' To the right of these controls are navigation links: '< First', 'Previous', 'Next', and 'Last >'. The main content is a table with the following columns: User ID, User Name, Resource ID, Resource Name, Resource Action, Client IP, and Create Date. The table contains six rows of data, all for user John Watson.

| User ID | User Name | Resource ID | Resource Name | Resource Action | Client IP | Create Date |
|---------|-------------|-------------|---------------|-----------------|-----------|----------------|
| 44336 | John Watson | 44281 | User | Assign | 127.0.0.1 | 6 Seconds Ago |
| 44336 | John Watson | 44336 | User | Update | 127.0.0.1 | 56 Seconds Ago |
| 44336 | John Watson | 44336 | User | Update | 127.0.0.1 | 2 Minutes Ago |
| 44336 | John Watson | 44336 | User | Update | 127.0.0.1 | 2 Minutes Ago |
| 44336 | John Watson | 44336 | User | Update | 127.0.0.1 | 2 Minutes Ago |
| 44336 | John Watson | 44336 | User | Login | 127.0.0.1 | 2 Minutes Ago |

Figure 48.1: The Audit app displays the events it captures in a searchable list.

As you can see, depending on how many users you have, this list can get populated very quickly. That's why page view events aren't displayed by default. They'll clutter up your audit report, since they'll definitely be the most frequent event.

Note: You can add page view events to your audit report, but keep in mind that doing so adds LOTS of events. If you're a glutton for this kind of punishment, add this property to your `portal-ext.properties` file:

```
audit.message.com.liferay.portal.kernel.model.Layout.VIEW=true
```

Liferay DXP's code refers to pages as *layouts*. Setting this property to true therefore records audit events for page views. It's turned off by default because this is too fine-grained for most installations.

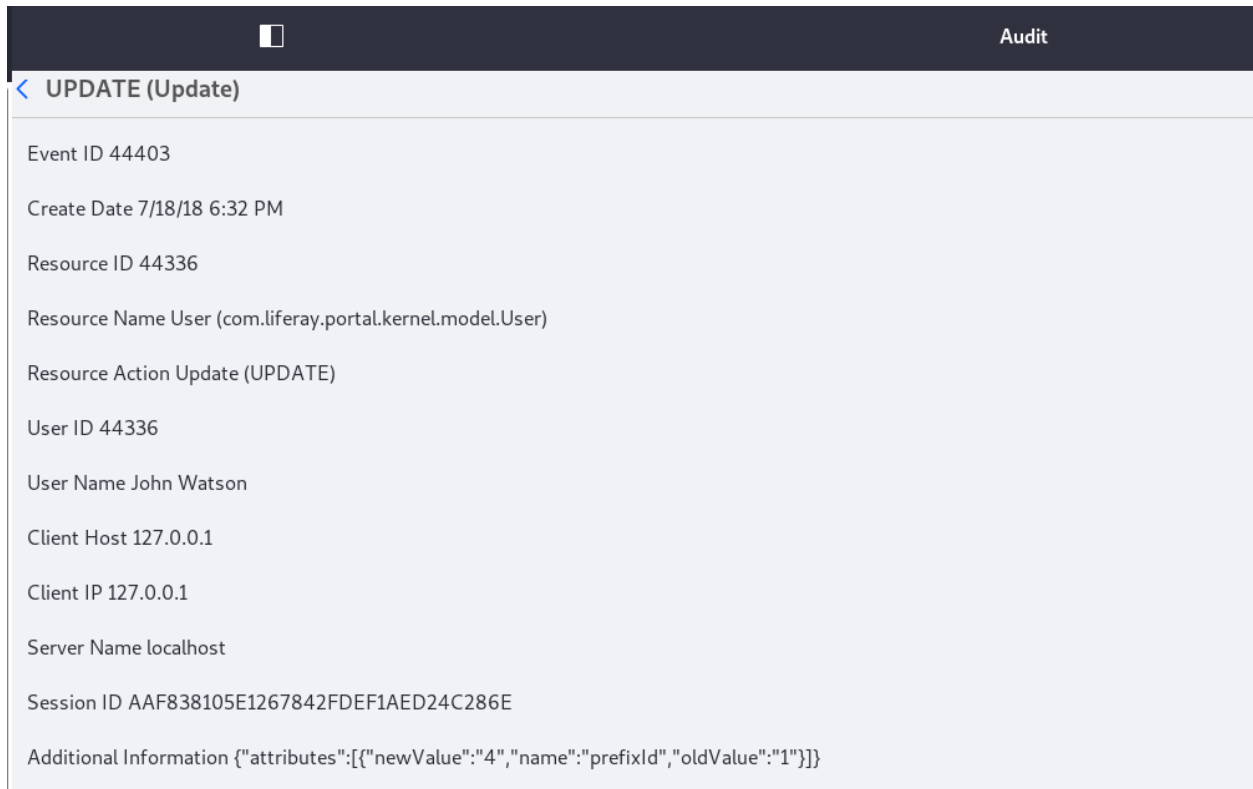


Figure 48.2: Click an event in the list to show its details. The details for this event show that John Watson updated his user account's `prefixId` from 1 to 4. The `prefixId` represents a name prefix like Dr., Mr., Mrs., or Ms.

Once you've added the property, restart your server.

48.1 Finding Audit Events

Finding what you want in a big list of events is like searching for a needle in a haystack. This is why the Audit app has a robust search mechanism. By default, there's only a single search field. Clicking the *magnifier* icon next to the search bar, however, reveals an advanced search dialog broken out by various fields you can use in your search.

Here are the available search options:

Match: Search for matches to *all* the fields you've specified or *any* single field.

User ID: The user ID to search for. This is usually the User who performed some action you'd like to audit.

User Name: The user name to search for. This is often easier than searching for a user ID, especially if you don't have access to the database containing the user ID.

Resource ID: The ID of the resource that was modified or viewed in this audit record.

Class Name: The name of the resource that was modified or viewed in this audit record. For example, you could search for user resources to see if someone modified a user's account.

Resource Action: An action performed on the resource. This could be any of these: add, assign, delete, impersonate, login, login_failure, logout, unassign, or update.

Session ID: The session ID to search for. You can use this to correlate a session ID from your web server logs with activity in Liferay DXP.

Client IP: The IP address of the client that performed the activity you wish to audit.

Client Host: The host name of the client that performed the activity you wish to audit.

Server Name: The name of the server in which the activity occurred. If you're using a cluster, each member of the cluster can be individually queried.

Server Port: The server port in which the activity occurred. You need this if you run a vertical cluster of multiple VMs on the same machine.

Start Date: The low end of the date range you wish to search for.

End Date: The high end of the date range you wish to search.

For example, to check if someone unassigned a User from a particular Role, you might search for a resource name of *user* and a resource action of *unassign*. Once you have the search results, you can click any of the returned records to see that record's detail page.

As you can see, the Audit app shows you what's happening as Users make changes. Use this information to troubleshoot problems, determine ownership of particular actions, or, as Harry (from the story in the introduction) is about to do, find out who made permission changes they weren't supposed to make.

The image shows a mobile application interface for searching audit events. At the top right, there is a 'Search' button. Below it, a close button 'x' is visible. The main section is titled 'Match' and contains a dropdown menu currently set to 'All'. Below the dropdown, the text 'of the following fields:' is displayed. There are five search criteria, each with a corresponding text input field: 'User ID', 'User Name', 'Resource ID', 'Class Name' (with 'User' entered), and 'Resource Action' (with 'Unassign' entered).

Search

x

Match

All

of the following fields:

User ID

User Name

Resource ID

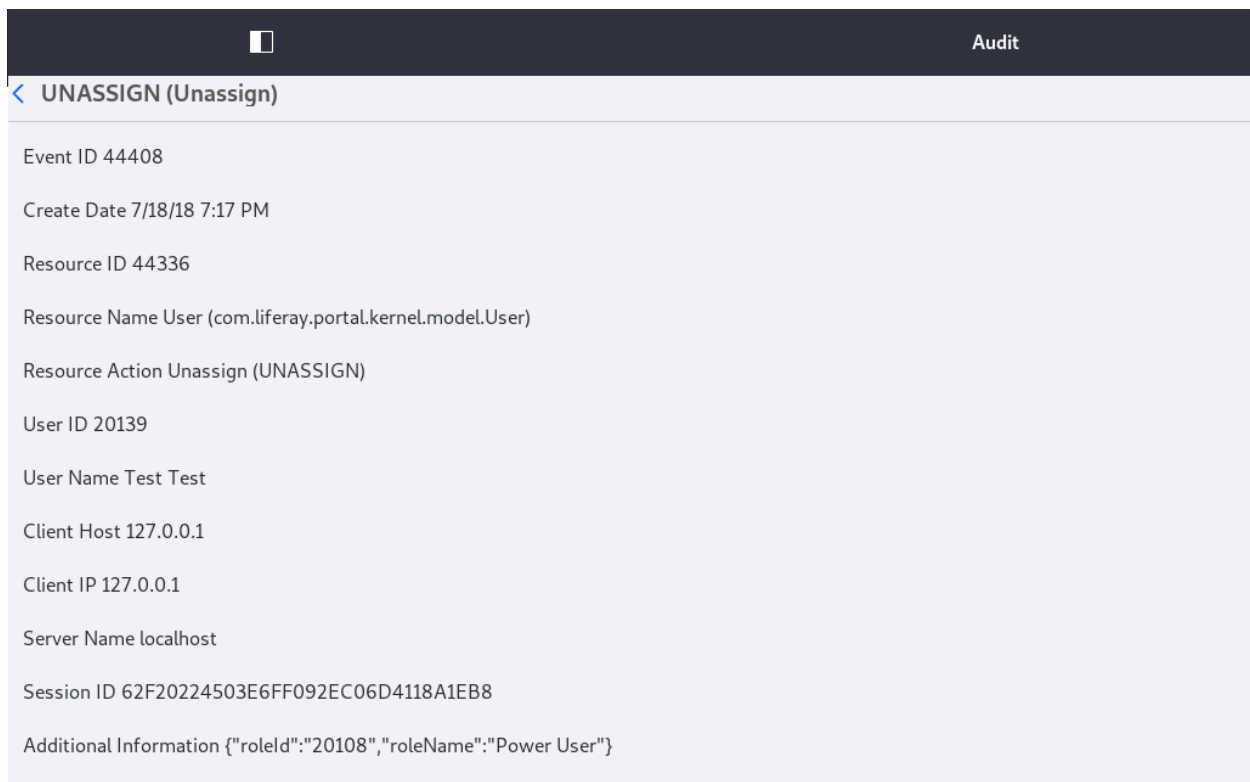
Class Name

User

Resource Action

Unassign

Figure 48.3: Searching for audit events is easy with the Audit app's advanced search form. You can specify various search criteria to find the types of events you want.



The screenshot displays an audit event record in a dark-themed interface. At the top, there is a header bar with a square icon on the left and the word "Audit" on the right. Below the header, the event title "UNASSIGN (Unassign)" is shown with a back arrow icon. The main content area lists the following details:

- Event ID 44408
- Create Date 7/18/18 7:17 PM
- Resource ID 44336
- Resource Name User (com.liferay.portal.kernel.model.User)
- Resource Action Unassign (UNASSIGN)
- User ID 20139
- User Name Test Test
- Client Host 127.0.0.1
- Client IP 127.0.0.1
- Server Name localhost
- Session ID 62F20224503E6FF092EC06D4118A1EB8
- Additional Information {"roleId":"20108","roleName":"Power User"}

Figure 48.4: This record shows that the default administrative user removed the Power User Role from the User Test Test.

CONFIGURING AUDITS

Audits are enabled by default. The Audit app reports audit events, but you can also report them in Liferay DXP's logs or console, enable them for scheduled jobs, or disable them entirely.

There are two main ways to configure Liferay DXP:

1. Edit a configuration via the Control Panel. This saves the configuration to the database.
2. Edit a configuration via an OSGi configuration file (.config file) in your [Liferay Home]/osgi/configs folder.

These methods apply to each of the audit configuration options explained below.

49.1 Reporting Audit Events in Liferay's Logs and Console

Follow these steps to use the Control Panel to configure the reporting of log events in Liferay DXP's log and console:

1. Go to *Control Panel* → *Configuration* → *System Settings* and select *Audit* from the *Security* section.
2. In the *SYSTEM SCOPE* column on the left, select *Logging Message Audit Message Processor*.
3. Select the *Enabled* checkbox to report audit events in Liferay DXP's log.
4. Select the *Output to Console* checkbox to report audit events in the console.
5. In the *Log Message Format* selector menu, select the format for the audit events (CSV or JSON).
6. Click *Save* when you're finished.

Alternatively, you can make the same configuration via an OSGi configuration file:

1. Create a file called `com.liferay.portal.security.audit.router.configuration.LoggingAuditMessageProcessorCon`

2. Add these properties to the file:

```
enabled="true"
logMessageFormat="CSV"
#logMessageFormat="JSON"
outputToConsole="true"
```

Note that these are the same options set in the Control Panel. Edit them as you see fit.

3. Deploy the file to the [Liferay Home]/osgi/configs folder. Note that the Liferay Home folder is typically the application server's parent folder.

Regardless of your configuration approach, you must also extend Liferay DXP's `log4j-ext.xml` file to configure Log4j (Liferay DXP's logging implementation) to log messages produced by the appropriate class to the appropriate file. To do so, create a `portal-log4j-ext.xml` file in [Liferay Home]/tomcat-[version]/webapps/ROOT/WEB-INF/classes/META-INF with this configuration:

```
<?xml version="1.0"?>
<!DOCTYPE log4j:configuration SYSTEM "log4j.dtd">

<log4j:configuration xmlns:log4j="http://jakarta.apache.org/log4j/">

  <!-- additional audit logging -->

  <appender name="auditFile" class="org.apache.log4j.rolling.RollingFileAppender">
    <rollingPolicy class="org.apache.log4j.rolling.TimeBasedRollingPolicy">
      <param name="FileNamePattern" value="@liferay.home@/logs/audit.%d{yyyy-MM-dd}.log" />
    </rollingPolicy>
    <layout class="org.apache.log4j.EnhancedPatternLayout">
      <param name="ConversionPattern" value="%d{ABSOLUTE} %-5p [%t][%c{1}:%L] %m%n" />
    </layout>
  </appender>

  <category name="com.liferay.portal.security.audit.router.internal.LoggingAuditMessageProcessor">
    <priority value="INFO" />
    <appender-ref ref="auditFile"/>
  </category>
</log4j:configuration>
```

This configures Log4j to record INFO level messages from the class `com.liferay.portal.security.audit.router.internal.LoggingAuditMessageProcessor` to a file called `audit.yyyy-MM-dd.log` in the [Liferay Home]/logs folder. Adjust the audit file properties or log level to your liking.

49.2 Configuring Audit Events for Scheduled Liferay Jobs

By default, scheduled jobs don't trigger audit events. Follow these steps to enable them via the Control Panel:

1. Go to *Control Panel* → *Configuration* → *System Settings*, and select *Infrastructure* from the *Platform* section.
2. In the *SYSTEM SCOPE* column on the left, select *Scheduler Engine Helper*.

3. Select the checkbox for *Audit scheduler job enabled* and click *Save*.

Alternatively, you can make the same configuration via an OSGi configuration file:

1. Create a file called `com.liferay.portal.scheduler.configuration.SchedulerEngineHelperConfiguration.config`.
2. Add this property to the file:

```
auditSchedulerJobEnabled=true
```

3. Deploy the file to the `[Liferay Home]/osgi/configs` folder. Note that the Liferay Home folder is typically the application server's parent folder.

Auditing scheduled jobs is a smart choice if there's a chance someone with a dubious competence level would try to schedule jobs, as you'll find out below in the conclusion of our story.

49.3 Enabling or Disabling Audit Events Entirely

Audit events are enabled by default. Follow these steps to disable them via the Control Panel:

1. Go to *Control Panel* → *Configuration* → *System Settings* and then click *Audit* in the *Security* section.
2. Uncheck the *Enabled* box. Note that when auditing is enabled, you can adjust the audit message max queue size from its default value.

Alternatively, you can enable or disable auditing via an OSGi configuration file:

1. Create a file called `com.liferay.portal.scheduler.configuration.SchedulerEngineHelperConfiguration.config`.
2. Add these properties to the file. You can adjust their values as desired:

```
enabled="true"  
auditMessageMaxQueueSize="200"
```

3. Deploy the file to the `[Liferay Home]/osgi/configs` folder. Note that the Liferay Home folder is typically the application server's parent folder.

49.4 The End of the Story

“Okay,” says Harry, “let’s fire up the audit system and see if we can figure out what happened.”

You and Dick stand behind Harry’s chair and watch as he enters a query into a form on the Audit app. After clicking *search*, the screen fills up with audit events.

“Wow, that’s a lot of unassign events.” Harry says. “And look who the culprit is,” he adds sarcastically.

“Who’s Melvin Dooitrong?” Dick asks.

“That’s my new intern,” Harry says. “He’s gonna be sorry.” Harry pushes out his chair and walks down the row of cubes to the end, where a kid no more than 20 years old with disheveled hair sits, earbuds in his ears.

“Hey Melvin,” Harry says as Melvin turns around to face him. “Didn’t I ask you to move that set of users from Site membership to Organization membership?”

“Yeah,” Melvin says, “I did that already.”

“How’d you do it?”

“It was going to take a while to do it manually, so I wrote a script and executed it in the scripting host,” Melvin replies, matter-of-factly.

“You did, did you? Well, guess what? Your script removed *everybody* from *all* Sites.”

“*What?*”

“Yeah, and now you’re going to start adding them back, one by one, manually, starting with Mr. Jones....”

Tom and Dick back away slowly from Melvin’s cube as Harry and Melvin continue to have their—let’s call it a discussion. One thing is clear: they’re having a better day than Melvin is.

WEB EXPERIENCE MANAGEMENT

Experience: consider that word for a moment. Not the type of experience you gain with repetition, but the contact or encounter you have with something. Suppose you're buying a new phone. It's easy to look at the phone with the biggest screen or the fastest processor and say that it's the best, but what will your experience be? The "best" phone might not be the fastest or the biggest; it might be the phone with the longest battery life or the most comprehensive suite of integrated apps. Or it might not be any of those things. Experience isn't always something that you can quantify with specs or features.

Liferay takes the experience factor very seriously when it comes to site and content management. The Web Experience Management suite is focused on providing the best experience for users building websites. When it comes to web experience, just like with your phone, everyone is looking for something different. Some smartphone users might love watching videos during their commute with a big beautiful screen, while on the other side, some people might be more excited about a small sleek phone with a great battery life that fits easily in their pocket and simply does all the basic communication they need.

Liferay's Content Management is the big beautiful phone and the sleek utilitarian one all in one. Marketers will find easy to use tools to build content without having to write any code or peek under the hood. Developers will find powerful tools like Structures and Templates that enable them to create dynamic content. And designers will love how Fragments and Content pages provide a way to perfectly realize their designs.

AUTHORING CONTENT: STRUCTURED AND INLINE CONTENT

The primary goal of Content Management isn't to show off the flashiest new features or follow all the latest trends in design, but to provide you with the tools you need to create digital content that communicates your message clearly and effectively. With this in mind, Liferay DXP offers two core approaches to help you accelerate and simplify creating and organizing content: Structured Web Content and Inline Content.

51.1 Structured Web Content

If you've entered content into a CMS before, you may be familiar with the process of filling content into various fields like this:

- Title
- Abstract
- Text Body

This is an example of Structured Content. Structured Content is created within a predefined content structure and then added to pages as needed. The structure defines the fields and then a template defines its styles. The content is then saved, ready to be added to a page later.

The structure defines what kind of content you are creating and provides different types of fields that can be used. A developer could create a format for publishing articles that contains a **Title**, **Header Image**, **Body Text**, and a **Key Quotation**. The template defines how the elements of the structure are rendered. You could style the elements in the structure with the **Title** as large bold text, the **Header Image** as a full page width block above the title, the **Text** as standard text, and the **Key Quotation** as large font italics with a thin border that displays within the main text section. A content writer or marketer could then create any number of articles, all having a uniform style based on the structure.

In Liferay DXP, those articles could be added to pages across the site, or displayed dynamically with tools like the Asset Publisher and Web Content Display Pages.

51.2 Inline Content

Inline Content is content that is created directly within a page. Rather than filling in fields to create content that's added to a page later, you have a completed page design where you edit the text and image content. Content Pages start with a design which is then created with Fragments. Inside the Fragments, a developer can define where text and images can be placed or edited. Marketers and content writers are then free to write or add images within the page and publish it.

With content pages, basic HTML and CSS define the primary design, while JavaScript and Liferay specific tags can add dynamic behavior to the Content Page. After a developer creates the page and a content writer or marketer provides the content, the content exists inline within the page, and is published with it.

51.3 What's Best for Your Use Case?

When you step back to look at the big picture, what you see are two different paradigms for building pages: content pages, where the content is built into the page; and widget pages, where content and other features can be added, removed, and rearranged as desired.

Often it is helpful to have reusable elements or content that can be moved around a page or placed anywhere on a site. For example, you might have a content based banner which you want to be able to drop onto any various pages with different layouts and styles, or you might have content that uses the same template to create similar items for different pages. Structured Content on widget pages is the tool you need to quickly create what you need and manage these cases. Widget pages with structured content are great for some cases, for example:

- Portal pages, where you provide users a gateway into several different services or providing aggregated information.
- Pages where the primary focus is widgets.
- Pages that are based around structured content.

In other cases, you need to create a page as a complete unit. For example, you have a series of marketing driven landing pages that must match a specific design and have associated content intended for use on that page or with that campaign. Content Pages provide the best tool for quickly bringing a design to life and empowering marketing with inline content. Content Pages are useful for pages like:

- Landing pages

- Front pages that provide marketing information or a direct path into the website.
- Pages with multiple variations and small graphical or textual changes across a large number of pages.

Most sites need a little bit of both. Read on to learn more about building sites with Liferay DXP and how Content Pages and Structured Content can help you do that.

BUILDING A SITE

A site is a set of pages where content or applications are published. Sites can be independent or serve as an associated organization's website. You can create as many different sites as you like within the context of a single Liferay instance.

You'll start with a tour of the site management user interface. Then you'll create a custom Lunar Resort Example instance and explore ways to create sites and pages for that Liferay instance. Finally, you'll learn how to change various settings for sites and pages to meet your needs. To begin building a site, continue on to the next section.

SITE MANAGEMENT

You can have many Sites on one Liferay instance, which work together to create one complete website, or you can simply have one Site which contains all of your pages and content—or anything in between. In this section you'll look at the interface for creating and managing Sites, create a Site, and learn how to use Site Templates for more efficient Site creation.

UNDERSTANDING SITE MANAGEMENT

Whether you're building a large corporate web Site or a small Site for facilitating collaboration among team members, supporting different kinds of collaboration and social scenarios is a must. Liferay's Sites provide three membership types:

Open: Users can become members of the Site at any time.

Restricted: Users can request Site membership but Site administrators must approve requests for users to become members.

Private: Users cannot join the Site or request Site membership. Site administrators must manually select users and assign them as Site members.

In addition to these memberships, when a Site is associated with an organization, all the users of that organization are automatically considered members of the Site.

You can view all the available open and restricted Sites by adding the My Sites application to a page and accessing the *Available Sites* tab. You can request access to any of the Sites you're not already a member of by selecting the Site's *Options* button

() and clicking *Join*.

54.1 Site Scope

Members of a Site can be given additional privileges in the Site by using permissions. It is also possible to assign different roles within the Site to different members. This can be done through *Site Roles*, which are defined equally for all Sites or *Teams* which are unique for each Site. These concepts are discussed later.

Liferay DXP separates Site-scoped information from the Control Panel by placing it in the Site menu. From this menu, you can select the specific Site to work on. The Site Administration panel is available for your Site, which includes Build, Content, Categorization, Recycle Bin, Members, Configuration, and Publishing.




Figure 54.1: Your Site's content resides in the Site Administration menu.

54.2 Site Hierarchies

Sites can also be organized hierarchically, just like Organizations. The difference between Sites and Organizations, of course, is that Sites organize pages, content, application data, and users (via Site memberships) whereas organizations only group users. Content sharing is available for Sites within the same hierarchy. For instance, if a parent Site has a document type called *Lunar Presentation* and all its child Sites should have a copy, the parent Site's administrator can enable content sharing to share the document type automatically with its child Sites. Also, content sharing privileges can be set to let every Site administrator share content across Sites they manage. You can share the following content across Sites:

- Web Content Structures
- Web Content Templates

- Document Types
- Vocabularies and Categories
- Widget Templates
- Data Definitions (Dynamic Data Lists)

Please refer to the Sites Admin Portlet section of Liferay's `portal.properties` file for a list of relevant configurable properties. For example, the `Sites.content.sharing.with.children.enabled` property can disable content sharing between Sites and child Sites, disable it by default while allowing Site administrators to enable it per Site, or to enable it by default while allowing administrators to disable it per Site.

The Sites Directory application is a configurable app that shows a hierarchy of Sites and child Sites. It enables users to navigate to any of the displayed Sites. To use this app to display Site hierarchies, add it to a page, open its Configuration window, and under Display Style, select *List Hierarchy*. The My Sites Directory application is similar to the Sites Directory application, except that it lists only the Sites a user belongs to.

Each child Site in the hierarchy has its own administrator, and the Site Administrator role permissions do not flow down to child Sites in the hierarchy. If a Site Administrator creates a child Site, he or she has the same permissions in that child Site. This is not, however, because of inheritance. It is only because creating a Site makes you the Owner of that Site. A Site Administrator or a parent Site has no default role in any child Sites created by other Site Administrators.

If you wanted a user to have administrative access to all Sites in a Site/child Site hierarchy, you must create a role based on the Site Administrator role that has the permission *Manage SubSites*.

The Site Map application helps users navigate a Site. A Site administrator can configure a root page and a display depth. Just as Sites can have hierarchies, so can the pages within a Site. The display depth of the Site Map application determines how many levels of nested pages to display.

SITE MAP

- [Welcome](#)
- [Community](#)
- [Activities](#)
 - [Rover Racing](#)
 - [Lunar Golf](#)

Figure 54.2: The Site Map application lets users navigate among pages of a Site organized hierarchically.

54.3 Site Members

Another useful administrative application is the Site Members application. This enables administrators to survey all the users, organizations, and user groups that reside in the Site. Similarly, Liferay provides the Portal Directory application, which functions the same as the Site Members app, but globally scoped for all Sites in the instance.

54.4 Page Sets

Sites have two categories of pages called page sets. There are two kinds of page sets: public pages and private pages. A Site can have only public pages, only private pages, or both. Private pages can only be accessed by Site members. Public pages can be accessed by anyone, including users who haven't logged in. It's possible to restrict access to pages at the page set level or at the level of individual pages through the permissions system. Public pages and private pages have different URLs and can have different content, applications, themes, and layouts.


Building a corporate intranet is a typical use case for Sites. A corporate intranet could have Sites for all the organizations in the company: Sales, Marketing, Information Technology, Human Resources and so on. But what about the corporate health and fitness center? That's something everybody in the company, regardless of organization, may want to join. This makes it a good candidate for an open and independent Site. Similarly, the home page for a corporate intranet should probably be placed in an open independent Site so any member of the instance can access it.

For other kinds of websites, you may want to use independent Sites to bring users together who share a common interest. If you were building a photo sharing website, you might have independent Sites based on the types of photos people want to share. For example, those who enjoy taking pictures of landscapes could join a Landscapes Site and those who enjoy taking pictures of sunsets could join a Sunsets Site.


There is always one default Site, which is also known as the main Site of the instance. This Site does not have its own name but rather takes the name of the instance. By default the instance name is *Liferay* but this value can be changed through the configuration of the setup wizard. The instance name can also be changed at any time through the Control Panel within *Configuration* ▢ Instance Settings*.

ADDING SITES

Sites can be created through the Control Panel by a Liferay administrator. The Control Panel provides an administrative interface for managing your Liferay instance. There are four main sections of the Liferay Control Panel: Users, Sites, Apps, and Configuration. In this section, you'll learn how to use the Control Panel to manage Sites. For information about the Apps, Users, and Configuration sections of the Control Panel, see the Using the Liferay Marketplace, Managing Users, and System Wide Settings sections, respectively.

Tip: If you're signed in as an administrator, you can access all Sites by navigating to the Site Administration menu from the Control Panel. To manage a single Site, navigate to the Site by going to the Menu and clicking the *Site Selector* button () from the Sites dropdown menu and selecting the appropriate Site name. Once finished, the Site administration options (i.e., Navigation, Content, Members, etc.) for that Site are available.

Now, you'll add a Site for the Lunar Resort.

1. Navigate to the Control Panel and select *Sites* → *Sites*.
2. Click the Add icon () at the top right of the page.
3. Select a *Blank Site*.

Any available Site templates appear for you to select. Site templates provide a preconfigured set of pages, applications, and content that can be used as the basis of a Site's public or private page set. To create a Site from scratch, select *Blank Site*. Otherwise, select the name of the Site template you want to use. If you opt to create a Site from a Site template, you have to choose whether to copy the Site template's pages as your new Site's public or private page set. If other Site templates are created, they will appear in the Add menu as they become available.

4. Name your Site "The Lunar Resort"

After you enter the name, you will be prompted to enter additional information about the Site and configure certain Site settings.

Name: names the Site you wish to create. You also have the option to translate the name for many different languages. This can be done by selecting the language flag under the Name field, and inserting the name in the selected language. Liferay saves the name translation for each language and displays the translated Site name when that specific language is selected for the instance. If a name translation is not provided, the default instance language's name is displayed.

Description: describes the Site's intended function. The description can also be translated to other languages; see the Name description for more information on translating the Site's description.

Active: determines whether a Site is active or inactive. Inactive Sites are inaccessible but can be activated whenever a Site administrator wishes.

Membership Type: can be open, restricted, or private. An open Site appears in the My Sites app and users can join and leave the Site whenever they want. A restricted Site is the same except users must request membership. A Site administrator must then explicitly grant or deny users' requests to join. A private Site does not appear in the My Sites app and users must be added to it manually by a Site administrator.

Allow Manual Membership Management: determines whether to allow or disallow users to be manually added or removed from the Site. By default, manual Site membership management is enabled. This allows administrators to manually assign users to the Site. It also allows users to join open Sites or request membership from restricted Sites using the My Sites app. For organization Sites, manual Site membership management is disabled, by default. This causes organization members to be automatically assigned membership following the organization's membership policy. Also, because manual membership management is disabled for organization Sites, by default, the *Users* section of *Sites* is unavailable. To activate the *Users* functionality for your organization Site, you'll need to check *Allow Manual Membership Management* after creating the organization Site by navigating to its *Site Settings* menu.

Note: It's possible for Site memberships to be handled automatically by a membership policy. The membership policy can check various pieces of information from each user, such as their first names, last names, birthdays, job titles, organizations, and user groups. Using this information, the Site membership policy can automatically assign members to the Site. If your Site will implement a membership policy, your Site administrators can disallow manual membership management for their Site. When the Allow Manual Membership Management option is disabled, the *Members* section of Site Administration (Site Memberships and Site Teams) is hidden, even from administrators.

Parent Site: lets you select a parent Site for the Site that's being created. Sites can be organized hierarchically. Using hierarchical Sites provides a simplified way to manage Site memberships and Site content sharing. For organizations that have attached Sites, the organization hierarchy should match the Site hierarchy. When you select a parent Site, an additional option appears: *Limit membership to members of the parent Site*. If this option is enabled, the Site's membership policy performs a check so that you can only assign members to the current Site if they're already members of the parent Site.

2. Set the *Membership Type* as *Restricted*.
3. Leave the remain defaults and click *Save*.

When creating a blank Site or organization Site, the Site is not immediately viewable. This is because Sites without a page are impossible to view. Therefore, before you can view your Site,

you must first create a page for it. To add a page for your temporarily invisible Site, navigate to the *Navigation* option from Site Administration. Then add a public page. After adding your Site's first page, it renders and your Site is viewable. For more information about adding pages, see the *Creating and Managing Pages* section.

You can also categorize your Site template using tags and categories by selecting the *Categorization* menu from the bottom of the page. To learn more about using tags and categories in Liferay, see the *Organizing Content with Tags and Categories* section. Lastly, at the top of the page is an additional tab named *Social*. This tab manages whether users of your Site can mention other users. You'll learn about mentioning users later in the *Social Collaboration* sections.

When creating a Site from a Site template, you're asked if you want to copy the pages from the template as public pages or as private pages. By default, the Site is linked to the Site template and changes to the Site template propagate to any Site based on it. A checkbox appears for unlinking the Site template if the User has permission to do so.

Once the Site has been created, you should configure its settings to fit your needs. You can learn more about Site Settings in *Configuring Sites*.

ADDING PAGES TO SITES

In the previous section, you learned how to create sites. You may have gathered from that section that sites aren't particularly useful without pages. In fact, sites primarily exist for the sake of organizing pages and content, so now you'll learn about the different types of pages in Liferay, and how to select the best tools based on your use cases. You'll also learn how to manage pages and use various configuration options.

Before diving into page creation, you should understand the two major page types provided in 7.0:

Content Pages: This new page type is flexible, especially for non-technical users. You can build pages using content created from pre-defined fragments, which themselves can contain widgets.

Widget Pages: Liferay DXP's traditional page type is made up of one or more widgets. There are some use cases (particularly if a page's sole purpose is to host an application) to prefer widget pages.

You should always opt for Content Pages, unless there's a concrete reason otherwise. Content Pages offer many key features of Widget Pages plus more. Some key features of Content Pages include personalized Experiences and A/B Testing. Furthermore, Content Pages are easier to use and are more powerful for non-technical users compared to Widget Pages.

Why would someone prefer Widget Pages? Widget Pages were once the only page type available in earlier versions of Liferay DXP, so they're more familiar than Content Pages.

Additionally, there are still a few things that Widget Pages provide that are not possible with Content Pages:

- *Developing an advanced custom layout:* Using Content Pages, authors can create their own page layouts. This prevents developers from creating pre-selectable, custom layouts with FreeMarker like Widget Pages allowed for. Though Content Pages let you create a layout visually (a more user-friendly approach), the programmatic approach of Widget Pages allows for more advanced capabilities.
- *User-Customizable columns:* This was a rarely used feature of Widget Pages that is not provided in Content Pages. If your page requires a user-customizable column, you must use a Widget Page.

- *Using Staging with Page Variations:* Content Pages do not support Staging's Page Variations. This avoids possible confusion with the similar capability to create variations of a page that are used in personalization and A/B Testing.
- *Widget permissions:* You can configure widget permissions on a Widget Page. This is not yet possible for Content Pages; it's planned for a future release.
- *Widget Look & Feel:* On Widget Pages you can access the Look and Feel tool for widgets, offering fine-grained control over its CSS. This is not available for widgets on Content Pages, since the look and feel of your content is defined in the theme or by using Fragments.

Continue on to learn more about creating pages!

CREATING PAGES

After you create a Site, you can add new pages and maintain them. You can do everything you need with pages from Site Administration.

1. If you're not currently on the Site you want to edit, click the *Site Selector* button (🔄) next to your current Site name in the Menu and select your desired Site.
2. Go to *Site Administration* → *Site Builder*.
3. Click on *Pages*

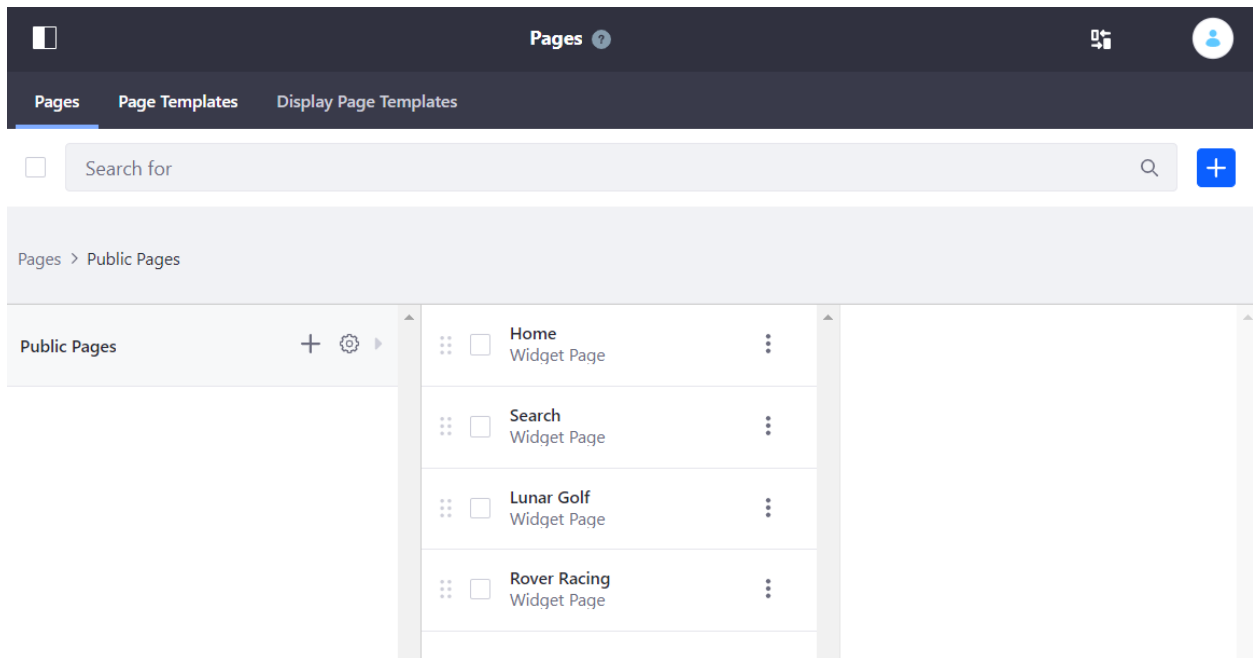


Figure 57.1: The Pages screen lets you edit your Site pages as a whole.

From here, you'll create pages and page templates.

Note: Pages are always part of page sets, and page sets are always associated with Sites. Even users' personal pages are part of their personal Sites. All pages belong to one of two types of page sets: public pages and private pages. By default, anyone can access public pages, even non-logged in users (guests). Only users who are members of the Site that owns the pages can access private pages. This means the private pages of an organization's Site are viewable only by Site members and members of the organization.

From *Pages* you can do several things:

1. Click the (+) button in the top right corner to add a new page.
2. Click options, icons, manage page, or page set settings.
3. Create child pages by clicking the + button next to an existing page.

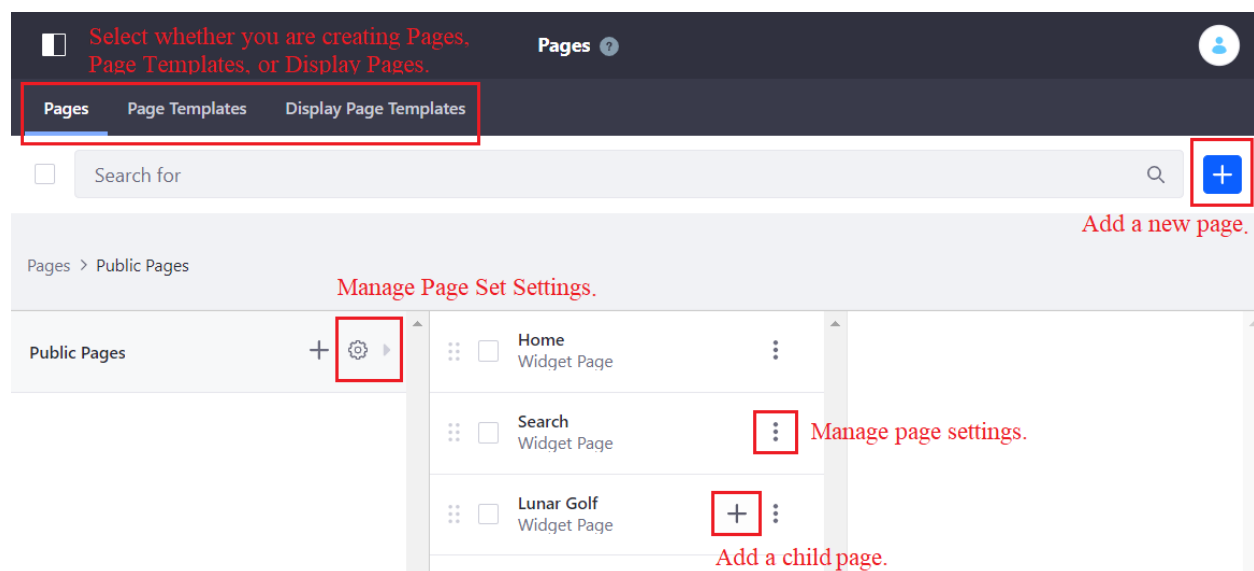


Figure 57.2: Understanding the options on Site Pages.

Adding a child page creates child pages in the hierarchy below the page you've selected. You can nest pages as deep as you like.

Note: You're not forced to define the page hierarchy in a page's friendly URL. Therefore, child page friendly URLs are not required to include their parent page. For example, a parent page named Parent and its child page named Child could have the URLs *SITE_URL/parent* and *SITE_URL/child*, respectively. The default friendly URL given to a page is based only on the page name and not the hierarchy. If you wish to modify a generated friendly URL, you can do so by following the Friendly URL configuration section.

Once you've clicked the + icon to add a page, you're asked to select the type of page you are creating. There are two top options followed by other page types:

Widget: Creates a page with a layout template that defines a number of rows and columns for adding widgets to your page.

Content: Creates a Content Page with inline editing based on Fragments.

Below those you have other options:

Full Page Application: Creates a page that displays a single full page application.

Page Set: Creates a container for subpages that is not actually a page itself.

Link to a Page of this Site: Links to a page within the same Site. This is often used to make a page available in multiple parts of a Sites hierarchy.

Panel: A page containing any number of applications as selected by an administrator, but only one is displayed at a time. Users select the portlet they want to use from a menu on the left side of the page, and the selected portlet takes up the entire page.

Embedded: Displays content from another website inside your instance. An administrator can set a URL from the page management interface and that page appears in the context and within the navigation of your Liferay instance. To embed an external website, you must provide the protocol in the URL (e.g. <https://www.liferay.com/>).

Link to URL: Creates a link to any URL. This could be an external page or a link across Sites in the same Liferay instance.

To the left, under Collections, you can choose to view the basic page types or a collection of page templates. By default, only *Global Templates* appears, but additional collections you create appear here as well.

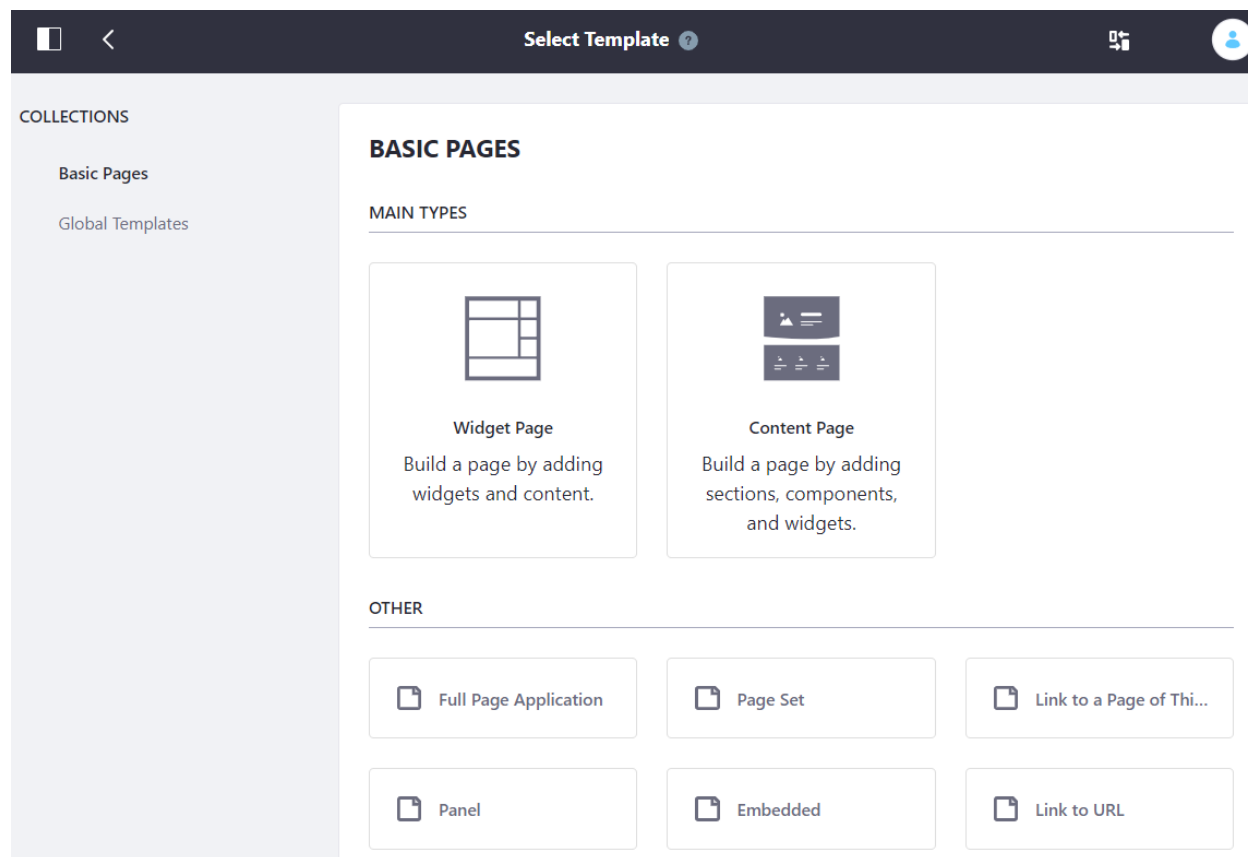


Figure 57.3: You must select a page type when adding pages.

After you've added a page, it may be difficult to track what kind of page you're currently viewing. The page type appears at the top of the page to help you determine the administration options you have and where you need to go to configure the page.




Figure 57.4: Here are three different page types as they're displayed in the heading.

Now that you know the basics of adding pages, you can start working on the Lunar Resort Site. If you're not currently on the right Site, navigate to Site Administration in the Menu, select the compass icon next to the current Site name, and select the Site you wish to edit.

If you must ever modify the page you've created for your Site, select *Configure* from the Options menu for the page from *Pages*. When configuring a specific page, you have more options than when you were creating a new page. You can also read *Configuring Sites*.

There are also configuration options that are only available for individual pages or page groups only. You'll learn about options available for both instances.

Next, you'll look at creating the main page types you'll use in Liferay DXP.

CREATING CONTENT PAGES

Content Pages provide a new paradigm for creating pages in Liferay DXP. They empower marketers and content creators to build pages that can be easily managed and have their content edited in-line and on the fly.

With Content Pages, you build pages from reusable, extensible elements called **Page Fragments**. Page Fragments contain things like images, content, or functionality and are added to a page to be edited and configured to match your design. Web Developers can create Page Fragments to expand the options you have available for building pages.

Some Fragments are completely static elements, like an image or text. These can be useful for building a landing page, or realizing a complicated design vision. Other Fragments have a number of editable fields which enable flexibility and experimentation when creating the perfect page for your needs.

58.1 Creating Page Fragments

The first thing you need to create a page, are the raw materials: Page Fragments. Fragments are organized into Collections which describe some aspect of the purpose of the Fragments and help keep them organized. For example, if you have a number of Fragments associated with a certain marketing campaign, you could create a Collection for that campaign, so that you have all of those Fragments readily available when you're working on a page for that campaign.

With all of the Fragments and Collections that are included in 7.0, many projects can be completed using only the out-of-the-box Fragments and capabilities. In many cases, though, you need to develop your own Fragments. You can learn about the Fragment creation process and editing interface in the Creating Page Fragments article. Developing Fragments is covered in more detail in the Page Fragment development guide.

58.2 Building Content Pages

After you've got your materials, you need to arrange them on a page. Content Pages are built with Page Fragments and Widgets. Widgets work just like they always have, drag them onto the page where you want to use them. Fragments are intended to be reusable, so each Fragment is like a single puzzle piece that can fit in many different puzzles. Using the tools you have available, you can create stunning pages through an intuitive, empowering interface. In the Building Content Pages guide you will create Content Pages using a variety of the features and capabilities of Content Pages.

58.3 The Content Page Interface

There are many features of the Content Page Interface—too many to cover in a practical exercise. For a complete overview of every nook and cranny of the Content Page management interface see the Content Page Interface reference.

58.4 Personalizing Content Pages

When you create Content Pages, you can create different **Experiences** for users based on User Segments. You can create a unique Experience on any Content Page for any existing User Segment. For more information, see the Content Page Personalization guide.

CONTENT PAGE ELEMENTS

Content Pages, like Widget Pages, are built by dragging and dropping elements onto the page and then configuring the way those elements appear. There are three kinds of elements:

Sections are fragments that define a space to place other elements. A section fills the entire width of the page. Sections can be thought of as *complete* Fragments that serve a purpose by themselves. A large banner image with a text overlay is an example of something you might build as a section.

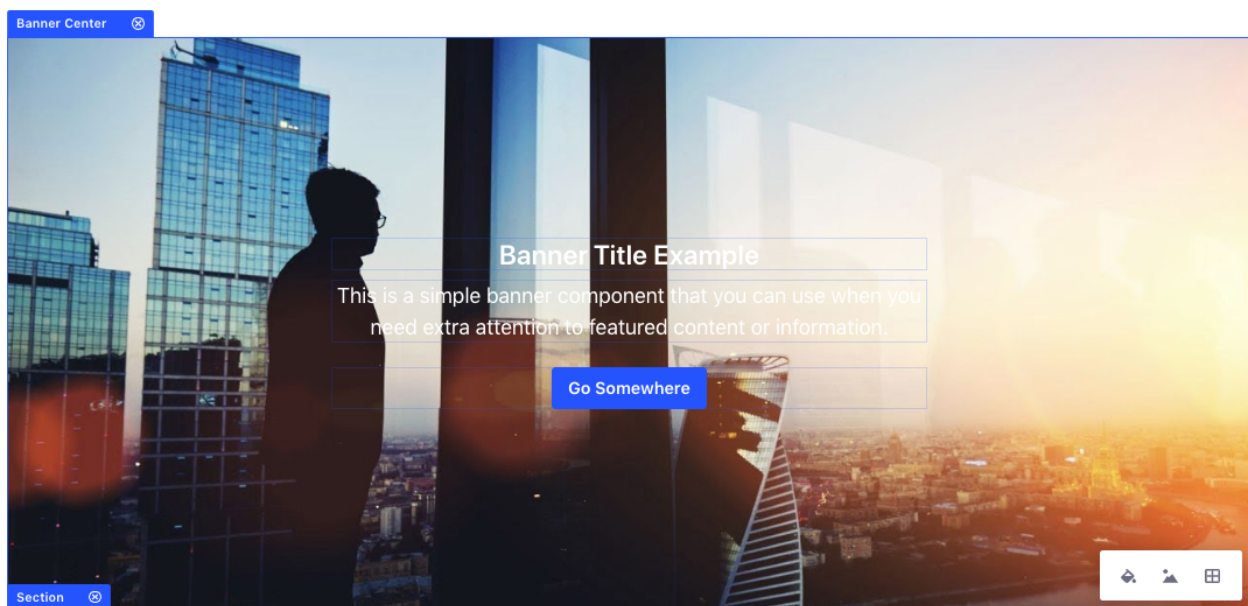


Figure 59.1: A Section named *Banner* being displayed while editing a Content Page.

Layouts are special Sections that define spaces where you can add fragments or widgets. Each layout you add fills the width of the page. You can add any number of layouts to the page.

Components are small design elements or pieces that add functionality to the page. A component might be an image with formatting or a block of text with styling pre-applied. Components must be added to the page inside a Layout. If you add a component outside an existing Layout, a




Figure 59.2: A 3 Column and 1 Column layout stacked on top of each other.

one column layout is automatically added to contain the Component. While Sections should be complete by themselves, Components work together to build pages piece by piece.

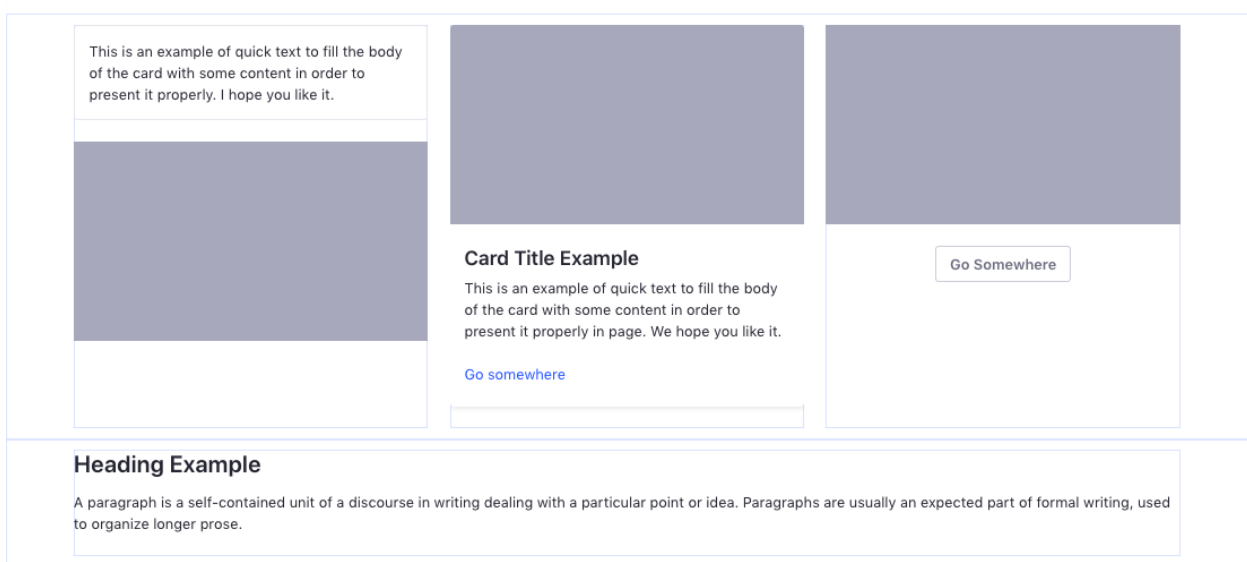


Figure 59.3: Here are several of Liferay's out of the box components arranged in the layout you saw previously.

Liferay DXP ships with a plethora of Layouts, Sections, and Components to use to build pages, and a web developer can create their own Fragments to add to these.

59.1 Editable Elements

Fragments can have editable elements. After a Fragment has been added to a page, you can click on an editable area to provide your own text, image, or links in place of the default defined in the Fragment.

You can also map these elements to content. You can set the *Content* for the element (web content article, document, or blog) and choose its applicable *Field* to display (e.g., title, author name, tags, etc.). You can configure this by selecting the element's *Map* button (⚡).

Note: Many mapping improvements were released in Liferay DXP 7.2 SP1+ and Liferay Portal 7.2 GA2+. For example, mapping editables to text/URL fields of existing content and mapping Fragment background images to image fields of existing content. You can also map custom fields. To ensure you leverage the latest editable element mapping features, upgrade to these versions.

For more information on developing these elements, see *Fragment Specific Tags*. Now you'll learn about each editable type.

59.2 Editable Text

Editable text can be plain or rich text. Plain text has no special styling. Rich text enables text styles, typographical emphasis, alignment, and list formatting.

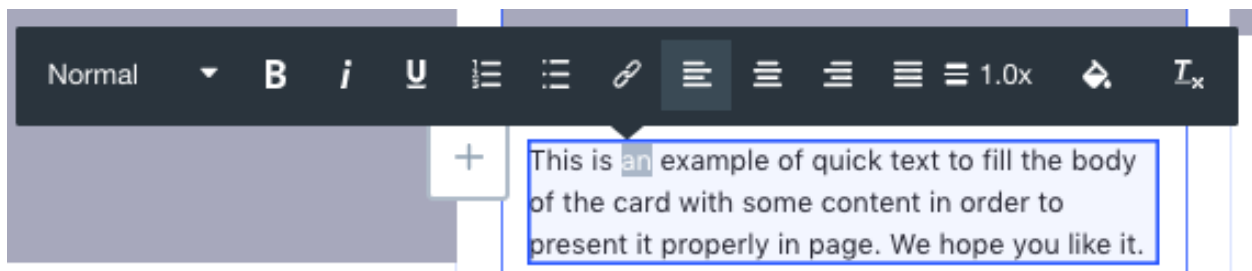


Figure 59.4: The rich text editor provides a simple WYSIWYG interface with a number of formatting options.

59.3 Editable Images

Editable image elements allow replacing the image URL or an image from your Documents and Media library. You can provide a link target for the image.

To edit an image from the Content Page editor,

1. Click on the image you want to replace.

2. Click .

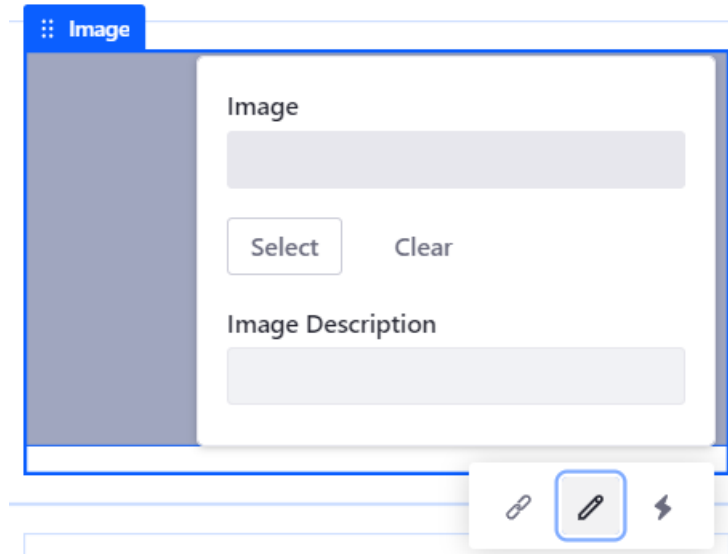


Figure 59.5: Editing an image allows you to enter a URL, select an image from Documents and Media or set a link for the image.

From here, you can click *Select* to upload an image from Docs and Media or define an image URL. Click *Clear* to reset the image. You can also specify an image description.

You can also specify a background image for a layout from Section Builder. Click the Layout, select *Layout Background Image*, and define the image to display.



Note: Mapping a Layout background image is available in Liferay DXP 7.2 SP1+ and Liferay Portal 7.2 GA2+.


For more information on developing editable images, see *Making Images Editable*. You can also define a link for your image. You'll learn about this next.

59.4 Editable Links

Editable links can be associated with entities that redirect you to a content type or Page (e.g., buttons).

To edit a link from the content page editor,

1. Click on the link or button that you want to edit.
2. Click on  to edit the link text.
3. Click on  to edit the link properties.

4. Click on  to edit the link mapping (described earlier).

From the Link Properties popup, you can define the following link options:

Manual: defines a manual link or map it to an existing content field.

- *URL:* sets the link's URL.
- *Target:* set the link's behavior.

From Content Field:

- *Content:* sets the content type.
- *Field:* sets the field to display for the selected content.

Some of the content fields include


- Categories
- Tags
- Display Page URL
- Description
- Publish Date
- Summary
- Title
- Last Editor Name
- Author Name
- Basic Web Content

For more information on developing editable links, see [Creating Editable Links](#). Next you'll learn about the [Content Page Editing Interface](#).

CONTENT PAGE MANAGEMENT INTERFACE

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Unlike Widget Pages, Content Pages can only be edited through the *Site Builder* and cannot be edited live on the page. Any edits that you make to a page are saved as a draft until you publish the page. Subsequent changes after the initial publication are again saved as a draft, without affecting the live page, until the page is published again. To create a Content Page,

1. Go to *Site Management* → *Site Builder* → *Pages*.
2. Click .
3. On the next page, select *Content Page* and provide a name for the page.
You will be brought to the Content Page management interface.




Figure 60.1: Each Content Page starts as a blank page.

To edit an existing Content Page,

1. Go to *Site Administration* → *Site Builder* → *Pages*.
2. Click *Actions* (⋮) → *Edit* next to the Content Page you want to edit.

You can also get to this page by selecting the *Edit* button (✎) from the Control Menu if you're viewing the published Content Page.

On this page you can view a preview of your page, add Fragments and Widgets, and manage the configuration for the page or any Fragments and Widgets currently residing on the page.

Your tools for building the page are all found on the right side of the page. From top to bottom, the options are

- Sections
- Section Builder
- Widgets
- Page Structure
- Look and Feel

60.1 Sections

When you open *Sections* you see a list of Collections available. Initially, you only have the **Basic Sections** Collection which is included with the product. You can open the Collection and drag Sections directly onto the page.

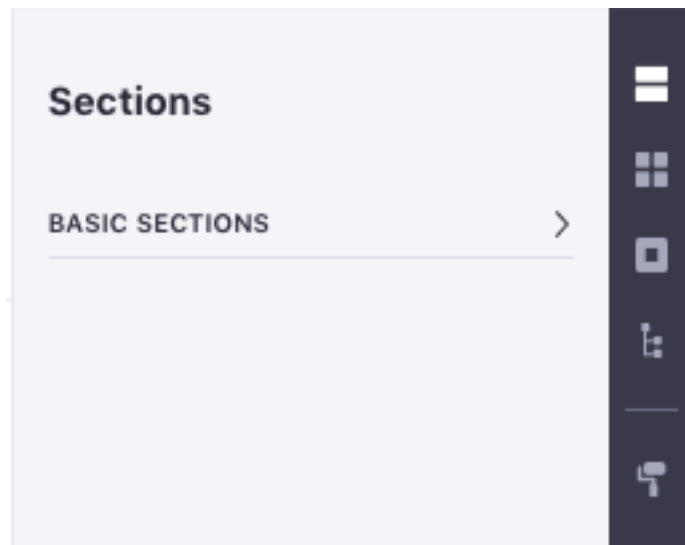


Figure 60.2: *Sections* contains Fragments that fully define spaces on your page.

Once a section is added to a page, you can edit its background color, background image, and spacing. Since these options are available to marketers and administrators editing a page, the options are limited, and the color palette can be set by the Fragment developer.




Figure 60.3: The Section management tool provide powerful tools, but with the training wheels still on.

60.2 Section Builder

In Section Builder, you start with *Layouts* and *Basic Components*. Add Layouts to the page to provide a spaces where you can add Components.

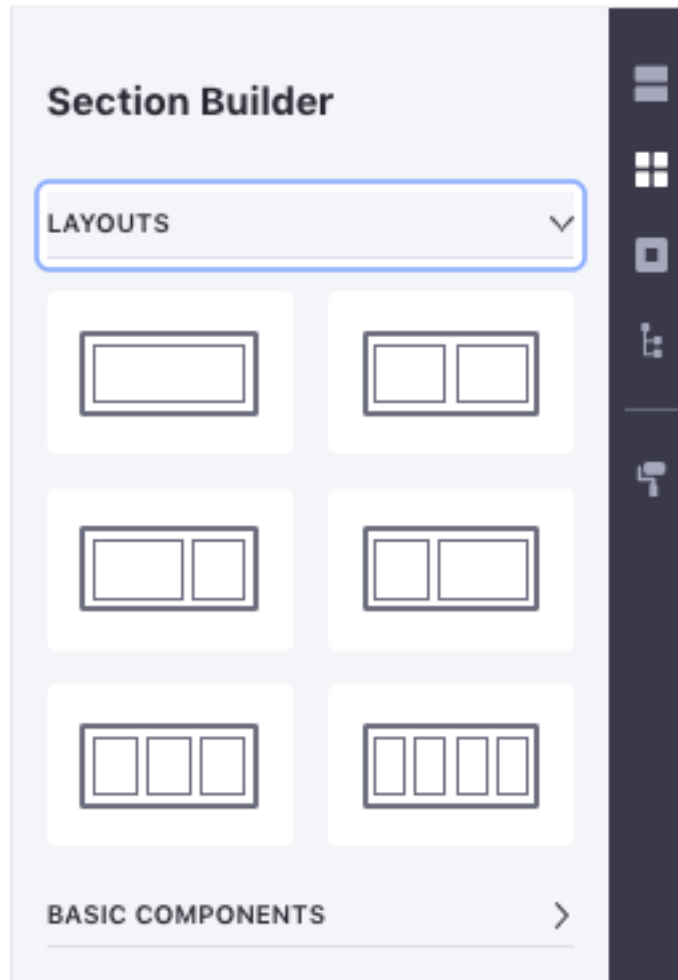


Figure 60.4: *Sections Builder* contains *Component Fragments* which are intended to be combined to create *Sections*.

60.3 Widgets

The Widgets section functions just like the *Add* menu on a Widget Page. The full list of available widgets is displayed, and you can add them to the page.

The main difference is that only the main configuration options are available for widgets on Content Pages. Various other configurations like *Look and Feel* are only available for widgets on Widget Pages.




Figure 60.5: The Widgets section provides a list of Widgets that can be added inside of a Layout.

60.4 Page Structure

Page Structure provides a high level view of every Fragment and every field within each Fragment on the page.

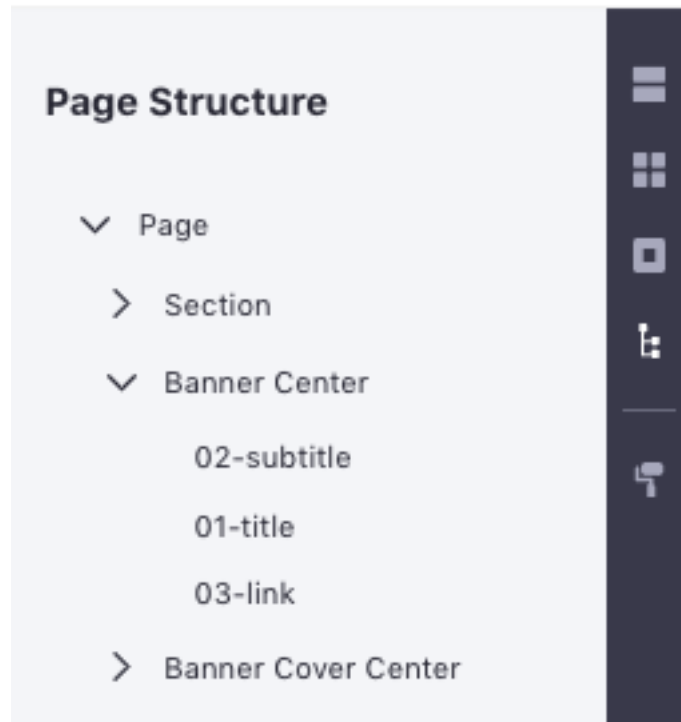



Figure 60.6: *Page Structure* shows you a hierarchy of your page.

Clicking on a field in Page Structure will highlight it on the page. On large complicated page, this helps you keep on top of where everything is and also access items that might be hard to click on directly.

60.5 Look and Feel

Click the *Look and Feel* icon () to change the theme or manage other options for the page. These options are fully documented in *Creating Pages*.

60.6 Comments

You can also comment on any page fragments. This allows discussion and collaboration for teams creating content pages.

Comments are disabled by default, but administrators can enable them from *Control Panel* → *Configuration* → *System Settings* → *Pages* → *Content Page Editor*. Select the *Comments Enabled* checkbox and click *Update*. This enables content page comments for all instances. To control this on an instance-by-instance basis, navigate to the same setting in *Instance Settings* (instead of *System Settings*).

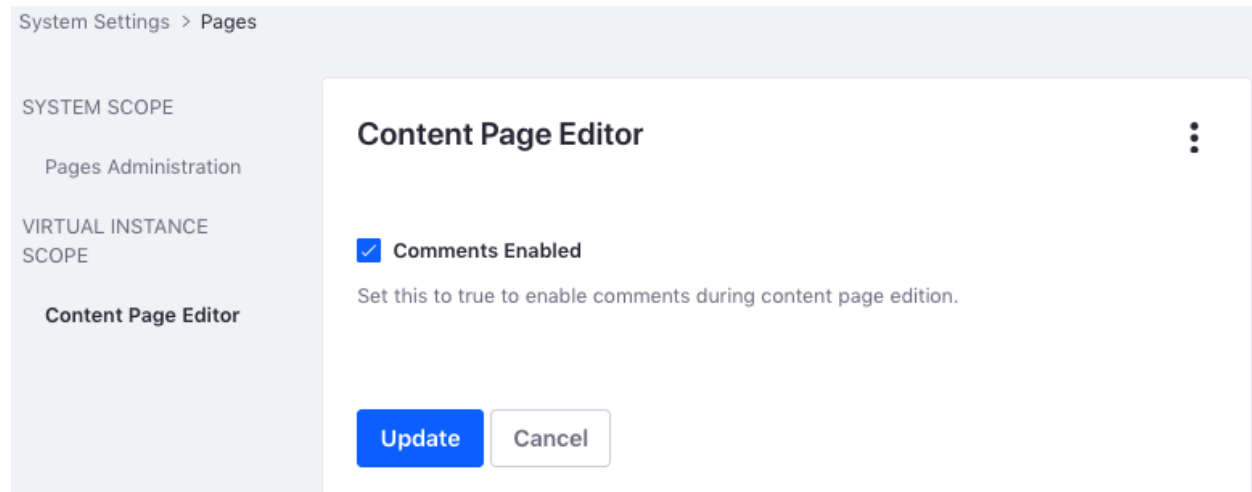


Figure 60.7: Administrators can enable comments for content pages.

If comments are enabled, you can access them via the *Comments* icon (🗨️). The comments appear for the selected fragment. You can take the following actions in the comments UI for a fragment:

- Add new comments and reply to any existing ones.
- Resolve comments by clicking the checkbox for each. Resolving a comment hides it from view, unless *Show Resolved Comments* is selected.
- Edit and delete your own comments via the Actions button (⋮) for each.

If you de-select a fragment or enter the comments UI without a fragment selected, a list of the fragments on the page appears with the number of comments for each. Selecting a fragment then shows its comments.

60.7 The Title Bar

The title bar provides navigation back to the Main Menu, a link to page configuration, and the ability to search for other pages. The title bar is covered in more detail in *Creating Pages*.

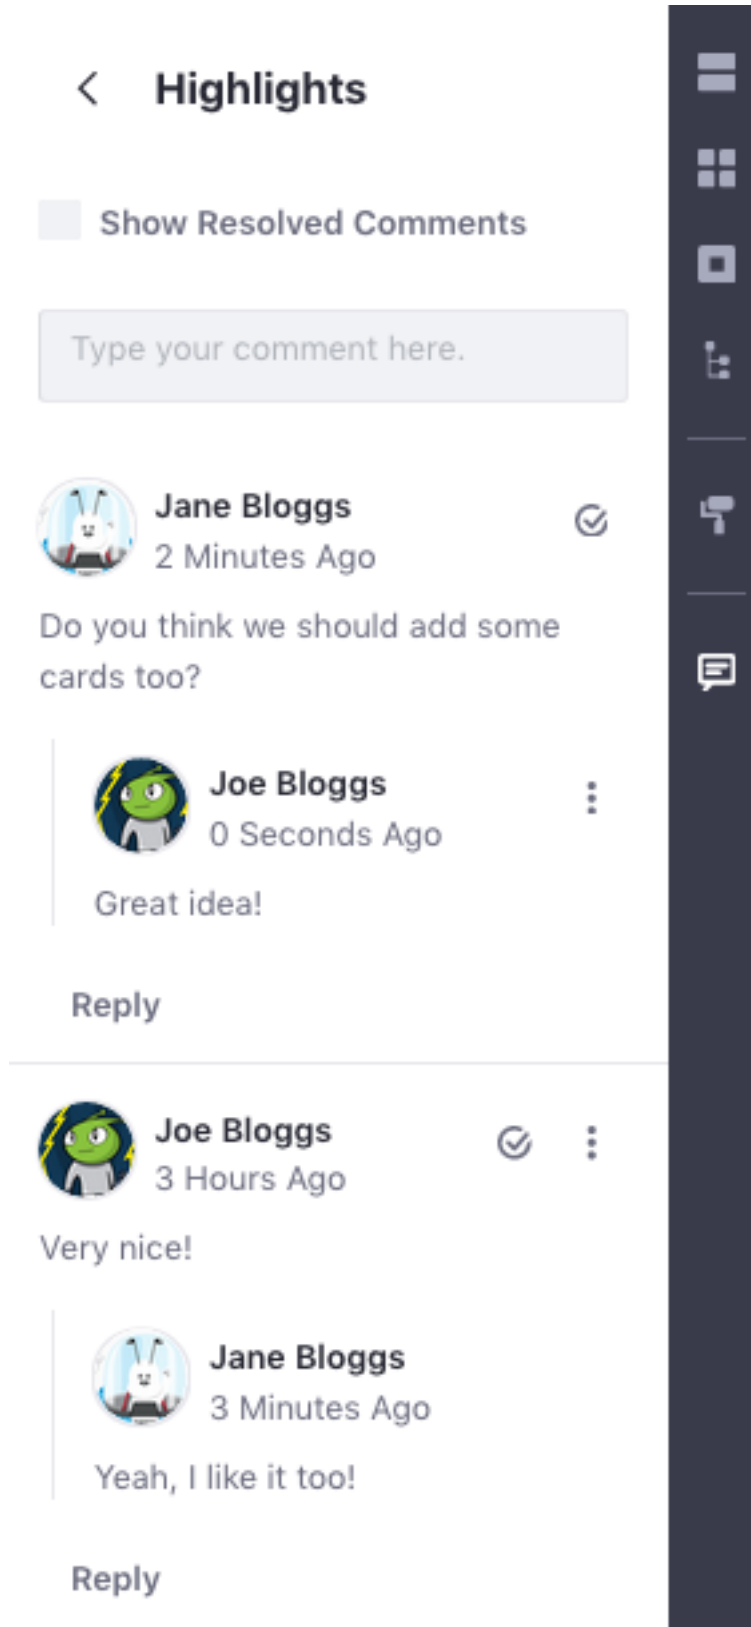


Figure 60.8: When creating content pages, you and your team can comment on any fragments.



Figure 60.9: The title bar has several tools built into it.

Great! Now you know how to use the content page interface!

BUILDING CONTENT PAGES

To demonstrate Content Page capabilities, create a *Landing Page* for a space shuttle tour based on this design:

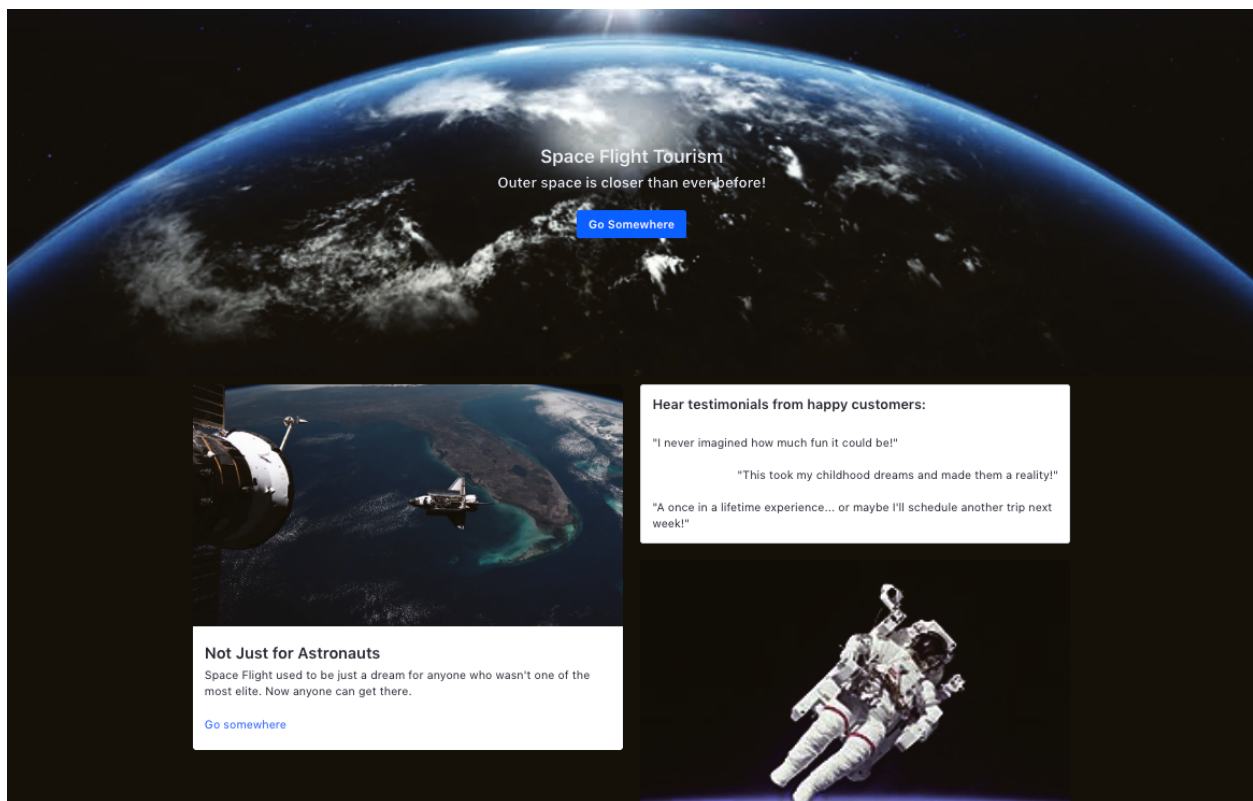



Figure 61.1: You have lots of flexibility when arranging Fragments on a page.

61.1 Creating a Content Page

To create the page, follow these general steps:

1. Add the Fragments you need to define the basic structure of the page.
2. Edit the text for the current page.
3. Edit the images.
4. Publish the page.


You'll complete this process next.

1. From Site Administration for your Site, go to *Site Builder* → *Pages*.
2. Click the  and select *Content Page*.
3. Set the *Name* as *Space Landing Page* and click *Save*.

Now you're in Content Page creation. As you work, a draft of the page is automatically saved, but you must click *Publish* to make it available for use.

4. Open *Sections* → *Basic Sections* and drag a *Banner Center* Fragment onto the page.
5. Open *Section Builder* → *Layouts* and add a 2 Column layout below that.
6. Open *Basic Components*.
7. Add a *Card* to the left column of the layout.
8. Add a *Paragraph* to the right side, and then an *Image* below that inside the same column.

Now that the structure is defined, start editing the text and images.

1. Click in each text box and edit the text to be relevant to your goal of directing potential customers to space shuttle tours.
2. Click on the main banner image, and then on the  icon.
3. Select one of the many space flight images that are surely filling up your Documents and Media library, upload a new one, or specify a URL.

Anytime during Fragment creation, you can remove, duplicate, or configure the Fragment.

Configuring a Fragment lets you modify the default options for a provided Fragment. This also lets you duplicate a Fragment and configure duplicates differently, so you can reuse base Fragments instead of developing new ones. When duplicating a Fragment, its configuration and editable elements are also copied.

Note: Fragment configuration and duplication are available in Liferay DXP 7.2 SP1+ and Liferay Portal GA2+.

This is looking good so far, but the one difference between the design mockup and the final result is that the background was black for the original. To finish it up, change the background color for the section to black.

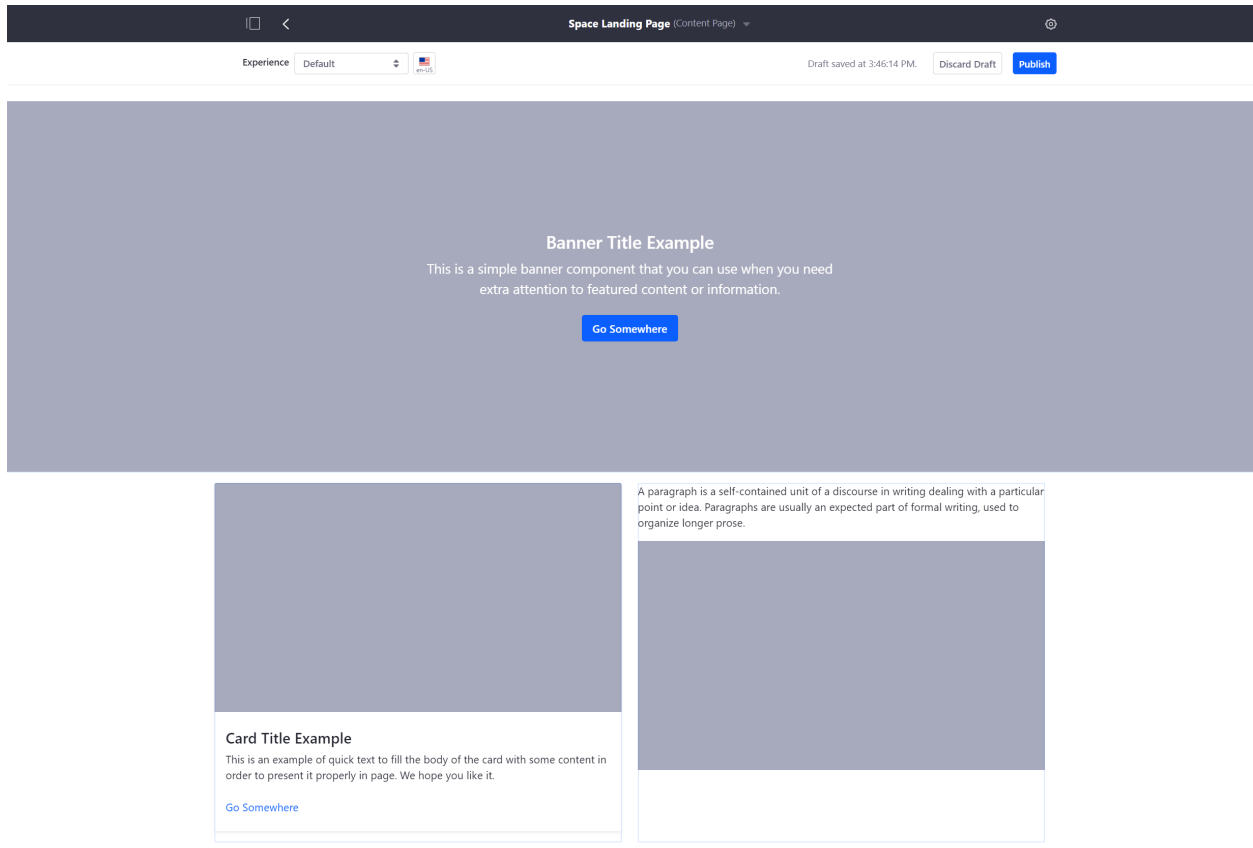


Figure 61.2: You have lots of flexibility when arranging Fragments on a page.




Figure 61.3: Edit the text and formatting as you see fit.

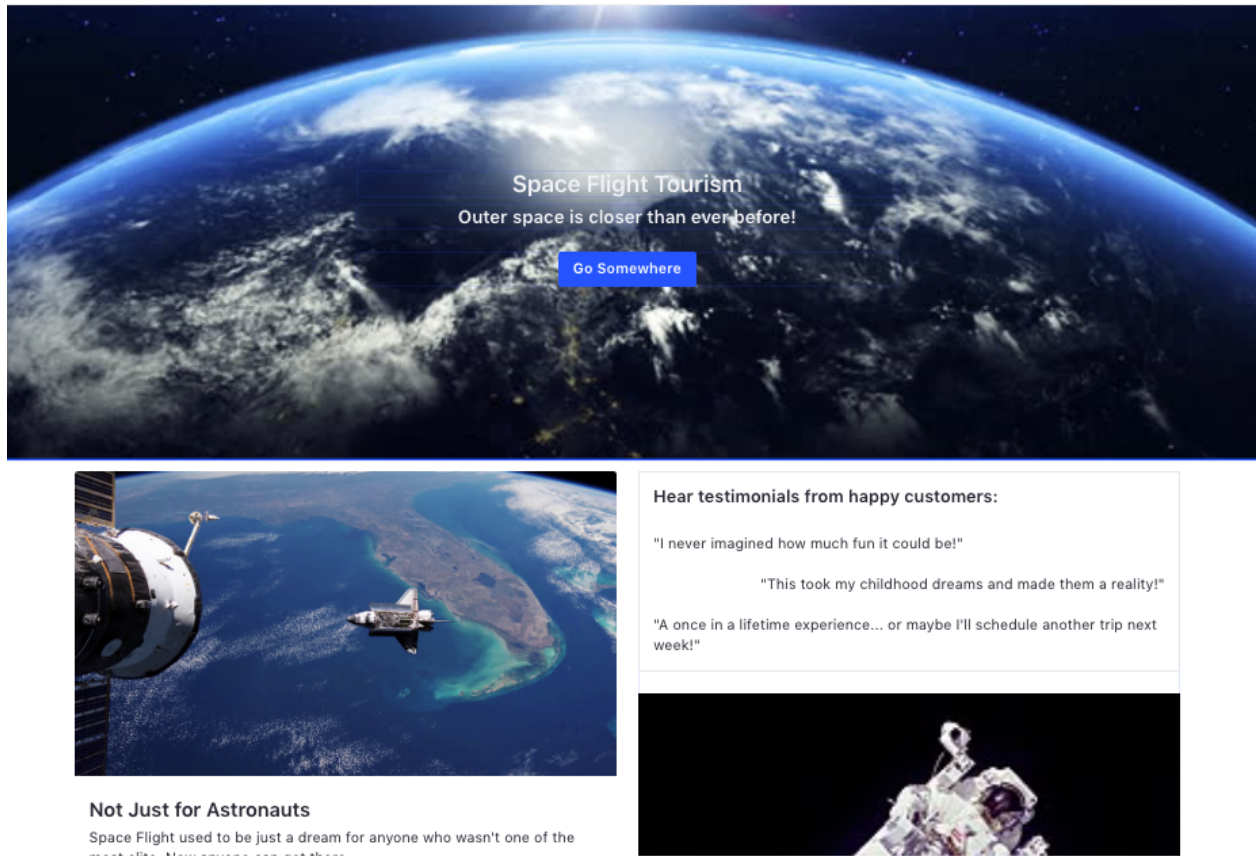


Figure 61.4: Add some images, and the big picture comes together.



Figure 61.5: Add some images, and the big picture comes together.






Figure 61.6: You can change the background color, image, or edit spacing and padding for a section. You can also remove it.

1. Click on the bottom Section and some additional icons will appear.
2. Click on *Background Color* , and select Black.
3. Finally, publish your page.

In just a few minutes, you used the power of Content Pages and Fragments to go from nothing to perfectly recreating a page design. To take it to the next level, head over to the Segmentation and Personalization guide.

PROPAGATION OF CHANGES


If you make an update to a Page Fragment or Content Page Template, it doesn't automatically propagate changes, but you can access the *Usages and Propagation* page to selectively propagate changes.

1. From the Site Administration menu, go to *Site Builder* → *Page Fragments*
2. Select the Collection containing the changed fragment.
3. Open the *Actions* menu () for the Fragment and select *View Usages*.

The *Usages and Propagation* page shows a list of every Page, Page Template, and Display Page that uses the selected Page Fragment. You can then selectively propagate Fragment changes to any or all of the pages listed. You can use the various filters and selection options to apply updates to pages quickly.

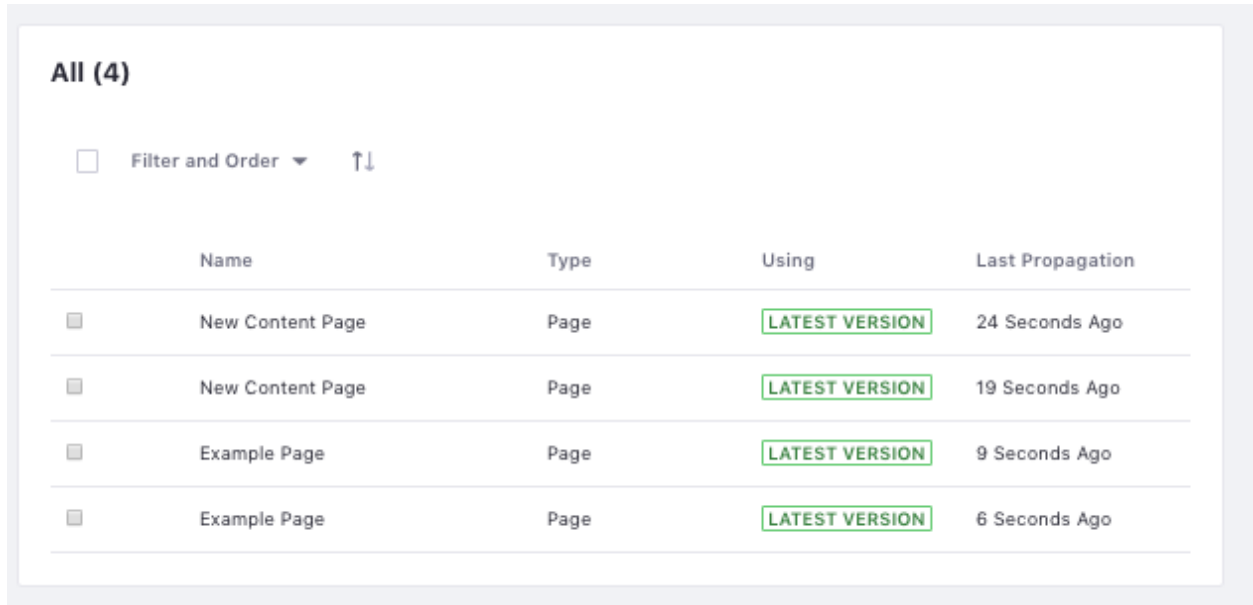
Note: Beginning in Liferay DXP 7.2 SP1+ and Liferay Portal 7.2 GA2+, you can propagate changes from global Fragments to their usages on child Sites. See *Creating and Managing Fragments* for more information on global Fragments.

To update a page or template,

1. Select the page or pages you want to update by checking the box next to the page name.
2. Click the *Propagate* icon ()

After you propagate changes, visit any affected page to verify there were no unexpected side effects of the changes.

****Note:**** Developers or others testing a Fragment can enable Fragment change propagation from the Control Panel. You can learn more about this [\[here\]\(/docs/7-2/frameworks/-/knowledge_base/f/managing-fragments-and-collections#propagating-fragment-changes-automatically\)](#). It's recommended to only leverage this functionality during testing, as automatic propagation on the production environment can cause unintended consequences.



The screenshot shows a web interface for viewing usages and propagation. At the top left, it says "All (4)". Below that is a "Filter and Order" dropdown menu with a sort icon. The main content is a table with the following columns: Name, Type, Using, and Last Propagation. There are four rows of data, each with a checkbox in the first column.

| | Name | Type | Using | Last Propagation |
|--------------------------|------------------|------|----------------|------------------|
| <input type="checkbox"/> | New Content Page | Page | LATEST VERSION | 24 Seconds Ago |
| <input type="checkbox"/> | New Content Page | Page | LATEST VERSION | 19 Seconds Ago |
| <input type="checkbox"/> | Example Page | Page | LATEST VERSION | 9 Seconds Ago |
| <input type="checkbox"/> | Example Page | Page | LATEST VERSION | 6 Seconds Ago |

Figure 62.1: Viewing the Usages and Propagation page.

Changes to existing `editable` fields are not propagated since this overwrites content currently in content pages. To force propagation to content in an `editable` field, a developer must change the field ID. Any content created in that field no longer appears in the Content Page when the changes are propagated, but it remains in the database and can be retrieved using the old ID.

Next you'll learn how to create your own Page Fragments.

CREATING PAGE FRAGMENTS

Fragments enable you to create rich content. With all of the included Fragments and Collections, many projects can be completed using just what ships with Liferay DXP. Sometimes, though, you must create your own Fragments. Fragments are built using HTML, CSS, and JavaScript. For a deeper dive into creating a custom Fragment, see *Developing Fragments*.

In this article, you'll learn about the Page Fragments interface.

63.1 Creating and Managing Fragments

To navigate to the Page Fragments interface,

1. Go to Site Administration. Make sure the Site where you want to work is selected. If you prefer creating a Fragment that's available for all Sites, navigate to the *Global* Site and create your Fragment there. Global Fragments are inherited by child Sites, so they can only be edited from the Global Site. Any resources the Global Fragment references (e.g., image) from the Global Site is copied to a Site that leverages the Fragment.

Note: Creating Global Fragments is available in Liferay DXP 7.2 SP1+ and Liferay Portal 7.2 GA2+. To expose this feature in the initial releases of those versions, however, you must create a `.config`` file. Create the ``com.liferay.fragment.web.internal.configuration.FragmentGlobalPanelAppConfiguration.config`` file and add the ``enabled="true"`` property. Then copy it to your Liferay DXP instance's ``osgi/configs`` folder.

2. Select *Site Builder* → *Page Fragments*

Fragments are organized in *Collections*. The main Page Fragments page shows available Collections (out-of-the-box Fragment Collections to start), allows Import and Export, and enables you to

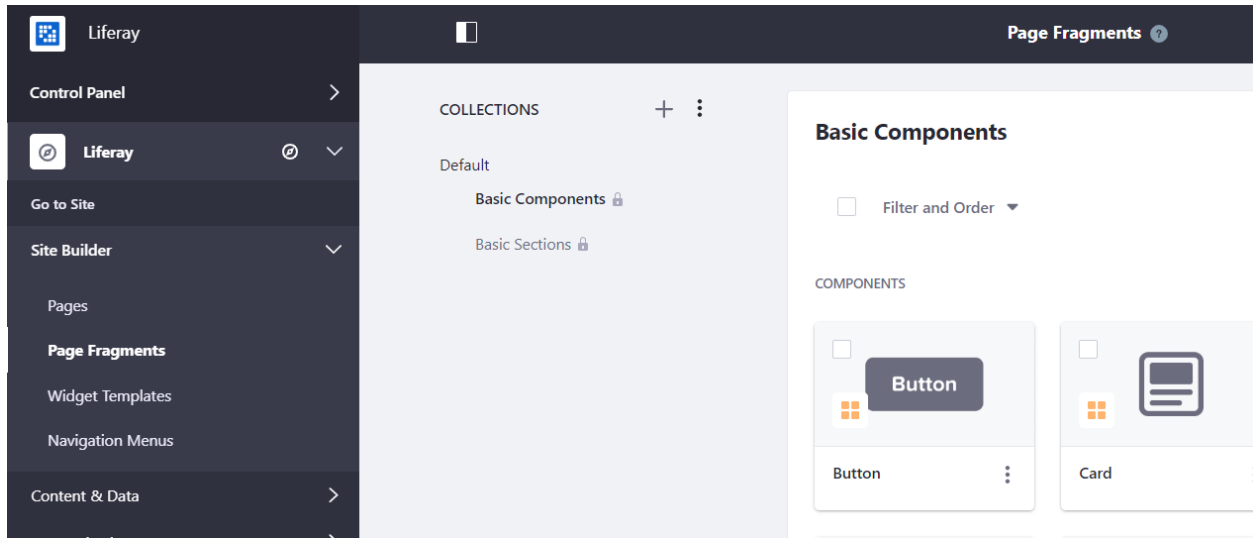



Figure 63.1: Here is the Page Fragments page with no custom Fragments or Collections created.


create Collections. You can also manage the organization and display of Fragments and Collections you have created.

Note: You cannot edit a default Fragment. If you'd like to provide an edited version of a default Fragment, you can copy it to your custom Collection and edit it there. To do this, navigate to the default Fragment Collection and click the Fragment's *Actions* () → *Copy To* button. Then select the Collection to copy the default Fragment to. Copying default Fragments is available in Liferay DXP 7.2 SP1+ and Liferay Portal 7.2 GA2+.

To create a Fragment, you must first create a Collection.

1. Click *New* → *Collection* to add a Collection.
2. Give the Collection a *Name* and *Description* and click *Save*.

Collections help you organize Fragments, and can be used to differentiate between different types of Fragments or Fragments used by different groups or departments. Next, create a Fragment inside the Collection you created.

1. Click on the Collection you created.
2. Click the *Add* icon () to create a Fragment.
3. Choose either *Section* or *Component*.
4. Give it a *Name* and click *Save*.

The Fragment development environment appears. Each pane in the editor has a different function:

- The top left pane is for entering HTML.

- The top right pane is for entering CSS.
- The bottom left pane is for entering JavaScript.
- The bottom right pane provides a live preview as you work in the other panes.




Figure 63.2: The Fragments editor provides an environment for creating all the parts of a Fragment.

In addition to standard HTML, CSS, and JavaScript, developers can also embed widgets and provide fields containing text and images that can be edited during the final Content Page publication process. Fragment development is covered in depth in *Developing Fragments*.

You can also include resources in your Fragment Collection that your Fragments can reference. This is helpful when exporting/importing Fragments: their resources are automatically included. If they're stored somewhere else (e.g., Documents and Media), you must export/import them separately. Click the *Resources* tab for your Collection and add resources (e.g., image, document, etc.) there.

Once added, you can reference resources from your Fragment's code without worrying about their availability. You can learn more about doing this in *Including Default Resources in Fragments*.

Next you'll learn how to import and export Page Fragments.




Figure 63.3: The Resources tab can be selected from the Fragment Collection.

EXPORTING AND IMPORTING FRAGMENTS


Often you'll want to reuse or re-purpose code from Fragments. To do this, export your Fragment Collections.

64.1 Exporting Fragments

There are two ways to export Fragments:



1. Export a Collection or set of Collections.
2. Export some Fragments outside of a Collection.

To export a single Collection,

1. Go to *Site Administration* → *Build* → *Page Fragments*.
2. Next to *Collections* click *Actions* () and select *Export*.
3. Select a Collection or multiple Collections to be exported. Each collection exports in a separate file.

Each Collection .zip contains all Collection data and Fragments.

To export individual Fragments,

1. Click on the Collection that you want to export.
2. To export all Fragments in the Collection without exporting any Collection data, click *Actions* () *Export* next to the Collection name. A .zip file is generated and downloaded automatically.
3. To export a single Fragment, click *Actions* () → *Export* next to the Fragment. A .zip file is generated and downloaded automatically.

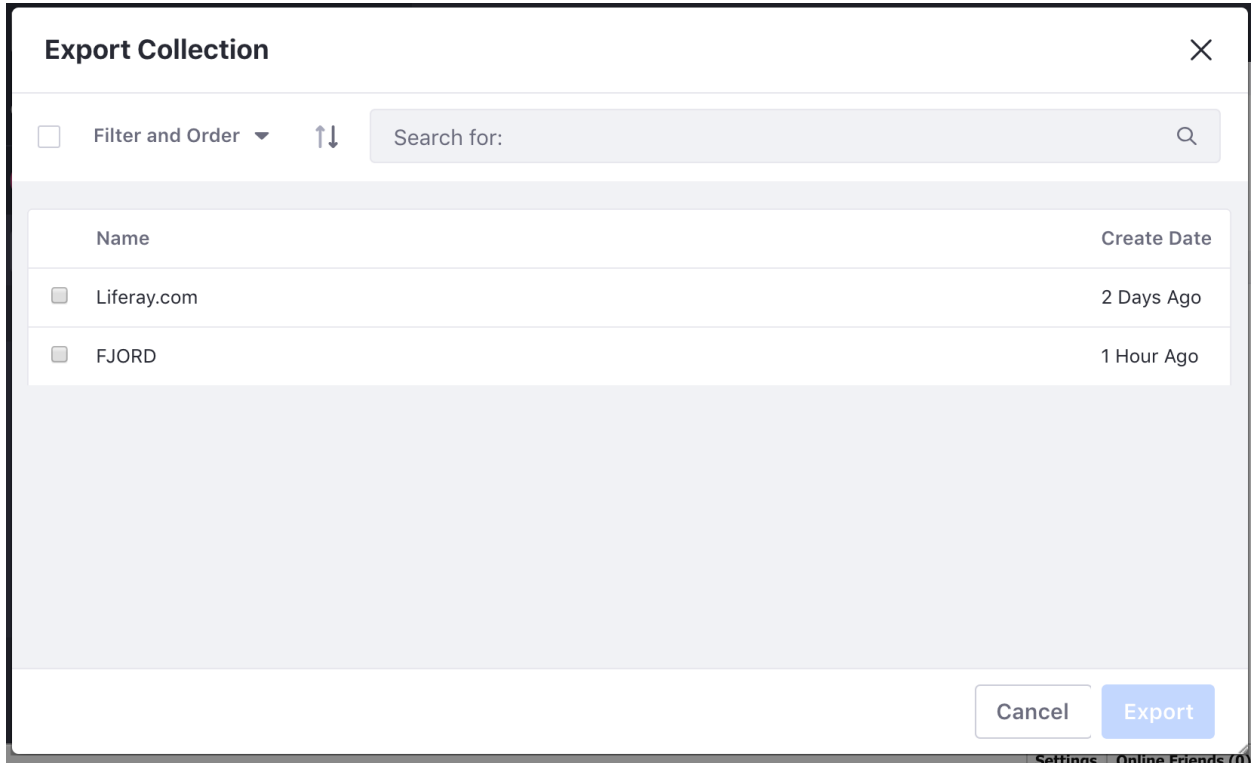


Figure 64.1: Select Collections to export.

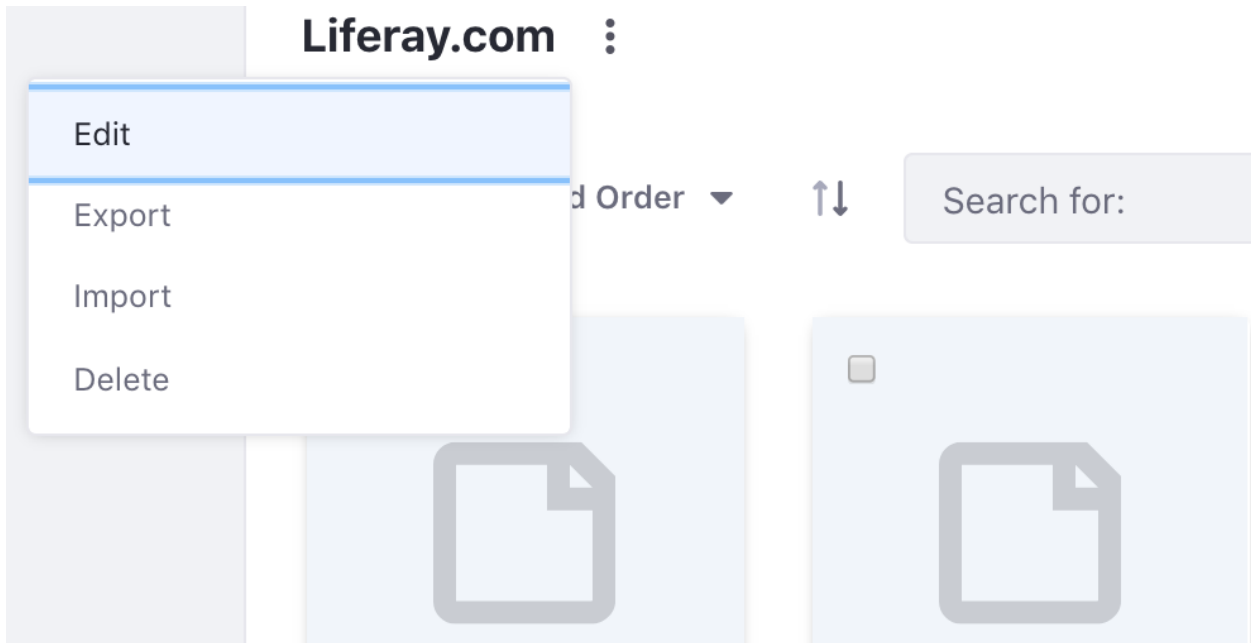


Figure 64.2: Exporting all of the Fragments in a Collection.

Note that if you export a single Fragment or a group of Fragments without a collection, they must be imported into an existing Collection.

You can also export Global Collections and single Fragments from your Site.

Now it's time to import your Fragments to where you need them.


64.2 Importing

You can import a Collection that was created in Liferay DXP, a Collection created using external tools, or Page Fragments without a collection. When you first import Page Fragments, they aren't available for use until you have approved them for use. This is to ensure that there are no errors in any imported fragments before they are added to a page.

See *Developing a Fragment Using Desktop Tools* for more information on creating and importing Fragments using other tools.

64.3 Importing Collections

To import a collection, follow these steps:

1. Go to *Site Administration* → *Build* → *Page Fragments*.
2. Next to *Collections*, click *Actions* () and select *Import*.





Figure 64.3: Importing and exporting Collections is accessed from a single menu.

3. On the next screen, click *Choose File* and select the file you want to import.
4. If you want to replace an existing collection, make sure the box is checked for *Overwrite Existing Files*.

5. Click *Import* and the collection is uploaded.

64.4 Importing Individual Page Fragments

You can also import a single Page Fragment or a set that was exported outside of a collection.

1. From the root level of the Fragments page, click on an existing Collection where you want to import the Fragments.
2. From inside the Collection click the *Actions* () button in the top right corner of the page.
3. Select *Import*.
4. Drag-and-drop or click *Select* to upload the Fragments .zip.
5. Click *Import*

The Fragments are imported into the Collection.

Exporting and importing fragments is the preferred way to share code or bring it into your Site.

USING WIDGET PAGES

Widget Pages are composed of *widgets*. A widget is any application that you can add to a page. Widget Pages are constructed by adding widgets to the page and filling them with content.


A widget could be a wiki display or a dynamic publishing tool like the Asset Publisher. The content you display with widgets could be long-form text or an image gallery, or anything in between. In this section, you'll learn to create widget pages and build content with them.

CREATING WIDGET PAGES

Widget Pages are the classic type of page in Liferay DXP. They're simple to create and fill with content or functionality. You create a blank page, define a layout, and then add widgets to the layout. Widgets can display content or provide some tool or function for Users.

66.1 Adding a Widget Page

When you first start Liferay DXP you get a widget page by default as your home page. To create a new widget page,

1. Go to *Site Administration* → *Site Builder* → *Pages*.
2. Click the *Add* icon () in the top right and select *Public Page* to add a new page.
3. Select *Basic Pages* if it is not selected by default.
4. Select the *Widget Page* type.
5. Name the page *Community* and click *Save*.
6. On the next screen, you can select a *Layout Template* or manage other options. Leave the defaults and click *Save*.

Creating a page by default also adds it to any *Navigation Menus* that are configured to have new pages added to them. If you don't want a new page added to a specific *Navigation Menu* that is listed during page creation, uncheck the box for that menu.

Your new page is now added to the navigation.

1. Click the logo in the top left of the page to go back to your Site's front page. The page you just created appears in the main navigation.
2. Click on *Community* to go to the page.

Currently the page is empty. Next you'll add some widgets to give it functionality.

General SEO Look and Feel Advanced

Name *
Community en-US

Hidden from Navigation Menu Widget
 NO ⓘ

Friendly URL ?
http://localhost:8080/web/guest
/community en-US

Create an empty page you can lay out manually.

| | | | |
|-----------------|-------------------|-------------------|------------------|
| 1 Column | 2 Columns (5...) | 2 Columns (3...) | 2 Columns (7...) |
| 3 Columns | 1-2 Columns (...) | 1-2 Columns (...) | 1-2-1 Column... |
| 1-2-1 Column... | 1-3-1 Columns | 1-3-2 Columns | 2-1-2 Columns |

Figure 66.1: Create a page called *Community* with two columns.




Figure 66.2: Your page has been added to the navigation automatically.

66.2 Adding Widgets to a Page

To add a widget to a page, go to the page and click the *Add* button (🧩) from the top menu and select the *Widgets* tab. You can either browse through the categories of available widgets until you find the one you want, or you can search for widgets by name. Once you've found a widget, click the *Add* button to add it to the current page. Once there, you can drag it to a new position. Alternatively, you can drag the widget directly from the Widgets menu to a specific location on the page. Follow the steps below to add some Collaboration apps to the Lunar Resort Site.

1. From the top menu, select *Add* → *Widgets*.
2. In the menu that appears, expand the *Collaboration* category.
3. Drag the *Blogs Aggregator* widget onto the right column of your page.
4. Next, drag the *Wiki* app to the left column.

See how easy it is to add applications to your pages? You've added the Wiki app and Blogs Aggregator app to a page.




Figure 66.3: Your page layout options are virtually limitless with a slew of application and layout combinations.

If the default layout options provided aren't enough, you can create your own. For more information about developing custom layout templates, see the tutorial *Layout Templates with the Liferay Theme Generator*.

Next, you'll look at creating reusable templates for widget pages.

CREATING WIDGET PAGES FROM TEMPLATES

Page templates provide pre-configured pages to reuse. There are two types of page templates in 7.0: Widget Page templates consist of a portlet layout and configuration. Content Page templates are constructed from Fragments. You can read about Content Page Templates in this article.

Three sample layout page templates are installed by default:

- **Search:** Contains a search bar and configuration to display various facets.
- **Wiki:** Provides a page with three applications related to authoring a wiki.
- **Blog:** Provides a page with three applications related to blogging.




Figure 67.1: The Blog page template is already available for use along with the Search and Wiki page templates.


To add a new widget page template,

1. Go to *Site Builder* → *Pages*.
2. Select the *Page Templates* tab.
3. Click *New* and create a collection named *Lunar Resort Templates*.
4. Click the *Add* icon (+) and select *Widget Page Template*.

5. Enter a name.
6. Click *Save*.

The editing page for the template appears. You can add widgets to the page or access page configuration now. The changes you make are instantly applied to the template.

If you want to edit the template again,


1. Go back to the *Page Templates* tab.
2. Click the *Actions* icon ().
3. Click *Configure*.

Note that after a new page template has been created, the default permissions only allow the creator to use the page template. To give other users access to it,

1. Use the Actions menu for the template and select *Permissions*.
2. In the matrix of Roles and permissions, check the *View* permission for the Roles needed to see the page template in the list of available page templates when creating a new page.

If you want any user who can create a page to be able to use the page template, check the *View* permission for the *User* Role.

To use your template to create a new page,

1. Go to Site Administration and select the *Pages* option from the *Site Builder* menu dropdown option.
2. Click the *Add* icon ().
3. Inside the Lunar Resort collection, select the page template that you created.
4. Enter the name of your page and click *Add*.

Pages based on templates can inherit changes from the page template:



Figure 67.2: You can choose whether or not to inherit changes made to the page template.



By default, when a Site administrator creates pages based on a page template, future changes to the template are automatically propagated to those pages. Site administrators can disable this behavior by disabling the *Inherit Changes* selector. Occasionally, propagation for page templates fails due to unintended errors. To learn how to manage a failed page template propagation, visit the [Propagating Changes from Site Templates to Sites](#) article.

If staging has been enabled, changes to the page template are automatically propagated to the staged page. These changes must still be approved before the page is published to live. For this reason, the automatic propagation of page template changes to the staged page cannot be turned off and the *Inherit Changes* selector does not appear.

You can read more about staging in the Staging1 articles.

67.1 Sharing Widget Page Templates

When importing pages to a new site or environment, you must also import templates associated with those pages. Generally templates are included automatically when an associated page is exported, but if not you can export the template collection separately so the page can be imported to the new environment. To export page templates,

1. Go to *Site Management* → *Site Builder* → *Pages*.
2. Select the *Page Templates* tab.
3. At the top right of the page, click *Options* () → *Export/Import*.
4. On the *Export/Import* page you can choose to export configuration data and change which Collections and templates are being exported.
5. When you're done configuring the export, click *Export* and save the exported LAR file.
6. On the target environment, go to *Site Management* → *Site Builder* → *Pages* and select *Page Templates*.
7. At the top right of the page, click *Options* () → *Export/Import*.
8. Select the *Import* tab.
9. Upload the LAR with your template data. If the LAR contains additional content you don't want to import, you can deselect it.

Once the template has been imported, the page can be imported normally to your new environment. For more information on exporting/importing content, visit the *Importing/Exporting Sites and Content* article.

BUILDING A RESPONSIVE SITE

Now more than half of all page views in the world come from mobile devices like phones and tablets. That means that if your pages don't look good on mobile devices, your pages don't look good for more than half the people looking at them. Liferay DXP can provide the best experience possible no matter what device you're using.

Most of the heavy lifting for mobile optimization comes from your developers. Themes must be designed to be responsive. Web Content Structures and Templates and Page Fragments must be created with code that transitions gracefully to fit smaller screens on mobile devices. But there's work to do for marketers and administrators as well.

BUILT-IN MOBILE SUPPORT

Out of the box, there are several features that help make your pages look just as good and have the same functionality on mobile devices as they do on a desktop:

- Liferay Widgets and custom widgets that use Liferay's UI frameworks automatically scale to fit the screen size.





Figure 69.1: A widget adjusts its size.

- UI elements like the navigation and Product Menu automatically adjust to remain usable on smaller screens.
- When the screen width is low, Liferay combines columns so that all content remains legible.
- For web developers, Liferay's theme tools provide a number of tools to help ensure optimum mobile performance.

For most business users, this means that all you need to do to display pages on Mobile device is to create a page. However, you also have tools available to verify that everything displays as intended. The Device Simulator  is a powerful tool that shows you how pages look on different devices.




Figure 69.2: The main navigation adjusts its size.

69.1 Using the Device Simulator

When creating a page or reviewing a page before it is published, one of your most important tools is the Device Simulator found in the top right corner of every page. The simulator lets you view the current page in a number of resolutions based on different display types. There are three predefined options:

Desktop: Fixes the width to display the page at full size.

Tablet: Puts your page in a box as if it is being displayed on a tablet. It also activates some of Liferay’s built-in mobile features.

Mobile: Puts your page in an even smaller box to demonstrate how the page looks to your average smartphone user.

There are also two options available to display

AutoSize: Provides another way to view the default behavior where the page shrinks and grows based on the width of the browser window.

Custom: Lets you enter a specific size for testing and fixes the height and width of the display.

Because modern mobile browsers are built on the same technology as desktop browsers, the behavior you see in the simulator should match the experience of users on mobile devices. In addition to making sure the basic layout looks good and that all functionality remains, it’s also important to make sure that automatic features—like how columns are combined at lower resolutions—don’t have unintended effects.

69.2 Designing Mobile Friendly Pages

Liferay DXP provides the tools you need, but building pages that provide a good experience across all kinds of devices still means working across all levels of web development and publishing. Theme developers must create themes that use Liferay’s frameworks to scale content well across all kinds of displays. Designers must have multiple screen sizes in mind when designing pages. And before anything it published it must be thoroughly reviewed to make sure that it provides the best experience to all of your users.

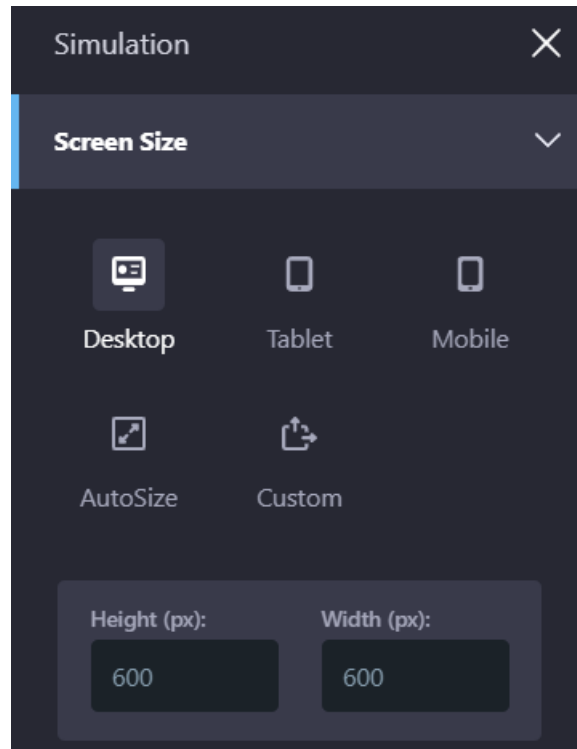


Figure 69.3: The Simulation panel defines multiple screen sizes.

Now that you've learned about Liferay's tools for making your website mobile friendly, let's look at your options for adapting to different types of mobile devices.

MOBILE DEVICE RULES

With mobile device rules, you can alter what gets displayed based on the device being used to access Liferay DXP. For instance, you can configure the look and feel of pages accessed by mobile device users differently from those accessed by users on a computer.

Whole Sites or individual pages can be configured for mobile device families. A family describes a group of devices. You can set rules that describe a category of devices, such as all Android devices or all iOS devices. You can define as many rules in a family as you need to classify all the devices for which you'd like to define actions. Families can be prioritized to determine which one applies to a given page request.

CREATING MOBILE DEVICE RULES

To configure mobile device rules, you need a way to find out the characteristics of the device. While some of the characteristics are provided by the device, most are not. For this reason, there are databases that contain information about thousands of devices. These databases make it possible to learn device details from the device type. Liferay DXP's Mobile Device Rules can connect to device databases so that you can use their device characteristics in your rules.

Important: For the features described in this article to work, you must install the Liferay Mobile Device Detection Lite (LMDD) app from Liferay Marketplace. This app provides the device detection database that's required to detect which mobile devices are accessing it.

You can develop plugins that integrate with other device databases. Even if you don't have a device database, you can still set up mobile device rules. They won't, however, be effective until a database is deployed, because the portal won't have enough information about the devices being used to make page requests.

To access the Mobile Device Families administrative page,

1. Open the *Product Menu*.
2. Use the *Site Selector* (🌐) to choose the Site that you want to define Mobile Device Rules for.
3. Select *Configuration* → *Mobile Device Families*.

You can also add families for all Sites by navigating to the Control Panel → *Sites* → *Sites* → *Global*. The Mobile Device Families administrative page displays a list of defined families and lets you add more. To add rules, you must first add a family.

1. Click *Add* button (+) to add a *New Device Family*.
2. Enter a *Name* and *Description*.
3. Click *Save*.
4. Click on the name of the Mobile Device Family to access the rules page.

The screenshot shows a form for creating a mobile device family. The 'Name' field is labeled 'Name *' and contains the text 'Android'. The 'Description' field is labeled 'Description' and contains the text 'All Android mobile devices'. Below the form are two buttons: 'Save' (blue) and 'Cancel' (white). At the bottom left, there is a breadcrumb trail 'Android' with a question mark icon. At the bottom right, there is a user profile icon.

Figure 71.1: Create a Mobile Device Family so you can create rules.

The rules defined for a family, along with the priorities of the families selected for a particular Site or page, determine which family's actions are applied to a given request. From the New Classification Rule page for a specific rule set, you can add a rule by specifying an operating system, rule type, physical screen size, and screen resolution. Remember that you can add as many rules to a family as you need in order to classify the devices on which you'd like to take actions.

1. Click on the *Add* button (+) to add a new rule.
2. Enter a *Name* and *Description*.
3. Select the classifications you want for this rule from *Operating System and Type*, *Physical Screen Size*, and *Screen Resolution*.
4. Click *Save*.

You'll notice after saving the classification rule that it's characterized as a *Simple Rule*. Only Simple Rules are included with Liferay DXP, but the rules are designed to be extensible, and additional rule types can be added by your developers.

Once you've created some mobile device families and added some rules to them, you're ready to create some actions. The actions defined for a family determine what happens to a request when the device is detected and the family has been found to apply.

You can add families to a Site, individual page, or page set from their respective configuration pages. To do it for a Page Set:

1. Go to *Site Builder* → *Pages* in your Site.
2. Click on *Configure* (⚙️) for the Public Pages.

Name *
Android

Description
Provide custom experience on Android

Operating System and Type

Operating System

- Any Operating System
- 3DS System Software
- AmigaOS
- Android

Device Type
Any

Physical Screen Size >

Screen Resolution >

Save Cancel

< Android (Android) ⓘ

Figure 71.2: Select the operating system and device type for your rule.

3. Select the *Advanced* tab and open the *Mobile Device Rules* option in the bottom menu.
4. Click *Select* to open the list of families that can be applied.

From the same page, you can access the configuration for an individual page, or you can configure Mobile Device Rules for an entire Site from *Configuration* → *Site Settings*. You can select multiple families for a particular Site or page and order them by priority. The families are checked in decreasing order of priority: the actions defined by the first family that applies are executed.

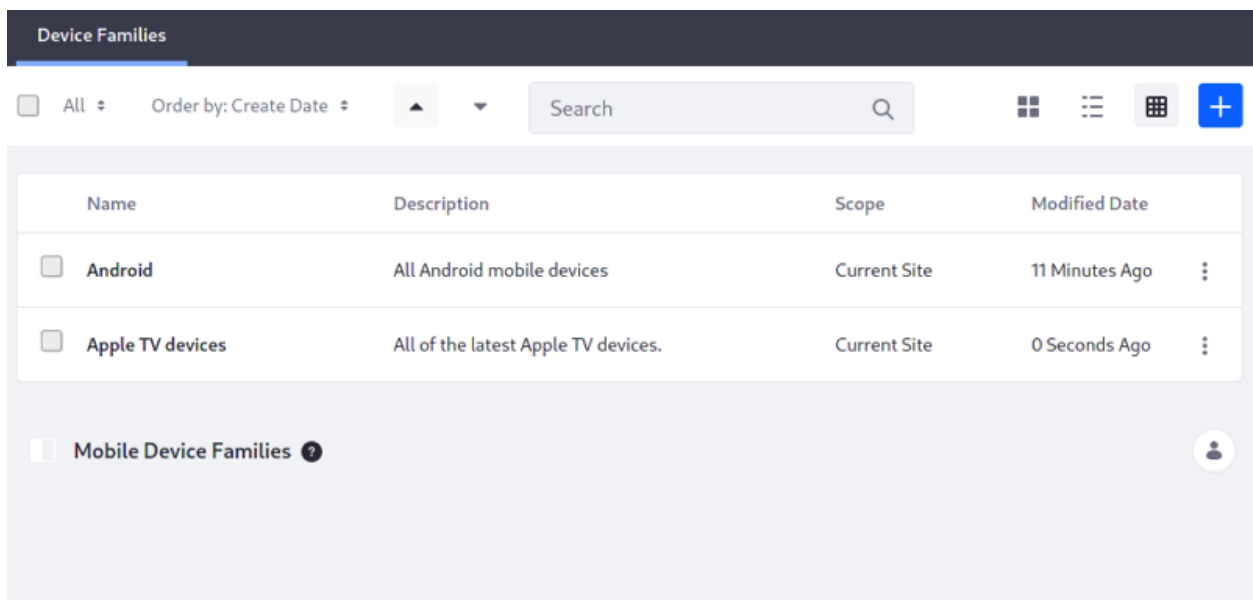


Figure 71.3: You can select a mobile device family to apply for a Site or page.

MOBILE DEVICE ACTIONS

After you've created families and applied rules to define those families, you can associate specific actions that occur when a user visits that Site on a device.

To add actions to a selected rule group:

1. Click *Configure* (⚙️) for the page, page set, or Site where you have configured a device family. The configuration page appears.
2. Click the *Advanced* tab.
3. Decline using the same mobile device rules from the public pages by setting the option to *No*. The device family section appears.

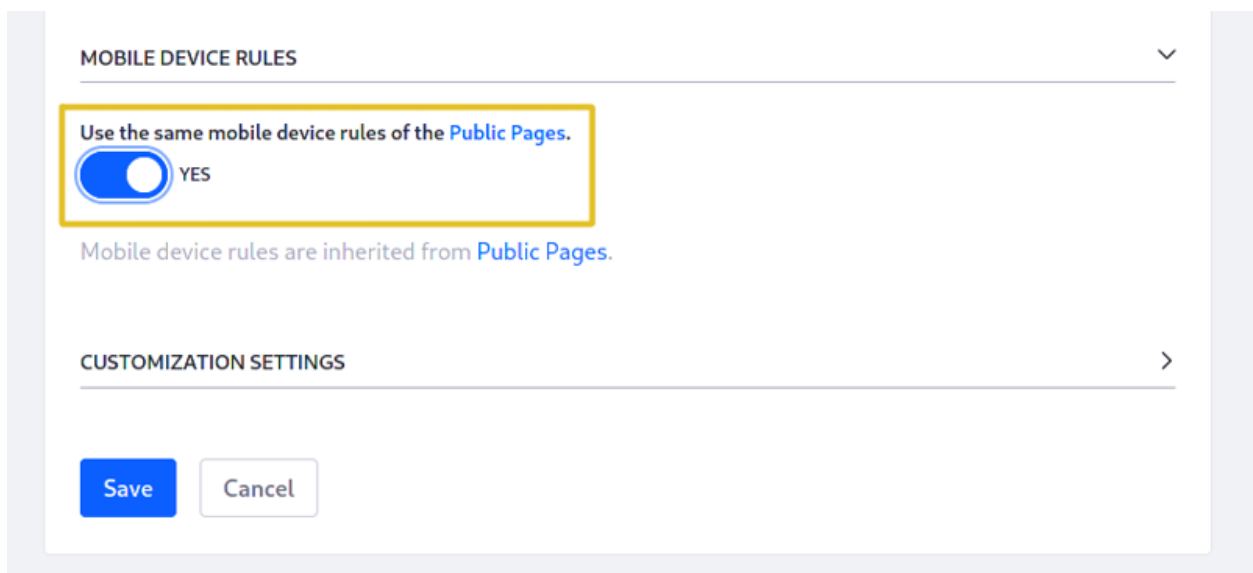



Figure 72.1: Disable using public pages device rules.

4. Click *Select* to open the Device Families page.

5. In the *Mobile Device Rules* section, click *Actions* () → *Manage Actions* next to the device family that you wish to add an action for.
6. Click *Add Action*.

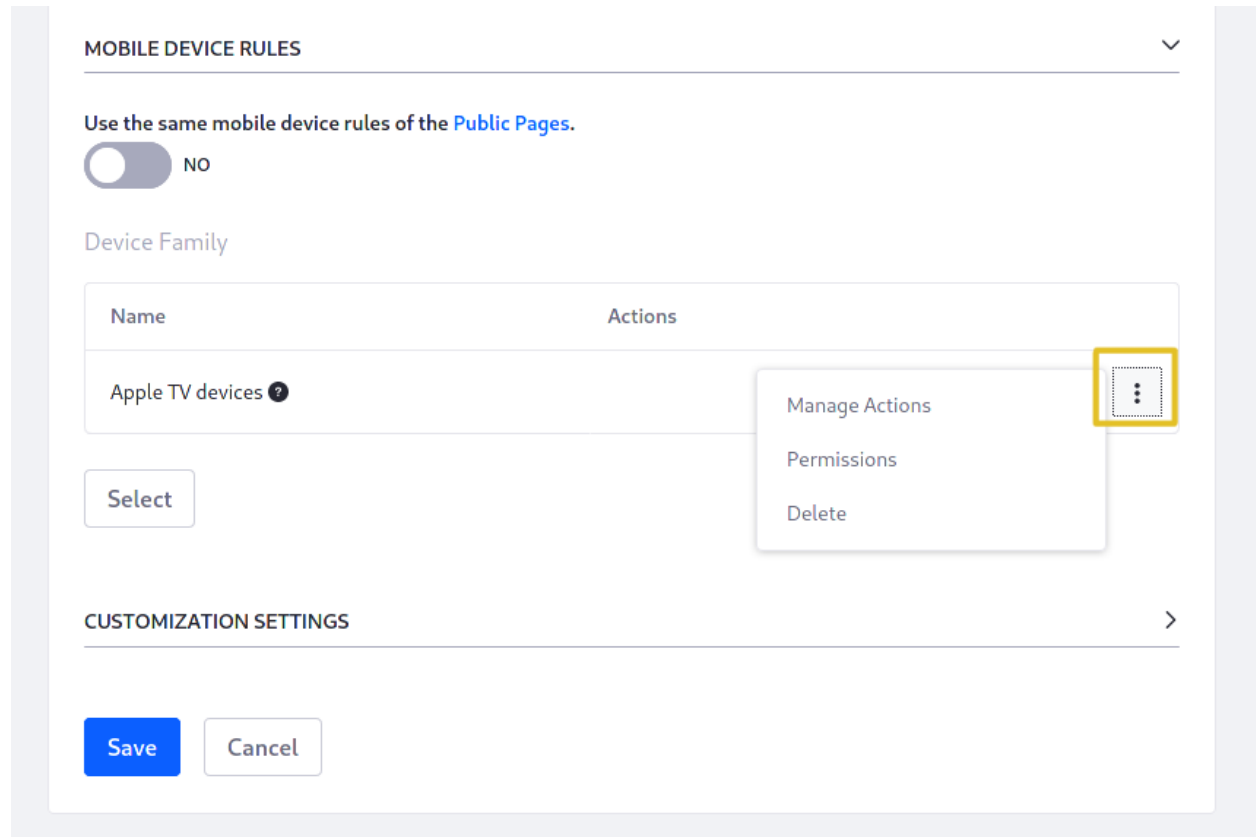


Figure 72.2: Getting to the Manage Actions page.

By default, there are four kinds of actions that can be configured for mobile families:

Layout Template Modification: Changes the way portlets are arranged on pages delivered to mobile devices. For example, you could have pages with more complex layouts automatically switch to a simpler template if it detects a mobile device—even if the resolution is theoretically high enough to support the standard layout.

Theme Modification: Selects a specific theme for different mobile device families. You’d have to have a mobile version of your Site’s theme that is automatically applied when a device hits your page.

URL Redirect: Sends mobile users to any URL. This can be used to direct mobile users to a mobile app download or a mobile version of the page.

Site Redirect: Sends mobile users to a different Site on your portal. In some cases, mobile content could be created on a mirror of your Site.

Tip: 7.0 was designed from the ground up to be responsive and adapt to any device that might be accessing it. Before creating new themes or forcing a layout template change, you should test

how the Site behaves using Liferay DXP out of the box. Certain features, like URL Redirects, can be disruptive and frustrating for users if used improperly.

Like mobile device rules, mobile device actions are extensible. Your developers can define custom actions in addition to the four actions provided by default.

To review, if you want to configure an action or actions that take place when mobile device requests are received, take the following steps:

1. Create a mobile device family to represent the group of devices for which to define an action or actions.
2. Define one or more rules for your family that describe the group of devices represented by your family.
3. Apply your family to an entire page set of a Site (all the public pages of a Site or all the private pages) or to a single page.
4. Define one or more actions for your family that describe how requests should be handled.

72.1 Mobile Device Rules Example

Now you'll look at an example of using mobile device rules. Suppose you want to create a rule so that when a Site is accessed by an Android device, a different layout is used. To set this up, you must follow the same four steps described above.

First create the Mobile Device Family:

1. Navigate to the *Mobile Device Families* page of *Site Administration*.
2. Click *Add Device Family* (+).
3. Enter *Android* for the *Name*.
4. Click *Save*.

Next create a rule for the family:

1. From the *Mobile Device Families* page, click on *Android*.
2. Click *Add Classification Rule* (+).
3. Name the rule *Rule 1*.
4. Under *Operating System* select *Android* (you can hold or to select multiple items).
5. Under *Device Type* select *All*,
6. Click *Save*.

The screenshot shows a mobile application interface for creating a new classification rule. The title bar at the top is dark blue with a white hamburger menu icon on the left, a back arrow, the text "New Classification Rule for Android and iOS Tablets" with a help icon, and a user profile icon on the right. Below the title bar, the form is organized into sections. The "Name" section has a text input field containing "Rule 1" and a language selector icon showing "en-US". The "Description" section has a large text area containing "Description" and another "en-US" language selector icon. The "Operating System and Type" section is expanded, showing a dropdown menu for "Operating System" with "Any Operating System" selected. Below it is a "Device Type" dropdown menu with "Any" selected. At the bottom of this section are two expandable options: "Physical Screen Size" and "Screen Resolution", each with a right-pointing chevron. At the very bottom of the form are two buttons: a blue "Save" button and a white "Cancel" button with a grey border.

Figure 72.3: Create the Classification rule.

As with the previous example, you only need one rule to describe your device family. Now you must apply the rule to some pages.

1. Go to *Site Builder* → *Pages* in Site Administration.
2. Click on the *Configure* icon for the *Public Pages*
3. Go to the *Advanced* tab.
4. Under *Mobile Device Rules*, select the *Android* device family.

Now you must define an action for your Android rule group to use a different layout.

1. Click *Actions* → *Manage Actions* for the *Android* rule.
2. Click *Add Action*.
3. Enter the name *Layout Template Modification*, and select the *Layout Template Modification* action type.
4. Select the *1 Column* layout template.
5. Click *Save*.

Now the Liferay Site's pages are presented to Android users with the 1 Column layout template.

Mobile Device Rules are a powerful way to manage the way pages and content appear on the various devices that access your Site. But remember to consider the power of modern devices and the experience of your users, and use this great power responsibly—to help users have a great experience on your website and to not interrupt or negatively impact that experience on whatever device they're using.

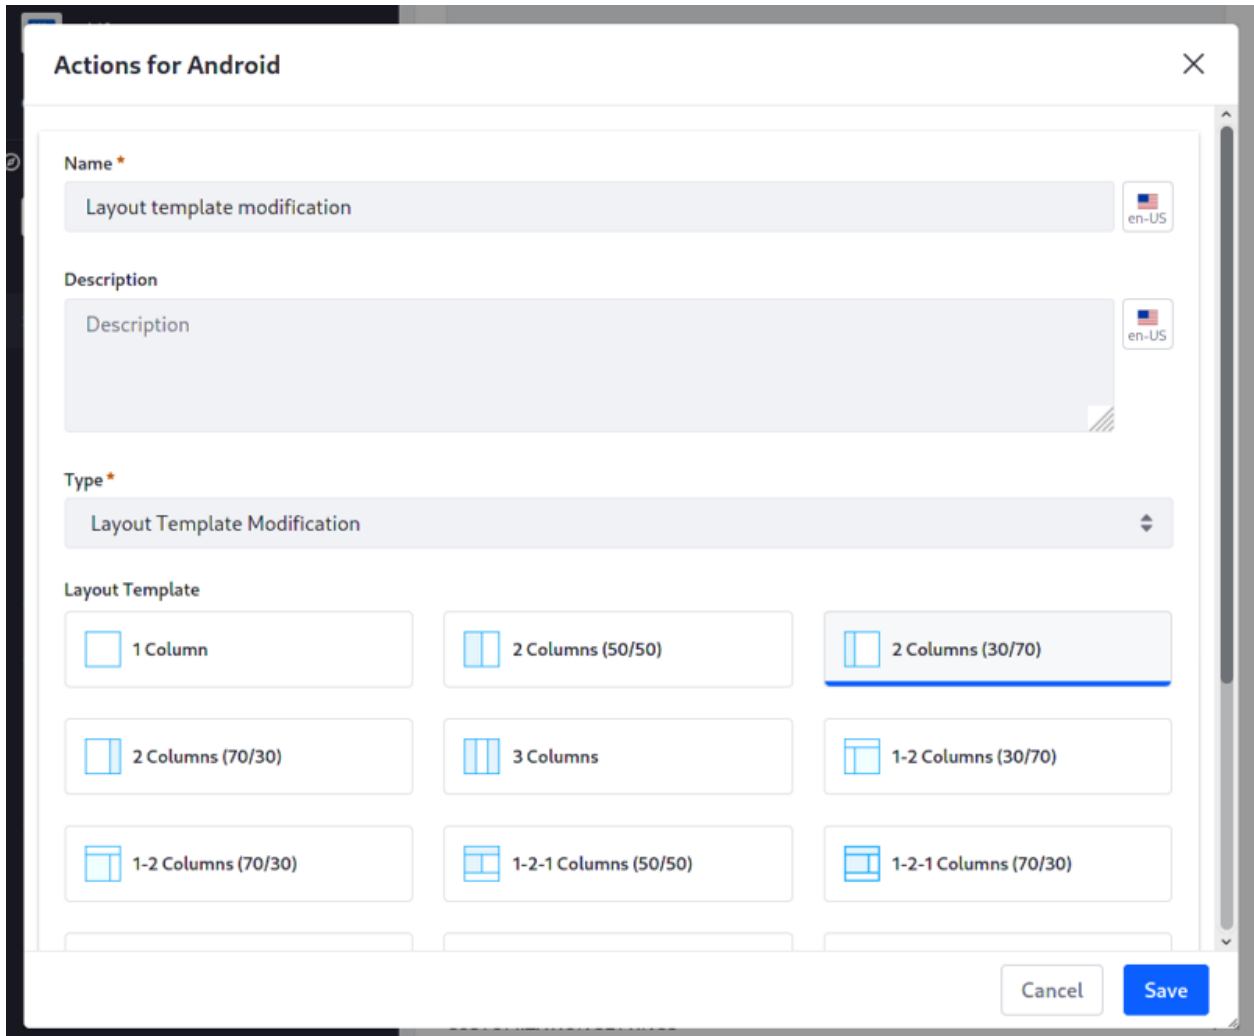



Figure 72.4: Create the Actions for Android.

USING FULL PAGE APPLICATIONS

Full Page Applications are the ideal way to display a Message Board, Wiki, or other application that demands a full page.

73.1 Configuring the Page

Creating a Full Page Application starts just like creating any other type of page.

1. Go to *Site Administration* → *Site Builder* → *Pages*.
2. Click the  icon.
3. Select the *Full Page Application* page type.
4. Give your page a *Name* and click *Save*.
5. Set the *Full Page Application* to *Wiki* and click *Save*.
6. Finish the page configuration and click *Save*.

Out of the box, you can set the Blogs, Wiki, Media Gallery, Message Boards, RSS, Hello Soy Portlet, Documents and Media, Form, or Application Authorization Request to be the sole application for the page. Developers can make their applications Full Page Applications.

7. Click *Go to Site* in the Site Administration menu, and then click on your page.

Now the page is configured to display the Wiki and only the Wiki. No other widgets can be added to the page, and the Wiki app cannot be removed.

Note that all of the applications that can be added to the page are non-instanceable and the content of whichever application you select is based on the instance for that site. So if you already had data in your Wiki it appears on this page.

If you want to configure the application to be scoped to this specific page, you can configure that through the application's settings.

General SEO Look and Feel Advanced

Name *

Test 2 en-US

Hidden from Navigation Menu Widget

NO ?

Friendly URL ?

http://localhost:8080/web/guest

/test-2 en-US

Create a page with one column that displays a single full page application.

Full Page Application

Wiki

Save Cancel

Figure 73.1: The Full Page Application configuration page.

1. From the page, click the (i) button for the Wiki and select *Configuration*.
2. From the *Wiki - Configuration* page, select the *Scope* tab.
3. Open the *Scope* menu and select *Space Wiki*.

Now the Wiki is scoped for this page, and doesn't share data with the Site or globally scoped Wiki.

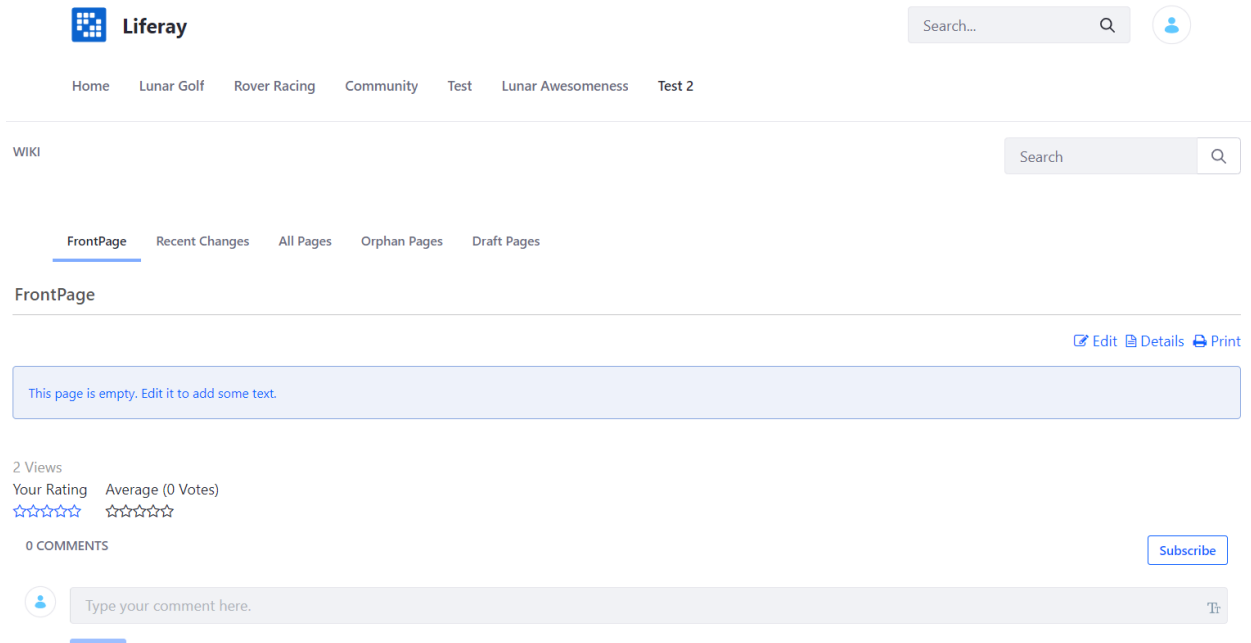


Figure 73.2: The Wiki displayed as a Full Page Application.

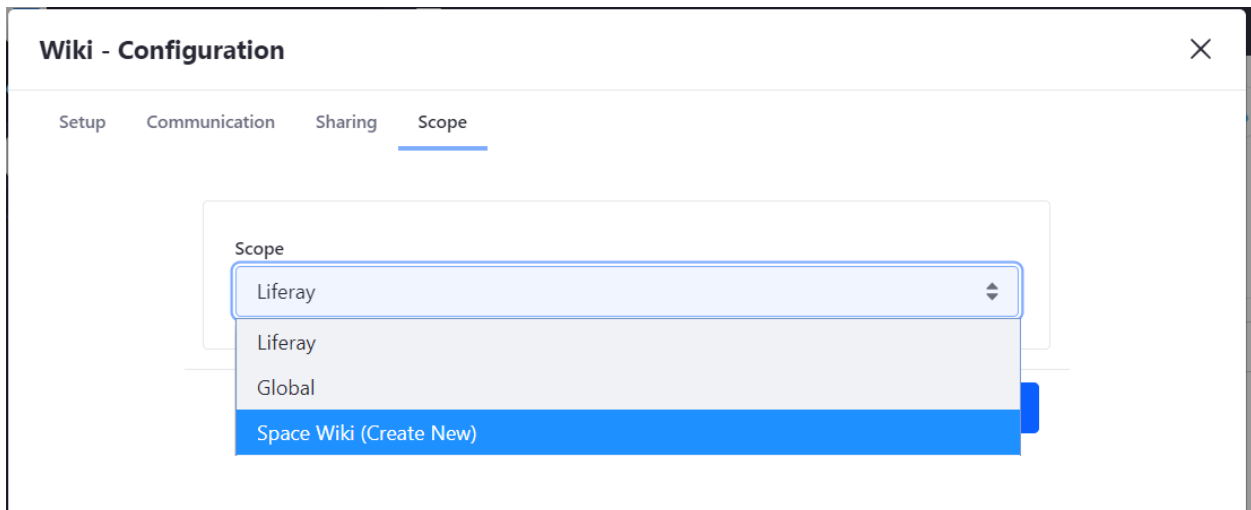


Figure 73.3: Configuring the scope.

MANAGING SITE NAVIGATION

Liferay provides powerful tools for creating and organizing pages. You can have anything from a simple, flat Site navigation to a complex hierarchy with tree of sub-pages nested down many levels.

7.0 lets you create Navigation Menus separate from your page hierarchy. Now you have the freedom to leave one-off marketing landing pages out of the navigation or create multiple navigation menus: a main menu, secondary menus, footer menus, and custom menus for anything that you can dream up.

Menus can differ by page: landing pages can show a simple list of frequently visited pages, and the rest can appear in secondary navigation. You can also create specific menus for different landing pages to direct users to content that is relevant to them.

Go to *Site Builder* → *Pages* to view the existing pages or create new pages. The Site hierarchy as displayed on *Pages* is the main reference for the organization of pages on that Site. While Navigation Menus can customize their organization and what appears and what doesn't appear, this menu is always the primary reference for the pages on your Site.

PAGE HIERARCHY

Using the Page Hierarchy, you create public and private pages and organize those pages in whatever order or structure that you see fit.

75.1 Creating a Page

New pages are created on the *Site Builder* → *Pages* page in Site Administration. Pages can be created as *Public Pages* which anyone can view, or *Private Pages* which can only be viewed by Site Members. To create a new page,

1. Go to *Site Administration* for the Site you want to work on, then *Site Builder* → *Pages*.




Figure 75.1: In the default site, initially only the *Home* and the hidden *Search* pages exist in the Public Pages Hierarchy.

2. Click *Add* (+) and select *Public Page*.
3. Select *Widget Page*.

4. Set the *Name* as *About Us*.
5. Click *Save*.

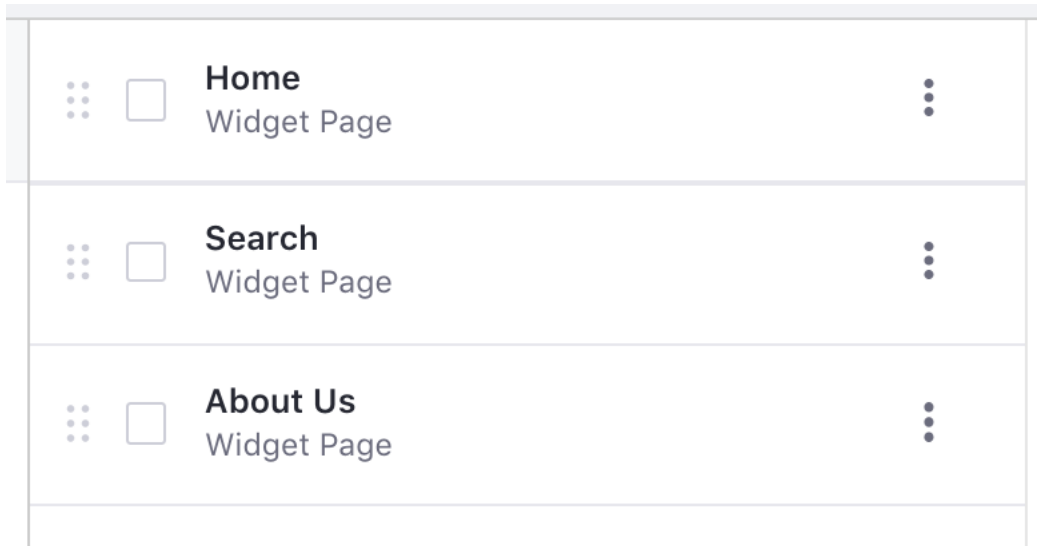


Figure 75.2: When you create a page, by default it is added to the site hierarchy.

Now that the page is created, it appears in the hierarchy, and you can move or organize its position there.

75.2 Organizing Pages

Drag and drop pages to reorder their position in the page hierarchy (and subsequently the default navigation that users see), and to nest them as subpages. The page at the top of the list is the *Home* page that users see automatically when visiting your Site.

1. Click *Go to Site* to view the organization of the navigation menu. You can see the order of the pages matches the order of the pages from Site Administration.
2. Drag the *About Us* page above the *Welcome* page in the list. It automatically becomes the *Home* page.
3. Click *Go to Site* to see how this affects your menu.
4. Drag *About Us* on top of *Welcome* to nest it.
5. Click *Go to Site* one more time to see how nested pages appear.

As you've just demonstrated, organizing pages in the default menu is simple, but very powerful.

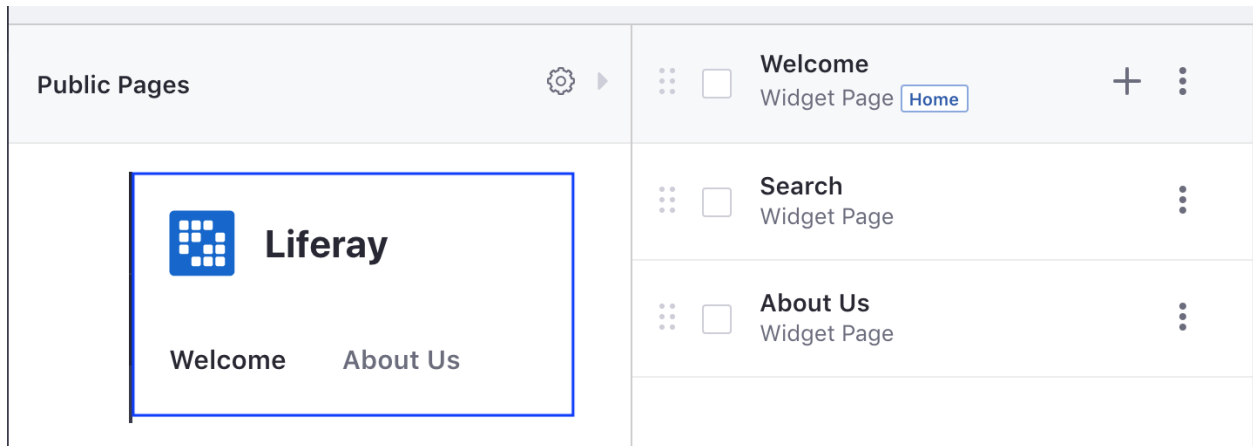


Figure 75.3: You can see the order of pages in Site Administration vs. how they appear on the site.




Figure 75.4: *About Us* is now the home page, and *Welcome* is second in the nav.




Figure 75.5: *About Us* is now nested under *Welcome* and appear when you mouse-over *Welcome*.

75.3 Public and Private Pages

As noted above, Private Pages work just like Public Pages, except they can be viewed only by registered members of a Site. In the default configuration, Public Pages are at the URL [web-address]/**web**/[site-name] while Private Pages are at [web-address]/**group**/[site-name]. Other than the membership distinctions, Public and Private Pages share the same behavior.

75.4 Page Options

While managing the default menu, you can also access page options. Clicking on the *Options* icon Option accesses several configuration tools:

- View** goes to the selected page on the Site.

- Configure** goes to page configuration.

- Copy Page** creates a new page in the current Site that duplicates the selected page.

- Permissions** opens the Permissions dialog.

- Orphan Widgets** clears data related to widgets that have been removed from the page.

- Delete** deletes the page and all its data.

Creating and managing pages using the page hierarchy is simple but very powerful. If you need more navigation options, however, Navigation Menus provide more flexibility.

CREATING AND MANAGING NAVIGATION MENUS

To better understand Navigation Menus, it's time to create a new menu.

76.1 Creating a Navigation Menu

1. Go to *Site Administration* → *Site Builder* → *Navigation Menus*.
2. Click the *Add* button (+) to add a new menu.
3. Give your menu a name and click *Save*.

On the next page appears a number of elements that you can add to a menu.

Page: Select an existing page from the current Site to add to the navigation menu.

Submenu: Create a second level of menu navigation.

URL: Create a link to any page anywhere by providing a URL. The link appears just like any other option in your menu.

Click on *Page* and you see a view of all the current pages on the current Site. Select a page and click *Add* to add that page to the menu.

Note: When you click on a page, you select that page. Multiple pages can be selected by clicking on each page one at a time. To deselect a page, click on the page again.

Now you see the menu management screen. From here, you can drag and drop menu elements to rearrange or nest them. You can also manage options for this menu by clicking the gear icon in the top right. Let's add another item to the menu.

1. Click the + icon.
2. Select *Submenu* in the menu that pops up.
3. Name your menu *External Links*.

4. Click *Add*.

Click the + button again and select *URL*. You're prompted to enter a page name and URL.

1. Enter a *Name*.
2. Enter the *URL* for an external Site.
3. Click *Add*.

Drag the URL item onto the *External Links* submenu. This nests the URL item in the submenu.

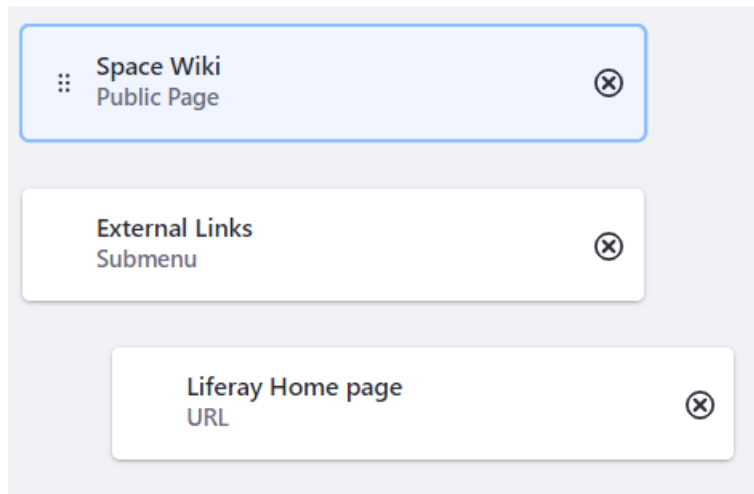


Figure 76.1: Menus can have a standard page, a submenu, and a URL link in the submenu.

Now that you can see how menus work, you can learn the details.

76.2 Managing Menus

After you create a menu, more configuration options appear on the main Navigation Menus page.

Title: Your menu's name.

Add New Pages: Determines if new pages added to the main navigation are added to the menu automatically.

Marked As: Can be set as *Primary Navigation*, *Secondary Navigation*, or *Social Navigation*.

Author: The user that created the menu.

Create Date: When the menu was created.

Theme and Fragment developers primarily use the menu types to determine how a menu should be styled. **Primary Navigation** is the main navigation for a page. **Secondary Navigation** is a second level of navigation, possibly a sidebar or a separate menu within a page. **Social Navigation** is for menus that contain links for sharing content on social media or similar tasks.

76.3 Modifying Menus

Next click on the options menu at the far right of your new navigation menu:




Figure 76.2: Menus with a standard page, a submenu, and a URL link in the submenu are created for different reasons.

Edit: Add, remove, or organize menu items.

Rename: Change the name of your menu.

Permissions: Define who can view, update, delete, and manage the permissions for the menu.

Mark As: Change the menu type.

Delete: Deletes the menu.

Next you'll learn about the tools for displaying menus on pages.

DISPLAYING NAVIGATION MENUS

You can display Navigation Menus in different ways on your Site. You may want to configure different display styles for a main menu, sidebar, and footer menu all on one page.

77.1 Navigation Menu Widget

The Navigation Menu widget lets you add navigation wherever you need it. You can place the widget on a page and then select a menu and style for the menu you are displaying.

1. Go to a Widget Page, open the Add menu on the right side of the page and add the *Content Management* → *Navigation Menu* to the page.
2. Open the *Configuration* menu.

From here you can configure three main categories:

- The Navigation Menu to be displayed
- The styling of the menu
- What level of navigation to display

77.2 Choosing a Navigation Menu

The Navigation Menu Widget has two ways to select a menu. You can choose to *Select Navigation* or *Choose Menu*.

NAVIGATION MENU

Select Navigation

Primary Navigation

Choose Menu

Select

Display Template

Bar minimally styled [Manage Templates](#)

MENU ITEMS TO SHOW

Start with Menu Items In

Level 0

Sublevels to Display **Expand Sublevels**

Unlimited Auto

Preview

Welcome Nav Menu Demo Page

Figure 77.1: Configuring the Navigation Menu Widget.

Select Navigation: Select from the three main menu types: *Primary Navigation*, *Secondary Navigation*, and *Social Navigation*.

Choose Menu: Choose any menu that was created for that Site. Once you select a menu, you must choose how to display it.

77.3 Display Template

The *Display Template* option lets you select an Widget Template for Navigation Menus. There ten included by default:

List Menu: Displays all the items in a vertical list.

Pills Horizontal: Displays the items horizontally and uses a button style for highlighting.

Pills Justified: Like Pills Horizontal, but pads the items to fill out the horizontal space.

Pills Stacked: A vertical version of the pills style.

Tabs: Displays the items like navigation tabs.

Tabs Justified: Navigation tabs that fill horizontal space.

Bar Minimally Styled: A lightweight version of the default display that you see in the embedded menu on your page.

Bar Minimally Justified Styled: Like Bar Minimally Styled with horizontal padding.

Bar Default Styled: The default embedded menu.

Split Button Dropdowns: Displays each item as a button with a dropdown for multiple navigation levels.

You can also add your own custom templates.

77.4 Menu Items to Show

Menu Items to Show configures which pages at what level from the menu are displayed in the widget. You can choose the starting level, how many levels deep to display, and how to display sub-levels.

Start with Menu Items In: Choose to start at a specific level of the navigation or a level relative to the current level (above or below).

Sublevels to Display: Select the number of levels to display in the navigation, from **1** down to **Infinite**.

Expand Sublevels: Choose if hovering your mouse over the navigation reveals navigation levels one at a time automatically or reveal all the levels at once.

Now you can see how there are a variety of customizations and configurations available for navigation menus that you can implement for your Site.

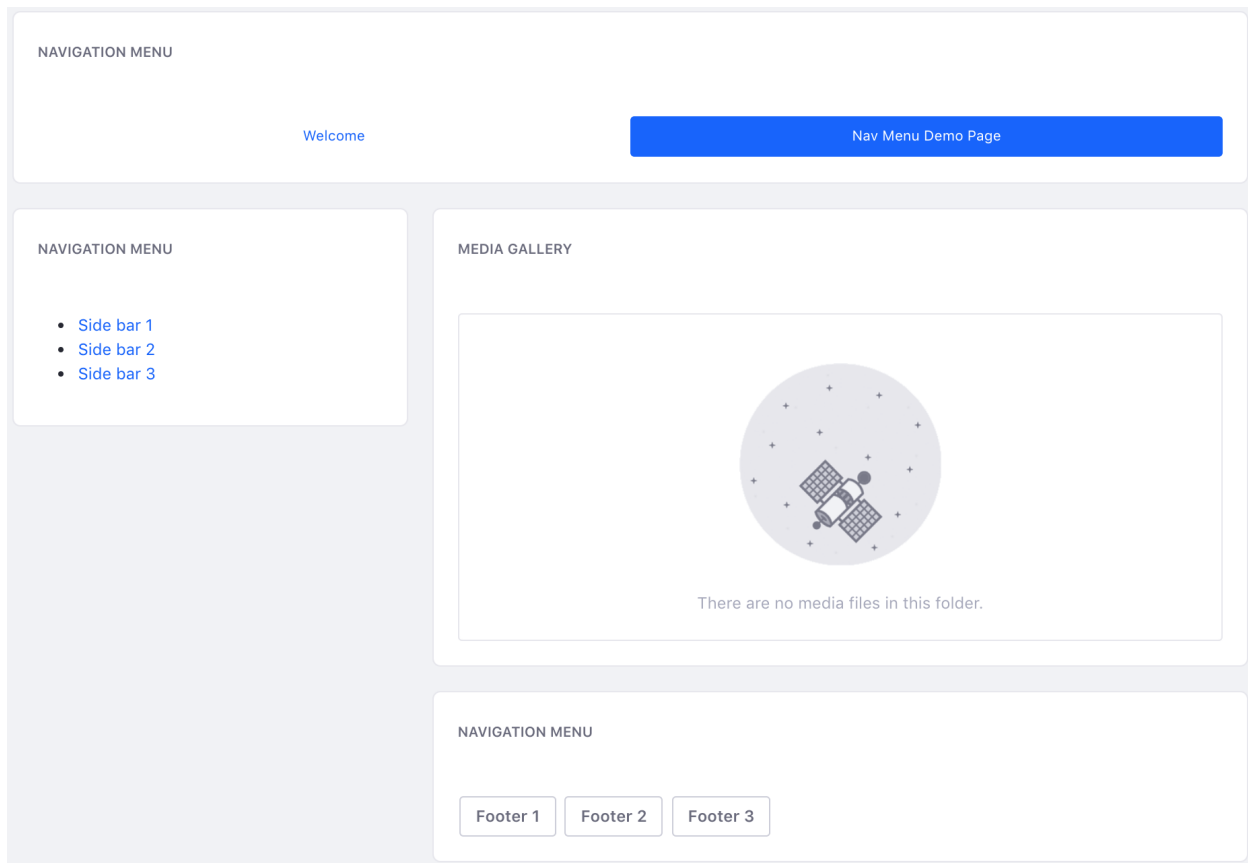


Figure 77.2: Navigation menus give you many ways to help users navigate your Site.


BUILDING SITES FROM TEMPLATES

Site Templates create a single Site structure that can be used for any new Site. They are created and administered from the Control Panel. In addition to creating multiple Sites with the same design, you can also use them to manage changes across multiple Sites with propagation of changes. Site templates can contain multiple pages, each with its own theme, layout template, applications, and app configurations.

Site templates can also contain content just like actual Sites. This allows administrators to use Site Templates to create new Sites that are each created with the same default pages, applications, and content. After they've been created, these Sites and their pages can be modified by Site administrators. Using Site templates can save Site administrators a lot of work even if each Site that was created from a given Site Template ends up being very different.

CREATING A SITE TEMPLATE



Suppose you need to create the following three private Sites for the Lunar Resort’s internal use: Engineering, Marketing, and Legal. These should be accessible only to members of these respective departments. You could design each Site separately, but you can save yourself time if you create a Site template instead.

1. Go to the Control Panel and click *Sites* → *Site Templates*.
2. Click the *Add* icon () and enter the name *Department* for your template.
3. Leave the *Active* and *Allow Site administrators to modify pages associated with this Site template...* boxes checked.
4. Click *Save* to create your Site template.

The *Active* box must be checked for your template to be usable. If your template is still a work in progress, uncheck it to ensure that no one uses it until it’s ready. Checking *Allow Site administrators to modify pages associated with this Site template...* allows Site administrators to modify or remove the pages and apps that the template introduces to their Sites—if you want the templates to be completely static, you should uncheck this.

Now it’s time to edit your Site template. This example, includes four pages.

Note: This section assumes knowledge of Liferay DXP 7.2 page management. For more information on how to create and manage pages in Liferay DXP 7.2, see the [Adding Pages to Sites](#) article.

1. Click the *Options* icon () and select *Manage*.
This brings you to the *Pages* page for the Site Template. You already have a home page. Create three more pages.
2. Create a *Full Page Application* page named *Documents*.
3. Click *Options* () → *Configure* and set the *Full Page Application* to *Documents and Media*.

4. Create a page using the Global Page Template *Wiki* and name it *Wiki*.
5. Create a widget page named *Message Boards*.
6. Click *Go to Site* in the menu to the left to go to the pages you just created.
7. On the *Home* page add the Activities, Announcements, and Calendar apps.
8. On the *Message Boards* page add the Message Boards and Tag Cloud apps.

The changes you made to your Site template above are completed in real time, so there's no *Save* button.

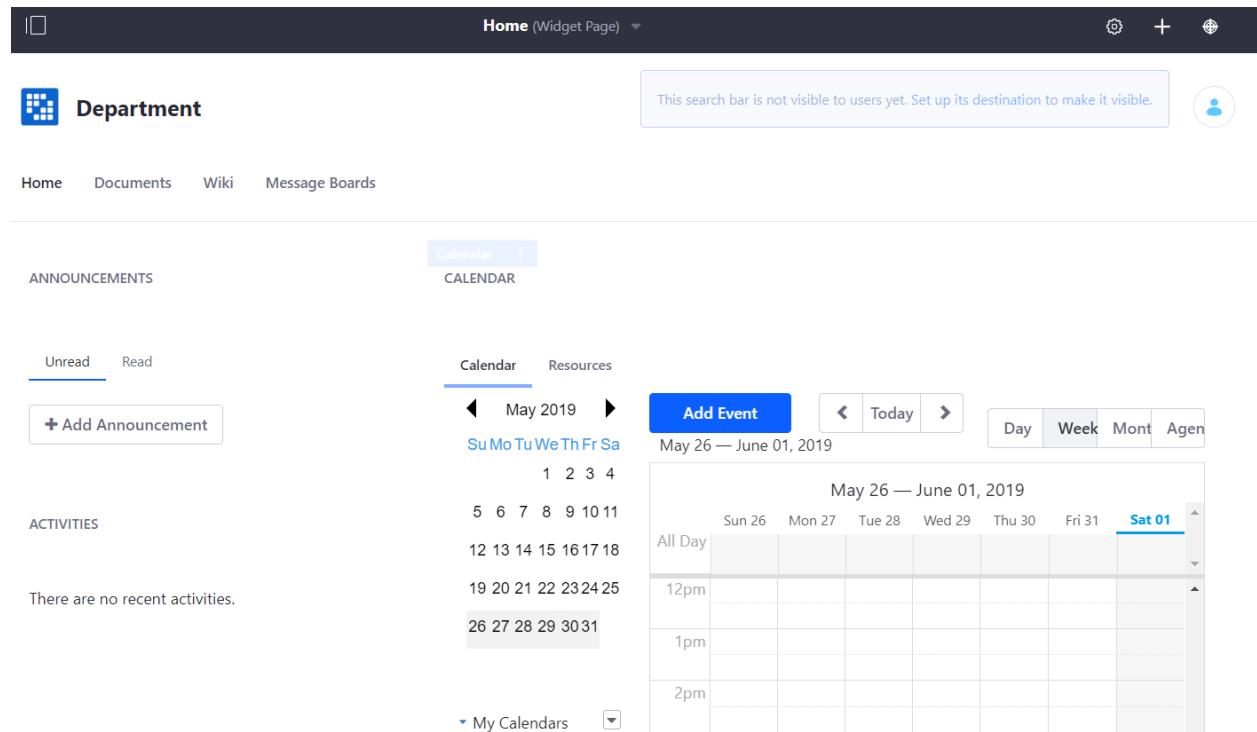


Figure 79.1: You can see the name of the Site template you're currently editing.

Next, you'll use your Site template to create the Engineering, Marketing and Legal Sites.

1. Go to the Control Panel and click on *Sites* → *Sites*.
2. Click the *Add* icon (+) → *Department*.
3. Enter *Engineering* for the Site name.
4. Check the *Create default pages as private (available only to members)*. If unchecked, they will be *public (available to anyone)* option since the Engineering Site is intended for internal use only.
5. Click *Save*.

6. In the next section, set the Membership Type to *Private*.

Recall that private Sites don't appear in the My Sites application so that regular users won't even know that the Engineering Site exists. Also, the only way users can be added to a private Site is via an invitation from a Site administrator.

7. Leave the *Active* selector enabled so that your Site can be used immediately.
8. Check the *Create default pages as private (available only to members)*. If unchecked, they will be *public (available to anyone)*. option since the Engineering Site is intended for internal use only.
9. Leave the *Enable propagation of changes from the Site template* box enabled so that the Engineering Site receives updates if the Department Site template is modified.
10. Click *Save* to create your Engineering Site.
11. Repeat these steps to create the Marketing and Legal Sites.


The new Sites have all the pages and apps you created in the Site template. To view the pages of the new Sites, click *Sites* → *Sites* in the Control Panel and then click *Actions* → *Go to Private Pages* next to one of your new Sites.

Using Site templates streamlines the Site creation process for administrators, making it easy to create Sites quickly. Now each department of the Lunar Resort has its own Calendar, Documents and Media Library, Wiki, and Message Boards on their Sites. Although the pages and apps of each department's Site are the same, each Site will quickly be filled with department-specific information as users add and share content within the Sites. Also, Site administrators can add new pages, apps, and content to their Sites, further differentiating each department's Site from the others.

MANAGING SITE TEMPLATES

To get started, click on *Site Templates* in the Sites section of the Control Panel. Here, you can add, manage, or delete Site templates. You can also configure the permissions of Site templates. As long as a Site is linked to the Site template it was created from, changes to the Site template's pages, apps, and app configurations are propagated to the Site. Changes to a Site template's content, however, are not propagated to existing Sites that are linked to the Site template. You'll learn about the propagation of changes between Site templates and Sites in more detail in the section on Site template use cases below.

To manage a Site Template's pages,

1. Click on *Site Templates* in the Control Panel.
2. Select the *Actions* icon () and then *Manage* for an existing template.

If you open the main Menu on the left side of your screen (if necessary), the Site Template is selected in the Site Administration dropdown menu. You're provided similar options as a regular Site, including *Build*, *Content*, *Configuration*, and *Publishing*. By default, the Manage Interface opens *Build* → *Pages*. From here, you can add or remove pages from a Site Template or select themes and layout templates to apply to the Site Template. You can also configure each page to have any theme, any layout template, and any number of applications, just like a page of a regular Site. As with Site pages, you can organize a Site Template's pages into hierarchies. When you create a Site using a Site template, the configuration of pages and apps is copied from the template to the Site. By default, all changes made to the Site template are automatically copied to Sites based on that template.

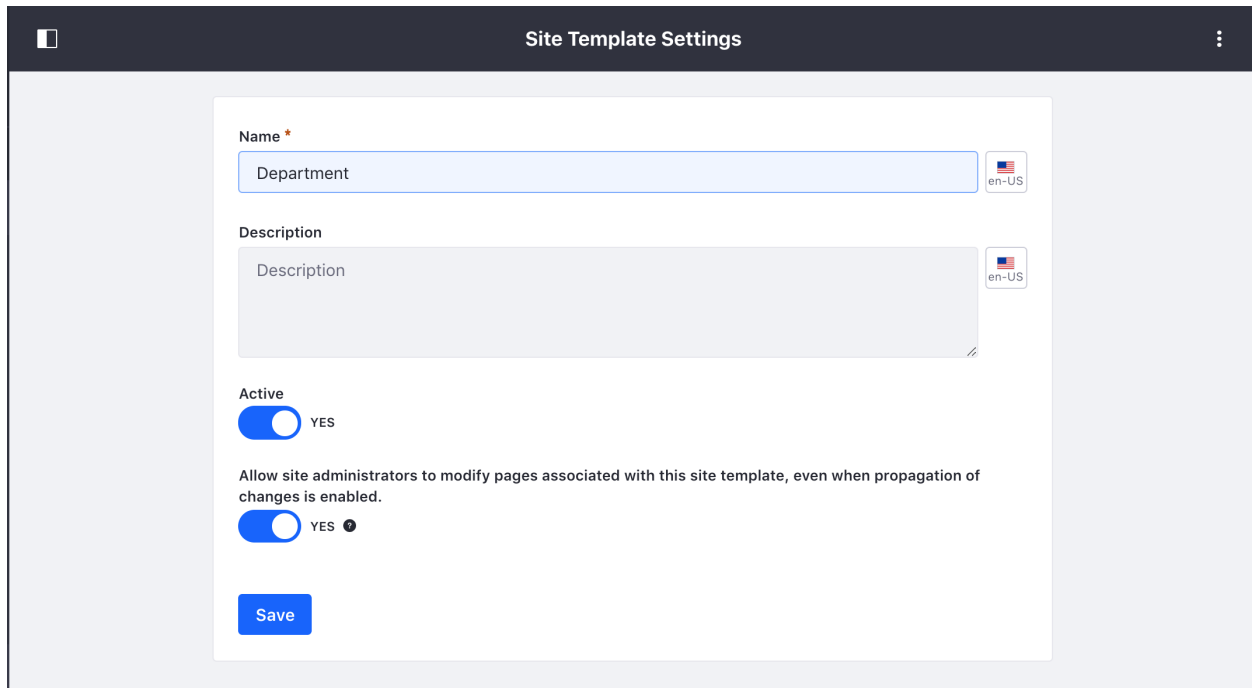
Tip: If you want to publish a piece of web content to many Sites and ensure modifications are applied to all, don't use Site template content for that purpose. Instead, place the content in the global scope and then reference it from a *Web Content Display* application in each Site.

The Content section offers separate repositories for content related apps based on your Site Template. For instance, by clicking *Polls* from the Content section, you can create a poll question that is only available for that specific Site template. Assets created within your template's Content section can only be accessed by Sites using the template.

The Configuration section includes Widget Templates and Mobile Device configuration options for your Site Template. Also, nested in the Configuration section is the *Site Template Settings*. This

edits the template's name and description while also offering boolean options for activating your Site template and allowing Site administrators to modify pages associated with your template.

The following figure displays the form shown when editing the *Department* template's settings:



The screenshot shows a web interface titled "Site Template Settings". The form contains the following elements:

- Name ***: A text input field containing "Department" and a language selector icon (en-US).
- Description**: A text area containing "Description" and a language selector icon (en-US).
- Active**: A toggle switch set to "YES".
- Allow site administrators to modify pages associated with this site template, even when propagation of changes is enabled.**: A toggle switch set to "YES" with a help icon.
- Save**: A blue button at the bottom.

Figure 80.1: Site templates have several configurable options including the option to allow Site administrators to modify pages associated with the Site template.

By default, the following Site templates are provided:

- **Intranet Site:** Provides a preconfigured Site for an intranet. The Home page displays the activities of the members of the Site, search, a language selector, and a list of the recent content created in the intranet. It also provides two additional pages for Documents and Media and external News obtained through public feeds.
- **Community Site:** Provides a preconfigured Site for building online communities. The Home page of a *community Site* provides message boards, search, a display of a poll and statistics of the activity of community members. The Site will also be created with a page for a wiki.

Now that you know the basics for creating and managing your Site templates, you can learn about propagating changes next.

PROPAGATING CHANGES FROM SITE TEMPLATES TO SITES

Site Template administrators can add, update, or delete Site Template pages. Changes made to a Site Template can be propagated to Sites whose page sets are linked to the Site Template. When you create a Site based on a Site Template with the *Enable propagation of changes from the Site template* box checked this link is created. To configure propagation of changes:


1. Select the Site from the Sites dropdown in the Menu by selecting the *Site Selector* button (🔍).
2. Navigate to the *Configuration* → *Settings* page and uncheck or recheck the *Enable propagation of changes from the Site template* checkbox.

In this section, you'll learn about the propagation of changes from Site templates to Sites and discuss the options available to Site administrators and Site template administrators.

81.1 Site Template Page Behavior

If a Site's page set has been created from a Site template and the propagation of changes from the Site template is enabled, Site administrators can add new pages but cannot remove or reorder the pages imported from the Site Template. If a Site has both pages imported from a Site template and custom Site pages, the Site Template pages always appear first in the Site page hierarchy; custom pages added by Site administrators appear after the Site template pages. Only Site template administrators can remove, reorder, or add Site template pages. Site administrators can add or remove custom Site pages. They can also reorder custom Site pages as long as they're all positioned after the Site template pages. Site template administrators cannot add, remove, or reorder custom Site pages.

Note: Pages containing a fragment (e.g., Content Pages) cannot propagate changes after a Site is *first* created based on a site template.

If a Site administrator changes a page that was imported from a Site Template and refreshes the page, the following Information icon () appears in the Control Menu with the following message:

This page has been changed since the last update from the Site template. No further updates from the Site template will be applied.

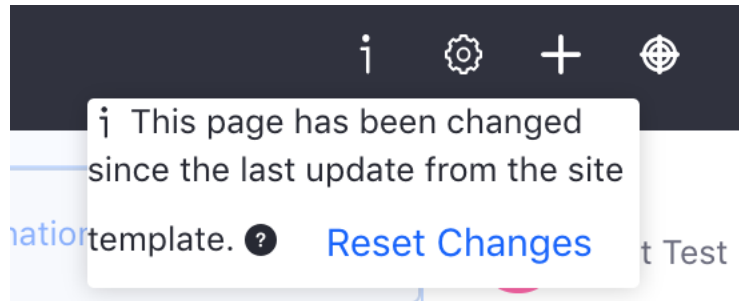


Figure 81.1: You can click the Information icon to view important information about your Site template.

81.2 Merging and Resetting Changes

If the Site administrator clicks the *Reset Changes* button, changes are propagated from the Site template page to the corresponding Site page that was imported from the Site template. Clicking the *Reset Changes* button makes two kinds of updates to a page. First, changes made by Site administrators to the Site page are undone. Second, changes made by Site template administrators to the Site template page are applied to the Site page. Note: clicking the *Reset Changes* button only resets one page. If multiple Site pages have been modified and you'd like to re-apply the Site template pages to them, you'll need to click the *Reset Changes* button for each page.

Site template administrators can set preferences for apps on Site template pages. When a Liferay administrator creates a Site from a Site template, the app preferences are copied from the Site template's apps, overriding any default app preferences. When merging Site template and Site changes (e.g., when resetting), app preferences are copied from Site template apps to Site apps. Only global app preferences or local app preferences which don't refer to IDs are overwritten.

In some cases, merging Site template and Site changes fails. For example, if pages from a Site template cannot be propagated because their friendly URLs are in conflict, Liferay DXP could try to continuously merge the Site changes. Instead of entering into an infinite loop of merge fails, Liferay DXP stops the merge after several unsuccessful attempts. Liferay DXP, however, doesn't stop there: your merge is temporarily paused, you're given an indication of the current merge fail, and then you have the opportunity to fix your merge conflicts. After you've squared away your conflict, navigate to your Site's *Site Administration* → *Configuration* → *Site Settings* and click the *Reset and Propagate* button.

The *Reset and Propagate* button resets the merge fail count and attempts to propagate your Site changes again. This process gives you the opportunity to detect and fix a merge fail when problems

Enable propagation of changes from the site template *Department*.



Some pages from the site template cannot be propagated because their friendly URLs conflict with the following pages. Modify the friendly URL of the pages to allow their propagation from the site template.

- [About Us](#)

Figure 81.2: This type of warning is given when there are friendly URL conflicts with Site template pages.

arise. This helpful process can also be done with page template merges, which follows similar steps.

Site administrators can also add data to Site template applications. For example, Site template administrators can add the Wiki app to a Site template page and use the Wiki to create lots of articles. When a Liferay administrator creates a Site from a Site template, data is copied from the Site template's apps to the Site's apps. The preferences of the Site's apps are updated with the IDs of the copied data. For example, if a Site is created from a Site template that has a Wiki app with lots of wiki articles, the wiki articles are copied from the Site template's scope to the Site's scope and the Site's Wiki app is updated with the IDs of the copied wiki articles.

Important: App data, fragment-based pages, related resources, and permissions on resources are only copied from a Site template to a Site when that Site is *first* created based on the template. No changes made to these entities are propagated to the Site after the Site is created. Neither are such changes propagated to a Site by the *Reset* or *Reset and Propagate* features.



For example, consider a Site template administrator who includes a Message Boards app as part of a Site template. They even create Message Board categories and configures permissions over the actions of the categories. The first time a Site is created based on the Site template, the categories (app data) and related permissions are copied to the Site. If the Site template administrator adds, removes, or deletes some categories, however, such changes *aren't* propagated to the Site.

Now that you've learned how Site templates work, you'll learn how to share Site templates.

SHARING SITE TEMPLATES

If you want to export a Site that uses Site or Page Templates to a different environment (through a LAR file or remote publication), the templates must be exported and imported manually in advance or the import fails.

To export a Site using a Site Template, use the following process:

1. Go to the *Control Panel* → *Sites* → *Site Templates* menu.
2. Click the *Actions* icon () and then *Manage* for the Site template you want to export.
3. Open the Site Template's management section and click on *Publishing* → *Export*.
4. Click the *Add* icon () to create a new Custom Export.
5. Select the content and pages you want to export and click *Export*.
6. Click on the *Download* icon for the template that you exported.
7. In your target environment, go to *Control Panel* → *Sites* → *Site Templates* and create a new Site template.
8. Click *Actions* → *Import* for that Site template and upload the LAR file containing your Site template's content.

Now the template can be used normally in the new environment. For more information on exporting/importing content, visit the [Importing/Exporting Pages and Content](#) article.

CONFIGURING SITES

Just like there's more than one way to cook an egg — or eat an egg for that matter — there's more than one way to build a site. Liferay is created with that in mind, and beyond just creating content and pages, Liferay has a wealth of configuration options and tools available to help you create the site that meets your needs and the needs of your users.

In this section, you'll explore all of Liferay's options for configuring sites, and gain a deeper understanding of why you might want to use various options and configurations.

CONFIGURING SITE SETTINGS

General settings range from core configuration, like your Site's Membership Type, to finer details like Documents and Media indexing options.

84.1 Details

Details provides the same menu you filled out when first creating your Site. This allows an administrator to change the description and membership type of a Site.

84.2 Membership Options

The membership type can be set as open, restricted, or private based on the privacy needs of the Site. Users can join and leave an open Site at will. To join a restricted Site, users must be added by the Site administrator, but they can request membership through the Sites section of the Control Panel. A private Site works like a restricted Site but is hidden from users who aren't members.

84.3 Site Hierarchies

Sites can be organized into hierarchies. At the bottom of the Details sub-section is the Parent Site section. When you select the parent Site for the Site you're currently on, a checkbox appears for limiting membership to members of the parent Site.

84.4 Pages

Under Pages you can view your Site's Public or Private Pages, if any exist. If they don't exist, a *Site Templates* selector appears for creating pages with a Site Template.

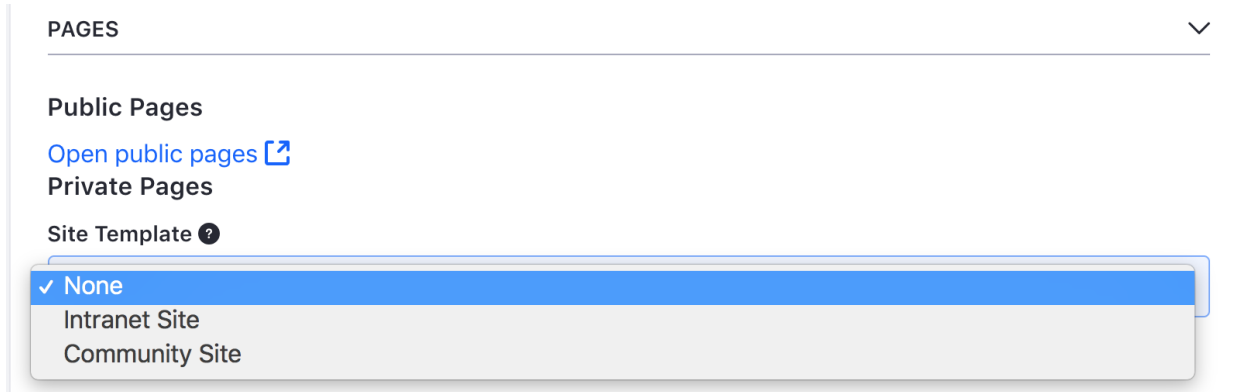


Figure 84.1: Selecting a Site Template.

84.5 Categorization

Categorization helps administrators organize the Site and allows for users to easily find your Site and its content through search and navigation. For more information on using tags and categories, visit the Organizing Content with Tags and Categories section.

84.6 Site URL

The *Friendly URL* option lets you set your Site's URL paths. Friendly URLs are used for both public and private pages. The public Site base URL is `https://localhost:8080/web`, and the private one is `https://localhost:8080/group`. Each friendly URL must be unique.

For example, setting the friendly URL of your default Site to `/lunar-resort` makes your Site's public home page's URL `https://localhost:8080/web/lunar-resort/home`. The private Site's URL is thus `https://localhost:8080/group/lunar-resort/home`.

Note that if you add a friendly URL for your instance's home page, you should update your instance's Home URL field so that page requests to `http://localhost:8080` redirect properly:

1. Go to *Configuration* → *Instance Settings* → *Instance Configuration* → *General* in the Control Panel.
2. Under *Navigation*, in the *Home URL* field enter your home URL (i.e. `/web/lunar-resort/home`).

Once you've entered this setting, page requests to `localhost:8080` redirect to the friendly URL of your Liferay instance's new home page.

You can also configure Virtual Hosts, which connects a domain name to a Site, under *Site URL*. You can use this to define a domain name (i.e., `www.lunar-resort.com`) for your Site. This can be a full domain or a subdomain. You can use this to host a number of web sites as separate Sites on one Liferay server.

Public Pages

`www.lunar-resort.com`

Private Pages

`private.lunar-resort.com`

Figure 84.2: When configuring virtual hosts, the public and private pages of a site can be configured to different domains.

For instance, if you set this up for the Lunar Resort's development network, users in that Site would access `developers.lunar-resort.com`, provided that the Lunar Resort instance's network administrators created the domain name and pointed it to the Liferay server.

1. With your provider, set the DNS name `developers.lunar-resort.com` to point to your Liferay instance's IP address.
2. In the Virtual Host tab for the Developers Site, set the URL to `http://developers.lunar-resort.com`

This helps users quickly access their Site without having to recall an extended URL. The *Site URL* option is listed under the General tab.

84.7 Documents and Media

Below that is *Documents and Media*, which lets you enable/disable Directory Indexing. If on, Site administrators can browse your Site's Documents and Media files and folders. For example, a Site administrator of a Site called *Lunar Resort* can browse documents at `http://localhost:8080/documents/lunar-resort` if this option is enabled.

84.8 Site Template

If you created your Site using a Site Template, this section appears and displays information about the link between the Site Template and the Site. Specifically, you can see which Site Template was used and whether or not it allows modifications to the pages inherited from it by Site administrators. To learn more about Site Templates and how to create your own, see [Building Sites from Templates](#).

84.9 Asset Auto Tagging

Asset Auto Tagging lets you enable or disable the use of any Asset Auto Tagging rules on your site. See [Asset Auto Tagging](#) to learn more about setting up auto tagging features.

84.10 Custom Fields

Custom Fields only appears if you've created them in Control Panel → *Configuration* → *Custom Fields*. For more information on Custom Fields, see [Custom Fields](#).

SOCIAL SETTINGS AND LANGUAGES

The Social tab provides options for managing the social interactions on your Site. Languages lets you configure language options and change the default language options for the Site.

85.1 Ratings

The *Ratings* option lets you select the ratings type to use for applications like Documents and Media, Web Content, Comments, etc. Ratings types include Stars, Likes, and Thumbs.

85.2 Mentions

At the bottom of the page is *Mentions*, which lets you enable/disable Mentioning functionality, which is used to *mention* (notify and/or draw attention to) friends and colleagues by entering the “@” character followed by their user names. See the Mentioning Users article for more information.

85.3 Languages

The *Languages* tab lets you configure the language options for your Site.

You can use the default language or define another supported language as the default for your Site.




Figure 85.1: In the Languages tab, you can configure the site to use the instance’s default language or another supported language.

ADVANCED SITE SETTINGS

Advanced Settings relate to security (like User Roles) or require external configuration (like creating a Google Analytics account) to use.

86.1 Default User Associations

Default User Associations configures Site roles and teams that newly assigned Site members have by default. If you'd like to learn more about creating roles and/or teams, visit the Roles and Permissions and Creating Teams for Advanced Site Membership Management. sections.

86.2 Analytics

Liferay DXP includes built-in support for Google Analytics for analyzing traffic on your Site. Google Analytics provides a snippet of code which you add to your pages enable tracking. Adding this code to every page on a Site would be tedious, especially if it's a large Site with a lot of user-generated content.

There are two ways to mitigate this problem:

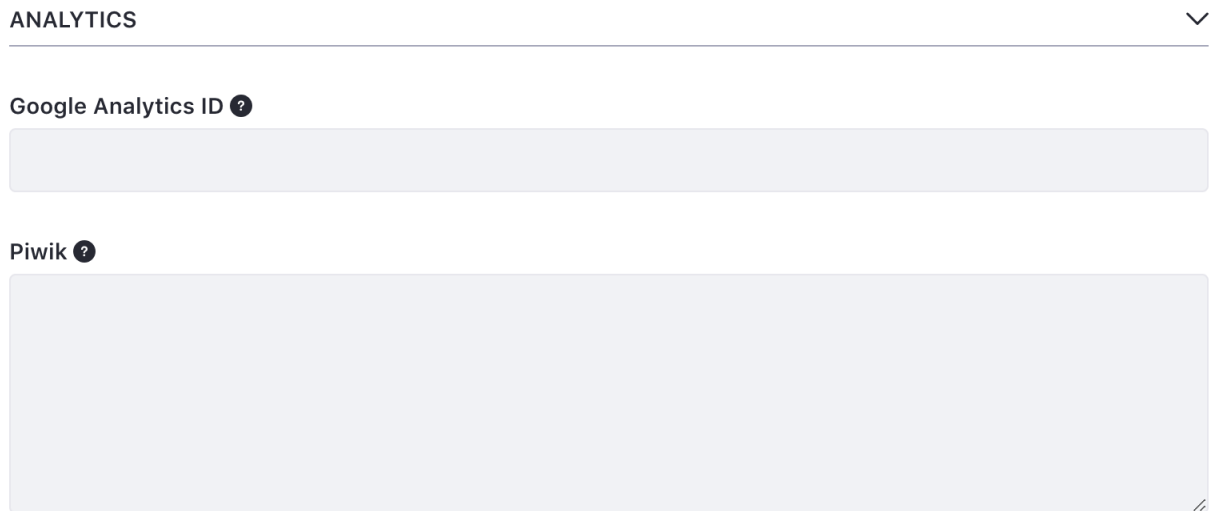
1. A web developer can hard-code the tracking code into a theme, which embeds it on every page.
2. An administrator can enter the tracking code in Site settings.

To use option #2:

1. Go to *Configuration* → *Settings* → *Advanced*.

2. Expand the *Analytics* section.
3. Enter your Google Analytics ID.
4. Click *Save*.

All the pages in the Site you selected now have the Google Analytics code and can be tracked.



ANALYTICS

Google Analytics ID ?

Piwik ?

Figure 86.1: To set up Google Analytics: sign up, receive an ID, and then enter it into the Google Analytics ID field.

To enable a different analytics service:

1. Go to *Configuration* in the Control Panel.
2. Go to *Instance Settings* → *Analytics*.
3. Enter the name of any additional service you want to add in the *Analytics* field provided.
4. Once you have entered the name, go to the *Site Settings* → *Advanced* → *Analytics* page for the Site where you wish to add analytics.
5. Copy the JavaScript tracking code provided by your analytics platform into the corresponding field for your service.

Now all pages on the selected Site contain the tracking script and send analytics data to your analytics platform.

86.3 Maps

The *Maps* option configures the maps API provider used by your Liferay instance when displaying geolocalized assets. Geolocalized assets can be displayed for documents, web content articles, DDL records, etc. Maps is available under the Advanced tab. You can read more about Geolocation in Geolocating Assets.

86.4 Recycle Bin

The *Recycle Bin* option enables or disables the Recycle Bin for your Site. You can also regulate the age (in minutes) for which content is able to be stored in the Recycle Bin until it is permanently deleted. For a full explanation of the Recycle Bin, see Restoring Deleted Assets.

86.5 Content Sharing

If you select the *Content Sharing* tab from the Advanced tab, you can configure whether sub-Sites can display content from this Site. Administrators of this Site's sub-Sites can use all structures, templates, categories, widget templates, and more from this parent Site. Even if you initially allowed content sharing between the parent Site and its sub-Sites, you can disable this option and immediately revoke content sharing from all sub-Sites.

You can manage this globally by navigating to the Control Panel → *Configuration* → *Instance Settings* → *Content & Data* → *Sharing* → *Content Sharing*. First, you can choose if Site administrators can display content in Sites from other Sites they administer. For example, suppose that a certain User is a Site administrator of two Sites: *Engineering* and *Marketing*. The checkbox in the Content Sharing section of Instance Settings determines if the Site administrator can display content from the Marketing Site in the Engineering Site and vice versa.

You can also choose if child Sites can display content from parent Sites and configure the defaults. There are three options:

Enabled by Default: Child Sites can display content from parent Sites by default, but this can be disabled by a Site administrator.

Disabled by Default: Child Sites cannot display content from parent Sites by default, but this can be enabled by a Site administrator.

Disabled: Child Sites cannot display content from parent Sites, and this behavior cannot be changed by a Site administrator.

That covers your Site's advanced settings. You're now equipped to manage all aspects of your Site's configuration.

CUSTOMIZING PERSONAL SITES

By default, newly created users are granted a personal Site.

- Users function as Site administrators of their personal Sites.
- Personal Sites are fully customizable but cannot have more than one member.
- Users can have publicly available content on their Site's Public Pages. This is often used for a user blog.
- Users can also have Private Pages where they can keep personal information or use Documents and Media to have their own private file repositories.

You can disable personal Sites by adding the following properties to your `portal-ext.properties` file:

```
layout.user.public.layouts.enabled=false  
layout.user.private.layouts.enabled=false
```

Note: The public and private page sets of personal Sites are handled separately. You can leave one page set enabled while disabling the other.

If you initially had user personal Sites enabled for your instance but then disabled them, existing personal Sites remain on your Liferay instance until the next time users log in, at which point they're removed.

You can allow users to create personal Sites but not have them automatically created for new users. To do this, add the following properties to your `portal-ext.properties` file:

```
layout.user.public.layouts.auto.create=false  
layout.user.private.layouts.auto.create=false
```

If the properties `layout.user.public.layouts.enabled`, `layout.user.private.layouts.enabled`, `layout.user.public.layouts.auto.create`, and `layout.user.private.layouts.auto.create` are all set to true, which is the default, users have personal Sites and public and private pages are created automatically for new users.

There are a number of portal properties you can use to customize the automatically created pages. You can customize the names of the default pages, the applications that appear on the pages, the themes and layout templates of the default pages, and more. Please refer to the Default User Public Layouts and Default User Private Layouts sections of the `portal.properties` file for details.

Note: By default, users can modify the pages and applications on their personal Sites. Administrators, however, can customize the modifiable portions of personal Sites through Liferay DXP's permissions system by removing permissions from Roles. To disallow all Liferay users from modifying something, remove the relevant permission from the User Role.

Great! Now you know how to customize a personal site!

IMPORTING/EXPORTING SITES AND CONTENT

Export/Import lets you backup and restore your Site and app data as a LAR (Liferay Archive). There are two primary places Export/Import is used: Sites and apps. You can learn more about exporting/importing app data in the Exporting/Importing Widget Content section. In this section, you'll learn how to export and import content for Sites.

88.1 Backing Up and Restoring Pages and Their Content

In *Site Administration* → *Publishing*, you can find the *Export* and *Import* option for pages. If you click on *Export*, you see an interface for exporting your public or private pages. The *Export* feature exports your Site's data as a single LAR file. Similarly, *Import* is a similar interface for importing public or private pages from a LAR file.

When importing data into a Site, you should use a newly created Site to avoid conflicts between the existing data and the data being imported. When exporting Site data, you can specify exactly what data should be included in the LAR:

- Site pages (you can select exactly which ones)
- Page settings
- Theme
- Theme settings
- Logo
- Application configurations
- Application content
- Archived setups
- User preferences

A LAR file can be imported into a Site on another Liferay server. You can take content from a Site in one environment (say, a development or QA environment) and move it all to a Site on another server with LARs. You can use LARs to import data onto production servers, but you should

not make this a regular occurrence. If you want to regularly move pages from one server to another, you should use Liferay DXP's staging environment. See the Staging section for more details.


You can export LARs to use them as a backup. If you ever have to restore your Site, you must only import the latest LAR file. However, please be careful! If there's content that exists both in the LAR and in the Site that's importing the data, there may be a conflict, and data could be corrupted. If you want to restore a Liferay Site using a LAR file, delete the Site entirely, create a new Site with the same name as the old one, and then import the LAR file into the new Site. This way, there's no chance for there to be a data conflict.

Some naming collisions are handled automatically. For example, a collision occurs if the LAR you're importing and the Site both have a page with the same friendly URL. Liferay DXP resolves the collision by adding a number to the end of the friendly URL and incrementing until there's no collision. Similarly, if importing a LAR into a Site causes a category name collision, the imported categories are automatically renamed.

Note: LAR files are version dependent. You can't import a LAR file that was exported from one version of Liferay into a Liferay server that's running a different version of Liferay. Also, note that periodically exporting LARs is *not* a complete backup solution; please refer to the Backing up a Liferay Installation section for information on backing up Liferay.

88.2 Page Export Example

Here's how the export process works:

1. Go the *Site Administration* → *Publishing*.
2. Click *Export*.
3. Click *Add* .

A *New Custom Export* page loads, so you can choose the pages and content you want to export from your Site.

4. Enter *Lunar Resort Version 1* for the *Title*.
5. Under *Pages*, select public or private pages and the settings you want to export.
6. Under the *Content* category, select *All*.

Note that if you select one of the *Choose* radio selectors or *Change* links, you're given checkboxes for options to choose. The applications' content can also be selected for export, including the Documents and Media Library, Message Boards, and Web Content assets. You can even export the theme you're using. Finally, you can select whether the permissions for your exported pages and content are included.

7. Click *Export*.

The screenshot shows the 'EXPORT' configuration interface, divided into three main sections: PAGES, CONTENT, and DELETIONS.

- PAGES:**
 - Pages Options:** A button labeled 'Change to Private Pages'.
 - Pages to Export:** A list of pages with checkboxes:
 - Public Pages
 - Home
 - Search
 - Look and Feel:** A list of settings with checkboxes:
 - Theme Settings
 - Logo
 - Site Pages Settings
 - Site Template Settings
- CONTENT:**
 - Radio buttons for selection: All, Date Range, Last...
 - A link for 'Refresh Counts' with a refresh icon.
- DELETIONS:**
 - Export Individual Deletions: If this option is chosen, the delete operations performed for content types will be exported to the LAR file.

Figure 88.1: You can configure your export options manually by selecting pages, content, and permissions.

Once you click *Export*, the menu automatically switches to the *Processes* tab, where you see the status of your exported LAR file. You can select the *Download* icon (↓) to download the export to your local machine. Once you have the file, you can copy it to a backup location for safekeeping or import it into another installation of Liferay. If you must rebuild or wish to revert back to this version of your Site, you can import this file by clicking the *Import* button from the Publishing menu, browsing to it, and selecting it. You can also drag a LAR file inside the dotted area, which also executes the import process.

88.3 Export Templates


Instead of manually customizing an export process every time you export pages/content, you can use an Export Template. This provides you the convenience of storing export process settings so they can be reused. If you export pages frequently and usually select the same options to export, you can create an export template to export with your standard options.

To create an export template,

1. Select the *Options* icon (ⓘ) from the top right corner of the screen and select *Export Templates*.
2. Click the *Add* button (+).
3. Assign the template a *Name* and *Description*.

4. Fill out the configuration options for your export process.
5. Click *Save*.

Your template is now available to use from the *Export Templates* menu. To use the template,

1. Click the *Actions* () next to the template.
2. Select *Export*.
This automatically fills the fields and options for exporting pages and their content.
3. Give the export a name.
4. Click *Export* and your LAR file is generated.

STYLING APPS AND ASSETS


Widget Template define custom display templates used to render widgets. For example, you may want to show blog entries horizontally instead of vertically, or list your assets in the asset publisher application in different sizes. In this section, you'll learn about the capabilities of widget templates and how to create and configure them.

STYLING WIDGETS WITH WIDGET TEMPLATES


Suppose you're customizing the Lunar Resort Site and want to allow users to use Facebook or Twitter to communicate with other interested travelers. You can add this functionality to an existing widget using widget templates: launch a template editor, create a custom template, and configure your app to host that template. Widget templates let you re-skin your widget and give you ultimate control over its appearance and functionality.

90.1 Creating a Widget Template

Here's the process of creating a widget template:

1. From Site Administration, click the *Site Selector* button () to choose the Site where you want to create the widget template.
2. Open *Site Builder* → *Widget Templates*.

If you selected the Global context, this page shows a list of sample templates available for your apps. These sample templates differ from the default templates already configured in the apps. If you choose a Site to host your template, you must create a custom template for that Site's apps.

3. Click the *Add* () button, and you're prompted to select the type of template to create.
 - *Asset Publisher*
 - *Blogs*
 - *Breadcrumb*
 - *Categories Navigation*
 - *Language Selector*
 - *Media Gallery*
 - *Navigation Menu*
 - *RSS Publisher*




Figure 90.1: The Site Administration dropdown menu lets you choose the context in which your widget template resides.

- *Site Map*
 - *Tags Navigation*
 - *Wiki*
4. Enter the name and, optionally, open *Details* to provide a description and a small image to use. You can select the language type for your template.
 5. Within *Details* select a scripting language to use. You can use FreeMarker or Velocity. FreeMarker is recommended.
 6. Use the *Script* section to create the widget template.
 7. Click *Save* when done.

90.2 The Template Editor

On the left side of the template editor is a palette of common variables used for making templates. This is a great reference when creating your template. To place one of the variables into the template editor, position your text cursor where you want it placed, and click the variable name.

Each variable also has a tooltip which displays a detailed description. Because there are multiple kinds of widget templates, there are also different variables for each widget template. Thus, each template has a different set of variables only applicable for that specific template.

You can also use the autocomplete feature to add variables to your template. It can be invoked by typing `#{` which opens a drop-down menu of available variables. By clicking one of the variables, the editor inserts the variable into the editor.

You can also embed same-type templates into other templates. For example, suppose you have an existing Wiki widget template and would like to create another similar Wiki widget template.

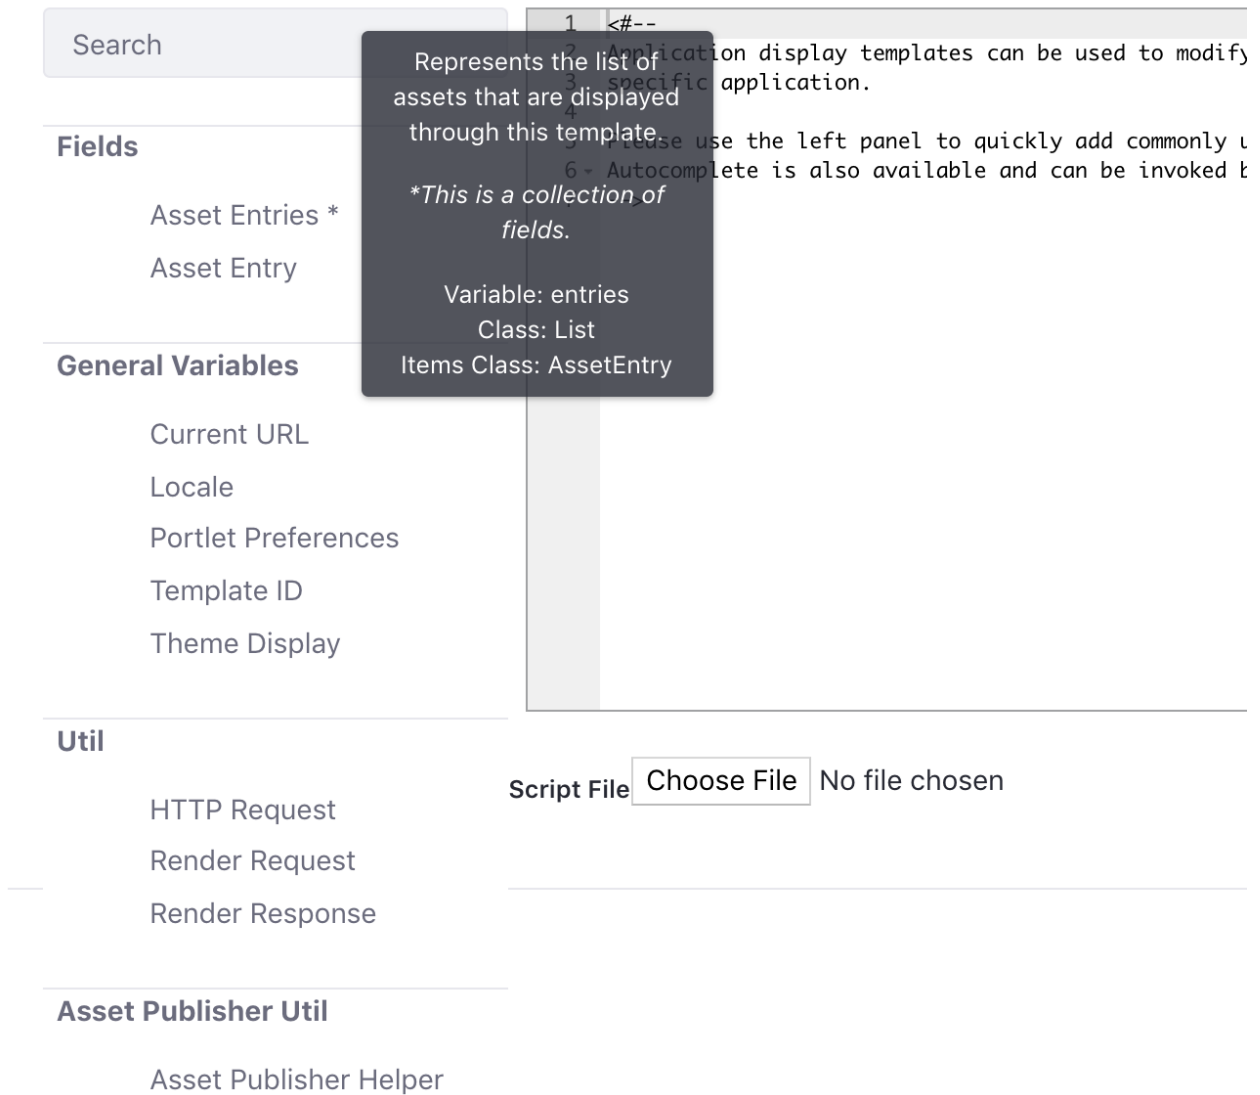



Figure 90.2: Liferay offers a versatile script editor to customize your widget template.

Instead of starting from scratch, you can import the existing Wiki widget template into your new one and build off of it. In other words, you can utilize widget templates as generic templates which allow for reusable code to be imported by Velocity or FreeMarker templates in the system.

90.3 Configuring Widget Templates

After you've saved your widget template, you can manage it through its *Actions* () button. This provides several options:

- *Edit*: lets you modify the widget template's setup properties.
- *Permissions*: lets you manage the permissions *Update*, *Permissions*, *Delete*, and *View* for the widget template.
- *Copy*: creates a copy of the widget template.
- *Delete*: deletes the widget template.

Additionally, your widget template generates a static URL and a WebDAV URL. These values access the XML source of your template. You can find these URLs by clicking the widget template from the menu and expanding the *Details* section. With the WebDAV URL, Site administrators can add, browse, edit, and delete widget templates on a remote server. If you want to learn more about what the WebDAV URL can do, visit the article on WebDAV access.

Note: Embedding widgets into widget templates, although possible, is not recommended because this could cause conflicts with other widgets or unexpected behavior (e.g., embedding a widget that aggregates data to the breadcrumb). If embedding a widget into a widget template is your only option, make sure it does not interfere with other widgets.

Next you must configure the widget to use the new widget template:

1. Go to the *Configuration* page for the widget you want to modify and open its *Display Settings*.
2. Under *Display Template*, select your widget template from the drop-down menu.

Also, you can manage Site-specific display templates for your app: do this by clicking the *Manage Display Templates for [SPECIFIC_SITE]* link next to the *Display Template* drop-down menu. A window appears with a list of your configured templates only available for your Site with options to add new templates or edit existing templates.

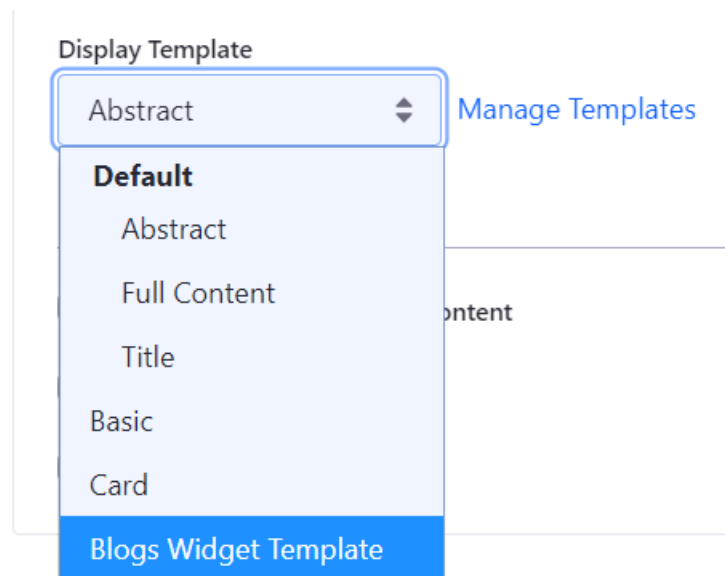


Figure 90.3: In the *Configuration* menu of an app, you can edit and manage available widget templates.

WIDGET TEMPLATE EXAMPLE

Now that you know the general functions of widget templates, you'll create your own. This brief demonstration will show you just how easy, yet powerful, widget templates can be for your Liferay instance.

1. Add the Media Gallery widget to a page by navigating to *Add* (+) → *Widgets* → *Content Management* → *Media Gallery*.
2. Click the widgets's *Add* button (+) → *Multiple Files Upload* and select two custom photos to display. Then click *Save*, and navigate back to the main application screen.
3. Notice the default format of the pictures. To change the display template for this widget, navigate to *Options* (⋮) → *Configuration*.
4. From the *Display Template* drop-down menu, select *Carousel*. Then click *Save*.

The Media Gallery application is transformed into a carousel slideshow. At this time, it's perfectly natural to be experiencing “I can conquer the world” feelings, just as Liferay's mascot, Ray, exudes in the image above. widget templates have that kind of power to transform your site into an enjoyable and convenient home for users.

Customizing the user interface of Liferay DXP's bundled widgets provides the ultimate customization experience for Liferay users.


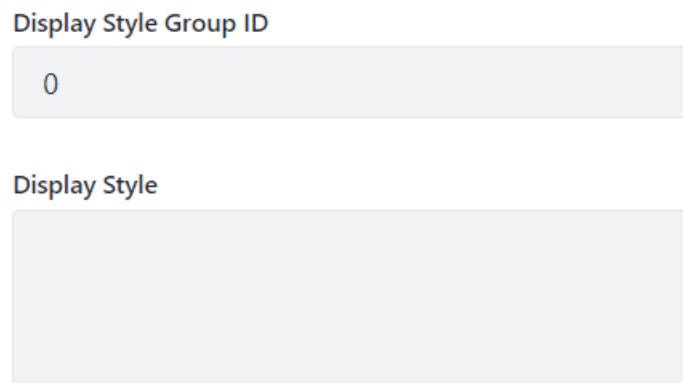


Figure 91.1: After applying the Carousel widget template, your pictures are displayed as a carousel slideshow.

SETTING A DEFAULT WIDGET TEMPLATE

You can change the widget template for an individual widget through its own configuration, but to configure the default widget template for all widgets of that type, you must go to *System Settings*. In the System Settings you can find a configuration for every widget in Liferay DXP. Any widget that supports widget templates has a *Display Style Group ID* and a *Display Style* option.



The image shows a configuration form with two sections. The first section is labeled 'Display Style Group ID' and contains a text input field with the value '0'. The second section is labeled 'Display Style' and contains an empty text input field.


Figure 92.1: The widget template configuration in System Settings lets you change the display style.

- **Display Style Group ID:** The Site ID where the widget template is located. For Global templates use *0* for the ID.
- **Display Style:** The widget template's key.

To enter a Display Style, you first need the *Template Key* for the template you want to use. To get the Template Key, go to the *Application Display Template* list for a given Site and retrieve it from the widget template listing. Then enter the display style as `ddmTemplate_[template-key]`.

92.1 Default Widget Template Example

For example, configure the Language Selector widget templates like this:

1. Open the *Product Menu*.
2. Using the *Site Selector*, select the *Global* site.
3. Go to *Site Builder* → *Widget Templates*
4. Create a *Language Selector Template*.
5. Click the *Actions* () button for the new widget template.
6. Open *Details* and find the *Template Key* - LANGUAGE-ICON-FTL

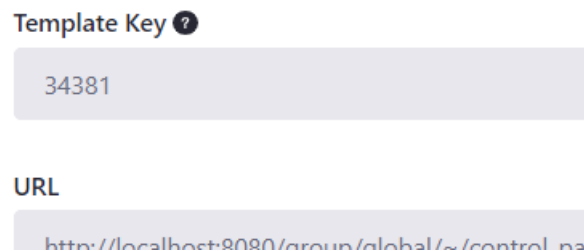


Figure 92.2: System Settings shows where you can find the Template Key.

Now that you have the ID, you can change the template from System Settings.

1. Go to the *Control Panel* → *Configuration* → *System Settings*.
2. Find *Localization* under the *Platform* heading and select *Language Selector* from the options on the left.
3. In the *Display Style* field, enter `ddmTemplate_LANGUAGE-ICON-FTL`.

Now any Language Selector widgets are added to a page use the new defaults. This doesn't affect widgets already added to a page and configured.




Figure 92.3: You can see the new default configuration.

CUSTOMIZING PAGE OPTIONS

Every page has options to give it a unique configuration. You can handle that configuration at the individual page level or configure a Page Set of public or private pages all at once.

When you configure options across Sites, Page Sets, and individual pages, there is a hierarchy that flows down. For example, you can set a theme at the Site level that applies to all pages within a Site. You could then set a different theme for the Private Pages set. You can even set a different theme for a specific page that is different than the master site configuration or the configuration for its page set.

When configuring pages, be aware of your context, so you don't configure an option that you intended for the whole site just on one page—or change the whole site's configuration, when you meant to change it for only a specific page.

CONFIGURING PAGE SETS

To configure options for the entire Page Set, select the *Configure* icon next to the Page Set in *Pages*. Options configured for the Page Set apply to all its pages. Page Set options override options set at the Site level, and customizations to an individual page override those for the Page Set.

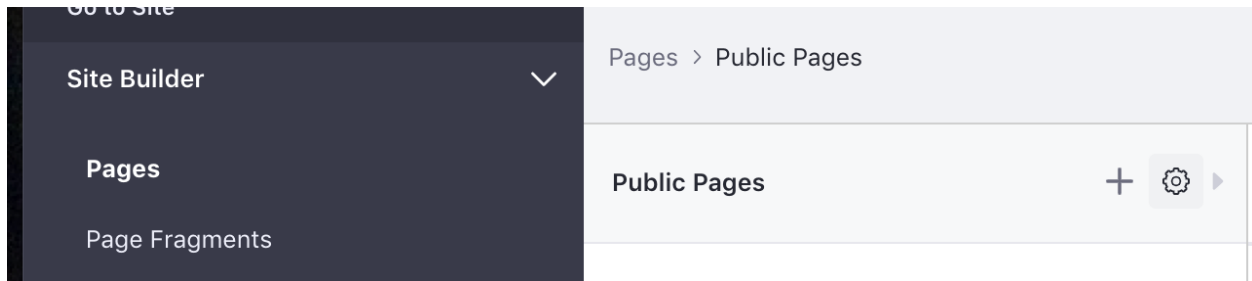


Figure 94.1: Selecting the Page Set configuration option.

CONFIGURING PAGE SETS

Page Set configuration starts with the *Look and Feel* tab. Here you have an interface for choosing a theme for the current Site.




Figure 95.1: The Look and Feel page set tab.

95.1 Themes

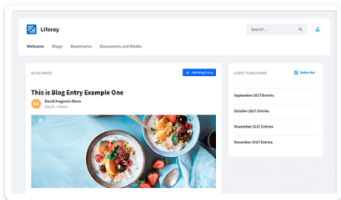
Themes can transform the entire look of your Site. They are created by developers and are easily installed using Liferay Marketplace.

You can apply themes to the entire Site or to individual pages. For the Site, go to *Pages* → the Site (public or private), and click the Gear icon. For individual pages, click *Configure* → *Define a specific*


Look and Feel Advanced

LOOK AND FEEL ▼

Current Theme



Name
Classic

Author
[Liferay, Inc.](#) 

Settings

Show Footer
 YES

Show Header
 YES

Show Header Search
 YES

Show Maximize/Minimize Application Links
 NO

Wrap Widget Page Content
 YES

look and feel for this page option under the page's *Look and Feel* category. See the Themes section for information on creating and developing your own custom themes.

Use the same look and feel of the [Public Pages](#).
 Define a specific look and feel for this page.

Current Theme

Name
Classic Theme

Author
[Liferay, Inc.](#)

Settings

Bullet Style: dots

Show Header Search: true

Show Maximize/Minimize Application Links: false

Figure 95.3: You can define a specific look and feel for a page.

Many themes include more than one color scheme, which keeps the existing look and feel while giving the Site a different flavor. The Color Schemes option is not available for the default theme.

There are a few more configurable settings for your Page Set look and feel. You can switch the bullet style between dots and arrows and you can choose whether or not to show maximize/minimize application links by default. The CSS section lets you enter custom CSS for tweaking your theme.

95.2 Using a Custom Logo for a Site

By default, the Liferay logo is used for your Site pages' logo. If you want to use your own logo for a specific Site, select the *Logo* tab from the *Configure* interface and browse to the location of your logo. Make sure your logo fits the space in the top left corner of the theme you're using for your website. If you don't, you could wind up with a Site that's difficult to navigate, as other page elements are pushed aside to make way for the logo.

In the logo tab, you can also choose whether or not to display the Site name on the Site. If you check the box labeled *Show Site Name*, the Site name appears next to the logo. This option is enabled by default and cannot be disabled if the *Allow Site Administrators to set their own logo* option is disabled in *Instance Settings*. Removing the Site name is not available for the default Site—you can configure this only for new Sites and user pages.

ADVANCED PAGE SET OPTIONS

There are some powerful options that should only be used by those with a firm command of the technology, or they could have major unintended side effects. You can find these options under the *Advanced* tab.

96.1 Executing JavaScript in Site Pages

At the top of the *Advanced* tab is a JavaScript editor. Code entered here is executed at the bottom of every page in the Site. If your Site's theme uses JavaScript (as is usually the case), it's best to add custom JavaScript code to the theme and *not* here. This way, all your Site's JavaScript code remains in one place.

This may be useful if your Site's theme does *not* use JavaScript. In this case, you can place *all* of your Site's JavaScript here.

96.2 Merge Public Pages

If you have more than one Site on a specific Liferay instance, one of those Sites will be the *Default Site* where visitors will be directed if they access the instance's root URL. By default, visitors will only see the pages of that Site in the navigation. To have another Site's public pages appear in the primary navigation for the Default Site, check the box to *Merge public pages* for that site. Be careful as adding too many pages to the main navigation can make it become unwieldy very quickly.

96.3 Rendering Pages for Mobile Devices

Mobile Device Rules lets you configure your page set to have specific behaviors for specific mobile devices or types. Mobile device rules are inherited from your Public Pages, but you can define specific rules per page. You can edit the Look and Feel of specific pages for mobile devices, including the theme. This is explained in *Mobile Device Rules*.

96.4 Robots

The *Robots* option lets you configure robots.txt rules for the domain: both its public and private pages. The robots.txt file provides instructions to search engines and other tools that are automatically crawling and indexing your Site. Common entries here include defining some pages not to be indexed.

96.5 Notifying Search Engines of Site Pages

The *Sitemap* option generates a sitemap you can send to some search engines so they can crawl your Site. It uses the industry standard sitemap protocol.

Select a search engine link to send the sitemap to it. It's only necessary to do this once per Site.

If you're interested in seeing what is sent to the search engines, select the *preview* link to see the generated XML.

CONFIGURING INDIVIDUAL PAGES

After you've configured your Page Set, you can reconfigure some options at the individual page level. When you decide to customize a single page, options that were not available when initially creating a page appear. You can customize a page by navigating to *Pages* under *Build* menu and selecting *Options* (🔧) → *Configure* next to the page you want to edit from the navigation tree. Alternatively, you can click the *Configure* icon on the top right of any page.

INDIVIDUAL PAGE SETTINGS

On the Configure page are four tabs: General, SEO, Look and Feel, and Advanced. Options selected here have no effect on the rest of the Site; just the page you've selected. Many of these options are the same as those that configure the complete page set, so you can view more details in the Configuring Page Sets article.

Note that many of the options are localizable, so you can provide translations based on the user's locale.

98.1 General

The *General* tab lets you configure the basic information and design for the page. You can change the *Name*, *Friendly URL*, and *Page Layout*.

98.2 Name and Friendly URL

The *Name* is the title that appears in the browser's title bar, and how the page is identified in the navigation. The *Friendly URL* defines the page's link. It is a best practice to have the URL match the name of the Page, so these two should generally be updated together.

98.3 Page Layout

For Widget Pages, you can select a Layout Template that defines droppable locations for widgets. Layout templates define a number of sections with columns and rows. Widgets added to a section expand (or contract) horizontally to fill the space and can be stacked vertically.

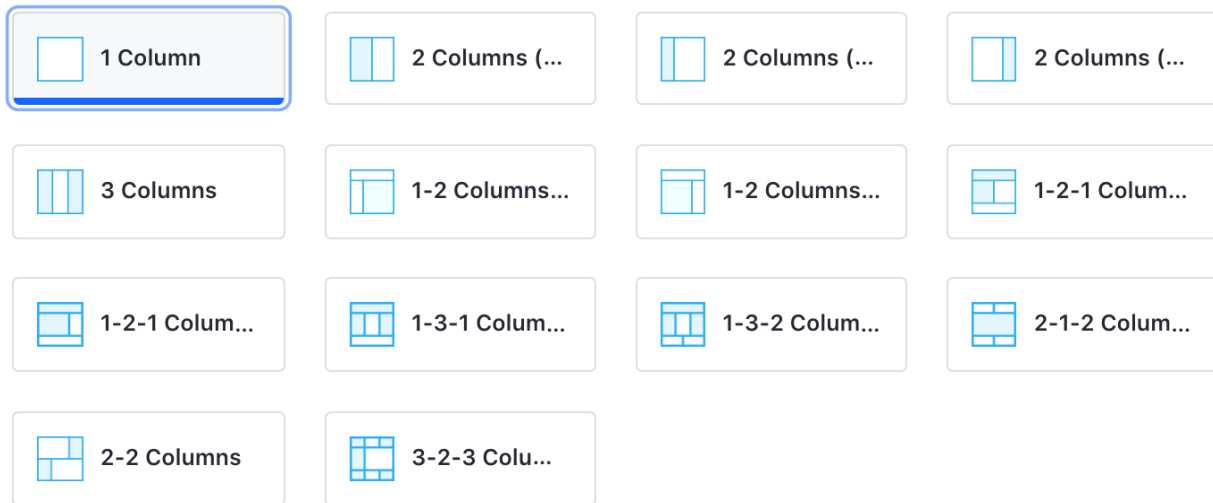


Figure 98.1: Setting a layout template for your page.

98.4 Categorization and SEO

Managing your page's content drastically improves your page's organization and user experience. The Site page's configuration options offers some opportunities to organize page content.

98.5 Categorization

The *Categorization* tab shows the categorization options. These tools help administrators organize the page so users can find your page and its content through search and navigation. For more information on using tags and categories, see *Organizing Content with Tags and Categories*.

98.6 SEO

SEO provides several ways to optimize the data the page provides to an indexer that's crawling the page. You can set the various meta tags for description, keywords and robots. There's also a separate Robots section for telling robots how frequently the page is updated and how it should be prioritized. If the page is localized, you can select a box to generate canonical links by language. If you want to set some of these settings for the entire Site, you can specify them from the Sitemaps and Robots tabs of the Manage Site Settings dialog box (see below).

Each asset (web content article, blog entry, etc.) has a unique URL. From the search engine's point of view, this makes your pages rank higher since any references to variations of a specific URL are considered references to the same page.

You can also configure the page to use a custom canonical URL. To do so, set the *Use Custom Canonical URL* toggle to *YES*, then enter your desired canonical URL in the field that appears. You can define a custom canonical URL for each language. If there's no value for a specific language, the canonical URL for that language is controlled by the global/instance-level setting.




Figure 98.2: Enter the custom canonical URL that you want to use for the page.

You can also configure canonical URLs at the global and instance levels.

Note: Any custom canonical URLs set for individual pages take precedent over the global and instance level settings.

Global: Control Panel → Configuration → System Settings → Pages → SEO

Instance: Control Panel → Configuration → Instance Settings → Pages → SEO

Navigate to the level (global/instance) on which you want to configure canonical URLs, then choose one of the following in the *Canonical URL* menu:

Use Default Language URL (default): When a user visits a page in any supported language, the default language's URL is used as the canonical URL.

Use Localized URL: The page's localized URL is used as the canonical URL.

98.7 Look and Feel

Look and Feel lets you set a page-specific theme. You can inherit what you already have configured for your Page Set's theme, or you can define a theme per page. See Customizing the Look and Feel of Site Pages for more details.

SEO

This configuration is not saved yet. The values shown are the default.

SEO Configuration

Canonical URL

Use Default Language URL

Choose an Option

Use Default Language URL

Use Localized URL

Save Cancel

Figure 98.3: You can also configure canonical URLs at the global and instance levels.

98.8 Advanced Settings

Advanced Settings contains options useful for specific cases. Some of these are the same as the options available at the Site or Page Set level, but *Custom Fields*, *Embedded Widgets*, and *Customization Settings* are unique to the individual page configuration.

98.9 Query String

You can set a query string to provide parameters to the page. This can become useful to web content templates. You can set a target for the page so that it pops up in a particularly named window or appears in a frameset. And you can set an icon for the page that appears in the navigation menu.

98.10 Custom Fields

Custom Fields lets you edit the custom fields you already have configured for the *Page* resource. If you don't have any custom fields configured in your Site, this option doesn't appear. In this case, navigate to the Control Panel → *Custom Fields* located under the *Configuration* tab. These are metadata about the page and can be anything you like, such as author or creation date. For more information on Custom Fields, see Custom Fields.

98.11 Embedded Widgets

This option only appears if you have embedded one or more widgets on the page.

Widgets can be embedded on a page via web content template or fragment. To learn more about this, see Adding Templates. You can embed a widget on a page layout or theme programmatically. If you're interested in learning more about this, visit the Embedding Portlets in Themes tutorial.

98.12 Customization Settings

This configuration option in the *Advanced* tab lets you mark specific sections of the page you want users to be able to customize. You can learn more about page customizations in Personalizing Pages.

98.13 JavaScript

This shows a JavaScript editor for code that's executed at the bottom of your page. If your Site's theme uses JavaScript (as is usually the case), it's best to add custom JavaScript code to the theme instead. This way, all your Site's JavaScript code remains in one place.

This configuration option is also available for Page Sets like Public Pages and Private Pages. Visit Executing JavaScript in Site Pages for more information on doing this for Page Sets.

98.14 Mobile Device Rules

Apply rules for how this page should render for various mobile devices here. Create them by navigating to Site Administration menu and selecting *Configuration* → *Mobile Device Families*.

PERSONALIZING PAGES

Administrators can designate pages or sections of Widget Pages as customizable. When a user visits such a page, a notification appears stating that the user can customize the page. Users can make customizations only in the sections of pages designated by administrators. Customizations are based on the rows and columns of a page layout. Page customizations are only visible to the user who made the customizations. By default, Site members can make page customizations but non-Site members and guests can't.

99.1 Enabling Page Customizations

To enable page customizations as an administrator,

1. Click *Configure Page* from the *Options* button next to the Page you want to let Site members modify.
2. Select the *Advanced* tab at the top of the page and expand the *Customization Settings* area.
3. Click the *Customizable* selector button to activate customizations.
4. Select the sections of the page that should be customizable.
5. Enable one or more of the *Customizable* sections so Site members can customize sections of the page. Regions that you've designated as customizable are colored blue.

When Site members visit your customizable page, they see an extended Control Menu with a notification saying *You can customize this page*. Site members can toggle whether to view or hide the customizable regions. If you toggle the selector to view customizable regions, the regions on the page are color-coded to help distinguish customizable vs. non-customizable sections of the page.

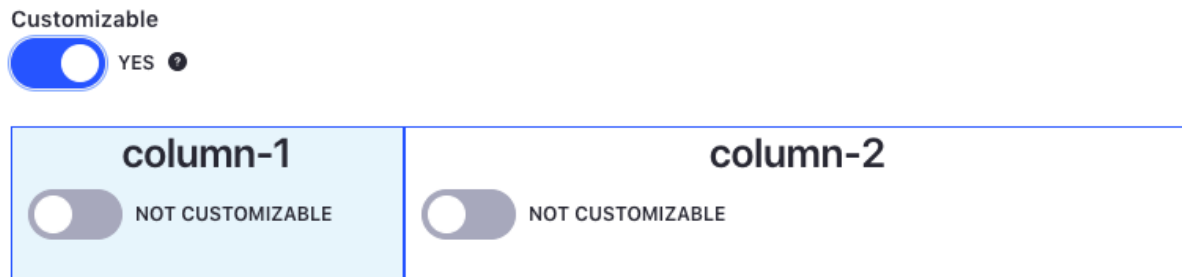


Figure 99.1: To enable page customizations, click on the *Configure Page* button next to the page, expand the *Customization Settings* area, and click on the *Customizable* button.

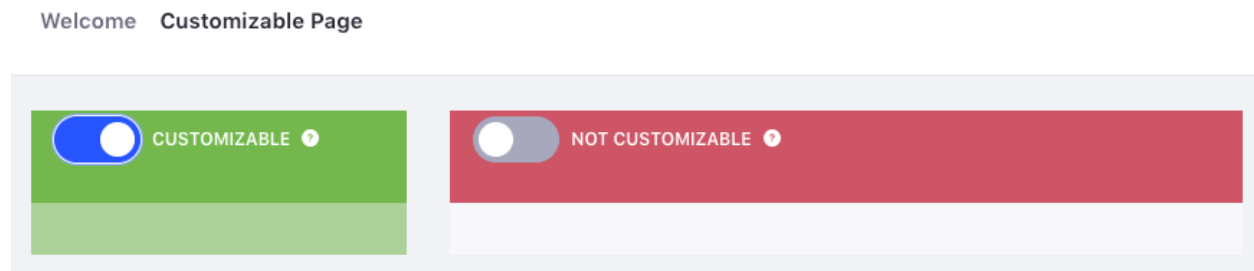


Figure 99.2: Customizable regions are colored green and non-customizable regions are colored red.

99.2 Customization Permissions

Administrators must grant users permission to customize pages under the Site section. This can be achieved by assigning permission to a Role, then assigning this Role to the appropriate users. For example, if you want users to be able to customize your customizable pages, assign the *Customize* permission to the Role *User*. If you want Site members to be able to customize their Sites' customizable pages, accept the default setting. By default, the *Customize* permission is assigned to the Role *Site Member*.

The *Customize* permission also lets users customize the look and feel of apps and import or export app settings.

99.3 Customizing Pages

With customization active, Site members can access the Add menu from the top right side of the screen when viewing their customizable page, which lets them add apps to the customizable sections of the page. If they click *View Page without my customizations*, the Add menu disappears.

Users can make two kinds of customizations to customizable regions:

1. They can configure applications within the customizable regions.
2. They can add apps to or remove apps from the customizable regions.

Reset My Customizations from the *Options* button restores a user's customized page to match the default page, discarding their customizations so they can start anew.

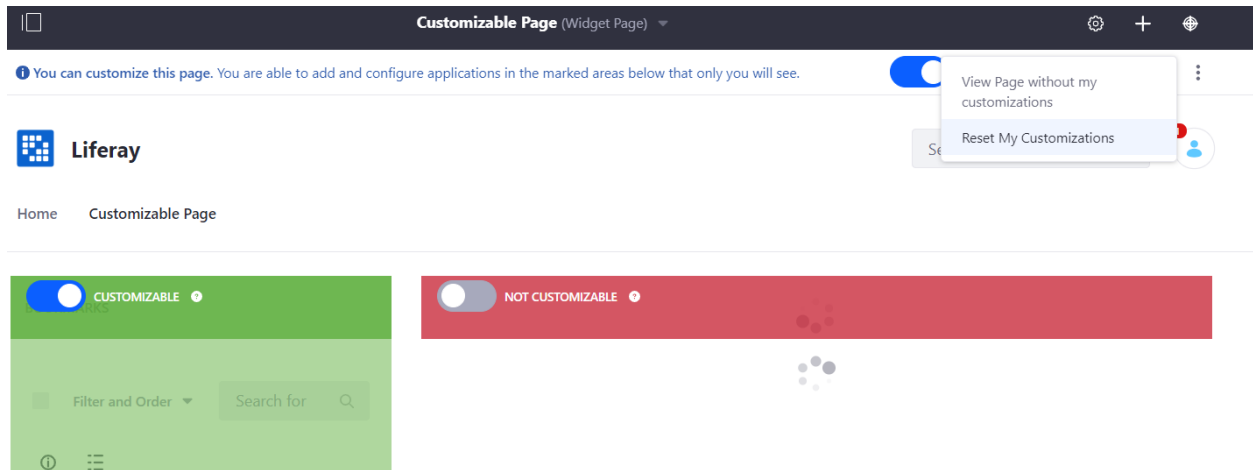


Figure 99.3: Customizable areas are highlighted green when organizing apps on the page.

Users can't change a non-instanceable app's configuration inside a customizable region since those apps are tied to the Site where they've been added.

99.4 Viewing Customized Pages

Site members can also choose between viewing their customized page and viewing the default page by selecting the *Options* button (ⓘ) from the Control Menu and clicking the *View Page without my customizations* or *View My Customized Page*.

Administrators of customizable pages have the same two views as Site members: the *default page* view and the *customized page* view. Changes made to the default page affect all users, whereas changes made to the customized page affect only the administrator who made the changes. Changes made by administrators to non-customizable sections in the default view are immediately applied for all users. Changes made by administrators to customizable sections, however, do *not* overwrite users' customizations.

99.5 Customization Example

As an administrator,

1. Go to *Configure* for the *Welcome* page.
2. Go to the *Advanced* tab and activate Customizations.
3. Set the main column of the Welcome page of the Lunar Resort Site to be customizable.

As a regular user,

1. Navigate to the *Welcome* page.
2. Click *Add* → *Widgets*.
3. Locate the *Language Selector* widget and add it to the page.

The Language Selector application lets users select their language to view a translation of your Site into their native language. After closing the Configuration dialog box of the Language Selector app, the customized Welcome page looks like this:

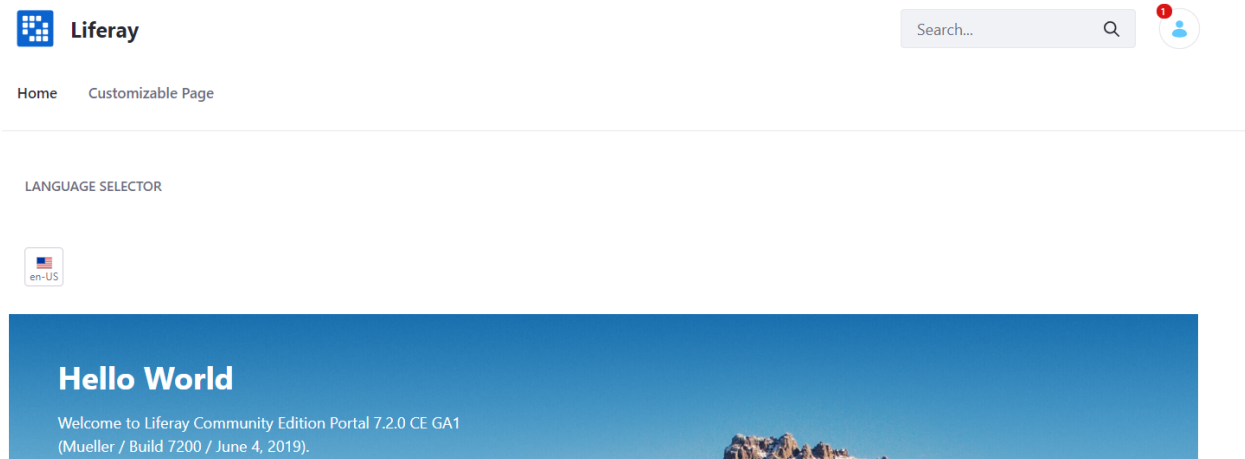



Figure 99.4: In this example, the user added the Language app, and changed the display style from icons to a select box.

CHANGING PAGE PERMISSIONS

Public pages are just that: public. They can be viewed by anybody, logged in or not. And private pages are only private from non-members of the Site. If someone has joined your Site or is a member of your organization, that person can see all the private pages. If you want to further protect some content, you can modify the permissions on individual pages in either page group so only certain users can view them.

Here's how to create a page only administrators can see:

1. Go to your Site's Site Administration dropdown and select *Site Builder* → *Pages* → *Private Pages*.
2. Create a page called *Admin Tips*.
3. Click *Configure* from the Options button dropdown for the page in the left menu.
4. Select *Permissions* from the Options icon () in the top right corner of the screen.
5. Uncheck the *View* and *Add Discussion* permissions next to the Site Member role.
6. Click the *Save* button.

Congratulations! You've changed the permissions for this page so only Site administrators can view it. Any users you add to this Role can now see the page. Other users, even members of this Site, don't have permission to see it.

Pages are as flexible as pages you'd create manually without Liferay. Using a point and click interface, you can define your Site any way you want. You can create and remove pages, export and import them, set their layouts, define how they are indexed by search engines, and more.

| Role | Update Discussion | Delete | Permissions | Customize | Add Page | Delete Discussion | Configure Applications | Update | View | Add Discussion |
|-------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| Guest | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| Owner | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Portal Content Reviewer | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Power User | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Site Content Reviewer | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Site Member | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| User | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

20 Entries ▾ Showing 1 to 7 of 7 entries. < 1 >

Cancel Save


Figure 100.1: The Permissions offer a plethora of options for each role.

CONFIGURING WIDGETS

Just like siblings have common features inherited from their parents, widgets that ship with Liferay DXP also share common features. These include look and feel, exporting/importing app data, communication, sharing, permissions, scoping, and configuration templates. These features work together to facilitate information flow within Liferay DXP and provide an enhanced experience for your users. You'll start with look and feel configuration options.

LOOK AND FEEL CONFIGURATION

To access the look and feel configuration menu of any widget,

1. Click *Options* () in the top right corner of the widget.
2. Select *Look and Feel Configuration*.

Look and Feel Configuration has six tabs:

- General
- Text Styles
- Background Styles
- Border Styles
- Margin and Padding
- Advanced Styling

After making customizations, click *Save* and refresh your page to apply your changes. If you don't like the effect of your changes, some tabs have a *Reset* button to discard changes.

102.1 General Settings

On the General tab are the following options:

Use Custom Title enables changes to your widget's title. The value in the title box is displayed on widget's decorator. The title is localizable, so you can provide translations of the title for different languages.

Application Decorators gives you the choice between three decorators: *Barebone*, *Borderless*, and *Decorate*. The Decorate application decorator is the default. Be careful about turning widget borders off; some themes assume widget borders are turned on and may not display correctly with them turned off.

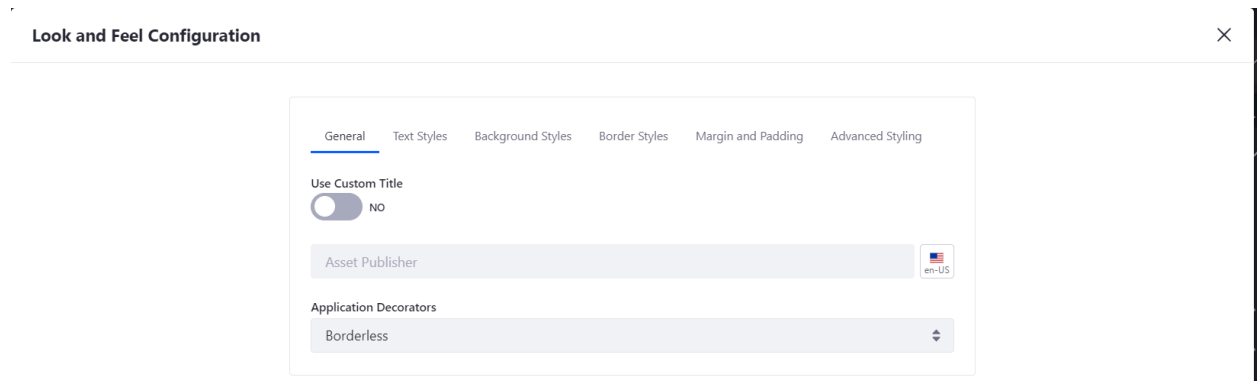


Figure 102.1: The General tab of the Look and Feel Configuration menu lets you define a custom widget title and select the widget contrast option using decorators.

102.2 Text Styles

Text Styles configures the format of the text that appears in the widget. The options include

Font: Choose various fonts. You can set the text to bold, italics, or both.

Size: Set the font size anywhere from 0.1 em to 12 em, with 0.1 em increments. 1 em is the default.

Color: Set to any six digit hex color code. Click on the text box to open the color palette.

Alignment: Set to *Left*, *Center*, *Right*, or *Justified*.

Text Decoration: Set to *Underline*, *Overline*, or *Strikethrough*. The default text decoration is *None*.

Word Spacing: Set from -1 em to 0.95 em, with 0.05 em increments. 0 em is the default.

Line Spacing: Set from 0 em to 12 em, with 0.1 em increments. 0 em is the default.

Letter Spacing: Set from -10 px to 50 px, with 1 px increments. 0 px is the default.

102.3 Background Styles

The Background Styles tab specifies the widget's background color. When you select the text space, you're given a color palette to choose your background color or you can manually enter any six digit hex color code.

102.4 Border Styles

The Border Styles tab, configures your widget's border width, style, and color. For each of these attributes, leave the *Same for All* selector enabled to apply the same settings to top, right, bottom,

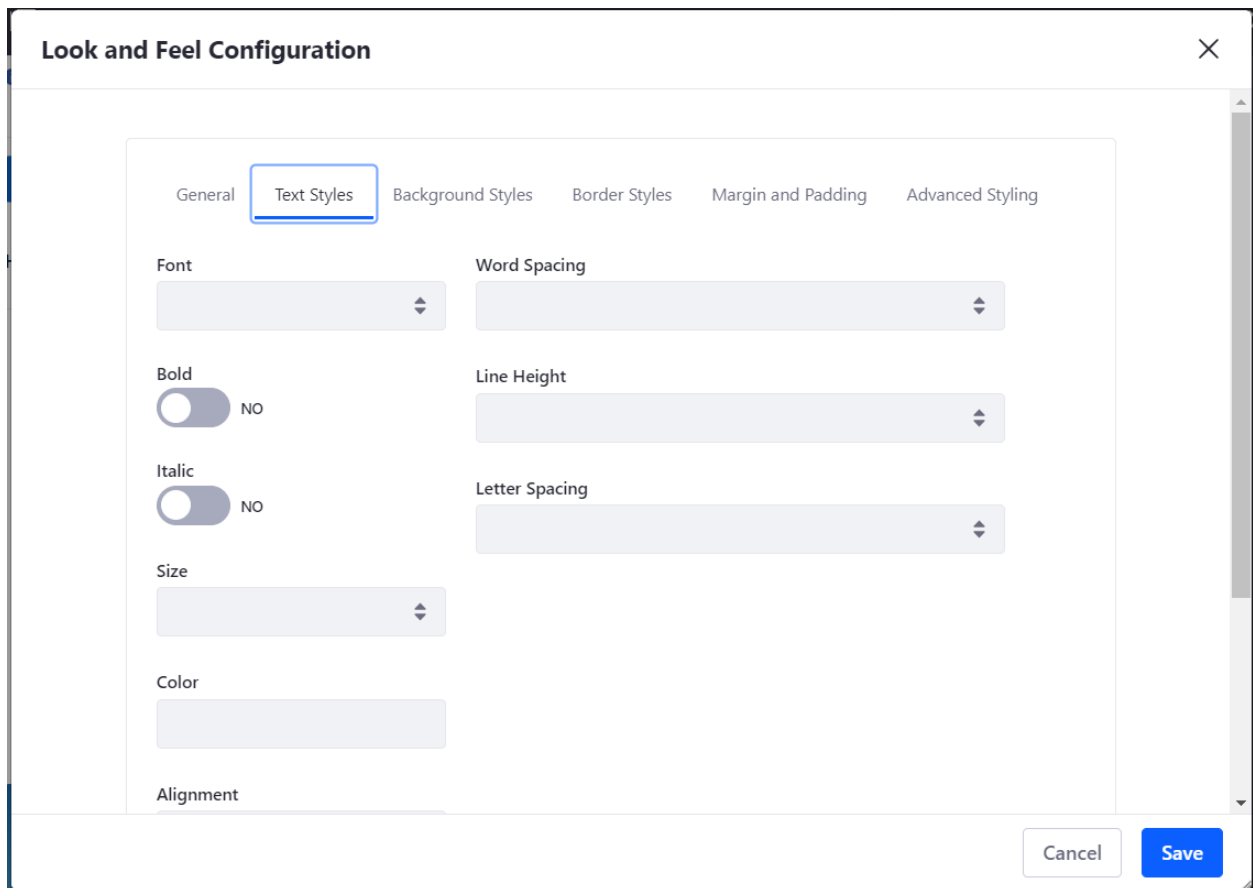


Figure 102.2: The Text Styles tab lets you configure the format of the text that appears in the widget.

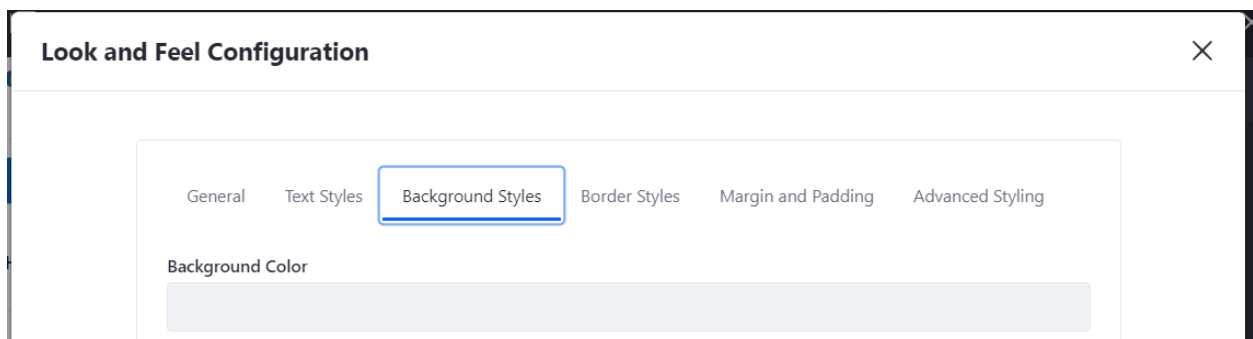


Figure 102.3: The Background Styles tab lets you specify the widget's background color.

and left borders.

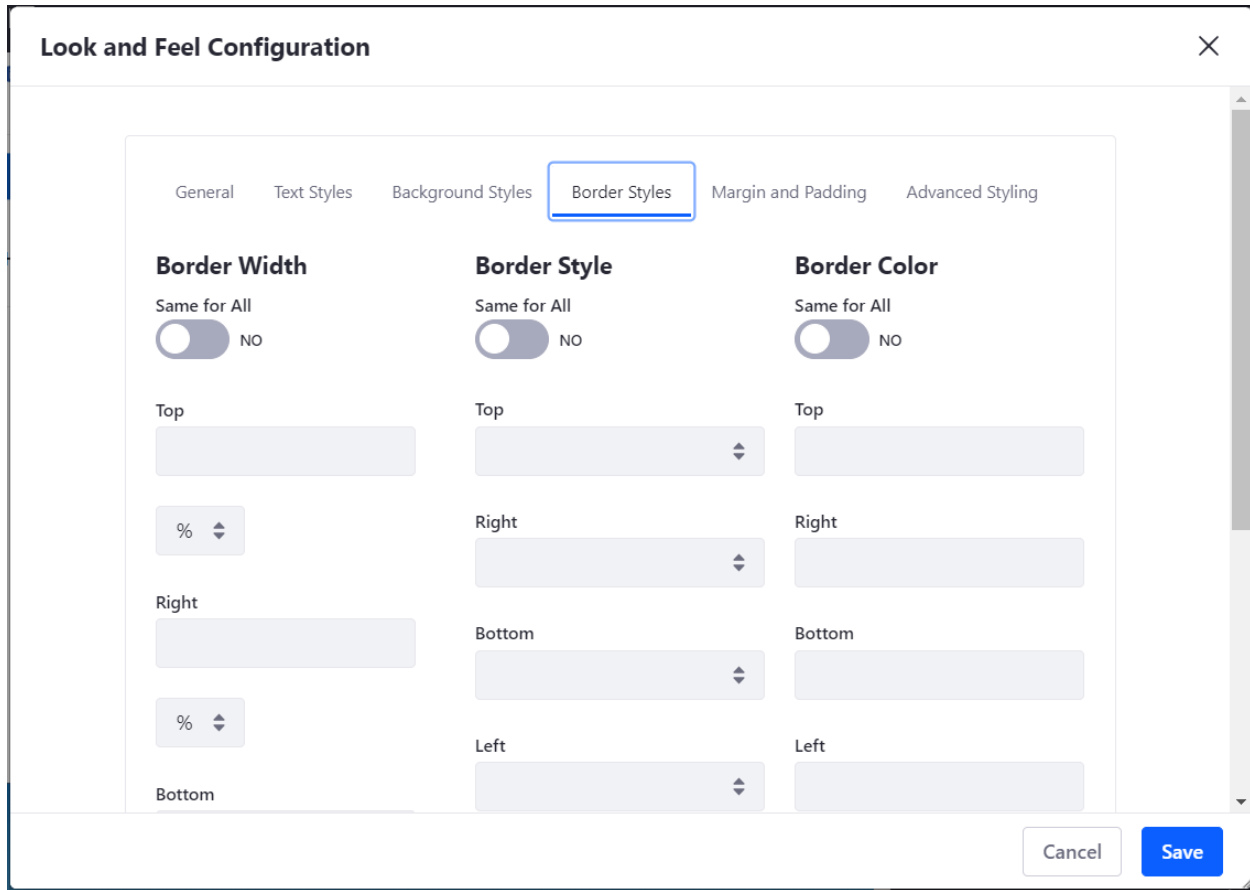


Figure 102.4: The Border Styles tab lets you specify a border width, style, and color for each side of the widget.

For border width, you can specify any % value, em value, or px value. For border style, you can select Dashed, Double, Dotted, Groove, Hidden, Inset, Outset, Ridge, or Solid. For border color, you can enter any six digit hex color code, just like for the text color and background color. You can also use the color palette.

102.5 Margin and Padding

The Margin and Padding tab specifies margin and padding lengths for the edges of your widget. Just like for border styles, leave the *Same for All* selector enabled to apply the same settings to each side (top, right, bottom, and left) of the widget.

For both padding and margin, you can specify any % value, em value, or px value.

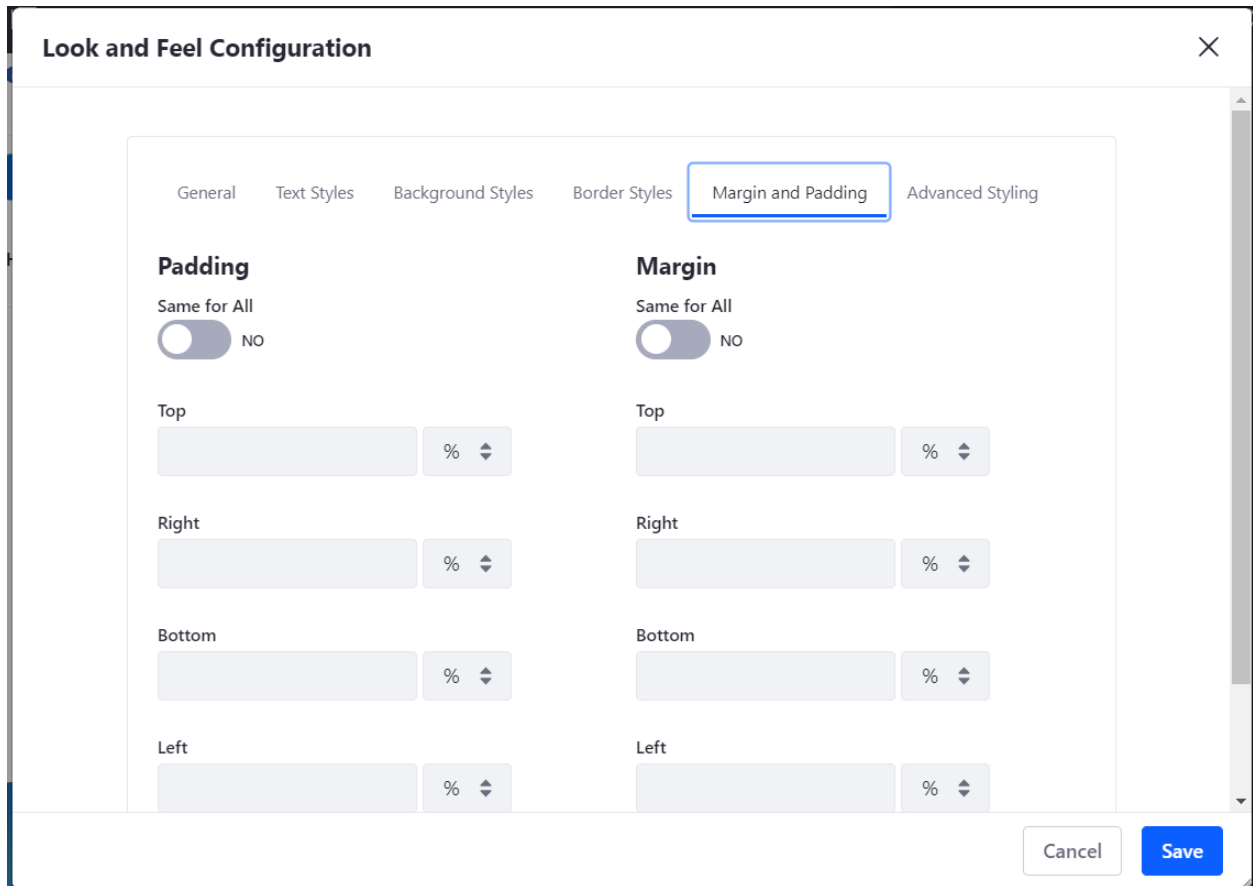


Figure 102.5: The Margin and Padding tab allows you to specify margin and padding lengths for the sides of your widget.

102.6 Advanced Styling

The Advanced Styling tab displays current information about your widget, including your widget's Liferay ID and CSS classes.

You can also enter custom CSS class names for your widget and custom CSS code. Clicking the *Add a CSS rule for just this portlet* or *Add a CSS rule for all portlets like this one* links adds the CSS code shells into your custom CSS text box. If you check the *Update my styles as I type* box, your CSS code is applied dynamically to your widget so you can see the effects of your edits.

Next, you'll learn about communication between widgets.




Figure 102.6: The Advanced Styling tab displays your widget's Liferay ID and allows you to enter CSS code to customize the look and feel of your widget.

EXPORTING/IMPORTING WIDGET DATA


You may need to export data from a specific widget instance, without regard to content on the rest of the Site. There are many widgets that let you export or import their data individually:

- Blogs
- Bookmarks
- Dynamic Data Lists
- Forms
- Knowledge Base
- Message Boards
- Web Content
- Wiki
- And more

Exporting widget data produces a .lar file that you can save and import into another widget of the same type. To import widget data, you must select a .lar file. Be careful not to confuse widget-specific .lar files with Site-specific .lar files. See the Importing/Exporting Pages and Content article for information on importing/exporting Site page data.

There are two ways to export/import widget content. You can navigate to the widget's administrative area located in the Product Menu, or you can visit the widget on its page. Both export/import menus work the same, but the administrative area may hold content different from its widget counterpart (e.g., Web Content Admin in Product Menu and Web Content Display widget do not offer same content for export/import), so be wary of your selection.

To export or import data from the widget's administrative area, follow the steps below.

1. Navigate to the widget's designated area in the Product Menu. For example, if you plan to export Web Content data, navigate to *Content* → *Web Content*.
2. Click the *Options* button () from the top right of the page and select *Export/Import*.
3. Select the *Export* or *Import* tab to begin configuring the respective process.

To export or import data from a widget, follow the steps below:

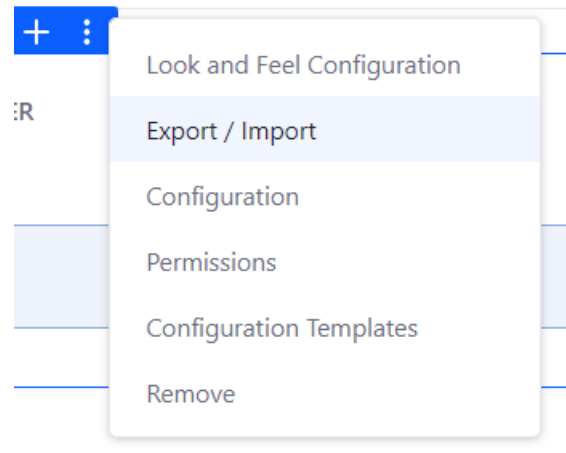


Figure 103.1: You can access a widget's administrative *Export/Import* feature by selecting its Options menu.

1. Ensure the widget you're exporting/importing from is available on a page. You can add widgets from the *Add* (☰) → *Widgets* menu.
2. Select the widget's *Options* button (⋮) and select *Export/Import*.
3. Select the *Export* or *Import* tab to begin configuring the respective process.

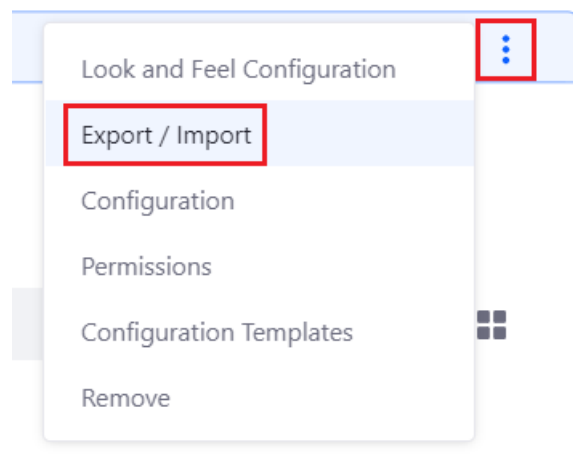


Figure 103.2: You can access a widget's *Export/Import* feature by selecting its Options menu.

Now that you know how to navigate to the *Export/Import* menus, you can explore the export process.

103.1 Exporting Widget Data

To export widget data, create a new export process by selecting the *New Export Process* tab (default). You have several export options to configure.

First, you can choose to export your widget's configuration settings. This exports your customized settings from your widget's *Options* → *Configuration* menu. For some widgets, the configuration export might also include content. For example, a Web Content Display widget that shows a web content article also exports the article when exported, even though no content is selected. This applies when publishing a Web Content Display widget too; the configured article is published with the widget.

Next, you can select a *Date Range* of content that you want to export. Content added to your widget within your specified date range is included in the .lar file. The following date range choices are available:

All: Publishes all content regardless of its creation or last modification date.

Date Range: Publishes content based on a specified date range. You can set a start and end date/time window. The content created or modified within that window of time is published.

Last...: Publishes content based on a set amount of time since the current time. For example, you can set the date range to the past 48 hours, starting from the current time.

By checking the *Content* box, you can choose specific content you want to export. When you check the *Content* box, more options appear, letting you choose specific kinds of metadata to include. For example, if you have a wiki page with referenced content that you don't want, check the *Wiki Pages* checkbox and uncheck the *Referenced Content* checkbox. Another option is the selection of content types. Two familiar content types in your Liferay instance are *Comments* and *Ratings*. If you want to include these entities in your .lar file, select *Change* and select them from the checklist. For more information on managing content types, see the *Managing Content Types in Staging* article.

Next, you can choose to export individual deletions. This lets delete operations performed for content types be exported to the LAR file.

Finally, you can choose whether to include permissions for your exported content. The permissions assigned for the exported widget window are included if you enable the *Export Permissions* selector.

After you've exported your widget's data, switch to the *Current and Previous* tab to view ongoing export processes and the history of past exports. You can also download the exported .lar file from this tab.

103.2 Importing Widget Data

To import widget data, you can select the LAR using your file explorer or by dragging and dropping the file between the dotted lines.

Your LAR file is uploaded and displayed to you for review. Click *Continue*.

Now that you've uploaded and confirmed your LAR file, you're given a similar screen to what you'd be offered during export. Several of these options are covered in great detail in the *Importing/Exporting Pages and Content* tutorial. There are some additional options available: *Update Data* and *Authorship of the Content*. Here's options and descriptions for each section:




Figure 103.3: When importing widget data, you can choose a LAR file using the file explorer or drag and drop the file between the dotted lines.

Update Data

Mirror: All data and content inside the imported LAR is newly created the first time while maintaining a reference to the source. Subsequent imports from the same source updates entries instead of creating new entries.

Mirror with overwriting: Same as the mirror strategy, but if a document or an image with the same name is found, it is overwritten.

Copy as New: All data and content inside the imported LAR is created as new entries within the current Site every time the LAR is imported.


Authorship of the Content

Use the Original Author: Keep authorship of imported content whenever possible. Use the current user as author if the original one is not found.

Use the Current User as Author: Assign the current user as the author of all imported content.

Once you've selected the appropriate options, select *Import* and your widget's data is imported and ready for use.

COMMUNICATION BETWEEN PORTLET WIDGETS

Portlet widgets can communicate with each other using public render parameters and events. Public render parameters are easy to use and can be quite powerful. Some Liferay portlets provide a configuration UI to help you get the most out of this communication mechanism. To access this UI, open your portlet's configuration window by clicking on the *Options* icon () and selecting *Configuration*. Then click on the *Communication* tab.

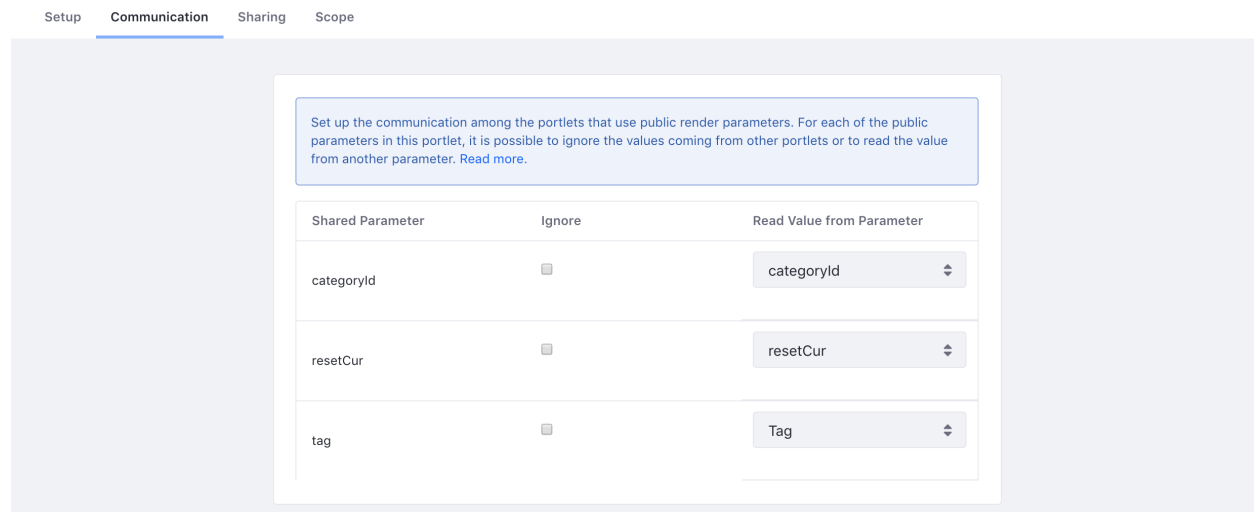


Figure 104.1: You can configure portlets to communicate with each other using public render parameters.

Note: If your widget isn't a portlet, this feature isn't available.

The screenshot above is for the Wiki, which has six public render parameters: `categoryId`, `nodeId`, `nodeName`, `resetCur`, `tag`, and `title`. For each of these parameters, you can configure the portlet to ignore the values coming from other portlets or read the value from another parameter.


Why might it be useful to ignore the values for certain parameters that come from other portlets?

Consider a common use case for the Wiki application. The Wiki portlet is often used along with the Tags Navigation portlet so that when a user clicks on a tag of the latter, the Wiki shows a list of pages with that tag. An administrator may want the Wiki to show the front page always independently of any tag navigation done through other portlets. Ignoring the values of the parameter coming from other widgets let this happen.

Reading the value of a parameter from another portlet is an advanced but very powerful option that allows portlets to communicate with each other even if their developers didn't intend them to. For example, imagine that the Wiki is used to publish information about certain countries, and there's another portlet that allows browsing countries for administrative reasons. The second portlet has a public render parameter called *country* with the name of the country. You'd like the Wiki to show the information from the country that's selected in the administration portlet. This can be achieved by setting the value of the title parameter of the Wiki portlet to be read from the country parameter of the administration portlet. Cool, isn't it?

SHARING WIDGETS WITH OTHER SITES

You can share widgets with other Sites by embedding an instance of a widget running on your Site into another website, such as Facebook. This opens up a whole new avenue of exposure to your web site that you would not have had otherwise. In fact, this is how all those Facebook games work.

To share one of your widgets, open the *Configuration* dialog box from the widget's *Options* icon () and select the *Sharing* tab. There are five sub-tabs under sharing: Any Website, Facebook, OpenSocial Gadget, and Netvibes.

105.1 Any Web Site

Copy and paste the provided snippet of JavaScript code into the web site where you want to add the widget. That's all you need to do. When a user loads the page on the other website, the code pulls the relevant widget from your Site and displays it.

105.2 Facebook

You can add any widget as a Facebook app. To do this, you must first get a developer key. A link for doing this is provided to you in the Facebook tab. You must create the application on Facebook and get the API key and canvas page URL from Facebook. Once you've done this, you can copy and paste their values into the Facebook tab. Save the configuration and navigate back to the Facebook tab in Liferay DXP. You're given the Callback URL, which you can copy and paste into Facebook. When opening your app in Facebook, the correct callback URL is used to render the application. You can also enable the *Allow users to add [application-name] to Facebook*. Then you can navigate to your app's Options menu and select *Add to Facebook*.

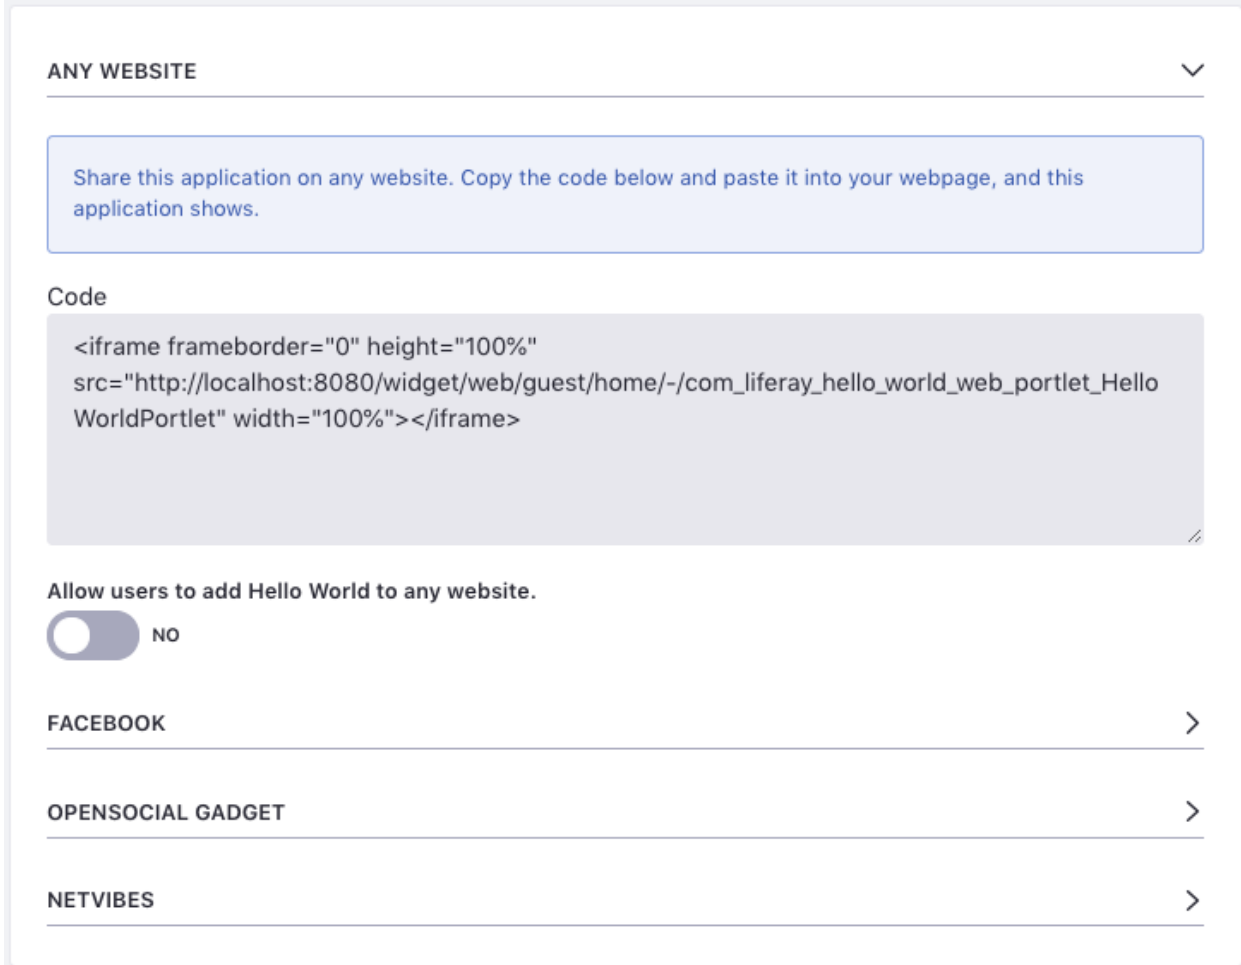


Figure 105.1: The Sharing tab in your widget's Configuration menu lets you share your widget in a variety of ways.

105.3 OpenSocial Gadget

OpenSocial comprises a container and a set of APIs for social networking and other web applications. Liferay DXP can serve up applications to be used as OpenSocial Gadgets on any OpenSocial-compatible pages.

To serve a Liferay widget on an OpenSocial platform, copy and paste the provided gadget URL and add it to the appropriate configuration page of the OpenSocial platform you're using. Your Liferay instance serves that widget directly onto that platform's page. The URL provided is unique to the specific instance of the widget, so you could serve multiple instances of the same widget as different OpenSocial Gadgets.

From the Sharing tab in the Configuration menu, you can also enable the selector *Allow users to add [application-name] to an OpenSocial platform*. Click *Save* and revisit the *Options* button of your widget. A new button appears named *Add to an OpenSocial Platform*. When selecting this new button, the URL is provided for sharing the widget to an OpenSocial platform.

105.4 Netvibes

Netvibes offers a similar environment where users can log in, create their own personal dashboard, and add customizable widgets to it. To set up Netvibes support for a widget, enable the *Allow users to add [application-name] to Netvibes pages* selector. You can then use the provided URL to create a custom Netvibes widget based on the instance of the Liferay widget that you're using.

Next, you'll learn how to set permissions for Liferay applications.

WIDGET PERMISSIONS

All of Liferay's widgets support Liferay DXP's robust, fine-grained permissions system. Some higher level permissions can be configured in the permissions tab of the widget's configuration dialog box. You can grant Roles permission to


- Add a display template
- Add the widget to a page
- Configure the app
- Modify the widget's permissions
- Modify the widget's preferences
- View the widget

| Role | Add Display Template | Permissions | Preferences | Configuration | View | Add to Page |
|-------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| Guest | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| Owner | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Portal Content Reviewer | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Power User | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| Site Content Reviewer | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Site Member | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| User | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |

20 Entries • Showing 1 to 7 of 7 entries. < 1 >

Cancel Save

Figure 106.1: Viewing the permissions configuration for a widget.

To set these permissions, go to the widget's *Options* icon () and click select *Permissions*. This shows you a table of Roles. Use the check boxes to grant certain permissions to different Roles. Click *Save* after you've made your selections.

Beyond this, specific permissions are generally defined for specific widgets. For example, Message Boards contains a *Ban User* permission. This makes no sense in the context of most other widgets. You'll go over permissions for specific widgets in the sections for those widgets. Next, you'll learn about widget scopes.

WIDGET SCOPE


As you learned earlier, Roles can be scoped by the instance, by Site, or by an Organization. A Role only takes effect within its scope. For example, a Message Boards Administrator Role with complete access to the Message Boards has different permissions based on the Role's scope. If it's a global Role, members have permission to administer message boards across the entire installation. If it's a Site Role, members only have permission to administer message boards within the Site where they've been assigned the Role. For Organizations with Sites, Site Roles are automatically assigned to Organization members based on the Organization Roles they have. For an Organization-scoped Message Boards administrator Role, members only have permission to administer message boards within the Site of the Organization that assigned the Role to them.

You've also heard the word *scope* refer to the data set of a widget. By default, when a widget is added to a page in a Site, it is *scoped* for that Site. This means its data belongs to that Site. If the widget is added to a page in a different Site, it employs a completely different data set. This enables you to place a Message Boards widget in one Site with one set of categories and threads, and place another Message Boards widget in different Site with a different set of categories and threads.

Scoping by Site means that you can only have one Message Boards widget per Site. If you add one Message Boards widget to a page in a Site and add another Message Boards widget to a different page in the same Site, the second Message Boards widget contains exactly the same data as the first. This is because, by default, the Message Boards widget is scoped by Site. Most of Liferay DXP's other widgets also default to being scoped by Site.

To avoid this limitation, many Liferay widgets can be scoped by page. The data sets of page-scoped widget serve a single page, not an entire Site. If you set the scope of a widget to *page* instead of *Site*, you can add any number of these widgets to different pages, and then they have different sets of data. Then you can have more than one message board per Site if you wish. Most widgets, however, default to the "native" configuration, and have their scopes set to the Site where they are placed.

Unless otherwise noted, all widgets support scoping by instance (global), Site (default), or page. This grants you some flexibility in how you want to set up your Liferay instance. You can configure the scope of a widget's content with just a few simple steps.

1. Click the widget's *Options* icon ()
2. Select *Configuration*.

3. Select the *Scope* tab.
4. Use the drop-down menu to set the scope.

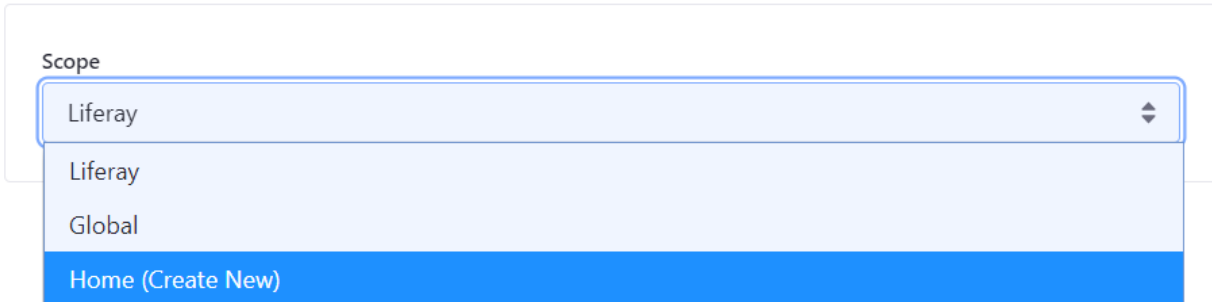



Figure 107.1: You can change the scope of your widget’s content by navigating to its Configuration menu.

Once you’ve created a new scope for widgets, a button () with a drop-down menu appears in the *Content & Data* menu for you to select which scope to manage content for. You can choose the default scope or any new scopes you created for your widgets. Your selection changes the content that appears when you manage each type.

That’s all it takes to change the scope for a particular widget. By setting the scope to the current page, you can add as many of these widgets to a Site as you want, provided they are all added to separate pages.

Another useful feature of Liferay’s widgets is Configuration Templates.





Figure 107.2: Use the drop-down menu under Content & Data to determine which scope to manage content for.

CONFIGURATION TEMPLATES

Once you've configured a widget, Configuration Templates can save those settings in a reusable template. If someone goes in and changes the settings of a particular widget, it then becomes easy to revert those changes back to the original configuration template. Configuration templates are only available for widgets placed on a page. Applications available from the Product Menu do not provide configuration templates.

To create a configuration template, click the *Options* icon () from the menu in the widget's title bar and select *Configuration Templates*. If widget's current settings are the ones you want to save, click the *Save Current Configuration as Template* button. If not, change the settings until it's configured the way you want it, and then click the button.

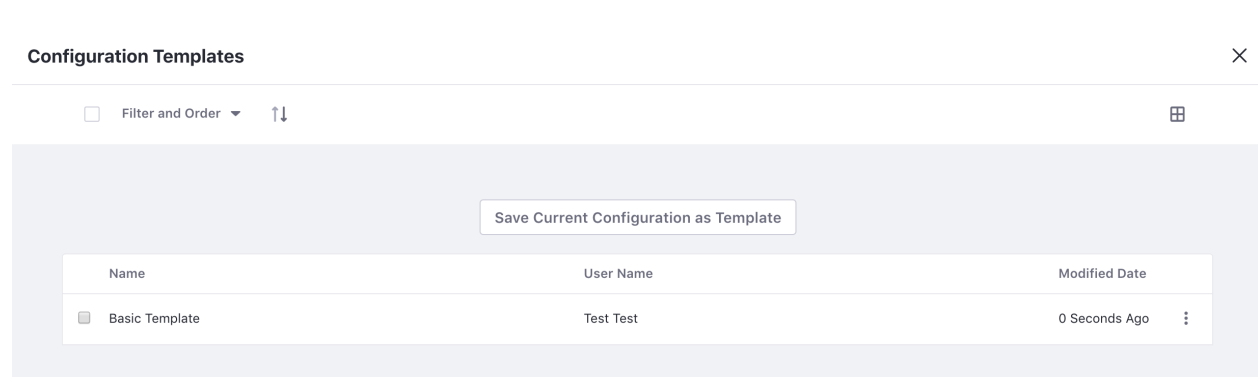



Figure 108.1: Create a configuration template to save your app's configuration settings.

There is only one field to fill out. Enter a name for your template and click *Save*. You should now see your configuration in the list. If you ever need to revert the app to these archived settings, you can click *Actions* () → *Apply* next to the configuration template you want to apply.

Unless otherwise noted, all widgets in Liferay DXP support this feature. This is particularly useful for widgets that have a lot of configuration options, such as the Message Boards application.

108.1 Summary

You've now explored the configuration options available for Liferay widgets. You learned how to customize your widgets, export/import data, communicate between widgets, take advantage of different scopes, and save configuration settings. You also examined the different uses of social applications like Facebook and Netvibes for your Liferay widgets. In all, Liferay DXP gives you an abundance of options to leverage the full capability of your widgets.

MANAGING MEMBERS IN YOUR SITE

Users and Sites are important concepts. Sites are where all your content and pages are stored, and Users access and create that content. While user management is covered in depth in our User Management tutorial, there are other user configuration options specific to Site Management:

- Adding members to Sites administratively
- Adding members to Sites automatically
- Creating Teams of Site members for various functions

This section of tutorials shows you how to manage and configure these Site options for Users.

ADDING MEMBERS TO SITES

In Adding Sites you learned the difference between Site Membership Types and about public and private pages within a Site. Now you'll learn how to add users manually to Sites and how to provide options for self management. For review, there are a few key reasons why Site Membership management is important:

1. Only Site Members can view a Site's Private Pages.
2. Site Members have more permissions than guests for many widgets like Message Boards and Wikis that enable them to create content and collaborate on your Site.
3. Site Members can be associated with Roles that grant Site privileges.

110.1 Administrating Site Membership

Administrators can manage Site members from that Site's *Site Membership* page.

1. Open *Site Administration* and select the Site that you want to manage members for.
2. Click on *People* → *Memberships*

From here you can manage Site Memberships, Organization, and User Group associations. You can learn more about those in the Users and Organizations tutorial. Here you see a list of all of the current Users of the Site and you can add or remove user memberships from the Site.

110.2 Adding Members to a Site

Follow these steps to make an existing user a member of the Site:

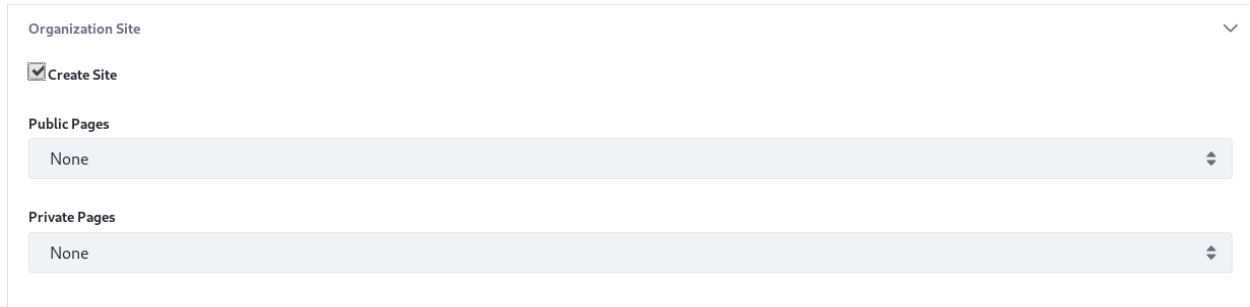


Figure 110.1: The current members of the Site as displayed on the *Site Memberships* page.

1. Click the *New* (+) button in the top right of the screen.
2. Use *Filter and Order* or the *Search* function to locate the User you want to add to the Site.
3. Select the User(s) you wish and click *Done*.

On the *Assign Users to This Site* screen, all Users eligible to be added to the Site appear. Deactivated Users do not appear. Site members also appear, but with a greyed-out checkbox.

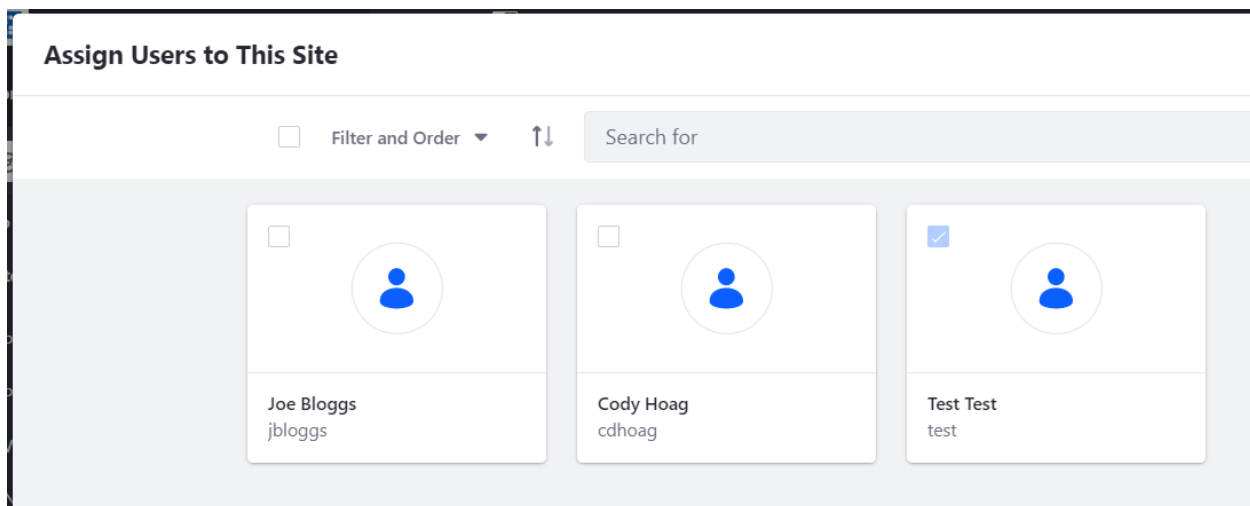



Figure 110.2: The list of users available to add to the current Site. Note that the current members are visible but cannot be added or removed here.

110.3 Removing User Membership from a Site

There are two ways to remove a user from a Site. You can remove an individual member like this:

1. Click the *Actions* () icon for the user that you want to remove.
2. Select *Remove Membership*.
3. In the pop-up that appears, confirm the removal.




Figure 110.3: Selecting to remove a user.

To remove several users at once, you can do this:

1. Click the checkbox for each user that you want to remove.
2. In the menu at the top of the page, click the X icon to remove the users from the Site.
3. In the pop-up that appears, confirm the removal.

Removed Users lose access to the Site's private pages and membership in any Site Roles or Teams they had.

110.4 Assigning Site Roles

Roles grant permissions in Liferay DXP. Roles can be assigned for the entire instance or just for one specific Site or Organization. Site Roles assign permissions for a specific Site.

You can use the same interface options that you used to remove Users from the Site to assign them to Site Roles. If you select a User or Users and click *Assign Site Roles* (either through the *Actions* menu or the menu at the top), you are taken to the *Assign Site Roles* screen. From here:

1. Select the Roles that you want to assign to the selected users.
2. Click *Done*.




Figure 110.4: Assigning Site Roles.

Site Roles are created at a global level, but when they're assigned they only provide privileges for the specific Site where they were assigned. Since Roles are created at a global level, they cannot be created by Site Administrators (since Site Administrators only have Administrator privileges for their Site). **Teams** allow Site Administrators to assign permissions to groups of Users within their Sites. Next, you'll look at more configuration for managing members of your Site.

CREATING TEAMS TO EMPOWER SITE MEMBERS

If you have an *ad hoc* group of Users who perform the same set of tasks in a Site, you can organize them into Site Teams, and then assign the team permissions for various site-specific functions. Site Teams are the preferred method for collecting permissions within a single Site. Some common functions to assign a Site Team include

- Moderating site Wiki content
- Managing Message Boards threads
- Writing blogs
- Editing a specific page in the site

If your Site has Message Boards, you might want to enable a subset of the Site's members to moderate the categories and threads, and perhaps to ban abusive/offensive posters. To do this, you could create a Site Team named *Lunar Resort Message Board Moderators*, define the team's permissions in the Message Boards application, and assign the desired Site members to the team.

The permissions assigned to a Site Team only apply to that Site. The two key features of Teams are that they are limited to their Sites and that they empower Site Administrators to manage permissions for their Sites since Site Administrators cannot create new Roles.

Note: To create and apply permissions for a group of users to use across multiple Sites or Organizations in your Liferay instance, consider aggregating the Users into a User Group and assigning the User Group permissions via Roles.

To create a team within a Site,

1. Go to the Site Administration page of your Site.
2. Select *People* → *Teams*.


It's important to note that configuring other Site membership groupings, such as *Users*, *Organizations*, and *User Groups* can be done in the *Site Memberships* app, which is also in the Members tab. You can visit User Management for more information on how Site memberships work.

3. Finally, click the *Add Team* icon ()


The screenshot shows a 'New Team' form. At the top, there is a dark blue header with the text 'New Team' and a small question mark icon. Below the header, the form is contained within a light gray border. The first section is labeled 'Name *' and has a text input field containing 'Lunar Team'. The second section is labeled 'Description' and has a larger text area with a placeholder 'Description'. At the bottom left of the form, there are two buttons: a blue 'Save' button and a white 'Cancel' button with a gray border.

Figure 111.1: Creating teams within your site can foster teamwork and collaboration, as team permissions enable team members to access the same resources and perform the same types of tasks.

4. Enter a name and a description and click *Save*. Your new team shows in the list.
5. To add members, click on the team name link and then select *Add Team Members*.

To manage a team's permissions, click the *Actions* icon () and select *Permissions* for that team. Setting permissions for the team assigns those permissions to all the team's members. Only administrators who can edit/manage the team can manage team permissions.

If you created a team whose task is to moderate the Message Boards, for example, you'd want to give the team all the permissions they'd need.

1. Go to *Site Administration* → *Content & Data* → *Message Boards*.
2. Select *Home Category Permissions* from the *Options* icon () in the top right of the screen.
3. Find the Team in the Role column and select the appropriate permissions.

That's it! It's easy to give groups of Site Users permissions to perform their tasks.

These tutorials have introduced you to Liferay DXP Site management. You've learned how to use Liferay DXP to create multiple Sites with different membership types. You've also seen how easy it is to create and manage Sites and to create and manage pages within Sites. Next, you'll begin working with web content.

The screenshot shows a table titled "Home Category Permissions" with a search bar at the top. The table lists four roles and their permissions across twelve categories. The "Lunar Resort Message Board Moderators" role has all permissions checked, while the other roles have a subset of permissions.

| Role | Permissions | Add File | Ban User | Add Category | Reply to Message | Lock Thread | Subscribe | View | Add Message | Move Thread | Update Thread Priority |
|---------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| Lunar Resort Message Board Moderators | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Guest | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Owner | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Portal Content Reviewer | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

Figure 111.2: The Lunar Resort Message Board Moderators Site Team has unlimited permissions on the Message Boards application.

MANAGING CONTENT

There are two primary ways to manage and display content in Liferay DXP

- Web Content
- Content Sets

The Web Content framework helps users who are not web developers publish content with a simple point and click interface, while enabling developers to create complex templates with dynamic elements.

A Content Set defines a list of content, and then that list can be displayed.
Continue on to learn more about managing content!

MANAGING WEB CONTENT

Web Content Management (WCM) helps users who are not web developers publish content with a simple point and click interface, while enabling developers to create complex templates with dynamic elements. Once these templates have been deployed into Liferay DXP, your non-technical users can use them to manage complex content as easily as they would manage basic content.

It has these components:

Web Content Editor: A complete HTML editor for modifying fonts, adding color, inserting images, and much more.

Structure Editor: Define fields for structured content and more advanced designs.

Template Editor: Import template script files or create your own template to inform the system how to display the content within the fields determined by the structure.

Web Content Display: Place web content on pages in your Site.

Asset Publisher: Aggregate and display different types of content together in one view. This is covered in more detail in Publishing Assets.

Scheduler: Schedule when content is reviewed, displayed or removed. This is covered in more detail in Scheduling Web Content Publication.

Workflow Integration: Run your content through a review process. This is covered in more detail in the Workflow section.

Staging: Use a separate staging server or stage your content locally so you can keep your changes separate from the live site. This is covered in more detail in the Staging section.

These tools streamline the content creation process for end users and are also integrated with Liferay's services so advanced template developers can use them to query for data stored elsewhere on your website.

To demonstrate Liferay DXP's Web Content Management features, you'll create and manage content on Liferay for the ambitious (and fictitious) *Lunar Resort* project. The Lunar Resort project specializes in facilitating lunar vacations. It provides space shuttle transportation from the Earth to the Moon and back, offers the use of a state-of-the-art recreational facility enclosed by a large, transparent habitat dome, and even rents out lunar rovers.

PUBLISHING BASIC WEB CONTENT

Web Content is one of many different kinds of assets, along with blog posts, wiki articles, message board posts, and other kinds of content. Like all of these assets, Liferay DXP handles Web Content using an asset framework that includes categories, tags, comments, ratings and more. Please see [Publishing Content Dynamically](#) for more information on Liferay's asset framework.

To start working with Web Content, publish some basic material using Web Content Management's WYSIWYG editor. Then you can cover the editor's features in greater depth and learn how to publish the content you create to a page.

CREATING WEB CONTENT

To start you'll create and publish some simple content using the WYSIWYG editor to the home page of the Lunar Resort's web site.

1. Go to *Site Administration* and click the *Site Selector* button (📍).

Content is created in whichever Site is selected, so always make sure that you're working on the right Site. Content can also be created as Global so that any Site can access it.




Figure 115.1: You can choose where to create content by navigating to the Site Administration menu and selecting your Site and page scope.

2. Select the *Lunar Resort* Site.
3. Open the *Content & Data* section and click on *Web Content*.

Here you can create web content and organize it into folders.

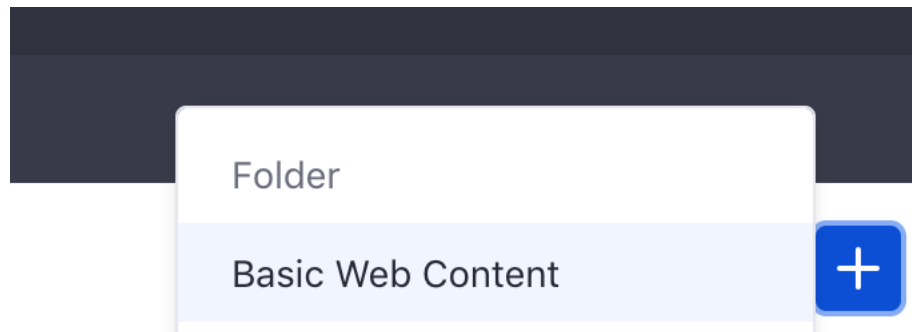


Figure 115.2: By default, *Basic Web Content* is the only article type available. The next tutorial covers how to create new types.

4. Click Add (+) → *Basic Web Content* to create a new web content article.
5. Type *Welcome to the Lunar Resort* in the top *Title* field.
6. In the *Summary* field under the Basic Information tab on the right side, give a short description of the Lunar Resort's facilities (be creative).
7. In the *Content* field, add the body of your web content article, which you'll dive into next.
8. Click *Publish*.

That's all it takes to create Web Content, but there's much more under the hood of the web content editor that you'll learn about next.



USING THE WEB CONTENT EDITOR


In the previous article, you created a simple web content article, but the Web Content editor can do much more than make plain text articles.

116.1 Basic Editor Functions

To explore these options, go back to the article you just created and make it better:

1. On the *Web Content* page in *Site Administration*, click on the title of the article.
2. Highlight the text you entered in the *Content* field.
A number of controls appear. These let you style the text or provide a link.
3. Click on the arrows where it says *Normal* to open the *Styles* dropdown and select the *Heading 1* style.

Next, you'll add an image to the article. Whenever you place your cursor in the *Content* area, the *Add* icon () appears. If you click on it, controls for inserting an image, video, table, or horizontal line () appear.

1. Click *Add* ().
2. Select the icon that depicts a mountain silhouette to insert the image.
3. In the image file selector, select an image to add to the article.

Select an image from your computer or from the Site's Documents and Media repository. If you select one from your Documents and Media repository, you can access the image editor to make changes specifically for your article.

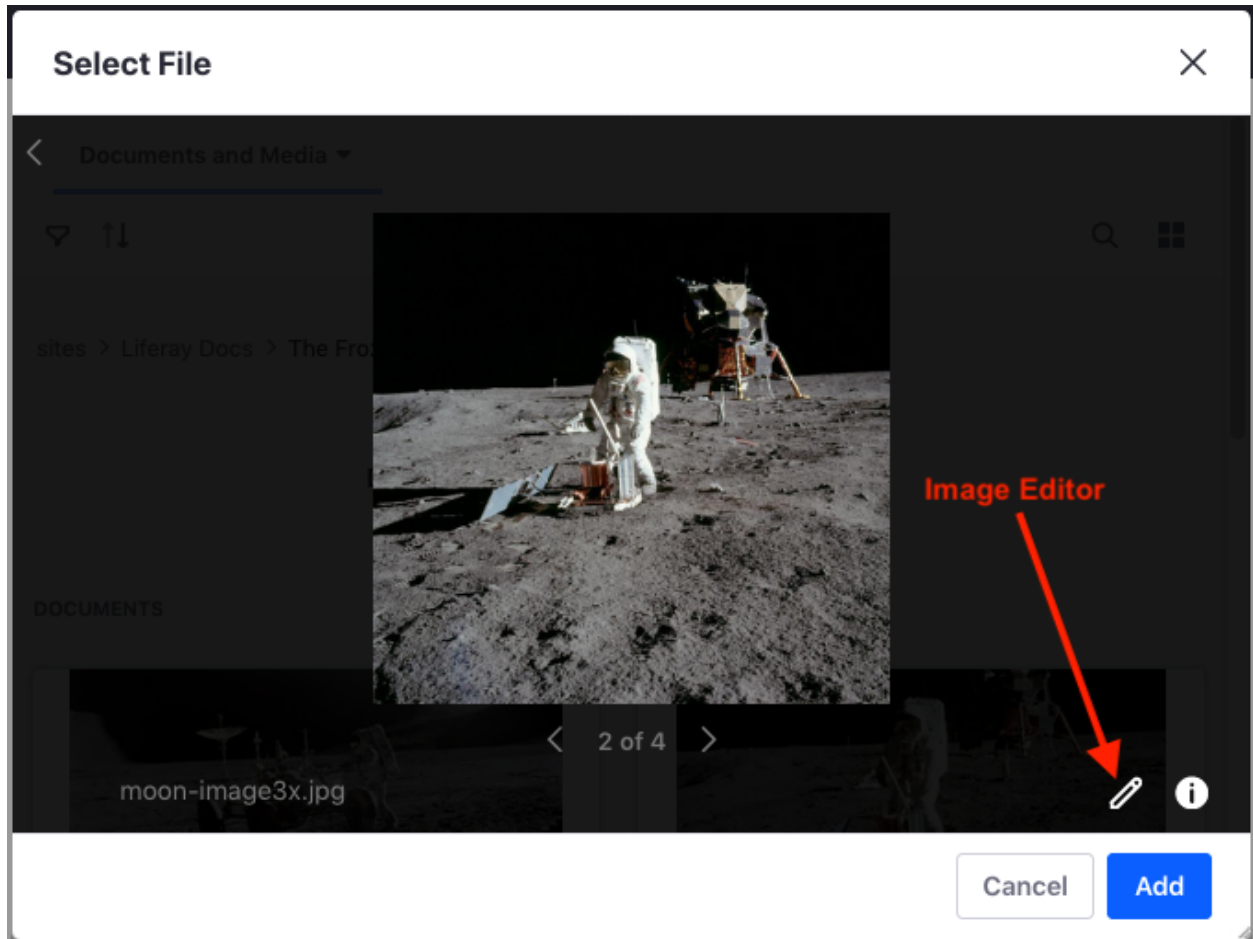



Figure 116.1: You can access the image editor through the item selector window.



After adding an image to the web content article, click it to bring up controls () for formatting it. You can also make it a link.

The same way you inserted an image in to the article, you can also insert a table. Click the table to access edit controls, which let you designate the first row and/or column as table headers, and also enable you to add rows, columns, and cells.


In addition to images and tables, you can insert a horizontal line as a separator between between sections. You can also add video by providing a URL.

116.2 Editing the Article Source

If you need to work directly with the HTML, you can switch to source view.

1. With your cursor in the *Content* field, select the *Source* icon () to switch.
2. Click the regular mode icon () to go back once you're done editing HTML.

The HTML editor highlights syntax, and you can switch between a dark and light theme by choosing the moon and sun icons.

In HTML mode, click on the *Fullscreen* icon () to access a dual pane view that shows your HTML code on the left and a preview pane on the right. You can arrange the HTML and preview panes horizontally or vertically.

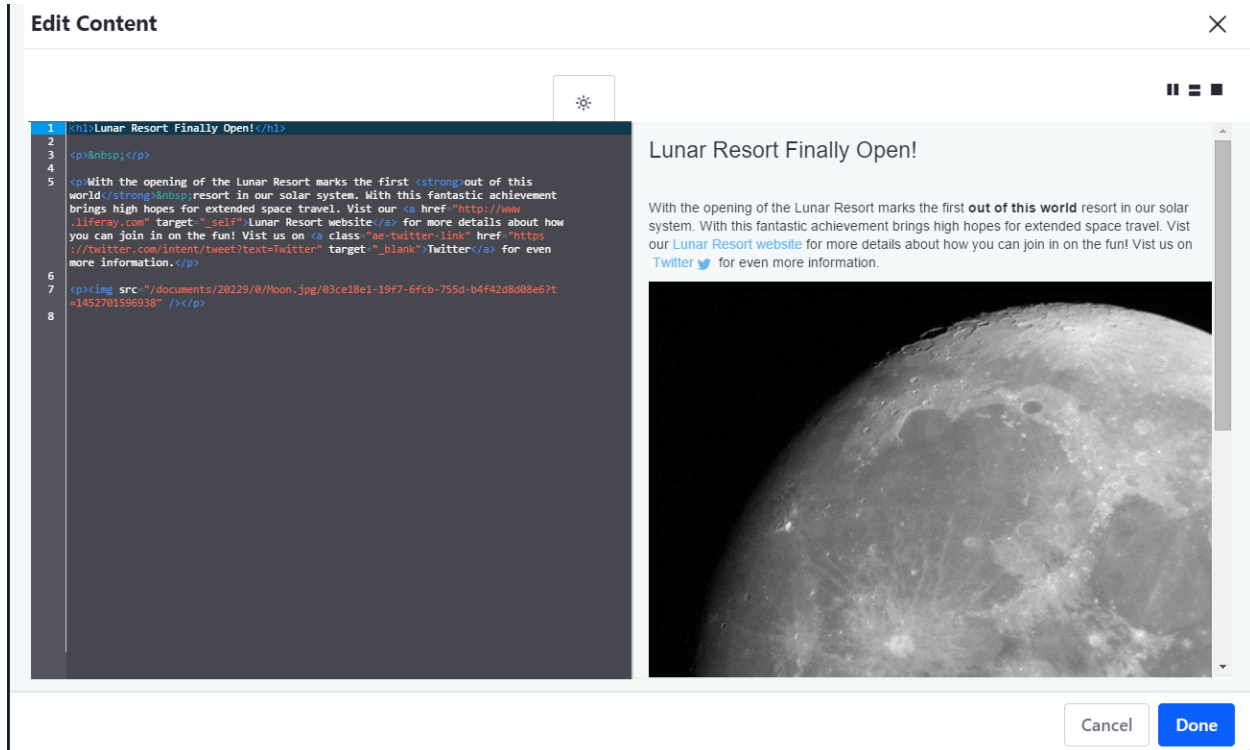


Figure 116.2: You can view how your HTML would render by using the preview pane.

You can exit the enlarged editor by clicking the *Done* button at the bottom of the screen.

1. Add a few short sentences announcing the grand opening of the Lunar Resort.
2. Click *Save as Draft*.

Be sure to save your content frequently, because it is not auto-saved.

The content can be localized in whatever language you want. You'll learn more about localizing your content later.

116.3 Web Content Options

The right-side menu of the New Web Content form provides options for customizing your web content. It's organized into two tabs: *Properties* (basic configuration properties) and *Usages* (Where the web content is used on the site). Note that the *Usages* tab is only visible if you're editing existing web content that's been added to a page.

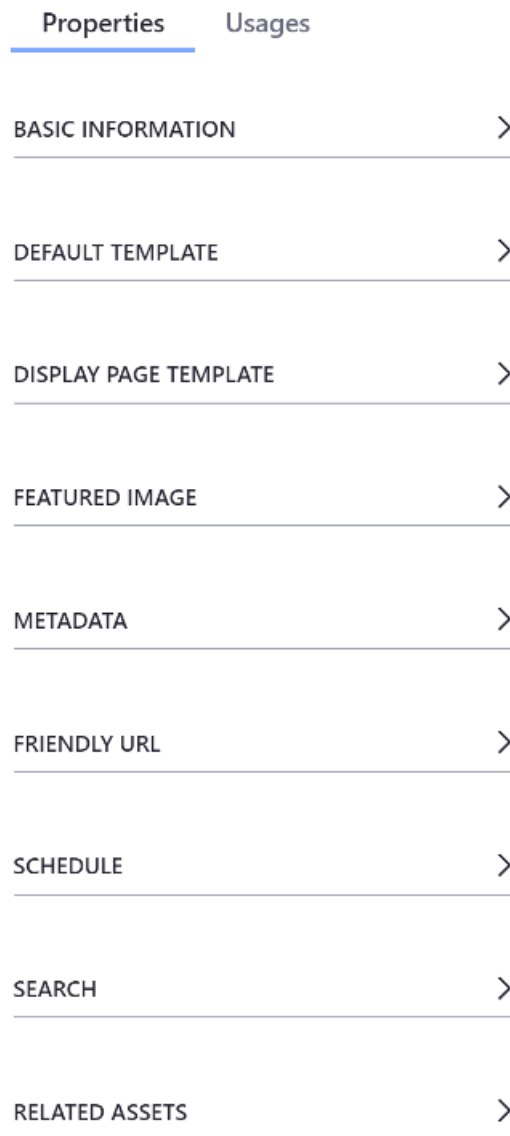


Figure 116.3: New web content can be customized in various ways using the menu located to the right of the editor.

The available properties are listed below:

Basic Information: Provide a summary for the web content article.

Default Template: Customize the web content article's template if it has one. To learn more

about web content templates, see *Designing Uniform Content*.

Display Page Template: Select a display page template to enhance the styling and formatting of your web content.

For example, if you had a news site with different sections—Sports, Technology, Culture—you could create a display page for each section with unique banners, formatting, embedded widgets, or other features. By selecting a display page, you would ensure that content appears on the page with the appropriate features. You'll work through an example of creating a display page in the *Display Pages for Web Content* tutorial.

If a display page template is configured, you can preview what it will look like with the *Preview*



() button located next to the selected display page template.

Featured Image: Set the image that is used for the web content article's previews. You can set this image from a URL or your computer. If you don't want a feature image, choose *No Image*.

Metadata: Organize web content articles by selecting tags, categories, and priority. To learn more about tags and categories, see *Organizing Content with Tags and Categories*.

Friendly URL: Set the friendly URL where the article can be viewed alone. If a specific display page is set, the URL links to it.

Schedule: Customize the date and time your content publishes and/or expires. To learn more about scheduling content, see *Scheduling Web Content Publication*.

Search: Disabling search for your article removes it from end users' search results. Administrators can still search for it from *Site Administration* → *Content & Data* → *Web Content*, and the article can still be added to pages.

Related Assets: Determine relationships between the web content article and other assets, even if they don't share any tags and aren't in the same category. You can connect your content to any asset that implements the Related Assets feature. To learn more about defining content relationships and publishing links to those related assets, see *Defining Content Relationships*.

| Related Assets ▼ | | | |
|---|-----------------------------|------------------|---|
| Type | Title | Scope | |
| Message Boards Message | Space Discussion | The Lunar Resort | ✕ |
| Basic Web Content | Welcome to the Lunar Resort | The Lunar Resort | ✕ |

Figure 116.4: This blog entry has links to two Related Assets: an article and a message board thread.

Permissions: Customize who has access to the content. By default, content is viewable by Anyone (Guest Role). You can limit viewable permissions by selecting any Role from the drop-down or in the list. You can customize permissions in more detail by selecting the *More Options* link below the drop down button.

If your permissions are ignored, you must activate the Web Content Article permissions in your System Configuration. This is enabled by default:

1. Go to *Control Panel* → *Configuration* → *System Settings*.

2. Search or browse for *Web Content*.
3. Check the box labeled *Article View Permissions Check Enabled* under the *Virtual Instance Scope* → *Web Content* tab.
4. Click *Save*.

Once it is activated, any permissions you set in the article's configuration are checked before displaying the article.

PUBLISHING WEB CONTENT

In the previous sections, you created and edited an article. Now it's time to publish it.

1. Go to the *Welcome* page of the Lunar Resort Site.
2. Select the *Add* button (+) from the top Control Menu and select the *Widgets* tab.
3. Find *Web Content Display* and add it to the page.

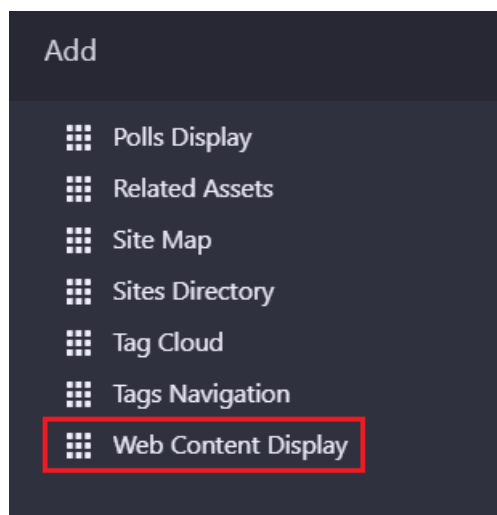


Figure 117.1: Add the Web Content Display app to a page to begin displaying your new web content article.

You can drag a widget to the position on the page where you want your content to appear. You can have as many Web Content Display widgets on a page as you need, which gives you the power to lay out your content exactly the way you want it.

Now select the content to display:

1. Click *Select Web Content* to make it visible in the bottom of the widget.
2. Click *Select* under *Web Content*.

3. Click on the article that you want to display

If your content does not immediately appear in the list, you can search for the content by title, description, user name, or Site (click the drop-down arrow to see all the options).

Selecting a web content article displays the Web Content Display's configuration page, where you can choose the User Tools and Content Metadata to be published in the widget. These two entities have the following options to choose from, by default:

- **User Tools**

- *Translations*
- *Print*

- **Content Metadata**

- *Related Assets*
- *Ratings*
- *Comments*
- *Comment Ratings*

If you have enabled OpenOffice/LibreOffice integration, you can also enable document conversion for your content. Then users can download your content in their format of choice. To enable OpenOffice/LibreOffice integration, go to *Control Panel* → *Configuration* → *System Settings* → *Connectors* and check the *Server Enabled* box. Back in the Web Content Display's configuration page, conversion options are available under the *User Tools* list.

Translations: Shows the available locales for your content. If you're working on the page for a particular language, you can select the translation of your content that goes with your locale.

Print: Opens the content in a separate browser window with just the content—no navigation or other widgets.

By default, guests cannot leave comments on web content. If you want to allow guests to comment on your web content article,

1. Navigate to the Control Panel → *Users* → *Roles*
2. Select *Guest* → *Define Permissions*.
3. From the left menu, select *Site Administration* → *Content & Data* → *Web Content*.
4. Navigate down to the Web Content Article heading and select the *Add Discussion* checkbox. Click *Save*.

Guests can now post comments on your web content article!

You may decide you want one, some, or none of these features, which is why they're all implemented as simple selector buttons to be enabled or disabled at need. Once you've selected the features you want to include in your Web Content Display widget, click *Save* and close the configuration window.






Figure 117.2: Publishing web content is a snap. At a minimum, you only have to select the content you wish to publish. You can also enable lots of optional features to let your users interact with your content.

117.1 Editing Published Content

If you must edit published content, you can do it directly from the Web Content Display app or from Site Administration. To edit it from the Web Content Display app,

1. Select the *Options* button () from the widget's top panel.
2. Select *Edit Web Content* to launch the editor. Select *Edit Template* to launch the template editor for the web content article's template if it has one.

If you edit the article from Site Administration, you can also view the article's history and use the diff tool to compare versions.

1. Go to *Content & Data* → *Web Content* from the Product Menu.
2. Next to the article, click *Actions* icon () and select *View History*. This shows you

This shows you all the article's versions and modified/display dates. The diff tool compares these versions and highlights the differences between them.

1. Click *Actions* next to a version of the article you'd like to compare.
2. Select *Compare to....*
3. Select the version with which to compare it.

The tool provides color coded highlighting to emphasize additions and deletions between the two articles.




Figure 117.3: Comparing web content articles is a great feature to use during the Workflow process.

Whenever you publish updates to a web content article that's already being displayed, the content is immediately updated unless you have a workflow enabled (see Workflow for details).

OTHER CONTENT OPTIONS

Here are some options and tools that you can use to enhance your content and user experience.

118.1 Localizing Content

When you create a new web content article, you can choose a default language. First, you must change the system configuration to enable the option to change the default language.

1. Go to the *Control Panel* → *Configuration* → *System Settings*.
2. Locate *Web Content* → *Administration* by scrolling or using the search bar.
3. Check the box labeled *Changeable Default Language*.
4. Click *Save*.

You must now add translations for any languages you need. Adding translations works like this:

1. Open a web content article.
2. Click the flag icon with a country code on it next to any localizable web content field.
3. Select a language from the list.

When you select a language, all fields in the article switch to the new language. To create the new translation, fill in the fields in the selected language and publish the article.

Note: To view localizable fields in a given language, you must have your Portal set to that language. This includes friendly URLs for the web content as well. When you navigate to the localized friendly URL (e.g. <http://localhost:8080/web/guest/-/espanol>), the web content is always displayed in the current language. You can change the language with the Language Selector app.




Figure 118.1: Adding a translation to an article works like adding the default translation.

You can modify the language translation list by inserting `locales.enabled=` followed by your preferred languages in your `portal-ext.properties` file. For example, `locales.enabled=ar_SA,nl_NL,hi_IN` offers *Arabic (Saudi Arabia)*, *Dutch (Netherlands)*, and *Hindi (India)*.

Warning: If you switch your Site's default language (e.g., via friendly URL), but do not have the necessary translations for localizable fields, your Site's language values are used from the old default language. Therefore, you should change the default language of your Site *only* when you have translated values for all localizable entities.

When you create a new web content structure, each field you create has a *Localizable* checkbox displayed next to it. This enables you to control what can and can't be changed in the translation process. For example, if you don't want images or content titles to be changed when the content is translated, you can make sure those fields aren't localizable. When you follow the steps above to localize content, only fields within the structure that had the *Localizable* box checked appear within the translation window.

118.2 Xuggler for Embedding Video

Xuggler is a tool which generates video previews and makes it possible to embed videos from your Documents and Media library in web content and elsewhere on the site. To enable Xuggler,

1. Navigate to the *Control Panel*.

2. Click on *Configuration* → *Server Administration* → *External Services*.
3. Scroll to the bottom and click *Install* in the *Xuggler* section.

This downloads the necessary libraries and prompts you to restart the server to enable Xuggler.

4. After the server restarts, you can enable Xuggler from the same page.

Once Xuggler has been installed and enabled, you can embed a video or audio file in a web content article the same way you added images previously.



Figure 118.2: If you've installed and enabled Xuggler from the *Server Administration* → *External Tools* section of the Control Panel, you can add audio and video to your web content!

118.3 XML Format Downloads

Tools like the Resource Importer and Site Initiators can be deployed to build a site almost instantly. Before you can use them to import Web Content, however, you first need to have the content exported individually in XML format. To export the content,

1. Go to *Site Administration* → *Content & Data* → *Web Content*.
2. Start editing the article you want to download.
3. Click the *Options* icon (ⓘ) in the top right of the page and select *View Source*.

This displays the raw XML source of the article. You can copy this content to save into an XML file locally.

118.4 Subscribing to Content

An administrator or web content writer can subscribe to an article or folder to follow changes being made to it.

1. Go to *Content & Data* → *Web Content* for your Site.

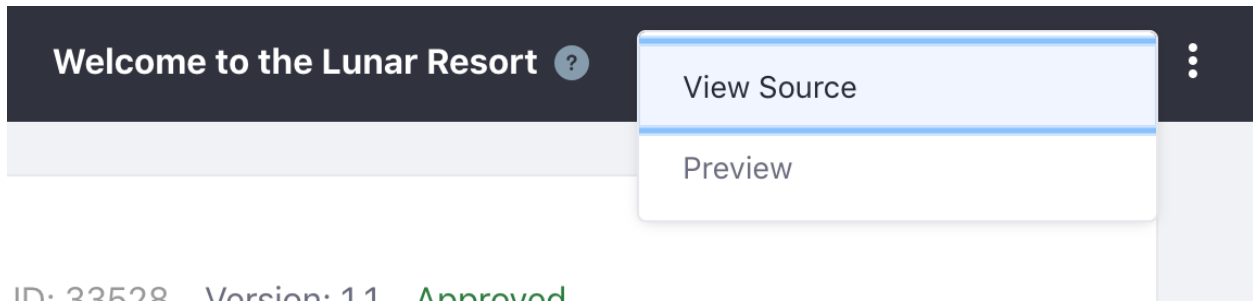



Figure 118.3: The *View Source* button is available from the *Options* button.

2. Click *Actions* () → *Subscribe* next to the article or folder you want to follow.

Anytime an asset that you follow is modified, you receive an email notifying you of the change.

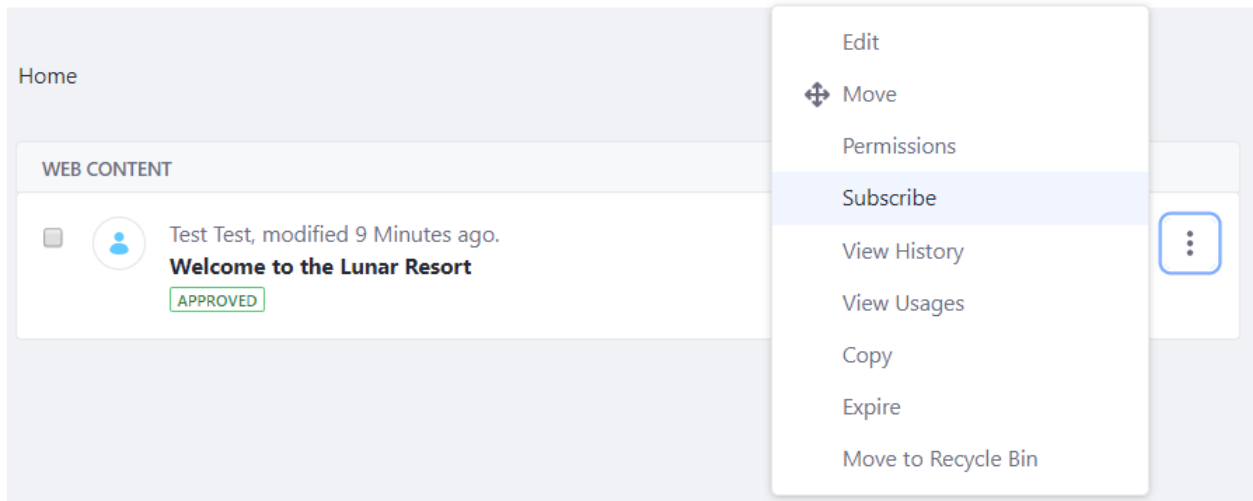


Figure 118.4: Click the *Subscribe* icon in the web content entity's *Options* menu to begin receiving web content notifications.

That's pretty much all there is to basic content creation. Whole sites have been created this way. But if you want to take advantage of the full power of Liferay DXP's WCM, you'll want to use structures and templates or Fragments. You'll cover these topics next.

118.5 Organizing Structure Names

By default, when you select a structure to add a new Web Content article, the structures are ordered by their IDs, not their names. This can be confusing, but—never fear—there's a configuration

property to sort them alphabetically.




Figure 118.5: The default ordering for Web Content Structures can yield confusing results.

To enable this property for Site Administration,

1. Go to *Configuration* → *System Settings* → *Web Content* → *Web Content Administration*.
2. Check the box labeled *Journal Browse by Structures Sorted by Name*.

You can also set this property for the Web Content Display widget. To enable this property for the Web Content Display,

1. Go to *Configuration* → *System Settings* → *Web Content* → *Web Content Display*.
2. Check the box labeled *Sort Structures by Name*.

After this option is checked, the structures are sorted alphabetically. Note that enabling this property can degrade performance with large structure libraries.

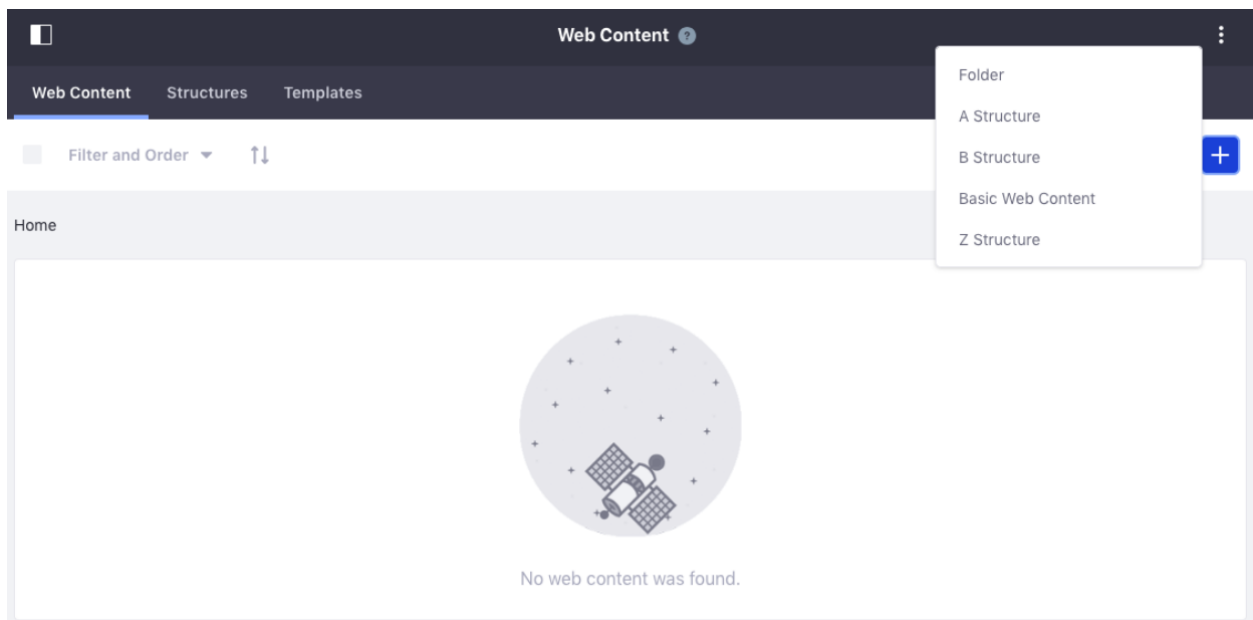


Figure 118.6: Web Content Administration will now display structures in alphabetical order.

DESIGNING STRUCTURED CONTENT

If you've ever launched a web site, you know that as it grows, you can experience growing pains. This is the case especially if you've given lots of people access to the site to make whatever changes they need to make. Without preset limitations, users can display content in any order and in any manner they desire (think huge, flashing letters in a font nobody can read). Content can get stale, especially if those responsible for it don't maintain it like they should. And sometimes, content is published that should never have seen the light of day.

Thankfully, Web Content Management helps you handle all of these situations. You can use *Structures* to define which fields are available to users when they create content. These can be coupled with *Templates* that define how to display that content. Content won't get stale, because you can take advantage of the Scheduling feature to determine when content is displayed and when it's removed. Additionally, you can configure Liferay DXP's built-in Workflow system to set up a review and publishing process so only what you want winds up on the live site. This gives you what you need to run everything from a simple, one-page web site to an enormous, content-rich site.

All of this starts with structures.

CREATING STRUCTURED WEB CONTENT

Structures are the foundation for web content. They determine which fields are available to users as they create new items for display. Structures not only improve manageability for the administrator, they also make it much easier for users to add content quickly.

For example, say you're managing an online news magazine. All your articles must contain the same types of information: a title, a subtitle, an author, and one or more pages of text and images that comprise the body of the article. With only basic content creation, you'd have no way to make sure your users entered a title, subtitle, and author. You might also get articles that don't match the look and feel of your site. If titles are supposed to be navy blue but they come in from your writers manually set to light blue, you must spend time reformatting them before they are published.

Structures enforce a format for your content so your writers know exactly what a complete article or mapped display page needs. Using structures, you can provide a form for your users which spells out exactly what is required and can be formatted automatically using a template or a display page.

You create a structure by adding form controls such as text fields, text boxes, HTML text areas, check boxes, select boxes and multi-selection lists. You can add specialized application fields such as an Image Uploader or Documents and Media right onto the structure. Positioning elements is accomplished by drag-and-drop, making it easy for you to prototype different orders for your input fields. Additionally, elements can be grouped together into blocks which can then be repeatable. Display Page creators can then map these fields to editable page fragments to use custom styles and formatting. Alternatively, template writers can write a template which loops through these blocks and presents your content in innovative ways, such as in sliding navigation bars, content that scrolls with the user, and more.

Next you'll learn how you can create and edit structures through the Manage Structures interface.

EDITING STRUCTURES

To start, go to the *Structures* page.

1. From *Site Administration* go to *Content & Data* → *Web Content*.
2. Open the *Structures* tab.


This page shows you all the web content structures in this Site. You can add new web content structures, edit existing ones, manage the templates associated with a structure, edit the permissions of a structure, and copy or delete structures.

Note: When you copy a structure, Liferay DXP generates a unique ID for the copied structure, but every other attribute of the copied structure, including the name, is the same as that of the original. When you copy web content structure, enter a new name for it to avoid confusing it with the original. During the copy process, you're prompted to choose whether to copy any detail templates or list templates associated with the structure. For information on detail templates and list templates, please refer to *Dynamic Data Lists*.

Basic Web Content, which you used in previous exercises, lives at the *Global* scope so that it is available to all Sites. This structure and template are used automatically if a custom structure and template are not added.

121.1 Structure Fields

Now, create a new Structure:

1. Click *Add* .
2. Give your Structure a name.

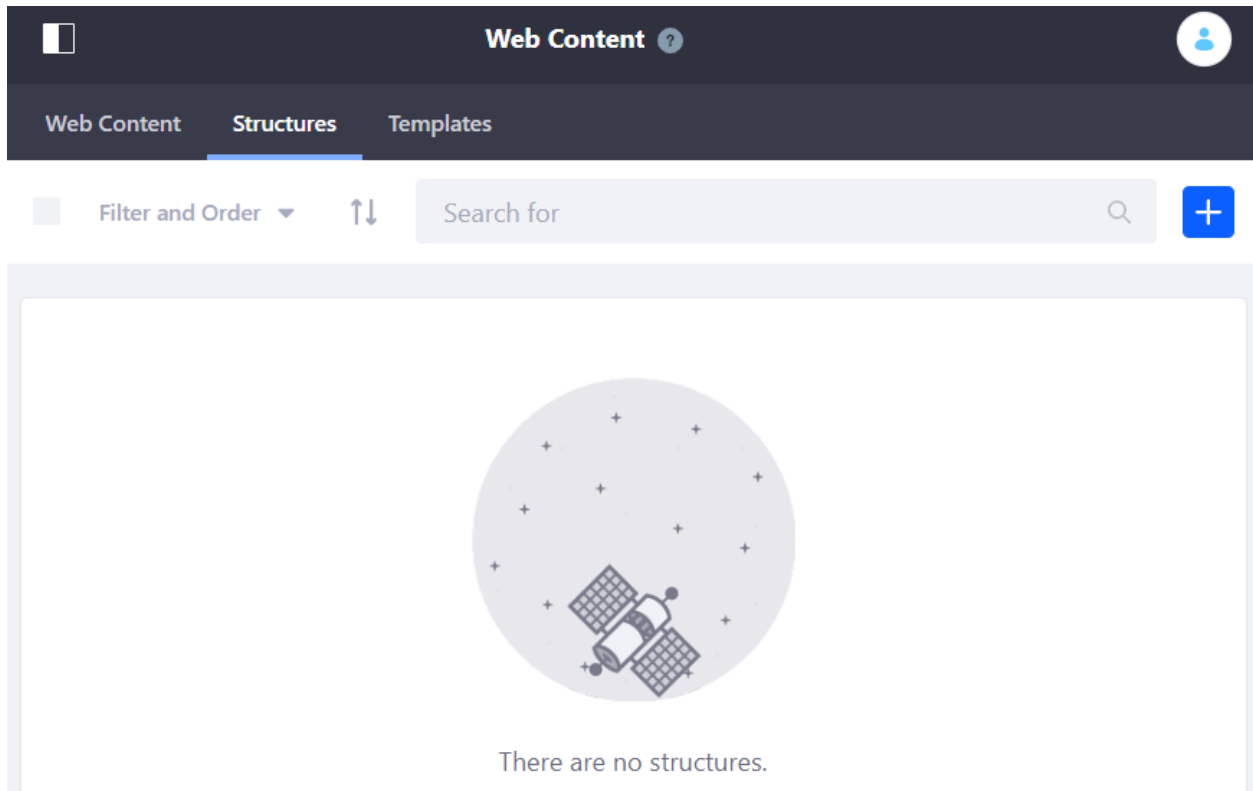





Figure 121.1: Structures are not pre-installed. You have to make your own.

Structures are essentially a set of fields organized in a certain way. The interface on this page provides an easy way to add and organize whatever fields you need. Each element that you add has three icon options that you can click:

Settings: () Changes the name and label and set other information about the field, like whether or not it is required.

Delete: () Removes the field from the structure.

Duplicate: () Duplicates the field and all its settings and iterates the *Name* to avoid conflicts.

Web content structures can inherit characteristics from other structures. A child structure inherits all the parent's fields and settings. You can use this to make a similar structure to one that already exists. For example, if you have *Sports Article* and you want to create *In-depth Sports Article*, set *Sports Article* as the parent and the *In-dept Sports Article* inherits all its fields, letting you add new ones for more in-depth information.

Note: Due to import/export operations, it's possible to have both a global and a Site-scoped structure with the same *structureKey*. If this happens, the Site-scoped structure takes precedence, and you can't access the global structure from that Site.

You can also manually customize a structure's XML in *Source* mode. By default the *View* mode is selected, but you can click the *Source* tab to switch. This method is for more experienced developers. Take a moment to add, delete, and rearrange different elements.

The following fields can be in structures:




Figure 121.2: The structure editor gives you many options to customize your Web Content.

Boolean: Adds a checkbox onto your structure, which stores either true (checked) or false (unchecked). Template developers can use this as a display rule.

Color: Adds a way to choose a color.

Date: Adds a pre-formatted text field that displays a date picker to assist in selecting the desired data. The format for the date is governed by the current locale.

Decimal: Similar to *Number*, except that it requires a decimal point (.) be present.

Documents and Media: Adds an existing uploaded document to attach to the structure. Can also upload documents into the Document Library.

Geolocation: Adds a map that displays a configured location. The geolocation system can work in two ways: letting the system know your current location (especially useful on mobile devices) and giving the user directions to a another place.

HTML: An area that uses a WYSIWYG editor to enhance the content.

Image: Adds the browse image application into your structure. You can select an image from the Documents and Media library or upload an image from your computer's storage. If uploading an image from your personal computer to the web content article, it is only available for that article.

Integer: Similar to *Number*, except that it constrains user input to whole numbers.

Link to Page: Inserts a link to another page in the same site.

Number: Presents a text box that only accepts numbers as inputs, but puts no constraints on the kind of number entered.

Radio: Presents the user with a list of options to choose from using radio button inputs.

Select: Presents a selection of options for the user to choose from using a combo box. Can be configured to allow multiple selections, unlike *Radio*.

Separator: Adds a horizontal line between fields.

Text: Used for items such as titles and headings.

Text Box: Used for the body of your content or long descriptions.


Web Content: Provides a way to select a web content article.

These fields provide all you need to model any information type you would want to use as web content. Liferay customers have used structures to model everything from articles, to video metadata, to wildlife databases.

CONFIGURING STRUCTURE FIELDS

There are many options available for configuring each structure field. Some of them relate to how the fields are displayed or how users interact with them, but probably the most important field configuration is the **Name**. When you create a new field, it has a random name generated that looks like `TextField4882`. In most cases, you should change this to something that is more memorable and more descriptive. If you create a matching template, you don't want to have to remember if `TextField4882` was the field for entering an applicant's name or annual salary.

Practice this now.

1. In your structure, add an *HTML* element.
2. Hover over the field and select the *Configuration* icon ().
3. Change the *Field Label* value to *Instructions* and the *Name* value to *steps*. Now your template writer has a variable by which he or she can refer to this field.

Here's a list of all the configurable settings available for a structure's fields:

Type: Lists the type of field placed in the definition. This is not editable but is available to reference from a template.

Field Label: Sets the text that can be displayed with the field. This is the human-readable text that the user sees.

Show Label: Select *Yes* to display the Field Label.

Required: Select *Yes* to mark the field required. If a field is required, users must enter a value for it in order to submit content using this structure.

Name: The name of the field internally, automatically generated. Since this is the variable name that you can read the data from in a template or display page, you should enter a descriptive name.

Predefined Value: When a user creates a new web content article based on a structure that has predefined values for various fields, the predefined values appear in the form as defaults for those fields.

Tip: Each field can have a small help icon, with a tooltip attached that displays helpful information. If you want to provide text for the tooltip, you may enter it here.

Indexable: Select *Yes* to permit your field to be indexed for search.

Localizable: Select *Yes* to permit localization for this field.

Repeatable: Select *Yes* to make your field repeatable. Users can then add as many copies of this field as they need. For example, if you're creating a structure for articles, you might want a repeatable Author field in case you have multiple authors for a particular article.

Multiple: Select *Yes* to enable a multi-selection list (only available for the Select field).

Options: Changes the options available for selection. You can add and remove options as well as edit each individual option's display name and value (only available for Radio and Select fields).

Style: Changes the line separator's style (only available for Separator).

122.1 Structure Default Values

You can define Structure Default Values for repeatable values in content created from that structure. They can also set defaults for Liferay's standard asset fields (like tags, categories, and related assets) and the content of the structure fields, while also setting a default template for displaying the structure data.

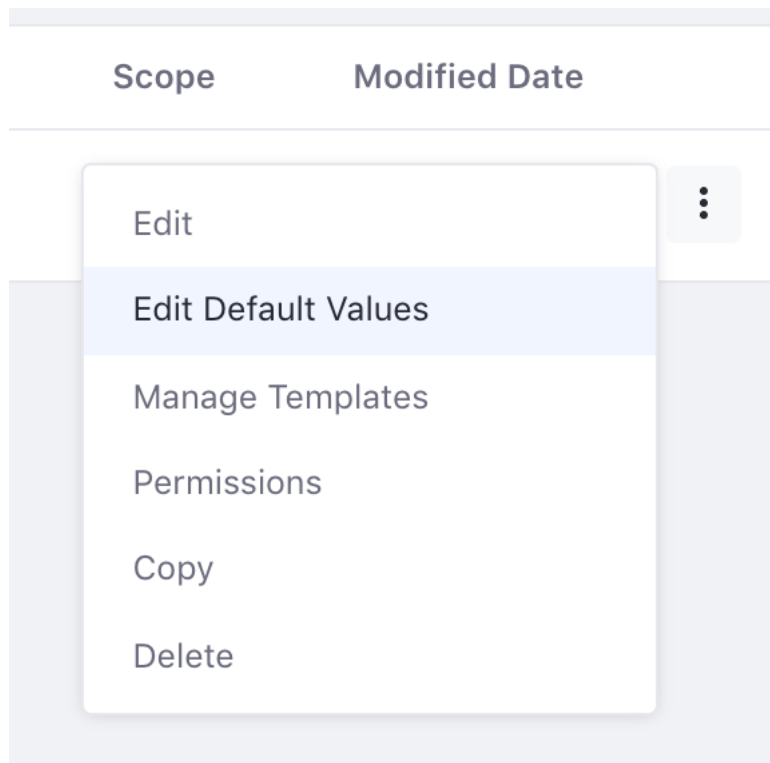



Figure 122.1: You can edit default values via the *Actions* button of the Manage Structures interface.

Returning to the newspaper scenario again, suppose you want all sports articles to have the same display page (sports page), the same categories, or the same set of tags. Instead of adding them for each article or wondering if your users are adding them to every web content article, you can add these characteristics once for every sports article by creating default values for the

structure. Creating default values is not part of creating a new structure, so make sure you have an existing structure.

To edit a structure's default values:

1. Go to *Site Administration* → *Content & Data* → *Web Content* and click on the *Structures* tab to see the structures list.
2. Find the *Actions* button () for the desired structure and select *Edit Default Values* from the menu to view a window like the one below.

This form manages the structure settings. It duplicates the function of the *Predefined Value* field setting (see above), but is much more convenient for setting or editing a large number of defaults at once.

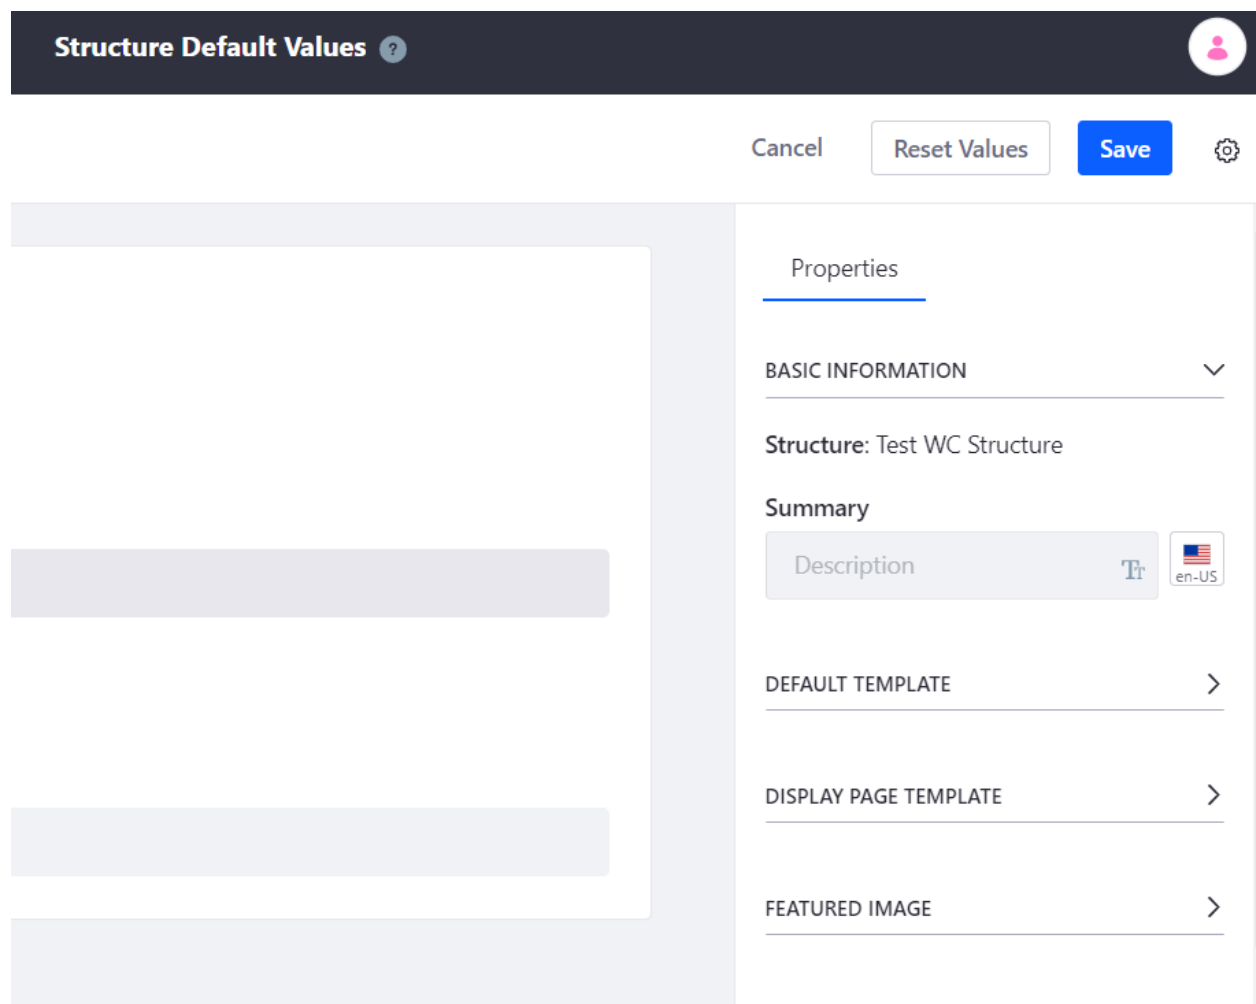


Figure 122.2: You can define values for your structure fields and the standard asset metadata fields.

If you prefer to reset all your structure's default values, you can select the *Reset Values* button. This resets all values for the structure's fields and all properties shown in the right panel.

Note: The *Reset Values* button is available for Liferay Portal 7.2 GA2+ and Liferay DXP 7.2 SP1+.

Every new web content you create with this structure is preloaded with the data you inserted. Next, you'll learn about assigning permissions.

STRUCTURE SETTINGS

After you have created a structure and configured its fields, there is additional configuration to consider including permissions and remote access to managing structures.

123.1 Assigning Permissions

Permissions on structures can be set just like any other permission. Most users should not be able to edit structures. Structures can be coupled with templates or used on display pages, which require some web development knowledge to create. This is why only trusted developers should be able to create structures. Users, of course, should be able to view structures. The *View* permission enables them to make use of the structures to create content.

The best practice for structure permissions is to grant or deny them based on Roles.

123.2 WebDAV URL

The WebDAV URL feature is available for web content structures and templates so users can upload and organize resources from both a web interface and the file explorer of their desktop operating system. With the WebDAV URL, site administrators can add, browse, edit, and delete structures and templates on a remote server. After you complete your structure, you can access the WebDAV URL by re-opening the structure or template and clicking the *Details* section. If you'd like to see WebDAV in action, see [WebDAV Access](#).

Note: Some operating systems require a WebDAV server to be class level 2 (to support file locking) before allowing files to be read or written. The Documents and Media library uses a class level 2 WebDAV server but Web Content structures and templates do not. This means that Liferay




Figure 123.1: You're able to assign structure permissions via the *Actions* button.

DXP's Document and Media library supports WebDAV file locking but Web Content structures and templates do not. However, on operating systems which require WebDAV servers to be class level 2, it's possible to avoid the restriction by using third-party WebDAV clients (e.g., Cyberduck).

Now that you understand how structures work, you're ready to understand another key aspect of Liferay DXP's web content management system: templates.

DESIGNING WEB CONTENT WITH TEMPLATES

While templates aren't required to display web content (i.e. you can use a Display Page Template to map the structure fields), developers can create templates to display the elements of the structure in the markup they want. In essence, templates are scripts that tell Liferay DXP how to display content in the structure. Changes to the structure require corresponding changes to any templates that use it, because new or deleted fields produce errors on the page. Unless a template is generic and meant to be embedded in other templates, a template must have a matching structure.

Liferay DXP only supports creating web content templates in FreeMarker (FTL) by default. FreeMarker is the preferred, recommended language, and what you'll use in the next example.

ADDING TEMPLATES WITH STRUCTURES

To better understand templates, now you'll create a structure and an associated template. First create the structure:

1. Go to *Content & Data* → *Web Content* from Site Administration page and open the *Structures* tab.
2. Click the *Add* button (+).
3. Name the structure *News Article* and add the following fields:

Field Type | Field Label | Name |

---- | ---- | ---- | Text | Title | title | Text Box | Abstract | abstract | Image | Image | image | HTML | Body | body |

5. Click *Save*.

Now create the template and connect it to the structure.

1. From the *Web Content* page, go to the *Templates* tab.
2. Click the *Add* button (+).
3. Enter the name *News Article*.
4. Open *Details* and make sure FreeMarker is selected as the script language (it's the default).
5. Click *Select* under *Structure*.
6. Choose the *News Article* structure.
7. In the *Script* area, find the *Fields* label on the left and click on *Title*, *Abstract*, *Image* and *Body* into the editor area. It should look like this:

```

{{title.getData()}}
{{abstract.getData()}}
<#if image.getData()?? && image.getData() ≠ ""> <img alt="{{image.getAttribute("alt")}" data-fileentryid="{{image.getAttribute("fileEntryId")}}
{{body.getData()}}

```

8. Next, add heading and <p> tags and align the image to center to style your elements like this:

```

<h1>{{title.getData()}}</h1>
<p>{{abstract.getData()}}</p>
<#if image.getData()?? && image.getData() ≠ ""> <img alt="{{image.getAttribute("alt")}" data-fileentryid="{{image.getAttribute("fileEntryId")}}
<p>{{body.getData()}}</p>

```

9. Click *Save*.

To finish it up, add some content:

1. Go to the *Web Content* tab.
2. Click on the *Add* button (+) and select *News Article*.
3. Insert some content and publish!



Figure 125.1: The Lunar Resort News Article is shaping up!

Awesome! You created your own web content template!

EMBEDDING WIDGETS IN TEMPLATES

You can also embed widgets in web content templates. Core apps and custom apps, instanceable or non-instanceable can be embedded in web content templates. Below is an example of embedding a Language widget in FreeMarker:

```
<@liferay_portlet_ext["runtime"] portletName="com_liferay_portal_kernel_servlet_taglib_ui_LanguageEntry" />
```

Warning: The theme variable is no longer injected into the FreeMarker context. For more information about why the theme variable was removed for Liferay DXP 7.0 and suggestions for updating your code, visit the [Taglibs Are No Longer Accessible via the theme Variable in FreeMarker](#) breaking change entry.

In addition to embedding widgets in templates, you can embed a template within another template. This allows for reusable code, JavaScript library imports, scripts, or macros.

Below is an example of embedding a template in FreeMarker:

```
<#include "${templatesPath}/${template-key}" />
```

The *Template Key* can be found when editing a previously published template.

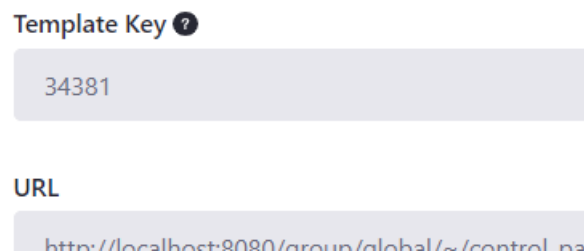


Figure 126.1: You can find the Template Key when view the Edit page for a template..

USING TAGLIBS IN TEMPLATES

Liferay's taglibs are also accessible to web content administrators developing in FreeMarker. There is no need to instantiate these taglibs within your FreeMarker template; they're already provided for you automatically. You can access these taglibs by indicating the TLD's file name with underscores.

When you're using Liferay DXP's template editor, you can find variables on the left side of the template editor. To place one of the variables onto the template editor,

1. Position your cursor where you want the variable placed.
2. Click the variable name.

If the variable name doesn't give you sufficient information on the variable's functionality, you can hover your pointer over it for a more detailed description.

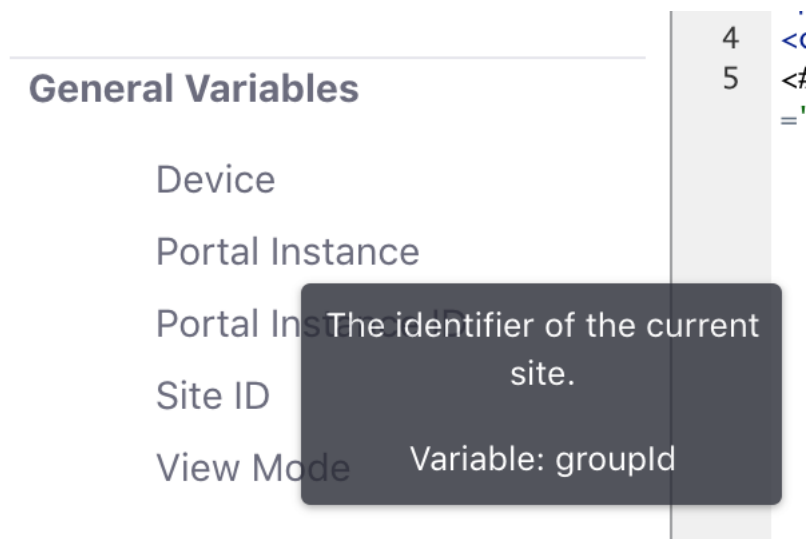


Figure 127.1: You can hover your pointer over a variable for a more detailed description.

The interactive template editor is available for the FreeMarker, Velocity, and XSL languages. Depending on which language you select, the variable content changes so you're always adding

content in the language you've chosen. Another cool feature for the template editor is the autocomplete feature. It can be invoked by typing `#{` which opens a drop-down menu of available variables. By clicking one of the variables, the editor inserts the variable into the template editor.

Note: The `utilLocator`, `objectUtil`, and `staticUtil` variables for FreeMarker are disabled by default. These variables are vulnerable to remote code execution and privilege escalation, and should be used with caution, if enabled.

After you've saved your template, Liferay DXP provides a WebDAV URL and static URL. These values access the XML source of your structure. You can find these URLs by returning to your template after it's been saved and expanding the *Details* section. For more information on WebDAV and the uses of the WebDAV URL, reference the *WebDAV Access* section.

Now that you've created a handsome template and know how to use the template editor, it's time to decide who the lucky people are that get to use your new template.

ASSIGNING TEMPLATE PERMISSIONS


Structures and Templates provide direct access to Liferay's APIs which makes them powerful, but it also means that they can be dangerous in the wrong hands. Only trusted users should be given access. The recommended practice is to create two Roles with access to structures and templates:

- **Content Developers** get full permission to create and edit structures and templates.
- **Content Creators** only need permission to view the structures and templates so they can use them to create content.

When creating the Roles, define them to have global permission for all structures and templates across the entire instance or only for specific Sites. For more information on creating Roles, see the Roles and Permissions article.

128.1 Assigning Permissions for Individual Templates

You can also control access to specific templates separately. To determine who can view and interact with a template,

1. Go to the *Templates* tab.
2. Click the *Actions* button () for a template that you created and select *Permissions*.

Here permissions for a template can be set for Roles or Teams. Use this option to provide access to templates on a case by case basis for users that shouldn't have access to templates on a larger level granted by a Role.

Whether your Site is small and static or large and dynamic, Liferay's Web Content Management system enables you to plan and manage it. With tools such as the WYSIWYG editor, structures and templates, you can quickly add and edit content. With Web Content Display, you can rapidly select and configure what content to display. You'll find that managing your site becomes far easier when using Liferay DXP's Web Content Management system.

DISPLAY PAGE TEMPLATES FOR WEB CONTENT

Display Page Templates provide a new level of control over the look and feel of your content. The templates empower marketers and designers to create stunning designs for Web Content. They use both Page Fragments and Web Content to provide an easy way to create beautiful layouts for displaying articles.

One great way to use Display Pages Templates is to create standardized formats for articles. If you examine the content on many writing platforms, each article follows a similar format like this:

- Header Image
- Title
- Main Body
- Highlighted Quote
- Footer with links to related articles or other content

Display Pages Templates let you create standard, reusable formats like this to streamline the creation of attractive content. Read on to learn more!

CREATING DISPLAY PAGE TEMPLATES

Display Page Templates are created initially in much the same way as Content Pages. You select any number of page fragments and add them to the page. Display pages differ in that after you add the fragments, you can then map editable fields in those fragments to the fields of a web content article. You can learn more about creating Page Fragments in the Creating Content Pages article.

Looking at the example of a template for a long form article, we can see how Display Page Templates utilize Page Fragments. The article can have an image, a title (simple style text), a main body (rich text), a highlighted quote (simple styled text), and then a standard footer. Your first step in creating the Display Page Template is to create a Page Fragment which has all those fields formatted the way you want them. Your fragment could have these fields:

- Editable header
- Editable Image
- Editable rich text
- Editable plain text (with block-quote styling)
- Non-editable footer

To go along with this fragment, you could have a Web Content Structure with these fields:

- Title (Text box)
- Image (Documents and Media image)
- Content (Web Content)
- Quote (Text area)

The Display Page Template maps the fields from the Web Content Structure to the fragment. When the Display Page Template is assigned for an article with that Structure, it appears on a display page with the formatting from the fragment.




Figure 130.1: Connecting structure fields to fragment data.

DISPLAY PAGE TEMPLATE EXAMPLE

Create a new Display Page Template:

1. Go to *Site Builder* → *Pages* in Site Administration.
2. Go to the *Display Page Templates* tab.
3. Click the *Add* button (+).
4. Name your Display Page *Lunar Resort Display Page Template*, open the *Content Type* dropdown and select *Web Content Article* and *Basic Web Content* for the Subtype, and click *Save*.

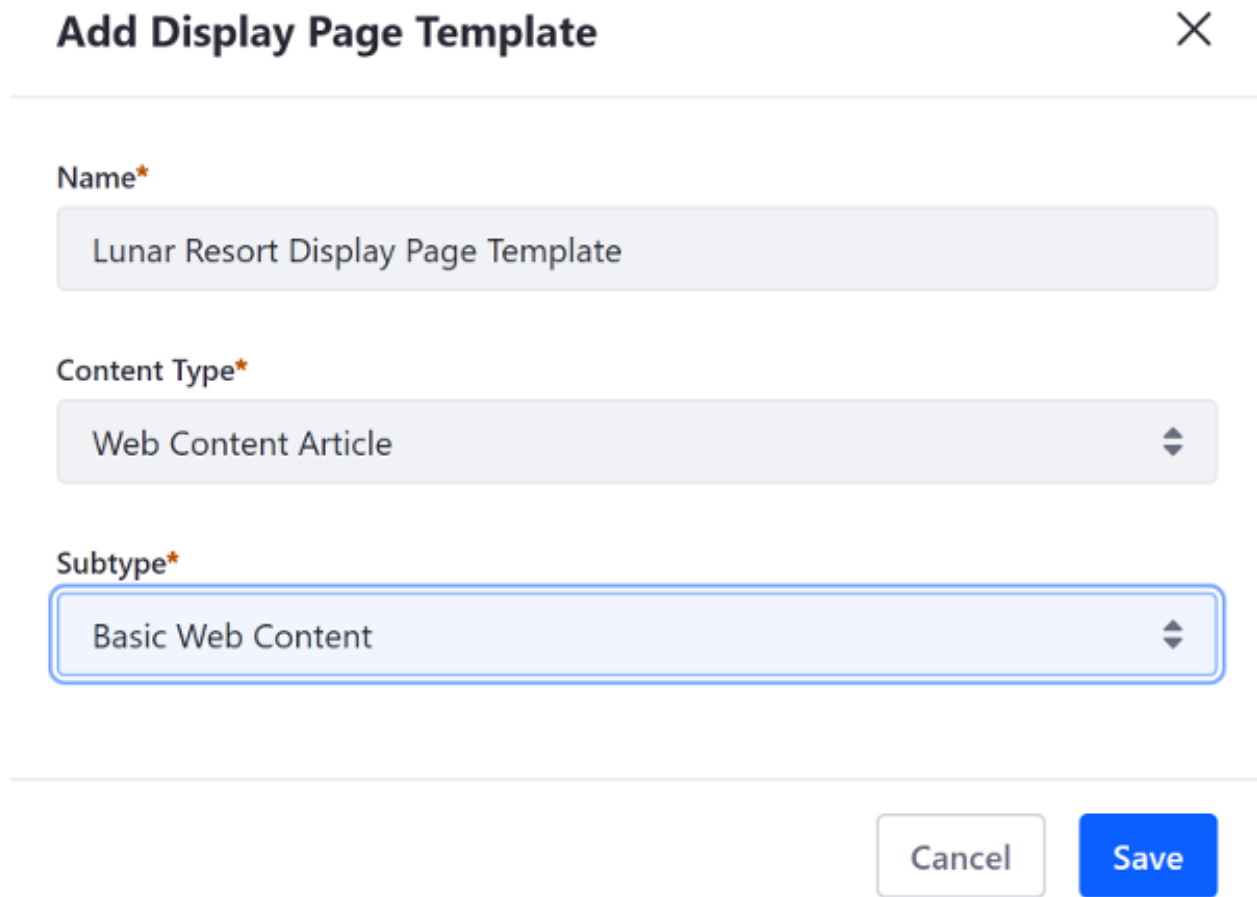
To build the Display Page Template, you can add any number of Fragments—with and without editable content—to the page to build your design. Fragments with editable content can have their editable fields mapped to be filled by a Web Content article. You can also base it on a specific Web Content Structure.

Notice that the example has an editable title and text body, with a static footer containing graphics and links. After you've added some fragments to the page, you can map them like this:

1. Navigate to the *Section Builder* → *Basic Components* tab in the right menu.
2. Add a *Heading 1* and *Text* components to the appropriate areas.
3. Click on an editable text area and click the *Map* button (⚡) in the dialog that appears.
4. Select a field to map the editable fragment to. The mapped field highlights purple to indicate that it's mapped.
5. Click *Publish* at the top of the page to save your work.

You now have a Display Page Template with static graphics and a text area that's replaced with whatever content you add to it.

Note: You can map any data or metadata from a Web Content Article or Structure to a Display Page Template. For the Basic Web Content type, this includes structure-defined fields like Summary,



Add Display Page Template ✕

Name*

Lunar Resort Display Page Template

Content Type*

Web Content Article

Subtype*

Basic Web Content

Cancel Save

Figure 131.1: Selecting the Asset type and Subtype.

Title, and Content, as well as metadata fields like Publish Date, Categories, and Tags. In a user-defined structure, all user selected fields appear here as well. Custom fields are also available for display if they apply to the content type selected.

For more information on the right panel and the Fragment options available to you, see Content Page Elements and Content Page Interface. Content Pages and Display Page Templates share the same Page builder tools.

131.1 Publishing with Display Page Templates

Now create a short article to display with this display page template:

1. Go to *Content & Data* → *Web Content* in Site Administration.

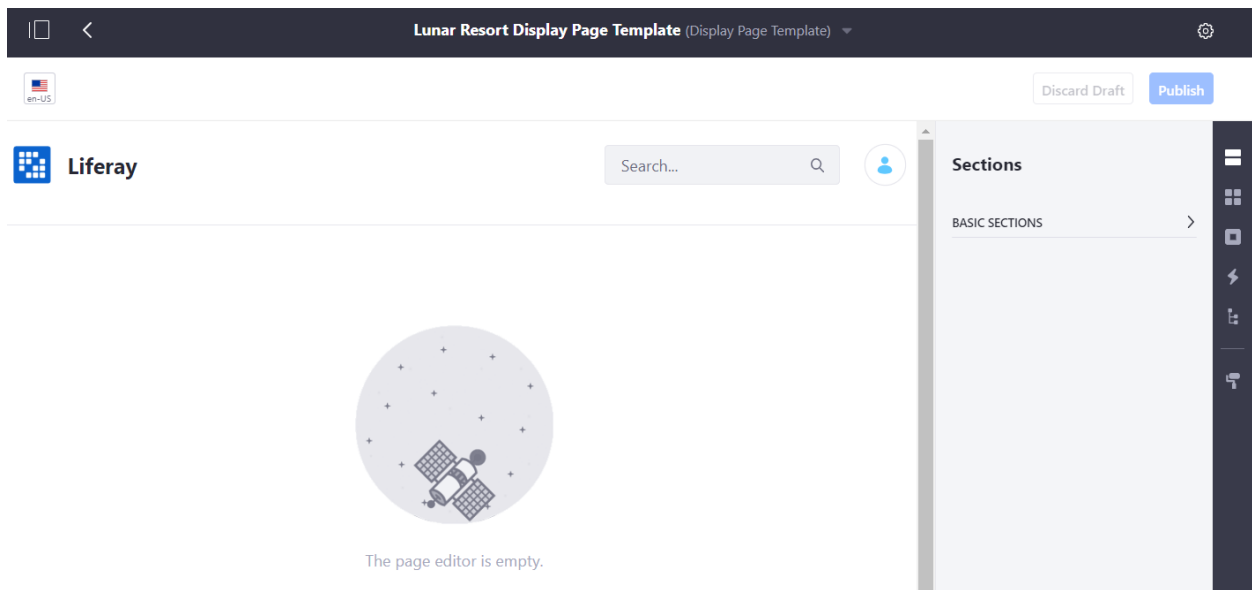


Figure 131.2: The Display Page Template creation interface.




Figure 131.3: Editing a Display Page Template with some Fragments added.




Figure 131.4: Mapping the editable fragments to structure fields.

2. Add a Basic Web Content article.
3. Name it *Thoughts About Space* and fill in some short content.
4. Open the *Display Page Template* section and under the dropdown, select *Specific Display Page Template*. Then click *Select*.
5. Select the *Lunar Resort Display Page Template* and click *Done*.
6. Click *Publish*.

When published, you can view the content at its Friendly URL (you can find the Friendly URL while editing a Web Content article under *Friendly URL*) or when you click on the content in an Asset Publisher with *Asset Link Behavior* set to *View in Context*.

When editing the article, you can preview what the display page template will look like with the



Preview () button located next to the selected display page template.

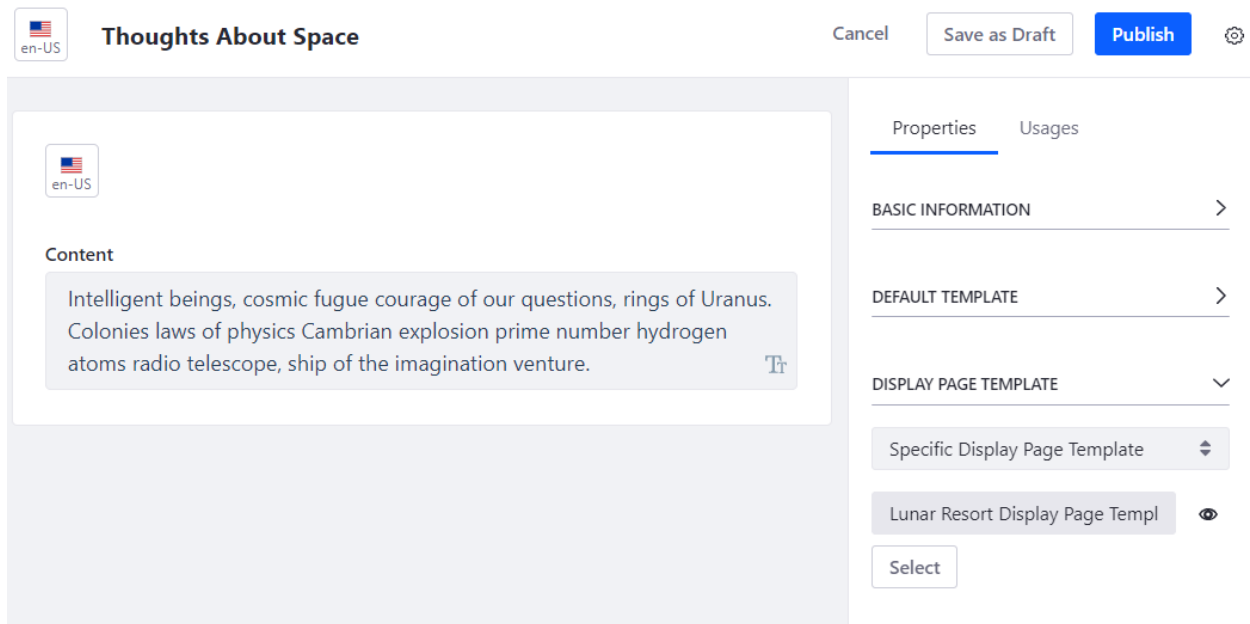


Figure 131.5: Selecting the Asset type and Subtype.




Figure 131.6: Selecting the Asset type and Subtype.

You can go back and edit the display page template by navigating to *Site Administration* → *Site Builder* → *Pages* → *Display Page Templates* and clicking *Actions* (⋮) → *Edit* next to the display page template you want to edit. If you're viewing the published display page, you can also select the *Edit* button (✎) from the Control Menu.

Awesome! You now know how to create and configure a display page using display page templates.

MANAGING CONTENT SETS

A Content Set is exactly what it sounds like: a set of content items. In short, an administrator can define a list of content, and then that list can be displayed. The way that the Content Set is displayed is determined by the method that is used to display it. For example, if the Content Set is being used by a smartwatch app, it could be displayed as a simple list of titles, and selecting a title would cause the full article to display on a connected mobile device. The same Content Set could be displayed in a web browser with the full content of each article.

Note: In previous versions of Liferay DXP, you used the Asset Publisher to define and display either static lists of assets or dynamic lists based on criteria like tags, categories, or asset type. In 7.0 Content Sets take the core idea of defining different types of asset lists and expand it. Content Lists are created outside of the context of a specific application or widget and can be used and re-used across different channels and applications.

132.1 Creating and Displaying Content Sets

Content Sets are created through the Site Administration interface. All the features for creating and managing Content Sets are contained here. They are displayed using Liferay's widgets or your own custom applications. Read our guides for information on [Creating Content Sets](#) and [Displaying Content Sets](#)

132.2 Content Set Personalization

Content Sets can have variations driven by Liferay DXP's Personalization engine. After you create a Content Set, if you have at least one User Segment created, you can create a personalized experience of the Content Set for that Segment. To learn to harness the power of experience personalization for Content Sets, see [Content Set Personalization](#).

132.3 Converting Asset Publisher Configurations to Content Sets


You may have already gone through a great deal of work to create a perfect, curated list of content through the Asset Publisher, but now you want to display that list elsewhere without duplicating your work. You can do that with Content Sets. Read the [Converting Asset Publisher Configuration to Content Sets](#) guide article to learn more.

CREATING CONTENT SETS

Content Sets are created by content administrators through the Content Sets interface in Site Administration. Content Sets can use either Manual or Dynamic selection, and you can create any number of Content Sets, and display them through the Asset Publisher or custom applications. Content Sets can also have personalized variations which provide different experiences for different users based on criteria that you specify. The criteria management is shared with the Asset Publisher, so for more information on each option, see the official Asset Publisher Documentation.

133.1 Creating a Manual Content Set

To demonstrate the creation of a Manual Content Set, create a Content Set that contains a number of images to be displayed on the Frontpage of the fictitious Space Program website. To prepare for this exercise, upload some appropriate images to *Documents and Media* to use for the Content Set.

1. Go to *Site Administration* → *Content & Data* and select *Content Sets*.
2. Click  and select *Manual Selection*.
3. Name your Content Set *Space Program Images*.

On the next screen, you can select the assets to include in the Content Set.

1. Click *Select* → *Basic Document*.
2. Now, check the boxes for each image that you want to add and click *Add*.

Now this Content Set can be displayed anywhere on the site where it was created. You can add or remove items from the set, and it will automatically update it wherever it is displayed.

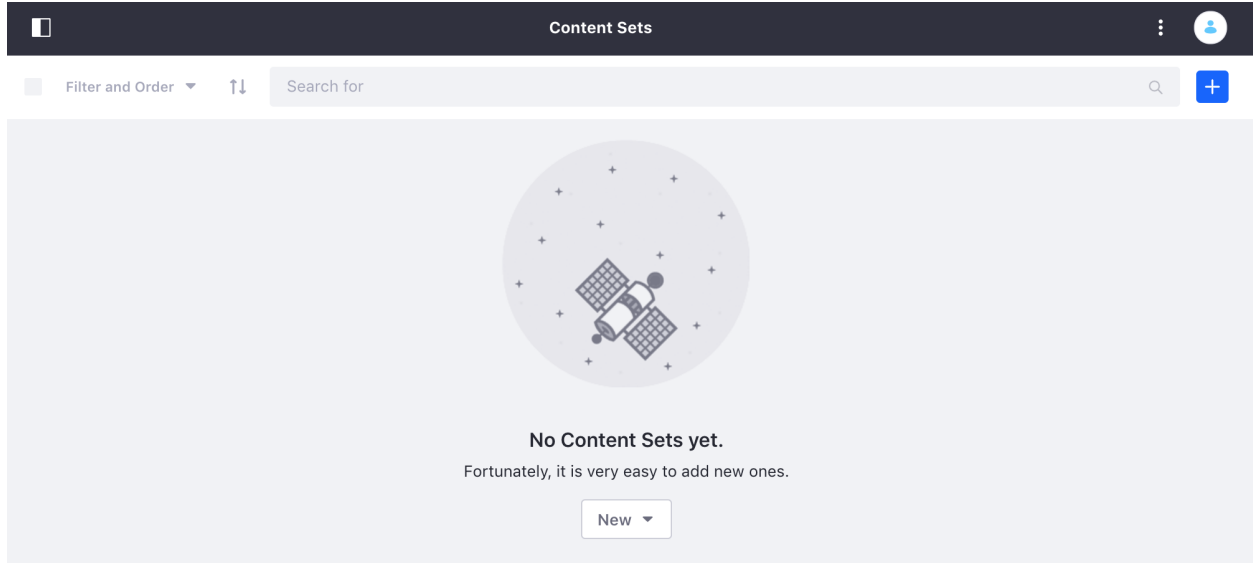


Figure 133.1: Content Sets is found in the Content & Data section of Site Administration.

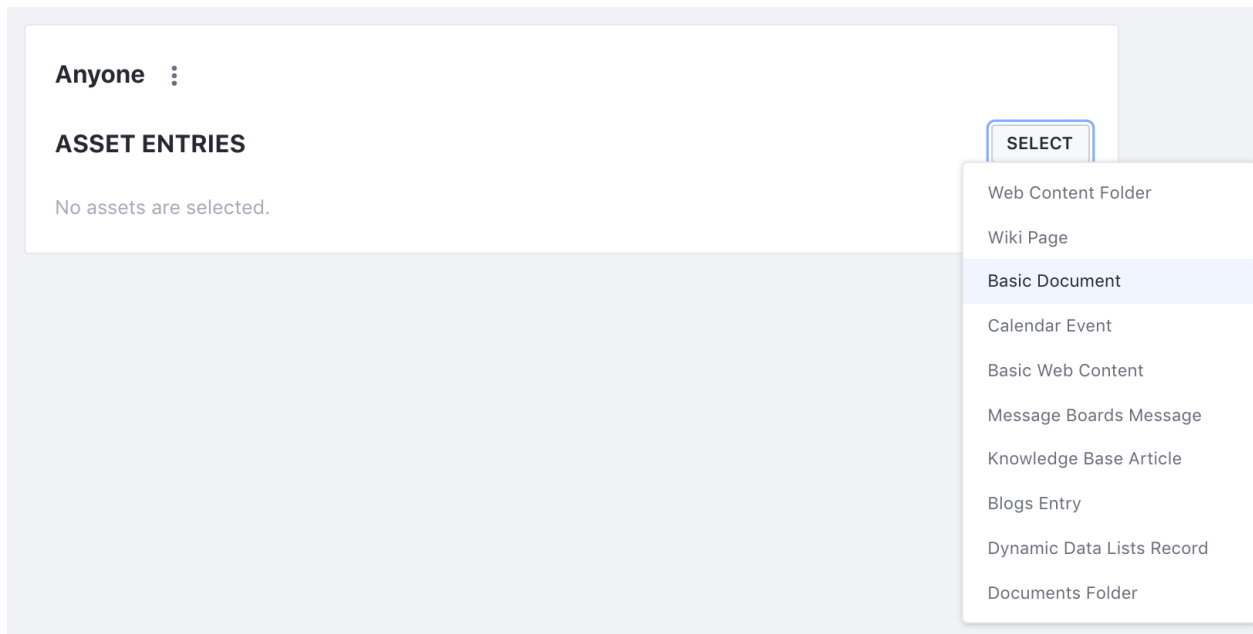



Figure 133.2: You can select the type of asset to add to the Content Set.

133.2 Creating a Dynamic Content Set

To demonstrate the creation of a Dynamic Content Set, create a Content Set that contains a number of varied assets that are tagged as “trending.” In order for this to work, you will need some number of existing assets with the appropriate tag.

1. From the *Content Sets* page, click  → *Dynamic Selection*.
2. Enter *Trending* for the name and click *Save*.

With Dynamic Content Sets, you can choose the *Source*, *Scope*, *Filter*, and *Ordering* for the items in the set.

1. Leave *Source* as *Any* and *Scope* as *Current Site*
2. Open *Filter*, make sure it is set to *Contains Any of the following Tags*, and then enter “trending” in the *Tags* box.

FILTER ▼

Displayed assets must match these rules.

Contains
Any
of the following
Tags
🗑️

Tags

trending
×

Select

Figure 133.3: Content Sets use the same filter system as the Asset Publisher.

3. Open *Ordering* and set it to *Order By: Publish Date, And Then By: Title*.
4. Click *Save*.

This will create a Content Set which will contain any items that are currently tagged as *trending* and any future items with the *trending* tag will be added to the Content Set automatically.


Now that you have your Content Sets created, you can display them on a page.

DISPLAYING CONTENT SETS

Content Sets are primarily displayed through the Asset Publisher. It is currently the only method to display them out of the box, but you can develop your own external applications or widgets to utilize Content Sets. In *Creating Content Sets* you created two Content Sets. Now display them on a page.

134.1 Configuring the Asset Publisher for Content Sets

To display the Content Sets, start with a blank page, and then add the necessary Asset Publishers and configure them to display the Content Sets.

1. Create a new *Home* page for your site as a *Widget Page* with a 1 column layout. If you're using a fresh Liferay DXP bundle, you can just remove the *Hello World* widget from the sample *Home* page.
2. Open the *Add* menu and add two *Content Management* → *Asset Publishers* to the page stacked vertically.
3. Click  → *Configuration* for the top Asset Publisher.
4. Under *Asset Selection* choose *Content Set*.
5. Open *Select Content Set* and click *Select*.
6. Click on the *Space Program Images* Content Set.
7. Click *Save*.

Now the images will appear at the top of the page. You can manage the way the content is displayed—like what metadata appears—or even create a *Widget Template* to style the content, but the items which display and the order in which they appear are determined by the Content Set.

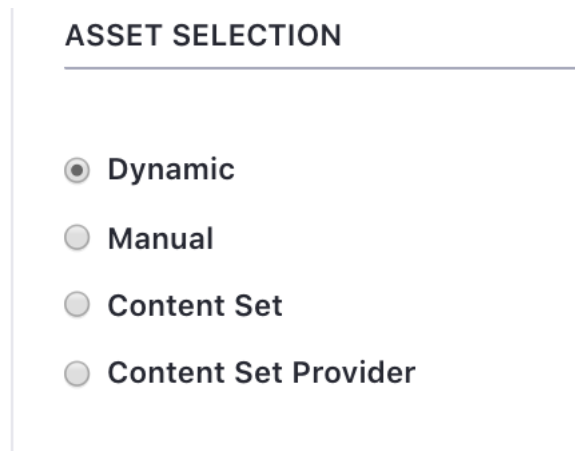



Figure 134.1: The Asset Publisher has a number of options available for selecting its source for content.

Now configure the bottom Asset Publisher with the other Content Set.

1. Click  → *Configuration* for the bottom Asset Publisher.
2. Under *Asset Selection* choose *Content Set*.
3. Open *Select Content Set* and click *Select*.
4. Click on the *Trending* Content Set.




Figure 134.2: Select the Content Set you want to use.

7. Click *Save*.

Again, you can manage various display settings, but the items which appear and their order are determined by the Content Set criteria.

134.2 Adding Items to an existing Content Set

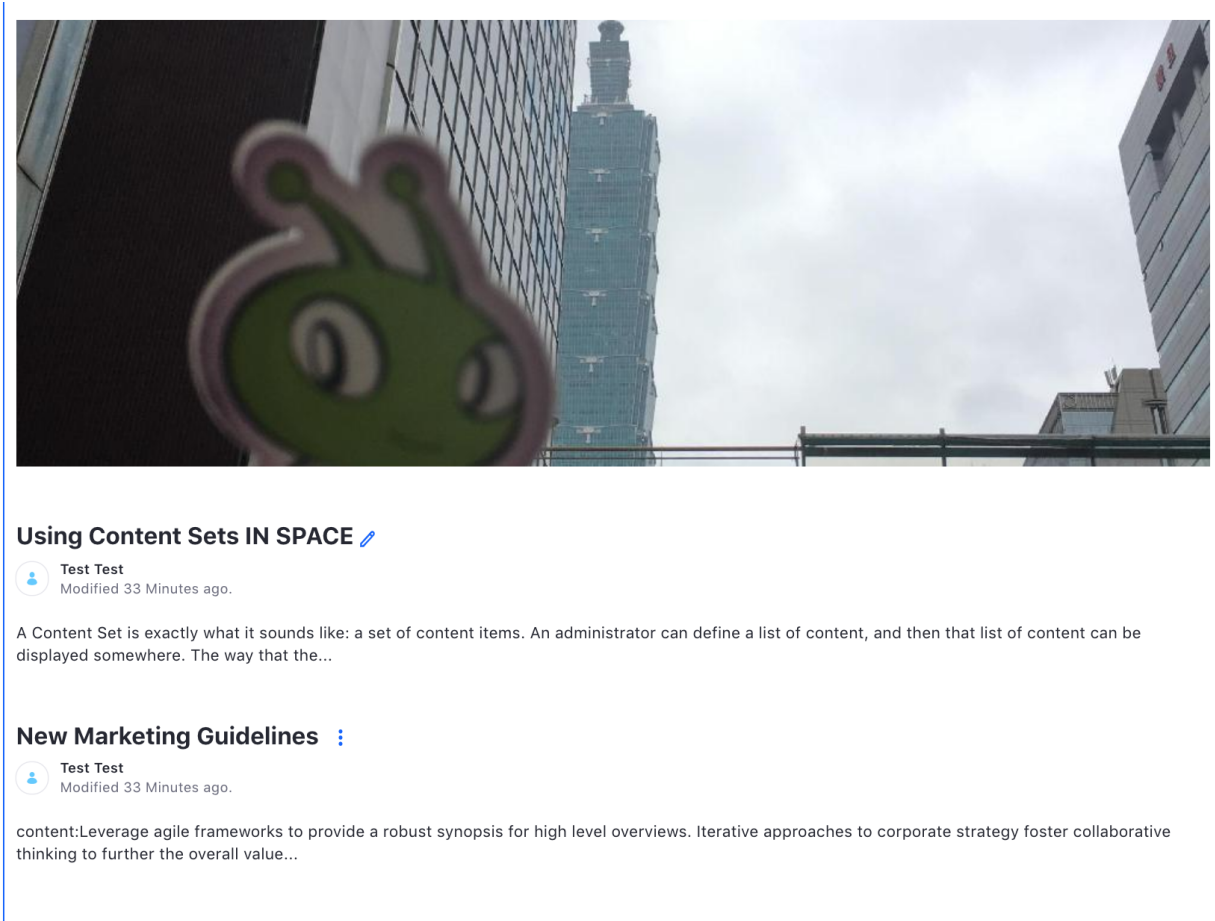




Figure 134.3: You can see the results as the standard Asset Publisher output. You can create Widget Templates to add more style and pizzazz here.

To demonstrate both the management of both static and dynamic Content Sets, upload a new image, tag it, and add it to the static set manually.

1. Upload a new image, and under *Categorization* tag it as *trending*.
2. Without lifting another finger, the image is added to the top of the *Trending* Content List.
3. To add it to the manual set, go back to *Site Administration* → *Content & Data* → *Site Builder*.
4. Click on *Space Program Images* or select  → *Edit* next to *Space Program Images*.
5. Next to *Asset Entries* click *Select* → *Basic Document*.
6. Select the new image and click *Add*.
7. Navigate back to the *Home* page to see your image added to the list.

cute-dog.jpg

 **Test Test**
Modified 1 Minute ago.



ray-taiwan.jpg

 **Test Test**
Modified 33 Minutes ago.



Using Content Sets IN SPACE

 **Test Test**
Modified 34 Minutes ago.

A Content Set is exactly what it sounds like: a set of content items. An administrator can define a list of content, and then that list of content can be displayed somewhere. The way that the Content Set is displayed is determined by the method that is used to display


Figure 134.4: The result is dynamically added to the Content List wherever it is displayed.

Content Sets are a powerful feature which provide one place to easily define content and other assets to be displayed all over your site. Their reusability also means less repeated work involved in getting great content delivered to your users.

CONVERTING ASSET PUBLISHER CONFIGURATIONS TO CONTENT SETS

In the previous two guides in this section, you've seen *Creating Content Sets* and *Displaying Content Sets* demonstrated. Next, try out converting an existing Asset Publisher configuration to a Content Set.

In this case, you have an Asset Publisher on a page, which is configured to display images tagged as *trending* in reverse alphabetical order by title. This might not be too hard to reproduce in the *Content Set* creator, but it's even easier to create the Content Set definition directly from the Asset Publisher.

1. Go to  → *Configuration* for the Asset Publisher.
2. Click *Create a content set from this configuration*.




Figure 135.1: You can generate a Content Set directly from the Asset Publisher configuration.

3. Enter the title and click *Save*.

And as quickly as that you have a new Content Set that you can use with Asset Publishers anywhere on the site.

Great! You converted your Asset Publisher configuration to a Content Set.

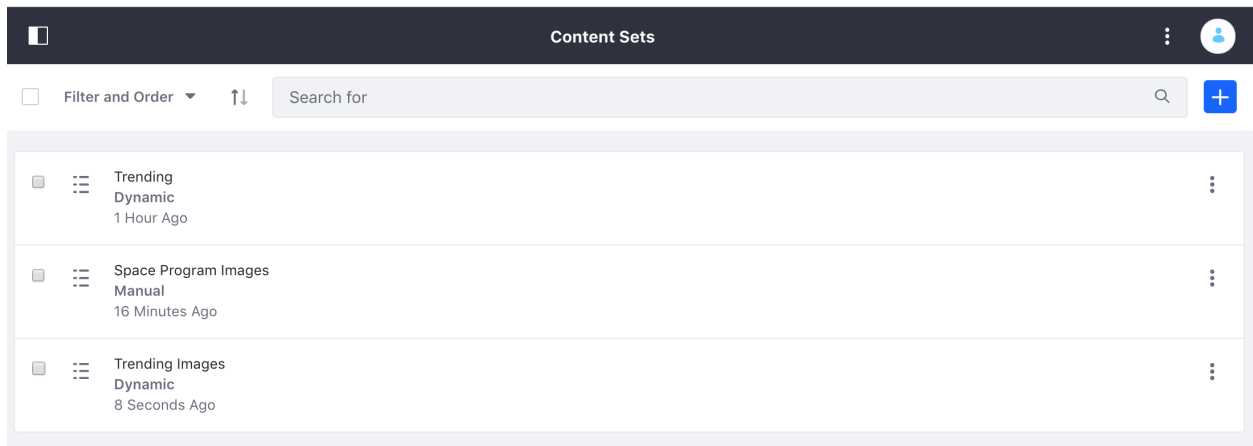
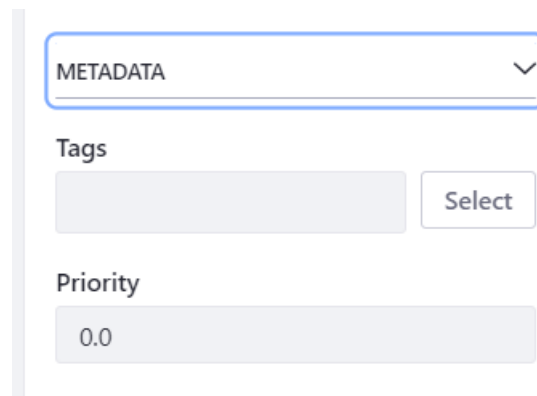


Figure 135.2: The Content Set is added right alongside any existing sets.

ORGANIZING CONTENT WITH TAGS AND CATEGORIES


Tags and categories are two important tools you can use to help organize information in Liferay DXP. These tools help users to easily find the content they're looking for through search or navigation. Tagging and categorizing assets is easy. You can tag or categorize an asset at creation time or when editing an existing asset. If you click on the *Metadata* section of the form when creating or editing an asset, you'll find an interface for adding tags and categories. If no categories are available to be added to the asset (e.g., if no categories have been created), the *Select* option doesn't appear.



The screenshot shows a form section titled 'METADATA'. Below the title, there is a 'Tags' field with a 'Select' button. Below that, there is a 'Priority' field with the value '0.0'.

Figure 136.1: Here is the Web Content application's metadata section.

Note: You'll notice in Figure 1 above that there is also a *Priority* field for web content. This field is not related to categories and tags, but rather, specifies the order in which the web content article is listed when displayed in the Asset Publisher. To learn more about the Asset Publisher, see the Publishing Assets section.

The Menu () contains interfaces for managing tags and categories for each site in Liferay DXP. Navigate to the Site Administration menu → *Categorization*, and you'll find the *Tags* and *Categories* options. These options can be used to manage all your site's tags and categories. It is


important that you both tag and categorize your content when you enter it. You'll take a closer look at tags and categories next.

TAGGING CONTENT

Tags are an important tool that can help organize information and make it easier for users to find the content they want. Tags are all-lowercase words or phrases that you can attach to any content. Tagging content makes your search results more accurate and enables you to use tools like the Asset Publisher to display content in an organized fashion on a web page.

There are two ways to create tags: through the administrative console in the Site Administration section of the Menu or on the fly as content is created. By default, tags can be created by regular users and users can apply them to any assets they have permission to create or edit.

Only site administrators can access the *Tags* application in the Content section of the Site Administration area of the Menu. Here, site administrators can create new tags and edit any existing site tags:

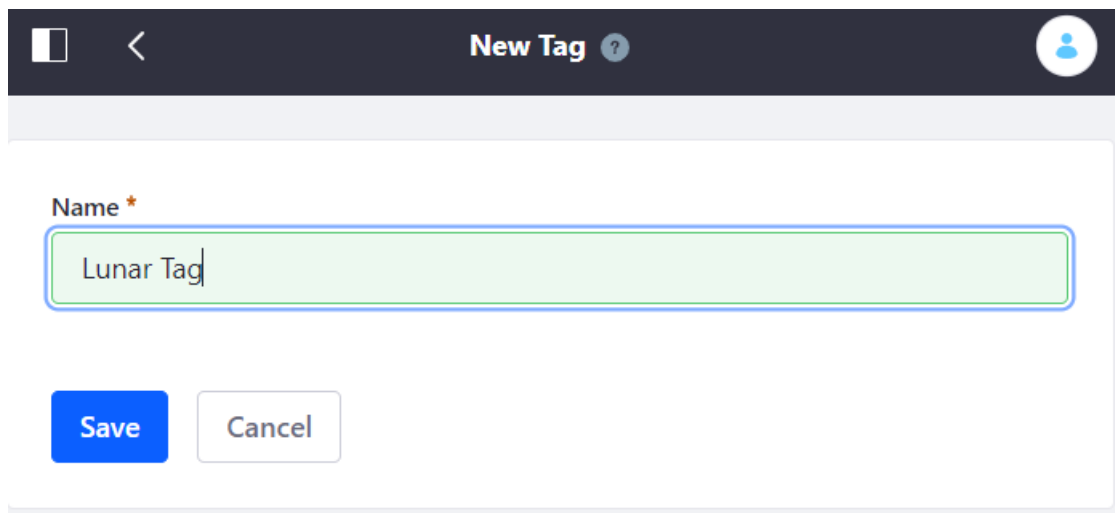
1. Go to the site you want to create tags for and click *Categorization* → *Tags*.
From this screen, you can view existing tags and create new ones.
2. To create a new tag, click the *Add Tag* icon () and enter a name for the tag.

The process for adding tags during content creation is similar. For example, to create tags for a new web content article:

1. Go to the *Metadata* dropdown in a New Web Content menu.
2. Add tags *lunar*, *moon*, and *spectacular*.

Once you've created the web content with these tags, the web content is associated with those tag words when they are searched or referenced.

Tags are not the only instance-wide mechanism for describing content: the next tutorial describes categories.



The image shows a mobile application interface for creating a new tag. At the top, there is a dark header bar with a hamburger menu icon on the left, a back arrow, the text "New Tag" with a help icon, and a user profile icon on the right. Below the header is a light gray container. Inside this container, the label "Name *" is positioned above a text input field. The input field has a light green background and a blue border, and it contains the text "Lunar Tag". Below the input field are two buttons: a blue "Save" button and a white "Cancel" button with a gray border.

Figure 137.1: The Add Tag interface is very simple, only requiring the name of your tag.

DEFINING CATEGORIES FOR CONTENT

Categories are similar in concept to tags, but are designed for use by administrators, not regular users. Hierarchies of categories can be created, and categories can be grouped together in *vocabularies*. While tags represent an *ad hoc* method for grouping content, categories exist to allow administrators to organize content in a more official, hierarchical structure. Think of tags like the index of a book and categories like its table of contents. Both serve the same purpose: to help users find the information they seek.

You can add properties to categories. Category properties are a way to add information to specific categories. You can think of category properties as tags for your categories. Structurally, category properties are just like tag properties: they are key-value pairs associated with specific categories that provide information about the categories.

Adding vocabularies and categories is similar to adding tags:

1. Go to the site where you want to create categories.
2. Click *Categorization* → *Categories* to view the Categories application.

Clicking on a vocabulary displays categories that have been created under that vocabulary. To create a new vocabulary,

1. Click on the *Add Vocabulary* button (+).
2. Enter a name and, optionally, a description.
3. Click *Save*.

By default, the *Allow Multiple Categories* option is enabled. This allows multiple categories from the vocabulary to be applied to an asset. If the box is disabled, only one category from the vocabulary can be applied to an asset. The *Associated Asset Types* lets you choose which asset types the categories of the vocabulary can be applied to and which asset types are *required* to have an associated asset from the vocabulary. Finally, you can configure the permissions of the vocabulary. By default, guests can view the vocabulary but only the owner can delete it, update it, or configure its permissions.

Creating new categories is similar to creating new tags except that categories must be added to an existing vocabulary and they can only be created by site administrators. Once created, however,

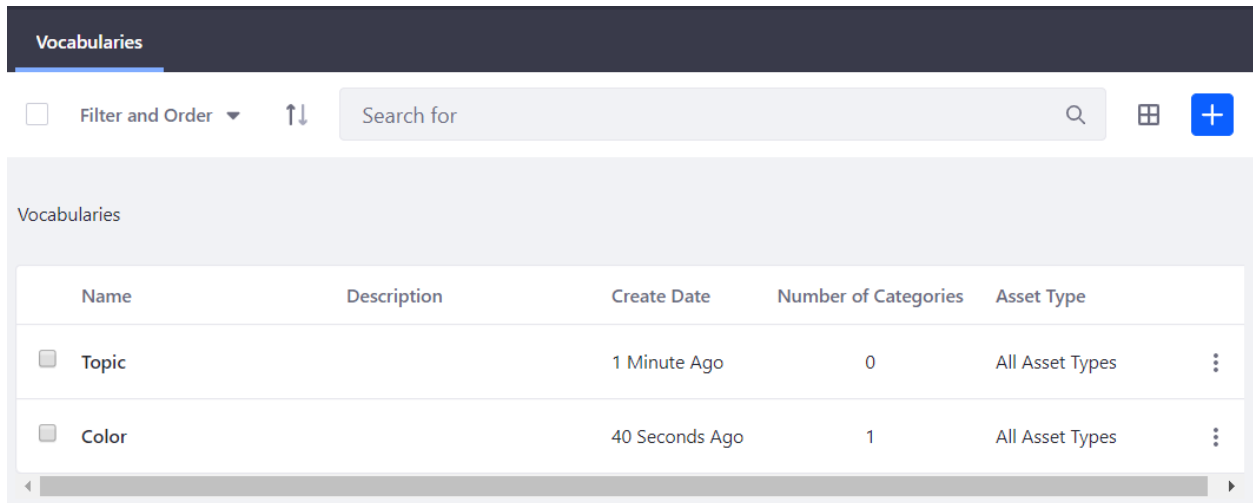


Figure 138.1: After adding new vocabularies, you'll notice your vocabularies indicate the amount of categories existing beneath them.

regular users can apply categories to any assets they have permission to create or edit. To create a new category:

1. Click the *Add Category* icon (+).

If you're already viewing a vocabulary:

1. Select the *Actions* button (⋮) next to an existing vocabulary and select *Add Category*.
2. Enter a name for the new category and, optionally, a description.
3. Click *Save*.

Just as with tags, you can configure the category's permissions, choosing which roles (guest, site member, owner) can view the category, apply it to an asset, delete it, update it, or configure its permissions. By default, categories are viewable by guests, and site members can apply categories to assets.

Once you have created some vocabularies and categories, you can take advantage of the full capabilities of categories by creating a nested hierarchy of categories. To nest categories, select the *Actions* button for the category you want to be the parent category. Then select *Add Subcategory*, which adds a child category to the selected parent.

After you've created a hierarchy of categories, they're available to apply to content:

1. Click on *Web Content* in the *Content & Data* section of Site Administration and click *Add* → *Basic Web Content*.
2. Click on *Metadata* from the right-side menu and click *Select* on the vocabulary you'd like to apply. A dialog box appears with your categories.
3. Select relevant categories by checking the box next to them, and they'll be applied to the content.

Suppose you're running a Lunar Resort shop called Lunar Fireworks and you have many web content articles describing the colors and types of fireworks you offer. The abundance of your articles is overwhelming, and as your shop grows, so too does the web content articles you're required to manage. You've decided to categorize your web content based on the color and type of firework, so the articles are easier to manage.

1. Go to Site Administration → *Categorization* → *Categories* and create vocabularies *Type* and *Color*.
2. Make sure both vocabularies are only used for web content articles by clicking the *Associated Asset Types* dropdown and selecting *Web Content Article*.
3. Create categories *Fire* and *Smoke* for the *Type* vocabulary and *Red*, *Yellow*, and *Blue* categories for the *Color* vocabulary.
4. Now navigate to *Content & Data* → *Web Content* in Site Administration and create an article called *Red Rocket*. This is your best selling product, so make sure to give it a detailed explanation and an awesome picture.
5. Select the *Metadata* dropdown for your web content article and select the *Type* → *Fire* and *Color* → *Red* categories.

When you publish your new web content article for your best selling product, it's organized by its type and color. Once you've organized all your articles, you'll always be able to reference the type and color of a firework, just in case you forget.

There are a few other cool features for vocabularies and categories. A few of them were mentioned already when the *Allow Multiple Categories* and *Required* selectors for vocabularies and categories were discussed. The three new features are targeted vocabularies, single/multi-valued vocabularies, and separated widgets for every vocabulary. They're in the next tutorial.

TARGETED VOCABULARIES

You can decide which vocabularies can be applied to an asset type and which vocabularies are required for an asset type with Targeted Vocabularies. To configure these settings, go to the Categories application in Site Administration and select a vocabulary's *Actions* → *Edit* button. Select the *Associated Asset Types* tab to reveal a dialog box like the one below.

The screenshot shows a dialog box titled "ASSOCIATED ASSET TYPES" with a dropdown arrow on the right. It contains two sections for configuring asset types:

- Section 1:** "Choose Asset Type" with two dropdown menus: "Web Content Article" and "All Asset Subtypes". Below this is a "Required" toggle switch set to "YES" (blue). To the right are "+" and "-" buttons.
- Section 2:** "Choose Asset Type" with one dropdown menu: "Blogs Entry". Below this is a "Required" toggle switch set to "NO" (grey). To the right are "+" and "-" buttons.

Figure 139.1: You can target vocabularies by checking the *Allow Multiple Categories* selector and then selecting the Asset Types.

The default value for *Associated Asset Types* is *All Asset Types*. You can fine tune your choices by using the + and - buttons, which narrows the scope of the vocabulary to specific assets. In the

screenshot above, notice that the vocabulary is configured to be available for Web Content articles and Blog entries, but it is not required. It is mandatory, however, for Bookmark entries.

139.1 Single and Multi-valued Vocabularies

You can also decide if users can choose one or more categories from the same vocabulary to apply to an asset. If a vocabulary is single-valued you can only choose one. If it allows more, you can choose several categories from the vocabulary to apply to an asset.




Figure 139.2: Multi-valued vocabularies allow multiple categories from the vocabulary to be applied to an asset. Single-valued vocabularies only allow one category from the vocabulary to be applied. Here, the *Dining* and *Nightlife* categories are selected to be applied but the *Scenic Adventures* category is not.

You can configure the single-valued or multi-valued status of a vocabulary through the Categories application. Edit a vocabulary and deselect the *Allow Multiple Categories* selector to create a single-valued vocabulary. Use the default option to create a multi-valued vocabulary.

139.2 Separated Entries

A third feature of vocabularies and categories is that every vocabulary has its own separated entry. These entries appear in the Categorization section of the form for editing an asset, and they allow users to easily select appropriate categories for that asset.

It's important to use tags and categories with all your content, so that content is easier for users to find.

Next, you'll learn how to geo-locate assets.

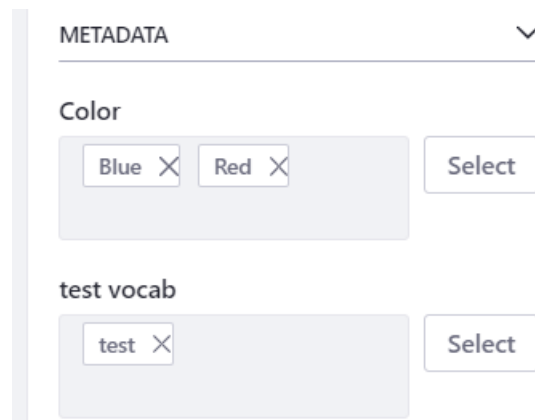


Figure 139.3: Vocabularies have their own entries, making it easy to select available categories.

GEOLOCATING ASSETS

Geolocation adds the geographic coordinates where an asset was created as metadata to your assets. You can add geolocation metadata to your web content, Data Lists, and Documents & Media. This feature is provided for you out-of-the-box. However, you must first enable it in your assets in order to use it.

Let's examine how you can enable geolocation in your web content.

140.1 Geolocating Web Content

To use geolocation in your web content, you must create a structure and template that includes a Geolocation field.

1. Create a structure with a Geolocation field like in the image above.
2. Create a new template and select the structure you just created with the geolocation field.
3. Scroll down to the *Script* heading and locate the *Fields* section. Here are *Content* and *Geolocation* snippets.
4. Click on the snippets to add them to the template and *Save*.

To set your location for the web content, you can share your location with the browser, type a specific address into the address bar on the map, or even drag the indicator and drop it in any point in the map and the address is automatically updated to reflect the new point. Once the web content is saved, the location is added as metadata to the web content.

Note: Depending on your browser settings, you may need to configure it to share your location.

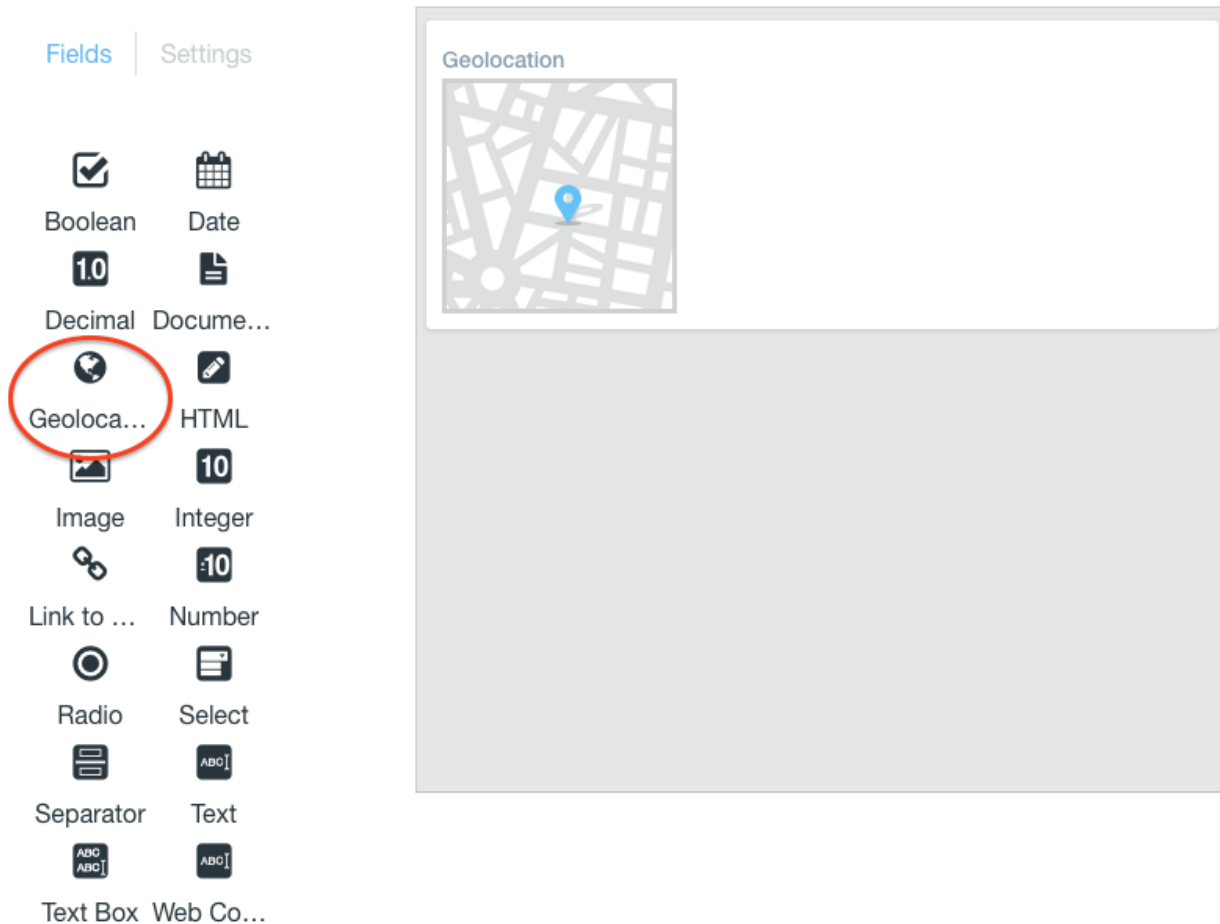


Figure 140.1: Add a geolocation field to your structure to enable geolocation in your web content.

140.2 Geolocating Data Lists

To use geolocation in your dynamic data lists, you must first create a data definition that includes a geolocation field.

1. Open the Product Menu and navigate to *Content & Data* → *Dynamic Data Lists*.
2. Click the *Options* menu and select *Manage Data Definitions*.
3. Click the *Add* button to create a new data definition.
4. Enter a name, optional description, and parent data definition if you have one.
5. Scroll down and add a *Geolocation* field to the data definition, along with any other fields you wish to add and *Save*.
6. Go back to the Dynamic Data Lists screen and click the *Add* button (+) to create a new list.

The screenshot shows the 'Lunar Resort template' editor. On the left, there are two panels: 'GENERAL VARIABLES' and 'FIELDS'. Under 'GENERAL VARIABLES', 'Web Content' and 'Geolocation' are highlighted with red boxes. Under 'FIELDS', 'Boolean', 'Color', 'Web Content', and 'Geolocation' are listed. The main editor area displays a code snippet with line numbers 1 through 24. Red arrows point from the 'Web Content' and 'Geolocation' boxes in the 'GENERAL VARIABLES' panel to lines 3 and 13 in the code editor, respectively. The code includes assignments for 'webContentData', 'latitude', and 'longitude', and an if-statement for geolocation data. A 'Script File' section at the bottom shows a 'Choose File' button and the text 'No file chosen'.

Figure 140.2: Add the Content and Geolocation snippets to create your web content template quickly.

7. Enter a name and optional description.
8. Finally, click the *Select* button and choose the newly created data definition.

Now that your data list is complete, you can use the Data List Display portlet to display it.

140.3 Geolocating Documents and Media

To enable geolocation in Documents and Media, you must first create a document type that includes geolocation metadata. You can add geolocation metadata as part of a Metadata Set or as part of the new document type. To add geolocation metadata as part of a Metadata Set:

1. Open the Product Menu and navigate to *Content & Data* → *Documents and Media*. Open the *Options* menu, and select *Metadata Sets*.
2. Click the *Add* (+) button and enter a name, optional description, and Parent Metadata Set if you have one.




Figure 140.3: You can enter your location in the address bar, move the indicator to a location, or share your location with the browser.




Figure 140.4: Make sure your browser is configured to share your location.

3. Scroll down and add a *Geolocation* field, along with any additional fields you wish to have, and *Save*.

To create the new document type with geolocation:

1. Navigate to *Documents and Media*, open the *Options* menu and select *Document Types*.
2. Click the *Add* button (+) and enter a name and optional description.
3. Scroll down to the *Main Metadata Fields* heading and add a *Geolocation* field along with any other fields you wish to have for the document type.
4. If you are using a *Metadata Set*, scroll down to the *Additional Metadata Fields* heading, click the *Select Metadata Set* button.
5. Choose your *Metadata Set* with the geolocation metadata and *Save*.
6. Navigate back to the *Documents and Media* screen and click the *Add* button (+) and select your newly created document type.
7. Fill out the information for the document, and just as with the web content, your location is automatically obtained from the browser and added to your document.

Once your assets are geolocation enabled, you can use the *Asset Publisher* to display the location of the assets on a map, using the map display template. Check out the *Configuring Display Settings* section to learn more.

Asset Publisher

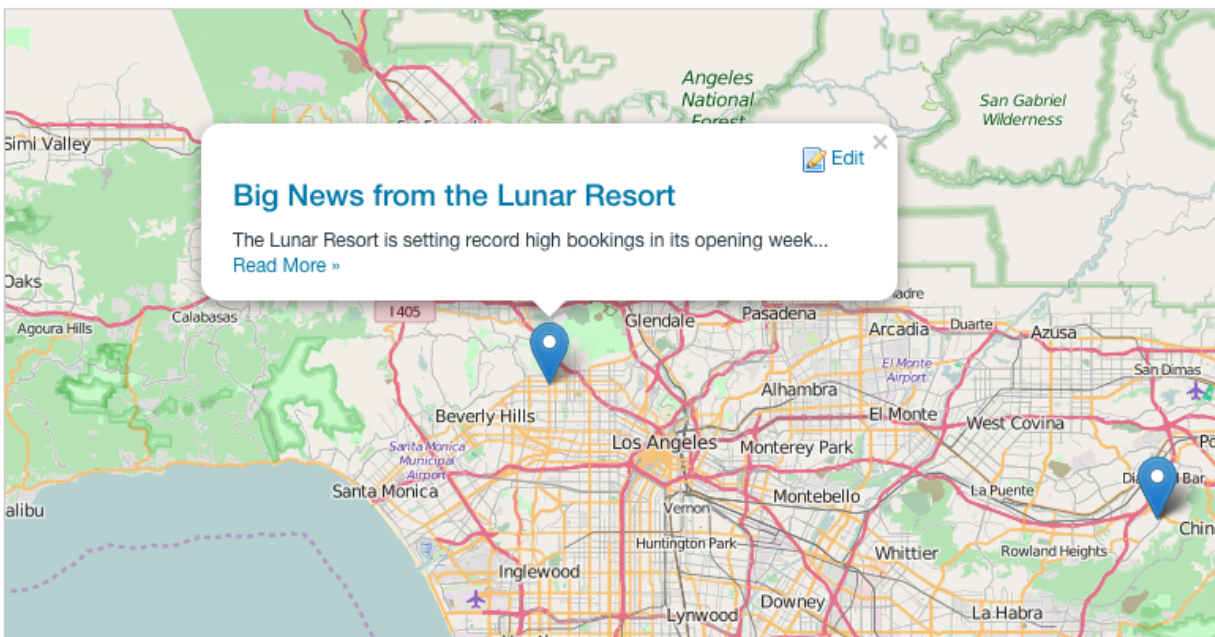


Figure 140.5: The *Asset Publisher* can display your geolocated assets on a map.

PUBLISHING CONTENT DYNAMICALLY

Most content types are Assets. In the Creating Web Content articles, you examined the most common type of asset: web content. Other types of assets include blog posts, wiki articles, message board posts, bookmarks, and documents. Developers can define custom asset types that use the asset framework, which provides support for tags, categories, vocabularies, comments, ratings, and asset relationships.

The Asset Publisher application displays assets. It has many configuration options which you'll cover in this chapter. By default, Asset Publisher displays abstracts (previews) of recently published assets with links to their full views. You can configure the Asset Publisher app to display a table of assets, a list of asset titles, or the full content of assets. You can also make it display only certain kinds of assets, and you choose how many items to display in a list.

You might use Asset Publisher to display chosen content types, recent content, or content by tags and categories.

This section covers the following topics:

- Adding relationships between assets
- Publishing assets
- Publishing RSS feeds
- Restoring deleted assets


The first thing you'll learn about is tagging and categorizing content.

DEFINING CONTENT RELATIONSHIPS

Related Assets are assets connected to other assets, even if they don't share any tags and aren't in the same category. Here you'll focus on how to define relationships between assets so when you begin publishing assets, the Related Assets widget can display those relationships.

142.1 Related Assets Widget

By default, the Related Assets widget displays any related asset of the asset selected in the Asset Publisher. If you don't want to show every related asset, you can configure what content relationships to display. To do this, follow these steps:

1. Go to the Related Assets app and select the *Options* icon () in the upper right corner of the application and click *Configuration*.
2. Under the *Setup* → *Asset Selection* tab, set the type of asset(s) to display using the *Asset Type* menu. The default value is set to *Any*.
3. You can narrow the scope of the app to display any single category of asset type or select multiple assets from the menu.

Filter options set minimum requirements for displaying assets by their categories, tags, and custom fields. Ordering and Grouping organizes assets using the same criteria. Display settings customize how the app shows assets: by title, in a table, by abstract, or full content. You can convert assets to different document types like ODT, PDF, and RTF. You can choose to show metadata fields such as author, modification date, tags, and view count. You can even enable RSS subscriptions and customize their display settings.

4. When you're finished setting the Source and Filter options, click *Save*.

Now that you've configured the Related Assets widget to display specific content types, you must define the relationships for your assets. Here's a simple example of defining related assets for a web content article and then displaying those related assets.

Suppose you own a gift shop at the Lunar Resort, and you want all your shop's assets to appear when an asset is clicked. You must define relationships between your content, so when an asset is clicked, its related assets are appear alongside the clicked asset. Here's how to do it:

1. Create a blog entry explaining your gift shop's new apparel and a photo of the moon, just so consumers are aware that you offer the *only* gift shop on a desolate rock orbiting the Earth!
2. Create a web content article describing your shop. Once you've given your article a title and some content, open the *Related Assets* dropdown menu. Click *Select*, choose *Blogs Entry*, and select the blog you created. Click *Select* again, choose *Basic Document*, and select the photo of the moon. Click *Publish* to publish your web content article.
3. Now that those assets are created, you can relate the blog entry and photo to your web content article. Navigate to your article in Site Administration → *Content & Data* → *Web Content*.
4. You've now defined relationship with your three assets. Click the *Add* icon (+) at the top of your page in the Control Menu, select *Widgets*, and add the Related Assets and Asset Publisher widgets to the page. Don't panic: related assets don't appear until you select an asset in the Asset Publisher.

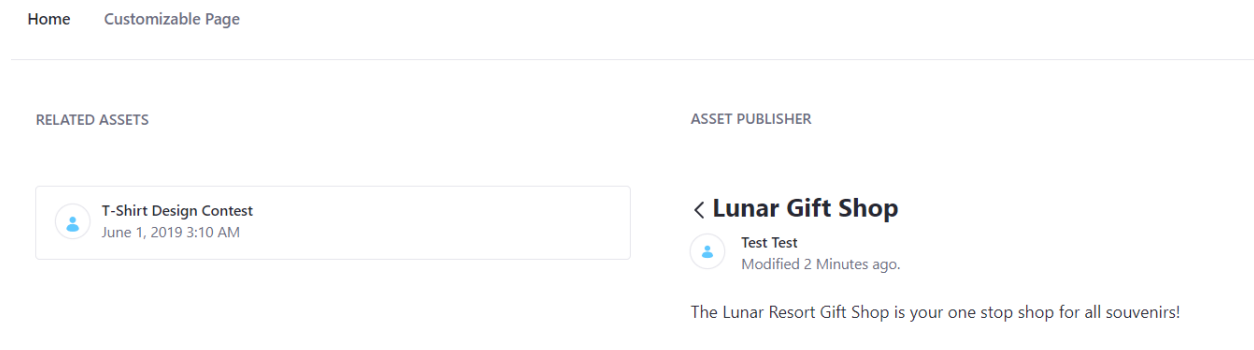


Figure 142.1: Select an asset in the Asset Publisher to see its related assets displayed in the Related Assets application.

Once you select an asset, its related assets appear in the Related Assets app, as in the image above. If you want more detail, you can place two Related Assets widgets on the page and name one *Related Blogs* and the other *Related Photos*.

Next, you'll learn more about how to use the Asset Publisher.





Figure 142.2: Related Assets applications can be configured to display specific content.

PUBLISHING ASSETS

As you create web content, remember that pieces of content are assets, just like message board entries and blog posts. Since the Asset Publisher publishes assets, it excels at publishing mixed content types like images, documents, blogs, and of course, web content. This helps in creating a more dynamic web site: you can place user-created wiki entries, blog posts, or message board messages in context with your web content. You'll examine some of the Asset Publisher's features next.

QUERYING FOR CONTENT

The Asset Publisher works by querying for mixed types of content on the fly. Since you can control what and how content is displayed from one location, the Asset Publisher helps to “bubble up” the most relevant content to your users.

To get to all the application’s options, click the *Options* icon () in the application’s menu. If you click the *Configuration* option and then *Setup* (if necessary), you can configure the Asset Publisher’s settings from the following three areas:

- Asset Selection
- Display Settings
- Subscriptions

Asset Selection configures which assets are displayed. You can set asset selection to

- *Dynamic*
- *Manual*
- *Content Set*
- *Content Set Provider*

Dynamic displays assets based on certain rules or filters. For example, you can set the Asset Publisher to display only assets of a certain type or to which certain tags or categories have been applied. Manual asset selection only displays assets that have been explicitly selected by an administrator. For more information on Content Sets and their place in the Asset Publisher, see the Managing Content Sets section.

The Asset Publisher supports a scope that restricts both dynamic and manual asset selection. The Asset Publisher can only display assets from its configured scope. By default, the Asset Publisher app is scoped to the site of the page to which it was added. You can, however, customize the scope from the Asset Selection section of the Asset Publisher configuration window. To extend your Asset Publisher’s scope,

1. Click *Select* under Scope.
2. Choose *Global* to add the global scope or *Other Site...* to add the scope of another site.


The Display Settings section of the Asset Publisher configuration window is for customizing how content is displayed. The Subscription section enables, disables, or configures email subscriptions and RSS subscriptions. In the following sections, you'll explore the available configurations for the Asset Selection, Display Settings, and Subscriptions sections of the Asset Publisher's configuration window. You'll start by learning how select content manually. You'll see that it's very similar to using the Web Content Display application except that you can select assets of any type, not just web content articles.

SELECTING ASSETS

You can configure Asset Publisher to select assets manually or dynamically through various criteria. Within those options there is flexibility in what assets are displayed and how they are displayed.

145.1 Selecting Assets Manually

To enable manual asset selection,

1. Click the *Options* icon () in the widget's menu.
2. Select *Configuration*.
3. In the Asset Publisher configuration, select *Manual* from the select box beneath *Asset Selection*.

Now you must select a *Scope* and specific *Asset Entries* from that scope to display. You can configure multiple scopes, including the global scope, from which to select assets.

When selecting assets manually, a list of configured scopes appears under the *Scope* heading. You can configure the scope like this:

1. Click the *X* button at the right to remove a scope from the list.
2. Click the *Select* button to add additional scopes to the Asset Publisher's configuration.
3. After you've added a scope, a new *Select* button appears under the *Asset Entries* heading. A list of assets selected for display appears in the *Asset Entries* section. You can select assets to be displayed by clicking the appropriate *Select* button. One button appears for each configured scope. By default, these are the available asset types:
 - Blogs Entry
 - Bookmarks Entry
 - Bookmarks Folder




Figure 145.1: Selecting assets in the Asset Publisher manually is similar to selecting assets in the Web Content Display application except that you can select assets of any type, not just web content. You can also add scopes to expand the list of assets that are available to be displayed in the Asset Publisher.

- Calendar Event
- Basic Document
- Google Drive Shortcut
- Documents Folder
- Dynamic Data Lists Record
- Knowledge Base Article
- Message Boards Message
- Basic Web Content
- Web Content Folder
- Wiki Page

You can select any number of assets to be displayed. Note, however, that there's a display setting called *Number of Items to Display* that determines the maximum number of items to display (or, if pagination is enabled, the maximum number of items to display per page). The Asset Publisher can mix and match different asset types in the same interface.

4. When you're done selecting items to display, click *Save*. Any selected assets are added to the list of assets that are displayed by the application.

Once you have your content selected, you can configure the display types to configure how the content appears. You'll explore the display settings in more detail after we finish discussing how to select assets for display.

Manual asset selection lets you select assets of various types from different scopes, but it can be time-consuming to update the assets that should be displayed. It's often more convenient to use the Asset Publisher to select content dynamically.

145.2 Selecting Assets Dynamically

The Asset Publisher's default behavior is to select assets dynamically according a set of customizable rules. These rules can be combined so that they compliment each other to create a nice, refined query for your content. Assets are filtered by permissions automatically, no matter how complicated your asset selection rules are. You have the following rule types:

Scope: Choose the sites containing the content that should be selected. This works the same way as with manual asset selection: assets can only be displayed if they belong to a configured scope. The following scope options are available:

- *Current Site*
- *Global*
- *Other Site*

The Other Site scope option is unavailable for Asset Publisher applications configured on a page template (e.g., Content Display Page).

Asset Type: Choose the asset types you want, from all assets, to only one, or any combination in between. For example, you could choose only web content, only wiki entries, or any combination of multiple types.

Filter: Add as many filters on tags or categories as you like. You can choose whether the content must contain or must not contain any or all of the tags or categories that you enter.

FILTER ▼

Displayed assets must match these rules.

Rules

Contains ▼

Any ▼

of the following

Tags ▼

Select

+

Figure 145.2: You can filter by tags and categories, and you can set up as many filter rules as you need.

Once you've set up your filter rules for dynamically selecting content, you can decide how the content is displayed.

If you've added custom User profile attributes, you can configure the Asset Publisher to display assets that match them. This setting retrieves assets that have matching categorization. These categories must be from the global context. For example, suppose a User has a custom field called *Location* with the type *Text*. If this attribute is set to *Moon*, you could create a vocabulary called *Location* and a category for the Location vocabulary called *Moon*. Then you could categorize content with *Moon* in the *Location* vocabulary. With this organizational setup, adding an Asset Publisher

and specifying *Location* as the Asset Publisher’s custom user attribute would only display content that had been categorized as *Moon*. Pretty cool, right?

See *Defining Categories for Content* for further information.

In addition, you can use these advanced filters:

- **Show only assets with *Welcome* as its display page** displays only assets specifically configured for the *Welcome* page.
- **Include tags specified in the URL?** lets you specify tags in the URL for the Asset Publisher to display.

The *Ordering* section of the Asset Publisher precisely controls how content is ordered and grouped when displayed. You can order the assets displayed by Asset Publisher in ascending or descending order by the following attributes:

- Title
- Create Date
- Modified Date
- Publish Date
- Expiration Date
- Priority

Say you have a series of “How To” articles that you want displayed in descending order based on whether the article was tagged with the *hammer* tag. Or suppose you want a series of video captures to appear in ascending order based on a category called *birds*. For these use cases, you can configure the ordering and grouping settings.

You can also configure a second ordering. The second ordering is applied to any assets for which the first ordering wasn’t sufficient. For example, if you ordered assets by title and there are multiple assets with the same title, the second ordering takes effect, perhaps the publication date.

You can establish grouping rules as well as ordering rules. You can group assets by type or by vocabulary. For example, suppose you have a vocabulary called *Membership Type* with two categories: *Premium* and *Regular*. If you group assets by Membership Type, all assets with the Premium category appear in one group and all assets with the Regular category appear in another group. Grouping rules are applied before any ordering rules: they’re a way to divide up the displayed assets into separate lists. The ordering rules are applied separately to each group of assets.

Note that ordering rules are only one way to control how your content appears. You can refine the display through many other display settings which you’ll examine next.


Note: The following actions have immediate effects in your Asset Publisher:

- Change the value of the *Asset Selection* option.
- Change the value of the *Scope* option.
- Select, add, sort or delete asset entries (only when selecting assets manually).

Other changes happen after clicking *Save*. Next you’ll learn about the Asset Publisher’s other configuration options.

145.3 Selecting a Content Set

7.0 adds the Content Set feature to create predefined lists of content to display in the Asset Publisher. To use a Content Set,

1. Click the click the *Options* icon () in the Asset Publisher's menu.
2. Select *Configuration* from menu.
3. Select *Manual* from the select box beneath the *Asset Selection* tab.
4. Choose the Content Set that you want to use.

For more information on using Content Sets, see [Creating Content Sets](#).

CONFIGURING DISPLAY SETTINGS

This document has been updated and ported to Liferay Learn and is no longer maintained here.

From the Asset Publisher's configuration page, open the Setup tab's *Display Settings* sub-tab. This section gives you precise control over the display of your assets. There are many options available to configure how you want your content to appear. Many of these, such as printing, flags, ratings, comments, comment ratings, and social bookmarks work the same way they do in the Web Content Display application.

Display Template: This selector lets you choose an application display template to customize how the Asset Publisher displays assets. These templates are in every site by default:

- *Abstracts:* Shows the first 200-500 characters of the content, defined by the **Abstract Length** field. This is the default display template.
- *Table:* Displays the content in an HTML table which can be styled by a theme developer.
- *Title List:* Displays the content's title as defined by the user who entered it.
- *Full Content:* This display template displays the entire content of the entry.

There's also the *Rich Summary* and *Map* display templates that belong to the global scope. The Rich Summary template provides a summary view of each asset along with a *Read More* link to the article's full content. The Map template displays geo-localized assets in either a Google Map or an Open Street Map provider. The map provider can be configured in Instance Settings, and Site Settings in the Advanced section.

Abstract Length: Select the number of characters to display for abstracts. The default is 200.

Asset Link Behavior: The default value is *Show Full Content*, which displays the full asset in the current Asset Publisher. *View in a Context* causes that asset to be displayed in the application where the asset belongs. For example, a blog entry is displayed in Blogs where it was created. See the section below on display pages for more information.

Number of Items to Display: Select the maximum number of assets that can be displayed by the Asset Publisher. If pagination is enabled, this number represents the maximum number of assets that can be displayed per page.

Pagination Type: This can be set to *None*, *Simple*, or *Regular*. *None* displays at most the number of assets specified in the **Number of Items to Display** property. *Simple* adds Previous and Next buttons for browsing through pages of assets in the Asset Publisher. *Regular* adds more options and information including First and Last buttons, a dropdown selector for pages, the number of items per page, and the total number of results (assets being displayed).

Show Add Content Button: When selected, an *Add New* button appears that lets users add new assets directly from the Asset Publisher application. This is checked by default.

Show Metadata Descriptions: Enables Metadata descriptions such as *Content Related to...* or *Content with tag...* to be displayed with the published assets.

Show Available Locales: Since content can be localized, you can have different versions of it based on locale. Enabling this option shows the locales available, so users can view the content in their languages.

Set as the Default Asset Publisher for This Page: The Asset Publisher app is an instanceable app: multiple Asset Publishers can be added to a page and each has an independent configuration. The default Asset Publisher for a page is the one used to display web content associated with the page.

Show only assets with Home as its display page template: Display assets that only exist on the Home page template.

Include tags specified in the URL: Incorporate tags specified in the URL.

Enable ...: Enable/disable the following options for displayed assets:

- Print
- Flags
- Ratings
- Related Assets
- Subscribe
- Comments
- Comment Ratings
- View Count Increment

The Print option adds a *Print* link to the full view of an asset displayed in the Asset Publisher. Clicking *Print* opens a new browser window with a print view of the asset. Enabling flags, related assets, ratings, comments, comment ratings, or social bookmarks add links to the corresponding social features to the view full of the asset in the Asset Publisher.

Tip: An alternate way to add flags, comments, and ratings to a page is through the *Page Flags*, *Page Comments*, and *Page Ratings* applications. Just add the applications in the appropriate location near the asset that should have feedback.


Metadata: Select various metadata types to be displayed (see below). For example, you can select tags and categories for display. Upon saving your configuration, the Asset Publisher displays tags and categories for each displayed asset. Then users can click on the tags and categories to filter the displayed assets manually.

Next you'll learn about configuring subscriptions for email and RSS through the Asset Publisher.

METADATA

Current

Author
Modified Date



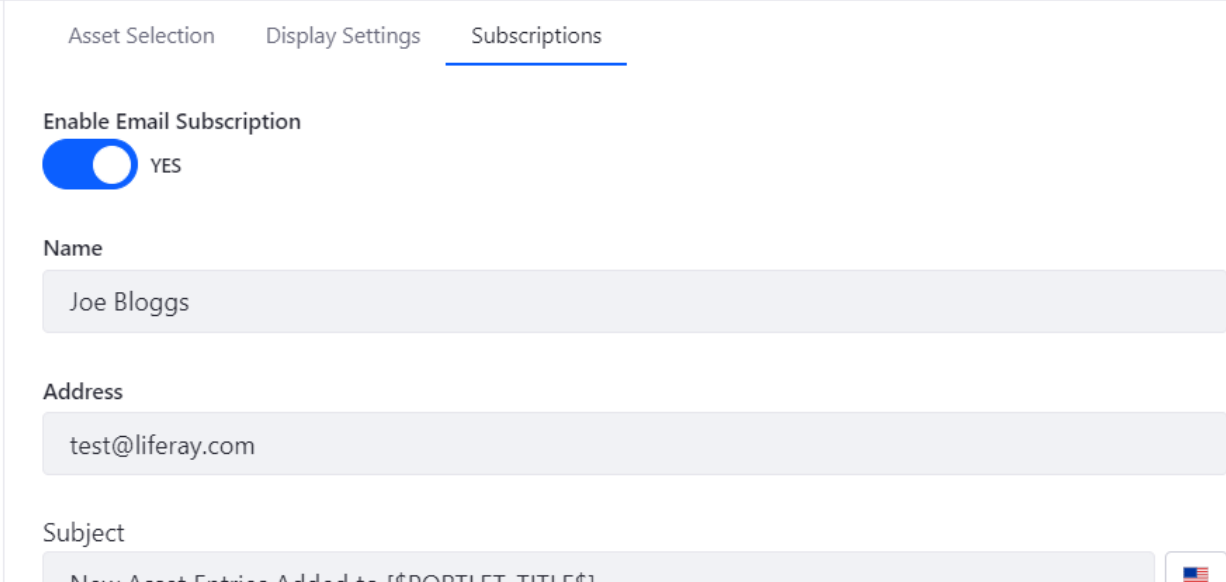
Available

Categories
Create Date
Expiration Date
Priority
Publish Date
Tags
View Count

Figure 146.1: You can configure the Asset Publisher to display various kinds of metadata about the displayed assets.

CONFIGURING ASSET PUBLISHER SUBSCRIPTIONS

The Asset Publisher application supports two kinds of subscriptions email subscriptions. To enable subscriptions, click the Asset Publisher's Options icon and select *Configuration*. In the configuration window, open the Setup tab's Subscriptions tab. There you can enable/disable the *Enable Email Subscription* selector.



The screenshot shows the configuration window for Asset Publisher Subscriptions. It has three tabs: 'Asset Selection', 'Display Settings', and 'Subscriptions' (which is selected). Under 'Subscriptions', there is a section 'Enable Email Subscription' with a blue toggle switch turned on and the text 'YES'. Below this are three text input fields: 'Name' with the value 'Joe Bloggs', 'Address' with the value 'test@liferay.com', and 'Subject' with the value 'New Asset Entries Added to {PORTLET_TITLE}'. There is also a small US flag icon in the bottom right corner of the subject field.

Figure 147.1: An email subscription notifies users when new assets are published.

Enabling Email Subscription adds a *Subscribe* link to the Asset Publisher. Users wishing to be notified of newly published assets can click on this link to be added to the subscription list. Liferay DXP periodically checks for new assets and sends emails to subscribed users informing them about the new assets. By default, Liferay performs this check every twenty-four hours.

PUBLISHING RSS FEEDS

RSS is a family of web feed formats used to publish frequently updated works such as blog entries and news articles. RSS allows users to stay up-to-date with your site's content without actually having to visit your site. Users can use their own RSS feed readers to aggregate content, and you can also use RSS to share and aggregate content across sites. Next, you'll see how to create RSS feeds from Asset Publisher configurations.

CONFIGURING RSS FEEDS

Note: RSS feeds are deprecated for Liferay DXP 7.2 and are disabled by default. To leverage RSS feeds, you must enable this feature. Go to the Control Panel → *Configuration* → *System Settings* → *Web Content*. Under the *System Scope* → *Administration* tab, check the *Show Feeds* box. For more information on deprecated apps, see this article.

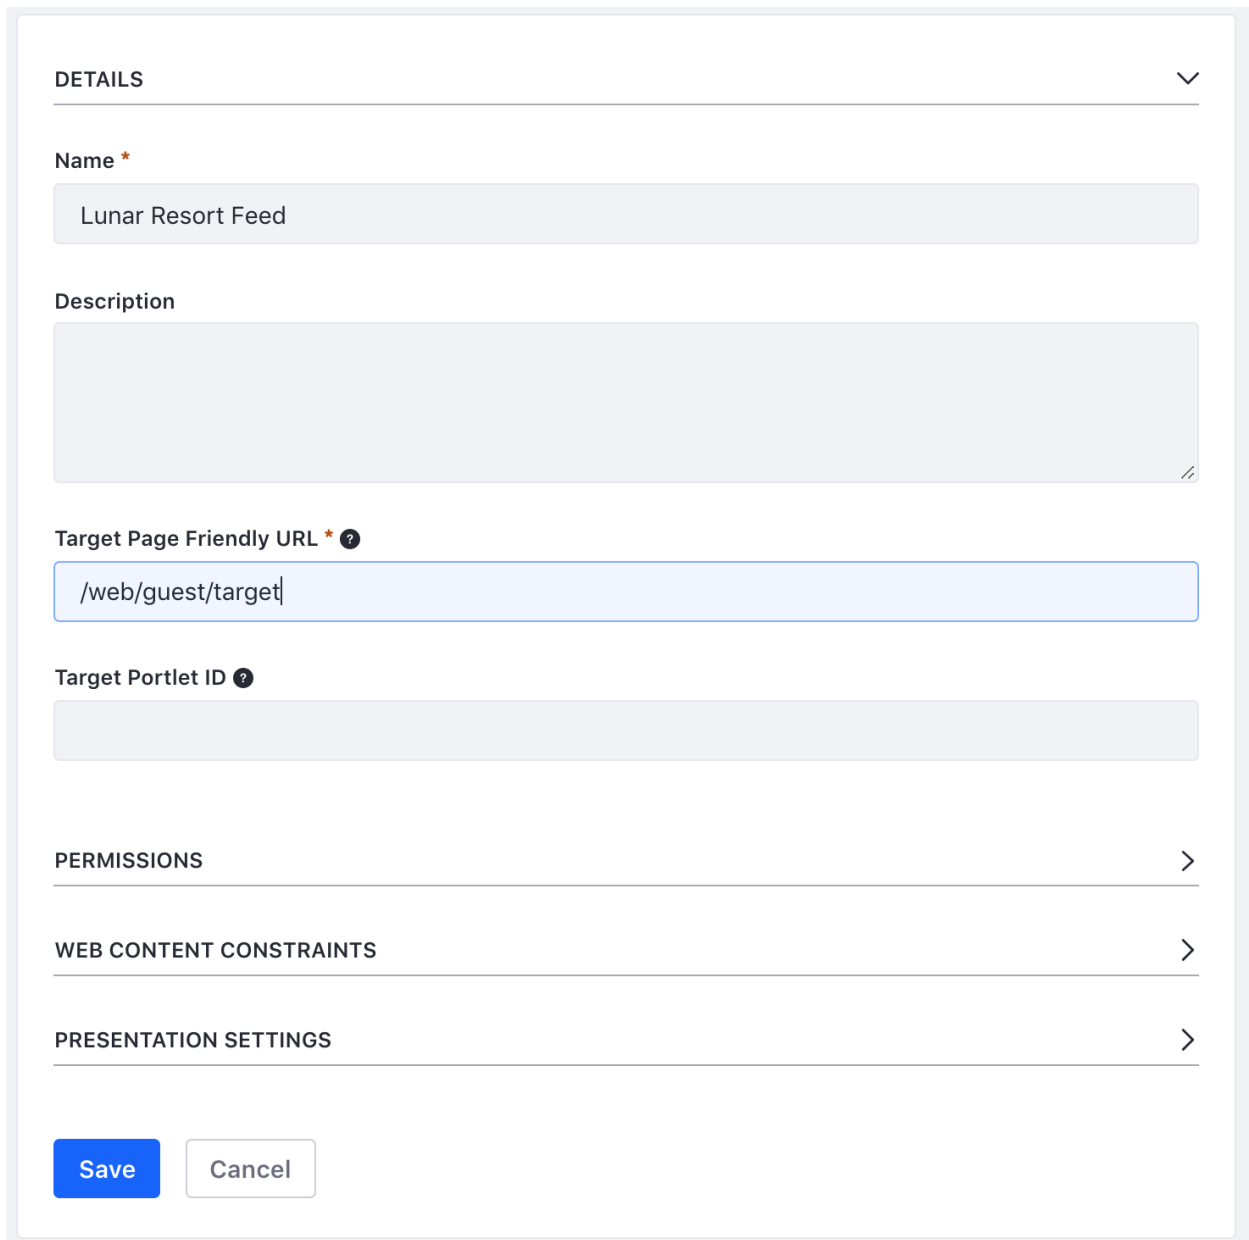
To manage a Liferay site's RSS feeds, navigate to your Site's Site Administration → *Content & Data* → *Web Content*. Site administrators can use this Web Content menu option to manage their site's web content, including web content structures and templates, which you learned in the Creating Web Content section. Site administrators can also use this option to manage their site's RSS feeds. To add a new feed:

1. Go to the *Feeds* tab.
2. Click the *Add Feed* button.
3. Enter a *Name*, select a *Target Page*, and select a *Web Content Structure* for the feed.

A feed's target page serves two purposes:

1. The site the target page belongs to determines which web content articles appear in the feed. For example, if the target page belongs to the Marketing site, only web content articles belonging to the Marketing site appear in the feed.
2. The target page is the page where "orphaned" web content articles are displayed. Orphaned web content articles have been published in your Site but are not displayed in specific Web Content Display applications. Liferay RSS feeds can provide links to any published web content articles, both orphaned articles and articles that have been configured to be displayed in specific Web Content Display applications. For articles that have been configured to be displayed, the RSS feeds' links point to the Liferay page of that app. For orphaned articles, the RSS feeds' links point to the feed's target page. When users click on such links for orphaned articles, the full content of the orphaned article is displayed on the target page.

To specify a target page, you must enter the target page's friendly URL. Note that friendly URLs don't include the host name. For example, the friendly URL of a public page called *Welcome*



The image shows a configuration form for an RSS feed. It is divided into several sections: 'DETAILS', 'PERMISSIONS', 'WEB CONTENT CONSTRAINTS', and 'PRESENTATION SETTINGS'. The 'DETAILS' section is currently expanded and contains the following fields:

- Name ***: A text input field containing 'Lunar Resort Feed'.
- Description**: A large, empty text area.
- Target Page Friendly URL * ?**: A text input field containing '/web/guest/target|'.
- Target Portlet ID ?**: An empty text input field.

At the bottom of the form, there are two buttons: a blue 'Save' button and a white 'Cancel' button with a grey border.

Figure 149.1: To create a new RSS feed, you only need to specify a name, target page, and web content structure. Of course, you can also configure other features of the feed such as its permissions, web content constraints, and presentation settings.

belonging to a Site called *Marketing* might look like this: `/web/marketing/welcome`. Optionally, you can specify a target portlet ID. This would be the portlet ID of a Web Content Display application on the target page in which orphaned web content should be displayed. The application must exist or else the content isn't displayed. The URL field contains the address of your RSS feed. It appears after you've actually created the feed by clicking *Save*.

The final two sections of the *Add Feed* form are for customizing the web content articles that appear in your feed.

1. **Web Content Constraints** selects a web content structure to filter the articles that appear in your feed. This is useful since all web content articles are created using web content structures.
2. **Presentation Settings** customizes additional details about your feed and how articles are displayed in your feed. Leave the **Feed Item Content** set to *Web Content Description* if you want a description of each article to appear in your feed. Set it to *Rendered Web Content: Use Default Template* if you want the full content of each article to appear in the feed. Customizing the **Feed Type** allows you to choose which web feed language to use for your feed. You can choose *Atom 1.0* (the default), *RSS 1.0*, or *RSS 2.0*. Customize the *Maximum Items to Display* to choose the maximum number of articles should appear in your feed at one time. Leave the **Order By Column** set to *Modified Date* to have articles arranged in order from the last time they were published or modified. You can set the **Order by Column** to *Display Date* if you want to have articles arranged in order from the time they were configured to be displayed in a specific Web Content Display application. Lastly, you can leave the **Order by Type** set to *Ascending* to have the oldest articles at the top of the feed or you can set it to *Descending* to have the newest articles at the top of the feed.

When you're done configuring your RSS feed, click *Save* to create your feed.

Once one or more feeds have been created, they'll appear in a list in the **Feeds** tab. You can edit existing feeds using the same form used for creating them. The main difference is that when you edit an existing feed, the URL field is populated. Copy this URL into a new browser tab or window to test your feed. From the **Feeds** popup window, you can also customize the permissions of feeds or delete feeds.

It's possible to completely disable RSS feeds at the instance level. You can do this by setting the `rss.feeds.enabled` property to `false` in your `portal-ext.properties` file. By default, it's set to `true`. If you keep the default, RSS enabled, you can make several other RSS property customizations. Please refer to the RSS section of your `portal.properties` file for details.

THE RSS PUBLISHER WIDGET

The RSS Publisher widget displays RSS feeds. If you're looking for a web-based RSS reader, look no further: just add the RSS Publisher widget to one your personal Site's private pages, and *voila!* You have your own personal RSS reader. You can select the RSS feeds the widget displays and how it displays them. The RSS Publisher widget can also be placed on Sites' public or private pages to make feeds available to guests or Site members, respectively. In these cases, make sure that only Site administrators have permission to customize the RSS widget and select feeds to be displayed.

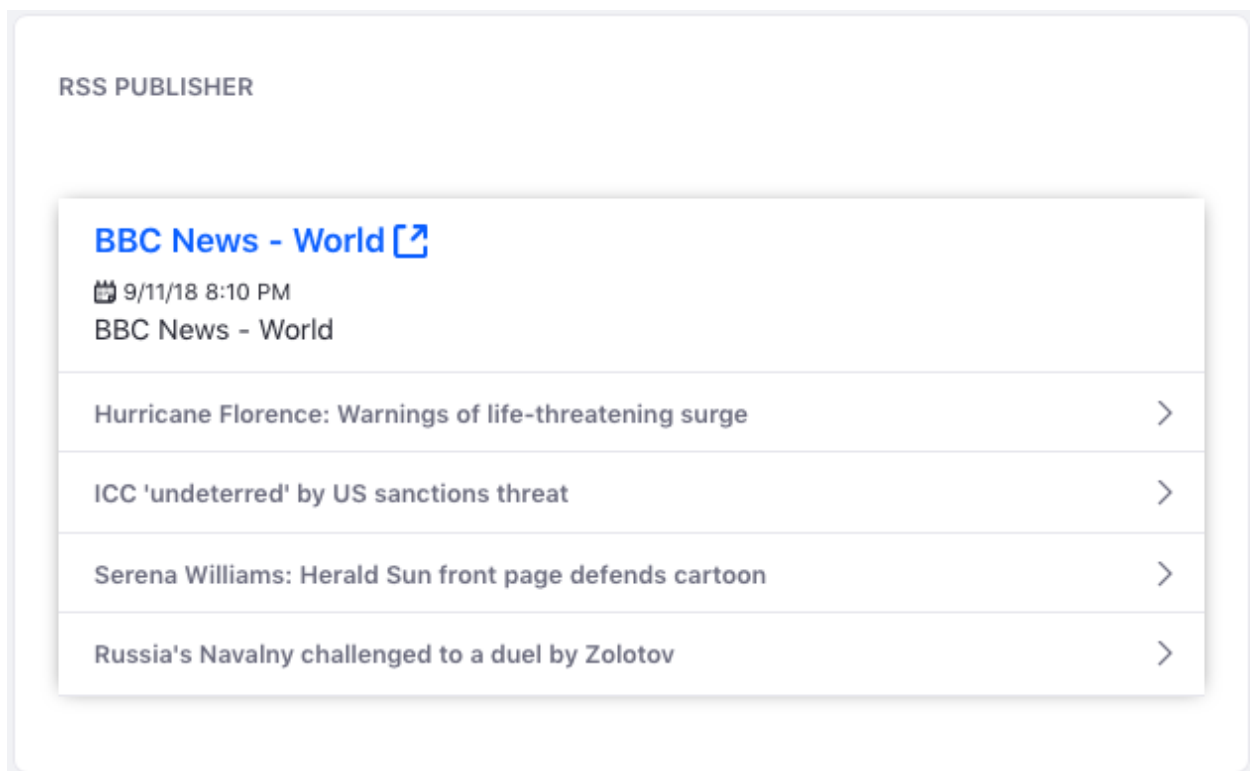


Figure 150.1: The RSS Publisher widget lets you display RSS feeds of your choosing.

Note: If you run your server behind a proxy, you must set the appropriate Java proxy settings (such as `http.proxyHost=` and `http.proxyPort=`) in your `setenv` script or in your `system-ext.properties`. Without these properties, the RSS Publisher widget can't access any RSS feeds.

Note that the RSS Publisher widget is deprecated. In Liferay CE Portal 7.1 GA2+, and Liferay DXP 7.1 FP4+, the widget is available from the *Add* (+) → *Widgets* → *News* menu. However, the widget is hidden in earlier releases of Liferay CE Portal 7.1 and Liferay DXP 7.1. In these releases, you must therefore make the widget visible via a configuration file. The next section shows you how to do this.

150.1 Using the RSS Publisher Widget

You can add the RSS Publisher widget to a page from the *Add* (+) → *Widgets* → *News* menu. Once you've done so, open the widget's Configuration menu by clicking on the *Options* icon (⋮) at the top-right corner of the widget and selecting *Configuration*.

By default, the RSS Publisher widget displays one feed. In the *Feeds* section, add or remove a feed via the plus or minus buttons, respectively. To add a feed, enter its URL and title in the respective fields. If you leave the *Title* field blank, the feed's default title is used (the *Title* field is for custom titles).

In the top section, use the following toggles to enable/disable the display of the feed's details:

- Show Feed Title
- Show Feed Published Date
- Show Feed Description
- Show Feed Image
- Show Feed Item Author

You can also select the number of entries and expanded entries that should be displayed per feed. Expanded entries show more of an article's actual content than regular entries. By default, each feed shows four entries per feed and eight expanded entries per feed. You can set the feed image alignment to control whether feed images appear to the right or left of the text. By default, the feed image alignment is set to *Right*.

DISPLAY SETTINGS ∨

Display Template
Default ∨ [Manage Templates](#)

Show Feed Title
 YES

Show Feed Published Date
 YES

Show Feed Description
 YES

Show Feed Image
 YES

Show Feed Item Author
 YES

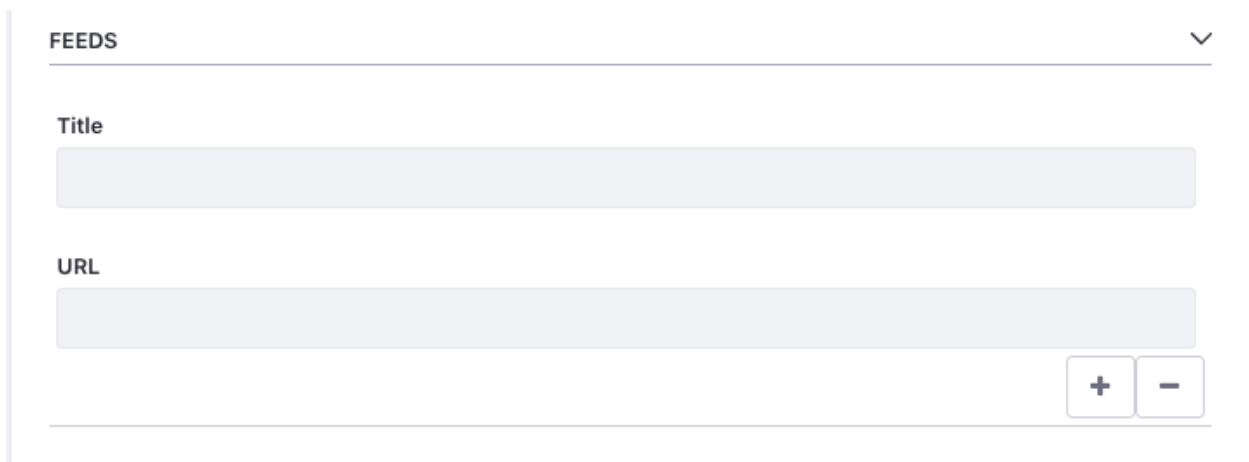
of Entries Per Feed
4 ∨

of Expanded Entries Per Feed
8 ∨

Feed Image Alignment
Right ∨

FEEDS >

Figure 150.2: The RSS Publisher widget's configuration lets you customize how the widget displays RSS feeds.



The image shows a configuration panel for the RSS Publisher widget. At the top, the word "FEEDS" is displayed in a bold, sans-serif font, followed by a horizontal line and a small downward-pointing chevron icon in the top right corner. Below this, the label "Title" is positioned above a wide, light gray rectangular input field. Further down, the label "URL" is positioned above another wide, light gray rectangular input field. At the bottom right of the configuration area, there are two small, square buttons: one with a plus sign (+) and one with a minus sign (-). A horizontal line is visible at the very bottom of the configuration panel.

Figure 150.3: You can also use the RSS Publisher widget's configuration to specify which feeds to display.

RESTORING DELETED ASSETS

Have you ever had that life-altering experience where you deleted an important file and immediately regretted deleting it? The deed is usually followed by a palm to the forehead or a sick feeling. Good news! Liferay DXP is here to turn that frown upside down with the *Recycle Bin* feature. With the Recycle Bin, the *Move to the Recycle Bin* action replaces *Delete* for certain asset types. Content is now temporarily stored in the Recycle Bin. This allows the content to be restored back to its original state. Recycled items can expire after a certain period of time, resulting in their permanent deletion. Before diving into how the Recycle Bin works, you'll look at how to configure it.

CONFIGURING THE RECYCLE BIN

The Recycle Bin supports instance-wide scope or site-specific scope. The instance-wide scope of the Recycle Bin is set by adding the `trash.enabled` property to your `portal-ext.properties` file. By default, the Recycle Bin is enabled instance-wide. You'll go into more detail for adding this property and several others to your properties file later in the section. First, you'll explore the UI and see what the Recycle Bin can do.

First, you'll configure the Recycle Bin for site-specific scoping.

1. Choose the Site you'd like configure for the Recycle Bin from the Site Administration menu.
2. Click *Configuration* → *Settings*.
3. Next, select the top *Advanced* tab and click *Recycle Bin*. You'll notice a few configurable options.

Enable Recycle Bin: enable and disable settings for the Recycle Bin's site-specific scope.

Trash Entries Max Age: customize the number of minutes a file is kept in the Recycle Bin until its permanent deletion (default is 43200 minutes, or 30 days).

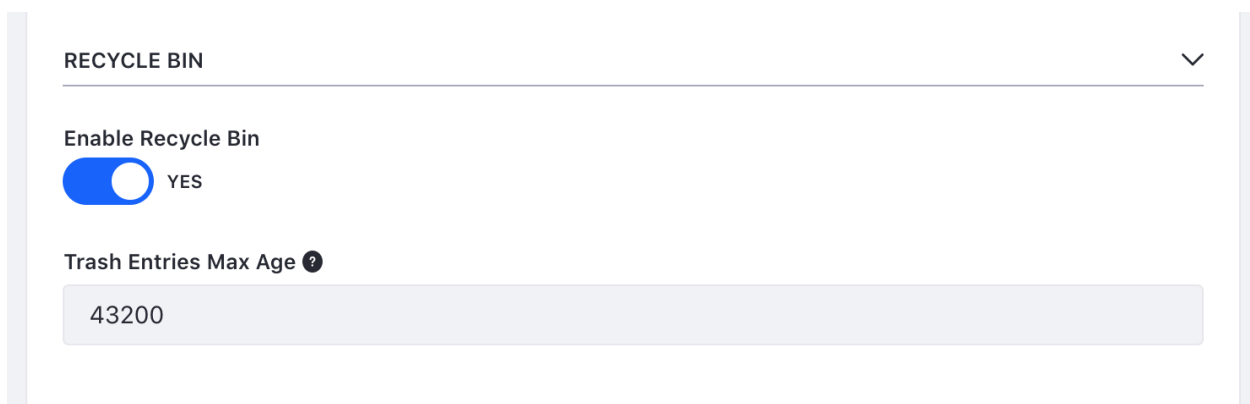


Figure 152.1: The Recycle Bin offers several configurable options for your site.

4. When you've finished configuring your Recycle Bin settings, click *Save*.

Note: If you disable the Recycle Bin while it's still holding recycled items, the recycled items remain stored and reappear in the Recycle Bin if it is re-enabled.

You can also configure the Recycle Bin via properties in the `portal.properties` file. Remember that it's a best practice not to edit the `portal.properties` directly, but to create a separate `portal-ext.properties` file containing the properties to override. There are some additional options not available in the GUI that you can set:

`trash.search.limit=500`: set the limit for results used when performing searches in the Recycle Bin (default is 500).

`trash.entry.check.interval=60`: set the interval in minutes for how often the trash handler runs to delete trash entries that have been in the Recycle Bin longer than the maximum age (default is 60).

Also, as was mentioned earlier, there are properties to enable the Recycle bin instance-wide and set trash entries' maximum age.

`trash.enabled=true`: set this property to *false* to disable the Recycle Bin for all sites in the portal (default is *true*).

`trash.entries.max.age=43200`: set the number of minutes trash entries should be held before being permanently deleted.

Visit the `portal.properties` file to view all of the configurable properties for the Recycle Bin.

Next, you should make sure permissions are set properly for users who can handle/view the assets in the Recycle Bin. Users who had *View* permissions on a document when it was recycled can also view that document in the Recycle Bin. Users who had *Update* or *Delete* permissions on a document when it was recycled can restore the document.

Now that you've successfully configured the Recycle Bin, you'll look at how to use it.

USING THE RECYCLE BIN

The Recycle Bin is temporary storage where assets go when you delete them. You can recycle several different types of assets:

- Blogs
- Bookmarks
- Documents and Media
- Message Boards (and attachments)
- Web Content
- Wiki (and attachments)

RECYCLE BIN

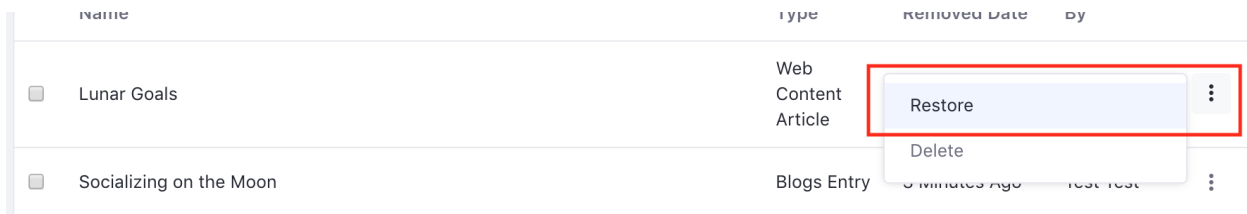
| Name | Type | Removed Date | Removed By | |
|---|---------------------|----------------|------------|---|
| <input type="checkbox"/> Lunar Goals | Web Content Article | 4 Minutes Ago | Test Test | ⋮ |
| <input type="checkbox"/> Socializing on the Moon | Blogs Entry | 3 Minutes Ago | Test Test | ⋮ |
| <input type="checkbox"/> Lunar Resort Homepage | Bookmarks Entry | 3 Minutes Ago | Test Test | ⋮ |
| <input type="checkbox"/> Lunar Resort Wiki | Wiki Node | 2 Minutes Ago | Test Test | ⋮ |
| <input type="checkbox"/> Lunar Resort Restaurant Menu | Document | 11 Seconds Ago | Test Test | ⋮ |

Figure 153.1: The Recycle Bin provides a seamless administrative experience for deleting and removing content.

Note: Attachments added to Wiki and Message Board entries do not go to the Recycle Bin when they are deleted. They can be restored in a similar fashion from the *Removed Attachments* menu within the application.

To demonstrate using the Recycle Bin let's delete a web content article and then restore it. You'll run through two different methods of restoring the file.

1. Go to Site Administration and select *Content & Data* → *Web Content*.
2. Select the *Add* button (+) and click *Basic Web Content*.
3. Enter some text for the Title and Content and click *Publish*.
4. Click the article's *Actions* button (⋮) and click *Move to the Recycle Bin*.
Note that the *Delete* button is not listed. Liferay DXP avoids the risk of accidental deletion of your files by funneling the content through the Recycle Bin.
5. After deleting the file, a success message appears, offering an *Undo* option. Click *Undo*. The web content is retrieved from the Recycle Bin and stored in its original place.
6. Click *Move to the Recycle Bin* again.
7. Go back to Site Administration and click *Recycle Bin* from the *Content* dropdown.
8. Find your sample web content and click its *Actions* button.
9. You can restore or delete the content. Select *Restore*.



| Name | Type | Removed Date | By |
|--|---------------------|---------------|-----------|
| <input type="checkbox"/> Lunar Goals | Web Content Article | | |
| <input type="checkbox"/> Socializing on the Moon | Blogs Entry | 5 minutes ago | test test |

Figure 153.2: In the Recycle Bin, you have the option of restoring or permanently deleting the content.

10. Navigate back to the Web Content screen and notice that your sample web content is back to its original place.

That covers the two general processes of sending and restoring content to/from the Recycle Bin. For other asset types, the Recycle Bin works similarly.

Some applications, such as Web Content and Documents and Media, support folders for organizing content. You can also send folders to the Recycle Bin. Keep in mind that this sends any sub-folders of the deleted folder all the files it contains to the Recycle Bin. Folders are restored and deleted the same way as a single file.

Delete within the Recycle Bin is the permanent delete button. Once you select this, your file cannot be retrieved and is gone forever. There is also an *Empty the Recycle Bin* option accessible from the (E) button at the top of the Recycle Bin screen. This permanently deletes all the files from the Recycle Bin.

153.1 Drag and Drop

You can also drag and drop items into the Recycle Bin. While you're in the Control Panel, select an asset and drag it to the Recycle Bin portlet on the Control Panel menu. When you click and begin dragging the asset, a message appears near your cursor notifying you of the number of files ready to be moved, and the Recycle Bin is highlighted, showing you where the files can be dropped. After you drop the asset onto the Recycle Bin portlet, the asset is removed from its original location and transferred to the Recycle Bin.

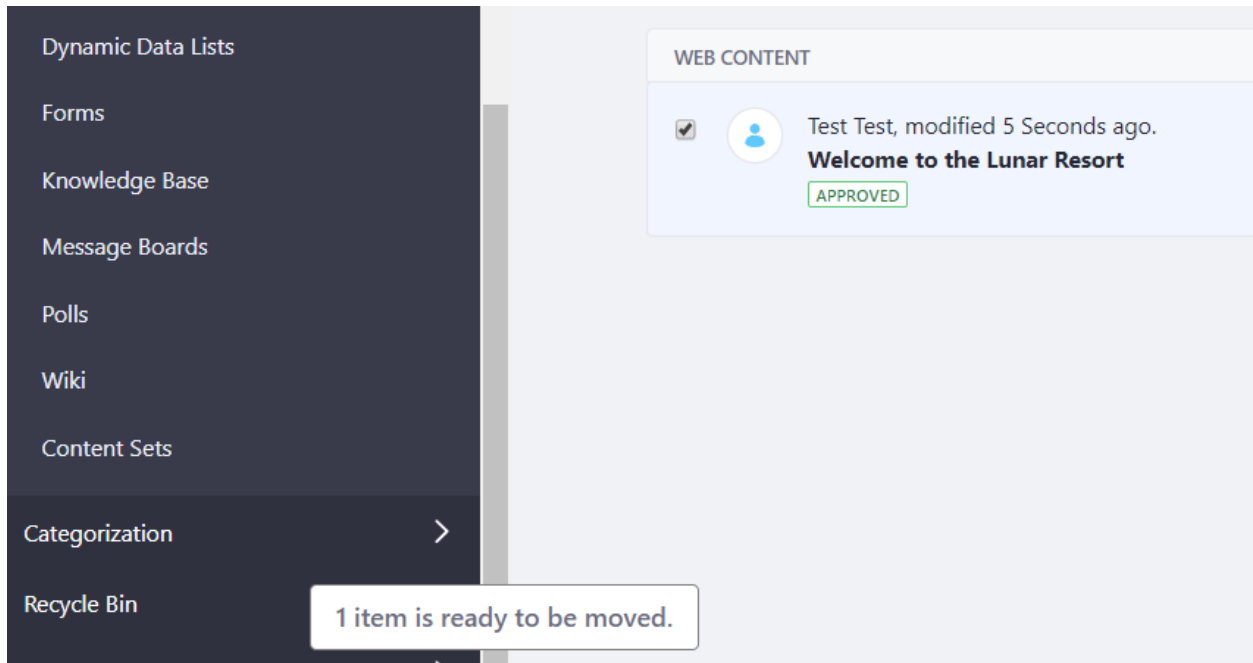


Figure 153.3: A quick and easy way of disposing your items is the drag and drop method.

Awesome! You now know how to use the Recycle Bin!

RECYCLE BIN INTELLIGENCE AND SUPPORT

Have you ever wondered what happens to file shortcuts if their linked assets are recycled? What if you restore a file that has the same name as another file currently stored in your site/instance? The Recycle Bin already knows how to handle these types of issues.

When documents with shortcuts are moved to the Recycle Bin, the shortcuts are removed. This ensures that all your links and shortcuts work and cuts down on maintenance time and backtracking.

Another important trait how recycled content is managed with the Staging framework.

Although you there is only one master Recycle Bin for all asset types, when staging is enabled a *Staging* Recycle Bin is created. The original Recycle Bin, or *Live* Recycle Bin, is still viewable while in staging; however, it is never used.

During staging, everything you recycle is sent to the Staging Recycle Bin. This prevents staged and unstaged recycled content from mixing. For example, if you have an unstaged document currently on your live site you can enable staging and delete that document. If you were to turn staging off and return to the live site, without separate Recycle Bins, the live document would be both on your site and in the Recycle Bin! Because of this, the separate Staging Recycle Bin is necessary and only used during the staging process. When you publish your staged material, the Staging Recycle Bin content is transferred to the Live Recycle Bin.

Note: The Staging Recycle Bin saves its contents until the staged material has been published to the live site. This means that you can turn the staging mode on and off without losing your recycled material.

The Recycle Bin saves you time by letting you restore content that's been recycled. Instead of recreating or re-uploading content, you'll be tailoring your Liferay instance to fully leverage its capabilities.

COLLABORATION

Liferay DXP contains an expansive collaboration suite that empowers users to create content and communities that they couldn't create alone. A robust document management system is a key component of this suite. As users produce digital assets—documents, videos, audio—they can store and share them using the Documents and Media Library. Documents and Media supports file check in and check out to prevent conflicting edits from multiple users, and maintains a version history of those files. It also contains its own repository, and for added flexibility can connect to external repositories. Once files exist in Documents and Media, users can insert them in other content like blog posts and wiki articles.

The collaboration suite also contains apps that let users share information and create active communities. The Message Boards app gives users a platform for discussions. The Blogs app lets users publish their ideas using rich content. Notifications keep users informed of what's happening. Social networking apps let users connect and share in ways that bolster friendship and productivity. And this is just scratching the surface—there are many more apps that help users communicate, produce, and present.

The guides that follow show you how to leverage these features, and more, in detail.






Figure 155.1: You can use the Documents and Media Library to manage and use documents in the portal.

BLOGS [RSS](#) [Subscribe](#) [New Entry](#)

Lunar Spelunking

 **Joe Bloggs**
4 Hours Ago



Lunar spelunking is a lot of fun, but it's not for the faint of heart. Nor is spelunking on Earth. A low-gravity environment, however, increases the difficulty. And although the science on this is unsettled (or perhaps just unsettling), there may be strange, fanged creatures dwelling beneath the lunar surface. An early member of the Apollo missions was said to have almost caught one, but its...

[Comment \(0\)](#) [👍 0](#) [👁️ 0](#) [Twitter](#) [Facebook](#) [LinkedIn](#) [Share](#)

Figure 155.2: You can also make your blog entries look great.

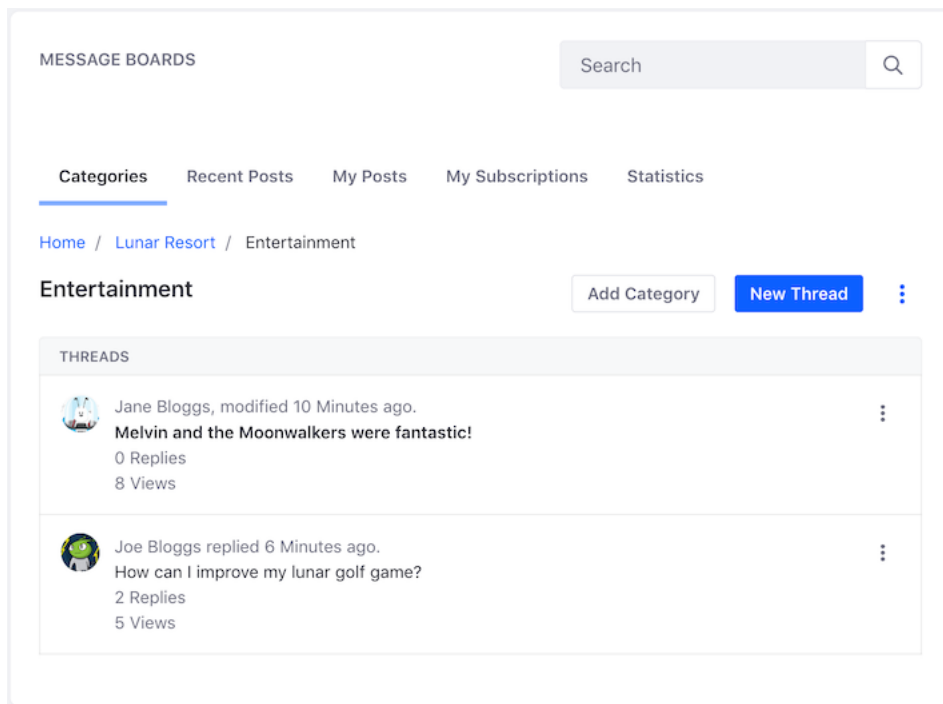


Figure 155.3: The Message Boards app is fantastic for facilitating discussions.

MANAGING DOCUMENTS AND MEDIA

The Documents and Media library stores files on the server using the same type of structure that you use to store files locally. It accepts files of any kind, can serve as a virtual shared drive, and can mount and browse external repositories. You can organize documents using customizable document types and metadata sets and display them with automatic document preview generation. Its companion app, the Media Gallery, displays selected content from the Documents and Media library. It can render image, audio, and video files.

Liferay Sync synchronizes Documents and Media folders with local folders on your devices, both your desktop machines and mobile devices.

You'll get started with Documents and Media by exploring how to publish files.

PUBLISHING FILES

As you create sites, you'll probably want to share files on them. The Documents and Media library (Document Library) lets you upload and publish all kinds of files on your sites. Pictures, videos, spreadsheets, slide presentations, and more can be stored in and shared from the Document Library. Document Library instances can be scoped to a portal instance, site, or page, so you can work with files where they're relevant.

Here, you'll learn how to add files, display them, and collaborate on them. You'll learn how to use both the Documents and Media Library and the Media Gallery. And lastly, you'll learn how to collaborate on files from within several environments, including your browser and local desktop file system.

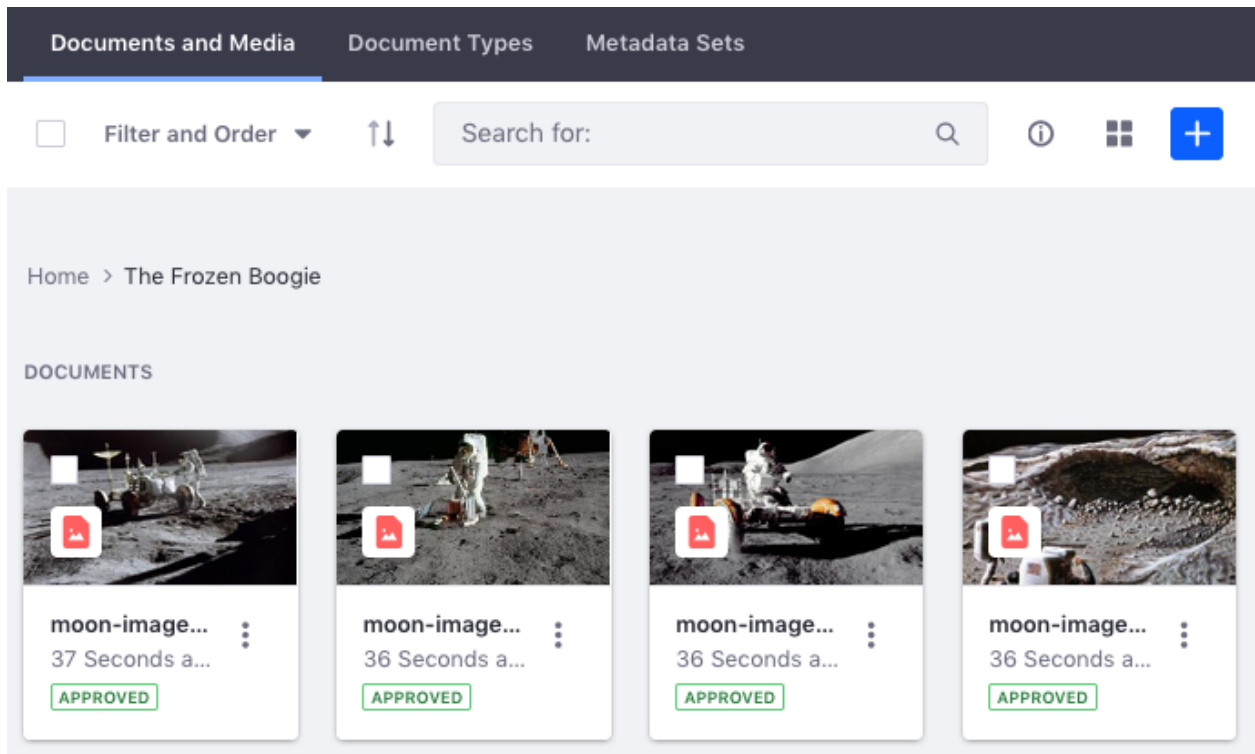


Figure 157.1: These documents are awesome.

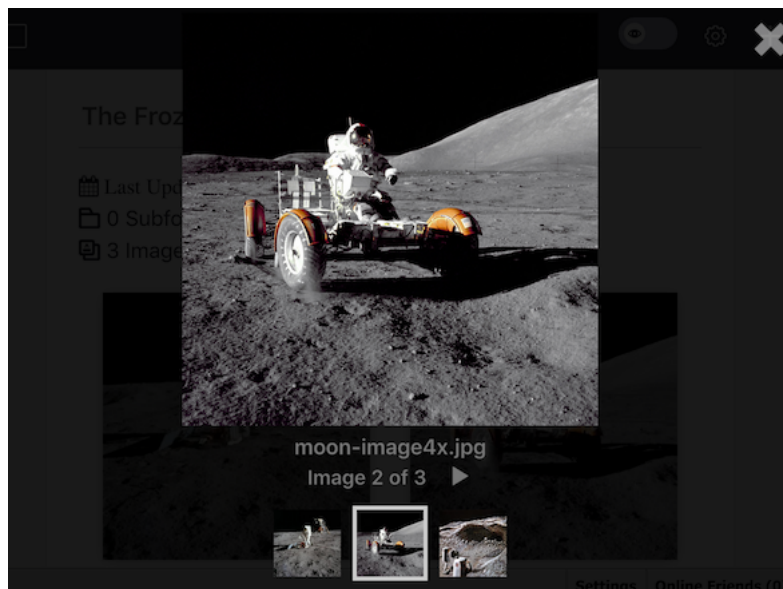



Figure 157.2: This slideshow rules.

moon-image3x.jpg


Share Download



moon-image3x.jpg

VERSION 1.0 APPROVED

Details Versions



Joe Bloggs
Owner

Download

Latest Version URL

http://localhost:8080/documents/20'

WebDAV URL ⓘ

http://localhost:8080/webdav/guest/

DOCUMENT TYPE
Basic Document

EXTENSION
jpg

SIZE
126KB

MODIFIED
4/1/19 7:51 PM by Joe Bloggs

CREATED
4/1/19 7:50 PM by Joe Bloggs

Comments

Subscribe

Type your comment here.

Reply

Figure 157.3: Viewing a file's details is fun.

ADDING FILES TO A DOCUMENT LIBRARY



This article covers the following topics to help you get started adding files to your Document Library:

1. **Granting File Permissions and Roles:** Determine who can add, view, and update files. Doing this before adding files ensures that only those you wish can access your Document Library.
2. **Adding Files:** Add specific types of files and their associated metadata to your Document Library.

158.1 Granting File Permissions and Roles

You should carefully manage who can add, view, and update files. You can store files of all kinds for various purposes. For example, you may have one set of files intended for only specific site members and another intended for everyone, including guests. You can use Roles and Permissions to control access to Document Library files. The Document Library's folder permissions also help you organize files.

Follow these steps to create a Role for managing files in your site's Documents and Media:

1. Open the *Menu* () and navigate to *Control Panel* ▾ *Users* ▾ *Roles*.
2. Select the *Site Roles* tab (or *Organization Roles*, for an Organization Role) and then click the *Add* button () to begin creating a role.
3. Give your Role a name and a description, then click *Save*.
4. Select your Role's *Define Permissions* tab. In the Role's permission definition screen, navigate to *Site Administration* → *Content & Data* → *Documents and Media*. In the *General Permissions* section, select *Access in Site Administration* and click *Save*.
5. Assign this Role to the Users that should manage media. For more information on this and other topics related to Roles, see *Roles and Permissions*.

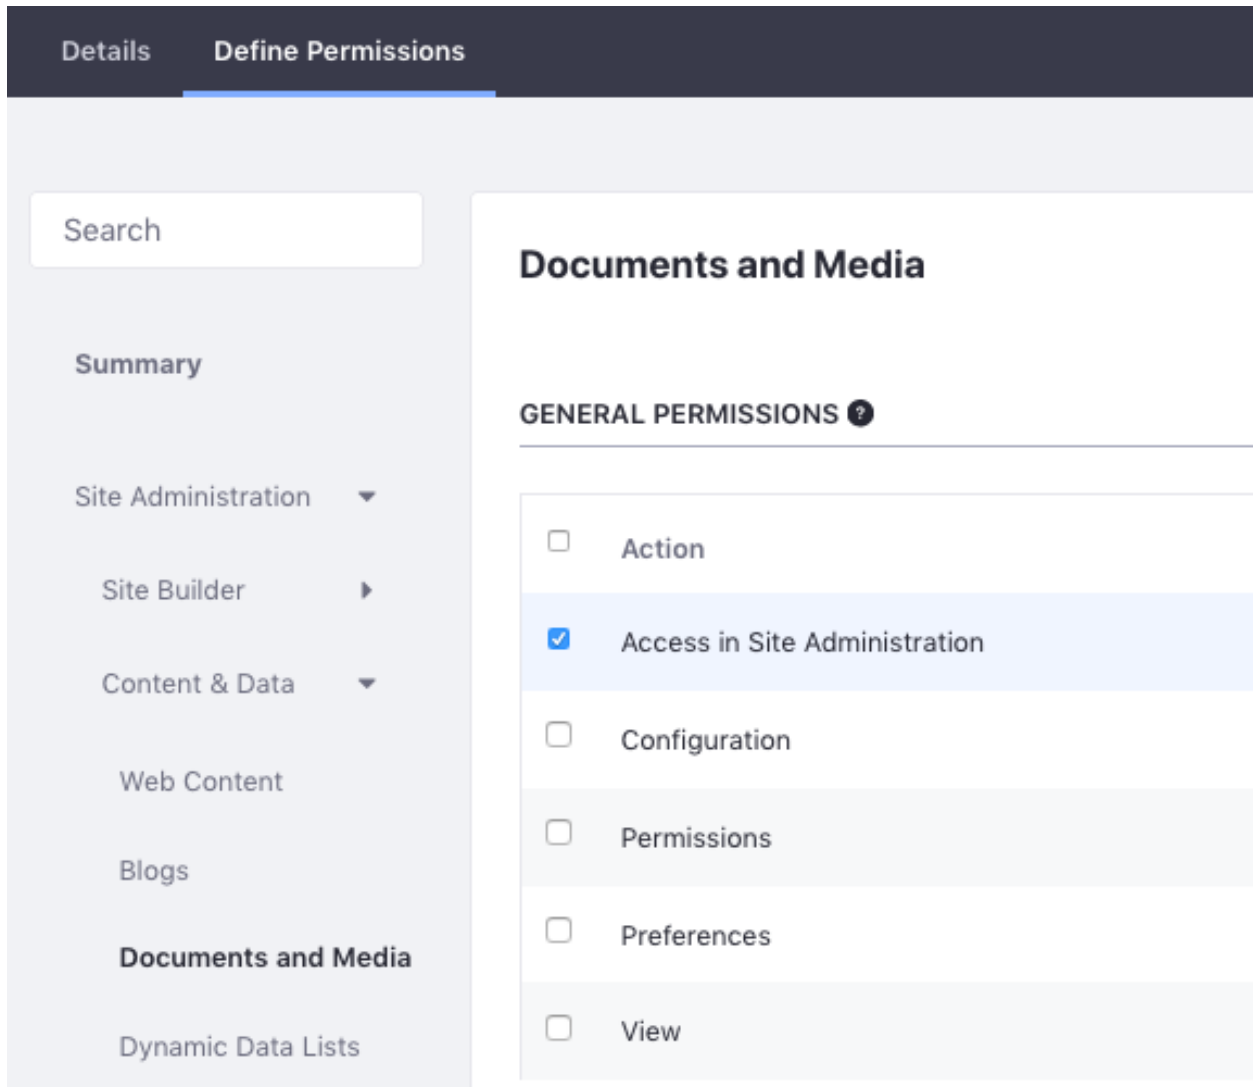




Figure 158.1: It's often helpful to define a role for specific users to access Documents and Media from Site Administration.

158.2 Using the Add Menu

Follow these steps to add files to your site's Document Library:

1. Open the *Menu* () , click on your site's name, and navigate to *Content & Data* → *Documents and Media*. The Documents and Media screen appears and displays the Documents and Media library's *Home* (its root folder). As you add files and folders to the Document Library, they're listed here.
2. Click the *Add* icon () and select the type of document to add to the Document Library. You can add documents, folders, and shortcuts much like you would on a desktop file system.

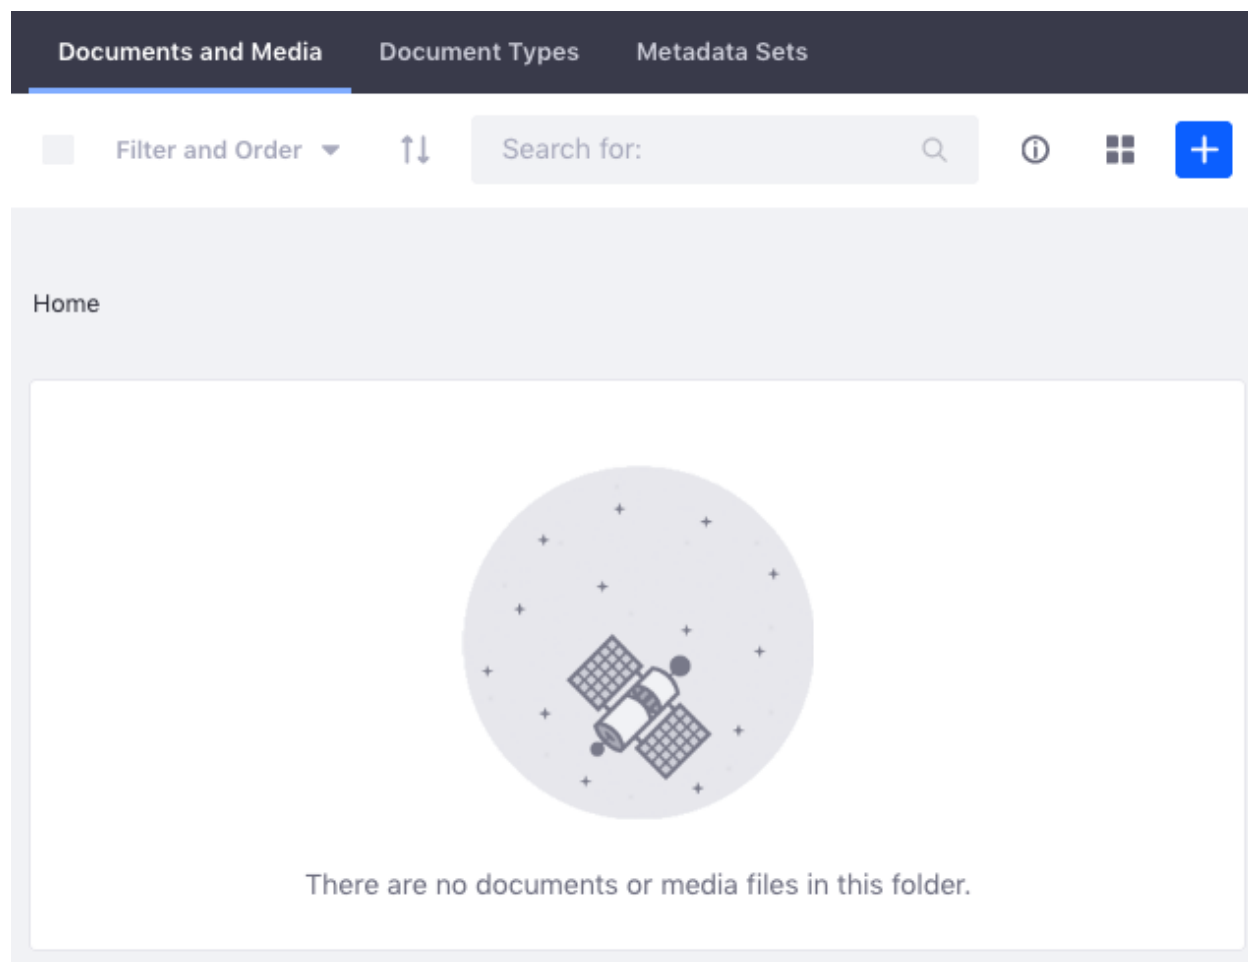


Figure 158.2: The Documents and Media's *Home* folder starts empty.

You can even configure access to an entirely different repository. The Add menu's options are described below.

3. When you're finished selecting the file to upload and filling out any document type fields that are necessary, click *Publish*.

File Upload: Upload a file to the Documents and Media library.

Folder: Create a new folder in the Documents and Media library's file system.

Multiple Files Upload: Upload several files at once. You can apply a single description and document type to all the files. You can also categorize and tag the files, and assign them default permissions.

Repository: Add access to an external repository.

Shortcut: Create a shortcut to any document that you can view. You can set permissions on the shortcut to specify who can access the original document via the shortcut.

Any additional items in the Add menu are document types described by a unique metadata set. When you add a document belonging to a document type, a form appears that lets you pick the file to upload and enter the data defined by the document type's metadata set.

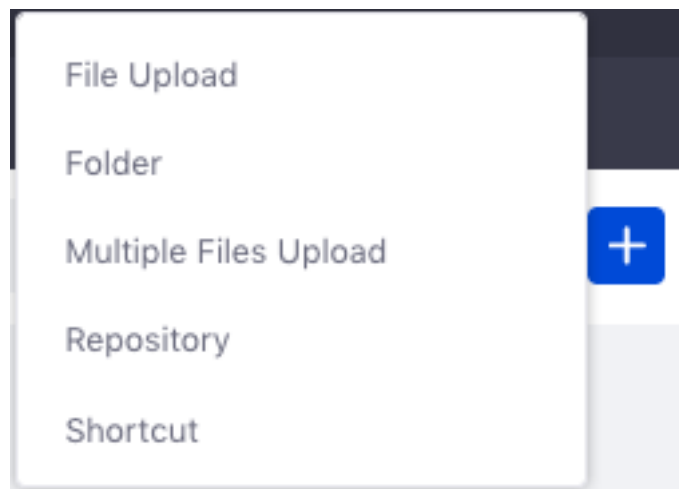


Figure 158.3: The Add menu lets you upload and add all kinds of documents to the library.



CREATING FOLDERS

You'll need folders to organize all but the most limited set of files. Here, you'll learn how to work with folders in a Document Library:

- Adding a Folder
- Document Type Restrictions and Workflow
- Setting Folder Permissions

159.1 Adding a Folder

Follow these steps to add a folder:

1. Open the *Menu* () , click on your site's name, and navigate to *Content & Data* → *Documents and Media* for your site. The Documents and Media screen appears and displays the Documents and Media library's *Home* (its root folder).
2. Click the *Add* icon () and select *Folder*. The New Folder form appears.
3. In the New Folder form, name and describe your folder. Then expand the *Permissions* section.
4. In the Permissions section, set the folder's permissions. The *Viewable by* menu lets you select who has View permission for the folder:
 - Anyone (the Guest role; this is the default option)
 - Site Members
 - Owner

Click the *More Options* link to choose the other folder permissions for the Guest and Site Member roles. By default, site members can add files, subfolders, shortcuts, and subscribe to changes to the folder's files. Guests don't have any such permissions, which is typically what you want.

Permissions ▼

Viewable by


Anyone (Guest Role) ⌵

[Hide Options](#)

| Role | Delete | Permissions | Add Subfolder | Add Shortcut | Update | Add Document | Subscribe | View | Access |
|-------------|--------------------------|--------------------------|-------------------------------------|-------------------------------------|--------------------------|-------------------------------------|-------------------------------------|-------------------------------------|--------------------------|
| Guest | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |
| Site Member | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input type="checkbox"/> |

Figure 159.1: Select your folder's permissions.

- To finish creating the folder, click *Save* after making your selections in the Permissions section.

Upon creating the folder, it appears in your Document Library. Opening the folder's *Actions* menu () presents several options for managing the folder. The following sections describe some of these options.

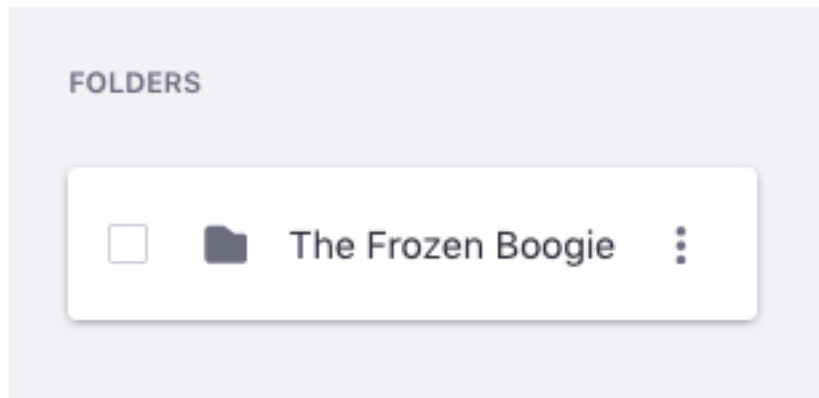


Figure 159.2: Your new folder appears in the Document Library.

159.2 Document Type Restrictions and Workflow

After creating a folder, you can restrict what document types are allowed in it. You can also choose what workflow (if any) to use for approving files added to or edited in the folder.

Follow these steps to change a folder's document type restrictions and workflow:


1. Click the folder's *Actions* menu () and select *Edit*.
2. Expand the *Document Type Restrictions and Workflow* section. In this section, choose from the following options:
 - Use Document Type Restrictions and Workflow of the Parent Folder (the parent folder)
 - Define Specific Document Type Restrictions and Workflow for this Folder (the current folder)
 - Default Workflow for this Folder (the current folder)
3. Click *Save* when you're finished.




Figure 159.3: You can set the document type restrictions and workflow to use for a folder's files.

159.3 Setting Folder Permissions

When creating a folder, you can set some of its permissions via the new folder form. Fine tuning a folder's permissions, however, can only be done after creating the folder.

Follow these steps to fine tune a folder's permissions:

1. Click the folder's *Actions* menu () and select *Permissions*. The Permissions window appears.
2. In the Permissions window, set the permissions you want to use for this folder. The permissions listed below are available for each role.
3. Click *Save* when you're finished setting permissions.

Here are the permissions you can set:

Delete: Move the folder to the Recycle Bin.

Permissions: View and modify the folder's permissions.

Add Subfolder: Create folders within the folder.

Add Shortcut: Create a shortcut (link) to any file in the folder that the role is authorized to view.

Update: Edit the folder's attributes and/or move the folder under a new parent folder.

Add Document: Add a new file to the folder.

Subscribe: Receive email notifications when files are added to or modified in the folder. Note that you can specify the email sender and template from the Documents and Media's *Options* (ⓘ) → *Configuration* menu.

View: View the folder.

Access: Access the folder's contents.

USING THE DOCUMENTS AND MEDIA MANAGEMENT BAR

The Documents and Media *Management Bar* is where people who manage documents go to unwind after a long day at work. Just kidding. The Management Bar, as its name implies, contains tools for managing the files and folders in your Document Library. It appears above the files and folders in Documents and Media.




Figure 160.1: The Management Bar is a great place to hang out if you're managing documents.

If you've added files or folders to your Document Library, then you're already familiar with the Management Bar's *Add* button (⊕). The sections that follow describe the rest of the Management Bar.

Note: If a Document Library contains more items than it can display at once, you can use the navigation tool that appears at the bottom of the window to switch your view to another page or configure the page to display more items per page.

160.1 View Types

The *View Types* button is to the left of the *Add* button. It lets you choose how to display the Document Library's items. The *View Types* button's icon depends on the selected view type:

Cards (■): Shows a card-like rendering of the item. If the item isn't an image, a generic image for the item's type is displayed. For files, each card also contains the file's suffix (e.g., JPG, PNG, etc.), timestamp, name, and workflow status (e.g., Approved, Draft, etc.).

List (☰): Shows the same information as the Cards view type, in a list with small file renderings.

Table (☒): Shows the same information as the other view types, in a list with no file renderings. Also, the file information is in columns.

The items in all view types have an Actions menu (⋮). These actions are also available in when viewing each item separately.

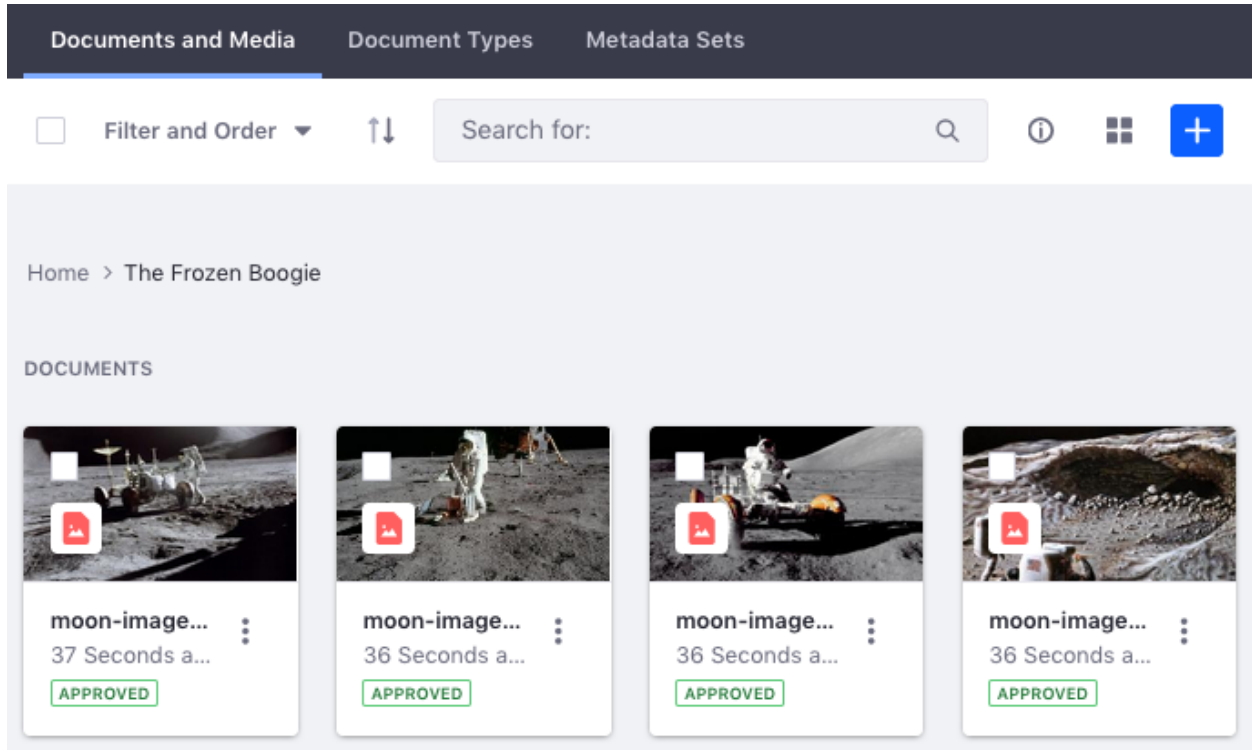


Figure 160.2: The Cards View type shows items in large card-like renderings.

160.2 The Info Panel

To display an info panel with the current folder's details, click the *Information* icon (ⓘ). The info panel slides out from the right side of the screen and contains the folder's name and number of items. It also has these buttons:

Subscribe (☆): Get notifications about files added to or modified in the folder.

Actions (⋮): Lists actions you can perform on the current folder.

160.3 Finding and Arranging Items

The Management Bar also contains tools that help you locate and arrange items in the Document Library. The most prominent of these tools is the *Search* bar, where you can find files by keywords.

To the left of the Search bar, the Sort button (↑↓) arranges items in ascending or descending order.

You can also arrange items via the *Filter and Order* selector using these criteria:

All: Shows all of the current folder's immediate subfolders and files (default).

Mine: Shows all the current user's files (no matter their folder).

Document Type: Shows the files of the selected document type. Upon choosing this option, you must select the document type you want from a popup.

You can also select from the following criteria for ordering items:

- Size
- Downloads
- Modified Date (default)
- Create Date
- Title

160.4 Selecting Items

The checkbox on the left-most side of the Management Bar selects all currently displayed items. Selecting multiple items lets you act on all of them at once. You can also select multiple items individually by using the checkboxes for each. When you select one or more items, the Management Bar changes to reflect the actions you can take on the selected items.




Figure 160.3: With items selected, the Management Bar changes.

Here are the actions you can take on the selected items:

- Download (↓)
- Move (↔)
- Edit Tags
- Move to Recycle Bin (🗑)

The Actions button (⋮) contains all the actions displayed in the Management Bar, plus actions for file checkin and checkout. File checkout and checkin is explained in *Checking out and Editing Files*.

VIEWING FILE PREVIEWS

File previews help users browse and find media efficiently. To view a preview of a file, click the file's name in the Document Library. If the file is an image, the image appears. If an app is installed that can render a preview of the file type, a representative image of the file appears (e.g., the opening frame of a video file or a presentation's first slide). If there are no such preview apps for the file, a generic image based on the file type appears.

161.1 File Preview Apps

Whenever possible, Liferay DXP generates previews of documents added to the Document Library. Out of the box, Java-based APIs generate previews. The only tool available that is 100% Java and has a compatible license to be distributed with Liferay DXP is PDFBox. A separate thread generates a preview for PDFs when uploaded. This process may last only a few seconds for a small file. The larger the file, the longer it takes.


While PDFBox provides a default implementation of image generation for document previews and thumbnails, you must install and configure additional tools to harness the full power of document previews. These tools include:

- **OpenOffice or LibreOffice:** Using one of these in server mode lets you generate thumbnails and previews for supported file types (.pdf, .docx, .odt, .ppt, .odp, etc.), view documents in your browser, and convert documents.
- **ImageMagick (also requires Ghostscript):** Enables faster and higher-quality previews and conversions.
- **Xuggler:** Enables audio and video previews, lets you play audio and video files in your browser, and extracts thumbnails from video files.

After installing these tools, you can configure them via portal properties in the Control Panel's Server Administration screen, or in a `portal-ext.properties` file. To learn how to use these tools, see [Configuring Liferay DXP](#).

moon-image3x.jpg


Share Download



moon-image3x.jpg

VERSION 1.0 APPROVED

Details Versions



Joe Bloggs
Owner

Download

Latest Version URL

<http://localhost:8080/documents/20>

WebDAV URL ⓘ

<http://localhost:8080/webdav/guest/>

DOCUMENT TYPE
Basic Document

EXTENSION
jpg

SIZE
126KB

MODIFIED
4/1/19 7:51 PM by Joe Bloggs

CREATED
4/1/19 7:50 PM by Joe Bloggs

Comments

Subscribe

Type your comment here.

Reply

Figure 161.1: File previews let you view and manage a file.

With these tools installed and configured, a customized viewer displays Documents and Media content, depending on the content type. For example, you can view a document with a customized viewer that lets you navigate through the document's pages. You can also view and play multimedia documents (audio or video). If the browser supports HTML5, the viewer uses the browser's native player. Otherwise it falls back to a Flash player.

161.2 Managing Files

You can also manage a file from its preview. The bar above the preview contains these buttons:

Info (ⓘ): Open/close the file's info panel. This panel contains more detailed information about the file. For more information on this, see [The Info Panel](#).

Share: Share the file with other users. For more information, see [Sharing Files](#).

Download: Download the file.

Actions (⋮): Opens a menu that lets you perform these actions on the file:

- **Download**
- **Edit**: Modify the file's name, description, document type, categorization, and related assets. You can even upload a new file to replace it. Note that modifying the file increments its version.
- **Edit with Image Editor**: Edit the image in the Image Editor. The Image Editor is explained in [Editing Images](#).
- **Checkout/Checkin**: Checkout prevents others from editing the document while you are working on it. Other users can still view the current version of the document, if they have permission. You can check in the document when you're done with it.
- **Move**: Relocate the file to a different parent folder.
- **Permissions**: Specify which actions each role can perform on the file.
- **Move to Recycle Bin**: Move the file from the Documents and Media library to the Recycle Bin.
- **Share**

Also note that the *Options* menu (⚙) at the top-right of the screen contains the same actions as the Actions menu.

The comments area (below the preview area) lets you comment on and subscribe to comments on the file.

161.3 The Info Panel

As mentioned above, clicking the *Info* icon (ⓘ) opens the info panel. The top of the info panel displays the file's name, version, and workflow status. There are two tabs in the info panel: *Details*, and *Versions*. *Details* is selected by default and shows the following:

Owner: The file's owner.

Download: A button to download the file.

Latest Version URL: A URL to access the newest version of the file.

WebDAV URL: A WebDAV URL for accessing the file via a desktop.

Document Type: The file's document type.

Extension: The file's extension (e.g., JPG, PDF, etc.).

Size: The file's size on disk.

Modified: The user that last modified the file, and when it was last modified.

Created: The user that created the file, and when it was created.


Ratings: The file's average user rating.


Automatically Extracted Metadata: Any and all metadata automatically extracted from the file.


When adding new documents or viewing existing documents, a process is triggered automatically that extracts the file's metadata. The library used by this process is TIKKA and it's included out of the box. Depending on your file's type and the metadata written with the file, you can find out all kinds of details. In the case of audio or video files, the media's duration is displayed.

To instead view the file's version history, select the *Versions* tab near the top of the info panel. The info panel then changes to list the different versions of the file and lets you view, download, remove, and revert to specific file versions. File version history actions are explained in *Checking Out and Editing Files*.

EDITING IMAGES

Editing and re-uploading images when you only need to apply simple edits is tedious. Docs & Media contains a simple built-in image editor for exactly this reason. To access the image editor, locate the image you want to edit. Click the Actions icon () and select *Edit With Image Editor*.

You can also access the image editor when selecting an image to insert in content (i.e., via an item selector). Anywhere an image is, you can edit it. For example, you can access the image editor via item selector preview windows in blog entries and web content articles. To do this, click the pencil icon () in the bottom-right corner of the preview window.

If you edit and save the image via the Documents and Media repository, the file version is incremented a minor version (e.g., from version 1.0 to version 1.1). You can view the image's version history (and previous versions) by clicking the image, clicking its *Info* button (), and then selecting the *Versions* tab. In contrast, if you edit and save an image via an item selector, a copy of the image is created and saved to the Document Library.


Liferay designed the image editor with quick editing in mind. It offers a minimal, user-friendly UI. The main toolbar consists of three buttons, each of which contains a subset of options:

Transform ()

- **Rotate:** Rotate the image to the left or right, in 90 degree increments.
- **Resize:** Resize the image. If the lock is closed, the aspect ratio is locked and changing width or height automatically adjusts the other dimension to maintain the aspect ratio. When the lock is opened, the width and height can be changed individually, letting the aspect ratio change (this isn't recommended because the image can become distorted).
- **Crop:** Crop the image.

Adjustment ()

- **Saturation:** Adjust the color saturation. The default value is 50. Values range from 0 (completely desaturated) to 100 (completely saturated).
- **Contrast:** Adjust the contrast. The default value is 50. Values range from 0 (no contrast) to 100 (full contrast).
- **Brightness:** Adjust the brightness. The default value is 50. Values range from 0 (completely black) to 100 (completely white).

Filter (): Apply a filter to the image.

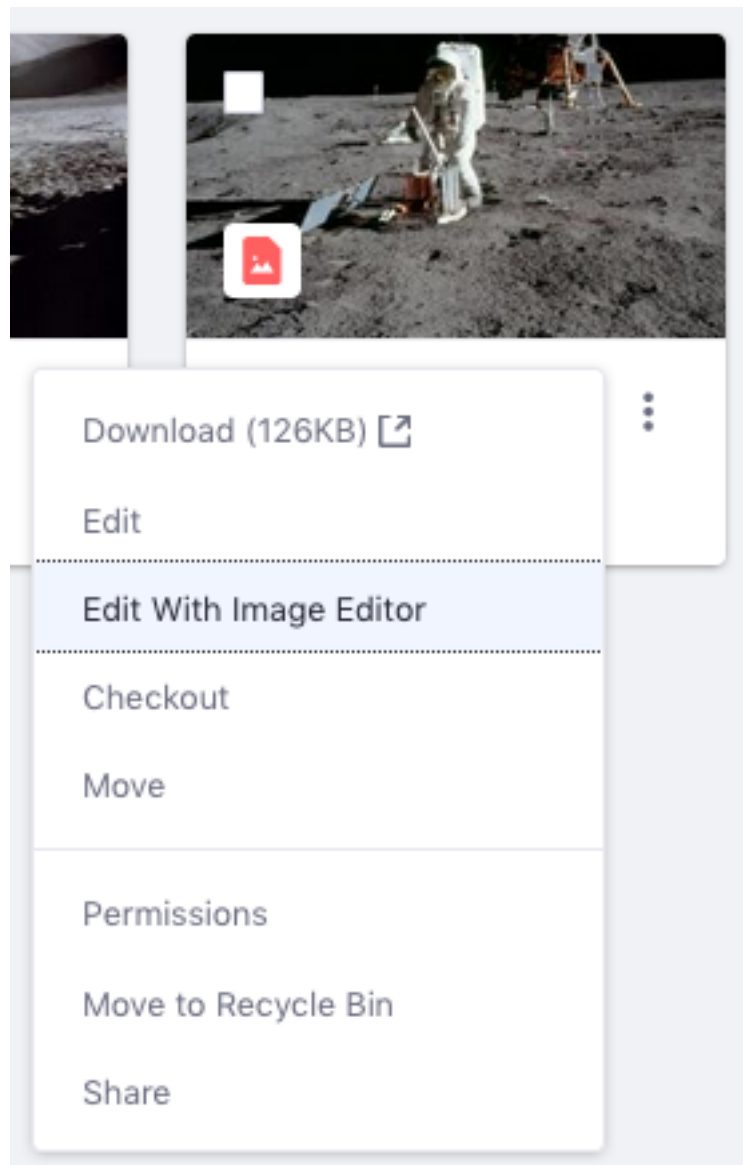


Figure 162.1: You can access the image editor through the Documents and Media repository.




Figure 162.2: You can also access the image editor through the item selector preview window.




Figure 162.3: The image editor's UI is clear and to the point, offering only what you need.




Figure 162.4: Select from a set of preset image filters.

Upon editing the image in the editor, you can click the *Cancel* button to cancel the changes, or the *Apply* button to apply them. Upon applying the changes, the history bar appears. It lets you undo, redo, or reset the changes. Use the Reset button with caution; it resets the image to its original state, reverting all changes made in the editor.




Figure 162.5: The history bar lets you undo, redo, and reset changes.

PUBLISHING FILES

Once your Document Library contains files, you may want to publish them in your site. Here are some ways to publish files:



- Show them in a Documents and Media app.
- Display them in a Media Gallery.
- Use the Asset Publisher.
- Insert them in an asset like a web content article or blog entry.

Here, you'll learn to use the Media Gallery.

163.1 Using the Media Gallery

The Media Gallery publishes your media files in a simple gallery-like style. It shows a large thumbnail of each media file, lets the user download files, and has slideshow capabilities. A common way to use the Media Gallery is to create a separate page for displaying media and add a Media Gallery widget to it. This way, your media takes center stage.

Follow these steps to create a page that contains a Media Gallery widget:

1. Create a page and navigate to it in your site.
2. At the top-right of the screen, click the *Add* icon () then navigate to *Widgets* → *Content Management* and select *Add* next to *Media Gallery* (alternatively, drag the Media Gallery onto your page). The Media Gallery widget appears on the page.
3. Configure the Media Gallery widget to show your files. By default, it shows files from the Home folder of your site's Documents Library. To choose a different folder, click the widget's Options icon () and select *Configuration*.

The Configuration window appears and shows the *Setup* tab. This tab contains these sections:

Display Settings: Lets you show each file's actions, filter the media types to display, and choose a display template for your media.

Folders Listing: Lets you select a Document Library folder to serve as the root folder from which to display files. The root folder you select becomes the highest-level folder the Media Gallery can access. For example, if you create a subfolder of a parent folder, and then set that subfolder as the Media Gallery's root folder, the Media Gallery can no longer access the parent folder.

****Note**:** To access the Carousel display template in Media Gallery, your role must have View access for that template. Since the Carousel template is in the Global scope, a Global-scope administrator must grant the role permission to view the template.

! [You can configure the Media Gallery to use any Documents and Media folder as its root folder.] (./images/dm-select-root-folder.png)

4. Configure the rest of the settings as desired in the Media Gallery app's other configuration tabs:

Communication: Lists public render parameters the widget publishes to other widgets on the page. Other widgets can take action on these parameters. For each shared parameter, you can specify whether to allow communication via the parameter and select which incoming parameter can populate it.

Sharing: Embed the widget instance as a widget on any website, Facebook, Netvibes, or as an OpenSocial Gadget.

Scope: Specify the Document Library instance the widget uses: the current site's instance (default), the global instance, or the page's instance. If the page doesn't already have an instance of the widget, you can select *Your Page (Create New)* to create a page-scoped instance for the widget to display.

5. Click *Save* when you're finished configuring the Media Gallery widget.

The Media Gallery now shows your files, with images appearing as thumbnails. When you click a thumbnail, a slideshow appears showing the selected image. Below that image, thumbnails of the folder's other images are displayed. The slideshow continues until you click pause or view the last image. Closing the slideshow window returns you to the page.

MEDIA GALLERY

The Frozen Boogie

Last Updated 4/1/19

0 Subfolders

4 Images




Figure 163.1: The Media Gallery renders large thumbnail images of media files.

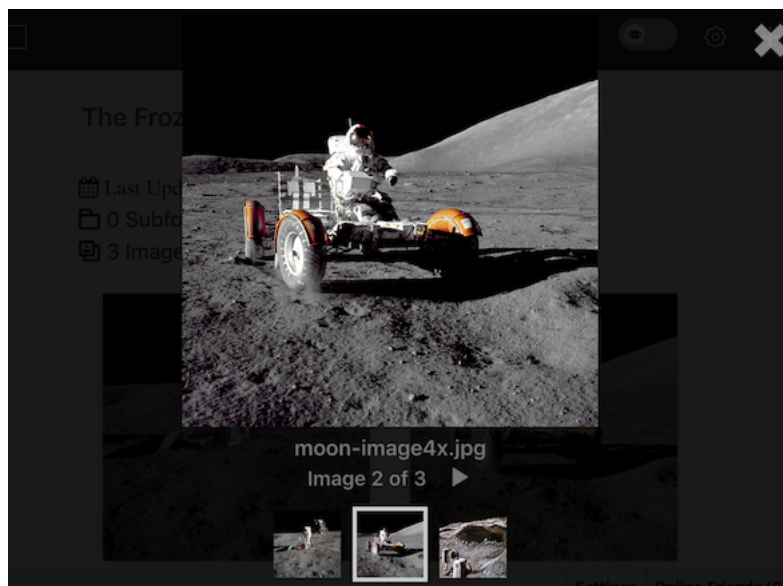



Figure 163.2: The Media Gallery's slideshow provides a nice way to view images.

CHECKING OUT AND EDITING FILES

When you check out a document in the Document Library, only you can make changes to it until you check it back in. This prevents conflicting edits on the same document by multiple users. When you check out a file, you can download it, replace it, move it to another Document Library folder, check it in, or cancel the checkout. Checking in a file also increments its version, which lets you keep track of changes.

Unless you're using Liferay Sync or a local drive mapped to the file's WebDAV URL, follow these steps to edit a Document Library file from your machine:

1. Checkout the file by clicking its Actions icon () → *Checkout*. Upon checkout, the file's status changes to Draft and a lock icon appears on the file.

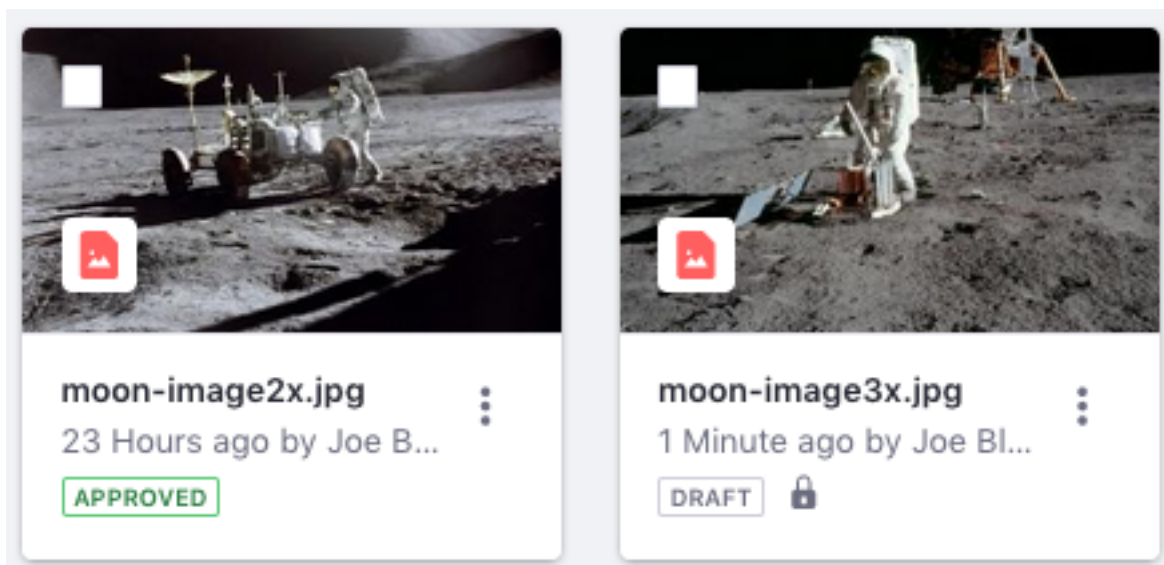





Figure 164.1: The file on the right in this image is checked out.


2. Download the file by clicking its Actions icon () → *Download*.
3. Edit the file locally.

4. Return to the Documents and Media Library and click the file's Actions icon () → *Edit*. The file's edit screen appears.
5. From the file's Edit screen, select the edited local file for upload.
6. Click *Save and Check In*. In the pop-up that appears, select whether your change is a major or minor version, add any version notes that you need, and click *Save*.

Note: If you edit a file without checking it out, the file's edit screen displays a toggle for *Customize the Version Number Increment and Describe My Changes*. Setting this to *YES* lets you specify the version increment's type and description.

Follow these steps to access a file's version history:

1. Click the file in the Documents and Media Library.
2. Click the file's *Info* button () at the top-right of the screen. This opens the file's info panel.
3. Select the *Versions* tab in the info panel.

Each file version has an Actions menu () that you can use to perform the following actions on that file version:

Download: Download the selected version of the file to your machine.

View: View the file entry screen for the selected version of the file.

Revert: Restores the selected file version as a new major file version. Note that this option isn't available for the newest file version.

Delete Version: Remove the file version from the Document Library. All other file versions remain intact.

moon-image3x.jpg

VERSION 1.1

APPROVED

Details

Versions

Version 1.1

By Joe Bloggs, on 4/2/19 7:02 PM

No Change Log



Version 1.0

By Joe Bloggs, on 4/1/19 7:50 PM

No Change Log



Figure 164.2: The version history actions let you inspect, delete, and reinstate file versions.

SHARING FILES

Liferay DXP's role-based permissions system defines which actions users can take on assets, including files. Administrators can let users collaborate on files by assigning the appropriate file permissions to a Role, and then assigning users to that Role. Similarly, non-administrative users can grant permissions to Roles for files they own.

This Role-based permissions system sometimes falls short. For example, if a Role appropriate for file collaboration doesn't exist, an administrator must create it and manage its users and permissions. Non-administrative users can't create or manage Roles. Also, if a user wants to share a file with one other user, it's not practical for an administrator to create and manage a Role for only two users.

Liferay DXP's sharing feature solves these problems by letting users share files directly with each other, without involving an administrator. This saves time and effort for everyone. After all, sharing is caring.

Note: Administrators can disable sharing. For instructions on this, see [Configuring Sharing](#).

When you share, you grant some of your own permissions for that file to the receiving user. However, there are some important caveats:

- You can only grant View, Comment, or Update permissions. For example, you can't grant Delete or Override Checkout permissions even if you have those permissions on the file.
- You can only grant permissions you have on the file. For example, you can't grant Update permission if you only have View and Comment permissions on the file.
- You must grant at least View permission.
- Traditional Role-based permissions always take precedence over sharing permissions. So although sharing can extend permissions, it can't remove those granted via Roles in the portal.
- By default, the Guest Role has Add Discussion permission. This overlaps with the Comment permission in sharing. Therefore, all users can comment on a file regardless of whether the Comment permission was granted via sharing. Administrators can change this by removing the Add Discussion permission from the Guest Role.

Also note that the receiving user must be part of the same instance, but doesn't have to be a member of the same Site.

165.1 Sharing Files in Documents and Media


To share a file, you must own that file or be an administrator. You must share files via the Documents and Media app in Site Administration or the Documents and Media widget on a page.


Follow these steps to share a file:

1. Using the Documents and Media widget on a page or in Site Administration, navigate to the file you want to share.

To navigate to the Documents and Media app in Site Administration, open the *Menu* () , click your Site's name, and go to *Content & Data* → *Documents and Media*.

To share a file via the Documents and Media widget on a page, actions must be enabled for the widget. Follow these steps to enable actions:

- Select *Configuration* from the widget's *Options* menu ().
- In the *Setup* tab's *Display Settings*, select *Show Actions*.
- Click *Save* and close the Configuration window.

2. Click the file's *Actions* button () and select *Share*. This opens the Share dialog.

Alternatively, click the file in Documents and Media and then click the *Share* button at the top-right. This opens the same Share dialog.

3. Enter the name or email address of the user you want to share the file with. To share the file with multiple users, enter each user's email address in a comma delimited list.
4. To let receiving users also share the file, select *Allow the document to be shared with other users*. Note, however, that administrators can share the file regardless of your selection here.
5. Select the file permissions to grant to receiving users. Because you can only grant your own permissions for the file, some of these options may be unavailable:

- **Update:** View, comment, and update.
- **Comments:** View and comment.
- **View:** View only.

If you enabled further sharing in the previous step, note that receiving users can only share the file with the permissions you grant here.

6. Click *Share*.

165.2 Working with Shared Files

You can access files shared with you in three places:

Share moon-image5.jpg ✕

Invite to Collaborate

You can use a comma to enter multiple collaborators.

Allow the document to be shared with other users.

SHARING PERMISSIONS

Update
Collaborators can view, comment, and update the document.

Comments
Collaborators can comment on the document.

View
Collaborators can only view the document.

Figure 165.1: To share a file, you must fill out the Share dialog as these steps describe.

1. **The Documents and Media Library:** Files shared with you are visible in their existing Documents and Media locations. For example, if someone shares a file with you that resides in the Documents and Media Library's Home folder, then you can access the file in that folder.
2. **The Notifications app:** When a file is shared with you, you get a notification in the Notifications app. Clicking the notification takes you to the file in Documents and Media. For information on notifications, see Managing Notifications and Requests.

NOTIFICATIONS

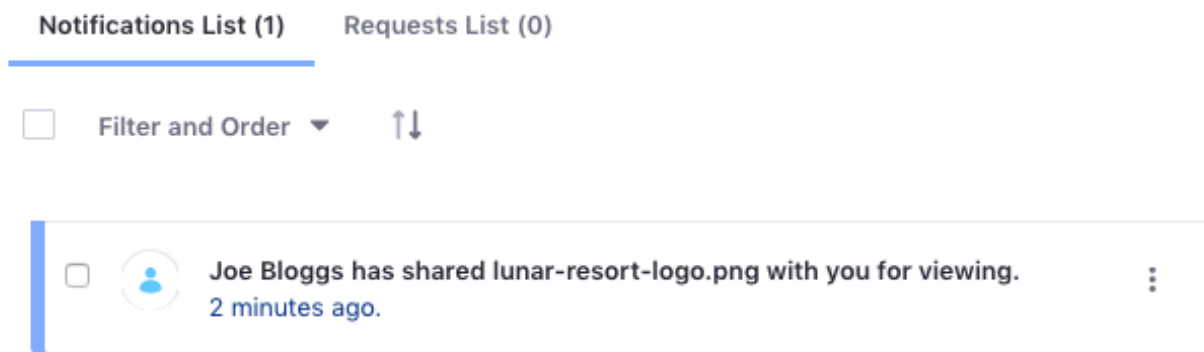


Figure 165.2: The Notifications app contains the notifications that are sent when a user shares a file with you.

3. **The Shared Content app:** This app lists all the content shared with you, and the content you shared. You can access this app from your user menu. Each file has an Actions button (⋮) for performing permitted actions on the file (e.g., view, comment, update).

165.3 Managing Shared Files

After sharing a file, you can unshare it or modify its permissions on a per-user basis. This can only be done by Administrators, the file's owner, or any user with Update permission and permission to share the file. You can take these actions from the file's Info panel in Documents and Media. Follow these steps:

1. Click the file in Documents and Media, then click the *Info* button (ⓘ) at the top-right. The file's Info panel slides out from the right.
2. Click the *Manage Collaborators* link. This shows a list of the users you shared the file with and their file permissions.

SHARED CONTENT

Shared with Me Shared by Me

Filter and Order ▼ ↑↓

| Title | Asset Type | Status | Shared Date | |
|-----------------------|------------|---------|---------------|---|
| lunar-resort-logo.png | Document | VISIBLE | 4 Minutes Ago | ⋮ |

Figure 165.3: The Shared Content app lists the files shared with you, and the files you shared.

3. Make any changes you want to the list of collaborators. To unshare the file with a user, click the x icon next to that user. You can also change the file permissions via the selector menu for each user.
4. Click *Save* and close the dialog.

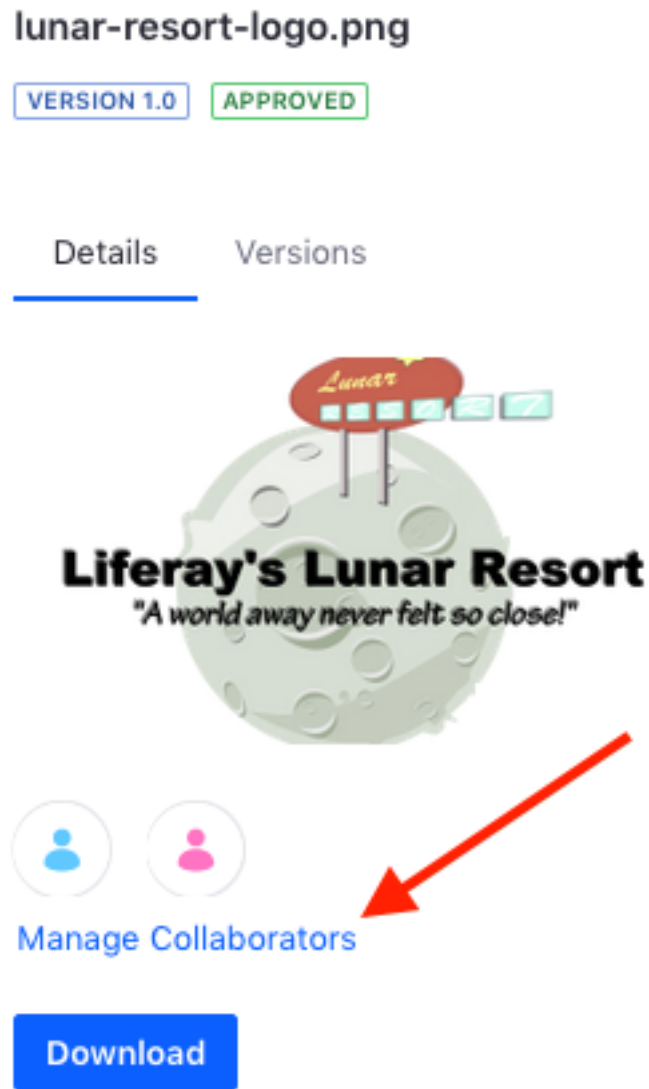


Figure 165.4: Click *Manage Collaborators* to open up the list of users you shared the file with.

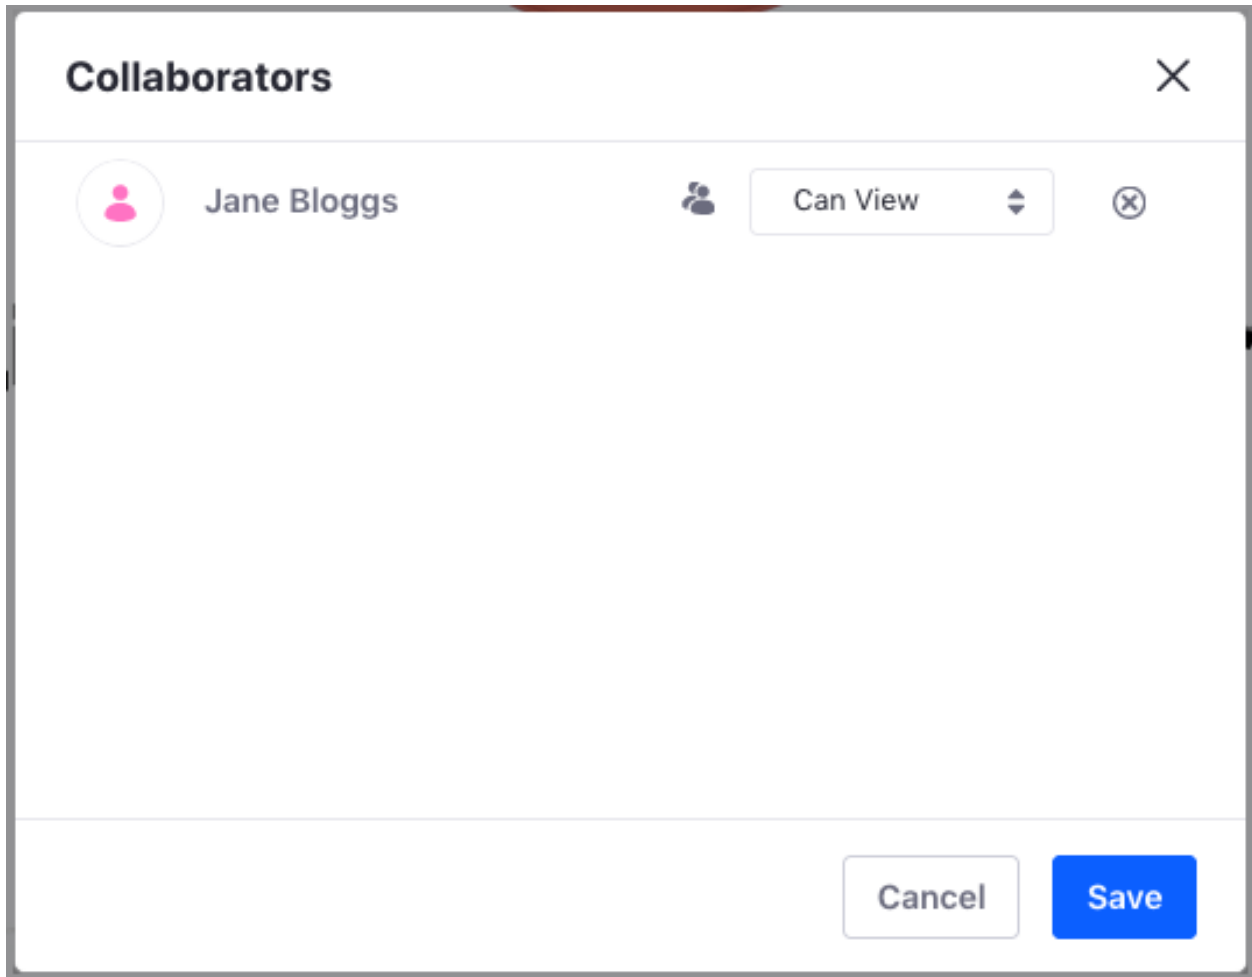


Figure 165.5: The Collaborators dialog lets you unshare a file or change the file permissions for each user.

CONFIGURING SHARING

Administrators can choose whether file sharing is enabled at the global, instance, and Site levels.

166.1 Global Configuration

Sharing is enabled globally by default. To configure sharing globally, follow these steps:

1. Go to *Control Panel* → *Configuration* → *System Settings* → *Sharing*.
2. Under *SYSTEM SCOPE*, select *Sharing*.
3. Configure sharing via these settings:

Expired Sharing Entries Check Interval: The interval in minutes for how often expired sharing entries are checked for deletion.

Enabled: Whether sharing is enabled globally.

When sharing is enabled globally, it's also enabled by default for all portal instances. You can change this from *Virtual Instance Sharing* under *VIRTUAL INSTANCE SCOPE*:

Enabled: Whether sharing is enabled by default for all instances in the portal.

166.2 Instance Configuration

To enable or disable sharing on a per-instance basis, follow these steps:

1. Go to *Control Panel* → *Configuration* → *Instance Settings* → *Sharing*.
2. Select *Virtual Instance Sharing* under *VIRTUAL INSTANCE SCOPE*.

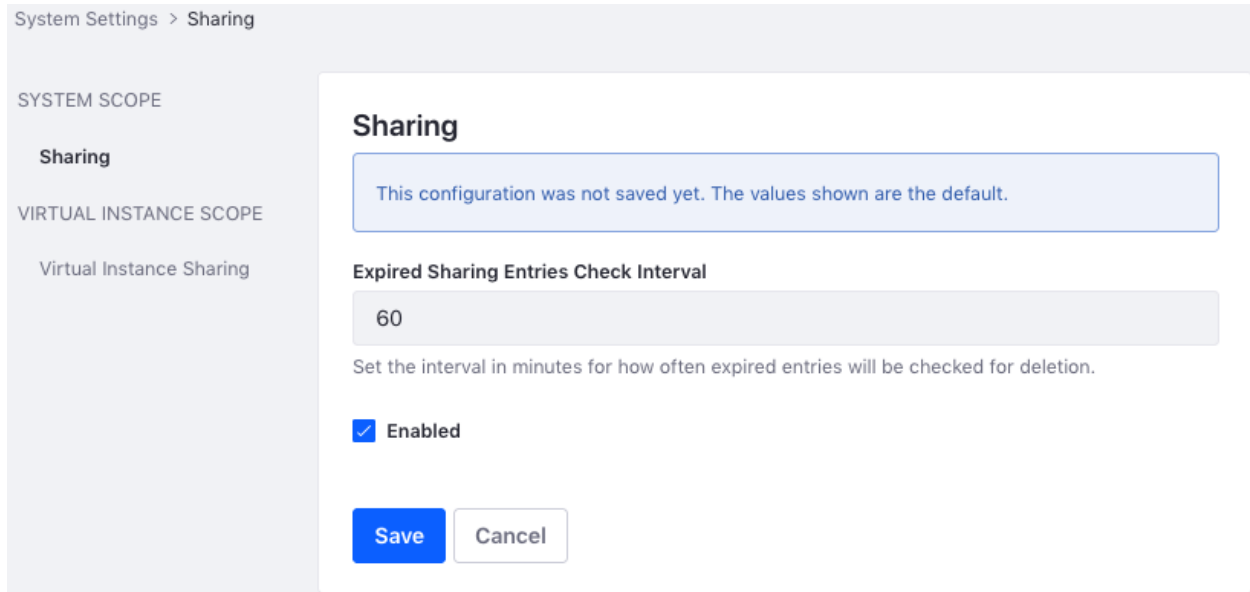


Figure 166.1: Configure sharing globally.

3. Check or uncheck the *Enabled* checkbox to enable or disable sharing, respectively.

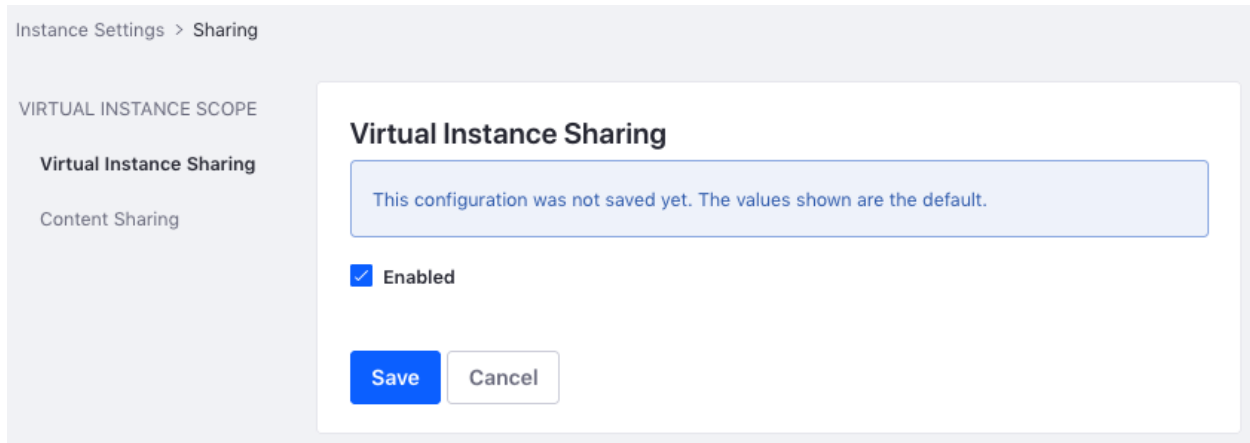


Figure 166.2: You can enable or disable sharing for each instance.

166.3 Site Configuration

To enable or disable sharing for a Site, follow these steps:

1. Go to Site Administration (your site's menu) → *Configuration* → *Settings*.
2. Select the *General* tab.
3. Expand the *Sharing* section and use the toggle to enable or disable sharing for the Site.




Figure 166.3: You can enable or disable sharing for each Site.

DESKTOP ACCESS TO DOCUMENTS AND MEDIA

You can access the Document Library from your desktop file manager via WebDAV. WebDAV is a set of methods based on HTTP that let users create, edit, move, or delete files stored on web servers. WebDAV is supported by most major operating systems and desktop environments, including Linux, macOS, and Windows. Using your file manager via WebDAV doesn't bypass the functionality of the web interface—Liferay DXP increments the version numbers of files edited and uploaded via WebDAV.

To access the Document Library folder from a file browser, you must use your log-in credentials and the WebDAV URL of the folder you want to access. Follow these steps:

1. Navigate to the Documents and Media app that contains the folder you want to access. Click the folder's Actions icon () and select *Access from Desktop*.
2. Copy the WebDAV URL and follow the instructions for your operating system:

Windows: Map a network drive to the WebDAV URL. Enter your credentials when prompted. The Document Library folder appears in the network drive. From your file browser, you can now add, edit, move, or delete files in this folder.

macOS: In the Finder, select *Go* → *Connect to Server*. In the Server Address field, enter the WebDAV URL of the folder you want to access, then click *Connect* and enter your credentials when prompted.

Linux: In your file manager, you must slightly modify the Document Library folder's WebDAV URL. For KDE's Dolphin, change the URL's protocol to `webdav://` instead of `http://`. For GNOME's Nautilus, change the URL's protocol to `dav://` instead of `http://`. Then press *Enter* and enter your credentials when prompted.

Now you can access the Document Library folder from your desktop file system. If you edit a file in this folder on your file system, the change also shows up in the same Document Library folder in the portal. What's more, the file's minor version is incremented due to the edit.

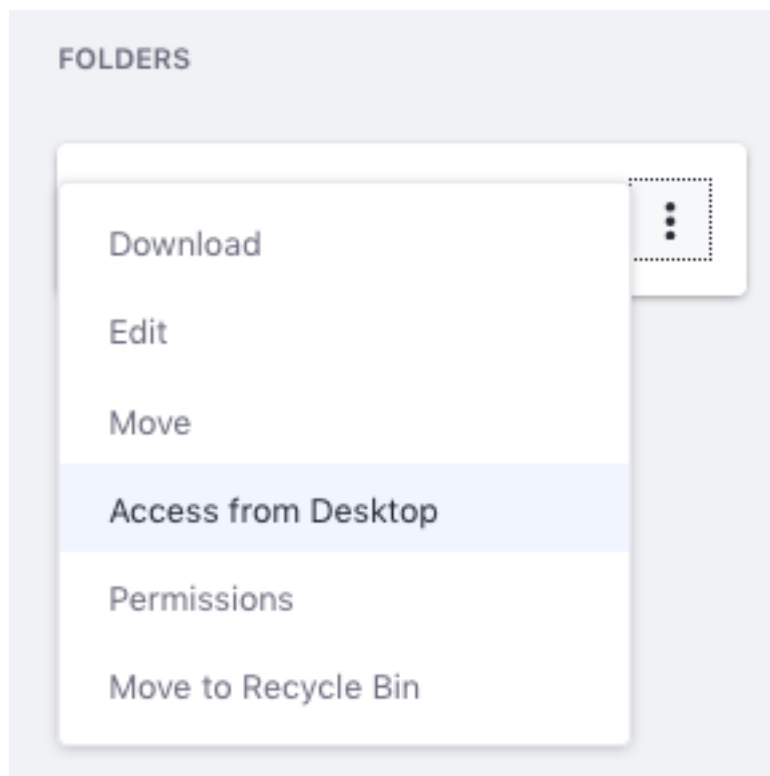


Figure 167.1: Select *Access from Desktop* to get the folder's WebDAV URL.

LINKING TO GOOGLE DRIVE™

You can create Document Library files that link to files in Google Drive™ and Google Photos™. This lets you access your Google files from the Document Library. Note that this functionality isn't available by default. To enable it, you must complete these steps:

1. Install the Liferay Plugin for Google Drive™ from Liferay Marketplace.
2. Create and/or configure a Google project capable of communicating with your Liferay DXP instance. The Google Picker API must be enabled for this project. This API lets you select Google files to link to. You must also create the credentials the Google project needs to communicate with your Liferay DXP instance.
3. Configure your portal to communicate with your Google project.

This article shows you how to complete these steps and finishes with an example of linking to a Google file from the Document Library.

Note: You can also use Google Docs™ for online file creation and editing. This doesn't require a plugin and is covered in a separate section of the documentation.

Important: The Liferay Plugin for Google Drive™ is a Labs application available for Liferay CE Portal and Liferay DXP. Labs apps are experimental and not supported by Liferay. They're released to accelerate the availability of useful and cutting-edge features. This status may change without notice. Use Labs apps at your own discretion.

168.1 Install the App

First, you must install the the Liferay Plugin for Google Drive™ from Liferay Marketplace. This app is available via the following links for Liferay CE Portal and Liferay DXP:

- Liferay Plugin for Google Drive - CE
- Liferay Plugin for Google Drive - DXP

If you need help installing apps from Marketplace, see the documentation on using Marketplace.

168.2 Configure Your Google Project

Follow these steps to create and/or configure your Google project so it can communicate with your Liferay DXP instance:

1. Go to the Google API Console. If you don't have a suitable project, create a new one.
2. Enable the Google Picker API for your project. For instructions, see the Google API Console documentation on enabling and disabling APIs.
3. Create an OAuth 2 client ID in your Google project. For instructions, see the Google API Console documentation on setting up OAuth 2.0. Enter these values when creating your client ID:
 - **Application type:** Web application
 - **Name:** Google Docs Hook
 - **Authorized JavaScript origins:** [liferay-instance-URL] (e.g., `http://localhost:8080` is the default for local development machines)
 - **Authorized redirect URIs:** [liferay-instance-URL]/oath2callback
4. Create a new API key in your Google project. For instructions, see the Google API Console documentation on creating API keys. Be sure to restrict the key to HTTP referrers (web sites), and set it to accept requests from your Liferay DXP instance's URL.

Your new OAuth client ID and public API access key now appear on your Google project's Credentials screen. Keep this screen open to reference these values as you specify them in Liferay DXP.

168.3 Configure Your Portal

Now that you have a Google project set up for use with Liferay DXP, you must connect your installation to that project. You can do this at two scopes:

1. Globally, for all instances in your Liferay DXP installation.
2. At the instance scope, for one or more instances in your Liferay DXP installation.

You can override the global configuration for one or more instances by configuring those instances separately. Similarly, you can configure only the instances you want to connect to your Google project and leave the global configuration empty.

Follow these steps to configure your Liferay DXP installation to connect to your Google project:

1. Note that the configuration options are the same in the global and instance-level configurations.


To access the global configuration, go to *Control Panel* → *Configuration* → *System Settings* → *Documents and Media*.


To access the instance-level configuration, go to *Control Panel* → *Configuration* → *Instance Settings* → *Documents and Media*.

2. Under *VIRTUAL INSTANCE SCOPE*, select *Google Drive*.
3. Enter your Google project's OAuth 2 client ID and client secret into the *Client ID* and *Client Secret* fields.
4. In the field *Picker API Key*, enter the API key you created in the previous section.
5. Click *Save*.

168.4 Creating Linked Files

With the preceding configuration steps complete, you can create files in your Document Library that link to files in Google Drive™ or images in Google Photos™. Follow these steps to do so:

1. In your Document Library, click the *Add* button () and select *Google Drive Shortcut*. The *New Google Drive Shortcut* screen appears.
2. Click the *Select File* button to open Google's file picker.
3. Use the file picker to select a file from Google Drive™ or Google Photos™.
4. Click *Publish*.

A new file entry appears for the Google document you linked to. You can view the file entry as you would any file entry. The Google document's contents show in the file entry's preview pane. As with any file entry, the *Options* button () gives you access to the Download, Edit, Move, Permissions, Move to Recycle Bin, and Checkin/Checkout/Cancel Checkout options.

System Settings > Documents and Media

SYSTEM SCOPE

- TensorFlow Image Auto Tagging Process
- Service
- File Entries
- File Rank
- Item Selector View

VIRTUAL INSTANCE SCOPE

- Google Cloud Natural Language Text Auto Tagging
- Google Cloud Vision Image Auto Tagging
- Microsoft Cognitive Services Image Auto Tagging
- OpenNLP Text Auto Tagging
- TensorFlow Image Auto Tagging
- Google Drive**

Google Drive

This configuration was not saved yet. The values shown are the default.

All fields marked with * are required.

Client ID *

Sets the OAuth2 client id for the Google Drive API. Go to <https://console.developers.google.com/apis/credentials/oauthclient> to create it.

Client Secret *

Sets the OAuth2 client secret for the Google Drive API. Go to <https://console.developers.google.com/apis/credentials/oauthclient> to create it.

Picker API Key

Sets the API Key for the Google Picker API. Go to https://console.developers.google.com/flows/enableapi?apiid=picker&credential=client_key to create it.

Save **Cancel**

Figure 168.1: Enter your Google project's OAuth 2 client ID, OAuth 2 client secret, and Picker API key.

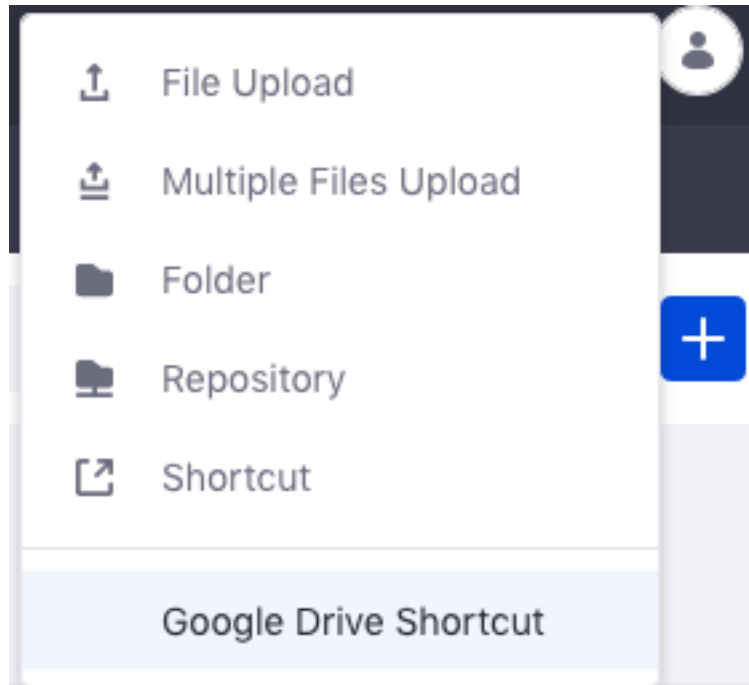


Figure 168.2: Select *New Google Drive Shortcut* from the *Add* menu in your Document Library.

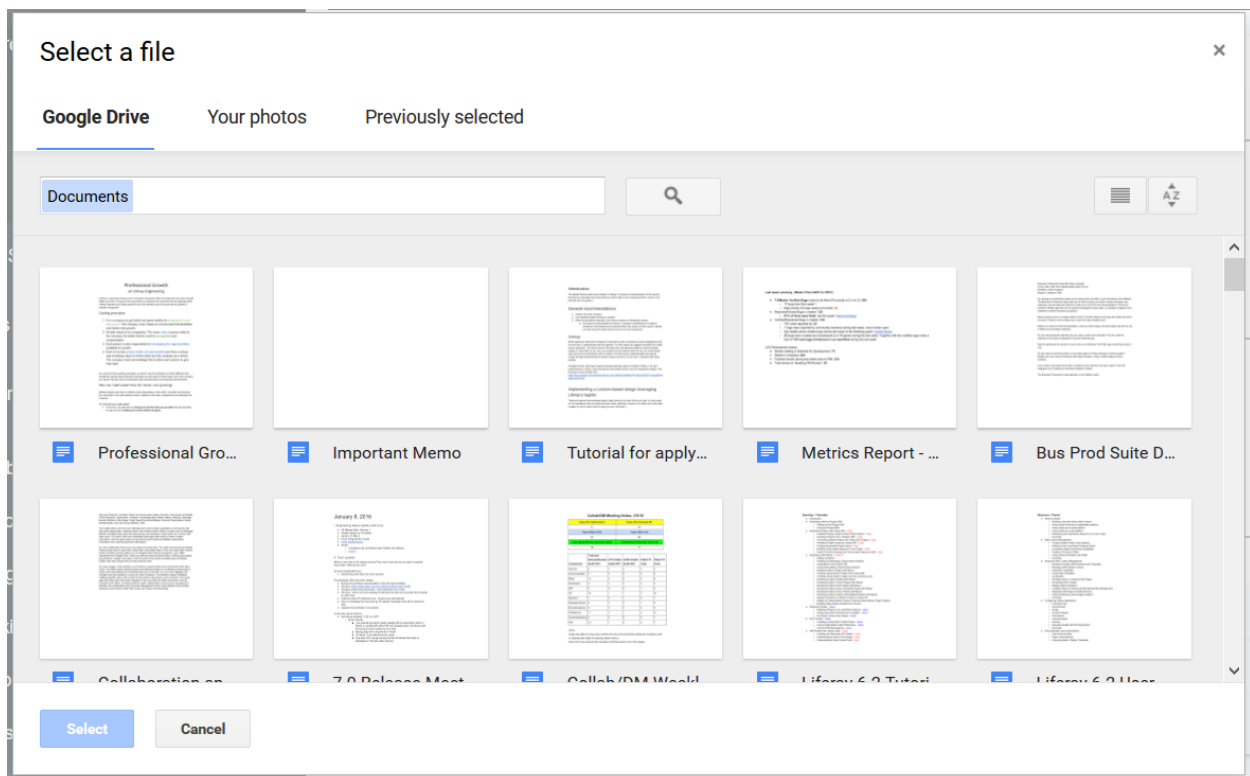


Figure 168.3: You can select files from Google Drive™ or your photos.

METADATA SETS


You can define metadata fields that users fill out when they create or edit Document Library files. You do this by creating *metadata sets* and then associating them with document types, which wrap Document Library files and thus apply your metadata fields to the files. Although you apply metadata sets via document types, metadata sets exist independently and you can apply them to any number of document types.

169.1 Managing Metadata Sets

To see the available metadata sets, open the *Menu* () , expand your site's menu, and navigate to *Content & Data* → *Documents and Media*. Then click the *Metadata Sets* tab. Any existing sets appear in a table.

To select a metadata set, select the checkbox to its left. To select all the sets, select the checkbox in the Management Bar. With one or more sets selected, an X icon appears in the Management Bar. Clicking it deletes the selected metadata set(s). Note that metadata sets don't support the Recycle Bin. If you delete a metadata set, it's gone forever.

The Management Bar also contains other options for managing the metadata sets. The selector menu to the right of the checkbox filters the sets the table displays (it's set to *All* by default). The *Order by* selector orders the sets by Modified Date or ID. The up and down arrows sort the sets in ascending or descending order, respectively. You can also use the Search bar to search for a set.

In the table, each metadata set has an Actions button () for performing the following actions on that set:

Edit: Edit the set. Alternatively, click the set's name in the table.

Permissions: Configure the set's permissions.

Copy: Copy the metadata set.



Delete: Delete the set.

| ID | Name | Description | Scope | Modified Date |
|--------------------------------|-------------------------|----------------------------------|--------------|---------------|
| <input type="checkbox"/> 34208 | Meeting Metadata | Metadata set for meetings | Current Site | 4 Minutes Ago |
| <input type="checkbox"/> 34216 | Travel Request Metadata | Metadata set for travel requests | Current Site | 0 Seconds Ago |


Figure 169.1: The Metadata Sets management window lets you view existing sets and create new ones for applying to document types.

169.2 Creating Metadata Sets

Follow these steps to create a metadata set:

1. From the *Menu* () , click your Site's name and navigate to *Content & Data* → *Documents and Media*. Then click the *Metadata Sets* tab.
2. Click the *Add* button () . The New Metadata Set form appears.
3. Give your metadata set a name.
4. Open the *Details* section of the form to give your metadata set a description or select a metadata set to extend (both are optional). To select a metadata set to extend, click the *Select* button for *Parent Metadata Set* and then select the metadata set. When a user creates a document of a document type that uses an extended metadata set, the parent metadata set's fields appear above the extended metadata set's.
5. Add the metadata fields that should be part of this metadata set. To do this, first select the editor's *View* tab and select the *Fields* tab within it. Icons representing the field types are listed on one side and the metadata set's canvas is on the other side. To add a field type to the metadata set, select its icon, drag, and drop it onto the canvas. The field appears on the canvas as it does for users. By dragging a field onto a field that's already on the canvas, you can nest the new field in the existing field. When you mouse over a field on the canvas, the



field action icons () appear. Clicking the + icon creates a duplicate of the current field and adds it below the current field. Clicking the trash can deletes the field.

The following metadata fields are available:

- **Boolean:** A check box.

- **Color:** Specifies a color.
- **Date:** Enter a date. A valid date format is required for the date field, but you don't have to enter a date manually. When you select the date field a mini-calendar pops up which you can use to select a date.
- **Decimal:** Enter a decimal number. The value is persisted as a double.
- **Documents and Media:** Select a file from a Documents and Media library.
- **Geolocation:** Specify a location to associate with the document.
- **HTML:** An area that uses a WYSIWYG editor to enhance the content.
- **Integer:** Enter an integer. The value is persisted as an int.
- **Link to Page:** Link to another page in the same site.
- **Number:** Enter a decimal number or an integer. The value is persisted either as a double or an int, depending on the input's type.
- **Radio:** Displays several clickable options. The default number of options is three but this is customizable. Only one option can be selected at a time.
- **Select:** This is just like the radio field except that the options are hidden and must be accessed from a drop-down menu.
- **Text:** Enter a single line of text.
- **Text Box:** This is just like the text field except you can enter multiple lines of text or separate paragraphs.
- **Web Content:** Select web content.

6. Edit your fields to reflect their intended metadata. For example, a text field's default label is *Text*. If you want to use the text field as a title, for instance, then you should change the field's label to *Title*. To do this, first select the field on the canvas. This automatically selects the *Settings* tab on the left. Alternatively, you can access the Settings tab by clicking the field's wrench icon. To edit a setting value, double-click it in the Settings table and enter the new value.

Labels, default values, variable names, mouse-over tips, widths, and other settings can be configured for most fields. Some fields have a *Required* setting for specifying whether users must populate the field. If a field's *Repeatable* setting is *Yes*, users can add multiple consecutive instances of the field to the document's metadata.

Also note that you can translate each of a metadata set's field values to any supported locales. To specify a field value for a translation, select the flag that represents the locale and enter the field value for the locale.

7. Click *Save* when you're done specifying your new metadata set.

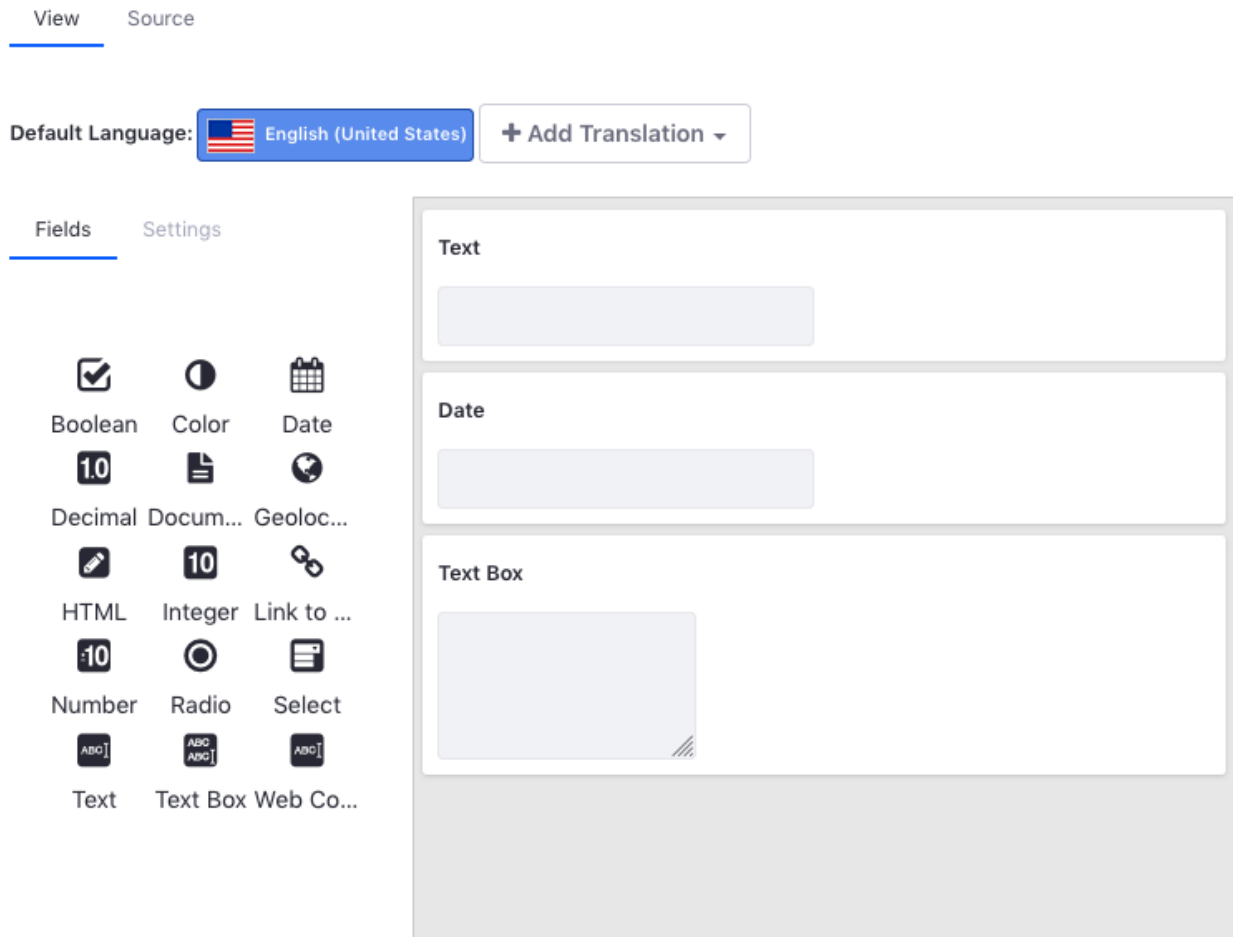






Figure 169.2: Add your metadata set's fields to the canvas.

View Source

Default Language:  English (United States) + Add Translation ▾

Fields Settings

| name | value |
|------------------|---------------------|
| Type | text |
| Field Label | Text |
| Show Label | Yes |
| Required | No |
| Name | Textdunh |
| Predefined Value | |
| Tip | |
| Indexable | Indexable - Keyword |
| Localizable | Yes |
| Repeatable | No |

Text   

Date



Text Box

Figure 169.3: Edit your metadata set's fields to match the metadata that you want each field to hold.

DOCUMENT TYPES

Document types are made of metadata fields and help users define the purpose of Document Library files. For example, a *Contract* document type may need metadata fields for the effective date, expiration date, contract type, legal reviewer, and more. When users create Document Library files of the Contract document type, they can then populate those metadata fields. Document types also help you integrate files with other features like search and workflow. Search works on file metadata so users can find files faster. You can also apply workflows to specific document types. And you can more cleanly organize document libraries by designating folders to hold particular document types exclusively.

170.1 Managing Document Types

To see the available document types, open the *Menu* () , expand your site's menu, and navigate to *Content & Data* → *Documents and Media*. Then click the *Document Types* tab. A searchable table lists any existing document types. The following actions are available for each document type via its *Actions* button ():

Edit: Edit the document type.

Permissions: Set the document type's permissions.

Delete: Delete the document type. Note that document types don't support the Recycle Bin. Once you delete a document type, it's gone forever.

170.2 Creating Document Types

Follow these steps to create a document type:






Figure 170.1: The Document Types management window lets you view existing document types and create new ones.


1. From the *Menu* () , expand your site's menu and navigate to *Content & Data* → *Documents and Media*. Then click the *Document Types* tab.
2. Click the *Add* button () . The *New Document Type* form appears.
3. Give your document type a name and a description.
4. Define the metadata to use with your document type. You do this via these sections in the form:

Main Metadata Fields: These are tied directly to the document type. They can be created only via the form and can't be used with other document types. You create and edit these metadata fields in the form the same way that you do when creating metadata sets.


Additional Metadata Fields: Select a metadata set to associate with the document type. Each document type must be associated with one or more metadata set. To differentiate document types that use the same metadata sets, define different main metadata fields.
5. Define your document type's permissions via the form's *Permissions* section. By default, anyone can view the document type, including site guests. You can restrict its view, update, delete, and permissions configuration to site members or the document type's owner.
6. Click *Save* when you're finished specifying your new document type.

Your document type is now available when adding a document via the Documents and Media's *Add* menu. When users create new files of the document type, they're presented with metadata fields to describe the document.

Name *

Video Team Meeting 


Description

Document type for the video team's meeting minutes. 

Main Metadata Fields >

Additional Metadata Fields ▾

Name

Meeting Metadata 

Select

Permissions >

Figure 170.2: Create your new document type.

ONLINE FILE CREATION AND EDITING WITH GOOGLE DOCS™

Although you can add and edit Documents and Media files via upload and download, Liferay DXP doesn't contain a UI that lets you author or edit documents directly. You can, however, create and edit Documents and Media files online in Google Docs™, Google Sheets™, and Google Slides™.

Note: For simplicity and readability, this documentation refers only to Google Docs™. The material, however, also applies to Google Sheets™ and Google Slides™.

Note that when you use Google Docs™ to create or edit a Documents and Media file, that file isn't permanently stored in Google Docs™. Google Docs™ is only used for its editing UI. Your edits are then saved to the Documents and Media Library.

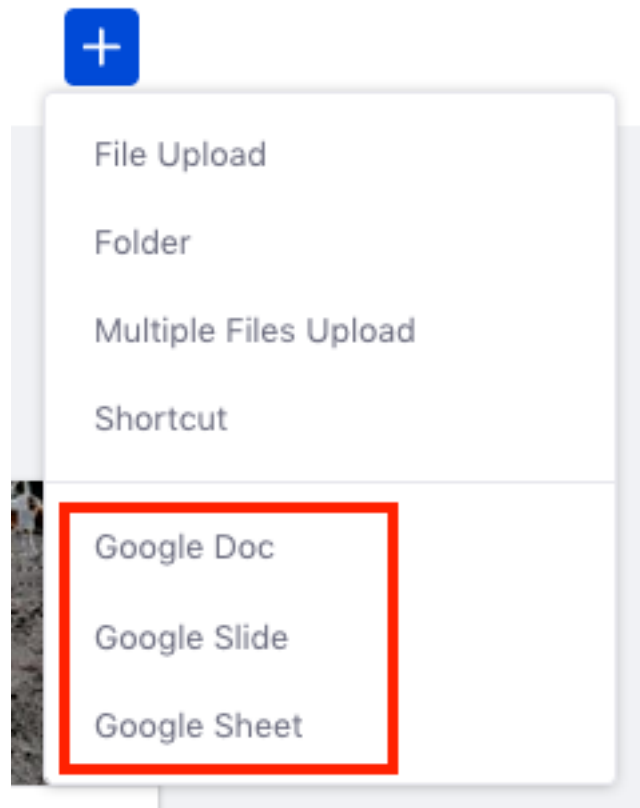


Figure 171.1: You can create new Google documents in Documents and Media.

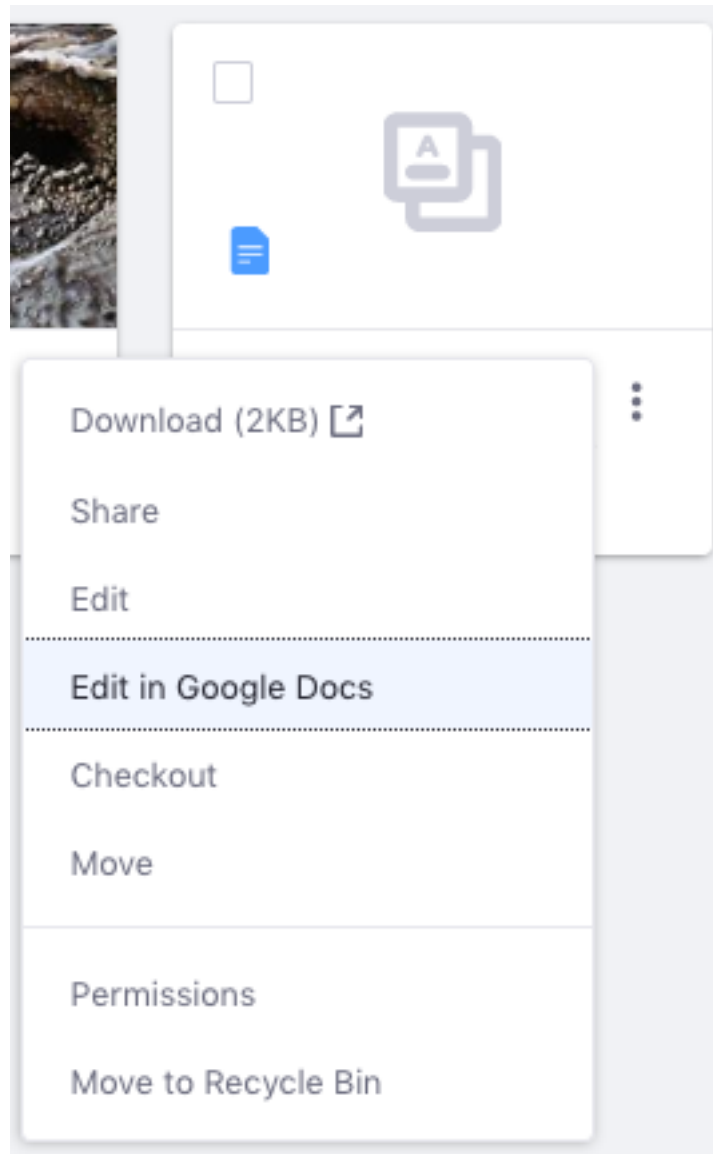


Figure 171.2: You can also use Google’s document editor to edit existing Documents and Media files.



Figure 171.3: When using Google’s document editor, you can save or discard your changes via the editor’s toolbar.

CONFIGURING GOOGLE DOCS™ INTEGRATION

Before you can use Google Docs™ to create and edit Documents and Media files, you must configure Liferay DXP to connect with an application in the Google API Console.

Note: You must be an administrator to complete these steps.

172.1 Configure Your Google Project

First, you must configure your Google project to use the Google Drive™ API and set up OAuth 2 for use with that project.

1. Go to the Google API Console. If you don't have a suitable project, create a new one.
2. Enable the Google Drive™ API for your project. For instructions, see the Google API Console documentation on enabling and disabling APIs.
3. Create an OAuth 2 client ID for your Google project. For instructions, see the Google API Console documentation on setting up OAuth 2.0. Select *Web application* when prompted to select your application type. Take note of the client ID and client secret that appear—you'll need them to configure the portal to use the Google Drive™ API.

172.2 Configuring the Portal

Now that you have a Google project set up for use with Liferay DXP, you must connect your installation to that project. You can do this at two scopes:

1. Globally, for all instances in your Liferay DXP installation.
2. At the instance scope, for one or more instances in your Liferay DXP installation.

You can override the global configuration for one or more instances by configuring those instances separately. Similarly, you can configure only the instances you want to connect to your Google project and leave the global configuration empty.

Follow these steps to configure your Liferay DXP installation to connect to your Google project:

1. Note that the configuration options are the same in the global and instance-level configurations.

To access the global configuration, go to *Control Panel* → *Configuration* → *System Settings* → *Documents and Media*.

To access the instance-level configuration, go to *Control Panel* → *Configuration* → *Instance Settings* → *Documents and Media*.

2. Under *VIRTUAL INSTANCE SCOPE*, select *Google Drive*.
3. Enter your Google project's OAuth 2 client ID and client secret into the *Client ID* and *Client Secret* fields.
4. Click *Save*.

Note: To turn this feature off, delete the client ID and client secret values from the form.

Note: You can ignore the *Picker API Key* field. This field is unrelated to the Google Docs™ online editing features in Liferay DXP.

System Settings > Documents and Media

SYSTEM SCOPE

- TensorFlow Image Auto Tagging Process
- Service
- File Entries
- File Rank
- Item Selector View

VIRTUAL INSTANCE SCOPE

- Google Cloud Natural Language Text Auto Tagging
- Google Cloud Vision Image Auto Tagging
- Microsoft Cognitive Services Image Auto Tagging
- OpenNLP Text Auto Tagging
- TensorFlow Image Auto Tagging
- Google Drive**

Google Drive

This configuration was not saved yet. The values shown are the default.

All fields marked with * are required.

Client ID *

Sets the OAuth2 client id for the Google Drive API. Go to <https://console.developers.google.com/apis/credentials/oauthclient> to create it.

Client Secret *

Sets the OAuth2 client secret for the Google Drive API. Go to <https://console.developers.google.com/apis/credentials/oauthclient> to create it.

Picker API Key

Sets-the API Key for the Google Picker API. Go to https://console.developers.google.com/flows/enableapi?apiid=picker&credential=client_key to create it.

Save **Cancel**

Figure 172.1: Enter your Google project's OAuth 2 client ID and client secret.

CREATING AND EDITING FILES WITH GOOGLE DOCS™

You can use Google Docs™ to create and edit text documents, spreadsheets, or presentations for storage in the Docs and Media library. When you finish your Google Docs™ editing session, your changes are automatically saved in the Documents and Media Library. You can


- Create Files
- Edit Files
- Manage Multiple Editing Sessions

173.1 Authentication

The first time you create or edit a Documents and Media file via Google Docs™, you must authenticate with your Google account. This links Google Drive™ to your portal account, so you only need to do this once. You can unlink your account at any time by navigating to User Menu → *Account Settings* → *General* → *Apps*, and clicking *Revoke* next to Google Drive™.

173.2 Creating Files

Follow these steps to create a new Documents and Media file via Google Docs™:

1. Open the *Menu* () , click your site's name, and navigate to *Content & Data* → *Documents and Media*.

ACCOUNT SETTINGS

General Contact Preferences

Information

Organizations

Memberships

Roles

Password

Apps

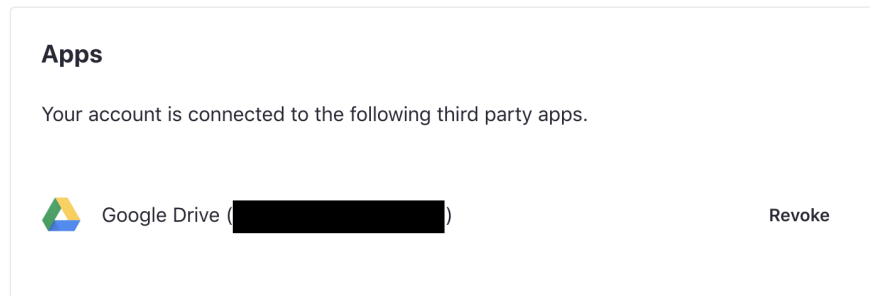


Figure 173.1: You can unlink your Google account from the portal.

2. Click the *Add* icon (+) and select the type of Google document to add to the Document Library.

- Google Docs™
- Google Slides™
- Google Sheets™

When you select one of these options, Liferay DXP creates a temporary Documents and Media file and links it to a new Google file. Your browser then redirects you to that Google file so you can create its content.

Note that some actions on the temporary Documents and Media file can affect its linked Google file. For more information, see [multiple editing sessions](#).

3. Use the Google Docs™ editor to create your document's content. All Google Docs™ features are available except for sharing.
4. Save or discard your changes by clicking one of these toolbar buttons in the Google Docs™ editor:

Save and Return to Liferay: Saves your document as a new file in the Documents and Media Library, deletes the Google file, and returns you to the portal. The saved file's format depends on the type of Google document you selected in step two above:

- Google Docs™: Microsoft Word (.docx)
- Google Slides™: Microsoft PowerPoint (.pptx)
- Google Sheets™: Microsoft Excel (.xlsx)

Discard Changes: Returns you to the portal without saving your file in the Documents and Media Library or Google Docs™.

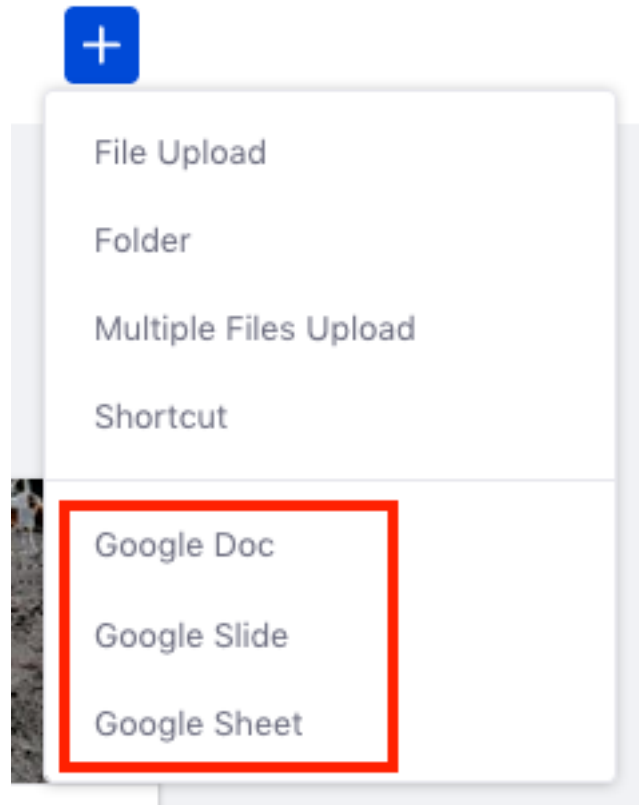


Figure 173.2: Select the type of Google document you want to create.

Note that it's also possible to close the Google Docs™ window without clicking either button. In this case, the editing session remains open even though the window that displayed it is closed. For more information, see the section below on multiple editing sessions.



Figure 173.3: Save or discard your changes by using the toolbar in the editor.

173.3 Editing Files


You can use Google Docs™ to edit the following types of Documents and Media files:

- Text files (.docx, .html, .txt, .rtf, .odt)
- Presentation files (.pptx, .odp)

- Spreadsheet files (.xlsx, .ods, .csv, .tsv)
- PDF files


Note: Google Docs™ doesn't support older, non-XML-based Microsoft Office file types (.doc, .ppt, .xls).

Follow these steps to edit a Documents and Media file in Google Docs™:

1. Navigate to the file in the Documents and Media Library.
2. Click the file's Actions icon () and select *Edit in Google Docs*. This automatically checks out the file, transfers its content to a new Google Docs™ file, and redirects you to that Google Docs™ file.
3. Edit the file in Google Docs™. The editing process is exactly the same as described above for creating files.

173.4 Multiple Editing Sessions

When you create or edit a Documents and Media file in Google Docs™, you can save or discard your changes by clicking *Save and Return to Liferay* or *Discard Changes*, respectively. If you instead close the window without clicking either, the editing session still exists. You can access it via the original file in Documents and Media. If the file didn't exist before (e.g., you were creating a new file), it appears in Documents and Media as a temporary file.

When an editing session already exists for a Documents and Media file, the following actions are available via that file's Actions icon ():

Edit in Google Docs: Resume editing the file in Google Docs™.

Check in: Saves the Google file (including any changes) to the Documents and Media file, then deletes the Google file. This is equivalent to clicking *Save and Return to Liferay* in a Google Docs™ editing window.

Cancel Checkout: Deletes the Google file, discarding any changes. This is equivalent to clicking *Discard Changes* in a Google Docs™ editing window.

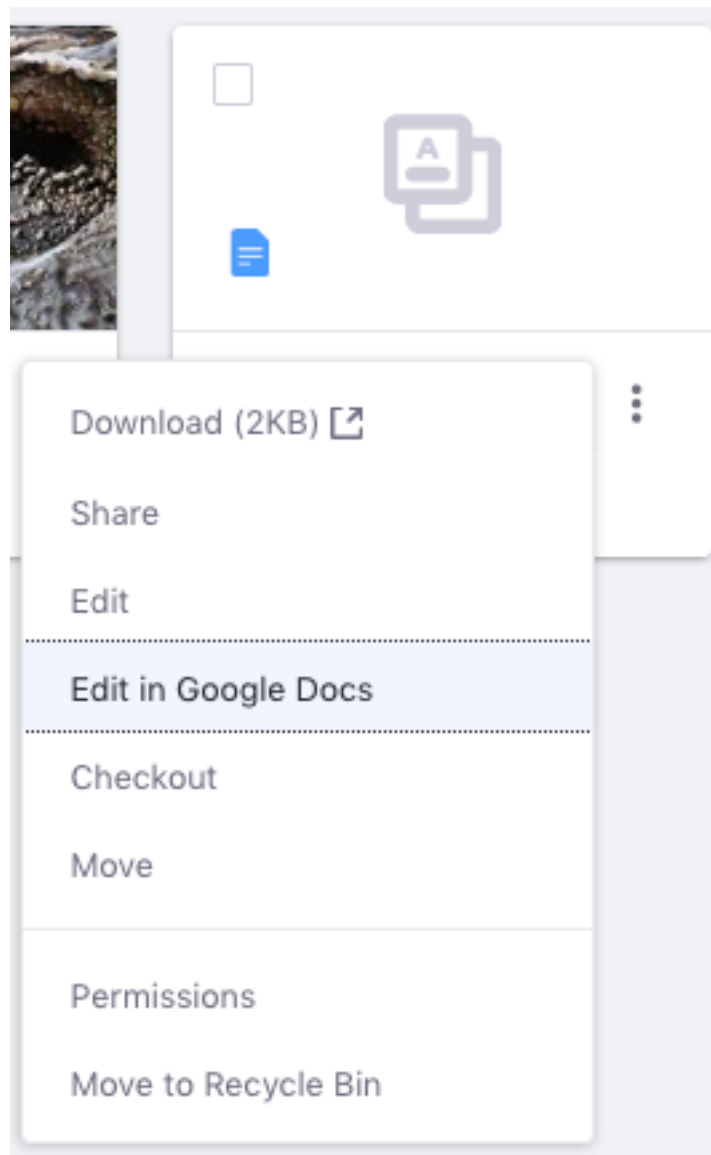


Figure 173.4: Select *Edit in Google Docs* from the file's Actions menu.

INTEGRATION WITH MICROSOFT OFFICE 365™

Although you can add and edit Documents and Media files via upload and download, Liferay DXP doesn't contain a UI that lets you author or edit documents directly. You can, however, create and edit Documents and Media files online in Word™, Excel™, and PowerPoint™.

Note that when you use Office 365™ to create or edit a Documents and Media file, that file isn't permanently stored in Office 365™. Office 365™ is only used for its editing UI. Your edits are then saved to the Documents and Media Library.

Here you can learn how to use it.

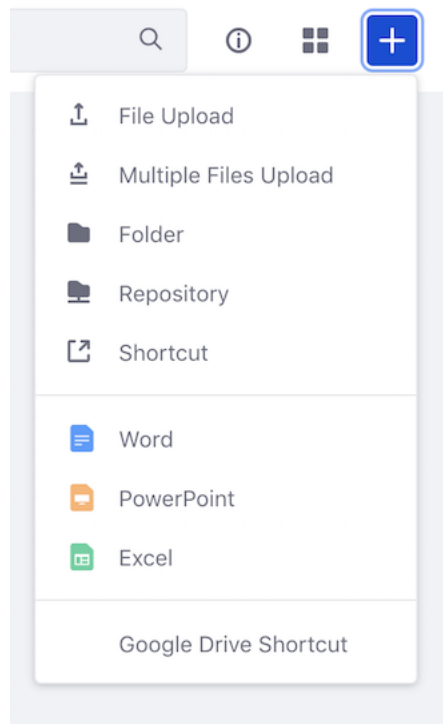


Figure 174.1: You can create new Office 365™ documents in Documents and Media.




Figure 174.2: You can also edit existing Documents and Media files in Office 365™.

CONFIGURING OFFICE 365™ INTEGRATION

Before you can use Office 365™ to create and edit Documents and Media files, you must configure Liferay DXP to connect with an application in the Azure portal.

Note: You must be an administrator to complete these steps.

175.1 Register an Application with the Microsoft Identity Platform

First, configure your application with the Microsoft identity platform™. To do so, follow the steps described in Microsoft's documentation.

To construct a URL for the *Redirect URI* parameter, follow this pattern:

```
https://[hostname]/o/document_library/onedrive/oauth2
```

Here's the minimum permission set needed to use Office 365™ integration:

- Files.Read.All
- Files.ReadWrite.All

For more information about permissions, see Microsoft's documentation.

175.2 Configuring Liferay DXP

Now you must connect your Liferay DXP installation with your Microsoft identity platform™ application. You can do this at two scopes:

1. Globally, for all instances in your Liferay DXP installation.
2. At the instance scope, for one or more instances in your Liferay DXP installation.

You can override the global configuration for one or more instances by configuring those instances separately. Similarly, you can configure only the instances you want to connect to your application and leave the global configuration empty.

Follow these steps to configure your Liferay DXP installation to connect to your application:

1. Note that the configuration options are the same in the global and instance-level configurations.

To access the global configuration, go to *Control Panel* → *Configuration* → *System Settings* → *Documents and Media*.

To access the instance-level configuration, go to *Control Panel* → *Configuration* → *Instance Settings* → *Documents and Media*.

2. Under *VIRTUAL INSTANCE SCOPE*, select *OneDrive*.
3. Enter your application's OAuth 2 client ID and client secret in the *Client ID* and *Client Secret* fields, respectively.
4. Enter your tenant ID in the *Tenant* field. To find your tenant ID, see Microsoft's documentation.
5. Click *Save*.

Note: Once enabled, you can disable this feature by deleting the client ID, client secret, and tenant values from the form.

SYSTEM SCOPE

- Service
- File Entries
- File Rank
- Item Selector View
- Recent Documents
- Sharepoint OAuth2
- Sharepoint Search

VIRTUAL INSTANCE SCOPE

- Google Drive
- OneDrive**

OneDrive

Client ID *

Set the OAuth 2.0 client id for the OneDrive API. Go to <https://docs.microsoft.com/graph/auth-register-app-v2> to create it.

Client Secret *

Set the OAuth 2.0 client secret for the OneDrive API. Go to <https://docs.microsoft.com/graph/auth-register-app-v2> to create it.

Tenant *

Set the tenant for the OneDrive API. Go to <https://docs.microsoft.com/onedrive/find-your-office-365-tenant-id> to obtain it.

Update **Cancel**

Figure 175.1: Enter your application’s client ID, client secret, and tenant.

CREATING AND EDITING DOCUMENTS AND MEDIA FILES WITH OFFICE 365™

You can use Office 365™ to create and edit text documents, spreadsheets, or presentations for storage in the Documents and Media library. When you finish your Office 365™ editing session, you must check in the document to save the changes in the Documents and Media Library.

Here, you'll learn how to do these things:


- Authenticate with OneDrive™
- Create Files
- Edit Files

176.1 Authentication

The first time you create or edit a Documents and Media file via Office 365™, you must authenticate with your Microsoft account. This links OneDrive™ to your Liferay DXP account, so you only need to do this once. You can unlink your account at any time by navigating to User Menu → *Account Settings* → *General* → *Apps*, and clicking *Revoke* next to OneDrive™.

176.2 Creating Files

Follow these steps to create a new Documents and Media file via Office 365™:

1. Open the *Menu* () , click your Site's name, and navigate to *Content & Data* → *Documents and Media*.




Figure 176.1: You can unlink your account from the portal.

2. Click the *Add* icon () and select the type of Office 365™ document to add to the Document Library:

- Word™
- Excel™
- PowerPoint™

When you select one of these options, a new window opens for you to enter the document's name.

3. Enter the document's name in the *Title* field, and click *Save*. When you click *Save*, Liferay DXP creates a temporary Documents and Media file and links it to the new Office 365™ file. Your browser then opens a new window with that Office 365™ file so you can create its content.
4. Use the Office 365™ editor to create your document's content.
5. Save or discard your changes to Liferay DXP:

Check in: Saves the Office 365™ file to Documents and Media, then deletes the file from Office 365™. The saved file's format depends on the document type you selected in step two above.

- Word: Microsoft Word™ (.docx)
- PowerPoint: Microsoft PowerPoint™ (.pptx)
- Excel: Microsoft Excel™ (.xlsx)

Cancel Checkout: Deletes the Office 365™ file, discarding any changes.

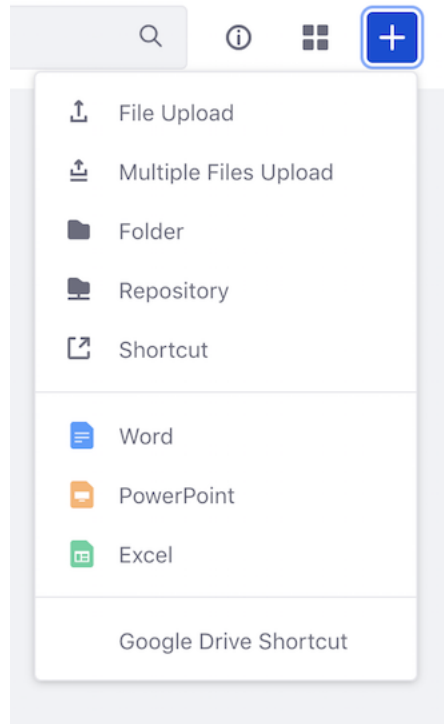


Figure 176.2: Select the type of document you want to create.





Figure 176.3: Give the document a name.

176.3 Editing Files

You can use Office 365™ to edit the following types of Documents and Media files:

- Text files (.doc, .docx, .docm, .dot, .dotx, .dotm, .html, .txt, .rtf, .odt)
- Presentation files (.ppt, .pptx, .pptm, .pps, .ppsx, .ppsm, .pot, .potx, .potm)
- Spreadsheet files (.xls, .xlsx, .xlsm, .xlt, .xltx, .xltm, .ods, .csv, .tsv, .txt, .tab)

Follow these steps to edit a Documents and Media file in Office 365™:

1. Navigate to the file in the Documents and Media Library.
2. Click the file's Actions icon () and select *Edit in Office 365*. This automatically checks out the file, transfers its content to a new Office 365™ file, and redirects you to that file Office 365™.
3. Edit the file in Office 365™. The editing process is exactly the same as described above for creating files.




Figure 176.4: Select *Edit in Office 365* from the file's Actions menu.

STORE TYPES

You can change the file system (the *store*) that the Documents and Media library uses to store files. This is configured in the `portal-ext.properties` file by setting the `dl.store.impl` property. Configuring stores is covered in the Document Repository Configuration guide. Here, you'll consider the ramifications of different stores:

Simple File System Store: Uses the file system (local or a mounted share) to store files. This is the default store.

Advanced File System Store: Nests files into directories by version, for faster performance and to store more files.

DBStore (Database Storage): Stores files in the Liferay DXP database. The file (stored as a blob) size limit is 1 GB. Use the Simple File System Store or Advanced File System Store to store larger files.

S3Store (Amazon Simple Storage): Uses Amazon's cloud-based storage solution.

If you must move your files from one store to another, use the migration utility in *Control Panel* → *Configuration* → *Server Administration* → *Data Migration*.

177.1 Simple File System Store

The Simple File System Store is the default store. It stores Documents and Media files on the server's file system (local or mounted). This store is heavily bound to Liferay DXP's database. The store's default root folder is `[Liferay Home]/data/document_library`. You can change this via the `dl.store.file.system.root.dir` property in a `portal-ext.properties` file, or in the Control Panel. For instructions on this, see the Document Repository Configuration guide.

The Simple File System Store uses a local folder to store files. You can use the file system for your clustered configuration, but you must make sure the folder you point the store at can handle things like concurrent requests and file locking. You must therefore use a Storage Area Network or a clustered file system.

The Simple File System Store creates a folder structure based on primary keys in Liferay DXP's database. If, for example, you upload a presentation with the file name `workflow.odp` to a folder

named stuff, the store creates a folder structure like this:

```
/companyId/folderId/numericFileName/versionNumber
```

companyId: The site's company ID.

folderId: The ID of the Documents and Media folder containing the document.

numericFileName: The document's numeric file entry name.

versionNumber: The document's version number.




Figure 177.1: The Simple File System Store creates a folder structure based on primary keys in Liferay DXP's database.

Note: Be careful not to confuse a document's numeric file entry name from its document ID. Each has an independent counter. The numeric file entry name is used in the folder path for storing the document, but the document ID is not. The numeric file entry name can be found in the name column of the `DLFileEntry` table in Liferay DXP's database; the document ID can be found in the `fileEntryId` column of the same table.

177.2 Using the Advanced File System Store

The Advanced File System Store, like the Simple File System Store, saves files to the local file system. It uses a slightly different folder structure, however, and can overcome operating system limitations on the number of files stored in a particular folder by programmatically creating a structure that can expand to millions of files. It alphabetically nests the files in folders. This also improves performance, as there are fewer files stored per folder.

The same rules apply to the Advanced File System Store as apply to the Simple File System Store. To cluster it, you must point the store to a network mounted file system that all the nodes

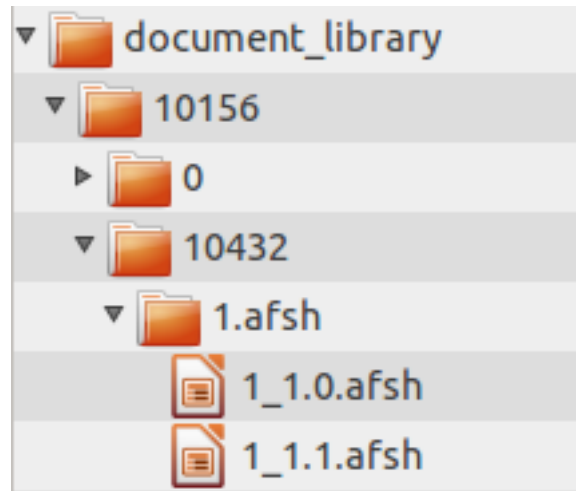


Figure 177.2: The Advanced File System Store creates a more nested folder structure than the Simple File System Store.

can access. That networked file system must also support concurrent requests and file locking. Otherwise, you may experience data corruption issues if two users attempt to write to the same file at the same time from two different nodes.

See the Document Repository Configuration guide for instructions on using the Advanced File System Store.

177.3 Using Amazon Simple Storage Service

Amazon's Simple Storage Service (S3) is a cloud-based storage solution that you can use with Liferay DXP. It lets you store your documents to the cloud seamlessly from all nodes.

When you sign up for the service, Amazon assigns you unique keys that link you to your account. In Amazon's interface, you can create *buckets* of data optimized by region. Once you create these to your specifications, follow these instructions to connect your repository to Liferay DXP.

Consult Amazon's S3 documentation for more information.

LIFERAY SYNC

Liferay Sync synchronizes files between your Liferay DXP server and users' desktop and mobile environments. With Liferay Sync, users can publish and access shared documents and files from their native environments without using a browser. Clients for Windows and Mac OS desktops and Android and iOS mobile platforms are supported. As users add and collaborate on documents and files, Liferay Sync automatically synchronizes them across all configured Sync clients and your server. Liferay Sync integrates completely with Liferay DXP so that features such as authentication and versioning function in the supported environments. The Liferay Sync desktop client stores files locally so they're always available, even when users are offline. It automatically synchronizes files upon client reconnection. The Liferay Sync mobile client saves storage space on users' devices by downloading only the files they choose.

ADMINISTERING LIFERAY SYNC

Before your users can use Liferay Sync with their Sites, you must install and configure it on your Liferay DXP server. The articles here walk you through this and also discuss important topics like preventing accidental file deletion and ensuring Sync security. As such, make sure you thoroughly read each article before letting your users connect to Sync.

Note: To install and configure Liferay Sync on your Liferay DXP server, you must be an administrator.

INSTALLING LIFERAY SYNC'S PREREQUISITES

Liferay Sync requires the *Liferay CE Sync Connector* app from Liferay Marketplace. This app enables and configures Sync in your Liferay DXP instance. For example, you can disable Sync across the instance or on a site-by-site basis. Note that Sync is enabled by default for all Sites.

For instructions on installing Marketplace apps, see the Liferay Marketplace documentation.

Note: The Liferay Sync Security module that Sync requires is included and enabled by default in Liferay DXP. You can verify this by ensuring that the `SYNC_DEFAULT` and `SYNC_TOKEN` entries are enabled in *Control Panel* → *Configuration* → *Service Access Policy*.

If you want to use Sync Connector's default settings and are fine with Sync being enabled for all your Sites, you can skip the articles that follow on configuring Sync. However, before directing your users to install and configure the Sync desktop and mobile clients, **make sure to read** this guide's articles on preventing accidental file deletion and ensuring Sync security. You should also **warn your users** about the potential for accidental data loss.

CONFIGURING LIFERAY SYNC

Sync Connector lets you manage how, or if, clients connect to your Liferay DXP server. You can also configure default file permissions on a per-Site basis, and manage the devices that connect to your Liferay DXP instance. To access Sync Connector, select *Control Panel* → *Configuration* → *Sync Connector Admin*.

Sync Connector Admin has three tabs:

1. **Settings:** Control Sync's general behavior. These settings apply globally to Sync.

Allow the use of Sync?: Whether Sync is enabled.

Allow users to sync their personal Sites?: Whether users can sync data with their personal Sites.

Allow LAN Syncing?: Whether desktop clients attempt to download updates from other desktop clients on the same local network before downloading from the server. This can help reduce server load and increase data transfer speeds. Note that LAN syncing only works with clients that also enable it.

Max Connections: The maximum number of simultaneous connections each client is allowed per account. For example, if Max Connections is three, a client can simultaneously upload or download up to three files for each account. Note, this setting operates on a per client basis. If Max Connections is set to three and a user has two clients connected to an account (which is possible if Sync is installed on two different machines), then the user is effectively allowed six simultaneous connections. While increasing Max Connections can speed up file transfers it also places a heavier load on the server. *Max Connections* is set to one by default.

Poll Interval: The frequency in seconds that clients automatically check the Liferay DXP instance for updates. For example, if set to ten, connected clients check the instance for updates every ten seconds. The default Poll Interval is five.

Max Download Rate: The maximum transfer rate, in bytes, at which clients can download. A value of 0 specifies no limit. This setting takes precedence over clients' download rate setting.

Max Upload Rate: The maximum transfer rate, in bytes, at which clients can upload. A value of 0 specifies no limit. This setting takes precedence over clients' upload rate setting.

Settings Sites Devices

General

Allow the use of Sync?
 YES

Allow users to sync their personal sites?
 YES

desktop

Allow LAN syncing?
 YES ?

Max Connections ?
1

Poll Interval ?
5

Max Download Rate ?
0

Max Upload Rate ?
0

Mobile

Force Security Mode
 NO ?

[Save](#)

Figure 181.1: The Control Panel's Configuration section contains Sync Connector Admin.

Force Security Mode: Whether to force security mode on mobile clients. Security mode encrypts Sync files on the device and requires a passcode when accessing the Sync mobile app.

2. **Sites:** Control Sync on a per-Site basis.

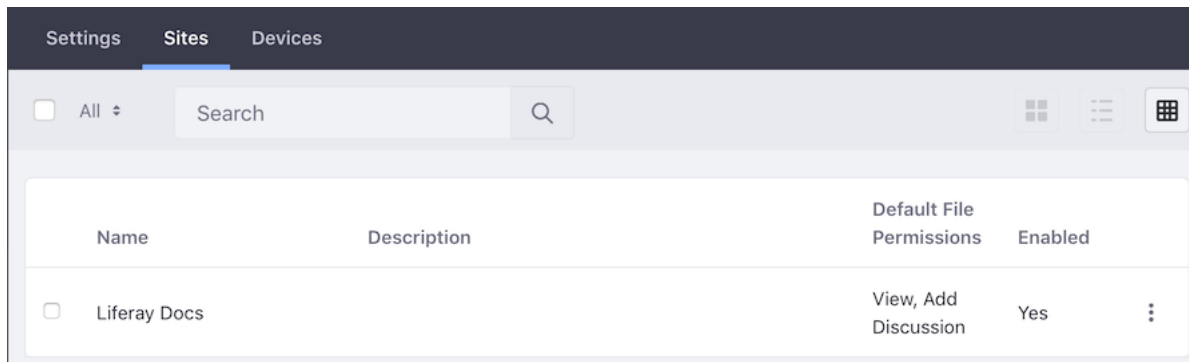



Figure 181.2: Sync Connector Admin's Sites tab lets you manage Sync on a per-Site basis.

For each Site in the Liferay DXP instance, the Sites tab lists each Site's default file permissions (more on this in a moment) and whether Sync is enabled for that Site. Sync is enabled by default for all Sites. To disable Sync for a Site, click the Site's *Actions* button () and select *Disable Sync Site*. **Please use caution** when disabling Sync for a Site, as doing so **deletes** files for that Site from the Sync clients. Disabling Sync for a Site, however, doesn't affect the Site's files on the server.

****Warning:**** Disabling Sync for specific Sites from Sync Connector Admin can result in data loss across clients. If Sync is disabled for a Site users are currently syncing, any files in the clients' sync folders for that Site are automatically deleted from their clients. If a user is offline when Sync is disabled for a Site, any offline changes or additions they make are deleted upon client reconnection.

You can enable Sync for a Site by selecting **Enable Sync Site** from its *Actions* button. Make sure that each Site for which Sync is enabled has a Documents and Media app on at least one of its pages. If a Site doesn't have the app on any of its pages and users click the **Open Website** link from their Sync menus, the error message **The requested resource was not found** appears.

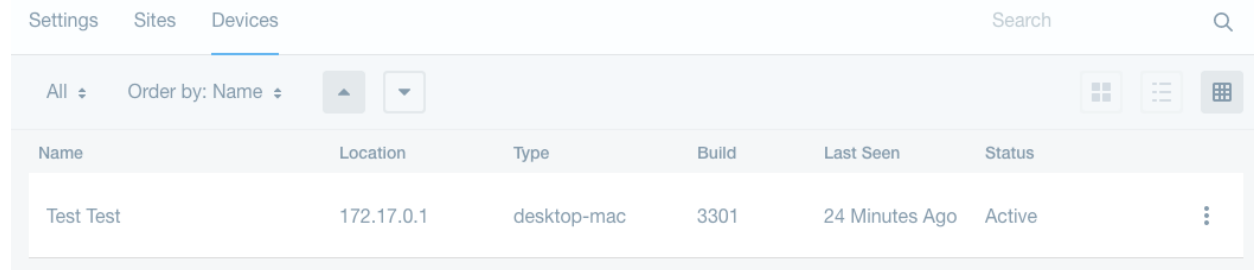
The Sites tab also sets default file permissions for files uploaded from Sync clients. The process for setting permissions is nearly the same as for enabling or disabling Sync for Sites. To set the default file permissions for a single Site, click its *Actions* button and select **Default File Permissions**. This lets you select the default file permissions for that Site. Click **Choose** for the permissions you want to use.

! [Click **Choose** to select the default file permissions for a Site in Sync.] (./images/sync-admin-03.png)

To set the default file permissions for several Sites, select the checkboxes for the Sites, click the **Default File Permissions** link that appears above the table, and select the permissions you want to use. Default file

permissions might behave differently than you'd expect. They control *only* the permissions for new files uploaded through the Sync clients; they don't affect permissions for uploading or restrict document owners (the user who originally uploaded a document) in any way. For example, even if you set a Site's default file permissions to View Only, that Site's users can still upload new documents to the Site. The file's owner has edit permission; the rest of the Site's users have the View Only permission.

3. **Devices:** View and manage the devices registered with Sync.



The screenshot shows the 'Devices' tab in the Sync Connector Admin interface. At the top, there are navigation tabs for 'Settings', 'Sites', and 'Devices', with 'Devices' being the active tab. To the right is a search bar. Below the navigation is a toolbar with 'All' (dropdown), 'Order by: Name' (dropdown), and two arrow buttons. There are also three view icons: a grid, a list, and a table view (which is selected). The main content is a table with the following data:

| Name | Location | Type | Build | Last Seen | Status | |
|-----------|------------|-------------|-------|----------------|--------|---|
| Test Test | 172.17.0.1 | desktop-mac | 3301 | 24 Minutes Ago | Active | ⋮ |

Figure 181.3: Sync Connector Admin's Devices tab lists all the devices Sync has registered.

Each row in the Devices tab's table represents a device. The *Name* column lists the user that registered the device. The remaining columns list each device's location, client type, client build number, last connection date, and status. Each device's Actions button (⋮) manages that device. You can change a device's status from Active to Inactive by selecting *Actions* → *Disable Sync Device*. Inactive devices can't sync. Inactive mobile devices also can't access local Sync files. Once a device is Inactive, you can erase Sync files from it by selecting *Actions* → *Wipe Sync Device*. This also signs the device out and removes the account from the client. If the device is offline, this happens when it tries to reconnect. The Actions menu also enables or deletes an Inactive device. Deleting a device only removes it from the list of registered devices; it can still reconnect and reregister.

PREVENTING ACCIDENTAL FILE DELETION IN LIFERAY SYNC

Liferay Sync's power rests in its ability to propagate between Liferay DXP and connected Sync clients. When a user deletes a file from a connected client, Sync also deletes the file in the instance and in any other connected clients. Likewise, if a user deletes a file in the instance, Sync also deletes the file in all connected clients. In other words, anywhere a user deletes a file, Sync deletes it *everywhere*. But don't worry: Liferay DXP's Recycle Bin is enabled by default and lets you recover deleted files. You can access the Recycle Bin from each Site's *Site Administration* menu.

Warning: Liferay Sync automatically propagates file and folder deletion through the Liferay DXP server and in all connected clients. If an instance or Site administrator disables the Recycle Bin, deleted files can't be recovered.

Liferay DXP instance and Site administrators can, of course, disable the Recycle Bin. Disabling the Recycle Bin in a Site, however, leaves the Site vulnerable to accidental file deletions that propagate through Sync.

ENSURING LIFERAY SYNC SECURITY

As an administrator, you have a stake in the security of all connections to and from your servers. As long as Liferay DXP is configured to use HTTPS, Sync clients use user-supplied credentials to communicate securely. Users can only access the documents and Sites they're permitted to access. To support Security Mode in the Sync mobile client and securely transmit files, your Liferay DXP server must also use SSL. The next section demonstrates how Sync's permissions work with your Liferay DXP instance's permissions.

183.1 Liferay Sync Permissions Demonstration

Sync uses Liferay DXP's default permissions to determine files and folders to sync with the user's devices. It can only sync files a user can access. After installing the desktop Sync client, follow the steps below to test this functionality.

First, enter classified information into a new text file and save it on your desktop as `secret.txt`. Then use your browser to sign into Liferay DXP and create a new user with the user name *secretagent* and the email address *secretagent@liferay.com*. Give this user a password and then create a new private Site called *Secret Site*. Create a page on the Site and add the Documents and Media app to it. Then add the *secretagent* user to the Secret Site and grant the *Site Administrator* Role to the user. Log in as *secretagent* and navigate to the Secret Site. Then upload the `secret.txt` document to the Documents and Media app. Make sure you also have a user that isn't a member of the Secret Site and therefore doesn't have access to any of its documents through Sync. If you don't have such a user, create one now.

Next, configure your Liferay Sync client to sign in with the *secretagent* user's credentials and sync with the Secret Site. Open the Liferay Sync menu from the system tray and select *Preferences*. In the *Accounts* tab, click the plus icon at the window's bottom left to add an account. Provide the *secretagent* user's credentials and uncheck all Sites except the Secret Site. Now confirm that Sync downloaded the `secret.txt` file to your new Sync folder. Open it and check that it contains the text classified information. Next, use Sync to connect to your Liferay DXP instance with the user that doesn't belong to the Secret Site. The file doesn't sync because this user isn't a Site member.

Now go to Sync Connector Admin and set the Secret Site's default file permissions to View Only. Create a new user, add it to the Secret Site, and add its account in your Liferay Sync client. As with the secretagent user, Sync downloads the secret.txt file to this user's local Sync folder because the user is a member of the Secret Site. Now edit and save this file. Even though you can edit and save it locally, the edits aren't synced because the Site's default file permissions are View Only. After attempting the sync, a red x appears next to the file in the local Sync folder. Right click the file to see the error. It confirms the user doesn't have the required permissions.

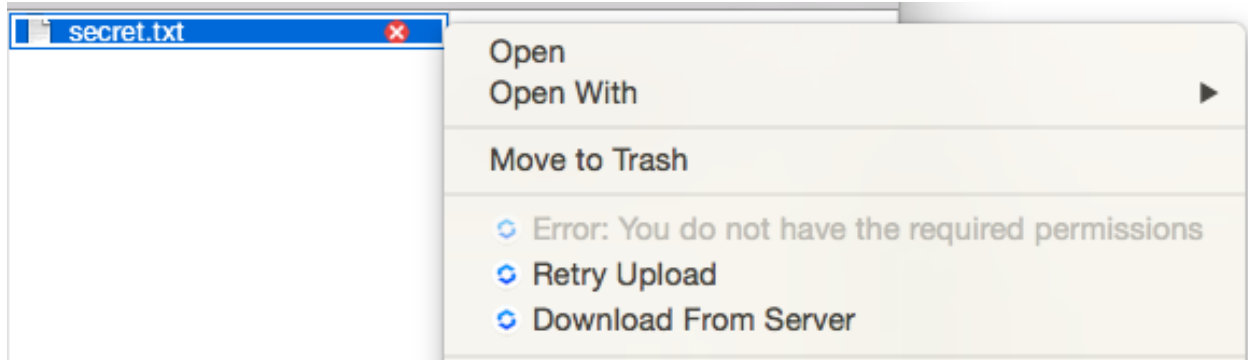


Figure 183.1: The upload error occurs because the user only has permission to view files.

To confirm that the error didn't propagate through Sync, open the file in the secretagent user's local Sync folder. It still contains the original text. Likewise, the original file remains in the Site's Documents and Media portlet. To get rid of the error in the other user's local Sync folder, return there and then right click the file and select *Download From Server*. This replaces the file with the latest file in the Liferay DXP instance.

Now edit secret.txt in the secretagent user's local Sync folder. When you check the file in the other user's local Sync folder and in the Liferay DXP instance, notice that Sync propagated the edits. The changes were propagated because the secretagent user owns the file in the instance. Owners can do anything with their files, even when the Site's default file permissions are set to View Only.

USING LIFERAY SYNC ON YOUR DESKTOP

Liferay Sync synchronizes files between your Liferay DXP Sites and desktop devices. It lets you work with your files without using a browser. The Sync clients also ensure that the files are updated with the latest changes made by other users. To use Liferay Sync in your desktop environment, you must install the Sync desktop client. It's currently available for Windows and Mac OS. The Sync client stores files locally so that they're always available, even when you're offline. Files are automatically synchronized upon your client's reconnection to your Liferay DXP server.

On your desktop devices, Liferay Sync creates a new folder structure that it uses to synchronize files. You can treat the files the same as you do any others. Credentials, Sync folder location, and other options are configured in the client. Also, native desktop notification events inform you of what Sync is doing. The native menu and task bar integration keep Sync controls within easy reach.

This guide walks you through setting up and using the Liferay Sync client on your desktop. Before proceeding, check your Liferay DXP instance or Site administrator to ensure that Sync is enabled for your Sites.

INSTALLING AND CONFIGURING THE DESKTOP LIFERAY SYNC CLIENT

You can download the desktop client from the Liferay customer downloads page. Note that you'll need a Liferay account for this. Once you've downloaded the appropriate desktop client for your operating system, installing Liferay Sync on Windows or Mac OS is straightforward.

185.1 Installing the Liferay Sync Desktop Client

To install the Liferay Sync client on Windows, you must have administrator privileges. Upon launching the Windows application installer, you're prompted to choose an install location. Select an appropriate location and click *Install*. Sync automatically starts after the installation finishes. The first time Sync runs, you must configure it to connect and sync with Liferay DXP. The configuration steps are shown below.

Note: You can upgrade previous versions of the desktop Liferay Sync client to version 3.0. When doing so, however, you must set up your account again in the new version of the client. Prior to upgrading, it's typically best to shut down Liferay Sync, backup files from your local Sync folder, and delete that folder.

The Liferay Sync client for Mac is packaged in a DMG file. Double-clicking on the DMG file mounts it as a disk image and opens a window showing the image's contents. To install Sync, drag the Liferay Sync icon to your Applications folder. Once it's installed, run it from the Applications folder. If you're running Mac OS X 10.9 or lower, you're prompted for your machine's administrator credentials to install the Finder icon/context menu tool. This prompt only appears when installing or upgrading the tool.

Next, you'll configure the Sync client.




Figure 185.1: Drag the Liferay Sync icon to the Applications folder.

185.2 Configuring the Liferay Sync Desktop Client

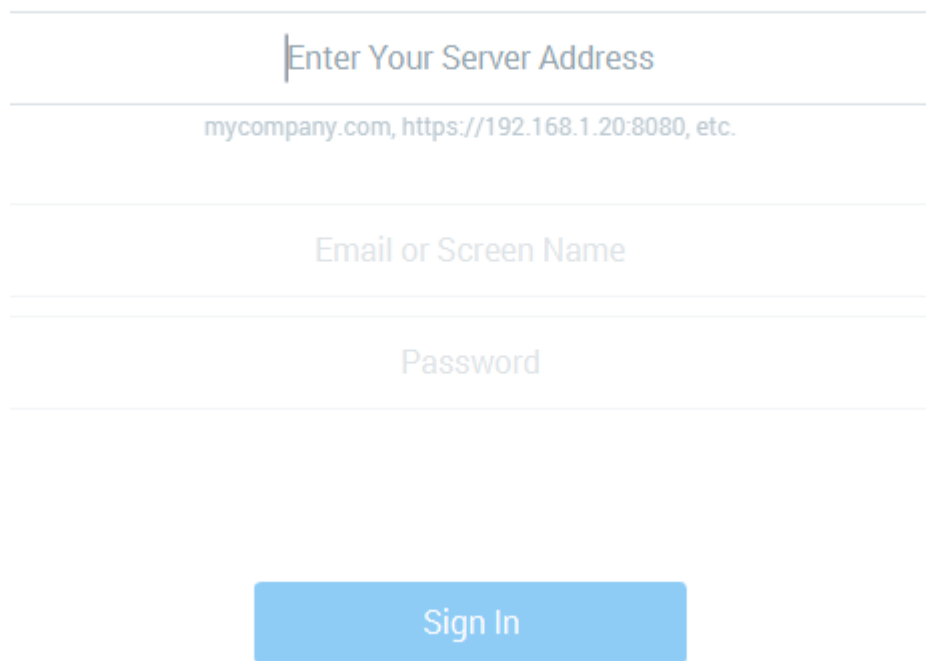
Now that you've installed Sync, you're ready to configure it! The configuration steps for Sync on Windows and Mac are identical.

1. Open Sync and enter your server's address along with your account credentials. Click *Sign In* when you're finished.

When connecting to a server via HTTPS, an error appears if the certificate can't be verified. Choosing *Proceed Anyway* bypasses verification and leaves the connection open to compromise. Liferay Sync attempts to read the certificates specified in the Java Control Panel (see section 20.4.5). If Java isn't installed, you can also put your certificates in `[user.home]/.liferay-sync-3/certificates`. Liferay Sync trusts all certificates in this folder.

2. Select the Sites you want to sync with. You can search for a Site in the *Search* bar above the Site list. If you want to sync all the subfolders of your selected Sites, click *Proceed* and move on to the next step.

To sync only specific folders in a Site, first click the Site's gear icon. In the window that appears, all folders are selected by default. Unselect the folders you don't want to sync with. Unselecting a subfolder causes the parent folder's checkbox to show a minus sign, indicating that you haven't selected all of the parent folder's subfolders. To sync only the documents at the top of a folder's hierarchy (no subfolders), unselect all of that folder's subfolders. You can also do this by clicking the folder's checkbox until the minus sign appears. Click *Select* when you're finished with your selections, and then click *Proceed* to move on to the next step.



The screenshot shows a sign-in form with three input fields and a button. The first field is labeled "Enter Your Server Address" and contains the placeholder text "mycompany.com, https://192.168.1.20:8080, etc.". The second field is labeled "Email or Screen Name". The third field is labeled "Password". Below the fields is a blue button labeled "Sign In".

Figure 185.2: The first time you run Liferay Sync, you must tell it how to communicate with your Liferay DXP server.




Figure 185.3: When connecting over HTTPS, Liferay Sync produces an error if it can't verify the security certificate. Choosing *Proceed Anyway* bypasses verification and leaves the connection open to compromise.

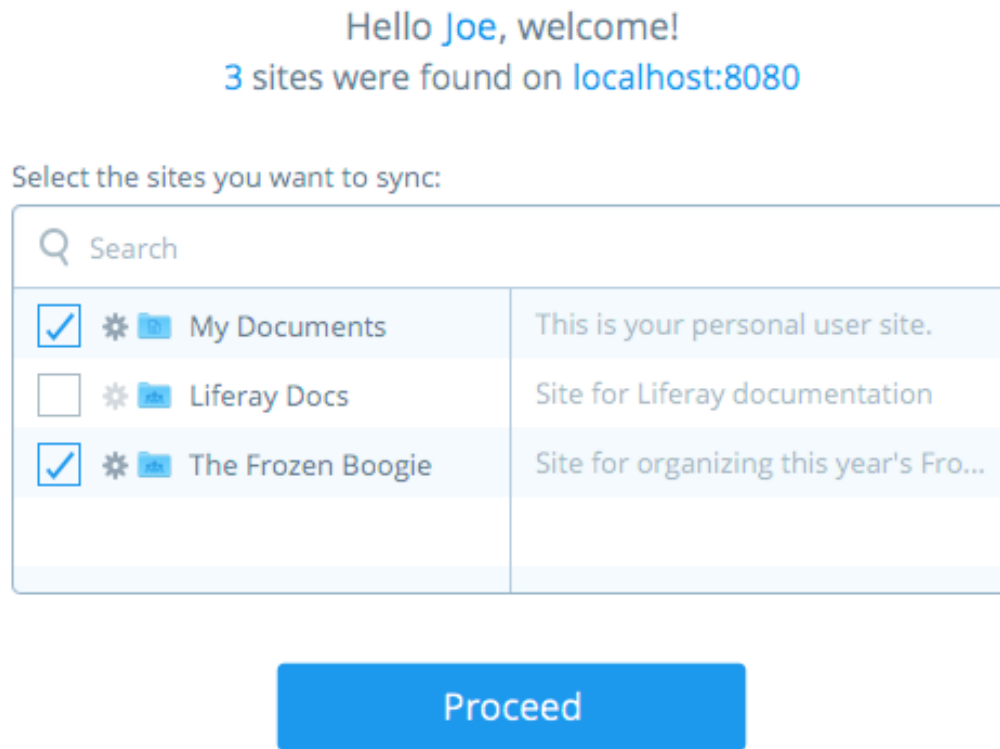


Figure 185.4: Select the Sites you want to sync with. Clicking a Site’s gear icon opens another window where you can choose to sync with only specific subfolders in that Site.

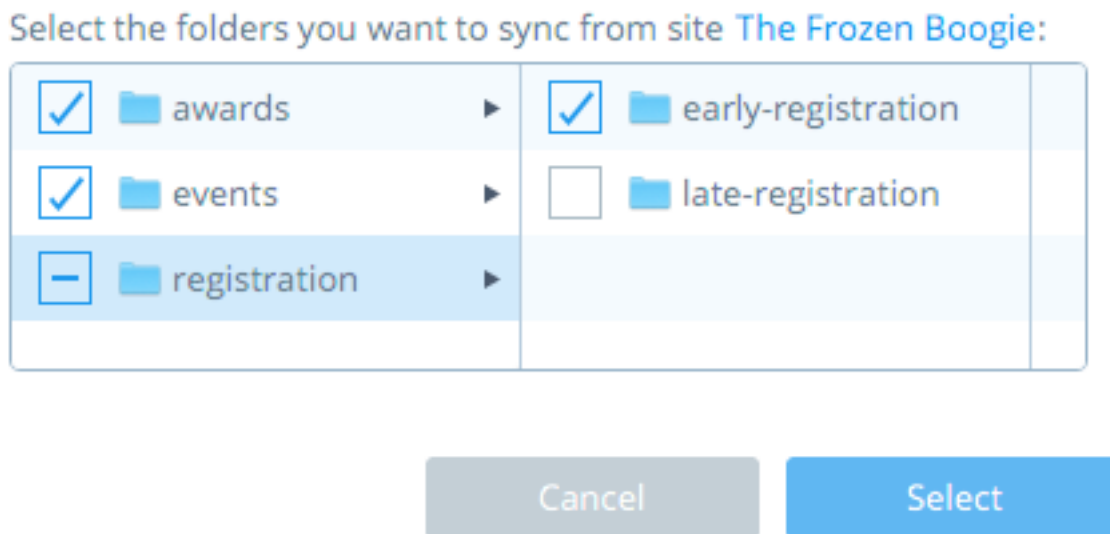


Figure 185.5: Choose the Site’s subfolders that you want to sync with. The checkbox with the minus sign indicates that not all of the *registration* folder’s subfolders are selected.

3. Specify the local folder to sync. This folder is used exclusively for Sync: Sync creates it and it must not conflict with any existing local folder. The Sync folder's default name is the instance's host name, and its default location is the user's documents folder. For example, since the instance in the following screenshots runs locally at the address `http://localhost:8080/`, Sync creates a Sync folder named *localhost* in the user's documents folder. You can, of course, specify any unique name and location for the Sync folder. Click *Start Syncing* to begin syncing files.

****Note:**** Syncing to network drives is not supported because Liferay Sync can't reliably detect local file changes on such drives.

![Specify your local Sync folder's name and location.](./images/sync-setup-03.png)

4. Celebrate! You've successfully set up Liferay Sync! Sync congratulates you on setting it up and begins to sync files from the Sites you selected to your local Sync folder. Note, completing the initial synchronization may take a significant amount of time, depending on the amount of data being transferred. You can safely close the window as syncing continues in the background. To view the local Sync folder, click *Open Folder*. To open Sync's preferences, click the small gray text *advanced setup* near the top-right.

advanced setup

Congratulations **Joe!**
Your sites are syncing in:



localhost

Open Folder

close this window

Figure 185.6: Congratulations, you've successfully set up Liferay Sync!

USING THE LIFERAY SYNC DESKTOP CLIENT

When Liferay Sync is running, its icon appears in your task bar (Windows) or menu bar (Mac). Clicking this icon opens a menu that lets you work with and manage Liferay Sync.

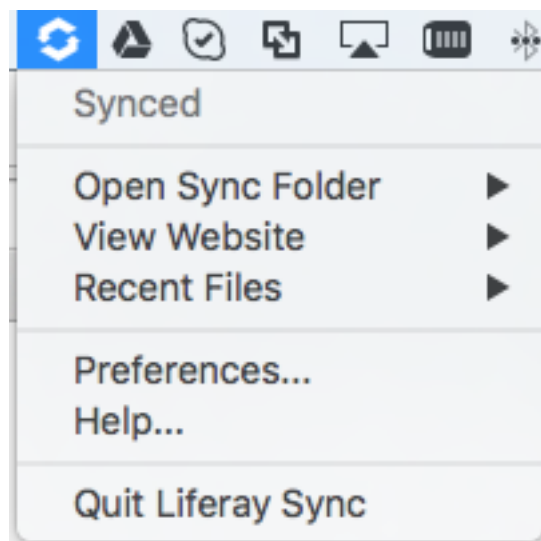


Figure 186.1: The Liferay Sync menu in the Windows task bar and Mac menu bar gives you quick access to Sync.

The top of this menu shows your Sync status. If all your selected Sites are synced, then your status is *Synced*.

Below your Sync status, the menu lists three shortcuts for accessing your files:

Open Sync Folder: Select a Site to open its local Sync folder.

View Website: Select a Site to view the page in Liferay DXP that contains its Documents and Media app.

Recent Files: Lists recently created and modified files in the repositories you can access.

Note that if you sync with two or more Liferay DXP instances, Sync shows each at the top of the menu instead of your Sync status. Mouse over each instance to reveal a submenu with that instance's Sync status and file shortcuts.

Finally, regardless of how many Liferay DXP instances you sync with, the menu lists the following three options:

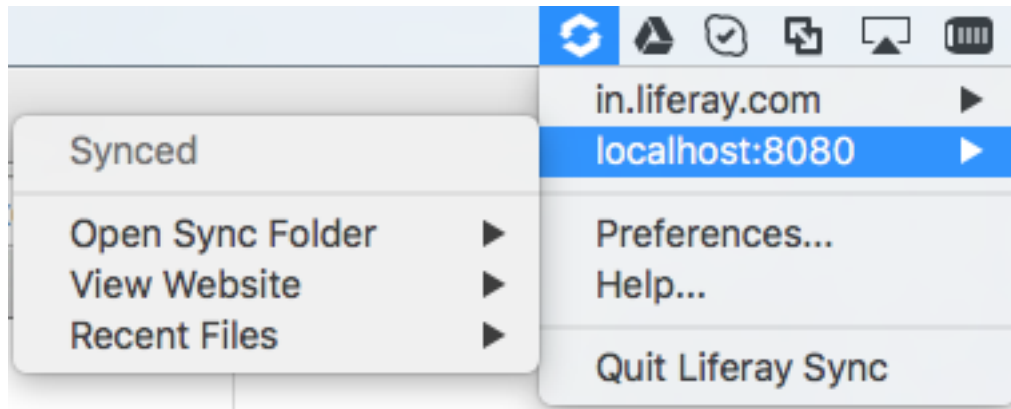


Figure 186.2: When you sync with more than one Liferay DXP instance, Sync shows submenus for each.

Preferences: Open Sync’s preferences.

Help: Open Sync’s documentation.

Quit: Shut down Sync on your machine.

Next, you’ll learn how to use Sync’s preferences to control how Sync functions on your machine.

186.1 Using Sync Preferences

You can use Sync’s preferences to add/remove Liferay DXP instances to sync with, edit connection settings, and control Sync’s basic behavior. Open Sync’s preferences by clicking the Sync icon in the task bar (Windows) or menu bar (Mac OS) and selecting *Preferences*. A preference screen for your instance accounts displays. This is the *Accounts* tab in *Preferences*.

The *Accounts* tab contains the following options:

Accounts: The accounts you sync with. When you select an account, the Sites you have permission to sync with are shown on the right under *Syncing Sites*. You can use the plus, minus, and pencil icons at the bottom of the account list to add, delete, or edit an account, respectively. You should use caution when deleting an account from your Sync client, as doing so also deletes any local files and folders for that account. Adding an account takes you through the same set of steps you used to set up the Sync client. Click here for instructions on this.

Syncing Sites: The Sites you have permission to sync with for the selected account. The Sites you currently sync with are shown under *Selected Sites*. Other Sites available for syncing are shown under *Unselected Sites*. To change the Sites you sync with, click the *Manage Sites* button. The window that appears lets you select and/or unselect Sites to sync with. This window is identical to the one that appeared when you first configured the client. Click here and see step two for instructions on using it. Use caution when de-selecting Sites. De-selecting a Site deletes its folder on your machine.

Location: The selected account’s local Sync folder location. Click the *Change* button to change this folder’s location.

The Preferences menu’s *General* tab contains settings for the Sync client’s general behavior. It lists the following options:

Launch Liferay Sync on startup: Starts Sync automatically each time your machine starts.

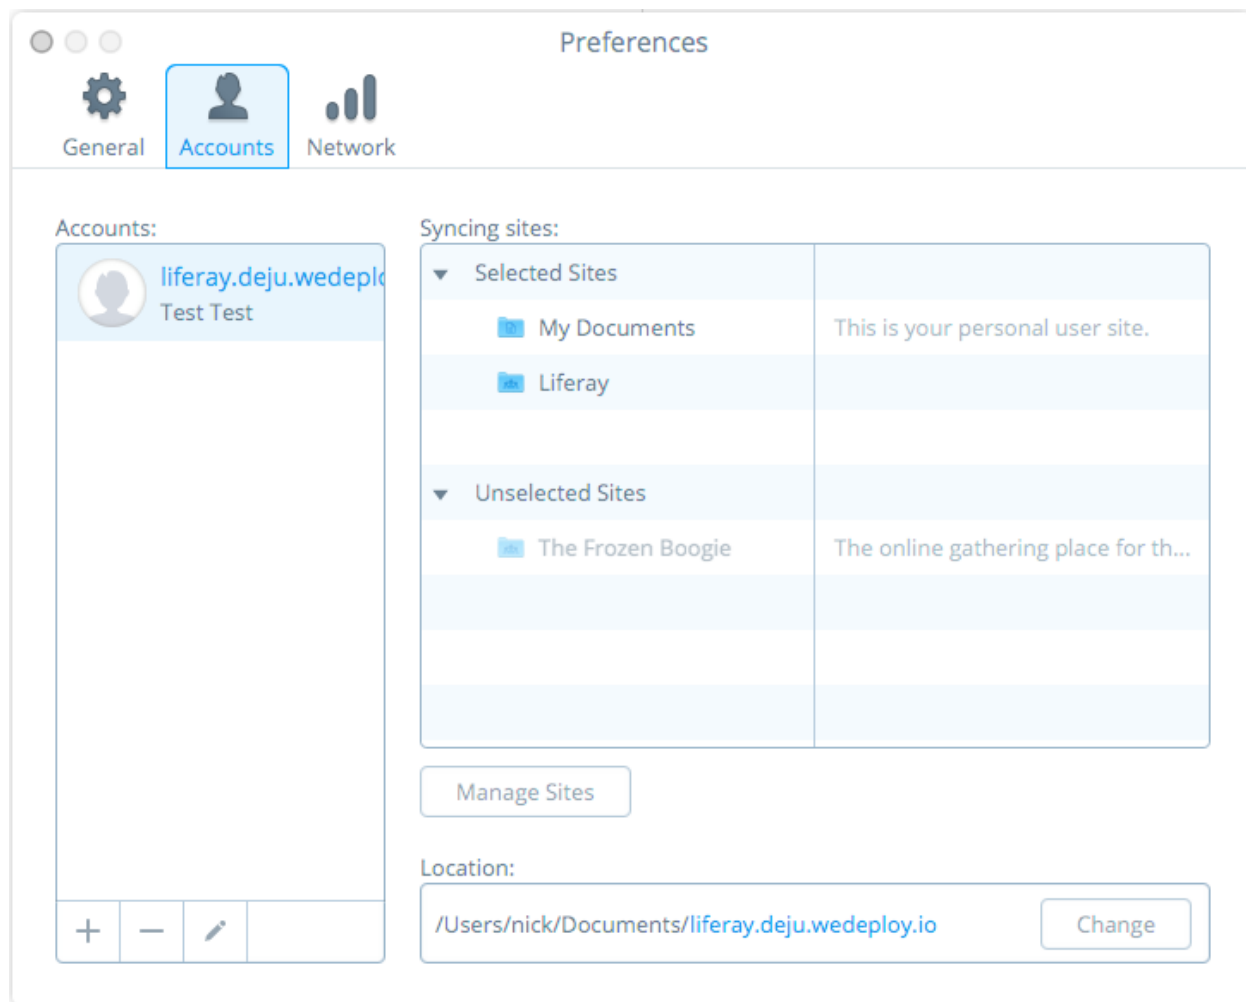


Figure 186.3: The Preferences menu's *Accounts* tab lets you manage syncing with Sites per account.

Show desktop notifications: Shows a small notification in the corner of your screen when a synced file changes.

Automatically check for updates: Automatically check for new client versions. You can click the *Check Now* button to check for updates manually.

Finally, the Preferences menu's *Network* tab controls how Sync transfers data with your Liferay DXP servers. It contains the following options:

Download Rate: To limit the rate at which Sync downloads data, select *Limit to* and then specify the rate.

Upload Rate: To limit the rate at which Sync uploads data, select *Limit to* and then specify the rate.

Enable LAN Syncing: Whether to download updates from other desktop clients on the same local network before downloading from the server. This can help reduce server load and increase data transfer speeds. Note that LAN syncing only works when enabled in the Liferay DXP instance by the administrator, and in other clients.

Note that your Liferay DXP administrator can also limit the download/upload rate. In this case,

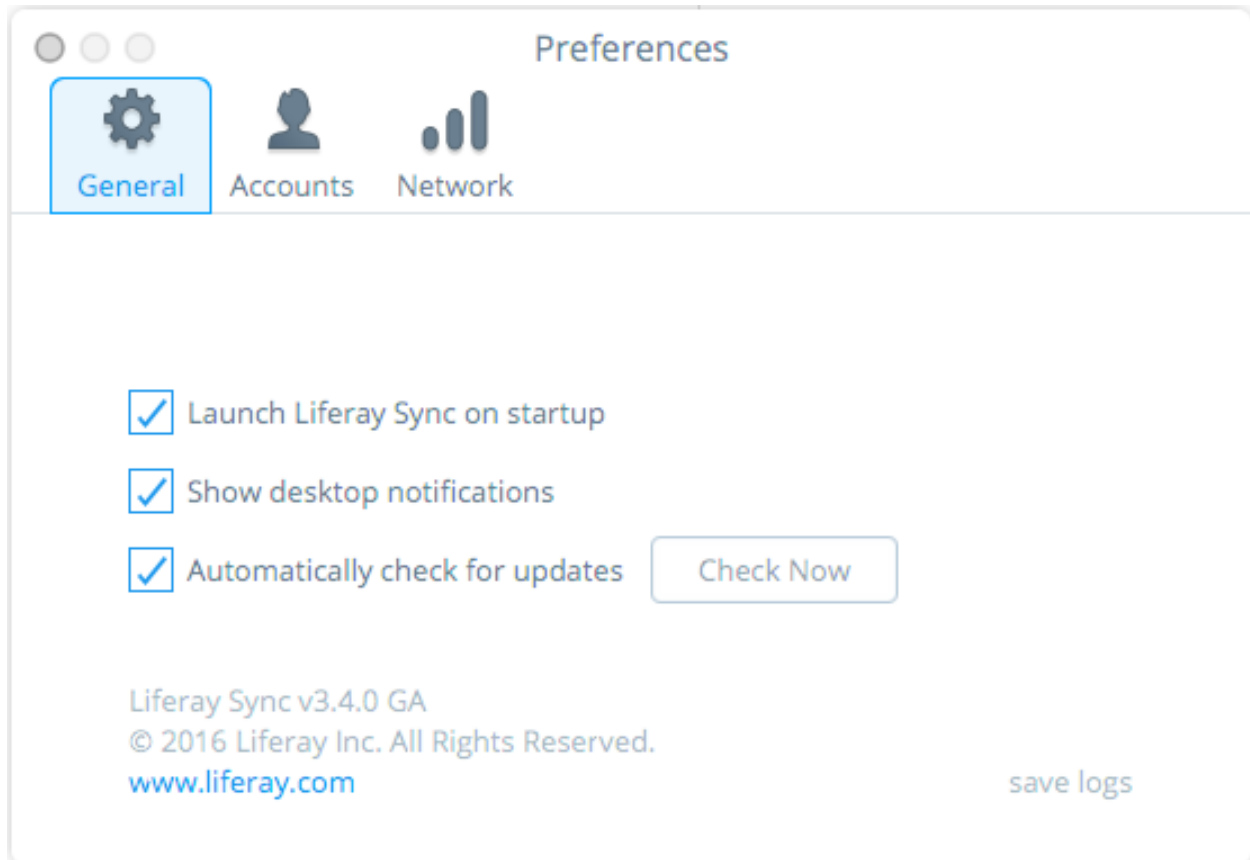


Figure 186.4: The Preferences menu's *General* tab contains settings for Sync's general behavior.

Liferay DXP's settings take precedent. For example, if you set a 5.0 MB/s download rate in the client but Liferay DXP's download limit is 2.0 MB/s, the latter takes precedence. Also, the client's rate applies across all its accounts. For example, if the client connects to three accounts and its download rate is 5.0 MB/s, then the sum of the download rate for all three accounts never exceeds 5.0 MB/s.




Figure 186.5: The Preferences menu's *Network* tab contains settings for Sync's data transfer behavior.

USING YOUR LOCAL LIFERAY SYNC FOLDER

Once you configure and run Sync, Sync automatically uploads to Liferay DXP any files you add or modify in your Sync folder. Sync also downloads to your Sync folder any file changes by other users. If you delete a file in your Sync folder, Sync also deletes it from the server and other clients. You should therefore use **extreme caution** when deleting files in your Sync folder. If, however, you accidentally delete a file, all is not lost! The file can still be recovered from the instance's Recycle Bin, which is enabled by default. Note, if the administrator has disabled the Recycle Bin, recovering deleted files is impossible.

Warning: Deleting a file in your Sync folder also deletes it in the Liferay DXP instance and in other clients. If you accidentally delete a file, it can be recovered from the instance's Recycle Bin. The Recycle Bin is enabled by default. File recovery is, however, impossible if the instance or Site administrator has disabled the Recycle Bin.

You can run through the following exercise to familiarize yourself with how to create, edit, download, and upload files with Sync. First, open the Sync folder in your file manager and create a new file called `README.txt`. Enter the word `test` in this file. Next, make sure you can access this file in your Site. Go to the Site you want to sync with and navigate to its Documents and Media app. It lists your `README.txt` file.

Download the `README.txt` file to a convenient location on your machine. Open the file and check that it still says `test`. Now open the `README.txt` file in your Sync folder and edit it so that it says `second test`. Once the changes are synced, go back to your browser and refresh the page with your Documents and Media app. Click on the `README.txt` file's name, look at the file information displayed, and check that the file's version number has been incremented.

If you download and open `README.txt` again, it now says `second test`. Your edit was uploaded to the Site! You can be confident that this edit was also downloaded by all other Sync clients connected to your Site.

Now delete the `README.txt` file from your local Sync folder. When the changes finish syncing, go back to your browser and refresh the page containing your Documents and Media app. The file is gone! The file is also deleted from the local Sync folders of all other Sync clients connected to the Site. Remember this very important rule: deleting files in your local Sync folder deletes them *everywhere*! Next, you'll learn how to use the Sync client for your mobile device.




Figure 187.1: Updating a file through Liferay Sync increments the file's version number. You can view a file's version number through the web interface.

USING LIFERAY SYNC ON YOUR MOBILE DEVICE

Liferay Sync for Android and iOS contains most of the desktop Sync client's functionality. The mobile client can, however, only be connected to one Liferay DXP account at a time. Also, mobile Sync doesn't automatically download files to your device. To save storage space on your device, the Sync mobile app lets you choose the files you want to work with. As with the Sync desktop clients, the latest versions of Sync on Android and iOS provide a consistent user experience across platforms. While this article details using Sync on Android, the instructions also apply to Sync on iOS.

You need to download and install Sync on your Android or iOS device through its respective app store, the same as you do any other mobile app. To find the app, search Google Play or the App Store for *Liferay*. You can also download Sync from the Liferay customer downloads page. Once you've installed the Sync app on your device, the rest of the articles in this guide show you how to use it.

CONNECTING LIFERAY SYNC MOBILE

When Liferay Sync first starts on your mobile device, press the *Get Started* button to begin setup. The setup screen asks for your login credentials and your server's address. Once you enter them, press *Sign In*. After signing in, you see a panel that shows your name, a gear icon for accessing the app's settings, and navigation options *My Sites* and *My Documents*. *My Sites* and *My Documents* encompass the Sites in Liferay DXP that you can sync with. *My Documents* is your personal user Site, while *My Sites* shows the other Sites with which you can sync. No matter how deep you are in the folder hierarchy of a Site, swiping to the right returns you to this panel. If you're in the first level of *My Sites* or *My Documents*, pressing the location bar at the top slides the screen slightly to the right to reveal a compact view of the panel. The following screenshots show both views of the panel.

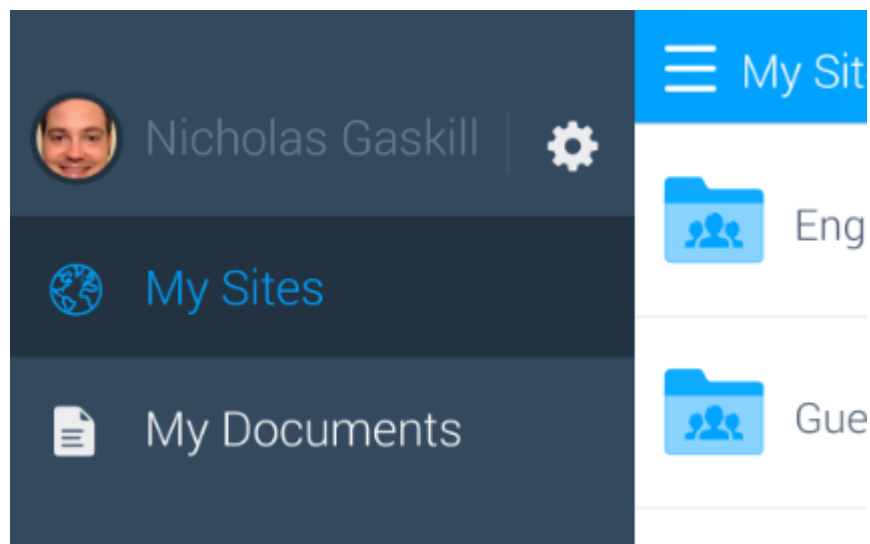


Figure 189.1: This panel lets you access the app's settings, as well as your Sites and documents.

Press the gear icon to access Sync's settings. Settings shows your account information and an option to sign out of your Liferay DXP instance. Settings also lets you toggle *Security Mode*. Security Mode protects files stored on your device by encrypting them. Using Security Mode requires you to set up a passcode to use when accessing the Sync app. Security Mode protects the files on




Figure 189.2: Tapping the title bar at the top of My Sites or My Documents opens the main Sync panel's compact view.

your device and Liferay DXP instance in the event your device is lost or stolen. You should note, however, downloading and opening files in Security Mode takes slightly longer than usual because the Liferay DXP server must use SSL—if it didn't, your files would be transmitting in the open. Below the Security Mode toggle are the app's version and a link to send app feedback to Liferay.




Figure 189.3: The Settings screen for the Sync app lets you sign out of your Liferay DXP instance, enable Security Mode, view the app's version, and send feedback.

MANAGING FILES AND FOLDERS IN LIFERAY SYNC MOBILE

Whether you're working in My Documents or My Sites, you manage files and folders the same way. Pressing a Site or folder shows you a list of its files and folders. It displays each file's size and modification date. You can refresh the list by pulling down from the top of the screen. Your current location in the navigation hierarchy also appears at the top of the screen alongside a plus icon. Pressing the plus icon launches an upload screen for adding content in the current location. You can add a new folder, upload a file, or launch your device's camera app to take and upload a picture or video. Pressing the X icon on the upload screen's top right corner cancels any action and returns you to the current file list.

To download a file to your device, press the file's name in the list. The label that previously showed the file's size and modification date is replaced by a download progress indicator. When the file finishes downloading, your device automatically opens it in the app you've configured to open files of that type. If you haven't configured your device to use a specific app for that file type, you're presented with a list of apps on your device that can open the file. If your device doesn't have an app that can open the file, Sync tells you to install one that can. Downloaded files appear in the list with the file size in blue instead of gray. For example, the screenshot below shows that `LiferayinAction.pdf` is on the device.

The Sync mobile app also lets you move, rename, and delete files and folders. To the right of each file and folder in the list is a circle icon with three dots. Pressing this icon slides open a context menu on the right that lets you move, rename, or delete that item. The screenshots below show these options. Note that you should use **extreme caution** when deleting files or folders. Deleting files or folders in the mobile Sync app also deletes them from Liferay DXP and across any synced clients. Accidentally deleted files can be restored from the Recycle Bin, which is enabled by default. If the instance or Site administrator disables the Recycle Bin, however, recovering deleted files is impossible.

What if you want to delete a file on your device without also deleting it in the instance? Currently, you can only do this by signing out of your account in the app's Settings menu. Doing so removes all downloaded files from your device, but preserves them in the instance. If you're on Android, it may be possible to use a system file browser app to remove downloaded files manually.

Warning: Deleting a file in the mobile Sync app deletes it in Liferay DXP and across any synced




Figure 190.1: Sync shows files and folders in a list.




Figure 190.2: Downloaded files appear in the list with their size in blue.

clients. If you accidentally delete a file, the instance or Site administrator can restore it from the instance's Recycle Bin. The Recycle Bin is enabled by default. If the instance or Site administrator disables the Recycle Bin, however, recovering deleted files is impossible.

The context menu also provides additional options for files. A small badge on the file icon's top-right corner indicates the file's version in the Liferay DXP instance. You can also use the context menu to share files you've downloaded. Pressing the *Share* icon opens a list of your device's apps capable of sharing the file. To close the context menu and return to the list of files and folders, swipe to the right. The following screenshot shows the options available in a file's context menu.




Figure 190.3: The badge on the file's icon shows the file's version in the Liferay DXP instance. You can also share files that you've downloaded.

ADAPTING YOUR MEDIA ACROSS MULTIPLE DEVICES

Media providers must consider differences between devices (phones, laptops, tablets, etc.) when delivering content: not only their screen size but also their bandwidth and processing capabilities. Liferay DXP's Adaptive Media app allows administrators to control image quality and dynamically adjusts uploaded media to best fit the screen being used.

Note: At this time, Adaptive Media only works for images in blog entries and web content articles.

Adaptive Media integrates with Documents and Media, Blogs, and Web Content. It generates a set of images for use on various screens. When the content is accessed, Adaptive Media checks the screen type and resolution and selects appropriate the appropriate image. Adaptive Media comes preinstalled in Liferay DXP.

In this section, you'll learn how to manage and use Adaptive Media.

INSTALLING ADAPTIVE MEDIA

The Adaptive Media app is installed in Liferay DXP by default. The following sections describe the Adaptive Media app's modules, and how to prepare Adaptive Media to handle animated GIFs.

Note: Since the Adaptive Media app is installed by default, it's updated via Liferay DXP Fix Packs and Liferay Portal CE GA releases. Using Liferay Marketplace to update the app causes an error.

192.1 Adaptive Media's Modules

Some modules in the Adaptive Media app are mandatory and must be enabled for Adaptive Media to function, while others can be disabled. The Adaptive Media API modules, which export packages for the other modules to consume, are mandatory; disabling one also disables any other modules that depend on it. Here's a list of the Adaptive Media API modules:

- Liferay Adaptive Media API
- Liferay Adaptive Media Content Transformer API
- Liferay Adaptive Media Image API
- Liferay Adaptive Media Image Item Selector API

The Adaptive Media core modules are also mandatory, and must be enabled to ensure that Adaptive Media works as expected:

- Liferay Adaptive Media Document Library
- Liferay Adaptive Media Document Library Item Selector Web
- Liferay Adaptive Media Document Library Web
- Liferay Adaptive Media Image Content Transformer
- Liferay Adaptive Media Image Implementation
- Liferay Adaptive Media Image Item Selector Implementation

- Liferay Adaptive Media Image JS Web
- Liferay Adaptive Media Image Service
- Liferay Adaptive Media Image Taglib
- Liferay Adaptive Media Image Web
- Liferay Adaptive Media Item Selector Upload Web
- Liferay Adaptive Media Web

The Adaptive Media Blogs modules, which ensure that images uploaded to blog entries can be processed and adapted, are optional. Here's a list of these modules:

- Liferay Adaptive Media Blogs Editor Configuration
- Liferay Adaptive Media Blogs Item Selector Web
- Liferay Adaptive Media Blogs Web
- Liferay Adaptive Media Blogs Web Fragment

The Adaptive Media Journal modules are optional. These modules apply Adaptive Media to web content articles:

- Liferay Adaptive Media Journal Editor Configuration
- Liferay Adaptive Media Journal Web

There are two more optional modules included in Adaptive Media:

Liferay Adaptive Media Image Content Transformer Backwards Compatibility: Ensures that content created before the Adaptive Media installation can use adapted images without the need to edit that content manually. It transforms the images both at startup and when a user views the content, which can negatively affect performance. We therefore recommend that you run some performance tests before using this module in production. You can disable this module if you don't have old content, are experiencing performance problems, or your old content doesn't need adapted images.

Liferay Adaptive Media Document Library Thumbnails: Lets thumbnails in Documents and Media use adapted images. For this to work, you must first migrate the original thumbnails to adapted images. We highly recommend that you enable this module, but it's not mandatory.

Great! Now you know the mandatory and optional modules that come with Adaptive Media. The next section discusses the installation requirements for using animated GIFs with Adaptive Media. If you don't need to use GIFs, you can skip ahead to the article on adding image resolutions to Adaptive Media.

192.2 Processing Animated GIFs

To process animated GIFs, Adaptive Media uses an external tool called Gifsicle. This tool ensures that the animation works when the GIF is scaled to different resolutions. You must manually install Gifsicle on the server and ensure that it's on the PATH. Once it's installed, you must enable it in Adaptive Media's advanced configuration options.

If Gifsicle isn't installed and image/gif is included as a supported MIME type in the advanced configuration options, Adaptive Media scales only a GIF's single frame. This results in a static image in place of the animated GIF.

ADDING IMAGE RESOLUTIONS

To use Adaptive Media, you must first define the resolutions for the images delivered to users' devices. Adaptive Media then generates new images scaled to fit those resolutions, while maintaining the original aspect ratio.

To access Adaptive Media settings, open the Control Panel and go to *Configuration* → *Adaptive Media*. Here you can create and manage resolutions.

Note: Adaptive Media configurations apply only to the current Liferay DXP instance.

Once you create a resolution, Adaptive Media automatically generates copies of newly uploaded images in that resolution. Images uploaded before you create the resolution aren't affected and must be adapted separately (see *Generating Missing Adapted Images*).

| Name | State | Adapted Images | Max Width | Max Height |
|--|---------|-----------------------------------|-----------|------------|
| <input type="checkbox"/> Preview 1000x0 | Enabled | <div style="width: 0%;"></div> 0% | 1000px | Auto |
| <input type="checkbox"/> Thumbnail 300x300 | Enabled | <div style="width: 0%;"></div> 0% | 300px | 300px |

Figure 193.1: Adaptive Media's image resolutions are listed in a table.

193.1 Adding a New Image Resolution

The number of image resolutions required and the values for each depend on the use case. More resolutions may optimize image delivery, but generating more images requires additional computational resources and storage space. To start, we recommend that you create resolutions to cover common device sizes like mobile phones, tablets, laptops, and desktops. If most users use one device (e.g., all Intranet users have the same company mobile phone), you can create a resolution to target that device.

To add a new resolution, click the *Add* icon (+) on the Adaptive Media configuration page and provide the following information:

Name: The resolution's name (this must be unique). This can be updated if a custom Identifier is defined.

Max Width: The generated image's maximum width. If a *Max Height* is given, this field is optional. This value must be at least 1.

Maximum Height: The generated image's maximum height. If a *Max Width* is given, this field is optional. This value must be at least 1.

Note: Adaptive Media generates images that fit the Max Width and Max Height, while retaining the original aspect ratio. If you only provide one value (either Max Width or Max Height), the generated image scales proportionally to fit within the specified dimension, while maintaining its original aspect ratio. This ensures that adapted images are not distorted.

Add a resolution for high density displays (2x): Defines a scaled up resolution for HIDPI displays. Selecting this option creates a new resolution double the size of the original with the same name and the suffix -2x. For example, if the original resolution is 400px by 300px (max width by max height), the high density resolution is 800px by 600px.

Identifier: The resolution's ID. By default, this is automatically generated from the name. You can specify a custom identifier by selecting the *Custom* option and entering a new *ID*. Third party applications can use this ID to obtain images for the resolution via Adaptive Media's APIs.

Note: Image resolutions and their identifiers can't be updated if the resolution has been used to adapt images. This prevents inconsistencies in generated images.

Name *

High Definition

Description

Image resolution for high definition

Size *

Please enter at least one of the following fields ?

Max Width (px) **Max Height (px)**

3200 2800

Add a resolution for high density displays (2x).

Identifier

Automatic ?

Custom

ID

high-definition

Save **Cancel**


Figure 193.2: The form for adding a new Adaptive Media resolution.

MANAGING IMAGE RESOLUTIONS


Adaptive Media lets you manage image resolutions and their resulting adapted images. For example, you can disable, enable, edit, and delete resolutions. You can also generate any adapted images that may be missing for a resolution. This article discusses these topics and more.

194.1 Disabling Image Resolutions

Disabling an image resolution prevents it from generating adapted images. Any images uploaded after the resolution is disabled use the most appropriate resolution that's still active. Adapted images previously generated by the disabled resolution are still available.

To disable an image resolution, click its *Actions* menu () and select *Disable*.

194.2 Enabling Image Resolutions

Image resolutions are enabled by default. If you need to enable a disabled resolution, click that resolution's *Actions* menu () and select *Enable*.

While a resolution is disabled, it doesn't generate adapted images for new image uploads. After enabling a resolution, you should generate the adapted images that weren't generated while it was disabled (see [Generating Missing Adapted Images](#) for instructions on this).

194.3 Editing Image Resolutions


You can't edit an image resolution that already has adapted images. This prevents odd behavior (of the adapted images—you're still free to be as odd as you want). This is because any changes would only be applied to images uploaded after the edit, creating an inconsistent set of adapted images. Odd indeed.

Therefore, editing an image resolution is only possible if Adaptive Media hasn't yet generated adapted images for it. If you must change the values of a resolution that already has adapted images, you must delete that resolution and create a new one with the new values. The next section discusses deleting resolutions.

194.4 Deleting Image Resolutions

Be careful when deleting an image resolution, as any adapted images it created are irretrievably lost and are not automatically replaced by new image resolutions you create.

Follow these steps to delete an image resolution:

1. Disable the resolution. You can't delete enabled resolutions. This prevents the accidental deletion of image resolutions.
2. To delete the resolution and all its adapted images, select *Delete* from the resolution's Actions menu ().


194.5 Generating Missing Adapted Images

If Adaptive Media hasn't generated all the images you need—say, if new images were uploaded before a new image resolution was created or while the resolution was disabled—you must generate the missing images manually.

| Name | State | Adapted Images | Max Width | Max Height | |
|--|---------|--|-----------|------------|---|
| <input type="checkbox"/> High Definition | Enabled | <div style="width: 0%; background-color: #ccc;"></div> | 3200px | 2800px | ⋮ |
| <input type="checkbox"/> Medium | Enabled | <div style="width: 0%; background-color: #ccc;"></div> | 700px | 700px | ⋮ |
| <input type="checkbox"/> Thumbnail 300x300 | Enabled | <div style="width: 100%; background-color: #008000;"></div> 100% | 300px | 300px | ⋮ |

Figure 194.1: The *Adapted Images* column shows the percentage of images that are adapted for each resolution.

To manually generate missing adapted images,

1. For a single resolution, select *Adapt Remaining* from the resolution's Actions menu ().
2. For all resolutions at once, select *Adapt All Images* from the Actions menu in the Control Menu at the top of the page.

194.6 The Recycle Bin and Adapted Images

You can't move adapted images directly to the Recycle Bin. But if the original image is in the Recycle Bin, the corresponding adapted images behave as if they are in the Recycle Bin and users can't view them.

Note: URLs that point to adapted images whose original image is in the Recycle Bin return an error code of 404 Not Found.

If the original image is restored from the Recycle Bin, the adapted images are accessible again. Awesome! Now you know how to manage image resolutions in Adaptive Media. Next, you'll learn about creating content with adapted images.

CREATING CONTENT WITH ADAPTED IMAGES

Adaptive Media is mostly invisible for blog and web content creators. Once an image is added to the content, the app works behind the scenes to deliver an adapted image appropriate to the device in use. Content creators select an image when adding it to their content—they don't have to (and can't) select an adapted image. Adaptive Media identifies each adapted image in the content's HTML with a `data-fileentryid` attribute that is replaced with the latest adapted image when the user views the content. This lets Adaptive Media deliver the latest adapted images to your content, even if the content existed prior to those images.

Note: If Adaptive Media is uninstalled, the original images are displayed in the blog entries and web content articles.

195.1 Including Adapted Images in Content

Since Adaptive Media delivers the adapted images behind the scenes, content creators should add images to blog entries and web content as usual: by clicking the image button in the editor and then selecting the image in the file selector.

However, there are some important caveats. When using the file selector to include an image for a blog entry, Adaptive Media works only with images added from the *Blog Images*, *Documents and Media*, and *Upload* tabs. Additionally, adapted images can only be applied to a blog entry's content—cover images excluded. Adaptive Media works for images added to a blog entry via drag and drop, as the image is automatically uploaded to the Blog Images repository, adapted, and then included in the blog entry's content. You can see this by inspecting the HTML and checking that the image contains the `` tag and `data-fileentryid` attribute.

For web content articles, Adaptive Media works only with images added from the file selector's *Documents and Media* tab. Unlike blogs, Adaptive Media doesn't deliver adapted images for images added to web content articles via drag and drop.

For both blog entries and media content articles, Adaptive Media doesn't work with images

added from the file selector's *URL* tab. This is because the image is linked directly from the URL and therefore provides no image file for Adaptive Media to copy.

Note that you can see the `` tag and `data-fileentryid` attribute in the HTML of a blog entry or a web content article while you're writing it. When the content is displayed, the HTML is automatically replaced and looks similar to this:

```
<picture>
  <source media="(max-width:850px)"
  srcset="/o/adaptive-media/image/44147/med/photo.jpeg">
  <source media="(max-width:1200px) and (min-width:850px)"
  srcset="/o/adaptive-media/image/44147/hd/photo.jpeg">
  <source media="(max-width:2000px) and (min-width:1200px)"
  srcset="/o/adaptive-media/image/44147/ultra-hd/photo.jpeg">
  
</picture>
```

This example uses three different images, each with a different resolution. A source tag defines each of these images. Also note the original image (`img`) is used as a fallback in case the adapted images aren't available.

195.2 Using Adapted Images in Structured Web Content

To use adapted images in structured web content, content creators must manually include an image field in the web content's structure. Then they can reference that image field in the matching template by selecting it on the left side of the editor. Here's an example snippet of an image field named `Imagecrrf` included in a template:

```
<#if Imagecrrf.getData()?? && Imagecrrf.getData() != "">
  
</#if>
```

This snippet includes the `data-fileentryid` attribute to ensure that Adaptive Media replaces the image with an adapted image. If you inspect the resulting web content's HTML in the editor's code view, you should see a tag like this:

```

```

Note the `` tag with a `data-fileentryid` attribute. Adaptive Media uses the file entry ID to replace the `` element automatically with a `<picture>` element that contains the available adapted images for each resolution (see the `<picture>` example above).

195.3 Staging Adapted Images

Adaptive Media is fully integrated with Liferay DXP's content staging and export/import functionality. Adaptive Media includes adapted images in staged content when published, and can update those images to match any new resolutions.

Similarly, when content that contains adapted images is exported, Adaptive Media exports those images in the LAR file. That LAR file can then be imported to restore or transfer that content, along with its adapted images.

Adaptive Media doesn't regenerate adapted images during export/import or the publication of staged content. To improve performance, Adaptive Media instead reuses the existing adapted images.

Awesome! Now you know how create content that contains adapted images. You also know how Adaptive Media includes adapted images in the content's HTML.

MIGRATING DOCUMENTS AND MEDIA THUMBNAILS TO ADAPTIVE MEDIA

Liferay DXP automatically generates thumbnails for images in Documents and Media. Once you deploy the Adaptive Media app, however, Liferay DXP doesn't display thumbnails until you migrate them to Adaptive Media. This article walks you through this migration process.

Note: You must be a Portal Administrator to perform the actions described here.

You'll get started by creating image resolutions for the thumbnails in Adaptive Media.

196.1 Adding the Replacement Image Resolutions

To migrate the existing Documents and Media thumbnails, you must add new image resolutions in Adaptive Media that have maximum height and maximum width values that match the values specified in the following portal properties:

`dl.file.entry.thumbnail.max.height`

`dl.file.entry.thumbnail.max.width`

`dl.file.entry.thumbnail.custom1.max.height`

`dl.file.entry.thumbnail.custom1.max.width`

`dl.file.entry.thumbnail.custom2.max.height`

`dl.file.entry.thumbnail.custom2.max.width`

Note: Some of these properties may not be enabled. You need only create image resolutions in Adaptive Media for the enabled properties.

To create the new Image Resolutions, follow the instructions found in the Adding Image Resolutions section of the Adaptive Media user guide.

Now you're ready to create the Adaptive Media images.

196.2 Creating the Adaptive Media Images

Once the required image resolutions exist, you can convert the Documents and Media thumbnails to Adaptive Media images. As mentioned in the Adaptive Media installation guide, the module *Liferay Adaptive Media Document Library Thumbnails* (which is included in the Adaptive Media app) enables this functionality.

There are two different ways to migrate the Documents and Media thumbnails to Adaptive Media:

Adapt the images for the thumbnail image resolution: This scales the existing thumbnails to the values in the Adaptive Media image resolutions, which can take time depending on the number of images. We only recommend this approach when there isn't a large number of thumbnails to process, or if you prefer to generate your images from scratch. This approach is covered in more detail in *Generating Missing Adapted Images*.

Execute a migrate process that reuses the existing thumbnails: This copies the existing thumbnails to Adaptive Media, which performs better because it avoids the computationally expensive scaling operation. The next section describes the steps to run this process.

196.3 Running the Migration Process

The migration process is a set of Gogo console commands. You can learn more about using the Gogo console in the *Felix Gogo Shell* article.

Follow these steps to migrate your thumbnails from the Gogo console:

1. Run the `thumbnails:check` command. For each instance, this lists how many thumbnails are pending migration.
2. Run the `thumbnails:migrate` command. This executes the migration process, which may take a while to finish depending on the number of images.
3. Run the `thumbnails:cleanUp` command. This deletes all the original Documents and Media thumbnails and updates the count returned by `thumbnails:check`. Therefore, you should **only** run `thumbnails:cleanUp` after running the migrate command and ensuring that the migration ran successfully and no images are pending migration.

Note: If you undeploy Adaptive Media at some point after running the migration process, you must regenerate the Documents and Media thumbnails. To do this, navigate to *Control Panel* →

Configuration → *Server Administration* and click *Execute* next to *Reset preview and thumbnail files for Documents and Media*.

Great! Now you know how to migrate your Documents and Media thumbnails to adapted images.

ADVANCED CONFIGURATION OPTIONS

Adaptive Media's advanced configuration options are available in System Settings. Open the Control Panel and go to *Configuration* → *System Settings*, then select *Adaptive Media*. There are two configurations listed under *SYSTEM SCOPE*:

- Images
- Processes

The *Images* configuration contains the following options:

Supported MIME Types: A list of the image MIME types that Adaptive Media supports. If an image is uploaded and its MIME type isn't in this list, Adaptive Media ignores the image. By default, this list contains many common MIME types.

Gifsicle: To scale animated GIFs, Adaptive Media uses an external tool called Gifsicle. First install Gifsicle on the server, ensure that it's on the PATH environment variable, and then click the box next to *Gifsicle Enabled*. If Gifsicle isn't installed and *image/gif* is included as a supported MIME type, Adaptive Media scales only one frame of the GIF, making a static GIF.

Max Image Size: Maximum size of the source images that Adaptive Media can use to generate adapted images. Adaptive Media will not generate adapted images for source images larger than this setting. The default value is 10 MB. To generate adapted images for all source images regardless of size, set this to -1. Since generating adapted images from large source images requires significant amounts of memory, you can specify a lower *Max Image Size* to avoid out of memory errors.

The *Processes* configuration is related to Adaptive Media's asynchronous processing. These values can be modified to improve performance for specific scenarios or use cases. The following options are available:

Max Processes: The maximum number of processes for generating adapted media. The default value is 5.

Core Processes: The number of processes always available for generating adapted media. The default value is 2. This setting can't exceed the *Max processes* setting.

Warning: Larger values for Max Processes and Core Processes may cause out of memory errors, as processing more images at once can consume large amounts of memory. Out of memory errors can also occur if the source images Adaptive Media uses to generate adapted images are large. You can restrict the maximum size of such images via the *Max Image Size* setting in the *Adaptive Media*

Gifsicle Enabled

Adaptive Media uses Gifsicle to scale animated GIF images. Enable this setting if Gifsicle has already been installed in the server. See <https://www.lcdf.org/gifsicle> for more information.

Max Image Size

10485760

Set the maximum size of the source images that Adaptive Media can use to generate adapted images.

Save

Cancel

Figure 197.1: You can configure Gifsicle and the maximum image size for Adaptive Media.

Image configuration, which is described next. You should run performance tests to optimize these settings for the amount of memory available on your system.

Processes

This configuration was not saved yet. The values shown are the default.

Max Processes

5

The maximum number of processes for generating adapted media.

Core Processes

2

The number of processes always available for generating adapted media. This setting cannot exceed the max processes setting.

Save Cancel

Figure 197.2: You can also configure Adaptive Media's image processing resources.

COLLABORATING


Your users can leverage a robust suite of collaboration apps to get things done and form extensive communities. These apps provide all the features that you would expect of standalone apps, but these apps are integrated: they share a common look and feel, security model, and architecture because of Liferay's development framework. You can use them in combination with user registration and content management features to build a well-integrated, feature-rich site.


This guide shows you how to administer and use the collaboration apps:

- Blogs
- Message Boards
- Wiki
- Alerts and Announcements
- Knowledge Base
- Bookmarks

BLOGS [RSS](#) [Subscribe](#) [New Entry](#)

Lunar Spelunking

 **Joe Bloggs**
4 Hours Ago




Lunar spelunking is a lot of fun, but it's not for the faint of heart. Nor is spelunking on Earth. A low-gravity environment, however, increases the difficulty. And although the science on this is unsettled (or perhaps just unsettling), there may be strange, fanged creatures dwelling beneath the lunar surface. An early member of the Apollo missions was said to have almost caught one, but its...

[Comment \(0\)](#) [👍 0](#) [👁️ 0](#) [Twitter](#) [Facebook](#) [LinkedIn](#) [Share](#)

Figure 198.1: This blog entry looks fascinating.

How can I improve my lunar golf game?


Joe Bloggs, modified 4 Hours ago. P 👍 0 🗨️ 0 ⋮

 [How can I improve my lunar golf game?](#)

Youngling Posts: 2 Join Date: 5/30/18 [🔍 Recent Posts](#)

Hey everyone. I thought lunar golf would be easier, seeing as there's less gravity and all, but I can't seem to make solid contact with the ball. The lunar surface is like a giant sand trap! Have any tips for this lunar duffer?


Jane Bloggs, modified 4 Hours ago. P 👍 0 🗨️ 0 ⋮

 [RE: How can I improve my lunar golf game?](#)

Youngling Posts: 2 Join Date: 5/30/18 [🔍 Recent Posts](#)

You may want to try a higher-lofted club than you would normally use. I find that this makes getting the ball up and flying much easier. Also, don't try to hit down on the ball like you would on Earth or another high-gravity planet. You want to try to "scoop" the ball up out of the moon dust, and let that low gravity do the work. And remember, don't swing too hard! Easy does it.

Figure 198.2: This is a great thread.

WIKI 

Lunar Resort [Main](#)

[FrontPage](#) [Recent Changes](#) [All Pages](#) [Orphan Pages](#) [Draft Pages](#)



FrontPage

[+ Add Child Page](#) [Edit](#) [Details](#) [Print](#)



Welcome to the Lunar Resort's wiki! Here, Lunar Resort guests can get together and write fun and informative content about things to do and see at the Lunar Resort.

6 Views

Your Rating [Average \(0 Votes\)](#)

0 COMMENTS [Subscribe](#)

[Reply](#)

Figure 198.3: The Wiki widget displays your wiki on a Site page.

PUBLISHING BLOGS


The Blogs app's editor has a complete set of WYSIWYG controls that appear when and where you need them. You can also switch to source mode to edit your content's HTML code. In source mode, you can work with light text on a dark background or dark text on a light background. You can even open a dual-screen HTML editor to see your code rendered in real time.


The Blogs app also contains a powerful set of tools for customizing how your blogs appear. For example, display templates like Abstract or Full Content let you choose how much of a blog post appears on a page. You can leverage the built-in display templates or create your own. You can also add a cover image to each of your blog entries. Let's face it—you might not be able to judge a book by its cover, but a blog post with a nice cover image is more likely to draw readers.

Read on to learn about Liferay's blogging platform.

BLOGS [RSS](#) [Subscribe](#) [New Entry](#)

Lunar Spelunking

 **Joe Bloggs**
4 Hours Ago





Lunar spelunking is a lot of fun, but it's not for the faint of heart. Nor is spelunking on Earth. A low-gravity environment, however, increases the difficulty. And although the science on this is unsettled (or perhaps just unsettling), there may be strange, fanged creatures dwelling beneath the lunar surface. An early member of the Apollo missions was said to have almost caught one, but its...

[Comment \(0\)](#) [👍 0](#) [👁️ 0](#) [Twitter](#) [Facebook](#) [LinkedIn](#) [Share](#)


Figure 199.1: This blog entry looks fascinating.

ADDING BLOG ENTRIES

Each site contains a built-in blog instance, so you can add blog entries to it right away. The easiest way to do this is in the Site Administration menu. Follow these steps to add a blog entry in Site Administration:

1. Click the Menu button () to open the product menu. Then expand the menu for your site and select *Content & Data* → *Blogs*. This takes you to the Blogs app for your site. The *Entries* tab is selected by default, which lists the site's blog entries.
2. Click the *Add* button () to create a new blog entry. This presents the blog entry editor. Note that the same form appears when editing a blog entry.
3. The first input field, *Drag & Drop to Upload*, is for optionally adding a cover image for your entry. By default, an Asset Publisher shows this cover image as part of the blog entry's abstract. You can insert any image you like in this field, either via drag and drop or the *Select File* button. The latter lets you choose an existing image in the blog, an image from Documents and Media, or an image that you upload from your machine.

If you select an image from Documents and Media, you can make changes to it with the Image Editor. Edits you make are applied automatically to a copy of the image, which you can then use as your cover photo.

Upon upload, the image appears in the pane. To center the image, drag it into place. You can also add a caption. If you want to select a different image, click the *Change* icon (). Clicking the trash can icon removes the image from the blog entry.

4. Enter a title for your blog entry. You can also add a subtitle if needed.
5. Enter your blog entry's content in the *Content* field. This field is small at first, but it expands as you add content. The editor displays the editing controls when you need them and hides them from view when you don't. When you select text in your blog post, for example, a bar with context-specific editing controls appears. This keeps your canvas uncluttered so you can focus on writing. You can also add images, videos, and tables in your blog entry's content. See *Using the Blog Entry Editor* for instructions on creating your blog entry's content.

Drag & Drop to Upload or

.gif, .jpeg, .jpg, .png Maximum Size 5MB

Title

Subtitle

Content *

| | |
|-----------------------|---|
| Categorization | > |
| Related Assets | > |
| Configuration | > |
| Display Page Template | > |
| Permissions | > |

Figure 200.1: This screenshot shows some of the blog entry editor's controls.

6. Expand the *Categorization* panel and associate your blog entry with tags and/or categories. Although this is optional, it improves search results for blog entries and gives your users more navigation options. For example, you can add the Tags Navigation app to another column on your blogs page, which lets users browse blog entries by tag.
7. Expand the *Related Assets* panel and choose any other content in your site that you want to associate with this blog entry. Although this is optional, related assets let you tie together content on your site. For example, you might want to write a blog entry about a discussion that happened on the forums. To link those two assets together, select the forum thread under Related Assets. For more information, see the related assets documentation.
8. Expand the *Configuration* panel if you want to customize your blog entry's URL, abstract, or display date. You can also set whether to allow pingbacks for your blog entry. For the URL, the default selection of *Automatic* generates the URL for you based on the blog entry's title. This URL appears in the *Blog Entry URL* text box. Selecting *Custom* lets you enter your own URL. Note that if you change the blog entry's URL after publishing the entry, the original URL redirects to the new URL.

You can also specify the blog entry's abstract. Enter a 400 character text-only abstract, or a custom abstract that contains a thumbnail image and a manually written description. The *Small Image* section lets you add a small image that appears when blog entries are displayed in list view. Below the abstract section, you can set the entry's display date and time.

Note that if you're editing an existing blog entry, the *Send Email Entry Updated* toggle appears. Setting this to *YES* sends an email to any subscribers when the blog entry is updated. You can customize this email when configuring the Blogs app.

Finally, you can allow *pingbacks* for the blog entry. Pingbacks are XML-RPC requests that are sent automatically when you link to another site. If you link to another site in your blog entry, Liferay DXP sends a pingback to the other site to notify that site that you linked to it. Similarly, if someone links to your blog entry, Liferay DXP can receive a pingback from that site and record the link.

9. Expand the *Display Page Template* panel if you want to select a display page template for displaying your blog entry. The following options are available:

Default Display Page Template: Use the default display page template.

Specific Display Page Template: Click the *Select* button to select the display page template you want to use.

No Display Page Template: Do not use a display page template.

10. Expand the *Permissions* panel to customize your blog entry's permissions. Use the *Viewable by* selector to set who can view the blog entry:

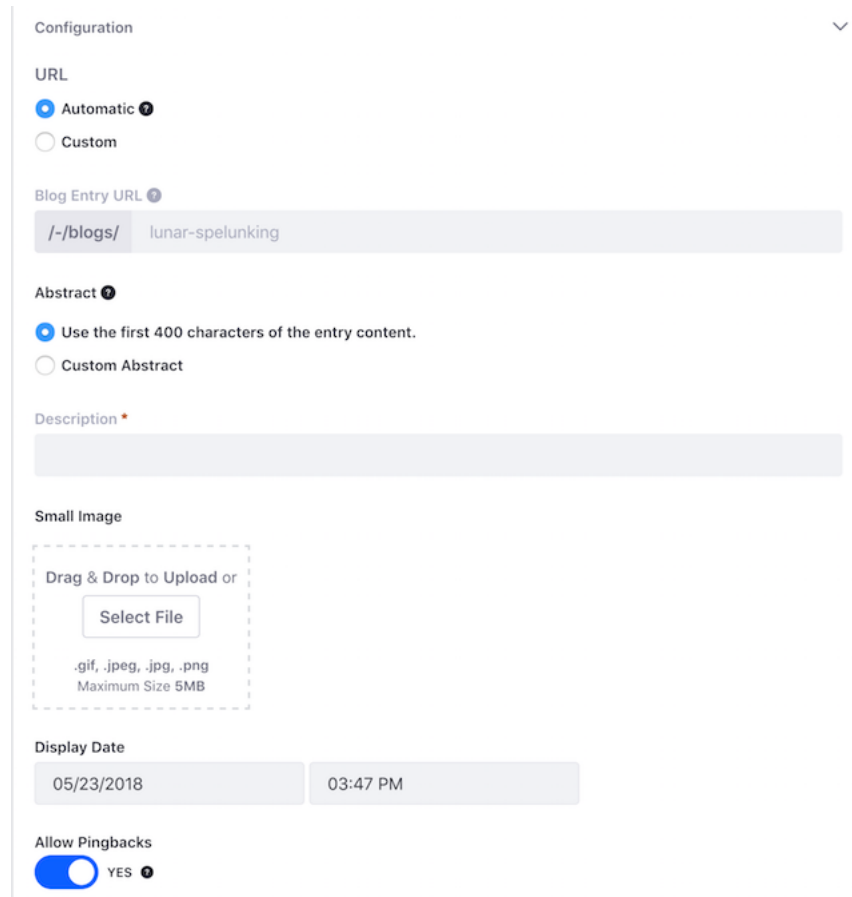
Anyone (Guest Role): Anyone, including guests, can view the entry.

Site members: Only site members can view the entry.

Owner: Only the entry's owner can view the entry.

Click the *More Options* link to bring up a permissions table that lets you grant or revoke the following permissions for guests and site members:

Update Discussion: Edit another user's comment on the blog entry.



The screenshot shows a 'Configuration' panel with the following settings:

- URL:** Automatic, Custom
- Blog Entry URL:**
- Abstract:** Use the first 400 characters of the entry content., Custom Abstract
- Description:**
- Small Image:** Drag & Drop to Upload or
Supported formats: .gif, .jpeg, .jpg, .png
Maximum Size 5MB
- Display Date:**
- Allow Pingbacks:** YES

Figure 200.2: When creating a blog entry, the Configuration panel lets you control when and where the blog entry appears, and what to use for the entry's abstract.

Delete: Move the blog entry to the Recycle Bin.

Permissions: View and modify the blog entry's permissions.

Delete Discussion: Delete any comments on the blog entry.

Update: Edit and modify the blog entry.

View: View the blog entry.

Add Discussion: Comment on the blog entry.

11. Click *Publish* to publish your blog entry. It now appears in the *Entries* tab.




Figure 200.3: The Blogs app in Site Administration lists the site's blog entries.

USING THE BLOG ENTRY EDITOR

When you create a new blog entry, you create its content with the blog entry editor. This editor is simple yet powerful. Its editing tools are context-aware, only appearing when you need them. They aren't visible while you're writing, which lets you focus on the task at hand. But when you select text, for example, a toolbar appears that contains buttons for styling the text. There are similar contextual options for adding and editing images, tables, and other types of content. And this is all in the editor's text view. You can switch to code view to edit the content's HTML. Regardless of which view you use, your entry is automatically saved as a draft every 25 seconds, so a browser crash or network interruption won't cause you to lose your work.

This guide shows you how to use this editor to create and edit blog entries.

Lunar Spelunking

Or: How I learned to stop worrying and love the dark

Normal **B** *i* U ☰ ☷ 🔗 🐦

+

Lunar spelunking is a lot of fun, but it's not for the faint of heart. Nor is spelunking on Earth. A low-gravity environment, however, increases the difficulty. And although the science on this is unsettled (or perhaps just unsettling), there may be strange, fanged creatures dwelling beneath the lunar surface. An early member of the Apollo missions was said to have almost caught one, but its slimy skin allowed it to escape. So why even bother with lunar spelunking? For the challenge, my dear friend! You can't let something that may or may not exist stop you from your endeavors. As for the existence of these creatures? Poppycock, I say! The astral cryptozoologists aren't going to keep me home!

Nonetheless, you need to come prepared for your excursion. You should bring these items:

| Item | Qty | Comments |
|-----------|----------|---|
| Duct Tape | 42 rolls | This stuff can save your life. I bet you could even construct an entire lunar rover from nothing but duct tape. |
| | | |

Figure 201.1: This screenshot shows some of the blog entry editor's controls.

201.1 Using the Editor's Text View

When you create or edit a blog entry, the editor is in text view by default. Text view is a WYSIWYG editor that lets you enter and edit text and other types of content. To enter text, place your cursor in any text field (e.g., Title, Subtitle, Content, etc.) and type or paste your text. Note that the Content area expands to fit its contents.

To style or format text in the Content area, first select the text. A toolbar appears above the text that contains the following options, as shown in the above screenshot:

Text Style: This selector menu, set to Normal by default, lets you choose the text's style. Normal is typical body text, but you can also select from different heading styles, alert or error message styles, code style, and more.

Typeface: Select bold, italic, or underline.

List Style: Select a numbered or bulleted list.

Link: Link the selected text to a specific URL, or to an item in the portal (e.g. a file in Documents and Media).

Twitter: Generates a link to tweet the selected text.


When you park your cursor in the entry's content area, the *Add* icon (+) appears. If you click this icon, it shows controls for inserting an image, video, table, or horizontal line (). Follow these instructions to insert each:

Image: Click the mountain icon, then select or upload an image in the image file selector screen that appears. Alternatively, you can drag-and-drop image files into the content area. You can also use the built-in Image Editor to apply simple edits to an image. Any edits you make are automatically applied to a copy of the image. After you add an image to the blog entry, clicking the image brings up controls for justifying it to the right or left side of the article.

Video: Click the play icon and insert the video's link. The video then appears in your content.


Table: Click the table (grid) icon and then choose the number of rows and columns in your table. When you click inside the table, table editing controls appear. They let you designate the first row and/or column as table headers. The controls also enable you to add rows, columns, and cells. As you type in a cell, the column width automatically adjusts to fit the content.


Line: Click the line icon. A simple, lightweight horizontal line then appears in your content. Such lines are good for separating sections of content in a large blog entry.

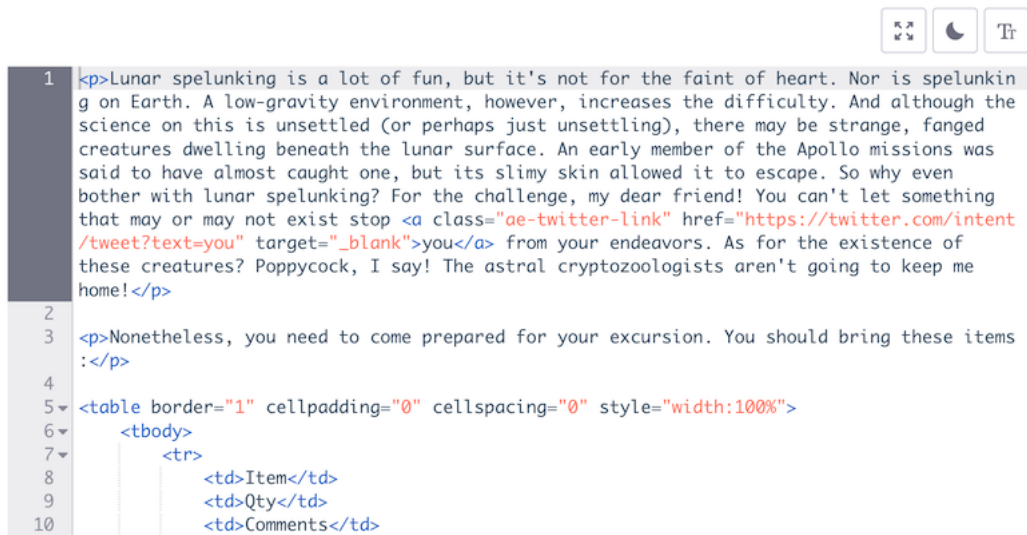
201.2 Using the Editor's Code View

To switch to code view, click the the *Source* icon (</>) that appears when you place your cursor in the Content area. The following buttons exist at the top-right of code view:

Text View (): Switch back to text view.

Dark/Light Theme (): Switch the code editor between dark and light theme.

Fullscreen (): Work in a dual-pane view that shows your HTML code on the left and a preview pane on the right. In this view, you can arrange the HTML and preview panes horizontally or vertically. You can also hide the preview pane so the HTML editor takes up the entire window space.



```
1 <p>Lunar spelunking is a lot of fun, but it's not for the faint of heart. Nor is spelunking on Earth. A low-gravity environment, however, increases the difficulty. And although the science on this is unsettled (or perhaps just unsettling), there may be strange, fanged creatures dwelling beneath the lunar surface. An early member of the Apollo missions was said to have almost caught one, but its slimy skin allowed it to escape. So why even bother with lunar spelunking? For the challenge, my dear friend! You can't let something that may or may not exist stop <a class="ae-twitter-link" href="https://twitter.com/intent/tweet?text=you" target="_blank">you</a> from your endeavors. As for the existence of these creatures? Poppycock, I say! The astral cryptozoologists aren't going to keep me home!</p>
2
3 <p>Nonetheless, you need to come prepared for your excursion. You should bring these items:</p>
4
5 <table border="1" cellpadding="0" cellspacing="0" style="width:100%">
6   <tbody>
7     <tr>
8       <td>Item</td>
9       <td>Qty</td>
10      <td>Comments</td>
```

Figure 201.2: Editing in code view lets you work with your blog entry's underlying HTML.

MANAGING BLOG ENTRIES

The Blogs app in Site Administration helps bloggers and blog administrators manage blog entries.

To access this app, click the Menu button (☰) to open the product menu, then expand the menu for your site and select *Content & Data* → *Blogs*. The *Entries* tab is selected by default, which lists the site's blog entries.

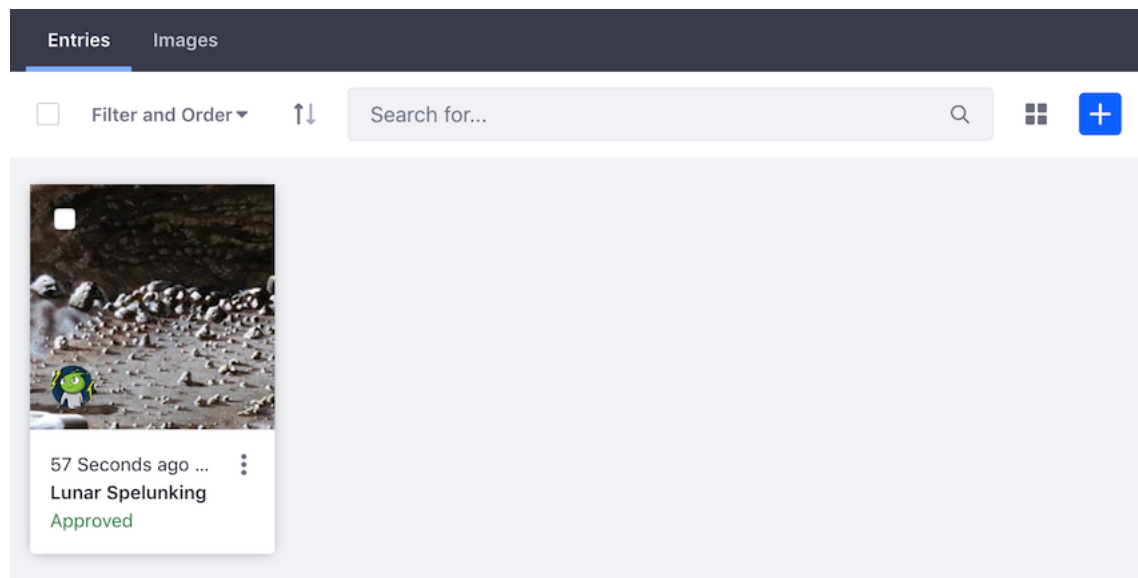


Figure 202.1: The Blogs app in Site Administration lists the site's blog entries.

You can use the Management Bar to manage your site's blog entries. If you've added blog entries via your site's Blogs app, then you're already familiar with the Management Bar's *Add* button (+). The sections that follow describe how to manage your blog entries.

- View Types
- Finding and Arranging Blog Entries
- Selecting Blog Entries
- Sharing Blog Entries

202.1 View Types

The *View Types* button is to the left of the Add button. It lets you choose how to display the blog entries in the Blogs app. The View Types button's icon depends on the selected view type:

Cards (☰): Shows a card-like rendering of the blog entry, with the author's profile picture. If the entry doesn't contain a cover image, a generic rendering of the entry is displayed. Each card also contains the entry's timestamp, title, workflow status (e.g., Approved, Draft, etc.), and an Actions menu (⋮).

List (☰): Shows the same information as the Cards view type, in a list with the author's profile picture instead of the blog entry's cover image.

Table (☰): Shows the same information as the other view types, in a list with no file renderings. Also, the blog entry information is in columns.

202.2 Finding and Arranging Blog Entries

The Management Bar also contains tools that help you locate and arrange blog entries. The most prominent of these tools is the *Search* bar, where you can find files by keywords.

To the left of the Search bar, the Sort button (↑↓) arranges entries in ascending or descending order.

You can also arrange entries via the *Filter and Order* selector using these criteria:

All: Shows all of the site's entries.

Mine: Shows only the current user's entries.

Display Date: Orders the entries by display date.

202.3 Selecting Blog Entries

The checkbox on the left-most side of the Management Bar selects all currently displayed blog entries. Selecting multiple entries lets you act on all of them at once. You can also select multiple entries individually by using the checkboxes for each. When you select one or more entries, the Management Bar changes to reflect the actions you can take on the selected entries.

Click the *Trash* button (🗑) to move the selected entries to the Recycle Bin. Unselect the checkbox to return the Management Bar to its normal view.

202.4 Sharing Blog Entries





Figure 202.2: With multiple blog entries selected, the management bar changes to reflect the actions you can take on the selected entries.

You can also share blog entries from the Blogs app in Site Administration. Sharing is enabled by default, as described in [Configuring Sharing](#). Sharing blog entries works the same as sharing files.

When sharing is enabled and you have sharing permission, you can share a blog entry by clicking its *Actions* menu () and selecting *Share*.

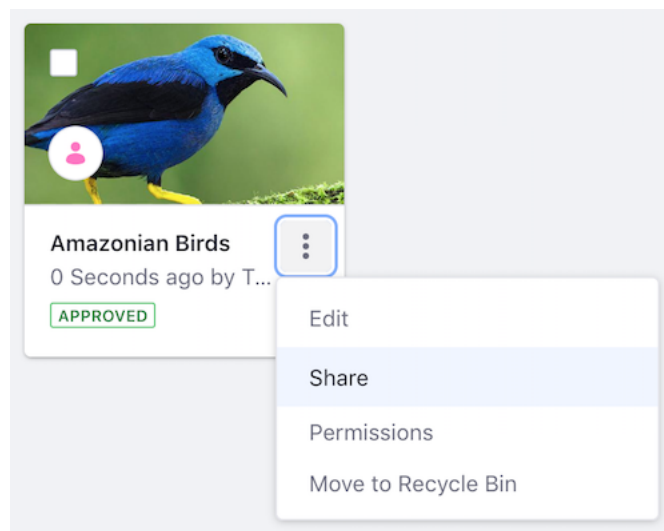


Figure 202.3: You can share a blog entry via its Actions menu.

CONFIGURING THE BLOGS APP

By configuring the Blogs app in Site Administration, you can control how the app behaves for all blogs in your site. To access this app, click the Menu button (☰) to open the product menu, then expand the menu for your site and select *Content* → *Blogs*. The *Options* menu (ⓘ) at the top-right of the Blogs app lets you configure permissions and notifications, or import/export the app's content.

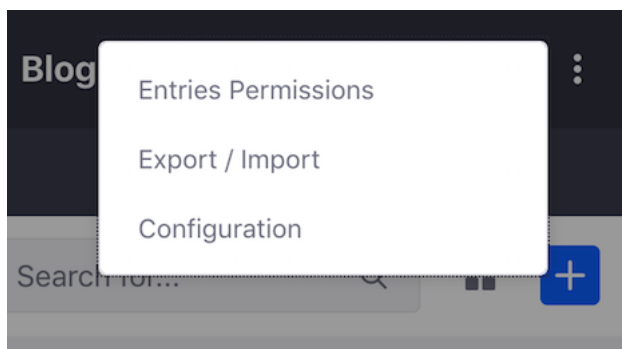


Figure 203.1: You can configure the options for your site's Blogs app.

Here are each of the options available in this menu:

Entries Permissions: Configure the permissions that can be applied to the Blogs app. You can control which roles can add an entry, configure entry permissions, and subscribe to entries.

Export/Import: Export or import a LAR file that contains the Blogs app's content.

Configuration: Configure the following options for the Blogs app, in these tabs:


- **Email From:** Define the *From* field in the email messages that users receive from Blogs.
- **Entry Added Email:** Define a subject and body for the emails sent when a new blog entry has been added.
- **Entry Updated Email:** Define a subject and body for the emails sent when a new blog entry has been updated.
- **RSS:** Lets you enable RSS subscription and choose how blogs are displayed to RSS readers. The *Maximum Items to Display* selector lets you choose the total number of RSS feed entries to display on the initial page. You can choose up to one hundred to be displayed. The *Display*

Style selector lets you choose between *Full Content*, *Abstract*, and *Title* for the entry display in the RSS feed. Lastly, the *Format* selector lets you choose which format the RSS feed uses to deliver the entries: Atom 1.0, RSS 1.0, or RSS 2.0.

DISPLAYING BLOGS


The Blogs app in Site Administration lets you add, manage, and configure your site's blogs. You can then display those blogs by adding a separate Blogs widget to a page. Adding the Blogs widget to a site page creates a shared blog for site members. Adding the widget to a user's personal site (dashboard) creates a blog just for that user. The widget works the same way in both cases. And of course, you can scope a blog to a page to produce a blog instance for just that page.

To add a Blogs widget to a page:

1. Navigate to the page.
2. From the *Add* menu () , open *Widgets* → *Collaboration*.
3. Drag a *Blogs* widget onto the page.

By default, the Blogs widget lists abstracts of the site's recent blog entries. The listing shows each entry's cover image prominently. Each abstract in the listing also shows the number of comments, thumbs up/down ratings, and links to share the entry on Twitter, Facebook, LinkedIn, and other social networking sites. Clicking a blog entry lets you view its full content, where you can also comment on the entry.

There are several display options that let you configure the listing to look the way you want.

To configure the widget, click the *Options* icon () in its title bar and select *Configuration*. The display settings are in the *Setup* tab:

Enable Ratings: Whether readers can rate blog entries.

Enable Comments: Whether readers can comment on blog entries.

Show View Count: Whether to show the number views for each entry.

Social Bookmarks: The social networking sites that users can share blog entries with. Only those in the *Current* column are displayed via the share buttons on each blog entry. To move social networking sites between the *Current* and *Available* columns, select the sites and use the arrows between those columns. Similarly, use the up/down arrows beneath the *Current* column to reorder the sites as they appear on each blog entry. For more information, see the social bookmarks documentation.

Display Style: The display style for social bookmarks. *Inline* is the default and displays the social bookmark icons in a row. *Menu* hides them inside a single share menu.

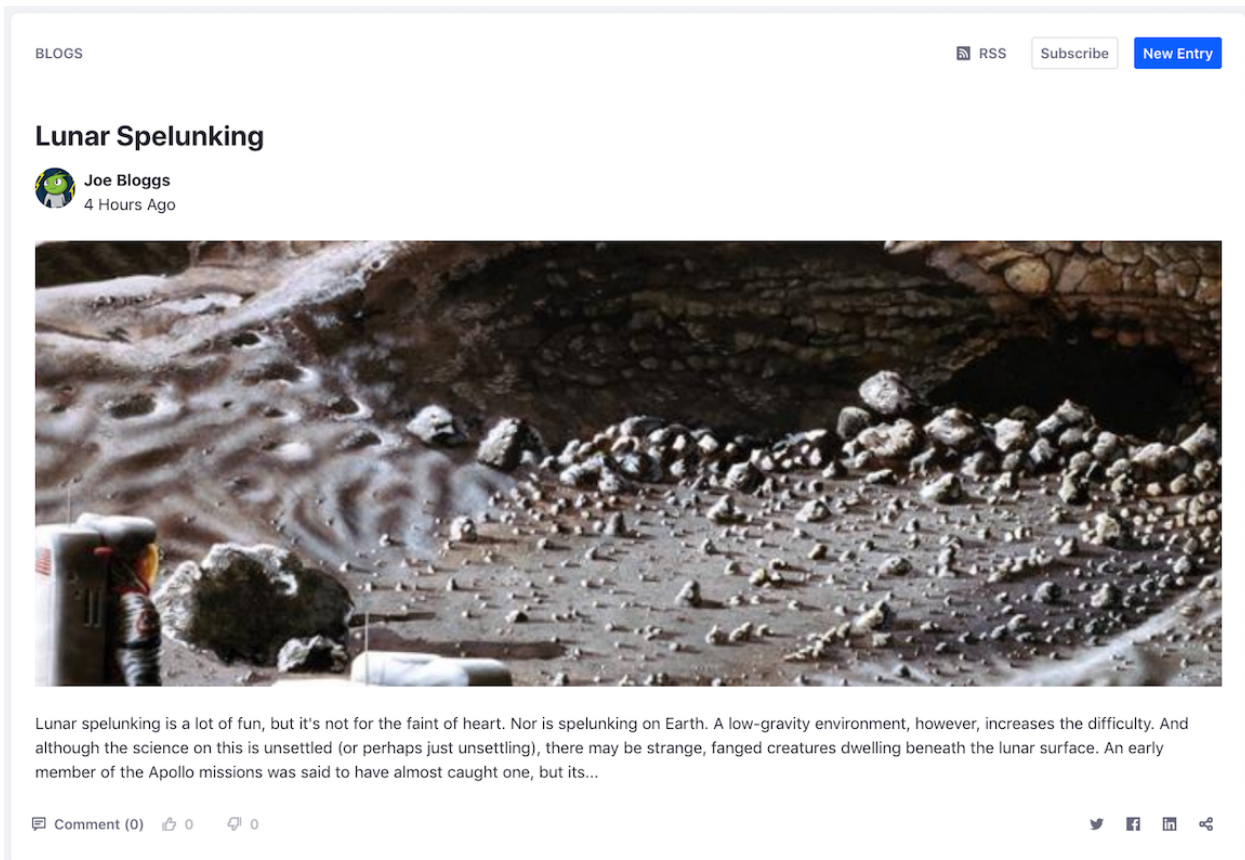


Figure 204.1: Fancy a lunar spelunking trip? This blog entry's abstract lets you know what you're getting into.

Maximum Items to Display: The total number of blog entries to display on the initial page. You can select up to 75 to display at once.

Display Template: The overall appearance of blog entries in the listing. *Abstract* is the default, and is shown in the above screenshot. You can also choose the following:

- **Full Content:** Displays each blog entry's full content.
- **Title:** Displays only the blog entry's title.
- **Basic:** A stripped-down version of the Abstract, with less text and no cover image.
- **Card:** Displays each blog entry in a card-like rectangle that shows the cover image, title, author, post date, and a few lines of text.

To select a different widget template or create your own, click *Manage Templates*. For more information, see the documentation on widget templates.

Enable Report Inappropriate Content: Whether to let users flag content as inappropriate, which sends an email to administrators.

Enable Ratings for Comments: Whether to let readers rate blog entry comments.

Show Related Assets: Whether to display related content from other apps/widgets in blog entries.

There are also other tabs in *Configuration*:

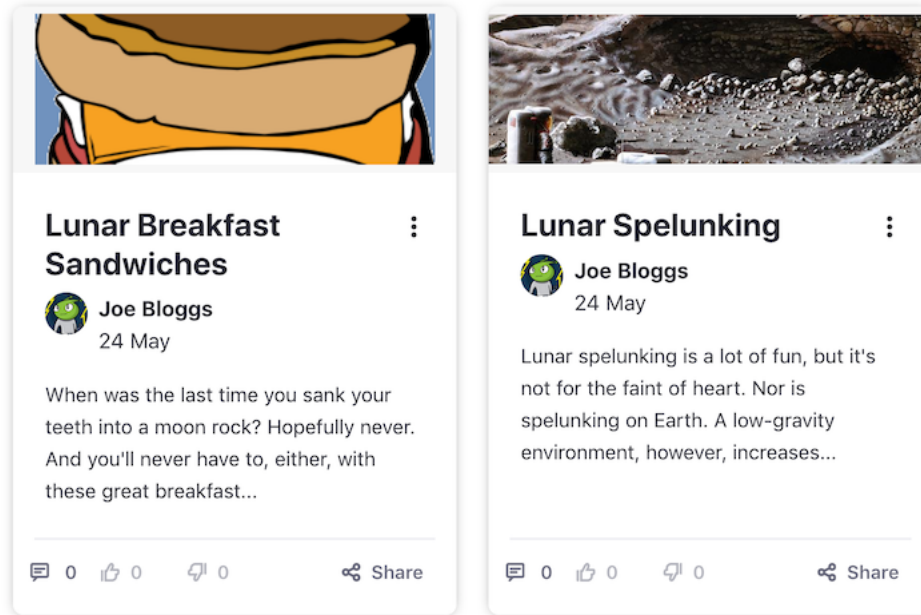


Figure 204.2: The Card display template makes your blog posts look like fun little trading cards.


Communication: Lists public render parameters the widget publishes to other widgets on the page. Other apps/widgets can read and take actions on these. For each shared parameter, you can specify whether to allow communication using the parameter and select which incoming parameter can populate it.


Sharing: Embed the widget instance as a widget on any website, Facebook, Netvibes, or as an OpenSocial Gadget.

Scope: Specify the blog instance the widget displays: the current site's blog (default), the global blog, or the page's blog. If the page doesn't already have a blog instance, you can select scope option *[Page Name] (Create New)* to create a page-scoped blog instance.

When you finish setting the options, click *Save* and then close the dialog box.

AGGREGATING BLOGS

You can set up a whole web site devoted just to blogging if you wish. The Blogs Aggregator lets you publish entries from multiple bloggers on one page, giving further visibility to blog entries. You can add it to a page from the *Collaboration* category in the *Add* () → *Widgets* menu.

If you click *Configuration* from the *Options* icon () in the widget's title bar, the Blogs Aggregator's configuration page appears. The *Setup* tab contains these options:

Selection Method: Set how the widget selects blogs for display. You can choose *Users* or *Scope*. If you select *Users*, the widget aggregates the entries of every blogger on your system. To refine the aggregation, you can select an organization by which to filter the users. If you select *Scope*, the widget aggregates the entries of users in the current scope. This limits the entries to members of the site where the widget resides.

Organization: The organization whose blogs you want to aggregate.

Display Style: Select the overall appearance for blog entries in the widget: *Body and Image*, *Body*, *Abstract* (default), *Abstract without Title*, *Quote*, *Quote without Title*, and *Title*.

Maximum Items to Display: The maximum number of entries the widget displays.

Enable RSS Subscription: Whether to enable an RSS feed of the aggregated entries. This lets users subscribe to an aggregate feed of all your bloggers. Below this option, you can configure how you want to display the RSS feed:

- **Maximum Items to Display:** The maximum number of RSS items to display.
- **Display Style:** The overall appearance of each entry in the RSS feed: *Abstract*, *Full Content*, or *Title*.
- **Format:** The language to use for your RSS feed: *Atom 1.0*, *RSS 1.0*, or *RSS 2.0*.
- **Show Tags:** Whether to display each entry's tags.

Here are descriptions for the other tabs in the Blogs Aggregator's configuration:

Sharing: Embed the widget instance as a widget on any website, Facebook, Netvibes, or as an OpenSocial Gadget.


Scope: Specify the blog instance the widget displays: the current site's blog (default), the global blog, or the page's blog. If the page doesn't already have a blog instance, you can select scope option *[Page Name] (Create New)* to create a page-scoped blog instance.

When you finish setting the options, click *Save* and then close the dialog box. You'll notice that the Blogs Aggregator looks very much like the Blogs widget, except that it shows entries from multiple blogs.

BLOGS AGGREGATOR

Joe Bloggs - 5/24/18

Lunar Breakfast Sandwiches




When was the last time you sank your teeth into a moon rock? Hopefully never. And you'll never have to, either, with these great breakfast sandwiches available at the Lunar Resort's Cafe!

0

Joe Bloggs - 5/24/18

Lunar Spelunking



Lunar spelunking is a lot of fun, but it's not for the faint of heart. Nor is spelunking on Earth. A low-gravity environment, however, increases the difficulty. And although the science on this is...

0

[RSS](#)

Figure 205.1: The Blogs Aggregator lets you display blog entries authored by multiple authors from different sites.

HIGHLIGHTING RECENT BLOGGERS

The Recent Bloggers widget lets you highlight the work of your site's most recent blog authors. This widget lists each recent author's name, profile picture, and number of posts. You can add the Recent Bloggers widget to a page from the *Collaboration* category in the *Add* (🛠️) → *Widgets* menu.

To access the widget's configuration options, click *Configuration* from the *Options* menu (⋮) in the widget's title bar. The *Setup* tab appears first:

Selection Method: Set how the widget selects blogs authors to highlight. You can choose *Users* or *Scope*. If you select *Users*, the widget aggregates every recent blogger on your system. To refine the aggregation, you can select an organization by which to filter the users. If you select *Scope*, the widget aggregates the recent bloggers in the current scope. This limits the entries to members of the site where the widget resides.

Organization: The organization whose recent bloggers you want to aggregate.

Display Style: Select how the widget displays recent bloggers: *User Name and Image*, or *User Name*.

Maximum Bloggers to Display: Select the maximum number of recent bloggers the widget displays.

Here are descriptions for the other tabs in the widget's configuration:

Sharing: Embed the widget instance as a widget on any website, Facebook, Netvibes, or as an OpenSocial Gadget.

Scope: Specify the blog instance the widget displays: the current site's blog (default), the global blog, or the page's blog. If the page doesn't already have a blog instance, you can select scope option *[Page Name] (Create New)* to create a page-scoped blog instance.

When you're finished setting the options, click *Save*. Then close the dialog box.

RECENT BLOGGERS




Figure 206.1: You can show off your site or organization's most recent bloggers from the Recent Bloggers app.

CREATING FORUMS WITH MESSAGE BOARDS

Although you're likely already familiar with what a modern forum can do, here's a sampling of what users and administrators can do with Liferay DXP's Message Boards app.

Users can:

- Start and reply to threads
- Mark a thread as a question and select an answer from the replies.
- Subscribe to threads
- Author posts in BBCode or with the standard WYSIWYG editor
- Assign thread priority (e.g., sticky, announcement, etc.)
- Attach files to a thread
- Rate a thread (e.g., like/dislike)
- And more

Administrators can:

- Organize threads into categories and subcategories
- Scope a message board to a page, a Site, or the entire portal.
- Publish threads via RSS
- Rank users by the number of messages they post and assign labels to these rankings (e.g., novice, legend, etc.)
- Create and modify thread priorities (e.g., sticky, announcement, etc.)
- Configure email notifications for thread activity
- And more

As you can see, there's something for everyone!

The Message Boards app also integrates with the rest of Liferay DXP's features. In many web sites, it's obvious that there's no link between the main Site and the message boards. In some cases, users are even required to register twice: once for the main Site and once for the message boards. Sometimes it's even three times: once for the Site, once for the message boards, and once for the shopping cart. By providing a message boards app along with all of the other apps and widgets, Liferay DXP provides a unique, integrated approach to building Sites. Administrators can concentrate on building their Site while the integration work is done for them.

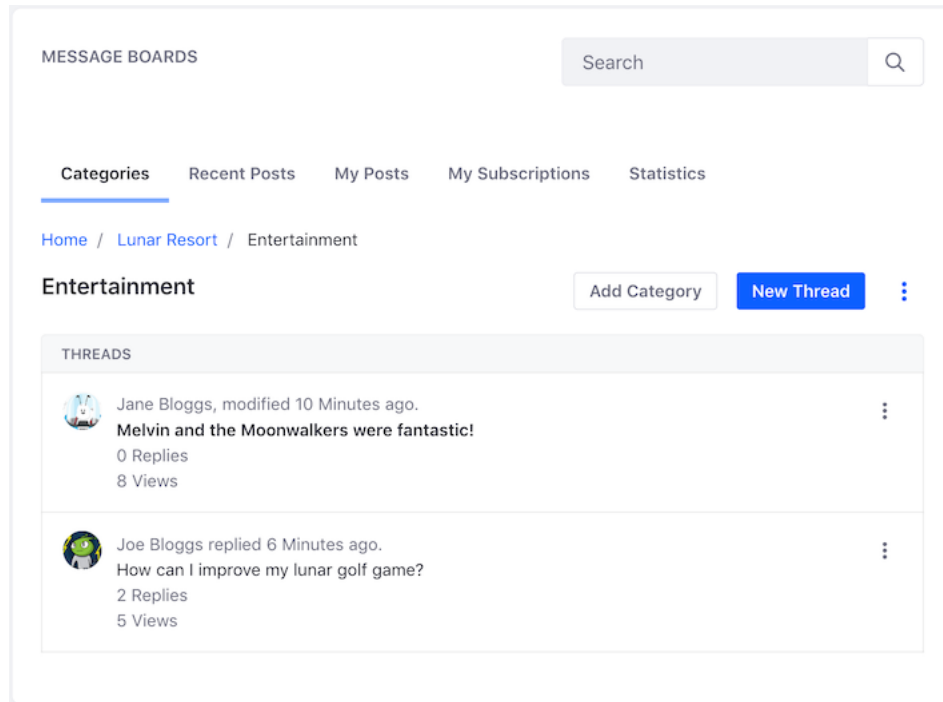


Figure 207.1: The Message Boards widget lets you explore its categories, interact with message threads, and post new messages.

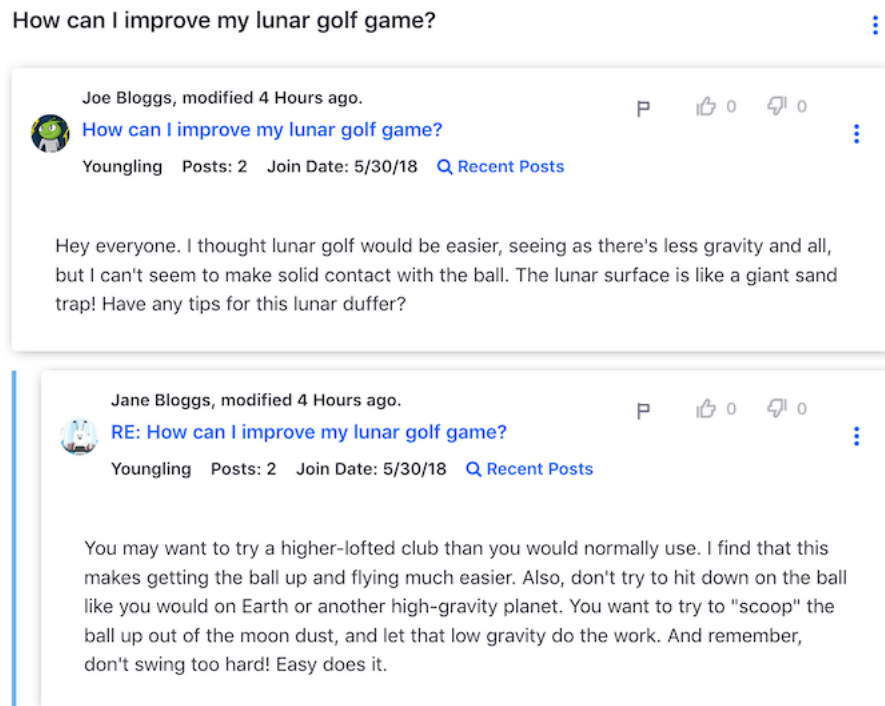



Figure 207.2: A thread's view displays author information and thread content, for the thread and all replies to the thread.

CREATING MESSAGE BOARDS



You can create and manage message boards in the Global, Site, and Page scopes. Regardless of scope, you manage a message board via the Site Administration menu. The following sections show you how to use this menu to manage message boards in each of these scopes.

208.1 Site-scoped Message Boards

By default, the Message Boards app in Site Administration is scoped to the current Site. To administer this message board, open the *Menu* () , expand the menu for your Site, then navigate to *Content & Data* → *Message Boards*. The Message Boards administration screen then appears. Note that the options available on this screen are the same regardless of scope. The next sections show you how to change scope and then access this screen.

208.2 Page-scoped Message Boards

If you need a page-scoped message board, you must add a Message Boards widget to that page and then set its scope to the page. Follow these steps:

1. Navigate to the page.
2. From the *Add* menu () , open *Widgets* → *Collaboration*.
3. Drag a *Message Boards* widget onto the page.
4. Click the *Options* icon () in the widget's title bar and select *Configuration*.

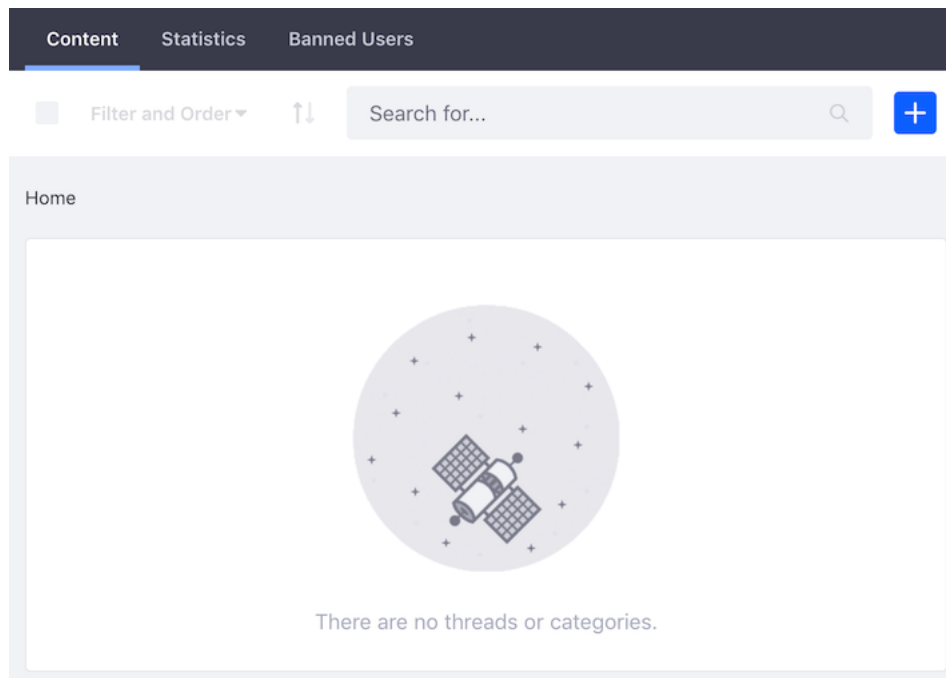




Figure 208.1: A Message Board instance starts empty, ready for you to configure for your purposes.

5. From the *Scope* menu in the *Scope* tab, select the page's name or *PageName (Create New)* if the page scope doesn't exist yet.
6. Click *Save*, and then close the dialog.

Note that you must still use the Site Administration menu to administer a page-scoped Message Boards widget. You do so by setting the Site Administration menu's active scope. Follow these steps to do this:

1. Open the *Menu* () , expand the menu for your Site, then expand *Content & Data*.
2. The current scope appears just below the *Content & Data* heading. *Default Scope* is the current Site. To change this, click the gear icon () and then select your desired scope. This changes the Site Administration menu to reflect scope you selected. To work in a page's scope, for example, select that page from the gear icon. That page's name then becomes the Site Administration menu's title.
3. Select *Message Boards* from the *Content & Data* menu. Any changes you make here apply to the scope that you selected in the previous step.

208.3 Globally-scoped Message Boards

To manage a message board in the global scope, follow these steps:

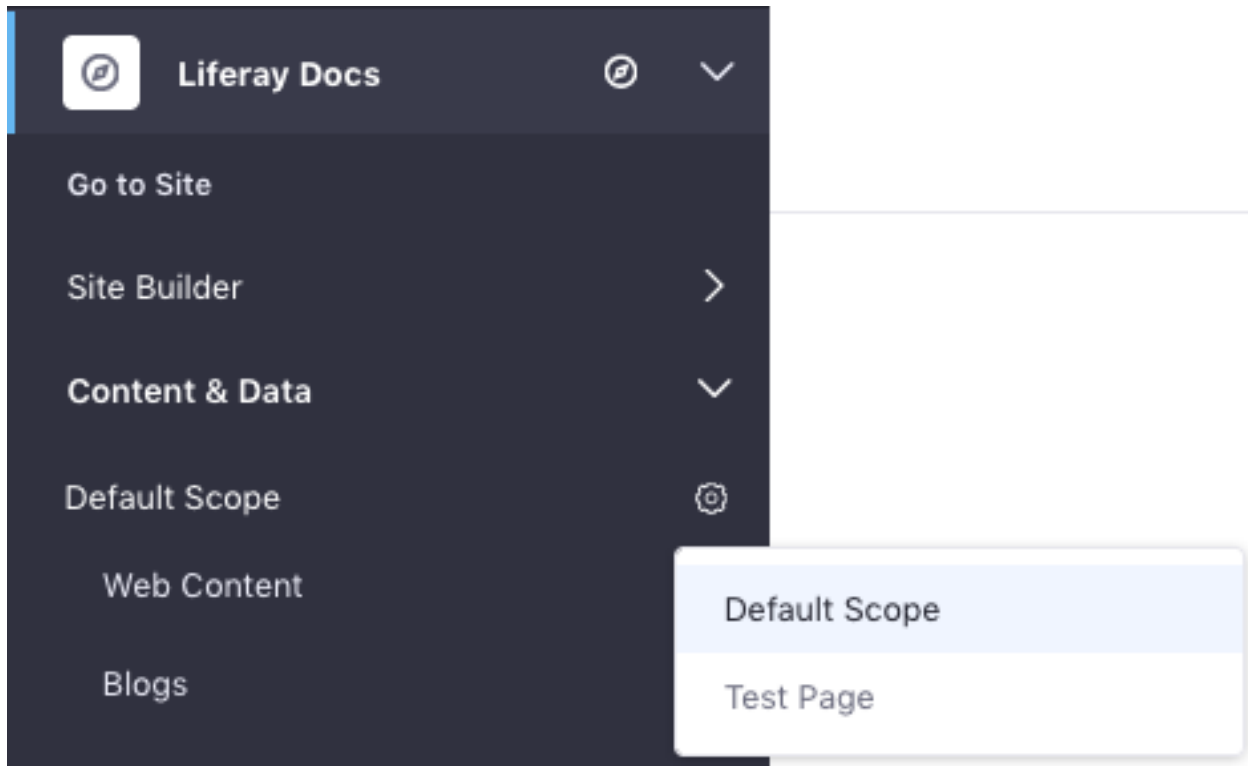


Figure 208.2: Select the page's scope under the *Content & Data* menu in Site Administration.

1. Open the *Menu* (☰), then click the compass icon (🧭) on the Site Administration menu. This opens the Select Site dialog.
2. Select the *My Sites* tab, then select *Global*. This closes the dialog and changes the Site Administration menu's title to *Global*.
3. Select *Message Boards* from the *Content & Data* menu. Any changes you make here apply to the global scope.




Figure 208.3: After changing to the global scope, select *Message Boards* from the *Content & Data* menu in Site Administration.

CONFIGURING MESSAGE BOARDS

Before using a message board, configure it to your needs. First, open the Message Boards app in your scope's Site Administration menu, as described earlier. To open the message board's configuration screen, click the message board's *Options* menu (⚙️) and select *Configuration*. The below sections cover these tabs.

209.1 General Setup

The *General* tab contains general settings:

Allow Anonymous Posting: Choose if users can post anonymously. Use this with caution—anonymous users tend to be mean.

Subscribe by Default: Choose if users are subscribed automatically to threads in which they've posted.

Message Format: Define the markup language of users' message board posts. You can choose BBCode or HTML. When creating posts, the type of WYSIWYG editor presented to users depends on which option is enabled. Both editors have a *Source* button that lets users view a message's underlying BBCode or HTML. Users can compose messages using either the WYSIWYG or Source view and can switch between views during message composition by clicking the *Source* button. For security reasons, BBCode is preferred.

Enable Report Inappropriate Content: Choose if users can report content as inappropriate. This sends a message to administrators so they can take action.

Enable Ratings: Choose if users can rate posts.

Thread as Question by Default: This automatically checks the *Mark as question* box in the new thread window. Threads marked as questions display *waiting for an answer*. Replies to the original message can be marked as an answer.

Show Recent Posts from Last: The *Recent Posts* tab shows posts from the following timeframes you define here:

- 24 hours

- 7 days (default)
- 30 days
- 365 days

After the allotted time has passed, the post expires from *Recent Posts*, but is still accessible everywhere else in the message board.

209.2 Email Setup

Use these tabs to configure how the Message Boards app handles email notifications:

Email From: The name and email address that sends email notifications. The default administrator account's name and email address. Default values, are from the `admin.email.from.name` and `admin.email.from.address` portal properties). These were set in the Basic Configuration Wizard when installing Liferay DXP. Make sure to update this email address to a valid one that can be dedicated to notifications.

HTML Format: Support HTML in these emails.

Definition of Terms: Shows variables you can use in the email templates you'll define next.

Message Added Email: Create a template for email users receive when a message is added to a topic they subscribe to.

- **Enabled:** Whether automatic emails are sent to subscribed users.
- **Subject:** Choose a prefix for the email's subject line. This lets users set up message filters in their email clients for these notifications.
- **Body:** The message body content. Use the variables defined in *Definition of Terms* to customize this content for users.
- **Definition of Terms:** Shows variables you can use in the email templates.

Message Updated Email: This tab is identical to the Message Added Email tab, except it defines the email that users receive when a post is updated.





209.3 Thread Priorities

The *Thread Priorities* tab defines custom priorities for message threads. This lets privileged Roles tag a thread with a certain priority, which highlights it for users. Three priorities are defined by default:

- Urgent
- Sticky
- Announcement

To define a thread priority, enter its name, a URL to its image icon, and a priority number. Threads with a higher priority are posted above threads with a lower priority.

Thread Icons

| Icon | Definition |
|---|--------------|
|  | Urgent |
|  | Announcement |
|  | Sticky |
|  | Question |

The localized language field lets you name the priorities in each locale. You can select the locale, update the priority names for it, and save your updates.

209.4 User Ranks

The User Ranks tab ranks users by the number of messages they have posted. Default ranks from 0 to 1000 are provided, but you can set custom ranks here as well.

You can also use this to define message boards labels that appear on user profiles. For example, you can use the message boards label *Moderator* for anyone who is a part of any of the Message Boards Administrator groups: the Site Role, the Organization, the Organization Role, the regular Role, or the User Group:

Moderator=organization:Message Boards Administrator

Moderator=organization-role:Message Boards Administrator

Moderator=regular-role:Message Boards Administrator

Moderator=site-role:Message Boards Administrator

Moderator=user-group:Message Boards Administrator

As with thread priority names, the *Localized Language* field localizes rank names.

209.5 RSS

Message board threads can be published as RSS feeds. The RSS tab enables/disables RSS subscriptions and defines how the feeds are generated:


Maximum Items to Display: The number of items to display in the feed.

Display Style: The feed's appearance. You can publish the full content, an abstract, or just the thread title.

Format: The feed's format: RSS 1.0, RSS 2.0, or Atom 1.0.

Once you've finished configuring your message board, make sure to *Save* your changes.

MESSAGE BOARD PERMISSIONS

Open the Message Boards app in your scope's Site Administration menu, as described in the article on creating message boards. Then click the *Options* icon () and select the *Home Category Permissions* option. This permissions screen is for granting and revoking access to message board functions.

The permissions enable a Role to perform the following actions:

Permissions: View and modify permissions.

Add File: Attach a file to a message.

Ban User: Forbid a user from participating in the message board.

Add Category: Add a new category to the message board.

Reply to Message: Respond to an existing message.

Lock Thread: Stop any further additions or modifications to a thread's messages.

Subscribe: Receive notifications on new and modified posts.


View: View all the contents of message threads.

Add Message: Post a new thread.

Move Thread: Move a thread to a different category or subcategory.

Update Thread Priority: Modify a thread's priority.

Configure the Roles with the permissions you want and *Save* your changes.

After adding a Message Boards widget to a page, you can access that widget instance's general permissions. To do so, select the widget's *Options* menu () and select *Permissions*. This permissions screen lets you control access to the widget instance's Permissions, Preferences, and Configuration menus.


MESSAGE BOARD CATEGORIES

Message Board categories organize threads by topic. This makes it easier to find the right topic for discussion, and can also help discussions stay on topic. For example, a tropical fishkeeping message board may have separate categories for freshwater and saltwater topics.

This article shows you how to create and manage message board categories.

211.1 Adding Categories

Follow these steps to create a message board category:

1. Open the Message Boards app in your scope's Site Administration menu, as described in [Creating Message Boards](#).
2. Click the *Add* icon () and select *Category*. This opens the Add Category form.
3. Enter a name and description for the category.
4. Select the category's *Display Style*. This controls how threads in the category appear. By default, you can choose these display styles:

Default: Classic display style for general purpose discussions.

Question: Threads appear in a question and answer style.

You can create custom display styles and make them available for selection in this form. You must set the available display styles via the portal property `message.boards.category.display.styles`. Similarly, you can set the default display style in `message.boards.category.display.styles.default`.

5. Open the *Mailing List* section of the form and set the mailing list options you want. To enable a mailing list for the category, set the *Active* toggle to *YES*. To enable anonymous emails in the list, set the *Allow Anonymous Emails* toggle to *YES*. The default for both toggles is *NO*. For an explanation of these features, see the documentation on mailing lists for Message Boards.

The screenshot shows a form for creating a message board category. The form is divided into several sections:

- Name ***: A text input field containing "Lunar Resort".
- Description**: A text area containing "Category for discussions about the Lunar Resort."
- Display Style**: A dropdown menu set to "Default".
- Mailing List**: A section with two toggle switches, both set to "NO":
 - Active**: A toggle switch set to "NO".
 - Allow Anonymous Emails**: A toggle switch set to "NO".
- Permissions**: A section with a dropdown menu set to "Anyone (Guest Role)" and a link for "More Options" with a question mark icon.

Figure 211.1: You have several options to create a message board category for your needs.

- Open the *Permissions* section and set the category's permissions. The *Viewable by* selector lets you pick who can view the category:
 - Anyone (Guest Role)
 - Site Members
 - Owner

To show more permissions options, click *More Options*. A table appears with the rest of the category's permissions, which you can assign to the Guest and Site Member roles:

Delete: Remove the category.

Permissions: View and modify permissions.

Add File: Attach a file to any of your messages.

Reply to Message: Respond to existing messages.

Lock Thread: Stop any further additions or modifications to a thread's messages.

Update: Edit the category.

Subscribe: Receive notifications on new and modified posts.


View: View the category.

Add Message: Post a new thread.

Move Thread: Move a thread to a different category or subcategory.

Add Subcategory: Add a new category within this category.

Update Thread Priority: Modify a thread's priority.

Note that after creating a category, you can revisit its permission options by clicking the category's *Actions* icon () and selecting *Permissions*.

7. Click *Save* when you're finished. Your category now appears in the table.

As you add categories to a message board, they appear on the message board's home screen. The list displays the category names and the numbers of subcategories, threads, and posts in each one.

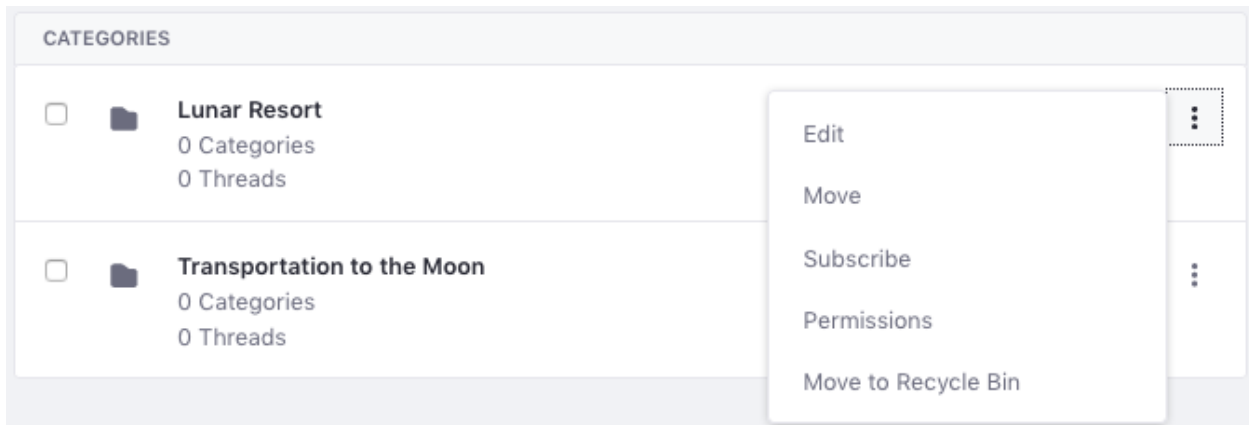



Figure 211.2: Categories help you organize threads so users can find topical threads that interest them.

211.2 Adding Subcategories

Categories can contain as many subcategories as you like. If, however, you nest categories too deep, users can have trouble finding them.


Follow these steps to add a subcategory to a category:

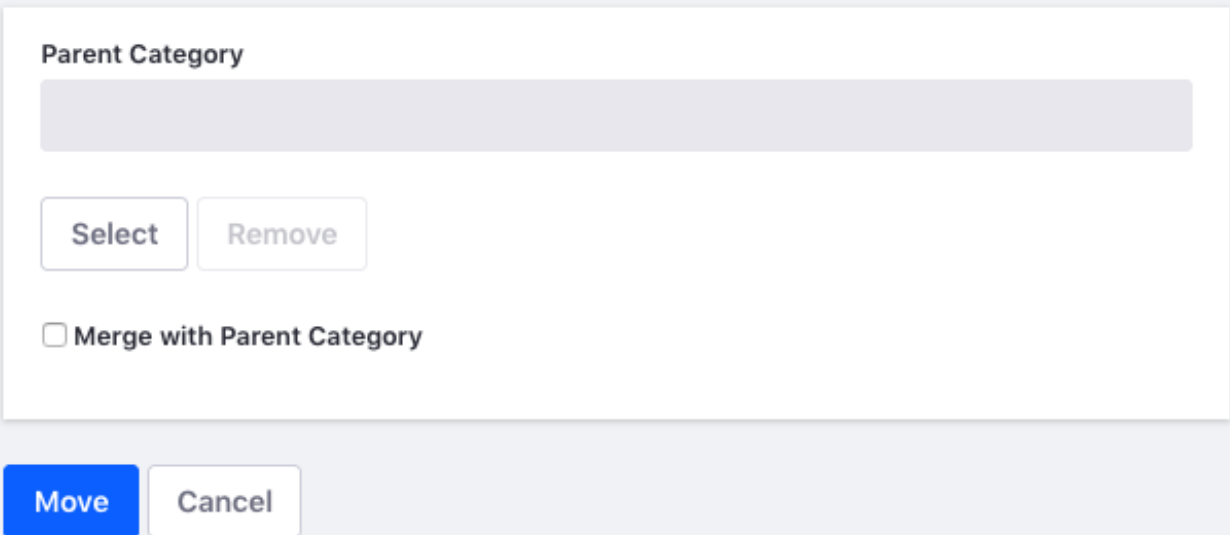
1. Click the category's name in the list, then click the *Add* icon () and select *Category*. The Add Category form appears.
2. Fill out the Add Category form with the values and settings you want to use for the subcategory. This form is populated with the parent category's properties by default.
3. Click *Save* when you're finished. Your subcategory now appears in the table.

211.3 Moving and Merging Categories

Each category can have any number of threads, and you can add as many categories and subcategories as you wish. You can also move and merge categories.

Follow these steps to move a category or merge it with another:

1. Click the category's *Actions* icon () and select *Move*. This brings up the Move Category form.
2. Select a new parent category via the *Select* button under the *Parent Category* field. Note that this field is empty for top-level categories.
3. If you want to merge the category with the selected parent category, select *Merge with Parent Category*.
4. Click *Move*.



Parent Category

Select Remove

Merge with Parent Category

Move Cancel

Figure 211.3: The Move Category form lets you move and merge categories.

USER SUBSCRIPTIONS AND MAILING LISTS

The Message Boards app notifies users of message boards activity via email in two ways:

- User subscriptions
- Mailing lists

Note: Since multiple sites can use a globally scoped message board, such message boards don't support user subscriptions or mailing lists. Make sure to use a site-scoped or page-scoped message board if you need user subscriptions or a mailing list.

212.1 User Subscriptions

In the user subscriptions mechanism, the Message Boards app uses its configured *Email From* address to send email notifications to subscribed users. The app can also import email replies to message board notifications directly into the message board. Then, users can interact on the message board via email without logging in to view the message board directly. This is disabled by default. To enable it, add the following line to your `portal-ext.properties` file:

```
pop.server.notifications.enabled=true
```

The user subscription mechanism uses the POP mail protocol. When the Message Boards app receives an email reply to a message board notification, it posts that reply to the message board and then deletes it from the mail server. Deleting the message from the mail server is the POP protocol's default behavior and the Message Boards app assumes that your POP mail server behaves this way. Most POP clients offer an option to leave mail on the mail server after it downloads, but you shouldn't exercise this option. If you configure mail to be left on the mail server, the Message Boards app sends copies of each retained message along with each new email notification it sends to subscribed users.

When enabling Message Boards to import replies to email notifications, you must decide whether to handle notifications with a mail server subdomain. By default, the following property setting is specified in the portal properties:

```
pop.server.subdomain=events
```

This property creates a special MX (mail exchange) subdomain to receive all virtual instance related email (e.g., events.liferay.com). If you don't want to use this approach, unset this value in a portal-ext.properties file:

```
pop.server.subdomain=
```

Doing so tells Message Boards to use the *Email From* address specified in the Message Board's configuration to receive message board notification email replies. For example, the *Email From* address could be set to *replies@liferay.com*.

If you're not using a mail subdomain, Message Boards parses the message headers of emails from the *Email From* address to determine the message board category and message ID. If you keep the pop.server.subdomain=events default, the email notification address takes the following form:

```
mb.[category_id][message_id]@events.liferay.com
```

In this case, Message Boards parses the email address to find the category and message ID. Parsing the email address is safer than parsing message headers, since different email clients treat message headers differently. This is why the events subdomain is enabled by default.

You can also configure the interval on which the POPNotificationListener runs. The value is set in one minute increments. The default setting is to check for new mail every minute, but you can set it to whatever you like:

```
pop.server.notifications.interval=1
```

Note: Depending on your mail provider, if you use multiple devices to access email through POP, you might need to configure in your POP settings something like Gmail's *recent mode*, which keeps the last 30 days of email available on the server. Then, more than just the first client can receive email. To enable recent mode in Gmail, for example, prefix the value of your POP client's Username or Email field with recent:.

If you don't use Gmail, IMAP may be a better solution for you.


212.2 Mailing Lists

Alternatively, the Message Boards app can use mailing lists to send email notifications. Any category in a message board can have its own mailing list. The mailing list mechanism, unlike the user subscription mechanism, supports both the POP and the IMAP protocols. POP is the default, but each message board's mailing list is configured independently. If you choose the IMAP protocol for a category's mailing list, make sure to configure the IMAP inbox to delete messages as they are

pulled by the email client that sends messages to the users on the mailing list. Otherwise, each email message retained on the server is sent to the mailing list each time there's a new post or update in the category.

When a mailing list is enabled for a message board category, Message Boards listens to the specific email inbox that's configured for the mailing list. Enabling the mailing list function lets users on the mailing list reply to the notification messages in their email clients. Message Boards pulls the messages from the email inbox it's configured to listen to and automatically copies those replies to the appropriate message board thread.

To enable the mailing list functionality for a category, follow these steps:

1. Set up a dedicated email address for the category.
2. Click the category's *Actions* icon () and select *Edit*.
3. In the *Mailing List* section of the form, set the *Active* slider to *YES*. Several options then appear. Fill these out as follows:
 - Email Address:** The email address of the account that receives the messages.
 - Protocol:** Select POP or IMAP.
 - Server Name:** Your mail server's host name.
 - Server Port:** The port on which your mail service is running.
 - Use a Secure Network Connection:** Whether to use an encrypted connection if your server supports it.
 - User Name:** The login name on the mail server.
 - Password:** The password for the account on the server.
 - Read Interval (Minutes):** How often to poll the server looking for new messages to post.
 - Email Address (Outgoing):** The email address originating messages from this category. If you want your users to be able to reply to the categories using email, this should be the same address as the incoming email address.
 - Use Custom Outgoing Server:** Use a different mail server than global default. Fields appear for configuring the server's name, port, user name, password, and secure connection.
4. If you want to let emails from anonymous users post to the message board category, set the *Allow Anonymous Emails* toggle to *YES*.
5. Click *Save* when you're finished.

USING THE MESSAGE BOARDS

You can add a Message Boards widget to a page from the *Add* (+) menu's *Widgets* → *Collaboration* section. The Message Boards interface is similar to other message boards that populate the Internet. In any case, it can't hurt to explore how to use Liferay DXP's Message Boards and discover its features.




Figure 213.1: The Message Boards widget lets you explore its categories, interact with message threads, and post new messages.

Threads can be viewed many ways. At the top of the widget is a set of tabs:

Categories: The message board's categories.

Recent Posts: Posts from all categories, sorted by date.

My Posts: The current user's posts.

My Subscriptions: Lets users view and manage their thread subscriptions.

Statistics: The number of categories, posts, participants, and a list of the top contributors. You can also access this from the same tab in Site Administration’s Message Board app.

You can also use the search bar at the top of the widget to search for threads and posts. Although search works on threads and posts within categories, it doesn’t work on categories themselves.

213.1 Posting New Threads

Follow these steps to post a new thread:

1. Click the Message Boards widget’s *New Thread* button. Alternatively, click the *Add* button (✚) and select *Thread* in the Message Boards app in Site Administration. Either way, the same *Add Message* form appears.

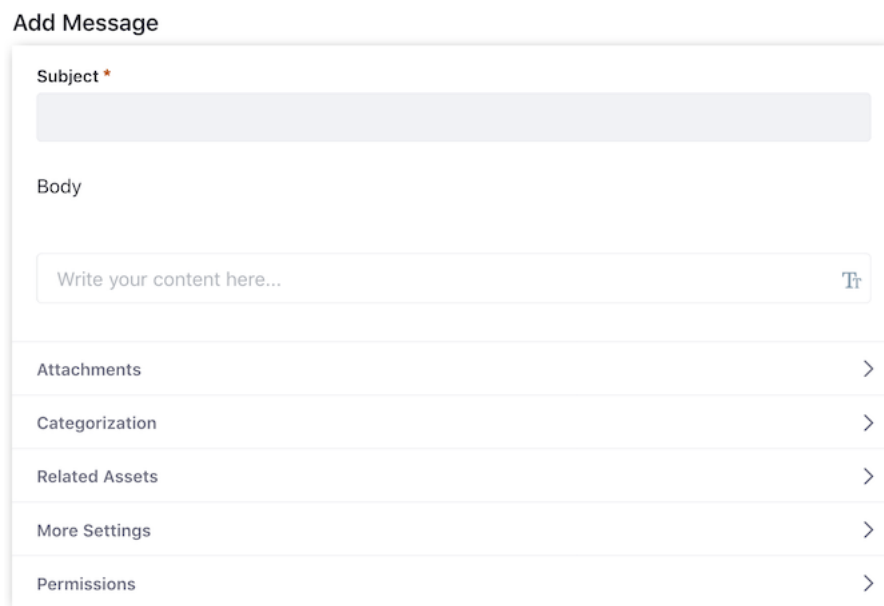


Figure 213.2: The Add Message form lets you create a new thread.

2. Give your thread a title in the *Subject* field.
3. Create your thread’s content in the *Body* field. This field uses the same editor as the Blogs app, except that it uses BBCode instead of HTML. For further instructions, see the documentation on using the editor. Also note that you can mention other users by entering the @ character and their user name.
4. If you want to add attachments, open the *Attachments* section and add them via drag and drop or the *Select Files* button.

5. If you want to associate a tag with the message, open the *Categorization* section and use the *Select* button to select an existing tag. You can also create a new tag by entering it in the *Tags* field and clicking *Add*. See the documentation on tags for more information.
6. If you want to select an existing asset in the portal (e.g., a media file, blog post, etc.) to relate to your thread, open the *Related Assets* section and use the *Select* button to select that asset.
7. Open the *More Settings* section and select the settings you want to use:
 - Mark as a Question:** Whether to mark this thread as a question. This lets you later select a post in the thread as the answer.
 - Anonymous:** Whether this thread is posted anonymously.
 - Subscribe Me:** Receive notifications for activity on the thread.
 - Priority:** The thread's priority in the Message Board. By default, you can choose *Urgent*, *Sticky*, or *Announcement*. Additional priorities can also be configured in the Message Boards app in Site Administration.
 - Allow Pingbacks:** Whether pingbacks are allowed for your thread.
8. Open the *Permissions* section and set the thread's permissions. Possible values in the *Viewable by* selector are
 - Anyone (Guest Role)
 - Site Members
 - Owner

You can also click the *More Options* link to select additional permissions:

- Delete:** Remove the thread.
- Permissions:** Grant/revoke thread permissions.
- Update:** Edit the thread.
- Subscribe:** Receive notifications for thread activity.
- View:** View the thread.

Note that you can revisit the thread's permissions after posting it. To do so, select the thread's Actions menu () and select *Permissions*.

9. Click *Publish* to publish the thread. Once it's published, it appears along with the other threads in the category.

213.2 Participating in Message Board Threads

To find message board threads that interest you, browse a message board's Categories or Recent Posts tabs. In the Categories tab, you can view a category's thread listing by clicking the category's name. Within a category, you can subscribe to an RSS feed and/or emails that inform you about

activity in that category. The Recent Posts tab also lists threads, except they're the latest threads across all categories.

Click a thread to view it. Messages appear in a threaded view so that replies are aligned under their parent thread. This makes it easy to follow conversations. Thread replies are indented under their parent thread.

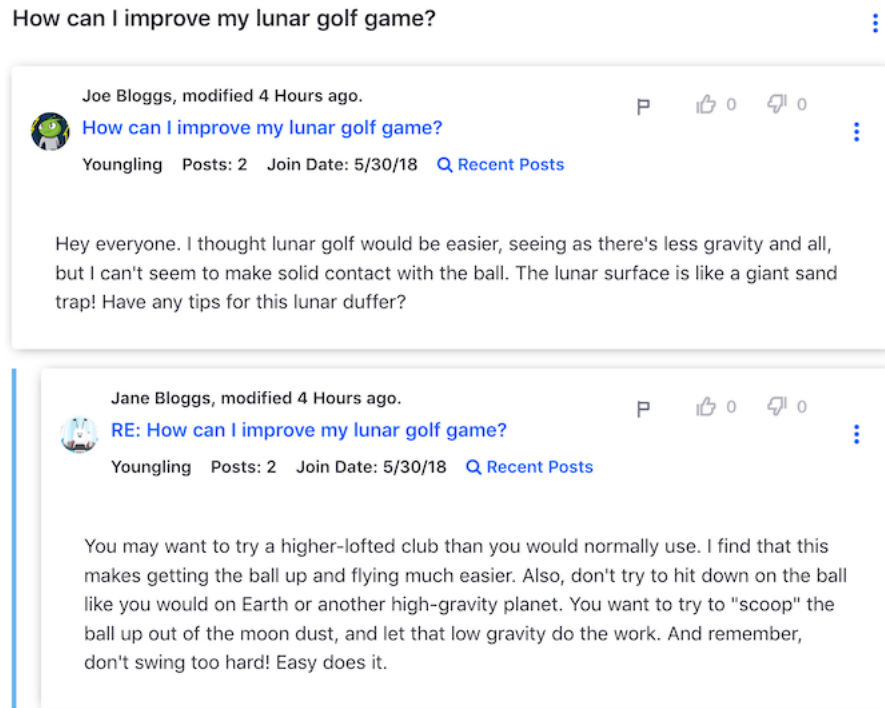


Figure 213.3: A thread's view displays author information and thread content, for the thread and all replies to the thread.

Subscribing to a thread causes Message Boards to send the user an email whenever a new message is posted to the thread. If you have enabled the mailing list feature for the thread's category, users can reply to these messages to post back to the thread without having to visit your site.

Most threads get more interesting as users reply to them. Follow these steps to reply to a message in a thread:

1. Click the *Reply* button. This opens the quick reply form, which only contains a text field for entering your reply.
2. Enter your reply in the text field. To access more options for your reply, click the *Advanced Reply* link. This opens the full editor from the add/edit thread form.
3. Click *Publish* to publish your reply.

In addition to replying to a message, you can rate it or flag it as objectionable. A message board moderator can evaluate flagged messages and decide how to handle the messages and their authors.


MANAGING MESSAGE BOARDS

Message boards can become unwieldy if left unmanaged. The Message Boards in Site Administration facilitates day-to-day thread administration. You may wish to assign this function to a Role that you give to one or more users. This frees you to concentrate on other areas of your Site. For example, you can create a Role called *Message Board Administrator* scoped to the portal (globally), an Organization, or a Site. Members of a global Role can administer Message Boards throughout the portal. Members of an Organization or Site-scoped Role can only administer Message Boards in that Organization or Site, respectively.

Follow these steps to create a global Role:

1. In the Control Panel, select *Users* → *Roles*.
2. Select or create the Role.
3. Select the Role's *Define Permissions* tab and navigate to *Site Administration* → *Content & Data* → *Message Boards*. A screen appears that lets you configure Message Board permissions.
4. Select the permissions you want message board administrators to have, then click *Save*.
5. Add users to this Role.

214.1 Locking Threads

You may encounter threads that you think should be preserved, but stopped. You can halt activity on a thread by selecting *Lock* from the thread's *Actions* menu ().

214.2 Moving Threads

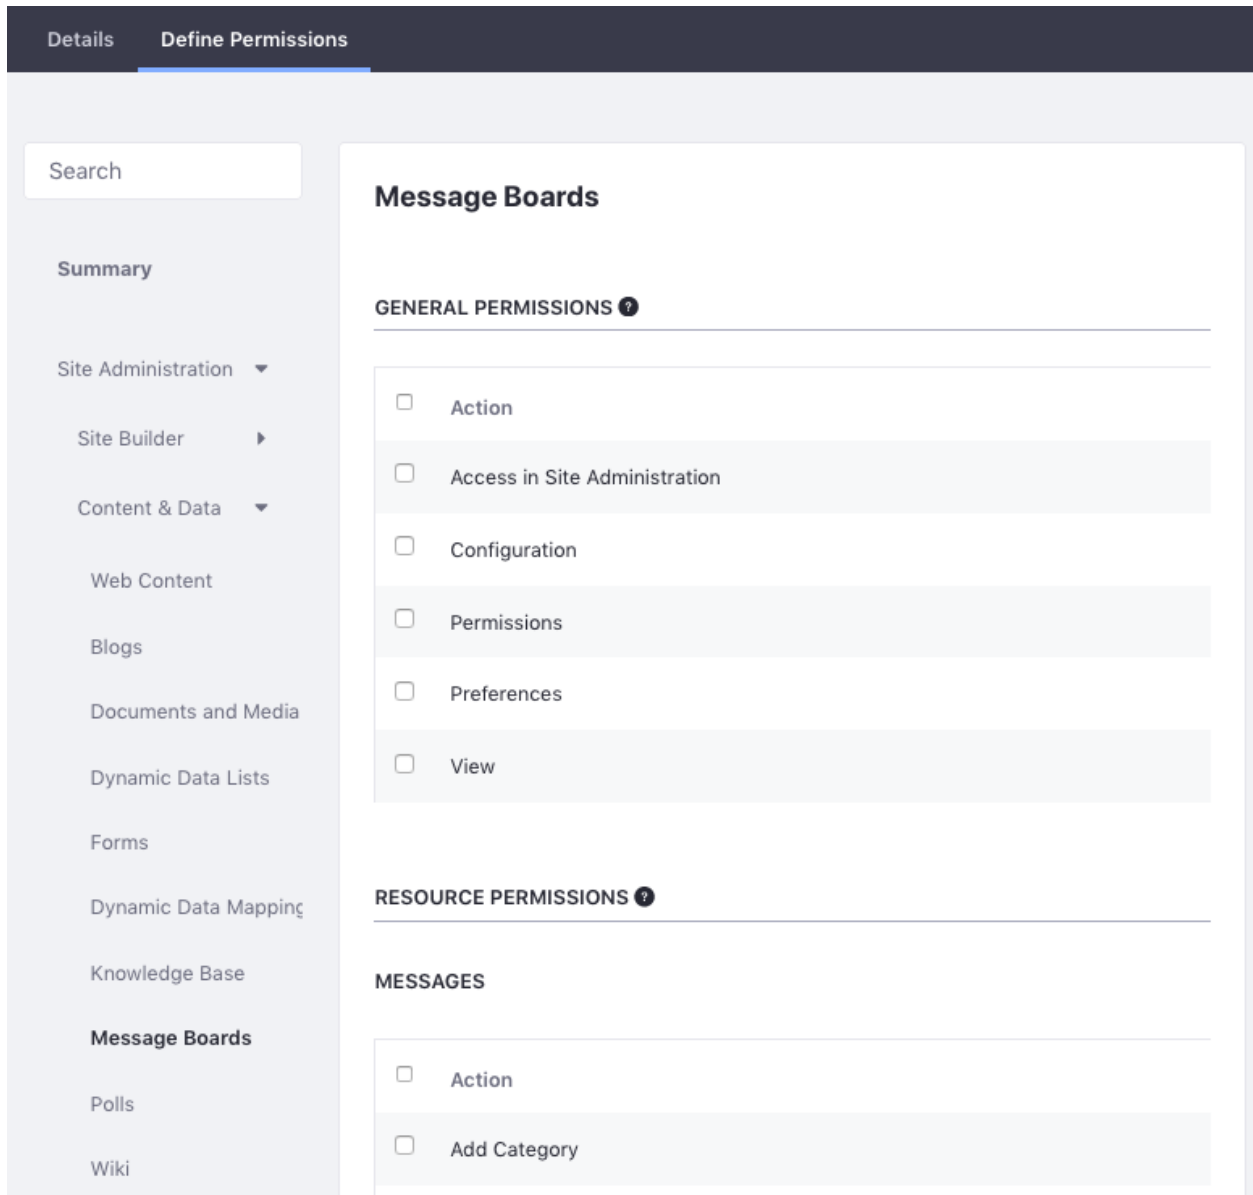




Figure 214.1: Define the permissions you want to use for the message boards administrators.


If someone posts a thread to the wrong category, you can move it to the proper one. Follow these steps:


1. Select *Move* from the thread's *Actions* menu (). This opens the *Move Thread* form.
2. Click the *Select* button and select the new category.
3. If you want to add a post explaining the move, select *Add explanation post*.
4. Click *Move* to move the thread.

214.3 Deleting Threads


Sometimes users begin discussing topics that are inappropriate or that reveal confidential information. In this case, administrators can delete the thread from the message boards. To do so, select *Move to Recycle Bin* from the thread's *Actions* menu ().

214.4 Banning Users

Unfortunately, message board users can be abusive. In this case, you can ban the user from the message boards. While viewing any of the user's posts in any thread, select the post's *Actions* menu () and select *Ban this User*.


To reinstate a banned user, you must use the Message Boards app in Site Administration. Navigate to this app and select the *Banned Users* tab. Select the user's *Actions* menu () and select *Unban this User*.

214.5 Splitting Threads


Sometimes a thread goes on for a while and the discussion completely changes into something else. In this case, you can split the thread where the discussion diverged and create a whole new thread for the new topic. To split a thread at a certain post, administrators can select that post's *Actions* menu () and select *Split Thread*. This brings up a form that lets you add an explanation post to the split thread. Click *OK* to split the thread.

214.6 Editing Posts

Administrative users can edit anyone's posts, not just their own. Sometimes users post links to copyrighted material or unsuitable pictures. By editing these posts, you can redact information that shouldn't be posted, or remove content not conforming to your terms of use. You can also update the thread's priority or mark a reply as an answer to a thread's question.

To edit a post, select its *Actions* menu () and select *Edit*.

214.7 Post Permissions

Permissions can be set not only on threads, but also on individual posts. You can choose to limit a particular conversation or post to only a select group of users. To do this, select the post's *Actions* menu () and select *Permissions*. You can then choose which Roles have the following permissions:

- Delete
- Permissions
- Subscribe
- Update
- View

Use this, for example, to let some privileged users post on a certain thread, while others are only allowed to view it. Other combinations of these permissions are also possible.

MENTIONING USERS

Have you ever wanted to include another user in a discussion on the Message Boards? Have you ever wanted to give kudos to a colleague in content you're writing? You can mention (notify and/or draw attention to) other users by entering the @ character in front of each user's user name.

When you mention a user, the user receives a site notification next to the user's profile icon and an email, alerting the user with a link to the content. You can mention users in a blog entry, a message boards post, or comments in any app that supports comments. A mention also links to the user's home page, so readers can find out more about that user.

Give me a break! They put me to sleep. I would've preferred something in the classic rock genre, like *Three Dog Night*; but I guess I'm partial, being a dog and all 🐕

@ray - What was your take on the band?



Figure 215.1: As you enter a user name after the @ character, links to users that match the text you enter are displayed. Select the user you want to mention and publish your content.

A selector appears after entering the @ character, listing users that match the name you're entering. In the selector, users are represented by their profile picture, name, and user name. Click the user you want to mention and finish editing your content.

On publishing the content, mentioned users receive a notification and an email informing them that they've been mentioned. The notification and email indicate the author's name and content type, and contain links to the content. You can access your notifications by selecting *Notifications* from your user menu.


Mentions are enabled by the Mentions app, which is a part of the Collaboration Suite. By default, the Mentions app is enabled globally. However, you can enable/disable it globally or per site. For a site to use Mentions, it must be enabled for the site's Virtual Instance.

To access the global Mentions settings for your Virtual Instance, first open the *Menu* () and



Figure 215.2: Your notifications are accessible from your user menu and appear in a list.

navigate to *Control Panel* → *Configuration* → *Instance Settings*. Then click *Community Tools* and select *Mentions* under *VIRTUAL INSTANCE SCOPE*. By default, all users are allowed to mention fellow site members and friends. To fine tune these options, select *Define Mentions Capability for Users* and specify the settings you want.

For Mentions to be available for a site, the app must be enabled for that site's Virtual Instance. Site administrators can enable or disable Mentions for a site. A site's Mentions app configuration is accessible from within the *Menu* () . Once in the menu, navigate to *[Site Name]* → *Configuration* → *Settings*. In the *Social* tab, expand the *Mentions* section. Enable or disable mentions via the toggle labeled *Allow Users to Mention Other Users*.

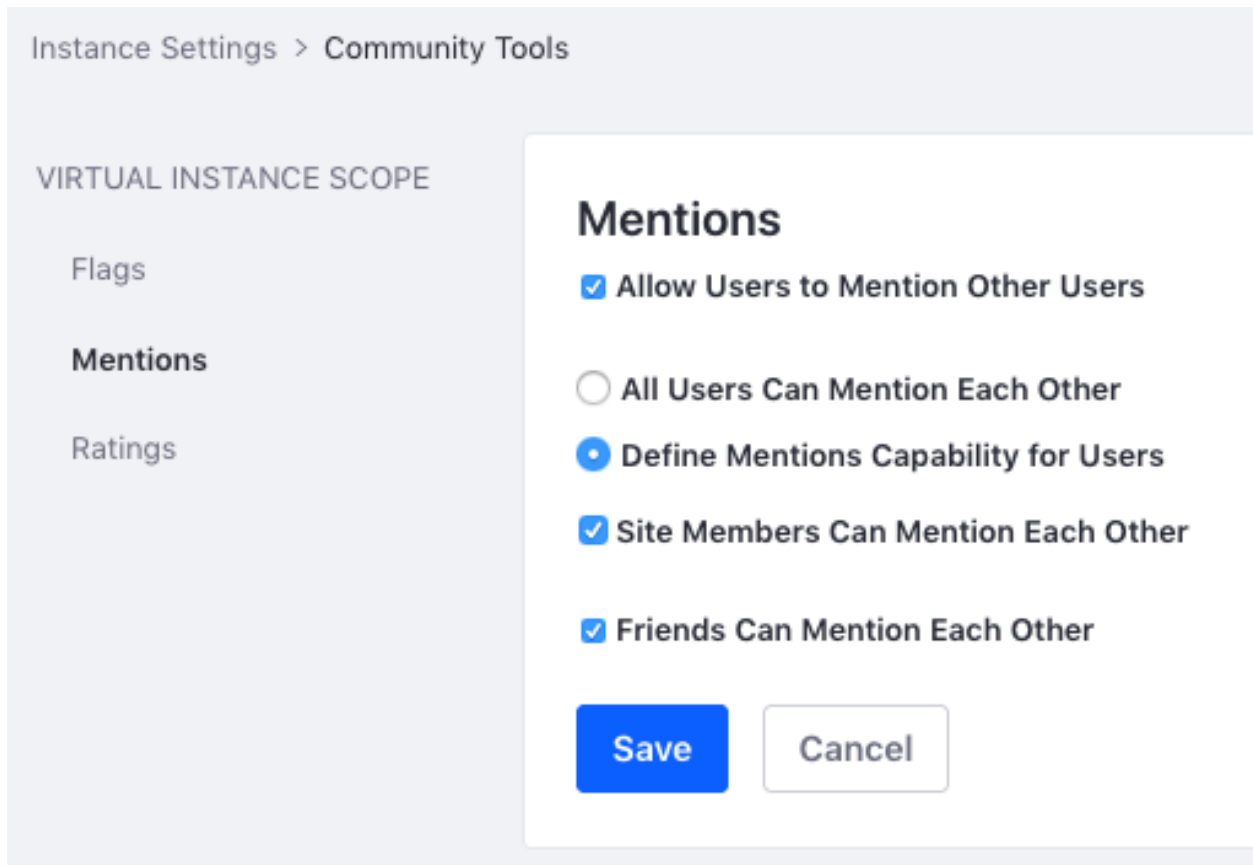


Figure 215.3: You can enable or disable the Mentions feature for all of the Virtual Instance's sites.

MENTIONS

Allow Users to Mention Other Users



Figure 215.4: Mentions can also be enabled or disabled per site.

WORKING TOGETHER WITH THE WIKI

Wikis are for collaboratively building a collection of information. The most famous wiki on the planet is Wikipedia. It's a full encyclopedia developed collaboratively by users from all over the world, using a wiki. Liferay DXP's wiki does these things:

- Creates multiple wikis in a single wiki app instance.
- Scopes wikis to a page, a site, or the entire portal.
- Creates and edit wikis in WikiCreole syntax.
- Attaches files to wiki articles.
- Associates wiki articles with other assets in the portal.
- And more.

As you can see, Liferay DXP's wiki is flexible and can be configured to fit nearly any use case. What's more, it's completely integrated with the portal's user management, tagging, and security features.

WIKI

Lunar Resort [Main](#)

[FrontPage](#) [Recent Changes](#) [All Pages](#) [Orphan Pages](#) [Draft Pages](#)

FrontPage

[+ Add Child Page](#) [Edit](#) [Details](#) [Print](#)

Welcome to the Lunar Resort's wiki! Here, Lunar Resort guests can get together and write fun and informative content about things to do and see at the Lunar Resort.

6 Views

Your Rating [Average \(0 Votes\)](#)

☆☆☆☆☆ ☆☆☆☆☆

0 COMMENTS

Figure 216.1: The Wiki widget displays your wiki on a Site page.

GETTING STARTED WITH WIKIS

The Menu (☰) is the best place to start working with your wikis. Click the Menu (☰), navigate to your site, and select the *Content & Data* section. If you're updating an existing page-scoped wiki app instance, you can select that page scope from the scope menu the Gear icon (⚙️) makes available. The site's wiki app instance is available in the Default scope. Once you're in the proper content scope, click *Wiki*. The Wiki administration screen lets you add, modify, and delete wiki nodes. A Wiki app instance can contain many wiki nodes. By default, it contains one node: *Main*.

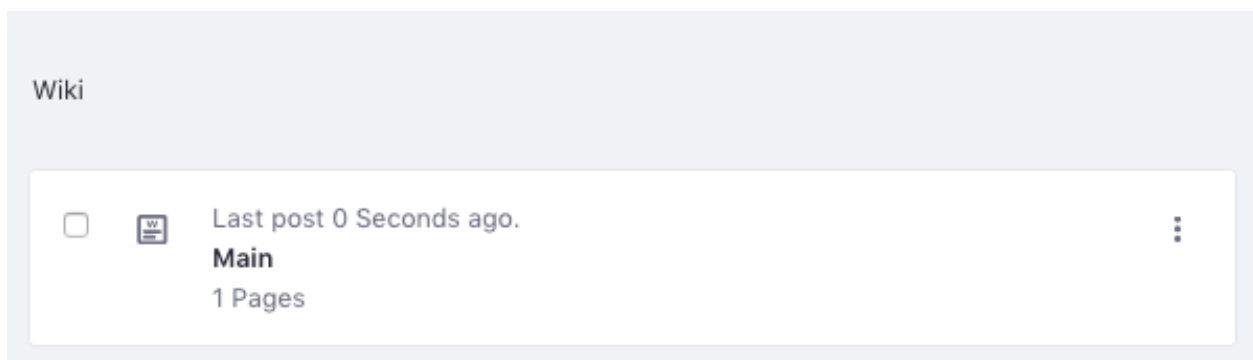


Figure 217.1: The Wiki app instance has a wiki node named *Main* with a single front page. You can build on the Main node or click the Add icon to create a new node.

217.1 Configuring Wikis

Before adding to your wiki instance, you should configure it. The instance's interfaces for permissions, export and import, configuration, and application templates are accessible from the Options menu. Click the *Options* icon (ⓘ) to open this menu.

The following options are available in this menu:


Wikis Permissions: Specify which Roles can create wiki nodes and access the Wikis Permissions screen. For example, if you've created a specific Role for creating wiki nodes and want to enable that Role to create new wiki nodes in this wiki application instance, select the Role's check box in the *Add Node* column and then click *Save*.

Export / Import: Import existing wiki content into your wiki app instance, or export wiki content to a file. For details, refer to Importing/Exporting Pages and Content.

Configuration: Configure email notifications and RSS feeds. The *Email From*, *Page Added Email*, and *Page Updated Email* tabs are similar to other apps' notification email settings tabs; they customize who wiki emails come from and the format and text of the email sent when a page is added or updated. The *RSS* tab lets you configure RSS feeds.

217.2 Adding Wikis


Follow these steps to create a new wiki node:

1. Click the *Add* icon () to start creating a new wiki node. The *New Wiki Node* form appears.
2. Add a name and description for the wiki node.
3. Open the form's *Permissions* section and define the wiki node's permissions. You can select the following permissions in the *Viewable by* menu:
 - Anyone (Guest Role)
 - Site Members
 - Owner

You can also click the *More Options* link to assign permissions to specific Roles.

4. Click *Save* when you're done creating the wiki node.

217.3 Wiki Node Options

Next to each listed wiki node is an *Actions* menu (). Here are the actions available in this menu:

Edit: Edit the wiki's name and description.

Permissions: Specify which roles can add attachments to wiki pages, add pages, delete pages, import pages, set permissions on the wiki node, subscribe to modifications, update existing pages, and view the wiki node.

Import Pages: Import data from other wikis. This lets you migrate from another wiki application to the Liferay DXP wiki. You might want to do this if you're migrating your site from a set of disparate

Figure 217.2: The New Wiki Node form lets you describe your new node, set view permissions, and set permissions for the Guest and Site Member roles.

applications (i.e., a separate forum, a separate wiki, a separate content management system) to Liferay DXP, which provides all of these features. Currently, MediaWiki is the only supported wiki.

RSS: Subscribe to an RSS feed using Live Bookmarks, Yahoo, Microsoft Outlook, or an application on your machine.

Subscribe: Subscribe to a wiki node. Any time a wiki page is added or updated, the portal sends you an email notification.

View Removed Attachments: Display attachments that have been removed from the wiki node.

Move to Recycle Bin: Moves the wiki node to the Recycle Bin.

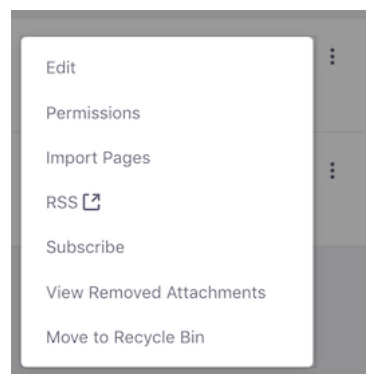


Figure 217.3: Each wiki node's Actions menu lists actions you can perform.

Before opening wiki nodes to contributors, you should consider whether to associate a workflow with them. For example, you could create a workflow that requires an administrator's approval to publish a wiki page modification (add, update, or delete). You can access your site's default *Wiki Page* workflow from within the Site Administration Menu, by navigating to *Configuration* → *Workflow* for your site. To learn how to use workflow, see the Workflow section.

ADDING AND EDITING WIKI PAGES

Wiki nodes initially have no pages. When you navigate into a node for the first time, a default page called *FrontPage* is created automatically. To view the page, click the wiki node's name and then click *FrontPage*. The *FrontPage* appears and shows a message that explains the page is empty and needs you to add content. That message is a link; click it to start editing the page. The wiki page editing form then appears.

Note: See the getting started article for instructions on accessing your wiki nodes.




Figure 218.1: Each empty wiki page presents a default message link you can click to edit the page.

Follow these steps to use the wiki page editing form:


1. Enter your content in the field that contains the text *Write your content here....* This is a rich-text, WYSIWYG editor that is almost identical to the one used in the Blogs app. The only difference is that the wiki editor uses Creole instead of HTML as its source. Click the link *Show Syntax Help* if you need help with Creole syntax (e.g., syntax for text styling, header formatting, link creation, etc.). For a detailed explanation of the rest of the editor, see the Blogs documentation.
2. If you want to attach files to the page, open the *Attachments* section of the form and add them via drag and drop or the *Select Files* button.
3. If you want to associate a tag with the page, open the *Categorization* section and enter a new or existing tag in the *Tags* field. See the documentation on tags for more information.

Figure 218.2: The wiki page editing form lets you create and edit your page's content.

4. If you want to select an existing asset in the portal (e.g., a media file, blog post, etc.) to relate to the page, open the *Related Assets* section and use the *Select* button to select that asset.
5. In the form's *Configuration* section, you can set the page to use Creole (default), plain text, or HTML. We recommend that you stick with the Creole format, as it allows for a much cleaner separation of content and code. You can also use the Configuration section to summarize your edit, and specify whether it's a minor edit.
6. Click *Publish* to publish the page when you're done editing it.

As is common with wikis in general, if you link to a page that doesn't exist, clicking that link opens the new page form with a note stating that the page doesn't exist and that you are creating it.

Note: When you create a page by clicking a link to a page that doesn't exist, the new page is **not** a child of the current page. The page is created at the wiki node's root. From Wiki in Site Administration, you can use the page's Move action to assign it a new parent page. Clicking the Move action brings up a window that lets you select a new parent for the wiki page.

Return to the wiki node view to see a list of the node's top-level pages. If you navigate to a page that has child pages, its child pages are listed. In these page listings, each page's Actions menu () lists the following actions you can take on the page:

Edit: Opens the page in the page editor.

Permissions: Lets you determine which roles can view, update, delete, subscribe to, or set permissions on the page, and add, update, or delete page discussions (comments).

Copy: Opens a page editor window with all the content from the source wiki page. You're prompted to specify a new title for it.



Move: Opens a dialog that lets you rename the page or assign the page to a new parent page within the wiki node.

Add Child Page: Create a new child page of the wiki page.

Subscribe (or Unsubscribe): Subscribes you to (or unsubscribes you from) notifications for the wiki page's modifications.

Print: Print the wiki page.

Move to Recycle Bin: Moves the wiki page to the Recycle Bin.

Each wiki page has a check box next to it. When you select a page's check box, the Management Bar changes to show an Info icon () and Recycle Bin icon (). To move the selected page to the Recycle Bin, click the Recycle Bin icon. To get additional information about the page via an info panel, click the Info icon. The info panel has a star icon that you can select to subscribe to the page's modifications. The info panel's Details section displays the page's summary, format, version, creation and modification dates, number of attachments, and RSS link.

There are several more features in the wiki node view's Management Bar. The *Filter and Order* menu orders the pages by title or modification date and filters them by page type. The arrows button sorts the pages in ascending or descending order. The search bar searches for pages.

The *View Types* button is next to the Info icon. It lets you choose how to display the pages. The View Types button's icon depends on the selected view type:



List (): Shows the pages in a list with an icon representing each page. Each page's entry contains the name of its author, when it was last modified, and its workflow status (e.g., Approved, Draft, etc.).

Table (): Shows the same information as the List view type, in a smaller list with no page icon. Also, the page's information is in columns and includes the revision number.

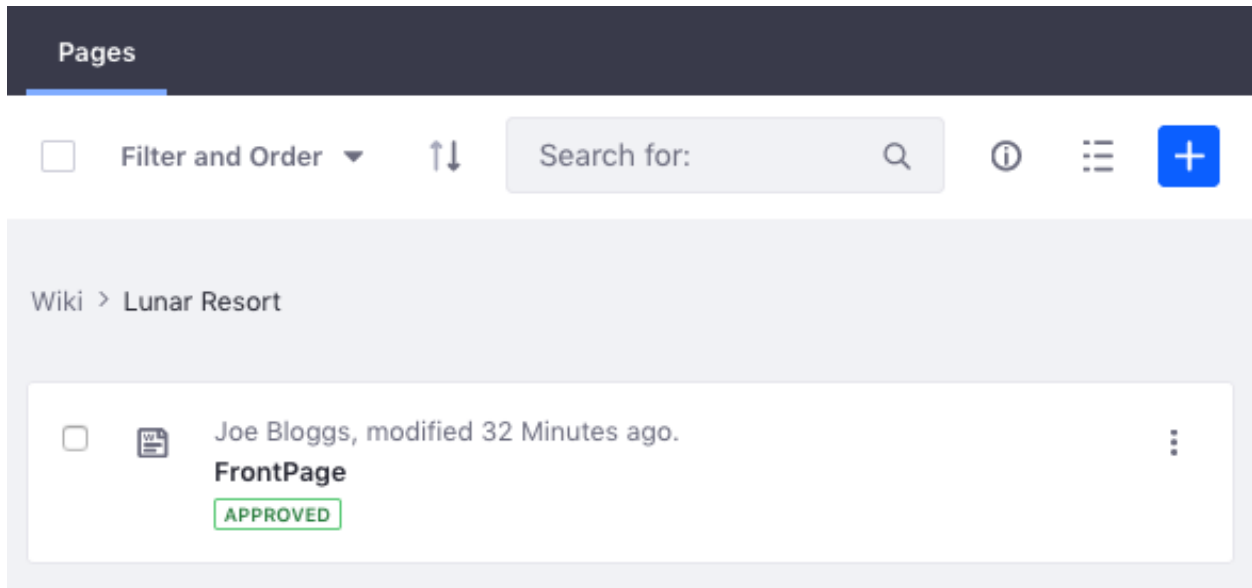




Figure 218.3: The wiki node's view in site administration has features that help you access and learn information about a wiki node's pages.

USING THE WIKI ON SITE PAGES

You can use the Wiki on Site pages via the Wiki widget. Follow these steps to add the Wiki widget to a page:

1. Navigate to the page where you want to place a wiki.
2. From the *Add* () menu, open *Widgets* → *Wiki* and add a *Wiki* to the page.

Your Site's wiki nodes appear in tabs across the top of the widget.

To view the Wiki widget's configuration options, click its *Options* icon () and select *Configuration*. The Configuration screen appears with these tabs:

Setup: Lets you choose wikis to display and gives you several options for displaying them. The *Show Related Assets*, *Enable Page Ratings*, *Enable Comments*, *Enable Ratings for Comments*, and *Enable Highlighting* check boxes enable or disable those features for the Wiki. You can set how you want users to interact with wiki documents. The *Display Template* selector menu lets you choose the Wiki's Widget Template. Below this, you can set which wiki nodes are visible. For example, you might host two wikis on a given site, exposing one to the public and keeping the other private for site members.

Communication: Configure communication across portlets, using predefined public render parameters. From here you can modify six public render parameters: *categoryId*, *nodeId*, *nodeName*, *resetCur*, *tag*, and *title*. You can perform these actions on each parameter:

- Ignore the values for this parameter that come from other portlets. For example, the wiki can be used along with the tags navigation app. When a user clicks on a tag in tags navigation, the wiki shows a list of pages with that tag. In some cases, an administrator may want the wiki to show the front page always, independently of any tag navigation done through other portlets. This can be achieved by selecting *Ignore*, so that the values of the parameter coming from those other portlets are ignored.
- Read the value of a parameter from another app. This is an advanced but very powerful option that lets portlets communicate without prior configuration. For example, imagine that the wiki is used to publish information about certain countries, and a custom app that allows browsing countries for administrative reasons was written and placed on the same page. You

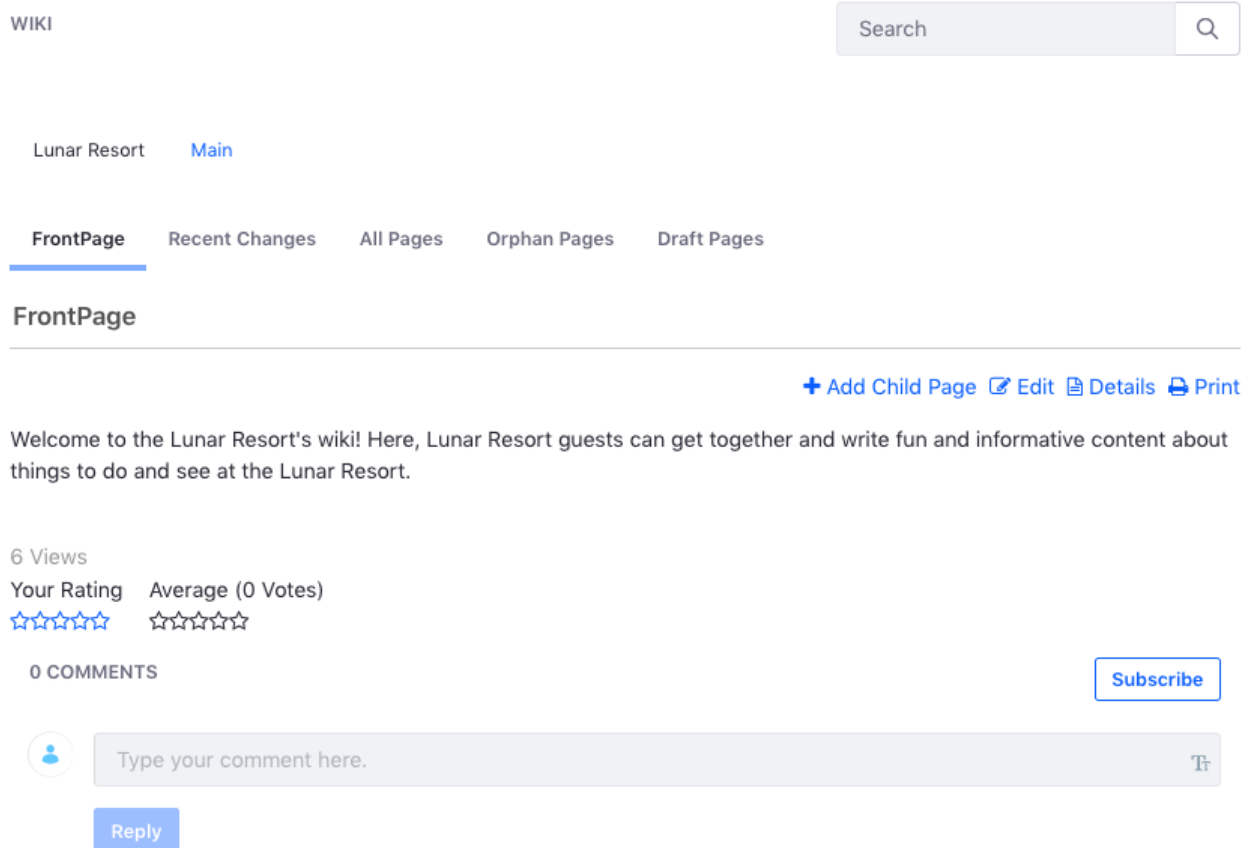


Figure 219.1: Users can interact with your Wiki nodes when you add the Wiki widget to a page.

could associate to this second app a public render parameter called *country* to designate the name of the country. Using this procedure, you can cause the wiki to show the information from the country being browsed in the other app. You can do this here for the wiki by setting the value for the title parameter to be read from the country parameter of the other app.

Sharing: Displays options you're likely to be familiar with such as the sections for sharing the Wiki with websites, Facebook, and NetVibes.

Scope: Set the wiki's scope. You can select the site-scoped or global-scoped instance, or select/create an instance for the page. If the page doesn't already have an instance scoped to it, you can click the *[page name] (Create New)* menu option to create a page-scoped wiki instance.

Once you set the wiki's configuration options the way you want them, click *Save*.

The Wiki's Options menu also contains the usual widget options:

Look and Feel Configuration: Set the widget's look and feel.

Export/Import: Export or import widget data.

Permissions: Set the widget's permissions.

Configuration Templates: Use configuration templates to store the widget's current setup or apply an existing archived setup.

Remove: Remove the widget from the page.




Figure 219.2: Here the user has selected to create a new Wiki instance scoped to the current page named *Welcome*

The Wiki displays links to all of the Wiki instance's nodes, and provides links for navigating around the wiki. Click on a wiki node's name to begin browsing that node's pages. The following navigation links are listed after the wiki nodes:

FrontPage: The wiki node's front page article. This is shown by default when the node is initially selected.

Recent Changes: Shows all of the recently updated pages.

All Pages: A flat, alphabetical list of all pages currently stored in the wiki.

Orphan Pages: A list of pages that have no links to them. This can happen if you remove a page link without realizing it's the only link to that page. This area lets you review such orphaned wiki pages so that you can re-link or delete them.

Draft Pages: A list of unpublished pages. Users can edit pages and save their changes as drafts. They can come back later to finish their changes and publish them.

The current wiki page's content shows in the wiki's main viewing area. Several features display above the wiki page content, depending on which wiki features are enabled and your permissions:

Add Child Page: Add a wiki page as a child of the current wiki page.

Edit: Edit the wiki page (if you have sufficient permissions).

Details: View the wiki page's details (if you have sufficient permissions). This is explained further in the documentation on page details.

Print: Print the wiki page.

Additional features appear below the wiki page's content. A view counter displays the wiki page's view count. Ratings and comments also appear if they're enabled.

WIKI PAGE DETAILS

When viewing a wiki page, you can view its details by clicking *Details* above the page content. Several tabs appear, to give you access to several categories of information about the page.




Figure 220.1: Click *Details* to view the wiki page's details.

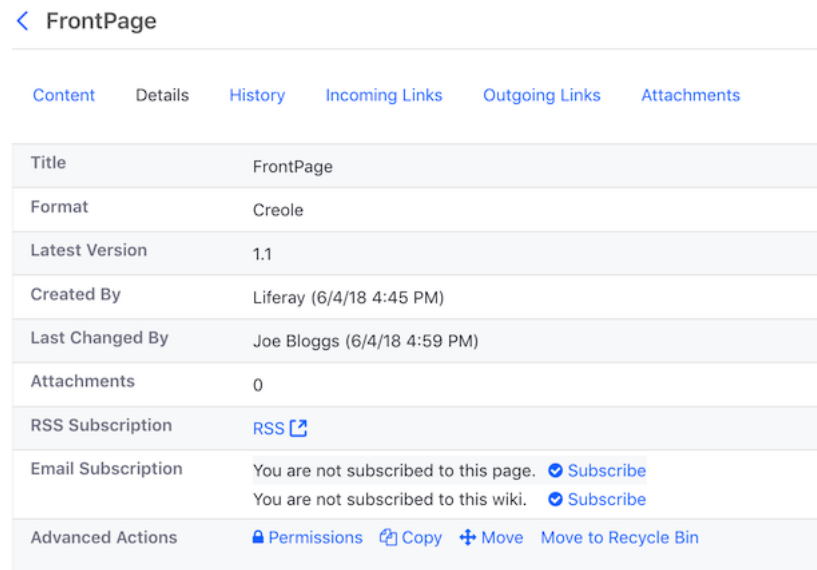


Figure 220.2: The wiki page's details.

220.1 Details

The Details tab shows page statistics and lets you perform some actions on the page:

Title: The page title.

Format: The page's format (Creole, HTML, MediaWiki, or plain text).

Latest Version: The page's latest version. The wiki automatically tracks page versions whenever a page is edited.

Created By: The user who created the page.

Last Changed By: The user who last modified the page.

Attachments: The number of attachments to the page.


RSS Subscription: An icon that opens a new page where you can subscribe to an RSS feed using Live Bookmarks, Yahoo, Microsoft Outlook, or an application you can choose from your machine.


Email Subscription: Links that let you to subscribe to or unsubscribe from modifications notifications for the page and the entire wiki node.

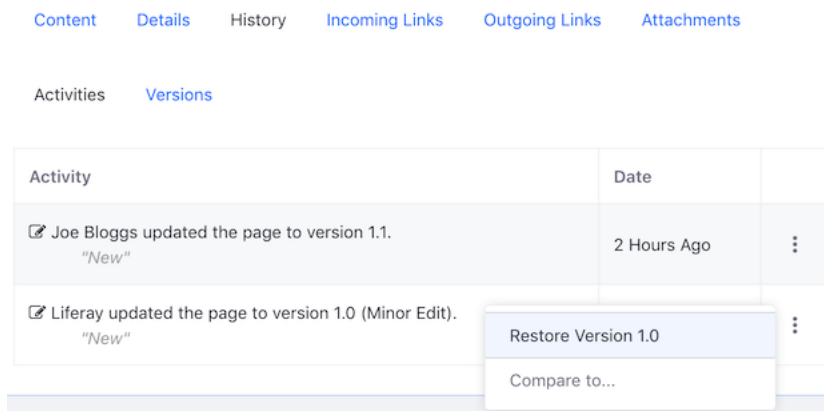
Advanced Actions: Links that let you modify the page's permissions, make a copy of the page, move (rename) the page, or move the page to the recycle bin.

220.2 History

The History tab lets you access the page's activities and versions via tabs:

Activities: Lists actions performed on the page. Each activity has an icon that represents the type of action, the name of the user, the action's description, date, and an *Actions* menu () to revert the action or compare its resulting version to that of another action.

Versions: Lists all the wiki page's versions. You can revert a page back to a previous version by selecting *Revert* from that version's *Actions* menu (). You can also compare the differences between versions by selecting two versions and then clicking the *Compare Versions* button.






| Activity | Date | |
|---|-------------|---|
|  Joe Bloggs updated the page to version 1.1. <i>"New"</i> | 2 Hours Ago |  |
|  Liferay updated the page to version 1.0 (Minor Edit). <i>"New"</i> | | <div style="border: 1px solid #ccc; padding: 5px; width: fit-content;"> Restore Version 1.0 Compare to... </div> |

Figure 220.3: The Activities tab displays the actions taken on the wiki page.

220.3 Incoming/Outgoing Links

The tabs *Incoming Links* and *Outgoing Links* list incoming and outgoing links, respectively. These are wiki links to and from the wiki page. You can use this tab to examine how this page links to other pages and how other pages link back to this page.

220.4 Attachments

The *Attachments* tab lists the name and size of each file attached to the page. You can attach any file to the wiki. Images are the most common type of file attached to a page. Referencing them using the proper WikiCreole syntax renders the image inline, which is a nice way to include illustrations in your wiki documents.

OTHER WIKI WIDGETS

The widgets that accompany the main Wiki widget help you display and navigate particular wiki nodes. The following widgets are available:

Page Menu: Displays a single wiki page's outgoing links.


Tree Menu: Displays a wiki's page hierarchy as a tree.

Wiki Display: Displays a single wiki node.

You can find these widgets in the *Add* () → *Widgets* → *Wiki* menu.

221.1 Page Menu

The Page Menu widget displays a wiki page's outgoing links. It answers the question, "What wiki pages can I access from this page?" After adding the Page Menu widget to a site page, you must set the wiki page it displays links from. Follow these steps to do so:

1. Click the widget's *Options* icon () and select *Configuration*.
2. In the configuration's *Setup* tab, choose the wiki node then click *Save*.
3. Still in the configuration's *Setup* tab, select the wiki page then click *Save* and close the configuration dialog.

When you click a Page Menu link, the site page's Wiki or Wiki Display widget displays the wiki page associated with that link.

221.2 Tree Menu





Figure 221.1: The Page Menu widget displays a wiki page's outgoing links.

The Tree Menu widget displays a wiki's page hierarchy as a tree that lets you navigate all the wiki's pages. Much like the Page Menu setup, you configure the Tree Menu widget to focus on a wiki node. You can also configure how deep users can navigate into the page hierarchy. You can set the *Depth* to a value from 1 to 5, or select *All* to allow navigation to all of the wiki node's pages.

Follow these steps to configure the Tree Menu widget after adding it to a site page:

1. Click the widget's *Options* icon () and select *Configuration*.
2. In the configuration's *Setup* tab, choose the wiki node then select the depth of wiki pages to display in the hierarchy.
3. Click *Save* and close the configuration dialog.

In the Tree Menu, folder icons represent parent wiki pages and document icons represent child wiki pages at the end of the nodes. When you click a parent wiki page or child wiki page, the Wiki or Wiki Display widgets on the site page display the respective wiki page.

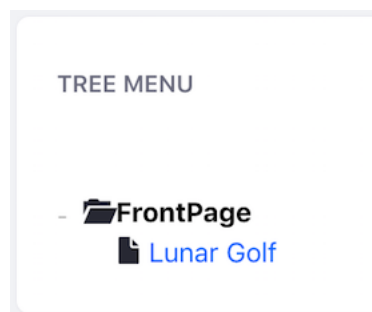



Figure 221.2: The Tree Menu widget displays a wiki node's hierarchy to the configured depth.

221.3 Wiki Display

The Wiki Display widget lets you focus user attention on one wiki node. After adding the widget to a page, follow these steps to configure it:

1. Click the widget's *Options* icon () and select *Configuration*.
2. In the configuration's *Setup* tab, choose the wiki node then click *Save*.
3. Still in the configuration's *Setup* tab, select the wiki page then click *Save* and close the configuration dialog. This page serves as the entry point for the wiki.

The configuration options and user interface for the Wiki Display are almost identical to that of the Wiki widget.


SENDING ALERTS AND ANNOUNCEMENTS

You can use the Alerts and Announcements widgets on Site pages to broadcast important information to users. The Alerts widget is designed for displaying high-priority information (e.g. planned downtime alerts, security alerts, etc.). Each alert is therefore labeled with a red *Important* tag. The Announcements widget displays all other information you want to broadcast on your site. Each announcement therefore lacks the red tag. To separate important alerts from more mundane announcements, you can place the Alerts and Announcements widgets on different pages. However, you can use either widget to display any information you wish. Besides the red tag, they function the same. You can also scope your alerts and announcements to specific groups of users.

These widgets have two tabs:



Unread: Non-expired alerts/announcements that you haven't read.

Read: Alerts/announcements that have expired, or that you've read.

Click an alert/announcement's *Actions* button () to edit or delete it.

222.1 Creating Alerts and Announcements

There are two places where you can create alerts and announcements:

1. The *Announcements and Alerts* app. Access this app at *Control Panel* → *Configuration* → *Announcements and Alerts*. Announcements and alerts are in separate tabs in this app. To begin creating an announcement or alert, select the appropriate tab and then click the *Add* button (). This app gives administrators a central location to create announcements and alerts that are then displayed on Site pages by the Announcements and Alerts widgets.
2. The Announcements and Alerts widgets, after adding them to a Site page from the *Add* () → *Widgets* → *News* menu. To begin creating an announcement or alert, click the widget's *Add Alert* or *Add Announcement* button.

Regardless of where you create the alert or announcement, the form for creating it is the same. Follow these steps to complete the form:

ALERTS

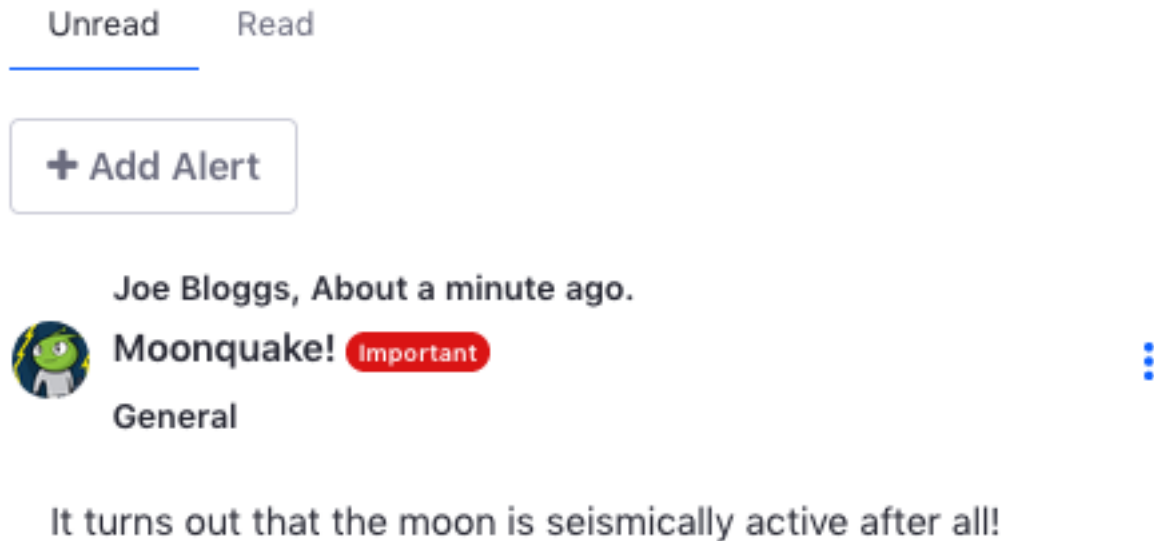


Figure 222.1: The Alerts widget provides administrators with an easy way to communicate important information to appropriate groups of users.

1. Use the *Title* field to give the alert or announcement a title. Then create your content in the field *Write your content here....* For a detailed explanation of the editor, see the Blogs documentation.
2. Open the *Configuration* section of the form and set the following options, if desired:

Distribution Scope: The scope where the alert/announcement is displayed. The default *General* scope sends the alert/announcements to everyone. Alternatively, you can select your site or specific roles as the scope.

URL: A URL (optional) to include with the alert/announcement. For example, an announcement about a news story could include a link to the news article. The URL must be valid and begin with `http://` or `https://`.

Type: The alert/announcement type. This can be *General*, *News*, or *Test*. Note that each user can specify a different delivery mechanism for each type of alert/announcement. See User Configuration for details.

Priority: The announcement's priority. This can be *Normal* or *Important*. Note that this is disabled for alerts because alerts are always high priority.

Display Date: The display date of the alert/announcement. This determines when the alert/announcement is sent to users and appears in the widget. By default, the *Display Immediately* box is checked. This sets the display date equal to the creation date. Uncheck this box to enter

The screenshot shows a mobile-style form titled 'Entry'. At the top left is a back arrow and the word 'Entry'. The form contains two main input areas: a 'Title' field and a larger text area with the placeholder 'Write your content here...'. Below these is a 'Configuration' section with a right-pointing arrow. At the bottom of the form are two buttons: a blue 'Save' button and a white 'Cancel' button with a grey border.

Figure 222.2: Enter your alert or announcement's title and content.

a custom display date. For example, administrators can create alerts/announcements for display on a later date. This date can be days, weeks, months, or years in the future. Once the *Display Immediately* box is unchecked, clicking the Display Date field opens the date-picker.

Expiration Date: The date and time the alert/announcement expires. Once an alert/announcement expires, the widget displays it in the Read tab. Clicking the Expiration Date field opens the date-picker.

3. Click *Save* when you're done. Your alert/announcement then appears in the widget.

222.2 User Configuration

Users can configure how they'd like to receive announcements.

1. Open your user menu and select *Account Settings*.
2. On the *Preferences* tab, select *Alerts and Announcements Delivery*. This shows options for customizing the delivery of alerts and announcements.
3. Select a configuration for each type of alert/announcement (General, News, or Test). For each type, you can enable delivery by email and SMS (text message). Note that the *Website* delivery option is selected and grayed out for each alert type. This means that each alert/announcement is always viewable in its respective widget on a site.
4. Click *Save* when you're finished.

Configuration ▼

Distribution Scope

General ⌵

URL

Type

General ⌵

Priority

Important ⌵

Display Date

04/10/2019 07:03 PM

Display Immediately

Expiration Date

05/10/2019 07:03 PM

Figure 222.3: Configure your new alert or announcement.

Alerts and Announcements Delivery

Select the delivery options for alerts and announcements.

| Type | Email | SMS | Website |
|---------|--------------------------|--------------------------|-------------------------------------|
| General | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| News | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| Test | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |

Figure 222.4: Each user can choose how they receive alerts and announcements.

222.3 Alert and Announcement Roles

You can also create roles for users to make general announcements. For instance, if you want someone specific to have strict control over announcements, give that person an Announcements Role. Follow these steps to create a simple Announcements Role:

1. Navigate to *Control Panel* → *Users* → *Roles*.
2. With the *Regular Roles* tab selected, click the *Add* button (+). This opens the *New Role* form.
3. Name your Role *Announcements*, give it a description, and click *Save*.
4. Select the Role's *Define Permissions* tab then grant these permissions:
 - In *Control Panel* → *General Permissions*, select *Add General Announcements*.
 - In *Site Administration* → *Applications* → *Announcements*, select all the resource permissions.

Click *Save* after selecting each permission. These permissions let the Role add alerts and announcements.

Now you have a simple Announcements Role that can manage your site's general announcements. Of course, you can adjust this Role's permissions as needed.

MANAGING NOTIFICATIONS AND REQUESTS

If you subscribed to a blog or message board, or if someone sent you a private message, invitation, event reminder, or mentioned you in a post, you received a notification or request.

To access notifications and requests, click your user menu and select *Notifications*. The *Notifications List* tab is selected by default. This is where all your notifications appear. Click the *Requests List* tab to view and manage your requests.



Figure 223.1: The *Notifications List* section displays all your notifications in a paginated list.

223.1 Managing Notifications

Notifications can pile up after some time, especially if you were away for a few days. The Management Bar gives you several ways to filter and sort your notifications.

The *Filter and Order* menu gives you the following options for viewing notifications:


All: The default option. Displays both read and unread notifications.

Unread: Displays notifications that haven't been marked as read. Unread notifications are indicated with a blue border on the left-hand side of the notification.

Read: Displays notifications that have been marked as read.




Date: Order notifications by date.

By default, notifications are listed by date in descending order. To sort notifications by ascending order, click the up/down arrow icon in the management bar. Clicking the button again reverses the sort.

Each notification's *Actions* menu () lets you mark the notification as read/unread, or delete the notification.

223.2 Managing Multiple Notifications

You can also manage multiple notifications at once. Select the checkbox next to notifications you want to manage and choose an option from the Management Bar. Select the checkbox above the notifications list to select all notifications on the current page. The Management Bar shows three actions for selected notifications:

- Mark as Read ()
- Mark as Unread ()
- Delete ()


223.3 Managing Requests

When you get a request, it appears in the *Requests List* tab. In each request's *Actions* menu (), you can click *Confirm* to accept, *Ignore* to decline, or *Delete* to remove the request.

USING THE KNOWLEDGE BASE

The Knowledge Base app can be used to display professional product documentation or form complete books or guides. It even lets you import article source files written in Markdown. It's workflow-enabled, so you can require articles to be approved before publication. Additionally, you can create article templates that help users follow a common outline.

Here are the Knowledge Base's key features:

- Navigation is built into Knowledge Base Display.
- The suggestions interface enables user feedback on articles.
- Articles are stored in folders.
- Metadata fields exist for the friendly URL, source URL, categorization, and related assets.
- The *Edit on GitHub* button () can take readers to an article's source repository location (if you choose to use it that way).
- Markdown source files can be imported to create and update articles.

The Knowledge Base has several widgets you can add to Site pages:

- Knowledge Base Article
- Knowledge Base Display
- Knowledge Base Search
- Knowledge Base Section

Lunar Excursion

[Wilma Carwin Takes Lunar 500](#)

[Analysts Bullish on Moon Pies](#)

Lunar Excursion

[RSS](#) [Subscribe](#) [History](#) [Print](#) [Edit](#) [+ Add Child Article](#) [Permissions](#) [+ Move](#) [X Delete](#)

Lunar Spelunking

We're going spelunking on the moon! Join us and be amazed!

When: 07 June 2184

Where to meet: The Lunar Resort's base camp.

What to wear or bring: Just your sense of adventure! The base camp has all the equipment you'll need.

Age requirements: 18+

[Share](#)

[1](#) [0](#)


[Wilma Carwin Takes Lunar 500](#)

Figure 224.1: Knowledge Base Display's navigation and viewing provide a great reading experience.

CREATING KNOWLEDGE BASE ARTICLES

The Knowledge Base app in Site Administration contains everything you need to create articles in the Knowledge Base. You can create articles by authoring them in the app's WYSIWYG editor or by importing them from Markdown files (.markdown, .md) in a ZIP archive. The sections below cover both ways of creating articles.

Note: To access Knowledge Base in Site Administration, a Role must have the permission *Knowledge Base → Access in Site Administration*. To add or act on articles, folders, or suggestions, the Site administrator must grant the appropriate permissions using the Permissions window in Knowledge Base.

To navigate to the Knowledge Base app, open the Menu () then go to Site Administration (the menu for your Site) → *Content & Data* → *Knowledge Base*. The Knowledge Base app has three tabs:

Articles: Create and manage articles and folders.


Templates: Create and manage templates.

Suggestions: Manage user-submitted feedback for articles.

Select the *Articles* tab, then proceed to the sections below for instructions on creating articles.

225.1 Authoring Articles in the Editor

Follow these steps to create an article in the editor:

1. In the Articles tab, click the *Add* button () and choose *Basic Article* or the name of an available template. This brings up the New Article form.
2. Enter a title for the article. A URL-safe version of the title you provide is added to the end of the article's friendly URL. You can manage the friendly URL in the *Configuration* section's *Friendly URL* field.

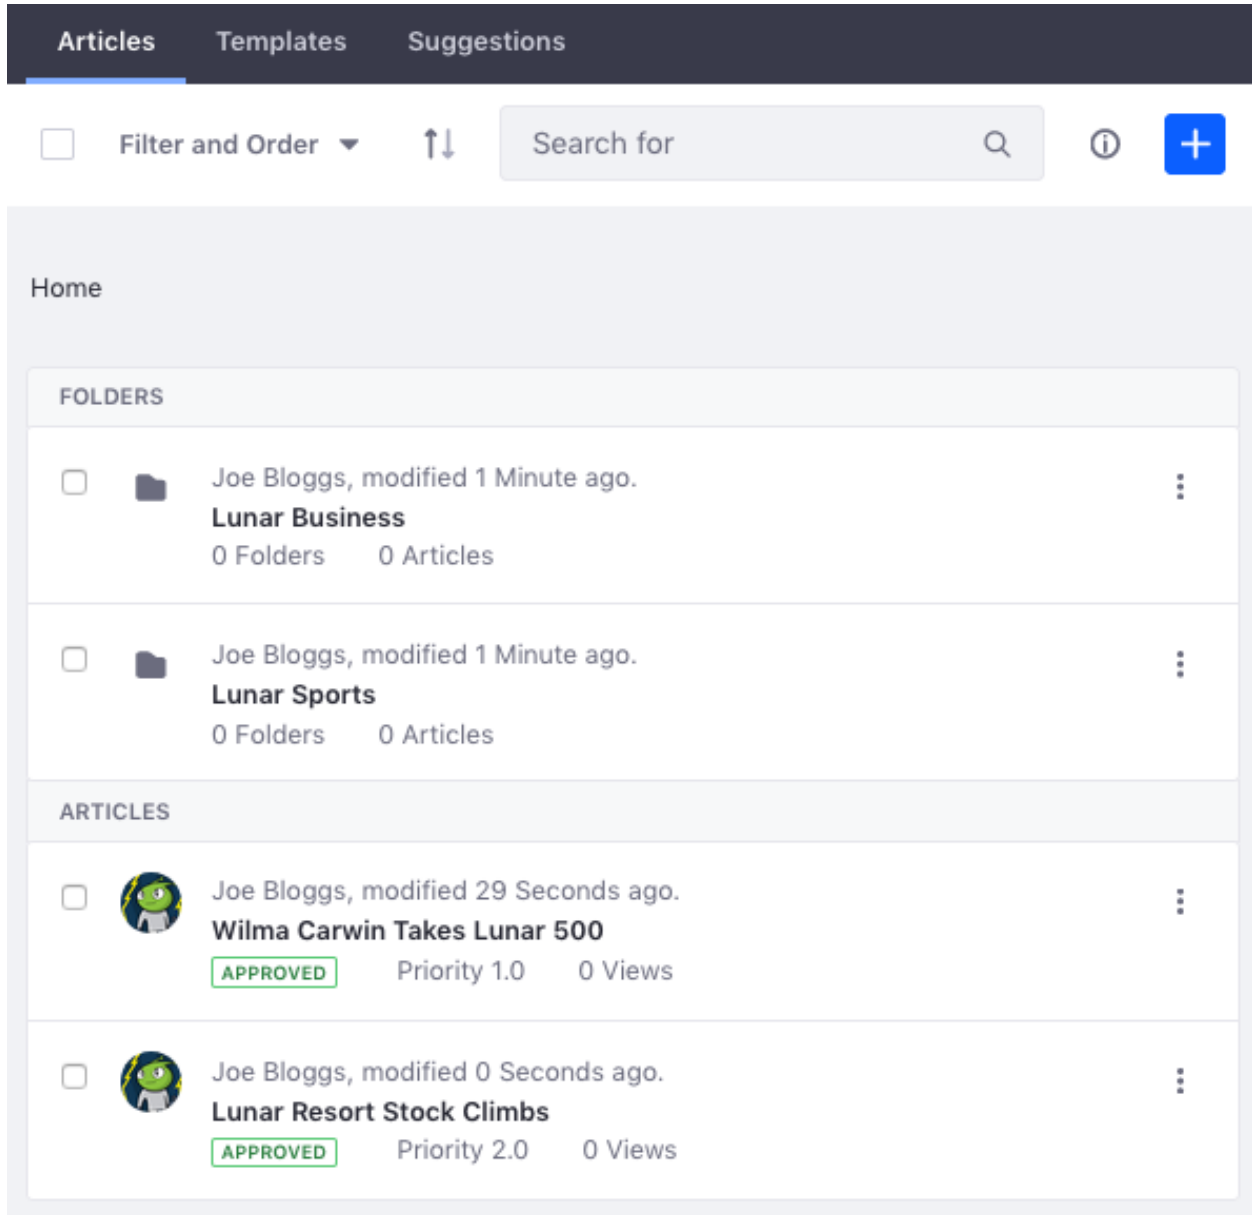


Figure 225.1: The Knowledge Base app in Site Administration lets you create Knowledge Base articles.

3. Use the WYSIWYG editor to create the article's content. To view or edit the article's HTML source, click the *Source* button in the editor. The sections below the editor let you add attachments and tags, specify related assets, and set permissions for the article. By default, View permission is granted to the Guest role, meaning anyone can view your article.

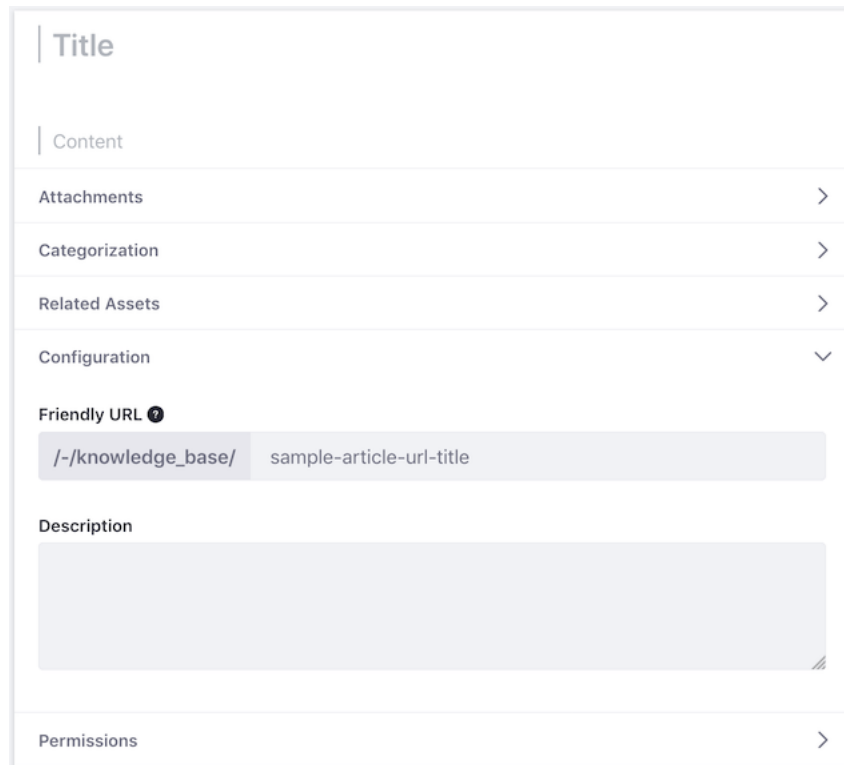
The image shows a vertical sidebar menu for editing a Knowledge Base article. The menu items are: Title, Content, Attachments, Categorization, Related Assets, Configuration, Friendly URL, Description, and Permissions. The 'Friendly URL' section is expanded, showing a text input field with the value '/-/knowledge_base/sample-article-url-title'. The 'Description' section is also expanded, showing a large text area with a diagonal slash icon in the bottom right corner.

Figure 225.2: You can create and modify a Knowledge Base article's content using the WYSIWYG editor.

4. Click *Publish* to submit the article for publication or click *Save as Draft* to continue working on it later. Note that if you've enabled workflow for the Knowledge Base, your article must be approved before publication.

Once the article is saved, it is converted automatically to HTML for the Knowledge Base. Articles are listed in a table in the Articles tab.

225.2 Importing Knowledge Base Articles

You can also create new Knowledge Base articles by importing them from a ZIP archive that contains articles in the Markdown format (.markdown, .md). For example, you could write articles in your favorite Markdown editor, package them in a ZIP file, and then import that ZIP file to create those articles in the Knowledge Base. The Knowledge Base can also prioritize articles by their filenames' numerical prefixes. For example, the Knowledge Base would list 01-article.markdown and 02-article.markdown in ascending order by their numerical prefix (01, 02). For more information on

article priority, see [Managing Knowledge Base Articles](#) For detailed information on the Knowledge Base importer, see the following topics:

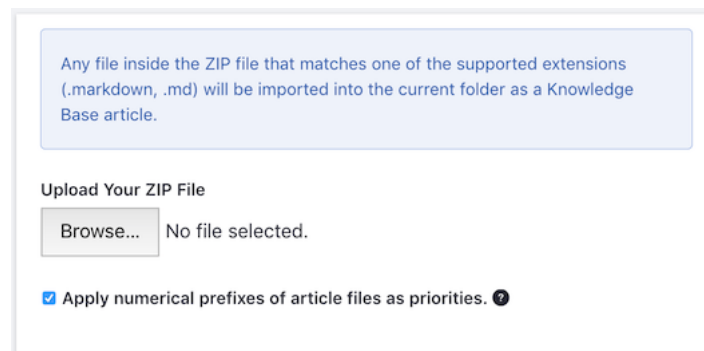
- [Importing Knowledge Base Articles](#)
- [Knowledge Base ZIP File Requirements](#)
- [Knowledge Base Importer FAQs](#)

Note: To import articles, your Role must have the permission *Knowledge Base → Resource Permissions: Import Articles*.

Follow these steps to import articles into the Knowledge Base:

1. In the Articles tab, click the *Add* button (+) and choose *Import*. This brings up the Import form.
2. Click *Browse* and select the ZIP file that contains the articles you want to import.
3. If you want to use the files' numerical prefixes to prioritize the imported articles in the Knowledge Base, select *Apply numerical prefixes of article files as priorities*.
4. Click *Save* when you're finished.

Like all articles, imported articles are automatically converted to HTML for the Knowledge Base and listed in a table with the rest of the articles in the Articles tab.



Any file inside the ZIP file that matches one of the supported extensions (.markdown, .md) will be imported into the current folder as a Knowledge Base article.


Upload Your ZIP File

No file selected.

Apply numerical prefixes of article files as priorities. ⓘ


Figure 225.3: You can import ZIP files that contain Knowledge Base articles in Markdown format.

MANAGING THE KNOWLEDGE BASE

The Knowledge Base app in Site Administration manages the Knowledge Base. To navigate to this app, open the Menu () then go to Site Administration (the menu for your site) → *Content & Data* → *Knowledge Base*.

Note: To access Knowledge Base in Site Administration, a Role must have the permission *Knowledge Base* → *Access in Site Administration*. To add or act on articles, folders, or suggestions, the site administrator must grant the appropriate permissions using the Permissions window in Knowledge Base.

226.1 Setting the Knowledge Base's Options

At the top-right of the Knowledge Base app, the Options menu () contains these options:

Subscribe: Get notified when Knowledge Base articles are created, updated, or deleted.

Home Folder Permissions: Define detailed permissions for the Knowledge Base app. You can choose the Roles that can perform the following tasks:

- Add/delete articles, folders, and templates
- Change the Knowledge Base app's permissions
- Subscribe to articles
- View templates and suggestions

Export/Import: Export or import the Knowledge Base app's configuration.

Configuration: Configure email notifications for article subscriptions and suggestions. You can also make the Knowledge Base app's articles available via RSS (enabled by default), and configure the RSS feed's options.

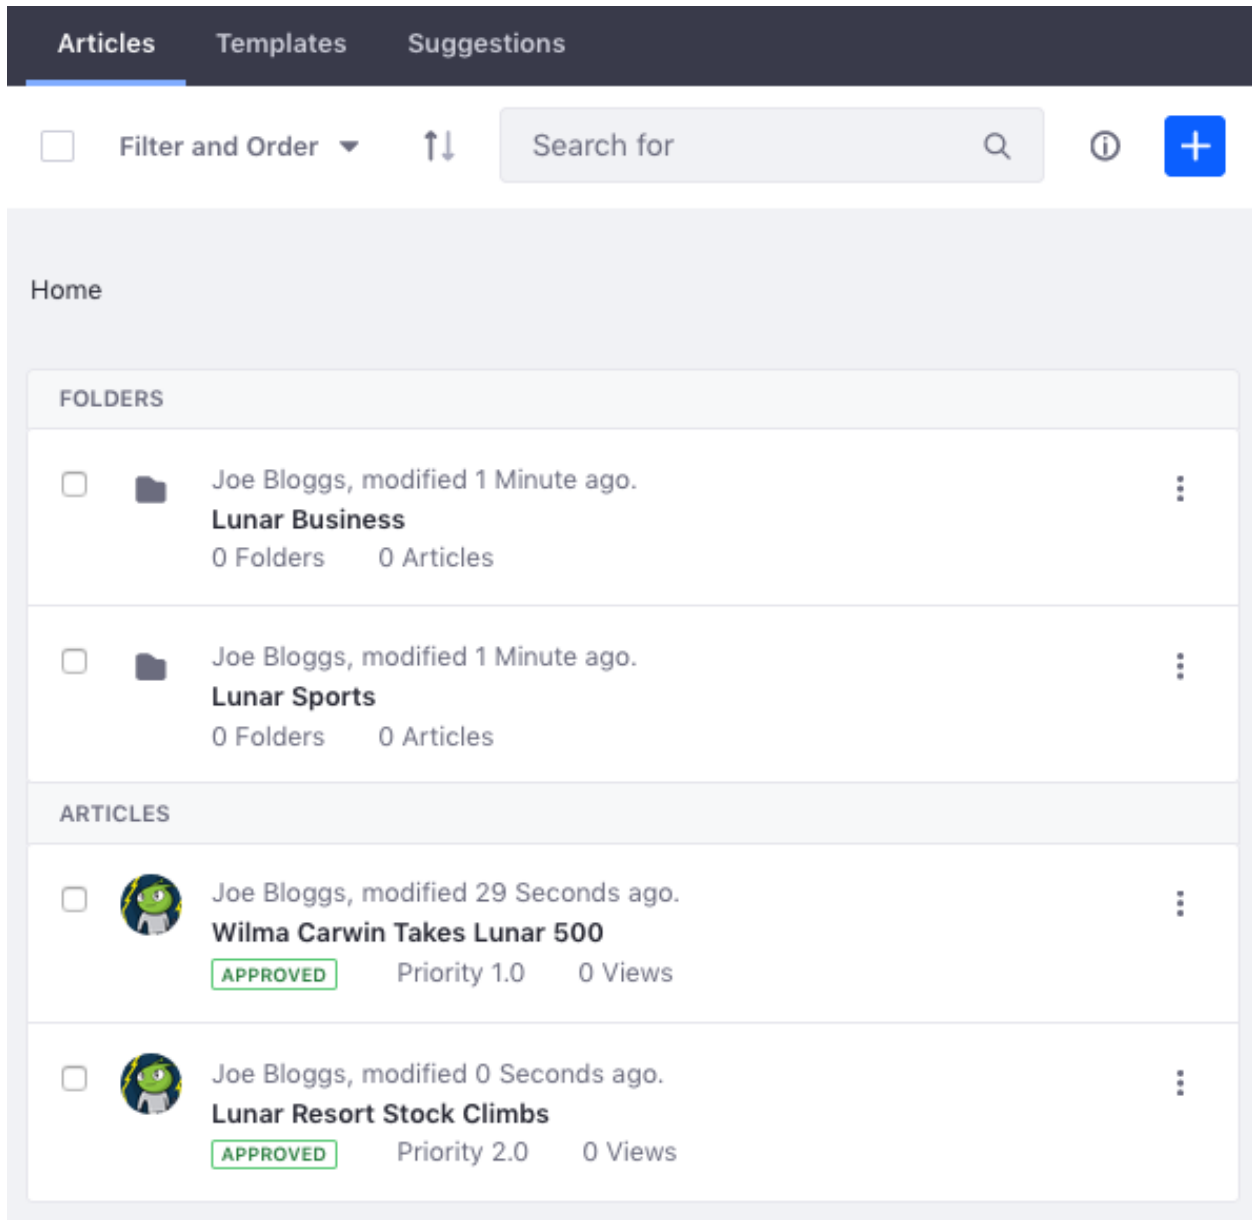


Figure 226.1: You can manage Knowledge Base articles, folders, and suggestions.





Figure 226.2: The Knowledge Base App's options.


226.2 Managing Knowledge Base Articles

Each article also has a *priority* value that determines its position in the Knowledge Base Display widget's navigation. Each article's priority value appears beneath the article's title. The Knowledge Base Display widget's navigation arranges articles in ascending priority. Priority 1 is the highest priority. The higher an article's priority, the higher it appears in the navigation. Articles are assigned the next lowest priority by default. This behavior can be changed via Knowledge Base System Settings.

To assign articles a new priority value, follow these steps:

1. Select *Move* from the Actions menu () next to the article.
2. Enter a new priority value for the article.
3. Click *Move* to apply the new priority.

You can also organize articles with folders. Follow these steps to create a folder:

1. Click the Add button () and select *Folder*. This opens a form for creating the new folder.
2. Enter a name and an optional description.

By default, anyone can view the folder. You can manage this setting along with the other permissions in the form's *Permissions* section.

3. Click *Save*. The folder is then listed in a table in the Articles tab.

The text immediately below the *Filter and Order* selector at the top of the app shows your position in the folder hierarchy. Click a folder's name in the hierarchy to navigate to it. You can also move articles into folders and create child articles. Knowledge Base also supports nested folders.





Figure 226.3: This screenshot uses a red box to highlight the text that indicates the current position in the folder hierarchy.


Each folder's Actions menu () lets you perform the following actions on the folder:

Edit: Change the folder's name and description.

Move: Relocate the folder under a new parent folder or update its priority.

Delete: Remove the folder and its articles from the Knowledge Base.

Permissions: Grant or revoke the following permissions: add an article to the folder, add a sub-folder to the folder, delete the folder, move the folder, set permissions on the folder, edit (update) the folder, and view the folder.

You can also delete multiple articles or folders at once. To do this, select the checkbox for each item that you want to delete and click the *Delete* button () that appears in the Management Bar.

You can also see the info for selected items by clicking the *Info* button () in the Management Bar.

KNOWLEDGE BASE TEMPLATES

Templates give users a starting point. For example, you can create templates that contain default headers or other content for articles. Templates help foster consistent formatting and content organization for articles. You can create and manage templates from the Knowledge Base app in Site Administration. To navigate to this app, open the Menu (☰) and go to Site Administration (the menu for your Site) → *Content & Data* → *Knowledge Base*.

Note: To access Knowledge Base in Site Administration, a Role must have the permission *Knowledge Base* → *Access in Site Administration*. To add or act on articles, folders, or suggestions, the Site administrator must grant the appropriate permissions using the Permissions window in Knowledge Base.

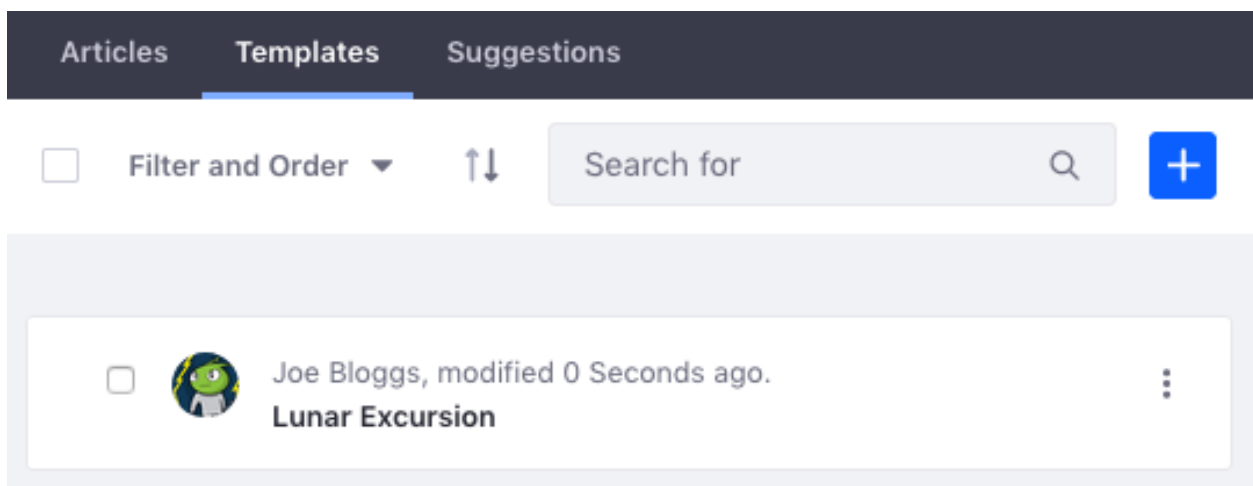
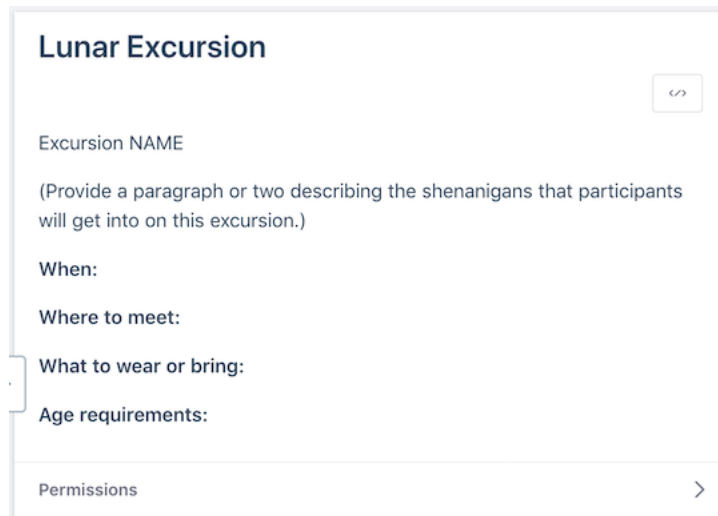


Figure 227.1: The Knowledge Base app's Templates tab.

227.1 Creating Templates

To create a new template, follow these steps:

1. Click the *Templates* tab.
2. Click the *Add* button (+). This brings up the *New Template* form.
3. Enter a title for the template.
4. Use the WYSIWYG editor to create the template's content. To view or edit the article's HTML source, click the *Source* button (</>) in the editor. You can also set the template's permissions via the form's *Permissions* section.
5. Click *Publish* to finish creating the template.



The screenshot shows a form titled "Lunar Excursion" with a "Source" button in the top right corner. The form contains several text input fields: "Excursion NAME", a descriptive paragraph field with the placeholder "(Provide a paragraph or two describing the shenanigans that participants will get into on this excursion.)", "When:", "Where to meet:", "What to wear or bring:", and "Age requirements:". At the bottom, there is a "Permissions" section with a right-pointing arrow.

Figure 227.2: The New Template form.

227.2 Managing Templates

Each template appears in a list in the *Templates* tab. You can take the following actions on each template via its *Actions* button (⋮):


View: Display the template. From here, you can print the template, use it to create an article, edit it, modify its permissions, or delete it.

Edit: Change the template's title and content.

Permissions: Configure the template's permissions. You can choose whether a Role can change permissions, update, view, or delete the template.

Delete: Remove the template from the Knowledge Base.

RESPONDING TO KNOWLEDGE BASE FEEDBACK

The Knowledge Base app's *Suggestions* tab shows user feedback on articles and lets you mark progress on addressing the feedback. To navigate to this app, open the Menu () then go to Site Administration (the menu for your Site) → *Content & Data* → *Knowledge Base*.

Note: To access Knowledge Base in Site Administration, a Role must have the permission *Knowledge Base* → *Access in Site Administration*. To add or act on articles, folders, or suggestions, the Site administrator must grant the appropriate permissions using the Permissions window in Knowledge Base.

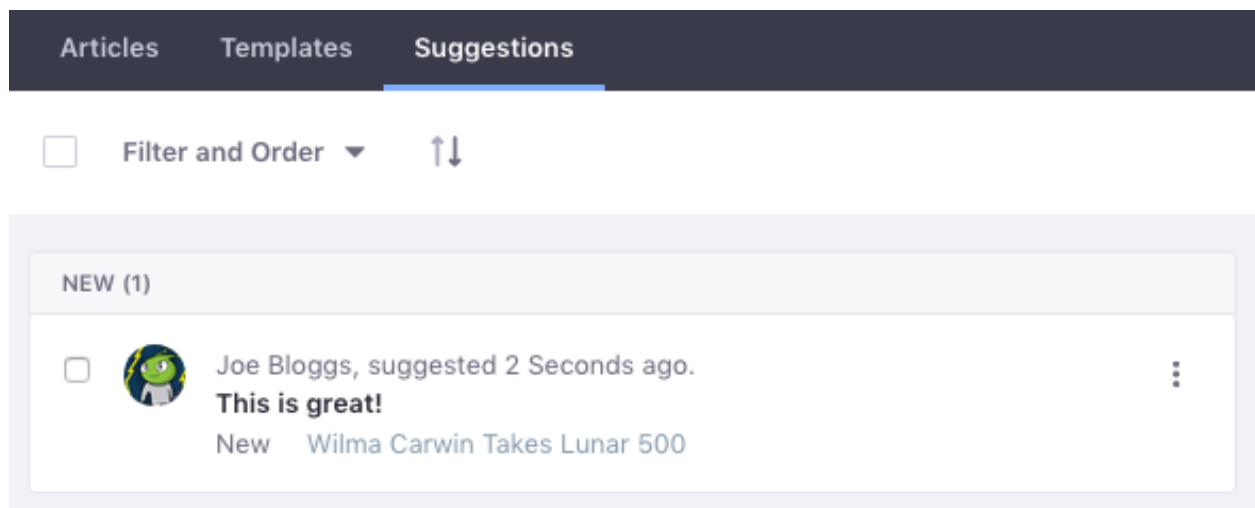



Figure 228.1: The Suggestions tab in Knowledge Base displays each piece of feedback that users leave on Knowledge Base articles.

Each suggestion provides the link to the associated article, the user's feedback, the user's name, the feedback's time stamp, and the status on addressing the suggestion. You can use each entry's Actions menu () to move the entry between the *New*, *In Progress*, and *Resolved* states.


Note: To view article suggestions, your Role must have the permission *Knowledge Base → Knowledge Base: View Suggestions*.

To move suggestions between the *New*, *In Progress*, and *Resolved* states, your Role must have the permission *Knowledge Base → Knowledge Base Article: Update*. Roles assigned this permission can also view and update the state of article suggestions from any of the other Knowledge Base widgets.


When you move the suggestion to a different state, an email is sent notifying the user of the change. You can view and configure the automated emails from the Knowledge Base app's *Options* (ⓘ) → *Configuration* menu.


KNOWLEDGE BASE DISPLAY

You can use the Knowledge Base Display widget to display your published Knowledge Base articles. You can customize how this widget displays articles, and which ones it displays. To get started, add the widget to the Site page you want to display articles on:

1. Navigate to the page and open the *Add* menu (.
2. Open the *Widgets* → *Content Management* section, then add *Knowledge Base Display* to the page.

By default, the Knowledge Base Display widget displays articles from the Knowledge Base's Home folder. To change the location of its articles, follow these steps:

1. Click the widget's *Options* icon () and select *Configuration*. This opens the Configuration dialog.
2. In the *Setup* tab, select the *General* tab then click *Select* in the *Article or Folder* field. This brings up the *Select Entry* form.
3. Click *Choose* next to the article or folder of articles you want to display. Alternatively, you can click the *Choose This Folder* button at the top of the *Select Entry* form to select the current folder.
4. Click *Save* and close the Configuration dialog.

The Knowledge Base Display widget's Options icon () also provides these common configuration options:

- Look and Feel Configuration
- Export/Import
- Permissions
- Configuration Templates

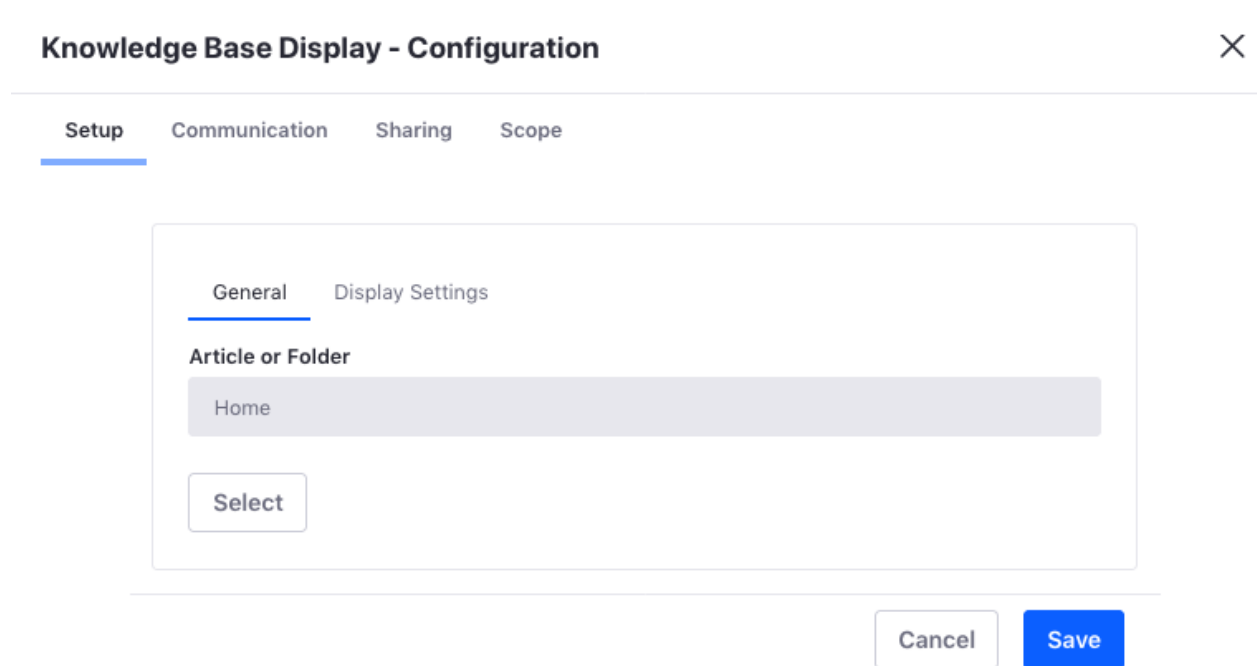


Figure 229.1: Select the article or folder of articles that the Knowledge Base Display widget displays.

For more information on these, see the section on configuring widgets in Web Experience Management.

The Knowledge Base Display's navigation menu and display options make it the perfect candidate for a full page widget. If you display a folder of articles, the navigation on the left side of the widget displays links to all the folder's articles. The viewing area on the right side of the widget displays the folder's leading article (the *priority one* article). Click an article in the navigation to display it in the viewing area. The currently displayed article's link appears in bold in the navigation. You can also move between articles by clicking the links with arrows at the bottom of the widget.

Knowledge Base Display can also show article hierarchies. Viewing an article that has child articles expands the navigation tree to show links to the child articles. Any expanded nodes collapse when you view a different top level article.

The links at the top of the widget allow users to perform the following actions on an article:

- Subscribe to an RSS feed of the Knowledge Base
- Subscribe to the current article
- View the current article's history
- Print the current article

Administrators have access to an additional set of links at the top of the widget that lets them perform the following actions:


- Edit the article
- Add a child article
- Set the article's permissions
- Move the article

The screenshot shows a side-by-side comparison of article navigation and viewing. On the left, under the heading 'Lunar Excursion', there are two article titles: 'Wilma Carwin Takes Lunar 500' and 'Analysts Bullish on Moon Pies'. On the right, under the heading 'Lunar Excursion', there is a horizontal menu with icons and labels for 'RSS', 'Subscribe', 'History', 'Print', 'Edit', 'Add Child Article', 'Permissions', 'Move', and 'Delete'. Below this menu, the article content is displayed, including the title 'Lunar Spelunking', a short paragraph 'We're going spelunking on the moon! Join us and be amazed!', and metadata: 'When: 07 June 2184', 'Where to meet: The Lunar Resort's base camp.', 'What to wear or bring: Just your sense of adventure! The base camp has all the equipment you'll need.', and 'Age requirements: 18+'. At the bottom of the article content, there is a 'Share' button and a rating interface showing '1' thumbs up and '0' thumbs down. To the right of the article content, there is a link 'Wilma Carwin Takes Lunar 500' with a small icon.

Figure 229.2: Knowledge Base Display’s navigation and viewing provide a great reading experience.

- Delete the article

Below the article’s content is the rating interface, showing thumbs up/down icons. Users can also submit suggestions or comments below the article in the text box labeled *Do you have any suggestions?*. Administrators can view the suggestions and mark progress on them.

If the administrator enables the Knowledge Base app’s source URL feature (more on this in a moment) and an article has an assigned source URL, an *Edit on GitHub* button () appears to the right of the article’s title. This button lets users access the article’s source in GitHub. You can use this feature to encourage users to contribute fixes or improvements to articles. If you’re interested in this feature, you can direct your administrator to follow the instructions in Knowledge Base System Settings.

229.1 Displaying Different Article Sets

As an administrator, say that you’ve used folders to aggregate similar articles, and you want to provide an easy way for users to switch between these sets of articles. The Knowledge Base Display’s content folder feature adds a selector to the top of the navigation that lets users switch between article sets.

Follow these steps to set up content folders:

1. Add a folder in the Knowledge Base app in Site Administration. Then create sub-folders in this folder. These sub-folders are the *content folders*.
2. Add articles to each content folder.

3. Select *Configuration* from Knowledge Base Display's *Options* menu (). In the *Setup* → *General* tab, select the content folders' parent folder and click *Save*.

The content selector's values reflect the names of your content folders. Select one to view its articles.

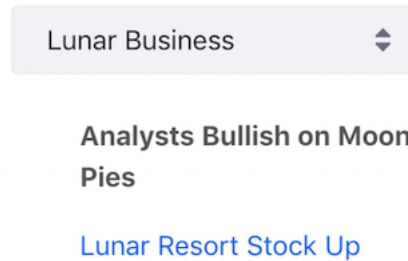



Figure 229.3: Knowledge Base Display's content folder feature lets users switch between different sets of articles.

You can also add a common prefix to the names shown in the selector:

1. Select *Configuration* from Knowledge Base Display's *Options* menu (). In the Configuration dialog, select the *Setup* → *Display Settings* tab.
2. Enter the prefix into the *Content Root Prefix* field and click *Save*.

OTHER KNOWLEDGE BASE WIDGETS

There are other Knowledge Base widgets you can add to Site pages besides Knowledge Base Display:

Knowledge Base Article: Display a single article's content.

Knowledge Base Section: Publish articles associated with a specific topic (section).

Knowledge Base Search: Search for articles.


You can add these widgets from *Add* () → *Widgets* → *Content Management*.

230.1 Knowledge Base Article

Knowledge Base Article displays a single article's content. It even shows abstracts of child articles. You can add multiple Knowledge Base Article instances to a page, and each one can show a different article.

After adding Knowledge Base Article to a page, follow these steps to configure the widget:

1. Click *Please configure this portlet to make it visible to all users*. This opens the Configuration dialog.
2. In the *Setup* → *General* tab, click *Select*, choose an article, click *Save*, and close the Configuration dialog.

You can change your selection from the widget's *Options* () → *Configuration* menu.

Knowledge Base Article shares the same UI as the Knowledge Base Display to display and manage its articles. Refer to the Knowledge Base Display documentation for a detailed description of the widget's UI.

230.2 Knowledge Base Section

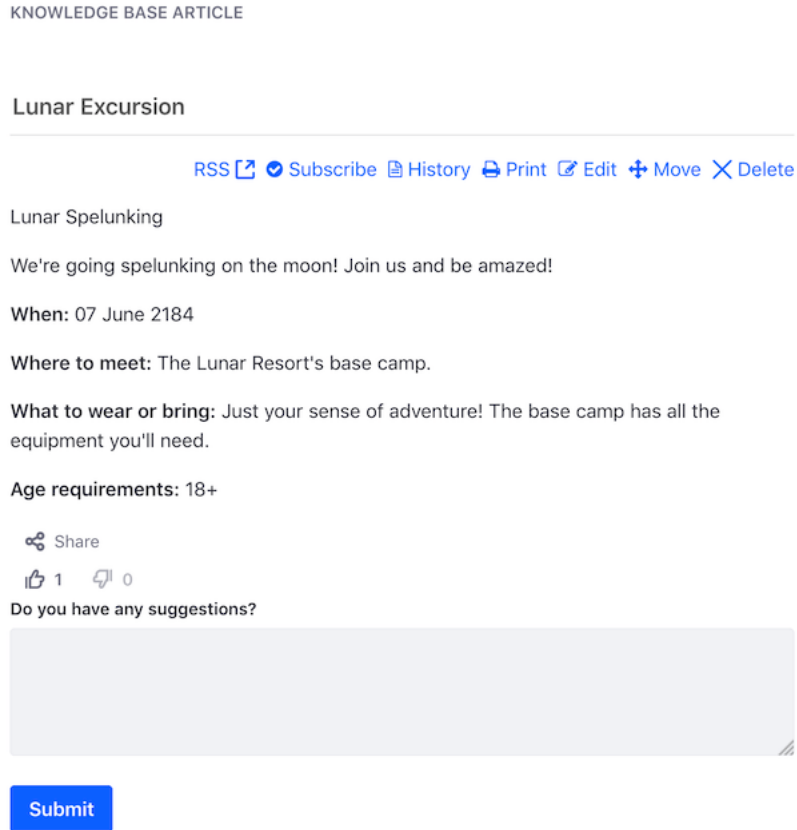


Figure 230.1: The Knowledge Base Article app is great at displaying individual articles.

Note: as of Knowledge Base 3.0.0, the Knowledge Base Sections widget is deprecated and replaced by categories.

The Knowledge Base Section widget lets you publish articles associated with a specific topic (section). For example, a news Site might have the sections *World*, *Politics*, *Business*, and *Entertainment*.





Figure 230.2: The Knowledge Base Section widget.

To use sections, an administrator must first configure the feature in System Settings, creating

the section names for use in the Knowledge Base Section widget. This process is covered in detail in Knowledge Base System Settings. When creating or editing a Knowledge Base article, authors can then select the article's section in the *Configuration* → *Section* field.

You can add multiple instances of the Knowledge Base Section widget to a page. Each widget can display articles from any number of sections. You can configure the widget to display article titles or abstracts. You can also define whether to show pagination or section titles.

Follow these steps to configure an instance of the Knowledge Base Section widget:

1. Select *Configuration* from the Knowledge Base Section widget's *Options* menu (). This opens the widget's Configuration dialog.
2. In the Configuration dialog's *Setup* → *General* tab, select the section or sections that you want to use and click *Save*.
3. Close the Configuration dialog to see the updates.

The matching articles are displayed in the app beneath their section heading.


230.3 Knowledge Base Search

Note: as of Knowledge Base 3.0.0, the Knowledge Base Search widget is deprecated and replaced by Liferay Search.

Even though the Knowledge Base can show the structure of its articles, it may be difficult to find exactly what you're looking for by browsing. That's where the Knowledge Base Search widget comes in.

Enter your search term and press the *Search* button. The results are displayed in a table with the following criteria for each matching article:

- Title
- Author
- Create date
- Modified date
- Number of views

You can select the criteria to display in the widget's *Options* () *Configuration* dialog.

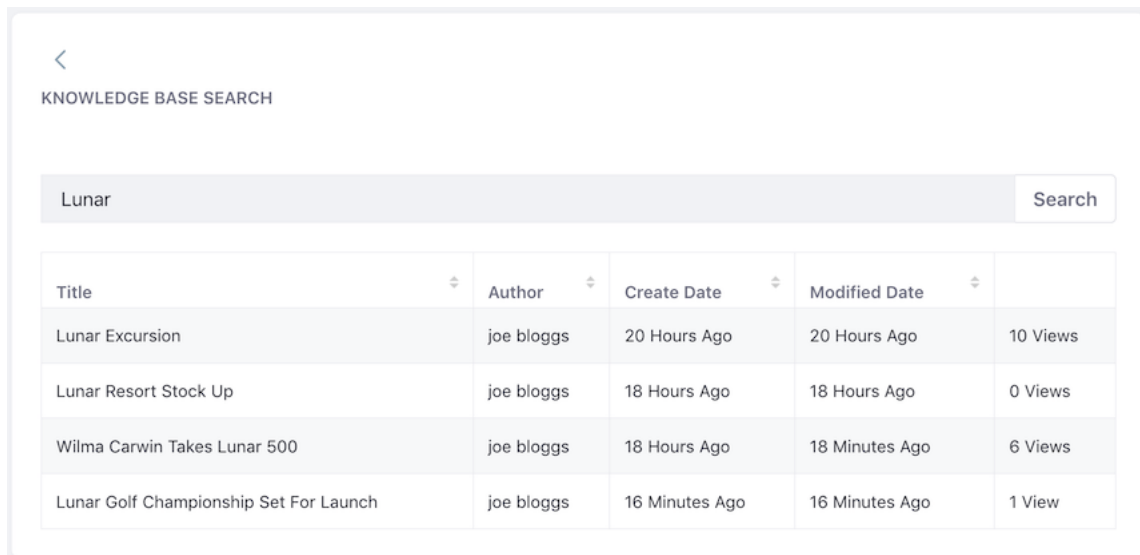


Figure 230.3: The Knowledge Base Search widget lets you search the Knowledge Base for keywords.

IMPORTING KNOWLEDGE BASE ARTICLES

As mentioned earlier, the Knowledge Base app can import articles in bulk. This lets you have an offline process where articles are prepared ahead of time before they're published. Articles are imported into the Knowledge Base as Markdown files. Markdown is a text-only file format that is easy to read, yet supports all the things you'd need to do to format your articles.

Note: To import articles, your Role must be granted the *Knowledge Base → Resource Permissions: Import Articles* permission.

The Knowledge Base supports a Markdown dialect known as Multi-Markdown. This dialect extends the original Markdown with features like table formatting, image captions, and footnotes.

For the Knowledge Base to import your Markdown articles, they must adhere to these requirements:

- All source files must use the `.markdown` or `.md` extensions.
- Articles must start with a top-level header (e.g., `# Some Heading ...`).
- Each header must have an associated, unique ID for the article's friendly URL title and for anchor tags in the article's sub headers. Here's an example of a top-level header that correctly specifies an ID:

```
# Some Heading [](id=some-heading)
```

Here's Markdown source text for a simple example article:

```
# The Moons of Mars [](id=the-moons-of-mars)
```

```
As you look up from your chaise lounge, you're sure to see our neighboring planet Mars. Did you know that Mars has two moons? You might have to break out a pair of binoculars to see them.
```

```
Its two moons are aptly named after the two sons of mythical Roman god Mars. Their names are Phobos and Deimos.
```

In the first line above, notice the header's ID assignment `id=the-moons-of-mars`. On import, the ID value becomes the Knowledge Base article's URL title.

Markdown is something of a standard: there's Github Flavored Markdown, a proposed common Markdown syntax, forums that support Markdown (reddit, StackExchange, and others), Markdown

editors, and an IETF draft for making it an official Internet media type (text/markdown). Why is there so much interest in Markdown?

1. It's readable. Even if you don't know Markdown, you can read it without having to filter out the syntax.
2. It gets out of a writer's way. You don't have to worry about mousing to various icons to change text into a heading or create bulleted lists. Just start typing. The syntax is very intuitive.
3. There are tools to convert it to many other formats, though it was designed to convert to HTML. If your articles are in Markdown, it's straightforward to publish them to the web, mobile formats (Kindle, ePub), and print.
4. Since it's only text, you can use existing tools to collaborate on that text. Using services like GitHub, people can contribute to your articles, and you can see all the changes that have been made to them.

KNOWLEDGE BASE ZIP FILE REQUIREMENTS

The Knowledge Base importer supports article hierarchies, so Markdown files can be specified anywhere in the ZIP file's directory structure. They can be nested in any number of folders. Image files are the only files supported for attachments.

Note: Imported articles are independent of the workflow settings. This means that **imported articles are automatically approved**.

Only users with the *Import Articles* permission assigned to their Role are able to import articles. This permission can be assigned manually through *Control Panel* → *Users* → *Roles*. If you've upgraded from Liferay Portal 6.2, you can also assign this Role to every Role that was already able to add articles with a command from the Gogo shell.

Open the Gogo shell. Type `knowledgeBase:addImportArticlePermissions` and hit enter.

The ZIP file's articles are imported in file order (alphanumerically). To designate an article's priority, add a numeric prefix to its file name. For example, the priorities for articles named `01-file.markdown` and `02-file.markdown` become `1.0` and `2.0`.

To designate an article to be the parent of all other articles in the same source folder, end its file name with `-intro.markdown`. This creates a parent-child hierarchy. You can use the prefix `00` for parent articles to place them at the top of the folder's file order. The importer uses the numeric prefix of an intro file's folder as its article priority.

Here's the underlying logic for the `00` prefix:

- A file prefix of `00` for a non-intro file assigns the resulting article's priority to `1.0`.
- A file prefix of `00` for a top-level intro file sets the article's priority to the first folder numeric prefix found that is `1.0` or greater.

This convention lets you specify priorities for top-level (non-child) articles in your hierarchy.

When importing, keep the checkbox labeled *Apply numerical prefixes of article files as priorities* selected. If a file doesn't have a prefix, its article gets the next available priority (the highest current priority, plus one).

Below is an example ZIP file structure that demonstrates the features mentioned so far:

ZIP File Structure Example:

- `01-winter-events/`

- 00-winter-excursions-intro.markdown
- 01-star-dust-snow-shoeing.markdown
- 02-lunar-alpine.markdown
- 02-summer-events/
 - 00-summer-excursions-intro.markdown
 - 01-lunar-rock-scrambling.markdown
 - 02-extra-terrestrial-mountain-biking.markdown
 - 03-lunar-olympics/
 - * 00-lunar-olympics-intro.markdown
 - * 01-zero-gravity-diving.markdown
- images/
 - some-image.png
 - another-image.jpeg

The above ZIP file specifies 00-winter-excursions-intro.markdown as the parent of its neighboring Markdown files: 01-star-dust-snow-shoeing.markdown and 02-lunar-alpine.markdown. Likewise, 00-lunar-olympics-intro.markdown is the parent of 01-zero-gravity-diving.markdown. 00-lunar-olympics-intro.markdown is also the peer of 01-lunar-rock-scrambling.markdown and 02-extra-terrestrial-mountain-biking.markdown, and the child of 00-summer-excursions-intro.markdown.

ZIP Example's Resulting Relationships and Priorities



- 01-winter-events/00-winter-excursions-intro.markdown
 - **Article:** Winter Excursions
 - **Relationship:** Peer of *Summer Excursions*
 - **Priority:** 1.0
- 01-winter-events/01-star-dust-snow-shoeing.markdown
 - **Article:** Star Dust Snow Shoeing
 - **Relationship:** Child of *Winter Excursions*
 - **Priority:** 1.0
- 01-winter-events/02-lunar-alpine.markdown
 - **Article:** Lunar Alpine
 - **Relationship:** Child of *Winter Excursions*
 - **Priority:** 2.0
- 02-summer-events/00-summer-excursions-intro.markdown
 - **Article:** Summer Excursions
 - **Relationship:** Peer of *Winter Excursions*
 - **Priority:** 2.0
- 02-summer-events/01-lunar-rock-scrambling.markdown
 - **Article:** Lunar Rock Scrambling
 - **Relationship:** Child of *Summer Excursions*
 - **Priority:** 1.0

- 02-summer-events/02-extra-terrestrial-mountain-biking.markdown
 - **Article:** Extra Terrestrial Mountain Biking
 - **Relationship:** Child of *Summer Excursions*
 - **Priority:** 2.0
- 02-summer-events/03-summer-olympics/00-lunar-olympics-intro.markdown
 - **Article:** Lunar Olympics
 - **Relationship:** Child of *Summer Excursions*
 - **Priority:** 3.0
- 02-summer-events/03-summer-olympics/01-zero-gravity-diving.markdown
 - **Article:** Zero Gravity Diving
 - **Relationship:** Grandchild of *Summer Excursions*
 - **Relationship:** Child of *Opening Ceremonies*
 - **Priority:** 1.0

ZIP files must meet the following requirements:

- Each ZIP file must end in the suffix `.zip`.
- Each ZIP file must contain at least one Markdown source file, optionally organized in folders.
- All referenced image files must be in a folder named `images` in the ZIP file's root.
- Image files must be in a supported format and must use the appropriate file extensions. Supported extensions are `.bmp`, `.gif`, `.jpeg`, `.jpg`, and `.png`. They're specified via an app system setting. For details, see Knowledge Base System Settings.

Once you have your article ZIP file, it's time to import it. Follow these steps to import your ZIP file:

1. In the Menu () , navigate to *Site Administration* (the menu for your site) → *Content* → *Knowledge Base* → *Articles*.
2. Click *Add* () → *Import* to bring up the importer page.
3. Browse to the location of your file, and in most cases leave the checkbox for the article priorities checked, and then click *Save*.

Your file is uploaded, and the importer converts each source file's Markdown text to HTML, applying the HTML to the resulting article. Any image files that are referenced in an article and included in the ZIP file are imported as attachments to the article.

In addition to source files and images, you can configure a base source URL system setting for the importer that specifies your source file's repository location. Each article's *Edit on GitHub* button (if enabled) takes the user to the source location. The importer prefixes each file's path with the base source URL. This constructs a URL to the article's repository source location; it looks like `[base URL]/[article file path]`. Here's an example base source URL:

```
https://github.com/liferay/liferay-docs/blob/master/develop/tutorials
```

The source URL constructed from this base URL and article source file `folder-1/some-article.markdown` would be:

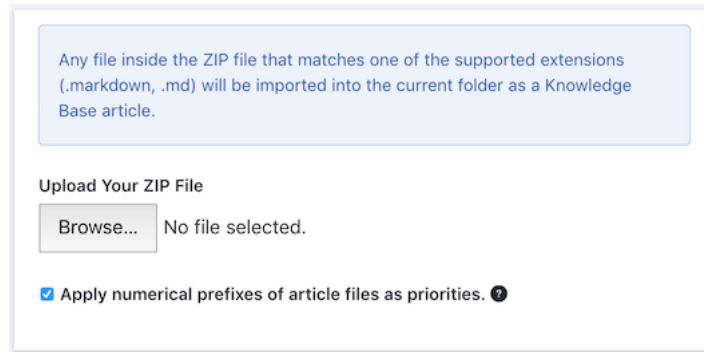


Figure 232.1: Selecting *Add* → *Import* in Knowledge Base brings up the interface for selecting a ZIP file of Markdown source files and images to produce and update articles in your Knowledge Base.

<https://github.com/liferay/liferay-docs/blob/master/develop/tutorials/folder-1/some-article.markdown>

You specify the base source URL in a file called `.METADATA` in the ZIP file's root folder. The importer treats the `.METADATA` file as a standard Java properties file and uses the base source URL to construct the source URL for all of the ZIP file's resulting articles.

To use the source URL feature, your administrator must enable it via the Knowledge Base System Settings.

KNOWLEDGE BASE IMPORTER FAQs

- **What happens when I import an existing article?**

The importer checks if the source file's leading header ID (e.g., # Some Heading [](id=some-heading)) matches the URL title of any existing article in the Knowledge Base folder. If a match is found, the importer replaces the article's content with the incoming content converted from the source file. If no match is found, a new article is created.

- **Do I need to import all of a Knowledge Base folder's articles, even if I only want to create a new article or update a subset of the folder's current articles?**

No. You can import as many or as few new or modified articles as you like.

- **Does the importer remove articles?**

No. The importer only creates and updates articles. It doesn't delete any existing articles. To delete an article, you must manually do so via the Knowledge Base app.

- **Can I update an article's priority?**

Yes. You can use the file/folder prefix convention and re-import the article to update its priority.

- **If I change an article's title, should I also change its header ID?**

It depends on whether you've already published your article. If it hasn't been published, then there are no public links to it, so it's fine to change the header ID. If the article is already published, you must decide whether it's worth breaking existing links to the article, and worth having search engines rediscover and re-rank your article based on its new friendly URL. The new friendly URL is based on the new header ID.

KNOWLEDGE BASE SYSTEM SETTINGS

Administrators can use the System Settings UI to set the Knowledge Base's global configuration (across sites). You can access this UI in *Control Panel* → *Configuration* → *System Settings* → *Knowledge Base*. There are five sections of Knowledge Base configuration settings:

- Service
- Knowledge Base Article
- Knowledge Base Display
- Knowledge Base Search
- Knowledge Base Section

The *Service* section's settings apply defaults to all the Knowledge Base widgets, and to the Knowledge Base app in Site Administration. The other sections apply to specific Knowledge Base widgets and override the *Service* defaults.

Important: Advanced configuration of the Knowledge Base application's system settings should only be performed by an Liferay DXP administrator.

The Knowledge Base has several optional features that are disabled by default, but can be enabled and configured from System Settings. These include source URL, import file conventions, new article priority increment, and sections.

234.1 Source URL Settings

The source URL settings define the source location of Markdown files for import. This should point to a source repository where the files are stored. GitHub is assumed as the default. Once defined, the Knowledge Base displays a button (default label is *Edit on GitHub*) above each displayed article. Users can click the button to navigate to an article's source location.

The source URL settings are accessible in the *Service* section of the Knowledge Base's System Settings. To enable the source URL, check the *Source URL Enabled* checkbox.

To change the source URL button's label, specify a new value for the setting *Source URL Edit Message Key*. Best practice is to specify the value as a language key. For example, if you create a language key `edit-on-bitbucket=Edit on Bitbucket`, you can specify that language key as the button's new label:

```
edit-on-bitbucket
```

Alternatively, you can specify the label explicitly:

```
Edit on Bitbucket
```

234.2 Importer File Convention Settings

These settings define the supported file extensions, the suffix for parent files, and the image folder's path within the imported ZIP files. These settings are accessible in the *Service* section of the Knowledge Base's System Settings.

The following settings specify the importer's supported file extensions:

Markdown Importer article extensions: Sets the supported article extensions. The default values are `.markdown` and `.md`.

Markdown Importer Image File Extensions: Sets the supported image file extensions. The default values are `.bmp`, `.gif`, `.jpeg`, `.jpg`, and `.png`.

Follow these steps to modify the supported file extensions:

1. Click the + or - button next to the setting to add or remove a supported file extension, respectively.
2. If adding an extension, enter a new value.
3. Click *Save*.

These settings define additional article configuration options for the importer:

Markdown Importer Article Intro: Sets the parent article's file suffix. The default value is `intro.markdown`.

Markdown Importer Image Folder: Sets the image folder path the importer looks for in the ZIP file. The default path is `/images`.

Article Increment Priority Enabled: Whether to increment new article priorities by 1.0. To disable this increment so that articles get a flat value of 1.0, deselect the checkbox.

Alternatively, you can enable or disable the article increment priority feature for each widget in the corresponding widget's configuration menu in System Settings.

234.3 Section Names Setting

The section names setting lets you specify names of arbitrary topics to attribute to articles. Using the Knowledge Base Section widget, you can display one or more *sections* (groups) of articles. To use sections, you must first define them in the System Settings for the *Knowledge Base Section* widget.

Follow these steps to make new sections available:

1. Navigate to the *Knowledge Base Section* configuration menu.
2. Click the plus button next to the *Admin Knowledge Base Article Sections* setting to add a new field for each section you want.
3. Enter a name for each new section and click *Update*.



Figure 234.1: Create the sections you want to use with the Knowledge Base Section widget.

Once your sections are added, you can follow the steps in the Knowledge Base Section documentation to learn how to use them.

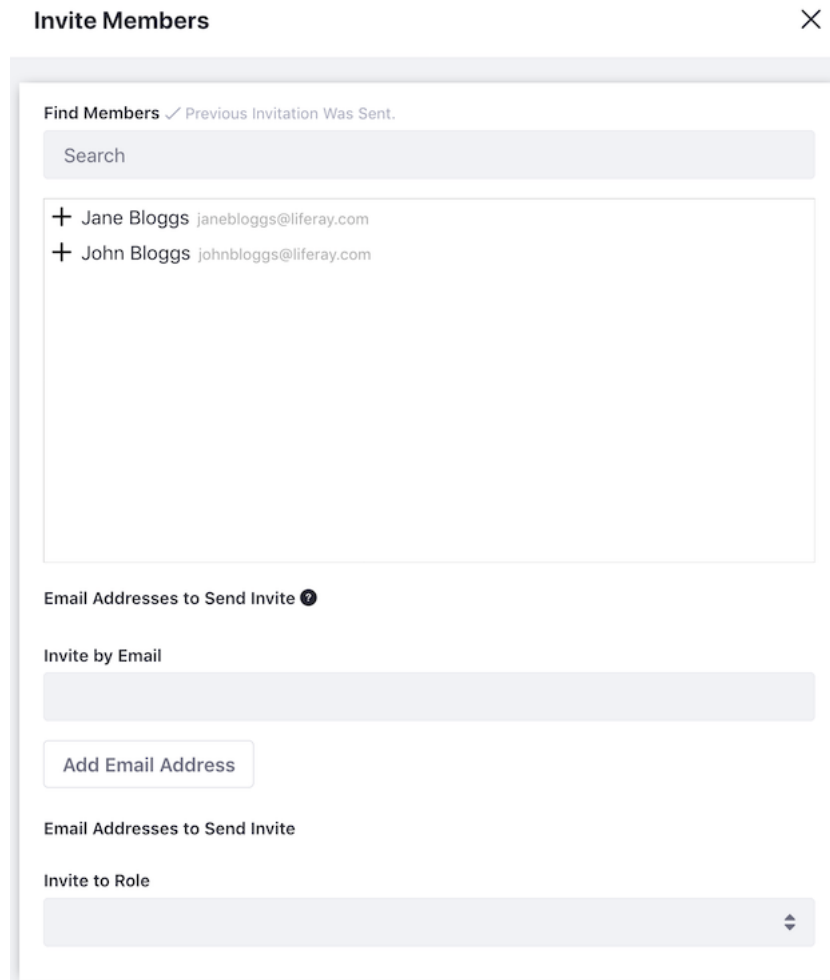
INVITING MEMBERS TO YOUR SITE

The Invite Members widget lets site administrators send invitations to join the Site. You can add this widget to a page from the menu *Add* (+) → *Widgets* → *Collaboration*. Click the *Invite members* button to bring up the interface for sending invitations.

Click the plus sign next to a User or click the *Add Email Address* button to add a User to the invite list. Users that have already been invited but have not yet responded appear with a check mark next to their names. You can also invite Users to the *Site Owner*, *Site Content Reviewer*, and *Site Administrator* Roles for your site by selecting that Role under the *Invite to Role* heading. Once you've added all the Users you want to invite and have selected their Roles, click the *Send Invitations* button to invite them. For more information on roles, see the Roles and Permissions documentation.

The Site invitation shows up under the *Requests List* tab on the User's *Notifications* page. The User can then choose to *Confirm* or *Ignore* the invitation.

When Users confirm such invitations, they become Site members assigned to the Roles you defined.



Invite Members ✕

Find Members ✓ Previous Invitation Was Sent.

Search

+ Jane Bloggs janebloggs@liferay.com

+ John Bloggs johnbloggs@liferay.com

Email Addresses to Send Invite ⓘ

Invite by Email

Add Email Address

Email Addresses to Send Invite

Invite to Role

Figure 235.1: You can invite users by clicking the add sign next to the user's name.

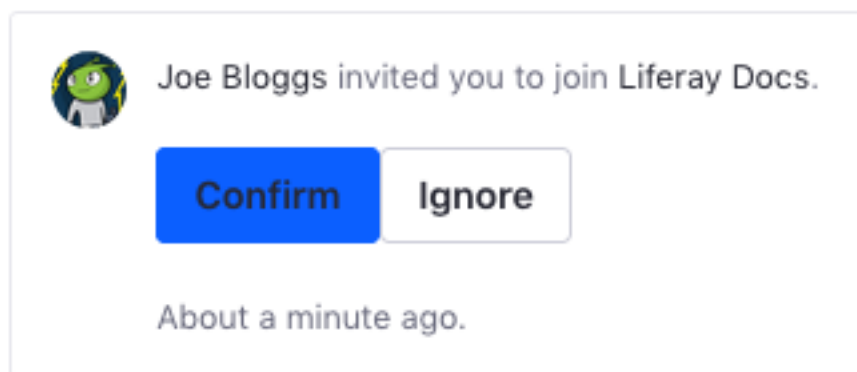


Figure 235.2: You can confirm or ignore the invitation.

ASSETS

When you think of assets, you probably think of things like houses, cars, money, or gold bricks. Although Liferay DXP is capable of holding many assets, they aren't financial assets. Assets in Liferay DXP are any kind of content. For example, images, documents, and web content are all assets. And although such assets may not have the tangible value of a gold brick, they could equal or even exceed that value when used to support an organization's day-to-day business activities. You might even think of Liferay DXP as your very own Fort Knox (don't worry, we won't tell anyone). Here, you'll learn to configure Liferay DXP to apply tags to assets automatically.




Figure 236.1: The tags *freight car* and *electric locomotive* were automatically applied to this image.

CONFIGURING ASSET AUTO TAGGING

Tagging assets is a great way to organize content. Typically, the content creator applies tags while creating the content. It's also possible, however, to tag content automatically. For example, Liferay DXP can scan an image on upload and apply tags that describe the image's content. This lets you leverage tags without requiring content creators to apply them manually.

Note: Auto-tagging currently works only for images, text-based documents, text-based web content, and blog entries.

Here, you'll learn how to configure asset auto-tagging in general. This is required prior to configuring auto-tagging for specific asset types, which is documented separately:

- Auto Tagging Images
- Auto Tagging Text

237.1 Configuration Levels

Auto-tagging is enabled by default. You can configure it at three levels:

Global (System): For auto-tagging to function on any level, it must be enabled globally. You can also set the default auto-tagging configuration for every portal instance.

Instance: When enabled globally, auto-tagging is also enabled by default for each portal instance. However, you can override the global auto-tagging configuration on a per-instance basis.

Site: When enabled for an instance, auto-tagging is also enabled by default for all that instance's sites. You can disable it for specific sites.

237.2 Global Configuration

Follow these steps to configure auto tagging globally:

1. Go to *Control Panel* → *Configuration* → *System Settings* → *Assets*.
2. Under *SYSTEM SCOPE*, select *Asset Auto Tagging*.
3. The following settings are available:

Enable Auto Tagging of Assets: Whether asset auto tagging is enabled.

Maximum Number of Tags: The maximum number of tags that can be automatically applied to each asset. The default value of 0 means that there is no limit.

4. Click *Save* to save your changes.

To set the default auto-tagging configuration for all instances, select *Asset Auto Tagging* under *VIRTUAL INSTANCE SCOPE*. The available settings are exactly the same as those in the *SYSTEM SCOPE*.




Figure 237.1: You can configure auto tagging globally in the Assets section of System Settings.

237.3 Instance-level Configuration

When enabled globally, auto-tagging is also enabled by default for each instance. You can, however, disable or configure it for each instance.

Follow these steps to configure auto tagging on the instance level:

1. Go to *Control Panel* → *Configuration* → *Instance Settings* → *Assets*.
2. Under *VIRTUAL INSTANCE SCOPE*, select *Asset Auto Tagging*.
3. The settings here are identical to those in the global configuration, but apply only to the current instance.
4. Click *Save* to save your changes.

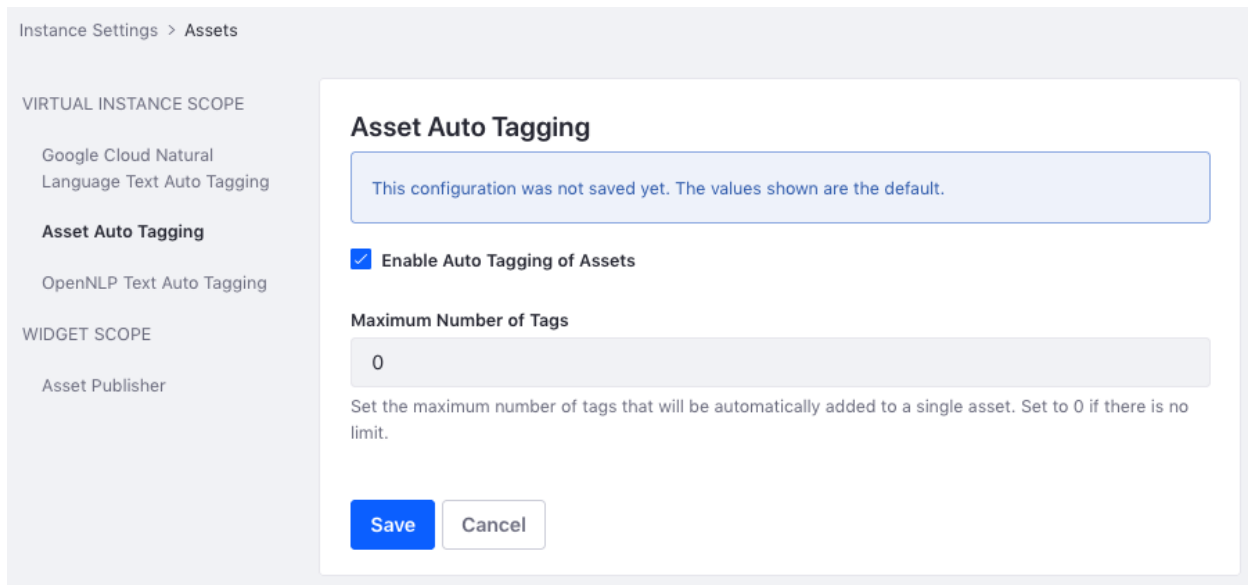



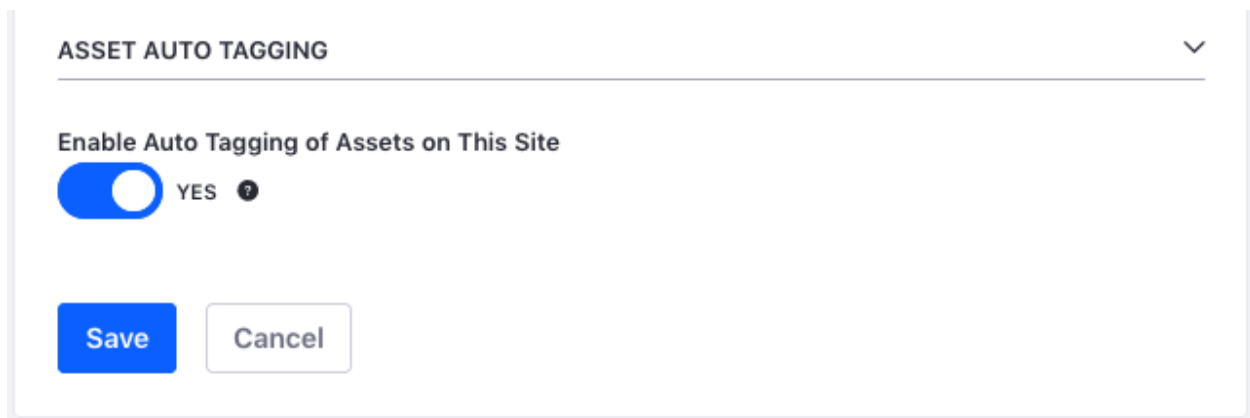
Figure 237.2: You can also configure auto tagging for each instance.

237.4 Site-level Configuration

When enabled for an instance, auto-tagging is also enabled by default for all that instance's sites. You can, however, enable or disable it for each site.

Follow these steps to configure auto tagging for a site:

1. Open the *Menu* () , click your site's name, and navigate to *Configuration* → *Settings*.
2. In the *General* tab, expand the *Asset Auto Tagging* section. Use the toggle to enable or disable auto tagging for the site.
3. Click *Save* to save your changes.



ASSET AUTO TAGGING ▼

Enable Auto Tagging of Assets on This Site

YES ⓘ

Save Cancel

Figure 237.3: You can enable or disable auto-tagging for a site.

AUTO TAGGING IMAGES

Tags help you find and organize files, including images. With asset auto tagging enabled, you can also enable image auto tagging. Image auto tagging automatically tags images uploaded to the Documents and Media Library. This lets you use tags without requiring anyone to apply them manually.

Note: Currently, tags applied automatically are English only.

Image auto tagging is disabled by default. To use it, you must do two things:

1. Ensure that asset auto tagging is enabled. Although it's enabled by default, administrators can disable it.
2. Ensure that an image auto tagging provider is enabled. These providers are disabled by default. Here, you'll learn how to enable/disable them.

Note: Prior to Liferay DXP Fix Pack 1 and Liferay Portal CE GA2, you must configure the providers shown here in *Documents and Media* instead of *Assets* (in System/Instance Settings).

Three such providers are available:

TensorFlow: An open-source library that provides machine learning capabilities. TensorFlow image auto-tagging in Liferay DXP is based on TensorFlow's LabelImage sample for Java, and uses the Inception5h model. Use this with caution, since its accuracy is limited.

Google Cloud Vision: Uses the Google Cloud Vision API to automatically tag images.

Microsoft Cognitive Services: Uses Microsoft Cognitive Services to automatically tag images.

238.1 Configuring TensorFlow Image Auto Tagging

Follow these steps to configure TensorFlow Image Auto Tagging:




Figure 238.1: The tags *freight car* and *electric locomotive* were automatically applied to this image.

1. Go to *Control Panel* → *Configuration* → *System Settings* → *Assets*.
2. Under *VIRTUAL INSTANCE SCOPE*, select *TensorFlow Image Auto Tagging*. The following settings are available:

Enable TensorFlow Image Auto Tagging: Check this box to select whether image auto-tagging is enabled by default in any instance that has asset auto tagging enabled. Note that you can override this value for specific instances, as the next section shows.

Confidence Threshold: TensorFlow assigns a confidence level between 0 and 1 for each tag, where 1 is the highest confidence and 0 is the lowest. This field sets the minimum confidence level that TensorFlow needs to apply a tag. Higher values yield fewer tags because TensorFlow needs more confidence before it applies a tag. Likewise, lower values yield more tags.

3. Click *Save* to save your changes.

You can override these settings for each instance from *Control Panel* → *Configuration* → *Instance Settings* → *Assets*.

To optimize performance, you can also control the process that runs TensorFlow image auto tagging:

1. Go to *Control Panel* → *Configuration* → *System Settings* → *Assets*.

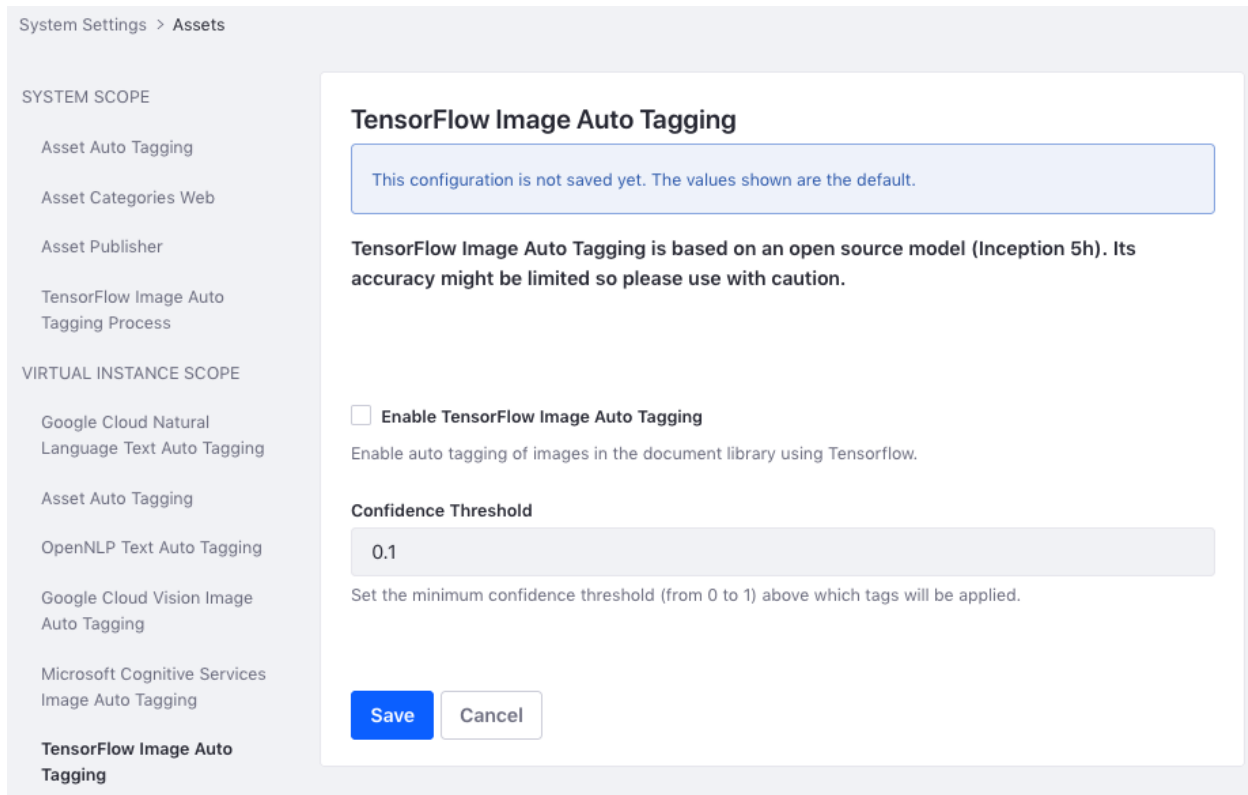


Figure 238.2: Configure TensorFlow image auto-tagging for your portal instances.

2. Under *SYSTEM SCOPE*, select *TensorFlow Image Auto Tagging Process*. The following settings are available:
 - Maximum Number of Relaunches:** The maximum number of times the process is allowed to crash before it is permanently disabled.
 - Maximum Number of Relaunches Time Interval:** The time in seconds after which the counter is reset.
3. Click *Save* to save your changes.

238.2 Configuring Google Cloud Vision

Follow these steps to configure Google Cloud Vision image auto tagging:

1. Go to *Control Panel* → *Configuration* → *System Settings* → *Assets*.
2. Under *VIRTUAL INSTANCE SCOPE*, select *Google Cloud Vision Image Auto Tagging*. The following settings are available:

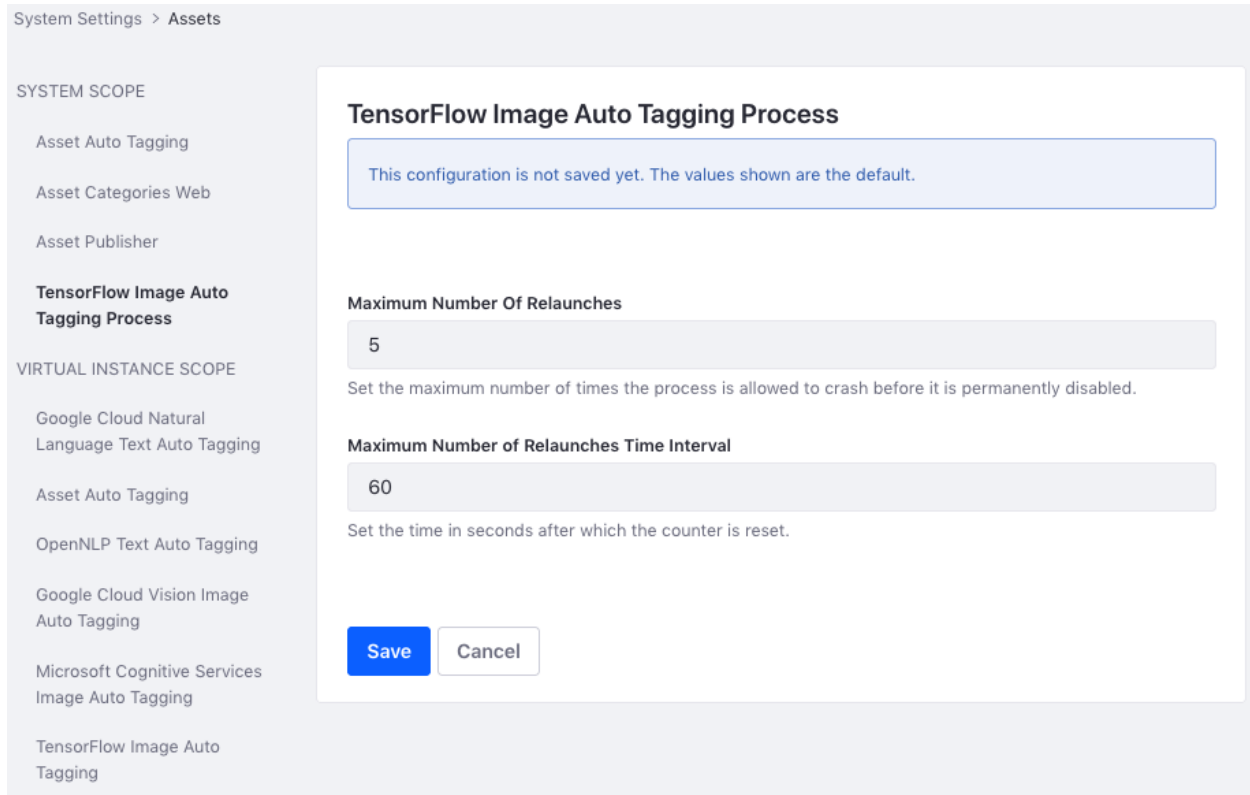


Figure 238.3: You can fine tune the process that runs the TensorFlow image auto tagging in the portal.

API Key: The API key to use for the Google Cloud Vision API. For more information, see Google’s documentation on API keys.

Enabled: Whether Google Cloud Vision image auto tagging is enabled.

3. Click *Save* to save your changes.

You can override these settings for each instance from *Control Panel* → *Configuration* → *Instance Settings* → *Assets*.

238.3 Configuring Microsoft Cognitive Services

Follow these steps to configure Microsoft Cognitive Services image auto tagging:

1. Go to *Control Panel* → *Configuration* → *System Settings* → *Assets*.
2. Under *VIRTUAL INSTANCE SCOPE*, select *Microsoft Cognitive Services Image Auto Tagging*. The following settings are available:

API Key: Your API key for the Computer Vision API V2.

The screenshot shows a web interface for configuring 'Google Cloud Vision Image Auto Tagging'. On the left is a sidebar with a navigation menu. The main content area is titled 'Google Cloud Vision Image Auto Tagging' and contains a message box stating 'This configuration is not saved yet. The values shown are the default.' Below this is a section 'Auto tags images using the Google Cloud Vision API.' with a sub-section 'API Key *' containing a large text input field. A note below the field says 'Set the API key for the Google Cloud Vision API. Billing must be enabled for this API to work. See <https://cloud.google.com/docs/authentication/api-keys> for more information.' There is an 'Enabled' checkbox which is currently unchecked. Below the checkbox is the text 'Enable auto tagging of images in the document library using the Google Cloud Vision API.' At the bottom of the form are 'Save' and 'Cancel' buttons.

System Settings > Assets

SYSTEM SCOPE

- Asset Auto Tagging
- Asset Categories Web
- Asset Publisher
- TensorFlow Image Auto Tagging Process

VIRTUAL INSTANCE SCOPE

- Google Cloud Natural Language Text Auto Tagging
- Asset Auto Tagging
- OpenNLP Text Auto Tagging
- Google Cloud Vision Image Auto Tagging**
- Microsoft Cognitive Services Image Auto Tagging
- TensorFlow Image Auto Tagging

WIDGET SCOPE

- Asset Publisher

Google Cloud Vision Image Auto Tagging

This configuration is not saved yet. The values shown are the default.

Auto tags images using the Google Cloud Vision API.

API Key *

Set the API key for the Google Cloud Vision API. Billing must be enabled for this API to work. See <https://cloud.google.com/docs/authentication/api-keys> for more information.

Enabled

Enable auto tagging of images in the document library using the Google Cloud Vision API.

Save **Cancel**

Figure 238.4: The Google Cloud Vision provider requires an API key.

API Endpoint: The endpoint for the Computer Vision API V2 (e.g., <https://westcentralus.api.cognitive.microsoft.com>)

Enabled: Whether Microsoft Cognitive Services image auto tagging is enabled.

For more information, see the Microsoft Cognitive Services documentation.

3. Click *Save* to save your changes.

You can override these settings for each instance from *Control Panel* → *Configuration* → *Instance Settings* → *Assets*.

System Settings > Assets

SYSTEM SCOPE

- Asset Auto Tagging
- Asset Categories Web
- Asset Publisher
- TensorFlow Image Auto Tagging Process

VIRTUAL INSTANCE SCOPE

- Google Cloud Natural Language Text Auto Tagging
- Asset Auto Tagging
- OpenNLP Text Auto Tagging
- Google Cloud Vision Image Auto Tagging
- Microsoft Cognitive Services Image Auto Tagging**
- TensorFlow Image Auto Tagging

WIDGET SCOPE

- Asset Publisher

Microsoft Cognitive Services Image Auto Tagging

This configuration is not saved yet. The values shown are the default.

Auto tags images using the Microsoft Cognitive Services Computer Vision API V2.

API Key *

Set the API key for the Computer Vision API V2. See <https://azure.microsoft.com/en-us/try/cognitive-services/my-apis/?apiSlug=computer-services> for more information.

API Endpoint *

Set the endpoint for the Computer Vision API V2 (e.g. <https://westcentralus.api.cognitive.microsoft.com/vision/v2.0>). See <https://azure.microsoft.com/en-us/try/cognitive-services/my-apis/?apiSlug=computer-services> for more information.

Enabled

Enable auto tagging of images in the document library using Microsoft Cognitive Services.

Save **Cancel**

Figure 238.5: The Microsoft Cognitive Services provider requires an API key and an endpoint.

AUTO TAGGING TEXT

With asset auto tagging enabled, you can also configure text auto tagging. Text auto tagging automatically tags text-based assets. This lets you use tags without requiring anyone to manually apply them.

Note: Currently, text auto tagging is only available for text-based documents, text-based web content, and blog entries. Tags applied automatically are English only.

Text auto tagging is disabled by default. To use it, you must enable it:

1. Ensure that asset auto tagging is enabled. Although it's enabled by default, administrators can disable it.
2. Ensure that a text auto tagging provider is configured and enabled for the asset types you want to auto tag. You'll learn how to do this here. Note that these providers aren't configured or enabled by default.

Note: Prior to Liferay DXP Fix Pack 1 and Liferay Portal CE GA2, you must enable these providers separately for each content type in System/Instance Settings. For example, you must enable text auto tagging for documents and web content in *Documents and Media* and *Web Content*, respectively.

There are two text auto-tagging providers in the portal:

Google Cloud Natural Language Text Auto Tagging: Uses the Google Cloud Natural Language API to analyze and automatically tag portal content.

OpenNLP Text Auto Tagging: Uses the open source Apache OpenNLP library to analyze and automatically tag portal content. Three models are used: location name finder, organization finder, and person name finder. Use this provider with caution, as its accuracy may be limited.

239.1 Configuring Google Cloud Natural Language Text Auto Tagging

Follow these steps to configure the auto-tagging provider for the Google Cloud Natural Language API:

1. Go to *Control Panel* → *Configuration* → *System Settings* → *Assets*.
2. Under *VIRTUAL INSTANCE SCOPE*, select *Google Cloud Natural Language Text Auto Tagging*. The following settings are available:



API Key: The API key to use for the Google Cloud Natural Language API. For more information, see Google's documentation on API keys.

Classification Endpoint Enabled: Whether to enable auto tagging of text using the Google Cloud Natural Language API Classification endpoint.

Confidence: Set the classifier's confidence of the category. This number represents how certain the classifier is that this category represents the given text.

Entity Endpoint Enabled: Whether to enable auto tagging of text using the Google Cloud Natural Language API Entity endpoint.

Saliency: The saliency score for an entity provides information about the importance or centrality of that entity to the entire text.

Enable Google Cloud Natural Language Text Auto Tagging For: The asset types to enable text auto tagging for. Use the menu to select *Document*, *Blogs Entry*, or *Web Content Article*. To add multiple asset types, click the *Add* icon () and select the asset type from the menu. You can delete any additional asset types by clicking the *Trash* icon ()

3. Click *Save* to save your changes.

You can override these settings for each instance from *Control Panel* → *Configuration* → *Instance Settings* → *Assets*.

239.2 Configuring OpenNLP Text Auto Tagging

Follow these steps to configure the OpenNLP Text Auto Tagging provider:

1. Go to *Control Panel* → *Configuration* → *System Settings* → *Assets*.
2. Under *VIRTUAL INSTANCE SCOPE*, select *OpenNLP Text Auto Tagging*. The following settings are available:

Confidence Threshold: Set the minimum confidence threshold (from 0 to 1, where 1 is the highest confidence) above which tags will be applied. Higher values yield fewer tags because the provider needs more confidence before it applies a tag. Likewise, lower values yield more tags.

Enable OpenNLP Text Auto Tagging For: The asset types to enable text auto tagging for. Use the menu to select *Document*, *Blogs Entry*, or *Web Content Article*. To add multiple asset types,

Google Cloud Natural Language Text Auto Tagging

This configuration is not saved yet. The values shown are the default.

Auto tag of text using the Cloud Natural Language API.

API Key *

Set the API key for the Google Cloud Natural Language API. Billing must be enabled for this API to work. See <https://cloud.google.com/docs/authentication/api-keys> for more information.

Classification Endpoint Enabled
Enable auto tagging of text using the Cloud Natural Language API Classification endpoint.

Confidence
0.5
Set the classifier's confidence of the category. This number represents how certain the classifier is that this category represents the given text.

Entity Endpoint Enabled
Enable auto tagging of text using the Cloud Natural Language API Entity endpoint.

Salience
0.02
The salience score for an entity provides information about the importance or centrality of that entity to the entire text.

Enable Google Cloud Natural Language Text Auto Tagging For +



Document

Enable Google Cloud Natural Language Text Auto Tagging For +

Choose an Option

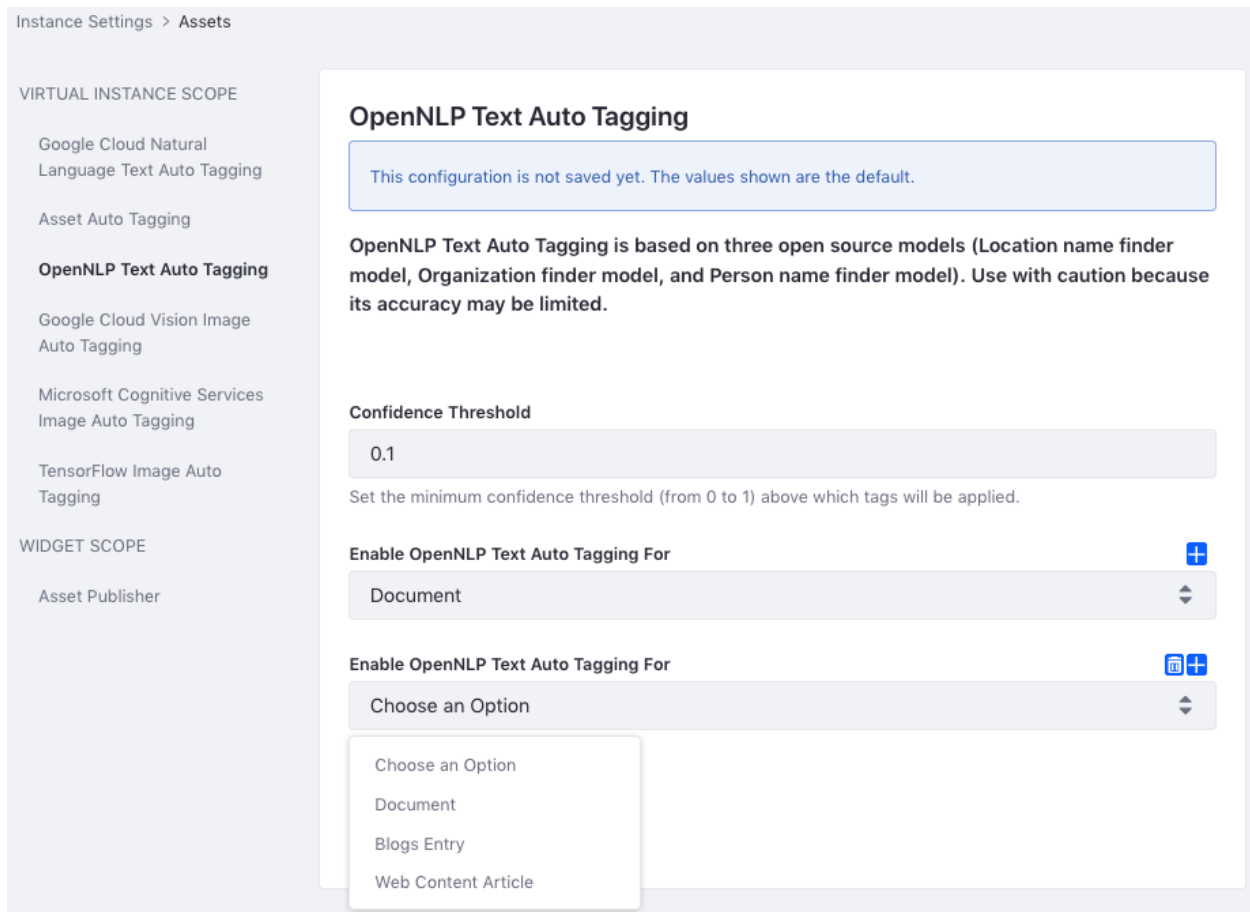
- Choose an Option
- Document
- Blogs Entry
- Web Content Article

Figure 239.1: Configure Google Cloud Natural Language text auto tagging for your portal instances.

click the *Add* icon () and select the asset type from the menu. You can delete any additional asset types by clicking the *Trash* icon ().

3. Click *Save* to save your changes.

You can override these settings for each instance from *Control Panel* → *Configuration* → *Instance Settings* → *Assets*.



Instance Settings > Assets

VIRTUAL INSTANCE SCOPE

- Google Cloud Natural Language Text Auto Tagging
- Asset Auto Tagging
- OpenNLP Text Auto Tagging**
- Google Cloud Vision Image Auto Tagging
- Microsoft Cognitive Services Image Auto Tagging
- TensorFlow Image Auto Tagging

WIDGET SCOPE

- Asset Publisher

OpenNLP Text Auto Tagging


This configuration is not saved yet. The values shown are the default.


OpenNLP Text Auto Tagging is based on three open source models (Location name finder model, Organization finder model, and Person name finder model). Use with caution because its accuracy may be limited.



Confidence Threshold


0.1

Set the minimum confidence threshold (from 0 to 1) above which tags will be applied.

Enable OpenNLP Text Auto Tagging For 

Document 

Enable OpenNLP Text Auto Tagging For  

Choose an Option 

- Choose an Option
- Document
- Blogs Entry
- Web Content Article

Figure 239.2: Configure OpenNLP text auto tagging for your portal instances.

CREATING A SOCIAL NETWORK

Liferay DXP contains several features and widgets for leveraging its social framework. The Activities widget lets you broadcast user activities on a Site. This is a good way for Site members to see what's going on in their communities. When placed on a user's private Dashboard, the Activities widget displays the activities of that user's social connections. Users can make those social connections via the Contacts Center app, which lets them establish connections and followers throughout the portal. What's more, widgets can be exported as OpenSocial gadgets and/or used with Facebook.

This guide shows you how to do these things, and more.

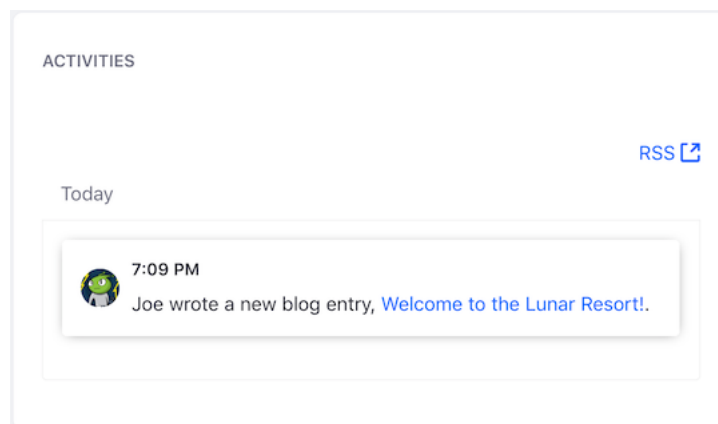


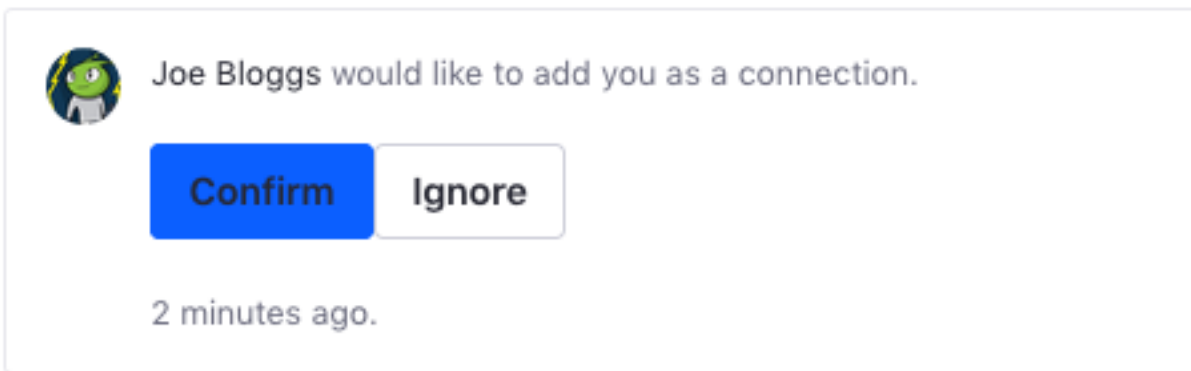
Figure 240.1: The Activities widget shows information about asset-related user activity in the current Site.

NOTIFICATIONS

Notifications List (0)

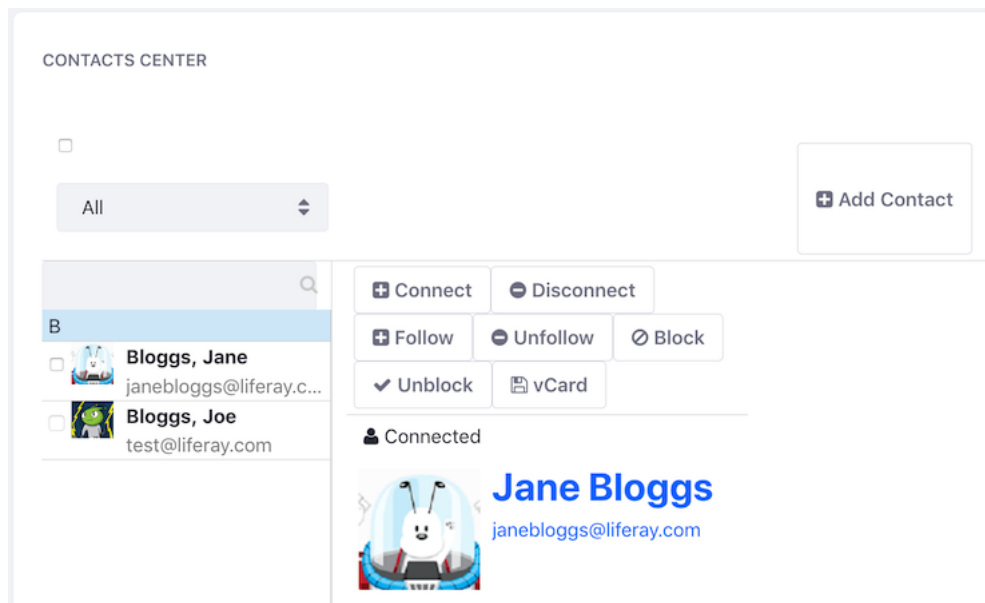
Requests List (1)

Filter and Order ▾



A notification card with a rounded border. On the left is a circular profile picture of a green alien. To its right, the text reads "Joe Bloggs would like to add you as a connection." Below this text are two buttons: a blue "Confirm" button and a white "Ignore" button with a grey border. At the bottom of the card, it says "2 minutes ago."

Figure 240.2: Users get a notification that lets them respond to connection requests.



The "CONTACTS CENTER" widget features a search bar at the top left and a filter dropdown menu currently set to "All". On the right side of the top section is a button with a plus sign and the text "Add Contact". Below the search bar is a list of contacts, with the letter "B" indicating the current filter. Two contacts are visible: "Bloggs, Jane" with email "janebloggs@liferay.c..." and "Bloggs, Joe" with email "test@liferay.com". To the right of the list is a panel of action buttons: "Connect" and "Disconnect" (disabled), "Follow", "Unfollow", and "Block" (disabled), "Unblock" (checked), and "vCard". Below these buttons, it shows a "Connected" status with a person icon. At the bottom right of the widget is a profile card for "Jane Bloggs" with her profile picture and email "janebloggs@liferay.com".

Figure 240.3: The Contacts Center widget lets users make connections.

USING THE ACTIVITIES WIDGET

The core social widget is Activities. It displays information about user activity on the Site where you added it. User activities tracked include updates to the Documents and Media library, blog posts, message boards posts, wiki pages, and bookmarks. Liferay DXP also tracks information about web content but only displays this information if the logged-in user is a Site administrator. This widget provides a summary of recent Site activity. You can use it on a Site's public or private pages to show what Site members have been up to, or you can use it on the public or private pages of a user's personal Site. When added to a personal Site, the Activities widget shows the activities of only that user.

Add the Activities widget to a page from the *Add* (+) → *Widgets* → *Social* menu.

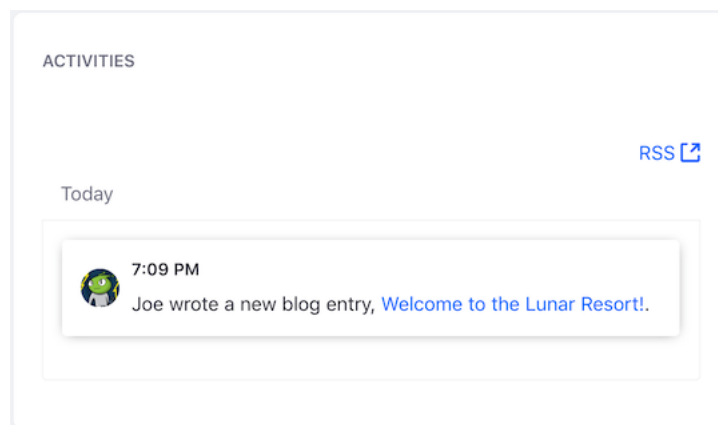


Figure 241.1: The Activities widget shows information about asset-related user activity in the current Site.

Note that the widget provides links to the assets listed in the feed. These links don't work, however, unless there's a way to display the assets on the page. For example, if you click a link to a blog post in the Activities widget, that page must have a Blogs widget to display that blog post.

CONNECTING USERS

By adding the Contacts Center and Activities widgets to users' private Dashboards, administrators can let users form social connections. The Contacts Center widget lets users form these connection types:

Connection: A two-way relationship formed by a user accepting a connection request from another user. The Activities widget on each user's private Dashboard displays the activities of the other user.

Following: A one-way relationship in which one user follows another user. The followed user's activities appear in the Activities widget on the follower's private Dashboard.

To ensure that all users have a page on their private Dashboard that contains the Contacts Center and Activities widgets, you can create a user group and then create a page via a Site Template. For instructions on this, see the user groups documentation. The widgets are in *Add* (+) → *Widgets* → *Social*.




Figure 242.1: The Contacts Center widget lets users make connections.

When you select a user in the Contacts Center, these buttons appear across the top of the widget:

Connect: Send a connection request to the user.

Disconnect: Disconnect from the user.

Follow: Follow the user.

Unfollow: Stop following the user.

Block: Block the user. Blocking a user only prevents that user from following you or adding you as a connection. A blocked user can still send messages to and view the public profile information of the blocking user.

Unblock: Stop blocking the user.

vCard: Export the user's vCard and save it as a VCF file. vCard is a file format standard for electronic business cards.

When you send a connection request, the user is notified and can confirm or ignore the request.


NOTIFICATIONS

Notifications List (0)

Requests List (1)

Filter and Order ▼





Joe Bloggs would like to add you as a connection.

Confirm Ignore

2 minutes ago.

Figure 242.2: Users get a notification that lets them respond to connection requests.



EXPORTING WIDGETS TO OTHER WEBSITES

You can publish its widgets to other websites via OpenSocial. This lets you provide your widget or content in the context of another website. Read on to find out how this is done.

243.1 Sharing OpenSocial Gadgets

OpenSocial consists of a set of APIs for social networking. It may be beneficial for you to share widgets from your server with other sites, such as ighome.com or igoogportal.com. These sites let users customize their own pages with gadgets. Users can share Liferay DXP widgets on any OpenSocial-compatible site. Liferay DXP does this by sharing the widget's URL with the OpenSocial platform. This URL is unique to the user's specific widget instance. Users can therefore share multiple instances of the same widget as different OpenSocial gadgets.

Follow these steps to share a widget with OpenSocial:

1. Add the widget you want to share to a site page.
2. Click the widget's *Options* icon () and select *Configuration*.
3. In the *Sharing* tab, open the *OpenSocial Gadget* section.
4. Set the selector to *YES* for *Allow users to add [widget name] to an OpenSocial platform*. In the *OpenSocial Gadget URL* field, replace `localhost:8080` with the name of your public domain and port.
5. Click *Save*.
6. Close the dialog and click the widget's *Options* icon (). There's a new option available named *Add to an OpenSocial Platform*. Select this option to add the widget to a page on an OpenSocial platform.

OPENSOCIAL GADGET ∨

Use the OpenSocial gadget URL to create an OpenSocial gadget.

OpenSocial Gadget URL

`http://localhost:8080/google_gadget/web/guest/home/-/blogs`

Allow users to add Blogs to an OpenSocial platform.

NO

Figure 243.1: You can share widgets via OpenSocial.

INTEGRATING WITH FACEBOOK

Liferay DXP provides tools for integrating your portal and its content with Facebook. For example, you can use Facebook for authentication and even export widgets as Facebook applications. This article shows you how.

244.1 Facebook Sign On

Like many websites you may visit, any site running on Liferay DXP can use Facebook for sign in. This makes it easier for users to sign in to your Site, since they don't need to remember another user name and password.

244.2 Using Your Widget as Facebook Applications

You can add any Liferay DXP widget as an app on Facebook. To do this, you must first get a developer key. A link for doing this is provided to you in the Facebook tab in any widget's Configuration screen. You must create the app on Facebook and get the key and canvas page URL from Facebook. You can then copy and paste their values into the Facebook tab. Once you do that, your widget is available on Facebook.

This integration lets you make things like Message Boards, Calendars, Wikis, and other content on your website available to a much larger audience (unless you already have a billion users on your site, in which case, kudos to you). If you're a developer, you can implement your widget on Liferay DXP and then publish it to Facebook.

USING SOCIAL BOOKMARKS

Social bookmarks appear below content as buttons for sharing that content on social networks. For example, social bookmarks appear by default in the Blogs widget below each blog post. For more information on configuring social bookmarks in the Blogs widget, see the documentation on displaying blogs.

These social bookmarks are available by default:

- Twitter
- Facebook
- LinkedIn

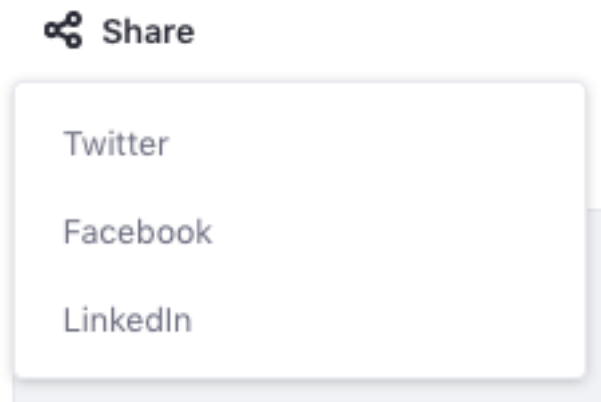


Figure 245.1: The default social bookmarks appear in a menu below content.

You can install the Social Bookmarks app from Liferay Marketplace. This app is available for Liferay CE Portal and Liferay DXP. It adds the following social bookmarks:

- AddThis
- Delicious
- Digg
- Evernote

- [Reddit](#)
- [Slashdot](#)

If you need help installing apps from Liferay Marketplace, see [Using the Liferay Marketplace](#).

SEARCH

Sites often feature lots of content split over lots of asset types. Web content articles, documents and media files, and blogs entries are just a few examples. Most content types are *assets*. Under the hood, assets use the Asset API and have an Indexer class. Any content that has these features can be searched.




Figure 246.1: The Type Facet configuration lists the searchable out-of-the-box asset types.

246.1 Elasticsearch

The default search engine is Elasticsearch, which is backed by the Lucene search library. There's an Elasticsearch server embedded in all bundles, which is handy for testing and development purposes. Production environments must install a separate, remote Elasticsearch server (or even better, cluster of servers). For information on how to install Elasticsearch, read the deployment

guide.

246.2 Search Features

Searching is simple and straightforward. Find a search bar (there's one embedded in every page by default), enter a term, and click *Enter*.




Figure 246.2: There's a search bar embedded on all pages by default.

After search is triggered, a results page appears. If there are hits to search engine documents, they appear as search results in the right hand column. In the left hand column are search facets.

5 results for **test**






| | |
|---|--|
|  | Leveraging Slow Time to Market Blogs Entry · By ziltoid omniscient On Jun 1, 2019 1:59 AM Capitalize on low hanging fruit to identify a ballpark value added activity to beta test |
|  | Capitalize on Content Blogs Entry · By test test On Jun 1, 2019 1:54 AM Capitalize on low hanging fruit to identify a ballpark value added activity to beta test |
|  | Keep Business Humming with Shiny New Office Toys Blogs Entry · By marvin smart On Jun 1, 2019 2:00 AM Capitalize on low hanging fruit to identify a ballpark value added activity to beta test |
|  | Don't Just Make Money: Make a Difference Blogs Entry · By marvin smart On Jun 1, 2019 2:01 AM Holistically predominate extensible testing procedures for reliable supply chains. |
|  | Test Test User · By test test |

Figure 246.3: Results are displayed in the Search Results portlet.

The search bar, search results, and search facets make up three powerful features in the search UI.

246.3 Search Bar

The search bar is simple: it's where you enter *search terms*. Search terms are the text you send to the search engine to match against the documents in the index.

246.4 Search Results and Relevance

The search term is processed by an algorithm in the search engine, and search results are returned to users in order of relevance. Relevance is determined by a document's *score*, generated against the search query. The higher the score, the more relevant a document is considered. The particular relevance algorithm used is dependent on algorithms provided by the search engine (Elasticsearch by default).

246.5 Search Facets

Facets allow users of the Search application to filter search results. Think of facets as buckets that hold similar search results. You might want to see the results in all the buckets, but after scanning the results, you might decide that the results of just one bucket better represent what you want. So what facets are included out of the box?

- Category
- Folder
- Site
- Tag
- Type
- User
- Modified
- Custom

You've probably used something similar on any number of sites. You search for an item, are presented with a list of results, and a list of buckets you can click to further drill down into the search results, without entering additional search terms. Search facets work the same way. Facets are, of course, configurable.




Figure 246.4: *Site* and *Type* are two of the facet sets you'll encounter. They let you drill down to results that contain the search terms you entered.

WHAT'S NEW IN SEARCH

Lots of new and improved search capabilities are present in 7.0, from new widgets to new APIs and infrastructure.

247.1 New and Improved Widgets

Add search widgets by clicking the Add () icon on the page. Then expand the Widgets → Search section.

247.2 Custom Filter

New Widget


Add a widget to the page for each of the filters you'd like applied to the search results. Let search page users see and manipulate the filters or make them invisible and/or immutable (this is just a cool word for “they can't be changed”).

For example, add a custom filter to ensure that all returned results have the keyword *street* in the content field.

247.3 Sort

New Widget

The Sort widget reorders the results based on the value of certain keyword fields in the index. For example, show results in alphabetic order of the Title field. The default order is determined by the search engine's *Relevance* calculation.

Add more fields to the sort widget if the default options aren't enough. Click the widget Options  menu → Configuration. Enter a human readable label and the fieldName to sort by. Just make sure it's a keyword field.

247.4 Search Insights

Improved

Past versions of Search Insights showed you the full query string sent to the search engine, but now it also displays the response from the search engine with an explanation of the score for each search hit.

247.5 New Search Admin Functionality

The Search Admin functionality is found in Control Panel → Configuration → Search.

247.6 Search Engine Info

The displayed Search Engine information is enhanced, showing the client and node information as well as the vendor and operation mode.

247.7 Field Mappings

The Field Mappings tab shows the field mappings for all indexes in the search engine.

247.8 Indexing Progress

Indexing now displays a progress bar so you can see in the UI when the re-indexing action has completed.

247.9 New System Settings

Access the Search System Settings at Control Panel → Configuration → System Settings → Search.

There's a new entry in the search category: *Title Field Query Builder*. Use it to configure how search responds to matches on the Title field of a document.

Exact Match boost: Give an additional boost when searched keywords exactly match the title field of a document.

Maximum Expansions: Limit the number of documents to return when matching searched keywords to the title field as a phrase prefix. See Elasticsearch's Match Phrase Query documentation for more information.

247.10 New Infrastructure

There are some important search infrastructure changes to know about.

247.11 Elasticsearch Support

7.0 supports Elasticsearch 6.5.x through 6.8.x, and 6.5.0 is included as the embedded version to use for testing out-of-the-box search behavior. See the deployment guide for more information.

247.12 Application-Specific Indexes

You'll notice more search indexes in 7.0. That's because you can now configure application-specific indexes. At the time of this writing, the additional indexes are all related to the DXP Workflow Metrics feature. More will likely appear in future versions, and third party developers can use the portal-search APIs to create their own indexes. It's under development, so visit the Search Framework documentation frequently to discover new search infrastructure changes that expose more functionality for developers.

247.13 API Enhancements

Enhancements to the search framework APIs include

- Low level indexing and queries
- Operations directly on indexed documents (no need for the Indexer framework)
- New Aggregation types

247.14 Multi-Language Search

Liferay DXP's support for multi-language search took a step forward, with improvements to Documents and Media and Web Content. More improvements are necessary in this area and will be prioritized in future releases. See the Multi-Language Search article for more information.

CONFIGURING SEARCH PAGES

There are multiple ways to skin the search cat (disclaimer: no actual cats were harmed during the writing of this article).

If you're unsure which approach to take, use the default configuration. It provides a sensible starting point that can be modified later, as needed.

If you've been using Liferay DXP for a long time and like the search experience you've always used, use the legacy approach.

If you're in need of a fully customized experience, manually configure the search experience.

After choosing your approach and reading here to get it up and running, find the articles on the Search Bar, Search Facets, and Search Results to understand the full suite of configuration options.

Search display pages are where users go to enter search terms and browse search results.

248.1 Search Page Templates

The default search page is backed by a Search Page Template, and manually configured search pages can use the template, too. The template can be used in two ways:

1. Enable inheriting changes to the template, if you want the search page to get any updates made to the template at a later date.
2. Create the page based on the template, but independently configured after the initial creation.

Out of the box, the Search Page Template includes these widgets:

- Search Bar
- Suggestions
- Search Results
- Search Options

- Site Facet: This widget is hidden unless the Search Bar is configured to search the *Everything* scope and results from multiple sites are returned.
- Type Facet
- Tag Facet
- Category Facet
- Folder Facet
- User Facet
- Modified Facet

Out of the box, widgets use the *Barebone* Application Decorators: unless there's content to render in the widget, the widget body is hidden. The header is displayed if you hover over it.




Figure 248.1: At first glance, not much is happening on the search page. But, there's more than meets the eye.

Because of this, when you visit a search page created from the default search page template, you won't see certain widgets fully rendered.

By contrast, when you add a search widget to a page manually, they use the *Borderless* decorator (by default), which shows more of the widget even when there is no content to display.

248.2 Default Search Pages

Using the default site and the default theme with the default search settings, the out-of-the-box search experience has two components for end users:

1. A search bar embedded on each page.
2. A default search page where search requests are routed.

Behind the scenes, The search bar widget points to a hidden search page with the friendly URL `/search`.

Enter a search term and you're redirected to the default search page, where results are displayed in the Search Results widget.

Destination Page

/search

Figure 248.2: By default, the embedded search bar points to the pre-configured /search destination page.

The screenshot shows a search interface with a navigation bar containing 'Home' and 'Search'. Below the navigation bar is a search input field containing the text 'test' and a search icon. To the left of the search results are three facet widgets: 'Type' with 'Blogs Entry (4)' and 'User (1)'; 'User' with 'marvin smart (2)', 'test test (2)', and 'ziltoid omniscient (1)'; and 'Last Modified' with options like 'Past Hour (5)', 'Past 24 Hours (5)', 'Past Week (5)', 'Past Month (5)', 'Past Year (5)', and 'Custom Range...'. The main content area displays '5 results for test' and lists five search results, each with a user icon, a title, a source, a date, and a snippet containing the word 'test'. At the bottom of the results area is a button that says 'Configure additional search options in this page.'

Figure 248.3: The default page is pre-configured with the Search Results widget and the various Facet widgets to provide a full search experience.

The default search page is based on a Search page template, but it doesn't inherit changes from the page template by default. That means you can customize the search page directly without changing the template's inheritance configuration.

General SEO Look and Feel Advanced

Name *

Search en-US

Hidden from Navigation Menu Widget

YES ⓘ

Friendly URL ⓘ

http://localhost:8080/web/guest

/search en-US

Inherit Changes

NO ⓘ

Create an empty page you can lay out manually.

| | | | |
|--|---|--|--|
| <input type="checkbox"/> 1 Column | <input type="checkbox"/> 2 Columns (5... | <input type="checkbox"/> 2 Columns (3... | <input type="checkbox"/> 2 Columns (7... |
| <input type="checkbox"/> 3 Columns | <input checked="" type="checkbox"/> 1-2 Columns (...) | <input type="checkbox"/> 1-2 Columns (...) | <input type="checkbox"/> 1-2-1 Column... |
| <input type="checkbox"/> 1-2-1 Column... | <input type="checkbox"/> 1-3-1 Columns | <input type="checkbox"/> 1-3-2 Columns | <input type="checkbox"/> 2-1-2 Columns |
| <input type="checkbox"/> 2-2 Columns | <input type="checkbox"/> 3-2-3 Columns | | |


Save Cancel

Figure 248.4: Configure the Search page. By default, it doesn't inherit changes from the page template.

If you require just a few changes to the default page, don't abandon it and create one manually. Just make the configuration changes you need, including adding, configuring, and removing widgets on the page. On the other hand, if you want a clean break from the default search page, starting from scratch is also an option.

248.3 Manual Search Page Configuration

It's reasonable to create the search experience from the ground up. If you're working from a newly created site, it's a necessity. These steps show you how to switch to a manually configured search experience in the default site, but you can skip the step on deleting the default search page if you're starting with a new site:

1. Delete the existent search page by navigating to the default site's menu and clicking *Site Builder* → *Pages*. Click the Search page's Actions menu icon () and select *Delete*. Confirm you want to delete the page, and it's gone.

Once deleted, the search bar disappears from your site pages, replaced by a warning message visible only to site administrators:



This search bar is not visible to users yet. Set up its destination to make it visible.

Figure 248.5: The search bar is only visible if it points to an existent page.

2. Create a new page named whatever you want (*Finders Keepers*, perhaps). Make it hidden or add it to the navigation as you please (the default search page is hidden from the navigation). If you want a pre-configured search page, create it from the Search page template. Find the template in the Add Site Page form. It's under *Global Templates*.





Figure 248.6: There's a handy page template for creating search pages.

3. If you're creating a page not backed by the template, add and configure all the widgets you need. You'll find all the available search widgets in the Add Widget menu's Search section. Lay them out however you want on the page.
4. Configure the search bar at the top of the page, making sure it points to your new search page's friendly URL (for example, `/finders-keepers`).

Click the Search Bar widget's Options menu ().


Click *Configuration* and set the Destination Page to the search page's friendly URL.

Click *Save*.

Now your search page is up and running.

248.4 Legacy Search Experience

In prior versions, the search experience was encapsulated in one application, *Search*. It was embedded in the default theme, just like the search bar is now. It looked very similar, with only the search bar visible in the default view of the application. Once a search term is entered, the maximized view of the application is presented, with all the search facets and results now in view. It looks a lot like the new search behavior, only its monolithic structure means it's difficult to customize. If you liked the old application, it's still available. Enable it with these steps:

1. Delete the default search page. From the site menu, click *Site Builder* → *Pages*. Click the Actions menu () for the Search page and choose *Delete*.
2. Enable the legacy search application. Go to Control Panel → Configuration → System Settings → Search → Search Web and check the box for *Classic Search Widget in Front Page*.

Now your portal's search is backed by the legacy Search application, and it's embedded on each page in the default theme. To add the legacy Search application to a page, open the Add Widget menu, find the Search widget under the Tools category, and drag and drop it onto the page.

Configure the portal's search behavior to suit your needs. Here you've seen three distinct search configurations.

SEARCHING FOR ASSETS

As explained in the Search introduction, all indexed assets can be returned as search results. Developers can create their own assets, so your installation might have additional asset types beyond the ones included by default.

Searching for Users: When you click an asset in the search results, it's displayed in an Asset Publisher (unless the *View in Context* option is selected in the Search Results portlet). Users are different, though. Think of them as invisible assets, not intended for display in the Asset Publisher application. While Users appear as search results with other indexed assets, when you click one you're taken to the User's profile page. If public personal pages are disabled, clicking on a User from the list of search results shows you a blank page.

249.1 Search Bar

Users enter the search context in the search bar. Users enter search terms, hit the *Enter* button (or click the magnifying glass icon), and they're taken to a search page with various search widgets deployed.

If using the Search Bar in the legacy search portlet, users see a maximized view of the search portlet displaying any results and facets that apply. See the article on configuring search pages to learn more about these options.




Figure 249.1: The default search configuration displays a search bar in its default view, beckoning users to enter the search context.

249.2 Entering Search Terms

Liferay’s search infrastructure supports full text search as implemented by its supported search engines (Elasticsearch and Solr).

Full text search compares all the words entered in a search query (for example, *space vacation*) to all the words in each index document. A search engine like Elasticsearch calculates relevance scores to ensure the best results are returned first (like a Blogs Entry titled *Is a vacation in space right for you?*) and lots of matching results are returned (anything with either the word *vacation* or *space* is returned).

In addition to full text search, advanced search syntax is supported. Liferay DXP relies on the underlying search engine for this behavior, so consult the Elasticsearch or Solr documentation for the details.

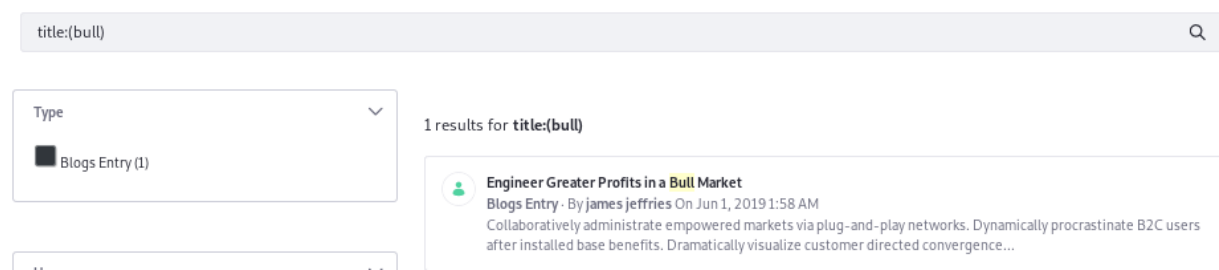


Figure 249.2: Search for text in a specific field using Elasticsearch’s Query String syntax.

249.3 Matching Exact Phrases: Quoted Search

What if users want their search terms (for example, *agile frameworks*) to produce only results with the exact phrase, as typed? In a regular full text search, searching *agile frameworks* returns search results containing just the terms *agile* and *frameworks*, and hits containing both terms but separated by other text, as well as results with the exact phrase match. To ensure that only hits with the exact phrase are returned, enclose it in quotes: “*agile frameworks*”.

249.4 Prefix Searching

If you’re searching in a site for classical musicians, you might search for the term *instrument*. This search returns documents with the full word in them, but it also returns variants with *instrument* as the prefix. For example, results with *instruments*, *instrumental*, and *instrumentation* are also returned.

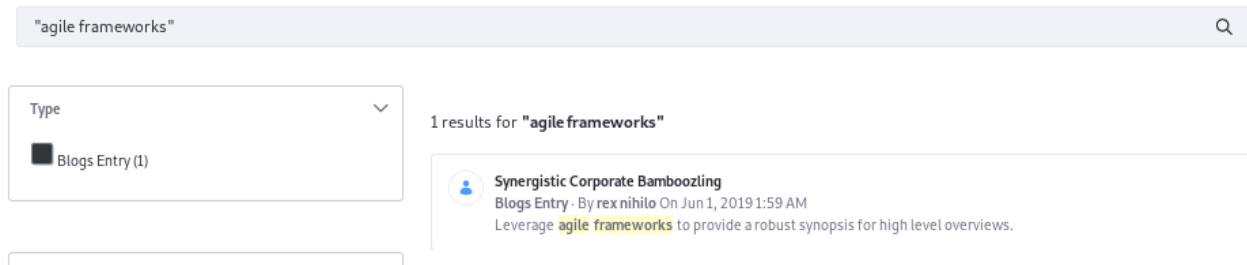


Figure 249.3: Search for exact phrase matches by enclosing search terms in quotes. If a user searched for “*agile frameworks*”, this result would not be returned.

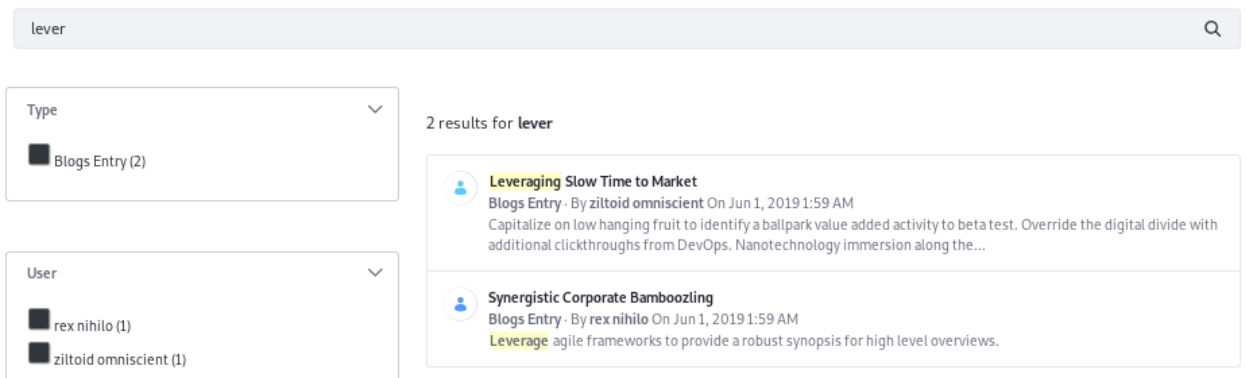


Figure 249.4: Searching for *lever* also returns *leverage* and *leveraging*.

Note: Prefix searching is available for many fields out of the box, but as with most things related to search behavior, it’s more complicated under the hood. The details of the field mapping, including the analyzer used on the field and any transformations performed, determine the final behavior.

Another way to ensure users see results is through search suggestions.

249.5 Configuring the Search Bar

Configure the Search Bar’s behavior via its portlet configuration screen.

Note: When you configure the globally embedded Search Bar widget at the top of one page, it configures the page-top Search Bar widget on all pages in the site. It also overrides the destination Search Page’s Search Bar portlet, if they’re configured differently. However, it does not override Search Bar widgets manually placed on other pages.

The screenshot shows a configuration interface for a search bar. It contains the following elements:

- Keywords Parameter Name:** A text input field containing the character 'q'.
- Scope:** A dropdown menu currently set to 'This Site'.
- Scope Parameter Name:** A text input field containing the word 'scope'.
- Destination Page:** A text input field containing the path '/search'.
- Use Advanced Search Syntax:** A checkbox that is currently unchecked.
- Invisible:** A checkbox that is currently unchecked.
- Federated Search Key:** A checkbox that is currently unchecked.

Figure 249.5: Configure the search bar behavior in its configuration screen.

There are several options:

Keywords Parameter Name Edit the parameter name for the keywords entered in the search. For example, the default URL when searching for the keyword term *data* looks like this:

`http://localhost:8080/web/guest/search?q=data`

If you change the Keywords Parameter Name to *keyword* it looks like this:

`http://localhost:8080/web/guest/search?keyword=data`

Scope Choose between three options: *This Site* (default), *Everything*, and *Let the User Choose*. *This Site* means only the assets associated with the site where the search is executed are searched. Expand the scope of the search to all sites by selecting *Everything*. To let users choose which scope they want to search, select *Let the User Choose*.

Scope Parameter Name : Set the URL parameter name for the scope where the search is taking place. This parameter only appears in the URL if the scope *Let the User Choose* is selected. The default value is *scope*, so searching for the word *data* produces the default URL of

`http://localhost:8080/web/guest/search?q=data&scope=this-site`




Figure 249.6: Let the user choose which scope the search is executed for.

Changing *scope* to *target* would produce this URL:

`http://localhost:8080/web/guest/search?q=data&target=this-site`


Destination Page Provide a friendly URL to the search page. If not configured or if it points to a page that doesn't exist, a message appears for administrators that the search bar must be configured for it to appear to users.

Use Advanced Search Syntax If using Elasticsearch, enabling this allows users to enter Query String Syntax into the Search Bar. If using Solr, consult its documentation for the proper syntax.

249.6 Search Suggestions

Suggest search terms to users when their initial queries are suboptimal. Spell check settings allow administrators to configure the Search application so that if a user types a search term that doesn't return many results (for example, a slightly misspelled word), the user can be prompted to improve the search.

To configure the spell check settings,

1. You must first reindex the spell check indexes. Go to *Control Panel* → *Configuration* → *Search*, then click *Execute* next to *Reindex all spell check indexes*.
2. Add the Suggestions widget to the search page.
3. Open its configuration screen. Click the widget Options button () and select *Configuration*.

There are three main settings here:

Display “Did you mean...” if the number of search results does not meet the threshold. : Present users alternate, spell checked search queries if their search did not return a minimum number of results (50 by default).

Display Related Queries If the number of search results doesn't meet the specified threshold (50 by default), display up to a maximum number of alternative queries (10 by default).

Add New Related Queries Based on Successful Queries Index a user's search query if it produces a minimum number of results (50 by default), so it can be displayed to users as a suggestion. If the Display Related Queries setting is enabled, it's used as a related query for similar search queries that don't produce enough results.

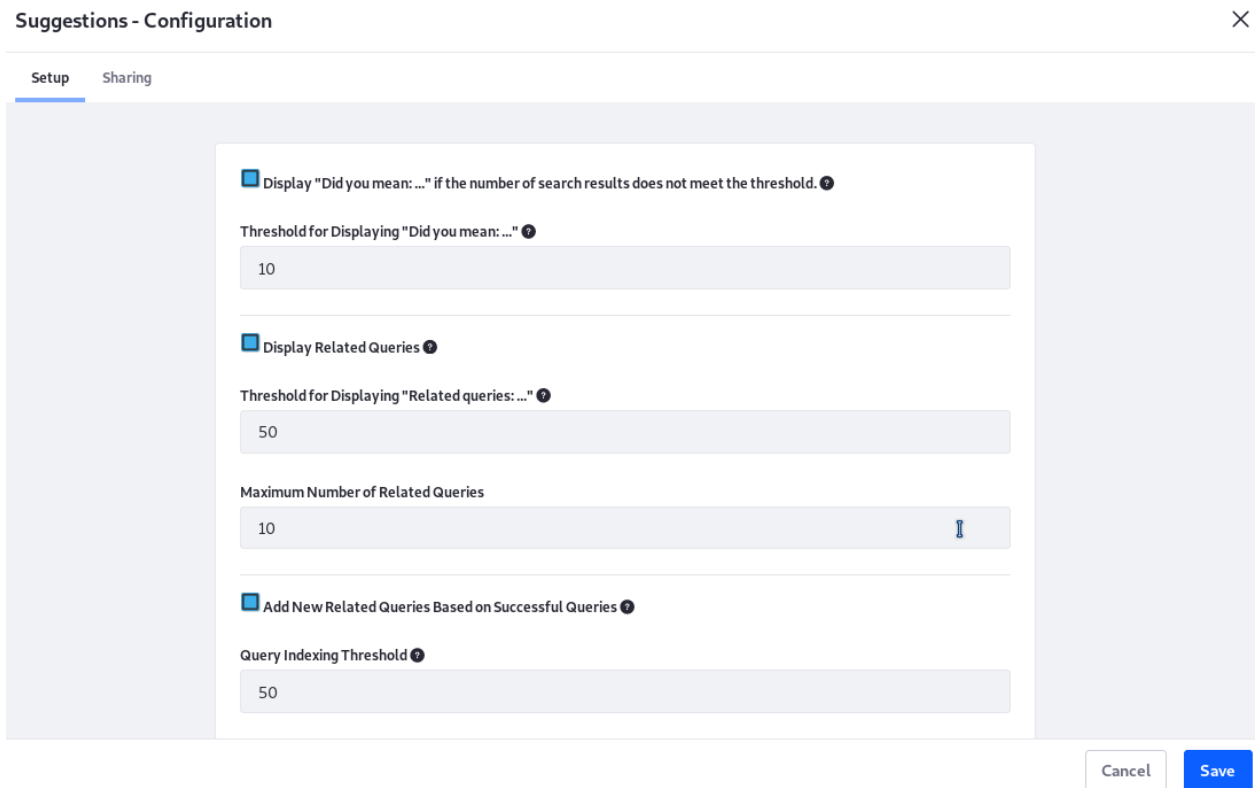


Figure 249.7: Configure the suggestion settings to allow for user input mistakes and help lead users to results.

FACETS

Enter a keyword in the Search Bar and click the Search button. The default search experience redirects to a page with results on the right and a collection of *facets* on the left.



Figure 250.1: *Site* and *Type* are two of the facet sets you'll encounter.

A Facet aggregates search results by some common characteristic, with each facet holding search results that share something in common. After scanning the full list of results, a User might decide the results from just one facet are more appropriate (for example, all the results from a particular site, or all the results that are Blogs Entries). So what facets are included by default?

- **Site Facet** for filtering results by their site.
- **Type Facet** for filtering results by the Asset Type.
- **Tag Facet** for filtering results by Tag.
- **Category Facet** for filtering results by Category.
- **Folder Facet** for filtering results by Folder.
- **User Facet** for filtering results by the content creator.
- **Modified Facet** for filtering results by the Last Modified Date.
- **Custom Facet** for filtering results by some other indexed field. See here for more information.

Each item in a facet (selected using the checkbox) is called a *Facet Term* (*term* for short).

In this tutorial, you'll explore how facets and their terms are used and how to find a facet's configuration. The remaining articles show the configurations available for each facet.

250.1 Using Facets

If you're not actually an accomplished oboe player, pretend for a moment. You're visiting a site for classical musicians. You remember reading a great technical analysis of Johann Bach's compositions, but you forgot to bookmark it (or would it be a *bachmark*?). You enter the keyword *bach* into the search bar and, because Johann Bach was a very important and famous composer, you get lots of results: too many, in fact. At first you're discouraged, but you remember that there's a site member who produces most of the site's good technical content, who's named *back2bach*. You see that his name is listed in the User facet, and there aren't many results in the facet count (the number in parentheses next to the facet). You click into the facet and quickly find the content you wanted.

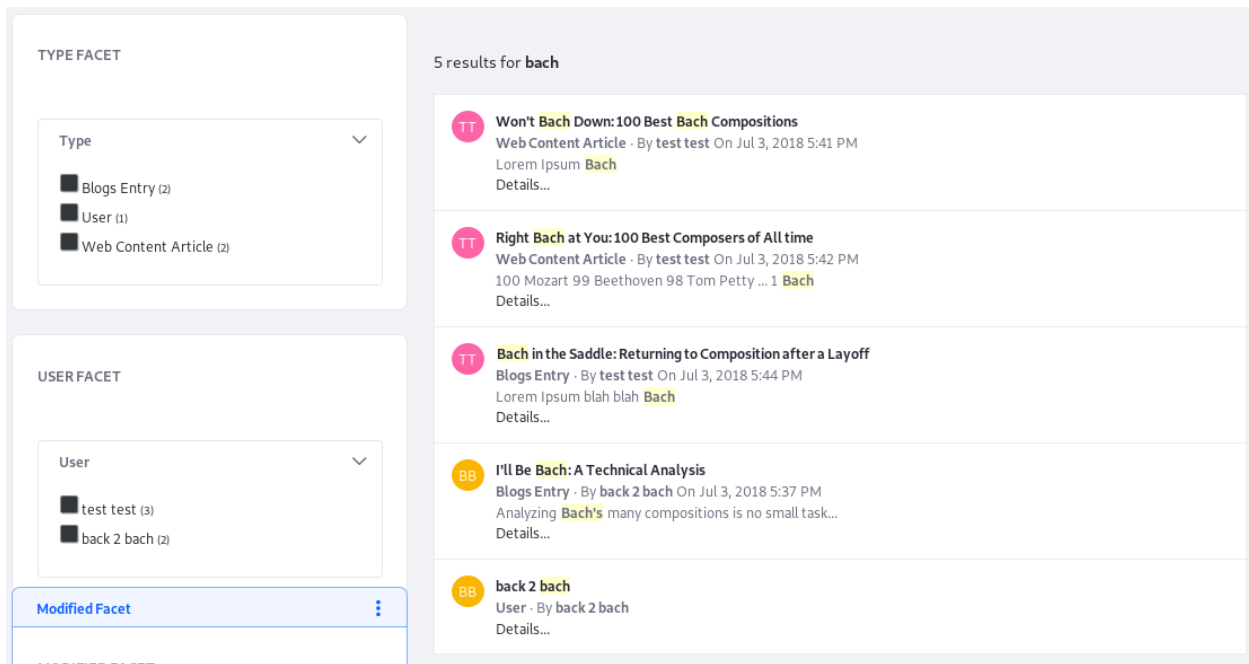



Figure 250.2: When presented lots of search results, facets narrow down the results list so users can find relevant content.

Clicking on a facet narrows down the search results. It's added to the filter list in the search query, and the results list is refined by the selected facets.



250.2 Multiple Facet Selection

Facet term selections within one facet are additive. Clicking more terms in the same facet expands the search results, because it's processed as if you want to see results matching *Term-1 OR Term-2, OR* etc. To remove all the term selections from a facet, click the *Clear* link.



Figure 250.3: Facet terms are additive when applied in the same facet. Any Blogs Entry OR Web Content article matching the keyword is shown here.

Facet term selections from different facets are exclusive. Clicking facet terms from multiple facets narrows the results because they're processed as if you want to see results matching *Facet-1 AND Facet-2, AND* etc. This is intuitive. The facets

Considering a case where you make two term selections in the Type Facet: (Blogs Entry and Web Content Article), and two term selections in the User Facet (James Jeffries and Marvin Smart). What results are displayed?

Blogs Entries OR Web Content Articles AND authored by James Jeffries OR Marvin Smart.

If Marvin and James each created four pieces of content (two blogs and two Web Content Articles), all eight would appear in the Search Results. Any Blogs or Web Content created by other Users are not shown, and assets of other Type created by Marvin and James are not displayed. Content that isn't Blog Entries or Web Content Articles created by other Users are obviously not searched.

Note: The new Search Facet widgets support the multiple selection of facet terms. Multiple facet selection is not supported in the classic Search portlet.




Figure 250.4: Facet terms selected from different facets are exclusive. These results must be of type Blogs Entry AND be from the User Marvin Smart.

250.3 Facets and Friendly URLs

In the classic, monolithic Search portlet, URLs like this were not uncommon:

```
http://localhost:8080/web/guest/home?_com_liferay_portal_search_web_portlet_SearchPortlet_formDate=1529671834606&p_p_id=com_liferay_portal_search_web_site
```

Search now uses friendly search URLs for facet filtering. With the default settings, here's the default main search URL when searching for keyword *test*:

```
http://localhost:8080/web/guest/search?q=test
```

Selecting a facet term causes a new parameter to the above URL. For example, selecting *Blogs Entry* from the Type facet results in this URL:

```
http://localhost:8080/web/guest/search?q=test&type=com.liferay.blogs.model.BlogsEntry
```

Selecting another facet term from the same facet category appends the same parameter again, but with the newly selected value:

```
http://localhost:8080/web/guest/search?q=test&type=com.liferay.blogs.model.BlogsEntry&type=com.liferay.portal.kernel.model.User
```

The rest of the facets work the same way. Filtering by the last hour option in the Last Modified facet portlet produces this URL:

```
http://localhost:8080/web/guest/search?q=test&modified=past-hour
```

The parameter names are configurable for each facet.

Now that you know how facets work, read about configuring each of the included facets.

The screenshot displays a search interface with three facets on the left and a list of search results on the right. The facets are: 'Type' with options 'Blogs Entry (4)', 'User (2)', and 'Web Content Article (4)'; 'User' with options 'james jeffries (4)', 'marvin smart (4)', and 'test test (1)'; and 'Last Modified' with options 'Past Hour (8)', 'Past 24 Hours (8)', 'Past Week (8)', 'Past Month (8)', 'Past Year (8)', and 'Custom Range...'. The search results are for the query 'lunar resort' and consist of 8 items. Each result includes a circular icon with initials, a title, a type, author, and timestamp, followed by a short text snippet.

Type

- Blogs Entry (4)
- User (2)
- Web Content Article (4)

Clear

User

- james jeffries (4)
- marvin smart (4)
- test test (1)

Clear

Last Modified

- Past Hour (8)
- Past 24 Hours (8)
- Past Week (8)
- Past Month (8)
- Past Year (8)
- Custom Range...

8 results for **lunar resort**

- JJ** **The Lunar Resort Construction Projects**
Blogs Entry · By james jeffries On Jul 2, 2018 7:42 PM
The re-imagining of The Lunar Resort continues, as construction continues on the...Resort-wide enhancements have now been...
- MS** **The Lunar Resort Medical Facility Disclaimer and Liability Waiver Notice**
Web Content Article · By marvin smart On Jul 2, 2018 8:04 PM
This medical disclaimer is designed to help obfuscate The Lunar Resort (and subsidiaries and their friends and relations) responsibilities and...
- JJ** **The Lunar Resort's Seasonal Activities**
Web Content Article · By james jeffries On Jul 2, 2018 7:53 PM
The Lunar Resort is where you can enjoy four distinct seasons. We recognize five seasons – Spring, Summer, Fall, Winter and The Holidays! Each...
- MS** **The Lunar Resort Staff's Picks: Quaint and Relaxing Locations at the Resort**
Blogs Entry · By marvin smart On Jul 2, 2018 7:56 PM
The Lunar Resort is great, but sometimes you want to escape the throngs....These staff pickc for best hidden places at the...
- MS** **The Lunar Resort Staff Picks: Best Library Books**
Blogs Entry · By marvin smart On Jul 2, 2018 7:59 PM
100 The Lord of the Rings by JRR Tolkien (1954-55) Hairy-toed hobbit Frodo leaves home to defend the world from dark forces by destroying a cursed ring, in Tolkien's epic trilogy. WH Auden thought...
- JJ** **The Lunar Resort's Main Attractions**
Web Content Article · By james jeffries On Jul 2, 2018 7:50 PM
Scenic, Sprawling, and Ever-expanding, The Lunar Resort has long been known intergalactically in the film and entertainment industry, particularly...
- JJ** **The Lunar Resort's Half-Moon Café Grand Opening**
Blogs Entry · By james jeffries On Jul 2, 2018 7:46 PM
If everything goes as smoothly as a Luna Sea diving trip, the Half-Moon Café will open by the end of this week. We're just waiting on final inspections, and the shuttle carrying our inspector was...
- MS** **The Lunar Resort Cancellation Policy**
Web Content Article · By marvin smart On Jul 2, 2018 8:00 PM
We require 14 days cancellation notice prior to your scheduled arrival, otherwise we will charge you cancellation fee as below. 10% of your total amount will be charged if you do not notify us...

Figure 250.5: Both intra-facet and inter-facet selection is possible.

SITE FACET

The Site Facet narrows search results down to those existing in a certain site. Each Site with content matching the searched keyword appears as a facet term.




Figure 251.1: Each Site with matching content is a facet term.

For the Site Facet to display multiple sites, the Search Bar must be configured to search *Everything*. See more about search scope here. If not searching for *Everything*, only the current site is searched, and the Site Facet has nothing to display. When this occurs, the Site Facet is hidden on the page.

Note: Configuring the globally embedded page-top Search Bar to search for *Everything* not only configures the embedded Search Bar on all pages. It also ensures that the Search Page's Search Bar searches *Everything*, because the page-top Search Bar's configuration overrides the Search Page's Search Bar configuration. The same does not apply to other Search Bar widgets in the site. Each of these must be configured as desired.

If the global Search Bar is disabled, configure the Search Page's Search Bar widget to search for *Everything*.

To configure the search scope,

1. Open the Search Bar's Options menu (ⓘ) and click *Configuration*.
2. Set the Scope option to *Everything*.
3. Click *Save* and close the pop-up.

The Site Facet contains several configuration options:

Site Parameter Name**Max Terms****Frequency Threshold** **Display Frequencies**

Figure 251.2: The Site Facet is configurable.

Site Parameter Name Set the URL parameter name for the Facet. The default is *site*. Searching for *lunar resort* and clicking on a site facet produces the URL

```
http://localhost:8080/web/guest/search?q=lunar+resort&site=20126
```

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the list of facet terms. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the term result list.

Display Frequencies Choose whether or not to display the term frequencies.

TYPE FACET

The Type Facet narrows search results down to those associated with a certain Asset Type. Each Type with content matching the searched keyword appears as a facet term.



Figure 252.1: Each Asset Type with matching content is a Type Facet term.

By default, all out of the box Asset Types are included as facet terms:

- Wiki Page
- Document
- User
- Blogs Entry
- Form Record
- Documents Folder
- Dynamic Data Lists Record
- Web Content Article
- Message Boards Message
- Calendar Event

- Knowledge Base Article

The Type Facet contains several configuration options:

Type Parameter Name

type

Frequency Threshold

1

Display Frequencies

Current

Wiki Page
Document
User
Bookmarks Folder
Blogs Entry
Form Record
Documents Folder
Dynamic Data Lists Record
Bookmarks Entry
Web Content Article

Available

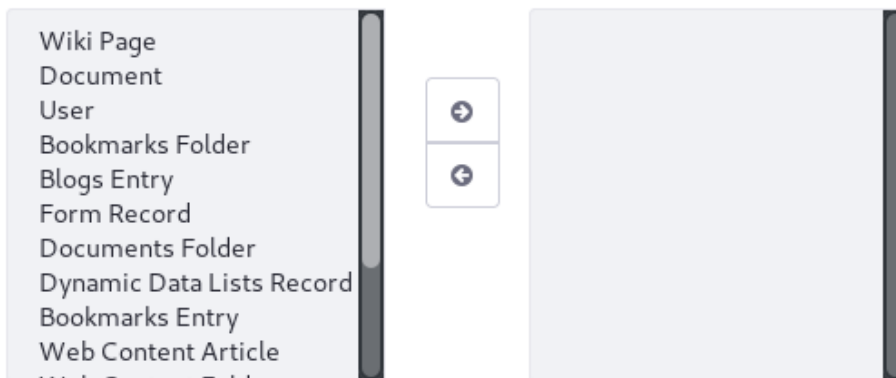


Figure 252.2: The Type Facet is configurable.

Type Parameter Name Set the URL parameter name for the Facet. The default is *type*. Searching for *lunar resort* and clicking on a site facet produces the URL

```
http://localhost:8080/web/guest/search?q=lunar+resort&type=com.liferay.blogs.model.BlogsEntry
```

Frequency Threshold Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Display Frequencies Choose whether or not to display the term frequencies.

Current and Available Add or remove Asset Types from the facet. To remove types, select from the Current section by clicking and highlighting. Click the right arrow and move the Asset Type from *Current* to *Available*. Add Asset Types by moving them to the Current section.

TAG AND CATEGORY FACETS

If tags or categories were applied to an asset appearing in the result set, they're displayed in the Tag or Category facet, respectively. Like other facets with the Frequency Threshold configuration option, not all tags necessarily appear. By default the top 10 tags or categories are listed.



Figure 253.1: Each Tag or Category with matching content is a facet term.

Tag and Category Facets contain identical configuration options:

Tag/Category Parameter Name Set the URL parameter name for the Facet. The default is *tag/category*. Searching for *lunar resort* and clicking on a *moon* Tag Facet term produces the URL

```
http://localhost:8080/web/guest/search?q=lunar resort&tag=moon
```

Display Style Choose whether to display the facet terms in Cloud or List style.

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the result list. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the term result list.

Display Frequencies Choose whether or not to display the term frequencies.

Tag Parameter Name

tag

Display Style

Cloud

Max Terms

10

Frequency Threshold

1

 Display Frequencies

Figure 253.2: Tag and Category Facets are configurable.

FOLDER FACET

The Folder Facet narrows search results down to those contained in a certain Asset Folder. If you search for *space*, a Folder titled *Space Images* doesn't necessarily show up here. The content inside the folder must match the keyword. Only if its content matches the searched keyword does the Folder appear in the Folder Facet.

Folders of these Types appear as Folder Facet terms:

- Documents and Media Folder
- Web Content Folder



Figure 254.1: Each Folder with matching content is a facet term.

The Folder Facet contains several configuration options:

Folder Parameter Name Set the URL parameter name for the Facet. The default is *folder*. Searching for *lunar resort* and clicking on a Folder Facet produces the URL

`http://localhost:8080/web/guest/search?q=lunar resort&folder=38716`

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the result list. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the term result list.

Display Frequencies Choose whether or not to display the term frequencies.

Folder Parameter Name

folder

Max Terms

10

Frequency Threshold

1

 Display Frequencies

Figure 254.2: The Folder Facet is configurable.

USER FACET

The User Facet narrows search results down to those created by a certain User.



Figure 255.1: Each User with matching content is a facet term.

The User Facet contains several configuration options:

User Parameter Name Set the URL parameter name for the Facet. The default is *user*. Searching for *lunar resort* and clicking on a User Facet produces the URL

`http://localhost:8080/web/guest/search?q=lunar+resort&user=38716`

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the result list. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the facet.

Display Frequencies Choose whether or not to display the term frequencies.

User Parameter Name**Max Terms****Frequency Threshold** **Display Frequencies**

Figure 255.2: The User Facet is configurable.

MODIFIED FACET

The Modified Facet narrows search results down to those that match the searched keyword and that were created or modified during a certain time period.



Figure 256.1: Each time period with matching content is a facet term.

In addition to selecting a pre-configured time period, Users can select a Custom Range, specifying a From and To date using a date picker:

The Modified Facet supports configuration actions:

- Modify existing time ranges
- Delete existing time ranges
- Create new time ranges

Edit or create time ranges using a time range alias.

The available time range aliases include:

```
past-hour
past-24-hours
past-week
past-month
past-year
```

Each Range has an alias and a Label.

By default, all the default ranges end in *, which evaluates to *now*. For example, the past-week range is

Custom Range... (0)

From

07/02/2018

to

07/03/2018

Search

Figure 256.2: Users can include a Custom Range in the Modified Facet.

Label

past-24-hours

Range

[past-24-hours TO *]

+

-

Label

past-week

Range

[past-week TO *]

+

-

Figure 256.3: The time ranges are set in the facet's configuration.

[past-week TO *]

You're not limited to ending Ranges. Instead of the *, specify another time range alias as the ending point.

To set up a range from 12 months ago to one month ago,

1. Click the plus button in one of the existing ranges.
2. Give it the label **1-12 Months Ago**.

Give it a Range value of

[past-year to past-month]

This gives you lots of flexibility in using alternative time ranges as Modified Facet terms.

CUSTOM FACET

All Facets are configurable, allowing you to narrow down search results based on a shared characteristic (all Blog Entries, for example). The Custom Facet lets you create entirely new Facets. The first thing to do is enter the Custom Facet's configuration screen.




Figure 257.1: Custom Facets must be configured first.

The screenshot above shows a Custom Facet with the Job Title of each User that matched the search. The next screenshot shows how it was configured.

Aggregation Field Specify the non-analyzed keyword field whose value is used to create the facet terms. If the value of the search result's `jobTitle_sortable` field is *upsale manager*, that's what appears in the Custom Facet as a term.

Custom Heading Enter a human readable heading for the Custom Facet.

Custom Parameter Name Set the URL parameter to use for the facet. With the configuration pictured above, searching for *jane* and clicking on *chief of security* produces the URL `http://localhost:8080/web/guest/search?q=jane&jobtitle=chief%20of%20security`

The image shows a configuration form for a custom facet. It contains the following elements:

- Aggregation Field**: A text input field containing "jobTitle_sortable".
- Custom Heading**: A text input field containing "Job title".
- Custom Parameter Name**: A text input field containing "jobtitle".
- Max Terms**: A text input field containing "10".
- Frequency Threshold**: A text input field containing "1".
- Display Frequencies**: A checkbox that is checked.
- Federated Search Key**: An empty text input field.

Figure 257.2: Configure a Custom Facet in no time.

Max Terms Set the maximum number of facet terms to display, regardless of how many matching terms are found for the facet.

Frequency Threshold Set the minimum frequency required for terms to appear in the result list. For example, if the frequency threshold of a facet is set to 3, a term with two matching results doesn't appear in the term result list.

Display Frequencies Choose whether to display the term frequencies.

257.1 Finding Indexed Fields

To use the Custom Facet, you must know which non-analyzed keyword field to specify.

To browse the entire list of available fields, inspect the field mappings from Control Panel → Configuration → Search. Alternatively, use your search engine's API.

For Elasticsearch, access the field mappings from your terminal using CURL to call the Get Mapping API:

```
curl -X GET "localhost:9200/_mapping/LiferayDocumentType"?pretty
```


Solr uses the ListFields API:

```
curl http://localhost:8983/solr/liferay/schema/
```

Here's a snippet of output from the Elasticsearch example:

```
"ddmStructureKey": {
  "store": true,
  "type": "keyword"
},
"ddmTemplateKey": {
  "store": true,
  "type": "keyword"
},
"defaultLanguageId": {
  "store": true,
  "type": "keyword"
},
"description": {
  "store": true,
  "term_vector": "with_positions_offsets",
  "type": "text"
},
"discussion": {
  "store": true,
  "type": "keyword"
},
```

Use Custom Fields to aggregate facet terms by shared non-analyzed keyword field values.

SEARCH RESULTS

The ideal search experience involves a User entering a search term, waiting an infinitesimal amount of time, and having the perfectly matching asset delivered at the top of a list of other extremely relevant hits. Like this:





Figure 258.1: The goal is to return the perfect results to Users searching your site.

The developers of an asset control much about how the asset's information is stored in the search engine (this process is called *indexing*, and how its information is searched and returned in the search results. Developers who dislike how a particular asset behaves in search can use an *Indexer Post Processor* to modify the asset's indexing behavior and how search queries are constructed to look up the assets in Liferay DXP.

The Search Results behavior configurable through the UI is covered in this section:

- Search Results Display Settings
- Filtering Results
- Sorting Results
- Search Results Behavior

DISPLAY SETTINGS

The Search Results widget's default display is a paginated list. Each list item is a summarized hit to a search query. Click on a specific result to look at it in more detail. Configure display options by opening the Search Results options menu () and selecting *Configuration*.

Enable Highlighting Highlight the search terms where they appear in the search result's title or summary.

Display Selected Result in Context When an asset is clicked, show it in its native application. For example, if you click on a blog post in the search results, you see where the Blogs Entry is posted in the Blogs application. Note that you're not in the search context after clicking on a search result. When this option is unchecked, the asset displays in an Asset Publisher window while still in the search context. If you have the right permissions, you can even edit the content directly from the Search context. Click the back arrow to return to the search results.

Display Results in Document Form Display results as search documents. Never use this in production. Developers use this feature to view search responses in their indexed, document-based format. Part of a developer's job when writing search indexers is to convert documents (the objects that get indexed) to the actual object and back again. Thus, developers can see how their objects are being indexed. Once enabled, click the *Details...* link below the result summary to expand the result's document view.

The next three configurations concern results pagination.

Pagination Start Parameter Name Set the name of the URL parameter for the results page. If the default value *start* is preserved, this URL displays when the User navigates to the second results page after searching for *test*:

```
http://localhost:8080/web/guest/search?q=test&start=2
```

Pagination Delta Set the number of results to display per results page. Defaults to 20 unless you customized the `search.container.page.default.delta` property in your `portal-ext.properties` file.

TT **Perfect Search Result**
 Web Content Article · By test test On Jun 27, 2018 11:45 PM
 Look no further. You have found the **perfect** content.
 Details...

| Key | Value |
|---------------------------|--|
| articleId | 37203 |
| articleId_String_sortable | 37203 |
| classNameId | 0 |
| classPK | 0 |
| classTypeId | 35601 |
| companyId | 20099 |
| content | Look no further. You have found the perfect content. |

Figure 259.1: Viewing a results document lets you inspect exactly what's being indexed for a particular asset. This is just a small portion of one document.

Pagination Delta Parameter Name Set the name of the URL parameter that stores the Pagination Delta value. This becomes visible in the browser if the User changes the number. If the User selects 10 results per page and searches for *test*, the Search Page is reloaded with this URL:






```
http://localhost:8080/web/guest/search?q=test&delta=10
```

Federated Search Key If this widget is participating in a search on a non-default index, enter the key of the alternate search index. If not set, the widget participates in the default search, against the default index (liferay-[companyId]). This value is usually the name of an application-defined index.

Fields to Display If searching an alternate index using the Federated Search Key configuration, specify what fields to search from that index.

For further reading, check out how to return suggestions for better search terms (for example, “Did you mean...”) when not enough results are returned initially.

16 results for **test**

-  **moon-image9.jpg**
Document · By test test On Jun 27, 2018 3:52 PM
[Details...](#)
-  **moon-image7.jpg**
Document · By test test On Jun 27, 2018 3:52 PM
[Details...](#)
-  **moon-image3.jpg**
Document · By test test On Jun 27, 2018 3:52 PM
[Details...](#)
-  **moon-image6.jpg**
Document · By test test On Jun 27, 2018 3:52 PM
[Details...](#)
-  **moon-image5.jpg**
Document · By test test On Jun 27, 2018 3:52 PM
[Details...](#)

5 Entries ▾ Showing 6 to 10 of 16 entries. [◀](#) [1](#) [2](#) [3](#) [4](#) [▶](#)

Figure 259.2: The number of results per page and the URL parameter names used to control pagination behavior are configurable.

FILTERING SEARCH RESULTS WITH THE CUSTOM FILTER WIDGET


You often need to exert control over the displayed search results. One viable approach is to develop your own search portlets using the Liferay DXP APIs. That can be overkill if you just want to make a slight modification to how the search is executed, so many of the out-of-the-box search widgets give you this type of control without coding anything (Search Options, Custom Facet, and more). In 7.0, new widgets have been added: Sort and Custom Filter.

With Custom Filters, you can contribute queries to the main search query, exerting control over the search results. Make the filter widgets visible or invisible to the search Users, and decide if they're changeable or immutable.

To explore all the options you have with the Custom Filter widget, you need one on the page.

260.1 Adding and Configuring Custom Filters

To get started with Custom Filters,

1. Open the Add menu () for the page and expand the Widgets section.
2. From the Search section, drag a Custom Filter onto the page.

Custom filters can do so many things, it's impossible to list them all. What follows is a widget configuration tour. Separate documentation will be written to provide a how-to demonstration of Custom Filters.

260.2 Custom Filter Configuration Options





Figure 260.1: A custom filter has no impact until it's configured.

Open the widget Options menu () and click *Configuration*.

Filter Field (text) Most often, filters operate on a specific field. Set the name of the indexed field to be filtered (for example, `title`). You won't need this if the Filter Query Type is set to a type that doesn't require a field, such as *Regex*.

The Query String and Script queries do not require a Filter Field to be set. All other queries require at least one field.

Filter Value (text) For most filters, you must enter a text value here that specifies the text to apply the filter on in the specified field (for example, set a *Match* query to the text `street` on the `title_en_US` field). Some Filter Query Types require special notation, as in the case of the *Regex* filter.

Filter Query Type (select list) Select the query type to filter results by. Available types include Bool, Exists, Fuzzy, Match, Match Phrase, Match Phrase Prefix, Multi Match, Prefix, Query String, Regex, Script, Simple Query String, Term, Wildcard. To learn more about these queries, visit the Elasticsearch documentation.

Occur (select list) Set the occurrence type for the query being contributed to the search. Options include Filter, must, must_not, and should. To understand each type, see the Elasticsearch documentation.

Query Name (text) Set the name of the contributed query. This is unnecessary unless this filter acts as a parent query to another filter that contributes child clauses; in that case set this filter's Query Name as the child filter's Parent Query Name. This parent/child behavior is only available for filters of type Bool.

Parent Query Name (text) When contributing a child clause to a Bool query, set this to match the Query Name configured in the parent Custom Filter widget. Otherwise, leave it blank.

Boost (number) Boost the score of the results matching this query. Specify any whole or decimal number here that makes sense. If you always want results matching this at the top, set the Boost value really high (e.g., `1000`).

Filter Field ?

Filter Value ?

Filter Query Type ?

Match ⌵

Occur ?

Filter ⌵

Query Name ?

Parent Query Name ?

Boost ?

Custom Heading ?

Custom Parameter Name ?

Invisible ?

Immutable ?

Disabled ?

Federated Search Key ?

Figure 260.2: Once the Custom Filter is added to the page, mold it like soft clay into the beautiful sculpture you've envisioned.

Custom Heading (text) Enter the heading to display for this filter. If not set, the Filter Field's value is displayed.

Custom Parameter Name (text) Specify a URL parameter name for the filter. If not set, the Filter Field's value is used.

Invisible (boolean) If checked, the widget is invisible to regular users. The Filter Value from the configuration is applied by default, but users can still filter for other values via URL Parameter. Don't worry, you can shut that down if you need to with the Immutable setting (see below).

Immutable (boolean) Enable this to ensure that the Filter Value cannot be changed by regular users. The widget becomes invisible to them *and* filter values set via URL parameters are not accepted. The Filter Value set in the widget configuration is applied at all times (unless it's disabled).

Disabled (boolean) If checked, the query is ignored and doesn't participate in searches. This gives you a quick way to stop the filter, but keep the configuration so it can be re-enabled later.

Federated Search Key (text) Enter the key of an alternate Search this widget is participating on. If it's set, be aware that the default Liferay DXP index isn't searched at all. If not set, this widget participates on the default search. Values in this field typically match the name of an application-defined index.

There are many use cases you can satisfy by incorporating Custom Filters into your search page. Two demonstrative articles are planned to show you some of the filtering capabilities at your disposal:

- *Refine to One (or Needle in a Haystack)* will show you how to add user-operated filters to the page so results can be refined down to just the result they were looking for.
- *Complex filtering* shows you some more advanced filters and how they work.

Check out the Custom Filter and see what it adds to your search page.

SORTING SEARCH RESULTS WITH THE SORT WIDGET

The Sort widget gives Users configurable control over the order of returned results: no code necessary.

Add it to a page and begin sorting results.

By default, results are sorted by the relevance score returned by the search engine. Users can choose from one of the out-of-the-box alternative sorting strategies, or one configured by a search administrator.

Out of the box, order results in these ways as an alternative to relevance sorting:


- alphabetically by Title
- by the Modified date (oldest first)
- by the Create date (newest first)
- by the Create date (oldest first)
- alphabetically by the User that created each matching asset

If the out-of-the-box alternatives aren't enough, an administrator can create additional sort options from the widget's configuration.

It's also possible to delete unwanted sort options from the widget.

261.1 Adding and Configuring the Sort Widget

To get started with the Sort widget,

1. Open the Add menu () for the page and expand the Widgets section.
2. From the Search section, drag a Sort widget onto the page.

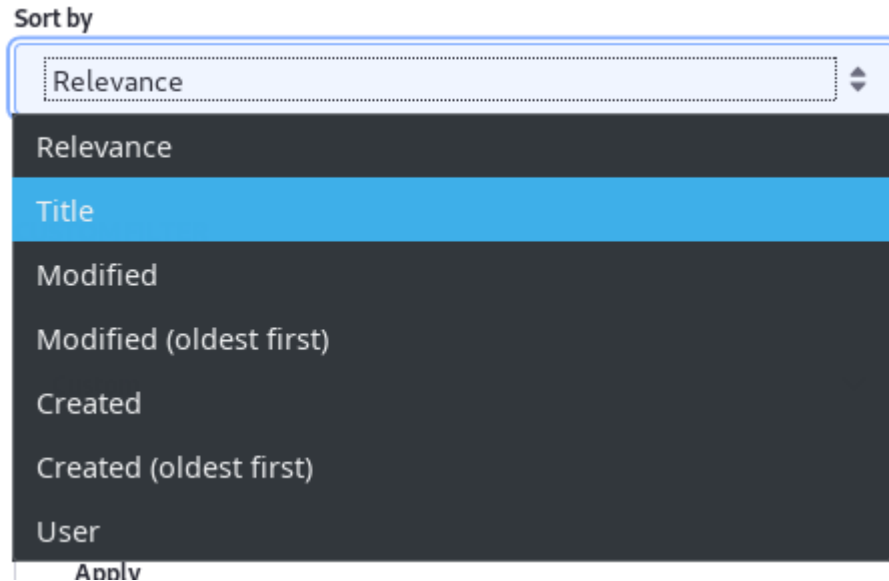



Figure 261.1: Users can re-order search results with the Sort widget.

261.2 Configuring the Sort Widget

Three things can be done from the Configuration screen:

- Editing existing Sort options
- Deleting options
- Adding options

To access the widget configuration screen, open the widget Options menu () and click *Configuration*.

Each Sort option has two fields: *Label* and *Field*.

Label Set the displayed label for the type of sort being configured.

Field The `fieldName` of the indexed field to sort. Most of the time this is a keyword field. Other acceptable options are date and any numeric datatype. There's even a way for persistent search administrators to coerce text fields into behaving with the Sort widget. Keep reading for details.

Finding Sortable Fields

To find the fields available for use in the Sort widget, Users with the proper permissions can navigate to *Control Panel* → *Configuration* → *Search*. From there, open the Field Mappings tab and browse the mappings for each index. Scroll to the properties section of the mapping, and find any keyword field, date field, or a field with any numeric datatype. The type field is instructive:

```
"type" : "keyword"
```

```
"type" : "date"
```

```
"type" : "long"
```

Sort - Configuration

Setup Sharing

Label

Field

+ -

Label

Field

+ -

Figure 261.2: From the Sort widget's configuration, add, edit, or remove Sort options.

What if you really need to sort by a text field? You can do it by adding a new version of the field to the index, with the type keyword. Don't worry; you don't need to code anything to do this. From the field mappings screen mentioned above, look at the `firstName` field in the index called `liferay-[companyID]`. In fact, look at the next entry as well:

```
"firstName" : {
  "type" : "text",
  "store" : true
},
"firstName_sortable" : {
  "type" : "keyword",
  "store" : true
},
```

There's a corresponding field with the suffix `_sortable`, and of the correct type for sorting (keyword). How did that get there? Via the portal property

```
index.sortable.text.fields=firstName,jobTitle,lastName,name,screenName,title
```

All the text fields listed here have a `fieldName_sortable` counterpart created automatically in the index. To add more, copy this value into a `portal-ext.properties` file into your Liferay Home folder, add any new field names you need to sort by, and restart the server.

Adding New Sort Options

To sort by the new field, use the plus symbol below any option's *Field* configuration make sure to use the `fieldName_sortable` version of the field in the widget configuration.

To add a new sort option that's already of the proper datatype, use the plus symbol below any option's *Field* configuration and fill in the fields. The order of options here in the configuration screen matches the order Users see in the select list while configuring the widget for their search.

Editing and Deleting Sort Options

To edit an existing option, edit the text in its configuration section.

To delete an existing option, use the minus symbol below its *Field* configuration.

261.3 Controlling the Sort Order

To control the order for the sort option, add a plus or minus symbol after the `fieldName`. Look how it's done for the existing sort options labeled *Created* and *Created (oldest first)* to understand how it works:

Label: *Created* **Field:** `createDate-`

The - sign following the field name indicates that the order is *descending*. Choosing to sort with this brings search results created most recently to the top of the list.

Label: *Created (oldest first)* **Field:** `createDate+`

The + sign following the field name indicates that the order is *ascending*. Choosing to sort with this brings the oldest (by creation date) results to the top of the list.

SEARCH RESULTS BEHAVIOR

The previous article covered ways to display search results. This article covers these additional Search Results concepts and configurations:

- Filtering search results with facets
- Understanding search results relevance
- The effect of permissions on search results
- Search results in the staging environment
- Search results summaries
- Search results term highlighting

262.1 Filtering Results with Facets

Results are filtered using *facets*. Most users have encountered similar filtering capabilities in other applications, particularly during commerce activities. Users enter a search term, are presented with a list of results and search facets, which you can think of as buckets that group results together if they share a common characteristic.

Administrators can configure facets. Read about configuring facets to learn more.

262.2 Search Results Relevance

The search engine decides which results appear at the top of the list using the concept of *relevance*. Relevance is a score calculated by the search engine. There are numerous factors contributing to the total score of a returned document, and all of the implementation details of how relevance scoring works are algorithms provided by the search engine.

262.3 Permissions and Search Results

Users lacking VIEW permission on an asset don't see it in the search results. A logged in User with the Site Administrator role likely sees more search results than a guest User to the site.

In the background, there are two rounds of permissions checks. The first permissions check, *pre-filtering*, happens in the search engine's index. It's faster than checking database permissions information, but occasionally the search index can have stale permissions information. To ensure the search engine's index has correct, up-to-date permissions information, a second, last-second permissions check, *post-filtering*, is performed on the results prior to their display.

262.4 Initial Permissions Checking

The first round of search results permissions filtering adds filter clauses to the search query. This ensures that results return from the search engine pre-filtered, containing results the current User can view.

This initial permission checking is configurable at *Control Panel* → *Configuration* → *System Settings* → *Search* → *Permission Checker*. It includes two system level settings to configure how search processes User permissions.

Include Inherited Permissions Ignore this setting. It's deprecated, no longer used anywhere, and will be removed in a future release.

Permissions Term Limit Limits the number of permission search clauses added to the search query before this level of permission checking is aborted. Permission checking then relies solely on the final permission filtering described below.

The only reason to limit permissions terms is performance. Users with administrative access to lots of sites and organizations generate many permissions terms added to the query. Too many terms in a query can make the search engine time out.

262.5 Final Permissions Checking

A final round of permission checking happens prior to presenting results in the UI. For example, the User searches for *liferay*, and the search engine returns all relevant forum posts. As the Search Results iterates through the list of relevant forum posts, it performs one last permission check of the post to ensure the User can view the post and its categories. If a matching forum post exists in a category the User doesn't have permission to view, it isn't displayed in the list of search results.

This final round of permission checking is configurable at *Control Panel* → *Configuration* → *System Settings* → *Search* → *Default Search Result Permission Filter*. It includes two settings:

1. The first setting, Permission Filtered Search Result Accurate Count Threshold, specifies the maximum number of search results to permissions-filter before results are counted. A higher threshold increases count accuracy, but decreases performance. Since results in the currently displayed page are always checked, any value below the search results pagination delta effectively disables this behavior.
2. The second setting, Search Query Result Window Limit, sets the maximum batch size for each permission checking request. . This is again impacted by pagination. For example, if there are 100 results per page, and a User wants to jump all the way to page 200 of the search results, all results between page one and 200 must be checked to ensure the User has permission. That's 20,000 results to permissions check. Doing this in one trip to and from the search engine can result in performance issues. Set the maximum batch size for each permission checking request.

262.6 Search and Staging

With staging, content is placed first in a preview and testing environment before being published for consumption by end Users (on the live site). Content added to the search index is marked so that the search API can decipher whether an item is live or not. In the live version of the site, only content that's marked for the live site is searchable.

In the staged version of the site, all content—live or staged—is searchable.

262.7 Result Summaries

A result summary includes the information from a document that the asset's developer felt is most useful to end Users searching for the asset. Each asset can have different fields included in the summary. For assets with text content, a common summary format includes the *title* and some of the *content*, with title displayed first. The asset type always appears on the second line, and a snippet of the content that matches the search term is on the last line. Assets without content fields, like Documents and Media documents, display the description instead.

Users are different. Only the User's full name and the asset type (User) appear in User result summaries.

For assets that contain other assets (Web Content and Documents & Media folders) or whose content is not amenable to display (Dynamic Data List Records and Calendar Events), it makes more sense to display the title, asset type, and description in results summaries. There'd never be anything in a content field for these assets.

The asset developer determines which fields are summary-enabled, but there's logic invoked at search time that determines precisely the part of the summary fields to display. For example, a content field can have a lot of text, but the summary doesn't show it all. Instead, it shows a relevant snippet of the field's text. If the keyword searched for is present in the summary field,




Figure 262.1: User summaries contain only the User’s full name.

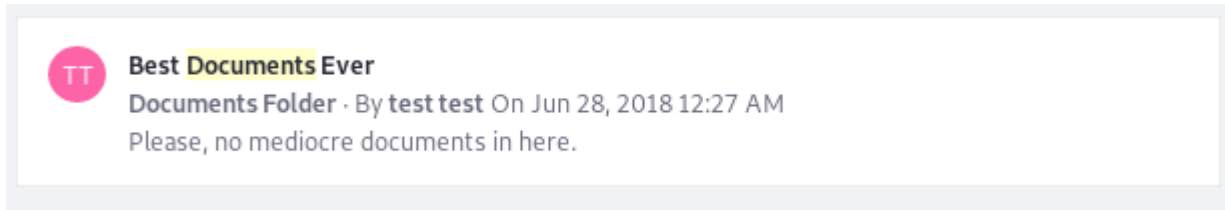


Figure 262.2: Documents and Media and Web Content folders include titles and descriptions in their summaries.

that portion of the field is used in the summary. In addition, the matching keyword is highlighted in the summary.

262.8 Highlighting

By now you’ve probably noticed that search terms appearing in the summary are highlighted by default. If this is undesirable, disable it in the widget configuration screen.

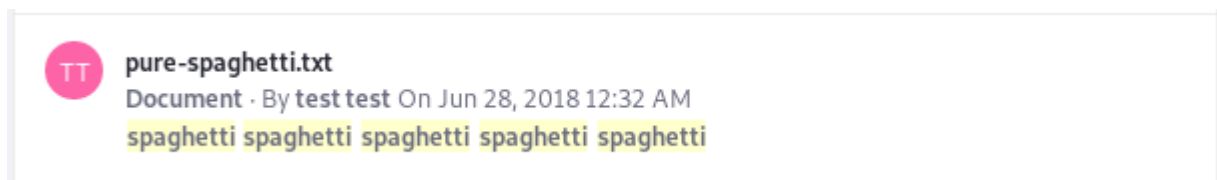


Figure 262.3: Some document summaries have lots of highlights if the search term matches text that appears in the summary.

Highlighting is a helpful visual cue that hints at why the result is returned, but beware. A hit can score well and thus be returned near the top of the results, without having any highlights in the summary. That’s because not all indexed fields appear in the summary. Consider a User named Arthur C. Clarke. He has an email address of *acc@authors.org*, which is searchable. Because results summaries for Users only contain the full name of the User, searching for Mr. Clarke by his email address returns the User, but no term is highlighted.

There are additional cases where search results don’t have highlighting.




Figure 262.4: Results that match the search term won't always have highlights.

SEARCH INSIGHTS

[Feature intended for testing and development only]

[Works with Elasticsearch only]

Add the Search Insights Widget to the Search Page to inspect two things:

- The query string that's constructed by the back-end search code when the User enters a keyword
- The response string returned from the search engine

Note: The JSON returned as a request string is pruned from several Elasticsearch query defaults for clarity. To see the full request JSON that Elasticsearch processed, adjust the Elasticsearch server's logging.

In 7.0, the Insights widget now adds the response string to the widget's output, and the *Enable Score Explanation* option (enabled by default) prints a relevance score explanation for each returned result.

When a search query is processed, results are returned. The concept of *Relevance* determines how well results match the query. The Score explanation for returned search documents helps clarify seemingly odd results, letting you adjust the relevancy scoring process by making matches in certain fields count for more (*boosting* the fields is the term for this).

263.1 Inspecting The Search Query String

To see the Search Insights widget in action, navigate to a Search Page in your test server and add it from the Add menu (⊕).

Once you search for keywords that return Search Results, the Search Insights portlet displays the returned query string in all its glory.


```

"description":"sum of:",
"details":[
  {
    "value":9.461341,
    "description":"sum of:",
    "details":[
      {
        "value":1.0,
        "description":"emailAddress:*test*",
        "details":[]
      }
    ],
  },
  {
    "value":5.0,
    "description":"userName:*test*^5.0",
    "details":[]
  },
  {
    "value":0.72928625,
    "description":"sum of:",
    "details":[]
  },
  {
    "value":1.0027686,
    "description":"sum of:",
    "details":[]
  },
  {
    "value":0.72928625,
    "description":"sum of:",
    "details":[]
  },
  {
    "value":1.0,
    "description":"screenName:*test*",
    "details":[]
  }
]
},
...
]
}}}
\`

```

Now you're able to see the entire query string, the response string, and how each returned Search Document was scored.

SEARCHING FOR LOCALIZED CONTENT

Liferay DXP supports setting a virtual instance-wide default language and setting a In addition, many out of the box assets support translation.

How an asset's fields are indexed in the search engine plays an important role in the end user's experience. Not all assets are indexed in a way that supports searching in a language other than the default language. Even assets that are translatable might not support searching for the content in that language.

In short, these assets contain text fields supporting localized search:

| Asset | Fields | Localized Search Approach |
|------------------------------|-----------------------------|---------------------------|
| Content Page | title | 2 |
| Documents and Media Document | content | 3 |
| Calendar | name, description | 1 |
| Calendar Booking | title, description | 1 |
| Dynamic Data List Record | content | 1 |
| Form Record | content | 1 |
| Web Content Article | title, content, description | 1 |
| Asset Category | title, description | 1 |
| Asset Tag | assetTagNames | 1 |
| Wiki Page | title, content | 2 |
| Blogs Entry | content, title | 2 |
| Message Boards Message | title, content | 2 |

* Asset tags and categories don't have dedicated documents in the index. Instead, their indexed fields are added to the tagged or categorized asset's document.

There are three localized search approaches represented in the table:

1. Fully Localized: The asset itself is localizable (in other words, translatable), and its translated fields are indexed into their respected locales.
2. Fully Localized for Search: Even though the asset itself is not localizable/translatable, its fields are indexed into *all* the supported locales in the site. This is a new approach, starting with 7.0.
3. Site-Localized for search: The asset's fields are indexed with the site's locale appended.

There are also assets with text fields and no localization support, meaning that they always index the plain field, without a locale appended (e.g., title is not localized, but title_en_EN is localized). That means they'll always be analyzed by the default language analyzer, and do not support localized search in any capacity.

264.1 What is Localized Search?

In localized search, fields are indexed with locale information appended (for example, en_US for English, making a localized title field indexed as `title_en_US`). It's then passed to the proper language analyzer in the search engine so that the analysis process is performed properly. Each localization approach is covered below.

264.2 Fully Localized

Fully localized search works like this:

1. One or more of an asset's fields are localizable in the user interface and database (the locale is appended based on the asset creator's language selection).
2. The fields are indexed with the appended locale and analyzed with the corresponding language analyzer.
3. At search time, properly indexed and analyzed content is returned according to search engine's relevance algorithms.

This is the ideal approach for assets that support translation of some or all fields outside of the search context.

264.3 Fully Localized for Search

Assets fully localized for search work like this:

1. The asset's fields are not localizable in the user interface or database.
2. For at least one text field being indexed, the asset has indexed localized fields for every locale available in the site.
3. At search time the result document is returned regardless of the search locale, because the content is available in all locales of the site.

264.4 Site-Localized Search

Site-localized search works like this:

1. The asset's indexed fields are appended with the locale of the site (set in Site Settings) and analyzed with the corresponding language analyzer.
2. If the site language is changed, reindexing is required to append the proper site locale to the indexed fields and analyze with the corresponding language analyzer.
3. At search time, if content exists matching the language of the site, it's properly returned according to the search engine's algorithms.

Not all assets support localized search, however. Refer to the table at the beginning of this article for which assets and fields are localized for search.

264.5 Assets Supporting Localized Search

Whether an asset supports localized search depends on how the asset was indexed in the search engine. At this time, no cohesive pan-asset approach to indexing assets for localized search exists. Localized search support is currently limited to the following assets:

264.6 Web Content Articles:

- The title, content, and description fields for each Web Content Article support fully localized search.

****Note:**** In 7.0 the default (non-localized) version of these fields are not indexed for Web Content Articles. Therefore, any custom ``IndexerPostProcessor``, ``ModelDocumentContributor`` or ``QueryPreFilterContributor`` relying on the presence of fields ``title``, ``content`` and ``description`` must be updated to use the localized version (e.g., ``title_en_US``).

- At search time, matching results (with any locale appended) can be returned.

264.7 Categories:

- The name and description fields support fully localized search.
- At search time, matching results (with any locale appended) can be returned.

264.8 Documents and Media File Entries:

- The content field (which contains the content of an uploaded file) supports site-localized search.
- No other fields are indexed with a locale. This means they're always analyzed using the default language analyzer.

264.9 Dynamic Data Mapping Fields:

- Dynamic Data Mapping (DDM) Fields include all form fields created in the Forms application and all fields created in Dynamic Data List Data Definitions and Web Content Structures.
- DDM Fields support fully localized search, with the exception that results can only be returned in the current display locale where the search is taking place.

264.10 Examples

To see localized search in action, refer to the examples below.

264.11 Fully Localized Search for Web Content Articles

1. Add a Basic Web Content article in English:

- Title: *What time is it?*
- Summary: *It's soccer time!*
- Content: *That's right, it's time for soccer. The 2018 World Cup is behind us, and teams all over the US are gearing up for soccer season. It's never too early to start practicing.*

2. Add a second article in English:

- Title: *What is the best soccer team ever?*
- Summary: *There are many good teams? Which is the best?*
- Content: *Here are the 10 best teams in the world: 1. The Lunar Resort's Club Team, Waxing Crescent FC...*

3. Add a Portuguese (*pt-BR*) translation for each field of the second article:

- Title: *Qual time de futebol é o melhor de todos os tempos?*
- Summary: *Existem muitas boas equipes. Qual é o melhor?*
- Content: *Aqui estao as 10 melhores equipes do mundo: 1. Selecao brasileira de Futebol 2. O time do Resort Lunar, Waxing Crescent FC...*

4. Find a search bar widget and enter *time* as the keyword.

The first article is returned, and so is the appropriate translation of the article about soccer teams (because *time* in Portuguese translates to the English word *team*). Note that if your search context is English, searching for the word *time* returns the English translation of the Web Content, which does not itself contain the matched keyword. The Portuguese translation contains the matching keyword, while the English translation is returned for English speaking search users.

In fully localized search, fields are appended with the proper locale, and even fields with a locale other than the User's display context are returned if they contain matches to the searched keyword.

264.12 Site-Localized Search for Documents and Media

1. Create a text file named `search-test.txt` with the following contents:

```
Meu time de futebol favorito é o melhor
```

2. Upload it as a Basic Document to the Documents and Media application.
3. If your site's language is currently set to English, adding this file appends its content field with the *en-US* locale.
4. Search in the site for the keyword *time*.




Figure 264.1: Even though the content of this DM File is written in Portuguese, it was appended with the *en* locale, so it's searchable in an English language site.

The file is returned because even though the text in the file is Portuguese, the locale appended to its *content* field is for English.

5. Now change the Site's default language to *Portuguese (Brazil)*. Use Site Settings → Languages to accomplish this.
6. Now search for *time* in the site, and the document is not returned in the results, because the search is looking for the *pt* locale.




Figure 264.2: The uploaded DM File doesn't appear when the site language is changed, because only fields with the site's locale are searched.

7. Now go to Control Panel → Configuration → Search, and click *Execute* next to *Reindex all search indexes*.
8. Search for *time* in the site's Search Bar again, and now the document is returned in the results, because the content field's locale was changed from *en_US* to *pt_BR* when reindexed.

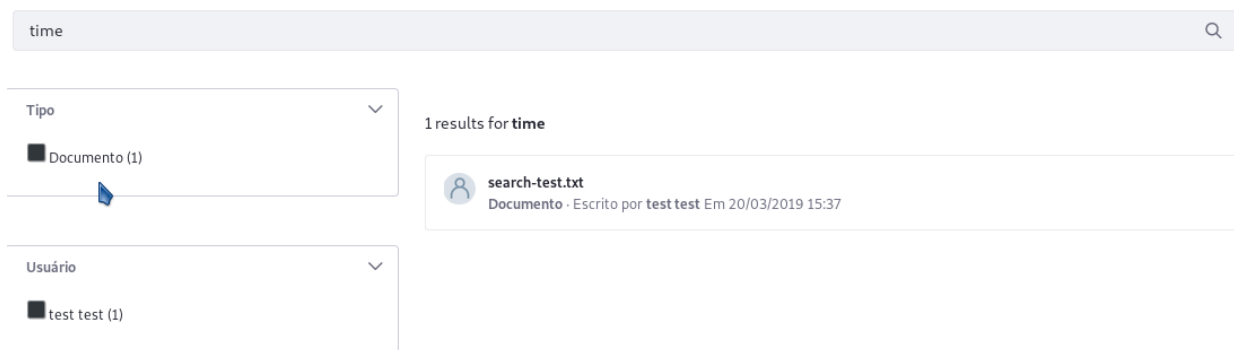


Figure 264.3: Once the field is reindexed with the site's locale, it can be returned as a search result in the site.

If an asset supports site-localized search, its fields must be reindexed after the site language is changed in order to be returned as search results.

CONFIGURING SEARCH

Configuring Search could mean lots of different things:

- System scoped search configuration
- Reindexing to make sure the search indexes are current with the database
- Tweaking the search widgets added to pages
- Creating new Search Pages
- Configuring the connectors that let Liferay DXP and the search engine communicate

In fact, *Configuring Search* means all those things. This is a high level overview of what search behavior is configurable out of the box, and importantly, *where* to find search configuration options.

265.1 System Scoped Search Configuration

System scoped search configurations are primarily found in System Settings.

1. Go to *Control Panel* → *Configuration* → *System Settings*.
2. Click the *Search* category under the Platform section.
Alternatively, search for *Search*.

These system scoped configurations are available in System Settings:

265.2 Default Keyword Query

The Default Keyword Query entry contains one setting:

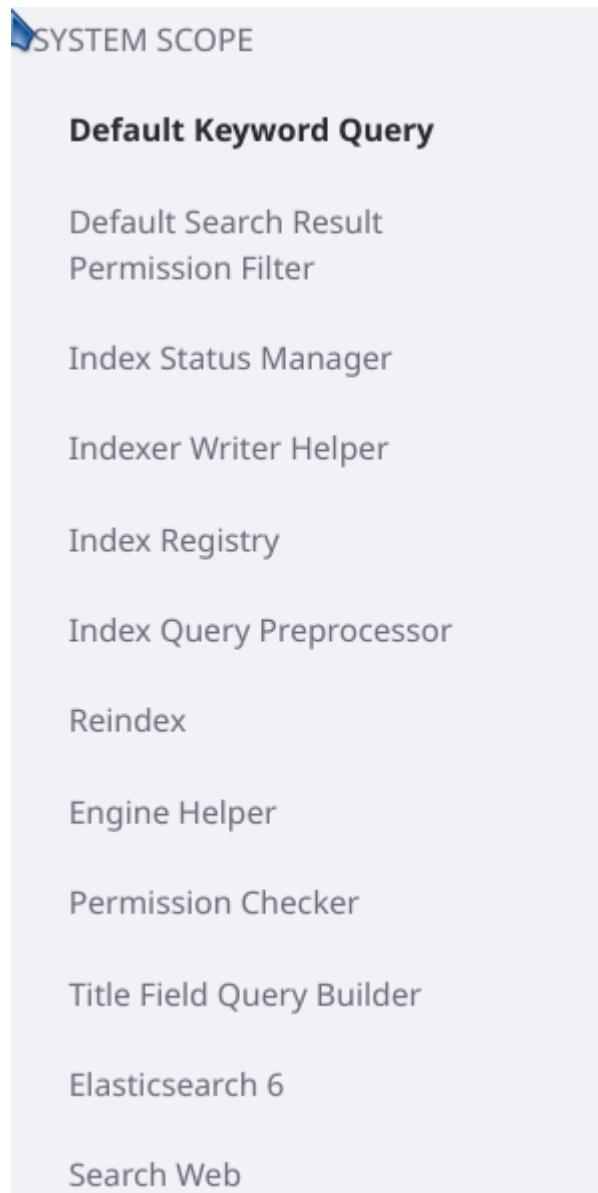


Figure 265.1: There are numerous system scoped entries for search in System Settings.

`disabledEntryClassNames`: The `DefaultKeywordQueryContributor` code automatically adds `description`, `userName`, and `title` fields to the keyword search query. Specify the entry class names `DefaultKeywordQueryContributor` should ignore.

265.3 Default Search Result Permission Filter

The Default Search Result Permission Filter entry allows configuration of *post-filtering permission checking* (database permission checking that occurs after the results are returned from the search index). Read here for more information on these settings:

- `permissionFilteredSearchResultAccurateCountThreshold`
- `searchQueryResultWindowLimit`

265.4 Index Status Manager

The Index Status Manager entry has one setting:

`indexReadOnly`: Suspends all indexing operations and writes to the search engine. Searches return only the documents already indexed. This is useful for speeding up large data imports, but it should be disabled and a full re-index executed once the import is finished.

265.5 Indexer Writer Helper

The Index Writer Helper entry contains only one valid entry. The second, `indexReadOnly`, is deprecated and unused, so setting it has no effect. Use `indexReadOnly` from the Index Status Manager instead.

`indexCommitImmediately`: When *true* (the default), each write request forces the search engine to refresh the index reader, potentially flushing transactions to disk. This may negatively impact search engine performance. The default behavior is to commit immediately for index writing on individual assets (e.g. add blog, update blog) but delay commits for bulk index writing operations (e.g. index all users, index all form entries) until all entries have been sent to the search engine. Setting this to false changes the behavior for individual index operations, and may cause applications like Asset Publisher to exhibit a delayed response when showing newly added content. See the Elasticsearch documentation for more information.

265.6 Index Registry

Configure the buffering of index requests:

`buffered`: Disable or configure the buffering of indexing requests. To stop the buffering of index requests, choose *Disabled*.

`bufferedExecutionMode`: Allows administrators to select a different `IndexerRequestBufferExecutor`, used to execute a `IndexerRequest`. One implementation of the executor is provided out of the box (*DEFAULT*). When a developer creates and deploys a new `IndexerRequestBufferExecutor` implementation, one of the properties they provided is a `buffered.execution.mode` which makes the implementation selectable from System Settings.

`maximumBufferSize`: If buffering is enabled, set the Maximum Buffer Size so that any additional indexing requests are executed immediately.

`minimumBufferAvailabilityPercentage`: When the capacity of the buffer has only the specified percent of space left, the existing requests in the buffer are executed in one batch and removed from the buffer.

265.7 Index Query Preprocessor

This entry has one repeatable property (use array syntax if you're defining via OSGi configuration file):

`fieldNamePatterns`: Fields with names matching the patterns set here are treated as non-analyzed keyword fields. Instead of scored full text queries, matching is performed by non-scored wildcard queries. This is a resource intensive operation that degrades search engine performance as indexes grow larger. For substring matching, relying on the ngram tokenizer usually performs better.

265.8 Reindex

This entry contains only one property:

`indexingBatchSizes`: Set the number of documents indexed per batch for model types that support batch indexing. Defaults to 10000. For models with large documents, decreasing this value may improve stability when executing a full re-index.

265.9 Engine Helper

This entry has one repeatable property (use array syntax if you're defining via OSGi configuration file):

excludedEntryClassNames: Exclude an asset type from being searched in the catchall query constructed for the Search application. For example, fields of the Organization asset must be indexed to be searchable from the Users and Organizations application, but should not be searched in the Search application. Thus, Organizations are added to `excludedEntryClassNames`.

265.10 Permission Checker

Configure *pre-filtering permission checking* (permission checking on the search index) behavior. See here for more information on these properties:

- `includeInheritedPermission`
- `permissionTermsLimit`

265.11 Title Field Query Builder

Configure how search responds to matches on the Title field of a document.

Exact Match boost: Give an additional boost when searched keywords exactly match the title field of a document.

Maximum Expansions: Limit the number of documents to return when matching searched keywords to the title field as a phrase prefix. See Elasticsearch's Match Phrase Query documentation for more information.

265.12 Elasticsearch 6

Configure the connection between Liferay DXP and Elasticsearch 6. See here for more information on these properties:

- `clusterName`
- `operationMode`
- `indexNamePrefix`
- `numberOfIndexReplicas`
- `numberOfIndexShards`

- bootstrapMlocakAll
- logExceptionsOnly
- retryOnConflict
- zenDiscoveryUnicastHostsPort
- networkHost
- networkBindHost
- networkPublishHost
- transportTcpPort
- transportAddresses
- clientTransportSniff
- clientTransprtIgnoreClusterName
- clientTransportPingTimeout
- clientTransportNodesSamplerInterval
- HttpEnabled
- HttpCrsEnabled
- HttpCorsAllowOrigin
- HttpCorsConfigurations
- additionalConfigurations
- additionalIndexConfigurations
- overrideTypeMappings
- synchronizedSearch

265.13 Search Web

This entry contains one property: `classicSearchPortletInFrontPage`: Revert the default search experience from using the new Search Widgets to the classic Search Portlet that was standard in past releases. See [here](#) for more information.

265.14 Search Administration

In *Control Panel* → *Configuration* → *Search* there are two administrative UIs: Index Actions and Field Mappings

Index Actions

In Index Actions, re-index one of these:

- All indexable assets
- An individual indexable asset
- All spell check indexes

Field Mappings

The Field Mappings tab shows you all field mappings that are effective in the system, by index. Currently, you can view the mappings, copy them, zoom in or out, and view them with a dark theme. Look for added functionality to this UI in future versions.

265.15 Portal Properties

Portal properties are system scoped configurations as well. The Lucene Search portal properties configure low level search behavior. Review the properties and their descriptions and determine if they apply to your search requirements.

265.16 Site Scoped Search Configuration

Search isn't configurable at the Site Scope by the strict definition of Site Scoped Configuration. However, Search Pages influence site-specific search behavior. Commonly, Search Pages contain search widgets configured to search for all content within a particular Site.

In addition, the Header Search (the Search Bar embedded in every Site page by default), whether populated by the new Search Bar widget or the legacy Search portlet, is Site scoped. Only one instance of the Header Search application exists per Site, and configuring it in one page context configures it for the entire Site.

Because of the modularity of Search, there are some important configuration nuances to be aware of when using the new Search widgets:

- If the Header Search uses the Search Bar widget, its configuration always requires a *destination page* to be set, where Users are redirected to complete their search activity, interacting with the other Search widgets (Results, Facets, Suggestions etc.). Search destination pages are ordinary pages holding the Search widgets. You can have as many pages with Search widgets across the Site as you want.
- Unlike the legacy Search portlet, the new Search Bar widget is instanceable, so one page can contain multiple Search Bar widgets configured differently. All Search Bar instances must point to a Search Page to be effective.
- **Important:** if the destination Search Page has a Search Bar widget instance besides the embedded Header Search, the configurations of the Header Search take precedence over the page's widget instance.

Conversely, searching from a Search Bar widget instance on other pages honors their configurations, even if they differ from the Header Search configuration.

See the documentation on configuring of a Search Bar for more information.

265.17 Widget Scoped Search Configuration

Several search widgets are available, and each one has its own configuration options:

Search Results Configure how search results are displayed. Read here for more information.

Search Bar Configure the behavior of how search keywords are processed. See here for more information.

Search Facets Configure each facet's behavior and URL parameters. See here for more information.

Search Options This is a special case, where configuring this widget defines page scoped behavior. Add the Search Option widget to a page and define two booleans for the Search Page:

- **Allow Empty Searches:** By default, failure to enter a keyword returns no results. Enabling this ensure that *all* results are returned when no keyword is entered in the Search Bar.
- **Basic Facet Selection:** By default, facet counts are recalculated after each facet selection. Enable this to turn off facet recounting.

Search Suggestions Suggest better queries and spell check queries. See here for more information.

Search Insights Add this to the Search Page to inspect the full query string that's constructed by the back-end search code when the User enters a keyword. Only useful for testing and development.

Custom Filter Add a widget to the page for each of the filters you'd like applied to the search results. Let search page users see and manipulate the filters or make them invisible and/or immutable.

Sort Let Users reorder the search results based on the value of certain keyword fields in the index. For example, show results in alphabetic order of the Title field. The default order is determined by the search engine's *Relevance* calculation.

LOW LEVEL SEARCH OPTIONS: SEARCHING ADDITIONAL OR ALTERNATE INDEXES


Low level search is a new concept in Liferay DXP version 7.2: it's a search that doesn't go through the Search and Indexing Framework, which is infrastructure used for searching documents in the Liferay Index.

A common use case for a low level search is to query an index other than the Liferay DXP index. By default, Search Pages search the Liferay DXP index, but you can also search another index, as long as it's in the same Elasticsearch cluster (this feature does not work with Solr).

Add the Low Level Search Options widget to a search page and configure it to direct the search to the alternate index. To search multiple indexes from the same page, you can add multiple Low Level Search Options widgets and configure each one with its own Index Name and Federated Search Key.


Searching alternate indexes is a low level operation that bypasses the Liferay DXP permission checking mechanisms, presenting whatever results the search engine returns. For this reason, only administrators can add and configure the Low Level Search Options widget.

To use the Low Level Search Options widget, add it to a Search Page:

1. Click the Add menu () on the page to open the Add Widgets menu.
2. Drag the Low Level Search Options widget (from the Search section), and drop it on the page.

It doesn't do anything unless you configure it.

266.1 Configuring Low Level Search

There are several configuration options for the widget. Access them by clicking the widget Options menu () → Configuration, or by clicking the hypertext URL in the widget body:

Configure additional low level search options in this page.

The image shows a configuration widget with five input fields, each with a question mark icon to its right. The fields are labeled as follows:

- Indexes**: A text input field for specifying alternative indexes to search.
- Fields to Return**: A text input field for specifying which stored fields to return from the search engine.
- Contributors to Include**: A text input field for specifying the IDs of registered search contributors to include.
- Contributors to Exclude**: A text input field for specifying the IDs of registered search contributors to exclude.
- Federated Search Key**: A text input field for specifying the key of an alternate search.

Figure 266.1: The Low Level Options widget has several configuration options.

Indexes: Enter the comma-separated names of the alternative indexes to search. Do not enter the standard Liferay index name.

Fields to Return: Enter the names of the stored fields to be returned from the search engine in a comma-separated list. Leave it blank to return all stored fields.

Contributors to Include: Enter the ids of registered search contributors to be included in this search in a comma-separated list of each `SearchRequestContributor`'s Fully Qualified Class Name (e.g., `com.liferay.docs.request.contributor.MySearchRequestContributor`). If not set, all registered search contributors are applied.

Contributors to Exclude: Enter the ids of registered search contributors to be excluded from this search, in a comma-separated list. If not set, all registered search contributors are applied.


Note: These *Contributors* are components implementing the `com.liferay.portal.search.spi.searcher.SearchRequestContributor` interface (provided by the `com.liferay.portal.search.spi` artifact), which is an extension point (SPI) that intercepts search requests and adds query parts.

Federated Search Key: Enter the key of an alternate search this widget is participating in. If not set, this widget participates in the default search. This value is usually the name of an application-defined index.

266.2 Example: Searching an Alternate Index

1. Whether testing on the default search page or creating a new Search Page, include the following widgets (removing extra widgets simplifies the exercise, but is not required for it to work):

- Low Level Search Options
- Custom Filter
- Search Bar
- Search Results

2. Configure all the widgets to participate in an alternate search, by opening the widget's Options menu () and clicking *Configuration*. For each, enter *liferay-0* in the Federated Search Key setting.

All the search widgets expected to react appropriately to the alternate search must be configured with the Federated Search Key. The following steps detail additional configuration.

3. Make an additional configuration in the Low Level Search Options widget, adding the index name of the alternate index:

Enter at least one index name in the *Indexes* setting. To follow this example, use *liferay-0*.

4. Configure the Custom Filter to use the search bar's default query parameter (*q*) and add a query to the search:

Enter *title* under field name to add the title field to the query.

Choose a Filter Query Type (e.g., Match) for the field.

Since you're overriding the default query to search an alternate index, there's nothing in the query by default. Add any query clauses using the Custom Filter widget(s).

If you're using *liferay-0* in your Federated Search Key and Indexes settings, search for *dynamic* in the search bar. You'll see results like this:

Now you're able to configure the out of the box search widgets to participate in searches against any Elasticsearch index in the cluster.

Dynamic Q

CUSTOM FILTER

Title v

Dynamic

4 results for **Dynamic**

| |
|---|
| Dynamic Data Mapping Service model.resource.com.liferay.configuration.admin.web.internal.model.ConfigurationModel |
| Dynamic Data Lists model.resource.com.liferay.configuration.admin.web.internal.model.ConfigurationModel |
| Dynamic Data Mapping Web model.resource.com.liferay.configuration.admin.web.internal.model.ConfigurationModel |
| Dynamic Data Lists Service model.resource.com.liferay.configuration.admin.web.internal.model.ConfigurationModel |

LOW LEVEL SEARCH OPTIONS

[Configure additional low level search options in this page.](#)

Figure 266.2: Configure the search page to search a different index.

SEARCH TUNING: SYNONYM SETS

Starting with 7.0 Service Pack 1, new search tuning features are available for administrative Users: Synonym Sets is one of them.

Synonym Sets are mappings that you (the admin) create, so that if a User searches for a certain keyword or phrase, the synonymous terms in your mapping are also searched. Matches to synonyms keywords are scored equally to matches with the exact keyword by the search engine.

Lunar Resort Use Case: Multiple content creators at the Lunar Resort write blogs about a variety of topics. Consistent terminology is a problem for some concepts. One writer might use the term “rover” for the vehicle that travels across the moonscape, while another uses “lunar cart” or “moon ATV”. As the portal administrator, you must ensure that the search experience is such that searching for any of those keywords returns all relevant results. Synonym Sets are a key ally in this pursuit.

267.1 Requirements and Limitations

Search tuning features like Synonym Sets are only supported when using Elasticsearch as the search engine. If you’re using Solr, make sure you disable the search tuning features (Synonym Sets and Result Rankings) when you upgrade your installation to Liferay DXP Service Pack 1 (Fix Pack 2).

As of the initial release (7.0 SP-1), Synonym Sets work with fields indexed in two locales: English and Spanish. Thus, the assets supporting localization out-of-the-box work with Synonym Sets. Technically, this means that synonym searches operate on fields indexed with the `en_*` and `es_*` suffixes. Read the multi-language search documentation to learn which native Liferay DXP assets/fields support localization in the search index. All asset types that index their data into English and Spanish are analyzed with a synonyms-aware analyzer and can be found during a synonym search.

The `⇒` format supported in Elasticsearch is not supported through the Synonyms Set UI.

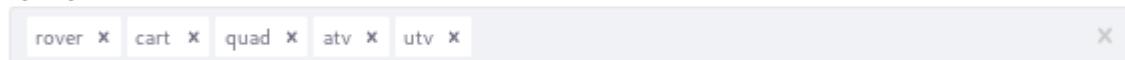
267.2 Creating and Managing Synonym Sets

Create a synonym set by adding as many synonymous keywords to a set as you want. Once the synonym set is saved, any searches in the same company scope (that's any site from the Virtual Instance where the synonyms were configured) take effect.

Create Synonym Set

Broaden the scope of search by treating terms equally using synonyms.

Synonyms



Type a comma or press enter to input a synonym.



Figure 267.1: Add as many synonymous keywords to a set as you'd like.

To create a synonym set,

1. Navigate to Control Panel → Search Tuning → Synonyms.
2. Click the Add button (+).
3. Enter the list of synonyms in the set. The input of a synonym is accomplished by clicking *Enter* or by entering a comma.
4. When the list is finished, click *Publish*.

The available synonym sets appear and can be managed in bulk or individually. The management options are to update a synonym set or delete one or more synonym sets.

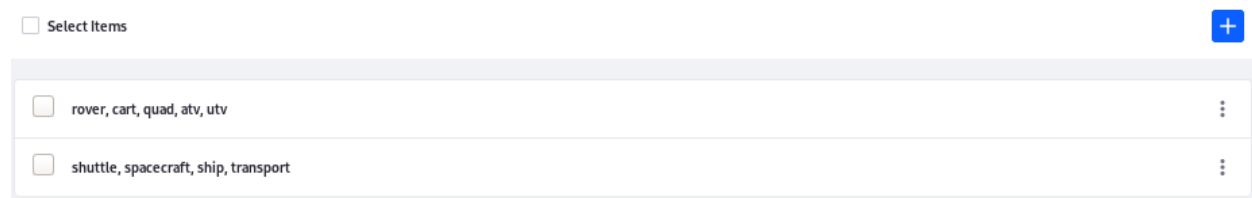


Figure 267.2: Synonym sets can be managed in bulk.

To edit or delete a single synonym set, click the Actions button (⋮) for the synonym set and choose Edit or Delete.

267.3 Using Synonym Sets

When you have a synonym set defined, the synonyms are ready for use. To test them, find a Search Bar anywhere in the virtual instance and enter a keyword from one of your synonym sets. Results matching the keyword and any synonym are returned in the Search Results widget.




Figure 267.3: The Blogs Entry does not contain the word “rover” but it can be matched because of a synonym set mapping “cart” as its synonym. The synonym is even highlighted.

267.4 Known Issues

There are several known issues for Synonym Sets. These are some of the most important ones:

LPS-100272: Reindexing permanently deletes all Synonym Sets. Please refer to the ticket for a way to backup and preserve (restore) Synonym Sets across reindex operations.

LPS-98126 Users can create duplicate Synonym Set entries and update other Synonym Sets unintentionally.

267.5 Related Resources

- <https://www.elastic.co/guide/en/elasticsearch/guide/current/synonyms.html>
- <https://www.elastic.co/guide/en/elasticsearch/reference/6.8/analyzer-anatomy.html>
- <https://www.elastic.co/guide/en/elasticsearch/reference/6.8/analysis-synonym-graph-token-filter.html>
- http://lucene.apache.org/core/7_7_0/analyzers-common/org/apache/lucene/analysis/en/EnglishPossessiveFilter.html

SEARCH TUNING: CUSTOMIZING SEARCH RESULTS

Starting with 7.0 Service Pack 1, new search tuning features are available for administrative Users: Custom Result Rankings is one of them.

Result Rankings provides a brute force method for intervening into the relevance scoring of the search engine, by doing these things:

1. Designate that certain results should appear at the top of the results if they are matched with a certain keyword. This is the idea of *pinning* results to the top of the results list.
2. By contrast, hide results that shouldn't appear in certain searches at all.
3. Add results that aren't normally returned by searching a certain keyword.
4. Re-order pinned results with a drag-and-drop interface.

Result Rankings let you pin, hide, and add search results for a given set of keywords. These rankings apply only to searches using the newer search widgets (Search Bar, Search Results, etc.). The rankings you customize do not apply to the legacy Search portlet results or to the individual application search bars.

Use Case: At the Lunar Resort website, visitors often search for activities, entering keywords like “rover races”, “atv rentals”, and “lunar golf”. For all of these, the Lunar Resort always wants a certain Content Page to appear at the top of the search results. This is the Activities page in the Lunar Resort where guests can find all of the resort’s adventurous offerings, including lunar rover races, ATV rentals, and information about golfing packages. By contrast, the Lunar Resort does not want the legal liability waiver form to appear during a search for fun activities: that’s a bridge to be crossed when guests sign up for the activity. It shouldn’t pollute a search for fun activities, even though it contains many of the keywords Users would search for. Result Rankings lets you *pin* the Activities Content Page to the top of the results and *hide* the liability waiver Web Content Article. In addition, a community member wrote a blog favorably reviewing the Lunar Resort, and you want that content added to searches for activities at the resort. This is a prime use case for Result Rankings.

The screenshot shows a search interface with a search bar containing 'atv'. On the left, there are three filter panels: 'Type' with options for Blogs Entry (1), Form Record (1), Page (1), and Web Content Article (2); 'User' with 'test test (5)'; and 'Last Modified' with options like Past Hour (3), Past 24 Hours (5), Past Week (5), Past Month (5), Past Year (5), and Custom Range... On the right, the search results are displayed under the heading '5 results for atv'. The results list includes: 'Announcing a New Cart Racing Course' (Blogs Entry), 'Lunar Rover Racing, ATV Rentals, Golf Packages, and More Lunacy!' (Web Content Article), 'Activities' (Page), 'Legal Waiver for Intentional or Accidental Dismemberment, Evisceration, and Everything Else' (Web Content Article), and 'New Entry for Form Rover Race Sign-Up' (Form Record). The 'Activities' and 'Legal Waiver' results are highlighted with red boxes, indicating they are the focus of the search tuning.

Figure 268.1: The Lunar Resort wants to tweak these results: pin the Activities page to the top, and hide the legal content entirely.

268.1 Availability

Search tuning features like Result Rankings are only supported when using Elasticsearch as the search engine. If you're using Solr, make sure you disable the search tuning features (Synonym Sets and Result Rankings) when you upgrade your installation to Liferay DXP Service Pack 1 (Fix Pack 2).

Results Rankings was added in 7.0 Service Pack 1.

268.2 Requirements and Limitations

Result Ranking entries are configured in a Virtual Instance, but are not applied only to that Virtual Instance. Instead, custom rankings made in one virtual Instance are shared across all Virtual Instances in the deployment, and even across separate deployments sharing an Elasticsearch cluster (in a multi-tenant scenario). Therefore, Result rankings shouldn't be used when connecting multiple Liferay DXP deployments to the same Elasticsearch cluster unless you intend for the same Result rankings to apply to every Liferay DXP deployment. See LPS-101291 for more information.

An existing Result Ranking cannot be renamed. Renaming requires recreating the ranking under a different name. See LPS-96357 for more information.

268.3 Creating and Managing Result Rankings

To manipulate result rankings, create a new *Alias* containing the keywords/search terms you want to intercept. Perform a search to get results (you can also do a separate search if you want to grab results that haven't even been returned during a natural search for the alias keywords). Once you have the results, choose to pin, hide, re-order, or add results as you please.

To create a new Result Rankings Alias:

1. Navigate to Control Panel → Search Tuning → Result Rankings.
2. Click the Add button (+).
3. On the New Ranking screen, enter one of the keywords or search phrases you want to intercept (it can be a phrase, instead of just one word; and don't worry, you can add more later) in the *Search Query* field.

Click *Customize Results*.

A search query is executed. The results are displayed and the tools for pinning, hiding, and adding results are made available. Re-ordering becomes possible after at least one result is pinned. First, consider whether to add one or more Aliases.

268.4 Adding Aliases

The Customize Rankings screen is ready to use, but any intervention only applies to the search query you initially entered in the New Ranking screen. To apply the customized rankings to additional search terms, add them as *Aliases*.

1. In the Aliases field, enter the search term to add as an Alias.

Warning: Do not use quotes in your alias terms.

2. To submit the search term as an alias, click Enter or a comma in the Aliases field. You can Add multiple aliases here.




Figure 268.2: Apply your custom rankings to matched results of additional search terms.

Note that results not manipulated manually here are returned as usual when the alias term is queried for in the Search Bar.

Now you can customize the rankings.

268.5 Activating and Deactivating Aliases


Available as of Liferay DXP 7.2, SP2

You can activate or deactivate existing aliases as you have need for them to take effect:

1. Find the Active alias and open its editor screen.
2. Click the toggle switch currently labeled Active.

Duplicate active aliases are not allowed, but you can deactivate an alias and then create a duplicate. After deactivating an alias, you can only reactivate it after deactivating any active duplicates.

268.6 Pinning and Hiding Results

To pin or hide rankings, hover over the result of interest: two icons appear, one for pinning and one for hiding. Click the one that applies. Otherwise click the Actions button (), and select *Pin Result* or *Hide Result*. Once you select either option, it's applied immediately. A pinned result moves to the top of the list, and a hidden result disappears. Repeat the action as many times as necessary.

If you're done customizing the results, click *Save*, and it's applied immediately.




Figure 268.3: Pin results to the top of the Search Results list.

268.7 Adding Results

To add a result that was not returned by searching for the first keyword or phrase, click the *Add Result* button and search for whichever asset you want to pin.

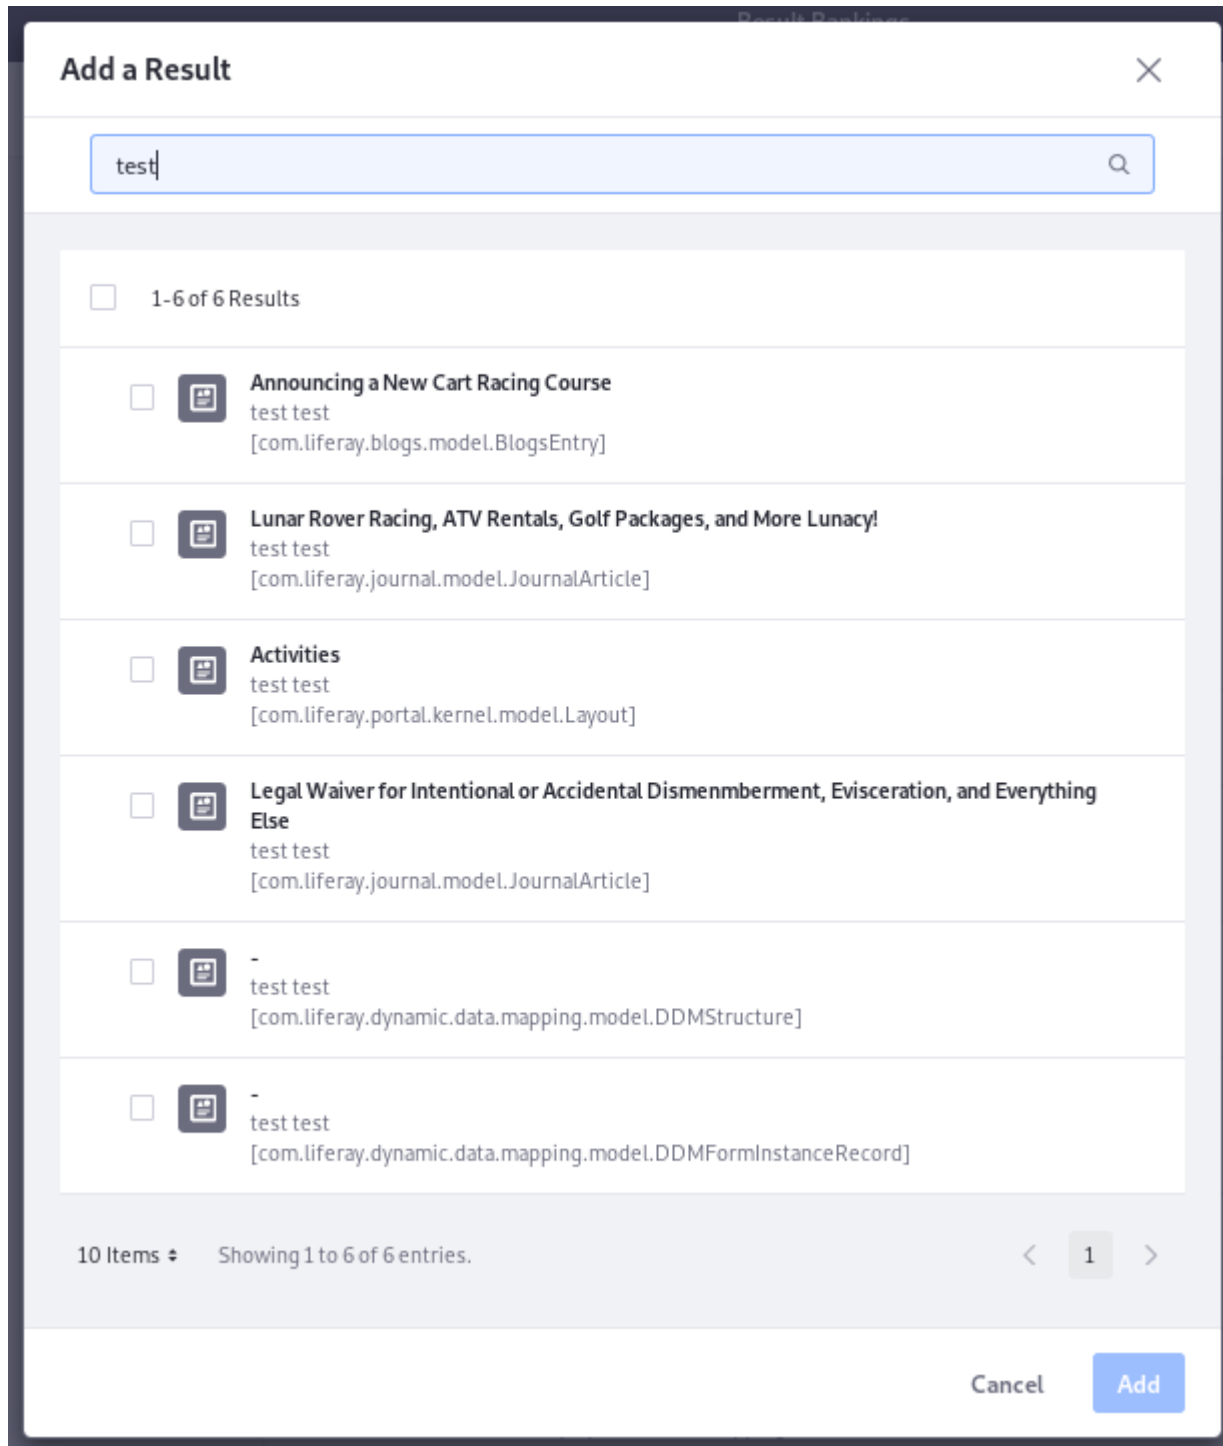


Figure 268.4: Add results that aren't normally returned.

Click *Save* if you're done customizing results.

268.8 Re-Ordering Pinned Results

To re-order pinned results (results that are not pinned cannot be re-ordered), click the handle icon, drag the result, and drop it in the preferred location in the list.

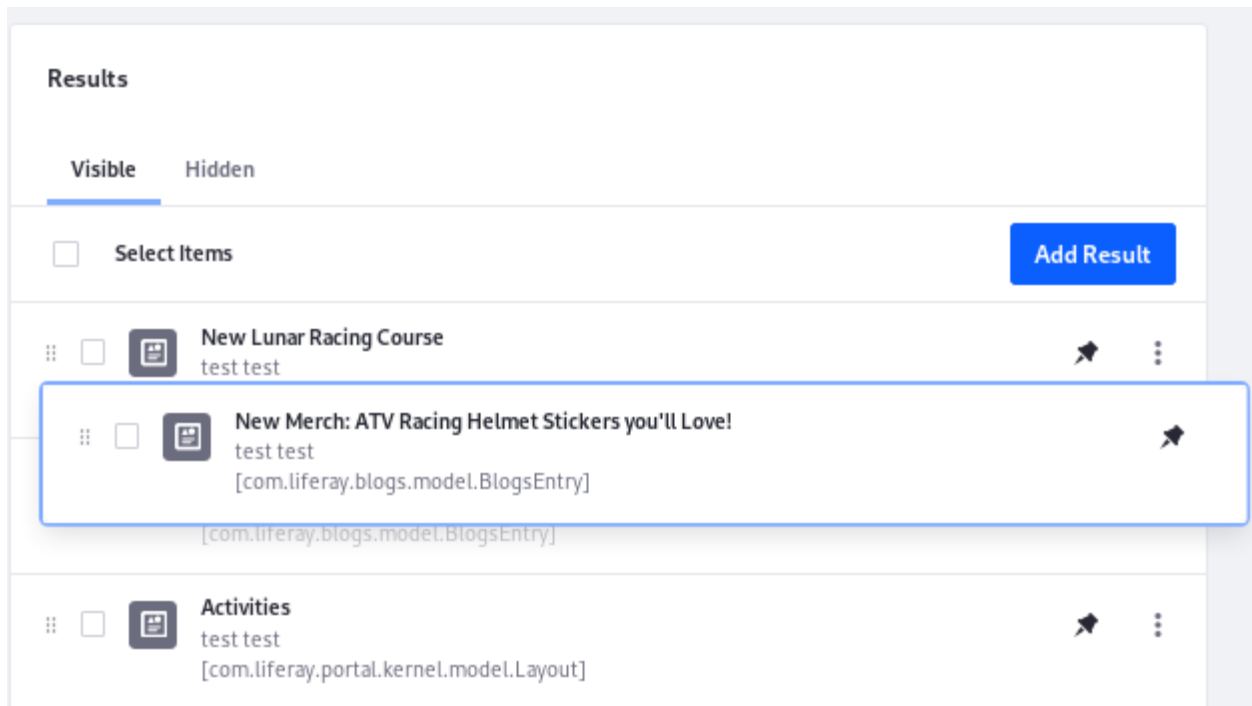


Figure 268.5: Re-order the pinned rankings if you want to emphasize or de-emphasize certain results.

Once finished customizing result rankings, click *Save*.

268.9 Result Rankings Scope and Permissions

Because configuration of Result rankings happens at the virtual instance scope, there are scoping and permissions behaviors to be aware of.

Scope is disregarded for pinned results: Pinned results existing in Site A always appear in searches from Site B, even if the Search Bar Scope is set to *This Site*.

Search from Result Rankings is global: When searching for results in Result Rankings admin, relevant results from all sites are returned.

Permissions are applied as usual: If a User doesn't have permission to see an asset, pinning it does not make it appear in the search results for that User.

268.10 Result Rankings Aliases versus Synonyms

Since both are new features without precedent in Liferay DXP, there can be confusion over Result Rankings Aliases and Synonyms. Synonyms expand the search to include results matched by additional (synonymous) keywords, so more results are returned if there are matches to the synonyms.

Result Rankings Aliases are just keywords that also have the particular ranking interventions applied to them. Only the searched keyword is matched to results, and then, the pins, hides, re-ordering, and additional results take effect after that.

These features don't interact in a predictable way. If you need synonym-like behavior in results rankings, define aliases for the keyword.

268.11 Known Issues

There are several known issues and planned improvements for Result Rankings. See LPS-99540 for the complete list.

FORMS

Liferay Forms gives you robust form building capability. For a complete list of the form fields available, visit the form fields reference article.

Because the complexity of use cases for Forms varies from a single input field to many pages of fields with different configurations, it makes sense to show you how to build and publish simple forms very quickly, and then show you all the additional features you can use for more complex use cases. Here's a sampling of the what the Forms application can do:

- Populate a Select or Radio field with a REST Data Provider
- Make a field appear based on the value of another field
- Add extra pages to the form
- Enable CAPTCHA for a form
- Store results in JSON
- Enable workflow for the form
- Redirect to a different URL after a successful form submission
- Send an email notification to administrators whenever a form is submitted
- Provide a default value (entered if left alone by the user) or a placeholder value (not entered if left alone by user) for each field
- Validate fields using a number of different criteria
- Redirect users to a success page after form submission
- Define Form Rules to create dynamic form behavior (for example, show or hide a field based on input in another field).
- Translate form text into any supported language.
- Create partial forms (with fields and other elements and specific configurations) and save them for reuse.
- Drag and drop fields onto the form layout.
- Duplicate a form instead of starting a similar form from scratch.

Despite this long list of more complex options, developing a simple, elegant form to suit basic needs takes little effort. The next article covers basic form building.

269.1 Forms and Lists

When you need a form, what you're really looking for is data. There are two applications for building forms to collect precisely the data you need:

1. **Liferay Forms:** The primary form building application is for the simplest one or two question survey to the most complex, multi-page, homeowners insurance application containing rules and lists populated by a REST data provider.
2. **Dynamic Data Lists (DDL):** Provides a user interface tool for building reusable form- and list-based applications intended for display on pages, using templates.

Kaleo Forms: If you're a Liferay Digital Enterprise customer, there's a third form building tool called Kaleo Forms. It integrates form building with workflow to create form-based business processes, like a Conference Room Checkout Form, or a Support Ticket Process so support tickets go through the proper channels on their way to resolution. Read more about Kaleo Forms in the workflow section.

269.2 Which Form Builder Should I Use?

Liferay Forms (also referred to as Forms) is a relatively new application, first appearing in Liferay DXP version 7.0. If you can use Liferay Forms for your use case, you should.

So the question "Which form builder should I use?" can be restated to "When should I use Dynamic Data Lists?"

- Use Dynamic Data Lists (DDL) if you need a way for users to enter data, *and* you need to display the data in the user interface.
- Use DDL if you need to style your lists and forms with templates.
- Use DDL if there's a field type you need that's not included (yet) in Liferay Forms. These are the field types included in DDL that *are not* in Liferay Forms at the time of this writing: - Color - Geolocation - Web Content - Link to Page

It's important to note that these (and more!) form field types will be included in future versions of the Liferay Forms application.

When all form building features are fully merged into Liferay Forms, the best features of DDL, all the new features of Liferay Forms, and all future improvements will be in one application. Now is the time to familiarize yourself with Liferay Forms and begin using it for all your form building needs, except for the narrow use cases described above.

CREATING AND MANAGING FORMS

The Forms widget can do a lot of things really well, but if you just need a simple form, how do you wade through all the features you don't need? Is your simple survey going to make you late for that lunch outing you've been planning with colleagues at that new shawarma place? No!

Let's create a simple yet elegant form, give access to the intended users, and get you on your way to lunch.

At The Lunar Resort, it's important to capture guests' feedback about their stay at the resort. After a (hopefully) safe journey home, guests should receive an email with a link to brief survey that prompts them to rate their stay from a list of selections, and add any additional comments they'd like in an optional field.

270.1 Viewing Forms

Whether creating a form or managing existing forms, it all starts in the same place: the Forms Application in your site's Content section. Access this in the Menu, first choosing the site to work in (for example, The Lunar Resort) and clicking *Content & Data* → *Forms*. The first thing you'll see is a list of existing forms (if there are any). This list is styled by the Display Style selector next to the Add button (+). By default, forms are displayed in List format.

There's also a Table format. Change the style for a single site right here in the Forms site menu application, or change the default display style for the system scope in Control Panel → Configuration → System Settings → Forms (in the Content section). Click the *Forms* entry and find the Default Display View property. Click *Update* and your changes are propagated to all sites.

270.2 Building a Form

Guest Survey

Tell us how your stay was!

All fields marked with * are required.

Rate your visit to The Lunar Resort*

Choose an Option

Want to tell us more?

Submit

Figure 270.1: Get feedback from guests of The Lunar Resort.

Filter and Order Search for: +

| | | | |
|--------------------------|--|---|---|
| <input type="checkbox"/> | | Liability Waiver Form Do not blame us for evisceration, lacerations, incineration, etc. ID: 42153 Modified Date: 4 Seconds Ago | ⋮ |
| <input type="checkbox"/> | | Employment Application Form ID: 42011 Modified Date: 1 Minute Ago | ⋮ |
| <input type="checkbox"/> | | The Lunar Resort Guest Survey ID: 45451 Modified Date: 1 Minute Ago | ⋮ |

Figure 270.2: Forms are displayed in List format by default.

To add a new form,

1. Click the *Add* button (+). The form builder view appears.
2. Name the form. Replace *Untitled Form* with *Guest Survey*.
3. For the description text enter *Tell us how your stay was!*
4. Add the fields. Click the *Add* button (+) to open the sidebar if it's not already opened.
5. Drag a *Select from List* field onto the form builder and configure it like this:

Label: *Rate your visit to The Lunar Resort.*

Help Text: Leave this blank for now. If you want a subheading for your field to provide additional guidance, this would be useful.

Turn on the *Required Field* selector. At a minimum, this form must capture whether guests like their stay or not.

Leave the manual option checked for creating the list of selections. To learn about populating the field with a data provider, read [here](#).

Add these options:

- *It was out of this world!*
- *I had a good time.*
- *I'd rather go to the beach.*
- *I'll never come back.*

Typing in one of the fields automatically adds another blank selection line. Just leave the last one blank when you're done.

6. To see additional options, click *Properties*. Close the sidebar when finished.
7. Add a text field, using the same procedure you used for the select field.

Label: *Want to tell us more?*

Help Text: Leave this blank again to give the form a consistent look.

My text field has: Choose *Multiple Lines*. Let guests prattle on about their stay if they want to.

Required Field: Leave this unselected. Only require guests to fill out the select field and leave this one as optional.
8. Close the sidebar.
9. In the form builder, you can see the way the fields are laid out on the form page.
10. When the form is finished, click *Save Form*. It's also auto-saved every minute by default.
11. Click *Publish Form*. A dedicated URL to the form is generated, but nobody has the URL yet.

Now your form can be added to a page, and Lunar Resort guests can be emailed and provided with a link to the page where the form is displayed.




Figure 270.3: You can choose from nine field types when creating forms.

The screenshot shows a form builder interface. At the top, there are tabs for 'Builder' and 'Rules', and a blue '+' button in the top right corner. Below the tabs, there is a 'Default Language' section with a dropdown menu showing 'English (United States)' and a '+ Add Translation' button. The main content area is titled 'Guest Survey' with the subtitle 'Tell us how your stay was!'. Below this, there is a preview of the form layout. The preview shows a page titled 'Untitled Page (1 of 1)' with a description 'Add a short description for this page. All fields marked with * are required.' The first field is a dropdown menu labeled 'Rate your visit to The Lunar Resort *' with the placeholder text 'Choose an Option'. Below this is a text area labeled 'Want to tell us more?'. At the bottom of the preview, there are three buttons: 'Unpublish Form' (blue), 'Save Form' (white), and 'Preview Form' (blue).

Figure 270.4: The form builder page lets you preview your form layout, add a page to the form, or add some more fields.

270.3 Accessing Forms

Once the form is developed and published, there are two options for getting the published form to targeted users:

1. Place the Form widget on a site page. This approach lets users navigate to the page in the site.
2. Copy the dedicated form URL and provide it to users (for example, via email). This limits access to the form to only those users who have the direct link.

To display the form on a site page in The Lunar Resort site:

1. Add a page to the site (choose Full Page Application for the page type if you only want the form application on the page). Call it *Guest Survey*.

General SEO Look and Feel Advanced

Name *

Guest Survey en-US

Friendly URL ?

http://localhost:8080/web/guest /guest-survey en-US

Create a page with one column that displays a single full page application.

Full Page Application

Form

Save Cancel

Figure 270.5: Add a page for guests to view and fill out your new form.

2. Add the Form widget to the page if you've chosen a Widget Page. If you used a full page application, use the page configuration to choose *Form* from the Full Page Application dropdown.
3. Once the Form widget is on the page, click *Select Form*, choose the *Guest Survey* form, and click *Save*.
4. Close the *Form—Configuration* dialog window and your form is ready for Lunar Resort site users.

To display the form on a dedicated page accessed only by its link:

1. In the form builder, click *Publish Form* if you haven't already.

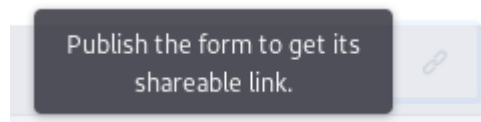


Figure 270.6: You must first publish a form before you can get a shareable link.

2. Once published, click the link icon at the top right of the builder.

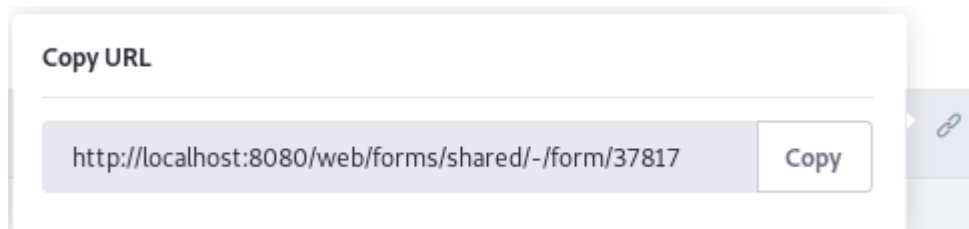


Figure 270.7: Copy the link to your form.

3. Once you get the link out to users, it's showtime.

Next you'll learn how to view the form entries. Since there aren't any yet, fill out and submit the form a few times.

Now you know the basics of creating and managing forms, but this presentation didn't do the Forms application justice. It's much more powerful than hinted at here. The remaining articles in this section immerse you in more advanced form building features.

Guest Survey

Tell us how your stay was!

All fields marked with * are required.

Rate your visit to The Lunar Resort *

Choose an Option

Want to tell us more?

Figure 270.8: Lunar Resort guests can use a simple form to record their feelings about the resort.

MANAGING FORM ENTRIES


Once users begin submitting form entries, you'll want to do these things with them:

- View form entries
- Export form entries
- Delete form entries

Start by learning how to access and view the entries.

271.1 Viewing Form Entries

When users fill out forms, they're generating data. You'll want to see that data at some point.

1. From the Menu, navigate back to the *Content* → *Forms* section of The Lunar Resort site.
2. Click the *Actions* () button for the form and select *View Entries*.










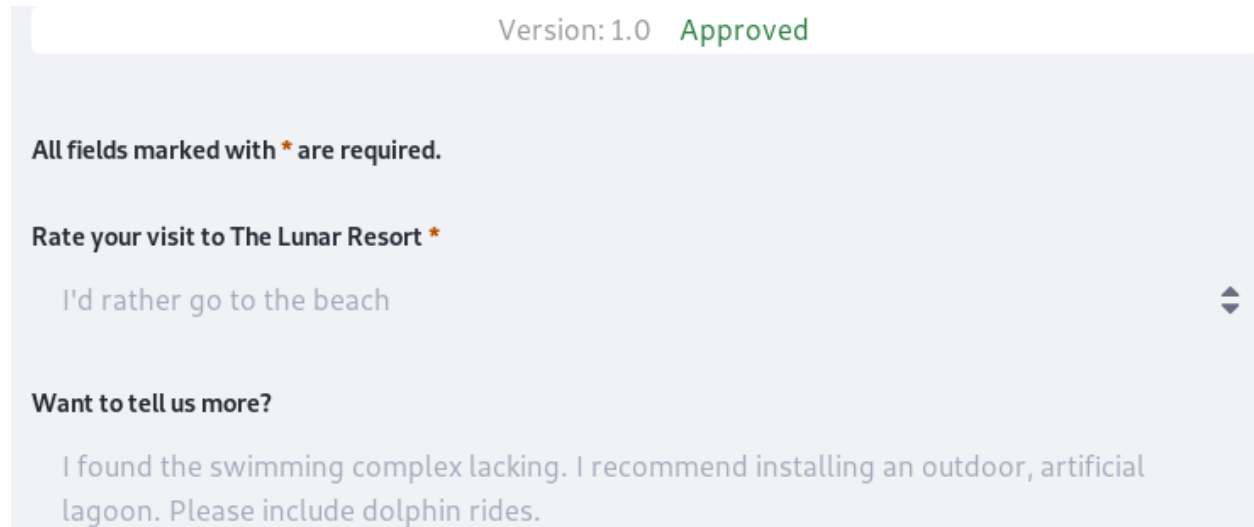
| | Rate your visit to The Lunar Resort | Want to tell us more? | Status | Modified Date | Author | |
|---|-------------------------------------|--|-----------------------|----------------|-----------|---|
|  | It was out of this world! | Best vacation ever! Marvin the Martian ... | Approved | 4 Minutes Ago | Test Test |  |
|  | I'd rather go to the beach | I found the swimming complex lacking. ... | Approved | 2 Minutes Ago | Test Test |  |
|  | I had a good time | Good time, but the earth kind of looks li... | Approved | 44 Seconds Ago | Test Test |  |
|  | I'll never come back | Having astrophobia didn't help matters.... | Approved | 4 Seconds Ago | Test Test |  |

Figure 271.1: You can view the entries right in the Forms application.

3. What if you have a lot of form fields and you can't see all the data for each entry in the search container? Just click the *Actions* () button for the entry and select *View*. You're shown all the specifics for that form entry.



Version: 1.0 Approved

All fields marked with * are required.

Rate your visit to The Lunar Resort *

I'd rather go to the beach ▾

Want to tell us more?


I found the swimming complex lacking. I recommend installing an outdoor, artificial lagoon. Please include dolphin rides.


Figure 271.2: You can view a single entry right in the Forms application.

Viewing entries is great, but this is serious data we're talking about. You might need to get all the entries into a spreadsheet so you can work with them.

271.2 Exporting Form Entries

So you need to put your form entries in a spreadsheet to do spreadsheet things with them? No problem.

1. Navigate to the Forms application in The Lunar Resort site's Content & Data section again.
2. Click the *Actions* () button and select *Export*.
3. Choose a File Extension. You can export entries in CSV, JSON, XLS, or XML formats by default. For this example, pick CSV.
4. Click *Okay*, and open the file or save it locally. Open it with your favorite spreadsheet program and verify your form entries.

Note: The Forms application itself has an *Import/Export* window accessible from the application's Configuration menu (). This is how you import and export the application configuration

| | A | B | C | D | E |
|---|-------------------------------------|---|----------|----------------|-----------|
| 1 | Rate your visit to The Lunar Resort | Want to tell us more? | Status | Modified Date | Author |
| 2 | I'll never come back | Having <u>astrophobia</u> didn't help matters. Recommend moving the resort to earth. | Approved | 5/9/18 7:50 PM | Test Test |
| 3 | I had a good time | Good time, but the earth kind of looks like a big pizza, which made me hungry. Please get a pizza shop at the resort. | Approved | 5/9/18 7:50 PM | Test Test |
| 4 | I'd rather go to the beach | I found the swimming complex lacking. I recommend installing an outdoor, artificial lagoon. Please include dolphin rides. | Approved | 5/9/18 7:48 PM | Test Test |
| 5 | It was out of this world! | Best vacation ever! Marvin the Martian was the best Lunar Sherpa we could have imagined! | Approved | 5/9/18 7:46 PM | Test Test |

Figure 271.3: You can export entries as CSV, JSON, XLS, or XML.

and its data (forms and form entries). The file format for this type of import and export is a LAR file. For more information, see the article on importing and exporting application content.


There's a system level setting to determine whether administrators can export entries in CSV format:

1. Go to Control Panel → Configuration → System Settings and click the *Forms* category in the CONTENT AND DATA section.
2. Click the *Forms* entry under SITE SCOPE.
3. The CSV Export property has three options:
 - *Enabled* to enable CSV Export without a warning
 - *Enabled (Show Warning)* to enable CSV Export with the following warning to administrators: This CSV file contains user supplied inputs. Opening a CSV file in a spreadsheet program may be dangerous.
 - *Disabled* to turn off CSV Export.

Once you export a batch of form entries, it can make sense to delete them from the database.

271.3 Deleting Form Entries

What if you export a form's entries and now you want to remove them from the Liferay database? It's easy to delete all of a form's entries at once.

1. Navigate back to the Forms application In The Lunar Resort Content section.
2. Click the *Actions* () button next to the Guest Survey form and select *View Entries* again.
3. Select all entries by checking the box next to *Filter and Order*. An X appears in the top right corner of the Form Entries screen. Click it.

If you just wanted to delete a single entry, select that entry by checking its box; then delete it.




Figure 271.4: Delete all form entries in one fell swoop.

If you're worried about deleting everything irrecoverably by accident, don't worry. You must confirm the deletion in a dialog box that pops up after clicking the trash can.

Now you can create basic forms and manage the entries. Keep reading in this section to learn about the many additional form building features available to you.

FORM FIELD TYPES

A form without fields is no form at all. To meet your form-building needs, Liferay Forms provides useful and highly configurable field types.

Paragraph: This is static text on the form. Users do not enter data into form text fields. The form creator enters text that form users see displayed on the form. This is useful for longer instructions.

Text Field: Users enter text into these fields. For example, a Full Name field is a text field. By default, a text field keeps all input on a single line of text. To accommodate longer responses, choose the multi-line setting when configuring the text field as in this example. You can put limits on the text users can enter (e.g., numbers from 1-10, email addresses, etc.) by using the text field's validation options.

Select from List: Users select one option (or more, if configured to allow it) from a list of choices. Choices are entered manually or are automatically populated by a data provider. For example, a Country of Residence field can be selected from list field populated by a Countries of the World data provider.

Single Selection: Using a radio button, users select one option from a list of options displayed on the form.

Date: Users select a date using a date picker. For example, a Birth Date field uses the Date field type.

Multiple Selection: Users select one or more options from check boxes (or toggles, if configured).

Grid: Using radio buttons, users select from options laid out in rows and columns. One selection can be made per row. This is useful when the same response metric is needed for multiple questions. For example, a product survey form might ask users to rate a list of their product's characteristics as *Wonderful*, *Pretty Good*, *Not So Good*, or *Awful*.

Numeric: Users enter numeric data (integers or decimals) into numeric fields. Non-number input is not accepted. For example, configure a numeric field that accepts integers to ask users how many pets they have.

Upload: Users can select a file from the Documents and Media library or upload a file from their local filesystems.

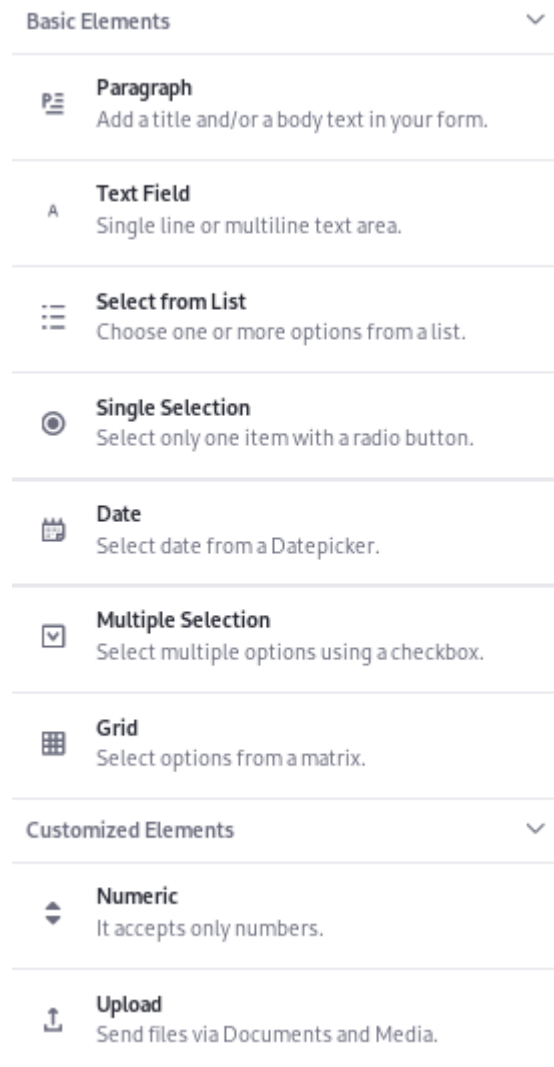


Figure 272.1: There are many useful out-of-the-box form field types.

Useful Form Text

Here is some text to help you understand the form that you are filling out.

Figure 272.2: Use Paragraph fields to enter longer instructions on Form Pages.

User Text Entry

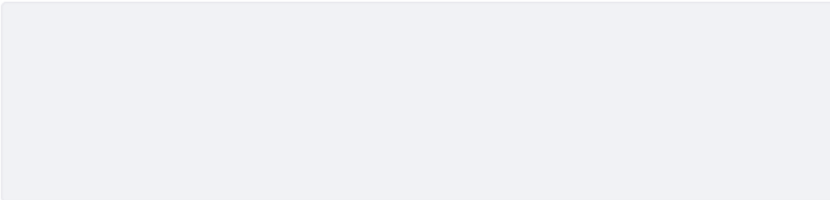


Figure 272.3: Text fields can be single line or multi-line.

Choose an Option

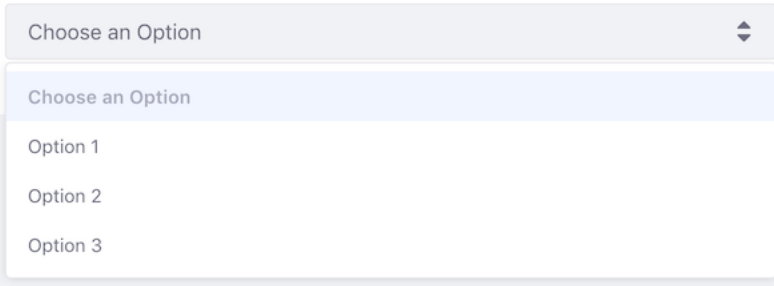


Figure 272.4: Use a select from list field to let Users choose predefined options.

Choose an option

Option 1 Option 2 Option 3

This time with radio buttons

Figure 272.5: Single selection fields allow only one selection.

What is today's date?

__/__/__

This way we know you don't have a head injury.

July 2018

| Su | Mo | Tu | We | Th | Fr | Sa |
|----|----|----|----|----|----|----|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| 8 | 9 | 10 | 11 | 12 | 13 | 14 |
| 15 | 16 | 17 | 18 | 19 | 20 | 21 |
| 22 | 23 | 24 | 25 | 26 | 27 | 28 |
| 29 | 30 | 31 | 1 | 2 | 3 | 4 |
| 5 | 6 | 7 | 8 | 9 | 10 | 11 |

Figure 272.6: Date fields show a date picker so Users enter a valid date.

We were going to ask you to check boxes



Check 1



Check 2



Check 3

But then we switched it on you

Figure 272.7: A multiple selection field can use a toggle.

Characteristics

| | Wonderful | Pretty | Good | Not so Good | Awful |
|--------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Paragraph | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Text Field | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Select from List | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Radio Buttons | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Date | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Multiple Selection | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Choose a characteristic for each row

Figure 272.8: Grid fields use the same options (columns) for multiple categories (rows).

How many pets do you have?

If it's more than 5 we might start judging you.

Figure 272.9: Numeric fields accept only numeric input.

Upload a picture of your pets

All pets will be rated according to the above grid criteria.

Figure 272.10: : Upload fields let Users attach files to the form.

FORM RULES

Chickens don't follow rules well, but dogs do. If you're skeptical, try teaching your chicken to sit on command or herd sheep. Better yet, get a team of chickens to pull a sled in the Iditarod. The Forms application is much more like the dog than the useful (southwestern omelet anyone?) but untrainable chicken, and it's only getting more trainable as time passes.

Form rules are a good example of the trainable nature of the Forms application. With form rules, you can train your form fields to behave as you wish. There are several things you can make them do:

Show/hide Based on a predefined condition, set the visibility of a form field.

Enable/disable Use a predefined condition to enable or disable a field.

Require Use a predefined condition to make a field required.

Jump to Page Based on user input, skip over some form pages directly to a relevant page. *This rule doesn't appear in the rule builder until a second page is added to the form.*

Autofill with Data Provider Use a data provider to populate fields when a condition is met in another field.

Calculate Populate a field with a calculated value using data entered in other fields.

Form rules are for changing fields and form elements by acting on conditions.

If [condition] do [action].

If you're not already familiar with the Forms application, start here. Once you know how to create forms, add and configure fields, and then publish forms, come back here and learn about form rules.

273.1 The Anatomy of a Form Rule

Each rule consists of one or more conditions and actions.

Conditions determine whether any actions are executed.

Actions determine what happens if the condition is met.

Rules are stored in the database in JSON format by default.

273.2 Creating Form Rules: Rule Builder

Once you create a form and lay out its fields, you're ready to set up rules in your form:

1. Save the form.
2. Open the Rule Builder by clicking the *Rules* tab at the top of the *Edit Form* screen.
3. In the rule builder view, you can now begin developing your form rule. Click the *Add* button (+) to get started.

Rule
Define condition(s) and action(s) to change fields and elements on the form.

Condition OR ▾

If Choose an Option ▾ Choose an Option ▾

Actions

Do Choose an Option ▾

Save Cancel

Figure 273.1: The Rule Builder gives you a handy interface for creating dynamic form rules.

Before looking at each type of rule condition and action you can use to develop rules, consider the *OR* selector box at the right side of the *Condition* (it's grayed out and unusable at first). You can choose *OR* or *AND* here, depending on what relationship the conditions should have with the action.

OR The action is triggered if *any* of the conditions you specify evaluates to *true*

AND The action is triggered only if *all* the conditions you specify evaluate to *true*

This box becomes usable once you click the Add button (+) to add an extra condition.

273.3 Conditions

Conditions are the gatekeepers of form rules. If the condition's *if statement* evaluates to *true*, the action is triggered. If it evaluates to *false*, no action happens.

A condition checks whether one field's value

- *Is equal to* a specific value or another field's value.
- *Is not equal to* a specific value or another field's value.
- *Contains* a specific value or another field's value.
- *Does not contain* a specific value or another field's value.
- *Is empty*. This assumes you want to do something if a field *is* empty.
- *Is not empty*. This assumes you want to do something as long as a field is *not* empty.

One exception to this is the User condition, which is the last option in the Condition dropdown menu.

The User condition doesn't act on a field at all. It checks whether a User belongs to a certain role. For example, if the condition

If User *belongs to* Administrator
evaluates to *true*, an action is triggered.

A condition is the gateway into a form rule, but actions define what actually happens when the condition evaluates to *true*. The remaining articles discuss the available actions and demonstrate their use.

ACTION: SHOW AND HIDE

With a show and hide rule, you use one or more conditions to determine whether to show or hide a field if the condition evaluates to *true*.

To set this example up, add these fields to a form:

- *I am 18 Years Old or Older*, a required single selection field with two options: *Yes* and *No*.
- *Legal Guardian Email Address*, a text field that accepts valid email addresses (use text field validation to dictate input type).

Example: If you're under 18 years old, you need the approval of a legal guardian to drive a sled in a sled dog race (even if you're racing chickens, not dogs). The form for registering your chicken team asks you the age of the driver. If you enter a number less than 18, the Legal Guardian Email Address field appears.

To configure a Show/Hide rule,

1. Open the Rules tab of the Edit Form page and click the Add (+) button.
2. Define the rule:
 - If the *I am 18 years old or older* field is equal to the Value *No*, show the *Legal Guardian Email Address* field.
 - Save the rule.

Now the *Legal Guardian Email Address* field is only displayed in the form if the user selects *No* in the *I am 18 years old or older* field.

Show rules let you keep a field hidden until some condition is met.

Rule
Define condition(s) and action(s) to change fields and elements on the form.

Condition OR ▾

If I am 18 Years Old o... ▾ Is equal to ▾ Value ▾ No ▾

Actions

Do Show ▾ Legal Guardian Em... ▾

Save Cancel

Figure 274.1: Build form rules quickly by defining your conditions and actions.

If Field I am 18 Years Old or Older *is equal to* Value No ⋮

Show Legal Guardian Email Address


Figure 274.2: Once a rule is saved, it is displayed so that you can easily understand what it does.

ACTION: REQUIRE

Use a require rule to make a field required based on one or more conditions.

Example: If you are following the example, you already set up a *show* rule, where a *Legal Guardian Email Address* field appears if the user selects *No* in the *I am 18 years old or older* field. You also want to make the *Legal Guardian Email Address* field required.

To configure a require rule,

1. Edit the Show Rule configured above. Open the Rules tab of the Edit Form page and click the kebab menu () for the rule, and then click *Edit*.
2. Add an Action to the rule:
 - If the *I am 18 years old or older* field is equal to *No*, show the *Legal Guardian Email Address* field and make it required.
 - Save the rule.

Require rules let you require fields based on input from other fields.

Rule

Define condition(s) and action(s) to change fields and elements on the form.

Condition AND ▾

If I am 18 Years Old o... Is equal to Value No

+

Actions

Do Show Legal Guardian Em... 🗑️

Do Require Legal Guardian Em...

+

Save
Cancel

Figure 275.1: Build form rules quickly by defining your conditions and actions.

If Field I am 18 Years Old or Older *is equal to* Value No ⋮

Require Legal Guardian Email Address ,

AND Show Legal Guardian Email Address

Figure 275.2: Once a rule is saved, it is displayed so that you can easily understand what it does.

ACTION: ENABLE AND DISABLE

Use an enable/disable rule to make a field editable based on one or more conditions.

Example: Part of the race registration fee pays for dog food. You don't have to feed your chicken team with the provided dog food though. There's a single selection field that asks *Would you like to use the provided dog food?*. If you select *Yes*, you can select how much food, in US pounds, you'll need for your team throughout the race. Since you're racing chickens, you'll select *No*, and the *Amount (US lb.)* field is disabled.

To follow the example, add a single selection field *Would you like to use the provided dog food?* with two options: *Yes* and *No*.

Add a numeric field called *Amount (US lb.)* and make it an Integer. Use field validation to make sure it's not greater than 100.

To set up the enable/disable rule,

1. Open the Rules tab of the Edit Form page and click the Add (+) button.
2. Define the rule:
 - If the *Would you like to use the provided dog food?* field is equal to *Yes*, enable the *Amount (US lb.)* field.
 - Save the rule.

Now users can fill out the amount of dog food they'll need only if they specify that they do indeed want to use the provided food.

Rule

Define condition(s) and action(s) to change fields and elements on the form.

Condition AND ▾

If

+

Actions

Do

+

Figure 276.1: Build form rules quickly by defining your conditions and actions.

If Field *is equal to* Value ⋮

Enable

Figure 276.2: Once a rule is saved, it is displayed so that you can easily understand what it does.

ACTION: JUMP TO PAGE

Use a Jump to Page rule to navigate automatically to a specific page in the form based on one or more conditions. This is useful if some pages don't apply to all the form's users. Even fields marked as required on the skipped pages can be successfully skipped using this rule.

This action doesn't appear in the rule builder unless the form has multiple pages.

To follow the example here,

1. Create a second form page called *Team Information*.
2. On the *first* form page, create a single selection field labeled *Are you a returning racer, with the same team?* with two options: *Yes* and *No*.
3. Create a text field on the *second* form page called *Animal Name*.
4. Create a third form page called *Final Confirmation*.

Example: There's a question on the *Team Information* page of the dog sled race registration form that asks *Are you a returning racer with the same team?* If you select *Yes*, when you click the form's *Next* button, you skip to the final page of the form, since there's no need to fill out your animal's name again. Their monogrammed T-shirts will be ready at the start of the race.

To configure the Jump to Page rule,

1. Open the Rules tab of the Edit Form page and click the Add (+) button.
2. Define the rule:
 - If field *Are you a returning racer, with the same team?* is equal to the Value *Yes*, Jump to Page *Final Confirmation*.
 - Save the rule.

Rule

Define condition(s) and action(s) to change fields and elements on the form.

Condition AND ▾

If Are you a returning... Is equal to Value Yes

+

Actions

Do Jump to Page 3 Final Confirmation

Save
Cancel

Figure 277.1: Build form rules quickly by defining your conditions and actions.

If Field Are you a returning racer with t... **is equal to Value** Yes ⋮

Jump to Page 3 Final Confirmation

Figure 277.2: Once a rule is saved, it is displayed so that you can easily understand what it does.

Once the form User fills out the first page and clicks *Next*, the rule condition will evaluate the answer to the field and either proceed to the next page or take the action of skipping to the page indicated in the rule.

If you use an *is not equal to* condition for form fields on two different pages, the condition is checked after leaving the page of the first form field, and evaluates to *true* since there's a value in the first field and no value in the second field. It's best to use this condition with fields existing on the same page.

ACTION: AUTOFILL

Autofill rules let you change the selection options of another field based on the value entered into a related field. A data provider's output is used to populate a field, as long as the condition is met.

Before configuring an autofill rule, set up a data provider. That's how autofilled fields are populated. Pay careful attention to the input and output parameters you choose when setting up the rule.

To follow this example:

1. Set up a data provider using the `get-countries` JSON web service. If you're running Liferay DXP at `localhost:8080`, you can access this web service here:

```
http://localhost:8080/api/jsonws?contextName=&signature=%2Fcountry%2Fget-countries-0
```

Make sure the output parameter is set to `$.nameCurentValue`. If you're unsure how to do this, first read the article on [Data Providers](#).

2. On the last form page, add two fields:
 - A Single Selection field called *If I win I'd like my award to be:*, with two choices: *Cash* and *All Expenses Paid Vacation*.
 - A Select from List field called *Choose a Destination Country*. Under *Create List*, choose *From Autofill*.

Example: Before submitting the race registration, let users decide whether they want a cash prize or an all-expenses-paid vacation. If they choose the vacation, populate the Choose a Destination Country with output from the data provider.

To configure an Autofill rule,

1. Open the Rules tab of the Edit Form page and click the Add (+) button.
2. Define the rule:

Figure 278.1: Build form rules quickly by defining your conditions and actions.

- If field *If I win I'd like my award to be* is equal to the Value *All Expenses Paid Vacation*, Autofill the *Choose a Destination Country* field from the *countries* data provider (note that you might have named this differently when setting it up).
- Save the rule.

Figure 278.2: Once a rule is saved, it is displayed so that you can easily understand what it does.

278.1 Using Inputs with Autofill

The above example is simple, using only an Output to autofill a Select from List field if another field has a certain value. Many times, the response from the REST provider must be filtered before display in the Select from List field. For this, a Data Provider Input field is required. For example,

to configure an autofill rule to display the countries of the world filtered by a Region field (for example, Americas, Europe, or Oceania),

1. Create a Data Provider.

Name: restcountries

URL: `https://restcountries.eu/rest/v2/region/{region}?fields=name`

Inputs: Fill in the Label (*region*), Parameter (*region*), and Type (*Text*).

Outputs: fill out a Label (*name*), Path (*\$.name*), and Type (*List*).

To understand more about these values, read the Data Provider documentation.

The screenshot displays the configuration interface for a Data Provider. It includes the following elements:

- URL:** A text field containing the URL `https://restcountries.eu/rest/v2/region/{region}?fields=name`.
- User Name:** A text field with the placeholder text "Enter a user name."
- Password:** A text field with masked characters (dots).
- Support filtering by keyword:** A toggle switch that is currently turned off.
- Cache data on the first request:** A toggle switch that is currently turned off.
- Timeout:** A text field containing the value "1000".
- INPUTS:** A table with columns for Label, Parameter, Type, and a Required checkbox.

| Label | Parameter | Type | Required |
|--------|-----------|------|--------------------------|
| region | region | Text | <input type="checkbox"/> |
- OUTPUTS:** A table with columns for Label, Path, and Type.

| Label | Path | Type |
|-------|---------|------|
| name | \$.name | List |

Figure 278.3: Create a data provider for the autofill rule.

2. Create a form with these fields:

Text: Use the Label *Region*.

Select from List: Label it *Country*, and choose *From Autofill* under Create List.

3. Configure the Autofill rule.

Condition: *If Region Is not Empty*

Action: Do **Autofill** From Data Provider restcountries, Data Provider's Input: region—*Region*, Data Provider's Output: name—*Country*.

The screenshot shows a form builder interface for a service named 'restcountries'. At the top left, there is a language selector set to 'English (United States)' with a '+ Add Translation' button. The main workspace is titled 'restcountries' and contains two form sections: 'Untitled Page (1 of 1)' with a description 'Add a short description for this page.' and a 'Region' text field, and a 'Country' select field with the placeholder text 'Dynamically Loaded Data'. At the bottom of the workspace are three buttons: 'Unpublish Form', 'Save Form', and 'Preview Form'. On the right side, there is a configuration panel for the 'Country' field. It includes a 'Label' section with the text 'Country', a 'Field Name' dropdown set to 'select', a 'Help Text' section with the text 'Enter help text.', a 'Required Field' toggle switch that is currently turned off, and a 'Create List' section with three radio button options: 'Manually', 'From Data Provider', and 'From Autofill' (which is selected).

Figure 278.4: Create a form with a text field and a select from list field. These are used to provide the input to the data provider and be autofilled by its output.

Once you're done, publish the form and try it out, by entering a valid region into the Region field, and observing that the options in the Select from List Field are filtered based on the Region. The `restcountries.eu` service has these regions you can use: Africa, Americas, Asia, Europe, Oceania, and Polar.

Autofill rules combine the power of data providers and form rules.

Condition

If

Actions

Do From Data Provider

Data Provider's Input: Assign the data provider parameter on the left column to a field in the select list on right column.

region

Data Provider's Output: Assign the data provider's parameters you want to a field and the information will autofill it automatically.

name

Figure 278.5: Create the autofill rule. Brag of your prowess.

ACTION: CALCULATE

Calculate rules let you automatically populate a numeric field by calculating its value based on other fields. Calculations are limited to numeric fields.

To follow the example below:

1. Create 5 Numeric fields called *Animal Weight 1 (US lb.)*, *Animal Weight 2 (US lb.)*, etc.
2. Create a Numeric field called *Total Food Required (US lb.)*.

Example: A 16-dog sled team can consume 2,000 US lb. of food during the Iditarod. This equates to about 0.25 lb. of food per lb. of animal, if the race lasts ten days. We'll use five numeric fields for animal weight instead of sixteen here, because it's tedious to create sixteen fields, even with the field duplication function. When the form user enters the weight of each animal the Total Food Required field should be calculated based on this simple formula:

Animal Weight 1, Animal Weight 2, ... = AW1, AW2, ...
Total Food Required = TFW

$(AW1 + AW2 + AW3 + AW4 + AW5) * 0.25 = TFW$

To configure a calculate rule:

1. Open the Rules tab of the Edit Form page and click the Add (+) button.
2. Define the rule:
 - If field *Animal Weight 1* is greater than 0, Calculate the sum of the Animal Weight fields, multiplied by 0.25.
 - Save the rule.

The calculation is defined using the embedded calculator. Use a mix of numeric field values, mathematical operators, and constants to define calculation rules.

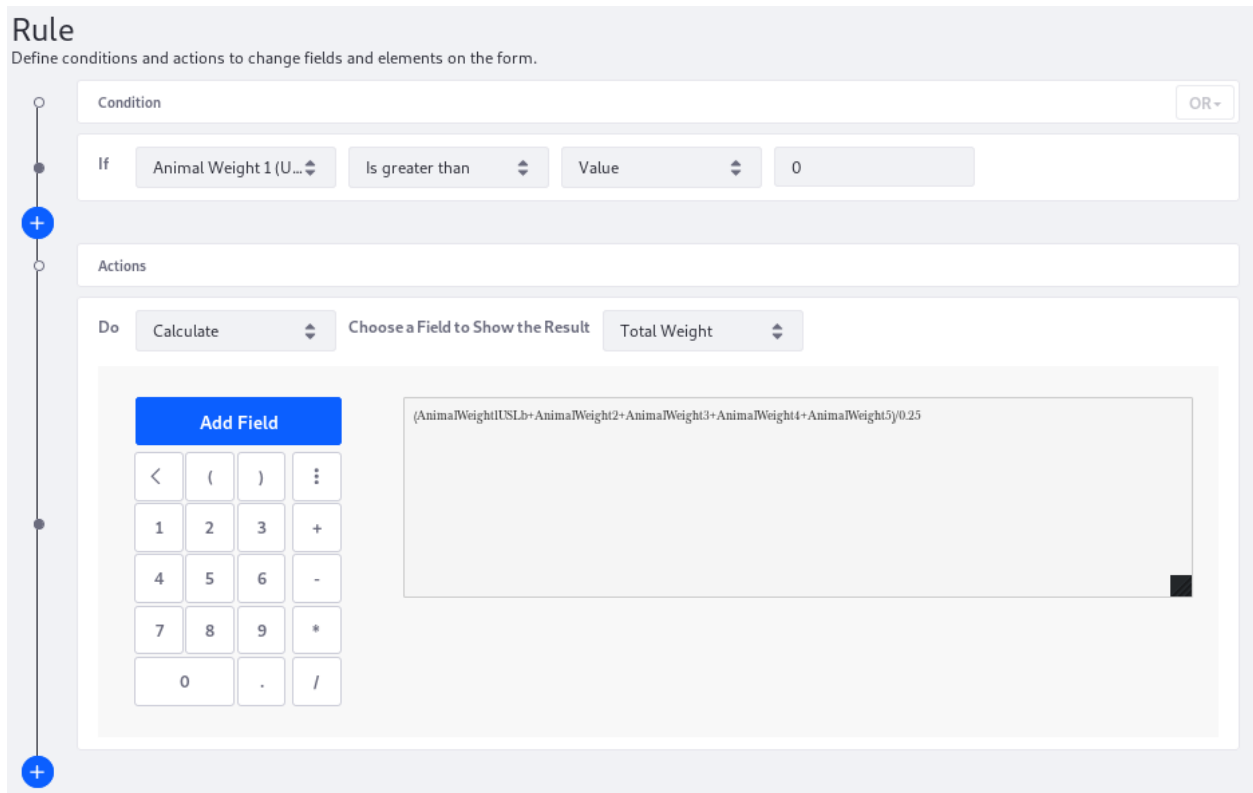


Figure 279.1: Build calculate actions with a handy calculator.




Figure 279.2: Once a rule is saved, it is displayed so that you can easily understand what it does.

FORM ELEMENT SETS

If you're here looking for information on reusable field sets, you're in the right place. We call them Element Sets in the Liferay Forms application because these sets include more than just fields: they include the layout and configuration of the fields as well. In the future, additional styling elements will be available here, too.

Element sets are more like composable Form fragments or reusable Form blocks.

Sometimes you might be able to create an entire form by composing existing Element Sets. Your colleagues might call you lazy, but we'd call you industrious.

280.1 Creating Element Sets

To create Element Sets, go to Site Menu → Content &Data → Forms. The Forms view is displayed by default. Click the *Element Sets* tab, and any existing Element Sets appear, just like existing Forms are displayed in the Forms view. Click the *Add* button (+).

Here's the thing. If you know how to create a Form, you already know how to create an Element Set. The process is identical. Drag and drop elements onto the form builder palette, configuring fields as you go.

When you're finished, click *Save*. Element Sets aren't publishable, so there's no button for that.

Once an Element Set is saved, it's instantly available for use, even in the same Element Set. That's right, you can use Element Sets to create Element Sets.

280.2 Using Element Sets

To use an Element Set in a Form:

1. Open the Form Builder.

Mailing Address

Add a short description for this element set.

Street Address *

City *

Postal Code *

Country *

Dynamically Loaded Data

Save Cancel

Figure 280.1: Creating Element Sets is just like creating Forms. You just can't publish them.

2. If the Add Elements sidebar isn't already displayed, open it by clicking the *Add* button (+).
3. The default view in the Add Elements sidebar is Elements. Instead click *Element Sets*.
4. Drag the Element Set onto the Form Builder, just like you would any single Form Element.

You probably already guessed that the process is identical for using Element Sets to build other Element Sets. That's all there is to it. There are just a couple more things to note:

- Once an Element Set is added to a Form, there's no connection with the root Element Set. You're free to move or configure the Fields and Elements however you want.
- Editing an Element Set doesn't retroactively affect the Forms where the Element Set was used.

Think ahead. Are there some common fields you'll commonly need to configure in your Forms? If so, create them as Element Sets once and save yourself repetitive work.

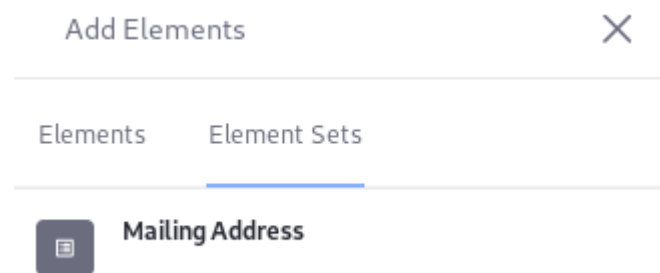


Figure 280.2: Add an Element Set the same way you add other Form Elements, like fields.

DATA PROVIDERS

Select from List fields can hold a lot of options. There are around 200 countries on Earth, for example. If you have unoccupied unpaid interns you could ask them to type each country into the Select from List field manually, or you could auto-populate your select fields using a REST web service. This saves you (or your interns) the trouble of typing all those options, and you can rely on someone else (hopefully a trustworthy expert) to keep the data updated.

When setting up a data provider, you're accessing a REST web service. Use the JSON web services registered in Liferay DXP, or any other REST web service you can access. To find a list of the registered JSON web services in Liferay DXP, navigate to <http://localhost:8080/api/jsonws> (assuming you're running a local server). Browse the available Liferay services. Many times, the services useful to you in the Forms application get a list of something. Find the `get-countries` JSON web service (there are two—use either one) and click on it, then click *Invoke*. The *Result* tab shows a list of countries using JSON syntax, like this:

```
[
  {
    "a2": "AF",
    "a3": "AFG",
    "countryId": "20",
    "idd": "093",
    "mvccVersion": "0",
    "name": "afghanistan",
    "nameCurrentValue": "Afghanistan",
    "number": "4"
  },
  ...
]
```

That's the record for the country Afghanistan. As you can see in the *URL Example* tab, the URL you entered into the data provider form is the same as the one generated for accessing the `get-countries` JSON web service. Find the URL for any registered JSON web service using this same procedure.

Note the field you want Users to select. With this service, it's most likely `nameCurrentValue`, because it contains the full, properly capitalized name of the country.

Enabling Access to Data on the Local Network: By default, you cannot configure data providers to use URLs on the local network. This is a good default for security in a production environment, but makes testing more difficult. To enable local network access from data providers, got to Control

Panel → Configuration → System Settings → Data Providers (under Content & Data), and enable *Access Local Network*. You'll need to configure this if you want to follow the basic example in the next section.

281.1 Adding a Basic Data Provider

To add a *Countries of the World* Data Provider for use in your Forms,

1. Go to the Forms application.
2. Click the *Data Providers* tab.
3. Click the Add button (+).
The REST Data Provider form loads.
4. Fill out the Name and Description fields.
Name: Countries of the World
5. Enter the URL and authentication tokens for the REST service. For the get-countries service:
URL: `http://localhost:8080/api/jsonws/country/get-countries/`
User Name: `adminuser@liferay.com`
Password: `adminuserpass`
6. In the Outputs fields, specify which field from the REST service populates your select list.
Label: Country Name
Path: `$.nameCurrentValue`
Type: List
7. Save the Data Provider.

What's that `$.` before `nameCurrentValue`? It's `JsonPath` syntax to navigate the JSON data structure and specify the path to the output. Learn more about `JsonPath` [here](#) and [here](#).

281.2 Using a Data Provider in a Select Field

Once the Data Provider is configured, use it to populate a Select from List field:

1. Go to the Form Builder (add a new form or edit an existing one)
2. Drag a Select from List field onto the form.

Name *

Countries of the World

Description

URL *

http://localhost:8080/api/jsonws/country/get-countries/

User Name ?

test@liferay.com

Password ?

....

Support filtering by keyword.

Cache data on the first request.

Timeout * ?

1000

INPUTS

| Label | Parameter | Type | <input type="checkbox"/> Required |
|----------------|----------------------|------------------|-----------------------------------|
| Enter a label. | Enter the parameter. | Choose an Option | |

OUTPUTS

| Label | Path | Type |
|--------------|---------------------|------|
| Country Name | \$.nameCurrentValue | List |


Figure 281.1: Set up a simple data provider in no time.

3. In the Create List section, choose *From Data Provider*.
4. Choose the Data Provider and its Output Parameter:
 Choose a Data Provider: Countries of the World
 Choose an Output Parameter: Country Name
5. Publish the form and test it.

Your Data Provider is now being used to populate a select field. However, this form should be submitted by Guest users, who don't currently have permission to see the list of results from the data provider. Arrgh! Now what?

Figure 281.2: Form users select an option from the list populated by the Data Provider.

281.3 Granting Data Provider Permissions

To configure the data provider's permissions, go to the Forms application (*Site Administration* → *Content & Data* → *Forms*). Open the Data Providers tab. For the data provider you want to configure, click the Actions button (), then *Permissions*.

Configure the permissions you need. If Guests are to fill out the form, they need the *View* permission, or else they won't be able to see the options provided by the data provider. Once you grant permissions, click *Save*.

281.4 Data Provider Configuration

The above instructions cover adding a basic Data Provider. Knowing more about each field in the Data Provider setup form opens up more possibilities.

URL The URL of an internal or external REST service endpoint. Consider the REST service at <https://restcountries.eu/>, which contains a REST API endpoint to find countries by region:
`https://restcountries.eu/rest/v2/region/{region}`

Data Provider URLs can take two parameter types: path parameters and query parameters.

Path parameters are part of the URL calling the REST web service, and are added using the pattern `https://service-url.com/service/{path_parameter_name}`:

The `restcountries.eu` service's region endpoint's path parameter is `{region}`. Path parameters are mandatory parts of the URL, so make sure you specify an Input (see below) with a *Parameter* field value matching the path parameter from the URL.

Query parameters are complementary parts of the URL that filter the output of the service call, following the pattern `?query_parameter=query_parameter_value`:

```
https://restcountries.eu/rest/v2/all?fields=capital
```

Unlike path parameters, query parameters are optional.

User Name and Password Credentials used to authenticate to the REST Web Service, if necessary.

Cache data on the first request. If the data is cached, a second load of the select list field is much faster, since a second call to the REST service provider is unnecessary.

Timeout The time (in ms) to allow the REST service call to process before aborting the request, if a response is not returned.

Inputs Configure path or query parameters from the REST service to filter the REST service's response. Specify the Label, Parameter, and Type (Text or Number), and choose whether the input is required to use the Data Provider. You can add multiple Inputs. To provide a way for users to specify the input value, use an *Autofill* Form Rule. A User enters input into one field, and their input is sent to the REST service. The REST service's response data is filtered by the input parameter.

Outputs The Parameter to display in Select from List or Text fields with autocomplete enabled. You can add multiple Outputs. Outputs can be filtered by inputs (see above) but can also be displayed without configuring input filtering. Specify the Label, Path, and Type (Text, Number, or List). The Path field is specified in JsonPath syntax, so it must always start with a `$`. The type of data returned by the Path must match the type you choose in the Type field. Using the `restcountries.eu` service, specify the name field as an Output by entering `enter $..name` in the Path field.

If you have a more complex JsonPath expression to construct (for example, you need the names of all countries with a population over 100 million—`$..[?(@.population>100000000)].name` with the `restcountries.eu` service), consider using an online JsonPath evaluator, like [this one](#) or [this one](#).

Hint: To display one value to the user, but persist another in the database, enter both into the Paths field, separated by a semicolon:

```
`$..name;$..numericCode`
```

If this is used with the `restcountries.eu` data provider, the name of the country is displayed to the User, while the numeric country code is stored in the database.

281.5 Troubleshooting Data Provider Errors

To uncover errors arising from Data Provider failures, configure log levels for these services:

Cache data on the first request.

Timeout ▲

INPUTS +

| Label | Parameter | Type | <input type="checkbox"/> Required |
|--|--|-------------------------------------|-----------------------------------|
| <input type="text" value="countryId"/> | <input type="text" value="countryId"/> | <input type="text" value="Number"/> | <input type="checkbox"/> |

OUTPUTS +

| Label | Path | Type |
|---|---|-----------------------------------|
| <input type="text" value="Country Name"/> | <input type="text" value="nameCurrentValue"/> | <input type="text" value="List"/> |
| <input type="text" value="Name-text"/> | <input type="text" value="nameCurrentValue"/> | <input type="text" value="Text"/> |

Figure 281.3: Set up Data Providers to display data retrieved from a REST service.

Category: `com.liferay.dynamic.data.mapping.data.provider.internal.DDMDataProviderInvokerImpl`
Level: WARN

Category: `com.liferay.dynamic.data.mapping.form.field.type.internal.DDMFormFieldOptionsFactoryImpl`
Level: DEBUG

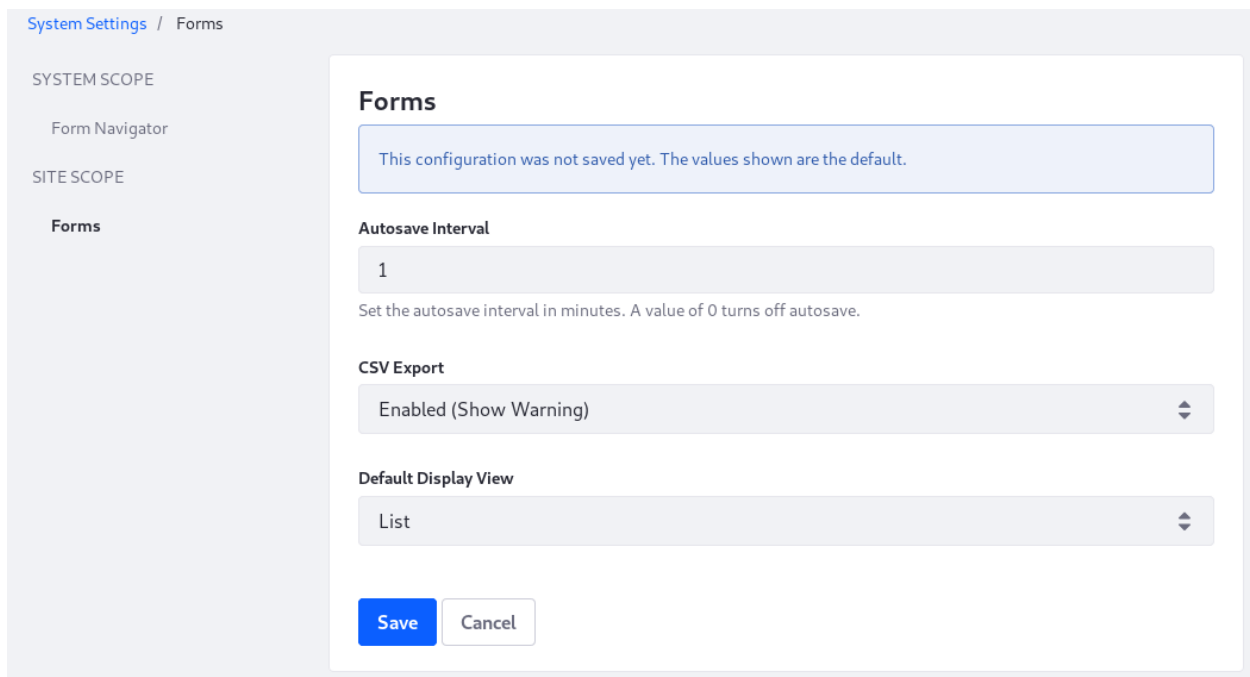
With Data Providers, the world’s (RESTful) data is at your disposal to use with the Forms application.

AUTO-SAVE

Losing progress on a partially created form is bad. Make sure to save your work frequently as you're creating forms. But if you forget to save your work, Liferay Forms has your back.

By default, a form is auto-saved every minute. You won't notice anything in the form builder while the back-end auto-saves the form. You can change the auto-save duration in *Control Panel* → *Configuration* → *System Settings* → *Forms* (in the Content and Data section). To disable auto-save, set the interval to 0.

For unpublished forms, an auto-save works just like a manual save. For a published form, however, auto-saved data isn't automatically propagated to the form. You must click the *Save* button in the form builder to publish the changes.



The screenshot shows the Liferay Forms configuration interface. The breadcrumb path is "System Settings / Forms". The left sidebar has "SYSTEM SCOPE" with "Form Navigator" and "SITE SCOPE" with "Forms" selected. The main content area is titled "Forms" and contains a warning message: "This configuration was not saved yet. The values shown are the default." Below this are three settings: "Autosave Interval" set to "1" with a subtext "Set the autosave interval in minutes. A value of 0 turns off autosave."; "CSV Export" set to "Enabled (Show Warning)"; and "Default Display View" set to "List". At the bottom are "Save" and "Cancel" buttons.

Figure 282.1: Configure the auto-save duration.

TRANSLATING FORMS

Forms can be translated to any locale that Liferay DXP supports. The form builder specifies a translation of the form's default language.

The form's default language and the available translations are set in the site's language configuration.

Follow these steps to create a form translation:

1. Go to *Site Administration* (your site's menu) → *Content & Data* → *Forms* and open the form to translate.
2. Click the + icon next to the current translation language and choose from the available languages.
3. Translate the form's title, field labels, field options, field placeholder text, and any other text visible to the user.
4. Save and publish the form.

To fill out a translated form in a translated language,

1. Access the form. If a signed-in user accesses the form and a translation is available in the user's language, the user sees the translated form by default.
2. To see the form in a different language, click the language icon and select a language.
3. Fill out the form as usual and click *Submit*.

Note: Translations work differently depending on how a User accesses a Form:

If accessed in the Form widget on a Liferay DXP page, the Form is displayed in the User's language automatically. If there's no translation available for the User's language, the default language of the Form is displayed.

If accessed via direct URL, the Form translation must be selected manually.




Figure 283.1: Add a translation for the form.




Figure 283.2: Translate as much of the form as possible into each language you expect users to need.

The image shows a web form with two language selection buttons at the top left: 'english-United States' (with a US flag icon) and 'pt-BR' (with a Brazilian flag icon). Below these is the question 'Qual a sua cor preferida?' in black text. Underneath the question is a label 'Cor' followed by a long, empty, light gray rectangular input field. At the bottom right of the form is a blue button with the text 'Submeter' in white.

Figure 283.3: Select the form's language.

AUTOCOMPLETING TEXT FIELDS

It's been scientifically proven that Internet users are lazy (not you, of course—other Internet users). For example, some users may not fill out your form if you make them type the entire title of their country in an employment application. This is especially true if they're filling out the form on their mobile devices. Make users' lives easier by configuring autocomplete on a form's text fields.

Why not just use a select field with a data provider to guide user input? Sometimes a data provider can't encompass all possible field entries. For example, if your data provider doesn't include mythical countries founded on old sea platforms, users the Principality of Sealand can't enter anything into the select field. Instead use a text field with autocomplete so users can begin typing their country's name and then select it from a list when it appears. Autocomplete combines a text field (accepting any response that meets your validation criteria) and common choices to select from. It's a win-win situation.

284.1 Configuring Autocomplete

Before configuring autocomplete for your text fields, create a form and add a text field. If you want the autocomplete options to be populated by a REST data provider, configure one before creating your form. Now you're ready to configure autocomplete for the field:

1. In the field configuration sidebar, click the *Properties* tab.
2. Click the *Autocomplete* switcher so it's enabled.
3. Select a data provider or create one manually. You can set up a data provider from a REST service, or manually enter the options users should see when they start typing in the text field.
4. Save and Publish the form.

Once users begin entering text into the field, a selection list of options appears. As they enter additional text, the list is refined to include only options that contain the currently entered text.

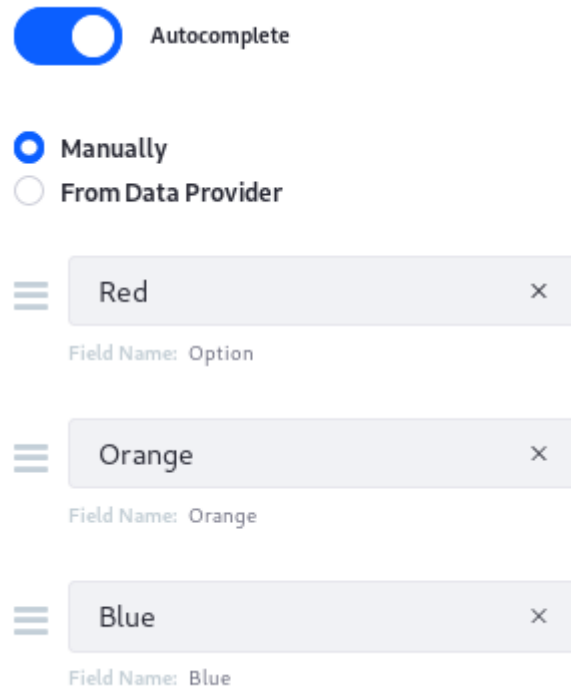


Figure 284.1: You can configure a manual data provider to specify the options users can select from.

For example, the imaginary users from Sealand (all two of them) begin reluctantly typing their country of origin by entering an *S*. They're delighted to see a selection list with a bunch of countries containing the letter *S* appear for their selection convenience. If they continue typing and enter *e*, the list is refined to options that have *se* in their name (for example Serbia and Senegal). If they continue typing and enter *a*, they'll now only see one option, Sealand, if it's in your data provider. Selecting it from the list after typing the first three letters is much easier than typing the remaining letters.

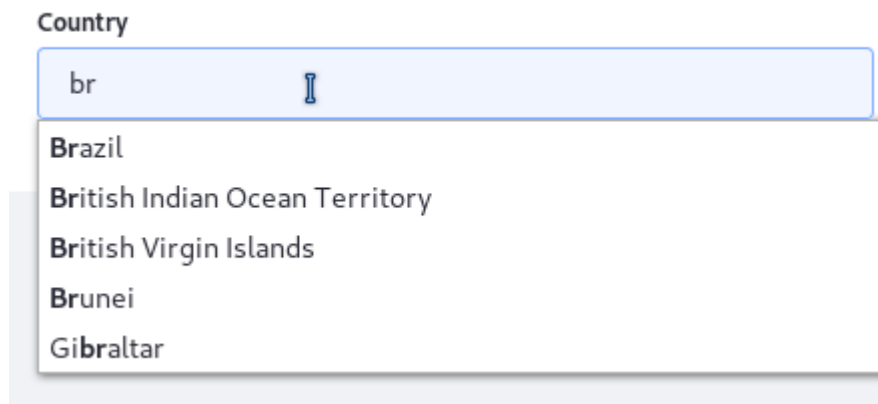


Figure 284.2: When typing in a field with autocomplete, users are presented a list of selections from the configured data provider. The displayed results are filtered to include only selections containing the text entered by the user.

What will the Forms team think of next? Configuring telepathic connections to the Forms

application would be nice. Then users could just think their form field entries into existence. Stay tuned.

FORM SUCCESS PAGES

After users submit one of your whiz-bang forms, what's next? How will they know they're done and can close the browser window or tab? What if they think their submission didn't go through and wonder if they need to fill out the whole form again? By default, submitting a form displays the default success message and returns users to the form's now empty first page. Don't leave users feeling equally empty. Instead, configure a *Success Page*. A Success Page is a terminal page showing users they've finished filling out the form and their submission was successfully received. A Success Page can even urge users to close the browser window or tab.

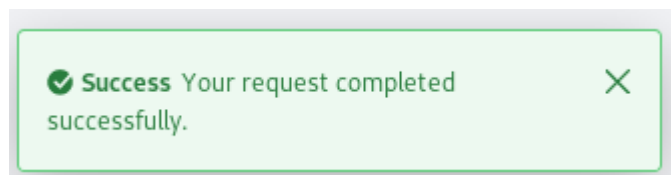



Figure 285.1: The default success message alerts users when their request completes successfully.

A Success Page is simple. It has a title in bold text and a description beneath the title. A common alternative to using a Success Page is to redirect users to a different page in your Site. What should you put in a Success Page? Whatever you want. If you can't think of anything important or creative to say, use the default message:

To configure a Success Page,

1. Add a form in *Site Administration* (your site's menu) → *Content & Data* → *Forms*.
2. Click on the form page's *Actions* button () and choose *Add Success Page*.

Once the Success Page is added to your form, fill in the *Title* and *Content* fields however you please. When the form is saved and published, the Success Page is live for your form users.

Note: You can't preview the Success Page. Success Pages can only be viewed once a form is submitted, and the *Submit* button isn't available in the form preview. The *Preview Form* link in the form builder only lets you preview the form's regular pages (use the *Next* button to navigate through the form).

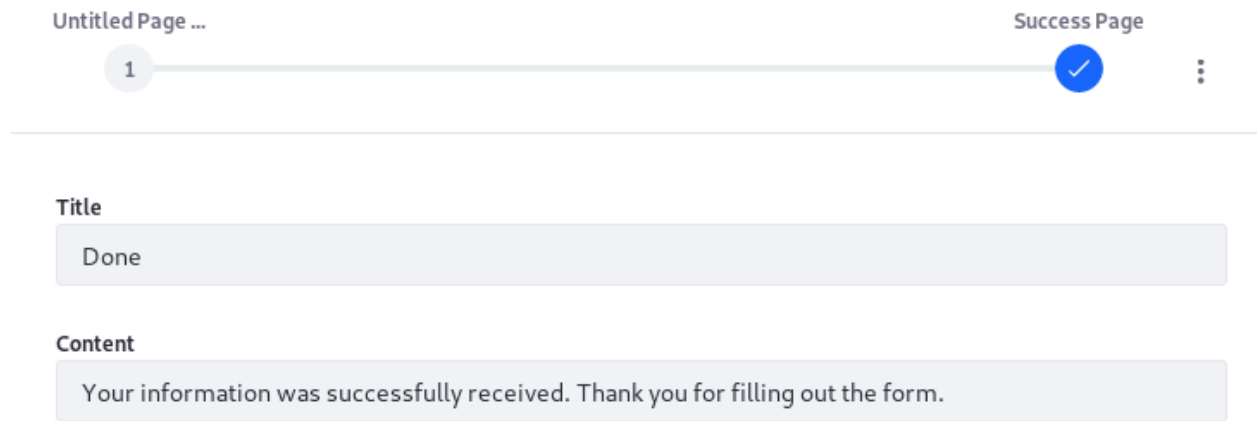


Figure 285.2: There's a default Success Page message if you can't think of anything else to say.



Figure 285.3: Add a Success Page using the edit menu for the form page.

To see what your Success Page looks like, submit a test entry of the form and then delete it if needed. For more information, see the documentation on viewing and managing form entries.

WORKFLOW AND FORMS

The workflow engine is for sending a submitted asset through a workflow process before it's published. Most assets are configured to use workflow at the instance or Site level.

| Asset Type | Workflow Assigned | |
|------------------------|-------------------|------|
| Blogs Entry | No Workflow | Edit |
| Calendar Event | No Workflow | Edit |
| Comment | No Workflow | Edit |
| Knowledge Base Article | No Workflow | Edit |
| Message Boards Message | No Workflow | Edit |
| Page Revision | No Workflow | Edit |
| User | No Workflow | Edit |
| Web Content Article | No Workflow | Edit |
| Wiki Page | No Workflow | Edit |

20 Entries ▾ Showing 1 to 9 of 9 entries. < 1 >

Figure 286.1: Workflow is enabled in the Control Panel or in Site Administration for most Liferay DXP assets.

Forms are different, so they don't appear in the above image. There are so many use cases for forms, and there could be so many per site, that a site- or instance-scoped workflow configuration won't serve your needs well. Instead, configure workflow for *each form* separately.

286.1 Enabling Workflow in a Form

To enable workflow in a form,

1. Open the form's editor by opening the Menu, selecting your Site, navigating to *Content & Data* → *Forms*, and clicking on the form you want.
2. Click the Options button (ⓘ) and choose *Settings*.
3. The Settings window has a *Select a Workflow* drop-down. Find the workflow you want, select it, and then click *Done*.

The screenshot shows the 'Form Settings' dialog box with a close button (X) in the top right corner. It has two tabs: 'Form Options' (selected) and 'Email Notifications'. Under 'Form Options', there are three toggle switches: 'Require User Authentication' (off), 'Require CAPTCHA' (off), and 'Save Answers Automatically' (on). Below these are three input fields: 'Redirect URL on Success' with a placeholder 'Enter a valid URL.', 'Select a Storage Type' with a dropdown menu showing 'json', and 'Select a Workflow' with a dropdown menu showing 'Single Approver'. At the bottom right, there are two buttons: 'Cancel' and 'Done'.


Figure 286.2: Enable workflow for each form in its Settings window.

286.2 Testing the Workflow



Test the workflow process:

1. Add the form to a page.
2. Click *Submit for Publication* to submit the form entry.

Next go find the form entry in the Forms application:

1. Go back to the Forms application in the Menu in your Site's *Content & Data* section.
2. Click the Form's *Actions* button () and select *View Entries*.
The entry is currently marked *Pending*.

Now approve the form record:

1. Navigate to *My Workflow Tasks* in the User Personal Menu.
2. Click the *Assigned to My Roles* tab.
3. Click on the form entry.
4. Click the *Actions* button () and choose *Assign to Me*.
5. Click *Done*.
6. Click the *Actions* button () again, then click *Approve*.
7. Click *Done* again.
8. Navigate back to the *View Entries* screen for the form, and now the entry is marked as *Approved*.




| Street Address | City | Postal code | Country | Country Text Field | Status | Modified Date | Author |
|--|----------|-------------|---------------|--------------------|----------|----------------|-----------|
|  | | | | Sea | DRAFT | 16 Minutes Ago | Test Test |
|  51 Alphonse Rd | Brockton | 2402 | United States | United States | APPROVED | 9 Seconds Ago | Test Test |

Figure 286.3: Each entry's status is visible in the Forms application's Form Entries screen.

DUPLICATING FORMS AND FORM FIELDS

Repetitive tasks are error prone. Instead of duplicating effort, duplicate forms and form fields.
To duplicate a form,

1. Go to *Site Administration* (your Site's menu) → *Content & Data* → *Forms*.
2. Click the *Actions* button () for the form to duplicate.
3. Click *Duplicate*.

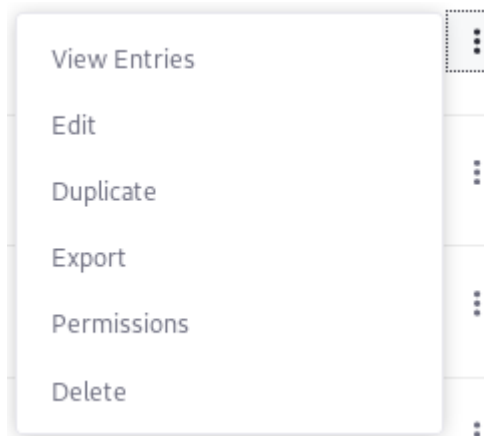



Figure 287.1: The Duplicate option works the same for forms and form fields.

The form is duplicated and automatically named *Copy of [Original Form Name]*. Once duplicated, you can edit the form however you want. When you duplicate a form, all configurations and form rules are duplicated as well.

To duplicate a form field,

1. Go to *Site Administration* (your Site's menu) → *Content & Data* → *Forms* and open or add a form.
2. In the Builder view, hover over the form field to duplicate and click the *Copy* icon ().




Figure 287.2: You can duplicate form fields.

The field is duplicated and labeled *Copy of [Original Field Label]*. All the form's properties, including its data provider configurations, are copied as well.

FORM PAGES

Are users more likely to abandon long forms with lots of scrolling? Are they more likely to see a multi-page form and abandon it without a second look, assuming that it gets longer and more tedious with every passing page? Such usability questions are worth thinking about. If you decide multiple pages are appropriate for your form, Liferay Forms supports two pagination styles: default, and alternate.

Multi-Page Form

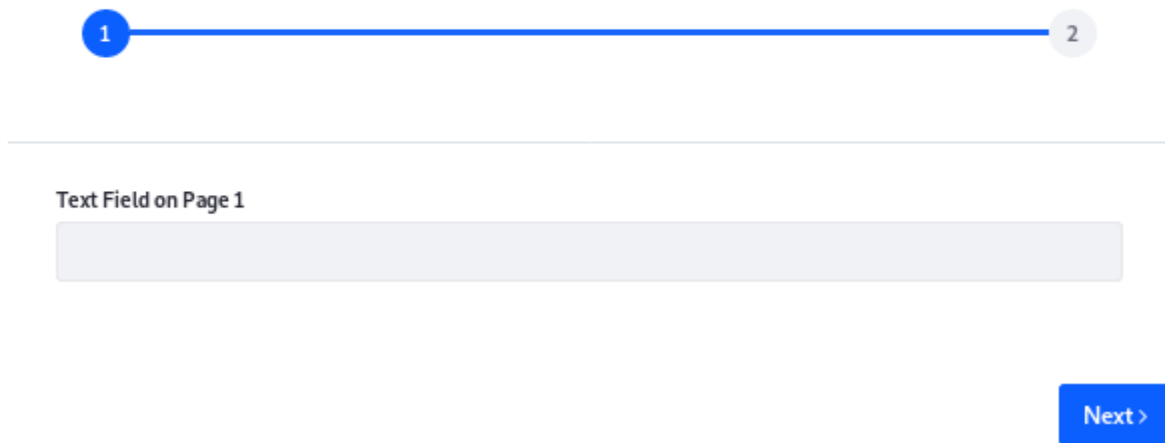


Figure 288.1: The default pagination style.

To add a form page,

1. Go to a form's builder view.
2. Click the *Actions* button () at the top-right corner of the form, then click *Add New Page*.

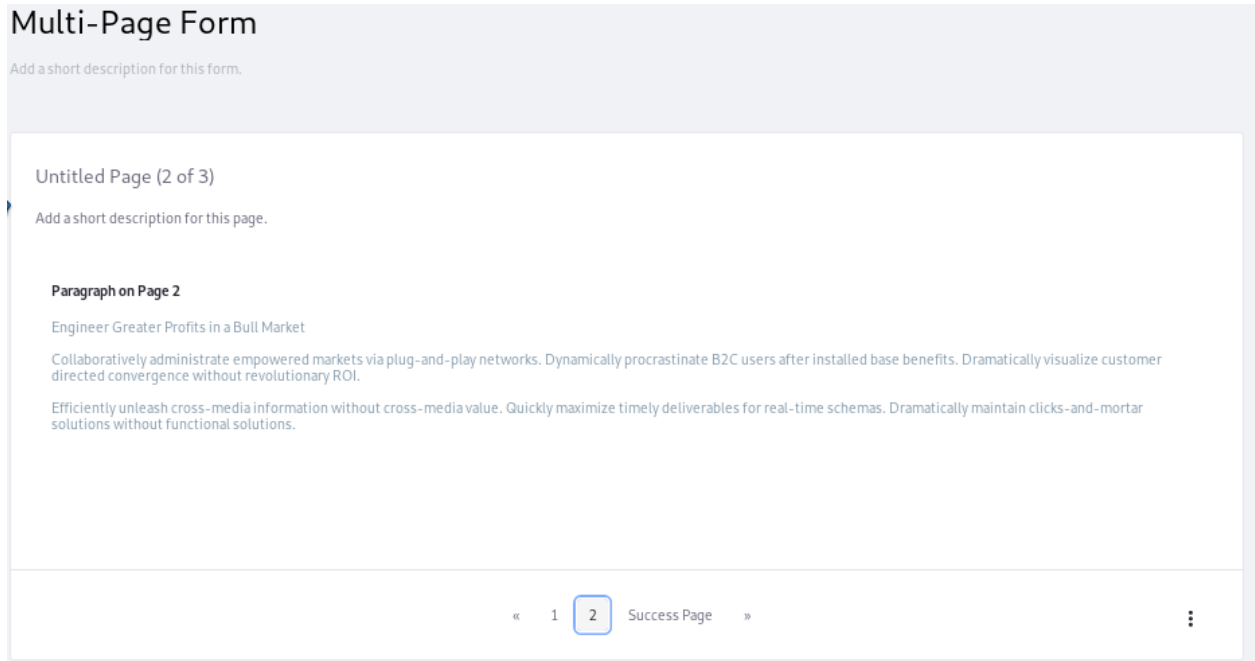


Figure 288.2: The alternate pagination style as seen in the Form Builder.

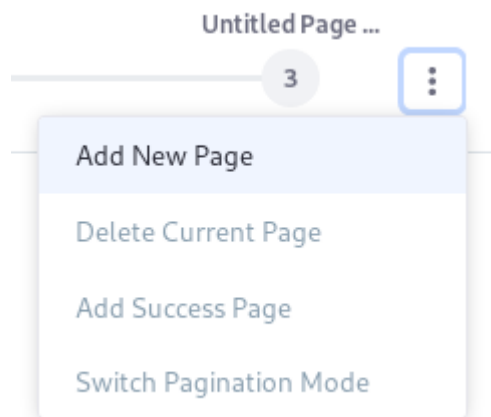


Figure 288.3: You can add new pages or reset the current page from the Page Actions menu.

You can also delete form pages, add success pages, and switch pagination modes.

HELP TEXT, PLACEHOLDER TEXT, AND PREDEFINED VALUES

Form fields can have help text, placeholder text, and predefined values.

- **Help Text:** Text that appears as a sub-heading to the field label, but doesn't appear in the field entry area. Enter help text in the Basic tab of the field's sidebar menu.
- **Placeholder Text:** Text in the field entry area that isn't submitted if the field is left untouched by the user.
- **Predefined Value:** Text in the field entry area that is submitted if the field is left untouched by the user.

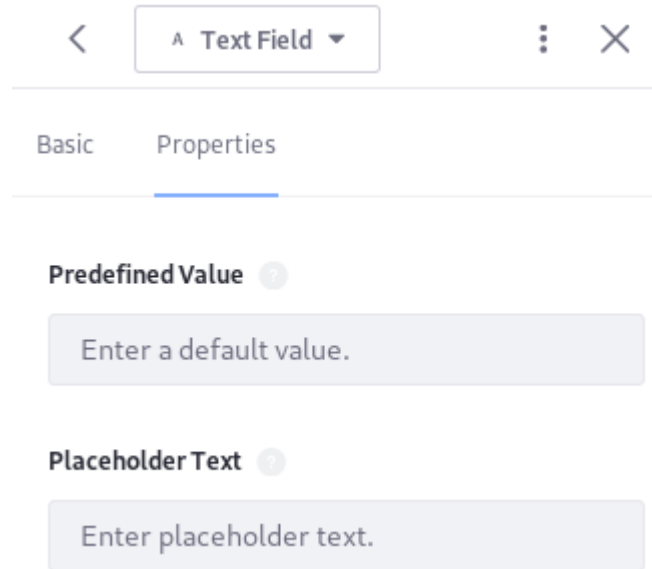
All form field types can have help text, and all form field types that accept user input can have predefined values. Only text and numeric fields can have placeholder text.

To enter placeholder text or predefined values,

1. Open a field's sidebar menu.
2. Open the *Properties* tab.

For example, many forms start with a *Full Name* field. You can use help text and/or placeholder text to inform users that you need their full names, regardless of length. Alternatively, if you're asking users to specify how many sandwiches they eat for lunch, a predefined value of *1* probably makes sense.

Remember, placeholder values aren't submitted if the field is left blank, so you don't have to worry about getting a bunch of submissions from *Maximillian Aurelius Piroux the 11th*.



The image shows a configuration interface for a text field. At the top, there is a navigation bar with a back arrow, a dropdown menu labeled 'A Text Field', and a close button. Below this, there are two tabs: 'Basic' and 'Properties', with 'Properties' being the active tab. Under the 'Properties' tab, there are two sections: 'Predefined Value' and 'Placeholder Text'. Each section has a text input field with a question mark icon to its right. The 'Predefined Value' field contains the text 'Enter a default value.' and the 'Placeholder Text' field contains the text 'Enter placeholder text.'

Figure 289.1: Predefined values and placeholder text are entered in the Properties tab.

Placeholders and Predefined Values



The image shows a form with two input fields. The first field is labeled 'Full Name' and contains the text 'Maximillian Aurelius Piroux the 11th'. Below the field is the text 'Really, please enter the whole thing.' The second field is labeled 'How Many Sandwiches Do You Require for Lunch?' and contains the number '1'. Below the field is the text 'Do you reall need more than 1?'. At the bottom right of the form is a blue 'Submit' button.

Figure 289.2: The Full Name field here uses help text and placeholder text, while the sandwiches field uses a predefined value.

VALIDATING TEXT AND NUMERIC FIELDS

Validation ensures that only certain values are entered in a field. Validation functionality is available for text and numeric fields.

To enable validation,

1. Add a Text or Numeric field to a form in the Builder view.
2. Open the field's *Properties* tab.
3. Turn on the *Validation* toggle to enable validation and open its configuration options.




Figure 290.1: Validate data to ensure you're collecting only useful information.

290.1 Validating Text Fields

Validation for text fields contains several options. You must first choose a list of available conditions to check:

- If Input *Contains*
- If Input *Does Not Contain*
- If Input *Is not URL*
- If Input *Is not Email*
- If Input *Does not Match*

If the condition isn't met, an error message is displayed to the user. Where you go from there depends on which condition you used.

290.2 If Input Contains/Does Not Contain

When you validate text data to check if it contains a certain value, there are two additional steps to take after selecting the condition:

1. Enter the text to check for.
2. Enter an error message so users understand why their submission failed.

Validation

If Input

Does not contain

The Value

Liferay

Show Error Message

Your entry must contain "Liferay"

Figure 290.2: If *Liferay* isn't part of the field's value, an error message is displayed.

290.3 If Input Is not URL/Email

Checking for properly formatted URLs and emails is easy. Just choose the condition from the *If Input* drop-down and enter the error message.

Valid URLs begin with `http://` or `https://`. Valid emails must contain `@`.

Validation

Email Address

Please enter a valid email address

Submit

Figure 290.3: Use text field validation to make sure users enter a valid email address or URL.

290.4 If Input Does Not Match

The *Does Not Match* condition is used for entering regular expressions to create custom validation criteria. For example, use this regular expression to ensure that ten consecutive numeric digits are entered in a phone number field:

```
^[0-9]{10}$
```

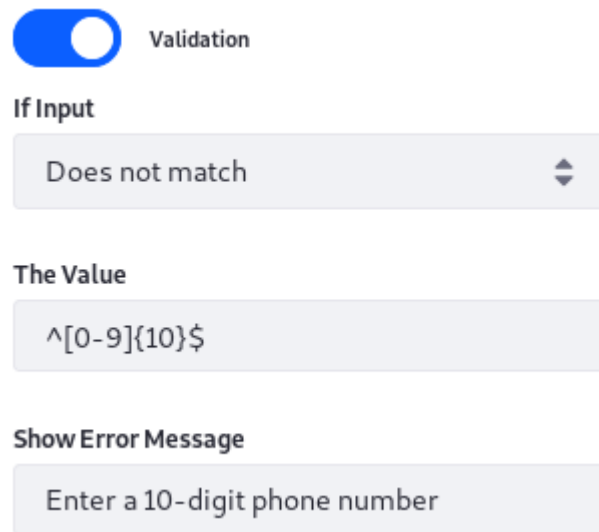
If you use regular expression validation, provide some explanatory text (e.g., help text, placeholder text, and a clear error message) to guide form users in entering the proper data.

290.5 Validating Numeric Fields

Numeric field validation is similar to text field validation, but the conditions compare the value of the number entered to some other value.

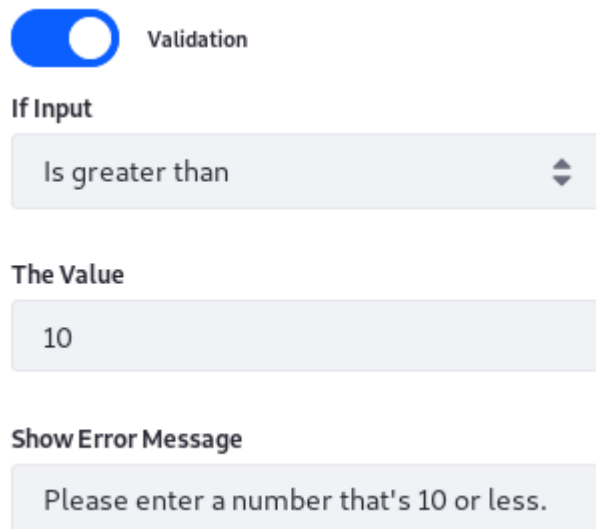
Available conditions to check include

- Is greater than or equal to
- Is greater than
- Is not equal to
- Is less than or equal to



A validation configuration interface with a blue toggle switch labeled "Validation" that is turned on. Below it, the "If Input" dropdown menu is set to "Does not match". The "The Value" text input field contains the regular expression `^[0-9]{10}$`. The "Show Error Message" text input field contains the message "Enter a 10-digit phone number".

Figure 290.4: Regular expression text validation opens up countless possibilities.



A validation configuration interface with a blue toggle switch labeled "Validation" that is turned on. Below it, the "If Input" dropdown menu is set to "Is greater than". The "The Value" text input field contains the number "10". The "Show Error Message" text input field contains the message "Please enter a number that's 10 or less."

Figure 290.5: Numeric conditions constrain user-entered numeric data.

- Is less than

For example, to make sure users don't enter a number over 10, enable validation and use *Is greater than* with a value of 10. Use the message *Please enter 10 or less*.

How many sandwiches do you need for lunch?

Please enter a number that's 10 or less.

Figure 290.6: Make sure user-entered numeric data is within reasonable bounds. Nobody needs 11 sandwiches for lunch.

Note that numeric fields are text fields validated to allow only numeric data entry. That's why they're in the Customized Elements section of the form fields list. In addition, the property *My numeric type is* (can be Integer or Decimal) on the Basic tab of a numeric form is another form of validation.

ENABLING CAPTCHA ON FORM SUBMISSIONS

CAPTCHA prevents a bot from submitting forms. It's often used in login apps, but you can also use it in the Forms app.

To enable CAPTCHA, click the form's *Options* button (ⓘ) and select *Settings*. Enable the *Require CAPTCHA* setting, click *Done*, and save the form. That's all there is to it!

Form Settings

Form Options

Email Notifications



Require user authentication.





Require CAPTCHA

Figure 291.1: You can enable CAPTCHA for your form in the Form Settings window.

capital

country

Choose an Option 




Text Verification *

[Submit for Publication](#)

Figure 291.2: Once you enable CAPTCHA, your form has protection against bot submissions.

FORM NOTIFICATIONS

You can configure the Forms app to send a notification email each time a form entry is submitted.

1. Open the form's *Form Settings* section by clicking the *Options* button () and selecting *Settings*.
2. Click the *Email Notifications* tab, enable the option *Send an Email Notification for Each Entry*, and fill out these fields:

From Name: The sender's name. This could be the Site name, the form name, or anything else informative to the recipient.

From Address: The sender's email address. You can use something like `noreply@example.com`, so that recipients don't try to reply.

To Address: The recipient's email address (e.g., `test@example.com`).

Subject: The email's subject. An informative subject line tells the recipient what happened. For example, `An application for employment was submitted in The Lunar Resort*`.

Note that if you enabled workflow for the form and it already sends a notification, you might not need to configure the Forms app to generate a notification.

Form Settings ✕

Form Options **Email Notifications**

Send an email notification for each entry.

From Name

From Address

To Address

Subject


[Cancel](#) [Done](#)

Figure 292.1: Configure email notifications each time a form entry is submitted.

REDIRECTING USERS

When users submit a form, you can present them with another page indicating success or some other information related to their submission. Sometimes all you need is a success page, but other times you might want to send users to a specific URL.

Whatever your use case is, follow these steps to set up a redirect URL:

1. Open the form's *Form Settings* section by clicking the *Options* button () and select *Settings*.
2. Enter the redirect URL in the *Redirect URL* field.

That's it! Now when users submit the form, they're not left wondering what to do next.

Form Settings ✕

Form Options Email Notifications

Require user authentication.

Require CAPTCHA


Save answers automatically.

Redirect URL on Success

`https://mysite.com/web/guest/home`

Figure 293.1: Redirect users after they submit a form.

FORM PERMISSIONS

To access a form's permissions, first navigate to the Forms app in *Site Administration* (your site's menu) → *Content & Data* → *Forms*. Then click the form's *Actions* button (), and select *Permissions*.

By default, you can grant these permissions for a form:

Delete: Delete the form.

Permissions: Access and configure the form's permissions.

Update: Update form entries.

Add Form Instance Record: Submit form entries.

View: View the form.

Note that guest users can view and fill out forms by default. The *Guest Role* has *View* and *Add Form Instance Record* permissions.

Note: By default, all users inherit the *Guest Role*'s permissions. The *Guest Role* represents unauthenticated visitors of your Site. If you want to let *Guest* users submit forms (the default setting), it makes sense that authenticated users can also. To disable automatic inheritance of the *Guest Role*'s permissions, set this property in your `portal-ext.properties` file:

```
`permissions.check.guest.enabled=false`
```

| Role | Delete | Permissions | Update | Add Form Instance Record | View |
|-------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| Guest [?] | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Owner | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Portal Content Reviewer | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Power User | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Site Content Reviewer | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Site Member | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input checked="" type="checkbox"/> |
| User | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

20 Entries ▾ Showing 1 to 7 of 7 entries. ◀ 1 ▶

Figure 294.1: You can configure a form’s permissions.

STYLING FORM PAGES

Let's face it: nobody likes an ugly, confusing form. Styling your form pages lets you make your forms user friendly. There are two features for styling your forms:

1. Create rows and columns for form fields.
2. Move fields from one location to another.

Sometimes it doesn't make sense to use the default single-column, vertically-oriented form. For example, a form with many fields can save space by putting them in different columns. You can also use a mixed approach, with each row broken into a different number of columns. The following screenshots show examples of these layouts.




Figure 295.1: This is the default single-column, vertically-oriented form.

By default, dragging a field onto a form page adds a field that occupies an entire row. Follow these steps to resize the field to make room for more fields in the row (columns):

1. Hover over the field to reveal its borders:
2. Drag the right or left edge of the field to resize it.
3. Add a new field to the remaining space in the row.

You can also move fields. To do so, follow these steps:

1. Hover over the field and the cursor becomes a hand.




Figure 295.2: Putting form fields in multiple columns can give you more space.




Figure 295.3: The first row is in two columns and the second row is in three columns.




Figure 295.4: Form field borders.

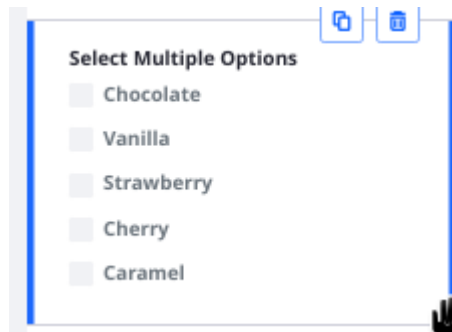


Figure 295.5: After resizing, the field is smaller.

Figure 295.6: There are now two fields in the row.

2. Drag the field to an open location in the builder. Available locations are highlighted in blue and outlined with a dotted line. A field moved to a new row fills that entire row. A field moved to an existing row fills that row's remaining space.

Styling Form

Add a short description for this form.

Untitled Page (1 of 1)
Add a short description for this page.

A Paragraph Field
By default, a field dragged onto a page in the Form Builder occupies an entire row.

Half-Row Text Field
Enter Some Text

Select from List
Choose an Option

Single Selection
 Choose This One Or This One
 Really, It Doesn't Matter

Select Multiple Options
 Chocolate
 Vanilla
 Strawberry
 Cherry
 Caramel

Unpublish Form Save Form Preview Form

Figure 295.7: You can also move fields on form pages.

DYNAMIC DATA LISTS

Dynamic data lists display forms created from field sets called *data definitions*. Data definitions consist of a form's field types (e.g., text, boolean, date, radio buttons, selector menus, etc.) and those fields' labels and settings. Data definitions effectively serve as data models for a dynamic data list. For example, you could create a data definition with two text fields: one for a user's name, and one for their comments. You could then display a form that gathers user feedback via a dynamic data list that uses that data definition.

To summarize:

- **Data Definitions:** Define a form's fields.
- **Dynamic Data Lists:** Display a form based on a data definition.

You can create one or multiple dynamic data lists from a single data definition. The user data entered for each dynamic data list is kept separate, even if the data definition is shared. For instance, you could use the example data definition above to create several dynamic data lists, and then place them anywhere you need to get feedback from users. Because each dynamic data list's form data is separate, you don't need to worry about trying to figure out which dynamic data list the user comment came from.

Dynamic data lists are flexible. You don't have to restrict dynamic data lists to simple input. You could create something as complex as an entire data entry system for real estate listings, or any other simple list-based application you can dream up.

You create data definitions and dynamic data lists in from the Site Menu's Content → Dynamic Data Lists application. Creating data definitions and lists doesn't require any coding. However, additional formatting can be added with FreeMarker templates.

These tutorials show you how to create and use data definitions and dynamic data lists:

- Creating data definitions.
- Creating dynamic data lists.
- Creating form and display templates.

296.1 System Configuration

There are two Dynamic Data Lists entries in System Settings. The Dynamic Data Lists Service entry contains one setting:

Add Default Structures This is enabled by default and pre-loads several embedded data definitions to base data lists on. Once loaded on portal startup, these definitions must be deleted manually from the Site Menu → Dynamic Data Lists application. This setting applies to the first start of a virtual instance.

The Dynamic Data Lists entry contains three settings:

Changeable Default Language If enabled, the default language of a data definition becomes changeable.



CSV Export Choose whether DDL records can be exported in CSV format with or without a warning, or disable this option. Here's what the warning says:

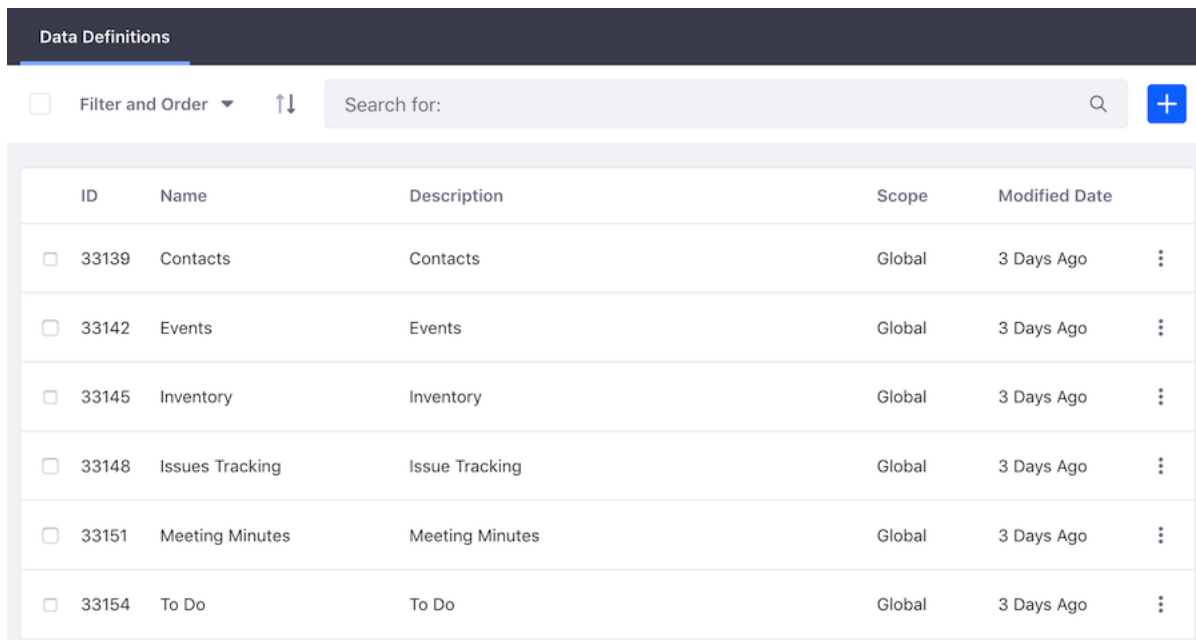
Warning: This CSV file contains user supplied inputs. Opening a CSV file in a spreadsheet program may be dangerous.

Default Display View Choose whether to use a table based default view or a list based default view.

CREATING DATA DEFINITIONS

Follow these steps to create a data definition:

1. Open the Menu () and expand your site's menu (the Site Administration menu). Then select *Content & Data* → *Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table lists any existing dynamic data lists.
2. Click the *Options* button at the top-right () and select *Manage Data Definitions*. The Data Definitions screen appears. A table lists any existing data definitions. Several definitions are embedded for common use cases like contacts, events, inventory, and more.



The screenshot shows the 'Data Definitions' screen. At the top, there is a dark header with the title 'Data Definitions'. Below the header, there is a toolbar with a 'Filter and Order' dropdown, a search bar labeled 'Search for:', and a blue plus button. The main content is a table with the following columns: ID, Name, Description, Scope, and Modified Date. The table contains six rows of data, each with a checkbox in the ID column and a vertical ellipsis in the Modified Date column.

| ID | Name | Description | Scope | Modified Date |
|--------------------------------|-----------------|-----------------|--------|---------------|
| <input type="checkbox"/> 33139 | Contacts | Contacts | Global | 3 Days Ago |
| <input type="checkbox"/> 33142 | Events | Events | Global | 3 Days Ago |
| <input type="checkbox"/> 33145 | Inventory | Inventory | Global | 3 Days Ago |
| <input type="checkbox"/> 33148 | Issues Tracking | Issue Tracking | Global | 3 Days Ago |
| <input type="checkbox"/> 33151 | Meeting Minutes | Meeting Minutes | Global | 3 Days Ago |
| <input type="checkbox"/> 33154 | To Do | To Do | Global | 3 Days Ago |

Figure 297.1: The Data Definitions screen.

3. Click the *Add* button (+) to begin creating a new data definition. This opens the *New Data Definition* form.
4. Give your data definition a name. Note that the definition's name appears for any users filling out a dynamic data list that uses the definition. Then expand the *Details* section of the form and give it a description.
5. The *Details* section of the form also contains the field *Parent Data Definition*. This optional field lets you select an existing data definition (the parent) to form the basis of the new one (the child). The child definition inherits the parent's fields and settings, which you can then customize. When you create a dynamic data list from a child definition, it includes the fields of the parent and child definitions. This lets you use a common definition (the parent) as the basis of a specialized definition (the child). For example, if you were planning a rock climbing trip, you could use the default *Events* definition as the parent of a *Rock Climbing Trip* definition that contains fields unique to rock climbing (e.g., climbing equipment availability, altitude, etc.).

To choose a parent definition, click the *Select* button below the *Parent Data Definition* field and then select an existing definition in the dialog that appears.

The screenshot shows a form for creating a new data definition. At the top, there is a 'Name' field with a red asterisk, containing the text 'Lunar Excursions'. Below this is a 'Details' section, which is expanded to show a 'Description' field containing 'Data definition for planning outings on the moon'. Underneath the description is a 'Parent Data Definition' field containing 'Events'. Below the parent definition field are two buttons: 'Select' and 'Remove'. The form has a light green header and a light blue footer.

Figure 297.2: After naming your data definition, expand the *Details* section of the form and give your definition a description and parent definition, if desired.

6. Add the data definition's fields in the data definition designer, below the form's *Details* section. The designer's default *View* tab lets you create the definition in a WYSIWYG editor. You can click the *Source* tab to work with the definition's underlying JSON, but it's much easier to stick with the WYSIWYG editor.

In the *View* tab select the *Fields* tab. Icons representing the field types are listed on one side and the data definition's canvas is on the other side. To add a field type to the definition, select its icon, drag, and drop it onto the canvas. By dragging a field onto a field that's already on the canvas, you can nest the new field in the existing field. When you mouse over a field on the canvas, the field action icons appear. Clicking the + icon creates a duplicate of the current field and adds it below the current field. Clicking the trash can deletes the field.

The following fields are available:

- **Boolean:** A check box.

- **Color:** Specifies a color.
- **Date:** Enter a date. A valid date format is required for the date field, but you don't have to enter a date manually. When you select the date field a mini-calendar pops up which you can use to select a date.
- **Decimal:** Enter a decimal number. The value is persisted as a double.
- **Documents and Media:** Select a file from a Documents and Media library.
- **Geolocation:** Associate a location with the User's form entry.
- **HTML:** An area that uses a WYSIWYG editor to write and display HTML content.
- **Integer:** Enter an integer. The value is persisted as an int.
- **Link to Page:** Link to another page in the same site.
- **Number:** Enter a decimal number or an integer. The value is persisted either as a double or an int, depending on the input's type.
- **Radio:** Displays several clickable options. The default number of options is three but this is customizable. Only one option can be selected at a time.
- **Select:** This is just like the radio field except that the options are hidden and must be accessed from a drop-down menu.
- **Text:** Enter a single line of text.
- **Text Box:** This is just like the text field except you can enter multiple lines of text or separate paragraphs.
- **Web Content:** Select web content.

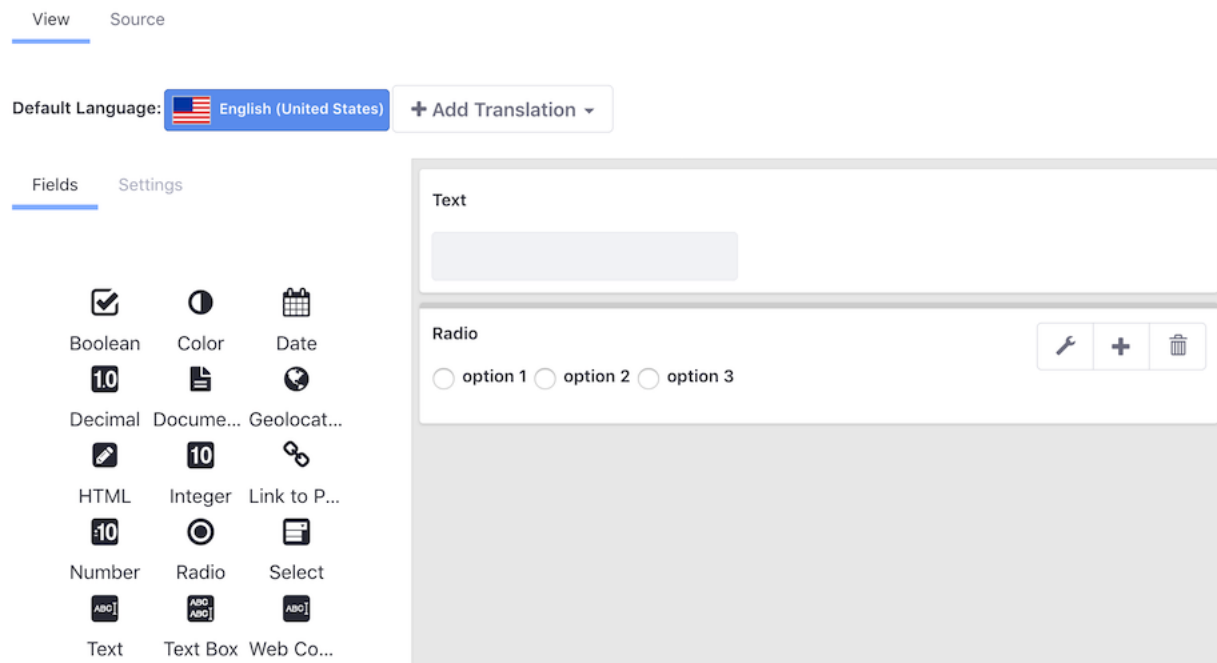


Figure 297.3: Use the data definition designer to add fields to the data definition.

7. Edit field labels to reflect their intended data. A text field's default label is *Text*. To use the text field as a title, then you should change the field's label to *Title*. First select the field on the canvas. This automatically selects the *Settings* tab on the left. Alternatively, you can access


the Settings tab by clicking the field's wrench icon. To edit a setting value, double-click it in the Settings table and enter the new value. The available settings are listed below.

You can translate each of a data definition's field values to any supported locales. To specify a field value for a translation, select the flag that represents the locale and enter the field value for the locale.

The following field settings are available. Note that some of these settings are only available for specific field types:




- **Type:** The field's type (e.g., text, radio, etc.). This setting can't be edited, but a display template can reference it.
 - **Field Label:** The field's display name.
 - **Show Label:** Whether the field label is shown.
 - **Required:** Whether users must fill out the field (not available for Boolean fields).
 - **Name:** The field's internal identifier. You can use this value in a display template to read the field's data. This value is automatically generated, but you can change it if you wish.
 - **Predefined Value:** The field's default value.
 - **Tip:** Text to display in a tooltip.
 - **Indexable:** Whether the field is indexed for search.
 - **Localizable:** Whether the field can be translated.
 - **Repeatable:** Whether users can make copies of the field.
 - **Multiple:** Whether the user can select more than one option. This is only available for Select fields.
 - **Options:** The options available for selection in Radio and Select fields. You can add and remove options, and edit each option's display name and value.
8. Click *Save* when you're done. Your new data definition then appears in the table with the pre-defined ones and any you've already added.

View Source

Default Language:  English (United States) + Add Translation

Fields Settings

| name | value |
|------------------|---------------------|
| Type | text |
| Field Label | Text |
| Show Label | Yes |
| Required | No |
| Name | Text48h2 |
| Predefined Value | |
| Tip | |
| Indexable | Indexable - Keyword |
| Localizable | Yes |
| Repeatable | No |

Text   

Radio

option 1 option 2 option 3

Figure 297.4: Configure the settings for each field in your data definition.

MANAGING DATA DEFINITIONS

There are several ways to manage your data definitions. Of course, you can edit a data definition, but you can also configure its permissions, manage its templates, copy it, or delete it.

Follow these steps to access your data definitions:

1. Open the Menu (☰) and expand your site's menu (the Site Administration menu). Then select *Content & Data* → *Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table lists any existing dynamic data lists.
2. Click the *Options* button at the top-right (ⓘ) and select *Manage Data Definitions*. The Data Definitions screen appears. A table lists the data definitions.




Figure 298.1: You can copy an existing data definition, manage its templates, and more.

You can manage your data definitions via the *Actions* menu (⋮) for each definition:

Edit: Edit the data definition. The edit screen uses the same form for creating data definitions. Note that if you edit a data definition referenced elsewhere (e.g., by a dynamic data list or display template), then you must update that reference.

Manage Templates: The *Manage Templates* screen creates and manages templates for the data definition. For details, see Using Templates to Display Forms and Lists.

Permissions: Configure the data definition's permissions. Note that these permissions are for an individual definition accessed through the Dynamic Data Lists application in *Site Administration* → *Content* → *Dynamic Data Lists*. For example, if Site members have View permission for a data definition, any Site member who also has a Role that can access the Dynamic Data Lists app and its data definitions can see this definition listed in the Manage Data Definitions screen. If you don't want this, remove the View permission for Site Member, and Site members won't see your data definition listed with the others.

Copy: The *Copy Data Definition* form copies the definition and its templates. In the form, give the copied definition a new name and description and select whether to also copy the original definition's templates. Click *Copy* when you're done. The copied definition then appears in the Data Definitions table with existing definitions. You can create new definitions based on existing ones, and then modify the copied one to suit your needs. You can, of course, edit any definition in the portal, but if you copy a definition instead, you can still access the original.

Delete: Delete the definition.

CREATING DATA LISTS

There are two places to create dynamic data lists:

1. **Site Administration:** Open the Menu and expand your Site's menu (the Site Administration menu). Then select *Content & Data* → *Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table contains any existing lists. Click the *Add* button (+) to open *New List* form. To add Dynamic Data Lists, you must have permission to access the Dynamic Data Lists app in Site Administration.
2. **Dynamic Data Lists Display widget:** Navigate to the Site page where you want this widget and add it to the page from *Add* (⊕) → *Widgets* → *Collaboration* → *Dynamic Data Lists Display*. Then click the widget's *Add List* button. This opens the *New List* form. To do this, you must have permission to create a new list in the widget.

Either option leads to the New List form:

1. Give the list a name and a description.
2. Select the list's data definition: click *Select* under the *Data Definition* field, then click the definition you want to use.
3. To use a workflow with this list, select it from the *Workflow* field.
4. To change the list's default permissions, expand the form's *Permissions* section and make your selections.
5. Click *Save* to finish creating the list. Your new list appears in the table.


299.1 Creating List Records

By default, only administrators have permission to create list records. Follow these steps to give other users this permission:


The screenshot shows a form titled 'New List' with the following fields and options:

- Name ***: A text input field containing 'Lunar Spelunking'.
- Description**: A text area containing 'Dynamic data list for lunar spelunking trips.'
- Data Definition ***: A dropdown menu showing 'Lunar Excursions' and a 'Select' button.
- Workflow**: A dropdown menu showing 'No Workflow'.
- Permissions**: A section with a right-pointing arrow.

Figure 299.1: The New List form.

1. Navigate to *Content & Data* → *Dynamic Data Lists* in Site Administration.
2. Click *Actions* () → *Permissions* for the list getting the new permissions.
3. Select *Add Record* for the Roles that should have that permission, then click *Save*. Allow unauthenticated Users to add records by giving Guest the Add Record permission.

Create new records in a list from the same places you can create the lists themselves:

1. **Site Administration:** In Site Administration, navigate to *Content & Data* → *Dynamic Data Lists*. Click a list in the table to view any existing records, then click the *Add* button (). This opens a form based on the list's data definition, which you can then fill out and submit to create a new record. To do this, you must have permission to access the Dynamic Data Lists app in Site Administration.
2. **Dynamic Data Lists Display widget:** See above for instructions on adding this widget to a page. You must then configure the widget to display a list's records.

To configure the widget to display a list's records:

- Click the widget's *Select List* button.
- In the dialog that appears, select a list, click *Save*, then close the dialog. The widget then displays the list's existing records.

To add a record:

- Click the widget's *Add* button ().

- Fill out the form that appears and click *Publish*.

See the section below for more information on configuring the widget.

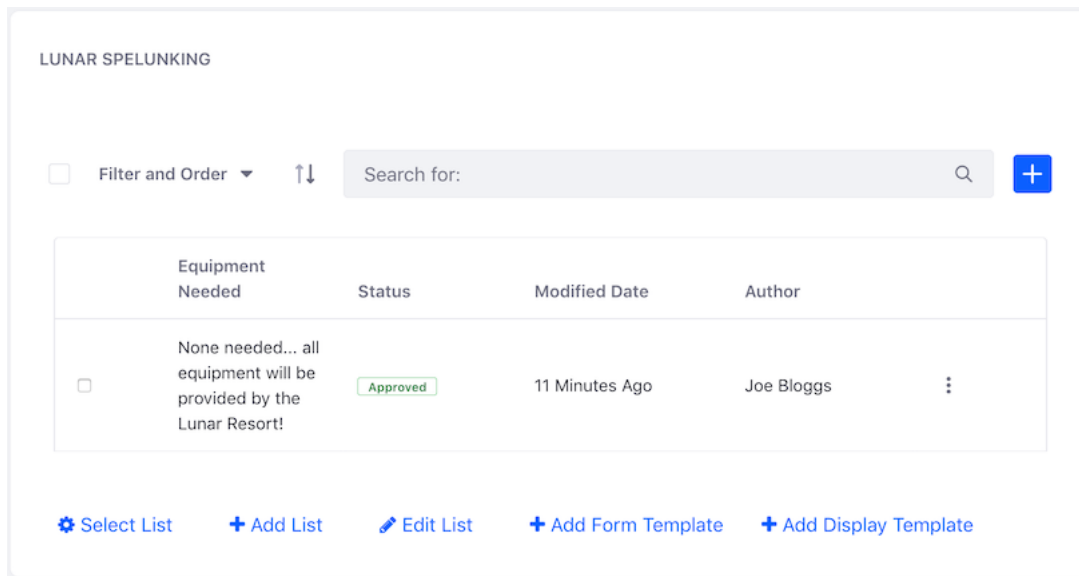


Figure 299.2: Dynamic Data Lists Display widget.

299.2 Configuring the Dynamic Data Lists Display Widget

The widget's default display template isn't exciting, but it shows the list's contents, and with permission, add and/or edit list items. To configure the widget, click its *Options* menu (⋮) and select *Configuration*. This opens the Configuration dialog, with the *Setup* tab selected by default. The Setup tab contains two other tabs:

Lists: Select the list that the widget displays. The currently displayed list appears at the top of the tab, while the available lists appear in a table. To change the widget's list, select the list from the table and click *Save*.

Optional Configuration:

Display Template: Select the display template for the list.

Form Template: Select the form template for the list.

Editable: Whether users can add records to the widget's list.

Form View: Whether to display the Add Record form by default, instead of the List View. Note that even without this selected, users can still add records via the widget's *Add* button (+).

Spreadsheet View: Whether the List View displays each record in a row, with columns for the record attributes.

When finished, click *Save* and close the Configuration dialog.




Figure 299.3: The Dynamic Data Lists Display widget's optional configuration.



USING TEMPLATES TO DISPLAY FORMS AND LISTS

After creating data definitions and lists, you can control how the form appears to users, and how the resulting list of records is displayed. You do this by creating templates for each view (form view for displaying the form and display view for the list of records) and selecting them in the DDL Display portlet. For example, you might need to create a form with a subset of a data definition's fields. Rather than creating a new definition, you can create a template that displays only the fields you want from the existing definition. You could also use a template to arrange fields differently, and/or with different labels and configuration options.

Data definitions can have as many form and display templates as you care to create (or none, if you're satisfied with the default templates). You then choose a list's template in the Dynamic Data Lists Display widget.





300.1 Managing Display and Form Templates

Since Display and Form Templates correspond to a particular data definition, they're accessed from the Data Definitions screen of the Dynamic Data Lists application in Site Administration. See the [Creating Data Definitions](#) article for instructions on accessing this screen.

The Data Definitions screen lists each definition in a table. To start working with a definition's templates, click the definition's Actions button () and select *Manage Templates*. This opens a screen that lists the definition's templates. You can edit, copy, delete, or configure permissions for a definition via its Actions button ().

CREATING FORM TEMPLATES

Form templates control how the data entry form appears for a data definition. Follow these steps to create a form template for a definition:

1. Open the Menu () and expand your Site's menu (the Site Administration menu). Then select *Content & Data* → *Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table lists any existing dynamic data lists.
2. Click the *Options* button at the top-right () and select *Manage Data Definitions*. The Data Definitions screen appears. A table lists any existing data definitions.
3. Click the definition's *Actions* button () and select *Manage Templates*. This lists the definition's templates.
4. Click the *Add* button () and select *Add Form Template*. This presents the same kind of graphical, drag-and-drop interface used to create definitions.
5. Give your template a name, then expand the *Details* section and give it a description. You can also use the *Mode* selector to select which mode the template applies to (*Create* or *Edit*).
6. Scroll down to the graphical designer in the *View* tab, and make your desired changes. For example, you can move or delete fields, change field labels, and more.
7. Click *Save* when you're finished.

Alternatively, you can create form templates from the Dynamic Data Lists Display widget:

1. Follow the instructions in the *Creating Data Lists* article for adding and configuring the widget in a Site page. Make sure to configure the widget to show the list you're creating a template for.
2. Click the widget's *Add Form Template* button. This opens the same form as above for creating a form template for the list's definition.





CREATING DISPLAY TEMPLATES

For every data definition, you can create as many displays as you need. If you've created a form template that doesn't show all the fields of a particular data definition in the data list's form view, you probably don't want to display those fields in the list view, either. Modify the list view using Display Templates.

Note: If you're familiar with web content templates, display templates customize the display of a list in the same way. Display templates are written in FreeMarker or Velocity, pulling data from the data definition in the same way that web content templates pull data from their structures. Also similar to web content templates, display templates can be embedded in other display templates. This allows for reusable code, JavaScript library imports, or macros imported by Velocity or FreeMarker templates in the system. Embedding display templates provides a more efficient process when you have a multitude of similar data definitions. Just import an embedded display template and work off of it for your new display template.

As with form templates, you can create display templates from the Dynamic Data Lists app in Site Administration or the Dynamic Data Lists Display widget.

Follow these steps to create a display template from Site Administration:

1. Open the Menu () and expand your Site's menu (the Site Administration menu). Then select *Content* → *Dynamic Data Lists*. This opens the Dynamic Data Lists screen. A table lists any existing dynamic data lists.
2. Click the *Options* button at the top-right () and select *Manage Data Definitions*. The Data Definitions screen appears. A table lists existing data definitions.
3. Click the definition's *Actions* button () and select *Manage Templates*. This lists the definition's templates.
4. Click the *Add* button () and select *Add Display Template*.
5. Give the template a name, expand the *Details* section of the form, and give it a description. Also in the details section of the form, select the templating language to use from the *Language* selector. You can choose FreeMarker, or Velocity.

6. In the *Script* section of the form, create your template in the editor using the templating language you chose in the previous step. The palette to the left of the editor contains common variables. Click a variable to insert it in the editor. The editor also autocompletes. In a FreeMarker template, type `${`, which opens an autocomplete list of common variables. Select a variable to insert it in the editor. Alternatively, you can upload a complete script file via the *Browse* button below the editor.
7. Click *Save* when you're done creating the template.




Figure 302.1: Create your display template in the editor.

Alternatively, you can use the Dynamic Data Lists Display widget to create display templates:

1. Follow the instructions in the [Creating Data Lists](#) article for adding and configuring the widget in a site page. Make sure to configure the widget to show the list's definition you're making into a template.
2. Click the widget's *Add Display Template* button. This opens the same form as above for creating a display template for the list's definition.

302.1 Display Template Editor

Helper variables are available in the template editor. These provide access to most of the data that you'll use in creating Display Templates. The variables under the heading *Data List Variables* let you inject specific information about the data definition the template is being created for:

Data Definition ID: `reserved_ddm_structure_id`

Data List Description: reserved_record_set_description

Data List ID: reserved_record_set_id

Data List Name: reserved_record_set_name

Template ID: reserved_ddm_template_id

Inside a template, these variables give the ID for the record set as well as the name, description, and data definition.

Display the list of records by retrieving them and assigning them to the handy records variable. Retrieve the list's records from `DDLDisplayTemplateHelper`, which contains these functions:

```
getDocumentLibraryPreviewURL
getHTMLContent
getLayoutFriendlyURL
getRecords
renderRecordFieldValue
```

`DDLDisplayTemplateHelper` performs common tasks. Use the `getRecords` method to access a data definition's entries and assign them to a records variable:

```
<#assign records = ddlDisplayTemplateHelper.getRecords(reserved_record_set_id)>
```

This *fetches* the records of the associated data list. You haven't done anything with them yet, so your display is still empty. To list all the records, use the *Data List Records* helper in the sidebar of the template editor. Remember to place your cursor in the proper place in the template editor window, then click *Data List Records*. This code appears at the cursor:

```
<#if records?has_content>
  <#list records as cur_record>
    ${cur_record}
  </#list>
</#if>
```

This default code snippet spits out everything in the database for the given data definition, which is ugly and practically useless:

```
{uuid=52c4ac1c-afe7-963c-49c6-5279b7030a99, recordId=35926, groupId=20126,
companyId=20099, userId=20139, userName=Test Test, versionUserId=20139,
versionUserName=Test Test, createDate=2018-07-26 14:31:51.056,
modifiedDate=2018-07-26 14:31:51.058, DDMSStorageId=35927, recordSetId=35922,
recordSetVersion=1.0, version=1.0, displayIndex=0, lastPublishDate=null}
```

Here's a simple example template that uses a list based on the embedded Contacts data definition, and only displays the Company Name and Email fields in a bulleted list:

```
<#assign records = ddlDisplayTemplateHelper.getRecords(reserved_record_set_id)>

<h1>Here are contacts by company name and email address.</h1>

<#if records?has_content>
  <#list records as cur_record>
    <ul>
      <li>
        <!-- The below gets the Company field and wraps it in an <em> tag -->
        Company Name: <em>${ddlDisplayTemplateHelper.renderRecordFieldValue(cur_record.getDDMFormFieldValues("company")?first, locale)}</em>
        <!-- The below gets the Email field and wraps it in an <em> tag -->
        Email: ${ddlDisplayTemplateHelper.renderRecordFieldValue(cur_record.getDDMFormFieldValues("email")?first, locale)}

```

```
</li>
</ul>
</#list>
</#if>
```

Here's what it looks like:

Here are contacts by company name and email address.

- Company Name: *Liferay*
Email: test@liferay.com
- Company Name: *The Lunar Resort*
Email: james@lunarresort.com

Figure 302.2: Extract appropriate display information, rather than spitting out the whole object.

Now you're prepared to make data lists beautiful using Display Templates.

WORKFLOW

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay's workflow engine is named *Kaleo*. In Greek, *Kaleo* means "called ones," which is appropriate for a workflow engine that calls users to participate in a review process designed for them. Workflow makes it possible to define any number of simple to complex business processes/workflows, deploy them, and manage them through a portal interface. The processes have knowledge of Users, Groups and Roles. You don't have to write a single line of code to accomplish this: it's defined in an XML document. If you're a Liferay Digital Experience Platform (DXP) customer, you also have access to a visual process builder.

There are several steps to effective workflowing:

- Designing review processes in XML
- (DXP only) Visually designing review processes
- Uploading workflow definitions
- Activating workflow for enabled assets
- Managing Workflow definitions
- Sending assets through review
- (DXP only) Using Workflow Metrics

After all that, you'll be familiar with using Liferay's workflow engine to set up approval processes for any workflow-enabled content.

303.1 What's New with Workflow?

There are some noteworthy enhancements to the workflow functionality:

303.2 DXP Feature: Workflow Metrics

For DXP subscribers, the *Workflow Metrics* feature was introduced. Configure one or more Service Level Agreements (SLAs; think of these as deadlines) on a workflow definition's events, and workflow submissions are accordingly tracked and reported as on time or overdue.

Service Level Agreements (SLAs)

SLAs define the deadlines on a workflow process's events. They're like a contract between the workflow participants and Users submitting workflow items.

SLAs can be formally agreed-upon deadlines between you and your customers, or informally created to meet internal goals, tracking events such as:

- Total time to resolution
- Time to complete a specific workflow task

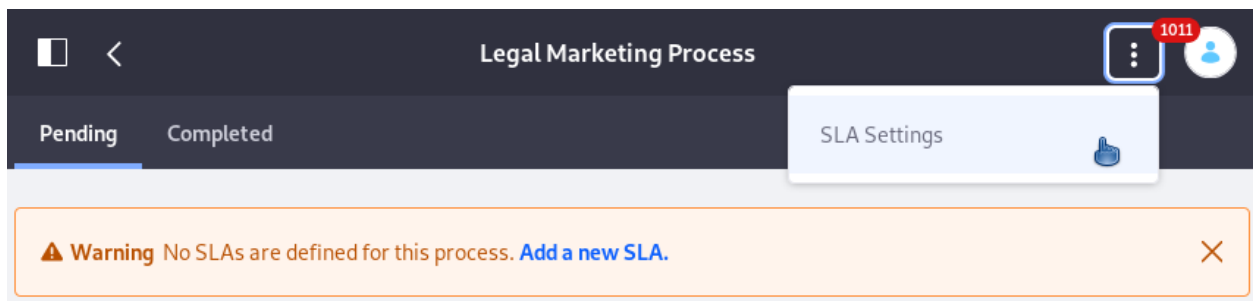


Figure 303.1: Use Service Level Agreements (SLAs) to define how workflow metrics are reported.

For each workflow event to track, set the SLA duration and when the timer should be paused, if at all.

Workflow Reports

Once an SLA is set, workflow submissions that trigger the SLA timer are automatically reported on by the workflow metrics framework, and given the status *on time* or *overdue*.

See the article on Workflow Metrics to learn more about SLAs and available reports.

303.3 Control Panel Reorganization

The Workflow section of the Control Panel is now a top-level section with its own subcategories: Process Builder and Submissions (Metrics, too, if you're a DXP subscriber). In Liferay DXP 7.1 Workflow was nested under Control Panel → Configuration.

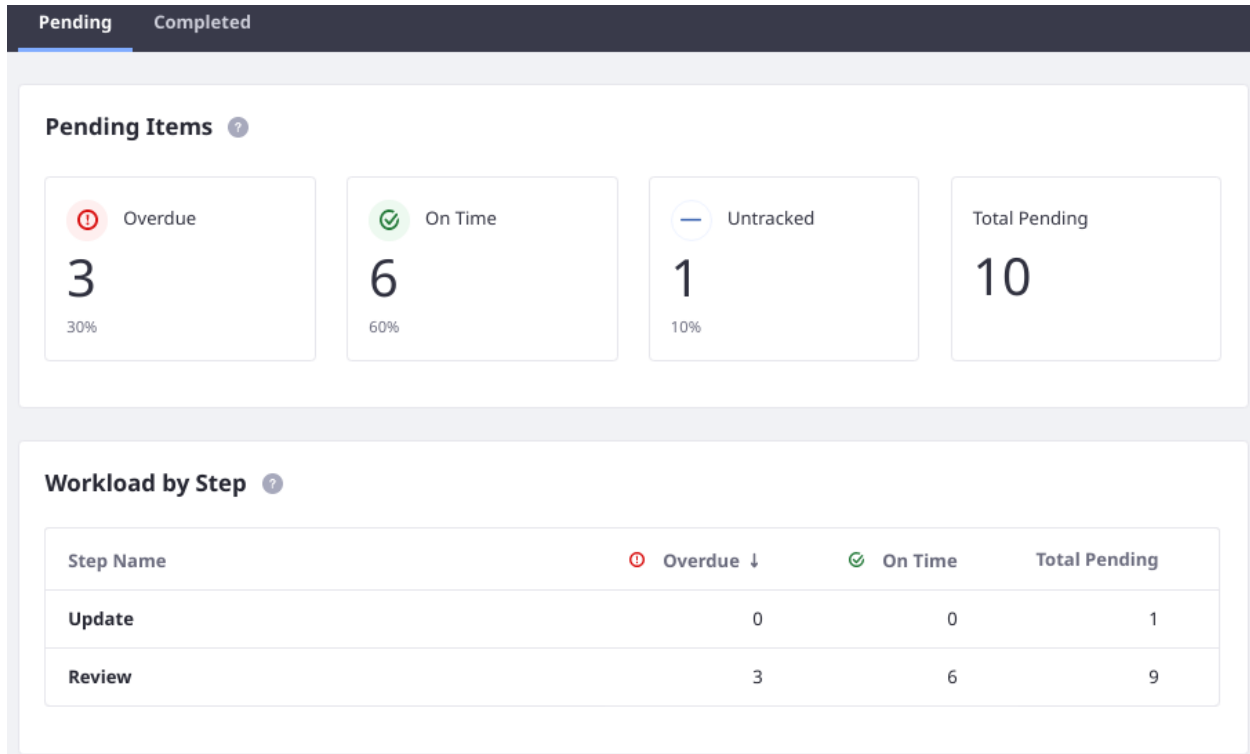


Figure 303.2: See Workflow Reports generated based on your SLAs.

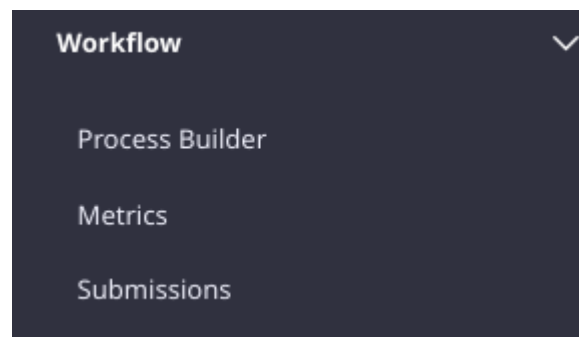


Figure 303.3: Workflow has a top-level entry in the Control Panel.

303.4 Workflow Definition Permissions: System Settings

The Workflow System Settings category (Control Panel → Configuration → System Settings) has a new system scoped configuration entry: *Workflow Definitions*. There's just one configuration option, but it's important: Enabling it gives administrators permission to publish workflows and scripts.

Workflow Definition

This configuration was not saved yet. The values shown are the default.

Allow Administrators to Publish and Edit Workflows

When selected, users with the Regular Administrator role in a virtual instance can publish new workflow definitions, including scripting content that is executed inside the portal (and potentially in other virtual instances).

Save

Cancel

Figure 303.4: Explicit permission must be granted before administrators are allowed to publish and edit workflow definitions.

Create your own workflows from scratch or leverage existing workflows.

303.5 Embedded Workflows

In addition to the Single Approver definition, there are some workflow definitions that ship with Liferay DXP but are not pre-installed, since they're primarily included for test cases. They can be found in the Liferay source code in

```
/modules/apps/portal-workflow/portal-workflow-kaleo-runtime-impl/src/main/resources/META-INF/definitions
```

They're also in your Liferay installation. Open your Liferay installation's `osgi/portal/com.liferay.portal.workflow` and then find and open the `com.liferay.workflow.kaleo.runtime.impl-[version].jar`. The definitions are in the `META-INF/definitions` folder.

ACTIVATING WORKFLOW

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Assets are integrated with the workflow framework if their content is meant to be sent through review processes. Since this is most often the case, there are many out of the box assets that integrate with workflow. In this article, learn how to select a workflow for use with each of these workflow-enabled assets.

304.1 Workflow Assets

Activate a workflow for these assets in Control Panel → Workflow → Process Builder, in the Configuration tab:

- Blogs Entry
- Calendar Event
- Comment
- Knowledge Base Article
- Message Boards Message
- Page Revision
- User
- Web Content Article
- Wiki Page

Activate workflow for these assets at the Site level in Site → Configuration → Workflow:

- Blogs Entry
- Calendar Event
- Comment
- Knowledge Base Article
- Message Boards Message
- Page Revision

- Web Content Article
- Wiki Page

What's the difference between the Site workflow configuration and the Control Panel Workflow configuration? As with most scoped configurations, the higher level setting (in the Control Panel) sets the default behavior. It's overridden at the more granular level (in the Site menu).

User doesn't appear on the Site list because adding users is strictly a portal-wide administrator activity. Only assets that can be added and configured at the Site level (for example, those that are accessed from the Site menu) have workflow configuration controls at the Site level.

304.2 Activating Workflow in Applications

Some assets that are workflow-enabled are activated in their respective application:

Activate workflow for Web Content Folders from the folder settings menu:

Activate workflow for Documents and Media Folders from the folder settings:

Enable workflow on Dynamic Data List entries in each list's Add form:

Activate workflow for each individual form's entries from the Form Settings screen:

304.3 Workflow Behavior

Most of the resources listed above behave just as you might expect with workflow activated: The Publish button for the resource's *Add* form is replaced by a *Submit for Publication* button, and instead of instant publication, the asset has its status set as *Pending* and must proceed through the workflow before publication.

Page revisions are slightly different. Page revisions only occur in staging environments that have Page Versioning enabled. When a Page Variation or Site Page Variation is created, its creator must click *Submit for Publication* at the top of the page, and the variation must be approved in the workflow before it can be published to the live Site.

DETAILS ▼

Name *

Lunar Resort Holiday Getaway Content

Description

PARENT FOLDER >

STRUCTURE RESTRICTIONS AND WORKFLOW ? ▼

Use Structure Restrictions and Workflow of the Parent Folder (Home)

Define Specific Restrictions and Workflow for This Folder (Lunar Resort Holiday Getaway Content)

Default Workflow for This Folder (Lunar Resort Holiday Getaway Content)

No Workflow ⌵

Save

Figure 304.1: Activate workflow on Web Content folders from the folder’s edit screen.

The screenshot shows a folder configuration interface. At the top, the 'Parent Folder' is set to 'Home'. The 'Name' field contains 'Lunar Resort Shuttle Images'. Below this is a large 'Description' text area. A section titled 'Document Type Restrictions and Workflow' is expanded, showing three radio button options: 'Use Document Type Restrictions and Workflow of the Parent Folder (Home)', 'Define Specific Document Type Restrictions and Workflow for This Folder (Lunar Resort Shuttle Images)', and 'Default Workflow for This Folder (Lunar Resort Shuttle Images)'. The third option is selected. Below these options is a dropdown menu for 'Default Workflow for all Document Types' currently set to 'No Workflow'. At the bottom are 'Save' and 'Cancel' buttons.

Figure 304.2: Activate workflow on Documents and Media folders from the folder’s edit screen.

The screenshot shows a Dynamic Data List configuration interface. The 'Name' field contains 'Emergency Contact List'. Below it is a 'Description' text area. The 'Data Definition' field is set to 'Contacts', with a 'Select' button below it. The 'Workflow' dropdown menu is set to 'Single Approver'. At the bottom, there is a 'Permissions' section with a right-pointing arrow and 'Save' and 'Cancel' buttons.

Figure 304.3: Activate workflow for each individual Dynamic Data List.

Form Settings ✕

Require user authentication

Require CAPTCHA

Save answers automatically.

Redirect URL on Success

Enter a valid URL.

Select a Storage Type

json

Select a Workflow

Single Approver

Cancel Done

Figure 304.4: Activate workflow on each form's entries from the Form Settings window.




Figure 304.5: Instead of a Publish button, a Submit for Publication button appears for workflow-enabled resources.



Figure 304.6: With workflow enabled on Page Revisions, the Site administrator must submit their page variation for publication before it can go live.

MANAGING WORKFLOWS

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Like other resources, workflow definitions can be added, edited, and deleted. But that's just the beginning of workflow management.

- Because workflow definitions can be complex works in progress, they can be versioned.
- Unpublished drafts can be saved.
- Because workflow definitions are XML files, they're portable. Thus, they can also be uploaded.

Start by learning basic workflow management.

305.1 Workflow Definition Publication Permissions

Users with permission to edit or publish workflow definitions can add Groovy scripts to the workflow. Access to the scripting engine means access to the Java Virtual Machine (JVM) of the server. Users who publish (or edit) workflow definitions containing scripts, therefore, can get access to any data within the reach of the JVM, such as data contained in a separate Virtual Instance of Liferay DXP itself.

Because of this far-reaching access, permission to create or edit workflow definitions is limited to Regular Administrators of the Default Virtual Instance.

To grant Users with these Roles the workflow publication access in additional Virtual Instances,

1. Make sure you understand the access you're granting these admins.
2. Navigate to Control Panel → System Settings → Workflow → Workflow Definition.
3. Check the box for the setting *Allow Administrators to Publish and Edit Workflows*.

This only applies to Virtual Instances that have been added to the system. The Default Virtual Instance provides workflow publication access to Regular Administrators (via Control Panel → Configuration → Workflow), and, if running Liferay DXP, to Site Administrators and other Users with access to the Kaleo Forms Admin applications.

305.2 Adding, Editing, and Deleting

To add a workflow definition,

1. Go to Control Panel → Workflow → Process Builder.
2. In the default view (Workflows), click the Add button (+).
3. From here you're either writing an XML definition, designing a definition in the visual designer (DXP only), or uploading an existing definition.

To edit a definition,

1. Go to Control Panel → Workflow → Process Builder.
2. Click the *Actions* button (⋮) for the workflow, and click *Edit*.

To delete a definition,

1. Go to Control Panel → Workflow → Process Builder.
2. Click the *Actions* button (⋮) for the workflow, and click *Unpublish*.

A published workflow cannot be deleted, so you must unpublish its workflow definition first.

You can't unpublish a definition if it's activated for an asset. First dissociate the workflow definition from any assets that use it. See here for more information.

305.3 Uploading Workflow Definitions

If you have a local XML definition file (perhaps you want to create a new workflow based on one of the embedded workflows), upload it to Liferay DXP:

1. Navigate to Control Panel → Workflow → Process Builder.
2. Click the *Add* button (+).
3. Name the workflow.
4. In the *Source* tab, click the hyperlink *import a file* in the sentence
Write your definition or import a file
5. Find the XML file and upload it. Once uploaded, the definition's XML appears in the workflow editor.
6. If it's ready to publish, click *Publish*. Otherwise, *Save* it and it stays Unpublished.

What's the difference between saving and publishing?

305.4 Published Versus Unpublished



The difference between a published and unpublished workflow is important:

Published: Validation is complete, and the workflow can be assigned to assets.

Unpublished: Validation is not performed on the unpublished workflow, and it cannot be assigned to assets until it's published.

305.5 Workflow Versions

You're making a simple edit to a workflow, when suddenly you remember you have a meeting with your boss. Quickly you save the workflow and hurry off to your meeting. Congratulations! You were promoted to Director of Business Productivity! You have no time to edit workflows now, so your colleague must finish editing and publishing the workflow. Unfortunately, in all the excitement of your promotion, you forgot what you changed in the workflow. It's best to revert to the prior version and start editing it from scratch.

1. Open the workflow editor. Go to *Control Panel* → *Workflow* → *Process Builder*, and select the workflow from the list (click its title to open the editor).
2. Click the *Information* button ()
3. There are two information panel sections: *Details* and *Revision History*.
The *Details* screen shows information about the creation of the workflow and its last modification, and a summary of the total modifications.
The *Revision History* screen shows the workflow's current and prior, restorable versions. To view an old workflow or to restore it if you're sure it's the right version, click the *Actions* button () and select either *Preview* or *Restore*.
4. When you click *Restore* and see the success message, the prior version is the current version of the workflow. You can now edit the restored version of the workflow.
5. If edits are necessary, edit and click *Update*. This creates another version of the workflow.

Alternatively, you can refer to the embedded definitions to get workflow definition ideas.

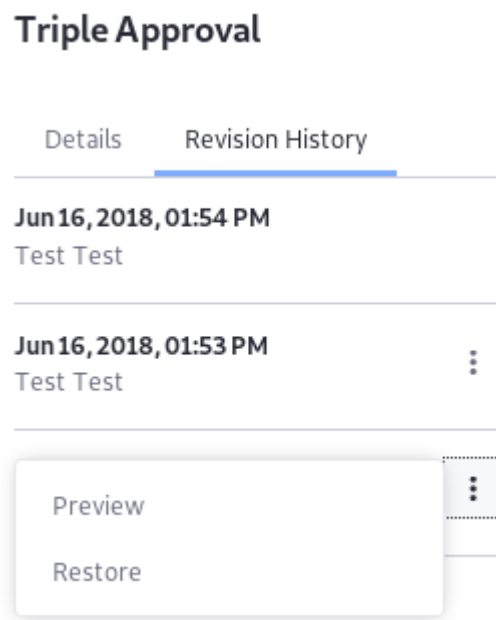


Figure 305.1: View and restore prior versions of a workflow.

REVIEWING ASSETS

This document has been updated and ported to Liferay Learn and is no longer maintained here.

User interaction is required at each workflow process Task Node. Where do users complete tasks? In the *My Workflow Tasks* application of the User menu.

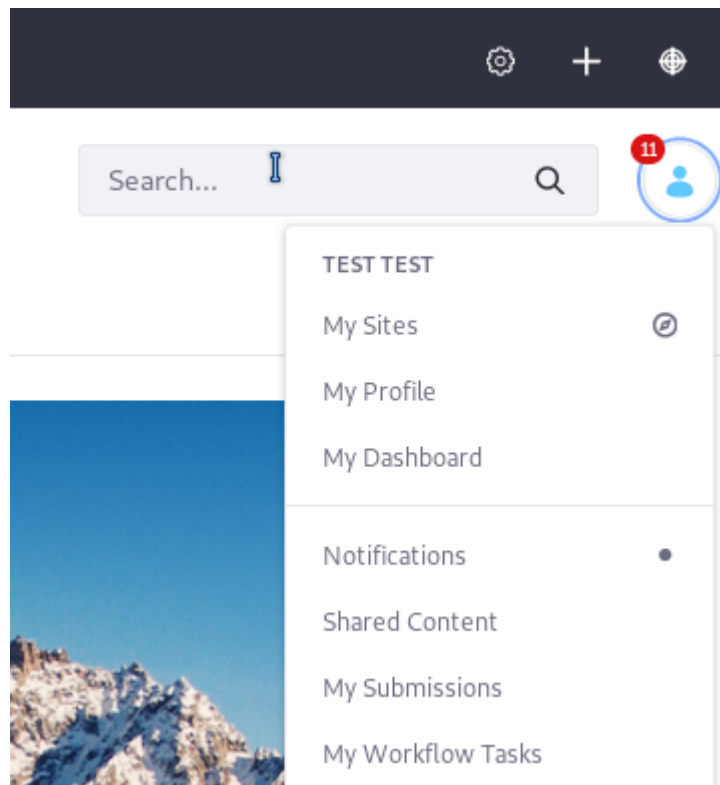


Figure 306.1: Users manage workflow tasks from their My Workflow Tasks widget.

306.1 Asset Submission to Workflow

If an Asset has a workflow activated, when a user clicks Submit for Publication, the workflow definition determines the next step. A user assigned a Role associated to the workflow task receives a Notification indicating that there is a workflow task to complete.

Workflow Assignees Have Implicit Asset Permissions: Users with permission to execute a workflow task (e.g., Users with the Portal Content Reviewer Role) have full resource action permissions over the assets they can review. These permissions apply in the My Workflow Tasks widget in the User's personal page and anywhere else actions on the Asset can be performed.

For example, consider a User with two permissions:

- The Portal Content Reviewer Role enables Users to review workflow submissions and grants edit and delete permissions on the content they're reviewing.
- Users also have permission to view Web Content Articles in the Site's *Content* section.

Neither permission explicitly grants the User management permissions on Web Content Articles. Users cannot normally edit or delete a Web Content Article, for example. However, if a Web Content Article is sent to the workflow, Users can access the Web Content Article for review (in their *Assigned to Me* or *Assigned to my Roles* section of My Workflow Tasks), and they can edit or delete the content while reviewing it in the workflow. While it's in the status *Pending*, they can also edit or delete the article from Site Administration → Content → Web Content because of their implicit permissions granted by the workflow system. This additional permission is temporary, and the normal resource permissions are activated once the Web Content Article exits the workflow process (for example, it's rejected or approved).

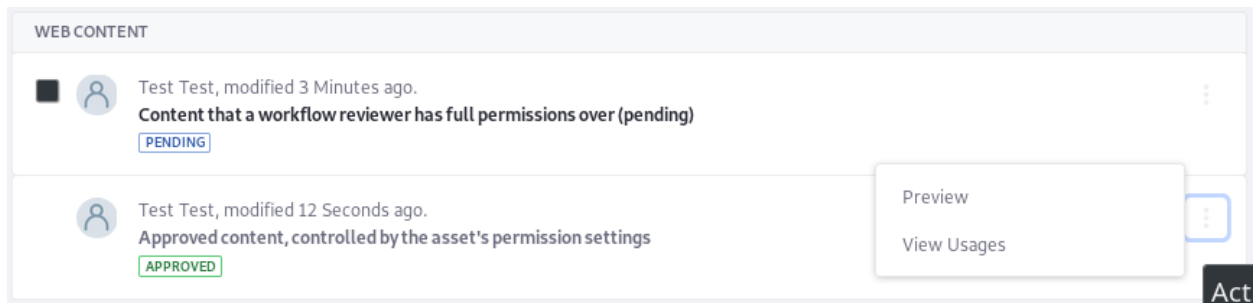


Figure 306.2: A User with VIEW permission on Web Content cannot manage Approved Articles.

306.2 Assigning the Task

Workflow Tasks can be completed only by certain users, based on the Assignment.

All workflow tasks assigned directly to a user are listed in the My Workflow Task widget's *Assigned to Me* tab.

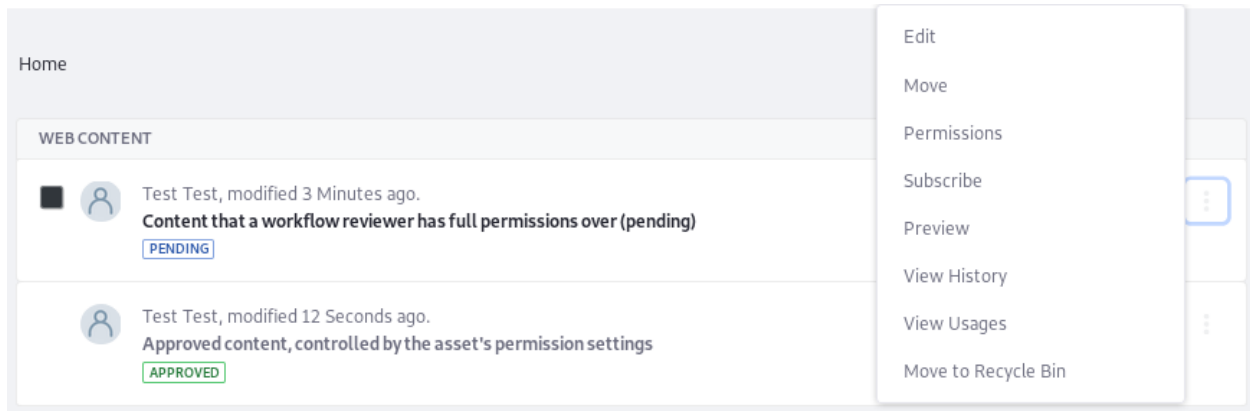


Figure 306.3: A User with access to Web Content in the Workflow can manage Pending Articles.

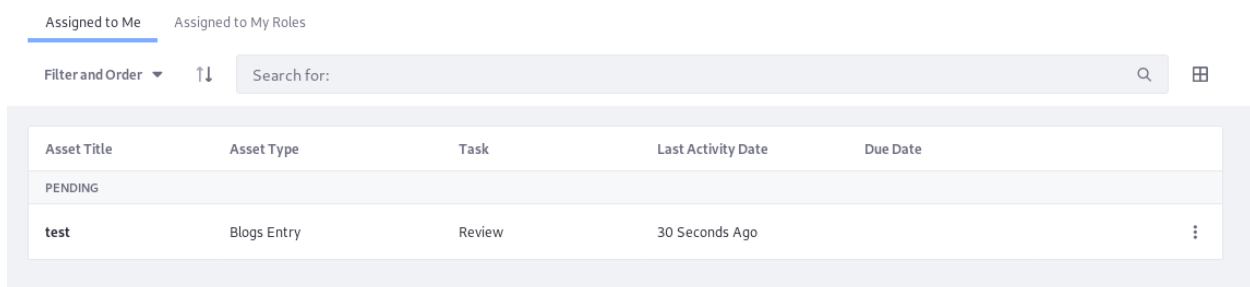


Figure 306.4: The assets assigned to a user are listed in *Assigned to Me*.

If a workflow was assigned to a Role that the user occupies, the workflow's tasks appear in the *Assigned to My Roles* tab.

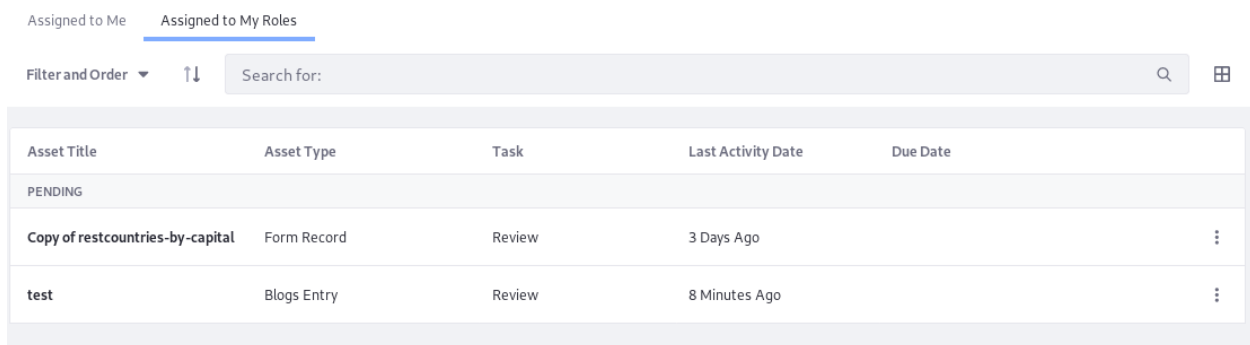



Figure 306.5: The Assets assigned to Roles are listed in each associated user's *Assigned to My Roles* tab.

To claim a task, the user must move the task into the *Assigned to Me* tab.

1. Click the Asset's *Actions* button () and select *Assign to Me*.

2. Add a comment in the pop-up box if necessary, and click *Done*.


Alternatively, assign the task to another user.

1. Click the Asset's *Actions* button () and select *Assign to*
2. Select the user to assign the task.
3. Add a comment in the pop-up box if necessary, and click *Done*.

306.3 Completing the Task

Once a task is assigned, it's ready to be completed.

There's a fast way to send an asset along in the review process:

1. From the *Assigned to Me* tab, click the *Actions* button () and select *Approve* or *Reject*.

Note that options names in this menu are identical to the workflow definition's Transition names. Your menu might have different options than the *Approve* and *Reject* options in the figure below.

2. Enter a comment if desired and click *Done*.





Figure 306.6: Complete Tasks right from the *Assigned to Me* list.

Here's how to get a closer look at the Asset before sending it along in the workflow:


1. From the *Assigned to Me* tab, click the title of the Asset to review.

The Task screen appears showing details about the Asset:

2. Inspect the Asset to your liking (or even edit it if you have permission) and click the *Actions* button ().
3. Send it along in the workflow by clicking one of the Transition names (for example, *Approve* or *Reject* in the Single Approver Definition).

And, you're done. Once you've completed your workflow tasks, kick back and wait for more to come in.

Assigned to ⋮

 Melvin Dooitrong

State

Review


Create Date




5/20/19 6:18 PM ⏏

Due Date

Never

Preview of Blogs Entry ∨


 en-US

 **Synergistic Corporate Brainwashing**  

Leverage agile frameworks to provide a robust synopsis for high level overviews. Iterative approaches to corporate strategy foster collaborative thinking to further the overall value...

Comments ∨

0 COMMENTS [Subscribe](#)

 Tr

[Reply](#)

Activities >

Figure 306.7: Inspect Assets before completing the Task.

WORKFLOW METRICS: THE SERVICE LEVEL AGREEMENT (SLA)

A brand new feature in Liferay DXP 7.2, *Workflow Metrics* gives insights into the time certain workflow events take to complete. To use it, set up deadlines on a workflow process's events. These deadline configurations are referred to as SLAs (Service Level Agreements). Once defined, Workflow Reports measure compliance with the SLAs.

Requires Elasticsearch: To use Workflow Metrics, you must be using Elasticsearch to index your @product data. Read [here](#) to learn about installing Elasticsearch.

SLAs define the deadlines on a workflow process's events. They're like a contract between the workflow participants and Users submitting workflow items.

Workflow Reports shows data for all processes with SLAs, including each workflow item's SLA status: on time or overdue.

Editing a Workflow with SLAs: Editing a workflow (e.g., removing nodes, editing a task name) with SLAs defined on it may invalidate the SLA for items already in the workflow/SLA pipeline.

Creating or Editing SLAs for Active Processes: Editing an SLA's duration or defining a new SLA while items are already in the workflow process causes a recalculation for all instances currently in the workflow. Completed workflow instances are not recalculated.

307.1 Adding SLAs

To add an SLA,

1. Go to Control Panel → Workflow → Metrics.

2. Click on the title of the Process.

The Reports UI for the process is displayed.

3. If there's no SLA for the process, a warning message stating as much appears. Click the *Add a new SLA* link from the warning to access the New SLA form directly.

Alternatively, click the Options (ⓘ) menu and select *SLA Settings*.

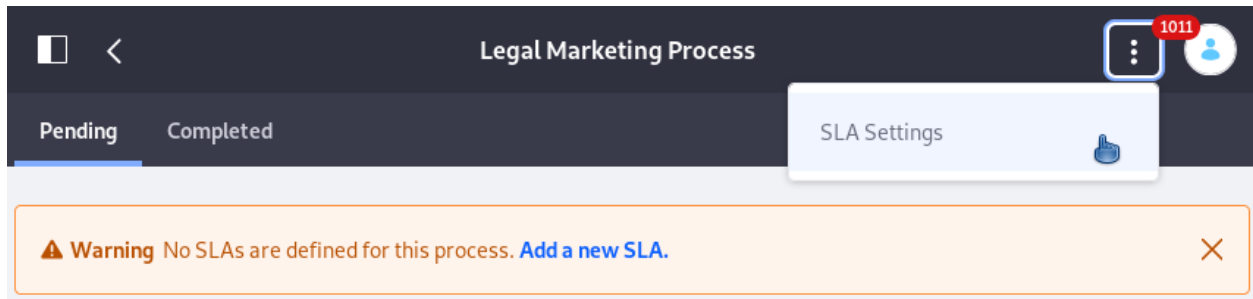


Figure 307.1: Add SLAs to a workflow definition from the Metrics application.

4. On the SLAs screen, click the *Add* button (+).
5. In the New SLA form, Give the SLA a Name and Description.
6. Define the time frame for the SLA, specifying three things:
 - **Start:** When the item makes it to the event defined here, the SLA timer begins counting.
 - **Pause:** If there's an event in the workflow when time should stop counting, enter it here. For the Single Approver workflow, you might choose to pause the SLA timer when the item is in the Update task.
 - **Stop:** Choose when the SLA is completed. If the item makes it to the Stop event before the defined SLA duration (the deadline), it's *On Time* according to the SLA. If it fails to make it to the Stop event in the specified duration, it's *Overdue*.
7. Define the duration (i.e., the deadline) for the SLA. Fill out at least one of the two time boxes.

Days: Enter a whole number of days.

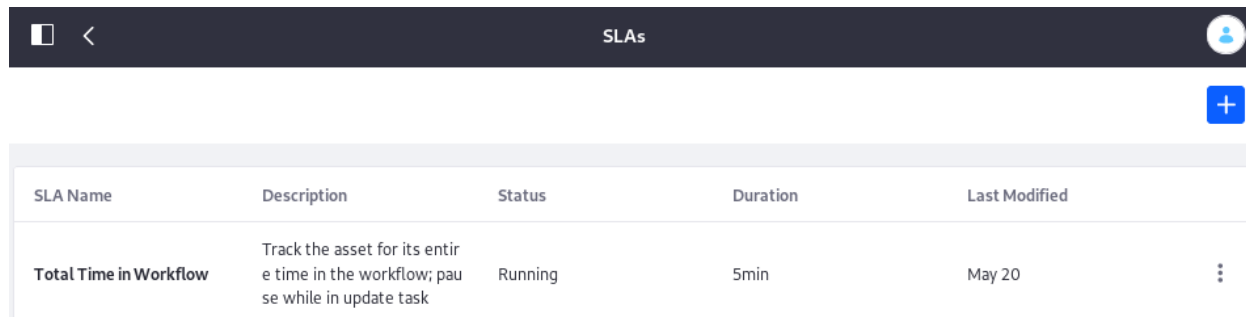
Hours: Enter hours and minutes in the format HH:MM
8. Once you click *Save*, you'll see the SLA listed on the SLAs screen.

System Calendar: By default, there's an internal calendar that assumes the SLA duration should continue counting all the time: in other words, 24 hours per day, seven days per week. If you need a different calendar format, find a friendly developer to create a custom calendar. Official docs will be written for this extension point, but the basic idea is to implement the `WorkflowMetricsSLACalendar` interface. New implementations of this service are picked up automatically by the Workflow Metrics applications, so it becomes available as soon as the module holding the service implementation is deployed. The interface has three methods to implement:


```
public Duration getDuration(    LocalDateTime startLocalDateTime, LocalDateTime endLocalDateTime);
public LocalDateTime getOverdueLocalDateTime(    LocalDateTime nowLocalDateTime, Duration
remainingDuration);
```

```
public String getTitle(Locale locale);
```

See the `DefaultWorkflowMetricsSLACalendar` from the `portal-workflow-metrics-service` module for example code. If you define a new calendar, a new option becomes available in the Add SLA form, allowing you to choose from the default 24/7 calendar or any custom ones you've provided (for example, make the timer run for 9 hours per day, 5 days per week).



| SLA Name | Description | Status | Duration | Last Modified |
|------------------------|---|---------|----------|---------------|
| Total Time in Workflow | Track the asset for its entire time in the workflow; pause while in update task | Running | 5min | May 20 |

Figure 307.2: Manage SLAs from the SLAs screen.

Valid Start and Stop Events

Any workflow task can be used as a start or end parameter for the SLA.

When defining the tasks to act as the SLA's Start Events, choose between three events:

- The start node
- Entry into a task
- Exit from a task

When defining the tasks to act as the SLA's Stop Events, choose between three events:

- Entry into a task
- Exit from a task
- The end node

The SLA can be paused at any task that falls between the start node and the end node, and it's defined by setting the node(s) when the SLA should be paused. *The SLA timer is paused the entire time a workflow item is in the specified node.*

Durations

Define the SLA durations in at least one of the available boxes (Days and Hours). Here are some examples:

Example Duration: 1 day, 24 hours Valid configuration — Days: 1

Invalid — Hours: 24:00. The Hours box must not exceed 23:59.

Example Duration: 36 hours Valid — Days: 1, Hours: 12:00

Invalid — Days: 1.5. Only whole numbers are accepted.

Example Duration: 6.5 hours Valid — Hours: 06:30

Once your SLAs are configured, activate the workflow on an asset, stretch your fingers, and get ready for the submissions to roll in if you're one of the workflow assignees. You're on the hook to get those workflow items through the process within the SLA duration!

WORKFLOW METRICS: REPORTS

As soon as you enter the Metrics screen (Control Panel → Workflow → Metrics) you see metrics on each workflow installed in the system.

| Process Name | Overdue ↑ | On Time | Total Pending | |
|-------------------------------------|-----------|---------|---------------|---|
| Single Approver | 6 | 0 | 7 | ⋮ |
| Legal/Marketing Workflow | 0 | 0 | 0 | ⋮ |
| Category Specific Definition | 0 | 0 | 0 | ⋮ |
| Scripted Assignment Single Approver | 0 | 0 | 0 | ⋮ |

Figure 308.1: In this view, the only process with pending items is the Single Approver.

A table view of all installed workflow processes shows you how many items are Overdue, how many are On Time, and how many are Pending in the workflow process.

There's more to Metrics than the overview report though. Get more detailed reports by clicking on one of the workflow processes.

Requires Elasticsearch: To use Workflow Metrics, you must be using Elasticsearch to index your @product data. Read here to learn about installing Elasticsearch.

308.1 Understanding Reports

The Reports UI has two main views, represented as tabs: *Pending* (changed to *Dashboard* as of 7.0 Service Pack 1—SP1) and *Completed* (changed to *Performance* in SP1).

Pending/Dashboard items are those currently in the workflow process, and include items untracked by the SLA. This might include items in the paused step of the workflow, or items that are outside the scope of the SLA duration.

Completed/Performance items show any item that has completed processing in the workflow. SP1 includes a new chart on this screen: *Completion Velocity*.

When you first click into the metrics for a specific process, you're presented with two valuable reports on pending items: the Pending Items overview and Workload by Step.

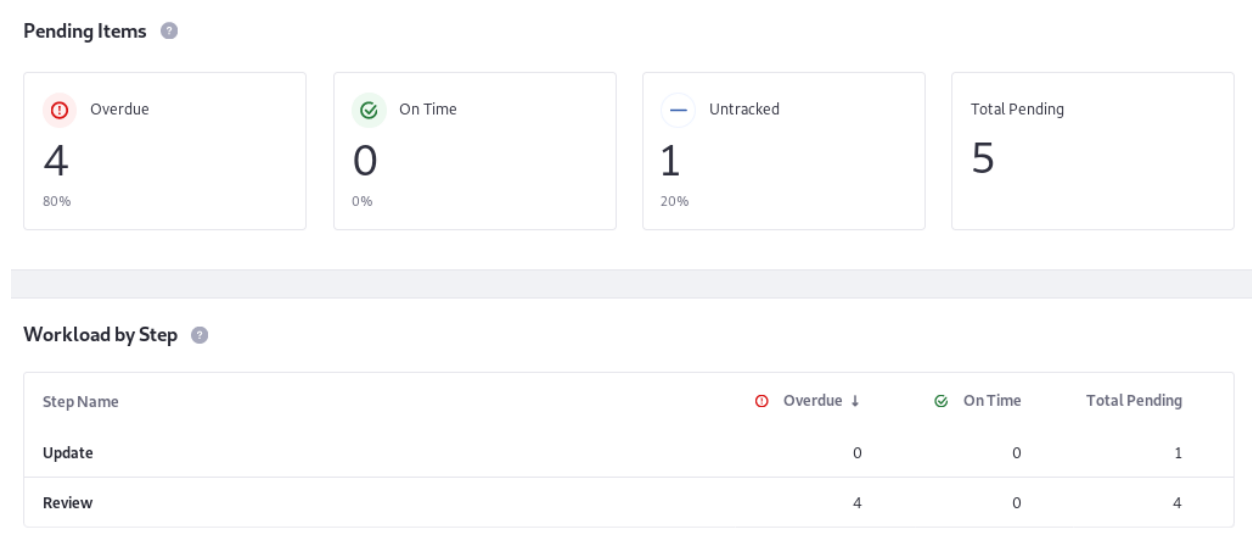


Figure 308.2: See data on the Pending Items and the Workload by Step for a process.

308.2 Pending Items

Pending Items shows you the overview of items by their SLA status. Drill down by clicking on any of the statuses to see the specific items that are enumerated in Pending Items.

308.3 Workload by Step

Workload by Step shows a breakdown of the items that are in each step of the workflow process, by their SLA status (Overdue or On Time).

308.4 Completed Items

Click the *Completed* tab (*Performance* on SP1) to see the items that have completed the workflow process. Workload by Step data doesn't make sense in this case, because by definition, these items are no longer in any workflow process step. Instead, there's a *Completion Velocity* chart that shows the performance over time.

308.5 Completion Velocity

A line chart shows you the completion rate for the workflow process. The default display shows the number of completed workflow instances per day, for the last 30 days.

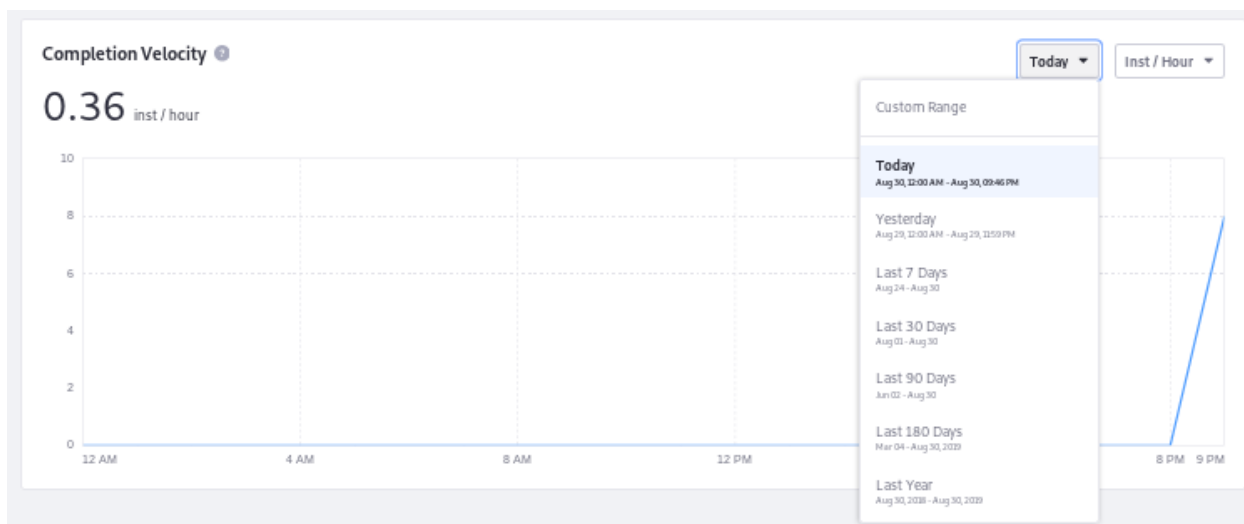


Figure 308.3: View the completion rate of items in a workflow process over time.

The overall completion rate for the time period is displayed in the top right corner of the chart (as *Inst/timeUnit*), while the trend-line is presented in the chart body. The overall performance metric and the chart body are updated when you select a new time period; the time unit changes depending on the total time period you're measuring. For some of the longer durations, the time unit is configurable:

Today Calculate *Inst/Hour* from 00:00, or 12:00 AM, of the current day until the current time (rounded to the nearest whole hour).

Yesterday Calculate *Inst/Hour* From 00:00-23:59, or 12:00 AM to 11:59 PM, of the previous day.

Last 7 Days Calculate *Inst/Day*. The current day counts as 1 day, so this is from 6 days ago to the current day.

Last 30 Days Calculate the *Inst/Week* or the *Inst/Day*. The current day counts as 1 day, so this is from 29 days ago to the current day.

Last 90 Days Calculate the *Inst/Month*, *Inst/Week*, or *Inst/Day*. The current day counts as 1 day, so this is from 89 days ago to the current day.

Last 180 Days Calculate the *Inst/Month* or *Inst/Week*. The current day counts as 1 day, so this is from 179 days ago to the current day.

Last Year Calculate the *Inst/Month* or *Inst/Week*. The current day counts as 1 day, so this is from 364 days (365 for a leap year) ago to the current day.

308.6 Items View

Hover over the status you're interested in, from either the *Pending* or the *Completed* tabs (on SP1, these tabs were renamed to *Dashboard* and *Completion*). Click into the All Items screen from the overview report and a more detailed table appears, including the following columns:

ID This is the workflow item's numeric identifier to the system. Importantly, you can click this to enter the Item Detail view.

Item Subject This shows a human readable summary of the item, to help administrators identify the item.

Process Step This identifies where the item is in the workflow.

Created By This shows the user name of the submitting User.

Creation Date This shows the date and time the item was submitted to the workflow.

The All Items view can be filtered so you can find the subset of items you want to analyze.

308.7 Filtering by SLA Status

Filter items based on whether they're Overdue, On Time, or Untracked.

Overdue Overdue items have breached at least one SLAs defined deadline.

On Time On Time items have not breached *any* SLA deadline.

Untracked Untracked items are items in the workflow process that aren't currently under the purview of an SLA. They can be in a task identified as a *Pause* in the SLA, or perhaps outside the scope of the SLA entirely, if the SLA isn't defined for the entire process (Process Begins to Process Ends in the SLA Definition screen).

308.8 Filtering by Process Status and Completion Period

Filter items based on whether they're Pending or Completed in the workflow process.

Filter By **SLA Status** Process Status Process Step

9 Results Overdue On Time Untracked Clear All

| | | Process Step | Created By | Creation Date | |
|-------------------------------------|-------|---|------------|---------------|-----------------|
| <input checked="" type="checkbox"/> | 40304 | Blogs Entry: Capitalize on Content | Review | Test Test | May 09, 6:37 PM |
| <input checked="" type="checkbox"/> | 40323 | Blogs Entry: Thinking Outside the Box: Collaborate Cross-Functionally | Review | Test Test | May 09, 6:37 PM |
| <input checked="" type="checkbox"/> | 40342 | Blogs Entry: Engineer Greater Profits in a Bull Market | Review | Test Test | May 09, 6:37 PM |
| <input checked="" type="checkbox"/> | 40362 | Blogs Entry: Synergistic Corporate Brainwashing | Review | Test Test | May 09, 6:38 PM |
| <input checked="" type="checkbox"/> | 40381 | Blogs Entry: Leveraging Slow Time to Market | Review | Test Test | May 09, 6:38 PM |
| <input checked="" type="checkbox"/> | 40400 | Blogs Entry: Secret Success Strategy: Over-delivering and Under-promising | Review | Test Test | May 09, 6:38 PM |
| <input checked="" type="checkbox"/> | 40420 | Blogs Entry: Keep Business Humming with Shiny New Office Toys | Review | Test Test | May 09, 6:39 PM |
| <input checked="" type="checkbox"/> | 40439 | Blogs Entry: Don't Just Make Money: Make a Difference | Review | Test Test | May 09, 6:39 PM |
| <input checked="" type="checkbox"/> | 40458 | Blogs Entry: Capitalize on you Competitors' Awesomeness | Review | Test Test | May 09, 6:39 PM |

Figure 308.4: Filter by SLA status: Overdue, On Time, or Untracked.

If you filter by the Completed status, you'll get an additional filtering option: filter items by the Completion Period. Select from these time periods:

- Today
- Yesterday
- Last 7 Days
- Last 30 Days (default)
- Last 90 Days
- Last 180 Days
- Last Year
- All Time

308.9 Filtering by Process Step

Filter items based on where they are in the workflow definition. For example, in the Single Approver workflow process, you can choose to see a report including all items in the Review task. This is different for each workflow definition.

Filter By SLA Status Process Status Completion Period Process Step

2 Results for SLA Status: Overdue Process Status

| ID | Item Subject | Process Step | Created By |
|-------|--|--------------|------------|
| 38512 | Blogs Entry: Capitalize c | Completed | Test Test |
| 38534 | Blogs Entry: Thinking O ss-Functionally | Completed | Test Test |

Completion Period options:

- Today: 20 May, 12 AM - 20 May, 08 PM
- Yesterday: 19 May, 12 AM - 19 May, 11 PM
- Last 7 Days: 14 May - 20 May
- Last 30 Days: 21 Apr - 20 May
- Last 90 Days: 20 Feb - 20 May
- Last 180 Days: 22 Nov, 2018 - 20 May, 2019
- Last Year: 20 May, 2018 - 20 May, 2019
- All Time

Figure 308.5: Filter by Process Status and Completion Period.

308.10 Combining Filters

Use a combination of filters to find just the items you need to see. For example, below are all items in the Single Approver process's Review task that have the status Completed or Pending, whether On time or Overdue. Untracked items aren't shown.

308.11 Item Details

To see the metrics for a single workflow process item, click the ID field while in the All Items view. A pop-up shows you more detailed information on the item.

From here you can view detailed information about the asset and even click *Go to Submission Page*, which redirects you to the item's view in the Submissions section of the Control Panel.










The top of the Item Detail view is important. It shows you the information about the due date for the item in the SLA, and its SLA completion status: *Open* or *Resolved*.

Open The defined SLA goals are not yet met. Open items can be of status Overdue or On Time.

Resolved The defined SLA goals are completed. Resolved items can be of status Overdue or On Time.

Filter By SLA Status ▾ Process Status ▾ Completion Period ▾ Process Step ▾


9 Results for SLA Status: Overdue X SLA Status: On Time X Process Status: Pending X Process Status: Completed X Process Step: Review X Clear All

| ID | Item Subject | Process Step | Created By | Creation Date |
|---|---|--------------|------------|-----------------|
|  38512 | Blogs Entry: Capitalize on Content | Review | Test Test | May 20, 2:17 PM |
|  38534 | Blogs Entry: Thinking Outside the Box: Collaborate Cross-Functionally | Review | Test Test | May 20, 2:17 PM |
|  38555 | Blogs Entry: Engineer Greater Profits in a Bull Market | Review | Test Test | May 20, 2:18 PM |
|  38575 | Blogs Entry: Synergistic Corporate Brainwashing | Review | Test Test | May 20, 2:18 PM |
|  38595 | Blogs Entry: Leveraging Slow Time to Market | Review | Test Test | May 20, 2:18 PM |
|  38616 | Blogs Entry: Secret Success Strategy: Over-delivering and Under-promising | Review | Test Test | May 20, 2:19 PM |
|  38636 | Blogs Entry: Keep Business Humming with Shiny New Office Toys | Review | Test Test | May 20, 2:19 PM |
|  38657 | Blogs Entry: Don't Just Make Money: Make a Difference | Review | Test Test | May 20, 2:20 PM |
|  38677 | Blogs Entry: Capitalize on you Competitors' Awesomeness | Review | Test Test | May 20, 2:20 PM |

20 Entries ▾ Showing 1 to 9 of 9 entries. < 1 >


Figure 308.6: Combine filters to see just the items you want.

From the overall metrics of a workflow process down to the details on a single item in the workflow, the new Workflow Metrics functionality gives you insights into the time it takes to *get things done* in Liferay DXP.


 Item #38595 ✕

Due Date by SLA

OPEN (1)

 Total time in workflow May 20, 2:23 PM (4d 01h 04min Overdue)

Process Details

| | |
|----------------|--|
| Process Status | Pending |
| Created By | Test Test |
| Creation Date | May 20, 2:18 PM |
| Asset Type | Blogs Entry |
| Asset Title | Leveraging Slow Time to Market  |
| Current Step | Review |

[Go to Submission Page !\[\]\(cc550a6475ae0a4b1b30a2c3ecc0f2f6_img.jpg\)](#)

Figure 308.7: Item Details include SLA status information and whether the item is Resolved or Open.

WORKFLOW DESIGNER

With the proper permissions, users can publish assets. Even if your enterprise has the greatest employees in the world, many of the items they want to publish must still be reviewed, for a variety of reasons. The Workflow Designer lets you design workflow definitions so your assets go through a review process before publication.

With the Workflow Designer, you develop workflow definitions using a convenient drag and drop user interface. You don't need to be familiar with writing XML definitions by hand. Some of the features can be enhanced, however, if you're familiar with Groovy, a supported Java-based scripting language. All that is to say, don't be scared off when you come to a block of code in these articles. Just decide if you need the feature and find someone familiar with Java or Groovy to help you out.

Note: By default, only one workflow definition is installed: the Single Approver Workflow definition. What you might not know is that you have access to several others too. Look in `[Liferay_Home]/osgi/marketplace/` and find the Liferay Forms and Workflow - Liferay Portal Workflow - Impl.lpkg. Open it, find the `com.liferay.portal.workflow.kaleo.runtime.impl-[version].jar`, and look in its `META-INF/definitions` folder. You'll see the following workflow definitions:

```
category-specific-definition.xml
legal-marketing-definition.xml
single-approver-definition.xml
single-approver-definition-scripted-assignment.xml
```

To work with any of these definitions in the Workflow Designer, extract them from the JAR file first. Once you have the XML files locally,

1. Add a new workflow. Go to Control Panel → Workflow → Process Builder, and click the Add button (+).
2. Go to the Source tab.
3. Click *import a file* and upload the XML file.
4. Name the definition appropriately, and click either *Save* (to save it as a draft) or *Publish* (see below for more information on saving and publishing).

Now you can begin exploring or modifying the definition.

It's time to start exploring the Workflow Designer and its features.

MANAGING WORKFLOWS WITH THE WORKFLOW DESIGNER


The Workflow Designer gives you an intuitive interface for creating workflow definitions, from the simplest approval processes to the most complex business processes you can imagine. It features a drag and drop interface, workflow definition versioning, and a graphical representation of definitions and their nodes. Without the Workflow Designer, you'd have to hand-craft your entire workflow definition in XML. With the Workflow Designer, you might never need to look at a single line of XML. Of course, the Workflow Designer also lets you directly manipulate the XML (using the *Source* tab) if you find it convenient.

310.1 Adding New Workflow Definitions with the Workflow Designer

Access the Workflow Designer by going to the Control Panel → Workflow → Process Builder. Click the Add icon ().

Give the workflow definition a title and you're ready to start designing your workflow.

310.2 Saving and Publishing Workflow Definitions

First, look below the canvas to see the buttons that let you *Save* or *Publish*. Saving the definition as a draft lets you save your work so it's not lost (due to a timeout, for example). It won't be published (and assignable to assets), and it won't be considered a version until the Publish button is clicked. Each time you save the workflow as a draft, a new revision is added to the Revision history. To see the Revision history and manage workflow versions, open the Info sidebar () and click *Revision*

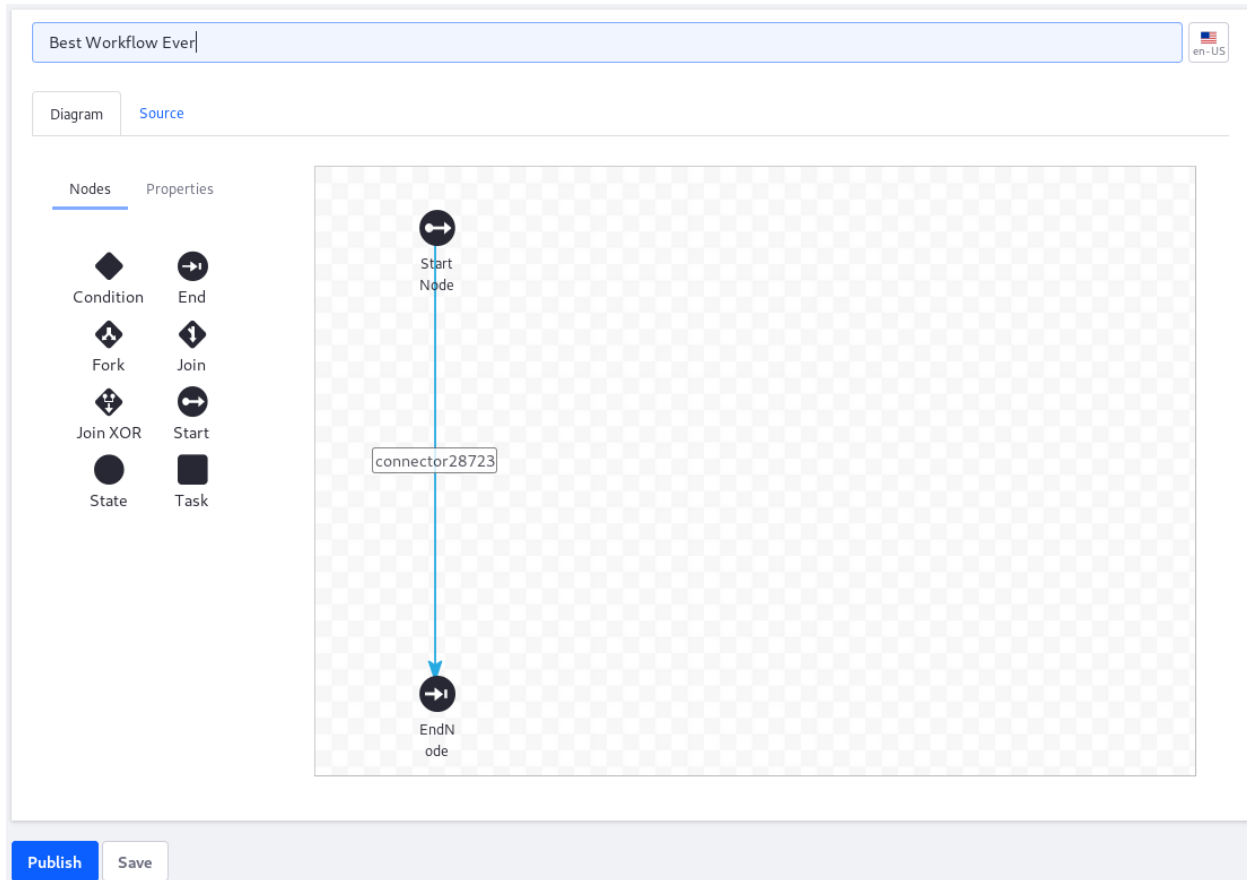


Figure 310.1: The Workflow Designer’s graphical interface makes designing workflows intuitive.

History.

310.3 Adding Nodes

A new workflow is already populated with a start node, an end node, and a transition between them. To make the workflow the way you want it, add nodes to the workflow.

1. Drag a node from the *Nodes* palette on the right of the designer and *drop* it on the canvas.
2. You’ll see it’s not connected to other nodes by a transition, so right now it can’t be used in the workflow. Delete the existing transition and then you can make new transitions to direct the *flow* of your workflow (see more about transitions below if you’re not sure what they’re for or how to use them in the Workflow Designer).

Alternatively, start by deleting the default transition, then click the edge of the start node, drag a new transition from the start node to a blank spot on the canvas, and release it. You’re prompted

| Title | Description | Last Modified | |
|-------------------------------------|---|----------------------------|---|
| PUBLISHED | | | |
| Single Approver | A single approver can approve a workflow content. | 29 Minutes Ago | ⋮ |
| Best Workflow Ever | | 1 Second ago by Test Test | ⋮ |
| NOT PUBLISHED | | | |
| single approver scripted assignment | A single approver can approve a workflow content. | 2 Minutes ago by Test Test | ⋮ |

Figure 310.2: View a list of the current workflows that can be edited in the Workflow Designer.

to create a node at that spot, because you can't have a transition without a starting point and an ending point on a node.



Figure 310.3: You can add a node by creating a transition that ends at a blank spot on your Designer canvas.

That's it. Of course, if you drag, say, a *Task* node onto the canvas, it must be configured.

310.4 Node Settings

Now you know how to add nodes to the workflow definition. By default you have three things added to your canvas: a start node, a transition, and an end node. Think of the *EndNode* as the point in

the workflow where an asset reaches the *Approved* status. The *StartNode* is where the asset goes from the *Draft* status to *Pending*. You might decide to name your nodes to reflect what's happening in each one. To name a node, double click it, and its *Settings* appear. Then double click the value of the *Name* property and you can edit the name. Click *Save* when you're done.

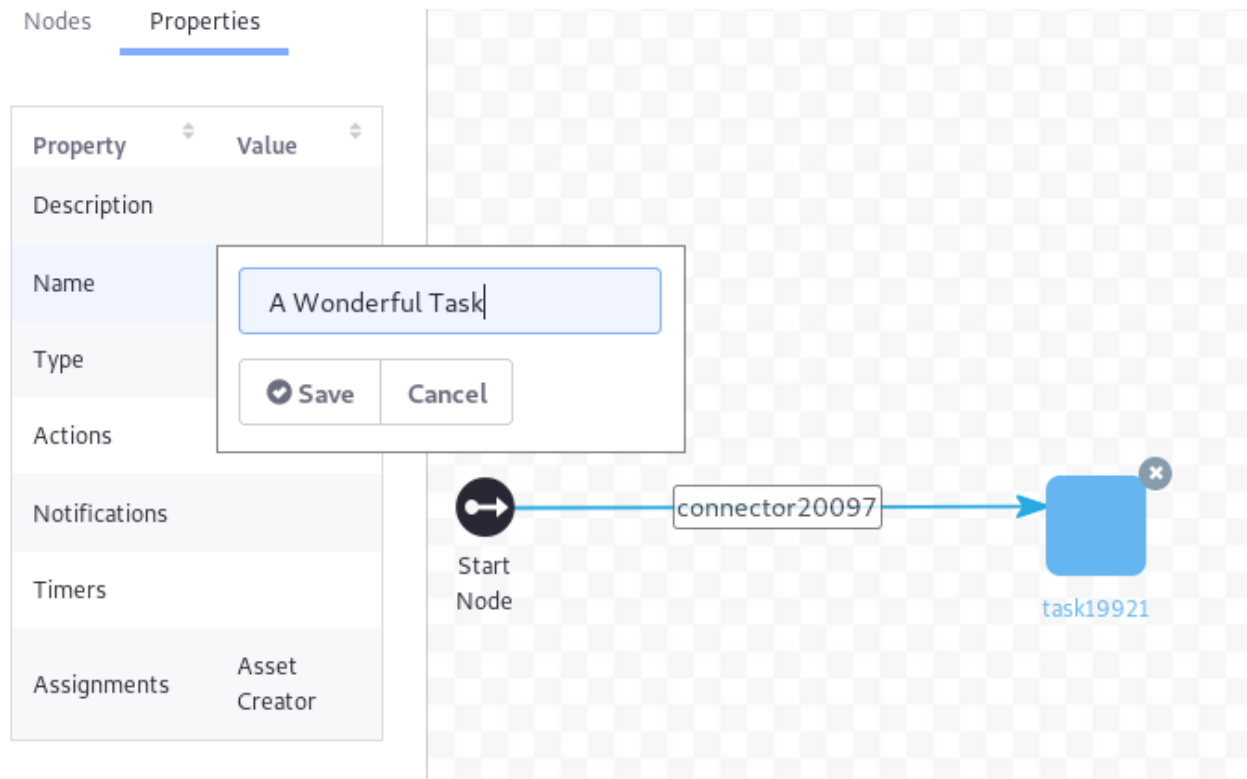


Figure 310.4: You can edit a node's settings.

Of course, there's more you can do besides changing node names. Actions, Notifications, and Assignments can be used to make your workflow definition useful and interactive. Keep reading to learn about these features.

310.5 Related Topics

Kaleo Forms

Activating Workflow

Liferay's Workflow Framework

Dynamic Data Lists

WORKFLOW DEFINITION NODES

Once you know the basics of creating workflow definitions with the workflow Designer, you can get into the details. In this tutorial you'll learn about Actions and Notifications, two important features your workflow nodes can use. You'll also learn how to affect the processing of the workflow using Transitions, Forks, Joins, and Conditions.

There are several node types you can use in workflow definitions:

- Task nodes
- Fork and Join nodes
- Condition nodes
- Start nodes
- End nodes
- State nodes

Because they're the most complex node, and often the meat of your workflow definitions, Task nodes are covered separately.

Fork, Join, and Condition nodes are discussed, along with Transitions, in a tutorial on workflow processing, since they're used for affecting the processing of the workflow.

This tutorial discusses State nodes, Start nodes, and End nodes, along with Actions and Notifications.

311.1 Node Actions and Notifications

Any node can have Actions and Notifications.

311.2 Actions

Actions do additional processing before entering the node, after exiting a node, or once a task node is assigned. They're configured by accessing a node's Properties tab, then double clicking *Actions*.

The screenshot shows a configuration form for an Action. It has the following sections:

- Name:** A text input field containing "The Action".
- Description:** A text area containing "This is an Action that will run at during the task's execution (exactly when is configured below)".
- Script:** A text area containing "Define what will happen here".
- Script Language:** A dropdown menu currently set to "Groovy".
- Execution Type:** A dropdown menu currently set to "On Assignment".
- Priority:** An empty text input field.
- Buttons:** Three buttons at the bottom: "Save" (with a checkmark icon), "Cancel", and "Add Section" (with a plus icon).

Figure 311.1: You can add an Action to a Task node.

The Single Approver workflow contains an Update task with an action written in Groovy that sets the status of the asset as *denied*, then sets it to *pending*.

```
import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("denied"), workflowContext);
WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("pending"), workflowContext);
```

Why would the action script first set the status to one thing and then to another like that? Because for some assets, the *denied* status sends the asset creator an email notification that the item has been denied.

The end node in your workflow definition has an action configured on it by default, on entry to the end node:

```
import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("approved"), workflowContext);
```

This is a Groovy script that updates the status to *approved*, since that's usually the goal of a workflow process.

You can do something simple like the actions above, or you can be as creative as you'd like.

It's good to assign a task to a user, and it's even more useful if the user is notified about workflow tasks.

311.3 Notifications

Notifications are sent to tell task assignees that the workflow needs attention or to update asset creators on the status of the process. They can be sent for tasks or any other type of node in the workflow. To set up notifications, double click on *Notifications* in a node's Properties tab and create a notification.

You must specify the Notification Type, and you can choose User Notification or Email. You can use Freemarker if you need a template, or you can choose to write a plain text message.

Note: Instant Messenger and Private Message also appear as Notification Type options, but these are non-functional and will be removed in a future version.

Here's a basic Freemarker template that reports the name of the asset creator and the type of asset in the notification:

```
`${userName}` sent you a `${entryType}` for review in the workflow.
```

You can also choose to link the sending of the notification to entry into the node (On Entry), when a task is assigned (On Assignment), or when the workflow processing is leaving a node (On Exit). You can configure multiple notifications on a node.

Commonly, the assignment and notification settings are teamed up so a user receives a notification when assigned a task in the workflow. To do this, choose *Task Assignees* under Recipient Type when configuring the notification.

Note: The *from name* and *from address* of an email notification are configurable via portal properties. Place these settings into a `portal-ext.properties` file, in your Liferay Home folder. Then restart the server:

```
workflow.email.from.name=  
workflow.email.from.address=
```

These can also be set programmatically into the `WorkflowContext`, and the programmatic setting always takes precedence over the system scoped portal property.

311.4 Start and End Nodes

Name ✕

The Notification

Description

Send a notification to users.

Template Language

Text

Template

Enter the Notification here.

Notification Type

- Email
- Instant Messenger
- Private Message
- User Notification**

Execution Type

On Assignment

Recipient Type

Asset Creator

+ Add Section

✓ Save Cancel + Add Section

Figure 311.2: You can send a Notification from a Task node.

Start and end nodes kick off the workflow processing and bring the asset to its final, approved state. Often you can use the default start and end nodes without modification. If you want to do some more processing (in the case of a start node), add an action to the node using the Properties tab as described in the section on Actions above.

End nodes have a default action that sets the workflow status to Approved using the Groovy scripting language:

```
import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("approved"), workflowContext);
```

Add more to the action script if you need to do additional processing.

By default, there's a transition connecting the start node and end node, but you'll probably want to delete it, since most workflows don't proceed straight from the initial state to approved.

311.5 State Nodes

State nodes can have Notifications and Actions. The default end node added by workflow Designer is a pre-configured state node that sets the workflow status to Approved. Perhaps you want to create a node that sets the status to *Expired*. You could create a state node for it by dragging one onto your workflow Designer canvas, then configuring an action in it that sets the status to Expired. Here's what it would look like in Groovy:

```
import com.liferay.portal.kernel.workflow.WorkflowStatusManagerUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

WorkflowStatusManagerUtil.updateStatus(WorkflowConstants.getLabelStatus("expired"), workflowContext);
```

Next, you'll learn to do parallel processing using fork and join nodes.

311.6 Related Topics

Kaleo Forms

Activating Workflow

Liferay's Workflow Framework

Dynamic Data Lists

AFFECTING THE PROCESSING OF WORKFLOW DEFINITIONS

Workflow definitions all contain nodes: a Start Node, and End node, and at least one Task node. You might know that for the workflow to progress from one node to the other, you need Transitions. In this tutorial you'll learn about using transitions to move the asset through the workflow from node to node, but you'll also learn about some other features that affect the processing of the workflow.

- Transitions
- Forks
- Joins
- Conditions

Start by learning about the ever important Transition.

312.1 Transitions

What are transitions? Workflow transitions connect one node to another. On exiting the first node, processing continues to the node pointed to by the transition. Every time you create an arrow from one node to another, Workflow Designer creates a transition.

Each node you add has a pop-up menu letting you delete the node. As you hover your mouse over the edges of a node, notice your mouse pointer changes to a cross. The cross indicates you can connect the current node to another node. Hold down your mouse button and drag the mouse to start drawing your transition to another node. If you stop before reaching the edge of the next node, a pop-up displays node types you can create and connect to on-the-fly. To connect with an existing node, continue dragging the connector to that node.

When developing workflows in the Workflow Designer, make sure you go through all the transitions and name them appropriately. By default, these transitions get system generated names, so rename them to something more human-readable, as they're displayed to workflow users as links that send the item to the next step in the workflow.

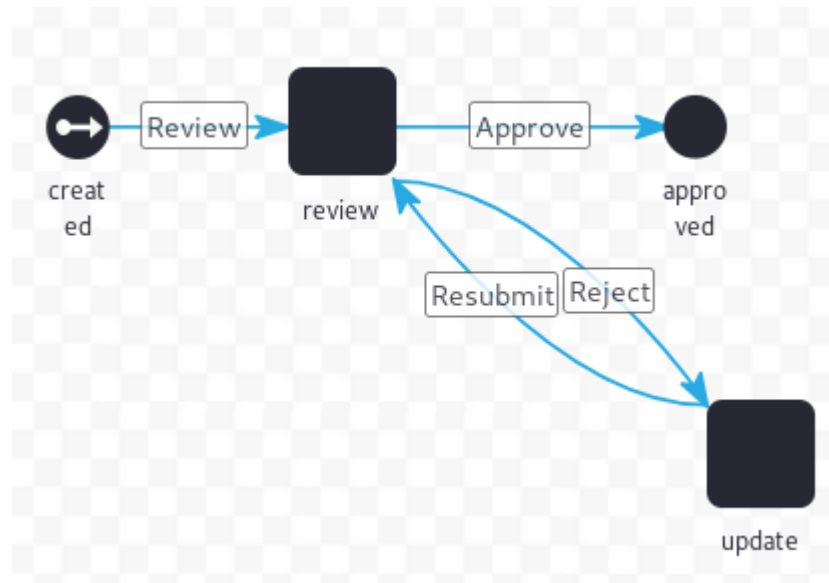



Figure 312.1: You connect nodes and direct workflow processing with transitions. The Single Approver workflow has transitions named Submit, Resubmit, Reject, and Approve.

Assigned to
 Test Test

State
 Review

Create Date
 6/17/18 2:13 AM

Due Date
 Never


Approve




Reject


Assign to...

Update Due Date

Preview of Blogs Entry

 en-us

 Is a Space Vacation Right for You?  

 Test Test
 1 Minute Ago

The moon is beautiful this time of year, but let's face it. Space is harsh...

Figure 312.2: In the Single Approver workflow, a user in the Review task can choose to Approve or Reject the asset, which sends the asset either to the EndNode or to the Update task.

To rename transitions, click on the arrow representing the transition and use the Properties tab to set the name just like you do for a node.

312.2 Forks and Joins

Sometimes you don't need to wait for one task to be completed before moving on to another one. Instead, you want to do two or more things at the same time. To do this, transition to a fork node, make two transitions from the fork to your parallel tasks, and then come back together using a join node.




Figure 312.3: Forks and Joins are used to enable parallel processing in the workflow.

With a regular Join node, for the workflow to proceed beyond the join, the transition from both parallel executions must be invoked. However, if you use a Join XOR node instead, the workflow proceeds as long as the transition from one of the parallel executions is invoked.

Keep in mind that you must balance your fork and join nodes. In other words, for every fork, there must be a join that brings the parallel workflow threads back together.

312.3 Conditions

Sometimes you must inspect an asset or its execution context, and depending on the result, send it to the appropriate transition. You need a node for a script that concludes by setting a value to one of your transitions.

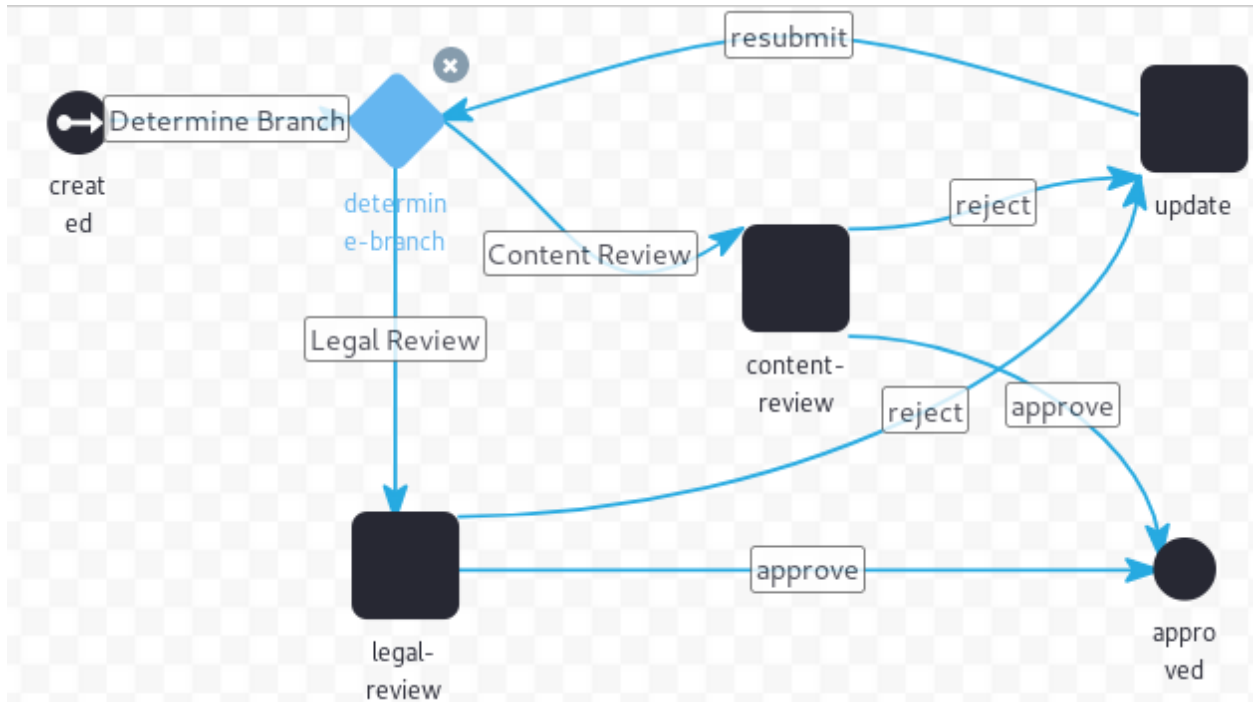


Figure 312.4: The Category Specific Approval definition starts with a Condition node.

From the *Category Specific Approval* (`category-specific-definition.xml`), this is the script in the condition node that starts the workflow (coming directly from the start node):

```

import com.liferay.asset.kernel.model.AssetCategory;
import com.liferay.asset.kernel.model.AssetEntry;
import com.liferay.asset.kernel.model.AssetRenderer;
import com.liferay.asset.kernel.model.AssetRendererFactory;
import com.liferay.asset.kernel.service.AssetEntryLocalServiceUtil;
import com.liferay.portal.kernel.util.GetterUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;
import com.liferay.portal.kernel.workflow.WorkflowHandler;
import com.liferay.portal.kernel.workflow.WorkflowHandlerRegistryUtil;

import java.util.List;

String className = (String)workflowContext.get(WorkflowConstants.CONTEXT_ENTRY_CLASS_NAME);

WorkflowHandler workflowHandler = WorkflowHandlerRegistryUtil.getWorkflowHandler(className);

AssetRendererFactory assetRendererFactory = workflowHandler.getAssetRendererFactory();

long classPK = GetterUtil.getLong((String)workflowContext.get(WorkflowConstants.CONTEXT_ENTRY_CLASS_PK));

AssetRenderer assetRenderer = workflowHandler.getAssetRenderer(classPK);

AssetEntry assetEntry = assetRendererFactory.getAssetEntry(assetRendererFactory.getClassName(), assetRenderer.getClassPK());

```

```
List<AssetCategory> assetCategories = assetEntry.getCategories();

returnValue = "Content Review";

for (AssetCategory assetCategory : assetCategories) {
    String categoryName = assetCategory.getName();

    if (categoryName.equals("legal")) {
        returnValue = "Legal Review";
    }
    return;
}
```

This example checks the asset category to choose the processing path, whether to transition to the *Legal Review* task or the *Content Review* task.

You may be wondering what that `returnValue` variable is. It's the variable that points from the condition to a transition, and its value must match a valid transition in the workflow definition. This script looks up the asset in question, retrieves its asset category, and sets an initial `returnValue`. Then it checks to see if the asset has been marked with the *legal* category. If not it goes through *Content Review* (the content-review task in the workflow), and if it does it goes through *Legal Review* (the legal-review task in the workflow).

Now you're equipped with the basic knowledge to design beautiful, effective workflows so that your assets can be properly reviewed before they're published in your sites.

312.4 Related Topics

Kaleo Forms

Activating Workflow

Liferay's Workflow Framework

Dynamic Data Lists

CREATING TASKS IN THE WORKFLOW DESIGNER

Task nodes have several parts and are the most complex parts of a workflow definition. Unlike other workflow nodes, task nodes have Assignments, because a User is expected to *do something* (often approve or reject the submitted asset) when a workflow process enters the task node: the assignment specifies who that User is.

Commonly, task nodes contain Notifications, Assignments, and Actions (defined in scripts). See more about Notifications and Actions in the article on workflow nodes. Task nodes and their assignments are more complex and deserve their own article (this one).

To get started, drag and drop a task node on your workflow canvas if you haven't already. Open its Properties and give it a name. Then double click *Actions* in the task's Properties pane.

You can define a notification (often Task Assignee is appropriate), or write a Groovy script defining an action that's triggered for your task.

Next learn about creating Assignments for your task nodes.

313.1 Assignments

Workflow tasks must be completed by a User. You can choose how you want to configure your assignments.

You can choose to add assignments to specific Roles, multiple Roles of a Role Type (Organization, Site, or regular Role types), to the Asset Creator, to Resource Actions, or to specific Users. Additionally, you can write a script to define the assignment.

Assigning tasks to Roles, Organizations, or Asset Creators is a straightforward concept, but what does it mean to assign a workflow task to a Resource Action? Imagine an *UPDATE* resource action. If your workflow definition specifies the UPDATE action in an assignment, anyone who has permission to update the type of asset being processed in the workflow is assigned to the task. You can configure multiple assignments for a task.

Figure 313.1: You can add an Assignment to a Task node.

313.2 Resource Action Assignments

Resource actions are operations performed by Users on an application or entity. For example, a User might have permission to update Message Boards Messages. This is called an UPDATE resource action, because the User can update the resource. If you're still uncertain about what resource actions are, refer to the developer tutorial on the permission system for a more detailed explanation.

To find all the resource actions that have been configured, you need access to the Roles Admin application in the Control Panel (in other words, you need permission for the VIEW action on the Roles resource).

- Navigate to Control Panel → Users → Roles.
- Add a new Regular Role. See the article on managing roles for more information.
- Once the Role is added, navigate to the Define Permissions interface for the Role.
- Find the resource whose action you want to use for defining your workflow assignment.

How do you go from finding the resource action to using it in the workflow? Use the Workflow Designer's interface for setting up a resource action assignment.

When configuring your task node's Assignment, select Resource Actions as the Assignment Type, then specify the Resource Actions to use for the assignment (for example, UPDATE).

Here's what the assignment looks like in the Source (Workflow XML) tab:

```
<assignments>
  <resource-actions>
    <resource-action>UPDATE</resource-action>
  </resource-actions>
</assignments>
```

As usual, assign the workflow to the appropriate workflow enabled asset.

Now when the workflow proceeds to the task with the resource action assignment, Users with UPDATE permission on the resource (for example, Message Boards Messages) are notified of the task

Figure 313.2: Configure resource action assignments in the Workflow Designer.

and can assign it to themselves (if the notification is set to Task Assignees). Specifically, Users see the tasks in their *My Workflow Tasks* application under the tab *Assigned to My Roles*.

Use all upper case letters for resource action names. Here are some common resource actions:

```
UPDATE
ADD
DELETE
VIEW
PERMISSIONS
SUBSCRIBE
ADD_DISCUSSION
```

You can determine the probable resource action name from the permissions screen for that resource. For example, in Message Boards, one of the permissions displayed on that screen is *Add Discussion*. Convert that to all uppercase and replace the space with an underscore, and you have the action name.

313.3 Scripted Assignments

You can also use a script to manage the assignment. Here's the script for the Review task assignment in the Scripted Single Approver workflow definition (`single-approver-definition-scripted-assignment.xml`):

```
import com.liferay.portal.kernel.model.Group;
import com.liferay.portal.kernel.model.Role;
import com.liferay.portal.kernel.service.GroupLocalServiceUtil;
import com.liferay.portal.kernel.service.RoleLocalServiceUtil;
import com.liferay.portal.kernel.util.GetterUtil;
import com.liferay.portal.kernel.workflow.WorkflowConstants;

long companyId = GetterUtil.getLong((String)workflowContext.get(WorkflowConstants.CONTEXT_COMPANY_ID));

long groupId = GetterUtil.getLong((String)workflowContext.get(WorkflowConstants.CONTEXT_GROUP_ID));

Group group = GroupLocalServiceUtil.getGroup(groupId);

roles = new ArrayList<Role>();
```

```
Role adminRole = RoleLocalServiceUtil.getRole(companyId, "Administrator");

roles.add(adminRole);

if (group.isOrganization()) {
    Role role = RoleLocalServiceUtil.getRole(companyId, "Organization Content Reviewer");

    roles.add(role);
}
else {
    Role role = RoleLocalServiceUtil.getRole(companyId, "Site Content Reviewer");

    roles.add(role);
}

user = null;
```

Don't let all that code intimidate you. It's just assigning the task to the *Administrator* Role, then checking whether the *group* of the asset is an Organization and assigning it to the *Organization Content Reviewer* Role if it is. If it's not, it's assigning the task to the *Site Content Reviewer* Role.

Note the `roles = new ArrayList<Role>();` line above. In a scripted assignment, the `roles` variable is where you specify any Roles the task is assigned to. For example, when `roles.add(adminRole);` is called, the Administrator role is added to the assignment.

313.4 Related Topics

Kaleo Forms

- Activating Workflow

- Liferay's Workflow Framework



- Dynamic Data Lists

KALEO FORMS

Business processes are often form-based and workflow-driven. They start with entered data and progress by sending that data to other people or groups. Then it's processed in some way (for example, further data is entered or approval is granted), and the process moves on until completion, when each interested party has seen and manipulated the data. To write an app for each of these processes is laborious. It's much better to have a tool for quickly defining a process to suit each use case. The process architect must define both the data that gets collected and the process the data moves through to reach its final state. To accomplish this, Liferay DXP already includes the Dynamic Data Lists app for defining forms, and the Workflow Designer app for designing workflows. The Kaleo Forms solution combines the features of these apps, letting you use a single UI to design an integrated process for sending forms through a workflow.

314.1 Creating Kaleo Forms Process

To start creating a Kaleo Forms Process you need to get to Kaleo Forms Admin:

1. Go to *Site Administration* (your Site's menu) → *Content & Data* → *Kaleo Forms Admin*. The Kaleo Forms app appears with a list of any defined processes.
2. Click the *Add* button () to open the New Process Wizard.
3. Give the process a name and a description and click *Next*.
4. Define the fields that can appear in your process's forms. There are two ways to do this:
 - Use an existing field set. Click the field set's Actions button () and select *Choose*.
 - Create a new field set/data definition. Click the *Add Field Set* button. If you need help with this, see the documentation on creating data definitions. After creating the field set, select it as you would an existing field set.

Click *Next* to move to the wizard's next step.

Details Fields Workflow Forms

1 2 3 4

Details
Please type a name for your process and a description of what it does.

Name *
Spa Order Process en-US

Description
Process for completing customer orders at the Lunar Resort Spa en-US

Cancel Next

Figure 314.1: Add a Kaleo Forms Process to link a form with a workflow definition.

Details Fields Workflow Forms

1 2 3 4


Fields
Please select or create a new field set containing all the fields that will be used by your forms.

Selected Field Set:


Add Field Set

| Name | Description | Modified Date | |
|------------------------------------|-------------|----------------|---|
| Meeting Minutes (View Fields) | | 2 Seconds Ago | ⋮ |
| To Do (View Fields) | | 14 Seconds Ago | ⋮ |
| Spa Order--Field Set (View Fields) | | 46 Seconds Ago | ⋮ |

Figure 314.2: Define and choose your form’s fields.

5. Select a workflow to use for your forms. To do this, click the workflow's *Actions* button () and select *Choose*.

You can also edit an existing workflow or create a new one:

- To edit a workflow, click its *Actions* button () and select *Edit*.
- To begin creating a new workflow, click the *Add Workflow* button.

In either case, you use the same UI to edit/create the workflow. This UI is called Workflow Designer. It lets you create your workflow graphically instead of via code.

Once you select a workflow to use with your forms, click *Next*.

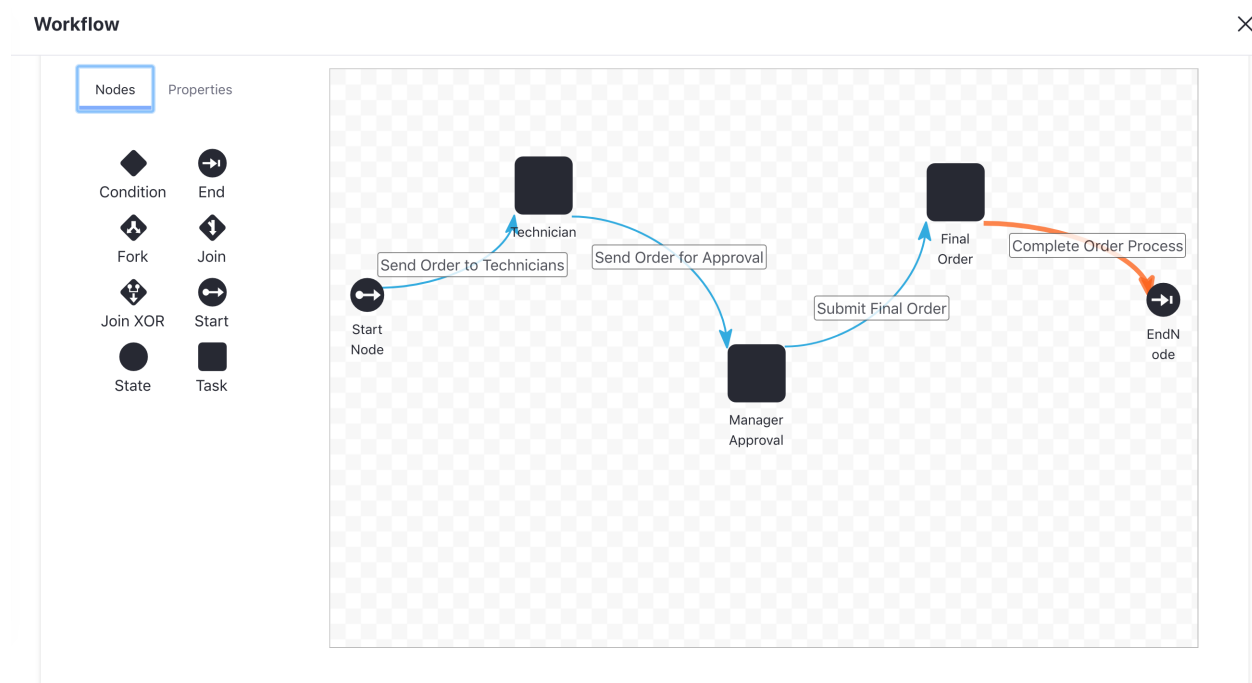




Figure 314.3: This example workflow has three tasks that happen sequentially.

6. Select or create a form to use for each workflow task. To do this, click each task's *Actions* button () and select *Assign Form*. On the screen that appears, select an existing form or click the *Add* button () and to create one.

Click *Save* when you're finished. Your process is done and appears in Kaleo Forms Admin's table.

314.2 Adding Records to a Process

Forms
Please select or create one form for each Workflow Task. Each form is a subset of the field set defined in step 2.

| Task | Form |
|------------------|--|
| StartNode | Spa Order Process--Order Form |
| Technician | Sp Order Process--Technician Claim Form |
| Manager Approval | Sp Order Process--Managerial Approval Form |
| Final Order | Sp Order Process--Final Order Form |

Previous Cancel Save

Figure 314.4: Assign a form to each task in the workflow, and for the initial state.

To add records to a process, click it in Kaleo Forms Admin and then click the *Add* button (+). This brings up the form you assigned to the workflow's initial state. Fill it out and click *Save*. Once submitting the initial form, the workflow engine then takes over and moves through each task in the workflow. Whatever Users or Roles you assigned to the tasks receive notifications, and the task appears in the *Assigned to My Roles* section of the *My Workflow Tasks* app. A notification also appears in the *Notifications* app. Once in the task, the user views and approves the form or clicks the *Edit* button. At this point, the workflow task forms you created come into play. Each assigned user fills out the form, saves it, and sends it along in the workflow.

SEGMENTATION AND PERSONALIZATION

Liferay's Segmentation and Personalization shows the right content to the right people at the right time. It provides the tools you need to manage different audiences and dynamically provide personalized experiences for people using your site. For example, if you're creating a campaign to promote new financial service products, you need a way to display offers to customers who are likely to be interested in those offers. You don't want to display information on a basic free checking account for an "advanced" customer who carries a high balance across several types of accounts, but you do want to show that information to a visitor who entered the site through a landing page from a promotion at a local college. At the same time, you probably don't want to recommend options for optimizing retirement account contributions to the college student, but the other customer might be a great target for that campaign. By using data like user attributes or visitor interactions, you can dynamically target relevant content to your site's guests.

315.1 Defining Segments

The first part of the equation is defining the types of segments that you need. You can create Segments to capture every case. Segments are composed of different criteria. In the previous example you might have a segment for *Free Checking Account Prospects* that contains criteria based on user data, like customers that don't currently have an open checking account; or based on user behavior, like visitors who came to the site through specific channels. To learn more about Segmentation options, see the overview of the Segment editor, practice creating basic Segments, or create more complicated Segments with custom fields and session data.

315.2 Integrating Segments with Analytics Cloud

There are two different stories that User Segments can tell. One is targeting content to specific audiences that encourages engagement and positive user experiences. The other is defining groups of users and visitors to analyze their behavior and interactions with your site. To tell the second story, you must integrate with Analytics Cloud.

Analytics Cloud is a Liferay service that provides in-depth information on who uses your site and how they use it. Analytics Cloud is a key component to fully utilizing Segments and Personalization, since it enables you to see the full picture of how users and visitors on your site behave and interact with both standard and targeted content. You can learn more about this in [Using Analytics Cloud with User Segments](#).

315.3 Personalizing Experiences

The most important piece of the puzzle isn't defining groups or analyzing user behavior. It's the final step of using the data to provide users and site visitors with the best possible experience, and driving campaigns and content engagement. If you strategically create segments, you can then use that to enhance user experiences, and make sure that users see content targeted to them. Content Page Personalization and Content List Personalization are two key aspects of this.

315.4 Content Page Personalization

Content Page Personalization dynamically changes the page layout and content based on who is viewing the page. You can create *Experiences* for any Content Page which provide different text, images, widgets, and even different layouts based on the segment criteria of the user viewing the page. You can see a step by step demonstration of this in [Content Page Personalization](#).

315.5 Content Set Personalization

Content Sets organize and display content. Content Set Personalization provides dynamic selection of Content Sets based on User Segments. This means the Content Set which displays in a given context is determined by their segment criteria. For example, you could use a content list to display "featured" articles at the top of a page. Then you could create Segments containing users who should receive more specialized content, rather than the default. Those Segments would then see content personalized to their interest rather than the default. You can see a step by step demonstration of this in [Content Set Personalization](#).

THE SEGMENT EDITOR

Segmentation and Personalization in 7.0 also provides an editor for defining User Segments.

1. Go to *Site Administration* and select the site that you wish to create segments for.
2. Click *People* → *Segments*.
3. Click the *Add User Segment* button (+).

At the top of the editor you can set the name, view the current members of the segment as it is defined, and choose to *Save* the Segment or *Cancel* to discard changes.




Figure 316.1: The top portion of the Segment Editor has the segment name and its members.

On the right side of the page, there's a Properties menu with the following options:

- User
- Organization
- Session

In addition to the various properties, there are operations and conjunctions that you use to define criteria.

316.1 User Properties

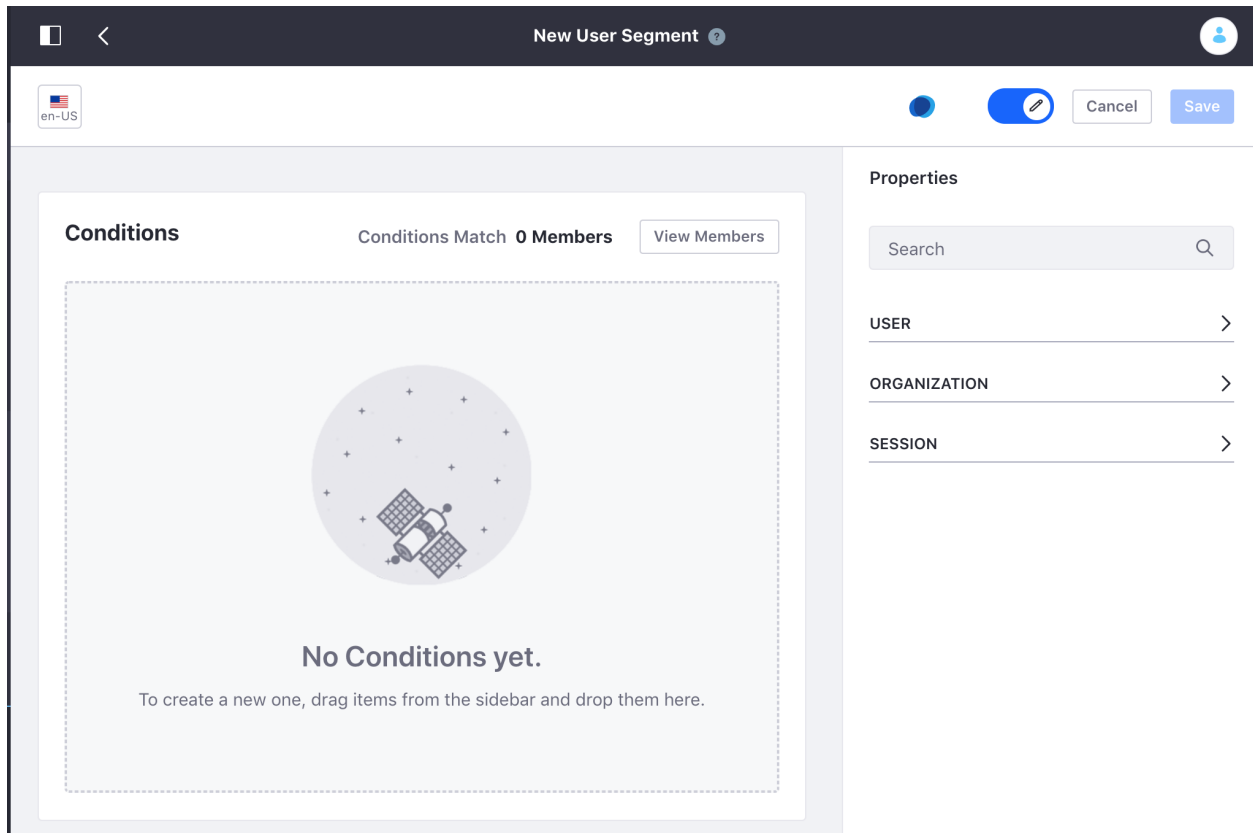


Figure 316.2: You use the Segment Editor to create new Segments.

User Properties are user attributes you want to capture. This is made up of user metadata as defined in their accounts, but also contains certain group memberships (like Roles and User Groups) as well as information like the date the user profile was last modified.

316.2 Organization Properties

Organization Properties is a selectable list of Organizations to include in your Segment. They contain similar criteria as the User selection, like *Name* and *Date Modified*.

316.3 Session Properties

Session Properties contains criteria based on the user's activity, browser, and system information. You can use this to target the user's device or OS, or for activity-based criteria like if a user entered the website through a specific campaign driven landing page.

316.4 Operations and Conjunctions

There are a number of different ways to evaluate properties, and different ways that you can relate different fields. For a comprehensive list, see the Defining Segment Criteria Reference.

CREATING USER SEGMENTS


To learn how to use segmentation, you'll step through an example for defining two segments for a site. The segments use user data like the *Job Title* field and organization membership for evaluating segments. The first Segment you'll create, *American Engineers*, uses standard fields as criteria.

To get started, navigate to the Segments page.

1. Go to *Site Administration* for the Bank site.
2. Open the *People* section and select *Segments*.

317.1 Creating a Custom Segment

On the Segments page, you'll see a list of all the currently available segments, if they're available. Create a segment named *American Engineers* for employees of your company.

1. Click the *Add* button () button.
2. Click in the top text area and enter the name *American Engineers*.
3. Open the *User* tab in the right side *Properties* menu and drag the *Job Title* property to the *Conditions* area.
4. Click on the operator field and set it to *contains*.
5. Click on the text field and enter *Engineer*.
6. Open the *Organization* properties and drag over an *Organization* field.
7. Set the comparator between *User Properties* and *Organization Properties* to *And*.
8. Set the *Organization* field to be *equals* and select the organization.
9. Click *Save* to save your Segment.




Figure 317.1: Setting the comparator to *contains* includes variations of “Engineer” like “Software Engineer” in the segment.





Figure 317.2: You can prevent typos by directly selecting Organizations through the interface.

As you edit, a count of members meeting the criteria appears at the top of the page. You can click on *View Members* to see the list. This helps you determine if you are correctly defining the Segment.

317.2 Managing Segments

After you create your Segment, you can see it in the list of Segments on the *People → Segments* page. From there you can edit the segment, delete it, or change the permissions for it. You cannot delete a segment if it’s being used in an experience. Also, changing permissions only affects who has access to manage the Segment; it doesn’t change Segment membership or criteria.


1. Go to *People → Segments*.
2. Click the *Actions* button () next to the *American Engineers*.

American Engineers Members

| Name | Email Address |
|---------------|-----------------------|
| Bill Engineer | bengineer@company.com |
| Dave Engineer | dengineer@company.com |

Figure 317.3: You can view the list of Segment members at any time.

You can manage the options here. You can also click on the Segment's name to edit it.

| Name | Active | Source | Scope | Modified Date |
|---|--------|---|-----------|---------------|
| <input type="checkbox"/> American Engineers | Yes |  | Current S | |

- Edit
- Permissions
- Delete

Figure 317.4: You can edit, delete or manage permissions from the options menu.


Next, you'll define a Custom Field and use it as part of your Segmentation criteria.

CREATING SEGMENTS WITH CUSTOM FIELDS AND SESSION DATA

Now that you created a segment, you can take things to the next level and use a Custom Field to define segment criteria.

318.1 Creating a Custom Field

First, create a custom field to use for the Segment:

1. In the *Control Panel* go to *Configuration* → *Custom Fields*.
2. Click on *User*.
3. Click the *Add* button () button in the top right.
4. On the next page, click on *Dropdown*.
5. For the *Field Name* enter *Cardholder Type*.
6. For values, enter *None*, *Basic*, *Gold*, and *Platinum* on four separate lines.
7. Click *Save*.

Now any time user is created, they are prompted to enter the *Cardholder Type*, and existing users can select it from their user profiles.

For more information on adding Custom Fields, see [Adding Custom Fields](#).

Edit Selection List

Field Name

Cardholder Type

The custom field key is used to access the field programmatically through the `<liferay-ui:custom-attribute />` tag.

Data Type

Text

Values * ?

None
Basic
Gold
Platinum

ADVANCED PROPERTIES >

Save **Cancel**

Figure 318.1: You can easily create custom fields to capture whatever kind of data you need.

318.2 Defining a Segment with a Custom Field

Next, use a custom field to define another segment.

1. From Segments Administration, click the *Add* button (+).
2. Click in the text area at the top of the page, enter the name *Premium Card Prospects*.
3. For User Properties select *Cardholder Type*.
4. Click on the operator field and set it to *equals*.
5. Select *Basic* from the select box.
9. Click *Save* to save your Segment.

As you can see, segment criteria can be easily defined using Liferay DXP's built-in criteria or your own custom fields. Now, you'll use session data to make this Segment definition even more robust.





Figure 318.2: The custom field you created is seamlessly integrated into segment creation.

318.3 Extending a Segment With Session Data

So far, you've used criteria derived from user profiles to determine if they should be members of a segment; now it's time to use session data to make your criteria more effective.

Note: For this exercise to work, you must have set a cookie on the specified page.

1. Click the *Actions* button () next to the *Premium Card Prospects* segment and select *Edit*.
2. Click the *Session* dropdown from the Properties menu.
3. Set the comparator for Session Properties to *Or*.
4. Drag *Cookies* into the Session Properties section.
5. Change the selection box to *contains*.
6. Enter *Cards* in the Key text box.

Now any user who has a cookie from visiting the *Cards* page is evaluated as part of the *Premium Card Prospects* segment.

USING ANALYTICS CLOUD WITH USER SEGMENTS

To use Analytics Cloud with User Segments, you must first connect your DXP data source to Analytics Cloud and enable synchronization of users and analytics. For more information about Analytics Cloud, including instructions for connecting it with DXP, see the official Analytics Cloud Documentation.

Once you're connected to Analytics Cloud, you can create Segments to analyze user behavior.

319.1 Creating a New Segment

Synchronization with Analytics Cloud is not instant, so once you have connected Analytics Cloud and Liferay DXP, you must first wait for the users and data to synchronize. After that completes, you can create Segments in Analytics Cloud to capture data in DXP.

Only Segments that contain at least one member are synchronized with Liferay DXP. This means that empty Segments created with Analytics Cloud are unavailable to use on Liferay DXP.

See Analytics Cloud's documentation on creating Segments for more information.

Next, you can use your new Segment to define behaviors on your server.

319.2 Getting Segment Analytics

After you create and sync a Segment in Analytics Cloud, you can view it and customize it in Liferay DXP.

1. Go to the *Segments* page.
2. Click on the new Segment.






| Name | Active | Source | Scope | Modified Date |
|---|--------|---|--------------|--|
| <input type="checkbox"/> Anyone | Yes |  | Global | 13 Days Ago |
| <input type="checkbox"/> American Engineers | Yes |  | Current Site | 13 Days Ago  |
| <input type="checkbox"/> Premium Card Prospects | Yes |  | Current Site | 13 Days Ago  |
| <input type="checkbox"/> Landing Page Visitors | Yes |  | Global | 56 Seconds Ago  |

Figure 319.1: When you see Analytics Cloud Segments in the list of Segments, they are marked with the Analytics Cloud icon.

Analytics are based on the criteria that you set on Analytics Cloud, but you can set additional criteria here to use this Segment for personalization in DXP. Changing the Segment criteria here doesn't affect the gathered analytics data, unless it is configured in some way that restricts its members from viewing content that you are using as an Analytics Cloud criteria.


When you put it all together to provide personalized experiences and analyze user behavior, you can see the true power of Segmentation.

PERSONALIZATION EXPERIENCE MANAGEMENT

There is no direct location for managing all aspects of Experience Personalization. All the different aspects are managed from whatever you're personalizing. The key integration points of personalization are Content Pages and Content Sets. You can see an overview of their personalization options below.


320.1 Managing Content Page Personalization



Before you can Personalize Content Pages, you first need to create Content Pages. When you edit a Content Page, you can click on the *Experience* to manage the options for that page.

1. Go to *Site Administration* → *Site Builder* → *Pages*.
2. Click the *Actions* button  → *Edit* for any Content Page.
3. Click on the *Default* Experience to manage experiences.

From here you have three options:

 changes the name or selected Segment for the Experience.

 deletes the Experience.

  changes the priority of the Experience. If a user meets the criteria for more than one Experience, the highest ordered one is displayed.

When creating a new Experience, you must define the audience by choosing a Segment. If your target audience for the Experience is not yet represented by a Segment, you can create one from the New Experience interface.

Note: Creating new Segments from the New Experience interface is available in Liferay DXP 7.2 Fix Pack 1+ and Liferay Portal GA2+.

Next you'll learn about managing Content Set personalization.

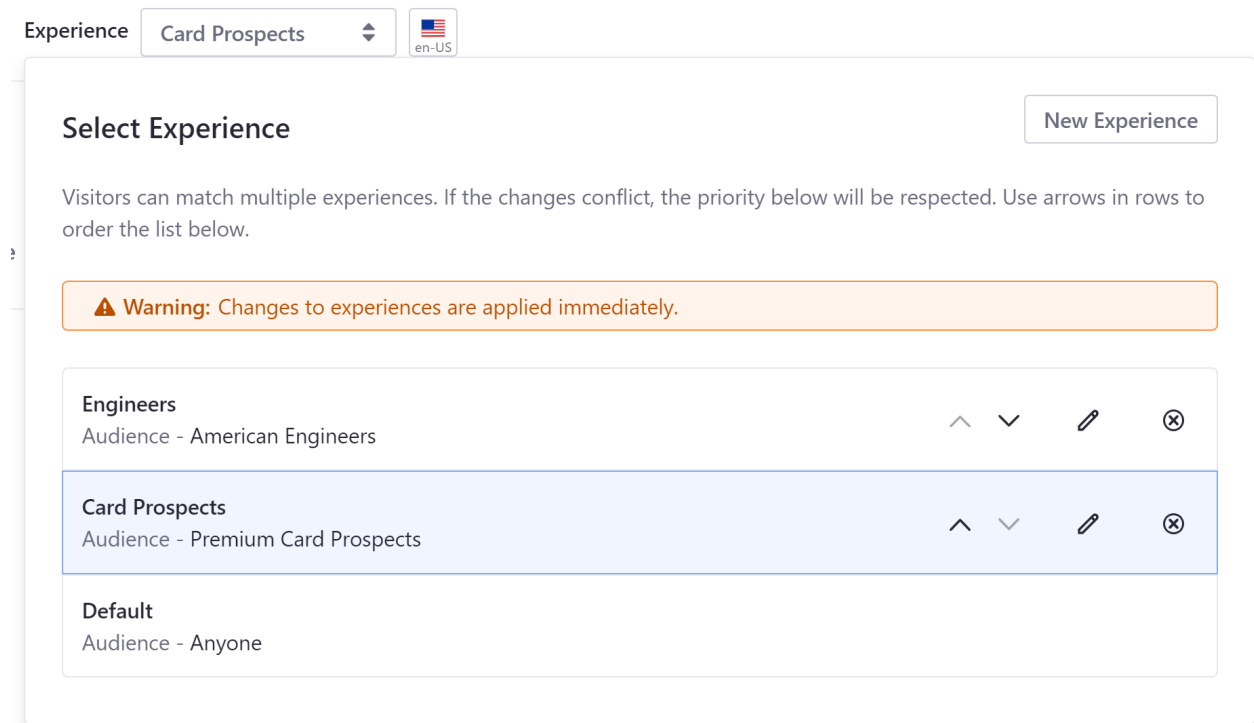


Figure 320.1: You can add, edit, delete, or change priority for Experiences.

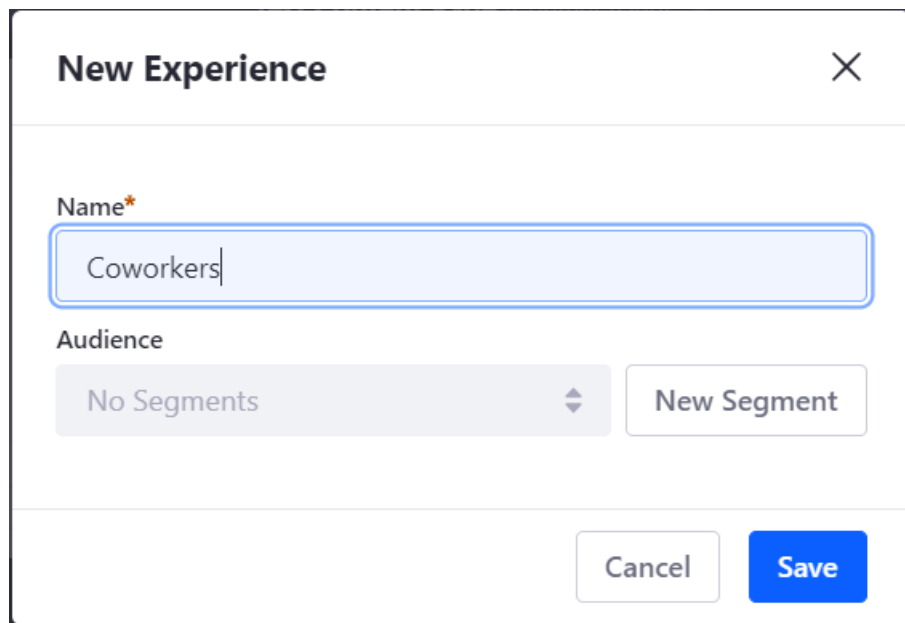



Figure 320.2: You can add a new Segment while creating a new Experience.

320.2 Managing Content Set Personalization

Managing Personalization options for Content Sets is similar. First you must create Content Sets then you can personalize them. To create a new Personalized Variation of a Content Set,

1. Click the *Actions* button  → *Edit* for the Content Set you wish to personalize.
2. Next to the content set is the message *No Personalized Variations yet*. Click the *New Personalized Variation* button.
3. Select a Segment from the next page to create a variation for that Segment.

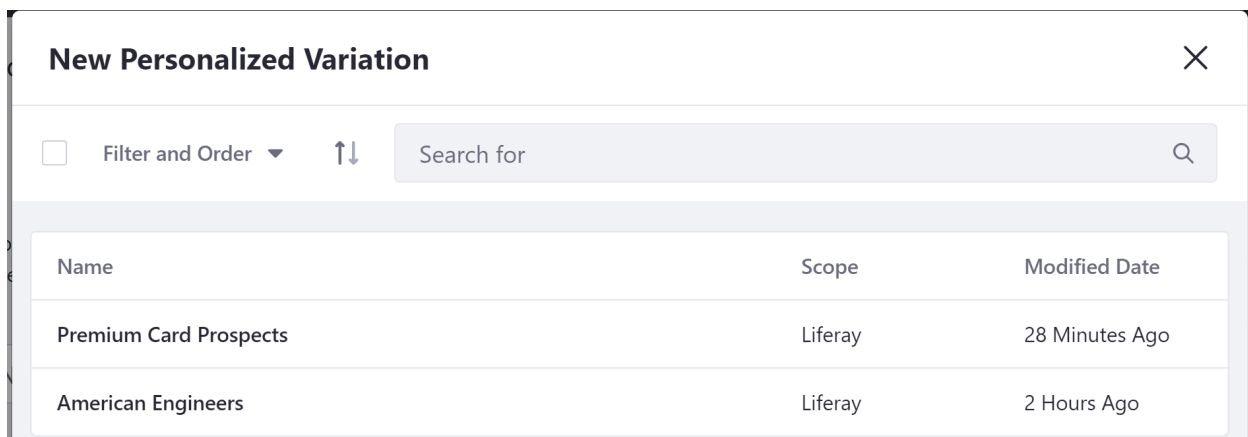



Figure 320.3: Select a Segment to create a variation for.

You can create a personalized variation for any existing Segment. Each new variation copies the default Content Set, but then essentially functions as its own Content Set after that. To edit or manage a Content Set Variation,

1. Click on the Segment name under *Personalized Variations*.
2. Click on the *Actions* button () and you can select *View Content* to preview the content in that set or *Delete* to remove it.

320.3 Previewing User Experiences

As an administrator, when you view a page, you can preview the different experiences that users can have on that page.

1. Click the *Simulation* button () icon from the top of any page.




Figure 320.4: You can preview or delete a Personalized Variation from the *Actions* menu.

2. Select a segment from the *Segments* selection to preview the page as a member of that segment.

Viewing the perspective of a segment previews for the administrator any personalizations for Content Pages or Content Sets for that Segment.

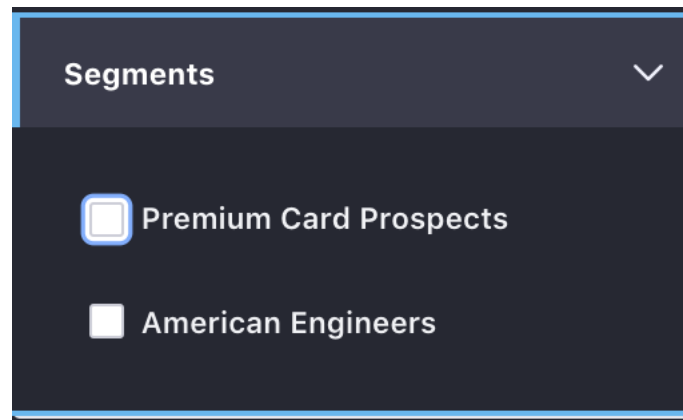


Figure 320.5: You can preview different experiences from the Preview Panel.

CONTENT PAGE PERSONALIZATION

In *Creating Segments with Customer Fields and Session Data* you created a Segment called *Premium Card Prospects*. Now, you'll use it to demonstrate Content Page Personalization.

If you're not familiar with Content Pages, see the *Creating Content Pages* article before you get started here.

321.1 Creating the Default Page

First, you need to create the *Credit Cards* page. To do this, create a new content page, add some fragments to it, and edit the content to match. First you need to build the bones of the page:

1. Go to *Site Administration* → *Site Builder* → *Pages*.
2. Click the *Add* button (+).
3. Select *Content Page* and name the page *Credit Cards*.
4. Open the *Sections* → *Basic Sections* from the build menu on the right side of the screen and add a *Banner* to the page.
5. Next, click on *Section Builder*, open *Layouts*, and add a three column layout to the page, above the banner.
6. Open the *Basic Components* tab and add a *Card* inside each of the columns.

Next, edit the content (this step is optional):

1. Click on the text for the cards and the banner to edit it and provide content relevant to someone looking for information about credit cards.
2. Click on each image and provide an appropriate image.

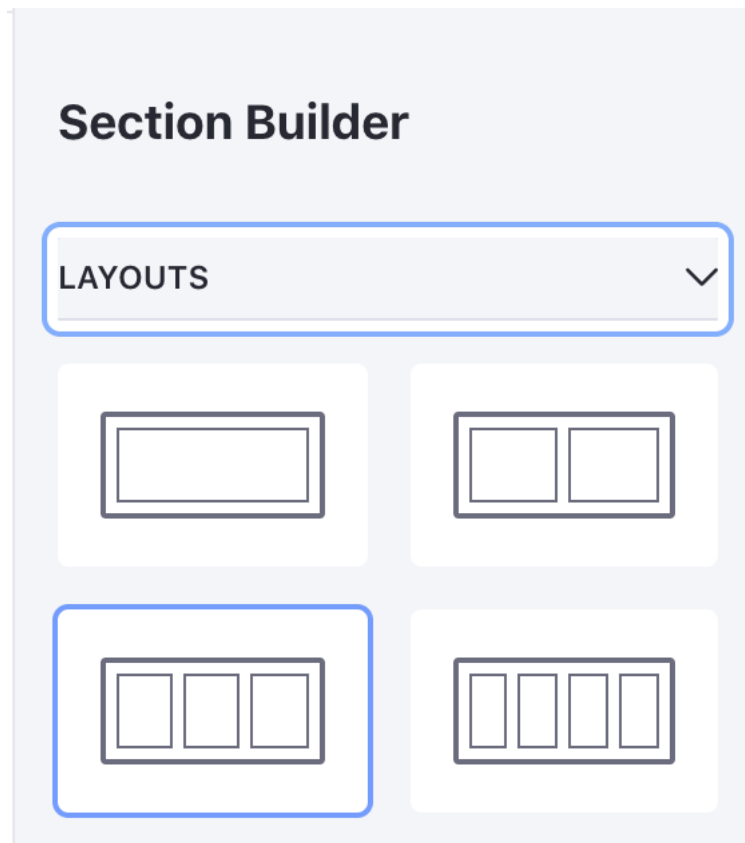


Figure 321.1: Open Layouts from the Section Builder.

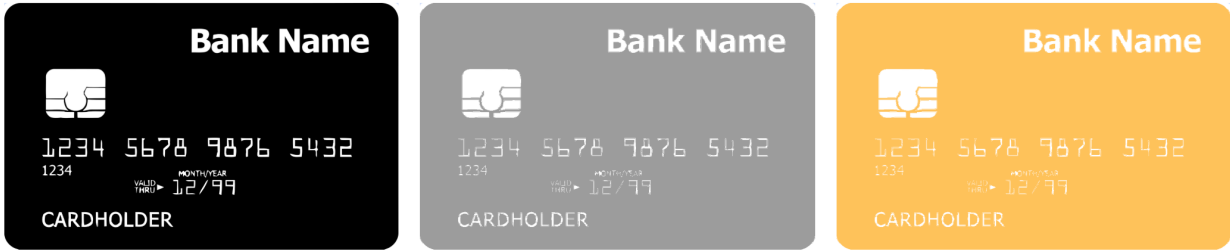
3. Click *Publish*.

This is the default page that anyone visiting the site sees.

321.2 Defining Custom Experiences

Next, define an experience specifically tailored to customers whom you have identified as *Premium Card Prospects* using User Segments.

1. At the top of the page, for the *Experience* click on *Default* to open the experience selection dialog.
2. Click on *New Experience*.
3. Name it *Card Prospects* and select *Premium Card Prospects* for the *Audience*.
4. Add a *Banner* fragment above the three columns you added earlier.
5. Edit the Banner to provide information specifically related to upgrading a card for an existing customer.



Bank Name

1234 5678 9876 5432
1234 VALID THRU MONTH/YEAR 12/99
CARDHOLDER

Basic Card
Our Basic Card provides competitive rates and our world class service.

[Apply Now!](#)

Bank Name

1234 5678 9876 5432
1234 VALID THRU MONTH/YEAR 12/99
CARDHOLDER

Silver Card
Get better rates and more rewards with the Silver Card!

[Apply Now!](#)

Bank Name



1234 5678 9876 5432
1234 VALID THRU MONTH/YEAR 12/99
CARDHOLDER

Gold Card
For our most elite and prestigious customers. The best card in the world!

[Apply Now!](#)

Find Out What Card is Right For You
Compare the features and benefits of all of our offerings to find out what's best for you.

Figure 321.2: Your final result might look something like this.

Experience Default   en-US

Select Experience [New Experience](#)

Create an experience to address a page to a specific audience. An experience is a copy of your current page. You can change content and layout for the new experience.

Warning: Changes to experiences are applied immediately.

Figure 321.3: Click on the current experience to create a new one or select a different existing experience.

6. Click *Publish*.

The *Default* version of the page appears for everyone except for those defined as *Premium Card Prospects*, but customers in that segment have an experience curated just for them.

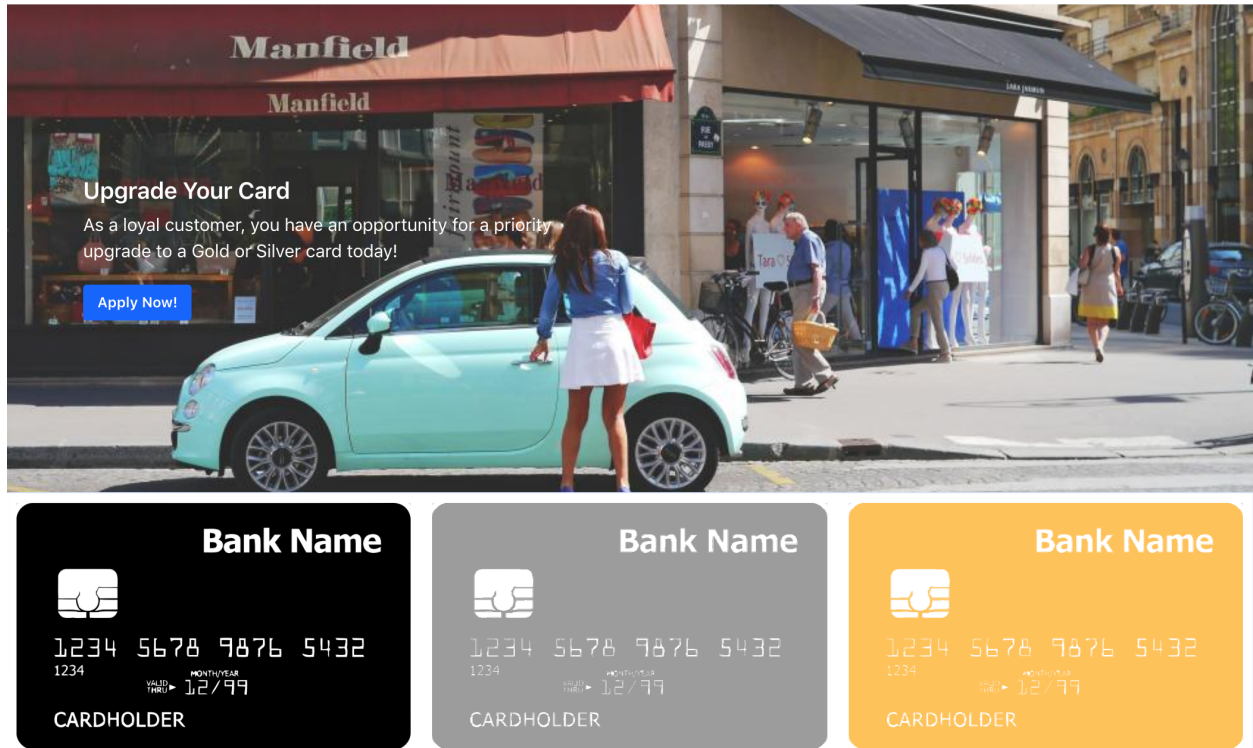


Figure 321.4: Your final result for the card prospects might look something like this.

Note: When you create a new experience, it copies the *Default* experience at the time that it is created. Any further changes to the *Default* experience do not effect any of experiences for that page.

CONTENT SET PERSONALIZATION

In *Creating User Segments* you created a Segment called *American Engineers*. Now, you'll use it to demonstrate Content Set Personalization. For this example, create a Content Set to be the default displayed on the *Home* page. Then you'll modify it to create a personalized variation containing technical articles for members of the *American Engineers* segment.

If you're not familiar with Content Set, see the *Creating Content Sets* article before you get started here.

322.1 Creating and Setting the Default Content Set

First create the default Content Set and configure it on the Home page using the Asset Publisher.

1. Go to *Site Administration* → *Content & Data* → *Content Sets*.
2. Click the Add button (+) and choose *Manual Selection*.
3. Name it *Home Page Content*.
4. For the new Content Set, click *Select* next to *Asset Entries* and select *Basic Web Content*.
5. On the *Select Basic Web Content* page, check the boxes next to the content you want to add and click *Add*.
6. Navigate to the *Home* page and add an Asset Publisher to the page.
7. Open *Configuration* for the Asset Publisher.
8. Under *Asset Selection* select *Content Set*.
9. Under *Select Content Set* click *Select*, choose *Home Page Content*, and click *Save*.

Now the Content Set that you configured appears in the Asset Publisher on the *Home Page*. Next configure the Content Set for Personalization.




Figure 322.1: Click *Select* to add a new Asset Entries.

322.2 Personalizing the Content Set

Now create the content set for engineers and configure its display.

1. Go back to the Content Set from *Site Administration*.
2. Click *New Personalized Variation* and select the *American Engineers* segment




Figure 322.2: Create a new Personalized Variation.

3. Click *Select* next to *Asset Entries* and select *Basic Web Content*.
4. Select articles appropriate to an engineering audience and click *Add*.

Now anytime a member of the *American Engineers* segment views this Content Set being displayed, they see the personalized version and not the default. Test this now, using the *Simulator*.

RECOMMENDING CONTENT BASED ON USER BEHAVIOR

Note: A/B Testing is available for Liferay DXP 7.2 SP1+.

A site's content generates clicks from users. For example, if someone visits a sporting goods store's site and clicks on several hunting promotional ads, you can deduce an interest in hunting products and can promote this type of content when this user visits the site again.

Accomplishing this with Segment-based personalization is possible, but that method is really targeted for vertical specific messaging or content with a preconceived audience. You may have to create hundreds of Segments to target all combinations of customer use cases. Instead, you need an infrastructure that tracks user views and displays the appropriate content based on behavior. You can accomplish this with *Content Recommendation*.

This is done by adding tags to content/Pages and monitoring the users who visit them. When a user views a specific content type or Page, its tags are attached to that user as *interests*. When the user visits other pages, content that matches their interests is displayed to them. The monitoring process is facilitated by Analytics Cloud, so you must have your DXP instance synced with it. If you haven't done this yet, start by adding your DXP instance as a data source.

Once your DXP instance is synced with Analytics Cloud and you're leveraging Content Recommendation, a user's interests are viewable by navigating to the left menu → *Individuals* → *Interests*.

You can learn more about Analytics Cloud's individual analytics here.

To begin recommending content to users, you must

- Add tags to content and/or Pages.
- Display content based on user behavior.

You'll step through these processes next.

323.1 Adding Tags to Track User Behavior

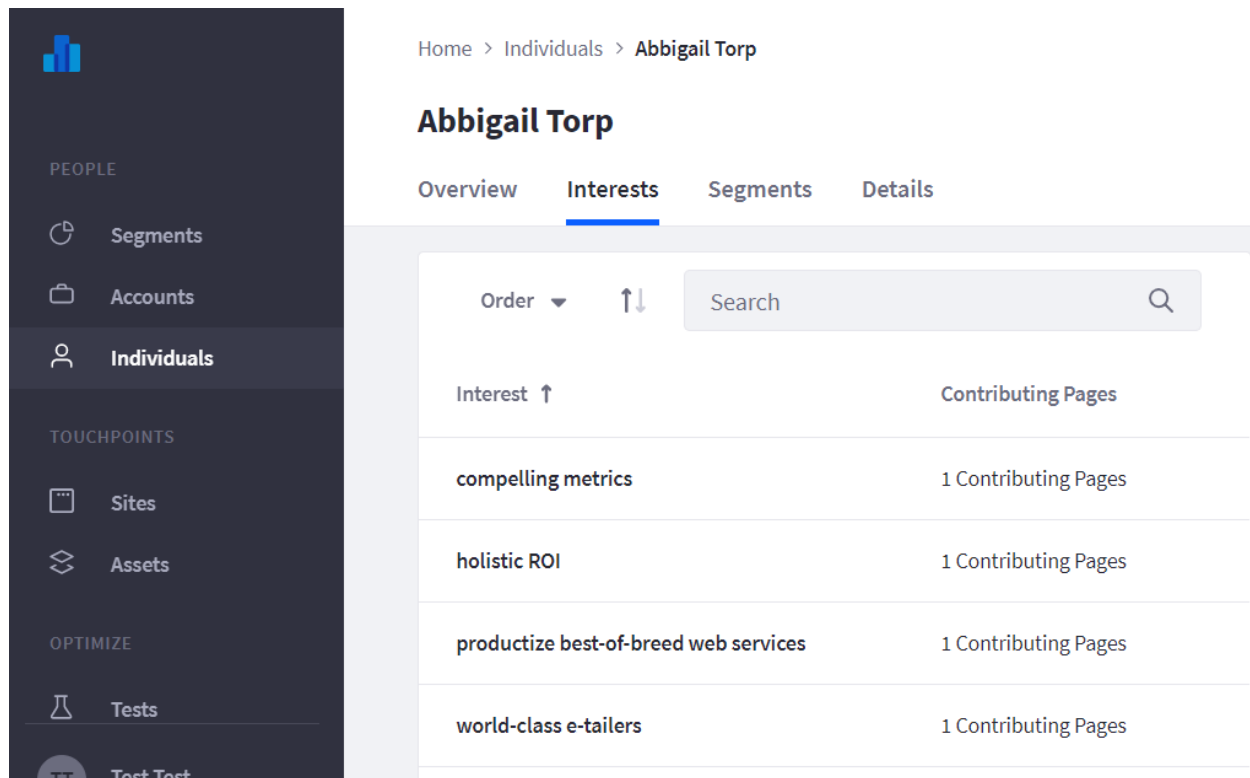



Figure 323.1: A user's interests are stored and accessible from Analytics Cloud.

To track user behavior and accumulate their interests, you must add tags to the content and Pages they visit. First, you'll add tags to web content and configure it to be viewable using a Display Page Template.

1. Go to Site Administration → Site Builder → Pages → Display Page Templates. Select the Add button (+), give it a name, and click Save.
2. In the right menu, select Mapping (🔗), choose the Web Content Article content type and choose Basic Web Content for the subtype. Then click Save.
3. Add a Fragment to the Display Page and map its field to Basic Web Content. For example, click Section Builder → Basic Components and drag the Paragraph Fragment to the page. Then click the Map button and select Basic Web Content for its Source and Field.
4. Publish the Display Page Template.
5. Begin creating Basic Web Content. Before publishing the content, navigate to the Display Page Template section and select Specific Display Page Template from the selector. Then select the Display Page Template you created previously.
6. Go to the Metadata section in the right menu. Assign tags that characterize the content. These are the tags that are referenced as interests when a user views the content. Then click Publish.

Now your web content is mapped to a Display Page, which allows the assigned tags to be tracked as interests when the web content is clicked. You can alter this process based on the asset types you want to recommend.


You can also assign tags to a Page's SEO configuration, which would then be assigned to users as interests when they visit the Page. Here's how to do this:

1. Navigate to Site Administration → *Site Builder* → *Pages*.
2. Click the *Actions* button () → *Configure* for a Page you want to add tags to.
3. Select *SEO* → *Categorization* and add relevant tags to the page. Then click *Save*.

Awesome! Now your content and Pages have tags that are assigned to users as interests when they visit them. These interests are assessed when recommending content, which you'll learn how to leverage next.

323.2 Displaying Content Based on User Behavior

Now that your Site's users' have their interests tracked using tags, you'll want to set up an Asset Publisher to display the content based on their behavior.

1. Navigate to the Site Administration → *Content & Data* → *Content Sets*.
2. Select the *Add* button () and click *Dynamic Selection*. Assign a name and click *Save*.
3. Under the Content Recommendation tab, enable Content Recommendation. Then save the Content Set.

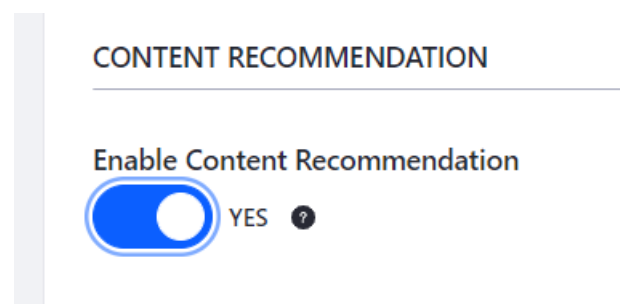



Figure 323.2: Enable Content Recommendation for your Content Set.

For more information on Content Sets, see [Creating Content Sets](#).

4. Add an Asset Publisher widget to a Page. Navigate to its *Options* () → *Configuration* menu and select the *Content Set* asset selection.
5. Select the Content Set you want to display. Then click *Save*.

In a realistic scenario, Content Sets have many assets with differing tags. That way, content similar to a user's interests is displayed over other content.

Great! Now when users have accumulated interests based on views, the Asset Publisher only shows content based on their interests.

A/B TESTING

Note: A/B Testing is available for Liferay DXP 7.2 SP1+.

Liferay DXP can help you hone your messages to customers by periodically testing different UX and messaging schemes. You can show different versions of your site to different users to see which is more effective.

A/B Testing lets you maintain a page's current UX and messaging, but provide alternative page variants for a select group of visitors. Then the current page and page variant(s) are tested based on algorithms to determine which pages perform better for a given goal (e.g., bounce rate, clicks, etc.).

For example, a Marketing team for a bank provides a Content Page advertising a new credit card. The page has been published for a few weeks, but a redesign might help promote the new credit card better. With A/B Testing, the team can create a new page variant and display both pages at random to visitors. Then they can analyze the clickthrough rate for the two pages and find which page is more effective. If the new variant is more effective than the original page, they can publish it and remove the old page.

If improving your site's UX and messaging is something you're interested in, continue on to learn more!

ENABLING A/B TESTING

Before creating an A/B test, you must ensure some conditions are met:


- You must have Liferay DXP connected to Analytics Cloud. To begin, add a Data Source.
- Your page must be a Content Page, since only Content Pages (not Widget Pages) support Experiences for different Segments.
- The Content Page you intend to test must be published.

If these conditions are met, A/B Testing is automatically enabled for any Content Page you navigate to, assuming you have the proper permissions. You'll learn more about how to assign permissions for A/B Testing next.

325.1 Setting A/B Testing Permissions

To use all the features of A/B Testing, you must have *Update* permissions for the Content Page.

To assign *Update* permissions for a Content Page,


1. Navigate to Site Administration → *Site Builder* → *Pages*.
2. Select the Actions button () next to the Content Page and click *Permissions*.
3. Enable the *Update* permission for the Role(s) you want to grant A/B Testing access for and click *Save*.

For more information, see [Page Permissions and Managing Roles](#).

CREATING A/B TESTS

To begin leveraging A/B Testing, you must first create an A/B test. You cannot create a test on an Experience that already has an active test running.

To create an A/B test,

1. Navigate to the Content Page you want to test.
2. Click the A/B Testing  button from the Control Menu.
3. Choose the Experience you want to test. This option is only available when you have a custom Experience (an Experience other than the default).

A test can be performed on the Default Experience as well as a personalized Experience mapped to a Segment. When an Experience is being used in an A/B test, it cannot be edited. Deleting a Page/Experience being used in an A/B test also deletes the test for that Page.

4. Click *Create Test*.
5. Assign the test a name and description (optional).
6. Assign the goal you want the test to track. There are two:
 - Bounce Rate*: the percentage of users who don't exhibit any activities on the page (click, scroll, etc.) and then navigate away from the site without visiting another page.
 - Click*: the percentage of users who clicked on the page (per session).
7. Click *Save*.

You now have an A/B test! Notice that the test's status is *Draft*. This means it's not yet visible to visitors.

Note: You can only create an A/B test for one page/experience at a time.


Note: The *Control* entity represents the currently published Content Page.

The screenshot shows a 'Create New Test' dialog box. At the top left is the title 'Create New Test' and a close button (X) at the top right. Below the title bar, there are three sections:

- Test Name ***: A text input field containing the text 'Prospect AB Test'.
- Description**: A text area containing the text 'Test if different wording for the paragraph decreases bounce rate.'.
- Select Goal ***: A dropdown menu with 'Bounce Rate' selected.


 At the bottom right of the dialog are two buttons: 'Cancel' and 'Save'.

Figure 326.1: Fill out the form to create your A/B test.

You can always edit or delete the new A/B test by clicking the *Actions* button () in the top right of the A/B Test menu. Deleted tests are not recoverable (i.e., not sent to the Recycle Bin). These options are not available for an active test. You can also, at any time, view your A/B Testing history by selecting the *History* tab. This displays all completed and terminated A/B tests.

Now it's time to create your test Variant(s). A test Variant is a customization of the Experience you want to optimize. An A/B test must contain at least one Variant before it can run.

To create a Variant,

1. In the A/B Test menu, click *Create Variant*.
2. Give the Variant a name and click *Save*.
3. Select the new Variant's Edit button . The current Control page's content/formatting is copied and displayed as the baseline for the Variant.
4. Edit the Variant as desired. Then click *Save Variant*.

You now have a Variant of the Control Page. You can create as many Variants as you want.

If you selected the Click goal, you must select the clickable element you want to target on the Control and Variant pages. If you selected a different goal, you can skip the steps below.

To configure the Click goal target,

1. Click *Set Target* under the Click Goal heading of your A/B Test. Any clickable element on the page is highlighted.

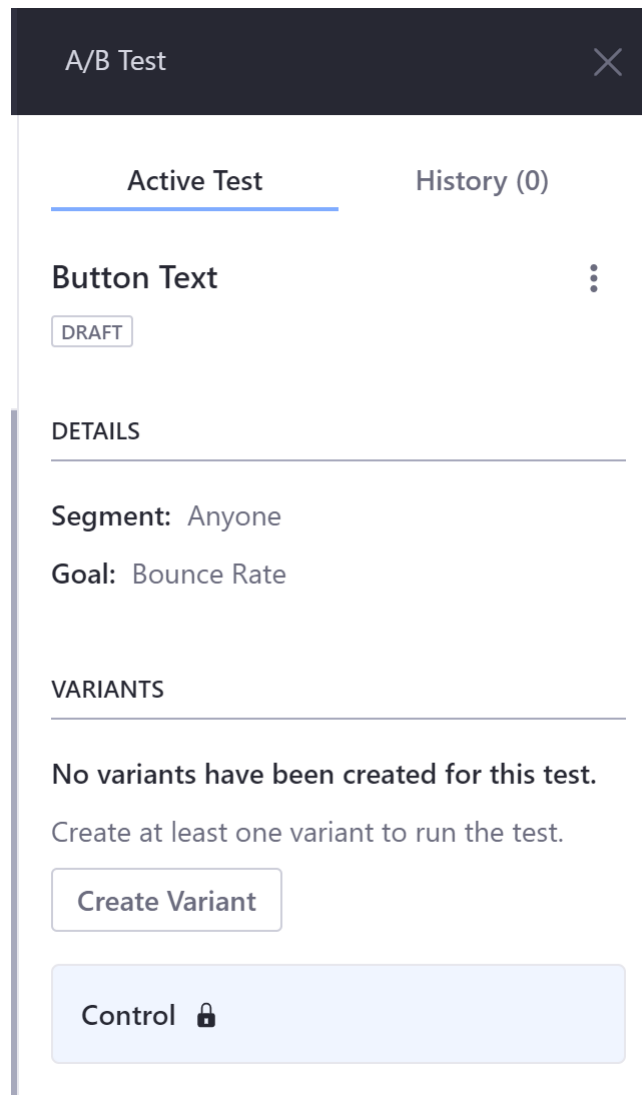


Figure 326.2: You now have an A/B test, but there are additional configurations you can apply.

****Note:**** Only links and buttons with an ID attribute can be selected as a target for the Click goal.

2. Select the element you want to set as the click target for your Control and Variant pages.

Your Click goal is now set! You can edit the target at any time before starting the test.

Once you're finished creating Variants and configuring goals for your A/B test, you're ready to run the A/B test. You'll learn how to do this next.




Figure 326.3: Set the click target to be tracked.

CLICK GOAL

Target: [#join-now-button](#)

Edit Target

Figure 326.4: Once the click target is set, you can run the A/B test.

RUNNING A/B TESTS

Once you've created and configured your A/B test, you'll want to run it to begin gathering data on your Control Page and Variants. To run an A/B test,

1. Click the *Review and Run Test* button from the A/B Test menu.
2. Configure how your test should run. There are two configurations:

Traffic Split: the percentage of visitors that are randomly split between the Variants when visiting the Page. A visitor randomly assigned a Variant always sees the same Variant until the test is finished.

Confidence Level Required: the accuracy of the test results (i.e., when the winning Variant truly outperforms the other Variants). Typically you want to have the highest confidence level possible, but this impacts test duration. The higher the required confidence level, the longer it takes to declare a winning Variant. Choose the percentage based on your expectations.

The *Estimated Time to Declare Winner* field is also displayed. This is the estimated duration the test runs. This is calculated based on the selected traffic split, confidence level, and projected page traffic. Your page's past traffic history is provided by Analytics Cloud.

3. Select *Run*.

Your A/B test is now running!

While an A/B test is running, you have two buttons available to help manage the test:

Terminate Test: terminates the test. To delete a test, you must terminate it first.

View Data in Analytics Cloud: redirects you to your A/B Testing dashboard hosted on Analytics Cloud. Here you can view your test's traffic, reports, statistics, etc. related to your test. See the [Monitoring A/B Test Results](#) article for more information.

Awesome! You now have a running A/B test accumulating data based on user interactions with your Page. Next, you'll learn how to monitor your A/B test's results.

Review and Run Test ✕

TRAFFIC SPLIT

Control 50%

Get Offer 50%

CONFIDENCE LEVEL

Confidence Level Required 95%

Estimated Time to Declare Winner 14 Days
Time depends on confidence level and traffic to the variants.

Figure 327.1: Configure the final parameters of your A/B test before running it.

MONITORING A/B TEST RESULTS

The information from your A/B test created in Liferay DXP is automatically synchronized with Analytics Cloud. The test results used to calculate the winning Variant based on user interaction data is viewable there. You can also view your A/B testing history, meaningful statistics, helpful graphs, etc. from Analytics Cloud.

To navigate to your test's Analytics Cloud dashboard from Liferay DXP, click the *View Data in Analytics Cloud* button in the A/B Test sidebar panel while it's running.

Liferay DXP only displays your test's status (draft, running, etc.) and the winning Variant once the test finishes. For more information on what you can view/manage from Analytics Cloud in relation to your A/B test, see the A/B Testing Analytics article.

PUBLISHING A/B TEST VARIANTS

Once the A/B test has concluded and Analytics Cloud has computed the test results, the status for the test displays as *Winner Declared* (if a winning Variant was found) or *No Winner* in the A/B Test sidebar panel. You're also alerted to the result via a notification, which you can view from the User Menu.

Note: When the required confidence level is not met during the time duration, there is no winning Variant.

Click the *Publish* button to publish your winning Variant Experience. Once you publish a Variant, the A/B test's status is *Completed* and the test is finished.

If you want to publish a Variant that was not declared a winner, select the Variant from the A/B Test panel and click *Publish*. You can select *Discard Test* to ignore the A/B test recommendations and keep the currently published Control Page.

The most productive Variant is now available to all users who visit the Page.

Awesome! You successfully ran an A/B test and published the Variant that is most effective for your Site's users.

The screenshot shows a mobile interface for an A/B test. At the top, a dark header contains the text "A/B Test" and a close icon. Below this, the "Select Experience" section shows a dropdown menu with "Blue Button" selected. A horizontal line separates this from the "Active Test" and "History (0)" tabs, with "Active Test" being the selected tab. The "Button Text" section displays "WINNER DECLARED" in a green-bordered box. Below this is a "DETAILS" section with a horizontal line, containing the text: "Segment: Engineers", "Goal: Click", and "Confidence Level: 95%". The "CLICK GOAL" section follows, with a horizontal line and the text "Element: #join-now-button". The "VARIANTS" section is at the bottom, featuring a light blue header with "Control" and a lock icon. Below the header are two variant rows: the first row shows "Get Offer" with a green checkmark and a blue "Publish" button; the second row shows "Learn More" with a grey "Publish" button. At the very bottom, there are two buttons: "Discard Test" and "View Data in Analytics Cloud" with an external link icon.

Figure 329.1: If you're satisfied with the A/B test's results, publish the winning Variant.




Figure 329.2: Once you've published a Variant, the A/B test is complete.

CONTENT PUBLICATION MANAGEMENT

Today's enterprises generate an enormous amount of content. You can use advanced publishing tools to manage content for a seamless publication experience.

Staging lets you change your Site behind the scenes without affecting the live Site. This is done by creating your Site in a temporary staging area (local or remote) and then publishing all the changes at once.

Continue on to begin managing content publication!

STAGING

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Staging is a tool used to manage content publication. The concept of staging is simple: you can modify your site behind the scenes and then publish all your updates in one shot. You don't want users seeing your web site change before their eyes as you're modifying it, do you? The staging environment lets you make changes to your site in a specialized *staging area* that's linked to a production environment. Typically the staging site is used only by content editors and site administrators, while the production environment is public. Content is published from staging to production all at once.

Site administrators can set up their staging environments locally or remotely. With Local Live staging, your staging environment and live environment are hosted on the same server. Remote Live staging has the staging and live environments on separate servers. You'll learn more about the differences between these two staging environments and how to enable them for your portal instance.

You can also leverage the Page Versioning feature. This feature works with both Local Live and Remote Live staging and lets site administrators create multiple variations of staged pages. This allows several different versions of sites and pages to be developed at the same time. Variations can be created, merged, and published using a Git-like versioning system. In the next section, you'll jump in to see how to enable staging.

ENABLING STAGING

This document has been updated and ported to Liferay Learn and is no longer maintained here. You have two different ways to set up staging:

- Local Live
- Remote Live

Whether you enable Local Live or Remote Live staging, the interface for managing and publishing staged pages is the same.

Local Live staging lets you publish site changes quickly, since the staged and live environments are on the same server. It's also easier to switch between the staged and live environments using Local Live staging. Since the staged content, however, is stored in the same database as the production content, your server must have more resources, and the content isn't as well protected or backed up as with Remote Live staging. Also, you can't install new versions of widgets for testing purposes in a Local Live staging environment, since only one version of an widget can be installed at any given time on a single Liferay server.

With Remote Live staging, your staging and live environments are hosted on separate servers, so your data is separated. This lets you deploy new versions of widgets and content to your staging environment without interfering with your live environment. Publishing is slower, however, with Remote Live staging since data must be transferred over a network. Of course, you also need more hardware to run a separate staging server.

Visit the staging environment article (Local or Remote) that most closely aligns with your goal for staging content.

ENABLING LOCAL LIVE STAGING

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Local Live staging places both your staging environment and your live environment on the same server. When it's enabled, a clone of the site is created containing copies of all of the site's existing pages. This means the staging and live environments share the same JVM, database, portlet data (depending on which portlets are selected when staging is enabled), and configurations, such as the properties set in the `portal-ext.properties` file. The cloned site becomes the staging environment and the original site becomes the live environment.

You can enable local staging for a site by navigating to the *Publishing* → *Staging* menu. To get some hands-on experience with enabling Local Live staging, you can complete a brief example which creates a Local Live staging environment for your site.

1. Navigate to the Product Menu (left side) and select *Publishing* → *Staging*.
2. Select *Local Live*. You can also enable page versioning and select staged content. For more information on these options, see the [Enabling Page Versioning and Staged Content](#) article.
3. Click *Save*.

You've officially begun the staging process!

Because Local Live staging creates a clone of your site, you should only activate staging on new, clean sites. Having a few pages and some widgets (like those of the example site you created) is no big deal. If you've already created a large amount of content, however, enabling staging can take a lot of time since it's a resource intensive operation. Also, if you intend to use page versioning to track the history of updates to your site, you should enable it as early as possible, *before* your site has many pages and lots of content. Your site's update history isn't saved until you enable page versioning. Page versioning requires staging (either Local Live or Remote Live) to be enabled.

If you ever need to turn off the staging environment, return back to *Staging* from the *Publishing* dropdown. For more information on this, see the [Disabling Staging](#) article.

Great! Now you're ready to use Local Live Staging.

ENABLING REMOTE LIVE STAGING

This document has been updated and ported to Liferay Learn and is no longer maintained here.

In Remote Live staging, a connection is established between the current site and another site on a remote Liferay server. The remote site becomes the live environment and the current site becomes the staging environment—an instance of Liferay used solely for staging. The remote (live) Liferay server and the local (staging) Liferay server should be completely separate systems. They should not, for example, share the same database. When Remote Live staging is enabled, all the necessary information is transferred over the network connecting the two servers. Content creators use the staging server to make their changes while the live server handles the incoming user traffic. When changes to the site are ready to be published, they are pushed over the network to the remote live server.

Before enabling Remote Live staging, ensure you've configured your Liferay server and remote server appropriately. Follow the [Configuring Servers for Remote Live Staging](#) article to do this.

You can enable remote staging for a site by navigating to the *Publishing* → *Staging* menu. Step through the instructions below to create a Remote Live staging environment for your site.

1. Navigate to the Product Menu (left side) and select *Publishing* → *Staging*.
2. Select *Remote Live*. Additional fields appear for Remote Live Connection Settings.
3. Enter your remote Liferay server's IP address into the Remote Host/IP field. This field should match the host you specified as your `tunnel.servlet.hosts.allowed` property in the `portal-ext.properties` file. If you're configuring an IPv6 address, it must contain brackets when entered into the *Remote Host/IP* field (e.g., `[0:0:0:0:0:0:1]`).

If the remote Liferay server is a cluster, you can set the Remote Host/IP to the load balanced IP address of the cluster to increase the availability of the publishing process. See the [Configuring Remote Staging in a Clustered Environment](#) for details.

4. Enter the port on which the remote Liferay instance is running into the Remote Port field. You only need to enter a Remote Path Context if a non-root portal servlet context is being used on the remote Liferay server.
5. Enter the ID of the site on the remote Liferay server that's for the Live environment. If a site hasn't already been prepared on the remote Liferay server, you can log in to the remote Liferay server and create a new blank site.

Select one of the options:

None

Local Live: A clone of the current site is created which contains copies of all existing pages and data of applications configured to be staged (see Staged Applications below). This clone becomes the Staging environment while the original becomes the Live environment.

Remote Live: A connection is made between this site and one existing in a remote server. This site becomes the Staging environment while the remote site becomes the Live environment.

REMOTE LIVE CONNECTION SETTINGS

Info: To publish changes to a *Remote Host/IP*, remote access between the servers must be configured. For detailed instructions on configuring this, refer to Liferay's official documentation. The *Remote Path Context* is only required if a non-root portal servlet context path is used on the target server. Access to this context must not be blocked by a proxy or firewall. Also, if the target server is a cluster, it is safe to set the *Remote Host/IP* to the cluster's load balanced address to increase the high availability of the publishing process.

Remote Host/IP

Remote Port

Remote Path Context

Remote Site ID

Use a Secure Network Connection

Figure 334.1: After your remote Liferay server and local Liferay server have been configured to communicate with each other, you have to specify a few Remote Live connection settings.

After the site has been created, note the site ID so you can enter it into the Remote Site ID field on your local Liferay server. You can find any site's ID by selecting the site's name on the Sites page of the Control Panel.

6. Check the *Use a Secure Network Connection* field to use HTTPS for the publication of pages from your local (staging) Liferay server to your remote (live) Liferay server.
7. Decide whether to enable page versioning and select staged content. For more information on these options, see the [Enabling Page Versioning and Staged Content](#) article.
8. Click *Save*.

You've officially begun the staging process!

If you fail to configure your current and remote server properly, you won't be able to enable staging and an error message appears. If you have issues, verify you've configured your servers properly.

When a user publishes changes from the local (staging) server to the remote (live) server, Liferay DXP passes the user's email address, screen name, or user ID to the remote server to perform a permission check. For a publishing operation to succeed, the operation must be performed by a user that has identical credentials and permissions on both the local (staging) and the remote (live) server. This is true regardless of whether the user attempts to publish the changes immediately or attempts to schedule the publication for later.

If only a few users should have permission to publish changes from staging to production, it's easy enough to create a few user accounts on the remote server that match a selected few on the local server. The more user accounts that you have to create, however, the more tedious this job becomes and the more likely you are to make a mistake. And you not only have to create identical user accounts, you also have to ensure that these users have identical permissions. For this reason, it's recommended that you use LDAP to copy selected user accounts from your local (staging) Liferay server to your remote (live) Liferay server. Liferay's Virtual LDAP Server application, available on Liferay Marketplace, makes this easy.

See the [Disabling Staging](#) article to learn how to turn off the staging environment.

Great! Now you're ready to use Remote Live Staging.

CONFIGURING SERVERS FOR REMOTE LIVE STAGING

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Before you can enable Remote Live staging for a site, you must satisfy some necessary requirements:

- Add the remote Liferay server to the current Liferay server's list of allowed servers, and vice versa.
- Specify an authentication key to be shared by your current and remote server.
- Enable each Liferay server's tunneling servlet authentication verifier.
- Update the Tunnel Auth Verifier Configuration of your remote Liferay instance.

Follow the steps below to configure your servers for Remote Live staging.

1. Add the following lines to your current Liferay server and remote Liferay server's `portal-ext.properties` file:

```
tunneling.servlet.shared.secret=[secret]
tunneling.servlet.shared.secret.hex=true
```

Liferay DXP's use of a pre-shared key between your staging and production environments helps secure the remote publication process. It also removes the need to send the publishing user's password to the remote server for web service authentication. Using a pre-shared key creates an authorization context (permission checker) from the provided email address, screen name, or user ID *without* the user's password.

2. Specify the values for the servers' `tunneling.servlet.shared.secret` property.

The values for these properties depend on the chosen configured encryption algorithm, since different encryption algorithms support keys of different lengths. See the HTTP Tunneling properties documentation for more information. Note that the following key lengths are supported by the available encryption algorithms:

- AES: 128, 192, and 256 bit keys
- Blowfish: 32 - 448 bit keys

- DESede (Triple DES): 56, 112, or 168 bit keys (However, Liferay places an artificial limit on the minimum key length and does not support the 56 bit key length)

To prevent potential character encoding issues, you can use one of the following two strategies:

2a. Use hexadecimal encoding (recommended). For example, if your password was *abcdefghijklm-nop*, you'd use the following settings in your `portal-ext.properties` file:

```
tunneling.servlet.shared.secret=6162636465666768696a6b6c6d6e6f70
tunneling.servlet.shared.secret.hex=true
```

2b. Use printable ASCII characters (less secure). This degrades the password entropy.

If you don't use hexadecimal encoding (i.e., if you use the default setting `tunneling.servlet.shared.secret.hex=true`), the `tunneling.servlet.shared.secret` property's value *must* be ASCII compliant.

Once you've chosen a key, make sure the value of your current server matches the value of your remote server.

Important: Do not share the key with any user. It is used exclusively for communication between staging and production environments. Any user with possession of the key can manage the production server, execute server-side Java code, etc.

3. Add the following line to your remote Liferay server's `portal-ext.properties` file:

```
tunnel.servlet.hosts.allowed=127.0.0.1,SERVER_IP,[STAGING_IP]
```

The `[STAGING_IP]` value must be replaced by the staging server's IP addresses. If the server has multiple interfaces, each IP address must also be added, which would show as a source address for the `http(s)` requests coming from the staging server. The `SERVER_IP` constant can remain set for this property; it's automatically replaced by the Liferay server's IP addresses.

If you're validating IPv6 addresses, you must configure the app server's JVM to not force the usage of IPv4 addresses. For example, if you're using Tomcat, add the `-Djava.net.preferIPv4Stack=false` attribute in the `$TOMCAT_HOME/bin/setenv.[bat|sh]` file.

4. Update the *TunnelAuthVerifierConfiguration* of your remote Liferay instance. To do this, navigate to the Control Panel → *Configuration* → *System Settings* → *API Authentication* → *Tunnel Authentication*. Click `/api/liferay/do*` and insert the additional IP addresses you're using in the *Hosts allowed** field. Then select *Update*.

Alternatively, you can also write this configuration into an OSGi file (e.g., `osgi/configs/com.liferay.portal.secret.default.config`) in your Liferay DXP instance:

```
enabled=true
hostsAllowed=127.0.0.1,SERVER_IP,[Local server IP address]
serviceAccessPolicyName=SYSTEM_USER_PASSWORD
urlsIncludes=/api/liferay/do
```

5. Restart both Liferay servers after making these configuration updates. After restarting, log back in to your local Liferay instance as a site administrator.

That's all you need to do to configure Remote Live Staging! You can now enable it! For additional information on configuring Remote Live staging, see the topics below.

335.1 Applying Patches When Using Remote Staging

When applying patches to a remote staging environment, you must apply them to all your servers. Having servers on different patch levels is not a good practice and can lead to import failures and data corruption. It is essential that all servers are updated to the same patch level to ensure remote staging works correctly.

335.2 Configuring Remote Staging's Buffer Size

Similar to Local Live staging, it is a good idea to turn remote staging on at the beginning of your site's development for good performance. When you're using Remote Live staging, and you are publishing a large amount of content, your publication could be slow and cause a large amount of network traffic. Liferay DXP's system is very fast for the amount of data being transferred over the network. This is because the data transfer is completed piecemeal, instead of one large data dump. You can control the size of data transactions by setting the following portal property in your `portal-ext.properties` file:


```
staging.remote.transfer.buffer.size
```

This property sets the file block sizes for remote staging. If a LAR file used for remote staging exceeds this size, the file will be split into multiple files prior to transmission and then reassembled on the remote server. The default buffer size is 10 megabytes.

ENABLING PAGE VERSIONING AND STAGED CONTENT

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Enabling page versioning for a site lets site administrators work in parallel on multiple versions of the site's pages. Page versioning also maintains a history of all updates to the site from the time page versioning was enabled. Site administrators can revert to a previous version of the site at any time. This flexibility is very important in cases where a mistake is found and it's important to publish a fix quickly.

You can enable page versioning for public pages or private pages on the Staging Configuration page below the menu for selecting your staging environment (Local or Remote). If you've already enabled staging, you can navigate to the Product Menu → *Publishing* → *Staging* and click the Options  button and select *Staging Configuration*.

You can also choose content for the staging environment to manage on the Staging Configuration page.

Choosing content to be staged may sound self-explanatory, but content must have specific attributes in Liferay DXP to use it in a staged environment. Content or an entity should be site-scoped, so they are always part of a site; otherwise, they are not eligible for staging. For example, page-scoped entities are only eligible for staging on published pages. When scoped data is on a page (e.g., Web Content Display widget) and the page is published, the scoped data is published with it.

Liferay DXP by default supports the following content groups for staging:

- Blogs
- Bookmarks
- Calendar
- Documents and Media
- Dynamic Data Lists
- Forms
- Knowledge Base
- Message Boards
- Mobile Device Families
- Polls

PAGE VERSIONING

Page versioning lets you work in parallel in different variations of the pages. It also lets you keep track of the history of changes in those pages.

- Enabled on Public Pages
- Enabled on Private Pages

STAGED CONTENT ?

When an application is checked, its data is copied to staging and it may not be possible to edit them directly in live. When unchecking an application, make sure that any changes done in staging are published first, because otherwise they might be lost.

- Select All
- Blogs
- Bookmarks
- Calendar
- Documents and Media
- Dynamic Data Lists
- Forms
- Knowledge Base
- Message Boards
- Mobile Device Families
- Polls
- Segments
- Web Content
- Widget Templates
- Wiki

Save

Figure 336.1: You can decide to use versioning and choose what content should be staged.

- Web Content
- Widget Templates
- Wiki

Before you activate staging, choose which of these widgets' data you'd like to copy to staging. You'll learn about many of the collaboration widgets listed under the Staged Content heading when you read the Collaboration Suite's section of articles. For now, be aware that you can enable or disable staging for any of these widgets.

Why might you want to enable staging for some widget types but not others? In the case of collaborative widgets, you probably *don't* want to enable staging since such widgets are designed for user interaction. If their content were staged, you'd have to publish your site manually whenever somebody posted a message on the message boards to make that message appear on the live site. Generally, you want web content to be staged because end users aren't creating that kind of content—web content is the stuff you publish to your site. But widgets like the Message Boards or Wiki should *not* be staged. Notice which widgets are marked for staging by default: if you enable staging and accept the defaults, staging is *not* enabled for the collaborative widgets.

The listed widgets, or content groups, contain one or more specific entity. For example, selecting the Web Content widget does not mean you're only selecting web content itself, but also web content folders.

Certain content types can be linked together and can reference each other on different levels. One of the responsibilities of staging is to discover and maintain these references when publishing. Site administrators and content creators have control over the process on different levels: staging can be enabled for a content group and a content group can be selected for publication.

Disabled staged content types can cause unintended problems if you're referring to them on a staged site. For example, the Asset Publisher portlet and its preferences are always staged. If the content types it's set to display are not enabled for staging, the Asset Publisher can't access them on a staged site. Make sure to plan for the content types you'll need in your staged site.

Turning Staging on and off for individual portlet data could cause data inconsistencies between the staging and live sites. Because of this, it's not possible to modify the individual portlet configuration once you enable staging. In case you need adjustments later on, you must turn staging off and re-enable it with your new configuration.

Besides managing the widget-specific content, Liferay DXP also has several special content types such as pages or users. For instance, pages are a part of the site and can reference other content types, but in a special way. The page references widgets, which means publishing a page also implies publishing its widgets. The content gives the backbone of the site; however, content alone is useless. To display content to the end user, you'll need widgets as the building blocks for your site.

Before you begin exploring the Staging UI, it's important to understand the publishing process for staging to make informed decisions so you use the staging environment efficiently and effectively.

PUBLISHING STAGED CONTENT EFFICIENTLY

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Now that you understand how staging works, you'll dive deeper into the publication process and some prerequisites you should follow before publishing. By understanding how the process works, you can make smart and informed decisions about how you want to publish your site's content.

337.1 Understanding the Publication Process

In simple terms, a publication is the process where all content, referenced content, apps and their preferences, pages, etc. are transferred from the staging scope to the live site. If you've enabled remote staging, this process involves network communication with another remote site. From a low level perspective, staging is an equivalence relation where entities are mirrored to a different location. From a high level perspective, the staging process happens in three phases:

1. **Export:** processes the publication configuration, which defines the site's content and apps. This phase also gathers the obligatory referenced entities that are required on the live site. Then everything according to the publication parameters is processed into the instance's own file format, and that file is stored locally or transferred to the remote live Liferay instance.
2. **Validation:** determines if it's possible to start the import process. This phase verifies the file's version and its integrity, checks for additional system information like language settings, and validates there is no missing content referenced.
3. **Import:** makes any necessary updates or additions to the site's content, layouts, and apps according to the publishing parameters. If everything is verified and correct, the staged content is published to your live Site.

These phases are executed sequentially.

A crucial factor for successfully publishing staged content is data integrity. If anything is not verified during the publication process, the transactional database reverts the site back to its

original state, discarding the current publication. This is a necessary action to safeguard against publishing incomplete information, which could break an otherwise well-designed live Site.

If the file system is not *database-stored* (e.g., DBStore), it's not transactional and isn't reverted if a staging failure occurs. This could potentially cause a discrepancy between a file and its reference in the database. Because of this, administrators should take great care with staging the document library, making sure that regular backups of both database and file system are being maintained.

Next, you'll learn about staging best practices and prerequisites to follow for a seamless staging experience.

337.2 Planning Ahead for Staging

Staging is a complex subsystem that's flexible and scalable. Before advanced users and administrators begin using it for their site, it's important to plan ahead and remember a few tips for a seamless process. There are several factors to evaluate.

- **Content (amount, type, and structure):** Depending on the content in your site, you can turn on staging for only the necessary content types, leaving others turned off to avoid unnecessary work. Publication can also be configured to publish only certain types of content. See the [Managing Content Types](#) article for more information.
- **Hardware Environment:** You should plan your environment according to your content types. If your site operates on large images and video files, decide if a shared network drive is the best option. Storing many large images in the Document Library usually requires a faster network or local storage. If you're dealing with web content, however, these are usually smaller and take up very little disk space.
- **Customizations and Custom Logic for Your Staging Environment:** Your organization's business logic is most likely implemented in an app, and if you want to support staging for that app, you must write some code to accomplish this. You can also consider changing default UI settings by writing new JSP code if you want your staging environment's look and feel to change.

Once you've finished planning for your site, you should turn on staging at the very beginning of the site creation process. This allows the site creator to avoid waiting for huge publications that can take long periods to execute. Taking smaller steps throughout the publication process forms an iterative creative process as the site is built from the ground up, where content creators can publish their changes immediately, avoiding long wait times.

Here are some JVM/network configuration recommendations for Staging:

- 4 GB of memory
- 20 MB/s transfer rate minimum (disk)

Now that you know how the staging environment works and how to enable it for your site, you'll begin using it in the next section.

USING THE STAGING ENVIRONMENT

This document has been updated and ported to Liferay Learn and is no longer maintained here.

After enabling staging (either Local Live or Remote Live) for a site, you'll notice additional options provided on the top Control Menu (Staging Bar) and also in the menu to the left. These new menus help you manage staged pages. Most of your page management options have also been removed; now you can't directly edit live pages. You now must use the staging environment to make changes.

Click the *Staging* button to view the staged area. Management options are restored and you can access some new options related to staging.




Figure 338.1: You can see the new staging options added to the top and left of your screen.

To test out the staging environment, add the Bookmarks widget and then click on *Live* from the top menu. Notice that the Bookmarks widget isn't there. That's because you've staged a change to the page but haven't published that change yet to the live site.

Next, you'll learn the basics of staging content.

STAGING CONTENT

This document has been updated and ported to Liferay Learn and is no longer maintained here.

When you're on the staged site, you have several options in the Staging Bar to help start your staging conquest.

Site Pages Variation: lets you work in parallel on multiple versions of a staged site page. You can choose the site page variation from the dropdown menu or manage them from the *Options* icon (⋮) in the Staging Bar. See the Using Site Pages Variations article for more information.

Page Variations: lets you work in parallel on multiple versions of a staged page. You can choose the page variation from the dropdown menu or manage them from the *Options* icon (⋮) in the Staging Bar. See the Using Site Pages Variations article for more information.

Undo/Redo: steps back/forward through recent changes to a page, which can save you the time of manually adding or removing apps if you make a mistake. To access *Undo/Redo*, select the *Options* icon (⋮) in the Staging Bar.

History: shows the list of revisions of the page, based on publication dates. You can go to any change in the revision history and see how the pages looked at that point. To access *History*, select the *Options* icon (⋮) in the Staging Bar.

Ready for Publication: After you're done making changes to the staged page, click this button. The status of the page changes from *Draft* to *Ready for Publication* and any changes you've made can be published to the Live Site. When you publish a page to live, only the version *Marked as Ready for Publication* is published.

Now you'll step through a brief example for using the Staging Bar to stage and publish content.

1. On the staged site, navigate to the Product Menu and select *Content & Data* → *Web Content*.
2. Create a Basic Web Content article and save it.
3. Go back to your staged site's main page and navigate to the *Add* (⊕) → *Widgets* → *Content Management* menu and drag the *Web Content Display* widget to the page.
4. Select the web content article you created to display in the new widget.
5. Select the *Ready for Publication* button to confirm you're ready to publish the content from the staged site to the live site. This prepares the staged content for publication. If workflow is enabled for any new resource, the resource must go through the workflow process before it can be published to the live site.




Figure 339.1: The staging toolbar indicates whether you're able to publish to the live site.

6. Click the *Publish to Live* button. A pop-up window appears with configuration options for your publication.
7. Enter the name of your publication.
8. Observe the changes listed in the menu. This lists the changed content planned for publication.

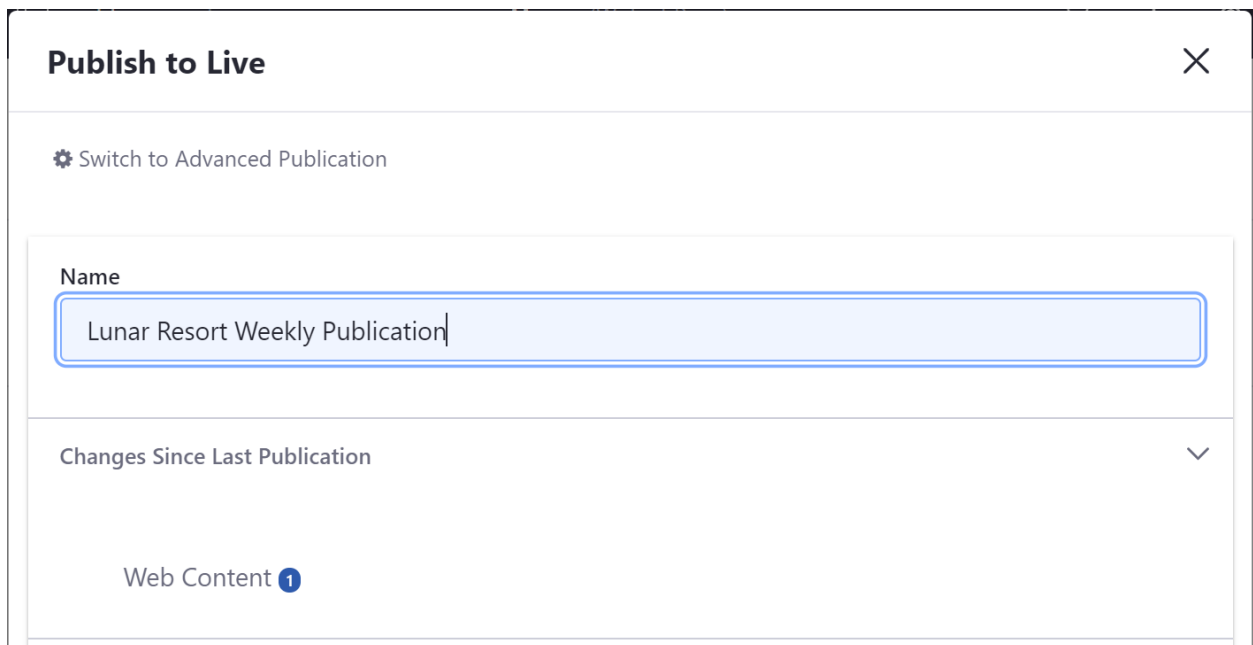


Figure 339.2: The Simple Publication menu displays the changes since last publication and a way to name your publication.

9. Click the *Publish to Live* button to publish your staged results to the live site.

Awesome! You've created content on the staged site and published it to your live site. It's now available for all your site users to see!

Note: Although publishing content is the more well-known function, publishing a portlet is also a viable option. You can publish portlets residing in the Control Panel and on pages. For example, you can modify a portlet's title and publish the change to live. This is possible because portlet configurations are always staged.

To publish a portlet that is on a page, you must publish the page first.

This example explored the Simple Publication menu. If your publication requires more advanced configuration like specifying specific content, dates, pages, etc., you should click the *Switch to Advanced Publication* button. You'll explore the more advanced configuration options next.

ADVANCED PUBLICATION WITH STAGING

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Once you've finished your changes on the staged site and want to publish them, select the *Publish to Live* button from the Staging Bar. To configure advanced publication options, select the *Switch to Advanced Publication* button. Opening the Advanced Publication menu presents options for scheduling a time to publish your content, editing the pages/content to include in the publication, managing permissions, etc. This lets you perform advanced editing to your publication process.

340.1 Date

You have two options for the Date category:

Now: immediately pushes any changes to the live Site.

Schedule: set a specific date to publish or to set up recurring publishing. You could use this to publish all changes made during the week every Monday morning without any further intervention.

These options let you plan staging schedules in advance.

340.2 Pages

This area of the menu gives you the option to choose which pages to include when you publish. You can choose the page group (Public or Private) to publish by selecting the *Change to Public Pages* or *Change to Private Pages*. You cannot publish both at the same time; you must complete their publication processes separately if you want to publish both page groups.

You can also choose specific pages to publish, and the look and feel of those pages.




Figure 340.1: You have several ways to specify the pages you want included in your publication.

Note: If you're publishing pages with a custom theme, you must check the *Theme Settings* option under the Look and Feel heading for your staging configuration. If it's not checked, the default theme is always applied.

The *Delete Missing Pages* option under the Look and Feel heading deletes all pages from the live Site that are not present on the staging Site. You can choose specific pages to publish (and remove) by manually selecting them under the Pages to Publish heading.

Note: The Simple Publication menu displays the number of page deletions tracked by the Staging framework. Keep in mind that this number counts the page deletions on the staging Site, not how many pages will be deleted on the live Site. There could be an inconsistency between the number of page deletions to be published and the actual number of pages present on either of the staging and live Sites. For example, pages that were deleted on the staging Site before they were published.

See the Deletions section for more information on content deletion.

340.3 Content

The Content area lets you select the content to be published. If you choose a page to be published from the Pages menu, the portlets and their references are always published, even if you specify differently in the Content section.

There are other filtering sub-options for certain content types. You first must choose what content to publish based on date. Specifying a date range lets you choose content to publish based on when it was created or last modified. Select the option that best fits your workflow. The available options are described in more detail below:

All: publishes all content regardless of its creation or last modification date.

From Last Publish Date: publishes content that was created or modified since the last publish date. This is the default option.

Date Range: publishes content based on a specified date range. You can set a start and end date/time window. The content created or modified within that window of time is published.

Last: publishes content based on a set amount of time since the current time. For example, you can set the date range to the past 48 hours, starting from the current time.

Under the date options are the different types of content that can be published. This list is populated based on the provided date range. For example, if at least one article is created or modified in the given date range, a Web Content section appears in the list, and the number of articles is shown next to the Web Content label. Otherwise, the Web Content section is absent.

The *Categories* content type is not dependent on the date range and is always shown in the list.

Note: Since comments and ratings are meant for the end user, they are not supported in staging and can only be added to the live site.

Unchecking the checkbox next to a certain content type excludes it from the current publication to the live site.

Some of the content types in the list, like Web Content and Documents and Media, have further filtering options. For instance, when the Web Content section is present and checked, it shows a comma-separated list of related items to be published, including the articles themselves. A sample list of related items for web content might look like this: *Web Content(12), Structures(3), Referenced Content, Include Always, Version History*. You can remove items by clicking the *Change* button next to the list.

See the *Managing Content Types in Staging* article for more information on managing content during the publication process.

340.4 Deletions

This portion of the menu lets you delete two things:

- portlet metadata before publishing
- operations performed for content types

You have two options to manage for deletions:

Delete Application Data Before Importing: all data created by the application is deleted before the import process. Ensure you understand the ramifications of this option before selecting it. Some applications in other pages may reference this data. This process cannot be undone. If you are unsure, complete an export first.

Replicate Individual Deletions: operations performed for content in the staging environment are replicated to the target site. This does not include page deletions.

340.5 Permissions

This area lets you include permissions for the pages and portlets when the changes are published. Select the *Publish Permissions* checkbox to enable this functionality.

Once you're finished configuring you advanced publication, select *Publish to Live* to publish or schedule your publication.

MANAGING CONTENT TYPES IN STAGING

This document has been updated and ported to Liferay Learn and is no longer maintained here.

When managing content in Staging's Advanced Publication menu, there are several factors to consider when preparing your content for publication. As described in Advanced Publication with Staging, you can navigate to the Content area of the Advanced Publication menu to select content you want to publish. There are options attached to each content group (e.g., Web Content) that you can manage too.

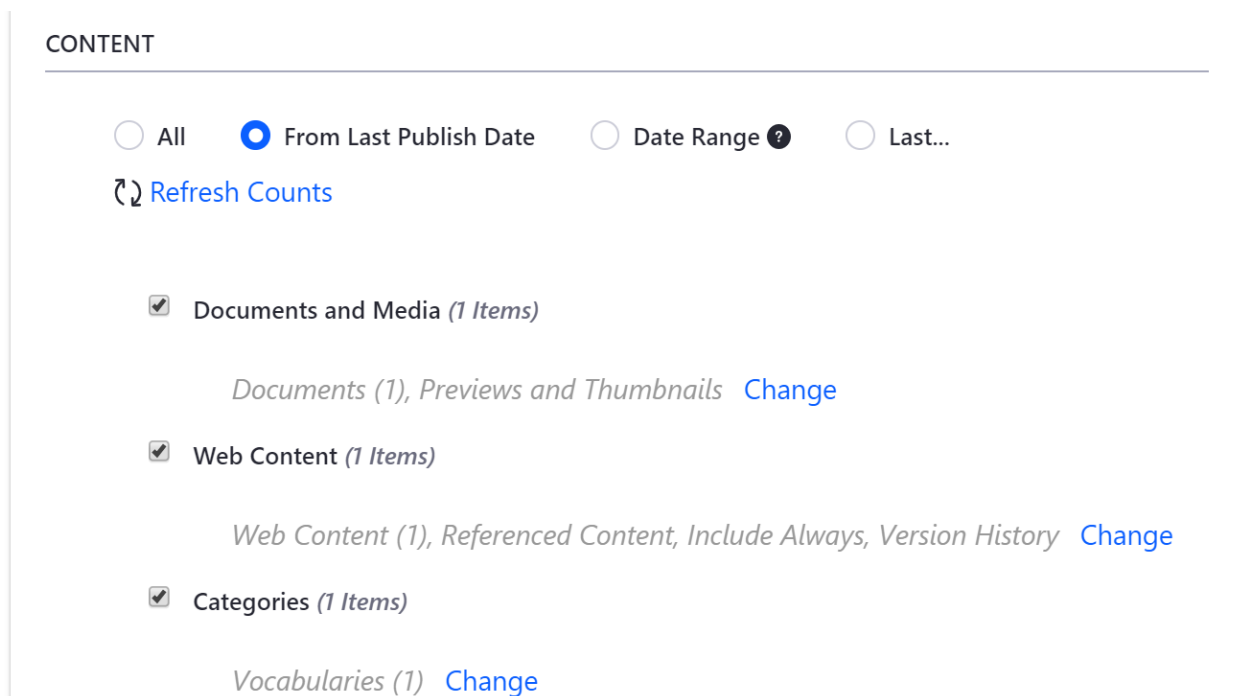


Figure 341.1: Click the *Change* button for a content group to manage its specific content.

You'll learn about some of these options and their best practices next.

341.1 Referenced Content

This is represented by

- Structures and templates included in web content.
- Documents and Media files (e.g., images) included in web content.
- etc.

You can exclude some of this content during publication or export to speed up the process. These references are validated during the publication process or an import, so the images must be published or imported first.

341.2 Version History

Web content tends to be updated frequently, often more so than other kinds of content. Sometimes this can result in high numbers of versions. If there are hundreds of versions, it can take a long time to publish these articles. You can bypass this by choosing to not publish the *Version History* (i.e., the past versions of the web content articles to be published). If you disable this, only the last **approved** version of each web content article is published to Live. This can significantly speed up the publication process.

You can set this option globally. If you navigate to the Control Panel → *Configuration* → *System Settings* → *Web Content* → *Virtual Instance Scope* → *Web Content*, you can toggle the *Version History by Default Enabled* checkbox. This sets the default behavior. When publishing content, it is selected by default, so site administrators must manually uncheck the *Version History* box to publish only the latest approved version of web content articles. To change the default behavior, enable the checkbox in System Settings.

341.3 Previews and Thumbnails

Previews and thumbnails are generated automatically for documents. Disabling this, though, can greatly increase your publishing speed in some cases. You should be careful about publishing previews and thumbnails to the live Site.

Imagine a scenario where a site has approximately 4000 images or documents. If the previews and thumbnails are turned on, this could end up in 28000 physical files on the disk. If staging is set up to publish the previews and thumbnails, this would mean that instead of taking care of the 4000 images, it would process seven times more files! If you still want to use the previews on your live environment, you can set up that Liferay instance to generate them automatically.

It depends on your environment for whether you can use the publishing of the previews and thumbnails. Publishing them is a heavy operation, and you must also transfer the LAR file over the

network if you use remote staging. If you decide to generate them on the live Site, understand that this could take some time, since it's a CPU intense operation.

341.4 Vocabularies

When working within a site, a user may select vocabularies from both the current site as well as the global site. While this doesn't pose an issue when creating content, it can cause issues when publishing.

For environments that use both global and local vocabularies, note that global vocabularies must be published to the live site through global site staging. One way to avoid confusion with vocabularies is to keep all vocabularies local or global.

If both must be used, you can resolve the issue by ensuring that dependencies (e.g., categories and vocabularies) are published before publishing the site that depends on them (whether the dependencies are local or global).


Assets like tags, categories, structures, templates, widget templates, document types, and dynamic data lists can also be shared by a parent to its child sites. In this case, ensure that the ancestor's dependencies are published before the site in question.

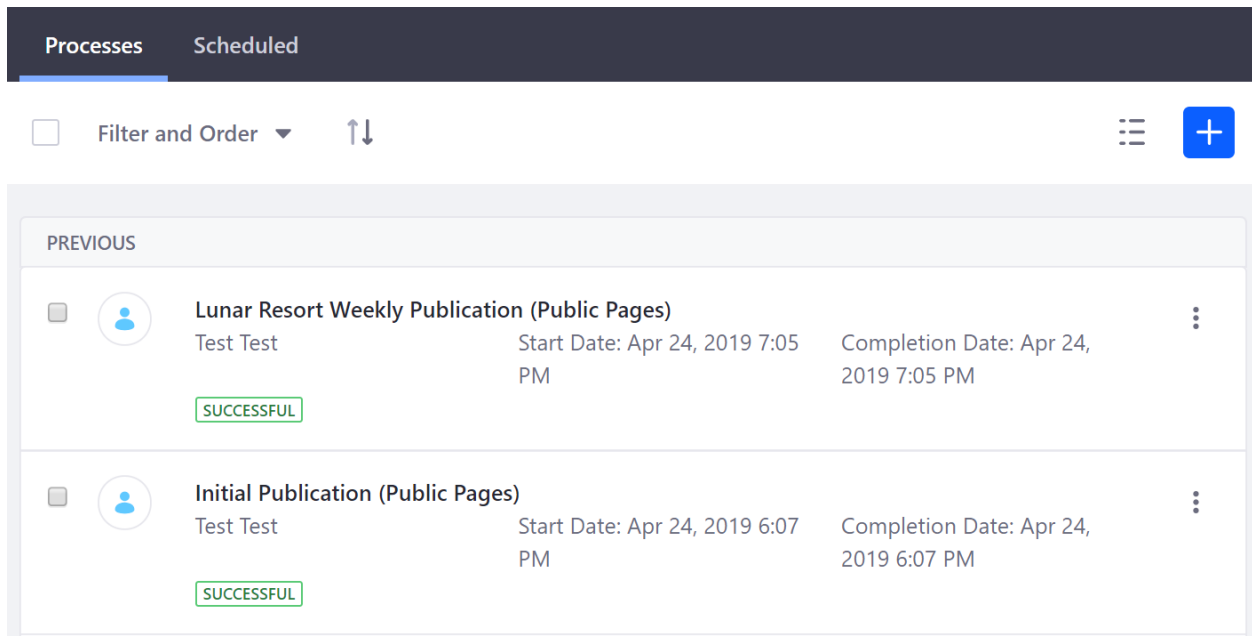
341.5 Deletions

The Staging framework gathers deletions (including trashed entities) in a site. These deletions can be published to clean up the live Site. If you plan to process it later, or if it's not a problem to have lingering data on live, this can be turned off as well to save execution time during the process.

STAGING PROCESSES AND TEMPLATES

This document has been updated and ported to Liferay Learn and is no longer maintained here.

When you make a staging publication, it's captured as a staging process and stored for future reference. You can manage these processes by navigating to the *Staging* option located in the Product Menu's Publishing tab. From there, you'll see a list of staging processes that have been completed. You can relaunch or clear any of these publications by clicking the *Actions* button () next to a process. If you click the *Scheduled* tab from above, you'll find staging processes that you've scheduled for future publication dates.



The screenshot shows the Liferay Staging Processes interface. At the top, there are two tabs: "Processes" (selected) and "Scheduled". Below the tabs, there is a "Filter and Order" dropdown menu and a sort icon. On the right side, there are three menu icons and a blue "+" button. The main content area is titled "PREVIOUS" and contains a list of two staging processes:




| Process Name | Start Date | Completion Date | Status |
|---|----------------------|----------------------|------------|
| Lunar Resort Weekly Publication (Public Pages) Test Test | Apr 24, 2019 7:05 PM | Apr 24, 2019 7:05 PM | SUCCESSFUL |
| Initial Publication (Public Pages) Test Test | Apr 24, 2019 6:07 PM | Apr 24, 2019 6:07 PM | SUCCESSFUL |

Figure 342.1: Your staging processes can be viewed at any time.

If you find yourself repeatedly creating similar staging processes, you should think about using Publish Templates.

Instead of manually having to customize a publication process every time you're looking to publish pages/content, you can use a publish template. With publish templates, you can select a custom template and immediately publish with the options you configured.

Follow the steps below to create and use a publish template.

1. Select the *Options* icon () from the top right corner of the Staging screen and select *Publish Templates*.
2. Click the *Add* button () and assign the template a title and description, and then fill out the configuration options as you would during a custom publication process.
3. Save your publish template. It's available to use from the *Publish Templates* tab in the *Publish to Live* menu's Advanced Publication area.
4. To use the template, click the *Actions* button () next to the template and select *Publish*.

This automatically sets the options for publishing pages and their content. All you have to do is give the publication process a name. Once you confirm the configuration settings, your staging settings are published.

Note: When staging is enabled, the options available from the *Publishing* tab are modified. When in the Live environment, you can only access the *Export* feature. When in the Staging environment, you can only access the *Import* and *Staging* features. The disabled features for each environment don't make sense in that context. For example, you shouldn't be able to import content when in the live environment; it must be imported into the staged environment and then published before it is available in the live site.

Now you know how to reference stored/scheduled staging processes and create publish templates to streamline publication.

DISABLING STAGING

Disabling staging doesn't take a lot of steps, but should not be done lightly. It's important to know the consequences of turning the staging environment off so you can decide if your circumstances really warrant it.

The consequences for disabling Local Live and Remote Live staging are slightly different, so you'll learn about both.

343.1 Disabling Local Live Staging

This material has been updated and ported to Liferay Learn and is no longer maintained here.

Conceptually, the live site is the final approved version of your site, whereas the staging site is a temporary workspace containing information that is not finalized.

Disabling local staging completely removes the staging environment, which means all content not published to your live site is erased. Therefore, before disabling staging, you must ensure all necessary information on the staged site is published or preserved elsewhere.

Keep in mind that draft content types are not published, so they can be lost too.

When you enabled staging there was an initial publication. Disabling staging does not start a publication; the staging site is deleted. If the staged site contains a large amount of content, however, those deletions could take a substantial amount of time to process. For this reason, don't disable staging when your portal instance is busy.

343.2 Disabling Remote Live Staging

This material has been updated and ported to Liferay Learn and is no longer maintained here.


Disabling remote staging does not delete the staged site; it only disables the connection between the live site and remote staging site. This means no data is deleted from the live or remote staging sites after the connection is disabled. Since no data is erased and no processes are started, disabling remote staging is almost instantaneous.

When Remote Live staging is enabled, certain information (e.g., which portlet is being staged) is recorded on both the live and staged Sites. For this reason, when you disable remote staging, you must ensure the live Site is still accessible so both sides can communicate that it's disabled. Do not shut down your live Site and then attempt to disable remote staging from your staged Site; this results in errors.

If there's ever a lost network connection between the remote staged Site and the live Site, a message appears, informing you of the error and a way to forcibly disable staging. This is only an option for the staged site; executing this option erases the staged site's staging information—not the content. On the contrary, the live site remains in a locked state. A possible workaround is to create a new live site and import content to it, if necessary.

343.3 Steps to Disable Staging

Follow the steps below to disable Local Live or Remote Live staging:

1. Navigate to the *Publishing* → *Staging* option, which is only available from the staged site.
2. Click the *Options* icon () from the upper right corner of the page and select *Staging Configuration*.
3. Select the *None* radio button and click *Save*.

That's it! Your staging environment is now turned off.

PUBLISHING SINGLE ASSETS FROM A STAGED SITE

This document has been updated and ported to Liferay Learn and is no longer maintained here. Sometimes, stepping through the entire publication process is not necessary and can be overkill. For example,

- What if you created a web content article and want only to publish it and nothing else?
- What if you want to publish a new folder of articles and their dependencies, but don't want the hassle of the publication process?
- What if you found a typo in a document and want to fix it quickly and push to the live site?


You're in luck! You can publish certain single assets from the staged Site to the live Site without creating a staging publication process, from their respective app menus:

- Web Content
 - Web Content
 - Folder
- Documents and Media
 - Document
 - Folder
 - Shortcut
 - Document Type
- Blogs
 - Blog
- Bookmarks
 - Bookmark
 - Folder

Important: Single asset publication is not supported for page-scoped content.

Note: When publishing a Web Content or Bookmarks folder, their respective entries and subfolders are included. Publishing a Documents and Media folder works the same way, but also includes shortcuts.

You'll step through an example to see how this is done.

1. Make sure the Staging framework is enabled and you're on the staged Site.
2. Create a Web Content Article in the Product Menu's *Content & Data* → *Web Content* menu.
3. Once you've saved the new Web Content Article, select its *Actions* button () next to the article and select *Publish to Live*.

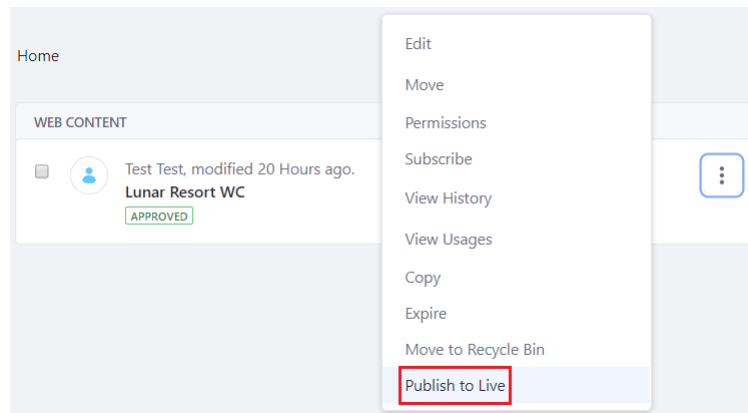


Figure 344.1: You can publish the single web content article to the live site.

4. You're presented a Process Details page where you can view the progress of your single asset publication request. Ensure the Web Content Article is published successfully.

****Note:**** Sometimes the publication process doesn't start immediately (e.g., if there's another publication running). You can check a specific asset's publication progress by navigating to the **Options** (![Options](./images/icon-options.png)) &arr; **Staging** &arr; **Current and Previous** tab in its Site Admin app.

There you have it! If you navigate to your live site's Web Content section, the new article is available.

Similar to the regular staging publication process, your single asset publications also include associated dependencies. For example, if your web content article contains an image, custom structure, and custom template, they are all published together. The same concept applies for folders—if you publish a folder containing several web content articles, all the articles and their associated dependencies are published too.

By default, only those with permissions to publish widgets can publish single assets. Follow the steps below to modify these permissions for a Role:

1. Navigate to the Control Panel → *Users* → *Roles*.

2. Select the Role you're updating.
3. Click the *Define Permissions* tab.
4. In the left menu, navigate to *Control Panel* → *Sites* → *Sites*.
5. Under the Resource Permissions heading, select the *Export/Import Application Info* option.

Also, make sure the *Publish Staging* permission is granted to the role. This is required to publish assets with staging. See the [Managing Permissions](#) article for more information.

Great! Now you know how to publish single assets and manage the permissions for who can do it.

ORGANIZING PAGES FOR STAGING

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Say you're working on a product-oriented Site where several major changes are required for a page or a set of pages over a short period of time. You must work on multiple versions of the Site at the same time to ensure everything has been properly reviewed before it goes live. With staging, you can do this using *Page Variations*.

In this section, you'll explore page variations and how they're useful.

USING MULTI AND SINGLE PAGE VARIATIONS

This document has been updated and ported to Liferay Learn and is no longer maintained here. There are two page variation options available from the Staging Bar:

- *Site Pages Variation*: Different variations for a set of Site pages. For instance, you could use this if you had three separate pages and wanted to modify these pages while keeping them together as a set.
- *Page Variations*: Different variations for a single page. A page variation can only exist inside a Site pages variation.

You must enable page versioning in the Staging Configuration menu before the above options are available in the Staging Bar. You'll learn more about this later.

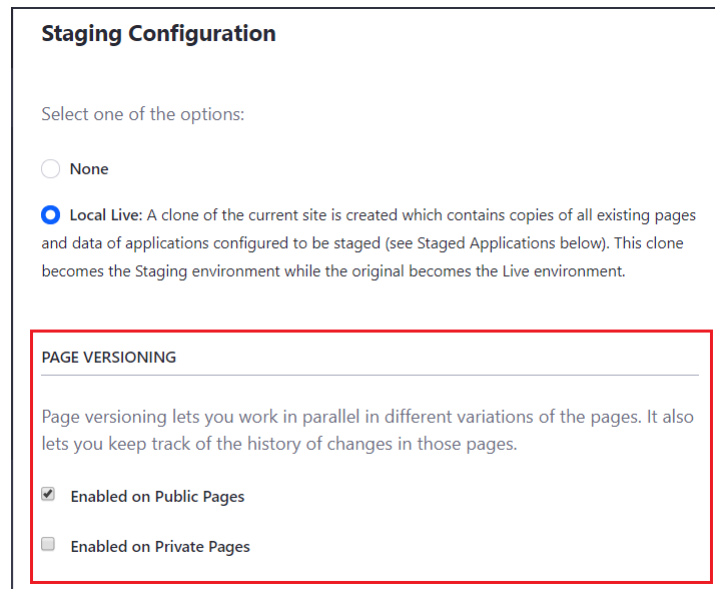
Variations only affect pages and not the content, which means all the existing content in your staging Site is shared by all your variations. The content, however, can be displayed in many different ways for each page variation. For example, in different Site page variations, you can have different logos, a different look and feel for your pages, different applications on these pages, different configuration of these applications, and even different pages. One page can exist in just one Site page variation or in several of them. Modifying the layout type (e.g., Layout, Panel, Embedded, etc.) or friendly URL of a page, however, **does** affect every Site page variation. For example, if a page template is modified, those modifications are propagated to the pages configured to inherit changes from the template, overriding Staging's Page Variations and Site Pages Variations.

Note: Page templates are not recognized by the Staging framework. This means that existing page templates are not viewable or editable on a staged Site. If they're created on a staged Site, they cannot be preserved once staging is disabled. You can, however, export and import page templates.

You'll learn about enabling page versioning next.

346.1 Enabling Page Versioning

Page Versioning is enabled on the Staging Configuration screen when first enabling staging.



Staging Configuration

Select one of the options:

None

Local Live: A clone of the current site is created which contains copies of all existing pages and data of applications configured to be staged (see Staged Applications below). This clone becomes the Staging environment while the original becomes the Live environment.

PAGE VERSIONING

Page versioning lets you work in parallel in different variations of the pages. It also lets you keep track of the history of changes in those pages.

Enabled on Public Pages

Enabled on Private Pages

Figure 346.1: You can enable page versioning for public and/or private pages.

You can enable page versioning for public and private pages. When page versioning is enabled, the page variation options are available in the Staging Bar. By default, you only have one Site pages variation and page variation which are both called *Main Variation*.

If you did not enable page versioning during the initial setup of your staging environment, navigate to the Product Menu → *Publishing* → *Staging* → *Options* (E) → *Staging Configuration*. You can enable the page versioning options there.

346.2 Using Site Pages Variations

Site pages variations are useful when you must plan multiple page sets for your Site at once. For example, consider this scenario:

If there were separate teams in your company that needed to create three drastically different page sets for your Site at the same time, they would need to create three Site pages variations. For example,

- The marketing team can give your Site a completely different look and feel for the Holidays.
- The product management team can work on a version that is planned to publish on the first of the year for a new product launch.
- The developer relations team can work on a version that is planned to publish on the upcoming Hack-a-thon day.

With this use case, having a Site pages variation for each planned page set allows three ideas to be fully planned and implemented before publication.

Another option for this scenario is to let the product management team experiment with two different ideas for the home page of the Site. This can be done by creating several page variations within the current Site pages variation of their product launch Site. You'll learn more about page variations later.

Once you've created a Site pages variation, you can now navigate to its home page and change the logo, apply a new theme, move applications around, change the order of the pages, configure different apps, and more. The other variations aren't affected. You can even delete existing pages or add new ones (remember to *Mark as Ready for Publication* when you are finished with your changes). When you delete a page, it is deleted only in the current variation. The same happens when you add a new page. If you try to access a page that was deleted in the current variation, Liferay informs you this page is not *enabled* in this variation and you must enable it.




Figure 346.2: Select the *Enable* button to create a missing page in the current Site pages variation.

To publish a variation to the live Site, click on *Publish to Live* in the staging menu and then select *Publish to Live*. Publications can also be scheduled independently for different variations. For example, you could have a variation called *Mondays* which is published to the live Site every Monday and another one called *Day 1* which is published to the live Site every first day of each month.

346.3 Using Page Variations

You can also have variations for a single page inside a Site pages variation, which lets you work in parallel on different versions of a page. For example, you might work on two different proposals for the design of the home page for a Holidays Site pages variation. Page variations only exist inside a Site pages variation.

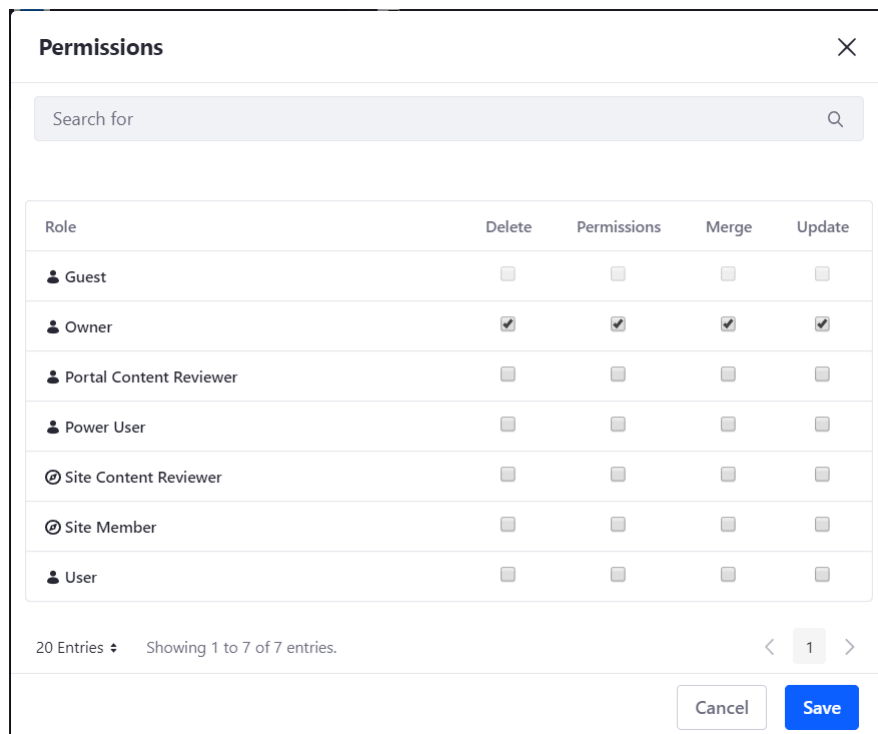
Once you've created a page variation, you can choose it from the dropdown menu on the Staging Bar. You can always switch between different variations by clicking on them from the Staging Bar.

Once you've modified the page variation to the way you want, mark it as *Ready for Publication*. Only one page variation can be marked as ready for publication and that is the one that gets published to the live Site. To publish a variation to the live Site, click on *Publish to Live* in the staging menu and then select *Publish to Live*.

346.4 Managing Variation Permissions

It's also possible to set permissions on each variation, so certain users have access to manage some, but not all variations. To do this,

1. Navigate to the Staging Bar's *Options* button (⋮) and select the variation type you want to configure.
2. Select the desired variation's *Actions* button (⋮) and select *Permissions*.
3. Configure the variation's permissions and then click *Save*.



The screenshot shows a 'Permissions' dialog box with a search bar and a table of roles. The table has columns for Role, Delete, Permissions, Merge, and Update. The 'Owner' role has all permissions checked.

| Role | Delete | Permissions | Merge | Update |
|-------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|
| Guest | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Owner | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |
| Portal Content Reviewer | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Power User | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Site Content Reviewer | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Site Member | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| User | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

20 Entries ▾ Showing 1 to 7 of 7 entries. < 1 >

Cancel Save

Figure 346.3: Configure the roles that can access and modify your variation.

Awesome! You now know how to manage variation permissions!

CREATING MULTI AND SINGLE PAGE VARIATIONS

This document has been updated and ported to Liferay Learn and is no longer maintained here. You can create two types of page variations:

- Site Pages Variation (multiple pages)
- Page Variation (single page)

You can learn more about these variations in the [Using Multi and Single Page Variations](#) article. As an example, you'll step through creating a Site pages variation next. Both variation processes, however, are similar.

1. Select the *Options* icon (⋮) in the Staging Bar and select the variation option. For example, select the *Site Pages Variation* option. This brings you to a list of the existing Site page variations for your Site.
2. Click *Add Site Pages Variation* to create a new one.

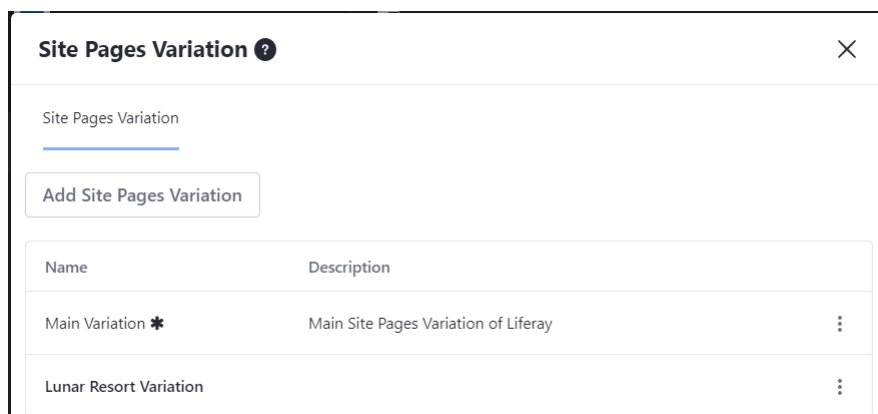


Figure 347.1: When selecting the *Site Pages Variation* link from the Staging Bar, you're able to add and manage your Site pages variations.

3. Set a name and description for your new Site pages variation.

4. Set how you want your variation created. From the *Copy Pages from Site Page Variation* field, you can copy content from an existing variation to create your new one. There are several options to choose in this selector.

All Site Pages Variations: Creates a new variation that contains the last version marked as ready for publication from any single page existing in any other variation.

None (Empty Site Pages Variation): Creates a new, empty variation.

[Existing Variations]: Creates a new Site page variation that contains only the last version of all the pages that exist in a specific variation (e.g., *Main Variation*). The current variation must be marked as ready for publication.

The copy option is not available when creating a page variation. A new page variation is a copy of the current Site pages variation.

5. Click *Add* to create the Site pages variation.

Note: You can rename any variation after it's created, if necessary. For example, edit the *Main Variation* and change its name to something that makes more sense in your Site, such as *Basic*, *Master*, or *Regular*.

Awesome! Your Site pages variation is created and available for modification.

MERGING SITE PAGES VARIATIONS

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Another powerful feature of Staging's Page Versioning framework is the possibility of *merging* Site Pages Variations. To merge two Site Pages Variations, follow the instructions below.

1. Click the Staging Bar's *Options* button (⋮) and select *Site Pages Variation*.
2. Click the Site Pages Variation's *Actions* button (⋮) you want to use as the base for merging and select *Merge*.
3. Select the Site Pages Variation to merge on top of the base Site Pages Variation.

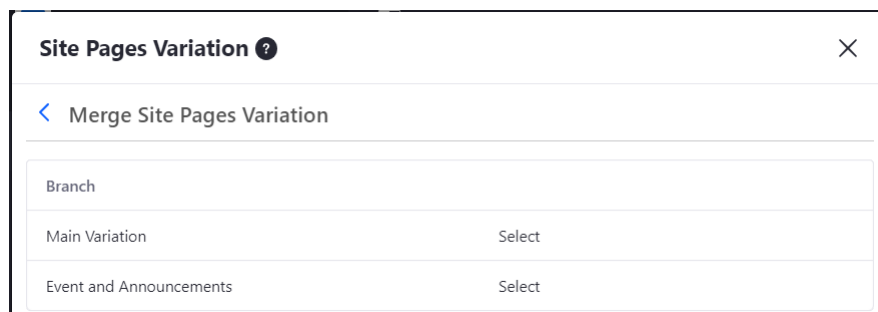


Figure 348.1: Select the site pages variation you'd like to merge with your base variation.

Merging works like this:



- New pages that don't exist in the base variation are added.
- If a page exists in both Site Pages Variations, and at least one version of the page was marked as ready for publication, the latest version marked as ready is added as a new page variation in the target page of the base variation. Note that older versions or page variations not marked as ready for publication aren't copied. Merging can be executed, however, as many times as needed and creates the needed page variations in the appropriate page of the base site pages variation.
- Merging does not affect content and doesn't overwrite anything in the base variation; it adds more versions, pages, and page variations as needed.

Great! You've merged site pages variations!

MANAGING PERMISSIONS

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The staging environment has many different options for building and managing a Site and its pages. Sometimes administrators want to limit access to staging's subset of powerful features. A Role with permissions can accomplish this. To create/modify a Role, complete the following steps:

1. Navigate to the *Control Panel* → *Users* → *Roles*.
2. To create a new Role, select the *Add* button () and complete the New Role menu. Once you have a new Role created, or you've decided on the Role you want to modify, select the Role's *Actions* icon () and select *Edit*.
3. From the top menu, select *Define Permissions*.

The most obvious permissions for staging are the general permissions that look similar to the permissions for most Liferay apps. These permissions are in the *Site Administration* → *Publishing* → *Staging* section of the Define Permissions menu. They include

- *Access in Site Administration*
- *Add to Page*
- *Configuration*
- *Permissions*
- *Preferences*
- *View*

Also, there are some Site resource permissions that deal directly with staging. These permissions are in the *Control Panel* → *Sites* → *Sites* section in the Define Permissions menu. The relevant Site resource permissions related to staging are listed below:

Add Page Variation: Hides/shows the *Add Page Variation* button on the Staging Bar's Manage Page Variations screen.

Add Site Pages Variation: Hides/shows the *Add Site Pages Variation* button on the Staging Bar's Manage Site Page Variations screen.

Export/Import Application Info: If the Publish Staging permission is not granted, hides/shows the application level Export/Import menu. The Configuration permission for the Export/Import app is also required.

Export/Import Pages: If the Publish Staging permission is not granted, hides/shows the Export/Import app in the Site Administration menu.

Manage Staging: Hides/shows the Staging Configuration menu in the Site Administration → *Publishing* → *Staging* → *Options* (ⓘ) menu.

Publish Application Info: Hides/shows the application level Staging menu.

Publish Staging: Hides/shows the *Publish to Live* button on the Staging Bar and hides/shows the *Add Staging Process* button (+) in the Site Administration menu's Staging app. This permission automatically applies the *Export/Import Application Info*, *Export/Import Pages*, and *Publish Application Info* permission functionality regardless of whether they're unselected.

View Staging: If Publish Staging, Manage Pages, Manage Staging, or Update permissions are not granted, hides/shows the Site Administration menu's Staging app.

Notice that some of the permissions above are related to the export/import functionality. Since these permissions are directly affected by the Publish Staging permission, they are important to note. Visit the Importing/Exporting Pages and Content section for more details on importing/exporting Site and page content.

SCHEDULING WEB CONTENT PUBLICATION

Liferay's Web Content framework lets you define when your content goes live. You can determine when the content is displayed, expired, and/or reviewed. This is an excellent way to keep your Site current and free from outdated (and perhaps incorrect) information.

The scheduler is built right into the Properties menu your users access when adding web content. To access this menu, click the *Options* gear (⚙️) and open the Schedule dropdown menu.

The screenshot shows a 'SCHEDULE' dropdown menu with the following options:

- Display Date:** 04/25/2019, 06:25 PM
- Expiration Date:** 04/25/2020, 06:25 PM
- Never Expire**
- Review Date:** 01/25/2020, 06:25 PM
- Never Review**

Figure 350.1: The web content scheduler can be easily accessed from the right panel of the page.

The scheduler offers several configurable options:

Display Date: Sets (within a minute) when content will be displayed.

Expiration Date: Sets a date to expire the content. The default is one year.

Never Expire: Sets your content to never expire.

Review Date: Sets a content review date.

Never Review: Sets the content to never be reviewed.

As an example, you'll step through the process of scheduling a web content article.

1. Navigate to the Product Menu → *Content & Data* → *Web Content*.
2. Create a new web content article by selecting the *Add Web Content* button (+) → *Basic Web Content*.
3. Add content for your web content article.

4. Select the *Schedule* tab from the web content's Properties menu. Configure the publication schedule.
5. Click *Publish*. Your web content article is now created and abides by the scheduling parameters you've set.

When you set a Display Date for an existing article it does not affect previous versions of the article. If a previous version is published, it remains the same until the new version is scheduled to display. However, the expiration date affects all versions of the article. Once an article has expired, no version of that article appears.

Tip: If you want only the latest version of articles to expire, and not every past version, go to *Control Panel* → *Configuration* → *System Settings* → *Web Content* → *Virtual Instance Scope* → *Web Content* and uncheck *Expire All Article Versions Enabled*. This makes the previously approved version of an article appear if the latest version expires.

The scheduling feature gives you great control in managing when, and for how long, your web content is displayed on your Site. Additionally, you can determine when your content should be reviewed for accuracy and/or relevance. This makes it possible to manage your growing inventory of content.

MANAGING APPS

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Apps run on the platform Liferay DXP provides. The web experience management, collaboration, and business productivity features all consist of apps. Even the Control Panel consists of configuration apps. You can also add to or change built-in functionality by installing other apps. There are several ways to manage, find, and install apps. This section covers these topics and more.

MANAGING AND CONFIGURING APPS

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay DXP is a platform for deploying apps that comprise modules and components. It has conveniences for managing apps and app management best practices for maximizing stability. Best practices in production environments involve stopping the server before applying changes, but in cases where this isn't feasible, you can “auto deploy” changes several different ways.

There are two places in the Control Panel where you can manage and configure apps: the *App Manager* and the *Components* listing. The App Manager manages apps in the OSGi framework. You can use the App Manager to install, activate, deactivate, and delete apps. You can manage apps at the app and module levels.

The Components listing views and manages apps at the OSGi component level. It differs from the App Manager by showing apps by type (portlet, theme, and layout template), and setting app permissions. You can use the Components listing to activate and deactivate apps, but it can't install or delete apps.

Start with learning app management best practices in production, or wherever you want to maximize stability.

352.1 Managing Apps in Production

Not all apps are designed to be “auto deployed”—deployed while the server is running. Deploying that way can cause instabilities, such as class loading leaks and memory leaks. On production systems, avoid “auto deploying” apps and configurations whenever possible.

If you're installing an app or a component configuration on a production system and stopping the server is feasible, follow these steps:

1. Stop your server.
2. Copy your app (.lpkg, module .jar, or plugin .war) to your [Liferay Home]/deploy folder, or copy your component configuration (.config file) to the [Liferay Home]/osgi/configs folder. The Liferay Home folder is typically the app server's parent folder.

3. Start your server.

If you're running in cluster, follow the instructions for updating a cluster.

Warning: Avoid repeatedly “auto deploying” new versions of apps that aren't designed for “auto deployment”.

If it's not feasible to stop your server or you're app is designed for “auto deployment”, Liferay DXP provides several “auto deployment” conveniences. Except where stopping/starting the server is explicitly mentioned, the practices described in the rest of this article and in the following articles involve “auto deployment”.

352.2 Using the App Manager

Access the App Manager by selecting *Control Panel* → *Apps* → *App Manager*. The App Manager lists your apps. The *Filter and Order* menu lets you filter and order by category, status, or title. Click the up or down arrows to perform an ascending or descending sort, respectively. To search for an app or module, use the search bar. This is often the quickest way to find something.

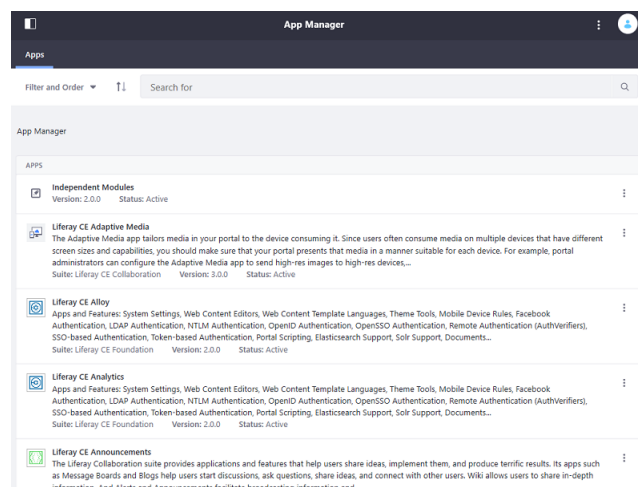



Figure 352.1: The App Manager lets you manage the apps, modules, and components installed in your Liferay DXP instance.

Each item listed in the table contains a description (if available), version, and status. Here are the statuses:

- **Installed:** The item is installed to Liferay DXP.
- **Resolved:** The item's dependencies are active. Resolved items can typically be activated. Some items, however, can't be activated and are intended to remain in the Resolved state (e.g., WSDD modules containing SOAP web services).
- **Active:** The item is running in Liferay DXP.


Clicking each item's Actions button () brings up a menu that lets you activate, deactivate, or uninstall that item.

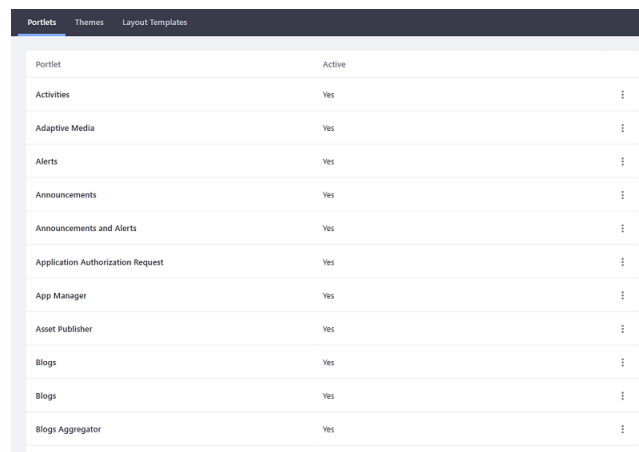
To view an item's contents, click its name in the table. If you click an app, the app's modules are listed. If you click a module, the module's components and portlets appear. The component level is as far down as you can go without getting into the source code. At any level in the App Manager, a link trail appears that lets you navigate back in the hierarchy.

For information on using the App Manager to install an app, see [Installing Apps Manually](#).

Next, you'll learn how to use the Components listing.

352.3 Using the Components Listing

Access the components listing by selecting *Control Panel* → *Configuration* → *Components*. The components listing first shows a table containing a list of installed portlets. Select the type of component to view—portlets, themes, or layout templates—by clicking the matching tab on top of the table. To configure a component, select its name in the table or select *Edit* from its Actions button (). Doing either opens the same configuration screen.














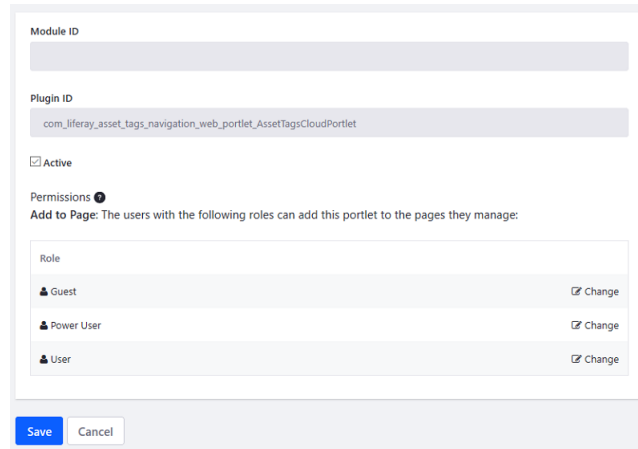
| Portlet | Active | |
|-----------------------------------|--------|---|
| Activities | Yes |  |
| Adaptive Media | Yes |  |
| Alerts | Yes |  |
| Announcements | Yes |  |
| Announcements and Alerts | Yes |  |
| Application Authorization Request | Yes |  |
| App Manager | Yes |  |
| Asset Publisher | Yes |  |
| Blogs | Yes |  |
| Blogs | Yes |  |
| Blog: Aggregator | Yes |  |

Figure 352.2: The components listing lets you manage the portlets, themes, and layout templates installed in your Liferay DXP instance.

The configuration screen lets you view a component's module ID and plugin ID, activate or deactivate the component, and change the component's Add to Page permission. The component's module ID and plugin ID appear at the top of the screen. You can activate or deactivate a component by checking or unchecking the *Active* checkbox, respectively. To change a component's Add to Page permission for a role, select the role's *Change* button in the permissions table. This takes you to *Control Panel* → *Users* → *Roles*, where you can change the component's permissions for the selected role.



The screenshot shows a configuration dialog for a component. It includes fields for Module ID and Plugin ID, an Active checkbox, and a Permissions table. The Active checkbox is checked. The Permissions table lists three roles: Guest, Power User, and User, each with a 'Change' link.

| Role | |
|------------|------------------------|
| Guest | Change |
| Power User | Change |
| User | Change |

Figure 352.3: You can activate or deactivate a component, and change its permissions.

USING THE LIFERAY MARKETPLACE

This document has been updated and ported to Liferay Learn and is no longer maintained here. Liferay Marketplace is a hub for sharing, browsing, and downloading apps. Marketplace leverages the entire Liferay ecosystem to release and share apps in a user-friendly, one-stop shop. There are two ways to access the Marketplace.

1. **Via the website:** Using your favorite browser, navigate to the Marketplace at web.liferay.com/marketplace. If you're new to Marketplace, this is the easiest way to access it. You can browse Marketplace without signing in with your liferay.com account.
2. **Via the Control Panel:** In the Control Panel, navigate to *Apps* → *Store*. To view Marketplace, you must sign in with your liferay.com account.

No matter how you access Marketplace, you'll see the same content. Note that to download apps, you must have a liferay.com account and agree to the Marketplace Terms of Use.

Here you'll learn how to,

- Find and purchase apps
- Manage purchased apps
- Renew purchased apps

Start with finding and purchasing the apps you want.

353.1 Finding and Purchasing Apps

If you've used an app store before, Marketplace should be familiar. You'll see apps in the center of the page, in the following sections:

- **Featured Apps:** Liferay features a different set of apps each month.

- **New and Interesting:** The latest apps added to Marketplace.
- **Most Viewed in the Past Month:** The top 5 most viewed apps in the last month.
- **Themes / Site Templates:** Apps that change your Liferay instance’s look and feel.
- **App categories:** Communication, productivity, security, etc.
- **Weekly Stats:** The newest apps, latest apps updated, and trend chart for app downloads and views.

Each section’s *See All* link shows more section info. At the top of the page, you can search Marketplace by category, Liferay DXP version, and price. To browse by category, click the *Categories* menu at the top of the page.

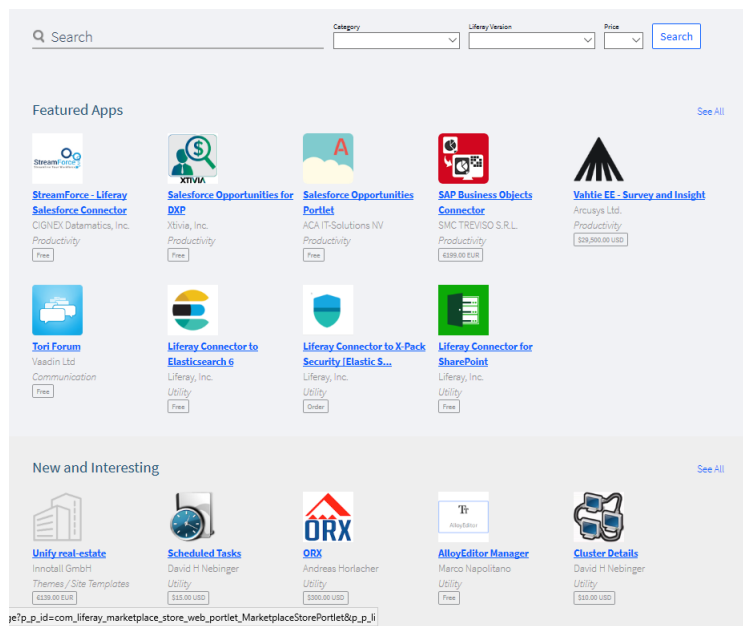


Figure 353.1: The Liferay Marketplace home page lets you browse and search for apps.

Click an app to view its details. This includes its description, screenshots, price, latest version, number of downloads, a link to the developer’s website, a link to the app’s license agreement, and a purchase button (labeled Free or Buy, depending on the price). You can also view the app’s version history, read reviews, or write your own review.

The purchase button prompts you to choose a purchase type. You can purchase an app for your personal account, or for a Liferay project associated with your company. If you have the necessary permissions, you can also create a new project for your company. Once you select a purchase type, accept the EULA and Terms of Service, and click *Purchase*.

Once you purchase an app, you can download and install it.

Warning: Not all apps are designed to be “auto deployed”—deployed while the server is running. Deploying that way can cause instabilities, such as class loading leaks and memory leaks. On production systems, avoid “auto deploying” apps whenever possible. See the best practices for managing apps in production.

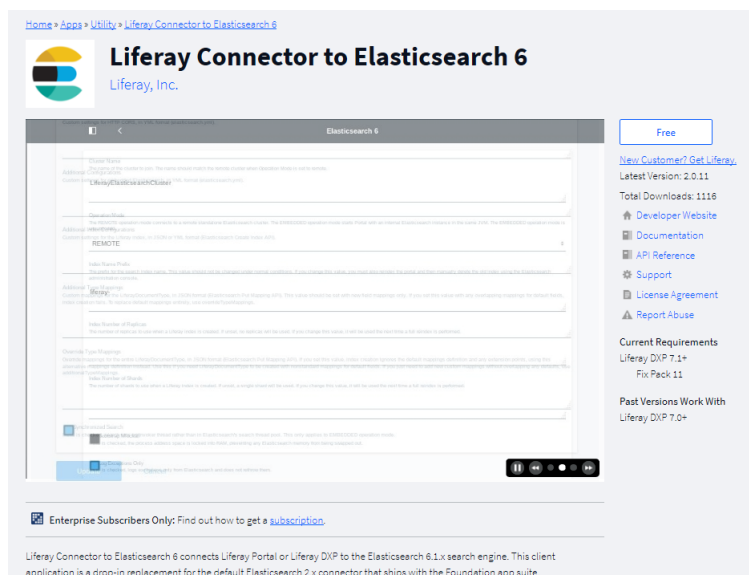


Figure 353.2: Click an app to view its details.

An app downloads and installs immediately if you purchase it from the Control Panel. If you purchase the app on the Marketplace website, however, your receipt is displayed immediately after purchase. To download the app, click the *See Purchased* button on the bottom of the receipt, and then click the *App* button to start the download. You must then install the app manually. Alternatively, you can use Marketplace from the Control Panel to download and install the app after purchase on the Marketplace website. The next section shows you how to do this.

Note that sometimes administrators disable automatic app installations so they can manage installations manually. In this case, Marketplace apps downloaded from the Control Panel are placed in the deploy folder in Liferay Home. Administrators must then manually install the app from this folder. Manual install is also required if the server is behind a corporate firewall or otherwise lacks direct Marketplace access. Regardless of how the app is downloaded, the manual install process is the same. For details, see the article *Installing Apps Manually*.

353.2 Managing Purchased Apps

Important: When uninstalling an app or module, make sure to use the same agent you used to install the app. For example, if you installed it with Marketplace, uninstall it with Marketplace. If you installed it with the file system, use the file system to uninstall it. If you installed it with the App Manager, however, use Blacklisting to uninstall it.

There are two places to manage your purchased apps:

1. Your liferay.com account's home page. After signing in, click the user menu at the top-right

and select *Account Home*. Note that your home page is distinct from your profile page. Your home page is private, while your profile page is public. On your home page, select *Apps* from the menu on the left to view your projects. Select a project to view its registered apps. Clicking an app lets you view its versions. You can download the version of the app that you need. This is especially useful if you need a previous version of the app, or can't download the app from the Control Panel.




Figure 353.3: You can manage your purchased apps from your liferay.com account's home page.

- From the Control Panel. Navigate to *Apps* → *Purchased* to see your purchased apps. A button next to each app lets you install or uninstall the app. If the app isn't compatible with your Liferay DXP version, *Not Compatible* is displayed in place of the button. Additional compatibility notes are also shown, such as whether a newer version of the app is available. You can also search for an app here by project, category, and title. Clicking the app takes you to its Marketplace entry.

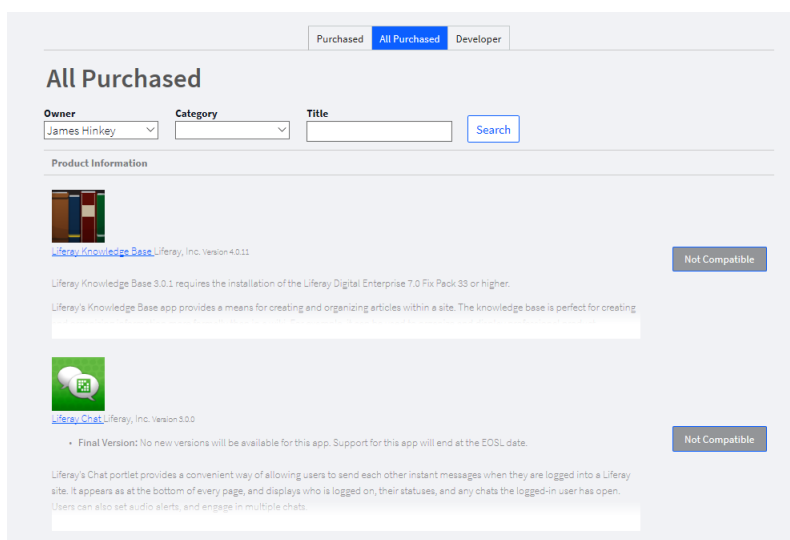


Figure 353.4: You can also manage your purchased apps from within a running Liferay instance.

353.3 Renewing a Purchased App

To continue using a purchased app whose license terms are non-perpetual, you must renew your app subscription, register your server to use the app, and generate a new activation key to use on your server. Here are the steps:

1. Go to <https://web.liferay.com/marketplace>.
2. Click your profile picture in the upper right corner and select *Purchased Apps*. The Purchased Apps page appears and shows your app icons organized by project.
3. Click your app's icon. Your app's details page appears.
4. Click *Manage Licenses*.
5. Select *Register New Server*.
6. Select the most recent *Order ID* (typically the order that has no registered servers).
7. Fill in your server's details.
8. Click *Register*.
9. Click *Download*. The new app activation key to use on your server downloads.
10. Copy the activation key file to your `deploy/` folder in your [Liferay Home].

You can continue using the application on your server.

INSTALLING APPS MANUALLY

This document has been updated and ported to Liferay Learn and is no longer maintained here.

By default, apps you download from the Control Panel via Liferay Marketplace install automatically. But what if the app you want to install isn't on Marketplace? What if all you have is the app's file? In this case, you must install the app manually. Here you'll learn how to install any app manually.

Warning: Not all apps are designed to be “auto deployed”—deployed while the server is running. Deploying that way can cause instabilities, such as class loading leaks and memory leaks. On production systems, avoid “auto deploying” apps whenever possible. See the best practices for managing apps in production.

Important: When uninstalling an app or module, make sure to use the same agent you used to install the app. For example, if you installed it with Marketplace, uninstall it with Marketplace. If you installed it with the file system, use the file system to uninstall it. If you installed it with the App Manager, however, use Blacklisting to uninstall it.

354.1 Using the Control Panel to Install Apps

To install an app manually from the Control Panel, navigate to *Control Panel* → *Apps* → *App Manager*, and select *Upload* from the options button (ⓘ). In the Upload dialog, choose the app on your machine and then click *Install*. When the install completes, close the dialog and you're ready to roll!

354.2 Using Your File System to Install Apps

To install an app manually on the Liferay DXP server, put the app in the [Liferay Home]/deploy folder (the Liferay Home folder is typically the app server's parent folder). That's it. The auto deploy mechanism takes care of the rest.

You might now be thinking, "Whoa there! What do you mean by 'the rest?' What exactly happens here? And what if my app server doesn't support auto deploy?" These are fantastic questions! When you put an app in the [Liferay Home]/deploy folder, the OSGi container deploys the app to the appropriate subfolder in [Liferay Home]/osgi. By default, the following subfolders are used for apps matching the indicated file type:

- marketplace: Marketplace LPKG packages
- modules: OSGi modules
- war: WAR files

You can, however, change these subfolders by setting the properties `module.framework.base.dir` and `module.framework.auto.deploy.dirs` in a `portal-ext.properties` file. These properties define the [Liferay Home]/osgi folder and its auto deploy subfolders, respectively. The default settings for these properties in the `portal.properties` file are as follows:

```
module.framework.base.dir=${liferay.home}/osgi

module.framework.auto.deploy.dirs=\
  ${module.framework.base.dir}/configs,\
  ${module.framework.base.dir}/marketplace,\
  ${module.framework.base.dir}/modules,\
  ${module.framework.base.dir}/war
```

Note that the `configs` subfolder isn't for apps: it's for configuration files imported from other Liferay DXP instances.

But what happens if your app server doesn't support "hot deploy"? No problem! Liferay DXP's module framework (OSGi) enables auto deploy. Any app server running Liferay DXP therefore also supports this auto deploy mechanism.

354.3 Manually Deploying an LPKG App

When manually installing an LPKG app, the installation may hang with a server log message like this:

```
14:00:15,789 INFO [com.liferay.portal.kernel.deploy.auto.AutoDeployScanner][AutoDeployDir:252] Processing Liferay Push 2.1.0.lpkg
```

This happens when LPKG apps have the `restart-required=true` property in their `liferay-marketplace.properties` file (inside the LPKG file). This property setting specifies that a server restart is required to complete the installation.

APP TYPES

This document has been updated and ported to Liferay Learn and is no longer maintained here.

There are several different kinds of apps. Some apps can even contain other apps. The types of apps you can install include:

- OSGi Modules
- Portlets
- Web Plugins
- Templates
- Themes

Read on to learn about these app types.

355.1 OSGi Modules

Since Liferay DXP runs on OSGi, apps can be implemented as OSGi modules. An OSGi module is a JAR file adapted to run on OSGi. Although it's possible for a single module to implement a single app, an app typically consists of multiple modules that are packaged together. Also note that apps in OSGi modules aren't required to have a UI. For example, Liferay DXP can run OSGi modules that expand built-in APIs without requiring any user interaction. This is crucial for developers that must leverage custom APIs. By providing such an API via one or more OSGi modules, you can let developers leverage your API.

OSGi modules can also contain apps that have a UI: portlets. The next section discusses these.

355.2 Portlets

Portlets are small web applications that run in a portion of a web page. For example, the built-in Blogs app is a portlet. Portlet applications, like servlet applications, are a Java standard implemented by various portal server vendors. The JSR-168 standard defines the portlet 1.0 specification, the JSR-286 standard defines the portlet 2.0 specification, and the JSR-362 standard defines the portlet 3.0 specification. A Java standard portlet should be deployable on any portlet container that supports the standard. Portlets are placed on the page in a certain order by the end user and are served up dynamically by the portal server. This means certain things that apply to servlet-based projects, such as control over URLs or access to the `HttpServletRequest` object, don't apply in portlet projects because the portal server generates these objects dynamically.

Portlets can be composed of OSGi modules (recommended), or contained in WAR files. For information on developing portlets see *Web Front-ends*.

355.3 Web Plugins

Web plugins are regular Java EE web modules designed to work with Liferay DXP. You can integrate with various Enterprise Service Bus (ESB) implementations, as well as Single Sign-On implementations, workflow engines, and so on. These are implemented as web modules used by Liferay DXP portlets to provide functionality.

355.4 Templates and Themes

Templates and themes are plugins that change Liferay DXP's appearance. Templates (layout templates) control how you can arrange portlets on a page. They make up a page's body (the large area into which you can drag and drop portlets). There are several built-in layout templates. If you have a complex page layout (especially for your home page), you may wish to create a custom layout template of your own.

Themes can completely transform Liferay DXP's look and feel. Most organizations have their own look and feel standards that apply to all of their web sites and applications. By using a theme plugin, an organization can apply these standards on Liferay DXP. There are many available theme plugins on Liferay's web site and more are being added every day. This makes it easy for theme developers, as they can customize existing themes instead of writing a new one from scratch.

355.5 Liferay Marketplace App Packages

Regardless of app type, each Liferay Marketplace app is distributed in an LPKG package. The LPKG package contains Marketplace metadata and the files the app needs to run. Note that it's possible for an LPKG package to contain multiple apps. For example, a single LPKG package can contain several portlets. This is common in cases where an app requires a Control Panel portlet for administrators, and another portlet for end users.

BLACKLISTING OSGI BUNDLES AND COMPONENTS


This document has been updated and ported to Liferay Learn and is no longer maintained here.

Blacklists are used for good and evil. An evil blacklist penalizes unfairly; a good blacklist protects. Liferay DXP's OSGi bundle and component blacklists are files that prevent particular bundles from installing and particular components from enabling. This saves you the trouble of uninstalling and disabling them individually with the Application Manager, Components list, or Gogo shell.

356.1 Blacklisting Bundles

Liferay DXP removes any installed OSGi bundles on the blacklist. Blacklisted bundles therefore can't be installed. The log reports each bundle uninstallation.

Follow these steps to blacklist bundles:

1. In the Control Panel, navigate to *Configuration* → *System Settings* → *Module Container*.
2. In the Bundle Blacklist screen, add the bundle symbolic names (see the table below) for the Module JARs, LPKG files, or WARs to uninstall. Click the *Save* button when you're finished.
3. To export the blacklist, click its Actions button () and then click *Export*. The blacklist configuration file then downloads (`com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.com`). Here are contents from an example file:

```
blacklistBundleSymbolicNames=["com.liferay.docs.greeting.api","Liferay\ Marketplace","classic-theme"]
```

4. Add the bundle symbolic names of any bundles not already listed that you want to prevent from installing.

Important: Configuration values can't contain extra spaces. Extra spaces can short-circuit lists or invalidate the configuration entry.



Figure 356.1: This blacklist uninstalls the `com.liferay.docs.greeting.api` bundle, Liferay Marketplace LPKG, and `classic-theme` WAR.

- To deploy the configuration file, copy it into the folder `[Liferay_Home]/osgi/configs`. The Liferay Home folder is typically the app server's parent folder.

Note: Blacklisting an LPKG uninstalls all of its internal bundles.

Blacklist Bundle Symbolic Names

| Type | Bundle Symbolic Name |
|------------|---|
| Module JAR | Bundle-SymbolicName in <code>bnd.bnd</code> or <code>MANIFEST.MF</code> file |
| LPKG | LPKG file name without the <code>.lpkg</code> extension |
| WAR | Servlet context name in <code>liferay-plugin-package.properties</code> file or the WAR file name (minus <code>.war</code>), if there is no servlet context name property |

356.2 Reinstalling Blacklisted Bundles

To reinstall and permit installation of blacklisted OSGi bundles, follow these steps:

1. Open the configuration file `com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.config`.
2. Remove the symbolic names of the module JARs, LPKGs, or WARs from the `blacklistBundleSymbolicNames` list and save the file.

To reinstall *all* the blacklisted bundles execute one of these options:


- Remove the configuration file.
- Uninstall the bundle `com.liferay.portal.bundle.blacklist` using the Application Manager or Felix Gogo Shell.

Note: To temporarily reinstall a bundle that's been blacklisted, you can remove its symbolic name from the Bundle Blacklist module in *System Settings* and click the *Update* button. If you want the bundle to install on subsequent server startup, make sure to remove the bundle's symbolic name from any existing blacklist configuration file in the `[Liferay_Home]/osgi/configs` folder.

The log reports each bundle installation.

356.3 Blacklisting Components

Follow these steps to blacklist components:

1. In the Control Panel, navigate to *Configuration* → *System Settings* → *Module Container*.
2. In the Component Blacklist screen, add the names of components to disable and click the *Save* button.
3. To export the blacklist, click on the Component Blacklist module's Actions button () and then click *Export*. The blacklist configuration file then downloads (`com.liferay.portal.component.blacklist.internal.ComponentBlacklistConfiguration.config`). Here are contents from an example file:

```
blacklistComponentNames=["com.liferay.portal.security.ldap.internal.authenticator.LDAPAuth", "com.liferay.ip.geocoder.sample.web.internal.port
```

4. Add the names of any components not already listed (e.g., components of bundles not yet installed) that you want to prevent from activating.

Important: Configuration values can't contain extra spaces. Extra spaces can short-circuit lists or invalidate the configuration entry.
5. To deploy the configuration file, copy it into the folder `[Liferay_Home]/osgi/configs`. The Liferay Home folder is typically the app server's parent folder.

Component Blacklist Configuration

This configuration was not saved yet. The values shown are the default.

Blacklist Component Names +

OSGi component name used to uniquely identify a component for the system to disable.

```
com.liferay.portal.security.ldap.internal.authenticator.LDAPAuth
```

Blacklist Component Names - +

OSGi component name used to uniquely identify a component for the system to disable.

```
com.liferay.ip.geocoder.sample.web.internal.portlet.IPGeocoderSamplePortlet
```

Figure 356.2: This blacklist disables the components `com.liferay.portal.security.ldap.internal.authenticator.LDAPAuth` and `com.liferay.ip.geocoder.sample.web.internal.portlet.IPGeocoderSamplePortlet`.

356.4 Re-enabling Blacklisted Components

To re-enable and permit enabling of blacklisted components, follow these steps:

1. Open the configuration file `com.liferay.portal.component.blacklist.internal.ComponentBlacklistConfiguration`.
2. Remove the names of the components from the `blacklistComponentNames` list and save the file.

To enable *all* the blacklisted components, remove the configuration file.

Note: To temporarily reactivate a blacklisted component, remove its name from the Component Blacklist Configuration module in System Settings and click *Update*. If you want the component to activate on subsequent server startup, make sure to remove the component's name from any existing component blacklist configuration file in the `[Liferay_Home]/osgi/configs` folder.

POLLS


How can The Lunar Resort stay connected with its earthbound clientèle from 239,000 miles away? Make them feel really involved and enthusiastic about the resort by asking them for feedback. You're not just creating a poll, you're making connections.

Use Polls to find out what your site visitors are thinking and keep them engaged with your site's content.

Two applications make and display a poll: the *Polls* application in the Site Menu and the *Polls Display* widget you add to a page.

357.1 Creating a Poll

From the Site Menu, go to *Content & Data* → *Polls*.

1. Click the  button and fill out the form.

Title: (Required) Enter the name of the poll question.

Polls Question: (Required) Enter the text of the poll question.

Expiration Date: Enter the date and time you want the poll to expire.

Choices: Enter at least two options for the poll question.

Add Choice: Enter additional answer options for the poll question.

Permissions: Manage who can view and edit the poll.

2. Click *Save* to add the poll to the Polls application.

Once a poll is created, the Polls Display portlet publishes it until it expires or is deleted. Set an expiration date for a poll by selecting the day and time in the Add Poll form. The default is set to *Never Expire*.

When a published poll expires, the poll results are displayed, but users can't add new entries. To remove an expired poll from a page, remove the Poll Display portlet or configure it to show another poll question. See the section below for more details about the Polls Display portlet.

The screenshot shows a mobile application interface for creating a new poll. At the top, there is a dark header with a back arrow, the text "New Poll", and a user profile icon. Below the header is a form with several sections:

- Title ***: A text input field with a red border and a "This field is required." message below it. A language selector icon (en-US) is on the right.
- Polls Question ***: A larger text input field with a red border and a "This field is required." message below it. A language selector icon (en-US) is on the right.
- Expiration Date**: Two input fields showing "07/01/2019" and "01:35 AM".
- Never Expire**: A checked checkbox.
- Choices**: Two input fields labeled "a." and "b.", each with a language selector icon (en-US) on the right.
- Add Choice**: A button below the choice fields.
- Permissions**: A section with a right-pointing arrow.
- Save** and **Cancel**: Two buttons at the bottom of the form.

Figure 357.1: Besides the Title and the Polls Question, you must enter data for each of the Choices fields when creating a new poll.

Permissions are set on individual polls. Use permissions, for example, to allow some privileged users to vote on a certain poll question, while others can only view it.

Creating a poll is fairly straightforward. Next, complete the two-step process and put your poll on a page.

357.2 Adding a Poll to a Page

Now that you have created your poll question, you can present it to your users:

1. Go to a page and add a Polls Display widget from *Add* → *Widgets* → *Content Management*.

2. Click *Please configure this portlet to make it visible to all users*.
3. In the dialog box that appears, select the poll to display.
4. Click *Save*.

Once the poll question has been placed on the page, perform other tasks using the icons at the bottom of the portlet.

LUNAR RESORT POLL

Which word best describes the Lunar Resort?

a. Fun

b. Wacky

c. Inspiring

d. Stunning

 [Select Poll](#)  [Add](#)  [Edit Question](#)

Figure 357.2: These buttons provide shortcuts to the widget's configuration, as well as to some of the Polls Application's functionality.

Edit Question: Displays a similar dialog box to the one used to create the poll.

Select Poll: Displays the same dialog box as Configuration, allowing you to choose different polls from the drop-down menu.

Add: Allows you to create a new poll.

357.3 Viewing Poll Results

All the polls you create appear in the Polls portlet in the Site Menu → *Content & Data* → *Polls*. When users vote in the poll, the data is collected here. Click on a poll to see a breakdown of the results.



Figure 357.3: Selecting a poll in the Polls portlet puts the data at your fingertips.

If you click on one of the listed *Charts*, the portlet generates an appropriate visualization of the data.

Below this is an item called *Charts*. This option shows the poll results represented in various graphs. The graphs are *Area*, *Horizontal Bar*, *Line*, *Pie*, and *Vertical Bar*.

There is also a listing of the users who voted in the poll, how they voted, and a time/date stamp of when their votes were cast. Registered users are represented by name. Guest users have a blank *User* field.





Figure 357.4: This is what the vertical bar graph for the Lunar Resort poll results looks like.

USING THE CALENDAR

The Calendar widget is an updated, digitized, 3D-printed sundial. Okay, it's really a tool for storing and sharing scheduled events. It's a personal planner for individual users, a shared calendar for an entire site, or both at the same time. It can be used to create multiple calendars for a single Site or User, to overlay the events stored in multiple calendars for simultaneous view, to send email reminders to users, and more.

Note: The calendar supports social activities. Whenever a calendar event is added or updated, a corresponding social activity notification is created. If the event was added or updated in a calendar that the current user has permission to view, the social activity is viewable in the Activities widget.

358.1 Configuring the Calendar

Once the Calendar widget is on a page, open the  menu in the widget's header and click *Configuration*.

From the *User Settings* tab, customize the calendar's default view and settings.

Time Format: Choose *Locale*, *AM/PM*, or *24 Hour*. *Locale* is a dynamic setting that chooses whether to display the time in *AM/PM* or *24 Hour* format, based on the preferences set by the User's locale. *AM/PM* displays times such as 8AM or 11PM. The *24 Hour* time format displays times such as 08:00 and 23:00.

Default Duration: Choose an event duration. When you add a new event to the calendar, the time you set here specifies how long events last by default.

Default View: Choose *Day*, *Week*, *Month* or *Agenda*. This sets the default for when the calendar is first displayed, but the view can be changed by clicking the appropriate button at the top-right of the widget.

Week Starts On: Choose *Sunday*, *Monday*, or *Saturday*.

Time Zone: Choose a time zone or check the *Use Global Time Zone* box.

The screenshot shows the 'User Settings' tab selected. The settings are as follows:

- Time Format:** Locale
- Default Duration:** 60 Minutes
- Default View:** Week
- Week Starts On:** Sunday
- Time Zone:** (UTC) Coordinated Universal Time
- Use Global Time Zone

Figure 358.1: The Setup → User Settings tab provides the options you need to get started quickly.

If you check *Use Global Time Zone*, the time displayed depends on whether it's being viewed by a logged-in user or a guest. If a user is logged in, the Calendar displays events using the time zone set for the user in *User Personal Menu* → *Account Settings* → *Preferences* → *Display Settings* → *Time Zone*. If the Calendar is viewed by a guest or a user who is not logged in, the Calendar displays events using the time zone set by the portal administrator in *Control Panel* → *Configuration* → *Instance Settings* → *Platform* → *Localization* → *VIRTUAL INSTANCE SCOPE* → *Time Zone*.

From the *Display Settings* tab, set the display behavior for the calendar.

Display Scheduler Only: By default, the list of calendars and a mini-calendar view (used for quickly navigating to a particular date) are displayed. Check this to display only the scheduler (the large calendar view showing the calendar and scheduled events).

Display User Events: Turns off the display of the current, logged in User's personal calendar and events.

Display Scheduler's Header: If disabled, removes the ability to toggle through the calendar views (for example, Day/Week/Month/Agenda) and access to the Add Event button.

Enabled Views: If one of the available views is disabled (Day, Week, Month, Agenda), it disappears from the scheduler's header.

Maximum Days to Display: Set the maximum number of days to display in the Agenda view.

Maximum Events to Display: Set the maximum number of events to display in the Agenda view.

Use the *RSS* tab to disable RSS subscription or configure the RSS behavior.

Enough with configuration. Next you'll learn how to use it.

USING THE CALENDAR WIDGET

The calendar widget displays a small monthly calendar showing an overview of upcoming events. A larger area shows the Scheduler, a more detailed calendar with a number of options: you can set it to display a *Day*, *Week*, or *Month*, or choose a more event-oriented *Agenda* setting.




Figure 359.1: The default view is set in configuration, but a user can change it at any time.

Two calendars are included by default when the widget is first added to a page: a personal calendar for the current user and a Site calendar for the current Site. These are displayed in the

widget's lower left. Next to each calendar is a colored box: click it to show/hide that calendar's events in the main viewing area.

359.1 Adding New Calendars

To create a new personal calendar,

1. Click on the arrow to the right of the *My Calendars* header and select *Add Calendar* from the menu.
2. Fill in the *Add Calendar* form. Give the calendar a name and description, set a time zone, and decide if it's your user's *default calendar*—the one that is shown automatically whenever the widget is displayed. You can also pick a color, which color codes events whenever multiple calendar's events are displayed at once. You can also decide to enable ratings or comments on the calendar's events, and configure permissions.

To edit an existing calendar instead of adding a new one, select *Manage Calendars* from the menu.

To add or edit a Site calendar, open the menu next to the header with the Site's name.

359.2 Adding Events to a Calendar

To add events to a calendar,

1. Click on any day in the main viewing area to open an event creation pop up. If you've selected the *Day* or *Week* view, you can click on the specific time when your event begins.
2. Name your event and assign it to a calendar. Click *Save* to create the event immediately or *Edit* to enter additional information.
3. If you clicked *Edit*, complete the *Edit Event* form. Enter start and end times and enter a description. To schedule an event that reoccurs, check the *Repeat* box and fill in the *Repeat* pop up.

359.3 Additional Event Functions

At the bottom of the *Edit Event* form, there are several collapsed sections: *Details*, *Invitations*, *Reminders*, *Categorization*, and *Related Assets*.

Calendar
Resources

◀
June 2019
▶

Su
Mo
Tu
We
Th
Fr
Sa

| | | | | | | |
|----|----|----|----|----|----|----|
| | | | | | | 1 |
| 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 9 | 10 | 11 | 12 | 13 | 14 | 15 |
| 16 | 17 | 18 | 19 | 20 | 21 | 22 |
| 23 | 24 | 25 | 26 | 27 | 28 | 29 |
| 30 | | | | | | |

- ▼ My Calendars ▼
- Test Test ▼
- Personal Schedule ▼
- ▼ Liferay's Calendars ▼
- Liferay ▼

▼ Other Calendars

Add Event

◀
Today
▶

| | |
|-----|-----|
| Sun | Mon |
| 26 | 27 |
| 2 | 3 |
| 9 | 10 |
| 16 | 17 |
| 23 | 24 |

Figure 359.2: Personal and Site calendars are shown in the lower left. This image shows calendars belonging to User *Test Test* and Site *Liferay DXP*.




Figure 359.3: When you click anywhere on the calendar, you'll see the event creation pop up appear. Click *Edit* to specify details for your event.

Details

In the Details section, you can move the event to another calendar and enter a location.

Invitations

In the invitations section, invite Users, Sites, or Calendar Resources (see the next tutorial for more on resources: in brief, a resource is anything you might need for an event—a conference room, a vehicle, etc.). Follow these steps:

1. Enter the name of an invitee (User, Site, or Resource) in the *Invitations* field. Hit *Enter* to add them to the *Pending* column.
2. Check the availability of invitees by clicking the arrow next to their names and selecting *Check Availability*. This displays their calendars (assuming you have permission to view them).

An automated email is sent to invitees who must navigate to the calendar widget to respond. See below to customize the content of the invitation.

When invitees respond to the invitation, their names move to the *Accepted*, *Declined*, or *Maybe* columns.

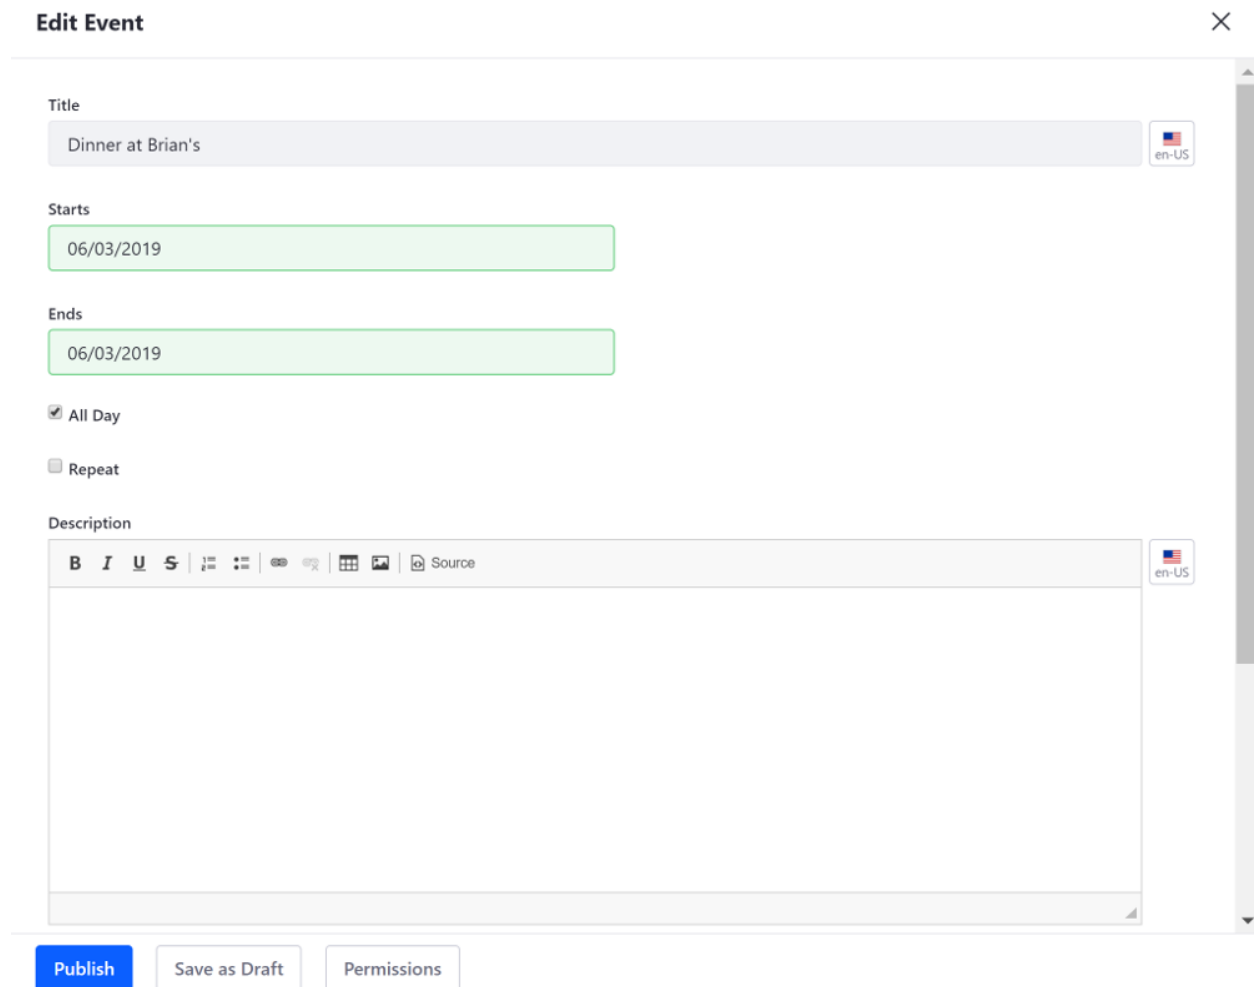
Reminders

Schedule up to two email reminders to send to attendees. Reminders translate the time of the event into the recipients own time zone.

See below to customize the content of the reminder email.

Categorization

Tag your event or assign it to a category so it appears in appropriate search results and is published by any asset publisher set to publish content assigned to the same category.



The screenshot shows the 'Edit Event' form with the following details:

- Title:** Dinner at Brian's
- Starts:** 06/03/2019
- Ends:** 06/03/2019
- All Day:**
- Repeat:**
- Description:** A rich text editor with a toolbar containing icons for Bold (B), Italic (I), Underline (U), Strikethrough (S), Bulleted List, Numbered List, Link, Unlink, Table, Image, and Source. The text area is currently empty.

At the bottom of the form are three buttons: **Publish** (highlighted in blue), **Save as Draft**, and **Permissions**. A language selector 'en-US' is visible in the top right corner of the form area.

Figure 359.4: You can specify event details such as the event title, start date, end date, description, location, and more.

Related Assets

List an asset—such as an agenda or supplementary material for a meeting—as related to your event. Links to related assets are displayed in the *Event Details* window.

Saving and Drafting Changes and Updating Permissions

At the very bottom of the Edit form is a set of buttons that let you publish the changes, save the changes as a draft, and configure the event's permissions.

Giving a user permission to add, delete, or update discussion allows them to make, edit and remove comments on the event. The *Permissions* permission allows a Role to update permissions for the event.

Repeat Daily ✕

Repeat:

Repeat Every:

Stop Repeating: Never
 After
occurrences
 On

Summary: Daily

Figure 359.5: The *Repeat* box allows you to specify whether an events repeats daily, weekly, monthly, or yearly, how often it repeats, and when (or if) it ends.

359.4 Customizing Email Notifications

To customize email notifications for event invitations and reminders,

1. From the Calendar widget, click on the arrow next to a calendar and select *Calendar Settings*.
2. Click on the *Notification Templates* tab. Then select either the *Invite Email* or the *Reminder Email* sub-tab.
3. Edit the email as desired. At the bottom of the screen is a glossary that specifies variables for terms that were set when you created the event. Use these variables to refer to event-specific information, such as the event's name, date or location. It's a good idea to include a link to the event (use the variable [`$EVENT_URL$`]) as users must navigate to the calendar widget in order to respond.

Click *Save*. Now your notifications contain the proper text.

The next article covers setting up calendar resources and porting data from one installation to another.

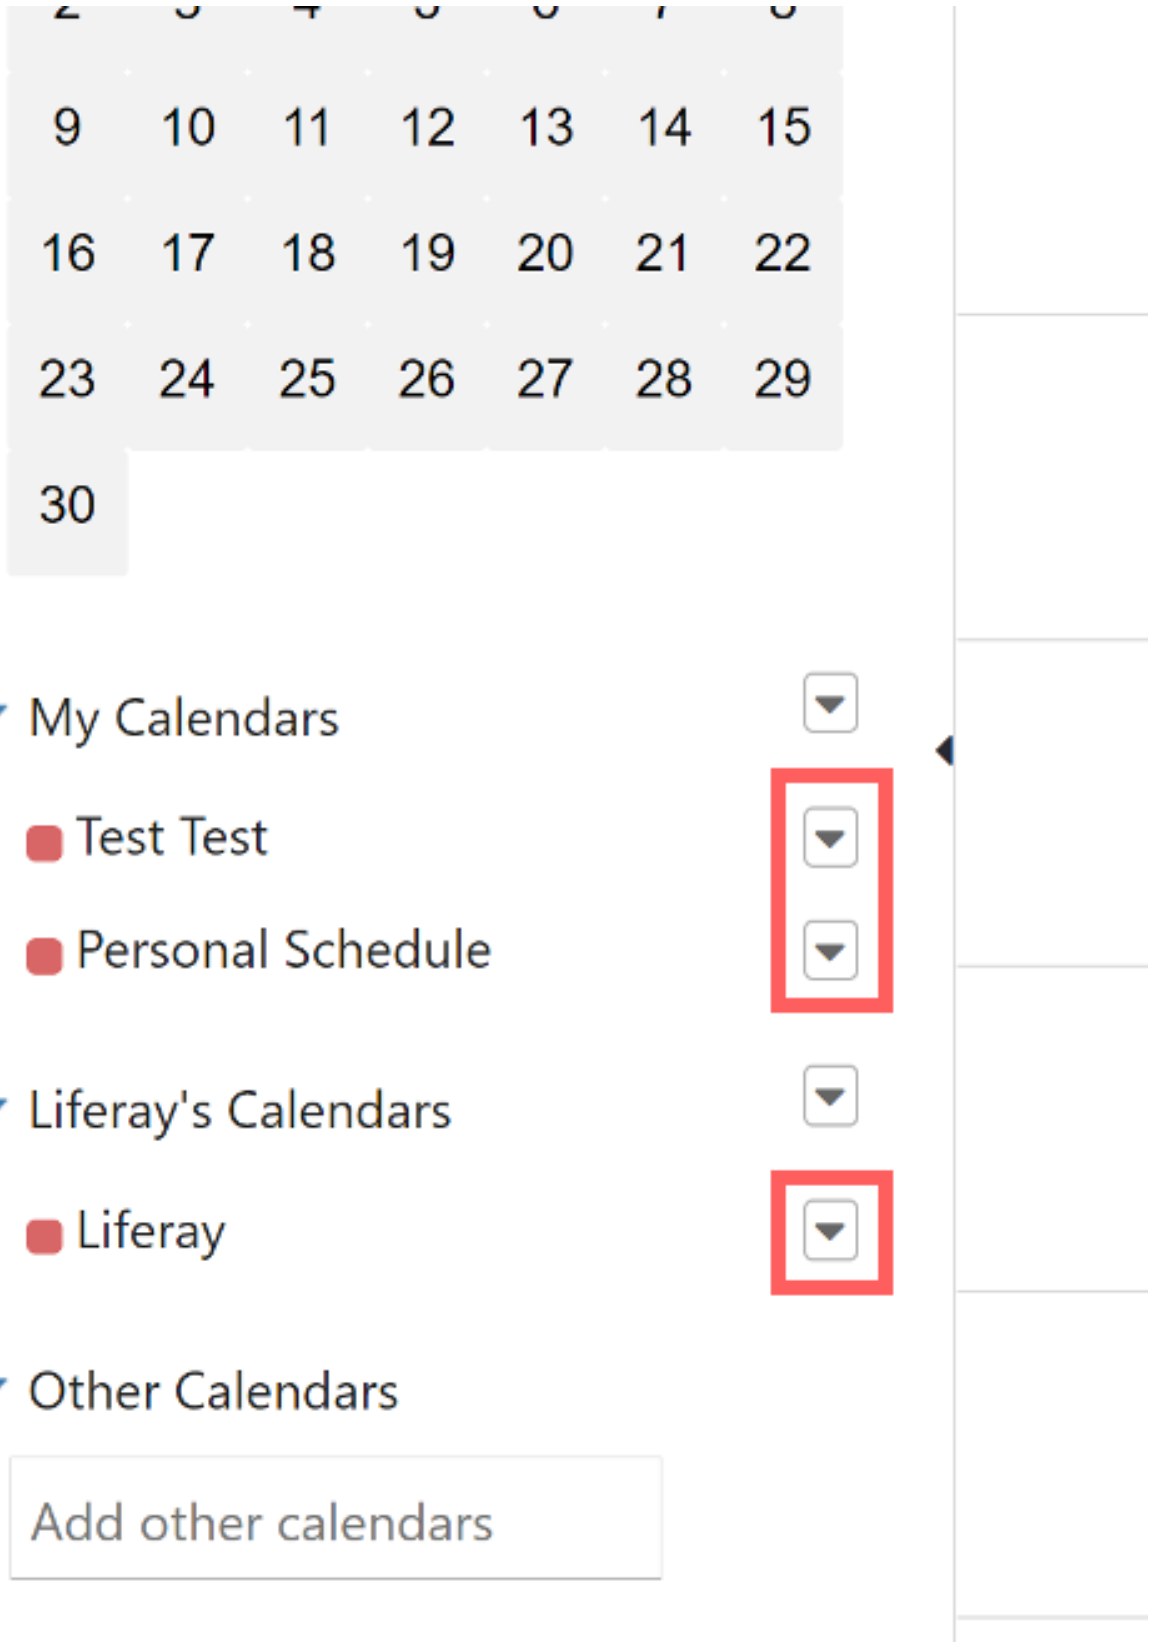


Figure 359.6: Email templates apply to a single calendar and all its events.


CALENDAR RESOURCES AND PORTING

With calendar resources, you can invite entities other than people to your events. This is beneficial for finding the availability of important resources your event requires, like a conference room, laptop, or, at The Lunar Resort, the Sasquatch Space Suit used to scare guests out on Lunar hikes.

Another important topic is porting your calendar's data from one installation of Liferay DXP to another.

360.1 Calendar Resources

Follow these steps to add a new calendar resource:

1. Click on the *Resources* tab and click the  button to add a new resource.

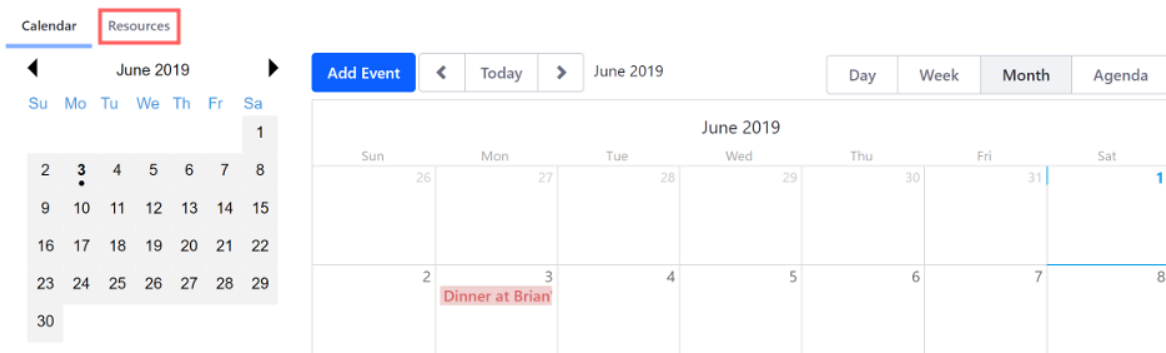



Figure 360.1: Resources are accessed from the tab menu at the top of the widget.

2. Fill in the *New Resource* form. Enter a name, give it a description, and choose whether to set it as active. You can also tag it, assign it to categories, and configure its permissions. Click *Save*.


The resource has its own calendar that was generated automatically (this is how users can check its availability when creating events). Just as with Users, however, resources can have more than one calendar. Follow these steps to assign a new calendar to the resource:

1. Go to the widget's *Resources* tab, click the  button next to the resource, and select *View Calendars*.
2. Click *Add Calendar* and continue just as if you were creating a calendar for a user or a Site.

Once a resource is created, invite it to your events just as you would an attendee.

360.2 Exporting and Importing Calendar Data

Like other Liferay Applications, the calendar allows data to be exported or imported as LAR files. As with all LAR files, data can only be ported between installations of the same version.

1. From the calendar widget, click the  button in the widget header and select *Export/Import*.
2. Enter a name for the LAR file (or use the default).

Under *Application*, choose whether to include the widget's configuration in the LAR.

Under *Content*, choose how much historical data to export and select the content types (calendars, resources, and events) to include. You can also choose whether to include comments and ratings.

Check the appropriate boxes to select whether to include deletions and permissions in the LAR.

3. Click *Export*. When a success message displays (this may take a few moments) you can click on the LAR's filename to download it.





| User | Status | Create Date | Completion Date | Download |
|---|------------|---------------|-----------------|---|
|  | Successful | 4 Seconds Ago | 3 Seconds Ago |  Calendar-201906010345.portlet.lar(8 KB)  |

Figure 360.2: This LAR is ready to be downloaded.

Follow these steps to import a LAR:

1. From the calendar widget, click the  button in the widget header and select *Export/Import*.
2. Click the *Import* tab.
3. Click *Choose File* or else drag-and-drop a LAR into the area surrounded by a dotted line. Click *Continue*.
4. Decide how much data you want to import:
Under *Application*, check the box to import the configuration stored in the LAR or leave in unchecked to keep your current configuration.
Under *Content*, decide which content types (calendars, resources, and events) to import, and whether to include comments and ratings.
Choose whether to import permissions and deletions, and decide whether to delete your widget's existing data before the import.
5. In the collapsible *Update Data* section, choose how data will be updated.
Mirror: The data will be imported along with a reference to its source. This allows data to be updated rather than duplicated if the same LAR is imported more than once.
Copy as New: All data is imported as new entries. Repeat imports will produce duplicates.
6. In the *Authorship of the Content* section, choose whether to keep the original author of the imported content (where available) or to list the current user as the author.
7. Click *Import*.

Your calendar is set up and ready to go! Better check it to see what's next on the agenda.

Part II

Deploying Liferay DXP

DEPLOYING LIFERAY DXP

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay DXP is one of the most flexible applications on the market today with respect to database and application server support. It supports a wide variety of databases and application servers, freeing you to use the ones you want. Liferay DXP also scales very well. You can install it on a shared hosting account, on a multi-node cluster running a commercial application server, or on anything in between. In fact, Liferay DXP is used successfully in all of these scenarios every day.

You'll find that because of Liferay DXP's flexibility in its deployment options, it is also easy to install. If you already have an application server, you can use your application server's deployment tools to install Liferay DXP. If you don't already have an application server, Liferay provides several application server bundles from which to choose. These are pre-configured application servers with Liferay DXP already installed on them. With a small amount of configuration, these can be made into production-ready systems.

There are some preparations to make before installing. You must create a database and install a supported Java Development Kit (JDK). It can also be worthwhile to pre-configure or gather information for configuring a data source, mail session, and more. You'll get guidance for these preparations.

Lastly, you'll install and deploy Liferay DXP for the first time and then set up Marketplace. You can continue configuring and tuning as you desire.

Read on to obtain the Liferay DXP installer that's right for you.

OBTAINING LIFERAY DXP

This document has been updated and ported to Liferay Learn and is no longer maintained here. Before you begin, you should answer a few questions.

- Which version of Liferay DXP will you install?
- Is your installation for development purposes or are you preparing to run in production?
- Are you installing Liferay DXP in a clustered environment?
- Which application server do you want Liferay DXP to run on?
- Are you installing on an existing application server?

Here you'll determine the installation that's best for you and download it.

Anyone can download Liferay DXP from liferay.com. Clicking *Resources* ▾ *Community Downloads* takes you to the Community Downloads page, where you can access Liferay Portal CE or a trial of the enterprise supported Liferay DXP. The installers are available in several different formats. The formats include a Liferay Tomcat bundle (Liferay DXP bundled with Tomcat) as well as a Liferay DXP .war (WAR) file for installing @product@ on an existing application server of choice.

Liferay enterprise subscribers can download Liferay DXP from the Help Center.

Liferay DXP runs on a wide variety of application servers (Check the Support page for a complete listing). Here are the ways to install Liferay DXP:

- Install a Liferay Tomcat bundle (Tomcat application server with Liferay DXP pre-installed).
- Install the Liferay WAR (and supporting libraries) onto an existing application server.

Since the Liferay Tomcat bundle is the easiest way, it's described first.

362.1 Liferay Tomcat Bundle

The Liferay Tomcat bundle includes the Tomcat application server with Liferay DXP pre-installed. If you prefer using another application server with Liferay DXP, you must install it manually. If you

don't currently have an application server preference, consider starting with the Tomcat bundle. Tomcat is one of the most lightweight and straightforward bundles to configure.

Note: Application server bundles for proprietary application servers such as WebLogic or WebSphere aren't available because the licenses for these servers don't allow for redistribution. Liferay DXP's commercial offering, however, runs just as well on these application servers as it does on the others.

Bundles are released as 7-Zip (.7z) and gzipped (.tar.gz) compressed file archive formats. Installing Liferay DXP demonstrates installing Liferay DXP from a bundle. Follow its instructions once you've prepared for installing Liferay DXP (see the next article).

Liferay DXP bundles might not be appropriate for you. Here are some reasons for installing the Liferay DXP WAR manually instead of using a bundle.

- There is no Liferay DXP bundle for the application server you want to use.
- You're installing Liferay DXP in a clustered environment.
- You're installing to an existing application server.

Manual installation is described next.

362.2 Downloading the Liferay WAR and Dependency JARs

Manual installation requires installing the Liferay DXP WAR and dependency JARs onto the application server. These files are available to download for DXP or Portal CE:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

After preparing for install (next), follow the Liferay DXP installation steps for your application server. There are specific configuration and script changes required on each application. The installation article titles follow this format, with *[Application Server]* replaced by the application server name.

Installing Liferay DXP on [Application Server]

PREPARING FOR INSTALL

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay DXP doesn't require much to deploy. You need a Java Development Kit (JDK) and a database. Several configuration topics (e.g., search engine integration, document repository configuration, security management, clustering, and more) can be addressed *after* deploying Liferay DXP.

Note: If you are installing Liferay DXP to multiple machines (e.g., in a cluster) or prefer centralizing configuration in a file, using portal properties in a [LIFERAY_HOME]/portal-ext.properties file is the recommended way to configure. The install preparation topics here and the configuration topics throughout this guide demonstrate using applicable portal properties.

Note: LIFERAY_HOME is the location from which Liferay DXP launches applications, applies configurations, loads JAR files, and generates logs. Liferay Home is customizable and can differ between application servers. The Liferay Home reference describes its folder structure.

Start preparing for Liferay DXP install by installing a supported Java Development Kit.

363.1 JDK Requirements

Liferay DXP deployment requires a JDK. The Support page lists the supported JDKs from various vendors. You must use one of these JDK versions:

- JDK 8
- JDK 11

JDK 11 is backwards compatible with JDK 8 applications. Applications and customizations developed on JDK 8 run on JDK 8 or JDK 11 runtimes. This makes JDK 8 best for developing on 7.0.

363.2 JVM Requirements

Liferay DXP requires that the application server JVM use the GMT time zone and UTF-8 file encoding. Include these JVM arguments to set the required values.

```
-Dfile.encoding=UTF8 -Duser.timezone=GMT
```

On JDK 11, it's recommended to add this JVM argument to display four-digit years.

```
-Djava.locale.providers=JRE,COMPAT,CLDR
```

Note: Since JDK 9, the Unicode Common Locale Data Repository (CLDR) is the default locales provider. CLDR, however, is not providing years in a four-digit format (see LPS-87191). The setting `java.locale.providers=JRE,COMPAT,CLDR` works around this issue by using JDK 8's default locales provider.

The recommended maximum heap size is 2GB. Setting the minimum heap size to the maximum heap size value minimizes garbage collections.

```
-Xms2560m -Xmx2560m
```

If you're using JDK 11, you may see *Illegal Access* warnings like these:

```
WARNING: An illegal reflective access operation has occurred
WARNING: Illegal reflective access by com.liferay.petra.reflect.ReflectionUtil (file:/Users/malei/dev/project/bundles/master-bundles/tomcat-9.0.10/lib/ext/com.liferay.petra.reflect.jar) to field java.lang.reflect.Field.modifiers
WARNING: Please consider reporting this to the maintainers of com.liferay.petra.reflect.ReflectionUtil
WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective access operations
WARNING: All illegal access operations will be denied in a future release
```

This is a known issue: LPS-87421. As a workaround, you can eliminate these warnings by adding these properties after your application server JVM options:

```
--add-opens=java.base/java.io=ALL-UNNAMED \
--add-opens=java.base/java.lang.reflect=ALL-UNNAMED \
--add-opens=java.base/java.lang=ALL-UNNAMED \
--add-opens=java.base/java.net=ALL-UNNAMED \
--add-opens=java.base/java.nio=ALL-UNNAMED \
--add-opens=java.base/java.text=ALL-UNNAMED \
--add-opens=java.base/java.util=ALL-UNNAMED \
--add-opens=java.base/sun.nio.ch=ALL-UNNAMED \
--add-opens=java.desktop/java.awt.font=ALL-UNNAMED \
--add-opens=java.rmi/sun.rmi.transport=ALL-UNNAMED \
--add-opens=java.xml/com.sun.org.apache.xerces.internal.parsers=ALL-UNNAMED
```

If you're using JDK 11 on Linux or UNIX and are activating Liferay DXP using an LCS 5.0.0 client, you may see an error like this:

```
ERROR [LCS Worker 2][BaseScheduledTask:92] java.lang.reflect.InaccessibleObjectException: Unable to make public long com.sun.management.internal.OperatingSystemImpl.getOpenFileDescriptorCount() opens com.sun.management.internal" to unnamed module @1a3325e5
java.lang.reflect.InaccessibleObjectException: Unable to make public long com.sun.management.internal.OperatingSystemImpl.getOpenFileDescriptorCount()
at java.base/java.lang.reflect.AccessibleObject.checkCanSetAccessible(AccessibleObject.java:
at java.base/java.lang.reflect.AccessibleObject.checkCanSetAccessible(AccessibleObject.java:
at java.base/java.lang.reflect.Method.checkCanSetAccessible(Method.java:198)
at java.base/java.lang.reflect.Method.setAccessible(Method.java:192)
```


To workaroud this issue, add this property after your application server JVM options:

```
--add-opens=jdk.management/com.sun.management.internal=ALL-UNNAMED
```

It's time to prepare your database.

363.3 Preparing a Database

The recommended way to set up your Liferay DXP database is also the simplest. Liferay DXP takes care of just about everything. Here are the steps:

1. Create a blank database encoded with the character set UTF-8. Liferay DXP is a multilingual application and needs UTF-8 encoding to display all of its supported character sets.

****Note:**** If you plan to migrate from one database vendor to another, [configure the database to use the default query result order you expect for entities Liferay DXP lists](/docs/7-2/deploy/-/knowledge_base/d/sort-order-changed-with-a-different-database).

2. Create a database user for accessing this database. Grant this database user all rights, including the rights to create and drop tables, to the blank Liferay DXP database.

Liferay DXP uses this database user's credentials to connect to the database either directly or through its application server. After you've configured the database connection, Liferay DXP creates its tables in the database automatically, complete with indexes.

This is the recommended way to set up Liferay DXP. It enables @product@ to maintain its database automatically during upgrades or when various Liferay DXP plugins that create database tables of their own are installed. This method is by far the best way to set up your database.

Warning: If you're using an Oracle database, use the ojdbc8.jar driver library with at least Oracle 12.2.0.1.0 JDBC 4.2 versioning because data truncation issues have been detected reading data from CLOB columns.

You can connect Liferay DXP with your database using @product@'s built-in data source (recommended) or using a data source you create on your app server.

363.4 Using the Built-in Data Source

You can configure the built-in data source from the Basic Configuration page (available when Liferay DXP starts up the first time) or by specifying it using portal properties.

Here's how set it using portal properties:

1. Create a `portal-ext.properties` file if you haven't created one already.
2. Copy a set of `jdbc.*` properties from one of the JDBC templates into your `portal-ext.properties` file.
3. Modify the `jdbc.*` property values to specify your database and database user credentials.
4. Put the `portal-ext.properties` file into your `LIFERAY_HOME` once you've established it based on your installation.

Liferay DXP connects to the data source on startup.

As an alternative to the built-in data source, you can use your application server's data source.

363.5 Using a Data Source on Your Application Server

Here's how to use your application server's data source:

1. Create your data source based on the instructions in the *Installing Liferay DXP on [Application Server]* article (for your application server) and your application server's documentation.
2. Create a `portal-ext.properties` file, if you haven't created one already.
3. Add the `jdbc.default.jndi.name` property set to the data source's JNDI name. Here's an example:

```
jdbc.default.jndi.name=jdbc/LiferayPool
```

4. Put the `portal-ext.properties` file into your `LIFERAY_HOME`, once you've established your `LIFERAY_HOME` based on your installation.

Liferay DXP connects to your data source on startup.

Allowing the database user you're using to initialize the Liferay DXP database to continue with all database rights is recommended. If you're fine with that user having the recommended permissions, skip the next section on limiting database access.

363.6 Limiting Database Access

Warning: The instructions below are not ideal for Liferay DXP installations. The following procedure is documented so that enterprises with more restrictive standards can install Liferay DXP with stricter (but sub-optimal) database settings. If it's at all possible, allow the database user that initializes the database to continue using the database with the same recommended permissions. The start of this section (Database Preparation) describes the recommended procedure for

initializing the Liferay DXP database and preserving that user's permissions for maintaining the Liferay DXP database and updating the database as plugin installations and plugin updates require.

Even though it's recommended for Liferay DXP to use the same database user to create and maintain its database automatically, your organizations might insist on revoking database initialization and maintenance permissions from that user once the database is initialized. If permissions for Select, Insert, Update and Delete operations are the only ones you allow for that user, you must initialize and maintain the database manually (even though it's not recommended). Here is the manual procedure:

1. Create a new, blank, database for Liferay DXP.
2. Grant full rights for the Liferay DXP database user to do anything to the database.
3. Install Liferay DXP and start it so that it automatically populates the database.
4. Once the database has been populated with the Liferay DXP tables, remove all permissions from that user except permissions to perform Select, Insert, Update and Delete operations.

There are some caveats to running Liferay DXP like this. Many plugins create new tables when they're deployed. Additionally, you must run the database upgrade function to upgrade Liferay DXP. If the @product@ database user doesn't have adequate rights to create/modify/drop tables in the database, you must grant those rights to that user before you deploy one of these plugins or start upgrading Liferay DXP. Once the tables are created or the upgrade completes, you can remove those rights until the next deploy or upgrade. Additionally, your own developers might create plugins that must create their own tables. These are just like Liferay DXP plugins that do the same thing, and they can only be installed if Liferay DXP can create database tables. Installing these plugins requires granting the Liferay DXP database user rights to create tables in the database before you attempt to install the plugins.

Liferay DXP has many more configurable features; but they can wait until *after* deployment. The Configuring Liferay DXP section explains them.

Now it's time to install Liferay DXP.

INSTALLING LIFERAY DXP

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Now that you've prepared for installing Liferay DXP, you can install and deploy it! This involves uncompressing the archive (the 7-Zip or gzip bundle file), possibly copying a JDBC driver to your application server, and starting the application server. Lastly, Liferay DXP provides a setup wizard to configure portal essentials.

Note: Since bundles are the easiest way to complete an installation, the installation steps here assume you're installing a Liferay DXP bundle. If you plan to install Liferay DXP manually, please refer to the *Installing @product@ on [Application Server]* article for the application server you're installing on.

364.1 Extracting a Liferay DXP Bundle

Extract your Liferay DXP bundle to the appropriate location on your server. This folder is the *Liferay Home*.

364.2 Installing the JDBC Driver

If you're using a supported open source database or if you're setting up Liferay DXP to use the embedded HSQL database for demo purposes, you can skip this step. Otherwise, copy your database's JDBC driver .jar file to the folder your application server documentation specifies. On Tomcat, for example, the driver belongs in the [Tomcat]/lib/ext folder.

364.3 Running Liferay DXP for the First Time

Start your application server using the start script bundled with your application server. For example, the Tomcat bundle provides the `startup.sh` script in `$CATALINA_HOME/bin`.

Note: Liferay DXP writes log files to folder `[Liferay Home]/logs`.

Important: Liferay DXP requires that the application server JVM use the GMT time zone and UTF-8 file encoding. They're preset in your Liferay DXP bundle.

The first time Liferay DXP starts, it creates all of its database tables. On completing startup, it launches a web browser that displays the Basic Configuration page (the setup wizard). If for some reason your browser doesn't load the Basic Configuration page automatically, open your browser and navigate to your app server's address and port (for example, `http://localhost:8080`).

364.4 Using the Setup Wizard

The Basic Configuration page provides a convenient way to configure these things:

- Portal name and default locale
- Administrator user
- Database

364.5 Portal

Supply this basic portal information:

Portal Name: name the installation you're powering with Liferay DXP.

Default Language: choose your portal's default locale and click the *Change* button. This immediately localizes your portal content, including the Basic Configuration page.

Time Zone: select your Liferay DXP instance's default time zone.

364.6 Administrator User

For the administrator, supply the following information:

Basic Configuration

| | |
|---|-----------------------------------|
| PORTAL | ADMINISTRATOR USER |
| Portal Name * | First Name * |
| <input type="text" value="Liferay"/> | <input type="text" value="Test"/> |
| Default Language | Last Name * |
| <input type="text" value="English (United States)"/> Change | <input type="text" value="Test"/> |
| Time Zone | Email * |
| <input type="text" value="(UTC) Coordinated Universal Time"/> | <input type="text"/> |

DATABASE

Default Database (Hypersonic)

This database is useful for development and demo'ing purposes, but it is not recommended for production use. ([Change](#))

[Finish Configuration](#)

Figure 364.1: Supply the information for your portal and your portal's default administrator user on the Basic Configuration page.

First Name: the default administrator user's first name

Last Name: the default administrator user's last name

Email: the default administrator user's email address

Note: the administrator user's email domain is used as the Liferay DXP instance's default domain (i.e., the `company.default.web.id` portal property).

364.7 Database

This section lets you connect to Liferay DXP's built-in data source.

Important: If you haven't created a database for Liferay DXP, create one now following database preparation instructions in the preceding article.

HSQL is selected as the default database, but it's primarily for demonstration or trial purposes. Click the *Change* link if you want to use Liferay DXP's built-in data source and configure it to use the database you created earlier.

The database configuration section also has an *Add Sample Data* checkbox for adding sample data to your database. This data includes Users, Sites, and Organizations for demo purposes. If you're installing Liferay DXP on your own machine to explore features, the sample data may be useful. If, however, you're installing Liferay DXP on a real server, start with a clean system by leaving this checkbox unselected.

Warning: HSQL should not be used in production Liferay DXP instances. Configure Liferay DXP to use a different database; specify that database via the Basic Configuration page here or using portal properties. See Database Preparation for details.

Once you've filled out the Basic Configuration form, click *Finish Configuration*. The setup wizard creates a `[LIFERAY_HOME]/portal-setup-wizard.properties` file which stores the settings that you entered. When you begin customizing your portal's configuration, however, use a `portal-ext.properties` file. The Portal properties reference documentation lists the default properties and describes all the properties you can set for Liferay DXP.

Tip: The wizard is a helpful tool, especially if you're setting up Liferay DXP for the first time. If you're a veteran and you already have your various properties set up, you can disable the setup wizard. If you disable the setup wizard, you must configure everything manually from a portal properties file (e.g., `[LIFERAY_HOME]/portal-ext.properties`). To disable the setup wizard, set `setup.wizard.enabled=false` in your portal properties file. Note that property values in `portal-setup-wizard.properties` (the file the setup wizards creates in Liferay Home) override property values in `portal-ext.properties`.

On finishing basic configuration, Liferay DXP prompts you to restart your server. When you restart your application server, Liferay DXP initializes the database you specified.

Now that Liferay DXP is up and running, you can continue configuring it as desired. Here are some suggestions:

1. Configure your mail session, if you haven't already configured it.
2. Install the Marketplace plugin, if it isn't already installed. If your machine has restricted access to the public network or if you restricted the Liferay DXP database user's permissions after initializing the database (not recommended), you can still set up Marketplace by following the Marketplace setup instructions.
3. Read the Configuring Liferay DXP articles for guidance in configuring Liferay DXP's default time zone, locales, logging, search engine, document repository, and more.

You're on your way to providing your organization with terrific experiences on Liferay DXP.

INSTALLING LIFERAY DXP ON TOMCAT

This document has been updated and ported to Liferay Learn and is no longer maintained here.

7.0 bundled with Tomcat 9 is available on the Help Center (DXP) or the Community Downloads page (Portal CE). The Tomcat bundle contains JARs, scripts, and configuration files required for deploying Liferay DXP. Copying these files from the @product@ Tomcat bundle facilitates installing Liferay DXP on an existing Tomcat application server.

Whether you copy bundle files (recommended) or download and create the files, you must download these files for DXP or Portal CE:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

Important: Prepare for the install before continuing.

Here are the basic steps for installing Liferay DXP on Tomcat:

- Installing dependencies to your application server
- Configuring your application server for Liferay DXP
- Deploying the Liferay DXP WAR file to your application server

Liferay Home is Tomcat server's parent folder. `$TOMCAT_HOME` refers to your Tomcat server folder. It is usually named `tomcat-[version]` or `apache-tomcat-[version]`.

365.1 Installing Dependencies

Liferay DXP depends on many JARs included by @product@ Tomcat bundle. Some of the bundle's JARs are not strictly required but can still be useful. If you don't have a bundle, you can download the Liferay JARs by downloading the *Dependencies* archive and the *OSGi Dependencies* archive, and you can download the third-party JARs as described below.

1. Create the folder `$TOMCAT_HOME/lib/ext` if it doesn't exist and extract the JARs from the dependencies ZIP to it.
2. Copy the following JARs from a Liferay DXP Tomcat bundle (or download them) to the `$TOMCAT_HOME/lib/ext` folder:
 - `activation.jar`
 - `ccpp.jar`
 - `jms.jar`
 - `jta.jar`
 - `jutf7.jar`
 - `mail.jar`
 - `persistence.jar`
 - `support-tomcat.jar`
3. Copy the JDBC driver for your database to the `$CATALINA_BASE/lib/ext` folder.

****Note:**** The [Liferay DXP Compatibility Matrix](<https://web.liferay.com/documents/14/21598941/Liferay+DXP+7.2+Compatibility+Matrix/b6e0f064-db31-49b4-8317-a29d1d76abf7?>) specifies supported databases and environments.

4. Create an `osgi` folder in your Liferay Home. Extract the folders (i.e., `configs`, `core`, and more) from OSGi ZIP file to the `osgi` folder. The `osgi` folder provides the necessary modules for Liferay DXP's OSGi runtime.

365.2 Configuring Tomcat

Configuring Tomcat to run Liferay DXP includes

- Setting environment variables
- Specifying a web application context for Liferay DXP
- Setting properties and descriptors

Optionally, if you're not using Liferay DXP's built-in data source or mail session, you can configure Tomcat to manage them:

- Data source
- Mail session

Start with configuring Tomcat to run Liferay DXP.

1. If you have a Liferay DXP Tomcat bundle, copy the `setenv.bat`, `setenv.sh`, `startup.bat`, `startup.sh`, `shutdown.bat`, and `shutdown.sh` files from it to your `$CATALINA_BASE/bin` folder. If not, create the `setenv.bat` and `setenv.sh` scripts.

The scripts set JVM options for Catalina, which is Tomcat's servlet container. Among these options is the location of the Java runtime environment. If this environment is not available on your server globally, you must set its location in these files so Tomcat can run. Do this by pointing the `JAVA_HOME` environment variable to a Liferay DXP-supported JRE:

```
export JAVA_HOME=/usr/lib/jvm/java-8-jdk
export PATH=$JAVA_HOME/bin:$PATH
```

Then configure Catalina's JVM options to support Liferay DXP.

Unix:

```
CATALINA_OPTS="$CATALINA_OPTS -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Dorg.apache.catalina.loader.WebappClassLoader.ENABLE_CLE
Duser.timezone=GMT -Xms2560m -Xmx2560m -XX:MaxMetaspaceSize=512m"
```

Windows:

```
set "CATALINA_OPTS=%CATALINA_OPTS% -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Dorg.apache.catalina.loader.WebappClassLoader.ENABLE
Duser.timezone=GMT -Xms2560m -Xmx2560m -XX:MaxMetaspaceSize=512m"
```

This sets the file encoding to UTF-8, prefers an IPv4 stack over IPv6, prevents Tomcat from working around garbage collection bugs relating to static or final fields (these bugs don't exist in Liferay DXP and working around them causes problems with the logging system), sets the time zone to GMT, gives the JVM 2GB of RAM, and limits Metaspace to 512MB.

****Important:**** Liferay DXP requires that the application server JVM use the GMT time zone and UTF-8 file encoding.

****Note:**** On JDK 11, it's recommended to add this JVM argument to display four-digit years.

```
``properties
-Djava.locale.providers=JRE,COMPAT,CLDR
``
```

After installation, tune your system (including these JVM options) for performance.

2. If you have a Liferay DXP Tomcat bundle, copy its `$CATALINA_BASE/conf/Catalina/localhost/ROOT.xml` file to the corresponding location in your application server. Create the file path if it doesn't exist. If you don't have a Liferay DXP Tomcat bundle, create a `ROOT.xml` file.

The `ROOT.xml` file specifies a web application context for Liferay DXP. `ROOT.xml` looks like this:

```

<Context crossContext="true" path="">

  <!-- JAAS -->

  <!--<Realm
    className="org.apache.catalina.realm.JAASRealm"
    appName="PortalRealm"
    userClassNames="com.liferay.portal.kernel.security.jaas.PortalPrincipal"
    roleClassNames="com.liferay.portal.kernel.security.jaas.PortalRole"
  />-->

  <!--
  Uncomment the following to disable persistent sessions across reboots.
  -->

  <!--<Manager pathname="" />-->

  <!--
  Uncomment the following to not use sessions. See the property
  "session.disabled" in portal.properties.
  -->

  <!--<Manager className="com.liferay.support.tomcat.session.SessionLessManagerBase" />-->

  <Resources>
    <PreResources
      base="{catalina.base}/lib/ext/portal"
      className="com.liferay.support.tomcat.webresources.ExtResourceSet"
      webAppMount="/WEB-INF/lib"
    />
  </Resources>
</Context>

```

Setting `crossContext="true"` lets multiple web applications use the same class loader. This configuration includes commented instructions and tags for configuring a JAAS realm, disabling persistent sessions, and disabling sessions entirely.

3. Provide Catalina access to the JARs in `$(CATALINA_BASE)/lib/ext` by opening your `$(CATALINA_BASE)/conf/catalina.p` file and appending this value to the `common.loader` property:

```

, "${catalina.home}/lib/ext/global", "${catalina.home}/lib/ext/global/*.jar", "${catalina.home}/lib/ext", "${catalina.home}/lib/ext/*.jar"

```

4. Make sure to use UTF-8 URI encoding consistently. If you have a Liferay DXP Tomcat bundle, copy the `$(CATALINA_BASE)/conf/server.xml` file to your server. If not, open your `$(CATALINA_BASE)/conf/server.xml` file and add the attribute `URIEncoding="UTF-8"` to HTTP and AJP connectors that use `redirectPort=8443`. Here are examples:

Old:

```

<Connector port="8080" protocol="HTTP/1.1" connectionTimeout="20000" redirectPort="8443" />

```

New:

```

<Connector port="8080" protocol="HTTP/1.1" connectionTimeout="20000" redirectPort="8443" URIEncoding="UTF-8" />

```

Old:

```

<Connector port="8009" protocol="AJP/1.3" redirectPort="8443" />

```

New:

```
<Connector port="8009" protocol="AJP/1.3" redirectPort="8443" URIEncoding="UTF-8" />
```

5. Refrain from writing access logs (optional) by commenting out the access log Valve element in `$CATALINA_BASE/conf/server.xml`. It's commented out here:

```
<!-- <Valve className="org.apache.catalina.valves.AccessLogValve"
      directory="logs"
      prefix="localhost_access_log" suffix=".txt"
      pattern="%h %l %u %t &quot;%r&quot; %s %b" /> -->
```

6. Optionally, set the following log levels in your `$CATALINA_HOME/conf/logging.properties` file:

```
org.apache.catalina.startup.Catalina.level=INFO
org.apache.catalina.startup.ClassLoaderFactory.level=SEVERE
org.apache.catalina.startup.VersionLoggerListener.level=WARNING
org.apache.level=WARNING
```

7. In `$CATALINA_HOME/conf/web.xml`, set the JSP compiler to Java 8 and set Liferay DXP's `TagHandlerPool` class to manage the JSP tag pool. Do this by adding the following elements above the `jsp` servlet element's `<load-on-startup>` element.

```
<init-param>
  <param-name>compilerSourceVM</param-name>
  <param-value>1.8</param-value>
</init-param>
<init-param>
  <param-name>compilerTargetVM</param-name>
  <param-value>1.8</param-value>
</init-param>
<init-param>
  <param-name>tagpoolClassName</param-name>
  <param-value>com.liferay.support.tomcat.jasper.runtime.TagHandlerPool</param-value>
</init-param>
```

8. In `$CATALINA_HOME/conf/web.xml`, specify whether the application server should look for extra metadata, such as annotations in the application's JARs and classes. Setting `web-app` element's attribute `metadata-complete="true"` tells the application server there's no extra metadata. The application server starts up faster this way. The default is to check for extra metadata.
9. If you're on Unix, Linux, or Mac OS, make the shell scripts in your `$CATALINA_HOME/bin` and `$CATALINA_BASE/bin` folders executable by running this command in each folder:

```
chmod a+x *.sh
```

Checkpoint:

Your application server is configured to run Liferay DXP.

1. The file encoding, user time-zone, and preferred protocol stack are set in your `setenv.sh`.
2. The default memory available and Metaspace limit are set.
3. `$CATALINA_BASE/conf/Catalina/localhost/ROOT.xml` declares the web application context.

4. The `common.loader` property in `/$CATALINA_BASE/conf/catalina.properties` grants Catalina access to the JARs in `/$CATALINA_BASE/lib/ext`.
5. `/$CATALINA_BASE/conf/server.xml` sets UTF-8 encoding.
6. `/$CATALINA_BASE/conf/server.xml` doesn't declare any valve for writing host access logs. (optional)
7. `/$CATALINA_HOME/conf/logging.properties` sets the desired log levels.
8. `/$CATALINA_HOME/conf/web.xml` sets the tag handler pool and sets Java 8 as the JSP compiler.
9. `/$CATALINA_HOME/conf/web.xml` specifies for the application server to refrain from looking for extra metadata. (optional)
10. The scripts in Tomcat's bin folders are executable.

365.3 Database Configuration

The easiest way to handle your database configuration is to let Liferay DXP manage your data source. If you want to use the built-in data source (recommended), skip this section.

If you want Tomcat to manage your data source, follow these steps:

1. Make sure your database server is installed and working. If it's installed on a different machine, make sure your Liferay DXP machine can access it.
2. Open `/$CATALINA_BASE/conf/Catalina/localhost/ROOT.xml` and add your data source as a Resource in your web application Context:

```
<Context...>
...
  <Resource
    name="jdbc/LiferayPool"
    auth="Container"
    type="javax.sql.DataSource"
    driverClassName="[place the database driver class here]"
    url="[place the URL to your database here]"
    username="[place your user name here]"
    password="[place your password here]"
    maxTotal="100"
    maxIdle="30"
    maxWaitMillis="10000"
  />
</Context>
```

Make sure to replace the database URL, user name, and password with the appropriate values. For example JDBC connection values, please see Database Templates

3. In a `portal-ext.properties` file in your Liferay Home, specify your data source:

```
jdbc.default.jndi.name=jdbc/LiferayPool
```

You created a data source for Tomcat to manage and configured Liferay DXP to use it. Mail session configuration is next.

365.4 Mail Configuration

As with database configuration, the easiest way to configure mail is to let Liferay DXP handle your mail session. If you want to use @product@’s built-in mail session, skip this section.

If you want to manage your mail session with Tomcat, follow these steps:

1. Open `$(CATALINA_BASE)/conf/Catalina/localhost/ROOT.xml` and add your mail session as a Resource in your web application Context. Make sure to replace the example mail session values with your own.

```
<Context...>
...
<Resource
  name="mail/MailSession"
  auth="Container"
  type="javax.mail.Session"
  mail.pop3.host="pop.gmail.com"
  mail.pop3.port="110"
  mail.smtp.host="smtp.gmail.com"
  mail.smtp.port="465"
  mail.smtp.user="user"
  mail.smtp.password="password"
  mail.smtp.auth="true"
  mail.smtp.starttls.enable="true"
  mail.smtp.socketFactory.class="javax.net.ssl.SSLSocketFactory"
  mail.imap.host="imap.gmail.com"
  mail.imap.port="993"
  mail.transport.protocol="smtp"
  mail.store.protocol="imap"
/>
</Context>
```

2. In your `portal-ext.properties` file in Liferay Home, reference your mail session:

```
mail.session.jndi.name=mail/MailSession
```

You’ve created a mail session for Tomcat to manage and configured Liferay DXP to use it.

365.5 Deploying Liferay DXP

Now you’re ready to deploy Liferay DXP using the @product@ WAR file.

1. If you are manually installing Liferay DXP on a clean Tomcat server, delete the contents of the `$(CATALINA_BASE)/webapps/ROOT` folder. This removes the default Tomcat home page.
2. Extract the Liferay DXP .war file contents to `$(CATALINA_BASE)/webapps/ROOT`.
It’s time to launch Liferay DXP on Tomcat!
3. Start Tomcat by navigating to `$(CATALINA_HOME)/bin` and executing `./startup.sh`. Alternatively, execute `./catalina.sh run` to tail Liferay DXP’s log file. The log audits startup activities and is useful for debugging deployment.

Congratulations on successfully installing and deploying Liferay DXP on Tomcat!

INSTALLING LIFERAY DXP ON WILDFLY

This document has been updated and ported to Liferay Learn and is no longer maintained here. Installing Liferay DXP on Wildfly 11 takes three steps:

- Installing dependencies to your application server
- Configuring your application server for Liferay DXP
- Deploying the Liferay DXP WAR file to your application server

Important: Before installing Liferay DXP, familiarize yourself with preparing for install.

Now, download the Liferay DXP WAR and Dependency JARs:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

Note that *Liferay Home* is the folder containing your Wildfly server folder. After installing and deploying Liferay DXP, the Liferay Home folder contains the Wildfly server folder as well as data, deploy, logs, and osgi folders. `$WILDFLY_HOME` refers to your Wildfly server folder. It is usually named `wildfly-[version]`.

366.1 Installing Dependencies

Liferay DXP depends on a driver for your database and the JARs in the Dependencies ZIP and OSGi Dependencies ZIP files you downloaded. Here's how to install them:

1. Create the folder `$WILDFLY_HOME/modules/com/liferay/portal/main` if it doesn't exist and extract the Dependencies ZIP JARs to it.
2. Download your database driver `.jar` file and copy it into the same folder.

****Note:**** The [Liferay DXP Compatibility Matrix](https://web.liferay.com/documents/14/21598941/Liferay+DXP+7.2+Compatibility+Matrix/b6e0f064-db31-49b4-8317-a29d1d76abf7?) specifies supported databases and environments.

3. Create the file `module.xml` in the `$WILDFLY_HOME/modules/com/liferay/portal/main` folder and insert this configuration:

```
<?xml version="1.0"?>

<module xmlns="urn:jboss:module:1.0" name="com.liferay.portal">
  <resources>
    <resource-root path="com.liferay.petra.concurrent.jar" />
    <resource-root path="com.liferay.petra.executor.jar" />
    <resource-root path="com.liferay.petra.function.jar" />
    <resource-root path="com.liferay.petra.io.jar" />
    <resource-root path="com.liferay.petra.lang.jar" />
    <resource-root path="com.liferay.petra.memory.jar" />
    <resource-root path="com.liferay.petra.nio.jar" />
    <resource-root path="com.liferay.petra.process.jar" />
    <resource-root path="com.liferay.petra.reflect.jar" />
    <resource-root path="com.liferay.petra.string.jar" />
    <resource-root path="com.liferay.registry.api.jar" />
    <resource-root path="hsql.jar" />
    <resource-root path="[place your database vendor's jar here]" />
    <resource-root path="portal-kernel.jar" />
    <resource-root path="portlet.jar" />
  </resources>
  <dependencies>
    <module name="javax.api" />
    <module name="javax.mail.api" />
    <module name="javax.servlet.api" />
    <module name="javax.servlet.jsp.api" />
    <module name="javax.transaction.api" />
  </dependencies>
</module>
```

Replace [place your database vendor's jar here] with the driver JAR for your database.

4. Create an `osgi` folder in your Liferay Home folder. Extract the OSGi Dependencies ZIP file that you downloaded into the [Liferay Home]/`osgi` folder.

The `osgi` folder provides the necessary modules for Liferay DXP's OSGi runtime.

366.2 Running Liferay DXP on Wildfly in Standalone Mode vs. Domain Mode

Wildfly can be launched in either *standalone* mode or *domain* mode. Domain mode allows multiple application server instances to be managed from a single control point. A collection of such application servers is known as a *domain*. For more information on standalone mode vs. domain mode, please refer to the section on this topic in the Wildfly Admin Guide. Liferay DXP fully supports Wildfly in standalone mode but not in domain mode.

You can run Liferay DXP on Wildfly in domain mode, but this method is not fully supported. In particular, Liferay DXP's hot-deploy does not work with a managed deployment, since Wildfly

manages the content of a managed deployment by copying files (exploded or non-exploded). This prevents JSP hooks and Ext plugins from working as intended. For example, JSP hooks don't work on Wildfly running in managed domain mode, since Liferay DXP's JSP override mechanism relies on the application server. Since JSP hooks and Ext plugins are deprecated, however, you may not be using them.

The command line interface is recommended for domain mode deployments.

Note: This does not prevent Liferay DXP from running in a clustered environment on multiple Wildfly servers. You can set up a cluster of Liferay DXP instances running on Wildfly servers running in standalone mode. Please refer to the Liferay DXP clustering articles for more information.

366.3 Configuring Wildfly

Configuring Wildfly to run Liferay DXP includes these things:

- Setting environment variables
- Setting properties and descriptors
- Removing unnecessary configurations

Optionally, you can configure Wildfly to manage Liferay DXP's data source and mail session. Start with configuring Wildfly to run Liferay DXP.

Make the following modifications to `WILDFLY_HOME/standalone/configuration/standalone.xml`:

1. Locate the closing `</extensions>` tag. Directly beneath that tag, insert the following system properties:

```
<system-properties>
  <property name="org.apache.catalina.connector.URI_ENCODING" value="UTF-8" />
  <property name="org.apache.catalina.connector.USE_BODY_ENCODING_FOR_QUERY_STRING" value="true" />
</system-properties>
```

2. Add the following `<filter-spec>` tag within the `<console-handler>` tag, directly below the `<level name="INFO"/>` tag:

```
<filter-spec value="not(any(match(&quot;WFLYSRV0059&quot;),match(&quot;WFLYEE0007&quot;)))" />
```

3. Add a timeout for the deployment scanner by setting `deployment-timeout="600"` as seen in the excerpt below.

```
<subsystem xmlns="urn:jboss:domain:deployment-scanner:2.0">
  <deployment-scanner deployment-timeout="600" path="deployments" relative-to="jboss.server.base.dir" scan-interval="5000" runtime-
failure-causes-rollback="{jboss.deployment.scanner.rollback.on.failure:false}"/>
</subsystem>
```

4. Add the following JAAS security domain to the security subsystem `<security-domains>` defined in element `<subsystem xmlns="urn:jboss:domain:security:2.0">`.

```
<security-domain name="PortalRealm">
  <authentication>
    <login-module code="com.liferay.portal.security.jaas.PortalLoginModule" flag="required" />
  </authentication>
</security-domain>
```

5. Remove the welcome content code snippets:

```
<location name="/" handler="welcome-content"/>
```

and

```
<handlers>
  <file name="welcome-content" path="${jboss.home.dir}/welcome-content"/>
</handlers>
```

6. Find the <jsp-config/> tag and set the development, source-vm, and target-vm attributes in the tag. Once finished, the tag should look like this:

```
<jsp-config development="true" source-vm="1.8" target-vm="1.8" />
```

Checkpoint:

Before continuing, verify the following properties have been set in the standalone.xml file:

1. The new <system-property> is added.
2. The new <filter-spec> is added.
3. The <deployment-timeout> is set to 600.
4. The new <security-domain> is created.
5. Welcome content is removed.
6. The <jsp-config> tag contains its new attributes.

Now you must configure your JVM and startup scripts.

In the \$WILDFLY_HOME/bin/ folder, modify your standalone domain's configuration script file standalone.conf (standalone.conf.bat on Windows):

- Set the file encoding to UTF-8
- Set the user time zone to GMT
- Set the preferred protocol stack
- Increase the default amount of memory available.

Important: For Liferay DXP to work properly, the application server JVM must use the GMT time zone and UTF-8 file encoding.

Make the following edits as applicable for your operating system:

Windows:

1. Comment out the initial JAVA_OPTS assignment like this:

```
rem set "JAVA_OPTS=-Xms64M -Xmx512M -XX:MetaspaceSize=96M -XX:MaxMetaspaceSize=2560m"
```

2. Add the following JAVA_OPTS assignment one line above the :JAVA_OPTS_SET line found at end of the file:

```
set "JAVA_OPTS=%JAVA_OPTS% -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Djboss.as.management.blocking.timeout=480 -
Duser.timezone=GMT -Xms2560m -Xmx2560m -XX:MaxMetaspaceSize=512m -XX:MetaspaceSize=200m"
```

Unix:

1. Below the if ["x\$JAVA_OPTS" = "x"]; statement, replace this JAVA_OPTS statement:

```
JAVA_OPTS="-Xms64m -Xmx512m -XX:MetaspaceSize=96M -XX:MaxMetaspaceSize=256m -Djava.net.preferIPv4Stack=true"
```

with this:

```
JAVA_OPTS="-Djava.net.preferIPv4Stack=true"
```

2. Add the following statement to the bottom of the file:

```
JAVA_OPTS="$JAVA_OPTS -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Djboss.as.management.blocking.timeout=480 -
Duser.timezone=GMT -Xms2560m -Xmx2560m -XX:MaxMetaspaceSize=512m -XX:MetaspaceSize=200m"
```

This sets the file encoding to UTF-8, prefers an IPv4 stack over IPv6, sets the time zone to GMT, gives the JVM 2GB of RAM, and limits Metaspace to 512MB.

On JDK 11, it's recommended to add this JVM argument to display four-digit years.

```
-Djava.locale.providers=JRE,COMPAT,CLDR
```

After installation, tune your system (including these JVM options) for performance.

Important: For Liferay DXP to work properly, the application server JVM must use the GMT time zone and UTF-8 file encoding.

Note: If you plan on using the IBM JDK with your Wildfly server, you must complete some additional steps. First, navigate to the `$WILDFLY_HOME/modules/com/liferay/portal/main/module.xml` file and insert the following dependency within the `<dependencies>` element:

```
<module name="ibm.jdk" />
```

Then navigate to the `$WILDFLY_HOME/modules/system/layers/base/sun/jdk/main/module.xml` file and insert the following path names inside the `<paths>...</paths>` element:

```
<path name="com/sun/crypto" />
<path name="com/sun/crypto/provider" />
<path name="com/sun/org/apache/xml/internal/resolver" />
<path name="com/sun/org/apache/xml/internal/resolver/tools" />
```

The added paths resolve issues with deployment exceptions and image uploading problems.

Checkpoint:

You've configured the application server's JVM settings.

1. The file encoding, user time-zone, preferred protocol stack have been set in the JAVA_OPTS in the standalone.conf.bat file.
2. The default amount of memory available has been increased.

The prescribed script modifications are now complete for your Liferay DXP installation on Wildfly. Next you'll configure your database.

366.4 Database Configuration

The easiest way to handle database configuration is to let Liferay DXP manage your data source. The Basic Configuration page lets you configure Liferay DXP's built-in data source. If you want to use the built-in data source, skip this section.

If using WildFly to manage the data source, follow these steps:

1. Add the data source inside the `$WILDFLY_HOME/standalone/configuration/standalone.xml` file's `<datasources>` element:

```
<datasource jndi-name="java:jboss/datasources/ExampleDS" pool-name="ExampleDS" enabled="true" jta="true" use-java-
context="true" use-ccm="true">
  <connection-url>[place the URL to your database here]</connection-url>
  <driver>[place your driver name here]</driver>
  <security>
    <user-name>[place your user name here]</user-name>
    <password>[place your password here]</password>
  </security>
</datasource>
```

Make sure to replace the database URL, user name, and password with the appropriate values. For example JDBC connection values, please see Database Templates

Note: If the data source `jndi-name` must be changed, edit the `datasource` element in the `<default-bindings>` tag.

1. Add the driver to the standalone.xml file's `<drivers>` element also found within the `<datasources>` element.

```
<drivers>
  <driver name="[name of database driver]" module="com.liferay.portal">
    <driver-class>[JDBC driver class]</driver-class>
  </driver>
</drivers>
```

A final data sources subsystem that uses MySQL should look like this:

```
<subsystem xmlns="urn:jboss:domain:datasources:1.0">
  <datasources>
    <datasource jndi-name="java:jboss/datasources/ExampleDS" pool-name="ExampleDS" enabled="true" jta="true" use-
java-context="true" use-ccm="true">
      <connection-url>jdbc:mysql://localhost/lportal</connection-url>
      <driver>mysql</driver>
      <security>
        <user-name>root</user-name>
        <password>root</password>
      </security>
    </datasource>
  </datasources>
  <drivers>
    <driver name="mysql" module="com.liferay.portal">
      <driver-class>com.mysql.cj.jdbc.Driver</driver-class>
    </driver>
  </drivers>
</datasources>
</subsystem>
```

2. In a `portal-ext.properties` file in your Liferay Home, specify your data source:

```
jdbc.default.jndi.name=java:jboss/datasources/ExampleDS
```

Now that you've configured your data source, the mail session is next.

366.5 Mail Configuration

As with database configuration, the easiest way to configure mail is to let Liferay DXP handle your mail session. If you want to use `@product@`'s built-in mail session, skip this section and configure the mail session in the Control Panel.

If you want to manage your mail session with Wildfly, follow these steps:

1. Specify your mail subsystem in the `$WILDFLY_HOME/standalone/configuration/standalone.xml` file like this:

```
<subsystem xmlns="urn:jboss:domain:mail:3.0">
  <mail-session jndi-name="java:jboss/mail/MailSession" name="mail-smtp">
    <smtp-server ssl="true" outbound-socket-binding-ref="mail-smtp" username="USERNAME" password="PASSWORD"/>
  </mail-session>
</subsystem>
...
<socket-binding-group name="standard-sockets" default-interface="public" port-offset="${jboss.socket.binding.port-
offset:0}">
  ...
  <outbound-socket-binding name="mail-smtp">
    <remote-destination host="[place SMTP host here]" port="[place SMTP port here]"/>
  </outbound-socket-binding>
</socket-binding-group>
```

2. In your `portal-ext.properties` file in Liferay Home, reference your mail session:

```
mail.session.jndi.name=java:jboss/mail/MailSession
```

Next, you'll deploy Liferay DXP to your Wildfly app server.

366.6 Deploying Liferay DXP

Now you're ready to deploy Liferay DXP using the @product@ WAR file.

1. If the folder `$WILDFLY_HOME/standalone/deployments/ROOT.war` already exists in your Wildfly installation, delete all of its subfolders and files. Otherwise, create a new folder called `$WILDFLY_HOME/standalone/deployments/ROOT.war`.
2. Unzip the Liferay DXP .war file into the `ROOT.war` folder.
3. To trigger deployment of `ROOT.war`, create an empty file named `ROOT.war.dodeploy` in your `$WILDFLY_HOME/standalone/deployments/` folder. On startup, Wildfly detects this file and deploys it as a web application.
4. Start the Wildfly application server by navigating to `$WILDFLY_HOME/bin` and running `standalone.bat` or `standalone.sh`.

Congratulations; you've deployed Liferay DXP on Wildfly!

Note: After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving PhaseOptimizer. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass gatherExternProperties
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass checkControlFlow
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array pattern,
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expression,
current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array pattern,
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expression]
```

INSTALLING LIFERAY DXP ON JBOSS EAP

This document has been updated and ported to Liferay Learn and is no longer maintained here. Installing Liferay DXP on JBoss EAP 7.1 takes three steps:

- Installing dependencies to your application server
- Configuring your application server for Liferay DXP
- Installing the Liferay DXP WAR file to your application server

Important: Before installing Liferay DXP, familiarize yourself with preparing for install.

Now, download the Liferay DXP WAR and Dependency JARs:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

Not that *Liferay Home* is the folder containing your JBoss server folder. After installing and deploying Liferay DXP, the Liferay Home folder contains the JBoss server folder as well as data, deploy, logs, and osgi folders. `$JBOSS_HOME` refers to your JBoss server folder. This folder is usually named `jboss-eap-[version]`.

367.1 Installing Dependencies

Liferay DXP depends on several Liferay-specific and third-party JARs. Download and install the required JARs as described below.

1. Create the folder `$JBOSS_HOME/modules/com/liferay/portal/main` if it doesn't exist and extract the JARs from the dependencies ZIP to it.
2. Download your database driver `.jar` file and copy it into the same folder.

Note: The [Liferay DXP Compatibility Matrix](<https://web.liferay.com/documents/14/21598941/Liferay+DXP+7.2+Compatibility+Matrix/b6e0f064-db31-49b4-8317-a29d1d76abf7?>) specifies supported databases and environments.

3. Create the file `module.xml` in the `$JBOSS_HOME/modules/com.liferay.portal/main` folder and insert this configuration:

```
<?xml version="1.0"?>

<module xmlns="urn:jboss:module:1.0" name="com.liferay.portal">
  <resources>
    <resource-root path="com.liferay.petra.concurrent.jar" />
    <resource-root path="com.liferay.petra.executor.jar" />
    <resource-root path="com.liferay.petra.function.jar" />
    <resource-root path="com.liferay.petra.io.jar" />
    <resource-root path="com.liferay.petra.lang.jar" />
    <resource-root path="com.liferay.petra.memory.jar" />
    <resource-root path="com.liferay.petra.nio.jar" />
    <resource-root path="com.liferay.petra.process.jar" />
    <resource-root path="com.liferay.petra.reflect.jar" />
    <resource-root path="com.liferay.petra.string.jar" />
    <resource-root path="com.liferay.registry.api.jar" />
    <resource-root path="hsqldb.jar" />
    <resource-root path="[place your database vendor's jar here]" />
    <resource-root path="portal-kernel.jar" />
    <resource-root path="portlet.jar" />
  </resources>
  <dependencies>
    <module name="javax.api" />
    <module name="javax.mail.api" />
    <module name="javax.servlet.api" />
    <module name="javax.servlet.jsp.api" />
    <module name="javax.transaction.api" />
  </dependencies>
</module>
```

Replace `[place your database vendor's jar here]` with the driver JAR for your database.

4. Create an `osgi` folder in your Liferay Home folder. Extract the OSGi Dependencies ZIP file that you downloaded into the `[Liferay Home]/osgi` folder.

The `osgi` folder provides the necessary modules for Liferay DXP's OSGi runtime.

Checkpoint:

1. The dependencies files have been unzipped into the `$JBOSS_HOME/modules/com.liferay.portal/main` folder and a database jar.
2. The `module.xml` contains all JARs in the `<resource-root-path>` elements.
3. The `osgi` dependencies have been unzipped into the `osgi` folder.

367.2 Running Liferay DXP on JBoss EAP in Standalone Mode vs. Domain Mode

JBoss EAP can be launched in either *standalone* mode or *domain* mode. Domain mode allows multiple application server instances to be managed from a single control point. A collection of such application servers is known as a *domain*. For more information on standalone mode vs. domain mode, please refer to the section on this topic in the JBoss EAP Product Documentation.

Liferay DXP supports JBoss EAP when it runs in standalone mode but not when it runs in domain mode. Liferay DXP's hot-deploy does not work with a managed deployment, since JBoss manages the content of a managed deployment by copying files (exploded or non-exploded). This prevents JSP hooks and Ext plugins from working as intended. For example, JSP hooks don't work on JBoss EAP running in managed domain mode, since Liferay DXP's JSP override mechanism relies on the application server. Since JSP hooks and Ext plugins are deprecated, however, you may not be using them.

The command line interface is recommended for domain mode deployments.

Note: This does not prevent Liferay DXP from running in a clustered environment on multiple JBoss servers. You can set up a cluster of Liferay DXP instances running on JBoss EAP servers running in standalone mode. Please refer to the Liferay DXP clustering articles for more information.

367.3 Configuring JBoss

Configuring JBoss to run Liferay DXP includes these things:

- Setting environment variables
- Setting properties and descriptors
- Removing unnecessary configurations

Optionally, you can configure JBoss to manage Liferay DXP's data source and mail session. Start with configuring JBoss to run Liferay DXP.

Make the following modifications to `$JBOSS_HOME/standalone/configuration/standalone.xml`:

1. Locate the closing `</extensions>` tag. Directly beneath that tag, insert these system properties:

```
<system-properties>
  <property name="org.apache.catalina.connector.URI_ENCODING" value="UTF-8" />
  <property name="org.apache.catalina.connector.USE_BODY_ENCODING_FOR_QUERY_STRING" value="true" />
</system-properties>
```

2. Add the following `<filter-spec>` tag within the `<console-handler>` tag, directly below the `<level name="INFO"/>` tag:

```
<filter-spec value="not(any(match(&quot;WFLYSRV0059&quot;),match(&quot;WFLYEE0007&quot;)))" />
```

3. Add a timeout for the deployment scanner by setting `deployment-timeout="600"` as seen in the excerpt below.

```
<subsystem xmlns="urn:jboss:domain:deployment-scanner:2.0">
  <deployment-scanner deployment-timeout="600" path="deployments" relative-to="jboss.server.base.dir" scan-interval="5000"/>
</subsystem>
```

4. Add the following JAAS security domain to the security subsystem <security-domains> defined in element <subsystem xmlns="urn:jboss:domain:security:2.0">.

```
<security-domain name="PortalRealm">
  <authentication>
    <login-module code="com.liferay.portal.security.jaas.PortalLoginModule" flag="required" />
  </authentication>
</security-domain>
```

5. Remove the welcome content code snippets:

```
<location name="/" handler="welcome-content"/>
```

and

```
<handlers>
  <file name="welcome-content" path="${jboss.home.dir}/welcome-content"/>
</handlers>
```

6. Find the <jsp-config/> tag and set the development, source-vm, and target-vm attributes in the tag. Once finished, the tag should look like this:

```
<jsp-config development="true" source-vm="1.8" target-vm="1.8"/>
```

Checkpoint:

Before continuing, verify the following properties have been set in the standalone.xml file:

1. The new <system-property> is added.
2. The new <filter-spec> is added.
3. The <deployment-timeout> is set to 360.
4. The new <security-domain> is created.
5. Welcome content is removed.
6. The <jsp-config> tag contains its new attributes.

Now you should configure your JVM and startup scripts.

In the \$WILDFLY_HOME/bin/ folder, modify your standalone domain's configuration script file standalone.conf (standalone.conf.bat on Windows):

- Set the file encoding to UTF-8
- Set the user time zone to GMT
- Set the preferred protocol stack
- Increase the default amount of memory available.

Important: For Liferay DXP to work properly, the application server JVM must use the GMT time zone and UTF-8 file encoding.

Make the following edits as applicable to your operating system:

Windows

1. Comment out the initial JAVA_OPTS assignment as demonstrated in the following line:

```
rem set "JAVA_OPTS=-Xms1G -Xmx1G -XX:MetaspaceSize=96M -XX:MaxMetaspaceSize=2560m"
```

2. Add the following JAVA_OPTS assignment one line above the :JAVA_OPTS_SET line found at end of the file:

```
set "JAVA_OPTS=%JAVA_OPTS% -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Djboss.as.management.blocking.timeout=480 -Duser.timezone=GMT -Xms2560m -Xmx2560m -XX:MaxMetaspaceSize=768m"
```

Unix

1. Below the if ["x\$JAVA_OPTS" = "x"]; statement, replace this JAVA_OPTS statement:

```
JAVA_OPTS="-Xms1303m -Xmx1303m -XX:MetaspaceSize=96M -XX:MaxMetaspaceSize=2560m -Djava.net.preferIPv4Stack=true"
```

with this:

```
JAVA_OPTS="-Djava.net.preferIPv4Stack=true"
```

2. Add the following statement to the bottom of the file:

```
JAVA_OPTS="$JAVA_OPTS -Dfile.encoding=UTF-8 -Djava.net.preferIPv4Stack=true -Djboss.as.management.blocking.timeout=480 -Duser.timezone=GMT -Xms2560m -Xmx2560m -XX:MaxMetaspaceSize=512m"
```

On JDK 11, it's recommended to add this JVM argument to display four-digit years.

```
-Djava.locale.providers=JRE,COMPAT,CLDR
```

Note: If you plan on using the IBM JDK with your JBoss server, you must complete some additional steps. First, navigate to the \$JBOSS_HOME/modules/com/liferay/portal/main/module.xml file and insert the following dependency within the <dependencies> element:

```
<module name="ibm.jdk" />
```

Then navigate to the \$JBOSS_HOME/modules/system/layers/base/sun/jdk/main/module.xml file and insert the following path names inside the <paths> ... </paths> element:

```
<path name="com/sun/crypto" />
<path name="com/sun/crypto/provider" />
<path name="com/sun/image/codec/jpeg" />
<path name="com/sun/org/apache/xml/internal/resolver" />
<path name="com/sun/org/apache/xml/internal/resolver/tools" />
```

The added paths resolve issues with portal deployment exceptions and image uploading problems.

Checkpoint:

At this point, you've finished configuring the application server's JVM settings.

1. The file encoding, user time-zone, preferred protocol stack have been set in the `JAVA_OPTS` in the `standalone.conf.bat` file.
2. The default amount of memory available has been increased.

The prescribed script modifications are now complete for your Liferay DXP installation on JBoss. Next you'll configure the database and mail.

367.4 Database Configuration

The easiest way to handle your database configuration is to let Liferay DXP manage your data source. The Basic Configuration page lets you configure Liferay DXP's built-in data source. If you want to use the built-in data source, skip this section.

This section demonstrates configuring a MySQL database. If you're using a different database, modify the data source and driver snippets as necessary.

If using JBoss to manage the data source, follow these steps:

1. Add the data source inside the `$JBOSS_HOME/standalone/configuration/standalone.xml` file's the `<datasources>` element.

```
<datasource jndi-name="java:jboss/datasources/ExampleDS" pool-name="ExampleDS" enabled="true" jta="true" use-java-
context="true" use-ccm="true">
  <connection-url>[place the URL to your database here]</connection-url>
  <driver>[place the driver name here]</driver>
  <security>
    <user-name>[place your user name here]</user-name>
    <password>[place your password here]</password>
  </security>
</datasource>
```

Make sure to replace the database URL, user name, and password with the appropriate values.

****Note:**** If the data source ``jndi-name`` must be changed, edit the ``datasource`` element in the ``<default-bindings>`` tag.

2. Add your driver to the `standalone.xml` file's `<drivers>` element also found within the `<datasources>` element.

```

<drivers>
  <driver name="[name of driver must match name above]" module="com.liferay.portal">
    <driver-class>[place your JDBC driver class here]</driver-class>
  </driver>
</drivers>

```

A final data sources subsystem that uses MySQL should look like this:

```

<subsystem xmlns="urn:jboss:domain:datasources:5.0">
  <datasources>
    <datasource jndi-name="java:jboss/datasources/ExampleDS" pool-name="ExampleDS" enabled="true" jta="true" use-
java-context="true" use-ccm="true">
      <connection-url>jdbc:mysql://localhost/lportal</connection-url>
      <driver>mysql</driver>
      <security>
        <user-name>root</user-name>
        <password>root</password>
      </security>
    </datasource>
  </datasources>
  <drivers>
    <driver name="mysql" module="com.liferay.portal"/>
  </drivers>
</subsystem>

```

3. In a portal-ext.properties file in your Liferay Home, specify your data source:

```
jdbc.default.jndi.name=java:jboss/datasources/ExampleDS
```

Now that you've configured your data source, the mail session is next.

367.5 Mail Configuration

As with database configuration, the easiest way to configure mail is to let Liferay DXP handle your mail session. If you want to use @product@'s built-in mail session, skip this section and configure the mail session in the Control Panel.

If you want to manage your mail session with JBoss, follow these steps:

1. Specify your mail subsystem in the \$JBOSS_HOME/standalone/configuration/standalone.xml file like this:

```

<subsystem xmlns="urn:jboss:domain:mail:3.0">
  <mail-session jndi-name="java:jboss/mail/MailSession" >
    <smtp-server ssl="true" outbound-socket-binding-ref="mail-smtp">
      <login username="[place user name here]" password="[place password here]"/>
    </smtp-server>
  </mail-session>
</subsystem>
...
<socket-binding-group name="standard-sockets" default-interface="public" port-offset="${jboss.socket.binding.port-
offset:0}">
  ...
  <outbound-socket-binding name="mail-smtp">
    <remote-destination host="[place SMTP mail host here]" port="[place mail port here]"/>
  </outbound-socket-binding>
</socket-binding-group>

```

2. In your `portal-ext.properties` file in Liferay Home, reference your mail session:

```
mail.session.jndi.name=java:jboss/mail/MailSession
```

You've got mail! Next, you'll deploy Liferay DXP to your JBoss app server.

367.6 Deploy Liferay

Now you're ready to deploy Liferay DXP using the `@product@` WAR file.

1. If the folder `$JBOSS_HOME/standalone/deployments/ROOT.war` already exists in your JBoss installation, delete all of its subfolders and files. Otherwise, create a new folder called `$JBOSS_HOME/standalone/deployments/ROOT.war`.
2. Unzip the Liferay DXP `.war` file into the `ROOT.war` folder.
3. To trigger deployment of `ROOT.war`, create an empty file named `ROOT.war.dodeploy` in your `$JBOSS_HOME/standalone/deployments/` folder. On startup, JBoss detects this file and deploys it as a web application.
4. Start the JBoss application server by navigating to `$JBOSS_HOME/bin` and running `standalone.bat` or `standalone.sh`.

Congratulations; you've now deployed Liferay DXP on JBoss!

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving `PhaseOptimizer`. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass gatherExternProperties
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass checkControlFlow
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array pattern,
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressions,
current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array pattern,
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressions]
```

INSTALLING LIFERAY DXP ON WEBLOGIC 12C R2

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Although you can install Liferay DXP in a WebLogic Admin Server, this isn't recommended. It's a best practice to install web apps, including Liferay DXP, in a WebLogic Managed server. Deploying to a Managed Server lets you start or shut down Liferay DXP more quickly and facilitates transitioning into a cluster configuration. This article therefore focuses on installing Liferay DXP in a Managed Server.

Before getting started, create your Admin and Managed Servers. See WebLogic's documentation for instructions on setting up and configuring Admin and Managed Servers.

Also familiarize yourself with preparing for install.

Now, download the Liferay DXP WAR and Dependency JARs:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

368.1 Configuring WebLogic's Node Manager

WebLogic's Node Manager starts and stops managed servers.

If you're running WebLogic on a UNIX system other than Solaris or Linux, use the Java Node Manager, instead of the native version of the Node Manager, by configuring these Node Manager properties in the `domains/your_domain_name/nodemanager/nodemanager.properties` file:

```
NativeVersionEnabled=false
```

```
StartScriptEnabled=true
```

Note: By default, SSL is used with Node Manager. If you want to disable SSL during development, for example, set `SecureListener=false` in your `nodemanager.properties` file.

See Oracle's Configuring Java Node Manager documentation for details.

368.2 Configuring WebLogic

Next, you must set some variables in two WebLogic startup scripts. These variables and scripts are as follows. Be sure to use `set` instead of `export` if you're on Windows.

1. `your-domain/startWebLogic.[cmd|sh]`: This is the Admin Server's startup script.
2. `your-domain/bin/startWebLogic.[cmd|sh]`: This is the startup script for Managed Servers.

Add the following variables to both `startWebLogic.[cmd|sh]` scripts:

```
export DERBY_FLAG="false"
export JAVA_OPTIONS="${JAVA_OPTIONS} -Dfile.encoding=UTF-8 -Duser.timezone=GMT -da:org.apache.lucene... -da:org.aspectj..."
export MW_HOME="/your/weblogic/directory"
export USER_MEM_ARGS="-Xmx2560m -Xms2560m"
```

****Important:**** For Liferay DXP to work properly, the application server JVM must use the ``GMT`` time zone and ``UTF-8`` file encoding.

The ``DERBY_FLAG`` setting disables the Derby server built in to WebLogic, as Liferay DXP doesn't require this server. The remaining settings support Liferay DXP's memory requirements, UTF-8 requirement, Lucene usage, and Aspect Oriented Programming via AspectJ. Also make sure to set ``MW_HOME`` to the directory containing your WebLogic server on your machine. For example:

```
``bash
export MW_HOME="/Users/ray/Oracle/wls12210"
````
```

3. Some of the settings are also found in the `your-domain/bin/SetDomainEnv.[cmd|sh]`. Add the following variables (Windows):

```
set WLS_MEM_ARGS_64BIT=-Xms2560m -Xmx2560m
set WLS_MEM_ARGS_32BIT=-Xms2560m -Xmx2560m
```

or on Mac or Linux:

```
WLS_MEM_ARGS_64BIT="-Xms2560m -Xmx2560m"
export WLS_MEM_ARGS_64BIT

WLS_MEM_ARGS_32BIT="-Xms2560m -Xmx2560m"
export WLS_MEM_ARGS_32BIT
```

4. Set the Java file encoding to UTF-8 in `your-domain/bin/SetDomainEnv.[cmd|sh]` by appending `-Dfile.encoding=UTF-8` ahead of your other Java properties:

```
JAVA_PROPERTIES="-Dfile.encoding=UTF-8 ${JAVA_PROPERTIES} ${CLUSTER_PROPERTIES}"
```

5. You must also ensure that the Node Manager sets Liferay DXP's memory requirements when starting the Managed Server. In the Admin Server's console UI, navigate to the Managed Server you want to deploy Liferay DXP to and select the *Server Start* tab. Enter the following parameters into the *Arguments* field:

```
-Xmx2560m -Xms2560m -XX:MaxMetaspaceSize=512m
```

Click *Save* when you're finished.

Next, you'll set some Liferay DXP-specific properties for your @product@ installation.

### 368.3 Setting Liferay DXP Properties

---

Before installing Liferay DXP, you must set the *Liferay Home* folder's location via the `liferay.home` property in a `portal-ext.properties` file. You can also use this file to override other Liferay DXP properties that you may need.

First, decide which folder you want to serve as Liferay Home. In WebLogic, your domain's folder is generally Liferay Home, but you can choose any folder on your machine. Then create your `portal-ext.properties` file and add the `liferay.home` property:

```
liferay.home=/full/path/to/your/liferay/home/folder
```

Remember to change this file path to the location on your machine that you want to serve as Liferay Home.

Now that you've created your `portal-ext.properties` file, you must put it inside the Liferay DXP WAR file. Expand the @product@ WAR file and place `portal-ext.properties` in the `WEB-INF/classes` folder. Later, you can deploy the expanded archive to WebLogic. Alternatively, you can re-WAR the expanded archive for later deployment. In either case, Liferay DXP reads your property settings once it starts up.

If you need to make any changes to `portal-ext.properties` after Liferay DXP deploys, you can find it in your domain's `autodeploy/ROOT/WEB-INF/classes` folder. Note that the `autodeploy/ROOT` folder contains the Liferay DXP deployment.

Next, you'll install Liferay DXP's dependencies.

### 368.4 Installing Liferay DXP Dependencies

---

You must now install Liferay DXP's dependencies. Recall that earlier you downloaded two ZIP files containing these dependencies. Install their contents now:

1. Unzip the Dependencies ZIP file and place its contents in your WebLogic domain's lib folder.
2. Unzip the OSGi Dependencies ZIP file and place its contents in the Liferay\_Home/osgi folder (create this folder if it doesn't exist).

You must also add your database's driver JAR file to your domain's lib folder.

---

**Note:** Although Hypersonic is fine for testing purposes, **do not** use it for production Liferay DXP instances.

---

---

**Note:** The Liferay DXP Compatibility Matrix specifies supported databases and environments.

---

Next, you'll configure your database.

### 368.5 Database Configuration

---

Use the following procedure if you want WebLogic to manage your database for Liferay DXP. You can skip this section if you want to use @product@'s built-in Hypersonic database.

1. Log in to your AdminServer console.
2. In the *Domain Structure* tree, find your domain and navigate to *Services* → *JDBC* → *Data Sources*.
3. To create a new data source, click *New*. Fill in the *Name* field with Liferay Data Source and the *JNDI Name* field with jdbc/LiferayPool. Select your database type and driver. For example, MySQL is *MySQL's Driver (Type 4) Versions:using com.mysql.cj.jdbc.Driver*. Click *Next* to continue.
4. Accept the default settings on this page and click *Next* to move on.
5. Fill in your database information for your MySQL database.
6. If using MySQL, add the text `?useUnicode=true&characterEncoding=UTF-8&\useFastDateParsing=false` to the URL line and test the connection.

---

**\*\*Tip:\*\*** For more example URLs, see the `\jdbc.default.url`` values in [Database Templates](/docs/7-2/deploy/-/knowledge\_base/d/database-templates).

---

If the connection works, click *\*Next\**.

7. Select the target for the data source and click *Finish*.
8. You must now tell Liferay DXP about the JDBC data source. Create a `portal-ext.properties` file in your Liferay Home directory, and add the line:

```
jdbc.default.jndi.name=jdbc/LiferayPool
```

Alternatively, you can make the above configuration strictly via properties in the `portal-ext.properties` file. Please see the Database Templates for example properties.

Next, you'll configure your mail session.

## 368.6 Mail Configuration

---

If you want WebLogic to manage your mail session, use the following procedure. If you want to use Liferay's built-in mail session (recommended), skip this section.

1. Start WebLogic and log in to your Admin Server's console.
2. Select *Services* → *Mail Sessions* from the *Domain Structure* box on the left hand side of your Admin Server's console UI.
3. Click *New* to begin creating a new mail session.
4. Name the session *LiferayMail* and give it the JNDI name `mail/MailSession`. Then fill out the *Session Username*, *Session Password*, *Confirm Session Password*, and *JavaMail Properties* fields as necessary for your mail server. See the WebLogic documentation for more information on these fields. Click *Next* when you're done.
5. Choose the Managed Server that you'll install Liferay DXP on, and click *Finish*. Then shut down your Managed and Admin Servers.
6. With your Managed and Admin servers shut down, add the following property to your `portal-ext.properties` file in Liferay Home:

```
mail.session.jndi.name=mail/MailSession
```

Liferay DXP references your WebLogic mail session via this property setting. If you've already deployed Liferay DXP, you can find your `portal-ext.properties` file in your domain's `autodeploy/ROOT/WEB-INF/classes` folder.

Your changes take effect upon restarting your Managed and Admin servers.

## 368.7 Deploying Liferay DXP

---

As mentioned earlier, although you can deploy Liferay DXP to a WebLogic Admin Server, you should instead deploy it to a WebLogic Managed Server. Dedicating the Admin Server to managing other servers that run your apps is a best practice.

Follow these steps to deploy Liferay DXP to a Managed Server:

1. Make sure the Managed Server you want to deploy Liferay DXP to is shut down.
2. In your Admin Server's console UI, select *Deployments* from the *Domain Structure* box on the left hand side. Then click *Install* to start a new deployment.
3. Select the Liferay DXP WAR file or its expanded contents on your file system. Alternatively, you can upload the WAR file by clicking the *Upload your file(s)* link. Click *Next*.
4. Select *Install this deployment as an application* and click *Next*.
5. Select the Managed Server you want to deploy Liferay DXP to and click *Next*.
6. If the default name is appropriate for your installation, keep it. Otherwise, give it a name of your choosing and click *Next*.
7. Click *Finish*. After the deployment finishes, click *Save* if you want to save the configuration.
8. Start the Managed Server where you deployed Liferay DXP. `@product@` precompiles all the JSPs and then launches.

Nice work! Now you're running Liferay DXP on WebLogic.

---

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving PhaseOptimizer. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass gatherExternProperties
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass checkControlFlow
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array pattern,
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressions,
current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array pattern,
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expressions]
```

---

# INSTALLING LIFERAY DXP ON WEBSPHERE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

IBM® WebSphere® is a trademark of International Business Machines Corporation, registered in many jurisdictions worldwide.

---

**Tip:** Throughout this installation and configuration process, WebSphere prompts you to click *Save* to apply changes to the Master Configuration. Do | so intermittently to save your changes.

---

For Liferay DXP to work correctly, WebSphere 9 (Fix Pack 11 is the latest) must be installed. You can find more information about this fix pack [here](#).

Please also note that Liferay DXP doesn't support the WebSphere Application Liberty Profile.

---

**Important:** Before installing Liferay DXP, familiarize yourself with preparing for install.

---

Now, download the Liferay DXP WAR and Dependency JARs:

- Liferay DXP WAR file
- Dependencies ZIP file
- OSGi Dependencies ZIP file

Note that the *Liferay Home* folder is important to the operation of Liferay DXP. In Liferay Home, @product@ creates certain files and folders that it needs to run. On WebSphere, Liferay Home is typically [Install Location]/WebSphere/AppServer/profiles/your-profile/liferay.

Without any further ado, get ready to install Liferay DXP in WebSphere!

---

## 369.1 Preparing WebSphere for Liferay DXP

---

When the application server binaries have been installed, start the *Profile Management Tool* to create a profile appropriate for Liferay DXP.

1. Click *Create...*, choose *Application Server*, and then click *Next*.
2. Click the *Advanced* profile creation option and then click *Next*. You need the advanced profile to specify your own values for settings such as the location of the profile and names of the profile, node and host, to assign your own ports, or to optionally choose whether to deploy the administrative console and sample application and also add web-server definitions for IBM HTTP Server. See the WebSphere documentation for more information about these options.

#### Optional Application Deployment



Select the applications to deploy to the WebSphere Application Server environment being created.

Deploy the administrative console (recommended).  
Install a Web-based administrative console that manages the application server. Deploying the administrative console is recommended, but if you deselect this option, the information center contains detailed steps for deploying it after the profile exists.

Deploy the default application.  
Install the default application that contains the Snoop, Hello, and HitCount servlets.

Deploy the Installation Verification Tool application.  
Install the Installation Verification Tool (IVT) to verify that the creation of the application server is successful. IVT scans product log files for errors and verifies core product functionality by starting and monitoring the server process and querying servlets.

Figure 369.1: Choose the Advanced profile option to specify your own settings.

3. Check the box *Deploy the administrative console*. This gives you a web-based UI for working with your application server. Skip the default applications. You'd only install these on a development machine. Click *Next*.
4. Set the profile name and location. Ensure you specify a performance tuning setting other than *Development*, since you're installing a production server. See the WebSphere documentation for more information about performance tuning settings. Click *Next*.
5. Choose node, server, and host names for your server. These are specific to your environment. Click *Next*.
6. Administrative security in WebSphere is a way to restrict who has access to the administrative tools. You may want to have it enabled in your environment so that a user name and password are required to administer the WebSphere server. See WebSphere's documentation for more information. Click *Next*.
7. Each profile needs a security certificate, which comes next in the wizard. If you don't have certificates already, choose the option to generate a personal certificate and a signing certificate and click *Next*.
8. Once the certificates are generated, set a password for your keystore. Click *Next*.



9. Now you can customize the ports this server profile uses. Be sure to choose ports that are open on your machine. When choosing ports, the wizard detects existing WebSphere installations and if it finds activity, it increments ports by one.
10. Choose whether to start this profile when the machine starts. Click *Next*.
11. WebSphere ships with IBM HTTP Server, which is a re-branded version of Apache. Choose whether you want a web server definition, so that this JVM receives requests forwarded from the HTTP server. See WebSphere's documentation for details on this. When finished, click *Next*.
12. The wizard then shows you a summary of what you selected, enabling you to keep your choices or go back and change something. When you're satisfied, click *Next*.

WebSphere then creates your profile and finishes with a message telling you the profile was created successfully. Awesome! Your profile is complete. Now there are a few things you need to configure in your application server.

## 369.2 Configuring the WebSphere Application Server

---

In this version of WebSphere, servlet filters are not initialized on web application startup, but rather, on first access. This can cause problems when deploying certain apps to Liferay DXP. To configure servlet filters to initialize on application startup (i.e., deployment), set the following webcontainer properties in your WebSphere application server:

```
com.ibm.ws.webcontainer.initFilterBeforeInitServlet = true
com.ibm.ws.webcontainer.invokeFilterInitAtStartup = true
```

To set webcontainer properties in the WebSphere application server, follow the instructions here in WebSphere's documentation.

## 369.3 Setting up JVM Parameters for Liferay DXP

---

Next, in the WebSphere profile you created for Liferay DXP, you must set an argument that supports Liferay DXP's Java memory requirements. You'll modify this file:

```
[Install Location]/WebSphere/AppServer/profiles/your-profile/config/cells/your-cell/nodes/your-node/servers/your-server/server.xml
```

Add `maximumHeapSize="2560"` inside the `jvmEntries` tag. For example:

```
<jvmEntries xmi:id="JavaVirtualMachine_1183122130078" ... maximumHeapSize="2560">
```




Figure 369.2: Example of the settings before creating the profile.

**Note:** The JVM parameters used here are defaults intended for initial deployment of production systems. Administrators should change the settings to values that best address their specific environments. These must be tuned depending on need.

Administrators can set the UTF-8 properties in the `<jvmEntries genericJvmArguments=.../>` attribute in `server.xml`. This is required or else special characters will not be parsed correctly. Set the maximum and minimum heap sizes to `2560m` there too. Add the following inside the `jvmEntries` tag:

```
<jvmEntries xmi:id="JavaVirtualMachine_1183122130078" ...genericJvmArguments="-Dfile.encoding=UTF-8 -Duser.timezone=GMT -Xms2560m -Xmx2560m">
```

**Important:** For Liferay DXP to work properly, the application server JVM must use the GMT time zone and UTF-8 file encoding.

---

Alternately, you can set the UTF-8 properties from the WebSphere Admin Console. (See below.)

### 369.4 Removing the secureSessionCookie Tag

---

In the same profile, you should delete a problematic secureSessionCookie tag that can cause Liferay DXP startup errors. Note that this is just a default setting; once Liferay DXP is installed, you should tune it appropriately based on your usage.

In [Install Location]/WebSphere/AppServer/profiles/your-profile/config/cells/your-cell/cell.xml, Delete the secureSessionCookie tag containing xmi:id="SecureSessionCookie\_1".

If this tag is not removed, an error similar to this may occur:

```
WSVR0501E: Error creating component com.ibm.ws.runtime.component.CompositionUnitMgrImpl@d74fa901
com.ibm.ws.exception.RuntimeWarning: com.ibm.ws.webcontainer.exception.WebAppNotLoadedException: Failed to load webapp: Failed to load webapp: SRVE8
```

### 369.5 Installing Liferay DXP's Dependencies

---

You must now install Liferay DXP's dependencies. Recall that earlier you downloaded two ZIP files containing these dependencies. Install their contents now:

1. Unzip the Dependencies ZIP file and place its contents in your WebSphere application server's [Install Location]/WebSphere/AppServer/lib/ext folder. If you have a JDBC database driver JAR, copy it to this location as well.

---

**\*\*Note:\*\*** The [Liferay DXP Compatibility Matrix](<https://web.liferay.com/documents/14/21598941/Liferay+DXP+7.2+Compatibility+Matrix/b6e0f064-db31-49b4-8317-a29d1d76abf7?>) specifies supported databases and environments.

- 
2. From the same archive, copy portlet.jar into [Install Location]/WebSphere/AppServer/javaext for WebSphere 9.0.0.x. WebSphere already contains an older version of portlet.jar which must be overridden at the highest class loader level. The new portlet.jar (version 3) is backwards-compatible.
  3. Unzip the OSGi Dependencies ZIP file and place its contents in the [Liferay Home]/osgi folder (create this folder if it doesn't exist). This is typically [Install Location]/WebSphere/AppServer/profiles/your-profile/liferay/osgi.

### 369.6 Ensuring that Liferay DXP's portlet.jar is loaded first

---

In addition to placing the `portlet.jar` in the correct folder, you must configure the `config.ini` file so that it is loaded first. Navigate to `/IBM/WebSphere/AppServer/configuration/config.ini`.

1. Find the property `com.ibm.CORBA,com.ibm`
2. Insert the property `javax.portlet,javax.portlet.filter,javax.portlet.annotations` after `com.ibm.CORBA` and before `com.ibm`.
3. Save the file.

Once you've installed these dependencies and configured the `config.ini` file, start the server profile you created for Liferay DXP. Once it starts, you're ready to configure your database.

### 369.7 Database Configuration

---

If you want WebSphere to manage the database connections, follow the instructions below. Note this is not necessary if you plan to use Liferay DXP's standard database configuration; in that case, skip this section. See the Using the Built-in Data Sources section for more article.

You'll set your database information in Liferay DXP's setup wizard after the install.

---

**Note:** Although Liferay DXP's embedded database is fine for testing purposes, you **should not** use it for production Liferay DXP instances.

---

1. Start WebSphere.
2. Open the Administrative Console and log in.
3. Click *Resources*  $\boxtimes$  *JDBC Providers*.
4. Select a scope and then click *New*.
5. Select your database type, provider type, and implementation type. If you select a predefined database, the wizard fills in the name and description fields for you. If the database you want to use isn't listed, select *User-defined* from the *Database type* field and then fill in the *Implementation Class Name*. For example, if you use MySQL, select *Database type*  $\rightarrow$  *User-defined*, and then enter `com.mysql.jdbc.jdbc2.optional.MysqlConnectionPoolDataSource` in *Implementation Class Name*. Click *Next* when you are finished.
6. Clear any text in the class path settings. You already copied the necessary JARs to a location on the server's class path. Click *Next*.
7. Review your settings and click *Finish*. The final configuration should look like this:

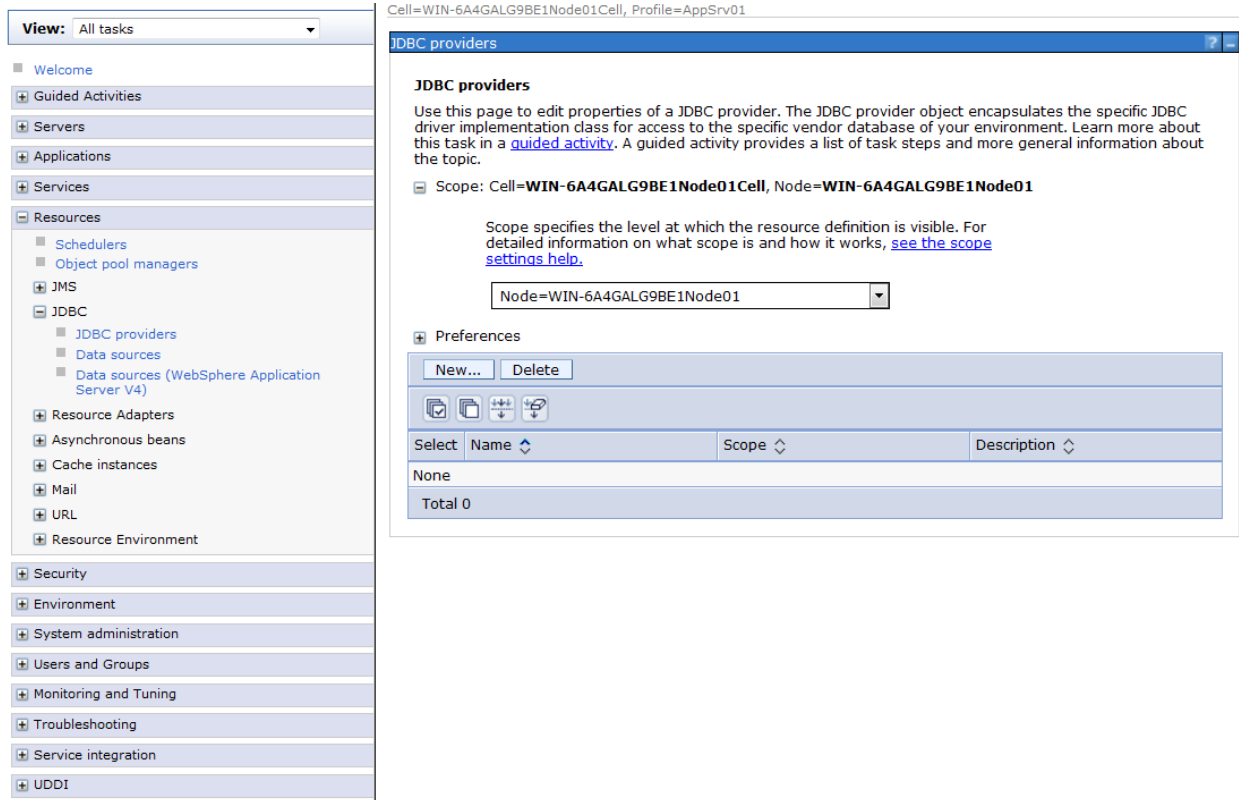


Figure 369.3: WebSphere JDBC providers

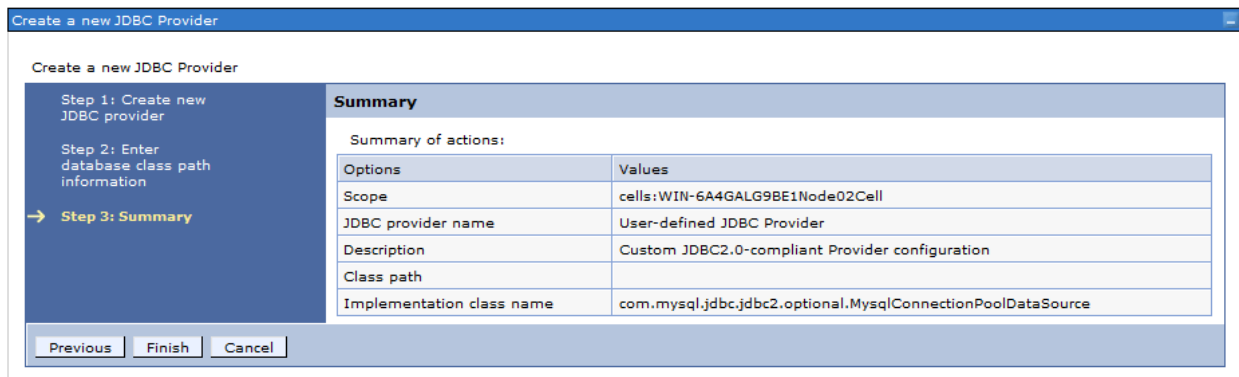


Figure 369.4: Completed JDBC provider configurations.

8. Click your new provider configuration when it appears in the table, and then click *Data Sources* under *Additional Properties*. Click *New*.
9. Enter *liferaydatabasesource* in the *Data source name* field and *jdbc/LiferayPool* in the *JNDI name* field. Click *Next*.
10. Click *Next* in the remaining screens of the wizard to accept the default values. Then review your changes and click *Finish*.
11. Click the data source when it appears in the table and then click *Custom Properties*. Now click the *Show Filter Function* button. This is the second from last of the small icons under the *New* and *Delete* buttons.
12. Type *user* into the search terms and click *Go*.

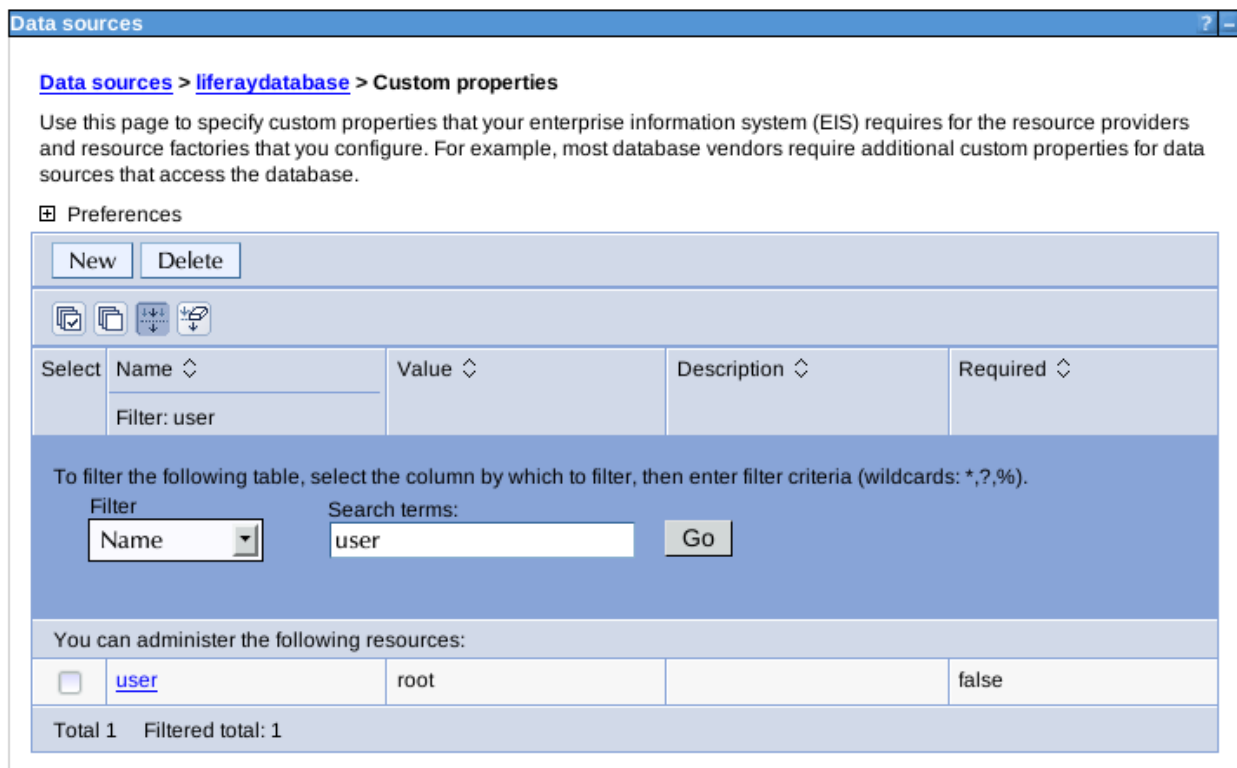


Figure 369.5: Modifying data source properties in WebSphere

13. Select the *user* property and give it the value of the user name to your database. Click *OK* and save to master configuration.
14. Do another filter search for the *url* property. Give this property a value that points to your database. For example, a MySQL URL would look like this:

```
jdbc:mysql://localhost/lportal?useUnicode=true&characterEncoding=UTF-8&serverTimezone=GMT&useFastDateParsing=false
```

**\*\*Tip:\*\*** For more example URLs, see the ``jdbc.default.url`` values in [Database Templates](/docs/7-2/deploy/-/knowledge\_base/d/database-templates).

---

Click *OK* and save to master configuration.

15. Do another filter search for the *password* property. Enter the password for the user ID you added earlier as the value for this property. Click *OK* and save to master configuration.
16. Go back to the data source page by clicking it in the breadcrumb trail. Click the *Test Connection* button. It should connect successfully.

Once you've set up your database, you can set up your mail session.

---

## 369.8 Mail Configuration

---

If you want WebSphere to manage your mail sessions, use the following procedure. If you want to use Liferay DXP's built-in mail sessions, you can skip this section. See the *Configuring Mail* article on how to use Liferay DXP's built-in mail sessions.

---

## 369.9 Creating a WebSphere-Managed Mail Session (Optional)

---

1. Click *Resources*  $\boxtimes$  *Mail*  $\boxtimes$  *Mail Providers*.
2. Click the Built-In Mail Provider for your node and server.
3. Click *Mail Sessions* and then click the *New* button.
4. Give your mail session a name of *liferaymail* and a JNDI name of *mail/MailSession*. Fill in the correct information for your mail server in the sections *Outgoing Mail Properties* and *Incoming Mail Properties*. Click *OK* and then save to the master configuration.
5. Click your mail session when it appears in the table and select *Custom Properties* under the *Additional Properties* section. Set any other JavaMail properties required by your mail server, such as the protocol, ports, whether to use SSL, and so on.
6. Click *Security*  $\boxtimes$  *Global Security* and de-select *Use Java 2 security to restrict application access to local resources* if it is selected. Click *Apply*.

Note that you may also need to retrieve a SSL certificate from your mail server and add it to WebSphere's trust store. See WebSphere's documentation for instructions on this.

### 369.10 Verifying WebSphere Mail Provider

---

To validate that the mail session has been configured correctly, there are a number of ways to test this once the WAR has been deployed, the server has started, and the user has signed in as the system administrator. One quick way to validate is to create a new user with a valid email account. The newly created user should receive an email notification. The logs should display that the SMTP server has been pinged with the correct port number listed.

### 369.11 Enable Cookies for HTTP Sessions

---

WebSphere restricts cookies to HTTPS sessions by default. If you're using HTTP instead, this prevents users from signing in to Liferay DXP and displays the following error in the console:

```
20:07:14,021 WARN [WebContainer : 1][SecurityPortletContainerWrapper:341]
User 0 is not allowed to access URL http://localhost:9081/web/guest/home and
portlet com.liferay.login_web_portlet_LoginPortlet
```

This occurs because Liferay DXP can't use the HTTPS cookie when you use HTTP. The end result is that new sessions are created on each page refresh. Follow these steps to resolve this issue in WebSphere:

1. Click *Application Servers* → *server1* → *Session Management* → *Enable Cookies*
2. De-select *Restrict cookies to HTTPS sessions*
3. Click *Apply*
4. Click *Save*

### 369.12 Enable UTF-8

---

If you did not add the `-Dfile.encoding=UTF-8` property in the `server.xml`, you can do so in the Administrative Console.

1. Click *Application Servers* → *server1* → *Process definition*.
2. Click *Java Virtual Machine* under *Additional Properties*.
3. Enter `-Dfile.encoding=UTF-8` in the *Generic JVM arguments* field.
4. Click *Apply* and then *Save* to master configuration.

Once the changes have been saved, Liferay DXP can parse special characters if there is localized content.



### 369.13 Deploy Liferay DXP

---

Now you're ready to deploy Liferay DXP!

1. In WebSphere's administrative console, click *Applications* → *New Application* → *New Enterprise Application*.
2. Browse to the Liferay DXP .war file, select it, and click *Next*.
3. Leave *Fast Path* selected and click *Next*. Ensure that *Distribute Application* has been checked and click *Next* again.
4. Choose the WebSphere runtimes and/or clusters where you want Liferay DXP deployed. Click *Next*.
5. Select the virtual host to deploy Liferay DXP on and click *Next*.
6. Map Liferay DXP to the root context (/) and click *Next*.
7. Select the *metadata-complete attribute* setting that you want to use and click *Next*.
8. Ensure that you have made all the correct choices and click *Finish*. When Liferay DXP has installed, click *Save to Master Configuration*.

You've now installed Liferay DXP!

### 369.14 Setting the JDK Version for Compiling JSPs

---

Liferay DXP requires that its JSPs are compiled to the Java 8 bytecode format. To ensure that WebSphere does this, shut down WebSphere after you've deployed the Liferay DXP .war file. Navigate to the WEB\_INF folder and add the following setting to the `ibm-web-ext.xml` or in most cases the `ibm-web-ext.xmi` file:

```
<jsp-attribute name="jdkSourceLevel" value="18" />
```

The exact path to the `ibm-web-ext.xmi` file depends on your WebSphere installation location and Liferay DXP version, but here's an example:

```
/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/config/cells/localhostNode01Cell/applications/liferayXX.ear/deployments/liferayXX/liferayXX.war/WEB-INF/ibm-web-ext.xmi
```

Note that the Liferay DXP .war comes pre-packaged with the `ibm-web-ext.xmi` file; this format is functionally the same as .xml and WebSphere recognizes both formats. For more general information on how WebSphere compiles JSPs see IBM's official documentation for WebSphere Application Server 9.0.0.x.

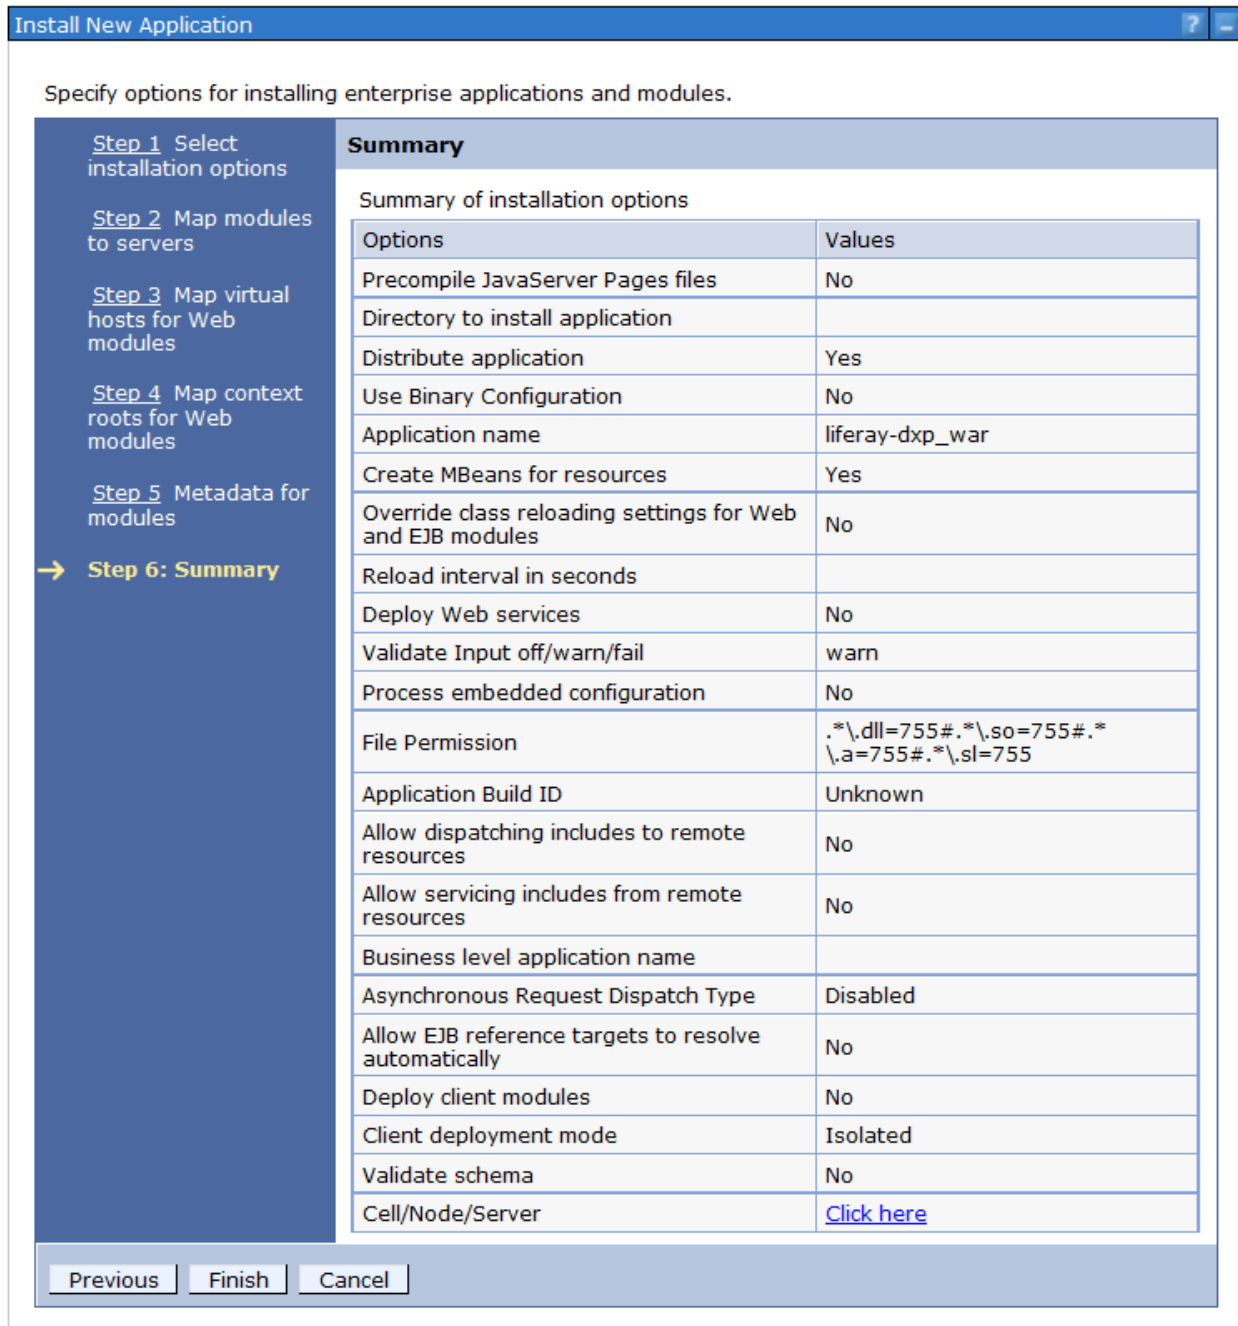


Figure 369.6: Review your deployment options before deploying.

## 369.15 Start Liferay DXP

---

1. If you plan to use Liferay DXP's [setup wizard](/docs/7-2/deploy/-/knowledge\_base/d/installing-product#using-the-setup-wizard), skip to the next step. If you wish to use WebSphere's data source and mail session, create a file called `portal-ext.properties` in your Liferay Home folder. Place the following configuration in the file:

```
jdbc.default.jndi.name=jdbc/LiferayPool
mail.session.jndi.name=mail/MailSession
setup.wizard.enabled=false
```
2. In the WebSphere administrative console, navigate to *Enterprise Applications*, select the Liferay DXP application, and click *Start*. While Liferay DXP is starting, WebSphere displays a spinning graphic.
3. In Liferay DXP's setup wizard, select and configure your database type. Click *Finish* when you're done. Liferay DXP then creates the tables it needs in the database.

Congratulations! You've installed Liferay DXP on WebSphere!

---

After deploying Liferay DXP, you may see excessive warnings and log messages, such as the ones below, involving PhaseOptimizer. These are benign and can be ignored. Make sure to adjust your app server's logging level or log filters to avoid excessive benign log messages.

```
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass gatherExternProperties
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
WARNING: Skipping pass checkControlFlow
May 02, 2018 9:12:27 PM com.google.javascript.jscomp.PhaseOptimizer$NamedPass process
INFO: pass supports: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array pattern,
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expression,
current AST contains: [ES3 keywords as identifiers, getters, reserved words as properties, setters, string continuation, trailing comma, array pattern,
scoped function declaration, class, computed property, const declaration, default parameter, destructuring, extended object literal, for-
of loop, generator, let declaration, member declaration, new.target, octal literal, RegExp flag 'u', RegExp flag 'y', rest parameter, spread expression]
```



## ACTIVATING LIFERAY DXP

---

This document has been updated and ported to Liferay Learn and is no longer maintained here. There are two ways to activate your Liferay DXP instance:

- With an XML activation key that you request and receive from Liferay Support.
- Online activation through Liferay Connected Services (LCS). Liferay DXP 7.0 introduced LCS as a way to activate Liferay DXP instances. LCS can also install fix packs, monitor each instance's performance, and help administrators automatically manage Liferay DXP subscriptions. See the LCS documentation for instructions on activating your instances with LCS.

---

**Note:** You must use LCS for activation of Elastic subscriptions. Otherwise, you don't have to use LCS for activation. You can instead request an XML activation key from Liferay Support.

---



---

## SETTING UP MARKETPLACE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay Marketplace is more than just a store for Liferay applications. Under the hood, it provides both the store and Liferay DXP's application deployment features. For this reason, you must ensure that Marketplace can run and configure itself.

Here are some scenarios to work around to ensure Marketplace works successfully:

- Server is Firewalled without Access to the Internet
- Limited Database Access

The firewall scenario is discussed first.

---

### **371.1 Server is Firewalled without Access to the Internet**

---

Your server might be behind a firewall that prevents access to the Internet. Or your security policy might not allow direct download and installation from the Internet. In these cases, you have two options:

1. From an Internet-enabled computer, download the Marketplace plugin. Then deploy the plugin (.lpkg file) by copying it into the deploy folder in Liferay Home.
2. Alternately, once you have the downloaded .lpkg file, deploy it using the App Manager.

Next you'll learn how to work around database access restrictions.

---

### **371.2 Limited Database Access**

---

Some production environments do not have the necessary database permissions for Liferay DXP, apps, modules, and plugins to maintain their tables. In these cases:

1. Grant the Liferay DXP database user temporary full rights to the database.
2. Install Liferay DXP and start it so that it populates its database.
3. Once the database is created, remove the permissions for creating tables and dropping tables from the Liferay DXP database user.

See the database preparation instructions for more information. Note that many sophisticated Liferay DXP apps—not just the Marketplace app—require new tables when deployed. If your environment restricts database access, you may need to repeat the above steps whenever you deploy a new app.

You've prepared Liferay DXP for installing Marketplace and additional apps.



# TRIAL PLUGIN INSTALLATION

This document has been updated and ported to <https://learn.liferay.com/dxp/latest/en/system-administration/installing-and-managing-apps/installing-apps/accessing-ee-plugins-during-a-trial-period.html>) and is no longer maintained here.

For Liferay customers who are evaluating Liferay DXP on a trial basis, **the plugins can be accessed from within the *Apps* → *Store* (i.e., Marketplace) section of the Control Panel in your product installation.**

## 372.1 Installation Process

Follow the steps below to install a trial plugin:

1. Register a `liferay.com` account (LRDC) account by visiting Liferay’s home page (if necessary). Do this by clicking *Sign In/Create Account* button from the top right Profile button.




Figure 372.1: Hover over the Profile button and click *Sign In/Create Account*.

2. Start your Liferay DXP instance (trial license is OK).
3. After signing in as an Admin in your Liferay DXP trial server, go to the Control Panel → *Apps* → *Store* and sign in to the Marketplace using your `liferay.com` (LRDC) account credentials. Authorize Marketplace to access your local account.
4. Once signed into the Store, click on the *Purchased* link, and then click on the *EE* tab.




Figure 372.2: Click the *Store* link and authorize Marketplace to access your local account.

Here you can see a list of Liferay DXP plugins that are installed, as well as options to update or install certain plugins. See *Using the Liferay Marketplace* for details.

Next are answers to some common questions.

## 372.2 FAQ

---

**Q:** Where are the *Liferay DXP Trial Plugins*?

**A:** There is no such thing. The Liferay DXP plugins in Liferay Marketplace are the same ones that you get to try out with your Liferay DXP trial license for your portal. The Liferay DXP license (trial or official @product@ subscriber) gives you access to the Liferay DXP plugins. Also, there is no difference code-wise or release-wise between a Liferay DXP trial installation and a regular @product@ non-trial installation. The only difference is the license.

**Q:** Why can't I go to [liferay.com/marketplace](http://liferay.com/marketplace)? Why can't I *purchase* from the Marketplace site?

**A:** DXP trial users must use the Marketplace from within the product's Control Panel (instructions above). You do not need to *purchase* any DXP plugins because if you access Marketplace from within the Control Panel, Marketplace sees that you have a DXP license installed and gives access to DXP plugins. Official DXP subscription customers (i.e., non trial) can log into [liferay.com](http://liferay.com) with their designated DXP subscriber login and access all DXP plugins through the Marketplace website.

**Q:** Why are the plugins under the *Purchased* tab? If I click on the *DXP Marketplace* link, it does not let me get the DXP plugins.

**A:** Once you're signed into the Store, click on the *Purchased* tab, then click on the *EE* tab.

**Q:** What happens when DXP trial customers become official Liferay Digital Experience subscribers?

**A:** They can still complete the above process, or they can also visit the Liferay Marketplace website.

**Q:** Do DXP trial customers get the DXP source code?

**A:** No, they can only install the plugin. The DXP source code becomes available once they are official Liferay DXP Enterprise subscribers.

**Q:** Can this process of installing DXP plugins be used from Liferay Portal CE (Community Edition)?

**A:** No, the Marketplace must detect that you are running Liferay DXP.



---


## DOCUMENT REPOSITORY CONFIGURATION

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

You can configure file storage in several ways. Each option is a *store* which can be configured through the `portal-ext.properties` file by setting the `dl.store.impl=` property.

The default store is called Simple File Store. It stores documents and media files on a file system (local or mounted). The store's default root folder is `[Liferay Home]/data/document_library`. You can specify a different root directory from within System Settings.

1. Access System Settings by opening the *Menu* () and navigating to *Control Panel* ▾ *Configuration* ▾ *System Settings*.
2. In the *Platform* section, click *File Storage*. The File Storage page appears.
3. Click *Simple File System Store*.
4. For the store's *Root directory* value, specify its absolute path or its path relative to Liferay Home.
5. Click the *Save* button.

The document library store switches immediately to the new folder.

You can also use an entirely different method for storing documents and media files:

**Simple File System Store:** uses the file system (local or a mounted share) to store files.

**Advanced File System Store:** in addition to using the file system (local or a mounted share) to store files, Advanced File System Store nests the files into folders by version, for faster performance and to store more files.

**S3 Store (Amazon Simple Storage):** uses Amazon's cloud-based storage solution.

**DBStore (Database Storage):** stores the files in the database. DBStore's file (stored as a blob) size is 1 gigabyte. To store files larger than 1 gigabyte, use Simple File System Store or Advanced File System Store.

These articles explain details for each one.

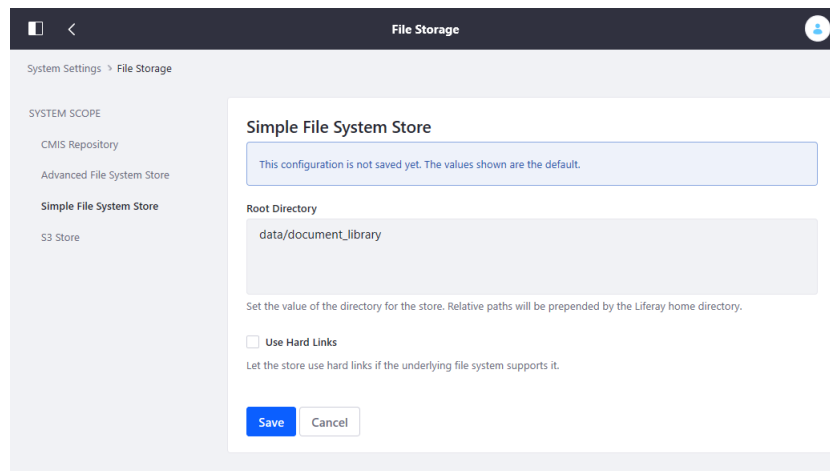


Figure 373.1: The File Storage page in System Settings lets you configure document repository storage.

---

## USING THE SIMPLE FILE SYSTEM STORE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The simple file storage implementation is the default store. It uses a local folder to store files. You can use the file system for your clustered configuration, but the folder you're pointing to must be shared by all nodes and handle concurrent requests and file locking. For this reason, you need to use a Storage Area Network or a clustered file system.

The file system store was the first store used in Liferay DXP and is heavily bound to the Liferay DXP database. By default, documents are stored in a `document_library` subfolder of the data folder. Of course, you can change this path to anything you want in System Settings.

The Simple File System store uses this folder path format for storing documents:

```
/companyId/folderId/numericFileEntryName/versionNumber
```

The first folder name is the site's company ID. The second folder name is the Documents and Media folder's ID where the document resides. The third folder name is the document's numeric file entry name. Finally, the fourth name is a version number used for storing multiple versions of the document.

---

**Note:** A document's numeric file entry name is distinct from the document ID; don't confuse the two! Each has an independent counter. The numeric file entry name is used in the folder path for storing the document but the document ID is not. The numeric file entry name is in the name column of the `DLFileEntry` table in Liferay DXP's database; the document ID is in the `fileEntryId` column of the same table.

---

**Warning:** If a database transaction rollback occurs in the Document Library, file system changes that have occurred since the start of the transaction aren't reversed. Inconsistencies between Document Library files and those in the file system store can occur and may require manual synchronization.

---

The Simple File System Store binds documents very closely to Liferay DXP, and may not be exactly what you want. If you've been using the default settings for a while and must migrate your documents, there's a migration utility in the Control Panel in *Server Administration* → *Data Migration*. The utility facilitates moving documents from one store implementation to another.





## USING THE ADVANCED FILE SYSTEM STORE

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The advanced file system store is similar to the simple file system store (the default store). Like that store, it saves files to the local file system—which, of course, could be a remote file system mount. It uses a slightly different folder structure to store files, which is pictured below.




Figure 375.1: The advanced file system store creates a more nested folder structure than the file system store.

So what makes the advanced file system store *advanced*? Several operating systems have limitations on the number of files that can be stored in a particular folder. The advanced file system store overcomes this limitation by programmatically creating a structure that can expand to millions of files, by alphabetically nesting the files in folders. This not only allows for more files to be stored, but also improves performance as there are fewer files stored per folder.

The same rules apply to the advanced file system store as apply to the simple file system store. To cluster this, you must point the store to a network mounted file system that all the nodes can access, and that networked file system must support concurrent requests and file locking. Otherwise, you may experience data corruption issues if two users attempt to write to the same file at the same time from two different nodes.

To use the advanced file system store, follow these steps:

1. Configure `portal-ext.properties` with this property:

```
dl.store.impl=com.liferay.portal.store.file.system.AdvancedFileSystemStore
```

2. Restart Liferay DXP.
3. In the Control Panel, navigate to *Configuration* → *System Settings* → *File Storage*.
4. In the *Advanced File System Store* screen, configure the store your way.
5. Click *Save*.

Liferay DXP is using the advanced file system store.

---

**Warning:** If a database transaction rollback occurs in the Document Library, file system changes that have occurred since the start of the transaction aren't reversed. Inconsistencies between Document Library files and those in the file system store can occur and may require manual synchronization.

---

You may decide the advanced file system store for whatever reason doesn't serve your needs. If this is the case, you can of course mount other file systems into the documents and media library. In addition to this, you can also redefine the Liferay DXP store to use one of the other supported protocols. S3 store is next.

---

## USING AMAZON SIMPLE STORAGE SERVICE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Amazon's simple storage service (S3) is a cloud-based storage solution that you can use with Documents and Media. All you need is an account, and you can store your documents to the cloud from all nodes, seamlessly.

When you sign up for the service, Amazon assigns you unique keys that link you to your account. In Amazon's interface, you can create "buckets" of data optimized by region.

Here are the steps for configuring Liferay DXP to use your S3 account for file storage:

1. Amazon S3 requires a SAXParser from the application server to operate. If you are using an app server like Apache Tomcat that have one, you must include this property in a `system-ext.properties` file:

```
org.xml.sax.driver=com.sun.org.apache.xerces.internal.parsers.SAXParser
```

2. Place your `system-ext.properties` file in a folder that resides in your Liferay DXP installation's class path (e.g., `/WEB-INF/classes/`).
3. Set the following property in a `portal-ext.properties` file in your Liferay Home folder:

```
dl.store.impl=com.liferay.portal.store.s3.S3Store
```

4. Restart Liferay DXP.
5. In the Control Panel, navigate to *Configuration* → *System Settings* → *File Storage*.
6. In the *S3 Store Configuration* screen, configure the store your way.
7. Click *Save*.

Your Liferay DXP instance is using the Amazon S3 store.

---

**Warning:** If a database transaction rollback occurs in a Document Library that uses a file system based store, file system changes that have occurred since the start of the transaction aren't reversed. Inconsistencies between Document Library files and those in the file system store can

occur and may require manual synchronization. All stores except DBStore are vulnerable to this limitation.

---

**Note:** No action is required to support AWS Signature Version 4 request authorization.

---

Consult the Amazon Simple Storage documentation for additional details on using Amazon's service.

## USING THE DBSTORE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

You can store Documents and Media files in your Liferay DXP database using DBStore. DBStore's maximum file (stored as a blob) size is 1 gigabyte. To store files larger than that, use Simple File System Store or Advanced File System Store.

Here are the DBStore configuration steps:

1. Set the following property in a `portal-ext.properties` file in your Liferay Home folder:

```
dl.store.impl=com.liferay.portal.store.db.DBStore
```

2. Restart Liferay DXP.

Documents and Media now uses Liferay DXP's database via DBStore.



## CONFIGURING LIFERAY DXP

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Once you have Liferay DXP installed, it's time to configure it to the specifics of your environment. This means doing things like setting the time zone and language, configuring mail, configuring a cluster, configuring a Content Delivery Network, tuning, and more. These topics and more are discussed here.





---

## CONFIGURING MAIL

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay DXP uses a mail server and SMTP to email notifications. @product@'s built-in mail session is the easiest way to configure mail and it's recommended. You can configure the built-in mail session before or after deploying Liferay DXP. Alternatively, you can configure Liferay DXP to use a mail session on the application server.

Creating a mail session in Liferay DXP or on the application server requires this information:

- Incoming POP Server and port
- POP User Name
- POP Password
- Outgoing SMTP Server and port
- SMTP User Name
- SMTP Password
- All JavaMail properties you want to use

Built-in mail session setup is recommended and easiest.

---

### **379.1 Configuring Liferay DXP's Built-in Mail Session**

---

The built-in mail session setup can be done using either of these methods:

- Control Panel
- Portal properties

---

### **379.2 Built-in Mail Session in the Control Panel**

---

After deploying Liferay DXP, you can configure the mail session from the Control Panel.

1. Sign in as the administrative user (the user you specified on the Basic Configuration page).
2. Navigate to *Control Panel* » *Configuration* » *Server Administration* » *Mail*.
3. Fill out the form. You're asked for the following information:

**Incoming POP Server:** The hostname for a server running the Post Office Protocol. Liferay DXP checks this mailbox for incoming messages, such as message board replies.

**Incoming Port:** The port on which the POP server is listening.

**Use a Secure Network Connection:** Use an encrypted connection when connecting to the POP server.

**User Name:** The user ID Liferay DXP should use to log into the POP server.

**Password:** The password Liferay DXP should use to log into the POP server.

**Outgoing SMTP Server:** The hostname for a server running the Simple Mail Transfer Protocol. Liferay DXP uses this server to send emails, such as password change emails and other notifications.

**Outgoing Port:** The port on which the SMTP server is listening.

**Use a Secure Network Connection:** Use an encrypted connection when connecting to the SMTP server.

**User Name:** The user ID Liferay DXP should use to log into the SMTP server.

**Password:** The password Liferay DXP should use to log into the SMTP server.

**Manually specify additional JavaMail properties to override the above configuration:** If there are additional properties you need to specify, supply them here.

4. Click *Save*.

Liferay DXP connects to the mail session immediately.

### 379.3 Built-in Mail Session Portal Properties

---

If you prefer specifying your mail session offline or before deploying Liferay DXP, use portal properties.

1. Create a `portal-ext.properties` file, if you haven't already created one.
2. Copy these default property settings into your `portal-ext.properties` file:

```
mail.session.mail=false
mail.session.mail.pop3.host=localhost
mail.session.mail.pop3.password=
mail.session.mail.pop3.port=110
mail.session.mail.pop3.user=
mail.session.mail.smtp.auth=false
mail.session.mail.smtp.host=localhost
mail.session.mail.smtp.password=
mail.session.mail.smtp.port=25
mail.session.mail.smtp.user=
mail.session.mail.store.protocol=pop3
mail.session.mail.transport.protocol=smtp
```

3. Replace the default mail session values with your own.
4. Put the `portal-ext.properties` file into your `LIFERAY_HOME`, once you've established it based on your installation.

Liferay DXP connects to the mail session on the next startup.

### 379.4 Configuring a Mail Session on the Application Server

---

You can manage a mail session for Liferay DXP on your application server. Here's how:

1. Create a mail session on your application server, following your application server documentation.
2. Point Liferay DXP to that mail session using the Control Panel or portal properties. Here are instructions for both:
  - Configure the JNDI name in the *Mail* page at *Control Panel* > *Configuration* > *Server Administration* > *Mail*.
  - Set a `mail.session.jndi.name` portal property in a `[LIFERAY_HOME]/portal-ext.properties` file. Here's an example property:

```
mail.session.jndi.name=mail/MailSession
```

Lastly, configure your instance's email senders.

### 379.5 Configuring default email senders

---

Email senders are the default name and email address Liferay DXP uses to send administrative emails and announcement emails.

Default email senders are configured in the `portal-ext.properties` file.

- Admin email configuration:

```
admin.email.from.name=Joe Bloggs
admin.email.from.address=test@domain.invalid
```

- Announcements email configuration:

```
announcements.email.to.name=
announcements.email.to.address=noreply@domain.invalid
```

1. Replace the names and email addresses above with your values.

---

**Note:** Following emails are blacklisted by default and cannot be used in any Liferay DXP installation:

- noreply@liferay.com
- test@liferay.com
- noreply@domain.invalid
- test@domain.invalid

If you use them, Liferay DXP logs a WARN trace:

Email xxx will be ignored because it is included in mail.send.blacklist

---

2. Restart your server.

Congratulations on configuring mail for Liferay DXP.

---

# LOCALES AND ENCODING CONFIGURATION

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

You can display content based on language, time zone, “right to left” (that is, languages such as Hebrew, Arabic, and Persian), and you can localize user names and titles. Administrators can localize specific core UI messages so that the messages display in certain languages.

---

## 380.1 Time Zones

---

You can set time zones in the Control Panel and theoretically in the JVM (but this must be set to GMT: see below).

Time zone configuration and default language customization are done in the Control Panel, at the Instance level.

1. Navigate to the *Control Panel* → *Configuration*.
2. Click *Instance Settings*.
3. Click on the *Miscellaneous* tab.

The central left and right arrows let you add or remove available languages and locales. You can also set these as properties in your `portal-ext.properties` file in your Liferay Home folder. The `portal.properties` reference document’s *Company* section defines the default locale. The *Languages and Time Zones* section defines the available and current locales.

```
company.default.locale=en_GB
```

---

**Note:** The `company.default.locale` portal property is only intended for use on initial startup. To change the language settings on an existing instance, open the Control Panel and navigate to *Configuration* → *Instance Settings* and select the *Localization* category under the *PLATFORM* heading.

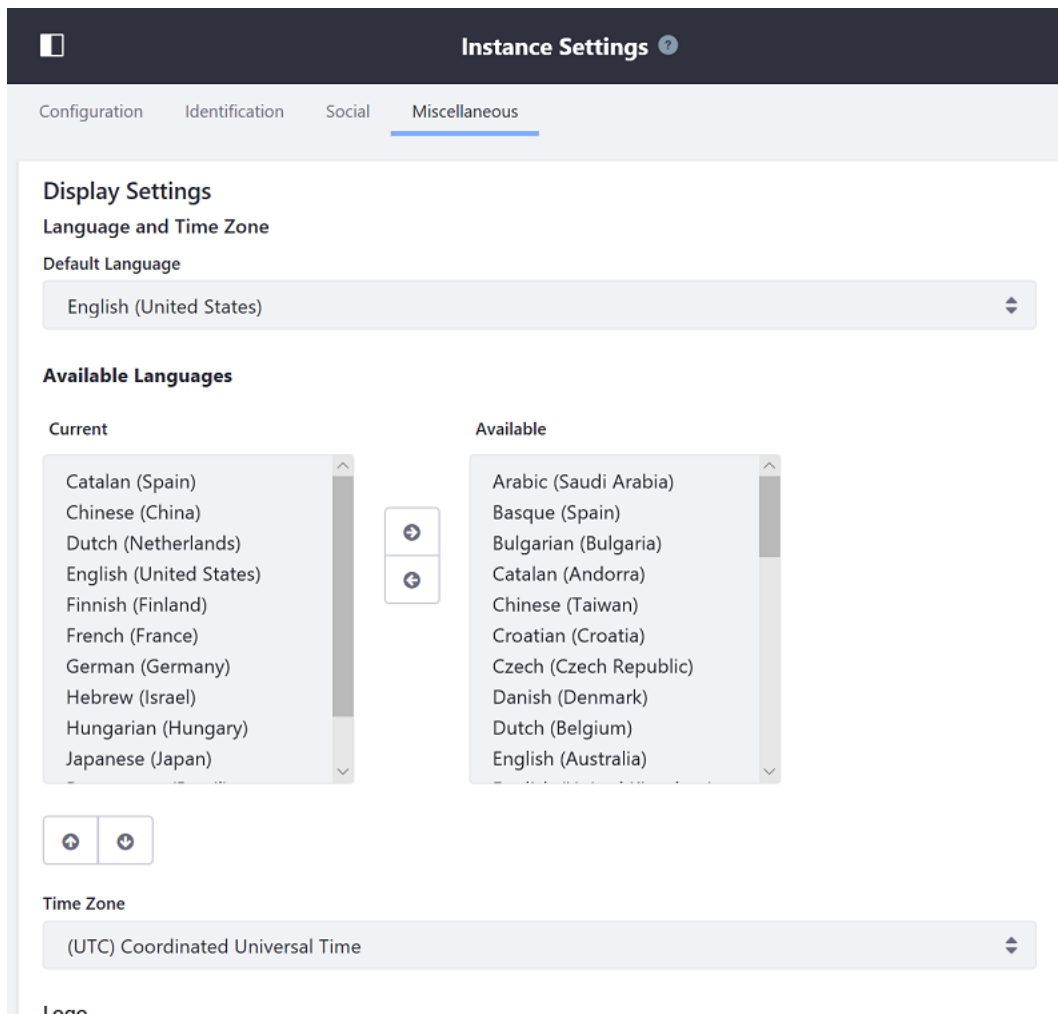


Figure 380.1: You can change the default and available languages and the time zone in Instance Settings.

Under the Language entry you can change the default language, as well as define the current locales.

---

As an example, the above property changes the locale to English, Great Britain.

## 380.2 Set the JVM Time Zone to GMT

---

If you set the time zone in the JVM, it causes issues such as Calendar Events and Web Content articles displaying the wrong dates. This happens because the system assumes each date stored in the database is stored in GMT time. When the system needs to display one stored date to the end users, the display date is calculated by the application server's current date. This date is affected by the configured JVM level time zone and the stored GMT format date. To make sure the display date

is calculated correctly, the time zone must be configured to GMT at the JVM level. Otherwise, an incorrect time zone offset at the JVM level causes the display date to be wrongly calculated and displayed.

### 380.3 Friendly URLs and Locales

---

In addition to configuring instance settings, you can also define unique URLs for specific languages using the `I18nServlet` by editing Portal's `web.xml` file:

```
<servlet-mapping>
 <servlet-name>I18n Servlet</servlet-name>
 <url-pattern>/ar/*</url-pattern>
</servlet-mapping>
.
.
<servlet-mapping>
 <servlet-name>I18n Servlet</servlet-name>
 <url-pattern>/de/*</url-pattern>
</servlet-mapping>
```

The defaults should be sufficient for nearly all circumstances. Because `web.xml` changes require stopping and possibly redeploying Liferay DXP (depending on your app server), test the defaults and make sure you really need to modify these settings. If you're clustered, you must make these changes on all nodes.

### 380.4 Modifying Language Keys

---

Developers can add or modify certain core UI messages (e.g. *Your request completed successfully.*) by modifying the language keys that ship by default.

### 380.5 Right to Left

---

For languages that are displayed right to left, use the following language properties settings:

```
lang.dir=rtl
lang.line.begin=right
lang.line.end=left
```

To display right to left by default, override these properties globally.

## 380.6 Localizing User Names

---

Users can change the prefix and suffix values for a locale. For example, for Spanish, the `language_es.properties` file contains these values:

```
lang.user.name.field.names=prefix,first-name,last-name
lang.user.name.prefix.values=Sr,Sra,Sta,Dr,Dra
lang.user.name.required.field.names=last-name
```

For more information, see [Using Liferay Language Settings](#).

## 380.7 Related Topics

---

- Using Liferay Language Settings
- Overriding Global Language Keys
- Overriding a Module's Language Keys



---

## LIFERAY DXP CLUSTERING

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay DXP can serve everything from the smallest to the largest web sites. Out of the box, it's configured optimally for a single server environment. If one server isn't sufficient to serve your site's high traffic needs, Liferay DXP scales to the size you need.

Liferay DXP works well in clusters of multiple machines (horizontal cluster) or in clusters of multiple VMs on a single machine (vertical cluster), or any mixture of the two. Once you have Liferay DXP installed on more than one application server node, there are several optimizations that must be made. At a minimum, Liferay DXP should be configured in the following way for a clustered environment:

1. All nodes should point to the same database or database cluster.
2. Documents and Media repositories must have the same configuration and be accessible to all nodes of the cluster.
3. Search should be on a separate search server that is optionally clustered.
4. Cluster Link must be enabled so the cache replicates across all nodes of the cluster.
5. Applications must be auto-deployed to each node individually.

Many of these configuration changes can be made by adding or modifying properties in your `portal-ext.properties` file. Remember that this file overrides the defaults in the `portal.properties` file. It's a best practice to copy the relevant section you want to modify from `portal.properties` into your `portal-ext.properties` file, and then modify the values there.

---

**Note:** This documentation describes a Liferay DXP-specific cluster configuration without getting into specific implementations of third party software, such as Java EE application servers, HTTP servers, and load balancers. Please consult your documentation for those components of your cluster to configure those components. Before creating a Liferay DXP cluster, make sure your OS is not defining the hostname of your box to the local network at 127.0.0.1.

---

Each step defined above is covered below to give you a step by step process for creating your cluster. Start with making all Nodes point to the same database.

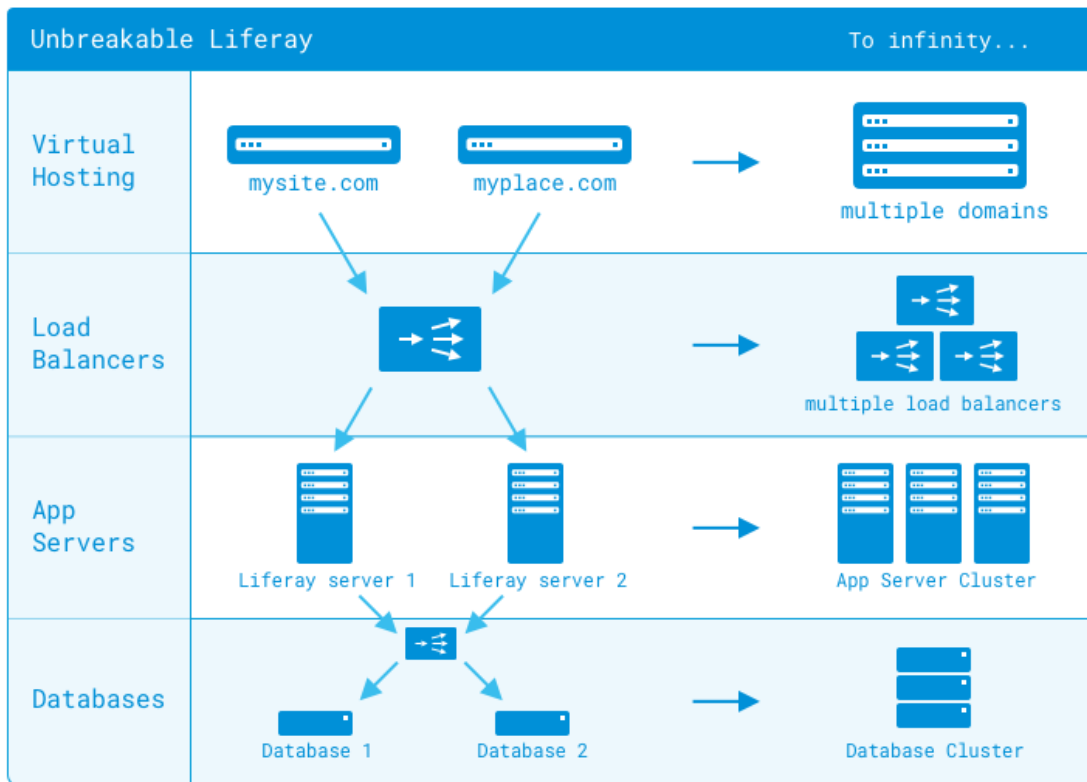


Figure 381.1: Liferay DXP is designed to scale to as large an installation as you need.

---

# POINT ALL NODES TO THE SAME LIFERAY DXP DATABASE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Each node should have a data source that points to one Liferay DXP database (or a database cluster) that all the nodes share. This means, of course, Liferay DXP cannot (and should not) use the embedded HSQL database that is shipped with the bundles (but you already knew that, right?). And, of course, the database server should be on a separate system from the Liferay DXP server.

---

## **382.1 Read-Writer Database Configuration**

---

To improve database performance, you can use a read-writer database configuration. Instead of using the same data source for read and read-write operations, this strategy uses a separate data source for each operation type. DXP's Aspect Oriented Programming (AOP) transaction infrastructure directs read transactions to the read data source and read-write transactions to the write data source.

Connections to separate read and read-write data sources are configured using JDBC or JNDI Portal Properties (e.g., in a `portal-ext.properties` file), as explained in the following sections. The data sources should use separate instances of the DXP database, where the read-write database instance is replicated to the read database instance.

---

## **382.2 JDBC**

---

Edit your `portal-ext.properties` file following these steps to connect directly to your separate read and write data sources using JDBC:

1. Set the default connection pool provider. For provider information, see the JDBC properties reference. The default setting specifies HikariCP as the pool provider:

```
jdbc.default.liferay.pool.provider=hikaricp
```

2. Configure JDBC connections to your separate read and write data sources. Here's an example:

```
jdbc.read.driverClassName=[place your driver name here]
jdbc.read.url=[place the URL to your "read" database here]
jdbc.read.username=[place your user name here]
jdbc.read.password=[place your password here]

jdbc.write.driverClassName=[place your driver name here]
jdbc.write.url=[place the URL to your "read-write" database here]
jdbc.write.username=[place your user name here]
jdbc.write.password=[place your password here]
```

For example JDBC connection values, please see Database Templates.

3. Configure DXP to use the write data source (the data source whose prefix is `jdbc.write.`) to create the Counter data source. A separate data source is always dedicated to the counter.

```
counter.jdbc.prefix=jdbc.write.
```

4. Optionally validate the data connections to make sure bad connections are handled gracefully.

Some connection pools used with JDBC4 (check your driver's JDBC version) validate connections automatically. Other connection pools may require additional, vendor-specific connection validation properties—specify them in a Portal Properties file. Refer to your connection pool provider documentation for connection validation details.

5. Enable the read-writer database configuration by copying the default `spring.configs` and `spring.infrastructure.configs` Portal Properties to your `portal-ext.properties` file and adding the following Spring configuration file paths to them.

Add to `spring.configs`:

```
META-INF/dynamic-data-source-spring.xml
```

Add to `spring.infrastructure.configs`:

```
META-INF/dynamic-data-source-infrastructure-spring.xml
```

For more information, see the Spring configuration Portal Properties.

### 382.3 JNDI

---

Edit your `portal-ext.properties` file following these steps to connect to your read and write data sources on your app server using JNDI:

1. Set your read and write JNDI data source user names and passwords.

```
jdbc.read.jndi.name=[place your "read" data source JNDI name here]

jdbc.read.username=[place your user name here]
jdbc.read.password=[place your password here]

jdbc.write.jndi.name=[place your "read-write" data source JNDI name here]

jdbc.write.username=[place your user name here]
jdbc.write.password=[place your password here]
```

2. Configure DXP to use the write data source (the data source whose prefix is `jdbc.write.`) to create the Counter data source. A separate data source is always dedicated to the counter.

```
counter.jdbc.prefix=jdbc.write.
```

3. Optionally validate the data connections to make sure bad connections are handled gracefully. Some connection pools used with JDBC4 (check your driver's JDBC version) validate connections automatically. Other connection pools may require additional, vendor-specific connection validation properties—specify them in a Portal Properties file. Refer to your connection pool provider documentation for connection validation details.
4. Enable the read-writer database configuration by copying the default `spring.configs` and `spring.infrastructure.configs` Portal Properties to your `portal-ext.properties` file and add the following Spring configuration file paths to them.

Add to `spring.configs`:

```
META-INF/dynamic-data-source-spring.xml
```

Add to `spring.infrastructure.configs`:

```
META-INF/dynamic-data-source-infrastructure-spring.xml
```

For more information, see the Spring configuration Portal Properties.

DXP uses a read data source, a write data source, and a counter data source the next time it starts.



## CONFIGURE DOCUMENTS AND MEDIA THE SAME FOR ALL NODES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

In a cluster, Documents and Media must use the same document repository configuration on all nodes.

Note if you are using the File System or Advanced File System stores, the file system must be accessible from all nodes (i.e., a network share), support concurrent requests, and file locking.

**Checkpoint:** Verify sharing works by executing these steps:

1. On Node 1 upload a document to the Documents and Media.
2. On Node 2 download the document. The download should be successful.
3. Repeat the test with reversed roles.





---

## CLUSTERING SEARCH

---

Search should always run on a separate environment from your Liferay DXP server. Liferay DXP supports Elasticsearch, and Solr. Both can also be clustered.

For more information on how to cluster Elasticsearch, see Elasticsearch's distributed cluster setup.

Once Liferay DXP servers have been properly configured as a cluster and the same for Elasticsearch, change Liferay DXP from *embedded* mode to *remote* mode. On the first connection, the two sets of clustered servers communicate with each other the list of all IP addresses; in case of a node going down, the proper failover protocols enable. Queries and indexes can continue to be sent for all nodes.

For more information on how to cluster Solr, see Apache Solr Cloud documentation.

Once Liferay DXP servers have been properly configured as a cluster, deploy the Liferay Connector to Solr on all nodes. (This app is available for download from Liferay Marketplace) Create a Solr Cloud (cluster) managed by *Apache Solr Zookeeper*. Connect the Liferay DXP cluster to Zookeeper and finish the final configurations to connect the two clusters.



---

## ENABLING CLUSTER LINK

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Enabling Cluster Link automatically activates distributed caching. The cache is distributed across multiple Liferay DXP nodes running concurrently. Cluster Link does Ehcache replication. The Ehcache global settings are in the `portal.properties` file.

By default Liferay does not copy cached entities between nodes. If an entity is deleted or changed, for example, Cluster Link sends a *remove* message to the other nodes to invalidate this entity in their local caches. Requesting that entity on another node results in a cache *miss*; the entity is then retrieved from the database and put into the local cache. Entities added to one node's local cache are not copied to local caches of the other nodes. An attempt to retrieve a new entity on a node which doesn't have that entity cached results in a cache *miss*. The miss triggers the node to retrieve the entity from the database and store it in its local cache.

Here are the Cluster Link topics:

- Enabling Cluster Link
- Multicast Over UDP
- Unicast Over TCP
- Using Different Control and Transport Channel Ports
- Modifying the Cache Configuration with a Module
- Conclusion

---

### 385.1 Enabling Cluster Link

---

To enable Cluster Link, add this portal property to a `portal-ext.properties` file:

```
cluster.link.enabled=true
```

The Cluster Link portal properties provide a default configuration that you can override to fit your needs.

Many of the defaults use `localhost`, instead of a real address. In some configurations, however, `localhost` is bound to the internal loopback network (`127.0.0.1` or `:::1`), rather than the host's real

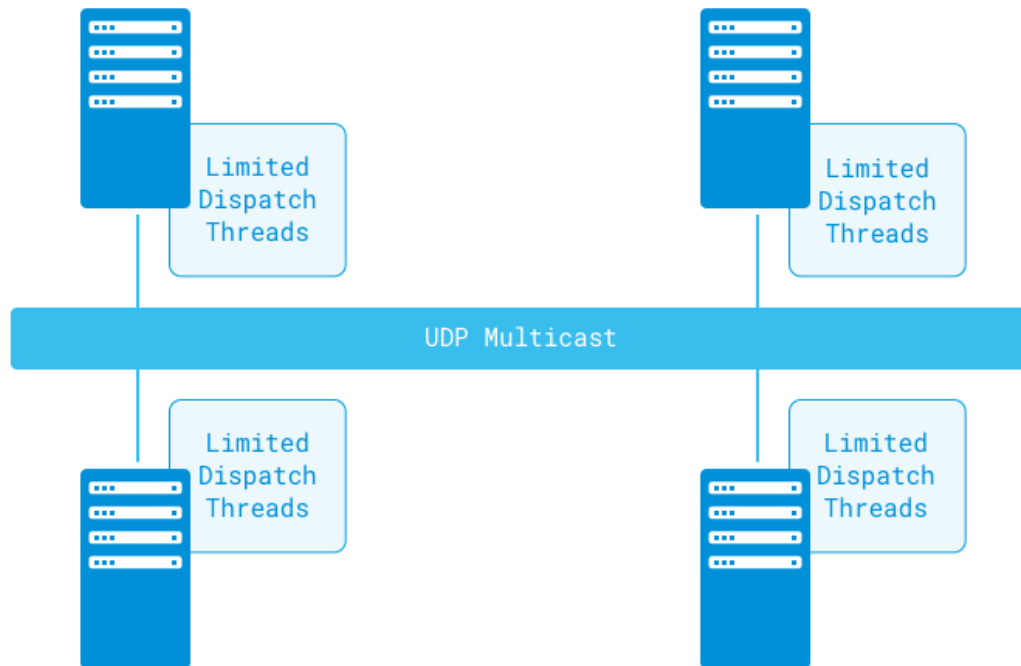


Figure 385.1: Liferay DXP's cache algorithm is extremely efficient.

address. If for some reason you need this configuration, you can make DXP auto detect the real address with this property:

```
cluster.link.autodetect.address=www.google.com:80
```

Set it to connect to some other host that's contactable by your server. By default, it points to Google, but this may not work if your server is behind a firewall. If you use each host's real address, you don't need to set the auto-detect address.

Cluster Link depends on JGroups and provides an API for nodes to communicate. It can

- Send messages to all nodes in a cluster
- Send messages to a specific node
- Invoke methods and retrieve values from all, some, or specific nodes
- Detect membership and notify when nodes join or leave

Cluster Link contains an enhanced algorithm that provides one-to-many type communication between the nodes. This is implemented by default with JGroups's UDP multicast, but unicast and TCP are also available.

## 385.2 Multicast Over UDP

---

When you enable Cluster Link, Liferay DXP's default clustering configuration is enabled. This configuration defines IP multicast over UDP. Liferay DXP uses two groups of channels from JGroups to implement this: a control group and a transport group. If you want to customize the channel properties, you can do so in `portal-ext.properties`:

```
cluster.link.channel.name.control=[your control channel name]
cluster.link.channel.properties.control=[your control channel properties]
```

Please see JGroups's documentation for channel properties. The default configuration sets many properties whose settings are discussed there.

Multicast broadcasts to all devices on the network. Clustered environments on the same network communicate with each other by default. Messages and information (e.g., scheduled tasks) sent between them can lead to unintended consequences. Isolate such cluster environments by either separating them logically or physically on the network, or by configuring each cluster's `portal-ext.properties` to use different sets of multicast group address and port values.

JGroups sets a bind address automatically. If you want to set a manual address, you can do this. By default, these are set to `localhost`:

```
cluster.link.bind.addr["cluster-link-control"]=localhost
cluster.link.bind.addr["cluster-link-udp"]=localhost
```

In some configurations, however, `localhost` is bound to the internal loopback network (127.0.0.1 or `::1`), rather than the host's real address. If for some reason you need this configuration, you can make Liferay DXP auto detect its real address with this property:

```
cluster.link.autodetect.address=www.google.com:80
```

Set it to connect to some other host that's contactable by your server. By default, it points to Google, but this may not work if your server is behind a firewall. If you set the address manually using the properties above, you don't need to set the auto-detect address.

Your network configuration may preclude the use of multicast over TCP, so below are some other ways you can get your cluster communicating. Note that these methods are all provided by JGroups.

## 385.3 Checkpoint:

---

1. If you are binding the IP address instead of using `localhost`, make sure the right IP addresses are declared using these properties:

```
cluster.link.bind.addr["cluster-link-control"]=localhost
cluster.link.bind.addr["cluster-link-udp"]=localhost
```

2. Test your load and then optimize your settings if necessary.

### 385.4 Unicast over TCP

---

If your network configuration or the geographical distance between nodes prevents you from using UDP Multicast clustering, you can configure TCP Unicast. You must use this if you have a firewall separating any of your nodes or if your nodes are in different geographical locations.

1. Add a parameter to your app server's JVM on each node:

```
-Djgroups.bind_addr=[node_ip_address]
```

Use the node's IP address.

2. Select a discovery protocol for the nodes to use to find each other. Here are the protocol choices:
  - TCPPing
  - JDBCPing
  - S3\_Ping
  - Rackspace\_Ping

If you aren't sure which one to choose, use TCPPing. The rest of these steps use TCPPing

3. Extract the tcp.xml file from \$LIFERAY.HOME/osgi/marketplace/Liferay Foundation - Liferay Portal - Impl.lpkg/com.liferay.portal.cluster.multiple-[version].jar/lib/jgroups-[version].Final.jar/tcp.xml to a location accessible to DXP, such as a folder called jgroups in the DXP web application's WEB-INF/classes folder.

```
WEB-INF/classes/jgroups/tcp.xml
```

4. In the tcp.xml file, set the TCP bind port to an unused port on your node. Here's an example:

```
<TCP bind_port="7800"/>
```

5. In the tcp.xml file, make each node discoverable to TCPPing by specifying its IP address and an unused port on that node. Building off of the previous step, here's an example <TCPPing> element:

```
<TCP bind_port="7800"/>
<TCPPING async_discovery="true"
 initial_hosts="192.168.224.154[7800],192.168.224.155[7800]"
 port_range="0"/>
```

#### Regarding Initial Hosts:

- An alternative to specifying initial hosts in a TCP XML file is to specify them to your app server using a JVM argument like this: `-Djgroups.tcpping.initial_hosts=192.168.224.154[7800],192.168.224.155[7800]`
- Make sure the initial hosts value accounts for all your nodes. If `initial_hosts` is not specified in a TCP XML file or in a JVM argument, `localhost` is the initial host.

6. Copy your `tcp.xml` file to each node, making sure to set the TCP bind port to the node's bind port. On the node with IP address 192.168.224.155, for example, configure TCPPing like this:

```
<TCP bind_port="7800"/>
<TCPPING async_discovery="true"
 initial_hosts="192.168.224.154[7800],192.168.224.155[7800]"
 port_range="0"/>
```

7. Modify the Cluster Link properties in the node's `portal-ext.properties` file to enable Cluster Link and point to the TCP XML file for each Cluster Link channel:

```
cluster.link.enabled=true
cluster.link.channel.properties.control=/jgroups/tcp.xml
cluster.link.channel.properties.transport.0=/jgroups/tcp.xml
```

The JGroups configuration demonstrated above is typically all that Unicast over TCP requires. However, in a very specific case, if (*and only if*) cluster nodes are deployed across multiple networks, then the parameter `external_addr` must be set on each host to the external (public IP) address of the firewall. This kind of configuration is usually only necessary when nodes are geographically separated. By setting this, clustered nodes deployed to separate networks (e.g. separated by different firewalls) can communicate together. This configuration may be flagged in security audits of your system. See JGroups documentation for more information.

**Note:** The `singleton_name` TCP attribute was deprecated in JGroups v4.0.0 and has therefore been removed since Liferay DXP 7.2 SP1 and Liferay Portal CE GA2 which use JGroups v 4.1.1-Final.

You're now set up for Unicast over TCP clustering! Repeat either TCPPING process for each node you want to add to the cluster.

## 385.5 JDBC Ping

---

Rather than use TCP Ping to discover cluster members, you can use a central database accessible by all the nodes to help them find each other. Cluster members write their own and read the other members' addresses from this database. To enable this configuration, replace the TCPPING tag with the corresponding JDBC\_PING tag:

```
<JDBC_PING
 connection_url="[place the URL to your database here]"
 connection_username="[place your user name here]"
 connection_password="[place your password here]"
 connection_driver="[place your driver name here]"/>
```

For example JDBC connection values, please see Database Templates. For further information about JDBC Ping, please see the JGroups Documentation.

### 385.6 S3 Ping

---

Amazon S3 Ping can be used for servers running on Amazon's EC2 cloud service. Each node uploads a small file to an S3 bucket, and all the other nodes read the files from this bucket to discover the other nodes. When a node leaves, its file is deleted.

To configure S3 Ping, replace the TCPPING tag with the corresponding S3\_PING tag:

```
<S3_PING
 secret_access_key="[SECRETKEY]"
 access_key="[ACCESSKEY]"
 location="ControlBucket"/>
```

Supply your Amazon keys as values for the parameters above. For further information about S3 Ping, please see the JGroups Documentation.

### 385.7 Other Pings

---

JGroups supplies other means for cluster members to discover each other, including Rackspace Ping, BPing, File Ping, and others. Please see the JGroups Documentation for information about these discovery methods.

The control and transport channels can be configured to use different ports.

### 385.8 Using Different Control and Transport Channel Ports

---

Using separate control and transport channel ports lets you monitor control and transport traffic and helps you isolate information to diagnose problems. These steps use Unicast over TCPping to demonstrate the approach.

1. Add a parameter to your app server's JVM on each node:

```
-Djgroups.bind_addr=[node_ip_address]
```

2. Extract the tcp.xml file from \$LIFERAY.HOME/osgi/marketplace/Liferay Foundation - Liferay Portal - Impl.lpkg/com.liferay.portal.cluster.multiple-[version].jar/lib/jgroups-[version].Final.jar/tcp.xml to a location accessible to DXP, such as a folder called jgroups in the DXP web application's WEB-INF/classes folder.
3. Make a copy of the tcp.xml in the same location and rename both files, designating one for the control channel and the other for the transport channel. For example, you could use these file names:

- tcp-control.xml



- tcp-transport.xml

4. Modify the Cluster Link properties in the node's portal-ext.properties file to enable Cluster Link and point to the TCP XML file for each Cluster Link channel:

```
cluster.link.enabled=true
cluster.link.channel.properties.control=/jgroups/tcp-control.xml
cluster.link.channel.properties.transport.0=/jgroups/tcp-transport.xml
```

5. Modify each tcp-\*.xml file's the TCP and TCPping elements to account for each node's IP address and bind port.

If you're vertically clustering (i.e., you have multiple servers running on the same physical or virtual system), every channel must use a unique unused bind port for discovery communication. In each tcp-\*.xml file, assign the TCP tag's bind\_port attribute to a unique, unused port.

For example, your first two nodes might assign these bind ports:

---

| Node   | Properties File     | Port   |
|--------|---------------------|--------|
| Node 1 | `tcp-control.xml`   | `7800` |
| Node 1 | `tcp-transport.xml` | `7801` |
| Node 2 | `tcp-control.xml`   | `7802` |
| Node 2 | `tcp-transport.xml` | `7803` |

---

Here are example TCP and TCPping elements using the bind ports on nodes running on the same system (i.e., same IP address):

```
Node 1 `tcp-control.xml`

```xml
<TCP bind_port="7800"/>
<TCPPING async_discovery="true"
  initial_hosts="192.168.224.154[7800],192.168.224.154[7802]"
  port_range="0"/>
...

**Node 1 `tcp-transport.xml`**

```xml
<TCP bind_port="7801"/>
<TCPPING async_discovery="true"
 initial_hosts="192.168.224.154[7801],192.168.224.154[7803]"
 port_range="0"/>
...

Node 2 `tcp-control.xml`

```xml
<TCP bind_port="7802"/>
<TCPPING async_discovery="true"
  initial_hosts="192.168.224.154[7800],192.168.224.154[7802]"
  port_range="0"/>
...

**Node 2 `tcp-transport.xml`**

```xml
```

```
<TCP bind_port="7803"/>
<TCPPING async_discovery="true"
 initial_hosts="192.168.224.154[7801],192.168.224.154[7803]"
 port_range="0"/>
...
```

If you have added entities that can be cached or you want to tune the cache configuration for your system, you can do so using a module.

### 385.9 Modifying the Cache Configuration with a Module

---

It's recommended to test your system under a load that best simulates the kind of traffic your system must handle. If you serve a lot of message board messages, your script should reflect that. If web content is the core of your site, your script should reflect that too.

As a result of a load test, you may find that the default distributed cache settings aren't optimized for your site. In this case, tweak the settings using a module. Follow instructions for [Overriding Cache](#).

You can install the module on each node and change the settings without taking down the cluster. This is a great benefit, but beware: since Ehcache doesn't allow for changes to cache settings while the cache is alive, reconfiguring a cache while the server is running flushes the cache.

### 385.10 Conclusion

---

Once you've configured your cluster, you can start it. A log file message shows your cluster's name (e.g., `cluster=liferay-channel-control`):

```
GMS: address=oz-52865, cluster=liferay-channel-control, physical address=192.168.1.10:50643
```

Congratulations! Your cluster is using Cluster Link.

## AUTO DEPLOY TO ALL NODES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

All modules and WAR files you deploy onto the cluster must be deployed to all cluster nodes. Because Liferay DXP installs applications as OSGi bundles, you cannot rely on your application server's means of installing WAR files (even if you only intend to install WAR files) to deploy an application to the entire cluster. Instead, place the application in each node's auto deploy folder (e.g., [Liferay Home]/deploy).

This, as you might imagine, can be done with a script. Write a shell script that uploads applications to each node using sftp or some other service. This way, when you deploy an application, it uploads to each node's auto deploy folder and installs to Liferay DXP on each node.



## UPDATING A CLUSTER

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Maintaining a cluster is a big responsibility. It includes deploying new and updated plugins and modules, applying fix packs, making configuration changes, and more. Maximizing server uptime and minimizing risks take priority when applying changes. Liferay DXP supports using standard cluster maintenance techniques.

- Rolling restarts: Nodes are shut down and updated one at a time.
- Blue-green deployment: Blue-green involves duplicating the current environment (*blue* environment), updating the duplicate (*green* environment), and cutting over users to the updated environment (green).

The techniques are compared below.

### Cluster Update Techniques

| Update                                                               | Rolling Restart | Blue-green |
|----------------------------------------------------------------------|-----------------|------------|
| Plugin/module installation                                           | Supported       | Supported  |
| Plugin/module update (backward-compatible data/schema changes)       | Supported       | Supported  |
| Plugin/module update (non-backward-compatible data/schema changes) 1 | Not supported   | Supported  |

| Update                                                  | Rolling Restart | Blue-green |
|---------------------------------------------------------|-----------------|------------|
| Fix pack installation and removal (revertable fix pack) | Supported       | Supported  |
| Fix pack installation (non-revertible fix pack)         | Not supported   | Supported  |
| Cluster code changes 2                                  | Not supported   | Supported  |
| Portal property changes                                 | Supported       | Supported  |
| System Setting changes via configuration admin files    | Supported       | Supported  |
| Application server updates                              | Supported       | Supported  |
| JVM setting changes                                     | Supported       | Supported  |
| New Java version (minor)                                | Supported       | Supported  |
| New Java version (major)                                | Not supported   | Supported  |

[1] Data and data schema changes that are not backward-compatible include, but are not limited to these:

- Modifying data in existing columns
- Dropping columns
- Changing column types
- Changing data formats used in columns (such as changing from XML to JSON)
- Updating a Service Builder service module's data schema to a version outside of the module's required data schema range. A module's Liferay-Require-SchemaVersion (specified in its `bnd.bnd`) must match the module's schema version value in the `Release_` table. Installing a module with a new schema version updates the `Release_` table with that schema version and triggers a data upgrade process. If you install such a module on one node, the schema version

in the `Release_` table no longer matches the Liferay-Require-SchemaVersion of the modules on the other nodes, and the module's Service Builder services become unavailable until the module is installed on the other nodes. Such changes cannot be reverted: the database must be restored from a backup. These schema version changes must be applied while all nodes are shut down.

[2] Cluster communication must stay intact. For this reason, cluster code must not be updated in rolling restarts. The Customer Portal identifies DXP fix packs that contain such changes as non-revertible. Here are packages you must not change in rolling restarts:

- `com.liferay.portal.kernel.cluster`
- `com.liferay.portal.kernel.cluster.*`
- `com.liferay.portal.kernel.exception.NoSuchClusterGroupException`
- `com.liferay.portal.kernel.scheduler.multiple`
- `com.liferay.portal.kernel.scheduler.multiple.*`
- `com.liferay.portal.cache.multiple`
- `com.liferay.portal.cache.multiple.*`
- `com.liferay.portal.scheduler.multiple`
- `com.liferay.portal.scheduler.multiple.*`

Since eligible changes should be done with rolling restarts, it's explained first.





---

## ROLLING RESTARTS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The rolling restart cluster maintenance process involves shutting down and updating nodes one at a time (while the other nodes are running) until they're all updated. It maximizes uptime while you update your cluster. Rolling restarts can be used in container and image based environments.

---

**Note:** Rolling restart does not include concepts for blue-green (separate, but identical environments) architectures, as these concepts specifically address multi-cluster style developments.

---

Here are the rolling restart steps:

1. Shut down one cluster node (JVM instance).
2. Update/modify the deployment for that node (see the maintenance scenarios that follow).
3. Start the node.
4. Repeat these steps for all other cluster nodes.

Maintenance scenarios vary in how they behave in rolling restarts. For example, UI changes in a plugin update are only visible on the updated nodes. Users on nodes that haven't been updated don't see the UI changes. Maintenance scenarios might have specific cases that cannot be performed in rolling restarts—the scenario descriptions mention these cases.

The maintenance scenarios eligible for rolling restart are described below.

---

### 388.1 New Modules and Plugins

---

For a new plugin or module (one that does not already exist in the cluster) to be eligible for rolling restart it must not modify data, or delete or rename database columns in a way that breaks compatibility with existing plugins or modules.

### **388.2 Updating Existing Modules and Plugins**

---

For a new version of an existing plugin or module to be eligible for rolling restart, it must not modify data or delete or rename database columns in a way that breaks compatibility with the existing version of the plugin or module.

### **388.3 Applying Fix Packs (DXP only)**

---

The Customer Portal identifies fix packs that are not revertible, and therefore ineligible for rolling restart. All other fix packs are eligible.

### **388.4 Reverting Fix Packs (DXP only)**

---

Revertible fix packs can be removed in rolling restarts.

### **388.5 Portal Properties controlled by portal-ext.properties**

---

Portal Properties file changes can be applied in rolling restarts.

### **388.6 System Settings controlled by Configuration Admin Files**

---

System configuration files can be applied in rolling restarts.

### **388.7 Application Server or JVM setting modifications**

---

Modifications to application server and JVM settings can be done in rolling restarts.

**388.8 Java Version Updates**

---

Minor version updates of Java can be applied in rolling restarts. Major version updates are not supported in rolling restarts, and should instead be done when all cluster nodes are shut down.

All rolling restart eligible updates can be applied using the rolling restart steps listed earlier. Other updates must be done differently as described next.

**388.9 Related Topics**

---

Liferay DXP Clustering

Maintaining Liferay DXP



---

## BLUE-GREEN DEPLOYMENT

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Blue-green is a deployment technique in which you duplicate your production environment (the *blue* environment) and modify the duplicate (the *green* environment) with software and data changes. When you've successfully tested the green environment, you cut users over from the blue environment to the green environment. Blue-green eliminates system down time.

Data schema and data changes require special attention. Custom plugin/module data schema changes that break compatibility with existing code must be introduced over several releases in which the data is transitioned and maintained in old and new columns until the old columns are unnecessary.

Data and schema changes require these steps:

1. Create a new column.
2. Copy the data to the new column.
3. Maintain both columns until the old column is no longer used by any cluster nodes.
4. Delete the column in the next release.

For more information, refer to these blue-green deployment articles:

- [BlueGreenDeployment](#)
- [Implementing Blue-Green Deployments with AWS](#)

---

### 389.1 Related Topics

---

Rolling Restarts

Liferay DXP Clustering

Maintaining Liferay DXP



---

## CONFIGURING REMOTE STAGING IN A CLUSTERED ENVIRONMENT

---

If you're running Liferay DXP as a clustered environment and you want to use remote staging, you must configure it properly for a seamless experience. In this tutorial, you'll learn how to set up remote staging in an example clustered environment scenario. The example environment assumes you have

- a Staging instance with database configurations and a file repository different from the cluster nodes.
- a balancer responsible for managing the traffic flow between the cluster's nodes.
- two nodes that call two Liferay app servers (e.g., *App Server 1* and *App Server 2*), both of which are connected to the same database.

The steps below also assume your web tier, application tier, and cluster environment are already configured. You may need to adjust the configurations in this tutorial to work with your specific configuration.

Let's begin!

1. You must secure the communication made between your nodes and Staging server. Add the following property to both app servers' and Staging server's `portal-ext.properties` file:

```
tunneling.servlet.shared.secret=[secret]
```

This secret key denies other portals access to your configured portal servers. If you'd like to set your secret key using hexadecimal encoding, also set the following property in your `portal-ext.properties` files:

```
tunneling.servlet.shared.secret.hex=true
```

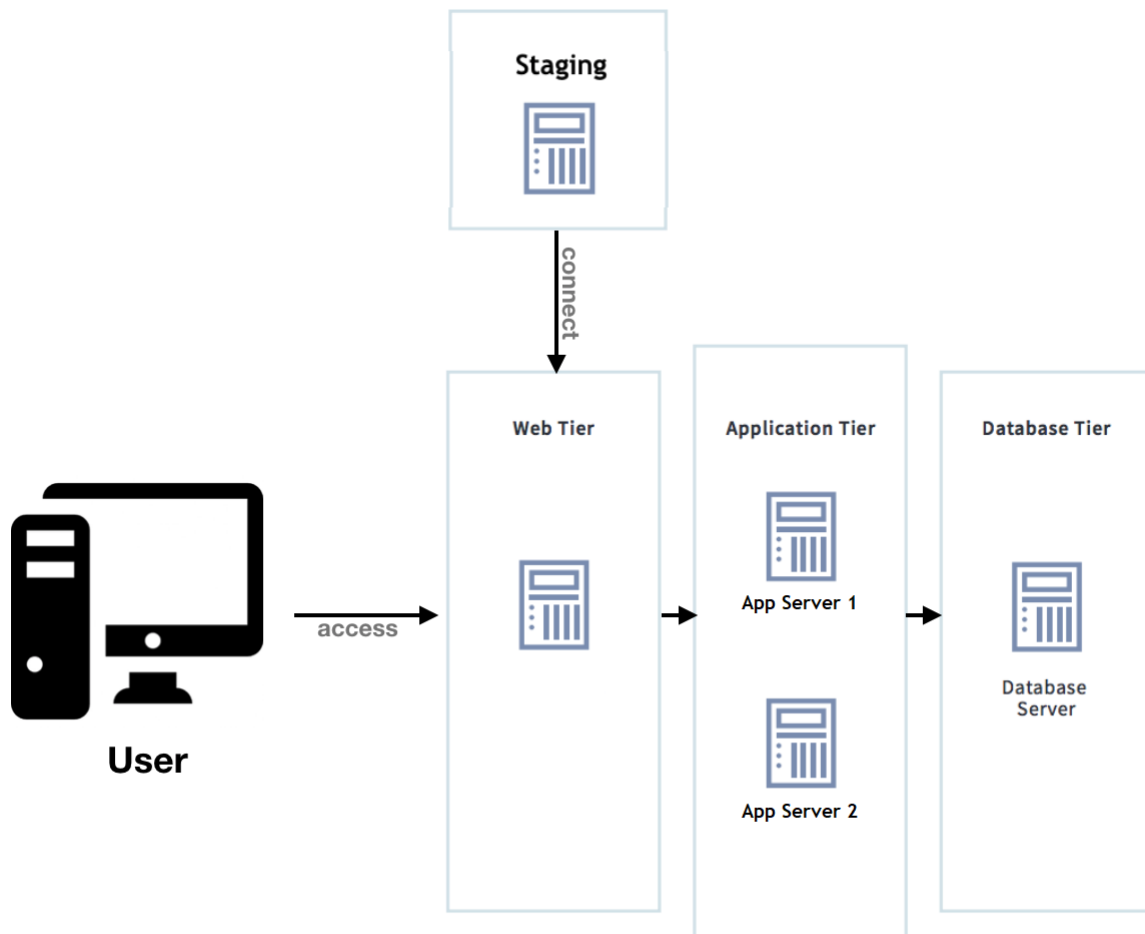


Figure 390.1: This is the assumed setup for your clustered environment.

**\*\*Note:\*\*** The following key lengths are supported by the available encryption algorithms:

- \*AES:\* 128, 192, and 256-bit keys
- \*Blowfish:\* 32-448 bit keys
- \*DESede (Triple DES):\* 56, 112, or 168 bit keys (Liferay places an artificial limit on the minimum key length and does not support the 56-bit key length)

For example, you can use [OpenSSL](https://www.openssl.org/) to generate a 128-bit AES key:

```
openssl enc -aes-128-cbc -k abc123 -P -md sha1
```

2. You must allow the connection between the configured IPs of your app servers and the Staging server. Open your remote Liferay server's `portal-ext.properties` file and add the following properties:

```
tunnel.servlet.hosts.allowed=127.0.0.1,SERVER_IP,[STAGING_IP]
tunnel.servlet.https.required=false
```



The [STAGING\_IP] variable must be replaced by the staging server's IP addresses. The SERVER\_IP constant can remain set for this property; it's automatically replaced by the Liferay server's IP addresses.

3. If you're validating IPv6 addresses, you must configure the app server's JVM to not force the usage of IPv4 addresses. For example, if you're using Tomcat, add the following attribute in the \$TOMCAT\_HOME\bin\setenv.[bat|sh] file.

```
`-Djava.net.preferIPv4Stack=false`
```

4. Restart both app servers for the new properties to take effect.
5. Configure the *TunnelAuthVerifier* property for your nodes' app servers. There are two ways to do this:

- **Use a .config file (recommended):** In the \$LIFERAY\_HOME/osgi/configs folder of one of your node Liferay DXP instances, create (if necessary) a `com.liferay.portal.security.auth.verifier.tunnel.default.config` file and insert the properties listed below. Creating one .config file configures all cluster nodes the same way. For more information on .config files, see the Understanding System Configuration Files article.

```
enabled=true
hostsAllowed=127.0.0.1,SERVER_IP,STAGING_IP
serviceAccessPolicyName=SYSTEM_USER_PASSWORD
urlsIncludes=/api/liferay/do
```

- **Via System Settings:** Navigate to the *Control Panel* → *Configuration* → *System Settings* → *Foundation* → *Tunnel Auth Verifiers*. Click on the `/api/liferay/do` configuration entry and add the Staging IP address to the *Hosts allowed* field. If you choose to configure the *TunnelAuthVerifier* this way, you **must** do this for all nodes (e.g., App Server 1 and App Server 2).

6. On your Staging instance, navigate to the Site Administration portion of the Product Menu and select *Publishing* → *Staging*. Then select *Remote Live*.
7. For the Remote Host/IP field, insert the balancer's IP of your web tier. Configuring the Staging instance with the balancer's IP ensures the availability of the environment at the time of publication from staging to live.
8. Enter the port on which the balancer is running into the Remote Port field.
9. Insert the remote site ID of your app servers into the Remote Site ID field. The site ID of all your app servers are the same since they are configured for the same database and are shared between nodes.  
 Navigate to the Site Administration portion of the Product Menu and select *Configuration* → *Site Settings* to find the site ID.
10. Save the Remote Live settings.

That's it! You've configured remote staging in your clustered environment.

## Staging

None

Local Live ?

Remote Live ?

Remote Live Connection Settings ▼

To publish changes to a *Remote Host/IP*, remote access between the servers must be configured. For detailed instructions on configuring this, refer to Liferay's official documentation. The *Remote Path Context* is only required if a non-root portal servlet context path is used on the target server. Access to this context must not be blocked by a proxy or firewall. Also, if the target server is a cluster, it is safe to set the *Remote Host/IP* to the cluster's load balanced address to increase the high availability of the publishing process.

Remote Host/IP

Remote Port

Remote Path Context

Remote Site ID

Use a Secure Network Connection  NO

Figure 390.2: When selecting the Remote Staging radio button, you're given a list of options to configure.

---

# CONTENT DELIVERY NETWORK

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

A Content Delivery Network (CDN) is an network of servers deployed in multiple data centers that contain your static content. When users hit your site, that static content is loaded from a server with geographical proximity to the user, speeding up requests.

Here, you'll first discover the perks of using a CDN and learn about general guidelines for using a CDN with Liferay DXP. Then, you'll configure a CDN. It's time to distribute your content around the world!

---

## **391.1 Using CDN for Performance Enhancements**

---

A CDN serves static web resources to users. These resources (images, CSS files, JavaScript files, etc.) are stored on multiple servers around the world. When requested, the resources are retrieved from the server nearest to the user.

The CDN functions as a caching proxy. This means that once static content is copied to a local server, it is stored in a cache for quick and easy retrieval. This drastically improves latency time, because browsers can download static resources from a local server down the street instead of halfway around the world. A user's request to the CDN for content is directed to specific server machine based on an algorithm that finds the server closest to the user. The figure below shows a visual representation of using geographical proximity to improve latency.

Because of the reduced wait time for requests and reduced load on your application server, a CDN is a great option to improve performance. Using a CDN with Liferay DXP, however, has some restrictions.

---

## **391.2 Liferay CDN Requirements**

---

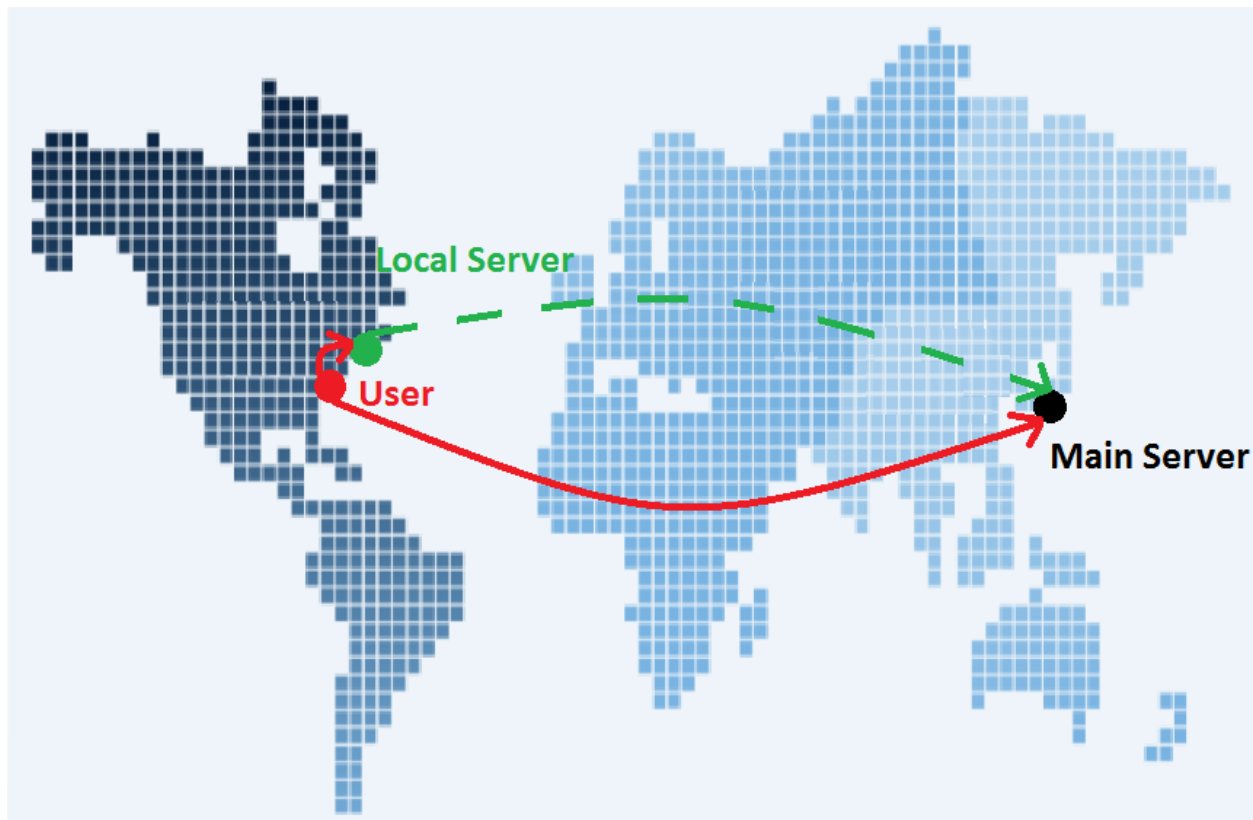


Figure 391.1: The red lines on the map represent the required distances traveled by requests from a server to the user. Using CDN allows a user to request static resources from a much closer local server, improving download times.

Liferay DXP only works with CDNs that can dynamically retrieve requested resources. Dynamic resources change over time or via interaction with end users and thus cannot be cached. For this reason, check with your CDN provider to make sure you don't have to upload anything manually in order for the CDN to work. The CDN must automatically fetch the content.

The CDN must work like a transparent proxy. A request first goes to the CDN. If the CDN doesn't have the requested resource, the CDN makes an identical request back to the origin (Liferay DXP), caches the resource, then serves the resource.

Once you're using a CDN (see below), it serves both portal resources and plugin resources (e.g., theme resources or JavaScript files referenced from a plugin's `liferay-portlet.xml` file). The CDN only serves resources that are included in a plugin. It does not serve resources that are dynamically loaded from external sources.

To get the CDN URL for a resource, developers should replace the portal host in the resource path with `themeDisplay.getCDNDynamicResourcesHost()`. Prefix resources with the CDN host name. Don't manually upload any resources to the CDN or put anything on the CDN which requires permission checking or complex policy access.

There are several portal properties for configuring your CDN to suit your needs. You'll learn how to do this next.

### 391.3 Configuring Liferay DXP to Use a CDN

Now that you understand what a CDN accomplishes and how it's used, it's time to set one up for yourself. You can set your CDN and its properties using two different methods:

1. By editing your portal properties file
2. By using the Control Panel

To configure your CDN via a properties file, create a `portal-ext.properties` file in your Liferay Home folder and set the appropriate Content Delivery Network properties.

Once you configure your CDN host, Liferay DXP generates URLs to the static assets that replace the old host with your new CDN host so they are automatically cached and served afterwards by the CDN.

To configure your CDN in the Control Panel, navigate to *Control Panel* → *Configuration* → *Instance Settings*. In the main configuration, there are three fields related to CDNs:

- *CDN Host HTTP*
- *CDN Host HTTPS*
- *Enable CDN Dynamic Resources*

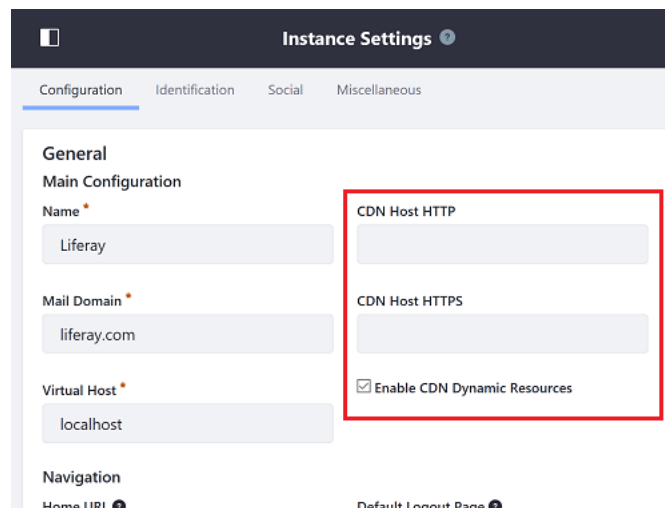


Figure 391.2: The Control Panel lets you configure your portal's CDN.

These properties are exactly the same as the ones you can specify in your `portal-ext.properties`. Make sure to visit the Content Delivery Network section of the `portal.properties` reference document if you don't know how to fill in the CDN fields. Once you're finished, click *Save* and your old host is replaced with your new CDN host for static content.

As you can see, configuring a CDN is easy and can drastically reduce latency time and improve performance.



---

# TUNING GUIDELINES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Although setting names may differ, these concepts apply to most application servers. To keep things simple, Tomcat is used as the example. For other application servers, consult the provider's documentation for specific settings.

Here are the tuning topics:

- Database Connection Pool
- Deactivating Development Settings in the JSP Engine
- Thread Pool

## 392.1 Database Connection Pool

---

The database connection pool should be roughly 30-40% of the thread pool size. It provides a connection whenever Liferay DXP needs to retrieve data from the database (e.g., user login). If the pool size is too small, requests queue in the server waiting for database connections. If the size is too large, however, idle database connections waste resources. As with thread pools, monitor these settings and adjust them based on your performance tests.

In Tomcat, the connection pools are configured in the Resource elements in `$CATALINA_HOME/conf/Catalina/local`. Liferay Engineering tests with this configuration:

```
<Resource auth="Container"
 description="Digital Enterprise DB Connection"
 driverClass="[place the driver class name here]"
 maxPoolSize="75"
 minPoolSize="10"
 acquireIncrement="5"
 name="jdbc/LiferayPool"
 user="[place your user name here]"
 password="[place your password here]"
 factory="org.apache.naming.factory.BeanFactory"
 type="com.mchange.v2.c3p0.ComboPooledDataSource"
 jdbcUrl="[place the URL to your database here]"/>
```

This configuration starts with 10 threads and increments by 5 as needed to a maximum of 75 connections in the pool.

There are a variety of database connection pool providers, including DBCP, C3P0, HikariCP, and Tomcat. You may also configure the Liferay JDBC settings in your `portal-ext.properties` file. For example JDBC connection values, please see Database Templates

## 392.2 Deactivating Development Settings in the JSP Engine

---

Many application servers' JSP Engines are in development mode by default. Deactivate these settings prior to entering production:

**Development mode:** This makes the JSP container poll the file system for changes to JSP files. Since you won't change JSPs on the fly like this in production, turn off this mode.

**Mapped File:** Generates static content with one print statement versus one statement per line of JSP text.

To disable these in Tomcat, for example, update the `$CATALINA_HOME/conf/web.xml` file's JSP servlet configuration to this:

```
<servlet>
 <servlet-name>jsp</servlet-name>
 <servlet-class>org.apache.jasper.servlet.JspServlet</servlet-class>
 <init-param>
 <param-name>development</param-name>
 <param-value>>false</param-value>
 </init-param>
 <init-param>
 <param-name>mappedFile</param-name>
 <param-value>>false</param-value>
 </init-param>
 <load-on-startup>3</load-on-startup>
</servlet>
```

Development mode and mapped files are disabled.

## 392.3 Thread Pool

---

Each request to the application server consumes a worker thread for the duration of the request. When no threads are available to process requests, the request is queued to wait for the next available worker thread. In a finely tuned system, the number of threads in the thread pool are balanced with the total number of concurrent requests. There should not be a significant amount of threads left idle to service requests.

Use an initial thread pool setting of 50 threads and then monitor it within your application server's monitoring consoles. You may wish to use a higher number (e.g., 250) if your average page times are in the 2-3 second range. Too few threads in the thread pool might queue excessive requests; too many threads can cause excessive context switching.



In Tomcat, the thread pools are configured in the `$CATALINA_HOME/conf/server.xml` file's Connector element. The Apache Tomcat documentation provides more details. Liferay Engineering tests with this configuration:

```
<Connector maxThreads="75" minSpareThreads="50"
 maxConnections="16384" port="8080"
 connectionTimeout="600000" redirectPort="8443"
 URIEncoding="UTF-8" socketBuffer="-1"
 maxKeepAliveRequests="-1" address="xxx.xxx.xxx.xxx"/>
```

Additional tuning parameters around Connectors are available, including the connector types, the connection timeouts, and TCP queue. Consult your application server's documentation for further details.



---

# JAVA VIRTUAL MACHINE TUNING

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Java Virtual Machine (JVM) tuning primarily focuses on adjusting the garbage collector and the Java memory heap. We used Oracle's 1.8 JVM for the reference architecture. You may choose other supported JVM versions and implementations. Please consult the Liferay DXP Compatibility Matrix for additional compatible JVMs.

Here are the JVM tuning topics:

- Garbage Collector
- Code Cache
- Java Heap
- JVM Advanced Options

Garbage collection is first.

---

## 393.1 Garbage Collector

---

Choosing the appropriate garbage collector (GC) helps improve the responsiveness of your Liferay DXP deployment. Use the concurrent low pause collectors:

```
-XX:+UseParNewGC -XX:ParallelGCThreads=16 -XX:+UseConcMarkSweepGC
-XX:+CMSParallelRemarkEnabled -XX:+CMSCompactWhenClearAllSoftRefs
-XX:CMSInitiatingOccupancyFraction=85 -XX:+CMSScavengeBeforeRemark
```

You may choose from other available GC algorithms, including parallel throughput collectors and G1 collectors. Start tuning using parallel collectors in the new generation and concurrent mark sweep (CMS) in the old generation.

**Note:** the `ParallelGCThreads` value (e.g., `ParallelGCThreads=16`) varies based on the type of CPUs available. Set the value according to CPU specification. On Linux machines, report the number of available CPUs by running `cat /proc/cpuinfo`.

**Note:** There are additional “new” algorithms like G1, but Liferay Engineering’s tests for G1 indicated that it does not improve performance. Since your application performance may vary, you should add G1 to your testing and tuning plans.

### 393.2 Code Cache

---

Java’s just-in-time (JIT) compiler generates native code to improve performance. The default size is 48m. This may not be sufficient for larger applications. Too small a code cache reduces performance, as the JIT isn’t able to optimize high frequency methods. For Liferay DXP, start with 64m for the initial code cache size.

```
-XX:InitialCodeCacheSize=64m -XX:ReservedCodeCacheSize=96m
```

Examine the efficacy of the parameter changes by adding the following logging parameters:

```
-XX:+PrintCodeCache -XX:+PrintCodeCacheOnCompilation
```

### 393.3 Java Heap

---

When most people think about tuning the Java memory heap, they think of setting the maximum and minimum memory of the heap. Unfortunately, most deployments require far more sophisticated heap tuning to obtain optimal performance, including tuning the young generation size, tenuring durations, survivor spaces, and many other JVM internals.

For most systems, it’s best to start with at least the following memory settings:

```
-server -XX:NewSize=700m -XX:MaxNewSize=700m -Xms2048m -Xmx2048m -XX:MetaspaceSize=512m
-XX:MaxMetaspaceSize=512m -XX:SurvivorRatio=6 -XX:TargetSurvivorRatio=9 -XX:MaxTenuringThreshold=15
```

On systems that require large heap sizes (e.g., above 4GB), it may be beneficial to use large page sizes. You can activate large page sizes like this:

```
-XX:+UseLargePages -XX:LargePageSizeInBytes=256m
```

You may choose to specify different page sizes based on your application profile.

**Note:** To use large pages in the JVM, you must configure your underlying operating system to activate them. In Linux, run `cat /proc/meminfo` and look at “huge page” items.

---

**Caution:** Avoid allocating more than 32GB to your JVM heap. Your heap size should be commensurate with the speed and quantity of available CPU resources.

---

### 393.4 JVM Advanced Options

---

The following advanced JVM options were also applied in the Liferay benchmark environment:

```
-XX:+UseLargePages -XX:LargePageSizeInBytes=256m
-XX:+UseCompressedOops -XX:+DisableExplicitGC -XX:-UseBiasedLocking
-XX:+BindGCThreadThreadsToCPUs -XX:UseFastAccessorMethods
```

Please consult your JVM documentation for additional details on advanced JVM options.  
Combining the above recommendations together, makes this configuration:

```
-server -XX:NewSize=1024m -XX:MaxNewSize=1024m -Xms4096m
-Xmx4096m -XX:MetaspaceSize=512m -XX:MaxMetaspaceSize=512m
-XX:SurvivorRatio=12 -XX:TargetSurvivorRatio=90
-XX:MaxTenuringThreshold=15 -XX:+UseLargePages
-XX:LargePageSizeInBytes=256m -XX:+UseParNewGC
-XX:ParallelGCThreads=16 -XX:+UseConcMarkSweepGC
-XX:+CMSParallelRemarkEnabled -XX:+CMSCompactWhenClearAllSoftRefs
-XX:CMSInitiatingOccupancyFraction=85 -XX:+CMSScavengeBeforeRemark
-XX:+UseLargePages -XX:LargePageSizeInBytes=256m
-XX:+UseCompressedOops -XX:+DisableExplicitGC -XX:-UseBiasedLocking
-XX:+BindGCThreadThreadsToCPUs -XX:+UseFastAccessorMethods
-XX:InitialCodeCacheSize=32m -XX:ReservedCodeCacheSize=96m
```

---

**Caution:** The above JVM settings should formulate a starting point for your performance tuning. Every system's final parameters vary due to many factors, including number of current users and transaction speed.

---

Monitor the garbage collector statistics to ensure your environment has sufficient allocations for metaspace and also for the survivor spaces. Using the configuration above in the wrong environment could result in dangerous runtime scenarios like out of memory failures. Improperly tuned survivor spaces also lead to wasted heap space.



---

## INSTALLING A SEARCH ENGINE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

A search engine is a critical component of your Liferay DXP installation. If you're here, you probably know the basics already and want to configure a search engine for your Liferay DXP deployment.

Elasticsearch, a highly scalable, full-text search engine, is installed by default, as an embedded server. Elasticsearch is well-supported and almost certainly meets any search and indexing need you have, but you must not use the embedded version in your production deployment.

Learn to configure a remote Elasticsearch server or cluster [here](#).

Solr is another capable and popular search engine supported in Liferay DXP.

Learn to configure a remote Solr server or cluster [here](#). But first, make sure you understand the disparity in functionality between the supported search engines.

---

### **394.1 Choosing a Search Engine**

---

Elasticsearch and Solr are both supported, but there are limitations to Liferay's Solr integration. To make use of some features, you must choose Elasticsearch.

---

### **394.2 End User Feature Limitations of Liferay's Solr Integration**

---

- Liferay Commerce
- Workflow Metrics
- Custom Filter search widget
- The Low Level Search Options widget
- Search Tuning: Customizing Search Results

- Search Tuning: Synonyms

### 394.3 Developer Feature Limitations of Liferay's Solr Integration

---

Implementation for the following APIs may be added in the future, but they are not currently supported by Liferay's Solr connector.

- From Portal Core (Module: `portal-kernel`, Artifact: `com.liferay.portal.kernel`):
  - `com.liferay.portal.kernel.search.generic.NestedQuery`
  - `com.liferay.portal.kernel.search.filter`:
    - \* `ComplexQueryPart`
    - \* `GeoBoundingBoxFilter`
    - \* `GeoDistanceFilter`
    - \* `GeoDistanceRangeFilter`
    - \* `GeoPolygonFilter`
- From the Portal Search API (Module: `portal-search-api`, Artifact: `com.liferay.portal.search.api`):
  - `com.liferay.portal.search.filter`:
    - \* `ComplexQueryPart`
    - \* `TermsSetFilter`
  - `com.liferay.portal.search.geolocation.*`
  - `com.liferay.portal.search.highlight.*`
  - `com.liferay.portal.search.query.function.*`
  - `com.liferay.portal.search.query.*`:
  - `com.liferay.portal.search.script.*`
  - `com.liferay.portal.search.significance.*`
  - `com.liferay.portal.search.sort.*`: only `Sort`, `FieldSort`, and `ScoreSort` are supported
- Portal Search Engine Adapter API (Module: `portal-search-engine-adapter-api`, Artifact: `com.liferay.portal.search.engine.adapter.api`)
  - `com.liferay.portal.search.engine.adapter.cluster.*`
  - `com.liferay.portal.search.engine.adapter.document.UpdateByQueryDocumentRequest`
  - `com.liferay.portal.search.engine.adapter.index.*`: only `RefreshIndexRequest` is supported
  - `com.liferay.portal.search.engine.adapter.search.*`:
    - \* `MultisearchSearchRequest`
    - \* `SuggestSearchRequest`
  - `com.liferay.portal.search.engine.adapter.snapshot.*`

Liferay Commerce requires the `TermsSetFilter` implementation, only available in the Elastic-search connector.



### **394.4 Elasticsearch Java Distribution Compatibility**

---

Another factor to consider in your search engine selection is JDK version. The search engine and Liferay DXP must use the same Java version and distribution (e.g., Oracle Open JRE 1.8.0\_201). Consult the Elasticsearch compatibility matrix and the Liferay DXP compatibility matrix to learn more about supported JDK distributions and versions. This consideration is not necessary for Solr, because no JVM level serialization happens between the servers. All communication occurs at the HTTP level.



## ELASTICSEARCH

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Elasticsearch is an open source, highly scalable, full-text search and analytics engine.

By default, Elasticsearch runs as an embedded search engine, which is useful for development and testing but is not supported in production. In production environments you must run Elasticsearch in remote mode, as a separate server or cluster. This guide walks you through the process of configuring Elasticsearch in remote mode.

Search Engine Vendor: **Elasticsearch(Embedded)**, Client Version: **6.5.0**, Nodes: **01BT8H4 (6.5.0)**

Figure 395.1: To see information about the currently connected search engine, go to *Control Panel* → *Configuration* → *Search*.

**Note:** Although Elasticsearch 6.5 is shipped as the embedded Elasticsearch server version, Elasticsearch 7 is the most recent supported Elasticsearch version for 7.0. Installing Elasticsearch 7 requires that you are running Service Pack 1/Fix Pack 2 or later (GA2 or later for CE users). Elasticsearch 6.8.x is also supported. See the compatibility matrix for exact versions.

If you'd rather use Solr, it's also supported. See the documentation on *Installing Solr* if you're interested.

To get up and running quickly with Elasticsearch as a remote server, refer to the *Installing Elasticsearch* article. Included there are basic instructions for installing and configuring Elasticsearch in a single server environment. Additional articles include more details and information on configuring and tuning Elasticsearch.

These terms are useful to understand as you read this guide:

- *Elasticsearch Home* refers to the root folder of your unzipped Elasticsearch installation (for example, `elasticsearch-7.4.1`).
- *Liferay Home* refers to the root folder of your Liferay DXP installation. It contains the `osgi`, `deploy`, `data`, and `license` folders, among others.

### 395.1 Embedded vs. Remote Operation Mode

---

When you start Liferay DXP, this message is displayed in the log:

```
2019-04-29 09:59:02.276 WARN [Elasticsearch initialization thread][EmbeddedElasticsearchConnection:288] Liferay is configured to use embedded Elast.
```

When you start Liferay DXP, Elasticsearch is already running in embedded mode. Liferay DXP runs an Elasticsearch node in the same JVM so it's easy to test-drive with minimal configuration. Running both servers in the same process has drawbacks:

- Elasticsearch must use the same JVM options as Liferay DXP.
- Liferay DXP and Elasticsearch compete for the same system resources.

---

**Note:** While it's not a supported production configuration, installing Kibana to monitor the embedded Elasticsearch server is useful during development and testing. Just be aware that you must install the OSS only Kibana build.

---

You wouldn't run an embedded database like HSQL in production, and you shouldn't run Elasticsearch in embedded mode in production either. Instead, run Elasticsearch in *remote operation mode*, as a standalone server or cluster of server nodes.

### 395.2 Troubleshooting Elasticsearch Integration

---

Sometimes things don't go as planned. If you've set up Liferay DXP with Elasticsearch in remote mode, but Liferay DXP can't connect to Elasticsearch, check these things:

**Cluster name:** The value of the `cluster.name` property in `elasticsearch.yml` must match the `clusterName` property configured in the Liferay DXP Elasticsearch connector.

**Transport address:** The value of the `transportAddresses` property in the Elasticsearch connector configuration must contain at least one valid host and port where an Elasticsearch node is running. If Liferay DXP is running in embedded mode, and you start a standalone Elasticsearch node or cluster, it detects that port 9300 is taken and switches to port 9301. If you then set Liferay's Elasticsearch connector to remote mode, it continues to look for Elasticsearch at the default port (9300).

The following articles cover the Liferay Connector to Elasticsearch's configuration options in more detail.

**Cluster Sniffing (Additional Configurations):** Elasticsearch clusters can have multiple node types. Cluster sniffing, enabled by default in the Liferay DXP connector, looks for data nodes configured in the `transportAddresses` property. If none are available, the connector may throw a `NoNodeAvailableException` in the console log. If cluster sniffing is to remain enabled, be sure that your configuration allows for at least one data node's transport address to be "sniffable" at all times to avoid this error.

To disable cluster sniffing, add `clientTransportSniff=false` to the `.config` file or un-check the Client Transport Sniff property in System Settings.



---

## PREPARING TO INSTALL ELASTICSEARCH

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

By default, 7.0 and its embedded Elasticsearch engine run in the same JVM. Although this enables out-of-the-box search, it's only supported for development. For production use, Elasticsearch must run in a separate JVM. See the installation guide for information on installing a remote Elasticsearch cluster.

Because search engines benefit heavily from caching, their JVM memory profiles differ substantially from those of a JVM running Liferay DXP. Therefore, the two applications should always be kept separate in production environments.

The following sections provide a synopsis of Elasticsearch configurations for 7.0. Prior to deployment, we strongly recommend reading Elastic's documentation on production deployment.

---

### 396.1 Sizing Your Deployment

---

When sizing your Elasticsearch deployment, carefully consider CPU, memory, disk, and network capacity. To scale effectively and avoid using lots of machines, deploy Elasticsearch on medium to large machines (for example, machines with two to eight CPUs). Avoid running multiple Elasticsearch JVMs on the same operating system.

---

### 396.2 CPU

---

We recommend allocating at least eight total CPU cores to the Elasticsearch engine, assuming only one Elasticsearch JVM is running on the machine.

### 396.3 Memory

---

At least 16 GB of memory is recommended, with 64 GB preferred. The precise memory allocation required depends on how much data is indexed. For index sizes 500 GB to 1 TB, 64 GB of memory suffices.

### 396.4 Disk

---

Search engines store their indexes on disk, so disk I/O capacity can impact search performance. Deploy Elasticsearch on SSD whenever possible. Otherwise use high-performance traditional hard disks (for example, 15k RPM). In either case, consider using RAID 0.

Avoid using Network Attached Storage (NAS) whenever possible as the network overhead can be large. If you're using public cloud infrastructure like Amazon Web Services, use instance local storage instead of network storage, such as Elastic Block Store (EBS).

Maintain 25 percent more disk capacity than the total size of your indexes. If your index is 60 GB, make sure you have at least 75 GB of disk space available. To estimate the disk space you need, you can index a representative sample of your production content and multiply that size by the fraction of your production content that it represents. For example, index 25 percent of your production content and then multiply the resulting index size by four. Keep in mind that indexing a 1 MB file doesn't result in 1 MB of disk space in the search index.

### 396.5 Cluster Size

---

While Liferay DXP can work with an Elasticsearch cluster comprised of one or two nodes, the minimum cluster size recommended by Elastic for fault tolerance is three nodes.

### 396.6 Networking

---

Elasticsearch relies on clustering and sharding to deliver fast, accurate search results, and thus requires a fast and reliable network. Most modern data centers provide 1 GbE or 10 GbE between machines.



---

## INSTALLING ELASTICSEARCH

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay DXP uses Elasticsearch to index its content. By default, it's installed as an embedded service. It works, but it's not a supported configuration for a production server. Feel free to use it while testing or developing, but when you're ready to put your site in production, you must run Elasticsearch as a standalone process. This is better anyway, because it frees you to design your infrastructure the way you want it. If you've got hardware or a VM to spare, you can separate your search infrastructure from Liferay DXP and reap some performance gains by putting search on a separate box. If you're more budget-conscious, you can still increase performance by running Elasticsearch in a separate, individually tunable JVM on the same box.

Before installing Elasticsearch, refer to [Preparing to Install Elasticsearch](#) for guidance on configuring the servers to support an Elasticsearch deployment properly.

Here's an overview of the installation steps:

1. Download a supported version of Elasticsearch. See [Elastic's website](#).
2. Install Elasticsearch by extracting its archive to the system where you want it to run.
3. Install some required Elasticsearch plugins.
4. Name your Elasticsearch cluster.
5. Configure Liferay DXP to connect to your Elasticsearch cluster.
6. Restart Liferay DXP and reindex your search and spell check indexes.

---

**Prerequisites:** Before continuing, make sure you have set the `JAVA_HOME` environment variable.

If you have multiple JDKs installed, make sure Elasticsearch and Liferay DXP are using the same version and distribution (e.g., Oracle Open JDK 1.8.0\_201). You can specify this in `[Elasticsearch Home]/bin/elasticsearch.in.sh`: `JAVA_HOME=/path/to/java`.

---

**Replacing the Default Elasticsearch 6 Connector:** If you're installing Elasticsearch 6, use the connector application installed by default. If you're installing Elasticsearch 7, you'll need to

download the connector from Liferay Marketplace for either CE and DXP. Always refer to the compatibility matrix to find the exact versions supported. Before installing the connector, blacklist the Elasticsearch 6 connector and APIs. The upgrade documentation holds detailed blacklisting steps.

---

When you perform these steps, you'll have a basic, production-ready instance of Liferay DXP and Elasticsearch up and running. But that's just the beginning of your server/connector configuration:

- Read about [Configuring Elasticsearch for Liferay DXP](#) in more detail.
- Learn how to [Secure Elasticsearch](#).
- [[Liferay Enterprise Search](#)] Learn how to configure [Monitoring](#).

For complete information on compatibility, check the [Liferay DXP compatibility matrix](#) and the [Liferay Enterprise Search compatibility matrix](#) if you have a subscription.

---

### 397.1 Step One: Download a Supported Version of Elasticsearch

---

If Liferay DXP isn't running, start it.

Visit port 9200 on localhost to access the embedded Elasticsearch:

```
http://localhost:9200
```

A JSON document is returned that looks similar to this:

```
{
 "name" : "01BT8H4",
 "cluster_name" : "LiferayElasticsearchCluster",
 "cluster_uuid" : "ziPGEBeSToGhc71VqaYHnA",
 "version" : {
 "number" : "6.5.0",
 "build_flavor" : "unknown",
 "build_type" : "unknown",
 "build_hash" : "816e6f6",
 "build_date" : "2018-11-09T18:58:36.352602Z",
 "build_snapshot" : false,
 "lucene_version" : "7.5.0",
 "minimum_wire_compatibility_version" : "5.6.0",
 "minimum_index_compatibility_version" : "5.0.0"
 },
 "tagline" : "You Know, for Search"
}
```

The version of Elasticsearch that's running is the value of the number field. In this example, it's 6.5.0. You can install the embedded version, but it might not be the most up-to-date version of Elasticsearch that's supported with Liferay DXP. Consult the [Compatibility Matrix](#) for definitive information on what's supported.

---

**Note:** Although the embedded server uses Elasticsearch 6.5, Elasticsearch 6.8.x has been tested with 7.0 GA1, and is fully supported. If you've upgraded to 7.0 Service Pack 1/Fix Pack 2 (or GA2 for CE users), Elasticsearch 7 is supported through the Liferay Connector to Elasticsearch 7, which can

be downloaded from Liferay Marketplace for both CE and DXP. Always refer to the compatibility matrix to find the exact versions supported.

---

Shut down the Liferay DXP server. In a local, single-machine testing environment, if you continue without shutting down, the Elasticsearch server you're about to install and start throws errors in the log if its cluster name and HTTP port match the already-running embedded Elasticsearch server. An alternative to shutting down Liferay DXP is to use a different cluster name (i.e., not `LiferayElasticsearchCluster`) and HTTP port (i.e., not 9200) in the remote Elasticsearch server.

When you know the version of Elasticsearch you need, go to Elastic's website and download that version.

---

### 397.2 Step Two: Install Elasticsearch

---

Most of this step entails deciding where you want to run Elasticsearch. Do you want to run it on the same machine as Liferay DXP, or do you want to run it on its own hardware? The answer to this question comes down to a combination of the resources you have available and the size of your installation. Regardless of what you decide, either way you get the benefit of a separately tunable search infrastructure.

Once you have a copy of the right version of Elasticsearch, extract it to a folder on the machine where you want it running. That's it!

---

### 397.3 Step Three: Install Elasticsearch Plugins

---

Install the following required Elasticsearch plugins:

- `analysis-icu`
- `analysis-kuromoji`
- `analysis-smartcn`
- `analysis-stempel`

To install these plugins, navigate to Elasticsearch Home and enter

```
./bin/elasticsearch-plugin install [plugin-name]
```

Replace *[plugin-name]* with the Elasticsearch plugin's name.

---

### 397.4 Step Four: Name Your Elasticsearch Cluster

---

A *cluster* in Elasticsearch is a collection of nodes (servers) identified as a cluster by a shared cluster name. The nodes work together to share data and workload. A one node cluster is discussed here; to create a multi-node cluster, please refer to Elastic’s documentation.

Now that you’ve installed Elastic, it sits in a folder on your machine, which is referred to here as [Elasticsearch Home]. To name your cluster, you’ll define the cluster name in both Elasticsearch and in Liferay DXP. First, define it in Elasticsearch. Edit the following file:

```
[Elasticsearch Home]/config/elasticsearch.yml
```

Uncomment the line that begins with `cluster.name`. Set the cluster name to whatever you want to name your cluster:

```
cluster.name: LiferayElasticsearchCluster
```

Of course, this isn’t a very imaginative name; you may choose to name your cluster `finders_keepers` or something else you can remember more easily. Save the file.

---

**Elasticsearch 6.x:** On Elasticsearch 6.x, you must also disable X-Pack Security unless you have a Liferay Enterprise Search subscription. Add this to `elasticsearch.yml`: `xpack.security.enabled: false`.

---

Now you can start Elasticsearch. Run the executable for your operating system from the [Elasticsearch Home]/bin folder:

```
./elasticsearch
```

Elasticsearch starts, and one of its status messages includes a transport address:

```
[2019-04-01T16:55:50,127][INFO][o.e.t.TransportService] [HfkqdKv] publish_address {127.0.0.1:9300}, bound_addresses {:::9300}, {127.0.0.1:9300}
```

Take note of this address; you’ll need to give it to your Liferay DXP server so it can find Elasticsearch on the network.

---

### 397.5 Step Five: Configure Liferay DXP to Connect to your Elasticsearch Cluster

---

Now that you’re ready to configure Liferay DXP, start it if you haven’t already, log in, and then click on *Control Panel* → *Configuration* → *System Settings* → *Search*. Enter the term *elasticsearch* in the search bar and click the *Elasticsearch [Version]* entry from the list of settings (at the time of writing, the version will either be 6 or 7). Now you can configure it. Here are the configuration options to change:

**Cluster Name:** Enter the name of the cluster as you defined it in Elasticsearch.

**Operation Mode:** Defaults to EMBEDDED. Change it to REMOTE to connect to a standalone Elasticsearch.

**Transport Addresses:** Enter a delimited list of transport addresses for Elasticsearch nodes. Here, you’ll enter the transport address from the Elasticsearch server you started. The default value is `localhost:9300`, which will work.

When finished, click *Save*. You’re almost done.

### 397.6 Step Six: Restart Liferay DXP and Reindex

If you're doing a local test installation, you probably only changed the Operation Mode in the connector configuration, so there's no need to restart; skip to re-indexing. If you've made more configuration changes in the connector's configuration, stop and restart Liferay DXP. When it's back up, log in as an administrative user and click on *Control Panel* → *Configuration* → *Search* and click the *Execute* button for *Reindex all search indexes* and then *Reindex all spell check indexes*. When you do that, you'll see some messages scroll up in the Elasticsearch log.

When restarting Liferay DXP, `update_mappings` messages will appear in the Elasticsearch logs:

```
[2019-04-01T17:08:57,462][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-0/m27eNsekTAyP27zD0jGojw] update_mapping [LiferayDocumentType]
[2019-04-01T17:08:57,474][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-0/m27eNsekTAyP27zD0jGojw] update_mapping [LiferayDocumentType]
[2019-04-01T17:08:58,393][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-0/m27eNsekTAyP27zD0jGojw] update_mapping [LiferayDocumentType]
[2019-04-01T17:08:58,597][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-0/m27eNsekTAyP27zD0jGojw] update_mapping [LiferayDocumentType]
[2019-04-01T17:09:07,040][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/0JidpYkgR50cCD5dgWB8Aw] update_mapping [LiferayDocumentType]
```

Once you reindex, more log messages appear in Elasticsearch:

```
[2019-04-01T17:11:17,338][INFO][o.e.c.m.MetaDataDeleteIndexService] [HfkqdKv] [liferay-20101/0JidpYkgR50cCD5dgWB8Aw] deleting index
[2019-04-01T17:11:17,389][INFO][o.e.c.m.MetaDataCreateIndexService] [HfkqdKv] [liferay-20101] creating index, cause [api], templates [], shards [1],
[2019-04-01T17:11:17,471][INFO][o.e.c.r.a.AllocationService] [HfkqdKv] Cluster health status changed from [YELLOW] to [GREEN] (reason: [shards starting
20101][0]) ...).
[2019-04-01T17:11:17,520][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:19,047][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:19,133][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:19,204][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:19,249][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:21,215][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:21,262][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:21,268][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:21,275][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:21,282][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
[2019-04-01T17:11:21,373][INFO][o.e.c.m.MetaDataMappingService] [HfkqdKv] [liferay-20101/Meacn_uXR06g0tCJonS4eA] update_mapping [LiferayDocumentType]
```

Reindexing the spell check dictionaries produces log messages like these:

```
2019-04-29 14:02:22.034 INFO [liferay/search_writer/SYSTEM_ENGINE-11][BaseSpellCheckIndexWriter:278] Start indexing dictionary for com.liferay/porta
2019-04-29 14:02:34.166 INFO [liferay/search_writer/SYSTEM_ENGINE-11][BaseSpellCheckIndexWriter:299] Finished indexing dictionary for com.liferay/po
2019-04-29 14:02:34.167 INFO [liferay/search_writer/SYSTEM_ENGINE-11][BaseSpellCheckIndexWriter:278] Start indexing dictionary for com.liferay/porta
2019-04-29 14:02:39.379 INFO [liferay/search_writer/SYSTEM_ENGINE-11][BaseSpellCheckIndexWriter:299] Finished indexing dictionary for com.liferay/po
```

For additional confirmation that Liferay DXP recognizes the remote search engine, navigate to the Search Control Panel application and note the subtle change there: the vendor name is now simply *Elasticsearch*, whereas prior to the installation of the remote Elasticsearch server, it said *Elasticsearch (Embedded)*.



Figure 397.1: To see information about the currently connected search engine, go to *Control Panel* → *Configuration* → *Search*.

For additional details refer to the Elasticsearch installation guide.



---

# CONFIGURING THE LIFERAY ELASTICSEARCH CONNECTOR

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

For detailed Elasticsearch configuration information, refer to the Elasticsearch documentation.

The name of your Elasticsearch cluster is important. When you're running Elasticsearch in remote mode, the cluster name is used by Liferay DXP to recognize the Elasticsearch cluster. To learn about setting the Elasticsearch cluster name on the Liferay DXP side, refer below to the section called Configuring the Liferay Elasticsearch Connector.

---

**Note:** The `http.enabled` setting in Elasticsearch corresponds to the `httpEnabled` setting in the Liferay Connector to Elasticsearch 6 application. As this setting was deprecated in Elasticsearch 6.3, the connector's corresponding setting is now also deprecated. This setting was only used for configuring the embedded Elasticsearch server, so its deprecation should have minimal impact to production deployments.

---

Elasticsearch's configuration files are written in YAML and kept in the `[Elasticsearch Home]/config` folder. The main configuration file is `elasticsearch.yml`, used for configuring Elasticsearch modules.

To set the name of the Elasticsearch cluster, open `[Elasticsearch Home]/config/elasticsearch.yml` and specify

```
cluster.name: LiferayElasticsearchCluster
```

Since `LiferayElasticsearchCluster` is the default name given to the cluster in the Liferay DXP Elasticsearch connector, this works just fine. Of course, you can name your cluster whatever you want (we humbly submit the recommendation `clustery_mcclusterface`).<sup>1</sup> Configure your node name using the same syntax (setting the `node.name` property). There's no client setting for this, it exists only in each Elasticsearch node's `elasticsearch.yml` file.

If you'd rather work from the command line than in the configuration file, navigate to `Elasticsearch Home` and enter

```
./bin/elasticsearch --cluster.name clustery_mcclusterface --node.name nody_mcnodeface
```

Feel free to change the node name or the cluster name. Once you configure Elasticsearch to your liking, start it up.

### 398.1 Starting Elasticsearch

---

Start Elasticsearch by navigating to Elasticsearch Home and typing

```
./bin/elasticsearch
```

if you run Linux, or

```
\bin\elasticsearch.bat
```

if you run Windows.

To run as a daemon in the background, add the `-d` switch to either command:

```
./bin/elasticsearch -d
```

Once both Elasticsearch and Liferay DXP are installed and running, introduce them to each other.

### 398.2 Configuring the Liferay Elasticsearch Connector

---

The Elasticsearch connector provides integration between Elasticsearch and the portal. Before configuring the connector, make sure Elasticsearch is running.

There are two ways to configure the connector:

1. Use the System Settings application in the Control Panel.
2. Manually create an OSGi configuration file.


It's convenient to configure the Elasticsearch connector from System Settings, but this is often only possible during development and testing. If you're not familiar with System Settings, read about it here. Remember that you can generate configuration files for deployment to other systems by configuring System Settings, and then exporting the `.config` file with your configuration.

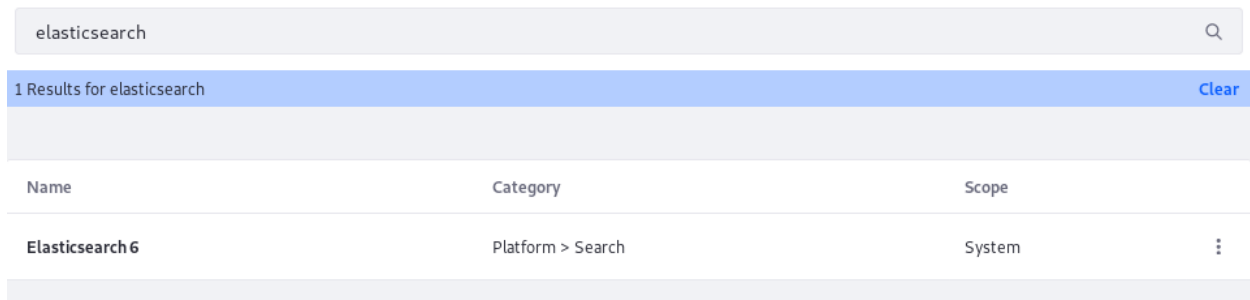
### 398.3 Configuring the Connector in the Control Panel

---

To configure the Elasticsearch connector from the System Settings application,



1. Start Liferay DXP.
2. Navigate to *Control Panel* → *Configuration* → *System Settings* → *Platform*.
3. Find the *Elasticsearch* entry (scroll down and browse to it or use the search box) and click the Actions icon (  ), then *Edit*.



The screenshot shows a search interface with a search bar containing 'elasticsearch'. Below the search bar, it indicates '1 Results for elasticsearch'. A table lists the search results with columns for Name, Category, and Scope. The result shown is 'Elasticsearch 6' under the category 'Platform > Search' and scope 'System'.

| Name            | Category          | Scope  |
|-----------------|-------------------|--------|
| Elasticsearch 6 | Platform > Search | System |

Figure 398.1: Use the System Settings application in Liferay DXP’s Control Panel to configure the Elasticsearch connector.

4. Make any edits to the configuration and click *Save*.

---

**Note:** If you switch operation modes (EMBEDDED → REMOTE), you must trigger a re-index. Navigate to *Control Panel* → *Configuration* → *Search*, and click *Execute* next to *Reindex all search indexes*.

---

### 398.4 Configuring the Connector with an OSGi .config File

When preparing a system for production deployment, you want to use a repeatable deployment process. Therefore, it’s best to use the OSGi configuration file, where your configuration is maintained in a controlled source.

Follow these steps to configure the Elasticsearch connector using a configuration file:

1. Create the following file:

```
[Liferay_Home]/osgi/configs/com.liferay.portal.search.elasticsearch7.configuration.ElasticsearchConfiguration.config
```

---

```
Elasticsearch 6: The name of the `.config` file for the Elasticsearch
6 connector is
`com.liferay.portal.search.elasticsearch6.configuration.ElasticsearchConfiguration.config`
```

---

2. Add configurations to the file, in the format `propertyName="Value"`. For example,

## Elasticsearch 6



### Cluster Name

LiferayElasticsearchCluster

The name of the cluster to join. The name should match the remote cluster when Operation Mode is set to remote.

### Operation Mode

REMOTE

Choose an Option

EMBEDDED

**REMOTE**

liferay-

The prefix for the search index name. This value should not be changed under normal conditions. If you change this value, you must also reindex the portal and then manually delete the old index using the Elasticsearch administration console.

### Number of Index Replicas

Figure 398.2: Configure the Elasticsearch connector's settings. Make sure you set the Operation Mode to *Remote*.

```
operationMode="REMOTE"
If running Elasticsearch from a different computer:
#transportAddresses="ip.of.elasticsearch.node:9300"
Highly recommended for all non-production usage (e.g., practice, tests, diagnostics):
#logExceptionsOnly="false"
```

### 3. Start Liferay DXP or re-index if already running.

As you can see from the System Settings entry for Elasticsearch, there are a lot more configuration options available that help you tune your system for optimal performance.

What follows here are some known good configurations for clustering Elasticsearch. These, however, can't replace the manual process of tuning, testing under load, and tuning again, so we encourage you to examine the Elasticsearch documentation and go through that process once you have a working configuration.

---

## 398.5 Configuring a Remote Elasticsearch Host

---

In production systems Elasticsearch and Liferay DXP are installed on different servers. To make Liferay DXP aware of the Elasticsearch cluster, set

```
transportAddresses=[IP address of Elasticsearch Node]:9300
```

Here's an example that sets the IP address of two nodes in the Elasticsearch cluster:

```
transportAddresses=["192.168.1.1:9300","192.168.1.2:9300"]
```

Set this in the Elasticsearch connector's OSGi configuration file. List as many or as few Elasticsearch nodes in this property as you want. This tells Liferay DXP the IP address or host name where search requests should be sent. If using System Settings, set the value in the *Transport Addresses* property.

---

**Note:** In an Elasticsearch cluster you can list the transport addresses for multiple Elasticsearch nodes as a comma-separated list in the `transportAddresses` property. If you set only one transport address, Liferay DXP loses contact with Elasticsearch if that node goes down.

---

On the Elasticsearch side, set the `network.host` property in your `elasticsearch.yml` file. This property simultaneously sets both the *bind host* (the host where Elasticsearch listens for requests) and the *publish host* (the host name or IP address Elasticsearch uses to communicate with other nodes). See here for more information.

---

## 398.6 Clustering Elasticsearch in Remote Operation Mode

---

To cluster Elasticsearch, first set `node.max_local_storage_nodes` to be something greater than 1. When you run the Elasticsearch start script, a new local storage node is added to the cluster. If you want four nodes running locally, for example, run `./bin/elasticsearch` four times. See here for more information.

Configure the number of shards and replicas in the Elasticsearch 6 connector, using the `indexNumberOfShards` and `indexNumberOfReplicas` properties to specify the number of primary shards and number of replica shards, respectively. Elasticsearch's default configuration works for a cluster of up to ten nodes, since the default number of shards is 5 and the default number of replica shards is 1.

---

**Note:** Elasticsearch uses the Zen Discovery Module by default, which provides unicast discovery. Additionally, nodes in the cluster communicate using the Transport Module, through TCP. See the Elasticsearch documentation for the available properties (to be set in the `elasticsearch.yml` file), and the Liferay DXP Elasticsearch connector's settings for the connector's available settings.

At a minimum, provide the list of hosts (as `host:port`) to act as gossip routers during unicast discovery in the `elasticsearch.yml`:

```
discovery.zen.ping.unicast.hosts: ["node1.ip.address", "node2.ip.address"]
```

For example,

```
discovery.zen.ping.unicast.hosts: ["10.10.10.5", "10.10.10.5:9305"]
```

For more information on configuring an Elasticsearch cluster, see the documentation on Elasticsearch Index Settings.

---

### 398.7 Elasticsearch Connector System Settings, By Operation Mode

---

Some of the settings available for the Elasticsearch connector are applicable for only one operation mode (REMOTE or EMBEDDED). Refer to the table below:

| Connector Setting/Operation Mode                 | EMBEDDED | REMOTE |
|--------------------------------------------------|----------|--------|
| <code>clusterName</code>                         | x        | x      |
| <code>operationMode</code>                       | x        |        |
| <code>indexNamePrefix</code>                     | x        | x      |
| <code>indexNumberOfReplicas*</code>              | x        | x      |
| <code>indexNumberOfShards*</code>                | x        | x      |
| <code>bootstrapMlockAll</code>                   | x        |        |
| <code>-logExceptionsOnly</code>                  | x        | x      |
| <code>retryOnConflict</code>                     | x        | x      |
| <code>discoveryZenPingUnicastHostsPort</code>    | x        |        |
| <code>-networkHost</code>                        | x        |        |
| <code>-networkBindHost</code>                    | x        |        |
| <code>-networkPublishHost</code>                 | x        |        |
| <code>-transportTcpPort</code>                   | x        |        |
| <code>-transportAddresses</code>                 | -        | x      |
| <code>clientTransportSniff</code>                | -        | x      |
| <code>clientTransportIgnoreClusterName</code>    | -        | x      |
| <code>clientTransportPingTimeout*</code>         | -        | x      |
| <code>clientTransportNodesSamplerInterval</code> | -        | x      |
| <code>httpEnabled</code>                         | x        |        |
| <code>httpCORSEnabled</code>                     | x        |        |
| <code>httpCORSAllowOrigin</code>                 | x        |        |
| <code>httpCORSConfigurations</code>              | x        |        |
| <code>additionalConfigurations</code>            | x        |        |
| <code>additionalIndexConfigurations</code>       | x        |        |
| <code>additionalTypeMappings</code>              | x        |        |
| <code>overrideTypeMappings</code>                | x        |        |

1 This is, of course, a nod to all those fans of Boaty Mcboatface.

---

# ADVANCED CONFIGURATION OF THE LIFERAY ELASTICSEARCH CONNECTOR

---

The default configuration for Liferay's Elasticsearch connector module is set in a Java class called `ElasticsearchConfiguration`.

While the Elasticsearch connector has a lot of configuration options out of the box, you might find an Elasticsearch configuration you need that isn't provided by default. In this case, add the configuration options you need. If something is configurable for Elasticsearch, it's configurable using the Elasticsearch connector.

## 399.1 Adding Settings and Mappings to the Liferay Elasticsearch Connector

---

Think of the available configuration options as being divided into two groups: the most common ones that are easily configured, and more complex configurations requiring a more brute-force approach: these include the `additionalConfigurations`, `additionalIndexConfigurations`, `additionalTypeMappings`, and `overrideTypeMappings` settings.

Figure 1: You can add Elasticsearch configurations to the ones currently available in System Settings.

## 399.2 Additional Configurations

---

The `additionalConfigurations` configuration defines extra settings (in YAML) for the embedded Elasticsearch. This is only useful for testing environments using the embedded Elasticsearch server. Any node settings normally set in `elasticsearch.yml` can be declared here. See the Elasticsearch documentation for a description of all possible node settings.

### 399.3 Adding Index Configurations

---

The `additionalIndexConfigurations` configuration defines extra settings (in JSON or YAML) that are applied to the Liferay DXP index when it's created. For example, you can create custom analyzers and filters using this setting. For a complete list of available settings, see the Elasticsearch reference.

Here's an example that shows how to configure analysis that can be applied to a field or a dynamic template (see below for an example application to a dynamic template).

```
{
 "analysis": {
 "analyzer": {
 "kuromoji_liferay_custom": {
 "filter": [
 "cjk_width",
 "kuromoji_baseform",
 "pos_filter"
],
 "tokenizer": "kuromoji_tokenizer"
 }
 },
 "filter": {
 "pos_filter": {
 "type": "kuromoji_part_of_speech"
 }
 }
 }
}
```

### 399.4 Adding Type Mappings

---

`additionalTypeMappings` defines extra mappings for the `LiferayDocumentType` type definition. These are applied when the index is created. Add the mappings using JSON syntax. For more information see [here](#) and [here](#). Use `additionalTypeMappings` for new field (properties) mappings and new dynamic templates, but don't try to override existing mappings. If any of the mappings set here overlap with existing mappings, index creation fails. Use `overrideTypeMappings` to replace default mappings.

As with dynamic templates, you can add sub-field mappings to Liferay DXP's type mapping. These are referred to as `properties` in Elasticsearch.

```
{
 "LiferayDocumentType": {
 "properties": {
 "fooName": {
 "index": "true",
 "store": "true",
 "type": "keyword"
 }
 }
 }
}
```

See [here](#) for more details on Elasticsearch's field datatypes.

The above example shows how a `fooName` field might be added to Liferay DXP's type mapping. Because `fooName` is not an existing property in the mapping, it works fine. If you try to override an existing property mapping, index creation fails. Instead use the `overrideTypeMappings` setting to override properties in the mapping.

To see that your additional mappings have been added to the `LiferayDocumentType`, use `curl` to access this URL after saving your additions and re-indexing:

```
curl http://[HOST]:[ES_PORT]/liferay-[COMPANY_ID]/_mapping/LiferayDocumentType?pretty
```

Here's what it would look like for an Elasticsearch instance running on `localhost:9200`, with a Liferay DXP Company ID of `20116`:

```
curl http://localhost:9200/liferay-20116/_mapping/LiferayDocumentType?pretty
```

In the above URL, `liferay-20116` is the index name. Including it indicates that you want to see the mappings that were used to create the index with that name.

### 399.5 Overriding Type Mappings

---

Use `overrideTypeMappings` to override Liferay DXP's default type mappings. This is an advanced feature that should be used only if strictly necessary. If you set this value, the default mappings used to define the Liferay Document Type in Liferay DXP source code (for example, `liferay-type-mappings.json`) are ignored entirely, so include the whole mappings definition in this property, not just the segment you're modifying. To make a modification, find the entire list of the current mappings being used to create the index by navigating to the URL

```
http://[HOST]:[ES_PORT]/liferay-[COMPANY_ID]/_mapping/LiferayDocumentType?pretty
```

Copy the contents in as the value of this property (either into System Settings or your OSGi configuration file). Leave the opening curly brace `{`, but delete lines 2-4 entirely:

```
"liferay-[COMPANY_ID]": {
 "mappings" : {
 "LiferayDocumentType" : {
```

Then, from the end of the mappings, delete the concluding three curly braces.

```
 }
 }
}
```

Now modify whatever mappings you'd like. The changes take effect once you save the changes and trigger a re-index from Server Administration.

Here's a partial example, showing a dynamic template that uses the analysis configuration from `additionalIndexConfigurations` to analyze all string fields that end with `_ja`. You'd include this with all the other default mappings, replacing the provided `template_ja` with this custom one:

```

{
 "LiferayDocumentType": {
 "dynamic_templates": [
 {
 "template_ja": {
 "mapping": {
 "analyzer": "kuromoji_liferay_custom",
 "index": "analyzed",
 "store": "true",
 "term_vector": "with_positions_offsets",
 "type": "string"
 },
 "match": "\\w+_ja\\b|\\w+_ja-[A-Z]{2}\\b",
 "match_mapping_type": "string",
 "match_pattern": "regex"
 }
 }
]
 }
}

```

### 399.6 Multi-line YAML Configurations

---

If you configure the settings from the last section using an OSGi configuration file, you might find yourself needing to write YAML snippets that span multiple lines. The syntax for that is straightforward and just requires appending each line with `\n`, like this:

```

additionalConfigurations=
 cluster.routing.allocation.disk.threshold_enabled: false\n
 cluster.service.slow_task_logging_threshold: 600s\n
 index.indexing.slowlog.threshold.index.warn: 600s\n
 index.search.slowlog.threshold.fetch.warn: 600s\n
 index.search.slowlog.threshold.query.warn: 600s\n
 monitor.jvm.gc.old.warn: 600s\n
 monitor.jvm.gc.young.warn: 600s

```

From simple configurations to overriding existing type mappings, Elasticsearch and Liferay's connector to Elasticsearch are configurable.



## ELASTICSEARCH CONNECTOR SETTINGS: REFERENCE

Elasticsearch is the default search engine for 7.0. The connection is managed through the *Liferay Connector to Elasticsearch [Version]*, and is configurable through System Settings or an OSGi configuration file named

```
com.liferay.portal.search.elasticsearch6.configuration.ElasticsearchConfiguration.config
```

If you are using Elasticsearch 7, your configuration file must be named

```
com.liferay.portal.search.elasticsearch7.configuration.ElasticsearchConfiguration.config
```

Deploy the file to [Liferay\_Home]/osgi/configs and a listener auto-detects it.

The list below is all the configuration settings for Liferay's default Elasticsearch connector, in the order they appear in the System Settings application (The *Elasticsearch [Version]* entry under the *Search* category):

**clusterName=LiferayElasticsearchCluster** A String value that sets the name of the cluster to integrate with. This name should match the remote cluster when Operation Mode is set to remote. (See also: remote operation mode)

**operationMode=EMBEDDED** There are two operation modes you can choose from: EMBEDDED or REMOTE. Set to REMOTE to connect to a remote standalone Elasticsearch cluster. Set to EMBEDDED to start Liferay with an internal Elasticsearch instance. Embedded operation mode is unsupported for production environments.

**indexNamePrefix=liferay-** Set a String value to use as the prefix for the search index name. The default value should not be changed under normal conditions. If you change it, you must also perform a *reindex all* operation for the portal and then manually delete the old index using the Elasticsearch administration console.

**indexNumberOfReplicas=** Set the number of replicas for each index. If left unset, no replicas are used. A full reindex is required to make changes take effect.

**indexNumberOfShards=** Set the number of index shards to use when a Liferay index is created. If left unset, a single shard is used. A full reindex is required to make changes take effect.

- bootstrapMlockAll=false** A boolean setting that, when set to true, tries to lock the process address space into RAM, preventing any Elasticsearch memory from being swapped out (see here) for more information)
- logExceptionsOnly=true** A boolean setting that, when set to true, only logs exceptions from Elasticsearch, and does not rethrow them.
- retryOnConflict=5** Set an int value for the number of retries to attempt if a version conflict occurs because the document was updated between getting it and updating it (see here for more information).
- discoveryZenPingUnicastHostsPort=9300-9400** Set a String value for the range of ports to use when building the value for `discovery.zen.ping.unicast.hosts`. Multiple Elasticsearch nodes on a range of ports can act as gossip routers at the same computer (see here for more information).
- networkHost=** Set this String value to instruct the node to bind to this hostname or IP address and publish (advertise) this host to other nodes in the cluster. This is a shortcut which sets the bind host and the publish host at the same time (see here for more information).
- networkBindHost=** Set the String value of the network interface(s) a node should bind to in order to listen for incoming requests (see here for more information).
- networkPublishHost=** Set the String value of a single interface that the node advertises to other nodes in the cluster, so that those nodes can connect to it (see here for more information).
- transportTcpPort=** Set the String value for the port to bind for communication between nodes. Accepts a single value or a range (see here for more information).
- transportAddresses=localhost:9300** Set the String values for the addresses of the remote Elasticsearch nodes to connect to. This value is required when Operation Mode is set to remote (see here for more information). Specify as many or few nodes as you see fit.
- clientTransportSniff=true** Set this boolean to true to enable cluster sniffing and dynamically discover available data nodes in the cluster (see here for more information).
- clientTransportIgnoreClusterName=false** Set this boolean to true to ignore cluster name validation of connected nodes (see here for more information).

**clientTransportPingTimeout=** The time (in seconds) the client node waits for a ping response from a node. If unset, the default Elasticsearch `client.transport.ping_timeout` is used.

- clientTransportNodesSamplerInterval=** Set this String value to instruct the client node on how often to sample / ping the nodes listed and connected (see here for more information).
- httpEnabled=true** Set this boolean to false to disable the http layer entirely on nodes which are not meant to serve REST requests directly. As this setting was deprecated in Elasticsearch 6.3, the connector's corresponding setting is now also deprecated. This setting was only used for configuring the embedded Elasticsearch server, so its deprecation should have minimal impact to production deployments.
- httpCORSEnabled=true** Set this boolean to false to disable cross-origin resource sharing, i.e. whether a browser on another origin can do requests to Elasticsearch. If disabled, web front end tools like `elasticsearch-head` may be unable to connect (see here for more information).
- httpCORSAllowOrigin=/https?:\\/\\*/localhost(:[0-9]+)?/** Set the String origins to allow when HTTP CORS is enabled (see here for more information).
- httpCORSConfigurations=** Set the String values for custom settings for HTTP CORS, in YML format (`elasticsearch.yml`) (see here for more information).
- additionalConfigurations=** Set the String values for custom settings for embedded Elasticsearch, in YML format. See: Adding Settings to the Liferay Elasticsearch Connector

**additionalIndexConfigurations**= Set the String values for custom settings for the Liferay index, in JSON or YML format (refer to the Elasticsearch Create Index API for more information). See: Adding Settings to the Liferay Elasticsearch Connector

**additionalTypeMappings**= Set the String values for custom mappings for the LiferayDocumentType, in JSON format (refer to the Elasticsearch Put Mapping API for more information) See: Adding Settings to the Liferay Elasticsearch Connector

**overrideTypeMappings**= Settings here override Liferay DXP's default type mappings. This is an advanced feature that should be used only if strictly necessary. If you set this value, the default mappings used to define the Liferay Document Type in Liferay DXP source code (for example, liferay-type-mappings.json) are ignored entirely, so include the whole mappings definition in this property, not just the segment you're modifying.

#### **400.1 Configurations only Affecting the Embedded Elasticsearch Server**

---

These settings (defined above) are only meant to use while configuring the embedded Elasticsearch server. Configuring these will elicit no effect on remote Elasticsearch installations:

- bootstrapMlockAll
- discoveryZenPingUnicastHostsPort
- networkHost
- networkBindHost
- networkPublishHost
- transportTcpPort
- httpEnabled
- httpCORSEnabled
- httpCORSAllowOrigin
- httpCORSConfigurations

You can easily configure these settings in the System Setting application, or as mentioned above, you can specify them in a deployable OSGi .config file.



# INSTALLING LIFERAY ENTERPRISE SEARCH SECURITY

---

This document has been updated and ported to Liferay Learn and is no longer maintained here. The very first thing you must do to secure Elasticsearch is enable X-Pack Security. After that you can begin configuring authentication and Transport Layer Security.

---

**Elasticsearch 6.x:** If you're using Elasticsearch 6, you'll need a Liferay Enterprise Search (LES) subscription to use X-Pack. Starting with the Liferay Connector to Elasticsearch 7 (available on Liferay Marketplace), X-Pack security is included by default. X-Pack monitoring still requires LES.

---

## 401.1 Enabling X-Pack Security

---

To enable security, add this setting in `elasticsearch.yml`:

```
xpack.security.enabled: true
```

Now you can set up X-Pack users.

## 401.2 Setting Up X-Pack Users

---

In a system using X-Pack Security and X-Pack Monitoring, these built-in X-Pack users are important:

- kibana
- elastic

Set the passwords for all X-Pack's built-in users. The `setup-passwords` command is the simplest method to set the built-in users' first-use passwords for the first time. To update a password subsequently, use Kibana's UI or the Change Password API.

The interactive argument sets the passwords for all built-in users. The configuration shown in these articles assumes you set all passwords to *liferay*. Of course, that's not recommended for production systems.

```
./bin/elasticsearch-setup-passwords interactive
```

Elastic's `setup-passwords` command documentation describes additional options. Since you're securing Elasticsearch, remember the elastic user's password. Enable transport layer security on each node.

### 401.3 Enabling Transport Layer Security

---

The following instructions for enabling TLS use *liferay* as the password whenever one is needed. Use your own passwords for your installation.

---

**Important:** Elasticsearch and Liferay DXP must share the keys and certificates used to configure TLS. Copy them between servers and point to the local copy in the corresponding configuration files.

---

### 401.4 Generate Node Certificates

---

Generate a node certificate for each node. Alternatively, use a Certificate Authority to obtain node certificates.

1. Create a certificate authority, using X-Pack's `certutil` command:

```
./bin/elasticsearch-certutil ca --pem --ca-dn CN=localhost
```

This generates a ZIP file. Unzip the contents in the `[Elasticsearch Home]/config/certs` folder.

2. Generate X.509 certificates and private keys using the CA from Step 1:

```
./bin/elasticsearch-certutil cert --pem --ca-cert /path/to/ca.crt --ca-key /path/to/ca.key --dns localhost --ip 127.0.0.1 --name localhost
```

This generates another ZIP file. Extract the contents in the `[Elasticsearch Home]/config/certs` folder.

---

**Note:** The `certutil` command defaults to using the `PKSC#12` format for certificate generation. Since Kibana does not work with `PKSC#12` certificates, the `--pem` option (generates the certificate in PEM format) is important if you're using X-Pack monitoring.

---

**Checkpoint:** You now have the following files in your `[Elasticsearch Home]/config/certs` folder:

```
ca.crt
ca.key
localhost.crt
localhost.key
```

---

## 401.5 Enable TLS for Elasticsearch 7

---

Enable TLS on each node via its `elasticsearch.yml`.

1. Enable transport layer TLS with these settings in `elasticsearch.yml` for inter-node communication:

```
xpack.security.transport.ssl.enabled: true
```

2. Add the certificate, key and certificate authority paths to each node's `elasticsearch.yml`:

```
xpack.security.transport.ssl.certificate: certs/localhost.key
xpack.security.transport.ssl.certificate_authorities: ["certs/ca.crt"]
xpack.security.transport.ssl.key: certs/localhost.crt
xpack.security.transport.ssl.verification_mode: certificate
```

The example paths above assume you added the certificate to `Elasticsearch Home/config/`.

3. Enable TLS on the HTTP layer to encrypt client communication:

```
xpack.security.http.ssl.enabled: true
```

4. Configure the certificate, key, and certificate authority paths to each node's `elasticsearch.yml`:

```
xpack.security.http.ssl.certificate_authorities: ["certs/ca.crt"]
xpack.security.http.ssl.certificate: certs/localhost.crt
xpack.security.http.ssl.key: certs/localhost.key
xpack.security.http.ssl.verification_mode: certificate
```

## 401.6 Elasticsearch 6 TLS

---

The settings on Elasticsearch 6 were slightly different than those presented above for Elasticsearch 7. Enable TLS on each node via its `elasticsearch.yml`.

1. Add the certificate, key and certificate authority paths to each node's `elasticsearch.yml`:

```
xpack.ssl.certificate: certs/localhost.crt
xpack.ssl.certificate_authorities: ["certs/ca.crt"]
xpack.ssl.key: certs/localhost.key
xpack.ssl.verification_mode: certificate
```

The example paths above assume you added the certificate to `Elasticsearch Home/config/`.

2. Enable transport layer TLS with these settings in `elasticsearch.yml`:

```
xpack.security.transport.ssl.enabled: true
```

3. Enable TLS on the HTTP layer to encrypt client communication:

```
xpack.security.http.ssl.enabled: true
```

After X-Pack is installed and TLS is enabled, configure the X-Pack Security adapter in Liferay DXP.

## 401.7 Example Elasticsearch Security Configuration

---

For ease of copying and pasting, here is the complete Elasticsearch configuration (`elasticsearch.yml`) used in this guide (with the Elasticsearch 6 example commented out):

```
For Elasticsearch 7.3/7.4
cluster.name: LiferayElasticsearchCluster

X-Pack Security
xpack.security.enabled: true

TLS/SSL settings for Transport layer
xpack.security.transport.ssl.enabled: true
xpack.security.transport.ssl.verification_mode: certificate
xpack.security.transport.ssl.key: certs/localhost.key
xpack.security.transport.ssl.certificate: certs/localhost.crt
xpack.security.transport.ssl.certificate_authorities : ["certs/ca.crt"]

TLS/SSL settings for HTTP layer
xpack.security.http.ssl.enabled: true
xpack.security.http.ssl.verification_mode: certificate
xpack.security.http.ssl.key: certs/localhost.key
xpack.security.http.ssl.certificate: certs/localhost.crt
xpack.security.http.ssl.certificate_authorities : ["certs/ca.crt"]

Comment out when Kibana and Liferay's X-Pack Monitoring are also configured
```



```

#xpack.monitoring.collection.enabled: true

For Elasticsearch 6.5/6.8
#cluster.name: LiferayElasticsearchCluster
#
X-Pack Security
#xpack.security.enabled: true
#
Enable TLS/SSL
#xpack.security.transport.ssl.enabled: true # To enable Transport level SSL for internode-communication
#xpack.security.http.ssl.enabled: true # To enable HTTP level SSL required by Kibana
#
General TLS/SSL settings for both Transport and HTTP levels
#xpack.ssl.verification_mode: certificate
#xpack.ssl.key: certs/localhost.key
#xpack.ssl.certificate: certs/localhost.crt
#xpack.ssl.certificate_authorities : ["certs/ca.crt"]
#
Comment out when Kibana and Liferay's X-Pack Monitoring are also configured
#xpack.monitoring.collection.enabled: true

```

For both Elasticsearch 6 and Elasticsearch 7, the Liferay Connector settings are the same.

## 401.8 Install and Configure the Liferay Enterprise Search Security app

---

If you have a Liferay Enterprise Search subscription, download the Liferay Enterprise Search Security app . Install the LPKG file by copying it into the Liferay Home/deploy folder.

To configure the X-Pack adapter, navigate to *Control Panel* → *Configuration* → *System Settings*. Find the *Search* category and click on the *X-Pack Security* entry. You can enter the property values here, but it's more common to use a configuration file deployed to [Liferay Home]/osgi/configs. For the X-Pack security connector, create a file called

```
com.liferay.portal.search.elasticsearch7.configuration.XPackSecurityConfiguration.config
```

The exact contents of the file depend on your X-Pack setup. To configure the adapter according to the Elasticsearch setup documented here, populate the file like this:

```

sslKeyPath="/path/to/localhost.key"
sslCertificatePath="/path/to/localhost.crt"
certificateFormat="PEM"
requiresAuthentication="true"
username="elastic"
password="liferay"
sslCertificateAuthoritiesPaths="/path/to/ca.crt"
transportSSLVerificationMode="certificate"
transportSSLEnabled="true"

```

The password should match what you set during the X-Pack password setup above.

The certificate and key files referenced here are the same ones used on the Elasticsearch server. Copy them to the Liferay DXP server and update their paths in the configuration accordingly.

Enable authentication by setting `requiresAuthentication` to `true` and providing the credentials for the Elasticsearch user. For TLS, enable transport TLS, set the certificate verification mode and certificate format, and provide the path to the certificate, key, and certificate authority. Of course, the exact values depend on your X-Pack configuration.

Here's the complete list of configuration options for the X-Pack Connector:

- `sslKeyPath`
- `sslCertificatePath`
- `sslCertificateAuthoritiesPaths`
- `certificateFormat`
- `requiresAuthentication`
- `username`
- `password`
- `transportSSLVerificationMode`
- `transportSSLEnabled`
- `sslKeystorePath`
- `sslKeystorePassword`
- `sslTruststorePath`
- `sslTruststorePassword`

When you're finished configuring X-Pack Security, restart Elasticsearch. These steps require a full cluster restart.

### 401.9 Disabling Elasticsearch Deprecation Logging

---

Some Elasticsearch APIs used by Liferay's Elasticsearch 6 connector were deprecated as of Elasticsearch 6.6 and 6.7. This can result WARN log entries in Elasticsearch's deprecation log when Liferay DXP is configured with Elasticsearch 6.8.x and X-Pack Security is enabled:

```
2019-07-16T14:47:05,779][WARN][o.e.d.c.j.Joda] [
ode_name]'y' year should be replaced with 'u'. Use 'y' for year-of-era. Prefix your date format with '8' to use the new specifier.
[2019-07-16T14:47:06,007][WARN][o.e.d.c.s.Settings] [
ode_name][xpack.ssl.certificate] setting was deprecated in Elasticsearch and will be removed in a future release! See the breaking changes documentat
[2019-07-16T14:47:06,007][WARN][o.e.d.c.s.Settings] [
ode_name][xpack.ssl.certificate_authorities] setting was deprecated in Elasticsearch and will be removed in a future release! See the breaking chang
[2019-07-16T14:47:06,008][WARN][o.e.d.c.s.Settings] [
ode_name][xpack.ssl.key] setting was deprecated in Elasticsearch and will be removed in a future release! See the breaking changes documentation for
[2019-07-16T14:47:06,463][WARN][o.e.d.x.c.s.SSLService] [
ode_name]SSL configuration [xpack.http.ssl] relies upon fallback to another configuration for [key configuration, trust configuration], which is dep
[2019-07-16T14:47:06,464][WARN][o.e.d.x.c.s.SSLService] [
ode_name]SSL configuration [xpack.security.transport.ssl.] relies upon fallback to another configuration for [key configuration, trust configuration]
```

These warnings do not signal any functional issues, and can be disabled (see here to learn how).

---

# BACKING UP ELASTICSEARCH

---

Elasticsearch replicas protect against a node going down, but they won't help you with a catastrophic failure. Only good backup practices can help you then.

Back up and restore your Elasticsearch cluster in three steps:

1. Configure a repository
2. Make a snapshot of the cluster
3. Restore from the snapshot

For more detailed information, refer to the Elasticsearch administration guide, and in particular to the documentation on the Snapshot and Restore module.

## 402.1 Creating a Repository

---

First create a repository to store your snapshots. Several repository types are supported:

- Shared file system, such as a Network File System or NAS
- Amazon S3
- HDFS (Hadoop Distributed File System)
- Azure Cloud

If using a shared file system repository type, first register the path to the shared file system in each node's `elasticsearch.yml` using the `path.repo` setting.

```
path.repo: ["path/to/shared/file/system/"]
```

Once the path to the folder hosting the repository is registered (make sure the folder exists), create the repository with a PUT command. For example,

```
curl -X PUT "localhost:9200/_snapshot/test_backup" -H 'Content-Type: application/json' -d'
{
 "type": "fs",
 "settings": {
 "location": "/path/to/shared/file/system/"
 }
}'
```

Replace `localhost:9200` with the proper `hostname:port` combination for your system, replace `test_backup` with the name of the repository to create, and use the absolute path to your shared file system in the location.

If the repository is set up successfully, you see this message:

```
{"acknowledged":true}
```

Once the repository exists, you can start creating snapshots.

## 402.2 Taking Snapshots of the Cluster

---

The easiest snapshot approach is to create a snapshot of all the indexes in your cluster. For example,

```
curl -XPUT localhost:9200/_snapshot/test_backup/snapshot_1
```

If `{"accepted":true}` appears in the terminal, the snapshot was a success.

It's possible to be more selective when taking snapshots. For example, if you use LES Monitoring, you can exclude the monitoring indexes. Explicitly declare the indexes to include in the snapshot. For example,

```
curl -XPUT localhost:9200/_snapshot/test_backup/snapshot_2
{ "indices": "liferay-0,liferay-20116" }
```

**Note:** For a list of all the Elasticsearch indexes, use this command:

```
curl -X GET "localhost:9200/_cat/indices?v"
```

This shows the index metrics:

| health | status | index         | uuid                   | pri | rep | docs.count | docs.deleted | store.size | pri.store.size |
|--------|--------|---------------|------------------------|-----|-----|------------|--------------|------------|----------------|
| green  | open   | liferay-20099 | obqiNE1_SDqfuz7rincrGQ | 1   | 0   | 195        | 0            | 303.1kb    | 303.1kb        |
| green  | open   | liferay-47206 | 3YEjtye1S90VT0i0EZcXcw | 1   | 0   | 7          | 0            | 69.7kb     | 69.7kb         |
| green  | open   | liferay-0     | shBWwpkXRxuAmGEaE475ug | 1   | 0   | 147        | 1            | 390.9kb    | 390.9kb        |

It's important to note that Elasticsearch uses a *smart snapshots* approach. To understand what that means, consider a single index. The first snapshot includes a copy of the entire index, while subsequent snapshots only include the delta between the first, complete index snapshot and the current state of the index.

Eventually you'll end up with a lot of snapshots in your repository, and no matter how cleverly you name the snapshots, you may forget what some snapshots contain. For this purpose, the Elasticsearch API provides getting information about any snapshot. For example:

```
curl -XGET localhost:9200/_snapshot/test_backup/snapshot_1
```

returns

```
{
 "snapshots": [
 {
 "snapshot": "snapshot_1",
 "uuid": "W1SjvJwHRh-xlAny7zeW3w",
 "version_id": 6.80399,
 "version": "6.8.2",
 "indices": ["liferay-20099", "liferay-0", "liferay-47206"],
 "state": "SUCCESS",
 "start_time": "2018-08-15T21:40:17.261Z",
 "start_time_in_millis": 1534369217261,
 "end_time": "2018-08-15T21:40:17.482Z",
 "end_time_in_millis": 1534369217482,
 "duration_in_millis": 221,
 "failures": [],
 "shards": {
 "total": 3,
 "failed": 0,
 "successful": 3
 }
 }
]
}
```

There's lots of useful information here, including which indexes were included in the snapshot. If you want to get rid of a snapshot, use the DELETE command.

```
curl -XDELETE localhost:9200/_snapshot/test_backup/snapshot_1
```

You might trigger creation of a snapshot and regret it (for example, you didn't want to include all the indexes in the snapshot). If your snapshots contain a lot of data, this can cost time and resources. To cancel the ongoing creation of a snapshot, use the same DELETE command. The snapshot process is terminated and the partial snapshot is deleted from the repository.

### 402.3 Restoring from a Snapshot

---

What good is a snapshot if you can't use it to restore your search indexes in case of catastrophic failure? Use the `_restore` API to restore all the snapshot's indexes:

```
curl -XPOST localhost:9200/_snapshot/test_backup/snapshot_1/_restore
```

Restore only specific indexes from a snapshot by passing in the `indices` option, and rename the indexes using the `rename_pattern` and `rename_replacement` options:

```
curl -XPOST
localhost:9200/_snapshot/test_backup/snapshot_1/_restore
{
 "indices": "liferay-20116",
 "rename_pattern": "liferayindex_(.+)",
 "rename_replacement": "restored_liferayindex_$1"
}
```

This restores only the index named `liferay-20116index_1` from the snapshot. The `rename...` settings specify that the beginning `liferayindex_` are replaced with `restored_liferayindex_`, so `liferay-20116index_1` becomes `restored_liferay-20116index_1`.

As with the process for taking snapshots, an errant restored index can be canceled with the `DELETE` command:

```
curl -XDELETE localhost:9200/restored_liferay-20116index_3
```

Nobody likes catastrophic failure on a production system, but Elasticsearch's API for taking snapshots and restoring indexes can help you rest easy knowing that your search cluster can be restored if disaster strikes. For more details and options, read Elastic's documentation on the [Snapshot and Restore Module](#).

---

## UPGRADING TO ELASTICSEARCH 6.5

---

Elasticsearch 6.5.x is the default, most up-to-date supported version of Elasticsearch for Liferay DXP. If you're upgrading @product@ and still running Elasticsearch 6.1, it's time to upgrade your Elasticsearch servers too. If you're setting up a new system and not already running a remote Elasticsearch 6.1.x server, follow the installation guide to install Elasticsearch and the configuration guide to configure the Elasticsearch adapter. Here, you'll learn to upgrade an existing Elasticsearch 6.1.x server (or cluster) to Elasticsearch 6.5.x:

1. Install and configure Elasticsearch 6.5.x.
2. Disable X-Pack Security in `elasticsearch.yml` unless you have an Liferay Enterprise Search subscription, which gives you access to the Liferay Enterprise Search Security app:

```
xpack.security.enabled: false
```

3. Configure the bundled Liferay Connector to Elasticsearch 6.
4. Re-index all search and spell check indexes.

---

**Before Proceeding**, back up your existing data before upgrading Elasticsearch. If something goes wrong during or after the upgrade, roll back to the previous version using the uncorrupted index snapshots. See here for more information.

---

Learn about configuring Elasticsearch here.

---

### 403.1 Re-index

---

Once the Elasticsearch adapter is installed and talking to the Elasticsearch cluster, navigate to *Control Panel* → *Configuration* → *Server Administration*, and click *Execute* for the *Reindex all search indexes* entry.

You must also re-index the spell check indexes.

### **403.2 Reverting to Elasticsearch 6.1**

---

Stuff happens. If that stuff involves an unrecoverable failure during the upgrade to Elasticsearch 6.5, roll back to Elasticsearch 6.1 and regroup.

Since your 6.1 and 6.5 are currently two separate installations, this procedure is straightforward:

1. Stop the Liferay Connector to Elasticsearch 6.
2. Stop Elasticsearch 6.5 and make sure that the Elasticsearch 6.1 `elasticsearch.yml` and the connector app are configured to use the same port (9200 by default).
3. Start the Elasticsearch server, and then restart the Liferay Connector to Elasticsearch 6.



---

## UPGRADING TO ELASTICSEARCH 6.8

---

Elasticsearch 6.8.x is supported for 7.0. If you're upgrading Liferay DXP and still running Elasticsearch 6.1, it's time to upgrade your Elasticsearch servers too. If you're setting up a new system and not already running a remote Elasticsearch 6.1.x server, follow the installation guide to install Elasticsearch and the configuration guide to configure the Elasticsearch adapter. Here, you'll learn to upgrade an existing Elasticsearch 6.1.x server (or cluster) to Elasticsearch 6.8.x:

1. Install and configure Elasticsearch 6.8.x.
2. Disable X-Pack Security in `elasticsearch.yml` unless you have an Liferay Enterprise Search subscription which gives you access to the LES Security app:

```
xpack.security.enabled: false
```

3. Configure the bundled Liferay Connector to Elasticsearch 6.
4. Re-index all search and spell check indexes.

---

**Before Proceeding**, back up your existing data before upgrading Elasticsearch. If something goes wrong during or after the upgrade, roll back to the previous version using the uncorrupted index snapshots. See here for more information.

---

Learn about configuring Elasticsearch here.

---

### 404.1 Re-index

---

Once the Elasticsearch adapter is installed and talking to the Elasticsearch cluster, navigate to *Control Panel* → *Configuration* → *Server Administration*, and click *Execute* for the *Reindex all search indexes* entry.

You must also re-index the spell check indexes.

## 404.2 Reverting to Elasticsearch 6.1

---

Stuff happens. If that stuff involves an unrecoverable failure during the upgrade to Elasticsearch 6.8, roll back to Elasticsearch 6.1 and regroup.

Since your 6.1 and 6.8 are currently two separate installations, this procedure is straightforward:

1. Stop the Liferay Connector to Elasticsearch 6.
2. Stop Elasticsearch 6.8 and make sure that the Elasticsearch 6.1 `elasticsearch.yml` and the connector app are configured to use the same port (9200 by default).
3. Start the Elasticsearch server, and then restart the Liferay Connector to Elasticsearch 6.

---

## UPGRADING TO ELASTICSEARCH 7

---

Elasticsearch 7 is supported for 7.0. If you're upgrading Liferay DXP and still running Elasticsearch 6, consider upgrading your Elasticsearch servers too. If you're setting up a new system and not already running a remote Elasticsearch 6 server, follow the installation guide to install Elasticsearch and the configuration guide to configure the Elasticsearch adapter.

---

**Before Proceeding**, back up your existing data before upgrading Elasticsearch. If something goes wrong during or after the upgrade, roll back to the previous version using the uncorrupted index snapshots. See [here](#) for more information.

---

Here, you'll learn to upgrade an existing Elasticsearch 6 server (or cluster) to Elasticsearch 7:

1. Install and configure Elasticsearch 7.
2. Back up the application specific indexes for Workflow Metrics and Result Rankings.
3. In 7.0, security is now provided out of the box. If you're using X-Pack security, enable it (it's disabled by default):

```
xpack.security.enabled: true
```

4. Blacklist the bundled Liferay Connector to Elasticsearch 6.
5. Install and configure the Liferay Connector to Elasticsearch 7.
6. Re-index all search and spell check indexes.

---

**Known Issue:** See LPS-103938. The Liferay Connector to Elasticsearch 7 throws an exception in the log when the LPKG file is deployed. There are no known functional impacts. If unexpected errors occur, re-start the Liferay DXP server.

---

Learn about configuring Elasticsearch [here](#).

### 405.1 Backing up Application-Specific Indexes

---

To preserve data stored in application-specific indexes, use a rolling upgrade for each index you need to preserve across the upgrade.

**Synonym Sets:** If you follow the workaround for the bug LPS-100272, your Synonym sets are preserved across the upgrade, as they are stored in the index settings directly, and not in their own index.

---

### 405.2 Blacklisting Elasticsearch 6

---

To blacklist Elasticsearch 6,

1. Create a configuration file named

```
com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.config
```

2. Give it these contents:

```
blacklistBundleSymbolicNames=[\
 "com.liferay.portal.search.elasticsearch6.api", \
 "com.liferay.portal.search.elasticsearch6.impl", \
 "com.liferay.portal.search.elasticsearch6.spi", \
 "com.liferay.portal.search.elasticsearch6.xpack.security.impl", \
 "Liferay Connector to X-Pack Security [Elastic Stack 6.x] - Impl", \
 "Liferay Enterprise Search Security - Impl" \
]
```

### 405.3 Re-index

---

Once the Elasticsearch adapter is installed and talking to the Elasticsearch cluster, navigate to *Control Panel* → *Configuration* → *Search*, and click *Execute* for the *Reindex all search indexes* entry.

You must also re-index the spell check indexes.

### 405.4 Reverting to Elasticsearch 6

---

Stuff happens. If that stuff involves an unrecoverable failure during the upgrade to Elasticsearch 7, roll back to Elasticsearch 6 and regroup.

Since your Elasticsearch 6 and 7 are currently two separate installations, this procedure takes only a few steps:

1. Stop the Liferay Connector to Elasticsearch 6.
2. Stop Elasticsearch 7 and make sure that the Elasticsearch 6 `elasticsearch.yml` and the connector app are configured to use the same port (9200 by default).
3. Start the Elasticsearch server, and then restart the Liferay Connector to Elasticsearch 6.



---

# INSTALLING LIFERAY ENTERPRISE SEARCH

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

A Liferay Enterprise Search (LES) subscription gets you additional features beyond what's available out of the box with your Liferay DXP subscription. It includes

- Liferay Enterprise Search Security\*
- Liferay Enterprise Search Monitoring
- Liferay Enterprise Search Learning to Rank

\* A LES subscription is not necessary if using Elasticsearch 7 via the \_Liferay Connector to Elasticsearch 7 as X-Pack's security features are bundled. See the LES compatibility matrix for more information.

X-Pack is an Elasticsearch extension for securing and monitoring Elasticsearch clusters. If you use Elasticsearch, you should secure it with X-Pack. The security features of X-Pack include authenticating access to the Elasticsearch cluster's data and encrypting Elasticsearch's internal and external communications. These are necessary security features for most production systems. A LES subscription gets you access to two connectors if you're using Elasticsearch 6: monitoring and security. Elasticsearch 7 bundles these security features, and Liferay has followed suit. Therefore, security is bundled with the Liferay Connector to Elasticsearch 7, and no LES subscription is necessary. Because of this, the documentation for installing Liferay Enterprise Search Security on Liferay DXP has been moved from the LES documentation section (this section) to the Elasticsearch installation and configuration guide. Contact Liferay's Sales department for more information.

Here's an overview of using the LES applications with Liferay DXP:

1. Get an Enterprise Search subscription.
2. You'll receive a license for X-Pack monitoring. Install it on your Elasticsearch servers.  
**Note:** If using Elasticsearch 6, you'll also need a LES subscription for X-Pack security.
3. Download and install the Liferay Enterprise Search apps you purchased. Find them in the Help Center Downloads page, choosing Enterprise Search from the Product drop-down menu.
4. Configure the connectors with the proper credentials, encryption information, and settings.
5. Restart Elasticsearch. These steps require a full cluster restart.

More detailed installation instructions are available in the article for each LES feature.

Elastic's documentation explains additional configuration options, features, and the architecture of X-Pack.

---

**Note:** Out of the box, X-Pack comes with a 30-day trial. This can be useful if there's a delay between your subscription and receipt of your production X-Pack license.

---

Now configure security, monitoring, and/or Learning to Rank, depending on your needs.



# INSTALLING LIFERAY ENTERPRISE SEARCH MONITORING

---

First configure security if you're using X-Pack's security features. Then come back here for instructions on installing and configuring Kibana (the monitoring server) with X-Pack so that Elasticsearch (secured with X-Pack), Kibana (secured with X-Pack), and Liferay DXP can communicate effortlessly and securely. A Liferay Enterprise Search (LES) subscription is necessary for this integration. Contact Liferay's Sales department for more information.

---

**The same monitoring connector is used for Elasticsearch 6 and 7:** When first created, the Liferay Enterprise Search Monitoring app was intended to be used only with Elasticsearch 6. However, testing revealed that the same connector works with Elasticsearch 7, so a new connector is not necessary if you're using Elasticsearch 7.

---

To install X-Pack monitoring,

1. Tell Elasticsearch to enable data collection.
2. Download and install Kibana.
3. Configure Kibana with the proper security settings.
4. Install the Liferay Enterprise Search Monitoring app.
5. Configure the connector to communicate with Elasticsearch.

This document assumes you're enabling security (with authentication and encrypted communication) *and* monitoring for Elasticsearch 7, but differences in the process for Elasticsearch 6 are noted where necessary.

---

## 407.1 Enable Encrypting Communication (TLS/SSL) in Elasticsearch and in Liferay DXP

---

Start by following the steps in this article to enable TLS/SSL in your Elasticsearch and Liferay DXP installation.

Then continue by enabling data collection in Elasticsearch.

## 407.2 Enable Data Collection

---

Monitoring is enabled on Elasticsearch by default, but data collection isn't. Enable data collection by adding this line to `elasticsearch.yml`.

```
xpack.monitoring.collection.enabled: true
```

Now install Kibana.

## 407.3 Install Kibana

---

Make sure to install the correct version of Kibana. Check the Liferay Enterprise Search compatibility matrix for details.

1. Download Kibana and extract it. The root folder is referred to as *Kibana Home*.
2. Tell Kibana where to send monitoring data by setting Elasticsearch's URL in `kibana.yml`:

```
elasticsearch.hosts: ["https://localhost:9200"]
```

On 6.5 and below, use

```
elasticsearch.url: "https://localhost:9200"
```

If TLS/SSL is not enabled on Elasticsearch, this is an `http` URL, otherwise use `https`.

3. If not using X-Pack security, start Kibana by opening a command prompt to Kibana Home and entering this command:

```
./bin/kibana
```

If you're using X-Pack's security features on the Elasticsearch server, there's additional configuration required before starting Kibana.

## 407.4 Configure Kibana with Authentication

---

If X-Pack requires authentication to access the Elasticsearch cluster, follow these steps or refer to Elastic's documentation.

1. Set the password for the built-in kibana user in [Kibana Home]/config/kibana.yml:

```
elasticsearch.username: "kibana"
elasticsearch.password: "liferay"
```

Use your kibana user password from your X-Pack setup. Once Kibana is installed, you can change the built-in user passwords from the *Management* user interface.

2. If you're not encrypting communication with the Elasticsearch cluster, start Kibana from Kibana home.

```
./bin/kibana
```

3. Go to <http://localhost:5601> and make sure you can sign in as a user who has the kibana\_user role or a superuser (like the elastic user).

## 407.5 Configuring Kibana with Encryption (TLS/SSL)

---

Follow these steps to configure Kibana if X-Pack encrypts communication with the Elasticsearch cluster. Consult Elastic's guide for more information.

1. Copy the [Elasticsearch Home]/config/certs folder into the [Kibana Home]/config/ folder.

This example reuses the certificate files created for Elasticsearch itself. If you wish to generate a separate certificate for your Kibana instance, make sure it is signed by the same CA as the Elasticsearch node certificates.

2. Add these settings to kibana.yml:

```
xpack.security.encryptionKey: "xsoomethingxatleastx32xcharactersx"
xpack.security.sessionTimeout: 600000

elasticsearch.hosts: ["https://localhost:9200"]

elasticsearch.ssl.verificationMode: certificate
elasticsearch.ssl.certificateAuthorities: ["config/certs/ca.crt"]
elasticsearch.ssl.certificate: config/certs/localhost.crt
elasticsearch.ssl.key: config/certs/localhost.key

server.ssl.enabled: true
server.ssl.certificateAuthorities: ["config/certs/ca.crt"]
server.ssl.certificate: config/certs/localhost.crt
server.ssl.key: config/certs/localhost.key
```

Elasticsearch/Kibana 6.5 and below use a different property for specifying the host URL. Replace the `elasticsearch.hosts` property with

```
elasticsearch.url: "https://localhost:9200"
```

For more information about monitoring and security best practices in a clustered environment, refer to Elastic's documentation.

After this step you can access Kibana at `https://localhost:5601` and sign in with a Kibana user. The last step is to connect Kibana to Liferay DXP.

## 407.6 Configuring the Liferay Enterprise Search Monitoring app

---

If you have a LES subscription, download the Liferay Enterprise Search Monitoring app . Install the LPKG file by copying it into the Liferay Home/Deploy folder.

1. Once the connector is installed and Kibana and Elasticsearch are securely configured, create a configuration file named

```
com.liferay.portal.search.elasticsearch6.xpack.monitoring.web.internal.configuration.XPackMonitoringConfiguration.config
```

2. Place these settings in the `.config` file:

```
kibanaPassword="liferay"
kibanaUserName="elastic"
kibanaURL="https://localhost:5601"
```

The values depend on your Kibana configuration. For example, use a URL such as `kibanaURL="http://localhost:5601"` if you are not using X-Pack Security TLS/SSL features.

Alternatively, configure the monitoring adapter from System Settings. Navigate to *Control Panel* → *Configuration* → *System Settings* and find the X-Pack Monitoring entry in the Search category. All the configuration options for the monitoring connector appear there.

3. Deploy this configuration file to Liferay Home/osgi/configs, and your running instance applies the settings. There's no need to restart the server.
4. There are two more settings to add to Kibana itself. The first forbids Kibana from rewriting requests prefixed with `server.basePath`. The second sets Kibana's base path for the Monitoring portlet to act as a proxy for Kibana's monitoring UI. Add this to `kibana.yml`:

```
server.rewriteBasePath: false
server.basePath: "/o/portal-search-elasticsearch-xpack-monitoring/xpack-monitoring-proxy"
```

Note that once you set the `server.basePath`, you cannot access the Kibana UI through Kibana's URL (e.g., `https://localhost:5601`). All access to the Kibana UI is through the Monitoring portlet, which is only accessible to signed in Liferay DXP users. Navigate directly to the portlet using this URL:

```
http://localhost:8080/o/portal-search-elasticsearch-xpack-monitoring/xpack-monitoring-proxy/app/monitoring
```

5. Because you're using the Monitoring portlet in Liferay DXP as a proxy to Kibana's UI, if you are using X-Pack Security with TLS/SSL, you must configure the application server's startup JVM parameters to recognize a valid *truststore* and *password*.

First, navigate to Elasticsearch Home and generate a PKCS#12 certificate from the CA you created when setting up X-Pack security:

```
./bin/elasticsearch-certutil cert --ca-cert /path/to/ca.crt --ca-key /path/to/ca.key --ip 127.0.0.1 --dns localhost --name localhost --out /path/to/Elasticsearch_Home/config/localhost.p12
```

Next use the `keytool` command to generate a truststore:

```
keytool -importkeystore -deststorepass liferay -destkeystore /path/to/truststore.jks -srckeystore /path/to/Elasticsearch_Home/config/localhost.p12 -srcstoretype PKCS12 -srcstorepass liferay
```

Add the truststore path and password to your application server's startup JVM parameters. Here are example truststore and path parameters for appending to a Tomcat server's `CATALINA_OPTS`:

```
-Djavax.net.ssl.trustStore=/path/to/truststore.jks -Djavax.net.ssl.trustStorePassword=liferay
```

Restart Liferay DXP and Kibana.

## 407.7 Monitoring in Liferay DXP

---

Once Kibana and X-Pack are successfully installed and configured and all the servers are running, add the X-Pack Monitoring portlet to a page:

1. Open the *Add* menu on a page and choose *Widgets*
2. Search for *monitoring* and drag the *X-Pack Monitoring* widget from the Search category onto the page.

See the Elastic documentation for information on monitoring Elasticsearch.

## 407.8 Example Kibana Configuration

---

Here are the `kibana.yml` properties demonstrated in this article, for convenient copy/pasting:

```
X-Pack Security enabled (Basic Auth)
elasticsearch.username: "kibana"
elasticsearch.password: "liferay"

When TLS/SSL is enabled in X-Pack Security
xpack.security.encryptionKey: "xsoemthingatleastx32xcharactersx"
```

```
xpack.security.sessionTimeout: 600000

If on Elasticsearch 6.5 or below, replace the next property with:
elasticsearch.url: "http://localhost:9200"
elasticsearch.hosts: ["https://localhost:9200"]

Enable TLS/SSL for out-bound traffic: from Kibana to Elasticsearch
elasticsearch.ssl.verificationMode: certificate
elasticsearch.ssl.certificateAuthorities: ["config/certs/ca.crt"]
elasticsearch.ssl.certificate: config/certs/localhost.crt
elasticsearch.ssl.key: config/certs/localhost.key

Enable TLS/SSL for in-bound traffic: from browser or
DXP (X-Pack Monitoring widget, proxy) to Kibana
server.ssl.enabled: true
server.ssl.certificateAuthorities: ["config/certs/ca.crt"]
server.ssl.certificate: config/certs/localhost.crt
server.ssl.key: config/certs/localhost.key

To use Kibana inside the X-Pack Monitoring widget
server.rewriteBasePath: false
server.basePath: "/o/portal-search-elasticsearch-xpack-monitoring/xpack-monitoring-proxy"
```

---

# LIFERAY ENTERPRISE SEARCH: LEARNING TO RANK

---

Search engines like Elasticsearch have well-tuned relevance algorithms, good for general search purposes. Sometimes, this “generally good” relevance scoring just isn’t good enough. You can attain more perfect search results by employing machine learning.

Learning to Rank harnesses machine learning to improve search result rankings. It combines the expertise of data scientists with machine learning to produce a smarter scoring function that’s applied to search queries.

7.0, Service Pack 1/Fix Pack 2 and later, supports Learning to Rank through its support of Elasticsearch versions 6.x and 7.4.x. It requires a LES subscription. It’s important to understand that the Elasticsearch Learning to Rank plugin is not produced by Elastic, and there is not a pre-built plugin for all of Liferay DXP’s supported Elasticsearch versions.

---

## 408.1 Disabling Learning to Rank on a Search Page

---

Learning to Rank does not work with the Sort widget.

If you must use Learning to Rank in your Liferay DXP instance, but want to disable it on a particular Search page (perhaps to use the Sort widget), you can:

1. Add a Low Level Search Options widget to the Search page.
2. Open the widget’s Configuration screen by clicking *Configure additional low level search options in this page*.
3. In the *Contributors to Exclude* field, enter `com.liferay.portal.search.learning.to.rank`

Now the Learning to Rank re-scoring process is skipped for queries entered into the page’s Search Bar, and results are sortable in the Sort widget and returned using the default relevance algorithm.

## 408.2 Prerequisites

---

There are some prerequisites for using Learning to Rank to re-score Liferay queries sent to Elasticsearch:

- If using Elasticsearch 7, 7.0 Service Pack 1/Fix Pack 2 or later is required, with the appropriate Elasticsearch Connector version installed.
- If using Elasticsearch 6, 7.0 Fix Pack 3 or later is required, with the appropriate Elasticsearch Connector version installed.
- A Liferay Enterprise Search (LES) subscription is required for Learning to Rank. Once you have a subscription, download the Liferay Enterprise Search Learning to Rank LPKG file and install it to your Liferay DXP server.
- A remote Elasticsearch server, with your data indexed into it.
- The corresponding version of the Elasticsearch Learning to Rank plugin installed into Elasticsearch.
- A trained model uploaded into the Learning to Rank plugin.

To understand more about the compatibility requirements for LES, see its compatibility matrix. How does Learning to Rank work?

## 408.3 Technical Overview

---

In a normal search, the User sends a query to the search engine via Liferay DXP's Search Bar. The order of returned results is dictated by the search engine's relevance scoring algorithm.

Here's where Learning to Rank intervenes and makes that process different:

1. User enters a query into the search bar.
2. Liferay sends the query to Elasticsearch, and retrieves the first 1000 results as usual, using the search engine's relevance algorithm.
3. The top 1000 results are not returned as search hits, but are used by Elasticsearch for re-scoring via the rescore functionality.
4. The results are re-scored by the SLTR query, which includes the keywords and the trained model to use for re-scoring.
5. Once the trained model re-ranks the results, they're returned in Liferay's Search Results in their new order.

Though it's just a sub-bullet point in the ordered list above, much of the work in this paradigm is in creating and honing the trained model. That's beyond the scope of Liferay's role, but we'll help you get all the pieces in place to orchestrate the magic of machine learning on your Liferay queries. Here's a brief overview of what constitutes *model training*.



## 408.4 Model Training

---

A useful trained model is produced when a good judgment list and a good feature set are fed to a Learning to Rank algorithm (this is the machine learning part of the puzzle). Therefore, it's incumbent on you to assemble

- The Learning to Rank algorithm you wish to use for creating a training model. This demonstration uses RankLib.
- A *judgment list*, containing a graded list of search results. The algorithm is designed to produce a model that honors the ordering of the judgment list.
- A feature set, containing all the *features* you're handing to the Learning to Rank algorithm, which it uses in conjunction with the judgment list to produce a reliable model. An example feature set for Liferay DXP data is shown in the next article.

Judgment lists are lists of graded search results.

Features are the variables that the algorithm uses to create a function that can score results in a smarter way. If you don't give enough, or the correct, relevant features, your model will not be "smart" enough to provide improved results.

In the next article you'll see the steps required to configure Learning to Rank with Liferay DXP.



---

## CONFIGURING LEARNING TO RANK

---

Before beginning, you must have a remote Elasticsearch 6 or 7 cluster communicating with 7.0. See the compatibility matrix for more information

---

**Helpful hint:** Use Suggestions to discover the most common queries (this can be one way to decide which queries to create Learning to Rank models for).

---

### 409.1 Step 1: Install the Learning to Rank Plugin on Elasticsearch

---

See the Elasticsearch Learning to Rank plugin documentation to learn about installing the Learning to Rank plugin.

You'll be running a command like this one, depending on the plugin version you're installing:

```
./bin/elasticsearch-plugin install http://es-learn-to-rank.labs.o19s.com/ltr-1.1.0-es6.5.4.zip
```

If using X-Pack security in your Elasticsearch cluster, there may be additional steps.

### 409.2 Step 2: Training and Uploading a Model

---

Detailed instructions on training models is outside the scope of this guide. This requires the intervention of data scientists, who can recommend appropriate tools and models. Use what works for you. In doing so, you'll almost certainly be compiling Judgment lists and feature sets that can be used by the training tool you select to generate a model that produces good search results. This can be a long journey, but once you get to the end of it, you'll want to upload the model to the Learning to Rank plugin.

### 409.3 Upload the Model to the Learning to Rank Plugin

---

You'll upload the model using a POST request, but first you need to make sure you have a `_ltr` index and a feature set uploaded to the Learning to Rank plugin. Use Kibana (or even better, the Monitoring widget), to make these tasks easier.

1. If you don't already have an `_ltr` index, create one:

```
PUT _ltr
```

2. Add a feature set to the `_ltr` index. In this example the set is called `liferay`:

```
POST _ltr/_featureset/liferay
{
 "featureset": {
 "name": "liferay",
 "features": [
 {
 "name": "title",
 "params": [
 "keywords"
],
 "template": {
 "match": {
 "title_en_US": "{{keywords}}"
 }
 }
 },
 {
 "name": "content",
 "params": [
 "keywords"
],
 "template": {
 "match": {
 "content_en_US": "{{keywords}}"
 }
 }
 }
],
 {
 "name": "asset tags",
 "params": [
 "keywords"
],
 "template": {
 "match": {
 "assetTagNames": "{{keywords}}"
 }
 }
 }
]
}
```

Take note of the syntax used here, since it's required.

3. Add the trained model to the feature set:

```
POST _ltr/_featureset/liferay/_createmodel
{
 "model": {
 "name": "linearregression",
 "model": {
 "type": "model/ranklib",
 "definition": ""
 }
 }
 # Linear Regression
 # Lambda = 1.0E-10
 0:-0.717621803830712 1:-0.717621803830712 2:-2.235841905601106 3:19.546816765721594
 ""
}
}
```

This is a very high level set of instructions, because there's not much to do in Liferay DXP. To learn in more detail about what's required, see the Learning to Rank plugin's documentation. Keep reworking those judgment lists!

#### 409.4 Step 3: Enable Learning to Rank

---

Enable Learning to Rank from Control Panel → Configuration → System Settings → Search → Learning to Rank. There's a simple on/off configuration and a text field where you must enter the name of the trained model to apply to search queries.

The model in the previous step was named `linearregression`, so that's what you'd enter.

That's all the configuration required to get the Elasticsearch Learning to Rank plugin ingesting a trained model, a feature set, and search queries from Liferay DXP.

## Learning to Rank

This configuration is not saved yet. The values shown are the default.

Enabled

By checking this option, all Elasticsearch queries will be rescored by the Learning to Rank model.

### Model

linearregression

The name of the model in the Learning to Rank configuration index.

Save

Cancel

Figure 409.1: Enable Learning to Rank in Liferay DXP from the System Settings entry.

---

# INSTALLING SOLR

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Solr is a popular enterprise search platform built on Apache Lucene. It's reliable, scalable, and fault tolerant. Read more about it [here](#).

Elasticsearch is the default search engine that ships with Liferay DXP, and some Liferay Search features are only available on Elasticsearch. It's valid, however, to use Solr instead. In particular, if you've already been using Solr with a previous version of Liferay DXP, or your platform (for example, your OS or JVM) isn't supported by Elasticsearch, you might choose to use Solr to search and index your Liferay DXP data.

There are circumstances that force you to use Elasticsearch instead of Solr. Read [here](#) for more information.

Liferay DXP 7.2, Fix Pack 1, supports Solr 7.5.x through the Liferay Connector to Solr 7 application, version 2.0.0.

Liferay DXP 7.2, Service Pack 1/Fix Pack 2 and later, supports Solr 7.5.x through the Liferay Connector to Solr 7 application, version 2.0.1.

Liferay Portal CE 7.2, GA2 and later (not available at time of writing), support Solr 7.5.x through the Liferay CE Connector to Solr 7 application.

---

**Upgrading to Service Pack 1 or Fix Pack 2 (or later) requires installation of a new Solr connector:** If you were running version 2.0.0 of the Liferay Connector to Solr 7 application, and you want to install Service Pack 1/Fix Pack 1 (or later), you must install version 2.0.1 of the Liferay Connector to Solr 7 application.

---

## 410.1 Blacklisting Elasticsearch-Only Features

---

Before installing Solr, you must blacklist certain DXP features that only work with Elasticsearch.

1. Create a configuration file named

```
com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.config
```

2. Give it these contents:

```
blacklistBundleSymbolicNames=["com.liferay.portal.search.tuning.web.api","com.liferay.portal.search.tuning.web","com.liferay.portal.search.tuning.web"]
```

3. Place the file in Liferay Home/osgi/configs.

It is required during the Solr installation process to also stop the Elasticsearch Connectors that ship with Liferay DXP. If you're ready to blacklist those bundles now, use these contents in the blacklist configuration file:

```
blacklistBundleSymbolicNames=["com.liferay.portal.search.tuning.web.api","com.liferay.portal.search.tuning.web","com.liferay.portal.search.tuning.web"]
```

The Liferay Enterprise Search bundles must be excluded if you don't have a LES subscription.



---

## INSTALLING SOLR: BASIC INSTALLATION

---

This document has been updated and ported to Liferay Learn and is no longer maintained here. There are two ways to install the Liferay Connector to Solr 7:

1. Navigate to Liferay Marketplace and download the app that corresponds to your portal.

Once the app LPKG is downloaded, copy it to `Liferay_Home/osgi/marketplace`.

2. In your running portal instance, navigate to *Control Panel* → *Apps* → *Store*. Sign in using your credentials, search for Solr Search Engine, and purchase (it's free) the Liferay Connector to Solr 7 entry.

As you proceed, remember these terms:

*Solr Home*: The center of the Solr system (pun intended). This directory is `solr-[version]/server/solr`.

*Liferay Home*: The root folder of your Liferay DXP installation. It contains the `osgi`, `deploy`, `data`, and `license` folders, among others.

There are two installation steps:

1. Installing and configuring Solr 7.
2. Installing and configuring the Solr 7 connector for Liferay DXP.

Before configuring Liferay DXP for Solr, install and set up Solr.

---

### 411.1 Installing and Configuring Solr 7

---

To install and properly configure Solr for Liferay DXP:

1. Download Solr and unzip it.

2. Navigate to `solr-[version]/server/solr`. This is Solr Home.
3. Create a new folder called `liferay`.
4. In the `liferay` folder, create two new folders: `conf` and `data`.
5. Copy the contents of `Solr_Home/configsets/_default/conf` to `Solr_Home/liferay/conf`.
6. Open the Liferay Connector to Solr 7's LPKG file with an archive manager.  
Open the `com.liferay.portal.search.solr7.impl.jar` file, and extract

```
META-INF/resources/solrconfig.xml
```

and

```
META-INF/resources/schema.xml
```

to

```
Solr_Home/liferay/conf
```

This replaces the current `solrconfig.xml` and `schema.xml` files with ones that tell Solr how to index data coming from Liferay DXP.

7. Create a `core.properties` file in `Solr_Home/liferay` and add this configuration:

```
config=solrconfig.xml
dataDir=data
name=liferay
schema=schema.xml
```

8. Checkpoint: your `Solr_Home/liferay` folder now has this structure:

```
liferay
├── conf
│ ├── lang
│ │ ├── contractions_ca.txt
│ │ └── ...txt
│ ├── managed-schema
│ ├── params.json
│ ├── protwords.txt
│ ├── schema.xml
│ ├── solrconfig.xml
│ ├── stopwords.txt
│ └── synonyms.txt
├── core.properties
└── data
```

9. Start the Solr server by entering

```
./bin/solr start -f
```

from the top-level folder of your Solr installation (`solr-[version]`).

10. The Solr server listens on port 8983 by default. Navigate to `http://localhost:8983/solr/#/~cores` (assuming you're testing locally with `localhost` as your host), and confirm that the `liferay` core is available.

Solr is now installed. Next install and configure the Solr connector.

## 411.2 Installing and Configuring the Liferay Solr Adapter

---


Since Elasticsearch is the default search engine, the Elasticsearch connector is already installed and running. You must stop it before configuring the Solr connector.

### 411.3 Stopping the Elasticsearch Connector

---

Stop the Elasticsearch connector bundle using the App Manager, the Felix Gogo shell, or the bundle blacklist. If you're a Liferay DXP customer, use the blacklist feature as described below. The App Manager and Gogo shell rely on the `osgi/state` folder to “remember” the state of the bundle. If you delete this folder (recommended during patching) the Elasticsearch connector is reinstalled and started automatically.

Navigate to Control Panel → Apps → App Manager.

Once in the App Manager, search for *elasticsearch*. Find the Liferay Connector to Elasticsearch 6 module and click the Actions ((  )) button. Choose the Deactivate option. This leaves the bundle installed, but stops it in the OSGi runtime.

Alternatively, use the Felix Gogo shell to stop the Elasticsearch connector. Enter

```
lb elasticsearch
```

You'll see two active bundles for the Liferay Connector to Elasticsearch 6: an API and an IMPL bundle.

```
ID|State |Level|Name
476|Active | 10|Liferay (CE) Connector to Elasticsearch 6 - API (3.0.0)
478|Active | 10|Liferay Portal Search Elasticsearch 6 API (3.0.4)
480|Active | 10|Liferay Portal Search Elasticsearch 6 SPI (3.2.1)
706|Active | 10|Liferay (CE) Connector to Elasticsearch 6 - Impl (3.0.0)
707|Active | 10|Liferay Portal Search Elasticsearch 6 Implementation (3.0.15)
```

Stop the API bundle by entering

```
stop [bundle ID]
```

In the example above, the `[bundle ID]` is 476.

**Liferay DXP:** DXP customers should blacklist the Elasticsearch, Shield, and Marvel plugins.

1. Create a

```
com.liferay.portal.bundle.blacklist.internal.BundleBlacklistConfiguration.config
```

file with these contents:

```
blacklistBundleSymbolicNames=["com.liferay.portal.search.elasticsearch6.spi","com.liferay.portal.search.elasticsearch6.api","com.liferay.port
```

If the LES Security and Monitoring LPKG files are installed, you must blacklist these too.

2. Place the file in Liferay Home/`osgi/configs`.

## 411.4 Install and Configure the Solr Connector

Now you're ready to install the connector:

1. Start Liferay DXP, then deploy the Solr connector by copying the LPKG you downloaded to Liferay\_Home/deploy.

You'll see a `STARTED` message in your Liferay DXP log once the Solr connector is installed. Here's what the log message looks like:

```
2018-11-06 19:59:49.396 INFO [pipe-start 943 944][BundleStartStopLogger:39] STARTED com.liferay.portal.search.solr7.api_2.0.5 [943]
2018-11-06 19:59:49.490 INFO [pipe-start 943 944][BundleStartStopLogger:39] STARTED com.liferay.portal.search.solr7.impl_2.0.11 [944]
```

2. To re-index against Solr, navigate to *Control Panel* → *Configuration* → *Search*, and click *Execute* next to the *Reindex all search indexes* option.

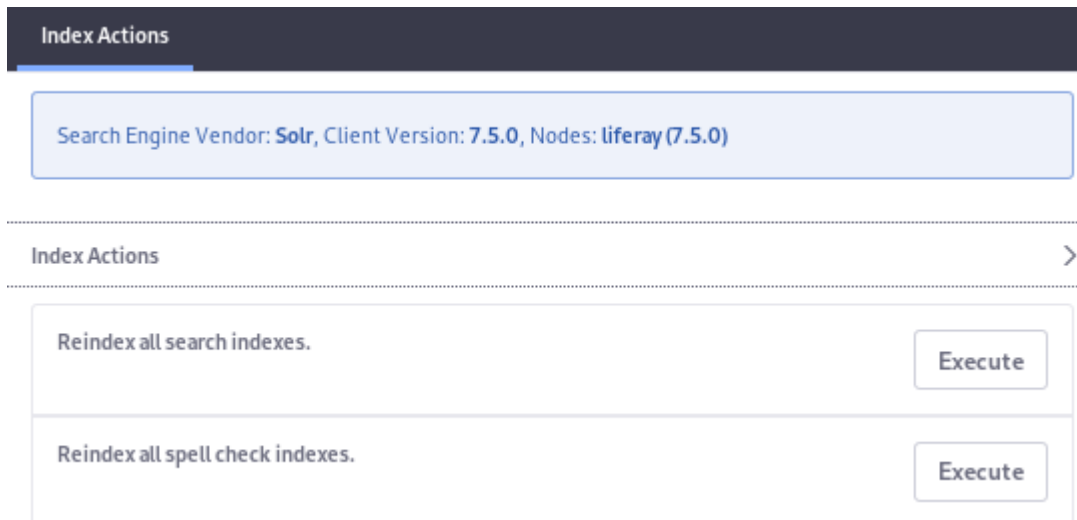


Figure 411.1: Once the Solr connector is installed, you can re-index your Liferay DXP data against your Solr server.

In production deployments, specify your edits to the Solr connector's default configurations using a configuration file deployed to the Liferay\_Home/osgi/configs folder. Name the file

```
com.liferay.portal.search.solr7.configuration.SolrConfiguration.config
```

During testing and development, use the Solr 7 System Settings entry *Control Panel* → *Configuration* → *System Settings* for editing the default configurations.

The next article covers clustering Solr with SolrCloud.

## SOLR 7

This configuration was not saved yet. The values shown are the default.

### Authentication mode

Basic

### Client type

Replicated

### Default collection

liferay

### Log exceptions only

Set to true to only log exceptions from Solr and not rethrow them.

### Read URL +

http://localhost:8983/solr/liferay

### Write URL +

http://localhost:8983/solr/liferay

### Zk host

localhost:9983

Figure 411.2: You can configure Solr from Liferay DXP's System Settings application. This is most useful during development and testing.



---

# INSTALLING SOLR: HIGH AVAILABILITY WITH SOLRCLOUD

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Use SolrCloud if you need a cluster of Solr servers. Note that to use SolrCloud in production, you should set up an external ZooKeeper ensemble. ZooKeeper is a centralized coordination service for managing distributed systems like your SolrCloud cluster.

The steps included here should be considered the bare minimum of what must be done to configure SolrCloud with Liferay DXP. For example, these instructions cover configuring SolrCloud on a single machine, whereas a production environment would feature multiple physical or virtual machines. These instructions also assume you've followed the earlier section on *Installing and Configuring Solr 7*. Refer to the SolrCloud guide for more information.

1. Stop the Solr server if it's running.
2. Navigate to the Solr\_Home/configsets folder and create a folder called

```
liferay_configs
```

3. Copy the conf folder from Solr\_Home/liferay to the liferay\_configs folder you just created.

The configset/liferay\_configs folder contains the SolrCloud Liferay DXP collection configuration and is uploaded to ZooKeeper. By copying the conf folder from the liferay server configured earlier, you're using the schema.xml and solrconfig.xml files provided with the Liferay Solr Adapter.

4. Next launch an interactive SolrCloud session to configure your SolrCloud cluster. Use this command:

```
./bin/solr -e cloud
```

5. Complete the setup wizard. These steps demonstrate creating a two-node cluster:
  - Enter 2 for the number of nodes.

- Specify ports 8983 and 7574 (the defaults). Both nodes are started with the start commands printed in the log:

```
Starting up Solr on port 8983 using command:
"bin/solr" start -cloud -p 8983 -s "example/cloud/node1/solr"
```

```
Waiting up to 180 seconds to see Solr running on port 8983 [[] [-]
Started Solr server on port 8983 (pid=8846). Happy searching!
```

```
Starting up Solr on port 7574 using command:
"bin/solr" start -cloud -p 7574 -s "example/cloud/node2/solr" -z localhost:9983
```

```
Waiting up to 180 seconds to see Solr running on port 7574 [[] [/]
Started Solr server on port 7574 (pid=9026). Happy searching!
```

- Name the collection *liferay*.
- Split the collection into two shards.
- Specify two replicas per shard.
- When prompted to choose a configuration, enter *liferay\_configs*. You should see a log message that concludes like this when the cluster has been started:

```
SolrCloud example running, please visit http://localhost:8983/solr
```

Now you have a new collection called *liferay* in your local SolrCloud cluster. Verify its status by running the *status* command:

```
./bin/solr status
```

You'll see log output like this:

```
Found 2 Solr nodes:
```

```
Solr process 12828 running on port 8983
```

```
INFO - 2019-07-18 16:46:35.137; org.apache.solr.util.configuration.SSLCredentialProviderFactory; Processing SSL Credential Provider chain: env;sysp
{
 "solr_home":"/home/russell/liferay-bundles/liferay-portal-7.2.10-gal/solr-7.5.0/example/cloud/node1/solr",
 "version":"7.5.0 b5bf70b7e32d7ddd9742cc821d471c5fabd4e3df - jimcz - 2018-09-18 13:07:55",
 "startTime":"2019-07-18T20:44:13.138Z",
 "uptime":"0 days, 0 hours, 2 minutes, 22 seconds",
 "memory":"56.4 MB (%11.5) of 490.7 MB",
 "cloud":{
 "ZooKeeper":"localhost:9983",
 "liveNodes":"2",
 "collections":"1"}}}
```

```
Solr process 12995 running on port 7574
```

```
INFO - 2019-07-18 16:46:35.848; org.apache.solr.util.configuration.SSLCredentialProviderFactory; Processing SSL Credential Provider chain: env;sysp
{
 "solr_home":"/home/russell/liferay-bundles/liferay-portal-7.2.10-gal/solr-7.5.0/example/cloud/node2/solr",
 "version":"7.5.0 b5bf70b7e32d7ddd9742cc821d471c5fabd4e3df - jimcz - 2018-09-18 13:07:55",
 "startTime":"2019-07-18T20:44:16.847Z",
 "uptime":"0 days, 0 hours, 2 minutes, 19 seconds",
 "memory":"108.2 MB (%22.1) of 490.7 MB",
 "cloud":{
 "ZooKeeper":"localhost:9983",
 "liveNodes":"2",
 "collections":"1"}}}
```

To stop Solr while running in SolrCloud mode, use the *stop* command, like this:

```
bin/solr stop -all
```



## 412.1 Configure the Solr Adapter for SolrCloud

---

There's only one thing left to do: specify the client type as *CLOUD* in Liferay's Solr connector.

1. From System Settings or your OSGi configuration file, set the *Client Type* to *CLOUD*.

```
clientType="CLOUD"
```

2. Start Liferay DXP if it's not running already.



Figure 412.1: From the Solr 7 System Settings entry, set the *Client Type* to *Cloud*.

Now you can configure Liferay DXP for Solr and Solr for @product@. Remember that Elasticsearch is the default search engine, so if you're not constrained to use Solr or already a Solr expert, consider Elasticsearch for your search engine requirements. If you do use Solr, tell all your colleagues that your Liferay DXP installation's search capability is Solr powered (pun intended).



---

## SOLR CONNECTOR SETTINGS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Solr can be configured for use with 7.0. Liferay Marketplace includes a Solr connector app called the Liferay Connector to Solr 7. The connector is configurable through System Settings or an OSGi configuration file named `com.liferay.portal.search.solr7.configuration.SolrConfiguration.config` and deployed to `[Liferay_Home]/osgi/configs`.

The list below is all the configuration settings for Liferay's Solr connector, in the order they appear in the System Settings application:

### 413.1 Solr 7

---

**authenticationMode=BASIC** A String with the value of *BASIC* or *CERT*. Use *BASIC* when connecting using the Basic Authentication plugin, otherwise select *CERT* to connect using 2-way SSL authentication.

**clientType=REPLICATED** A String with the value of *REPLICATED* or *CLOUD*. Use the default (*REPLICATED*) when connecting to a single-node Solr server. Specify *CLOUD* to connect to SolrCloud (see the next section, titled *High Availability with SolrCloud* for more information).

**logExceptionsOnly=true** A boolean value that, when set to true, only logs exceptions from Solr, without rethrowing them.

**readURL=http://localhost:8983/solr/liferay** A String array with the URLs to which Liferay will send search requests. This will be different from the `writeURL` if you use separate servers for indexing (write) and searching (read).

**writeURL=http://localhost:8983/solr/liferay** A String array with the URLs to which Liferay will send indexing requests. This is different from the `readURL` if you use separate servers for indexing (write) and searching (read).

**zkHost=localhost:9983** A String with the ZooKeeper host and port. This is required when using the adapter in *CLOUD* mode.



---

## SECURING LIFERAY DXP

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay follows the OWASP Top 10 (2013) and CWE/SANS Top 25 lists to ensure Liferay DXP is as secure as possible. Following these recommendations protects the product against known kinds of attacks and security vulnerabilities. For example, Liferay DXP's persistence layer is generated and maintained by the Service Builder framework which prevents SQL Injection using Hibernate and parameter based queries.

To prevent Cross Site Scripting (XSS), user-submitted values are escaped on output. To support integration features, Liferay DXP doesn't encode input. Data is stored in the original form as submitted by the user. Liferay DXP includes built-in protection against CSRF attacks, Local File Inclusion, Open Redirects, Uploading and serving files of dangerous types, Content Sniffing, Clickjacking, Path Traversal, and many other common attacks.

To stay on top, Liferay DXP also contains fixes for state-of-the-art attacks and techniques to improve product security. For example, Liferay DXP uses PBKDF2 to store passwords. Liferay DXP also contains mitigation for Quadratic Blowup XXE attack, Rosetta Flash vulnerability, Reflected File Download, and other kinds of attacks.

This section of tutorials shows you how to configure various security and login features, such as LDAP, single sign-on, Service Access Policies, and more. What follows is an overview of what's available.

### 414.1 Authentication Overview

---

Liferay DXP user authentication can take place using any of a variety of prepared solutions:

- Form authentication using the Sign In Portlet with extensible adapters for checking and storing credentials (portal database, LDAP)
- Single-Sign-On (SSO) solutions - NTLM, CAS, SiteMinder, OpenSSO, OpenID, Facebook
- SAML plugin
- JAAS integration with application server

Note: Although Liferay's SSO solutions are incompatible with WebDAV, they can be used with Liferay Sync. See the Publishing Files article for more information on WebDAV and Liferay Sync. You can authenticate and authorize apps remotely using the AuthVerifier layer:

- Password based HTTP Basic + Digest authentication
- Token based OAuth plugin
- Portal session based solution for JavaScript applications

Both user authentication and remote application authentication are extensible. Developers can create custom Login portlets and plugins, extend the default Login portlet auth.pipeline, create AutoLogin extensions for SSO, or create custom AuthVerifier implementations.

## 414.2 Authorization and Permission Checking

---

There are several adjustable authorization layers in place to prevent unauthorized or unsecured access to data:

- Remote IP and HTTPS transport check to limit access to Liferay DXP's Java servlets
- Extensible Access Control Policies layer to perform any portal service related authorization check
- Extensible role-based permission framework for almost any portal entity or data (stored in the portal database or elsewhere)
- Portlet Container security checks to control portlet access
- Remote IP check for portal remote API authentication methods
- Service Access Policies to control access to portal remote API

## 414.3 Additional Security Features

---

Users can be assigned to sites, teams, user groups, or organizations. Custom roles can be created, permissions can be assigned to those roles, and those roles can be applied to users. Roles are scoped to apply only with a specific context like a site, an organization, or globally. See the Roles and Permissions documentation for more information.

---

Note: Liferay DXP relies on the application server for sanitizing CRLF in HTTP headers. You should, therefore, make sure this is configured properly in your application server, or you may experience false positives in reports from automatic security verification software such as Veracode. There is one exception to this for Resin, which does not have support for this feature. In this case, Liferay DXP sanitizes HTTP headers.

---

There are additional security plugins available from Liferay Marketplace. For example, you can find an Audit plugin for tracking user actions or an AntiSamy plugin for clearing HTML from XSS vectors.

There are many ways to fine-tune or disable various security features. Here are a few examples of these kinds of configuration actions:

- Disable the Sign In portlet's *Create Account* link
- Configure Liferay DXP's HTTPS web server address
- Configure the list of allowed servers to which users can be redirected
- Configure the list of portlets that can be accessed from any page
- Configure the file types allowed to be uploaded and downloaded

#### **414.4 Secure Configuration and Run Recommendations**

---

Liferay DXP is built using the “secure by default” concept in mind. It's not recommended to disable built-in protections or to allow all values in security white-lists. Such acts may lead to security misconfiguration and an insecure deployment.

Also, customers are advised to deploy security patches as described on the customer portal.

For community and CE deployments, the stay secure by always using the latest community version, which contains all previous security patches.





---

# LOGGING INTO LIFERAY DXP

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

One of the primary functions of a security system is to make pages, content, and web applications accessible only to the appropriate users. A student logging into a university portal should not be able to access the same resources a professor can. A patient logging into a health care portal should not be able to access a doctor's resources. Some content (at least a login page) should be available to everybody, including unauthenticated users (called *guest* users). To learn more about how Liferay DXP restricts access to portal resources to different users, please see the Roles and Permissions documentation.

## 415.1 Authentication Types

---

There are three authentication types: by email address, screen name, or user ID. To choose an authentication type, navigate to the Control Panel, click on *Configuration* → *Instance Settings* → *Platform* → *User Authentication* and use the *How do users authenticate?* selector to make a selection. Alternatively, add the following lines to your `portal-ext.properties` file, uncomment the appropriate line, comment out the others, and restart your server.

```
company.security.auth.type=emailAddress
#company.security.auth.type=screenName
#company.security.auth.type=userId
```

The default authentication type is by email address, but you can choose screen names or user IDs instead. Users choose screen names when they create their accounts or administrators can choose them. User IDs are auto-generated when the account is created. Regardless of which authentication type is configured, users must always enter a password. For information on adding restrictions on the kinds of passwords that are allowed or required (e.g., to require a minimum password length or require special characters), please see the Password Policies documentation.

## 415.2 The Sign In Portlet

---

Sign In portlet is how users log in. By default, the Sign In portlet can create new accounts or request a password reset. The default home page contains a Sign In portlet. You can access this page at <http://localhost:8080/web/guest/home>.

The image shows a user interface for signing in. At the top left, there is a blue arrow pointing left followed by the text "SIGN IN". Below this, there is a section titled "Email Address" with a light blue rounded rectangular input field containing the text "@liferay.com". Underneath the email field is a section titled "Password" with a light gray rounded rectangular input field. Below the password field is a checkbox labeled "Remember Me". At the bottom left, there is a blue button with the text "Sign In". At the bottom right, there are two links: "+ Create Account" and "? Forgot Password", both in blue text.

Figure 415.1: By default, the Sign In portlet allows users to log in, create a new account, or request a password reset.

If the Sign In portlet doesn't appear on any page, you can still access it here:  
[localhost:8080/c/portal/login](http://localhost:8080/c/portal/login)

By default, guest users can create accounts by clicking on the *Create Account* link in the Sign In portlet, completing the form, and submitting it. If a user has an account but has forgotten its

password, the user can click the *Forgot Password* link to request a password reset.

Both the *Create Account* form and the *Forgot Password* form include a CAPTCHA-based text verification field. Using CAPTCHA prevents bots from submitting these forms.

You can use reCAPTCHA instead of CAPTCHA. One advantage of reCAPTCHA is that it can allow visually impaired users to pass the test. To use reCAPTCHA, navigate to the Control Panel, then click on *Configuration* → *System Settings* → *CAPTCHA*.

You can prevent guest users from creating new user accounts, if your site requires users be registered by administrators. Navigate to the Control Panel → *Configuration* → *Instance Settings* → *Platform* → *User Authentication* and uncheck the *Allow strangers to create accounts?* box. You can also prevent users from requesting forgotten passwords or from requesting password reset links by unchecking the appropriate boxes. With these options, the Create Account and Forgot Password links no longer appear in the Sign-In portlet.

Remember that the Sign In portlet is the default way for users to log in, but it's not the only way. User accounts can be imported from and exported to LDAP directories. You can use single-sign-on (SSO) solutions or token-based authentication, which allows remote web applications to authenticate. Please refer to the other articles in this section for more information. Finally, remember that user authentication and remote application authentication mechanisms are extensible.



---

## SERVICE ACCESS POLICIES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

*Service access policies* comprise a layer of web service security that defines services or service methods that can be invoked remotely. You can apply many of them at once to produce a combined effect. To help you understand how service access policies fit into the big picture, here's a summary of Liferay DXP's web service security layers:

**IP permission layer:** The IP address from which a web service invocation request originates must be white-listed in the portal properties file. A web service invocation coming from a non-whitelisted IP address automatically fails.

**Service access policy layer:** Methods corresponding to a web service invocation request must be whitelisted by each service access policy that's in effect. You can use wildcards to reduce the number of service classes and methods that must be explicitly whitelisted.

**Authentication/verification layer (browser-only):** If a web service invocation request comes from a browser, the request must include an authentication token. This authentication token is the value of the `p_auth` URL parameter. The token is generated by Liferay DXP and associated with your browser session. The `p_auth` parameter is automatically supplied when you invoke a Liferay DXP web service via the JSON web services API page or via JavaScript using `Liferay.Service(...)`. If Liferay DXP cannot associate the caller's authentication token with a `User`, the web service invocation request fails.

**User permission layer:** Properly implemented web services have permission checks. The user invoking a web service must have permission to invoke the service.

Note that service access policies respect the permissions system. If a service access policy grants a user access to a remote service, the user must still have the appropriate permissions to invoke that service.

Service access policies are especially useful when remote applications such as mobile devices or Liferay Sync instances must access web services. Administrators can use service access policies to ensure that these devices can only invoke remote services from approved lists that can be modified at runtime.

---

### 416.1 Managing Service Access Policies

---

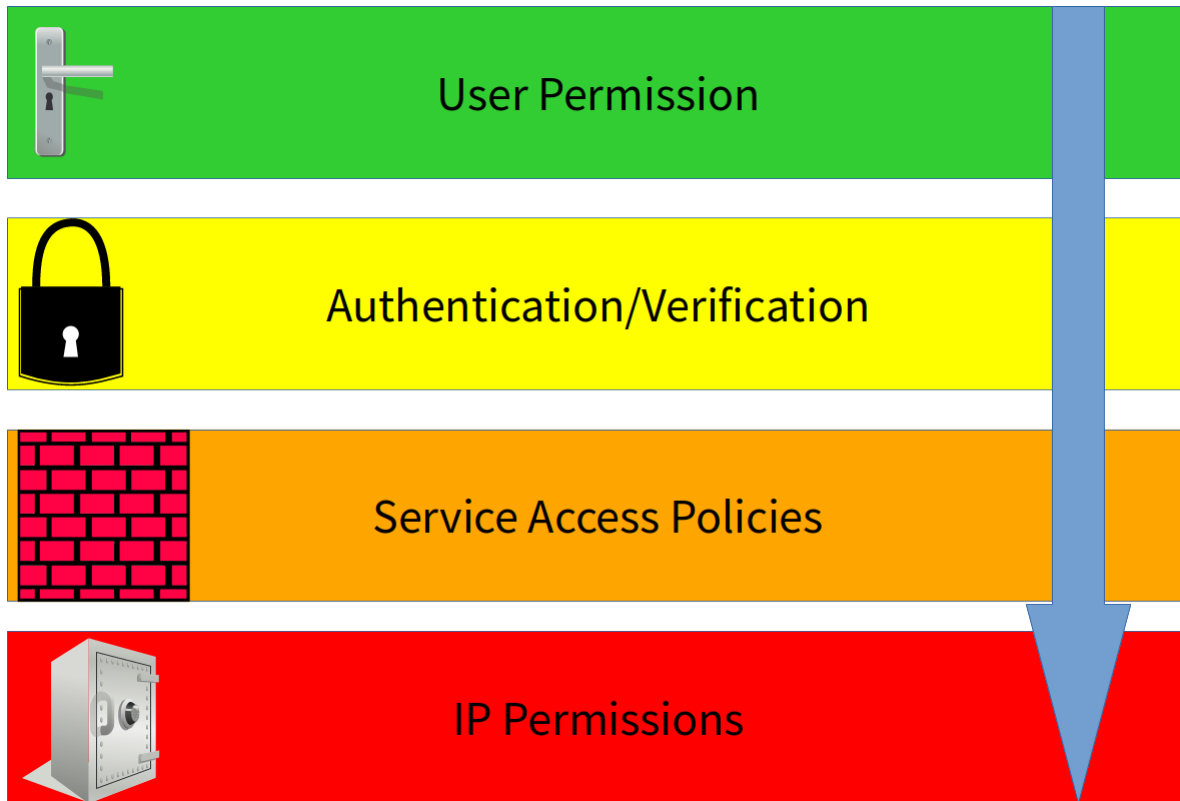


Figure 416.1: To get to a service, a request must pass through the door lock of user permissions, the padlock of the verification layer, the brick wall of service access policies, and finally the safe of predefined IP permissions.

Navigate to the Control Panel and click on *Service Access Policy* under the Configuration heading. Here, you can see the default service access policies and add new ones. When creating or editing service access policies, keep these points in mind:

- Service access policy names must be unique per portal instance.
- Service access policy names can include only these allowed characters:

```
0123456789ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz#@-./_
```

- Service access policy titles can be localized; service access policy names cannot be localized.
- Allowed service signatures must be entered one per line. Wildcards (\*) are allowed for both class names and method names. The # symbol must be used to separate a class name from a method name.

For example, `com.liferay.portal.kernel.service.UserService` allows any method from the `UserService` class to be invoked. `com.liferay.document.library.kernel.service.DLAppService#get*` allows any method from the `DLAppService` that starts with `get` to be invoked. Thus,

```
com.liferay.portal.kernel.service.UserService
com.liferay.document.library.kernel.service.DLAppService#get*
```

allows any method from the `UserService` class to be invoked and any method from the `DLAppService` whose name starts with `get` to be invoked.

There are 16 service access policies that are enabled by default. Six of these have to do with the system:

**ASSET\_ENTRY\_DEFAULT:** Allows the view counter for assets to be updated when an asset is retrieved.

**CALENDAR\_DEFAULT:** Makes it possible to search public events in the calendar.

**SYNC\_DEFAULT:** Allows only the `com.liferay.sync.service.SyncDLObjectService.getSyncContext` method. It applies to every Liferay Sync request, including unauthenticated Sync requests.

**SYNC\_TOKEN:** Allows `com.liferay.sync.service.*`, meaning that any API function that's a method of a class in this package can be invoked. It applies to Sync requests which are accompanied by an authentication token.

**SYSTEM\_DEFAULT:** Allows access to country/region services by JavaScript calls, so users can switch languages on the fly. Applies to every request, including unauthenticated requests.

**SYSTEM\_USER\_PASSWORD:** Allows any method to be invoked. Of course, since API functions include permission checks, this call works only if the user has the required permission. It applies to requests for which `AuthVerifierResult.isPasswordBasedAuthentication` is `true`: i.e., whenever user authentication took place using a password. If you want to completely disallow certain API functions from being invoked, you can change the `SYSTEM_USER_PASSWORD` policy to something more restrictive than `*`.

`SYNC_DEFAULT` and `SYSTEM_DEFAULT`, as their names suggest, are default service access policies. Default service access policies are applied to all incoming requests, including unauthenticated requests.

The other 10 policies have to do with OAuth and JSON web services:

**OAUTH2\_analytics.read/write:** Integrates with Liferay Analytics Cloud, allowing it access to JSON web services.

**OAUTH2\_everything/read/documents/userprofile/write:** The Everything policies grant access to all the JSON web services for various reasons. Everything is everything: all JSON web services (matches `*`). The others match method signatures appropriate to their description. For example, `OAUTH2_everything.read` matches all methods starting with `fetch`, `get`, `has`, `is`, or `search`.

**OAUTH\_READ/WRITE:** These provide access to JSON web services via the OAuth 1.0a plugin.

The default configuration makes available corresponding scopes that provide access to all web services shipped with the system. The scopes must be assigned to OAuth 1 or 2 applications before they become usable. Administrators should review the ones you want to use and disable the others.

You can create new default service access policies:

1. Navigate to the *Configuration* → *Service Access Policy* section of the Control Panel.
2. Click *Add* (+).
3. Give your policy a name.
4. Flip the switch to enable your policy.
5. If you want the policy applied to unauthenticated requests as well as authenticated requests, flip the switch labeled *Default*.
6. Give your policy a localized title.

7. Under Allowed Service Signatures, start typing the fully qualified name of a service class that's installed. Code completion helps you find the class. For example, if you're creating a policy for Liferay's Knowledge Base application, you could enter `com.liferay.knowledge.base.service.KBArticleService`.
8. Under Method Name, start typing a service method call. Again, code completion helps you. For Knowledge Base, you could enter `getKBArticle`.
9. To specify another service or method, click the plus icon to add another entry.
10. When done, click *Save*.

---

**Note:** If you know all the method signatures ahead of time, you can click *Switch to Advanced Mode* and enter them all in one field on separate lines.

---

Liferay applications can declare their own default policies (the `SYNC_DEFAULT` policy is a good example). This policy can then be changed or disabled by administrators. In this case, the plugin can still verify that the policy exists so there is no need to redefine or update it.

By default, Liferay's tunneling servlet uses the `SYSTEM_USER_PASSWORD` service access policy. You can, however, create your own policy for the tunneling servlet and use the property `service.access.policy.name` for the `TunnelingServletAuthVerifier` to specify that your policy should be used instead.

## 416.2 Service Access Policy Module

---

Liferay's service access policy functionality is provided by the Service Access Policy module. This module includes the following important classes:

- `com.liferay.portal.kernel.security.service.access.policy.ServiceAccessPolicy`: defines the public interface for `ServiceAccessPolicy`.
- `com.liferay.portal.kernel.security.service.access.policy.ServiceAccessPolicyManager`: defines the public interface for retrieving instances of `ServiceAccessPolicy`.
- `com.liferay.portal.kernel.security.service.access.policy.ServiceAccessPolicyManagerUtil`: bridges service access policy functionality to the parts of Liferay's core that have not yet been modularized.
- `com.liferay.portal.kernel.security.service.access.policy.ServiceAccessPolicyThreadLocal`: makes `ServiceAccessPolicy` instances active.

Liferay's Service Access Policy module resides in the `modules/apps/service-access-policy` folder in the source code. When running, these three bundles provide the service access policy functionality (they're in the `[Liferay Home]/osgi/modules` folder):

- `com.liferay.service.access.policy.api.jar`
- `com.liferay.service.access.policy.service.jar`
- `com.liferay.service.access.policy.web.jar`



These modules provide the service access policy management UI that's accessible from the Control Panel. They also provide the interface and default implementation for `ServiceAccessPolicy`.

To configure the Service Access Policy module, navigate to the Control Panel, click on *System Settings*, and find the *Service Access Policies* module in the Security section. Click on its name to edit it. Here, you can edit the default service access policy configuration. You can also force a default policy to be applied even when no policies are applied by the `AuthVerifier`.

There's also an `AuthenticatedAccessControlPolicy`. This policy doesn't do anything if a `ServiceAccessPolicyManager` implementation is present. If the service access policy module is disabled, however, the `AuthenticatedAccessControlPolicy` provides a fallback that still requires authenticated access for web services.

### **416.3 Summary**

---

Great! Now you know service access policies can restrict access to Liferay DXP's web services. Custom service access policies can be created by portal administrators. They are applied by the portal's token authenticator, e.g., by OAuth.

### **416.4 Related Topics**

---

Creating Service Access Policies



---

# AUTHENTICATION VERIFIERS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The Authentication Verification Layer is a centralized and extensible way to authenticate remote invocations of Liferay DXP's API.

The main responsibilities of the authentication verification layer are to

1. Verify provided credentials using registered AuthVerifier instances
2. Create portal authorization contexts based on verification results

If no available AuthVerifier can verify request credentials, an authorization context supporting non-authenticated access is created for a guest user. This allows each API to expose only a single API endpoint. In contrast, legacy (prior to 6.2) versions of Liferay DXP exposed two API endpoints for each API: the `/api/endpoint` URI was for non-authenticated access and the URI `/api/secure/endpoint` was for authenticated access.

There are built-in AuthVerifier implementations for the most common situations, such as when remote clients use HTTP Basic or HTTP Digest authentication, send credentials in request parameters, send authenticated JSESSIONIDs, or use shared secrets to establish trust. Other AuthVerifier implementations can be deployed as modules containing implementations of the AuthVerifier interface that are registered as services in the OSGi runtime.

Note: The authentication verification layer's focus is on verifying authentication, not on providing credentials. It does NOT issue tokens, credentials, or display Sign In portlets. Instead, the layer verifies existing credentials and authenticated sessions and is therefore a complement to authentication endpoints. To ensure backwards compatibility, however, the default implementations support requests providing user name and password credentials. Thus, the authentication verification layer stands on the border between authentication and authorization.

---

## 417.1 Authentication Verification Process Overview

---

This layer and surrounding processes are provided by the AuthVerifierFilter class that implements the `javax.servlet.Filter` interface.

**Step 1: Verify Request Credentials**

The layer uses the chain of responsibility design pattern to support both built-in and third party AuthVerifier implementations. Each AuthVerifier can provide configurations where it specifies mapped URLs and other properties.

Each incoming request is matched against all registered AuthVerifiers to select the final list of AuthVerifiers that is used to process the request. It's the responsibility of each AuthVerifier to verify the incoming request credentials.

**Step 2: Create an Authorization Context**

When a request is processed by all matching AuthVerifiers, Liferay DXP creates an authorization context for the resolved user.

This encompasses setting the `HttpServletRequest` `remoteUser` to return the resolved user ID setting `ThreadLocals` to the resolved user.

The resolved user can be the user returned by one of the AuthVerifier instances or a guest user if no instance was able to verify the provided credentials. AuthVerifiers are created by developers, and are processed automatically as long as they're registered in the OSGi runtime. Each Auth Verifier gets its own configuration in *Control Panel* → *System Settings* → *Security* → *API Authentication*. Configuration for Auth Verifiers that ship with the product include

- Basic Auth Header
- Digest Authentication
- HTTP Tunnel Extender
- Image Request
- Portal Sessions
- Request Parameter
- Tunnel Auth

The following Auth Verifiers are enabled by default and can be used to access remote API out-of-the-box:

- Basic Auth Header
- Portal Sessions

---

**417.2 Basic Auth Header**

This Auth Verifier allows the remote client to authenticate using HTTP Basic Authentication. Configure it by providing URL paths that should be authenticated this way. When Force Basic Authentication field is checked then HTTP Basic Authentication is required.

The default URLs are `/api/*`, `/xmlrpc*` for web services. The mapping excludes `/api/liferay*` to prevent accessing `TunnelServlet`. For more information please see Tunnel Authentication Verifiers.

---

**417.3 Digest Auth Header**

This Auth Verifier allows the remote client to authenticate using HTTP Digest Authentication. Configure it by providing URL paths that should be authenticated this way. When Force Digest Authentication field is checked then HTTP Basic Authentication is required.

This Auth Verifier is not enabled by default.

#### **417.4 HTTP Tunnel Extender**

---

As Liferay embraced modularity, this extender was written to enable modules to be part of TunnelServlet. It maps TunnelServlet and TunnelingServletAuthVerifier to the module servlet context. Modules with Http-Tunnel in the manifest can make use of the Tunnel Servlet, and can expose the API via `/o/_module_/api/liferay/do`.

Configure it by setting client IP addresses allowed to tunnel. For more information, please see the properties documentation as well as remote staging.

Note that this is not a recommended way to export remote APIs; it's far better to expose remote services using JAX-RS or Liferay JSON Web Service technologies.

#### **417.5 Image Request Authentication Verifier**

---

When connected to LibreOffice/OpenOffice, the Office process must download images from Liferay DXP to render docs with images. To do this, a JWT Token is created to access the images securely.

Configure this by setting the Hosts Allowed, URLs included, and URLs excluded if necessary.

This Auth Verifier is not enabled by default.

#### **417.6 Portal Sessions Auth Verifiers**

---

Enables JavaScript in a browser to access Liferay JSON Web Services using an existing portal session.

In the default configuration, the URLs included field shields access to the legacy JSON remote services layer: `/api/json*`, `/api/jsonws*`, `/c/portal/json_service*`.

#### **417.7 Request Parameter Auth Verifiers**

---

For backwards compatibility with `RequestParamerAutoLogin` you can authenticate and access portal endpoints with credentials inside HTTP request parameters `parameterAutoLoginLogin` and `parameterAutoLoginPassword`.

This Auth Verifier is not enabled by default.

### **417.8 Tunnel Authentication Verifiers**

---

`TunnelServlet` is a legacy remote API endpoint mapped at `/api/liferay/do` to provide access to the portal remote services. The Tunnel Auth Verifier allows trusted remote clients authenticated access using any user ID provided, on behalf of the user.

An example of a trusted remote client is the Staging remote publishing feature.

Trusted remote clients authenticate using a shared secret stored in the portal property `tunneling.servlet.shared.secret`. The default value is empty and forbids all access.

Even though the default configuration is enabled by default, access is limited to localhost only. Configure it by setting client IP addresses allowed to tunnel. For more information, please see the properties documentation as well as remote staging.

### **417.9 Related Topics**

---

Service Access Policies

---

# LDAP

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

LDAP is a common user store for Liferay DXP. You can configure LDAP at the system scope in System Settings or at the instance scope in Instance settings. Users can be imported from LDAP or exported to LDAP. To access LDAP configuration settings, navigate to *Control Panel* ☒ *Configuration* → *Instance Settings*. At the bottom of the list on the left, click *Servers*.

Click the *Add* button to add an LDAP server connection. If you have more than one LDAP server, you can arrange the servers by order of preference using the up/down arrows. Regardless of how many LDAP servers you add, each server has the same configuration options.

**Server Name:** Enter a name for your LDAP server.

**Default Values:** Several common directory servers appear here. If you use one of these, select it. The rest of the form is populated with default values for that directory.

---

## 418.1 Connection

---

These settings cover the connection to LDAP.

**Base Provider URL:** The link to the LDAP server. Make sure the Liferay DXP server can communicate with the LDAP server. If there is a firewall between the two systems, make sure the appropriate ports are opened.

**Base DN:** The Base Distinguished Name for your LDAP directory. It is usually modeled after your organization. It may look similar to this: `dc=companynamehere,dc=com`.

**Principal:** The default LDAP administrator user ID is populated here. If your administrator ID differs, use that credential instead. You need an administrative credential because Liferay DXP uses this ID to synchronize user accounts to and from LDAP.

**Credentials:** This is the password for the LDAP administrative user.

This is all you need to make a regular connection to an LDAP directory. The rest of the configuration, however, may need to be customized, as it represents “best guesses” as to correct defaults. The default attribute mappings usually provide enough data to synchronize back to the Liferay

DXP database when a user attempts to log in. To test the connection to your LDAP server, click the *Test LDAP Connection* button.

## 418.2 Checkpoint

---

Before proceeding to fine tune Liferay DXP's LDAP connections, ensure the following steps have been taken:

1. The LDAP connection is enabled. Depending on your needs, LDAP authentication may be required so that only users who have been bound may log in.
2. *Export/Import*: for users in a clustered environment, Enable Import/Export on Startup should be disabled so that there are no massive imports on every node upon start up.
3. When adding the LDAP server, the *Server Name*, *Default Values*, *Connection* values are correct. It is always a good idea to click the *Test LDAP Connection* before saving.

## 418.3 Instance Settings vs. System Settings

---

You can define an LDAP server connection at the System Settings scope as well. Because this user interface is auto-generated, it's not as helpful as the one in Instance Settings. For this reason, you should define and troubleshoot your settings in Instance Settings first. If you decide you want your LDAP connection at the system scope, you can copy your configuration from Instance Settings and then delete the server from Instance Settings.

Of course, you can also configure LDAP servers at the system scope using OSGi `.config` files. The easiest way to do this is to use the GUI and export the configuration. Then you can use the resulting `.config` file anywhere you need it (such as other nodes in a cluster).

**Note:** To use config files for LDAP server configuration, you must specify the Virtual Instance ID (in the source, the variable name is `companyId`) in the exported configuration file, because servers are defined at the instance scope, not the system scope. To do this, specify the virtual instance ID somewhere in the file like this:

```
companyId=1234
```

You can find your Virtual Instance ID in Control Panel → Configuration → Virtual Instances.

## 418.4 Security

---



If you run your LDAP directory in SSL mode to encrypt credential information on the network, you must perform extra steps to share the encryption key and certificate between the two systems.

For example, if your LDAP directory is Microsoft Active Directory on Windows Server 2003, you'd share the certificate like this:

Click *Start* → *Administrative Tools* → *Certificate Authority*. Highlight the machine that is the certificate authority, right-click on it, and click *Properties*. From the General menu, click *View Certificate*. Select the Details view, and click *Copy To File*. Use the resulting wizard to save the certificate as a file. You must also import the certificate into the *cacerts keystore* like this:

```
keytool -import -trustcacerts -keystore /some/path/java-8-jdk/jre/lib/security/cacerts -storepass changeit -noprompt -
alias MyRootCA -file /some/path/MyRootCA.cer
```

The `keytool` utility ships as part of the Java SDK.

Once this is done, go back to the LDAP page in the Control Panel. Modify the LDAP URL in the Base DN field to the secure version by changing the protocol to `ldaps` and the port to `636` like this:

```
ldaps://myLdapServerHostname:636
```

Save the changes. Communication to LDAP is now encrypted.

## 418.5 Configuring LDAP Import/Export

---

The other settings configure mappings between LDAP and Liferay DXP so users and groups can be imported.

## 418.6 Users

---

This section contains settings for finding users in your LDAP directory.

**Authentication Search Filter:** Use this search filter box to determine the search criteria for user logins. By default, Liferay DXP uses users' email addresses for their login names. The value here must use the authentication method you use. For example, if you changed Liferay DXP's authentication method to use screen names instead of the email addresses, you would modify the search filter so it can match the entered log in name:

```
(cn=@screen_name@)
```

**Import Search Filter:** Depending on the LDAP schema, there are different ways to identify the user. The default setting is usually fine:

```
(objectClass=inetOrgPerson)
```

If you want to search for only a subset of users or users that have different LDAP object classes, you can change this.

**User Mapping:** Next, you can define mappings from LDAP attributes to Liferay fields. Though LDAP user attributes may be different from LDAP server to LDAP server, there are five fields Liferay DXP requires to be mapped for the user to be recognized:

- *Screen Name* (e.g., uid or cn)
- *Password* (e.g., userPassword)
- *Email Address* (e.g., mail or email)
- *First Name* (e.g., name or givenName)
- *Last Name* (e.g., sn)

**Note:** If you intend to create or import users with no email addresses, you must set `users.email.address.required=false` in `portal-ext.properties`. With this set, Liferay auto-generates an email address combining the user ID plus the suffix defined in the property `users.email.address.auto.suffix=`. Finally, make sure to set Liferay and LDAP authentication to something other than email address.

If you want to import LDAP groups as Liferay DXP user groups, make sure define a mapping for the Liferay DXP group field so that membership information is preserved:

- *Group* (e.g., *member*)

The other LDAP user mapping fields are optional.

The Control Panel provides default mappings for commonly used LDAP attributes. You can also add your own mappings.

**Test LDAP Users:** Once you have your attribute mappings set up (see above), click the *Test LDAP Users* button and Liferay DXP attempts to pull LDAP users and match them with their mappings as a preview.

## 418.7 Groups

---

This section contains settings for mapping LDAP groups to Liferay DXP user groups.

**Import Search Filter:** This is the filter for mapping LDAP groups to Liferay DXP user groups. For example,

```
(objectClass=groupOfNames)
```

Enter the LDAP group attributes you want retrieved for this mapping. The following attributes can be mapped. The *Group Name* and *User* fields are required, the *Description* is optional.

- *Group Name* (e.g., cn or o)
- *Description* (e.g., description)
- *User* (e.g., member)

**Test LDAP Groups:** Click the *Test LDAP Groups* button to display a list of the groups returned by your search filter.

**LDAP** ×

---

Test LDAP Users

A subset of users has been displayed for you to review.

| screenName | emailAddress                | First Name | Last Name   | Password | Job Title         | Group |
|------------|-----------------------------|------------|-------------|----------|-------------------|-------|
| captain    | spectacular@lunarresort.com | Captain    | Spectacular | *****    | Shuttle Captain   | 0     |
| ziltoid    | ziltoid@lunarresort.com     | Ziltoid    | Omniscient  | *****    | Coffee Reviewer   | 0     |
| james      | james@lunarresort.com       | James      | Jeffries    | *****    | Lunar Associate   | 0     |
| marvin     | marvin@lunarresort.com      | Marvin     | Smart       | *****    | Lunar Associate   | 0     |
| melvin     | melvin@lunarresort.com      | Melvin     | Dooitrong   | *****    | Intern            | 0     |
| rex        | rex@lunarresort.com         | Rex        | Nihilo      | *****    | Upsale Manager    | 0     |
| jane       | jane@lunarresort.com        | Jane       | Slaughter   | *****    | Chief of Security | 0     |

Page 1 of 1 ▾
20 Items per Page ▾
Showing 7 results.

[← First](#)
[Previous](#)
[Next](#)
[Last →](#)

Figure 418.1: You should see a list of users when you click the *Test LDAP Users* button.

## 418.8 Export

---

This section contains settings for exporting user data from LDAP.

**Users DN:** Enter the location in your LDAP tree where the users are stored. Liferay DXP exports the users to this location.

**User Default Object Classes:** Users are exported with the listed default object classes. To find out what your default object classes are, use an LDAP browser tool such as Apache Directory Studio to locate a user and view the Object Class attributes stored in LDAP for that user.

**Groups DN:** Enter the location in your LDAP tree where the groups are stored. When Liferay DXP does an export, it exports the groups to this location.

**Group Default Object Classes:** When a group is exported, the group is created with the listed default object classes. To find out what your default object classes are, use an LDAP browser tool such as Apache Directory Studio to locate a group and view the Object Class attributes stored in LDAP for that group.

When you've set all your options and tested your connection, click *Save*.

---

**Note:** If a user changes a value like a password in Liferay DXP, that change is passed to the LDAP server, provided Liferay DXP has enough schema access to make the change.

---

Now you know how to connect an LDAP server to Liferay DXP and how to configure user import behavior, export behavior, and other LDAP settings.

**418.9 Related Topics**

---

Liferay DXP Security Overview Logging into Liferay DXP

---

# CONFIGURING LDAP

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

In this article, you'll learn how to configure import settings, export settings, and related LDAP configuration settings.

To access LDAP configuration settings, navigate to *Control Panel*  $\boxtimes$  *Configuration*  $\rightarrow$  *Instance Settings*  $\rightarrow$  *Security*  $\rightarrow$  *LDAP*. There are four categories on the left: Export, General, Import, and Servers. The Servers category was covered in the last article. The rest are covered below.

---

## 419.1 Export

---

**Enable Export:** Check this box to export user accounts to LDAP. A listener tracks changes made to the User object and pushes updates to the LDAP server whenever a User object is modified. Note that by default on every login, fields such as lastLoginDate are updated. When export is enabled, this causes a user export every time the user logs in. You can prevent updates to users' lastLoginDate fields from triggering LDAP user exports by setting the following property in your portal-ext.properties file:

```
users.update.last.login=false
```

**Enable Group Export:** Export groups to LDAP.

---

## 419.2 General

---

**Enabled:** Check this box to enable LDAP Authentication.

**Required:** Check this box if LDAP authentication is required. Users can't log in unless they can bind to the LDAP directory successfully. Uncheck this box if users with Liferay DXP accounts but no LDAP accounts can log in.

**Use LDAP Password Policy:** Liferay DXP uses its own password policy by default. This can be configured on the Control Panel's Password Policies page. Check the *Use LDAP Password Policy* box if you want to use the password policies defined by your LDAP directory. Once this is enabled, the Password Policies tab states that you are not using a local password policy. You must now use your LDAP directory's mechanism for setting password policies. Liferay DXP cannot enforce these policies; the best it can do is pass through the messages returned by your LDAP server. It does this by parsing the messages in the LDAP controls the server returns. By default, Liferay DXP is configured to parse the messages returned by the Fedora Directory Server. If you use a different LDAP server, you must customize the messages in *System Settings* → *Security* → *LDAP* → *Connection*.

**Method:** Choose *Bind* (the default) or *Password Compare*. Bind does a standard LDAP bind; Password Compare attempts to compare Liferay and LDAP passwords using the encryption algorithm specified in the field below. Password Compare is rarely used.

**Password Encryption Algorithm:** Choose the password encryption algorithm your LDAP server uses to encrypt passwords so they can be compared if using the Password Compare bind method. This is rarely used.

### 419.3 Import

---

You can import user data from LDAP directories using the following options:

**Enable Import:** Check this box to do a mass import from your LDAP directories. Otherwise, Users are imported as they log in.

| Name                                        | Screen Name | Job Title       | Organizations | User Groups |
|---------------------------------------------|-------------|-----------------|---------------|-------------|
| <input type="checkbox"/> Rex Nihilo         | rex         | Upsale Manager  |               | ⋮           |
| <input type="checkbox"/> Ziltoid Omniscient | ziltoid     | Coffee Reviewer |               | ⋮           |
| <input type="checkbox"/> Test Test          | test        |                 |               | ⋮           |

Figure 419.1: Ziltoid and Rex have been imported because they logged in.

**Enable Import on Startup:** Check this box to do the mass import when Liferay DXP starts. Note: this box only appears if you check **Enable Import**, described above. Definitely leave this unchecked if you have a Liferay DXP cluster, or all your nodes will do a mass import when each of them starts up.

**Import Interval:** When mass importing users, import users every X minutes.

**Import Method:** Set either User or Group. If you set this to User, Liferay DXP imports all users from the location specified in the server connection. If you set this to Group, Liferay DXP searches all the groups and imports the users in each group. If you have users who do not belong to any groups, they are not imported.

**Lock Expiration Time:** Set the account lock expiration time for LDAP User import. The default is one day.

**Import User Sync Strategy:** Set the strategy used to sync user accounts. Options are Auth Type (i.e., the way the user authenticates, like with screen name) and UUID (requires a UUID attribute in LDAP).

**Enable User Password on Import:** Assign a default password (see below) when users are imported, so they can be synced between the two systems.

**Autogenerate User Password on Import:** Create a random password on user import.

**Default User Password:** Enter the default password users are assigned when they first log in via LDAP.

**Enable Group Cache on Import:** Cache the imported groups so import isn't slowed by database access.

**Create Role per Group on Import:** For every LDAP group, create a corresponding Liferay Role.

---

## 419.4 Servers

**LDAP Servers:** Liferay DXP supports connections to multiple LDAP servers. Use the *Add* button to add LDAP servers. Each LDAP server has several configuration options explained here.

Once you've finished configuring LDAP, click the *Save* button.

---

## 419.5 LDAP Options Available in System Settings

Although most LDAP configuration can be done from Instance Settings, there are several configuration parameters that are only available in System Settings. There are also settings duplicated from the ones in Instance Settings. These change the *default* settings for new virtual instances (see note below).

---

**Note:** When you make a change in System Settings, it takes effect for the virtual instance you're in. If after changing a setting you create a new virtual instance, that virtual instance inherits the settings of the one it was created from as defaults. For example, say you have virtual instances named A, B, and C. From A, you modify *Error password history keywords*. This change appears only in A, not in B or C. Then from A, you create virtual instance D. The change to *Error password history keywords* appears in D (not B or C), since D defaults to A's settings because you created it from A.

---

If you must change any of these options, navigate to *Control Panel* → *Configuration* → *System Settings*. Go to the *Security* section and find the entries with LDAP in the title. The only new settings here are in the *Connection* entry.

Use the *Connection* entry to manage error properties like *Error password age keywords* which lets you set a list of phrases from error messages which can possibly be returned by the LDAP server. When a user binds to LDAP, the server returns *controls* with its response of success or failure.

These controls contain a message describing the error or the information that is returned with the response. Though the controls are the same across LDAP servers, the messages can be different. The properties described here contain snippets of words from those messages and work with Red Hat's Fedora Directory Server. If you are not using that server, the word snippets may not work with your LDAP server. If they don't, you can replace the values of these properties with phrases from your server's error messages. This enables Liferay DXP to recognize them.



---

# TOKEN-BASED SINGLE SIGN ON AUTHENTICATION

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Token-based SSO authentication was introduced in Liferay DXP 7.0 to standardize support for Shibboleth, SiteMinder, Oracle OAM, or any other SSO product that works by propagating a token via one of the following mechanisms:

- HTTP request parameter
- HTTP request header
- HTTP cookie
- Session attribute

Since these providers have a built-in web server module, you should use the Token SSO configuration.

The authentication token contains the Liferay DXP user's screen name or email address, whichever Liferay DXP has been configured to use for the particular company (portal instance). Recall that Liferay DXP supports three authentication methods:

- By email address
- By screen name
- By user ID

Token-based authentication only supports email address and screen name. If Liferay DXP is configured to use user ID when a token-based authentication is attempted, the `TokenAutoLogin` class logs this warning:

```
Incompatible setting for: company.security.auth.type
```

Please note that the above sources are fully trusted.

Furthermore, you must use a security mechanism external to Liferay DXP, such as a fronting web server like Apache. The chosen fronting solution must prevent malicious Liferay DXP user impersonation that otherwise might be possible by sending HTTP requests directly to Liferay DXP from the client's web browser.

Token-based authentication is disabled by default. To manage token-based SSO authentication, navigate to Control Panel → *System Settings*, → *Security* → *SSO*. Token Based SSO appears in Virtual Instance Scope at the bottom. Here are the configuration options for the Token Based SSO module:

**Enabled:** Check this box to enable token-based SSO authentication.

**Import from LDAP:** Check this box to import users automatically from LDAP if they don't exist.

**User token name:** Set equal to the name of the token. This is retrieved from the specified location. (Example: SM\_USER)

**Token location:** Set this to the location of the user token. As mentioned earlier, the options are:

- HTTP request parameter
- HTTP request header
- HTTP cookie
- Session attribute

**Authentication cookies:** Set this to the cookie names that must be removed after logout. (Example: SMIDENTITY, SMSESSION)

**Logout redirect URL:** When user logs out of Liferay DXP, the user is redirected to this URL. Remember to click *Save* to activate Token Based SSO.

### 420.1 Required SiteMinder Configuration

---

If you use SiteMinder, note that Liferay DXP sometimes uses the tilde character in its URLs. By default, SiteMinder treats the tilde character (and others) as bad characters and returns an HTTP 500 error if it processes a URL containing any of them. To avoid this issue, change this default setting in the SiteMinder configuration to this one:

```
BadUrlChars //,./,/,/*,*,\,%00-%1f,%7f-%ff,%25
```

The configuration above is the same as the default except the ~ was removed from the bad URL character list. Restart SiteMinder to make your configuration update take effect. For more information, please refer to SiteMinder's documentation

### 420.2 Summary

---

Liferay DXP's token-based SSO authentication mechanism is highly flexible and compatible with any SSO solution that provides it with a valid Liferay DXP user's screen name or email address. These include Shibboleth and SiteMinder.

---

# AUTHENTICATING WITH OPENID CONNECT

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

OpenID Connect is a lightweight authentication layer built on top of the OAuth 2.0 authorization protocol. It compliments local accounts by enabling users to authenticate using accounts they have on other systems. Users who avoid signing up for new accounts can then use an account they already have to sign into your website. By using OpenID Connect, you *delegate* user authentication to other providers, making it easy for users with existing accounts to authenticate to your system.

---

**Note:** You can add multiple providers to your installation, but Liferay DXP can't be an OpenID Connect provider.

---

Because OpenID Connect is built on OAuth 2.0, its token flow is similar. OAuth 2.0 is only an authorization protocol, so it sends an *access token* that grants access to particular APIs. OpenID Connect adds to this an *identity token* that passes user information like name and email, provided the user has authenticated and granted permission.

---

## 421.1 Creating a Client in OpenID Connect Provider

---

To use OpenID Connect, you must first register it as a client in your provider. This is an OAuth 2.0 client. The process varies by provider:

1. Navigate to the provider's website and create a client.
2. During the creation process, you must supply an *authorized redirect URL* that can process the tokens sent from the provider. Liferay DXP's URL is

`https://[server.domain]/c/portal/login/openidconnect`

3. The provider sends several pieces of information. Some of these, like the Discovery Endpoint, Authorization Endpoint, or Issuer URL are the same regardless of the client. The two pieces of information unique to your request are the `client_id` and the `client_secret`.

Collect the information from the provider. You'll need it create the provider next.

## 421.2 Configuring an OpenID Connect Provider Connection

---

Go to *Control Panel* → *Configuration* → *System Settings* → *Security* → *SSO* and select **OpenID Connect Provider** under the *System Scope* and follow these steps:

1. Add the provider by clicking the *Add* button.
2. Use the information you received from the provider to fill out the form:

**Provider Name:** This name appears in the Sign-In Portlet when users use OpenID Connect to log in.

**OpenID Client ID:** Provide the OAuth 2.0 Client ID you received from your provider.

**OpenID Connect Client Secret:** Provide the OAuth 2.0 Client Secret you received from your provider.

**Scopes:** Leave the default, which requests the user name and the email. Your provider may offer other scopes of user information.

**Discovery Endpoint:** Other URLs may be obtained from this URL, and they vary by provider.

**Discovery Endpoint Cache in Milliseconds:** Cache the endpoints (URLs) discovered for this amount of time.

**Authorization Endpoint:** This URL points to the provider's URL for authorizing the user (i.e., signing the user in).

**Issuer URL:** The provider's URL that points to information about the provider who is issuing the user information.

**JWKS URI:** A URL that points to the provider's JSON Web Key Set that contains the public keys that can verify the provider's tokens.

**ID Token Signing Algorithms:** Set the supported ID token algorithms manually. Normally, this is "discovered" at the discovery endpoint. You can add as many of these as you need.

**Subject Types:** A Subject Identifier is a unique and never reassigned identifier the provider uses to establish who the user is, and is consumed by the client (i.e., Liferay DXP). There are two types: public (provides the same value to all clients) and private (provides a different value to each client).

**Token Endpoint:** The provider's URL where tokens can be requested.

**User Information Endpoint:** The OAuth 2.0 protected URL from which user information can be obtained.

Once you've filled out the form, click *Save*, and you're ready to enable OpenID Connect authentication.

---

System Settings configuration file:

```
com.liferay.portal.security.sso.openid.connect.internal.configuration.OpenIdConnectProviderConfiguration-[name].config
```

where [name] is a descriptive, but unique name for example provider1.

---

### 421.3 Enabling OpenID Connect Authentication

---

1. Go to *Control Panel* → *Configuration* → *System Settings* → *Security* → *SSO* and select **OpenID Connect** under *Virtual Instance Scope*.
2. Click the *Enabled* check box, and then click *Save*.

**Note:** You can also enable OpenID Connect authentication for the given virtual instance through the *Control Panel* → *Configuration* → *Instance Settings* → *OpenID Connect* tab.

---

System Settings configuration file:

```
com.liferay.portal.security.sso.openid.connect.configuration.OpenIdConnectConfiguration.config
```

---

Now users can sign in with OpenID Connect.

### 421.4 Signing In With OpenID Connect

---

There's a new link in the Sign-In Portlet for signing in with OpenID Connect:

1. From the Sign-In Portlet, click the OpenID Connect link at the bottom.
2. Choose a provider and click *Sign In*.
3. This takes you to your provider's sign in page. Enter your credentials and log in.
4. Upon successful authentication, you're redirected back to Liferay DXP in an authenticated state.

OpenID is a standards-based, secure way to authenticate users from other systems.



---

## OPENAM SINGLE SIGN ON AUTHENTICATION

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

OpenAM is an open source single sign-on solution that comes from the code base of Sun's System Access Manager product. Liferay DXP integrates with OpenAM, allowing you to use OpenAM to integrate Liferay DXP into an infrastructure that contains a multitude of different authentication schemes against different repositories of identities.

Note that OpenAM relies on cookie sharing between applications. Thus, in order for OpenAM to work, **all applications that require SSO must be in the same web domain**. You should also add the following property if you have enabled HTTPOnly cookies due to the way some web containers (like Apache Tomcat™) parse cookies that contain special characters:

```
com.iplanet.am.cookie.encode=true
```

You can install OpenAM on the same or different server as Liferay DXP. Be sure to review the context path and server hostname for your OpenAM server.

If you want to install OpenAM on the same server as Liferay DXP, you must deploy the OpenAM .war, downloadable from here. Otherwise, follow the instructions at the OpenAM 13 site to install OpenAM.

---

**Note:** OpenAM 12 and below work with Liferay DXP, but are at end of life. Because of this, we recommend only OpenAM 13 for production use.

---

Once you have it installed, create the Liferay DXP administrative user in it. Users are mapped back and forth by screen names. By default, the Liferay DXP administrative user has a screen name of *test*, so if you were to use that account, register the user in OpenAM with the ID of *test* and the email address specified in the `admin.email.from.address` portal property). Once you have the user set up, log in to OpenAM using this user.

In the same browser window, log in to Liferay DXP as the administrative user (using the previous admin email address). Go to the Control Panel and click *Configuration* → *Instance Settings* → *Security* → *SSO*. Then choose *OpenSSO* in the list on the left.

Modify the three URL fields (Login URL, Logout URL, and Service URL) so they point to your OpenAM server (in other words, only modify the host name portion of the URLs), check the *Enabled* box, and click *Save*. Liferay DXP then redirects users to OpenAM when they request the `/c/portal/login` URL for example, when they click on the Sign In\* link).

**Enabled**

**Version**

Openam 13

**Import from LDAP**

Set this to true if you want to import the users from LDAP after successful login through OpenSSO. The LDAP settings need to be configured properly in the LDAP section. If this is set to false, the users will be created from OpenSSO provided data.

**Login URL**

```
http://openssohost.example.com:8080/opensso/UI/Login?goto=http://portalhost.example.com:8080/c/portal/login
```

**Logout on Session Expiration**

Set this to true to log a user out of OpenSSO when the portal session expires.

**Logout URL**

```
http://openssohost.example.com:8080/opensso/UI/Logout?goto=http://portalhost.example.com:8080/web/guest/home
```

**Service URL**

---

Figure 422.1: OpenSSO Configuration.



Liferay DXP's OpenAM configuration can be applied at either the system scope or at the instance scope. To configure the OpenAM SSO module at the system scope, navigate to the Control Panel, click on *Configuration* → *System Settings* → *Security* → *SSO* → *OpenSSO*. Click on it and you'll find these settings to configure. The values configured here provide the default values for all portal instances. Enter them in the same format as you would when initializing a Java primitive type with a literal value.

| Property Label   | Property Key     | Description                                                                                                                                                                     | Type    | Version                        |
|------------------|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------|--------------------------------|
| version          | version          | OpenAM version to use (12 and below or 13)                                                                                                                                      | String  | <b>Enabled</b>                 |
| enabled          | enabled          | Check this box to enable OpenAM authentication. Note that OpenAM works only if LDAP authentication is also enabled and Liferay DXP's authentication type is set to screen name. | boolean | <b>Import from LDAP</b>        |
| importFromLDAP   | importFromLDAP   | If this is checked, users authenticated from OpenAM that do not exist in Liferay DXP are imported from LDAP. LDAP must be enabled.                                              | boolean | <b>Login URL</b>               |
| loginURL         | loginURL         | The URL to the login page of the OpenAM server                                                                                                                                  | String  | <b>Logout URL</b>              |
| logoutURL        | logoutURL        | The URL to the logout page of the OpenAM server                                                                                                                                 | String  | <b>Service URL</b>             |
| serviceURL       | serviceURL       | The URL by which OpenAM can be accessed to use the authenticated web services. If you are using OpenAM Express 8 or higher, you need to have the server running Java 6.         | String  | <b>Screen Name Attribute</b>   |
| screenNameAttr   | screenNameAttr   | The name of the attribute on the OpenAM representing the user's screen name                                                                                                     | String  | <b>Email Address Attribute</b> |
| emailAddressAttr | emailAddressAttr | The name of the attribute on the OpenAM representing the user's email address                                                                                                   | String  | <b>First Name Attribute</b>    |
| firstNameAttr    | firstNameAttr    | The name of the attribute on the OpenAM representing the user's first name                                                                                                      | String  | <b>Last Name Attribute</b>     |
| lastNameAttr     | lastNameAttr     | The name of the attribute on the OpenAM representing the user's last name                                                                                                       | String  |                                |

To override these default settings for a particular portal instance, navigate to the Control Panel and click *Configuration* → *Instance Settings* → *Security* → *SSO*. Then choose *OpenSSO* in the list on the left.



---

# CAS (CENTRAL AUTHENTICATION SERVICE) SINGLE SIGN ON AUTHENTICATION

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

CAS is an authentication system originally created at Yale University. It is a widely used open source single sign-on solution and was the first SSO product to be supported by Liferay DXP. @product@'s CAS module includes the CAS client, so there's no need to install it separately.

---

**Note:** Liferay DXP supports CAS 3.3.x. If you use a later version of CAS, it is best to use CAS's support for standards such as OpenID Connect or SAML to interface with Liferay DXP.

---

The CAS Server application requires your server to have a properly configured Secure Socket Layer (SSL) certificate. To generate one yourself, use the `keytool` utility that comes with the JDK. First generate the key, then export the key into a file. Finally, import the key into your local Java key store. For public, Internet-based production environments, you must purchase a signed key from a recognized certificate authority or have your key signed by a recognized certificate authority. For Intranets, you should have your IT department pre-configure users' browsers to accept the certificate so they don't get warning messages about the certificate.

To generate a key, use the following command:

```
keytool -genkey -alias tomcat -keypass changeit -keyalg RSA
```

Instead of the password in the example (`changeit`), use a password you can remember. If you are not using Tomcat, you may want to use a different alias as well. For first and last names, enter `localhost` or the host name of your server. It cannot be an IP address.

To export the key to a file, use the following command:

```
keytool -export -alias tomcat -keypass changeit -file server.cert
```

Finally, to import the key into your Java key store, use the following command:

```
keytool -import -alias tomcat -file server.cert -keypass changeit -keystore $JAVA_HOME/jre/lib/security/cacerts
```

If you are on a Windows system, replace `$JAVA_HOME` above with `%JAVA_HOME%`. Of course, all of this must be done on the system where CAS is running.

Once your CAS server is up and running, configure Liferay DXP to use it. CAS configuration can be applied either at the system scope or at the scope of a portal instance. To configure the CAS SSO module at the system or instance scope, navigate to the Control Panel, click on *Configuration* → *System Settings* (or *Instance Settings*) → *Security* → *SSO*. The values configured in System Settings provide the default values for all portal instances. Enable CAS authentication and then modify the URL properties to point to your CAS server.

**Enabled:** Check this box to enable CAS single sign-on.

**Import from LDAP:** A user may be authenticated from CAS and not yet exist in Liferay DXP. Select this to automatically import users from LDAP if they do not exist in Liferay DXP. For this to work, LDAP must be enabled.

The rest of the settings are various URLs with defaults included. Change *localhost* in the default values to point to your CAS server. When you are finished, click *Save*. After this, when users click the *Sign In* link, they are directed to the CAS server to sign in to Liferay DXP.

For some situations, it might be more convenient to specify the system configuration via files on the disk. To do so, create the following file:

```
{LIFERAY_HOME}/osgi/configs/com.liferay.portal.security.sso.cas.configuration.CASConfiguration.cfg
```

The format of this file is the same as any properties file. The key to use for each property that can be configured is shown below. Enter values in the same format as you would when initializing a Java primitive type with a literal value.

Property Label | Property Key | Description | Type **Enabled** | enabled | Check this box to enable CAS SSO authentication. | boolean **Import from LDAP** | importFromLDAP | Users authenticated from CAS that do not exist in Liferay DXP are imported from LDAP. LDAP must be enabled separately. | boolean **Login URL** | loginURL | Set the CAS server login URL. | String **Logout on session expiration** | logoutOnSessionExpiration | If checked, browsers with expired sessions are redirected to the CAS logout URL. | boolean **Logout URL** | logoutURL | The CAS server logout URL. Set this if you want Liferay DXP's logout function to trigger a CAS logout | String **Server Name** | serverName | The name of the Liferay DXP instance (e.g., liferay.com). If the provided name includes the protocol (https://, for example) then this will be used together with the path /c/portal/login to construct the URL to which the CAS server will provide tickets. If no scheme is provided, the scheme normally used to access the @product@ login page will be used. | String **Server URL** | serviceURL | If provided, this will be used as the URL to which the CAS server provides tickets. This overrides any URL constructed based on the Server Name as above. | String **No Such User Redirect URL** | noSuchUserRedirectURL | Set the URL to which to redirect the user if the user can authenticate with CAS but cannot be found in Liferay DXP. If import from LDAP is enabled, the user is redirected if the user could not be found or could not be imported from LDAP. | String

To override system defaults for a particular portal instance, navigate to the Control Panel, click on *Configuration* → *Instance Settings*, click on *Authentication* on the right and then on *CAS* at the top.

---

# AUTHENTICATING USING SAML

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The SAML (Security Assertion Markup Language) adapter provides Single Sign On (SSO) and Single Log Off (SLO) in your deployment. Each Liferay DXP instance serves as either the Service Provider (SP) or the Identity Provider (IdP). An identity provider is a trusted provider that provides single sign-on for users to access other websites. A service provider is a website that hosts applications and grants access only to identified users with proper credentials.

---

**Note:** A single Liferay DXP instance is *either* the SP or the IdP in your SSO setup; it can't be both. You can, however, use separate instances for both purposes (for example, one instance is the SP and another is the IdP).

---

Below is background on how SAML works. To jump right to its configuration, see the articles on [Setting Up SAML as an Identity Provider](#) or [Setting Up SAML as a Service Provider](#) for instructions on using the SAML adapter. Use the instructions to make the conceptual magic from this article come to life!

---

## 424.1 What's new in Liferay Connector to SAML 2.0

---

The 5.0.0 version of the application for Liferay DXP brings some long-awaited improvements:

- Liferay DXP acting as a Service Provider (SP) can now connect to multiple Identity Providers (IdP).
- Developers have an extension point for customizing which Identity Providers to users can use to sign in.
- Support for other Signature Algorithms (like SHA-256)
- Signature method algorithm URL's can now be blacklisted from the metadata (for example, disabling SHA-1: <http://www.w3.org/2000/09/xmldsig#rsa-sha1>)

**Note:** If you're migrating from a Liferay SAML adapter prior to version 3.1.0, your portal properties are automatically migrated to System Settings configurations. Please see the Configuring SAML article for details on settings.

---

## 424.2 Important SAML Paths

---

For reference, here are a few important SAML paths.

This URL is the default location of the metadata XML file:

```
[host]:[port]/c/portal/saml/metadata
```

Note that when configuring SAML, no importing of SAML certificates is required. Liferay DXP reads certificates from the SAML metadata XML file. If you want a third-party application like Salesforce to read a Liferay SAML certificate, you can export the Liferay DXP certificate from the keystore. The default keystore file is

```
[Liferay Home]/data/keystore.jks
```

You can change this path in System Settings → SSO → SAML Configuration → Key Store Path.

## 424.3 Single Sign On

---

Both the IdP and the SP can initiate the Single Sign On process, and the SSO flow is different depending on each one. Regardless of how it's initiated, SSO is configured for HTTPS between the SP and IdP, so all transport-level communication is encrypted. SAML requests are signed using certificates configured in Liferay DXP, using the SAML Web Browser SSO profile as defined in the SAML 2.0 specification. In all cases, responses are sent using HTTP-POST or HTTP-Redirect. HTTP-POST is preferred because it reduces the risk that the URL is too long for a browser to handle.

Consider IdP initiated SSO first.

## 424.4 Identity Provider Initiated SSO

---

Sometimes a user enters the SSO cycle by sending a request directly from the browser to the IdP.




Figure 424.1: Identity Provider Initiated SSO

### The SSO Request to the IdP

If Liferay DXP is the IdP, the IdP initiated SSO URL

- Must specify the path as `/c/portal/saml/sso`.
- Must include the `entityId` parameter which is the identifier to a previously configured Service Provider Connection (SPC).
- May include a `RelayState` parameter which contains a URL encoded value where the user is redirected upon successful authentication. This URL should point to a location on the desired SPC (according to the SAML 2.0 standards section 3.4.3, this value *must not* exceed 80 bytes in length). It is useful to specify a landing page after SSO has been executed.

For non-Liferay DXP IdPs (Siteminder, ADFS, etc.), consult the vendor's documentation on constructing IdP initiated SSO URLs.

If the IdP determines that the user isn't authenticated, it prompts the user with the appropriate login screen.

### The SSO Response from the IdP

Upon successful authentication, the IdP constructs a SAML Response. It includes attribute statements configured in the designated Service Provider Connection (SPC; see the next article on setting up the SPC in Liferay DXP's SAML adapter).

The IdP sends the response to the Assertion Consumer Service URL. The request contains two parameters: `SAMLResponse` and `RelayState`.

**Note:** The method for sending the SAML response (for example, HTTP-POST) and the Assertion Consumer Service URL are usually imported as part of the SAML metadata XML provided by the SP. In Liferay DXP, you import the SP's metadata in the SAML Adapter's Service Provider Connections tab.

---

### **The SP Processes the SSO Response**

The SP validates and processes the SAML Response. Liferay DXP's SAML solution requires signed SAMLResponse messages. This signature process ensures proper identification for the IdP and prevents potential SAML Response spoofing.

- If one Liferay DXP instance is the IdP and another is the SP, make sure the SAML metadata XML file imported into the SP contains the IdP's certificate.
- If Liferay DXP is the IdP and another application is the SP, export the certificate from the IdP and import it into the SP's certificate store.

If a RelayState is included in the SAML Response, the user is redirected to it. Otherwise the home page of the SP is served.

## **424.5 Service Provider Initiated SSO**

---

Most of the time, authentication requests come from the Service Provider.

### **The SSO Request to the SP**

When the user's browser requests a protected resource or login URL on the SP, it triggers the SP initiated SSO process. When Liferay DXP is the SAML SP, SSO is initiated either by requesting `/c/portal/login` URL or a protected resource that requires authentication (for example, a document not viewable by the Guest Role). If the user requests a protected resource, its URL is recorded in the RelayState parameter. If the user requested `/c/portal/login`, the RelayState can be set by providing the `redirect` parameter. Otherwise, if the portal property `auth.forward.by.last.path` is set to true, the last accessed path is set as the RelayState. For non-Liferay DXP SPs, consult the vendor documentation on initiating SSO.

### **The AuthnRequest to the IdP**

The SP looks up the IdP's Single Sign On service URL and sends an AuthnRequest. When Liferay DXP is the SP it looks up the configured SAML Identity Provider Connection and sends a SAML AuthnRequest to the IdP's Single Sign On service URL as defined in the SAML metadata XML document. Liferay DXP supports sending and receiving the AuthnRequest using HTTP-POST or HTTP-Redirect binding. HTTP-POST is preferred.

If the user doesn't have an active session or if ForceAuthn was requested by the SP, the user must authenticate by providing credentials. When Liferay DXP is the IdP, authentication occurs in the Login Portlet. Liferay DXP decodes and verifies the AuthnRequest before requesting the user to authenticate.



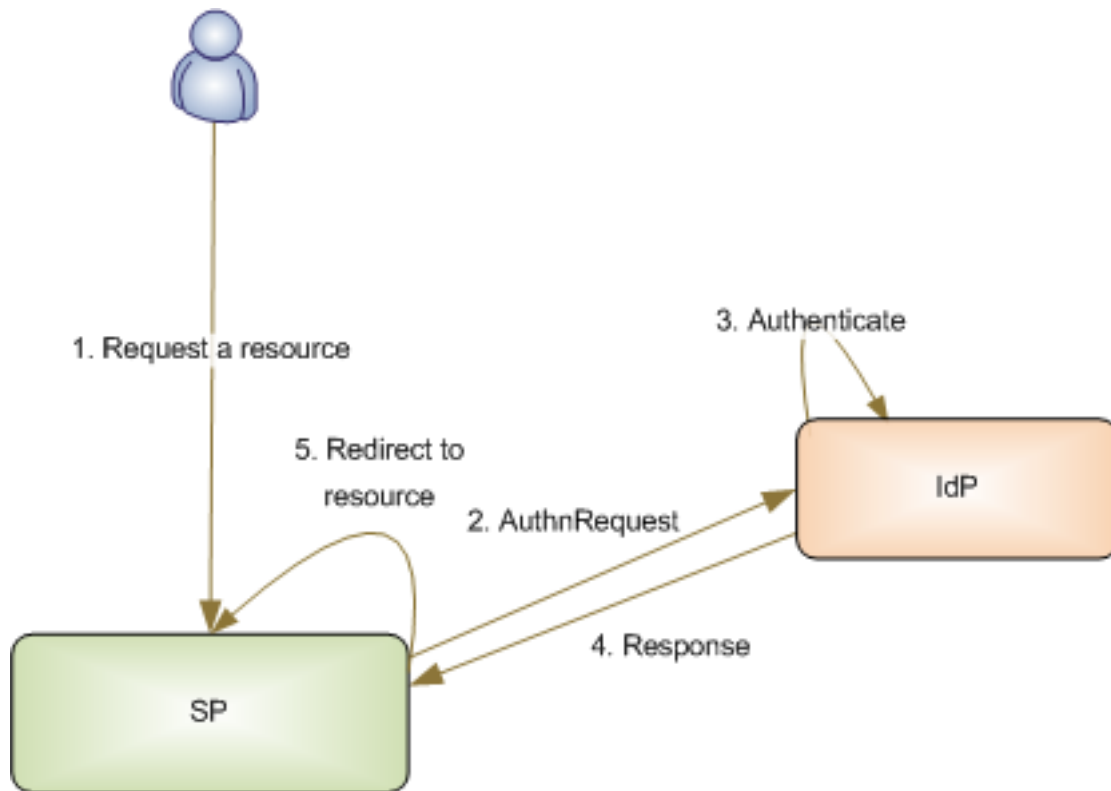


Figure 424.2: Service Provider Initiated SSO

### The SSO Response from the IdP

After authentication, a SAML Response is constructed, sent to the Assertion Consumer Service URL of the SP, and verified. The IdP automatically makes this choice based on the SP metadata.

When Liferay DXP is configured as the IdP, any attributes configured on the Service Provider Connection are included in the response as attribute statements. The Assertion Consumer Service URL is looked up from the SAML metadata XML of the SP.

When Liferay DXP is configured as the SP, response and assertion signatures are verified. Liferay DXP requires the sender to be authenticated. This is done via whole message signature from the issuing IdP. Responses missing the signature are considered unauthenticated and the response is rejected. For non-Liferay DXP SP or IdP vendors, consult their documentation.

The user is redirected to the requested resource or to the URL contained in the RelayState parameter (for example, the last page the user accessed before initiating SSO).

## 424.6 Single Log Off

The Single Log Off request is sent from the user's browser to the IdP or an SP, and the SLO flow differs in each case. First consider IdP initiated SLO.

### 424.7 Identity Provider Initiated SLO

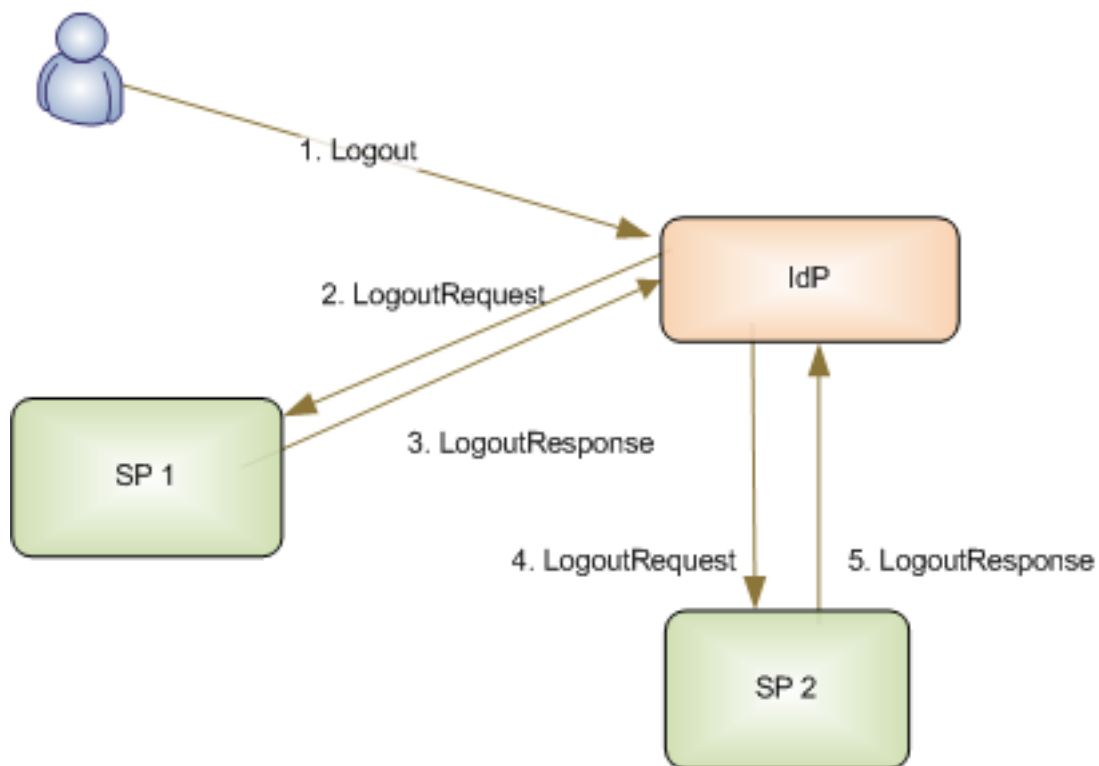


Figure 424.3: Identity Provider Initiated SLO

#### The SLO Request to the IdP

An IdP initiated SLO request is sent directly to the IdP by the user's browser. When Liferay DXP is the IdP, the IdP initiated SSO URL must specify the URL path as

`/c/portal/logout`

If the user is signed on to any configured SP, the SAML plugin takes over the logout process, displaying all the signed on services. The single logout screen displays the authentication status of each SP and whether any SPs can't be logged out of (for example, if the SP is down or doesn't support SLO). For non-Liferay DXP IdPs (Siteminder, ADFS, etc.) consult the vendor's documentation on constructing IdP initiated SLO URLs.

The IdP sends a SAML LogoutRequest to the SP.

- When Liferay DXP is configured as the IdP, the LogoutRequest is sent using either HTTP-POST, HTTP-Redirect, or SOAP binding. HTTP-Post binding is preferred but in its absence, the first available SLO endpoint with supported binding is selected.
- When Liferay DXP is configured as the SP, supported bindings for LogoutRequest are HTTP-Post, HTTP-Redirect, or SOAP.
- For other IdPs or SPs, please consult the vendor's documentation.

### The SLO Response from the SP

The SP delivers a LogoutResponse to the IdP.

The IdP sends a SAML LogoutRequest to the second SP.

The second SP then delivers the LogoutResponse to the IdP. The process is repeated for all SPs the user is logged into. When Liferay DXP is the IdP, Liferay DXP logs the user out after the last SP has delivered its LogoutResponse or has timed out.

## 424.8 Service Provider Initiated SLO

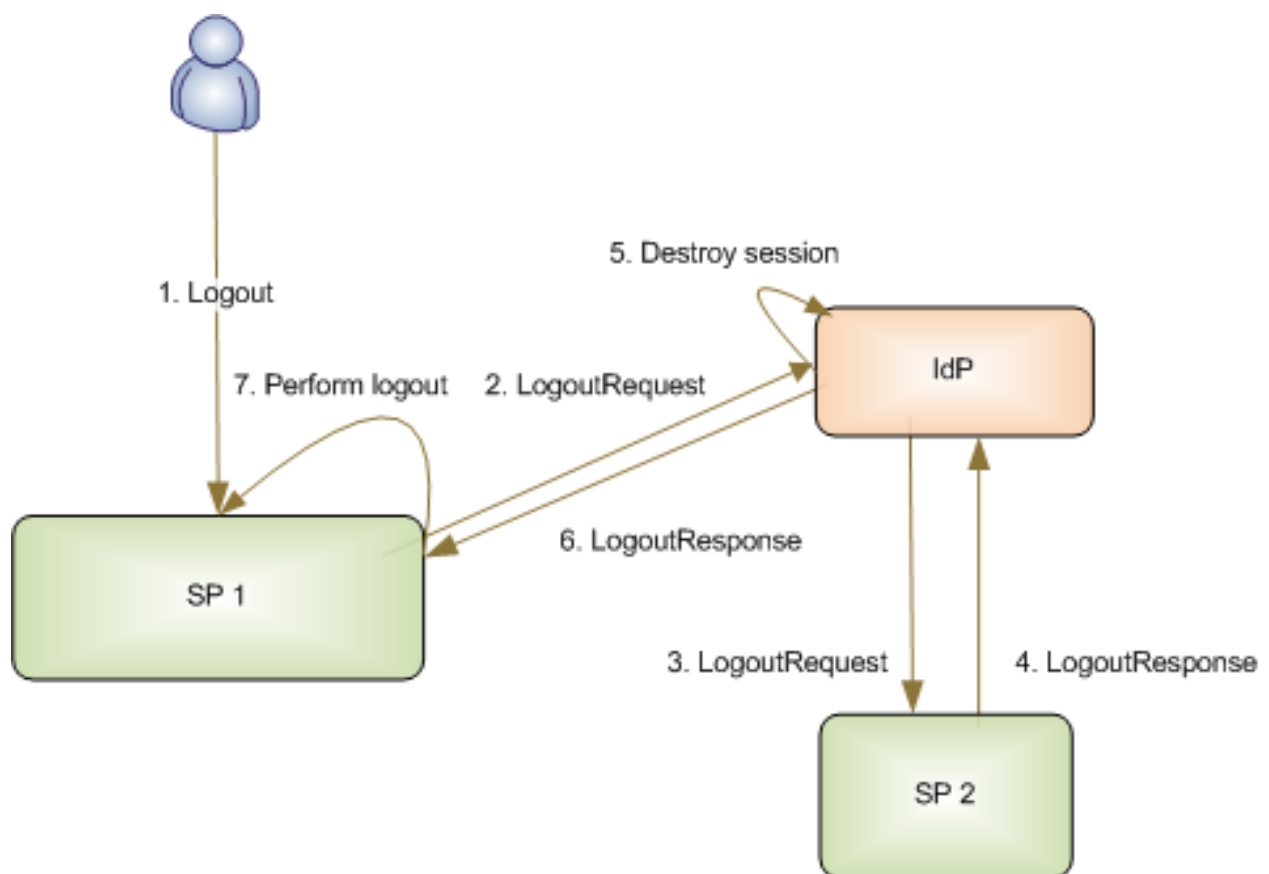


Figure 424.4: Service Provider Initiated SLO

### The SLO Request to the SP

In SP initiated SLO, the user's browser sends a logout request directly to the SP. When Liferay DXP is configured as the SP, the SLO is initiated by requesting this logout URL:

`/c/portal/logout`

For other SPs, consult the vendor's documentation on initiating SLO. A SAML LogoutRequest is sent to the Single Log Out service URL of the IdP.

- If Liferay DXP serves as the SP, the LogoutRequest is sent to the IdP configured by the IdP Connections tab of the SAML provider (see the next article to set up the IdP Connection) and the SLO service URL defined in the SAML metadata.
- When Liferay DXP is the IdP, if the user has logged on to other SPs, the user is presented with a single logout screen with the status of each SP logout, flagging any that can't be logged out of (some SPs might not support SLO or are currently down). If there are no other SPs to log out of, the SAML session terminates and the IdP destroys its session.

### **The SLO Response from the SP**

If the user is logged in to additional SPs (beyond just the initiating SP), the IdP sends the SAML LogoutRequest to each one. When Liferay DXP is the IdP, the LogoutResponse is sent using either HTTP-Post, HTTP-Redirect, or SOAP binding.

Each SP delivers its LogoutResponse to the IdP. When Liferay DXP is the SP, the LogoutResponse is sent using either HTTP-Post, HTTP-Redirect or direct response to SOAP request.

After all additional SPs deliver their LogoutResponses to the IdP, the IdP destroys its SSO session. When Liferay DXP is the IdP, once the last SP has delivered its LogoutResponse or has timed out, the IdP destroys the Liferay DXP session, logging out the user.

Finally, the IdP sends a LogoutResponse to the SP that initiated SLO. The initiating SP terminates its SAML session and logs the user out.

## **424.9 Related Topics**

---

- [Setting Up SAML as an Identity Provider](#)
- [Setting Up SAML as a Service Provider](#)
- [Token-Based SSO Authentication](#)

---

# SETTING UP LIFERAY DXP AS A SAML IDENTITY PROVIDER

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

An identity provider is a trusted provider that provides single sign-on for users to access other websites. A service provider is a website that hosts applications and grants access only to identified users with proper credentials. Liferay Portal 6.1 EE and later versions support SAML 2.0 integration via the Liferay Connector to SAML 2.0 application. It is provided from Liferay Marketplace and allows Liferay DXP to act as a SAML 2.0 identity provider or as a service provider.

---

**Important:** You can set Liferay DXP up as an Identity Provider or as a Service Provider. Each single Liferay DXP instance can serve as an identity provider or as a service provider, but **not both**. Both configurations are covered in this article.

---

## 425.1 Storing Your Keystore

---

Your first step is to determine where to store your keystore. You have two options:

- In the file system
- In the Documents and Media library

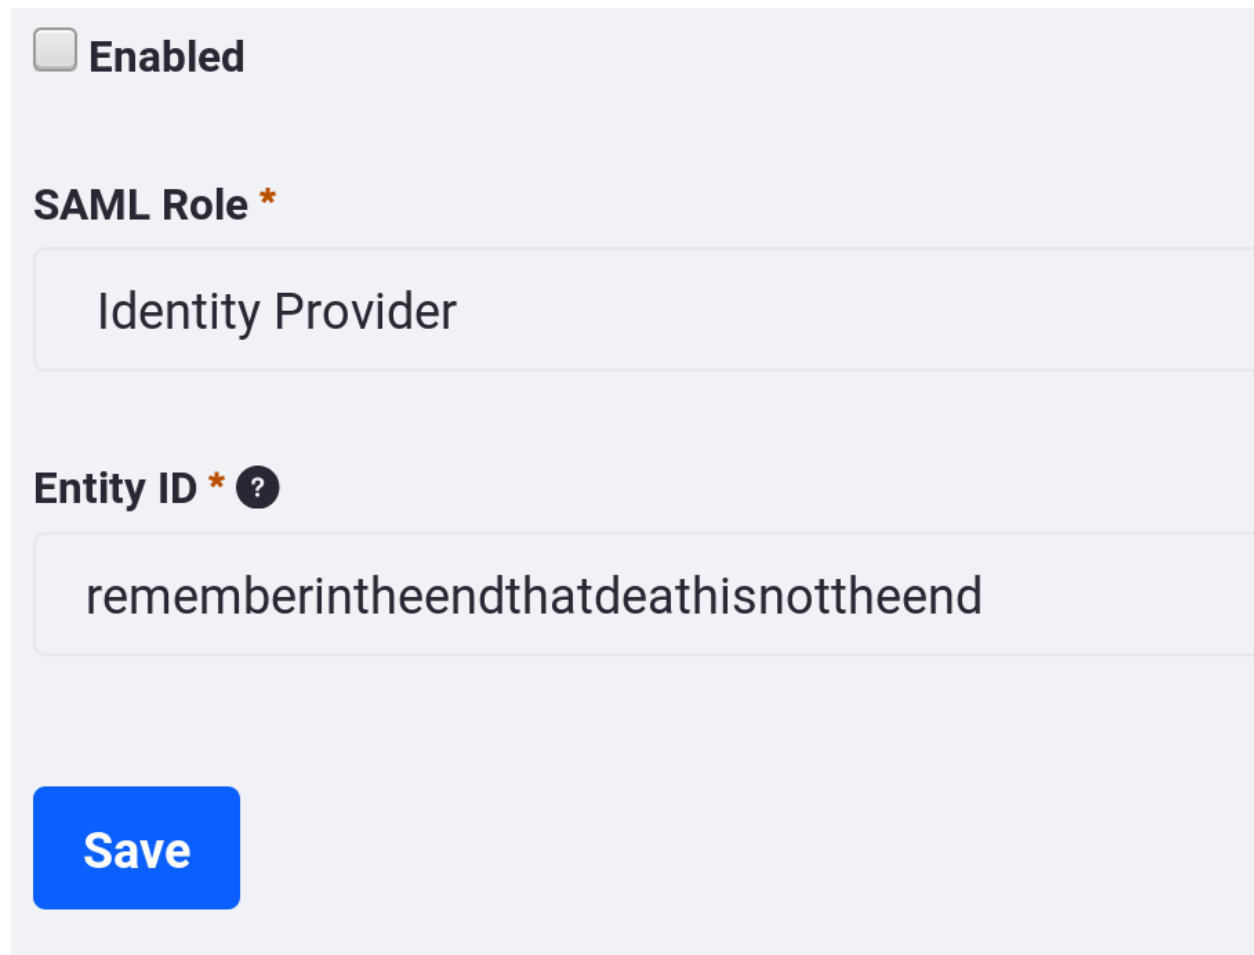
The file system keystore manager is used by default and the default location is the [Liferay Home]/data directory (you can change the location in System Settings → SSO → SAML Configuration → Key Store Path). To use Documents and Media library storage for your keystore instead of file system storage, go to *Control Panel* → *System Settings* → *Security* → *SSO* → *SAML KeyStoreManager Implementation Configuration*. Select from the two options: *Filesystem Keystore Manager* or *Document Library Keystore Manager*.

If you use Document Library storage, you can use any number of back-end file stores. These are protected not only by the system where the key is stored, but also by Liferay DXP's permissions system.

## 425.2 Configuring Liferay DXP as a SAML Identity Provider

---

1. To access the SAML Admin interface, click on *Control Panel* → *Configuration* and then on *SAML Admin*.
2. To begin configuring Liferay DXP to use SAML, select a SAML role for @product@ and choose an entity ID.



Enabled

**SAML Role \***

Identity Provider

**Entity ID \* ?**

rememberintheendthatdeathisnottheend

**Save**

Figure 425.1: Select a SAML role for Liferay and enter an entity ID.

Select the *Identity Provider* SAML role. Enter your own entity ID. Then click *Save*. A new Certificate and Private Key section appears.

3. The Certificate and Private Key section lets you create a keystore for SAML. Click on *Create Certificate* and enter the following information:
  - Your common name (your first and last name)
  - The name of your organization
  - The name of your organizational unit
  - The name of your city or locality
  - The name of your state or province
  - The name of your country
  - The length in days that your keystore will remain valid (how long before the keystore expires)
  - The key algorithm (RSA is the default)
  - The key length in bits (2048 is the default)
  - The key password

Click *Save*.

When you create the certificate and private key, you also create a keystore if one doesn't already exist. As described above, this keystore has two storage options: file system storage (the default) and Documents and Media storage. By default, the certificate uses the SHA256 algorithm for encryption and is fingerprinted and self-signed via RSA and SHA256.

4. After you click *Save*, you can click *Replace Certificate* at any time to replace the current certificate with a new one if your old one has expired or if you want to change the key's password.

Three more tabs now appear:

**General:** For enabling or disabling a SAML IdP and managing the required keystore.

**Identity Provider:** Contains IdP options, such as whether to enable SSL. If SSL has been enabled, then SAML requests are not approved unless they are also encrypted.

**Service Provider Connections:** This tab manages any Service Providers connected to this Liferay DXP instance.

See below for more information on the Identity Provider and Service Provider Connections tabs.

5. After you save your certificate and private key information, check the *Enabled* box at the top of the General tab and click *Save*. You successfully set Liferay DXP up as a SAML Identity Provider!

### 425.3 Changing the Identity Provider Settings

---

To configure Liferay DXP's SAML Identity Provider Settings, navigate to the *Identity Provider* tab of the SAML Admin Control Panel entry.

The *Identity Provider* tab includes these options:

**Sign Metadata?:** Check this box to ensure the metadata XML file that's produced is signed.

The screenshot shows the 'General' tab of the SAML Admin portlet. It features three tabs: 'General', 'Identity Provider', and 'Service Provider Connections'. The 'General' tab is active. Below the tabs, there is a 'Enabled' checkbox which is checked. A 'SAML Role' field contains the text 'Identity Provider'. An 'Entity ID' field contains the text 'rememberintheendthatdeathisnottheend'. A blue 'Save' button is located below the fields. Underneath, there is a section titled 'Certificate and Private Key' which displays the following information: 'Subject DN Optional[C=USA, ST=NJ, L=Somewhere, OU=Knowledge Management, O=Liferay, CN=Joe Bloggs]', 'Serial Number 16a69e287f8', and 'Valid from Mon Apr 29 16:16:54 GMT 2019 until Sun Apr 19 16:16:54 GMT 2020.'. Below this, there is a section for 'Certificate Fingerprints' with 'MD5' and 'SHA1' values. The 'Signature Algorithm' is listed as 'SHA256withRSA'. At the bottom of this section are two buttons: 'Replace Certificate' and 'Download Certificate'.

Figure 425.2: The General tab of the SAML Admin portlet displays information about the current certificate and private key and allows administrators to download the certificate or replace the certificate.

**SSL Required:** Check this box to reject any SAML messages that are *not* sent over SSL. This affects URLs in the generated metadata.

**Require Authn Request Signature?:** When this box is checked, each Authn Request must be signed by the sending Service Provider. In most cases, this should be enabled.

**Session Maximum Age:** Specify the maximum duration of the SAML SSO session in seconds. If this property is not set or is set to 0, the SSO session has an unlimited duration. The SSO session maximum duration can be longer than the portal session maximum duration. If the portal session expires before the SSO session expires, the user is logged back in to Liferay DXP automatically. SSO session expiration does not trigger a single logout from all service providers. You can use the session maximum age, for example, to force users to sign in again after a certain period of time.

**Session Timeout:** Specify the maximum idle time of the SAML SSO session. Even if the session maximum age is unlimited, the SSO session expires whenever the user's idle time reaches the limit set by the session timeout property.



## 425.4 Checkpoint

---

Before adding a Service Provider (SP), verify you've completed these tasks:

1. A SAML keystore has been generated. It can be stored in one of two locations: the data folder or in the Documents and Media library.
2. On the *Identity Provider* tab, the following settings have been set:
  - a. **Sign Metadata** has been checked.
  - b. **SSL Required** - checked if SSL is active elsewhere. SSL is disabled by default.
  - c. **Authn Request Signature Required:** has been checked.
  - d. **Session Maximum Age:** has been set. If set to 0, then the SSO has an unlimited duration.
  - e. **Session Timeout:** Specify the maximum idle time of the SAML SSO session.
3. Once the *Enabled* checkbox has been checked, the IdP is live, and you can generate the required metadata. This URL is the default location of Liferay DXP's metadata XML file:

`[host]:[port]/c/portal/saml/metadata`

If this URL does not display correctly, then the SAML instance has not been enabled. Use the URL or click *Save* in the browser to generate an actual XML file.

Your Identity Provider is now set up. Next, you must register a Service Provider.



---

## REGISTERING A SAML SERVICE PROVIDER

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Setting up Liferay DXP as a SAML Identity Provider is only useful if you can connect to one or more SAML Service Providers. Navigate to the Service Provider Connections tab of the SAML Admin Control Panel entry and click the *Add Service Provider* button to add a SAML Service Provider.

The New Service Provider page includes these options:

**Name:** The name of the Service Provider with which to connect. The name can be anything; it's purely cosmetic.

**Entity ID:** The Service Provider's entity ID. This value must match the entity ID declared in the Service Provider metadata.

**Enabled:** Check this box to activate the Service Provider connection.

**Assertion Lifetime:** Defines the number of seconds after which the SAML assertion issued by the Identity Provider should be considered expired.

**Force Encryption:** If the SP does not provide a public key for encrypting the assertions, abort the single sign-on.

**Metadata:** Provide a URL to the Service Provider metadata XML file or manually upload the Service Provider metadata XML file. If you provide a URL, the XML file is retrieved and periodically polled for updates. The update interval can be configured in System Settings with the Runtime Metadata Refresh Interval (`saml.metadata.refresh.interval` if using a config file) property which specifies a number of seconds. If fetching the metadata XML file by URL fails, you can't enable the Service Provider connection. If the Identity Provider server cannot access the metadata via URL, you can upload the XML file manually. In this case, the metadata XML file is not updated automatically.

**Name Identifier Format:** Choose the Name Identifier Format used in the SAML Response. This should be set according to what the Service Provider expects to receive. For Liferay Service Providers, any selection other than email address indicates that the Name Identifier refers to screen name. The formats don't have any special meaning to Liferay Identity Providers. The NameID value is defined by the Name Identifier attribute. See the next option.

**Name Identifier Attribute Name:** This specifies which attribute of the Liferay DXP User object to use as the NameID value. Possible values include `emailAddress`, `screenName` and `uuid`. Additionally, you can prefix the name with `static:` or `expando:`. If you use the prefix `static:`, the value is whatever comes after `static:`. If you use the prefix `expando:`, the value is whatever custom field is specified after `expando:`. For example, `expando:SSN` would look up the User custom field with the name `SSN`.

**Attributes Enabled:** Include and resolve assertion attributes.

**Attributes Namespace Enabled:** When this box is checked, the attribute names are namespaced like this:

```
urn:liferay:user:expando:
urn:liferay:user:
urn:liferay:groups:
urn:liferay:organizationRole:
urn:liferay:organization:
urn:liferay:roles:
urn:liferay:siteRole:
urn:liferay:userGroupRole:
urn:liferay:userGroups:
```

**Attributes:** Enter a list of attributes to include in the assertion, one per line. Each line is an expression that gets parsed. Examples:

```
organizations
organizationRoles
roles
siteRoles
userGroups
static:[attributeName]=[attributeValue]
expando:[userCustomFieldName]
```

Note that the full namespace depends on the attribute name. Attribute namespaces can be useful. Use them when attribute names from different namespaces might conflict. For example, `expando:user` vs `urn:liferay:roles:user`.

**Keep Alive URL:** If users are logged into several Liferay DXP SP instances via a Liferay DXP IdP, their sessions can be kept alive as long as they keep a browser window open to one of them. Configure this only if the SP is Liferay DXP. The URL is `https://[SP host name]/c/portal/saml/keep_alive`.

## 426.1 Checkpoint

---

Verify your settings are correct by connecting the Liferay DXP-based IdP to its first SP. SPs connect to only one IdP, so if the first one doesn't work, the rest won't either.

1. Provide a general name for the SP.
2. The Entity ID name must be identical to the one declared in the Service Provider metadata.
3. Check the *Enabled* checkbox.
4. Set a value for the *Assertion Lifetime*.
5. Choose whether encryption should be required (recommended).
6. Make sure the SP's metadata has been provided either as a URL or an XML file has been uploaded.
7. Make sure *Name Identifier Format* and *Name Identifier Attribute Name* have been set.

8. Make sure *Attributes Namespace Enabled* has been set.

If you don't have a Service Provider to add right now, that's fine. In the next section, you'll learn how to set Liferay DXP up as a SAML Service Provider. The same instance can't be both, but after you set up another Liferay DXP instance as a Service Provider, come back to this one and add the Service Provider according to the instructions above.



---

## SETTING UP LIFERAY DXP AS A SAML SERVICE PROVIDER

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Many of these steps are similar to configuring Liferay DXP as a SAML Identity Provider. As a reminder, a single Liferay DXP installation can be configured as a SAML Identity Provider *or* as a SAML Service Provider but not as both. If your Liferay DXP installation is already a SAML Identity Provider, use a *different* Liferay DXP installation as a SAML Service Provider.

---

**Note:** If you have a third party IdP with Liferay DXP as the SP, all messages coming from the IdP must be signed. If they're not, an error message appears and communication between the IdP and Liferay DXP fails.

---

1. Install the Liferay Connector to SAML 2.0 app. To confirm successful deployment, look for the *SAML Admin* entry in the Configuration section of the Control Panel.
2. To begin configuring Liferay DXP to use SAML, you must select a SAML role for Liferay DXP and choose an entity ID. Select the *Service Provider* SAML role. Choose your own entity ID. Then click *Save* and a new section entitled Certificate and Private Key appears.
3. The Certificate and Private Key section is for creating a keystore for SAML. Click *Create Certificate* and enter the following information:
  - Your common name (your first and last name)
  - The name of your organization
  - The name of your organizational unit
  - The name of your city or locality
  - The name of your state or province
  - The name of your country
  - The length in days that your keystore will remain valid (how long before the keystore expires)
  - The key algorithm (RSA is the default)
  - The key length in bits (2048 is the default)

- The key password

Remember that the keystore has two storage options: file system storage (the default) and Documents and Media storage. By default, the certificate uses the SHA256 algorithm for encryption and is fingerprinted and self-signed via MD5 and SHA1. When you enter all the required information, click *Save*.

4. After you clicked *Save*, check that you can view information about your certificate or download your certificate. If you can, you successfully created a keystore. After you create a keystore, additional options appear. There are three tabs:

**General:** Enables or disables SAML SP and manages the required keystore.

**Service Provider:** This tab manages basic and advanced configurations for the SP.

**Identity Provider Connection:** This tab manages connections to the IdP. There can be multiple IdP connections.

5. You can also generate an encryption certificate. This is a separate key for encrypting assertions. If you want assertions encrypted, you must generate a key for this. The procedure is exactly the same as generating your certificate in step 3 above.
6. Next, you must configure an Identity Provider connection. Click on the *Identity Provider Connections* tab. Enter a name for the Identity Provider, enter its entity ID, and enter its metadata URL. If you have already followed the previous instructions and configured a separate Liferay DXP installation as an Identity provider, you'd enter the following information:

- Name: *Liferay IdP*
- Entity ID: [ID of IdP]
- Clock Skew: Set the tolerance in milliseconds between SP and IdP.
- Force Authn: Whether the IdP should force reauthentication regardless of context.
- Metadata URL: `http://localhost:8080/c/portal/saml/metadata` (test this URL first)
- Name Identifier Format: See settings.
- Attribute Mapping: See settings.
- Keep Alive URL: See settings.

**Important:** The Liferay Connector to SAML 2.0 app supports using *either* a URL to a SAML IdP metadata file *or* an actual (uploaded) SAML metadata XML file. The value entered in the *Metadata URL* field is persisted to the database only when there a metadata URL and there is no specified metadata XML file. Otherwise, Liferay DXP keeps the original metadata URL in the database. This behavior ensures that once a metadata URL has been specified, there is always a metadata URL saved in the database. This way, if a portal administrator forgets the previously entered metadata URL or its format, he or she can look at the displayed metadata URL and choose to modify the displayed metadata URL or overwrite the previously saved metadata URL by specifying a metadata XML file.

7. Finally, after you save your certificate and private key information and configure an Identity Provider connection, check the *Enabled* box at the top of the General tab and click *Save*. Liferay DXP is now a SAML Service Provider!



Note that the SAML Service Provider session is tied to the normal session on the application server. Session expiration on the application server terminates the session on the Service Provider but does not initiate single logout.

You can add multiple IdP connections. To add another Identity Provider, click *Add Identity Provider* again and enter the details for the other provider. When users log in, they are asked to choose an identity provider, so be sure to name the providers so users can recognize them.

## 427.1 Checkpoint

---

1. A SAML keystore has been generated.
2. Verify the connection to the IdP.
  - a. *Name*: generic name for the IdP.
  - b. *Entity ID*: the same name of the IdP. If the IdP is another Liferay DXP instance, then it is the same name as the above example.
  - c. *Metadata URL*: The IdP's metadata as a URL or as an XML file.
  - d. If the IdP is another @product instance, ensure its corresponding Service Provider Connection for this SP is enabled.
3. On the *General* tab, the *Enabled* checkbox has been checked.
4. Once *Enabled* checkbox has been checked, the service provider's metadata becomes available:

[host]:[port]/c/portal/saml/metadata

## 427.2 Setting Up Liferay DXP as a SAML Service Provider in a Clustered Environment

---

You can use the Liferay Connector to SAML 2.0 app as an SSO solution for a clustered Liferay DXP environment. If your multi-node cluster is behind a load balancer, you must enable all the nodes as SPs, and they must share the same keystore manager.

If using the Filesystem Keystore Manager (the default):

1. Configure each node of your Liferay DXP cluster as a SAML service provider as above.
2. Copy the keystore file ([Liferay Home]/data/keystore.jks, by default) from the first node to the remaining nodes. This file is the Java keystore that's created by the SAML Provider app. The keystore contains the valid or self-signed certificate managed by the SAML connector app.

3. Verify that the service provider metadata has been generated to be used either as a URL or an XML file. The metadata is the same for all nodes because of the same database back-end. The IdP's request goes through the load balancer.
4. At this point, all nodes have the same SAML SP configuration and each of them can respond to web requests and handle the SAML protocol. To test your SSO solution, sign into Liferay DXP via your load balancer, navigate to a few pages of a few different sites, and then log out.

If using the Document Library Keystore Manager, skip step 2 because the keystore file is stored in the database shared by all the nodes.

Now you know how to configure Liferay DXP either as a SAML identity provider or a service provider. You also know how to configure SAML in a clustered environment.

---

## CHANGING THE SETTINGS FOR SERVICE PROVIDER AND IDENTITY PROVIDER CONNECTIONS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

To change the SAML Service Provider Settings, navigate to the Service Provider tab.

The Service Provider tab includes these options:

**Require Assertion Signature?:** Check this box to require SAML assertions to be individually signed in addition to the entire SAML message.

---

**Note:** Individual assertions need not be signed as long as the SAML response itself is signed. The SP and IdP should always communicate over https to have encryption at the transport level.

If you believe man-in-the-middle attacks are possible, the SAML response can be signed. The only reason to sign the assertions is if the SAML response is not signed. In this case, assertions should not only be signed but also encrypted.

---

**Clock Skew:** Clock skew is a tolerance in milliseconds used by the Service Provider for mitigating time differences between the clocks of the Identity Provider and the Service Provider. This usually only matters when assertions have been made to expire very quickly.

**LDAP Import Enabled:** Check this box to import user information from the configured LDAP connection based on the resolved NameID. LDAP connections can be configured from Instance Settings.

**Sign Authn Requests:** Check this box to sign the AuthnRequest even if the Identity Provider metadata indicates that it's not required.

**Sign Metadata:** Check this box to sign the metadata XML file.

**SSL Required:** Check this box to reject SAML messages that are not sent over HTTPS. This does not affect how URLs are generated.

---

### 428.1 Changing the SAML Identity Provider Connection Settings

---

To configure Liferay DXP's SAML Identity Provider Settings, navigate to the Identity Provider Connection tab of the SAML Admin portlet and click the *Edit* action button on the IdP you want to configure.

**Name:** The name of the Identity Provider with which to connect.

**Entity ID:** The Identity Provider's entity ID. This value must match the entity ID declared in the Identity Provider metadata.

**Enabled:** Check the box to enable this IdP.

**Clock Skew:** Clock skew is a tolerance in milliseconds used by the Service Provider for mitigating time differences between the clocks of the Identity Provider and the Service Provider. This usually only matters when assertions have been made to expire very quickly.

**Force Authn:** Check this box to have the Service Provider ask the Identity Provider to re-authenticate the user before verifying the user.

**Metadata:** You can provide a URL to the Identity Provider metadata XML file or you can manually upload it. If you provide a URL, the XML file is automatically retrieved and periodically polled for updates. You can change the update interval in System Settings by modifying the Runtime Metadata Refresh Interval property which specifies a number of seconds. If fetching the metadata XML file by URL fails, you can't enable the Identity Provider connection. If the metadata is inaccessible via URL, you can upload the XML file manually. In this case, the metadata XML file is not updated automatically.

**Name Identifier Format:** Choose the Name Identifier Format used in the SAML Response. Set this according to what the Service Provider expects to receive. For Liferay Service Providers, selections other than email address indicate that the Name Identifier refers to screen name. The formats don't have any special meaning to Liferay Identity Providers. The Name Identifier attribute defines the NameID value.

**Attribute Mapping:** Attribute mapping is done from the attribute name or friendly name in the SAML Response to the Liferay DXP attribute name. For example, if you want to map a response attribute named mail to the Liferay DXP attribute emailAddress, enter the following mapping:

```
mail=emailAddress
```

Available Liferay DXP attributes are: emailAddress, screenName, firstName, lastName, modifiedDate, and uuid.

**Keep Alive URL:** If users are logged into several Liferay DXP SP instances via a Liferay DXP IdP, their sessions can be kept alive as long as they keep a browser window open to one of them. Configure this only if the IdP is Liferay DXP. The URL is `https://[IdP host name]/c/portal/saml/keep_alive`. On the Liferay DXP IdP, configure this URL the same way, but point back to this SP.

Save your changes when you are finished configuring the Liferay DXP instance as a service provider. There is no need to restart the server: the changes are applied immediately.

Make the above configurations through the SAML Control Panel interface and not via properties. Some features of the Liferay Connector to SAML 2.0 app are not available as properties.

---

**Limitation:** The Liferay SAML app can only be used with a single virtual host. Technically, this means that in the SAML metadata for Liferay DXP, only one binding can be added in this form:

```
<md:EntityDescriptor>
...
<md:SPSSODescriptor>
...
```

```
<md:AssertionConsumerService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST" Location="https://portal.domain.com/c/portal/saml/acs" index="1" ...
</md:SPSSODescriptor>
</md:EntityDescriptor>
```

```
\noindent\hrulefill
```

#### # Configuring SAML

```
<aside class="alert alert-info">
 This document has been updated and ported to Liferay Learn and is no longer maintained here.
</aside>
```

There are some ways of configuring the SAML plugin outside the UI. This is done via OSGi configuration files and by uploading metadata XML to configure how connections are negotiated.

#### ## OSGi Configuration Properties

As noted in the previous tutorials, anything related to configuring SP connections must be done through the SAML Admin UI where configurations are saved to Liferay's database. SP connections can no longer be made via properties files as they were in versions prior to 3.1.0.

```
\noindent\hrulefill
```

**Note:** Don't use OSGi `.config` files or Liferay DXP's System Settings Control Panel application to configure SAML providers (IdP or SP). The System Settings UI is auto-generated, and is for advanced administrators. It does not perform the enhanced validation on the fields that the SAML Admin UI performs, so it could allow administrators to create invalid configurations.

```
\noindent\hrulefill
```

This is a portal instance scoped configuration which can be managed via OSGi Configuration Admin. The affected properties are those in the ``SAMLProviderConfiguration`` metatype:

```
- `keyStoreCredentialPassword()`
- `keyStoreEncryptionCredentialPassword()`
- `assertionSignatureRequired()`
- `authnRequestSignatureRequired()`
- `clockSkew()`
- `defaultAssertionLifetime()`
- `entityId()`
- `enabled()`
- `ldapImportEnabled`
- `role()`
- `sessionMaximumAge`
- `sessionTimeout()`
- `signAuthnRequest()`
- `signMetadata()`
- `sslRequired()`
- `allowShowingTheLoginPortlet()`
```

The SAML Admin UI remains the place for creating the portal instance scoped configuration instances.

Note that there is also a system wide configuration, represented by the ``SamlConfiguration`` metatype.

If you used Liferay 6.2, please note that the following system wide properties were removed:

```
`saml.metadata.paths` (served no purpose after removal of SP connection defaults)
`saml.runtime.metadata.max.refresh.delay`
`saml.runtime.metadata.min.refresh.delay`
```

The latter two properties were replaced with the single property  
`com.liferay.saml.runtime.configuration.SamlConfiguration.getMetadataRefreshInterval()`.

Note also the introduction of the *SAML KeyStoreManager Implementation Configuration* in *Control Panel* & *System Settings* & *Security* & *SSO*. The options for this configuration are explained above in the *Setting up Liferay DXP as a SAML Identity Provider* section.

In the latest version of the plugin, the `SHA256` algorithm is the default encryption algorithm used in the configuration and to generate keys. The default configuration tries `SHA256`, then `SHA384`, then `SHA512` before falling back to `SHA1`. Because `SHA1` is potentially vulnerable, you can blacklist it using this property:

```
``properties
blacklisted.algorithms=["blacklisted_algorithm_url", "another_blacklisted_algorithm_url"]
```

To blacklist SHA1, therefore, you'd have this configuration:

```
blacklisted.algorithms=["http://www.w3.org/2000/09/xmldsig#sha1"]
```

Place these in a config file with this name:

```
com.liferay.saml.opensaml.integration.internal.bootstrap.SecurityConfigurationBootstrap.config
```

There's a lot more granularity in how connections are negotiated if you configure the metadata XML.

## 428.2 Configuring Negotiation Via metadata.xml

---

If the default negotiation configuration doesn't work for you, you can craft your own configuration and upload it. Before doing this, visit your host's metadata URL and save a copy of the configuration in case you need it later:

```
http://[hostname]/c/portal/saml/metadata
```

For example, if you're stuck connecting to a legacy IdP that only supports SHA1, you can upload a configuration that disables the other algorithms:

```
<?xml version="1.0" encoding="UTF-8"?>
<md:EntityDescriptor xmlns:md="urn:oasis:names:tc:SAML:2.0:metadata" entityID="samlidp">
 <md:IDPSSODescriptor WantAuthnRequestsSigned="true" protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
 <md:Extensions>
 <alg:SigningMethod xmlns:alg="urn:oasis:names:tc:SAML:metadata:algsupport" Algorithm="http://www.w3.org/2000/09/xmldsig#rsa-sha1"/>
 </md:Extensions>
 <md:KeyDescriptor use="signing">
 <ds:KeyInfo xmlns:ds="http://www.w3.org/2000/09/xmldsig#">
 <ds:X509Data>
 <ds:X509Certificate>... omitted ...</ds:X509Certificate>
 </ds:X509Data>
 </ds:KeyInfo>
 </md:KeyDescriptor>
 </md:IDPSSODescriptor>
</md:EntityDescriptor>
```

```
</ds:KeyInfo>
</md:KeyDescriptor>
<md:SingleLogoutService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST" Location="http://localhost:8080/c/portal/saml/slo"/>
<md:SingleLogoutService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect" Location="http://localhost:8080/c/portal/saml/slo"/>
<md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect" Location="http://localhost:8080/c/portal/saml/sso"/>
<md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST" Location="http://localhost:8080/c/portal/saml/sso"/>
</md:IDPSSODescriptor>
</md:EntityDescriptor>
```

Notice that in the configuration above, the `<md:Extensions>` block has only one signing algorithm: SHA1.

---

**Note:** Since the default configuration falls back to SHA1, you shouldn't need to do this unless your legacy system can't negotiate via the fallback mechanism. Also note that if you blacklisted SHA1, this won't work. Due to vulnerabilities in SHA1, it's best to avoid using it altogether if possible.

---

If you've changed your metadata configuration, you can go back to the default configuration if you saved it before making the change. If you didn't, you can provide a URL instead of an uploaded XML file to one of your peers' metadata configurations.





---

## NTLM SINGLE SIGN ON AUTHENTICATION

---

NTLM (NT LAN Manager) is a suite of Microsoft protocols that provide authentication, integrity, and confidentiality for users. Though Microsoft has adopted Kerberos in modern versions of Windows server, NTLM is still used when authenticating to a workgroup. Liferay DXP now supports NTLM v2 authentication. NTLM v2 is more secure and has a stronger authentication process than NTLMv1.

---

**Note:** NTLM authentication was deprecated in 7.0 and was removed. You can still install it from Marketplace [here](#) or [here](#).

---

Note that in order to use NTLM SSO, Liferay DXP's portal instance authentication type must be set to screen name.

---

**Note:** To USE NTLM with Liferay DXP, you must configure your browser. Consult your browser vendor's documentation for the details.

---

Most importantly, all users *must* be imported from an Active Directory server. NTLM (and Kerberos) works only if the users are in the AD; otherwise any SSO requests initiated by Liferay DXP fail.

NTLM configuration can be applied either at the system scope or at the scope of a portal instance. To configure the NTLM SSO module at the system scope, navigate to the Control Panel, click on *Configuration* → *System Settings* → *Security* → *SSO* → *NTLM*. The values configured there provide the default values for all portal instances. Enter values in the same format as you would when initializing a Java primitive type with a literal value.

Property Label | Property Key | Description | Type **Enabled** | enabled | Check this box to enable NTLM SSO authentication. Note that NTLM will only work if Liferay DXP's authentication type is set to screen name. | boolean **Domain Controller** | domainController | Enter the IP address of your domain controller. This is the server that contains the user accounts you want to use with Liferay DXP. | String **Domain Controller Name** | domainControllerName | Specify the domain controller NetBIOS name. | String **Domain** | domain | Enter the domain / workgroup name | String **Service Account** | serviceAccount | You need to create a service account for NTLM. This account will be a computer account, not a user account. | String **Service Password** | serviceAccount | Enter the password for the service account. | String **Negotiate Flags** | negotiateFlags | Only available at system

level. Set according to the client's requested capabilities and the server's ServerCapabilities. See here | String

Note the AD's name and IP address correspond to the domainControllerName and domainController settings. The Service Account is for the *NTLM* account (registered with NTLM), not the Liferay DXP user account.

To override system defaults for a particular portal instance, navigate to the Control Panel, click on *Configuration* → *Instance Settings*, click on *Authentication* and then on *NTLM*.

## 429.1 Summary

---

NTLM authentication is often highly desirable in Intranet scenarios where the IT department has control over what software is running on client devices and thus can ensure NTLM compatibility. In an Active Directory based network / domain, it is hard to beat the user experience that NTLM authentication can provide.

Please remember that in order to use NTLM SSO, your Liferay DXP instance authentication type must be set to screen name *and* that all users have been imported from your active directory. If this is not acceptable for your Liferay DXP implementation, then another SSO solution (such as CAS) can be used as a broker between your portal and the NTLM authentication process.

---

## OPENID SINGLE SIGN ON AUTHENTICATION

---

OpenID is a single sign-on standard implemented by multiple vendors. Users can register for an ID with the vendor they trust. The credential issued by that vendor can be used by all the web sites that support OpenID. Some high profile OpenID vendors are Google, Paypal, Amazon, and Microsoft. Please see the OpenID site for a more complete list.

---

**Note:** OpenID is deprecated in 7.0 and has been removed. You can still install it from Marketplace [here](#) or [here](#).

---

With OpenID, users don't have to register for a new account on every site which requires an account. Users register on *one* site (the OpenID provider's site) and then use those credentials to authenticate to many web sites which support OpenID. Web site owners sometimes struggle to build communities because users are reluctant to register for *another* account. Supporting OpenID removes that barrier, making it easier for site owners to build their communities. All the account information is kept with the OpenID provider, making it much easier to manage this information and keep it up to date.

Liferay DXP can act as an OpenID consumer, allowing users to automatically register and sign in with their OpenID accounts. Internally, the product uses OpenID4Java to implement the feature.

---

### 430.1 OpenID at the System Scope

---

OpenID is enabled by default in Liferay DXP but can be disabled or enabled at either the system scope or portal instance scope. To configure the OpenID SSO module at the system level, navigate to the Control Panel and click on *Configuration* → *System Settings* → *Security* → *SSO*. There's only a single configuration setting. Check the *Enabled* box to enable OpenID at the system scope (for all portal instances), uncheck it to disable it at the system scope.

## 430.2 OpenID at the Instance Scope

---

To configure the OpenID SSO module at the portal instance scope, navigate to the Control Panel and click on *Configuration* → *Instance Settings*, then on *Authentication* → *OpenID*. There's only a single configuration setting. Check the *Enabled* box to enable OpenID for the current portal instance, or uncheck it to disable it for the current portal instance.

Regardless of whether OpenID is enabled at the System or Instance scope, users can see the OpenID icon when they sign into Liferay DXP. Click *Sign In*. The OpenID icon is displayed at the lower left.

---

## AUTHENTICATING WITH KERBEROS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

You can use Kerberos to authenticate Microsoft Windows™ accounts with Liferay DXP. This is done completely through configuration by using a combination of Liferay DXP's LDAP support and a web server that supports the Kerberos protocol.

Note that this configuration is preferred above NTLM because security vulnerabilities persist.

While it's beyond the scope of this article to explain how to set up Kerberos and Active Directory on a Windows™ server, we can describe the minimum prerequisites for setting up Liferay authentication:

1. A Windows™ server with Active Directory and DNS set up so the AD server and Liferay DXP can resolve each other on the network. In other words, they must be able to ping each other *by name*.
2. An administrative user in AD Liferay DXP can use to bind to AD.
3. A Kerberos keytab file exported via the `ktpass` command containing the cryptographic information the Liferay DXP server needs to bind to AD.
4. A web server in front of Liferay DXP that supports Kerberos, such as Apache, NGNIX, or IIS. The web server must also support injecting a header to be used as a token in the Liferay DXP configuration (see below).
5. Of course, you need a Liferay DXP installation that can also resolve by name the other servers. It should never run on the Active Directory server.

When you have all of these prerequisites in place, you're ready to configure Kerberos authentication.

---

### 431.1 How Kerberos Authentication Works

---

From the prerequisites, you may be able to guess that there are several moving parts to how SSO works with Kerberos.

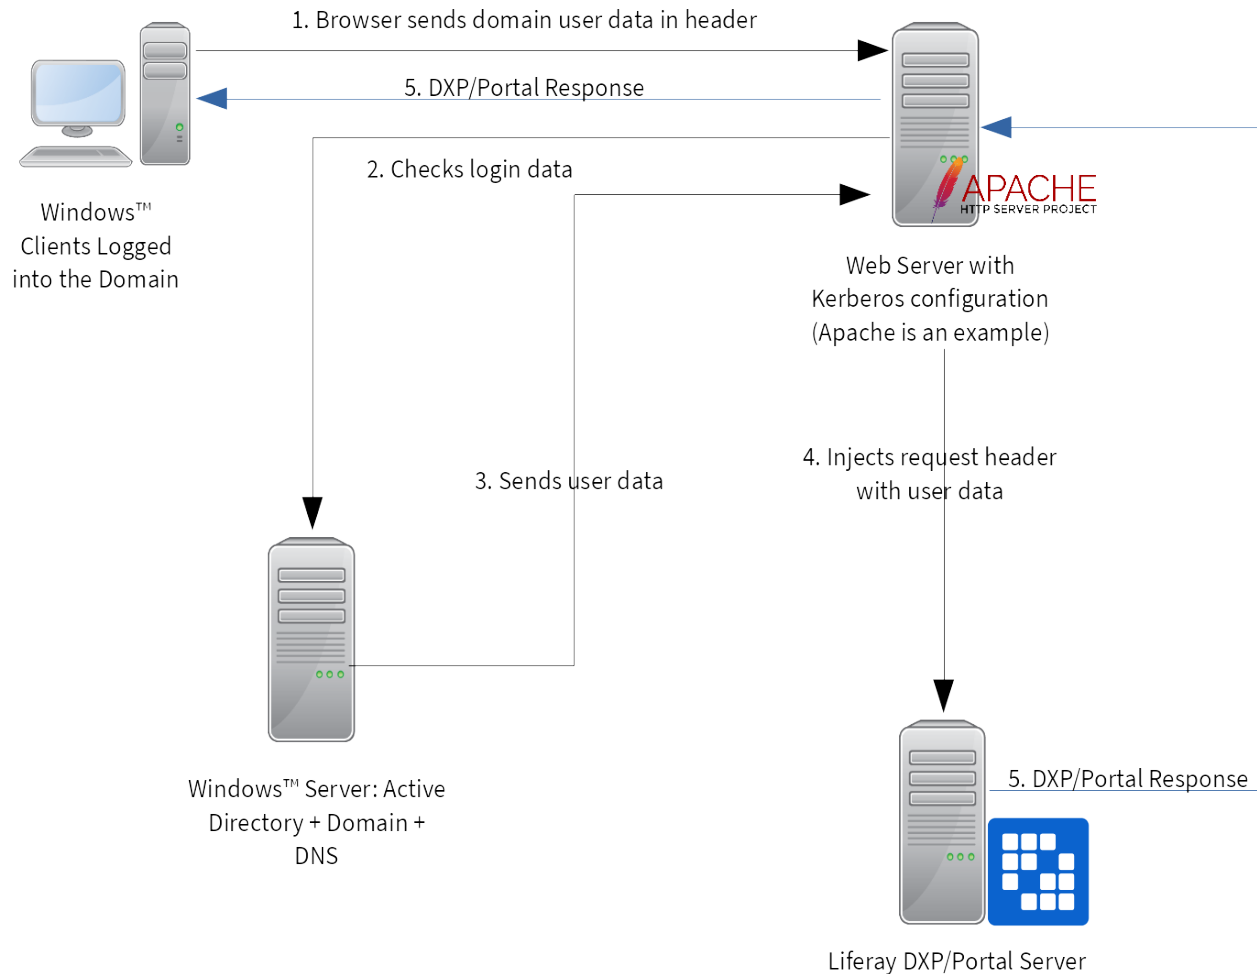


Figure 431.1: Kerberos authentication requires a web server in front of your Liferay DXP server.

First, a properly configured web browser sends a negotiate request using encrypted Windows user data. To configure this, the browser must recognize the site as a trusted site (explained below). The web server's Kerberos module uses the keytab file to bind over the Kerberos protocol to AD and verify the user information. If all is okay, the AD server confirms the connection with a valid response.

The web server you choose must support both the Kerberos protocol and the injection of a custom header into the request that Liferay DXP can later read. When the web server forwards the request to Liferay DXP, it reads the header to obtain the user data and authenticate the user.

Next, you'll learn how to get all of this working.

## 431.2 Configuring Kerberos Authentication

---

There are four components to configure: a user keytab from Active Directory, a web server in front of your application server, Liferay DXP, and your Windows™ clients.

### 431.3 Creating the User Keytab

---

1. Create a user so Liferay DXP can bind to Active Directory.
2. Generate a Kerberos keytab file using ktpass:

```
ktpass -princ HTTP/[web server host name][domain] -mapuser [user name][domain] -crypto ALL -ptype KRB5_NT_PRINCIPAL -
pass [password] -out c:\kerberos.keytab
```

For example:

```
ktpass -princ HTTP/mywebserver.intdomain.local@INTDOMAIN.LOCAL -mapuser Marta@INTDOMAIN.LOCAL -crypto ALL -ptype KRB5_NT_PRINCIPAL -
pass password-for-Marta -out c:\kerberos.keytab
```

3. Ensure that the AD domain controller and the web server can see each other on the network via DNS configuration or hosts file.

### 431.4 Configuring Your Web Server

---

1. Configure Kerberos authentication. On Linux, this involves installing krb5 and configuring it to match your realm that's already configured for Active Directory. The example domain for the user configured in step two above would look like this:

```
[libdefaults]
 default_realm = INTDOMAIN.LOCAL

[domain_realm]
 mywebserver.intdomain.local = INTDOMAIN.LOCAL
 intdomain.local = INTDOMAIN.LOCAL
 .intdomain.local = INTDOMAIN.LOCAL

[realms]
 INTDOMAIN.LOCAL = {
 admin_server = winserver.intdomain.local
 kdc = winserver.intdomain.local
 }
```

2. Copy the keytab file you generated on your AD server to your web server.

3. Configure your web server, making sure you set the correct server name, Kerberos service name, Kerberos authentication realms, and the path to the keytab file. For example, if you're using the Apache HTTP server, the configuration might look like this:

```

LoadModule headers_module /usr/lib/apache2/modules/mod_headers.so
LoadModule rewrite_module /usr/lib/apache2/modules/mod_rewrite.so
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so
LoadModule proxy_ajp_module /usr/lib/apache2/modules/mod_proxy_ajp.so
LoadModule auth_kerb_module /usr/lib/apache2/modules/mod_auth_kerb.so

<VirtualHost *:10080>
 <Proxy *>
 Order deny,allow
 Allow from all
 </Proxy>
 ProxyRequests Off
 ProxyPreserveHost On
 ProxyPass / ajp://localhost:8009/
 ProxyPassReverse / ajp://localhost:8009/
 ServerName mywebserver.intdomain.local
 <Location />
 Order allow,deny
 Allow from all
 AuthType Kerberos
 KrbServiceName HTTP/mywebserver.intdomain.local@INTDOMAIN.LOCAL
 AuthName "Domain login"
 KrbAuthRealms INTDOMAIN.LOCAL
 Krb5KeyTab /etc/apache2/kerberos.keytab
 require valid-user
 KrbMethodNegotiate On
 KrbMethodK5Passwd Off
 #KrbLocalUserMapping On

 # Below directives put logon name of authenticated user into http header X-User-Global-ID
 RequestHeader unset X-User-Global-ID
 RewriteEngine On
 RewriteCond %{LA-U:REMOTE_USER} (.+)
 RewriteRule /* - [E=RU:%1,L,NS]
 RequestHeader set X-User-Global-ID %{RU}e

 # Remove domain suffix to get the simple logon name
 # RequestHeader edit X-User-Global-ID "@INTDOMAIN.LOCAL$" ""

 </Location>
</VirtualHost>
Listen 10080

```

The last line is commented out based on user preference. If you want the domain removed from the user name when saved in Liferay DXP, uncomment it. Otherwise, leave it commented out to store the domain with the user name.

### 431.5 Connecting Liferay DXP to Active Directory over LDAP

---

1. Finally, configure Liferay DXP to access Active Directory via the LDAP protocol. Change authentication to be by Screen Name by selecting it in Configuration → Instance Settings → Authentication → General.



2. Connect Liferay DXP to AD over LDAP by going to Configuration → Instance Settings → Authentication → LDAP and adding an LDAP server. Provide the information appropriate to your installation:

**Base Provider URL:** Your AD server on the proper port.

**Base DN:** Your domain configuration. The example above might be DC=INTDOMAIN.DC=LOCAL.

**Principal/Credentials:** Supply the credentials for the user exported to the keytab file.

**Authentication Search Filter:** Supply the appropriate search filter to return user objects. For example, (&(objectCategory=person)(sAMAccountName=\*))

**UUID:** Supply what uniquely identifies a user, such as sAMAccountName.

**Screen Name:** Supply the field that should be mapped to Liferay DXP's screen name field, such as sAMAccountName.

**Password:** Supply the field that contains the user's password, such as userPassword.

3. Be sure to test the connection, save, and enable the configuration.
4. Finally, configure the token for single sign-on at Configuration → System Settings → Security → SSO → Token Based SSO. Make sure the User Token Name matches *exactly* the token you configured in your web server. Click the *Enabled* and *Import from LDAP* boxes and click *Save*.

Excellent! You've configured your servers. All that's left is to configure your clients.

## 431.6 Configuring your Clients

---

You must do two things: make your computer log into the domain and configure your Liferay DXP server as a trusted Internet site.

1. Join your computer to your domain. In keeping with the example above, you'd make your computer a member of the INTDOMAIN.LOCAL domain.
2. Log in as a user in that domain.
3. Internet Explorer, Edge, and Chrome use the Windows™ settings for trusted sites. If you use these browsers, go to Internet Options → Security → Local Intranet Sites and add your Liferay DXP server's URL. For example, add `http://mywebserver.intdomain.local:10080`.
4. Firefox can be configured by typing `about:config` in its address bar. Search for the below two preferences and add the Liferay DXP server's URL as the value for both:
  - `network.negotiate-auth.delegation-uris`
  - `network.negotiate-auth.trusted-uris`

After configuring these things, test your configuration by accessing Liferay DXP through the web server's URL. Since you are already logged into your client machine, you should be automatically logged into Liferay DXP without a user/password prompt.

Congratulations on configuring Kerberos with Liferay DXP!



---

## CONFIGURING CORS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

CORS stands for Cross-Origin Resource Sharing. An Origin is a web server at a different domain, and a Resource is some asset stored on the server, like an image, PDF, or HTML file. Sometimes you must request resources stored on another origin. This is called a cross-origin request, and web servers have policies to allow or deny such requests.

For example, browsers themselves don't allow cross-origin AJAX-style requests from scripts to help mitigate cross-site scripting attacks. These APIs follow a *same origin* policy. But for certain resources, it can be convenient to allow Liferay DXP to serve them to different origins.

For example, if you manage images in Docs & Media, you may want to allow cross-origin requests for them. You can enable CORS for matching URLs in Liferay DXP or for JAX-RS application resources.

### 432.1 Enabling CORS for Liferay DXP Services

---

You'll find the settings in Configuration → System Settings → Security → Security Tools → Portal Cross Resource Origin Sharing (CORS):

1. Click *Add* to create a configuration entry.
2. Fill out the fields on the form. When finished, click *Save*.

**Enabled:** Check this box to enable the entry.

**Name:** Give the configuration entry a name.

**URL Pattern:** Use the Plus button to add as many patterns as you need. Define patterns that match URLs to the resources you want to share. For example, if you have many attachments in the Knowledge Base application, you could define this pattern:

`/knowledge_base/*`

The screenshot displays the Liferay Security Tools configuration interface. The top navigation bar includes a hamburger menu, a back arrow, the title "Security Tools", and a user profile icon. The breadcrumb trail reads "System Settings > Security Tools > Portal Cross Resource Origin Sharing (CO...".

The left sidebar, titled "SYSTEM SCOPE", lists several settings: CAPTCHA, Portal Cross Resource Origin Sharing (CORS) (which is selected), Web Contexts Cross Resource Origin Sharing (CORS), AntiSamy Sanitizer, Inline Permission, and WeDeploy Auth Configuration.

The main configuration area is titled "Add" and contains a blue box with the message: "This configuration is not saved yet. The values shown are the default." Below this, the status is "Enabled" with a checked checkbox. The "Name" field is empty, with a description: "Identifies the configuration. The name is not used directly in the application." The "URL Pattern" field contains "/o/api\*" and is described as "Set URL pattern where CORS is enabled." There are four "CORS Response Headers" sections, each with a description: "Set the headers to be sent in response to CORS requests in the form of 'Name: Value'". The headers are:
 

- Access-Control-Allow-Credentials: true
- Access-Control-Allow-Headers: \*
- Access-Control-Allow-Methods: \*
- Access-Control-Allow-Origin: \*

 At the bottom of the configuration area are "Save" and "Cancel" buttons.

Figure 432.1: The CORS system settings provide a way to configure CORS headers for Liferay services.

This would define resources stored in the Knowledge Base as applicable to the policy you set in the response header below.

**CORS Response Headers:** Use the Plus button to add as many headers as you need. Define policies for any of the CORS headers here.

You can also use a configuration file to configure CORS.

## 432.2 Enabling CORS for JAX-RS Applications

---

You'll find the settings in Configuration → System Settings → Security → Security Tools → Web Contexts Resource Origin Sharing (CORS):

1. Click *Add* to create a configuration entry.
2. Fill out the fields on the form. When finished, click *Save*.

**Dynamic Web Context OSGi Filter:** Define an LDAP-style filter to define which JAX-RS whiteboard applications the CORS headers in this entry apply to. This is the default filter:

```
(&(!(liferay.cors=false))(osgi.jaxrs.name=*)
```

It applies CORS headers to all deployed JAX-RS whiteboard applications without a `liferay.cors=false` property. This helps during development, but in production you should use the narrowest configuration possible.

**URL Pattern:** Use the Plus button to add as many patterns as you need. Define patterns that match URLs to the web services you want to access.

**CORS Response Headers:** Use the Plus button to add as many headers as you need. Define policies for any of the CORS headers here.

JAX-RS developers can use the `@CORS` annotation to set policies for their deployed applications.

System Settings > Security Tools > Web Contexts Cross Resource Origin Shar...

**SYSTEM SCOPE**

- CAPTCHA
- Portal Cross Resource Origin Sharing (CORS)
- Web Contexts Cross Resource Origin Sharing (CORS)**
- AntiSamy Sanitizer
- Inline Permission
- WeDeploy Auth Configuration

### Add

This configuration is not saved yet. The values shown are the default.

**Enables CORS for matching URLs inside dynamic web contexts.**

**Dynamic Web Context OSGi filter**

(osgi.jaxrs.name=MyApplicationName)

Set the OSGi filter to match dynamic web contexts to enable this CORS configuration.

**URL Pattern**

/demo/get

Set URL pattern where CORS is enabled.

**URL Pattern**

/demo/post

Set URL pattern where CORS is enabled.

**CORS Response Headers**

Access-Control-Allow-Credentials: true

Set the headers to be sent in response to CORS requests in the form of "Name: Value".

**CORS Response Headers**

Access-Control-Allow-Headers: \*

Set the headers to be sent in response to CORS requests in the form of "Name: Value".

**CORS Response Headers**

Access-Control-Allow-Methods: \*

Set the headers to be sent in response to CORS requests in the form of "Name: Value".

**CORS Response Headers**

Access-Control-Allow-Origin: \*

Set the headers to be sent in response to CORS requests in the form of "Name: Value".

**Save** **Cancel**

Figure 432.2: There's a separate system settings category for CORS web contexts.

---

# ANTISAMY

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay DXP includes an AntiSamy module that protects against user-entered malicious code. If your site allows users to post content, such as in message boards, blogs, or other applications, they could include malicious code either intentionally or unintentionally. The AntiSamy module filters HTML/CSS fragments and removes suspect JavaScript code from them.

The module leverages the powerful OWASP AntiSamy library to enforce a content policy that's been effective for the auction site eBay. The AntiSamy module adds an OWASP AntiSamy implementation to your portal's list of existing sanitizer implementations. Liferay DXP uses the AntiSamy sanitizer and any existing configured sanitizers to scrub user input to blogs entries, calendar events, message boards posts, wiki pages, and web content articles.

AntiSamy is enabled by default.

---

## 433.1 Configuring AntiSamy

---

AntiSamy uses both a blacklist and a whitelist, so you can define subsets of entities that should be sanitized or not sanitized. The whitelist prevents content of that type from being filtered, while the blacklist filters content of that type.

By default, everything is sanitized except for `JournalArticle`, `BlogsEntry`, and `FragmentEntry`. The assumption is that users posting these kinds of content are trusted, while users posting message boards or wiki articles may not be trusted. If this is not the configuration you want, you can change it:

1. Navigate to *Control Panel* → *System Settings* → *Security Tools* → *AntiSamy Sanitizer*.
2. Enter a package path you want to sanitize into the *Blacklist* field.
3. Use the plus (+) button to add further Blacklist fields if you need them.
4. Use the plus (+) button to add further Whitelist fields if you need them.

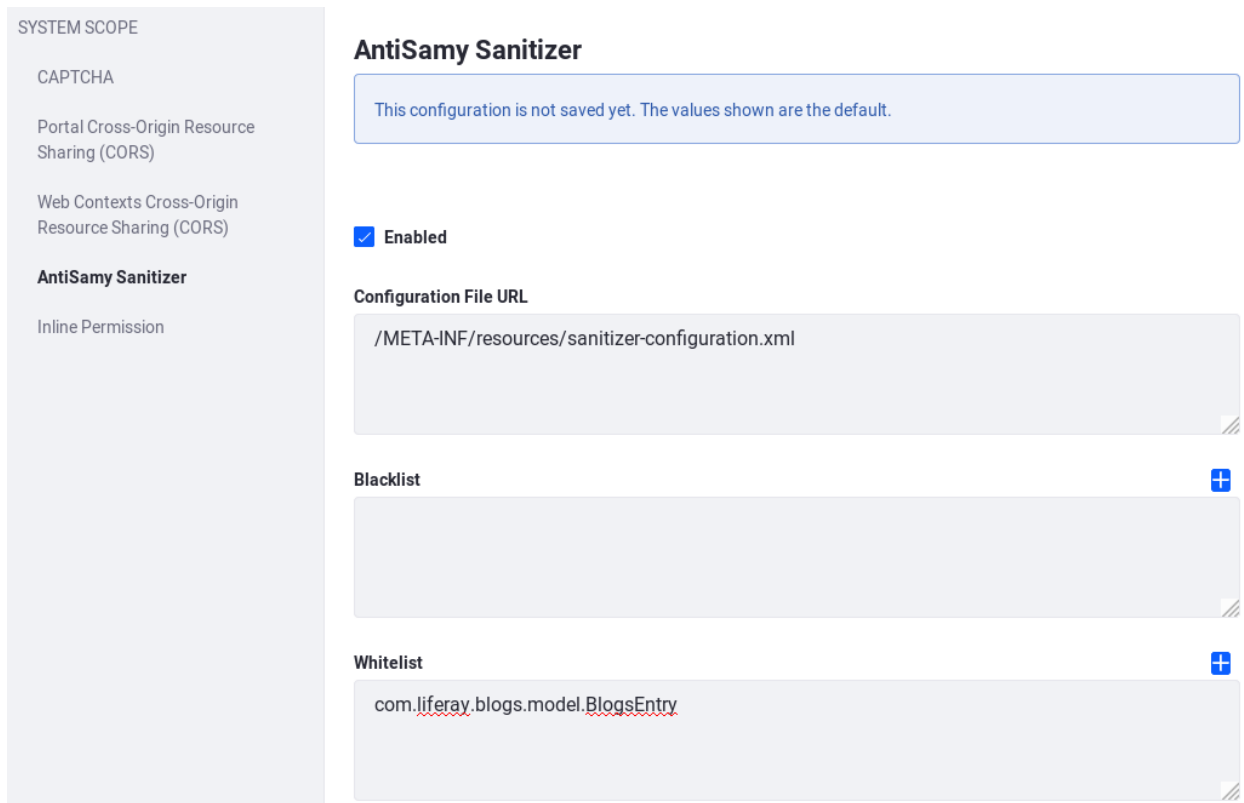


Figure 433.1: Liferay DXP’s AntiSamy configuration options allow you to specify both a blacklist and a whitelist.

5. Enter a package path you don’t want sanitized into a *Whitelist* field.
6. If you want to remove a package path from the configuration, click the trash can icon.
7. When finished, click *Save*.

## 433.2 Using Wildcards

---

You can use wildcards in the configuration. For example, if you only want to sanitize message board posts and nothing else, you can

1. Configure the whitelist to \*
2. Configure the blacklist to `com.liferay.message.boards.*`

The whitelist and the blacklist work together. Without the blacklist, the above configuration’s whitelist must include every content type except `com.liferay.message.boards`, which would be a daunting task to configure.

Use AntiSamy to ensure user-generated content stays safe for other users to view.



---

## OAUTH 2.0

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

OAuth 2.0 is an industry-standard authorization protocol. Users can seamlessly share select credentials from your Liferay-based website with various clients. You've seen this before: any time you see a "This site wants to access:" button (followed by a list of things like email address, friends list, etc.) from Google or Facebook, or you authorize a third-party Twitter client, that's OAuth 2.0 in action. It works by authorizing password-less access to portions of user-owned resources (such as an email address, a user profile picture, or something else from your account) and other permissioned resources.

OAuth 2.0's design encrypts all authorization transport through HTTPS, which prevents data passed between the systems from being intercepted.

### 434.1 Flow of OAuth 2.0

---

OAuth 2.0 takes advantage of web standards wherever possible: transport is encrypted with HTTPS; tokens are implemented as HTTP headers; data is passed via web services.

Here's how OAuth 2.0 works:

1. A user accesses a third-party application that supports authorization via credentials from a Liferay-based website. In the application (web or mobile), the user requests authorization via OAuth, sending the browser or app to the Liferay-based website. When using PKCE (explained below), the application also generates a code verifier and sends a code challenge that is created by applying a transformation to it.
2. The user authenticates and is shown the resources the application wants permission to access. When the user gives permission by clicking *Allow*, Liferay generates an authorization code that's sent to the application over HTTPS.
3. The application then requests a more permanent authorization token and sends the code with the request (along with the PKCE code verifier).

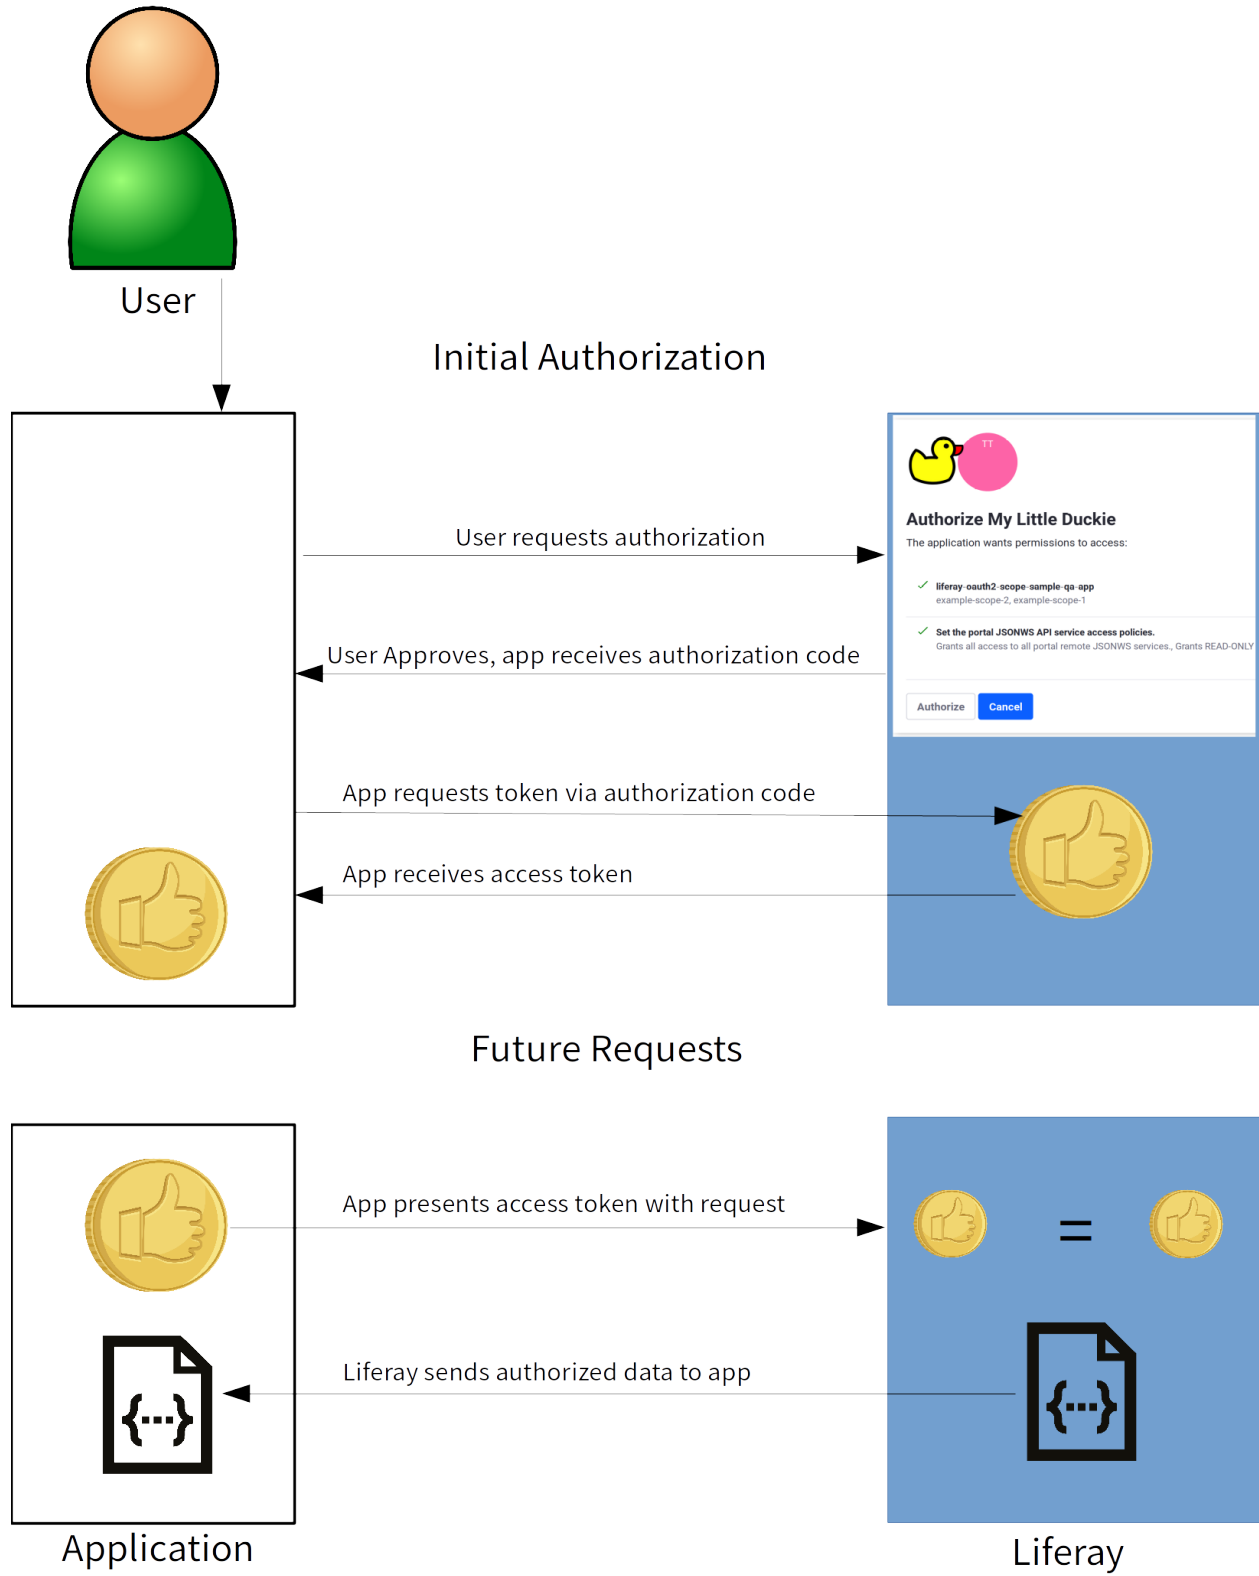


Figure 434.1: OAuth 2.0 takes advantage of web standards.

4. If the authorization code matches (and the transformed PKCE code verifier matches the previously sent code challenge), Liferay cryptographically generates an authorization token for this user and application combination. It sends the token to the application over HTTPS. Initial authorization is now complete!
5. When the application must retrieve data, it sends the token with the request to prove it's authorized to have that data.
6. Provided the token matches what Liferay has for that user and application, access is granted to retrieve the data.

That description throws around a lot of terms. Definitions provided below.

## 434.2 OAuth 2.0 Terminology

---

**Authentication:** Providing credentials so a system can verify who you are by matching those credentials with what it has stored. OAuth is not an authentication protocol.

**Authorization:** Granting access to resources stored on another system. OAuth is an authorization protocol.

**Application:** Any client (web, mobile, etc.) that must be authorized to have access to resources. Applications must be registered by administrators before users can authorize access to their resources.

**Client:** Almost synonymous with *application*, except that applications can have variants, such as web and mobile. These variants are clients.

**Client ID:** An identifier given to a client so it can be recognized.

**Client Secret:** A previously agreed-upon text string that identifies a client as a legitimate client.

**Access Token:** A cryptographically generated text string that identifies a user/client combination for access to that User's resources.

**Response Type:** OAuth 2.0 supports several response types. Pictured and described above is the most common: the authorization code. Other response types are *password* (logging in with a user name and password), and *client credentials* (headless predefined application access).

**Scope:** A list of items that define what the application wants to access. This list is sent during the initial authorization request (or otherwise defaults to scopes selected in the application registration) so users can grant or deny access to their resources.

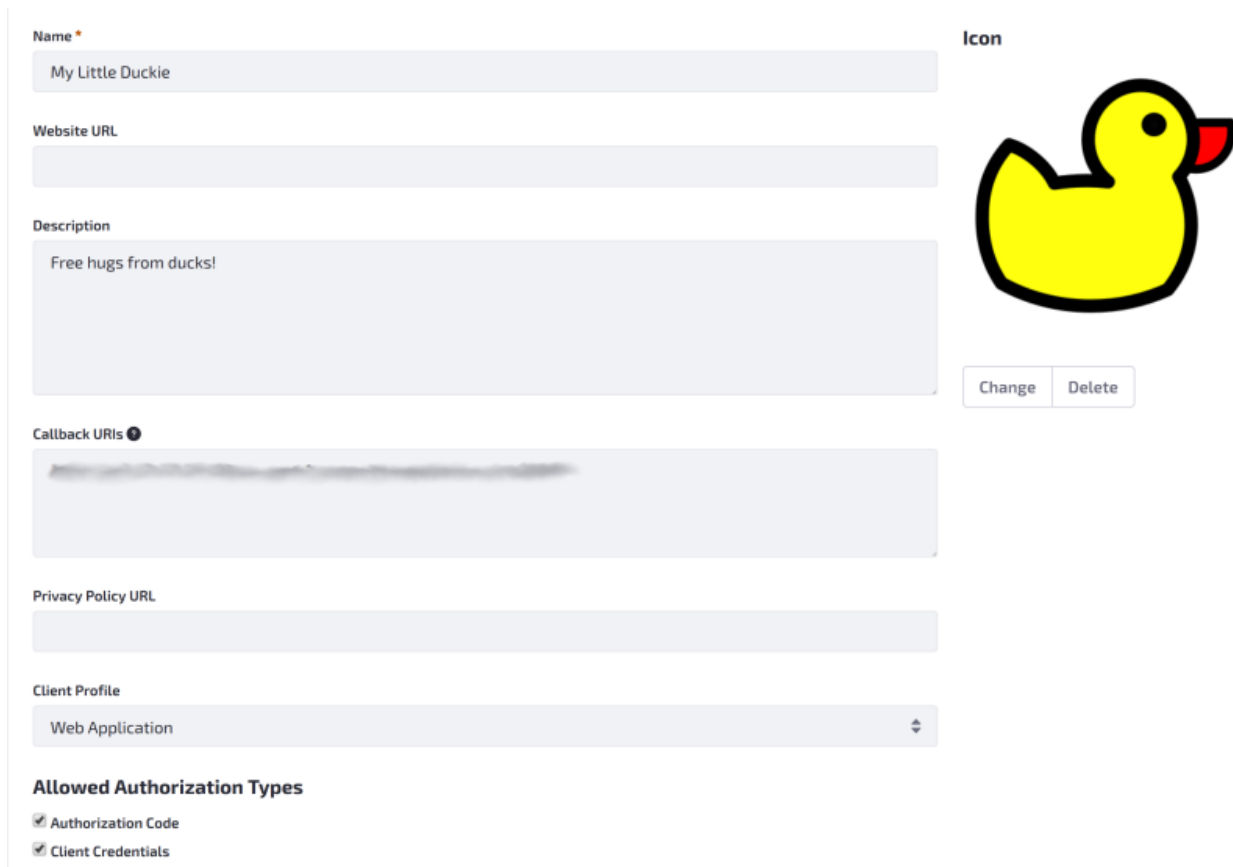
**Callback URI:** Also called a Redirection Endpoint URI. After authorization is complete, the authorization server (i.e., Liferay) sends the client to this location.

## 434.3 Creating an Application

---

When you have an application that can use OAuth 2.0 for authorization, you must register that application so Liferay DXP can recognize it. Do this by accessing *Control Panel* → *Configuration* → *OAuth2 Administration*:

1. Click the *Add* (+) button.
2. Fill out the form (description below).
3. Click *Save* to save the application.



The screenshot shows the OAuth2 Administration form in Liferay DXP. The form is titled "Name" and contains the following fields:

- Name:** My Little Duckie
- Website URL:** (empty)
- Description:** Free hugs from ducks!
- Callback URIs:** (empty)
- Privacy Policy URL:** (empty)
- Client Profile:** Web Application
- Allowed Authorization Types:** Authorization Code, Client Credentials

To the right of the form is an **Icon** section showing a yellow duck icon. Below the icon are two buttons: **Change** and **Delete**.

Figure 434.2: Adding an application registers it so users can authorize access to their data.

**Name:** Give the application a recognizable title.

**Website URL:** Add a link to the application's website.

**Callback URIs:** Enter at least one (line-separated) URI where users should be redirected after they authorize (or refuse to authorize) access to their accounts. This should link to a handler for whichever Allowed Authorization Types you support (see below).

**Client Profile:** Choose a template that filters the authorization types that are appropriate (secure) for that profile. For example, if your application is a web application, choose *Web Application*, and these authorization types are available and selected automatically: Authorization Code, Client

Credentials, Refresh Token, and Resource Owner Password Credentials. These are OAuth 2 “flows” documented in the OAuth2 RFC 6749 Standards Document. If you want to select authorization types manually, select *Other*.

**Allowed Authorization Types:** Select the defined OAuth 2 protocol flows your application supports. Several common combinations are defined for you in the various Client Profiles above.

After you save the form, it reappears with additional fields:

**Client ID:** The system generates this for you; it’s an identifier for your application, so that Liferay DXP knows what application is being authorized to access user data.

**Client Secret:** Click the *pencil* icon to generate a client secret. The secret identifies the client during the authorization process (see figure 1 above). Not all client profiles require a client secret, because some are incapable of keeping it secret! This is when the aforementioned PKCE code challenge and verifier is needed.

**Icon:** Upload an icon that your application’s users identify with your application. This is displayed on the authorization screen.

**Privacy Policy URL:** Add a link to your application’s privacy policy.

**Token Introspection:** Allow your application to retrieve metadata from the token by requesting it from Liferay DXP. This implements RFC 7662.

Excellent! Now you know how to add OAuth2 authorization for your application to Liferay DXP! Next, you must define scopes of user data the application can access.



---

## OAuth2 Scopes

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

In OAuth 2.0, applications are granted access to limited subsets of user data. These are called *scopes* (not to be confused with Liferay scopes). They are created in two ways:

1. By administrators, by creating a Service Access Policy for the scope
2. By developers, by creating a JAX-RS endpoint. By default, scopes are generated based on the HTTP verbs supported by the JAX-RS endpoint. A special annotation overrides this behavior and registers specific scopes.

### 435.1 Creating a Scope for a JSONWS Service

---

The most common way to create a scope is to create a Service Access Policy prefixed with the name `OAuth2_`. This naming convention causes the policy to appear in the OAuth application configuration screen as a scope.

For example, say the application needs access to a user's profile information to retrieve the email address. To grant the application access to this, go to *Control Panel* → *Configuration* → *Service Access Policy*, and create the policy pictured below.

Note that the policy is not a default policy, and that it grants access only to one method in the `UserService`. This is a JSONWS web service generated by Service Builder. You can view a list of all available services in your installation at this URL:

`http://[host]:[port]/api/jsonws/`

Once you create a policy and name it with the `OAuth2_` prefix, it appears in the *Scopes* tab in OAuth2 Administration.

Now you can select it and save your application.

The screenshot shows a configuration page for a Service Access Policy. The fields are as follows:

- Name \***: OAUTH2\_userprofile
- Enabled**: YES (toggle switch)
- Default**: NO (toggle switch)
- Title \***: Read your personal user data. (with a language selector for en-US)
- Allowed Service Signatures**:
  - Service Class**: com.liferay.portal.kernel.service.UserService
  - Method Name**: getCurrentUser

At the bottom, there is a blue **Save** button and a **Switch to Advanced Mode** button.

Figure 435.1: A Service Access Policy defines a scope for OAuth 2.0 applications.

## 435.2 Creating the Authorization Page

This step is optional. Users need an interface to authorize access to their accounts, and one is provided automatically. If, however, you want to customize the page, you can create an authorization page in your Site.

1. Go to *Control Panel* → *System Settings* → *Security* → *OAuth2*. Click the bottom item on the left, labeled *Authorize Screen*.
2. Two defaults appear. The first is the URL to the authorize page. By default, it's `/group/guest/authorize-oauth2-application`. This corresponds to the default site's URL and a page on that site called `authorize-oauth2-application`.
3. If you have customized the name and URL of your default site, make the appropriate change here so the URL matches the page you'll create in that site next. Click *Save*.
4. Go to your Site's *Build* → *Pages* screen. Click the **+** button and choose *Private Page*. This forces users to log in.
5. Choose the *Full Page Application* type.
6. Give the page the same name you configured in step 2.
7. Uncheck the box labeled *Add this Page to the following Menus:*. You don't want this page showing up in your Site navigation.



Resource scopes    Global scopes

Set the portal JSONWS API service access policies.

- Grants READ-ONLY access to all portal remote JSONWS services.**  
everything.readonly  
[CAN be assigned via a global scope](#)
- Grants all access to all portal remote JSONWS services.**  
[MUST be assigned via a global scope](#)
- Read your personal user data.**  
userprofile

liferay-oauth2-scope-sample-qa-app

**Save**    Cancel

Figure 435.2: Scopes named with the proper prefix appear in the Scopes tab of your application configuration.

8. On the page that appears next, verify the Friendly URL matches the URL you configured in step 2.
9. Under *Full Page Application*, choose *OAuth2 Authorize Portlet*.
10. Click *Save*.

Excellent! Users can use the default or the UI of your design to go through the authorization process. Now that you have the UI and you understand scopes, it's time to make the authorization process happen in your application.

---

# AUTHORIZING ACCOUNT ACCESS WITH OAUTH2

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Once you have an application registered, you can start authorizing users. To do that, you must construct the URL to the authorization server (Liferay DXP). The authorization server asks users to authorize the requested permissions to their resources, defined as you saw in the previous tutorial as scopes.

---

## 436.1 Authorization Code Flow

---

The most common OAuth flow is the Authorization Code flow, used for web applications. The URL for this requires the following request parameters:

- `response_type`
- `client_id`

To construct a URL for this authorization, follow this pattern:

```
https://[hostname]/o/oauth2/authorize?response_type=code&client_id=[client ID]
```

The client ID comes from registering the application. It's automatically generated (you can change it if you edit the application).

**IMPORTANT:** Sometimes the phrase “web application” is used loosely, implying applications where the above URL is requested from the web browser directly. If this happens, you'd leak the client secret, compromising the security of the grant flow and the application. In such cases, select the “User Agent Application” client profile instead when registering your application. This makes a secure alternative available to your application: PKCE Extended Authorization Code flow (see below).

Once the user has authorized the requested permissions to their resources, the authorization server returns an authorization code to your application at its registered callback URI (A.K.A. redirect URI) as a query string parameter.

```
[your callback URI]?code=[authorization server generated code]
```

Your application must then exchange this authorization code for an access token by sending a POST request following this pattern:

```
http://localhost:8080/o/oauth2/token
```

With the following parameters in the body (encoded as `application/x-www-form-urlencoded`):

```
client_id=[client ID]
client_secret=[client secret]
grant_type=authorization_code
code=[authorization server generated code]
redirect_uri=[registered callback URI]
```

In the body of HTTP response to this request, you receive JSON like this:

```
{
 "access_token": "[authorization server generated access token]",
 "token_type": "Bearer",
 "expires_in": 600,
 "scope": "[the scopes that were authorized by the user]",
 "refresh_token": "[authorization server generated refresh token]"
}
```

From this you should extract and persist the access token. If you intend to use the token for an indefinite amount of time (beyond 600 seconds from the above example) you also need the refresh token. This can be used in conjunction with the Refresh Token Flow to obtain a new access token with the same permissions, without further user authorization. The authorization server only issues Refresh Tokens if your application registration is registered for this flow.

## 436.2 PKCE Extended Authorization Code Flow

---

This flow is the same as above with the addition of the Proof Key for Code Exchange (PKCE). It requires another request parameter: `code_challenge`. This flow is for clients like smartphone applications that may not have sole access to the URL (and thus the request parameters) redirected to by the authorization server after the user authorization. It protects against a malicious application on the same system authorizing itself by reading the response code. To do this, the client application sends a *code challenge* with the authorization request: a string it has generated and which it only knows. To generate this string it must first create another secret string known as the *Code Verifier*, and then apply a transformation to it. After authorization, the code verifier is sent with the authorization code, validating the client.

For more detail on how to do this, please refer to the PKCE specification.

To support this flow, you must have defined PKCE as an Allowed Authorization Type when you created the application. This is part of the Native Application and User Agent Application client profiles. To request an authorization code using PKCE, use a URL containing the `code_challenge` request parameter:

```
https://[hostname]/o/oauth2/authorize?response_type=code&client_id=[client ID]&code_challenge=[PKCE code challenge]
```

The rest of the process is identical to Authorization Code flow, except that when making the final request to get the access token, you must also provide the following parameter:

```
code_verifier=[Code Verifier that was transformed and sent as code_challenge previously]
```

### 436.3 Client Credentials and Resource Owner Flows

---

There are two other, less used flows. If you have a scenario where two servers exchange agreed upon, non user-centric data, you can bypass the Allow/Deny screen for users and authorize the client. This is called the Client Credentials flow, and you'd use this URL pattern:

```
https://[hostname]/o/oauth2/token?grant_type=client_credentials&client_id=[client ID]&client_secret=[client secret]
```

A final flow, where users trust the application with their passwords is rare, but possible. This is called the Resource Owner Password flow, and its URL pattern looks like this:

```
https://[hostname]/o/oauth2/token?grant_type=password&client_id=[client ID]&client_secret=[client secret]&username=[user@emailaddress.com]&password=
```

Users are prompted for their passwords, and upon successful log in, receive an authorization code.

### 436.4 Token Use

---

All flows above result in an access token that's sent by the authorization server (Liferay DXP) to the client application. This token is sent in the response for the client application to store and send along with any future request for data.

For example, say the authorization code 946856e2b5ddf0928f6fc55f657bab73 was sent to the client application. When the client requests data, this code must be sent in each request header. Using a command line HTTP client such as Curl, you could send a request like this:

```
curl -H 'Authorization: Bearer 946856e2b5ddf0928f6fc55f657bab73' 'https://[hostname]/o/api/sample2'
```

OAuth 2.0 provides a convenient way for client applications to be granted access to particular services (scopes) by users without sharing credential information.

### 436.5 Revoking Access

---

Once access is granted, users or administrators are free to revoke access whenever they wish. If this happens to a client, the token becomes invalid and the client must ask the user for authorization

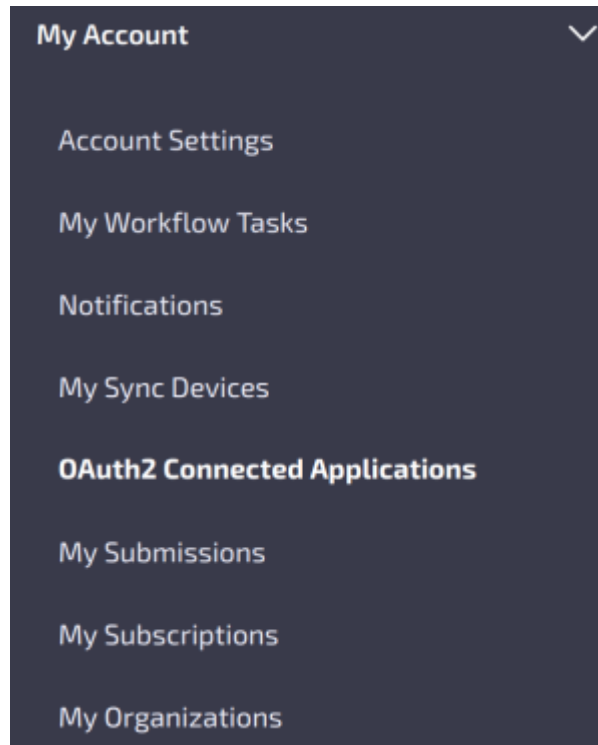


Figure 436.1: Users have complete control over what applications have access to their data in their account profiles.

again. This puts users in control of what has access to their data, and they can exercise this control at any time.

In their account areas, users can click *OAuth2 Connected Applications* and see a list of applications they've allowed to access their accounts. From here, they can revoke access by clicking the *Remove Access* item in the Action menu or the *Remove Access* button in the detail screen for the application.

Administrators can view the authorizations in the Authorizations tab of any app in *Control Panel* → *Configuration* → *OAuth2 Administration*.

|                          | User ID | User Name | Authorization | Last Access | Expiration | Remote IP  |        |
|--------------------------|---------|-----------|---------------|-------------|------------|------------|--------|
| <input type="checkbox"/> | 20139   | Test Test | 7 Days Ago    | 7 Days Ago  | 7 Days Ago | [REDACTED] | Revoke |
| <input type="checkbox"/> | 20139   | Test Test | 7 Days Ago    | 7 Days Ago  | 7 Days Ago | [REDACTED] | Revoke |
| <input type="checkbox"/> | 20139   | Test Test | 5 Days Ago    | 5 Days Ago  | 5 Days Ago | [REDACTED] | Revoke |

Figure 436.2: All authorizations for an app appear in the Authorizations tab for the app.

Clicking the *Revoke* button on any listed authorization revokes that application's access to that user's account.

**436.6 Summary**

---

OAuth 2.0 provides a complete and secure authorization flow for users, without their having to share any credential information. Once applications are created in the system, secure tokens provide access to particular scopes of information, and this access can be revoked at any time, making OAuth 2.0 a convenient method for users and developers alike to access the information they need.





## UPGRADING TO 7.0

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Upgrading to 7.0 involves migrating your installation and code (your theme and custom apps) to the new version. Here you'll learn how to upgrade your installation.

Here are the installation upgrade paths:

| Upgrade Path                                                        | Description                                                  |
|---------------------------------------------------------------------|--------------------------------------------------------------|
| Liferay Portal 5.x and 6.0.x → Liferay Portal 6.2 → Liferay DXP 7.2 | Support life ended for Liferay Portal 5.0, 5.1, 5.2, and 6.0 |
| Liferay Portal 6.1.x → Liferay DXP 7.1 → @product@ 7.2              | Support life ended for Liferay Portal 6.1                    |
| Liferay Portal 6.2+ → Liferay DXP 7.2                               |                                                              |
| Liferay DXP 7.0+ → @product@ 7.2                                    |                                                              |

**Note:** Themes and custom apps from Liferay Portal 6.0 through Liferay DXP 7.1 can be upgraded directly to Liferay DXP 7.2. See the code upgrade instructions for details.

Here are the upgrade steps:

**Note:** You can prepare a new Liferay server for data upgrade in parallel with the steps up to and including the step to upgrading the Liferay DXP data.

1. `{.root}`If You're Upgrading to Liferay Portal 6.2, Follow the Liferay Portal 6.2 Upgrade Instructions First
2. If You're Upgrading a Sharded Environment, Follow the Instructions for Upgrading It  
Upgrading a sharded installation to 7.0 requires migrating it to as many non-sharded Liferay DXP installations (servers) as you have shards.`{.summary}`

### 3. If You're a Upgrading a Cluster, Read Those Instructions First

If you're updating a cluster, read those instructions first and apply them to your upgrade.{summary}

### 4. Plan for Handling the Deprecated Applications

Every application deprecation has different ramifications. Learn how the deprecations might affect your site and decide how to replace the functionality you use from those applications.{summary}

### 5. Test Upgrading a Liferay DXP Backup Copy

Here you'll prune a backup copy of your database and upgrade the data. You'll learn how to use the upgrade tool and resolve upgrade issues. The notes and scripts you assemble as you prune and upgrade the database copy are invaluable for correctly and efficiently upgrading the Liferay DXP database you'll use with 7.0.{summary}

#### 1. Copy the Production Installation to a Test Server

You'll use the installation copy to test data changes.{summary}

#### 2. Copy the Production Database Backup

Copy the production backup to the test database and save the copy logs for analysis.{summary}

#### 3. Remove Duplicate Web Content and Structure Field Names

#### 4. Find and Remove Unused Objects

You may have intermediate versions of objects (e.g., JournalArticle objects) that you don't need. Remove them and objects that only reference them.{summary}

#### 5. Test Liferay DXP with its Pruned Database Copy

Make sure Liferay DXP continues to work successfully. If it's broken, start over with a fresh database backup and prune it more carefully.{summary}

#### 6. Install the New Liferay DXP Version on a Test Server

Install the Liferay DXP version you're upgrading to, to use its upgrade tool.{summary}

#### 7. Tune Your Database for the Upgrade

#### 8. Upgrade the Liferay Data, then Return Here

#### 9. If the Upgrade Took too Long, Prune a Fresh Database Backup More and Upgrade Its Data

#### 10. Test the Upgraded Instance

Make sure Liferay DXP continues to work successfully. If it's broken, start over with a fresh database backup and prune it more carefully.{summary}

#### 11. Checkpoint: You've pruned and upgraded your production database copy. You're ready to prepare for upgrading the production database.

### 6. Prepare to Upgrade the Liferay DXP Database

Preparing for the production database upgrade involves pruning and testing it, upgrading your Marketplace apps, publishing staged changes, and synchronizing a complete data and configuration backup.{summary}

1. Remove All Noted Unused Objects  
Remove all unused objects you noted from pruning your test database.{.summary}
  2. Test Liferay DXP
  3. Upgrade Your Marketplace Apps
  4. Publish All Staged Changes
  5. Synchronize a Complete Liferay DXP Backup  
Synchronize a complete backup of your production Liferay DXP server installation and pruned production database.{.summary}
  6. Checkpoint: You're ready to prepare a 7.0 server for upgrading a production database.
7. Prepare a New Liferay DXP Server for Data Upgrade
- Set up a production server with 7.0, configured to use your document repository and Liferay DXP database. You'll migrate your portal and system properties too. (Note, this step can be done in parallel with any of the previous steps.){.summary}
1. Request an Upgrade Patch From Liferay Support (Liferay DXP Only)
  2. Install the Liferay DXP Version You're Upgrading To
  3. Install the Latest Upgrade Patch or Fix Pack (Liferay DXP Only)
  4. Migrate Your OSGi Configurations (Liferay DXP 7.0+)
  5. Migrate Your Portal Properties  
Migrate your portal properties to your new 7.0 server.{.summary}
    1. Update Your Portal Properties  
Some of the portal properties have new values or have been removed or replaced. Update your properties for 7.0.{.summary}
    2. Convert Applicable Properties to OSGi Configurations  
Many applications are configured using OSGi Configuration (Config Admin) instead of portal properties. Convert your existing properties to their OSGi Configuration replacements.{.summary}
  6. Configure Your Documents and Media File Store  
The upgrade tool upgrades your Documents and Media file store too. Update your Documents and Media file store configuration and specify it for the upgrade tool.{.summary}
  7. Disable Indexing  
Improve the data upgrade performance by disabling indexing.{.summary}
  8. Checkpoint: You've prepared a new Liferay DXP server for executing the data upgrade
8. Upgrade the Liferay DXP data
- This section explains the data upgrade options, upgrade configuration, and the upgrade process.{.summary}
1. Tune Your Database for the Upgrade
  2. Configure the Data Upgrade  
Configure the data upgrade, including the data store and whether to automatically upgrade the modules.{.summary}

### 3. Upgrade the Core

#### 1. Run the Data Upgrade Tool

Run the data upgrade tool. Resolve any core upgrade issues.{.summary}

#### 2. Issues Upgrading to 7.0 or Lower? Restore the Database Backup

If the issues were with upgrades to Liferay 7.0 or lower, get a clean start by restoring the pruned production database backup.{.summary}

#### 3. Upgrade Your Resolved Issues

If there were issues upgrading to 7.2, resolve them and restart the data upgrade tool; continue if there were no issues.{.summary}

### 4. Upgrade the Liferay Modules

Learn how to use Gogo Shell to upgrade the Liferay modules, if you didn't upgrade them automatically with the core.{.summary}

#### 1. Upgrade Modules that are Ready for Upgrade

Discover which modules are ready for upgrade and upgrade them.{.summary}

#### 2. Check Module Upgrade Status and Resolve Any Module Upgrade Issues

#### 3. Checkpoint: You've completed upgrading the Liferay data. It's time to get your server ready for production.{.summary}

### 9. Execute the Post-Upgrade Tasks

Now that your database is upgraded, clean up remnants of upgrading by restoring your database optimizations, enabling and regenerating your search indexes, and more.{.summary}

#### 1. Remove the Database Tuning

#### 2. Re-enable and Re-Index the Search Indexes

#### 3. Update Web Content Permissions (7.0 and lower)

#### 4. Address Any Deprecated Apps That Still Need Handling

### 10. Checkpoint: You've completed the upgrade and post-upgrade tasks

Follow the steps above to upgrade your Liferay DXP installation to 7.0. They link upgrade topic details to help complete a safe, successful upgrade.

## PLANNING FOR DEPRECATED APPLICATIONS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

As Liferay DXP evolves so do Liferay applications. They may, for example, be deprecated in favor of newer and better Liferay applications. The deprecations might call for migrating application data to a new application or completely removing the application and data. Before upgrading, examine the deprecations:

- 7.2 deprecations
- 7.1 deprecations

Determine how and when to address the deprecations in your upgrade plan.



---

## TEST UPGRADING A LIFERAY DXP BACKUP COPY

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Before upgrading your production Liferay instance, you should do a trial run (even multiple runs) to make sure that you upgrade successfully and efficiently. Here's the process:

- **Preparing a test server and database:** This involves copying your current production installation to a test server and copying your production data backup to a test database. After you prune data from the test database (next step) you'll test against it.
- **Pruning the database:** Free your database of duplicate and unused objects. By removing them you can reduce upgrade time and improve your server's performance.
- **Upgrading your test server and database:** First you'll optimize your database for the data upgrade. Taking time to do this can save upgrade time. Then you'll do an upgrade test run (or several test runs) on a the pruned database copy. After going through the upgrade process, resolving any issues, and testing the upgraded server successfully, you can confidently upgrade your production database.

---

**Tip:** These steps and preparing a new Liferay DXP server can be done in parallel to save time.

---

Now prepare your test environment.

---

### 439.1 Preparing a Test Server and Database

---

Using a new separate server and database let's you safely test upgrading.

---

### 439.2 Copy the Production Installation to a Test Server

---

Prepare a test server to use a copy of your production installation. Your test server must use the same Liferay version you're using on production. Configure your server to use a new empty database for testing data upgrades.

### 439.3 Copy the Production Backup to the Test Database

---

Import data from your production database backup to the new empty database.

---

**Important:** Make sure to save the data import log—you'll examine it in the next steps.

---

Next you'll prune your database of unneeded data.



---

## PRUNING THE DATABASE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Accumulating unneeded site data is common. For example, you may have many unused versions of Web Content articles or Documents and Media files. If you're done revising them and don't need the intermediate revisions, you can remove them. This saves you space and upgrade time. Here you'll remove unneeded data and then test your server.

---

### 440.1 Remove Duplicate Web Content Structure Field Names

---

If you've used Web Content Management extensively, you might have structures without unique field names. Find and remove duplicate field names before upgrading. If you upgraded to Liferay Portal 6.2 previously and skipped doing this, you'll encounter this error:

```
19:29:35,298 ERROR [main][VerifyProcessTrackerOSGiCommands:221] com.liferay.portal.verify.VerifyException: com.liferay.dynamic.data.mapping.validator.com.liferay.portal.verify.VerifyException: com.liferay.dynamic.data.mapping.validator.DDMFormValidationException$MustNotDuplicateFieldName: The field
```

If this is the case, roll back to your previous backup of Liferay Portal 6.2 and find and remove duplicate field names.

---

### 440.2 Find and Remove Unused Objects

---

In the UI or using database queries, identify unused objects. Then remove them via Liferay's UI or using Liferay's API through the script console or a portlet you create.

---

**Important:** You should only use Liferay's UI or API because they account for relationships between Liferay DXP objects.

Never use SQL directly on your database to remove records. Your SQL might miss object relationships, orphaning objects and causing performance problems.

---

Here are some common places to check for unused objects.

### 440.3 Objects From the Large/Populated Tables

---

Table rows are mapped to Liferay DXP objects. Large tables with many records might contain lots of unused objects. The greater the table size and the records per table, the longer upgrading takes.

Finding and removing unused objects associated with such tables reduces upgrade times. Your data import log (from the previous step) can provide valuable table information. Database engines show this information in different ways. Your database import log might look like this:

```
Processing object type SCHEMA\EXPORT\TABLE\TABLE_DATA
imported "LIFERAY"."JOURNALARTICLE" 13.33 GB 126687 rows
imported "LIFERAY"."RESOURCEPERMISSION" 160.9 MB 1907698 rows
imported "LIFERAY"."PORTLETPREFERENCES" 78.13 MB 432285 rows
imported "LIFERAY"."LAYOUT" 52.05 MB 124507 rows
imported "LIFERAY"."ASSETENTRY" 29.11 MB 198809 rows
imported "LIFERAY"."MBMESSAGE" 24.80 MB 126185 rows
imported "LIFERAY"."PORTALPREFERENCES" 4.091 MB 62202 rows
imported "LIFERAY"."USER__" 17.32 MB 62214 rows
...
```

Several items stand out in the example database import:

- The JOURNALARTICLE table makes up 98% of the database size.
- There are many RESOURCEPERMISSION records.
- There are many PORTLETPREFERENCES records.

Search for unused objects associated with the tables that stand out and use Liferay's API (e.g., the UI or script console) to delete the objects.

### 440.4 Common Object Types Worth Checking

---

Some object types should be checked for unused objects. Here are some reasons for checking them:

- Removing them frees related unused objects for removal
- They're version objects that aren't worth keeping

Check these object types:

- **Sites:** Remove sites you don't need. When you remove a site, remove its related objects:
  - Layouts
  - Portlet preferences
  - File entries (document library objects)
  - Asset Entries
  - Tags
  - Vocabularies and categories
  - Expando fields and their values
  - ResourcePermission objects
  - (and everything else)
- **Instances:** Unused instances are rare, but since they are the highest object in the hierarchy, removing their objects can optimize upgrades considerably:
  - Sites (and all their related content)
  - Users
  - Roles
  - Organizations
  - Global ResourcePermission objects
- **Intermediate web content versions:** Liferay DXP generates a new web content version after any modification (including translations). Consider removing versions you don't need. Removing a Journal Article, for example, also removes related objects such as image files (JournalArticleImage) that are part of the content. Removing unneeded image files frees space in your database and file system. For more details, see [Example: Removing Intermediate Journal Article Versions](#).
- **Document versions:** As with Journal Articles, if you don't need intermediate document versions, delete them. This saves space both in the database and on the file system, space that no longer needs to be upgraded.
- **Layouts:** Layouts are site pages, and they affect upgrade performance because they relate to other entities such as portlet preferences, permissions, assets, ratings, and more. Remove unneeded layouts.
- **Roles:** Remove any Roles you don't need. Deleting them also deletes related ResourceBlockPermission and ResourcePermission objects.
- **Users:** If you have Users that aren't active anymore, remove them.

- **Vocabularies:** Remove any unused vocabularies. Note that removing a vocabulary also removes its categories.
- **Orphaned data:** Check for unused objects that are not connected to anything. Here are some examples:
  - DLFileEntries with no file system data.
  - ResourcePermission objects associated to a Role, Layout, User, portlet instance, etc. that no longer exists.
  - PortletPreference objects associated with a portlet or layout that no longer exists. This is common in environments with many embedded portlets. These portlet instances have a different lifecycle and aren't deleted when the portlet is removed from a template.

If you want to see an example of removing intermediate object versions, read [Example: Removing Intermediate Journal Article Versions](#) and then return here.

Next, you'll test Liferay DXP with its pruned database.

#### 440.5 Test with the Pruned Database Copy

---

Find and resolve any issues related to the objects you removed. You can always restart pruning a new copy of your production database if you can't resolve an issue.

---

**Warning:** the upgrade to Liferay DXP 7.2 moves Web Content images to the Document Library and then deletes their former table `JournalArticleImage`. Make sure the images show in the upgraded Web Content articles.

---

Once you've successfully tested Liferay DXP with its pruned database copy, you can upgrade the database to 7.0.

---

## EXAMPLE: REMOVING INTERMEDIATE JOURNAL ARTICLE VERSIONS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

These instructions and code samples demonstrate removing intermediate Journal Article versions. In the script console, you can remove unneeded object versions by executing Java or Groovy code.

Here are example steps for removing intermediate Journal Article versions:

1. Decide how many of the latest versions to keep. You must keep the original version and the most recent version, but you may keep older recent versions too. For example, you may want to keep the two latest versions or just the latest.
2. Find a method for deleting the entity versions. Liferay DXP app APIs and `com.liferay.portal.kernel` API are available at `@platform-ref@`.

If it's a Service Builder entity, examine the `delete*` methods in the entity's `*LocalServiceUtil` class. For example, this `deleteArticle` in `JournalArticleLocalServiceUtil` deletes a Journal Article version:

```
deleteArticle(long groupId, java.lang.String articleId, double version,
 java.lang.String articleURL,
 com.liferay.portal.kernel.service.ServiceContext serviceContext)
```

3. Aggregate the entity versions to delete and the information required to delete them. For example, get all the Journal Article versions in range that match your removal criteria and associate their entity IDs and group IDs with them—the `deleteArticle` method requires the entity ID and group ID.

The entity object (e.g., `JournalArticle`) typically has a version field. `JournalArticleResource` has each Journal Article's article ID (the entity's ID) and group ID. `JournalArticleResource` is our key to getting each `JournalArticle`, which can have multiple versions. Here are steps for identifying the Journal Article versions to delete:

1. Get all the `JournalArticleResource` objects.

```
List<JournalArticleResource> journalArticleResources =
 JournalArticleLocalServiceUtil.getJournalArticleResources(start, end);
```

2. Get each Journal Article version's workflow status via the JournalArticle object associated with each JournalArticleResource. Dynamic Query is an efficient way to get exactly the data you want (and nothing more) from each object.

```
for (JournalArticleResource
 journalArticleResource : journalArticleResources) {

 List<Double> journalArticlesVersionsToDelete =
 new ArrayList<Double>();

 DynamicQuery dq =
 DynamicQueryFactoryUtil.forClass(JournalArticle.class)
 .setProjection(ProjectionFactoryUtil.projectionList()
 .add(ProjectionFactoryUtil.property("id"))
 .add(ProjectionFactoryUtil.property("version"))
 .add(ProjectionFactoryUtil.property("status")))
 .add(PropertyFactoryUtil.forName("groupId")
 .eq(journalArticleResource.getGroupId()))
 .add(PropertyFactoryUtil.forName("articleId")
 .eq(journalArticleResource.getArticleId()))
 .addOrder(OrderFactoryUtil.asc("version"));

 List<Object[]> result =
 JournalArticleLocalServiceUtil.dynamicQuery(dq);

 // See the next step for the sample code that goes here
}
```

3. For each JournalArticleResource (there's one for each Journal Article entity), build a list of intermediate versions in range of the first or latest versions you want to keep and whose status qualifies them for deletion. For example, you may want to delete intermediate article versions that are approved or expired (i.e., WorkflowConstants.STATUS\_APPROVED or WorkflowConstants.STATUS\_EXPIRED). The MIN\_NUMBER\_FIRST\_VERSIONS\_KEPT and MIN\_NUMBER\_LATEST\_VERSIONS\_KEPT variables here mark the minimum and maximum number of first (oldest) and latest (newest) versions to keep.

```
List<Double> journalArticlesVersionsToDelete =
 new ArrayList<Double>();

for (int i=0; i < result.size(); i++) {
 long id = (long) result.get(i)[0];
 double version = (double) result.get(i)[1];
 int status = (int) result.get(i)[2];

 if ((status == WorkflowConstants.STATUS_APPROVED) || (status == WorkflowConstants.STATUS_EXPIRED) {

 if (i < MIN_NUMBER_FIRST_VERSIONS_KEPT) {
 continue;
 }

 if (i ≥ (result.size() -
 MIN_NUMBER_LATEST_VERSIONS_KEPT)) {
 continue;
 }

 journalArticlesVersionsToDelete.add(version);
 }
}

// See the next step for the sample code that goes here
```

4. Lastly, delete each Journal Article matching the versions you aggregated.

```
for (double version : journalArticlesVersionsToDelete) {
 {
 JournalArticleLocalServiceUtil.deleteArticle(journalArticleResource.getGroupId(),
 journalArticleResource.getArticleId(),
 journalArticlesVersionsToDelete(i), null, null);
 }
}
```

You can write similar code to remove intermediate versions of other entities.

---

**Tip:** Print the version (and any other information of interest) of each object you're removing. You can also comment out the object deletion call and read the printout of versions to be removed as a test before committing to deleting them.

---

After you've pruned your database, test it with Liferay DXP.





---

## UPGRADING YOUR TEST SERVER AND DATABASE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

After you've pruned your database and tested it successfully, it's ready for upgrade. Here you'll install 7.0 and migrate your current installation files to it and upgrade them. Then you'll optimize your database for the upgrade and upgrade your data. Lastly, you'll test this upgraded test environment. You may run into issues that require you to start again with backup of your pruned database. After you're satisfied with the test upgrade, you can prepare for upgrading production. Start with preparing 7.0 on a test server.

---

### **442.1 Install Liferay on a Test Server and Configure It to Use the Pruned Database**

---

Prepare a new test server with 7.0. Configure it to use the pruned database copy—keep the original backup in case you want to restart test upgrades on a copy of it. You'll use the new test server's Liferay upgrade tool next.

---

### **442.2 Tune Your Database for the Upgrade**

---

Tune your database for the upgrade.

---

### **442.3 Upgrade the Database**

---

Upgrade the database to 7.0 (see Upgrade the Database); then return here.

If the upgrade took too long, search the upgrade log to identify more unused objects. Then retry these steps with a fresh copy of the production database.

#### **442.4 Test the Upgraded Portal and Resolve Any Issues**

---

Test this upgraded 7.0 instance and resolve any issues. If you can't resolve an issue, retry these steps with a fresh copy of the production database.

#### **442.5 Checkpoint: You've Pruned and Upgraded a Production Database Copy**

---

By removing unused objects from Liferay DXP in your test environment, you've made upgrading feasible to do in production. You identified unused objects, documented/scripted removing them, and successfully upgraded the Liferay DXP database copy.

It's time to prepare your production environment for upgrading.

# PREPARING TO UPGRADE THE LIFERAY DXP DATABASE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

After testing the upgrade on a copy of your production database, you can apply what you learned to your production database.

---

**Tip:** This step and preparing a new Liferay DXP server can be done in parallel to save time.

---

## **443.1 Remove All Unused Objects You Identified Earlier**

---

Previously you identified and removed unused objects from a copy of your Liferay DXP production database backup. In the same way (in the script console or UI) you removed the unused objects from the backup, remove them from your pre-upgrade production database.

## **443.2 Test Using the Pruned Database**

---

Find and resolve any issues related to the objects you removed. By removing the objects from production and testing your changes before upgrading, you can more easily troubleshoot issues, knowing that they're not related to upgrade processes.

### **443.3 Upgrade Your Marketplace Apps**

---

Upgrade each Marketplace app (Kaleo, Calendar, Notifications, etc.) that you're using to its latest version for your Liferay DXP installation. Before proceeding with the upgrade, troubleshoot any issues regarding these apps.

### **443.4 Publish all Staged Changes to Production**

---

If you have local/remote staging enabled and have content or data saved on the staged site, publish it to the live site. If you skip this step, you must run a full publish (or manually publish changes) after the upgrade, since the system won't know what content changed since the last publishing date.

### **443.5 Synchronize a Complete Backup**

---

Completely back up your Liferay DXP installation, pruned production database, and document repository.

It's time to prepare a new Liferay DXP server.

# PREPARING A NEW LIFERAY DXP SERVER FOR DATA UPGRADE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

To upgrade your Liferay DXP database, prepare a new server for hosting 7.0. You'll use this server to run the database upgrade and run 7.0. Then you can run your production server while you're configuring a new server to host 7.0 exactly the way you want.

---

**Note:** these steps can be done in parallel with any of the upgrade preparation steps: planning for deprecated apps, testing upgrades on a Liferay DXP backup copy, or preparing to upgrade the @product@ database.

---

Get the latest fixes for 7.0 by requesting an upgrade patch.

---

## 444.1 Request an Upgrade Patch from Liferay Support (Liferay DXP only)

---

An *upgrade patch* contains the latest fix pack and hot fixes planned for the next service pack. Upgrade patches provide the latest fixes available for your data upgrade.

---

## 444.2 Install Liferay

---

Install Liferay DXP on your application server or use Liferay DXP bundled with your application server of choice.

---

**Important:** Do not start your application server. It's not ready to start until after the Liferay DXP database upgrade.

---

### **444.3 Install the Latest Upgrade Patch or Fix Pack (Liferay DXP only)**

---

Install the upgrade patch (if you requested it from Liferay Support) or the latest Fix Pack.

### **444.4 Migrate Your OSGi Configurations (7.0+)**

---

Copy your OSGi configuration files (i.e., .config files) to your new server's [Liferay Home]/osgi/configs folder.

### **444.5 Migrate Your Portal Properties**

---

It is likely that you have overridden portal properties to customize your installation. If so, you must update the properties files (e.g., portal-setup-wizard.properties and portal-ext.properties) to 7.0. For features that use OSGi Config Admin, you must convert your properties to OSGi configurations. As you do this, you must account for property changes in all versions of Liferay DXP since your current version up to and including 7.0. Start with updating your portal properties.

### **444.6 Update Your Portal Properties**

---

If you're coming from a version prior to Liferay Portal 6.2, start with these property-related updates:

- If you're on Liferay Portal 6.1, adapt your properties to the new defaults that Liferay Portal 6.2 introduced.
- If you're on Liferay 6.0.12, migrate the Image Gallery.
- If you have a sharded environment, configure your upgrade to generate a non-sharded environment.

- Liferay’s image sprite framework is deprecated as of 7.2 and is disabled by default. The framework requires scanning plugins for image sprites. If you don’t use the framework, there’s no need for it to scan for images sprites. If you use the framework, enable it by overriding the default `sprite.enabled` portal property (new in 7.2) value with the following setting in a `portal-ext.properties` file:

```
sprite.enabled=true
```

---

**Note:** You can build image sprites using any framework you like and deploy them in your plugins.

---

When a new version of Liferay DXP is released, there are often changes to default settings, and this release is no different. If you rely on the defaults from your old version, you should review the changes and decide to keep the defaults from your old version or accept the defaults of the new.

Because no existing properties changed from 7.1 to 7.2, here’s a list of the 6.2 properties that have changed in 7.2:

```
users.image.check.token=false
organizations.types=regular-organization,location
organizations.rootable[regular-organization]=true
organizations.children.types[regular-organization]=regular-organization,location
organizations.country.enabled[regular-organization]=false
organizations.country.required[regular-organization]=false
organizations.rootable[location]=false
organizations.country.enabled[location]=true
organizations.country.required[location]=true
layout.set.prototype.propagate.logo=true
editor.wysiwyg.portal-web.docroot.html.taglib.ui.discussion.jsp=simple
web.server.servlet.check.image.gallery=true
blogs.trackback.enabled=true
discussion.comments.format=bbcode
discussion.max.comments=0
dl.file.entry.thumbnail.max.height=128
dl.file.entry.thumbnail.max.width=128
```

This property was removed:

```
organizations.children.types[location]
```

The latest portal properties reference provides property details and examples. Some properties are replaced by OSGi configurations.

---

## 444.7 Convert Applicable Properties to OSGi Configurations

Properties in modularized features have changed and must now be deployed separately in OSGi configuration files (OSGi Config Admin).

Use the `blade upgradeProps` command to scan your `portal-ext.properties` file to discover which properties are now set via OSGi Config Admin. You can also check the upgrade log from previous attempts for traces like these:

```
2019-03-09 17:05:17.678 ERROR [main][VerifyProperties:161] Portal property "layout.first.pageable[link_to_layout]" is obsolete
2019-03-09 17:05:17.679 ERROR [main][VerifyProperties:136] Portal property "journal.article.check.interval" was modularized to com.liferay.journal.w
```

---

**Tip:** The Control Panel’s *Configuration* & *System Settings* screens are the most accurate way to create `.config` files. Use them to export a screen’s configuration to a `.config` file.

---

#### 444.8 Update Your Database Driver

---

Install the recommended database driver and update your database connection driver specified in your `portal-ext.properties`. See the Database Templates.

#### 444.9 Configure Your Documents and Media File Store

---

General document store configuration (e.g., `dl.store.impl=[File Store Impl Class]`) continues to be done using `portal-ext.properties`. But here’s what’s changed for document storage:

- Store implementation class package names changed from `com.liferay.portlet.documentlibrary.store.*` in Liferay Portal 6.2 to `com.liferay.portal.store.*` in Liferay DXP 7.0+. Make sure your `portal-ext.properties` file sets `dl.store.impl` in one of these ways:

```
dl.store.impl=com.liferay.portal.store.file.system.FileSystemStore
dl.store.impl=com.liferay.portal.store.db.DBStore
dl.store.impl=com.liferay.portal.store.file.system.AdvancedFileSystemStore
dl.store.impl=com.liferay.portal.store.s3.S3Store
```

- JCR Store was deprecated in Liferay DXP 7.0. The Document Repository Configuration documentation describes other store options. Migrate to a supported document store before upgrading your data.
- CMIS Store was deprecated since 7.0.10 Fix Pack 14 and was removed in Liferay DXP 7.2. The Document Repository Configuration documentation describes other store options. Migrate to a supported document store before upgrading your data.
- Since Liferay DXP 7.0, document store type-specific configuration (e.g., specific to Simple File Store, Advanced File Store, S3, etc.) is done in the Control Panel at *Configuration* & *System Settings* & *File Storage* or using OSGi configuration files (`.config` files). Type specific configuration is no longer done using `portal-ext.properties`.

For example, these steps to create a `.config` file specifying a root file location for a Simple File Store or Advanced File Store:



1. Create a `.config` file named after your store implementation class.

Simple File Store: `com.liferay.portal.store.file.system.configuration.FileSystemStoreConfiguration.config`

Advanced File Store: `com.liferay.portal.store.file.system.configuration.AdvancedFileSystemStoreConfigura`

2. Set the following `rootDir` property and replace `{document_library_path}` with your file store's path.

```
rootDir="{document_library_path}"
```

3. Copy the `.config` file to your `[Liferay Home]/osgi/configs` folder.

The Document Repository Configuration provides more document store configuration details.

## 444.10 Configure Kerberos in place of NTLM

---

If you're using NTLM to authenticate Microsoft Windows™ accounts with Liferay DXP, switch to using Kerberos. Security vulnerabilities persist with NTLM. NTLM has been deprecated and removed from the bundle, but you can still build and deploy the module.

## 444.11 Disable Indexing

---

Before starting the upgrade process in your new installation, you must disable indexing to prevent upgrade process performance issues that arise when the indexer attempts to re-index content.

To disable indexing, create a file called `com.liferay.portal.search.configuration.IndexStatusManagerConfigurati` in your `[Liferay Home]/osgi/configs` folder and add the following content:

```
indexReadOnly="true"
```

After you complete the upgrade, re-enable indexing by removing the `.config` file or setting `indexReadOnly="false"`.

Your new 7.0 server is ready for upgrading your database.



---

## UPGRADING THE LIFERAY DXP DATA

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Now you're ready to upgrade the Liferay DXP data. The upgrade processes update the database schema for the core and your installed modules. Verification processes test the upgrade. Configured verifications for the core and modules run afterwards, but can be run manually too.

Here are the ways to upgrade:

- **Upgrade everything in one shot:** Use the upgrade tool to upgrade the core and all the modules.
- **Upgrade the core and the modules separately:** Use the upgrade tool (recommended) or Gogo shell to upgrade the core. Then use Gogo shell to upgrade each module.

If you are upgrading from Liferay Portal 6.2 or earlier, use the upgrade tool to upgrade everything. It's the easiest, most comprehensive way to upgrade from those versions. Since version 7.0, however, Liferay DXP's modular framework lets you upgrade modules—even the core—individually. A helpful practice for large databases is to focus first on upgrading the core and your most important modules; then back up your database before continuing upgrades. Upgrading is a flexible process that adjusts to your preferences.

---

**Note:** Liferay enterprise subscribers can use the upgrade tool to execute upgrades for fix packs. Since Liferay DXP 7.1, a fix pack's micro upgrade processes in the core (database schema micro version changes) are not mandatory. This means you can install a fix pack (i.e., core code) without having to execute the database schema micro version changes. You can execute micro version changes when you want, even outside of major or minor version upgrades. Before using the upgrade tool to execute a fix pack's micro upgrade process, however, you must shut down the server, install the fix pack, and back up the Liferay DXP database, installation, and Document Library store.

Module micro database schema version changes in fix packs execute automatically on server startup unless the `autoUpgrade` setting is `false` (the default is `true`).



---

## TUNING FOR THE DATA UPGRADE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Upgrading impacts the database differently from daily running in production. Because of this, you should tune your database for the upgrade process before you run it, and then re-apply your production settings after the upgrade completes.

- Data upgrades execute many more update statements (INSERT, UPDATE, and DELETE) and less SELECT statements than production instances. When upgrading, tune your database for executing updates.
- Data upgrades should be done in safe environments completely separate from production servers and should use database backup copies. If upgrade errors occur or you make mistakes, they don't impact production, and you can always restart using your database backup copy.

The data upgrade tuning instructions given here are a starting point for tuning your Liferay DXP data upgrade. They account for data upgrade activities and a safe data upgrade environment:

- Deactivate data integrity measures that impact performance. Restore the backup if failures occur.
- Make commit-related transaction I/O operations asynchronous.
- Increase the interval to flush commits to disk.
- Minimize transaction logging.

---

**Note:** These options worked well for us on specific versions of each database. Please consult your database vendor's documentation for information on how to optimize executing updates on your specific database version.

---

**Important:** Test your database configuration to determine tuning that's best for your system, and consult your DBA as appropriate. **Never** use database upgrade configurations in production.

Always restore your production database settings before starting your Liferay DXP server for production use with the database.

---

**Warning:** Some database properties and configurations are global and affect schemas in the same database.

---

These configurations were optimal for upgrading data in a Liferay 6.2 EE installation that had these characteristics:

- 3.2 GB database
- 15 GB Document Library
- Content translated in 3 languages
- Record count for most populated entities:
  - 1,694,000 rating entries
  - 1,605,000 permissions (ResourcePermission objects)
  - 871,000 assets (AssetEntry objects)
  - 400,000 users
  - 400,000 sites (Group objects)
  - 402,000 images
  - 259,000 message forum threads and posts
  - 200,000 documents
  - 193,000 portlet preferences
  - 103,000 web content pieces (JournalArticle objects)
  - 50,600 pages
  - 3,276 journal article images
  - 3,100 document folders

Start with configuring the database upgrade tool's Java process.

---

### 446.1 Tuning the Database Upgrade Java Process

---

Make sure to provide adequate memory for the database upgrade tool's Java process. 15GB was appropriate for the test scenario. Also make sure to set the file encoding to UTF-8 and the time zone to GMT. Here are the Java process settings:

- Xmx 15 GB RAM
- File encoding UTF-8
- User time zone GMT

Here is the `db_upgrade.sh` command:

```
db_upgrade.sh -j "-Xmx15000m -Dfile.encoding=UTF-8 -Duser.timezone=GMT"
```

It's time to tune your database transaction engine.

## **446.2 Tuning the Database Transaction Engine for Executing Updates**

---

Many more update statements are executed during data upgrade than in production. Here's how to optimize each database's transaction engine for the updates.

### **446.3 IBM DB2**

---

Please consult IBM's official DB2 documentation.

### **446.4 MariaDB**

---

In addition to the default database configuration, turn off InnoDB double-write.

### **446.5 Microsoft SQL Server**

---

In addition to the default database configuration, set transaction durability to FORCED.

### **446.6 MySQL**

---

In addition to the default database configuration, turn off InnoDB double-write.

### **446.7 Oracle Database**

---

The default configuration works well. It configures asynchronous I/O to disk automatically.

### **446.8 PostgreSQL**

---

In addition to the default database configuration, turn off synchronous commits.

### **446.9 Tuning the Database Transaction Log**

---

In production, transaction logs mark safe states to roll back to. In data upgrades, however, the safe state is the original data backup. Since transaction logging is insignificant for data upgrades, it should be disabled or minimized. Here are log tuning instructions for each database.

### **446.10 IBM DB2**

---

Please consult IBM's official DB2 documentation.

### **446.11 MariaDB**

---

In addition to the default database configuration, set the InnoDB flush log at transaction commit to 0.

### **446.12 Microsoft SQL Server**

---

Use the default database configuration.



**446.13 MySQL**

---

In addition to the default database configuration, set the InnoDB flush log at transaction commit to 0.

**446.14 Oracle Database**

---

Use the default database configuration.

**446.15 PostgreSQL**

---

In addition to the default database configuration, Set the write ahead log writer delay to 1000 milliseconds.

Congratulations! You have a starting point to plan your own Liferay DXP data upgrade project. Remember, optimal tuning depends on your data, infrastructure conditions, and database vendor. Analyze your data, tune for upgrade, and time your test upgrades. Use this information to determine the best database and Java process configuration for your Liferay DXP data upgrade.



---

## CONFIGURING THE DATA UPGRADE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The upgrade tool provides the easiest way to upgrade the core and installed modules. You can use text files or the tool's command line interface to configure your upgrade. The upgrade tool can upgrade everything—the core and all the modules—together or separately.

7.0 bundles include the upgrade tool. If you installed @product-ver@ manually, you can download the upgrade tool separately.

- *Liferay DXP 7.2*: Go to the *Downloads* page and select the *DXP 7.2* product and the *Product/Service Packs* file type. In the listing that appears, click *Download* for the *Liferay DXP Upgrade Client*.
- *Liferay Portal CE 7.2*: Go to the *Downloads* page and select *Download* for *Liferay Portal Tools for 7.2*.

Before starting the data upgrade process, configure the upgrade tool for the core upgrade and specify whether the upgrade tool should upgrade non-core module data automatically.

### 447.1 Configuring the Core Upgrade

---

The core upgrade requires configuration. You can configure it at runtime via the command line interface or pre-configure it in these files in [Liferay Home]/tools/portal-tools-db-upgrade-client/:

- `app-server.properties`: Specifies the server's location and libraries.
- `portal-upgrade-database.properties`: Configures the database connection.
- `portal-upgrade-ext.properties`: Sets the rest of the portal properties that the upgrade requires. You might want to copy your current portal properties (except your database properties) into this file. Before copying your current properties, make sure you've updated the portal properties for 7.0.

Each file's properties are described next.

## 447.2 Configuring app-server.properties

---

Specify the following information to configure 7.0's app server:

- `dir`: the absolute path of the application server folder. *(required)*
- `extra.lib.dirs`: a comma delimited list of extra directories containing any binaries or resources to add to the class path. Use all absolute paths OR all paths relative to `dir`. *(required)*
- `global.lib.dir`: the application server's global library directory. Use the absolute path or a path relative to `dir`. *(required)*
- `portal.dir`: the directory where portal is installed in your app server. Use the absolute path or a path relative to `dir`. *(required)*
- `server.detector.server.id`: ID of a supported application server. *(required)* Here are the IDs:

- jboss
- jonas
- resin
- tomcat
- weblogic
- websphere
- wildfly

Relative paths must use Unix style format. The following properties, for example, are for Windows and use relative paths:

```
dir=D:\
extra.lib.dirs=Liferay/liferay-portal-master/tomcat-9.0.10/bin
global.lib.dir=Liferay/liferay-portal-master/tomcat-9.0.10/lib
portal.dir=Liferay/liferay-portal-master/tomcat-9.0.10/webapps/ROOT
server.detector.server.id=tomcat
```

These properties, for example, are for Linux and use all absolute paths:

```
dir=/
extra.lib.dirs=/home/user/liferay/liferay-portal-master/tomcat-9.0.10/bin
global.lib.dir=/home/user/liferay/liferay-portal-master/tomcat-9.0.10/lib
portal.dir=/home/user/liferay/liferay-portal-master/tomcat-9.0.10/webapps/ROOT
server.detector.server.id=tomcat
```

## 447.3 Configuring portal-upgrade-database.properties

---

Specify the following information to configure the database you're upgrading. Note that these properties correspond exactly to the JDBC portal properties you'd use in a `portal-ext.properties` file.

- `jdbc.default.driverClassName` *(required)*
- `jdbc.default.url` *(required)*
- `jdbc.default.username` *(required)*
- `jdbc.default.password` *(required)*

## 447.4 Configuring portal-upgrade-ext.properties

---

Specify the following information to configure the upgrade:

`liferay.home`: The Liferay home folder (*required*)

`dl.store.impl`: The implementation for persisting documents to the document library store. This property is mandatory if you're using a `*FileSystemStore` implementation. If you updated this property in your `portal-ext.properties`, copy it here. Otherwise, set the property one of these ways:

```
dl.store.impl=com.liferay.portal.store.file.system.FileSystemStore
dl.store.impl=com.liferay.portal.store.db.DBStore
dl.store.impl=com.liferay.portal.store.file.system.AdvancedFileSystemStore
dl.store.impl=com.liferay.portal.store.s3.S3Store
```

`hibernate.jdbc.batch_size`: The JDBC batch size used to improve performance; set to 250 by default (*optional*)

## 447.5 Example Upgrade Configuration

---

Here's an example interaction with the upgrade tool's command line interface:

```
Please enter your application server (tomcat):
tomcat
Please enter your application server directory (../../tomcat-8.0.32):

Please enter your extra library directories (../../tomcat-8.0.32/bin):

Please enter your global library directory (../../tomcat-8.0.32/lib):

Please enter your portal directory (../../tomcat-8.0.32/webapps/ROOT):

[db2 mariadb mysql oracle postgresql sqlserver sybase]
Please enter your database (mysql):
mariadb
Please enter your database host (localhost):

(etc.)
```

The command line interface creates the configuration files based on your input. You can put this information into configuration files to configure the tool manually.

Here are example upgrade configuration files that you can customize and copy into `[Liferay Home]/tools/portal-tools-db-upgrade-client/`:

- `app-server.properties`:

```
dir=../../tomcat-8.0.32
global.lib.dir=/lib
portal.dir=/webapps/ROOT
server.detector.server.id=tomcat
extra.lib.dirs=/bin
```

- `portal-upgrade-database.properties`:

```
jdbc.default.url=jdbc:mysql://portal62?characterEncoding=UTF-8&dontTrackOpenResources=true&holdResultsOpenOverStatementClose=true&serverTime
jdbc.default.driverClassName=com.mysql.cj.jdbc.Driver
jdbc.default.username=root
jdbc.default.password=
```

- portal-upgrade-ext.properties:

```
liferay.home=/home/user/servers/liferay7
module.framework.base.dir=/home/user/servers/liferay7/osgi
dl.store.impl=com.liferay.portal.store.file.system.FileSystemStore
```

Next, decide if the upgrade tool should upgrade non-core modules automatically.

## 447.6 Configuring Non-Core Module Data Upgrades

---

You can configure the upgrade tool to upgrade all installed modules automatically or to open a Gogo shell (after core upgrade completes) for you to execute module upgrades manually.

If the upgrade tool's `autoUpgrade` property is set to `true` (the default setting), upgrade processes for all installed modules are run too.

If you set `autoUpgrade="false"` in a file called `com.liferay.portal.upgrade.internal.configuration.ReleaseManager` and copy the file into the `[Liferay Home]/osgi/configs` folder, the upgrade tool opens Gogo shell after the core upgrade. In the Gogo shell, you can administer module upgrades.

It's time to run the upgrade tool.

---

# UPGRADING THE CORE USING THE UPGRADE TOOL

---

This document has been updated and ported to Liferay Learn and is no longer maintained here. The upgrade tool provides the easiest way to upgrade the core and installed modules. Here's how to use it.

## 448.1 Upgrade Tool Usage

---

The `db_upgrade.sh` script in the `[Liferay Home]/tools/portal-tools-db-upgrade-client` folder (`db_upgrade.bat` on Windows) invokes the upgrade tool.

This command prints the upgrade tool usage:

```
db_upgrade.sh --help
```

This configuration prevents automatic module upgrade, but causes the upgrade tool to open a Gogo shell for upgrading modules after finishing the core upgrade.

Here are the tool's default Java parameters:

```
-Dfile.encoding=UTF-8 -Duser.country=US -Duser.language=en -Duser.timezone=GMT -Xmx2048m
```

The `-j` option overrides the JVM parameters. For example, these options set the JVM memory to 10GB, which is a good starting point for this process type:

```
db_upgrade.sh -j "-Dfile.encoding=UTF-8 -Duser.country=US -Duser.language=en -Duser.timezone=GMT -Xmx10240m"
```

The `-l` option specifies the tool's log file name:

```
db_upgrade.sh -l "output.log"
```

Here are all the upgrade tool command line options:

**-help** or **-h**: Prints the tool's help message.

**-jvm-opts** or **-j + [arg]**: Sets any JVM options for the upgrade process.

**-log-file** or **-l + [arg]**: Specifies the tool's log file name—the default name is `upgrade.log`.

**-shell** or **-s**: Automatically connects you to the Gogo shell after finishing the upgrade process.

---

**Note:** Only execute the upgrade process on a server with ideal memory, CPU, and database connection configurations. If executing an upgrade remotely using ssh, make sure to guard against interruptions:

- If you're executing the upgrade using ssh, ignore hangups (connection loss) by using `nohup` or something similar.
- On the machine you're connecting from, disable settings that shutdown or sleep that machine.

The upgrade process continues on the server even if you lose connection to it. If you lose connection, reconnect and monitor upgrade status via the log (default log file is `upgrade.log`). If you're using an earlier version of 7.0 and upgrade execution is interrupted, check your log file for where execution stopped.

- If execution stopped during an upgrade process for Core 7.1 or higher, or any module upgrade process, restart the upgrade tool to continue the upgrade from that point. You can also use Gogo shell to check module upgrade status and continue upgrading modules.
- If execution stopped during an upgrade process for Core 7.0 or lower, you must restore the data from a backup and start the upgrade again.

---

**Warning:** To prevent the tool's expanded command from growing too large for Windows, execute the upgrade tool script from the `[Liferay Home]/tools/portal-tools-db-upgrade-client` folder.

---

It's time to upgrade your core data using the upgrade tool.

---

## 448.2 Running and Managing the Core Upgrade

---

Start the upgrade tool, as the previous section explains. Here are the core upgrade stages:

1. Show the upgrade patch level
2. Execute the core upgrade processes
3. Execute the core verifiers

Monitor the upgrade via the upgrade tool log file (default file is `upgrade.log`). If a core upgrade process fails, analyze the failure and resolve it. If a core upgrade step for Liferay DXP 7.1 (or newer) fails, executing the upgrade tool again starts it from that step.

If you configured the upgrade tool to upgrade non-core modules, the tool opens a Gogo shell and starts upgrading them. The Gogo shell lets you upgrade modules, check module upgrade status, verify upgrades, and restart module upgrades. Read on to learn how to use Gogo shell commands to complete Liferay DXP upgrades.



---

# UPGRADING MODULES USING GOGO SHELL

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Liferay's Gogo shell can upgrade and verify individual modules. It's a fine-grained approach to upgrading the core and non-core modules. If you haven't already upgraded your non-core modules using the upgrade tool or if there are modules you need to revisit upgrading, you can upgrade them using Gogo Shell.

---

**Note:** You must Configure the core upgrade before using Gogo shell commands to upgrade the core.

---

Below is a list of commands.

---

## 449.1 Command Usage

---

If you ran the upgrade tool and it opened Gogo shell, you're already connected. Otherwise, you can execute commands using the Gogo Shell portlet.

Here are the commands:

`exit` or `quit`: Exits the Gogo shell

`upgrade:help`: Displays upgrade commands

`upgrade:check`: Lists upgrades pending execution because they failed in the past or the module hasn't reached its final version

`upgrade:execute {module_name}`: Executes upgrades for that module

`upgrade:executeAll`: Executes all pending module upgrade processes

`upgrade:list`: Lists all registered upgrades

`upgrade:list {module_name}`: Lists the module's required upgrade steps

`upgrade:list | grep Registered`: Lists registered upgrades and their versions

`verify:help`: Displays verify commands

`verify:check {module_name}`: Lists the latest execution result for the module's verify process

`verify:checkAll`: Lists the latest execution results for all verify processes

verify:execute {module\_name}: Executes the module's verifier  
 verify:executeAll: Executes all verifiers  
 verify:list: Lists all registered verifiers

There are many useful Liferay commands and standard commands available in Gogo shell. The following sections describe Liferay upgrade commands.

## 449.2 Listing module upgrade processes

Before upgrading modules, you should find which have unresolved dependencies, which are resolved and available to upgrade, and examine the module upgrade processes.

Executing `upgrade:list` in the Gogo shell lists the modules whose upgrade dependencies are satisfied. These modules can be upgraded.

If a module is active but not listed, its dependencies must be upgraded. The Gogo shell command `scr:info [upgrade_step_class_qualified_name]` shows the upgrade step class's unsatisfied dependencies. Here's an example `scr:info` command:

```
scr:info com.liferay.portal.upgrade.JournalServiceUpgrade
```

Invoking `upgrade:list [module_name]` lists the module's upgrade processes, in no particular order. For example, executing `upgrade:list com.liferay.bookmarks.service` (for the Bookmarks Service module), lists this:

```
Registered upgrade processes for com.liferay.bookmarks.service 1.0.0
 {fromSchemaVersionString=0.0.0, toSchemaVersionString=1.0.0, upgradeStep=com.liferay.portal.spring.extender.internal.context.ModuleApplication
 {fromSchemaVersionString=0.0.1, toSchemaVersionString=1.0.0-step-3, upgradeStep=com.liferay.bookmarks.upgrade.v1_0_0.UpgradePortletId@5f41b7
 {fromSchemaVersionString=1.0.0-step-1, toSchemaVersionString=1.0.0, upgradeStep=com.liferay.bookmarks.upgrade.v1_0_0.UpgradePortletSettings@
 {fromSchemaVersionString=1.0.0-step-2, toSchemaVersionString=1.0.0-step-1, upgradeStep=com.liferay.bookmarks.upgrade.v1_0_0.UpgradeLastPublis
 {fromSchemaVersionString=1.0.0-step-3, toSchemaVersionString=1.0.0-step-2, upgradeStep=com.liferay.bookmarks.upgrade.v1_0_0.UpgradeClassNames
```

An application's upgrade step class names typically reveal their intention. For example, the example's `com.liferay.bookmarks.upgrade.v1_0_0.UpgradePortletId` upgrade step class updates the app's portlet ID. The other example upgrade step classes update class names, the `LastPublishDate`, and `PortletSettings`. The example's step from `0.0.0` to `1.0.0` upgrades the module from an empty database.

To examine a module's upgrade process better, you can sort the listed upgrade steps mentally or in a text editor. Here's the upgrade step order for a Bookmarks Service module to be upgraded from Liferay Portal 6.2 (the module's database exists) to schema version `1.0.0`:

- `0.0.1` to `1.0.0-step-3`
- `0.0.1-step-3` to `1.0.0-step-2`
- `0.0.1-step-2` to `1.0.0-step-1`
- `0.0.1-step-1` to `1.0.0`

The overall module upgrade process starts at version `0.0.1` and finishes at version `1.0.0`. The first step starts on the initial version (`0.0.1`) and finishes on the target version's highest step (`step-3`). The last step starts on the target version's lowest step (`step-1`) and finishes on the target version (`1.0.0`).

Once you understand the module's upgrade process, you can execute it with confidence.

### 449.3 Executing module upgrades

---

Executing `upgrade:execute [module_name]` upgrades the module. You might run into upgrade errors that you must resolve. Executing the command again starts the upgrade from the last successful step.

You can check upgrade status by executing `upgrade:list [module_name]`. For example, entering `upgrade:list com.liferay.iframe.web` outputs this:

```
Registered upgrade processes for com.liferay.iframe.web 0.0.1
 {fromSchemaVersionString=0.0.1, toSchemaVersionString=1.0.0, upgradeStep=com.liferay.iframe.web.upgrade.IFrameWebUpgrade$1@1537752d}
```

The first line lists the module's name and current version. The example module's current version is 0.0.1. The `toSchemaVersionString` value is the target version.

Executing `upgrade:list [module_name]` on the module after successfully upgrading it shows the module's name followed by the version you targeted.

For example, if you successfully upgraded `com.liferay.iframe.web` to version 1.0.0, executing `upgrade:list com.liferay.iframe.web` shows the module's version is 1.0.0:

```
Registered upgrade processes for com.liferay.iframe.web 1.0.0
 {fromSchemaVersionString=0.0.1, toSchemaVersionString=1.0.0, upgradeStep=com.liferay.iframe.web.upgrade.IFrameWebUpgrade$1@1537752d}
```

For module upgrades that don't complete, you can check their status and resolve their issues.

### 449.4 Checking upgrade status

---

The command `upgrade:check` lists modules that have impending upgrades.

For example, if module `com.liferay.dynamic.data.mapping.service` failed in a step labeled 1.0.0-step-2, executing `upgrade:check` shows this:

```
Would upgrade com.liferay.dynamic.data.mapping.service from 1.0.0-step-2 to
1.0.0 and its dependent modules
```

Modules often depend on other modules to complete upgrading. Executing `scr:info [upgrade_step_class_qualified_name]` shows the upgrade step class's dependencies. You must upgrade modules on which your module depends.

To resolve and activate a module, its upgrade must complete. The Apache Felix Dependency Manager Gogo shell command `dm wtf` reveals unresolved dependencies. If your module requires a certain data schema version (e.g., its `bnd.bnd` specifies `Liferay-Require-SchemaVersion: 1.0.2`) but the module hasn't completed upgrade to that version, `dm wtf` shows that the schema version is not registered.

```
1 missing dependencies found.
```

```

The following service(s) are missing:
```

```
* com.liferay.portal.kernel.model.Release (&(release.bundle.symbolic.name=com.liferay.journal.service)(release.schema.version=1.0.2)) is not found
```

The `dm wtf` command can also help detect errors in portlet definitions and custom portlet `schemaVersion` fields.

Browsing the Liferay DXP database `Release_` table can help you determine a module's upgrade status too. The core's `servletContextName` field value is `portal`. If the core's `schemaVersion` field matches your new Liferay DXP version (e.g., 7.2.1 for Liferay Portal CE GA2) and the `verified` field is 1 (true), the core upgrade completed successfully.

Each module has one `Release_` table record, and the value for its `schemaVersion` field must be 1.0.0 or greater (1.0.0 is the initial version for 7.0 modules, except for those that were previously traditional plugins intended for Liferay Portal version 6.2 or earlier).

### **449.5 Executing verify processes**

---

Some modules have verify processes. These make sure the upgrade executed successfully. Verify processes in the core are automatically executed after upgrading Liferay DXP. You can also execute them by configuring the `verify.* portal` properties and restarting your server.

To check for available verify processes, enter the Gogo shell command `verify:list`. To run a verify process, enter `verify:execute [verify_qualified_name]`.

---

## EXECUTING POST-UPGRADE TASKS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here. Since you optimized your system for upgrading, after the upgrade is complete you must re-optimize it for production.

---

### 450.1 Tuning Your Database for Production

---

Prior to upgrading your Liferay DXP database, you tuned it for upgrade. Now that upgrade is complete, restore the production database tuning you used previously.

---

### 450.2 Re-enabling Search Indexing and Re-indexing Search Indexes

---

Make sure to re-enable search indexing by removing the `com.liferay.portal.search.configuration.IndexStatusManager` file from your `[Liferay Home]/osgi/configs` folder or setting this property in it:

```
indexReadOnly="false"
```

Then re-index Liferay DXP's search indexes. Don't just do this blindly, however. By default, Liferay DXP ships with an embedded configuration for Elasticsearch. This configuration works great for demo purposes, but is not supported in production. Make sure to install and configure a standalone Elasticsearch instance to run in production.

---

### 450.3 Enabling Web Content View Permissions

---

Prior to Liferay DXP 7.1, all users could view web content articles by default. Now view permissions are checked by default. Here are options for opening view permissions:

Option 1: Edit view permissions per web content article per role.

Option 2: Open view permissions for all web content articles by going to *System Settings*  *Web Experience*  *Web Content* and deselecting *Article view permissions check enabled*.

Once you've configured search, re-indexed your search index, and set web content view permissions, your upgraded system is ready for action! Congratulations!

---

## UPGRADING A SHARDED ENVIRONMENT

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Since Liferay DXP 7.0, Liferay removed its own physical partitioning implementation (also known as sharding) in favor of the capabilities provided natively by database vendors. Upgrading a sharded installation to DXP 7.0 or higher requires migrating it to as many non-sharded Liferay DXP installations (servers) as you have shards. These steps guide you through configuring the new Liferay DXP servers to use your formerly sharded data.

---

**Note:** Liferay continues to support its logical partitioning capabilities (also known as virtual instances) for the foreseeable future.

---

For any further assistance with sharding contact your Liferay account manager or Liferay Support.

---

### 451.1 Add Configurations Before the Data Upgrade

---

In addition to other configurations, you will need to set more properties to migrate your shards to virtual instances for your data upgrade.

Here is how to configure the upgrade to migrate from sharding:

1. Copy all of the shard JDBC connection properties from `portal-ext.properties` to `portal-upgrade-database.properties`. For example, JDBC connections for a default shard and two non-default shards might look like this:

```
jdbc.default.driverClassName=[the database driver class name]
jdbc.default.url=[the URL to the default database shard]
jdbc.default.username=[the user name]
jdbc.default.password=[the password]
```

```

jdbc.one.driverClassName=[the database driver class name]
jdbc.one.url=[the URL to database shard one]
jdbc.one.username=[the user name]
jdbc.one.password=[the password]

jdbc.two.driverClassName=[the database driver class name]
jdbc.two.url=[the URL to database shard two]
jdbc.two.username=[the user name]
jdbc.two.password=[the password]

```

2. Set the JDBC *default* connection properties in each server's `portal-upgrade-database.properties` to specify the associated shard.

- Add the original JDBC properties for the respective non-default shard database. For example, shard one's original properties might start with `jdbc.one`:

```

jdbc.one.driverClassName=[the database driver class name]
jdbc.one.url=[the URL to database shard one]
jdbc.one.username=[the user name]
jdbc.one.password=[the password]

```

- Rename the properties to start with `jdbc.default`. For example:

```

jdbc.default.driverClassName=[the database driver class name]
jdbc.default.url=[the URL to database shard one]
jdbc.default.username=[the user name]
jdbc.default.password=[the password]

```

## 451.2 Upgrade and Update Properties

---

When you perform the database upgrade, upgrade the default shard first, and then each of the non-default shards. See *Using the Database Upgrade Tool* for more information on running the database upgrade.

After the database upgrade has been completed, make the following configuration changes to your application servers:

1. In each server's `portal-ext.properties`, use the JDBC *default* properties you specified in the `portal-upgrade-database.properties` (see the *default* properties above).
2. Remove the non-default shard JDBC properties from the default shard server's `portal-ext.properties` file, leaving only the default shard database `jdbc.default` properties. For example:

Old JDBC properties:

```

jdbc.default.driverClassName=[the database driver class name]
jdbc.default.url=[the URL to the default database shard]
jdbc.default.username=[the user name]
jdbc.default.password=[the password]

jdbc.one.driverClassName=[the database driver class name]
jdbc.one.url=[the URL to database shard one]

```



```
jdbc.one.username=[the user name]
jdbc.one.password=[the password]

jdbc.two.driverClassName=[the database driver class name]
jdbc.two.url=[the URL to database shard two]
jdbc.two.username=[the user name]
jdbc.two.password=v[the password]
```

#### New JDBC properties:

```
jdbc.default.driverClassName=[the database driver class name]
jdbc.default.url=[the URL to your database]
jdbc.default.username=[the user name]
jdbc.default.password=[the password]
```

Once you have completed all of these steps, you have migrated off of a sharded environment to virtual instances on separate Liferay DXP servers together with your DXP upgrade.

Congratulations! You have migrated off of a sharded environment to virtual instances on separate Liferay DXP servers. You have also upgraded to 7.0. Your virtual instances are ready for action.



---

## MIGRATING FROM AUDIENCE TARGETING TO SEGMENTATION AND PERSONALIZATION

---

7.0 integrates all the Audience Targeting app's features into Liferay's core as Segmentation and Personalization. This enables better integration with other applications and provides developers with easier access to Segmentation and Personalization features. Audience Targeting users must migrate their user segments into the new Segments application. There are three steps to the migration process:

1. Upgrade to 7.0.
2. Migrate custom rules.
3. Migrate behavior-based features.

First, to upgrade to the latest version of Liferay DXP, follow the upgrade guide. Most of your Audience Targeting configuration is automatically transferred into the new engine.

Next, any custom rules that were created in Audience Targeting must be re-evaluated. Some custom rules may have an out-of-the-box equivalent now, while others must be migrated. If a rule must be re-implemented, follow the Segmentation and Personalization development guide. You can check the list of rules that are automatically migrated to see how much additional work you have in store. You must also migrate display widgets since the new Personalization features use different tools.

Finally, you must migrate behavior-based features, but since Audience Targeting's analytics features are now part of Analytics Cloud, there isn't a direct path to upgrade. See the Analytics Cloud documentation.



## MIGRATING USER SEGMENTS

In Audience Targeting, a user segment represents a subset of users. A user segment is defined by one or more rules that users must match to belong to that user segment. In 7.0, segments work in a similar way, but they are defined by criteria instead of rules. Segment criteria are sets of fields defined by different user actions or properties (profile information, organization information, session information) that can be combined through operations (like equals, not equals, contains, not contains, greater than, and less than) and conjunctions (AND, OR) to define complex filters.

Due to the similarities between Audience Targeting user segments and 7.0 Segments, certain data can be migrated automatically as part of the upgrade process.

### 453.1 Upgrade Process

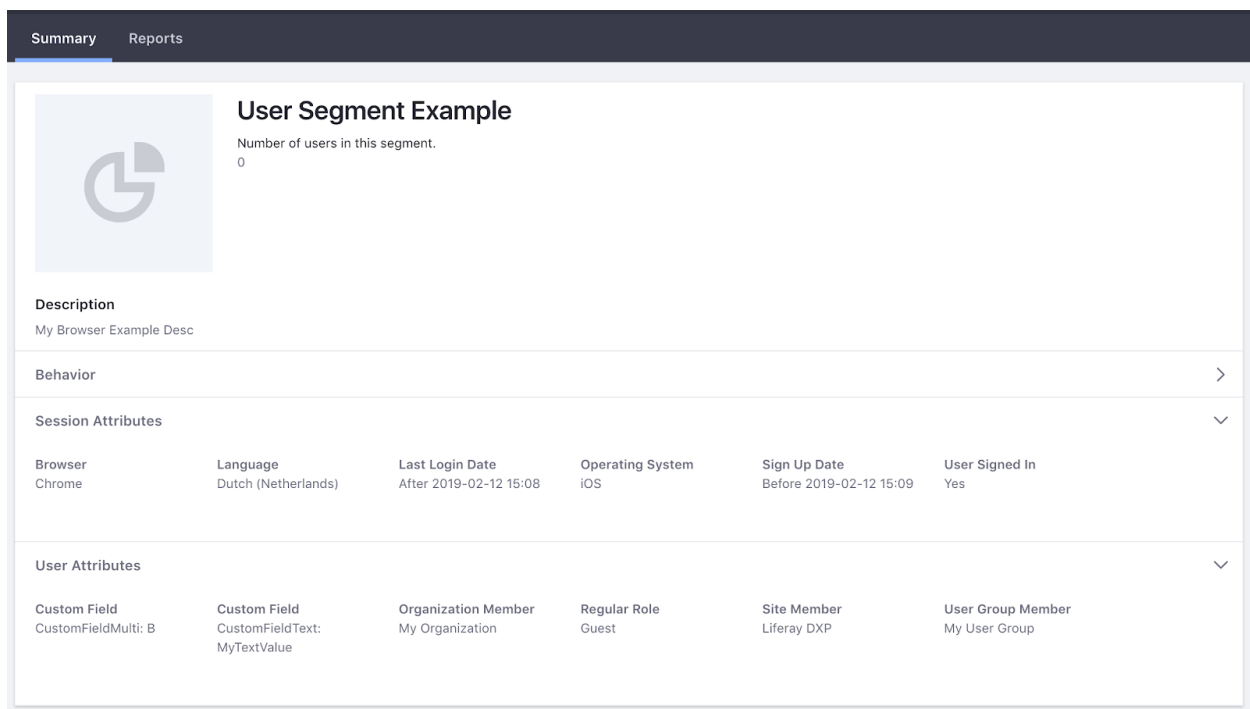
As a result of the upgrade process,

- All Audience Targeting User Segments appear under the new Segments administration in 7.2, with the same name.
- For every segment, those Audience Targeting rules with an equivalent in Liferay DXP 7.2 have been migrated into the corresponding criteria fields (see Table below).
- Audience Targeting tables have been removed from your Liferay DXP Database.

| Audience Targeting Rule | Liferay DXP 7.2. Segment Criteria Field | Upgrade Path                                                              |
|-------------------------|-----------------------------------------|---------------------------------------------------------------------------|
| Browser                 | Browser                                 | Automated. Use user agent field with contains operation as an alternative |
| Custom Field            | Custom Field                            | Automated                                                                 |
| Language                | Language                                | Automated                                                                 |

| Audience Targeting Rule | Liferay DXP 7.2. Segment Criteria Field | Upgrade Path            |
|-------------------------|-----------------------------------------|-------------------------|
| Last Login Date         | Last Sign In Date                       | Automated               |
| Organization Member     | Organization                            | Automated               |
| OS                      | User Agent                              | Automated               |
| Previous Visited Site   | Not Available                           | Automated               |
| Regular Role            | Role                                    | Automated               |
| Site Member             | Site                                    | Automated               |
| User Group Member       | User Group                              | Automated               |
| Age                     | Not Available                           | Suggested: custom field |
| Facebook (various)      | Not Available                           | Suggested: custom field |
| Gender                  | Not Available                           | Suggested: custom field |
| Score Points            | Not Available                           | Suggested: cookie       |
| Visited Page/Content    | Not Available                           | Suggested: cookie       |

Here's an example user segment as it would appear in Audience Targeting for Liferay DXP 7.1:



**User Segment Example**  
Number of users in this segment.  
0

**Description**  
My Browser Example Desc

**Behavior** >

**Session Attributes** ∨

| Browser | Language            | Last Login Date        | Operating System | Sign Up Date            | User Signed In |
|---------|---------------------|------------------------|------------------|-------------------------|----------------|
| Chrome  | Dutch (Netherlands) | After 2019-02-12 15:08 | iOS              | Before 2019-02-12 15:09 | Yes            |

**User Attributes** ∨

| Custom Field        | Custom Field                 | Organization Member | Regular Role | Site Member | User Group Member |
|---------------------|------------------------------|---------------------|--------------|-------------|-------------------|
| CustomFieldMulti: B | CustomFieldText: MyTextValue | My Organization     | Guest        | Liferay DXP | My User Group     |

Figure 453.1: A Liferay DXP 7.1 Audience Targeting Segment.

And here is the same segment migrated to Liferay 7.2:

For those Audience Targeting rules without a direct equivalent, a manual migration is required. If you have any these rules, you can learn about your next steps in Manual Migration.

The image shows a configuration interface for a Liferay DXP 7.2 Segment, divided into three sections: User Properties, Organization Properties, and Session Properties. Each section has a title, a list of conditions, and an edit icon in the top right corner.

- User Properties:** Contains five conditions connected by "- AND -" operators:
  - User with property **Group** equals **Liferay DXP**
  - User with property **Role** equals **Guest**
  - User with property **Customfieldmulti** equals **B**
  - User with property **Customfieldtext** equals **MyTextValue**
  - User with property **User Group** equals **My User Group**
- Organization Properties:** Contains one condition:
  - Organization with property **Organization** equals **MyOrganization**
- Session Properties:** Contains three conditions connected by "- AND -" operators:
  - Session with property **Browser** equals **Chrome**
  - Session with property **Language** equals **nl\_NL**
  - Session with property **Last Sign In Date** greater than **12 February 2019**

Each section is separated by a horizontal bar with the word "And" on the left side, indicating that the conditions within each section are combined with the conditions of the other sections using an AND logic.

Figure 453.2: A Liferay DXP 7.2 Segment





## MANUALLY MIGRATING FROM AUDIENCE TARGETING

---

As explained in the previous article, some Audience Targeting rules do not have a direct equivalent in Liferay DXP 7.2 and, therefore, they cannot be migrated automatically. Here are the recommended solutions for each rule type.

### **454.1 User Attribute Rules**

---

Some User Attributes, like Gender or Age, do not have a direct equivalent in 7.0. User Attributes retrieved from external sources like Facebook also do not have a replacement. To replace these, you must create a custom user field and use that to define your new Segment.

### **454.2 Session Rules**

---

For Session attributes that do not have a direct equivalent, the recommended solution is to use a URL field for the current URL or a previously visited URL on your site as criteria, or to use a Cookie for more advanced session tracking needs.

### **454.3 Behavior Rules**

---

In 7.0 analytics is now managed through Analytics Cloud. You can learn more about creating behavior based rules in the Analytics Cloud documentation.

#### **454.4 Migrating Custom Rules**

---

Audience Targeting segmentation features could be extended using custom rules. As part of the upgrade planning process, the function of any such rules should be re-evaluated with the new Segmentation features of 7.0 in mind.

First, check the Segmentation reference if any new criteria fields can replace their function. In particular, custom fields, URL fields, and cookies might help you migrate your custom rules with little to no additional development.

If none of them cover your requirements, follow the development guide for instructions on how to add new criteria fields and contributors.

#### **454.5 Migrating Display Portlets**

---

With Audience Targeting, you could display personalized content with the User Segment Display Content portlet or by using Asset Publisher with the Segments filter enabled. In 7.0, you must choose the most appropriate personalization option for your use cases.

#### **454.6 User Segment Content Display**

---

The User Segment Content Display portlet was used to display existing content based on segment membership rules. In 7.0, you can cover the same use case by defining manual content sets with variations for your different audiences and applying it to an asset publisher. See the documentation for creating personalized Content Sets. With this feature, you can assign any number of assets to the Content List for the given audience, and then use the Asset Publisher to define how content is displayed on the page.

#### **454.7 Asset Publisher Personalization**

---

Finally, if you want to display a dynamic list of content for your different audiences based on a filter in the same way you did with in Audience Targeting with the Segments filter in the Asset Publisher, you can create a dynamic content set with variations for your audiences and apply it to an asset publisher.

In addition, the new Experience-based Content Page personalization feature may fulfill a use case that you were previously solving with one of the methods previously available.

---

## DEPRECATED APPS IN 7.2: WHAT TO DO

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

During the development of any software product, it's sometimes necessary to stop development on or remove outdated or unpopular features. 7.0 is no different. In 7.0, Liferay has deprecated several apps and features.

There are three types of deprecated apps:

1. Deprecated apps that remain in Liferay DXP, but will be removed in a future release. (Availability: *Bundled*)
2. Deprecated apps that have been removed from Liferay DXP, yet are still available for download via Liferay Marketplace (Availability: *Marketplace*)
3. Deprecated apps that have been removed from Liferay DXP and aren't available for download. (Availability: *Removed*)

---

**Note:** All apps deprecated by Liferay are no longer in active development. You should therefore plan to stop using these apps. Such apps, however, may still be available for download.

---

---

**Note:** For information on apps deprecated in Liferay DXP 7.1, please see [Deprecated Apps in 7.1: What to Do](#)

---

Here are the apps deprecated in 7.0.

---

### 455.1 Foundation

---

| App                             | Availability | Notes                                                                                                                                                                                                                          |
|---------------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AlloyUI                         | Bundled      | Replaced by MetalJS (temporary) exposed as ClayUI tag equivalents.                                                                                                                                                             |
| CMIS Store                      | Removed      | Migrate to another Document Repository Store option. Before upgrading to 7.0, migrate your document store data using Data Migration in Server Administration.                                                                  |
| JCRStore                        | Removed      | Migrate to another Document Repository Store option. Before upgrading to 7.0, migrate your document store data using Data Migration in Server Administration.                                                                  |
| Legacy Search Portlet           | Bundled      | Will be removed in a future release. Replaced by the Search widgets.                                                                                                                                                           |
| Liferay Mobile Device Detection | Removed      | Contact 51Degrees for up-to-date definitions.                                                                                                                                                                                  |
| Enterprise Sprite framework     | Bundled      | Liferay's image sprite framework is deprecated and is disabled by default via the <code>sprite.enabled</code> portal property. You can still build image sprites using any framework you like and deploy them in your plugins. |

---

## 455.2 Personalization

---

| App                | Availability | Notes                        |
|--------------------|--------------|------------------------------|
| Audience Targeting | Removed      | Replaced by Personalization. |

### 455.3 Web Experience

| App                          | Availability | Notes                                                                      |
|------------------------------|--------------|----------------------------------------------------------------------------|
| RSS Publisher                | Bundled      | See the article on enabling and using this widget.                         |
| User Group Pages (Copy Mode) | Bundled      | See the Legacy User Group Sites Behavior instructions on how to enable it. |
| Resources Importer           | Bundled      | Deprecated as of 7.1 with no direct replacement                            |

### 455.4 Forms

| App      | Availability | Notes                           |
|----------|--------------|---------------------------------|
| Web Form | Removed      | Final version released for 7.0. |

### 455.5 Security

| App                                  | Availability | Notes                                 |
|--------------------------------------|--------------|---------------------------------------|
| Central Authentication Service (CAS) | Bundled      | Migrate to SAML based authentication. |
| Google Login                         | Marketplace  | Replaced by OpenID Connect.           |
| NTLM                                 | Marketplace  | Replaced by Kerberos.                 |
| OpenAM / OpenSSO                     | Bundled      | Migrate to SAML based authentication. |
| OpenID                               | Marketplace  | Replaced by OpenID Connect.           |

---

#### 455.6 User and System Management

---

| App        | Availability                                                         |
|------------|----------------------------------------------------------------------|
| Live Users | Enabled through the <code>live.users.enabled</code> portal property. |

---

#### 455.7 Related Topics

---

Apps in Maintenance Mode

---

## APPS IN MAINTENANCE MODE

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

At a designated time, Liferay may cease enhancing a product or capability. This is called *maintenance mode*. During this mode, Liferay actively supports and provides bug fixes for the product or capability in accordance with the subscribers' subscription level and the end of service life policies of the compatible Liferay DXP version. Maintenance mode does not necessarily mean that deprecation in a future Liferay DXP version is planned for the product or capability; it only means that enhancements aren't being made for the current Liferay DXP development cycle.

As of Liferay DXP 7.2, these products and capabilities have transitioned into maintenance mode:

- Liferay Connected Services
- Liferay Connector to OAuth 1.0a
- Liferay Drools
- Liferay Mobile Experience (Liferay Screens, Liferay Mobile SDK, Liferay Push)
- Liferay Reports
- Liferay Sync
- Staging

### **456.1 Related Topics**

---

Deprecated Apps in 7.2: What to do?





## MAINTAINING LIFERAY DXP

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Once you have a Liferay DXP installation, there are some things you must do to keep it running smoothly. Backing up your installation in case of a hardware failure protects your data and helps you get your system back in working order quickly. And if you're a DXP customer, patching your system regularly brings the latest bug fixes to your running instance.

---

Upgrading 7.0 to a new GA version can require data upgrade. Until you perform all required data upgrades (if any), Liferay DXP startup fails with messages like these:

```
2019-03-06 17:22:35.025 INFO [main][StartupHelper:72] There are no patches installed
You must first upgrade to Liferay DXP 7210
2019-03-06 17:22:35.098 ERROR [main][MainServlet:277] java.lang.RuntimeException: You must first upgrade to Liferay DXP 7201
```

---

Read on to learn about how to keep your system running well.



---

## PATCHING LIFERAY DXP

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

While we strive for perfection with every Liferay DXP release, the reality of the human condition dictates that releases may not be as perfect as originally intended. But we've planned for that. Included with every Liferay DXP bundle is a Patching Tool that handles installing two types of patches: fix packs and hotfixes.

---

**Important:** Make sure to back up your Liferay DXP installation and database regularly, especially before patching. The patching tool installs code changes and some of these make data changes (if necessary) automatically on startup.

Certain fix packs (service packs) can include data/schema micro changes—they're optional and revertible. Module upgrades and any micro changes they include are applied at server startup by default, or can be applied manually by disabling the `autoUpgrade` property. Server startup skips all Core micro changes. Instead, you can apply them using the upgrade tool before server startup.

---

**Important:** Installing the latest service pack on top of Liferay DXP 7.2 GA1/FP1 requires running the data upgrade tool. Examine the Liferay DXP upgrade instructions to determine preparations, testing, and post upgrade steps that are appropriate for you. To eliminate system down time during upgrade, consider using the Blue-green deployment technique.

Liferay DXP 7.2 FP2/SP1's data schema change adds version columns to several tables. Hibernate's optimistic locking system uses the columns to preserve data integrity during concurrent data modifications.

---

**Note:** Patching a cluster requires additional considerations.



---

## PATCHING BASICS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here. Liferay ships 7.0 fixes through three different channels:

- Fix packs
- Hotfixes
- Service Packs

---

### **459.1 Fix Packs**

---

The latest fixes that patch the core are bundled together weekly into fix packs that are provided to all of Liferay's customers. Fix packs include fixes for both the core and the applications and modules that ship with Liferay DXP. The fixes address regressions or obvious bugs and don't require you to make additional changes. Each fix pack contains all previous fix packs since the last service pack.

Security Fix Packs are a special fix pack type for deploying critical security fixes quickly without changing fix pack levels.

Fixes that don't fit these requirements are considered for service packs or hot fixes.

---

### **459.2 Hotfixes**

---

A hotfix is provided to customers when they contact Liferay about an emergency situation, and Liferay's support team—working with the customer—determines that the problem is a product issue that must be fixed very quickly. Support fixes the bug and provides a hotfix to the customer immediately. This is a short-term fix. Hotfixes can patch the core, the applications, and modules.

### 459.3 Service Packs

---

Service packs are built on top of the original Liferay DXP release and repackaged with the latest fix pack, Patching Tool, and modules. Since a service pack is a fix pack, it contains all previous fix packs since the last service pack. Each one includes the most recent patches and updates.

Service packs can also include changes that have these characteristics:

- Require more extensive testing.
- Require some of your attention, such as updating your documentation.
- Improve the product.

Rather than updating existing Liferay DXP systems with service packs, you should

1. Keep systems up-to-date with fix packs (according to your own deployment schedule).
2. Install the latest Marketplace updates frequently.
3. Update the Patching Tool when necessary.

This method updates the installation to the service pack levels, while allowing scheduled deployments and avoiding full environment rebuilds.

### 459.4 How Patches are Tested

---

Liferay extensively tests service packs, fix packs, and hotfixes to ensure high quality. Fixes in fix packs go through both automated regression testing and manual testing. Hotfixes receive similar automated testing, and the support engineer who fixes a reported issue tests it.

Before releasing a service pack, Liferay runs test suites on the packaged service pack.

---

## USING THE PATCHING TOOL

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The Patching Tool installs, removes, compares, and prepares Liferay DXP patches. It is pre-installed in Liferay DXP bundles, easy to install into @product@ manual installations, and easy to update. The Patching Tool's executable scripts facilitate patching.

Here are the essentials to get started using the Patching Tool:

- Installing the Patching Tool (for manual installations only)
- Executables

---

### **460.1 Installing the Patching Tool**

---

Liferay DXP bundles come with the Patching Tool pre-installed (in [Liferay Home]/patching-tool) and pre-configured with the default settings. Skip this section if you're using a bundle.

If you installed Liferay DXP manually, however, you must also install the Patching Tool manually.

1. Download the Patching Tool from the Customer Portal.
2. Unzip the Patching Tool to your Liferay Home folder (recommended) or to another folder.

After installing the Patching Tool, you must configure it to use your Liferay DXP installation. The patching-tool folder you extracted from the Patching Tool ZIP file contains the Patching Tool, including its executable scripts.

---

### **460.2 Executables**

---

The Unix shell and Windows batch scripts distributed with the Patching Tool make it easier to use. On Unix systems, run

```
./patching-tool.sh parameters
```

On Windows, run

```
patching-tool parameters
```

The Windows command `patching-tool` is used in the examples that follow. On Unix, replace the name of the executable before running the commands.

Installing patches is next!



---

## INSTALLING PATCHES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Before installing any patches, you must shut down your server. On Windows operating systems, files in use are locked by the OS, and can't be patched. On Unix-style systems, you can usually replace files that are running, but the old ones reside in memory. For these reasons, it is best to shut down Liferay DXP before installing patches.

Liferay distributes all patches (fix packs and hotfixes) as ZIP files. When you download a patch, either from a Help Center ticket (hotfix) or from the Customer Portal (fix pack), place it in the Patching Tool's patches folder (e.g., [Liferay Home]/patching-tool/patches) without unzipping it. To list your installed patches and available local patches, execute this command:

```
patching-tool info
```

This displays a list of patches you've already installed, along with a list of patches that *can* be installed from what's in the patches folder.

To install the available patches, use the following steps. First, issue the following command:

```
patching-tool install
```

To make sure the all changed OSGi bundles replace the existing ones, delete the `osgi/state` folder from the Liferay Home folder.

---

**Note:** The `osgi/state` folder contains OSGi bundle state information. If an OSGi bundle's changes in a hot fix or fix pack are internal only, they are invisible to the OSGi framework, that OSGi bundle stays installed, and its state information stays unchanged. Hot fixes, for example, may contain in-place changes that do not use the API. The framework cannot detect such changes. A fix pack's changes may also be transparent to the framework. For these reasons, deleting the `osgi/state` folder after applying fix packs and hot fixes is recommended.

---

**Important:** The `osgi/state` folder should **ONLY** be deleted when working in a development environment or when applying a fix pack or hot fix.

---

If there are new database indexes created by the patch, the Patching Tool tells you to update them. To get the list, run this command:

```
patching-tool index-info
```

Since there's no database connection at patching time, the indexes must be created at portal startup. If the server has permissions to modify the database indexes, instruct Liferay DXP to create the indexes automatically at startup by adding this setting to your `portal-ext.properties` file:

```
database.indexes.update.on.startup=true
```

Otherwise, you must create the indexes manually. Check the `patching-tool index-info` command output for more details.

After installing patches, you can execute the `patching-tool info` command to verify them.

---

**Note:** If there are any issues with the installed patches, verify that there aren't any remaining files from the previous patch installation of a fix pack or hotfix within the application server cache.

---

During the installation, `patching-backup-deps.zip` and `patching-backup.zip` files are created and stored in the web application's `WEB-INF` folder. These files are required to restore the Liferay DXP's original state; removing them disables patching.

---

**Note:** When installing patches, Liferay DXP's `web.xml` is always overwritten by the one contained in the patch. If you've customized `web.xml`, you must re-implement your customizations after installing a patch.

---

The `patching-backup.zip` file is necessary for installing future patches, because the Patching Tool reverts the installed fix pack before installing a new one. To revert the installed fix pack, it examines the contents of the `patching-backup.zip` to determine the changes that it needs to revert.

---

### 461.1 Handling Hotfixes and Patches

---

As stated previously, hotfixes are short term fixes provided as quickly as possible, and fix packs are larger bundles of hotfixes provided to all customers at regular intervals. If you already have a hotfix installed and the fix pack that contains that hotfix is released, the Patching Tool can manage this for you. Fix packs always supersede hotfixes; so when you install your fix pack, the hotfix it contains is uninstalled and the fix pack version is installed in its place.

The Patching Tool applies fixes to fix packs automatically. If a new (fixed) version of a fix pack is released, install it with the Patching Tool. The Patching Tool uninstalls the old fix pack and installs the new version in its place.

---

### 461.2 Fix Pack Dependencies

---

Some hotfixes depend on fix packs. If you attempt to install a hotfix that depends on a fix pack, the Patching Tool notifies you. Go to the Customer Portal and obtain the hotfix dependency. Once all the necessary patches are available in the patches folder, the Patching Tool installs them.

### **461.3 Updating the Patching Tool**

---

When a patch you're trying to install requires a Patching Tool update, the Patching Tool tells you. To update the Patching Tool, download the latest one from the Customer Portal. Overwrite the existing Patching Tool by unzipping the new one to the patching-tool folder's parent folder.

### **461.4 Cleaning Up**

---

After you've performed your patching procedure (whether you've installed or removed patches), it's important to clean up Liferay DXP's cache of deployed code. This ensures that you're using the revision you've just installed the patches for when you start the server. This is really easy to do.

To clear out the cached code, remove the contents of the [Liferay Home]/work folder. Now you're ready to start your server.



---

## WORKING WITH PATCHES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here. Here are some things you might need to do with patches:

- Report Patch Levels to Liferay Support
- Uninstall Patches
- Show collisions between patches and deployed plugins
- Separate Patches from your Installation

Start with reporting patch levels to Liferay Support.

---

### 462.1 Including support-info in Support Tickets

---

Providing your environment's information (e.g., hardware architecture) and patch level to Liferay Support is critical for reproducing your issues. Write your support information (including your patch level) to a file by executing this command:

```
patching-tool support-info
```

The support information is written to file `patching-tool-support-info-actual-timestamp.txt` in your `patching-tool` folder. Please upload this file to the Help Center ticket.

---

### 462.2 Uninstalling Patches

---

Have you noticed that the Patching Tool only seems to have an install command? This is because patches are managed not by the command, but by what appears in the patches folder. You manage the patches you have installed by adding or removing patches from this folder.

Here's how to uninstall (remove) a patch:

1. Remove the patch from your patches folder.
2. Run the `patching-tool install` command.

To revert ALL patches, run this command:

```
patching-tool revert
```

Now you know how to remove and revert patches. You can also compare patch levels. See the reference guide for a list of the available commands.

---

### 462.3 Showing collisions between patches and deployed plugins

---

Some patches update files you might have customized via a plugin. The `patching-tool list-collisions` command lists differences (collisions) between installed patch files and your plugin's version of them. Here's the command:

```
patching-tool list-collisions
```

It is an alias for the following `diff` command:

```
patching-tool diff collisions files _base
```

`_base` is the literal patch level name. Collisions are only listed for installed patches that contain source code files.

---

**Note:** As of Patching Tool 2.0.9, `patching-tool list-collisions` lists only JSP file collisions in fragment bundles.

---

---

### 462.4 Separating Patches from the Installation

---

The Patching Tool's `separate` command helps reduce the patched Liferay DXP installation size. If the installation has been patched, you can make it smaller by moving the restore files out of it.

Patched installations are large because the restore files are stored inside the web application's `WEB-INF` folder by default. These files are required for patching the installation again.

If these files are removed, subsequent patching processes fail. Because of this, Liferay added an option to separate the patching files from the installation while still preserving and restoring them safely when new patches arrive. To do this, use this command:

```
patching-tool separate [separation_name]
```

This command produces a `liferay-patching-files-[separation-name].zip` file in the Patching Tool's `patches` folder. It contains the necessary files and metadata for patching, verification, and validation. Once you create this file, the patch files are removed from their default location and are now only available in this file. You can store this file elsewhere to reduce your installation's size.

**WARNING:** If the product is separated from its patches in this way, you cannot run most of the Patching Tool commands until the patches are restored.

After the separation process only the following commands can be used:

- `auto-discovery`
- `info`
- `setup`

Any other command returns this:

This installation does not include data for patching. Please copy the `liferay-patching-files-[separation-name].zip` file into the `'patches'` directory and run `patching-tool setup`.

This is how you restore the patch files to your system. Details below.

## 462.5 Restoring the Separated Patch Files

---

When you need to patch Liferay DXP again, you must restore the separated patch artifact. To do this, copy the `liferay-patching-files-[separation-name].zip` back to the Patching Tool's `patches` folder and run `patching-tool setup` command.

The command finds the necessary patching artifact and restores the patch files to the installation. After that, the Patching Tool works like it did prior to separating the patches.





---

## CONFIGURING THE PATCHING TOOL

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The Patching Tool installs Liferay DXP patches. It ships with prepackaged Liferay DXP bundles. If any of the following scenarios describes your @product@ installation, however, you must configure the Patching Tool manually:

- Installed Liferay DXP manually on an existing application server
- Customized your Liferay DXP folder structure
- Running in a cluster

If none of the above scenarios describe your installation, you can skip this section.

If you installed Liferay DXP manually, you must also install the Patching Tool manually. Download it from the Customer Portal. Unzipping it to your Liferay Home folder is the easiest way to use it.

Read on to configure the Patching Tool for your environment.



---

## PATCHING TOOL BASIC CONFIGURATION

---

This document has been updated and ported to Liferay Learn and is no longer maintained here. There are two ways to configure the Patching Tool:

1. Automatically by executing the auto-discovery command
2. Manually by editing the configuration file (see Patching Tool Advanced Configuration)

Automatic configuration generates the configuration files by looking for Liferay DXP files in the local file system. By default the Patching Tool looks for them in its parent folder. To start the process, run this command in your Patching Tool folder (patching-tool):

```
patching-tool auto-discovery
```

If Liferay DXP is not installed in the parent folder, specify its location:

```
patching-tool auto-discovery /opt/liferay-dxp
```

If you specified the wrong location of Liferay DXP or it is not in the parent folder, the Patching Tool can't find the Liferay Home and reports an error like this:

```
The .liferay-home has not been detected in the given directory tree.
```

```
Configuration:
patching.mode=binary
war.path=../tomcat-9.0.17/webapps/ROOT/
global.lib.path=../tomcat-9.0.17/lib/ext/
liferay.home=**[please enter manually]**
```

```
The configuration hasn't been saved. Please save this to the default.properties file.
```

Here are ways to resolve the Liferay Home issue:

- Specify the Liferay Home path in the default.properties file.
- If the Liferay Home is in the Patching Tool's tree, create a .liferay-home file in the Liferay Home folder and re-run the auto-discovery process.

When the Patching Tool is configured, running `patching-tool info` reports product version information.

That's it! Now that you've installed and configured the Patching Tool, you're ready to download and install patches.



---

## PATCHING TOOL ADVANCED CONFIGURATION

---

This document has been updated and ported to Liferay Learn and is no longer maintained here. By default, the Patching Tool's configuration file called `default.properties` is in the tool's folder. A Patching Tool configuration file typically looks like this:

```
patching.mode=binary
war.path=../tomcat-9.0.17/webapps/ROOT/
global.lib.path=../tomcat-9.0.17/lib/ext/
liferay.home=../
```

The properties above (described fully in [Patching Tool Configuration Properties](#)) define the location of Liferay Home, the patching mode (binary or source), the path to where WAR files are deployed in the app server, and the global library path. The tool's auto-discovery bases the OSGi module framework paths on the Liferay Home. If, however, you changed the OSGi module framework paths to something different than those under the default folder `[Liferay Home]/osgi`, you must manually specify the following properties:

```
module.framework.core.path=path_to_modules_core_dir
module.framework.marketplace.path=path_to_modules_marketplace_dir
module.framework.modules.path=path_to_modules_modules_dir
module.framework.portal.path=path_to_modules_portal_dir
module.framework.static.path=path_to_modules_static_dir
```

Using auto-discovery and working with the default profile `default.properties` is the easiest way to use the Patching Tool, and is great for smaller, single server installations. But many Liferay DXP installations serve millions of pages per day, and the Patching Tool has been designed for this as well. So if you're running a small, medium, or large cluster of Liferay DXP machines, you can use the Patching Tool profiles to manage patching for all of them.

---

### 465.1 Using Profiles with the Patching Tool

---

You can create profiles for multiple runtimes by running auto-discovery or creating them manually. To auto-discover other runtimes, run the Patching Tool with parameters like this:

```
./patching-tool.sh [name of profile] auto-discovery [path/to/Liferay Home]
```

This runs the same discovery process, but on the path you specify. It writes the profile information to a file called [name of profile].properties. Alternatively, you can manually create profile property files in your patching-tool folder.

See Patching Tool configuration properties (profile properties) for a complete list of the available configuration properties.

You can have as many profiles as you want and use the same Patching Tool to patch all of them. This helps to keep all your installations in sync.

---

## INSTALLING PATCHES ON THE 7.0 WAR

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

If you installed Liferay DXP manually as a WAR file on a supported application server, you must apply patches to the WAR file and supporting files and re-deploy them. This article shows you how to do that.

### 466.1 Prerequisites

---

Download the necessary artifacts from the Customer Portal:

- Liferay DXP WAR file (liferay-dxp-[version].war)
- Dependencies ZIP file (liferay-dxp-dependencies-[version].zip)
- OSGi JARs ZIP file (liferay-dxp-osgi-[version].zip)
- Latest Patching Tool

### 466.2 Install the patch on the Liferay DXP WAR and artifacts

---

1. Create an arbitrary folder. Unzip the dependency artifacts and the Patching Tool into it. The folder contents should look like this:

- [patching-home]/
  - liferay-dxp-dependencies-[version]/ ← Unzipped Dependencies
  - osgi/ ← Unzipped OSGi JARs
  - patching-tool/ ← Unzipped Patching Tool
  - liferay-dxp-[version].war/ ← Liferay DXP WAR File

2. Create the default profile configuration file in the Patching Tool folder: `patching-home/patching-tool/default.properties`. The contents should look like this:

```
patching.mode=binary
war.path=../../patching-home/liferay-dxp-[version].war
global.lib.path=../../patching-home/liferay-dxp-dependencies-[version]
liferay.home=../../patching-home
```

If you're using a different OSGi folder structure, you can specify it as the [Patching Tool Advanced Configuration](/docs/7-2/deploy/-/knowledge\_base/d/patching-tool-advanced-configuration) documentation describes:

```
module.framework.core.path=/osgi-home/osgi/core
module.framework.marketplace.path=/osgi-home/osgi/marketplace
module.framework.modules.path=/osgi-home/osgi/modules
module.framework.portal.path=/osgi-home/osgi/portal
module.framework.static.path=/osgi-home/osgi/static
```

3. Download the patch (fix pack or hotfix) to install and put it in a folder called `patches` in your Patching Tool folder (i.e. `[patching-home]/patching-tool/patches`).

4. Execute the Patching Tool's `info` command:

```
/patching-home/patching-tool> patching-tool info
Loading product and patch information...
Product information:
* installation type: binary
* build number: 7210
* service pack version:
 - available SP version: Not available
 - installable SP version: Not available
* patching-tool version: 2.0.12
* time: 2019-06-03 18:30Z
* host: 91WRQ72 (8 cores)
* plugins: no plugins detected
```

```
Currently installed patches: -
...
```

5. Install the patch.

```
/patching-home/patching-tool> patching-tool.sh install
One patch is ready to be installed. Applying dxp...
Cleaning up: [1%..10%..20%..30%..40%..50%..60%..70%..80%..90%..100%]
Installing patches: [1%..10%..20%..30%..40%..50%..60%..70%..80%..90%...100%]
The installation was successful. One patch is installed on the system.
```

Great! You have successfully patched the artifacts, and they are ready to be deployed on any supported Application Server.

### 466.3 Related Topics

---

Patching Tool Advanced Configuration  
Deploying Liferay DXP



## KEEPING UP WITH FIX PACKS AND SERVICE PACKS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

The *Announcements* section on Liferay's Help Center page lists all fix pack updates, security alerts, product releases, and system updates. The approximate frequency of fix pack and service pack releases is explained here. The *Receive Notifications* sidebar lets you subscribe to the latest updates on products, patches, and system improvements.

Click *Downloads* on the Liferay Digital Experience Platform page to access:

- Latest Release
- Fix Packs
- Service Packs Archive
- Security Advisories
- Patching Tool

Click *Support Information* to access the compatibility matrix, support FAQs, and more.



---

## BACKING UP A LIFERAY DXP INSTALLATION

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Once you have an installation of Liferay DXP running, you should implement a comprehensive backup plan. In case some kind of catastrophic hardware failure occurs, you'll be thankful to have backups and procedures for restoring Liferay DXP from one of them. @product@ isn't very different from other Java web applications that might be running on your application server. Nevertheless, there are some specific components you should include in your backup plan.

The recommended backup plan includes backing up these things:

- Source code
- Liferay DXP's file System
- Liferay DXP's database

---

### **468.1 Backing up Source Code**

---

If you have extended Liferay DXP or have written any plugins, they should be stored in a source code repository such as Git, Subversion, or CVS, unless you're Linus Torvalds, and then tarballs are okay too (that's a joke). You should back up your source code repository on a regular basis to preserve your ongoing work. This probably goes without saying in your organization since nobody wants to lose source code that's taken months to produce. Thus you should include source code in your Liferay DXP backup plan.

Next, let's examine the Liferay DXP installation items you should back up.

---

### **468.2 Backing up Liferay DXP's File System**

---

The Liferay Home folder stores Liferay DXP's properties configuration files, such as `portal-setup-wizard.properties` and `portal-ext.properties`. You should absolutely back them up. In fact, it's best to back up your entire application server and Liferay Home folder contents.

Liferay DXP stores configuration files, search indexes, and cache information in Liferay Home's `/data` folder. If you're using the File System store or the Advanced File System store, the documents and media repository is also stored here by default. It's always important to back up your `/data` folder.

The files that comprise Liferay DXP's OSGi runtime are stored in Liferay Home's `/osgi` folder. It contains all of the app and module JAR files deployed to Liferay DXP. The `/osgi` folder also contains other required JAR files, configuration files, and log files. It's also important to back up your `/osgi` folder.

Liferay Home's `/logs` folder contains Liferay DXP's log files. If a problem occurs on Liferay DXP, the `@product@` log files often provide valuable information for determining what went wrong. The `/data`, `/osgi`, and `/logs` folders are all contained in the Liferay Home folder. Thus, if you're backing up both your application server folder and your Liferay Home folder, you're in good shape.

Remember that if you've configured the document library to store files to a location other than the default location, you should also back up that location.

That covers the Liferay DXP file system locations you should back up. Next, let's discuss how to back up Liferay DXP's database.

### 468.3 Backing up Liferay DXP's Database

---

Liferay DXP's database is the central repository for all of the portal's information. It's the most important component to back up. You can back up the database live (if your database allows this) or by exporting (dumping) the database into a file and then backing up the exported file. For example, MySQL ships with a `mysqldump` utility which lets you export the entire database and data into a large SQL file. This file can then be backed up. On restoring the database you can import this file into the database to recreate the database state to that of the time you exported the database.

If you're storing Liferay DXP's Documents and Media Library files to a Jackrabbit JSR-170 repository database, you should back it up. If you've placed your search index into a database (not recommended; see the Liferay DXP Clustering article for information on using Cluster Link or Solr), you should back up that database too.

If you wish to avoid re-indexing your content after restoring your database, back up your search indexes. This is easiest to do if you have a separate Elastic or Solr environment on which your index is stored. If you're in a clustered configuration and you're replicating indexes, you'll need to back up each index replica.

Restoring your application server, your Liferay Home folder, the locations of any file system-based media repositories, and your database from a backup system should give you a functioning portal. Restoring search indexes should avoid the need to re-index when you bring your site back up after a catastrophic failure. Good, consistent backup procedures are key to recovering successfully from a hardware failure.

## MONITORING LIFERAY DXP

---

These articles show you how to monitor Liferay DXP. Monitoring vital statistics such as Java memory heaps, garbage collection, database connection pools, and the application server helps you optimize performance. Better monitoring means better tuning and thus avoids dangerous runtime scenarios like out of memory errors and wasted heap space.

You'll learn basic monitoring techniques, such as

- Using the Visual VM tool and the JMX Console
- Garbage Collection

Read on to learn more about monitoring Liferay DXP!



---

# MONITORING GARBAGE COLLECTION AND THE JVM

---

Although the tuning parameters give you a good start to JVM tuning, you must monitor GC performance to ensure you have the best settings to meet your needs. There are several tools to help you monitor Oracle JVM performance.

## 470.1 VisualVM

---

VisualVM provides a centralized console for viewing Oracle JVM performance information and its Visual GC plugin shows garbage collector activities.

**Note:** Oracle's JDK has VisualVM bundled (`$JAVA_HOME/bin/jvisualvm`). However, always download and use the latest version from VisualVM's official website.

---

## 470.2 JMX Console

---

This tool helps display various statistics like Liferay DXP's distributed cache performance, application server thread performance, JDBC connection pool usage, and more.

**Note:** The JMX Console is the preferred tool for monitoring application server performance.

---

To enable JMX connections, add these JVM arguments:

```
-Dcom.sun.management.jmxremote=true
-Dcom.sun.management.jmxremote.port=5000
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.ssl=false
```

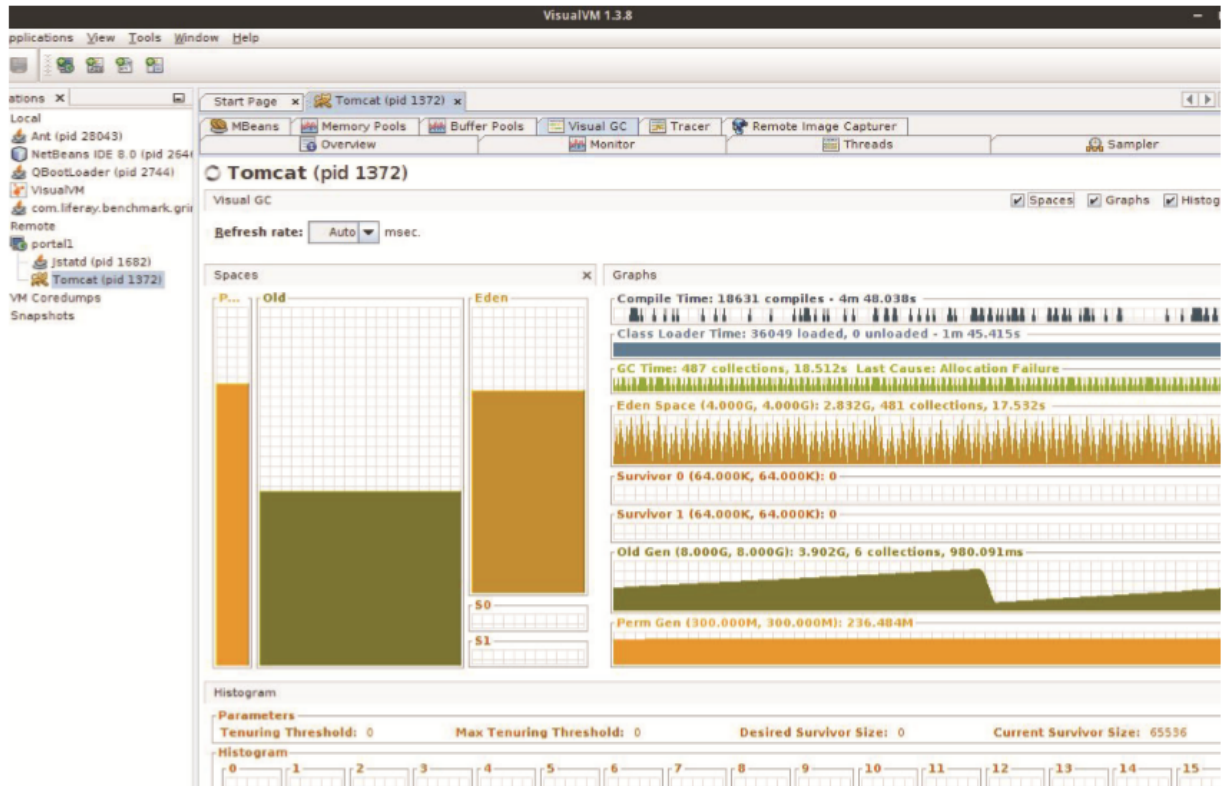


Figure 470.1: VisualVM's Visual GC plugin shows the garbage collector in real-time.

If you're running JMX Console from a another machine, add these JVM arguments too:

```
-Dcom.sun.management.jmxremote.local.only=false
-Dcom.sun.management.jmxremote.rmi.port=5000
-Djava.rmi.server.hostname=[place IP address here]
```

### 470.3 Garbage Collector Verbose Logging

Add these JVM arguments to activate verbose logging for the JVM garbage collector.

```
-verbose:gc -Xloggc:/tmp/liferaygc1.log -XX:+PrintGCDetails
-XX:+PrintGCCause -XX:+PrintGCApplicationConcurrentTime
-XX:+PrintGCApplicationStoppedTime
```

Examining these logs helps you tune the JVM properly.

**Note:** Adding these JVM arguments generates a heap dump if an `OutOfMemoryError` occurs. The dump is written to the heap dump path specified. Specify the path to use:



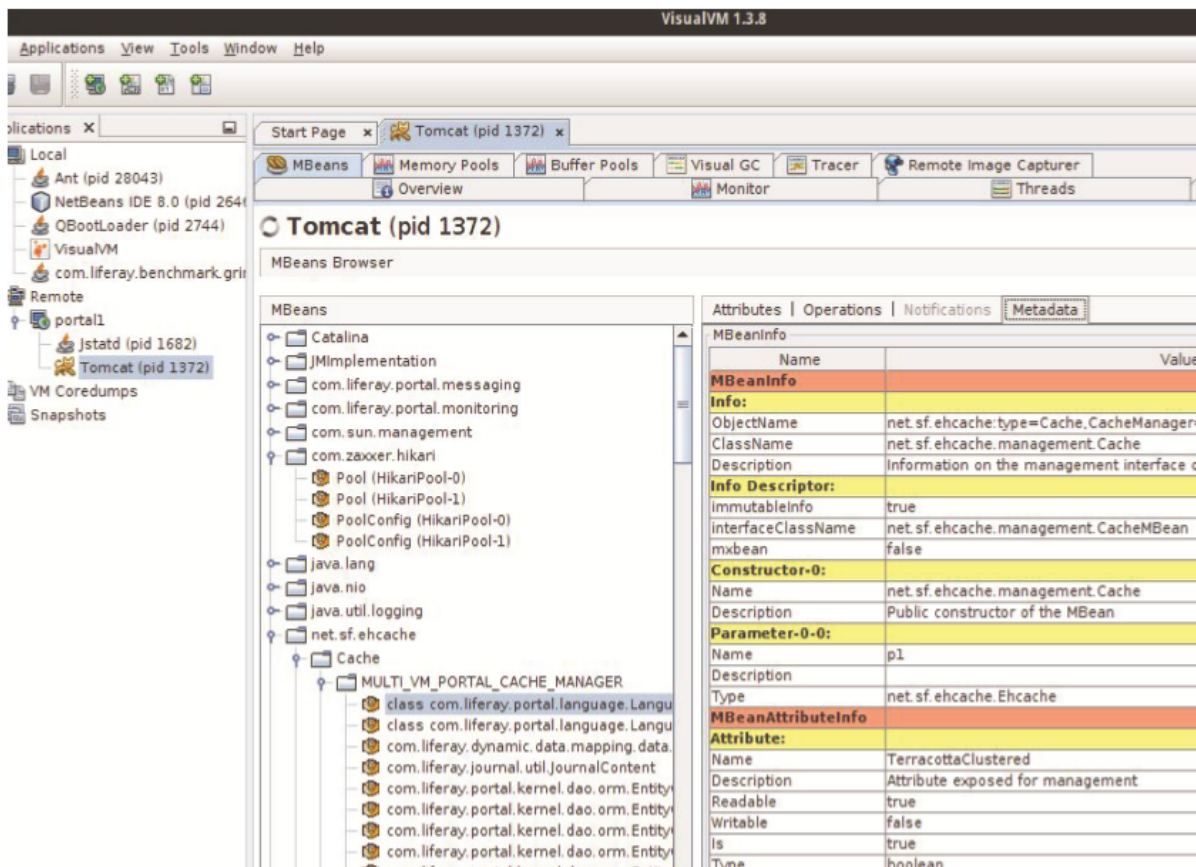


Figure 470.2: VisualVM monitors the JVM using Java Management Extensions.

`-XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/heap/dump/path/`

Garbage collector log files can grow huge. You can use additional arguments like the following ones to rotate the log to a new log file when the current log file reaches a maximum size:

`-XX:+PrintGCDateStamps -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=10  
-XX:GCLogFileSize=50M`

These arguments rotate the logs to up to 10 log files with a maximum size of 50M each. Now you can monitor garbage collection in the JVM and tune it for top performance.



---

## MANAGING LIFERAY DXP WITH LIFERAY CONNECTED SERVICES

---

Liferay Connected Services (LCS) is a set of tools and services for managing and monitoring your Liferay DXP instances. LCS can help you install fix packs, monitor your instances' performance, activate your instances, and help you manage your subscriptions. In other words, LCS is like a butler for the mansion that is Liferay DXP. It's like having a single butler that can serve several mansions at once!

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

---

Before going any further, you should take note of a few key terms used throughout this guide:

**Project:** Represents a group of users belonging to a company or organization. For example, a project can consist of all the users from a project team or business unit, or it can include the entire company.

**Environment:** Represents a physical cluster of servers or a virtual or logical aggregation of servers.

**Server:** Describes a concrete Liferay DXP instance. It can be a standalone server or a cluster node.

As you go through this guide, you'll cover the following topics:

- Getting Started
- LCS Preconfiguration
- Registering Your Liferay DXP Server with LCS
- Using LCS
- Troubleshooting Your LCS Connection

You'll get started with the configuration steps required to use LCS with Liferay DXP.



---

## GETTING STARTED WITH LCS

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).


To use LCS, you must register a server in an LCS environment. An LCS environment represents a physical cluster of servers or a virtual or logical aggregation of servers. Each environment is part of an LCS project. An LCS project represents a group of users belonging to a company or organization. For example, a project can consist of all the users from a project team or business unit, or it can include the entire company.

LCS projects don't initially contain any environments. You must therefore create one before you can register any servers in LCS. The first time you log in to [lcs.liferay.com](https://lcs.liferay.com), LCS presents you with a wizard that walks you through the environment creation process. Click *Get Started* to begin.

Each of these steps corresponds to a step in the wizard:

1. Select the LCS project for your new environment. You can select any of your available LCS projects. Note that each project lists its available subscriptions and whether it supports elastic subscriptions.
2. Name and describe the environment. The name is mandatory, but the description is optional. Although you can enter anything you wish in these fields, it's best to choose a name and description that accurately identify the environment (e.g., Development, Production, Test, etc.). Note that you can change these values after creating the environment.
3. Select the environment's subscription type from the project's available subscriptions. Even if you won't use LCS to activate the servers defined for this environment, you must still select a subscription type. Also note that you can't change this selection after creating the environment.
4. Select whether servers that connect to this environment are part of a cluster. LCS provides additional tools in clustered environments that help you manage the cluster. For example, clustered environments show cluster-specific metrics, and fix packs apply to all cluster nodes. There are a few things to keep in mind if you set the environment as clustered:

# Welcome to Liferay Connected Services



Liferay Connected Services offers tools and services that help you manage your Liferay projects. Click Get Started to begin.

[See more about Liferay Connected Services.](#)

Figure 472.1: Click *Get Started* to begin the wizard.

- You can't change this setting after creating the environment.
  - Each clustered environment can only support nodes that belong to a single cluster. To connect a different cluster's nodes, you must create a separate clustered environment exclusively for those nodes.
  - You must set the portal property `cluster.link.enabled` to `true` in any servers that connect to a clustered environment.
5. Select whether the environment allows elastic subscriptions. Elastic subscriptions let you register an unlimited amount of servers. This is critical for auto-scaling situations in which servers are created and destroyed automatically in response to demand. Elastic environments are also useful for bringing additional servers online on a temporary basis for any other purpose, such as business continuity planning. For more information, see the documentation on elastic subscriptions. Also note that you can't change this selection after creating the environment.
  6. Enable the LCS service you want to use with servers that connect to this environment. The following service is available:  
**Liferay Instance Activation:** Enabling this lets LCS activate any Liferay DXP instances that connect to the environment. If you disable this service, you must activate via an XML file

## Select a Project

✓ Docs Test - Nick

AVAILABLE SUBSCRIPTIONS  
Production, Production.

● Has Elastic Capability

◀ Previous

Next ▶

Figure 472.2: Select the LCS project you want to create the environment in, and click *Next*.

from Liferay support, and such instances must run version 5.0.0 or newer of the LCS client app.

Note that you **must** use LCS for activation of Elastic subscriptions. Otherwise, you don't have to use LCS for activation.

Portal Analytics, Fix Pack Management and Portal Properties Analysis have been removed from the list of available services. For more information about this change, please read this [article](#)

7. A completed form presents your selections. Review them and make any changes that you want. When you're finished, click *Create Environment*.

After creating your environment, the wizard shows a screen that lets you download the LCS client app, download the environment's token file, and go to your project's dashboard in LCS. Before registering a server in your new environment, however, you must complete the preconfiguration steps for that server.

# Describe the Environment

Environment Name

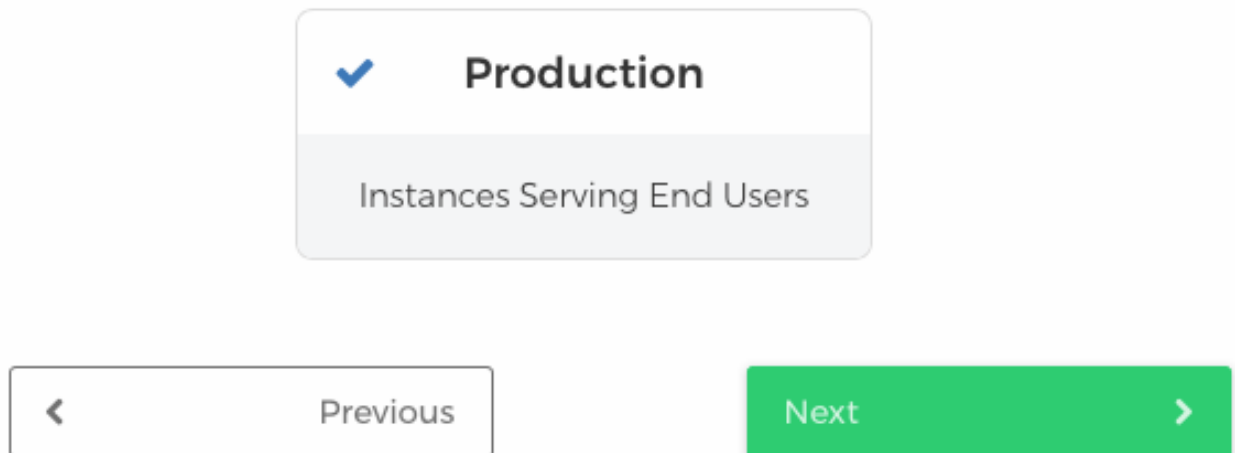
Environment Description

< Previous **Next** >

Figure 472.3: Name and describe the environment, then click *Next*.



## Select the Subscription Type



**Production**

Instances Serving End Users

< Previous

Next >

Figure 472.4: Select the environment's subscription type, then click *Next*.

## Prepare the Environment for Clustering

Will servers connecting to this environment be in a cluster?

NO

[< Previous](#)

[Next >](#)

Figure 472.5: Select whether this is a clustered environment, then click *Next*.

# Enable Elastic Subscriptions

Do you want an elastic environment?  **NO**

[< Previous](#) [Next >](#)

Figure 472.6: Select whether this is an elastic environment, then click *Next*.

## Enable Services

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>LIFERAY INSTANCE ACTIVATION</b> <input type="checkbox"/> <b>No</b></p> <p>A Liferay instance equals to one copy of Liferay DXP or Liferay Portal. Enabling this service lets LCS activate instances installed on the servers you connect to this environment. Note that even with this service disabled, you still must have an active subscription of the selected subscription type to connect servers to this environment. To disable this service, servers must be running version 5.0.0 or newer of the Liferay Connected Services Client app.</p> | <p><b>PORTAL ANALYTICS</b> <input checked="" type="checkbox"/> <b>Yes</b></p> <p>Enabling this service lets LCS collect and display performance metrics for your portal and its sites. This includes page and portlet load times, JVM garbage collector and memory metrics, and more.</p>                                                                                                                       |
| <p><b>FIX PACK MANAGEMENT</b> <input checked="" type="checkbox"/> <b>Yes</b></p> <p>Enabling this service lets you use LCS to check for and download portal fix pack updates. This requires LCS to determine your servers' currently installed fix packs.</p>                                                                                                                                                                                                                                                                                                 | <p><b>PORTAL PROPERTIES ANALYSIS</b> <input checked="" type="checkbox"/> <b>Yes</b></p> <p>Enabling this service lets you use LCS to compare your portal's current property settings to their default values. Note that LCS doesn't access security sensitive properties. Such properties are blacklisted and their data isn't collected by LCS.</p> <p><input type="text" value="BLACKLISTED PROPERTIES"/></p> |

[< Previous](#) [Next >](#)

Figure 472.7: Enable or disable the LCS services you want to use for servers that connect to the environment, then click *Next*.

## Review Your Environment

|                                                     |                                                                  |
|-----------------------------------------------------|------------------------------------------------------------------|
| Environment Name                                    | <input type="text" value="Production"/>                          |
| Description                                         | <input type="text" value="This is the production environment."/> |
| Subscription Type                                   | <input type="text" value="Production"/>                          |
| Cluster                                             | <input type="checkbox"/> No                                      |
| Elastic <span>(i)</span>                            | <input type="checkbox"/> No                                      |
| <b>Services</b>                                     |                                                                  |
| Liferay Instance Activation <span>(i)</span>        | <input type="checkbox"/> No                                      |
| Portal Analytics <span>(i)</span>                   | <input checked="" type="checkbox"/> Yes                          |
| Fix Pack Management <span>(i)</span>                | <input checked="" type="checkbox"/> Yes                          |
| Portal Properties Analysis <span>(i)</span>         | <input checked="" type="checkbox"/> Yes                          |
| <input type="text" value="BLACKLISTED PROPERTIES"/> |                                                                  |
| <input type="button" value="Create Environment"/>   |                                                                  |

Figure 472.8: This form contains each of your selections from the previous steps. Make any changes you want, then click *Create Environment*.

## LCS PRECONFIGURATION

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

---

Before registering your server with LCS, there are a few things you must configure. The sections in this guide walk you through these steps:

1. Downloading the LCS Client App
2. Preconfiguring LCS to Connect Through a Proxy
3. Ensuring Access to LCS
4. NTP Server Synchronization
5. Configuring WebSphere: This is only necessary if you're running Liferay DXP on the WebSphere application server.
6. Installing the LCS Client App

The last section in this guide shows you how to upgrade the LCS client app once your server is registered with LCS. We highly recommend that you upgrade the app whenever Liferay releases a new version of it.

---

**Note:** You must use LCS for activation of Elastic subscriptions. Otherwise, you don't have to use LCS for activation. You can instead request an XML activation key from Liferay Support.

---

### 473.1 Downloading the LCS Client App

---

The LCS client app is included in each Liferay DXP bundle and autodeploys when the bundle starts. The included version of the app, however, may be outdated. To get the latest version of the LCS client app, you must first download it via Liferay Marketplace.

**Note:** Even though Liferay Marketplace and this guide use the term *purchase*, the LCS client app is free of charge. The purchase process for a free app in Liferay Marketplace adds the app to your Liferay project, much like downloading a free app in a mobile app store adds the app to your account.

Use these steps to purchase and download the app (if you've already purchased the app, you can skip to step 3 to download it):

1. Navigate to the LCS client app in Liferay Marketplace. Sign in to Marketplace, then click the LCS client app's *Free* button.




Figure 473.1: Click the app's *Free* button to begin the purchase process.

2. Select your project, accept the license agreement, and then click the *Purchase* button. Marketplace then displays your receipt.
3. On the receipt, click *See Purchased*. This shows where you can download the LCS client app. To download the app, click the *App* button next to the latest version of the app.

**Note:** If you must download the LCS client app later, such as when [upgrading it](#upgrading-the-lcs-client-app), select *Purchased Apps* from the User menu at the top-right of Liferay Marketplace. On the Purchased Apps screen, select the project you associated with the LCS client app and then select the app. This takes you to the same downloads page shown in the screenshot.

[ Click the *App* button next to the version of the app you want to download.](./images/lcs-client-download-page.png){./images/lcs-client-download-page.png}

## Receipt

---

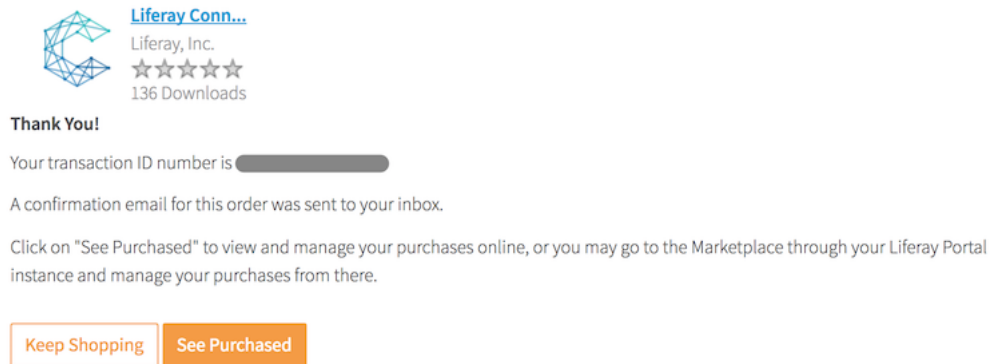


Figure 473.2: Liferay Marketplace displays your receipt for the LCS client app.

Great! You've successfully downloaded the LCS client app. Before installing it, however, there are a few additional pre-configuration steps you should complete. These appear next; then you'll learn how to install the app.

---

**Note:** If your server connects to the Internet through a proxy, you must configure your server or the LCS client app **before** deploying the app. The following section contains instructions on this. If your server doesn't connect through a proxy, skip this section.

---

### 473.2 Preconfiguring LCS to Connect Through a Proxy

---

If your server connects to the Internet through a proxy, you must set some properties **before** deploying the LCS client app. There are two ways to do so—choose only one:

1. As JVM app server arguments.
2. As LCS client app properties.

---

**Note:** Use only one of these methods to configure your server to connect through a proxy.

---

### 473.3 JVM App Server Arguments

---

To set the proxy properties in your server, set them as JVM app server arguments. Set these properties to the appropriate values for your proxy:

```
-Dhttp.proxyHost=
-Dhttp.proxyPort=
-Dhttp.proxyUser=
-Dhttp.proxyPassword=
-Dhttps.proxyHost=
-Dhttps.proxyPort=
```

Note that the user, password, and https properties are only needed if your proxy requires authentication.

#### 473.4 LCS Client App Properties

---

To set the proxy properties via the LCS client app, you must create and deploy a config file containing the properties. Follow these steps to do so:

1. Create the config file `com.liferay.lcs.client.configuration.LCSConfiguration.config`. In the steps that follow, you'll set the proxy properties in this file.
2. Set these proxy\* properties to the appropriate values for your proxy:

```
proxyHostName=""
proxyHostPort=""
```

3. If your proxy requires authentication, pass the credentials via these properties:

```
proxyHostLogin=""
proxyHostPassword=""
```

4. If your proxy requires NTLM authentication, you must also populate these properties:

```
proxyAuthType="ntlm"
proxyDomain=""
proxyWorkstation=""
```

Be sure to set `proxyDomain` and `proxyWorkstation` to the appropriate values for your proxy. Note that you can leave `proxyWorkstation` blank if you don't need it.

5. Deploy the config file to `osgi/configs`.



### 473.5 Ensuring Access to LCS

---

For the LCS client app to work, it must be able to access the following DNS names. If your server is behind a proxy and/or a firewall, then you must open access to these:

- lcs.liferay.com
- lcs-gateway.liferay.com

As an added security measure, you can also restrict traffic to HTTPS. The next section discusses NTP server synchronization.

### 473.6 NTP Server Synchronization

---

For LCS to work properly, the application server running Liferay DXP should be synchronized with a time server. If it's not, you may get log errors similar to these:

```
ERROR [pool-6-thread-3][HandshakeTask:68] java.lang.RuntimeException: Handshake expired.
Check that the server is synchronized with an NTP server.

WARN [liferay/hot_deploy-1][LCSHotDeployMessageListener:186] LCS portlet is not connected
java.lang.RuntimeException: com.liferay.jsonwebserviceclient.JSONWebServiceInvocationException:
com.fasterxml.jackson.core.JsonParseException: Unrecognized token 'oauth_problem': was expecting
('true', 'false' or 'null')_ at [Source: oauth_problem=timestamp_refused&oauth_acceptable_timestamps=1477311475-1477312075;
line: 1, column: 14] [Sanitized]
```

For information on how to synchronize your application server with a time server, see your application server's documentation.

### 473.7 Configuring WebSphere

---

IBM® WebSphere® is a trademark of International Business Machines Corporation, registered in many jurisdictions worldwide.

If you're running the WebSphere application server, then there are some additional configuration steps you must take before deploying the LCS client app:

1. Shut down the application server.
2. Add these properties in a portal-ext.properties file:

```
module.framework.properties.org.osgi.framework.bootdelegation=\
 __redirected,\
 com.sun.ccpp,\
 com.sun.ccpp.*,\
 com.liferay.aspectj,\
 com.liferay.aspectj.*,\
```

```

com.liferay.portal.servlet.delegate,\
com.liferay.portal.servlet.delegate*,\
com.sun.crypto.*,\
com.sun.image.*,\
com.sun.jmx.*,\
com.sun.jna,\
com.sun.jndi.*,\
com.sun.mail.*,\
com.sun.management.*,\
com.sun.media.*,\
com.sun.msv.*,\
com.sun.org.*,\
com.sun.syndication,\
com.sun.tools.*,\
com.sun.xml.*,\
com.yourkit.*,\
com.ibm.crypto.*,\
sun.*,\
javax.validation,\
javax.validation.*,\
jdk.*,\
weblogic.jndi,\
weblogic.jndi.*\

```

3. In your Liferay DXP installation, delete the `osgi/state` folder.
4. Start the application server.
5. Navigate to the WebSphere console in a browser.
6. Select your server and navigate to *Java and Process Management* → *Process Definition* → *Additional Properties*.
7. Select *Java Virtual Machine* → *Custom Properties*.
8. Click *New*, and enter the following:
  - Name: `com.ibm.crypto.provider.DoRSATypeChecking`
  - Value: `false`
9. Click *Save*, then *OK* to apply changes to the master configuration.

Note that for LCS client app versions prior to 5.0.0, you must also change the value of the `digital.signature.algorithm.provider` property in the app's `portlet.properties` file to `IBMJCE`:

```
digital.signature.algorithm.provider=IBMJCE
```

### 473.8 Installing the LCS Client App

---

Once you've addressed the above pre-configuration steps, you're ready to install the LCS client app. Follow these steps to install the app:

1. In your Liferay Home folder (usually the parent folder of the application server's folder), delete this file:

osgi/marketplace/Liferay Connected Services Client.lpkg

2. Place the new Liferay Connected Services Client.lpkg in osgi/marketplace.

Great! Now you're all set to register your server with LCS.

The next section shows you how to upgrade the LCS client app. We highly recommend that you do this whenever Liferay releases a new version of the app.

## 473.9 Upgrading the LCS Client App

---

Your server should always be running the latest version of the LCS client app. There are two ways to upgrade the app, depending on the exact LCS pre-configuration steps you followed:

1. Via Liferay Marketplace *inside* Liferay DXP. Use this method if you don't need to configure the LCS client app (e.g., to connect through a proxy) before it deploys.

---

**\*\*Note:\*\*** If you choose this method and have a clustered environment, you must perform the upgrade separately on each node in your cluster. Therefore, you may prefer to upgrade manually as detailed in the next step to ensure that all your nodes are running the exact same version of the LCS client app.

---

To perform the upgrade, first navigate to *Control Panel* &arr; *Apps* &arr; *Purchased*. Apps needing an update are listed first. Click *Update* next to the LCS client app. Note that you may need to restart your server for the upgrade to complete.

2. Manually, after downloading the LCS client app's LPKG file to your machine. Use this method if you must pre-configure the LCS client app to connect through a proxy.

---

**\*\*Note:\*\*** If you used JVM app server arguments to configure your server to connect through a proxy, then you don't need to pre-configure the LCS client app to connect through the same proxy.

---

To update the LCS client app manually, follow the previous sections in this guide for downloading and pre-configuring the app. Then deploy it to ``[Liferay Home]/deploy`` as you would any other app.

Contact Liferay Support if you need additional assistance with the upgrade process.



## REGISTERING YOUR LIFERAY DXP SERVER WITH LCS

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

Follow these steps to register your Liferay DXP server with LCS:

1. Ensure that you've completed the LCS preconfiguration steps.
2. Log in to [lcs.liferay.com](https://lcs.liferay.com). This takes you to your company's LCS project. If your company has multiple projects, from the menu to the right of the Dashboard tab select the project that's getting a new server.

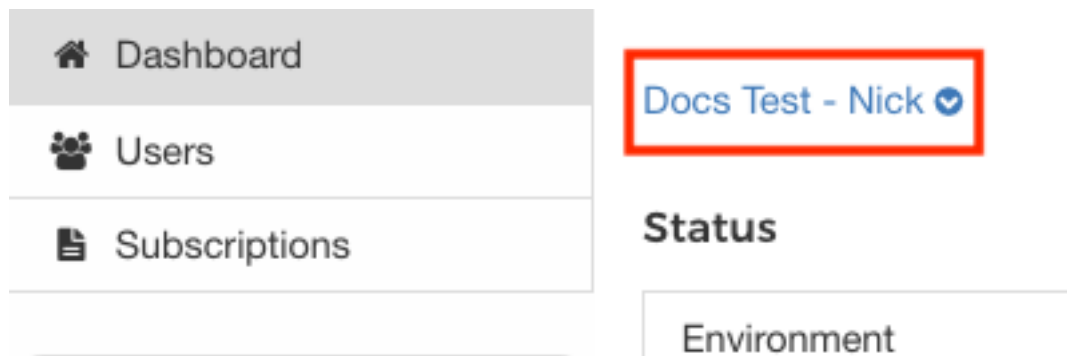


Figure 474.1: Select your LCS project from the menu highlighted by the red box in this screenshot.

3. Select or create the environment in which to register this server. If you're using LCS for activation, upon connection to LCS your server consumes an activation key from the subscription type assigned to the environment. Note that a subscription type can only be assigned to an environment when creating the environment. If you have sufficient permissions in your company's project, you can create a new environment by selecting *Add Environment*.

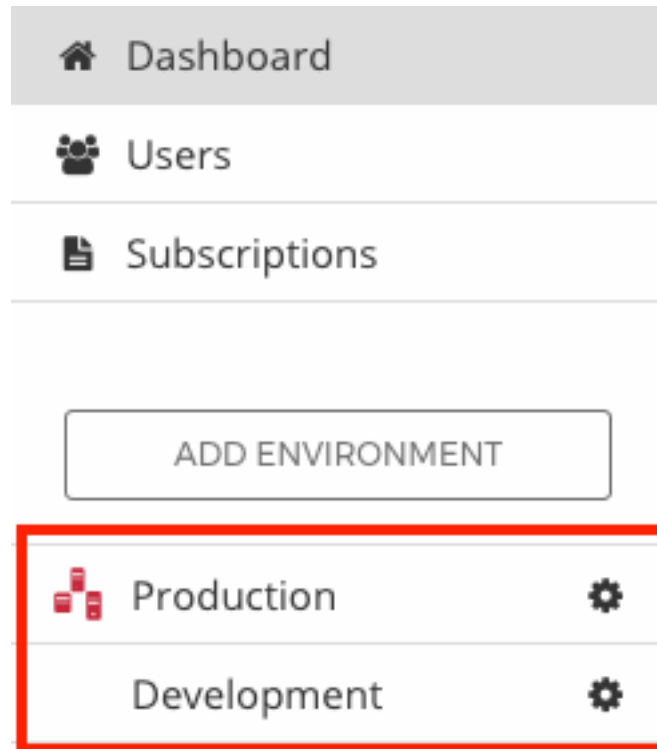


Figure 474.2: You must register your server in an LCS environment. The red box in this screenshot highlights environments.

4. Select the environment's *Registration* tab. This is where you manage and download the environment's token file, that registers servers in the environment.

In the Registration tab's *Services* section, change the Liferay Instance Activation setting, if needed. Note that if you change this option and there are servers already registered in the environment, you must regenerate the token file and use it to reconnect those servers to LCS. You'll regenerate and/or download the token in the next step.

Additionally, If you disable this service, you must activate via an XML file from Liferay support, and such instances must run version 5.0.0 or newer of the LCS client app.

Liferay Instance Activation is either enabled or disabled for all servers that connect to this environment. If Portal Property Analysis is selected, you can prevent LCS environment.

5. What you do now depends on what you did in the previous step:

**Changes to Liferay Instance Activation:** Regenerate and download the token. Regenerating a token causes all servers using the old token to disconnect from LCS. You must reconnect them using the new token.

**No changes to LCS service selections:** Download the token.

6. Place the token file in your server's [Liferay Home]/data folder. Note that Liferay Home is usually the parent folder of the application server's folder. If your server is running, it should connect to LCS in about one minute. If your server isn't running, it connects to LCS on startup.
7. Celebrate! Your Liferay DXP server is registered in LCS. If for some reason it isn't, see the LCS troubleshooting article.

[Fix Packs](#)
[Registration](#)
[Environment Settings](#)

### Liferay Connected Services Client

Please make sure that Liferay Connected Services client app is installed in your Liferay instance. We highly recommend downloading the latest version from Liferay Marketplace.

[Download Client](#) 


### Environment Token File


To connect a server to this environment, download the environment's token file and place it in your Liferay instance's data folder. This file ensures that the server connects to this environment. It also contains the credentials of the user who generated it, and the selection of LCS services to enable. Note that the connection occurs automatically, shortly after placing the file in the data folder. You don't need to restart the server or redeploy the LCS client.


[Download Token](#)


### Token Regeneration

Regenerate the token if you wish to enable or disable services, or if the user who created the token is no longer assigned to the project.


**Liferay Instance Activation** 


**Portal Analytics** 

**Fix Pack Management** 

**Portal Properties Analysis** 

Liferay Connected Services does not access sensitive data in your portal properties. Properties containing sensitive data are blacklisted and not sent to Liferay Connected Services.

[Show Blacklisted Properties](#) 

 Any servers using the previous token will be disconnected from Liferay Connected Services and will not be able to reconnect until the new token is installed.

[Regenerate Token](#)

Figure 474.3: An environment's Registration tab lets you manage the token file used to register your server in the environment.

**Note:** You may be wondering what happens if LCS goes offline. Don't worry, this doesn't cause a rift in the space-time continuum. LCS is deployed on a global cloud infrastructure set up for automatic failure recovery. The potential for non-availability is very low. In the event of an outage, however, registered servers maintain a local copy of their uptime information to transmit to LCS when it comes back online. If you use LCS for activation, active subscriptions have a 30-day grace period to re-establish connectivity and remain valid. This is ample time for LCS to come back online.

---

#### **474.1 Determining Your Server's LCS Connection Status**

---

In Liferay DXP, you can view your LCS connection status in the LCS client app. Access the client by clicking *Control Panel* → *Configuration* → *Liferay Connected Services*.

Here's a full description of what a connected LCS client app displays:

**Connection Uptime:** The duration of the client's connection with LCS.

**Last Message Received:** The time the LCS client received the latest connection message from LCS. These messages occur only upon connection/reconnection and are unrelated to server metrics. It's therefore common for a long period of time to pass before the client receives another such message for a reconnection event.

**Services:** The LCS services enabled for this server. Note that all servers in an environment use the same set of LCS services. LCS services can't be controlled on a server-by-server basis.

Note: Portal Analytics, Fix Pack Management and Portal Properties Analysis have been removed from the list of available services. For more information about this change, please read this article

**Project Home:** A link to this server's LCS project.

**Environment:** A link to this server's LCS environment.

**Server Dashboard:** A link to the server on LCS.



**Connected**

---

**Connection**

|                                      |                         |
|--------------------------------------|-------------------------|
| Connection Uptime <span>?</span>     | 00:38:37                |
| Last Message Received <span>?</span> | Oct 19, 2018 4:02:30 PM |

---

**Services**

- Portal Properties Analysis
- Fix Packs Management
- Portal Analytics

To enable or disable services, you need to [regenerate the token](#).

---

**Liferay Connected Services Sites**




|                  |                                                                                                                        |
|------------------|------------------------------------------------------------------------------------------------------------------------|
| Project Home     | <a href="#">Docs Test - Nick</a>  |
| Environment      | <a href="#">Production</a>          |
| Server Dashboard | <a href="#">Server 02</a>           |

Figure 474.4: The server is connected to LCS.



---

## USING LCS

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

Once your Liferay DXP server is connected to LCS, you can get down to the business that LCS is designed for—managing your servers. If you’re not already there, log in with your account on [lcs.liferay.com](https://lcs.liferay.com). This is where you’ll manage environments, register servers and more.

This articles in this section each detail one or more of LCS’s features:

- **What LCS Stores About Your Liferay DXP Servers:** For LCS to work, the LCS servers must store certain information about your servers. Sensitive data, however, isn’t stored on the LCS servers. This article describes the data that LCS does and doesn’t store.
- **Managing LCS Users in Your Project:** Learn how to manage your LCS project’s users by assigning them roles.
- **Using the Dashboard:** Learn how to manage your LCS projects and access your environments and servers in LCS.
- **Managing LCS Environments:** Learn how to create and manage your LCS project’s environments. This includes instructions on generating tokens for an environment’s servers.
- **Managing LCS Servers:** Learn how to manage your servers in LCS. This includes viewing server status and editing server settings.
- **Managing Your LCS Account:** Learn how to manage your LCS account. This includes setting general account preferences, managing LCS web notifications, and configuring LCS to send you notification emails when specific events occur in your LCS projects.
- **Managing Liferay DXP Subscriptions:** Learn how to view and manage your Liferay DXP subscriptions for the servers in your LCS project.
- **Understanding Environment Tokens:** Learn about the environment tokens that you use to connect your servers to LCS.



---

# WHAT LCS STORES ABOUT YOUR LIFERAY DXP SERVERS

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

To protect your users' privacy, LCS only stores system-specific data. LCS doesn't gather or store data on your users.

By default, LCS stores the following information about your server:

- Portal build number and edition
- Patching Tool Version
- LCS Client Build Number
- Application Server Name
- Database Name
- File Encoding
- OS Name and Version
- Timezone
- IP Address
- Java Version and Java Options
- Number of Processor Cores
- File System Usage
- Memory Usage

The other data LCS stores depends on the services you enabled in your environment token, and whether your server was connected before certain services were removed. For more information on this, see [Registering Servers with LCS](#). If you enabled the following services, LCS gathered and stored the data listed for each:

- **Portal analytics:**
  - Portal and portlet metrics

- JVM metrics
- Cache and server metrics
- **Fix pack management:**
  - Patches installed on the server
- **Portal properties analysis:**
  - portal.properties (except sensitive data)

Sensitive data is any key-value pair that contains user names or passwords. For example, LCS did not store the following properties because they contain sensitive data:

```
omniadmin.users
ldap.security.credentials.0, ldap.security.credentials.1, ldap.security.credentials.2 ...
facebook.connect.app.secret
auth.token.shared.secret
auth.mac.shared.key
captcha.engine.recaptcha.key.private
amazon.secret.access.key
tunneling.servlet.shared.secret
microsoft.translator.client.secret
dl.store.s3.secret.key
auto.deploy.glassfish.jee.dm.passwd
```

LCS also did not store properties that end in .password, besides the following non-sensitive properties:

```
portal.jaas.plain.password
portal.jaas.strict.password
login.create.account.allow.custom.password
```

LCS also allowed you to prevent it from analyzing specific properties of your choosing, by defining blacklisted properties.

LCS is no longer gathering or storing the data listed above, that was associated with enabled services.

## MANAGING LCS USERS IN YOUR PROJECT

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

The Users section of LCS is where you manage the LCS users that are part of your project. It's here that you can grant or revoke LCS Roles. To manage users, first click the *Users* tab just below the Dashboard tab on the upper-left of your screen.

**Note:** You can't add users to your project via the LCS UI or the LCS client app. To add users to your project, you must contact Liferay support.

The *Users* tab displays a list of the users in your project. This list includes each user's name, email, image, LCS Roles, and a *Manage Roles* button. Each LCS user must have an assigned Role. The following Roles are available:

**LCS Administrator:** All LCS functionality is available to administrators. This is the only Role that can manage other users' Roles.

**LCS Environment Manager:** All LCS functionality is available in the scope of an environment, with the exception of managing other users.

**LCS Environment Viewer:** Has read-only access in the scope of an environment.

You should note that each of these LCS Roles assume users already have the LCS User Role in their Liferay.com accounts. The LCS User Role is granted automatically the first time a user logs into LCS. The actions that can be performed by each of the LCS Roles are detailed in the below permissions matrix.

### LCS Permissions Matrix

| Action                                          | LCS Administrator | LCS Environment Manager | LCS Environment Viewer | Access LCS |
|-------------------------------------------------|-------------------|-------------------------|------------------------|------------|
| Access Any Environment                          | true              | false                   | false                  | true       |
| Access a Particular Environment                 | true              | true                    | true                   | true       |
| Manage Users                                    | true              | false                   | false                  | true       |
| Create and Delete Environments                  | true              | false                   | false                  | true       |
| Edit Any Environment                            | true              | false                   | false                  | true       |
| Edit a Particular Environment                   | true              | true                    | false                  | true       |
| Server Registration in Any Environment          | true              | false                   | false                  | true       |
| Server Registration in a Particular Environment | true              | true                    | false                  | true       |
| Install Fix Packs in Any Environment            | true              | false                   | false                  | true       |
| Install Fix Packs in a Particular Environment   | true              | true                    | false                  | true       |

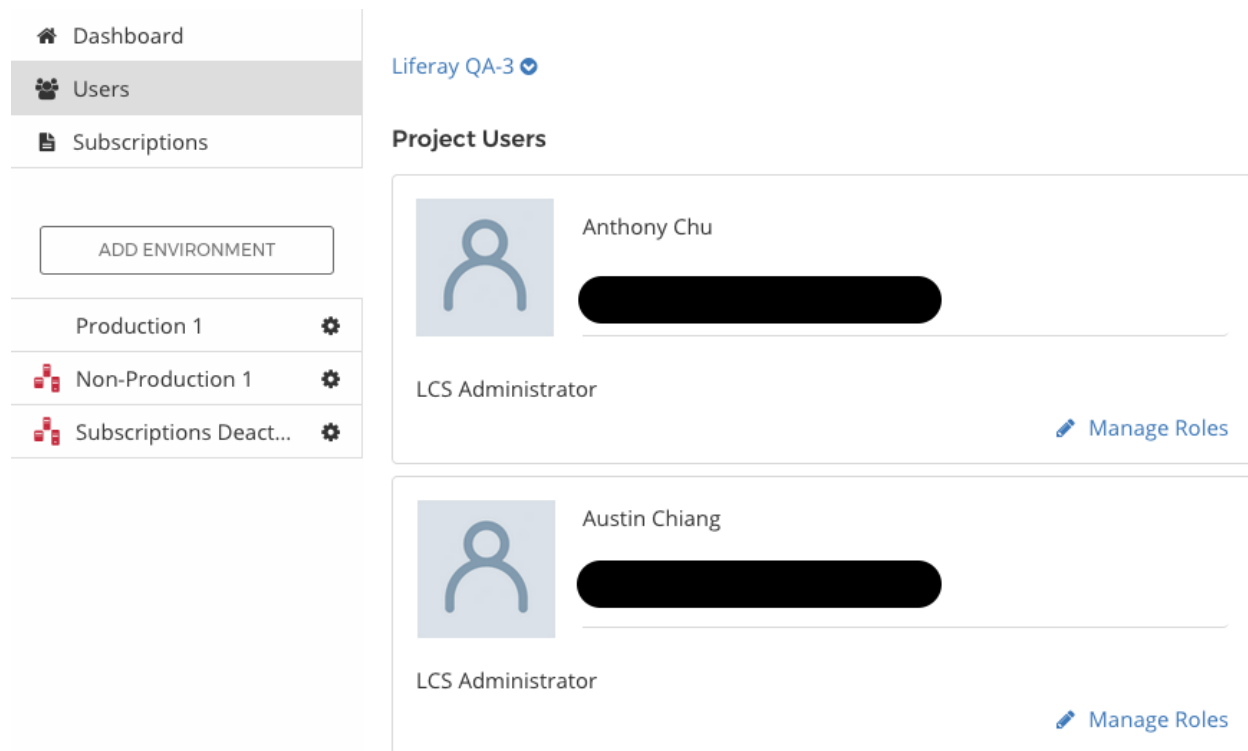


Figure 477.1: The Users tab lets you manage the LCS users in your project.

Now that you know what Roles are available in an LCS project and what they do, you're ready to learn how to manage them.

### 477.1 Managing LCS Roles

---

Follow these steps to manage a user's LCS Roles:

1. Click the user's *Manage Roles* button.
2. To revoke a Role, click *Revoke Role* for that Role.
3. To assign a Role, choose the Role (and environment, if applicable) and click *Assign*.

---

**Note:** A user can't have an environment Role (e.g., LCS Environment Manager, LCS Environment Viewer) and the LCS Administrator Role at the same time.

---



**Bialar Crais's Roles**

| Role                    | Environment    | Action                        |
|-------------------------|----------------|-------------------------------|
| LCS Environment Manager | Environment 11 | <a href="#">✕ Revoke Role</a> |
| LCS Environment Viewer  | Environment 12 | <a href="#">✕ Revoke Role</a> |

**Assign Role**

Role

---

Environment

Figure 477.2: You can assign or revoke a user's LCS Roles.



---

## USING THE DASHBOARD

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

The LCS Dashboard shows a project's environments and servers. If you're not already at the Dashboard, click it near the upper left-hand corner of your LCS site. Clicking *Dashboard* takes you to the project view. From there, you can get to the environment view and the server view. Each of these views gives you a different look into certain aspects of your LCS project. You'll start with the project view.

---

### 478.1 Using the Project View

---

You can get to the project view at any time by clicking the *Dashboard* tab near the upper left-hand corner of your LCS site. The project appears to the right of this tab, with a drop-down arrow for switching between projects if you have more than one. You can also switch between projects from the user menu at the top right of the Dockbar. The project view contains a Status table that lists status messages for each server in your project. For example, a status message appears for a server when the server is offline. Status messages also appear for servers when fix packs are available, monitoring is unavailable, the patching tool is unavailable, or other events occur that relate to LCS.

LCS lists the environments in your project on the left side of the screen. You can also create new environments here by clicking the *Add Environment* tab (more on this shortly). To view an environment's settings, click the environment's gear icon. Clicking an environment shows more information about it. This takes you to the environment view. Also note that each environment's icon indicates the environment's type and status:

**Red icon:** There's a problem with one or more of the environment's servers.

**Green icon:** The environment's servers are operating properly.

**Icon with a circle:** The environment's servers are clustered.

The screenshot displays the LCS project dashboard. On the left is a navigation sidebar with 'Dashboard' selected, and options for 'Users' and 'Subscriptions'. Below these is an 'ADD ENVIRONMENT' button and two environment entries: 'Production' and 'Development', each with a gear icon. The main content area shows the user 'Docs Test - Nick' and a 'Status' section. The status section contains a table with columns for Environment, Server, and Message.

| Environment | Server    | Message                                                           |
|-------------|-----------|-------------------------------------------------------------------|
| Production  | Server 01 | The server is offline.                                            |
| Production  | Server 02 | The server data is not initialized. Please try again in a minute. |

Figure 478.1: The LCS project view shows an overview of your LCS project.

## MANAGING LCS ENVIRONMENTS

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

---

Environments are the key components of your LCS project. When you register a server in LCS, you do so in an environment. An environment is therefore the gateway to managing and monitoring your servers in LCS.

### 479.1 Creating Environments

---

The first time you log in to LCS, a wizard walks you through each step required to create your project's first environment. The getting started article explains this in detail. You can create additional environments via the same wizard or a simple form.

To create an environment, click the *Add Environment* button from the Dashboard. This opens the New Environment form. Each section in this form corresponds to a step in the wizard. If you want to use the wizard instead, click the *Open Wizard* link at the top of the form. See the getting started article for a description of each setting in the form and wizard.

**Note:** When creating an environment, make your selections carefully for the *Subscription Type*, *Cluster*, and *Elastic* fields. You can't change them after creating the environment.

---

### 479.2 Working with Environments

---

## New Environment

**i** You can instead create an environment with a wizard, which describes each setting of this form in detail. [Open Wizard](#)

|                                                     |                                                                  |
|-----------------------------------------------------|------------------------------------------------------------------|
| Environment Name                                    | <input type="text" value="Production, Development, Testing..."/> |
| Description                                         | <input type="text" value="This is the production environment."/> |
| Subscription Type                                   | <input type="text" value="Production"/>                          |
| Cluster                                             | <input type="checkbox"/> No                                      |
| Elastic <b>i</b>                                    | <input type="checkbox"/> No                                      |
| Services                                            |                                                                  |
| Liferay Instance Activation <b>i</b>                | <input checked="" type="checkbox"/> Yes                          |
| Portal Analytics <b>i</b>                           | <input checked="" type="checkbox"/> Yes                          |
| Fix Pack Management <b>i</b>                        | <input checked="" type="checkbox"/> Yes                          |
| Portal Properties Analysis <b>i</b>                 | <input checked="" type="checkbox"/> Yes                          |
| <input type="text" value="BLACKLISTED PROPERTIES"/> |                                                                  |
| <input type="button" value="Create Environment"/>   |                                                                  |

Figure 479.1: The New Environment form lets you create environments.

Clicking an environment on the left-hand side of the Dashboard takes you to the environment view, which lets you manage an environment in your LCS project.

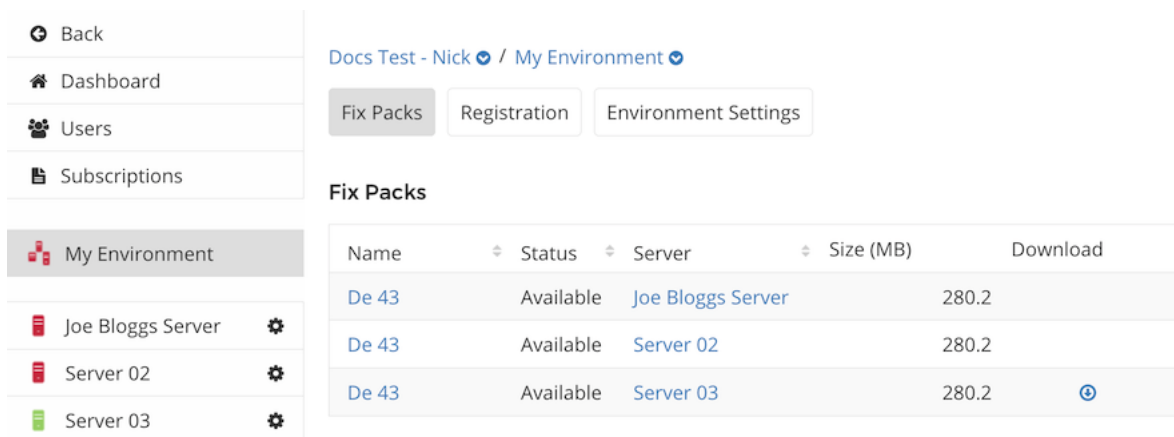
The UI is segmented into three tabs:

1. **Fix Packs:** View and apply fix packs for the environment's servers. This tab only appears if a server is registered in the environment. A table displays the following information for each fix pack:

- **Name:** The fix pack's name.
- **Status:** The fix pack's status.
- **Server:** The server the fix pack can be applied to.
- **Size:** The fix pack's size. This only appears if the server is running.
- **Download:** A button to download the fix pack to the server. This only appears if the server is running.

Once a fix pack downloads, LCS prompts you to restart your server, which installs any downloaded fix packs. Note that you must start your server with the privileges required to write to the disk location where patches are stored and processed (the `patching-tool` folder). To use LCS to install fix packs across a cluster, follow the same procedure. LCS downloads and installs fix packs simultaneously across all nodes—you don't have to handle each separately.

2. **Registration:** Generate and download *environment tokens* that connect your servers to LCS.
3. **Environment Settings:** Change the environment's name, location, and description. You can also see if the environment allows clustered servers and view the environment's subscription type. Click the *Save* button to save any changes you make in the Environment Settings tab. You can also delete the environment by clicking *Delete Environment*.



The screenshot shows the LCS environment view. On the left is a sidebar with navigation options: Back, Dashboard, Users, Subscriptions, and My Environment (selected). The main content area has a breadcrumb 'Docs Test - Nick / My Environment' and three tabs: Fix Packs (selected), Registration, and Environment Settings. Below the tabs is a 'Fix Packs' section with a table:


| Name  | Status    | Server            | Size (MB) | Download                                                                              |
|-------|-----------|-------------------|-----------|---------------------------------------------------------------------------------------|
| De 43 | Available | Joe Bloggs Server | 280.2     |                                                                                       |
| De 43 | Available | Server 02         | 280.2     |                                                                                       |
| De 43 | Available | Server 03         | 280.2     |  |

Figure 479.2: The LCS environment view shows an overview of an LCS environment.





---

## MANAGING LCS SERVERS

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

---

Clicking a server in the Dashboard or environment view takes you to the server view. Server view provides detailed information about a server. To protect your users' privacy, LCS doesn't gather, store, or analyze user data.

Server view is segmented into six tabs:

**Page Analytics:** This service has been disabled, if you enabled it earlier you can see here the past history for metrics on page views and load times.

**Snapshot Metrics:** This service has been disabled, if you enabled it earlier you can see here the past history for application, JVM, and server metrics.

**Fix Packs:** This service has been disabled, if you enabled it earlier you can see here the past history for the server's available and installed fix packs.

**Portal Properties:** This service has been disabled, if you enabled it earlier you can see here the past history for your portal's properties and their settings.

**Details:** Displays general information about your Liferay DXP installation, Java version, and hardware.

**Server Settings:** View or change your server's name, location, and description. You can also unregister the server from LCS.

---


**Note:** LCS only supported Snapshot Metrics for servers running on Tomcat or WebLogic. On other application servers you may see a console message indicating that LCS doesn't support server metrics for your application server. You may also see a benign `NullPointerException` for the `LCS TaskSchedulerServiceImpl` and `ScheduleTasksCommand`.

---

## 480.1 Details

---

The *Details* tab shows general information about your server. There are three tabs under Details: *Software*, *Java*, and *Hardware*. Each shows information, respectively, about your Liferay DXP installation, Java installation, and hardware. This information is useful to the Liferay Support team in the event you need their assistance.



### Liferay Instance

|                       |      |
|-----------------------|------|
| Portal Build Number   | 7110 |
| Portal Edition        | GA1  |
| Patching Tool Version | 2007 |

### Liferay Connected Services

|                         |                                      |
|-------------------------|--------------------------------------|
| LCS Server Key          | 6701bf61-de2d-4913-ad4e-f56abbef13e4 |
| LCS Client Build Number | 500                                  |
| Registration Date       | Oct 18, 2018 10:39:10 PM             |
| Last Heartbeat Date     | Oct 19, 2018 6:07:22 PM              |

Figure 480.1: The Details tab shows information about your server.

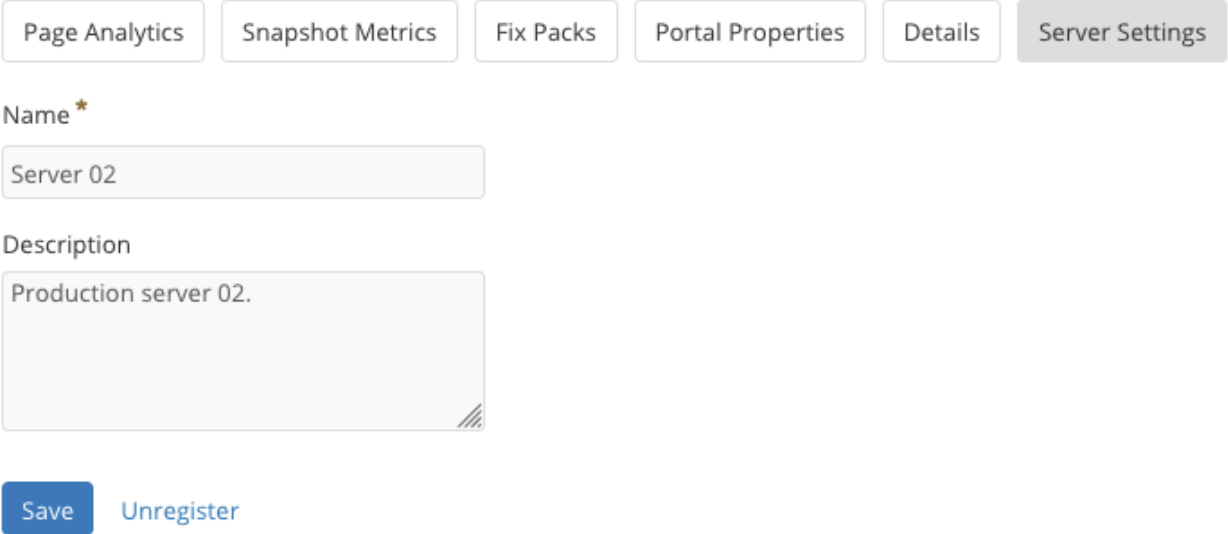
## 480.2 Server Settings

---

Finally, the *Server Settings* tab lets you view and edit your server's name, location, and description. You can also unregister your server from LCS.

## 480.3 Page Analytics

---



Page Analytics   Snapshot Metrics   Fix Packs   Portal Properties   Details   **Server Settings**

Name \*

Server 02

Description

Production server 02.

Save   Unregister

Figure 480.2: You can use the Server Settings tab to give your server a fun name.

Page Analytics appears by default when you enter server view. Page Analytics shows page views and load times for the selected site and time period. By default, all sites are selected. You can select a specific site from the *Site* selector menu. You can also select a different time period in the *Period* and *Ending At* fields. The arrows next to the *Ending At* field move the selected time period up or down, respectively, by one period. For example, if you select *One Hour* in the *Period* field, pressing the right arrow next to *Ending At* moves the selected time period up by one hour. Note that at the beginning of the current time period, it can take up to 15 minutes for data to become available. Also note that data is available for three months from the time LCS collected it.

By default, load times and page views for all pages are plotted against time in separate graphs. Below these graphs, a table displays summary statistics of data over the same time period, for each page. If you click a page in the table, the graphs plot the data for just that page. If you can't find the page you're looking for, you can search for it in the *Search* box at the top of the table. To plot data for all pages again, click the *All Pages* row at the bottom of the table.

Load times are also color coded to indicate speed. The *Load Times* graph's background is red for values above 3,000 ms, orange for values from 2,000 to 3,000 ms, and green for values less than 2,000 ms. Likewise, the table displays all load times greater than 3,000 ms in red text.

#### 480.4 Snapshot Metrics

---

To view other metrics and statistics of your server's performance, click the *Snapshot Metrics* tab near the top of the page. These metrics are broken down into three main categories: *Application*, *JVM*, and *Server*. *Application* is selected by default when you click the *Snapshot Metrics* button.

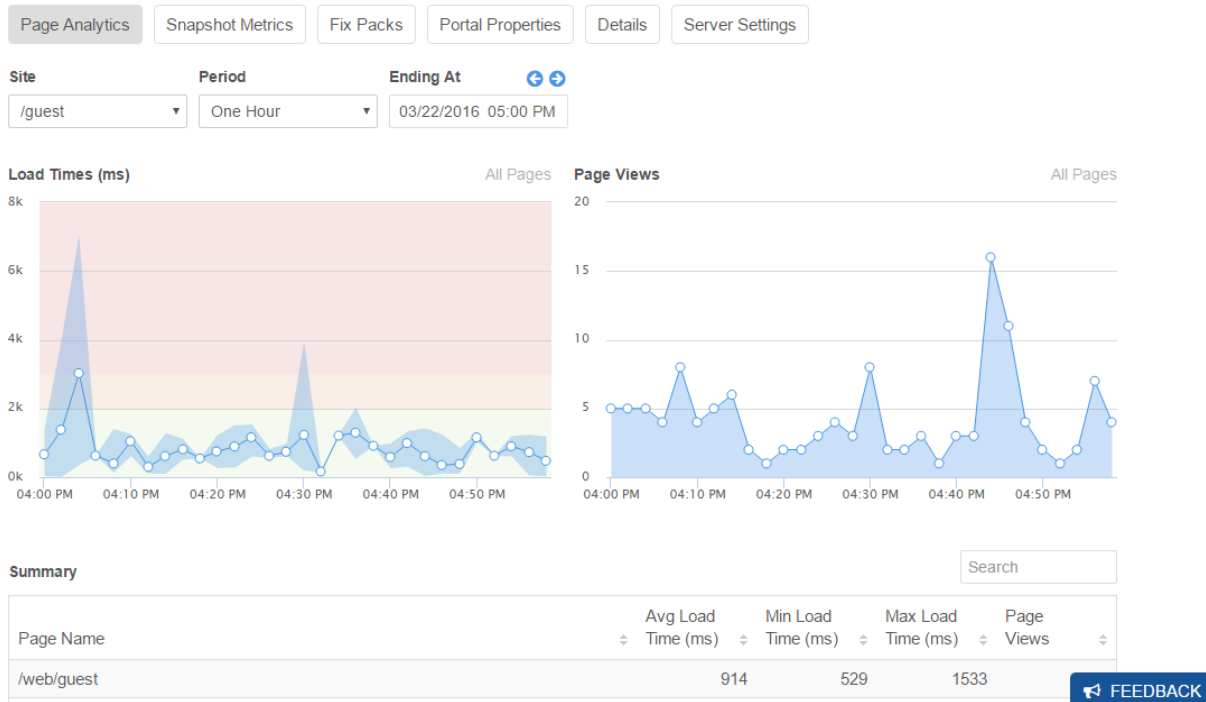


Figure 480.3: The Page Analytics interface in the LCS Server view.

The Application category also has three other categories: *Pages*, *Portlets*, and *Cache*. Pages lists the frequency that specific pages load, along with their average load times. Portlets lists the same statistics, but for specific portlets in your server. The Cache category lists Liferay Single VM metrics and Hibernate metrics. The following screenshot shows the statistics in the Portlets category.

The JVM category, as its name indicates, shows statistics about the JVM running on your server. This includes data on the garbage collector and memory. The number of runs, total time, and average time are listed for each garbage collector item. The memory metrics are presented in a bar chart that shows the usage of the PS Eden Space, Code Cache, Compressed Class Space, PS Old Gen, PS Survivor Space, and Metaspace.

Server is the third category in Snapshot Metrics. The Server category shows additional information about how your server is running. For example, horizontal bar graphs show the number of current threads running on your server, as well as the JDBC connection pools.

Note that in Snapshot Metrics, the application and garbage collector metrics are based on data collected by LCS from server registration to the present. Memory and server metrics, however, show only the current state.

## 480.5 Fix Packs




Figure 480.4: The LCS application metrics show portlet performance statistics, like frequency of use and average load time.

To view your server's fix packs, click the Fix Packs tab near the top of the page. The available and installed fix packs appear in separate tables. The available fix packs table functions exactly like the Fix Packs table in environment view for downloading and installing fix packs.

## 480.6 Portal Properties

The *Portal Properties* tab lets you view your portal's property values in a searchable table. This gives you a convenient display for your portal property settings. The properties in this table are organized into the following categories:

**Default Values:** The default values for your portal's properties.

**Custom Values:** Any custom values you've set for your portal's properties. This includes any property values you change via a portal-ext.properties file.

**Dynamic Properties:** Any property values set at runtime. For example, the Liferay Home folder's location depends on your configuration. To specify this folder when setting any properties that require it, use `${liferay.home}` instead of an absolute directory path.

You can display any combination of these categories by selecting the corresponding checkboxes from the gear icon next to the search box at the top-right of the table. For example, by checking

Application **JVM** Server

### Garbage Collector Metrics

| Name         | Runs | Total Time (ms) | Average Time (ms) |
|--------------|------|-----------------|-------------------|
| PS MarkSweep | 6    | 2242            | 373               |
| PS Scavenge  | 175  | 2378            | 13                |

### Memory Metrics




Figure 480.5: The LCS JVM metrics show performance data for memory and the garbage collector.




Figure 480.6: The LCS server metrics show current threads and JDBC connection pools.





Figure 480.7: The Fix Packs tab displays your server's available and installed fix packs.

the *Show Default Values* and *Show Custom Values* checkboxes, the table shows your portal's default and custom property values. To show only the custom values, select only *Show Custom Values*.

Page Analytics Snapshot Metrics Fix Packs **Portal Properties** Details Server Settings

Properties Differences ?  

|               |                                                                                                                                                                                                                                                                                                                  |                                                         |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------|
| Property Name | admin.obfuscated.properties                                                                                                                                                                                                                                                                                      | <input checked="" type="checkbox"/> Show Default Values |
| Default Value | jdbc.default.password                                                                                                                                                                                                                                                                                            | <input checked="" type="checkbox"/> Show Custom Values  |
| Custom Value  | auth.mac.shared.key,auth.token.shared.secret,captcha.er<br>ivate,default.admin.password,dl.file.entry.preview.generation.decrypt.passwo<br>rds.pdfbox,dl.repository.guest.password,jdbc.default.password,mail.session.m<br>ail.pop3.password,mail.session.mail.smtp.password,tunneling.servlet.shared.s<br>ecret | <input type="checkbox"/> Show Dynamic Properties        |

Figure 480.8: Click the gear icon to select the type of portal properties to show in the table.



---

## MANAGING YOUR LCS ACCOUNT

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

---

To manage your LCS account, select *My Account* from the user menu in the Dockbar. This takes you to a UI with four tabs:

**Projects:** Displays your LCS projects in a searchable table that lists the administrator's email address for each project.

**Email Notifications:** Configure LCS to send you emails when specific events occur in your LCS projects by adding rules that define what events trigger a notification. There are no rules by default. Click the *Add Rule* button to define one.

First specify the project, environment, and server for the notification. Note that you have the option of selecting all environments and servers in a project. Then check the checkbox for each event that you want to trigger an email notification. For example, if you create a notification rule with *The server unexpectedly shuts down* selected for all servers and environments in your project, then LCS sends you an email if any server in your project goes offline without a normal shut down event. Click *Save* when you're done defining the notification rule. It then appears in a table along with other existing rules. Each has Edit and Delete Action buttons.

**Notification History:** Displays your web notification history in a searchable table. You can also select the date range from which to display notifications.

**Preferences:** Manage your LCS account's preferences. You can change your account's language, time zone, and default LCS project. Your default LCS project is the one shown each time you log in to LCS.

---

### 481.1 Using Web Notifications

---

## ← Add Notification Rule

Project

Environment

Server

### Server Notifications

Send me an email when:

- Liferay Connected Services loses connection to the server.
- Liferay Connected Services restores connection to the server.
- The server is manually shut down.
- New Liferay Connected Services client is available.
- New fix pack is available.
- A connection between nodes in a cluster is broken.

### Project Notifications

Send me an email when:

- New project member is added.

Save

Cancel

Figure 481.1: You can add rules to determine the events that trigger notifications.

The image shows a user interface for account preferences. At the top, there are four tabs: 'Projects', 'Email Notifications', 'Notifications History', and 'Preferences'. The 'Preferences' tab is highlighted in grey. Below the tabs, there are three sections, each with a title and a dropdown menu:

- Language**: The dropdown menu shows 'English (United States)'.
- Time Zone**: The dropdown menu shows '(UTC) Coordinated Universal Time'.
- Default Project**: The dropdown menu shows 'Docs Test - Nick'. There is a small question mark icon to the right of the title.

At the bottom left of the form area, there is a blue button labeled 'Save'.

Figure 481.2: You can change your LCS account's general preferences.

LCS also displays web notifications under the bell icon in the Dockbar. A red badge on this icon shows your unread notification count. LCS and Liferay Support send these notifications. For example, LCS generates notifications when a server shuts down or some other event requiring your attention occurs. To mark a notification as read, click the small *x* icon next to it. To mark all notifications as read, click the *Mark All as Read* button. To mark notifications as unread again, click the *Undo* button that appears. To see your notification history, click the *Notifications History* button. You can also access your notification history by selecting *My Account* from the user menu in the Dockbar.

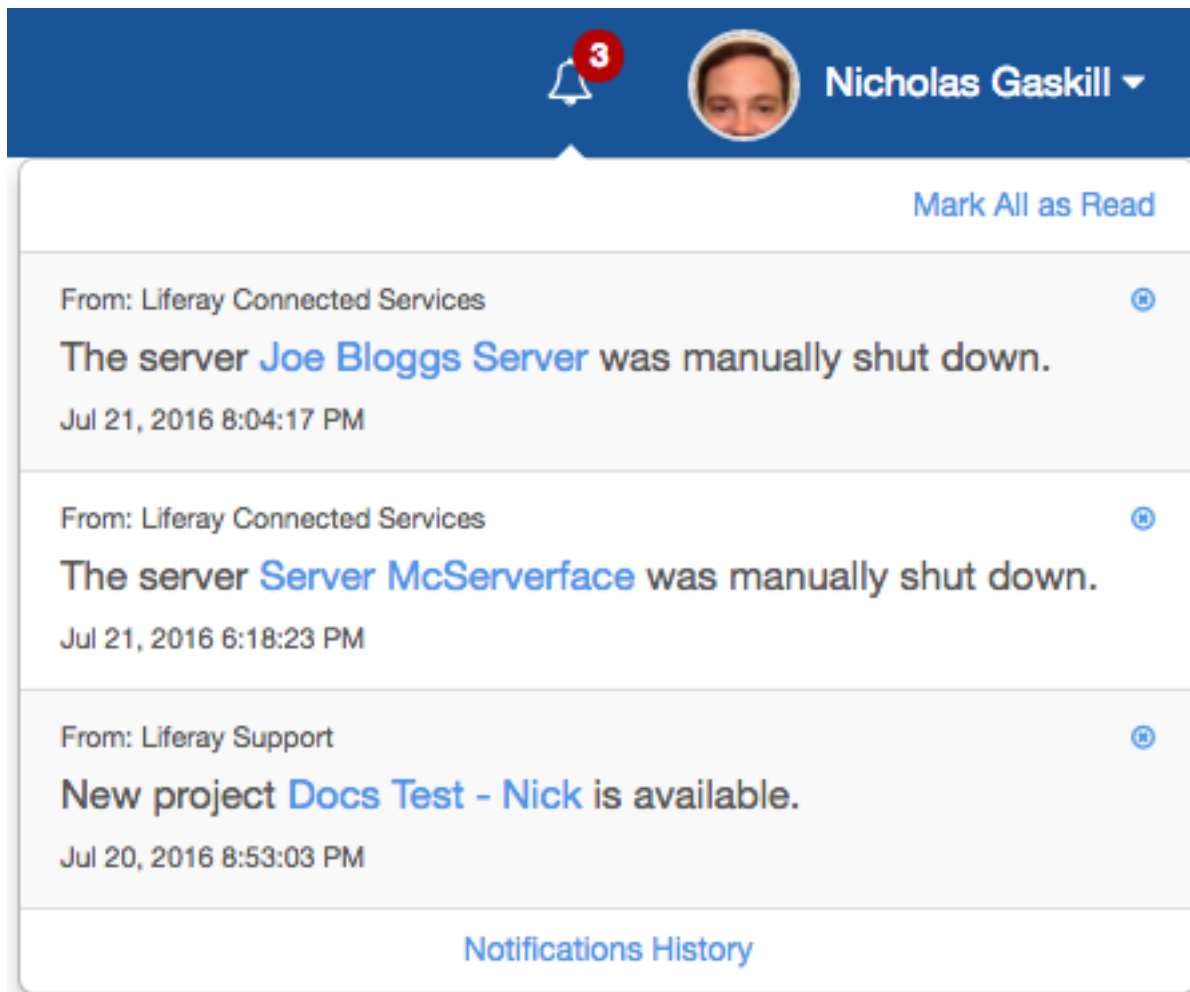


Figure 481.3: Web notifications let you know what's happening in your LCS projects.

---

## MANAGING LIFERAY DXP SUBSCRIPTIONS

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

LCS lets you use and view your Liferay DXP subscriptions. Recall that when you create an environment, you assign its subscription type and choose whether LCS activates servers that connect to that environment. If you use LCS for activation, registering a server in that environment consumes an activation key from the environment's subscription type. You can also view your project's available activation keys and see how they're being used.

Depending on your subscription agreement, LCS also lets you register servers via *elastic subscriptions*. Elastic subscriptions let you register an unlimited number servers. This is invaluable in auto-scaling environments, where servers are automatically created and destroyed in response to load. Note that to use elastic subscriptions, you must set the environment as elastic when you create it. Also note that LCS only uses elastic subscriptions for servers that exceed the number that the environment's subscription type allows. In other words, LCS uses the environment's regular subscriptions before any elastic subscriptions.

You can access these features from the *Subscriptions* tab on the upper-left of the LCS site. This tab contains two other tabs: *Details* and *Elastic Subscriptions*.

There are four tables in the *Details* tab:

1. **Subscriptions:** Shows a list of the available subscriptions in your LCS project. For each subscription, this table shows the following information:
  - Subscription Type
  - Start Date
  - Expiration Date
  - Support End Date
  - Platform
  - Product
  - Processor Cores Allowed
  - Activation Keys

Dashboard  
Users  
Subscriptions

ADD ENVIRONMENT

Production  
Development

Docs Test - Nick  
Details Elastic Subscription

When creating an environment, you need to assign an available subscription type to that environment. The servers will be able to connect to the environment as long as there is a valid subscription of the selected type.

If you need to activate a local developer workstation, you can download an activation key for your project on the [Liferay DXP activation page](#) or the [Liferay Portal activation page](#).

Subscriptions

| Subscription Type | Start Date | Expiration Date | Support End Date | Platform | Product                | Processor Cores Allowed | Activation Keys | Used Activation Keys |
|-------------------|------------|-----------------|------------------|----------|------------------------|-------------------------|-----------------|----------------------|
| Production        | 2016-07-20 | Indefinite      | 2017-07-20       |          | Gold Portal Production | 8                       | 100             | 0                    |

Subscriptions Summary

| Subscription Type | Activation Keys | Used Activation Keys | Available Activation Keys |
|-------------------|-----------------|----------------------|---------------------------|
| Production        | 100             | 0                    | 100                       |

Project Environments

| Environment | Subscription Type |
|-------------|-------------------|
| Production  | Production        |
| Development | Production        |

Project Servers

| Server    | Environment | Environment Subscription Type |
|-----------|-------------|-------------------------------|
| Server 01 | Production  | Production                    |
| Server 02 | Production  | Production                    |

Figure 482.1: LCS lets you view and manage your subscriptions.

- Used Activation Keys

Note that *Processor Cores Allowed* shows the number of processor cores that the subscription allows for each server.

2. **Subscriptions Summary:** Shows how your subscriptions are currently used in your project. For each subscription type, this table shows the number of activation keys allowed, used, and available.
3. **Project Environments:** Shows your project's environments and their assigned subscription types. Each environment must have a subscription type.
4. **Project Servers:** Shows the environment and subscription type for each server in your LCS project.

If any of the information in these tables is missing or incorrect, contact Liferay Support.

---

**Note:** If you don't use LCS for activating your servers, then you can register as many servers as you want in LCS.

---



---

**Note:** If you try to activate a server that exceeds the number of processor cores that your subscription allows per server, the activation fails and the server is locked down. A console error also indicates the server's core count. You can compare this with your subscription's processor

cores allowed in LCS's Subscriptions table. To activate the server, you can either reduce the number of cores it uses (e.g., by deploying to different server hardware, or reducing the number of virtual processors in a VM or container), or contact Liferay Sales to increase the number of processor cores that your subscription allows per server.

---

## 482.1 Decommissioning Servers

---

To decommission a server and free its activation key for reuse, select the server's environment on the left and then select the server. In the server's *Server Settings* tab, select *Unregister*. Also note that when you shut down a server normally, its activation key is immediately freed for reuse. If the server crashes or its shutdown is forced (e.g., kill), its activation key is freed for reuse within six minutes.

## 482.2 Elastic Subscriptions

---

Elastic subscriptions let you register an unlimited number of servers. This is crucial for auto-scaling environments where servers are created and destroyed automatically. You can view data on your elastic servers from the *Subscriptions* tab's *Elastic Subscriptions* tab.

---

**Note:** To register elastic servers in an environment, that environment must be set as elastic when it's created. For more information, see the documentation on creating environments.

---

The *Elastic Subscriptions* tab displays the number of elastic servers online and the uptime details for each. A graph shows the number of elastic servers online each day, while a table lists each elastic server's start time, end time, and duration. The total duration for servers is below the table. To download a report of the table's data, click *Download Report*. Also, you can use the *Environment* and *Month* selectors above the graph to select the environment and month to show data from, respectively. The data in both the graph and the table reflect your selections here.

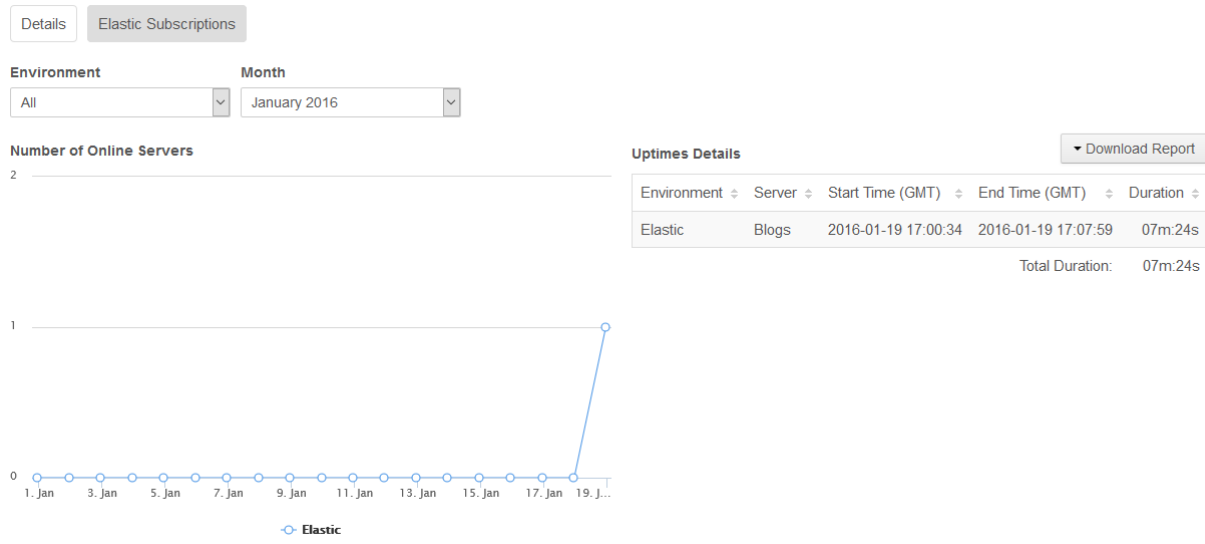


Figure 482.2: The *Elastic Subscriptions* tab shows details about your project’s elastic servers.



---

## UNDERSTANDING ENVIRONMENT TOKENS

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

To register a server in an environment, you must use that environment's token file. LCS Administrators and Environment Managers can generate and distribute this file. It contains all the information the LCS client app needs to register the server in the environment. When the server starts up, it uses the token to connect to LCS. If you use LCS for activation, the server automatically consumes an activation key from the environment's subscription upon connection. This makes it possible to activate servers automatically on startup with no interaction required.

**Note:** For instructions on using and managing your environment tokens, see the instructions on registering your server with LCS.

There are a few things to keep in mind when using environment tokens:

- Each environment can have only one token file. If you regenerate the token, servers using the old file are disconnected from LCS and can't reconnect until receiving the new file. If the server disconnects due to token regeneration and is running version 4.0.2 or later of the LCS client app, the server enters a 30-day grace period during which it functions normally. This gives the administrator time to use the new token file to reconnect to LCS. Servers running earlier versions of the LCS client app present users with an error page until the administrator reconnects with the new token.
- Use caution when distributing the token file, as anyone can use it to connect to your environment (and consume an activation key in your subscription if you're using LCS for activation).
- Minimal information (server name, location, etc...) is used to register a server with LCS. You can change this information from the server view in LCS at any time.
- Environment tokens connect using OAuth. Using an environment token overrides the OAuth authorization cycle. If LCS Administrators or Environment Managers have never registered

servers in LCS, the first time they do so an OAuth authorization entry is created in LCS. If they've previously registered servers in LCS, their existing credentials are used when they create a token file.

- If the credentials of the LCS user who generated the token become invalid, you must generate a new token and use it to reconnect to LCS. An LCS user's credentials become invalid if the user leaves the LCS project or becomes an LCS Environment Manager or LCS Environment Viewer in a different environment.

So why bother with environment tokens at all? Besides simplifying the LCS connection process, environment tokens are valuable in auto-scaling environments where algorithms create and destroy servers automatically. In this situation, having clients that activate and configure themselves is crucial.

---

**Note:** If your auto-scaling environment creates new server nodes from a server in a system image, that server can't require human interaction during setup. When creating such an image, you must change any portal property settings that prevent automatic setup. By default, Liferay DXP's setup wizard requires human interaction. You must therefore set the `setup.wizard.enabled` property to false if you want your auto-scaling environment to create new nodes from the server.

## TROUBLESHOOTING YOUR LCS CONNECTION

---

**Note:** LCS is deprecated and will be shut down on December 31, 2021. Customers who activate LCS are advised to replace it with our latest activation key type which is suitable for virtualized environments.

For further information, please see [Changes to Liferay Product Activation](#).

---

If you use LCS to activate Liferay DXP, your server must maintain its connection to LCS at all times. If this connection is interrupted, your server enters a grace period to allow for reconnection. Lengthy interruptions, however, can affect your server's uptime.

---

**Note:** You must use LCS for activation of Elastic subscriptions. Otherwise, you don't have to use LCS for activation. You can instead request an XML activation key from Liferay Support.

---

The following sections in this document provide some background information and help you troubleshoot problems with your server's LCS connection:

**LCS Grace Periods:** Describes how grace periods work in LCS. You should read this section before attempting any troubleshooting steps.

**Troubleshooting:** Presents troubleshooting steps for specific problems.

**Increasing Log Levels:** If you contact Liferay Support, you'll be asked to increase your server's log levels and then provide your log files. This section shows you how to do this.

---

**Note:** The odds of LCS being unavailable are low. LCS is deployed on a global cloud infrastructure set up for automatic failure recovery. Notifications also let the LCS team react quickly to any downtime. During LCS updates and new version releases, however, LCS is unavailable for a few minutes while changes are applied.

---

### 484.1 LCS Grace Periods

---

There are 2 grace period types in LCS:

1. **Connection Grace Period:** Occurs when your activated LCS connection is interrupted. This gives you time to re-establish the connection.
2. **Subscription Grace Period:** Occurs when your subscription is about to expire. This gives you time to renew the subscription.

---

**Note:** These grace periods only apply to servers previously connected and activated in LCS. If the subscription check or connection fails when a server attempts to connect to LCS for the first time, that server doesn't enter a grace period. It's therefore important to verify that an active subscription is available before connecting a new server to LCS. To do this, check the Subscriptions tab in LCS.

---

## 484.2 Connection Grace Period

---

If your server's LCS connection is interrupted, the server continues to run and enters a grace period that lasts for up to 30 days to allow for reconnection. During this grace period, Liferay DXP displays a warning message to administrators. Upon seeing this message, administrators should contact Liferay Support and follow the troubleshooting steps below. LCS automatically restores your server's activation upon reconnection (you shouldn't need to restart the server). If for some reason the connection can't be restored, Liferay Support will provide an alternative way to activate your server.

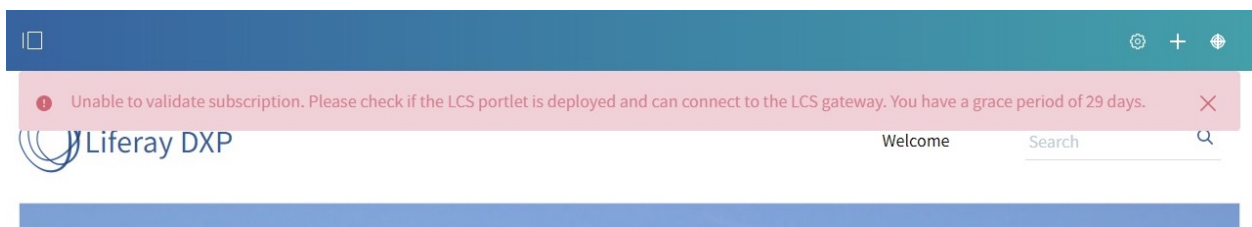


Figure 484.1: A warning message is displayed to administrators if the server can't connect to LCS to validate the subscription.

While disconnected from LCS, the LCS client app continually attempts to reconnect. If reconnection continues to fail, ensure that your server can access `lcs.liferay.com` and `lcs-gateway.liferay.com`. If the LCS client app stops attempting to reconnect, there will be no activity in the logs. In this case, you can force reconnection by redeploying the app. Follow these steps to do so:

1. In your server's Liferay Home folder (usually the parent folder of the application server's folder), remove this file:

```
osgi/marketplace/Liferay Connected Services Client.lpkg
```

2. Place Liferay Connected Services Client.lpkg in [Liferay Home]/deploy. If you connect to LCS through a proxy, and configured this inside the LCS client app, make sure the app you deploy is also configured to do so.

You should also ensure that you've enabled email notifications in LCS for server disconnection events. To do this, you must create a notification rule that sends an email whenever the server shuts down unexpectedly. The documentation on managing your LCS account explains how to do this.

### 484.3 Subscription Grace Period

---

At least 90 days before the subscription expires, Liferay will reach out to begin the renewal process. 30 days before expiration, Liferay Support sends warning messages through the Help Center, the LCS site, and the Customer Portal. After the expiration date, your servers may be placed in an additional grace period, which is communicated through the same support channels. If the renewal isn't completed during this grace period, then the subscription becomes inactive and the Liferay DXP instance enters the 30-day grace period. As soon as the renewal is processed, the instance activates and any error or warning messages disappear within 24 hours. Note that by using XML activation keys (provided by Liferay Support upon request), you can continue to use your Liferay DXP instances even after a subscription has expired.

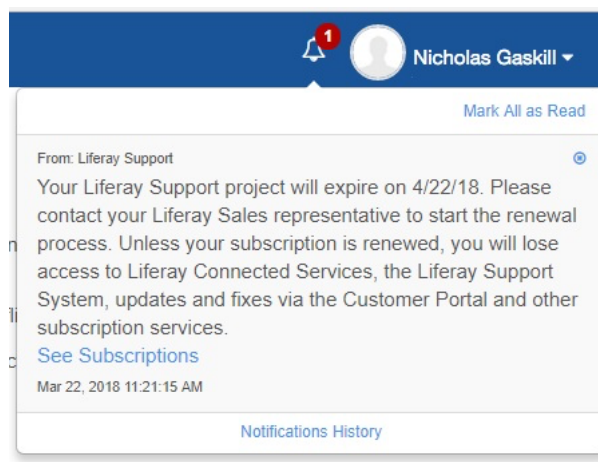


Figure 484.2: LCS sends you a notification prior to the expiration of your subscription.

#### 484.4 Troubleshooting

---

If you encounter issues with LCS, the Liferay Support team is here to help. If you need support, open a Help Center ticket. You can begin troubleshooting the following scenarios, which the Liferay Support team can also assist you with.

---

**Note:** Before troubleshooting specific issues or contacting Liferay Support, make sure that you've followed the LCS preconfiguration and registration steps correctly.

---

#### 484.5 Server Can't Reach LCS

---

If your server can't reach LCS, verify that you can access the public sites required by LCS:

- `lcs.liferay.com` should be viewable in a browser.
- `lcs-gateway.liferay.com` should respond on port 443:

```
curl -vk -I "https://lcs-gateway.liferay.com"
telnet lcs-gateway.liferay.com 443
```

#### 484.6 Subscription Issues

---

For issues related to your subscription, first review the documentation on managing your subscription. Subscription errors usually involve one of these problems:

- Your server can reach LCS, but can't locate a subscription.
- Your server can reach LCS and locate a subscription, but activating your server would exceed the subscription's number of activation keys or cores.

In either case, you must verify that a subscription is available and that you're not exceeding its number of activation keys or cores. You can find this information on the LCS site's Subscriptions page, as described in the documentation on managing subscriptions. If the environment in which you're trying to activate a server isn't assigned the subscription you want to use, then you must create a new environment and assign it the correct subscription. Once assigned, you can't change an environment's subscription. Follow the initial registration steps for instructions on creating a new environment and activating a new server.

---

**Note:** When shutting down servers, you must ensure that the LCS site receives the server shutdown commands. Otherwise, LCS may not release that server's activation key for reuse and attempts to activate additional servers may exceed the subscription's number of activation keys.

There's a higher likelihood of this happening in rolling deployments and/or when using containers. For more information, see the KB article on properly unregistering subscriptions.

---

#### 484.7 Invalid Token

---

If the token is invalid, first review the documentation on environment tokens. The following table lists causes and solutions for invalid tokens.

| Cause                                                                                                                                                                                                         | Solution                           |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------|
| The LCS user who generated the token no longer has permissions. This happens when the user leaves the LCS project or becomes an LCS Environment Manager or LCS Environment Viewer in a different environment. | Regenerate the token.              |
| The token's file name is changed after download.                                                                                                                                                              | Download the token again from LCS. |
| The token is regenerated.                                                                                                                                                                                     | Use the regenerated token.         |

---

#### 484.8 Increasing Log Levels

---

If you contact Liferay Support, you're asked to increase your server's log levels and then provide your log files. You can find these log files in [Liferay Home]/logs (Liferay Home is usually the parent folder of the application server's folder). There are 2 types of log files in this folder:

1. **Liferay log files:** The files liferay.[date].log and liferay.[date].xml are the logs for your Liferay DXP installation. Note that LOG and XML files for the same date contain the same information—the only difference is the file format.
2. **LCS log files:** The lcs-portlet-[date].log files are the LCS client app's logs. Note that if there's only a single LCS log file, it may appear without a date as lcs-portlet.log. When you increase the log levels as described in the following sections, the additional log messages are written to these LCS log files.

There are 2 ways to increase the log levels:

1. **In your Liferay DXP instance's Control Panel:** This is a temporary configuration that resets upon shutting down the server. Note that if the server isn't activated, you can't access the Control Panel. In that case, Liferay Support can provide an XML activation key.
2. **In a Log4j configuration:** This is a permanent configuration that persists through server shutdown and restart.

The following sections cover both options.

## 484.9 Control Panel

---

Follow these steps to increase the log levels via the Control Panel:

1. Navigate to *Control Panel* → *Configuration* → *Server Administration*.
2. Click the *Log Levels* tab.
3. Search for "lcs".
4. Change the log level for each matching entry to DEBUG.
5. While in the Control Panel, you should also navigate to *Configuration* → *Liferay Connected Services* and take a screenshot of what you see there. This is useful to Liferay Support.

## 484.10 Log4j

---

Follow these steps to increase the log levels via Log4j:

1. Download the latest LCS client as instructed in the LCS preconfiguration article. The app downloads as Liferay Connected Services Client.lpkg. If you don't want to download the latest client, you can use the one already installed in your server: it's in [Liferay Home]/osgi/marketplace (just make sure to shut down your server before following the rest of the steps in this section). Recall that the Liferay Home folder is usually the parent folder of the application server's folder.
2. Expand the LPKG file, then expand the lcs-portlet-[version].war file inside it.
3. Inside the WAR file, replace the contents of WEB-INF\classes\META-INF\portal-log4j.xml with the following configuration:

```
<?xml version="1.0"?>
<!DOCTYPE log4j:configuration SYSTEM "log4j.dtd">

<log4j:configuration xmlns:log4j="http://jakarta.apache.org/log4j/">
```



```

<appender class="org.apache.log4j.rolling.RollingFileAppender" name="RollingFileAppender">
 <rollingPolicy class="org.apache.log4j.rolling.TimeBasedRollingPolicy">
 <param name="ActiveFileName" value="@liferay.home@/logs/lcs-portlet.log" />
 <param name="FileNamePattern" value="@liferay.home@/logs/lcs-portlet.%d{yyyy-MM-dd}.log.zip" />
 </rollingPolicy>

 <layout class="org.apache.log4j.EnhancedPatternLayout">
 <param name="ConversionPattern" value="%d{yyyy/MM/dd HH:mm:ss} %-5p [%t][%c{1}:%L] %m%n" />
 </layout>
</appender>

<category name="com.liferay.lcs.task.scheduler">
 <priority value="ALL" />
</category>

<logger additivity="false" name="com.liferay.lcs">
 <level value="ALL" />
 <appender-ref ref="RollingFileAppender" />
</logger>
</log4j:configuration>

```

4. Save the file and repackage the WAR and LPKG (make sure not to change the names of these files).
5. Make sure your server is shut down.
6. In your installation's Liferay Home folder, delete the existing LCS client app:

```
osgi/marketplace/Liferay Connected Services Client.lpkg
```

7. Place the Liferay Connected Services Client.lpkg that you repackaged in step 4 in osgi/marketplace.
8. Start your server.

If you need assistance with the issues in this guide, or any other issues with LCS, contact Liferay Support.



## DEPLOYMENT REFERENCE

---

Here you'll find definitions, default settings, templates, and more. Here are some of the topics:

- **Liferay Home:** From this location, Liferay DXP launches applications, applies configurations, loads JAR files, and generates logs.
- **Portal properties:** Use a `portal-ext.properties` file (or another qualified properties file) to configure Liferay DXP and override features.
- **System properties:** Examine Liferay DXP default system configuration.
- **Database templates:** Use these portal properties templates to specify the Liferay DXP database.
- **Elasticsearch settings:** Examine the configuration settings for Liferay's default Elasticsearch adapter.



---

## LIFERAY HOME

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

*Liferay Home* is the location from which Liferay DXP launches applications, applies configurations, loads JAR files, and generates logs.

- *In a Liferay DXP bundle*, Liferay Home is the installation's top-level folder and it contains the application server.
- *In a manual installation*, the Liferay Home folder varies by application server. If you're doing a manual installation, please refer to the article covering that app server (e.g., *Installing Liferay DXP on [app server]*) for the Liferay Home location.

Bundles contain this folder structure regardless of application server:

- **Liferay Home**
  - **[Application Server]**: This folder is named after the application server where Liferay DXP is installed.
  - **data** (if HSQL database is selected): Stores an embedded HSQL database, Liferay DXP's file repository, and search indexes. The embedded HSQL database is configured by default, but it's intended for demonstration and trial purposes only. The Portal property `jdbc.default.url` sets the Hypersonic embedded HSQL database location.
  - **deploy**: To auto-deploy plugins, copy them to this folder. It supports application `.lpkg` files from Liferay Marketplace, plugin `.war` files, and plugin `.jar` files. The Portal property `auto.deploy.deploy.dir` sets the auto-deploy location.
  - **license**: Liferay DXP's copyright and version files are here.
  - **logs**: Log files go here. Examine them as you diagnose problems. `portal-impl.jar`'s `portal-impl/src/META-INF/portal-log4j.xml` file sets the log file location. To override the log file location, you must use an `ext-impl/src/META-INF/portal-log4j-ext.xml` file in an Ext plugin.
  - **osgi**: All the JAR files and a few configuration files for the OSGi runtime belong in this folder. The Portal property `module.framework.base.dir` sets the OSGi folder location. Here are its subfolders:

- \* configs: Component configuration files.
  - \* core: Liferay DXP's core modules.
  - \* marketplace: Marketplace applications and application suites.
  - \* modules: Modules you've deployed.
  - \* portal: Liferay DXP's non-core modules.
  - \* state: Contains OSGi internal state files for such things as OSGi bundle installation, bundle storage, and more.
  - \* target-platform: Target platform index.
  - \* test: Modules that support test integration.
  - \* war: WAR plugins you've deployed.
- patching-tool: (Liferay DXP only) This folder contains patches and a utility for installing the patches.
  - tools: For Liferay DXP upgrade and target platform indexer.
  - work: Module Jasper work files.

---

**Note:** If Liferay DXP cannot create resources in the Liferay Home folder or if it finds itself running on certain application servers, it creates a folder called `liferay` in the home folder of the operating system user that is running Liferay DXP. In this case, that `liferay` folder becomes Liferay Home. For example, if the operating system user's name is `jbloggs`, the `liferay` folder path is `/home/jbloggs/liferay` or `C:\Users\jbloggs\liferay`.

---

## PORTAL PROPERTIES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

You can set portal properties to configure and override Liferay DXP features. Your installation's `portal-impl.jar` embeds the default properties file:

`portal.properties` (Opens New Window)

Overriding a portal property requires creating an *extension* portal properties file that specifies the properties you're overriding.

---

**Note:** In a portal properties extension file, specify only the properties you're overriding.

---

Here's an example of setting Portal's data source to a MySQL database by adding override properties in a `[Liferay Home]/portal-ext.properties` file:

```
jdbc.default.driverClassName=com.mysql.cj.jdbc.Driver
jdbc.default.url=jdbc:mysql://localhost/lportal?characterEncoding=UTF-8&dontTrackOpenResources=true&holdResultsOpenOverStatementClose=true&serverTimezone=UTC
jdbc.default.username=jbloggs
jdbc.default.password=pass123
```

The `include-and-override` property specifies portal property files that override the defaults. It specifies the order the files are read—the last file read takes highest priority.

Properties file prioritization (highest to lowest):

1. `[Liferay Home]/portal-setup-wizard.properties`
2. `[user home]/portal-setup-wizard.properties`
3. `[Liferay Home]/portal-ext.properties`
4. `[user home]/portal-ext.properties`
5. `[Liferay Home]/portal-bundle.properties`
6. `[user home]/portal-bundle.properties`
7. `[Liferay Home]/portal.properties`
8. `portal-impl.jar/portal.properties`

`[Liferay Home]/portal-ext.properties` is the most commonly used file for overriding the defaults.





## SYSTEM PROPERTIES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.  
System properties configure the Liferay DXP system. Your installation's `portal-impl.jar` embeds the default properties file:  
`system.properties` (Opens New Window)



---

# DATABASE TEMPLATES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Below are templates (example portal properties) for configuring various databases as a built-in data source for Liferay DXP.

---

**Note:** The Liferay DXP Compatibility Matrix specifies supported databases and environments.

---

## 489.1 MariaDB

---

```
jdbc.default.driverClassName=org.mariadb.jdbc.Driver
jdbc.default.url=jdbc:mariadb://localhost/lportal?useUnicode=true&characterEncoding=UTF-8&useFastDateParsing=false
jdbc.default.username=
jdbc.default.password=
```

## 489.2 MySQL

---

---

**Note:** MySQL Connector/J 8.0 is highly recommended for use with MySQL Server 8.0 and 5.7.

---

```
jdbc.default.driverClassName=com.mysql.cj.jdbc.Driver
jdbc.default.url=jdbc:mysql://localhost/lportal?characterEncoding=UTF-8&dontTrackOpenResources=true&holdResultsOpenOverStatementClose=true&serverTime
jdbc.default.username=
jdbc.default.password=
```

### 489.3 PostgreSQL

---

```
jdbc.default.driverClassName=org.postgresql.Driver
jdbc.default.url=jdbc:postgresql://localhost:5432/lportal
jdbc.default.username=sa
jdbc.default.password
```

See the default portal properties for more database templates.

---

## COMPARING PATCH LEVELS

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

If you're a developer, the Patching Tool can show you what changed between different Liferay DXP patches and versions. These commands show you information about the different patch levels:

`patching-tool diff`: Prints the differences between two patch levels. At least one stored patch level must be available. This command accepts options for filtering the output:

- `source`: Shows the source differences between the two patch levels.
- `files`: Shows a list of the modified files.
- `fixed-issues`: Shows a list of LPS/LPE issues from our issue tracking system.
- `html`: Specify this along with one of the filtering options (`source`, `files`, or `fixed-issues`) and after the patch levels, to write the differences to an HTML file (`<stored-name-1>-<stored-name-2>-[type]-diff.html`) in the `diffs` folder. Additions are colored green and deletions are colored red.
- `collisions`: Shows a list of modified files which collide with deployed plugins.

For detailed usage information, run `patching-tool help diff`.

`patching-tool store`: Manages patching level information for the `diff` command. Your patches must contain source code to store the patch level and to prepare usable information for the `diff` command. Here are the store command options:

- `info`: Prints the list of patches which make up the stored patch level.
- `add`: Stores the patch level that can be found in the patches directory.
- `update`: Adds or updates patch level information.
- `rm`: Removes previously stored patch level information.

For detailed usage information, run `patching-tool help store`.



---

## PATCHING TOOL CONFIGURATION PROPERTIES

---

This document has been updated and ported to Liferay Learn and is no longer maintained here.

Here are the Patching Tool configuration properties. See [Configuring the Patching Tool](#) for more information on configuring the Patching Tool.

**patching.mode:** This can be binary (the default) or source if you're patching a source tree. Patches contain both binary and source patches. If your development team extends Liferay DXP, have them patch their source tree.

**patches.folder:** Specify where to store patches. The default location is `./patches`.

**war.path:** Specify the location of the Liferay DXP installation inside your application server. Alternatively, you can specify a `.war` file here, and you can patch a Liferay DXP `.war` for installation to your application server.

**global.lib.path:** Specify the location for storing `.jar` files on the global classpath. If you're not sure, search for `portal-kernel.jar`; it's on the global classpath. This property is only valid if your `patching.mode` is binary.

**liferay.home:** Specify the default location for the data, osgi, and tools folders.

**source.path:** Specify the location of your Liferay DXP source tree. This property is only valid if your `patching.mode` is source.

Service Pack detection is available behind a proxy server. To configure your proxy, use the following settings, making sure to replace `[PROXY_IP_ADDRESS]` with your proxy server's IP address and replace the port numbers with yours:

```
Proxy settings

HTTP Proxy

#proxy.http.host=[PROXY_IP_ADDRESS]
#proxy.http.port=80
#proxy.http.user=user
#proxy.http.password=password

HTTPS Proxy

proxy.https.host=[PROXY_IP_ADDRESS]
proxy.https.port=80
proxy.https.user=user
proxy.https.password=password

SOCKS Proxy
```

```
#proxy.socks.host=[PROXY_IP_ADDRESS]
#proxy.socks.port=1080
#proxy.socks.user=user
#proxy.socks.password=password
```



---

# TROUBLESHOOTING DEPLOYMENTS

---

When coding on any platform, you can sometimes run into issues that have no clear resolution. This can be particularly frustrating. If you have issues building, deploying, or running apps and modules, you want to resolve them fast. These frequently asked questions and answers help you troubleshoot and correct problems.

Click a question to view the answer.

Why did the entity sort order change when I migrated to a new database type?

[Your new database uses a different default query result](/docs/7-2/deploy/-/knowledge_base/d/sort-order-changed-with-a-different-database) you should be able to configure a different order.

How can I use files to configure components?

See [Using Files to Configure Module Components](/docs/7-2/deploy/-/knowledge_base/d/using-files-to-configure-product-modules).

The application server and database started, but Liferay DXP failed to connect to the database. What happened and how can I fix this?

Liferay DXP initialization can fail while attempting to connect to a database server that isn't ready. [Configuring startup to retry JDBC connections](/docs/7-2/deploy/-/knowledge_base/d/portal-failed-to-initialize-because-the-database-wasnt-ready) facilitates



---

## LIFERAY DXP FAILED TO INITIALIZE BECAUSE THE DATABASE WASN'T READY

---

If you start your database server and application server at the same time, Liferay DXP might try connecting to the data source before the database is ready. By default, Liferay DXP doesn't retry connecting to the database; it just fails. But there is a way to avoid this situation: database connection retries.

1. Create a `portal-ext.properties` file in your Liferay Home folder.
2. Set the property `retry.jdbc.on.startup.max.retries` equal to the number of times to retry connecting to the data source.
3. Set property `retry.jdbc.on.startup.delay` equal to the number of seconds to wait before retrying connection.

If at first the connection doesn't succeed, Liferay DXP uses the retry settings to try again.

### **493.1 Related Topics**

---

Connecting to JNDI Data Sources



## SORT ORDER CHANGED WITH A DIFFERENT DATABASE

---

If you've been using Liferay DXP, but are switching it to use a different database type, consult your database vendor documentation to understand your old and new database's default query result order. The default order is either case-sensitive or case-insensitive. This affects entity sort order in Liferay DXP.

Here are some examples of ascending alphabetical sort order.

Case-sensitive:

```
111
222
AAA
BBB
aaa
bbb
```

Case-insensitive:

```
111
222
AAA
aaa
BBB
bbb
```

Your new database's default query result order might differ from your current database's order. Consult your vendor's documentation to configure the order the way you want.



---

# USING FILES TO CONFIGURE MODULE COMPONENTS

---

Liferay DXP uses Felix File Install to monitor file system folders for new/updated configuration files, and the Felix OSGi implementation of Configuration Admin to let you use files to configure module service components.

To learn how to work with configuration files, first review [Understanding System Configuration Files](#).

## 495.1 Configuration File Formats

---

There are two different configuration file formats:

- `.cfg`: An older, simple format that only supports String values as properties.
- `.config`: A format that supports strings, type information, and other non-string values in its properties.

Although Liferay DXP supports both formats, use `.config` files for their flexibility and ability to use type information. Since `.cfg` files lack type information, if you want to store anything but a String, you must use properties utility classes to cast Strings to intended types (and you must carefully document properties that aren't Strings). `.config` files eliminate this need by allowing type information. The articles below explain the file formats:

- [Understanding System Configuration Files](#)
- [Configuration file \(`.config`\) syntax](#)
- [Properties file \(`.cfg`\) syntax](#)

## 495.2 Naming Configuration Files

Before you create a configuration file, follow these steps to determine whether multiple instances of the component can be created or if the component is intended to be a singleton:

1. Deploy the component's module if you haven't done so already.
2. In Liferay DXP's UI, go to *Control Panel* → *Configuration* → *System Settings*.
3. Find the component's settings by searching or browsing for the component.
4. If the component's settings page has a section called *Configuration Entries*, you can create multiple instances of the component configured however you like. Otherwise, you should treat the component as a singleton.




Figure 495.1: You can create multiple instances of components whose System Settings page has a *Configuration Entries* section.

All configuration file names must start with the component's PID (PID stands for *persistent identity*) and end with `.config` or `.cfg`.

For example, this class uses Declarative Services to define a component:

```
package com;
@Component
class Foo {}
```

The component's PID is `com.Foo`. All the component's configuration files must start with the PID `com.Foo`.

For each non-singleton component instance you want to create or update with a configuration, you must use a uniquely named configuration file that starts with the component's PID and ends with `.config` or `.cfg`. Creating configurations for multiple component instances requires that the configuration files use different *subnames*. A subname is the part of a configuration file name after the PID and before the suffix `.config` or `.cfg`. Here's the configuration file name pattern for non-singleton components:

- `[PID]-[subname1].config`
- `[PID]-[subname2].config`
- etc.

For example, you could configure two different instances of the component `com.Foo` by using configuration files with these names:

- `com.Foo-one.config`



- `com.Foo-two.config`

Each configuration file creates and/or updates an instance of the component that matches the PID. The subname is arbitrary—it doesn't have to match a specific component instance. This means you can use whatever subname you like. For example, these configuration files are just as valid as the two above:

- `com.Foo-puppies.config`
- `com.Foo-kitties.config`

Using the subname default, however, is Liferay DXP's convention for configuring a component's first instance. The file name pattern is therefore

```
[PID]-default.config
```

A singleton component's configuration file must also start with [PID] and end with `.config` or `.cfg`. Here's the common pattern used for singleton component configuration file names:

```
[PID].config
```

When you're done creating a configuration file, you can deploy it.

### 495.3 Resolving Configuration File Deployment Failures

---

The following `IOException` hints that the configuration file has a syntax issue:

```
Failed to install artifact: [path to .config or .cfg file]
java.io.IOException: Unexpected token 78; expected: 61 (line=0, pos=107)
```

To resolve this, fix the configuration file's syntax.

Great! Now you know how to configure module components using configuration files.

### 495.4 Related Articles

---

Understanding System Configuration Files