

Liferay Digital Experience Platform 7.0

Deployment Checklist



Table of Contents

Introduction	1	Document Library Storage	17
Reference Architecture	1	Default Timezones	18
Virtualized and Cloud Deployments	3	Direct Servlet Context Reload	18
Fault Tolerance	4	Enabled Locales	18
Performance	4	Encryption Algorithms	19
Scalability	4	Groups Complex SQL	19
Security	4	Message Bus	20
Liferay Digital Enterprise Configuration and Tuning Guidelines	5	Portlet CSS	22
Application Server Tuning	5	Servlet Filters	22
Database Connection Pool	5	Session Timeout	23
Deactivate Development Settings in the JSP Engine	6	Template Caching	23
Thread Pool	7	User Session Tracker	24
Java Virtual Machine Tuning	8	Liferay Enterprise Search	24
Garbage Collector	8	Sizing Your Deployment	25
Code Cache	9	Configuring Elasticsearch	27
Java Heap	9	Tuning Your Deployment	28
JVM Advanced Options	10	Monitoring Your Deployment	29
Monitoring GC and JVM	11	Securing Your Deployment	29
Liferay Digital Enterprise Tuning Parameters	13	Scaling Your Deployment	30
Caching	13	Summary	31
Cache Replication	16	Disclaimer	31
Counter Increment	16	Moving Forward	32
Document Library Previews	16	Liferay Connected Services	32
		Liferay and Dynatrace	32
		Liferay Global Services	32

Introduction

The Liferay Engineering and Global Services teams have performed intensive performance and scalability testing on Liferay Digital Experience Platform (DXP) and the accumulated knowledge has been condensed into this checklist. Although a deployment's performance profile may differ depending on several factors, this checklist will provide a list of critical parameters to monitor for tuning and some initial settings to use as a starting point. These initial settings have withstood heavy testing by the Liferay Engineering scalability team.

Liferay Digital Experience Platform contains both on-premise and cloud-based components. This deployment checklist focuses on the on-premise deployed components, specifically Liferay Digital Enterprise and Liferay Enterprise Search. Liferay DXP's cloud-based components are automatically tuned and optimized by Liferay's engineering teams.

The Liferay Global Services team also has a specialized [Go Live package](#) that can help with your pre-production tuning and configuration.

Reference Architecture

The selection of an appropriate architecture is one of the first decisions in your deployment path. To select an appropriate architecture, you must consider:

- **Information Security:** Securing sensitive hardware and information from malicious attack and intrusion
- **Performance:** Supporting the desired number of total users, concurrent transactions, etc.
- **Fault Tolerance:** Maintaining uptime during unexpected failure or scheduled maintenance
- **Flexibility and Scalability:** Designing an expandable architecture to support additional features and users without significant redesign

Although appearing somewhat complex, the reference architecture depicted in Figure 1 provides high levels of fault tolerance and flexibility.

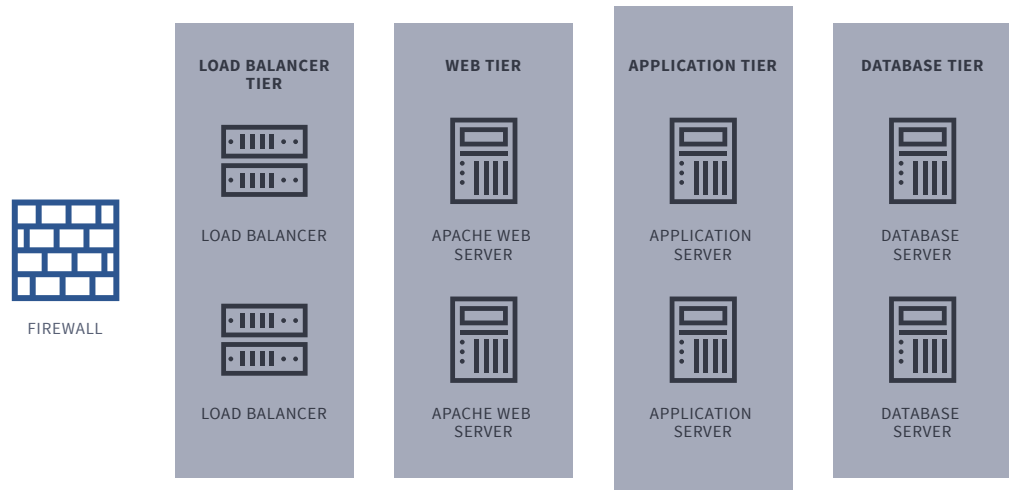


Figure 1 - Digital Enterprise Reference Architecture

The architecture contains the following tiers:

- **Firewall:** Intrusion detection and prevention
- **Load Balancer Tier:** Ensures smooth distribution of load between multiple web server resources
- **Web Server Tier:** Delivers static content elements like images, rich media, CSS files, etc. Also provides integration modules to single sign-on solutions like CA Siteminder, Oracle Identity, Ping, etc.
- **Application Tier:** Hosts Liferay supported application servers like Tomcat, JBoss, Oracle Weblogic, and IBM Websphere (please see [Liferay Digital Enterprise Compatibility Matrix](#) for additional platforms). Also hosts search engines like Solr and Elasticsearch.
- **Database Tier:** Hosts Liferay supported database servers like MySQL, Oracle, MS SQL, IBM DB2, Postgres (please see [Liferay Digital Enterprise Compatibility Matrix](#) for additional platforms)

The hardware deployed within each tier varies depending on the type of transactions. We will use Liferay Engineering's benchmarking environment as a hardware specification guide:

- **Load Balancer Tier:** Cisco Load Director or Cisco Content Services Switch (CSS) or F5 Big-IP
- **Web Tier:** Provides caching, compression, and other capabilities using Apache, Nginx, Varnish, etc.
 - 1 – Intel Core 2 Duo E6405 2.13GHz CPU, 4GB memory, 1-146GB 10k RPM SCSI
- **Application Tier:** Represents the workhorse of the architecture.
 - 2 – Intel Core 2 Quad X5677 3.46GHz CPU, 64GB memory, 2-300GB 15K RPM SATA 6Gbps - used for Liferay Portal
 - 2 – Intel Core 2 Quad X5677 3.46GHz CPU, 64GB memory, 2-300GB 15K RPM SATA 6Gbps - used for Elasticsearch
- **Database Tier**
 - 2 – Intel Core 2 Quad X5677 3.46GHz CPU, 64GB memory, 2-300GB 15K RPM SATA 6Gbps

Although the application servers have 64GB of physical memory, you may choose to deploy with less if your Java Virtual Machine (JVM) does not utilize large heap sizes. Modern operating systems will also use any available physical memory for file system caches.

Virtualized and Cloud Deployments

While the reference architecture describes a physical deployment, the same concepts may be applied to a cloud based or virtualized deployment. Many Liferay customers choose to deploy on either public clouds (e.g. Amazon EC2) or their own private clouds (e.g. VMWare VSX based private cloud). Each physical machine may be replaced by appropriate quantities of virtual machines.

In the virtualized deployments, it is critical to allocate sufficient CPU resources. For instance, for systems deployed to Amazon AWS, allocated CPUs are calculated using Amazon EC2 Compute Units. However, 1 Compute Unit does not equal to 1 physical CPU or even 1 core on a CPU. In Amazon's terms, each application server used in the reference architecture equates to roughly a "Cluster Compute Quadruple Extra Large Instance," or 33.5 EC2 Compute Units. Thus, to properly plan the virtualized/cloud deployment, customers must account for not only virtualization overhead, but also ensure allocation of sufficient CPU resources.

Fault Tolerance

The reference architecture is fault tolerant at every level. With clusters at the web, application, and database tier, you may suffer a catastrophic hardware failure of any node and continue to service users with little performance degradation.

The depicted reference architecture represents the minimum deployment units to ensure proper fault tolerance within a single data center. You may increase the number of servers within each tier according to your load patterns to achieve a multiplier effect in the number of users the deployment can support while maintaining sufficient fault tolerance.

Multi-data-center fault tolerant architectures are not provided as part of the reference architecture.

Performance

Each deployment's performance characteristics will vary depending on the type of activity and the performance of custom application elements. Liferay Engineering has created a series of scenarios to benchmark Liferay Digital Enterprise's out of box performance characteristics for content management, collaboration and social enterprise scenarios. Results from these reference architectures have indicated Liferay Digital Enterprise can support over 22,000 virtual collaboration users and over 150,000 logins per minute with an average login time of 300 milliseconds. Liferay Digital Enterprise accomplished this load within the reference architecture while utilizing no more than 40% of CPU resources in the Web Tier, 86% of CPU resources in the Application Tier, and 50% of CPU resources in the Database Tier.

Scalability

Liferay Engineering's testing has shown Liferay Digital Enterprise to scale linearly. Thus, if you know a single application server supports X virtual users and assuming sufficient database and web server resources, you may calculate the total number of application servers required.

Security

The firewall preceding the Load Balancer Tier will provide sufficient intrusion detection and prevention. However, depending on your organization's information security requirements, you may introduce additional firewall layers between each tier to further secure the infrastructure.

Liferay Digital Enterprise Configuration and Tuning Guidelines

When tuning Digital Enterprise installation, there are several factors to take into consideration, some specific to Liferay Digital Enterprise, while others are concepts that apply to all Java and Java enterprise applications. The following guidelines are meant to serve as an initial baseline from which to tune your specific deployment.

Application Server Tuning

Although the actual setting names may differ, the concepts are applicable across most application servers. We will use Tomcat as an example to demonstrate application server tuning concepts. You should also consult your application server provider's documentation for additional specific settings that they may advise.

Database Connection Pool

The database connection pool is generally sized at roughly 30-40% of the thread pool size. The connection pool provides a connection whenever Digital Enterprise needs to retrieve data from the database (e.g. user login, etc). If this size is too small, requests will queue in the server waiting for database connections. However, too large a setting will mean wasting resources with idle database connections.

As with thread pools, you should monitor these settings and adjust them based on your performance tests.

Tomcat 8.0.x

In Tomcat, the connection pools are configured in the Resource elements in \$CATALINA_HOME/conf/Catalina/localhost/ROOT.xml. Liferay Engineering used the following configuration during testing:

```
<Resource
    auth="Container"
    description="Digital Enterprise DB Connection"
    driverClass="com.mysql.jdbc.Driver"
    maxPoolSize="75"
    minPoolSize="10"
    acquireIncrement="5"
    name="jdbc/LiferayPool"
    user="XXXXXX"
    password="XXXXXXXXX"
    factory="org.apache.naming.factory.BeanFactory"
    type="com.mchange.v2.c3p0.ComboPooledDataSource"

    jdbcUrl="jdbc:mysql://someServer:3306/liferay_dxp?useUnicode=true&characterEncoding=UTF-8&useFastDateParsing=false"/>
```

In this configuration, we start with 10 threads, and increment by 5 as needed to a maximum of 75 connections in the pool.

You may choose from a variety of database connection pool providers, including DBCP, C3P0, HikariCP and Tomcat. You may also choose to configure the Liferay JDBC settings in your portal.properties.

Deactivate Development Settings in the JSP Engine

Most application servers have their JSP Engine configured for development mode. Liferay recommends deactivating many of these settings prior to entering production:

- **Development mode:** This will enable the JSP container to poll the file system for changes to JSP files.
- **Mapped File:** Generates static content with one print statement versus one statement per line of JSP text.

Tomcat 8.0.x

In the `$CATALINA_HOME/conf/web.xml`, update the JSP servlet to look like the following:

```
<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>
```

Thread Pool

Each incoming 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 will be queued waiting for the next available worker thread.

In a finely tuned system, the number of threads in the thread pool should be relatively balanced with the total number of concurrent requests. There should not be a significant amount of threads left idle waiting to service requests.

Liferay Engineering recommends setting this initially to 50 threads and then monitoring 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 2-3s range. Too few threads in the thread pool may lead to excessive request queuing while too many threads may lead to excessive context switching.

Tomcat 8.0.x

In Tomcat, the thread pools are configured in the Connector element in `$CATALINA_HOME/conf/server.xml`. Further information can be found in the [Apache Tomcat documentations](#). Liferay Engineering used the following configuration during testing:

```
<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. You should consult the appropriate Tomcat documentation for further details.

Java Virtual Machine Tuning

Tuning the JVM primarily focuses on tuning the garbage collector and the Java memory heap. These parameters look to optimize the throughput of your application. We used Oracle's 1.8 JVM for the reference architecture. You may also choose other supported JVM versions and implementations. Please consult the [Liferay Digital Enterprise Compatibility Matrix](#) for additional compatible JVMs.

Garbage Collector

Choosing the appropriate garbage collector (GC) will help improve the responsiveness of your Liferay Digital Enterprise deployment. Liferay recommends using 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. Liferay recommends first starting your tuning using parallel collectors in the new generation and concurrent mark sweep (CMS) in the old generation.

Note: the value 16 in “ParallelGCThreads=16” will vary based on the type of CPUs available. We recommend setting the value according to CPU specification. On Linux machines, you may find 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 have indicated that it does not improve performance. Your application performance may vary and you should add it to your testing and tuning plans.

Code Cache

Java uses a just-in-time (JIT) compiler to generate native code to improve performance. The default size is 48M. This may not be sufficient for larger applications. Too small a code cache will reduce performance as the JIT would not be able to optimize high frequency methods. For Digital Enterprise, we recommend starting with 64M for the initial code cache size.

```
-XX:InitialCodeCacheSize=32m -XX:ReservedCodeCacheSize=96m
```

You can examine the efficacy of the parameter changes by adding the following parameters:

```
-XX:+PrintCodeCache -XX:+PrintCodeCacheOnCompilation
```

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, Liferay recommends starting with at least the following memory settings:

```
-server -XX:NewSize=1024m -XX:MaxNewSize=1024m  
-Xms4096m -Xmx4096m -XX:MetaspaceSize=300m  
-XX:MaxMetaspaceSize=300m -XX:SurvivorRatio=12  
-XX:TargetSurvivorRatio=90 -XX:MaxTenuringThreshold=15
```

On systems that require large heap sizes (e.g., above 4GB), it may be beneficial to use large page sizes. You may activate large page sizes using the following JVM options:

```
-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 operation system to activate them. In Linux, run “cat /proc/meminfo” and look at “huge page” items.

Caution: You should avoid allocating more than 32GB to your JVM heap. Your heap size should be commensurate with the speed and quantity of available CPU resources.

JVM Advanced Options

The following advanced JVM options were also applied in the Liferay benchmark environment:

```
-XX:+UseCompressedOops -XX:+DisableExplicitGC  
-XX:-UseBiasedLocking -XX:+BindGCTaskThreadsToCPUs  
-XX:+UseFastAccessorMethods -XX:InitialCodeCacheSize=32m  
-XX:ReservedCodeCacheSize=96m
```

Please consult your JVM documentation for additional details on advanced JVM options.

Combining the above parameters together, we have:

```
-server -XX:NewSize=1024m -XX:MaxNewSize=1024m -Xms4096m  
-Xmx4096m -XX:MetaspaceSize=300m -XX:MaxMetaspaceSize=300m  
-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:+UseCompressedOops -XX:+DisableExplicitGC -XX:-UseBiasedLocking  
-XX:+BindGCTaskThreadsToCPUs -XX:+UseFastAccessorMethods  
-XX:InitialCodeCacheSize=32m -XX:ReservedCodeCacheSize=96m
```

Caution: The above JVM settings should formulate a starting point for your performance tuning. Each system’s final parameters will vary due to a variety of factors including number of current users and transaction speed.

Liferay recommends monitoring the garbage collector statistics to ensure your environment has sufficient allocations for metaspace and also for the survivor spaces. Simply using the guideline numbers above may result in dangerous runtime scenarios like out of memory failures. Improperly tuned survivor spaces also lead to wasted heap space.

Monitoring GC and JVM

Although the previously introduced parameters give you a good start to tuning your JVM, 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 including:

VISUAL VM

This tool provides a centralized console for viewing Oracle JVM performance information, including garbage collector activities.

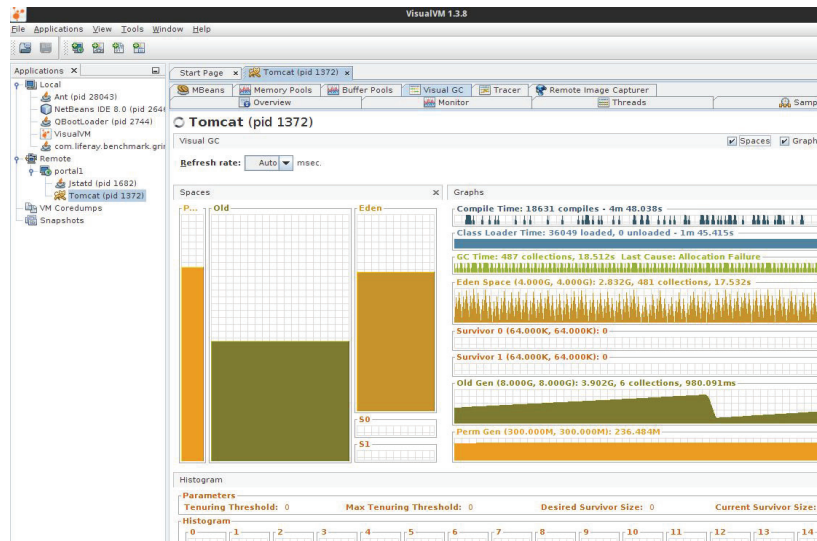


Figure 2 - Visual VM's Visual GC

JMX CONSOLE

This tool helps display various statistics like Liferay's distributed cache performance, the performance of application server threads, JDBC connection pool usage, etc.¹

Add the following to your application server's JVM arguments to enable JMX connections:

```
-Dcom.sun.management.jmxremote=true -Dcom.sun.management.jmxremote.port=5000  
-Dcom.sun.management.jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false
```

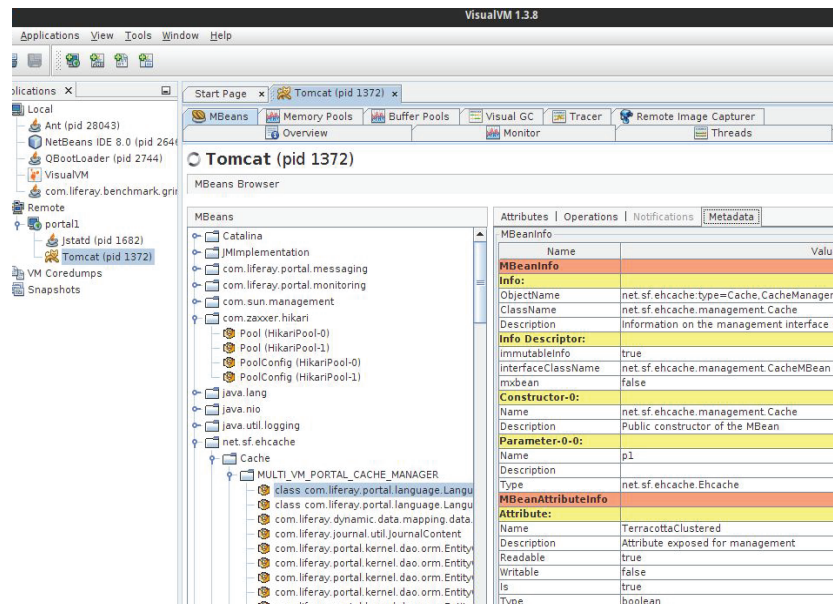


Figure 3 - Visual VM's JMX Console

Garbage Collector Verbose Logging

Add the following to your 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
```

¹ The JMX Console is the preferred tool to use when observing Tomcat performance information.

You will need these logs to properly tune the JVM.

Note: To ensure you do have sufficient debugging information should your JVM encounter out of memory scenarios, you should consider adding:

```
-XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/tmp/dumps
```

Liferay Digital Enterprise Tuning Parameters

Liferay Digital Enterprise ships with many parameters already optimized for performance. However, depending on your specific use case, you may wish to tune additional settings. Unless otherwise stated, you may tune these parameters in the `portal.properties` file¹ or using System Settings.

Caching

WARNING: Care should be taken when tuning caching. The size of the cache will have a direct impact on the amount of memory your application server has to work with. The larger the cache, the less memory available for processing requests.

If your application caching requirements exceed what the out of the box, in-process provides in scalability, Liferay recommends investigating the use of Terracotta as a clustered, out of process cache.

Liferay Digital Enterprise relies upon both content and object caching to minimize database interaction and excessive object creation. Out of the box, Liferay Digital Enterprise leverages EHCACHE for its caching needs. You may configure additional caches including Terracotta, Oracle Coherence Cache, etc. For the purposes of our discussion, we will focus on the out-of-box EHCACHE.

To monitor the caches, you will need to rely upon the JMX Console.² In the preceding figure, we see the JMX Console view for monitoring a cache used to store user information:

- **CacheHits:** Displays the number of requests that successfully retrieved from the cache rather than going to the database. This includes objects stored in memory and stored on disk.
- **CacheMisses:** Displays the number of requests that could not find its object in cache and thus had to retrieve from the database.

¹ We recommend you consult the official Liferay Digital Enterprise documentation for instructions on how to override properties stored in `portal.properties`, how to change settings using System Settings, and how to export and import settings using System Settings.

² You may also use Liferay Connected Services to help monitor cache utilization

- **InMemoryHits:** Displays the number of requests that successfully retrieved from the in-memory cache. This does not include requests that triggered retrievals from on disk storage.
- **ObjectCount:** The total number of objects in cache.
- **OnDiskHits:** The total number of requests that could not find their objects in memory, but successfully located the objects on the local file system. This is only applicable if you have enabled cache overflow to disk.³

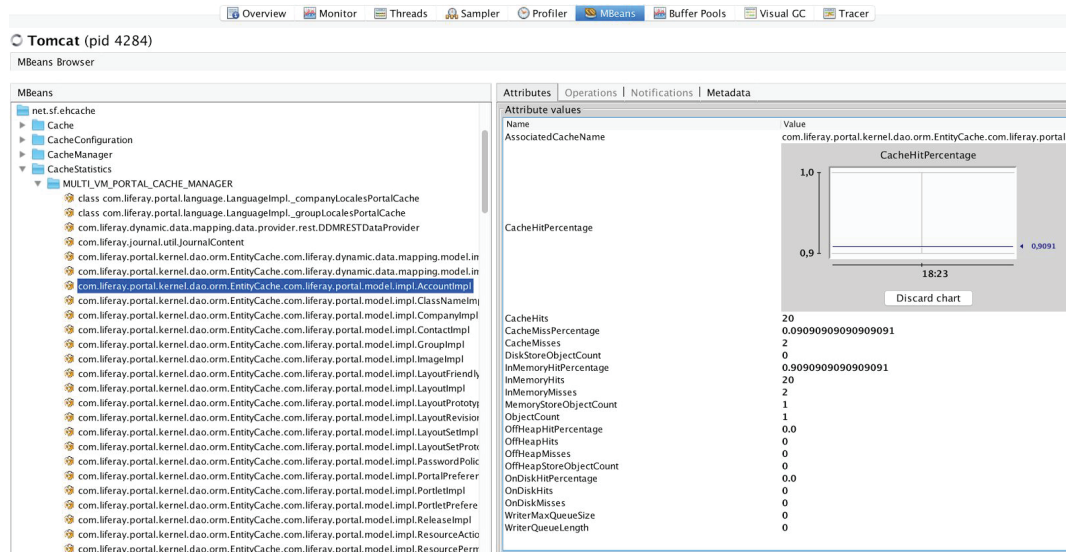


Figure 4 - Monitoring Liferay Data Caches

By default, Liferay Digital Enterprise’s cache configuration uses the memory store and allows for a maximum of 10,000 elements. Although using a disk store may increase the size of the cache, Liferay does not recommend this approach due to the increased dependency upon disk IO operations.

Let us use the user cache as our tuning example. After monitoring, you have determined that the user cache requires tuning due to the number of cache misses.

³ Cache overflow to disk means allow Ehcache to serialize (write) objects to the local file system. This is akin to using virtual memory page at the operating system level.

You may tune and override Liferay's cache settings either via a plugin or via changes to `portal-ext.properties`.

Basic Liferay Digital Enterprise data caches of interest include:

```
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portal.  
model.impl.AccountImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portal.  
model.impl.ContactImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portal.  
model.impl.GroupImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portal.  
model.impl.LayoutImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portal.  
model.impl.LayoutSetImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portal.  
model.impl.ResourceImpl
```

The following caches may be of interest, depending on the features you are leveraging in the product:

```
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portlet.  
blogs.model.impl.BlogsEntryImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portlet.  
wiki.model.impl.WikiPageImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portlet.  
messageboards.model.impl.MBCategoryImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portlet.  
messageboards.model.impl.MBThreadImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portlet.  
journal.model.impl.JournalArticleImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portlet.  
journal.model.impl.JournalStructureImpl  
com.liferay.portal.kernel.dao.orm.EntityCache#com.liferay.portlet.  
journal.model.impl.JournalTemplateImpl
```

Cache Replication

Liferay Digital Enterprise ships with an enhanced algorithm that uses more efficient thread pooling for cached object and event replication. To configure the Digital Enterprise to use this algorithm, simply activate clustering using the following `portal.properties`:

```
cluster.link.enabled=true
```

Counter Increment

Liferay Digital Enterprise uses an internal sequence generator for generating object IDs. Increasing the increment size of this counter will reduce the number of times it must communicate with the database to reserve IDs. We recommend setting this number at roughly 2000, depending on the load of your system. You may do so by adding the following to `portal-ext.properties`:

```
counter.increment=2000
```

Document Library Previews

Liferay Digital Enterprise's Document Library's preview capabilities enable users to view assets stored in the repository without downloading the file. The document preview feature relies upon PDFBox + OpenOffice Server or ImageMagick.

The PDFBox + OpenOffice Server approach uses OpenOffice Server to first convert documents to PDF and then uses PDFBox to generate images from the PDFs. This approach may be less accurate and certain types of files may not render properly. To configure this approach, please set the following properties in your `portal-ext.properties`:

```
openoffice.server.enabled=true  
openoffice.server.host=<IP_ADDRESS>  
openoffice.server.port=<PORT>  
openoffice.cache.enabled=true
```

You must also install OpenOffice in “server” or “headless” mode.

The ImageMagick approach requires installation of ImageMagick on your OS hosting Liferay Digital Enterprise. You may obtain ImageMagick from imagemagick.org/. To configure Digital Enterprise to use ImageMagick, you must configure:

```
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\\gs\\
bin;C:\\Program Files\\ImageMagick
```

You may wish to further tune ImageMagick depending on your document sizes and other parameters. You may find more details at: imagemagick.org/script/architecture.php.

In addition, preview generation can be quite expensive. Thus, by default, Liferay Digital Enterprise executes preview generation in a separate JVM:

```
dl.file.entry.preview.fork.process.enabled=true
```

This helps improve stability for the Digital Enterprise Java Virtual Machine.

Document Library Storage

Liferay Digital Enterprise's Document Library can be configured with a variety of storage engines including Amazon S3, database, and file system.

Liferay provides 2 file system based storage facilities:

FileSystemStore and AdvancedFileSystemStore. The FileSystemStore stores files using a path structure of `${liferay.home}/document_library/<companyId>/<groupId>/<fileName>/<version>`. The AdvancedFileSystemStore stores files using a path structure of `${liferay.home}/document_library/<companyId>/<groupId>/<fileName_with_extension>/<fileName_without_extension>_<version>/<version>`.

The additional directory hierarchy will help improve performance and scalability. Consequently, Liferay recommends using the AdvancedFileSystemStore for production use.

```
dl.store.impl=com.liferay.portal.store.file.
system.AdvancedFileSystemStore
```

It is important to note that Liferay does not handle file system backups, replication and file locking. Liferay Digital Enterprise relies upon your selected OS, backup and file system replication tools for those facilities.

If you do not have file system replication or backup capabilities, then the DBStore may be a better choice. DBStore stores the asset binaries into a relational database. This allows you to use the database's replication and backup facilities to also backup your document assets. However, you will experience slower download performance. This option is not recommended for Liferay Digital Enterprise deployments with large amounts of document downloads.

Liferay also provides JCR and CMIS adapters for storage. These are not recommended for production use.

Default Timezones

Liferay is capable of support many different locales. However, it is strongly recommend to use GMT for the Java Virtual Machine's default timezone. This allows Liferay to accurately display time when users select their display time zone. Please make sure the following value is set in the JVM system properties:

```
-Duser.timezone=GMT
```

Direct Servlet Context Reload

In many production systems, we do not expect to need the servlet and JSP containers to reload JSPs, because there will be no administrative operations that require JSP dynamic reloading at runtime, such as:

- Manual stop / start of a module that contains a JSP file.
- Deploy a new version of an existing OSGi module.
- Deploy a new fragment to override a JSP file of an existing OSGi module.

Consequently, in production environments where we can assume that none of the above occur, we should be able to set the following value in portal-ext.properties:

```
direct.servlet.context.reload=false
```

Enabled Locales

Liferay Digital Enterprise is capable of supporting many different languages. You may wish to reduce the number of available languages in your solution. The fully supported locales can be found in the following property:

```
locales.enabled=ca_ES,zh_CN,nl_NL,en_US,fi_FI,fr_FR,de_DE,iw_IL,hu_HU,ja_JP,pt_BR,es_ES
```

There are also languages that currently are undergoing translation and thus under beta support:

```
locales.beta=ar_SA,eu_ES,bg_BG,ca_AD,zh_TW,hr_HR,cs_CZ,da_DK,nl_BE,en_GB,en_AU,et_EE,gl_ES,el_GR,hi_IN,in_ID,it_IT,ko_KR,lo_LA,lt_LT,nb_NO,fa_IR,pl_PL,pt_PT,ro_RO,ru_RU,sr_RS,sr_RS_latin,sl_SI,sk_SK,sv_SE,tr_TR,uk_UA,vi_VN
```

To use these locales, simply add them to `locales.enabled` in your `portal-ext.properties`. For improved user experience for your content managers, you may wish to reduce the `locales.enabled` to only those that your solution needs to support. You can add more locales in the future by simply adding to the `locales.enabled` property.

Encryption Algorithms

If you intend to use Liferay's out of box login facilities and store user passwords in the database, you may wish to tune the encryption algorithm used for passwords. By default, the value is:

```
passwords.encryption.algorithm=PBKDF2WithHmacSHA1/160/128000
```

It is a Password-Based Key Derivation Function algorithm that generates a 160 bit hash after 128000 encryption rounds (2014 OWASP recommendations).

However, in deployments less powerful CPUs (e.g., vCPUs or older CPUs), you may wish to use a different algorithm or reduce the number of rounds. For instance, you can reduce the number of encryption rounds. Fewer rounds of encryption will improve performance by reducing CPU usage. You should consult your information security organization for proper guidance on password hashing guidelines.

Groups Complex SQL

Liferay Digital Enterprise uses a Group object to help track certain entities that require site-like features. For instance, organizations with sites, users with private and public profile pages, etc. The number of groups can become quite large and complex as the number of organizations and users grow. Liferay has implemented two ways to optimize querying for these groups:

1. Complex SQL which will increase utilization on the database side or
2. More in-memory processing which will increase memory consumption and CPU usage on the Digital Enterprise JVM.

Option 1 is suitable when you have a large number of groups of a particular type (e.g., users) while option 2 will help reduce database CPU usage but increase CPU and memory usage on the Digital Enterprise JVM. Thus, there is a tradeoff decision to be made.

To configure additional types for complex SQL, you may update `groups.complex.sql.class.names` in `portal-ext.properties`. The default only includes `com.liferay.portal.kernel.model.User`. You may wish to include other available models including:

1. **`com.liferay.portal.kernel.model.Organization`**: If you have a large number of organizations
2. **`com.liferay.portal.kernel.model.UserGroup`**: If you have a large number of user groups
3. **`com.liferay.portal.kernel.model.LayoutPrototype` and `com.liferay.portal.kernel.model.LayoutSetPrototype`**: If you are using staging with Layout versions and/or branching

For instance:

```
groups.complex.sql.class.names=com.liferay.portal.kernel.model.User,com.liferay.portal.kernel.model.UserGroup
```

Message Bus

WARNING: This is an advanced configuration and you should modify sparingly. The Liferay Message Bus has been tuned to handle most scenarios.

Liferay Digital Enterprise leverages asynchronous messaging, via the Liferay Message Bus, for many of its services like mail and indexing. The bus provides a loosely coupled, pluggable architecture that helps improve user experience by performing system tasks (e.g., email notifications) in the background.

Liferay recommends monitoring and tuning the message bus according to your use case. You may monitor message bus statistics via the JMX Console discussed previously.

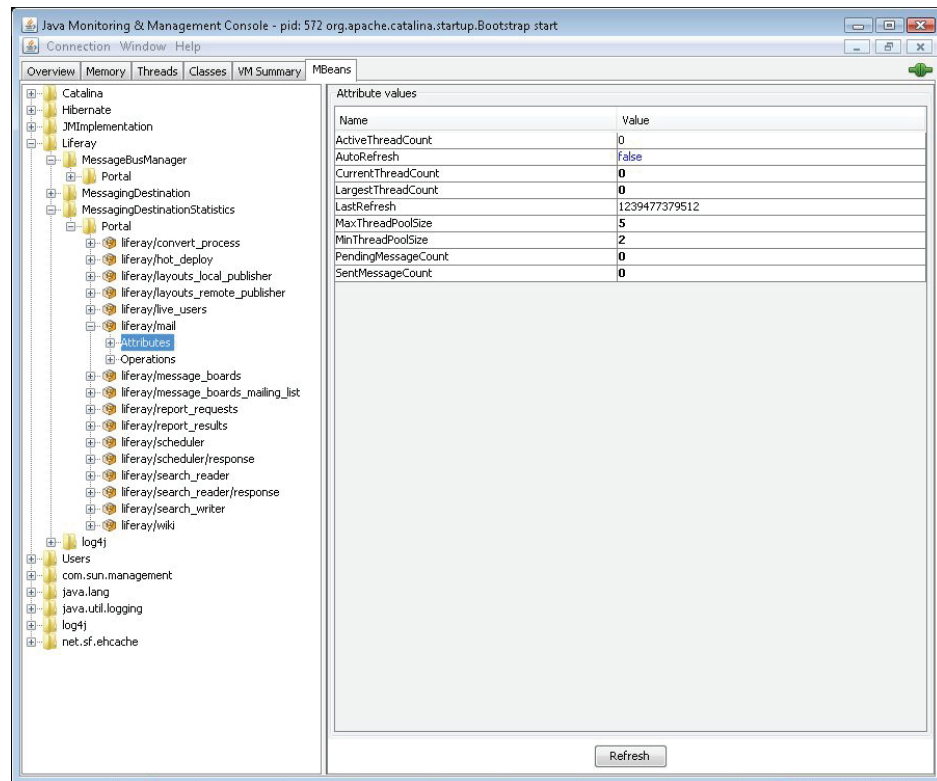


Figure 5 - Monitoring Liferay Message Bus

In above picture, we see the statistics for the messaging destination “liferay/mail”:

- **ActiveThreadCount:** Displays the current number of active worker threads for this destination.
- **CurrentThreadCount:** Displays the total number of worker threads.
- **LargestThreadCount:** Displays the maximum number of worker threads. This is the high water mark reached when the destination was the busiest.
- **MaxThreadPoolSize:** Displays the maximum number of worker threads allowed for this destination. The destination will not allocate more than this number of threads to service requests.
- **MinThreadPoolSize:** Displays the minimum number of worker threads that should be started when the destination is activated.
- **PendingMessageCount:** Displays the number of messages waiting for worker threads to deliver them.
- **SentMessageCount:** Displays the total number of messages delivered via this destination.

Portlet CSS

Liferay Digital Enterprise provides you the ability to assign custom style sheets within your portlets. You may choose to deactivate this feature, especially if you are not planning on deploying any custom portlets in Liferay Digital Enterprise. You may do so by adding the following to `portal-ext.properties`:

```
portlet.css.enabled=false
```

Servlet Filters

Liferay Digital Enterprise ships with a large collection of servlet filters to implement features like compression, SharePoint support, SSO, etc. You can improve performance by disabling those that you are not using.

You may disable the following by modifying `portal.properties`:

GZIP Filter: Used for GZIP compression of the generated pages. You should allow the Apache Web Server perform this operation. To deactivate, add the following to `portal-ext.properties`

```
com.liferay.portal.servlet.filters.gzip.GZipFilter=false
```

Strip Filter: Used to remove extraneous whitespaces in generated HTML. Web servers should be used to handle this processing. To deactivate, add the following to `portal-ext.properties`

```
com.liferay.portal.servlet.filters.strip.StripFilter=false1
```

SharePoint Filter: Used to enable Liferay Digital Enterprise to understand SharePoint protocols for integration with MS Office applications. To deactivate, add the following to `portal-ext.properties`

```
com.liferay.portal.sharepoint.SharepointFilter=false
```

Valid HTML Filter: Used to ensure HTML generated from Digital Enterprise is XHTML compliant. To deactivate, add the following to `portal-ext.properties`

```
com.liferay.portal.servlet.filters.validhtml.ValidHtmlFilter=false
```

¹ In place of the GZIP and Strip filters, we recommend using the PageSpeed Apache Module (). This module does require deploying the Apache HTTP server in your Liferay architecture.

The following filters have been moved to OSGi modules and are disabled by default:

- **SSO CAS Filter:** Used to implement single sign-on using CAS.
- **SSO NTLM Filter:** Used to implement single sign-on using NTLM.
- **SSO OpenSSO Filter:** Used to implement single sign-on using OpenSSO.

You can check the status of these filters in System Settings.

Session Timeout

With most web applications, the default session timeout is configured for 30 minutes. Although this may be friendly for the user, it does create added resource consumption for the application. Liferay Digital Enterprise provides several techniques to help you reduce the session timeout for idle users while minimally impacting usability. To reduce the lifespan of the session, you should modify the `web.xml`¹ and change the timeout:

```
<session-config>
  <session-timeout>10</session-timeout>
</session-config>
```

Template Caching

Liferay utilizes Velocity and Freemarker templating engines as part of its WCM, Theming, ADT and other features. By default, they are configured to frequently check whether template files have been modified. These settings allows for developers to work quickly, but the settings should be modified when going into production.

By default, the settings designate templates to remain in cache for 60ms. For production use, you may increase this for instance to 600000ms (10min), 3600000ms (1 hour) or -1 (indefinite).

To change the values, you may use System Settings to find settings for FreeMarker Engine and Velocity Engine. The property to modify will be labeled as “resource modification check interval.”

¹ On Tomcat, the Digital Enterprise `web.xml` is located in `$CATALINA_HOME/webapps/ROOT/WEB-INF/web.xml`

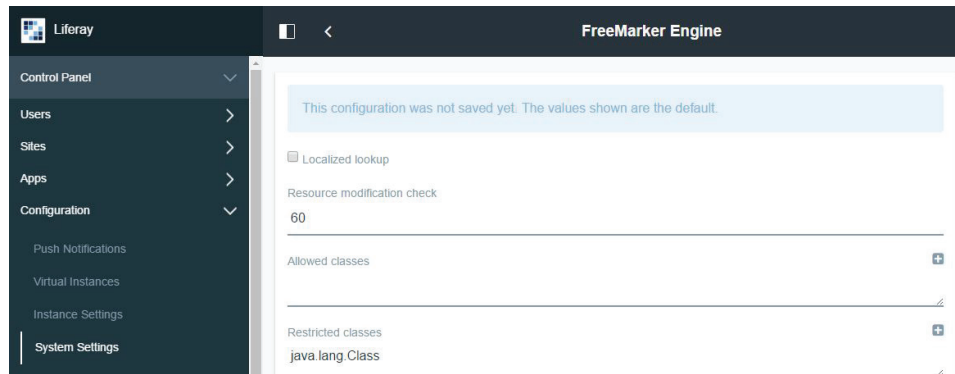


Figure 6 - Freemarker Engine System Settings

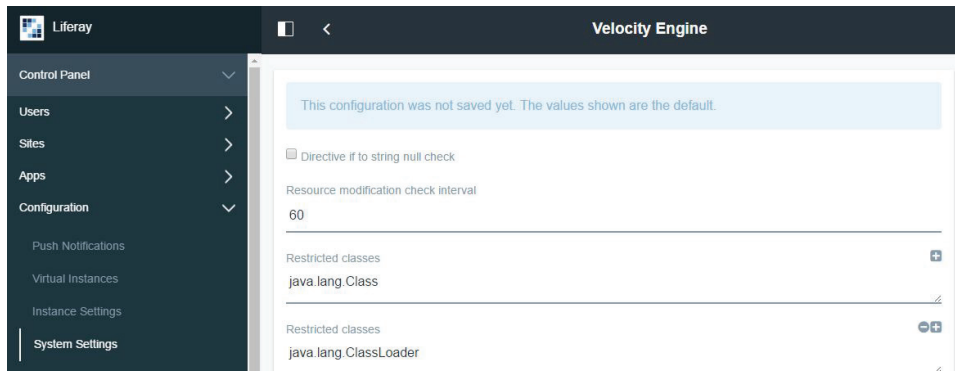


Figure 7 - Velocity Engine System Settings

User Session Tracker

Liferay Digital Enterprise enables administrators to view the activities of the users currently using the system. While useful for troubleshooting, this feature may decrease performance. In general, we recommend deactivating this feature. You may do so by adding the following to `portal-ext.properties`:

```
session.tracker.memory.enabled=false
```

Liferay Enterprise Search

In Liferay Digital Enterprise, Liferay ships with an embedded Elasticsearch search engine (i.e., the Elasticsearch engine runs in the same JVM as Liferay Digital Enterprise). Although this solution is great for having out-of-the-box search in Liferay, it will not officially be supported by Liferay for production use, only for development. For production usage, Liferay will only support use of the

Elasticsearch search engine running outside of the Digital Enterprise JVM (i.e. 1 JVM for Liferay Digital Enterprise and a separate JVM for the Elasticsearch search engine).

Search engines benefit heavily from caching and their JVM memory profiles are substantially different from a JVM focused on serving content and web views (e.g., Liferay JVM). For these reasons, the two applications should always be kept separate in production environments.

The following sections provide a synopsis of Elasticsearch configurations. Prior to deployment, we strongly recommend reading Elastic’s documentation on production deployment in “[Elasticsearch – the Definitive Guide](#).”

Sizing Your Deployment

When sizing your Elasticsearch deployment, you must carefully consider your CPU, memory, disk, and network capacity. As a general rule of thumb, you should strive to deploy Elasticsearch on medium to large machines² to scale effectively while avoiding large quantities of machines. Also, you should avoid running multiple Elasticsearch JVMs on the same operating system.

CPU

Liferay recommends at least 4 CPU cores allocated to the Elasticsearch engine. This assumes only 1 Elasticsearch JVM running on the machine.

MEMORY

Liferay recommends at least 16GB of memory. The preference is for 64GB of memory. The memory allocation depends upon the amount of index data. 64GB of memory should be sufficient for index sizes of 500GB to 1TB.

DISK

Search engines store their indices on disk and thus disk IO capacity can greatly impact search performance. Liferay recommends deploying Elasticsearch on SSD when possible. If you are unable to use SSD, Liferay recommends high performance disks (15k RPM drives). You should also consider using RAID 0 for both SSD and traditional hard disks.

² Medium machines generally have at least 2-4 vCPUs. Large machines generally have at least 4-8 vCPUs.

In general, you should avoid using NAS (network attached storage) for Elasticsearch as the network overhead can be quite large. If you are using public cloud infrastructure like Amazon Web Services, this means you should rely upon instance local storage and avoid network storage like Elastic Block Store (EBS).

You must ensure you have at least 25% more disk capacity than the total size of your indices. For instance, if you have 50G of data to be indexed, you should plan for having at least 75GB disk space available. Index sizes will vary based on the indexed content.

To ensure you have sufficient disk space, Liferay recommends indexing a representative sample of your production content. For instance, indexing 25% of production content to determine index size and then multiplying index size by 4. It is important to keep in mind that indexing a 1MB PDF document does not translate into a 1MB document in the search index.

Networking

Elasticsearch relies upon clustering and sharding to deliver fast, accurate search results. Consequently, it relies upon a fast and reliable network. Most modern data centers provide 1GbE or 10GbE between machines. You should avoid spreading Elasticsearch clusters across multiple data centers. Elasticsearch does not support multi-data center deployments, especially data centers spread across large distances (e.g., cross continents). To support multi-data center deployments, you will need to create a custom solution that distributes index requests (update, delete, add document) to each data center.

Shards and Replicas

Elasticsearch uses shards and replicas to scale.

Shards are used to divide an index into more manageable chunks. For instance, if you have 500GB index, you may wish to split this into 10 shards. For best results, each shard should not exceed 50GB. More shards generally mean faster indexing (write) performance but slower search (read) performance. In Elasticsearch, an index with multiple shards will result in a distributed search and a subsequent result merge.

Replicas are used to provide resiliency and improve search performance. A replica will help load balance search operations across the cluster.

Configuring Elasticsearch

Prior to starting Elasticsearch, you must configure a few properties in your `elasticsearch.yml`:

- `cluster.name`
 - By default Liferay Digital Enterprise expects this to be “LiferayElasticsearchCluster”.
 - If you wish to configure with a different cluster name (e.g., “elasticsearch_production”), then you must also modify Liferay Digital Enterprise’s Elasticsearch configuration in System Settings.
The cluster name configured in Liferay must match the one configured on the Elasticsearch server.
- `node.name`
 - Each node in your Elasticsearch cluster should be given a unique name.
- `discovery.zen.ping.unicast.hosts`
 - We strongly recommend using unicast for discovery versus multicast. This prevents nodes accidentally joining the cluster.
 - Not every Elasticsearch node in the cluster needs to see each other at startup. If a new node can connect to one member of the cluster at startup, then it will automatically receive the topology information and communicate with other nodes.
- `threadpool: bulk: queue_size`
 - Liferay uses bulk requests to reduce the number of network calls between Liferay and Elasticsearch. It is advisable to increase the `queue_size` for the bulk request thread pool to at least 100.

In addition to the above settings, you may choose to configure the following for debugging purposes:

```
index.indexing.slowlog.threshold.index.debug: 2s
index.indexing.slowlog.threshold.index.info: 5s
index.indexing.slowlog.threshold.index.trace: 500ms
index.indexing.slowlog.threshold.index.warn: 10s

index.search.slowlog.threshold.fetch.debug: 500ms
index.search.slowlog.threshold.fetch.info: 800ms
index.search.slowlog.threshold.fetch.trace: 200ms
index.search.slowlog.threshold.fetch.warn: 1s

index.search.slowlog.threshold.query.debug: 2s
index.search.slowlog.threshold.query.info: 5s
index.search.slowlog.threshold.query.trace: 500ms
index.search.slowlog.threshold.query.warn: 10s

monitor.jvm.gc.old.debug: 2s
monitor.jvm.gc.old.info: 5s
monitor.jvm.gc.old.warn: 10s

monitor.jvm.gc.young.debug: 400ms
monitor.jvm.gc.young.info: 700ms
monitor.jvm.gc.young.warn: 1000ms
```

Tuning Your Deployment

JVM

In general, 45% of the available system memory should be allocated to Elasticsearch, up to a maximum of 31GB. In general, no other JVM settings should be adjusted within Elasticsearch. You should configure heap sizing by setting the environment variable: `ES_HEAP_SIZE`.

The JVM vendor and version used for the Elasticsearch server must be identical to the version used for Liferay Digital Enterprise.

FILE SYSTEM

You should configure your OS for at least 64,000 file descriptors. The default Linux value is 1024.

Elasticsearch also uses NioFS and MMapFS. Consequently, you must ensure there is sufficient virtual memory available for memory-mapped files.

Please consult your system administrator on how to configure these values.

Monitoring Your Deployment

Elasticsearch provides a monitoring tool called Marvel. Marvel is comprised of two components:

- A server side component residing in your Elasticsearch nodes that feeds performance data to another Elasticsearch node
- A series of Kibana dashboards to help you visualize Elasticsearch performance information

These dashboards will help you monitor the health and performance of your Elasticsearch cluster in order to better tune and manage your search cluster.

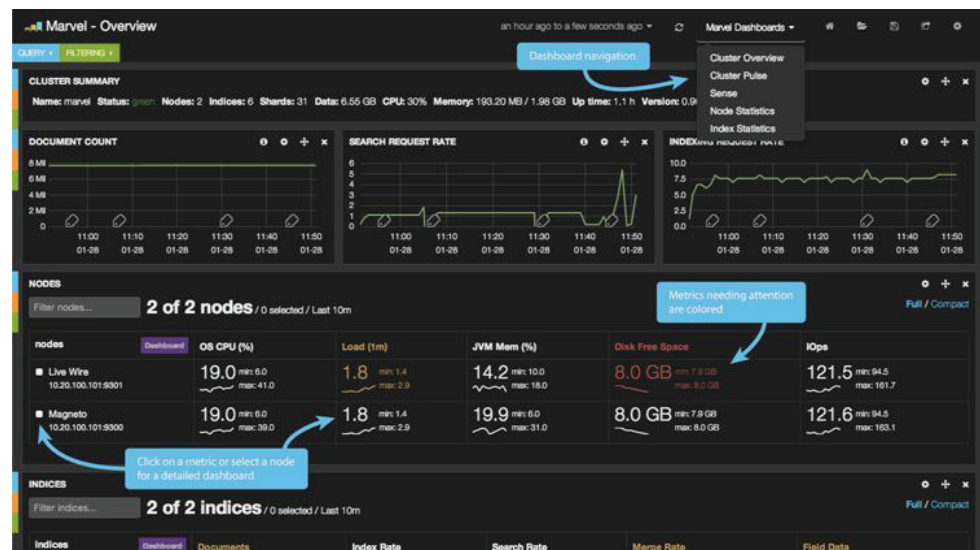


Figure 8 - Sample Marvel Dashboard

To use Marvel with Liferay's Elasticsearch cluster, you must have Liferay's Enterprise Search subscription. For more information, please contact your account executive.

Securing Your Deployment

In order to perform search operations, a search engine must index and store a large quantity of data in its indices. By default, access to this data is not secured. Unlike databases that require login for access, search engines do not require logins to perform searches and browse its stored information.

Elasticsearch's security add-on, Shield, adds over-the-wire encryption and authenticated search engine access. Over-the-wire encryption ensures all communications between Elasticsearch nodes are encrypted and the

communication between Liferay and the Elasticsearch servers are protected.³ Authenticated search engine access ensures all connections with Elasticsearch require authentication, much like a database.

These security capabilities are only available with a Liferay Enterprise Search subscription. For more information, please contact your account executive.

Scaling Your Deployment

For deployments with less than 1000 users, you may choose to deploy the Elasticsearch JVM and the Digital Enterprise JVMs on the same physical or virtual machines. With this configuration, the Elasticsearch and Liferay Digital Enterprise JVMs will share the same CPU and memory resources. Consequently, you should ensure there are at least 8 vCPU/cores and 16GB of memory allocated. You must take care to properly monitor resources to avoid over-allocating between the Liferay Digital Enterprise and Elasticsearch JVMs.

For most deployments, Liferay recommends deploying 2 Elasticsearch server with at least 4 vCPU/cores. This deployment is sufficient if you are not exceeding 50GB index size. The Elasticsearch JVMs should be configured to hold with one (1) index shard and one (1) index replica.

If you are planning a more search intensive site (e.g., 50000 users with 250000 documents), Liferay recommends moving to a three (3) server cluster. In this 3-server configuration, you should look to have 2 Elasticsearch JVMs on each server: a “master-eligible” JVM and a “data” JVM. A “master-eligible” JVM handles clustering operations, including determine which JVM holds which shard. A “data” JVM hold the shards containing the indexed documents.⁴ A dedicated “master-eligible” Elasticsearch JVM would only require 1GB of heap. A dedicated “data” JVM should be sized according to previously provided guidelines.

As the amount of indexed data and number of search requests grow, you may add additional data nodes while using the following guidelines:

- Each shard should not exceed 50GB
- Each data JVM or data + master JVM should not exceed 31GB heap.

Please keep in mind that you should only use these guidelines for initial planning. You must perform tests to ensure you have the optimal configurations.

³ Liferay Digital Enterprise communicates with the Elasticsearch cluster over Elasticsearch's binary API and not over its HTTPS/JSON API. The binary protocol provides better performance over the HTTP/JSON API.

⁴ For more detailed description, see [elastic.co/guide/en/elasticsearch/reference/current/modules-node.html](https://www.elastic.co/guide/en/elasticsearch/reference/current/modules-node.html).

Summary

In the preceding sections, we outlined the steps to design a fully fault-tolerant Liferay Digital Enterprise and Liferay Enterprise Search deployment. The architecture described builds a solid foundation for future growth. In addition, the Liferay Engineering team has shared several key factors to successfully tune a Liferay Digital Enterprise and Liferay Enterprise Search deployment. Although these parameters have wide application and have withstood many different load testing scenarios, Liferay Engineering recommends continuously monitoring your DXP deployment to ensure long-term performance.

Disclaimer

Liferay can only give you an initial tuning recommendation based on benchmarks that have been performed on the out of the box product. It is up to you as system architects and business analysts to come up with the utilization scenarios that your system will need to service. It is your responsibility to run the appropriate load tests on your system before production deployment, so that you can identify significant bottlenecks due to custom applications/portlets, and other unforeseen system and network issues, and implement appropriate configurations.

Moving Forward

Liferay Connected Services

Liferay Connected Services (LCS) provides a collection of basic performance and system monitoring tools that helps you manage and monitor your Liferay Digital Enterprise solution. For more information, visit liferay.com/liferay-connected-services.

Liferay and Dynatrace

Liferay has a partnership with Dynatrace, an industry leading application performance management (APM) solution. With Dynatrace and the Fastpack for Liferay, customers can gain deeper performance insight into their Liferay Digital Enterprise deployment. Learn more at liferay.com/dynatrace-apm.

Liferay Global Services

Liferay Global Services has specialists that focus on Liferay Digital Enterprise Platform Go Live and performance tuning consultation. Learn more at liferay.com/consulting.



Liferay makes software that helps companies create digital experiences on web, mobile and connected devices. Our platform is open source, which makes it more reliable, innovative and secure. We try to leave a positive mark on the world through business and technology. Hundreds of organizations in financial services, healthcare, government, insurance, retail, manufacturing and multiple other industries use Liferay. Visit us at [liferay.com](https://www.liferay.com).

© 2020 Liferay, Inc. All rights reserved.