



JBoss Enterprise Application Platform Continuous Delivery 18

Getting Started with JBoss EAP for OpenShift Online

Guide to developing with Red Hat JBoss Enterprise Application Platform for
OpenShift Online

JBoss Enterprise Application Platform Continuous Delivery 18 Getting Started with JBoss EAP for OpenShift Online

Guide to developing with Red Hat JBoss Enterprise Application Platform for OpenShift Online

Legal Notice

Copyright © 2019 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

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

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Guide to using Red Hat JBoss Enterprise Application Platform for OpenShift Online

Table of Contents

PREFACE	4
CHAPTER 1. INTRODUCTION	5
1.1. WHAT IS RED HAT JBOSS ENTERPRISE APPLICATION PLATFORM (JBOSS EAP)?	5
1.2. HOW DOES JBOSS EAP WORK ON OPENSIFT?	5
1.3. COMPARISON: JBOSS EAP AND JBOSS EAP CD FOR OPENSIFT	5
1.4. VERSION COMPATIBILITY AND SUPPORT	6
1.4.1. OpenShift 4.1 Support	6
CHAPTER 2. BUILD AND RUN A JAVA APPLICATION ON THE JBOSS EAP CD FOR OPENSIFT IMAGE ..	7
2.1. PREREQUISITES	7
2.2. PREPARE OPENSIFT FOR APPLICATION DEPLOYMENT	7
2.3. IMPORT THE LATEST JBOSS EAP CD FOR OPENSIFT IMAGE STREAMS AND TEMPLATES	8
2.4. DEPLOY A JBOSS EAP SOURCE-TO-IMAGE (S2I) APPLICATION TO OPENSIFT	9
2.5. POST DEPLOYMENT TASKS	10
CHAPTER 3. CONFIGURING THE JBOSS EAP CD FOR OPENSIFT IMAGE FOR YOUR JAVA APPLICATION .	12
3.1. HOW THE JBOSS EAP CD FOR OPENSIFT S2I PROCESS WORKS	12
3.2. CONFIGURING JBOSS EAP CD FOR OPENSIFT USING ENVIRONMENT VARIABLES	13
3.2.1. JVM Memory Configuration	13
3.2.1.1. JVM Default Memory Settings	13
3.2.1.2. JVM Garbage Collection Settings	14
3.2.1.3. Resource Limits in Default Settings	14
3.2.1.4. JVM Environment Variables	14
3.3. BUILD EXTENSIONS AND PROJECT ARTIFACTS	31
3.3.1. S2I Artifacts	32
3.3.1.1. Modules, Drivers, and Generic Deployments	32
3.3.2. Runtime Artifacts	35
3.3.2.1. Datasources	35
3.3.2.2. Resource Adapters	36
3.4. DEPLOYMENT CONSIDERATIONS FOR THE JBOSS EAP CD FOR OPENSIFT IMAGE	38
3.4.1. Scaling Up and Persistent Storage Partitioning	38
3.4.2. Scaling Down and Transaction Recovery	38
CHAPTER 4. MIGRATING TO JBOSS EAP CD FOR OPENSIFT JDK 11 IMAGE	39
4.1. PREPARE OPENSIFT FOR APPLICATION DEPLOYMENT USING JDK 11 IMAGE	39
4.2. IMPORT JDK 11 IMAGE	39
4.3. DEPLOY A JBOSS EAP S2I APPLICATION TO OPENSIFT USING JDK 11 IMAGE	39
4.4. CONFIGURE JBOSS EAP CD FOR OPENSIFT USING ENVIRONMENT VARIABLES FOR JDK 11 IMAGE	40
CHAPTER 5. TROUBLESHOOTING	41
5.1. TROUBLESHOOTING POD RESTARTS	41
5.2. TROUBLESHOOTING USING THE JBOSS EAP MANAGEMENT CLI	41
CHAPTER 6. REFERENCE INFORMATION	43
6.1. PERSISTENT TEMPLATES	43
6.2. INFORMATION ENVIRONMENT VARIABLES	43
6.3. CONFIGURATION ENVIRONMENT VARIABLES	44
6.4. APPLICATION TEMPLATES	49
6.5. EXPOSED PORTS	49
6.6. DATASOURCES	50

6.6.1. JNDI Mappings for Datasources	50
6.6.1.1. Database Drivers	50
6.6.1.2. Datasource Configuration Environment Variables	51
6.6.1.3. Examples	53
6.6.1.3.1. Single Mapping	53
6.6.1.3.2. Multiple Mappings	53
6.7. CLUSTERING	54
6.7.1. Configuring a JGroups Discovery Mechanism	54
6.7.1.1. Configuring KUBE_PING	54
6.7.1.2. Configuring DNS_PING	55
6.7.2. Configuring JGroups to Encrypt Cluster Traffic	56
6.7.2.1. Configuring SYM_ENCRYPT	56
6.7.2.2. Configuring ASYM_ENCRYPT	57
6.8. HEALTH CHECKS	58
6.9. MESSAGING	58
6.9.1. Configuring External Red Hat AMQ Brokers	58
Example OpenShift Application Definition	58
6.10. SECURITY DOMAINS	59
6.11. HTTPS ENVIRONMENT VARIABLES	60
6.12. ADMINISTRATION ENVIRONMENT VARIABLES	60
6.13. S2I	60
6.13.1. Custom Configuration	61
6.13.1.1. Custom Modules	61
6.13.2. Deployment Artifacts	61
6.13.3. Artifact Repository Mirrors	61
6.13.4. Scripts	62
6.13.5. Custom Scripts	62
6.13.5.1. Mounting a configmap to execute custom scripts	62
6.13.5.2. Using install.sh to execute custom scripts	63
6.13.6. Environment Variables	63
6.14. SSO	65
6.15. TRANSACTION RECOVERY	66
6.15.1. Unsupported Transaction Recovery Scenarios	66
6.15.2. Manual Transaction Recovery Process	66
6.15.2.1. Caveats	66
6.15.2.2. Prerequisite	67
6.15.2.3. Procedure	68
6.15.2.3.1. Resolving In-doubt Branches	68
6.15.2.3.2. Extract the Global Transaction ID and Node Identifier from Each XID	69
6.15.2.3.3. Obtain the List of Node Identifiers of All Running JBoss EAP Instances in Any Cluster that Can Contact the Resource Managers	71
6.15.2.3.4. Find the Transaction Logs	71
6.15.2.3.5. Cleaning Up the Transaction Logs for Reconciled In-doubt Branches	72
6.16. INCLUDED JBOSS MODULES	73

PREFACE

This document is intended for use with the JBoss Enterprise Application Platform continuous delivery release 18, which is a Technology Preview release available in the cloud only.

Some features described in this document might not work or might not be available on Red Hat OpenShift Online and Red Hat OpenShift Container Platform. For specific details about the feature differences in the JBoss EAP CD release, see the [Release Limitations](#) section in the *JBoss EAP Continuous Delivery 18 Release Notes*.



IMPORTANT

This continuous delivery release for JBoss EAP is provided as Technology Preview only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs), might not be functionally complete, and Red Hat does not recommend to use them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

See [Technology Preview Features Support Scope](#) on the Red Hat Customer Portal for information about the support scope for Technology Preview features.

CHAPTER 1. INTRODUCTION

1.1. WHAT IS RED HAT JBOSS ENTERPRISE APPLICATION PLATFORM (JBOSS EAP)?

Red Hat JBoss Enterprise Application Platform 7 (JBoss EAP) is a middleware platform built on open standards and compliant with the Java Enterprise Edition 7 specification. It provides preconfigured options for features such as high-availability clustering, messaging, and distributed caching. It includes a modular structure that allows you to enable services only when required, which results in improved startup speed.

The web-based management console and management command line interface (CLI) make editing XML configuration files unnecessary and add the ability to script and automate tasks. In addition, JBoss EAP includes APIs and development frameworks that allow you to quickly develop, deploy, and run secure and scalable Java EE applications. JBoss EAP 7 is a certified implementation of the Java EE 8 full and web profile specifications.

1.2. HOW DOES JBOSS EAP WORK ON OPENSHIFT?

Red Hat offers a containerized image for JBoss EAP that is designed for use with OpenShift. Using this image, developers can quickly and easily build, scale, and test applications that are deployed across hybrid environments.

1.3. COMPARISON: JBOSS EAP AND JBOSS EAP CD FOR OPENSHIFT

There are some notable differences when comparing the JBoss EAP product with the JBoss EAP CD for OpenShift image. The following table describes these differences and notes which features are included or supported in the current version of JBoss EAP CD for OpenShift.

Table 1.1. Differences between JBoss EAP and JBoss EAP CD for OpenShift

JBoss EAP Feature	Status in JBoss EAP CD for OpenShift	Description
JBoss EAP management console	Not included	The JBoss EAP management console is not included in this release of JBoss EAP CD for OpenShift.
JBoss EAP management CLI	Not recommended	The JBoss EAP management CLI is not recommended for use with JBoss EAP running in a containerized environment. Any configuration changes made using the management CLI in a running container will be lost when the container restarts. The management CLI is accessible from within a pod for troubleshooting purposes.
Managed domain	Not supported	Although a JBoss EAP managed domain is not supported, creation and distribution of applications are managed in the containers on OpenShift.

JBoss EAP Feature	Status in JBoss EAP CD for OpenShift	Description
Default root page	Disabled	The default root page is disabled, but you can deploy your own application to the root context as ROOT.war .
Remote messaging	Supported	Red Hat AMQ for inter-pod and remote messaging is supported. JBoss EAP CD releases only support client messaging, and Red Hat AMQ provides the messaging broker.
Transaction recovery	Partially supported	There are some unsupported transaction recovery scenarios and caveats when undertaking transaction recovery with the JBoss EAP CD for OpenShift image.

1.4. VERSION COMPATIBILITY AND SUPPORT

This guide covers the following JBoss EAP CD for OpenShift image: **jboss-eap-7-tech-preview/eap-cd-openshift-rhel8**.

You can see information on the latest tag for this image in the [Red Hat Container Catalog](#).

The Hawkular agent is not active in JDK 11, and will be ignored if configured.

JBoss EAP CD for OpenShift is updated frequently. Therefore, it is important to understand which versions of the images are compatible with which versions of OpenShift. See [OpenShift and Atomic Platform Tested Integrations](#) on the Red Hat Customer Portal for more information on version compatibility and support.

1.4.1. OpenShift 4.1 Support

Changes in OpenShift 4.1 affect access to Jolokia, and the Open Java Console is no longer available in the OpenShift 4.1 web console.

In previous releases of OpenShift, certain kube-apiserver proxied requests were authenticated and passed through to the cluster. This behavior is now considered insecure, and so, accessing Jolokia in this manner is no longer supported.

Due to changes in codebase for the OpenShift console, the link to the Open Java Console is no longer available.

CHAPTER 2. BUILD AND RUN A JAVA APPLICATION ON THE JBOSS EAP CD FOR OPENSIFT IMAGE

The following workflow demonstrates using the Source-to-Image (S2I) process to build and run a Java application on the JBoss EAP CD for OpenShift image.

As an example, the **kitchensink** quickstart is used in this procedure. It demonstrates a Java EE web-enabled database application using JSF, CDI, EJB, JPA, and Bean Validation. See the **kitchensink** quickstart that ships with JBoss EAP 7 for more information.

2.1. PREREQUISITES

This workflow assumes that you already have an [active OpenShift Online subscription](#) and that you have [installed the OpenShift CLI](#).

2.2. PREPARE OPENSIFT FOR APPLICATION DEPLOYMENT

1. Log in to your OpenShift instance using the **oc login** command.
2. Create a new project in OpenShift.
A project allows a group of users to organize and manage content separately from other groups. You can create a project in OpenShift using the following command.

```
$ oc new-project PROJECT_NAME
```

For example, for the **kitchensink** quickstart, create a new project named **eap-demo** using the following command.

```
$ oc new-project eap-demo
```

3. **Optional:** Create a keystore and a secret.



NOTE

Creating a keystore and a secret is required if you are using any HTTPS-enabled features in your OpenShift project. For example, if you are using the **eap-cd-https-s2i** template, you must create a keystore and secret.

This workflow demonstration for the **kitchensink** quickstart does not use an HTTPS template, so a keystore and secret are not required.

- a. Create a keystore.

**WARNING**

The following commands generate a self-signed certificate, but for production environments Red Hat recommends that you use your own SSL certificate purchased from a verified Certificate Authority (CA) for SSL-encrypted connections (HTTPS).

You can use the Java **keytool** command to generate a keystore using the following command.

```
$ keytool -genkey -keyalg RSA -alias ALIAS_NAME -keystore
KEYSTORE_FILENAME.jks -validity 360 -keysize 2048
```

For example, for the **kitchensink** quickstart, use the following command to generate a keystore.

```
$ keytool -genkey -keyalg RSA -alias eapdemo-selfsigned -keystore keystore.jks -validity
360 -keysize 2048
```

- b. Create a secret from the keystore.
Create a secret from the previously created keystore using the following command.

```
$ oc secrets new SECRET_NAME KEYSTORE_FILENAME.jks
```


For example, for the **kitchensink** quickstart, use the following command to create a secret.

```
$ oc secrets new eap7-app-secret keystore.jks
```

2.3. IMPORT THE LATEST JBOSS EAP CD FOR OPENSIFT IMAGE STREAMS AND TEMPLATES

Use the following command to import the latest JBoss EAP CD for OpenShift image streams and templates into your OpenShift project's namespace.

```
for resource in \
  eap-cd-amq-persistent-s2i.json \
  eap-cd-amq-s2i.json \
  eap-cd-basic-s2i.json \
  eap-cd-https-s2i.json \
  eap-cd-image-stream.json \
  eap-cd-sso-s2i.json \
  eap-cd-starter-s2i.json \
do
  oc replace --force -f \
  https://raw.githubusercontent.com/jboss-container-images/jboss-eap-7-openshift-image/eap-
  cd/templates/${resource}
done
```



NOTE

The JBoss EAP image streams and templates imported using the above command are only available within that OpenShift project.

If you want to import the image streams and templates into a different project, add the **-n PROJECT_NAME** to the **oc replace** line of the command. For example:

```
...
oc replace -n PROJECT_NAME --force -f
...
```

2.4. DEPLOY A JBOSS EAP SOURCE-TO-IMAGE (S2I) APPLICATION TO OPENSIFT

1. Create a new OpenShift application using the JBoss EAP CD for OpenShift image and your Java application's source code. Red Hat recommends using one of the provided JBoss EAP CD for OpenShift templates for S2I builds.

For example, for the **kitchensink** quickstart, use the following command to use the **eap-cd-basic-s2i** template in the **eap-demo** project, created in [Prepare OpenShift for Application Deployment](#), with the **kitchensink** source code on GitHub.

```
oc new-app --template=eap-cd-basic-s2i \ 1
-p IMAGE_STREAM_NAMESPACE=eap-demo \ 2
-p SOURCE_REPOSITORY_URL=https://github.com/jboss-developer/jboss-eap-quickstarts \ 3
-p SOURCE_REPOSITORY_REF=openshift \ 4
-p CONTEXT_DIR=kitchensink 5
```

- 1 The template to use.
- 2 The latest images streams and templates [were imported into the project's namespace](#), so you must specify the namespace of where to find the image stream. This is usually the project's name.
- 3 URL to the repository containing the application source code.
- 4 The Git repository reference to use for the source code. This can be a Git branch or tag reference.
- 5 The directory within the source repository to build.



NOTE

A template can specify default values for many template parameters, and you might have to override some, or all, of the defaults. To see template information, including a list of parameters and any default values, use the command **oc describe template TEMPLATE_NAME**.



NOTE

You might also want to [configure environment variables](#) when creating your new OpenShift application.

For example, if you are using an HTTPS template such as **eap-cd-https-s2i**, you must specify the required [HTTPS environment variables](#) **HTTPS_NAME**, **HTTPS_PASSWORD**, and **HTTPS_KEYSTORE** to match your keystore details.

2. Retrieve the name of the build configuration.

```
$ oc get bc -o name
```

3. Use the name of the build configuration from the previous step to view the Maven progress of the build.

```
$ oc logs -f buildconfig/BUILD_CONFIG_NAME
```

For example, for the **kitchensink** quickstart, the following command shows the progress of the Maven build.

```
$ oc logs -f buildconfig/eap-app
```

2.5. POST DEPLOYMENT TASKS

Depending on your application, some tasks might need to be performed after your OpenShift application has been built and deployed. This might include exposing a service so that the application is viewable from outside of OpenShift, or scaling your application to a specific number of replicas.

1. Get the service name of your application using the following command.

```
$ oc get service
```

2. Expose the main service as a route so you can access your application from outside of OpenShift. For example, for the **kitchensink** quickstart, use the following command to expose the required service and port.

```
$ oc expose service/eap-app --port=8080
```



NOTE

If you used a template to create the application, the route might already exist. If it does, continue on to the next step.

3. Get the URL of the route.

```
$ oc get route
```

4. Access the application in your web browser using the URL. The URL is the value of the **HOST/PORT** field from previous command's output.

If your application does not use the JBoss EAP root context, append the context of the application to the URL. For example, for the **kitchensink** quickstart, the URL might be **http://HOST_PORT_VALUE/kitchensink/**.

5. Optionally, you can also scale up the application instance by running the following command. This increases the number of replicas to **3**.

```
$ oc scale deploymentconfig DEPLOYMENTCONFIG_NAME --replicas=3
```

For example, for the **kitchensink** quickstart, use the following command to scale up the application.

```
$ oc scale deploymentconfig eap-app --replicas=3
```

CHAPTER 3. CONFIGURING THE JBOSS EAP CD FOR OPENSIFT IMAGE FOR YOUR JAVA APPLICATION

The JBoss EAP CD for OpenShift image is preconfigured for basic use with your Java applications. However, you can configure the JBoss EAP instance inside the image. The recommended method is to use the OpenShift S2I process, together with application template parameters and environment variables.



IMPORTANT

Any configuration changes made on a running container will be lost when the container is restarted or terminated.

This includes any configuration changes made using scripts that are included with a traditional JBoss EAP installation, for example **add-user.sh** or the management CLI.

It is strongly recommended that you use the OpenShift S2I process, together with application template parameters and environment variables, to make any configuration changes to the JBoss EAP instance inside the JBoss EAP CD for OpenShift image.

3.1. HOW THE JBOSS EAP CD FOR OPENSIFT S2I PROCESS WORKS



NOTE

The variable **EAP_HOME** is used to denote the path to the JBoss EAP installation inside the JBoss EAP CD for OpenShift image.

The S2I process for JBoss EAP CD for OpenShift works as follows:

1. If a **pom.xml** file is present in the source code repository, a Maven build process is triggered that uses the contents of the **\$MAVEN_ARGS** environment variable. Although you can specify custom Maven arguments or options with the **\$MAVEN_ARGS** environment variable, Red Hat recommends that you use the **\$MAVEN_ARGS_APPEND** environment variable to do this. The **\$MAVEN_ARGS_APPEND** variable takes the default arguments from **\$MAVEN_ARGS** and appends the options from **\$MAVEN_ARGS_APPEND** to it.

By default, the OpenShift profile uses the Maven **package** goal, which includes system properties for skipping tests (**-DskipTests**) and enabling the Red Hat GA repository (**-Dcom.redhat.xpaas.repo**).



NOTE

To use Maven behind a proxy on JBoss EAP CD for OpenShift image, set the **\$HTTP_PROXY_HOST** and **\$HTTP_PROXY_PORT** environment variables. Optionally, you can also set the **\$HTTP_PROXY_USERNAME**, **HTTP_PROXY_PASSWORD**, and **HTTP_PROXY_NONPROXYHOSTS** variables.

2. The results of a successful Maven build are copied to the **EAP_HOME/standalone/deployments/** directory inside the JBoss EAP CD for OpenShift image. This includes all JAR, WAR, and EAR files from the source repository specified by the

\$ARTIFACT_DIR environment variable. The default value of **\$ARTIFACT_DIR** is the Maven target directory.

3. All files in the **configuration** source repository directory are copied to the **EAP_HOME/standalone/configuration/** directory inside the JBoss EAP CD for OpenShift image. If you want to use a custom JBoss EAP configuration file, it should be named **standalone-openshift.xml**.
4. All files in the **modules** source repository directory are copied to the **EAP_HOME/modules/** directory inside the JBoss EAP CD for OpenShift image.

See [Artifact Repository Mirrors](#) for additional guidance on how to instruct the S2I process to utilize the custom Maven artifacts repository mirror.

3.2. CONFIGURING JBOSS EAP CD FOR OPENSIFT USING ENVIRONMENT VARIABLES

Using environment variables is the recommended method of configuring the JBoss EAP CD for OpenShift image. See the OpenShift documentation for instructions on [specifying environment variables](#) for application containers and build containers.

For example, you can set the JBoss EAP instance's management username and password using environment variables when creating your OpenShift application:

```
oc new-app --template=eap-cd-basic-s2i \
-p IMAGE_STREAM_NAMESPACE=eap-demo \
-p SOURCE_REPOSITORY_URL=https://github.com/jboss-developer/jboss-eap-quickstarts \
-p SOURCE_REPOSITORY_REF=openshift \
-p CONTEXT_DIR=kitchensink \
-e ADMIN_USERNAME=myspecialuser \
-e ADMIN_PASSWORD=myspecialp@ssw0rd
```

Available environment variables for the JBoss EAP CD for OpenShift image are listed in [Reference Information](#).

3.2.1. JVM Memory Configuration

The OpenShift EAP image has a mechanism to automatically calculate the default JVM memory settings based on the current environment, but you can also configure the JVM memory settings using environment variables.

3.2.1.1. JVM Default Memory Settings

If a memory limit is defined for the current container, and the limit is lower than the total available memory, the default JVM memory settings are calculated automatically. Otherwise, the default JVM memory settings are the default defined in the **standalone.conf** file of the EAP version used as the base server for the image.

The container memory limit is retrieved from the file **/sys/fs/cgroup/memory/memory.limit_in_bytes**. The total available memory is retrieved using the **/proc/meminfo** command.

When memory settings are calculated automatically, the following formulas are used:

- Maximum heap size (-Xmx): fifty percent (50%) of user memory

- Initial heap size (-Xms): twenty-five percent (25%) of the calculated maximum heap size.

For example, the defined memory limit is 1 GB, and this limit is lower than the total available memory reported by `/proc/meminfo`, then the memory settings will be: `-Xms128m -Xmx512`

You can use the following environment variables to modify the JVM settings calculated automatically. Note that these variables are only used when default memory size is calculated automatically (in other words, when a valid container memory limit is defined).

- **JAVA_MAX_MEM_RATIO**
- **JAVA_INITIAL_MEM_RATIO**
- **JAVA_MAX_INITIAL_MEM**

You can disable automatic memory calculation by setting the value of the following two environment variables to 0.

- **JAVA_INITIAL_MEM_RATIO**
- **JAVA_MAX_MEM_RATIO**

3.2.1.2. JVM Garbage Collection Settings

The EAP image for OpenShift includes settings for both garbage collection and garbage collection logging

Garbage Collection Settings

```
-XX:+UseParallelOldGC -XX:MinHeapFreeRatio=10 -XX:MaxHeapFreeRatio=20 -XX:GCTimeRatio=4 -XX:AdaptiveSizePolicyWeight=90 -XX:+ExitOnOutOfMemoryError
```

Garbage Collection Logging Settings for Java 8 (non-modular JVM)

```
-verbose:gc -Xloggc:/opt/eap/standalone/log/gc.log -XX:+PrintGCDetails -XX:+PrintGCDateStamps -XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=5 -XX:GCLogFileSize=3M -XX:-TraceClassUnloading
```

Garbage Collection Logging Settings for Java 11 (modular JVM)

```
-Xlog:gc*:file=/opt/eap/standalone/log/gc.log:time,uptimemillis:filecount=5,filesize=3M
```

3.2.1.3. Resource Limits in Default Settings

If set, additional default settings are included in the image.

```
-XX:ParallelGCThreads={core-limit} -Djava.util.concurrent.ForkJoinPool.common.parallelism={core-limit} -XX:CICompilerCount=2
```

The value of `{core-limit}` is defined using the **JAVA_CORE_LIMIT** environment variable, or by the CPU core limit imposed by the container.

The value of **CICompilerCount** is always fixed as 2.



3.2.1.4. JVM Environment Variables

Use these environment variables to configure the JVM in the EAP for OpenShift image.

Table 3.1. JVM Environment Variables


Variable Name	Example	Default Value	JVM Settings	Description
JAVA_OPTS	-verbose:class	No default	Multiple	<p>JVM options to pass to the java command.</p> <div data-bbox="1034 499 1145 2128" style="background-image: repeating-linear-gradient(45deg, transparent, transparent 2px, black 2px, black 4px); background-size: 15px 15px; width: 100%; height: 100%;"></div> <div data-bbox="1219 506 1246 2119" style="writing-mode: vertical-rl; text-orientation: mixed;"> <p>N O T E Use JAVA_OPTS_APPEND to configure additional JVM settings</p> </div>


Variable Name	Example	Default Value	JVM Settings	Description
				<p>... if you use JAVA_OPTS, some unconfigurable default settings are not added to the</p>

Variable Name	Example	Default Value	JVM Settings	Description
				<p>  JVM settings. You must explicitly add these settings. </p> <p>  I M P O R T A N T </p>

Variable Name	Example	Default Value	JVM Settings	Description
				<p> This is a table describing the JVM settings added by default by the container. </p>

Variable Name	Example	Default Value	JVM Settings	Description
				<p>Disables the following settings include</p> <ul style="list-style-type: none"> - XX:MetaspaceSize=96M - Djava.net.pre

Variable Name	Example	Default Value	JVM Settings	Description	f e r r
					P v 4 S t a c k = t r u e - D j b o s s . m o d u l e s . s y s t e m . p k g s = o r g . j b o s s . l o g m a n a g e r

Variable Name	Example	Default Value	JVM Settings	Description	
					<pre> jdk . n a s h o r n . . a p i , c o m . s u n . c r y p t o . p r o v i d e r - D j a v a . a w t . h e a d l e s s = t r u e </pre>

Variable Name	Example	Default Value	JVM Settings	Description
				<p>When a condition is not enabled, the initial JVM</p>

Variable Name	Example	Default Value	JVM Settings	Description
				<p> y (- X m s) a n d m a x i m u m J a v a m e m o r y (- X m s) a r e n o t d e f i n e d . A d d t h e s e d e f a u l </p>

Variable Name	Example	Default Value	JVM Settings	Description
				<p>You use JAVA_OPTS to configure additional settings.</p>
<p>JAVA_OPTS_APPEND</p>	<p>- Dsome.property=value</p>	<p>No default</p>	<p>Multiple</p>	<p>User-specified Java options to append to generated options in JAVA_OPTS.</p>

Variable Name	Example	Default Value	JVM Settings	Description
JAVA_MAX_MEMORY_RATIO	50	50	-Xmx	Use this variable when the -Xmx option is not specified in JAVA_OPTS . The value of this variable is used to calculate a default maximum heap memory size based on the restrictions of the container. If this variable is used in a container without a memory constraint, the variable has no effect. If this variable is used in a container that does have a memory constraint, the value of -Xmx is set to the specified ratio of the container's available memory. The default value, 50 means that 50% of the available memory is used as an upper boundary. To skip calculation of maximum memory, set the value of this variable to 0. No -Xmx

Variable Name	Example	Default Value	JVM Settings	Description	
JAVA_INITIAL_MEM_RATIO	25	25	-Xms	<p>option will be added to JAVA_OPTS.</p> <p>Use this variable when the -Xms option is not specified in JAVA_OPTS. The value of this variable is used to calculate the default initial heap memory size based on the maximum heap memory. If this variable is used in a container without a memory constraint, the variable has no effect. If this variable is used in a container that does have a memory constraint, the value of -Xms is set to the specified ratio of the -Xmx memory. The default value, 25 means that 25% of the maximum memory is used as the initial heap size. To skip calculation of initial memory, set the value of this variable to 0. No -Xms option will be added to JAVA_OPTS.</p>	

Variable Name	Example	Default Value	JVM Settings	Description
JAVA_MAX_INITIAL_MEM	4096	4096	-Xms	Use this variable when the -Xms option is not specified in JAVA_OPTS . The value of this variable is used to calculate the maximum size of the initial memory heap. The value is expressed in megabytes (MB). If this variable is used in a container without a memory constraint, the variable has no effect. If this variable is used in a container that does have a memory constraint, the value of -Xms is set to the value specified in the variable. The default value, 4096, specifies that the maximum initial heap will never be larger than 4096MB.

Variable Name	Example	Default Value	JVM Settings	Description
JAVA_DIAGNOSTICS	true	false (disabled)	<p>The settings depend on the JDK used by the container.</p> <ul style="list-style-type: none"> Open JDK8 and Open JDK10: -XX:NativeMemoryTracking=summary -XX:+PrintGC -XX:+PrintGCDateStamps -XX:+PrintGCCTimeStamps -XX:+UnlockDiagnosticVMOptions Open JDK11: -Xlog:gc:utctime -XX:NativeMemoryTracking=summary Other JDK: TBD 	<p>Set the value of this variable to <i>true</i> to include diagnostic information in standard output when events occur. If this variable is defined as <i>true</i> in an environment where JAVA_DIAGNOSTICS has already been defined as <i>true</i>, diagnostics are still included.</p>

Variable Name	Example	Default Value	JVM Settings	Description
DEBUG	true	false	- agentlib:jdwp=transport=dt_socket,address=\$DEBUG_PORT,server=y,suspend=n	Enables remote debugging.
DEBUG_PORT	8787	8787	- agentlib:jdwp=transport=dt_socket,address=\$DEBUG_PORT,server=y,suspend=n	Specifies the port used for debugging.
JAVA_CORE_LIMIT		Undefined	- XX:parallelGCThreads - Djava.util.concurrent.ForkJoinPool.common.parallelism - XX:CICompilerCount	A user-defined limit on the number of cores. If the container reports a limit constraint, the value of the JVM settings is limited to the container core limit. The value of - XXCICompilerCount is always 2. By default, this variable is undefined. In that case, if a limit is not defined on the container, the JVM settings are not set.
GC_MIN_HEAP_FREE_RATIO	20	10	- XX:MinHeapFreeRatio	Minimum percentage of heap free after garbage collection to avoid expansion.

Variable Name	Example	Default Value	JVM Settings	Description
GC_MAX_HEAP_FREE_RATIO	40	20	-XX:MaxHeapFreeRatio	Maximum percentage of heap free after garbage collection to avoid shrinking.
GC_TIME_RATIO	4	4	-XX:GCTimeRatio	Specifies the ratio of the time spent outside of garbage collection (for example, time spent in application execution) to the time spent in garbage collection.
GC_ADAPTIVE_SIZE_POLICY_WEIGHT	90	90	-XX:AdaptiveSizePolicyWeight	The weighting given to the current garbage collection time versus the previous garbage collection times.
GC_METASPACE_SIZE	20	96	-XX:MetaspaceSize	The initial metaspace size.
GC_MAX_METASPACE_SIZE	100	256	-XX:MaxMetaspaceSize	The maximum metaspace size.

Variable Name	Example	Default Value	JVM Settings	Description
GC_CONTAINER_OPTIONS	- XX:+UserG1GC	-XX:- UseParallelOld GC	-XX:- UseParallelOld GC	Specifies the Java garbage collection to use. The value of the variable should be the JRE command-line options to specify the required garbage collection. The JRE command specified overrides the default.

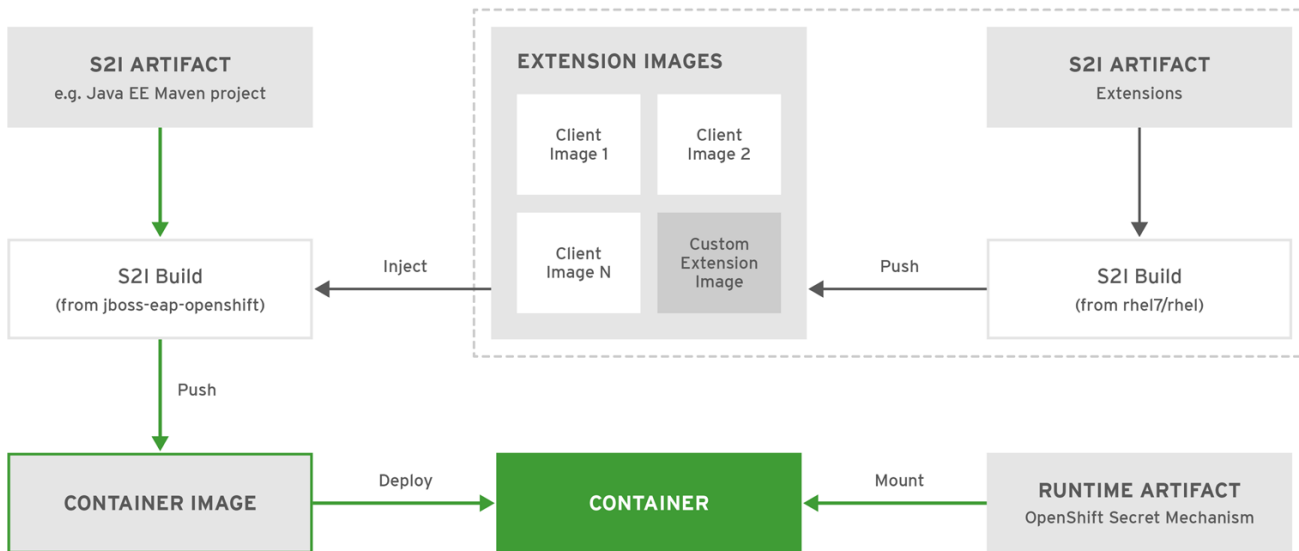
The following environment variables are deprecated:

- **JAVA_OPTIONS**: Use **JAVA_OPTS**.
- **INITIAL_HEAP_PERCENT**: Use **JAVA_INITIAL_MEM_RATIO**.
- **CONTAINER_HEAP_PERCENT**: Use **JAVA_MAX_MEM_RATIO**.

3.3. BUILD EXTENSIONS AND PROJECT ARTIFACTS

The JBoss EAP CD for OpenShift image extends database support in OpenShift using various artifacts. These artifacts are included in the built image through different mechanisms:

- [S2I artifacts](#) that are injected into the image during the S2I process.
- [Runtime artifacts](#) from environment files provided through the OpenShift Secret mechanism.



JBOSS_409952_0617



IMPORTANT

Support for using the Red Hat-provided internal datasource drivers with the JBoss EAP CD for OpenShift image is now deprecated. Red Hat recommends that you use JDBC drivers obtained from your database vendor for your JBoss EAP applications.

The following internal datasources are no longer provided with the JBoss EAP CD for OpenShift image:

- MySQL
- PostgreSQL

For more information about installing drivers, see [Modules, Drivers, and Generic Deployments](#).

For more information on configuring JDBC drivers with JBoss EAP, see [JDBC drivers](#) in the *JBoss EAP Configuration Guide*.

3.3.1. S2I Artifacts

The S2I artifacts include modules, drivers, and additional generic deployments that provide the necessary configuration infrastructure required for the deployment. This configuration is built into the image during the S2I process so that only the datasources and associated resource adapters need to be configured at runtime.

See [Artifact Repository Mirrors](#) for additional guidance on how to instruct the S2I process to utilize the custom Maven artifacts repository mirror.

3.3.1.1. Modules, Drivers, and Generic Deployments

There are a few options for including these S2I artifacts in the JBoss EAP CD for OpenShift image:

1. Include the artifact in the application source deployment directory. The artifact is downloaded during the build and injected into the image. This is similar to deploying an application on the JBoss EAP CD for OpenShift image.

2. Include the **CUSTOM_INSTALL_DIRECTORIES** environment variable, a list of comma-separated list of directories used for installation and configuration of artifacts for the image during the S2I process. There are two methods for including this information in the S2I:
 - An **install.sh** script in the nominated installation directory. The install script executes during the S2I process and operates with impunity.

install.sh Script Example

```
#!/bin/bash

injected_dir=$1
source /usr/local/s2i/install-common.sh
install_deployments ${injected_dir}/injected-deployments.war
install_modules ${injected_dir}/modules
configure_drivers ${injected_dir}/drivers.env
```

The **install.sh** script is responsible for customizing the base image using APIs provided by **install-common.sh**. **install-common.sh** contains functions that are used by the **install.sh** script to install and configure the modules, drivers, and generic deployments.

Functions contained within **install-common.sh**:

- **install_modules**
- **configure_drivers**
- **install_deployments**

Modules

A module is a logical grouping of classes used for class loading and dependency management. Modules are defined in the **EAP_HOME/modules/** directory of the application server. Each module exists as a subdirectory, for example **EAP_HOME/modules/org/apache/**. Each module directory then contains a slot subdirectory, which defaults to **main** and contains the **module.xml** configuration file and any required JAR files.

For more information about configuring **module.xml** files for MySQL and PostgreSQL JDBC drivers, see the [Datasource Configuration Examples](#) in the JBoss EAP Configuration Guide.

Example module.xml File

```
<?xml version="1.0" encoding="UTF-8"?>
<module xmlns="urn:jboss:module:1.0" name="org.apache.derby">
  <resources>
    <resource-root path="derby-10.12.1.1.jar"/>
    <resource-root path="derbyclient-10.12.1.1.jar"/>
  </resources>
  <dependencies>
    <module name="javax.api"/>
    <module name="javax.transaction.api"/>
  </dependencies>
</module>
```

Example module.xml File for PostgreSQL Datasource

```
<?xml version="1.0" encoding="UTF-8"?>
<module xmlns="urn:jboss:module:1.0" name="org.postgresql">
  <resources>
    <resource-root path="postgresql-jdbc.jar"/>
  </resources>
  <dependencies>
    <module name="javax.api"/>
    <module name="javax.transaction.api"/>
  </dependencies>
</module>
```

Example module.xml File for MySQL Connect/J 8 Datasource

```
<?xml version="1.0" encoding="UTF-8"?>
<module xmlns="urn:jboss:module:1.0" name="com.mysql">
  <resources>
    <resource-root path="mysql-connector-java-8.0.Z.jar" />
  </resources>
  <dependencies>
    <module name="javax.api"/>
    <module name="javax.transaction.api"/>
  </dependencies>
</module>
```



NOTE

The ".Z" in **mysql-connector-java-8.0.Z.jar** indicates the version of the **JAR** file downloaded. The file can be renamed, but the name must match the name in the **module.xml** file.

The **install_modules** function in **install.sh** copies the respective JAR files to the modules directory in JBoss EAP, along with the **module.xml**.

Drivers

Drivers are installed as modules. The driver is then configured in **install.sh** by the **configure_drivers** function, the configuration properties for which are defined in a [runtime artifact](#) environment file.

Example drivers.env File

```
#DRIVER
DRIVERS=DERBY
DERBY_DRIVER_NAME=derby
DERBY_DRIVER_MODULE=org.apache.derby
DERBY_DRIVER_CLASS=org.apache.derby.jdbc.EmbeddedDriver
DERBY_XA_DATASOURCE_CLASS=org.apache.derby.jdbc.EmbeddedXADataSource
```

Adding Datasource Drivers

The MySQL and PostgreSQL datasources are no longer provided as pre-configured

internal datasources. It is recommended that JDBC drivers should be obtained from the database vendor for your JBoss EAP applications. However, these drivers can still be installed as modules as described in [Modules, Drivers, and Generic Deployments](#)

The mechanism follows the **Derby** driver example and uses S2I artifacts. Create a **drivers.env** file for each datasource to be installed.

Example drivers.env File for MySQL Datasource

```
#DRIVER
DRIVERS=MYSQL
MYSQL_DRIVER_NAME=mysql
MYSQL_DRIVER_MODULE=org.mysql
MYSQL_DRIVER_CLASS=com.mysql.cj.jdbc.Driver
MYSQL_XA_DATASOURCE_CLASS=com.mysql.jdbc.jdbc2.optional.MysqlXADataSource
```

Example drivers.env File for PostgreSQL Datasource

```
#DRIVER
DRIVERS=POSTGRES
POSTGRES_DRIVER_NAME=postgres
POSTGRES_DRIVER_MODULE=org.postgresql
POSTGRES_DRIVER_CLASS=org.postgresql.Driver
POSTGRES_XA_DATASOURCE_CLASS=org.postgresql.xa.PGXADatasource
```

For information about download locations for various drivers, such as MySQL or PostgreSQL, see [JDBC Driver Download Locations](#) in the Configuration Guide.

Generic Deployments

Deployable archive files, such as JARs, WARs, RARs, or EARs, can be deployed from an injected image using the **install_deployments** function supplied by the API in **install-common.sh**.

- If the **CUSTOM_INSTALL_DIRECTORIES** environment variable has been declared but no **install.sh** scripts are found in the custom installation directories, the following artifact directories will be copied to their respective destinations in the built image:
 - **modules/*** copied to **\$JBOSS_HOME/modules/system/layers/openshift**
 - **configuration/*** copied to **\$JBOSS_HOME/standalone/configuration**
 - **deployments/*** copied to **\$JBOSS_HOME/standalone/deployments**

This is a basic configuration approach compared to the **install.sh** alternative, and requires the artifacts to be structured appropriately.

3.3.2. Runtime Artifacts

3.3.2.1. Datasources

There are two types of datasources:

1. Internal datasources. These datasources run on OpenShift, but are not available by default through the Red Hat Registry. Configuration of these datasources is provided by environment files added to OpenShift Secrets.
2. External datasources. These datasources do not run on OpenShift. Configuration of external datasources is provided by environment files added to OpenShift Secrets.

Example: Datasource Environment File

```
# derby datasource
ACCOUNTS_DERBY_DATABASE=accounts
ACCOUNTS_DERBY_JNDI=java:/accounts-ds
ACCOUNTS_DERBY_DRIVER=derby
ACCOUNTS_DERBY_USERNAME=derby
ACCOUNTS_DERBY_PASSWORD=derby
ACCOUNTS_DERBY_TX_ISOLATION=TRANSACTION_READ_UNCOMMITTED
ACCOUNTS_DERBY_JTA=true

# Connection info for xa datasource
ACCOUNTS_DERBY_XA_CONNECTION_PROPERTY_DataSourceName=/home/jboss/source/data/databases/derby/accounts

# _HOST and _PORT are required, but not used
ACCOUNTS_DERBY_SERVICE_HOST=dummy
ACCOUNTS_DERBY_SERVICE_PORT=1527
```

The **DATASOURCES** property is a comma-separated list of datasource property prefixes. These prefixes are then appended to all properties for that datasource. Multiple datasources can then be included in a single environment file. Alternatively, each datasource can be provided in separate environment files.

Datasources contain two types of properties: connection pool-specific properties and database driver-specific properties. Database driver-specific properties use the generic **XA_CONNECTION_PROPERTY**, because the driver itself is configured as a driver S2I artifact. The suffix of the driver property is specific to the particular driver for the datasource.

In the above example, **ACCOUNTS** is the datasource prefix, **XA_CONNECTION_PROPERTY** is the generic driver property, and **DataSourceName** is the property specific to the driver.

The datasources environment files are added to the OpenShift Secret for the project. These environment files are then called within the template using the **ENV_FILES** environment property, the value of which is a comma-separated list of fully qualified environment files as shown below.

```
{
  "Name": "ENV_FILES",
  "Value": "/etc/extensions/datasources1.env,/etc/extensions/datasources2.env"
}
```

3.3.2.2. Resource Adapters

Configuration of resource adapters is provided by environment files added to OpenShift Secrets.

Table 3.2. Resource Adapter Properties

Attribute	Description
<i>PREFIX_ID</i>	The identifier of the resource adapter as specified in the server configuration file.
<i>PREFIX_ARCHIVE</i>	The resource adapter archive.
<i>PREFIX_MODULE_SLOT</i>	The slot subdirectory, which contains the module.xml configuration file and any required JAR files.
<i>PREFIX_MODULE_ID</i>	The JBoss Module ID where the object factory Java class can be loaded from.
<i>PREFIX_CONNECTION_CLASS</i>	The fully qualified class name of a managed connection factory or admin object.
<i>PREFIX_CONNECTION_JNDI</i>	The JNDI name for the connection factory.
<i>PREFIX_PROPERTY_ParentDirectory</i>	Directory where the data files are stored.
<i>PREFIX_PROPERTY_AllowParentPaths</i>	Set AllowParentPaths to false to disallow .. in paths. This prevents requesting files that are not contained in the parent directory.
<i>PREFIX_POOL_MAX_SIZE</i>	The maximum number of connections for a pool. No more connections will be created in each sub-pool.
<i>PREFIX_POOL_MIN_SIZE</i>	The minimum number of connections for a pool.
<i>PREFIX_POOL_PREFILL</i>	Specifies if the pool should be prefilled. Changing this value requires a server restart.
<i>PREFIX_POOL_FLUSH_STRATEGY</i>	How the pool should be flushed in case of an error. Valid values are: FailingConnectionOnly (default), IdleConnections , and EntirePool .

The **RESOURCE_ADAPTERS** property is a comma-separated list of resource adapter property prefixes. These prefixes are then appended to all properties for that resource adapter. Multiple resource adapter can then be included in a single environment file. In the example below, **MYRA** is used as the prefix for a resource adapter. Alternatively, each resource adapter can be provided in separate environment files.

Example: Resource Adapter Environment File

```
#RESOURCE_ADAPTER
RESOURCE_ADAPTERS=MYRA
MYRA_ID=myra
MYRA_ARCHIVE=myra.rar
```

```
MYRA_CONNECTION_CLASS=org.javaee7.jca.connector.simple.connector.outbound.MyManagedConnectionFactory
MYRA_CONNECTION_JNDI=java:/eis/MySimpleMFC
```

The resource adapter environment files are added to the OpenShift Secret for the project namespace. These environment files are then called within the template using the **ENV_FILES** environment property, the value of which is a comma-separated list of fully qualified environment files as shown below.

```
{
  "Name": "ENV_FILES",
  "Value": "/etc/extensions/resourceadapter1.env,/etc/extensions/resourceadapter2.env"
}
```

3.4. DEPLOYMENT CONSIDERATIONS FOR THE JBOSS EAP CD FOR OPENSIFT IMAGE

3.4.1. Scaling Up and Persistent Storage Partitioning

There are two methods for deploying JBoss EAP with persistent storage: single-node partitioning, and multi-node partitioning.

Single-node partitioning stores the JBoss EAP data store directory, including transaction data, in the storage volume.

Multi-node partitioning creates additional, independent **split-*n*** directories to store the transaction data for each JBoss EAP pod, where *n* is an incremental integer. This communication is not altered if a JBoss EAP pod is updated, goes down unexpectedly, or is redeployed. When the JBoss EAP pod is operational again, it reconnects to the associated split directory and continues as before. If a new JBoss EAP pod is added, a corresponding **split-*n*** directory is created for that pod.

To enable the multi-node configuration you must set the **SPLIT_DATA** parameter to **true**. This results in the server creating independent **split-*n*** directories for each instance within the persistent volume which are used as their data store.



IMPORTANT

Due to the different storage methods of single-node and multi-node partitioning, changing a deployment from single-node to multi-node results in the application losing all data previously stored in the data directory, including messages, transaction logs, and so on. This is also true if changing a deployment from multi-node to single-node, as the storage paths will not match.

3.4.2. Scaling Down and Transaction Recovery

When the JBoss EAP CD for OpenShift image is deployed using a **multi-node** configuration, it is possible for unexpectedly terminated transactions to be left in the data directory of a terminating pod if the cluster is scaled down.

See [manual transaction recovery](#) to complete these branches.

CHAPTER 4. MIGRATING TO JBOSS EAP CD FOR OPENSIFT JDK 11 IMAGE

4.1. PREPARE OPENSIFT FOR APPLICATION DEPLOYMENT USING JDK 11 IMAGE

Preparing OpenShift for application deployment using the JDK 11 image stream follows the same procedure as described in [Prepare OpenShift for Application Deployment](#).

4.2. IMPORT JDK 11 IMAGE

Use the following command to import the JBoss EAP CD for OpenShift JDK 11 image stream and templates into your OpenShift project's namespace:

```
for resource in \
  eap73-openjdk11-image-stream.json \
  eap73-openjdk11-amq-persistent-s2i.json \
  eap73-openjdk11-amq-s2i.json \
  eap73-openjdk11-basic-s2i.json \
  eap73-openjdk11-https-s2i.json \
  eap73-openjdk11-sso-s2i.json \
  eap73-openjdk11-starter-s2i.json \
  eap73-openjdk11-third-party-db-s2i.json \
  eap73-openjdk11--tx-recovery-s2i.json
do
  oc replace --force -f \
  https://raw.githubusercontent.com/jboss-container-images/jboss-eap-7-openshift-image/eap73-
  openjdk11-ubi8/templates/${resource}
done
```

4.3. DEPLOY A JBOSS EAP S2I APPLICATION TO OPENSIFT USING JDK 11 IMAGE

Deploying a JBoss EAP S2I application to OpenShift follows the same procedures as described in [Deploy a JBoss EAP Source-to-Image \(S2I\) Application to OpenShift](#). The JDK 11 stream uses the **eap73-openjdk11-basic-s2i** for S2I builds, instead of the **eap73-basic-s2i** template used in JDK 8.

To deploy the **kitchensink** quickstart, use the following command to use the **eap73-openjdk11-basic-s2i** template with the **kitchensink** source code on GitHub:

```
oc new-app --template=eap73-openjdk11-basic-s2i
-p IMAGE_STREAM_NAMESPACE=eap-demo
-p SOURCE_REPOSITORY_URL=https://github.com/jboss-developer/jboss-eap-quickstarts.git
-p SOURCE_REPOSITORY_REF=openshift
-p CONTEXT_DIR=kitchensink
```

The **eap73-openjdk11-basic-s2i** template in the **eap-demo** project was created in [Prepare OpenShift for Application Deployment](#).

4.4. CONFIGURE JBOSS EAP CD FOR OPENSIFT USING ENVIRONMENT VARIABLES FOR JDK 11 IMAGE

Configuring JBoss EAP CD for OpenShift using environment variables for JDK 11 follows the same procedures as described in [Configuring JBoss EAP for OpenShift Using Environment Variables](#).

To set the JBoss EAP instance's management username and password using environment variables, use the following command when creating your OpenShift application.

```
oc new-app --template=eap73-openjdk11-basic-s2i \  
-p IMAGE_STREAM_NAMESPACE=eap-demo \  
-p SOURCE_REPOSITORY_URL=https://github.com/jboss-developer/jboss-eap-quickstarts \  
-p SOURCE_REPOSITORY_REF=openshift \  
-p CONTEXT_DIR=kitchensink \  
-e ADMIN_USERNAME=myspecialuser \  
-e ADMIN_PASSWORD=myspecialp@ssw0rd
```

CHAPTER 5. TROUBLESHOOTING

5.1. TROUBLESHOOTING POD RESTARTS

Pods can restart for a number of reasons, but a common cause of JBoss EAP pod restarts might include OpenShift resource constraints, especially out-of-memory issues. See the OpenShift documentation for more information on [OpenShift pod eviction](#).

By default, JBoss EAP CD for OpenShift templates are configured to automatically restart affected containers when they encounter situations like out-of-memory issues. The following steps can help you diagnose and troubleshoot out-of-memory and other pod restart issues.

1. Get the name of the pod that has been having trouble.
You can see pod names, as well as the number times each pod has restarted with the following command.

```
$ oc get pods
```

2. To diagnose why a pod has restarted, you can examine the JBoss EAP logs of the previous pod, or the OpenShift events.
 - a. To see the JBoss EAP logs of the previous pod, use the following command.

```
oc logs --previous POD_NAME
```

- b. To see the OpenShift events, use the following command.

```
$ oc get events
```

3. If a pod has restarted because of a resource issue, you can attempt to modify your OpenShift pod configuration to increase its [resource requests and limits](#). See the OpenShift documentation for more information on [configuring pod compute resources](#).

5.2. TROUBLESHOOTING USING THE JBOSS EAP MANAGEMENT CLI

The JBoss EAP management CLI, ***EAP_HOME/bin/jboss-cli.sh***, is accessible from within a container for troubleshooting purposes.



IMPORTANT

It is not recommended to make configuration changes in a running pod using the JBoss EAP management CLI. Any configuration changes made using the management CLI in a running container will be lost when the container restarts.

To make configuration changes to JBoss EAP CD for OpenShift, see [Configuring the JBoss EAP CD for OpenShift Image for Your Java Application](#).

1. First open a remote shell session to the running pod.

```
$ oc rsh POD_NAME
```

2. Run the following command from the remote shell session to launch the JBoss EAP management CLI:

```
█ $ /opt/eap/bin/jboss-cli.sh
```

CHAPTER 6. REFERENCE INFORMATION



NOTE

The content in this section is derived from the engineering documentation for this image. It is provided for reference as it can be useful for development purposes and for testing beyond the scope of the product documentation.

6.1. PERSISTENT TEMPLATES

The JBoss EAP database templates, which deploy JBoss EAP and database pods, have both ephemeral and persistent variations.

Persistent templates include an environment variable to provision a persistent volume claim, which binds with an available persistent volume to be used as a storage volume for the JBoss EAP CD for OpenShift deployment. Information, such as timer schema, log handling, or data updates, is stored on the storage volume, rather than in ephemeral container memory. This information persists if the pod goes down for any reason, such as project upgrade, deployment rollback, or an unexpected error.

Without a persistent storage volume for the deployment, this information is stored in the container memory only, and is lost if the pod goes down for any reason.

For example, an EE timer backed by persistent storage continues to run if the pod is restarted. Any events triggered by the timer during the restart process are enacted when the application is running again.

Conversely, if the EE timer is running in the container memory, the timer status is lost if the pod is restarted, and starts from the beginning when the pod is running again.

6.2. INFORMATION ENVIRONMENT VARIABLES

The following environment variables are designed to provide information to the image and should not be modified by the user:

Table 6.1. Information Environment Variables

Variable Name	Description and Value
JBOSS_IMAGE_NAME	The image name. Value: jboss-eap-7-tech-preview/eap-cd-openshift-rhel8
JBOSS_IMAGE_VERSION	The image version. Value: This is the image version number. See the Red Hat Container Catalog for the latest value.
JBOSS_MODULES_SYSTEM_PKGS	A comma-separated list of JBoss EAP system modules packages that are available to applications. Value: org.jboss.logmanager,jdk.nashorn.api

Variable Name	Description and Value
STI_BUILDER	Provides OpenShift S2I support for jee project types. Value: jee


6.3. CONFIGURATION ENVIRONMENT VARIABLES

You can configure the following environment variables to adjust the image without requiring a rebuild.

Table 6.2. Configuration Environment Variables

Variable Name	Description
AB_JOLOKIA_AUTH_OPENSIFT	<p>Switch on client authentication for OpenShift TLS communication. The value of this parameter can be true, false, or a relative distinguished name, which must be contained in a presented client's certificate. The default CA cert is set to /var/run/secrets/kubernetes.io/serviceaccount/ca.crt.</p> <ul style="list-style-type: none"> ● Set to false to disable client authentication for OpenShift TLS communication. ● Set to true to enable client authentication for OpenShift TLS communication using the default CA certificate and client principal. ● Set to a relative distinguished name, for example cn=someSystem, to enable client authentication for OpenShift TLS communication but override the client principal. This distinguished name must be contained in a presented client's certificate.
AB_JOLOKIA_CONFIG	<p>If set, uses this fully qualified file path for the Jolokia JVM agent properties, which are described in the Jolokia reference documentation. If you set your own Jolokia properties config file, the rest of the Jolokia settings in this document are ignored.</p> <p>If not set, /opt/jolokia/etc/jolokia.properties is created using the settings as defined in the Jolokia reference documentation.</p> <p>Example value: /opt/jolokia/custom.properties</p>
AB_JOLOKIA_DISCOVERY_ENABLED	<p>Enable Jolokia discovery.</p> <p>Defaults to false.</p>
AB_JOLOKIA_HOST	<p>Host address to bind to.</p> <p>Defaults to 0.0.0.0.</p> <p>Example value: 127.0.0.1</p>

Variable Name	Description
AB_JOLOKIA_HTTPS	<p>Switch on secure communication with HTTPS.</p> <p>By default self-signed server certificates are generated if no serverCert configuration is given in AB_JOLOKIA_OPTS.</p> <p>Example value: true</p>
AB_JOLOKIA_ID	<p>Agent ID to use.</p> <p>The default value is the \$HOSTNAME, which is the container id.</p> <p>Example value: openjdk-app-1-xqlsj</p>
AB_JOLOKIA_OFF	<p>If set to true, disables activation of Jolokia, which echos an empty value.</p> <p>Jolokia is enabled by default.</p>
AB_JOLOKIA_OPTS	<p>Additional options to be appended to the agent configuration. They should be given in the format key=value, key=value, ...</p> <p>Example value: backlog=20</p>
AB_JOLOKIA_PASSWORD	<p>The password for basic authentication.</p> <p>By default, authentication is switched off.</p> <p>Example value: mypassword</p>
AB_JOLOKIA_PASSWORD_RANDOM	<p>Determines if a random AB_JOLOKIA_PASSWORD should be generated.</p> <p>Set to true to generate a random password. The generated value is saved in the /opt/jolokia/etc/jolokia.pw file.</p>
AB_JOLOKIA_PORT	<p>The port to listen to.</p> <p>Defaults to 8778.</p> <p>Example value: 5432</p>
AB_JOLOKIA_USER	<p>The name of the user to use for basic authentication.</p> <p>Defaults to jolokia.</p> <p>Example value: myusername</p>

Variable Name	Description
AB_PROMETHEUS_ENABLE	<p>If set to true, this variable activates the jmx-exporter java agent that exposes Prometheus format metrics. Default is set to false.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>The MicroProfile Metrics subsystem is the preferred method to expose data in the Prometheus format. For more information about the MicroProfile Metrics subsystem, see Eclipse MicroProfile in the Configuration Guide.</p> </div> </div>
AB_PROMETHEUS_JMX_EXPORTER_CONFIG	<p>The path within the container to a user-specified configuration.yaml for the jmx-exporter agent to use instead of the default configuration.yaml file. To find out more about the S2I mechanism to incorporate additional configuration files, see S2I Artifacts.</p>
AB_PROMETHEUS_JMX_EXPORTER_PORT	<p>The port on which the jmx-exporter agent listens for scrapes from the Prometheus server. Default is 9799. The agent listens on localhost. Metrics can be made available outside of the container by configuring the DeploymentConfig file for the application to include the service exposing this endpoint.</p>
CLI_GRACEFUL_SHUTDOWN	<p>If set to any non-zero length value, the image will prevent shutdown with the TERM signal and will require execution of the shutdown command using the JBoss EAP management CLI.</p> <p>Example value: true</p>
CONTAINER_HEAP_PERCENT	<p>Set the maximum Java heap size, as a percentage of available container memory.</p> <p>Example value: 0.5</p>
CUSTOM_INSTALL_DIRECTORIES	<p>A list of comma-separated directories used for installation and configuration of artifacts for the image during the S2I process.</p> <p>Example value: custom,shared</p>
DEFAULT_JMS_CONNECTION_FACTORY	<p>This value is used to specify the default JNDI binding for the JMS connection factory, for example jms-connection-factory='java:jboss/DefaultJMSConnectionFactory'.</p> <p>Example value: java:jboss/DefaultJMSConnectionFactory</p>

Variable Name	Description
ENABLE_ACCESS_LOG	<p>Enable logging of access messages to the standard output channel.</p> <p>Logging of access messages is implemented using following methods:</p> <ul style="list-style-type: none"> • The JBoss EAP 6.4 OpenShift image uses a custom JBoss Web Access Log Valve. • The JBoss EAP CD for OpenShift image uses the Undertow AccessLogHandler. <p>Defaults to false.</p>
INITIAL_HEAP_PERCENT	<p>Set the initial Java heap size, as a percentage of the maximum heap size.</p> <p>Example value: 0.5</p>
JAVA_OPTS_APPEND	<p>Server startup options.</p> <p>Example value: -Dfoo=bar</p>
JBOSS_MODULES_SYSTEM_PKGS_APPEND	<p>A comma-separated list of package names that will be appended to the JBOSS_MODULES_SYSTEM_PKGS environment variable.</p> <p>Example value: org.jboss.byteman</p>
JGROUPS_CLUSTER_PASSWORD	<p>Password used to authenticate the node so it is allowed to join the JGroups cluster. Required, when using ASYM_ENCRYPT JGroups cluster traffic encryption protocol. If not set, authentication is disabled, cluster communication is not encrypted and a warning is issued. Optional, when using SYM_ENCRYPT JGroups cluster traffic encryption protocol.</p> <p>Example value: mypassword</p>
JGROUPS_ENCRYPT_KEYSTORE	<p>Name of the keystore file within the secret specified via JGROUPS_ENCRYPT_SECRET variable, when using SYM_ENCRYPT JGroups cluster traffic encryption protocol. If not set, cluster communication is not encrypted and a warning is issued.</p> <p>Example value: jgroups.jceks</p>

Variable Name	Description
JGROUPS_ENCRYPT_KEystore_DIR	Directory path of the keystore file within the secret specified via JGROUPS_ENCRYPT_SECRET variable, when using SYM_ENCRYPT JGroups cluster traffic encryption protocol. If not set, cluster communication is not encrypted and a warning is issued. Example value: /etc/jgroups-encrypt-secret-volume
JGROUPS_ENCRYPT_NAME	Name associated with the server's certificate, when using SYM_ENCRYPT JGroups cluster traffic encryption protocol. If not set, cluster communication is not encrypted and a warning is issued. Example value: jgroups
JGROUPS_ENCRYPT_PASSWORD	Password used to access the keystore and the certificate, when using SYM_ENCRYPT JGroups cluster traffic encryption protocol. If not set, cluster communication is not encrypted and a warning is issued. Example value: mypassword
JGROUPS_ENCRYPT_PROTOCOL	JGroups protocol to use for encryption of cluster traffic. Can be either SYM_ENCRYPT or ASYM_ENCRYPT . Defaults to SYM_ENCRYPT . Example value: ASYM_ENCRYPT
JGROUPS_ENCRYPT_SECRET	Name of the secret, containing the JGroups keystore file used for securing the JGroups communications, when using SYM_ENCRYPT JGroups cluster traffic encryption protocol. If not set, cluster communication is not encrypted and a warning is issued. Example value: eap7-app-secret
JGROUPS_PING_PROTOCOL	JGroups protocol to use for node discovery. Can be either openshift.DNS_PING or openshift.KUBE_PING .
MQ_SIMPLE_DEFAULT_PHYSICAL_DESTINATION	For backwards compatibility, set to true to use MyQueue and MyTopic as physical destination name defaults instead of queue/MyQueue and topic/MyTopic .
OPENSIFT_DNS_PING_SERVICE_NAME	Name of the service exposing the ping port on the servers for the DNS discovery mechanism. Example value: eap-app-ping

Variable Name	Description
OPENSIFT_DNS_PING_SERVICE_PORT	The port number of the ping port for the DNS discovery mechanism. If not specified, an attempt will be made to discover the port number from the SRV records for the service, otherwise the default 8888 will be used. Defaults to 8888 .
OPENSIFT_KUBE_PING_LABELS	Clustering labels selector for the Kubernetes discovery mechanism. Example value: app=eap-app
OPENSIFT_KUBE_PING_NAMESPACE	Clustering project namespace for the Kubernetes discovery mechanism. Example value: myproject
SCRIPT_DEBUG	If set to true , ensures that the Bash scripts are executed with the -x option, printing the commands and their arguments as they are executed.

**NOTE**

Other environment variables not listed above that can influence the product can be found in the [JBoss EAP documentation](#).

6.4. APPLICATION TEMPLATES

Table 6.3. Application Templates

Variable Name	Description
AUTO_DEPLOY_EXPLODED	Controls whether exploded deployment content should be automatically deployed. Example value: false

6.5. EXPOSED PORTS

Table 6.4. Exposed Ports

Port Number	Description
8443	HTTPS
8778	Jolokia Monitoring

6.6. DATASOURCES

Datasources are automatically created based on the value of some of the environment variables.

The most important environment variable is **DB_SERVICE_PREFIX_MAPPING**, as it defines JNDI mappings for the datasources. The allowed value for this variable is a comma-separated list of **POOLNAME-DATABASETYPE=PREFIX** triplets, where:

- **POOLNAME** is used as the **pool-name** in the datasource.
- **DATABASETYPE** is the database driver to use.
- **PREFIX** is the prefix used in the names of environment variables that are used to configure the datasource.

6.6.1. JNDI Mappings for Datasources

For each **POOLNAME-DATABASETYPE=PREFIX** triplet defined in the **DB_SERVICE_PREFIX_MAPPING** environment variable, the launch script creates a separate datasource, which is executed when running the image.



NOTE

The first part (before the equal sign) of the **DB_SERVICE_PREFIX_MAPPING** should be lowercase.

The **DATABASETYPE** determines the driver for the datasource.

For more information about configuring a driver, see [Modules, Drivers, and Generic Deployments](#). The JDK 8 image has drivers for **postgresql** and **mysql** configured by default.



WARNING

Do not use any special characters for the **POOLNAME** parameter.

6.6.1.1. Database Drivers



IMPORTANT

Support for using the Red Hat-provided internal datasource drivers with the JBoss EAP CD for OpenShift image is now deprecated. Red Hat recommends that you use JDBC drivers obtained from your database vendor for your JBoss EAP applications.


The following internal datasources are no longer provided with the JBoss EAP CD for OpenShift image:

- MySQL
- PostgreSQL

For more information about installing drivers, see [Modules, Drivers, and Generic Deployments](#).

For more information on configuring JDBC drivers with JBoss EAP, see [JDBC drivers](#) in the *JBoss EAP Configuration Guide*.

Every JDK 8 image contains Java drivers for MySQL, PostgreSQL, and MongoDB databases deployed. Datasources are generated only for MySQL and PostgreSQL databases.



IMPORTANT

JDK 11 image streams do not contain drivers for MySQL, PostgreSQL, and MongoDB databases and the datasources are not generated.

For more information about installing drivers, see [Modules, Drivers, and Generic Deployments](#).




NOTE

For MongoDB database there are no JNDI mappings created because MongoDB is not a SQL database.

6.6.1.2. Datasource Configuration Environment Variables

To configure other datasource properties, use the following environment variables.



IMPORTANT

Be sure to replace the values for ***POOLNAME***, ***DATABASETYPE***, and ***PREFIX*** in the following variable names with the appropriate values. These replaceable values are described in this section and in the [Datasources](#) section.

Variable Name	Description
<code>POOLNAME_DATABASETYPE_SERVICE_HOST</code>	<p>Defines the database server's host name or IP address to be used in the datasource's connection-url property.</p> <p>Example value: 192.168.1.3</p>

Variable Name	Description
<code>POOLNAME_DATABASETYPE_SERVICE_PORT</code>	<p>Defines the database server's port for the datasource.</p> <p>Example value: 5432</p>
<code>PREFIX_BACKGROUND_VALIDATION</code>	<p>When set to true database connections are validated periodically in a background thread prior to use. Defaults to false, meaning the validate-on-match method is enabled by default instead.</p>
<code>PREFIX_BACKGROUND_VALIDATION_MILLIS</code>	<p>Specifies frequency of the validation, in milliseconds, when the background-validation database connection validation mechanism is enabled (<code>PREFIX_BACKGROUND_VALIDATION</code> variable is set to true). Defaults to 10000.</p>
<code>PREFIX_CONNECTION_CHECKER</code>	<p>Specifies a connection checker class that is used to validate connections for the particular database in use.</p> <p>Example value: org.jboss.jca.adapters.jdbc.extensions.postgres.PostgreSQLValidConnectionChecker</p>
<code>PREFIX_DATABASE</code>	<p>Defines the database name for the datasource.</p> <p>Example value: myDatabase</p>
<code>PREFIX_DRIVER</code>	<p>Defines Java database driver for the datasource.</p> <p>Example value: postgresql</p>
<code>PREFIX_EXCEPTION_SORTER</code>	<p>Specifies the exception sorter class that is used to properly detect and clean up after fatal database connection exceptions.</p> <p>Example value: org.jboss.jca.adapters.jdbc.extensions.mysql.MySQLExceptionSorter</p>
<code>PREFIX_JNDI</code>	<p>Defines the JNDI name for the datasource. Defaults to java:jboss/datasources/POOLNAME_DATABASETYPE, where POOLNAME and DATABASETYPE are taken from the triplet described above. This setting is useful if you want to override the default generated JNDI name.</p> <p>Example value: java:jboss/datasources/test-postgresql</p>
<code>PREFIX_JTA</code>	<p>Defines Java Transaction API (JTA) option for the non-XA datasource. The XA datasources are already JTA capable by default.</p> <p>Defaults to true.</p>

Variable Name	Description
<i>PREFIX_MAX_POOL_SIZE</i>	Defines the maximum pool size option for the datasource. Example value: 20
<i>PREFIX_MIN_POOL_SIZE</i>	Defines the minimum pool size option for the datasource. Example value: 1
<i>PREFIX_NONXA</i>	Defines the datasource as a non-XA datasource. Defaults to false .
<i>PREFIX_PASSWORD</i>	Defines the password for the datasource. Example value: password
<i>PREFIX_TX_ISOLATION</i>	Defines the java.sql.Connection transaction isolation level for the datasource. Example value: TRANSACTION_READ_UNCOMMITTED
<i>PREFIX_URL</i>	Defines connection URL for the datasource. Example value: jdbc:postgresql://localhost:5432/postgresdb
<i>PREFIX_USERNAME</i>	Defines the username for the datasource. Example value: admin

When running this image in OpenShift, the ***POOLNAME_DATABASETTYPE_SERVICE_HOST*** and ***POOLNAME_DATABASETTYPE_SERVICE_PORT*** environment variables are set up automatically from the database service definition in the OpenShift application template, while the others are configured in the template directly as **env** entries in container definitions under each pod template.

6.6.1.3. Examples

These examples show how value of the **DB_SERVICE_PREFIX_MAPPING** environment variable influences datasource creation.


6.6.1.3.1. Single Mapping

Consider value **test-postgresql=TEST**.

This creates a datasource with **java:jboss/datasources/test_postgresql** name. Additionally, all the required settings like password and username are expected to be provided as environment variables with the **TEST_** prefix, for example **TEST_USERNAME** and **TEST_PASSWORD**.

6.6.1.3.2. Multiple Mappings

You can specify multiple datasource mappings.



NOTE

Always separate multiple datasource mappings with a comma.

Consider the following value for the **DB_SERVICE_PREFIX_MAPPING** environment variable: **cloud-postgresql=CLOUD,test-mysql=TEST_MYSQL**.

This creates the following two datasources:

1. **java:jboss/datasources/test_mysql**
2. **java:jboss/datasources/cloud_postgresql**

Then you can use **TEST_MYSQL** prefix for configuring things like the username and password for the MySQL datasource, for example **TEST_MYSQL_USERNAME**. And for the PostgreSQL datasource, use the **CLOUD_** prefix, for example **CLOUD_USERNAME**.


6.7. CLUSTERING

6.7.1. Configuring a JGroups Discovery Mechanism

To enable JBoss EAP clustering on OpenShift, configure the JGroups protocol stack in your JBoss EAP configuration to use either the **kubernetes.KUBE_PING** or **openshift.DNS_PING** discovery mechanism.

Although you can use a custom **standalone-openshift.xml** configuration file, it is recommended that you [use environment variables](#) to configure JGroups in your image build.

The instructions below use environment variables to configure the discovery mechanism for the JBoss EAP CD for OpenShift image.



IMPORTANT

If you use one of the available application templates to deploy an application on top of the JBoss EAP CD for OpenShift image, the default discovery mechanism is **openshift.DNS_PING**.

The **openshift.DNS_PING** and **kubernetes.KUBE_PING** discovery mechanisms are not compatible with each other. It is not possible to form a supercluster out of two independent child clusters, with one using the **openshift.DNS_PING** mechanism for discovery and the other using the **kubernetes.KUBE_PING** mechanism. Similarly, when performing a rolling upgrade, the discovery mechanism needs to be identical for both the source and the target clusters.

6.7.1.1. Configuring KUBE_PING

To use the **KUBE_PING** JGroups discovery mechanism:

1. The JGroups protocol stack must be configured to use **KUBE_PING** as the discovery mechanism.
You can do this by setting the **JGROUPS_PING_PROTOCOL** environment variable to **kubernetes.KUBE_PING**:

```
JGROUPS_PING_PROTOCOL=kubernetes.KUBE_PING
```

- The **KUBERNETES_NAMESPACE** environment variable must be set to your OpenShift project name. If not set, the server behaves as a single-node cluster (a "cluster of one"). For example:

```
KUBERNETES_NAMESPACE=PROJECT_NAME
```

- The **KUBERNETES_LABELS** environment variable should be set. This should match the [label set at the service level](#). If not set, pods outside of your application (albeit in your namespace) will try to join. For example:

```
KUBERNETES_LABELS=application=APP_NAME
```

- Authorization must be granted to the service account the pod is running under to be allowed to access Kubernetes' REST API. This is done using the OpenShift CLI. The following example uses the **default** service account in the current project's namespace:

```
oc policy add-role-to-user view system:serviceaccount:$(oc project -q):default -n $(oc project -q)
```

Using the **eap-service-account** in the project namespace:

```
oc policy add-role-to-user view system:serviceaccount:$(oc project -q):eap-service-account -n $(oc project -q)
```



NOTE

See [Prepare OpenShift for Application Deployment](#) for more information on adding policies to service accounts.

6.7.1.2. Configuring DNS_PING

To use the **DNS_PING** JGroups discovery mechanism:

- The JGroups protocol stack must be configured to use **DNS_PING** as the discovery mechanism. You can do this by setting the **JGROUPS_PING_PROTOCOL** environment variable to **openshift.DNS_PING**:

```
JGROUPS_PING_PROTOCOL=openshift.DNS_PING
```

- The **OPENSIFT_DNS_PING_SERVICE_NAME** environment variable must be set to the name of the ping service for the cluster. If not set, the server will act as if it is a single-node cluster (a "cluster of one").

```
OPENSIFT_DNS_PING_SERVICE_NAME=PING_SERVICE_NAME
```

- The **OPENSIFT_DNS_PING_SERVICE_PORT** environment variable should be set to the port number on which the ping service is exposed. The **DNS_PING** protocol attempts to discern the port from the SRV records, otherwise it defaults to **8888**.

```
OPENSIFT_DNS_PING_SERVICE_PORT=PING_PORT
```

- A ping service which exposes the ping port must be defined. This service should be headless (ClusterIP=None) and must have the following:

- a. The port must be named.
- b. The service must be annotated with **service.alpha.kubernetes.io/tolerate-unready-endpoints** set to **"true"**.



NOTE

Omitting this annotation will result in each node forming their own "cluster of one" during startup, then merging their cluster into the other nodes' clusters after startup, as the other nodes are not detected until after they have started.

```
kind: Service
apiVersion: v1
spec:
  clusterIP: None
  ports:
  - name: ping
    port: 8888
  selector:
    deploymentConfig: eap-app
metadata:
  name: eap-app-ping
  annotations:
    service.alpha.kubernetes.io/tolerate-unready-endpoints: "true"
    description: "The JGroups ping port for clustering."
```



NOTE

DNS_PING does not require any modifications to the service account and works using the default permissions.

6.7.2. Configuring JGroups to Encrypt Cluster Traffic

To encrypt cluster traffic for JBoss EAP on OpenShift, you must configure the JGroups protocol stack in your JBoss EAP configuration to use either the **SYM_ENCRYPT** or **ASYM_ENCRYPT** protocol.

Although you can use a custom **standalone-openshift.xml** configuration file, it is recommended that you [use environment variables](#) to configure JGroups in your image build.

The instructions below use environment variables to configure the protocol for cluster traffic encryption for the JBoss EAP CD for OpenShift image.



IMPORTANT

The **SYM_ENCRYPT** and **ASYM_ENCRYPT** protocols are not compatible with each other. It is not possible to form a supercluster out of two independent child clusters, with one using the **SYM_ENCRYPT** protocol for the encryption of cluster traffic and the other using the **ASYM_ENCRYPT** protocol. Similarly, when performing a rolling upgrade, the protocol needs to be identical for both the source and the target clusters.

6.7.2.1. Configuring SYM_ENCRYPT

To use the **SYM_ENCRYPT** protocol to encrypt JGroups cluster traffic:

1. The JGroups protocol stack must be configured to use **SYM_ENCRYPT** as the encryption protocol.
You can do this by setting the **JGROUPS_ENCRYPT_PROTOCOL** environment variable to **SYM_ENCRYPT**:

```
JGROUPS_ENCRYPT_PROTOCOL=SYM_ENCRYPT
```

2. The **JGROUPS_ENCRYPT_SECRET** environment variable must be set to the name of the secret containing the JGroups keystore file used for securing the JGroups communications. If not set, cluster communication is not encrypted and a warning is issued. For example:

```
JGROUPS_ENCRYPT_SECRET=eap7-app-secret
```

3. The **JGROUPS_ENCRYPT_KEYSTORE_DIR** environment variable must be set to the directory path of the keystore file within the secret specified via **JGROUPS_ENCRYPT_SECRET** variable. If not set, cluster communication is not encrypted and a warning is issued. For example:

```
JGROUPS_ENCRYPT_KEYSTORE_DIR=/etc/jgroups-encrypt-secret-volume
```

4. The **JGROUPS_ENCRYPT_KEYSTORE** environment variable must be set to the name of the keystore file within the secret specified via **JGROUPS_ENCRYPT_SECRET** variable. If not set, cluster communication is not encrypted and a warning is issued. For example:

```
JGROUPS_ENCRYPT_KEYSTORE=jgroups.jceks
```

5. The **JGROUPS_ENCRYPT_NAME** environment variable must be set to the name associated with the server's certificate. If not set, cluster communication is not encrypted and a warning is issued. For example:

```
JGROUPS_ENCRYPT_NAME=jgroups
```

6. The **JGROUPS_ENCRYPT_PASSWORD** environment variable must be set to the password used to access the keystore and the certificate. If not set, cluster communication is not encrypted and a warning is issued. For example:

```
JGROUPS_ENCRYPT_PASSWORD=mypassword
```

6.7.2.2. Configuring ASYM_ENCRYPT

To use the **ASYM_ENCRYPT** protocol to encrypt JGroups cluster traffic:

1. The JGroups protocol stack must be configured to use **ASYM_ENCRYPT** as the encryption protocol.
You can do this by setting the **JGROUPS_ENCRYPT_PROTOCOL** environment variable to **ASYM_ENCRYPT**:

```
JGROUPS_ENCRYPT_PROTOCOL=ASYM_ENCRYPT
```

2. The **JGROUPS_CLUSTER_PASSWORD** environment variable must be set to the password

used to authenticate the node so it is allowed to join the JGroups cluster. If not set, authentication is disabled, cluster communication is not encrypted and a warning is issued. For example:

```
JGROUPS_CLUSTER_PASSWORD=mypassword
```



NOTE

If you configure **ASYM_ENCRYPT** but also define any of the environment variables required for **SYM_ENCRYPT**, the **SYM_ENCRYPT** environment variables are ignored and a warning is issued. In this situation, cluster communication is still encrypted using **ASYM_ENCRYPT**.

6.8. HEALTH CHECKS

The JBoss EAP CD for OpenShift image utilizes the [liveness and readiness probes](#) included in OpenShift by default. In addition, this image includes [Eclipse MicroProfile Health](#), as discussed in the *Configuration Guide*.

The following table demonstrates the values necessary for these health checks to pass. If the status is anything other than the values found below, then the check is failed and the image is restarted per the image’s restart policy.

Table 6.5. Liveness and Readiness Checks

Performed Test	Liveness	Readiness
Server Status	Any status	Running
Boot Errors	None	None
Deployment Status ^[a]	N/A or no failed entries	N/A or no failed entries
Eclipse MicroProfile Health ^[b]	N/A or UP	N/A or UP
<p>^[a] N/A is only a valid state when no deployments are present.</p> <p>^[b] N/A is only a valid state when the microprofile-health-smallrye subsystem has been disabled.</p>		

6.9. MESSAGING

6.9.1. Configuring External Red Hat AMQ Brokers

You can configure the JBoss EAP CD for OpenShift image with environment variables to connect to external Red Hat AMQ brokers.

Example OpenShift Application Definition

The following example uses a template to create a JBoss EAP application connected to an external Red Hat AMQ 7 broker.

Example: JDK 8

```
oc new-app eap-cd-amq-s2i \
-p APPLICATION_NAME=eap-cd-mq \
-p MQ_USERNAME=MY_USERNAME \
-p MQ_PASSWORD=MY_PASSWORD
```

Example: JDK 11

```
oc new-app eap73-openjdk11-amq-s2i \
-p APPLICATION_NAME=eap-cd-mq \
-p MQ_USERNAME=MY_USERNAME \
-p MQ_PASSWORD=MY_PASSWORD
```

**IMPORTANT**

The template used in this example provides valid default values for the required parameters. If you do not use a template and provide your own parameters, be aware that the **MQ_SERVICE_PREFIX_MAPPING** name must match the **APPLICATION_NAME** name, appended with "-amq7=MQ".

6.10. SECURITY DOMAINS

To configure a new Security Domain, the user must define the **SECDOMAIN_NAME** environment variable.

This results in the creation of a security domain named after the environment variable. The user may also define the following environment variables to customize the domain:

Table 6.6. Security Domains

Variable name	Description
SECDOMAIN_NAME	Defines an additional security domain. Example value: myDomain
SECDOMAIN_PASSWORD_STACKING	If defined, the password-stacking module option is enabled and set to the value useFirstPass . Example value: true
SECDOMAIN_LOGIN_MODULE	The login module to be used. Defaults to UsersRoles
SECDOMAIN_USERS_PROPERTIES	The name of the properties file containing user definitions. Defaults to users.properties

Variable name	Description
SECDOMAIN_ROLES_PROPERTIES	The name of the properties file containing role definitions. Defaults to roles.properties

6.11. HTTPS ENVIRONMENT VARIABLES

Variable name	Description
HTTPS_NAME	If defined along with HTTPS_PASSWORD and HTTPS_KEYSTORE , enables HTTPS and sets the SSL name. This should be the value specified as the alias name of your keystore if you created it with the keytool -genkey command. Example value: example.com
HTTPS_PASSWORD	If defined along with HTTPS_NAME and HTTPS_KEYSTORE , enables HTTPS and sets the SSL key password. Example value: passw0rd
HTTPS_KEYSTORE	If defined along with HTTPS_PASSWORD and HTTPS_NAME , enables HTTPS and sets the SSL certificate key file to a relative path under EAP_HOME/standalone/configuration Example value: ssl.key

6.12. ADMINISTRATION ENVIRONMENT VARIABLES

Table 6.7. Administration Environment Variables

Variable name	Description
ADMIN_USERNAME	If both this and ADMIN_PASSWORD are defined, used for the JBoss EAP management user name. Example value: eapadmin
ADMIN_PASSWORD	The password for the specified ADMIN_USERNAME . Example value: passw0rd

6.13. S2I

The image includes S2I scripts and Maven.

Maven is currently only supported as a build tool for applications that are supposed to be deployed on JBoss EAP-based containers (or related/descendant images) on OpenShift.

Only WAR deployments are supported at this time.

6.13.1. Custom Configuration

It is possible to add custom configuration files for the image. All files put into **configuration/** directory will be copied into **EAP_HOME/standalone/configuration/**. For example to override the default configuration used in the image, just add a custom **standalone-openshift.xml** into the **configuration/** directory. See [example](#) for such a deployment.

6.13.1.1. Custom Modules

It is possible to add custom modules. All files from the **modules/** directory will be copied into **EAP_HOME/modules/**. See [example](#) for such a deployment.

6.13.2. Deployment Artifacts

By default, artifacts from the source **target** directory will be deployed. To deploy from different directories set the **ARTIFACT_DIR** environment variable in the BuildConfig definition. **ARTIFACT_DIR** is a comma-delimited list. For example: **ARTIFACT_DIR=app1/target,app2/target,app3/target**

6.13.3. Artifact Repository Mirrors

A repository in Maven holds build artifacts and dependencies of various types, for example, all of the project JARs, library JARs, plug-ins, or any other project specific artifacts. It also specifies locations from where to download artifacts while performing the S2I build. Besides using central repositories, it is a common practice for organizations to deploy a local custom mirror repository.

Benefits of using a mirror are:

- Availability of a synchronized mirror, which is geographically closer and faster.
- Ability to have greater control over the repository content.
- Possibility to share artifacts across different teams (developers, CI), without the need to rely on public servers and repositories.
- Improved build times.

Often, a repository manager can serve as local cache to a mirror. Assuming that the repository manager is already deployed and reachable externally at **http://10.0.0.1:8080/repository/internal/**, the S2I build can then use this manager by supplying the **MAVEN_MIRROR_URL** environment variable to the build configuration of the application as follows:

1. Identify the name of the build configuration to apply **MAVEN_MIRROR_URL** variable against.

```
oc get bc -o name
buildconfig/eap
```

2. Update build configuration of **eap** with a **MAVEN_MIRROR_URL** environment variable.

```
oc env bc/eap MAVEN_MIRROR_URL="http://10.0.0.1:8080/repository/internal/"
buildconfig "eap" updated
```

-
3. Verify the setting.

```
oc env bc/eap --list
# buildconfigs eap
MAVEN_MIRROR_URL=http://10.0.0.1:8080/repository/internal/
```

4. Schedule new build of the application.



NOTE

During application build, you will notice that Maven dependencies are pulled from the repository manager, instead of the default public repositories. Also, after the build is finished, you will see that the mirror is filled with all the dependencies that were retrieved and used during the build.

6.13.4. Scripts

run

This script uses the **openshift-launch.sh** script that configures and starts JBoss EAP with the **standalone-openshift.xml** configuration.

assemble

This script uses Maven to build the source, create a package (WAR), and move it to the **EAP_HOME/standalone/deployments** directory.

6.13.5. Custom Scripts

You can add custom scripts to run when starting a pod, before JBoss EAP is started.

You can add any script valid to run when starting a pod, including CLI scripts.

Two options are available for including scripts when starting JBoss EAP from an image:

- Mount a configmap to be executed as `postconfigure.sh`
- Add an `install.sh` script in the nominated installation directory

6.13.5.1. Mounting a configmap to execute custom scripts

Mount a configmap when you want any changes from the script to be implemented when the `s2i` image is built by the build controller.

To mount a configmap:

1. Create a configmap with content you want to include in the `postconfigure.sh`.
For example, if a subdirectory called **extensions** included the scripts **postconfigure.sh** and **extensions.cli**, you would run the following command:

```
$ oc create configmap jboss-cli --from-file=postconfigure.sh=extensions/postconfigure.sh --
from-file=extensions.cli=extensions/extensions.cli
```

2. Mount the configmap into the pods via the deployment controller (dc).

```
$ oc set volume dc/eap-app --add --name=jboss-cli -m /opt/eap/extensions -t configmap --
configmap-name=jboss-cli --default-mode='0755' --overwrite
```

Example `postconfigure.sh`

```
#!/usr/bin/env bash
set -x
echo "Executing postconfigure.sh"
$JBOSS_HOME/bin/jboss-cli.sh --file=$JBOSS_HOME/extensions/extensions.cli
```

Example `extensions.cli`

```
embed-server --std-out=echo --server-config=standalone-openshift.xml
:whoami
quit
```

6.13.5.2. Using `install.sh` to execute custom scripts

Use this option when you want the script to run each time a pod is started after building the image.

To execute custom scripts using `install.sh`:

1. In the git repository of the project that will be used during `s2i` build, create a directory called **.s2i**.
2. Inside the **.s2i** directory, add a file called `environment`, with the following content:

```
$ cat .s2i/environment
CUSTOM_INSTALL_DIRECTORIES=extensions
```

3. Create a directory called **extensions**.
4. In the **extensions** directory, create the file `postconfigure.sh` with contents similar to the following (replace placeholder code with appropriate code for your environment):

```
$ cat extensions/postconfigure.sh
#!/usr/bin/env bash
echo "Executing patch.cli"
$JBOSS_HOME/bin/jboss-cli.sh --file=$JBOSS_HOME/extensions/some-cli-example.cli
```

5. In the `extensions` directory, create the file `install.sh` with contents similar to the following (replace placeholder code with appropriate code for your environment):

```
$ cat extensions/install.sh
#!/usr/bin/env bash
set -x
echo "Running $PWD/install.sh"
injected_dir=$1
# copy any needed files into the target build.
cp -rf ${injected_dir} $JBOSS_HOME/extensions
```

6.13.6. Environment Variables

You can influence the way the build is executed by supplying environment variables to the **s2i build** command. The environment variables that can be supplied are:

Table 6.8. s2i Environment Variables

Variable name	Description
ARTIFACT_DIR	The .war , .ear , and .jar files from this directory will be copied into the deployments/ directory. Example value: target
HTTP_PROXY_HOST	Host name or IP address of a HTTP proxy for Maven to use. Example value: 192.168.1.1
HTTP_PROXY_PORT	TCP Port of a HTTP proxy for Maven to use. Example value: 8080
HTTP_PROXY_USERNAME	If supplied with HTTP_PROXY_PASSWORD , use credentials for HTTP proxy. Example value: myusername
HTTP_PROXY_PASSWORD	If supplied with HTTP_PROXY_USERNAME , use credentials for HTTP proxy. Example value: mypassword
HTTP_PROXY_NONPROXYHOSTS	If supplied, a configured HTTP proxy will ignore these hosts. Example value: some.example.org *.example.net
MAVEN_ARGS	Overrides the arguments supplied to Maven during build. Example value: -e -Popenshift -DskipTests -Dcom.redhat.xpaas.repo.redhatga package
MAVEN_ARGS_APPEND	Appends user arguments supplied to Maven during build. Example value: -Dfoo=bar
MAVEN_MIRROR_URL	URL of a Maven Mirror/repository manager to configure. Example value: http://10.0.0.1:8080/repository/internal/
MAVEN_CLEAR_REPO	Optionally clear the local Maven repository after the build. Example value: true

Variable name	Description
APP_DATADIR	If defined, directory in the source from where data files are copied. Example value: mydata
DATA_DIR	Directory in the image where data from \$APP_DATADIR will be copied. Example value: EAP_HOME/data

**NOTE**

For more information, see [Build and Run a Java Application on the JBoss EAP CD for OpenShift Image](#), which uses Maven and the S2I scripts included in the JBoss EAP CD for OpenShift image.

6.14. SSO

This image contains support for Red Hat JBoss SSO-enabled applications.

**NOTE**

See the [Red Hat JBoss SSO for OpenShift documentation](#) for more information on how to deploy the Red Hat JBoss SSO for OpenShift image with the JBoss EAP CD for OpenShift image.

Table 6.9. SSO Environment Variables

Variable name	Description
SSO_URL	URL of the SSO server.
SSO_REALM	SSO realm for the deployed applications.
SSO_PUBLIC_KEY	Public key of the SSO Realm. This field is optional but if omitted can leave the applications vulnerable to man-in-middle attacks.
SSO_USERNAME	SSO User required to access the SSO REST API. Example value: mySsoUser
SSO_PASSWORD	Password for the SSO user defined by the SSO_USERNAME variable. Example value: 6fedmL3P
SSO_SAML_KEYSTORE	Keystore location for SAML. Defaults to /etc/ss0-saml-secret-volume/keystore.jks .

Variable name	Description
SSO_SAML_KEYSTORE_PASSWORD	Keystore password for SAML. Defaults to mykeystorepass .
SSO_SAML_CERTIFICATE_NAME	Alias for keys/certificate to use for SAML. Defaults to jboss .
SSO_BEARER_ONLY	SSO Client Access Type. (Optional) Example value: true
SSO_CLIENT	Path for SSO redirects back to the application. Defaults to match module-name .
SSO_ENABLE_CORS	If true , enable CORS for SSO applications. (Optional)
SSO_SECRET	The SSO Client Secret for Confidential Access. Example value: KZ1Qylq4
SSO_DISABLE_SSL_CERTIFICATE_VALIDATION	If true the SSL/TLS communication between JBoss EAP and the RH-SSO server will be unsecure, for example, the certificate validation is disabled with curl . Not set by default. Example value: true

6.15. TRANSACTION RECOVERY

When a cluster is scaled down, it is possible for transaction branches to be in doubt. In these cases, [manual transaction recovery](#) might be necessary.

6.15.1. Unsupported Transaction Recovery Scenarios

- JTS transactions
Because the network endpoint of the parent is encoded in recovery coordinator IORs, recovery cannot work reliably if either the child or parent node recovers with either a new IP address, or if it is intended to be accessed using a virtualized IP address.
- XTS transactions
XTS does not work in a clustered scenario for recovery purposes. See [JBTM-2742](#) for details.
- Transactions propagated over [JBoss Remoting](#)
- Transactions propagated over XATerminator
Because the EIS is intended to be connected to a single instance of a Java EE application server, there are no well-defined ways to couple these processes.

6.15.2. Manual Transaction Recovery Process

6.15.2.1. Caveats

This procedure only describes how to manually recover transactions that were wholly self-contained within a single JVM. The procedure does not describe how to recover JTA transactions that have been propagated to other JVMs.



IMPORTANT

There are various network partition scenarios in which OpenShift might start multiple instances of the same pod with the same IP address and same node name and where, due to the partition, the old pod is still running. During manual recovery, this might result in a situation where you might be connected to a pod that has a stale view of the object store. If you think you are in this scenario, it is recommended that all JBoss EAP pods be shut down to ensure that none of the resource managers or object stores are in use.

When you enlist a resource in an XA transaction, it is your responsibility to ensure that each resource type is supported for recovery. For example, it is known that PostgreSQL and MySQL are well-behaved with respect to recovery, but for others, such as A-MQ and JDV resource managers, you should check documentation of the specific OpenShift release.

The deployment must use a [JDBC object store](#).



IMPORTANT

The transaction manager relies on the uniqueness of node identifiers. The maximum byte length of an XID is set by the XA specification and cannot be changed. Due to the data that the JBoss EAP CD for OpenShift image must include in the XID, this leaves room for 23 bytes in the node identifier.

OpenShift coerces the node identifier to fit this 23 byte limit:

- For all node names, even those under 23 bytes, the - (dash) character is stripped out.
- If the name is still over 23 bytes, characters are truncated from the beginning of the name until length of the name is within the 23 byte limit.

However, this process might impact the uniqueness of the identifier. For example, the names **aaa123456789012345678m0jwh** and **bbb123456789012345678m0jwh** are both truncated to **123456789012345678m0jwh**, which breaks the uniqueness of the names that are expected. In another example, **this-pod-is-m0jwh** and **thispod-is-m0jwh** are both truncated to **thispodism0jwh**, again breaking the uniqueness of the names.

It is your responsibility to ensure that the node names you configure are unique, keeping in mind the above truncation process.

6.15.2.2. Prerequisite

It is assumed the OpenShift instance has been configured with a JDBC store, and that the store tables are partitioned using a table prefix corresponding to the pod name. This should be automatic whenever a JBoss EAP deployment is in use. You can verify that the JBoss EAP instance is using a JDBC object store by looking at the configuration of the transactions subsystem in a running pod:

1. Determine if the `/opt/eap/standalone/configuration/openshift-standalone.xml` configuration file contains an element for the transaction subsystem:

```
<subsystem xmlns="urn:jboss:domain:transactions:3.0">
```

- If the JDBC object store is in use, then there is an entry similar to the following:


```
<jdbc-store datasource-jndi-name="java:jboss/datasources/jdbcstore_postgresql"/>
```



NOTE

The JNDI name identifies the datasource used to store the transaction logs.

6.15.2.3. Procedure



IMPORTANT

The following procedure details the process of manual transaction recovery solely for datasources.

- Use the database vendor tooling to list the XIDs (transaction branch identifiers) for in-doubt branches. It is necessary to list XIDs *for all datasources that were in use by any deployments running on the pod* that failed or was scaled down. Refer to the vendor documentation for the database product in use.
- For each such XID, determine which pod created the transaction and check to see if that pod is still running.
 - If it is running, then leave the branch alone.
 - If the pod is not running, assume it was removed from the cluster and you must apply the manual resolution procedure described here. Look in the transaction log storage that was used by the failed pod to see if there is a corresponding transaction log:
 - If there is a log, then manually commit the XID using the vendor tooling.
 - If there is not a log, assume it is an orphaned branch and roll back the XID using the vendor tooling.

The rest of this procedure explains in detail how to carry out each of these steps.

6.15.2.3.1. Resolving In-doubt Branches

First, find all the resources that the deployment is using.

It is recommended that you do this using the JBoss EAP management CLI. Although the resources should be defined in the JBoss EAP **standalone-openshift.xml** configuration file, there are other ways they can be made available to the transaction subsystem within the application server. For example, this can be done using a file in a deployment, or dynamically using the management CLI at runtime.

- Open a terminal on a pod running a JBoss EAP instance in the cluster of the failed pod. If there is no such pod, scale up to one.
- Create a management user using the **/opt/eap/bin/add-user.sh** script.
- Log into the management CLI using the **/opt/eap/bin/jboss-cli.sh** script.
- List the datasources configured on the server. These are the ones that may contain in-doubt transaction branches.


```

/subsystem=datasources:read-resource
{
  "outcome" => "success",
  "result" => {
    "data-source" => {
      "ExampleDS" => undefined,
      ...
    },
    ...
  }
}

```

- Once you have the list, find the connection URL for each of the datasources. For example:

```

/subsystem=datasources/data-source=ExampleDS:read-attribute(name=connection-url)
{
  "outcome" => "success",
  "result" => "jdbc:h2:mem:test;DB_CLOSE_DELAY=-1;DB_CLOSE_ON_EXIT=FALSE",
  "response-headers" => {"process-state" => "restart-required"}
}

```

- Connect to each datasource and list any in-doubt transaction branches.



NOTE

The table name that stores in-doubt branches will be different for each datasource vendor.

JBoss EAP has a default SQL query tool (H2) that you can use to check each database. For example:

```

java -cp /opt/eap/modules/system/layers/base/com/h2database/h2/main/h2-1.3.173.jar \
-url "jdbc:postgresql://localhost:5432/postgres" \
-user sa \
-password sa \
-sql "select gid from pg_prepared_xacts;"

```

Alternatively, you can use the resource's native tooling. For example, for a PostgreSQL datasource called **sampledb**, you can use the OpenShift client tools to remotely log in to the pod and query the in-doubt transaction table:

```

$ oc rsh postgresql-2-vwf9n # rsh to the named pod
sh-4.2$ psql sampledb
psql (9.5.7)
Type "help" for help.

sampledb=# select gid from pg_prepared_xacts;
131077_AAAAAAAAAAAAAAP//rBEAB440GK1aJ72oAAAAGHAtanRhLWNyYXNoLXJlYy0zLXAj
Y2N3_AAAAAAAAAAAAAAP//rBEAB440GK1aJ72oAAAAGgAAAAEAAAAA

```

6.15.2.3.2. Extract the Global Transaction ID and Node Identifier from Each XID

When all XIDs for in-doubt branches are identified, convert the XIDs into a format that you can compare to the logs stored in the transaction tables of the transaction manager.

For example, the following Bash script can be used to perform this conversion. Assuming that **\$PG_XID** holds the XID from the [select statement](#) above, then the JBoss EAP transaction ID can be obtained as follows:

```
PG_XID="$1"
IFS='_' read -ra lines <<< "$PG_XID"
[[ "${lines[0]}" = 131077 ]] || exit 0; # this script only works for our own FORMAT ID
PG_TID=${lines[1]}

a=$(echo "$PG_TID" | base64 -d | xxd -ps | tr -d '\n' | while read -N16 i ; do echo 0x$i ; done)
b=$(echo "$PG_TID" | base64 -d | xxd -ps | tr -d '\n' | while read -N8 i ; do echo 0x$i ; done)
c=("${b[@]:4}") # put the last 3 32-bit hexadecimal numbers into array c
# the negative elements of c need special handling since printf below only works with positive
# hexadecimal numbers
for i in "${c[@]}"; do
    arg=${c[$i]}
    # inspect the MSB to see if arg is negative - if so convert it from a 2's complement number
    [[ (($arg >> 31)) = 1 ]] && x=$(echo "obase=16; $((($arg - 0x100000000))" | bc) || x=$arg
    if [[ ${x:0:1} = \- ]] ; then # see if the first character is a minus sign
        neg[$i]="-";
        c[$i]=0x${x:1} # strip the minus sign and make it hex for use with printf below
    else
        neg[$i]=""
        c[$i]=$x
    fi
done
EAP_TID=$(printf %x:%x:%x:%x:%x:%x:${neg[0]}%x:${neg[1]}%x:${neg[2]}%x ${a[0]} ${a[1]} ${c[0]} ${c[1]} ${c[2]})
```

After completion, the **\$EAP_TID** variable holds the global transaction ID of the transaction that created this XID. The node identifier of the pod that started the transaction is given by the output of the following bash command:

```
echo "$PG_TID" | base64 -d | tail -c +29
```



NOTE


The node identifier starts from the 29th character of the PostgreSQL global transaction ID field.

- If this pod is [still running](#), then leave this in-doubt branch alone since the transaction is still in flight.
- If this pod is not running, then you need to search the relevant transaction log storage for the transaction log. The log storage is located in a JDBC table, which is named following the **os<node-identifier>jbosststxtable** pattern.
 - If there is no such table, leave the branch alone as it is owned by some other transaction manager. The URL for the datasource containing this table is defined in the transaction subsystem description shown below.
 - If there is such a table, look for an entry that matches the global transaction ID.
 - If there is an entry in the table that matches the global transaction ID, then the in-doubt branch needs to be committed using the datasource vendor tooling as described below.

- If there is no such entry, then the branch is an orphan and can safely be rolled back.

An example of how to commit an in-doubt PostgreSQL branch is shown below:

```
$ oc rsh postgresql-2-vwf9n
sh-4.2$ psql sampledb
psql (9.5.7)
Type "help" for help.
psql sampledb
commit prepared '131077_AAAAAAAAAAAAAAP//rBEAB440GK1aJ72oAAAAGHAtanRh
----
LWNyYXNoLXJlYy0zLXAyY2N3_AAAAAAAAAAAAAAP//rBEAB440GK1aJ72oAAAAGgAAAAEAAAAA';
```



IMPORTANT

Repeat this procedure for all datasources and in-doubt branches.


6.15.2.3.3. Obtain the List of Node Identifiers of All Running JBoss EAP Instances in Any Cluster that Can Contact the Resource Managers

Node identifiers are configured to be the same name as the pod name. You can obtain the pod names in use using the **oc** command. Use the following command to list the running pods:

```
$ oc get pods | grep Running
eap-manual-tx-recovery-app-4-26p4r 1/1 Running 0 23m
postgresql-2-vwf9n 1/1 Running 0 41m
```

For each running pod, look in the output of the pod's log and obtain the node name. For example, for first pod shown in the above output, use the following command:

```
$ oc logs eap-manual-tx-recovery-app-4-26p4r | grep "jboss.node.name" | head -1
jboss.node.name = tx-recovery-app-4-26p4r
```



IMPORTANT

The aforementioned [JBoss node name identifier](#) will always be truncated to the maximum length of 23 characters in total by removing characters from the beginning and retaining the trailing characters until the maximum length of 23 characters is reached.

6.15.2.3.4. Find the Transaction Logs

1. The transaction logs reside in a JDBC-backed object store. The JNDI name of this store is defined in the **transaction** subsystem definition of the JBoss EAP configuration file.
2. Look in the configuration file to find the datasource definition corresponding to the above JNDI name.
3. Use the JNDI name to derive the connection URL.
4. You can use the URL to connect to the database and issue a **select** query on the relevant in-doubt transaction table.

Alternatively, if you know which pod the database is running on, and you know the name of the database, it might be easier to open an OpenShift remote shell into the pod and use the database tooling directly.

For example, if the JDBC store is hosted by a PostgreSQL database called **sampledb** running on pod **postgresql-2-vwf9n**, then you can find the transaction logs using the following commands:



NOTE

The `ostxrecoveryapp426p4rjbossstxttable` table name listed in the following command has been chosen since it follows the [pattern for JDBC table names holding the log storage entries](#). In your environment the table name will have similar form:

- Starting with **os** prefix.
- The part in the middle is derived from the [JBoss node name above](#), possibly deleting the "-" (dash) character if present.
- Finally the **jbossstxttable** suffix is appended to create the final name of the table.

```
$ oc rsh postgresql-2-vwf9n
sh-4.2$ psql sampledb
psql (9.5.7)
Type "help" for help.

sampledb=# select uidstring from ostxrecoveryapp426p4rjbossstxttable where
TYPENAME='StateManager/BasicAction/TwoPhaseCoordinator/AtomicAction'
;
      uidstring
-----
0:ffff0a81009d:33789827:5a68b2bf:40
(1 row)
```

6.15.2.3.5. Cleaning Up the Transaction Logs for Reconciled In-doubt Branches



WARNING

Do not delete the log unless you are certain that there are no remaining in-doubt branches.

When all the branches for a given transaction are complete, and all potential resources managers have been checked, including A-MQ and JDV, it is safe to delete the transaction log.

Issue the following command, specify the transaction log to be removed using the appropriate **uidstring**:

```
DELETE FROM ostxrecoveryapp426p4rjbossststxtable where uidstring = UIDSTRING
```



IMPORTANT


If you do not delete the log, then completed transactions which failed after prepare, but which have now been resolved, will never be removed from the transaction log storage. The consequence of this is that unnecessary storage is used and future manual reconciliation will be more difficult.

6.16. INCLUDED JBOSS MODULES

The table below lists included JBoss Modules in the JBoss EAP CD for OpenShift image.

Table 6.10. Included JBoss Modules

JBoss Module
org.jboss.as.clustering.common
org.jboss.as.clustering.jgroups
org.jboss.as.ee
org.jboss.logmanager.ext
org.jgroups
org.mongodb
org.openshift.ping
org.postgresql
com.mysql
net.oauth.core



NOTE

The following modules are not included in the JDK 11 image:

- **org.mongodb**
- **org.postgresql**
- **com.mysql**

Revised on 2019-12-06 06:34:48 UTC