



Red Hat JBoss Middleware for OpenShift 3

Red Hat JBoss Middleware for OpenShift Release Notes

Information about the release and known issues

Red Hat JBoss Middleware for OpenShift 3 Red Hat JBoss Middleware for OpenShift Release Notes

Information about the release and known issues

Legal Notice

Copyright © 2018 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, 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 Software Collections 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

Release Note Information for Red Hat JBoss Middleware for OpenShift

Table of Contents

CHAPTER 1. RELEASE NOTES	3
CHAPTER 2. RED HAT JBOSS DATA VIRTUALIZATION FOR OPENSIFT	4
CHAPTER 3. RED HAT JBOSS BPM SUITE INTELLIGENT PROCESS SERVER FOR OPENSIFT	5
CHAPTER 4. RED HAT JBOSS SSO FOR OPENSIFT	6
CHAPTER 5. RED HAT JBOSS ENTERPRISE APPLICATION PLATFORM FOR OPENSIFT	7
CHAPTER 6. RED HAT JBOSS WEB SERVER FOR OPENSIFT	8
CHAPTER 7. RED HAT JBOSS A-MQ FOR OPENSIFT	9
CHAPTER 8. RED HAT JBOSS FUSE INTEGRATION SERVICES FOR OPENSIFT	10
CHAPTER 9. RED HAT JBOSS BRMS DECISION SERVER FOR OPENSIFT	11
CHAPTER 10. RED HAT JBOSS DATA GRID FOR OPENSIFT	12
CHAPTER 11. KNOWN ISSUES	13
11.1. ISSUES COMMON TO MULTIPLE IMAGES	13
11.1.1. Fixed issues	13
11.2. A-MQ FOR OPENSIFT	15
11.2.1. Known issues	15
11.3. DECISION SERVER FOR OPENSIFT AND IPS FOR OPENSIFT	16
11.3.1. Known issues	16
11.4. EAP FOR OPENSIFT	16
11.4.1. Fixed issues	16
11.4.2. Known issues	17
11.5. FIS FOR OPENSIFT	19
11.5.1. Known issues	19
11.6. JDG FOR OPENSIFT	20
11.6.1. Fixed issues	20
11.6.2. Known issues	20
11.7. JDV FOR OPENSIFT AND JDG FOR OPENSIFT INTEGRATION	21
11.7.1. Known issues	21
11.8. JWS FOR OPENSIFT	21
11.8.1. Fixed issues	21
11.8.2. Known issues	22
11.9. RH-SSO FOR OPENSIFT	22
11.9.1. Fixed issues	22
11.9.2. Known issues	23

CHAPTER 1. RELEASE NOTES

Starting in OpenShift Enterprise 3.0, Red Hat JBoss Middleware for OpenShift images are provided for the following:

- Red Hat JBoss Enterprise Application Platform (EAP)
- Red Hat JBoss Web Server (JWS)
- Red Hat JBoss A-MQ

Starting in OpenShift Enterprise 3.1, Red Hat JBoss Middleware for OpenShift images are also provided for the following:

- Red Hat JBoss Fuse (Fuse Integration Services)
- Red Hat JBoss BRMS (Decision Server)
- Red Hat JBoss Data Grid (JDG)

Starting in OpenShift Enterprise 3.2, Red Hat JBoss Middleware for OpenShift images are also provided for the following:

- Red Hat JBoss SSO for OpenShift
- Red Hat JBoss BPM Suite Intelligent Process Server for OpenShift

Red Hat JBoss Middleware for OpenShift images are now also provided for:

– Red Hat JBoss Data Virtualization (JDV)

See enterprise.openshift.com/middleware-services for additional information.

CHAPTER 2. RED HAT JBOSS DATA VIRTUALIZATION FOR OPENSIFT

Red Hat JBoss Data Virtualization is available as a [containerized image](#) that is designed for use with OpenShift Enterprise 3.2 and later.

However, there are [significant differences](#) in supported configurations and functionality in the JDV for OpenShift image compared to the regular release of JDV. Documentation for other JDV functionality not specific to the JDV for OpenShift image can be found in the [Red Hat JBoss Data Virtualization documentation on the Red Hat Customer Portal](#).

See the [Red Hat JBoss Data Virtualization for OpenShift](#) product page on the Red Hat Container Catalog for the current image versions and tags.

CHAPTER 3. RED HAT JBOSS BPM SUITE INTELLIGENT PROCESS SERVER FOR OPENSIFT

Red Hat JBoss BPM Suite (Process Server) is available as a [containerized image](#) that is designed for use with OpenShift Enterprise 3.2 and later.

However, there are [significant differences](#) in supported configurations and functionality in the IPS for OpenShift image compared to the regular release of Red Hat JBoss BPM Suite. Documentation for other Red Hat JBoss BPM Suite functionality not specific to the IPS for OpenShift image can be found in the [Red Hat JBoss BPM Suite documentation on the Red Hat Customer Portal](#) .

See the [Red Hat JBoss BPM Suite Intelligent Process Server for OpenShift](#) product page on the Red Hat Container Catalog for the current image versions and tags.

CHAPTER 4. RED HAT JBOSS SSO FOR OPENSIFT

Red Hat Single Sign-On is available as a [containerized image](#) that is designed for use with OpenShift Enterprise 3.2 and later.

However, there are [significant differences](#) in supported configurations and functionality in the RH-SSO for OpenShift image compared to the regular release of Red Hat Single Sign-On. Documentation for other Red Hat Single Sign-On functionality not specific to the RH-SSO for OpenShift image can be found in the [Red Hat Single Sign-On documentation on the Red Hat Customer Portal](#) .

See the [Red Hat JBoss SSO for OpenShift](#) product page on the Red Hat Container Catalog for the current image versions and tags.

CHAPTER 5. RED HAT JBOSS ENTERPRISE APPLICATION PLATFORM FOR OPENSIFT

Red Hat JBoss EAP is available as a [containerized image](#) that is designed for use with OpenShift Enterprise 3.0 and later.

However, there are [significant differences](#) in supported configurations and functionality in the EAP for OpenShift image compared to the regular release of JBoss EAP. Documentation for other JBoss EAP functionality not specific to the EAP for OpenShift image can be found in the [JBoss EAP documentation on the Red Hat Customer Portal](#).

See the [Red Hat JBoss Enterprise Application Platform for OpenShift](#) product page on the Red Hat Container Catalog for the current image versions and tags.

See the [Software Downloads](#) page of the Red Hat Customer Portal for a list of bug fixes or changes for Red Hat JBoss Enterprise Application Platform.

CHAPTER 6. RED HAT JBOSS WEB SERVER FOR OPENSIFT

The Apache Tomcat 7 and Apache Tomcat 8 components of Red Hat JBoss Web Server(JWS) 3.1 are available as [containerized images](#) that are designed for use with OpenShift Enterprise 3.0 and later.

However, there are [significant differences](#) in the functionality between the JWS for OpenShift images and the regular release of JBoss Web Server. Documentation for other JBoss Web Server functionality not specific to the JWS for OpenShift images can be found in the [JBoss Web Server documentation on the Red Hat Customer Portal](#).

See the [Red Hat JBoss Web Server for OpenShift](#) product page on the Red Hat Container Catalog for the current image versions and tags.

CHAPTER 7. RED HAT JBOSS A-MQ FOR OPENSIFT

Red Hat JBoss A-MQ is available as a [containerized image](#) that is designed for use with OpenShift Enterprise 3.0 and later. It allows developers to quickly deploy an A-MQ message broker in a hybrid cloud environment.

However, there are [significant differences](#) in supported configurations and functionality in the A-MQ for OpenShift image compared to the regular release of JBoss A-MQ. Release notes and documentation for other JBoss A-MQ functionality not specific to the A-MQ for OpenShift image can be found in the [JBoss A-MQ documentation on the Red Hat Customer Portal](#) .

See the [Red Hat JBoss A-MQ for OpenShift](#) product page on the Red Hat Container Catalog for the current image versions and tags.

See the [Security Patch](#) page of the Red Hat Customer Portal for a list of bug fixes or changes for this specific version of Red Hat JBoss A-MQ.

CHAPTER 8. RED HAT JBOSS FUSE INTEGRATION SERVICES FOR OPENSIFT

Red Hat JBoss Fuse is available as a [containerized image](#), known as Fuse Integration Services(FIS), that is designed for use with OpenShift Enterprise 3.1. It allows developers to quickly deploy applications in a hybrid cloud environment. In Fuse Integration Services, application runtime is dynamic.

However, there are [significant differences](#) in supported configurations and functionality in the FIS for OpenShift image compared to the regular release of JBoss Fuse. Documentation for other JBoss Fuse functionality not specific to the FIS for OpenShift image can be found in the [JBoss Fuse documentation on the Red Hat Customer Portal](#).

See the [Red Hat JBoss Fuse Integration Services for OpenShift](#) product page on the Red Hat Container Catalog for the current image versions and tags.

CHAPTER 9. RED HAT JBOSS BRMS DECISION SERVER FOR OPENSIFT

Red Hat JBoss BRMS is available as a [containerized Red Hat JBoss Middleware for OpenShift image](#) , known as Decision Server, that is designed for use with OpenShift Enterprise 3.1 as an execution environment for business rules. Developers can quickly build, scale, and test applications deployed across hybrid environments.

However, there are [significant differences](#) in supported configurations and functionality in the Decision Server for OpenShift image compared to the regular release of JBoss BRMS. Documentation for other JBoss BRMS functionality not specific to the Decision Server for OpenShift image can be found in the [JBoss BRMS documentation on the Red Hat Customer Portal](#) .

See the [Red Hat JBoss BRMS Decision Server for OpenShift](#) product page on the Red Hat Container Catalog for the current image versions and tags.

CHAPTER 10. RED HAT JBOSS DATA GRID FOR OPENSIFT

Red Hat JBoss Data Grid is available as a [containerized image](#) that is designed for use with OpenShift Enterprise 3.1. This image provides an in-memory distributed database so that developers can quickly access large amounts of data in a hybrid environment.

However, there are [significant differences](#) in supported configurations and functionality in the JDG for OpenShift image compared to the full release of JBoss Data Grid. Documentation for other JBoss Data Grid functionality not specific to the JDG for OpenShift image can be found in the [JBoss Data Grid documentation on the Red Hat Customer Portal](#).

See the [Red Hat JBoss Data Grid for OpenShift](#) product page on the Red Hat Container Catalog for the current image versions and tags.

CHAPTER 11. KNOWN ISSUES

The following are some fixed issues and current known issues along with any known workarounds:

11.1. ISSUES COMMON TO MULTIPLE IMAGES

11.1.1. Fixed issues

- <https://issues.jboss.org/browse/CLOUD-1956>: [TEMPLATE] Add initialDelaySeconds to liveness probe in templates to avoid restarts
Previously, the [liveness probe](#) checking if the container, the probe was configured for is still running, was started immediately together with its associated container. When an application, deployed on top of the related Red Hat JBoss Middleware for OpenShift image was deployed using some of the available application templates, the liveness probe might detect application's failure to run in cases, when the initialization process of the container did not finish prior the first check of the liveness probe was made. In such cases this behaviour resulted into unnecessary container restarts.

With this update, an intermediary period of time was inserted between container start and launch of the liveness probe checking process. This allows proper start also for containers using more advanced initialization process without continuous restarts.

- <https://issues.jboss.org/browse/CLOUD-2181>: [TEMPLATES] All template VOLUME_CAPACITY defaults should be 1Gi
Previously, the default volume capacity of persistent storage was set to 512 (five hundred twelve) megabits when using application templates to configure also persistent storage together with provisioning of the application using the related Red Hat JBoss Middleware for OpenShift image.

With this update in these deployment scenarios, the default volume capacity of persistent storage was increased to 1 (one) gigabit.

- <https://issues.jboss.org/browse/CLOUD-2201>: [TEMPLATE] Add option to specify memory limit and default to 1Gi for EAP based templates
Previously, when using some of the application templates to deploy an application on top of the related Red Hat JBoss Middleware for OpenShift image it was not possible in the template to specify a memory limit to constraint the amount of memory the container can use. In environments, restricting allowed memory and CPU usage based on specified memory and CPU limits (e.g. in OpenShift Online), the absence of this memory limit setting on the container resulted the default value specific to that environment was used to configure the memory limit for the container. Since the environment might use lower amount of memory by default, than actually needed for proper work of the Red Hat JBoss Middleware for OpenShift image, in cases like this the container incorporating the application failed to start.

With this update, a new, application templates specific **MEMORY_LIMIT** variable was introduced. It sets the amount of memory the container can use to 1 (one) gigabit by default. This allows applications provisioned via application templates to run correctly also within environments placing restrictions on the system resources, the container in question is allowed to use (value of **MEMORY_LIMIT** will be used instead of the default memory limit specific to that environment).

- <https://issues.jboss.org/browse/CLOUD-2208>: [EAP7] Version 1.6-9 container image contains /etc/yum.repos.d/jboss-rhel-ose.repo file

Previously, the EAP for OpenShift and JWS for OpenShift images contained unnecessary definition of an additional yum repository. This caused a potential local rebuild of aforementioned images to end up with a failure.

With this update, the unnecessary yum repository definition was removed, leading to the rebuild of these images now to work correctly.

- <https://issues.jboss.org/browse/CLOUD-2209>: Remove service account fields from templates
Previously, when using application templates to deploy an application on top of the related Red Hat JBoss Middleware for OpenShift image and if such an application required OpenShift [secrets](#) in order to work correctly, it was necessary first to create a respective OpenShift [service account](#) to allow the pod to reference any secret in the project.

With this update, when provisioning an application using secrets via application templates it is not necessary to create corresponding OpenShift service account any more if using [the default](#) OpenShift Container Platform configuration (`limitSecretReferences` field set to value of `false` in `/etc/origin/master/master-config.yml` configuration file).



NOTE

If using the `limitSecretReferences=true` OpenShift Container Platform setting, it is still necessary to create appropriate service account in order the application using secret(s) to be able to reference them properly in the project.

- <https://issues.jboss.org/browse/CLOUD-1784>: [COMMON] Make the accessLog valve configurable
- <https://issues.jboss.org/browse/CLOUD-2054>: [JDV][SSO][JDG65][JDG71] Make accessLog enabled iff `ENABLE_ACCESS_LOG=true`
Previously, logging of access messages to the standard output channel was enabled by default for Red Hat JBoss Middleware for OpenShift images, although this did not correspond to the default configuration of the respective upstream products.

To align the logging configuration of Red Hat JBoss Middleware for OpenShift images with the configuration used in relevant upstream products, with this update, logging of access messages to the standard output channel is now disabled by default. It can be re-enabled by setting `ENABLE_ACCESS_LOG` variable to `true` if desired. Implementation of the access logging facility varies across the Red Hat JBoss Middleware for OpenShift images, and is implemented as follows:

- Using the [Access Log Valve](#) for EAP 6.4 for OpenShift, Decision Server for OpenShift, IPS for OpenShift, JWS for OpenShift, JDG 6.5 for OpenShift, and JDV 6.3 for OpenShift images,
 - Using the `Infinispan RestAccessLoggingHandler` set to `TRACE` level for JDG 7.1 for OpenShift image, and
 - Using the [Undertow AccessLogHandler](#) for EAP 7.0 for OpenShift, EAP 7.1 for OpenShift, and RH-SSO 7.1 for OpenShift images.
- <https://issues.jboss.org/browse/CLOUD-2024>: Enable Remote IP Valve by default
Previously, the [Remote IP Valve](#) was not included by default in the respective configuration file of the following images:
 - Decision Server for OpenShift,

- IPS for OpenShift,
- EAP 6.4 for OpenShift,
- JDG 6.5 for OpenShift, and
- JDV for OpenShift.

Because of this issue, certain request headers data (for example, the content of "X-Forwarded-For" header) were not available to the Access Log Valve.

With this update, the Remote IP Valve is included in the respective configuration file of these images. This allows the remote client IP address and hostname data to be available to the Access Log Valve.

- <https://issues.jboss.org/browse/CLOUD-2020>: [EAP7] no_proxy configuration not properly translated to maven settings
Previously, the S2I build process failed to correctly translate the specified **no_proxy** and **NO_PROXY** settings to corresponding Maven `<nonProxyHosts>` configuration, when comma-separated list of entire domains that can be accessed directly (e.g. ".example1.com,.example2.com") was provided as **no_proxy** or **NO_PROXY** value, respectively. Because of this issue, an attempt to connect to Maven proxy servers was still performed through the configured HTTP proxy, instead of making a direct connection to some of the ensemble hosts.

With this update, the **no_proxy** and **NO_PROXY** values are properly translated to Maven `<nonProxyHosts>` settings even when they contain specification of entire domains, which allows the direct connection to the Maven proxy servers to be made, when necessary.

- <https://issues.jboss.org/browse/CLOUD-1982>: Custom extensions are broken with OCP3.6
Previously, the EAP S2I build process failed to copy modules from the client image to the EAP **modules** directory on OpenShift Container Platform version 3.6. This issue affected S2I functionality of the following images:
 - EAP for OpenShift,
 - JDG 6.5 for OpenShift,
 - JDG 7.1 for OpenShift,
 - Decision Server for OpenShift,
 - IPS for OpenShift,
 - RH-SSO for OpenShift, and
 - JDV for OpenShift.

With this update, the module directories gets injected with group permissions, which allows the EAP S2I build process to properly copy modules on OpenShift Container Platform version 3.6.

11.2. A-MQ FOR OPENSIFT

11.2.1. Known issues

- <https://issues.jboss.org/browse/CLOUD-1952>: Implement message migration on pod shutdown
When a cluster scales down, only messages sent through Message Queues are migrated to other instances in the cluster. Messages sent through Topics will remain in storage until the cluster scales up. Support for migrating messages sent through Virtual Topics will be introduced in a future release.

There is no known workaround.

11.3. DECISION SERVER FOR OPENSIFT AND IPS FOR OPENSIFT

11.3.1. Known issues

- <https://issues.jboss.org/browse/CLOUD-609>: KIE.SERVER.REQUEST queue consumer not transacted
The first request after an A-MQ message broker restart produces no response when using the A-MQ interface to the IPS. The request message is consumed, even though there is an exception.

There is no known workaround.

- <https://issues.jboss.org/browse/CLOUD-610>: KIE server consumes request messages before all KIE containers are deployed
The kie-server starts handling messages from the `KIE . SERVER . REQUEST` queue before the KIE containers are in the `STARTED` state. This issue causes problems with scaling and pod restarts.

There is no known workaround.

- <https://issues.jboss.org/browse/CLOUD-874>: kie-server deployment sometimes fails with A-MQ clients
The IPS pod sometimes fails to start up with A-MQ clients running, during an IPS pod restart.

There is no known workaround.

- <https://issues.jboss.org/browse/CLOUD-875>: A-MQ responses to clients are PERSISTENT and have unlimited TTL, may clog the response queue
The kie-server replies to A-MQ clients are persistent and have unlimited TTL. This causes the response queue to grow if the A-MQ client disconnects, which results in the A-MQ message broker to eventually block any more messages pushed to the response queue.

To get around this issue, purge the response queue manually.

- <https://issues.jboss.org/browse/CLOUD-878>: Empty responses from kie-server for JMS clients during startup
JMS client calls for `UserTaskService` methods can get lost while restarting a pod. The client does not get an error message or notification of failure.

There is no known workaround.

11.4. EAP FOR OPENSIFT

11.4.1. Fixed issues

- <https://issues.jboss.org/browse/CLOUD-2055>: [EAP7] Configure proxy-address-forwarding="true" by default so that `HttpServletRequest.getRemoteHost()` returns client IP address instead of the OpenShift HAProxy router address
Previously, the Undertow HTTP listener was not configured with proxy-address-forwarding attribute enabled in `standalone.xml` configuration of EAP 7.0 for OpenShift and EAP 7.1 for OpenShift images. Because of this setting, calling the `HttpServletRequest.getRemoteHost()` method from an application deployed on top of EAP 7.0 for OpenShift or EAP 7.1 for OpenShift image returned IP address of the HAProxy router, instead of the IP address of the client, the HTTP Servlet's request was made from.

With this update, the proxy-address-forwarding attribute of the Undertow HTTP listener is enabled in `standalone.xml` of EAP 7.0 for OpenShift and EAP 7.1 for OpenShift images, which allows the original IP address of the remote client to be reported.

- <https://issues.jboss.org/browse/CLOUD-1864>: [EAP7] Support clean shutdown on TERM
Previously, the TERM signal handler for EAP 7.0 for OpenShift and EAP 7.1 for OpenShift images was implemented utilizing the `preStop` OpenShift lifecycle hook. This resulted the clean EAP server shutdown to be possible only when EAP 7.0 for OpenShift and EAP 7.1 for OpenShift images were deployed using some of the available application templates.

To support clean EAP server shutdown, regardless of the way the EAP 7.0 for OpenShift and EAP 7.1 for OpenShift images were deployed (both for deployments using the application templates, but also direct deployments specifying an image stream using the CLI), the TERM signal handler was re-implemented to be part of the `openshift-launch` script.

11.4.2. Known issues

- <https://issues.jboss.org/browse/CLOUD-61>: JPA application fails to start when the database is not available
JPA applications fail to deploy in the EAP OpenShift Enterprise 3.0 image if an underlying database instance that the EAP instance relies on is not available at the start of the deployment. The EAP application tries to contact the database for initialization, but because it is not available, the server starts but the application fails to deploy.

There are no known workarounds available at this stage for this issue.

- <https://issues.jboss.org/browse/CLOUD-158>: Continuous HornetQ errors after scale down "Failed to create netty connection"
In the EAP image, an application not using messaging complains about messaging errors related to HornetQ when being scaled.

Since there are no configuration options to disable messaging to work around this issue, simply include the `standalone-openshift.xml` file within the source of the image and remove or alter the following lines related to messaging:

```
Line 18:
<!-- ##MESSAGING_EXTENSION## -->

Line 318:
<!-- ##MESSAGING_SUBSYSTEM## -->
```

- <https://issues.jboss.org/browse/CLOUD-161>: EAP pod serving requests before it joins cluster, some sessions reset after failure

In a distributed web application deployed on an EAP image, a new container starts serving requests before it joins the cluster.

There are no known workarounds available at this stage for this issue.

- <https://issues.jboss.org/browse/CLOUD-159>: Database pool configurations should contain validation SQL setting
In both the EAP and JWS images, when restarting a crashed database instance, the connection pools contain stale connections.

To work around this issue, restart all instances in case of a database failure.

- <https://issues.jboss.org/browse/CLOUD-2155>: [EAP7] EAP7 access log valve logs IP of HAProxy router instead of client IP
Even when the proxy-address-forwarding attribute of the Undertow HTTP listener is enabled in **standalone.xml** of EAP 7.0 for OpenShift and EAP 7.1 for OpenShift images, the IP address of the HAProxy router instead of the IP address of the original remote client, the request was made from, is logged to the standard output channel in the case when the HTTP Servlet request was issued over a secured route with TLS re-encryption termination.

There is no known workaround.

- <https://issues.jboss.org/browse/CLOUD-2254>: [EAP6] XA (eXtended Architecture) transactions ending in heuristic state after scale down, which leads to recovery cycle
XA distributed transactions may end up in a heuristic decision, when scaling down a cluster of pods utilizing the EAP 6.4 for OpenShift image. Upon receiving a scale down request, the EAP 6.4 for OpenShift pod receives a SIGTERM signal and the shutdown starts. This is followed by closing all the connections till the final shutdown of the container occurs. There is an intermediary period of time between the scale down request and the actual shutdown of the container. During this period the transaction recovery mechanism attempts to finish the unfinished transactions. As connection is already down, the JDBC driver throws an `XAException` informing about the situation. But if the JDBC driver returns an incorrect error code in this case, then the transaction manager is not able to determine the outcome of the operation and reports the heuristic state back.



IMPORTANT

Both JDBC driver implementations (the MySQL Connector/J JDBC Java driver and the PostgreSQL JDBC Java driver) that are shipped with EAP 6.4 for OpenShift image:

- `mysql-connector-java-5.1.25-3.e17.noarch.rpm`
- `postgresql-jdbc-9.2.1002-5.e17.noarch.rpm`

are known to throw an `XAException` with such incorrect return codes in these cases. For more information, see the upstream issue-tracking system entries, depending on the database vendor:

- <https://bugs.mysql.com/bug.php?id=72890>
- <https://github.com/pgjdbc/pgjdbc/issues/236>

To work around this issue Red Hat recommends that you perform a manual transaction recovery, as described in the official documentation for the EAP for OpenShift image.

For the MySQL Connector/J JDBC Java driver, use a more recent version of the driver (versions **5.1.28** and newer), which addresses this issue. To use a newer version of the MySQL Connector/J JDBC Java driver (`mysql-connector-java`), together with EAP 6.4 for OpenShift image:

1. In the application source create the `.s2i/environment` file containing definition of the **`CUSTOM_INSTALL_DIRECTORIES`** environment variable. Point this variable to the directory containing the JAR with updated `mysql-connector-java` and definition of the `install.sh` script:

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

```
$ tree extensions/
extensions/
├── mysql
│   ├── install.sh
│   └── mysql-connector-java.jar
```

```
1 directory, 2 files
```

2. Define `install.sh` as follows:

```
#!/bin/sh

INSTALL_DIR="$1"

unlink
/opt/eap/modules/system/layers/openshift/com/mysql/main/mysql-
connector-java.jar
cp -p "${INSTALL_DIR}"/mysql-connector-java.jar
"${JBoss_HOME}"/modules/system/layers/openshift/com/mysql/main/my
sql-connector-java.jar
```

See the [S2I Artifacts](#) section of [the EAP for OpenShift image documentation](#) for more information.

11.5. FIS FOR OPENSIFT

11.5.1. Known issues

- <https://issues.jboss.org/browse/OSFUSE-112>: karaf /deployments/karaf/bin/client CNFE org.apache.sshd.agent.SshAgent Attempting to run the karaf client in the container to locally SSH to the karaf console fails.

Workaround: Adding both `shell` and `ssh` features make the client work. It will log the warning errors in the logs.

```
$ oc exec karaf-shell-1-bb9zu -- /deployments/karaf/bin/client
osgi:list
```

These warnings are logged when trying to use the JBoss Fuse bin/client script to connect to the JBoss Fuse micro-container. This is an unusual case, since the container is supposed to

contain only bundles and features required for a micro-service, and hence does not need to be managed extensively like a traditional JBoss Fuse install. Any changes made using commands in the remote shell will be temporary and not recorded in the micro-service's docker image.

- <https://issues.jboss.org/browse/OSFUSE-190>: cdi-camel-jetty S2I template incorrect default service name, breaking cdi-camel-http

The cdi-camel-http quickstart expects the cdi-camel-jetty service to be named `qs-cdi-camel-jetty`. In the cdi-camel-jetty template however, the service is named `s2i-qs-cdi-camel-jetty` instead by default. This causes the cdi-camel-http to output this error when both are deployed using the S2I with default values.

Workaround: Set the cdi-camel-jetty `SERVICE_NAME` template parameter to `qs-cdi-camel-jetty`.

- <https://issues.jboss.org/browse/OSFUSE-193>: karaf-camel-rest-sql template service name too long

oc process karaf-camel-rest-sql template fails with the following error: The Service "s2i-qs-karaf-camel-rest-sql" is invalid. SUREFIRE-859: metadata.name: invalid value 's2i-qs-karaf-camel-rest-sql', Details: must be a DNS 952 label (at most 24 characters, matching regex `[a-z]([-a-z0-9]*[a-z0-9])?`): e.g. "my-name" deploymentconfig "s2i-quickstart-karaf-camel-rest-sql" created

Workaround: Set `SERVICE_NAME` template parameter to `karaf-camel-rest-sql`.

- <https://issues.jboss.org/browse/OSFUSE-195>: karaf-camel-amq template should have parameter to configure A-MQ service name
The application template for A-MQ deployments uses a suffix for every transport type to distinguish. Hence there should be a configurable parameter for setting the service name as environment parameter `A_MQ_SERVICE_NAME`.

11.6. JDG FOR OPENSIFT

11.6.1. Fixed issues

- <https://issues.jboss.org/browse/CLOUD-1513>: [JDG] Verify that memcached cache exists
Previously, when JDG for OpenShift image was configured with the `memcached` connector (`INFINISPAN_CONNECTORS` contained `memcached` type) and the name of the cache to expose through the `memcached` connector (value of `MEMCACHED_CACHE` variable) was not listed in the `CACHE_NAMES` variable (comma-separated list of cache names to configure) too, the definition of this cache was not properly created, leading the memcached-connector to fail to start.

With this update, the cache listed in `MEMCACHED_CACHE` gets defined correctly even when not present in the `CACHE_NAMES` list, which allows the memcached-connector to start properly.

- <https://issues.jboss.org/browse/CLOUD-1855>: JDG 6.5 did not support multiple roles due to a bug.
This issue is fixed and the JDG 7.1 image supports multiple roles.

11.6.2. Known issues

- <https://issues.jboss.org/browse/CLOUD-1947>: Cluster upgrade from JDG 6.5 to JDG 7.1 is not possible

It is not possible to upgrade from JDG 6.5 to JDG 7.1 without data loss.

There is no known workaround.

- <https://issues.jboss.org/browse/CLOUD-1949>: protobuf indexing is not working due to security changes. The metadata protobuf cache cannot be accessed without authentication. To get around this add the `__schema_manager` role to the users which require access to protobuf cache.
 1. For remote clients, specify the **USERNAME** and **PASSWORD** environment variables and use the same user account to access the caches remotely.
 2. Verify that the `__schema_manager` role is set for the user (by default, only the `admin` and **REST** roles are set for the user). To do this, set the **ADMIN_GROUP** environment variable on the DeploymentConfig, for example:

```
$ oc env dc/datagrid-app ADMIN_GROUP=REST,admin,__schema_manager
```

11.7. JDV FOR OPENSIFT AND JDG FOR OPENSIFT INTEGRATION

11.7.1. Known issues

- <https://issues.jboss.org/browse/CLOUD-1498>: When using JDG as a datasource for JDV, the JDV connector does not reinitialize metadata for distributed caches after the JDG pod has restarted.
To get around this, use replicated caches with the **CACHE_TYPE_DEFAULT=replicated** environment variable when using JDG as a datasource for JDV.
- [BZ#1437212](#) Querying indexed caches of a restarted JDG pod does not return complete entries.
To get around this issue, use non-indexed caches by supplying cache names with the **DATAVIRT_CACHE_NAMES** environment variable.

11.8. JWS FOR OPENSIFT

11.8.1. Fixed issues

- <https://issues.jboss.org/browse/CLOUD-1814>: Enable Remote IP Valve by default
Previously, the [Remote IP Valve](#) was not included by default in the `server.xml` configuration file of the JWS image. Due to this issue, certain request headers data (for example, the content of "X-Forwarded-For" header) were not available to the Access Log Valve.

With this update, the Remote IP Valve is included in the `server.xml` configuration file of the JWS image. This allows the remote client IP address and hostname data to be available to the Access Log Valve.

- <https://issues.jboss.org/browse/CLOUD-1820>: StdoutAccessLogValve should also log X-Forwarded-For header
Previously, the apparent client remote IP address for the request presented by a proxy, was not replaced with the actual IP address of the remote client in the access log file of the JWS container.

With this update, the access log file of the JWS container contains the actual IP address of the remote client, as presented by a proxy via the "X-Forwarded-For" request header.

11.8.2. Known issues

- <https://issues.jboss.org/browse/CLOUD-57>: Tomcat's access log valve logs to file in container instead of stdout

Due to this issue, the logging data is not available for the central logging facility. To work around this issue, use the `oc exec` command to get the contents of the log file.

- <https://issues.jboss.org/browse/CLOUD-153>: `mvn clean` in JWS STI can fail
Cleaning up after a build in JWS STI is not possible, because the Maven command `mvn clean` fails. This is due to Maven not being able to build the object model during startup.

To work around this issue, add Red Hat and JBoss repositories into the *pom.xml* file of the application if the application uses dependencies from there.

- <https://issues.jboss.org/browse/CLOUD-156>: Datasource realm configuration is incorrect for JWS
It is not possible to do correct JNDI lookup for datasources in the current JWS image if an invalid combination of datasource and realm properties is defined. If a datasource is configured in the *context.xml* file and a realm in the *server.xml* file, then the *server.xml* file's `localDataSource` property should be set to `true`.

- <https://issues.jboss.org/browse/CLOUD-159>: Database pool configurations should contain validation SQL setting
In both the EAP and JWS images, when restarting a crashed database instance, the connection pools contain stale connections.

To work around this issue, restart all instances in case of a database failure.

11.9. RH-SSO FOR OPENSIFT

- The following image streams are not supported in Red Hat JBoss SSO 7.1 image release:
 - `jboss-eap64-openshift:1.4`
 - `jboss-eap70-openshift:1.4`
 - `jboss-datavirt63-openshift:1.1`

11.9.1. Fixed issues

- <https://issues.jboss.org/browse/CLOUD-1965>: [SSO7] Support clean shutdown on TERM
Previously, the TERM signal handler for RH-SSO 7.1 for OpenShift image was implemented utilizing the `preStop OpenShift lifecycle hook`. This resulted the clean Red Hat Single Sign-On 7.1 server shutdown to be possible only when RH-SSO 7.1 for OpenShift image was deployed using some of the available application templates.

To support clean Red Hat Single Sign-On 7.1 server shutdown, regardless of the way the RH-SSO 7.1 for OpenShift image was deployed (both for deployments using the application templates, but also direct deployments specifying an image stream using the CLI), the TERM signal handler was re-implemented to be part of the `openshift-launch` script.

11.9.2. Known issues

- <https://issues.jboss.org/browse/CLOUD-1599>: RH-SSO for OpenShift image still uses `keycloak-server.json`
Although, the Red Hat JBoss SSO Server 7.1 has removed the use of the `keycloak-server.json` configuration and moved it to the `standalone.xml` file, the RH-SSO 7.1 image still uses `keycloak-server.json` for configuration.

There is no fix for this issue at this point.

- <https://issues.jboss.org/browse/CLOUD-1646>: RH-SSO for OpenShift 7.1 image self-registration script issue
The RH-SSO 7.1 image supports multiple certificates exposed on different endpoints. Because of this, the Red Hat JBoss JDV for OpenShift image is not able to self-register. This image exits without any message.

There is no known workaround at this point.

- <https://issues.jboss.org/browse/CLOUD-1601>: SAML client public key support for Red Hat JBoss SSO for OpenShift 7.1 image
This release of the RH-SSO 7.1 image supports multiple certificates exposed on different endpoints. Because of this, the Red Hat EAP6 and EAP7 images fail to create proper SAML client adapter configuration. These server images will start, but with errors.

There is no known workaround at this point.

- <https://issues.jboss.org/browse/CLOUD-747>: SSO server starts before DB is available.
Due to the lack of a dependency mechanism in OpenShift, the SSO server starts before the corresponding database is available. The liveness probe in OpenShift eventually re-starts the SSO pod, but this can take up to 3 minutes.
There is no known workaround.
- <https://issues.jboss.org/browse/CLOUD-712>: Keycloak admin client doesn't use SNI extension
Since the KeyCloak admin Java client does not use SNI, it is not possible to use this client from outside of OpenShift using the HTTPS protocol.
- <https://issues.jboss.org/browse/CLOUD-637>: Router node IP address, instead of the end-user client IP address, are logged on login failures
On invalid login requests, the SSO image logs the OpenShift router's node IP address instead of the end-user's client IP. This information is not helpful in debugging log in failure requests.
- <https://issues.jboss.org/browse/CLOUD-674>: Not giving certificates to SAML deployment causes failure
If the following environment variables are not defined, the SSO deployment will fail even if no SAML application is deployed: `SSO_SAML_KEYSTORE`, `SSO_SAML_CERTIFICATE_NAME`, and `SSO_SAML_KEYSTORE_PASSWORD`.

As a workaround, you may pass the HTTPS secret using the `SSO_SAML_KEYSTORE_SECRET` variable.

- <https://issues.jboss.org/browse/CLOUD-799>: Support for EAR
The SSO auto-registration scripts do not work with EAR archives and there is no current workaround for this issue.

