



OpenShift Container Platform 3.3

Using Images

OpenShift Container Platform 3.3 Guide to Using Images

OpenShift Container Platform 3.3 Using Images

OpenShift Container Platform 3.3 Guide to Using Images

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

Use these topics to find out what different S2I (Source-to-Image), database and Docker images are available for OpenShift Container Platform 3.3 users.

Table of Contents

CHAPTER 1. OVERVIEW	5
CHAPTER 2. SOURCE-TO-IMAGE (S2I)	6
2.1. OVERVIEW	6
2.2. .NET CORE	6
2.2.1. Benefits of Using .NET Core	6
2.2.2. Supported Versions	6
2.2.3. Images	6
2.2.4. Configuration	7
2.2.5. Quickly Deploying Applications from .NET Core Source	9
2.2.6. .NET Core Templates	9
2.3. NODE.JS	10
2.3.1. Overview	10
2.3.2. Versions	10
2.3.3. Images	10
2.3.4. Configuration	10
2.3.5. Hot Deploying	11
2.4. PERL	12
2.4.1. Overview	12
2.4.2. Versions	12
2.4.3. Images	12
2.4.4. Configuration	12
2.4.5. Accessing Logs	13
2.4.6. Hot Deploying	13
2.5. PHP	14
2.5.1. Overview	14
2.5.2. Versions	14
2.5.3. Images	14
2.5.4. Configuration	15
2.5.4.1. Apache Configuration	16
2.5.5. Accessing Logs	17
2.5.6. Hot Deploying	17
2.6. PYTHON	17
2.6.1. Overview	17
2.6.2. Versions	17
2.6.3. Images	17
2.6.4. Configuration	18
2.6.5. Hot Deploying	19
2.7. RUBY	20
2.7.1. Overview	20
2.7.2. Versions	20
2.7.3. Images	20
2.7.4. Configuration	21
2.7.5. Hot Deploying	22
CHAPTER 3. DATABASE IMAGES	24
3.1. OVERVIEW	24
3.2. MYSQL	24
3.2.1. Overview	24
3.2.2. Versions	24
3.2.3. Images	24
3.2.4. Configuration and Usage	25

3.2.4.1. Initializing the Database	25
3.2.4.2. Running MySQL Commands in Containers	25
3.2.4.3. Environment Variables	26
3.2.4.4. Volume Mount Points	28
3.2.4.5. Changing Passwords	28
3.2.5. Creating a Database Service from a Template	30
3.2.6. Using MySQL Replication	30
3.2.6.1. Creating the Deployment Configuration for the MySQL Master	31
3.2.6.2. Creating a Headless Service	33
3.2.6.3. Scaling the MySQL Slaves	34
3.2.7. Troubleshooting	34
3.2.7.1. Linux Native AIO Failure	34
3.3. POSTGRESQL	35
3.3.1. Overview	35
3.3.2. Versions	35
3.3.3. Images	35
3.3.4. Configuration and Usage	36
3.3.4.1. Initializing the Database	36
3.3.4.2. Running PostgreSQL Commands in Containers	36
3.3.4.3. Environment Variables	37
3.3.4.4. Volume Mount Points	38
3.3.4.5. Changing Passwords	38
3.3.5. Creating a Database Service from a Template	40
3.4. MONGODB	40
3.4.1. Overview	40
3.4.2. Versions	40
3.4.3. Images	40
3.4.4. Configuration and Usage	41
3.4.4.1. Initializing the Database	41
3.4.4.2. Running MongoDB Commands in Containers	41
3.4.4.3. Environment Variables	42
3.4.4.4. Volume Mount Points	43
3.4.4.5. Changing Passwords	43
3.4.5. Creating a Database Service from a Template	44
3.4.6. Using MongoDB Replication	45
3.4.6.1. Creating the Deployment Configuration	45
3.4.6.2. Creating the Service Pod	47
3.4.6.3. Creating a Headless Service	48
3.4.6.4. Creating the Final Replication Set	48
3.4.6.5. Scaling the MongoDB Replication Set	49
3.5. MARIADB	49
3.5.1. Overview	49
3.5.2. Versions	49
3.5.3. Images	49
3.5.4. Configuration and Usage	50
3.5.4.1. Initializing the Database	50
3.5.4.2. Running MariaDB Commands in Containers	50
3.5.4.3. Environment Variables	51
3.5.4.4. Volume Mount Points	53
3.5.4.5. Changing Passwords	53
3.5.5. Creating a Database Service from a Template	54
3.5.6. Using MariaDB Replication	55
3.5.6.1. Creating the Deployment Configuration	56

3.5.6.1.1. Creating the MySQL Slaves	57
3.5.6.2. Creating a Headless Service	58
3.5.6.3. Scaling the MariaDB Slaves	59
3.5.7. Troubleshooting	59
3.5.7.1. Linux Native AIO Failure	59
CHAPTER 4. DOCKER IMAGES	61
4.1. OVERVIEW	61
CHAPTER 5. OTHER IMAGES	62
5.1. OVERVIEW	62
5.2. JENKINS	62
5.2.1. Overview	62
5.2.2. Versions	62
5.2.3. Images	62
5.2.4. Configuration and Usage	63
5.2.4.1. Initializing Jenkins	63
5.2.4.1.1. OpenShift Container Platform OAuth authentication	63
5.2.4.1.2. Jenkins Standard Authentication	64
5.2.4.2. Environment Variables	64
5.2.4.3. Volume Mount Points	64
5.2.5. Creating a Jenkins Service from a Template	64

CHAPTER 1. OVERVIEW

Use these topics to discover the different [S2I \(Source-to-Image\)](#), database, and other container images that are available for OpenShift Container Platform users.

Red Hat's official container images are provided in the Red Hat Registry at registry.access.redhat.com. OpenShift Container Platform's supported S2I, database, and Jenkins images are provided in the [openshift3 repository](#) in the Red Hat Registry. For example, **registry.access.redhat.com/openshift3/nodejs-010-rhel7** for the Node.js image.

The xPaaS middleware images are provided in their respective product repositories on the Red Hat Registry, but suffixed with a **-openshift**. For example, **registry.access.redhat.com/jboss-eap-6/eap64-openshift** for the JBoss EAP image.

CHAPTER 2. SOURCE-TO-IMAGE (S2I)

2.1. OVERVIEW

This topic group includes information on the different [S2I \(Source-to-Image\)](#) supported images available for OpenShift Container Platform users.

2.2. .NET CORE

2.2.1. Benefits of Using .NET Core

[.NET Core](#) is a general purpose development platform featuring automatic memory management and modern programming languages. It allows users to build high quality applications efficiently. .NET Core is available on Red Hat Enterprise Linux (RHEL 7) and OpenShift Container Platform via certified containers. .NET Core offers:

- The ability to follow a microservices-based approach, where some components are built with .NET and others with Java, but all can run on a common, supported platform in Red Hat Enterprise Linux and OpenShift Container Platform.
- The capacity to more easily develop new .NET Core workloads on Windows; customers are able to deploy and run on either Red Hat Enterprise Linux or Windows Server.
- A heterogeneous data center, where the underlying infrastructure is capable of running .NET applications without having to rely solely on Windows Server.
- Access to many of the popular development frameworks such as .NET, Java, Ruby, and Python from within OpenShift Container Platform.

2.2.2. Supported Versions

- .NET Core version 1.0
- .NET Core version 1.1
- Supported on Red Hat Enterprise Linux (RHEL) 7 and OpenShift Container Platform versions 3.3 and later

2.2.3. Images

Image stream definitions for the .NET Core on RHEL S2I image are now added during OpenShift Container Platform installations.

The RHEL 7 images are available through Red Hat's subscription registry using:

```
$ docker pull registry.access.redhat.com/dotnet/dotnetcore-10-rhel7
$ docker pull registry.access.redhat.com/dotnet/dotnetcore-11-rhel7
```

To use these images, you can either access them directly from these [image registries](#), or push them into your [OpenShift Container Platform Docker registry](#). Additionally, you can create an [image stream](#) that points to the image, either in your Docker registry or at the

external location. Your OpenShift Container Platform resources can then reference the [image stream definition](#).

2.2.4. Configuration

The .NET Core images support a number of environment variables, which you can set to control the build behavior of your .NET Core application.



NOTE

You must set environment variables that control build behavior in the S2I build configuration or in the **.s2i/environment** file in order to make them available to the build steps.

Table 2.1. NET Core Environment Variables

Variable Name	Description	Default
DOTNET_STARTUP_PROJECT	Used to select the project to run. This must be the folder in the source repository containing project.json .	.
DOTNET_PUBLISH	Used to control whether the application should be built by executing dotnet build or dotnet publish . To publish the application, set the value to true . It is recommended to publish your application.	For backwards compatibility, the default is false . In the next major release, this variable will be removed and the builder will always publish the application.

Variable Name	Description	Default
DOTNET_ASSEMBLY_NAME	<p>Used to select the assembly to run. This must <i>not</i> include the .dll extension. Set this to the output assembly name specified in project.json (name, buildOptions/outputName). For project.json, the assembly name defaults to the project.json parent folder. When project.json is at the context-dir, the parent folder name will be src. So, by default, this generates a src.dll assembly. Setting DOTNET_ASSEMBLY_NAME will cause:</p> <ul style="list-style-type: none"> • The assembly to be <DOTNET_ASSEMBLY_NAME>.dll • The application sources to be in subfolder DOTNET_ASSEMBLY_NAME in the deployed container. 	The name of the DOTNET_STARTUP_PROJECT folder.
DOTNET_RESTORE_SOURCES	Used to specify the space-separated list of NuGet package sources used during the restore operation. This overrides all of the sources specified in the NuGet.config file.	Unset
DOTNET_NPM_TOOLS	Used to specify a list of NPM packages to install before building the application.	Unset
DOTNET_TEST_PROJECTS	Used to specify the space-separated list of test projects to run. This must be folders containing project.json . dotnet test is invoked for each folder.	Unset
DOTNET_CONFIGURATION	Used to run the application in Debug or Release mode. This value should be either Release or Debug .	Release

Variable Name	Description	Default
ASPNETCORE_URLS	This variable is set to http://*:8080 to configure ASP.NET Core to use the port exposed by the image. It is <i>not</i> recommended to change this.	http://*:8080

2.2.5. Quickly Deploying Applications from .NET Core Source



IMPORTANT

The [.NET image stream](#) must first be installed. If you ran a standard installation, the image stream will be present.

An image can be used to build an application by running **oc new-app** against a sample repository:

```
$ oc new-app registry.access.redhat.com/dotnet/dotnetcore-10-
rhel7~https://github.com/redhat-developer/s2i-dotnetcore-ex#dotnetcore-1.0
--context-dir=app
$ oc new-app registry.access.redhat.com/dotnet/dotnetcore-11-
rhel7~https://github.com/redhat-developer/s2i-dotnetcore-ex#dotnetcore-1.1
--context-dir=app
```



NOTE

The **oc new-app** command can detect .NET Core source starting in OpenShift Container Platform 3.3.

2.2.6. .NET Core Templates



IMPORTANT

The [.NET image templates](#) and the .NET images streams must first be [installed](#). If you ran a standard installation, the templates and image streams will be present. This can be checked with:

```
$ (oc get -n openshift templates; oc get -n openshift is) |
grep dotnet
```

OpenShift Container Platform includes templates for the .NET Core images to help easily deploy a sample application.

The [.NET Core sample application](#) running on **dotnet/dotnetcore-10-rhel7** can be deployed with:

```
$ oc new-app --template dotnet-example
```

The [.NET Core sample application](#) running on **dotnet/dotnetcore-11-rhel7** can be deployed with:

```
$ oc new-app --template dotnet-example -p  
DOTNET_IMAGE_STREAM_TAG=dotnet:1.1 -p SOURCE_REPOSITORY_REF=dotnetcore-1.1
```

The [.NET Core MusicStore application](#) using PostgreSQL as database can be deployed with:

```
$ oc new-app --template=dotnet-pgsql-persistent
```

2.3. NODE.JS

2.3.1. Overview

OpenShift Container Platform provides [S2I](#) enabled Node.js images for building and running Node.js applications. The Node.js S2I builder image assembles your application source with any required dependencies to create a new image containing your Node.js application. This resulting image can be run either by OpenShift Container Platform or by Docker.

2.3.2. Versions

Currently, OpenShift Container Platform provides versions [0.10](#) and [4](#) of Node.js.

2.3.3. Images

These images come in two flavors, depending on your needs:

- RHEL 7
- CentOS 7

RHEL 7 Based Images

The RHEL 7 images are available through the Red Hat Registry:

```
$ docker pull registry.access.redhat.com/openshift3/nodejs-010-rhel7  
$ docker pull registry.access.redhat.com/rhsc1/nodejs-4-rhel7
```

CentOS 7 Based Image

This image is available on Docker Hub:

```
$ docker pull openshift/nodejs-010-centos7
```

To use these images, you can either access them directly from these [image registries](#), or push them into your [OpenShift Container Platform Docker registry](#). Additionally, you can create an [image stream](#) that points to the image, either in your Docker registry or at the external location. Your OpenShift Container Platform resources can then reference the ImageStream. You can find [example image stream definitions](#) for all the provided OpenShift Container Platform images.

2.3.4. Configuration

The Node.js image supports a number of environment variables, which can be set to control the configuration and behavior of the Node.js runtime.

To set these environment variables as part of your image, you can place them into a [.s2i/environment](#) file inside your source code repository, or define them in the [environment](#) section of the build configuration's **sourceStrategy** definition.

You can also set environment variables to be used with an existing image when [creating new applications](#), or by [updating environment variables for existing objects](#) such as deployment configurations.



NOTE

Environment variables that control build behavior must be set as part of the s2i build configuration or in the **.s2i/environment** file to make them available to the build steps.

Table 2.2. Development Mode Environment Variables

Variable name	Description
DEV_MODE	When set to true , enables hot deploy and opens the debug port. Additionally, indicates to tooling that the image is in development mode. Default is false .
DEBUG_PORT	The debug port. Only valid if DEV_MODE is set to true. Default is 5858.
NPM_MIRROR	The custom NPM registry mirror URL. All NPM packages will be downloaded from the mirror link during the build process.

2.3.5. Hot Deploying

Hot deployment allows you to quickly make and deploy changes to your application without having to generate a new S2I build. In order to immediately pick up changes made in your application source code, you must run your built image with the **DEV_MODE=true** environment variable.

You can set new environment variables when [creating new applications](#), or [updating environment variables for existing objects](#).



WARNING

Only use the **DEV_MODE=true** environment variable while developing or debugging. Using this in your production environment is not recommended.

To change the source code of a running pod, [open a remote shell into the container](#):

```
$ oc rsh <pod_id>
```

Entering into a running container changes your current directory to ***/opt/app-root/src***, where the source code is located.

2.4. PERL

2.4.1. Overview

OpenShift Container Platform provides [S2I](#) enabled Perl images for building and running Perl applications. The Perl S2I builder image assembles your application source with any required dependencies to create a new image containing your Perl application. This resulting image can be run either by OpenShift Container Platform or by Docker.

2.4.2. Versions

Currently, OpenShift Container Platform supports versions [5.16](#), [5.20](#), and [5.24](#) of Perl.

2.4.3. Images

Images comes in two flavors, depending on your needs:

- RHEL 7
- CentOS 7

RHEL 7 Based Images

The RHEL 7 images are available through the Red Hat Registry:

```
$ docker pull registry.access.redhat.com/openshift3/perl-516-rhel7
$ docker pull registry.access.redhat.com/rhsc1/perl-520-rhel7
$ docker pull registry.access.redhat.com/rhsc1/perl-524-rhel7
```

CentOS 7 Based Image

A CentOS image for Perl 5.16 is available on Docker Hub:

```
$ docker pull openshift/perl-516-centos7
```


To use these images, you can either access them directly from these [image registries](#) or push them into your [OpenShift Container Platform Docker registry](#). Additionally, you can create an [image stream](#) that points to the image, either in your Docker registry or at the external location. Your OpenShift Container Platform resources can then reference the ImageStream. You can find [example image stream definitions](#) for all the provided OpenShift Container Platform images.

2.4.4. Configuration

The Perl image supports a number of environment variables which can be set to control the configuration and behavior of the Perl runtime.

To set these environment variables as part of your image, you can place them into a [.s2i/environment](#) file inside your source code repository, or define them in the [environment](#) section of the build configuration's **sourceStrategy** definition.

You can also set environment variables to be used with an existing image when [creating new applications](#), or by [updating environment variables for existing objects](#) such as deployment configurations.



NOTE

Environment variables that control build behavior must be set as part of the s2i build configuration or in the **.s2i/environment** file to make them available to the build steps.

Table 2.3. Perl Environment Variables

Variable name	Description
ENABLE_CPAN_TEST	When set to true , this variable installs all the cpan modules and runs their tests. By default, the testing of the modules is turned off.
CPAN_MIRROR	This variable specifies a mirror URL which cpanminus uses to install dependencies. By default, this URL is not specified.
PERL_APACHE2_RELOAD	Set this to true to enable automatic reloading of modified Perl modules. By default, automatic reloading is turned off.
HTTPD_START_SERVERS	The StartServers directive sets the number of child server processes created on startup. Default is 8.
HTTPD_MAX_REQUEST_WORKERS	Number of simultaneous requests that will be handled by Apache. The default is 256, but it will be automatically lowered if memory is limited.

2.4.5. Accessing Logs

Access logs are streamed to standard output and as such they can be viewed using the **oc logs** command. Error logs are stored in the **/tmp/error_log** file, which can be viewed using the **oc rsh** command to access the container.

2.4.6. Hot Deploying

Hot deployment allows you to quickly make and deploy changes to your application without having to generate a new S2I build. To enable hot deployment in this image, you must set the **PERL_APACHE2_RELOAD** environment variable to **true**. For example, see the **oc new-app** command. You can use the **oc set env** command to update environment variables of existing objects.

**WARNING**

You should only use this option while developing or debugging; it is not recommended to turn this on in your production environment.

To change your source code in a running pod, use the **oc rsh** command to enter the container:

```
$ oc rsh <pod_id>
```

After you enter into the running container, your current directory is set to **/opt/app-root/src**, where the source code is located.

2.5. PHP

2.5.1. Overview

OpenShift Container Platform provides [S2I](#) enabled PHP images for building and running PHP applications. The PHP S2I builder image assembles your application source with any required dependencies to create a new image containing your PHP application. This resulting image can be run either by OpenShift Container Platform or by Docker.

2.5.2. Versions

Currently, OpenShift Container Platform provides versions [5.5](#), [5.6](#), and [7.0](#) of PHP.

2.5.3. Images

These images come in two flavors, depending on your needs:

- RHEL 7
- CentOS 7

RHEL 7 Based Images

The RHEL 7 images are available through the Red Hat Registry:

```
$ docker pull registry.access.redhat.com/openshift3/php-55-rhel7
$ docker pull registry.access.redhat.com/rhsccl/php-56-rhel7
$ docker pull registry.access.redhat.com/rhsccl/php-70-rhel7
```

CentOS 7 Based Images

CentOS images for PHP 5.5 and 5.6 are available on Docker Hub:

```
$ docker pull openshift/php-55-centos7
$ docker pull openshift/php-56-centos7
```

To use these images, you can either access them directly from these [image registries](#) or push them into your [OpenShift Container Platform Docker registry](#). Additionally, you can create an [image stream](#) that points to the image, either in your Docker registry or at the external location. Your OpenShift Container Platform resources can then reference the image stream.


You can find [example image stream definitions](#) for all the provided OpenShift Container Platform images.

2.5.4. Configuration

The PHP image supports a number of environment variables which can be set to control the configuration and behavior of the PHP runtime.

To set these environment variables as part of your image, you can place them into a [.s2i/environment](#) file inside your source code repository, or define them in the [environment section](#) of the build configuration's **sourceStrategy** definition.

You can also set environment variables to be used with an existing image when [creating new applications](#), or by [updating environment variables for existing objects](#) such as deployment configurations.



NOTE

Environment variables that control build behavior must be set as part of the s2i build configuration or in the **.s2i/environment** file to make them available to the build steps.

The following environment variables set their equivalent property value in the **php.ini** file:

Table 2.4. PHP Environment Variables

Variable Name	Description	Default
ERROR_REPORTING	Informs PHP of the errors, warnings, and notices for which you would like it to take action.	E_ALL & ~E_NOTICE
DISPLAY_ERRORS	Controls if and where PHP outputs errors, notices, and warnings.	ON
DISPLAY_STARTUP_ERRORS	Causes any display errors that occur during PHP's startup sequence to be handled separately from display errors.	OFF
TRACK_ERRORS	Stores the last error/warning message in \$php_errormsg (boolean).	OFF
HTML_ERRORS	Links errors to documentation that is related to the error.	ON

Variable Name	Description	Default
INCLUDE_PATH	Path for PHP source files.	<i>./opt/openshift/src:/opt/rh/php55/root/usr/share/pear</i>
SESSION_PATH	Location for session data files.	<i>/tmp/sessions</i>
DOCUMENTROOT	Path that defines the document root for your application (for example, <i>/public</i>).	<i>/</i>

The following environment variable sets its equivalent property value in the ***opcache.ini*** file:

Table 2.5. Additional PHP settings

Variable Name	Description	Default
OPCACHE_MEMORY_CONSUMPTION	The OPcache shared memory storage size.	16M
OPCACHE_REVALIDATE_FREQ	How often to check script time stamps for updates, in seconds. 0 results in OPcache checking for updates on every request.	2

You can also override the entire directory used to load the PHP configuration by setting:

Table 2.6. Additional PHP settings

Variable Name	Description
PHPRC	Sets the path to the <i>php.ini</i> file.
PHP_INI_SCAN_DIR	Path to scan for additional <i>.ini</i> configuration files

You can use a custom composer repository mirror URL to download packages instead of the default 'packagist.org':

Table 2.7. Composer Environment Variables

Variable Name	Description	COMPOSER_MIRROR
---------------	-------------	-----------------

2.5.4.1. Apache Configuration

If the **DocumentRoot** of the application is nested in the source directory ***/opt/openshift/src***, you can provide your own ***.htaccess*** file to override the default Apache behavior and specify how application requests should be handled. The ***.htaccess*** file must be located at the root of the application source.

2.5.5. Accessing Logs

Access logs are streamed to standard out and as such they can be viewed using the **oc logs** command. Error logs are stored in the **/tmp/error_log** file, which can be viewed using the **oc rsh** command to access the container.

2.5.6. Hot Deploying

Hot deployment allows you to quickly make and deploy changes to your application without having to generate a new S2I build. In order to immediately pick up changes made in your application source code, you must run your built image with the **OPCACHE_REVALIDATE_FREQ=0** environment variable.

For example, see the **oc new-app** command. You can use the **oc env** command to update environment variables of existing objects.



WARNING

You should only use this option while developing or debugging; it is not recommended to turn this on in your production environment.

To change your source code in a running pod, use the **oc rsh** command to enter the container:

```
$ oc rsh <pod_id>
```

After you enter into the running container, your current directory is set to **/opt/app-root/src**, where the source code is located.

2.6. PYTHON

2.6.1. Overview

OpenShift Container Platform provides **S2I** enabled Python images for building and running Python applications. The Python S2I builder image assembles your application source with any required dependencies to create a new image containing your Python application. This resulting image can be run either by OpenShift Container Platform or by Docker.

2.6.2. Versions

Currently, OpenShift Container Platform provides versions **2.7**, **3.3**, **3.4**, and **3.5** of Python.

2.6.3. Images

These images come in two flavors, depending on your needs:

- RHEL 7
- CentOS 7

RHEL 7 Based Images

The RHEL 7 images are available through the Red Hat Registry:

```
$ docker pull registry.access.redhat.com/rhsc1/python-27-rhel7
$ docker pull registry.access.redhat.com/openshift3/python-33-rhel7
$ docker pull registry.access.redhat.com/rhsc1/python-34-rhel7
$ docker pull registry.access.redhat.com/rhsc1/python-35-rhel7
```

CentOS 7 Based Images

These images are available on Docker Hub:

```
$ docker pull centos/python-27-centos7
$ docker pull openshift/python-33-centos7
$ docker pull centos/python-34-centos7
$ docker pull centos/python-35-centos7
```

To use these images, you can either access them directly from these [image registries](#) or push them into your [OpenShift Container Platform Docker registry](#). Additionally, you can create an [image stream](#) that points to the image, either in your Docker registry or at the external location. Your OpenShift Container Platform resources can then reference the ImageStream. You can find [example image stream definitions](#) for all the provided OpenShift Container Platform images.

2.6.4. Configuration

The Python image supports a number of environment variables which can be set to control the configuration and behavior of the Python runtime.

To set these environment variables as part of your image, you can place them into a [.s2i/environment](#) file inside your source code repository, or define them in [the environment section](#) of the build configuration's **sourceStrategy** definition.

You can also set environment variables to be used with an existing image when [creating new applications](#), or by [updating environment variables for existing objects](#) such as deployment configurations.



NOTE

Environment variables that control build behavior must be set as part of the s2i build configuration or in the **.s2i/environment** file to make them available to the build steps.

Table 2.8. Python Environment Variables

Variable name	Description
APP_FILE	This variable specifies the file name passed to the python interpreter which is responsible for launching the application. This variable is set to app.py by default.

Variable name	Description
APP_MODULE	This variable specifies the WSGI callable. It follows the pattern <code>\$(MODULE_NAME):\$(VARIABLE_NAME)</code> , where the module name is a full dotted path and the variable name refers to a function inside the specified module. If you use setup.py for installing the application, then the module name can be read from that file and the variable defaults to application . There is an example setup-test-app available.
APP_CONFIG	This variable indicates the path to a valid Python file with a gunicorn configuration .
DISABLE_COLLECTSTATIC	Set it to a nonempty value to inhibit the execution of manage.py collectstatic during the build. Only affects Django projects.
DISABLE_MIGRATE	Set it to a nonempty value to inhibit the execution of manage.py migrate when the produced image is run. Only affects Django projects.
PIP_INDEX_URL	Set this variable to use a custom index URL or mirror to download required packages during build process. This only affects packages listed in the requirements.txt file.
WEB_CONCURRENCY	Set this to change the default setting for the number of workers . By default, this is set to the number of available cores times 4.

2.6.5. Hot Deploying

Hot deployment allows you to quickly make and deploy changes to your application without having to generate a new S2I build. If you are using Django, hot deployment works out of the box.

To enable hot deployment while using Gunicorn, ensure you have a Gunicorn configuration file inside your repository with [the reload option](#) set to **true**. Specify your configuration file using the **APP_CONFIG** environment variable. For example, see the [oc new-app](#) command. You can use the [oc set env](#) command to update environment variables of existing objects.

**WARNING**

You should only use this option while developing or debugging; it is not recommended to turn this on in your production environment.

To change your source code in a running pod, use the **oc rsh** command to enter the container:

```
$ oc rsh <pod_id>
```

After you enter into the running container, your current directory is set to **/opt/app-root/src**, where the source code is located.

2.7. RUBY

2.7.1. Overview

OpenShift Container Platform provides [S2I](#) enabled Ruby images for building and running Ruby applications. The Ruby S2I builder image assembles your application source with any required dependencies to create a new image containing your Ruby application. This resulting image can be run either by OpenShift Container Platform or by Docker.

2.7.2. Versions

Currently, OpenShift Container Platform provides versions [2.0](#), [2.2](#), and [2.3](#) of Ruby.

2.7.3. Images

These images come in two flavors, depending on your needs:

- RHEL 7
- CentOS 7

RHEL 7 Based Images

The RHEL 7 images are available through the Red Hat registry:

```
$ docker pull registry.access.redhat.com/openshift3/ruby-20-rhel7
$ docker pull registry.access.redhat.com/rhscl/ruby-22-rhel7
$ docker pull registry.access.redhat.com/rhscl/ruby-23-rhel7
```

CentOS 7 Based Images

These images are available on Docker Hub:

```
$ docker pull openshift/ruby-20-centos7
$ docker pull openshift/ruby-22-centos7
$ docker pull centos/ruby-23-centos7
```



To use these images, you can either access them directly from these [image registries](#) or push them into your [OpenShift Container Platform Docker registry](#). Additionally, you can create an [image stream](#) that points to the image, either in your Docker registry or at the external location. Your OpenShift Container Platform resources can then reference the ImageStream. You can find [example image stream definitions](#) for all the provided OpenShift Container Platform images.

2.7.4. Configuration

The Ruby image supports a number of environment variables which can be set to control the configuration and behavior of the Ruby runtime.

To set these environment variables as part of your image, you can place them into a [.s2i/environment](#) file inside your source code repository, or define them in the [environment](#) section of the build configuration's **sourceStrategy** definition.

You can also set environment variables to be used with an existing image when [creating new applications](#), or by [updating environment variables for existing objects](#) such as deployment configurations.



NOTE

Environment variables that control build behavior must be set as part of the s2i build configuration or in the **.s2i/environment** file to make them available to the build steps.

Table 2.9. Ruby Environment Variables

Variable name	Description
RACK_ENV	This variable specifies the environment within which the Ruby application is deployed; for example, production , development , or test . Each level has different behavior in terms of logging verbosity, error pages, and ruby gem installation. The application assets are only compiled if RACK_ENV is set to production ; the default value is production .
RAILS_ENV	This variable specifies the environment within which the Ruby on Rails application is deployed; for example, production , development , or test . Each level has different behavior in terms of logging verbosity, error pages, and ruby gem installation. The application assets are only compiled if RAILS_ENV is set to production . This variable is set to \${RACK_ENV} by default.
DISABLE_ASSET_COMPILATION	When set to true , this variable disables the process of asset compilation. Asset compilation only happens when the application runs in a production environment. Therefore, you can use this variable when assets have already been compiled.

Variable name	Description
PUMA_MIN_THREADS, PUMA_MAX_THREADS	This variable indicates the minimum and maximum number of threads that will be available in Puma's thread pool.
PUMA_WORKERS	This variable indicates the number of worker processes to be launched in Puma's clustered mode (when Puma runs more than two processes). If not explicitly set, the default behavior sets PUMA_WORKERS to a value that is appropriate for the memory available to the container and the number of cores on the host.
RUBYGEM_MIRROR	Set this variable to use a custom RubyGems mirror URL to download required gem packages during the build process. Note: This environment variable is only available for Ruby 2.2+ images.

2.7.5. Hot Deploying

Hot deployment allows you to quickly make and deploy changes to your application without having to generate a new S2I build. The method for enabling hot deployment in this image differs based on the application type.

Ruby on Rails Applications

For Ruby on Rails application, run the built Rails application with the **RAILS_ENV=development** environment variable passed to the running pod. For an existing deployment configuration, you can use [the `oc set env` command](#):

```
$ oc set env dc/rails-app RAILS_ENV=development
```

Other Types of Ruby Applications (Sinatra, Padrino, etc.)

For other types of Ruby applications, your application must be built with a gem that can reload the server every time a change to the source code is made inside the running container. Those gems are:

- [Shotgun](#)
- [Rerun](#)
- [Rack-livereload](#)

In order to be able to run your application in development mode, you must modify the [S2I `run` script](#) so that the web server is launched by the chosen gem, which checks for changes in the source code.

After you build your application image with your version of the [S2I `run` script](#), run the image with the **RACK_ENV=development** environment variable. For example, see the [oc `new-app` command](#). You can use the [oc `set env` command](#) to update environment variables of

existing objects.

**WARNING**

You should only use this option while developing or debugging; it is not recommended to turn this on in your production environment.

To change your source code in a running pod, use the **oc rsh** command to enter the container:

```
$ oc rsh <pod_id>
```

After you enter into the running container, your current directory is set to **/opt/app-root/src**, where the source code is located.

CHAPTER 3. DATABASE IMAGES

3.1. OVERVIEW

This topic group includes information on the different database images available for OpenShift Container Platform users.



NOTE

Enabling clustering for database images is currently in Technology Preview and not intended for production use.

3.2. MYSQL

3.2.1. Overview

OpenShift Container Platform provides a container image for running MySQL. This image can provide database services based on username, password, and database name settings provided via configuration.

3.2.2. Versions

Currently, OpenShift Container Platform provides versions [5.5](#), [5.6](#), and [5.7](#) of MySQL.

3.2.3. Images

This image comes in two flavors, depending on your needs:

- RHEL 7
- CentOS 7

RHEL 7 Based Images

The RHEL 7 image is available through the Red Hat Registry:

```
$ docker pull registry.access.redhat.com/openshift3/mysql-55-rhel7
$ docker pull registry.access.redhat.com/rhsc1/mysql-56-rhel7
$ docker pull registry.access.redhat.com/rhsc1/mysql-57-rhel7
```

CentOS 7 Based Images

CentOS images for MySQL 5.5 and 5.6 are available on Docker Hub:

```
$ docker pull openshift/mysql-55-centos7
$ docker pull openshift/mysql-56-centos7
```

To use these images, you can either access them directly from these registries or push them into your OpenShift Container Platform Docker registry. Additionally, you can create an ImageStream that points to the image, either in your Docker registry or at the external

location. Your OpenShift Container Platform resources can then reference the ImageStream. You can find [example](#) ImageStream definitions for all the provided OpenShift Container Platform images.

3.2.4. Configuration and Usage

3.2.4.1. Initializing the Database

The first time you use the shared volume, the database is created along with the database administrator user and the MySQL root user (if you specify the **MYSQL_ROOT_PASSWORD** environment variable). Afterwards, the MySQL daemon starts up. If you are re-attaching the volume to another container, then the database, database user, and the administrator user are not created, and the MySQL daemon starts.

The following command creates a new database [pod](#) with MySQL running in a container:

```
$ oc new-app \
  -e MYSQL_USER=<username> \
  -e MYSQL_PASSWORD=<password> \
  -e MYSQL_DATABASE=<database_name> \
  registry.access.redhat.com/openshift3/mysql-55-rhel7
```

3.2.4.2. Running MySQL Commands in Containers

OpenShift Container Platform uses [Software Collections](#) (SCLs) to install and launch MySQL. If you want to execute a MySQL command inside of a running container (for debugging), you must invoke it using `bash`.

To do so, first identify the name of the pod. For example, you can view the list of pods in your current project:

```
$ oc get pods
```

Then, open a remote shell session to the pod:

```
$ oc rsh <pod>
```

When you enter the container, the required SCL is automatically enabled.

You can now run the **mysql** command from the bash shell to start a MySQL interactive session and perform normal MySQL operations. For example, to authenticate as the database user:

```
bash-4.2$ mysql -u $MYSQL_USER -p$MYSQL_PASSWORD -h $HOSTNAME
$MYSQL_DATABASE
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 4
Server version: 5.5.37 MySQL Community Server (GPL)
...
mysql>
```

When you are finished, enter **quit** or **exit** to leave the MySQL session.

3.2.4.3. Environment Variables

The MySQL user name, password, and database name must be configured with the following environment variables:

Table 3.1. MySQL Environment Variables

Variable Name	Description
MYSQL_USER	Specifies the user name for the database user that is created for use by your application.
MYSQL_PASSWORD	Password for the MYSQL_USER .
MYSQL_DATABASE	Name of the database to which MYSQL_USER has full rights.
MYSQL_ROOT_PASSWORD	Optional password for the root user. If this is not set, then remote login to the root account is not possible. Local connections from within the container are always permitted without a password.
MYSQL_SERVICE_HOST	Service host variable automatically created by Kubernetes.
MYSQL_SERVICE_PORT	Service port variable automatically created by Kubernetes.



WARNING

You must specify the user name, password, and database name. If you do not specify all three, the pod will fail to start and OpenShift Container Platform will continuously try to restart it.

MySQL settings can be configured with the following environment variables:

Table 3.2. Additional MySQL Settings

Variable Name	Description	Default
MYSQL_LOWER_CASE_TABLE_NAMES	Sets how the table names are stored and compared.	0
MYSQL_MAX_CONNECTIONS	The maximum permitted number of simultaneous client connections.	151

Variable Name	Description	Default
MYSQL_MAX_ALLOWED_PACKET	The maximum size of one packet or any generated/intermediate string.	200M
MYSQL_FT_MIN_WORD_LEN	The minimum length of the word to be included in a FULLTEXT index.	4
MYSQL_FT_MAX_WORD_LEN	The maximum length of the word to be included in a FULLTEXT index.	20
MYSQL_AIO	Controls the innodb_use_native_aio setting value if the native AIO is broken.	1
MYSQL_TABLE_OPEN_CACHE	The number of open tables for all threads.	400
MYSQL_KEY_BUFFER_SIZE	The size of the buffer used for index blocks.	32M (or 10% of available memory)
MYSQL_SORT_BUFFER_SIZE	The size of the buffer used for sorting.	256K
MYSQL_READ_BUFFER_SIZE	The size of the buffer used for a sequential scan.	8M (or 5% of available memory)
MYSQL_INNODB_BUFFER_POOL_SIZE	The size of the buffer pool where InnoDB caches table and index data.	32M (or 50% of available memory)

Variable Name	Description	Default
MYSQL_INNODB_LOG_FILE_SIZE	The size of each log file in a log group.	8M (or 15% of available memory)
MYSQL_INNODB_LOG_BUFFER_SIZE	The size of the buffer that InnoDB uses to write to the log files on disk.	8M (or 15% of available memory)

Some of the memory-related parameters have two default values. The fixed value is used when a container does not have [memory limits](#) assigned. The other value is calculated dynamically during a container's startup based on available memory.

3.2.4.4. Volume Mount Points

The MySQL image can be run with mounted volumes to enable persistent storage for the database:

- **`/var/lib/mysql/data`** - This is the data directory where MySQL stores database files.

3.2.4.5. Changing Passwords

Passwords are part of the image configuration, therefore the only supported method to change passwords for the database user (**MYSQL_USER**) and **root** user is by changing the environment variables **MYSQL_PASSWORD** and **MYSQL_ROOT_PASSWORD**, respectively.

You can view the current passwords by viewing the pod or deployment configuration in the web console or by listing the environment variables with the CLI:

```
$ oc set env pod <pod_name> --list
```

Whenever **MYSQL_ROOT_PASSWORD** is set, it enables remote access for the **root** user with the given password, and whenever it is unset, remote access for the **root** user is disabled. This does not affect the regular user **MYSQL_USER**, who always has remote access. This also does not affect local access by the **root** user, who can always log in without a password in **localhost**.

Changing database passwords through SQL statements or any way other than through the environment variables aforementioned causes a mismatch between the values stored in the variables and the actual passwords. Whenever a database container starts, it resets the passwords to the values stored in the environment variables.

To change these passwords, update one or both of the desired environment variables for

the related deployment configuration(s) using the **oc set env** command. If multiple deployment configurations utilize these environment variables, for example in the case of an application created from a template, you must update the variables on each deployment configuration so that the passwords are in sync everywhere. This can be done all in the same command:

```
$ oc set env dc <dc_name> [<dc_name_2> ...] \
  MYSQL_PASSWORD=<new_password> \
  MYSQL_ROOT_PASSWORD=<new_root_password>
```

IMPORTANT

Depending on your application, there may be other environment variables for passwords in other parts of the application that should also be updated to match. For example, there could be a more generic **DATABASE_USER** variable in a front-end pod that should match the database user's password. Ensure that passwords are in sync for all required environment variables per your application, otherwise your pods may fail to redeploy when triggered.

Updating the environment variables triggers the redeployment of the database server if you have a [configuration change trigger](#). Otherwise, you must manually start a new deployment in order to apply the password changes.

To verify that new passwords are in effect, first open a remote shell session to the running MySQL pod:

```
$ oc rsh <pod>
```

From the bash shell, verify the database user's new password:

```
bash-4.2$ mysql -u $MYSQL_USER -p<new_password> -h $HOSTNAME
$MYSQL_DATABASE -te "SELECT * FROM (SELECT database()) db CROSS JOIN
(SELECT user()) u"
```

If the password was changed correctly, you should see a table like this:

```
+-----+-----+
| database() | user() |
+-----+-----+
| sampled_b | user0PG@172.17.42.1 |
+-----+-----+
```

To verify the **root** user's new password:

```
bash-4.2$ mysql -u root -p<new_root_password> -h $HOSTNAME $MYSQL_DATABASE
-te "SELECT * FROM (SELECT database()) db CROSS JOIN (SELECT user()) u"
```

If the password was changed correctly, you should see a table like this:

```
+-----+-----+
| database() | user() |
+-----+-----+
```

```
| sampled | root@172.17.42.1 |
+-----+-----+
```

3.2.5. Creating a Database Service from a Template

OpenShift Container Platform provides a [template](#) to make creating a new database service easy. The template provides parameter fields to define all the mandatory environment variables (user, password, database name, etc) with predefined defaults including auto-generation of password values. It will also define both a [deployment configuration](#) and a [service](#).

The MySQL templates should have been registered in the default **openshift** project by your cluster administrator during the initial cluster setup. See [Loading the Default Image Streams and Templates](#) for more details, if required.

There are two templates available:

- **mysql-ephemeral** is for development or testing purposes only because it uses ephemeral storage for the database content. This means that if the database pod is restarted for any reason, such as the pod being moved to another node or the deployment configuration being updated and triggering a redeploy, all data will be lost.
- **mysql-persistent** uses a persistent volume store for the database data which means the data will survive a pod restart. Using persistent volumes requires a persistent volume pool be defined in the OpenShift Container Platform deployment. Cluster administrator instructions for setting up the pool are located [here](#).

You can find instructions for instantiating templates by following these [instructions](#).

Once you have instantiated the service, you can copy the user name, password, and database name environment variables into a deployment configuration for another component that intends to access the database. That component can then access the database via the service that was defined.

3.2.6. Using MySQL Replication



NOTE

Enabling clustering for database images is currently in Technology Preview and not intended for production use.

Red Hat provides a proof-of-concept [template](#) for MySQL master-slave replication (clustering); you can obtain the [example template from GitHub](#).

To upload the example template into the current project's template library:

```
$ oc create -f \
https://raw.githubusercontent.com/openshift/mysql/master/5.5/examples/replica/mysql_replica.json
```

The following sections detail the objects defined in the example template and describe how they work together to start a cluster of MySQL servers implementing master-slave replication. This is the recommended replication strategy for MySQL.

3.2.6.1. Creating the Deployment Configuration for the MySQL Master

To set up MySQL replication, a [deployment configuration](#) is defined in the example template that defines a [replication controller](#). For MySQL master-slave replication, two deployment configurations are needed. One deployment configuration defines the MySQL *master* server and second the MySQL *slave* servers.

To tell a MySQL server to act as the master, the **command** field in the container's definition in the deployment configuration must be set to **run-mysqld-master**. This script acts as an alternative entrypoint for the MySQL image and configures the MySQL server to run as the master in replication.

MySQL replication requires a special user that relays data between the master and slaves. The following environment variables are defined in the template for this purpose:

Variable Name	Description	Default
MYSQL_MASTER_USER	The user name of the replication user	master
MYSQL_MASTER_PASSWORD	The password for the replication user	generated

Example 3.1. MySQL Master Deployment Configuration Object Definition in the Example Template

```
kind: "DeploymentConfig"
apiVersion: "v1"
metadata:
  name: "mysql-master"
spec:
  strategy:
    type: "Recreate"
  triggers:
    - type: "ConfigChange"
  replicas: 1
  selector:
    name: "mysql-master"
  template:
    metadata:
      labels:
        name: "mysql-master"
    spec:
      volumes:
        - name: "mysql-master-data"
          persistentVolumeClaim:
            claimName: "mysql-master"
      containers:
```

```

- name: "server"
  image: "openshift/mysql-55-centos7"
  command:
    - "run-mysqld-master"
  ports:
    - containerPort: 3306
      protocol: "TCP"
  env:
    - name: "MYSQL_MASTER_USER"
      value: "${MYSQL_MASTER_USER}"
    - name: "MYSQL_MASTER_PASSWORD"
      value: "${MYSQL_MASTER_PASSWORD}"
    - name: "MYSQL_USER"
      value: "${MYSQL_USER}"
    - name: "MYSQL_PASSWORD"
      value: "${MYSQL_PASSWORD}"
    - name: "MYSQL_DATABASE"
      value: "${MYSQL_DATABASE}"
    - name: "MYSQL_ROOT_PASSWORD"
      value: "${MYSQL_ROOT_PASSWORD}"
  volumeMounts:
    - name: "mysql-master-data"
      mountPath: "/var/lib/mysql/data"
  resources: {}
  terminationMessagePath: "/dev/termination-log"
  imagePullPolicy: "IfNotPresent"
  securityContext:
    capabilities: {}
    privileged: false
  restartPolicy: "Always"
  dnsPolicy: "ClusterFirst"

```

Since we claimed a persistent volume in this deployment configuration to have all data persisted for the MySQL master server, you must ask your cluster administrator to create a persistent volume that you can claim the storage from.

After the deployment configuration is created and the pod with MySQL master server is started, it will create the database defined by **MYSQL_DATABASE** and configure the server to replicate this database to slaves.

The example provided defines only one replica of the MySQL master server. This causes OpenShift Container Platform to start only one instance of the server. Multiple instances (multi-master) is not supported and therefore you can not scale this replication controller.

To replicate the database created by the [MySQL master](#), a deployment configuration is defined in the template. This deployment configuration creates a replication controller that launches the MySQL image with the **command** field set to **run-mysqld-slave**. This alternative endpoints skips the initialization of the database and configures the MySQL server to connect to the **mysql-master** service, which is also defined in example template.

Example 3.2. MySQL Slave Deployment Configuration Object Definition in the Example Template

```

kind: "DeploymentConfig"
apiVersion: "v1"

```

```

metadata:
  name: "mysql-slave"
spec:
  strategy:
    type: "Recreate"
  triggers:
    - type: "ConfigChange"
  replicas: 1
  selector:
    name: "mysql-slave"
  template:
    metadata:
      labels:
        name: "mysql-slave"
    spec:
      containers:
        - name: "server"
          image: "openshift/mysql-55-centos7"
          command:
            - "run-mysqld-slave"
          ports:
            - containerPort: 3306
              protocol: "TCP"
          env:
            - name: "MYSQL_MASTER_USER"
              value: "${MYSQL_MASTER_USER}"
            - name: "MYSQL_MASTER_PASSWORD"
              value: "${MYSQL_MASTER_PASSWORD}"
            - name: "MYSQL_DATABASE"
              value: "${MYSQL_DATABASE}"
          resources: {}
          terminationMessagePath: "/dev/termination-log"
          imagePullPolicy: "IfNotPresent"
          securityContext:
            capabilities: {}
            privileged: false
          restartPolicy: "Always"
          dnsPolicy: "ClusterFirst"

```

This example deployment configuration starts the replication controller with the initial number of replicas set to **1**. You can [scale this replication controller](#) in both directions, up to the resources capacity of your account.

3.2.6.2. Creating a Headless Service

The pods created by the MySQL slave replication controller must reach the MySQL master server in order to register for replication. The example template defines a headless service named **mysql-master** for this purpose. This service is not used only for replication, but the clients can also send the queries to **mysql-master:3306** as the MySQL host.

To have a headless service, the **portaIP** parameter in the service definition is set to **None**. Then you can use a DNS query to get a list of the pod IP addresses that represents the current endpoints for this service.

Example 3.3. Headless Service Object Definition in the Example Template

```

kind: "Service"
apiVersion: "v1"
metadata:
  name: "mysql-master"
  labels:
    name: "mysql-master"
spec:
  ports:
    - protocol: "TCP"
      port: 3306
      targetPort: 3306
      nodePort: 0
  selector:
    name: "mysql-master"
  portalIP: "None"
  type: "ClusterIP"
  sessionAffinity: "None"
status:
  loadBalancer: {}

```

3.2.6.3. Scaling the MySQL Slaves

To [increase the number of members](#) in the cluster:

```
$ oc scale rc mysql-slave-1 --replicas=<number>
```

This tells [the replication controller](#) to create a new MySQL slave pod. When a new slave is created, the slave endpoint first attempts to contact the **mysql-master** service and register itself to the replication set. Once that is done, the MySQL master server sends the slave the replicated database.

When scaling down, the MySQL slave is shut down and, because the slave does not have any persistent storage defined, all data on the slave is lost. The MySQL master server then discovers that the slave is not reachable anymore, and it automatically removes it from the replication.

3.2.7. Troubleshooting

This section describes some troubles you might encounter and presents possible resolutions.

3.2.7.1. Linux Native AIO Failure**Symptom**

The MySQL container fails to start and the logs show something like:

```

151113 5:06:56 InnoDB: Using Linux native AIO
151113 5:06:56 InnoDB: Warning: io_setup() failed with EAGAIN. Will make
5 attempts before giving up.
InnoDB: Warning: io_setup() attempt 1 failed.

```

```
InnoDB: Warning: io_setup() attempt 2 failed.
Waiting for MySQL to start ...
InnoDB: Warning: io_setup() attempt 3 failed.
InnoDB: Warning: io_setup() attempt 4 failed.
Waiting for MySQL to start ...
InnoDB: Warning: io_setup() attempt 5 failed.
151113 5:06:59 InnoDB: Error: io_setup() failed with EAGAIN after 5
attempts.
InnoDB: You can disable Linux Native AIO by setting innodb_use_native_aio
= 0 in my.cnf
151113 5:06:59 InnoDB: Fatal error: cannot initialize AIO sub-system
151113 5:06:59 [ERROR] Plugin 'InnoDB' init function returned error.
151113 5:06:59 [ERROR] Plugin 'InnoDB' registration as a STORAGE ENGINE
failed.
151113 5:06:59 [ERROR] Unknown/unsupported storage engine: InnoDB
151113 5:06:59 [ERROR] Aborting
```

Explanation

MySQL's storage engine was unable to use the kernel's AIO (Asynchronous I/O) facilities due to resource limits.

Resolution

Turn off AIO usage entirely by setting environment variable **MYSQL_AIO** to have value **0**. On subsequent deployments, this arranges for the MySQL configuration variable **innodb_use_native_aio** to have value **0**.

Alternatively, increase the **aio-max-nr** kernel resource. The following example examines the current value of **aio-max-nr** and doubles it.

```
$ sysctl fs.aio-max-nr
fs.aio-max-nr = 1048576
# sysctl -w fs.aio-max-nr=2097152
```

This is a per-node resolution and lasts until the next node reboot.

3.3. POSTGRESQL

3.3.1. Overview

OpenShift Container Platform provides a container image for running PostgreSQL. This image can provide database services based on username, password, and database name settings provided via configuration.

3.3.2. Versions

Currently, OpenShift Container Platform supports versions [9.2](#), [9.4](#), and [9.5](#) of PostgreSQL.

3.3.3. Images

These images come in two flavors, depending on your needs:

- RHEL 7

- CentOS 7

RHEL 7 Based Image

The RHEL 7 images are available through the Red Hat Registry:

```
$ docker pull registry.access.redhat.com/openshift3/postgresql-92-rhel7
$ docker pull registry.access.redhat.com/rhsc1/postgresql-94-rhel7
$ docker pull registry.access.redhat.com/rhsc1/postgresql-95-rhel7
```

CentOS 7 Based Image

These images are available on Docker Hub:

```
$ docker pull openshift/postgresql-92-centos7
$ docker pull centos/postgresql-94-centos7
$ docker pull centos/postgresql-95-centos7
```

or

```
$ docker pull centos/postgresql-94-centos7
```

To use these images, you can either access them directly from these registries or push them into your OpenShift Container Platform Docker registry. Additionally, you can create an ImageStream that points to the image, either in your Docker registry or at the external location. Your OpenShift Container Platform resources can then reference the ImageStream. You can find [example](#) ImageStream definitions for all the provided OpenShift Container Platform images.

3.3.4. Configuration and Usage

3.3.4.1. Initializing the Database

The first time you use the shared volume, the database is created along with the database administrator user and the PostgreSQL postgres user (if you specify the **POSTGRESQL_ADMIN_PASSWORD** environment variable). Afterwards, the PostgreSQL daemon starts up. If you are re-attaching the volume to another container, then the database, the database user, and the administrator user are not created, and the PostgreSQL daemon starts.

The following command creates a new database [pod](#) with PostgreSQL running in a container:

```
$ oc new-app \
  -e POSTGRESQL_USER=<username> \
  -e POSTGRESQL_PASSWORD=<password> \
  -e POSTGRESQL_DATABASE=<database_name> \
  registry.access.redhat.com/rhsc1/postgresql-94-rhel7
```

3.3.4.2. Running PostgreSQL Commands in Containers

OpenShift Container Platform uses [Software Collections](#) (SCLs) to install and launch PostgreSQL. If you want to execute a PostgreSQL command inside of a running container (for debugging), you must invoke it using bash.

To do so, first identify the name of the running PostgreSQL pod. For example, you can view the list of pods in your current project:

```
$ oc get pods
```

Then, open a remote shell session to the desired pod:

```
$ oc rsh <pod>
```

When you enter the container, the required SCL is automatically enabled.

You can now run the **psql** command from the bash shell to start a PostgreSQL interactive session and perform normal PostgreSQL operations. For example, to authenticate as the database user:

```
bash-4.2$ PGPASSWORD=$POSTGRESQL_PASSWORD psql -h postgresql
$POSTGRESQL_DATABASE $POSTGRESQL_USER
psql (9.2.8)
Type "help" for help.

default=>
```

When you are finished, enter **\q** to leave the PostgreSQL session.

3.3.4.3. Environment Variables

The PostgreSQL user name, password, and database name must be configured with the following environment variables:

Table 3.3. PostgreSQL Environment Variables

Variable Name	Description
POSTGRESQL_USER	User name for the PostgreSQL account to be created. This user has full rights to the database.
POSTGRESQL_PASSWORD	Password for the user account.
POSTGRESQL_DATABASE	Database name.
POSTGRESQL_ADMIN_PASSWORD	Optional password for the postgres administrator user. If this is not set, then remote login to the postgres account is not possible. Local connections from within the container are always permitted without a password.

**WARNING**

You must specify the user name, password, and database name. If you do not specify all three, the pod will fail to start and OpenShift Container Platform will continuously try to restart it.

PostgreSQL settings can be configured with the following environment variables:

Table 3.4. Additional PostgreSQL settings

Variable Name	Description	Default
POSTGRESQL_MAX_CONNECTIONS	Maximum number of client connections allowed.	100
POSTGRESQL_MAX_PREPARED_TRANSACTIONS	Maximum number of transactions that can be in the "prepared" state. If using prepared transactions, the value should be at least as large as POSTGRESQL_MAX_CONNECTIONS .	0
POSTGRESQL_SHARED_BUFFERS	Amount of memory dedicated to PostgreSQL for caching data.	32M
POSTGRESQL_EFFECTIVE_CACHE_SIZE	Estimated amount of memory available for disk caching by the operating system and within PostgreSQL itself.	128M

3.3.4.4. Volume Mount Points

The PostgreSQL image can be run with mounted volumes to enable persistent storage for the database:

- **/var/lib/pgsql/data** - This is the database cluster directory where PostgreSQL stores database files.

3.3.4.5. Changing Passwords

Passwords are part of the image configuration, therefore the only supported method to change passwords for the database user (**POSTGRESQL_USER**) and **postgres** administrator user is by changing the environment variables **POSTGRESQL_PASSWORD** and **POSTGRESQL_ADMIN_PASSWORD**, respectively.

You can view the current passwords by viewing the pod or deployment configuration in the web console or by listing the environment variables with the CLI:

```
$ oc set env pod <pod_name> --list
```

Changing database passwords through SQL statements or any way other than through the

environment variables aforementioned will cause a mismatch between the values stored in the variables and the actual passwords. Whenever a database container starts, it resets the passwords to the values stored in the environment variables.

To change these passwords, update one or both of the desired environment variables for the related deployment configuration(s) using the **oc set env** command. If multiple deployment configurations utilize these environment variables, for example in the case of an application created from a template, you must update the variables on each deployment configuration so that the passwords are in sync everywhere. This can be done all in the same command:

```
$ oc set env dc <dc_name> [<dc_name_2> ...] \
  POSTGRESQL_PASSWORD=<new_password> \
  POSTGRESQL_ADMIN_PASSWORD=<new_admin_password>
```



IMPORTANT

Depending on your application, there may be other environment variables for passwords in other parts of the application that should also be updated to match. For example, there could be a more generic **DATABASE_USER** variable in a front-end pod that should match the database user's password. Ensure that passwords are in sync for all required environment variables per your application, otherwise your pods may fail to redeploy when triggered.

Updating the environment variables triggers the redeployment of the database server if you have a [configuration change trigger](#). Otherwise, you must manually start a new deployment in order to apply the password changes.

To verify that new passwords are in effect, first open a remote shell session to the running PostgreSQL pod:

```
$ oc rsh <pod>
```

From the bash shell, verify the database user's new password:

```
bash-4.2$ PGPASSWORD=<new_password> psql -h postgresql
$POSTGRESQL_DATABASE $POSTGRESQL_USER -c "SELECT * FROM (SELECT
current_database()) cdb CROSS JOIN (SELECT current_user) cu"
```

If the password was changed correctly, you should see a table like this:

```
current_database | current_user
-----+-----
default         | django
(1 row)
```

From the bash shell, verify the **postgres** administrator user's new password:

```
bash-4.2$ PGPASSWORD=<new_admin_password> psql -h postgresql
$POSTGRESQL_DATABASE postgres -c "SELECT * FROM (SELECT
current_database()) cdb CROSS JOIN (SELECT current_user) cu"
```

If the password was changed correctly, you should see a table like this:

```
current_database | current_user
-----+-----
default         | postgres
(1 row)
```

3.3.5. Creating a Database Service from a Template

OpenShift Container Platform provides a [template](#) to make creating a new database service easy. The template provides parameter fields to define all the mandatory environment variables (user, password, database name, etc) with predefined defaults including auto-generation of password values. It will also define both a [deployment configuration](#) and a [service](#).

The PostgreSQL templates should have been registered in the default **openshift** project by your cluster administrator during the initial cluster setup. See [Loading the Default Image Streams and Templates](#) for more details, if required.

There are two templates available:

- **PostgreSQL-ephemeral** is for development or testing purposes only because it uses ephemeral storage for the database content. This means that if the database pod is restarted for any reason, such as the pod being moved to another node or the deployment configuration being updated and triggering a redeploy, all data will be lost.
- **PostgreSQL-persistent** uses a persistent volume store for the database data which means the data will survive a pod restart. Using persistent volumes requires a persistent volume pool be defined in the OpenShift Container Platform deployment. Cluster administrator instructions for setting up the pool are located [here](#).

You can find instructions for instantiating templates by following these [instructions](#).

Once you have instantiated the service, you can copy the user name, password, and database name environment variables into a deployment configuration for another component that intends to access the database. That component can then access the database via the service that was defined.

3.4. MONGODB

3.4.1. Overview

OpenShift Container Platform provides a container image for running MongoDB. This image can provide database services based on username, password, and database name settings provided via configuration.

3.4.2. Versions

Currently, OpenShift Container Platform provides versions [2.4](#), [2.6](#), and [3.2](#) of MongoDB.

3.4.3. Images

These images come in two flavors, depending on your needs:

- RHEL 7

- CentOS 7

RHEL 7 Based Images

The RHEL 7 images are available through the Red Hat Registry:

```
$ docker pull registry.access.redhat.com/openshift3/mongodb-24-rhel7
$ docker pull registry.access.redhat.com/rhsc1/mongodb-26-rhel7
$ docker pull registry.access.redhat.com/rhsc1/mongodb-32-rhel7
```

CentOS 7 Based Images

These images are available on Docker Hub:

```
$ docker pull openshift/mongodb-24-centos7
$ docker pull centos/mongodb-26-centos7
$ docker pull centos/mongodb-32-centos7
```

To use these images, you can either access them directly from these registries or push them into your OpenShift Container Platform Docker registry. Additionally, you can create an ImageStream that points to the image, either in your Docker registry or at the external location. Your OpenShift Container Platform resources can then reference the ImageStream. You can find [example](#) ImageStream definitions for all the provided OpenShift Container Platform images.

3.4.4. Configuration and Usage

3.4.4.1. Initializing the Database

You can configure MongoDB with an ephemeral volume or a persistent volume. The first time you use the volume, the database is created along with the database administrator user. Afterwards, the MongoDB daemon starts up. If you are re-attaching the volume to another container, then the database, database user, and the administrator user are not created, and the MongoDB daemon starts.

The following command creates a new database [pod](#) with MongoDB running in a container with an ephemeral volume:

```
$ oc new-app \
  -e MONGODB_USER=<username> \
  -e MONGODB_PASSWORD=<password> \
  -e MONGODB_DATABASE=<database_name> \
  -e MONGODB_ADMIN_PASSWORD=<admin_password> \
  registry.access.redhat.com/rhsc1/mongodb-26-rhel7
```

3.4.4.2. Running MongoDB Commands in Containers

OpenShift Container Platform uses [Software Collections](#) (SCLs) to install and launch MongoDB. If you want to execute a MongoDB command inside of a running container (for debugging), you must invoke it using bash.

To do so, first identify the name of the running MongoDB pod. For example, you can view the list of pods in your current project:

```
$ oc get pods
```

Then, open a remote shell session to the desired pod:

```
$ oc rsh <pod>
```

When you enter the container, the required SCL is automatically enabled.

You can now run **mongo** commands from the bash shell to start a MongoDB interactive session and perform normal MongoDB operations. For example, to switch to the **sampledb** database and authenticate as the database user:

```
bash-4.2$ mongo -u $MONGODB_USER -p $MONGODB_PASSWORD $MONGODB_DATABASE
MongoDB shell version: 2.4.9
connecting to: sampledb
>
```

When you are finished, press **CTRL+D** to leave the MongoDB session.

3.4.4.3. Environment Variables

The MongoDB user name, password, database name, and **admin** password must be configured with the following environment variables:

Table 3.5. MongoDB Environment Variables

Variable Name	Description
MONGODB_USER	User name for MongoDB account to be created.
MONGODB_PASSWORD	Password for the user account.
MONGODB_DATABASE	Database name.
MONGODB_ADMIN_PASSWORD	Password for the admin user.



WARNING

You must specify the user name, password, database name, and **admin** password. If you do not specify all four, the pod will fail to start and OpenShift Container Platform will continuously try to restart it.




NOTE

The administrator user name is set to **admin** and you must specify its password by setting the **MONGODB_ADMIN_PASSWORD** environment variable. This process is done upon database initialization.

MongoDB settings can be configured with the following environment variables:

Table 3.6. Additional MongoDB Settings

Variable Name	Description	Default
MONGODB_NOPREALLOC	Disable data file preallocation.	true
MONGODB_SMALLFILES	Set MongoDB to use a smaller default data file size.	true
MONGODB_QUIET	Runs MongoDB in a quiet mode that attempts to limit the amount of output.	true
MONGODB_TEXT_SEARCH_ENABLED	<p>(MongoDB version 2.4 only) Enables the text search feature.</p> <div>  <div> <p>NOTE</p> <p>Text search is enabled by default in MongoDB versions 2.6 and higher, and therefore has no configurable parameter.</p> </div> </div>	false

3.4.4.4. Volume Mount Points

The MongoDB image can be run with mounted volumes to enable persistent storage for the database:

- **/var/lib/mongodb** - This is the database directory where MongoDB stores database files.

3.4.4.5. Changing Passwords

Passwords are part of the image configuration, therefore the only supported method to change passwords for the database user (**MONGODB_USER**) and **admin** user is by changing the environment variables **MONGODB_PASSWORD** and **MONGODB_ADMIN_PASSWORD**, respectively.

You can view the current passwords by viewing the pod or deployment configuration in the web console or by listing the environment variables with the CLI:

```
$ oc set env pod <pod_name> --list
```

Changing database passwords directly in MongoDB causes a mismatch between the values stored in the variables and the actual passwords. Whenever a database container starts, it resets the passwords to the values stored in the environment variables.

To change these passwords, update one or both of the desired environment variables for the related deployment configuration(s) using the **oc set env** command. If multiple deployment configurations utilize these environment variables, for example in the case of an application created from a template, you must update the variables on each deployment configuration so that the passwords are in sync everywhere. This can be done all in the same command:

```
$ oc set env dc <dc_name> [<dc_name_2> ...] \
  MONGODB_PASSWORD=<new_password> \
  MONGODB_ADMIN_PASSWORD=<new_admin_password>
```



IMPORTANT

Depending on your application, there may be other environment variables for passwords in other parts of the application that should also be updated to match. For example, there could be a more generic **DATABASE_USER** variable in a front-end pod that should match the database user's password. Ensure that passwords are in sync for all required environment variables per your application, otherwise your pods may fail to redeploy when triggered.

Updating the environment variables triggers the redeployment of the database server if you have a [configuration change trigger](#). Otherwise, you must manually start a new deployment in order to apply the password changes.

To verify that new passwords are in effect, first open a remote shell session to the running MongoDB pod:

```
$ oc rsh <pod>
```

From the bash shell, verify the database user's new password:

```
bash-4.2$ mongo -u $MONGODB_USER -p <new_password> $MONGODB_DATABASE --
eval "db.version()"
```

If the password was changed correctly, you should see output like this:

```
MongoDB shell version: 2.6.9
connecting to: sampled
2.6.9
```

To verify the **admin** user's new password:

```
bash-4.2$ mongo -u admin -p <new_admin_password> admin --eval
"db.version()"
```

If the password was changed correctly, you should see output like this:

```
MongoDB shell version: 2.4.9
connecting to: admin
2.4.9
```

3.4.5. Creating a Database Service from a Template

OpenShift Container Platform provides a [template](#) to make creating a new database service easy. The template provides parameter fields to define all the mandatory environment variables (user, password, database name, etc) with predefined defaults including auto-generation of password values. It will also define both a [deployment configuration](#) and a [service](#).

The MongoDB templates should have been registered in the default **openshift** project by your cluster administrator during the initial cluster setup. See [Loading the Default Image Streams and Templates](#) for more details, if required.

There are two templates available:

- **mongodb-ephemeral** is for development/testing purposes only because it uses ephemeral storage for the database content. This means that if the database pod is restarted for any reason, such as the pod being moved to another node or the deployment configuration being updated and triggering a redeploy, all data will be lost.
- **mongodb-persistent** uses a persistent volume store for the database data which means the data will survive a pod restart. Using persistent volumes requires a persistent volume pool be defined in the OpenShift Container Platform deployment. Cluster administrator instructions for setting up the pool are located [here](#).

You can find instructions for instantiating templates by following these [instructions](#).

Once you have instantiated the service, you can copy the user name, password, and database name environment variables into a deployment configuration for another component that intends to access the database. That component can then access the database via the service that was defined.

3.4.6. Using MongoDB Replication



NOTE

Enabling clustering for database images is currently in Technology Preview and not intended for production use.

Red Hat provides a proof-of-concept [template](#) for MongoDB replication (clustering); you can obtain the [example template from GitHub](#).

For example, to upload the example template into the current project's template library:

```
$ oc create -f \
https://raw.githubusercontent.com/openshift/mongodb/master/2.4/examples/re
plica/mongodb-clustered.json
```



IMPORTANT

The example template does not use persistent storage. When you lose all members of the replication set, your data will be lost.

The following sections detail the objects defined in the example template and describe how they work together to start a cluster of MongoDB servers implementing master-slave replication and automated failover. This is the recommended replication strategy for MongoDB.

3.4.6.1. Creating the Deployment Configuration

To set up MongoDB replication, a [deployment configuration](#) is defined in the example template that defines a [replication controller](#). The replication controller manages the members of the MongoDB cluster.

To tell a MongoDB server that the member will be part of the cluster, additional environment variables are provided for the container defined in the replication controller pod template:

Variable Name	Description	Default
MONGODB_REPLICA_NAME	Specifies the name of the replication set.	rs0
MONGODB_KEYFILE_VALUE	See: Generate a Key File	generated

Example 3.4. Deployment Configuration Object Definition in the Example Template

```
kind: DeploymentConfig
apiVersion: v1
metadata:
  name: "${MONGODB_SERVICE_NAME}"
spec:
  strategy:
    type: Recreate
    resources: {}
  triggers:
    - type: ConfigChange
  replicas: 3
  selector:
    name: mongodb-replica
  template:
    metadata:
      labels:
        name: mongodb-replica
    spec:
      containers:
        - name: member
          image: openshift/mongodb-24-centos7
          env:
            - name: MONGODB_USER
              value: "${MONGODB_USER}"
            - name: MONGODB_PASSWORD
              value: "${MONGODB_PASSWORD}"
            - name: MONGODB_DATABASE
              value: "${MONGODB_DATABASE}"
            - name: MONGODB_ADMIN_PASSWORD
              value: "${MONGODB_ADMIN_PASSWORD}"
            - name: MONGODB_REPLICA_NAME
              value: "${MONGODB_REPLICA_NAME}"
            - name: MONGODB_SERVICE_NAME
              value: "${MONGODB_SERVICE_NAME}"
            - name: MONGODB_KEYFILE_VALUE
```

```

        value: "${MONGODB_KEYFILE_VALUE}"
    ports:
      - containerPort: 27017
        protocol: TCP
    restartPolicy: Never
    dnsPolicy: ClusterFirst

```

After the deployment configuration is created and the pods with MongoDB cluster members are started, they will not be initialized. Instead, they start as part of the **rs0** replication set, as the value of **MONGODB_REPLICA_NAME** is set to **rs0** by default.

3.4.6.2. Creating the Service Pod

To initialize members created by [the deployment configuration](#), the pods are started with the **initiate** argument, which instructs the startup script to behave [slightly differently](#) than a regular, stand-alone MongoDB database.

Example 3.5. Deployment Configuration Object Definition in the Example Template

```

- kind: DeploymentConfig
  apiVersion: v1
  metadata:
    name: "${MONGODB_SERVICE_NAME}"
  spec:
    strategy:
      type: Recreate
      recreateParams:
        post:
          failurePolicy: Retry
          execNewPod:
            command: ["run-mongod", "initiate"]
            containerName: mongodb
            env:
              - name: MONGODB_INITIAL_REPLICA_COUNT
                value: '3'
    triggers:
      - type: ConfigChange
    replicas: 3
    selector:
      name: mongodb-replica
    template:
      metadata:
        labels:
          name: mongodb-replica
      spec:
        containers:
          - name: mongodb
            image: openshift/mongodb-24-centos7
            readinessProbe:
              tcpSocket:
                port: 27017
              initialDelaySeconds: 15
              timeoutSeconds: 1

```

```

env:
  - name: MONGODB_USER
    value: "${MONGODB_USER}"
  - name: MONGODB_PASSWORD
    value: "${MONGODB_PASSWORD}"
  - name: MONGODB_DATABASE
    value: "${MONGODB_DATABASE}"
  - name: MONGODB_ADMIN_PASSWORD
    value: "${MONGODB_ADMIN_PASSWORD}"
  - name: MONGODB_REPLICA_NAME
    value: "${MONGODB_REPLICA_NAME}"
  - name: MONGODB_SERVICE_NAME
    value: "${MONGODB_SERVICE_NAME}"
  - name: MONGODB_KEYFILE_VALUE
    value: "${MONGODB_KEYFILE_VALUE}"
ports:
  - containerPort: 27017

```

3.4.6.3. Creating a Headless Service

The **initiate** argument in the [container specification above](#) instructs the container to first discover all running member pods within the MongoDB cluster. To achieve this, a *headless service* is defined named **mongodb** in the example template.

To have a headless service, the **portalIP** parameter in the service definition is set to **None**. Then you can use a DNS query to get a list of the pod IP addresses that represents the current endpoints for this service.

Example 3.6. Headless Service Object Definition in the Example Template

```

kind: "Service"
apiVersion: "v1"
metadata:
  name: "${MONGODB_SERVICE_NAME}"
  labels:
    name: "${MONGODB_SERVICE_NAME}"
spec:
  ports:
    - protocol: "TCP"
      port: 27017
      targetPort: 27017
      nodePort: 0
  selector:
    name: "mongodb-replica"
  portalIP: "None"
  type: "ClusterIP"
  sessionAffinity: "None"
status:
  loadBalancer: {}

```

3.4.6.4. Creating the Final Replication Set

When the script that runs as the container entrypoint has the IP addresses of all running MongoDB members, it creates a MongoDB replication set configuration where it lists all member IP addresses. It then initiates the replication set using `rs.initiate(config)`. The script waits until MongoDB elects the **PRIMARY** member of the cluster.

Once the **PRIMARY** member has been elected, the entrypoint script starts creating MongoDB users and databases.

Clients can then start using the MongoDB instance by sending the queries to the **mongodb** service. As this service is a headless service, they do not need to provide the IP address. Clients can use **mongodb:27017** for connections. The service then sends the query to one of the members in the replication set.

3.4.6.5. Scaling the MongoDB Replication Set

To [increase the number of members](#) in the cluster:

```
$ oc scale rc mongodb-1 --replicas=<number>
```

This tells [the replication controller](#) to create a new MongoDB member pod. When a new member is created, the member entrypoint first attempts to discover other running members in the cluster. It then chooses one and adds itself to the list of members. Once the replication configuration is updated, the other members replicate the data to a new pod and start a new election.

3.5. MARIADB

3.5.1. Overview

OpenShift Container Platform provides a container image for running MariaDB. This image can provide database services based on username, password, and database name settings provided in a configuration file.

3.5.2. Versions

Currently, OpenShift Container Platform provides versions [10.0](#) and [10.1](#) of MariaDB.

3.5.3. Images

These images come in two flavors, depending on your needs:

- RHEL 7
- CentOS 7

RHEL 7 Based Images

The RHEL 7 images are available through the Red Hat Registry:

```
$ docker pull registry.access.redhat.com/rhsc1/mariadb-100-rhel7
$ docker pull registry.access.redhat.com/rhsc1/mariadb-101-rhel7
```

CentOS 7 Based Images

These images are available on Docker Hub:

```
$ docker pull openshift/mariadb-100-centos7
$ docker pull centos/mariadb-101-centos7
```

To use these images, you can either access them directly from these registries or push them into your OpenShift Container Platform Docker registry. Additionally, you can create an ImageStream that points to the image, either in your Docker registry or at the external location. Your OpenShift Container Platform resources can then reference the ImageStream. You can find [example](#) ImageStream definitions for all the provided OpenShift Container Platform images.

3.5.4. Configuration and Usage

3.5.4.1. Initializing the Database

The first time you use the shared volume, the database is created along with the database administrator user and the MariaDB root user (if you specify the **MYSQL_ROOT_PASSWORD** environment variable). Afterwards, the MariaDB daemon starts up. If you are re-attaching the volume to another container, then the database, database user, and the administrator user are not created, and the MariaDB daemon starts.

The following command creates a new database [pod](#) with MariaDB running in a container:

```
$ oc new-app \
  -e MYSQL_USER=<username> \
  -e MYSQL_PASSWORD=<password> \
  -e MYSQL_DATABASE=<database_name> \
  registry.access.redhat.com/rhsc1/mariadb-101-rhel7
```

3.5.4.2. Running MariaDB Commands in Containers

OpenShift Container Platform uses [Software Collections](#) (SCLs) to install and launch MariaDB. If you want to execute a MariaDB command inside of a running container (for debugging), you must invoke it using bash.

To do so, first identify the name of the running MariaDB pod. For example, you can view the list of pods in your current project:

```
$ oc get pods
```

Then, open a remote shell session to the pod:

```
$ oc rsh <pod>
```

When you enter the container, the required SCL is automatically enabled.

You can now run **mysql** commands from the bash shell to start a MariaDB interactive session and perform normal MariaDB operations. For example, to authenticate as the database user:

```
bash-4.2$ mysql -u $MYSQL_USER -p$MYSQL_PASSWORD -h $HOSTNAME
$MYSQL_DATABASE
```

```
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 4
Server version: 5.5.37 MySQL Community Server (GPL)
...
mysql>
```

When you are finished, enter **quit** or **exit** to leave the MySQL session.

3.5.4.3. Environment Variables

The MariaDB user name, password, and database name must be configured with the following environment variables:

Table 3.7. MariaDB Environment Variables

Variable Name	Description
MYSQL_USER	User name for MySQL account to be created.
MYSQL_PASSWORD	Password for the user account.
MYSQL_DATABASE	Database name.
MYSQL_ROOT_PASSWORD	Password for the root user (optional).



WARNING

You must specify the user name, password, and database name. If you do not specify all three, the pod will fail to start and OpenShift Container Platform will continuously try to restart it.

MariaDB settings can be configured with the following environment variables:

Table 3.8. Additional MariaDB Settings

Variable Name	Description	Default
MYSQL_LOWER_CASE_TABLE_NAMES	Sets how the table names are stored and compared.	0
MYSQL_MAX_CONNECTIONS	The maximum permitted number of simultaneous client connections.	151
MYSQL_MAX_ALLOWED_PACKET	The maximum size of one packet or any generated/intermediate string.	200M

Variable Name	Description	Default
MYSQL_FT_MIN_WORD_LENGTH	The minimum length of the word to be included in a FULLTEXT index.	4
MYSQL_FT_MAX_WORD_LENGTH	The maximum length of the word to be included in a FULLTEXT index.	20
MYSQL_AIO	Controls the innodb_use_native_aio setting value if the native AIO is broken.	1
MYSQL_TABLE_OPEN_CACHE	The number of open tables for all threads.	400
MYSQL_KEY_BUFFER_SIZE	The size of the buffer used for index blocks.	32M (or 10% of available memory)
MYSQL_SORT_BUFFER_SIZE	The size of the buffer used for sorting.	256K
MYSQL_READ_BUFFER_SIZE	The size of the buffer used for a sequential scan.	8M (or 5% of available memory)
MYSQL_INNODB_BUFFER_POOL_SIZE	The size of the buffer pool where InnoDB caches table and index data.	32M (or 50% of available memory)
MYSQL_INNODB_LOG_FILE_SIZE	The size of each log file in a log group.	8M (or 15% of available memory)

Variable Name	Description	Default
MYSQL_INNODB_LOG_BUFFER_SIZE	The size of the buffer that InnoDB uses to write to the log files on disk.	8M (or 15% of available memory)
MYSQL_DEFAULTS_FILE	Point to an alternative configuration file.	/etc/my.cnf
MYSQL_BINLOG_FORMAT	Set sets the binlog format, supported values are row and statement .	statement

3.5.4.4. Volume Mount Points

The MariaDB image can be run with mounted volumes to enable persistent storage for the database:

- **/var/lib/mysql/data** - The MySQL data directory is where MariaDB stores database files.



NOTE

When mounting a directory from the host into the container, ensure that the mounted directory has the appropriate permissions. Also verify that the owner and group of the directory match the user name running inside the container.

3.5.4.5. Changing Passwords

Passwords are part of the image configuration, therefore the only supported method to change passwords for the database user (**MYSQL_USER**) and **admin** user is by changing the environment variables **MYSQL_PASSWORD** and **MYSQL_ROOT_PASSWORD**, respectively.

You can view the current passwords by viewing the pod or deployment configuration in the web console or by listing the environment variables with the CLI:

```
$ oc set env pod <pod_name> --list
```

Changing database passwords through SQL statements or any way other than through the environment variables aforementioned causes a mismatch between the values stored in the variables and the actual passwords. Whenever a database container starts, it resets the passwords to the values stored in the environment variables.

To change these passwords, update one or both of the desired environment variables for the related deployment configuration(s) using the **oc set env** command. If multiple deployment configurations utilize these environment variables, for example in the case of an application created from a template, you must update the variables on each deployment configuration so that the passwords are in sync everywhere. This can be done all in the same command:

```
$ oc set env dc <dc_name> [<dc_name_2> ...] \
  MYSQL_PASSWORD=<new_password> \
  MYSQL_ROOT_PASSWORD=<new_root_password>
```



IMPORTANT

Depending on your application, there may be other environment variables for passwords in other parts of the application that should also be updated to match. For example, there could be a more generic **DATABASE_USER** variable in a front-end pod that should match the database user's password. Ensure that passwords are in sync for all required environment variables per your application, otherwise your pods may fail to redeploy when triggered.

Updating the environment variables triggers the redeployment of the database server if you have a [configuration change trigger](#). Otherwise, you must manually start a new deployment in order to apply the password changes.

To verify that new passwords are in effect, first open a remote shell session to the running MariaDB pod:

```
$ oc rsh <pod>
```

From the bash shell, verify the database user's new password:

```
bash-4.2$ mysql -u $MYSQL_USER -p<new_password> -h $HOSTNAME
$MYSQL_DATABASE -te "SELECT * FROM (SELECT database()) db CROSS JOIN
(SELECT user()) u"
```

If the password was changed correctly, you should see a table like this:

```
+-----+-----+
| database() | user() |
+-----+-----+
| sampledbs | user0PG@172.17.42.1 |
+-----+-----+
```

To verify the **root** user's new password:

```
bash-4.2$ mysql -u root -p<new_root_password> -h $HOSTNAME $MYSQL_DATABASE
-te "SELECT * FROM (SELECT database()) db CROSS JOIN (SELECT user()) u"
```

If the password was changed correctly, you should see a table like this:

```
+-----+-----+
| database() | user() |
+-----+-----+
| sampledbs | root@172.17.42.1 |
+-----+-----+
```

3.5.5. Creating a Database Service from a Template

OpenShift Container Platform provides a [template](#) to make creating a new database service

easy. The template provides parameter fields to define all the mandatory environment variables (user, password, database name, etc) with predefined defaults including auto-generation of password values. It will also define both a [deployment configuration](#) and a [service](#).

The MariaDB templates should have been registered in the default **openshift** project by your cluster administrator during the initial cluster setup. See [Loading the Default Image Streams and Templates](#) for more details, if required.

There are two templates available:

- **mariadb-ephemeral** is for development or testing purposes only because it uses ephemeral storage for the database content. This means that if the database pod is restarted for any reason, such as the pod being moved to another node or the deployment configuration being updated and triggering a redeploy, all data will be lost.
- **mariadb-persistent** uses a persistent volume store for the database data which means the data will survive a pod restart. Using persistent volumes requires a persistent volume pool be defined in the OpenShift Container Platform deployment. Cluster administrator instructions for setting up the pool are located [here](#).

You can find instructions for instantiating templates by following these [instructions](#).

Once you have instantiated the service, you can copy the user name, password, and database name environment variables into a deployment configuration for another component that intends to access the database. That component can then access the database through the service that was defined.

3.5.6. Using MariaDB Replication




NOTE

Enabling clustering for database images is currently in Technology Preview and not intended for production use.

Red Hat provides a proof-of-concept [template](#) for MariaDB replication (clustering); you can obtain the [example template from GitHub](#).

To upload the example template into the current project's template library:

```
$ oc create -f \
  https://github.com/openshift/origin/blob/master/examples/db-
  templates/mariadb-ephemeral-template.json
```



IMPORTANT

The example template does not use persistent storage. When you lose all members of the replication set, your data will be lost.

The following sections detail the objects defined in the example template and describe how they work together to start a cluster of MariaDB servers implementing master-slave replication and automated failover. This is the recommended replication strategy for MariaDB.

3.5.6.1. Creating the Deployment Configuration

To set up MariaDB replication, a [deployment configuration](#) is defined in the example template that defines a [replication controller](#). The replication controller manages the members of the MariaDB cluster.

To tell a MariaDB server to act as the master, the **command** field in the container's definition in the deployment configuration must be set to **run-mysqld-master**. This script acts as an alternative entrypoint for the MariaDB image and configures the MariaDB server to run as the master in replication.

MariaDB replication requires a special user that relays data between the master and slaves. The following environment variables are defined in the template for this purpose:

Variable Name	Description	Default
MYSQL_MASTER_USER	The user name of the replication user	master
MYSQL_MASTER_PASSWORD	The password for the replication user	generated

Example 3.7. MySQL Master Deployment Configuration Object Definition in the Example Template

```
kind: "DeploymentConfig"
apiVersion: "v1"
metadata:
  name: "mysql-master"
spec:
  strategy:
    type: "Recreate"
  triggers:
    - type: "ConfigChange"
  replicas: 1
  selector:
    name: "mysql-master"
  template:
    metadata:
      labels:
        name: "mysql-master"
    spec:
      volumes:
        - name: "mysql-master-data"
          persistentVolumeClaim:
            claimName: "mysql-master"
      containers:
        - name: "server"
          image: "openshift/mariadb-100-centos7"
          command:
            - "run-mysqld-master"
          ports:
            - containerPort: 3306
```

```

        protocol: "TCP"
    env:
      - name: "MYSQL_MASTER_USER"
        value: "${MYSQL_MASTER_USER}"
      - name: "MYSQL_MASTER_PASSWORD"
        value: "${MYSQL_MASTER_PASSWORD}"
      - name: "MYSQL_USER"
        value: "${MYSQL_USER}"
      - name: "MYSQL_PASSWORD"
        value: "${MYSQL_PASSWORD}"
      - name: "MYSQL_DATABASE"
        value: "${MYSQL_DATABASE}"
      - name: "MYSQL_ROOT_PASSWORD"
        value: "${MYSQL_ROOT_PASSWORD}"
    volumeMounts:
      - name: "mysql-master-data"
        mountPath: "/var/lib/mysql/data"
    resources: {}
    terminationMessagePath: "/dev/termination-log"
    imagePullPolicy: "IfNotPresent"
    securityContext:
      capabilities: {}
      privileged: false
    restartPolicy: "Always"
    dnsPolicy: "ClusterFirst"

```

Since we claimed a persistent volume in this deployment configuration to have all data persisted for the MySQL master server, you must ask your cluster administrator to create a persistent volume that you can claim the storage from.

After the deployment configuration is created and the pod with MariaDB master server is started, it will create the database defined by **MYSQL_DATABASE** and configure the server to replicate this database to slaves.

The example provided defines only one replica of the MariaDB master server. This causes OpenShift Container Platform to start only one instance of the server. Multiple instances (multi-master) is not supported and therefore you can not scale this replication controller.

3.5.6.1.1. Creating the MySQL Slaves

To replicate the database created by the [MySQL master](#), a deployment configuration is defined in the template. This deployment configuration creates a replication controller that launches the MySQL image with the **command** field set to **run-mysqld-slave**. This alternative entrypoints skips the initialization of the database and configures the MySQL server to connect to the **mysql-master** service, which is also defined in example template.

Example 3.8. MariaDB Slave Deployment Configuration Object Definition in the Example Template

```

kind: "DeploymentConfig"
apiVersion: "v1"
metadata:
  name: "mysql-slave"
spec:

```

```

strategy:
  type: "Recreate"
triggers:
  - type: "ConfigChange"
replicas: 1
selector:
  name: "mysql-slave"
template:
  metadata:
    labels:
      name: "mysql-slave"
  spec:
    containers:
      - name: "server"
        image: "openshift/mariadb-100-centos7"
        command:
          - "run-mysqld-slave"
        ports:
          - containerPort: 3306
            protocol: "TCP"
        env:
          - name: "MYSQL_MASTER_USER"
            value: "${MYSQL_MASTER_USER}"
          - name: "MYSQL_MASTER_PASSWORD"
            value: "${MYSQL_MASTER_PASSWORD}"
          - name: "MYSQL_DATABASE"
            value: "${MYSQL_DATABASE}"
        resources: {}
        terminationMessagePath: "/dev/termination-log"
        imagePullPolicy: "IfNotPresent"
        securityContext:
          capabilities: {}
          privileged: false
        restartPolicy: "Always"
        dnsPolicy: "ClusterFirst"

```

This example deployment configuration starts the replication controller with the initial number of replicas set to **1**. You can [scale this replication controller](#) in both directions, up to the resources capacity of your account.

3.5.6.2. Creating a Headless Service

The **initiate** argument in the container specification instructs the container to first discover all running member pods within the MariaDB cluster. To achieve this, a *headless service* is defined named **mariadb** in the example template.

To have a headless service, the **portName** parameter in the service definition is set to **None**. Then you can use a DNS query to get a list of the pod IP addresses that represents the current endpoints for this service.

Example 3.9. Headless Service Object Definition in the Example Template

```

kind: "Service"
apiVersion: "v1"

```

```

metadata:
  name: "mysql-master"
  labels:
    name: "mysql-master"
spec:
  ports:
    - protocol: "TCP"
      port: 3306
      targetPort: 3306
      nodePort: 0
  selector:
    name: "mysql-master"
  portalIP: "None"
  type: "ClusterIP"
  sessionAffinity: "None"
status:
  loadBalancer: {}

```

3.5.6.3. Scaling the MariaDB Slaves

To [increase the number of members](#) in the cluster:

```
$ oc scale rc mysql-slave-1 --replicas=<number>
```

This tells [the replication controller](#) to create a new MySQL slave pod. When a new slave is created, the slave endpoint first attempts to contact the **mysql-master** service and register itself to the replication set. Once that is done, the MySQL master server sends the slave the replicated database.

When scaling down, the MySQL slave is shut down and, because the slave does not have any persistent storage defined, all data on the slave is lost. The MySQL master server then discovers that the slave is not reachable anymore, and it automatically removes it from the replication.

3.5.7. Troubleshooting

This section describes some troubles you might encounter and presents possible resolutions.

3.5.7.1. Linux Native AIO Failure

Symptom

The MySQL container fails to start and the logs show something like:

```

151113 5:06:56 InnoDB: Using Linux native AIO
151113 5:06:56 InnoDB: Warning: io_setup() failed with EAGAIN. Will make
5 attempts before giving up.
InnoDB: Warning: io_setup() attempt 1 failed.
InnoDB: Warning: io_setup() attempt 2 failed.
Waiting for MySQL to start ...
InnoDB: Warning: io_setup() attempt 3 failed.
InnoDB: Warning: io_setup() attempt 4 failed.

```

```
Waiting for MySQL to start ...
InnoDB: Warning: io_setup() attempt 5 failed.
151113  5:06:59  InnoDB: Error: io_setup() failed with EAGAIN after 5
attempts.
InnoDB: You can disable Linux Native AIO by setting innodb_use_native_aio
= 0 in my.cnf
151113  5:06:59 InnoDB: Fatal error: cannot initialize AIO sub-system
151113  5:06:59 [ERROR] Plugin 'InnoDB' init function returned error.
151113  5:06:59 [ERROR] Plugin 'InnoDB' registration as a STORAGE ENGINE
failed.
151113  5:06:59 [ERROR] Unknown/unsupported storage engine: InnoDB
151113  5:06:59 [ERROR] Aborting
```

Explanation

MariaDB's storage engine was unable to use the kernel's AIO (Asynchronous I/O) facilities due to resource limits.

Resolution

Turn off AIO usage entirely, by setting environment variable **MYSQL_AIO** to have value **0**. On subsequent deployments, this arranges for the MySQL configuration variable **innodb_use_native_aio** to have value **0**.

Alternatively, increase the **aio-max-nr** kernel resource. The following example examines the current value of **aio-max-nr** and doubles it.

```
$ sysctl fs.aio-max-nr
fs.aio-max-nr = 1048576
# sysctl -w fs.aio-max-nr=2097152
```

This is a per-node resolution and lasts until the next node reboot.

CHAPTER 4. DOCKER IMAGES

4.1. OVERVIEW

You can use arbitrary container images in your OpenShift Container Platform instance, for example those found on the [Docker Hub](#). For instructions on how to enable images to run with **USER** in the Dockerfile, see [Managing Security Context Constraints](#).

CHAPTER 5. OTHER IMAGES

5.1. OVERVIEW

This topic group includes information on other container images available for OpenShift Container Platform users.

5.2. JENKINS

5.2.1. Overview

OpenShift Container Platform provides a container image for running Jenkins. This image provides a Jenkins server instance which can be used to set up a basic flow for continuous testing, integration, and delivery.

This image also includes a sample Jenkins job which triggers a new build of a **BuildConfig** defined in OpenShift Container Platform, tests the output of that build, and then on successful build, retags the output to indicate the build is ready for production.

5.2.2. Versions

OpenShift Container Platform follows the [LTS](#) releases of Jenkins. Currently, OpenShift Container Platform provides versions 1.x and 2.x.

5.2.3. Images

These images come in two flavors, depending on your needs:

- RHEL 7
- CentOS 7

RHEL 7 Based Images

The RHEL 7 images are available through the Red Hat Registry:

```
$ docker pull registry.access.redhat.com/openshift3/jenkins-1-rhel7
$ docker pull registry.access.redhat.com/openshift3/jenkins-2-rhel7
```

CentOS 7 Based Images

This image is available on Docker Hub:

```
$ docker pull openshift/jenkins-1-centos7
$ docker pull openshift/jenkins-2-centos7
```

To use these images, you can either access them directly from these registries or push them into your OpenShift Container Platform Docker registry. Additionally, you can create an ImageStream that points to the image, either in your Docker registry or at the external location. Your OpenShift Container Platform resources can then reference the ImageStream. You can find [example](#) ImageStream definitions for all the provided OpenShift Container Platform images.

5.2.4. Configuration and Usage

5.2.4.1. Initializing Jenkins

The first time you start Jenkins, the configuration is created along with the administrator user and password. The default login is **admin/password**. The default password can be configured by setting the **JENKINS_PASSWORD** environment variable.

- OpenShift Container Platform OAuth authentication provided by the OpenShift Login plug-in.
- Standard authentication provided by Jenkins

5.2.4.1.1. OpenShift Container Platform OAuth authentication

[OAuth authentication](#) is activated by configuring the **Configure Global Security** panel in the Jenkins UI, or by setting the **OPENSIFT_ENABLE_OAUTH** environment variable on the Jenkins **Deployment Config** to anything other than **false**. This activates the OpenShift Login plug-in, which retrieves the configuration information from pod data or by interacting with the OpenShift Container Platform API server.

Valid credentials are controlled by the OpenShift Container Platform identity provider. For example, if **Allow All** is the default identity provider, you can provide any non-empty string for both the user name and password.

For non-browser access, the OpenShift Login plug-in also supports using the HTTP bearer token authorization header to supply valid credentials for accessing Jenkins. Ensure to use the token associated with the serviceaccount for the project in which Jenkins is running, which, if you started Jenkins using the [jenkins-ephemeral](#) template or [jenkins-persistent](#) template (see below), are found by using:

```
$ oc describe serviceaccount jenkins
$ oc describe secret <serviceaccount secret name>
```

Valid users are automatically added to the Jenkins authorization matrix at log in, where OpenShift Container Platform **Roles** dictate the specific Jenkins permissions the user will have.

Users with the **admin** role will have the traditional Jenkins administrative user permissions. Users with the **edit** or **view** role will have progressively less permissions. See the [Jenkins image source repository README](#) for the specifics on the OpenShift roles to Jenkins permissions mappings.



NOTE

The **admin** user that is pre-populated in the OpenShift Container Platform Jenkins image with administrative privileges will not be given those privileges when OpenShift Container Platform OAuth is used, unless the OpenShift Container Platform cluster administrator explicitly defines that user in the OpenShift Container Platform identity provider and assigns the **admin** role to the user.

Jenkins' users permissions can be changed after the users are initially established. The OpenShift Login plug-in polls the OpenShift Container Platform API server for permissions and updates the permissions stored in Jenkins for each user with the permissions retrieved

from OpenShift Container Platform. If the Jenkins UI is used to update permissions for a Jenkins user, the permission changes are overwritten the next time the plug-in polls OpenShift Container Platform.

You can control how often the polling occurs with the **OPENSIFT_PERMISSIONS_POLL_INTERVAL** environment variable. The default polling interval is five minutes.

The easiest way to create a new Jenkins service using OAuth authentication is to use a template as described below.

5.2.4.1.2. Jenkins Standard Authentication

Jenkins authentication is used by default if the image is run directly, without using a template.

The first time Jenkins starts, the configuration is created along with the administrator user and password. The default user credentials are **admin** and **password**. Configure the default password by setting the **JENKINS_PASSWORD** environment variable when using (and only when using) standard Jenkins authentication.

To create a new Jenkins application using standard Jenkins authentication:

```
$ oc new-app -e \
  JENKINS_PASSWORD=<password> \
  openshift/jenkins-1-centos7
```

5.2.4.2. Environment Variables

The Jenkins password can be configured with the following environment variable:

Table 5.1. Jenkins Environment Variables

Variable name	Description
JENKINS_PASSWORD	Password for the admin user.

5.2.4.3. Volume Mount Points

The Jenkins image can be run with mounted volumes to enable persistent storage for the configuration:

- **/var/lib/jenkins** - This is the data directory where Jenkins stores configuration files including job definitions.

5.2.5. Creating a Jenkins Service from a Template

[Templates](#) provide parameter fields to define all the environment variables (password) with predefined defaults. OpenShift Container Platform provides templates to make creating a new Jenkins service easy. The Jenkins templates should have been registered in the default **openshift** project by your cluster administrator during the initial cluster setup. See [Loading the Default Image Streams and Templates](#) for more details, if required.

The two available templates both define a [deployment configuration](#) and a [service](#). The templates differ in their storage strategy, which affects whether or not the Jenkins content persists across a pod restart.

**NOTE**

A pod may be restarted when it is moved to another node, or when an update of the deployment configuration triggers a redeployment.

- **jenkins-ephemeral** uses ephemeral storage. On pod restart, all data is lost. This template is useful for development or testing only.
- **jenkins-persistent** uses a persistent volume store. Data survives a pod restart. To use a persistent volume store, the cluster administrator must define a persistent volume pool in the OpenShift Container Platform deployment.