



# Red Hat CodeReady Workspaces 2.4

## Administration Guide

Administering Red Hat CodeReady Workspaces 2.4



# Red Hat CodeReady Workspaces 2.4 Administration Guide

---

Administering Red Hat CodeReady Workspaces 2.4

Robert Kratky  
rkratky@redhat.com

Michal Maléř  
mmaler@redhat.com

Fabrice Flore-Thébault  
ffloreth@redhat.com

Yana Hontyk  
yhontyk@redhat.com

## Legal Notice

Copyright © 2020 Red Hat, Inc.

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

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

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

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

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

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

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

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

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

The OpenStack<sup>®</sup> 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

Information for administrators operating Red Hat CodeReady Workspaces.

## Table of Contents

<b>CHAPTER 1. CODEREADY WORKSPACES ARCHITECTURE OVERVIEW</b> .....	<b>6</b>
1.1. CODEREADY WORKSPACES WORKSPACE CONTROLLER	6
1.1.1. CodeReady Workspaces server	7
1.1.2. CodeReady Workspaces user dashboard	7
1.1.3. Devfile registry	7
1.1.4. CodeReady Workspaces plug-in registry	8
1.1.5. CodeReady Workspaces and PostgreSQL	8
1.1.6. CodeReady Workspaces and RH-SSO	8
1.2. CODEREADY WORKSPACES WORKSPACES ARCHITECTURE	9
1.2.1. CodeReady Workspaces workspace components	10
1.2.1.1. Che Editor plug-in	10
1.2.1.2. CodeReady Workspaces user runtimes	11
1.2.1.3. CodeReady Workspaces workspace JWT proxy	11
1.2.1.4. CodeReady Workspaces plug-ins broker	11
1.2.2. CodeReady Workspaces workspace configuration	12
1.2.2.1. Storage strategies for codeready-workspaces workspaces	12
1.2.2.1.1. The common PVC strategy	13
1.2.2.1.2. The per-workspace PVC strategy	13
1.2.2.1.3. The unique PVC strategy	14
1.2.2.1.4. How subpaths are used in PVCs	14
1.2.2.2. Configuring a CodeReady Workspaces workspace with a persistent volume strategy	14
1.2.2.2.1. Configuring a PVC strategy using the Operator	15
1.2.2.3. Workspace OpenShift project configuration	16
1.2.3. CodeReady Workspaces workspace creation flow	16
<b>CHAPTER 2. CALCULATING CODEREADY WORKSPACES RESOURCE REQUIREMENTS</b> .....	<b>18</b>
2.1. CONTROLLER REQUIREMENTS	18
2.2. WORKSPACES REQUIREMENTS	18
2.3. A WORKSPACE EXAMPLE	21
<b>CHAPTER 3. CUSTOMIZING THE DEVFILE AND PLUG-IN REGISTRIES</b> .....	<b>24</b>
3.1. BUILDING AND RUNNING A CUSTOM REGISTRY IMAGE	24
3.1.1. Building a custom devfile registry	24
3.1.2. Building a custom plug-in registry	25
3.1.3. Deploying the registries	26
3.1.3.1. Deploying registries in OpenShift	27
3.2. INCLUDING THE PLUG-IN BINARIES IN THE REGISTRY IMAGE	29
3.3. EDITING A DEVFILE AND PLUG-IN AT RUNTIME	30
3.3.1. Adding a plug-in at runtime	30
3.3.2. Adding a devfile at runtime	31
3.4. USING A VISUAL STUDIO CODE EXTENSION IN CODEREADY WORKSPACES	32
3.4.1. Publishing a VS Code extension into the CodeReady Workspaces plug-in registry	33
3.4.1.1. Writing a meta.yaml file and adding it to a plug-in registry	33
3.4.2. Adding a plug-in registry VS Code extension to a workspace	34
3.4.2.1. Adding a VS Code extension using the CodeReady Workspaces Plugins panel	35
3.4.2.2. Adding a VS Code extension using the workspace configuration	35
3.5. TESTING A VISUAL STUDIO CODE EXTENSION IN CODEREADY WORKSPACES	36
3.5.1. Testing a VS Code extension using GitHub gist	36
3.5.2. Verifying the VS Code extension API compatibility level	40
<b>CHAPTER 4. RETRIEVING CODEREADY WORKSPACES LOGS</b> .....	<b>41</b>
4.1. ACCESSING OPENSIFT EVENTS ON OPENSIFT	41

4.2. VIEWING THE STATE OF THE CODEREADY WORKSPACES CLUSTER DEPLOYMENT USING OPENSIFT 4 CLI TOOLS	41
4.3. VIEWING CODEREADY WORKSPACES SERVER LOGS	42
4.3.1. Viewing the CodeReady Workspaces server logs using the OpenShift CLI	42
4.4. VIEWING EXTERNAL SERVICE LOGS	43
4.4.1. Viewing RH-SSO logs	43
4.4.1.1. Viewing the RH-SSO server logs	43
4.4.1.2. Viewing the RH-SSO client logs on Firefox	43
4.4.1.3. Viewing the RH-SSO client logs on Google Chrome	44
4.4.2. Viewing the CodeReady Workspaces database logs	44
4.5. VIEWING CODEREADY WORKSPACES WORKSPACES LOGS	44
4.5.1. Viewing Che-Theia IDE logs	44
4.5.1.1. Viewing Che-Theia editor logs using the OpenShift CLI	45
4.5.2. Viewing logs from language servers and debug adapters	46
4.5.2.1. Checking important logs	46
4.5.2.2. Detecting memory problems	46
4.5.2.3. Logging the client-server traffic for debug adapters	46
4.5.2.4. Viewing logs for Python	47
4.5.2.5. Viewing logs for Go	47
4.5.2.5.1. Finding the GOPATH	47
4.5.2.5.2. Viewing the Debug Console log for Go	47
4.5.2.5.3. Viewing the Go logs output in the Output panel	48
4.5.2.6. Viewing logs for the NodeDebug NodeDebug2 adapter	49
4.5.2.7. Viewing logs for Typescript	49
4.5.2.7.1. Enabling the label switched protocol (LSP) tracing	49
4.5.2.7.2. Viewing the Typescript language server log	49
4.5.2.7.3. Viewing the Typescript logs output in the Output panel	49
4.5.2.8. Viewing logs for Java	50
4.5.2.8.1. Verifying the state of the Eclipse JDT Language Server	50
4.5.2.8.2. Verifying the Eclipse JDT Language Server features	50
4.5.2.8.3. Viewing the Java language server log	50
4.5.2.8.4. Logging the Java language server protocol (LSP) messages	50
4.5.2.9. Viewing logs for Intelphense	51
4.5.2.9.1. Logging the Intelphense client-server communication	51
4.5.2.9.2. Viewing Intelphense events in the Output panel	51
4.5.2.10. Viewing logs for PHP-Debug	51
4.5.2.11. Viewing logs for XML	51
4.5.2.11.1. Verifying the state of the XML language server	51
4.5.2.11.2. Checking XML language server feature flags	52
4.5.2.11.3. Enabling XML Language Server Protocol (LSP) tracing	52
4.5.2.11.4. Viewing the XML language server log	52
4.5.2.12. Viewing logs for YAML	52
4.5.2.12.1. Verifying the state of the YAML language server	53
4.5.2.12.2. Checking the YAML language server feature flags	53
4.5.2.12.3. Enabling YAML Language Server Protocol (LSP) tracing	54
4.5.2.13. Viewing logs for Dotnet with Omnisharp-Theia plug-in	54
4.5.2.13.1. Omnisharp-Theia plug-in	54
4.5.2.13.2. Verifying the state of the Omnisharp-Theia plug-in language server	54
4.5.2.13.3. Checking Omnisharp Che-Theia plug-in language server features	54
4.5.2.13.4. Viewing Omnisharp-Theia plug-in logs in the Output panel	55
4.5.2.14. Viewing logs for Dotnet with NetcoredebugOutput plug-in	55
4.5.2.14.1. NetcoredebugOutput plug-in	55
4.5.2.14.2. Verifying the state of the NetcoredebugOutput plug-in	55

4.5.2.14.3. Viewing NetcoredebugOutput plug-in logs in the Output panel	55
4.5.2.15. Viewing logs for Camel	56
4.5.2.15.1. Verifying the state of the Camel language server	56
4.5.2.15.2. Viewing Camel logs in the Output panel	56
4.6. VIEWING THE PLUG-IN BROKER LOGS	57
4.7. COLLECTING LOGS USING CRWCTL	57
<b>CHAPTER 5. MONITORING CODEREADY WORKSPACES</b>	<b>58</b>
5.1. ENABLING AND EXPOSING CODEREADY WORKSPACES METRICS	58
5.2. COLLECTING CODEREADY WORKSPACES METRICS WITH PROMETHEUS	59
5.3. EXTENDING CODEREADY WORKSPACES MONITORING METRICS	61
<b>CHAPTER 6. TRACING CODEREADY WORKSPACES</b>	<b>62</b>
6.1. TRACING API	62
6.2. TRACING BACK END	62
6.3. INSTALLING THE JAEGER TRACING TOOL	62
6.3.1. Installing Jaeger using OperatorHub on OpenShift 4	62
6.3.2. Installing Jaeger using CLI on OpenShift 4	63
6.4. ENABLING METRICS COLLECTION	64
6.5. VIEWING CODEREADY WORKSPACES TRACES IN JAEGER UI	66
6.6. CODEREADY WORKSPACES TRACING CODEBASE OVERVIEW AND EXTENSION GUIDE	67
6.6.1. Tagging	67
<b>CHAPTER 7. MANAGING USERS</b>	<b>68</b>
7.1. CONFIGURING AUTHORIZATION	68
7.1.1. Authorization and user management	68
7.1.2. Configuring CodeReady Workspaces to work with RH-SSO	68
7.1.3. Configuring RH-SSO tokens	68
7.1.4. Setting up user federation	69
7.1.5. Enabling authentication with social accounts and brokering	69
7.1.6. Using protocol-based providers	70
7.1.7. Managing users using RH-SSO	71
7.1.8. Configuring CodeReady Workspaces to use an external RH-SSO installation	71
7.1.9. Configuring SMTP and email notifications	73
7.2. REMOVING USER DATA	73
7.2.1. Removing user data according to GDPR	73
<b>CHAPTER 8. SECURING CODEREADY WORKSPACES</b>	<b>77</b>
8.1. AUTHENTICATING USERS	77
8.1.1. Authenticating to the CodeReady Workspaces server	77
8.1.1.1. Authenticating to the CodeReady Workspaces server using OpenID	77
8.1.1.1.1. Obtaining the token from credentials through RH-SSO	79
8.1.1.1.2. Obtaining the token from the OpenShift token through RH-SSO	79
8.1.1.2. Authenticating to the CodeReady Workspaces server using other authentication implementations	80
8.1.1.3. Authenticating to the CodeReady Workspaces server using OAuth	80
8.1.1.4. Using Swagger or REST clients to execute queries	81
8.1.2. Authenticating in a CodeReady Workspaces workspace	81
8.1.2.1. Creating secure servers	82
8.1.2.2. Workspace JWT token	82
8.1.2.3. Machine token validation	83
8.2. AUTHORIZING USERS	83
8.2.1. CodeReady Workspaces workspace permissions	84
8.2.2. CodeReady Workspaces system permissions	84
8.2.3. manageSystem permission	85

8.2.4. monitorSystem permission	85
8.2.5. Listing CodeReady Workspaces permissions	86
8.2.6. Assigning CodeReady Workspaces permissions	86
8.2.7. Sharing CodeReady Workspaces permissions	87
<b>CHAPTER 9. BACKUP AND DISASTER RECOVERY</b> .....	<b>88</b>
9.1. EXTERNAL DATABASE SETUP	88
9.1.1. Configuring external PostgreSQL	88
9.1.2. Configuring CodeReady Workspaces to work with an external PostgreSQL	89
9.2. PERSISTENT VOLUMES BACKUPS	90
9.2.1. Recommended backup tool: Velero	91
<b>CHAPTER 10. CACHING IMAGES FOR FASTER WORKSPACE START</b> .....	<b>92</b>
10.1. IMAGE PULLER OVERVIEW	92
10.2. DEPLOYING IMAGE PULLER USING THE OPERATOR	93
10.2.1. Installing the Image Puller on OpenShift using OperatorHub	93
10.3. DEPLOYING IMAGE PULLER USING OPENSIFT TEMPLATES	94






# CHAPTER 1. CODEREADY WORKSPACES ARCHITECTURE OVERVIEW

Red Hat CodeReady Workspaces components are:

- A central workspace controller: an always running service that manages users workspaces through the OpenShift API.
- Users workspaces: container-based IDEs that the controller stops when the user stops coding.

Figure 1.1. High-level CodeReady Workspaces architecture



When CodeReady Workspaces is installed on an OpenShift cluster, the workspace controller is the only component that is deployed. A CodeReady Workspaces workspace is created immediately after a user requests it.

## Additional resources

- [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/administration\\_guide/index#codeready-workspaces-workspace-controller\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/administration_guide/index#codeready-workspaces-workspace-controller_crw)
- [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/administration\\_guide/index#codeready-workspaces-workspaces-architecture\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/administration_guide/index#codeready-workspaces-workspaces-architecture_crw)


## 1.1. CODEREADY WORKSPACES WORKSPACE CONTROLLER

The workspaces controller manages the container-based development environments: CodeReady Workspaces workspaces. It can be deployed in the following distinct configurations:

- **Single-user:** No authentication service is set up. Development environments are not secured. This configuration requires fewer resources. It is more adapted for local installations, such as when using Minikube.
- **Multi-user:** This is a multi-tenant configuration. Development environments are secured, and this configuration requires more resources. Appropriate for cloud installations.

The different services that are a part of the CodeReady Workspaces workspaces controller are shown in the following diagram. Note that RH-SSO and PostgreSQL are only needed in the multi-user configuration.

Figure 1.2. CodeReady Workspaces workspaces controller



### 1.1.1. CodeReady Workspaces server

The CodeReady Workspaces server, also known as **wsmaster**, is the central service of the workspaces controller. It is a Java web service that exposes an HTTP REST API to manage CodeReady Workspaces workspaces and, in multi-user mode, CodeReady Workspaces users.

Source code	<a href="#">Red Hat CodeReady Workspaces GitHub</a>
Container image	<b>eclipse/che-server</b>
Environment variables	<a href="https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/installation_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component_crw">https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/installation_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component_crw</a>

### 1.1.2. CodeReady Workspaces user dashboard

The user dashboard is the landing page of Red Hat CodeReady Workspaces. It is an Angular front-end application. CodeReady Workspaces users create, start, and manage CodeReady Workspaces workspaces from their browsers through the user dashboard.

Source code	<a href="#">CodeReady Workspaces Dashboard</a>
Container image	<b>eclipse/che-server</b>

### 1.1.3. Devfile registry

The CodeReady Workspaces devfile registry is a service that provides a list of CodeReady Workspaces stacks to create ready-to-use workspaces. This list of stacks is used in the **Dashboard → Create Workspace** window. The devfile registry runs in a container and can be deployed wherever the user dashboard can connect.

For more information about devfile registry customization, see the Customizing devfile registry section.

Source code	<a href="#">CodeReady Workspaces Devfile registry</a>
Container image	<b>quay.io/eclipse/che-devfile-registry</b>

#### 1.1.4. CodeReady Workspaces plug-in registry

The CodeReady Workspaces plug-in registry is a service that provides the list of plug-ins and editors for the CodeReady Workspaces workspaces. A devfile only references a plug-in that is published in a CodeReady Workspaces plug-in registry. It runs in a container and can be deployed wherever wsmaster connects.

For more information about plug-in registry customization, see [Section 3.1, “Building and running a custom registry image”](#).

Source code	<a href="#">CodeReady Workspaces plug-in registry</a>
Container image	<b>quay.io/eclipse/che-plugin-registry</b>

#### 1.1.5. CodeReady Workspaces and PostgreSQL

The PostgreSQL database is a prerequisite to configure CodeReady Workspaces in multi-user mode. The CodeReady Workspaces administrator can choose to connect CodeReady Workspaces to an existing PostgreSQL instance or let the CodeReady Workspaces deployment start a new dedicated PostgreSQL instance.

The CodeReady Workspaces server uses the database to persist user configurations (workspaces metadata, Git credentials). RH-SSO uses the database as its back end to persist user information.

Source code	<a href="#">CodeReady Workspaces Postgres</a>
Container image	<b>registry.redhat.io/rhel8/postgresql-96:1</b>

#### 1.1.6. CodeReady Workspaces and RH-SSO

RH-SSO is a prerequisite to configure CodeReady Workspaces in multi-user mode. The CodeReady Workspaces administrator can choose to connect CodeReady Workspaces to an existing RH-SSO instance or let the CodeReady Workspaces deployment start a new dedicated RH-SSO instance.

The CodeReady Workspaces server uses RH-SSO as an OpenID Connect (OIDC) provider to authenticate CodeReady Workspaces users and secure access to CodeReady Workspaces resources.

Source code	<a href="#">CodeReady Workspaces RH-SSO</a>
-------------	---

Container image	<code>registry.redhat.io/rh-ss0-7/ssh074-openshift-rhel8:7.4</code>
-----------------	---

## 1.2. CODEREADY WORKSPACES WORKSPACES ARCHITECTURE

A CodeReady Workspaces deployment on the cluster consists of the CodeReady Workspaces server component, a database for storing user profile and preferences, and a number of additional deployments hosting workspaces. The CodeReady Workspaces server orchestrates the creation of workspaces, which consist of a deployment containing the workspace containers and enabled plug-ins, plus related components, such as:


- ConfigMaps
- services
- endpoints
- ingresses/routes
- secrets
- PVs

The CodeReady Workspaces workspace is a web application. It is composed of microservices running in containers that provide all the services of a modern IDE (an editor, language auto-completion, debugging tools). The IDE services are deployed with the development tools, packaged in containers and user runtime applications, which are defined as OpenShift resources.

The source code of the projects of a CodeReady Workspaces workspace is persisted in an OpenShift **PersistentVolume**. Microservices run in containers that have read-write access to the source code (IDE services, development tools), and runtime applications have read-write access to this shared directory.

The following diagram shows the detailed components of a CodeReady Workspaces workspace.

Figure 1.3. CodeReady Workspaces workspace components



In the diagram, there are three running workspaces: two belonging to **User A** and one to **User C**. A fourth workspace is getting provisioned where the plug-in broker is verifying and completing the workspace configuration.

Use the devfile format to specify the tools and runtime applications of a CodeReady Workspaces workspace.

## 1.2.1. CodeReady Workspaces workspace components

This section describes the components of a CodeReady Workspaces workspace.

### 1.2.1.1. Che Editor plug-in

A **Che Editor** plug-in is a CodeReady Workspaces workspace plug-in. It defines the web application that is used as an editor in a workspace. The default CodeReady Workspaces workspace editor is Che-Theia.

The Che-Theia source-code repository is at [Che-Theia Github](#). It is based on the [Eclipse Theia open-source project](#).

Che-Theia is written in TypeScript and is built on the [Microsoft Monaco editor](#). It is a web-based source-code editor similar to [Visual Studio Code](#) (VS Code). It has a plug-in system that supports VS Code extensions.

Source code	<a href="#">Che-Theia</a>
Container image	<b>eclipse/che-theia</b>
Endpoints	<b>theia, webviews, theia-dev, theia-redirect-1, theia-redirect-2, theia-redirect-3</b>

### 1.2.1.2. CodeReady Workspaces user runtimes

Use any non-terminating user container as a user runtime. An application that can be defined as a container image or as a set of OpenShift resources can be included in a CodeReady Workspaces workspace. This makes it easy to test applications in the CodeReady Workspaces workspace.

To test an application in the CodeReady Workspaces workspace, include the application YAML definition used in stage or production in the workspace specification. It is a 12-factor app dev/prod parity.

Examples of user runtimes are Node.js, SpringBoot or MongoDB, and MySQL.

### 1.2.1.3. CodeReady Workspaces workspace JWT proxy

The JWT proxy is responsible for securing the communication of the CodeReady Workspaces workspace services. The CodeReady Workspaces workspace JWT proxy is included in a CodeReady Workspaces workspace only if the CodeReady Workspaces server is configured in multi-user mode.

An HTTP proxy is used to sign outgoing requests from a workspace service to the CodeReady Workspaces server and to authenticate incoming requests from the IDE client running on a browser.

Source code	<a href="#">JWT proxy</a>
Container image	<b>eclipse/che-jwtproxy</b>

### 1.2.1.4. CodeReady Workspaces plug-ins broker

Plug-in brokers are special services that, given a plug-in **meta.yaml** file:

- Gather all the information to provide a plug-in definition that the CodeReady Workspaces server knows.
- Perform preparation actions in the workspace project (download, unpack files, process configuration).

The main goal of the plug-in broker is to decouple the CodeReady Workspaces plug-ins definitions from the actual plug-ins that CodeReady Workspaces can support. With brokers, CodeReady Workspaces can support different plug-ins without updating the CodeReady Workspaces server.

The CodeReady Workspaces server starts the plug-in broker. The plug-in broker runs in the same OpenShift project as the workspace. It has access to the plug-ins and project persistent volumes.

A plug-ins broker is defined as a container image (for example, **eclipse/che-plugin-broker**). The plug-in type determines the type of the broker that is started. Two types of plug-ins are supported: **Che Plugin** and **Che Editor**.

Source code	<a href="#">CodeReady Workspaces Plug-in broker</a>
Container image	<b>quay.io/eclipse/che-plugin-artifacts-broker</b> <b>eclipse/che-plugin-metadata-broker</b>

## 1.2.2. CodeReady Workspaces workspace configuration

This section describes the properties of the CodeReady Workspaces server that affect the provisioning of a CodeReady Workspaces workspace.

### 1.2.2.1. Storage strategies for codeready-workspaces workspaces

Workspace Pods use Persistent Volume Claims (PVCs), which are bound to the physical Persistent Volumes (PVs) with [ReadWriteOnce access mode](#). It is possible to configure how the CodeReady Workspaces server uses PVCs for workspaces. The individual methods for this configuration are called PVC strategies:

strategy	details	pros	cons
<b>unique</b>	One PVC per workspace volume or user-defined PVC	Storage isolation	An undefined number of PVs is required
<b>per-workspace (default)</b>	One PVC for one workspace	Easier to manage and control storage compared to unique strategy	PV count still is not known and depends on workspaces number
<b>common</b>	One PVC for all workspaces in one OpenShift namespace	Easy to manage and control storage	If PV does not support ReadWriteMany (RWX) access mode then workspaces must be in a separate OpenShift namespaces  Or there must not be more than 1 running workspace per namespace at the same time  See how to configure <a href="#">namespace strategy</a>

Red Hat CodeReady Workspaces uses the **common** PVC strategy in combination with the "one project per user" project strategy when all CodeReady Workspaces workspaces operate in the user's project, sharing one PVC.



### 1.2.2.1.1. The **common** PVC strategy

All workspaces inside a OpenShift project use the same Persistent Volume Claim (PVC) as the default data storage when storing data such as the following in their declared volumes:

- projects
- workspace logs
- additional Volumes defined by a use

When the **common** PVC strategy is in use, user-defined PVCs are ignored and volumes that refer to these user-defined PVCs are replaced with a volume that refers to the common PVC. In this strategy, all CodeReady Workspaces workspaces use the same PVC. When the user runs one workspace, it only binds to one node in the cluster at a time.

The corresponding containers volume mounts link to a common volume, and sub-paths are prefixed with **<workspace-ID>** or **<original-PVC-name>**. For more details, see [Section 1.2.2.1.4, "How subpaths are used in PVCs"](#).

The CodeReady Workspaces Volume name is identical to the name of the user-defined PVC. It means that if a machine is configured to use a CodeReady Workspaces volume with the same name as the user-defined PVC has, they will use the same shared folder in the common PVC.

When a workspace is deleted, a corresponding subdirectory (**/\${ws-id}**) is deleted in the PV directory.

### Restrictions on using the **common** PVC strategy

When the **common** strategy is used and a workspace PVC access mode is ReadWriteOnce (RWO), only one node can simultaneously use the PVC.

If there are several nodes, you can use the **common** strategy, but:

- The workspace PVC access mode must be reconfigured to **ReadWriteMany** (RWM), so multiple nodes can use this PVC simultaneously.
- Only one workspace in the same project may be running. See [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/installation\\_guide/index#running-more-than-one-workspace-at-a-time\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/installation_guide/index#running-more-than-one-workspace-at-a-time_crw).

The **common** PVC strategy is not suitable for large multi-node clusters. Therefore, it is best to use it in single-node clusters. However, in combination with the **per-workspace** project strategy, the **common** PVC strategy is usable for clusters with not more than 75 nodes. The PVC used with this strategy must be large enough to accommodate all projects to prevent a situation in which one project depletes the resources of others.

### 1.2.2.1.2. The **per-workspace** PVC strategy

The **per-workspace** strategy is similar to the **common** PVC strategy. The only difference is that all workspace Volumes, but not all the workspaces, use the same PVC as the default data storage for:

- projects
- workspace logs
- additional Volumes defined by a user

With this strategy, CodeReady Workspaces keeps its workspace data in assigned PVs that are allocated by a single PVC.

The **per-workspace** PVC strategy is the most universal strategy out of the PVC strategies available and acts as a proper option for large multi-node clusters with a higher amount of users. Using the **per-workspace** PVC strategy, users can run multiple workspaces simultaneously, results in more PVCs being created.

### 1.2.2.1.3. The unique PVC strategy

When using the `unique` PVC strategy, every CodeReady Workspaces Volume of a workspace has its own PVC. This means that workspace PVCs are:

Created when a workspace starts for the first time. Deleted when a corresponding workspace is deleted.

User-defined PVCs are created with the following specifics:

- They are provisioned with generated names to prevent naming conflicts with other PVCs in a project.
- Subpaths of the mounted Physical persistent volumes that reference user-defined PVCs are prefixed with **<workspace-ID>** or **<PVC-name>**. This ensures that the same PV data structure is set up with different PVC strategies. For details, see [Section 1.2.2.1.4, “How subpaths are used in PVCs”](#).

The **unique** PVC strategy is suitable for larger multi-node clusters with a lesser amount of users. Since this strategy operates with separate PVCs for each volume in a workspace, vastly more PVCs are created.

### 1.2.2.1.4. How subpaths are used in PVCs

Subpaths illustrate the folder hierarchy in the Persistent Volumes (PV).

```

/pv0001
  /workspaceID1
  /workspaceID2
  /workspaceIDn
  /che-logs
  /projects
  /<volume1>
  /<volume2>
  /<User-defined PVC name 1 | volume 3>
  ...

```

When a user defines volumes for components in the devfile, all components that define the volume of the same name will be backed by the same directory in the PV as **<PV-name>**, **<workspace-ID>**, or **<original-PVC-name>**. Each component can have this location mounted on a different path in its containers.

#### Example

Using the **common** PVC strategy, user-defined PVCs are replaced with subpaths on the common PVC. When the user references a volume as **my-volume**, it is mounted in the common-pvc with the **/workspace-id/my-volume** subpath.

## 1.2.2.2. Configuring a CodeReady Workspaces workspace with a persistent volume strategy

A persistent volume (PV) acts as a virtual storage instance that adds a volume to a cluster.

A persistent volume claim (PVC) is a request to provision persistent storage of a specific type and configuration, available in the following CodeReady Workspaces storage configuration strategies:

- Common
- Per-workspace
- Unique

The mounted PVC is displayed as a folder in a container file system.

### 1.2.2.2.1. Configuring a PVC strategy using the Operator

The following section describes how to configure workspace persistent volume claim (PVC) strategies of a CodeReady Workspaces server using the Operator.



#### WARNING

It is not recommended to reconfigure PVC strategies on an existing CodeReady Workspaces cluster with existing workspaces. Doing so causes data loss.

[Operators](#) are software extensions to OpenShift that use [Custom Resources](#) to manage applications and their components.

When deploying CodeReady Workspaces using the Operator, configure the intended strategy by modifying the **spec.storage.pvcStrategy** property of the CheCluster Custom Resource object YAML file.

#### Prerequisites

- The **oc** tool is available.

#### Procedure

The following procedure steps are available for OpenShift command-line tool, `'oc'`.

To do changes to the CheCluster YAML file, choose one of the following:

- Create a new cluster by executing the **oc apply** command. For example:

```
$ oc apply -f <my-cluster.yaml>
```

- Update the YAML file properties of an already running cluster by executing the **oc patch** command. For example:

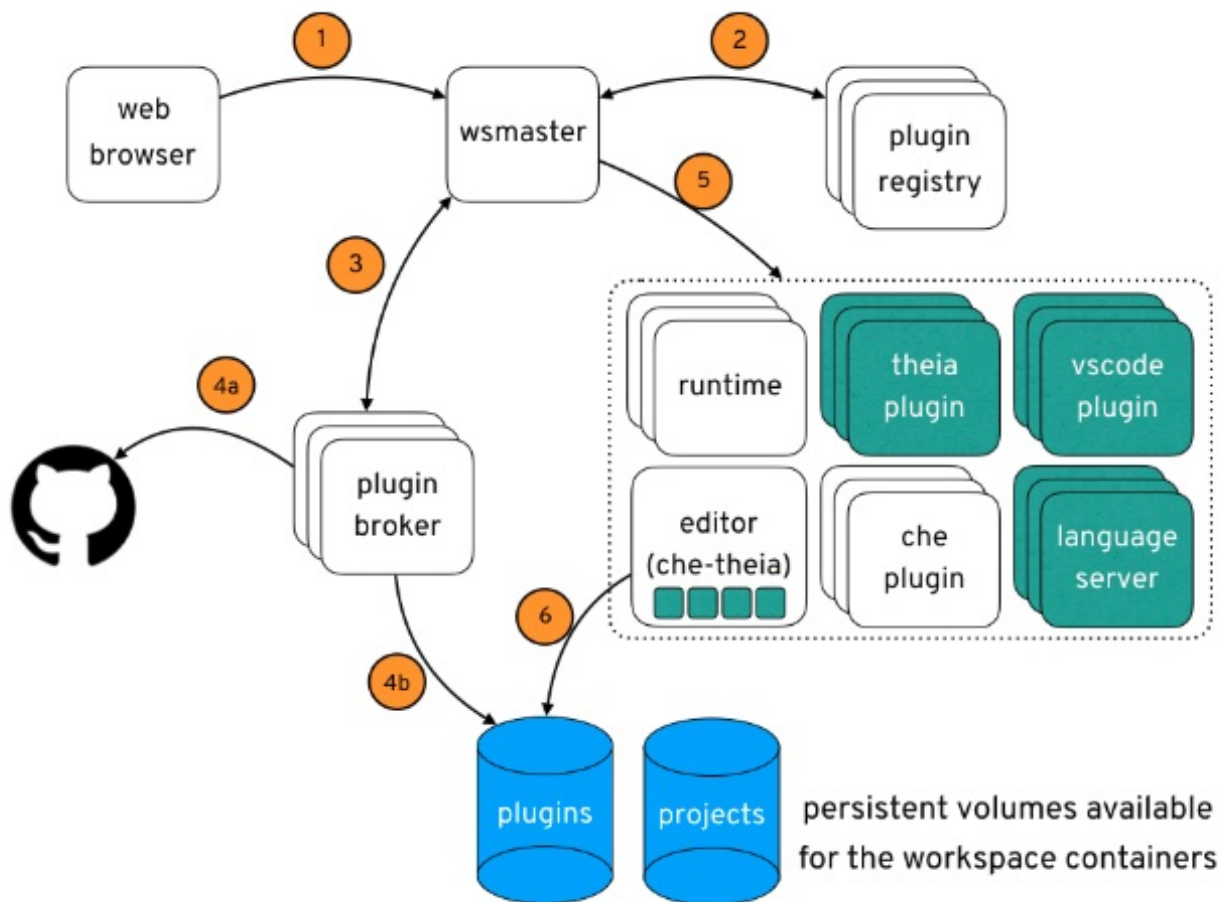
```
$ oc patch checluster codeready-workspaces --type=json \
-p [{"op": "replace", "path": "/spec/storage/pvcStrategy", "value": "<per-workspace>"}]
```

Depending on the strategy used, replace the **<per-workspace>** option in the above example with **unique** or **common**.

### 1.2.2.3. Workspace OpenShift project configuration

The OpenShift project where a new workspace Pod is deployed depends on the CodeReady Workspaces server configuration. By default, every workspace is deployed in a distinct OpenShift project, but the user can configure the CodeReady Workspaces server to deploy all workspaces in one specific OpenShift project. The name of a OpenShift project must be provided as a CodeReady Workspaces server configuration property and cannot be changed at runtime.

### 1.2.3. CodeReady Workspaces workspace creation flow



The following is a CodeReady Workspaces workspace creation flow:

1. A user starts a CodeReady Workspaces workspace defined by:
  - An editor (the default is Che-Theia)
  - A list of plug-ins (for example, Java and OpenShift tools)
  - A list of runtime applications
2. **wsmaster** retrieves the editor and plug-in metadata from the plug-in registry.
3. For every plug-in type, **wsmaster** starts a specific plug-in broker.

4. The CodeReady Workspaces plug-ins broker transforms the plug-in metadata into a Che Plugin definition. It executes the following steps:
  - a. Downloads a plug-in and extracts its content.
  - b. Processes the plug-in **meta.yaml** file and sends it back to **wsmaster** in the format of a Che Plugin.
5. **wsmaster** starts the editor and the plug-in sidecars.
6. The editor loads the plug-ins from the plug-in persistent volume.

## CHAPTER 2. CALCULATING CODEREADY WORKSPACES RESOURCE REQUIREMENTS

This section describes how to calculate resources (memory and CPU) required to run Red Hat CodeReady Workspaces.

Both the CodeReady Workspaces central controller and user workspaces consist of a set of containers. Those containers contribute to the resources consumption in terms of CPU and RAM limits and requests.

### 2.1. CONTROLLER REQUIREMENTS

The Workspace Controller consists of a set of five services running in five distinct containers. The following table presents the default resource requirements of each of these services.

**Table 2.1. ControllerServices**

Pod	Container name	Default memory limit	Default memory request
CodeReady Workspaces Server and Dashboard	che	1 GiB	512 MiB
PostgreSQL	postgres	1 GiB	512 MiB
RH-SSO	keycloak	2 GiB	512 MiB
Devfile registry	che-devfile-registry	256 MiB	16 MiB
Plug-in registry	che-plugin-registry	256 MiB	16 MiB

These default values are sufficient when the CodeReady Workspaces Workspace Controller manages a small amount of CodeReady Workspaces workspaces. For larger deployments, increase the memory limit. See the [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/installation\\_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/installation_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component_crw) article for instructions on how to override the default requests and limits. For example, the hosted version of CodeReady Workspaces that runs on <https://che.openshift.io> uses 1 GB of memory.

#### Additional resources

- [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/administration\\_guide/index#codeready-workspaces-workspace-controller\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/administration_guide/index#codeready-workspaces-workspace-controller_crw).

### 2.2. WORKSPACES REQUIREMENTS

This section describes how to calculate the resources required for a workspace. It is the sum of the resources required for each component of this workspace.

These examples demonstrate the necessity of a proper calculation:

- A workspace with 10 active plug-ins requires more resources than the same workspace with fewer plug-ins.
- A standard Java workspace requires more resources than a standard Node.js workspace because running builds, tests, and application debugging requires more resources.

## Procedure

1. Identify the workspace components explicitly specified in the **components** section of the [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/end-user\\_guide/index#making-a-workspace-portable-using-a-devfile\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/end-user_guide/index#making-a-workspace-portable-using-a-devfile_crw).
2. Identify the implicit workspace components:
  - a. CodeReady Workspaces implicitly loads the default **cheEditor**: **che-theia**, and the **chePlugin** that allows commands execution: **che-machine-exec-plugin**. To change the default editor, add a **cheEditor** component section in the devfile.
  - b. When CodeReady Workspaces is running in multiuser mode, it loads the **JWT Proxy** component. The JWT Proxy is responsible for the authentication and authorization of the external communications of the workspace components.
3. Calculate the requirements for each component:
  - a. Default values:  
The following table presents the default requirements for all workspace components. It also presents the corresponding CodeReady Workspaces server property to modify the defaults cluster-wide.

**Table 2.2. Default requirements of workspace components by type**

Component types	CodeReady Workspaces server property	Default memory limit	Default memory request
<b>chePlugin</b>	<b>che.workspace.sidebar.default_memory_limit_mb</b>	128 MiB	128 MiB
<b>cheEditor</b>	<b>che.workspace.sidebar.default_memory_limit_mb</b>	128 MiB	128 MiB
<b>kubernetes, openshift, dockerimage</b>	<b>che.workspace.default_memory_limit_mb, che.workspace.default_memory_request_mb</b>	1 Gi	512 MiB
<b>JWT Proxy</b>	<b>che.server.secure_exposer.jwtproxy.memory_limit</b>	128 MiB	128 MiB

- b. Custom requirements for **chePlugins** and **cheEditors** components:

## i. Custom memory limit and request:

If present, the **memoryLimit** and **memoryRequest** attributes of the **containers** section of the **meta.yaml** file define the memory limit of the **chePlugins** or **cheEditors** components. CodeReady Workspaces automatically sets the memory request to match the memory limit in case it is not specified explicitly.

**Example 2.1. The chePlugin che-incubator/typescript/latest****meta.yaml spec section:**

```
spec:
  containers:
    - image: docker.io/eclipse/che-remote-plugin-node:next
      name: vscode-typescript
      memoryLimit: 512Mi
      memoryRequest: 256Mi
```

It results in a container with the following memory limit and request:

Memory limit	512 MiB
Memory request	256 MiB

**TIP****How to find the meta.yaml file of chePlugin**

Community plug-ins are available in the [che-plugin-registry GitHub repository](#) in folder **v3/plugins/{organization}/{name}/{version}/**.

For non-community or customized plug-ins, the **meta.yaml** files are available on the local OpenShift cluster at

**`{pluginRegistryEndpoint}/v3/plugins/{organization}/{name}/{version}/meta.yaml`**.

For example, on a local Minikube cluster, the URL for the **che-incubator/typescript/latest meta.yaml** is **`http://plugin-registry-che.192.168.64.78.nip.io/v3/plugins/che-incubator/typescript/latest/meta.yaml`**.

## ii. Custom CPU limit and request:

CodeReady Workspaces does not set CPU limits and requests by default. However, it is possible to configure CPU limits for the **chePlugin** and **cheEditor** types in the **meta.yaml** file or in the devfile in the same way as it done for memory limits.

**Example 2.2. The chePlugin che-incubator/typescript/latest****meta.yaml spec section:**

```
spec:
  containers:
    - image: docker.io/eclipse/che-remote-plugin-node:next
```



```
name: vscode-typescript
cpuLimit: 2000m
cpuRequest: 500m
```

It results in a container with the following CPU limit and request:

CPU limit	2 cores
CPU request	0.5 cores

To set CPU limits and requests globally, use the following dedicated environment variables:

<b>CPU Limit</b>	<b>CHE_WORKSPACE_SIDECAR_DEFAULT_CPU_LIMIT_CORES</b>
<b>CPU Request</b>	<b>CHE_WORKSPACE_SIDECAR_DEFAULT_CPU_REQUEST_CORES</b>

See also [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/installation\\_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/installation_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component_crw).

Note that the **LimitRange** object of the OpenShift project may specify defaults for CPU limits and requests set by cluster administrators. To prevent start errors due to resources overrun, limits on application or workspace levels must comply with those settings.

- a. Custom requirements for **dockerimage** components  
If present, the **memoryLimit** and **memoryRequest** attributes of the devfile define the memory limit of a **dockerimage** container. CodeReady Workspaces automatically sets the memory request to match the memory limit in case it is not specified explicitly.

```
- alias: maven
  type: dockerimage
  image: eclipse/maven-jdk8:latest
  memoryLimit: 1536M
```

- b. Custom requirements for **kubernetes** or **openshift** components:  
The referenced manifest may define the memory requirements and limits.
  1. Add all requirements previously calculated.

### Additional resources

- [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/administration\\_guide/index#codeready-workspaces-workspaces-architecture\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/administration_guide/index#codeready-workspaces-workspaces-architecture_crw).

## 2.3. A WORKSPACE EXAMPLE

This section describes a CodeReady Workspaces workspace example.

The following devfile defines the CodeReady Workspaces workspace:

```

apiVersion: 1.0.0
metadata:
  generateName: guestbook-nodejs-sample-
projects:
  - name: guestbook-nodejs-sample
    source:
      type: git
      location: "https://github.com/l0rd/nodejs-sample"
components:
  - type: chePlugin
    id: che-incubator/typescript/latest
  - type: kubernetes
    alias: guestbook-frontend
    reference: https://raw.githubusercontent.com/l0rd/nodejs-sample/master/kubernetes-manifests/guestbook-frontend.deployment.yaml
    mountSources: true
  entrypoints:
    - command: ['sleep']
      args: ['infinity']

```

This table provides the memory requirements for each workspace component:

**Table 2.3. Total workspace memory requirement and limit**

Pod	Container name	Default memory limit	Default memory request
Workspace	theia-ide (default <b>cheEditor</b> )	512 MiB	512 MiB
Workspace	machine-exec (default <b>chePlugin</b> )	128 MiB	128 MiB
Workspace	vscode-typescript ( <b>chePlugin</b> )	512 MiB	512 MiB
Workspace	frontend ( <b>kubernetes</b> )	1 GiB	512 MiB
JWT Proxy	verifier	128 MiB	128 MiB
<b>Total</b>		<b>2.25 GiB</b>	<b>1.75 GiB</b>

- The **theia-ide** and **machine-exec** components are implicitly added to the workspace, even when not included in the devfile.
- The resources required by **machine-exec** are the default for **chePlugin**.
- The resources for **theia-ide** are specifically set in the **cheEditor meta.yaml** to **512 MiB** as **memoryLimit**.
- The Typescript VS Code extension has also overridden the default memory limits. In its **meta.yaml** file, the limits are explicitly specified to **512 MiB**.

- CodeReady Workspaces is applying the defaults for the **kubernetes** component type: a memory limit of **1 GiB** and a memory request of **512 MiB**. This is because the **kubernetes** component references a **Deployment** manifest that has a container specification with no resource limits or requests.
- The JWT container requires **128 MiB** of memory.

Adding all together results in **1.75 GiB** of memory requests with a **2.25 GiB** limit.

### Additional resources

- [Chapter 1, \*CodeReady Workspaces architecture overview\*](#)
- [Kubernetes compute resources management documentation](#)
- [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/installation\\_guide/index#configuring-the-codeready-workspaces-installation\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/installation_guide/index#configuring-the-codeready-workspaces-installation_crw)
- [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/installation\\_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/installation_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component_crw)
- [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/end-user\\_guide/index#making-a-workspace-portable-using-a-devfile\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/end-user_guide/index#making-a-workspace-portable-using-a-devfile_crw)
- [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/end-user\\_guide/index#making-a-workspace-portable-using-a-devfile\\_crw#a-minimal-devfile\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/end-user_guide/index#making-a-workspace-portable-using-a-devfile_crw#a-minimal-devfile_crw)
- [Section 8.1, "Authenticating users"](#)
- [che-plugin-registry GitHub repository](#)

## CHAPTER 3. CUSTOMIZING THE DEVFILE AND PLUG-IN REGISTRIES

CodeReady Workspaces 2.4 introduces two registries: the plug-ins registry and the devfile registry. They are static websites where the metadata of CodeReady Workspaces plug-ins and CodeReady Workspaces devfiles is published.

The plug-in registry makes it possible to share a plug-in definition across all the users of the same instance of CodeReady Workspaces. Only plug-ins that are published in a registry can be used in a devfile.

The devfile registry holds the definitions of the CodeReady Workspaces stacks. These are available on the CodeReady Workspaces user dashboard when selecting **Create Workspace**. It contains the list of CodeReady Workspaces technological stack samples with example projects.

The devfile and plug-in registries run in two separate pods and are deployed when the CodeReady Workspaces server is deployed (that is the default behavior of the CodeReady Workspaces Operator). The metadata of the plug-ins and devfiles are versioned on GitHub and follow the CodeReady Workspaces server life cycle.

In this document, the following two ways to customize the default registries that are deployed with CodeReady Workspaces (to modify the plug-ins or devfile metadata) are described:

1. Building a custom image of the registries
2. Running the default images but modifying them at runtime
  - [Section 3.1, "Building and running a custom registry image"](#)
  - [Section 3.2, "Including the plug-in binaries in the registry image"](#)
  - [Section 3.3, "Editing a devfile and plug-in at runtime"](#)

### 3.1. BUILDING AND RUNNING A CUSTOM REGISTRY IMAGE

This section describes the building of registries and updating a running CodeReady Workspaces server to point to the registries.

#### 3.1.1. Building a custom devfile registry

This section describes how to build a custom devfiles registry. Following operations are covered:

- Getting a copy of the source code necessary to build a devfiles registry.
- Adding a new devfile.
- Building the devfiles registry.

#### Procedure

1. Clone the devfile registry repository:

```
$ git clone git@github.com:redhat-developer/codeready-workspaces.git
$ cd codeready-workspaces/dependencies/che-devfile-registry
```

- In the `./che-devfile-registry/devfiles/` directory, create a subdirectory `<devfile-name>/` and add the `devfile.yaml` and `meta.yaml` files.

### File organization for a devfile

```
./che-devfile-registry/devfiles/
├── <devfile-name>
│   ├── devfile.yaml
│   └── meta.yaml
```

- Add valid content in the `devfile.yaml` file. For a detailed description of the devfile format, see [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/end-user\\_guide/index#making-a-workspace-portable-using-a-devfile\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/end-user_guide/index#making-a-workspace-portable-using-a-devfile_crw).
- Ensure that the `meta.yaml` file conforms to the following structure:

**Table 3.1. Parameters for a devfilemeta.yaml**

Attribute	Description
<b>description</b>	Description as it appears on the user dashboard.
<b>displayName</b>	Name as it appears on the user dashboard.
<b>globalMemoryLimit</b>	The sum of the expected memory consumed by all the components launched by the devfile. This number will be visible on the user dashboard. It is informative and is not taken into account by the CodeReady Workspaces server.
<b>icon</b>	Link to an <code>.svg</code> file that is displayed on the user dashboard.
<b>tags</b>	List of tags. Tags usually include the tools included in the stack.

### Example devfile meta.yaml

```
displayName: Rust
description: Rust Stack with Rust 1.39
tags: ["Rust"]
icon: https://www.eclipse.org/che/images/logo-eclipseche.svg
globalMemoryLimit: 1686Mi
```

- Build the containers for the custom devfile registry:

```
$ docker build -t my-devfile-registry .
```

## 3.1.2. Building a custom plug-in registry

This section describes how to build a custom plug-in registry. Following operations are covered:

- Getting a copy of the source code necessary to build a custom plug-in registry.
- Adding a new plug-in.

- Building the custom plug-in registry.

## Procedure

1. Clone the plug-in registry repository:

```
$ git clone git@github.com:redhat-developer/codeready-workspaces.git
$ cd codeready-workspaces/dependencies/che-plugin-registry
```

2. In the `./che-plugin-registry/v3/plugins/` directory, create new directories `<publisher>/<plugin-name>/<plugin-version>/` and a `meta.yaml` file in the last directory.

### File organization for a plugin

```
./che-plugin-registry/v3/plugins/
├── <publisher>
│   └── <plugin-name>
│       ├── <plugin-version>
│       │   ├── meta.yaml
│       │   └── latest.txt
│       └── latest.txt
```

3. Add valid content to the `meta.yaml` file. See the “Using a Visual Studio Code extension in CodeReady Workspaces” section or the README.md file in the `eclipse/che-plugin-registry` repository for a detailed description of the `meta.yaml` file format.
4. Create a file named `latest.txt` with content the name of the latest `<plugin-version>` directory.

### EXAMPLE

```
$ tree che-plugin-registry/v3/plugins/redhat/java/
che-plugin-registry/v3/plugins/redhat/java/
├── 0.38.0
│   └── meta.yaml
├── 0.43.0
│   └── meta.yaml
├── 0.45.0
│   └── meta.yaml
├── 0.46.0
│   └── meta.yaml
├── 0.50.0
│   └── meta.yaml
└── latest.txt
$ cat che-plugin-registry/v3/plugins/redhat/java/latest.txt
0.50.0
```

5. Build the containers for the custom plug-in registry:

```
./build.sh
```

## 3.1.3. Deploying the registries

### Prerequisites

The **my-plugin-in-registry** and **my-devfile-registry** images used in this section are built using the **docker** command. This section assumes that these images are available on the OpenShift cluster where CodeReady Workspaces is deployed.

This is true on Minikube, for example, if before running the **docker build** commands, the user executed the **eval  $\{\text{minikube docker-env}\}$**  command (or, the **eval  $\{\text{minishift docker-env}\}$**  command for Minishift).

Otherwise, these images can be pushed to a container registry (public, such as **quay.io**, or the DockerHub, or a private registry).

### 3.1.3.1. Deploying registries in OpenShift

#### Procedure

An OpenShift template to deploy the plug-in registry is available in the **openshift/** directory of the GitHub repository.

1. To deploy the plug-in registry using the OpenShift template, run the following command:

```

NAMESPACE=<namespace-name> 1
IMAGE_NAME="my-plugin-in-registry"
IMAGE_TAG="latest"
oc new-app -f openshift/che-plugin-registry.yml \
-n "${NAMESPACE}" \
-p IMAGE="${IMAGE_NAME}" \
-p IMAGE_TAG="${IMAGE_TAG}" \
-p PULL_POLICY="IfNotPresent"

```

- 1 If installed using **crwctl**, the default CodeReady Workspaces project is **openshift-workspaces**. The OperatorHub installation method deploys CodeReady Workspaces to the users current project.

2. The devfile registry has an OpenShift template in the **deploy/openshift/** directory of the GitHub repository. To deploy it, run the command:

```

NAMESPACE=<namespace-name> 1
IMAGE_NAME="my-devfile-registry"
IMAGE_TAG="latest"
oc new-app -f openshift/che-devfile-registry.yml \
-n "${NAMESPACE}" \
-p IMAGE="${IMAGE_NAME}" \
-p IMAGE_TAG="${IMAGE_TAG}" \
-p PULL_POLICY="IfNotPresent"

```

- 1 If installed using **crwctl**, the default CodeReady Workspaces project is **openshift-workspaces**. The OperatorHub installation method deploys CodeReady Workspaces to the users current project.

3. Check if the registries are deployed successfully on OpenShift.
  - a. To verify that the new plug-in is correctly published to the plug-in registry, make a request to the registry path **/v3/plugins/index.json** (or **/devfiles/index.json** for the devfile registry).

```

$ URL=$(oc get -o 'custom-columns=URL:.spec.rules[0].host' \
  -l app=che-plugin-registry route --no-headers)
$ INDEX_JSON=$(curl -sSL http://${URL}/v3/plugins/index.json)
$ echo ${INDEX_JSON} | grep -A 4 -B 5 "\"name\": \"my-plug-in\""
,\{
  "id": "my-org/my-plug-in/1.0.0",
  "displayName": "This is my first plug-in for CodeReady Workspaces",
  "version": "1.0.0",
  "type": "VS Code extension",
  "name": "my-plug-in",
  "description": "This plugin shows that we are able to add plugins to the registry",
  "publisher": "my-org",
  "links": \{"self": "/v3/plugins/my-org/my-plug-in/1.0.0" }
}
--
--
,\{
  "id": "my-org/my-plug-in/latest",
  "displayName": "This is my first plug-in for CodeReady Workspaces",
  "version": "latest",
  "type": "VS Code extension",
  "name": "my-plug-in",
  "description": "This plugin shows that we are able to add plugins to the registry",
  "publisher": "my-org",
  "links": \{"self": "/v3/plugins/my-org/my-plug-in/latest" }
}

```

- b. Verify that the CodeReady Workspaces server points to the URL of the registry. To do this, compare the value of the **CHE\_WORKSPACE\_PLUGIN\_REGISTRY\_\_URL** parameter in the **che** ConfigMap (or **CHE\_WORKSPACE\_DEVFILE\_\_REGISTRY\_\_URL** for the devfile registry):

```

$ oc get \
  -o "custom-columns=URL:.data['CHE_WORKSPACE_PLUGINREGISTRYURL']" \
  --no-headers cm/che
URL
http://che-plugin-registry-che.192.168.99.100.nip.io/v3

```

with the URL of the route:

```

$ oc get -o 'custom-columns=URL:.spec.rules[0].host' \
  -l app=che-plugin-registry route --no-headers
che-plugin-registry-che.192.168.99.100.nip.io

```

- c. If they do not match, update the ConfigMap and restart the CodeReady Workspaces server.

```

$ oc edit cm/che
(...)
$ oc scale --replicas=0 deployment/che
$ oc scale --replicas=1 deployment/che

```

When the new registries are deployed and the CodeReady Workspaces server is configured to use them, the new plug-ins are available in the **Plugin** view of a workspace and the new stacks are displayed in the **New Workspace** tab of the user dashboard.



## 3.2. INCLUDING THE PLUG-IN BINARIES IN THE REGISTRY IMAGE

The plug-in registry of CodeReady Workspaces differs from the Eclipse Che version. The Eclipse Che only hosts plug-in metadata, but the CodeReady Workspaces plug-in registry also hosts the corresponding binaries, and it is built in an offline mode by default. This means the binaries are already hosted in the plug-in-registry image.

This section describes how to add a new plug-in or reference a different version of a plug-in. This is achieved by modifying the plug-in **meta.yaml** file to point to a new plug-in and building a new registry in offline mode that contains the modified plug-in **meta.yaml** file and the plug-in binary file.

### Prerequisites

- An instance of CodeReady Workspaces is available.
- The **oc** tool is available.

### Procedure

1. Clone the **codeready-workspaces** repository

```
$ git clone https://github.com/redhat-developer/codeready-workspaces
$ cd codeready-workspaces/dependencies/che-plugin-registry
```

2. Identify the binaries you wish to change in the plug-in registry

The **meta.yaml** file includes the **extension** section, which defines the URLs of required extensions for the plug-in. For example, the **redhat/java11/0.63.0** plug-in lists the following two extensions:

#### meta.yaml

```
extensions:
  - https://download.jboss.org/jbosstools/vscode/3rdparty/vscode-java-debug/vscode-java-debug-0.26.0.vsix
  - https://download.jboss.org/jbosstools/static/jdt.ls/stable/java-0.63.0-2222.vsix
```

Change the first extension to reference the version hosted on GitHub and rebuild the plug-in registry. When using the **redhat/java11/0.63.0** plug-in, the binary will be fetched from the custom plug-in-registry server. Set the following environment variables to help with the subsequent commands:

```
ORG=redhat
NAME=java11
CHE_PLUGIN_VERSION=0.63.0
VSCODE_JAVA_DEBUG_VERSION=0.26.0
VSCODE_JAVA_DEBUG_URL="https://github.com/microsoft/vscode-java-debug/releases/download/0.26.0/vscjava.vsix"
OLD_JAVA_DEBUG_META_YAML_URL="https://download.jboss.org/jbosstools/vscode/3rdparty/vscode-java-debug/vscode-java-debug-0.26.0.vsix"
```


3. Get the plug-in registry URL:

```
$ oc get route plugin-registry -o jsonpath='{.spec.host}' -n ${CHE_NAMESPACE}
```

Save this value in a variable called **PLUGIN\_REGISTRY\_URL**.

- Update the URLs in the **meta.yaml** file to point to the VS Code extension binaries that are saved in the registry container:

```
$ sed -i -e
"s#${OLD_JAVA_DEBUG_META_YAML_URL}#${VSCODE_JAVA_DEBUG_URL}#g" \
.v3/plugins/${ORG}/${NAME}/${CHE_PLUGIN_VERSION}/meta.yaml
.v3/plugins/${ORG}/${NAME}/${CHE_PLUGIN_VERSION}/meta.yaml
```



### IMPORTANT

By default, CodeReady Workspaces is deployed with TLS enabled. For installations that do not use TLS, use **http://** in the **NEW\_JAVA\_DEBUG\_URL** and **NEW\_JAVA\_LS\_URL** variables.

- Confirm that the **meta.yaml** has the correctly substituted URLs:

```
$ cat .v3/plugins/${ORG}/${NAME}/${CHE_PLUGIN_VERSION}/meta.yaml
```

#### meta.yaml

```
extensions:
- https://plugin-registry-che.apps-crc.testing/v3/plugins/redhat/java11/0.63.0/vscode-java-
debug-0.26.0.vsix
- https://plugin-registry-che.apps-crc.testing/v3/plugins/redhat/java11/0.63.0/java-0.63.0-
2222.vsix
```

- Build and deploy the plug-in registry using the instructions in the [Section 3.1, "Building and running a custom registry image"](#) section.

## 3.3. EDITING A DEVFILE AND PLUG-IN AT RUNTIME

An alternative to building a custom registry image is to:

- Start a registry
- Modify its content at runtime

This approach is simpler and faster. But the modifications are lost as soon as the container is deleted.

### 3.3.1. Adding a plug-in at runtime

#### Procedure

To add a plug-in:

- Check out the plugin registry sources.

```
$ git clone https://github.com/redhat-developer/codeready-workspaces; \
cd codeready-workspaces/dependencies/che-plugin-registry
```

2. Create a **meta.yaml** in some local folder. This can be done from scratch or by copying from an existing plug-in's **meta.yaml** file.

```
$ PLUGIN="v3/plugins/new-org/new-plugin/0.0.1"; \
  mkdir -p ${PLUGIN}; cp v3/plugins/che-incubator/cpptools/0.1/* ${PLUGIN}/
  echo "${PLUGIN##*/}" > ${PLUGIN}/../latest.txt
```

3. If copying from an existing plug-in, make changes to the **meta.yaml** file to suit your needs. Make sure the new plug-in has a unique **title**, **displayName** and **description**. Update the **firstPublicationDate** to today's date.
4. These fields in **meta.yaml** must match the path defined in **PLUGIN** above.

```
publisher: new-org
name: new-plugin
version: 0.0.1
```

5. Get the name of the Pod that hosts the plug-in registry container. To do this, filter the **component=plugin-registry** label:

```
$ PLUGIN_REG_POD=$(oc get -o custom-columns=NAME:.metadata.name \
  --no-headers pod -l component=plugin-registry)
```

6. Regenerate the registry's **index.json** file to include the new plug-in.

```
$ cd codeready-workspaces/dependencies/che-plugin-registry; \
  "$(pwd)/build/scripts/generate_latest metas.sh" v3 && \
  "$(pwd)/build/scripts/check_plugins_location.sh" v3 && \
  "$(pwd)/build/scripts/set_plugin_dates.sh" v3 && \
  "$(pwd)/build/scripts/check_plugins_viewer_mandatory_fields.sh" v3 && \
  "$(pwd)/build/scripts/index.sh" v3 > v3/plugins/index.json
```

7. Copy the new **index.json** and **meta.yaml** files from the new local plug-in folder to the container.

```
$ cd codeready-workspaces/dependencies/che-plugin-registry; \
  LOCAL_FILES="$(pwd)/${PLUGIN}/meta.yaml $(pwd)/v3/plugins/index.json"; \
  oc exec ${PLUGIN_REG_POD} -i -t -- mkdir -p /var/www/html/${PLUGIN}; \
  for f in $LOCAL_FILES; do e=${f/$(pwd)/}; echo "Upload ${f} -> /var/www/html/${e}"; \
  oc cp "${f}" ${PLUGIN_REG_POD}:/var/www/html/${e}; done
```

8. The new plug-in can now be used from the existing CodeReady Workspaces instance of the plug-in registry. To discover it, go to the CodeReady Workspaces dashboard, then click the **Workspaces** link. From there, click the gear icon to configure one of your workspaces. Select the **Plugins** tab to see the updated list of available plug-ins.

### 3.3.2. Adding a devfile at runtime

#### Procedure

To add a devfile:

1. Check out the devfile registry sources.

```
$ git clone https://github.com/redhat-developer/codeready-workspaces; \
cd codeready-workspaces/dependencies/che-devfile-registry
```

2. Create a **devfile.yaml** and **meta.yaml** in some local folder. This can be done from scratch or by copying from an existing devfile.

```
$ STACK="new-stack"; \
mkdir -p devfiles/${STACK}; cp devfiles/03_web-nodejs-simple/* devfiles/${STACK}/
```

3. If copying from an existing devfile, make changes to the devfile to suit your needs. Make sure the new devfile has a unique **displayName** and **description**.
4. Get the name of the Pod that hosts the devfile registry container. To do this, filter the **component=devfile-registry** label:

```
$ DEVFILE_REG_POD=$(oc get -o custom-columns=NAME:.metadata.name \
--no-headers pod -l component=devfile-registry)
```

5. Regenerate the registry's **index.json** file to include the new devfile.

```
$ cd codeready-workspaces/dependencies/che-devfile-registry; \
"${(pwd)}/build/scripts/check_mandatory_fields.sh" devfiles; \
"${(pwd)}/build/scripts/index.sh" > index.json
```

6. Copy the new **index.json**, **devfile.yaml** and **meta.yaml** files from the new local devfile folder to the container.

```
$ cd che-devfile-registry; \
oc exec ${DEVFILE_REG_POD} -i -t -- mkdir -p /var/www/html/devfiles/${STACK}; \
oc cp $(pwd)/devfiles/${STACK}/meta.yaml \
${DEVFILE_REG_POD}:/var/www/html/devfiles/${STACK}/meta.yaml; \
oc cp $(pwd)/devfiles/${STACK}/devfile.yaml \
${DEVFILE_REG_POD}:/var/www/html/devfiles/${STACK}/devfile.yaml; \
oc cp $(pwd)/index.json ${DEVFILE_REG_POD}:/var/www/html/devfiles/index.json
```

7. The new devfile can now be used from the existing CodeReady Workspaces instance of the devfile registry. To discover it, go to the CodeReady Workspaces dashboard, then click the **Workspaces** link. From there, click **Add Workspace** to see the updated list of available devfiles.

### 3.4. USING A VISUAL STUDIO CODE EXTENSION IN CODEREADY WORKSPACES

In Red Hat CodeReady Workspaces, Visual Studio Code (VS Code) extensions can be installed to extend the functionality of a CodeReady Workspaces workspace. VS Code extensions can run in the Che-Theia editor container, or they can be packaged in their own isolated and pre-configured containers with their prerequisites.

This document describes:

- Use of a VS Code extension in CodeReady Workspaces with workspaces.
- CodeReady Workspaces Plug-ins panel.

- How to publish a VS Code extension in the CodeReady Workspaces plug-in registry (to share the extension with other CodeReady Workspaces users).
  - The extension-hosting sidecar container and the use of the extension in a devfile are optional for this.
  - How to review the compatibility of the VS Code extensions to be informed whether a specific API is supported or has not been implemented yet.

### 3.4.1. Publishing a VS Code extension into the CodeReady Workspaces plug-in registry

The user of CodeReady Workspaces can use a workspace devfile to use any plug-in, also known as Visual Studio Code (VS Code) extension. This plug-in can be added to the plug-in registry, then easily reused by anyone in the same organization with access to that workspace installation.

Some plug-ins need a runtime dedicated container for code compilation. This fact makes those plug-ins a combination of a runtime sidecar container and a VS Code extension.

The following section describes the portability of a plug-in configuration and associating an extension with a runtime container that the plug-in needs.

#### 3.4.1.1. Writing a meta.yaml file and adding it to a plug-in registry

The plug-in meta information is required to publish a VS Code extension in an Red Hat CodeReady Workspaces plug-in registry. This meta information is provided as a **meta.yaml** file. This section describes how to create a **meta.yaml** file for an extension.

#### Procedure

1. Create a **meta.yaml** file in the following plug-in registry directory:  
**<apiVersion>/plugins/<publisher>/<plug-inName>/<plug-inVersion>/**.
2. Edit the **meta.yaml** file and provide the necessary information. The configuration file must adhere to the following structure:

```

apiVersion: v2
publisher: myorg
name: my-vscode-ext
version: 1.7.2
type: value
displayName:
title:
description:
icon: https://www.eclipse.org/che/images/logo-eclipseche.svg
repository:
category:
spec:
  containers:
    - image:
      memoryLimit:
      memoryRequest:
      cpuLimit:

```

```

cpuRequest:
extensions:
  - https://github.com/redhat-developer/vscode-
yaml/releases/download/0.4.0/redhat.vscode-yaml-0.4.0.vsix
  - https://github.com/SonarSource/sonarlint-vscode/releases/download/1.16.0/sonarlint-
vscode-1.16.0.vsix

```

- 1 Version of the file structure.
- 2 Name of the plug-in publisher. Must be the same as the publisher in the path.
- 3 Name of the plug-in. Must be the same as in path.
- 4 Version of the plug-in. Must be the same as in path.
- 5 Type of the plug-in. Possible values: **Che Plugin, Che Editor, Theia plugin, VS Code extension**.
- 6 A short name of the plug-in.
- 7 Title of the plug-in.
- 8 A brief explanation of the plug-in and what it does.
- 9 The link to the plug-in logo.
- 10 Optional. The link to the source-code repository of the plug-in.
- 11 Defines the category that this plug-in belongs to. Should be one of the following: **Editor, Debugger, Formatter, Language, Linter, Snippet, Theme, or Other**.
- 12 If this section is omitted, the VS Code extension is added into the Che-Theia IDE container.
- 13 The Docker image from which the sidecar container will be started. Example: **theia-endpoint-image**.
- 14 The maximum RAM which is available for the sidecar container. Example: "512Mi". This value might be overridden by the user in the component configuration.
- 15 The RAM which is given for the sidecar container by default. Example: "256Mi". This value might be overridden by the user in the component configuration.
- 16 The maximum CPU amount in cores or millicores (suffixed with "m") which is available for the sidecar container. Examples: "500m", "2". This value might be overridden by the user in the component configuration.
- 17 The CPU amount in cores or millicores (suffixed with "m") which is given for the sidecar container by default. Example: "125m". This value might be overridden by the user in the component configuration.
- 18 A list of VS Code extensions run in this sidecar container.

### 3.4.2. Adding a plug-in registry VS Code extension to a workspace

When the required VS Code extension is added into a CodeReady Workspaces plug-in registry, the user can add it into the workspace through the **CodeReady Workspaces Plugins** panel or through the workspace configuration.

### 3.4.2.1. Adding a VS Code extension using the CodeReady Workspaces Plugins panel

#### Prerequisites

- A running instance of Red Hat CodeReady Workspaces. To install an instance of Red Hat CodeReady Workspaces, see [Installing CodeReady Workspaces on OpenShift Container Platform](#).

#### Procedure

To add a VS Code extension using the **CodeReady Workspaces Plugins** panel:

1. Open the **CodeReady Workspaces Plugins** panel by pressing **CTRL+SHIFT+J** or navigate to **View/Plugins**.
2. Change the current registry to the registry in which the VS Code extension was added.
3. In the search bar, click the **Menu** button and then click **Change Registry** to choose the registry from the list. If the required registry is not in the list, add it using the **Add Registry** menu option. The registry link points to the **plugins** segment of the registry, for example: **https://my-registry.com/v3/plugins/index.json**.
4. To update the list of plug-ins after adding a new registry link, use **Refresh** command from the search bar menu.
5. Search for the required plug-in using the filter, and then click the **Install** button.
6. Restart the workspace for the changes to take effect.

### 3.4.2.2. Adding a VS Code extension using the workspace configuration

#### Prerequisites

- A running instance of Red Hat CodeReady Workspaces. To install an instance of Red Hat CodeReady Workspaces, see [Installing CodeReady Workspaces on OpenShift Container Platform](#).
- An existing workspace defined on this instance of Red Hat CodeReady Workspaces [https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/end-user\\_guide/index#creating-and-configuring-a-new-workspace\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/end-user_guide/index#creating-and-configuring-a-new-workspace_crw).

#### Procedure

To add a VS Code extension using the workspace configuration:

1. Click the **Workspaces** tab on the **Dashboard** and select the workspace in which you want to add the plug-in. The **Workspace <workspace-name>** window is opened showing the details of the workspace.
2. Click the **devfile** tab.
3. Locate the **components** section, and add a new entry with the following structure:
  -

```
- type: chePlugin
  id: 1
```

1 ID format: `<publisher>/<plug-inName>/<plug-inVersion>`

CodeReady Workspaces automatically adds the other fields to the new component.

Alternatively, you can link to a **meta.yaml** file hosted on GitHub, using the dedicated reference field.

```
- type: chePlugin
  reference: 1
```


1 `https://raw.githubusercontent.com/<username>/<registryRepository>/v3/plugins/<publisher>/<plug-inName>/<plug-inVersion>/meta.yaml`

4. Restart the workspace for the changes to take effect.

## 3.5. TESTING A VISUAL STUDIO CODE EXTENSION IN CODEREADY WORKSPACES

Visual Studio Code (VS Code) extensions work in a workspace. VS Code extensions can run in the Che-Theia editor container, or in their own isolated and preconfigured containers with their prerequisites.

This section describes how to test a VS Code extension in CodeReady Workspaces with workspaces and how to review the compatibility of VS Code extensions to check whether a specific API is available.



### NOTE

The extension-hosting sidecar container and the use of the extension in a devfile are optional.

### 3.5.1. Testing a VS Code extension using GitHub gist

Each workspace can have its own set of plug-ins. The list of plug-ins and the list of projects to clone are defined in the **devfile.yaml** file.

For example, to enable an AsciiDoc plug-in from the Red Hat CodeReady Workspaces dashboard, add the following snippet to the devfile:

```
components:
- id: joapinto/vscode-asciidoctor/latest
  type: chePlugin
```

To add a plug-in that is not in the default plug-in registry, build a custom plug-in registry. See [Section 3.1.2, "Building a custom plug-in registry"](#), or, alternatively, use GitHub and the gist service.

#### Prerequisites

- A running instance of Red Hat CodeReady Workspaces. To install an instance of Red Hat CodeReady Workspaces, see [Installing CodeReady Workspaces on OpenShift Container Platform](#).



- A GitHub account.

## Procedure

1. Go to the [gist](#) webpage and create a **README.md** file with the following description: **Try Bracket Pair Colorizer extension in Red Hat CodeReady Workspaces** and content: **Example VS Code extension**. ([Bracket Pair Colorizer](#) is a popular VS Code extension.)
2. Click the **Create secret gist** button.
3. Clone the gist repository by using the URL from the navigation bar of the browser:

```
$ git clone https://gist.github.com/<your-github-username>/<gist-id>
```

### Example of the output of the git clone command

```
git clone https://gist.github.com/benoitf/85c60c8c439177ac50141d527729b9d9 1
Cloning into '85c60c8c439177ac50141d527729b9d9'...
remote: Enumerating objects: 3, done.
remote: Counting objects: 100% (3/3), done.
remote: Total 3 (delta 0), reused 0 (delta 0), pack-reused 0
Unpacking objects: 100% (3/3), done.
```

- 1** Each gist has a unique ID.

4. Change the directory:

```
$ cd <gist-directory-name> 1
```

- 1** Directory name matching the gist ID.

5. Download the plug-in from the [VS Code marketplace](#) or from its [GitHub page](#), and store the plug-in file in the cloned directory.
6. Create a **plugin.yaml** file in the cloned directory to add the definition of this plug-in.

### Example of the plugin.yaml file referencing the .vsix binary file extension

```
apiVersion: v2
publisher: CoenraadS
name: bracket-pair-colorizer
version: 1.0.61
type: VS Code extension
displayName: Bracket Pair Colorizer
title: Bracket Pair Colorizer
description: Bracket Pair Colorizer
icon: https://raw.githubusercontent.com/redhat-developer/codeready-workspaces/master/dependencies/che-plugin-registry/resources/images/default.svg?
sanitize=true
repository: https://github.com/CoenraadS/BracketPair
category: Language
firstPublicationDate: '2020-07-30'
```

```
spec:
  extensions:
    - "{{REPOSITORY}}/CoenraadS.bracket-pair-colorizer-1.0.61.vsix"
  latestUpdateDate: "2020-07-30"
```

1 This extension requires a basic Node.js runtime, so it is not necessary to add a custom runtime image in **plugin.yaml**.

2 **{{REPOSITORY}}** is a macro for a pre-commit hook.

7. Define a memory limit and volumes:

```
spec:
  containers:
    - image: "quay.io/eclipse/che-sidecar-java:8-0cfbacb"
      name: vscode-java
      memoryLimit: "1500Mi"
      volumes:
        - mountPath: "/home/theia/.m2"
          name: m2
```

8. Create a **devfile.yaml** that references the **plugin.yaml** file:

```
apiVersion: 1.0.0
metadata:
  generateName: java-maven-
projects:
  -
    name: console-java-simple
    source:
      type: git
      location: "https://github.com/che-samples/console-java-simple.git"
      branch: java1.11
components:
  -
    type: chePlugin
    id: redhat/java11/latest
  -
    type: chePlugin
    reference: "{{REPOSITORY}}/plugin.yaml"
  -
    type: dockerimage
    alias: maven
    image: quay.io/eclipse/che-java11-maven:nightly
    memoryLimit: 512Mi
    mountSources: true
    volumes:
      - name: m2
        containerPath: /home/user/.m2
commands:
  -
    name: maven build
    actions:
      -
```

```

type: exec
component: maven
command: "mvn clean install"
workdir: ${CHE_PROJECTS_ROOT}/console-java-simple
-
name: maven build and run
actions:
-
type: exec
component: maven
command: "mvn clean install && java -jar ./target/*.jar"
workdir: ${CHE_PROJECTS_ROOT}/console-java-simple

```

- 1 Any other devfile definition is also accepted. The important information in this devfile are the lines defining this external component. It means that an external reference defines the plug-in (instead of an ID pointing to a definition in the default plug-in registry).

9. Verify there are 4 files in the current Git directory:

```

$ ls -la
.git
CoenraadS.bracket-pair-colorizer-1.0.61.vsix
README.md
devfile.yaml
plugin.yaml

```

10. Before committing the files, add a pre-commit hook to update the **{{REPOSITORY}}** variable to the public external raw gist link:
- Create a **.git/hooks/pre-commit** file with this content:

```

#!/bin/sh

# get modified files
FILES=$(git diff --cached --name-only --diff-filter=ACMR "*.yaml" | sed 's| |\ |g')

# exit fast if no files found
[ -z "$FILES" ] && exit 0

# grab remote origin
origin=$(git config --get remote.origin.url)
url="${origin}/raw"

# iterate on files and add the good prefix pattern
for FILE in ${FILES}; do
  sed -e "s#{{REPOSITORY}}#${url}#g" "${FILE}" > "${FILE}.back"
  mv "${FILE}.back" "${FILE}"
done

# Add back to staging
echo "$FILES" | xargs git add

exit 0

```

The hook replaces the **{{REPOSITORY}}** macro and adds the external raw link to the gist.

- b. Make the script executable:

```
$ chmod u+x .git/hooks/pre-commit
```

11. Commit and push the files:

```
# Add files
$ git add *

# Commit
$ git commit -m "Initial Commit for the test of our extension"
[master 98dd370] Initial Commit for the test of our extension
3 files changed, 61 insertions(+)
create mode 100644 CoenraadS.bracket-pair-colorizer-1.0.61.vsix
create mode 100644 devfile.yaml
create mode 100644 plugin.yaml

# and push the files to the main branch
$ git push origin
```

12. Visit the gist website and verify that all links have the correct public URL and do not contain any **{{REPOSITORY}}** variables. To reach the devfile:

```
$ echo "$(git config --get remote.origin.url)/raw/devfile.yaml"
```

or:

```
$ echo "https://<che-server>/f?url=$(git config --get remote.origin.url)/raw/devfile.yaml"
```

### 3.5.2. Verifying the VS Code extension API compatibility level

Che-Theia does not fully support the VS Code extensions API. The [vscode-theia-comparator](#) is used to analyze the compatibility between the Che-Theia plug-in API and the VS Code extension API. This tool runs nightly, and the results are published on the [vscode-theia-comparator](#) GitHub page.

#### Prerequisites

- Personal GitHub access token. See [Creating a personal access token for the command line](#) . A GitHub access token is required to increase the GitHub download limit for your IP address.

#### Procedure

To run the **vscode-theia comparator** manually:

1. Clone the [vscode-theia-comparator](#) repository, and build it using the **yarn** command.
2. Set the **GITHUB\_TOKEN** environment variable to your token.
3. Execute the **yarn run generate** command to generate a report.
4. Open the **out/status.html** file to view the report.

## CHAPTER 4. RETRIEVING CODEREADY WORKSPACES LOGS

For information about obtaining various types of logs in CodeReady Workspaces, see the following sections:

- [Section 4.1, “Accessing OpenShift events on OpenShift”](#)
- [Section 4.3, “Viewing CodeReady Workspaces server logs”](#)
- [Section 4.4, “Viewing external service logs”](#)
- [Section 4.5, “Viewing CodeReady Workspaces workspaces logs”](#)
- [Section 4.6, “Viewing the plug-in broker logs”](#)
- [Section 4.7, “Collecting logs using crwctl”](#)

### 4.1. ACCESSING OPENSIFT EVENTS ON OPENSIFT

For high-level monitoring of OpenShift projects, view the OpenShift events that the project performs.

This section describes how to access these events in the OpenShift web console.

#### Prerequisites

- A running OpenShift web console.

#### Procedure

1. In the left panel of the OpenShift web console, click the **Home → Events**.
2. To view the list of all events for a particular project, select the project from the list.
3. The details of the events for the current project are displayed.

#### Additional resources

- For a list of OpenShift events, see [Comprehensive List of Events in OpenShift documentation](#).

### 4.2. VIEWING THE STATE OF THE CODEREADY WORKSPACES CLUSTER DEPLOYMENT USING OPENSIFT 4 CLI TOOLS

This section describes how to view the state of the CodeReady Workspaces cluster deployment using OpenShift 4 CLI tools.

#### Prerequisites

- An instance of Red Hat CodeReady Workspaces running on OpenShift.
- An installation of the OpenShift command-line tool, **oc**.

#### Procedure

1. Run the following commands to select the **crw** project:

```
$ oc project <project_name>
```

- Run the following commands to get the name and status of the Pods running in the selected project:

```
$ oc get pods
```

- Check that the status of all the Pods is **Running**.

#### Example 4.1. Pods with status Running

NAME	READY	STATUS	RESTARTS	AGE
codeready-8495f4946b-jrzdc	0/1	Running	0	86s
codeready-operator-578765d954-99szc	1/1	Running	0	42m
keycloak-74fbfb9654-g9vp5	1/1	Running	0	4m32s
postgres-5d579c6847-w6wx5	1/1	Running	0	5m14s

- To see the state of the CodeReady Workspaces cluster deployment, run:

```
$ oc logs --tail=10 -f `(oc get pods -o name | grep operator)`
```

#### Example 4.2. Logs of the Operator:

```
time="2019-07-12T09:48:29Z" level=info msg="Exec successfully completed"
time="2019-07-12T09:48:29Z" level=info msg="Updating eclipse-che CR with status:
provisioned with OpenShift identity provider: true"
time="2019-07-12T09:48:29Z" level=info msg="Custom resource eclipse-che updated"
time="2019-07-12T09:48:29Z" level=info msg="Creating a new object: ConfigMap, name:
che"
time="2019-07-12T09:48:29Z" level=info msg="Creating a new object: ConfigMap, name:
custom"
time="2019-07-12T09:48:29Z" level=info msg="Creating a new object: Deployment,
name: che"
time="2019-07-12T09:48:30Z" level=info msg="Updating eclipse-che CR with status:
CodeReady Workspaces API: Unavailable"
time="2019-07-12T09:48:30Z" level=info msg="Custom resource eclipse-che updated"
time="2019-07-12T09:48:30Z" level=info msg="Waiting for deployment che. Default
timeout: 420 seconds"
```

## 4.3. VIEWING CODEREADY WORKSPACES SERVER LOGS

This section describes how to view the CodeReady Workspaces server logs using the command line.

### 4.3.1. Viewing the CodeReady Workspaces server logs using the OpenShift CLI

This section describes how to view the CodeReady Workspaces server logs using the OpenShift CLI (command line interface).

#### Procedure

1. In the terminal, run the following command to get the Pods:

```
$ oc get pods
```

#### Example

```
$ oc get pods
NAME          READY STATUS RESTARTS AGE
codeready-11-j4w2b 1/1  Running 0      3m
```

2. To get the logs for a deployment, run the following command:

```
$ oc logs <name-of-pod>
```

#### Example

```
$ oc logs codeready-11-j4w2b
```

## 4.4. VIEWING EXTERNAL SERVICE LOGS

This section describes how to view the logs from external services related to CodeReady Workspaces server.

### 4.4.1. Viewing RH-SSO logs

The RH-SSO OpenID provider consists of two parts: Server and IDE. It writes its diagnostics or error information to several logs.

#### 4.4.1.1. Viewing the RH-SSO server logs

This section describes how to view the RH-SSO OpenID provider server logs.

##### Procedure

1. In the OpenShift Web Console, click **Deployments**.
2. In the **Filter by label** search field, type **keycloak** to see the RH-SSO logs.
3. In the **Deployment Configs** section, click the **keycloak** link to open it.
4. In the **History** tab, click the **View log** link for the active RH-SSO deployment.
5. The RH-SSO logs are displayed.

##### Additional resources

- See the [Section 4.3, “Viewing CodeReady Workspaces server logs”](#) for diagnostics and error messages related to the RH-SSO IDE Server.

#### 4.4.1.2. Viewing the RH-SSO client logs on Firefox

This section describes how to view the RH-SSO IDE client diagnostics or error information in the Firefox **WebConsole**.

#### Procedure

- Click **Menu** > **WebDeveloper** > **WebConsole**.

#### 4.4.1.3. Viewing the RH-SSO client logs on Google Chrome

This section describes how to view the RH-SSO IDE client diagnostics or error information in the Google Chrome **Console** tab.

#### Procedure

1. Click **Menu** > **More Tools** > **Developer Tools**.
2. Click the **Console** tab.

#### 4.4.2. Viewing the CodeReady Workspaces database logs

This section describes how to view the database logs in CodeReady Workspaces, such as PostgreSQL server logs.

#### Procedure

1. In the OpenShift Web Console, click **Deployments**.
2. In the **Find by label** search field, type:
  - **app=che** and press **Enter**
  - **component=postgres** and **Enter**The OpenShift Web Console now searches base on those two keys and displays PostgreSQL logs.
3. Click **postgres** deployment to open it.
4. Click the **View log** link for the active PostgreSQL deployment.  
The OpenShift Web Console displays the database logs.

#### Additional resources

- Some diagnostics or error messages related to the PostgreSQL server can be found in the active CodeReady Workspaces deployment log. For details to access the active CodeReady Workspaces deployments logs, see the [Section 4.3, “Viewing CodeReady Workspaces server logs”](#) section.

## 4.5. VIEWING CODEREADY WORKSPACES WORKSPACES LOGS

This section describes how to view CodeReady Workspaces workspaces logs.

#### 4.5.1. Viewing Che-Theia IDE logs

This section describes how to view Che-Theia IDE logs.



### 4.5.1.1. Viewing Che-Theia editor logs using the OpenShift CLI

Observing Che-Theia editor logs helps to get a better understanding and insight over the plug-ins loaded by the editor. This section describes how to access the Che-Theia editor logs using the OpenShift CLI (command-line interface).

#### Prerequisites

- CodeReady Workspaces is deployed in an OpenShift cluster.
- A workspace is created.
- User is located in a CodeReady Workspaces installation project.

#### Procedure

1. Obtain the list of the available Pods:

```
$ oc get pods
```

#### Example

```
$ oc get pods
NAME                                READY STATUS RESTARTS AGE
codeready-9-xz6g8                   1/1   Running 1    15h
workspace0zqb2ew3py4srthh.go-cli-549cdcf69-9n4w2 4/4   Running 0    1h
```

2. Obtain the list of the available containers in the particular Pod:

```
$ oc get pods <name-of-pod> --output jsonpath='{.spec.containers[*].name}'
```

#### Example:

```
$ oc get pods workspace0zqb2ew3py4srthh.go-cli-549cdcf69-9n4w2 -o
jsonpath='{.spec.containers[*].name}'
> go-cli che-machine-exechr7 theia-idexzb vscode-gox3r
```

3. Get logs from the **theia/ide** container:

```
$ oc logs --follow <name-of-pod> --container <name-of-container>
```

#### Example:

```
$ oc logs --follow workspace0zqb2ew3py4srthh.go-cli-549cdcf69-9n4w2 -container
theia-idexzb
>root INFO unzipping the plug-in 'task_plugin.theia' to directory: /tmp/theia-
unpacked/task_plugin.theia
root INFO unzipping the plug-in 'theia_yeoman_plugin.theia' to directory: /tmp/theia-
unpacked/theia_yeoman_plugin.theia
root WARN A handler with prefix term is already registered.
root INFO [nsfw-watcher: 75] Started watching: /home/theia/.theia
root WARN e.onStart is slow, took: 367.460000013015 ms
```

```
root INFO [nsfw-watcher: 75] Started watching: /projects
root INFO [nsfw-watcher: 75] Started watching: /projects/.theia/tasks.json
root INFO [4f9590c5-e1c5-40d1-b9f8-ec31ec3bdac5] Sync of 9 plugins took:
62.26000000242493 ms
root INFO [nsfw-watcher: 75] Started watching: /projects
root INFO [hosted-plugin: 88] PLUGIN_HOST(88) starting instance
```

## 4.5.2. Viewing logs from language servers and debug adapters

### 4.5.2.1. Checking important logs

This section describes how to check important logs.

#### Procedure

1. In the OpenShift web console, click **Applications** → **Pods** to see a list of all the active workspaces.
2. Click on the name of the running Pod where the workspace is running. The Pod screen contains the list of all containers with additional information.
3. Choose a container and click the container name.

#### TIP

The most important logs are the **theia-ide** container and the plug-ins container logs.

4. On the container screen, navigate to the **Logs** section.

### 4.5.2.2. Detecting memory problems

This section describes how to detect memory problems related to a plug-in running out of memory. The following are the two most common problems related to a plug-in running out of memory:

#### The plug-in container runs out of memory

This can happen during plug-in initialization when the container does not have enough RAM to execute the entrypoint of the image. The user can detect this in the logs of the plug-in container. In this case, the logs contain **OOMKilled**, which implies that the processes in the container requested more memory than is available in the container.

#### A process inside the container runs out of memory without the container noticing this

For example, the Java language server (Eclipse JDT Language Server, started by the `vscode-java` extension) throws an **OutOfMemoryException**. This can happen any time after the container is initialized, for example, when a plug-in starts a language server or when a process runs out of memory because of the size of the project it has to handle.

To detect this problem, check the logs of the primary process running in the container. For example, to check the log file of Eclipse JDT Language Server for details, see the relevant plug-in-specific sections.

### 4.5.2.3. Logging the client-server traffic for debug adapters

This section describes how to log the exchange between Che-Theia and a debug adapter into the **Output** view.

## Prerequisites

- A debug session must be started for the **Debug adapters** option to appear in the list.

## Procedure

1. Click **File** → **Settings** and then **open Preferences**.
2. Expand the **Debug** section in the **Preferences** view.
3. Set the **trace** preference value to **true** (default is **false**).
4. All the communication events are now logged.
5. To watch these events, click **View** → **Output** and select **Debug adapters** from the drop-down list at the upper right corner of the **Output** view.

### 4.5.2.4. Viewing logs for Python

This section describes how to view logs for the Python language server.

## Procedure

- Navigate to the **Output** view and select **Python** in the drop-down list.



```

Output x Python
Starting Microsoft Python language server.
Downloading https://pvsc.azureedge.net/python-language-server-stable/Python-Language-Server-linux-x64.0.2.96.nupkg... #####Linting 0
***** Module test
12,0,error,import-error:Unable to import 'demonstrate'
-----
Your code has been rated at -2.50/10
complete
Unpacking archive... done

```

### 4.5.2.5. Viewing logs for Go

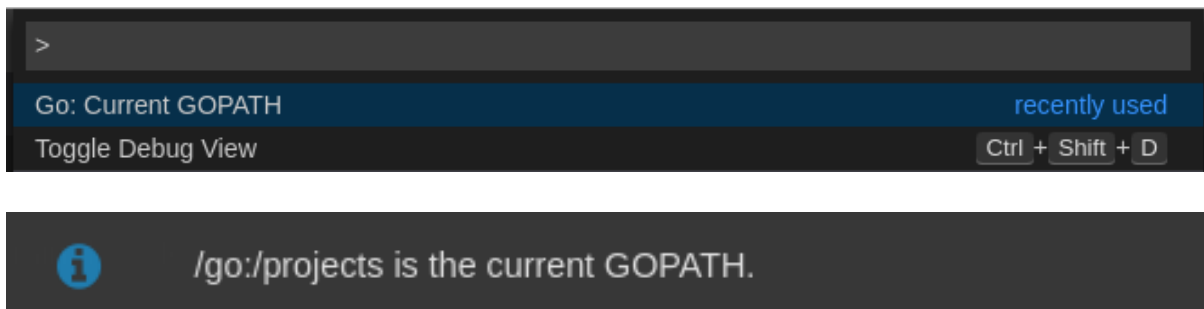
This section describes how to view logs for the Go language server.

#### 4.5.2.5.1. Finding the GOPATH

This section describes how to find where the **GOPATH** variable points to.

## Procedure

- Execute the **Go: Current GOPATH** command.



```

>
Go: Current GOPATH recently used
Toggle Debug View Ctrl + Shift + D

/go:/projects is the current GOPATH.

```

#### 4.5.2.5.2. Viewing the Debug Console log for Go

This section describes how to view the log output from the Go debugger.

## Procedure

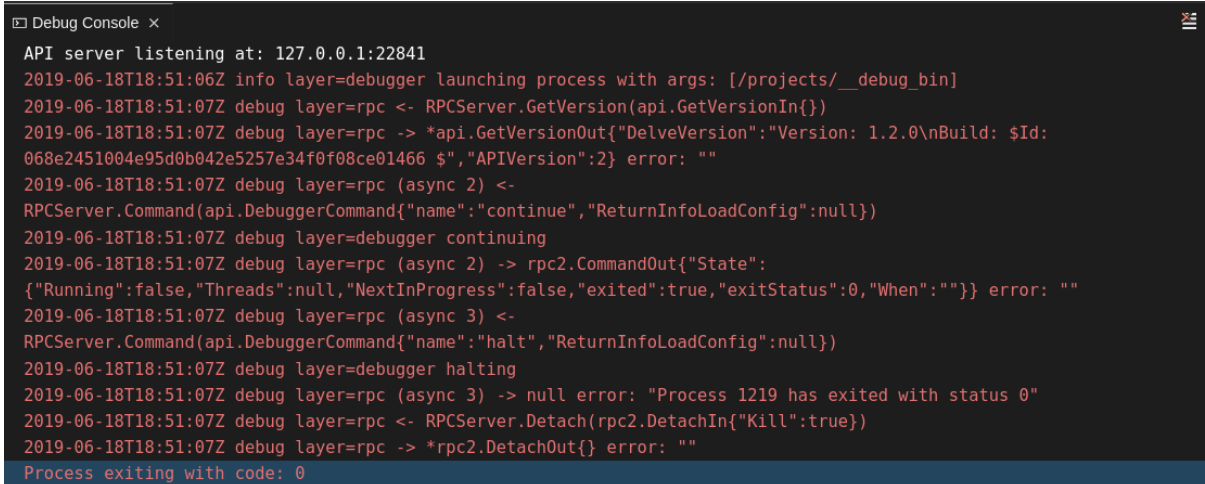
1. Set the **showLog** attribute to **true** in the debug configuration.

```
{
  "version": "0.2.0",
  "configurations": [
    {
      "type": "go",
      "showLog": true
      ....
    }
  ]
}
```

2. To enable debugging output for a component, add the package to the comma-separated list value of the **logOutput** attribute:

```
{
  "version": "0.2.0",
  "configurations": [
    {
      "type": "go",
      "showLog": true,
      "logOutput": "debugger,rpc,gdbwire,lldbout,debuglineerr"
      ....
    }
  ]
}
```

3. The debug console prints the additional information in the debug console.



```
Debug Console x
API server listening at: 127.0.0.1:22841
2019-06-18T18:51:06Z info layer=debugger launching process with args: [/projects/__debug_bin]
2019-06-18T18:51:07Z debug layer=rpc <- RPCServer.GetVersion(api.GetVersionIn{})
2019-06-18T18:51:07Z debug layer=rpc -> *api.GetVersionOut{"DelveVersion":"Version: 1.2.0\nBuild: $Id: 068e2451004e95d0b042e5257e34f0f08ce01466 $","APIVersion":2} error: ""
2019-06-18T18:51:07Z debug layer=rpc (async 2) <-
RPCServer.Command(api.DebuggerCommand{"name":"continue","ReturnInfoLoadConfig":null})
2019-06-18T18:51:07Z debug layer=debugger continuing
2019-06-18T18:51:07Z debug layer=rpc (async 2) -> rpc2.CommandOut{"State":
{"Running":false,"Threads":null,"NextInProgress":false,"exited":true,"exitStatus":0,"When":""}} error: ""
2019-06-18T18:51:07Z debug layer=rpc (async 3) <-
RPCServer.Command(api.DebuggerCommand{"name":"halt","ReturnInfoLoadConfig":null})
2019-06-18T18:51:07Z debug layer=debugger halting
2019-06-18T18:51:07Z debug layer=rpc (async 3) -> null error: "Process 1219 has exited with status 0"
2019-06-18T18:51:07Z debug layer=rpc <- RPCServer.Detach(rpc2.DetachIn{"Kill":true})
2019-06-18T18:51:07Z debug layer=rpc -> *rpc2.DetachOut{} error: ""
Process exiting with code: 0
```

### 4.5.2.5.3. Viewing the Go logs output in the Output panel

This section describes how to view the Go logs output in the **Output** panel.

## Procedure

1. Navigate to the **Output** view.
2. Select **Go** in the drop-down list.

```

Output x
Starting linting the current package at /projects
Starting "go vet" under the folder /projects
Starting building the current package at /projects
Not able to determine import path of current package by using cwd: /projects and Go workspace:
/projects>Finished running tool: /go/bin/golint
/projects>Finished running tool: /usr/local/go/bin/go vet ./...
/projects>Finished running tool: /usr/local/go/bin/go build -i -o /tmp/vscode-goGJoFLE/go-code-check .

```

#### 4.5.2.6. Viewing logs for the NodeDebug NodeDebug2 adapter



#### NOTE

No specific diagnostics exist other than the general ones.

#### 4.5.2.7. Viewing logs for Typescript

##### 4.5.2.7.1. Enabling the label switched protocol (LSP) tracing

#### Procedure

1. To enable the tracing of messages sent to the Typescript (TS) server, in the **Preferences** view, set the **typescript.tsserver.trace** attribute to **verbose**. Use this to diagnose the TS server issues.
2. To enable logging of the TS server to a file, set the **typescript.tsserver.log** attribute to **verbose**. Use this log to diagnose the TS server issues. The log contains the file paths.

##### 4.5.2.7.2. Viewing the Typescript language server log

This section describes how to view the Typescript language server log.

#### Procedure

1. To get the path to the log file, see the Typescript **Output** console:

```

Problems Output x TypeScript
Info - 11:14:26 AM] Using tsserver from: /tmp/vscode-unpacked/che-incubator.typescript.latest.dvuuojoyht.che-typescript-language-1.35.1.vsix/extension/node_modules/typescript/Lib
Info - 11:14:26 AM] TSServer_log file: /home/theia/.theia/logs/20190621T111312/host/vscode.typescript-language-features/tsserver-log-cdBAj1/tsserver.log
Info - 11:14:26 AM] Forking TSServer
Info - 11:14:26 AM] Started TSServer

```

2. To open log file, use the **Open TS Server log** command.

```

>Typescript
TypeScript: Open TS Server log
TypeScript: Restart TS server
TypeScript: Select TypeScript Version...

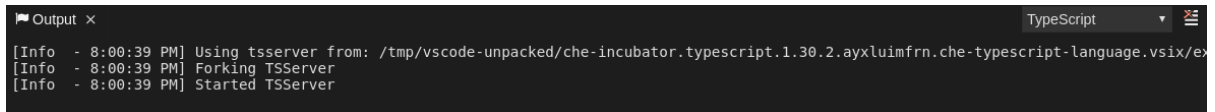
```

##### 4.5.2.7.3. Viewing the Typescript logs output in the Output panel

This section describes how to view the Typescript logs output in the **Output** panel.

#### Procedure

1. Navigate to the **Output** view
2. Select **TypeScript** in the drop-down list.



```

Output x TypeScript
[Info - 8:00:39 PM] Using tsserver from: /tmp/vscode-unpacked/che-incubator.typescript.1.30.2.ayxluimfrn.che-typescript-language.vsix/ex
[Info - 8:00:39 PM] Forking TSServer
[Info - 8:00:39 PM] Started TSServer

```

#### 4.5.2.8. Viewing logs for Java

Other than the general diagnostics, there are [Language Support for Java \(Eclipse JDT Language Server\)](#) plug-in actions that the user can perform.

##### 4.5.2.8.1. Verifying the state of the Eclipse JDT Language Server

###### Procedure

Check if the container that is running the Eclipse JDT Language Server plug-in is running the Eclipse JDT Language Server main process.

1. Open a terminal in the container that is running the Eclipse JDT Language Server plug-in (an example name for the container: **vscode-javaxxx**).
2. Inside the terminal, run the **ps aux | grep jdt** command to check if the Eclipse JDT Language Server process is running in the container. If the process is running, the output is:

```
usr/lib/jvm/default-jvm/bin/java --add-modules=ALL-SYSTEM --add-opens java.base/java.util
```

This message also shows the VSCode Java extension used. If it is not running, the language server has not been started inside the container.

3. Check all logs described in [Checking important logs](#)

##### 4.5.2.8.2. Verifying the Eclipse JDT Language Server features

###### Procedure

If the Eclipse JDT Language Server process is running, check if the language server features are working:

1. Open a Java file and use the hover or autocomplete functionality. In case of an erroneous file, the user sees Java in the **Outline** view or in the **Problems** view.

##### 4.5.2.8.3. Viewing the Java language server log

###### Procedure

The Eclipse JDT Language Server has its own workspace where it logs errors, information about executed commands, and events.

1. To open this log file, open a terminal in the container that is running the Eclipse JDT Language Server plug-in. You can also view the log file by running the **Java: Open Java Language Server log file** command.
2. Run **cat <PATH\_TO\_LOG\_FILE>** where **PATH\_TO\_LOG\_FILE** is **/home/theia/.theia/workspace-storage/<workspace\_name>/redhat.java/jdt\_ws/.metadata/.log**.

##### 4.5.2.8.4. Logging the Java language server protocol (LSP) messages

## Procedure

To log the LSP messages to the VS Code **Output** view, enable tracing by setting the **java.trace.server** attribute to **verbose**.

## Additional resources

For troubleshooting instructions, see the [VS Code Java Github repository](#).

### 4.5.2.9. Viewing logs for Intelephense

#### 4.5.2.9.1. Logging the Intelephense client-server communication

## Procedure

To configure the PHP Intelephense language support to log the client-server interexchange in the **Output** view:

1. Click **File → Settings**.
2. Open the **Preferences** view.
3. Expand the **Intelephense** section and set the **trace.server.verbose** preference value to **verbose** to see all the communication events (the default value is **off**).

#### 4.5.2.9.2. Viewing Intelephense events in the Output panel

This procedure describes how to view Intelephense events in the **Output** panel.

## Procedure

1. Click **View → Output**
2. Select **Intelephense** in the drop-down list for the **Output** view.

### 4.5.2.10. Viewing logs for PHP-Debug

This procedure describes how to configure the PHP Debug plug-in to log the PHP Debug plug-in diagnostic messages into the **Debug Console** view. Configure this before the start of the debug session.

## Procedure

1. In the **launch.json** file, add the **"log": true** attribute to the selected launch configuration.
2. Start the debug session.
3. The diagnostic messages are printed into the **Debug Console** view along with the application output.

### 4.5.2.11. Viewing logs for XML

Other than the general diagnostics, there are XML plug-in specific actions that the user can perform.

#### 4.5.2.11.1. Verifying the state of the XML language server

## Procedure

1. Open a terminal in the container named **vscode-xml-<xxx>**.
2. Run **ps aux | grep java** to verify that the XML language server has started. If the process is running, the output is:

```
java ***/org.eclipse.ls4xml-uber.jar`
```

If is not, see the Checking important logs chapter.

### 4.5.2.11.2. Checking XML language server feature flags

## Procedure

1. Check if the features are enabled. The XML plug-in provides multiple settings that can enable and disable features:
  - **xml.format.enabled**: Enable the formatter
  - **xml.validation.enabled**: Enable the validation
  - **xml.documentSymbols.enabled**: Enable the document symbols
2. To diagnose whether the XML language server is working, create a simple XML element, such as **<hello></hello>**, and confirm that it appears in the **Outline** panel on the right.
3. If the document symbols do not show, ensure that the **xml.documentSymbols.enabled** attribute is set to **true**. If it is **true**, and there are no symbols, the language server may not be hooked to the editor. If there are document symbols, then the language server is connected to the editor.
4. Ensure that the features that the user needs, are set to **true** in the settings (they are set to **true** by default). If any of the features are not working, or not working as expected, file an issue against the [Language Server](#).

### 4.5.2.11.3. Enabling XML Language Server Protocol (LSP) tracing

## Procedure

To log LSP messages to the VS Code **Output** view, enable tracing by setting the **xml.trace.server** attribute to **verbose**.

### 4.5.2.11.4. Viewing the XML language server log

## Procedure

The log from the language server can be found in the plug-in sidecar at **/home/theia/.theia/workspace-storage/<workspace\_name>/redhat.vscode-xml/lsp4xml.log**.

### 4.5.2.12. Viewing logs for YAML

This section describes the YAML plug-in specific actions that the user can perform, in addition to the general diagnostics ones.



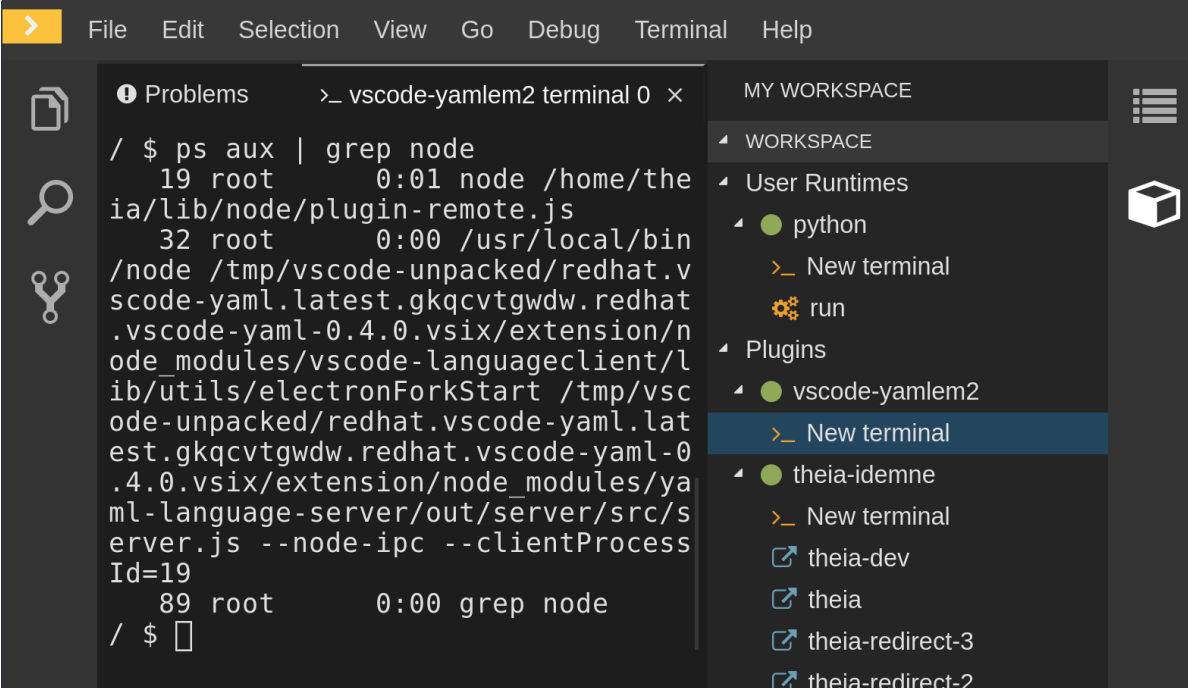
#### 4.5.2.12.1. Verifying the state of the YAML language server

This section describes how to verify the state of the YAML language server.

##### Procedure

Check if the container running the YAML plug-in is running the YAML language server.

1. In the editor, open a terminal in the container that is running the YAML plug-in (an example name of the container: **vscode-yaml-*<xxx>***).
2. In the terminal, run the **ps aux | grep node** command. This command searches all the node processes running in the current container.
3. Verify that a command **node \*\*/server.js** is running.



```

> File Edit Selection View Go Debug Terminal Help
Problems >_ vscode-yamlem2 terminal 0 x
/ $ ps aux | grep node
 19 root      0:01 node /home/the
ia/lib/node/plugin-remote.js
 32 root      0:00 /usr/local/bin
/node /tmp/vscode-unpacked/redhat.v
scode-yaml.latest.gkqcvtdgdw.redhat
.vscode-yaml-0.4.0.vsix/extension/n
ode_modules/vscode-languageclient/l
ib/utis/electronForkStart /tmp/vsc
ode-unpacked/redhat.vscode-yaml.la
test.gkqcvtdgdw.redhat.vscode-yaml-0
.4.0.vsix/extension/node_modules/ya
ml-language-server/out/server/src/s
erver.js --node-ipc --clientProcess
Id=19
 89 root      0:00 grep node
/ $ █
  
```

The **node \*\*/server.js** running in the container indicates that the language server is running. If it is not running, the language server has not started inside the container. In this case, see [Checking important logs](#).

#### 4.5.2.12.2. Checking the YAML language server feature flags

##### Procedure

To check the feature flags:

1. Check if the features are enabled. The YAML plug-in provides multiple settings that can enable and disable features, such as:
  - **yaml.format.enable**: Enables the formatter
  - **yaml.validate**: Enables validation
  - **yaml.hover**: Enables the hover function
  - **yaml.completion**: Enables the completion function

2. To check if the plug-in is working, type the simplest YAML, such as `hello: world`, and then open the Outline panel on the right side of the editor.
3. Verify if there are any document symbols. If yes, the language server is connected to the editor.
4. If any feature is not working, make sure that the settings listed above are set to `true` (they are set to `true` by default). If a feature is not working, file an issue against the [Language Server](#).

#### 4.5.2.12.3. Enabling YAML Language Server Protocol (LSP) tracing

##### Procedure

To log LSP messages to the VS Code Output view, enable tracing by setting the `yaml.trace.server` attribute to `verbose`.

#### 4.5.2.13. Viewing logs for Dotnet with Omnisharp-Theia plug-in

##### 4.5.2.13.1. Omnisharp-Theia plug-in

CodeReady Workspaces uses the Omnisharp-Theia plug-in as a remote plug-in. It is located at [github.com/redhat-developer/omnisharp-theia-plugin](https://github.com/redhat-developer/omnisharp-theia-plugin). In case of an issue, report it, or contribute your fix in the repository.

This plug-in registers `omnisharp-roslyn` as a language server and provides project dependencies and language syntax for C# applications.

The language server runs on .NET SDK 2.2.105.

##### 4.5.2.13.2. Verifying the state of the Omnisharp-Theia plug-in language server

##### Procedure

To check if the container running the Omnisharp-Theia plug-in is running OmniSharp, execute the `ps aux | grep OmniSharp.exe` command. If the process is running, the following is an example output:

```
/tmp/theia-unpacked/redhat-developer.che-omnisharp-  
plugin.0.0.1.zcpaqpczwb.omnisharp_theia_plugin.theia/server/bin/mono  
/tmp/theia-unpacked/redhat-developer.che-omnisharp-  
plugin.0.0.1.zcpaqpczwb.omnisharp_theia_plugin.theia/server/omnisharp/OmniSharp.exe
```

If the output is different, the language server has not started inside the container. Check the logs described in [Checking important logs](#).

##### 4.5.2.13.3. Checking Omnisharp Che-Theia plug-in language server features

##### Procedure

- If the `OmniSharp.exe` process is running, check if the language server features are working by opening a `.cs` file and trying the hover or completion features, or opening the `Problems` or `Outline` view.

#### 4.5.2.13.4. Viewing Omnisharp-Theia plug-in logs in the Output panel

##### Procedure

If **Omnisharp.exe** is running, it logs all information in the Output panel. To view the logs, open the Output view and select **C#** from the drop-down list.

#### 4.5.2.14. Viewing logs for Dotnet with NetcoredebugOutput plug-in

##### 4.5.2.14.1. NetcoredebugOutput plug-in

The NetcoredebugOutput plug-in provides the **netcoredbg** tool. This tool implements the VS Code Debug Adapter protocol and allows users to debug .NET applications under the .NET Core runtime.

The container where the NetcoredebugOutput plug-in is running contains Dotnet SDK v.2.2.105.

##### 4.5.2.14.2. Verifying the state of the NetcoredebugOutput plug-in

##### Procedure

To test the plug-in initialization:

1. Check if there is a netcoredbg debug configuration in the **launch.json** file. The following is an example debug configuration:

```
{
  "type": "netcoredbg",
  "request": "launch",
  "program": "${workspaceFolder}/bin/Debug/<target-framework>/<project-name.dll>",
  "args": [],
  "name": ".NET Core Launch (console)",
  "stopAtEntry": false,
  "console": "internalConsole"
}
```

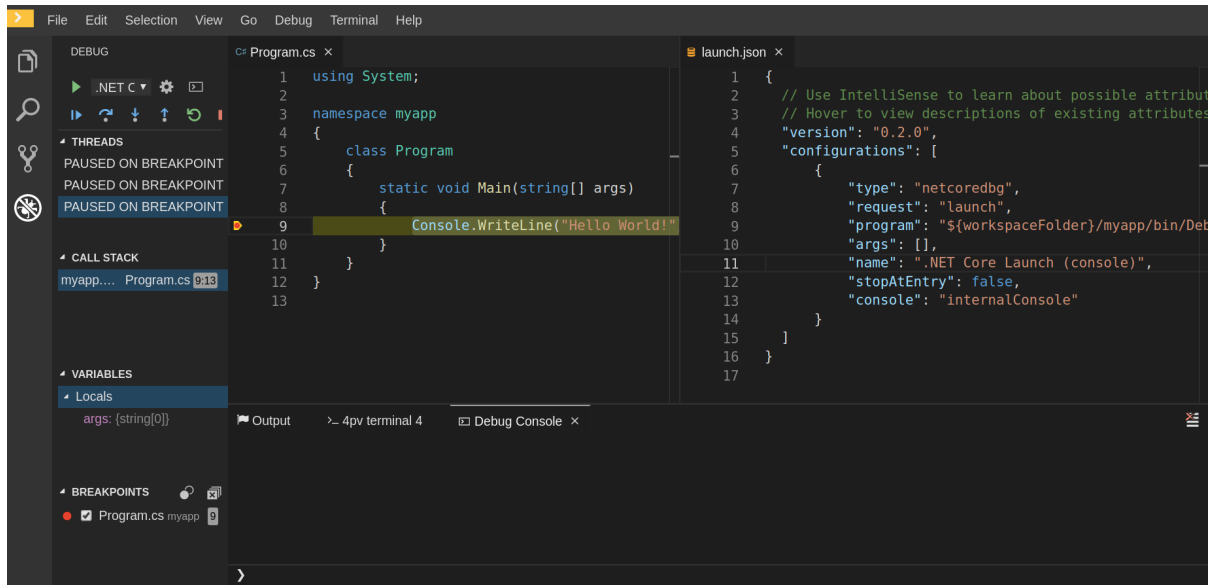
2. To test if it exists, test the autocompletion feature within the braces of the **configuration** section of the **launch.json** file. If you can find **netcoredbg**, the Che-Theia plug-in is correctly initialized. If not, see [Checking important logs](#).

##### 4.5.2.14.3. Viewing NetcoredebugOutput plug-in logs in the Output panel

This section describes how to view NetcoredebugOutput plug-in logs in the Output panel.

##### Procedure

- Open the Debug console.



### 4.5.2.15. Viewing logs for Camel

#### 4.5.2.15.1. Verifying the state of the Camel language server

##### Procedure

The user can inspect the log output of the sidecar container using the Camel language tools that are stored in the `vscode-apache-camel<xxx>` Camel container.

To verify the state of the language server:

1. Open a terminal inside the `vscode-apache-camel<xxx>` container.
2. Run the `ps aux | grep java` command. The following is an example language server process:

```
java -jar /tmp/vscode-unpacked/camel-tooling.vscode-apache-camel.latest.euqhbmeplx.d.camel-tooling.vscode-apache-camel-0.0.14.vsix/extension/jars/language-server.jar
```

3. If you cannot find it, see [Checking important logs](#).

#### 4.5.2.15.2. Viewing Camel logs in the Output panel

The Camel language server is a SpringBoot application that writes its log to the `$(java.io.tmpdir)/log-camel-lsp.out` file. Typically, `$(java.io.tmpdir)` points to the `/tmp` directory, so the filename is `/tmp/log-camel-lsp.out`.

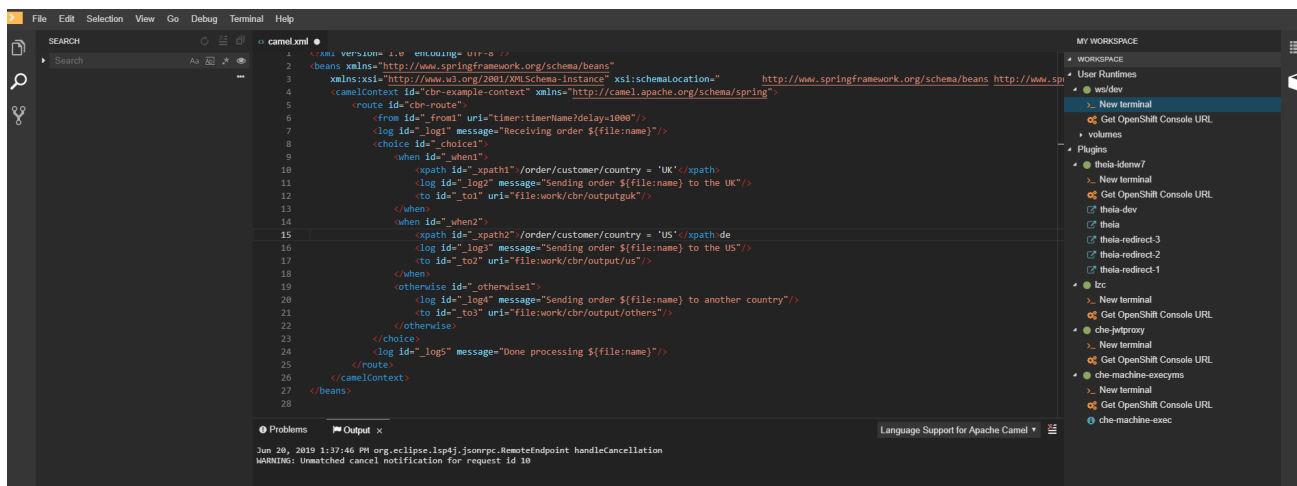
##### Procedure

The Camel language server logs are printed in the Output channel named `Language Support for Apache Camel`.



##### NOTE

The output channel is created only at the first created log entry on the client side. It may be absent when everything is going well.



## 4.6. VIEWING THE LOG-IN BROKER LOGS

This section describes how to view the plug-in broker logs.

The `che-plugin-broker` Pod itself is deleted when its work is complete. Therefore, its event logs are only available while the workspace is starting.

### Procedure

To see logged events from temporary Pods:

1. Start a CodeReady Workspaces workspace.
2. From the main OpenShift Container Platform screen, go to Workload → Pods.
3. Use the OpenShift terminal console located in the Pod's Terminal tab

### Verification step

- OpenShift terminal console displays the plug-in broker logs while the workspace is starting

## 4.7. COLLECTING LOGS USING CRWCTL

It is possible to get all Red Hat CodeReady Workspaces logs from a OpenShift cluster using the `crwctl` tool.

- `crwctl server:start` automatically starts collecting Red Hat CodeReady Workspaces servers logs during installation of Red Hat CodeReady Workspaces
- `crwctl server:logs` collects existing Red Hat CodeReady Workspaces server logs
- `crwctl workspace:logs` collects workspace logs

## CHAPTER 5. MONITORING CODEREADY WORKSPACES

This chapter describes how to configure CodeReady Workspaces to expose metrics and how to build an example monitoring stack with external tools to process data exposed as metrics by CodeReady Workspaces.

### 5.1. ENABLING AND EXPOSING CODEREADY WORKSPACES METRICS

This section describes how to enable and expose CodeReady Workspaces metrics.

#### Procedure

1. Set the **CHE\_METRICS\_ENABLED=true** environment variable, which will expose the **8087** port as a service on the che-master host.

When Red Hat CodeReady Workspaces is installed from the OperatorHub, the environment variable is set automatically if the default **CheCluster** CR is used:

[Eclipse Che](#) > Create Che Cluster

## Create Che Cluster

Create by manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.

```

1  apiVersion: org.eclipse.che/v1
2  kind: CheCluster
3  metadata:
4    name: eclipse-che
5    namespace: che-metrics
6  spec:
7    server:
8      cheImageTag: nightly
9      devfileRegistryImage: 'quay.io/eclipse/che-devfile-registry:nightly'
10     pluginRegistryImage: 'quay.io/eclipse/che-plugin-registry:nightly'
11     tlsSupport: true
12     selfSignedCert: false
13   database:
14     externalDb: false
15     chePostgresHostName: ''
16     chePostgresPort: ''
17     chePostgresUser: ''
18     chePostgresPassword: ''
19     chePostgresDb: ''
20   auth:
21     openShiftoAuth: true
22     identityProviderImage: 'quay.io/eclipse/che-keycloak:nightly'
23     externalIdentityProvider: false
24     identityProviderURL: ''
25     identityProviderRealm: ''
26     identityProviderClientId: ''
27   storage:
28     pvcStrategy: per-workspace
29     pvcClaimSize: 1Gi
30     preCreateSubPaths: true
31   metrics:
32     enable: true
33

```

```

spec:
  metrics:
    enable: true

```

## 5.2. COLLECTING CODEREADY WORKSPACES METRICS WITH PROMETHEUS

This section describes how to use the Prometheus monitoring system to collect, store and query metrics about CodeReady Workspaces.

## Prerequisites

- CodeReady Workspaces is exposing metrics on port **8087**. See [Enabling and exposing che metrics](#).
- Prometheus 2.9.1 or higher is running. The Prometheus console is running on port **9090** with a corresponding service and route. See [First steps with Prometheus](#).

## Procedure

- Configure Prometheus to scrape metrics from the **8087** port:

### Prometheus configuration example

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: prometheus-config
data:
  prometheus.yml: |-
    global:
      scrape_interval: 5s 1
      evaluation_interval: 5s 2
    scrape_configs: 3
      - job_name: 'che'
        static_configs:
          - targets: ['[che-host]:8087'] 4
```

- 1 Rate, at which a target is scraped.
- 2 Rate, at which recording and alerting rules are re-checked (not used in the system at the moment).
- 3 Resources Prometheus monitors. In the default configuration, there is a single job called **che**, which scrapes the time series data exposed by the CodeReady Workspaces server.
- 4 Scrape metrics from the **8087** port.

## Verification steps

- Use the Prometheus console to query and view metrics.  
Metrics are available at: <http://<che-server-url>:9090/metrics>.

For more information, see [Using the expression browser](#) in the Prometheus documentation.

## Additional resources

- [First steps with Prometheus](#).
- [Configuring Prometheus](#).
- [Querying Prometheus](#).



- [Prometheus metric types](#).

### 5.3. EXTENDING CODEREADY WORKSPACES MONITORING METRICS

This section describes how to create a metric or a group of metrics to extend the monitoring metrics that CodeReady Workspaces is exposing.

CodeReady Workspaces has two major modules metrics:

- **che-core-metrics-core** – contains core metrics module
- **che-core-api-metrics** – contains metrics that are dependent on core CodeReady Workspaces components, such as workspace or user managers

Procedure

- Create a class that extends the **MeterBinder** class. This allows to register the created metric in the overridden **bindTo(MeterRegistry registry)** method. The following is an example of a metric that has a function that supplies the value for it:

Example metric

```
public class UserMeterBinder implements MeterBinder {

    private final UserManager userManager;

    @Inject
    public UserMeterBinder(UserManager userManager) {
        this.userManager = userManager;
    }

    @Override
    public void bindTo(MeterRegistry registry) {
        Gauge.builder("che.user.total", this::count)
            .description("Total amount of users")
            .register(registry);
    }

    private double count() {
        try {
            return userManager.getTotalCount();
        } catch (ServerException e) {
            return Double.NaN;
        }
    }
}
```

Alternatively, the metric can be stored with a reference and updated manually in other place in the code.

Additional resources

- [Metric and label naming for Prometheus](#)
- [Metric types for Prometheus](#)

## CHAPTER 6. TRACING CODEREADY WORKSPACES

Tracing helps gather timing data to troubleshoot latency problems in microservice architectures and helps to understand a complete transaction or workflow as it propagates through a distributed system. Every transaction may reflect performance anomalies in an early phase when new services are being introduced by independent teams.

Tracing the CodeReady Workspaces application may help analyze the execution of various operations, such as workspace creations, workspace startup, breaking down the duration of sub-operations executions, helping finding bottlenecks and improve the overall state of the platform.

Tracers live in applications. They record timing and metadata about operations that take place. They often instrument libraries, so that their use is transparent to users. For example, an instrumented web server records when it received a request and when it sent a response. The trace data collected is called a span. A span has a context that contains information such as trace and span identifiers and other kinds of data that can be propagated down the line.

### 6.1. TRACING API

CodeReady Workspaces utilizes [OpenTracing API](#) - a vendor-neutral framework for instrumentation. This means that if a developer wants to try a different tracing back end, then instead of repeating the whole instrumentation process for the new distributed tracing system, the developer can simply change the configuration of the tracer back end.

### 6.2. TRACING BACK END

By default, CodeReady Workspaces uses Jaeger as the tracing back end. Jaeger was inspired by Dapper and OpenZipkin, and it is a distributed tracing system released as open source by Uber Technologies. Jaeger extends a more complex architecture for a larger scale of requests and performance.

### 6.3. INSTALLING THE JAEGER TRACING TOOL

The following sections describe the installation methods for the Jaeger tracing tool. Jaeger can then be used for gathering metrics in CodeReady Workspaces.

Installation methods available:

- [Section 6.3.1, "Installing Jaeger using OperatorHub on OpenShift 4"](#)
- [Section 6.3.2, "Installing Jaeger using CLI on OpenShift 4"](#)

For tracing a CodeReady Workspaces instance using Jaeger, version 1.12.0 or above is required. For additional information about Jaeger, see the [Jaeger website](#).

#### 6.3.1. Installing Jaeger using OperatorHub on OpenShift 4

This section provide information about using Jaeger tracing tool for testing an evaluation purposes in production.

To install the Jaeger tracing tool from the OperatorHub interface in OpenShift Container Platform, follow the instructions below.

Prerequisites

- The user is logged in to the OpenShift Container Platform Web Console.
- A CodeReady Workspaces instance is available in a project.

### Procedure

1. Open the OpenShift Container Platform console.
2. From the left menu of the main OpenShift Container Platform screen, navigate to Operators → OperatorHub.
3. In the Search by keyword search bar, type **Jaeger Operator**.
4. Click the **Jaeger Operator** tile.
5. Click the **Install** button in the **Jaeger Operator** pop-up window.
6. Select the installation method: **A specific project on the cluster** where the CodeReady Workspaces is deployed and leave the rest in its default values.
7. Click the **Subscribe** button.
8. From the left menu of the main OpenShift Container Platform screen, navigate to the Operators → Installed Operators section.
9. Red Hat CodeReady Workspaces is displayed as an Installed Operator, as indicated by the **InstallSucceeded** status.
10. Click the **Jaeger Operator** name in the list of installed Operators.
11. Navigate to the **Overview** tab.
12. In the **Conditions** sections at the bottom of the page, wait for this message: **install strategy completed with no errors**.
13. **Jaeger Operator** and additional **Elasticsearch Operator** is installed.
14. Navigate to the Operators → Installed Operators section.
15. Click **Jaeger Operator** in the list of installed Operators.
16. The **Jaeger Cluster** page is displayed.
17. In the lower left corner of the window, click **Create Instance**
18. Click **Save**.
19. OpenShift creates the Jaeger cluster **jaeger-all-in-one-inmemory**.
20. Follow the steps in [Enabling metrics collection](#) to finish the procedure.

### 6.3.2. Installing Jaeger using CLI on OpenShift 4

This section provide information about using Jaeger tracing tool for testing an evaluation purposes.

To install the Jaeger tracing tool from a CodeReady Workspaces project in OpenShift Container Platform, follow the instructions in this section.

### Prerequisites

- The user is logged in to the OpenShift Container Platform web console.
- A instance of CodeReady Workspaces in an OpenShift Container Platform cluster.

### Procedure

1. In the CodeReady Workspaces installation project of the OpenShift Container Platform cluster, use the `oc` client to create a new application for the Jaeger deployment.

```
$ oc new-app -f / ${CHE_LOCAL_GIT_REPO}/deploy/openshift/templates/jaeger-all-in-one-template.yml:
```

```
--> Deploying template "<project_name>/jaeger-template-all-in-one" for "/home/user/crw-projects/crw/deploy/openshift/templates/jaeger-all-in-one-template.yml" to project <project_name>
```

```
Jaeger (all-in-one)
```

```
-----
```

```
Jaeger Distributed Tracing Server (all-in-one)
```

```
* With parameters:
```

```
* Jaeger Service Name=jaeger
```

```
* Image version=latest
```

```
* Jaeger Zipkin Service Name=zipkin
```

```
--> Creating resources ...
```

```
deployment.apps "jaeger" created
```

```
service "jaeger-query" created
```

```
service "jaeger-collector" created
```

```
service "jaeger-agent" created
```

```
service "zipkin" created
```

```
route.route.openshift.io "jaeger-query" created
```

```
--> Success
```

```
Access your application using the route: 'jaeger-query-<project_name>.apps.ci-ln-whx0352-d5d6b.origin-ci-int-aws.dev.rhcloud.com'
```

```
Run 'oc status' to view your app.
```

2. Using the Workloads → Deployments from the left menu of main OpenShift Container Platform screen, monitor the Jaeger deployment until it finishes successfully.
3. Select Networking → Routes from the left menu of the main OpenShift Container Platform screen, and click the URL link to access the Jaeger dashboard.
4. Follow the steps in [Enabling metrics collection](#) to finish the procedure.

## 6.4. ENABLING METRICS COLLECTION

### Prerequisites

- Installed Jaeger v1.12.0 or above. See instructions at [Section 6.3, “Installing the Jaeger tracing tool”](#)

## Procedure

For Jaeger tracing to work, enable the following environment variables in your CodeReady Workspaces deployment:

```
# Activating CodeReady Workspaces tracing modules
CHE_TRACING_ENABLED=true

# Following variables are the basic Jaeger client library configuration.
JAEGER_ENDPOINT="http://jaeger-collector:14268/api/traces"

# Service name
JAEGER_SERVICE_NAME="che-server"

# URL to remote sampler
JAEGER_SAMPLER_MANAGER_HOST_PORT="jaeger:5778"

# Type and param of sampler (constant sampler for all traces)
JAEGER_SAMPLER_TYPE="const"
JAEGER_SAMPLER_PARAM="1"

# Maximum queue size of reporter
JAEGER_REPORTER_MAX_QUEUE_SIZE="10000"
```

To enable the following environment variables:

1. In the `yaml` source code of the CodeReady Workspaces deployment, add the following configuration variables under `spec.server.customCheProperties`.

```
customCheProperties:
  CHE_TRACING_ENABLED: 'true'
  JAEGER_SAMPLER_TYPE: const
  DEFAULT_JAEGER_REPORTER_MAX_QUEUE_SIZE: '10000'
  JAEGER_SERVICE_NAME: che-server
  JAEGER_ENDPOINT: 'http://jaeger-collector:14268/api/traces'
  JAEGER_SAMPLER_MANAGER_HOST_PORT: 'jaeger:5778'
  JAEGER_SAMPLER_PARAM: '1'
```

2. Edit the `JAEGER_ENDPOINT` value to match the name of the Jaeger collector service in your deployment.  
From the left menu of the main OpenShift Container Platform screen, obtain the value of `JAEGER_ENDPOINT` by navigation to Networking → Services. Alternatively, execute the following `oc` command:

```
$ oc get services
```

The requested value is included in the service name that contains the `collector` string.

## Additional resources

- For additional information about custom environment properties and how to define them in CheCluster Custom Resource, see <https://access.redhat.com/documentation/en->

[us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/installation\\_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component\\_crw](https://www.redhat.com/en/technologies/cloud-management/containers/red-hat-codeready-workspaces/2.4/html-single/installation_guide/index#advanced-configuration-options-for-the-codeready-workspaces-server-component_crw).

- For custom configuration of Jaeger, see the list of [Jaeger client environment variables](#).

## 6.5. VIEWING CODEREADY WORKSPACES TRACES IN JAEGER UI

This section demonstrates how to utilize the Jaeger UI to overview traces of CodeReady Workspaces operations.

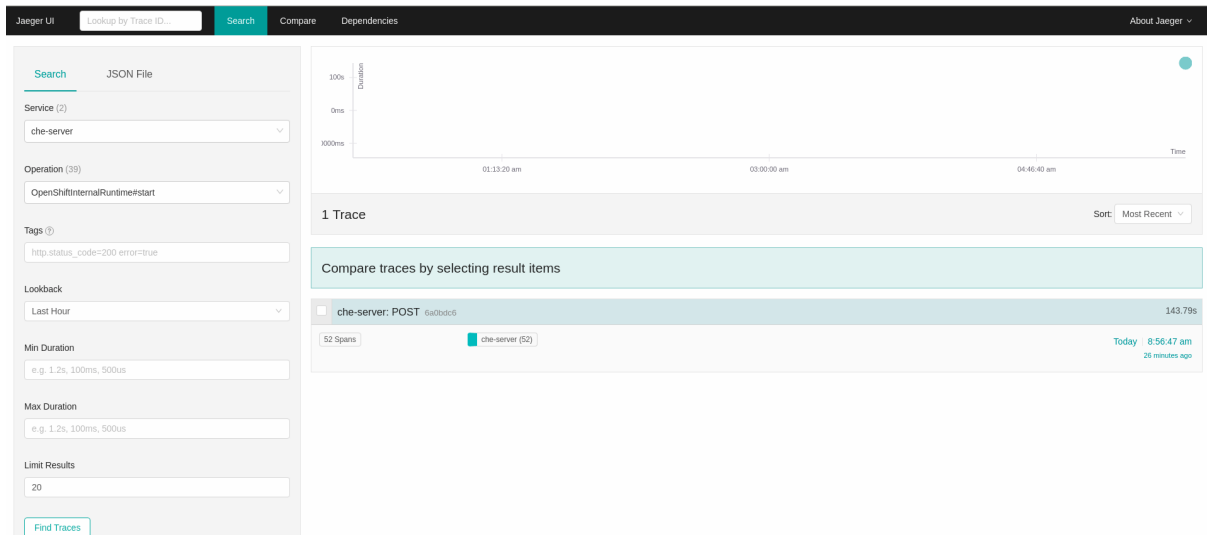
### Procedure

In this example, the CodeReady Workspaces instance has been running for some time and one workspace start has occurred.

To inspect the trace of the workspace start:


1. In the Search panel on the left, filter spans by the operation name (span name), tags, or time and duration.

Figure 6.1. Using Jaeger UI to trace CodeReady Workspaces



2. Select the trace to expand it and show the tree of nested spans and additional information about the highlighted span, such as tags or durations.

Figure 6.2. Expanded tracing tree



## 6.6. CODEREADY WORKSPACES TRACING CODEBASE OVERVIEW AND EXTENSION GUIDE

The core of the tracing implementation for CodeReady Workspaces is in the `che-core-tracing-core` and `che-core-tracing-web` modules.

All HTTP requests to the tracing API have their own trace. This is done by `TracingFilter` from the [OpenTracing library](#), which is bound for the whole server application. Adding `@Traced` annotation to methods causes the `TracingInterceptor` to add tracing spans for them.

### 6.6.1. Tagging

Spans may contain standard tags, such as operation name, span origin, error, and other tags that may help users with querying and filtering spans. Workspace-related operations (such as starting or stopping workspaces) have additional tags, including `userId`, `workspaceId`, and `stackId`. Spans created by `TracingFilter` also have an HTTP status code tag.

Declaring tags in a traced method is done statically by setting fields from the `TracingTags` class:

```
TracingTags.WORKSPACE_ID.set(workspace.getId());
```

`TracingTags` is a class where all commonly used tags are declared, as respective `AnnotationAware` tag implementations.

#### Additional resources

For more information about how to use Jaeger UI, visit Jaeger documentation: [Jaeger Getting Started Guide](#).

## CHAPTER 7. MANAGING USERS

This section describes how to configure authorization and authentication in Red Hat CodeReady Workspaces and how to administer user groups and users.

- [Section 7.1, “Configuring authorization”](#)
- [Section 7.2, “Removing user data”](#)

### 7.1. CONFIGURING AUTHORIZATION

#### 7.1.1. Authorization and user management

Red Hat CodeReady Workspaces uses [RH-SSO](#) to create, import, manage, delete, and authenticate users. RH-SSO uses built-in authentication mechanisms and user storage. It can use third-party identity management systems to create and authenticate users. Red Hat CodeReady Workspaces requires a RH-SSO token when you request access to CodeReady Workspaces resources.

Local users and imported federation users must have an email address in their profile.

The default RH-SSO credentials are **admin:admin**. You can use the **admin:admin** credentials when logging into Red Hat CodeReady Workspaces for the first time. It has system privileges.

#### Identifying the RH-SSO URL

Go to the OpenShift web console and navigate to the RH-SSO project.

#### 7.1.2. Configuring CodeReady Workspaces to work with RH-SSO

The deployment script configures RH-SSO. It creates a **codeready-public** client with the following fields:

- Valid Redirect URIs: Use this URL to access CodeReady Workspaces.
- Web Origins

The following are common errors when configuring CodeReady Workspaces to work with RH-SSO:

#### Invalid `redirectURI` error

Occurs when you access CodeReady Workspaces at **myhost**, which is an alias, and your original **CHE\_HOST** is **1.1.1.1**. If this error occurs, go to the RH-SSO administration console and ensure that the valid redirect URIs are configured.

#### CORS error

Occurs when you have an invalid web origin.

#### 7.1.3. Configuring RH-SSO tokens

A user token expires after 30 minutes by default.

You can change the following RH-SSO token settings:



Che 

General

Login

Keys

Email


Themes

Cache

Tokens

Client Registration

Security Defenses

Revoke Refresh Token  OFFSSO Session Idle  SSO Session Max  Offline Session Idle  Access Token Lifespan  Access Token Lifespan For Implicit Flow  Client login timeout  Login timeout  Login action timeout  User-Initiated Action Lifespan  Default Admin-Initiated Action Lifespan   

### 7.1.4. Setting up user federation

RH-SSO federates external user databases and supports LDAP and Active Directory. You can test the connection and authenticate users before choosing a storage provider.

See the [User storage federation](#) page in RH-SSO documentation to learn how to add a provider.

See the [LDAP and Active Directory](#) page in RH-SSO documentation to specify multiple LDAP servers.

### 7.1.5. Enabling authentication with social accounts and brokering

RH-SSO provides built-in support for GitHub, OpenShift, and most common social networks such as Facebook and Twitter.

See Instructions to [enable Login with GitHub](#).

You can also enable the SSH key and upload it to the CodeReady Workspaces users' GitHub accounts.

To enable this feature when you register a GitHub identity provider:

1. Set scope to `repo,user,write:public_key`.
2. Set store tokens and stored tokens readable to ON.

## GitHub

Settings Mappers

Redirect URI [?](#)

\* Client ID [?](#)

\* Client Secret [?](#)

Default Scopes [?](#)

Store Tokens [?](#)  ON

Stored Tokens Readable [?](#)  ON

Enabled [?](#)  ON

Disable User Info [?](#)  OFF

Trust Email [?](#)  OFF

Account Linking Only [?](#)  OFF

Hide on Login Page [?](#)  OFF

GUI order [?](#)

First Login Flow [?](#)  ▼

Post Login Flow [?](#)

## 3. Add a default read-token role.

Che ▼

Configure

- Realm Settings
- Clients
- Client
- Templates
- Roles**
- Identity
- Providers
- User
- Federation
- Authentication

Roles

Realm Roles **Default Roles**

Realm Roles Available Roles [?](#) Realm Default Roles [?](#)

user

Add selected »

offline\_access  
uma\_authorization

« Remove selected

Client Roles Available Roles [?](#) Client Default Roles [?](#)

broker ▼

read-token

Add selected »

« Remove selected

This is the default **delegated** OAuth service mode for multiuser CodeReady Workspaces. You can configure the OAuth service mode with the property `che.oauth.service_mode`.

## 7.1.6. Using protocol-based providers

RH-SSO supports [SAML v2.0](#) and [OpenID Connect v1.0](#) protocols. You can connect your identity provider systems if they support these protocols.

### 7.1.7. Managing users using RH-SSO

You can add, delete, and edit users in the user interface. See [RH-SSO User Management](#) for more information.

### 7.1.8. Configuring CodeReady Workspaces to use an external RH-SSO installation

By default, CodeReady Workspaces installation includes the deployment of a dedicated RH-SSO instance. However, using an external RH-SSO is also possible. This option is useful when a user has an existing RH-SSO instance with already-defined users, for example, a company-wide RH-SSO server used by several applications.

Table 7.1. Placeholders used in examples

<b>&lt;provider-realm-name&gt;</b>	Identity provider realm name intended for use by CodeReady Workspaces
<b>&lt;oidc-client-name&gt;</b>	Name of the <b>oidc</b> client defined in <b>&lt;provider-realm-name&gt;</b>
<b>&lt;auth-base-url&gt;</b>	Base URL of the external RH-SSO server

#### Prerequisites

- In the administration console of the external installation of RH-SSO, define a [realm](#) containing the users intended to connect to CodeReady Workspaces:

The screenshot shows the 'Realm-for-users' configuration page in the RH-SSO administration console. The left sidebar contains a navigation menu with options: Realm, Settings, Clients, Client Scopes, Roles, Identity, Providers, User, Federation, Authentication, and Manage. The main content area is titled 'Realm-for-users' and has tabs for 'General', 'Login', 'Keys', 'Email', 'Themes', 'Cache', 'Tokens', 'Client Registration', and 'Security Defenses'. The 'General' tab is active, showing the following configuration fields:

- Name:** realm-for-users
- Display name:** (empty)
- HTML Display name:** (empty)
- Frontend URL:** (empty)
- Enabled:** ON
- User-Managed Access:** OFF
- Endpoints:** OpenID Endpoint Configuration, SAML 2.0 Identity Provider Metadata

At the bottom of the configuration area are 'Save' and 'Cancel' buttons.

- In this realm, define an [OIDC client](#) that CodeReady Workspaces will use to authenticate the users. This is an example of such a client with the correct settings:

The screenshot shows the 'Public-client' configuration page in the administration console. The left sidebar contains navigation options like 'Configure', 'Manage', and 'Clients'. The main content area is titled 'Public-client' and has tabs for 'Settings', 'Roles', 'Client Scopes', 'Mappers', 'Scope', 'Revocation', and 'Sessions'. The 'Settings' tab is active, showing various configuration fields:

- Client ID: public-client
- Name: (empty)
- Description: (empty)
- Enabled: ON
- Consent Required: OFF
- Login Theme: (empty)
- Client Protocol: openid-connect
- Access Type: public
- Standard Flow Enabled: ON
- Implicit Flow Enabled: OFF
- Direct Access Grants Enabled: ON
- Root URL: (empty)
- Valid Redirect URIs:
 

http://che-eclipse-che.apps-crc.testing/*	-
https://che-eclipse-che.apps-crc.testing/*	-
	+
- Base URL: (empty)
- Admin URL: (empty)
- Web Origins:
 

http://che-eclipse-che.apps-crc.testing	-
https://che-eclipse-che.apps-crc.testing	-

## NOTE

- CodeReady Workspaces only supports **public** OIDC clients. Therefore, selecting the **openid-connect** Client Protocol option and the **public** Access Type option is recommended.
  - The list of Valid Redirect URIs must contain at least two URIs related to the CodeReady Workspaces server, one using the **http** protocol and the other **https**. These URIs must contain the base URL of the CodeReady Workspaces server, followed by **/\*** wildcards.
  - The list of Web Origins must contain at least two URIs related to the CodeReady Workspaces server, one using the **http** protocol and the other **https**. These URIs must contain the base URL of the CodeReady Workspaces server, without any path after the host. The number of URIs depends on the number of installed product tools.
- With CodeReady Workspaces that uses the default OpenShift OAuth support, user authentication relies on the integration of RH-SSO with OpenShift OAuth. This allows users to log in to CodeReady Workspaces with their OpenShift login and have their workspaces

created under personal OpenShift projects.

This requires setting up an OpenShift identity provider in RH-SSO. When using an external RH-SSO, set up the identity provider manually. For instructions, see the appropriate RH-SSO documentations for either link:[OpenShift 3](#)[OpenShift 3] or link:[OpenShift 4](#)[OpenShift 4].

- The configured identity provider has the options Store Tokens and Stored Tokens Readable enabled.

## Procedure

1. Set the following properties in the **CheCluster** Custom Resource (CR):

```
spec:
  auth:
    externalIdentityProvider: true
    identityProviderURL: <auth-base-url>
    identityProviderRealm: <provider-realm-name>
    identityProviderClientId: <oidc-client-name>
```

2. When installing CodeReady Workspaces with OpenShift OAuth support enabled, set the following properties in the **CheCluster** Custom Resource (CR):

```
spec:
  auth:
    openShifttoAuth: true
  # Note: only if the OpenShift identity provider alias is different from 'openshift-v3' or
  # 'openshift-v4'
  server:
    customCheProperties:
      CHE_INFRA_OPENSHIFT_OAUTH_IDENTITY_PROVIDER: <OpenShift identity
      provider alias>
```

### 7.1.9. Configuring SMTP and email notifications

Red Hat CodeReady Workspaces does not provide any pre-configured SMTP servers.

To enable SMTP servers in RH-SSO:

1. Go to **che realm settings > Email**.
2. Specify the host, port, username, and password.

Red Hat CodeReady Workspaces uses the default theme for email templates for registration, email confirmation, password recovery, and failed login.

## 7.2. REMOVING USER DATA

### 7.2.1. Removing user data according to GDPR

The General Data Protection Regulation ([GDPR](#)) law enforces the right for individuals to have personal data erased.

The following procedure describes how to remove a user's data from a cluster and the RH-SSO database.



## NOTE

The following commands use the default OpenShift project, `openshift-workspaces`, as a user's example for the `-n` option.

## Prerequisites

- A user or an administrator authorization token. To delete any other data except the data bound to a user account, `admin` privileges are required. The `admin` is a special CodeReady Workspaces administrator account pre-created and enabled using the `CHE_SYSTEM_ADMIN_NAME` and `CHE_SYSTEM_SUPER_PRIVILEGED_MODE = true` Custom Resource definitions.

```
spec:
  server:
    customCheProperties:
      CHE_SYSTEM_SUPER_PRIVILEGED_MODE: 'true'
      CHE_SYSTEM_ADMIN_NAME: '<admin-name>'
```

If needed, use commands below for creating the `admin` user:

```
$ oc patch checluster codeready-workspaces \
  --type merge \
  -p '{"spec": {"server": {"customCheProperties":
{"CHE_SYSTEM_SUPER_PRIVILEGED_MODE": "true"} } }' \
  -n openshift-workspaces
```

```
$ oc patch checluster codeready-workspaces \
  --type merge \
  -p '{"spec": {"server": {"customCheProperties": {"CHE_SYSTEM_ADMIN_NAME":
"<admin-name>" } } }' \
  -n openshift-workspaces
```



## NOTE

All system permissions are granted to the administrative user who is configured in the `CHE_SYSTEM_ADMIN_NAME` property (the default is `admin`). The system permissions are granted when the CodeReady Workspaces server starts. If the user is not present in the CodeReady Workspaces user database, it happens after the first user's login.

Authorization token privileges:

- `admin` - Can delete all personal data of all users
- `user` - Can delete only the data related to the user

- A user or an administrator is logged in the OpenShift cluster with deployed CodeReady Workspaces.
- A user ID is obtained. Get the user ID using the commands below:

- o For the current user:

```
$ curl -X GET \
  --header 'Authorization: Bearer <user-token>' \
  'https://<codeready-<openshift_deployment_name>.<domain_name>>/api/user'
```

- o To find a user by name:

```
$ curl -X GET \
  --header 'Authorization: Bearer <user-token>' \
  'https://<codeready-<openshift_deployment_name>.<domain_name>>/api/user/find?
  name=<username>'
```

- o To find a user by email:

```
$ curl -X GET \
  --header 'Authorization: Bearer <user-token>' \
  'https://<codeready-<openshift_deployment_name>.<domain_name>>/api/user/find?
  email=<email>'
```

### Example of obtaining a user ID

This example uses `vparfono` as a local user name.

```
$ curl -X GET \
  --header 'Authorization: Bearer <user-token>' \
  'https://che-vp-che.apps.che-dev.x6e0.p1.openshiftapps.com/api/user/find?
  name=vparfono'
```

The user ID is at the bottom of the curl command output.

```
{
  "name": "vparfono",
  "links": [
    {
      .
      .
      .
    }
  ],
  "email": "vparfono@redhat.com",
  "id": "921b6f33-2657-407e-93a6-fb14cf2329ce"
}
```

### Procedure

1. Update the **codeready-workspaces CheCluster Custom Resource (CR)** definition to permit the removal of a user's data from the RH-SSO database:

```
$ oc patch checluster/codeready-workspaces \
  --patch '{"spec":{"server":{"customCheProperties":
{"CHE_KEYCLOAK_CASCADE__USER__REMOVAL__ENABLED": "true"}}}}' \
  --type=merge -n openshift-workspaces
```

## 2. Remove the data using the API:

```
$ curl -i -X DELETE \  
  --header 'Authorization: Bearer <user-token>' \  
  https://<codeready-<openshift_deployment_name>.<domain_name>/api/user/<user-id>
```

### Verification

Running the following command returns code **204** as the API response:

```
$ curl -i -X DELETE \  
  --header 'Authorization: Bearer <user-token>' \  
  https://<codeready-<openshift_deployment_name>.<domain_name>/api/user/<user-id>
```

### Additional resources

To remove the data of all users, follow the instructions for

[https://access.redhat.com/documentation/en-us/red\\_hat\\_codeready\\_workspaces/2.4/html-single/installation\\_guide/index#uninstalling-codeready-workspaces\\_crw](https://access.redhat.com/documentation/en-us/red_hat_codeready_workspaces/2.4/html-single/installation_guide/index#uninstalling-codeready-workspaces_crw).



## CHAPTER 8. SECURING CODEREADY WORKSPACES

This section describes all aspects of user authentication, types of authentication, and permissions models on the CodeReady Workspaces server and its workspaces.

- [Section 8.1, “Authenticating users”](#)
- [Section 8.2, “Authorizing users”](#)

### 8.1. AUTHENTICATING USERS

This document covers all aspects of user authentication in Red Hat CodeReady Workspaces, both on the CodeReady Workspaces server and in workspaces. This includes securing all REST API endpoints, WebSocket or JSON RPC connections, and some web resources.

All authentication types use the [JWT open standard](#) as a container for transferring user identity information. In addition, CodeReady Workspaces server authentication is based on the [OpenID Connect](#) protocol implementation, which is provided by default by [RH-SSO](#).

Authentication in workspaces implies the issuance of self-signed per-workspace JWT tokens and their verification on a dedicated service based on [JWTProxy](#).

#### 8.1.1. Authenticating to the CodeReady Workspaces server

##### 8.1.1.1. Authenticating to the CodeReady Workspaces server using OpenID

OpenID authentication on the CodeReady Workspaces server implies the presence of an external OpenID Connect provider and has the following main steps:

- Authenticate the user through a JWT token that is retrieved from an HTTP request or, in case of a missing or invalid token, redirect the user to the RH-SSO login page.
- Send authentication tokens in an Authorization header. In limited cases, when it is impossible to use the Authorization header, the token can be sent in the token query parameter. Example: OAuth authentication initialization.
- Compose an internal **subject** object that represents the current user inside the CodeReady Workspaces server code.



#### NOTE

The only supported and tested OpenID provider is RH-SSO.

#### Procedure

To authenticate to the CodeReady Workspaces server using OpenID authentication:

1. Request the OpenID settings service where clients can find all the necessary URLs and properties of the OpenId provider, such as **jwtks.endpoint**, **token.endpoint**, **logout.endpoint**, **realm.name**, or **client\_id** returned in the JSON format.
2. The service URL is **https://codeready-<openshift\_deployment\_name>.<domain\_name>/api/keycloak/settings**, and it is only available in the CodeReady Workspaces multiuser mode. The presence of the service in the URL confirms that the authentication is enabled in the current deployment.

Example output:

```
{
  "che.keycloak.token.endpoint":
  "http://172.19.20.9:5050/auth/realms/che/protocol/openid-connect/token",
  "che.keycloak.profile.endpoint": "http://172.19.20.9:5050/auth/realms/che/account",
  "che.keycloak.client_id": "che-public",
  "che.keycloak.auth_server_url": "http://172.19.20.9:5050/auth",
  "che.keycloak.password.endpoint":
  "http://172.19.20.9:5050/auth/realms/che/account/password",
  "che.keycloak.logout.endpoint":
  "http://172.19.20.9:5050/auth/realms/che/protocol/openid-connect/logout",
  "che.keycloak.realm": "che"
}
```

The service allows downloading the JavaScript client library to interact with the provider using the `https://codeready-<openshift_deployment_name>.<domain_name>/api/keycloak/OIDCKeycloak.js` URL.

3. Redirect the user to the appropriate provider's login page with all the necessary parameters, including `client_id` and the return redirection path. This can be done with any client library (JS or Java).
4. When the user is logged in to the provider, the client side-code is obtained, and the JWT token has validated the token, the creation of the **subject** begins.

The verification of the token signature occurs in two main steps:

1. Authentication: The token is extracted from the Authorization header or from the `token` query parameter and is parsed using the public key retrieved from the provider. In case of expired, invalid, or malformed tokens, a **403** error is sent to the user. The minimal use of the query parameter is recommended, due to its support limitations or complete removal in upcoming versions.

If the validation is successful, the parsed form of the token is passed to the environment initialization step:

2. Environment initialization: The filter extracts data from the JWT token claims, creates the user in the local database if it is not yet available, and constructs the **subject** object and sets it into the per-request `EnvironmentContext` object, which is statically accessible everywhere.

If the request was made using only a JWT token, the following single authentication filter is used:

`org.eclipse.che.multiuser.machine.authentication.server.MachineLoginFilter`: The filter finds the user that the `userId` token belongs to, retrieves the user instance, and sets the principal to the session. The CodeReady Workspaces server-to-server requests are performed using a dedicated request factory that signs every request with the current subject token obtained from the `EnvironmentContext` object.



## NOTE

## Providing user-specific data

Since RH-SSO may store user-specific information (first and last name, phone number, job title), there is a special implementation of the ProfileDao that can provide this data to consumers. The implementation is read-only, so users cannot perform create and update operations.

## 8.1.1.1.1. Obtaining the token from credentials through RH-SSO

Clients that cannot run JavaScript or other clients (such as command-line clients or Selenium tests) must request the authorization token directly from RH-SSO.

To obtain the token, send a request to the token endpoint with the username and password credentials. This request can be schematically described as the following cURL request:

```
$ curl --insecure --data "grant_type=password&client_id=codeready-
public&username=<USERNAME>&password=<PASSWORD>" \ 1 2
https://<keycloak_host>/auth/realms/codeready/protocol/openid-connect/token 3
```

- 1 Red Hat CodeReady Workspaces username
- 2 Red Hat CodeReady Workspaces user's password
- 3 RH-SSO host

The CodeReady Workspaces dashboard uses a customized RH-SSO login page and an authentication mechanism based on `grant_type=authorization_code`. It is a two-step authentication process:

1. Logging in and obtaining the authorization code.
2. Obtaining the token using this authorization code.

## 8.1.1.1.2. Obtaining the token from the OpenShift token through RH-SSO

When CodeReady Workspaces was installed on OpenShift using the Operator, and the OpenShift OAuth integration is enabled, as it is by default, the user's CodeReady Workspaces authentication token can be retrieved from the user's OpenShift token.

To retrieve the authentication token from the OpenShift token, send a schematically described cURL request to the OpenShift token endpoint:

```
$ curl --insecure -X POST \
-d "client_id=codeready-public" \
-d "subject_token=<USER_OPENSHIFT_TOKEN>" \ 1
-d "subject_issuer=<OPENSHIFT_IDENTITY_PROVIDER_NAME>" \ 2
--data-urlencode "grant_type=urn:ietf:params:oauth:grant-type:token-exchange" \
--data-urlencode "subject_token_type=urn:ietf:params:oauth:token-type:access_token" \
https://<KEYCKLOAK_HOST>/auth/realms/codeready/protocol/openid-connect/token 3
```

- 1 The token retrieved by the end-user with the command `oc whoami --show-token`

2 **openshift-v4** for OpenShift 4.x and **openshift-v3** for OpenShift 3.11

3 **RH-SSO** host



#### WARNING

Before using this token exchange feature, it is required for an end user to be interactively logged in at least once to the CodeReady Workspaces Dashboard using the OpenShift login page. This step is needed to link the OpenShift and RH-SSO user accounts properly and set the required user profile information.

### 8.1.1.2. Authenticating to the CodeReady Workspaces server using other authentication implementations

This procedure describes how to use an OpenID Connect (OIDC) authentication implementation other than RH-SSO.

#### Procedure

1. Update the authentication configuration parameters that are stored in the **multiuser.properties** file (such as client ID, authentication URL, realm name).
2. Write a single filter or a chain of filters to validate tokens, create the user in the CodeReady Workspaces dashboard, and compose the **subject** object.
3. If the new authorization provider supports the OpenID protocol, use the OIDC JS client library available at the settings endpoint because it is decoupled from specific implementations.
4. If the selected provider stores additional data about the user (first and last name, job title), it is recommended to write a provider-specific ProfileDao implementation that provides this information.

### 8.1.1.3. Authenticating to the CodeReady Workspaces server using OAuth

For easy user interaction with third-party services, the CodeReady Workspaces server supports OAuth authentication. OAuth tokens are also used for GitHub-related plug-ins.

OAuth authentication has two main flows:

#### delegated

Default. Delegates OAuth authentication to RH-SSO server.

#### embedded

Uses built-in CodeReady Workspaces server mechanism to communicate with OAuth providers.

To switch between the two implementations, use the **che.oauth.service\_mode=<embedded/delegated>** configuration property.

The main REST endpoint in the OAuth API is **/api/oauth**, which contains:

- An authentication method, `/authenticate`, that the OAuth authentication flow can start with.
- A callback method, `/callback`, to process callbacks from the provider.
- A token GET method, `/token`, to retrieve the current user's OAuth token.
- A token DELETE method, `/token`, to invalidate the current user's OAuth token.
- A GET method, `/`, to get the list of configured identity providers.

#### 8.1.1.4. Using Swagger or REST clients to execute queries

The user's RH-SSO token is used to execute queries to the secured API on the user's behalf through REST clients. A valid token must be attached as the Request header or the `?token=$token` query parameter.


Access the CodeReady Workspaces Swagger interface at `https://codeready-  
<openshift_deployment_name>.<domain_name>/swagger`. The user must be signed in through RH-SSO, so that the access token is included in the Request header.

#### 8.1.2. Authenticating in a CodeReady Workspaces workspace

Workspace containers may contain services that must be protected with authentication. Such protected services are called secure. To secure these services, use a machine authentication mechanism.

JWT tokens avoid the need to pass RH-SSO tokens to workspace containers (which can be insecure). Also, RH-SSO tokens may have a relatively shorter lifetime and require periodic renewals or refreshes, which is difficult to manage and keep in sync with the same user session tokens on clients.

Figure 8.1. Authentication inside a workspace



### 8.1.2.1. Creating secure servers

To create secure servers in CodeReady Workspaces workspaces, set the `secure` attribute of the endpoint to `true` in the `dockerimage` type component in the devfile.

Devfile snippet for a secure server

```
components:
- type: dockerimage
  endpoints:
  - attributes:
    secure: 'true'
```

### 8.1.2.2. Workspace JWT token

Workspace tokens are JSON web tokens ([JWT](#)) that contain the following information in their claims:

- **uid:** The ID of the user who owns this token
- **uname:** The name of the user who owns this token

- **wsid**: The ID of a workspace which can be queried with this token

Every user is provided with a unique personal token for each workspace. The structure of a token and the signature are different than they are in RH-SSO. The following is an example token view:

```
# Header
{
  "alg": "RS512",
  "kind": "machine_token"
}
# Payload
{
  "wsid": "workspacekrh99xjenek3h571",
  "uid": "b07e3a58-ed50-4a6e-be17-fcf49ff8b242",
  "uname": "john",
  "jti": "06c73349-2242-45f8-a94c-722e081bb6fd"
}
# Signature
{
  "value": "RSASHA256(base64UrlEncode(header) + . + base64UrlEncode(payload))"
}
```

The SHA-256 cipher with the RSA algorithm is used for signing JWT tokens. It is not configurable. Also, there is no public service that distributes the public part of the key pair with which the token is signed.

### 8.1.2.3. Machine token validation

The validation of machine tokens (JWT tokens) is performed using a dedicated per-workspace service with **JWTProxy** running on it in a separate Pod. When the workspace starts, this service receives the public part of the SHA key from the CodeReady Workspaces server. A separate verification endpoint is created for each secure server. When traffic comes to that endpoint, **JWTProxy** tries to extract the token from the cookies or headers and validates it using the public-key part.

To query the CodeReady Workspaces server, a workspace server can use the machine token provided in the **CHE\_MACHINE\_TOKEN** environment variable. This token is the user's who starts the workspace. The scope of such requests is restricted to the current workspace only. The list of allowed operations is also strictly limited.

## 8.2. AUTHORIZING USERS

User authorization in CodeReady Workspaces is based on the permissions model. Permissions are used to control the allowed actions of users and establish a security model. Every request is verified for the presence of the required permission in the current user subject after it passes authentication. You can control resources managed by CodeReady Workspaces and allow certain actions by assigning permissions to users.

Permissions can be applied to the following entities:

- Workspace
- System

All permissions can be managed using the provided REST API. The APIs are documented using Swagger at [### 8.2.1. CodeReady Workspaces workspace permissions](https://codeready-<code><openshift_deployment_name></code>.<code><domain_name>/swagger/#!/permissions</code>.</a></p>
</div>
<div data-bbox=)

The user who creates a workspace is the workspace owner. By default, the workspace owner has the following permissions: **read**, **use**, **run**, **configure**, **setPermissions**, and **delete**. Workspace owners can invite users into the workspace and control workspace permissions for other users.

The following permissions are associated with workspaces:

Table 8.1. CodeReady Workspaces workspace permissions

Permission	Description
read	Allows reading the workspace configuration.
use	Allows using a workspace and interacting with it.
run	Allows starting and stopping a workspace.
configure	Allows defining and changing the workspace configuration.
setPermissions	Allows updating the workspace permissions for other users.
delete	Allows deleting the workspace.

### 8.2.2. CodeReady Workspaces system permissions

CodeReady Workspaces system permissions control aspects of the whole CodeReady Workspaces installation. The following permissions are applicable to the system:

Table 8.2. CodeReady Workspaces system permission

Permission	Description
manageSystem	Allows control of the system and workspaces.
setPermissions	Allows updating the permissions for users on the system.
manageUsers	Allows creating and managing users.
monitorSystem	Allows accessing endpoints used for monitoring the state of the server.

All system permissions are granted to the administrative user who is configured in the



**CHE\_SYSTEM\_ADMIN\_\_NAME** property (the default `isadmin`). The system permissions are granted when the CodeReady Workspaces server starts. If the user is not present in the CodeReady Workspaces user database, it happens after the first user's login.

### 8.2.3. manageSystem permission

Users with the `manageSystem` permission have access to the following services:

Path	HTTP Method	Description
<code>/resource/free/</code>	GET	Get free resource limits.
<code>/resource/free/{accountId}</code>	GET	Get free resource limits for the given account.
<code>/resource/free/{accountId}</code>	POST	Edit free resource limit for the given account.
<code>/resource/free/{accountId}</code>	DELETE	Remove free resource limit for the given account.
<code>/installer/</code>	POST	Add installer to the registry.
<code>/installer/{key}</code>	PUT	Update installer in the registry.
<code>/installer/{key}</code>	DELETE	Remove installer from the registry.
<code>/logger/</code>	GET	Get logging configurations in the CodeReady Workspaces server.
<code>/logger/{name}</code>	GET	Get configurations of logger by its name in the CodeReady Workspaces server.
<code>/logger/{name}</code>	PUT	Create logger in the CodeReady Workspaces server.
<code>/logger/{name}</code>	POST	Edit logger in the CodeReady Workspaces server.
<code>/resource/{accountId}/details</code>	GET	Get detailed information about resources for the given account.
<code>/system/stop</code>	POST	Shutdown all system services, prepare CodeReady Workspaces to stop.

### 8.2.4. monitorSystem permission

Users with the `monitorSystem` permission have access to the following services.

Path	HTTP Method	Description
<code>/activity</code>	GET	Get workspaces in a certain state for a certain amount of time.

### 8.2.5. Listing CodeReady Workspaces permissions

To list CodeReady Workspaces permissions that apply to a specific resource, perform the **GET** `/permissions` request.

To list the permissions that apply to a user, perform the **GET** `/permissions/{domain}` request.

To list the permissions that apply to all users, perform the **GET** `/permissions/{domain}/all` request. The user must have `manageSystem` permissions to see this information.

The suitable domain values are:

- `system`
- `organization`
- `workspace`



#### NOTE

The domain is optional. If no domain is specified, the API returns all possible permissions for all the domains.

### 8.2.6. Assigning CodeReady Workspaces permissions

To assign permissions to a resource, perform the **POST** `/permissions` request. The suitable domain values are:

- `system`
- `organization`
- `workspace`

The following is a message body that requests permissions for a user with a `userId` to a workspace with a `workspaceID`:

Requesting CodeReady Workspaces user permissions

```
{
  "actions": [
    "read",
    "use",
    "run",
    "configure",
    "setPermissions"
  ],
```

```
"userId": "userID",  
"domainId": "workspace",  
"instanceId": "workspaceID"  
}
```

- 1 The `userId` parameter is the ID of the user that has been granted certain permissions.
- 2 The `instanceId` parameter is the ID of the resource that retrieves the permission for all users.

### 8.2.7. Sharing CodeReady Workspaces permissions

A user with `setPermissions` privileges can share a workspace and `grantread`, `use`, `run`, `configure`, or `setPermissions` privileges for other users.

#### Procedure

To share workspace permissions:

1. Select a workspace in the user dashboard.
2. Navigate to the Share tab and enter the email IDs of the users. Use commas or spaces as separators for multiple emails.

## CHAPTER 9. BACKUP AND DISASTER RECOVERY

This section describes aspects of the CodeReady Workspaces backup and disaster recovery.

- [Section 9.1, “External database setup”](#)
- [Section 9.2, “Persistent Volumes backups”](#)

### 9.1. EXTERNAL DATABASE SETUP

The PostgreSQL database is used by the CodeReady Workspaces server for persisting data about the state of CodeReady Workspaces. It contains information about user accounts, workspaces, preferences, and other details.

By default, the CodeReady Workspaces Operator creates and manages the database deployment.

However, the CodeReady Workspaces Operator does not support full life-cycle capabilities, such as backups and recovery.

For a business-critical setup, configure an external database with the following recommended disaster-recovery options:

- High Availability (HA)
- Point In Time Recovery (PITR)

Configure an external PostgreSQL instance on-premises or use a cloud service, such as Amazon Relational Database Service (Amazon RDS). With Amazon RDS, it is possible to deploy production databases in a Multi-Availability Zone configuration for a resilient disaster recovery strategy with daily and on-demand snapshots.

The recommended configuration of the example database is:

Parameter	Value
Instance class	db.t2.small
vCPU	1
RAM	2 GB
Multi-az	true, 2 replicas
Engine version	9.6.11
TLS	enabled
Automated backups	enabled (30 days)

#### 9.1.1. Configuring external PostgreSQL

## Procedure

1. Use the following SQL script to create user and database for the CodeReady Workspaces server to persist workspaces metadata etc:

```
CREATE USER <database-user> WITH PASSWORD '<database-password>' 1 2
CREATE DATABASE <database> 3
GRANT ALL PRIVILEGES ON DATABASE <database> TO <database-user>
ALTER USER <database-user> WITH SUPERUSER
```

- 1 CodeReady Workspaces server database username
- 2 CodeReady Workspaces server database password
- 3 CodeReady Workspaces server database name

2. Use the following SQL script to create database for RH-SSO back end to persist user information:

```
CREATE USER <identity-database-user> WITH PASSWORD '<identity-database-
password>' 1 2
CREATE DATABASE <identity-database> 3
GRANT ALL PRIVILEGES ON DATABASE <identity-database> TO <identity-database-
user>
```

- 1 RH-SSO database username
- 2 RH-SSO database password
- 3 RH-SSO database name

## 9.1.2. Configuring CodeReady Workspaces to work with an external PostgreSQL

### Prerequisites

- The `oc` tool is available.

### Procedure

1. Pre-create a project for CodeReady Workspaces:

```
$ oc create namespace openshift-workspaces
```

2. Create a secret to store CodeReady Workspaces server database credentials:

```
$ oc create secret generic <server-database-credentials> \ 1
--from-literal=user=<database-user> \ 2
--from-literal=password=<database-password> \ 3
-n openshift-workspaces
```

- 1 Secret name to store CodeReady Workspaces server database credentials

- 2 CodeReady Workspaces server database username
- 3 CodeReady Workspaces server database password

### 3. Create a secret to store RH-SSO database credentials:

```
$ oc create secret generic <identity-database-credentials> \ 1
--from-literal=user=<identity-database-user> \ 2
--from-literal=password=<identity-database-password> \ 3
-n openshift-workspaces
```

- 1 Secret name to store RH-SSO database credentials
- 2 RH-SSO database username
- 3 RH-SSO database password

### 4. To make the Operator skip deploying a database and pass connection details of an existing database to a CodeReady Workspaces server set the following values in the Custom Resource:

```
spec:
  database:
    externalDb: true
    chePostgresHostName: <hostname> 1
    chePostgresPort: <port> 2
    chePostgresSecret: <server-database-credentials> 3
    chePostgresDb: <database> 4
  spec:
    auth:
      identityProviderPostgresSecret: <identity-database-credentials> 5
```

- 1 External database hostname
- 2 External database port
- 3 Secret name with CodeReady Workspaces server database credentials
- 4 CodeReady Workspaces server database username
- 5 Secret name with RH-SSO database credentials

#### Additional resources

- [PostgreSQL](#)
- [RDS](#)

## 9.2. PERSISTENT VOLUMES BACKUPS

Persistent Volumes (PVs) store the CodeReady Workspaces workspace data similarly to how workspace data is stored for desktop IDEs on the local hard disk drive.

To prevent data loss, back up PVs periodically. The recommended approach is to use storage-agnostic tools for backing up and restoring OpenShift resources, including PVs.

### 9.2.1. Recommended backup tool: Velero

Velero is an open-source tool for backing up OpenShift applications and their PVs. Velero allows you to:

- Deploy in the cloud or on premises.
- Back up the cluster and restore in case of data loss.
- Migrate cluster resources to other clusters.
- Replicate a production cluster to development and testing clusters.



#### NOTE

Alternatively, you can use backup solutions dependent on the underlying storage system. For example, solutions that are Gluster or Ceph-specific.

#### Additional resources

- [Persistent Volumes documentation](#)
- [Gluster documentation](#)
- [Ceph documentation](#)
- [Velero on GitHub](#)

## CHAPTER 10. CACHING IMAGES FOR FASTER WORKSPACE START

This section describes installing the [Image Puller](#) on a CodeReady Workspaces cluster to cache images on cluster nodes.

### 10.1. IMAGE PULLER OVERVIEW

Slow starts of Red Hat CodeReady Workspaces workspaces may be caused by waiting for the underlying cluster to pull images used in workspaces from remote registries. As such, pre-pulling images can improve start times significantly. The *Image Puller* can be used to pre-pull images and shorten workspace start times.

The Image Puller is an additional deployment that runs alongside Red Hat CodeReady Workspaces. Given a list of images to pre-pull, the application runs inside a cluster and creates a *DaemonSet* that pulls the images on each node.



#### NOTE

The minimal requirement for an image to be pre-pulled is the availability of the `sleep` command, which means that **FROM scratch** images (for example, 'che-machine-exec') are currently not supported. Also, images that mount volumes in the dockerfile are not supported for pre-pulling on OpenShift.

The application can be deployed:

- using OperatorHub or installing the [Kubernetes Image Puller Operator](#)
- by processing and applying OpenShift templates.

The Image Puller loads its configuration from a **ConfigMap** with the following available parameters:

Table 10.1. Image Puller default parameters

Parameter	Usage	Default
<b>CACHING_INTERVAL_HOURS</b>	Interval, in hours, between checking health of DaemonSets	<b>"1"</b>
<b>CACHING_MEMORY_REQUEST</b>	The memory request for each cached image when the puller is running	<b>10Mi</b>
<b>CACHING_MEMORY_LIMIT</b>	The memory limit for each cached image when the puller is running	<b>20Mi</b>
<b>CACHING_CPU_REQUEST</b>	The CPU request for each cached image when the puller is running	<b>.05</b>
<b>CACHING_CPU_LIMIT</b>	The CPU limit for each cached image when the puller is running	<b>.2</b>



Parameter	Usage	Default
<b>DAEMONSET_NAME</b>	Name of DaemonSet to be created	<b>kubernetes-image-puller</b>
<b>NAMESPACE</b>	Namespace where DaemonSet is to be created	<b>k8s-image-puller</b>
<b>IMAGES</b>	List of images to be cached, in the format <b>&lt;name&gt;=&lt;image&gt;;</b> ...	Contains a default list of images. Before deploying, fill this with the images that fit the current requirements
<b>NODE_SELECTOR</b>	Node selector applied to the Pods created by the DaemonSet	<b>'{}'</b>

The default memory requests and limits ensure that the container has enough memory to start. When changing **CACHING\_MEMORY\_REQUEST** or **CACHING\_MEMORY\_LIMIT**, you will need to consider the total memory allocated to the DaemonSet Pods in the cluster:

**(memory limit) \* (number of images) \* (number of nodes in the cluster)**

For example, running the Image Puller that caches 5 images on 20 nodes, with a container memory limit of **20Mi** requires **2000Mi** of memory.

## 10.2. DEPLOYING IMAGE PULLER USING THE OPERATOR

The recommended way to deploy the Image Puller is through the [Operator](#).

### 10.2.1. Installing the Image Puller on OpenShift using OperatorHub

#### Prerequisites

- A project in your cluster to host the Image Puller. This document uses the project **image-puller** as an example.

#### Procedure

1. Navigate to your OpenShift cluster console, navigate to **Operators** → **OperatorHub**.
2. Use the **Filter by keyword** box to search for **Kubernetes Image Puller Operator**. Click the **Kubernetes Image Puller Operator**.
3. Read the description of the Operator. Click **Continue** → **Install**.
4. Select **A specific project on the cluster** for the **Installation Mode**. In the drop-down find the project you created to install the Image Puller. Click **Subscribe**.
5. Wait for the Kubernetes Image Puller Operator to install. Click the **KubernetesImagePuller** → **Create instance**.

6. In a redirected window with a YAML editor, make modifications to the **KubernetesImagePuller Custom Resource** and click **Create**.
7. Navigate to the **Workloads** and **Pods** menu in the project and verify that the Image Puller is available.

### 10.3. DEPLOYING IMAGE PULLER USING OPENSIFT TEMPLATES

The Image Puller repository contains OpenShift templates for deploying on OpenShift.

#### Prerequisites

- A running OpenShift cluster.
- The **oc** tool is available.

The following parameters are available to further configure the OpenShift templates:

Table 10.2. Parameters for installing with OpenShift templates

Value	Usage	Default
<b>DAEMONSET_NAME</b>	The value of <b>DAEMONSET_NAME</b> to set in the ConfigMap	<b>kubernetes-image-puller</b>
<b>IMAGE</b>	Image used for the <b>kubernetes-image-puller</b> deployment	<b>registry.redhat.io/codeready-workspaces/imagepuller-rhel8:2.4</b>
<b>IMAGE_TAG</b>	The image tag to pull	<b>2.4</b>
<b>SERVICEACCOUNT_NAME</b>	The name of the ServiceAccount used by the deployment (created as part of installation)	<b>k8s-image-puller</b>
<b>CACHING_INTERVAL_HOURS</b>	The value of <b>CACHING_INTERVAL_HOURS</b> to set in the ConfigMap	<b>"1"</b>
<b>CACHING_INTERVAL_REQUEST</b>	The value of <b>CACHING_MEMORY_REQUEST</b> to set in the ConfigMap	<b>"10Mi"</b>
<b>CACHING_INTERVAL_LIMIT</b>	The value of <b>CACHING_MEMORY_LIMIT</b> to set in the ConfigMap	<b>"20Mi"</b>
<b>NODE_SELECTOR</b>	The value of <b>NODE_SELECTOR</b> to set in the ConfigMap	<b>"{}"</b>

See [Table 10.1, “Image Puller default parameters”](#) for more information about configuration values, such as `DAEMONSET_NAME`, `CACHING_INTERVAL_HOURS`, and `CACHING_MEMORY_REQUEST`.

Table 10.3. List of recommended images to pre-pull

Image	URL	Tag
theia-rhel8	codeready-workspaces/theia-rhel8	2.4
theia-endpoint-rhel8	theia-endpoint-image	2.4
pluginbroker-metadata-rhel8	registry.redhat.io/codeready-workspaces/pluginbroker-metadata-rhel8:2.4	2.4
pluginbroker-artifacts-rhel8	registry.redhat.io/codeready-workspaces/pluginbroker-artifacts-rhel8:2.4	2.4
plugin-java8-rhel8	registry.redhat.io/codeready-workspaces/plugin-java8-rhel8:2.4	2.4
plugin-java11-rhel8	registry.redhat.io/codeready-workspaces/plugin-java11-rhel8:2.4	2.4
plugin-kubernetes-rhel8	registry.redhat.io/codeready-workspaces/plugin-kubernetes-rhel8:2.4	2.4
plugin-openshift-rhel8	registry.redhat.io/codeready-workspaces/plugin-openshift-rhel8:2.4	2.4
stacks-cpp-rhel8	registry.redhat.io/codeready-workspaces/stacks-cpp-rhel8:2.4	2.4
stacks-dotnet-rhel8	registry.redhat.io/codeready-workspaces/stacks-dotnet-rhel8:2.4	2.4
stacks-golang-rhel8	registry.redhat.io/codeready-workspaces/stacks-golang-rhel8:2.4	2.4
stacks-php-rhel8	registry.redhat.io/codeready-workspaces/stacks-php-rhel8:2.4	2.4

See [Table 10.1, “Image Puller default parameters”](#) for more information about configuration values, such as `DAEMONSET_NAME`, `CACHING_INTERVAL_HOURS`, and `CACHING_MEMORY_REQUEST`.

## Procedure

### Installing

1. Clone the Kubernetes Image Puller repository:

```
$ git clone https://github.com/che-incubator/kubernetes-image-puller
$ cd kubernetes-image-puller
```

2. Create a new OpenShift project to deploy the puller into:

```
$ oc new-project k8s-image-puller
```

3. Process and apply the templates to deploy the puller:

In CodeReady Workspaces you must use custom values to deploy the image puller. To set custom values, add to the `oc process` an option: `-p <parameterName>=<value>`:

```
$ oc process -f deploy/openshift/serviceaccount.yaml \
| oc apply -f -
$ oc process -f deploy/openshift/configmap.yaml \
-p IMAGES='plugin-java8-rhel8=registry.redhat.io/codeready-workspaces/plugin-java8-
rhel8:2.4;\
theia-rhel8=registry.redhat.io/codeready-workspaces/theia-rhel8:2.4;\
stacks-golang-rhel8=registry.redhat.io/codeready-workspaces/stacks-golang-rhel8:2.4;\
plugin-java11-rhel8=registry.redhat.io/codeready-workspaces/plugin-java11-rhel8:2.4;\
theia-endpoint-rhel8=registry.redhat.io/codeready-workspaces/theia-rhel8:2.4;\
pluginbroker-metadata-rhel8=registry.redhat.io/codeready-workspaces/pluginbroker-
metadata-rhel8:2.4;\
pluginbroker-artifacts-rhel8=registry.redhat.io/codeready-workspaces/pluginbroker-
artifacts-rhel8:2.4;' \
| oc apply -f -
$ oc process -f deploy/openshift/app.yaml \
-p IMAGE=registry.redhat.io/codeready-workspaces/imagepuller-rhel8 \
-p IMAGE_TAG='2.4' \
| oc apply -f -
```

### Verifying the installation

1. Confirm that a new deployment, `kubernetes-image-puller`, and a DaemonSet (named based on the value of the `DAEMONSET_NAME` parameter) exist. The DaemonSet needs to have a Pod for each node in the cluster:

```
$ oc get deployment,daemonset,pod --namespace k8s-image-puller
deployment.extensions/kubernetes-image-puller 1/1 1 1 2m19s
```

NAME	DESIRED	CURRENT	READY	UP-TO-DATE
daemonset.extensions/kubernetes-image-puller	1	1	1	1
<none>	2m10s			

NAME	READY	STATUS	RESTARTS	AGE
pod/kubernetes-image-puller-5495f46497-mkd4p	1/1	Running	0	2m18s
pod/kubernetes-image-puller-n8bmf	3/3	Running	0	2m10s

2. Check that the **ConfigMap** named **k8s-image-puller** has the values you specified in your parameter substitution, or that they contain the default values:

```
$ oc get configmap k8s-image-puller --output yaml
apiVersion: v1
data:
  CACHING_INTERVAL_HOURS: "1"
  CACHING_MEMORY_LIMIT: 20Mi
  CACHING_MEMORY_REQUEST: 10Mi
  DAEMONSET_NAME: kubernetes-image-puller
  IMAGES: |
    theia-rhel8=registry.redhat.io/codeready-workspaces/theia-rhel8:{prod-ver};
    theia-endpoint-rhel8=registry.redhat.io/codeready-workspaces/theia-rhel8:{prod-ver};
    pluginbroker-metadata-rhel8=registry.redhat.io/codeready-workspaces/pluginbroker-
metadata-rhel8:{prod-ver};
    pluginbroker-artifacts-rhel8=registry.redhat.io/codeready-workspaces/pluginbroker-
artifacts-rhel8:{prod-ver};
    plugin-java8-rhel8=registry.redhat.io/codeready-workspaces/plugin-java8-rhel8:{prod-ver};
    plugin-java11-rhel8=registry.redhat.io/codeready-workspaces/plugin-java11-rhel8:{prod-
ver};
    stacks-golang-rhel8=registry.redhat.io/codeready-workspaces/stacks-golang-rhel8:{prod-
ver};
  NAMESPACE: k8s-image-puller
  NODE_SELECTOR: '{}'
kind: ConfigMap
metadata:
  annotations:
    kubectl.kubernetes.io/last-applied-configuration: |
      {"apiVersion":"v1","data":
{"CACHING_INTERVAL_HOURS":"1","CACHING_MEMORY_LIMIT":"20Mi","CACHING_ME
MEMORY_REQUEST":"10Mi","DAEMONSET_NAME":"kubernetes-image-
puller","IMAGES":"theia-rhel8=registry.redhat.io/codeready-workspaces/theia-rhel8:{prod-
ver}; theia-endpoint-rhel8=registry.redhat.io/codeready-workspaces/theia-rhel8:{prod-ver};
pluginbroker-metadata-rhel8=registry.redhat.io/codeready-workspaces/pluginbroker-
metadata-rhel8:{prod-ver}; pluginbroker-artifacts-rhel8=registry.redhat.io/codeready-
workspaces/pluginbroker-artifacts-rhel8:{prod-ver}; plugin-java8-
rhel8=registry.redhat.io/codeready-workspaces/plugin-java8-rhel8:{prod-ver}; plugin-java11-
rhel8=registry.redhat.io/codeready-workspaces/plugin-java11-rhel8:{prod-ver}; stacks-
golang-rhel8=registry.redhat.io/codeready-workspaces/stacks-golang-rhel8:{prod-
ver};\n","NAMESPACE":"k8s-image-puller","NODE_SELECTOR":"
{}"},"kind":"ConfigMap","metadata":{"annotations":{},"name":"k8s-image-
puller","namespace":"k8s-image-puller"},"type":"Opaque"}
      creationTimestamp: 2020-02-17T22:40:13Z
      name: k8s-image-puller
      namespace: k8s-image-puller
      resourceVersion: "72250"
      selfLink: /api/v1/namespaces/k8s-image-puller/configmaps/k8s-image-puller
      uid: 76430ed6-51d6-11ea-9c19-52fdcf072182
```

