



## Red Hat Quay 3.5

### Use Red Hat Quay

Use Red Hat Quay



## Red Hat Quay 3.5 Use Red Hat Quay

---

Use Red Hat Quay

## Legal Notice

Copyright © 2021 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

Learn to use Red Hat Quay

# Table of Contents

<b>PREFACE</b> .....	<b>4</b>
<b>CHAPTER 1. USERS AND ORGANIZATIONS IN RED HAT QUAY</b> .....	<b>5</b>
1.1. RED HAT QUAY TENANCY MODEL	5
1.2. CREATING USER ACCOUNTS	5
1.3. CREATING ORGANIZATION ACCOUNTS	6
<b>CHAPTER 2. CREATING A REPOSITORY</b> .....	<b>8</b>
2.1. CREATING AN IMAGE REPOSITORY VIA THE UI	8
2.2. CREATING AN IMAGE REPOSITORY VIA DOCKER OR PODMAN	9
<b>CHAPTER 3. MANAGING ACCESS TO REPOSITORIES</b> .....	<b>10</b>
3.1. ALLOWING ACCESS TO USER REPOSITORIES	10
3.1.1. Allowing user access to a user repository	10
3.2. ALLOWING ROBOT ACCESS TO A USER REPOSITORY	11
3.3. ALLOWING ACCESS TO ORGANIZATION REPOSITORIES	12
3.3.1. Adding a Team to an organization	12
3.3.2. Setting a Team role	13
3.3.3. Adding users to a Team	14
<b>CHAPTER 4. WORKING WITH TAGS</b> .....	<b>15</b>
4.1. VIEWING AND MODIFYING TAGS	15
4.1.1. Adding a new tag to a tagged image	15
4.1.2. Moving a tag	15
4.1.3. Deleting a tag	15
4.1.4. Viewing tag history and going back in time	15
4.1.4.1. Viewing tag history	15
4.1.4.2. Going back in time	16
4.1.5. Fetching an image by tag or digest	16
4.2. TAG EXPIRATION	16
4.2.1. Setting tag expiration from a Dockerfile	17
4.2.2. Setting tag expiration from the repository	17
4.3. SECURITY SCANNING	18
<b>CHAPTER 5. VIEWING AND EXPORTING LOGS</b> .....	<b>19</b>
5.1. VIEWING LOGS	19
5.2. EXPORTING REPOSITORY LOGS	20
<b>CHAPTER 6. AUTOMATICALLY BUILDING DOCKERFILES WITH BUILD WORKERS</b> .....	<b>22</b>
6.1. ARCHITECTURE OVERVIEW	22
6.1.1. Build manager	22
6.1.2. Build workers' control plane	22
6.1.3. Orchestrator	22
6.2. OPENSIFT REQUIREMENTS	22
6.3. ORCHESTRATOR REQUIREMENTS	23
6.4. SETTING UP RED HAT QUAY BUILDERS WITH OPENSIFT	23
6.4.1. Prepare OpenShift for Red Hat Quay Builds	23
6.4.2. Enable Builders and add Build Configuration to Red Hat Quay's Configuration Bundle	25
6.5. OPENSIFT ROUTES LIMITATION	27
6.6. TROUBLESHOOTING BUILDS	28
6.6.1. DEBUG config flag	28
6.7. SETTING UP GITHUB BUILDS (OPTIONAL)	29

<b>CHAPTER 7. BUILDING DOCKERFILES</b> .....	<b>30</b>
7.1. VIEWING AND MANAGING BUILDS	30
7.2. MANUALLY STARTING A BUILD	30
7.3. BUILD TRIGGERS	30
7.3.1. Creating a new build trigger	30
7.3.2. Manually triggering a build trigger	30
7.3.3. Build Contexts	30
<b>CHAPTER 8. SETTING UP A CUSTOM GIT TRIGGER</b> .....	<b>32</b>
8.1. CREATING A TRIGGER	32
8.2. POST TRIGGER-CREATION SETUP	32
8.2.1. SSH public key access	32
8.2.2. Webhook	33
<b>CHAPTER 9. SKIPPING A SOURCE CONTROL-TRIGGERED BUILD</b> .....	<b>34</b>
<b>CHAPTER 10. SET UP GITHUB BUILD TRIGGER TAGS</b> .....	<b>35</b>
10.1. UNDERSTANDING TAG NAMING FOR BUILD TRIGGERS	35
10.2. SETTING TAG NAMES FOR BUILD TRIGGERS	35
<b>CHAPTER 11. CREATING AN OAUTH APPLICATION IN GITHUB</b> .....	<b>38</b>
11.1. CREATE NEW GITHUB APPLICATION	38
<b>CHAPTER 12. REPOSITORY NOTIFICATIONS</b> .....	<b>39</b>
12.1. REPOSITORY EVENTS	39
12.1.1. Repository Push	39
12.1.2. Dockerfile Build Queued	39
12.1.3. Dockerfile Build Started	40
12.1.4. Dockerfile Build Successfully Completed	41
12.1.5. Dockerfile Build Failed	42
12.1.6. Dockerfile Build Cancelled	43
12.1.7. Vulnerability Detected	44
12.2. NOTIFICATION ACTIONS	44
12.2.1. Quay Notification	44
12.2.2. E-mail	44
12.2.3. Webhook POST	44
12.2.4. Flowdock Notification	45
12.2.5. Hipchat Notification	45
12.2.6. Slack Notification	45
<b>CHAPTER 13. USING THE RED HAT QUAY API</b> .....	<b>46</b>
13.1. ACCESSING THE QUAY API FROM QUAY.IO	46
13.2. CREATE OAUTH ACCESS TOKEN	46
13.3. ACCESSING YOUR QUAY API FROM A WEB BROWSER	47
13.4. ACCESSING THE RED HAT QUAY API FROM THE COMMAND LINE	47
13.4.1. Get superuser information	47
13.4.2. Creating a superuser using the API	48
13.4.3. Directory synchronization	49
13.4.4. Create a repository build via API	49
13.4.5. Create an org robot	50
13.4.6. Trigger a build	50
13.4.7. Create a private repository	50



## PREFACE

Red Hat Quay container image registries let you store container images in a central location. As a regular user of a Red Hat Quay registry, you can create repositories to organize your images and selectively add read (pull) and write (push) access to the repositories you control. A user with administrative privileges can perform a broader set of tasks, such as the ability to add users and control default settings.

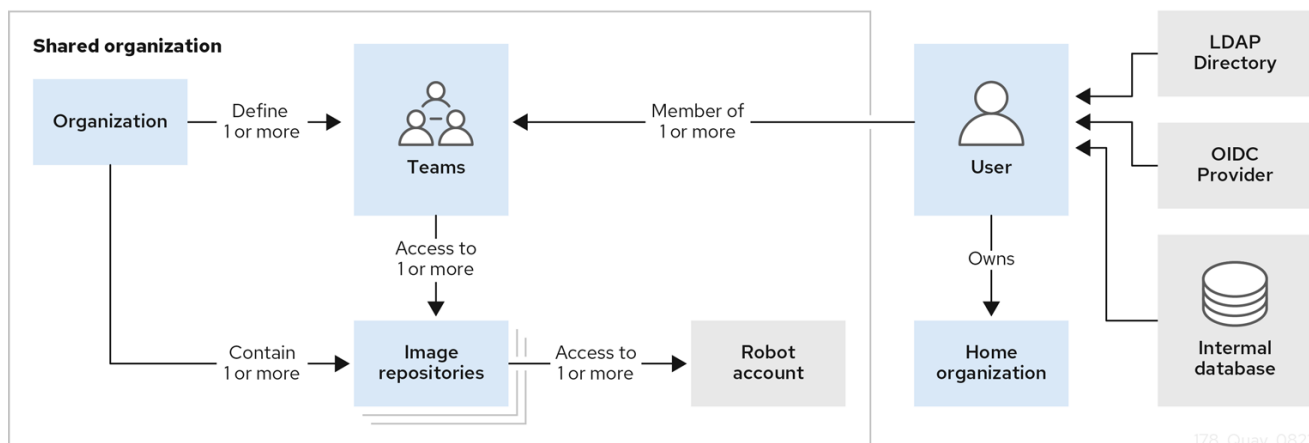
This guide assumes you have a Red Hat Quay deployed and are ready to start setting it up and using it.



## CHAPTER 1. USERS AND ORGANIZATIONS IN RED HAT QUAY

Before you begin creating repositories to hold your container images in Red Hat Quay, you should consider how you want to organize those repositories. Every repository in a Red Hat Quay instance must be associated with either an Organization or a User.

### 1.1. RED HAT QUAY TENANCY MODEL



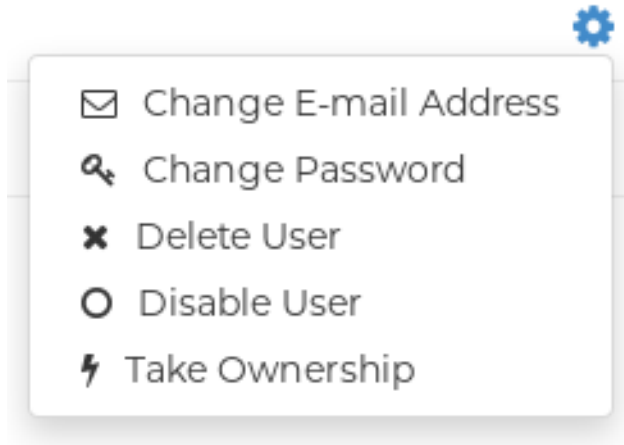
- **Organizations** provide a way of sharing repositories under a common namespace that does not belong to a single user, but rather to many users in a shared setting (such as a company).
- **Teams** provide a way for an organization to delegate permissions (both global and on specific repositories) to sets or groups of users
- **Users** can log in to a registry through the Quay web UI or a client (such as **podman login**). Each users automatically gets a user namespace, for example, **quay-server.example.com/user/<username>**
- **Super users** have enhanced access and privileges via the Super User Admin Panel in the user interface and through Super User API calls that are not visible or accessible to normal users
- **Robot accounts** provide automated access to repositories for non-human users such as pipeline tools and are similar in nature to OpenShift service accounts. Permissions can be granted to a robot account in a repository by adding that account like any other user or team.

### 1.2. CREATING USER ACCOUNTS

To create a new user for your Red Hat Quay instance:

1. Log in to Red Hat Quay as the superuser (quay by default).
2. Select your account name from the upper right corner of the home page and choose Super User Admin Panel.
3. Select the Users icon from the left column.
4. Select the Create User button.
5. Enter the new user's Username and Email address, then select the Create User button.

6. Back on the Users page, select the Options icon to the right of the new Username. A drop-down menu appears, as shown in the following figure:







7. Choose Change Password from the menu.
8. Add the new password and verify it, then select the Change User Password button.

The new user can now use that username and password to log in via the web ui or through some container client.

### 1.3. CREATING ORGANIZATION ACCOUNTS

Any user can create their own organization to share repositories of container images. To create a new organization:

1. While logged in as any user, select the plus sign (+) from the upper right corner of the home page and choose New Organization.
2. Type the name of the organization. The name must be alphanumeric, all lower case, and between 2 and 255 characters long
3. Select Create Organization. The new organization appears, ready for you to begin adding repositories, teams, robot accounts and other features from icons on the left column. The following figure shows an example of the new organization's page with the settings tab selected.

EXPLORE REPOSITORIES TUTORIAL  + -   adminis... **clairv4-org**[+ Create New Repository](#)

## Organization Settings



**Namespace:** clairv4-org  
Organization names cannot be changed once set.



**Avatar:**   
Avatar is generated based off the organization's name.



**Delete organization:** [Begin deletion >](#)



**Time Machine:**   
The amount of time, after a tag is deleted, that the tag is accessible in time machine before being garbage collected.  
[Save Expiration Time](#)

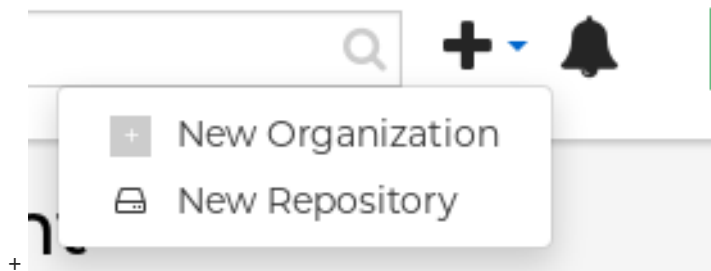


## CHAPTER 2. CREATING A REPOSITORY

A repository provides a central location for storing a related set of container images. There are two ways to create a repository in Red Hat Quay: via a push (from **docker** or **podman**) and via the Red Hat Quay UI. These are essentially the same, whether you are using Quay.io or your own instance of Red Hat Quay.

### 2.1. CREATING AN IMAGE REPOSITORY VIA THE UI

To create a repository in the Red Hat Quay UI under a user account: . Log in to the user account through the web UI. . Click the + icon in the top right of the header on the home page (or other page related to the user) and choose New Repository, as shown in the following figure:



1. On the Create New Repository page that appears
  - Add the new repository name to your user name
  - Click Repository Description and type a description of the repository
  - In Repository Visibility, select whether you want the repository to be public or private
  - Click the Create Repository button.

The new repository is created, starting out empty. A docker pull command you could use to pull an image from this repository (minus the image name) appears on the screen.

To create a repository in the Red Hat Quay UI under an organization:

1. Log in as a user that has Admin or Write permission to the organization.
2. From the Repositories view, select the organization name from the right column under Users and Organizations. The page for the organization appears, similar to the page shown in Figure 2.x:
3. Click +Create New Repository in the upper-right part of the page.
4. On the Create New Repository page that appears:
  - Add the new repository name to the organization name
  - Click Repository Description and type a description of the repository
  - In Repository Visibility, select whether you want the repository to be public or private
  - Click the Create Repository button.

The new repository is created, starting out empty. A docker pull command you could use to pull an image from this repository (minus the image name) appears on the screen.

## 2.2. CREATING AN IMAGE REPOSITORY VIA DOCKER OR PODMAN

Assuming you have the proper credentials, pushing an image to a repository that does not yet exist in your Red Hat Quay instance will create that repository as it pushes the image to that repository. Either the **docker** or **podman** commands will work for these examples.

1. Tag the image: With an image available from **docker** or **podman** on your local system, tag that image with the new repository name and image name. Here are examples for pushing images to Quay.io or your own Red Hat Quay setup (for example, reg.example.com). For the examples, replace namespace with your Red Hat Quay user name or organization and repo\_name with the name of the repository you want to create:

```
# sudo podman tag myubi-minimal quay.io/namespace/repo_name  
# sudo podman tag myubi-standard reg.example.com/namespace/repo_name
```

2. Push to the appropriate registry. For example:

```
# sudo podman push quay.io/namespace/repo_name  
# sudo podman push reg.example.com/namespace/repo_name
```



### NOTE

To create an application repository, follow the same procedure you did for creating a container image repository.

## CHAPTER 3. MANAGING ACCESS TO REPOSITORIES

As a Red Hat Quay user, you can create your own repositories and make them accessible to other users on your Red Hat Quay instance. As an alternative, you can create organizations to allow access to repositories based on teams. In both user and organization repositories, you can allow access to those repositories by creating credentials associated with robot accounts. Robot accounts make it easy for a variety of container clients (such as docker or podman) to access your repos, without requiring that the client have a Red Hat Quay user account.

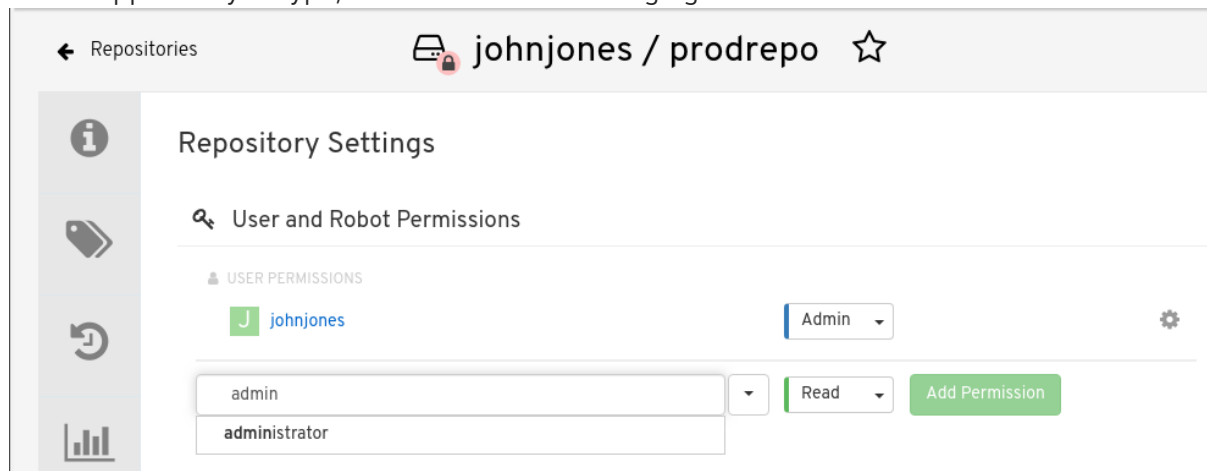
### 3.1. ALLOWING ACCESS TO USER REPOSITORIES

When you create a repository in a user namespace, you can add access to that repository to user accounts or through robot accounts.

#### 3.1.1. Allowing user access to a user repository

To allow access to a repository associated with a user account, do the following:

1. Log into your Red Hat Quay user account.
2. Select a repository under your user namespace to which you want to share access.
3. Select the Settings icon from the left column.
4. Type the name of the user to which you want to grant access to your repository. The user name should appear as you type, as shown in the following figure:



5. In the permissions box, select one of the following:
  - Read - Allows the user to view the repository and pull from it.
  - Write - Allows the user to view the repository, as well as pull images from or push images to the repository.
  - Admin - Allows all administrative settings to the repository, as well as all Read and Write permissions.
6. Select the Add Permission button. The user now has the assigned permission.

To remove the user permissions to the repository, select the Options icon to the right of the user entry, then select Delete Permission.

## 3.2. ALLOWING ROBOT ACCESS TO A USER REPOSITORY

Robot accounts are used to set up automated access to the repositories in your Red Hat Quay registry. They are similar to OpenShift service accounts. When you set up a robot account, you:

- Generate credentials that are associated with the robot account
- Identify repositories and images that the robot can push images to or pull images from
- Copy and paste generated credentials to use with different container clients (such as Docker, podman, Kubernetes, Mesos and others) to access each defined repository

Keep in mind that each robot account is limited to a single user namespace or organization. So, for example, the robot could provide access to all repositories accessible to a user `jsmith`, but not to any that are not in the user's list of repositories.

The following procedure steps you through setting up a robot account to allow access to your repositories.

1. Select Robot icon: From the Repositories view, select the Robot icon from the left column.
2. Create Robot account: Select the Create Robot Account button.
3. Set Robot name: Enter the name and description, then select the Create robot account button. The robot name becomes a combination of your user name, plus the robot name you set (for example, `jsmith+myrobot`)
4. Add permission to the robot account: From the Add permissions screen for the robot account, define the repositories you want the robot to access as follows:
  - Put a check mark next to each repository the robot can access
  - For each repository, select one of the following, and click Add permissions:
    - None - Robot has no permission to the repository
    - Read - Robot can view and pull from the repository
    - Write - Robot can read (pull) from and write (push) to the repository
    - Admin - Full access to pull from and push to the repository, plus the ability to do administrative tasks associated with the repository
  - Select the Add permissions button to apply the settings
5. Get credentials to access repositories via the robot: Back on the Robot Accounts page, select the Robot account name to see credential information for that robot.
6. Get the token: Select Robot Token, as shown in the following figure, to see the token that was generated for the robot. If you want to reset the token, select Regenerate Token. It is important to understand that regenerating a token makes any previous tokens for this robot invalid.

7. Get credentials: Once you are satisfied with the generated token, get the resulting credentials in the following ways:
- **Kubernetes Secret:** Select this to download credentials in the form of a Kubernetes pull secret yaml file.
  - **rkt Configuration:** Select this to download credentials for the rkt container runtime in the form of a json file.
  - **Docker Login:** Select this to copy a full **docker login** command line that includes the credentials.
  - **Docker Configuration:** Select this to download a file to use as a Docker config.json file, to permanently store the credentials on your client system.
  - **Mesos Credentials:** Select this to download a tarball that provides the credentials that can be identified in the uris field of a Mesos configuration file.

### 3.3. ALLOWING ACCESS TO ORGANIZATION REPOSITORIES

Once you have created an organization, you can associate a set of repositories directly to that organization. To add access to the repositories in that organization, you can add Teams (sets of users with the same permissions) and individual users. Essentially, an organization has the same ability to create repositories and robot accounts as a user does, but an organization is intended to set up shared repositories through groups of users (in teams or individually).

Other things to know about organizations:

- You cannot have an organization in another organization. To subdivide an organization, you use teams.
- Organizations can't contain users directly. You must first add a team, then add one or more users to each team.
- Teams can be set up in organizations as just members who use the repos and associated images or as administrators with special privileges for managing the organization

#### 3.3.1. Adding a Team to an organization

When you create a team for your organization you can select the team name, choose which repositories to make available to the team, and decide the level of access to the team.



1. From the Organization view, select the Teams and Membership icon from the left column. You will see that an owners Team exists with Admin privilege for the user who created the Organization.
2. Select Create New Team. You are prompted for the new team name to be associated with the organization. Type the team name, which must start with a lowercase letter, with the rest of the team name as any combination of lowercase letters and numbers (no capitals or special characters allowed).
3. Select the Create team button. The Add permissions window appears, displaying a list of repositories in the organization.
4. Check each repository you want the team to be able to access. Then select one of the following permissions for each:
  - Read - Team members are able to view and pull images
  - Write - Team members can view, pull, and push images
  - Admin - Team members have full read/write privilege, plus the ability to do administrative tasks related to the repository
5. Select Add permissions to save the repository permissions for the team.

### 3.3.2. Setting a Team role

After you have added a team, you can set the role of that team within the organization. From the Teams and Membership screen within the organization, select the TEAM ROLE drop-down menu, as shown in the following figure:

The screenshot shows the 'alldevelopers' organization's 'Teams and Membership' page. The page has a header with the organization name 'alldevelopers' and a '+ Create New Repository' button. Below the header is a navigation bar with 'Teams View', 'Members View', and 'Collaborators View' tabs. The main content area has a 'Teams and Membership' title and a '+ Create New Team' button. A table lists the teams:

TEAM NAME	MEMBERS	REPOSITORIES	TEAM ROLE
owners	1 member	No repositories	Admin
testers	0 members	No repositories	Member

The 'testers' team's role dropdown menu is open, showing the following options:

- Member: Inherits all permissions of the team
- Creator: Member and can create new repositories
- Admin: Full admin access to the organization

For the selected team, choose one of the following roles:

- Member - Inherits all permissions set for the team
- Creator - All member permissions, plus the ability to create new repositories
- Admin - Full administrative access to the organization, including the ability to create teams, add members, and set permissions.

### 3.3.3. Adding users to a Team

As someone with Admin privilege to an organization, you can add users and robots to a team. When you add a user, it sends an email to that user. The user remains pending until that user accepts the invitation.

To add users or robots to a team, start from the organization's screen and do the following:

1. Select the team you want to add users or robots to.
2. In the Team Members box, type one of the following:
  - A username from an account on the Red Hat Quay registry
  - The email address for a user account on the registry
  - The name of a robot account. The name must be in the form of `orgname+robotname`
3. In the case of the robot account, it is immediately added to the team. For a user account, an invitation to join is mailed to the user. Until the user accepts that invitation, the user remains in the INVITED TO JOIN state.

Next, the user accepts the email invitation to join the team. The next time the user logs in to the Red Hat Quay instance, the user moves from the INVITED TO JOIN list to the MEMBERS list for the organization.

## CHAPTER 4. WORKING WITH TAGS

Tags provide a way to identify the version of an image, as well as offering a means of naming the same image in different ways. Besides an image's version, an image tag can identify its uses (such as devel, testing, or prod) or the fact that it is the most recent version (latest).

From the **Tags** tab of an image repository, you can view, modify, add, move, delete, and see the history of tags. You also can fetch command-lines you can use to download (pull) a specific image (based on its name and tag) using different commands.

### 4.1. VIEWING AND MODIFYING TAGS

The tags of a repository can be viewed and modified in the tags panel of the repository page, found by clicking on the **Tags** tab.

Repository Tags Compact Expanded

Actions 1 - 25 of 287 < > Filter Tags...

TAG	LAST MODIFIED ↓	SECURITY SCAN	SIZE	IMAGE
<input checked="" type="checkbox"/> latest	16 hours ago	70 Medium · 10 fixable	711.0 MB	SHA256 9a347939468e
<input type="checkbox"/> master	16 hours ago	70 Medium · 10 fixable	711.0 MB	SHA256 014514e8ef9b
<input type="checkbox"/> dbb57f7	18 hours ago	70 Medium · 10 fixable	696.1 MB	SHA256 2592c71fe8f5
<input type="checkbox"/> 3e28797	a day ago	75 Medium · 15 fixable	693.5 MB	SHA256 0d37d281173e

#### 4.1.1. Adding a new tag to a tagged image

A new tag can be added to a tagged image by clicking on the gear icon next to the tag and choosing **Add New Tag**. Red Hat Quay will confirm the addition of the new tag to the image.

#### 4.1.2. Moving a tag

Moving a tag to a different image is accomplished by performing the same operation as adding a new tag, but giving an existing tag name. Red Hat Quay will confirm that you want the tag moved, rather than added.

#### 4.1.3. Deleting a tag

A specific tag and all its images can be deleted by clicking on the tag's gear icon and choosing **Delete Tag**. This will delete the tag and any images unique to it. Images will not be deleted until no tag references them either directly or indirectly through a parent child relationship.

#### 4.1.4. Viewing tag history and going back in time

##### 4.1.4.1. Viewing tag history

To view the image history for a tag, click on the **View Tags History** menu item located under the **Actions** menu. The page shown will display each image to which the tag pointed in the past and when it pointed to that image.


#### 4.1.4.2. Going back in time

To revert the tag to a previous image, find the history line where your desired image was overwritten, and click on the Restore link.

#### 4.1.5. Fetching an image by tag or digest

From the **Tags** tab, you can view different ways of pulling images from the clients that are ready to use those images.

1. Select a particular repository/image
2. Select Tags in the left column
3. Select the Fetch Tag icon for a particular image/tag combination
4. When the Fetch Tag pop-up appears, select the Image format box to see a drop-down menu that shows different ways that are available to pull the image. The selections offer full command lines for pulling a specific container image to the local system:



You can select to pull a regular of an image by tag name or by digest name using the **docker** command. . Choose the type of pull you want, then select **Copy Command**. The full command-line is copied into your clipboard. These two commands show a **docker pull** by tag and by digest:

```
docker pull quay.io/cnegus/whatever:latest
docker pull
quay.io/cnegus/whatever@sha256:e02231a6aa8ba7f5da3859a359f99d77e371cb47e643ce78e101958
782581fb9
```

Paste the command into a command-line shell on a system that has the **docker** command and service available, and press Enter. At this point, the container image is ready to run on your local system.

On RHEL and Fedora systems, you can substitute **podman** for **docker** to pull and run the selected image.

## 4.2. TAG EXPIRATION

Images can be set to expire from a Red Hat Quay repository at a chosen date and time using a feature called **tag expiration**. Here are a few things to know about about tag expiration:

- When a tag expires, the tag is deleted from the repository. If it is the last tag for a specific image, the image is set to be deleted.
- Expiration is set on a per-tag basis, not for a repository on the whole.
- When a tag expires or is deleted, it is not immediately removed from the registry. The value of Time Machine (in User settings) defines when the deleted tag is actually removed and garbage collected. By default, that value is 14 days. Up until that time, a tag can be repointed to an expired or deleted image.
- The Red Hat Quay superuser has no special privilege related to deleting expired images from user repositories. There is no central mechanism for the superuser to gather information and act on user repositories. It is up to the owners of each repository to manage expiration and ultimate deletion of their images.

Tag expiration can be set in different ways:

- By setting the **quay.expires-after=** LABEL in the Dockerfile when the image is created. This sets a time to expire from when the image is built.
- By choosing the expiration date from the EXPIRES column for the repository tag and selecting a specific date and time to expire.

The following figure shows the Options entry for changing tag expiration and the EXPIRES field for when the tag expires. Hover over the EXPIRES field to see the expiration date and time that is currently set.

The screenshot shows the Quay Registry interface for the repository 'johnjones / prodrepo'. The 'Repository Tags' section is active, displaying a table with the following columns: TAG, LAST MODIFIED, SECURITY SCAN, SIZE, EXPIRES, and MANIFEST. A single tag 'latest' is listed with a size of 34.5 MB and a security scan status of 'Passed'. The 'EXPIRES' column shows 'Mon, Aug 31, 2020 6:59 PM' and a link 'in 4 months'. A dropdown menu is open over the 'EXPIRES' column, showing options: '+ Add New Tag', 'Edit Labels', 'Delete Tag', and 'Change Expiration'.

#### 4.2.1. Setting tag expiration from a Dockerfile

Adding a label like **quay.expires-after=20h** via the Dockerfile LABEL command will cause a tag to automatically expire after the time indicated. The time values could be something like **1h**, **2d**, **3w** for hours, days, and weeks, respectively, from the time the image is built.

#### 4.2.2. Setting tag expiration from the repository

On the Repository Tag page there is a UI column titled **EXPIRES** that indicates when a tag will expire. Users can set this by clicking on the time that it will expire or by clicking the Settings button (gear icon) on the right and choosing **Change Expiration**.

Choose the date and time when prompted and select **Change Expiration**. The tag will be set to be deleted from the repository when the expiration time is reached.

## 4.3. SECURITY SCANNING

By clicking the on the vulnerability or fixable count next to a tab you can jump into the security scanning information for that tag. There you can find which CVEs your image is susceptible to, and what remediation options you may have available.


Keep in mind that image scanning only lists vulnerabilities found by the Clair image scanner. What each user does about the vulnerabilities that are uncovered is completely up to that user. The Red Hat Quay superuser does not act on those vulnerabilities found.

## CHAPTER 5. VIEWING AND EXPORTING LOGS

Activity logs are gathered for all repositories and namespaces (users and organizations) in Red Hat Quay. There are multiple ways of accessing log files, including:

- Viewing logs through the web UI
- Exporting logs so they can be saved externally.
- Accessing log entries via the API

To access logs, you must have Admin privilege to the selected repository or namespace.



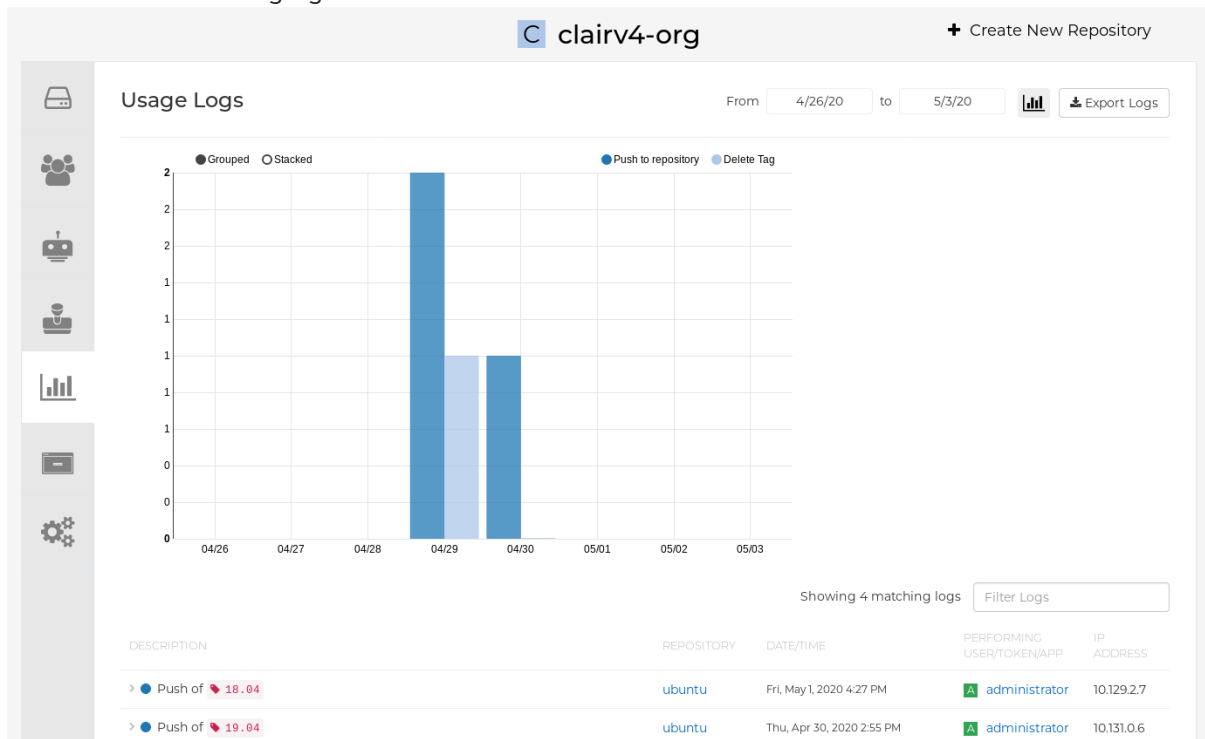
### NOTE

A maximum of 100 log results are available at a time via the API. To gather more results that that, you must use the log exporter feature described in this chapter.

## 5.1. VIEWING LOGS

To view log entries for a repository or namespace from the web UI, do the following:

1. Select a repository or namespace (organization or user) for which you have Admin privileges.
2. Select the Usage Logs icon from the left column. A Usage Logs screen appears, like the one shown in the following figure:



3. From the Usage Logs page, you can:
  - Set the date range for viewing log entries by adding dates to the From and to boxes. By default, the most recent one week of log entries is displayed.
  - Type a string into the Filter Logs box to display log entries that contain the given string.

- Toggle the arrow to the left of any log entry to see more or less text associated with that log entry.

## 5.2. EXPORTING REPOSITORY LOGS

To be able to grab a larger number of log files and save them outside of the Red Hat Quay database, you can use the Export Logs feature. Here are a few things you should know about using Export Logs:

- You can choose a range of dates for the logs you want to gather from a repository.
- You can request that the logs be sent to you via an email attachment or directed to a callback URL.
- You need Admin privilege to the repository or namespace to export logs
- A maximum of 30 days of log data can be exported at a time
- Export Logs only gathers log data that was previously produced. It does not stream logging data.
- Your Red Hat Quay instance must be configured for external storage for this feature (local storage will not work).
- Once the logs are gathered and available, you should immediately copy that data if you want to save it. By default, the data expires in an hour.

To use the Export Logs feature:

1. Select a repository for which you have Admin privileges.
2. Select the Usage Logs icon from the left column. A Usage Logs screen appears.
3. Choose the From and to date range of the log entries you want to gather.
4. Select the Export Logs button. An Export Usage Logs pop-up appears, as shown

### Export Usage Logs ✕

---

Enter an e-mail address or callback URL (must start with `http://` or `https://`) at which to receive the exported logs once they have been fully processed:

johnjones@example.com

Note: The export process can take **up to an hour** to process if there are many logs. As well, only a **single** export process can run at a time for each namespace. Additional export requests will be queued.

Start Logs Export
Cancel

5. Enter the email address or callback URL you want to receive the exported logs. For the callback URL, you could use a URL to a place such as `webhook.site`.
6. Select Start Logs Export. This causes Red Hat Quay to begin gathering the selected log entries. Depending on the amount of logging data being gathered, this can take anywhere from one minute to an hour to complete.
7. When the log export is completed you will either:



- Receive an email, alerting you to the availability of your requested exported log entries.
- See a successful status of your log export request from the webhook URL. A link to the exported data will be available for you to select to download the logs.

Keep in mind that the URL points to a location in your Red Hat Quay external storage and is set to expire within an hour. So make sure you copy the exported logs before that expiration time if you intend to keep them.

## CHAPTER 6. AUTOMATICALLY BUILDING DOCKERFILES WITH BUILD WORKERS

Red Hat Quay supports building Dockerfiles using a set of worker nodes on OpenShift or Kubernetes. Build triggers, such as GitHub webhooks can be configured to automatically build new versions of your repositories when new code is committed. This document will walk you through enabling builds with your Red Hat Quay installation and setting up one or more OpenShift/K8s clusters to accept builds from Red Hat Quay. With Red Hat Quay 3.4, the underlying Build Manager has been completely re-written as part of Red Hat Quay's migration from Python 2 to Python 3. As a result, builder nodes are now dynamically created as Kubernetes Jobs versus builder nodes that ran continuously in Red Hat Quay 3.3 and earlier. This greatly simplifies how Red Hat Quay manages builds and provides the same mechanism quay.io utilizes to handle thousands of container image builds daily. Customers who are currently running static ("Enterprise" builders under Red Hat Quay 3.3) will be required to migrate to a Kubernetes-based build mechanism.

### 6.1. ARCHITECTURE OVERVIEW

The Red Hat Quay Build system is designed for scalability (since it is used to host all builds at quay.io). The Build Manager component of Red Hat Quay provides an orchestration layer that tracks build requests and ensures that a Build Executor (OpenShift/K8s cluster) will carry out each request. Each build is handled by a Kubernetes Job which launches a small virtual machine to completely isolate and contain the image build process. This ensures that container builds do not affect each other or the underlying build system. Multiple Executors can be configured to ensure that builds are performed even in the event of infrastructure failures. Red Hat Quay will automatically send builds to a different Executor if it detects that one Executor is having difficulties.



#### NOTE

The upstream version of Red Hat Quay provides instructions on how to configure an AWS/EC2 based Executor. This configuration is not supported for Red Hat Quay customers.

#### 6.1.1. Build manager

The build manager is responsible for the lifecycle of scheduled build. Operations requiring updating the build queue, build phase and running jobs' status is handled by the build manager.

#### 6.1.2. Build workers' control plane

Build jobs are run on separate worker nodes, and are scheduled on separate control planes (executor). Currently, Red Hat Quay supports running jobs on AWS and Kubernetes. Builds are executed using quay.io/quay/quay-builder. On AWS, builds are scheduled on EC2 instances. On k8s, the builds are scheduled as job resources.

#### 6.1.3. Orchestrator

The orchestrator is used to store the state of currently running build jobs, and publish events for the build manager to consume. e.g expiry events. Currently, the supported orchestrator backend is Redis.

### 6.2. OPENSIFT REQUIREMENTS

Red Hat Quay builds are supported on Kubernetes and OpenShift 4.5 and higher. A bare metal (non-virtualized) worker node is required since build pods require the ability to run kvm virtualization. Each

build is done in an ephemeral virtual machine to ensure complete isolation and security while the build is running. In addition, your OpenShift cluster should permit the ServiceAccount associated with Red Hat Quay builds to run with the necessary SecurityContextConstraint to support privileged containers.

## 6.3. ORCHESTRATOR REQUIREMENTS

The Red Hat Quay builds need access to a Redis instance to track build status information. It is acceptable to use the same Redis instance already deployed with your Red Hat Quay installation. All build queues are managed in the Red Hat Quay database so there is no need for a highly available Redis instance.

## 6.4. SETTING UP RED HAT QUAY BUILDERS WITH OPENSIFT

### 6.4.1. Prepare OpenShift for Red Hat Quay Builds

There are several actions that are needed on an OpenShift cluster before it can accept builds from Red Hat Quay.

1. Create a project where builds will be run (e.g. 'builder')

```
$ oc new-project builder
```

2. Create a **ServiceAccount** in this **Project** that will be used to run builds. Ensure that it has sufficient privileges to create **Jobs** and **Pods**. Copy the **ServiceAccount**'s token for use later.

```
$ oc create sa -n builder quay-builder
$ oc policy add-role-to-user -n builder edit system:serviceaccount:builder:quay-builder
$ oc sa get-token -n builder quay-builder
```

3. Identify the URL for the OpenShift cluster's API server. This can be found from the OpenShift Console.
4. Identify a worker node label to be used when scheduling build **Jobs**. Because build pods need to run on bare metal worker nodes, typically these are identified with specific labels. Check with your cluster administrator to determine exactly which node label should be used.
5. If the cluster is using a self-signed certificate, get the kube apiserver's CA to add to Red Hat Quay's extra certs.

- a. Get the name of the secret containing the CA:

```
$ oc get sa openshift-apiserver-sa --namespace=openshift-apiserver -o json | jq
'.secrets[] | select(.name | contains("openshift-apiserver-sa-token"))'.name
```

- b. Get the **ca.crt** key value from the secret in the Openshift console. The value should begin with "-----BEGIN CERTIFICATE-----"
- c. Import the CA in Red Hat Quay using the ConfigTool. Ensure the name of this file matches **K8S\_API\_TLS\_CA**.

6. Create the necessary security contexts/role bindings for the **ServiceAccount**:

```
apiVersion: security.openshift.io/v1
```

```
kind: SecurityContextConstraints
metadata:
  name: quay-builder
priority: null
readOnlyRootFilesystem: false
requiredDropCapabilities: null
runAsUser:
  type: RunAsAny
seLinuxContext:
  type: RunAsAny
seccompProfiles:
- '*'
supplementalGroups:
  type: RunAsAny
volumes:
- '*'
allowHostDirVolumePlugin: true
allowHostIPC: true
allowHostNetwork: true
allowHostPID: true
allowHostPorts: true
allowPrivilegeEscalation: true
allowPrivilegedContainer: true
allowedCapabilities:
- '*'
allowedUnsafeSysctls:
- '*'
defaultAddCapabilities: null
fsGroup:
  type: RunAsAny
---
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: quay-builder-scc
  namespace: builder
rules:
- apiGroups:
  - security.openshift.io
  resourceNames:
  - quay-builder
  resources:
  - securitycontextconstraints
  verbs:
  - use
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: quay-builder-scc
  namespace: builder
subjects:
- kind: ServiceAccount
  name: quay-builder
roleRef:
```

```
apiGroup: rbac.authorization.k8s.io
kind: Role
name: quay-builder-scc
```

## 6.4.2. Enable Builders and add Build Configuration to Red Hat Quay's Configuration Bundle

1. Ensure that you've got Builds enabled in your Red Hat Quay configuration.

```
FEATURE_BUILD_SUPPORT: True
```

1. Add the following to your Red Hat Quay configuration bundle, replacing each value with a value specific to your installation.



### NOTE

Currently only the Build feature itself can be enabled via the Red Hat Quay Config Tool. The actual configuration of the Build Manager and Executors must be done manually in the config.yaml file.

```
BUILD_MANAGER:
- ephemeral
- ALLOWED_WORKER_COUNT: 1
ORCHESTRATOR_PREFIX: buildman/production/
ORCHESTRATOR:
  REDIS_HOST: quay-redis-host
  REDIS_PASSWORD: quay-redis-password
  REDIS_SSL: true
  REDIS_SKIP_KEYSPACE_EVENT_SETUP: false
EXECUTORS:
- EXECUTOR: kubernetes
  BUILDER_NAMESPACE: builder
  K8S_API_SERVER: api.openshift.somehost.org:6443
  K8S_API_TLS_CA: /conf/stack/extra_ca_cert_build_cluster.crt
  VOLUME_SIZE: 8G
  KUBERNETES_DISTRIBUTION: openshift
  CONTAINER_MEMORY_LIMITS: 5120Mi
  CONTAINER_CPU_LIMITS: 1000m
  CONTAINER_MEMORY_REQUEST: 3968Mi
  CONTAINER_CPU_REQUEST: 500m
  NODE_SELECTOR_LABEL_KEY: beta.kubernetes.io/instance-type
  NODE_SELECTOR_LABEL_VALUE: n1-standard-4
  CONTAINER_RUNTIME: podman
  SERVICE_ACCOUNT_NAME: *****
  SERVICE_ACCOUNT_TOKEN: *****
  QUAY_USERNAME: quay-username
  QUAY_PASSWORD: quay-password
  WORKER_IMAGE: <registry>/quay-quay-builder
  WORKER_TAG: some_tag
  BUILDER_VM_CONTAINER_IMAGE: <registry>/quay-quay-builder-qemu-rhcos:v3.4.0
  SETUP_TIME: 180
  MINIMUM_RETRY_THRESHOLD: 0
```

**SSH\_AUTHORIZED\_KEYS:**

- ssh-rsa 12345 someuser@email.com
- ssh-rsa 67890 someuser2@email.com

Each configuration field is explained below.

**ALLOWED\_WORKER\_COUNT**

Defines how many Build Workers are instantiated per Red Hat Quay Pod. Typically this is '1'.

**ORCHESTRATOR\_PREFIX**

Defines a unique prefix to be added to all Redis keys (useful to isolate Orchestrator values from other Redis keys).

**REDIS\_HOST**

Hostname for your Redis service.

**REDIS\_PASSWORD**

Password to authenticate into your Redis service.

**REDIS\_SSL**

Defines whether or not your Redis connection uses SSL.

**REDIS\_SKIP\_KEYSPACE\_EVENT\_SETUP**

By default, Red Hat Quay does not set up the keyspace events required for key events at runtime. To do so, set REDIS\_SKIP\_KEYSPACE\_EVENT\_SETUP to **false**.

**EXECUTOR**

Starts a definition of an Executor of this type. Valid values are 'kubernetes' and 'ec2'

**BUILDER\_NAMESPACE**

Kubernetes namespace where Red Hat Quay builds will take place

**K8S\_API\_SERVER**

Hostname for API Server of OpenShift cluster where builds will take place

**K8S\_API\_TLS\_CA**

The filepath in the **Quay** container of the build cluster's CA certificate for the Quay app to trust when making API calls.

**KUBERNETES\_DISTRIBUTION**

Indicates which type of Kubernetes is being used. Valid values are 'openshift' and 'k8s'.

**CONTAINER\_\***

Define the resource requests and limits for each build pod.

**NODE\_SELECTOR\_\***

Defines the node selector label name/value pair where build Pods should be scheduled.

**CONTAINER\_RUNTIME**

Specifies whether the builder should run **docker** or **podman**. Customers using Red Hat's **quay-builder** image should set this to **podman**.

**SERVICE\_ACCOUNT\_NAME/SERVICE\_ACCOUNT\_TOKEN**

Defines the Service Account name/token that will be used by build Pods.

**QUAY\_USERNAME/QUAY\_PASSWORD**

Defines the registry credentials needed to pull the Red Hat Quay build worker image that is specified in the WORKER\_IMAGE field. Customers should provide a Red Hat Service Account credential as defined in the section "Creating Registry Service Accounts" against registry.redhat.io in the article at <https://access.redhat.com/RegistryAuthentication>.

**WORKER\_IMAGE**

Image reference for the Red Hat Quay builder image. `registry.redhat.io/quay/quay-builder`

**WORKER\_TAG**

Tag for the builder image desired. The latest version is `v3.4.0`.

**BUILDER\_VM\_CONTAINER\_IMAGE**

The full reference to the container image holding the internal VM needed to run each Red Hat Quay build (**`registry.redhat.io/quay/quay-builder-qemu-rhcos:v3.4.0`**).

**SETUP\_TIME**

Specifies the number of seconds at which a build times out if it has not yet registered itself with the Build Manager (default is 500 seconds). Builds that time out are attempted to be restarted three times. If the build does not register itself after three attempts it is considered failed.

**MINIMUM\_RETRY\_THRESHOLD**

This setting is used with multiple Executors; it indicates how many retries are attempted to start a build before a different Executor is chosen. Setting to 0 means there are no restrictions on how many tries the build job needs to have. This value should be kept intentionally small (three or less) to ensure failovers happen quickly in the event of infrastructure failures. E.g Kubernetes is set as the first executor and EC2 as the second executor. If we want the last attempt to run a job to always be executed on EC2 and not Kubernetes, we would set the Kubernetes executor's

**MINIMUM\_RETRY\_THRESHOLD** to 1 and EC2's **MINIMUM\_RETRY\_THRESHOLD** to 0 (defaults to 0 if not set). In this case, `kubernetes' MINIMUM_RETRY_THRESHOLD > retries_remaining(1)` would evaluate to `False`, thus falling back to the second executor configured

**SSH\_AUTHORIZED\_KEYS**

List of ssh keys to bootstrap in the ignition config. This allows other keys to be used to ssh into the EC2 instance or QEMU VM

## 6.5. OPENSIFT ROUTES LIMITATION

**NOTE**

This section only applies if you are using the Quay Operator on OpenShift with managed **route** component.

Due to a limitation of OpenShift **Routes** to only be able to serve traffic to a single port, additional steps are required to set up builds. Ensure that your **kubectl** or **oc** CLI tool is configured to work with the cluster where the Quay Operator is installed and that your **QuayRegistry** exists (not necessarily the same as the bare metal cluster where your builders run).

- Ensure that HTTP/2 ingress is enabled on the OpenShift cluster by following [these steps](#).
- The Quay Operator will create a **Route** which directs gRPC traffic to the build manager server running inside the existing Quay pod(s). If you want to use a custom hostname (such as a subdomain like **`builder.registry.example.com`**), ensure that you create a CNAME record with your DNS provider which points to the **`status.ingress[0].host`** of the created **Route**:

```
$ kubectl get -n <namespace> route <quayregistry-name>-quay-builder -o jsonpath={.status.ingress[0].host}
```

- Using the OpenShift UI or CLI, update the **Secret** referenced by **`spec.configBundleSecret`** of the **QuayRegistry** with the build cluster CA certificate (name the key **`extra_ca_cert_build_cluster.cert`**), and update the **`config.yaml`** entry with the correct values

referenced in the builder config above (depending on your build executor) along with the **BUILDMAN\_HOSTNAME** field:

```
BUILD_MANAGER:
- ephemeral
- ALLOWED_WORKER_COUNT: 1
ORCHESTRATOR_PREFIX: buildman/production/
ORCHESTRATOR:
  REDIS_HOST: quay-redis-host
  REDIS_PASSWORD: quay-redis-password
  REDIS_SSL: true
  REDIS_SKIP_KEYSPACE_EVENT_SETUP: false
EXECUTORS:
- EXECUTOR: kubernetes
  BUILDER_NAMESPACE: builder
  BUILDMAN_HOSTNAME: <build-manager-hostname>
...
```

The extra configuration field is explained below:

### BUILDMAN\_HOSTNAME

The externally accessible server hostname which the build jobs use to communicate back to the build manager. Default is the same as **SERVER\_HOSTNAME**. For OpenShift **Route**, it is either **status.ingress[0].host** or the CNAME entry if using a custom hostname. **BUILDMAN\_HOSTNAME** needs to include the port number, e.g **somehost:443** for Openshift Route, as the gRPC client used to communicate with the build manager does not infer any port if omitted.

## 6.6. TROUBLESHOOTING BUILDS

The builder instances started by the build manager are ephemeral. This means that they will either get shut down by Red Hat Quay on timeouts/failure or garbage collected by the control plane (EC2/K8s). This means that in order to get the builder logs, one needs to do so **while** the builds are running.

### 6.6.1. DEBUG config flag

A DEBUG flag can be set in order to prevent the builder instances from getting cleaned up after completion/failure. To do so, in the desired executor configuration, set DEBUG to true. For example:

```
EXECUTORS:
- EXECUTOR: ec2
  DEBUG: true
...
- EXECUTOR: kubernetes
  DEBUG: true
...
```

When set to true, DEBUG will prevent the build nodes from shutting down after the quay-builder service is done or fails, and will prevent the build manager from cleaning up the instances (terminating EC2 instances or deleting k8s jobs). This will allow debugging builder node issues, and **should not** be set in a production environment. The lifetime service will still exist. i.e The instance will still shutdown after approximately 2 hours (EC2 instances will terminate, k8s jobs will complete) Setting DEBUG will also affect ALLOWED\_WORKER\_COUNT, as the unterminated instances/jobs will still count towards the total number of running workers. This means the existing builder workers will need to manually be deleted if ALLOWED\_WORKER\_COUNT is reached to be able to schedule new builds.



Use the following steps:

1. The guest VM forwards its SSH port (22) to its host's (the pod) port 2222. Port forward the builder pod's port 2222 to a port on localhost. e.g

```
$ kubectl port-forward <builder pod> 9999:2222
```

2. SSH into the VM running inside the container using a key set from SSH\_AUTHORIZED\_KEYS:

```
$ ssh -i /path/to/ssh/key/set/in/ssh_authorized_keys -p 9999 core@localhost
```

3. Get the quay-builder service logs:

```
$ systemctl status quay-builder  
$ journalctl -f -u quay-builder
```

- Step 2-3 can also be done in a single SSH command:

```
$ ssh -i /path/to/ssh/key/set/in/ssh_authorized_keys -p 9999 core@localhost 'systemctl  
status quay-builder'  
$ ssh -i /path/to/ssh/key/set/in/ssh_authorized_keys -p 9999 core@localhost 'journalctl -f  
-u quay-builder'
```

## 6.7. SETTING UP GITHUB BUILDS (OPTIONAL)

If your organization plans to have builds be conducted via pushes to GitHub (or GitHub Enterprise), continue with *Creating an OAuth application in GitHub*.

## CHAPTER 7. BUILDING DOCKERFILES

Red Hat Quay supports the ability to build [Dockerfiles](#) on our build fleet and push the resulting image to the repository.

### 7.1. VIEWING AND MANAGING BUILDS

Repository Builds can be viewed and managed by clicking the Builds tab in the **Repository View**.

### 7.2. MANUALLY STARTING A BUILD

To manually start a repository build, click the **+** icon in the top right of the header on any repository page and choose **New Dockerfile Build**. An uploaded **Dockerfile**, **.tar.gz**, or an HTTP URL to either can be used for the build.



#### NOTE

You will not be able to specify the Docker build context when manually starting a build.

### 7.3. BUILD TRIGGERS

Repository builds can also be automatically triggered by events such as a push to an SCM (GitHub, BitBucket or GitLab) or via [a call to a webhook](#).

#### 7.3.1. Creating a new build trigger

To setup a build trigger, click the **Create Build Trigger** button on the Builds view page and follow the instructions of the dialog. You will need to grant Red Hat Quay access to your repositories in order to setup the trigger and your account *requires admin access on the SCM repository*.

#### 7.3.2. Manually triggering a build trigger

To trigger a build trigger manually, click the icon next to the build trigger and choose **Run Now**.


#### 7.3.3. Build Contexts

When building an image with Docker, a directory is specified to become the build context. This holds true for both manual builds and build triggers because the builds conducted by Red Hat Quay are no different from running **docker build** on your own machine.

Red Hat Quay build contexts are always the specified *subdirectory* from the build setup and fallback to the root of the build source if none is specified. When a build is triggered, Red Hat Quay build workers clone the git repository to the worker machine and enter the build context before conducting a build.

For builds based on tar archives, build workers extract the archive and enter the build context. For example:

```
example
├── .git
└── Dockerfile
```



Imagine the example above is the directory structure for a GitHub repository called "example". If no subdirectory is specified in the build trigger setup or while manually starting a build, the build will operate in the example directory.

If **subdir** is specified to be the subdirectory in the build trigger setup, only the Dockerfile within it is visible to the build. This means that you cannot use the **ADD** command in the Dockerfile to add **file**, because it is outside of the build context.

Unlike the Docker Hub, the Dockerfile is part of the build context on Red Hat Quay. Thus, it must not appear in the **.dockerignore** file.

## CHAPTER 8. SETTING UP A CUSTOM GIT TRIGGER

A Custom Git Trigger is a generic way for any git server to act as a build trigger. It relies solely on SSH keys and webhook endpoints; everything else is left to the user to implement.

### 8.1. CREATING A TRIGGER

Creating a Custom Git Trigger is similar to the creation of any other trigger with a few subtle differences:

- It is not possible for Red Hat Quay to automatically detect the proper robot account to use with the trigger. This must be done manually in the creation process.
- There are extra steps after the creation of the trigger that must be done in order to use the trigger. These steps are detailed below.

### 8.2. POST TRIGGER-CREATION SETUP

Once a trigger has been created, **there are 2 additional steps required** before the trigger can be used:

- Provide read access to the *SSH public key* generated when creating the trigger.
- Setup a *webhook* that POSTs to the Red Hat Quay endpoint to trigger a build.

The key and the URL are both available at all times by selecting **View Credentials** from the gear located in the trigger listing.

#### Trigger Credentials ×

In order to use this trigger, the following first requires action:

- You must give the following public key read access to the git repository.
- You must set your repository to POST to the following URL to trigger a build.

For more information, refer to the [Custom Git Triggers documentation](#).

SSH Public Key:

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDv2pbbxUd8ii1wCExfL3LMUEwze8xm3CV9 
```

Webhook Endpoint URL:

```
http://%24token:NJKMIE8A2597KBPV2W2TJ2R6VNX3X2E3ZK5I3T6JEKRHKSSA5VKD64EP0 
```

Done

#### 8.2.1. SSH public key access

Depending on the Git server setup, there are various ways to install the SSH public key that Red Hat Quay generates for a custom git trigger. For example, [Git documentation](#) describes a small server setup in which simply adding the key to **\$HOME/.ssh/authorize\_keys** would provide access for builders to clone the repository. For any git repository management software that isn't officially supported, there is usually a location to input the key often labeled as **Deploy Keys**.

## 8.2.2. Webhook

In order to automatically trigger a build, one must POST a JSON payload to the webhook URL with the following format:

```
{
  "commit": "1c002dd",                // required
  "ref": "refs/heads/master",        // required
  "default_branch": "master",        // required
  "commit_info": {                  // optional
    "url": "gitsoftware.com/repository/commits/1234567", // required
    "message": "initial commit",     // required
    "date": "timestamp",             // required
    "author": {                      // optional
      "username": "user",            // required
      "avatar_url": "gravatar.com/user.png", // required
      "url": "gitsoftware.com/users/user" // required
    },
    "committer": {                  // optional
      "username": "user",            // required
      "avatar_url": "gravatar.com/user.png", // required
      "url": "gitsoftware.com/users/user" // required
    }
  }
}
```



### NOTE

This request requires a **Content-Type** header containing **application/json** in order to be valid.

Once again, this can be accomplished in various ways depending on the server setup, but for most cases can be done via a [post-receive git hook](#).

## CHAPTER 9. SKIPPING A SOURCE CONTROL-TRIGGERED BUILD

To specify that a commit should be ignored by the Red Hat Quay build system, add the text **[skip build]** or **[build skip]** anywhere in the commit message.

## CHAPTER 10. SET UP GITHUB BUILD TRIGGER TAGS

Red Hat Quay supports using GitHub or GitHub Enterprise as a trigger to building images. If you have not yet done so, go ahead and [enable build support in Red Hat Quay](#) .

### 10.1. UNDERSTANDING TAG NAMING FOR BUILD TRIGGERS

Prior to Red Hat Quay 3.3, how images created from build triggers were named was limited. Images built by build triggers were named:

- With the branch or tag whose change invoked the trigger
- With a **latest** tag for images that used the default branch

As of Red Hat Quay 3.3 and later, you have more flexibility in how you set image tags. The first thing you can do is enter custom tags, to have any string of characters assigned as a tag for each built image. However, as an alternative, you could use the following tag templates to tag images with information from each commit:

- `${commit_info.short_sha}`: The commit's short SHA
- `${commit_info.date}`: The timestamp for the commit
- `${commit_info.author}`: The author from the commit
- `${commit_info.committer}`: The committer of the commit
- `${parsed_ref.branch}`: The branch name

The following procedure describes how you set up tagging for build triggers.

### 10.2. SETTING TAG NAMES FOR BUILD TRIGGERS

Follow these steps to configure custom tags for build triggers:

1. From the repository view, select the Builds icon from the left navigation.
2. Select the Create Build Trigger menu, and select the type of repository push you want (GitHub, Bitbucket, GitLab, or Custom Git repository push). For this example, *GitHub Repository Push* is chosen, as illustrated in the following figure.



- When the *Setup Build Trigger* page appears, select the repository and namespace in which you want the trigger set up.
- Under *Configure Trigger*, select either *Trigger for all branches and tags* or *Trigger only on branches and tags matching a regular expression*. Then select *Continue*. The *Configure Tagging* section appears, as shown in the following figure:

### Configure Tagging

#### Confirm basic tagging options

- Tag manifest with the branch or tag name**

Tags the built manifest the name of the branch or tag for the git commit.

- Add `latest` tag if on default branch**

Tags the built manifest with `latest` if the build occurred on the default branch for the repository.

#### Add custom tagging templates

- foobar x

Enter a tag template:

By default, all built manifests will be tagged with the name of the branch or tag in which the commit occurred.

To modify this default, as well as the default to add the `latest` tag, change the corresponding options on the left.

Need more control over how the built manifest is tagged? Add one or more custom tag templates.

For example, if you want all built manifests to be tagged with the commit's short SHA, add a template of `${commit_info.short_sha}`.

As another example, if you want on those manifests committed to a branch to be tagged with the branch name, you can add a template of `${parsed_ref.branch}`.

A full reference of for these templates can be found in the [Tag template documentation](#).

- Scroll down to *Configure Tagging* and select from the following options:
  - **Tag manifest with the branch or tag name** Check this box to use the name of the branch or tag in which the commit occurred as the tag used on the image. This is enabled by default.
  - **Add latest tag if on default branch** Check this box to use the `latest` tag for the image if it is on the default branch for the repository. This is enabled by default.
  - **Add custom tagging templates** Enter a custom tag or a template into the *Enter a tag template* box. There are multiple tag templates you can enter here, as described earlier in this section. They include ways of using short SHA, timestamps, author name, committer, and branch name from the commit as tags.
- Select *Continue*. You are prompted to select the directory build context for the Docker build. The build context directory identifies the location of the directory containing the Dockerfile, along with other files needed when the build is triggered. Enter `/` if the Dockerfile is in the root



of the git repository.

7. Select Continue. You are prompted to add an optional Robot Account. Do this if you want to pull a private base image during the build process. The robot account would need access to the build.
8. Select Continue to complete the setup of the build trigger.

If you were to return to the Repository Builds page for the repository, the build triggers you set up will be listed under the Build Triggers heading.

Build Triggers					Create Build Trigger +
TRIGGER NAME	DOCKERFILE LOCATION	CONTEXT LOCATION	BRANCHES/TAGS	PULL ROBOT	TAGGING OPTIONS
 Push to GitHub repository <a href="#">dongboyan77:ruby-hello-world</a>	/Dockerfile	/	All	(None)	Branch/tag name foo
 Push to GitHub repository <a href="#">dongboyan77:ruby-hello-world</a>	/Dockerfile	/	All	(None)	Branch/tag name latest: if default branch latest
 Push to repository <a href="#">git@github.com:dongboyan77:ruby-hello-world.git</a>	/Dockerfile	/	All	(None)	Branch/tag name latest: if default branch

# CHAPTER 11. CREATING AN OAUTH APPLICATION IN GITHUB

You can authorize your registry to access a GitHub account and its repositories by registering it as a GitHub OAuth application.

## 11.1. CREATE NEW GITHUB APPLICATION

1. Log into GitHub (Enterprise)
2. Visit the Applications page under your organization's settings.
3. Click [Register New Application](#). The **Register a new OAuth application** configuration screen is displayed:

Applications / **Register a new OAuth application**

**Application name**  
  
 Something users will recognize and trust

**Homepage URL**  
  
 The full URL to your application homepage


**Application description**  
  
 This is displayed to all potential users of your application

**Authorization callback URL**  
  
 Your application's callback URL. Read our [OAuth documentation](#) for more information

[or choose an image](#)

**Register application**

4. Set Homepage URL: Enter the Quay Enterprise URL as the **Homepage URL**



### NOTE

If using public GitHub, the Homepage URL entered must be accessible by your users. It can still be an internal URL.

5. Set Authorization callback URL: Enter <https://my.registry.url/oauth2/github/callback> as the Authorization callback URL.
6. Save your settings by clicking the Register application button. The new application's summary is shown:
7. Record the Client ID and Client Secret shown for the new application.

## CHAPTER 12. REPOSITORY NOTIFICATIONS

Quay supports adding *notifications* to a repository for various events that occur in the repository's lifecycle. To add notifications, click the **Settings** tab while viewing a repository and select **Create Notification**. From the **When this event occurs** field, select the items for which you want to receive notifications:

The screenshot shows the 'Create repository notification' page. At the top, there's a breadcrumb 'cnegus / whatever' and a bell icon. The main heading is 'Create repository notification'. Below this, there's a section 'When this event occurs' with a dropdown menu that is open, showing the following options: 'Push to Repository', 'Dockerfile Build Queued', 'Dockerfile Build Started', 'Dockerfile Build Successfully Completed', 'Dockerfile Build Failed', 'Docker Build Cancelled', and 'Package Vulnerability Found'. To the right of the dropdown, there's a text box explaining that Quay Container Registry supports various events around repositories, such as push completed, building (build queued, build completed, etc), and security (vulnerability detected). Some events allow for filtering. At the bottom left, there's a 'Create Notification' button.

After selecting an event, further configure it by adding how you will be notified of that event.



### NOTE

Adding notifications requires *repository admin permission*.

The following are examples of repository events.

## 12.1. REPOSITORY EVENTS

### 12.1.1. Repository Push

A successful push of one or more images was made to the repository:

```
{
  "name": "repository",
  "repository": "dgangaia/test",
  "namespace": "dgangaia",
  "docker_url": "quay.io/dgangaia/test",
  "homepage": "https://quay.io/repository/dgangaia/repository",
  "updated_tags": [
    "latest"
  ]
}
```

### 12.1.2. Dockerfile Build Queued

Here is a sample response for a Dockerfile build has been queued into the build system. The response can differ based on the use of optional attributes.

```
{
  "build_id": "296ec063-5f86-4706-a469-f0a400bf9df2",
```

```

"trigger_kind": "github", //Optional
"name": "test",
"repository": "dgangaia/test",
"namespace": "dgangaia",
"docker_url": "quay.io/dgangaia/test",
"trigger_id": "38b6e180-9521-4ff7-9844-acf371340b9e", //Optional
"docker_tags": [
  "master",
  "latest"
],
"repo": "test",
"trigger_metadata": {
  "default_branch": "master",
  "commit": "b7f7d2b948aacbe844ee465122a85a9368b2b735",
  "ref": "refs/heads/master",
  "git_url": "git@github.com:dgangaia/test.git",
  "commit_info": { //Optional
    "url": "https://github.com/dgangaia/test/commit/b7f7d2b948aacbe844ee465122a85a9368b2b735",
    "date": "2019-03-06T12:48:24+11:00",
    "message": "adding 5",
    "author": { //Optional
      "username": "dgangaia",
      "url": "https://github.com/dgangaia", //Optional
      "avatar_url": "https://avatars1.githubusercontent.com/u/43594254?v=4" //Optional
    },
    "committer": {
      "username": "web-flow",
      "url": "https://github.com/web-flow",
      "avatar_url": "https://avatars3.githubusercontent.com/u/19864447?v=4"
    }
  }
},
"is_manual": false,
"manual_user": null,
"homepage": "https://quay.io/repository/dgangaia/test/build/296ec063-5f86-4706-a469-f0a400bf9df2"
}

```

### 12.1.3. Dockerfile Build Started

Here is an example of a Dockerfile build being started by the build system. The response can differ based on some attributes being optional.

```

{
  "build_id": "a8cc247a-a662-4fee-8dcb-7d7e822b71ba",
  "trigger_kind": "github", //Optional
  "name": "test",
  "repository": "dgangaia/test",
  "namespace": "dgangaia",
  "docker_url": "quay.io/dgangaia/test",
  "trigger_id": "38b6e180-9521-4ff7-9844-acf371340b9e", //Optional
  "docker_tags": [
    "master",
    "latest"
  ],
}

```

```

"build_name": "50bc599",
"trigger_metadata": {
  "commit": "50bc5996d4587fd4b2d8edc4af652d4cec293c42",
  "ref": "refs/heads/master",
  "default_branch": "master",
  "git_url": "git@github.com:dgangaia/test.git",
  "commit_info": {
    "url": "https://github.com/dgangaia/test/commit/50bc5996d4587fd4b2d8edc4af652d4cec293c42",
    "date": "2019-03-06T14:10:14+11:00",
    "message": "test build",
    "committer": {
      "username": "web-flow",
      "url": "https://github.com/web-flow",
      "avatar_url": "https://avatars3.githubusercontent.com/u/19864447?v=4"
    },
    "author": {
      "username": "dgangaia",
      "url": "https://github.com/dgangaia",
      "avatar_url": "https://avatars1.githubusercontent.com/u/43594254?v=4"
    }
  },
  "homepage": "https://quay.io/repository/dgangaia/test/build/a8cc247a-a662-4fee-8dcb-7d7e822b71ba"
}

```

#### 12.1.4. Dockerfile Build Successfully Completed

Here is a sample response of a Dockerfile build that has been successfully completed by the build system.



#### NOTE

This event will occur **simultaneously** with a *Repository Push* event for the built image(s)

```

{
  "build_id": "296ec063-5f86-4706-a469-f0a400bf9df2",
  "trigger_kind": "github",
  "name": "test",
  "repository": "dgangaia/test",
  "namespace": "dgangaia",
  "docker_url": "quay.io/dgangaia/test",
  "trigger_id": "38b6e180-9521-4ff7-9844-acf371340b9e",
  "docker_tags": [
    "master",
    "latest"
  ],
  "build_name": "b7f7d2b",
  "image_id": "sha256:0339f178f26ae24930e9ad32751d6839015109eabdf1c25b3b0f2abf8934f6cb",
  "trigger_metadata": {
    "commit": "b7f7d2b948aacbe844ee465122a85a9368b2b735",
    "ref": "refs/heads/master",
    "default_branch": "master",
    "git_url": "git@github.com:dgangaia/test.git",

```

```

"commit_info": {
    "url": "https://github.com/dgangaia/test/commit/b7f7d2b948aacbe844ee465122a85a9368b2b735",
    "date": "2019-03-06T12:48:24+11:00",
    "message": "adding 5",
    "committer": {
        "username": "web-flow",
        "url": "https://github.com/web-flow",
        "avatar_url": "https://avatars3.githubusercontent.com/u/19864447?v=4"
    }
},
"author": {
    "username": "dgangaia",
    "url": "https://github.com/dgangaia",
    "avatar_url": "https://avatars1.githubusercontent.com/u/43594254?v=4"
}
},
"homepage": "https://quay.io/repository/dgangaia/test/build/296ec063-5f86-4706-a469-
f0a400bf9df2",
"manifest_digests": [
    "quay.io/dgangaia/test@sha256:2a7af5265344cc3704d5d47c4604b1efcbd227a7a6a6ff73d6e4e08a27f
d7d99",
    "quay.io/dgangaia/test@sha256:569e7db1a867069835e8e97d50c96eccafde65f08ea3e0d5deba16e25
45d9d1"
]
}

```

### 12.1.5. Dockerfile Build Failed

A Dockerfile build has failed

```

{
    "build_id": "5346a21d-3434-4764-85be-5be1296f293c",
    "trigger_kind": "github",
    "name": "test",
    "repository": "dgangaia/test",
    "docker_url": "quay.io/dgangaia/test",
    "error_message": "Could not find or parse Dockerfile: unknown instruction: GIT",
    "namespace": "dgangaia",
    "trigger_id": "38b6e180-9521-4ff7-9844-acf371340b9e",
    "docker_tags": [
        "master",
        "latest"
    ],
    "build_name": "6ae9a86",
    "trigger_metadata": {
        "commit": "6ae9a86930fc73dd07b02e4c5bf63ee60be180ad",
        "ref": "refs/heads/master",
        "default_branch": "master",
        "git_url": "git@github.com:dgangaia/test.git",
        "commit_info": {
            "url": "https://github.com/dgangaia/test/commit/6ae9a86930fc73dd07b02e4c5bf63ee60be180ad",
            "date": "2019-03-06T14:18:16+11:00",

```

```

"message": "failed build test",
"committer": {                                     //Optional
  "username": "web-flow",
  "url": "https://github.com/web-flow",           //Optional
  "avatar_url": "https://avatars3.githubusercontent.com/u/19864447?v=4" //Optional
},
"author": {                                       //Optional
  "username": "dgangaia",
  "url": "https://github.com/dgangaia",          //Optional
  "avatar_url": "https://avatars1.githubusercontent.com/u/43594254?v=4" //Optional
}
},
"homepage": "https://quay.io/repository/dgangaia/test/build/5346a21d-3434-4764-85be-5be1296f293c"
}

```

### 12.1.6. Dockerfile Build Cancelled

A Dockerfile build was cancelled

```

{
  "build_id": "cbd534c5-f1c0-4816-b4e3-55446b851e70",
  "trigger_kind": "github",
  "name": "test",
  "repository": "dgangaia/test",
  "namespace": "dgangaia",
  "docker_url": "quay.io/dgangaia/test",
  "trigger_id": "38b6e180-9521-4ff7-9844-acf371340b9e",
  "docker_tags": [
    "master",
    "latest"
  ],
  "build_name": "cbce83c",
  "trigger_metadata": {
    "commit": "cbce83c04bfb59734fc42a83aab738704ba7ec41",
    "ref": "refs/heads/master",
    "default_branch": "master",
    "git_url": "git@github.com:dgangaia/test.git",
    "commit_info": {
      "url": "https://github.com/dgangaia/test/commit/cbce83c04bfb59734fc42a83aab738704ba7ec41",
      "date": "2019-03-06T14:27:53+11:00",
      "message": "testing cancel build",
      "committer": {
        "username": "web-flow",
        "url": "https://github.com/web-flow",
        "avatar_url": "https://avatars3.githubusercontent.com/u/19864447?v=4"
      },
      "author": {
        "username": "dgangaia",
        "url": "https://github.com/dgangaia",
        "avatar_url": "https://avatars1.githubusercontent.com/u/43594254?v=4"
      }
    }
  }
},

```

```
"homepage": "https://quay.io/repository/dgangaia/test/build/cbd534c5-f1c0-4816-b4e3-55446b851e70"
}
```

### 12.1.7. Vulnerability Detected

A vulnerability was detected in the repository

```
{
  "repository": "dgangaia/repository",
  "namespace": "dgangaia",
  "name": "repository",
  "docker_url": "quay.io/dgangaia/repository",
  "homepage": "https://quay.io/repository/dgangaia/repository",

  "tags": ["latest", "othertag"],

  "vulnerability": {
    "id": "CVE-1234-5678",
    "description": "This is a bad vulnerability",
    "link": "http://url/to/vuln/info",
    "priority": "Critical",
    "has_fix": true
  }
}
```

## 12.2. NOTIFICATION ACTIONS

### 12.2.1. Quay Notification

A notification will be added to the Quay.io notification area. The notification area can be found by clicking on the bell icon in the top right of any Quay.io page.

Quay.io notifications can be setup to be sent to a *User*, *Team*, or the *organization* as a whole.

### 12.2.2. E-mail

An e-mail will be sent to the specified address describing the event that occurred.



#### NOTE

All e-mail addresses will have to be verified on a *per-repository* basis

### 12.2.3. Webhook POST

An HTTP POST call will be made to the specified URL with the event's data (see above for each event's data format).

When the URL is HTTPS, the call will have an SSL client certificate set from Quay.io. Verification of this certificate will prove the call originated from Quay.io. Responses with status codes in the 2xx range are considered successful. Responses with any other status codes will be considered failures and result in a retry of the webhook notification.



### **12.2.4. Flowdock Notification**

Posts a message to Flowdock.

### **12.2.5. Hipchat Notification**

Posts a message to HipChat.

### **12.2.6. Slack Notification**

Posts a message to Slack.

## CHAPTER 13. USING THE RED HAT QUAY API

Red Hat Quay provides a full [OAuth 2](#), RESTful API that:

- Is available from endpoints of each Red Hat Quay instance from the URL <https://<yourquayhost>/api/v1>
- Lets you connect to endpoints, via a browser, to get, delete, post, and put Red Hat Quay settings by enabling the Swagger UI
- Can be accessed by applications that make API calls and use OAuth tokens
- Sends and receives data as JSON

The following text describes how to access the Red Hat Quay API and use it to view and modify setting in your Red Hat Quay cluster. Appendix A lists and describes API endpoints.

### 13.1. ACCESSING THE QUAY API FROM QUAY.IO

If you don't have your own Red Hat Quay cluster running yet, you can explore the Red Hat Quay API available from Quay.io from your web browser:

<https://docs.quay.io/api/swagger/>

The API Explorer that appears shows Quay.io API endpoints. You will not see superuser API endpoints or endpoints for Red Hat Quay features that are not enabled on Quay.io (such as Repository Mirroring).

From API Explorer, you can get, and sometimes change, information on:

- Billing, subscriptions, and plans
- Repository builds and build triggers
- Error messages and global messages
- Repository images, manifests, permissions, notifications, vulnerabilities, and image signing
- Usage logs
- Organizations, members and OAuth applications
- User and robot accounts
- and more...

Select to open an endpoint to view the Model Schema for each part of the endpoint. Open an endpoint, enter any required parameters (such as a repository name or image), then select the **Try it out!** button to query or change settings associated with a Quay.io endpoint.

### 13.2. CREATE OAUTH ACCESS TOKEN

To create an OAuth access token so you can access the API for your organization:

1. Log in to Red Hat Quay and select your Organization (or create a new one).

2. Select the Applications icon from the left navigation.
3. Select Create New Application and give the new application a name when prompted.
4. Select the new application.
5. Select Generate Token from the left navigation.
6. Select the checkboxes to set the scope of the token and select Generate Access Token.
7. Review the permissions you are allowing and select Authorize Application to approve it.
8. Copy the newly generated token to use to access the API.

### 13.3. ACCESSING YOUR QUAY API FROM A WEB BROWSER

By enabling Swagger, you can access the API for your own Red Hat Quay instance through a web browser. This URL exposes the Red Hat Quay API explorer via the Swagger UI and this URL:

```
https://<yourquayhost>/api/v1/discovery.
```

That way of accessing the API does not include superuser endpoints that are available on Red Hat Quay installations. Here is an example of accessing a Red Hat Quay API interface running on the local system by running the `swagger-ui` container image:

```
# export SERVER_HOSTNAME=<yourhostname>
# sudo podman run -p 8888:8080 -e API_URL=https://$SERVER_HOSTNAME:8443/api/v1/discovery
docker.io/swaggerapi/swagger-ui
```

With the `swagger-ui` container running, open your web browser to localhost port 8888 to view API endpoints via the `swagger-ui` container.

To avoid errors in the log such as "API calls must be invoked with an X-Requested-With header if called from a browser," add the following line to the `config.yaml` on all nodes in the cluster and restart Red Hat Quay:

```
BROWSER_API_CALLS_XHR_ONLY: false
```

### 13.4. ACCESSING THE RED HAT QUAY API FROM THE COMMAND LINE

You can use the `curl` command to GET, PUT, POST, or DELETE settings via the API for your Red Hat Quay cluster. Replace `<token>` with the OAuth access token you created earlier to get or change settings in the following examples.

#### 13.4.1. Get superuser information

```
$ curl -X GET -H "Authorization: Bearer <token_here>" \
  "https://<yourquayhost>/api/v1/superuser/users/"
```

For example:

```
$ curl -X GET -H "Authorization: Bearer mFCdgS7SAIoMcnTsHCGx23vcNsTgziAa4CmmHlsg"
http://quay-server:8080/api/v1/superuser/users/ | jq
```

```
{
  "users": [
    {
      "kind": "user",
      "name": "quayadmin",
      "username": "quayadmin",
      "email": "quayadmin@example.com",
      "verified": true,
      "avatar": {
        "name": "quayadmin",
        "hash": "357a20e8c56e69d6f9734d23ef9517e8",
        "color": "#5254a3",
        "kind": "user"
      },
    },
    "super_user": true,
    "enabled": true
  ]
}
```

### 13.4.2. Creating a superuser using the API

- Configure a superuser name, as described in the Deploy Quay book:
  - Use the configuration editor UI or
  - Edit the **config.yaml** file directly, with the option of using the configuration API to validate (and download) the updated configuration bundle
- Create the user account for the superuser name:
  - Obtain an authorization token as detailed above, and use **curl** to create the user:

```
$ curl -H "Content-Type: application/json" -H "Authorization: Bearer
Fava2kV9C92p1eXnMawBZx9vTqVnksvwNm0ckFKZ" -X POST --data '{
  "username": "quaysuper",
  "email": "quaysuper@example.com"
}' http://quay-server:8080/api/v1/superuser/users/ | jq
```

- The returned content includes a generated password for the new user account:

```
{
  "username": "quaysuper",
  "email": "quaysuper@example.com",
  "password": "EH67NB3Y6PTBED8H0HC6UVHGGGA3ODSE",
  "encrypted_password":
  "fn37AZAUQH0PTsU+vIO9IS0QxPW9A/boXL4ovZjIFtUPrBz9i4j9UDOqMjuxQ/0HTfy38go
KEpG8zYXVeQh3IOFzuOjSvKic2Vq7xdtQsU="
}
```

Now, when you request the list of users, it will show **quaysuper** as a superuser:

```
$ curl -X GET -H "Authorization: Bearer mFCdgS7SAIoMcnTsHCGx23vcNsTgziAa4CmmHlsg"
http://quay-server:8080/api/v1/superuser/users/ | jq
```

```

{
  "users": [
    {
      "kind": "user",
      "name": "quayadmin",
      "username": "quayadmin",
      "email": "quayadmin@example.com",
      "verified": true,
      "avatar": {
        "name": "quayadmin",
        "hash": "357a20e8c56e69d6f9734d23ef9517e8",
        "color": "#5254a3",
        "kind": "user"
      },
      "super_user": true,
      "enabled": true
    },
    {
      "kind": "user",
      "name": "quaysuper",
      "username": "quaysuper",
      "email": "quaysuper@example.com",
      "verified": true,
      "avatar": {
        "name": "quaysuper",
        "hash": "c0e0f155afcef68e58a42243b153df08",
        "color": "#969696",
        "kind": "user"
      },
      "super_user": true,
      "enabled": true
    }
  ]
}

```

### 13.4.3. Directory synchronization

To enable directory synchronization for the team **newteam** in organization **testadminorg**, where the corresponding group name in LDAP is **ldapgroup**:

```

$ curl -X POST -H "Authorization: Bearer 9rJYBR3v3pXcj5XqlA2XX6Thkww4gld4TCYLLWDF" \
  -H "Content-type: application/json" \
  -d '{"group_dn": "cn=ldapgroup,ou=Users"}' \
  http://quay1-server:8080/api/v1/organization/testadminorg/team/newteam/syncing

```

To disable synchronization for the same team:

```

$ curl -X DELETE -H "Authorization: Bearer 9rJYBR3v3pXcj5XqlA2XX6Thkww4gld4TCYLLWDF" \
  http://quay1-server:8080/api/v1/organization/testadminorg/team/newteam/syncing

```

### 13.4.4. Create a repository build via API

In order to build a repository from the specified input and tag the build with custom tags, users can use requestRepoBuild endpoint. It takes the following data:

```
{
  "docker_tags": [
    "string"
  ],
  "pull_robot": "string",
  "subdirectory": "string",
  "archive_url": "string"
}
```

The **archive\_url** parameter should point to a **tar** or **zip** archive that includes the Dockerfile and other required files for the build. The **file\_id** parameter was a part of our older build system. It cannot be used anymore. If Dockerfile is in a sub-directory it needs to be specified as well.

The archive should be publicly accessible. OAuth app should have "Administer Organization" scope because only organization admins have access to the robots' account tokens. Otherwise, someone could get robot permissions by simply granting a build access to a robot (without having access themselves), and use it to grab the image contents. In case of errors, check the json block returned and ensure the archive location, pull robot, and other parameters are being passed correctly. Click "Download logs" on the top-right of the individual build's page to check the logs for more verbose messaging.

### 13.4.5. Create an org robot

```
$ curl -X PUT https://quay.io/api/v1/organization/{orgname}/robots/{robot shortname} \
-H 'Authorization: Bearer <token>'
```

### 13.4.6. Trigger a build

```
$ curl -X POST https://quay.io/api/v1/repository/YOURORNAME/YOURREPONAME/build/ \
-H 'Authorization: Bearer <token>'
```

Python with requests

```
import requests
r = requests.post('https://quay.io/api/v1/repository/example/example/image', headers={'content-type':
'application/json', 'Authorization': 'Bearer <redacted>'}, data={<request-body-contents>})
print(r.text)
```

### 13.4.7. Create a private repository

```
$ curl -X POST https://quay.io/api/v1/repository \
-H 'Authorization: Bearer {token}' \
-H 'Content-Type: application/json' \
-d '{"namespace": "yournamespace", "repository": "yourreponame",
"description": "descriptionofyourrepo", "visibility": "private"}' | jq
```