



# Red Hat Quay 3.5

## Manage Red Hat Quay

Manage Red Hat Quay



# Red Hat Quay 3.5 Manage Red Hat Quay

---

Manage Red Hat Quay

## Legal Notice

Copyright © 2022 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

Manage Red Hat Quay

# Table of Contents

<b>PREFACE</b> .....	<b>5</b>
<b>CHAPTER 1. ADVANCED RED HAT QUAY CONFIGURATION</b> .....	<b>6</b>
1.1. USING RED HAT QUAY CONFIG TOOL TO MODIFY RED HAT QUAY	6
1.1.1. Running the Config Tool from the Red Hat Quay Operator	6
1.1.2. Running the Config Tool from the command line	7
1.2. USING THE API TO MODIFY RED HAT QUAY	8
1.3. EDITING THE CONFIG.YAML FILE TO MODIFY RED HAT QUAY	8
1.3.1. Add name and company to Red Hat Quay sign-in	8
1.3.2. Disable TLS Protocols	8
1.3.3. Rate limit API calls	8
1.3.4. Adjust database connection pooling	8
1.3.4.1. Database connection arguments	9
1.3.4.2. Database SSL configuration	9
1.3.4.2.1. PostgreSQL SSL connection arguments	9
1.3.4.2.2. MySQL SSL connection arguments	10
1.3.4.3. HTTP connection counts	10
1.3.4.4. Dynamic process counts	10
1.3.4.5. Environment variables	11
1.3.4.6. Turning off connection pooling	11
<b>CHAPTER 2. USING THE CONFIGURATION API</b> .....	<b>12</b>
2.1. RETRIEVING THE DEFAULT CONFIGURATION	12
2.2. RETRIEVING THE CURRENT CONFIGURATION	12
2.3. VALIDATING CONFIGURATION USING THE API	13
2.4. DETERMINING THE REQUIRED FIELDS	14
<b>CHAPTER 3. GETTING RED HAT QUAY RELEASE NOTIFICATIONS</b> .....	<b>15</b>
<b>CHAPTER 4. USING SSL TO PROTECT CONNECTIONS TO RED HAT QUAY</b> .....	<b>16</b>
4.1. INTRODUCTION TO USING SSL	16
4.2. CREATE A CERTIFICATE AUTHORITY AND SIGN A CERTIFICATE	16
4.2.1. Create a Certificate Authority	16
4.2.2. Sign a certificate	16
4.3. CONFIGURING SSL USING THE COMMAND LINE	17
4.4. CONFIGURING SSL USING THE UI	18
4.5. TESTING SSL CONFIGURATION USING THE COMMAND LINE	18
4.6. TESTING SSL CONFIGURATION USING THE BROWSER	19
4.7. CONFIGURING PODMAN TO TRUST THE CERTIFICATE AUTHORITY	19
4.8. CONFIGURING THE SYSTEM TO TRUST THE CERTIFICATE AUTHORITY	20
<b>CHAPTER 5. ADDING TLS CERTIFICATES TO THE RED HAT QUAY CONTAINER</b> .....	<b>22</b>
5.1. ADD TLS CERTIFICATES TO RED HAT QUAY	22
5.2. ADD CERTS WHEN DEPLOYED ON KUBERNETES	22
<b>CHAPTER 6. CONFIGURING ACTION LOG STORAGE FOR ELASTICSEARCH</b> .....	<b>24</b>
<b>CHAPTER 7. CLAIR SECURITY SCANNING</b> .....	<b>26</b>
7.1. SETTING UP CLAIR ON A RED HAT QUAY OPENSIFT DEPLOYMENT	26
7.1.1. Deploying Via the Quay Operator	26
7.1.2. Manually Deploying Clair	26
7.2. SETTING UP CLAIR ON A NON-OPENSIFT RED HAT QUAY DEPLOYMENT	31
7.3. USING CLAIR	32

---

7.4. CONFIGURING CLAIR FOR DISCONNECTED ENVIRONMENTS	33
7.5. CLAIR UPDATER URLS	34
7.6. ADDITIONAL INFORMATION	34
<b>CHAPTER 8. SCAN POD IMAGES WITH THE CONTAINER SECURITY OPERATOR</b> .....	<b>35</b>
8.1. RUN THE CSO IN OPENSIFT	35
8.2. QUERY IMAGE VULNERABILITIES FROM THE CLI	37
<b>CHAPTER 9. INTEGRATE RED HAT QUAY INTO OPENSIFT WITH THE BRIDGE OPERATOR</b> .....	<b>38</b>
9.1. RUNNING THE QUAY BRIDGE OPERATOR	38
9.1.1. Prerequisites	38
9.1.2. Setting up and configuring OpenShift and Red Hat Quay	38
9.1.2.1. Red Hat Quay setup	39
9.1.2.2. OpenShift Setup	39
<b>CHAPTER 10. REPOSITORY MIRRORING</b> .....	<b>43</b>
10.1. REPOSITORY MIRRORING	43
10.2. REPOSITORY MIRRORING VERSUS GEO-REPLICATION	43
10.3. USING REPOSITORY MIRRORING	44
10.4. MIRRORING CONFIGURATION UI	45
10.5. MIRRORING CONFIGURATION FIELDS	45
10.6. MIRRORING WORKER	46
10.7. CREATING A MIRRORED REPOSITORY	46
10.7.1. Repository mirroring settings	46
10.7.2. Advanced settings	47
10.7.3. Synchronize now	48
10.8. EVENT NOTIFICATIONS FOR MIRRORING	49
10.9. MIRRORING TAG PATTERNS	49
10.9.1. Pattern syntax	49
10.9.2. Example tag patterns	49
10.10. WORKING WITH MIRRORED REPOSITORIES	50
10.11. REPOSITORY MIRRORING RECOMMENDATIONS	52
<b>CHAPTER 11. BACKING UP AND RESTORING RED HAT QUAY ON AN OPENSIFT CONTAINER PLATFORM DEPLOYMENT</b> .....	<b>53</b>
11.1. BACKING UP RED HAT QUAY	53
11.2. RESTORING RED HAT QUAY	56
<b>CHAPTER 12. LDAP AUTHENTICATION SETUP FOR RED HAT QUAY</b> .....	<b>60</b>
12.1. CONSIDERATIONS PRIOR TO ENABLING LDAP	60
12.1.1. Existing Quay deployments	60
12.1.2. Manual User Creation and LDAP authentication	60
12.2. SET UP LDAP CONFIGURATION	60
12.2.1. Full LDAP URI	60
12.2.2. Team Synchronization	61
12.2.3. Base and Relative Distinguished Names	61
12.2.4. Additional User Filters	62
12.2.5. Administrator DN	63
12.2.6. UID and Mail attributes	63
12.2.7. Validation	64
12.3. COMMON ISSUES	64
12.4. CONFIGURE AN LDAP USER AS SUPERUSER	64
<b>CHAPTER 13. PROMETHEUS AND GRAFANA METRICS UNDER RED HAT QUAY</b> .....	<b>66</b>
13.1. EXPOSING THE PROMETHEUS ENDPOINT	66

---

13.1.1. Setting up Prometheus to consume metrics	66
13.1.2. DNS configuration under Kubernetes	66
13.1.3. DNS configuration for a manual cluster	66
<b>CHAPTER 14. GEO-REPLICATION</b> .....	<b>67</b>
14.1. GEO-REPLICATION FEATURES	67
14.2. GEO-REPLICATION REQUIREMENTS AND CONSTRAINTS	67
14.3. GEO-REPLICATION ARCHITECTURE	68
14.4. ENABLE STORAGE REPLICATION	68
14.4.1. Run Red Hat Quay with storage preferences	69
<b>CHAPTER 15. RED HAT QUAY TROUBLESHOOTING</b> .....	<b>70</b>
<b>CHAPTER 16. SCHEMA FOR RED HAT QUAY CONFIGURATION</b> .....	<b>71</b>
ADDITIONAL RESOURCES	88





## PREFACE

Once you have deployed a Red Hat Quay registry, there are many ways you can further configure and manage that deployment. Topics covered here include:

- Advanced Red Hat Quay configuration
- Setting notifications to alert you of a new Red Hat Quay release
- Securing connections with SSL and TLS certificates
- Directing action logs storage to Elasticsearch
- Configuring image security scanning with Clair
- Scan pod images with the Container Security Operator
- Integrate Red Hat Quay into OpenShift with the Quay Bridge Operator
- Mirroring images with repository mirroring
- Sharing Quay images with a BitTorrent service
- Authenticating users with LDAP
- Enabling Quay for Prometheus and Grafana metrics
- Setting up geo-replication
- Troubleshooting Quay

# CHAPTER 1. ADVANCED RED HAT QUAY CONFIGURATION

You can configure your Red Hat Quay after initial deployment using several different interfaces:

- The Red Hat Quay Config Tool: Running the **Quay** container in **config** mode presents a Web-based interface for configuring the Red Hat Quay cluster. This is the recommended method for most configuration of the Red Hat Quay service itself.
- Editing the **config.yaml**: The **config.yaml** file holds most of the configuration information for the Red Hat Quay cluster. Editing that file directly is possible, but it is only recommended for advanced tuning and performance features that are not available through the Config Tool.
- Red Hat Quay API: Some Red Hat Quay configuration can be done through the API.

While configuration for specific features is covered in separate sections, this section describes how to use each of those interfaces and perform some more advanced configuration.

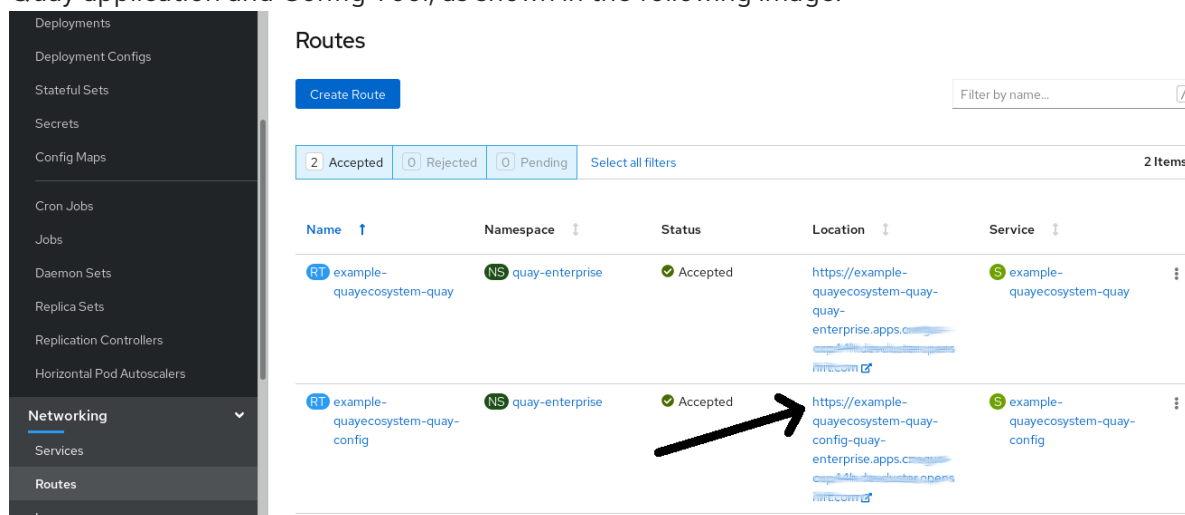
## 1.1. USING RED HAT QUAY CONFIG TOOL TO MODIFY RED HAT QUAY

The Red Hat Quay Config Tool is made available by running a **Quay** container in **config** mode alongside the regular Red Hat Quay service. Running the Config Tool is different for Red Hat Quay clusters running on OpenShift than it is for those running directly on host systems.

### 1.1.1. Running the Config Tool from the Red Hat Quay Operator

If you are running the Red Hat Quay Operator from OpenShift, the Config Tool is probably already available for you to use. To access the Config Tool, do the following:

1. From the OpenShift console, select the project in which Red Hat Quay is running. For example, quay-enterprise.
2. From the left column, select Networking → Routes. You should see routes to both the Red Hat Quay application and Config Tool, as shown in the following image:



3. Select the route to the Config Tool (for example, example-quayecosystem-quay-config) and select it. The Config tool Web UI should open in your browser.
4. Select **Modify configuration for this cluster**. You should see the Config Tool, ready for you to change features of your Red Hat Quay cluster, as shown in the following image:

## Red Hat Quay Setup

**Almost done!**  
Configure your Redis database and other settings below

① — 🗄️ — ② — ③ — ④ — 📄

### Custom SSL Certificates

This section lists any custom or self-signed SSL certificates that are installed in the Project Quay container on startup after being read from the `extra_ca_certs` directory in the configuration volume.

Custom certificates are typically used in place of publicly signed certificates for corporate-internal services.

Please **make sure** that all custom names used for downstream services (such as Clair) are listed in the certificates below.

**Upload certificates:**  Select file

Select custom certificate to add to configuration. Must be in PEM format and end extension '.crt'

CERTIFICATE FILENAME	STATUS	NAMES HANDLED	⚙️
quay.crt	✔️ Certificate is valid	example-quayecosystem-quay-quay-enterprise.apps.cnegus-ocp44h.devcluster.openshift.com quay-enterprise quay-enterprise@1599743640	⚙️
clair.crt	✔️ Certificate is valid	example-quayecosystem-clair example-quayecosystem-clair.quay-enterprise.svc example-quayecosystem-clair.quay-enterprise.svc.local example-quayecosystem-clair@1599743641	⚙️

Save Configuration Changes

5. When you have made the changes you want, select **Save Configuration Changes**. The Config Tool will validate your changes.
6. Make any corrections as needed by selecting **Continue Editing** or select **Next** to continue on.
7. When prompted, it is recommended that you select **Download Configuration**. That will download a tarball of your new **config.yaml**, as well as any certificates and keys used with your Red Hat Quay setup.
8. Select **Go to deployment rollout**, then **Populate the configuration to deployments**. The Red Hat Quay pods will be restarted and the changes will take effect.

The **config.yaml** file you saved can be used to make advanced changes to your configuration or just kept for future reference.

### 1.1.2. Running the Config Tool from the command line

If you are running Red Hat Quay directly from a host system, using tools such as the **podman** or **docker** commands, after the initial Red Hat Quay deployment, you can restart the Config Tool to modify your Red Hat Quay cluster. Here's how:

1. **Start quay in config mode** On the first **quay** node run the following, replacing **my-secret-password** with your password. If you would like to modify an existing config bundle, you can simply mount your configuration directory into the **Quay** container as you would in registry mode.

```
# podman run --rm -it --name quay_config -p 8080:8080 \
-v path/to/config-bundle:/conf/stack \
registry.redhat.io/quay/quay-rhel8:v3.5.7 config my-secret-password
```

2. **Open browser:** When the quay configuration tool starts up, open a browser to the URL and port 8080 of the system you are running the configuration tool on (for example <https://myquay.example.com:8080>). You are prompted for a username and password.

At this point, you can begin modifying your Red Hat Quay cluster as described earlier.

## 1.2. USING THE API TO MODIFY RED HAT QUAY

See the [Red Hat Quay API Guide](#) for information on how to access Red Hat Quay API.

## 1.3. EDITING THE `CONFIG.YAML` FILE TO MODIFY RED HAT QUAY

Some advanced Red Hat Quay configuration that is not available through the Config Tool can be achieved by editing the `config.yaml` file directly. Available settings are described in the [Schema for Red Hat Quay configuration](#). The following are examples of settings you can change directly in the `config.yaml` file.

### 1.3.1. Add name and company to Red Hat Quay sign-in

Setting the following will cause users to be prompted for their name and company when they first sign in. Although this is optional, it can provide you with extra data about your Red Hat Quay users:

```
+ FEATURE_USER_METADATA: true
```

### 1.3.2. Disable TLS Protocols

You can change the `SSL_PROTOCOLS` setting to remove SSL protocols that you do not want to support in your Red Hat Quay instance. For example, to remove TLS v1 support from the default `SSL_PROTOCOLS : ['TLSv1','TLSv1.1','TLSv1.2']`, change it as follows:

```
+ SSL_PROTOCOLS : ['TLSv1.1','TLSv1.2']
```

### 1.3.3. Rate limit API calls

Adding the `FEATURE_RATE_LIMITS` parameter to the `config.yaml` causes `nginx` to limit certain API calls to 30 per second. If that feature is not set, API calls are limited to 300 per second (effectively unlimited). Rate limiting can be an important feature, if you need to make sure the resources available are not overwhelmed with traffic.

Some namespace may require unlimited access (perhaps they are important to CI/CD and take priority, for example). In this case, those namespace may be placed in a list in `config.yaml` for `NON_RATE_LIMITED_NAMESPACES`.

### 1.3.4. Adjust database connection pooling

Red Hat Quay is composed of many different processes which all run within the same container. Many of these processes interact with the database.

If enabled, each process that interacts with the database will contain a connection pool. These per-process connection pools are configured to maintain a maximum of 20 connections. Under heavy load, it is possible to fill the connection pool for every process within a Red Hat Quay container. Under certain deployments and loads, this may require analysis to ensure Red Hat Quay does not exceed the database's configured maximum connection count.

Overtime, the connection pools will release idle connections. To release all connections immediately, Red Hat Quay requires a restart.

Database connection pooling may be toggled by setting the environment variable **DB\_CONNECTION\_POOLING={true|false}**

If database connection pooling is enabled, it is possible to change the maximum size of the connection pool. This can be done through the following **config.yaml** option:

```
DB_CONNECTION_ARGS:
  max_connections: 10
```

### 1.3.4.1. Database connection arguments

You can customize Red Hat Quay database connection settings within the **config.yaml** file. These are entirely dependent upon the underlying database driver, such as **psycopg2** for Postgres and **pymysql** for MySQL. It is also possible to pass in arguments used by Peewee's Connection Pooling mechanism as seen below.

```
DB_CONNECTION_ARGS:
  max_connections: n # Max Connection Pool size. (Connection Pooling only)
  timeout: n # Time to hold on to connections. (Connection Pooling only)
  stale_timeout: n # Number of seconds to block when the pool is full. (Connection Pooling only)
```

### 1.3.4.2. Database SSL configuration

Some key-value pairs defined under **DB\_CONNECTION\_ARGS** are generic while others are database-specific. In particular, SSL configuration depends on the database you are deploying.

#### 1.3.4.2.1. PostgreSQL SSL connection arguments

A sample PostgreSQL SSL configuration is given below:

```
DB_CONNECTION_ARGS:
  sslmode: verify-ca
  sslrootcert: /path/to/cacert
```

The **sslmode** option determines whether or with what priority a secure SSL TCP/IP connection will be negotiated with the server. There are six modes:

- **disable:** only try a non-SSL connection
- **allow:** first try a non-SSL connection; if that fails, try an SSL connection
- **prefer:** (default) first try an SSL connection; if that fails, try a non-SSL connection
- **require:** only try an SSL connection. If a root CA file is present, verify the certificate in the same way as if **verify-ca** was specified
- **verify-ca:** only try an SSL connection, and verify that the server certificate is issued by a trusted certificate authority (CA)
- **verify-full:** only try an SSL connection, verify that the server certificate is issued by a trusted CA and that the requested server host name matches that in the certificate

More information on the valid arguments for PostgreSQL is available at <https://www.postgresql.org/docs/current/libpq-connect.html>.

### 1.3.4.2.2. MySQL SSL connection arguments

A sample MySQL SSL configuration follows:

```
DB_CONNECTION_ARGS:
  ssl:
    ca: /path/to/cacert
```

Information on the valid connection arguments for MySQL is available at <https://dev.mysql.com/doc/refman/8.0/en/connecting-using-uri-or-key-value-pairs.html>.

### 1.3.4.3. HTTP connection counts

It is possible to specify the quantity of simultaneous HTTP connections using environment variables. These can be specified as a whole, or for a specific component. The default for each is 50 parallel connections per process.


Environment variables:

```
WORKER_CONNECTION_COUNT_REGISTRY=n
WORKER_CONNECTION_COUNT_WEB=n
WORKER_CONNECTION_COUNT_SECSCAN=n
WORKER_CONNECTION_COUNT=n
```

Specifying a count for a specific component will override any value set in `WORKER_CONNECTION_COUNT`.

### 1.3.4.4. Dynamic process counts

To estimate the quantity of dynamically sized processes, the following calculation is used by default.



#### NOTE

Red Hat Quay queries the available CPU count from the entire machine. Any limits applied using kubernetes or other non-virtualized mechanisms will not affect this behavior; Red Hat Quay will make its calculation based on the total number of processors on the Node. The default values listed are simply targets, but shall not exceed the maximum or be lower than the minimum.

Each of the following process quantities can be overridden using the environment variable specified below.

- registry - Provides HTTP endpoints to handle registry action
  - minimum: 8
  - maximum: 64
  - default:  $\$CPU\_COUNT \times 4$
  - environment variable: `WORKER_COUNT_REGISTRY`
- web - Provides HTTP endpoints for the web-based interface
  - minimum: 2

- maximum: 32
- default: \$CPU\_COUNT x 2
- environment\_variable: WORKER\_COUNT\_WEB
- seccan - Interacts with Clair
  - minimum: 2
  - maximum: 4
  - default: \$CPU\_COUNT x 2
  - environment variable: WORKER\_COUNT\_SECSCAN

### 1.3.4.5. Environment variables

Red Hat Quay allows overriding default behavior using environment variables. This table lists and describes each variable and the values they can expect.

**Table 1.1. Worker count environment variables**

Variable	Description	Values
WORKER_COUNT_REGISTRY	Specifies the number of processes to handle Registry requests within the <b>Quay</b> container.	Integer between 8 and 64
WORKER_COUNT_WEB	Specifies the number of processes to handle UI/Web requests within the container.	Integer between 2 and 32
WORKER_COUNT_SECSCAN	Specifies the number of processes to handle Security Scanning (e.g. Clair) integration within the container.	Integer between 2 and 4
DB_CONNECTION_POOLING	Toggle database connection pooling. In 3.4, it is disabled by default.	"true" or "false"

### 1.3.4.6. Turning off connection pooling

Red Hat Quay deployments with a large amount of user activity can regularly hit the 2k maximum database connection limit. In these cases, connection pooling, which is enabled by default for Red Hat Quay, can cause database connection count to rise exponentially and require you to turn off connection pooling.

If turning off connection pooling is not enough to prevent hitting that 2k database connection limit, you need to take additional steps to deal with the problem. In this case you might need to increase the maximum database connections to better suit your workload.

## CHAPTER 2. USING THE CONFIGURATION API

The configuration tool exposes 4 endpoints that can be used to build, validate, bundle and deploy a configuration. The config-tool API is documented at <https://github.com/quay/config-tool/blob/master/pkg/lib/editor/API.md>. In this section, you will see how to use the API to retrieve the current configuration and how to validate any changes you make.

### 2.1. RETRIEVING THE DEFAULT CONFIGURATION

If you are running the configuration tool for the first time, and do not have an existing configuration, you can retrieve the default configuration. Start the container in config mode:

```
$ sudo podman run --rm -it --name quay_config \  
-p 8080:8080 \  
registry.redhat.io/quay/quay-rhel8:v3.5.7 config secret
```

Use the **config** endpoint of the configuration API to get the default:

```
$ curl -X GET -u quayconfig:secret http://quay-server:8080/api/v1/config | jq
```

The value returned is the default configuration in JSON format:

```
{  
  "config.yaml": {  
    "AUTHENTICATION_TYPE": "Database",  
    "AVATAR_KIND": "local",  
    "DB_CONNECTION_ARGS": {  
      "autorollback": true,  
      "threadlocals": true  
    },  
    "DEFAULT_TAG_EXPIRATION": "2w",  
    "EXTERNAL_TLS_TERMINATION": false,  
    "FEATURE_ACTION_LOG_ROTATION": false,  
    "FEATURE_ANONYMOUS_ACCESS": true,  
    "FEATURE_APP_SPECIFIC_TOKENS": true,  
    ....  
  }  
}
```

### 2.2. RETRIEVING THE CURRENT CONFIGURATION

If you have already configured and deployed the Quay registry, stop the container and restart it in configuration mode, loading the existing configuration as a volume:

```
$ sudo podman run --rm -it --name quay_config \  
-p 8080:8080 \  
-v $QUAY/config:/conf/stack:Z \  
registry.redhat.io/quay/quay-rhel8:v3.5.7 config secret
```

Use the **config** endpoint of the API to get the current configuration:



```
$ curl -X GET -u quayconfig:secret http://quay-server:8080/api/v1/config | jq
```

The value returned is the current configuration in JSON format, including database and Redis configuration data:

```
{
  "config.yaml": {
    ...
    "BROWSER_API_CALLS_XHR_ONLY": false,
    "BUILDLOGS_REDIS": {
      "host": "quay-server",
      "password": "strongpassword",
      "port": 6379
    },
    "DATABASE_SECRET_KEY": "4b1c5663-88c6-47ac-b4a8-bb594660f08b",
    "DB_CONNECTION_ARGS": {
      "autorollback": true,
      "threadlocals": true
    },
    "DB_URI": "postgresql://quayuser:quaypass@quay-server:5432/quay",
    "DEFAULT_TAG_EXPIRATION": "2w",
    ...
  }
}
```

## 2.3. VALIDATING CONFIGURATION USING THE API

You can validate a configuration by posting it to the **config/validate** endpoint:

```
curl -u quayconfig:secret --header 'Content-Type: application/json' --request POST --data '
{
  "config.yaml": {
    ...
    "BROWSER_API_CALLS_XHR_ONLY": false,
    "BUILDLOGS_REDIS": {
      "host": "quay-server",
      "password": "strongpassword",
      "port": 6379
    },
    "DATABASE_SECRET_KEY": "4b1c5663-88c6-47ac-b4a8-bb594660f08b",
    "DB_CONNECTION_ARGS": {
      "autorollback": true,
      "threadlocals": true
    },
    "DB_URI": "postgresql://quayuser:quaypass@quay-server:5432/quay",
    "DEFAULT_TAG_EXPIRATION": "2w",
    ...
  }
}
' http://quay-server:8080/api/v1/config/validate | jq
```

The returned value is an array containing the errors found in the configuration. If the configuration is valid, an empty array [] is returned.

## 2.4. DETERMINING THE REQUIRED FIELDS

You can determine the required fields by posting an empty configuration structure to the **config/validate** endpoint:

```
curl -u quayconfig:secret --header 'Content-Type: application/json' --request POST --data '{
  "config.yaml": {
  }
}' http://quay-server:8080/api/v1/config/validate | jq
```


The value returned is an array indicating which fields are required:

```
[
  {
    "FieldGroup": "Database",
    "Tags": [
      "DB_URI"
    ],
    "Message": "DB_URI is required."
  },
  {
    "FieldGroup": "DistributedStorage",
    "Tags": [
      "DISTRIBUTED_STORAGE_CONFIG"
    ],
    "Message": "DISTRIBUTED_STORAGE_CONFIG must contain at least one storage location."
  },
  {
    "FieldGroup": "HostSettings",
    "Tags": [
      "SERVER_HOSTNAME"
    ],
    "Message": "SERVER_HOSTNAME is required"
  },
  {
    "FieldGroup": "HostSettings",
    "Tags": [
      "SERVER_HOSTNAME"
    ],
    "Message": "SERVER_HOSTNAME must be of type Hostname"
  },
  {
    "FieldGroup": "Redis",
    "Tags": [
      "BUILDLOGS_REDIS"
    ],
    "Message": "BUILDLOGS_REDIS is required"
  }
]
```


## CHAPTER 3. GETTING RED HAT QUAY RELEASE NOTIFICATIONS

To keep up with the latest Red Hat Quay releases and other changes related to Red Hat Quay, you can sign up for update notifications on the [Red Hat Customer Portal](#). After signing up for notifications, you will receive notifications letting you know when there is new a Red Hat Quay version, updated documentation, or other Red Hat Quay news.

1. Log into the [Red Hat Customer Portal](#) with your Red Hat customer account credentials.
2. Select your user name (upper-right corner) to see Red Hat Account and Customer Portal selections:



3. Select Notifications. Your profile activity page appears.
4. Select the Notifications tab.
5. Select Manage Notifications.
6. Select Follow, then choose Products from the drop-down box.
7. From the drop-down box next to the Products, search for and select Red Hat Quay:



8. Select the SAVE NOTIFICATION button. Going forward, you will receive notifications when there are changes to the Red Hat Quay product, such as a new release.

## CHAPTER 4. USING SSL TO PROTECT CONNECTIONS TO RED HAT QUAY

### 4.1. INTRODUCTION TO USING SSL

To configure Red Hat Quay with a [self-signed certificate](#), you need to create a Certificate Authority (CA) and then generate the required key and certificate files.

The following examples assume you have configured the server hostname **quay-server.example.com** using DNS or another naming mechanism, such as adding an entry in your **/etc/hosts** file:

```
$ cat /etc/hosts
...
192.168.1.112 quay-server.example.com
```

### 4.2. CREATE A CERTIFICATE AUTHORITY AND SIGN A CERTIFICATE

At the end of this procedure, you will have a certificate file and a primary key file named **ssl.cert** and **ssl.key**, respectively.

#### 4.2.1. Create a Certificate Authority

1. Generate the root CA key:

```
$ openssl genrsa -out rootCA.key 2048
```

2. Generate the root CA cert:

```
$ openssl req -x509 -new -nodes -key rootCA.key -sha256 -days 1024 -out rootCA.pem
```

3. Enter the information that will be incorporated into your certificate request, including the server hostname, for example:

```
Country Name (2 letter code) [XX]:IE
State or Province Name (full name) []:GALWAY
Locality Name (eg, city) [Default City]:GALWAY
Organization Name (eg, company) [Default Company Ltd]:QUAY
Organizational Unit Name (eg, section) []:DOCS
Common Name (eg, your name or your server's hostname) []:quay-server.example.com
```

#### 4.2.2. Sign a certificate

1. Generate the server key:

```
$ openssl genrsa -out ssl.key 2048
```

2. Generate a signing request:

```
$ openssl req -new -key ssl.key -out ssl.csr
```

3. Enter the information that will be incorporated into your certificate request, including the server hostname, for example:

```
Country Name (2 letter code) [XX]:IE
State or Province Name (full name) []:GALWAY
Locality Name (eg, city) [Default City]:GALWAY
Organization Name (eg, company) [Default Company Ltd]:QUAY
Organizational Unit Name (eg, section) []:DOCS
Common Name (eg, your name or your server's hostname) []:quay-server.example.com
```

4. Create a configuration file **openssl.cnf**, specifying the server hostname, for example:

#### openssl.cnf

```
[req]
req_extensions = v3_req
distinguished_name = req_distinguished_name
[req_distinguished_name]
[v3_req ]
basicConstraints = CA:FALSE
keyUsage = nonRepudiation, digitalSignature, keyEncipherment
subjectAltName = @alt_names
[alt_names]
DNS.1 = quay-server.example.com
IP.1 = 192.168.1.112
```

5. Use the configuration file to generate the certificate **ssl.cert**:

```
$ openssl x509 -req -in ssl.csr -CA rootCA.pem -CAkey rootCA.key -CAcreateserial -out
ssl.cert -days 356 -extensions v3_req -extfile openssl.cnf
```

## 4.3. CONFIGURING SSL USING THE COMMAND LINE

Another option when configuring SSL is to use the command line interface.

1. Copy the certificate file and primary key file to your configuration directory, ensuring they are named **ssl.cert** and **ssl.key** respectively:

```
$ cp ~/ssl.cert $QUAY/config
$ cp ~/ssl.key $QUAY/config
$ cd $QUAY/config
```

2. Edit the **config.yaml** file and specify that you want Quay to handle TLS:

#### config.yaml

```
...
SERVER_HOSTNAME: quay-server.example.com
...
PREFERRED_URL_SCHEME: https
...
```

3. Stop the **Quay** container and restart the registry:

```
$ sudo podman rm -f quay
$ sudo podman run -d --rm -p 80:8080 -p 443:8443 \
  --name=quay \
  -v $QUAY/config:/conf/stack:Z \
  -v $QUAY/storage:/datastorage:Z \
  registry.redhat.io/quay/quay-rhel8:v3.5.7
```

## 4.4. CONFIGURING SSL USING THE UI

This section configures SSL using the Quay UI. To configure SSL using the command line interface, see the following section.

1. Start the **Quay** container in configuration mode:

```
$ sudo podman run --rm -it --name quay_config -p 80:8080 -p 443:8443
registry.redhat.io/quay/quay-rhel8:v3.5.7 config secret
```

2. In the Server Configuration section, select **Red Hat Quay handles TLS** for TLS. Upload the certificate file and private key file created earlier, ensuring that the Server Hostname matches the value used when creating the certs. Validate and download the updated configuration.
3. Stop the **Quay** container and then restart the registry:

```
$ sudo podman rm -f quay
$ sudo podman run -d --rm -p 80:8080 -p 443:8443 \
  --name=quay \
  -v $QUAY/config:/conf/stack:Z \
  -v $QUAY/storage:/datastorage:Z \
  registry.redhat.io/quay/quay-rhel8:v3.5.7
```

## 4.5. TESTING SSL CONFIGURATION USING THE COMMAND LINE

- Use the **podman login** command to attempt to log in to the Quay registry with SSL enabled:

```
$ sudo podman login quay-server.example.com
Username: quayadmin
Password:
```

```
Error: error authenticating creds for "quay-server.example.com": error pinging docker registry
quay-server.example.com: Get "https://quay-server.example.com/v2/": x509: certificate
signed by unknown authority
```

- Podman does not trust self-signed certificates. As a workaround, use the **--tls-verify** option:

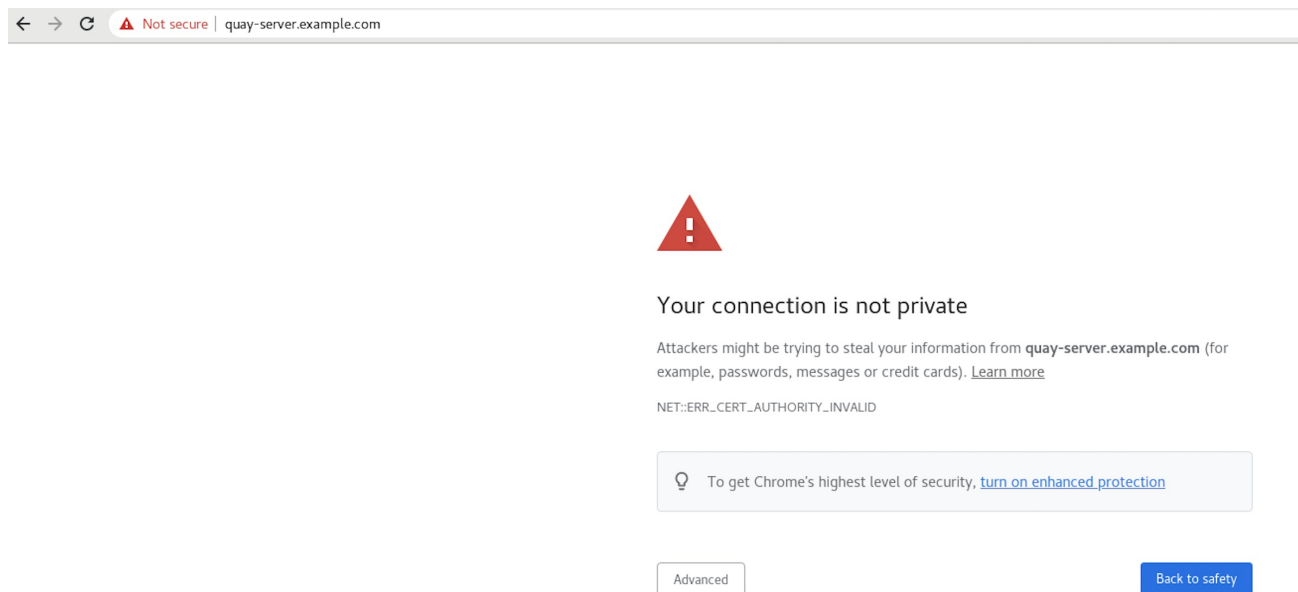
```
$ sudo podman login --tls-verify=false quay-server.example.com
Username: quayadmin
Password:
```

```
Login Succeeded!
```


Configuring Podman to trust the root Certificate Authority (CA) is covered in a subsequent section.

## 4.6. TESTING SSL CONFIGURATION USING THE BROWSER

When you attempt to access the Quay registry, in this case, <https://quay-server.example.com>, the browser warns of the potential risk:



Proceed to the log in screen, and the browser will notify you that the connection is not secure:



Configuring the system to trust the root Certificate Authority (CA) is covered in the subsequent section.

## 4.7. CONFIGURING PODMAN TO TRUST THE CERTIFICATE AUTHORITY

Podman uses two paths to locate the CA file, namely, `/etc/containers/certs.d/` and `/etc/docker/certs.d/`.

- Copy the root CA file to one of these locations, with the exact path determined by the server hostname, and naming the file **ca.crt**:

```
$ sudo cp rootCA.pem /etc/containers/certs.d/quay-server.example.com/ca.crt
```

- Alternatively, if you are using Docker, you can copy the root CA file to the equivalent Docker directory:

```
$ sudo cp rootCA.pem /etc/docker/certs.d/quay-server.example.com/ca.crt
```

You should no longer need to use the **--tls-verify=false** option when logging in to the registry:

```
$ sudo podman login quay-server.example.com
```

```
Username: quayadmin
```

```
Password:
```

```
Login Succeeded!
```

## 4.8. CONFIGURING THE SYSTEM TO TRUST THE CERTIFICATE AUTHORITY

1. Copy the root CA file to the consolidated system-wide trust store:

```
$ sudo cp rootCA.pem /etc/pki/ca-trust/source/anchors/
```

2. Update the system-wide trust store configuration:

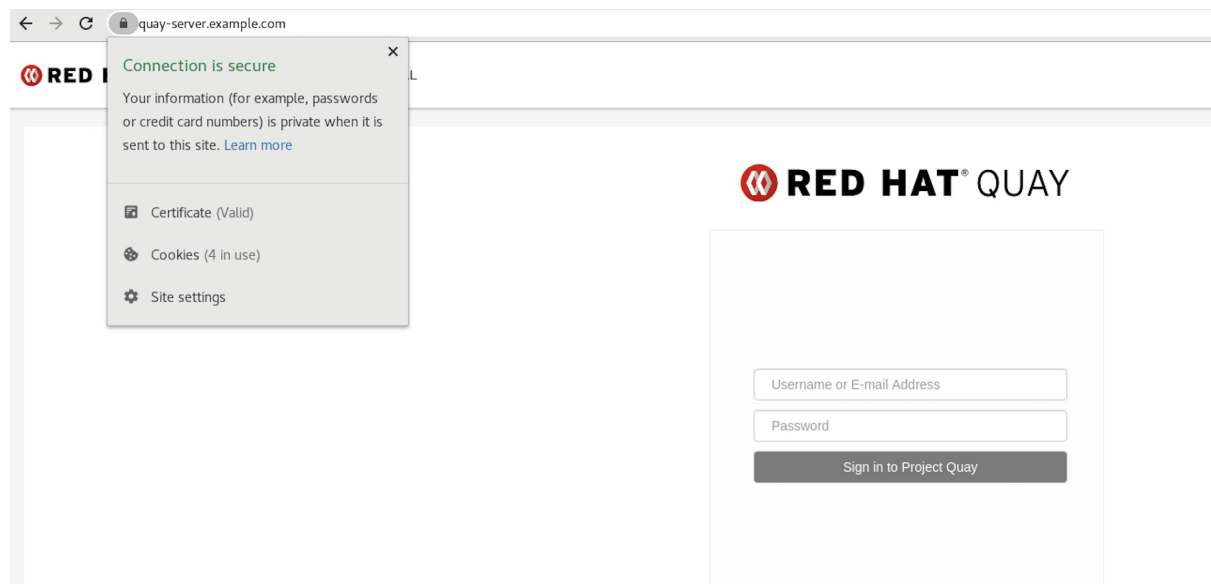
```
$ sudo update-ca-trust extract
```

3. You can use the **trust list** command to ensure that the Quay server has been configured:

```
$ trust list | grep quay  
label: quay-server.example.com
```

Now, when you browse to the registry at <https://quay-server.example.com>, the lock icon shows that the connection is secure:





4. To remove the root CA from system-wide trust, delete the file and update the configuration:

```
$ sudo rm /etc/pki/ca-trust/source/anchors/rootCA.pem
$ sudo update-ca-trust extract
$ trust list | grep quay
$
```

More information can be found in the RHEL 8 documentation in the chapter [Using shared system certificates](#).

## CHAPTER 5. ADDING TLS CERTIFICATES TO THE RED HAT QUAY CONTAINER

To add custom TLS certificates to Red Hat Quay, create a new directory named **extra\_ca\_certs/** beneath the Red Hat Quay config directory. Copy any required site-specific TLS certificates to this new directory.

### 5.1. ADD TLS CERTIFICATES TO RED HAT QUAY

1. View certificate to be added to the container

```
$ cat storage.crt
-----BEGIN CERTIFICATE-----
MIIDTTCCAjWgAwIbAgIJAMVr9ngjJhzbMA0GCSqGSIb3DQEBCwUAMD0xCzAJBgNV
[...]
-----END CERTIFICATE-----
```

2. Create certs directory and copy certificate there

```
$ mkdir -p quay/config/extra_ca_certs
$ cp storage.crt quay/config/extra_ca_certs/
$ tree quay/config/
|
|--- config.yaml
|--- extra_ca_certs
|    |--- storage.crt
```

3. Obtain the **Quay** container's **CONTAINER ID** with **podman ps**:

```
$ sudo podman ps
CONTAINER ID      IMAGE                                COMMAND                                CREATED
STATUS           PORTS                                GRAVEYARD                                UPTIME
5a3e82c4a75f     <registry>/<repo>/quay:v3.5.7 "/sbin/my_init" 24 hours ago    Up
18 hours        0.0.0.0:80->80/tcp, 0.0.0.0:443->443/tcp, 443/tcp  grave_keller
```

4. Restart the container with that ID:

```
$ sudo podman restart 5a3e82c4a75f
```

5. Examine the certificate copied into the container namespace:

```
$ sudo podman exec -it 5a3e82c4a75f cat /etc/ssl/certs/storage.pem
-----BEGIN CERTIFICATE-----
MIIDTTCCAjWgAwIbAgIJAMVr9ngjJhzbMA0GCSqGSIb3DQEBCwUAMD0xCzAJBgNV
```

### 5.2. ADD CERTS WHEN DEPLOYED ON KUBERNETES

When deployed on Kubernetes, Red Hat Quay mounts in a secret as a volume to store config assets. Unfortunately, this currently breaks the upload certificate function of the superuser panel.

To get around this error, a base64 encoded certificate can be added to the secret *after* Red Hat Quay has been deployed. Here's how:

1. Begin by base64 encoding the contents of the certificate:

```
$ cat ca.crt
-----BEGIN CERTIFICATE-----
MIIDljCCAn6gAwIBAgIBATANBgkqhkiG9w00BAQsFADA5MRcwFQYDVQQKDA5MQUlu
TEICQ09SRS5TTzEeMBwGA1UEAwwVQ2VydGlmaWNhdGUgQXV0aG9yaXR5MB4XDTE2
MDExMjA2NTkxMfoXDTM2MDExMjA2NTkxMFowOTEXMBUGA1UECgwOTEFCLkxJQkNP
UkUuU08xHjAcBgNVBAMMFUNlcnRpZmljYXRlIEF1dGhvcml0eTCCASlwDQYJKoZI
[...]
-----END CERTIFICATE-----

$ cat ca.crt | base64 -w 0
[...]
c1psWGpqeGIPQmNEWkJPMjJ5d0pDemVnR2QNCnRsbW9JdEF4YnFSdVd3PT0KLS0tLS1F
TkQgQ0VSVEIGSUNBVEUtLS0tLQo=
```

2. Use the **kubectl** tool to edit the quay-enterprise-config-secret.

```
$ kubectl --namespace quay-enterprise edit secret/quay-enterprise-config-secret
```

3. Add an entry for the cert and paste the full base64 encoded string under the entry:

```
custom-cert.crt:
c1psWGpqeGIPQmNEWkJPMjJ5d0pDemVnR2QNCnRsbW9JdEF4YnFSdVd3PT0KLS0tLS1F
TkQgQ0VSVEIGSUNBVEUtLS0tLQo=
```

4. Finally, recycle all Red Hat Quay pods. Use **kubectl delete** to remove all Red Hat Quay pods. The Red Hat Quay Deployment will automatically schedule replacement pods with the new certificate data.

## CHAPTER 6. CONFIGURING ACTION LOG STORAGE FOR ELASTICSEARCH

By default, the past three months of usage logs are stored in the Red Hat Quay database and exposed via the web UI on organization and repository levels. Appropriate administrative privileges are required to see log entries. For deployments with a large amount of logged operations, you can now store the usage logs in Elasticsearch instead of the Red Hat Quay database backend. To do this, you need to provide your own Elasticsearch stack, as it is not included with Red Hat Quay as a customizable component.

Enabling Elasticsearch logging can be done during Red Hat Quay deployment or post-deployment using the Red Hat Quay Config Tool. The resulting configuration is stored in the **config.yaml** file. Once configured, usage log access continues to be provided the same way, via the web UI for repositories and organizations.

Here's how to configure action log storage to change it from the default Red Hat Quay database to use Elasticsearch:

1. Obtain an Elasticsearch account.
2. Open the Red Hat Quay Config Tool (either during or after Red Hat Quay deployment).
3. Scroll to the *Action Log Storage Configuration* setting and select *Elasticsearch* instead of *Database*. The following figure shows the Elasticsearch settings that appear:

**Action Log Storage Configuration**

Action logs can be stored in the database or Elasticsearch. In the latter case, the actions logs can (optionally) be sent to a data stream first.

**Action Logs Storage:**

**Elasticsearch hostname:**

**Elasticsearch port:**  Access to this port and hostname must be allowed from all hosts running the enterprise registry

**Elasticsearch access key:**

**Elasticsearch secret key:**

**AWS region:**

**Index prefix:**

**Logs Producer:**

4. Fill in the following information for your Elasticsearch instance:
  - **Elasticsearch hostname:** The hostname or IP address of the system providing the Elasticsearch service.
  - **Elasticsearch port:** The port number providing the Elasticsearch service on the host you just entered. Note that the port must be accessible from all systems running the Red Hat Quay registry. The default is TCP port 9200.

- **Elasticsearch access key.** The access key needed to gain access to the Elastic search service, if required.
- **Elasticsearch secret key.** The secret key needed to gain access to the Elastic search service, if required.
- **AWS region:** If you are running on AWS, set the AWS region (otherwise, leave it blank).
- **Index prefix** Choose a prefix to attach to log entries.
- **Logs Producer:** Choose either Elasticsearch (default) or Kinesis to direct logs to an intermediate Kinesis stream on AWS. You need to set up your own pipeline to send logs from Kinesis to Elasticsearch (for example, Logstash). The following figure shows additional fields you would need to fill in for Kinesis:

The screenshot shows a configuration form with the following fields:

- AWS region:** The AWS region
- Index prefix:** logentry\_
- Logs Producer:** Kinesis (highlighted in a red box)
- Stream name:** The Kinesis stream name
- AWS access key:** The AWS access key
- AWS secret key:** The AWS secret key
- AWS region:** The AWS region

Below the form is a yellow button with a downward arrow and the text "9 configuration fields remaining".

5. If you chose Elasticsearch as the Logs Producer, no further configuration is needed. If you chose Kinesis, fill in the following:
  - **Stream name:** The name of the Kinesis stream.
  - **AWS access key.** The name of the AWS access key needed to gain access to the Kinesis stream, if required.
  - **AWS secret key.** The name of the AWS secret key needed to gain access to the Kinesis stream, if required.
  - **AWS region:** The AWS region.
6. When you are done, save the configuration. The Config Tool checks your settings. If there is a problem connecting to the Elasticsearch or Kinesis services, you will see an error and have the opportunity to continue editing. Otherwise, logging will begin to be directed to your Elasticsearch configuration after the cluster restarts with the new configuration.

## CHAPTER 7. CLAIR SECURITY SCANNING

Clair is a set of micro services that can be used with Red Hat Quay to perform vulnerability scanning of container images associated with a set of Linux operating systems. The micro services design of Clair makes it appropriate to run in a highly scalable configuration, where components can be scaled separately as appropriate for enterprise environments.

Clair uses the following vulnerability databases to scan for issues in your images:

- Alpine SecDB database
- AWS UpdateInfo
- Debian Oval database
- Oracle Oval database
- RHEL Oval database
- SUSE Oval database
- Ubuntu Oval database
- Pyup.io (python) database

For information on how Clair does security mapping with the different databases, see [ClairCore Severity Mapping](#).



### NOTE

With the release of Red Hat Quay 3.4, the new Clair V4 (image `registry.redhat.io/quay/clair-rhel8`) fully replaces the prior Clair V2 (image `quay.io/redhat/clair-jwt`). See below for how to run V2 in read-only mode while V4 is updating.

## 7.1. SETTING UP CLAIR ON A RED HAT QUAY OPENSIFT DEPLOYMENT

### 7.1.1. Deploying Via the Quay Operator

To set up Clair V4 on a new Red Hat Quay deployment on OpenShift, it is highly recommended to use the Quay Operator. By default, the Quay Operator will install or upgrade a Clair deployment along with your Red Hat Quay deployment and configure Clair security scanning automatically.

### 7.1.2. Manually Deploying Clair

To configure Clair V4 on an existing Red Hat Quay OpenShift deployment running Clair V2, first ensure Red Hat Quay has been upgraded to at least version 3.4.0. Then use the following steps to manually set up Clair V4 alongside Clair V2.

1. Set your current project to the name of the project in which Red Hat Quay is running. For example:

```
$ oc project quay-enterprise
```

2. Create a Postgres deployment file for Clair v4 (for example, **clairv4-postgres.yaml**) as follows.

### clairv4-postgres.yaml

```

---
apiVersion: apps/v1
kind: Deployment
metadata:
  name: clairv4-postgres
  namespace: quay-enterprise
  labels:
    quay-component: clairv4-postgres
spec:
  replicas: 1
  selector:
    matchLabels:
      quay-component: clairv4-postgres
  template:
    metadata:
      labels:
        quay-component: clairv4-postgres
    spec:
      volumes:
        - name: postgres-data
          persistentVolumeClaim:
            claimName: clairv4-postgres
      containers:
        - name: postgres
          image: postgres:11.5
          imagePullPolicy: "IfNotPresent"
          ports:
            - containerPort: 5432
          env:
            - name: POSTGRES_USER
              value: "postgres"
            - name: POSTGRES_DB
              value: "clair"
            - name: POSTGRES_PASSWORD
              value: "postgres"
            - name: PGDATA
              value: "/etc/postgres/data"
          volumeMounts:
            - name: postgres-data
              mountPath: "/etc/postgres"
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: clairv4-postgres
  labels:
    quay-component: clairv4-postgres
spec:
  accessModes:
    - "ReadWriteOnce"
  resources:
    requests:

```

```

    storage: "5Gi"
    volumeName: "clairv4-postgres"
---
apiVersion: v1
kind: Service
metadata:
  name: clairv4-postgres
  labels:
    quay-component: clairv4-postgres
spec:
  type: ClusterIP
  ports:
    - port: 5432
      protocol: TCP
      name: postgres
      targetPort: 5432
  selector:
    quay-component: clairv4-postgres

```

3. Deploy the postgres database as follows:

```
$ oc create -f ./clairv4-postgres.yaml
```

4. Create a Clair **config.yaml** file to use for Clair v4. For example:

### config.yaml

```

introspection_addr: :8089
http_listen_addr: :8080
log_level: debug
indexer:
  connstring: host=clairv4-postgres port=5432 dbname=clair user=postgres
  password=postgres sslmode=disable
  scanlock_retry: 10
  layer_scan_concurrency: 5
  migrations: true
matcher:
  connstring: host=clairv4-postgres port=5432 dbname=clair user=postgres
  password=postgres sslmode=disable
  max_conn_pool: 100
  run: ""
  migrations: true
  indexer_addr: clair-indexer
notifier:
  connstring: host=clairv4-postgres port=5432 dbname=clair user=postgres
  password=postgres sslmode=disable
  delivery: 1m
  poll_interval: 5m
  migrations: true
auth:
  psk:
    key: MTU5YzA4Y2ZkNzJoMQ== 1
    iss: ["quay"]
# tracing and metrics
trace:

```



```

name: "jaeger"
probability: 1
jaeger:
  agent_endpoint: "localhost:6831"
  service_name: "clair"
metrics:
  name: "prometheus"

```

- 1 To generate a Clair pre-shared key (PSK), enable **scanning** in the Security Scanner section of the User Interface and click **Generate PSK**.

More information about Clair's configuration format can be found in [upstream Clair documentation](#).

1. Create a secret from the Clair **config.yaml**:

```
$ oc create secret generic clairv4-config-secret --from-file=./config.yaml
```

2. Create the Clair v4 deployment file (for example, **clair-combo.yaml**) and modify it as necessary:

### clair-combo.yaml

```

---
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  labels:
    quay-component: clair-combo
  name: clair-combo
spec:
  replicas: 1
  selector:
    matchLabels:
      quay-component: clair-combo
  template:
    metadata:
      labels:
        quay-component: clair-combo
    spec:
      containers:
        - image: registry.redhat.io/quay/clair-rhel8:v3.5.7 1
          imagePullPolicy: IfNotPresent
          name: clair-combo
          env:
            - name: CLAIR_CONF
              value: /clair/config.yaml
            - name: CLAIR_MODE
              value: combo
      ports:
        - containerPort: 8080
          name: clair-http
          protocol: TCP
        - containerPort: 8089
          name: clair-intro

```

```

    protocol: TCP
    volumeMounts:
      - mountPath: /clair/
        name: config
    imagePullSecrets:
      - name: redhat-pull-secret
    restartPolicy: Always
    volumes:
      - name: config
        secret:
          secretName: clairv4-config-secret
  ---
  apiVersion: v1
  kind: Service
  metadata:
    name: clairv4 2
  labels:
    quay-component: clair-combo
  spec:
    ports:
      - name: clair-http
        port: 80
        protocol: TCP
        targetPort: 8080
      - name: clair-introspection
        port: 8089
        protocol: TCP
        targetPort: 8089
    selector:
      quay-component: clair-combo
  type: ClusterIP

```

- 1 Change image to latest clair image name and version.
- 2 With the Service set to clairv4, the scanner endpoint for Clair v4 is entered later into the Red Hat Quay config.yaml in the **SECURITY\_SCANNER\_V4\_ENDPOINT** as <http://clairv4>.

3. Create the Clair v4 deployment as follows:

```
$ oc create -f ./clair-combo.yaml
```

4. Modify the **config.yaml** file for your Red Hat Quay deployment to add the following entries at the end:

```

FEATURE_SECURITY_SCANNER: true
SECURITY_SCANNER_V4_ENDPOINT: http://clairv4 1

```

- 1 Identify the Clair v4 service endpoint
5. Redeploy the modified **config.yaml** to the secret containing that file (for example, **quay-enterprise-config-secret**):

```
$ oc delete secret quay-enterprise-config-secret
$ oc create secret generic quay-enterprise-config-secret --from-file=./config.yaml
```

6. For the new **config.yaml** to take effect, you need to restart the Red Hat Quay pods. Simply deleting the **quay-app** pods causes pods with the updated configuration to be deployed.

At this point, images in any of the organizations identified in the namespace whitelist will be scanned by Clair v4.

## 7.2. SETTING UP CLAIR ON A NON-OPENSIFT RED HAT QUAY DEPLOYMENT

For Red Hat Quay deployments not running on OpenShift, it is possible to configure Clair security scanning manually. Red Hat Quay deployments already running Clair V2 can use the instructions below to add Clair V4 to their deployment.

1. Deploy a (preferably fault-tolerant) Postgres database server. Note that Clair requires the **uuid-oss** extension to be added to its Postgres database. If the user supplied in Clair's **config.yaml** has the necessary privileges to create the extension then it will be added automatically by Clair itself. If not, then the extension must be added before starting Clair. If the extension is not present, the following error will be displayed when Clair attempts to start.

```
ERROR: Please load the "uuid-oss" extension. (SQLSTATE 42501)
```

2. Create a Clair config file in a specific folder, for example, **/etc/clairv4/config/config.yaml**).

### config.yaml

```
introspection_addr: :8089
http_listen_addr: :8080
log_level: debug
indexer:
  connstring: host=clairv4-postgres port=5432 dbname=clair user=postgres
  password=postgres sslmode=disable
  scanlock_retry: 10
  layer_scan_concurrency: 5
  migrations: true
matcher:
  connstring: host=clairv4-postgres port=5432 dbname=clair user=postgres
  password=postgres sslmode=disable
  max_conn_pool: 100
  run: ""
  migrations: true
  indexer_addr: clair-indexer
notifier:
  connstring: host=clairv4-postgres port=5432 dbname=clair user=postgres
  password=postgres sslmode=disable
  delivery_interval: 1m
  poll_interval: 5m
  migrations: true

# tracing and metrics
trace:
  name: "jaeger"
```

```

probability: 1
jaeger:
  agent_endpoint: "localhost:6831"
  service_name: "clair"
metrics:
  name: "prometheus"

```

More information about Clair's configuration format can be found in [upstream Clair documentation](#).

1. Run Clair via the container image, mounting in the configuration from the file you created.

```

$ podman run -p 8080:8080 -p 8089:8089 -e CLAIR_CONF=/clair/config.yaml -e
CLAIR_MODE=combo -v /etc/clair4/config:/clair -d registry.redhat.io/quay/clair-rhel8:v3.5.7

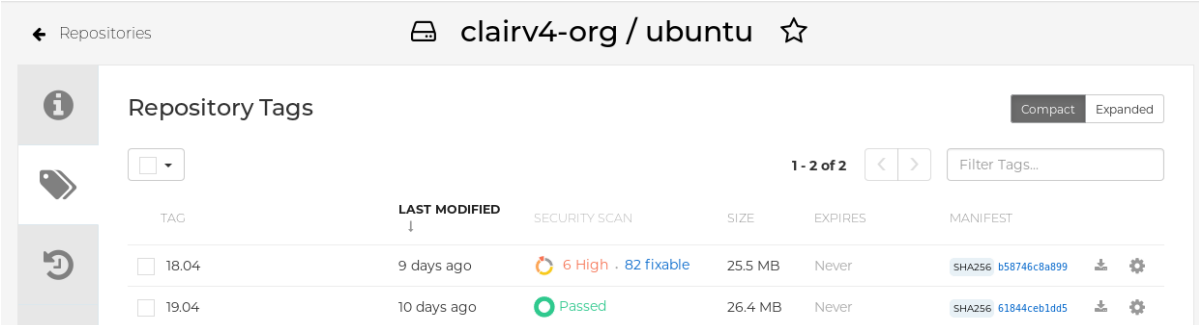
```

2. Follow the remaining instructions from the previous section for configuring Red Hat Quay to use the new Clair V4 endpoint.

Running multiple Clair containers in this fashion is also possible, but for deployment scenarios beyond a single container the use of a container orchestrator like Kubernetes or OpenShift is strongly recommended.

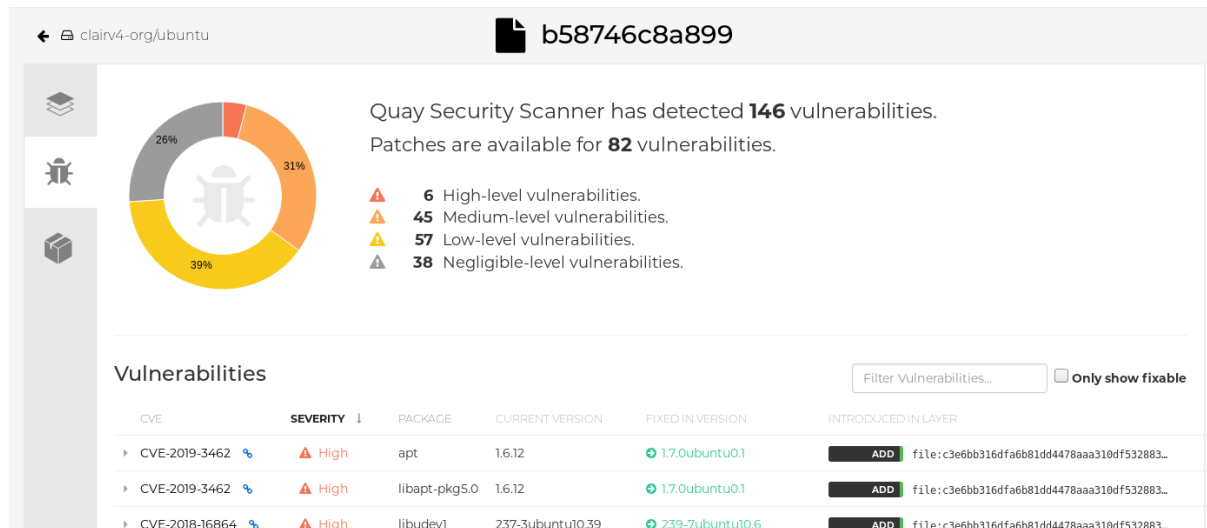
## 7.3. USING CLAIR

1. Log in to your Red Hat Quay cluster and select an organization for which you have configured Clair scanning.
2. Select a repository from that organization that holds some images and select Tags from the left navigation. The following figure shows an example of a repository with two images that have been scanned:



TAG	LAST MODIFIED	SECURITY SCAN	SIZE	EXPIRES	MANIFEST
18.04	9 days ago	6 High · 82 fixable	25.5 MB	Never	SHA256: b58746c8a899
19.04	10 days ago	Passed	26.4 MB	Never	SHA256: 61844ceb1dd5

3. If vulnerabilities are found, select to under the Security Scan column for the image to see either all vulnerabilities or those that are fixable. The following figure shows information on all vulnerabilities found:



## 7.4. CONFIGURING CLAIR FOR DISCONNECTED ENVIRONMENTS

Clair utilizes a set of components called Updaters to handle the fetching and parsing of data from various vulnerability databases. These Updaters are set up by default to pull vulnerability data directly from the internet and work out of the box. For customers in disconnected environments without direct access to the internet this poses a problem. Clair supports these environments through the ability to work with different types of update workflows that take into account network isolation. Using the **clairctl** command line utility, any process can easily fetch Updater data from the internet via an open host, securely transfer the data to an isolated host, and then import the Updater data on the isolated host into Clair itself.

The steps are as follows.

1. First ensure that your Clair configuration has disabled automated Updaters from running.

### config.yaml

```
matcher:
  disable_updaters: true
```

2. Export out the latest Updater data to a local archive. This requires the **clairctl** tool which can be run directly as a binary, or via the Clair container image. Assuming your Clair configuration is in **/etc/clairv4/config/config.yaml**, to run via the container image:

```
$ podman run -it --rm -v /etc/clairv4/config:/cfg:Z -v /path/to/output/directory:/updaters:Z --
  entrypoint /bin/clairctl registry.redhat.io/quay/clair-rhel8:v3.5.7 --config /cfg/config.yaml
  export-updaters /updaters/updaters.gz
```

Note that you need to explicitly reference the Clair configuration. This will create the Updater archive in **/etc/clairv4/updaters/updaters.gz**. If you want to ensure the archive was created without any errors from the source databases, you can supply the **--strict** flag to **clairctl**. The archive file should be copied over to a volume that is accessible from the disconnected host running Clair. From the disconnected host, use the same procedure now to import the archive into Clair.

```
$ podman run -it --rm -v /etc/clairv4/config:/cfg:Z -v /path/to/output/directory:/updaters:Z --
  entrypoint /bin/clairctl registry.redhat.io/quay/clair-rhel8:v3.5.7 --config /cfg/config.yaml
  import-updaters /updaters/updaters.gz
```

## 7.5. CLAIR UPDATER URLS

The following are the HTTP hosts and paths that Clair will attempt to talk to in a default configuration. This list is non-exhaustive, as some servers will issue redirects and some request URLs are constructed dynamically.

- <https://secdb.alpinelinux.org/>
- [http://repo.us-west-2.amazonaws.com/2018.03/updates/x86\\_64/mirror.list](http://repo.us-west-2.amazonaws.com/2018.03/updates/x86_64/mirror.list)
- [https://cdn.amazonlinux.com/2/core/latest/x86\\_64/mirror.list](https://cdn.amazonlinux.com/2/core/latest/x86_64/mirror.list)
- <https://www.debian.org/security/oval/>
- <https://linux.oracle.com/security/oval/>
- [https://packages.vmware.com/photon/photon\\_oval\\_definitions/](https://packages.vmware.com/photon/photon_oval_definitions/)
- <https://github.com/pyupio/safety-db/archive/>
- <https://catalog.redhat.com/api/containers/>
- <https://www.redhat.com/security/data/>
- <https://support.novell.com/security/oval/>
- <https://people.canonical.com/~ubuntu-security/oval/>

## 7.6. ADDITIONAL INFORMATION

For detailed documentation on the internals of Clair, including how the microservices are structured, please see the [Upstream Clair](#) and [ClairCore](#) documentation.

## CHAPTER 8. SCAN POD IMAGES WITH THE CONTAINER SECURITY OPERATOR

Using the [Container Security Operator](#), (CSO) you can scan container images associated with active pods, running on OpenShift (4.2 or later) and other Kubernetes platforms, for known vulnerabilities. The CSO:

- Watches containers associated with pods on all or specified namespaces
- Queries the container registry where the containers came from for vulnerability information provided an image's registry supports image scanning (such as a Quay registry with Clair scanning)
- Exposes vulnerabilities via the ImageManifestVuln object in the Kubernetes API

Using the instructions here, the CSO is installed in the **marketplace-operators** namespace, so it is available to all namespaces on your OpenShift cluster.



### NOTE

To see instructions on installing the CSO on Kubernetes, select the Install button from the [Container Security OperatorHub.io](#) page.

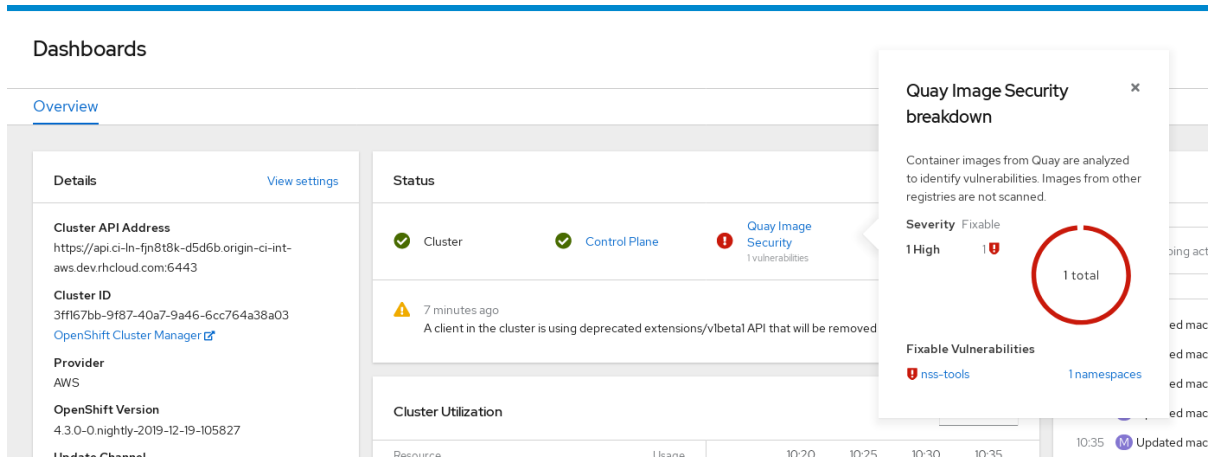
### 8.1. RUN THE CSO IN OPENSIFT

To start using the CSO in OpenShift, do the following:

1. Go to Operators → OperatorHub (select Security) to see the available **Container Security** Operator.
2. Select the **Container Security** Operator, then select **Install** to go to the Create Operator Subscription page.
3. Check the settings (all namespaces and automatic approval strategy, by default), and select **Subscribe**. The **Container Security** appears after a few moments on the **Installed Operators** screen.
4. Optionally, you can add custom certificates to the CSO. In this example, create a certificate named quay.crt in the current directory. Then run the following command to add the cert to the CSO (restart the Operator pod for the new certs to take effect):

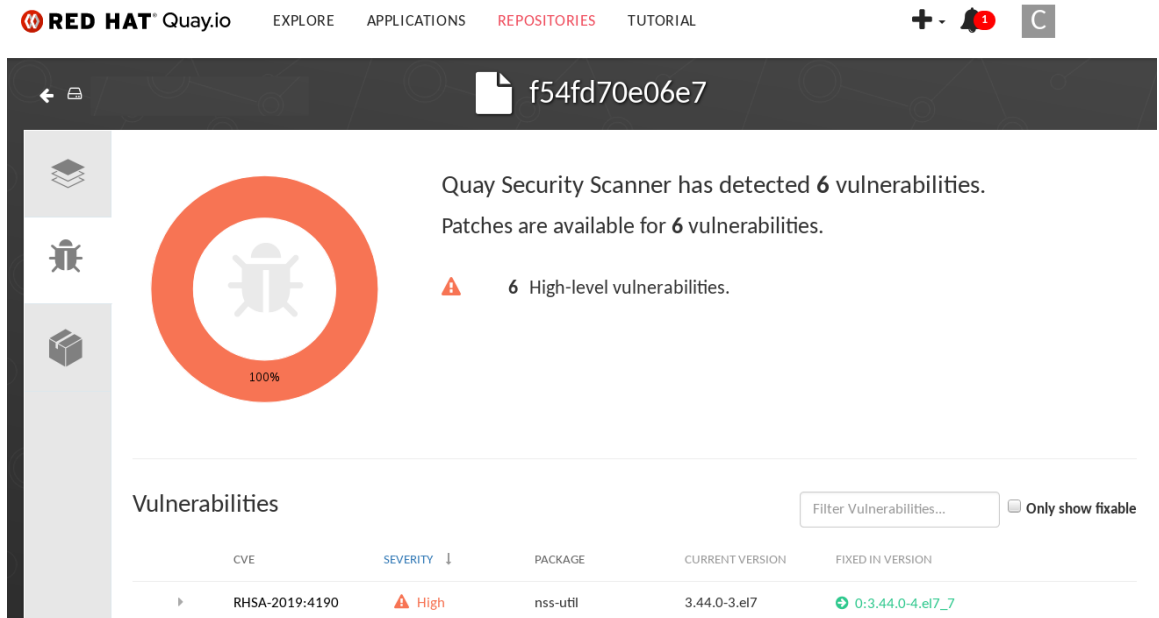
```
$ oc create secret generic container-security-operator-extra-certs --from-file=quay.crt -n openshift-operators
```

5. Open the OpenShift Dashboard (Home → Dashboards). A link to Image Security appears under the status section, with a listing of the number of vulnerabilities found so far. Select the link to see a Security breakdown, as shown in the following figure:

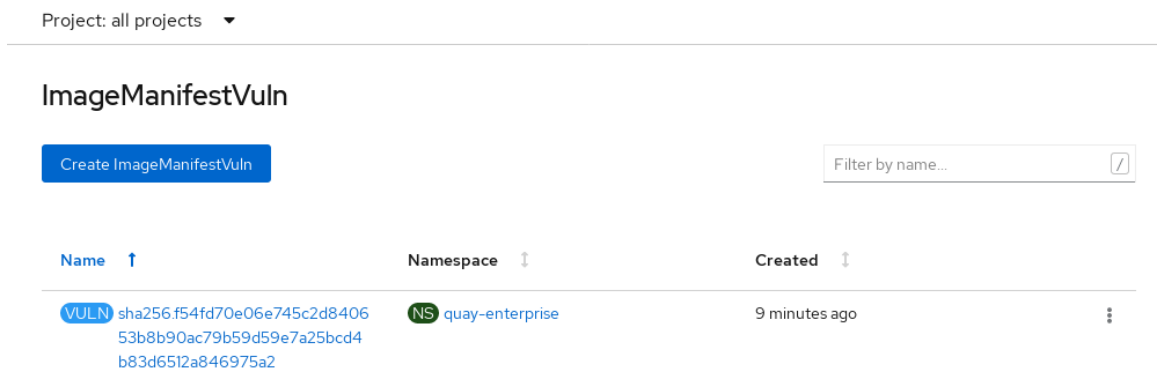


6. You can do one of two things at this point to follow up on any detected vulnerabilities:

- Select the link to the vulnerability. You are taken to the container registry, Red Hat Quay or other registry where the container came from, where you can see information about the vulnerability. The following figure shows an example of detected vulnerabilities from a Quay.io registry:



- Select the namespaces link to go to the ImageManifestVuln screen, where you can see the name of the selected image and all namespaces where that image is running. The following figure indicates that a particular vulnerable image is running in two namespaces:



At this point, you know what images are vulnerable, what you need to do to fix those vulnerabilities, and every namespace that the image was run in. So you can:



- Alert anyone running the image that they need to correct the vulnerability
- Stop the images from running (by deleting the deployment or other object that started the pod the image is in)

Note that if you do delete the pod, it may take a few minutes for the vulnerability to reset on the dashboard.

## 8.2. QUERY IMAGE VULNERABILITIES FROM THE CLI

You can query information on security from the command line. To query for detected vulnerabilities, type:

```
$ oc get vuln --all-namespaces
NAMESPACE  NAME                AGE
default    sha256.ca90...     6m56s
skynet     sha256.ca90...     9m37s
```

To display details for a particular vulnerability, identify one of the vulnerabilities, along with its namespace and the **describe** option. This example shows an active container whose image includes an RPM package with a vulnerability:

```
$ oc describe vuln --namespace mynamespace sha256.ac50e3752...
Name:      sha256.ac50e3752...
Namespace: quay-enterprise
...
Spec:
  Features:
    Name:      nss-util
    Namespace Name: centos:7
    Version:   3.44.0-3.el7
    Versionformat: rpm
  Vulnerabilities:
    Description: Network Security Services (NSS) is a set of libraries...
```

## CHAPTER 9. INTEGRATE RED HAT QUAY INTO OPENSIFT WITH THE BRIDGE OPERATOR

Using the Quay Bridge Operator, you can replace the integrated container registry in OpenShift with a Red Hat Quay registry. By doing this, your integrated OpenShift registry becomes a highly available, enterprise-grade Red Hat Quay registry with enhanced role based access control (RBAC) features.

The primary goals of the Bridge Operator is to duplicate the features of the integrated OpenShift registry in the new Red Hat Quay registry. The features enabled with this Operator include:

- Synchronizing OpenShift namespaces as Red Hat Quay organizations.
  - Creating Robot accounts for each default namespace service account
  - Creating Secrets for each created Robot Account (associating each Robot Secret to a Service Account as Mountable and Image Pull Secret)
  - Synchronizing OpenShift ImageStreams as Quay Repositories
- Automatically rewriting new Builds making use of ImageStreams to output to Red Hat Quay
- Automatically importing an ImageStream tag once a build completes

Using this procedure with the Quay Bridge Operator, you enable bi-directional communication between your Red Hat Quay and OpenShift clusters.



### WARNING

You cannot have more than one OpenShift Container Platform cluster pointing to the same Red Hat Quay instance from a Quay Bridge Operator. If you did, it would prevent you from creating namespaces of the same name on the two clusters.

## 9.1. RUNNING THE QUAY BRIDGE OPERATOR

### 9.1.1. Prerequisites

Before setting up the Bridge Operator, have the following in place:

- An existing Red Hat Quay environment for which you have superuser permissions
- A Red Hat OpenShift Container Platform environment (4.2 or later is recommended) for which you have cluster administrator permissions
- An OpenShift command line tool (**oc** command)

### 9.1.2. Setting up and configuring OpenShift and Red Hat Quay

Both Red Hat Quay and OpenShift configuration is required:

### 9.1.2.1. Red Hat Quay setup

Create a dedicated Red Hat Quay organization, and from a new application you create within that organization, generate an OAuth token to be used with the Quay Bridge Operator in OpenShift

1. Log in to Red Hat Quay as a user with superuser access and select the organization for which the external application will be configured.
2. In the left navigation, select Applications.
3. Select **Create New Application** and entering a name for the new application (for example, **openshift**).
4. With the new application displayed, select it.
5. In the left navigation, select **Generate Token** to create a new OAuth2 token.
6. Select all checkboxes to grant the access needed for the integration.
7. Review the assigned permissions and then select **Authorize Application**, then confirm it.
8. Copy and save the generated Access Token that appears to use in the next section.

### 9.1.2.2. OpenShift Setup

Setting up OpenShift for the Quay Bridge Operator requires several steps, including:

- **Creating an OpenShift secret** Using the OAuth token created earlier in Quay, create an OpenShift secret.
- **Adding MutatingWebhookConfiguration support:** To support Red Hat Quay integration with OpenShift, any new Build requests should be intercepted so that the output can be modified to target Red Hat Quay instead of OpenShift's integrated registry.

Support for dynamic interception of API requests that are performed as part of OpenShift's typical build process is facilitated through a MutatingWebhookConfiguration. A MutatingWebhookConfiguration allows for invoking an API running within a project on OpenShift when certain API requests are received.

Kubernetes requires that the webhook endpoint is secured via SSL using a certificate that makes use of the certificate authority for the cluster. Fortunately, OpenShift provides support for generating a certificate signed by the cluster.

1. Using the OpenShift **oc** command line tool, log in to OpenShift as a cluster administrator.
2. Choose an OpenShift namespace to use, such as **openshift-operators** or create a new one.
3. Create an OpenShift secret, replacing `<access_token>` with the Access Token obtained earlier from Red Hat Quay. For example, this creates a secret with your `<access_token>` called **quay-integration** with a key called **token**:

```
$ oc create secret generic quay-integration --from-literal=token=<access_token>
```

The result places the newly created private key and certificate within a secret specified. The secret will be mounted into the appropriate located within the operator as declared in the Deployment of the Operator.

4. Create a Service for the Operator's webhook endpoint:

**quay-webhook.yaml**

```

apiVersion: v1
kind: Service
metadata:
  labels:
    name: quay-bridge-operator
    name: quay-bridge-operator
    namespace: openshift-operators
spec:
  ports:
    - name: https
      port: 443
      protocol: TCP
      targetPort: 8443
  selector:
    name: quay-bridge-operator
  sessionAffinity: None
  type: ClusterIP

```

5. Create the webhook service as follows:

```
$ oc create -f quay-webhook.yaml
```

6. Download the [webhook-create-signed-cert.sh](#) script, so you can use it to generate a certificate signed by a Kubernetes certificate authority.
7. Execute the following command to request the certificate:

```
$ ./webhook-create-signed-cert.sh --namespace openshift-operators \
  --secret quay-bridge-operator-webhook-certs \
  --service quay-bridge-operator
```

8. Execute the following command to retrieve the CA and format the result as a single line so that it can be entered into the MutatingWebhookConfiguration resource:

```
$ oc get configmap -n kube-system \
  extension-apiserver-authentication \
  -o=jsonpath='{.data.client-ca-file}' | base64 | tr -d '\n'
```

9. Replace the `#{CA_BUNDLE}` variable in the following MutatingWebhookConfiguration YAML:

**quay-mutating-webhook.yaml**

```

apiVersion: admissionregistration.k8s.io/v1
kind: MutatingWebhookConfiguration
metadata:
  name: quay-bridge-operator
webhooks:
  - name: quayintegration.redhatcop.redhat.io
    clientConfig:
      service:
        namespace: openshift-operators
        name: quay-bridge-operator

```

```

    path: "/admissionwebhook"
    caBundle: "${CA_BUNDLE}" ❶
  rules:
  - operations: [ "CREATE" ]
    apiGroups: [ "build.openshift.io" ]
    apiVersions: ["v1" ]
    resources: [ "builds" ]
  failurePolicy: Fail
  matchPolicy: Exact
  timeoutSeconds: 30
  sideEffects: None
  admissionReviewVersions: [v1beta1]

```

- ❶ Replace `${CA_BUNDLE}` with the output of the previous step. It will appear as one long line that you copy and paste to replace `${CA_BUNDLE}`.

10. Create the `MutatingWebhookConfiguration` as follows:

```
$ oc create -f quay-mutating-webhook.yaml
```

Until the operator is running, new requests for builds will fail since the webserver the `MutatingWebhookConfiguration` invokes is not available and a proper response is required in order for the object to be persisted in etcd.

11. Go to the OpenShift console and install the Quay Bridge Operator as follows:
- Select OperatorHub and search for Quay Bridge Operator.
  - Select Install
  - Choose Installation Mode (all namespaces), Update Channel, and Approval Strategy (Automatic or Manual).
  - Select Subscribe
12. Create the custom resource (CR) called **QuayIntegration**. For example:

### quay-integration.yaml

```

apiVersion: redhatcop.redhat.io/v1alpha1
kind: QuayIntegration
metadata:
  name: example-quayintegration
spec:
  clusterID: openshift ❶
  credentialsSecretName: openshift-operators/quay-integration ❷
  quayHostname: https://<QUAY_URL> ❸
  whitelistNamespaces: ❹
  - default
  insecureRegistry: false ❺

```

- ❶ The `clusterID` value should be unique across the entire ecosystem. This value is optional and defaults to `openshift`.

- 2 For `credentialsSecretName`, replace **openshift-operators/quay-integration** with the name of the namespace and the secret containing the token you created earlier.
- 3 Replace `QUAY_URL` with the hostname of your Red Hat Quay instance.
- 4 The `whitelistNamespaces` is optional. If not used, the Bridge Operator will sync all namespaces to Red Hat Quay except the `openshift` prefixed project. In this example, the white listed namespace (default) will now have an associated Red Hat Quay organization. Use any namespace you like here.
- 5 If Quay is using self signed certificates, set the property **`insecureRegistry: true`**.

The result is that organizations within Red Hat Quay should be created for the related namespaces in OpenShift.

13. Create the **QuayIntegration** as follows:

```
┆ $ oc create -f quay-integration.yaml
```

At this point a Quay integration resource is created, linking the OpenShift cluster to the Red Hat Quay instance.

The whitelisted namespace you created should now have a Red Hat Quay organization. If you were to use a command such as **`oc new-app`** to create a new application in that namespace, you would see a new Red Hat Quay repository created for it instead of using the internal registry.

## CHAPTER 10. REPOSITORY MIRRORING

### 10.1. REPOSITORY MIRRORING

Red Hat Quay repository mirroring lets you mirror images from external container registries (or another local registry) into your Red Hat Quay cluster. Using repository mirroring, you can synchronize images to Red Hat Quay based on repository names and tags.

From your Red Hat Quay cluster with repository mirroring enabled, you can:

- Choose a repository from an external registry to mirror
- Add credentials to access the external registry
- Identify specific container image repository names and tags to sync
- Set intervals at which a repository is synced
- Check the current state of synchronization

To use the mirroring functionality, you need to:

- Enable Repository Mirroring in the Red Hat Quay configuration
- Run a repository mirroring worker
- Create mirrored repositories

All repository mirroring configuration can be performed using the configuration tool UI or via the Quay API

### 10.2. REPOSITORY MIRRORING VERSUS GEO-REPLICATION

Quay geo-replication mirrors the entire image storage backend data between 2 or more different storage backends while the database is shared (one Quay registry with two different blob storage endpoints). The primary use cases for geo-replication are:

- Speeding up access to the binary blobs for geographically dispersed setups
- Guaranteeing that the image content is the same across regions

Repository mirroring synchronizes selected repositories (or subsets of repositories) from one registry to another. The registries are distinct, with registry is separate database and image storage. The primary use cases for mirroring are:

- Independent registry deployments in different datacenters or regions, where a certain subset of the overall content is supposed to be shared across the datacenters / regions
- Automatic synchronization or mirroring of selected (whitelisted) upstream repositories from external registries into a local Quay deployment



#### NOTE

Repository mirroring and geo-replication can be used simultaneously.

**Table 10.1. Red Hat Quay Repository mirroring versus geo-replication**

Feature / Capability	Geo-replication	Repository mirroring
What is the feature designed to do?	A shared, global registry	Distinct, different registries
What happens if replication or mirroring hasn't been completed yet?	The remote copy is used (slower)	No image is served
Is access to all storage backends in both regions required?	Yes (all Red Hat Quay nodes)	No (distinct storage)
Can users push images from both sites to the same repository?	Yes	No
Is all registry content and configuration identical across all regions (shared database)	Yes	No
Can users select individual namespaces or repositories to be mirrored?	No, by default	Yes
Can users apply filters to synchronization rules?	No	Yes

### 10.3. USING REPOSITORY MIRRORING

Here are some features and limitations of Red Hat Quay repository mirroring:

- With repository mirroring, you can mirror an entire repository or selectively limit which images are synced. Filters can be based on a comma-separated list of tags, a range of tags, or other means of identifying tags through regular expressions.
- Once a repository is set as mirrored, you cannot manually add other images to that repository.
- Because the mirrored repository is based on the repository and tags you set, it will hold only the content represented by the repo/tag pair. In other words, if you change the tag so that some images in the repository no longer match, those images will be deleted.
- Only the designated robot can push images to a mirrored repository, superseding any role-based access control permissions set on the repository.
- With a mirrored repository, a user can pull images (given read permission) from the repository but not push images to the repository.
- Changing settings on your mirrored repository is done from the Mirrors tab on the Repositories page for the mirrored repository you create.
- Images are synced at set intervals, but can also be synced on demand.




## 10.4. MIRRORING CONFIGURATION UI

1. Start the **Quay** container in configuration mode and select the Enable Repository Mirroring check box. If you want to require HTTPS communications and verify certificates during mirroring, select the HTTPS and cert verification check box.

### Repository Mirroring

If enabled, scheduled mirroring of repositories from remote registries will be available.

Enable Repository Mirroring

 A repository mirror service must be running to use this feature. Documentation on setting up and running this service can be found at [Running Repository Mirroring Service](#).

Require HTTPS and verify certificates of Quay registry during mirror.

2. Validate and download the **configuration** file, and then restart Quay in registry mode using the updated config file.

## 10.5. MIRRORING CONFIGURATION FIELDS

Table 10.2. Mirroring configuration

Field	Type	Description
FEATURE_REPO_MIRROR	Boolean	Enable or disable repository mirroring  <b>Default: false</b>
REPO_MIRROR_INTERVAL	Number	The number of seconds between checking for repository mirror candidates  <b>Default: 30</b>
REPO_MIRROR_SERVER_HOSTNAME	String	Replaces the <b>SERVER_HOSTNAME</b> as the destination for mirroring.  <b>Default: None</b>  <b>Example:</b> <b>openshift-quay-service</b>
REPO_MIRROR_TLS_VERIFY	Boolean	Require HTTPS and verify certificates of Quay registry during mirror.  <b>Default: false</b>

## 10.6. MIRRORING WORKER

- To run the repository mirroring worker, start by running a **Quay** pod with the **repomirror** option:

```
$ sudo podman run -d --name mirroring-worker \
-v $QUAY/config:/conf/stack:Z \
registry.redhat.io/quay/quay-rhel8:v3.5.7 repomirror
```

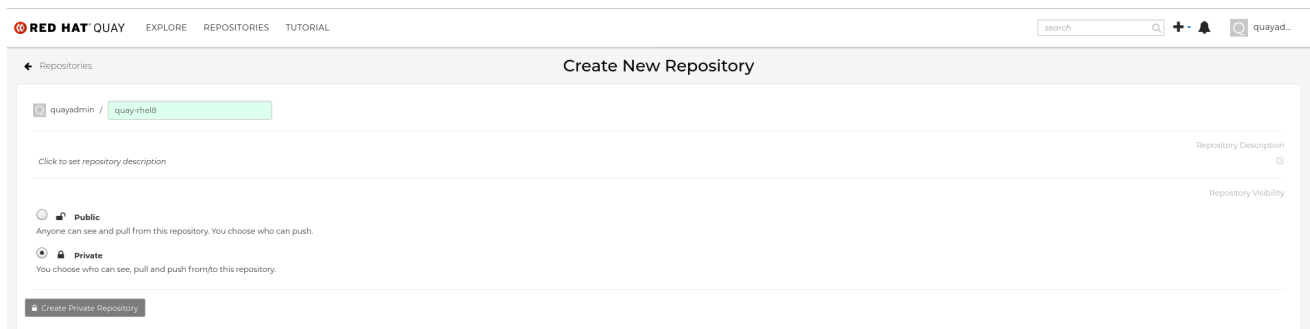
- If you have configured TLS communications using a certificate **/root/ca.crt**, then the following example shows how to start the mirroring worker:

```
$ sudo podman run -d --name mirroring-worker \
-v $QUAY/config:/conf/stack:Z \
-v /root/ca.crt:/etc/pki/ca-trust/source/anchors/ca.crt \
registry.redhat.io/quay/quay-rhel8:v3.5.7 repomirror
```

## 10.7. CREATING A MIRRORED REPOSITORY

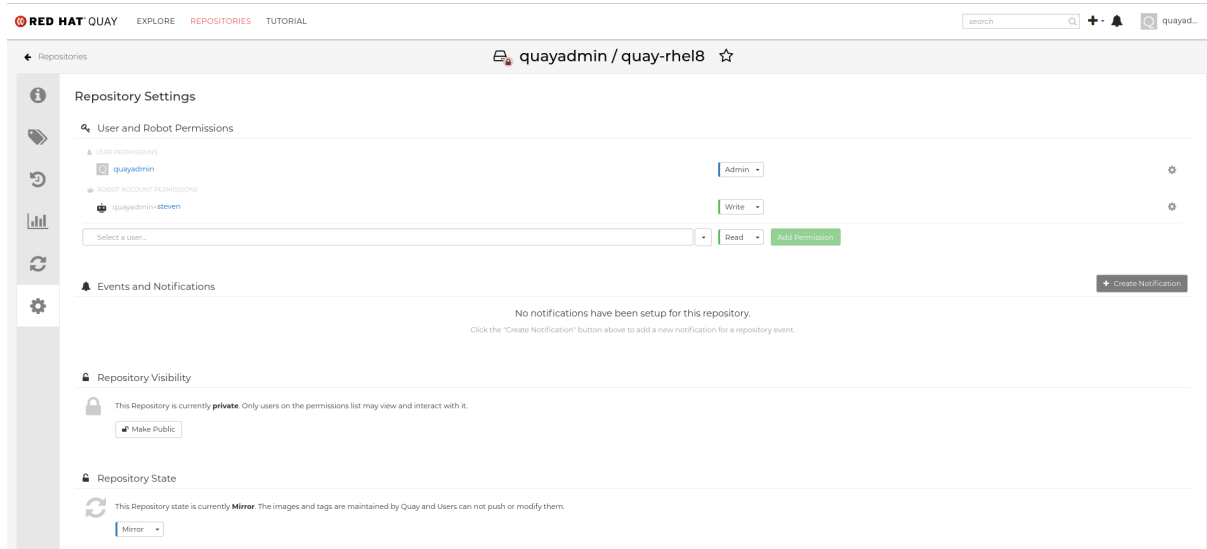
The steps shown in this section assume you already have enabled repository mirroring in the configuration for your Red Hat Quay cluster and that you have a deployed a mirroring worker.

When mirroring a repository from an external container registry, create a new private repository. Typically the same name is used as the target repository, for example, **quay-rhel8**:



### 10.7.1. Repository mirroring settings

- In the Settings tab, set the Repository State to **Mirror**:



- In the Mirror tab, enter the details for connecting to the external registry, along with the tags, scheduling and access information:

The screenshot shows the 'Repository Mirroring' configuration page in the Quay web interface. The page is titled 'Repository Mirroring' and includes a warning: 'This feature will convert quayadmin/quay-rhel8 into a mirror. Changes to the external repository will be duplicated here. While enabled, users will be unable to push images to this repository.' The configuration is divided into several sections:

- External Repository:**
  - Registry Location:** quay.io/redhat/quay
  - Tags:** comma-separated list of tag patterns to synchronize (Example: latest, 3.2)
  - Start Date:** May 27, 2021 4:17 PM
  - Sync Interval:** seconds
  - Robot User:** Select a user
- Credentials:** (Required if the external repository is private)
  - Username:**
  - Password:**
- Advanced Settings:**
  - Verify TLS:** Requires HTTPS and verify certificates when talking to the external registry (checkbox)
  - HTTP Proxy:** proxy.example.com
  - HTTPs Proxy:** proxy.example.com
  - No Proxy:** example.com

- Enter the details as required in the following fields:

- Registry Location:** The external repository you want to mirror, for example, **registry.redhat.io/quay/quay-rhel8**
- Tags:** This field is required. You may enter a comma-separated list of individual tags or tag patterns. (See *Tag Patterns* section for details.)



#### NOTE

In order for Quay to get the list of tags in the remote repository, one of the following requirements must be met:

- An image with the "latest" tag must exist in the remote repository *OR*
  - At least one explicit tag, without pattern matching, must exist in the list of tags that you specify
- Start Date:** The date on which mirroring begins. The current date and time is used by default.
  - Sync Interval:** Defaults to syncing every 24 hours. You can change that based on hours or days.
  - Robot User:** Create a new robot account or choose an existing robot account to do the mirroring.
  - Username:** The username for accessing the external registry holding the repository you are mirroring.
  - Password:** The password associated with the Username. Note that the password cannot include characters that require an escape character (\).

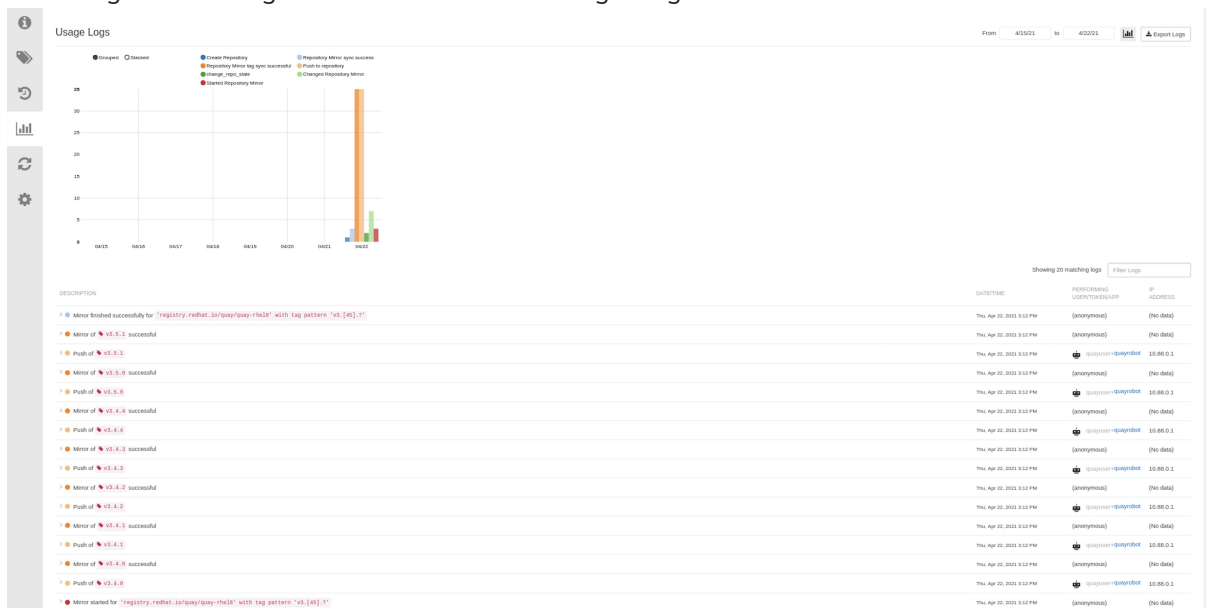
### 10.7.2. Advanced settings

- In the Advanced Settings section, configure TLS and proxy, if required:

- **Verify TLS:** Check this box if you want to require HTTPS and to verify certificates, when communicating with the target remote registry.
- **HTTP Proxy:** Identify the HTTP proxy server needed to access the remote site, if one is required.
- **HTTPS Proxy:** Identify the HTTPS proxy server needed to access the remote site, if one is required.
- **No Proxy:** List of locations that do not require proxy

### 10.7.3. Synchronize now


- To perform an immediate mirroring operation, press the Sync Now button on the repository's Mirroring tab. The logs are available on the Usage Logs tab:



When the mirroring is complete, the images will appear in the Tags tab:

The screenshot shows the 'Repository Tags' page for the repository 'quayuser / quay-rhel8'. It features a table with the following columns: TAG, LAST MODIFIED, SIZE, EXPIRES, and MANIFEST. The table lists several tags from v3.5.1 to v3.4.0. Each row includes a 'Last Modified' timestamp (e.g., 'a minute ago'), a 'Size' (e.g., 'N/A'), an 'Expires' status (e.g., 'Never'), and a 'Manifest' link (e.g., 'SHA256: 3d4d97d375'). The interface also includes a search bar and pagination controls.

Below is an example of a completed Repository Mirroring screen:



## 10.8. EVENT NOTIFICATIONS FOR MIRRORING

There are three notification events for repository mirroring:

- Repository Mirror Started
- Repository Mirror Success
- Repository Mirror Unsuccessful

The events can be configured inside the Settings tab for each repository, and all existing notification methods such as email, slack, Quay UI and webhooks are supported.

## 10.9. MIRRORING TAG PATTERNS

As noted above, at least one Tag must be explicitly entered (ie. not a tag pattern) or the tag "latest" must exist in the report repository. (The tag "latest" will not be synced unless specified in the tag list.). This is required for Quay to get the list of tags in the remote repository to compare to the specified list to mirror.

### 10.9.1. Pattern syntax

Pattern	Description
*	Matches all characters
?	Matches any single character
[seq]	Matches any character in seq
[!seq]	Matches any character not in seq

### 10.9.2. Example tag patterns

Example Pattern	Example Matches
v3*	v32, v3.1, v3.2, v3.2-4beta, v3.3
v3.*	v3.1, v3.2, v3.2-4beta
v3.?	v3.1, v3.2, v3.3
v3.[12]	v3.1, v3.2
v3.[12]*	v3.1, v3.2, v3.2-4beta
v3.[!1]*	v3.2, v3.2-4beta, v3.3

## 10.10. WORKING WITH MIRRORED REPOSITORIES

Once you have created a mirrored repository, there are several ways you can work with that repository. Select your mirrored repository from the Repositories page and do any of the following:

- **Enable/disable the repository.** Select the Mirroring button in the left column, then toggle the Enabled check box to enable or disable the repository temporarily.
- **Check mirror logs** To make sure the mirrored repository is working properly, you can check the mirror logs. To do that, select the Usage Logs button in the left column. Here's an example:

← Repositories
johnjones / ubi7repo ☆

📄  
🔄  
📊  
🔄  
⚙️

### Usage Logs

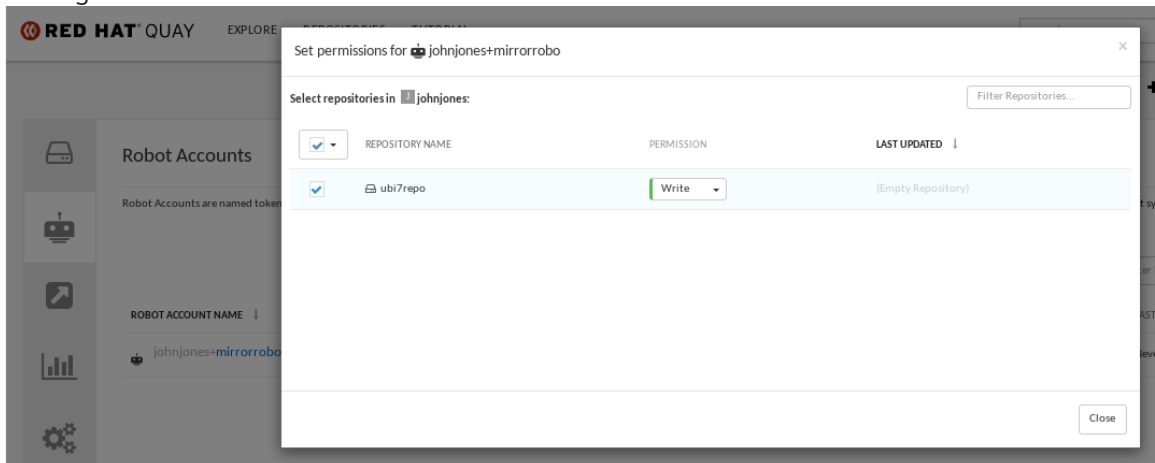
From  to 
📊
📄 Export Logs

● Grouped ○ Stacked
● Changed Repository Mirror
● Repository Mirror sync success
● Create Repository
● Started Repository Mirror


Showing 6 matching logs

DESCRIPTION	DATE/TIME	PERFORMING USER/TOKEN/APP	IP ADDRESS
● Mirror finished successfully for 'registry.access.redhat.com/ubi7/ubi-minimal' with tag pattern 'latest'	Fri, Aug 16, 2019 10:49 AM	(anonymous)	(No data)
● Mirror started for 'registry.access.redhat.com/ubi7/ubi-minimal' with tag pattern 'latest'	Fri, Aug 16, 2019 10:49 AM	(anonymous)	(No data)
● Immediate mirror scheduled	Fri, Aug 16, 2019 10:48 AM	quay	(No data)
● Mirror finished successfully for 'registry.access.redhat.com/ubi7/ubi-minimal' with tag pattern 'latest'	Fri, Aug 16, 2019 10:18 AM	(anonymous)	(No data)
● Mirror started for 'registry.access.redhat.com/ubi7/ubi-minimal' with tag pattern 'latest'	Fri, Aug 16, 2019 10:18 AM	(anonymous)	(No data)
● CreateRepository quay / ubi7minimal	Fri, Aug 16, 2019 10:01 AM	quay	(No data)

- **Sync mirror now:** To immediately sync the images in your repository, select the Sync Now button.
- **Change credentials:** To change the username and password, select DELETE from the Credentials line. Then select None and add the username and password needed to log into the external registry when prompted.
- **Cancel mirroring:** To stop mirroring, which keeps the current images available but stops new ones from being synced, select the CANCEL button.
- **Set robot permissions:** Red Hat Quay robot accounts are named tokens that hold credentials for accessing external repositories. By assigning credentials to a robot, that robot can be used across multiple mirrored repositories that need to access the same external registry. You can assign an existing robot to a repository by going to Account Settings, then selecting the Robot Accounts icon in the left column. For the robot account, choose the link under the REPOSITORIES column. From the pop-up window, you can:
  - Check which repositories are assigned to that robot.
  - Assign read, write or Admin privileges to that robot from the PERMISSION field shown in this figure:



- **Change robot credentials:** Robots can hold credentials such as Kubernetes secrets, Docker login information, and Mesos bundles. To change robot credentials, select the Options gear on the robot's account line on the Robot Accounts window and choose View Credentials. Add the appropriate credentials for the external repository the robot needs to access.



- **Check and change general setting:** Select the Settings button (gear icon) from the left column on the mirrored repository page. On the resulting page, you can change settings associated with the mirrored repository. In particular, you can change User and Robot Permissions, to specify exactly which users and robots can read from or write to the repo.

## 10.11. REPOSITORY MIRRORING RECOMMENDATIONS

- Repository mirroring pods can run on any node including other nodes where Quay is already running
- Repository mirroring is scheduled in the database and run in batches. As a result, more workers could mean faster mirroring, since more batches will be processed.
- The optimal number of mirroring pods depends on:
  - The total number of repositories to be mirrored
  - The number of images and tags in the repositories and the frequency of changes
  - Parallel batches
- You should balance your mirroring schedule across all mirrored repositories, so that they do not all start up at the same time.
- For a mid-size deployment, with approximately 1000 users and 1000 repositories, and with roughly 100 mirrored repositories, it is expected that you would use 3-5 mirroring pods, scaling up to 10 if required.



## CHAPTER 11. BACKING UP AND RESTORING RED HAT QUAY ON AN OPENSIFT CONTAINER PLATFORM DEPLOYMENT

Use the content within this section to back up and restore Red Hat Quay on an OpenShift Container Platform deployment.

### 11.1. BACKING UP RED HAT QUAY

This procedure is exclusively for OpenShift Container Platform and NooBaa deployments.

#### Prerequisites

- A Red Hat Quay deployment on OpenShift Container Platform.

#### Procedure

1. Backup the **QuayRegistry** custom resource by exporting it:

```
$ oc get quayregistry <quay-registry-name> -n <quay-namespace> -o yaml > quay-registry.yaml
```

2. Edit the resulting **quayregistry.yaml** and remove the status section and the following metadata fields:

```
metadata.creationTimestamp
metadata.finalizers
metadata.generation
metadata.resourceVersion
metadata.uid
```

3. Backup the managed keys secret:



#### NOTE

If you are running a version older than Red Hat Quay 3.7.0, this step can be skipped. Some secrets are automatically generated while deploying Quay for the first time. These are stored in a secret called **<quay-registry-name>-quay-registry-managed-secret-keys** in the QuayRegistry namespace.

```
$ oc get secret -n <quay-namespace> <quay-registry-name>-quay-registry-managed-secret-keys -o yaml > managed-secret-keys.yaml
```

4. Edit the the resulting **managed-secret-keys.yaml** file and remove all owner references. Your **managed-secret-keys.yaml** file should look similar to the following:

```
apiVersion: v1
kind: Secret
type: Opaque
metadata:
  name: <quayname>-quay-registry-managed-secret-keys
  namespace: <quay-namespace>
data:
```

```
CONFIG_EDITOR_PW: <redacted>
DATABASE_SECRET_KEY: <redacted>
DB_ROOT_PW: <redacted>
DB_URI: <redacted>
SECRET_KEY: <redacted>
SECURITY_SCANNER_V4_PSK: <redacted>
```

All information under the **data** property should remain the same.

5. Backup the current Quay configuration:

```
$ oc get secret -n <quay-namespace> $(oc get quayregistry <quay-registry-name> -n
<quay-namespace> -o jsonpath='{.spec.configBundleSecret}') -o yaml > config-bundle.yaml
```

6. Backup the **/conf/stack/config.yaml** file mounted inside of the Quay pods:

```
$ oc exec -it quay-pod-name -- cat /conf/stack/config.yaml > quay-config.yaml
```

7. Scale down the Quay the Quay Operator:

```
$ oc scale --replicas=0 deployment $(oc get deployment -n <quay-operator-namespace>
|awk '/^quay-operator/ {print $1}') -n <quay-operator-namespace>
```

8. Scale down the Quay namespace:

```
$ oc scale --replicas=0 deployment $(oc get deployment -n <quay-namespace> -l quay-
component=quay -o jsonpath='{.items[0].metadata.name}') -n <quay-namespace>
```

9. Wait for the **registry-quay-app** pods to disappear. You can check their status by running the following command:

```
$ oc get pods -n <quay-namespace>
```

Example output:

```
registry-quay-config-editor-77847fc4f5-nsbbv 1/1 Running 0 9m1s
registry-quay-database-66969cd859-n2ssm 1/1 Running 0 6d1h
registry-quay-mirror-758fc68ff7-5wxlp 1/1 Running 0 8m29s
registry-quay-mirror-758fc68ff7-lbl82 1/1 Running 0 8m29s
registry-quay-redis-7cc5f6c977-956g8 1/1 Running 0 5d21h
```

10. Identify the Quay PostgreSQL pod name:

```
$ oc get pod -l quay-component=postgres -n <quay-namespace> -o
jsonpath='{.items[0].metadata.name}'
```

Example output:

```
quayregistry-quay-database-59f54bb7-58xs7
```

1. Obtain the Quay database name:

```
$ oc -n <quay-namespace> rsh $(oc get pod -l app=quay -o NAME -n <quay-namespace>
|head -n 1) cat /conf/stack/config.yaml|awk -F"/" '/^DB_URI/ {print $4}'
quayregistry-quay-database
```

2. Download a backup database:

```
$ oc exec quayregistry-quay-database-59f54bb7-58xs7 -- /usr/bin/pg_dump -C quayregistry-
quay-database > backup.sql
```

3. Decode and export the **AWS\_ACCESS\_KEY\_ID**:

```
$ export AWS_ACCESS_KEY_ID=$(oc get secret -l app=noobaa -n <quay-namespace> -o
jsonpath='{.items[0].data.AWS_ACCESS_KEY_ID}' |base64 -d)
```


4. Decode and export the **AWS\_SECRET\_ACCESS\_KEY\_ID**:

```
$ export AWS_SECRET_ACCESS_KEY=$(oc get secret -l app=noobaa -n <quay-
namespace> -o jsonpath='{.items[0].data.AWS_SECRET_ACCESS_KEY}' |base64 -d)
```

5. Create a new directory and copy all blobs to it:

```
$ mkdir blobs
```

```
$ aws s3 sync --no-verify-ssl --endpoint https://$(oc get route s3 -n openshift-storage -o
jsonpath='{.spec.host}') s3://$(oc get cm -l app=noobaa -n <quay-namespace> -o
jsonpath='{.items[0].data.BUCKET_NAME}') ./blobs
```



#### NOTE

You can also use [rclone](#) or [sc3md](#) instead of the AWS command line utility.

1. Scale up the Quay the Quay Operator:

```
$ oc scale --replicas=1 deployment $(oc get deployment -n <quay-operator-namespace>
|awk '/^quay-operator/ {print $1}') -n <quay-operator-namespace>
```

2. Scale up the Quay namespace:

```
$ oc scale --replicas=1 deployment $(oc get deployment -n <quay-namespace> -l quay-
component=quay -o jsonpath='{.items[0].metadata.name}') -n <quay-namespace>
```

3. Check the status of the Operator:

```
$ oc get quayregistry <quay-registry-name> -n <quay-namespace> -o yaml
```

Example output:

```
apiVersion: quay.redhat.com/v1
kind: QuayRegistry
metadata:
  ...
```

```

name: example-registry
namespace: <quay-namespace>
...
spec:
  components:
  - kind: quay
    managed: true
  ...
  - kind: clairpostgres
    managed: true
  configBundleSecret: init-config-bundle-secret
status:
  configEditorCredentialsSecret: example-registry-quay-config-editor-credentials-fg2gdgtm24
  configEditorEndpoint: https://example-registry-quay-config-editor-quay-
enterprise.apps.docs.gcp.quaydev.org
  currentVersion: 3.7.0
  lastUpdated: 2022-05-11 13:28:38.199476938 +0000 UTC
  registryEndpoint: https://example-registry-quay-quay-enterprise.apps.docs.gcp.quaydev.org
0      5d21h

```

## 11.2. RESTORING RED HAT QUAY

This procedure is used to restore Red Hat Quay when the Red Hat Quay Operator manages the database. It should be performed after a backup of your Quay registry has been performed.

### Prerequisites


- Red Hat Quay is deployed on OpenShift Container Platform using the Quay Operator.
- Your Red Hat Quay database has been backed up.

### Procedure

1. Restore the backed up Quay configuration and the randomly generated keys:

```
$ oc create -f ./config-bundle.yaml
```

```
$ oc create -f ./managed-secret-keys.yaml
```



### NOTE

If you receive the error **Error from server (AlreadyExists): error when creating "/>**

2. Restore the QuayRegistry custom resource:

```
$ oc create -f ./quay-registry.yaml
```

3. Scale down the Quay the Quay Operator:

```
$ oc scale --replicas=0 deployment $(oc get deployment -n <quay-operator-namespace>
|awk '/^quay-operator/ {print $1}') -n <quay-operator-namespace>
```

- Scale down the Quay namespace:

```
$ oc scale --replicas=0 deployment $(oc get deployment -n <quay-namespace> -l quay-
component=quay -o jsonpath='{.items[0].metadata.name}') -n <quay-namespace>
```

- Identify your Quay database pod:

```
$ oc get pod -l quay-component=postgres -n <quay-namespace> -o
jsonpath='{.items[0].metadata.name}'
```

Example output:

```
quayregistry-quay-database-59f54bb7-58xs7
```

- Upload the backup by copying it from the local environment and into the pod:

```
$ oc cp ./backup.sql -n <quay-namespace> registry-quay-database-66969cd859-
n2ssm:/tmp/backup.sql
```

- Open a remote terminal to the database:

```
$ oc rsh -n <quay-namespace> registry-quay-database-66969cd859-n2ssm
```

- Enter psql:

```
bash-4.4$ psql
```

- You can list the database by running the following command:

```
postgres=# \l
```

Example output:

```

                                List of databases
   Name          | Owner          | Encoding | Collate  | Ctype    | Access
privileges
-----+-----+-----+-----+-----+-----
 postgres       | postgres      | UTF8     | en_US.utf8 | en_US.utf8 |
quayregistry-quay-database | quayregistry-quay-database | UTF8     | en_US.utf8 | en_US.utf8 |
en_US.utf8 |
```

- Drop the database:

```
postgres=# DROP DATABASE "quayregistry-quay-database";
```

Example output:

```
DROP DATABASE
```

- 11. Exit the postgres CLI to re-enter bash-4.4:

```
\q
```

- 12. Redirect your PostgreSQL database to your backup database:

```
sh-4.4$ psql < /tmp/backup.sql
```

- 13. Exit bash:

```
sh-4.4$ exit
```

- 14. Export the **AWS\_ACCESS\_KEY\_ID**:

```
$ export AWS_ACCESS_KEY_ID=$(oc get secret -l app=noobaa -n <quay-namespace> -o jsonpath='{.items[0].data.AWS_ACCESS_KEY_ID}' |base64 -d)
```

- 15. Export the **AWS\_SECRET\_ACCESS\_KEY**:

```
$ export AWS_SECRET_ACCESS_KEY=$(oc get secret -l app=noobaa -n <quay-namespace> -o jsonpath='{.items[0].data.AWS_SECRET_ACCESS_KEY}' |base64 -d)
```

- 16. Upload all blobs to the bucket by running the following command:

```
$ aws s3 sync --no-verify-ssl --endpoint https://$(oc get route s3 -n openshift-storage -o jsonpath='{.spec.host}') ./blobs s3://$(oc get cm -l app=noobaa -n <quay-namespace> -o jsonpath='{.items[0].data.BUCKET_NAME}')
```

- 17. Scale up the Quay the Quay Operator:

```
$ oc scale --replicas=1 deployment $(oc get deployment -n <quay-operator-namespace> |awk '/^quay-operator/ {print $1}') -n <quay-operator-namespace>
```

- 18. Scale up the Quay namespace:

```
$ oc scale --replicas=1 deployment $(oc get deployment -n <quay-namespace> -l quay-component=quay -o jsonpath='{.items[0].metadata.name}') -n <quay-namespace>
```

- 19. Check the status of the Operator and ensure it has come back online:

```
$ oc get quayregistry -n <quay-namespace> <registry-name> -o yaml
```

Example output:

```
apiVersion: quay.redhat.com/v1
kind: QuayRegistry
metadata:
  ...
  name: example-registry
  namespace: quay-enterprise
  ...
```

```
spec:
  components:
  - kind: quay
    managed: true
  ...
  - kind: clairpostgres
    managed: true
  configBundleSecret: init-config-bundle-secret
status:
  configEditorCredentialsSecret: example-registry-quay-config-editor-credentials-fg2gdgtm24
  configEditorEndpoint: https://example-registry-quay-config-editor-quay-
enterprise.apps.docs.gcp.quaydev.org
  currentVersion: 3.7.0
  lastUpdated: 2022-05-11 13:28:38.199476938 +0000 UTC
  registryEndpoint: https://example-registry-quay-quay-enterprise.apps.docs.gcp.quaydev.org
    0      5d21h
```

## CHAPTER 12. LDAP AUTHENTICATION SETUP FOR RED HAT QUAY

The Lightweight Directory Access Protocol (LDAP) is an open, vendor-neutral, industry standard application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network. Red Hat Quay supports using LDAP as an identity provider.

### 12.1. CONSIDERATIONS PRIOR TO ENABLING LDAP

#### 12.1.1. Existing Quay deployments

Conflicts between user names can arise when you enable LDAP for an existing Quay deployment that already has users configured. Consider the scenario where a particular user, **alice**, was manually created in Quay prior to enabling LDAP. If the user name **alice** also exists in the LDAP directory, Quay will create a new user **alice-1** when **alice** logs in for the first time using LDAP, and will map the LDAP credentials to this account. This might not be what you want, for consistency reasons, and it is recommended that you remove any potentially conflicting local account names from Quay prior to enabling LDAP.

#### 12.1.2. Manual User Creation and LDAP authentication

When Quay is configured for LDAP, LDAP-authenticated users are automatically created in Quay's database on first log in, if the configuration option **FEATURE\_USER\_CREATION** is set to **true**. If this option is set to **false**, the automatic user creation for LDAP users will fail and the user is not allowed to log in. In this scenario, the superuser needs to create the desired user account first. Conversely, if **FEATURE\_USER\_CREATION** is set to **true**, this also means that a user can still create an account from the Quay login screen, even if there is an equivalent user in LDAP.

### 12.2. SET UP LDAP CONFIGURATION

In the config tool, locate the Authentication section and select "LDAP" from the drop-down menu. Update LDAP configuration fields as required.

#### Internal Authentication

Authentication for the registry can be handled by either the registry itself, LDAP, Keystone, or external JWT endpoint.

Additional **external** authentication providers (such as GitHub) can be used in addition for **login into the UI**.

 It is **highly recommended** to require encrypted client passwords. External passwords used in the Docker client will be stored in **plaintext!** [Enable this requirement now.](#)

Authentication:

LDAP

- Here is an example of the resulting entry in the *config.yaml* file:

```
AUTHENTICATION_TYPE: LDAP
```

#### 12.2.1. Full LDAP URI



**LDAP URI:**   
 The full LDAP URI, including the `ldap://` or `ldaps://` prefix.

**Custom TLS Certificate:** Please select a file to upload as **ldap.crt**:  No file chosen  
 If specified, the certificate (in PEM format) for the LDAP TLS connection.

**Allow insecure:**  **Allow fallback to non-TLS connections**  
 If enabled, LDAP will fallback to insecure non-TLS connections if TLS does not succeed.

- The full LDAP URI, including the `ldap://` or `ldaps://` prefix.
- A URI beginning with `ldaps://` will make use of the provided SSL certificate(s) for TLS setup.
- Here is an example of the resulting entry in the `config.yaml` file:

```
LDAP_URI: ldaps://ldap.example.org
```


## 12.2.2. Team Synchronization

**Team synchronization:**  **Enable Team Synchronization Support**  
 If enabled, organization administrators who are also superusers can set teams to have their membership synchronized with a backing group in LDAP.

- If enabled, organization administrators who are also superusers can set teams to have their membership synchronized with a backing group in LDAP.

**Team synchronization:**  **Enable Team Synchronization Support**  
 If enabled, organization administrators who are also superusers can set teams to have their membership synchronized with a backing group in LDAP.

**Resynchronization duration:**   
 The duration before a team must be re-synchronized. Must be expressed in a duration string form: `30m`, `1h`, `1d`.

**Self-service team syncing setup:**  If enabled, this feature will allow \*any organization administrator\* to read the membership of any LDAP group.

**Allow non-superusers to enable and manage team syncing**  
 If enabled, non-superusers will be able to enable and manage team syncing on teams under organizations in which they are administrators.

- The resynchronization duration is the period at which a team must be re-synchronized. Must be expressed in a duration string form: `30m`, `1h`, `1d`.
- Optionally allow non-superusers to enable and manage team syncing under organizations in which they are administrators.
- Here is an example of the resulting entries in the `config.yaml` file:

```
FEATURE_TEAM_SYNCING: true
TEAM_RESYNC_STALE_TIME: 60m
FEATURE_NONSUPERUSER_TEAM_SYNCING_SETUP: true
```

## 12.2.3. Base and Relative Distinguished Names

**Base DN:**

A Distinguished Name path which forms the base path for looking up all LDAP records.  
Example: dc=my,dc=domain,dc=com

**User Relative DN:**

A Distinguished Name path which forms the base path for looking up all user LDAP records, relative to the Base DN defined above.  
Example: ou=employees

**Secondary User Relative DNs:**

- ou=SFO [Remove](#)

A list of Distinguished Name path(s) which forms the secondary base path(s) for looking up all user LDAP records, relative to the Base DN defined above. These path(s) will be tried if the user is not found via the primary relative DN.  
Example: [ou=employees]

- A Distinguished Name path which forms the base path for looking up all LDAP records. Example: *dc=my,dc=domain,dc=com*
- Optional list of Distinguished Name path(s) which form the secondary base path(s) for looking up all user LDAP records, relative to the Base DN defined above. These path(s) will be tried if the user is not found via the primary relative DN.
- User Relative DN is relative to BaseDN. Example: *ou=NYC* not *ou=NYC,dc=example,dc=org*
- Multiple "Secondary User Relative DNs" may be entered if there are multiple Organizational Units where User objects are located at. Simply type in the Organizational Units and click on Add button to add multiple RDNs. Example: *ou=Users,ou=NYC* and *ou=Users,ou=SFO*
- The "User Relative DN" searches with subtree scope. For example, if your Organization has Organizational Units NYC and SFO under the Users OU (*ou=SFO,ou=Users* and *ou=NYC,ou=Users*), Red Hat Quay can authenticate users from both the NYC and SFO Organizational Units if the User Relative DN is set to Users (*ou=Users*).
- Here is an example of the resulting entries in the *config.yaml* file:

```
LDAP_BASE_DN:
- dc=example
- dc=com
LDAP_USER_RDN:
- ou=users
LDAP_SECONDARY_USER_RDNS:
- ou=bots
- ou=external
```

## 12.2.4. Additional User Filters

**Additional User Filter Expression:**

**NOTE:** This query is added **unescaped** to user lookups, so be VERY careful with the query you specify.

If specified, the additional filter used for all user lookup queries. Note that all Distinguished Names used in the filter must be **full paths**; the **base\_dn** is not added automatically here. **Must** be wrapped in parens.

Example: (someOtherField=someOtherValue)

Example: (memberOf=some.full.path.to.a.group)

Example: ((someFirstField=someValue)(someOtherField=someOtherValue))

Example: (&(someFirstField=someValue)(someOtherField=someOtherValue))

- If specified, the additional filter used for all user lookup queries. Note that all Distinguished Names used in the filter must be **full paths**; the Base DN is not added automatically here. **Must** be wrapped in parens. Example: (&(someFirstField=someValue))

(someOtherField=someOtherValue))

- Here is an example of the resulting entry in the *config.yaml* file:


```
LDAP_USER_FILTER: (memberof=cn=developers,ou=groups,dc=example,dc=com)
```

### 12.2.5. Administrator DN

**Administrator DN:**

The Distinguished Name for the Administrator account. This account must be able to login and view the records for all user accounts.  
Example: uid=admin,ou=employees,dc=my,dc=domain,dc=com

**Administrator DN Password:**

 Note: This will be stored in **plaintext** inside the config.yaml, so setting up a dedicated account or using a [password hash](#) is **highly** recommended.

The password for the Administrator DN.

- The Distinguished Name and password for the administrator account. This account must be able to login and view the records for all user accounts. Example:  
uid=admin,ou=employees,dc=my,dc=domain,dc=com
- The password will be stored in **plaintext** inside the config.yaml, so setting up a dedicated account or using a password hash is highly recommended.
- Here is an example of the resulting entries in the *config.yaml* file:

```
LDAP_ADMIN_DN: cn=admin,dc=example,dc=com
LDAP_ADMIN_PASSWD: changeme
```

### 12.2.6. UID and Mail attributes

**UID Attribute:**

The name of the property field in your LDAP user records that stores your users' username. Typically "uid".

**Mail Attribute:**

The name of the property field in your LDAP user records that stores your users' e-mail address(es). Typically "mail".


- The UID attribute is the name of the property field in LDAP user record to use as the **username**. Typically "uid".
- The Mail attribute is the name of the property field in LDAP user record that stores user e-mail address(es). Typically "mail".
- Either of these may be used during login.
- The logged in username must exist in User Relative DN.
- *sAMAccountName* is the UID attribute for against Microsoft Active Directory setups.
- Here is an example of the resulting entries in the *config.yaml* file:

```
LDAP_UID_ATTR: uid
LDAP_EMAIL_ATTR: mail
```

## 12.2.7. Validation

Once the configuration is completed, click on "Save Configuration Changes" button to validate the configuration.

### Validating configuration



✓ CONFIGURATION VALIDATED

✓ Configuration Validated

Continue Editing

Download

All validation must succeed before proceeding, or additional configuration may be performed by selecting the "Continue Editing" button.

## 12.3. COMMON ISSUES

### *Invalid credentials*

Administrator DN or Administrator DN Password values are incorrect

***Verification of superuser %USERNAME% failed: Username not found The user either does not exist in the remote authentication system OR LDAP auth is misconfigured.***

Red Hat Quay can connect to the LDAP server via Username/Password specified in the Administrator DN fields however cannot find the current logged in user with the UID Attribute or Mail Attribute fields in the User Relative DN Path. Either current logged in user does not exist in User Relative DN Path, or Administrator DN user do not have rights to search/read this LDAP path.

## 12.4. CONFIGURE AN LDAP USER AS SUPERUSER

Once LDAP is configured, you can log in to your Red Hat Quay instance with a valid LDAP username and password. You are prompted to confirm your Red Hat Quay username as shown in the following figure:

### Confirm Username

The username `testadmin` was automatically generated to conform to the Docker CLI guidelines for use as a namespace in .

Please confirm the selected username or enter a different username below:

testadmin

Confirm Username

✓ Username valid

To attach superuser privilege to an LDAP user, modify the `config.yaml` file with the username. For example:

SUPER\_USERS:

- testadmin

Restart the Red Hat **Quay** container with the updated config.yaml file. The next time you log in, the user will have superuser privileges.

## CHAPTER 13. PROMETHEUS AND GRAFANA METRICS UNDER RED HAT QUAY

Red Hat Quay exports a [Prometheus](#)- and Grafana-compatible endpoint on each instance to allow for easy monitoring and alerting.

### 13.1. EXPOSING THE PROMETHEUS ENDPOINT

The Prometheus- and Grafana-compatible endpoint on the Red Hat Quay instance can be found at port **9091**. See [Monitoring Quay with Prometheus and Grafana](#) for details on configuring Prometheus and Grafana to monitor Quay repository counts.

#### 13.1.1. Setting up Prometheus to consume metrics

Prometheus needs a way to access all Red Hat Quay instances running in a cluster. In the typical setup, this is done by listing all the Red Hat Quay instances in a single named DNS entry, which is then given to Prometheus.

#### 13.1.2. DNS configuration under Kubernetes


A simple [Kubernetes service](#) can be configured to provide the DNS entry for Prometheus. Details on running Prometheus under Kubernetes can be found at [Prometheus and Kubernetes](#) and [Monitoring Kubernetes with Prometheus](#).

#### 13.1.3. DNS configuration for a manual cluster

[SkyDNS](#) is a simple solution for managing this DNS record when not using Kubernetes. SkyDNS can run on an [etcd](#) cluster. Entries for each Red Hat Quay instance in the cluster can be added and removed in the etcd store. SkyDNS will regularly read them from there and update the list of Quay instances in the DNS record accordingly.

## CHAPTER 14. GEO-REPLICATION

Geo-replication allows multiple, geographically distributed Quay deployments to work as a single registry from the perspective of a client or user. It significantly improves push and pull performance in a globally-distributed Quay setup. Image data is asynchronously replicated in the background with transparent failover / redirect for clients.



### NOTE

Deploying Red Hat Quay with geo-replication on OpenShift is not supported by the Operator.

### 14.1. GEO-REPLICATION FEATURES


- When geo-replication is configured, container image pushes will be written to the preferred storage engine for that Red Hat Quay instance (typically the nearest storage backend within the region).
- After the initial push, image data will be replicated in the background to other storage engines.
- The list of replication locations is configurable and those can be different storage backends.
- An image pull will always use the closest available storage engine, to maximize pull performance.
- If replication hasn't been completed yet, the pull will use the source storage backend instead.

### 14.2. GEO-REPLICATION REQUIREMENTS AND CONSTRAINTS

- A single database, and therefore all metadata and Quay configuration, is shared across all regions.
- A single Redis cache is shared across the entire Quay setup and needs to be accessible by all Quay pods.
- The exact same configuration should be used across all regions, with exception of the storage backend, which can be configured explicitly using the **QUAY\_DISTRIBUTED\_STORAGE\_PREFERENCE** environment variable.
- Geo-Replication requires object storage in each region. It does not work with local storage or NFS.
- Each region must be able to access every storage engine in each region (requires a network path).
- Alternatively, the storage proxy option can be used.
- The entire storage backend (all blobs) is replicated. This is in contrast to repository mirroring, which can be limited to an organization or repository or image.
- All Quay instances must share the same endpoint, typically via load balancer.
- All Quay instances must have the same set of superusers, as they are defined inside the common configuration file.

If the above requirements cannot be met, you should instead use two or more distinct Quay deployments and take advantage of repository mirroring functionality.

## 14.3. GEO-REPLICATION ARCHITECTURE



178\_Quay\_0821

In the example shown above, Quay is running in two separate regions, with a common database and a common Redis instance. Localized image storage is provided in each region and image pulls are served from the closest available storage engine. Container image pushes are written to the preferred storage engine for the Quay instance, and will then be replicated, in the background, to the other storage engines.

## 14.4. ENABLE STORAGE REPLICATION

1. Scroll down to the section entitled **Registry Storage**.
2. Click **Enable Storage Replication**.



3. Add each of the storage engines to which data will be replicated. All storage engines to be used must be listed.
4. If complete replication of all images to all storage engines is required, under each storage engine configuration click **Replicate to storage engine by default**. This will ensure that all images are replicated to that storage engine. To instead enable per-namespace replication, please contact support.
5. When you are done, click **Save Configuration Changes**. Configuration changes will take effect the next time Red Hat Quay restarts.
6. After adding storage and enabling “Replicate to storage engine by default” for Georeplications, you need to sync existing image data across all storage. To do this, you need to **oc exec** (or **docker/kubect exec**) into the container and run:

```
# scl enable python27 bash
# python -m util.backfillreplication
```

This is a one time operation to sync content after adding new storage.

#### 14.4.1. Run Red Hat Quay with storage preferences

1. Copy the config.yaml to all machines running Red Hat Quay
2. For each machine in each region, add a **QUAY\_DISTRIBUTED\_STORAGE\_PREFERENCE** environment variable with the preferred storage engine for the region in which the machine is running.

For example, for a machine running in Europe with the config directory on the host available from **\$QUAY/config**:

```
$ sudo podman run -d --rm -p 80:8080 -p 443:8443 \
  --name=quay \
  -v $QUAY/config:/conf/stack:Z \
  -e QUAY_DISTRIBUTED_STORAGE_PREFERENCE=europestorage \
  registry.redhat.io/quay/quay-rhel8:v3.5.7
```



#### NOTE

The value of the environment variable specified must match the name of a Location ID as defined in the config panel.

3. Restart all Red Hat Quay containers

## CHAPTER 15. RED HAT QUAY TROUBLESHOOTING

Common failure modes and best practices for recovery.

- [I'm receiving HTTP Status Code 429](#)
- [I'm authorized but I'm still getting 403s](#)
- [Base image pull in Dockerfile fails with 403](#)
- [Cannot add a build trigger](#)
- [Build logs are not loading](#)
- [I'm receiving "Cannot locate specified Dockerfile" \\* Could not reach any registry endpoint](#)
- [Cannot access private repositories using EC2 Container Service](#)
- [Docker is returning an i/o timeout](#)
- [Docker login is failing with an odd error](#)
- [Pulls are failing with an odd error](#)
- [I just pushed but the timestamp is wrong](#)
- [Pulling Private Quay.io images with Marathon/Mesos fails](#)

## CHAPTER 16. SCHEMA FOR RED HAT QUAY CONFIGURATION

Most Red Hat Quay configuration information is stored in the **config.yaml** file that is created using the browser-based config tool when Red Hat Quay is first deployed. This chapter describes the schema of those settings that are available to use in the `config.yaml` file.

The following fields required (all other are optional):

```
AUTHENTICATION_TYPE
BUILDLOGS_REDIS
DATABASE_SECRET_KEY
DB_URI
DEFAULT_TAG_EXPIRATION
DISTRIBUTED_STORAGE_CONFIG
DISTRIBUTED_STORAGE_PREFERENCE
PREFERRED_URL_SCHEME
SECRET_KEY
SERVER_HOSTNAME
TAG_EXPIRATION_OPTIONS
USER_EVENTS_REDIS
```

- **ACTION\_LOG\_ARCHIVE\_LOCATION** [string]: If action log archiving is enabled, the storage engine in which to place the archived data.
  - Example: **s3\_us\_east**
- **ACTION\_LOG\_ARCHIVE\_PATH** [string]: If action log archiving is enabled, the path in storage in which to place the archived data.
  - Example: **archives/actionlogs**
- **ACTION\_LOG\_ROTATION\_THRESHOLD** [string]: If action log archiving is enabled, the time interval after which to rotate logs.
  - Example: **30d**
- **ALLOW\_PULLS\_WITHOUT\_STRICT\_LOGGING** [boolean]: If true, pulls in which the pull audit log entry cannot be written will still succeed. Useful if the database can fallback into a read-only state and it is desired for pulls to continue during that time. Defaults to False.
  - Example: **True**
- **APP\_SPECIFIC\_TOKEN\_EXPIRATION** [string, null]: The expiration for external app tokens. Defaults to None.
  - Pattern: **^[0-9]+(w|m|d|h|s)\$**
- **AUTHENTICATION\_TYPE** [string] required: The authentication engine to use for credential authentication.
  - **enum**: Database, LDAP, JWT, Keystone, OIDC.
  - Example: **Database**
- **AVATAR\_KIND** [string]: The types of avatars to display, either generated inline (local) or Gravatar (gravatar)
  - **enum**: local, gravatar

- **BITBUCKET\_TRIGGER\_CONFIG** ['object', 'null']: Configuration for using BitBucket for build triggers.
  - **consumer\_key** [string] required: The registered consumer key(client ID) for this Red Hat Quay instance.
    - **Example:** **0e8dbe15c4c7630b6780**
- **BLACKLISTED\_EMAIL\_DOMAINS** [array]: The array of email-address domains that is used if `FEATURE_BLACKLISTED_EMAILS` is set to true.
  - **Example:** **"example.com", "example.org"**
- **BLACKLIST\_V2\_SPEC** [string]: The Docker CLI versions to which Red Hat Quay will respond that V2 is **unsupported**. Defaults to **<1.6.0**.
  - **Reference:**  
[http://pythonhosted.org/semantic\\_version/reference.html#semantic\\_version.Spec](http://pythonhosted.org/semantic_version/reference.html#semantic_version.Spec)
  - **Example:** **<1.8.0**
- **BRANDING** [object]: Custom branding for logos and URLs in the Red Hat Quay UI.
  - **Required:** logo
  - **properties:**
    - **logo** [string]: Main logo image URL.
      - **Example:** **/static/img/quay-horizontal-color.svg**
    - **footer\_img** [string]: Logo for UI footer.
      - **Example:** **/static/img/RedHat.svg**
    - **footer\_url** [string]: Link for footer image.
      - **Example:** **<https://redhat.com>**
- **BROWSER\_API\_CALLS\_XHR\_ONLY** [boolean]: If enabled, only API calls marked as being made by an XHR will be allowed from browsers. Defaults to True.
  - **Example:** False
- **BUILDLOGS\_REDIS** [object] required: Connection information for Redis for build logs caching.
  - **HOST** [string] required: The hostname at which Redis is accessible.
    - **Example:** **my.redis.cluster**
  - **PASSWORD** [string]: The password to connect to the Redis instance.
    - **Example:** **mypassword**
  - **PORT** [number]: The port at which Redis is accessible.
    - **Example:** **1234**

- **CONTACT\_INFO** [array]: If specified, contact information to display on the contact page. If only a single piece of contact information is specified, the contact footer will link directly.
  - **Min Items:** 1
  - **Unique Items:** True
    - **array item 0** [string]: Adds a link to send an e-mail
    - **Pattern:** `^mailto:(.)+$`
    - **Example:** `mailto:support@quay.io`
  - **array item 1** [string]: Adds a link to visit an IRC chat room
    - **Pattern:** `^irc://(.)+$`
    - **Example:** `irc://chat.freenode.net:6665/quay`
  - **array item 2** [string]: Adds a link to call a phone number
    - **Pattern:** `^tel:(.)+$`
    - **Example:** `tel:+1-888-930-3475`
  - **array item 3** [string]: Adds a link to a defined URL
    - **Pattern:** `^http(s)?://(.)+$`
    - **Example:** `https://twitter.com/quayio`
- **DB\_CONNECTION\_ARGS** [object]: If specified, connection arguments for the database such as timeouts and SSL.
  - **threadlocals** [boolean] required: Whether to use thread-local connections. Should **ALWAYS** be **true**.
  - **autorollback** [boolean] required: Whether to use auto-rollback connections. Should **ALWAYS** be **true**.
  - **ssl** [object]: SSL connection configuration
    - **ca** [string] required: Absolute container path to the CA certificate to use for SSL connections.
    - **Example:** `conf/stack/ssl-ca-cert.pem`
- **DATABASE\_SECRET\_KEY** [string] required: Key used to encrypt sensitive fields within the database. It is imperative that once set, this value is never changed. The consequence of changing this is invalidating all reliant fields (repository mirror username and password configurations, for example).
  - **Example:**  
`40157269433064266822674401740626984898972632465622168464725100311621640999470`
- **DB\_URI** [string] required: The URI at which to access the database, including any credentials.
  - **Reference:** <https://www.postgresql.org/docs/9.3/static/libpq-connect.html#AEN39495>

- Example: `mysql+pymysql://username:password@dns.of.database/quay`
- **DEFAULT\_NAMESPACE\_MAXIMUM\_BUILD\_COUNT** [number, **null**]: If not None, the default maximum number of builds that can be queued in a namespace.
  - Example: **20**
- **DEFAULT\_TAG\_EXPIRATION** [string] required: The default, configurable tag expiration time for time machine. Defaults to **2w**.
  - Pattern: `^[0-9]+(w|m|d|h|s)$`
- **DIRECT\_OAUTH\_CLIENTID\_WHITELIST** [array]: A list of client IDs of **Red Hat Quay-managed** applications that are allowed to perform direct OAuth approval without user approval.
  - **Min Items**: None
  - **Unique Items**: True
  - **Reference**: <https://coreos.com/quay-enterprise/docs/latest/direct-oauth.html>
    - **array item** [string]
- **DISTRIBUTED\_STORAGE\_CONFIG** [object] required: Configuration for storage engine(s) to use in Red Hat Quay. Each key represents a unique identifier for a storage engine. The value consists of a tuple of (key, value) forming an object describing the storage engine parameters.

- **OCS / Noobaa:**

```
rhocsStorage:
- RHOCSSStorage
- access_key: access_key_here
  secret_key: secret_key_here
  bucket_name: quay-datastore-9b2108a3-29f5-43f2-a9d5-2872174f9a56
  hostname: s3.openshift-storage.svc.cluster.local
  is_secure: 'true'
  port: '443'
  storage_path: /datastorage/registry
```

- **Ceph / RadosGW Storage / Hitachi HCP**

```
radosGWStorage:
- RadosGWStorage
- access_key: access_key_here
  secret_key: secret_key_here
  bucket_name: bucket_name_here
  hostname: hostname_here
  is_secure: 'true'
  port: '443'
  storage_path: /datastorage/registry
```

- **AWS S3 Storage:**

```
s3Storage:
- S3Storage
- host: s3.ap-southeast-2.amazonaws.com
  s3_access_key: s3_access_key_here
```

```
s3_secret_key: s3_secret_key_here
s3_bucket: s3_bucket_here
storage_path: /datastorage/registry
```

- **Azure Storage:**

```
azureStorage:
- AzureStorage
- azure_account_name: azure_account_name_here
  azure_account_key: azure_account_key_here
  azure_container: azure_container_here
  sas_token: some/path/
  storage_path: /datastorage/registry
```

- **Google Cloud Storage**

```
googleCloudStorage:
- GoogleCloudStorage
- access_key: access_key_here
  secret_key: secret_key_here
  bucket_name: bucket_name_here
  storage_path: /datastorage/registry
```

- **Swift Storage:**

```
swiftStorage:
- SwiftStorage
- swift_user: swift_user_here
  swift_password: swift_password_here
  swift_container: swift_container_here
  auth_url: https://example.org/swift/v1/quay
  auth_version: 1
  ca_cert_path: /conf/stack/swift.cert"
  storage_path: /datastorage/registry
```

- **DISTRIBUTED\_STORAGE\_DEFAULT\_LOCATIONS** [array]: The list of storage engine(s) (by ID in DISTRIBUTED\_STORAGE\_CONFIG) whose images should be fully replicated, by default, to all other storage engines.
  - **Min Items:** None
  - **Example:** **s3\_us\_east, s3\_us\_west**
    - **array item** [string]
- **DISTRIBUTED\_STORAGE\_PREFERENCE** [array] required: The preferred storage engine(s) (by ID in DISTRIBUTED\_STORAGE\_CONFIG) to use. A preferred engine means it is first checked for pulling and images are pushed to it.
  - **Min Items:** None
    - **Example:** **[u's3\_us\_east', u's3\_us\_west']**
    - **array item** [string]

- **preferred\_url\_scheme** [string] required: The URL scheme to use when hitting Red Hat Quay. If Red Hat Quay is behind SSL **at all**, this **must** be **https**.
  - enum: **http, https**
  - Example: **https**
- **DOCUMENTATION\_ROOT** [string]: Root URL for documentation links.
- **ENABLE\_HEALTH\_DEBUG\_SECRET** [string, **null**]: If specified, a secret that can be given to health endpoints to see full debug info when not authenticated as a superuser.
  - Example: **somesecrethere**
- **EXPIRED\_APP\_SPECIFIC\_TOKEN\_GC** [string, **null**]: Duration of time expired external app tokens will remain before being garbage collected. Defaults to 1d.
  - pattern: **^[0-9]+(w|m|d|h|s)\$**
- **EXTERNAL\_TLS\_TERMINATION** [boolean]: If TLS is supported, but terminated at a layer before Red Hat Quay, must be true.
  - Example: **True**
- **FEATURE\_ACI\_CONVERSION** [boolean]: Whether to enable conversion to ACIs. Defaults to False.
  - Example: **False**
- **FEATURE\_ACTION\_LOG\_ROTATION** [boolean]: Whether or not to rotate old action logs to storage. Defaults to False.
  - Example: **False**
- **FEATURE\_ADVERTISE\_V2** [boolean]: Whether the v2/ endpoint is visible. Defaults to True.
  - Example: **True**
- **FEATURE\_AGGREGATED\_LOG\_COUNT\_RETRIEVAL** [boolean]: Whether to allow retrieval of aggregated log counts. Defaults to True.
  - Example: **True**
- **FEATURE\_ANONYMOUS\_ACCESS** [boolean]: Whether to allow anonymous users to browse and pull public repositories. Defaults to True.
  - Example: **True**
- **FEATURE\_APP\_REGISTRY** [boolean]: Whether to enable support for App repositories. Defaults to False.
  - Example: **False**
- **FEATURE\_APP\_SPECIFIC\_TOKENS** [boolean]: If enabled, users can create tokens for use by the Docker CLI. Defaults to True.
  - Example: **False**
- **FEATURE\_BITBUCKET\_BUILD** [boolean]: Whether to support Bitbucket build triggers. Defaults to False.



- Example: **False**
- **FEATURE\_BLACKLISTED\_EMAIL**
- **FEATURE\_BUILD\_SUPPORT** [boolean]: Whether to support Dockerfile build. Defaults to True.
  - Example: **True**
- **FEATURE\_CHANGE\_TAG\_EXPIRATION** [boolean]: Whether users and organizations are allowed to change the tag expiration for tags in their namespace. Defaults to True.
  - Example: **False**
- **FEATURE\_DIRECT\_LOGIN** [boolean]: Whether users can directly login to the UI. Defaults to True.
  - Example: **True**
- **FEATURE\_GARBAGE\_COLLECTION** [boolean]: Whether garbage collection of repositories is enabled. Defaults to True.
  - Example: **True**
- **FEATURE\_GITHUB\_BUILD** [boolean]: Whether to support GitHub build triggers. Defaults to False.
  - Example: **False**
- **FEATURE\_GITHUB\_LOGIN** [boolean]: Whether GitHub login is supported. Defaults to False.
  - Example: **False**
- **FEATURE\_GITLAB\_BUILD**[boolean]: Whether to support GitLab build triggers. Defaults to False.
  - Example: **False**
- **FEATURE\_GOOGLE\_LOGIN** [boolean]: Whether Google login is supported. Defaults to False.
  - Example: **False**
- **FEATURE\_INVITE\_ONLY\_USER\_CREATION** [boolean]: Whether users being created must be invited by another user. Defaults to False.
  - Example: **False**
- **FEATURE\_LIBRARY\_SUPPORT** [boolean]: Whether to allow for "namespace-less" repositories when pulling and pushing from Docker. Defaults to True.
  - Example: **True**
- **FEATURE\_LOG\_EXPORT** [boolean]: Whether to allow exporting of action logs. Defaults to True.
  - Example: **True**
- **FEATURE\_MAILING** [boolean]: Whether emails are enabled. Defaults to True.
  - Example: **True**

- **FEATURE\_NONSUPERUSER\_TEAM\_SYNCING\_SETUP** [boolean]: If enabled, non-superusers can setup syncing on teams to backing LDAP or Keystone. Defaults To False.
  - **Example: True**
- **FEATURE\_PARTIAL\_USER\_AUTOCOMPLETE** [boolean]: If set to true, autocompletion will apply to partial usernames. Defaults to True.
  - **Example: True**
- **FEATURE\_PERMANENT\_SESSIONS** [boolean]: Whether sessions are permanent. Defaults to True.
  - **Example: True**
- **FEATURE\_PROXY\_STORAGE** [boolean]: Whether to proxy all direct download URLs in storage via the registry nginx. Defaults to False.
  - **Example: False**
- **FEATURE\_PUBLIC\_CATALOG** [boolean]: If set to true, the **\_catalog** endpoint returns public repositories. Otherwise, only private repositories can be returned. Defaults to False.
  - **Example: False**
- **FEATURE\_RATE\_LIMITS** [boolean]: Whether to enable rate limits on API and registry endpoints. Defaults to False.
  - **Example: False**
- **FEATURE\_READER\_BUILD\_LOGS** [boolean]: If set to true, build logs may be read by those with read access to the repo, rather than only write access or admin access. Defaults to False.
  - **Example: False**
- **FEATURE\_READONLY\_APP\_REGISTRY** [boolean]: Whether to App repositories are read-only. Defaults to False.
  - **Example: True**
- **FEATURE\_RECAPTCHA** [boolean]: Whether Recaptcha is necessary for user login and recovery. Defaults to False.
  - **Example: False**
  - **Reference:** <https://www.google.com/recaptcha/intro/>
- **FEATURE\_REPO\_MIRROR** [boolean]: If set to true, enables repository mirroring. Defaults to False.
  - **Example: False**
- **FEATURE\_REQUIRE\_ENCRYPTED\_BASIC\_AUTH** [boolean]: Whether non-encrypted passwords (as opposed to encrypted tokens) can be used for basic auth. Defaults to False.
  - **Example: False**
- **FEATURE\_REQUIRE\_TEAM\_INVITE** [boolean]: Whether to require invitations when adding a user to a team. Defaults to True.

- **Example: True**
- **FEATURE\_RESTRICTED\_V1\_PUSH** [boolean]: If set to true, only namespaces listed in V1\_PUSH\_WHITELIST support V1 push. Defaults to True.
  - **Example: True**
- **FEATURE\_SECURITY\_NOTIFICATIONS** [boolean]: If the security scanner is enabled, whether to turn on/off security notifications. Defaults to False.
  - **Example: False**
- **FEATURE\_SECURITY\_SCANNER** [boolean]: Whether to turn on/off the security scanner. Defaults to False.
  - **Reference:** [https://access.redhat.com/documentation/en-us/red\\_hat\\_quay/3.5/html-single/manage\\_red\\_hat\\_quay/#clair-initial-setup](https://access.redhat.com/documentation/en-us/red_hat_quay/3.5/html-single/manage_red_hat_quay/#clair-initial-setup)
  - **Example: False**
- **FEATURE\_STORAGE\_REPLICATION** [boolean]: Whether to automatically replicate between storage engines. Defaults to False.
  - **Example: False**
- **FEATURE\_SUPER\_USERS** [boolean]: Whether superusers are supported. Defaults to True.
  - **Example: True**
- **FEATURE\_TEAM\_SYNCING** [boolean]: Whether to allow for team membership to be synced from a backing group in the authentication engine (LDAP or Keystone).
  - **Example: True**
- **FEATURE\_USER\_CREATION** [boolean]: Whether users can be created (by non-superusers). Defaults to True.
  - **Example: True**
- **FEATURE\_USER\_LAST\_ACCESSED** [boolean]: Whether to record the last time a user was accessed. Defaults to True.
  - **Example: True**
- **FEATURE\_USER\_LOG\_ACCESS** [boolean]: If set to true, users will have access to audit logs for their namespace. Defaults to False.
  - **Example: True**
- **FEATURE\_USER\_METADATA** [boolean]: Whether to collect and support user metadata. Defaults to False.
  - **Example: False**
- **FEATURE\_USERNAME\_CONFIRMATION** [boolean]: If set to true, users can confirm their generated usernames. Defaults to True.
  - **Example: False**

- **FEATURE\_USER\_RENAME** [boolean]: If set to true, users can rename their own namespace. Defaults to False.
  - **Example: True**
- **FRESH\_LOGIN\_TIMEOUT** [string]: The time after which a fresh login requires users to reenter their password
  - **Example: 5m**
- **GITHUB\_LOGIN\_CONFIG** [object, 'null']: Configuration for using GitHub (Enterprise) as an external login provider.
  - **Reference:** <https://coreos.com/quay-enterprise/docs/latest/github-auth.html>
  - **allowed\_organizations** [array]: The names of the GitHub (Enterprise) organizations whitelisted to work with the **ORG\_RESTRICT** option.
    - **Min Items:** None
    - **Unique Items:** True
      - **array item** [string]
  - **API\_ENDPOINT** [string]: The endpoint of the GitHub (Enterprise) API to use. Must be overridden for github.com.
    - **Example:** <https://api.github.com/>
  - **CLIENT\_ID** [string] required: The registered client ID for this Red Hat Quay instance; cannot be shared with **GITHUB\_TRIGGER\_CONFIG**.
    - **Reference:** <https://coreos.com/quay-enterprise/docs/latest/github-app.html>
    - **Example:** **0e8dbe15c4c7630b6780**
  - **CLIENT\_SECRET** [string] required: The registered client secret for this Red Hat Quay instance.
    - **Reference:** <https://coreos.com/quay-enterprise/docs/latest/github-app.html>
    - **Example:** **e4a58ddd3d7408b7aec109e85564a0d153d3e846**
  - **GITHUB\_ENDPOINT** [string] required: The endpoint of the GitHub (Enterprise) being hit.
    - **Example:** <https://github.com/>
  - **ORG\_RESTRICT** [boolean]: If true, only users within the organization whitelist can login using this provider.
    - **Example: True**
- **GITHUB\_TRIGGER\_CONFIG** [object, **null**]: Configuration for using GitHub (Enterprise) for build triggers.
  - **Reference:** <https://coreos.com/quay-enterprise/docs/latest/github-build.html>
  - **API\_ENDPOINT** [string]: The endpoint of the GitHub (Enterprise) API to use. Must be overridden for github.com.

- Example: <https://api.github.com/>
- **CLIENT\_ID** [string] required: The registered client ID for this Red Hat Quay instance; cannot be shared with GITHUB\_LOGIN\_CONFIG.
  - Reference: <https://coreos.com/quay-enterprise/docs/latest/github-app.html>
  - Example: **0e8dbe15c4c7630b6780**
- **CLIENT\_SECRET** [string] required: The registered client secret for this Red Hat Quay instance.
  - Reference: <https://coreos.com/quay-enterprise/docs/latest/github-app.html>
  - Example: **e4a58ddd3d7408b7aec109e85564a0d153d3e846**
- **GITHUB\_ENDPOINT** [string] required: The endpoint of the GitHub (Enterprise) being hit.
  - Example: <https://github.com/>
- **GITLAB\_TRIGGER\_CONFIG** [object]: Configuration for using Gitlab (Enterprise) for external authentication.
  - **CLIENT\_ID** [string] required: The registered client ID for this Red Hat Quay instance.
    - Example: **0e8dbe15c4c7630b6780**
  - **CLIENT\_SECRET** [string] required: The registered client secret for this Red Hat Quay instance.
    - Example: **e4a58ddd3d7408b7aec109e85564a0d153d3e846**
    - **gitlab\_endpoint** [string] required: The endpoint at which Gitlab(Enterprise) is running.
      - Example: <https://gitlab.com>
- **GOOGLE\_LOGIN\_CONFIG** [object, **null**]: Configuration for using Google for external authentication
  - **CLIENT\_ID** [string] required: The registered client ID for this Red Hat Quay instance.
    - Example: **0e8dbe15c4c7630b6780**
  - **CLIENT\_SECRET** [string] required: The registered client secret for this Red Hat Quay instance.
    - Example: **e4a58ddd3d7408b7aec109e85564a0d153d3e846**
- **GPG2\_PRIVATE\_KEY\_FILENAME** [string]: The filename of the private key used to decrypt ACIs.
  - Example: **/path/to/file**
- **GPG2\_PRIVATE\_KEY\_NAME** [string]: The name of the private key used to sign ACIs.
  - Example: **gpg2key**
- **GPG2\_PUBLIC\_KEY\_FILENAME** [string]: The filename of the public key used to encrypt ACIs.
  - Example: **/path/to/file**

- **HEALTH\_CHECKER** [string]: The configured health check.
  - Example: ('RDSAwareHealthCheck', {'access\_key': 'foo', 'secret\_key': 'bar'})
- **JWT\_AUTH\_ISSUER** [string]: The endpoint for JWT users.
  - Example: <http://192.168.99.101:6060>
  - Pattern: `^http(s)?://(.)+$`
- **JWT\_GETUSER\_ENDPOINT** [string]: The endpoint for JWT users.
  - Example: <http://192.168.99.101:6060>
  - Pattern: `^http(s)?://(.)+$`
- **JWT\_QUERY\_ENDPOINT** [string]: The endpoint for JWT queries.
  - Example: <http://192.168.99.101:6060>
  - Pattern: `^http(s)?://(.)+$`
- **JWT\_VERIFY\_ENDPOINT** [string]: The endpoint for JWT verification.
  - Example: <http://192.168.99.101:6060>
  - Pattern: `^http(s)?://(.)+$`
- **LDAP\_ADMIN\_DN** [string]: The admin DN for LDAP authentication.
- **LDAP\_ADMIN\_PASSWD** [string]: The admin password for LDAP authentication.
- **LDAP\_ALLOW\_INSECURE\_FALLBACK** [boolean]: Whether or not to allow SSL insecure fallback for LDAP authentication.
- **LDAP\_BASE\_DN** [string]: The base DN for LDAP authentication.
- **LDAP\_EMAIL\_ATTR** [string]: The email attribute for LDAP authentication.
- **LDAP\_UID\_ATTR** [string]: The uid attribute for LDAP authentication.
- **LDAP\_URI** [string]: The LDAP URI.
- **LDAP\_USER\_FILTER** [string]: The user filter for LDAP authentication.
- **LDAP\_USER\_RDN** [array]: The user RDN for LDAP authentication.
- **LOGS\_MODEL** [string]: Logs model for action logs.
  - **enum**: database, transition\_reads\_both\_writes\_es, elasticsearch
  - Example: **database**
- **LOGS\_MODEL\_CONFIG** [object]: Logs model config for action logs
  - **elasticsearch\_config** [object]: Elasticsearch cluster configuration
    - **access\_key** [string]: Elasticsearch user (or IAM key for AWS ES)

- Example: **some\_string**
- **host** [string]: Elasticsearch cluster endpoint
  - Example: **host.elasticsearch.example**
- **index\_prefix** [string]: Elasticsearch's index prefix
  - Example: **logentry\_**
- **index\_settings** [object]: Elasticsearch's index settings
- **use\_ssl** [boolean]: Use ssl for Elasticsearch. Defaults to True
  - Example: **True**
- **secret\_key** [string]: Elasticsearch password (or IAM secret for AWS ES)
  - Example: **some\_secret\_string**
- **aws\_region** [string]: Amazon web service region
  - Example: **us-east-1**
- **port** [number]: Elasticsearch cluster endpoint port
  - Example: **1234**
- **kinesis\_stream\_config** [object]: AWS Kinesis Stream configuration
  - **aws\_secret\_key** [string]: AWS secret key
    - Example: **some\_secret\_key**
  - **stream\_name** [string]: Kinesis stream to send action logs to
    - Example: **logentry-kinesis-stream**
  - **aws\_access\_key** [string]: AWS access key
    - Example: **some\_access\_key**
  - **retries** [number]: Max number of attempts made on a single request
    - Example: **5**
  - **read\_timeout** [number]: Number of seconds before timeout when reading from a connection
    - Example: **5**
  - **max\_pool\_connections** [number]: The maximum number of connections to keep in a connection pool
    - Example: **10**
  - **aws\_region** [string]: AWS region
    - Example: **us-east-1**

- **connect\_timeout** [number]: Number of seconds before timeout when attempting to make a connection
  - **Example: 5**
- **producer** [string]: Logs producer if logging to Elasticsearch
  - **enum:** kafka, elasticsearch, kinesis\_stream
  - **Example: kafka**
- **kafka\_config** [object]: Kafka cluster configuration
  - **topic** [string]: Kafka topic to publish log entries to
    - **Example: logentry**
  - **bootstrap\_servers** [array]: List of Kafka brokers to bootstrap the client from
  - **max\_block\_seconds** [number]: Max number of seconds to block during a **send()**, either because the buffer is full or metadata unavailable
    - **Example: 10**
- **LOG\_ARCHIVE\_LOCATION** [string]: If builds are enabled, the storage engine in which to place the archived build logs.
  - **Example: s3\_us\_east**
- **LOG\_ARCHIVE\_PATH** [string]: If builds are enabled, the path in storage in which to place the archived build logs.
  - **Example: archives/buildlogs**
- **LOGS\_MODEL** [string]: Logs model for action logs.
- **enum:** database, transition\_reads\_both\_writes\_es, elasticsearch
- **Example: database**
- **MAIL\_DEFAULT\_SENDER** [string, null]: If specified, the e-mail address used as the **from** when Red Hat Quay sends e-mails. If none, defaults to **support@quay.io**.
  - **Example: support@myco.com**
- **MAIL\_PASSWORD** [string, null]: The SMTP password to use when sending e-mails.
  - **Example: mypassword**
- **MAIL\_PORT** [number]: The SMTP port to use. If not specified, defaults to 587.
  - **Example: 588**
- **MAIL\_SERVER** [string]: The SMTP server to use for sending e-mails. Only required if **FEATURE\_MAILING** is set to true.
  - **Example: smtp.somedomain.com**
- **MAIL\_USERNAME** [string, 'null']: The SMTP username to use when sending e-mails.



- Example: **myuser**
- **MAIL\_USE\_TLS** [boolean]: If specified, whether to use TLS for sending e-mails.
  - Example: **True**
- **MAXIMUM\_LAYER\_SIZE** [string]: Maximum allowed size of an image layer. Defaults to 20G.
  - Pattern: **^[0-9]+(G|M)\$**
  - Example: **100G**
- **PREFERRED\_URL\_SCHEME** [string]: The URL scheme to use when hitting Red Hat Quay. If Red Hat Quay is behind SSL **at all**, this **must** be **https**
  - enum: **http** or **https**
  - Example: **https**
- **PROMETHEUS\_NAMESPACE** [string]: The prefix applied to all exposed Prometheus metrics. Defaults to **quay**.
  - Example: **myregistry**
- **PUBLIC\_NAMESPACES** [array]: If a namespace is defined in the public namespace list, then it will appear on **all** user's repository list pages, regardless of whether that user is a member of the namespace. Typically, this is used by an enterprise customer in configuring a set of "well-known" namespaces.
  - Min Items: None
  - Unique Items: True
    - array item [string]
- **RECAPTCHA\_SECRET\_KEY** [string]: If recaptcha is enabled, the secret key for the Recaptcha service.
- **RECAPTCHA\_SITE\_KEY** [string]: If recaptcha is enabled, the site key for the Recaptcha service.
- **REGISTRY\_STATE** [string]: The state of the registry.
  - enum: **normal** or **read-only**
  - Example: **read-only**
- **REGISTRY\_TITLE** [string]: If specified, the long-form title for the registry. Defaults to **Quay Enterprise**.
  - Example: **Corp Container Service**
- **REGISTRY\_TITLE\_SHORT** [string]: If specified, the short-form title for the registry. Defaults to **Quay Enterprise**.
  - Example: **CCS**
- **REPO\_MIRROR\_INTERVAL** [number]: The number of seconds between checking for repository mirror candidates. Defaults to 30.

- **Example: 30**
- **REPO\_MIRROR\_SERVER\_HOSTNAME** [string]: Replaces the SERVER\_HOSTNAME as the destination for mirroring. Defaults to unset.
  - **Example: openshift-quay-service**
- **REPO\_MIRROR\_TLS\_VERIFY** [boolean]: Require HTTPS and verify certificates of Quay registry during mirror. Defaults to True.
  - **Example: True**
- **SEARCH\_MAX\_RESULT\_PAGE\_COUNT** [number]: Maximum number of pages the user can paginate in search before they are limited. Defaults to 10.
  - **Example: 10**
- **SEARCH\_RESULTS\_PER\_PAGE** [number]: Number of results returned per page by search page. Defaults to 10.
  - **Example: 10**
- **SECRET\_KEY** [string] required: Key used to encrypt sensitive fields within the database and a run time. It is imperative that once set, this value is never changed. The consequence of changing this is invalidating all reliant fields (encrypted password credentials, for example).
  - **Example:**  
**40157269433064266822674401740626984898972632465622168464725100311621640999470**
- **SECURITY\_SCANNER\_ENDPOINT** [string]: The endpoint for the security scanner.
  - **Pattern: `^http(s)?://(.)+$`**
  - **Example: <http://192.168.99.101:6060>**
- **SECURITY\_SCANNER\_INDEXING\_INTERVAL** [number]: The number of seconds between indexing intervals in the security scanner. Defaults to 30.
  - **Example: 30**
- **SECURITY\_SCANNER\_NOTIFICATIONS** [boolean]: Whether or not to the security scanner notification feature
  - **Example: false**
- **SECURITY\_SCANNER\_V4\_ENDPOINT** [string]: The endpoint for the V4 security scanner.
  - **Pattern: `^http(s)?://(.)+$`**
  - **Example: <http://192.168.99.101:6060>**
- **SECURITY\_SCANNER\_V4\_PSK** [string]: The generated pre-shared key (PSK) for Clair.
- **SERVER\_HOSTNAME** [string] required: The URL at which Red Hat Quay is accessible, without the scheme.
  - **Example: quay.io**

- **SESSION\_COOKIE\_SECURE** [boolean]: Whether the **secure** property should be set on session cookies. Defaults to False. Recommended to be True for all installations using SSL.
  - **Example:** True
  - **Reference:** [https://en.wikipedia.org/wiki/Secure\\_cookies](https://en.wikipedia.org/wiki/Secure_cookies)
- **SSL\_CIPHERS** [array]: If specified, the nginx-defined list of SSL ciphers to enabled and disabled.
  - **Example:** **CAMELLIA, !3DES**
- **SSL\_PROTOCOLS** [array]: If specified, nginx is configured to enabled a list of SSL protocols defined in the list. Removing an SSL protocol from the list disables the protocol during Red Hat Quay startup.
  - **SSL\_PROTOCOLS:** ['TLSv1','TLSv1.1','TLSv1.2']
- **SUCCESSIVE\_TRIGGER\_FAILURE\_DISABLE\_THRESHOLD** [number]: If not None, the number of successive failures that can occur before a build trigger is automatically disabled. Defaults to 100.
  - **Example:** 50
- **SUCCESSIVE\_TRIGGER\_INTERNAL\_ERROR\_DISABLE\_THRESHOLD** [number]: If not None, the number of successive internal errors that can occur before a build trigger is automatically disabled. Defaults to 5.
- **SUPER\_USERS** [array]: Red Hat Quay usernames of those users to be granted superuser privileges.
  - **Min Items:** None
  - **Unique Items:** True
    - **array item** [string]
- **TAG\_EXPIRATION\_OPTIONS** [array] required: The options that users can select for expiration of tags in their namespace (if enabled).
  - **Min Items:** None
  - **array item** [string]
  - **Pattern:** **^[0-9]+(w|m|d|h|s)\$**
- **TEAM\_RESYNC\_STALE\_TIME** [string]: If team syncing is enabled for a team, how often to check its membership and resync if necessary (Default: 30m).
  - **Pattern:** **^[0-9]+(w|m|d|h|s)\$**
  - **Example:** 2h
- **USERFILES\_LOCATION** [string]: ID of the storage engine in which to place user-uploaded files
  - **Example:** **s3\_us\_east**
- **USERFILES\_PATH** [string]: Path under storage in which to place user-uploaded files

- **Example: userfiles**
- **USER\_EVENTS\_REDIS** [object] required: Connection information for Redis for user event handling.
  - **HOST** [string] required: The hostname at which Redis is accessible.
    - **Example: my.redis.cluster**
  - **PASSWORD** [string]: The password to connect to the Redis instance.
    - **Example: mypassword**
  - **PORT** [number]: The port at which Redis is accessible.
    - **Example: 1234**
  - **CONSUMER\_SECRET** [string] required: The registered consumer secret(client secret) for this Red Hat Quay instance
    - **Example: e4a58ddd3d7408b7aec109e85564a0d153d3e846**
- **USERFILES\_LOCATION** [string]: ID of the storage engine in which to place user-uploaded files.
  - **Example: s3\_us\_east**
- **USERFILES\_PATH** [string]: Path under storage in which to place user-uploaded files.
  - **Example: userfiles**
- **USER\_RECOVERY\_TOKEN\_LIFETIME** [string]: The length of time a token for recovering a user accounts is valid. Defaults to 30m.
  - **Example: 10m**
  - **Pattern: `^[0-9]+(w|m|d|h|s)$`**
- **V1\_PUSH\_WHITELIST** [array]: The array of namespace names that support V1 push if `FEATURE_RESTRICTED_V1_PUSH` is set to true.
  - **Example: some, namespaces**
- **V2\_PAGINATION\_SIZE** [number]: The number of results returned per page in V2 registry APIs.
  - **Example: 100**
- **WEBHOOK\_HOSTNAME\_BLACKLIST** [array]: The set of hostnames to disallow from webhooks when validating, beyond localhost.
  - **Example: someexternaldomain.com**

## ADDITIONAL RESOURCES