



# Red Hat Quay 3.13

## Configure Red Hat Quay

Customizing Red Hat Quay using configuration options



## Red Hat Quay 3.13 Configure Red Hat Quay

---

Customizing Red Hat Quay using configuration options

## Legal Notice

Copyright © 2024 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

Configure Red Hat Quay

# Table of Contents

<b>CHAPTER 1. GETTING STARTED WITH RED HAT QUAY CONFIGURATION</b> .....	<b>6</b>
<b>CHAPTER 2. RED HAT QUAY CONFIGURATION DISCLAIMER</b> .....	<b>7</b>
2.1. CONFIGURATION UPDATES FOR RED HAT QUAY 3.12	7
2.1.1. Registry auto-pruning configuration fields	7
2.1.2. OAuth access token reassignment configuration field	8
2.1.3. Vulnerability detection notification configuration field	8
2.1.4. OCI referrers API configuration field	9
2.1.5. Disable strict logging configuration field	9
2.1.6. Notification interval configuration field	9
2.1.7. Clair indexing layer size configuration field	10
2.2. EDITING THE CONFIGURATION FILE	10
2.3. LOCATION OF CONFIGURATION FILE IN A STANDALONE DEPLOYMENT	10
2.4. MINIMAL CONFIGURATION	11
2.4.1. Sample minimal configuration file	11
2.4.2. Local storage	12
2.4.3. Cloud storage	12
<b>CHAPTER 3. CONFIGURATION FIELDS</b> .....	<b>14</b>
3.1. REQUIRED CONFIGURATION FIELDS	14
3.2. AUTOMATION OPTIONS	14
3.3. OPTIONAL CONFIGURATION FIELDS	14
3.4. GENERAL REQUIRED FIELDS	15
3.5. DATABASE CONFIGURATION	16
3.5.1. Database URI	16
3.5.2. Database connection arguments	17
3.5.2.1. PostgreSQL SSL/TLS connection arguments	17
3.5.2.2. MySQL SSL/TLS connection arguments	18
3.6. IMAGE STORAGE	18
3.6.1. Image storage features	18
3.6.2. Image storage configuration fields	19
3.6.3. Local storage	20
3.6.4. OpenShift Container Storage/NooBaa	20
3.6.5. Ceph Object Gateway/RadosGW storage	21
3.6.6. AWS S3 storage	21
3.6.6.1. AWS STS S3 storage	22
3.6.7. Google Cloud Storage	22
3.6.8. Azure Storage	23
3.6.9. Swift storage	23
3.6.10. Nutanix object storage	24
3.6.11. IBM Cloud object storage	24
3.6.12. NetApp ONTAP S3 object storage	24
3.6.13. Hitachi Content Platform object storage	25
3.7. REDIS CONFIGURATION FIELDS	25
3.7.1. Build logs	25
3.7.2. User events	26
3.7.3. Example Redis configuration	27
3.8. MODELCACHE CONFIGURATION OPTIONS	28
3.8.1. Memcache configuration option	28
3.8.2. Single Redis configuration option	28
3.8.3. Clustered Redis configuration option	28

3.9. TAG EXPIRATION CONFIGURATION FIELDS	29
3.9.1. Example tag expiration configuration	30
3.9.2. Registry-wide auto-prune policies examples	30
3.10. QUOTA MANAGEMENT CONFIGURATION FIELDS	31
3.10.1. Example quota management configuration	32
3.11. PROXY CACHE CONFIGURATION FIELDS	32
3.12. ROBOT ACCOUNT CONFIGURATION FIELDS	32
3.13. PRE-CONFIGURING RED HAT QUAY FOR AUTOMATION	33
3.13.1. Allowing the API to create the first user	33
3.13.2. Enabling general API access	33
3.13.3. Adding a superuser	33
3.13.4. Restricting user creation	33
3.13.5. Enabling new functionality in Red Hat Quay 3.13	34
3.13.6. Suggested configuration for automation	34
3.13.7. Deploying the Red Hat Quay Operator using the initial configuration	34
3.13.8. Using the API to create the first user	35
3.13.8.1. Using the OAuth token	36
3.13.8.2. Using the API to create an organization	37
3.14. BASIC CONFIGURATION FIELDS	38
3.15. SSL CONFIGURATION FIELDS	39
3.15.1. Configuring SSL	42
3.16. ADDING ADDITIONAL CERTIFICATE AUTHORITIES TO THE RED HAT QUAY CONTAINER	42
3.17. LDAP CONFIGURATION FIELDS	44
3.17.1. LDAP configuration references	46
3.17.1.1. Basic LDAP configuration	46
3.17.1.2. LDAP restricted user configuration	47
3.17.1.3. LDAP superuser configuration reference	48
3.18. MIRRORING CONFIGURATION FIELDS	49
3.19. SECURITY SCANNER CONFIGURATION FIELDS	49
3.19.1. Re-indexing with Clair v4	52
3.19.2. Example security scanner configuration	52
3.20. HELM CONFIGURATION FIELDS	52
3.20.1. Configuring Helm	53
3.21. OPEN CONTAINER INITIATIVE CONFIGURATION FIELDS	53
3.22. UNKNOWN MEDIA TYPES	54
3.22.1. Configuring unknown media types	54
3.23. ACTION LOG CONFIGURATION FIELDS	54
3.23.1. Action log storage configuration	54
3.23.1.1. Elasticsearch configuration fields	55
3.23.1.2. Splunk configuration fields	56
3.23.1.3. Splunk HEC configuration fields	57
3.23.2. Action log rotation and archiving configuration	58
3.23.3. Action log audit configuration	58
3.24. BUILD LOGS CONFIGURATION FIELDS	59
3.25. DOCKERFILE BUILD TRIGGERS FIELDS	59
3.25.1. GitHub build triggers	60
3.25.2. BitBucket build triggers	61
3.25.3. GitLab build triggers	61
3.26. BUILD MANAGER CONFIGURATION FIELDS	62
3.27. OAUTH CONFIGURATION FIELDS	65
3.27.1. GitHub OAuth configuration fields	66
3.27.2. Google OAuth configuration fields	67
3.28. OIDC CONFIGURATION FIELDS	67

3.28.1. OIDC configuration	69
3.29. NESTED REPOSITORIES CONFIGURATION FIELDS	70
3.30. QUAYINTEGRATION CONFIGURATION FIELDS	70
3.31. MAIL CONFIGURATION FIELDS	71
3.32. USER CONFIGURATION FIELDS	72
3.32.1. User configuration fields references	74
3.32.1.1. FEATURE_SUPERUSERS_FULL_ACCESS configuration reference	74
3.32.1.2. GLOBAL_READONLY_SUPER_USERS configuration reference	75
3.32.1.3. FEATURE_RESTRICTED_USERS configuration reference	75
3.32.1.4. RESTRICTED_USERS_WHITELIST configuration reference	75
3.33. RECAPTCHA CONFIGURATION FIELDS	75
3.34. ACI CONFIGURATION FIELDS	76
3.35. JWT CONFIGURATION FIELDS	76
3.36. APP TOKENS CONFIGURATION FIELDS	77
3.37. MISCELLANEOUS CONFIGURATION FIELDS	77
3.38. LEGACY CONFIGURATION FIELDS	82
3.39. USER INTERFACE V2 CONFIGURATION FIELDS	83
3.39.1. v2 user interface configuration	84
3.40. IPV6 CONFIGURATION FIELD	85
3.41. BRANDING CONFIGURATION FIELDS	85
3.41.1. Example configuration for Red Hat Quay branding	85
3.42. SESSION TIMEOUT CONFIGURATION FIELD	86
3.42.1. Example session timeout configuration	86
<b>CHAPTER 4. ENVIRONMENT VARIABLES</b> .....	<b>87</b>
4.1. GEO-REPLICATION	87
4.2. DATABASE CONNECTION POOLING	87
4.3. HTTP CONNECTION COUNTS	88
4.4. WORKER COUNT VARIABLES	89
4.5. DEBUG VARIABLES	89
<b>CHAPTER 5. CLAIR SECURITY SCANNER</b> .....	<b>91</b>
5.1. CLAIR CONFIGURATION OVERVIEW	91
5.1.1. Information about using Clair in a proxy environment	91
5.1.2. Clair configuration reference	92
5.1.3. Clair general fields	93
Example configuration for general Clair fields	93
5.1.4. Clair indexer configuration fields	94
Example indexer configuration	95
5.1.5. Clair matcher configuration fields	95
Example matcher configuration	97
5.1.6. Clair matchers configuration fields	97
Example matchers configuration	98
5.1.7. Clair updaters configuration fields	98
Example updaters configuration	99
5.1.8. Clair notifier configuration fields	99
Example notifier configuration	100
5.1.8.1. Clair webhook configuration fields	101
Example webhook configuration	101
5.1.8.2. Clair amqp configuration fields	101
Example AMQP configuration	103
5.1.8.3. Clair STOMP configuration fields	103
Example STOMP configuration	105

5.1.9. Clair authorization configuration fields	105
Example authorization configuration	106
5.1.10. Clair trace configuration fields	106
Example trace configuration	107
5.1.11. Clair metrics configuration fields	107
Example metrics configuration	107





# CHAPTER 1. GETTING STARTED WITH RED HAT QUAY CONFIGURATION

Red Hat Quay can be deployed by an independent, standalone configuration, or by using the Red Hat Quay Operator on OpenShift Container Platform.

How you create, retrieve, update, and validate the Red Hat Quay configuration varies depending on the type of deployment you are using. However, the core configuration options are the same for either deployment type. Core configuration is primarily set through a **config.yaml** file, but can also be set by using the configuration API.

For standalone deployments of Red Hat Quay, you must supply the minimum required configuration parameters before the registry can be started. The minimum requirements to start a Red Hat Quay registry can be found in the "Retrieving the current configuration" section.

If you install Red Hat Quay on OpenShift Container Platform using the Red Hat Quay Operator, you do not need to supply configuration parameters because the Red Hat Quay Operator supplies default information to deploy the registry.

After you have deployed Red Hat Quay with the desired configuration, you should retrieve, and save, the full configuration from your deployment. The full configuration contains additional generated values that you might need when restarting or upgrading your system.

## CHAPTER 2. RED HAT QUAY CONFIGURATION DISCLAIMER

With both standalone and Operator-based deployments of Red Hat Quay certain features and configuration parameters are not actively used or implemented. As a result, feature flags, such as those that enable or disable certain features, and configuration parameters that are not explicitly documented or requested for documentation by Red Hat Support, should only be modified with caution. Unused features or parameters might not be fully tested, supported, or compatible with Red Hat Quay. Modifying unused features parameters might lead to unexpected issues or disruptions with your deployment.

For information about configuring Red Hat Quay in standalone deployments, see [Advanced Red Hat Quay configuration](#)

For information about configuring Red Hat Quay Operator deployments, see [Configuring Red Hat Quay on OpenShift Container Platform](#)

### 2.1. CONFIGURATION UPDATES FOR RED HAT QUAY 3.12

The following sections detail new configuration fields added in Red Hat Quay 3.12.

#### 2.1.1. Registry auto-pruning configuration fields

The following configuration fields have been added to Red Hat Quay auto-pruning feature:

Field	Type	Description
<code>NOTIFICATION_TASK_RUN_MINIMUM_INTERVAL_MINUTES</code>	Integer	The interval, in minutes, that defines the frequency to re-run notifications for expiring images.  <b>Default: 300</b>
<code>DEFAULT_NAMESPACE_AUTOPRUNE_POLICY</code>	Object	The default organization-wide auto-prune policy.
<code>.method: number_of_tags</code>	Object	The option specifying the number of tags to keep.
<code>.value: &lt;integer&gt;</code>	Integer	When used with <b>method: number_of_tags</b> , denotes the number of tags to keep.  For example, to keep two tags, specify <b>2</b> .
<code>.method: creation_date</code>	Object	The option specifying the duration of which to keep tags.

<code>.value: &lt;integer&gt;</code>	Integer	When used with <code>creation_date</code> , denotes how long to keep tags.  Can be set to seconds ( <b>s</b> ), days ( <b>d</b> ), months ( <b>m</b> ), weeks ( <b>w</b> ), or years ( <b>y</b> ). Must include a valid integer. For example, to keep tags for one year, specify <b>1y</b> .
<code>AUTO_PRUNING_DEFAULT_POLICY_POLL_PERIOD</code>	Integer	The period in which the auto-pruner worker runs at the registry level. By default, it is set to run one time per day (one time per 24 hours). Value must be in seconds.

### 2.1.2. OAuth access token reassignment configuration field

The following configuration field has been added for reassigning OAuth access tokens:

Field	Type	Description
<code>FEATURE_ASSIGN_OAUTH_TOKEN</code>	Boolean	Allows organization administrators to assign OAuth tokens to other users.

#### Example OAuth access token reassignment YAML

```
# ...
FEATURE_ASSIGN_OAUTH_TOKEN: true
# ...
```

### 2.1.3. Vulnerability detection notification configuration field

The following configuration field has been added to notify users on detected vulnerabilities based on security level:

Field	Type	Description
<code>NOTIFICATION_MIN_SEVERITY_ON_NEW_INDEX</code>	String	Set minimal security level for new notifications on detected vulnerabilities. Avoids creation of large number of notifications after first index. If not defined, defaults to <b>High</b> . Available options include <b>Critical, High, Medium, Low, Negligible</b> , and <b>Unknown</b> .

## Example image vulnerability notification YAML

```
NOTIFICATION_MIN_SEVERITY_ON_NEW_INDEX: High
```

### 2.1.4. OCI referrers API configuration field

The following configuration field allows users to list OCI referrers of a manifest under a repository by using the v2 API:

Field	Type	Description
<code>FEATURE_REFERRERS_API</code>	Boolean	Enables OCI 1.1's referrers API.

## Example OCI referrers enablement YAML

```
# ...
FEATURE_REFERRERS_API: true
# ...
```

### 2.1.5. Disable strict logging configuration field

The following configuration field has been added to address when external systems like Splunk or ElasticSearch are configured as audit log destinations but are intermittently unavailable. When set to **True**, the logging event is logged to the stdout instead.

Field	Type	Description
<code>ALLOW_WITHOUT_STRICT_LOGGING</code>	Boolean	When set to <b>True</b> , if the external log system like Splunk or ElasticSearch is intermittently unavailable, allows users to push images normally. Events are logged to the stdout instead. Overrides <b>ALLOW_PULLS_WITHOUT_STRICT_LOGGING</b> if set.

## Example strict logging YAML

```
# ...
ALLOW_WITHOUT_STRICT_LOGGING: True
# ...
```

### 2.1.6. Notification interval configuration field

The following configuration field has been added to enhance Red Hat Quay notifications:

Field	Type	Description
-------	------	-------------

NOTIFICATION_TASK_RUN_MINIMUM_INTERVAL_MINUTES	Integer	The interval, in minutes, that defines the frequency to re-run notifications for expiring images. By default, this field is set to notify Red Hat Quay users of events happening every 5 hours.
--	---------	---

### Example notification re-run YAML

```
# ...
NOTIFICATION_TASK_RUN_MINIMUM_INTERVAL_MINUTES: 10
# ...
```

### 2.1.7. Clair indexing layer size configuration field

The following configuration field has been added for the Clair security scanner, which allows Red Hat Quay administrators to set a maximum layer size allowed for indexing.

Field	Type	Description
SECURITY_SCANNER_V4_INDEX_MAX_LAYER_SIZE	String	The maximum layer size allowed for indexing. If the layer size exceeds the configured size, the Red Hat Quay UI returns the following message: <b>The manifest for this tag has layer(s) that are too large to index by the Quay Security Scanner.</b> The default is <b>8G</b> , and the maximum recommended is <b>10G</b> . <b>Example: 8G</b>

## 2.2. EDITING THE CONFIGURATION FILE

To deploy a standalone instance of Red Hat Quay, you must provide the minimal configuration information. The requirements for a minimal configuration can be found in "Red Hat Quay minimal configuration."

After supplying the required fields, you can validate your configuration. If there are any issues, they will be highlighted.

For changes to take effect, the registry must be restarted.

## 2.3. LOCATION OF CONFIGURATION FILE IN A STANDALONE DEPLOYMENT

For standalone deployments of Red Hat Quay, the **config.yaml** file must be specified when starting the Red Hat Quay registry. This file is located in the configuration volume. For example, the configuration file is located at **\$QUAY/config/config.yaml** when deploying Red Hat Quay by the following command:

```
$ 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:v3.13.1
```

## 2.4. MINIMAL CONFIGURATION

The following configuration options are required for a standalone deployment of Red Hat Quay:

- Server hostname
- HTTP or HTTPS
- Authentication type, for example, Database or Lightweight Directory Access Protocol (LDAP)
- Secret keys for encrypting data
- Storage for images
- Database for metadata
- Redis for build logs and user events
- Tag expiration options

### 2.4.1. Sample minimal configuration file

The following example shows a sample minimal configuration file that uses local storage for images:

```
AUTHENTICATION_TYPE: Database
BUILDLOGS_REDIS:
  host: quay-server.example.com
  password: strongpassword
  port: 6379
  ssl: false
DATABASE_SECRET_KEY: 0ce4f796-c295-415b-bf9d-b315114704b8
DB_URI: postgresql://quayuser:quaypass@quay-server.example.com:5432/quay
DEFAULT_TAG_EXPIRATION: 2w
DISTRIBUTED_STORAGE_CONFIG:
  default:
    - LocalStorage
    - storage_path: /datastorage/registry
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS: []
DISTRIBUTED_STORAGE_PREFERENCE:
  - default
PREFERRED_URL_SCHEME: http
SECRET_KEY: e8f9fe68-1f84-48a8-a05f-02d72e6eccba
SERVER_HOSTNAME: quay-server.example.com
SETUP_COMPLETE: true
TAG_EXPIRATION_OPTIONS:
  - 0s
  - 1d
  - 1w
  - 2w
```

```

- 4w
- 3y
USER_EVENTS_REDIS:
  host: quay-server.example.com
  port: 6379
  ssl: false

```

## 2.4.2. Local storage

Using local storage for images is only recommended when deploying a registry for *proof of concept* purposes.

When configuring local storage, storage is specified on the command line when starting the registry.

The following command maps a local directory, **\$QUAY/storage** to the **datastorage** path in the container:

```

$ 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:v3.13.1

```

## 2.4.3. Cloud storage

Storage configuration is detailed in the [Image storage](#) section. For some users, it might be useful to compare the difference between Google Cloud Platform and local storage configurations. For example, the following YAML presents a Google Cloud Platform storage configuration:

### \$QUAY/config/config.yaml

```

DISTRIBUTED_STORAGE_CONFIG:
  default:
    - GoogleCloudStorage
    - access_key: GOOGQIMFB3ABCDEFGHIJKLMN
      bucket_name: quay_bucket
      secret_key: FhDAYe2HeuAKfvZCAGyOioNaaRABCDEFGHIJKLMN
      storage_path: /datastorage/registry
      boto_timeout: 120 1
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS: []
DISTRIBUTED_STORAGE_PREFERENCE:
  - default

```

- 1** Optional. The time, in seconds, until a timeout exception is thrown when attempting to read from a connection. The default is **60** seconds. Also encompasses the time, in seconds, until a timeout exception is thrown when attempting to make a connection. The default is **60** seconds.

When starting the registry using cloud storage, no configuration is required on the command line. For example:

```

$ sudo podman run -d --rm -p 80:8080 -p 443:8443 \
  --name=quay \
  -v $QUAY/config:/conf/stack:Z \

```



registry.redhat.io/quay/quay-rhel8:v3.13.1

## CHAPTER 3. CONFIGURATION FIELDS

This section describes the both required and optional configuration fields when deploying Red Hat Quay.

### 3.1. REQUIRED CONFIGURATION FIELDS

The fields required to configure Red Hat Quay are covered in the following sections:

- [General required fields](#)
- [Storage for images](#)
- [Database for metadata](#)
- [Redis for build logs and user events](#)
- [Tag expiration options](#)

### 3.2. AUTOMATION OPTIONS

The following sections describe the available automation options for Red Hat Quay deployments:

- [Pre-configuring Red Hat Quay for automation](#)
- [Using the API to create the first user](#)

### 3.3. OPTIONAL CONFIGURATION FIELDS

Optional fields for Red Hat Quay can be found in the following sections:

- [Basic configuration](#)
- [SSL](#)
- [LDAP](#)
- [Repository mirroring](#)
- [Quota management](#)
- [Security scanner](#)
- [Helm](#)
- [Action log](#)
- [Build logs](#)
- [Dockerfile build](#)
- [OAuth](#)
- [Configuring nested repositories](#)

- [Adding other OCI media types to Quay](#)
- [Mail](#)
- [User](#)
- [Recaptcha](#)
- [ACI](#)
- [JWT](#)
- [App tokens](#)
- [Miscellaneous](#)
- [User interface v2](#)
- [IPv6 configuration field](#)
- [Legacy options](#)

### 3.4. GENERAL REQUIRED FIELDS

The following table describes the required configuration fields for a Red Hat Quay deployment:

**Table 3.1. General required fields**

Field	Type	Description
<b>AUTHENTICATION_TYPE</b> (Required)	String	The authentication engine to use for credential authentication.  <b>Values:</b> One of <b>Database, LDAP, JWT, Keystone, OIDC</b>  <b>Default: Database</b>
<b>PREFERRED_URL_SCHEME</b> (Required)	String	The URL scheme to use when accessing Red Hat Quay.  <b>Values:</b> One of <b>http, https</b>  <b>Default: http</b>
<b>SERVER_HOSTNAME</b> (Required)	String	The URL at which Red Hat Quay is accessible, without the scheme.  <b>Example:</b> <b>quay-server.example.com</b>

Field	Type	Description
<b>DATABASE_SECRET_KEY</b> (Required)	String	Key used to encrypt sensitive fields within the database. This value should never be changed once set, otherwise all reliant fields, for example, repository mirror username and password configurations, are invalidated. This value is set automatically by the Red Hat Quay Operator for Operator-based deployments. For standalone deployments, administrators can provide their own key using Open SSL or a similar tool. Key length should not exceed 63 characters.
<b>SECRET_KEY</b> (Required)	String	Key used to encrypt the session cookie and the CSRF token needed for correct interpretation of the user session. The value should not be changed when set. Should be persistent across all Red Hat Quay instances. If not persistent across all instances, login failures and other errors related to session persistence might occur.
<b>SETUP_COMPLETE</b> (Required)	Boolean	This is an artifact left over from earlier versions of the software and currently it <b>must</b> be specified with a value of <b>true</b> .

## 3.5. DATABASE CONFIGURATION

This section describes the database configuration fields available for Red Hat Quay deployments.

### 3.5.1. Database URI

With Red Hat Quay, connection to the database is configured by using the required **DB\_URI** field.

The following table describes the **DB\_URI** configuration field:

**Table 3.2. Database URI**

Field	Type	Description
<b>DB_URI</b> (Required)	String	The URI for accessing the database, including any credentials.  Example <b>DB_URI</b> field:  <code>postgresql://quayuser:quaypass@quay-server.example.com:5432/quay</code>

### 3.5.2. Database connection arguments

Optional connection arguments are configured by the **DB\_CONNECTION\_ARGS** parameter. Some of the key-value pairs defined under **DB\_CONNECTION\_ARGS** are generic, while others are database specific.

The following table describes database connection arguments:

**Table 3.3. Database connection arguments**

Field	Type	Description
<b>DB_CONNECTION_ARGS</b>	Object	Optional connection arguments for the database, such as timeouts and SSL/TLS.
<b>.autorollback</b>	Boolean	Whether to use thread-local connections. Should always be <b>true</b>
<b>.threadlocals</b>	Boolean	Whether to use auto-rollback connections. Should always be <b>true</b>

#### 3.5.2.1. PostgreSQL SSL/TLS connection arguments

With SSL/TLS, configuration depends on the database you are deploying. The following example shows a PostgreSQL SSL/TLS configuration:

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

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

**Table 3.4. SSL/TLS options**

Mode	Description
<b>disable</b>	Your configuration only tries non-SSL/TLS connections.
<b>allow</b>	Your configuration first tries a non-SSL/TLS connection. Upon failure, tries an SSL/TLS connection.
<b>prefer</b> (Default)	Your configuration first tries an SSL/TLS connection. Upon failure, tries a non-SSL/TLS connection.
<b>require</b>	Your configuration only tries an SSL/TLS connection. If a root CA file is present, it verifies the certificate in the same way as if <code>verify-ca</code> was specified.
<b>verify-ca</b>	Your configuration only tries an SSL/TLS connection, and verifies that the server certificate is issued by a trusted certificate authority (CA).
<b>verify-full</b>	Only tries an SSL/TLS connection, and verifies that the server certificate is issued by a trusted CA and that the requested server hostname matches that in the certificate.

For more information on the valid arguments for PostgreSQL, see [Database Connection Control Functions](#).

### 3.5.2.2. MySQL SSL/TLS connection arguments

The following example shows a sample MySQL SSL/TLS configuration:

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

Information on the valid connection arguments for MySQL is available at [Connecting to the Server Using URI-Like Strings or Key-Value Pairs](#).

## 3.6. IMAGE STORAGE

This section details the image storage features and configuration fields that are available with Red Hat Quay.

### 3.6.1. Image storage features

The following table describes the image storage features for Red Hat Quay:

**Table 3.5. Storage config features**

Field	Type	Description
<b>FEATURE_REPO_MIRROR</b>	Boolean	If set to true, enables repository mirroring.  <b>Default: false</b>
<b>FEATURE_PROXY_STORAGE</b>	Boolean	Whether to proxy all direct download URLs in storage through NGINX.  <b>Default: false</b>
<b>FEATURE_STORAGE_REPLICATION</b>	Boolean	Whether to automatically replicate between storage engines.  <b>Default: false</b>

### 3.6.2. Image storage configuration fields

The following table describes the image storage configuration fields for Red Hat Quay:

**Table 3.6. Storage config fields**

Field	Type	Description
<b>DISTRIBUTED_STORAGE_CONFIG</b> (Required)	Object	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.  <b>Default: []</b>
<b>DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS</b> (Required)	Array of string	The list of storage engine(s) (by ID in <b>DISTRIBUTED_STORAGE_CONFIG</b> ) whose images should be fully replicated, by default, to all other storage engines.

Field	Type	Description
<b>DISTRIBUTED_STORAGE_PREFERENCE</b> (Required)	Array of string	The preferred storage engine(s) (by ID in <b>DISTRIBUTED_STORAGE_CONFIG</b> ) to use. A preferred engine means it is first checked for pulling and images are pushed to it.  <b>Default: false</b>
<b>MAXIMUM_LAYER_SIZE</b>	String	Maximum allowed size of an image layer.  <b>Pattern: <math>^[0-9]+(G M)\$</math></b>  <b>Example: 100G</b>  <b>Default: 20G</b>

### 3.6.3. Local storage

The following YAML shows a sample configuration using local storage:

```
DISTRIBUTED_STORAGE_CONFIG:
  default:
    - LocalStorage
    - storage_path: /datastorage/registry
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS: []
DISTRIBUTED_STORAGE_PREFERENCE:
  - default
```

### 3.6.4. OpenShift Container Storage/NooBaa

The following YAML shows a sample configuration using an OpenShift Container Storage/NooBaa instance:


```
DISTRIBUTED_STORAGE_CONFIG:
  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
      maximum_chunk_size_mb: 100 1
      server_side_assembly: true 2
```



- 1 Defines the maximum chunk size, in MB, for the final copy. Has no effect if **server\_side\_assembly** is set to **false**.
- 2 Optional. Whether Red Hat Quay should try and use server side assembly and the final chunked copy instead of client assembly. Defaults to **true**.

### 3.6.5. Ceph Object Gateway/RadosGW storage

The following YAML shows a sample configuration using Ceph/RadosGW.



#### NOTE

RadosGW is an on-premises S3-compatible storage solution. Note that this differs from general **AWS S3Storage**, which is specifically designed for use with Amazon Web Services S3. This means that RadosGW implements the S3 API and requires credentials like **access\_key**, **secret\_key**, and **bucket\_name**. For more information about Ceph Object Gateway and the S3 API, see [Ceph Object Gateway and the S3 API](#).

#### RadosGW with general s3 access

```
DISTRIBUTED_STORAGE_CONFIG:
  radosGWStorage: 1
  - RadosGWStorage
  - access_key: <access_key_here>
    bucket_name: <bucket_name_here>
    hostname: <hostname_here>
    is_secure: true
    port: '443'
    secret_key: <secret_key_here>
    storage_path: /datastorage/registry
    maximum_chunk_size_mb: 100 2
    server_side_assembly: true 3
```

- 1 Used for general s3 access. Note that general s3 access is not strictly limited to Amazon Web Services (AWS) s3, and can be used with RadosGW or other storage services. For an example of general s3 access using the AWS S3 driver, see "AWS S3 storage".
- 2 Optional. Defines the maximum chunk size in MB for the final copy. Has no effect if **server\_side\_assembly** is set to **false**.
- 3 Optional. Whether Red Hat Quay should try and use server side assembly and the final chunked copy instead of client assembly. Defaults to **true**.

### 3.6.6. AWS S3 storage

The following YAML shows a sample configuration using AWS S3 storage.

```
# ...
DISTRIBUTED_STORAGE_CONFIG:
  default:
    - S3Storage 1
```

```

- host: s3.us-east-2.amazonaws.com
s3_access_key: ABCDEFGHIJKLMN
s3_secret_key: OL3ABCDEFGHIJKLMN
s3_bucket: quay_bucket
s3_region: <region> 2
storage_path: /datastorage/registry
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS: []
DISTRIBUTED_STORAGE_PREFERENCE:
- default
# ...

```

- 1 The **S3Storage** storage driver should only be used for AWS S3 buckets. Note that this differs from general S3 access, where the RadosGW driver or other storage services can be used. For an example, see "Example B: Using RadosGW with general S3 access".
- 2 Optional. The Amazon Web Services region. Defaults to **us-east-1**.

### 3.6.6.1. AWS STS S3 storage

The following YAML shows an example configuration for using Amazon Web Services (AWS) Security Token Service (STS) with Red Hat Quay on OpenShift Container Platform configurations.

```

# ...
DISTRIBUTED_STORAGE_CONFIG:
  default:
    - STSS3Storage
    - sts_role_arn: <role_arn> 1
      s3_bucket: <s3_bucket_name>
      storage_path: <storage_path>
      sts_user_access_key: <s3_user_access_key> 2
      sts_user_secret_key: <s3_user_secret_key> 3
      s3_region: <region> 4
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS: []
DISTRIBUTED_STORAGE_PREFERENCE:
- default
# ...

```

- 1 The unique Amazon Resource Name (ARN).
- 2 The generated AWS S3 user access key.
- 3 The generated AWS S3 user secret key.
- 4 Optional. The Amazon Web Services region. Defaults to **us-east-1**.

### 3.6.7. Google Cloud Storage

The following YAML shows a sample configuration using Google Cloud Storage:

```

DISTRIBUTED_STORAGE_CONFIG:
  googleCloudStorage:
    - GoogleCloudStorage

```

```

- access_key: GOOGQIMFB3ABCDEFGHIJKLMN
  bucket_name: quay-bucket
  secret_key: FhDAYe2HeuAKfvZCAGyOioNaaRABCDEFGHIJKLMN
  storage_path: /datastorage/registry
  boto_timeout: 120 1
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS: []
DISTRIBUTED_STORAGE_PREFERENCE:
- googleCloudStorage

```

- <sup>1</sup> Optional. The time, in seconds, until a timeout exception is thrown when attempting to read from a connection. The default is **60** seconds. Also encompasses the time, in seconds, until a timeout exception is thrown when attempting to make a connection. The default is **60** seconds.

### 3.6.8. Azure Storage

The following YAML shows a sample configuration using Azure Storage:

```

DISTRIBUTED_STORAGE_CONFIG:
  azureStorage:
    - AzureStorage
    - azure_account_name: azure_account_name_here
      azure_container: azure_container_here
      storage_path: /datastorage/registry
      azure_account_key: azure_account_key_here
      sas_token: some/path/
      endpoint_url: https://[account-name].blob.core.usgovcloudapi.net 1
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS: []
DISTRIBUTED_STORAGE_PREFERENCE:
- azureStorage

```

- <sup>1</sup> The **endpoint\_url** parameter for Azure storage is optional and can be used with Microsoft Azure Government (MAG) endpoints. If left blank, the **endpoint\_url** will connect to the normal Azure region.

As of Red Hat Quay 3.7, you must use the Primary endpoint of your MAG Blob service. Using the Secondary endpoint of your MAG Blob service will result in the following error:

**AuthenticationErrorDetail:Cannot find the claimed account when trying to GetProperties for the account whusc8-secondary.**

### 3.6.9. Swift storage

The following YAML shows a sample configuration using Swift storage:

```

DISTRIBUTED_STORAGE_CONFIG:
  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: 3
      os_options:
        tenant_id: <osp_tenant_id_here>

```

```

    user_domain_name: <osp_domain_name_here>
    ca_cert_path: /conf/stack/swift.cert"
    storage_path: /datastorage/registry
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS: []
DISTRIBUTED_STORAGE_PREFERENCE:
  - swiftStorage

```

### 3.6.10. Nutanix object storage

The following YAML shows a sample configuration using Nutanix object storage.

```

DISTRIBUTED_STORAGE_CONFIG:
  nutanixStorage: #storage config name
  - RadosGWStorage #actual driver
  - access_key: access_key_here #parameters
    secret_key: secret_key_here
    bucket_name: bucket_name_here
    hostname: hostname_here
    is_secure: 'true'
    port: '443'
    storage_path: /datastorage/registry
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS: []
DISTRIBUTED_STORAGE_PREFERENCE: #must contain name of the storage config
  - nutanixStorage

```

### 3.6.11. IBM Cloud object storage

The following YAML shows a sample configuration using IBM Cloud object storage.

```

DISTRIBUTED_STORAGE_CONFIG:
  default:
  - IBMCloudStorage #actual driver
  - access_key: <access_key_here> #parameters
    secret_key: <secret_key_here>
    bucket_name: <bucket_name_here>
    hostname: <hostname_here>
    is_secure: 'true'
    port: '443'
    storage_path: /datastorage/registry
    maximum_chunk_size_mb: 100mb 1
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS:
  - default
DISTRIBUTED_STORAGE_PREFERENCE:
  - default

```

**1** Optional. Recommended to be set to **100mb**.

### 3.6.12. NetApp ONTAP S3 object storage

The following YAML shows a sample configuration using NetApp ONTAP S3.

```

DISTRIBUTED_STORAGE_CONFIG:

```

```

local_us:
- RadosGWStorage
- access_key: <access_key>
  bucket_name: <bucket_name>
  hostname: <host_url_address>
  is_secure: true
  port: <port>
  secret_key: <secret_key>
  storage_path: /datastorage/registry
  signature_version: v4
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS:
- local_us
DISTRIBUTED_STORAGE_PREFERENCE:
- local_us

```

### 3.6.13. Hitachi Content Platform object storage

The following YAML shows a sample configuration using HCP for object storage.

#### Example HCP storage configuration

```

DISTRIBUTED_STORAGE_CONFIG:
  hcp_us:
  - RadosGWStorage
  - access_key: <access_key>
    bucket_name: <bucket_name>
    hostname: <hitachi_hostname_example>
    is_secure: true
    secret_key: <secret_key>
    storage_path: /datastorage/registry
    signature_version: v4
DISTRIBUTED_STORAGE_DEFAULT_LOCATIONS:
- hcp_us
DISTRIBUTED_STORAGE_PREFERENCE:
- hcp_us

```

## 3.7. REDIS CONFIGURATION FIELDS

This section details the configuration fields available for Redis deployments.

### 3.7.1. Build logs

The following build logs configuration fields are available for Redis deployments:

**Table 3.7. Build logs configuration**

Field	Type	Description
<b>BUILDLGOS_REDIS</b> (Required)	Object	Redis connection details for build logs caching.

Field	Type	Description
<b>.host</b> (Required)	String	The hostname at which Redis is accessible. <b>Example:</b> <b>quay-server.example.com</b>
<b>.port</b> (Required)	Number	The port at which Redis is accessible. <b>Example:</b> <b>6379</b>
<b>.password</b>	String	The password to connect to the Redis instance. <b>Example:</b> <b>strongpassword</b>
<b>.ssl</b> (Optional)	Boolean	Whether to enable TLS communication between Redis and Quay. Defaults to false.

### 3.7.2. User events

The following user event fields are available for Redis deployments:

**Table 3.8. User events config**

Field	Type	Description
<b>USER_EVENTS_REDIS</b> (Required)	Object	Redis connection details for user event handling.
<b>.host</b> (Required)	String	The hostname at which Redis is accessible. <b>Example:</b> <b>quay-server.example.com</b>
<b>.port</b> (Required)	Number	The port at which Redis is accessible. <b>Example:</b> <b>6379</b>
<b>.password</b>	String	The password to connect to the Redis instance. <b>Example:</b> <b>strongpassword</b>
<b>.ssl</b>	Boolean	Whether to enable TLS communication between Redis and Quay. Defaults to false.

Field	Type	Description
<code>.ssl_keyfile</code> (Optional)	String	The name of the key database file, which houses the client certificate to be used. <b>Example:</b> <b>ssl_keyfile:</b> <b>/path/to/server/privatekey.pem</b>
<code>.ssl_certfile</code> (Optional)	String	Used for specifying the file path of the SSL certificate. <b>Example:</b> <b>ssl_certfile:</b> <b>/path/to/server/certificate.pem</b>
<code>.ssl_cert_reqs</code> (Optional)	String	Used to specify the level of certificate validation to be performed during the SSL/TLS handshake. <b>Example:</b> <b>ssl_cert_reqs:</b> <b>CERT_REQUIRED</b>
<code>.ssl_ca_certs</code> (Optional)	String	Used to specify the path to a file containing a list of trusted Certificate Authority (CA) certificates. <b>Example:</b> <b>ssl_ca_certs:</b> <b>/path/to/ca_certs.pem</b>
<code>.ssl_ca_data</code> (Optional)	String	Used to specify a string containing the trusted CA certificates in PEM format. <b>Example:</b> <b>ssl_ca_data: &lt;certificate&gt;</b>
<code>.ssl_check_hostname</code> (Optional)	Boolean	Used when setting up an SSL/TLS connection to a server. It specifies whether the client should check that the hostname in the server's SSL/TLS certificate matches the hostname of the server it is connecting to. <b>Example:</b> <b>ssl_check_hostname: true</b>

### 3.7.3. Example Redis configuration

The following YAML shows a sample configuration using Redis with optional SSL/TLS fields:

```
BUILDLOGS_REDIS:
  host: quay-server.example.com
  password: strongpassword
  port: 6379
  ssl: true
```

```
USER_EVENTS_REDIS:
  host: quay-server.example.com
  password: strongpassword
  port: 6379
  ssl: true
  ssl_*: <path_location_or_certificate>
```



#### NOTE

If your deployment uses Azure Cache for Redis and **ssl** is set to **true**, the port defaults to **6380**.

## 3.8. MODELCACHE CONFIGURATION OPTIONS

The following options are available on Red Hat Quay for configuring ModelCache.

### 3.8.1. Memcache configuration option

Memcache is the default ModelCache configuration option. With Memcache, no additional configuration is necessary.

### 3.8.2. Single Redis configuration option

The following configuration is for a single Redis instance with optional read-only replicas:

```
DATA_MODEL_CACHE_CONFIG:
  engine: redis
  redis_config:
    primary:
      host: <host>
      port: <port>
      password: <password if ssl is true>
      ssl: <true | false >
    replica:
      host: <host>
      port: <port>
      password: <password if ssl is true>
      ssl: <true | false >
```

### 3.8.3. Clustered Redis configuration option

Use the following configuration for a clustered Redis instance:

```
DATA_MODEL_CACHE_CONFIG:
  engine: rediscluster
  redis_config:
```



```

startup_nodes:
- host: <cluster-host>
  port: <port>
password: <password if ssl: true>
read_from_replicas: <true|false>
skip_full_coverage_check: <true | false>
ssl: <true | false >

```

### 3.9. TAG EXPIRATION CONFIGURATION FIELDS

The following tag expiration configuration fields are available with Red Hat Quay:

Table 3.9. Tag expiration configuration fields

Field	Type	Description
<b>FEATURE_GARBAGE_COLLECTION</b>	Boolean	Whether garbage collection of repositories is enabled.  <b>Default:</b> True
<b>TAG_EXPIRATION_OPTIONS</b> (Required)	Array of string	If enabled, the options that users can select for expiration of tags in their namespace.  <b>Pattern:</b> <b><code>^[0-9]+(y w m d h s)\$</code></b>
<b>DEFAULT_TAG_EXPIRATION</b> (Required)	String	The default, configurable tag expiration time for time machine.  <b>Pattern:</b> <b><code>^[0-9]+(w m d h s)\$</code></b> <b>Default:</b> <b>2w</b>
<b>FEATURE_CHANGE_TAG_EXPIRATION</b>	Boolean	Whether users and organizations are allowed to change the tag expiration for tags in their namespace.  <b>Default:</b> True
<b>FEATURE_AUTO_PRUNE</b>	Boolean	When set to <b>True</b> , enables functionality related to the auto-pruning of tags. <b>Default:</b> <b>False</b>
<b>NOTIFICATION_TASK_RUN_MINIMUM_INTERVAL_MINUTES</b>	Integer	The interval, in minutes, that defines the frequency to re-run notifications for expiring images.  <b>Default:</b> <b>300</b>

Field	Type	Description
DEFAULT_NAMESPACE_AUTOPRUNE_POLICY	Object	The default organization-wide auto-prune policy.
.method: number_of_tags	Object	The option specifying the number of tags to keep.
.value: <integer>	Integer	When used with <b>method: number_of_tags</b> , denotes the number of tags to keep.  For example, to keep two tags, specify <b>2</b> .
.creation_date	Object	The option specifying the duration of which to keep tags.
.value: <integer>	Integer	When used with <b>creation_date</b> , denotes how long to keep tags.  Can be set to seconds ( <b>s</b> ), days ( <b>d</b> ), months ( <b>m</b> ), weeks ( <b>w</b> ), or years ( <b>y</b> ). Must include a valid integer. For example, to keep tags for one year, specify <b>1y</b> .
AUTO_PRUNING_DEFAULT_POLICY_POLL_PERIOD	Integer	The period in which the auto-pruner worker runs at the registry level. By default, it is set to run one time per day (one time per 24 hours). Value must be in seconds.

### 3.9.1. Example tag expiration configuration

The following YAML example shows you a sample tag expiration configuration.

```
# ...
DEFAULT_TAG_EXPIRATION: 2w
TAG_EXPIRATION_OPTIONS:
  - 0s
  - 1d
  - 1w
  - 2w
  - 4w
  - 3y
# ...
```

### 3.9.2. Registry-wide auto-prune policies examples

The following YAML examples show you registry-wide auto-pruning examples by both number of tags and creation date.

### Example registry auto-prune policy by number of tags

```
# ...
DEFAULT_NAMESPACE_AUTOPRUNE_POLICY:
  method: number_of_tags
  value: 10 1
# ...
```

**1** In this scenario, ten tags remain.

### Example registry auto-prune policy by creation date

```
# ...
DEFAULT_NAMESPACE_AUTOPRUNE_POLICY:
  method: creation_date
  value: 1y
# ...
```

## 3.10. QUOTA MANAGEMENT CONFIGURATION FIELDS

Table 3.10. Quota management configuration

Field	Type	Description
FEATURE_QUOTA_MANAGEMENT	Boolean	Enables configuration, caching, and validation for quota management feature.  <b>**Default:**</b> `False`
DEFAULT_SYSTEM_REJECT_QUOTA_BYTES	String	Enables system default quota reject byte allowance for all organizations.  By default, no limit is set.
QUOTA_BACKFILL	Boolean	Enables the quota backfill worker to calculate the size of pre-existing blobs.  <b>Default: True</b>

Field	Type	Description
QUOTA_TOTAL_DELAY_SECONDS	String	The time delay for starting the quota backfill. Rolling deployments can cause incorrect totals. This field <b>must</b> be set to a time longer than it takes for the rolling deployment to complete.  <b>Default: 1800</b>
PERMANENTLY_DELETE_TAGS	Boolean	Enables functionality related to the removal of tags from the time machine window.  <b>Default: False</b>
RESET_CHILD_MANIFEST_EXPIRATION	Boolean	Resets the expirations of temporary tags targeting the child manifests. With this feature set to <b>True</b> , child manifests are immediately garbage collected.  <b>Default: False</b>

### 3.10.1. Example quota management configuration

The following YAML is the suggested configuration when enabling quota management.

#### Quota management YAML configuration

```
FEATURE_QUOTA_MANAGEMENT: true
FEATURE_GARBAGE_COLLECTION: true
PERMANENTLY_DELETE_TAGS: true
QUOTA_TOTAL_DELAY_SECONDS: 1800
RESET_CHILD_MANIFEST_EXPIRATION: true
```

## 3.11. PROXY CACHE CONFIGURATION FIELDS

Table 3.11. Proxy configuration

Field	Type	Description
FEATURE_PROXY_CACHE	Boolean	Enables Red Hat Quay to act as a pull through cache for upstream registries.  <b>Default: false</b>

## 3.12. ROBOT ACCOUNT CONFIGURATION FIELDS

Table 3.12. Robot account configuration fields

Field	Type	Description
ROBOTS_DISALLOW	Boolean	When set to <b>true</b> , robot accounts are prevented from all interactions, as well as from being created <b>Default: False</b>

## 3.13. PRE-CONFIGURING RED HAT QUAY FOR AUTOMATION

Red Hat Quay supports several configuration options that enable automation. Users can configure these options before deployment to reduce the need for interaction with the user interface.

### 3.13.1. Allowing the API to create the first user

To create the first user, users need to set the **FEATURE\_USER\_INITIALIZE** parameter to **true** and call the `/api/v1/user/initialize` API. Unlike all other registry API calls that require an OAuth token generated by an OAuth application in an existing organization, the API endpoint does not require authentication.

Users can use the API to create a user such as **quayadmin** after deploying Red Hat Quay, provided no other users have been created. For more information, see [Using the API to create the first user](#).

### 3.13.2. Enabling general API access

Users should set the **BROWSER\_API\_CALLS\_XHR\_ONLY** configuration option to **false** to allow general access to the Red Hat Quay registry API.

### 3.13.3. Adding a superuser

After deploying Red Hat Quay, users can create a user and give the first user administrator privileges with full permissions. Users can configure full permissions in advance by using the **SUPER\_USER** configuration object. For example:

```
# ...
SERVER_HOSTNAME: quay-server.example.com
SETUP_COMPLETE: true
SUPER_USERS:
  - quayadmin
# ...
```

### 3.13.4. Restricting user creation

After you have configured a superuser, you can restrict the ability to create new users to the superuser group by setting the **FEATURE\_USER\_CREATION** to **false**. For example:

```
# ...
FEATURE_USER_INITIALIZE: true
BROWSER_API_CALLS_XHR_ONLY: false
SUPER_USERS:
```

```
- quayadmin
FEATURE_USER_CREATION: false
# ...
```

### 3.13.5. Enabling new functionality in Red Hat Quay 3.13

To use new Red Hat Quay 3.13 functions, enable some or all of the following features:

```
# ...
FEATURE_UI_V2: true
FEATURE_UI_V2_REPO_SETTINGS: true
FEATURE_AUTO_PRUNE: true
ROBOTS_DISALLOW: false
# ...
```

### 3.13.6. Suggested configuration for automation

The following **config.yaml** parameters are suggested for automation:

```
# ...
FEATURE_USER_INITIALIZE: true
BROWSER_API_CALLS_XHR_ONLY: false
SUPER_USERS:
- quayadmin
FEATURE_USER_CREATION: false
# ...
```

### 3.13.7. Deploying the Red Hat Quay Operator using the initial configuration

Use the following procedure to deploy Red Hat Quay on OpenShift Container Platform using the initial configuration.

#### Prerequisites

- You have installed the **oc** CLI.

#### Procedure

1. Create a secret using the configuration file:

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

2. Create a **quayregistry.yaml** file. Identify the unmanaged components and reference the created secret, for example:

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

```
namespace: quay-enterprise
spec:
  configBundleSecret: init-config-bundle-secret
```

3. Deploy the Red Hat Quay registry:

```
$ oc create -n quay-enterprise -f quayregistry.yaml
```

## Next Steps

- [Using the API to create the first user](#)

### 3.13.8. Using the API to create the first user

Use the following procedure to create the first user in your Red Hat Quay organization.

#### Prerequisites

- The config option **FEATURE\_USER\_INITIALIZE** must be set to **true**.
- No users can already exist in the database.



#### PROCEDURE

This procedure requests an OAuth token by specifying "**access\_token**": **true**.

1. Open your Red Hat Quay configuration file and update the following configuration fields:

```
FEATURE_USER_INITIALIZE: true
SUPER_USERS:
  - quayadmin
```

2. Stop the Red Hat Quay service by entering the following command:

```
$ sudo podman stop quay
```

3. Start the Red Hat Quay service by entering the following command:

```
$ sudo podman run -d -p 80:8080 -p 443:8443 --name=quay -v $QUAY/config:/conf/stack:Z
-v $QUAY/storage:/datastorage:Z {productrepo}/{quayimage}:{productminv}
```

4. Run the following **CURL** command to generate a new user with a username, password, email, and access token:

```
$ curl -X POST -k http://quay-server.example.com/api/v1/user/initialize --header 'Content-Type: application/json' --data '{"username": "quayadmin", "password": "quaypass12345", "email": "quayadmin@example.com", "access_token": true}'
```

If successful, the command returns an object with the username, email, and encrypted password. For example:

```
{"access_token": "6B4QTRSTSD1HMIG915VPX7BMEZBVB9GPNY2FC2ED",
```

```
"email":"quayadmin@example.com","encrypted_password":"1nZMLH57RIE5UGdL/yYpDOHL
qiNCgimb6W9kfF8MjZ1xrfDpRyRs9NUnUuNuAitW","username":"quayadmin"} #
gitleaks:allow
```

If a user already exists in the database, an error is returned:

```
{"message":"Cannot initialize user in a non-empty database"}
```

If your password is not at least eight characters or contains whitespace, an error is returned:

```
{"message":"Failed to initialize user: Invalid password, password must be at least 8
characters and contain no whitespace."}
```

5. Log in to your Red Hat Quay deployment by entering the following command:

```
$ sudo podman login -u quayadmin -p quaypass12345 http://quay-server.example.com --tls-
verify=false
```

### Example output

```
Login Succeeded!
```

#### 3.13.8.1. Using the OAuth token

After invoking the API, you can call out the rest of the Red Hat Quay API by specifying the returned OAuth code.

#### Prerequisites

- You have invoked the `/api/v1/user/initialize` API, and passed in the username, password, and email address.

#### Procedure

- Obtain the list of current users by entering the following command:

```
$ curl -X GET -k -H "Authorization: Bearer
6B4QTRSTSD1HMIG915VPX7BMEZBVB9GPNY2FC2ED" https://example-registry-quay-
quay-enterprise.apps.docs.quayteam.org/api/v1/superuser/users/
```

Example output:

```
{
  "users": [
    {
      "kind": "user",
      "name": "quayadmin",
      "username": "quayadmin",
      "email": "quayadmin@example.com",
      "verified": true,
      "avatar": {
        "name": "quayadmin",
        "hash":
```



```

"3e82e9cbf62d25dec0ed1b4c66ca7c5d47ab9f1f271958298dea856fb26adc4c",
  "color": "#e7ba52",
  "kind": "user"
},
"super_user": true,
"enabled": true
}
]
}

```

In this instance, the details for the **quayadmin** user are returned as it is the only user that has been created so far.

### 3.13.8.2. Using the API to create an organization

The following procedure details how to use the API to create a Red Hat Quay organization.

#### Prerequisites

- You have invoked the **/api/v1/user/initialize** API, and passed in the username, password, and email address.
- You have called out the rest of the Red Hat Quay API by specifying the returned OAuth code.

#### Procedure

1. To create an organization, use a POST call to **api/v1/organization/** endpoint:

```

$ curl -X POST -k --header 'Content-Type: application/json' -H "Authorization: Bearer
6B4QTRSTSD1HMIG915VPX7BMEZBVB9GPNY2FC2ED" https://example-registry-quay-
quay-enterprise.apps.docs.quayteam.org/api/v1/organization/ --data '{"name": "testorg",
"email": "testorg@example.com"}'

```

Example output:

```
"Created"
```

2. You can retrieve the details of the organization you created by entering the following command:

```

$ curl -X GET -k --header 'Content-Type: application/json' -H "Authorization: Bearer
6B4QTRSTSD1HMIG915VPX7BMEZBVB9GPNY2FC2ED" https://min-registry-quay-quay-
enterprise.apps.docs.quayteam.org/api/v1/organization/testorg

```

Example output:

```

{
  "name": "testorg",
  "email": "testorg@example.com",
  "avatar": {
    "name": "testorg",
    "hash": "5f113632ad532fc78215c9258a4fb60606d1fa386c91b141116a1317bf9c53c8",
    "color": "#a55194",
    "kind": "user"
  },
}

```

```

"is_admin": true,
"is_member": true,
"teams": {
  "owners": {
    "name": "owners",
    "description": "",
    "role": "admin",
    "avatar": {
      "name": "owners",
      "hash":
"6f0e3a8c0eb46e8834b43b03374ece43a030621d92a7437beb48f871e90f8d90",
      "color": "#c7c7c7",
      "kind": "team"
    },
    "can_view": true,
    "repo_count": 0,
    "member_count": 1,
    "is_synced": false
  }
},
"ordered_teams": [
  "owners"
],
"invoice_email": false,
"invoice_email_address": null,
"tag_expiration_s": 1209600,
"is_free_account": true
}

```

### 3.14. BASIC CONFIGURATION FIELDS

Table 3.13. Basic configuration

Field	Type	Description
REGISTRY_TITLE	String	If specified, the long-form title for the registry. Displayed in frontend of your Red Hat Quay deployment, for example, at the sign in page of your organization. Should not exceed 35 characters. <b>Default:</b> <b>Red Hat Quay</b>
REGISTRY_TITLE_SHORT	String	If specified, the short-form title for the registry. Title is displayed on various pages of your organization, for example, as the title of the tutorial on your organization's <b>Tutorial</b> page. <b>Default:</b> <b>Red Hat Quay</b>

Field	Type	Description
CONTACT_INFO	Array of String	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.
[0]	String	Adds a link to send an e-mail.  <b>Pattern:</b> <b><code>^mailto:(.)+\$</code></b> <b>Example:</b> <b><code>mailto:support@quay.io</code></b>
[1]	String	Adds a link to visit an IRC chat room.  <b>Pattern:</b> <b><code>^irc://(.)+\$</code></b> <b>Example:</b> <b><code>irc://chat.freenode.net:6665/quay</code></b>
[2]	String	Adds a link to call a phone number.  <b>Pattern:</b> <b><code>^tel:(.)+\$</code></b> <b>Example:</b> <b><code>tel:+1-888-930-3475</code></b>
[3]	String	Adds a link to a defined URL.  <b>Pattern:</b> <b><code>^http(s)?://(.)+\$</code></b> <b>Example:</b> <b><code>https://twitter.com/quayio</code></b>

### 3.15. SSL CONFIGURATION FIELDS

Table 3.14. SSL configuration

Field	Type	Description
-------	------	-------------

Field	Type	Description
<b>PREFERRED_URL_SCHEME</b>	String	One of <b>http</b> or <b>https</b> . Note that users only set their <b>PREFERRED_URL_SCHEME</b> to <b>http</b> when there is no TLS encryption in the communication path from the client to Quay. Users must set their <b>PREFERRED_URL_SCHEME</b> to <b>https</b> when using a TLS-terminating load balancer, a reverse proxy (for example, Nginx), or when using Quay with custom SSL certificates directly. In most cases, the <b>PREFERRED_URL_SCHEME</b> should be <b>https</b> . <b>Default: http</b>
<b>SERVER_HOSTNAME</b> (Required)	String	The URL at which Red Hat Quay is accessible, without the scheme  <b>Example:</b> <b>quay-server.example.com</b>

Field	Type	Description
SSL_CIPHERS	Array of String	<p>If specified, the nginx-defined list of SSL ciphers to enabled and disabled</p> <p>Example:  <b>[ECDHE-RSA-AES128-GCM-SHA256, ECDHE-ECDSA-AES128-GCM-SHA256, ECDHE-RSA-AES256-GCM-SHA384, ECDHE-ECDSA-AES256-GCM-SHA384, DHE-RSA-AES128-GCM-SHA256, DHE-DSS-AES128-GCM-SHA256, kEDH+AESGCM, ECDHE-RSA-AES128-SHA256, ECDHE-ECDSA-AES128-SHA256, ECDHE-RSA-AES128-SHA, ECDHE-ECDSA-AES128-SHA, ECDHE-RSA-AES256-SHA384, ECDHE-ECDSA-AES256-SHA384, ECDHE-RSA-AES256-SHA, ECDHE-ECDSA-AES256-SHA, DHE-RSA-AES128-SHA256, DHE-RSA-AES128-SHA, DHE-DSS-AES128-SHA256, DHE-RSA-AES256-SHA256, DHE-DSS-AES256-SHA, DHE-DSS-AES256-SHA, AES128-GCM-SHA256, AES256-GCM-SHA384, AES128-SHA256, AES256-SHA256, AES128-SHA, AES256-SHA, AES, !3DES", !aNULL, !eNULL, !EXPORT, DES, !RC4, MD5, !PSK, !aECDH, !EDH-DSS-DES-CBC3-SHA, !EDH-RSA-DES-CBC3-SHA, !KRB5-DES-CBC3-SHA]</b></p>

Field	Type	Description
<b>SSL_PROTOCOLS</b>	Array of String	<p>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.</p> <p><b>Example:</b>  <code>['TLSv1','TLSv1.1','TLSv1.2','TLSv1.3']`</code></p>
<b>SESSION_COOKIE_SECURE</b>	Boolean	<p>Whether the <b>secure</b> property should be set on session cookies</p> <p><b>Default:</b> False</p> <p><b>Recommendation:</b> Set to True for all installations using SSL</p>

### 3.15.1. Configuring SSL

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

## 3.16. ADDING ADDITIONAL CERTIFICATE AUTHORITIES TO THE RED HAT QUAY CONTAINER

The **extra\_ca\_certs** directory is the directory where additional Certificate Authorities (CAs) can be stored to extend the set of trusted certificates. These certificates are used by Red Hat Quay to verify SSL/TLS connections with external services. When deploying Red Hat Quay, you can place the necessary CAs in this directory to ensure that connections to services like LDAP, OIDC, and storage systems are properly secured and validated.

For standalone Red Hat Quay deployments, you must create this directory and copy the additional CA certificates into that directory.

## Prerequisites

- You have a CA for the desired service.

## Procedure

1. View the certificate to be added to the container by entering the following command:

```
$ cat storage.crt
```

### Example output

```
-----BEGIN CERTIFICATE-----
MIIDTTCCAjWgAwIBAgIJAMVr9ngjJhzbMA0GCSqGSIb3DQEBCwUAMD0xCzAJBgNV...
-----END CERTIFICATE-----
```

2. Create the **extra\_ca\_certs** in the **/config** folder of your Red Hat Quay directory by entering the following command:

```
$ mkdir -p /path/to/quay_config_folder/extra_ca_certs
```

3. Copy the CA file to the **extra\_ca\_certs** folder. For example:

```
$ cp storage.crt /path/to/quay_config_folder/extra_ca_certs/
```

4. Ensure that the **storage.crt** file exists within the **extra\_ca\_certs** folder by entering the following command:

```
$ tree /path/to/quay_config_folder/extra_ca_certs
```

### Example output

```
/path/to/quay_config_folder/extra_ca_certs
├── storage.crt----
```

5. Obtain the **CONTAINER ID** of your **Quay** consider by entering the following command:

```
$ podman ps
```

### Example output

CONTAINER ID	IMAGE	COMMAND	CREATED
STATUS	PORTS		
5a3e82c4a75f	<registry>/<repo>/quay:{productminv}	"/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	

6. Restart the container by entering the following command

```
$ podman restart 5a3e82c4a75f
```

- Confirm that the certificate was copied into the container namespace by running the following command:

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

### Example output

```
-----BEGIN CERTIFICATE-----
MIIDTTCCAjWgAwIBAgIJAMVr9ngjJhzbMA0GCSqGSIb3DQEBCwUAMD0xCzAJBgNV...
-----END CERTIFICATE-----
```

## 3.17. LDAP CONFIGURATION FIELDS

Table 3.15. LDAP configuration

Field	Type	Description
<b>AUTHENTICATION_TYPE</b> (Required)	String	Must be set to <b>LDAP</b> .
<b>FEATURE_TEAM_SYNCING</b>	Boolean	Whether to allow for team membership to be synced from a backing group in the authentication engine (OIDC, LDAP, or Keystone).  <b>Default: true</b>
<b>FEATURE_NONSUPERUSER_TEAM_SYNCING_SETUP</b>	Boolean	If enabled, non-superusers can setup team synchronization.  <b>Default: false</b>
<b>LDAP_ADMIN_DN</b>	String	The admin DN for LDAP authentication.
<b>LDAP_ADMIN_PASSWD</b>	String	The admin password for LDAP authentication.
<b>LDAP_ALLOW_INSECURE_FALLBACK</b>	Boolean	Whether or not to allow SSL insecure fallback for LDAP authentication.
<b>LDAP_BASE_DN</b>	Array of String	The base DN for LDAP authentication.
<b>LDAP_EMAIL_ATTR</b>	String	The email attribute for LDAP authentication.
<b>LDAP_UID_ATTR</b>	String	The uid attribute for LDAP authentication.



Field	Type	Description
<b>LDAP_URI</b>	String	The LDAP URI.
<b>LDAP_USER_FILTER</b>	String	The user filter for LDAP authentication.
<b>LDAP_USER_RDN</b>	Array of String	The user RDN for LDAP authentication.
<b>LDAP_SECONDARY_USER_RDNS</b>	Array of String	Provide Secondary User Relative DNs if there are multiple Organizational Units where user objects are located.
<b>TEAM_RESYNC_STALE_TIME</b>	String	<p>If team syncing is enabled for a team, how often to check its membership and resync if necessary.</p> <p><b>Pattern:</b>  <code>^[0-9]+(w m d h s)\$</code></p> <p><b>Example:</b>  <b>2h</b></p> <p><b>Default:</b>  <b>30m</b></p>
<b>LDAP_SUPERUSER_FILTER</b>	String	<p>Subset of the <b>LDAP_USER_FILTER</b> configuration field. When configured, allows Red Hat Quay administrators the ability to configure Lightweight Directory Access Protocol (LDAP) users as superusers when Red Hat Quay uses LDAP as its authentication provider.</p> <p>With this field, administrators can add or remove superusers without having to update the Red Hat Quay configuration file and restart their deployment.</p> <p>This field requires that your <b>AUTHENTICATION_TYPE</b> is set to <b>LDAP</b>.</p>
<b>GLOBAL_READONLY_SUPER_USERS</b>	String	When set, grants users of this list read access to all repositories, regardless of whether they are public repositories. Only works for those superusers defined with the <b>LDAP_SUPERUSER_FILTER</b> configuration field.

Field	Type	Description
<code>LDAP_RESTRICTED_USER_FILTER</code>	String	<p>Subset of the <b>LDAP_USER_FILTER</b> configuration field. When configured, allows Red Hat Quay administrators the ability to configure Lightweight Directory Access Protocol (LDAP) users as restricted users when Red Hat Quay uses LDAP as its authentication provider.</p> <p>This field requires that your <b>AUTHENTICATION_TYPE</b> is set to <b>LDAP</b>.</p>
<code>FEATURE_RESTRICTED_USERS</code>	Boolean	<p>When set to <b>True</b> with <b>LDAP_RESTRICTED_USER_FILTER</b> active, only the listed users in the defined LDAP group are restricted.</p> <p><b>Default: False</b></p>
<code>LDAP_TIMEOUT</code>	Integer	<p>Specifies the time limit, in seconds, for LDAP operations. This limits the amount of time an LDAP search, bind, or other operation can take. Similar to the <b>-l</b> option in <b>ldapsearch</b>, it sets a client-side operation timeout.</p> <p><b>Default: 10</b></p>
<code>LDAP_NETWORK_TIMEOUT</code>	Integer	<p>Specifies the time limit, in seconds, for establishing a connection to the LDAP server. This is the maximum time Red Hat Quay waits for a response during network operations, similar to the <b>-o nettimeout</b> option in <b>ldapsearch</b>.</p> <p><b>Default: 10</b></p>

### 3.17.1. LDAP configuration references

Use the following references to update your **config.yaml** file with the desired LDAP settings.

#### 3.17.1.1. Basic LDAP configuration

Use the following reference for a basic LDAP configuration.

```
---
AUTHENTICATION_TYPE: LDAP 1
---
```

```

LDAP_ADMIN_DN: uid=<name>,ou=Users,o=<organization_id>,dc=
<example_domain_component>,dc=com 2
LDAP_ADMIN_PASSWD: ABC123 3
LDAP_ALLOW_INSECURE_FALLBACK: false 4
LDAP_BASE_DN: 5
- dc=example
- dc=com
LDAP_EMAIL_ATTR: mail 6
LDAP_UID_ATTR: uid 7
LDAP_URI: ldap://<example_url>.com 8
LDAP_USER_FILTER: (memberof=cn=developers,ou=Users,dc=<domain_name>,dc=com) 9
LDAP_USER_RDN: 10
- ou=people
LDAP_SECONDARY_USER_RDNS: 11
- ou=<example_organization_unit_one>
- ou=<example_organization_unit_two>
- ou=<example_organization_unit_three>
- ou=<example_organization_unit_four>

```

- 1 Required. Must be set to **LDAP**.
- 2 Required. The admin DN for LDAP authentication.
- 3 Required. The admin password for LDAP authentication.
- 4 Required. Whether to allow SSL/TLS insecure fallback for LDAP authentication.
- 5 Required. The base DN for LDAP authentication.
- 6 Required. The email attribute for LDAP authentication.
- 7 Required. The UID attribute for LDAP authentication.
- 8 Required. The LDAP URI.
- 9 Required. The user filter for LDAP authentication.
- 10 Required. The user RDN for LDAP authentication.
- 11 Optional. Secondary User Relative DN's if there are multiple Organizational Units where user objects are located.

### 3.17.1.2. LDAP restricted user configuration

Use the following reference for an LDAP restricted user configuration.

```

# ...
AUTHENTICATION_TYPE: LDAP
# ...
FEATURE_RESTRICTED_USERS: true 1
# ...
LDAP_ADMIN_DN: uid=<name>,ou=Users,o=<organization_id>,dc=
<example_domain_component>,dc=com
LDAP_ADMIN_PASSWD: ABC123

```

```

LDAP_ALLOW_INSECURE_FALLBACK: false
LDAP_BASE_DN:
  - o=<organization_id>
  - dc=<example_domain_component>
  - dc=com
LDAP_EMAIL_ATTR: mail
LDAP_UID_ATTR: uid
LDAP_URI: ldap://<example_url>.com
LDAP_USER_FILTER: (memberof=cn=developers,ou=Users,o=<example_organization_unit>,dc=
<example_domain_component>,dc=com)
LDAP_RESTRICTED_USER_FILTER: (<filterField>=<value>) 2
LDAP_USER_RDN:
  - ou=<example_organization_unit>
  - o=<organization_id>
  - dc=<example_domain_component>
  - dc=com
# ...

```

- 1 Must be set to **true** when configuring an LDAP restricted user.
- 2 Configures specified users as restricted users.

### 3.17.1.3. LDAP superuser configuration reference

Use the following reference for an LDAP superuser configuration.

```

# ...
AUTHENTICATION_TYPE: LDAP
# ...
LDAP_ADMIN_DN: uid=<name>,ou=Users,o=<organization_id>,dc=
<example_domain_component>,dc=com
LDAP_ADMIN_PASSWD: ABC123
LDAP_ALLOW_INSECURE_FALLBACK: false
LDAP_BASE_DN:
  - o=<organization_id>
  - dc=<example_domain_component>
  - dc=com
LDAP_EMAIL_ATTR: mail
LDAP_UID_ATTR: uid
LDAP_URI: ldap://<example_url>.com
LDAP_USER_FILTER: (memberof=cn=developers,ou=Users,o=<example_organization_unit>,dc=
<example_domain_component>,dc=com)
LDAP_SUPERUSER_FILTER: (<filterField>=<value>) 1
LDAP_USER_RDN:
  - ou=<example_organization_unit>
  - o=<organization_id>
  - dc=<example_domain_component>
  - dc=com
# ...

```

- 1 Configures specified users as superusers.

### 3.18. MIRRORING CONFIGURATION FIELDS

Table 3.16. 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: openshift-quay-service</b>
REPO_MIRROR_TLS_VERIFY	Boolean	Require HTTPS and verify certificates of Quay registry during mirror.  <b>Default: true</b>
REPO_MIRROR_ROLLBACK	Boolean	When set to <b>true</b> , the repository rolls back after a failed mirror attempt.  <b>Default: false</b>

### 3.19. SECURITY SCANNER CONFIGURATION FIELDS

Table 3.17. Security scanner configuration

Field	Type	Description
FEATURE_SECURITY_SCANNER	Boolean	Enable or disable the security scanner  <b>Default: false</b>

Field	Type	Description
<b>FEATURE_SECURITY_NOTIFICATIONS</b>	Boolean	If the security scanner is enabled, turn on or turn off security notifications  <b>Default: false</b>
<b>SECURITY_SCANNER_V4_REINDEX_THRESHOLD</b>	String	This parameter is used to determine the minimum time, in seconds, to wait before re-indexing a manifest that has either previously failed or has changed states since the last indexing. The data is calculated from the <b>last_indexed datetime</b> in the <b>manifestsecuritystatus</b> table. This parameter is used to avoid trying to re-index every failed manifest on every indexing run. The default time to re-index is 300 seconds.
<b>SECURITY_SCANNER_V4_ENDPOINT</b>	String	The endpoint for the V4 security scanner  <b>Pattern:</b> <b><code>^http(s)?://(.)+\$</code></b>  <b>Example:</b> <b><a href="http://192.168.99.101:6060">http://192.168.99.101:6060</a></b>
<b>SECURITY_SCANNER_V4_PSK</b>	String	The generated pre-shared key (PSK) for Clair
<b>SECURITY_SCANNER_ENDPOINT</b>	String	The endpoint for the V2 security scanner  <b>Pattern:</b> <b><code>^http(s)?://(.)+\$</code></b>  <b>Example:</b> <b><a href="http://192.168.99.100:6060">http://192.168.99.100:6060</a></b>

Field	Type	Description
SECURITY_SCANNER_INDEXING_INTERVAL	Integer	<p>This parameter is used to determine the number of seconds between indexing intervals in the security scanner. When indexing is triggered, Red Hat Quay will query its database for manifests that must be indexed by Clair. These include manifests that have not yet been indexed and manifests that previously failed indexing.</p> <p><b>Default: 30</b></p>
FEATURE_SECURITY_SCANNER_NOTIFY_ON_NEW_INDEX	Boolean	<p>Whether to allow sending notifications about vulnerabilities for new pushes.</p> <p><b>Default: True</b></p>
SECURITY_SCANNER_V4_MANIFEST_CLEANUP	Boolean	<p>Whether the Red Hat Quay garbage collector removes manifests that are not referenced by other tags or manifests.</p> <p><b>Default: True</b></p>
NOTIFICATION_MIN_SEVERITY_ON_NEW_INDEX	String	<p>Set minimal security level for new notifications on detected vulnerabilities. Avoids creation of large number of notifications after first index. If not defined, defaults to <b>High</b>. Available options include <b>Critical, High, Medium, Low, Negligible</b>, and <b>Unknown</b>.</p>
SECURITY_SCANNER_V4_INDEX_MAX_LAYER_SIZE	String	<p>The maximum layer size allowed for indexing. If the layer size exceeds the configured size, the Red Hat Quay UI returns the following message: <b>The manifest for this tag has layer(s) that are too large to index by the Quay Security Scanner</b>. The default is <b>8G</b>, and the maximum recommended is <b>10G</b>. Accepted values are <b>B, K, M, T</b>, and <b>G</b>.</p> <p><b>Default: 8G</b></p>

### 3.19.1. Re-indexing with Clair v4

When Clair v4 indexes a manifest, the result should be deterministic. For example, the same manifest should produce the same index report. This is true until the scanners are changed, as using different scanners will produce different information relating to a specific manifest to be returned in the report. Because of this, Clair v4 exposes a state representation of the indexing engine (`/indexer/api/v1/index_state`) to determine whether the scanner configuration has been changed.

Red Hat Quay leverages this index state by saving it to the index report when parsing to Quay's database. If this state has changed since the manifest was previously scanned, Red Hat Quay will attempt to re-index that manifest during the periodic indexing process.

By default this parameter is set to 30 seconds. Users might decrease the time if they want the indexing process to run more frequently, for example, if they did not want to wait 30 seconds to see security scan results in the UI after pushing a new tag. Users can also change the parameter if they want more control over the request pattern to Clair and the pattern of database operations being performed on the Red Hat Quay database.

### 3.19.2. Example security scanner configuration

The following YAML is the suggested configuration when enabling the security scanner feature.

#### Security scanner YAML configuration

```
FEATURE_SECURITY_NOTIFICATIONS: true
FEATURE_SECURITY_SCANNER: true
FEATURE_SECURITY_SCANNER_NOTIFY_ON_NEW_INDEX: true
...
SECURITY_SCANNER_INDEXING_INTERVAL: 30
SECURITY_SCANNER_V4_MANIFEST_CLEANUP: true
SECURITY_SCANNER_V4_ENDPOINT: http://quay-server.example.com:8081
SECURITY_SCANNER_V4_PSK: MTU5YzA4Y2ZkNzJoMQ==
SERVER_HOSTNAME: quay-server.example.com
SECURITY_SCANNER_V4_INDEX_MAX_LAYER_SIZE: 8G 1
...
```

1 Recommended maximum is **10G**.

## 3.20. HELM CONFIGURATION FIELDS

Table 3.18. Helm configuration fields

Field	Type	Description
<code>FEATURE_GENERAL_OCI_SUPPORT</code>	Boolean	Enable support for OCI artifacts.  <b>Default:</b> True

The following Open Container Initiative (OCI) artifact types are built into Red Hat Quay by default and are enabled through the `FEATURE_GENERAL_OCI_SUPPORT` configuration field:



Field	Media Type	Supported content types
Helm	<code>application/vnd.cncf.helm.config.v1+json</code>	<code>application/tar+gzip</code> , <code>application/vnd.cncf.helm.chart.content.v1.tar+gzip</code>
Cosign	<code>application/vnd.oci.image.config.v1+json</code>	<code>application/vnd.dev.cosign.simplesigning.v1+json</code> , <code>application/vnd.dsse.envelope.v1+json</code>
SPDX	<code>application/vnd.oci.image.config.v1+json</code>	<code>text/spdx</code> , <code>text/spdx+xml</code> , <code>text/spdx+json</code>
Syft	<code>application/vnd.oci.image.config.v1+json</code>	<code>application/vnd.syft+json</code>
CycloneDX	<code>application/vnd.oci.image.config.v1+json</code>	<code>application/vnd.cyclonedx</code> , <code>application/vnd.cyclonedx+xml</code> , <code>application/vnd.cyclonedx+json</code>
In-toto	<code>application/vnd.oci.image.config.v1+json</code>	<code>application/vnd.in-toto+json</code>
Unknown	<code>application/vnd.cncf.openpolicyagent.policy.layer.v1+rego</code>	<code>application/vnd.cncf.openpolicyagent.policy.layer.v1+rego</code> , <code>application/vnd.cncf.openpolicyagent.data.layer.v1+json</code>

### 3.20.1. Configuring Helm

The following YAML is the example configuration when enabling Helm.

#### Helm YAML configuration

```
FEATURE_GENERAL_OCI_SUPPORT: true
```

## 3.21. OPEN CONTAINER INITIATIVE CONFIGURATION FIELDS

Table 3.19. Additional OCI artifact configuration field

Field	Type	Description
<code>FEATURE_REFERRERS_API</code>	Boolean	Enables OCI 1.1's referencers API.

#### Example OCI referencers enablement YAML

```
# ...
FEATURE_REFERRERS_API: True
# ...
```

## 3.22. UNKNOWN MEDIA TYPES

Table 3.20. Unknown media types configuration field

Field	Type	Description
IGNORE_UNKNOWN_MEDIATYPES	Boolean	When enabled, allows a container registry platform to disregard specific restrictions on supported artifact types and accept any unrecognized or unknown media types.  <b>Default: false</b>

### 3.22.1. Configuring unknown media types

The following YAML is the example configuration when enabling unknown or unrecognized media types.

#### Unknown media types YAML configuration

```
IGNORE_UNKNOWN_MEDIATYPES: true
```

## 3.23. ACTION LOG CONFIGURATION FIELDS

### 3.23.1. Action log storage configuration

Table 3.21. Action log storage configuration

Field	Type	Description
FEATURE_LOG_EXPORT	Boolean	Whether to allow exporting of action logs.  <b>Default: True</b>
LOGS_MODEL	String	Specifies the preferred method for handling log data.  <b>Values:</b> One of <b>database</b> , <b>transition_reads_both_writes</b> , <b>elasticsearch</b> , <b>splunk</b> <b>Default: database</b>
LOGS_MODEL_CONFIG	Object	Logs model config for action logs.

Field	Type	Description
<code>ALLOW_WITHOUT_STRICT_LOGGING</code>	Boolean	When set to <b>True</b> , if the external log system like Splunk or Elasticsearch is intermittently unavailable, allows users to push images normally. Events are logged to the stdout instead. Overrides <b>ALLOW_PULLS_WITHOUT_STRICT_LOGGING</b> if set.  <b>Default:</b> False

### 3.23.1.1. Elasticsearch configuration fields

The following fields are available when configuring Elasticsearch for Red Hat Quay.

- `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**

### 3.23.1.2. Splunk configuration fields

The following fields are available when configuring Splunk for Red Hat Quay.

- **producer** [string]: **splunk**. Use when configuring Splunk.

- **splunk\_config** [object]: Logs model configuration for Splunk action logs or the Splunk cluster configuration.
  - **host** [string]: Splunk cluster endpoint.
  - **port** [integer]: Splunk management cluster endpoint port.
  - **bearer\_token** [string]: The bearer token for Splunk.
  - **verify\_ssl** [boolean]: Enable (**True**) or disable (**False**) TLS/SSL verification for HTTPS connections.
  - **index\_prefix** [string]: Splunk's index prefix.
  - **ssl\_ca\_path** [string]: The relative container path to a single **.pem** file containing a certificate authority (CA) for SSL validation.

### Example Splunk configuration

```
# ...
LOGS_MODEL: splunk
LOGS_MODEL_CONFIG:
  producer: splunk
  splunk_config:
    host: http://<user_name>.remote.csb
    port: 8089
    bearer_token: <bearer_token>
    url_scheme: <http/https>
    verify_ssl: False
    index_prefix: <splunk_log_index_name>
    ssl_ca_path: <location_to_ssl-ca-cert.pem>
# ...
```

#### 3.23.1.3. Splunk HEC configuration fields

The following fields are available when configuring Splunk HTTP Event Collector (HEC) for Red Hat Quay.

- **producer** [string]: **splunk\_hec**. Use when configuring Splunk HEC.
- **splunk\_hec\_config** [object]: Logs model configuration for Splunk HTTP event collector action logs configuration.
  - **host** [string]: Splunk cluster endpoint.
  - **port** [integer]: Splunk management cluster endpoint port.
  - **hec\_token** [string]: HEC token for Splunk.
  - **url\_scheme** [string]: The URL scheme for access the Splunk service. If Splunk is behind SSL/TLS, must be **https**.
  - **verify\_ssl** [boolean]: Enable (**true**) or disable (**false**) SSL/TLS verification for HTTPS connections.
  - **index** [string]: The Splunk index to use.

- `splunk_host` [string]: The host name to log this event.
- `splunk_sourcetype` [string]: The name of the Splunk **sourcetype** to use.

```
# ...
LOGS_MODEL: splunk
LOGS_MODEL_CONFIG:
  producer: splunk_hec
  splunk_hec_config: 1
    host: prd-p-aaaaaq.splunkcloud.com 2
    port: 8088 3
    hec_token: 12345678-1234-1234-1234-1234567890ab 4
    url_scheme: https 5
    verify_ssl: False 6
    index: quay 7
    splunk_host: quay-dev 8
    splunk_sourcetype: quay_logs 9
# ...
```

### 3.23.2. Action log rotation and archiving configuration

Table 3.22. Action log rotation and archiving configuration

Field	Type	Description
<code>FEATURE_ACTION_LOG_ROTATION</code>	Boolean	Enabling log rotation and archival will move all logs older than 30 days to storage.  <b>Default: false</b>
<code>ACTION_LOG_ARCHIVE_LOCATION</code>	String	If action log archiving is enabled, the storage engine in which to place the archived data.  <b>Example: s3_us_east</b>
<code>ACTION_LOG_ARCHIVE_PATH</code>	String	If action log archiving is enabled, the path in storage in which to place the archived data.  <b>Example: archives/actionlogs</b>
<code>ACTION_LOG_ROTATION_THRESHOLD</code>	String	The time interval after which to rotate logs.  <b>Example: 30d</b>

### 3.23.3. Action log audit configuration

Table 3.23. Audit logs configuration field

Field	Type	Description
<code>ACTION_LOG_AUDIT_LOGINS</code>	Boolean	When set to <b>True</b> , tracks advanced events such as logging into, and out of, the UI, and logging in using Docker for regular users, robot accounts, and for application-specific token accounts.  <b>Default: True</b>

### 3.24. BUILD LOGS CONFIGURATION FIELDS

Table 3.24. Build logs configuration fields

Field	Type	Description
<code>FEATURE_READER_BUILD_LOGS</code>	Boolean	If set to true, build logs can be read by those with <b>read</b> access to the repository, rather than only <b>write</b> access or <b>admin</b> access.  <b>Default: False</b>
<code>LOG_ARCHIVE_LOCATION</code>	String	The storage location, defined in <b>DISTRIBUTED_STORAGE_CONFIG</b> , in which to place the archived build logs.  <b>Example: s3_us_east</b>
<code>LOG_ARCHIVE_PATH</code>	String	The path under the configured storage engine in which to place the archived build logs in <b>.JSON</b> format.  <b>Example: archives/buildlogs</b>

### 3.25. DOCKERFILE BUILD TRIGGERS FIELDS

Table 3.25. Dockerfile build support

Field	Type	Description
<code>FEATURE_BUILD_SUPPORT</code>	Boolean	Whether to support Dockerfile build.  <b>Default: False</b>

Field	Type	Description
<code>SUCCESSIVE_TRIGGER_FAILURE_DISABLE_THRESHOLD</code>	Number	If not set to <b>None</b> , the number of successive failures that can occur before a build trigger is automatically disabled.  <b>Default: 100</b>
<code>SUCCESSIVE_TRIGGER_INTERNAL_ERROR_DISABLE_THRESHOLD</code>	Number	If not set to <b>None</b> , the number of successive internal errors that can occur before a build trigger is automatically disabled.  <b>Default: 5</b>

### 3.25.1. GitHub build triggers

Table 3.26. GitHub build triggers

Field	Type	Description
<code>FEATURE_GITHUB_BUILD</code>	Boolean	Whether to support GitHub build triggers.  <b>Default: False</b>
<code>GITHUB_TRIGGER_CONFIG</code>	Object	Configuration for using GitHub Enterprise for build triggers.
<code>.GITHUB_ENDPOINT</code> (Required)	String	The endpoint for GitHub Enterprise.  <b>Example: <a href="https://github.com/">https://github.com/</a></b>
<code>.API_ENDPOINT</code>	String	The endpoint of the GitHub Enterprise API to use. Must be overridden for <b>github.com</b> .  <b>Example: <a href="https://api.github.com/">https://api.github.com/</a></b>
<code>.CLIENT_ID</code> (Required)	String	The registered client ID for this Red Hat Quay instance; this cannot be shared with <b>GITHUB_LOGIN_CONFIG</b> .



Field	Type	Description
<code>.CLIENT_SECRET</code> (Required)	String	The registered client secret for this Red Hat Quay instance.

### 3.25.2. BitBucket build triggers

Table 3.27. BitBucket build triggers

Field	Type	Description
<code>FEATURE_BITBUCKET_BUILD</code>	Boolean	Whether to support Bitbucket build triggers.  <b>Default:</b> False
<code>BITBUCKET_TRIGGER_CONFIG</code>	Object	Configuration for using BitBucket for build triggers.
<code>.CONSUMER_KEY</code> (Required)	String	The registered consumer key (client ID) for this Red Hat Quay instance.
<code>.CONSUMER_SECRET</code> (Required)	String	The registered consumer secret (client secret) for this Red Hat Quay instance.

### 3.25.3. GitLab build triggers

Table 3.28. GitLab build triggers

Field	Type	Description
<code>FEATURE_GITLAB_BUILD</code>	Boolean	Whether to support GitLab build triggers.  <b>Default:</b> False
<code>GITLAB_TRIGGER_CONFIG</code>	Object	Configuration for using Gitlab for build triggers.
<code>.GITLAB_ENDPOINT</code> (Required)	String	The endpoint at which Gitlab Enterprise is running.

Field	Type	Description
<b>.CLIENT_ID</b> (Required)	String	The registered client ID for this Red Hat Quay instance.
<b>.CLIENT_SECRET</b> (Required)	String	The registered client secret for this Red Hat Quay instance.

## 3.26. BUILD MANAGER CONFIGURATION FIELDS

Table 3.29. Build manager configuration fields

Field	Type	Description
<b>ALLOWED_WORKER_COUNT</b>	String	Defines how many Build Workers are instantiated per Red Hat Quay pod. Typically set to <b>1</b> .
<b>ORCHESTRATOR_PREFIX</b>	String	Defines a unique prefix to be added to all Redis keys. This is useful to isolate Orchestrator values from other Redis keys.
<b>REDIS_HOST</b>	Object	The hostname for your Redis service.
<b>REDIS_PASSWORD</b>	String	The password to authenticate into your Redis service.
<b>REDIS_SSL</b>	Boolean	Defines whether or not your Redis connection uses SSL/TLS.
<b>REDIS_SKIP_KEYSPACE_EVENT_SETUP</b>	Boolean	By default, Red Hat Quay does not set up the keyspace events required for key events at runtime. To do so, set <b>REDIS_SKIP_KEYSPACE_EVENT_SETUP</b> to <b>false</b> .
<b>EXECUTOR</b>	String	Starts a definition of an Executor of this type. Valid values are <b>kubernetes</b> and <b>ec2</b> .
<b>BUILDER_NAMESPACE</b>	String	Kubernetes namespace where Red Hat Quay Builds will take place.

Field	Type	Description
K8S_API_SERVER	Object	Hostname for API Server of the OpenShift Container Platform cluster where Builds will take place.
K8S_API_TLS_CA	Object	The filepath in the <b>Quay</b> container of the Build cluster's CA certificate for the <b>Quay</b> application to trust when making API calls.
KUBERNETES_DISTRIBUTION	String	Indicates which type of Kubernetes is being used. Valid values are <b>openshift</b> and <b>k8s</b> .
CONTAINER_*	Object	Define the resource requests and limits for each <b>build</b> pod.
NODE_SELECTOR_*	Object	Defines the node selector label name-value pair where <b>build</b> Pods should be scheduled.
CONTAINER_RUNTIME	Object	Specifies whether the Builder should run <b>docker</b> or <b>podman</b> . Customers using Red Hat's <b>quay-builder</b> image should set this to <b>podman</b> .
SERVICE_ACCOUNT_NAME/SERVICE_ACCOUNT_TOKEN	Object	Defines the Service Account name or token that will be used by <b>build</b> pods.
QUAY_USERNAME/QUAY_PASSWORD	Object	Defines the registry credentials needed to pull the Red Hat Quay build worker image that is specified in the <b>WORKER_IMAGE</b> field. Customers should provide a Red Hat Service Account credential as defined in the section "Creating Registry Service Accounts" against registry.redhat.io in the article at <a href="https://access.redhat.com/RegistryAuthentication">https://access.redhat.com/RegistryAuthentication</a> .

Field	Type	Description
<b>WORKER_IMAGE</b>	Object	Image reference for the Red Hat Quay Builder image. registry.redhat.io/quay/quay-builder
<b>WORKER_TAG</b>	Object	Tag for the Builder image desired. The latest version is 3.13.
<b>BUILDER_VM_CONTAINER_IMAGE</b>	Object	The full reference to the container image holding the internal VM needed to run each Red Hat Quay Build. <b>(registry.redhat.io/quay/quay-builder-qemu-rhcos:3.13).</b>
<b>SETUP_TIME</b>	String	Specifies the number of seconds at which a Build times out if it has not yet registered itself with the Build Manager. Defaults at <b>500</b> seconds. Builds that time out are attempted to be restarted three times. If the Build does not register itself after three attempts it is considered failed.

Field	Type	Description
<code>MINIMUM_RETRY_THRESHOLD</code>	String	This setting is used with multiple Executors. It indicates how many retries are attempted to start a Build before a different Executor is chosen. Setting to <b>0</b> means there are no restrictions on how many tries the build job needs to have. This value should be kept intentionally small (three or less) to ensure failovers happen quickly during infrastructure failures. You must specify a value for this setting. For example, <b>Kubernetes</b> is set as the first executor and <b>EC2</b> as the second executor. If you want the last attempt to run a job to always be executed on EC2 and not Kubernetes, you can set the Kubernetes executor's <b>MINIMUM_RETRY_THRESHOLD</b> to <b>1</b> and EC2's <b>MINIMUM_RETRY_THRESHOLD</b> to <b>0</b> (defaults to <b>0</b> if not set). In this case, the Kubernetes' <b>MINIMUM_RETRY_THRESHOLD retries_remaining(1)</b> would evaluate to <b>False</b> , therefore falling back to the second executor configured.
<code>SSH_AUTHORIZED_KEYS</code>	Object	List of SSH keys to bootstrap in the <b>ignition</b> config. This allows other keys to be used to SSH into the EC2 instance or QEMU virtual machine (VM).

### 3.27. OAUTH CONFIGURATION FIELDS

Table 3.30. OAuth fields

Field	Type	Description
<code>DIRECT_OAUTH_CLIENTID_WHITELIST</code>	Array of String	A list of client IDs for <b>Quay-managed</b> applications that are allowed to perform direct OAuth approval without user approval.

Field	Type	Description
FEATURE_ASSIGN_OAUTH_TOKEN	Boolean	Allows organization administrators to assign OAuth tokens to other users.

### 3.27.1. GitHub OAuth configuration fields

Table 3.31. GitHub OAuth fields

Field	Type	Description
FEATURE_GITHUB_LOGIN	Boolean	Whether GitHub login is supported  **Default: <b>False</b>
GITHUB_LOGIN_CONFIG	Object	Configuration for using GitHub (Enterprise) as an external login provider.
.ALLOWED_ORGANIZATIONS	Array of String	The names of the GitHub (Enterprise) organizations whitelisted to work with the <code>ORG_RESTRICT</code> option.
.API_ENDPOINT	String	The endpoint of the GitHub (Enterprise) API to use. Must be overridden for github.com  <b>Example:</b> <a href="https://api.github.com/">https://api.github.com/</a>
.CLIENT_ID (Required)	String	The registered client ID for this Red Hat Quay instance; cannot be shared with <b>GITHUB_TRIGGER_CONFIG</b> .  <b>Example:</b> <b>0e8dbe15c4c7630b6780</b>
.CLIENT_SECRET (Required)	String	The registered client secret for this Red Hat Quay instance.  <b>Example:</b> <b>e4a58ddd3d7408b7aec109e85564a0d153d3e846</b>

Field	Type	Description
<code>.GITHUB_ENDPOINT</code> (Required)	String	The endpoint for GitHub (Enterprise).  Example: <a href="https://github.com/">https://github.com/</a>
<code>.ORG_RESTRICT</code>	Boolean	If true, only users within the organization whitelist can login using this provider.

### 3.27.2. Google OAuth configuration fields

Table 3.32. Google OAuth fields

Field	Type	Description
<code>FEATURE_GOOGLE_LOGIN</code>	Boolean	Whether Google login is supported.  **Default: <b>False</b>
<code>GOOGLE_LOGIN_CONFIG</code>	Object	Configuration for using Google for external authentication.
<code>.CLIENT_ID</code> (Required)	String	The registered client ID for this Red Hat Quay instance.  Example: <b>0e8dbe15c4c7630b6780</b>
<code>.CLIENT_SECRET</code> (Required)	String	The registered client secret for this Red Hat Quay instance.  Example: <b>e4a58ddd3d7408b7aec109e85564a0d153d3e846</b>

## 3.28. OIDC CONFIGURATION FIELDS

Table 3.33. OIDC fields

Field	Type	Description
-------	------	-------------

<code>&lt;string&gt;_LOGIN_CONFIG</code> (Required)	String	The parent key that holds the OIDC configuration settings. Typically the name of the OIDC provider, for example, <b>AZURE_LOGIN_CONFIG</b> , however any arbitrary string is accepted.
<code>.CLIENT_ID</code> (Required)	String	The registered client ID for this Red Hat Quay instance.  <b>Example:</b> <b>0e8dbe15c4c7630b6780</b>
<code>.CLIENT_SECRET</code> (Required)	String	The registered client secret for this Red Hat Quay instance.  <b>Example:</b> <b>e4a58ddd3d7408b7aec109e85564a0d153d3e846</b>
<code>.DEBUGLOG</code>	Boolean	Whether to enable debugging.
<code>.LOGIN_BINDING_FIELD</code>	String	Used when the internal authorization is set to LDAP. Red Hat Quay reads this parameter and tries to search through the LDAP tree for the user with this username. If it exists, it automatically creates a link to that LDAP account.
<code>.LOGIN_SCOPES</code>	Object	Adds additional scopes that Red Hat Quay uses to communicate with the OIDC provider.
<code>.OIDC_ENDPOINT_CUSTOM_PARAMS</code>	String	Support for custom query parameters on OIDC endpoints. The following endpoints are supported: <b>authorization_endpoint</b> , <b>token_endpoint</b> , and <b>user_endpoint</b> .



<code>.OIDC_ISSUER</code>	String	Allows the user to define the issuer to verify. For example, JWT tokens contain a parameter known as <b>iss</b> which defines who issued the token. By default, this is read from the <b>.well-known/openid/configuration</b> endpoint, which is exposed by every OIDC provider. If this verification fails, there is no login.
<code>.OIDC_SERVER</code> (Required)	String	The address of the OIDC server that is being used for authentication.  <b>Example:</b> <b>https://sts.windows.net/6c878.../</b>
<code>.PREFERRED_USERNAME_CLAIM_NAME</code>	String	Sets the preferred username to a parameter from the token.
<code>.SERVICE_ICON</code>	String	Changes the icon on the login screen.
<code>.SERVICE_NAME</code> (Required)	String	The name of the service that is being authenticated.  <b>Example: Microsoft Entra ID</b>
<code>.VERIFIED_EMAIL_CLAIM_NAME</code>	String	The name of the claim that is used to verify the email address of the user.
<code>.PREFERRED_GROUP_CLAIM_NAME</code>	String	The key name within the OIDC token payload that holds information about the user's group memberships.
<code>.OIDC_DISABLE_USER_ENDPOINT</code>	Boolean	Whether to allow or disable the <b>/userinfo</b> endpoint. If using Azure Entra ID, this field must be set to <b>true</b> because Azure obtains the user's information from the token instead of calling the <b>/userinfo</b> endpoint.  <b>Default: false</b>

### 3.28.1. OIDC configuration

The following example shows a sample OIDC configuration.

### Example OIDC configuration

```
AUTHENTICATION_TYPE: OIDC
# ...
AZURE_LOGIN_CONFIG:
  CLIENT_ID: <client_id>
  CLIENT_SECRET: <client_secret>
  OIDC_SERVER: <oidc_server_address_>
  DEBUGGING: true
  SERVICE_NAME: Microsoft Entra ID
  VERIFIED_EMAIL_CLAIM_NAME: <verified_email>
  OIDC_DISABLE_USER_ENDPOINT: true
  OIDC_ENDPOINT_CUSTOM_PARAMS":
    "authorization_endpoint":
      "some": "param",
# ...
```

## 3.29. NESTED REPOSITORIES CONFIGURATION FIELDS

Support for nested repository path names has been added under the **FEATURE\_EXTENDED\_REPOSITORY\_NAMES** property. This optional configuration is added to the config.yaml by default. Enablement allows the use of / in repository names.

Table 3.34. OCI and nested repositories configuration fields

Field	Type	Description
<b>FEATURE_EXTENDED_REPOSITORY_NAMES</b>	Boolean	Enable support for nested repositories  <b>Default:</b> True

### OCI and nested repositories configuration example

```
FEATURE_EXTENDED_REPOSITORY_NAMES: true
```

## 3.30. QUAYINTEGRATION CONFIGURATION FIELDS

The following configuration fields are available for the QuayIntegration custom resource:

Name	Description	Schema
allowlistNamespaces (Optional)	A list of namespaces to include.	Array
clusterID (Required)	The ID associated with this cluster.	String

Name	Description	Schema
credentialsSecret.key (Required)	The secret containing credentials to communicate with the Quay registry.	Object
denylistNamespaces (Optional)	A list of namespaces to exclude.	Array
insecureRegistry (Optional)	Whether to skip TLS verification to the Quay registry	Boolean
quayHostname (Required)	The hostname of the Quay registry.	String
scheduledImageStreamImport (Optional)	Whether to enable image stream importing.	Boolean

### 3.31. MAIL CONFIGURATION FIELDS

Table 3.35. Mail configuration fields

Field	Type	Description
FEATURE_MAILING	Boolean	Whether emails are enabled  <b>Default: False</b>
MAIL_DEFAULT_SENDER	String	If specified, the e-mail address used as the <b>from</b> when Red Hat Quay sends e-mails. If none, defaults to <b>support@quay.io</b>  <b>Example:</b> <b>support@example.com</b>
MAIL_PASSWORD	String	The SMTP password to use when sending e-mails
MAIL_PORT	Number	The SMTP port to use. If not specified, defaults to 587.

Field	Type	Description
MAIL_SERVER	String	The SMTP server to use for sending e-mails. Only required if FEATURE_MAILING is set to true.  <b>Example: smtp.example.com</b>
MAIL_USERNAME	String	The SMTP username to use when sending e-mails
MAIL_USE_TLS	Boolean	If specified, whether to use TLS for sending e-mails  <b>Default: True</b>

### 3.32. USER CONFIGURATION FIELDS

Table 3.36. User configuration fields

Field	Type	Description
FEATURE_SUPER_USERS	Boolean	Whether superusers are supported  <b>Default: true</b>
FEATURE_USER_CREATION	Boolean	Whether users can be created (by non-superusers)  <b>Default: true</b>
FEATURE_USER_LAST_ACCESSED	Boolean	Whether to record the last time a user was accessed  <b>Default: true</b>
FEATURE_USER_LOG_ACCESS	Boolean	If set to true, users will have access to audit logs for their namespace  <b>Default: false</b>
FEATURE_USER_METADATA	Boolean	Whether to collect and support user metadata  <b>Default: false</b>

Field	Type	Description
FEATURE_USERNAME_CONFIRMATION	Boolean	If set to true, users can confirm and modify their initial usernames when logging in via OpenID Connect (OIDC) or a non-database internal authentication provider like LDAP. <b>Default: true</b>
FEATURE_USER_RENAME	Boolean	If set to true, users can rename their own namespace <b>Default: false</b>
FEATURE_INVITE_ONLY_USER_CREATION	Boolean	Whether users being created must be invited by another user <b>Default: false</b>
FRESH_LOGIN_TIMEOUT	String	The time after which a fresh login requires users to re-enter their password <b>Example: 5m</b>
USERFILES_LOCATION	String	ID of the storage engine in which to place user-uploaded files <b>Example: s3_us_east</b>
USERFILES_PATH	String	Path under storage in which to place user-uploaded files <b>Example: userfiles</b>
USER_RECOVERY_TOKEN_LIFETIME	String	The length of time a token for recovering a user accounts is valid <b>Pattern: <code>^[0-9]+(w m d h s)\$</code></b> <b>Default: 30m</b>
FEATURE_SUPERUSERS_FULL_ACCESS	Boolean	Grants superusers the ability to read, write, and delete content from other repositories in namespaces that they do not own or have explicit permissions for. <b>Default: False</b>

Field	Type	Description
<code>FEATURE_SUPERUSERS_ORG_CREATION_ONLY</code>	Boolean	Whether to only allow superusers to create organizations.  <b>Default: False</b>
<code>FEATURE_RESTRICTED_USERS</code>	Boolean	When set to <b>True</b> with <b>RESTRICTED_USERS_WHITELIST</b> : <ul style="list-style-type: none"> <li>All normal users and superusers are restricted from creating organizations or content in their own namespace unless they are allowlisted via <b>RESTRICTED_USERS_WHITELIST</b>.</li> <li>Restricted users retain their normal permissions within organizations based on team memberships.</li> </ul> <b>Default: False</b>
<code>RESTRICTED_USERS_WHITELIST</code>	String	When set with <b>FEATURE_RESTRICTED_USERS: true</b> , specific users are excluded from the <b>FEATURE_RESTRICTED_USERS</b> setting.
<code>GLOBAL_READONLY_SUPER_USERS</code>	String	When set, grants users of this list read access to all repositories, regardless of whether they are public repositories. Only works for those superusers defined with the <b>SUPER_USERS</b> configuration field.

### 3.32.1. User configuration fields references

Use the following references to update your `config.yaml` file with the desired configuration field.

#### 3.32.1.1. FEATURE\_SUPERUSERS\_FULL\_ACCESS configuration reference

```
---
SUPER_USERS:
- quayadmin
```

```
FEATURE_SUPERUSERS_FULL_ACCESS: True
```

```
---
```

### 3.32.1.2. GLOBAL\_READONLY\_SUPER\_USERS configuration reference

```
---
```

```
GLOBAL_READONLY_SUPER_USERS:
```

```
- user1
```

```
---
```

### 3.32.1.3. FEATURE\_RESTRICTED\_USERS configuration reference

```
---
```

```
AUTHENTICATION_TYPE: Database
```

```
---
```

```
---
```

```
FEATURE_RESTRICTED_USERS: true
```

```
---
```

### 3.32.1.4. RESTRICTED\_USERS\_WHITELIST configuration reference

#### Prerequisites

- **FEATURE\_RESTRICTED\_USERS** is set to **true** in your **config.yaml** file.

```
---
```

```
AUTHENTICATION_TYPE: Database
```

```
---
```


```
---
```

```
FEATURE_RESTRICTED_USERS: true
```

```
RESTRICTED_USERS_WHITELIST:
```

```
- user1
```

```
---
```



#### NOTE

When this field is set, whitelisted users can create organizations, or read or write content from the repository even if **FEATURE\_RESTRICTED\_USERS** is set to **true**. Other users, for example, **user2**, **user3**, and **user4** are restricted from creating organizations, reading, or writing content

## 3.33. RECAPTCHA CONFIGURATION FIELDS

Table 3.37. Recaptcha configuration fields

Field	Type	Description
-------	------	-------------

Field	Type	Description
FEATURE_RECAPTCHA	Boolean	Whether Recaptcha is necessary for user login and recovery  <b>Default:</b> False
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

### 3.34. ACI CONFIGURATION FIELDS

Table 3.38. ACI configuration fields

Field	Type	Description
FEATURE_ACI_CONVERSION	Boolean	Whether to enable conversion to ACIs  <b>Default:</b> False
GPG2_PRIVATE_KEY_FILENAME	String	The filename of the private key used to decrypt ACIs
GPG2_PRIVATE_KEY_NAME	String	The name of the private key used to sign ACIs
GPG2_PUBLIC_KEY_FILENAME	String	The filename of the public key used to encrypt ACIs

### 3.35. JWT CONFIGURATION FIELDS

Table 3.39. JWT configuration fields

Field	Type	Description
JWT_AUTH_ISSUER	String	The endpoint for JWT users  <b>Pattern:</b> <code>^http(s)?://(.)+\$</code> <b>Example:</b> <a href="http://192.168.99.101:6060">http://192.168.99.101:6060</a>
JWT_GETUSER_ENDPOINT	String	The endpoint for JWT users <b>Pattern:</b> <code>^http(s)?://(.)+\$</code> <b>Example:</b> <a href="http://192.168.99.101:6060">http://192.168.99.101:6060</a>



Field	Type	Description
JWT_QUERY_ENDPOINT	String	The endpoint for JWT queries  Pattern: <b>^http(s)?://(.)+\$</b> Example: <a href="http://192.168.99.101:6060">http://192.168.99.101:6060</a>
JWT_VERIFY_ENDPOINT	String	The endpoint for JWT verification  Pattern: <b>^http(s)?://(.)+\$</b> Example: <a href="http://192.168.99.101:6060">http://192.168.99.101:6060</a>

### 3.36. APP TOKENS CONFIGURATION FIELDS

Table 3.40. App tokens configuration fields

Field	Type	Description
FEATURE_APP_SPECIFIC_TOKENS	Boolean	If enabled, users can create tokens for use by the Docker CLI  <b>Default:</b> True
APP_SPECIFIC_TOKEN_EXPIRATION	String	The expiration for external app tokens.  <b>Default:</b> None <b>Pattern:</b> <b>^[0-9]+(w m d h s)\$</b>
EXPIRED_APP_SPECIFIC_TOKEN_GC	String	Duration of time expired external app tokens will remain before being garbage collected  <b>Default:</b> 1d

### 3.37. MISCELLANEOUS CONFIGURATION FIELDS

Table 3.41. Miscellaneous configuration fields

Field	Type	Description
-------	------	-------------

Field	Type	Description
<code>ALLOW_PULLS_WITHOUT_STRICT_LOGGING</code>	String	<p>If true, pulls will still succeed even if the pull audit log entry cannot be written. This is useful if the database is in a read-only state and it is desired for pulls to continue during that time.</p> <p><b>Default:</b> False</p>
<code>AVATAR_KIND</code>	String	<p>The types of avatars to display, either generated inline (local) or Gravatar (gravatar)</p> <p><b>Values:</b> local, gravatar</p>
<code>BROWSER_API_CALLS_XHR_ONLY</code>	Boolean	<p>If enabled, only API calls marked as being made by an XHR will be allowed from browsers</p> <p><b>Default:</b> True</p>
<code>DEFAULT_NAMESPACE_MAXIMUM_BUILD_COUNT</code>	Number	<p>The default maximum number of builds that can be queued in a namespace.</p> <p><b>Default:</b> None</p>
<code>ENABLE_HEALTH_DEBUG_SECRET</code>	String	<p>If specified, a secret that can be given to health endpoints to see full debug info when not authenticated as a superuser</p>
<code>EXTERNAL_TLS_TERMINATION</code>	Boolean	<p>Set to <b>true</b> if TLS is supported, but terminated at a layer before Quay. Set to <b>false</b> when Quay is running with its own SSL certificates and receiving TLS traffic directly.</p>
<code>FRESH_LOGIN_TIMEOUT</code>	String	<p>The time after which a fresh login requires users to re-enter their password</p> <p><b>Example:</b> 5m</p>

Field	Type	Description
HEALTH_CHECKER	String	The configured health check  <b>Example:</b> <code>('RDSAwareHealthCheck', {'access_key': 'foo', 'secret_key': 'bar'})</code>
PROMETHEUS_NAMESPACE	String	The prefix applied to all exposed Prometheus metrics  <b>Default: quay</b>
PUBLIC_NAMESPACES	Array of String	If a namespace is defined in the public namespace list, then it will appear on <b>all</b> users' repository list pages, regardless of whether the user is a member of the namespace. Typically, this is used by an enterprise customer in configuring a set of "well-known" namespaces.
REGISTRY_STATE	String	The state of the registry  <b>Values: normal or read-only</b>
SEARCH_MAX_RESULT_PAGE_COUNT	Number	Maximum number of pages the user can paginate in search before they are limited  <b>Default: 10</b>
SEARCH_RESULTS_PER_PAGE	Number	Number of results returned per page by search page  <b>Default: 10</b>
V2_PAGINATION_SIZE	Number	The number of results returned per page in V2 registry APIs  <b>Default: 50</b>
WEBHOOK_HOSTNAME_BLACKLIST	Array of String	The set of hostnames to disallow from webhooks when validating, beyond localhost

Field	Type	Description
CREATE_PRIVATE_REPO_ON_PUSH	Boolean	Whether new repositories created by push are set to private visibility  <b>Default:</b> True
CREATE_NAMESPACE_ON_PUSH	Boolean	Whether new push to a non-existent organization creates it  <b>Default:</b> False
NON_RATE_LIMITED_NAMESPACES	Array of String	If rate limiting has been enabled using <b>FEATURE_RATE_LIMITS</b> , you can override it for specific namespace that require unlimited access.
<a href="#">FEATURE_UI_V2</a>	Boolean	When set, allows users to try the beta UI environment.  <b>Default: True</b>
FEATURE_REQUIRE_TEAM_INVITE	Boolean	Whether to require invitations when adding a user to a team  <b>Default:</b> True
FEATURE_REQUIRE_ENCRYPTED_BASIC_AUTH	Boolean	Whether non-encrypted passwords (as opposed to encrypted tokens) can be used for basic auth  <b>Default:</b> False
FEATURE_RATE_LIMITS	Boolean	Whether to enable rate limits on API and registry endpoints. Setting <b>FEATURE_RATE_LIMITS</b> to <b>true</b> causes <b>nginx</b> 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).  <b>Default:</b> False

Field	Type	Description
FEATURE_FIPS	Boolean	If set to true, Red Hat Quay will run using FIPS-compliant hash functions  <b>Default:</b> False
FEATURE_AGGREGATED_LOG_COUNT_RETRIEVAL	Boolean	Whether to allow retrieval of aggregated log counts  <b>Default:</b> True
FEATURE_ANONYMOUS_ACCESS	Boolean	Whether to allow anonymous users to browse and pull public repositories  <b>Default:</b> True
FEATURE_DIRECT_LOGIN	Boolean	Whether users can directly login to the UI  <b>Default:</b> True
FEATURE_LIBRARY_SUPPORT	Boolean	Whether to allow for "namespace-less" repositories when pulling and pushing from Docker  <b>Default:</b> True
FEATURE_PARTIAL_USER_AUTOCOMPLETE	Boolean	If set to true, autocompletion will apply to partial usernames+ <b>Default:</b> True
FEATURE_PERMANENT_SESSIONS	Boolean	Whether sessions are permanent  <b>Default:</b> True
FEATURE_PUBLIC_CATALOG	Boolean	If set to true, the <b>_catalog</b> endpoint returns public repositories. Otherwise, only private repositories can be returned.  <b>Default:</b> False

Field	Type	Description
<b>DISABLE_PUSHES</b>	Boolean	<p>Disables pushes of new content to the registry while retaining all other functionality. Differs from <b>read-only</b> mode because database is not set as <b>read-only</b>. When <b>DISABLE_PUSHES</b> is set to <b>true</b>, the Red Hat Quay garbage collector is disabled. As a result, when <b>PERMANENTLY_DELETE_TAGS</b> is enabled, using the Red Hat Quay UI to permanently delete a tag does not result in the immediate deletion of a tag. Instead, the image stays in the backend storage until <b>DISABLE_PUSHES</b> is set to <b>false</b>, which re-enables the garbage collector. Red Hat Quay administrators should be aware of this caveat when using <b>DISABLE_PUSHES</b> and <b>PERMANENTLY_DELETE_TAGS</b> together.</p> <p><b>Default:</b> False</p>

### 3.38. LEGACY CONFIGURATION FIELDS

The following fields are deprecated or obsolete.

Table 3.42. Legacy configuration fields

Field	Type	Description
<b>FEATURE_BLACKLISTED_EMAILS</b>	Boolean	If set to true, no new User accounts may be created if their email domain is blacklisted
<b>BLACKLISTED_EMAIL_DOMAINS</b>	Array of String	<p>The list of email-address domains that is used if <b>FEATURE_BLACKLISTED_EMAILS</b> is set to true</p> <p><b>Example:</b> <b>"example.com"</b>, <b>"example.org"</b></p>

Field	Type	Description
BLACKLIST_V2_SPEC	String	The Docker CLI versions to which Red Hat Quay will respond that V2 is <b>unsupported</b>  <b>Example: &lt;1.8.0</b> <b>Default: &lt;1.6.0</b>
DOCUMENTATION_ROOT	String	Root URL for documentation links. This field is useful when Red Hat Quay is configured for disconnected environments to set an alternatively, or allowlisted, documentation link.
SECURITY_SCANNER_V4_NAMESPACE_WHITELIST	String	The namespaces for which the security scanner should be enabled
FEATURE_RESTRICTED_V1_PUSH	Boolean	If set to true, only namespaces listed in V1_PUSH_WHITELIST support V1 push  <b>Default: False</b>
V1_PUSH_WHITELIST	Array of String	The array of namespace names that support V1 push if FEATURE_RESTRICTED_V1_PUSH is set to true
FEATURE_HELM_OCI_SUPPORT	Boolean	Enable support for Helm artifacts.  <b>Default: False</b>
ALLOWED_OCI_ARTIFACT_TYPES	Object	The set of allowed OCI artifact MIME types and the associated layer types.

### 3.39. USER INTERFACE V2 CONFIGURATION FIELDS


Table 3.43. User interface v2 configuration fields

Field	Type	Description
-------	------	-------------

Field	Type	Description
<code>FEATURE_UI_V2</code>	Boolean	When set, allows users to try the beta UI environment.  + <b>Default: False</b>
<code>FEATURE_UI_V2_REPO_SETTINGS</code>	Boolean	When set to <b>True</b> , enables repository settings in the Red Hat Quay v2 UI.  + <b>Default: False</b>

### 3.39.1. v2 user interface configuration

With `FEATURE_UI_V2` enabled, you can toggle between the current version of the user interface and the new version of the user interface.



#### IMPORTANT


- This UI is currently in beta and subject to change. In its current state, users can only create, view, and delete organizations, repositories, and image tags.
- When running Red Hat Quay in the old UI, timed-out sessions would require that the user input their password again in the pop-up window. With the new UI, users are returned to the main page and required to input their username and password credentials. This is a known issue and will be fixed in a future version of the new UI.
- There is a discrepancy in how image manifest sizes are reported between the legacy UI and the new UI. In the legacy UI, image manifests were reported in mebibytes. In the new UI, Red Hat Quay uses the standard definition of megabyte (MB) to report image manifest sizes.

#### Procedure

1. In your deployment's `config.yaml` file, add the `FEATURE_UI_V2` parameter and set it to `true`, for example:

```
---
FEATURE_TEAM_SYNCING: false
FEATURE_UI_V2: true
FEATURE_USER_CREATION: true
---
```

2. Log in to your Red Hat Quay deployment.
3. In the navigation pane of your Red Hat Quay deployment, you are given the option to toggle between **Current UI** and **New UI**. Click the toggle button to set it to new UI, and then click **Use Beta Environment**, for example:





## 3.40. IPV6 CONFIGURATION FIELD

Table 3.44. IPv6 configuration field

Field	Type	Description
FEATURE_LISTEN_IP_VERSION	String	<p>Enables IPv4, IPv6, or dual-stack protocol family. This configuration field must be properly set, otherwise Red Hat Quay fails to start.</p> <p><b>Default: IPv4</b></p> <p><b>Additional configurations: IPv6, dual-stack</b></p>

## 3.41. BRANDING CONFIGURATION FIELDS

Table 3.45. Branding configuration fields

Field	Type	Description
BRANDING	Object	Custom branding for logos and URLs in the Red Hat Quay UI.
.logo (Required)	String	<p>Main logo image URL.</p> <p>The header logo defaults to 205x30 PX. The form logo on the Red Hat Quay sign in screen of the web UI defaults to 356.5x39.7 PX.</p> <p><b>Example:</b> <b><code>/static/img/quay-horizontal-color.svg</code></b></p>
.footer_img	String	<p>Logo for UI footer. Defaults to 144x34 PX.</p> <p><b>Example:</b> <b><code>/static/img/RedHat.svg</code></b></p>
.footer_url	String	<p>Link for footer image.</p> <p><b>Example:</b> <b><a href="https://redhat.com">https://redhat.com</a></b></p>

### 3.41.1. Example configuration for Red Hat Quay branding

Branding config.yaml example

-

**BRANDING:**logo: [https://www.mend.io/wp-content/media/2020/03/5-tips\\_small.jpg](https://www.mend.io/wp-content/media/2020/03/5-tips_small.jpg)footer\_img: [https://www.mend.io/wp-content/media/2020/03/5-tips\\_small.jpg](https://www.mend.io/wp-content/media/2020/03/5-tips_small.jpg)footer\_url: <https://opensourceworld.org/>

## 3.42. SESSION TIMEOUT CONFIGURATION FIELD


The following configuration field relies on on the Flask API configuration field of the same name.

**Table 3.46. Session logout configuration field**

Field	Type	Description
PERMANENT_SESSION_LIFETIME	Integer	A <b>timedelta</b> which is used to set the expiration date of a permanent session. The default is 31 days, which makes a permanent session survive for roughly one month.  <b>Default: 2678400</b>

### 3.42.1. Example session timeout configuration

The following YAML is the suggest configuration when enabling session lifetime.

**IMPORTANT**

Altering session lifetime is not recommended. Administrators should be aware of the allotted time when setting a session timeout. If you set the time too early, it might interrupt your workflow.

#### Session timeout YAML configuration

```
PERMANENT_SESSION_LIFETIME: 3000
```

## CHAPTER 4. ENVIRONMENT VARIABLES

Red Hat Quay supports a limited number of environment variables for dynamic configuration.

### 4.1. GEO-REPLICATION

The same configuration should be used across all regions, with exception of the storage backend, which can be configured explicitly using the **QUAY\_DISTRICTED\_STORAGE\_PREFERENCE** environment variable.

**Table 4.1. Geo-replication configuration**

Variable	Type	Description
<b>QUAY_DISTRICTED_STORAGE_PREFERENCE</b>	String	The preferred storage engine (by ID in <code>DISTRICTED_STORAGE_CONFIG</code> ) to use.

### 4.2. 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.

Database connection pooling is enabled by default, and each process that interacts with the database contains 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 might require analysis to ensure that Red Hat Quay does not exceed the configured database's maximum connection count.

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

For standalone Red Hat Quay deployments, database connection pooling can be toggled off when starting your deployment. For example:

```
$ sudo podman run -d --rm -p 80:8080 -p 443:8443 \
  --name=quay \
  -v $QUAY/config:/conf/stack:Z \
  -v $QUAY/storage:/datastorage:Z \
  -e DB_CONNECTION_POOLING=false
registry.redhat.io/quay/quay-rhel8:v3.12.1
```

For Red Hat Quay on OpenShift Container Platform, database connection pooling can be configured by modifying the **QuayRegistry** custom resource definition (CRD). For example:

#### Example QuayRegistry CRD

```
spec:
  components:
  - kind: quay
    managed: true
  overrides:
```

```
env:
- name: DB_CONNECTION_POOLING
  value: "false"
```

Table 4.2. Database connection pooling configuration

Variable	Type	Description
DB_CONNECTION_POOLING	String	Whether to enable or disable database connection pooling. Defaults to true. Accepted values are <b>"true"</b> or <b>"false"</b>

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:

#### config.yaml

```
...
DB_CONNECTION_ARGS:
  max_connections: 10
...
```

## 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.

Table 4.3. HTTP connection counts configuration

Variable	Type	Description
WORKER_CONNECTION_COUNT	Number	Simultaneous HTTP connections  <b>Default:</b> 50
WORKER_CONNECTION_COUNT_REGISTRY	Number	Simultaneous HTTP connections for registry  <b>Default:</b> WORKER_CONNECTION_COUNT
WORKER_CONNECTION_COUNT_WEB	Number	Simultaneous HTTP connections for web UI  <b>Default:</b> WORKER_CONNECTION_COUNT

Variable	Type	Description
WORKER_CONNECTION_COUNT_SECSCAN	Number	Simultaneous HTTP connections for Clair  <b>Default:</b> WORKER_CONNECTION_COUNT

## 4.4. WORKER COUNT VARIABLES

Table 4.4. Worker count variables

Variable	Type	Description
WORKER_COUNT	Number	Generic override for number of processes
WORKER_COUNT_REGISTRY	Number	Specifies the number of processes to handle Registry requests within the <b>Quay</b> container  <b>Values:</b> Integer between <b>8</b> and <b>64</b>
WORKER_COUNT_WEB	Number	Specifies the number of processes to handle UI/Web requests within the container  <b>Values:</b> Integer between <b>2</b> and <b>32</b>
WORKER_COUNT_SECSCAN	Number	Specifies the number of processes to handle Security Scanning (e.g. Clair) integration within the container  <b>Values:</b> Integer. Because the Operator specifies 2 vCPUs for resource requests and limits, setting this value between <b>2</b> and <b>4</b> is safe. However, users can run more, for example, <b>16</b> , if warranted.

## 4.5. DEBUG VARIABLES

The following debug variables are available on Red Hat Quay.

Table 4.5. Debug configuration variables

Variable	Type	Description
<b>DEBUGLOG</b>	Boolean	Whether to enable or disable debug logs.
<b>USERS_DEBUG</b>	Integer. Either <b>0</b> or <b>1</b> .	<p>Used to debug LDAP operations in clear text, including passwords. Must be used with <b>DEBUGLOG=TRUE</b>.</p> <div style="display: flex; align-items: flex-start;"> <div style="background-color: black; width: 20px; height: 100%; margin-right: 10px;"></div> <div> <p><b>IMPORTANT</b></p> <p>Setting <b>USERS_DEBUG=1</b> exposes credentials in clear text. This variable should be removed from the Red Hat Quay deployment after debugging. The log file that is generated with this environment variable should be scrutinized, and passwords should be removed before sending to other users. Use with caution.</p> </div> </div>

## CHAPTER 5. CLAIR SECURITY SCANNER

### 5.1. CLAIR CONFIGURATION OVERVIEW

Clair is configured by a structured YAML file. Each Clair node needs to specify what mode it will run in and a path to a configuration file through CLI flags or environment variables. For example:

```
$ clair -conf ./path/to/config.yaml -mode indexer
```

or

```
$ clair -conf ./path/to/config.yaml -mode matcher
```

The aforementioned commands each start two Clair nodes using the same configuration file. One runs the indexing facilities, while other runs the matching facilities.

If you are running Clair in **combo** mode, you must supply the indexer, matcher, and notifier configuration blocks in the configuration.

#### 5.1.1. Information about using Clair in a proxy environment

Environment variables respected by the Go standard library can be specified if needed, for example:

- **HTTP\_PROXY**

```
$ export HTTP_PROXY=http://<user_name>:<password>@<proxy_host>:<proxy_port>
```

- **HTTPS\_PROXY.**

```
$ export HTTPS_PROXY=https://<user_name>:<password>@<proxy_host>:<proxy_port>
```

- **SSL\_CERT\_DIR**

```
$ export SSL_CERT_DIR=/<path>/<to>/<ssl>/<certificates>
```

- **NO\_PROXY**

```
$ export NO_PROXY=<comma_separated_list_of_hosts_and_domains>
```

If you are using a proxy server in your environment with Clair's updater URLs, you must identify which URL needs to be added to the proxy allowlist to ensure that Clair can access them unimpeded. For example, the **osv** updater requires access to **https://osv-vulnerabilities.storage.googleapis.com** to fetch ecosystem data dumps. In this scenario, the URL must be added to the proxy allowlist. For a full list of updater URLs, see "Clair updater URLs".

You must also ensure that the standard Clair URLs are added to the proxy allowlist:

- **https://search.maven.org/solrsearch/select**
- **https://catalog.redhat.com/api/containers/**
- **https://access.redhat.com/security/data/metrics/repository-to-cpe.json**

- <https://access.redhat.com/security/data/metrics/container-name-repos-map.json>

When configuring the proxy server, take into account any authentication requirements or specific proxy settings needed to enable seamless communication between Clair and these URLs. By thoroughly documenting and addressing these considerations, you can ensure that Clair functions effectively while routing its updater traffic through the proxy.

### 5.1.2. Clair configuration reference

The following YAML shows an example Clair configuration:

```
http_listen_addr: ""
introspection_addr: ""
log_level: ""
tls: {}
indexer:
  connstring: ""
  scanlock_retry: 0
  layer_scan_concurrency: 5
  migrations: false
  scanner: {}
  airgap: false
matcher:
  connstring: ""
  indexer_addr: ""
  migrations: false
  period: ""
  disable_updaters: false
  update_retention: 2
matchers:
  names: nil
  config: nil
updaters:
  sets: nil
  config: nil
notifier:
  connstring: ""
  migrations: false
  indexer_addr: ""
  matcher_addr: ""
  poll_interval: ""
  delivery_interval: ""
  disable_summary: false
  webhook: null
  amqp: null
  stomp: null
auth:
  psk: nil
trace:
  name: ""
  probability: null
  jaeger:
    agent:
      endpoint: ""
    collector:
      endpoint: ""
```



```

    username: null
    password: null
    service_name: ""
    tags: nil
    buffer_max: 0
metrics:
  name: ""
  prometheus:
    endpoint: null
  dogstatsd:
    url: ""

```



## NOTE

The above YAML file lists every key for completeness. Using this configuration file as-is will result in some options not having their defaults set normally.

### 5.1.3. Clair general fields

The following table describes the general configuration fields available for a Clair deployment.

Field	Type	Description
<code>http_listen_addr</code>	String	Configures where the HTTP API is exposed.  Default: <b>:6060</b>
<code>introspection_addr</code>	String	Configures where Clair's metrics and health endpoints are exposed.
<code>log_level</code>	String	Sets the logging level. Requires one of the following strings: <b>debug-color, debug, info, warn, error, fatal, panic</b>
<code>tls</code>	String	A map containing the configuration for serving the HTTP API of TLS/SSL and HTTP/2.
<code>.cert</code>	String	The TLS certificate to be used. Must be a full-chain certificate.

#### Example configuration for general Clair fields

The following example shows a Clair configuration.

#### Example configuration for general Clair fields

```
# ...
```

```

http_listen_addr: 0.0.0.0:6060
introspection_addr: 0.0.0.0:8089
log_level: info
# ...

```

### 5.1.4. Clair indexer configuration fields

The following table describes the configuration fields for Clair's **indexer** component.

Field	Type	Description
<b>indexer</b>	Object	Provides Clair indexer node configuration.
<b>.airgap</b>	Boolean	Disables HTTP access to the internet for indexers and fetchers. Private IPv4 and IPv6 addresses are allowed. Database connections are unaffected.
<b>.connstring</b>	String	A Postgres connection string. Accepts format as a URL or libpq connection string.
<b>.index_report_request_concurrency</b>	Integer	Rate limits the number of index report creation requests. Setting this to <b>0</b> attempts to auto-size this value. Setting a negative value means unlimited. The auto-sizing is a multiple of the number of available cores.  The API returns a <b>429</b> status code if concurrency is exceeded.
<b>.scanlock_retry</b>	Integer	A positive integer representing seconds. Concurrent indexers lock on manifest scans to avoid clobbering. This value tunes how often a waiting indexer polls for the lock.
<b>.layer_scan_concurrency</b>	Integer	Positive integer limiting the number of concurrent layer scans. Indexers will match a manifest's layer concurrently. This value tunes the number of layers an indexer scans in parallel.
<b>.migrations</b>	Boolean	Whether indexer nodes handle migrations to their database.

Field	Type	Description
<code>.scanner</code>	String	Indexer configuration.  Scanner allows for passing configuration options to layer scanners. The scanner will have this configuration pass to it on construction if designed to do so.
<code>.scanner.dist</code>	String	A map with the name of a particular scanner and arbitrary YAML as a value.
<code>.scanner.package</code>	String	A map with the name of a particular scanner and arbitrary YAML as a value.
<code>.scanner.repo</code>	String	A map with the name of a particular scanner and arbitrary YAML as a value.

### Example indexer configuration

The following example shows a hypothetical indexer configuration for Clair.

### Example indexer configuration

```
# ...
indexer:
  connstring: host=quay-server.example.com port=5433 dbname=clair user=clairuser
  password=clairpass sslmode=disable
  scanlock_retry: 10
  layer_scan_concurrency: 5
  migrations: true
# ...
```

### 5.1.5. Clair matcher configuration fields

The following table describes the configuration fields for Clair's **matcher** component.



#### NOTE

Differs from **matchers** configuration fields.

Field	Type	Description
<code>matcher</code>	Object	Provides Clair matcher node configuration.

Field	Type	Description
<code>.cache_age</code>	String	Controls how long users should be hinted to cache responses for.
<code>.connstring</code>	String	A Postgres connection string. Accepts format as a URL or libpq connection string.
<code>.max_conn_pool</code>	Integer	Limits the database connection pool size.  Clair allows for a custom connection pool size. This number directly sets how many active database connections are allowed concurrently.  This parameter will be ignored in a future version. Users should configure this through the connection string.
<code>.indexer_addr</code>	String	A matcher contacts an indexer to create a vulnerability report. The location of this indexer is required.  Defaults to <b>30m</b> .
<code>.migrations</code>	Boolean	Whether matcher nodes handle migrations to their databases.
<code>.period</code>	String	Determines how often updates for new security advisories take place.  Defaults to <b>30m</b> .
<code>.disable_updaters</code>	Boolean	Whether to run background updates or not.  Default: <b>False</b>

Field	Type	Description
<code>.update_retention</code>	Integer	<p>Sets the number of update operations to retain between garbage collection cycles. This should be set to a safe MAX value based on database size constraints.</p> <p>Defaults to <b>10m</b>.</p> <p>If a value of less than <b>0</b> is provided, garbage collection is disabled. <b>2</b> is the minimum value to ensure updates can be compared to notifications.</p>


### Example matcher configuration

### Example matcher configuration

```
# ...
matcher:
  connstring: >-
    host=<DB_HOST> port=5432 dbname=<matcher> user=<DB_USER> password=D<B_PASS>
    sslmode=verify-ca sslcert=/etc/clair/ssl/cert.pem sslkey=/etc/clair/ssl/key.pem
    sslrootcert=/etc/clair/ssl/ca.pem
  indexer_addr: http://clair-v4/
  disable_updaters: false
  migrations: true
  period: 6h
  update_retention: 2
# ...
```

### 5.1.6. Clair matchers configuration fields

The following table describes the configuration fields for Clair's **matchers** component.



#### NOTE

Differs from **matcher** configuration fields.

Table 5.1. Matchers configuration fields

Field	Type	Description
<code>matchers</code>	Array of strings	Provides configuration for the in-tree <b>matchers</b> .

Field	Type	Description
<code>.names</code>	String	A list of string values informing the matcher factory about enabled matchers. If value is set to <b>null</b> , the default list of matchers run. The following strings are accepted: <b>alpine-matcher</b> , <b>aws-matcher</b> , <b>debian-matcher</b> , <b>gobin</b> , <b>java-maven</b> , <b>oracle</b> , <b>photon</b> , <b>python</b> , <b>rhel</b> , <b>rhel-container-matcher</b> , <b>ruby</b> , <b>suse</b> , <b>ubuntu-matcher</b>
<code>.config</code>	String	Provides configuration to a specific matcher.  A map keyed by the name of the matcher containing a sub-object which will be provided to the matchers factory constructor. For example:

### Example matchers configuration

The following example shows a hypothetical Clair deployment that only requires only the **alpine**, **aws**, **debian**, **oracle** matchers.

### Example matchers configuration

```
# ...
matchers:
  names:
    - "alpine-matcher"
    - "aws"
    - "debian"
    - "oracle"
# ...
```

## 5.1.7. Clair updaters configuration fields

The following table describes the configuration fields for Clair's **updaters** component.

Table 5.2. Updaters configuration fields

Field	Type	Description
<code>updaters</code>	Object	Provides configuration for the matcher's update manager.

Field	Type	Description
<code>.sets</code>	String	<p>A list of values informing the update manager which updaters to run.</p> <p>If value is set to <b>null</b>, the default set of updaters runs the following: <b>alpine, aws, clair.cvss, debian, oracle, photon, osv, rhel, rhcc suse, ubuntu</b></p> <p>If left blank, zero updaters run.</p>
<code>.config</code>	String	<p>Provides configuration to specific updater sets.</p> <p>A map keyed by the name of the updater set containing a sub-object which will be provided to the updater set's constructor. For a list of the sub-objects for each updater, see "Advanced updater configuration".</p>

### Example updaters configuration

In the following configuration, only the **rhel** set is configured. The **ignore\_unpatched** variable, which is specific to the **rhel** updater, is also defined.

### Example updaters configuration

```
# ...
updaters:
  sets:
    - rhel
  config:
    rhel:
      ignore_unpatched: false
# ...
```

## 5.1.8. Clair notifier configuration fields

The general notifier configuration fields for Clair are listed below.

Field	Type	Description
<code>notifier</code>	Object	Provides Clair notifier node configuration.
<code>.connstring</code>	String	Postgres connection string. Accepts format as URL, or libpq connection string.

Field	Type	Description
<code>.migrations</code>	Boolean	Whether notifier nodes handle migrations to their database.
<code>.indexer_addr</code>	String	A notifier contacts an indexer to create or obtain manifests affected by vulnerabilities. The location of this indexer is required.
<code>.matcher_addr</code>	String	A notifier contacts a matcher to list update operations and acquire diffs. The location of this matcher is required.
<code>.poll_interval</code>	String	The frequency at which the notifier will query a matcher for update operations.
<code>.delivery_interval</code>	String	The frequency at which the notifier attempts delivery of created, or previously failed, notifications.
<code>.disable_summary</code>	Boolean	Controls whether notifications should be summarized to one per manifest.

### Example notifier configuration

The following **notifier** snippet is for a minimal configuration.

### Example notifier configuration

```
# ...
notifier:
  connstring: >-
    host=DB_HOST port=5432 dbname=notifier user=DB_USER password=DB_PASS
    sslmode=verify-ca sslcert=/etc/clair/ssl/cert.pem sslkey=/etc/clair/ssl/key.pem
    sslrootcert=/etc/clair/ssl/ca.pem
  indexer_addr: http://clair-v4/
  matcher_addr: http://clair-v4/
  delivery_interval: 5s
  migrations: true
  poll_interval: 15s
  webhook:
    target: "http://webhook/"
    callback: "http://clair-notifier/notifier/api/v1/notifications"
    headers: ""
  amqp: null
  stomp: null
# ...
```



### 5.1.8.1. Clair webhook configuration fields

The following webhook fields are available for the Clair notifier environment.

**Table 5.3. Clair webhook fields**

<code>.webhook</code>	Object	Configures the notifier for webhook delivery.
<code>.webhook.target</code>	String	URL where the webhook will be delivered.
<code>.webhook.callback</code>	String	The callback URL where notifications can be retrieved. The notification ID will be appended to this URL.  This will typically be where the Clair notifier is hosted.
<code>.webhook.headers</code>	String	A map associating a header name to a list of values.

#### Example webhook configuration

#### Example webhook configuration

```
# ...
notifier:
# ...
  webhook:
    target: "http://webhook/"
    callback: "http://clair-notifier/notifier/api/v1/notifications"
# ...
```

### 5.1.8.2. Clair amqp configuration fields

The following Advanced Message Queuing Protocol (AMQP) fields are available for the Clair notifier environment.

<code>.amqp</code>	Object	Configures the notifier for AMQP delivery.  [NOTE] ==== Clair does not declare any AMQP components on its own. All attempts to use an exchange or queue are passive only and will fail. Broker administrators should setup exchanges and queues ahead of time. ====
--------------------	--------	---

<code>.amqp.direct</code>	Boolean	If <b>true</b> , the notifier will deliver individual notifications (not a callback) to the configured AMQP broker.
<code>.amqp.rollup</code>	Integer	When <b>amqp.direct</b> is set to <b>true</b> , this value informs the notifier of how many notifications to send in a direct delivery. For example, if <b>direct</b> is set to <b>true</b> , and <b>amqp.rollup</b> is set to <b>5</b> , the notifier delivers no more than 5 notifications in a single JSON payload to the broker. Setting the value to <b>0</b> effectively sets it to <b>1</b> .
<code>.amqp.exchange</code>	Object	The AMQP exchange to connect to.
<code>.amqp.exchange.name</code>	String	The name of the exchange to connect to.
<code>.amqp.exchange.type</code>	String	The type of the exchange. Typically one of the following: <b>direct</b> , <b>fanout</b> , <b>topic</b> , <b>headers</b> .
<code>.amqp.exchange.durability</code>	Boolean	Whether the configured queue is durable.
<code>.amqp.exchange.auto_delete</code>	Boolean	Whether the configured queue uses an <b>auto_delete_policy</b> .
<code>.amqp.routing_key</code>	String	The name of the routing key each notification is sent with.
<code>.amqp.callback</code>	String	If <b>amqp.direct</b> is set to <b>false</b> , this URL is provided in the notification callback sent to the broker. This URL should point to Clair's notification API endpoint.
<code>.amqp.uris</code>	String	A list of one or more AMQP brokers to connect to, in priority order.
<code>.amqp.tls</code>	Object	Configures TLS/SSL connection to an AMQP broker.

<code>.amqp.tls.root_ca</code>	String	The filesystem path where a root CA can be read.
<code>.amqp.tls.cert</code>	String	The filesystem path where a TLS/SSL certificate can be read.  [NOTE] ==== Clair also allows <b>SSL_CERT_DIR</b> , as documented for the Go <b>crypto/x509</b> package. ====
<code>.amqp.tls.key</code>	String	The filesystem path where a TLS/SSL private key can be read.

### Example AMQP configuration

The following example shows a hypothetical AMQP configuration for Clair.

### Example AMQP configuration

```
# ...
notifier:
# ...
  amqp:
    exchange:
      name: ""
      type: "direct"
      durable: true
      auto_delete: false
    uris: ["amqp://user:pass@host:10000/vhost"]
    direct: false
    routing_key: "notifications"
    callback: "http://clair-notifier/notifier/api/v1/notifications"
    tls:
      root_ca: "optional/path/to/rootca"
      cert: "mandatory/path/to/cert"
      key: "mandatory/path/to/key"
# ...
```

### 5.1.8.3. Clair STOMP configuration fields

The following Simple Text Oriented Message Protocol (STOMP) fields are available for the Clair notifier environment.

<code>.stomp</code>	Object	Configures the notifier for STOMP delivery.
<code>.stomp.direct</code>	Boolean	If <b>true</b> , the notifier delivers individual notifications (not a callback) to the configured STOMP broker.

.stomp	Object	Configures the notifier for STOMP delivery.
.stomp.rollup	Integer	If <b>stomp.direct</b> is set to <b>true</b> , this value limits the number of notifications sent in a single direct delivery. For example, if <b>direct</b> is set to <b>true</b> , and <b>rollup</b> is set to <b>5</b> , the notifier delivers no more than 5 notifications in a single JSON payload to the broker. Setting the value to <b>0</b> effectively sets it to <b>1</b> .
.stomp.callback	String	If <b>stomp.callback</b> is set to <b>false</b> , the provided URL in the notification callback is sent to the broker. This URL should point to Clair's notification API endpoint.
.stomp.destination	String	The STOMP destination to deliver notifications to.
.stomp.uris	String	A list of one or more STOMP brokers to connect to in priority order.
.stomp.tls	Object	Configured TLS/SSL connection to STOMP broker.
.stomp.tls.root_ca	String	The filesystem path where a root CA can be read.  [NOTE] ==== Clair also respects <b>SSL_CERT_DIR</b> , as documented for the Go <b>crypto/x509</b> package. ====
.stomp.tls.cert	String	The filesystem path where a TLS/SSL certificate can be read.
.stomp.tls.key	String	The filesystem path where a TLS/SSL private key can be read.
.stomp.user	String	Configures login details for the STOMP broker.
.stomp.user.login	String	The STOMP login to connect with.

<code>.stomp</code>	Object	Configures the notifier for STOMP delivery.
<code>.stomp.user.passcode</code>	String	The STOMP passcode to connect with.

### Example STOMP configuration

The following example shows a hypothetical STOMP configuration for Clair.

### Example STOMP configuration

```
# ...
notifier:
# ...
stomp:
  desitnation: "notifications"
  direct: false
  callback: "http://clair-notifier/notifier/api/v1/notifications"
  login:
    login: "username"
    passcode: "passcode"
  tls:
    root_ca: "optional/path/to/rootca"
    cert: "madatory/path/to/cert"
    key: "madatory/path/to/key"
# ...
```

### 5.1.9. Clair authorization configuration fields

The following authorization configuration fields are available for Clair.

Field	Type	Description
<code>auth</code>	Object	Defines Clair's external and intra-service JWT based authentication. If multiple <b>auth</b> mechanisms are defined, Clair picks one. Currently, multiple mechanisms are unsupported.
<code>.psk</code>	String	Defines pre-shared key authentication.
<code>.psk.key</code>	String	A shared base64 encoded key distributed between all parties signing and verifying JWTs.
<code>.psk.iss</code>	String	A list of JWT issuers to verify. An empty list accepts any issuer in a JWT claim.

### Example authorization configuration

The following **authorization** snippet is for a minimal configuration.

### Example authorization configuration

```
# ...
auth:
  psk:
    key: MTU5YzA4Y2ZkNzJoMQ== 1
    iss: ["quay"]
# ...
```

### 5.1.10. Clair trace configuration fields

The following trace configuration fields are available for Clair.

Field	Type	Description
<b>trace</b>	Object	Defines distributed tracing configuration based on OpenTelemetry.
<b>.name</b>	String	The name of the application traces will belong to.
<b>.probability</b>	Integer	The probability a trace will occur.
<b>.jaeger</b>	Object	Defines values for Jaeger tracing.
<b>.jaeger.agent</b>	Object	Defines values for configuring delivery to a Jaeger agent.
<b>.jaeger.agent.endpoint</b>	String	An address in the <b>&lt;host&gt;: &lt;port&gt;</b> syntax where traces can be submitted.
<b>.jaeger.collector</b>	Object	Defines values for configuring delivery to a Jaeger collector.
<b>.jaeger.collector.endpoint</b>	String	An address in the <b>&lt;host&gt;: &lt;port&gt;</b> syntax where traces can be submitted.
<b>.jaeger.collector.username</b>	String	A Jaeger username.
<b>.jaeger.collector.password</b>	String	A Jaeger password.
<b>.jaeger.service_name</b>	String	The service name registered in Jaeger.

Field	Type	Description
<code>.jaeger.tags</code>	String	Key-value pairs to provide additional metadata.
<code>.jaeger.buffer_max</code>	Integer	The maximum number of spans that can be buffered in memory before they are sent to the Jaeger backend for storage and analysis.

### Example trace configuration

The following example shows a hypothetical trace configuration for Clair.

### Example trace configuration

```
# ...
trace:
  name: "jaeger"
  probability: 1
  jaeger:
    agent:
      endpoint: "localhost:6831"
      service_name: "clair"
# ...
```

### 5.1.11. Clair metrics configuration fields

The following metrics configuration fields are available for Clair.

Field	Type	Description
<code>metrics</code>	Object	Defines distributed tracing configuration based on OpenTelemetry.
<code>.name</code>	String	The name of the metrics in use.
<code>.prometheus</code>	String	Configuration for a Prometheus metrics exporter.
<code>.prometheus.endpoint</code>	String	Defines the path where metrics are served.

### Example metrics configuration

The following example shows a hypothetical metrics configuration for Clair.

### Example metrics configuration

■

```
# ...  
metrics:  
  name: "prometheus"  
  prometheus:  
    endpoint: "/metricsz"  
# ...
```