



Red Hat Subscription Management 1

Using Discovery

Understanding Discovery

Red Hat Subscription Management 1 Using Discovery

Understanding Discovery

Legal Notice

Copyright © 2019 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[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

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

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

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

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Table of Contents

PART I. ABOUT DISCOVERY	4
CHAPTER 1. WHAT IS DISCOVERY?	5
CHAPTER 2. WHAT PRODUCTS DOES DISCOVERY FIND?	6
CHAPTER 3. IS DISCOVERY RIGHT FOR ME?	7
PART II. ACCESSING THE DISCOVERY USER INTERFACE	8
CHAPTER 4. LOGGING IN TO THE DISCOVERY USER INTERFACE	9
CHAPTER 5. LOGGING OUT OF THE DISCOVERY USER INTERFACE	10
PART III. ADDING SOURCES AND CREDENTIALS	11
CHAPTER 6. ADDING NETWORK SOURCES AND CREDENTIALS	12
6.1. ADDING NETWORK SOURCES	12
6.2. ADDING NETWORK CREDENTIALS	13
6.3. ABOUT SOURCES AND CREDENTIALS	14
6.4. NETWORK AUTHENTICATION	14
6.4.1. Commands that are used in scans of remote network assets	15
6.4.1.1. Basic commands that do not need elevated privileges	15
6.4.1.2. Commands that need elevated privileges	16
CHAPTER 7. ADDING SATELLITE SOURCES AND CREDENTIALS	18
7.1. ADDING SATELLITE SOURCES	18
7.2. ADDING SATELLITE CREDENTIALS	19
7.3. ABOUT SOURCES AND CREDENTIALS	19
7.4. SATELLITE SERVER AUTHENTICATION	20
CHAPTER 8. ADDING VCENTER SOURCES AND CREDENTIALS	21
8.1. ADDING VCENTER SOURCES	21
8.2. ADDING VCENTER CREDENTIALS	22
8.3. ABOUT SOURCES AND CREDENTIALS	22
8.4. VCENTER SERVER AUTHENTICATION	23
PART IV. RUNNING AND MANAGING SCANS	24
CHAPTER 9. RUNNING AND MANAGING STANDARD SCANS	25
9.1. RUNNING STANDARD SCANS	25
9.2. RUNNING A NEW SCAN JOB	26
9.3. PAUSING, RESUMING, AND CANCELING SCANS	26
9.4. ABOUT SCANS AND SCAN JOBS	27
9.5. SCAN JOB PROCESSING	27
9.5.1. Scan job connection and inspection tasks	27
9.5.2. How these tasks are processed	27
9.6. SCAN JOB LIFE CYCLE	28
CHAPTER 10. RUNNING AND MANAGING DEEP SCANS	29
10.1. RUNNING SCANS WITH DEEP SCANNING	29
10.2. RUNNING A NEW SCAN JOB	30
10.3. PAUSING, RESUMING, AND CANCELING SCANS	30
10.4. ABOUT SCANS AND SCAN JOBS	31
10.5. SCAN JOB PROCESSING	31

10.5.1. Scan job connection and inspection tasks	31
10.5.2. How these tasks are processed	32
10.6. SCAN JOB LIFE CYCLE	32
PART V. MERGING AND DOWNLOADING REPORTS	33
CHAPTER 11. DOWNLOADING REPORTS	34
11.1. DOWNLOADING REPORTS	34
11.2. HOW REPORTS ARE CREATED	35
11.2.1. Facts and fingerprints	35
11.2.2. System deduplication and merging	35
11.2.2.1. System deduplication	35
11.2.2.2. System merging	36
11.2.3. System post-processing	36
11.2.4. Report creation	36
11.2.5. A fingerprint example	37
CHAPTER 12. MERGING REPORTS	39
12.1. MERGING REPORTS	39
12.2. HOW REPORTS ARE CREATED	40
12.2.1. Facts and fingerprints	40
12.2.2. System deduplication and merging	40
12.2.2.1. System deduplication	41
12.2.2.2. System merging	41
12.2.3. System post-processing	41
12.2.4. Report creation	42
12.2.5. A fingerprint example	42

PART I. ABOUT DISCOVERY

The product discovery tool is designed to help users collect data about their usage of specific Red Hat software. By using discovery, users can reduce the amount of time and effort that is required to calculate and report usage of those Red Hat products.

Learn more

To learn more about the purpose, benefits, and characteristics of discovery, see the following information:

- [What is discovery?](#)

To learn more about the products and product versions that discovery can find and inspect, see the following information:

- [What products does discovery find?](#)

To evaluate whether discovery is a correct solution for you, see the following information:

- [Is discovery right for me?](#)

CHAPTER 1. WHAT IS DISCOVERY?

The product discovery tool, also known as discovery, is an inspection and reporting tool. It is designed to find, identify, and report environment data, or facts, such as the number of physical and virtual systems on a network, their operating systems, and other configuration data. In addition, it is designed to find, identify, and report more detailed facts for some versions of key Red Hat packages and products for the IT resources in that network.

The ability to inspect the software and systems that are running on your network improves your ability to understand and report on your entitlement usage. Ultimately, this inspection and reporting process is part of the larger system administration task of managing your inventories.

The product discovery tool requires the configuration of two basic structures to access IT resources and run the inspection process. A *credential* contains user access data, such as the user name and password or SSH key of a user with sufficient authority to run the inspection process on a particular source or some of the assets on that source. A *source* contains data about a single asset or multiple assets that are to be inspected. These assets can be physical machines, virtual machines, or containers, identified as host names, IP addresses, IP ranges, or subnets. These assets can also be a systems management solution such as vCenter Server or Red Hat Satellite Server.

You can save multiple credentials and sources to use with discovery in various combinations as you run inspection processes, or *scans*. When you have completed a scan, you can access these facts in the output as a collection of formatted data, or *report*, to review the results.

By default, the credentials and sources that are created during the use of discovery are encrypted in a database. The values are encrypted with AES-256 encryption. They are decrypted when the discovery server runs a scan with the use of a vault password to access the encrypted values that are stored in the database.

The product discovery tool is an agentless inspection tool, so there is no need to install the tool on every source that is to be inspected. However, the system that discovery is installed on must have access to the systems to be discovered and inspected.

CHAPTER 2. WHAT PRODUCTS DOES DISCOVERY FIND?

The product discovery tool finds the following Red Hat products. For each version or release, the earliest version is listed, with later releases indicated as applicable.

If a product has changed names recently so that you might be more familiar with the current name for that product, that name is provided as additional information. No later version is implied by the inclusion of a newer product name unless specific versions of that product are also listed.

Red Hat Enterprise Linux

- Red Hat Enterprise Linux version 5 and later
- Red Hat Enterprise Linux version 6 and later
- Red Hat Enterprise Linux version 7 and later
- Red Hat Enterprise Linux version 8 and later

Red Hat Middleware products

- Red Hat JBoss BRMS version 5.0.1 and later, version 6.0.0 and later (current product name is Red Hat Decision Manager)
- JBoss Enterprise Web Server version 1 and later, version 2.0 and later, version 2.1.0 and later; Red Hat JBoss Web Server 3.0.1 and later, 3.1 and later, version 5.0.0
- Red Hat JBoss Enterprise Application Platform version 4.2 and later, version 4.3 and later, version 5 and later, version 6 and later, version 7
- Red Hat Fuse version 6.0 and later, version 6.1 and later, version 6.2 and later, version 6.3.0

CHAPTER 3. IS DISCOVERY RIGHT FOR ME?

The product discovery tool is intended to help you find and understand your Red Hat product inventory, including unknown product usage across complex networks. The reports generated by discovery are best understood through your partnership with a Red Hat Solution Architect (SA) or Technical Account Manager (TAM) or through the analysis and assistance supplied by the Subscription Education and Awareness Program (SEAP). Pilot programs that are currently underway include discovery as one tool that can help integrate your software inventory with other new and established Red Hat management offerings.

PART II. ACCESSING THE DISCOVERY USER INTERFACE

You access the discovery graphical user interface through a browser connection.

To use the discovery user interface, the system on which you want to run the user interface must be able to communicate with the system on which the discovery server is installed.

Learn more

To learn more about the requirements and steps to log in to and out of the discovery graphical user interface, see the following information:

- [Logging in to the discovery user interface](#)
- [Logging out of the discovery user interface](#)

CHAPTER 4. LOGGING IN TO THE DISCOVERY USER INTERFACE

Prerequisites

To log in to the discovery user interface, you need the IP address of the system where the discovery server is installed, the port number for the connection if the default port was changed during server installation, and the user name and password to use when logging in. If you do not have this information, contact the administrator who installed the discovery server.

Procedure

1. In a browser, enter the URL for the discovery server in the following format:
https://IPaddress:port, where **IPaddress** is the IP address of the discovery server and **port** is the exposed server port.

The following examples show two different ways to enter the URL, based on the system that you are logging in from and whether the default port is used:

- If you log in from the system where the server is installed and the default port is used, you can use the loopback address (also known as localhost) as the IP address, as shown in the following example:

```
https://127.0.0.1:9443
```

- If you log in from a system that is remote from the server, the server is running on the IP address **192.0.2.0**, and the default port was changed during installation to **8443**, you would log in as shown in the following example:

```
https://192.0.2.0:8443
```

After you enter the URL for the server, the discovery login page displays.

2. On the login page, enter the user name and password and then click **Log in** to log in to the server.

Verification steps

If this is the first time that you have logged in to discovery, the Welcome page displays. You can begin by adding sources and credentials that can be used in scans. If you have previously logged in to discovery, the Welcome page is skipped and you can interact with your previously created sources, credentials, and scans.

CHAPTER 5. LOGGING OUT OF THE DISCOVERY USER INTERFACE

Procedure

1. In the application toolbar, click the person icon or your user name.
2. Click **Logout**.

PART III. ADDING SOURCES AND CREDENTIALS

To prepare discovery to run scans, you must add the parts of your IT infrastructure that you want to scan as one or more sources. You must also add the authentication information, such as a username and password or SSH key, that is required to access those sources as one or more credentials. Because of differing configuration requirements, you add sources and credentials according to the type of source that you are going to scan.

Learn more

As part of the general process of adding sources and credentials that encompass the different parts of your IT infrastructure, you might need to complete a number of tasks.

Add network sources and credentials to scan assets such as physical machines, virtual machines, or containers in your network. To learn more, see the following information:

- [Adding network sources and credentials](#)

Add satellite sources and credentials to scan your deployment of Red Hat Satellite Server to find the assets that it manages. To learn more, see the following information:

- [Adding satellite sources and credentials](#)

Add vcenter sources and credentials to scan your deployment of vCenter Server to find the assets that it manages. To learn more, see the following information:

- [Adding vcenter sources and credentials](#)

CHAPTER 6. ADDING NETWORK SOURCES AND CREDENTIALS

To run a scan on one or more of the physical machines, virtual machines, or containers on your network, you must add a source that identifies each of the assets to scan. Then you must add credentials that contain the authentication data to access each asset.

Learn more

Add one or more network sources and credentials to provide the information needed to scan the assets in your network. To learn more, see the following information:

- To add a network source, see [Adding network sources](#).
- To add a network credential, see [Adding network credentials](#).

To learn more about sources and credentials and how discovery uses them, see the following information:

- [About sources and credentials](#)

To learn more about how discovery authenticates with assets on your network, see the following information. This information includes guidance about running commands with elevated privileges, a choice that you might need to make during network credential configuration:

- [Network authentication](#)
- [Commands that are used in scans of remote network assets](#)

6.1. ADDING NETWORK SOURCES

You can add sources from the initial Welcome page or from the Sources view.

Procedure

1. Click the option to add a new credential based on your location:
 - From the Welcome page, click **Add Source**.
 - From the Sources view, click **Add**.

The Add Source wizard opens.

2. On the Type page, select **Network Range** as the source type and click **Next**.
3. On the Credentials page, enter the following information.
 - a. In the **Name** field, enter a descriptive name.
 - b. In the **Search Addresses** field, enter one or more network identifiers separated by commas. You can enter host names, IP addresses, and IP ranges.
 - Enter host names as DNS host names.
 - Enter IP ranges in CIDR or Ansible notation.

- c. Optional: In the **Port** field, enter a different port if you do not want a scan for this source to run on the default port 22.
 - d. In the **Credentials** list, select the credentials that are required to access the network resources for this source. If a required credential does not exist, click the **Add a credential** icon to open the Add Credential wizard.
 - e. If your network resources require the Ansible connection method to be the Python SSH implementation, Paramiko, instead of the default OpenSSH implementation, select the **Connect using Paramiko instead of OpenSSH** check box.
4. Click **Save** to save the source and then click **Close** to close the Add Source wizard.

6.2. ADDING NETWORK CREDENTIALS

You can add credentials from the Credentials view or from the Add Source wizard during the creation of a source. You might need to add several credentials to authenticate to all of the assets that are included in a single source.

Prerequisites

- If you want to use the SSH key authentication type for network credentials, each SSH private key that you are going to use must be copied into the directory that was mapped to **/sshkeys** during discovery server installation. The default path for this directory is **~/discovery/server/volumes/sshkeys**.
For more information about the SSH keys that are available for use in the **/sshkeys** directory, or to request the addition of a key to that directory, contact the administrator who manages your discovery server.

Procedure

1. Click the option to add a new credential based on your location:
 - From the Credentials view, click **Add → Network Credential**.
 - From the Add Source wizard, click the **Add a credential** icon for the **Credentials** field.

The Add Credential wizard opens.
2. In the **Credential Name** field, enter a descriptive name.
3. In the **Authentication Type** field, select the type of authentication that you want to use. You can select either **Username and Password** or **SSH Key**.
4. Enter the authentication data in the appropriate fields, based on the authentication type.
 - For username and password authentication, enter a username and password for a user. This user must have root-level access to your network or to the subset of your network that you want to scan. Alternatively, this user must be able to obtain root-level access with the selected become method.
 - For SSH key authentication, enter a user name and the path to an SSH key file, where the path to the key file is a path that is local to the discovery server. For example, if the key file is in the **~/discovery/server/volumes/sshkeys** path on the server, enter that path in the **SSH Key File** field. Entering a passphrase is optional.

5. Enter the become method for privilege elevation. Privilege elevation is required to run some commands during a network scan. Entering a username and password for the become method is optional.
6. Click **Save** to save the credential and close the Add Credential wizard.

6.3. ABOUT SOURCES AND CREDENTIALS

To run a scan, you must configure data for two basic structures, sources and credentials. The type of source that you are going to inspect during the scan determines the type of data that is required for both source and credential configuration.

A *source* contains a single asset or a set of multiple assets that are to be inspected during the scan. You can configure three types of sources:

- A network source: One or more physical machines, virtual machines, or containers. These assets can be expressed as host names, IP addresses, IP ranges, or subnets.
- A vcenter source: A vCenter Server systems management solution that is managing all or part of your IT infrastructure.
- A satellite source: A Satellite systems management solution that is managing all or part of your IT infrastructure.

When you are working with network sources, you determine how many individual assets you should group within a single source. The following list contains some of the factors that you should consider when you are adding sources:

- Whether assets are part of a development, testing, or production environment, and if demands on computing power and similar concerns are a consideration for those assets.
- Whether you want to scan a particular entity or group of entities more often because of internal business practices such as frequent changes to the installed software.

A *credential* contains data such as the user name and password or SSH key of a user with sufficient authority to run the scan on all or part of the assets that are contained in that source. As with sources, credentials are configured as the network, vcenter, or satellite type. Typically, a network source might require multiple network credentials because it is expected that many credentials would be needed to access all of the assets in a broad IP range. Conversely, a vcenter or satellite source would typically use a single vcenter or satellite credential, as applicable, to access a particular system management solution server.

You can add new sources from the Sources view and you can add new credentials from the Credentials view. You can also add new or select previously existing credentials during source creation. It is during source creation that you associate a credential directly with a source. Because sources and credentials must have matching types, any credential that you add during source creation shares the same type as the source. In addition, if you want to use an existing credential during source creation, the list of available credentials contains only credentials of the same type. For example, during network source creation, only network credentials are available for selection.

6.4. NETWORK AUTHENTICATION

The discovery server inspects the remote systems in a network scan by using the SSH remote connection capabilities of Ansible. When you add a network credential, you configure the SSH connection by using either a user name and password or a user name and SSH keyfile pair. If remote systems are accessed with SSH key authentication, you can also supply a passphrase.

Also during network credential configuration, you can enable a become method. The become method is used during a scan to elevate privileges. These elevated privileges are needed to run commands and obtain data on the systems that you are scanning. For more information about the commands that do and do not require elevated privileges during a scan, see [Commands that are used in scans of remote network assets](#).

6.4.1. Commands that are used in scans of remote network assets

When you run a network scan, discovery must use the credentials that you provide to run certain commands on the remote systems in your network. Some of those commands must run with elevated privileges. This access is typically acquired through the use of the **sudo** command or similar commands. The elevated privileges are required to gather the types of facts that discovery uses to build the report about your installed products and consumed entitlements.

Although it is possible to run a scan for a network source without elevated privileges, the results of that scan will be incomplete. The incomplete results from the network scan will affect the quality of the generated report for the scan.

The following information lists the commands that discovery runs on remote hosts during a network scan. The information includes the basic commands that can run without elevated privileges and the commands that must run with elevated privileges to gather the most accurate and complete information for the report.



NOTE

In addition to the following commands, discovery also depends on standard shell facilities, such as those provided by the Bash shell.

6.4.1.1. Basic commands that do not need elevated privileges

The following commands do not require elevated privileges to gather facts during a scan:

- cat
- egrep
- sort
- uname
- ctime
- grep
- rpm
- virsh
- date
- id
- test
- whereis

- echo
- sed
- tune2fs
- xargs

6.4.1.2. Commands that need elevated privileges

The following commands require elevated privileges to gather facts during a scan. Each command includes a list of individual facts or categories of facts that discovery attempts to find during a scan. These facts cannot be included in reports if elevated privileges are not available for that command.

- **chkconfig**
 - EAP
 - Fuse on Karaf
- **command**
 - see dmicode
 - see subscription-manager
- **dmidecode**
 - cpu_socket_count
 - dmi_bios_vendor
 - dmi_bios_version
 - dmi_system_manufacturer
 - dmi_processor_family
 - dmi_system_uuid
 - virt_type
- **find**
 - BRMS
 - EAP
 - Fuse
 - Fuse on Karaf
- **ifconfig**
 - IP address
 - MAC address

- **java**
 - EAP info
- **locate**
 - BRMS
 - EAP
 - Fuse on Karaf
- **ls**
 - date_machine_id
 - EAP
 - Fuse on Karaf
 - BRMS
 - redhat_packages_certs
 - subman_consumed
- **ps**
 - EAP
 - Fuse on Karaf
 - virt type
- **subscription-manager**
 - subman_consumed
- **systemctl**
 - EAP
 - Fuse on Karaf
- **unzip**
 - EAP detection
- **virt-what**
 - virt_what_type
- **yum**
 - date_yum_history
 - yum_enabled_repolist

CHAPTER 7. ADDING SATELLITE SOURCES AND CREDENTIALS

To run a scan on a Red Hat Satellite Server deployment, you must add a source that identifies the Satellite Server server to scan. Then you must add a credential that contains the authentication data to access that server.

Learn more

Add a satellite source and credential to provide the information needed to scan Satellite Server. To learn more, see the following information:

- To add a satellite source, see [Adding satellite sources](#).
- To add a satellite credential, see [Adding satellite credentials](#).

To learn more about sources and credentials and how discovery uses them, see the following information:

- [About sources and credentials](#)

To learn more about how discovery authenticates with your Satellite Server server, see the following information. This information includes guidance about certificate validation and SSL communication choices that you might need to make during satellite credential configuration.

- [Satellite Server authentication](#)

7.1. ADDING SATELLITE SOURCES

You can add sources from the initial Welcome page or from the Sources view.

Procedure

1. Click the option to add a new credential based on your location:
 - From the Welcome page, click **Add Source**.
 - From the Sources view, click **Add**.

The Add Source wizard opens.

2. On the Type page, select **Satellite** as the source type and click **Next**.
3. On the Credentials page, enter the following information.
 - a. In the **Name** field, enter a descriptive name.
 - b. In the **IP Address or Hostname** field, enter the IP address or host name of the Satellite server for this source. Enter a different port if you do not want a scan for this source to run on the default port 443. For example, if the IP address of the Satellite server is 192.0.2.15 and you want to change the port to 80, you would enter **192.0.2.15:80**.
 - c. In the **Credentials** list, select the credential that is required to access the Satellite server for this source. If a required credential does not exist, click the **Add a credential** icon to open the Add Credential wizard.

- d. In the **Connection** list, select the SSL protocol to be used for a secure connection during a scan of this source.

**NOTE**

Satellite Server does not support the disabling of SSL. If you select the **Disable SSL** option, this option is ignored.

- e. If you need to upgrade the SSL validation for the Satellite server to check for a verified SSL certificate from a certificate authority, select the **Verify SSL Certificate** check box.
4. Click **Save** to save the source and then click **Close** to close the Add Source wizard.

7.2. ADDING SATELLITE CREDENTIALS

You can add credentials from the Credentials view or from the Add Source wizard during the creation of a source.

Procedure

1. Click the option to add a new credential based on your location:
 - From the Credentials view, click **Add → Satellite Credential**.
 - From the Add Source wizard, click the **Add a credential** icon for the **Credentials** field.

The Add Credential wizard opens.

2. In the **Credential Name** field, enter a descriptive name.
3. Enter the username and password for a Satellite Server administrator.
4. Click **Save** to save the credential and close the Add Credential wizard.

7.3. ABOUT SOURCES AND CREDENTIALS

To run a scan, you must configure data for two basic structures, sources and credentials. The type of source that you are going to inspect during the scan determines the type of data that is required for both source and credential configuration.

A *source* contains a single asset or a set of multiple assets that are to be inspected during the scan. You can configure three types of sources:

- A network source: One or more physical machines, virtual machines, or containers. These assets can be expressed as host names, IP addresses, IP ranges, or subnets.
- A vcenter source: A vCenter Server systems management solution that is managing all or part of your IT infrastructure.
- A satellite source: A Satellite systems management solution that is managing all or part of your IT infrastructure.

When you are working with network sources, you determine how many individual assets you should group within a single source. The following list contains some of the factors that you should consider when you are adding sources:

- Whether assets are part of a development, testing, or production environment, and if demands on computing power and similar concerns are a consideration for those assets.
- Whether you want to scan a particular entity or group of entities more often because of internal business practices such as frequent changes to the installed software.

A *credential* contains data such as the user name and password or SSH key of a user with sufficient authority to run the scan on all or part of the assets that are contained in that source. As with sources, credentials are configured as the network, vcenter, or satellite type. Typically, a network source might require multiple network credentials because it is expected that many credentials would be needed to access all of the assets in a broad IP range. Conversely, a vcenter or satellite source would typically use a single vcenter or satellite credential, as applicable, to access a particular system management solution server.

You can add new sources from the Sources view and you can add new credentials from the Credentials view. You can also add new or select previously existing credentials during source creation. It is during source creation that you associate a credential directly with a source. Because sources and credentials must have matching types, any credential that you add during source creation shares the same type as the source. In addition, if you want to use an existing credential during source creation, the list of available credentials contains only credentials of the same type. For example, during network source creation, only network credentials are available for selection.

7.4. SATELLITE SERVER AUTHENTICATION

For a satellite scan, the connectivity and access to Satellite Server derives from basic authentication (user name and password) that is encrypted over HTTPS. By default, the satellite scan runs with certificate validation and secure communication through the SSL (Secure Sockets Layer) protocol. During source creation, you can select from several different SSL and TLS (Transport Layer Security) protocols to use for the certificate validation and secure communication.

You might need to adjust the level of certificate validation to connect properly to the Satellite server during a scan. For example, your Satellite server might use a verified SSL certificate from a certificate authority. During source creation, you can upgrade SSL certificate validation to check for that certificate during a scan of that source. Conversely, your Satellite server might use self-signed certificates. During source creation, you can leave the SSL validation at the default so that scan of that source does not check for a certificate. This choice, to leave the option at the default for a self-signed certificate, could possibly avoid scan errors.

Although the option to disable SSL is currently available in the interface, Satellite Server does not support the disabling of SSL. If you select the **Disable SSL** option when you create a satellite source, this option is ignored.

CHAPTER 8. ADDING VCENTER SOURCES AND CREDENTIALS

To run a scan on a vCenter Server deployment, you must add a source that identifies the vCenter Server server to scan. Then you must add a credential that contains the authentication data to access that server.

Learn more

Add a vcenter source and credential to provide the information needed to scan vCenter Server. To learn more, see the following information:

- To add a vcenter source, see [Adding vcenter sources](#).
- To add a vcenter credential, see [Adding vcenter credentials](#).

To learn more about sources and credentials and how discovery uses them, see the following information:

- [About sources and credentials](#)

To learn more about how discovery authenticates with your vCenter Server server, see the following information. This information includes guidance about certificate validation and SSL communication choices that you might need to make during vcenter credential configuration:

- [vCenter Server authentication](#)

8.1. ADDING VCENTER SOURCES

You can add sources from the initial Welcome page or from the Sources view.

Procedure

1. Click the option to add a new credential based on your location:

- From the Welcome page, click **Add Source**.
- From the Sources view, click **Add**.

The Add Source wizard opens.

2. On the Type page, select **vCenter Server** as the source type and click **Next**.
3. On the Credentials page, enter the following information.
 - a. In the **Name** field, enter a descriptive name.
 - b. In the **IP Address or Hostname** field, enter the IP address or host name of the vCenter Server for this source. Enter a different port if you do not want a scan for this source to run on the default port 443. For example, if the IP address of the vCenter Server is 192.0.2.15 and you want to change the port to 80, you would enter **192.0.2.15:80**.
 - c. In the **Credentials** list, select the credential that is required to access the vCenter Server for this source. If a required credential does not exist, click the **Add a credential** icon to open the Add Credential wizard.

- d. In the **Connection** list, select the SSL protocol to be used for a secure connection during a scan of this source. Select **Disable SSL** to disable secure communication during a scan of this source.
 - e. If you need to upgrade the SSL validation for the vCenter Server to check for a verified SSL certificate from a certificate authority, select the **Verify SSL Certificate** check box.
4. Click **Save** to save the source and then click **Close** to close the Add Source wizard.

8.2. ADDING VCENTER CREDENTIALS

You can add credentials from the Credentials view or from the Add Source wizard during the creation of a source.

Procedure

1. Click the option to add a new credential based on your location:
 - From the Credentials view, click **Add → VCenter Credential**.
 - From the Add Source wizard, click the **Add a credential** icon for the **Credentials** field.

The Add Credential wizard opens.

2. In the **Credential Name** field, enter a descriptive name.
3. Enter the username and password for a vCenter Server administrator.
4. Click **Save** to save the credential and close the Add Credential wizard.

8.3. ABOUT SOURCES AND CREDENTIALS

To run a scan, you must configure data for two basic structures, sources and credentials. The type of source that you are going to inspect during the scan determines the type of data that is required for both source and credential configuration.

A *source* contains a single asset or a set of multiple assets that are to be inspected during the scan. You can configure three types of sources:

- A network source: One or more physical machines, virtual machines, or containers. These assets can be expressed as host names, IP addresses, IP ranges, or subnets.
- A vcenter source: A vCenter Server systems management solution that is managing all or part of your IT infrastructure.
- A satellite source: A Satellite systems management solution that is managing all or part of your IT infrastructure.

When you are working with network sources, you determine how many individual assets you should group within a single source. The following list contains some of the factors that you should consider when you are adding sources:

- Whether assets are part of a development, testing, or production environment, and if demands on computing power and similar concerns are a consideration for those assets.

- Whether you want to scan a particular entity or group of entities more often because of internal business practices such as frequent changes to the installed software.

A *credential* contains data such as the user name and password or SSH key of a user with sufficient authority to run the scan on all or part of the assets that are contained in that source. As with sources, credentials are configured as the network, vcenter, or satellite type. Typically, a network source might require multiple network credentials because it is expected that many credentials would be needed to access all of the assets in a broad IP range. Conversely, a vcenter or satellite source would typically use a single vcenter or satellite credential, as applicable, to access a particular system management solution server.

You can add new sources from the Sources view and you can add new credentials from the Credentials view. You can also add new or select previously existing credentials during source creation. It is during source creation that you associate a credential directly with a source. Because sources and credentials must have matching types, any credential that you add during source creation shares the same type as the source. In addition, if you want to use an existing credential during source creation, the list of available credentials contains only credentials of the same type. For example, during network source creation, only network credentials are available for selection.

8.4. VCENTER SERVER AUTHENTICATION

For a vcenter scan, the connectivity and access to vCenter Server derives from basic authentication (user name and password) that is encrypted over HTTPS. By default, the vcenter scan runs with certificate validation and secure communication through the SSL (Secure Sockets Layer) protocol. During source creation, you can select from several different SSL and TLS (Transport Layer Security) protocols to use for the certificate validation and secure communication.

You might need to adjust the level of certificate validation to connect properly to the vCenter server during a scan. For example, your vCenter server might use a verified SSL certificate from a certificate authority. During source creation, you can upgrade SSL certificate validation to check for that certificate during a scan of that source. Conversely, your vCenter server might use self-signed certificates. During source creation, you can leave the SSL validation at the default so that scan of that source does not check for a certificate. This choice, to leave the option at the default for a self-signed certificate, could possibly avoid scan errors.

You might also need to disable SSL as the method of secure communication during the scan if the vCenter server is not configured to use SSL communication for web applications. For example, your vCenter server might be configured to communicate with web applications by using HTTP with port 80. If so, then during source creation you can disable SSL communication for scans of that source.

PART IV. RUNNING AND MANAGING SCANS

After you add sources and credentials for the parts of your IT infrastructure that you want to scan, you can create and run scans. When you create a scan, you can choose to scan a single source or combine multiple sources of any type. You can also choose whether to run a standard scan for products that are installed with default installation processes and locations or to run a deep scan if products might be installed with nonstandard processes or locations.

After a scan is created, you can run that scan multiple times. Each instance of that scan is saved as a scan job.

Learn more

To learn more about running a standard scan that does not use deep scanning for products, see the following information:

- [Running and managing standard scans](#)

To learn more about running a deep scan, a scan that can find products that might have been installed with a nonstandard process or in a nonstandard location, see the following information:

- [Running and managing deep scans](#)

CHAPTER 9. RUNNING AND MANAGING STANDARD SCANS

After you add sources and credentials for the parts of your IT infrastructure that you want to scan, you can begin running scans. In most situations, you can run a standard scan to find the environment and product data that is required to report on your Red Hat products.

Learn more

Run a standard scan to find products in standard locations. To learn more, see the following information:

- [Running standard scans](#)

When you begin running scans, there are several tasks that you can do to manage your scans. These tasks include updating the data for a scan by running a new scan job and managing active scans by pausing, resuming, and canceling them. To learn more, see the following information:

- [Running a new scan job](#)
- [Pausing, resuming, and canceling scans](#)

To learn more about how scans and scan jobs work, including how a scan job is processed by discovery and the states a scan job moves through during its life cycle, see the following information:

- [About scans and scan jobs](#)
- [Scan job processing](#)
- [Scan job life cycle](#)

9.1. RUNNING STANDARD SCANS

You can run a new scan from the Sources view. You can run a scan for a single source or select multiple sources to combine into a single scan. Each time that you use the Sources view to run a scan, you are prompted to save it as a new scan.

After you run a scan for the first time, the scan is saved to the Scans view. From that view, you can run that scan again to update its data. Each time that you run a scan from the Scans view, it is saved as a new scan job for that scan.

Prerequisites

- To run a scan, you must first add the sources that you want to scan and the credentials to access those sources.

Procedure

1. From the Sources view, select one or more sources. You can select sources of different types to combine them into a single scan.
2. Click the **Scan** button that is appropriate for the selected sources:
 - For a single source, click **Scan** on the row for that source. Selecting the check box for the source is optional.
 - If you selected multiple sources, click **Scan** in the toolbar. The Scan wizard opens.

3. In the **Name** field, enter a descriptive name for the scan.
4. If you want to change the default number of maximum concurrent scans, set a new value in the **Maximum concurrent scans** field. This value is the maximum number of machines or virtual machines that are scanned in parallel during a scan.
5. To use the default scanning process, allow the **Deep scan for these products** check boxes to remain in the default, cleared state.
6. To begin the scan process, click **Scan**.

Verification steps


When the scan process begins, a notification displays in the Sources view. The running scan also displays in the Scans view, with a message about the progress of the scan.

9.2. RUNNING A NEW SCAN JOB

After you name a scan and run it for the first time, it is added to the Scans view. You can then run a new instance of that scan, known as a scan job, to update the data that is gathered for that scan.

Procedure

1. From the Scans view, click the **Run Scan** icon in the scan details.



NOTE

In the scan details, if the most recent scan job did not complete successfully, this icon is labeled **Retry Scan**.

Verification steps

When the scan process begins, a notification displays with a message about the progress of the scan.

1. In the scan details, expand **Previous** to view all previous scan jobs.

9.3. PAUSING, RESUMING, AND CANCELING SCANS

As you begin running scans, you might need to stop a scan job that is currently running. There might be various business reasons that require you to do this, for example, the need to do an emergency fix due to an alert from your IT health monitoring system or the need to run a higher priority scan that consumes more CPU resources if a lower priority scan is currently running.

You can stop a scan job by either pausing it or canceling it. You can resume a paused scan job, but you cannot resume a canceled scan job.

Procedure

To pause a scan job that is running:

1. From the Scans view, find the scan that contains the scan job that you want to pause.
2. Click **Pause Scan**.



NOTE

If you have multiple scans running at the same time, it might take several moments after starting a scan for the **Pause Scan** icon to appear.

To resume a scan job that is paused:

1. From the Scans view, find the scan that contains the scan job that you want to resume.
2. Click **Resume Scan**.

To cancel a scan job that is running:

1. From the Scans view, find the scan that contains the scan job that you want to cancel.
2. Click **Cancel Scan**.

9.4. ABOUT SCANS AND SCAN JOBS

After you create sources and credentials, you can create scans. A *scan* is an object that groups sources into a unit that can be inspected, or scanned, in a reproducible way. Each time that you run a saved scan, that instance is saved as a *scan job*. The output of a scan job is a *report*, the collection of facts gathered for all IT resources that are contained in that source.

A scan includes at least one source and the credentials that were associated with that source at source creation time. When the scan job runs, it uses the provided credentials to contact the assets contained in the source and then it inspects the assets to gather facts about those assets for the report. You can add multiple sources to a single scan, including a combination of different types of sources into a single scan.

9.5. SCAN JOB PROCESSING

A scan job moves through two phases, or tasks, while it is being processed. These two tasks are the connection task and the inspection task.

9.5.1. Scan job connection and inspection tasks

The first task that runs during a scan job is a connection task. The *connection task* determines the ability to connect to the source and finds the number of systems that can be inspected for the defined source. The second task that runs is an inspection task. The *inspection task* is the task that gathers data from each of the reachable systems in the defined source to output the scan results into a report.

If the scan is configured so that it contains several sources, then when the scan job runs, these two tasks are created for each source. First, all of the connection tasks for all of the sources run to establish connections to the sources and find the systems that can be inspected. Then all of the inspection tasks for all of the sources run to inspect the contents of the reachable systems that are contained in the sources.

9.5.2. How these tasks are processed

When the scan job runs the connection task for a source, it attempts to connect to the network, the server for vCenter Server, or the server for Satellite. If the connection to vCenter Server or Satellite fails, then the connection task fails. For a network scan, if the network is not reachable or the credentials are

invalid, the connection task reports zero (0) successful systems. If only some of the systems for a network scan are reachable, the connection task reports success on the systems that are reachable, and the connection task does not fail.

You can view information about the status of the connection task in the Scans view. The row for a scan displays the connection task results as the number of successful system connections for the most recent scan job. You can also expand the previous scan jobs to see the connection task results for a previous scan job.

When the scan job runs the inspection task for a source, it checks the state of the connection task. If the connection task shows a failed state or if there are zero (0) successful connections, the scan job transitions to the failed state. However, if the connection task reports at least one successful connection, the inspection task continues. The results for the scan job then show success and failure data for each individual system. If the inspection task is not able to gather results from the successful systems, or if another unexpected error occurs during the inspection task, then the scan job transitions to the failed state.

If a scan contains multiple sources, each source has its own connection and inspection tasks. These tasks are processed independently from the tasks for the other sources. If any task for any of the sources fails, the scan job transitions to the failed state. The scan job transitions to the completed state only if all scan job tasks for all sources complete successfully.

If a scan job completes successfully, the data for that scan job is generated as a report. In the Scans view, you can download the report for each successful scan job.

9.6. SCAN JOB LIFE CYCLE

A *scan job*, or individual instance of a scan, moves through several states during its life cycle.

When you start a scan, a scan job is created and the scan job is in the *created* state. The scan job is then queued for processing and the scan job transitions to the *pending* state. Scan jobs run serially, in the order that they are started.

As the discovery server reaches a specific scan job in the queue, that scan job transitions from the *pending* state to the *running* state as the processing of that scan job begins. If the scan process completes successfully, the scan job transitions to the *completed* state and the scan job produces results that can be viewed in a report. If the scan process results in an error that prevents successful completion of the scan, the scan job halts and the scan job transitions to the *failed* state. An additional status message for the failed scan contains information to help determine the cause of the failure.

Other states for a scan job result from user action that is taken on the scan job. You can pause or cancel a scan job while it is pending or running. A scan job in the *paused* state can be resumed. A scan job in the *anceled* state cannot be resumed.

CHAPTER 10. RUNNING AND MANAGING DEEP SCANS

After you add sources and credentials for the parts of your IT infrastructure that you want to scan, you can begin running scans. In a few situations, running standard scans is not sufficient to find the environment and product data that is required to report on your Red Hat products.

By default, discovery searches for and fingerprints products by using known metadata that relates to those products. However, it is possible that you have installed these products with a process or in an installation location that makes the search and fingerprinting algorithms less effective. In that case, you need to use deep scanning to find those products.

Learn more

Run a deep scan to find products in nonstandard locations. To learn more, see the following information:

- [Running scans with deep scanning](#)

When you begin running scans, there are several tasks that you can do to manage your scans. These tasks include updating the data for a scan by running a new scan job and managing active scans by pausing, resuming, and canceling them. To learn more, see the following information:

- [Running a new scan job](#)
- [Pausing, resuming, and canceling scans](#)

To learn more about how scans and scan jobs work, including how a scan job is processed by discovery and the states a scan job moves through during its life cycle, see the following information:

- [About scans and scan jobs](#)
- [Scan job processing](#)
- [Scan job life cycle](#)

10.1. RUNNING SCANS WITH DEEP SCANNING

You can run a new scan from the Sources view. You can run a scan for a single source or select multiple sources to combine into a single scan. As part of the scan configuration, you might choose to use the deep scanning process to search for products in nonstandard locations.

The deep scanning process uses the **find** command, so the search process could be CPU resource intensive for the systems that are being scanned. Therefore, you should use discretion when selecting a deep scan for systems that require continuous availability, such as production systems.

After you run a scan for the first time, the scan is saved to the Scans view. From that view, you can run the scan again to update its data.

Prerequisites

- To run a scan, you must first add the sources that you want to scan and the credentials to access those sources.

Procedure

1. From the Sources view, select one or more sources. You can select sources of different types to combine them into a single scan.

2. Click the **Scan** button that is appropriate for the selected sources:
 - For a single source, click **Scan** on the row for that source. Selecting the check box for the source is optional.
 - If you selected multiple sources, click **Scan** in the toolbar. The Scan wizard opens.
3. In the **Name** field, enter a descriptive name for the scan.
4. If you want to change the default number of maximum concurrent scans, set a new value in the **Maximum concurrent scans** field. This value is the maximum number of machines or virtual machines that are scanned in parallel during a scan.
5. To use the deep scanning process on one or more products, supply the following information:
 - Select the applicable **Deep scan for these products** check boxes.
 - Optionally, enter the directories that you want discovery to scan. The default directories that are used in a deep scan are the `/`, `/opt`, `/app`, `/home`, and `/usr` directories.
6. To begin the scan process, click **Scan**.

Verification steps

When the scan process begins, a notification displays in the Sources view. The running scan also displays in the Scans view, with a message about the progress of the scan.

10.2. RUNNING A NEW SCAN JOB

After you name a scan and run it for the first time, it is added to the Scans view. You can then run a new instance of that scan, known as a scan job, to update the data that is gathered for that scan.

Procedure

1. From the Scans view, click the **Run Scan** icon in the scan details.



NOTE

In the scan details, if the most recent scan job did not complete successfully, this icon is labeled **Retry Scan**.

Verification steps

When the scan process begins, a notification displays with a message about the progress of the scan.

1. In the scan details, expand **Previous** to view all previous scan jobs.

10.3. PAUSING, RESUMING, AND CANCELING SCANS

As you begin running scans, you might need to stop a scan job that is currently running. There might be various business reasons that require you to do this, for example, the need to do an emergency fix due to an alert from your IT health monitoring system or the need to run a higher priority scan that consumes more CPU resources if a lower priority scan is currently running.

You can stop a scan job by either pausing it or canceling it. You can resume a paused scan job, but you cannot resume a canceled scan job.

Procedure

To pause a scan job that is running:

1. From the Scans view, find the scan that contains the scan job that you want to pause.
2. Click **Pause Scan**.



NOTE

If you have multiple scans running at the same time, it might take several moments after starting a scan for the **Pause Scan** icon to appear.

To resume a scan job that is paused:

1. From the Scans view, find the scan that contains the scan job that you want to resume.
2. Click **Resume Scan**.

To cancel a scan job that is running:

1. From the Scans view, find the scan that contains the scan job that you want to cancel.
2. Click **Cancel Scan**.

10.4. ABOUT SCANS AND SCAN JOBS

After you create sources and credentials, you can create scans. A *scan* is an object that groups sources into a unit that can be inspected, or scanned, in a reproducible way. Each time that you run a saved scan, that instance is saved as a *scan job*. The output of a scan job is a *report*, the collection of facts gathered for all IT resources that are contained in that source.

A scan includes at least one source and the credentials that were associated with that source at source creation time. When the scan job runs, it uses the provided credentials to contact the assets contained in the source and then it inspects the assets to gather facts about those assets for the report. You can add multiple sources to a single scan, including a combination of different types of sources into a single scan.

10.5. SCAN JOB PROCESSING

A scan job moves through two phases, or tasks, while it is being processed. These two tasks are the connection task and the inspection task.

10.5.1. Scan job connection and inspection tasks

The first task that runs during a scan job is a connection task. The *connection task* determines the ability to connect to the source and finds the number of systems that can be inspected for the defined source. The second task that runs is an inspection task. The *inspection task* is the task that gathers data from each of the reachable systems in the defined source to output the scan results into a report.

If the scan is configured so that it contains several sources, then when the scan job runs, these two tasks are created for each source. First, all of the connection tasks for all of the sources run to establish connections to the sources and find the systems that can be inspected. Then all of the inspection tasks for all of the sources run to inspect the contents of the reachable systems that are contained in the sources.

10.5.2. How these tasks are processed

When the scan job runs the connection task for a source, it attempts to connect to the network, the server for vCenter Server, or the server for Satellite. If the connection to vCenter Server or Satellite fails, then the connection task fails. For a network scan, if the network is not reachable or the credentials are invalid, the connection task reports zero (0) successful systems. If only some of the systems for a network scan are reachable, the connection task reports success on the systems that are reachable, and the connection task does not fail.

You can view information about the status of the connection task in the Scans view. The row for a scan displays the connection task results as the number of successful system connections for the most recent scan job. You can also expand the previous scan jobs to see the connection task results for a previous scan job.

When the scan job runs the inspection task for a source, it checks the state of the connection task. If the connection task shows a failed state or if there are zero (0) successful connections, the scan job transitions to the failed state. However, if the connection task reports at least one successful connection, the inspection task continues. The results for the scan job then show success and failure data for each individual system. If the inspection task is not able to gather results from the successful systems, or if another unexpected error occurs during the inspection task, then the scan job transitions to the failed state.

If a scan contains multiple sources, each source has its own connection and inspection tasks. These tasks are processed independently from the tasks for the other sources. If any task for any of the sources fails, the scan job transitions to the failed state. The scan job transitions to the completed state only if all scan job tasks for all sources complete successfully.

If a scan job completes successfully, the data for that scan job is generated as a report. In the Scans view, you can download the report for each successful scan job.

10.6. SCAN JOB LIFE CYCLE

A *scan job*, or individual instance of a scan, moves through several states during its life cycle.

When you start a scan, a scan job is created and the scan job is in the *created* state. The scan job is then queued for processing and the scan job transitions to the *pending* state. Scan jobs run serially, in the order that they are started.

As the discovery server reaches a specific scan job in the queue, that scan job transitions from the *pending* state to the *running* state as the processing of that scan job begins. If the scan process completes successfully, the scan job transitions to the *completed* state and the scan job produces results that can be viewed in a report. If the scan process results in an error that prevents successful completion of the scan, the scan job halts and the scan job transitions to the *failed* state. An additional status message for the failed scan contains information to help determine the cause of the failure.

Other states for a scan job result from user action that is taken on the scan job. You can pause or cancel a scan job while it is pending or running. A scan job in the *paused* state can be resumed. A scan job in the *canceled* state cannot be resumed.

PART V. MERGING AND DOWNLOADING REPORTS

After you run a scan, you can download the reports for that scan to view the data that was gathered and processed during that scan.

If you choose to divide the scanning tasks for your IT infrastructure into several discrete scans, you can merge the output of those scans into a single report. Merged reports are also downloaded at the time of merging.

Learn more

To learn more about downloading reports, see the following information:

- [Downloading reports](#)

To learn more about merging two or more reports and downloading the result of the merge, see the following information:

- [Merging reports](#)

CHAPTER 11. DOWNLOADING REPORTS

After you run a scan, you can download the reports for that scan to view the data that was gathered and processed during that scan.

Reports for a scan are available in two formats, a comma-separated variable (CSV) format and a JavaScript Object Notation (JSON) format. They are also available in two content types, raw output from the scan as a details report and processed content as a deployments report.

Learn more

To learn more about merging and downloading reports, see the following information:

- [Downloading reports](#)

To learn more about how reports are created, see the following information. This information includes a chronology of the processes that change the raw facts of a details report into fingerprint data, and then change fingerprint data into the deduplicated and merged data of a deployments report. This information also includes a partial fingerprint example to show the types of data that are used to create a discovery report.

- [How reports are created](#)
- [A fingerprint example](#)

11.1. DOWNLOADING REPORTS

From the Scans view, you can select one or more reports and download them to view the report data.

Prerequisites

If you want to download a report for a scan, the most recent scan job for that scan must have completed successfully.

Procedure

1. From the Scans view, navigate to the row of the scan for which you want to download the report.
2. Click **Download** for that scan.

Verification steps

The downloaded report is saved to the downloads location for your browser as a **.tar.gz** file, for example, **report_id_224_20190702_173309.tar.gz**. The file name format is **report_id_ID_DATE_TIME.tar.gz**, where **ID** is the unique report ID assigned by the server, **DATE** is the date in *yyyymmdd* format, and **TIME** is the time in the *hhmmss* format, based on the 24-hour system. The date and time data is determined by the interaction of the browser that is running the client with the server APIs.

To view the report, uncompress the **.tar.gz** file into a **report_id_ID** directory. *The uncompressed report bundle includes four report files, two details reports in CSV and JSON formats, and two deployments reports in CSV and JSON formats.*

While you can view and use the output of these reports for your own internal processes, their format is designed for use by the Red Hat Subscription Education and Awareness Program (SEAP) team during customer engagements and for other Red Hat internal processes.

11.2. HOW REPORTS ARE CREATED

The scan process is used to discover the systems in your IT infrastructure, to inspect and gather information about the nature and contents of those systems, and to create a report from the information that it gathers during the inspection of each system.

A system is any entity that can be interrogated by the inspection tasks through an SSH connection, vCenter Server data, or the Satellite Server API. Therefore, a system can be a machine, such as a physical or virtual machine, and it can also be a different type of entity, such as a container.

11.2.1. Facts and fingerprints

During a scan, a collection of facts is gathered for each system that is contained in each source. A fact is a single piece of data about a system, such as the version of the operating system, the number of CPU cores, or a consumed entitlement for a Red Hat product.

Facts are processed to create a summarized set of data for each system, data that is known as a fingerprint. A fingerprint is the set of facts that identifies a unique system and its characteristics, including the architecture, operating system, the different products that are installed on that system and their versions, the entitlements that are in use on that system, and so on.

Fingerprinting data is generated when you run a scan job, but the data is used to create only one type of report. When you request a details report, you receive the raw facts for that scan without any fingerprinting. When you request a deployments report, you receive the fingerprinting data that includes the results from the deduplication, merging, and post-processing processes. These processes include identifying installed products and versions from the raw facts, finding consumed entitlements, finding and merging duplicate instances of products from different sources, and finding products installed in nondefault locations, among other steps.

11.2.2. System deduplication and merging

A single system can be found in multiple sources during a scan. For example, a virtual machine on vCenter Server could be running a Red Hat Enterprise Linux operating system installation that is also managed by Satellite. If you construct a scan that contains each type of source, vcenter, satellite, and network, that single system is reported by all three sources during the scan.

To resolve this issue and build an accurate fingerprint, discovery feeds unprocessed system facts from the scan into a fingerprint engine. The fingerprint engine matches and merges data for systems that are found in more than one source by using the deduplication and merge processes.

11.2.2.1. System deduplication

The system deduplication process uses specific facts about a system to identify duplicate systems. The process moves through several phases, using these facts to combine duplicate systems in successively broader sets of data:

- All systems from network sources are combined into a single network system set. Systems are considered to be duplicates if they have the same value for the **subscription_manager_id** or **bios_uuid** facts.
- All systems from vcenter sources are combined into a single vcenter system set. Systems are considered to be duplicates if they have the same value for the **vm_uuid** fact.
- All systems from satellite sources are combined into a single satellite system set. Systems are considered to be duplicates if they have the same value for the **subscription_manager_id** fact.

- The network system set is merged with the satellite system set to form a single network-satellite system set. Systems are considered to be duplicates if they have the same value for the **subscription_manager** fact or matching MAC address values in the **mac_addresses** fact.
- The network-satellite system set is merged with the vcenter system set to form the complete system set. Systems are considered to be duplicates if they have matching MAC address values in the **mac_addresses** fact or if the vcenter value for the **vm_uuid** fact matches the network value for the **bios_uuid** fact.

11.2.2.2. System merging

After the deduplication process determines that two systems are duplicates, the next step is to perform a merge of those two systems. The merged system has a union of system facts from each source. When a fact that appears in two systems is merged, the merge process uses the following order of precedence to merge that fact, from highest to lowest:

1. network source fact
2. satellite source fact
3. vcenter source fact

A system fingerprint contains a **metadata** dictionary that captures the original source of each fact for that system.

11.2.3. System post-processing

After deduplication and merging are complete, there is a post-processing phase that creates derived system facts. A derived system fact is a fact that is generated from the evaluation of more than one system fact. The majority of derived system facts are related to product identification data, such as the presence of a specific product and its version.

The following example shows how the derived system fact **system_creation_date** is created.

The **system_creation_date** fact is a derived system fact that contains the real system creation time. The value for this fact is determined by the evaluation of the following facts. The value for each fact is examined in the following order of precedence, with the order of precedence determined by the accuracy of the match to the real system creation time. The highest non-empty value is used to determine the value of the **system_creation_date** fact.

1. **date_machine_id**
2. **registration_time**
3. **date_anaconda_log**
4. **date_filesystem_create**
5. **date_yum_history**

11.2.4. Report creation

After the processing of the report data is complete, the report creation process builds two reports in two different formats, JavaScript Object Notation (JSON) and comma-separated variable (CSV). The details report for each format contains the raw facts with no processing, and the deployments report for each

format contains the output after the raw facts have passed through the fingerprinting, deduplication, merge, and post-processing processes.

While you can view and use the output of these reports for your own internal processes, their format is designed for use by the Red Hat Subscription Education and Awareness Program (SEAP) team during customer engagements and for other Red Hat internal processes.

11.2.5. A fingerprint example

A fingerprint is composed of a set of facts about a single system in addition to facts about products, entitlements, sources, and metadata on that system. The following example shows fingerprint data. A fingerprint for a single system, even with very few Red Hat products installed on it, can be many lines. Therefore, only a partial fingerprint is used in this example.

Example

```
{
  "os_release": "Red Hat Enterprise Linux Atomic Host 7.4",
  "cpu_count": 4,
  "products": [
    {
      "name": "JBoss EAP",
      "version": null,
      "presence": "absent",
      "metadata": {
        "source_id": 5,
        "source_name": "S62Source",
        "source_type": "satellite",
        "raw_fact_key": null
      }
    }
  ],
  "entitlements": [
    {
      "name": "Satellite Tools 6.3",
      "entitlement_id": 54,
      "metadata": {
        "source_id": 5,
        "source_name": "S62Source",
        "source_type": "satellite",
        "raw_fact_key": "entitlements"
      }
    }
  ],
  "metadata": {
    "os_release": {
      "source_id": 5,
      "source_name": "S62Source",
      "source_type": "satellite",
      "raw_fact_key": "os_release"
    },
    "cpu_count": {
      "source_id": 4,
      "source_name": "NetworkSource",
      "source_type": "network",
      "raw_fact_key": "os_release"
    }
  }
}
```

```

    }
  },
  "sources": [
    {
      "id": 4,
      "source_type": "network",
      "name": "NetworkSource"
    },
    {
      "id": 5,
      "source_type": "satellite",
      "name": "S62Source"
    }
  ]
}

```

The first several lines of a fingerprint show facts about the system, including facts about the operating system and CPUs. In this example, the **os_release** fact describes the installed operating system and release as **Red Hat Enterprise Linux Atomic Host 7.4**.

Next, the fingerprint lists the installed products in the **products** section. A product has a name, version, presence, and metadata field. Because the presence field shows **absent** as the value for JBoss EAP, the system in this example does not have JBoss EAP installed.

The fingerprint also lists the consumed entitlements for that system in the **entitlements** section. Each entitlement in the list has a name, ID, and metadata that describes the original source of that fact. In the example fingerprint, the system has the **Satellite Tools 6.3** entitlement.

In addition to the metadata fields that are in the **products** and **entitlements** sections, the fingerprint contains a **metadata** section that is used for system fact metadata. For each system fact, there is a corresponding entry in the **metadata** section of the fingerprint that identifies the original source of that system fact. In the example, the **os_release** fact was found in Satellite Server, during the scan of the satellite source.

Lastly, the fingerprint lists the sources that contain this system in the **sources** section. A system can be contained in more than one source. For example, for a scan that includes both a network source and a satellite source, a single system can be found in both parts of the scan.

CHAPTER 12. MERGING REPORTS

If you choose to divide the scanning tasks for your IT infrastructure into several discrete scans, you can merge the output of those scans into a single report.

Merged reports are also downloaded at the time of merging so that you can view the data that was gathered and processed during that scan. Reports for a scan are available in two formats, a comma-separated variable (CSV) format and a JavaScript Object Notation (JSON) format. They are also available in two content types, raw output from the scan as a details report and processed content as a deployments report.

Learn more

To learn more about merging reports and downloading the merged reports, see the following information:

- [Merging reports](#)

To learn more about how reports are created, see the following information. This information includes a chronology of the processes that change the raw facts of a details report into fingerprint data, and then change fingerprint data into the deduplicated and merged data of a deployments report. This information also includes a partial fingerprint example to show the types of data that are used to create a discovery report.

- [How reports are created](#)
- [A fingerprint example](#)

12.1. MERGING REPORTS

From the Scans view, you can select two or more reports and merge them into a single report.

There might be several reasons that you would choose to merge reports, including the following examples:

- You might have a large IT infrastructure with many different administrators, each of whom can access and scan only a part of that infrastructure.
- You might run multiple separate scans to limit the demand on CPU resources of the machines that are being scanned, especially in situations where deep scanning is required.
- You might want to run scans on a single type of source, isolating your network, satellite, and vcenter data into separate reports for your own internal purposes and then merge these reports later.

For these and similar reasons, it is possible that you will run multiple scans to provide complete scan coverage of your entire IT infrastructure. Merging reports enables you to combine the data from multiple scans into a single report.

Prerequisites

To merge reports, you must select at least two scans for which the most recent scan jobs have completed successfully.

Procedure

1. From the Scans view, select the check box for one or more scans.
2. Click **Merge reports**. A confirmation dialog box shows the selected scans.

3. Click **Merge** to merge the scans into a single report and download the merged reports.

Verification steps

The merged report is saved to the downloads location for your browser as a **.tar.gz** file, for example, **report_id_110_20190529_095005.tar.gz**. The file name format is **report_id_ID_DATE_TIME.tar.gz**, where **ID** is the unique report ID assigned by the server, **DATE** is the date in `yyyymmdd` format, and **TIME** is the time in the `hhmmss` format, based on the 24-hour system. The date and time data is determined by the interaction of the browser that is running the client with the server APIs.

To view the reports, uncompress the **.tar.gz** file into a **report_id_ID** directory. The uncompressed report bundle includes four report files, two details reports in CSV and JSON formats, and two deployments reports in CSV and JSON formats.

While you can view and use the output of these reports for your own internal processes, their format is designed for use by the Red Hat Subscription Education and Awareness Program (SEAP) team during customer engagements and for other Red Hat internal processes.

12.2. HOW REPORTS ARE CREATED

The scan process is used to discover the systems in your IT infrastructure, to inspect and gather information about the nature and contents of those systems, and to create a report from the information that it gathers during the inspection of each system.

A system is any entity that can be interrogated by the inspection tasks through an SSH connection, vCenter Server data, or the Satellite Server API. Therefore, a system can be a machine, such as a physical or virtual machine, and it can also be a different type of entity, such as a container.

12.2.1. Facts and fingerprints

During a scan, a collection of facts is gathered for each system that is contained in each source. A fact is a single piece of data about a system, such as the version of the operating system, the number of CPU cores, or a consumed entitlement for a Red Hat product.

Facts are processed to create a summarized set of data for each system, data that is known as a fingerprint. A fingerprint is the set of facts that identifies a unique system and its characteristics, including the architecture, operating system, the different products that are installed on that system and their versions, the entitlements that are in use on that system, and so on.

Fingerprinting data is generated when you run a scan job, but the data is used to create only one type of report. When you request a details report, you receive the raw facts for that scan without any fingerprinting. When you request a deployments report, you receive the fingerprinting data that includes the results from the deduplication, merging, and post-processing processes. These processes include identifying installed products and versions from the raw facts, finding consumed entitlements, finding and merging duplicate instances of products from different sources, and finding products installed in nondefault locations, among other steps.

12.2.2. System deduplication and merging

A single system can be found in multiple sources during a scan. For example, a virtual machine on vCenter Server could be running a Red Hat Enterprise Linux operating system installation that is also managed by Satellite. If you construct a scan that contains each type of source, vcenter, satellite, and network, that single system is reported by all three sources during the scan.

To resolve this issue and build an accurate fingerprint, discovery feeds unprocessed system facts from the scan into a fingerprint engine. The fingerprint engine matches and merges data for systems that are found in more than one source by using the deduplication and merge processes.

12.2.2.1. System deduplication

The system deduplication process uses specific facts about a system to identify duplicate systems. The process moves through several phases, using these facts to combine duplicate systems in successively broader sets of data:

- All systems from network sources are combined into a single network system set. Systems are considered to be duplicates if they have the same value for the **subscription_manager_id** or **bios_uuid** facts.
- All systems from vcenter sources are combined into a single vcenter system set. Systems are considered to be duplicates if they have the same value for the **vm_uuid** fact.
- All systems from satellite sources are combined into a single satellite system set. Systems are considered to be duplicates if they have the same value for the **subscription_manager_id** fact.
- The network system set is merged with the satellite system set to form a single network-satellite system set. Systems are considered to be duplicates if they have the same value for the **subscription_manager** fact or matching MAC address values in the **mac_addresses** fact.
- The network-satellite system set is merged with the vcenter system set to form the complete system set. Systems are considered to be duplicates if they have matching MAC address values in the **mac_addresses** fact or if the vcenter value for the **vm_uuid** fact matches the network value for the **bios_uuid** fact.

12.2.2.2. System merging

After the deduplication process determines that two systems are duplicates, the next step is to perform a merge of those two systems. The merged system has a union of system facts from each source. When a fact that appears in two systems is merged, the merge process uses the following order of precedence to merge that fact, from highest to lowest:

1. network source fact
2. satellite source fact
3. vcenter source fact

A system fingerprint contains a **metadata** dictionary that captures the original source of each fact for that system.

12.2.3. System post-processing

After deduplication and merging are complete, there is a post-processing phase that creates derived system facts. A derived system fact is a fact that is generated from the evaluation of more than one system fact. The majority of derived system facts are related to product identification data, such as the presence of a specific product and its version.

The following example shows how the derived system fact **system_creation_date** is created.

The **system_creation_date** fact is a derived system fact that contains the real system creation time. The value for this fact is determined by the evaluation of the following facts. The value for each fact is

examined in the following order of precedence, with the order of precedence determined by the accuracy of the match to the real system creation time. The highest non-empty value is used to determine the value of the **system_creation_date** fact.

1. **date_machine_id**
2. **registration_time**
3. **date_anaconda_log**
4. **date_filesystem_create**
5. **date_yum_history**

12.2.4. Report creation

After the processing of the report data is complete, the report creation process builds two reports in two different formats, JavaScript Object Notation (JSON) and comma-separated variable (CSV). The details report for each format contains the raw facts with no processing, and the deployments report for each format contains the output after the raw facts have passed through the fingerprinting, deduplication, merge, and post-processing processes.

While you can view and use the output of these reports for your own internal processes, their format is designed for use by the Red Hat Subscription Education and Awareness Program (SEAP) team during customer engagements and for other Red Hat internal processes.

12.2.5. A fingerprint example

A fingerprint is composed of a set of facts about a single system in addition to facts about products, entitlements, sources, and metadata on that system. The following example shows fingerprint data. A fingerprint for a single system, even with very few Red Hat products installed on it, can be many lines. Therefore, only a partial fingerprint is used in this example.

Example

```
{
  "os_release": "Red Hat Enterprise Linux Atomic Host 7.4",
  "cpu_count": 4,
  "products": [
    {
      "name": "JBoss EAP",
      "version": null,
      "presence": "absent",
      "metadata": {
        "source_id": 5,
        "source_name": "S62Source",
        "source_type": "satellite",
        "raw_fact_key": null
      }
    }
  ],
  "entitlements": [
    {
      "name": "Satellite Tools 6.3",
      "entitlement_id": 54,
      "metadata": {
```

```

        "source_id": 5,
        "source_name": "S62Source",
        "source_type": "satellite",
        "raw_fact_key": "entitlements"
    }
}
],
"metadata": {
    "os_release": {
        "source_id": 5,
        "source_name": "S62Source",
        "source_type": "satellite",
        "raw_fact_key": "os_release"
    },
    "cpu_count": {
        "source_id": 4,
        "source_name": "NetworkSource",
        "source_type": "network",
        "raw_fact_key": "os_release"
    }
},
"sources": [
    {
        "id": 4,
        "source_type": "network",
        "name": "NetworkSource"
    },
    {
        "id": 5,
        "source_type": "satellite",
        "name": "S62Source"
    }
]
}

```

The first several lines of a fingerprint show facts about the system, including facts about the operating system and CPUs. In this example, the **os_release** fact describes the installed operating system and release as **Red Hat Enterprise Linux Atomic Host 7.4**.

Next, the fingerprint lists the installed products in the **products** section. A product has a name, version, presence, and metadata field. Because the presence field shows **absent** as the value for JBoss EAP, the system in this example does not have JBoss EAP installed.

The fingerprint also lists the consumed entitlements for that system in the **entitlements** section. Each entitlement in the list has a name, ID, and metadata that describes the original source of that fact. In the example fingerprint, the system has the **Satellite Tools 6.3** entitlement.

In addition to the metadata fields that are in the **products** and **entitlements** sections, the fingerprint contains a **metadata** section that is used for system fact metadata. For each system fact, there is a corresponding entry in the **metadata** section of the fingerprint that identifies the original source of that system fact. In the example, the **os_release** fact was found in Satellite Server, during the scan of the satellite source.

Lastly, the fingerprint lists the sources that contain this system in the **sources** section. A system can be contained in more than one source. For example, for a scan that includes both a network source and a satellite source, a single system can be found in both parts of the scan.

