Red Hat Subscription Management 1

Installing and Configuring Discovery

Installing Discovery

Last Updated: 2020-06-22

Red Hat Subscription Management 1 Installing and Configuring Discovery

Installing Discovery

Legal Notice

Abstract

Table of Contents

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.

Although you can install and use discovery independently and then generate and view report data, the discovery documentation does not provide any information to help you interpret report results. In addition, although Red Hat Support can provide some basic assistance related to installation and usage of the product discovery tool, the support team does not provide any assistance to help you understand the reports. The graduated pilot programs are designed for subsets of Red Hat customers who qualify based on their Red Hat product profile and other factors. These pilot programs are helping to refine the ingestion, rendering, analysis, and usage of report information as part of a management solution to help you understand your product inventory across your environments, whether on-premise, cloud, or containers.

PART II. INSTALLING PREREQUISITES FOR DISCOVERY

Before you begin the installation process, review the information about discovery prerequisites. Then complete any installation or configuration tasks for the prerequisites.

Procedure

Install the following requirements for hardware, software, and the environment in which you are going to install and use discovery.

CHAPTER 4. HARDWARE PREREQUISITES

The system on which you are going to install discovery must meet or exceed the following hardware requirements:

- **CPU:** 2 core minimum, with a recommended 4 cores
- **Disk Storage:** 30 GB
- **RAM:** 1 GB minimum, with a recommended 2 GB

CHAPTER 5. SOFTWARE PREREQUISITES

The system on which you are going to install discovery must have the following software requirements installed:

- **Operating system:** The latest release of the following operating system versions:
    - Red Hat Enterprise Linux 8

In addition to these software requirements, discovery has dependencies on other software. However, in some cases these dependencies are installed for you during installation, or instructions for their installation is integrated as part of the discovery installation process. The process to install dependencies can vary according several factors, including your operating system and whether you are doing an offline or online installation. Therefore, install the dependencies only as instructed by the discovery installation steps.

The discovery dependencies include the following software:

- Ansible 2.4 or later, depending on the requirements of your operating system.
- Podman container software.
- The PostgreSQL database.
- Python, 3.4 or 3.6, depending on the requirements of your operating system.

CHAPTER 6. OTHER ENVIRONMENT PREREQUISITES

The environment in which you are going to use discovery must meet the following requirements. Some of these requirements affect the systems on which you are going to run discovery. Others affect the systems in your IT infrastructure that you are going to scan with discovery.

On the machine where discovery is installed:

- If you install the discovery server and CLI packages on different machines, discovery-tools must be installed on both machines to run each installation.
- If you install the discovery server and CLI packages on different machines, the CLI machine must be able to connect to the server machine.
- The discovery server must have access to the IT infrastructure assets that are to be scanned.

On the systems where discovery runs scans:

- Any network sources that are targeted for scanning must be running the Secure Shell (SSH) protocol.
- A user account that is used as a credential for a scan requires the **bash** shell. The shell cannot be the **/sbin/nologin** shell or the **/bin/false** shell.
- A user account that is used as a credential for a network scan must have adequate permissions to run commands and read certain files on those systems. For example, some commands that run during a scan require privilege elevation to gather the complete set of facts for the scan. The *Using product discovery* guide has additional information about the creation of credentials for network scans and the privileges that must be associated with those credentials to enable a more complete scan of network assets.
- A user account that is used as a credential for a network scan where authentication is done with an SSH key must have a copy of the private key on the discovery server. The private key must be stored in the **~/discovery/server/volumes/sshkeys** directory, the default location for this directory at the time of server installation.

PART III. INSTALLING DISCOVERY WITH THE ONLINE INSTALLATION PROCESS

The online installation process runs an installer that uses Ansible to install the command line interface tool, the server image, and the database image.

When you run the online installation process, the server and CLI packages are installed with the default options. However, you can change some of the defaults used by the installation process by setting options that are defined in Options for the discovery-tools install command.

Prerequisites

- Ensure that all prerequisites are installed and configured.
- To use the online installation process, the machine on which you are downloading the installer and installing discovery must be connected to the Internet.

Procedure

To install discovery with the online installation process, you do the following tasks:

1. Enable the Ansible and discovery-tools repositories and install discovery-tools. For more information, see Enabling and installing Ansible and discovery-tools.
2. Run the discovery installation with discovery-tools. For more informa tion, see Installing with the online installation process.

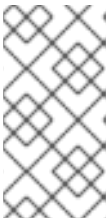

NOTE

This installation process installs discovery with default settings for all of the options. However, other settings are available for you to customize your installation. For more information about these settings, see Options for the discovery-tools install command.

CHAPTER 7. ENABLING AND INSTALLING ANSIBLE AND DISCOVERY-TOOLS

The discovery-tools package includes tools that you use to install and maintain discovery.

The discovery-tools installation process integrates with an Ansible playbook to complete the installation. Therefore, Ansible version 2.4 or later (depending on the requirements of your operating system) is a depenency for discovery-tools. As part of the installation process, discovery-tools installs Ansible, but you must manually enable the required repositories for Ansible.

Procedure

1. Enable the RHEL Ansible Engine repositories. For more information about this step, see the instructions for enabling repositories for the limited support version of Ansible in the "How do I Download and Install Red Hat Ansible Engine? Customer Portal article: https://access.redhat.com/articles/3174981 For more detailed information about Ansible, see https://docs.ansible.com/#coreversionselect.

   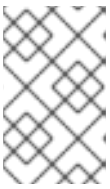

   NOTE

   The procedure at this link includes a step to install Ansible. Because discovery-tools installs Ansible for you, running the command to install Ansible is not required.

2. Register your system to Red Hat Subscription Manager:

   ```
   # subscription-manager register
   ```

3. Use the following command to help you find the discovery subscription and then note the pool ID of the subscription:

   ```
   # subscription-manager list --available
   ```

4. Attach the subscription, where ***pool_ID*** is the pool ID for the discovery subscription.

   ```
   # subscription-manager attach --pool=pool_ID
   ```

5. Enable the discovery repository:

   ```
   # subscription-manager repos --enable discovery-0-for-rhel-8-x86_64-rpms
   ```

6. Install discovery-tools:

   ```
   # yum install discovery-tools
   ```

CHAPTER 8. INSTALLING WITH THE ONLINE INSTALLATION PROCESS

During the discovery installation process, you enter commands to install the server and the command line interface (CLI). The simplest method for installing discovery is to run the online installation with all of the default options. With this method, the install process prompts you to do the following actions:

- Enter your user name and password for the Red Hat Container Catalog.
- Set a password for the discovery server administrator.
- Set a password for the database user.

The following information includes commands for the default installation process. However, you can change some of the defaults used by the installation process by setting options that are defined in Options for the discovery-tools install command.

Procedure

To install the discovery server and CLI packages, use the following steps:

8.1. INSTALLING THE SERVER

Install the server. The default server installation process prompts you to enter your user name and password for the Red Hat Container Catalog (the registry.redhat.io website). It also prompts you to supply a password for the discovery server administrator and a password for the PostgreSQL database user.

IMPORTANT

As a best practice, note the discovery server administrator and the PostgreSQL database user passwords in the password management system that is used by your organization. The product discovery tool does not offer a method to recover these passwords.

In addition, if you later use discovery-tools to upgrade discovery, you must use the same database user name and password during the upgrade. The failure to use the same database credentials could result in data loss.

1. Install the server by entering the following command:

```
# dsc-tools server install
```

2. At the prompt, enter your user name for the Red Hat Container Catalog, also known as the registry.redhat.io image registry website.
3. At the prompt, enter your password for the Red Hat Container Catalog.
4. At the prompt, set a password for the discovery server administrator.
5. At the prompt, set a password for the PostgreSQL database user.

Verification steps

1. For a successful installation of the server package, the output appears similar to the following example:

```
Installation of the server was successful.
```

8.2. INSTALLING THE COMMAND LINE INTERFACE

Install the command line interface.

The recommended method is to include the options to set the port and IP address of the discovery server. These options must be set so that the discovery command line interface can communicate with the server.

1. Install the command line interface by entering the following command, where **server_port** is the port that the discovery server is using for communication and ***server_host*** is the IP address of the server:

```
# dsc-tools cli install --server-port=server_port --server-host=server_host
```

Verification steps

1. For a successful installation command line interface packages the output appears similar to the following example:

```
Installation of the CLI was successful.
```

CHAPTER 9. OPTIONS FOR THE DISCOVERY-TOOLS INSTALL COMMAND

The installation procedures in this guide install the discovery server and CLI packages with all of the default options. However, the **install** subcommand that installs the discovery server and command line interface packages includes options to customize the installation process.

The following information lists the **install** subcommand options, along with applicable usage information and default values, for the server and CLI packages.

9.1. SERVER INSTALLATION OPTIONS

--version=*version.patch.release*

Enables the installation of a specific discovery server version. Contains the semantic versioning format (version.release.patch, such as 0.9.0) of the discovery server that you want to install. This option has a default of the latest version.

--home-dir=*server_home_dir*

Sets the fully qualified path to the installation directory for the discovery server. Defaults to ~/**discovery**/.

--port=*server_port*

Sets the port number for the discovery server. Defaults to **9443**.

--open-port=true|false

Determines whether to open the port in the firewall during the installation. This option enables communication between the discovery server and any remote clients over the port defined in the **port** option. Contains a true or false value. Defaults to true. Supply **false** to install without opening the server port in the firewall. The installation script must run with elevated privileges to open the server port.

--registry-user=*registry_website_username*

Specifies your user name for the Red Hat Container Catalog, also known as the **registry.redhat.io** image registry website. You are prompted for this value during server installation.

--registry-password=*registry_website_password*

Specifies your password for the Red Hat Container Catalog, also known as the **registry.redhat.io** image registry website. You are prompted for this value during server installation.

--db-user=*database_username*

Sets the database user name for PostgreSQL. Defaults to **postgres**.

IMPORTANT

If you later use discovery-tools to upgrade discovery, you must use the same database user name and password during the upgrade. The failure to use the same database credentials could result in data loss.

--db-password=*database_password*

Sets the database user password for PostgreSQL. This option has no default value. If omitted, discovery-tools prompts for a password.

IMPORTANT

If you later use discovery-tools to upgrade discovery, you must use the same database user name and password during the upgrade. The failure to use the same database credentials could result in data loss.

--username=*server_username*

Sets the discovery server administrator user name. Defaults to **admin**.

--password=*server_password*

Sets the discovery server administrator password. This option has no default value. If omitted, discovery-tools prompts for a password.

9.2. COMMAND LINE INTERFACE INSTALLATION OPTIONS

--version=*version.patch.release*

Enables the installation of a specific discovery CLI version. Contains the semantic versioning format (version.release.patch, such as 0.9.0) of the discovery CLI that you want to install. This option has a default of the latest version.

--home-dir=*cli_home_dir*

Sets the fully qualified path to the installation directory for the discovery CLI. Defaults to ~/**discovery**/.

--server-host=*server_IP_address*

Sets the IP address to use to connect to the discovery server. This option has no default value. You must either set a value during the CLI package installation or set a value by configuring the server connection for the command line interface.

--server-port=*server_port*

Sets the port number to use to connect to the discovery server. This option has no default value. You must either set a value during the CLI package installation or set a value by configuring the server connection for the command line interface.

PART IV. CONFIGURING AND MAINTAINING DISCOVERY

After installation is complete, you might have to complete other steps to configure or maintain discovery. The options that you choose during installation and the way in which you use discovery can determine the types of configuration and maintenance tasks that you need to perform.

Learn more

If you did not provide values to configure the connection between the discovery server and command line interface during installation, then you must configure those values before you can begin using discovery. To learn more, see the following information:

- Configuring discovery connections

If you are going to run network scans with credentials that include SSH keys as the authentication method, then the discovery server must have access to the keyfile information. To learn more about adding SSH keys to the discovery server, see the following information:

- Adding SSH keys to the discovery server for network scans

CHAPTER 10. CONFIGURING DISCOVERY CONNECTIONS

The discovery command line interface must communicate with the discovery server through a specific port and IP address.

During the installation of the CLI package, you might have provided this information by setting values for the **--server-port** and **--server-host** options. If you did not provide those values during CLI package installation, you must configure those values. You can also use this procedure to edit these values, as needed.

Procedure

To configure the command line interface to communicate with the server, run the following command. If the server and CLI packages are installed on different machines, run this command from the machine where the CLI package is installed:

1. Enter the following command, where ***server_port*** is the port that the discovery server is using for HTTPS communication ***server_host*** is the IP address of the server:

```
# dsc server config --port=server_port --host=server_host
```

CHAPTER 11. ADDING SSH KEYS TO THE DISCOVERY SERVER FOR NETWORK SCANS

When you configure sources and credentials for a network scan, you select the type of credential to use to authenticate to the network assets that are being scanned. One of the available options for the credential is to authenticate with a user name and SSH keyfile. If you choose this option, you must add a copy of the the private key to a specific directory on the server so that discovery can authenticate to those assets and complete the processes that occur during a scan.

You might have to perform these steps as an ongoing maintenance task as you create and refine the credentials needed for your network scans.

Procedure

To add an SSH keyfile to the discovery server:

1. Copy the private key from the keyfile, using the copy method of your choice.
2. Add the private key to the **~/discovery/server/volumes/sshkeys** directory on the discovery server, the default location for this directory at the time of server installation.
3. Repeat these steps as needed for all credentials that use SSH keyfiles as the authentication method.

PART V. ACCESSING THE DISCOVERY USER INTERFACE

You access the discovery graphical user interface through a browser connection. You access the command line interface by running a command to connect to the server.

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

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

- Logging in to the discovery command line interface
- Logging out of the discovery command line interface

CHAPTER 12. 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 13. LOGGING OUT OF THE DISCOVERY USER INTERFACE

Procedure

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

CHAPTER 14. LOGGING IN TO THE DISCOVERY COMMAND LINE INTERFACE

Prerequisites

To log in to the discovery command line interface, you need the IP address of the system where the discovery server is installed and the port number for the connection. These values are set for the command line interface either during the installation of the CLI package or after installation as a configuration step. You also need 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.

The login command retrieves a token that is used for authentication with subsequent command line interface commands. That token is removed when you log out of the server, and expires daily.

Procedure

1. To log in to the command line interface, enter the following command, where **server_port** is the port that the discovery server is using for HTTPS communication and ***server_host*** is the IP address of the server:

```
# dsc server login --port=server_port --host=server_host
```

2. At the prompts, enter the user name and password for discovery.

CHAPTER 15. LOGGING OUT OF THE DISCOVERY COMMAND LINE INTERFACE

The command to log out of the server removes the token that was created when you logged in to the server. This token also expires daily.

Procedure

1. To log out of the command line interface, enter the following command:

```
# dsc server logout
```