Subscription Central 1-latest

Installing and Configuring Discovery

Installing Discovery

Last Updated: 2024-06-25

Subscription Central 1-latest Installing and Configuring Discovery

Installing Discovery

Legal Notice

Abstract

Table of Contents

CHAPTER 1. ABOUT DISCOVERY

Discovery 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?

1.1. WHAT IS DISCOVERY?

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 subscription usage. Ultimately, this inspection and reporting process is part of the larger system administration task of managing your inventories.

Discovery 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 username 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 hostnames, IP addresses, IP ranges, or subnets. These assets can also be a systems management solution such as vCenter Server or Red Hat Satellite Server, or can be clusters deployed on Red Hat OpenShift Container Platform.

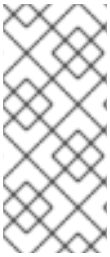

NOTE

Currently, the only virtualized deployment that discovery can scan with a specialized source for virtualization infrastructure is VMware vCenter. No other virtualization infrastructure that is supported by Red Hat can be scanned with a specialized scan. General scans of your network might still find these assets, without the precise metadata returned by a specialized scan.

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.

Discovery 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.

1.2. WHAT PRODUCTS DOES DISCOVERY FIND?

Discovery 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 Enterprise Linux version 9 and later

Red Hat Application Services products (formerly Red Hat Middleware)

- Red Hat JBoss BRMS version 5.0.1 and later, version 6.0.0 and later (also known as Red Hat Decision Manager, and currently part of Red Hat Process Automation Manager)
- JBoss Enterprise Web Server version 1 and later; Red Hat JBoss Web Server 3.0.1 and later
- 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 and later
- Red Hat Fuse version 6.0 and later

Red Hat Ansible Automation Platform

- Ansible Automation Platform version 2 and later

Red Hat OpenShift Container Platform

- Red Hat OpenShift Container Platform version 4 and later

Red Hat Advanced Cluster Security for Kubernetes

- Red Hat Advanced Cluster Security for Kubernetes version 4 and later

Red Hat Advanced Cluster Management for Kubernetes

- Red Hat Advanced Cluster Management for Kubernetes version 2 and later

1.3. IS DISCOVERY RIGHT FOR ME?

Discovery 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).

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 Discovery, the support team does not provide any assistance to help you understand the reports.

The Discovery tool does not automatically share data directly with Red Hat. Instead, you choose whether to prepare and send report data to Red Hat for ingestion by Red Hat tools and services. You can use the Discovery tool locally to scan your network for the Red Hat products that Discovery currently supports and then use the generated reports for your own internal purposes.

CHAPTER 2. INSTALLING PREREQUISITES FOR DISCOVERY

Discovery is a containerized solution that can be deployed on any RHEL or OpenShift Container Platform platform. The following documentation assumes Discovery is installed on a dedicated RHEL system. Adoption of this implementation will minimize impact on production servers, avoid complications of UI-to-server port mapping, and adhere to the officially documented and supported path for installation. You must install Discovery on a dedicated system that does not run any other workloads. No warranty or support is offered for installation on a nondedicated system.

Procedure

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

2.1. 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
- **RAM:** 1 GB minimum, with a recommended 2 GB
- **Disk Storage:** 30 GB

NOTE

Discovery uses rootless Podman, which, by default, stores containers in the following filepath: **${HOME}/.local/share/containers**. Ensure that the partition that contains this directory has at least 30 GB of storage. If Podman is configured to use a non-standard directory for container storage, ensure that the configured directory has 30 GB of storage.

2.2. SOFTWARE PREREQUISITES

The system on which you are going to install Discovery must meet the following software requirements:

- **Operating system:** The latest version of Red Hat Enterprise Linux 8 or Red Hat Enterprise Linux 9, installed as a clean install and specifically not upgraded from RHEL 7
- **File system:** Must run with d_type (Podman requirement)

In addition to these software requirements, Discovery has dependencies on other software that is required to install and run Discovery, primarily the Podman container tool. The Podman package is included with Red Hat Enterprise Linux 8 and later, so you are not required to obtain the Podman package separately. You will need sudo privileges to install Podman.

Related Information

- For more information about the definition of Red Hat Enterprise Linux clean installs, see the key migration terminology section in the *Upgrading from RHEL 8 to RHEL 9* guide.

2.3. OTHER ENVIRONMENT PREREQUISITES

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

In your network:

- If you want to use Discovery to scan a network that contains multiple air-gapped environments, you must install a Discovery server for each of those air-gapped environments.

On the system where Discovery is installed and running:

- The system should be a system that is dedicated to Discovery work only and should not be used for either development or production work.
- The system must have an internet connection to register to Red Hat subscription management tools and access Red Hat content.
  - For a disconnected installation, the connected system on which you are downloading the containers must meet this requirement.
- The system must have access to the Red Hat Ecosystem Catalog with your Red Hat Customer Portal credentials.
  - For a disconnected installation, the connected system on which you are downloading the containers must meet this requirement.

NOTE

If you cannot use the Red Hat Ecosystem Catalog to download Red Hat content, join the Red Hat Developer program to create a Red Hat Customer Portal account, obtain a Red Hat Developer subscription, and gain access to Red Hat content. For more information, see the Red Hat Developer website.

- The Discovery server must have access, through the SSH protocol, to the IT infrastructure assets that are to be scanned.

On the systems in your network where Discovery runs scans:

- Any network sources that are targeted for scanning must be running the 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 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 **"${HOME}"/.local/share/discovery/sshkeys** directory, the default location for this directory at the time of server installation.

- The user account that runs the **podman** commands to install and run Discovery must not be the **root** user, and it must not invoke these commands by using **sudo** or **su** to grant elevated user privileges. These actions are not supported.

Additional resources

- For more information about the creation of credentials for network scans and their association with elevated privileges, see the topic about adding network sources and credentials in the Installing and Configuring Discovery guide.

CHAPTER 3. INSTALLING DISCOVERY CONTAINERS WITH THE CONNECTED INSTALLATION PROCESS

You use the connected installation process when you want to run Discovery from a system that has internet access, also commonly referred to as a connected or online environment. During the Discovery connected installation process, you complete all steps from the system that you intend to use as the Discovery server. You will enter commands to obtain, install, configure the environment for, and start the Discovery containers from the container images on the Red Hat Ecosystem Catalog website.

The connected installation process uses Podman to perform the container installation tasks for the Discovery server and its supporting PostgreSQL database. You will need sudo permissions to install Podman on your system. However, running the Discovery application with **sudo** access or as the **root** user is not supported.

Prerequisites

- Before you begin, ensure that all hardware, software, and environment prerequisites are installed and configured. For more information about the prerequisites, see Installing prerequisites for Discovery.

Procedure

To install Discovery with the connected installation process, you do the following tasks:

- From a single connected system, run a series of commands:
  - Access and download the container images.
  - Install, configure, and start the Discovery server and database containers.

3.1. INSTALLING THE DISCOVERY DATABASE AND SERVER CONTAINER IMAGES FOR THE CONNECTED INSTALLATION

During the connected installation of the Discovery database and server container images, you obtain the Discovery containers from the container images on the Red Hat Ecosystem Catalog website, configure the containers and configure Discovery, and start the containers.

While you are completing the steps to configure the containers and configure Discovery, you must enter values for a number of environment variables. The example code provided in the following procedure includes example values for these environment variables. All of the environment variable values can be customized to suit the needs of your environment. For example, due to heavy traffic on your network, you might need to set a longer connection timeout than the default provided in the example command. To do so, you would change the value of the **-e NETWORK_CONNECT_JOB_TIMEOUT** environment variable.

Included in the environment variables that you set are the Discovery server administrator username, server administrator password, server hostname, and server port. These values can also be changed, and it is possible that security standards at your organization might require that you change them. At a minimum, you should change the password value to align with common security best practices.

The server administrator password must be:

- at least ten characters
- cannot be a word found in the dictionary

- cannot be the previously provided Discovery default passwords
- cannot be numeric only

The values for the username, password, and hostname are set with **-e** option for the **podman run** command. The value for the port is set with the **--publish 9443:443** option for the **podman run** command. The following list contains the defaults for these Discovery environment variables:

- server administrator username: **admin**
- server hostname: **localhost**
- server port: **9443**

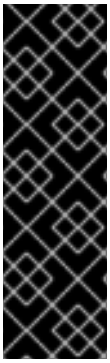

IMPORTANT

Note any changes to any usernames and passwords in the environment variables in the password management system that is used by your organization. Red Hat does not provide support for lost passwords for Discovery, and Discovery does not offer a method to recover these usernames and passwords.

In addition, if you upgrade Discovery, you must use the same database username and password during the upgrade. The failure to use the same database credentials could result in data loss of previously stored Discovery data.

Prerequisites

- The system on which you are downloading the Discovery container images must be connected to the internet.
- That system must be able to connect to the Red Hat Ecosystem Catalog.

Procedure

To install the Discovery server and database containers with the connected installation process, use the following steps:

1. Log in to the Red Hat Ecosystem Catalog (the registry.redhat.io website). When prompted, supply your Red Hat Customer Portal account credentials.

```
$ podman login registry.redhat.io
```

2. Create directories for the discovery server.

```
$ mkdir -p "${HOME}"/.local/share/discovery/data
$ mkdir -p "${HOME}"/.local/share/discovery/log
$ mkdir -p "${HOME}"/.local/share/discovery/sshkeys
```

3. Pull the container images for the Discovery database and server.

```
$ podman pull registry.redhat.io/rhel9/postgresql-15:latest
$ podman pull registry.redhat.io/discovery/discovery-server-rhel9:latest
```

4. Run the following command to start and configure the Discovery database container in a new **discovery-pod** pod. This command is shown on multiple lines for readability, but note the continuation characters at the end of each line.

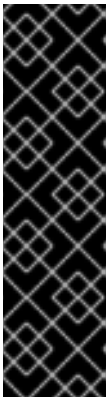

IMPORTANT

The following sample command contains the settings for multiple Discovery environment variables. The values of these environment variables can be changed as needed for your environment. In particular, the value for the Discovery server administrator password is used in multiple environment variables in this step and the following step and is shown as a variable because it is strongly recommended that you change this value. As you do this step, preserve any changed values for the environment variables related to usernames and passwords in your password management system. Discovery does not offer a method to recover these usernames and passwords.

```
$ podman run --name dsc-db \
  --pod new:discovery-pod \
  --publish 9443:443 \
  --restart on-failure \
  -e POSTGRESQL_USER=dsc \
  -e POSTGRESQL_PASSWORD=server_administrator_password \
  -e POSTGRESQL_DATABASE=dsc-db \
  -v dsc-data:/var/lib/pgsql/data:z \
  -d registry.redhat.io/rhel9/postgresql-15:latest
```

5. Run the following command to start and configure the Discovery server container in the **discovery-pod** pod. This command is shown on multiple lines for readability, but note the continuation characters at the end of each line.

IMPORTANT

The following sample command contains the settings for multiple Discovery environment variables. The values of these environment variables can be changed as needed for your environment. As in the previous step, the value for the Discovery server administrator password is shown as a variable. As you do this step, preserve any changed values for the environment variables related to usernames and passwords in your password management system. Discovery does not offer a method to recover these usernames and passwords.

```
$ podman run \
  --name discovery \
  --restart on-failure \
  --pod discovery-pod \
  -e DJANGO_DEBUG=False \
  -e NETWORK_CONNECT_JOB_TIMEOUT=60 \
  -e NETWORK_INSPECT_JOB_TIMEOUT=600 \
  -e PRODUCTION=True \
  -e QPC_DBMS_HOST=localhost \
  -e QPC_DBMS_PASSWORD=server_administrator_password \
  -e QPC_DBMS_USER=dsc \
  -e QPC_DBMS_DATABASE=dsc-db \
  -e QPC_SERVER_PASSWORD=server_administrator_password \
  -e QPC_SERVER_TIMEOUT=120 \
```

```
-e QPC_SERVER_USERNAME=admin \
-e QPC_SERVER_USER_EMAIL=admin@example.com \
-v "${HOME}"/.local/share/discovery/data/:/var/data:z \
-v "${HOME}"/.local/share/discovery/log/:/var/log:z \
-v "${HOME}"/.local/share/discovery/sshkeys/:/sshkeys:z \
-d registry.redhat.io/discovery/discovery-server-rhel9:latest
```

Additional resources

- For more information about the optional procedure to install the Discovery command line interface, see Installing the Discovery command line interface .
- For more information about logging in to Discovery, see Accessing the Discovery user interface.

CHAPTER 4. INSTALLING DISCOVERY CONTAINERS WITH THE DISCONNECTED INSTALLATION PROCESS

The disconnected installation process is required when you want to run Discovery from a system that does not have internet access, also commonly referred to as a disconnected, offline, or air-gapped environment. During the Discovery disconnected installation process, you must complete some steps on a *connected* system, one that has internet connectivity. These steps include obtaining the container images from the Red Hat Ecosystem Catalog website and then transferring the images to the disconnected system that you intend to use as the Discovery server. You also complete steps on the *disconnected* system to install, configure the environment for, and start the Discovery containers.

The disconnected installation process uses Podman to install the containers for the Discovery server and its supporting PostgreSQL database. You will need sudo access to install Podman on your system if it is not already present. However, running the Discovery application with **sudo** access or as the **root** user is not supported.

Prerequisites

- Before you begin, ensure that all hardware, software, and environment prerequisites are installed and configured. For more information about the prerequisites, see Installing prerequisites for Discovery.

Procedure

To install Discovery with the disconnected installation process, you do the following tasks:

- From the connected system, run a series of commands:
  - Access and download the container images.
- From the disconnected system, run a series of commands:
  - Transfer the container images to the disconnected system.
  - Install, configure, and start the Discovery server and database containers.

4.1. INSTALLING THE DISCOVERY DATABASE AND SERVER CONTAINER IMAGES FOR THE DISCONNECTED INSTALLATION

During the disconnected installation of the Discovery database and server container images, you begin with a connected system to install the Podman container management tool and obtain the Discovery containers from the container images on the Red Hat Ecosystem Catalog website. You then transfer those container images to the disconnected system where you want to run Discovery, install Podman on the disconnected system, configure the containers and configure Discovery, and start the containers.

While you are completing the steps to configure the containers and configure Discovery, you must enter values for a number of environment variables. The example code provided in the following procedure includes example values for these environment variables. All of the environment variable values can be customized to suit the needs of your environment. For example, due to heavy traffic on your network, you might need to set a longer connection timeout than the default provided in the example command. To do so, you would change the value of the **-e NETWORK_CONNECT_JOB_TIMEOUT** environment variable.

Included in the environment variables that you set are the Discovery server administrator username, server administrator password, server hostname, and server port. These values can also be changed, and

it is possible that security standards at your organization might require that you change them. At a minimum, you should change the password value to align with common security best practices.

The server administrator password must be:

- at least ten characters
- cannot be a word found in the dictionary
- cannot be the previously provided Discovery default passwords
- cannot be numeric only

The values for the username, password, and hostname are set with **-e** option for the **podman run** command. The value for the port is set with the **--publish 9443:443** option for the **podman run** command. The following list contains the defaults for these Discovery environment variables:

- server administrator username: **admin**
- server hostname: **localhost**
- server port: **9443**

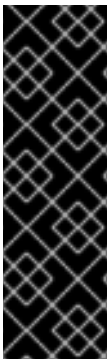

IMPORTANT

Note any changes to any usernames and passwords in the environment variables in the password management system that is used by your organization. Red Hat does not provide support for lost passwords for Discovery, and Discovery does not offer a method to recover these usernames and passwords.

In addition, if you upgrade Discovery, you must use the same database username and password during the upgrade. The failure to use the same database credentials could result in data loss of previously stored Discovery data.

Prerequisites

- The connected system on which you are downloading the Discovery container images must be connected to the internet.
- The connected system must be able to connect to the Red Hat Ecosystem Catalog.

Procedure

To install the Discovery server and database containers with the disconnected installation process, use the following steps:

On the connected system

1. Log in to the Red Hat Ecosystem Catalog (the registry.redhat.io website). When prompted, supply your Red Hat Customer Portal account credentials.

```
$ podman login registry.redhat.io
```

2. Pull the Discovery database and server containers and save them as **.tar.gz** files.

```
$ podman pull registry.redhat.io/rhel9/postgresql-15:latest
```

```
$ podman save registry.redhat.io/rhel9/postgresql-15:latest -o postgres.tar.gz
$ podman pull registry.redhat.io/discovery/discovery-server-rhel9:latest
$ podman save registry.redhat.io/discovery/discovery-server-rhel9 -o discovery.tar.gz
```

On the disconnected system

1. Use the transfer method of your choice to transfer the **postgres.tar.gz** and **discovery.tar.gz** files to the disconnected system where you are going to install Discovery.

2. Load the container images into container storage.

```
$ podman load -i postgres.tar.gz
$ podman load -i discovery.tar.gz
```

3. Create directories for the discovery server.

```
$ mkdir -p "${HOME}"/.local/share/discovery/data
$ mkdir -p "${HOME}"/.local/share/discovery/log
$ mkdir -p "${HOME}"/.local/share/discovery/sshkeys
```

4. Run the following command to start and configure the Discovery database container in a new **discovery-pod** pod. This command is shown on multiple lines for readability, but note the continuation characters at the end of each line.

IMPORTANT

The following sample command contains the settings for multiple Discovery environment variables. The values of these environment variables can be changed as needed for your environment. In particular, the value for the Discovery server administrator password is used in multiple environment variables in this step and the following step and is shown as a variable because it is strongly recommended that you change this value. As you do this step, preserve any changed values for the environment variables related to usernames and passwords in your password management system. Discovery does not offer a method to recover these usernames and passwords.

```
$ podman run --name dsc-db \
  --pod new:discovery-pod \
  --publish 9443:443 \
  --restart on-failure \
  -e POSTGRESQL_USER=dsc \
  -e POSTGRESQL_PASSWORD=server_administrator_password \
  -e POSTGRESQL_DATABASE=dsc-db \
  -v dsc-data:/var/lib/pgsql/data:z \
  -d registry.redhat.io/rhel9/postgresql-15:latest
```

5. Run the following command to start and configure the Discovery server container in the **discovery-pod** pod. This command is shown on multiple lines for readability, but note the continuation characters at the end of each line.

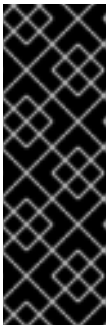

IMPORTANT

The following sample command contains the settings for multiple Discovery environment variables. The values of these environment variables can be changed as needed for your environment. As in the previous step, the value for the Discovery server administrator password is shown as a variable. As you do this step, preserve any changed values for the environment variables related to usernames and passwords in your password management system. Discovery does not offer a method to recover these usernames and passwords.

```
$ podman run \
  --name discovery \
  --restart on-failure \
  --pod discovery-pod \
  -e DJANGO_DEBUG=False \
  -e NETWORK_CONNECT_JOB_TIMEOUT=60 \
  -e NETWORK_INSPECT_JOB_TIMEOUT=600 \
  -e PRODUCTION=True \
  -e QPC_DBMS_HOST=localhost \
  -e QPC_DBMS_PASSWORD=server_administrator_password \
  -e QPC_DBMS_USER=dsc \
  -e QPC_DBMS_DATABASE=dsc-db \
  -e QPC_SERVER_PASSWORD=server_administrator_password \
  -e QPC_SERVER_TIMEOUT=120 \
  -e QPC_SERVER_USERNAME=admin \
  -e QPC_SERVER_USER_EMAIL=admin@example.com \
  -v "${HOME}"/.local/share/discovery/data/:/var/data:z \
  -v "${HOME}"/.local/share/discovery/log/:/var/log:z \
  -v "${HOME}"/.local/share/discovery/sshkeys/:/sshkeys:z \
  -d registry.redhat.io/discovery/discovery-server-rhel9:latest
```

Additional resources

- For more information about the optional procedure to install the Discovery command line interface, see Installing the Discovery command line interface .
- For more information about logging in to Discovery, see Accessing the Discovery user interface.

CHAPTER 5. INSTALLING THE DISCOVERY COMMAND LINE INTERFACE

After you have finished the procedure to install the Discovery server and database containers, you can install the Discovery command line interface (CLI) as an optional step.

The Discovery command line interface is available within the server container. Therefore, the code needed to install the command line interface is already available on the system where you installed the server and database containers, regardless of whether you used the connected or disconnected procedure for the installation.

Prerequisites

- You must install the Discovery command line interface on the same system where you installed the Discovery server and database containers.

5.1. INSTALLING THE DISCOVERY COMMAND LINE INTERFACE

The installation process for the Discovery command line interface includes subscribing to the discovery content delivery network (CDN) repository and installing the discovery RPM package on your system.

Prerequisites

- You must have root privileges or their equivalent to run some of the commands in the following procedure.

Procedure

To install the Discovery command line interface, use the following steps:

1. To subscribe to the Discovery CDN repositories for your operating system version, enter the following command:

```
# subscription-manager repos --enable <repository>
```

NOTE

In the preceding command, replace *<repository>* with the following values:

- In Red Hat Enterprise Linux 8, replace *<repository>* with **discovery-1-for-rhel-8-x86_64-rpms**
- In Red Hat Enterprise Linux 9, replace *<repository>* with **discovery-1-for-rhel-9-x86_64-rpms**

2. To install the RPM package, enter the following command as the root user:

```
# dnf install discovery-cli
```

3. Configure the hostname by using the **hostname** command and port that the Discovery command line interface uses to communicate with the Discovery server. For the **server_port** value, enter the port that is being used for HTTPS communication.

NOTE

In the procedure to install the Discovery server and database containers, the value for the *server_port* was set by the following option in the **podman run** command: **--publish 9443:443**. If you did not use **9443** as the server port, enter the value that you used.

```
$ dsc server config --host hostname --port server_port
```

Additional resources

- For more information about logging in to Discovery, see Accessing the Discovery user interface.

CHAPTER 6. ACCESSING THE DISCOVERY USER INTERFACE

You access the Discovery graphical user interface through a browser. You access the Discovery command line interface by running a terminal session command to connect to the server.

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

6.1. LOGGING IN TO THE DISCOVERY USER INTERFACE

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 server administrator username and password to use when logging in. If you do not have this information, contact the administrator who installed the Discovery server.

Prerequisites

- To use the Discovery graphical 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.

Procedure

1. In a browser, enter the URL for the Discovery server in the following format:
   https://*IPaddress*:*server_port*, where ***IPaddress*** is the IP address of the Discovery server and ***server_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 **9443** 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 username and password for the Discovery server administrator account 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.

6.2. LOGGING OUT OF THE DISCOVERY USER INTERFACE

Procedure

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

6.3. LOGGING IN TO THE DISCOVERY COMMAND LINE INTERFACE

To log in to the Discovery command line interface, you need the username and password for the Discovery server administrator. 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.

Prerequisites

- The Discovery command line interface must be installed on the same system where the server container is installed.
- You must access the command line interface on the same system where the server container is installed.

Procedure

1. To log in to the command line interface, enter the following command, where ***server_username*** is the username for the Discovery server administrator and ***server_password*** is the password for the server administrator:

```
$ dsc server login --username server_administrator_username --password
server_administrator_password
```

After you log in to the command line interface, you can begin entering commands to create sources, credentials, and scans.

6.4. 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
```

CHAPTER 7. 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 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

7.1. 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 username and SSH keyfile. If you choose this option, you must add a copy of 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.

IMPORTANT

Each SSH private key provided must be copied into the directory that was mapped to **/sshkeys** path on the server container during the Discovery server installation. In other words, the SSH private key must be copied to your file system at the mount point where the container will look for it. The default path for this mount point directory is **"${HOME}"/.local/share/discovery/sshkeys** on the system where Discovery is installed. That file path is a shared volume to the container at **discovery:/sshkeys**.

This process is required because the container must have a standardized mount point to map to the container volume during container initialization. Using a standardized mount point is required for security reasons. A container should never have full access to your entire file system. When you are using the Discovery command line interface or the graphical user interface, using the full path to a resource will result in an error stating that the file is not a valid file on the file system. This message occurs because the container searches for the path on its own file system, not at the full path that is passed to it.

When you or other Discovery users are using the graphical user interface to create network credentials that use SSH, the field that requires the mount point directory location is the **SSH Key File** field. For the command line interface, it is the **--sshkeyfile** argument. For both of these options, the default value of the mount point directory is **"${HOME}"/.local/share/discovery/sshkeys**.

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 **"${HOME}"/.local/share/discovery/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, including when relevant new credentials are added.