



Subscription Central 1-latest

Troubleshooting Discovery

Troubleshooting Discovery

Subscription Central 1-latest Troubleshooting Discovery

Troubleshooting Discovery

Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] 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

CHAPTER 1. ABOUT DISCOVERY	3
1.1. WHAT IS DISCOVERY?	3
1.2. WHAT PRODUCTS DOES DISCOVERY FIND?	4
1.3. IS DISCOVERY RIGHT FOR ME?	5
CHAPTER 2. TROUBLESHOOTING DISCOVERY	6
2.1. DETERMINING THE VERSION OF THE DISCOVERY SERVER	6
2.2. UNINSTALLING DISCOVERY	6
2.3. GETTING HELP WITH THE COMMAND LINE INTERFACE	6
2.4. SSH CREDENTIAL CONFIGURATION	7
2.5. LOG FILE LOCATIONS	7
2.6. BACKING UP OR RESTORING THE SERVER ENCRYPTION KEY	7
2.7. RESTARTING THE DISCOVERY SERVER AFTER A REBOOT	8
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	9

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)

- 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. TROUBLESHOOTING DISCOVERY

2.1. DETERMINING THE VERSION OF THE DISCOVERY SERVER

Prerequisites

- You must be logged in to the command line interface as the Discovery server administrator.

Procedure

To determine the version of the Discovery server, use the following steps:

- Enter the **dsc server status** command. The expected output provides the version of the server that you are using:

```
"server_address": "127.0.0.1:9443", "server_id":  
"45a8ea20-2ec4-4113-b459-234fed505b0d", "server_version":  
"1.0.0.3e15fa8786a974c9eafe6376ff31ae0211972c36"
```

If you cannot get the server status command to run, or you cannot log in to the server, use the following Podman images command:

```
podman images --filter 'reference=registry.redhat.io/discovery/discovery-server-rhel9' --  
format '{{.Labels.url}}'
```

2.2. UNINSTALLING DISCOVERY

Prerequisites

- You must be logged in to the system that is running Discovery.
- You will need sudo access to perform certain functions in **dnf**.

Procedure

To uninstall Discovery server, use the following steps:

1. Run the uninstall command.

```
discovery-installer uninstall
```

2. Uninstall the installer package.

```
sudo dnf remove discovery-installer
```

3. Uninstall the command line interface, if installed.

```
sudo dnf remove discovery-cli
```

2.3. GETTING HELP WITH THE COMMAND LINE INTERFACE

Prerequisites

- You must be logged in to the command line interface as the Discovery server administrator.

Procedure

- For help on general topics, see the man page information.
- For help on a specific subcommand, use the **-h** option. For example:

```
dsc cred -h
dsc source -h
dsc scan -h
```

2.4. SSH CREDENTIAL CONFIGURATION

If you receive an error message that includes text similar to **not a valid file on the filesystem**, that message might indicate an issue with the mount point on the file system that enables access to the SSH keyfiles.

When you are creating your network credentials with SSH keyfiles, make sure that the copy of the private key has been correctly added to the **"\${HOME}"/.local/share/discovery/sshkeys** directory on the server.

2.5. LOG FILE LOCATIONS

Prerequisites

- You must be logged in to the system that is running Discovery.
- You will need sudo access to perform certain functions in **dnf**.

Procedure

Log files for the Discovery server that are on the local file system are located in the following path: **"\${HOME}"/.local/share/discovery/log**.

Log data is also copied to **stdout** and can be accessed through Podman logs. To follow the log output, include the **-f** option as shown in the following command:

```
podman logs -f discovery-server
podman logs -f discovery-celery-worker
```

2.6. BACKING UP OR RESTORING THE SERVER ENCRYPTION KEY

Passwords are not stored as plain text. They are encrypted and decrypted by using the content of the **secret.txt** file as a secret key. If you need to back up and restore the **secret.txt** file, use these steps.

Prerequisites

- You must be logged in to the system that is running Discovery.
- You will need sudo access to perform certain functions in **dnf**.

Procedure

- To back up the encrypted SSH credentials, navigate to `"${HOME}"/.local/share/discovery/data` directory and copy the **secret.txt** file.
- To restore the **secret.txt** file, enter the following command, where *path_to_backup* is the path where the **secret.txt** file is backed up:

```
cp -p __path_to_backup__ /secret.txt "${HOME}"/.local/share/discovery/data/
```

2.7. RESTARTING THE DISCOVERY SERVER AFTER A REBOOT

Prerequisites

- You must be logged in to the system that is running Discovery.
- You will need sudo access to perform certain functions in **dnf**.



NOTE

If you installed Discovery using the standard process, the system should start automatically after a reboot. If it does not automatically restart, use the following procedure:

Procedure

- To restart the Discovery application after a reboot, use the following command:

```
systemctl --user restart discovery-app
```

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your feedback on our documentation. To provide feedback, open a Jira issue that describes your concerns. Provide as much detail as possible so that your request can be addressed quickly.

Prerequisites

- You have a Red Hat Customer Portal account. This account enables you to log in to the Red Hat Jira Software instance. If you do not have an account, you will be prompted to create one.

Procedure

To provide your feedback, perform the following steps:

1. Click the following link: [Create Issue](#).
2. In the **Summary** text box, enter a brief description of the issue.
3. In the **Description** text box, provide more details about the issue. Include the URL where you found the issue.
4. Provide information for any other required fields. Allow all fields that contain default information to remain at the defaults.
5. Click **Create** to create the Jira issue for the documentation team.

A documentation issue will be created and routed to the appropriate documentation team. Thank you for taking the time to provide feedback.