



Set up and configure Keystone

Keystone

NetApp
January 20, 2026

Table of Contents

Set up and configure Keystone	1
Requirements	1
Virtual infrastructure requirements for Keystone Collector	1
Linux requirements for Keystone Collector	3
Requirements for ONTAP and StorageGRID for Keystone	5
Install Keystone Collector	8
Deploy Keystone Collector on VMware vSphere systems	8
Install Keystone Collector on Linux systems	10
Automatic validation of Keystone software	12
Configure Keystone Collector	12
Configure HTTP Proxy on Keystone Collector	14
Limit collection of private data	14
Trust a custom root CA	15
Create Performance Service Levels	16
Install ITOM Collector	20
Installation requirements for Keystone ITOM Collector	21
Install Keystone ITOM Collector on Linux systems	22
Install Keystone ITOM Collector on Windows systems	23
Configure AutoSupport for Keystone	24
Monitor and upgrade	25
Monitor the health of Keystone Collector	25
Manually upgrade Keystone Collector	30
Keystone Collector security	32
Security hardening	32
Types of user data that Keystone collects	33
ONTAP data collection	33
StorageGRID data collection	40
Telemetry data collection	41
Keystone in private mode	42
Learn about Keystone (private mode)	42
Prepare for Keystone Collector installation in private mode	44
Install Keystone Collector in private mode	45
Configure Keystone Collector in private mode	46
Monitor Keystone Collector health in private mode	51

Set up and configure Keystone Requirements

Virtual infrastructure requirements for Keystone Collector

Your VMware vSphere system must meet several requirements before you can install Keystone Collector.

Prerequisites for the Keystone Collector server VM:

- Operating system: VMware vCentre server and ESXi 8.0 or later
- Core: 1 CPU
- RAM: 2 GB RAM
- Disk space: 20 GB vDisk

Other requirements

Ensure that the following generic requirements are met:

Networking Requirements

The networking requirements of Keystone Collector are listed in the following table.



Keystone Collector requires internet connectivity. You can provide internet connectivity by direct routing through default Gateway (via NAT), or through HTTP Proxy. Both variants are described here.

Source	Destination	Service	Protocol and Ports	Category	Purpose
Keystone Collector (for Keystone ONTAP)	Active IQ Unified Manager (Unified Manager)	HTTPS	TCP 443	Mandatory (if using Keystone ONTAP)	Keystone Collector usage metrics collection for ONTAP
Keystone Collector (for Keystone StorageGRID)	StorageGRID Admin Nodes	HTTPS	TCP 443	Mandatory (if using Keystone StorageGRID)	Keystone Collector usage metrics collection for StorageGRID
Keystone Collector (generic)	Internet (as per URL requirements given later)	HTTPS	TCP 443	Mandatory (internet connectivity)	Keystone Collector software, OS updates, and metrics upload

Keystone Collector (generic)	Customer HTTP Proxy	HTTP Proxy	Customer Proxy Port	Mandatory (internet connectivity)	Keystone Collector software, OS updates, and metrics upload
Keystone Collector (generic)	Customer DNS Servers	DNS	TCP/UDP 53	Mandatory	DNS resolution
Keystone Collector (generic)	Customer NTP Servers	NTP	UDP 123	Mandatory	Time synchronization
Keystone Collector (for Keystone ONTAP)	Unified Manager	MYSQL	TCP 3306	Optional Functionality	Performance metrics collection for Keystone Collector
Keystone Collector (generic)	Customer Monitoring System	HTTPS	TCP 7777	Optional Functionality	Keystone Collector health reporting
Customer's Operations Workstations	Keystone Collector	SSH	TCP 22	Management	Access to the Keystone Collector Management
NetApp ONTAP Cluster and Node Management Addresses	Keystone Collector	HTTP_8000, PING	TCP 8000, ICMP Echo Request/Reply	Optional Functionality	Webserver for ONTAP firmware updates

 The default port for MySQL, 3306, is restricted only to localhost during a fresh installation of Unified Manager, which prevents the collection of performance metrics for Keystone Collector. For more information, see [ONTAP requirements](#).

URL access

Keystone Collector needs access to the following internet hosts:

Address	Reason
https://keystone.netapp.com	Keystone Collector software updates and usage reporting

Linux requirements for Keystone Collector

Preparing your Linux system with the required software ensures precise installation and data collection by Keystone Collector.

Ensure that your Linux and Keystone Collector server VM have these configurations.

Linux server:

- Operating system: Any one of the following:
 - Debian 12
 - Red Hat Enterprise Linux 8.6 or later 8.x versions
 - Red Hat Enterprise Linux 9.0 or later versions
 - CentOS 7 (for existing environments only)
- Chronyd time synchronized
- Access to the standard Linux software repositories

The same server should also have the following third-party packages:

- podman (POD Manager)
- sos
- chrony
- python 3 (3.9.14 to 3.11.8)

Keystone Collector server VM:

- Core: 2 CPUs
- RAM: 4 GB RAM
- Disk space: 50 GB vDisk

Other requirements

Ensure that the following generic requirements are met:

Networking Requirements

The networking requirements of Keystone Collector are listed in the following table.



Keystone Collector requires internet connectivity. You can provide internet connectivity by direct routing through default Gateway (via NAT), or through HTTP Proxy. Both variants are described here.

Source	Destination	Service	Protocol and Ports	Category	Purpose
--------	-------------	---------	--------------------	----------	---------

Keystone Collector (for Keystone ONTAP)	Active IQ Unified Manager (Unified Manager)	HTTPS	TCP 443	Mandatory (if using Keystone ONTAP)	Keystone Collector usage metrics collection for ONTAP
Keystone Collector (for Keystone StorageGRID)	StorageGRID Admin Nodes	HTTPS	TCP 443	Mandatory (if using Keystone StorageGRID)	Keystone Collector usage metrics collection for StorageGRID
Keystone Collector (generic)	Internet (as per URL requirements given later)	HTTPS	TCP 443	Mandatory (internet connectivity)	Keystone Collector software, OS updates, and metrics upload
Keystone Collector (generic)	Customer HTTP Proxy	HTTP Proxy	Customer Proxy Port	Mandatory (internet connectivity)	Keystone Collector software, OS updates, and metrics upload
Keystone Collector (generic)	Customer DNS Servers	DNS	TCP/UDP 53	Mandatory	DNS resolution
Keystone Collector (generic)	Customer NTP Servers	NTP	UDP 123	Mandatory	Time synchronization
Keystone Collector (for Keystone ONTAP)	Unified Manager	MYSQL	TCP 3306	Optional Functionality	Performance metrics collection for Keystone Collector
Keystone Collector (generic)	Customer Monitoring System	HTTPS	TCP 7777	Optional Functionality	Keystone Collector health reporting
Customer's Operations Workstations	Keystone Collector	SSH	TCP 22	Management	Access to the Keystone Collector Management

NetApp ONTAP Cluster and Node Management Addresses	Keystone Collector	HTTP_8000, PING	TCP 8000, ICMP Echo Request/Reply	Optional Functionality	Webserver for ONTAP firmware updates
--	--------------------	-----------------	-----------------------------------	------------------------	--------------------------------------



The default port for MySQL, 3306, is restricted only to localhost during a fresh installation of Unified Manager, which prevents the collection of performance metrics for Keystone Collector. For more information, see [ONTAP requirements](#).

URL access

Keystone Collector needs access to the following internet hosts:

Address	Reason
https://keystone.netapp.com	Keystone Collector software updates and usage reporting
https://support.netapp.com	NetApp HQ for billing information and AutoSupport delivery

Requirements for ONTAP and StorageGRID for Keystone

Before you get started with Keystone, you need to ensure that ONTAP clusters and StorageGRID systems meet a few requirements.

ONTAP

Software versions

1. ONTAP 9.8 or later
2. Active IQ Unified Manager (Unified Manager) 9.10 or later

Before you begin

Meet the following requirements if you intend to collect usage data only through ONTAP:

1. Ensure that ONTAP 9.8 or later is configured. For information about configuring a new cluster, see these links:
 - [Configure ONTAP on a new cluster with System Manager](#)
 - [Set up a cluster with the CLI](#)
2. Create ONTAP login accounts with specific roles. To learn more, refer to [Learn about creating ONTAP login accounts](#).
 - **Web UI**
 - a. Log in to ONTAP System Manager using your default credentials. To learn more, refer to [Cluster management with System Manager](#).
 - b. Create an ONTAP user with the "readonly" role and "http" application type, and enable the password authentication by navigating to **Cluster > Settings > Security > Users**.
 - **CLI**
 - a. Log in to ONTAP CLI using your default credentials. To learn more, refer to [Cluster management with CLI](#).
 - b. Create an ONTAP user with the "readonly" role and "http" application type, and enable the password authentication. To learn more about authentication, refer to [Enable ONTAP account password access](#).

Meet the following requirements if you intend to collect usage data through Active IQ Unified Manager:

1. Ensure that Unified Manager 9.10 or later is configured. For information about installing Unified Manager, see these links:
 - [Installing Unified Manager on VMware vSphere systems](#)
 - [Installing Unified Manager on Linux systems](#)
2. Ensure that the ONTAP cluster has been added to Unified Manager. For information about adding clusters, see [Adding clusters](#).
3. Create Unified Manager users with specific roles for usage and performance data collection. Perform these steps. For information about user roles, see [Definitions of user roles](#).
 - a. Log into the Unified Manager web UI with the default application administrator user credentials that are generated during installation. See [Accessing the Unified Manager web UI](#).
 - b. Create a service account for Keystone Collector with `operator` user role. The Keystone Collector service APIs use this service account to communicate with Unified Manager and collect usage data. See [Adding users](#).
 - c. Create a Database user account, with the `Report Schema` role. This user is required for performance data collection. See [Creating a database user](#).



The default port for MySQL, 3306, is restricted only to localhost during a fresh installation of Unified Manager, which prevents the collection of performance data for Keystone ONTAP. This configuration can be modified, and the connection can be made available to other hosts using the Control access to MySQL port 3306 option on the Unified Manager maintenance console. For information, see [Additional menu options](#).

4. Enable API Gateway in Unified Manager. Keystone Collector makes use of the API Gateway feature to communicate with ONTAP clusters. You can enable API Gateway either from the web UI, or by running a few commands through Unified Manager CLI.

Web UI

To enable API Gateway from the Unified Manager web UI, log into the Unified Manager web UI and enable API Gateway. For information, see [Enabling API Gateway](#).

CLI

To enable API Gateway through Unified Manager CLI, follow these steps:

- a. On the Unified Manager server, begin an SSH session and log into Unified Manager CLI.
`um cli login -u <umadmin>`
For information about CLI commands, see [Supported Unified Manager CLI commands](#).
- b. Verify whether API Gateway is already enabled.
`um option list api.gateway.enabled`
A true value indicates that the API Gateway is enabled.
- c. If the value returned is false, run this command:
`um option set api.gateway.enabled=true`
- d. Restart the Unified Manager server:
 - Linux: [Restarting Unified Manager](#).
 - VMware vSphere: [Restarting the Unified Manager virtual machine](#).

StorageGRID

The following configurations are required for installing Keystone Collector on StorageGRID.

- StorageGRID 11.6.0 or later should be installed. For information about upgrading StorageGRID, see [Upgrade StorageGRID software: Overview](#).
- A StorageGRID local admin user account should be created for usage data collection. This service account is used by the Keystone Collector service for communicating with StorageGRID through administrator node APIs.

Steps

1. Log into the Grid Manager. See [Sign in to the Grid Manager](#).
2. Create a local admin group with Access mode: Read-only. See [Create an admin group](#).
3. Add the following permissions:
 - Tenant Accounts
 - Maintenance
 - Metrics Query
4. Create a Keystone service account user and associate it with the admin group. See [Manage](#)

Install Keystone Collector

Deploy Keystone Collector on VMware vSphere systems

Deploying Keystone Collector on VMware vSphere systems includes downloading the OVA template, deploying the template by using the **Deploy OVF Template** wizard, verifying the integrity of the certificates, and verifying the readiness of the VM.

Deploying the OVA template

Follow these steps:

Steps

1. Download the OVA file from [this link](#) and store it on your VMware vSphere system.
2. On your VMware vSphere system, navigate to the **VMs and Templates** view.
3. Right click on the required folder for the virtual machine (VM) (or data center, if not using VM folders) and select **Deploy OVF Template**.
4. On *Step 1* of the **Deploy OVF Template** wizard, click **Select and OVF template** to select the downloaded `KeystoneCollector-latest.ova` file.
5. On *Step 2*, specify the VM name and select the VM folder.
6. On *Step 3*, specify the required compute resource that is to run the VM.
7. On *Step 4: Review details*, verify the correctness and authenticity of the OVA file.

The vCenter root trust store contains only VMware certificates. NetApp uses Entrust as a certifying authority, and those certificates need to be added to the vCenter trust store.

- a. Download the code-signing CA certificate from Sectigo [here](#).
- b. Follow the steps in the Resolution section of this knowledge base (KB) article:
<https://kb.vmware.com/s/article/84240>.



For vCenter versions 7.x and earlier, you must update vCenter and ESXi to version 8.0 or later. Earlier versions are no longer supported.

When the integrity and authenticity of the Keystone Collector OVA are validated, you can see the text (Trusted certificate) with the publisher.

Deploy OVF Template

- 1 Select an OVF template
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details**
- 5 Select storage
- 6 Select networks
- 7 Customize template
- 8 Ready to complete

Review details

Verify the template details.

Publisher	Sectigo Public Code Signing CA R36 (Trusted certificate)
Product	Keystone-Collector
Version	3.12.31910
Vendor	NetApp
Download size	1.7 GB
Size on disk	3.9 GB (thin provisioned) 19.5 GB (thick provisioned)

CANCEL
BACK
NEXT

8. On **Step 5** of the **Deploy OVF Template** wizard, specify the location for storing the VM.
9. On **Step 6**, select the destination network for the VM to use.
10. On **Step 7 Customize template**, specify the initial network address and password for the admin user account.

i

The admin password is stored in a reversible format in vCentre and should be used as a bootstrap credential to gain initial access to the VMware vSphere system. During the initial software configuration, this admin password should be changed. The subnet mask for the IPv4 address should be supplied in CIDR notation. For example, use the value of 24 for a subnet mask of 255.255.255.0.

11. On **Step 8 Ready to complete** of the **Deploy OVF Template** wizard, review the configuration and verify that you have correctly set the parameters for the OVA deployment.

After the VM has been deployed from the template and powered on, open an SSH session to the VM and log in with the temporary admin credentials to verify that the VM is ready for configuration.

Initial system configuration

Perform these steps on your VMware vSphere systems for an initial configuration of the Keystone Collector servers deployed through OVA:

i

On completing the deployment, you can use the Keystone Collector Management Terminal User Interface (TUI) utility to perform the configuration and monitoring activities. You can use various keyboard controls, such as the Enter and arrow keys, to select the options and navigate across this TUI.

1. Open an SSH session to the Keystone Collector server. When you connect, the system will prompt you to update the admin password. Complete the admin password update as required.
2. Log in using the new password to access the TUI. On login, the TUI appears.

Alternatively, you can launch it manually by running the `keystone-collector-tui` CLI command.

3. If required, configure the proxy details in the **Configuration > Network** section on the TUI.
4. Configure the system hostname, location, and NTP server in the **Configuration > System** section.
5. Update the Keystone Collectors using the **Maintenance > Update Collectors** option. After the update, restart the Keystone Collector management TUI utility to apply the changes.

Install Keystone Collector on Linux systems

You can install the Keystone Collector software on a Linux server using an RPM or a Debian package. Follow the installation steps depending on your Linux distribution.

Using RPM

1. SSH to the Keystone Collector server and elevate to `root` privilege.

2. Import the Keystone public signing signature:

```
# rpm --import https://keystone.netapp.com/repo1/RPM-GPG-NetApp-Keystone-20251020
```

3. Ensure that the correct public certificate has been imported by checking the fingerprint for Keystone Billing Platform in the RPM database:

```
# rpm -qa gpg-pubkey --qf '%{Description}' | gpg --show-keys --fingerprint  
The correct fingerprint looks like this:
```

```
9297 0DB6 0867 22E7 7646 E400 4493 5CBB C9E9 FEDC
```

4. Download the `kestonerepo.rpm` file:

```
curl -O https://keystone.netapp.com/repo1/kestonerepo.rpm
```

5. Verify the authenticity of the file:

```
rpm --checksig -v kestonerepo.rpm
```

A signature for an authentic file looks like this:

```
Header V4 RSA/SHA512 Signature, key ID c9e9fedc: OK
```

6. Install the YUM software repository file:

```
# yum install kestonerepo.rpm
```

7. When the Keystone repo is installed, install the `keystone-collector` package through the YUM package manager:

```
# yum install keystone-collector
```

For Red Hat Enterprise Linux 9, run the following command to install the `keystone-collector` package:

```
# yum install keystone-collector-rhel9
```

Using Debian

1. SSH to the Keystone Collector server and elevate to `root` privilege.

```
sudo su
```

2. Download the `keystone-sw-repo.deb` file:

```
curl -O https://keystone.netapp.com/downloads/keystone-sw-repo.deb
```

3. Install the Keystone software repository file:

```
# dpkg -i keystone-sw-repo.deb
```

4. Update the package list:

```
# apt-get update
```

5. When the Keystone repo is installed, install the `keystone-collector` package:

```
# apt-get install keystone-collector
```

 On completing the installation, you can use the Keystone Collector Management Terminal User Interface (TUI) utility to perform the configuration and monitoring activities. You can use various keyboard controls, such as the Enter and arrow keys, to select the options and navigate across this TUI. See [Configure Keystone Collector](#) and [Monitor system health](#) for information.

Automatic validation of Keystone software

The Keystone repository is configured to automatically validate the integrity of Keystone software so that only valid and authentic software is installed at your site.

The Keystone YUM repository client configuration provided in `keystonerepo.rpm` makes use of enforced GPG checking (`gpgcheck=1`) on all software downloaded through this repository. Any RPM downloaded through the Keystone repository that fails signature validation is prevented from being installed. This functionality is used in the scheduled auto-update capability of Keystone Collector to ensure only valid and authentic software is installed at your site.

Configure Keystone Collector

You need to complete a few configuration tasks to enable Keystone Collector to collect usage data in your storage environment. This is a one-time activity to activate and associate the required components with your storage environment.

- Keystone Collector provides you with the Keystone Collector Management Terminal User Interface (TUI) utility to perform configuration and monitoring activities. You can use various keyboard controls, such as the Enter and arrow keys, to select the options and navigate across this TUI.
- Keystone Collector can be configured for organizations that do not have internet access, also known as a *dark site* or *private mode*. To learn more about, refer to [Keystone in private mode](#).

Steps

- Start the Keystone Collector management TUI utility:
\$ `keystone-collector-tui`
- Go to **Configure > KS-Collector** to open the Keystone Collector configuration screen to view the available options for update.
- Update the required options.

For ONTAP

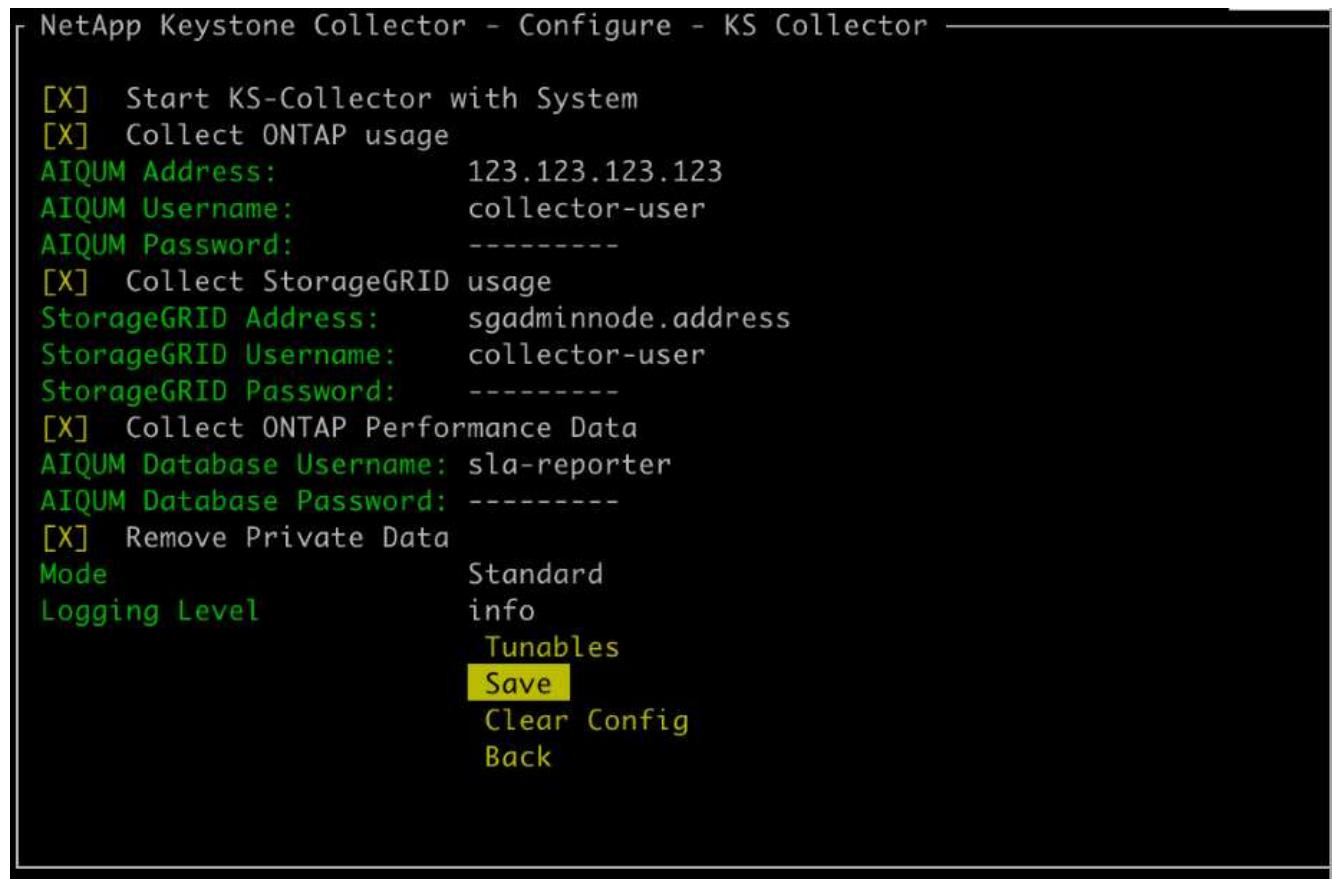
- Collect ONTAP usage:** This option enables collection of usage data for ONTAP. Add the details of the Active IQ Unified Manager (Unified Manager) server and service account.
- Collect ONTAP Performance Data:** This option enables collection of performance data for ONTAP. This is disabled by default. Enable this option if performance monitoring is required in your environment for SLA purposes. Provide the Unified Manager Database user account details. For information about creating database users, see [Create Unified Manager users](#).
- Remove Private Data:** This option removes specific private data of customers and is enabled by default. For information about what data is excluded from the metrics if this option is enabled, see [Limit collection of private data](#).

For StorageGRID

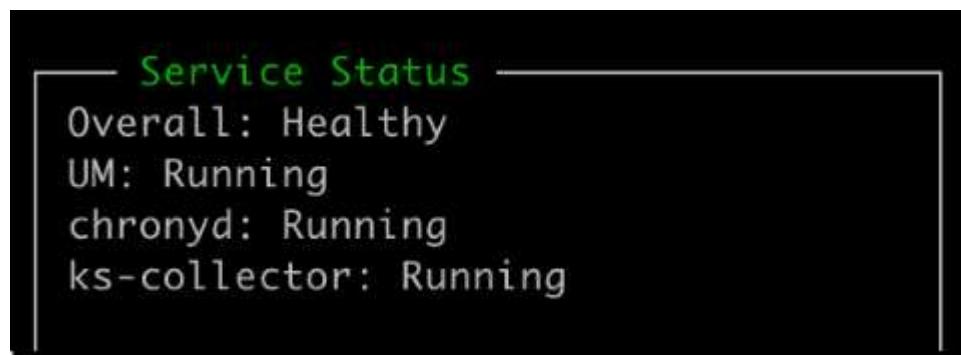
- **Collect StorageGRID usage:** This option enables collection of node usage details. Add the StorageGRID node address and user details.
- **Remove Private Data:** This option removes specific private data of customers and is enabled by default. For information about what data is excluded from the metrics if this option is enabled, see [Limit collection of private data](#).

4. Toggle the **Start KS-Collector with System** field.

5. Click **Save**.



6. Ensure that Keystone Collector is in a healthy state by returning to the main screen of the TUI and verifying the **Service Status** information. The system should show that the services are in an **Overall: Healthy** status.



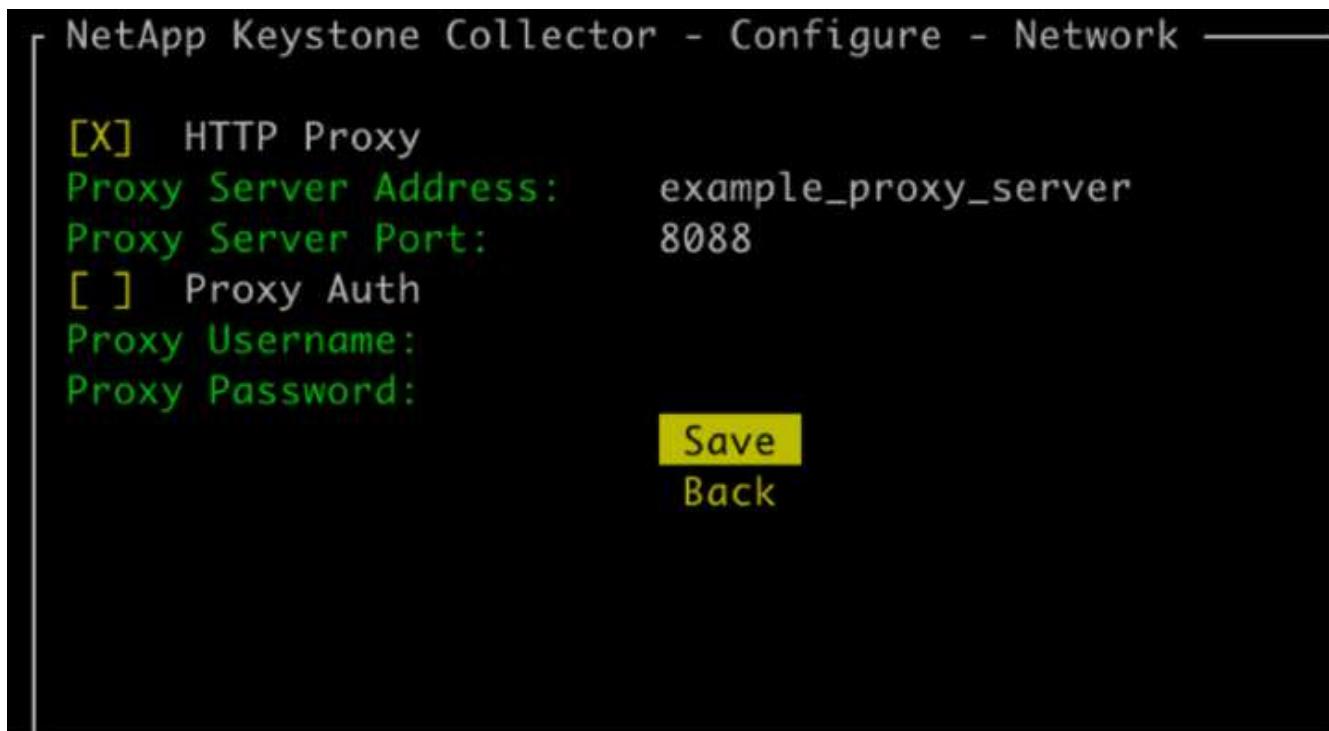
7. Exit the Keystone Collector management TUI by selecting the **Exit to Shell** option on the home screen.

Configure HTTP Proxy on Keystone Collector

The Collector software supports using a HTTP proxy to communicate with the internet. This can be configured in the TUI.

Steps

1. Restart the Keystone Collector management TUI utility if already closed:
\$ keystone-collector-tui
2. Toggle on the **HTTP Proxy** field, and add the details for the HTTP proxy server, port, and credentials, if authentication is required.
3. Click **Save**.



Limit collection of private data

Keystone Collector gathers limited configuration, status, and performance information required to perform subscription metering. There is an option to further limit the information collected by masking sensitive information from the content uploaded. This does not impact billing calculation. However, limiting the information might impact usability of the reporting information, as some elements, which can be easily identified by users, such as volume name, is replaced with UUIDs.

Limiting the collection of specific customer data is a configurable option on the Keystone Collector TUI screen. This option, **Remove Private Data**, is enabled by default.

NetApp Keystone Collector - Configure - KS Collector

```
[X] Start KS-Collector with System
[X] Collect ONTAP usage
AIQUM Address: 123.123.123.123
AIQUM Username: collector
AIQUM Password: -----
[ ] Collect StorageGRID usage
```

```
[ ] Collect ONTAP Performance Data
```

```
[X] Remove Private Data
Mode Standard
Logging Level info
      Tunables
      Save
      Clear Config
      Back
```

For information about the items removed on limiting private data access in both ONTAP and StorageGRID, see [List of items removed on limiting private data access](#).

Trust a custom root CA

Verification of certificates against a public root certificate authority (CA) is a part of the Keystone Collector security features. However, if required, you can configure Keystone Collector to trust a custom root CA.

If you use SSL/TLS inspection in your system firewall, it results in the internet-based traffic to be re-encrypted with your custom CA certificate. It is necessary to configure the settings to verify the source as a trusted CA before accepting the root certificate and allowing connections to occur. Follow these steps:

Steps

1. Prepare the CA certificate. It should be in *base64-encoded X.509* file format.



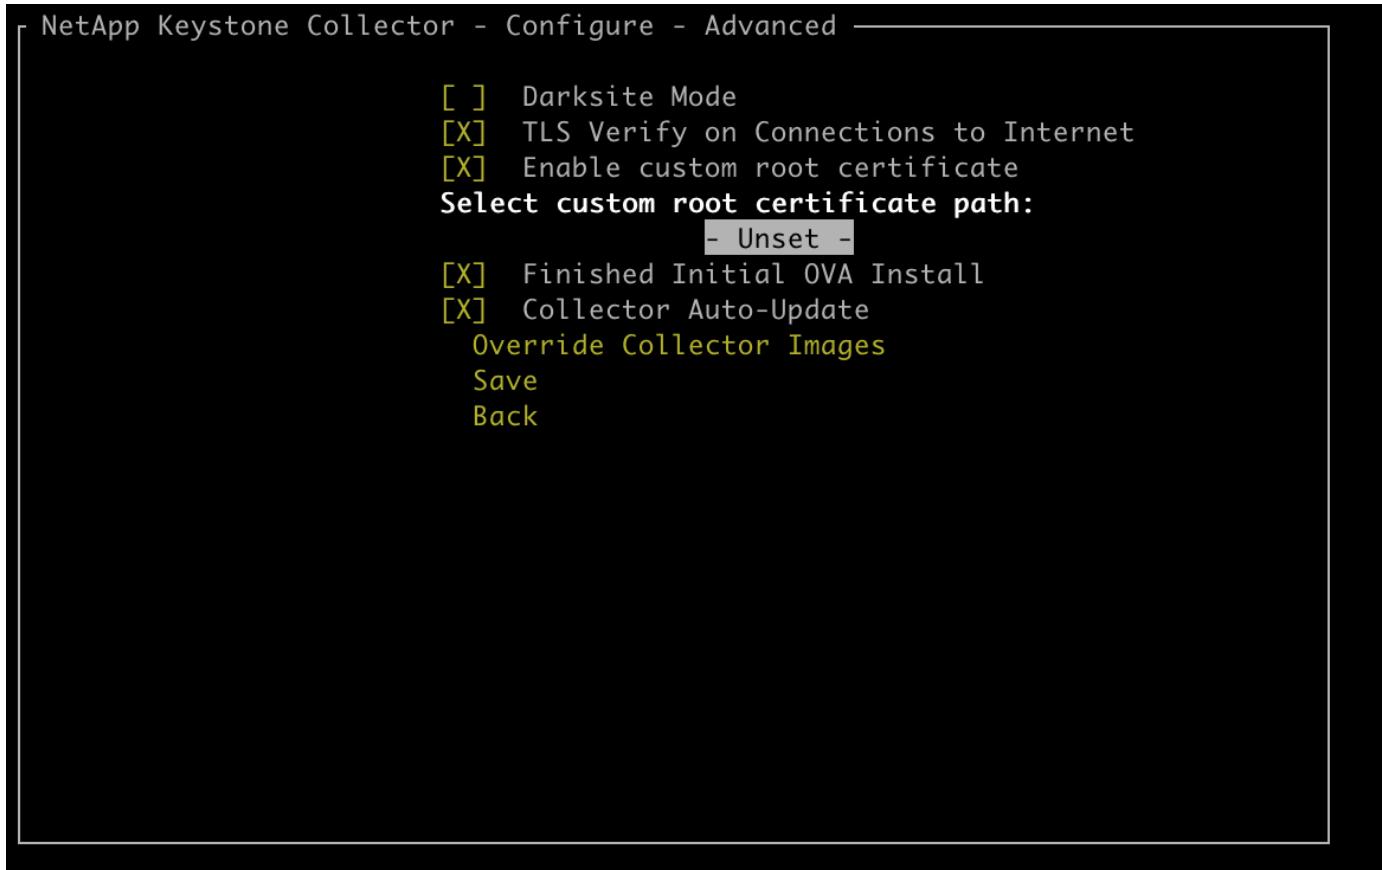
The supported file extensions are `.pem`, `.crt`, `.cert`. Ensure that the certificate is in one of these formats.

2. Copy the certificate to the Keystone Collector server. Make a note of the location where the file is copied.
3. Open a terminal on the server and run the management TUI utility.
 \$ `keystone-collector-tui`
4. Go to **Configuration > Advanced**.
5. Enable the option **Enable custom root certificate**.
6. For **Select custom root certificate path**: select – `Unset` –

7. Press Enter. A dialog box for selecting the certificate path is displayed.
8. Select the root certificate from the file system browser or enter the exact path.
9. Press Enter. You return to the **Advanced** screen.
10. Select **Save**. The configuration is applied.



The CA certificate is copied to `/opt/netapp/ks-collector/ca.pem` on the Keystone Collector server.



Create Performance Service Levels

You can create Performance Service Levels (PSLs) using the Keystone Collector management TUI utility. Creating PSLs through the TUI automatically selects the default values set for each performance service level, reducing the chance of errors that might occur when manually setting these values while creating PSLs through Active IQ Unified Manager.

To learn more about PSLs, refer to [Performance Service Levels](#).

To learn more about service levels, refer to [Service levels in Keystone](#).

Steps

1. Start the Keystone Collector management TUI utility:
`$ keystone-collector-tui`
2. Go to **Configure>AIQUM** to open the AIQUM screen.
3. Enable the option **Create AIQUM Performance Profiles**.

4. Enter the details of the Active IQ Unified Manager server and user account. These details are required to create PSLs and will not be stored.

NetApp Keystone Collector – Configure – AIQUM

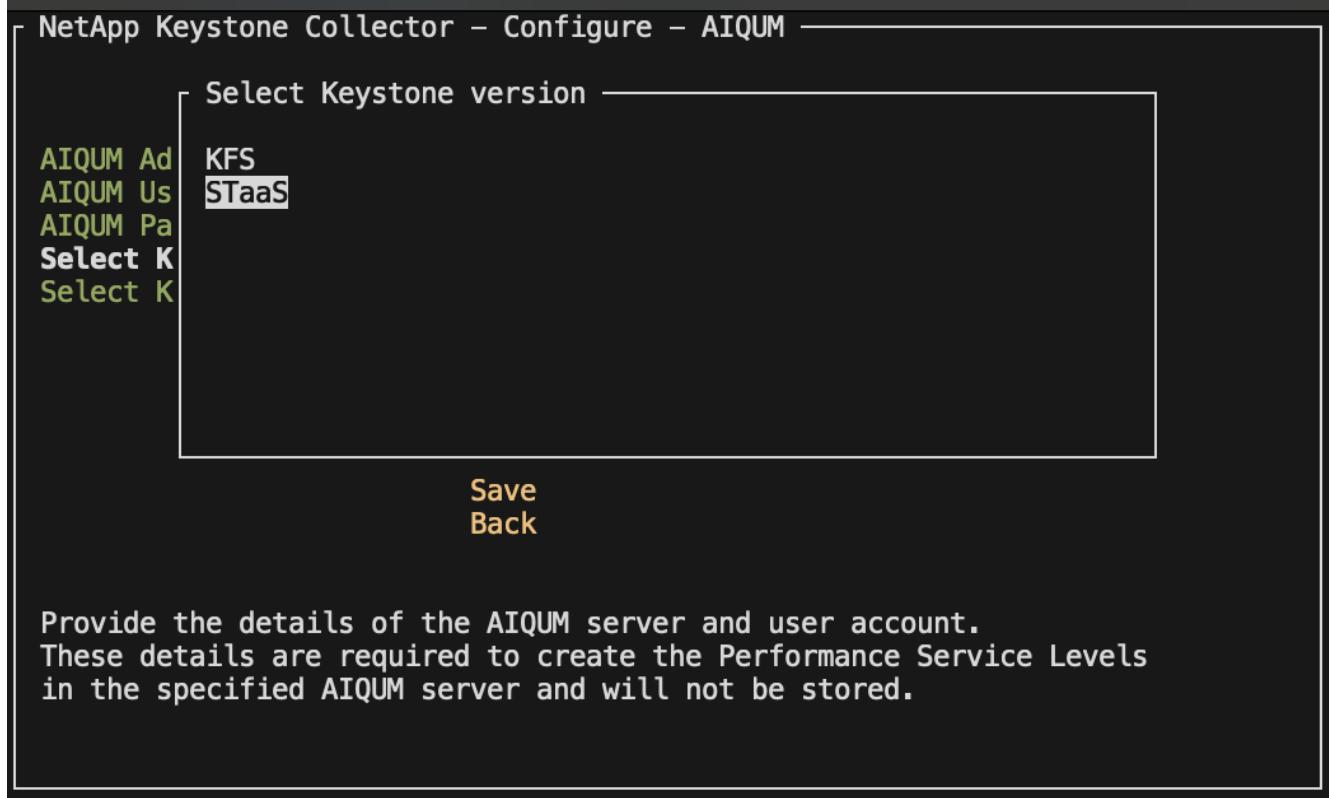
[] Enable Embedded UM
[X] Create AIQUM Performance Profiles

AIQUM Address:
AIQUM Username:
AIQUM Password:
Select Keystone version -unset-
Select Keystone Service Levels

Save
Back

Provide the details of the AIQUM server and user account.
These details are required to create the Performance Service Levels
in the specified AIQUM server and will not be stored.

5. For **Select Keystone version**, select `-unset-`.
6. Press Enter. A dialog box for selecting the Keystone version is displayed.
7. Highlight **STaaS** to specify the Keystone version for Keystone STaaS, and then press Enter.



You can highlight the **KFS** option for Keystone subscription services version 1. Keystone subscription services differ from Keystone STaaS in the constituent performance service levels, service offerings, and billing principles. To learn more, refer to [Keystone subscription services | Version 1](#).

8. All supported Keystone performance service levels will be displayed within the **Select Keystone Service Levels** option for the specified Keystone version. Enable the desired performance service levels from the list.

NetApp Keystone Collector – Configure – AIQUM

[] Enable Embedded UM
[X] Create AIQUM Performance Profiles

AIQUM Address:
AIQUM Username:
AIQUM Password: -----

Select Keystone version
Select Keystone Service Levels

STaaS
[X] Extreme
[X] Premium
[] Performance
[] Standard
[] Value

Save
Back

Provide the details of the AIQUM server and user account.
These details are required to create the Performance Service Levels in the specified AIQUM server and will not be stored.



You can select multiple performance service levels simultaneously to create PSLs.

9. Select **Save** and press Enter. Performance Service Levels will be created.

You can view the created PSLs, such as Premium-KS-STaaS for STaaS or Extreme KFS for KFS, on the **Performance Service Levels** page in Active IQ Unified Manager. If the created PSLs do not meet your requirements, then you can modify PSLs to meet your needs. To learn more, refer to [Creating and editing Performance Service Levels](#).

Performance Service Levels

View and manage the Performance Service Levels that you can assign to workloads.

Search Performance Service Levels  Filter 

 Add  Modify  Remove 

	Name	Type	Expected IOPS/TB	Peak IOPS/TB	Absolute Minim...	Expected Latency	Capacity	Workloads
 	Extreme - KFS	User-defined	6144	12288	1000	1	 Used: 0 bytes Available: 283.85 TiB	0
 	Extreme - KS-STaaS	User-defined	6144	12288	1000	1	 Used: 0 bytes Available: 283.85 TiB	0
Overview Description: Extreme - KS-STaaS Added Date: 1 Aug 2024, 18:08 Last Modified Date: 1 Aug 2024, 18:08								
 	Premium ...S-STaaS	User-defined	2048	4096	500	2	 Used: 0 bytes Available: 283.85 TiB	0
Overview Description: Premium - KS-STaaS Added Date: 1 Aug 2024, 18:08 Last Modified Date: 1 Aug 2024, 18:08								

If a PSL for the selected performance service level already exist on the specified Active IQ Unified Manager server, then you cannot create it again. If you attempt to do so, you will receive an error message.

NetApp Keystone Collector – Configure – AIQUM

Warning

AIQUM Ad
AIQUM Us
AIQUM Pa
Select K
Select K

Failed to create Performance Service Level for:
Extreme. Error: <Response [400]>

OK

> Save <
Back

Provide the details of the AIQUM server and user account.
These details are required to create the Performance Service Levels
in the specified AIQUM server and will not be stored.

Install ITOM Collector

Installation requirements for Keystone ITOM Collector

Before installing ITOM Collector, ensure that your systems are prepared with the necessary software and meet all required prerequisites.

Prerequisites for the ITOM Collector server VM:

- Supported operating system:
 - Debian 12 or later
 - Windows Server 2016 or later
 - Ubuntu 20.04 LTS or later
 - Red Hat Enterprise Linux (RHEL) 8.x
 - Red Hat Enterprise Linux 9.0 or later
 - Amazon Linux 2023 or later



The recommended operating systems are Debian 12, Windows Server 2016, or newer versions.

- Resource requirement: The VM resource requirements based on the number of NetApp nodes monitored are as follows:
 - 2-10 nodes: 4 CPUs, 8 GB RAM, 40 GB Disk
 - 12-20 nodes: 8 CPUs, 16 GB RAM, 40 GB Disk
- Configuration requirement: Ensure that a read-only account and SNMP are configured on the monitored devices. The ITOM Collector server VM also needs to be configured as an SNMP trap host and Syslog server on the NetApp cluster and cluster switches, if applicable.

Networking requirements

The networking requirements of ITOM Collector are listed in the following table.

Source	Destination	Protocol	Ports	Description
ITOM Collector	NetApp ONTAP cluster management IPs	HTTPS, SNMP	TCP 443, UDP 161	Monitoring of the ONTAP controllers
NetApp ONTAP cluster and node management IPs	ITOM Collector	SNMP, Syslog	UDP 162, UDP 514	SNMP traps and Syslogs from controllers
ITOM Collector	Cluster switches	SNMP	UDP 161	Monitoring of switches
Cluster switches	ITOM Collector	SNMP, Syslog	UDP 162, UDP 514	SNMP traps and Syslogs from switches
ITOM Collector	StorageGRID nodes IPs	HTTPS, SNMP	TCP 443, UDP 161	SNMP monitoring of StorageGRID
StorageGRID nodes IPs	ITOM Collector	SNMP, Syslog	UDP 162, UDP 514	SNMP traps from StorageGRID

ITOM Collector	Keystone Collector	SSH, HTTPS, SNMP	TCP 22, TCP 443, UDP 161	Keystone Collector monitoring and remote management
ITOM Collector	Local DNS	DNS	UDP 53	Public or private DNS services
ITOM Collector	NTP server(s) of choice	NTP	UDP 123	Time keeping

Install Keystone ITOM Collector on Linux systems

Complete a few steps to install ITOM Collector, which collects metrics data in your storage environment. You can install it on either Windows or Linux systems, depending on your requirements.



Keystone support team provides a dynamic link to download the ITOM Collector setup file, which expires in two hours.

To install ITOM Collector on Windows systems, refer to [Install ITOM Collector on Windows systems](#).

Follow these steps to install software on your Linux server:

Before you begin

- Verify that the Bourne shell is available for the Linux installation script.
- Install the `vim-common` package to get the `xxd` binary required for the ITOM Collector setup file.
- Ensure the `sudo` package is installed if planning to run ITOM Collector as a non-root user.

Steps

- Download the ITOM collector setup file to your Linux server.
- Open a terminal on the server and run the following command to change the permissions and make the binaries executable:
`# chmod +x <installer_file_name>.bin`
- Run the command to start the ITOM collector setup file:
`./<installer_file_name>.bin`
- Running the setup file prompts you to:
 - Accept the end-user license agreement (EULA).
 - Enter the user details for the installation.
 - Specify the installation parent directory.
 - Select the collector size.
 - Provide proxy details, if applicable.

For each prompt, a default option is displayed. It is recommended to select the default option unless you have specific requirements. Press the **Enter** key to choose the default option. When the installation completes, a message confirms that the ITOM Collector is installed successfully.

- The ITOM Collector setup file makes additions to `/etc/sudoers` to handle service restarts and memory dumps.
- Installing ITOM Collector on the Linux server creates a default user called **ITOM** to run ITOM Collector without root privileges. You can choose a different user or run it as root, but it is recommended to use the ITOM user created by the Linux installation script.

What's next?

On successful installation, contact the Keystone support team to validate the successful installation of ITOM Collector through the ITOM support portal. After verification, the Keystone support team will configure the ITOM Collector remotely, including further device discovery and monitoring setup, and will send a confirmation once the configuration is complete. For any queries or additional information, contact keystone.services@netapp.com.

Install Keystone ITOM Collector on Windows systems

Install ITOM Collector on a Windows system by downloading the ITOM Collector setup file, running the InstallShield wizard, and entering the required monitoring credentials.

 Keystone support team provides a dynamic link to download the ITOM Collector setup file, which expires in two hours.

You can install it on Linux systems based on your requirements. To install ITOM Collector on Linux systems, refer to [Install ITOM Collector on Linux systems](#).

Follow these steps to install ITOM collector software on your Windows server:

Before you begin

Ensure ITOM Collector service is granted **Log on as a service** under Local Policy/User Rights Assignment in the Windows server's local security policy settings.

Steps

1. Download the ITOM collector setup file to your Windows server.
2. Open the setup file to start the InstallShield wizard.
3. Accept the end-user license agreement (EULA). The InstallShield wizard extracts the necessary binaries and prompts you to enter credentials.
4. Enter the credentials for the account that ITOM Collector will run under:
 - If ITOM Collector is not monitoring other Windows servers, use local system.
 - If ITOM Collector is monitoring other Windows servers in the same domain, use a domain account with local administrator permissions.
 - If ITOM Collector is monitoring other Windows servers that are not part of the same domain, use a local administrator account and connect to each resource with local administrator credentials. You may choose to set the password so that it does not expire, to reduce authentication issues between ITOM Collector and its monitored resources.
5. Select the collector size. The default is the recommended size based on the setup file. Proceed with the suggested size unless you have specific requirements.
6. Select *Next* to begin the installation. You can use the populated folder or choose a different one. A status box displays the installation progress, followed by the InstallShield Wizard Completed dialog box.

What's next?

On successful installation, contact the Keystone support team to validate the successful installation of ITOM Collector through the ITOM support portal. After verification, the Keystone support team will configure the ITOM Collector remotely, including further device discovery and monitoring setup, and will send a confirmation once the configuration is complete. For any queries or additional information, contact keystone.services@netapp.com.

Configure AutoSupport for Keystone

When using the AutoSupport telemetry mechanism, Keystone calculates the usage based on the AutoSupport telemetry data. To achieve the necessary level of granularity, you should configure AutoSupport to incorporate Keystone data in the daily support bundles sent by the ONTAP clusters.

About this task

You should note the following before configuring AutoSupport to include Keystone data.

- You edit the AutoSupport telemetry options by using ONTAP CLI. For information about managing AutoSupport services and system (cluster) administrator role, see [Manage AutoSupport overview](#) and [Cluster and SVM administrators](#).
- You include the subsystems in the daily and weekly AutoSupport bundles to ensure precise data collection for Keystone. For information about AutoSupport subsystems, see [What AutoSupport subsystems are](#).

Steps

1. As a system administrator user, log in to the Keystone ONTAP cluster by using SSH. For information, see [Access the cluster by using SSH](#).
2. Modify the log content.
 - For ONTAP 9.16.1 and above, run this command to modify the daily log content:

```
autosupport trigger modify -node * -autosupport-message
management.log -basic-additional
wafl,performance,snapshot,object_store_server,san,raid,snapmirror
-troubleshooting-additional wafl
```

If the cluster is in a MetroCluster configuration, run this command:

```
autosupport trigger modify -node * -autosupport-message
management.log -basic-additional
wafl,performance,snapshot,object_store_server,san,raid,snapmirror,met
rocluster -troubleshooting-additional wafl
```

- For earlier ONTAP versions, run this command to modify the daily log content:

```
autosupport trigger modify -node * -autosupport-message
management.log -basic-additional
wafl,performance,snapshot,platform,object_store_server,san,raid,snapm
irror -troubleshooting-additional wafl
```

If the cluster is in a MetroCluster configuration, run this command:

```
autosupport trigger modify -node * -autosupport-message
management.log -basic-additional
wafl,performance,snapshot,platform,object_store_server,san,raid,snapm
irror,metrocluster -troubleshooting-additional wafl
```

- Run this command to modify the weekly log content:

```
autosupport trigger modify -autosupport-message weekly
-troubleshooting-additional wafl -node *
```

For more information about this command, see [system node autosupport trigger modify](#).

Monitor and upgrade

Monitor the health of Keystone Collector

You can monitor the health of Keystone Collector by using any monitoring system that supports HTTP requests. Monitoring the health can help to ensure that data is available on the Keystone dashboard.

By default, Keystone health services do not accept connections from any IP other than localhost. The Keystone health endpoint is `/uber/health`, and it listens on all interfaces of the Keystone Collector server on port 7777. On query, an HTTP request status code with a JSON output is returned from the endpoint as a response, describing the status of the Keystone Collector system.

The JSON body provides an overall health status for the `is_healthy` attribute, which is a boolean; and a detailed list of statuses per-component for the `component_details` attribute.

Here is an example:

```
$ curl http://127.0.0.1:7777/uber/health
{"is_healthy": true, "component_details": {"vicmet": "Running", "ks-
collector": "Running", "ks-billing": "Running", "chrony": "Running"}}
```

These status codes are returned:

- **200**: indicates that all monitored components are healthy
- **503**: indicates that one or more components are unhealthy

- **403**: indicates that the HTTP client querying the health status is not on the *allow* list, which is a list of allowed network CIDRs. For this status, no health information is returned.

The *allow* list uses the network CIDR method to control which network devices are allowed to query the Keystone health system. If you receive this error, add your monitoring system to the *allow* list from **Keystone Collector management TUI > Configure > Health Monitoring**.

Linux users, note this known issue:

Issue description: Keystone Collector runs a number of containers as part of the usage metering system. When the Red Hat Enterprise Linux 8.x server is hardened with USA Defense Information Systems Agency (DISA) Security Technical Implementation Guides (STIG) policies, a known issue with fapolicyd (File Access Policy Daemon) has been seen intermittently. This issue is identified as [bug 1907870](#).

Workaround: Until resolved by Red Hat Enterprise, NetApp recommends that you work around this issue by putting fapolicyd into permissive mode. In `/etc/fapolicyd/fapolicyd.conf`, set the value of `permissive = 1`.



View system logs

You can view Keystone Collector system logs to review system information and perform troubleshooting by using those logs. Keystone Collector uses the host's *journald* logging system, and the system logs can be reviewed through the standard *journalctl* system utility. You can avail the following key services to examine the logs:

- `ks-collector`
- `ks-health`
- `ks-autoupdate`

The main data collection service *ks-collector* produces logs in JSON format with a `run-id` attribute associated with each scheduled data collection job. The following is an example of a successful job for standard usage data collection:

```

{"level":"info","time":"2022-10-31T05:20:01.831Z","caller":"light-
collector/main.go:31","msg":"initialising light collector with run-id
cdflm0f74cgphgfon8cg","run-id":"cdflm0f74cgphgfon8cg"}

{"level":"info","time":"2022-10-
31T05:20:04.624Z","caller":"ontap/service.go:215","msg":"223 volumes
collected for cluster a2049dd4-bfcf-11ec-8500-00505695ce60","run-
id":"cdflm0f74cgphgfon8cg"}


{"level":"info","time":"2022-10-
31T05:20:18.821Z","caller":"ontap/service.go:215","msg":"697 volumes
collected for cluster 909cbacc-bfcf-11ec-8500-00505695ce60","run-
id":"cdflm0f74cgphgfon8cg"}


{"level":"info","time":"2022-10-
31T05:20:41.598Z","caller":"ontap/service.go:215","msg":"7 volumes
collected for cluster f7b9a30c-55dc-11ed-9c88-005056b3d66f","run-
id":"cdflm0f74cgphgfon8cg"}


{"level":"info","time":"2022-10-
31T05:20:48.247Z","caller":"ontap/service.go:215","msg":"24 volumes
collected for cluster a9e2dcff-ab21-11ec-8428-00a098ad3ba2","run-
id":"cdflm0f74cgphgfon8cg"}


{"level":"info","time":"2022-10-
31T05:20:48.786Z","caller":"worker/collector.go:75","msg":"4 clusters
collected","run-id":"cdflm0f74cgphgfon8cg"}


{"level":"info","time":"2022-10-
31T05:20:48.839Z","caller":"reception/reception.go:75","msg":"Sending file
65a71542-cb4d-bdb2-e9a7-a826be4fdcb7_1667193648.tar.gz type=ontap to
reception","run-id":"cdflm0f74cgphgfon8cg"}


{"level":"info","time":"2022-10-
31T05:20:48.840Z","caller":"reception/reception.go:76","msg":"File bytes
123425","run-id":"cdflm0f74cgphgfon8cg"}


{"level":"info","time":"2022-10-
31T05:20:51.324Z","caller":"reception/reception.go:99","msg":"uploaded
usage file to reception with status 201 Created","run-
id":"cdflm0f74cgphgfon8cg"}

```

The following is an example of a successful job for optional performance data collection:

```
{"level":"info","time":"2022-10-  
31T05:20:51.324Z","caller":"sql/service.go:28","msg":"initialising MySql  
service at 10.128.114.214"}  
  
{"level":"info","time":"2022-10-  
31T05:20:51.324Z","caller":"sql/service.go:55","msg":"Opening MySql db  
connection at server 10.128.114.214"}  
  
{"level":"info","time":"2022-10-  
31T05:20:51.324Z","caller":"sql/service.go:39","msg":"Creating MySql db  
config object"}  
  
{"level":"info","time":"2022-10-  
31T05:20:51.324Z","caller":"sla_reporting/service.go:69","msg":"initialising  
SLA service"}  
  
{"level":"info","time":"2022-10-  
31T05:20:51.324Z","caller":"sla_reporting/service.go:71","msg":"SLA  
service successfully initialised"}  
  
{"level":"info","time":"2022-10-  
31T05:20:51.324Z","caller":"worker/collector.go:217","msg":"Performance  
data would be collected for timerange: 2022-10-31T10:24:52~2022-10-  
31T10:29:52"}  
  
{"level":"info","time":"2022-10-  
31T05:21:31.385Z","caller":"worker/collector.go:244","msg":"New file  
generated: 65a71542-cb4d-bdb2-e9a7-a826be4fdcb7_1667193651.tar.gz"}  
  
{"level":"info","time":"2022-10-  
31T05:21:31.385Z","caller":"reception/reception.go:75","msg":"Sending file  
65a71542-cb4d-bdb2-e9a7-a826be4fdcb7_1667193651.tar.gz type=ontap-perf to  
reception","run-id":"cdf0f74cgphgfon8cg"}  
  
{"level":"info","time":"2022-10-  
31T05:21:31.386Z","caller":"reception/reception.go:76","msg":"File bytes  
17767","run-id":"cdf0f74cgphgfon8cg"}  
  
{"level":"info","time":"2022-10-  
31T05:21:33.025Z","caller":"reception/reception.go:99","msg":"uploaded  
usage file to reception with status 201 Created","run-  
id":"cdf0f74cgphgfon8cg"}  
  
{"level":"info","time":"2022-10-31T05:21:33.025Z","caller":"light-  
collector/main.go:88","msg":"exiting","run-id":"cdf0f74cgphgfon8cg"}
```

Generate and collect support bundles

The Keystone Collector TUI enables you to generate support bundles and add them to service requests for resolving support issues. Follow this procedure:

Steps

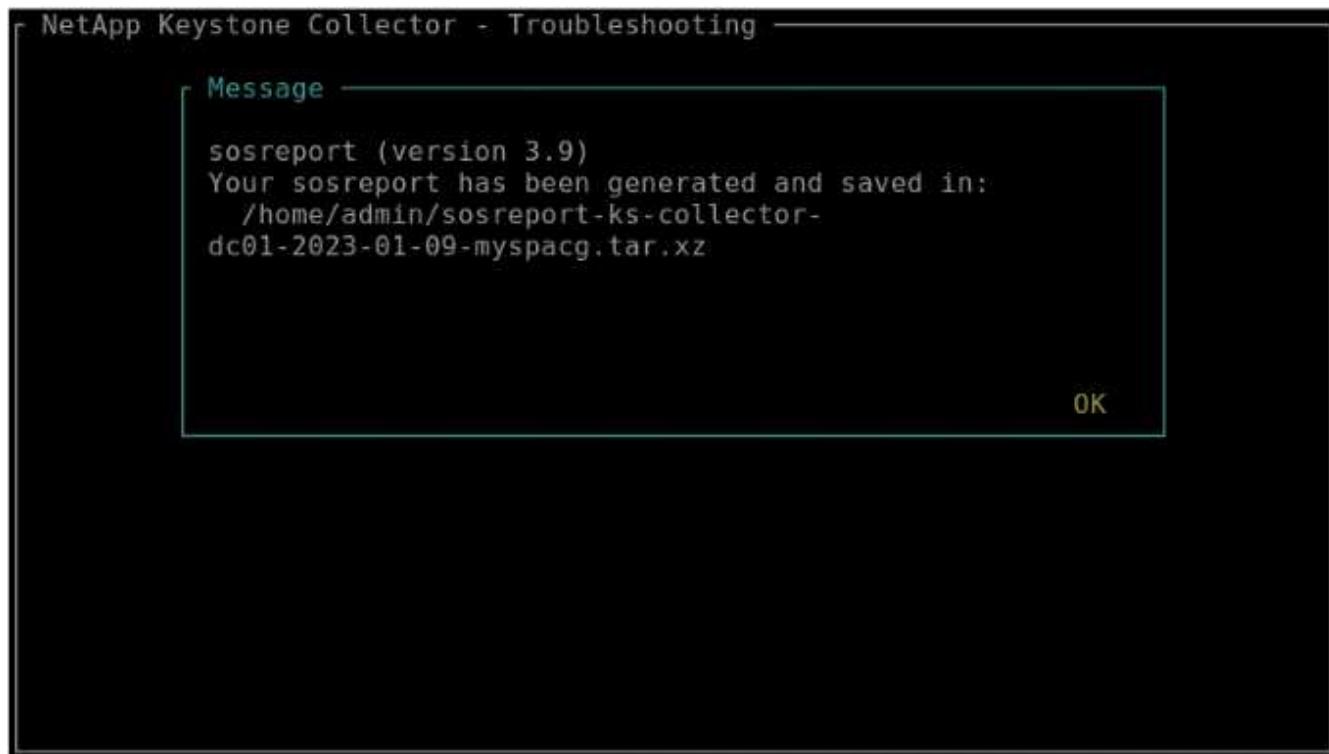
1. Start the Keystone Collector management TUI utility:

```
$ keystone-collector-tui
```

2. Go to **Troubleshooting > Generate Support Bundle**.



3. When generated, the location where the bundle is saved is displayed. Use FTP, SFTP, or SCP to connect to the location and download the log file to a local system.



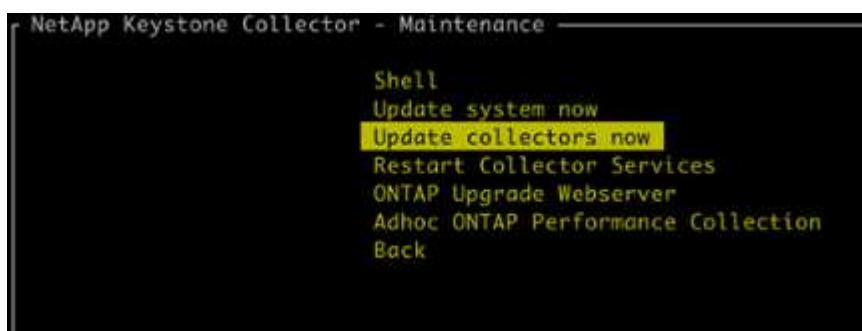
4. When the file is downloaded, you can attach it to the Keystone ServiceNow support ticket. For information about raising tickets, see [Generating service requests](#).

Manually upgrade Keystone Collector

The auto-update feature in Keystone Collector is enabled by default, which automatically upgrades the Keystone Collector software with every new release. You can, however, disable this feature and manually upgrade the software.

Steps

1. Start the Keystone Collector management TUI utility:
\$ keystone-collector-tui
2. On the maintenance screen, selecting the **Update collectors now** option.



Alternately, run these commands to upgrade the version:

For CentOS:

```
sudo yum clean metadata && sudo yum install keystone-collector
```

```
[admin@rhel8-serge-dev ~]$ sudo yum clean metadata && sudo yum install keystone-collector
Updating Subscription Management repositories.
Unable to read consumer identity

This system is not registered with an entitlement server. You can use subscription-manager to register.

Cache was expired
0 files removed
Updating Subscription Management repositories.
Unable to read consumer identity

This system is not registered with an entitlement server. You can use subscription-manager to register.

Netapp Keystone
Red Hat Enterprise Linux 8 - BaseOS
Red Hat Enterprise Linux 8 - AppStream
Package keystone-collector-1.3.0-1.noarch is already installed.
Dependencies resolved.

=====
Package           Architecture      Version       Repository      Size
=====
Upgrading:
  keystone-collector      noarch        1.3.2-1      keystone       411 M

Transaction Summary
=====
Upgrade 1 Package

Total download size: 411 M
Is this ok [y/N]: y
Downloading Packages:
keystone-collector-1.3.2-1.noarch.rpm
=====
8.3 MB/s | 411 MB  00:49
Total
Running transaction check
Transaction check succeeded.
Running transaction test
Transaction test succeeded.
Running transaction
  Preparing           : 1/1
  Running scriptlet: keystone-collector-1.3.2-1.noarch 1/1
  Running scriptlet: keystone-collector-1.3.2-1.noarch 1/2
  Upgrading          : keystone-collector-1.3.2-1.noarch 1/2
  Running scriptlet: keystone-collector-1.3.2-1.noarch 1/2
=====
*   Keystone Collector package installation complete!  *
*   Run command 'keystone-collector-tui' to configure . *
*                                                       *
=====
Running scriptlet: keystone-collector-1.3.0-1.noarch 2/2
Cleanup           : keystone-collector-1.3.0-1.noarch 2/2
Running scriptlet: keystone-collector-1.3.0-1.noarch 2/2
Verifying          : keystone-collector-1.3.2-1.noarch 1/2
Verifying          : keystone-collector-1.3.0-1.noarch 1/2
Installed products updated.

Upgraded:
  keystone-collector-1.3.2-1.noarch

Complete!
[admin@rhel8-serge-dev ~]$ rpm -q keystone-collector
keystone-collector-1.3.2-1.noarch
```

For Debian:

```
sudo apt-get update && sudo apt-get upgrade keystone-collector
```

3. Restart Keystone Collector management TUI, you can see the latest version on the upper left portion of the home screen.

Alternately, run these commands to view the latest version:

For CentOS:

```
rpm -q keystone-collector
```

For Debian:

```
dpkg -l | grep keystone-collector
```

Keystone Collector security

Keystone Collector includes security features that monitor the performance and usage metrics of Keystone systems, without risking the security of customer data.

The functioning of Keystone Collector is based on the following security principles:

- **Privacy by design**-Keystone Collector collects minimum data to perform usage metering and performance monitoring. For more information, see [Data collected for billing](#). The [Remove Private Data](#) option is enabled by default, which masks and protects sensitive information.
- **Least privilege access**-Keystone Collector requires minimum permissions to monitor the storage systems, which minimizes security risks and prevents any unintended modifications to the data. This approach aligns with the principle of least privilege, enhancing the overall security posture of the monitored environments.
- **Secure software development framework**- Keystone uses a secure software development framework throughout the development cycle, which mitigates risks, reduces vulnerabilities, and protects the system against potential threats.

Security hardening

By default, Keystone Collector is configured to use security-hardened configurations. The following are the recommended security configurations:

- The operating system of the Keystone Collector virtual machine:
 - Complies with the CIS Debian Linux 12 Benchmark standard. Making any changes to the OS configuration outside the Keystone Collector management software may reduce the system security. For more information, see [CIS Benchmark guide](#).
 - Automatically receives and installs security patches that are verified by Keystone Collector through the auto-update feature. Disabling this functionality may lead to unpatched vulnerable software.
 - Authenticates updates received from Keystone Collector. Disabling APT repository verification can lead to the automatic installation of unauthorized patches, potentially introducing vulnerabilities.
- Keystone Collector automatically validates HTTPS certificates to ensure connection security. Disabling this feature could lead to impersonation of external endpoints and usage data leakage.
- Keystone Collector supports [Custom Trusted CA](#) certification. By default, it trusts certificates that are signed by public root CAs recognized by the [Mozilla CA Certificate program](#). By enabling additional Trusted CAs, Keystone Collector enables HTTPS certificate validation for connections to endpoints that present these certificates.
- Keystone collector enables the **Remove Private Data** option by default, which masks and protects sensitive information. For more information, see [Limit collection of private data](#). Disabling this option results in additional data being communicated to the Keystone system. For example, it can include user-entered information such as volume names which may be considered sensitive information.

Related information

- [Keystone Collector overview](#)

- [Virtual infrastructure requirements](#)
- [Configure Keystone Collector](#)

Types of user data that Keystone collects

Keystone collects configuration, status, and usage information from Keystone ONTAP and Keystone StorageGRID subscriptions, as well as telemetry data from the virtual machine (VM) hosting Keystone Collector. It can collect performance data for ONTAP only, if this option is enabled in Keystone Collector.

ONTAP data collection

Usage data collected for ONTAP: Learn more

The following list is a representative sample of the capacity consumption data collected for ONTAP:

- Clusters
 - ClusterUUID
 - ClusterName
 - SerialNumber
 - Location (based on value input in ONTAP cluster)
 - Contact
 - Version
- Nodes
 - SerialNumber
 - Node name
- Volumes
 - Aggregate name
 - Volume Name
 - VolumeInstanceUUID
 - IsCloneVolume flag
 - IsFlexGroupConstituent flag
 - IsSpaceEnforcementLogical flag
 - IsSpaceReportingLogical flag
 - LogicalSpaceUsedByAfs
 - PercentSnapshotSpace
 - PerformanceTierInactiveUserData
 - PerformanceTierInactiveUserDataPercent
 - QoSAdaptivePolicyGroup Name
 - QoSPolicyGroup Name
 - Size
 - Used
 - PhysicalUsed
 - SizeUsedBySnapshots
 - Type
 - VolumeStyleExtended
 - Vserver name
 - IsVsRoot flag
- VServers
 - VserverName

- VserverUUID
- Subtype
- Storage aggregates
 - StorageType
 - Aggregate Name
 - Aggregate UUID
 - Physical Used
 - Available Size
 - Size
 - Used Size
- Aggregate object stores
 - ObjectStoreName
 - ObjectStoreUUID
 - ProviderType
 - Aggregate Name
- Clone volumes
 - FlexClone
 - Size
 - Used
 - Vserver
 - Type
 - ParentVolume
 - ParentVserver
 - IsConstituent
 - SplitEstimate
 - State
 - FlexCloneUsedPercent
- Storage LUNs
 - LUN UUID
 - LUN Name
 - Size
 - Used
 - IsReserved flag
 - IsRequested flag
 - LogicalUnit Name
 - QoS Policy UUID
 - QoS Policy Name

- VolumeUUID
- VolumeName
- SVMUUID
- SVM Name
- Storage volumes
 - VolumeInstanceUUID
 - VolumeName
 - SVMName
 - SVMUUID
 - QoS PolicyUUID
 - QoS PolicyName
 - CapacityTierFootprint
 - PerformanceTierFootprint
 - TotalFootprint
 - TieringPolicy
 - IsProtected flag
 - IsDestination flag
 - Used
 - PhysicalUsed
 - CloneParentUUID
 - LogicalSpaceUsedByAfs
- QoS policy groups
 - PolicyGroup
 - QoS PolicyUUID
 - MaxThroughput
 - MinThroughput
 - MaxThroughputIOPS
 - MaxThroughputMBps
 - MinThroughputIOPS
 - MinThroughputMBps
 - IsShared flag
- ONTAP adaptive QoS policy groups
 - QoS PolicyName
 - QoS PolicyUUID
 - PeakIOPS
 - PeakIOPSAllocation
 - AbsoluteMinIOPS

- ExpectedIOPS
- ExpectedIOPSAccOUNT
- BlockSize
- Footprints
 - Vserver
 - Volume
 - TotalFootprint
 - VolumeBlocksFootprintBin0
 - VolumeBlocksFootprintBin1
- MetroCluster
 - Node
 - Aggregate
 - LIFs
 - Config Replication
 - Connections
 - Clusters
 - Volumes
- MetroCluster clusters
 - ClusterUUID
 - ClusterName
 - RemoteClusterUUID
 - RemoteClusterName
 - LocalConfigurationState
 - RemoteConfigurationState
- MetroCluster nodes
 - DR Mirroring State
 - Intercluster LIF
 - Node reachability
 - DR Partner node
 - DR Aux Partner node
 - DR, DR Aux, and HA node symmetric relationship
 - Automatic Unplanned Switchover
- MetroCluster Config Replication
 - Remote Heartbeat
 - Last Heartbeat Sent
 - Last Heartbeat Received
 - Vserver Stream

- Cluster Stream
- Storage
- Storage in use volume
- MetroCluster mediators
 - Mediator Address
 - Mediator Port
 - Mediator Configured
 - Mediator Reachable
 - Mode
- Collector Observability Metrics
 - Collection Time
 - Active IQ Unified Manager API endpoint queried
 - Response time
 - Number of records
 - AIQUMInstance IP
 - CollectorInstance ID

Performance data collected for ONTAP: Learn more

The following list is a representative sample of the performance data collected for ONTAP:

- Cluster Name
- Cluster UUID
- ObjectID
- VolumeName
- Volume Instance UUID
- Vserver
- VserverUUID
- Node Serial
- ONTAPVersion
- AIQUM version
- Aggregate
- AggregateUUID
- ResourceKey
- TimeStamp
- IOPSPerTb
- Latency
- ReadLatency
- WriteMBps
- QoSMinThroughputLatency
- QoSNBlaeLatency
- UsedHeadRoom
- CacheMissRatio
- OtherLatency
- QoSAggregateLatency
- IOPS
- QoSNetworkLatency
- AvailableOps
- WriteLatency
- QoSCloudLatency
- QoSClusterInterconnectLatency
- OtherMBps
- QoSCopLatency
- QoSDBlaeLatency
- Utilization

- ReadIOPS
- MBps
- OtherIOPS
- QoSPolicyGroupLatency
- ReadMBps
- QoSSyncSnapmirrorLatency
- System level data
 - Write/Read/Other/Total IOPS
 - Write/Read/Other/Total Throughput
 - Write/Read/Other/Total Latency
- WriteIOPS

List of items removed on limiting private data access: [Learn more](#)

When the **Remove Private Data** option is enabled on Keystone Collector, the following usage information is eliminated for ONTAP. This option is enabled by default.

- Cluster Name
- Cluster Location
- Cluster Contact
- Node Name
- Aggregate name
- Volume Name
- QoSAdaptivePolicyGroup Name
- QoSPolicyGroup Name
- Vserver name
- Storage LUN name
- Aggregate Name
- LogicalUnit Name
- SVM Name
- AIQUMInstance IP
- FlexClone
- RemoteClusterName

StorageGRID data collection

Usage data collected for StorageGRID: [Learn more](#)

The following list is a representative sample of the Logical Data collected for StorageGRID:

- StorageGRID ID
- Account ID
- Account Name
- Account Quota Bytes
- Bucket Name
- Bucket Object Count
- Bucket Data Bytes

The following list is a representative sample of the Physical Data collected for StorageGRID:

- StorageGRID ID
- Node ID
- Site ID
- Site Name
- Instance
- StorageGRID storage utilization Bytes
- StorageGRID storage utilization metadata Bytes

The following list is a representative sample of the Availability/Uptime Data collected for StorageGRID:

- SLA Uptime percentage

List of items removed on limiting private data access: [Learn more](#)

When the **Remove Private Data** option is enabled on Keystone Collector, the following usage information is eliminated for StorageGRID. This option is enabled by default.

- AccountName
- BucketName
- SiteName
- Instance/NodeName

Telemetry data collection

Telemetry data collected from Keystone Collector VM: [Learn more](#)

The following list is a representative sample of the telemetry data collected for Keystone systems:

- System information
 - Operating system name
 - Operating system version
 - Operating system ID
 - System hostname
 - System default IP address
- System resource usage
 - System uptime
 - CPU core count
 - System load (1 min, 5 min, 15 min)
 - Total memory
 - Free memory
 - Available memory
 - Shared memory
 - Buffer memory
 - Cached memory
 - Total swap
 - Free swap
 - Cached swap
 - Disk filesystem name
 - Disk size
 - Disk used
 - Disk available
 - Disk usage percentage
 - Disk mount point
- Installed packages
- Collector configuration
- Service logs
 - Service logs from Keystone services

Keystone in private mode

[Learn about Keystone \(private mode\)](#)

Keystone offers a *private* deployment mode, also known as a *dark site*, to meet your

business and security requirements. This mode is available for organizations with connectivity restrictions.

NetApp offers a specialized deployment of Keystone STaaS tailored for environments with limited or no internet connectivity (also known as dark sites). These are secure or isolated environments where external communication is restricted due to security, compliance, or operational requirements.

For NetApp Keystone, offering services for dark sites means providing the Keystone flexible storage subscription service in a way that respects the constraints of these environments. This involves:

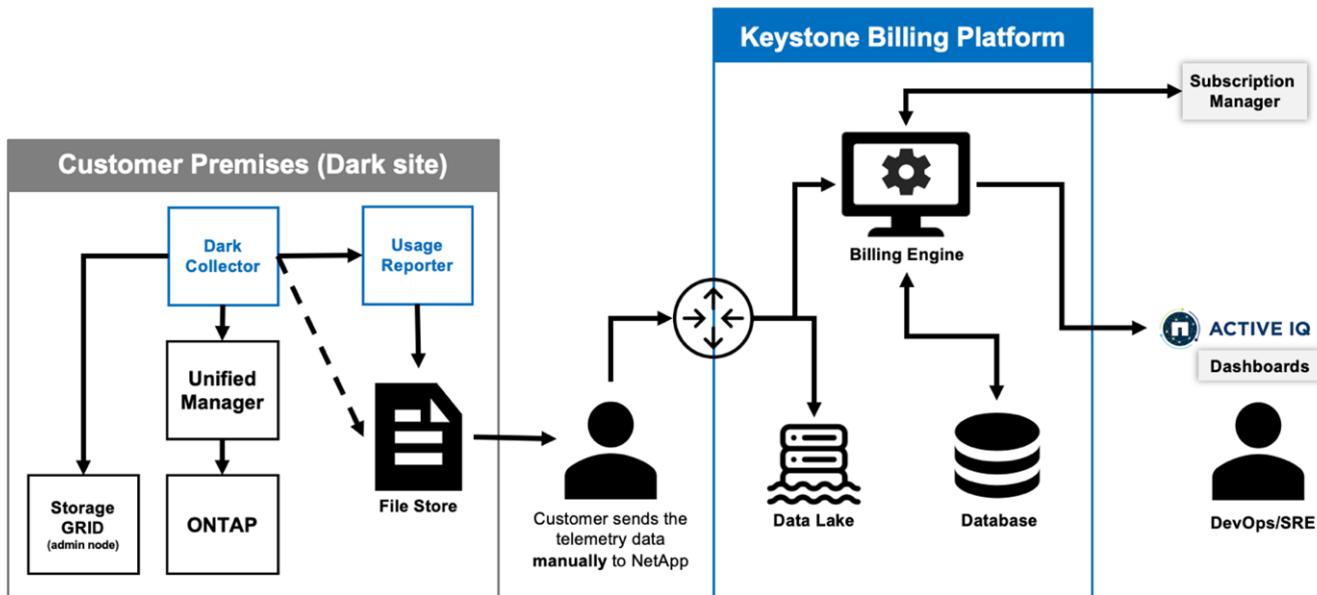
- **Local deployment:** Keystone can be configured within isolated environments independently, ensuring no need for internet connectivity or external personnel for setup access.
- **Offline operations:** All storage management capabilities with health checks and billing are available offline for operations.
- **Security and compliance:** Keystone ensures that the deployment meets the security and compliance requirements of dark sites, which may include advanced encryption, secure access controls, and detailed auditing capabilities.
- **Help and Support:** NetApp provides 24/7 global support with a dedicated Keystone success manager assigned to each account for assistance and troubleshooting.



Keystone Collector can be configured without connectivity restrictions, also known as *standard* mode. To learn more, refer to [Learn about Keystone Collector](#).

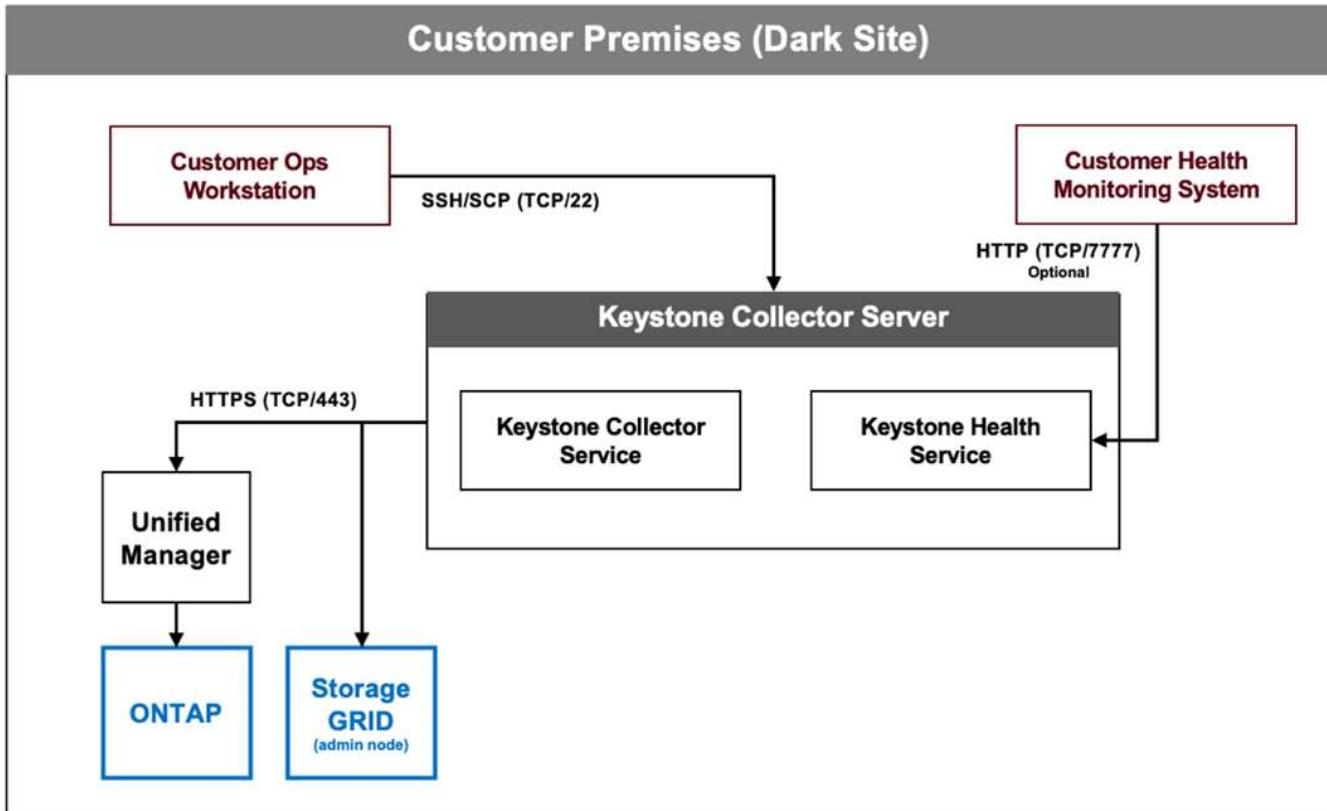
Keystone Collector in private mode

Keystone Collector is responsible for periodically collecting usage data from storage systems and exporting the metrics to an offline usage reporter and a local file store. The generated files, which are created in both encrypted and plain text formats, are then manually forwarded to NetApp by the user after the validation checks. Upon receipt, NetApp's Keystone billing platform authenticates and processes these files, integrating them into the billing and subscription management systems to calculate the monthly charges.



The Keystone Collector service on the server is tasked with periodically gathering usage data, processing this information, and generating a usage file locally on the server. The health service conducts system health

checks and is designed to interface with health monitoring systems used by the customer. These reports are available for offline access by users, allowing for validation and aiding in troubleshooting issues.



Prepare for Keystone Collector installation in private mode

Before installing Keystone Collector in an environment without internet access, also known as a *dark site* or *private mode*, ensure your systems are prepared with the necessary software and meet all required prerequisites.

Requirements for VMware vSphere

- Operating system: VMware vCenter server and ESXi 8.0 or later
- Core: 1 CPU
- RAM: 2 GB
- Disk space: 20 GB vDisk

Requirements for Linux

- Operating system (choose one):
 - Red Hat Enterprise Linux (RHEL) 8.6 or any later 8.x series
 - Red Hat Enterprise Linux 9.0 or later versions
 - Debian 12
- Core: 2 CPU
- RAM: 4 GB

- Disk space: 50 GB vDisk
 - At least 2 GB free in `/var/lib/`
 - At least 48 GB free in `/opt/netapp`

The same server should also have the following third-party packages installed. If available through the repository, these packages will be automatically installed as prerequisites:

- RHEL 8.6+ (8.x)
 - `python3 >=v3.6.8, python3 <=v3.9.13`
 - `podman`
 - `sos`
 - `yum-utils`
 - `python3-dnf-plugin-versionlock`
- RHEL 9.0+
 - `python3 >= v3.9.0, python3 <= v3.9.13`
 - `podman`
 - `sos`
 - `yum-utils`
 - `python3-dnf-plugin-versionlock`
- Debian v12
 - `python3 >= v3.9.0, python3 <= v3.12.0`
 - `podman`
 - `sosreport`

Networking requirements

The networking requirements for Keystone Collector include the following:

- Active IQ Unified Manager (Unified Manager) 9.10 or later, configured on a sever with the API Gateway functionality enabled.
- The Unified Manager server should be accessible by the Keystone Collector server on port 443 (HTTPS).
- A service account with Application User permissions should be set up for the Keystone Collector on the Unified Manager server.
- External internet connectivity is not required.
- Each month, export a file from Keystone Collector and email it to the NetApp support team. For more information on how to contact the support team, refer to [Get help with Keystone](#).

Install Keystone Collector in private mode

Complete a few steps to install Keystone Collector in an environment that does not have internet access, also known as a *dark site* or *private mode*. This type of installation is perfect for your secure sites.

You can either deploy Keystone Collector on VMware vSphere systems or install it on Linux systems,

depending on your requirements. Follow the installation steps that correspond to your selected option.

Deploy on VMware vSphere

Follow these steps:

1. Download the OVA template file from [NetApp Keystone web portal](#).
2. For steps to deploy Keystone collector with OVA file, refer to the section [Deploying the OVA template](#).

Install on Linux

Keystone Collector software is installed on the Linux server using the provided .deb or .rpm files, based on the Linux distribution.

Follow these steps to install the software on your Linux server:

1. Download or transfer the Keystone Collector installation file to the Linux server:

```
keystone-collector-<version>.noarch.rpm
```

2. Open a terminal on the server and run the following commands to begin the installation.

- **Using Debian package**

```
dpkg -i keystone-collector_<version>_all.deb
```

- **Using RPM file**

```
yum install keystone-collector-<version>.noarch.rpm
```

or

```
rpm -i keystone-collector-<version>.noarch.rpm
```

3. Enter **y** when prompted to install the package.

Configure Keystone Collector in private mode

Complete a few configuration tasks to enable Keystone Collector to collect usage data in an environment that does not have internet access, also known as a *dark site* or *private mode*. This is a one-time activity to activate and associate the required components with your storage environment. Once configured, Keystone Collector will monitor all ONTAP clusters managed by Active IQ Unified Manager.



Keystone Collector provides you with the Keystone Collector Management Terminal User Interface (TUI) utility to perform configuration and monitoring activities. You can use various keyboard controls, such as the Enter and arrow keys, to select the options and navigate across this TUI.

Steps

1. Start the Keystone Collector management TUI utility:

keystone-collector-tui

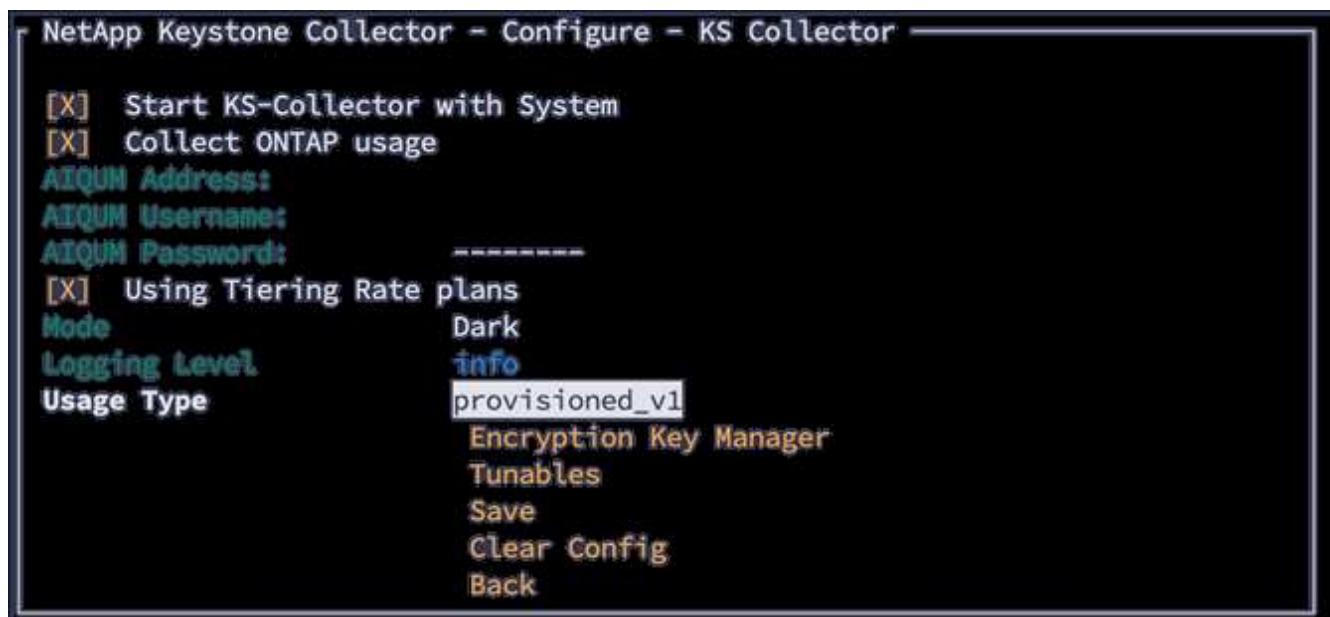
2. Go to **Configure > Advanced**.
3. Toggle the **Darksite Mode** option.



4. Select **Save**.
5. Go to **Configure > KS-Collector** to configure Keystone Collector.
6. Toggle the **Start KS Collector with System** field.
7. Toggle the **Collect ONTAP Usage** field. Add the details of the Active IQ Unified Manager (Unified Manager) server and user account.
8. **Optional:** Toggle the **Using Tiering Rate plans** field if data tiering is required for the subscription.
9. Based on the subscription type purchased, update the **Usage Type**.

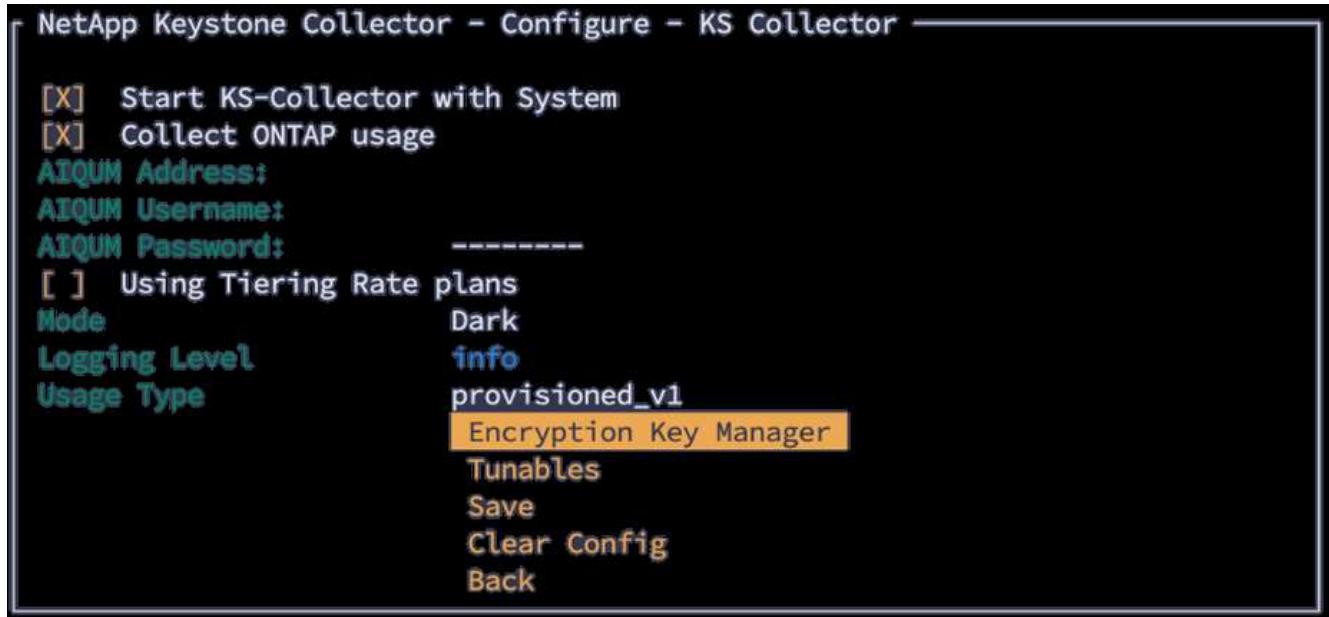


Before configuring, confirm the usage type associated with the subscription from NetApp.



10. Select **Save**.

11. Go to **Configure > KS-Collector** to generate the Keystone Collector keypair.
12. Go to **Encryption Key Manager** and press Enter.



13. Select **Generate Collector Keypair** and press Enter.



14. Ensure that the Keystone Collector is in a healthy state by returning to the main screen of the TUI and verifying the **Service Status** information. The system should show that the services are in an **Overall: Healthy** status. Wait up to 10 minutes, if the overall status remains unhealthy after this period, review the previous configuration steps and contact the NetApp support team.

Service Status

```
Overall: Healthy
UM-Dark: Running
ks-billing: Running
ks-collector-dark: Running
Recent collector data: Healthy
ONTAP REST response time: Healthy
DB Disk space: Healthy
DB Disk space 30d: Healthy
DB API responses: Healthy
DB Concurrent flushes: Healthy
DB Slow insert rate: Healthy
```

15. Exit the Keystone Collector management TUI by selecting **Exit to Shell** option on the home screen.

16. Retrieve the generated public key:

```
~/collector-public.pem
```

17. Send an email with this file to ng-keystone-secure-site-upload@netapp.com for secure non-USPS sites, or to ng-keystone-secure-site-usps-upload@netapp.com for secure USPS sites.

Export usage report

You should send the monthly usage summary report to NetApp at the end of every month. You can generate this report manually.

Follow these steps to generate the usage report:

1. Go to to **Export Usage** on the Keystone Collector TUI home screen.
2. Collect the files and send them to ng-keystone-secure-site-upload@netapp.com for secure non-USPS sites, or to ng-keystone-secure-site-usps-upload@netapp.com for secure USPS sites.

Keystone Collector generates both a clear file and an encrypted file, which should be manually sent to NetApp. The clear file report contains the following details that can be validated by the customer.

```
node_serial,derived_service_level,usage_tib,start,duration_seconds
123456781,extreme,25.0,2024-05-27T00:00:00,86400
123456782,premium,10.0,2024-05-27T00:00:00,86400
123456783,standard,15.0,2024-05-27T00:00:00,86400

<Signature>
31b3d8eb338ee319ef1

-----BEGIN PUBLIC KEY-----
31b3d8eb338ee319ef1
-----END PUBLIC KEY-----
```

Upgrade ONTAP

Keystone Collector supports ONTAP upgrades through TUI.

Follow these steps to upgrade ONTAP:

1. Go to **Maintenance > ONTAP Upgrade Webserver**.
2. Copy the ONTAP upgrade image file to `/opt/netapp/ontap-upgrade/`, then select **Start Webserver** to start the web server.



3. Go to <http://<collector-ip>:8000> using a web browser for upgrade assistance.

Restart Keystone Collector

You can restart the Keystone Collector service through the TUI. Go to **Maintenance > Restart Collector Services** in the TUI. This will reboot all collector services, and their status can be monitored from the TUI home screen.



Monitor Keystone Collector health in private mode

You can monitor the health of Keystone Collector by using any monitoring system that supports HTTP requests.

By default, Keystone health services do not accept connections from any IP other than localhost. The Keystone health endpoint is `/uber/health`, and it listens on all interfaces of the Keystone Collector server on port 7777. On query, an HTTP request status code with a JSON output is returned from the endpoint as a response, describing the status of the Keystone Collector system.

The JSON body provides an overall health status for the `is_healthy` attribute, which is a boolean; and a detailed list of statuses per-component for the `component_details` attribute.

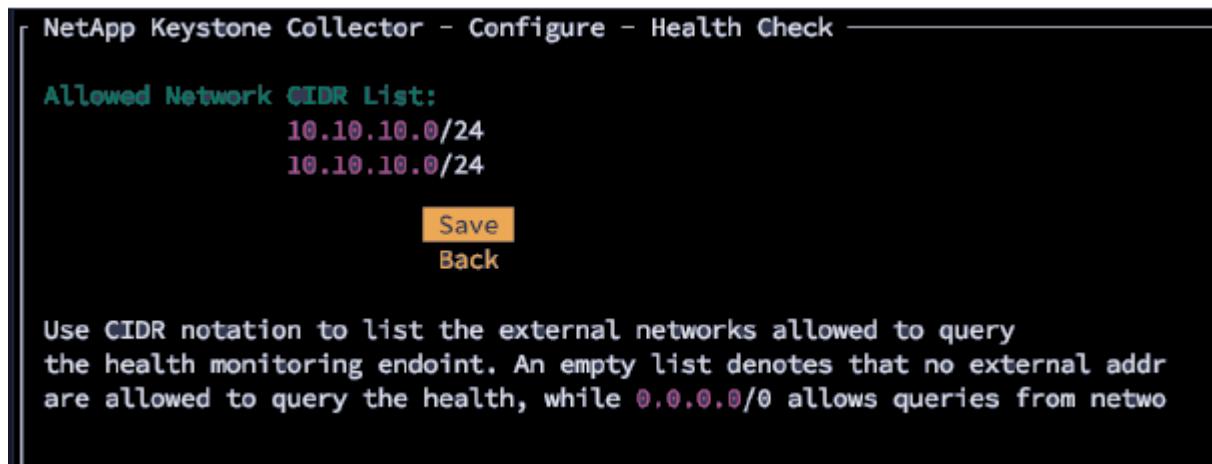
Here is an example:

```
$ curl http://127.0.0.1:7777/uber/health
{"is_healthy": true, "component_details": {"vicmet": "Running", "ks-collector": "Running", "ks-billing": "Running", "chronyd": "Running"}}
```

These status codes are returned:

- **200**: indicates that all monitored components are healthy
- **503**: indicates that one or more components are unhealthy
- **403**: indicates that the HTTP client querying the health status is not on the `allow` list, which is a list of allowed network CIDRs. For this status, no health information is returned.

The `allow` list uses the network CIDR method to control which network devices are allowed to query the Keystone health system. If you receive the 403 error, add your monitoring system to the `allow` list from **Keystone Collector management TUI > Configure > Health Monitoring**.

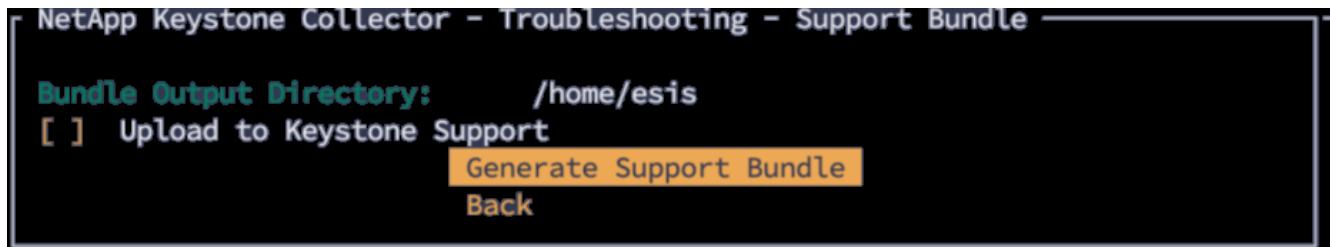


Generate and collect support bundles

To troubleshoot issues with Keystone Collector, you can work with NetApp Support who might ask for a *.tar* file. You can generate this file through the Keystone Collector management TUI utility.

Follow these steps to generate a *.tar* file:

1. Go to **Troubleshooting > Generate Support Bundle**.
2. Select the location to save the bundle, then click **Generate Support Bundle**.



This process creates a *tar* package at the mentioned location which can be shared with NetApp for troubleshooting issues.

3. When the file is downloaded, you can attach it to the Keystone ServiceNow support ticket. For information about raising tickets, see [Generating service requests](#).

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.