



Set up and configure Keystone

Keystone

NetApp
March 22, 2023

Table of Contents

- Set up and configure Keystone 1
- Keystone Collector overview 1
- Prerequisites for installing Keystone Collector 2
- Installation procedure 9
- Configure Keystone Collector 12
- Monitor system health 15
- Types of user data that Keystone collects 20

Set up and configure Keystone

Keystone Collector overview


For availing your Keystone services and viewing your usage data, you should install Keystone Collector on a VMware vSphere or Linux system at your site.



The installation and configuration information available here is applicable for ONTAP and StorageGRID. The requirements and procedures are generic to both ONTAP and StorageGRID, with exceptions pointed out.

Keystone Collector is the usage acquisition component of the Keystone billing platform that leverages Active IQ Unified Manager and other applications to connect to ONTAP and StorageGRID systems to collect metadata required for usage and SLA performance metering of your Keystone subscriptions. It provides you with the ability to monitor system health, while sending your billing data for reporting.

This architecture diagram outlines the constituent components and their connectivity in a typical Keystone environment.



Prerequisites for installing Keystone Collector

Virtual infrastructure requirements

There are a few virtual infrastructure configurations that are required for installing Keystone Collector on your VMware vSphere system.

The Keystone Collector server virtual machine should have the following:

- Operating system: VMware vCentre server and ESXi 6.5 or later
- Core: 4 CPUs
- RAM: 12 GB RAM
- Disk space: 200 GB vDisk

Other requirements

Endure that the following generic requirements are met:

Networking Requirements

The following table describes the networking requirements of Keystone Collector.



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)	HTTP, HTTPS	TCP 80, 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	HTTP	TCP 777	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

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
https://mirror.centos.org https://mirrorlist.centos.org	OS updates for Keystone Collector (OVA deployments only)
http://repo.mysql.com	Download and update mysql dependencies (for Linux deployments only)

Linux system requirements

There are a few configurations that are required for installing Keystone Collector on Linux systems.

The following configurations are required for the Linux server:

- Operating system: CentOS 7 or Red Hat Enterprise Linux 8.6 or later
- Chronyd time synchronized

- Access to the standard Linux standard software repositories

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

- podman (POD Manager)
- sos
- chrony
- python 3 (3.6.8 to 3.9.13)

The following configurations are required for the Keystone Collector server virtual machine (SVM):

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

Other requirements

Endure that the following generic requirements are met:

Networking Requirements

The following table describes the networking requirements of Keystone Collector.



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)	HTTP, HTTPS	TCP 80, 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	HTTP	TCP 777	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

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
https://mirror.centos.org https://mirrorlist.centos.org	OS updates for Keystone Collector (OVA deployments only)

<http://repo.mysql.com>

Download and update mysql dependencies (for Linux deployments only)

Additional prerequisites for ONTAP and StorageGRID

You should complete a few additional prerequisites for ONTAP and StorageGRID. Ensure that you have completed these specific prerequisites in addition to the Linux/VMware vSphere system requirements. Click the required tab to learn more.

ONTAP

For installing Keystone Collector in ONTAP, perform these steps as prerequisites:

1. Ensure that an Active IQ Unified Manager (Unified Manager) server is configured with Unified Manager 9.7 or later. 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](#).
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 11.4.0 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 users](#).

Installation procedure

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 <https://keystone.netapp.com/downloads/KeystoneCollector-latest.ova> 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. The product name should appear as *NetApp Keystone Collector*.

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	Entrust Code Signing CA - OVCS2 (Trusted certificate)
Product	NetApp Keystone Collector
Version	20220405
Vendor	NetApp
Download size	8.3 GB
Size on disk	12.1 GB (thin provisioned)
	200.0 GB (thick provisioned)

CANCEL
BACK
NEXT

vCentre validates the integrity of the OVA contents and that a valid code-signing digest is provided for the files contained in the OVA. However, it does not validate the authenticity of the code-signing certificate. For verifying the integrity, you should download the full signing digest certificate, and verify it against the public certificate published by Keystone.

- a. Click the **Publisher** link to download the full signing digest certificate.
- b. Download the *Keystone Billing* public certificate: <https://keystone.netapp.com/downloads/OVA-SSL-NetApp-Keystone-20221101.pem>.
- c. Verify the authenticity of the OVA signing certificate against the public certificate by using OpenSSL:


```
openssl verify -CAfile OVA-SSL-NetApp-Keystone-20221101.pem keystone-collector.cert
```

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 utilize.
10. On *Step 7 Customize template*, specify the initial network address and password for the admin user account.



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 Keystone Collector servers deployed through OVA:



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. On login, the TUI appears. You can also launch the TUI manually by running the `keystone-collector-tui` CLI command.
2. If required, configure the proxy details in the **Configuration > Network** section on the TUI.
3. Update the system/collector by using the **Maintenance > Update System** option. Some mirrors that are selected might be unavailable, and the system details are updated after a few retries.
4. Configure the system hostname, location, and NTP server in the **Configuration > System** section.
5. Update the admin password in the **Maintenance > User** section.
6. Mark the initial OVA configuration as complete in the **Configuration > Advanced** section.

Installing Keystone Collector on Linux systems

The Keystone Collector software is distributed by an online YUM software repository. You need to import and install the file on a Linux server.

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

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/repo/RPM-GPG-NetApp-Keystone-20221101
```
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:
90B3 83AF E07B 658A 6058 5B4E 76C2 45E4 33B6 C17D
4. Download the `keystonerepo.rpm` file:

```
curl -O https://keystone.netapp.com/repo/keystonerepo.rpm
```
5. Verify the authenticity of the file:

```
rpm --checksig -v keystonerepo.rpm
```

A signature for an authentic file looks like this:
Header V4 RSA/SHA512 Signature, key ID 33b6c17d: OK
6. Install the YUM software repository file:

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

```
# yum 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](#).

Automatic validation of software integrity

There is a reiterative process of validating the integrity of the Keystone software.

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 the 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 collector component with your storage environment.



The Keystone Collector provides you with 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.

Steps

1. Open the Keystone Collector management TUI utility:

```
$ keystone-collector-tui
```
2. Go to **Configure > KS-Collector** to open the Keystone Collector configuration screen to view the available options for update.
3. 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**.

```
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-user
AIQUM Password:        -----
[X] Collect StorageGRID usage
StorageGRID Address:   sgadminnode.address
StorageGRID Username:  collector-user
StorageGRID Password:  -----
[X] Collect ONTAP Performance Data
AIQUM Database Username: sla-reporter
AIQUM Database Password: -----
[X] Remove Private Data
Mode                   Standard
Logging Level         info
                      Tunables
                      Save
                      Clear Config
                      Back
```

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.

```
Service Status
Overall: Healthy
UM: Running
chronyd: Running
ks-collector: Running
```


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. Open 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

The 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](#).

Monitor system health

You can monitor your system health through Keystone Collector services by using any monitoring system that supports HTTP requests.

By default, the 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 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 cdf1m0f74cgphgfon8cg", "run-id": "cdf1m0f74cgphgfon8cg" }
{ "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": "cdf1m0f74cgphgfon8cg" }

{ "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": "cdf1m0f74cgphgfon8cg" }

{ "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": "cdf1m0f74cgphgfon8cg" }

{ "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": "cdf1m0f74cgphgfon8cg" }

{ "level": "info", "time": "2022-10-31T05:20:48.786Z", "caller": "worker/collector.go:75", "msg": "4 clusters collected", "run-id": "cdf1m0f74cgphgfon8cg" }

{ "level": "info", "time": "2022-10-31T05:20:48.839Z", "caller": "reception/reception.go:75", "msg": "Sending file 65a71542-cb4d-bdb2-e9a7-a826be4fdb7_1667193648.tar.gz type=ontap to reception", "run-id": "cdf1m0f74cgphgfon8cg" }

{ "level": "info", "time": "2022-10-31T05:20:48.840Z", "caller": "reception/reception.go:76", "msg": "File bytes 123425", "run-id": "cdf1m0f74cgphgfon8cg" }

{ "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": "cdf1m0f74cgphgfon8cg" }
```

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-a826be4fdb7_1667193651.tar.gz"}

{"level":"info","time":"2022-10-31T05:21:31.385Z","caller":"reception/reception.go:75","msg":"Sending file 65a71542-cb4d-bdb2-e9a7-a826be4fdb7_1667193651.tar.gz type=ontap-perf to reception","run-id":"cdf1m0f74cgphgfon8cg"}

{"level":"info","time":"2022-10-31T05:21:31.386Z","caller":"reception/reception.go:76","msg":"File bytes 17767","run-id":"cdf1m0f74cgphgfon8cg"}

{"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":"cdf1m0f74cgphgfon8cg"}

{"level":"info","time":"2022-10-31T05:21:33.025Z","caller":"light-collector/main.go:88","msg":"exiting","run-id":"cdf1m0f74cgphgfon8cg"}
```

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. Open 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](#).

Types of user data that Keystone collects

Keystone collects configuration, status, and usage information for your Keystone ONTAP and Keystone StoragGRID subscriptions. It can also collect performance data for only ONTAP, if the option is enabled in Keystone Collector.

ONTAP data collection

Usage data collected for ONTAP

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
- 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
 - QoSPolicyUUID
 - QoSPolicyName
 - VolumeUUID
 - VolumeName
 - SVMUUID
 - SVM Name

- Storage volumes
 - VolumeInstanceUUID
 - VolumeName
 - SVMName
 - SVMUUID
 - QoSPolicyUUID
 - QoSPolicyName
 - CapacityTierFootprint
 - PerformanceTierFootprint
 - TotalFootprint
 - TieringPolicy
 - IsProtected flag
 - IsDestination flag
 - Used
 - PhysicalUsed
 - CloneParentUUID
 - LogicalSpaceUsedByAfs
- QoS policy groups
 - PolicyGroup
 - QoSPolicyUUID
 - MaxThroughput
 - MinThroughput
 - MaxThroughputIOPS
 - MaxThroughputMBps
 - MinThroughputIOPS
 - MinThroughputMBps
 - IsShared flag
- ONTAP adaptive QoS policy groups
 - QoSPolicyName
 - QoSPolicyUUID
 - PeakIOPS
 - PeakIOPSAllocation
 - AbsoluteMinIOPS
 - ExpectedIOPS
 - ExpectedIOPSAllocation
 - BlockSize
- Footprints

- Vserver
- Volume
- TotalFootprint
- VolumeBlocksFootprintBin0
- VolumeBlocksFootprintBin1
- MetroCluster clusters
 - ClusterUUID
 - ClusterName
 - RemoteClusterUUID
 - RemoteClusterName
 - LocalConfigurationState
 - RemoteConfigurationState
 - 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

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
- QoSNBladeLatency
- UsedHeadRoom
- CacheMissRatio
- OtherLatency
- QoSAggregateLatency
- IOPS
- QoSNetworkLatency
- AvailableOps
- WriteLatency
- QoSCLatency
- QoSClusterInterconnectLatency
- OtherMBps
- QoSCopLatency
- QoSDBladeLatency
- Utilization

- ReadIOPS
- MBps
- OtherIOPS
- QoSPolicyGroupLatency
- ReadMBps
- QoSSyncSnapmirrorLatency
- WriteIOPS

List of items removed on limiting private data access

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

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

List of items removed on limiting private data access

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

Copyright information

Copyright © 2023 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—without 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.