

Cloud Insights documentation

Cloud Insights

NetApp April 16, 2024

This PDF was generated from https://docs.netapp.com/us-en/cloudinsights/index.html on April 16, 2024. Always check docs.netapp.com for the latest.

Table of Contents

Cloud Insights documentation	
What can Cloud Insights do for me?	1
Getting Started	1
What's New with Cloud Insights	2
April 2024	2
March 2024	
February 2024	4
January 2024	7
December 2023	8
November 2023	10
October 2023	11
September 2023.	
August 2023	15
July 2023	18
June 2023	
May 2023	
April 2023	22
March 2023	25
January 2023	26
December 2022	26
November 2022	28
October 2022	28
September 2022	
August 2022	30
June 2022	
May 2022	38
April 2022	40
March 2022	42
February 2022	43
December 2021	44
November 2021	
October 2021	47
September 2021	49
August 2021	50
June 2021	51
May 2021	54
April 2021	55
February 2021	58
January 2021	59
December 2020	61
November 2020	62
October 2020	63
September 2020	63

August 2020	65
July 2020	65
June 2020	73
May 2020	74
April 2020	77
February 2020	79
January 2020	80
December 2019	82
November 2019	82
October 2019	83
September 2019	83
August 2019	85
July 2019	85
June 2019	86
May 2019	86
April 2019	87
March 2019	87
February 2019	88
January 2019	88
December 2018	88
November 2018	89
Cloud Insights Onboarding.	90
Creating your NetApp BlueXP account	90
Starting your Cloud Insights free trial	90
Sign in and go	
Logging Out	91
Security	92
Cloud Insights Security	92
Information and Region	94
SecurityAdmin Tool.	96
Getting Started	105
Feature Tutorials	105
Collecting Data	
Importing from the Dashboard Gallery	
User Accounts and Roles	
Cloud Insights Data Collector List	143
Subscribing to Cloud Insights	
Trial Edition	
Module Trials	
Subscription Options	
How Do I Subscribe?	
View Your Subscription Status	
View your Usage Management	
Subscribe Directly and Skip the Trial	
Adding an Entitlement ID	152

Automatic Device Resolution	153
Automatic Device Resolution Overview	153
Device Resolution rules	155
Fibre Channel device resolution	158
IP device resolution	160
Setting options in the Preferences tab	162
Regular expression examples	163
Creating Dashboards	170
Dashboards Overview	170
Dashboard Features	173
Sample Dashboards	204
Best Practices for Dashboards and Widgets	209
Kubernetes	214
Kubernetes Cluster Overview	214
Before Installing or Upgrading the NetApp Kubernetes Monitoring Operator	215
Kubernetes Monitoring Operator Installation and Configuration	220
NetApp Kubernetes Monitoring Operator Configuration Options	238
Allow server status reports generated by mod_status,	974
with the URL of http://servername/server-status	974

Cloud Insights documentation

NetApp Cloud Insights is a cloud infrastructure monitoring tool that gives you visibility into your complete infrastructure. With Cloud Insights, you can monitor, troubleshoot and optimize all your resources including your public clouds and your private data centers.

What can Cloud Insights do for me?

Cloud Insights provides hybrid multicloud monitoring, giving you full-stack observability of infrastructure and workloads.

- · Data collectors for heterogeneous infrastructure and workloads, including Kubernetes
- Open Telegraf collector and open APIs for easy integration
- · Comprehensive alerting and notifications
- · Machine learning for intelligent insights
- · Optimize resource utilization
- · Built-in or customizable dashboards with advanced filters to minimize display noise to answer questions
- Discover the health of your ONTAP storage operations
- Protect your most valuable business asset data from ransomware or data destruction attack

Getting Started

- · How do I get started with Cloud Insights?
- I'm signed up. Now what do I do? Acquiring Data Setting up users
- Awesome! What's Next?
 Preparing Assets: Annotating
 Finding the Assets You Want: Querying
 Seeing the Data You want: Dashboards
 Monitoring and Alerts
 Securing Data
- This is great stuff! I'm ready to subscribe.

What's New with Cloud Insights

NetApp is continually improving and enhancing its products and services. Here are some of the latest features and functionalities available in the commercial Editions of Cloud Insights.



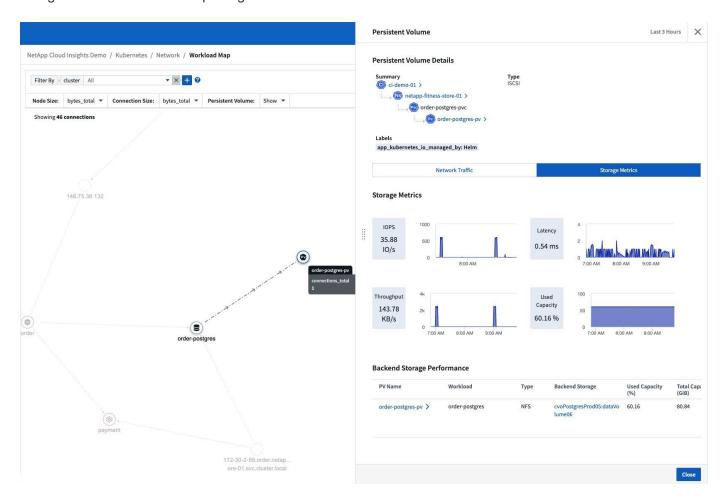
Some of these features may not be available in Cloud Insights Federal Edition, or may be available with reduced functionality. Wherever possible, these differences are noted in the documentation.

April 2024

iSCSI SAN Support for Kubernetes Network Map

Identify Kubernetes underlying storage dependencies for NFS and now iSCSI.

Accelerate issue resolution by gaining insight into the Kubernetes application's storage network activity and undelineated storage metrics. Understand how the Kubernetes application uses the storage infrastructure for chargeback and showback reporting.



Operating System support

The following operating systems are supported with Cloud Insights Acquisition Units, in addition to those already supported:

• Oracle Enterprise Linux 8.8

- Red Hat Enterprise Linux 8.8
- Rocky 9.3
- OpenSUSE Leap 15.1 through 15.5
- SUSE Enterprise Linux Server 15, 15 SP2 through 15 SP5

March 2024

Workload Security Agent Details

Each of your Workload Security Agents has its own landing page, where you can easily see summary information about the Agent as well as the installed Data and User Directory Collectors associated with that Agent.

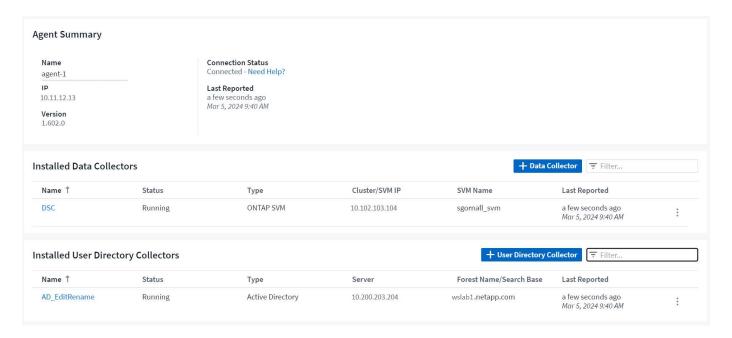
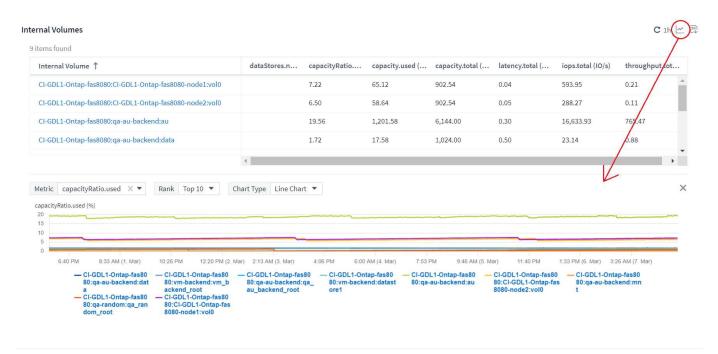


Chart more data more quickly

When analyzing data on an asset's landing page, adding additional data to the Expert View charts is a snap. For each table on the landing page, if an object type has relevant data, hover over that object to display the "Add to Expert View" icon. Selecting this icon adds that object to the Additional Resources and displays it in the Expert View charts.

Storage Node ↑ Add to Expert View CI-GDL1-Ontap-fas8080-node1 □ ② CI-GDL1-Ontap-fas8080-node2

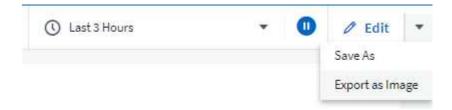
Or maybe you want to see a landing page table's data in its own chart. Simply select the *Show Chart* icon to open the chart below the table:



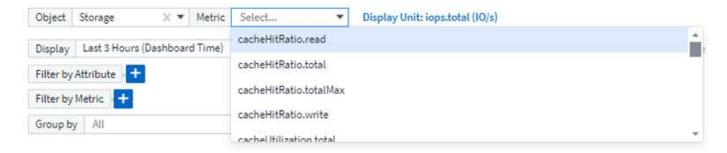
February 2024

Usability Improvements

Save a **snapshot** of your current dashboard by selecting *Export as Image* from the right-corner drop-down. Cloud Insights creates a .PNG of the current widget states.



Object and Metric selection is easier than ever for Widgets, Monitors, etc. Choose the object type you want, and then select a metric relevant to that object in the separate drop-down.



Export Data Collector and Acquisition Unit lists to .CSV by selecting the icon at the top of those pages.



We've **re-organized the Help > Support** page so it's easier to find what you're looking for, and because you asked for them, we added direct links on this page to **API Swagger** and user documentation.

API Access:

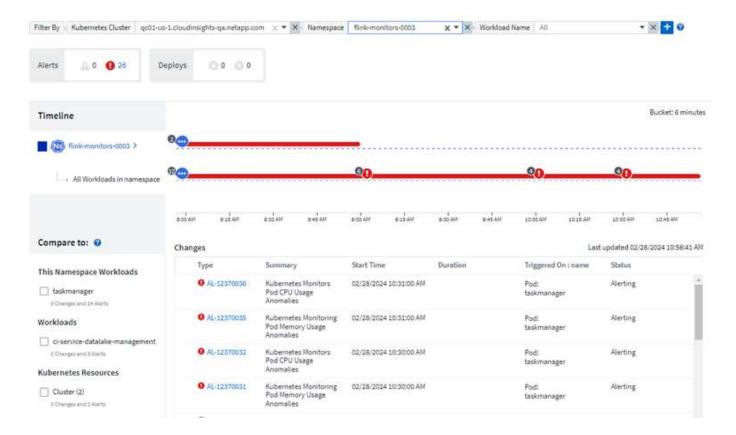
To integrate Cloud Insights with other applications see the Cloud Insights API List and documentation.

Links in the "triggeredOn" column on the Alerts list page will navigate to the appropriate Landing Page, if a Landing Page is available for that object.



See all changes in your namespace

Kubernetes Change Analysis now allows you to see a timeline of changes when selecting Cluster and Namespace. Previously, Workload must also have been selected. When filtering on Cluster and Namespace, the timeline of all workload changes in that namespace are shown on one line.



Related Logs for Alerts

When viewing a log alert, related log entries are shown in a new table. A log entry is related if it occurs in the same source and timeframe as the alert, and is subject to the same conditions. Select "Analyze Logs" to explore further.

Related Logs Analyze Logs timestamp ↓ message message 02/28/2024 11:07:21 AM iscsi.loginFailure: ISCSI: iSCSI login failure, "Invalid TargetName iqn.1992-08.com.netapp:sn.6ed012db378611ee8f24005056b3dcd8:vs.3 from Initiator iqn.1994-05.com.redhat:dc7292e4b936 at IP address 10.192.38.34" 02/28/2024 11:06:24 AM iscsi.loginFailure: ISCSI: iSCSI login failure, "Invalid TargetName iqn.1992-08.com.netapp:sn.091b27ae993c11ee9765005056b3f163:vs.3 from Initiator iqn.1994-05.com.redhat:e861299d2ffc at IP address 10.192.33.88" 02/28/2024 11:06:24 AM iscsi.loginFailure: ISCSI: iSCSI login failure, "Invalid TargetName iqn.1992-08.com.netapp:sn.091b27ae993c11ee9765005056b3f163:vs.3 from Initiator iqn.1994-05.com.redhat:e861299d2ffc at IP address 10.192.33.88"

Collect ONTAP Switch Data

Cloud Insights can collect data from the ONTAP system's back-end switches; simply enable the collection in the data collector's *Advanced Configuration* section, and ensure the ONTAP system is configured to provide switch information and has the appropriate permissions set.

Workload Security Data Collector API

In large environments, you can automate Workload Security collector creation using the new Data Collectors API. Navigate to **Admin > API Access > API Documentation** and select the *Workload Security* API type to learn more.

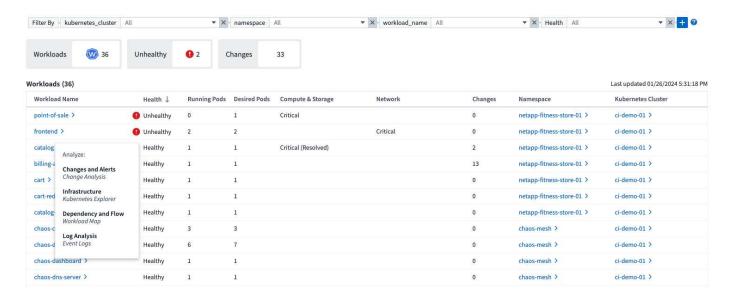
January 2024

Try Cloud Insights Features you haven't used yet

In addition to your initial trial of Cloud Insights, you may also take advantage of Module Trials. For example, if you are subscribed to Cloud Insights and have been monitoring storage and virtual machines, when you add Kubernetes to your environment, you will automatically enter into a 30-day trial of Kubernetes Observability. Kubernetes Observability managed unit usage will not count against your subscribed entitlement until after the trial period ends.

How healthy are my workloads?

Workload health is available at a glance on the **Kubernetes > Explore > Workloads** page, so you can quickly see which workloads are performing well and which may need some help. Easily identify if the health issue is related to infrastructure, network, or configuration changes, and drill down to analyze the root cause.



Data Collector Updates

Data Domain Identification

The Data Domain collector has been improved to better identify HA systems for durability across failover events This change will cause a **one time** re-identification of Data Domain appliances in HA systems, which will subsequently cause any annotations on those assets to be removed (because these arrays will be re-identified). You will need to re-attach annotations to your Data Domain objects.

Enhanced Ransomware Detection ML Algorithm

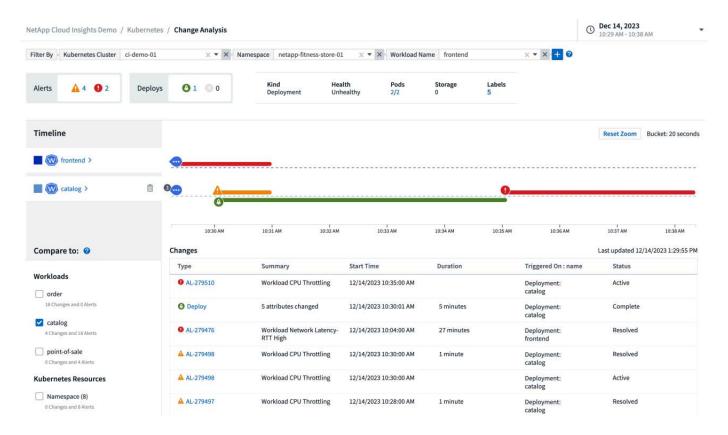
Workload Security includes a new 2nd-generation ransomware detection ML algorithm to detect the most sophisticated attacks faster and more accurately.

"Seasonality" of behaviors: weekend behavior may follow different patterns from weekday, or morning behavior from afternoon. Workload Security algorithms take this seasonality into account.

December 2023

Change Analytics at a glance

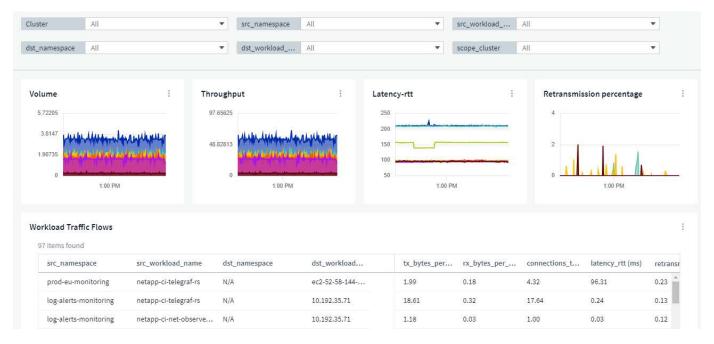
Kubernetes Change Analytics provide you with an all-in-one view of recent changes to your Kubernetes environment. Alerts and deployment status are at your fingertips. With Change Analytics, you can track every deployment and configuration change, and correlate it with the health and performance of K8s services, infrastructure, and clusters.



Kubernetes Workload Performance Dashboard

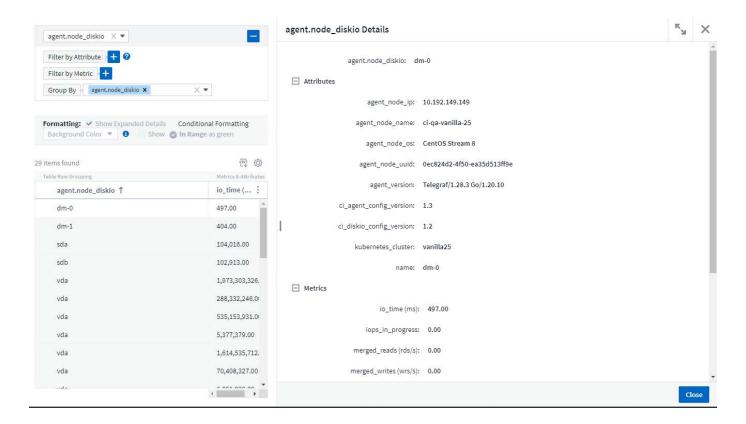
Workload performance is available at a glance in the comprehensive Kubernetes Workload Performance dashboard. Quickly view graphs of Volume, Throughput, Latency, and Retransmission trends, as well as a table of workload traffic for each namespace in your environment. Filters allow easy focus into areas of interest.





Query Details on one screen

In a query, selecting a row opens a side panel showing attribute, annotation, and metric details for the selected row, providing helpful information without needing to drill into the object's landing page. Links in the row or side panel allow for easy navigation.



Data Collector updates:

- **Brocade FOS REST**: This collector is moved out of "Preview" and is now generally available. Some things to note:
 - FOS introduced their REST API with FOS 8.2. But some features like routing only received REST API capabilities with 9.0.
 - If you have a fabric consisting of mixed FOS assets 8.2 higher, as well as some < 8.2, the Cloud Insights FOS REST collector will fail to discover those older assets. You can edit the FOS REST collector and build a comma-delimited list of the IPv4 address of those devices for exclusion from that collector.
- **SELinux**: Cloud Insights includes enhancements to the Linux Acquisition Unit initial installation to ensure robustness of operation within Linux environments with SELinux enforcement enabled. These enhancements only impact *new* AU deployments; if you have any SELinux issues relating to AU upgrades, contact NetApp Support to remediate your SELinux configuration.

November 2023

Workload Security: Pause/Resume a Collector

In Workload Security, you can Pause a Data Collector if the collector is in *Running* state. Open the "three dots" menu for the collector and select PAUSE. While the collector is paused, no data is gathered from ONTAP, and no data is sent from the collector to ONTAP. Select Resume to begin collecting again.

Storage Node Support Information

On a storage node landing page, the User Data section provides at-a-glance information about your support

offering, current status, support status, and warranty end date. Note that Cloud Insights currently only autopublishes this information for NetApp devices. Note also that these support fields are annotations, so they can be used in queries and dashboards.

User Data	+ Annotation
Serial Number Active	
Yes	
Serial Number Support Status	
Y	
Support Offering	
WARRANTY	
Warranty End Date	
12/31/2023	

Map VMWare tags to Cloud Insights annotations

The VMWare data collector allows you to populate Cloud Insights text annotations with same-name tags that are configured on VMWare.

Brocade CLI collector reliability enhancements for FOS 9.1.1c and higher firmware

On some Brocade Fibre Channel switches running 9.1.1c firmware, certain CLI commands' output may be prepended with the "motd" login banner text, or warnings for users to change default passwords. The Brocade CLI collector has been enhanced to ignore these two types of extraneous text.

Prior to this enhancement, only FOS 9.1.1c switches without Virtual Fabrics present were likely discoverable with this collector type.

October 2023

Enhanced Workload Security

Workload Security has been improved with the following:

- Access Denied: Workload Security integrates with ONTAP to receive "Access Denied" events and provide an additional analytics and automatic responses layer.
- **Allowed File Types**: If a ransomware attack is detected for a known file extension, that file extension can be added to an allowed file types list to prevent unnecessary alerting.

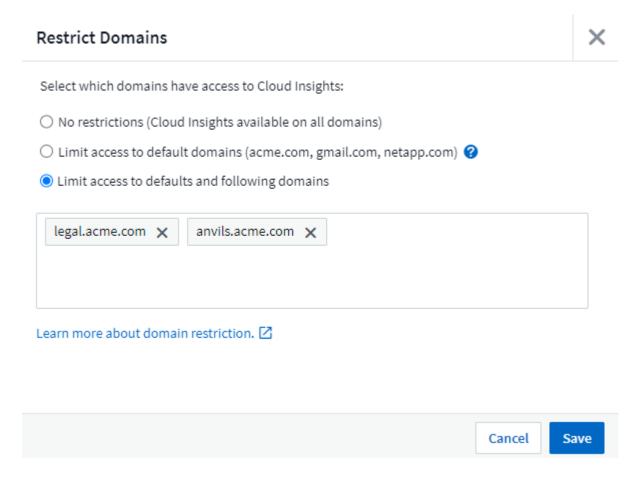
Module Trials

In addition to your initial trial of Cloud Insights, you may also take advantage of Module Trials. For example, if you are already subscribed to Infrastructure Observability but are adding Kubernetes to your environment, you

will automatically enter into a 30-day trial of Kubernetes Observability. You will only be charged for your Kubernetes Observability managed unit usage at the end of the trial period.

Restrict access to specified domains

Admins and Account Owners now have the ability to restrict Cloud Insights access to email domains they specify. Go to **Admin > User Management** and select the *Restrict Domains* button.



Data Collector Updates

The following Data Collector/Acquisition Unit changes are in place:

- Isilon / PowerScale REST: Various new attributes and metrics have been added to Cloud Insights enhanced analytics capabilities under the *emc_isilon.node_pool.** name. These counters and attributes will empower users to build dashboards and monitors for *node_pool* capacity consumption; users with Isilon clusters built from dissimilar hardware node models will have multiple node pools, and understanding your HDD/SSD/total capacity consumption at a node pool level is useful for both monitoring and planning.
- **Rubrik** "Service account" authentication support: Cloud Insights' Rubrik collector now supports both traditional HTTP Basic Authentication (username and password), and Rubrik's Service Account approach, which requires a username + secret + Organization ID.

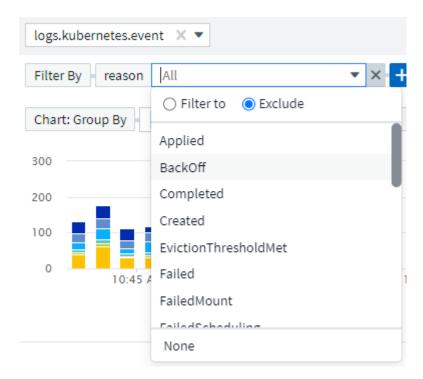
September 2023

Easily Find What You Want in the Logs

Log Query (**Observability > Log Queries > +New Log Query**) includes a number of enhancements to make log exploration easier and more informative.

Include/Exclude

When filtering for a value, you can easily choose whether to **Include** or **Exclude** results matching the filter. Selecting "Exclude" creates a "NOT <value>" filter. You can combine Include and Exclude values in a single filter.



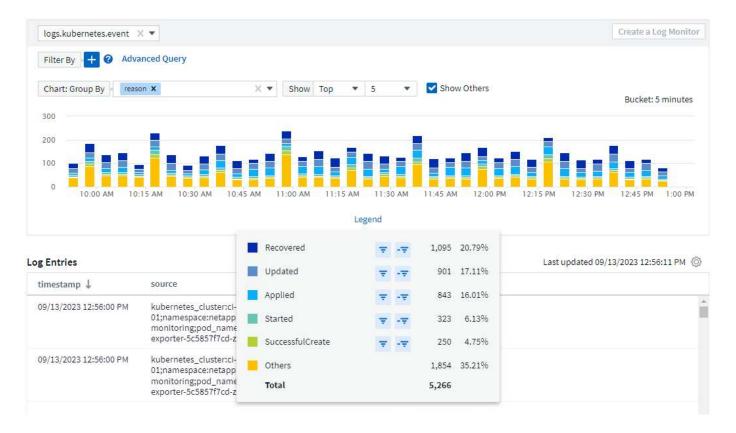
Advanced Query

Advanced Querying gives you the opportunity to create "free form" filters, combining or excluding values using AND, NOT, OR, wildcards, etc.

The "Filter By" and Advanced Query are "AND"ed together to form a single query. The results are displayed in the results list and the chart.

Grouping in the Chart

When you select a log attribute to **Group By**, the list and chart show the results of the current filter. In the chart, columns grouped into colors. Hovering over a column in the chart will display details about the specific entries, similar to the overall information shown when you expand the chart Legend. In the legend, you can also choose to set an Include or Exclude filter for a specific grouping.

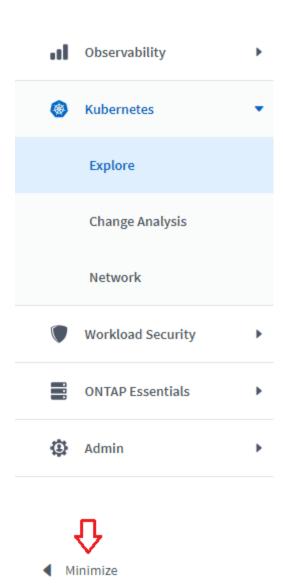


"Floating" Log Detail Panel

When exploring logs using the Log Query, selecting an entry in the list opens a detail panel for that entry. You can now choose to display that slideout panel "Floating" (i.e. displayed over the rest of the screen) or 'In Page" (i.e. displayed as its own frame within the page). To switch between these views, select the "In Page / Floating" button in the upper-right corner of the panel.

Collapse the Menu

You can collapse the left-side Cloud Insights navigation menu by selecting the "Minimize" button below the menu. While the menu is minimized, hover over an icon to see which section it opens; selecting the icon opens the menu and takes you directly to that section.



Data Collector Improvements

Cloud Insights has made it easier to show and find data collector information:

- **Processing of data collector lists** is more efficient, which means the time it takes to display and navigate these lists is greatly reduced. If you have a large environment with many data collectors, you will see a significant improvement when listing your data collectors.
- The **Data Collector Support Matrix** has moved from a .PDF file to an .HTML-based page, quicker to navigate and easier to maintain. Check out the new Matrix here: https://docs.netapp.com/us-en/cloudinsights/reference_data_collector_support_matrix.html

August 2023

Collecting Isilon/PowerScale Logs and Advanced Analytics Data

The Isilon REST and PowerScale REST collectors contain the following improvements:

- · Isilon log events are available for use in queries and alerts
- Isilon Advanced Analytic attributes are available for use in queries, dashboards, and alerts:
 - emc_isilon.cluster
 - emc_isilon.node
 - emc isilon.node disk
 - · emc isilon.net iface

These are enabled by default for users of the Isilon REST and/or PowerScale REST collectors. NetApp strongly encourages users of the Isilon CLI-based collector to migrate to the new REST API-based collector to receive enhancements such as the above.

Improved Workload Map

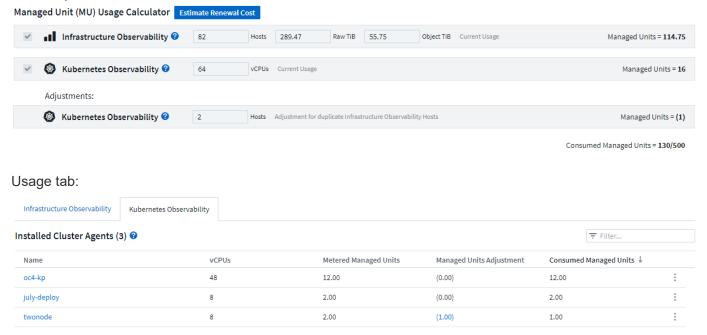
The workload map is more usable and less noisy; it groups all similar external services into one node if they communicate with the same workloads, reducing the complexity of the graph and making it easier to understand how services are interconnected.

Choosing a grouped node will display a detailed table with the network traffic metrics for each external service relevant to that node.

Kubernetes Managed Unit usage adjustment

In the event of a compute resource in your Kubernetes cluster environment being counted by both the NetApp Kubernetes Monitoring Operator and an underlying infrastructure data collector (for example, VMware), your usage of these resources will be adjusted to ensure the most efficient counting of managed units. You can view the Kubernetes MU adjustments on the Admin > Subscription page, in both the Summary and Usage tabs.





Collector/Acquisition changes:

The following Data Collector/Acquisition Unit changes are in place:

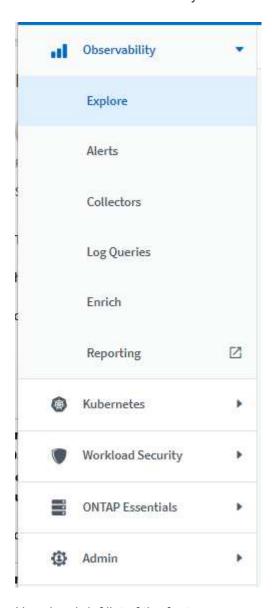
• Acquisition Units now support RHEL 8.7.

Improved Menus

We have updated the left hand navigation menu to better support our customers' workflows. New top level items such as *Kubernetes* provide accelerated access to what the customer needs, and a consolidated administrators console supports the tenant owner role.

Here are some additional examples of the changes:

- The top level Observability menu showcases data discovery, alerting and log queries
- · 'API Access' functionality for Observability and Workload Security are under one menu
- · Likewise for Observability and Workload Security 'Notifications' functionality, also now under one menu



Here is a brief list of the features you can find under each menu:

Observability:

• Explore (Dashboards, Metric Queries, Infrastructure Insights)

- Alerts (Monitors and Alerting)
- Collectors (Data Collectors and Acquisition Units)
- · Log Queries
- Enrich (Annotations and Annotation Rules, Applications, Device Resolution)
- Reporting

Kubernetes:

· Cluster Exploration and Network Map

Workload Security:

- Alerts
- Forensics
- Collectors
- · Policies

ONTAP Essentials:

- Data Protection
- · Security
- Alerts
- Infrastructure
- Networking
- Workloads
 - *VMware

Admin:

- API Access
- Auditing
- Notifications
- · Subscription Information
- · User Management

July 2023

Show Recent Changes

Data Collector landing pages now include a list of recent changes. Simply click the "Recent Changes" button at the bottom of any data collector landing page to display recent data collector changes.

Changes Reported by This Data Collector (1)

Time ↓	Change
07/06/2023 6:39:12 PM	☐ Storage CI-GDL1-Ontap-fas8080 configuration changed Property Display IP is changed from "10.192.122.10" to "10.192.122.12" Property Manage URL is changed from "HTTPS://10.192.122.10:443" to "HTTPS://10.192.122.12:443"
Hide Recent Changes	

Operator Improvements

The following improvements have been made to Kubernetes Operator deployment:

- Option to bypass docker metric collection
- · Ability to add and customize tolerations to telegraf Daemonsets and Replicasets

Insight: Reclaim Cold Storage

The Reclaim ONTAP Cold Storage Insight now supports FlexGroups, and is now available to all customers.

Operator Image Signature

For customers who use a private repository for their NetApp Kubernetes Monitoring Operator, you can now copy the Image Signature Public Key during Operator installation, allowing you to confirm authenticity of the downloaded software. Select the *Copy Image Signature Public Key* button during the optional step to *Upload the operator image to your private repository*.



Aggregation, Conditional Formatting, and more for Queries

Aggregation, Unit Selection, Conditional Formatting, and Column Renaming are among the most useful features of a dashboard table widget, and now those same features are available for Queries.

These features are available now for integration-type data (Kubernetes, ONTAP Advanced Metrics, etc.), and will be coming soon for Infrastructure objects (storage, volume, switch, etc.).

API for Audit

You can now use an API to query or export Audited events. Go to Admin > API Access and select the *API Documentation* link for information.



Data Collector: Trident Economy

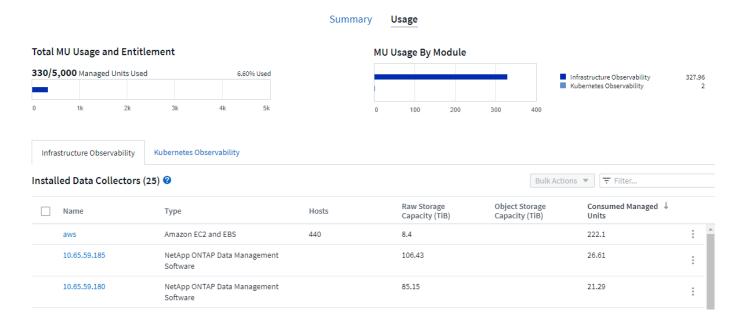
Cloud Insights now supports the Trident Economy Driver, realizing these benefits:

- Get visibility into pod-to-ONTAP Qtree mapping and performance metrics.
- · Provide seamless troubleshooting and easy navigation from Kubernetes pods to backend storage
- Proactively detect backend performance issues with monitors

June 2023

Check out your Usage

Beginning in June, 2023, Cloud Insights provides a breakdown of Managed Unit usage based on Feature Set. Now you can quickly view and monitor managed unit (MU) usage for your Infrastructure as well as MU usage tied to Kubernetes.



Kubernetes Network Monitoring and Map is available for all

The *Kubernetes Network Performance and Map* simplifies troubleshooting by mapping dependencies between Kubernetes workloads, providing real-time visibility into Kubernetes network performance latencies and anomalies to identify performance issues before they affect users. Many customers found it helpful during

Preview, and now it's available for everyone to enjoy.

Collector/Acquisition changes:

The following Data Collector/Acquisition Unit changes are in place:

- Data Domain and Cohesity MUs are metered at 40 TiB: 1 MU.
- Acquisition Units now support RHEL and Rocky 9.0 and 9.1.

New ONTAP Essentials dashboards

The following ONTAP Essentials dashboards have been available in Preview environments, and now they are available for everyone:

- · Security Dashboard
- Data Protection Dashboard (includes Local and Remote Protection overviews)

Additional System Monitors

The following System Monitors are included with Cloud Insights:

- Storage VM FCP Service Unavailable
- Storage VM iSCSI Service Unavailable

May 2023

Improved Kubernetes Monitoring Operator Installation

Installation and configuration of the NetApp Kubernetes Monitoring Operator is easier than ever with the following improvements:

- Environment configuration settings are held in a single, self-documented config file.
- · Step-by step instructions for uploading Kubernetes Monitoring Operator images to your private repository.
- Simple to upgrade with a single command to upgrade your Kubernetes Monitoring while keeping custom configurations.
- More secured: API keys are securely managing secrets.
- Easy to integrate and deploy with your CI/CD automation tools.

Storage Virtualization

Cloud Insights can differentiate between a storage array having local storage or virtualization of other storage arrays. This gives you the ability to relate cost and distinguish performance from the front-end all the way to the back-end of your infrastructure.

Storage Summary		
Model:	Virtualized Type:	IOPS - Total:
V-Series	Virtual	N/A
Vendor:	Backend Storage:	Throughput - Total:
NetApp	Sym-000050074300343	N/A
Family:	Microcode Version:	Management:
V-Series	8.0.2 7-Mode	
		FC Fabrics Connected:
Serial Number:	Raw Capacity:	7
1306894	0.0 GiB	
		Alert Monitors:
IP:	Latency - Total:	
192.168.7.41	N/A	

New Webhook Parameters

When creating a Webhook notification, you can now include these parameters in your webhook definition:

- %%TriggeredOnKeys%%
- %%TriggeredOnValues%%

Reporting on Kubernetes data

Kubernetes data collected by Cloud Insights—including Persistent Volumes (PV), PVC, Workloads, Clusters, and Namespaces—is now available for use in Reporting, enabling chargeback, trending, forecasting, TTF calculations, and other business reporting on metrics for Kubernetes.

Default ONTAP System Monitors Enabled for New Customers

Many ONTAP System Monitors are enabled (i.e. *Resumed*) by default in new Cloud Insights environments. Previously, most monitors defaulted to *Paused* state. Because business needs vary from company to company, we always recommend taking a look at the system monitors in your environment and pausing or resuming each based on your alerting needs.

April 2023

Kubernetes Performance Monitoring and Map

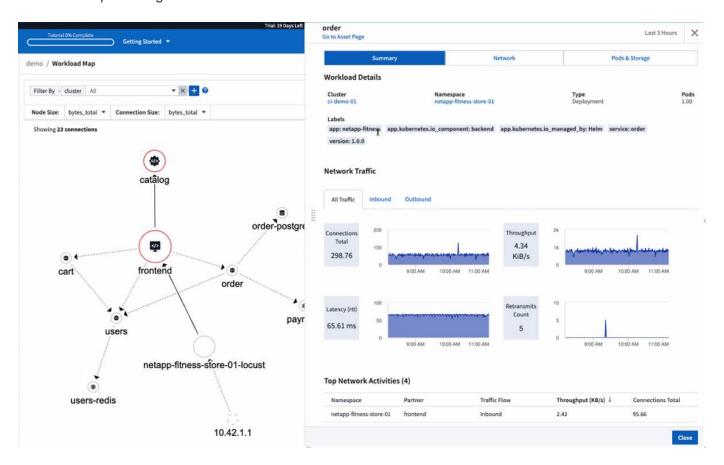
The *Kubernetes Network Performance and Map* feature simplifies troubleshooting by mapping dependencies between Kubernetes workloads. It provides real-time visibility into Kubernetes network performance latencies and anomalies to identify performance issues before they affect users.

This capability helps organizations reduce overall costs by analyzing and auditing Kubernetes traffic flows.

Kev Features:

- The Workload Map presents Kubernetes workload dependencies and flows and highlights network and performance issues.
- Monitor network traffic between Kubernetes pods, workloads, and nodes; identifies the source of traffic and latency problems.
- Reduce overall costs by analyzing ingress, egress, cross-region, and cross-zone network traffic.

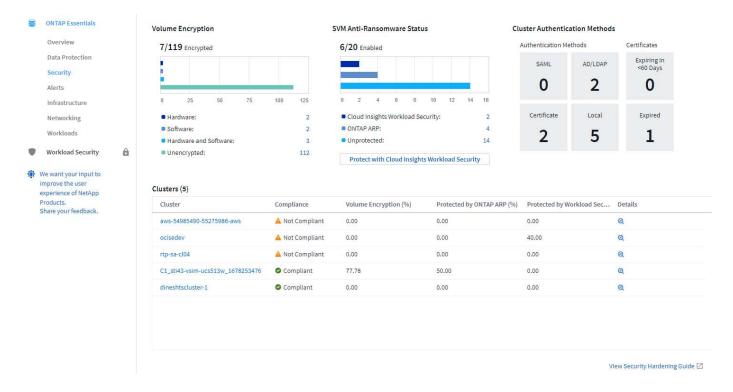
Workload Map showing "Slideout" details:



Kubernetes Performance Monitoring and Map is available as a Preview feature.

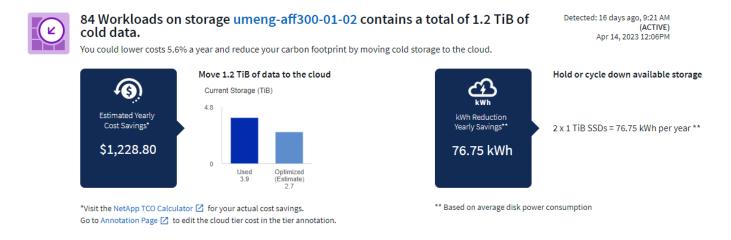
ONTAP Essentials Security Dashboard

The Security Dashboard gives you an instant view of your current security situation, showing charts for hardware and software volume encryption, anti-ransomware status, and cluster authentication methods. The Security Dashboard is available as a Preview feature.



Reclaim ONTAP Cold Storage

The *Reclaim ONTAP Cold Storage* Insight provides data about cold capacity, potential cost/power savings, and recommended action items for volumes on ONTAP systems.



With this Insight, you can answer such questions as:

- What amount of cold data on a storage cluster are sitting on (a) high-cost SSD disks, (b) HDD disks, and
 (c) virtual disks?
- What workloads are the highest contributors in regards of the non-optimized storage?
- What is the duration (in days) the data has been cold on a given workload?

Reclaim ONTAP Cold Storage is considered a Preview feature and is therefore subject to change.

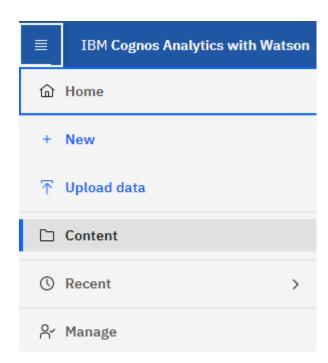
Subscription Notification also controls banner messages

Setting recipients for Subscription Notifications (Admin > Notifications) now also controls who will see subscription-related in-product banner notifications.

Your subscription is expiring in 2 days. View Subscription

Reporting has a new look

You will notice that Cloud Insights Reporting screens have a new look, and that some of the menu navigation have changed. These screens and navigation changes have been updated in the current Reporting Documentation.



Monitors Paused by Default

For new Cloud Insights environments, be aware that system-defined monitors do not send alert notifications by default. You will need to enable notifications for any monitor that you want alerting you, by adding one or more delivery methods for the monitor.

For existing Cloud Insights environments, the default global notification recipient list has been removed for any system-defined monitors currently in Paused state. User-defined notifications remain unchanged, as do notification settings for currently active system-defined monitors.

Looking for the API Metering tab?

API Metering has moved from the Subscription page to the **Admin > API Access** page.

March 2023

Cloud Connection for ONTAP 9.9+ deprecated

The Cloud Connection for ONTAP 9.9+ data collector is being deprecated. Starting April 4, 2023, Cloud Connection data collectors in your environment will no longer collect data, and will instead present an error when polling. The Cloud Connection data collector will be removed altogether from Cloud Insights in a subsequent update.

Prior to April 4, 2023, it is mandatory to configure a new NetApp ONTAP Data Management Software data collector for any ONTAP systems currently collected by Cloud Connection. Learn More.

January 2023

New Log Monitors

We've added almost two dozen additional system monitors to alert for broken interconnect links, heartbeat problems, and more. Additionally, three new Data Protection log monitors have been added, to alert on SnapMirror Auto Resync, MetroCluster Mirroring, and FabricPool Mirror Resync changes.

Note that some of these monitors will be *enabled* by default; you must *pause* them if you do not wish to alert on them. Also note that these monitors are not configured to deliver notifications; you must configure notification recipients on these monitors if you want to send alerts via email or webhook.

.CSV Export for all Dashboard Table Widgets

Ensuring accessibility to your data is essential, so we've made .CSV export [.csv export icon] available for all metric queries, dashboard table widgets, and object landing pages, regardless of the type of data (asset or integration) you're querying.

Data customizations like column selection, renaming columns, and unit conversions are also now included in the new export functionality.

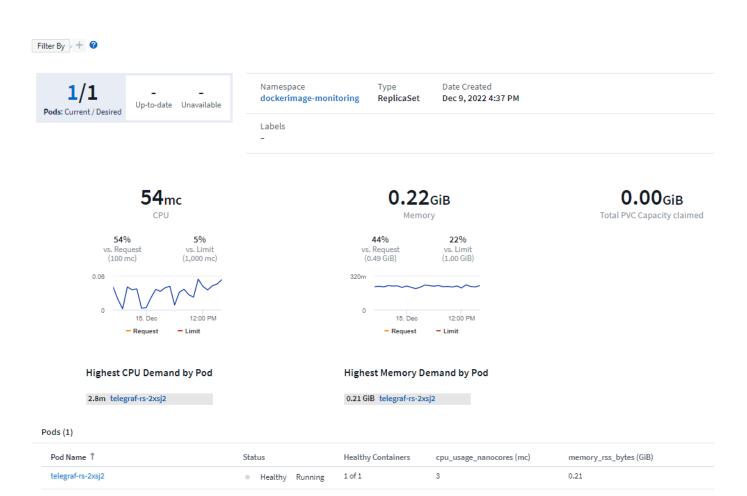
December 2022

Explore Ransomware Protection and other security features during Cloud Insights Trial

Starting today, signing up for a new Trial of Cloud Insights allows you to explore Security features such as Ransomware detection and automated user-blocking response policy. If you haven't signed up for your Trial, do it today!

Kubernetes Workloads have their own landing page

Workloads are a key part of your Kubernetes environment, so Cloud Insights now provides landing pages for those workloads. From here, you can view, explore, and troubleshoot issues that affect your Kubernetes workloads.



Check your Checksums

You asked us to provide you with checksum values during installation of the agent for Windows and Linux and we think that's a great idea. So here they are:

☐ Manually Verifying Telegraf Checksums

The Cloud Insights agent installer performs integrity checks, but some users may want to perform their own verifications before installing or applying downloaded artifacts.

For more information, read about verifying checksums before proceeding to the next step.

The SHA256 checksum for this telegraf.pkg is:

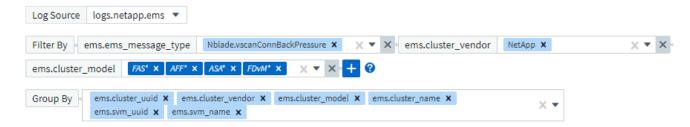
cbd0d8d0512b65fbcd0c786d8d0512b651de0e1cf003e0a0d9df01d8d0512b65

Log Alerting Improvements

Group By

When creating or editing a Log Monitor, you can now set "Group By" attributes to allow for more focused alerting. Look for the "Group By" attributes below the "filter" settings in your monitor definition.

Select the log to monitor



This change brings Metric Monitors and Log Monitors into feature parity by normalizing the "Group By" aspect of Monitor Definitions. This parity will allow customers to clone/duplicate **all** system-defined default Monitors for further customization.

Duplicating

You can now clone (duplicate) the Change Log, Kubernetes Log, and Data Collector Log monitors. This creates a new custom log monitor that you can modify to your specific definitions.



11 New Default ONTAP Monitors covering SnapMirror for Business Continuity

We've added almost a dozen new system monitors for SnapMirror for Business Continuity (SMBC), which alert on changes to SMBC certificates and ONTAP Mediators.

November 2022

More than 40 new Security, Data Collection, and CVO monitors!

We've added dozens of new system-defined monitors to alert you to potential issues with Cloud Volumes, Security, and Data Protection. Read more about these monitors here.

October 2022

Better and more accurate Ransomware detection with ONTAP Autonomous Ransomware Protection integration

Cloud Secure improves ransomware detection through integration with ONTAP Autonomous Ransomware Protection (ARP).

Cloud Secure receives ONTAP ARP events on potential volume file encryption activity, and

- · Correlates volume encryption events with user activity to identify who is causing the damage,
- Implements automatic response policies to block the attack,
- Identifies which files were affected, helping to recover faster and conduct data breach investigations.

September 2022

Monitors available in Basic Edition

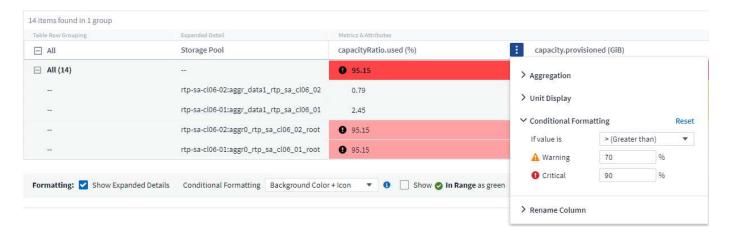
ONTAP Default monitors now available to use in Cloud Insights Basic Edition. This includes more than 70 infrastructure monitors and 30 workload examples.

ONTAP Power and StorageGRID dashboards

The dashboard gallery includes a new dashboard for ONTAP Power and Temperature as well as four dashboards for StorageGRID. If your environment is collecting ONTAP power metrics and/or StorageGRID data, import these dashboards by selecting **+From Gallery**.

At-a-glance threshold visibility in tables

Conditional Formatting allows you to set and highlight Warning-level and Critical-level thresholds in table widgets, bringing instant visibility to outliers and exceptional data points.

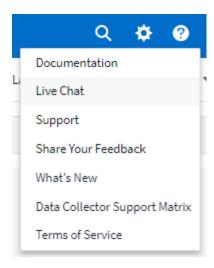


Security Monitor

Cloud Insights can alert you when it detects that FIPS mode is disabled on the ONTAP system. Read more about System Monitors, and watch this space for more Security Monitors, coming soon!

Chat from Anywhere

Chat with a NetApp Support specialist from any Cloud Insights screen by selecting the new **Help > Live Chat** link. Help is available from the "?" icon in the upper right of the screen.



More visible Insights

If your environment is experiencing an Insight such as *Shared resources Under Stress* or *Kubernetes Namespaces Running Out of Space*, asset landing pages for resources affected now include links to the Insight itself, providing quicker exploration and troubleshooting.

New Data Collectors

- Amazon S3 (available in Preview)
- Brocade FOS 9.0.x
- Dell/EMC PowerStore 3.0.0.0

Other Data Collector Updates

All data sources are now optimized to resume performance polling after Acquisition Unit updates and/or patches.

Operating System support

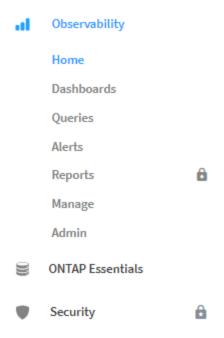
The following operating systems are supported with Cloud Insights Acquisition Units, in addition to those already supported:

• Red Hat Enterprise Linux 8.5, 8.6

August 2022

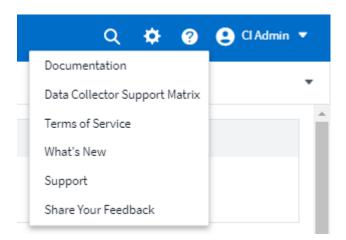
Cloud Insights has a new look!

Starting this month, "Monitor and Optimize" has been renamed **Observability**. You'll find all your favorite features like Dashboards, Queries, Alerts, and Reporting here. In addition, look for Cloud Secure under the new **Security** menu. Note that only the menus have changed; feature functionality remains the same.



Looking for the **Help** menu?

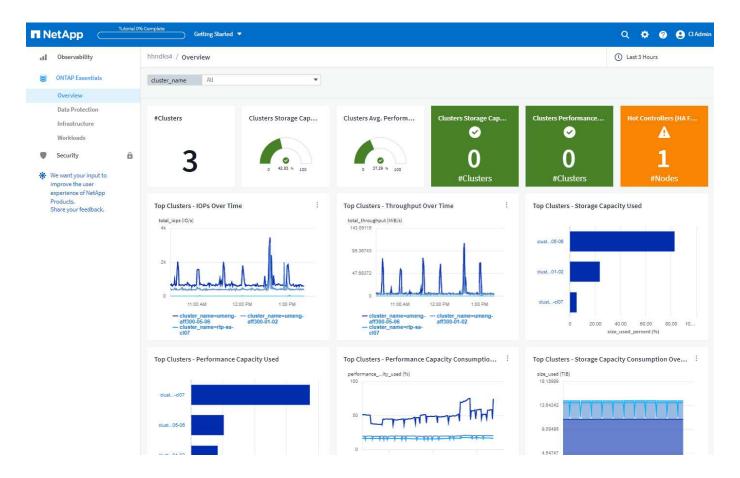
Help now lives in the upper right of the screen.



Not sure where to start? Check out ONTAP Essentials!

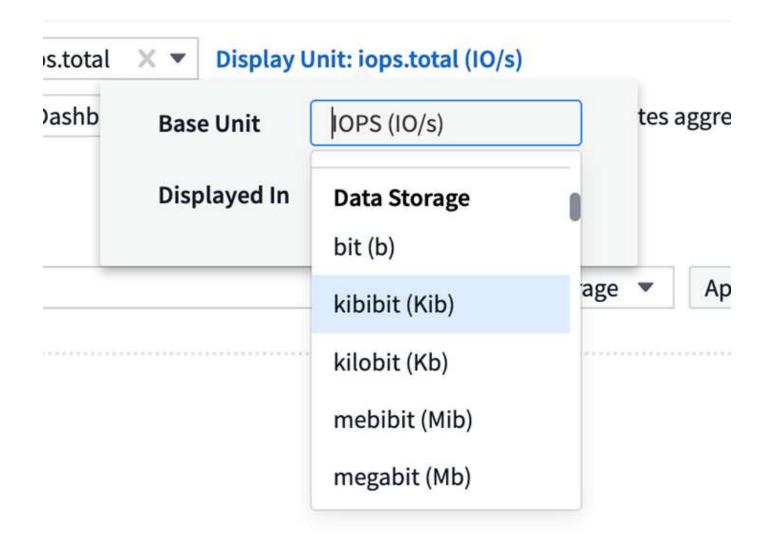
ONTAP Essentials is a set of dashboards and workflows that provide detailed views into your NetApp ONTAP inventories, workloads, and data protection, including days-to-full predictions for storage capacity and performance. You can even see if any controllers are running at high utilization. ONTAP Essentials is your ideal place for all of your NetApp ONTAP monitoring needs!

ONTAP Essentials—available in all Editions—is designed to be intuitive to existing ONTAP operators and administrators, easing the transition from ActiveIQ Unified Manager to service-based management tools.



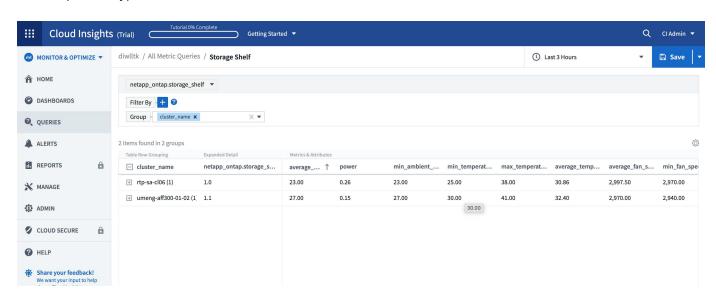
Storage Data families are merged

You asked for it, and now you've got it. Storage base-2 and base-10 data units are now combined into one family, from bits and bytes to tebibits and terabytes, making it easier to display data your way on your dashboards. Data Rates are also now one big family of their own.



How much power is my storage using?

Display and monitor your ONTAP storage shelf and node power consumption, temperature, and fan speed, using the netapp_ontap.storage_shelf, netapp_ontap.system_node and netapp_ontap.cluster (power consumption only) metrics.



Features graduated from Preview

The following features have moved out of Preview and are now available to all customers:

Feature	Description
Kubernetes Namespaces Running out of Space	The Kubernetes Namespaces Running Out of Space Insight gives you a view into workloads on your Kubernetes namespaces that are at risk of running out of space, with an estimate for the number of days remaining before each space becomes full. Read More
Shared Resource Under Stress	The Shared Resource Under Stress insight uses AI/ML to automatically identify where resource contention is causing performance degradation in your environment, highlights any workloads impacted by it, and provides recommended actions to remediate, letting you solve performance issues more quickly. Read More
Cloud Secure – Block user access on attack	Greater protection for your business-critical data with the ability to block user access when an attack is detected. Access can be blocked automatically, using Automated Response Policies, or manually from the alert or user details pages. Read More

How's my data collection health?

Cloud Insights provides two new heartbeat monitors for your Acquisition Units, as well as two monitors to alert you to data collector failures. These can be used to alert you quickly to data collection issues.

The following monitors are now available in the *Data Collection* monitor group:

- Acquisition Unit Heartbeat-Critical
- · Acquisition Unit Heartbeat-Warning
- · Collector Failed
- Collector Warning

Note that these monitors are in *Paused* state by default. Activate them to be alerted about data collection issues.

Auto-Renewing API Tokens

API Access Tokens can now be set for auto-renewal. By enabling this feature, new/refreshed API Access Tokens will automatically be generated for expiring tokens. Cloud Insights agents using an expiring token will automatically be updated to use the corresponding new/refreshed API Access Token, allowing them to continue to operate seamlessly. Simply check the "Renew token automatically" box when creating your token. This feature is currently supported on Cloud Insights agents running on the Kubernetes platform with the latest NetApp Kubernetes Monitoring Operator.

Basic Edition gives you more than before

Your trial is ending but you're not yet sure whether a subscription is right for you? Basic Edition has always given you a chance to continue using Cloud Insights with your current ONTAP data collector, but now you can continue capturing VMWare version, topology, and IOPS/Throughput/Latency data as well. NetApp customers with premium support on their storage systems will also be entitled to support for Cloud Insights.

Ready to learn more?

Check out the **Learning Center** section of the Help > Support page for links to NetApp University Cloud Insights course offerings!

Operating System support

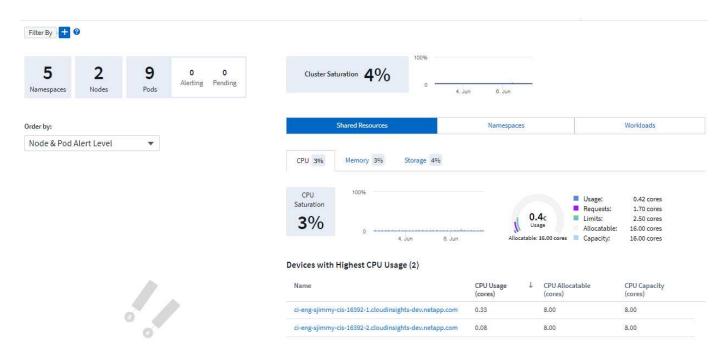
The following operating system is supported with Cloud Insights Acquisition Units, in addition to those already supported:

• Windows 11

June 2022

Kubernetes cluster saturation and other details

Cloud Insights makes it easier than ever to explore your Kubernetes environment, with an improved cluster detail page that provides Saturation details as well as a cleaner view into Namespaces and Workloads.



The Cluster list page also gives you a quick view of saturation, in addition to Node, Pod, Namespace, and Workload counts:



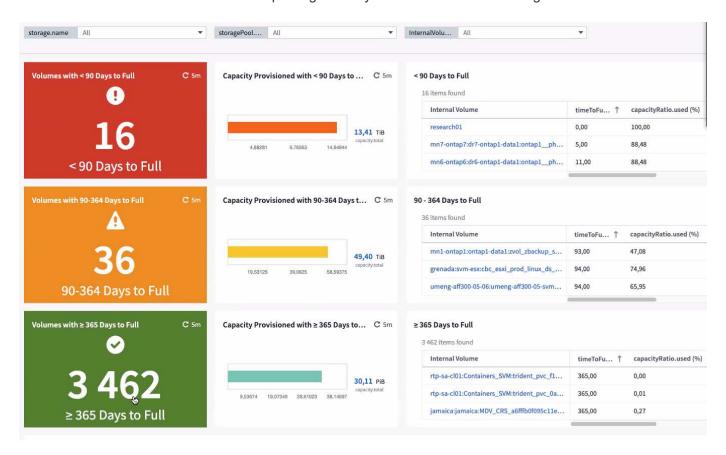
How old is your Kubernetes cluster?

Is your cluster just starting in the world, or has it experienced a long digital life? *Age* has been added as a time metric collected for Kubernetes Nodes.



Capacity Time-to-Full forecasting

Cloud Insights provides a dashboard to forecast the number of days until capacity runs out for each Internal Volume monitored. These values can help to significantly reduce the risk of an outage.

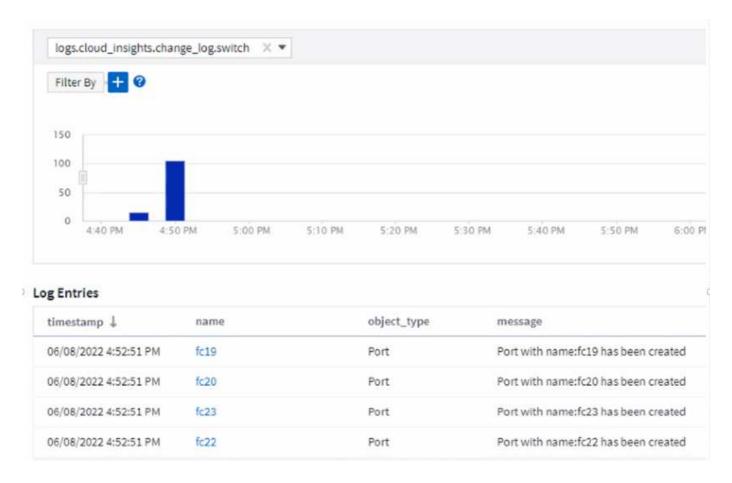


TTF counters are also available for Storage, Storage Pool, and Volume. Keep watching this space for additional dashboards for these objects.

Note that Time-to-Full forecasting is moving out of *Preview* and will be rolled out to all customers.

What's changed in my environment?

ONTAP change log entries can be viewed in the log explorer.



Operating System support

The following operating systems are supported with Cloud Insights Acquisition Units, in addition to those already supported:

- · CentOS Stream 9
- Windows 2022

Updated Telegraf Agent

The agent for ingestion of telegraf integration data has been updated to version **1.22.3**, with performance and security improvements.

Users wishing to update can refer to the appropriate upgrade section of the Agent Installation documentation. Previous versions of the agent will continue to function with no user action required.

Preview Features

Cloud Insights regularly highlights a number of exciting new preview features. If you are interested in previewing one or more of these features, contact your NetApp Sales Team for more information.

Feature	Description
Kubernetes Namespaces Running out of Space	The Kubernetes Namespaces Running Out of Space Insight gives you a view into workloads on your Kubernetes namespaces that are at risk of running out of space, with an estimate for the number of days remaining before each space becomes full. Read More
Cloud Secure – block user access on attack	Greater protection for your business-critical data with the ability to block user access when an attack is detected. Access can be blocked automatically, using Automated Response Policies or manually from the alert or user details pages. Read More
Shared Resource Under Stress	The Shared Resource Under Stress insight uses AI/ML to automatically identify where resource contention is causing performance degradation in your environment, highlights any workloads impacted by it, and provides recommended actions to remediate, letting you solve performance issues more quickly. Read More

May 2022

Chat live with NetApp Support

You can now chat live with NetApp Support personnel! On the Help > Support page, simply click the Chat icon or click *Chat* in the "Contact Us" section to start a chat session. Chat support is available US weekdays for Standard and Premium Edition users.



Kubernetes Operator

We've made it easier to get you up and running with Cloud Insights' advanced Kubernetes monitoring and cluster explorer.

The NetApp Kubernetes Monitoring Operator (NKMO) is the preferred method for installing Kubernetes for Cloud Insights Insights, for more flexible configuration of monitoring in fewer steps, as well as enhanced opportunities for monitoring other software running in the K8s cluster.

Click the link above for more information and pre-requisites

Manage Users and Invites with API

You can now manage users and invites using Cloud Insights' powerful API. Read more in the API Swagger Documentation.

Data Collection Alerts

Don't miss out on critical metrics due to a failed collector!

It's easier than ever to keep track of your data collectors with new alerts for data collector and acquisition unit failures.

Note that these Monitors are *Paused* by default. To enable, navigate to your monitors page and locate and resume "Acquisition Unit Shutdown" and "Collector Failed"

Alert on ONTAP storage changes

Don't let unexpected storage changes lead to outages!

You can now configure Cloud Insights to alert when modification or removal of FlexVols, nodes and SVMs are detected on ONTAP systems.

Preview Features

Cloud Insights regularly highlights a number of exciting new preview features. If you are interested in previewing one or more of these features, contact your NetApp Sales Team for more information.

Feature	Description
Kubernetes Namespaces Running out of Space	The Kubernetes Namespaces Running Out of Space Insight gives you a view into workloads on your Kubernetes namespaces that are at risk of running out of space, with an estimate for the number of days remaining before each space becomes full. Read More
Internal Volume and Volume Capacity Time-to-Full forecasting	Cloud Insights is able to prognose the number of days until capacity runs out for each Internal Volume and Volume monitored. This value can help to significantly reduce the risk of an outage.
Cloud Secure – block user access on attack	Greater protection for your business-critical data with the ability to block user access when an attack is detected. Access can be blocked automatically, using Automated Response Policies or manually from the alert or user details pages. Read More
Shared Resource Under Stress	The Shared Resource Under Stress insight uses AI/ML to automatically identify where resource contention is causing performance degradation in your environment, highlights any workloads impacted by it, and provides recommended actions to remediate, letting you solve performance issues more quickly. Read More

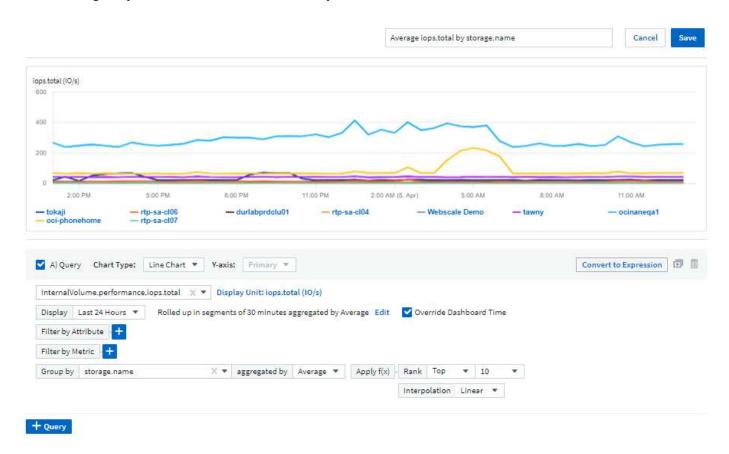
April 2022

Share your Feedback!

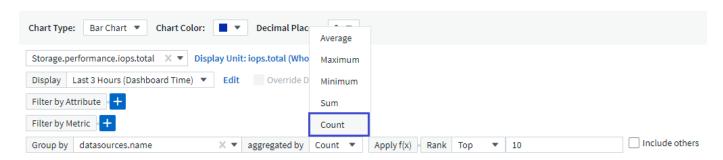
We want your input to help shape Cloud Insights. Earn points and prizes by participating in NetApp's **Insights to Action** program. **Sign up now!**

Updated Dashboard Editor

We've overhauled our dashboard creation tools to make it easier for you to visualize your data even more quickly. Navigate to the "Dashboards" page of Cloud Insights to edit an existing dashboard, add one from our dashboard gallery, or create a new dashboard of your own to check it out.



A new Count aggregation method has also been introduced. When grouping data in bar chart, column chart, and pie chart widgets, you can quickly and easily show the number of relevant objects for the selected metric.



Additionally, line charts now allow you to select one of three interpolation methods:

- · None No interpolation is done
- · Linear Interpolates a data point between the existing points
- · Stair Uses the previous data point as the interpolated data point

Enhanced Monitoring for Your Kubernetes Infrastructure

Cloud Insights keeps you on top of changes in your Kubernetes environment by alerting you when pods, daemonsets, and replicasets are created or removed, as well as when new deployments are created. Kubernetes monitors default to *paused* state, so you should enable only the specific ones you need.

Preview Features

Cloud Insights regularly highlights a number of exciting new preview features. If you are interested in previewing one or more of these features, contact your NetApp Sales Team for more information.

Feature	Description
Internal Volume and Volume Capacity Time-to-Full forecasting	Cloud Insights is able to prognose the number of days until capacity runs out for each Internal Volume and Volume monitored. This value can help to significantly reduce the risk of an outage.
Cloud Secure – block user access on attack	Greater protection for your business-critical data with the ability to block user access when an attack is detected. Access can be blocked automatically, using Automated Response Policies or manually from the alert or user details pages. Read More
Shared Resource Under Stress	The Shared Resource Under Stress insight uses AI/ML to automatically identify where resource contention is causing performance degradation in your environment, highlights any workloads impacted by it, and provides recommended actions to remediate, letting you solve performance issues more quickly. Read More

New Data Collector

• Cohesity SmartFiles - This REST API-based collector will acquire a Cohesity cluster, discovering the "Views" (as CI Internal Volumes), the various nodes, as well as collecting performance metrics.

Other Data Collector Updates

Collection and display of performance data has been improved on the following data collectors:

- Brocade CLI
- Dell/EMC VPlex, PowerStore, Isilon/PowerScale, VNX Block/Clariion CLI, XtremIO, Unity/VNXe
- Pure FlashArray

These performance enhancements are already available in all NetApp data collectors as well as VMware and

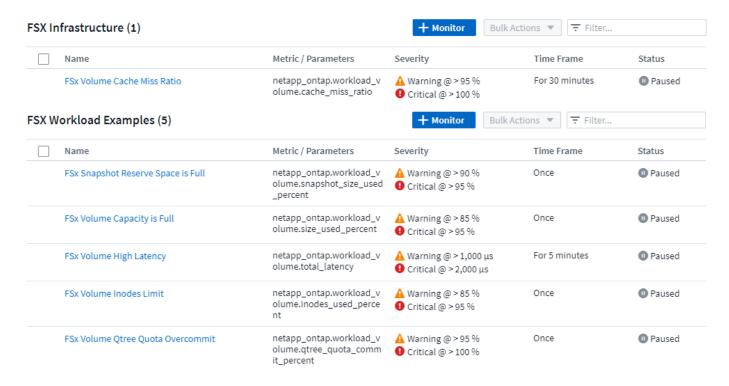
March 2022

Cloud Connection for ONTAP 9.9+

The NetApp Cloud Connection for ONTAP 9.9+ data collector eliminates the need to install an external acquisition unit, thereby simplifying troubleshooting, maintenance, and initial deployment.

New FSx for NetApp ONTAP Monitors

Monitoring your FSx for NetApp ONTAP environment is easy with new system-defined monitors for both infrastructure (metrics) and workloads (logs).



New Cloud Secure features available to all

Your environment is more secure than ever with the following Cloud Secure features now generally available:

Feature	Description
Data Destruction – File Deletion attack detection	Detect abnormal large-scale file deletion activity, block malicious file access by malicious users, and take automatic snapshots with automatic response policies.
Separate notifications for Warnings and Alerts	Warning and Alert notifications can be sent to separate recipients, ensuring the right team can stay informed

Updated Telegraf Agent

The agent for ingestion of telegraf integration data has been updated to version **1.21.2**, with performance and security improvements.

Users wishing to update can refer to the appropriate upgrade section of the Agent Installation documentation. Previous versions of the agent will continue to function with no user action required.

Data Collector Updates

• The Broadcom Fibre Channel Switches data collector has been optimized to reduce the number of CLI commands issued with each inventory poll.

February 2022

Cloud Insights addresses Apache Log4j vulnerabilities

Customer security is a top priority at NetApp. Cloud Insights includes updates to its software libraries to address the recent Apache Log4j vulnerabilities.

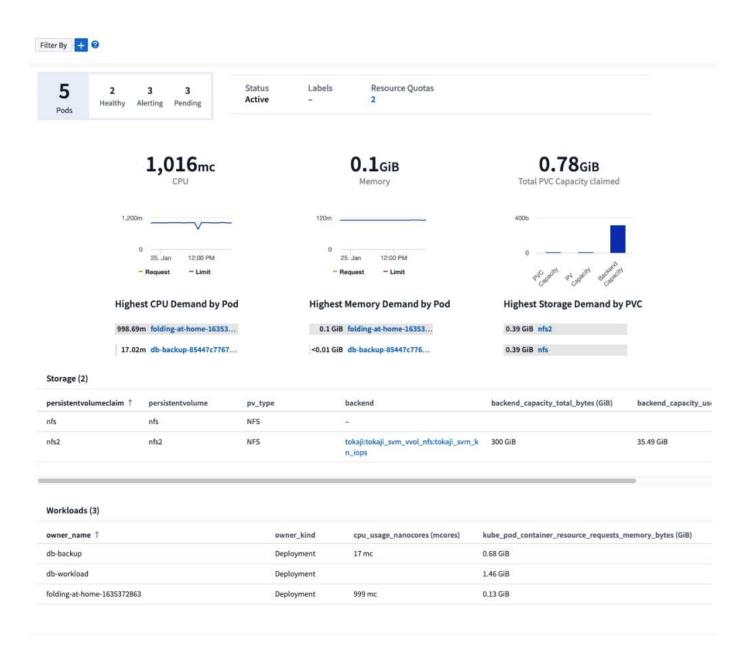
Please refer to the following on NetApp's Product Security Advisory website:

CVE-2021-44228 CVE-2021-45046 CVE-2021-45105

You can read more about these vulnerabilities and NetApp's response at the NetApp Newsroom.

Kubernetes Namespace Detail Page

Exploring your Kubernetes environment is now better than ever, with informative detail pages for your cluster's namespaces. The namespace detail page provides a summary of all the assets used by a namespace, including all the backend storage resources and their capacity utilizations.



December 2021

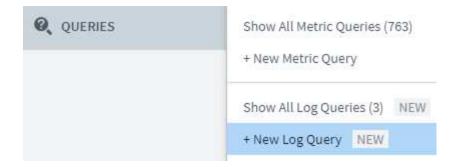
Deeper integration for ONTAP systems

Simplify alerting for ONTAP hardware failures and more with new integration with NetApp Event Management System (EMS).

Explore and alert on low-level ONTAP messages in Cloud Insights to inform and improve troubleshooting workflows and further reduce reliance on ONTAP element management tooling.

Querying Logs

For ONTAP systems, Cloud Insights Queries include a powerful Log Explorer, allowing you to easily investigate and troubleshoot EMS log entries.



Data Collector-level notifications.

In addition to system-defined and custom-created Monitors for alerting, you can also set alert notifications for ONTAP data collectors, allowing you to specify recipients for collector-level alerts, independent of other monitor alerts.

Greater flexibility of Cloud Secure roles

Users can be granted access to Cloud Secure features based on roles set by an administrator:

Role	Cloud Secure Access
Administrator	Can perform all Cloud Secure functions, including those for Alerts, Forensics, data collectors, automated response policies, and APIs for Cloud Secure. An Administrator can also invite other users but can only assign Cloud Secure roles.
User	Can view and manage Alerts and view Forensics. User role can change alert status, add a note, take snapshots manually, and block user access.
Guest	Can view Alerts and Forensics. Guest role cannot change alert status, add a note, take snapshots manually, or block user access.

Operating System support

CentOS 8.x support is being replaced with **CentOS 8 Stream** support. CentOS 8.x will reach End-of-Life on December 31, 2021.

Data Collector Updates

A number of Cloud Insights data collector names have been added to reflect vendor changes:

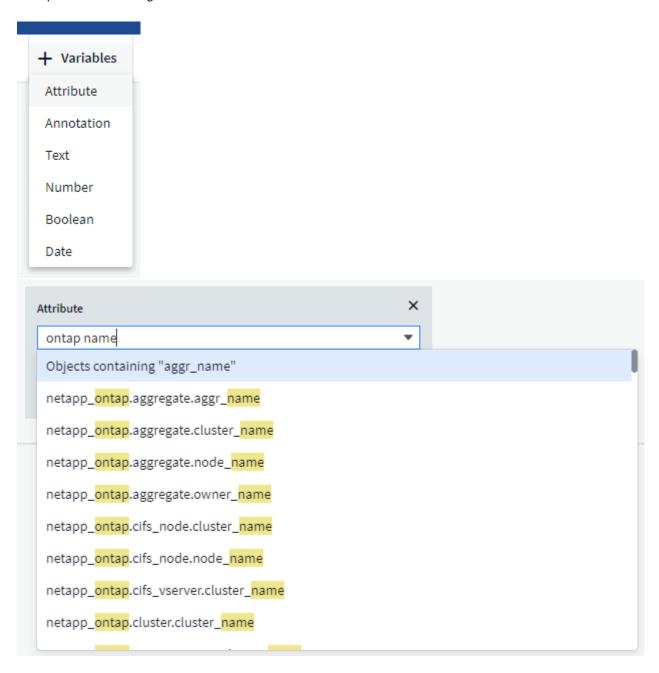
Vendor/Model	Previous Name
Dell EMC PowerScale	Isilon
HPE Alletra 9000 / Primera	3PAR
HPE Alletra 6000	Nimble

November 2021

Adaptive Dashboards

New variables for attributes and the ability to use variables in widgets.

Dashboards are now more powerful and flexible than ever. Build adaptive dashboards with attribute variables to quickly filter dashboards on the fly. Using these and other pre-existing variables you can now create one high level dashboard to see metrics for your entire environment, and seamlessly filter down by resource name, type, location, and more. Use number variables in widgets to associate raw metrics with costs, for example cost per GB for storage as a service.



Access the Reporting Database via API

Enhanced capabilities for integration with third party reporting, ITSM, and automation tools: Cloud Insights' powerful API allows users to query the Cloud Insights Reporting database directly, without going through the

Cognos Reporting environment.

Pod tables on VM Landing Page

Seamless navigation between VMs and the Kubernetes Pods using them: for improved troubleshooting and performance headroom management, a table of associated Kubernetes Pods will now appear on VM landing pages.



Data Collector Updates

- · ECS now reports firmware for storage and node
- · Isilon has improved prompt detection
- Azure NetApp Files collects performance data more quickly
- StorageGRID now supports Single Sign-On (SSO)
- Brocade CLI properly reports model for X&-4

Additional Operating Systems supported

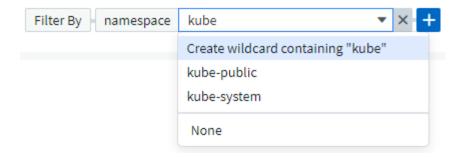
The Cloud Insights Acquisition Unit supports the following operating systems, in addition to those already supported:

- Centos (64-bit) 8.4
- Oracle Enterprise Linux (64-bit) 8.4
- Red Hat Enterprise Linux (64-bit) 8.4

October 2021

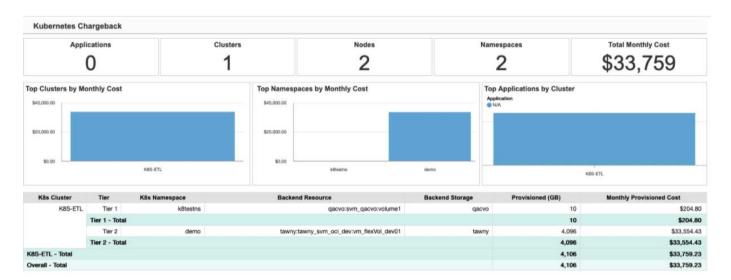
Filters on K8S Explorer pages

Kubernetes Explorer page filters give you focused control of the data displayed for your Kubernetes cluster, node, and pod exploration.



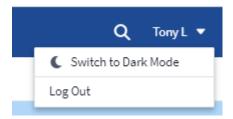
K8s Data for Reporting

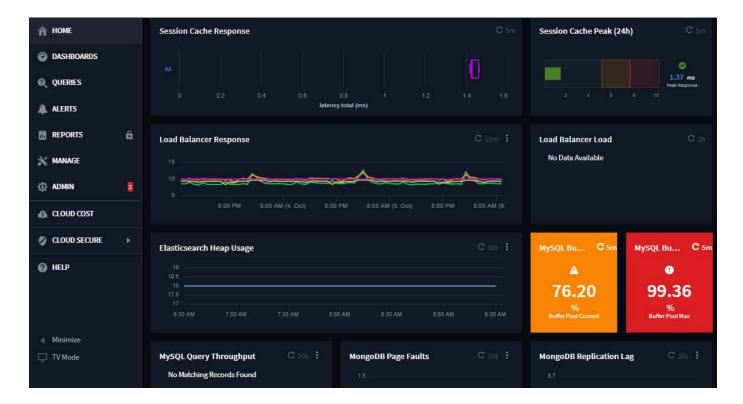
Kubernetes data is now available for use in Reporting, allowing you to create chargeback or other reports. In order for Kubernetes chargeback data to be passed to Reporting, you must have an active connection to, and Cloud Insights must be receiving data from, your Kubernetes cluster as well as its back-end storage. If there is no data received from the back-end storage, Cloud Insights can not send Kubernetes object data to Reporting.



Dark Theme has arrived

Many of you asked for a dark theme, and Cloud Insights has answered. To switch between light and dark theme, click the drop-down next to your user name.





Data Collector Support

We've made some improvements in Cloud Insights Data Collectors. Here are some highlights:

· New collector for Amazon FSx for ONTAP

September 2021

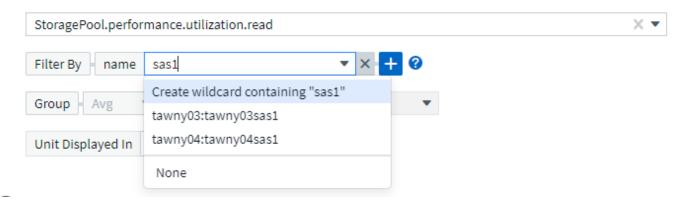
Performance Policies are now Monitors

Monitors and Alerts have supplanted Performance Policies and Violations throughout Cloud Insights. Alerting with Monitors provides greater flexibility and insight into potential problems or trends in your environment.

Autocomplete Suggestions, Wildcards, and Expressions in Monitors

When creating a monitor for alerting, typing in a filter is now predictive, allowing you to easily search for and find the metrics or attributes for your monitor. Additionally, you are given the option to create a wildcard filter based on the text you type.

Select a metric to monitor



Updated Telegraf Agent

The agent for ingestion of telegraf integration data has been updated to version **1.19.3**, with performance and security improvements.

Users wishing to update can refer to the appropriate upgrade section of the Agent Installation documentation. Previous versions of the agent will continue to function with no user action required.

Data Collector Support

We've made some improvements in Cloud Insights Data Collectors. Here are some highlights:

- Microsoft Hyper-V collector now uses PowerShell instead of WMI
- Azure VMs and VHD collector is now up to 10 times faster due to parallel calls
- HPE Nimble now supports federated and iSCSI configurations

And since we're always improving Data Collection, here are some other recent changes of note:

- New collector for EMC Powerstore
- · New collector for Hitachi Ops Center
- New collector for Hitachi Content Platform
- Enhanced ONTAP collector to report Fabric Pools
- Enhanced ANF with Storage Pool and Volume performance
- Enhanced EMC ECS with Storage Nodes and Storage performance as well as the Object Count in buckets
- Enhanced EMC Isilon with Storage Node and Qtree metrics
- Enhanced EMC Symetrix with volume QOS limit metrics
- Enhanced IBM SVC and EMC PowerStore with Storage Nodes parent serial number

August 2021

New Audit Page User Interface

The Audit page provides a cleaner interface and now allows the export of audit events to .CSV file.

Enhanced User Role Management

Cloud Insights now allows even greater freedom for assigning user roles and access controls. Users can now be assigned granular permissions for monitoring, reporting, and Cloud Secure separately.

This means you can allow more users administrative access to monitoring, optimization, and reporting functions whilst restricting access to your sensitive Cloud Secure audit and activity data to only those that need it.

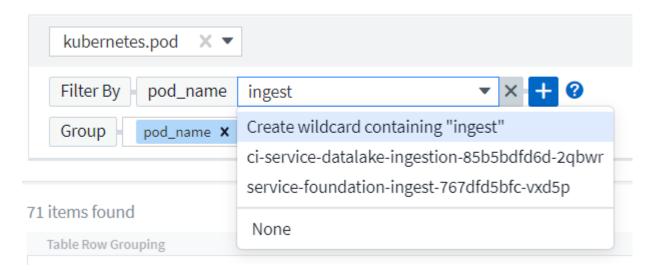
Find out more about the different levels of access in the Cloud Insights documentation.

June 2021

Autocomplete Suggestions, Wildcards, and Expressions in Filters

With this release of Cloud Insights, you no longer need to know all the possible names and values on which to filter in a query or widget. When filtering, you can simply start typing and Cloud insights will suggest values based on your text. No more looking up Application names or Kubernetes attributes ahead of time just to find the ones you want to show in your widget.

As you type in a filter, the filter displays a smart list of results from which you can choose, as well as the option to create a **wildcard filter** based on the current text. Selecting this option will return all results that match the wildcard expression. You can of course also select multiple individual values that you want added to the filter.



Additionally, you can create **expressions** in a filter using NOT or OR, or you can select the "None" option to filter for null values in the field.

Read more about filtering options in queries and widgets.

APIs available by Edition

Cloud Insights' powerful APIs are more accessible than ever, with Alerts APIs now available in Standard and Premium Editions.

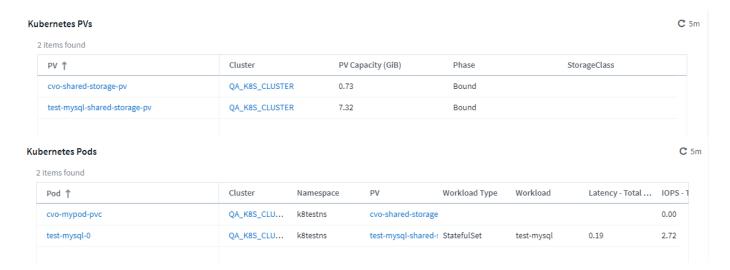
The following APIs are available for each Edition:

API Category	Basic	Standard	Premium
Acquisition Unit	✓	✓	✓
Data Collection	✓	✓	✓
Alerts		✓	✓
Assets		✓	✓
Data Ingestion		✓	✓

Kubernetes PV and Pod Visibility

Cloud Insights provides visibility into the back-end storage for your Kubernetes environments, giving you insight to your Kubernetes Pods and Persistent Volumes (PVs). You can now track PV counters such as IOPS, latency, and throughput from a single Pod's usage through a PV counter to a PV and all the way to the back-end storage device.

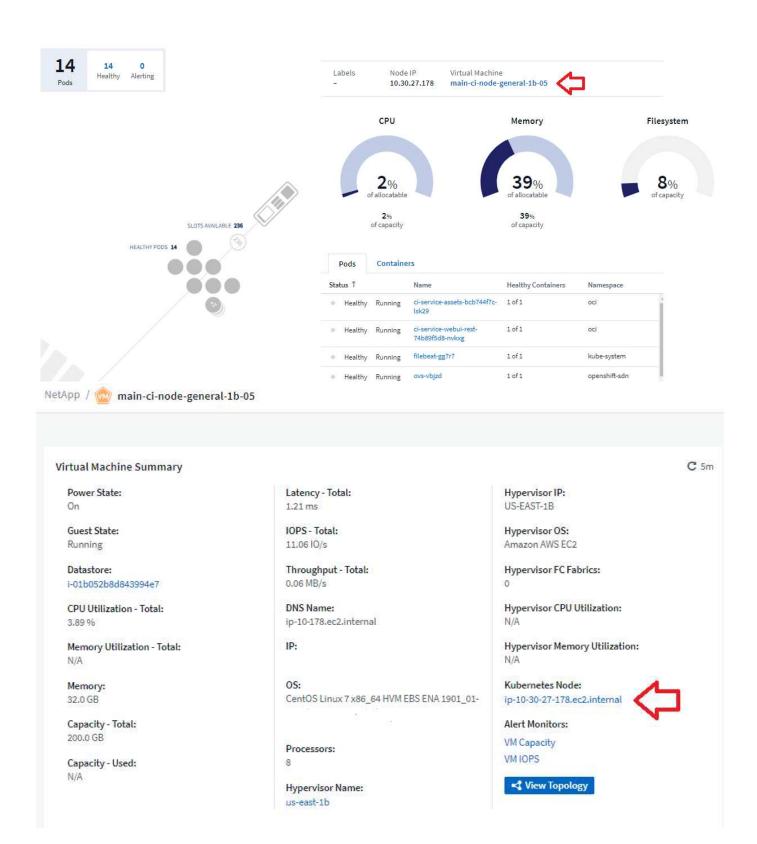
On a Volume or Internal Volume landing page, two new tables are displayed:



Note that to take advantage of these new tables, it is recommended to uninstall your current Kubernetes agent, and install it fresh. You must also install Kube-State-Metrics version 2.1.0 or later.

Kubernetes Node to VM links

On a Kubernetes Node page, you can now click to open the Node's VM page. The VM page also includes a link back to the Node itself.



Alert Monitors replacing Performance Policies

To enable the added benefits of multiple thresholds, webhook and email alert delivery, alerting on all metrics using a single interface, and more, Cloud Insights will be converting Standard and Premium Edition customers from **Performance Policies** to **Monitors** during the months of July and August, 2021. Learn more about Alerts and Monitors, and stay tuned for this exciting change.

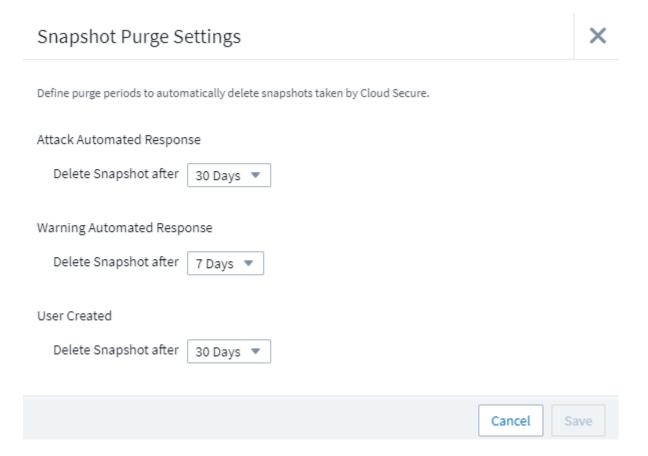
Cloud Secure supports NFS

Cloud Secure now supports NFS for ONTAP data collection. Monitor SMB and NFS user access to protect your data from ransomware attacks.

Additionally, Cloud Secure supports Active-Directory and LDAP user directories for collection of NFS user attributes.

Cloud Secure snapshot purge

Cloud Secure automatically deletes snapshots based on the Snapshot Purge Settings, to save storage space and reduce the need for manual snapshot deletion.



Cloud Secure data collection speed

A single data collector agent system can now post up to 20,000 events per second to Cloud Secure.

May 2021

Here are some of the changes we've made in April:

Updated Telegraf Agent

The agent for ingestion of telegraf integration data has been updated to version 1.17.3, with performance and security improvements.

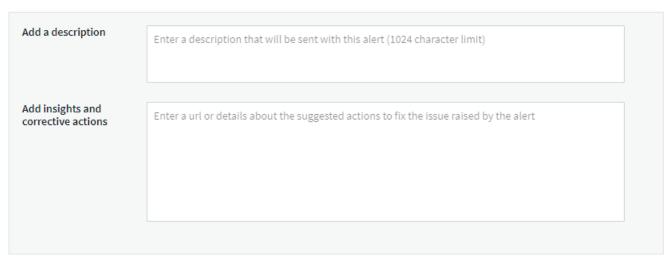
Users wishing to update can refer to the appropriate upgrade section of the Agent Installation documentation.

Previous versions of the agent will continue to function with no user action required.

Add Corrective Actions to an Alert

You can now add an optional description as well as additional insights and/or corrective actions when creating or modifying a Monitor by filling in the **Add an Alert Description** section. The description will be sent with the alert. The *insights and corrective actions* field can provide detailed steps and guidance for dealing with alerts and will be displayed in the summary section of the alert landing page.





Cloud Insights APIs for All Editions

API access is now available in all editions of Cloud Insights.

Users of Basic edition can now automate actions for Acquisition Units and Data Collectors, and Standard Edition users can query metrics and ingest custom metrics.

Premium edition continues to allow full use of all API categories.

API Category	Basic	Standard	Premium
Acquisition Unit	✓	✓	✓
Data Collection	✓	✓	✓
Assets		✓	✓
Data Ingestion		✓	✓
Data Warehouse			✓

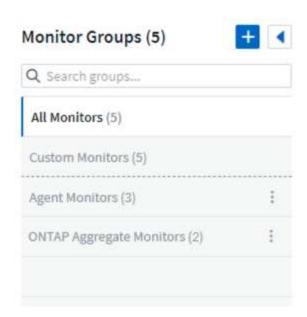
For details on API usage, please refer to the API documentation.

April 2021

Easier Management of Monitors

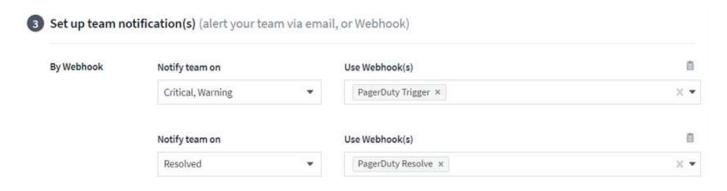
Monitor Grouping simplifies the management of monitors in your environment. Multiple monitors can now be grouped together and paused as one. For example, if you have an update occurring on a stack of infrastructure, you can pause alerts from all those devices via one click.

Monitor groups is the first part of an exciting new feature bringing improved management of ONTAP devices to Cloud Insights.



Enhanced Alerting Options Using Webhooks

Many commercial applications support Webhooks as a standard input interface. Cloud Insights now supports many of these delivery channels, providing default templates for Slack, PagerDuty, Teams, and Discord, in addition to providing customizable generic webhooks to support many other applications.



Improved Device Identification

To improve monitoring and troubleshooting as well as deliver accurate reporting, it is helpful to understand the names of devices rather than their IP addresses or other identifiers. Cloud Insights now incorporates an automatic way to identify the names of storage and physical host devices in the environment, using a rule-based approach called **Device Resolution**, available in the **Manage** menu.

You asked for more!

A popular ask by customers has been for more default options for visualizing the range of data, so we have added the following five new choices that are now available throughout the service via the time range picker:

- Last 30 Minutes
- Last 2 Hours
- · Last 6 Hours
- · Last 12 Hours
- · Last 2 Days

Multiple Subscriptions in one Cloud Insights Environment

Starting April 2, Cloud Insights supports multiple subscriptions of the same edition type for a customer in a single Cloud Insights instance. This enables customers to co-term parts of their Cloud Insights subscription with infrastructure purchases. Contact NetApp Sales for assistance with multiple subscriptions.

Choose Your Path

While setting up Cloud Insights, you can now choose whether to start with Monitoring and Alerting or Ransomware and Insider Threat Detection. Cloud Insights will configure your starting environment based on the path you choose. You can configure the other path at any time afterward.

Easier Cloud Secure Onboarding

And it is easier than ever to start using Cloud Secure, with a new step-by-step setup checklist.



CLOUD SECURE

Secure Your Data from Ransomware & Insider Threat

- · Ransomware & insider threat detection
- · User data access auditing

Setting up Cloud Secure

- Add an Agent on server or VM to collect data (system requirements □).
- Configure a User Directory Collector to collect user attributes from active directories (optional step).
- Configure a Data Collector to collect file access activity on your storage devices.
- Define Automated Response Policies to take automatic action in the event of an attack.

User activity data will appear in the Forensics section

As always, we love to hear your suggestions! Send them to ng-cloudinsights-customerfeedback@netapp.com.

February 2021

Updated Telegraf Agent

The agent for ingestion of telegraf integration data has been updated to version 1.17.0, which includes vulnerability and bug fixes.

Cloud Cost Analyzer

Experience the power of Spot by NetApp with Cloud Cost, which provides a detailed cost analysis of past, present, and estimated spending, providing visibility into cloud usage in your environment. The Cloud Cost dashboard delivers a clear view of cloud expenses and a drill down into individual workloads, accounts, and services.

Cloud Cost can help with these major challenges:

- Tracking and monitoring your cloud expenses
- · Identifying waste and potential optimization areas
- Delivering executable action items

Cloud Cost is focused on monitoring. Upgrade to the full Spot by NetApp account to enable automatic cost

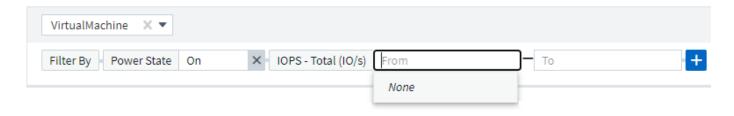
saving and environment optimization.

Querying for objects having null values using filters

Cloud Insights now allows searching for attributes and metrics having null/none values through the use of filters. You can perform this filtering on any attributes/metrics in the following places:

- · On the Query page
- In Dashboard widgets and page variables
- · On the Alerts list page
- · When creating Monitors

To filter for null/none values, simply select the *None* option when it appears in the appropriate filter drop-down.



Multi-Region Support

Starting today we offer the Cloud Insights service in different regions across the globe, which facilitates performance and increases security for customers based outside the United States. Cloud Insights/Cloud Secure stores information according to the region in which your environment is created.

Click here for more information.

January 2021

Additional ONTAP Metrics Renamed

As part of our continuing effort to improve efficiency of data-gathering from ONTAP systems, the following ONTAP metrics have been renamed.

If you have existing dashboard widgets or queries using any of these metrics, you will need to edit or re-create them to use the new metric names.

Previous Metric Name	New Metric Name
netapp_ontap.disk_constituent.total_transfers	netapp_ontap.disk_constituent.total_iops
netapp_ontap.disk.total_transfers	netapp_ontap.disk.total_iops
netapp_ontap.fcp_lif.read_data	netapp_ontap.fcp_lif.read_throughput
netapp_ontap.fcp_lif.write_data	netapp_ontap.fcp_lif.write_throughput
netapp_ontap.iscsi_lif.read_data	netapp_ontap.iscsi_lif.read_throughput
netapp_ontap.iscsi_lif.write_data	netapp_ontap.iscsi_lif.write_throughput

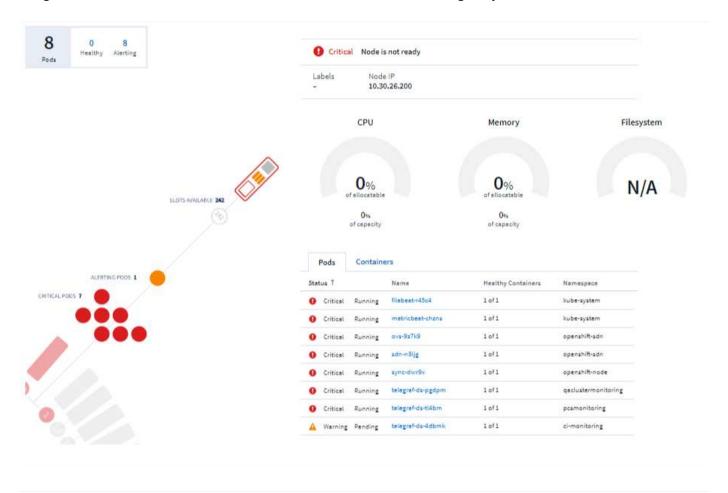
Previous Metric Name	New Metric Name
netapp_ontap.lif.recv_data	netapp_ontap.lif.recv_throughput
netapp_ontap.lif.sent_data	netapp_ontap.lif.sent_throughput
netapp_ontap.lun.read_data	netapp_ontap.lun.read_throughput
netapp_ontap.lun.write_data	netapp_ontap.lun.write_throughput
netapp_ontap.nic_common.rx_bytes	netapp_ontap.nic_common.rx_throughput
netapp_ontap.nic_common.tx_bytes	netapp_ontap.nic_common.tx_throughput
netapp_ontap.path.read_data	netapp_ontap.path.read_throughput
netapp_ontap.path.write_data	netapp_ontap.path.write_throughput
netapp_ontap.path.total_data	netapp_ontap.path.total_throughput
netapp_ontap.policy_group.read_data	netapp_ontap.policy_group.read_throughput
netapp_ontap.policy_group.write_data	netapp_ontap.policy_group.write_throughput
netapp_ontap.policy_group.other_data	netapp_ontap.policy_group.other_throughput
netapp_ontap.policy_group.total_data	netapp_ontap.policy_group.total_throughput
netapp_ontap.system_node.disk_data_read	netapp_ontap.system_node.disk_throughput_read
netapp_ontap.system_node.disk_data_written	netapp_ontap.system_node.disk_throughput_written
netapp_ontap.system_node.hdd_data_read	netapp_ontap.system_node.hdd_throughput_read
netapp_ontap.system_node.hdd_data_written	netapp_ontap.system_node.hdd_throughput_written
netapp_ontap.system_node.ssd_data_read	netapp_ontap.system_node.ssd_throughput_read
netapp_ontap.system_node.ssd_data_written	netapp_ontap.system_node.ssd_throughput_written
netapp_ontap.system_node.net_data_recv	netapp_ontap.system_node.net_throughput_recv
netapp_ontap.system_node.net_data_sent	netapp_ontap.system_node.net_throughput_sent
netapp_ontap.system_node.fcp_data_recv	netapp_ontap.system_node.fcp_throughput_recv
netapp_ontap.system_node.fcp_data_sent	netapp_ontap.system_node.fcp_throughput_sent
netapp_ontap.volume_node.cifs_read_data	netapp_ontap.volume_node.cifs_read_throughput
netapp_ontap.volume_node.cifs_write_data	netapp_ontap.volume_node.cifs_write_throughput
netapp_ontap.volume_node.nfs_read_data	netapp_ontap.volume_node.nfs_read_throughput
netapp_ontap.volume_node.nfs_write_data	netapp_ontap.volume_node.nfs_write_throughput
netapp_ontap.volume_node.iscsi_read_data	netapp_ontap.volume_node.iscsi_read_throughput
netapp_ontap.volume_node.iscsi_write_data	netapp_ontap.volume_node.iscsi_write_throughput
netapp_ontap.volume_node.fcp_read_data	netapp_ontap.volume_node.fcp_read_throughput
netapp_ontap.volume_node.fcp_write_data	netapp_ontap.volume_node.fcp_write_throughput
netapp_ontap.volume.read_data	netapp_ontap.volume.read_throughput
netapp_ontap.volume.write_data	netapp_ontap.volume.write_throughput

Previous Metric Name	New Metric Name
netapp_ontap.workload.read_data	netapp_ontap.workload.read_throughput
netapp_ontap.workload.write_data	netapp_ontap.workload.write_throughput
netapp_ontap.workload_volume.read_data	netapp_ontap.workload_volume.read_throughput
netapp_ontap.workload_volume.write_data	netapp_ontap.workload_volume.write_throughput

New Kubernetes Explorer

The Kubernetes Explorer provides a simple topology view of Kubernetes Clusters, allowing even non-experts to quickly identify issues & dependencies, from the cluster level down to the container and storage.

A wide variety of information can be explored using the Kubernetes Explorer's drill-down details for status, usage, and health of the Clusters, Nodes, Pods, Containers, and Storage in your Kubernetes environment.



December 2020

Simpler Kubernetes Installation

Kubernetes Agent installation has been streamlined to require fewer user interactions. Installing the Kubernetes Agent now includes Kubernetes data collection.

November 2020

Additional Dashboards

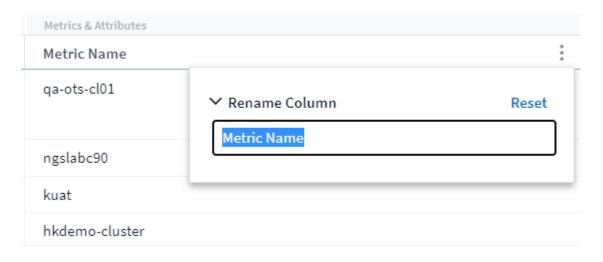
The following new ONTAP-focused dashboards have been added to the gallery and are available for import:

- ONTAP: Aggregate Performance & Capacity
- ONTAP FAS/AFF Capacity Utilization
- · ONTAP FAS/AFF Cluster Capacity
- ONTAP FAS/AFF Efficiency
- ONTAP FAS/AFF FlexVol Performance
- ONTAP FAS/AFF Node Operational/Optimal Points
- ONTAP FAS/AFF PrePost Capacity Efficiencies
- · ONTAP: Network Port Activity
- ONTAP: Node Protocols Performance
- ONTAP: Node Workload Performance (Frontend)
- ONTAP: Processor
- ONTAP: SVM Workload Performance (Frontend)
- ONTAP: Volume Workload Performance (Frontend)

Column Rename in Table Widgets

You can rename columns in the *Metrics and Attributes* section of a table widget by opening the widget in Edit mode and clicking the menu at the top of the column. Enter the new name and click *Save*, or click *Reset* to set the column back to the original name.

Note that this only affects the column's display name in the table widget; the metric/attribute name does not change in the underlying data itself.



October 2020

Default Expansion of Integration Data

Table widget grouping now allows for default expansions of Kubernetes, ONTAP Advanced Data, and Agent Node metrics. For example, if you group Kubernetes *Nodes* by *Cluster*, you will see a row in the table for each cluster. You could then expand each cluster row to see a list of the Node objects.

Basic Edition Technical Support

Technical Support is now available for subscribers to Cloud Insights Basic Edition in addition to Standard and Premium Editions. Additionally, Cloud Insights has simplified the workflow for creating a NetApp support ticket.

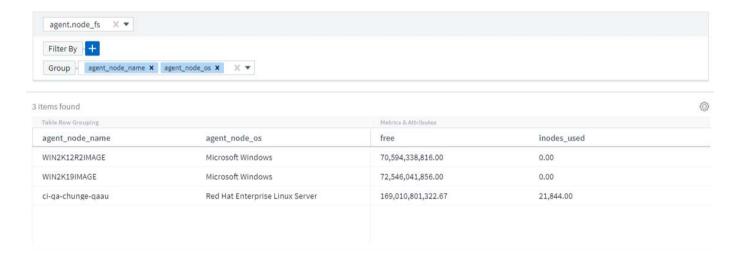
Cloud Secure Public API

Cloud Secure supports REST APIs for accessing Activity and Alert information. This is accomplished through the use of API Access Tokens, created through the Cloud Secure Admin UI, which are then used to access the REST APIs. Swagger documentation for these REST APIs is integrated with Cloud Secure.

September 2020

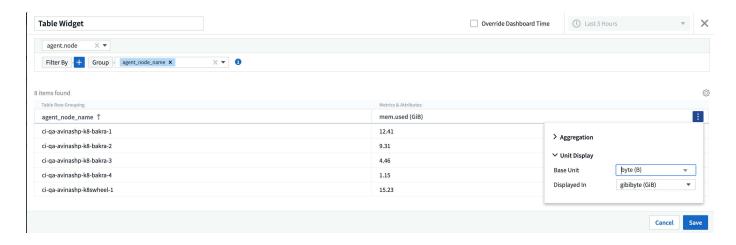
Query Page with Integration Data

The Cloud Insights Query page supports integration data (i.e. from Kubernetes, ONTAP Advanced Metrics, etc.). When working with integration data, the query results table displays a "Split-Screen" view, with object/grouping on the left side, and object data (attributes/metrics) on the right. You can also choose multiple attributes for grouping integration data.



Unit Display Formatting in Table Widget

Unit display formatting is now available in Table widgets for columns that display metric/counter data (for example, gigabytes, MB/second, etc.). To change a metric's display unit, click the "three dots" menu in the column header and select "Unit Display". You can choose from any of the available units. Available units will vary according to the type of metric data in the display column.



Acquisition Unit Detail Page

Acquisition Units now have their own landing page, providing useful detail for each AU as well as information to help with troubleshooting. The AU detail page provides links to the AU's data collectors as well as helpful status information.

Cloud Secure Docker Dependency Removed

Cloud Secure's dependency on Docker has been removed. Docker is no longer required for Cloud Secure agent installation.

Reporting User Roles

If you have Cloud Insights Premium Edition with Reporting, every Cloud Insights user in your environment also has a Single Sign-On (SSO) login to the Reporting application (i.e. Cognos); by clicking the **Reports** link in the menu, they will automatically be logged in to Reporting.

Their user role in Cloud Insights determines their Reporting user role:

Cloud Insights Role	Reporting Role	Reporting Permissions
Guest	Consumer	Can view, schedule, and run reports and set personal preferences such as those for languages and time zones. Consumers cannot create reports or perform administrative tasks.
User	Author	Can perform all Consumer functions as well as create and manage reports and dashboards.
Administrator	Administrator	Can perform all Author functions as well as all administrative tasks such configuration of reports and the shutdown and restart of reporting tasks.



Cloud Insights Reporting is available for environments of 500 MUs or more.



If you are a current Premium Edition customer and wish to retain your reports, read this important note for existing customers.

New API Category for Data Ingestion

Cloud Insights has added a **Data Ingestion** API category, giving you greater control over custom data and agents. Detailed documentation for this and other API Categories can be found in Cloud Insights by navigating to **Admin > API Access** and clicking the *API Documentation* link. You can also attach a comment to the AU in the Note field, which is displayed on the AU detail page as well as the AU list page.

August 2020

Monitoring and Alerting

In addition to the current ability to set performance policies for storage objects, VMs, EC2, and ports, Cloud Insights Standard Edition now includes the ability to configure monitors for thresholds on Integration data for Kubernetes, ONTAP advanced metrics, and Telegraf plugins. You simply create a monitor for each object metric you want to trigger alerts, set the conditions for warning-level or critical-level thresholds, and specify the email recipient(s) desired for each level. You can then view and manage alerts to track trends or troubleshoot issues.



July 2020

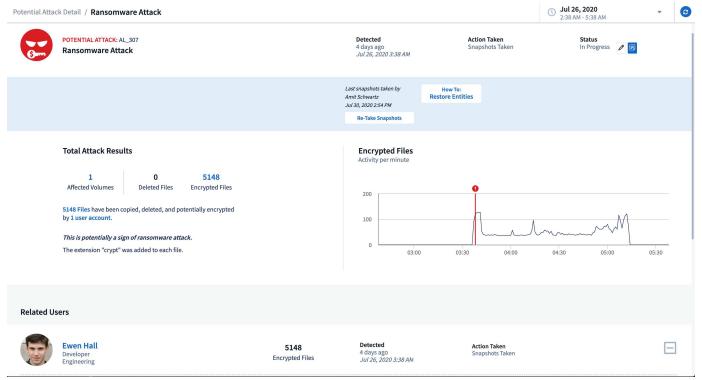
Cloud Secure Take a Snapshot Action

Cloud Secure protects your data by automatically taking a snapshot when malicious activity is detected, ensuring that your data is safely backed up.

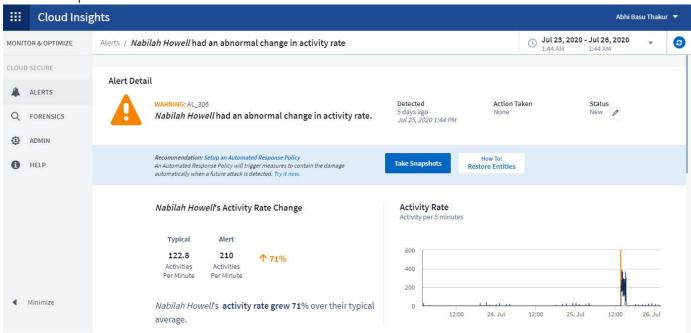
You can define automated response policies that take a snapshot when ransomware attack or other abnormal user activity is detected.

You can also take a snapshot manually from the alert page.

Automatic Snapshot taken:



Manual Snapshot:



Metric/Counter updates

The following capacity counters are available for use in Cloud Insights UI and REST API. Previously these counters were only available for the Data Warehouse / Reporting.

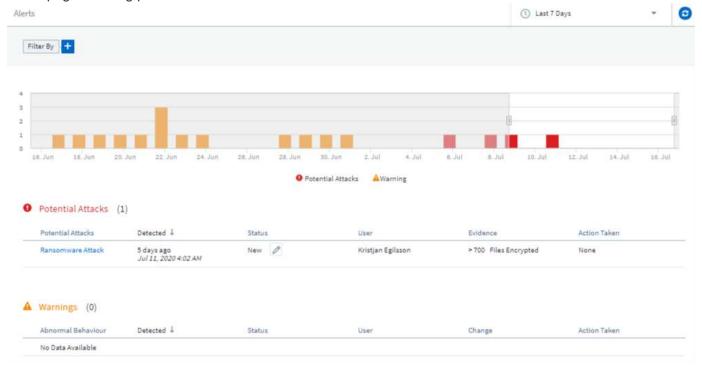
Object Type	Counter
Storage	Capacity - Spare Raw Capacity - Failed Raw
Storage Pool	Data Capacity - Used Data Capacity - Total Other Capacity - Used Other Capacity - Total Capacity - Raw Capacity - Soft Limit
Internal Volume	Data Capacity - Used Data Capacity - Total Other Capacity - Used Other Capacity - Total Clone Saved Capacity - Total

Cloud Secure Potential Attack Detection

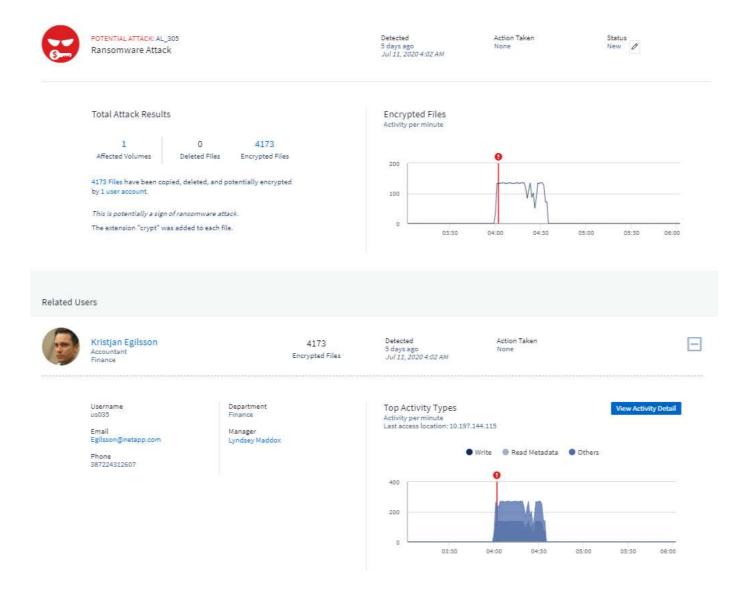
Cloud Secure now detects potential attacks such as ransomware. Click on an alert in the Alerts list page to open a detail page showing the following:

- · Time of attack
- · Associated user and file activity
- · Action taken
- · Additional information to assist with tracking down possible security breaches

Alerts page showing potential ransomware attack:



Detail page for potential ransomware attack:



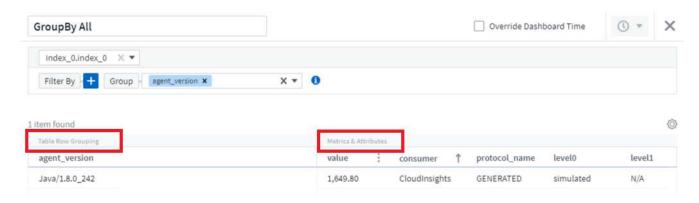
Subscribe to Premium Edition through AWS

During your trial of Cloud Insights, you can self-subscribe through AWS Marketplace to either Cloud Insights Standard Edition or Premium Edition. Previously, you could only self-subscribe through AWS Marketplace to Standard Edition only.

Enhanced Table Widget

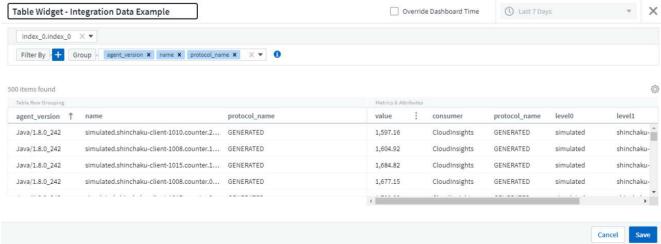
The dashboard/asset page Table widget includes the following enhancements:

• "Split-Screen" view: Table widgets display the object/grouping on the left side, and the object data (attributes/metrics) on the right.



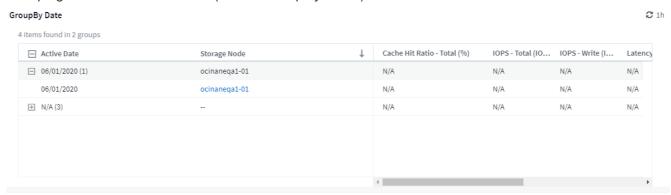
Multiple attribute grouping: For Integration data (Kubernetes, ONTAP Advanced Metrics, Docker, etc.), you
can choose multiple attributes for grouping. Data is displayed according to the grouping attributes/you
choose.





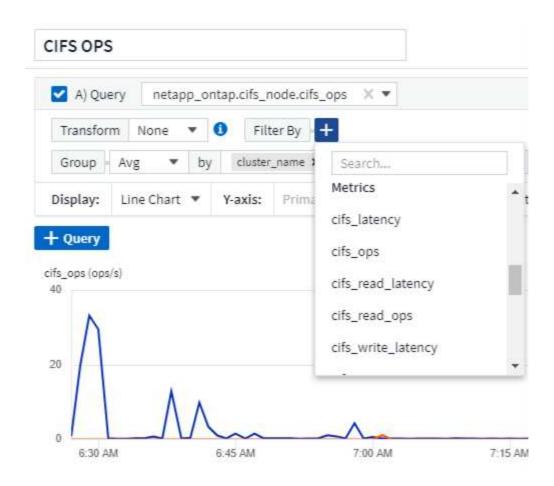
• Grouping for Infrastructure data (storage, EC2, VM, ports, etc.) is by a single attribute as before. When grouping by an attribute which is not the object, the table will allow you to expand the group row to see all the objects within the group.

Grouping with Infrastructure data (shown in display mode):

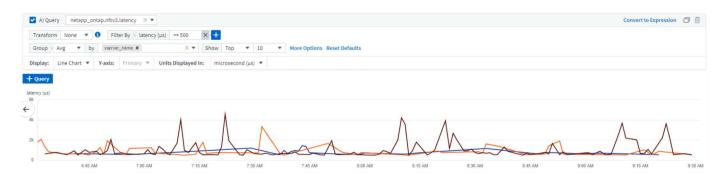


Metrics Filtering

In addition to filtering on an object's attributes in a widget, you can now filter on metrics as well.



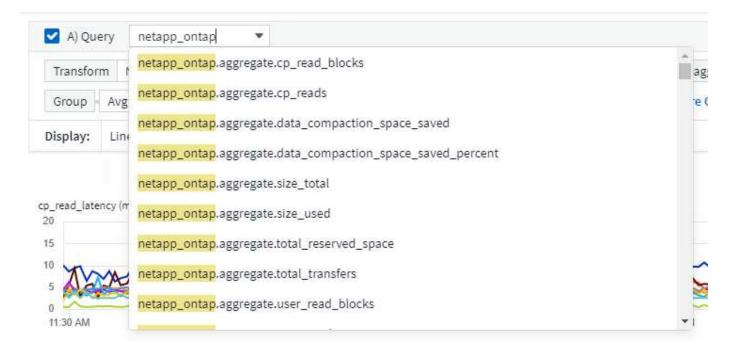
When working with integration data (Kubernetes, ONTAP Advanced Data, etc.), metric filtering removes the individual/unmatched data points from the plotted data series, unlike infrastructure data (storage, VM, ports etc.) where filters work on the aggregated value of the data series and potentially remove the entire object from the chart.



ONTAP Advanced Counter Data

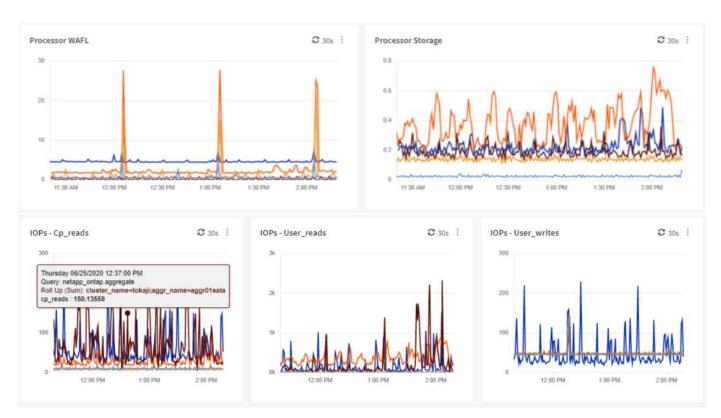
Cloud Insights takes advantage of NetApp's ONTAP-specific **Advanced Counter Data**, which provides a host of counters and metrics collected from ONTAP devices. ONTAP Advanced Counter Data is available to all NetApp ONTAP customers. These metrics enable customized and wide-ranging visualization in Cloud Insights widgets and dashboards.

ONTAP Advanced Counters can be found by searching for "netapp_ontap" in the widget's query, and selecting from among the counters.



You can refine your search by typing additional parts of the counter name. For example:

- lif
- · aggregate
- offbox_vscan_server
- · and more



Please note the following:

- Advanced Data collection will be enabled by default for new ONTAP data collectors. To enable Advanced
 Data collection for your existing ONTAP data collectors, edit the data collector and expand the Advanced
 Configuration section.
- Advanced Data collection is not available for 7-mode ONTAP.

Advanced Counter Dashboards

Cloud Insights comes with a variety of pre-designed dashboards to help get you started on visualizing ONTAP Advanced Counters for topics such as *Aggregate Performance*, *Volume Workload*, *Processor Activity*, and more. If you have at least one ONTAP data collector configured, these can be imported from the Dashboard Gallery on any dashboard list page.

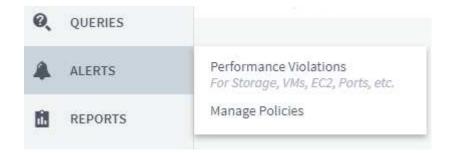
Learn More

More information on ONTAP Advanced Data can be found at the following links:

- https://mysupport.netapp.com/site/tools/tool-eula/netapp-harvest (Note: You will need to sign in to NetApp Support)
- https://nabox.org/faq/

Policies and Violations Menu

Performance Policies and Violations are now found under the **Alerts** menu. Policy and Violation functionality are unchanged.



Updated Telegraf Agent

The agent for ingestion of telegraf integration data has been updated to version 1.14, which includes bugs fixes, security fixes, and new plugins.

Note: When configuring a Kubernetes data collector on the Kubernetes platform, you may see an "HTTP status 403 Forbidden" error in the log, due to insufficient permissions in the "clusterrole" attribute.

To work around this issue, add the following highlighted lines to the *rules:* section of the endpoint-access clusterrole, and then restart the Telegraf pods.

```
rules:
- apiGroups:
  _ ""
  - apps
  - autoscaling
  - batch
  - extensions
  - policy
  - rbac.authorization.k8s.io
 attributeRestrictions: null
  resources:
  - nodes/metrics
  - nodes/proxy
                  <== Add this line
  - nodes/stats
  - pods
                    <== Add this line
 verbs:
  - get
  - list
                    <== Add this line
```

June 2020

Simplified Data Collector Error Reporting

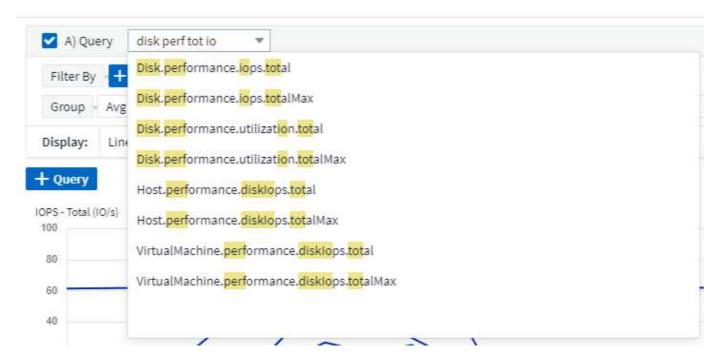
Reporting a data collector error is easier with the *Send Error Report* button on the data collector page. Clicking the button sends basic information about the error to NetApp and prompts investigation into the problem. Once pressed, Cloud Insights acknowledges that NetApp has been notified, and the Error Report button is disabled to indicate that an error report for that data collector has been sent. The button remains disabled until the browser page is refreshed.



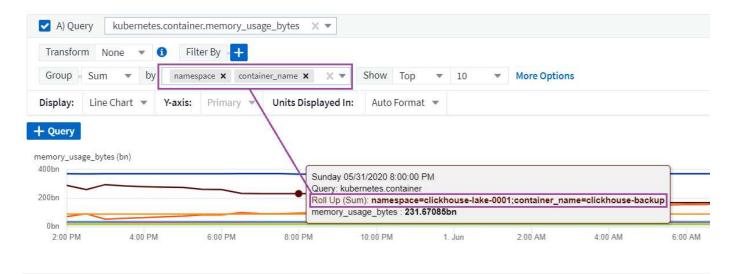
Widget Improvements

The following improvements have been made in dashboard widgets. These improvements are considered Preview functionality and may not be available for all Cloud Insights environments.

 New object/metric chooser: Objects (Storage, Disk, Ports, Nodes, etc.) and their associated metrics (IOPS, Latency, CPU Count, etc.) are now available in widgets in a single inclusive drop-down with powerful search capability. You can enter multiple partial terms in the drop-down, and Cloud Insights will list all object metrics meeting those terms.



 Multiple tags grouping: When working with integration data (Kubernetes, etc.), you may group the data by multiple tags/attributes. For example, Sum memory usage by Kubernetes Namespace and Container name.



May 2020

Reporting User Roles

The following roles have been added for Reporting:

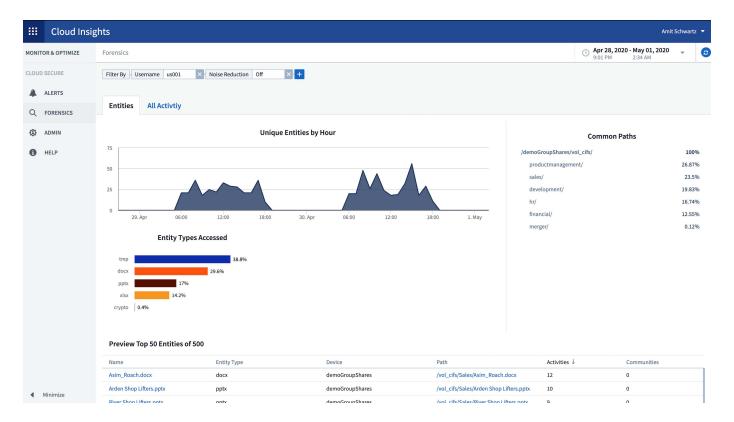
- · Cloud Insights Consumers: can run and view reports
- Cloud Insights Authors: can perform the Consumer functions as well as create and manage reports and dashboards
- Cloud Insights Administrators: can perform the Author functions as well as all administrative tasks

Cloud Secure Updates

Cloud Insights includes the following recent Cloud Secure changes.

In the Forensics > Activity Forensics page, we provide two views to analyze and investigate user activity:

- · Activity view, focused on user activity (What operation? Where performed?)
- · Entities view, focused on what files the user accessed.

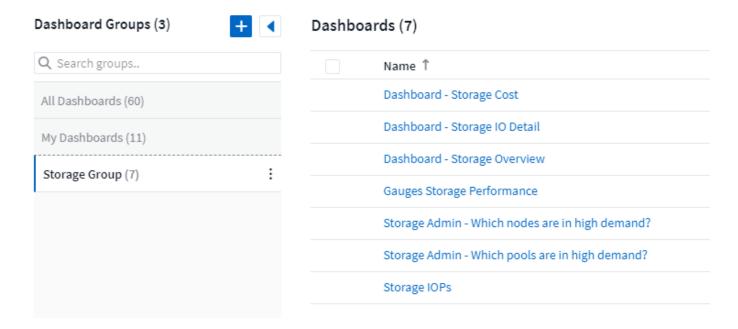


Additionally, the Alert email notification now contains a direct link to the alert page.

Dashboard Grouping

Dashboard grouping allows better management of dashboards that are relevant to you. You can add related dashboards to a group for "one-stop" management of, for example, your storage or virtual machines.

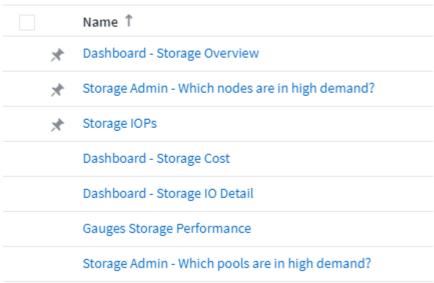
Groups are customized per user, so one person's groups can be different from someone else's. You can have as many groups as you need, with as few or as many dashboards in each group as you like.



Dashboard Pinning

You can pin dashboards so favorites always appear at the top of the list.

Dashboards (7)



TV Mode and Auto-Refresh

TV Mode and Auto-Refresh allow for near-real-time display of data on a dashboard or asset page:

- **TV Mode** provides an uncluttered display; the navigation menu is hidden, providing more screen real estate for your data display.
- Data in widgets on Dashboards and Asset Landing Pages Auto-Refresh according a refresh interval (as
 little as every 10 seconds) determined by the Dashboard Time Range selected (or widget time range, if set
 to override the dashboard time).

Combined, TV Mode and Auto-Refresh provide a live view of your Cloud Insights data, perfect for seamless demonstrations or in-house monitoring.

April 2020

New Dashboard Time Range Choices

Time range choices for dashboards and other Cloud insights pages now include *Last 1 Hour* and *Last 15 Minutes*.

Cloud Secure Updates

Cloud Insights includes the following recent Cloud Secure changes.

- Better file and folder metadata change recognition to detect if the user changed Permission, Owner, or Group Ownership.
- Export user activity report to CSV.

Cloud Secure monitors and audits all user access operations on files and folders. Activity auditing allows you to comply with internal security policies, meet external compliance requirements such as PCI, GDPR, and HIPAA, and conduct data breach and security incident investigations.

Default Dashboard Time

The default time range for dashboards is now 3 Hours instead of 24 hours.

Optimized Aggregation Times

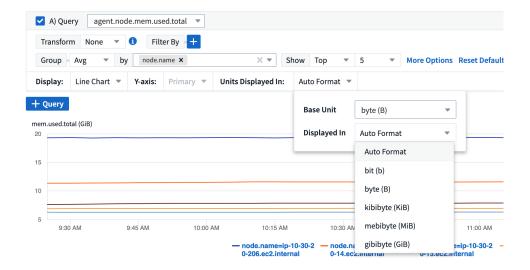
Optimized time aggregation intervals in time-series widgets (Line, Spline, Area, and Stacked Area charts) are more frequent for 3-hour and 24-hour dashboard/widget time ranges, allowing for faster charting of data.

- 3 hour time range optimizes to a 1 minute aggregation interval. Previously this was 5 minutes.
- 24 hour time range optimizes to a 30 minute aggregation interval. Previously this was 1 hour.

You can still override the optimized aggregation by setting a custom interval.

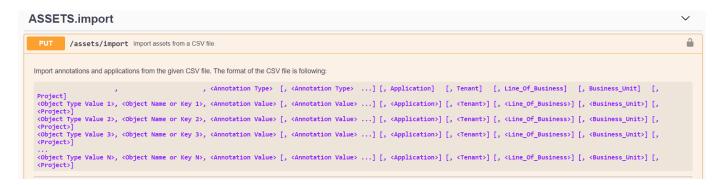
Display Unit Auto-Format

In most widgets, Cloud Insights knows the base unit in which to display values, for example *Megabytes*, *Thousands*, *Percentage*, *Milliseconds* (*ms*), etc., and now automatically formats the widget to the most readable unit. For example a data value of 1,234,567,890 bytes would be auto formatted to 1.23 gibibytes. In many cases, Cloud Insights knows the best format for the data being acquired. In cases where the best format is not known, or in widgets where you want to override the automatic formatting, you can choose the format you want.



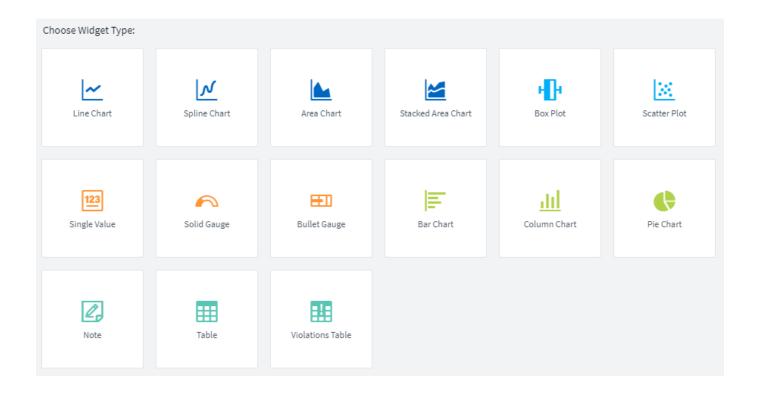
Import Annotations Using API

With Cloud Insights Premium Edition's powerful API, you can now import annotations and assign them to objects using a .CSV file. You can also import applications and assign business entities in the same way.



Simpler Widget Selector

Adding widgets to dashboards and asset landing pages is easier with a new widget selector that shows all widget types in a single all-at-once view, so the user no longer needs to scroll through a list of widget types to find the one they want to add. Related widgets are color-coordinated and grouped by proximity in the new selector.



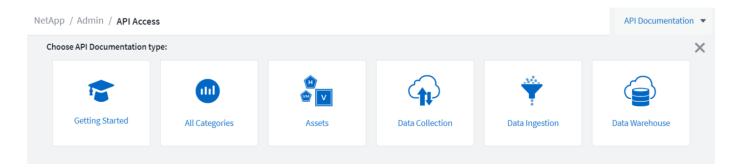
February 2020

API with Premium Edition

Cloud Insights Premium Edition comes with a powerful API that can be used to integrate Cloud Insights with other applications, such as CMDB's or other ticketing systems.

Detailed, Swagger-based information is found in **Admin > API Acccess**, under the **API Documentation** link. Swagger provides a brief description and usage information for the API, and allows you to try each API out in your environment.

The Cloud Insights API uses Access Tokens to provide permission-based access to categories of API, such as ASSETS or COLLECTION.



Initial Polling After Adding A Data Collector

Previously, after configuring a new data collector, Cloud Insights would poll the data collector immediately to gather *inventory* data, but would wait until the configured performance poll interval (typically 15 minutes) to gather initial *performance* data. It would then wait for another interval before initiating the second performance

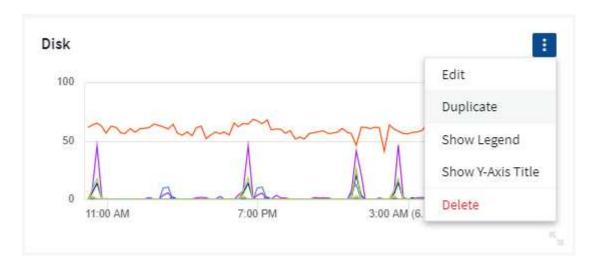
poll, which meant it would take up to 30 minutes before meaningful data was acquired from a new data collector.

Data collector polling has been greatly improved, such that the initial performance poll occurs immediately after the inventory poll, with the second performance poll occurring within a few seconds after completion of the first performance poll. This allows Cloud Insights to begin showing useful data on dashboards and graphs within a very short time.

This poll behavior also occurs after editing the configuration of an existing data collector.

Easier Widget Duplication

It is easier than ever to create a copy of a widget on a dashboard or landing page. In dashboard Edit mode, click the menu on the widget and select **Duplicate**. The widget editor is launched, pre-filled with the original widget's configuration and with a "copy" suffix in the widget name. You can easily make any necessary changes and Save the new widget. The widget will be placed at the bottom of your dashboard, and you can position it as needed. Remember to Save your dashboard when all changes are complete.



Single Sign-On (SSO)

With Cloud Insights Premium Edition, administrators can enable **Single Sign-On** (SSO) access to Cloud Insights for all users in their corporate domain, without having to invite them individually. With SSO enabled, any user with the same domain email address can log into Cloud Insights using their corporate credentials.



SSO is only available in Cloud Insights Premium Edition, and must be configured before it can be enabled for Cloud Insights. SSO configuration includes Identity Federation through NetApp Cloud Central. Federation allows single sign-on users to access your NetApp Cloud Central accounts using credentials from your corporate directory.

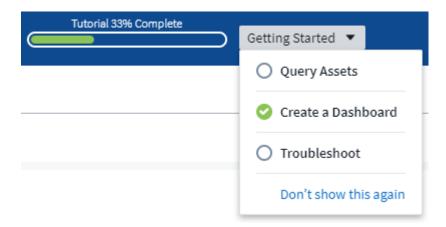
January 2020

Swagger documentation for REST API

Swagger explains each available REST API in Cloud Insights, as well as its usage and syntax. Information on Cloud Insights APIs is available in documentation.

Feature Tutorials Progress Bar

The feature tutorials checklist has been moved to the top banner and now features a progress indicator. Tutorials are available for each user until dismissed, and are always available in Cloud Insights documentation.



Acquisition Unit Changes

When installing an Acquisition Unit (AU) on a host or VM that has the same name as an already-installed AU, Cloud Insights assures a unique name by appending the AU name with "_1", "_2", etc. This is also the case when uninstalling and reinstalling an AU from the same VM without first removing it from Cloud Insights. Want a different AU name altogether? No problem; AU's can be renamed after installation.

Optimized Time Aggregation in Widgets

In widgets, you can choose between an *Optimized* time aggregation interval or a *Custom* interval that you set. Optimized aggregation automatically selects the right time interval based on the selected dashboard time range (or widget time range, if overriding the dashboard time). The interval dynamically changes as the dashboard or widget time range is changed.

Simplified "Getting Started with Cloud Insights" process

The process for getting started using Cloud Insights has been simplified to make your first-time setup smoother and easier. Simply select an initial data collector and follow the instructions. Cloud Insights will walk you through configuring the data collector and any agent or acquisition unit required. In most cases it will even import one or more initial dashboards so you can start gaining insight into your environment quickly (but please allow up to 30 minutes for Cloud Insights to collect meaningful data).

Additional improvements:

- Acquisition Unit installation is simpler and runs faster.
- Alphabetical Data Collectors choices make it easier to find the one you're looking for.
- · Improved Data Collector setup instructions are easier to follow.
- Experienced users can skip the getting started process with the click of a button.
- A new Progress bar shows you where you are in the process.

December 2019

Business Entity can be used in filters

Business Entity annotations can be used in filters for queries, widgets, performance policies, and landing pages.

Drill-down available for Single-Value and Gauge widgets, and any widgets rolled to by "All"

Clicking the value in a single-value or gauge widget opens a query page showing the results of the first query used in the widget. Additionally, clicking the legend for any widget whose data is rolled up by "All" will also open a query page showing the results of the first query used in the widget.

Trial period extended

New users who sign up for a free trial of Cloud Insights now have 30 days to evaluate the product. This is an increase from the previous 14-day trial period.

Managed Unit calculation

The calculation of Managed Units (MUs) in Cloud Insights has been changed to the following:

- 1 Managed Unit = 2 hosts (any virtual or physical machine)
- 1 Managed Unit = 4 TB of unformatted capacity of physical or virtual disks

This change effectively doubles the environment capacity that you can monitor using your existing Cloud Insights subscription.

November 2019

Editions Feature Comparison Table

The **Admin > Subscription** page comparison table has been updated to list the feature sets available in Basic, Standard, and Premium Editions of Cloud Insights. NetApp is constantly improving its Cloud Services, so check this page often to find the Edition that's right for your evolving business needs.

October 2019

Reporting

Cloud Insights Reporting is a business intelligence tool that enables you to view pre-defined reports or create custom reports. With Reporting you can perform the following tasks:

- · Run a pre-defined report
- · Create a custom report
- · Customize the report format and delivery method
- · Schedule reports to run automatically
- Email reports
- · Use colors to represent thresholds on data

Cloud Insights Reporting can generate custom reports for areas like chargeback, consumption analysis, and forecasting, and can help answer questions such as the following:

- · What inventory do I have?
- · Where is my inventory?
- · Who is using our assets?
- · What is the chargeback for allocated storage for a business unit?
- How long until I need to acquire additional storage capacity?
- · Are business units aligned along the proper storage tiers?
- How is storage allocation changing over a month, quarter, or year?

Reporting is available with Cloud Insights **Premium Edition**.

Active IQ Enhancements

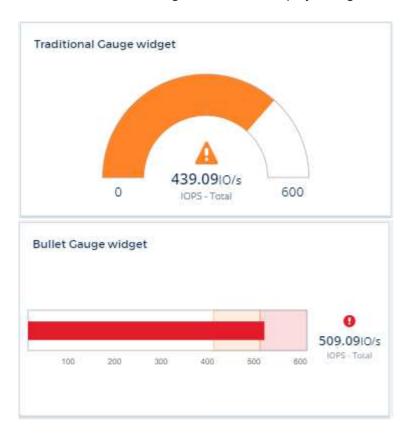
Active IQ Risks are now available as objects that can be queried as well as used in dashboard table widgets. The following Risks object attributes are included:

- * Category
- * Mitigation Category
- * Potential Impact
- * Risk Detail
- * Severity
- * Source
- * Storage
- * Storage Node
- * UI Category

September 2019

New Gauge Widgets

Two new widgets are available for displaying single-value data on your dashboards in eye-catching colors based on thresholds you specify. You can display values using either a **Solid Gauge** or **Bullet Gauge**. Values that land inside the Warning range are displayed in orange. Values in the Critical range are displayed in red. Values below the Warning threshold are displayed in green.



Conditional Color Formatting for Single Value Widget

You can now display the Single-Value widget with a colored background based on thresholds you set.



Invite Users During Onboarding

At any point during the onboarding process, you can click on Admin > User Management > +User to invite additional users to your Cloud Insights environment. Be aware that users with *Guest* or *User* roles will see greater benefit once onboarding is complete and data has been collected.

Data Collector Detail Page improvement

The data collector detail page has been improved to display errors in a more readable format. Errors are now displayed in a separate table on the page, with each error displayed on a separate line in the case of multiple

August 2019

All vs. Available Data Collectors

When adding data collectors to your environment, you can set a filter to show only the data collectors available to you based on your subscription level, or all data collectors.

ActiveIQ Integration

Cloud Insights collects data from NetApp ActiveIQ, which provides a series of visualizations, analytics, and other support related services to NetApp customers and their hardware / software systems. Cloud Insights integrates with ONTAP Data Management systems. See Active IQ for more information.

July 2019

Dashboard Improvements

Dashboards and Widgets have been improved with the following changes:

- In addition to Sum, Min, Max, and Avg, **Count** is now an option for roll up in Single-Value widgets. When rolling up by "Count", Cloud Insights checks if an object is active or not, and only adds the active ones to the count. The resulting number is subject to aggregation and filters.
- In the Single-Value widget, you now have a choice to display the resulting number with 0, 1, 2, 3, or 4 decimal places.
- Line charts show an axis label and units when a single counter is being plotted.
- **Transform** option is available for Services integration data now in all time-series widgets for all metrics. For any services integration (Telegraf) counter or metric in time-series widgets (Line, Spline, Area, Stacked Area), you are given a choice of how you want to **Transform the values**. None (display value as-is), Sum, Delta, Cumulative, etc.

Downgrading to Basic Edition

Downgrade to Basic Edition fails with an error message if there is no available NetApp device configured that has successfully completed a poll in the last 7 days.

Collecting Kube-State-Metrics

The Kubernetes Data Collector now collects objects and counters from the kube-state-metrics plugin, greatly expanding the number and scope of metrics available for monitoring in Cloud Insights.

June 2019

Cloud Insights Editions

Cloud Insights is available in different Editions to fit your budget and business needs. Existing NetApp customers with an active NetApp support account can enjoy 7 days of data retention and access to NetApp data collectors with the free **Basic Edition**, or get increased data retention, access to all supported data collectors, expert technical support and more with **Standard Edition**. For more information on available features, see NetApp's Cloud Insights site.

New Infrastructure Data Collector: NetApp HCI

 NetApp HCI Virtual Center has been added as an Infrastructure data collector. The HCI Virtual Center data collector collects NetApp HCI Host information and requires read-only privileges on all objects within the Virtual Center.

Note that the HCl data collector acquires from the HCl Virtual Center only. To collect data from the storage system, you must also configure the NetApp SolidFire data collector.

May 2019

New Service Data Collector: Kapacitor

Kapacitor has been added as a data collector for services.

Integration with Services via Telegraf

In addition to acquisition of data from infrastructure devices such as switches and storage, Cloud Insights now collects data from a variety of Operating Systems and Services, using Telegraf as its agent for collection of integration data. Telegraf is a plugin-driven agent that can be used to collect and report metrics. Input plugins are used to collect the desired information into the agent by accessing the system/OS directly, by calling third-party APIs, or by listening to configured streams.

Documentation for currently supported integrations can be found in the menu to the left under **Reference and Support**.

Storage Virtual Machine Assets

 Storage Virtual Machines (SVMs) are available as assets in Cloud Insights. SVMs have their own Asset Landing Pages, and can be displayed and used in searches, queries, and filters. SVMs can also be used in dashboard widgets as well as associated with annotations.

Reduced Acquisition Unit System Requirements

• The system CPU and memory requirements for the Acquisition Unit (AU) software have been reduced. The new requirements are:

Component	Old Requirement	New Requirement
CPU Cores	4	2

Memory	16 GB	8 GB	

Additional Platforms Supported

• The following platforms have been added to those currently supported for Cloud Insights:

Linux	Windows
CentOS 7.3 64-bit CentOS 7.4 64-bit CentOS 7.6 64-bit Debian 9 64-bit Red Hat Enterprise Linux 7.3 64-bit Red Hat Enterprise Linux 7.4 64-bit Red Hat Enterprise Linux 7.6 64-bit Ubuntu Server 18.04 LTS	Microsoft Windows 10 64-bit Microsoft Windows Server 2008 R2 Microsoft Windows Server 2019

April 2019

Filter Virtual Machines by Tags

When configuring the following data collectors, you can filter to include or exclude virtual machines from data collection according to their Tags or Labels.

- Amazon EC2
- Azure
- Google Cloud Platform

March 2019

Email Notifications for Subscription-related Events

- You can select recipients for email notifications when subscription-related events occur, such as upcoming trial expiration or subscribed account changes. You can choose recipients for these notifications from among following:
 - · All Account Owners
 - All Administrators
 - Additional Email Addresses that you specify

Additional Dashboards

- The following new AWS-focused dashboards have been added to the gallery and are available for import:
 - AWS Admin Which EC2 are in high demand?
 - AWS EC2 Instance Performance by Region

February 2019

Collecting from AWS Child Accounts

Cloud Insights supports collection from AWS child accounts within a single data collector. Your AWS
environment must be configured to allow Cloud Insights to collect from child accounts.

Data Collector Naming

• Data Collector names can now include periods (.), hyphens (-), and spaces () in addition to letters, numbers, and underscores. Names may not begin or end with a space, period, or hyphen.

Acquisition Unit for Windows

 You can configure a Cloud Insights Acquisition Unit on a Windows server/VM. Review the Windows prerequisites before installing the Acquisition Unit software.

January 2019

"Owner" field is more readable

 In Dashboard and Query lists, the data for the "Owner" field was previously an authorization ID string, instead of a user-friendly owner name. The "Owner" field now shows a simpler and more readable owner name.

Managed Unit Breakdown on Subscription Page

 For each data collector listed on the Admin > Subscription page, you can now see a breakdown of Managed Unit (MU) counts for hosts and storage, as well as the total.

December 2018

Improvement of UI Load Time

• The initial loading time for the Cloud Insights user interface (UI) has been significantly improved. Refresh time for the UI also benefits from this improvement in circumstances where metadata is loaded.

Bulk Edit Data Collectors

You can edit information for multiple data collectors at the same time. On the Observability > Collectors
page, select the data collectors to modify by checking the box to the left of each and click the Bulk Actions
button. Choose Edit and modify the necessary fields.

The data collectors selected must be the same vendor and model, and reside on the same Acquisition Unit.

Support and Subscription pages are Available During Onboarding

During the onboarding workflow, you can navigate to the Help > Support and Admin > Subscription
pages. Returning from those pages returns you to the onboarding workflow, providing you have not closed
the browser tab.

November 2018

Subscribe through NetApp Sales or AWS Marketplace

Cloud Insights subscription and billing is now available directly through NetApp. This is in addition to the
self-serve subscription available through AWS Marketplace. A new Contact Sales link is presented on the
Admin > Subscription page. For customers whose environments have or are expected to have 1,000 or
more Managed Units (MUs), it is recommended to contact NetApp sales via the Contact Sales link.

Text Annotation Hyperlinks

Text-type annotations can now include hyperlinks.

Onboarding Walkthrough

 Cloud Insights now features an onboarding walkthrough for the first user (administrator or account owner) to log in to a new environment. The walkthrough takes you through installing an Acquisition Unit, configuring an initial data collector, and selecting one or more useful dashboards.

Import Dashboards from the Gallery

In addition to selecting dashboards during onboarding, you can import dashboards via Dashboards >
Show All Dashboards and clicking +From Gallery.

Duplicating Dashboards

• The ability to duplicate a dashboard has been added to the dashboard list page as a choice in the options menu for each dashboard, and on a dashboard's main page itself from the *Save* menu.

Cloud Central products menu

• The menu allowing you to switch to other NetApp Cloud Central products has moved to the upper right corner of the screen.

Cloud Insights Onboarding

Before you can start working with Cloud Insights, you must sign up on the **NetApp BlueXP** portal. If you already have a BlueXP login, you can start a free trial of Cloud Insights with a few quick steps.

Creating your NetApp BlueXP account

To get started with NetApp's cloud services, go to NetApp BlueXP and click Get Started.

- If you have not already signed up, select Sign Up
- Enter a valid business email address and choose a password.
- Enter your company name, and your full name.
- Accept the terms and conditions and select Continue.
- · BlueXP will walk you through getting started.

What if I already have a NetApp BlueXP login?

Once you have a NetApp BlueXP account, simply choose Log In on the NetApp BlueXPI portal page.

Enter your email address and password. You will then be taken to NetApp's cloud offerings page.

Select Cloud Insights.



Starting your Cloud Insights free trial

If this is your first time logging in to Cloud Insights, under the Cloud Insights offering, click on **Start Free Trial**. Cloud Insights will walk you through creating your company's environment.

Once the creation of your environment is complete, you can use your BlueXP credentials to log in and start your free, 30-day trial of Cloud Insights. During this trial you can explore the features that Cloud Insights has to offer.

During the free trial, you can start your subscription to Cloud Insights at any time. When you are subscribed, You can use the Cloud Insights features based on your current subscription.

Sign in and go

Once your environment has been created, at any time you can simply log in to the NetApp BlueXP Portal and click **Go to Cloud Insights**. You will be taken directly to your Cloud Insights environment.

You can also open a browser directly to your Cloud Insights environment URL, for example:

https://<environment-prefix>.c01.cloudinsights.netapp.com/

The URL will also be included in each user's invitation email for simple access and bookmarking. If the user is not already logged in to BlueXP, they will be prompted to log in.



New users must still sign up for access to BlueXP before they can access their environment URL.

The first time you log in to a new environment, you will be guided through setting up to begin gathering data.

Logging Out

To log out of Cloud Insights, click your **User Name** and select **Log Out**. You will be taken back to the BlueXP sign-in screen.



Logging out of Cloud Insights logs you out of BlueXP. You will also be logged out of other NetApp Cloud services that use the BlueXP sign in.

Inactivity Timeout

By default, BlueXP will log a user out if there is no activity for six hours (360 minutes). Regardless of activity, users will be logged out after seven days.

Security

Cloud Insights Security

Product and customer data security is of utmost importance at NetApp. Cloud Insights follows security best practices throughout the release life cycle to make sure customer information and data is secured in the best possible way.

Security Overview

Physical security

The Cloud Insights production infrastructure is hosted in Amazon Web Services (AWS). Physical and environmental security-related controls for Cloud Insights production servers, which include buildings as well as locks or keys used on doors, are managed by AWS. As per AWS: "Physical access is controlled both at the perimeter and at building ingress points by professional security staff utilizing video surveillance, intrusion detection systems, and other electronic means. Authorized staff utilize multi-factor authentication mechanisms to access data center floors."

Cloud Insights follows the best practices of the Shared Responsibility model described by AWS.

Product security

Cloud Insights follows a development lifecycle in line with Agile principles, thus allowing us to address any security-oriented software defects more rapidly, compared to longer release cycle development methodologies. Using continuous integration methodologies, we are able to rapidly respond to both functional and security changes. The change management procedures and policies define when and how changes occur and help to maintain the stability of the production environment. Any impactful changes are formally communicated, coordinated, properly reviewed, and approved prior to their release into the production environment.

Network security

Network access to resources in the Cloud Insights environment is controlled by host-based firewalls. Each resource (such as a load balancer or virtual machine instance) has a host-based firewall that restricts inbound traffic to only the ports needed for that resource to perform its function.

Cloud Insights uses various mechanisms including intrusion detection services to monitor the production environment for security anomalies.

Risk Assessment

Cloud Insights team follows a formalized Risk Assessment process to provide a systematic, repeatable way to identify and assess the risks so that they can be appropriately managed through a Risk Treatment Plan.

Data protection

The Cloud Insights production environment is set up in a highly redundant infrastructure utilizing multiple availability zones for all services and components. Along with utilizing a highly available and redundant compute infrastructure, critical data is backed up at regular intervals and restores are periodically tested. Formal backup policies and procedures minimize the impact of interruptions of business activities and protects business processes against the effects of failures of information systems or disasters and ensures their timely and adequate resumption.

Authentication and access management

All customer access to Cloud Insights is done via browser UI interactions over https. Authentication is accomplished via the 3rd party service, Auth0. NetApp has centralized on this as the authentication layer for all Cloud Data services.

Cloud Insights follows industry best practices including "Least Privilege" and "Role-based access control" around logical access to the Cloud Insights production environment. Access is controlled on a strict need basis and is only granted for select authorized personnel using multi-factor authentication mechanisms.

Collection and protection of customer data

All customer data is encrypted in transit across public networks and encrypted at rest. Cloud Insights utilizes encryption at various points in the system to protect customer data using technologies that includes Transport Layer Security (TLS) and the industry-standard AES-256 algorithm.

Customer deprovisioning

Email notifications are sent out at various intervals to inform the customer their subscription is expiring. Once the subscription has expired, the UI is restricted and a grace period begins for data collection. The customer is then notified via email. Trial subscriptions have a 14-day grace period and paid subscription accounts have a 28-day grace period. After the grace period has expired, the customer is notified via email that the account will be deleted in 2 days. A paid customer can also request directly to be off the service.

Expired tenants and all associated customer data are deleted by the Cloud Insights Operations (SRE) team at the end of the grace period or upon confirmation of a customer's request to terminate their account. In either case, the SRE team runs an API call to delete the account. The API call deletes the tenant instance and all customer data. Customer deletion is verified by calling the same API and verifying that the customer tenant status is "DELETED."

Security incident management

Cloud Insights is integrated with NetApp's Product Security Incident Response Team (PSIRT) process to find, assess, and resolve known vulnerabilities. PSIRT intakes vulnerability information from multiple channels including customer reports, internal engineering, and widely recognized sources such as the CVE database.

If an issue is detected by the Cloud Insights engineering team, the team will initiate the PSIRT process, assess, and potentially remediate the issue.

It is also possible that a Cloud Insights customer or researcher may identify a security issue with the Cloud Insights product and report the issue to Technical Support or directly to NetApp's incident response team. In these cases, the Cloud Insights team will initiate the PSIRT process, assess, and potentially remediate the issue.

Vulnerability and Penetration testing

Cloud Insights follows industry best practices and performs regular vulnerability and penetration testing using internal and external security professionals and companies.

Security awareness training

All Cloud Insights personnel undergo security training, developed for individual roles, to make sure each employee is equipped to handle the specific security-oriented challenges of their roles.

Compliance

Cloud Insights performs independent third-party Audit and validations from external Licensed CPA firm of its security, processes, and services, including completion of the SOC 2 Audit.

NetApp Security Advisories

You can view NetApp's available security advisories here.

Information and Region

NetApp takes the security of customer information very seriously. Here is how and where Cloud Insights stores your information.

What information does Cloud Insights store?

Cloud Insights stores the following information:

· Performance data

Performance data is time-series data providing information about the performance of the monitored device/source. This includes, for example, the number of IOs delivered by a storage system, the throughput of a FibreChannel port, the number of pages delivered by a web server, the response time of a database, and more.

· Inventory data

Inventory data consists of metadata describing the monitored device/source and how it is configured. This includes, for example, hardware and software versions installed, disks and LUNs in a storage system, CPU cores, RAM and disks of a virtual machine, the tablespaces of a database, the number and type of ports on a SAN switch, directory/file names (if Storage Workload Security is enabled), etc.

· Configuration data

This summarizes customer-provided configuration data used to manage customer inventory and operations, e.g. hostnames or IP addresses of the monitored devices, polling intervals, timeout values, etc.

Secrets

Secrets consist of the credentials used by the Cloud Insights Acquisition Unit to access customer devices and services. These credentials are encrypted using strong asymmetric encryption, and the private keys are stored only on the Acquisition Units and never leave the customer environment. Even privileged Cloud Insights SREs are unable to access customer secrets in plain-text due to this design.

Functional Data

This is data generated as a result of NetApp providing the Cloud Data Service, which informs NetApp in the development, deployment, operations, maintenance, and securing of the Cloud Data Service. Functional Data does not contain Customer Information or Personal Information.

· User Access data

Authentication and access information that allows NetApp Cloud Central to communicate with regional Cloud Insights sites, including data related to user Authorization.

Storage Workload Security User Directory Data

In cases where the Workload Security functionality is enabled AND the customer chooses to enable the User Directory collector, the system will store user display names, corporate email addresses, and other information collected from Active Directory.



User Directory data refers to user directory information collected by the Workload Security User Directory data collector, not to data about the users of Cloud Insights/Workload Security themselves.

No explicit personal data is collected from infrastructure and services resources. Collected information consists of performance metrics, configuration information and infrastructure metadata only, much like many vendor phone-homes, including NetApp auto-support and ActivelQ. However, depending on a customer's naming conventions, data for shares, volumes, VMs, qtrees, applications, etc. may contain personally identifiable information.

If Workload Security is enabled, the system additionally looks at file and directory names on SMB or other shares, which may contain personally identifiable information. Where customers enable the Workload Security User Directory Collector (which essentially maps Windows SIDs to usernames through Active Directory), the display name, corporate email address and any additional attributes selected will be collected and stored by Cloud Insights.

Additionally, access logs to Cloud Insights are maintained and contain users' IP and email addresses used to log into the service.

Where is my information stored?

Cloud Insights stores information according to the region in which your environment is created.

The following information is stored in the host region:

- · Telemetry and asset/object information, including counters and performance metrics
- · Acquisition Unit information
- · Functional data
- · Audit information on user activities inside Cloud Insights
- Workload Security Active Directory information
- · Workload Security Audit information

The following information resides in the United States, regardless of the region hosting your Cloud Insights environment:

- Environment site (sometimes called "tenant") information such as site/account owner.
- Information that allows NetApp Cloud Central to communicate with regional Cloud Insights sites, including anything to do with user Authorization.
- Information related to the relation between the Cloud Insights user and the tenant.

Host Regions

Host regions include:

US: us-east-1

• EMEA: eu-central-1

APAC: ap-southeast-2

More Information

You can read more about NetApp's privacy and security at the following links:

- Trust Center
- Cross-Border Data Transfers
- Binding Corporate Rules
- Responding to Third-Party Data Requests
- NetApp Privacy Principles

SecurityAdmin Tool

Cloud Insights Includes security features that allow your environment to operate with enhanced security. The features include improvements to encryption, password hashing, and the ability to change internal user passwords as well as key pairs that encrypt and decrypt passwords.

To protect sensitive data, NetApp recommends you change the default keys and the *Acquisition* user password after an installation or upgrade.

Data source encrypted passwords are stored in Cloud Insights, which uses a a public key to encrypt passwords when a user enters them in a data collector configuration page. Cloud Insights does not have the private keys required to decrypt the data collector passwords; only Acquisition Units (AUs) have the data collector private key required to decrypt data collector passwords.

Upgrade and installation considerations

When your Insight system contains non-default security configurations (i.e. you have rekeyed passwords), you must back up your security configurations. Installing new software, or in some cases upgrading software, reverts your system to a default security configuration. When your system reverts to the default configuration, you must restore the non-default configuration in order for the system to operate correctly.

Managing security on the acquisition unit

The SecurityAdmin tool allows you to manage security options for Cloud Insights, and is run on the acquisition unit system. Security management includes managing keys and passwords, saving and restoring security configurations you create or restoring configurations to the default settings.

Before you begin

- You must have admin privileges on the AU system in order to install the Acquisition Unit software (which includes the SecurityAdmin tool).
- If you have non-admin users who will subsequently need to access the SecurityAdmin tool, they must be added to the *cisys* group. The *cisys* group is created during AU installation.

After AU install, the SecurityAdmin tool is found on the acquisition unit system at either of these locations:

```
Windows - C:\Program Files\SANscreen\securityadmin\bin\securityadmin.bat
Linux - /bin/oci-securityadmin.sh
```

Using the SecurityAdmin Tool

Start the SecurityAdmin tool in interactive mode (-i).



It is recommended to use the SecurityAdmin tool in interactive mode, to avoid passing secrets on the command line, which can be captured in logs.

The following options are displayed:

```
[root@ci-qa-xitij-cis2-28594linau bin]# ./securityadmin -i
Select Action:

1 - Backup

2 - Restore

3 - Register / Update External Key Retrieval Script

4 - Rotate Encryption Keys

5 - Reset to Default Keys

6 - Change Truststore Password

7 - Change Keystore Password

8 - Encrypt Collector Password

9 - Exit
Enter your choice:
```

1. Backup

Creates a backup zip file of the vault containing all passwords and keys and places the file in a location specified by the user, or in the following default locations:

```
Windows - C:\Program Files\SANscreen\backup\vault
Linux - /var/log/netapp/oci/backup/vault
```

It is recommended that vault backups be kept secure, as they include sensitive information.

2. Restore

Restores the zip backup of the vault that was created. Once restored, all passwords and keys are reverted to values existing at the time of the backup creation.

Restore can be used to synchronize passwords and keys on multiple servers, for example using these steps: 1) Change encryption keys on the AU. 2) Create a backup of the vault. 3) Restore the vault backup to each of the AUs.

3. Register / Update External Key Retrieval Script

Use an external script to register or change the AU encryption keys used to encrypt or decrypt device passwords.

When you change encryption keys, you should back up your new security configuration so that you can restore it after an upgrade or installation.

Note this option is only available on Linux.

When using your own key retrieval script with the SecurityAdmin tool, keep the following in mind:

- The current supported Algorithm is RSA with minimum 2048 bits.
- The script must return the private and public keys in plain text. The script must not return encrypted private and public keys.
- The script should return raw, encoded contents (PEM format only).
- The external script must have execute permissions.

4. Rotate Encryption Keys

Rotate your encryption keys (un-registers current keys and registers new keys). To use a key from an external key management system, you must specify the public key id and private key id.

1. Reset to Default Keys

Resets acquisition user password and acquisition user encryption keys to default values, Default values are those provided during installation.

2. Change Truststore Password

Change the password of the truststore.

3. Change Keystore Password

Change the password of the keystore.

4. Encrypt Collector Password

Encrypt data collector password.

5. Exit

Exit the SecurityAdmin tool.

Chose the option you want to configure and follow the prompts.

Specifying a user to run the tool

If you are in a controlled, security-conscious environment, you may not have the *cisys* group but may still want specific users to run the SecurityAdmin tool.

You can achieve this by manually installing the AU software and specifying the user/group for whom you want access.

- Using the API, download the CI Installer to the AU system and unzip it.
 - You will need a one-time authorization token. See the API Swagger documentation (Admin > API
 Access and select the API Documentation link) and find the GET /au/oneTimeToken API section.
 - Once you have the token, use the *GET /au/installers/{platform}/{version}* API to download the installer file. You will need to provide platform (Linux or Windows) as well as installer version.
- · Copy the downloaded installer file to the AU system and unzip it.
- Navigate to the folder containing the files, and run the installer as root, specifying the user and group:

```
./cloudinsights-install.sh <User> <Group>
```

If the specified user and/or group do not exist, they will be created. The user will have access to the SecurityAdmin tool.

Updating or Removing Proxy

The SecurityAdmin tool can be used to set or remove proxy information for the Acquisition Unit by running the tool with the *-pr* parameter:

```
[root@ci-eng-linau bin]# ./securityadmin -pr
usage: securityadmin -pr -ap <arg> | -h | -rp | -upr <arg>
The purpose of this tool is to enable reconfiguration of security aspects
of the Acquisition Unit such as encryption keys, and proxy configuration,
etc. For more information about this tool, please check the Cloud Insights
Documentation.
-ap, --add-proxy <arg>
                            add a proxy server. Arguments: ip=ip
                             port=port user=user password=password
                             domain=domain
                             (Note: Always use double quote(") or single
                             quote(') around user and password to escape
                             any special characters, e.g., <, >, ~, `, ^,
                             For example: user="test" password="t'!<@1"
                             Note: domain is required if the proxy auth
                             scheme is NTLM.)
-h, --help
                            remove proxy server
-rp, --remove-proxy
-upr, --update-proxy <arg>
                            update a proxy. Arguments: ip=ip port=port
                             user=user password=password domain=domain
                             (Note: Always use double quote(") or single
                             quote(') around user and password to escape
                             any special characters, e.g., <, >, ~, `, ^,
                             For example: user="test" password="t'!<@1"
                             Note: domain is required if the proxy auth
                             scheme is NTLM.)
```

For example, to remove the proxy, run this command:

```
[root@ci-eng-linau bin]# ./securityadmin -pr -rp
```

You must restart the Acquisition Unit after running the command.

To update a proxy, the command is

```
./securityadmin -pr -upr <arg>
```

External Key Retrieval

If you provide a UNIX shell script, it can be executed by the acquisition unit to retrieve the **private key** and the **public key** from your key management system.

To retrieve the key, Cloud Insights will execute the script, passing in two parameters: key id and key type. Key id can be used to identify the key in your key management system. Key type is either "public" or "private". When the key type is "public", the script must return the public key. When the key type is "private", the private key must be returned.

To send the key back to the acquisition unit, the script must print the key to standard output. The script must print *only* the key to standard output; no other text must be printed to standard output. Once the requested key is printed to the standard output, the script must exit with an exit code of 0; any other return code is considered an error.

The script must be registered with the acquisition unit using the SecurityAdmin tool, which will execute the script along with the acquisition unit. The script must have *read* and *execute* permission for the root and "cisys" user. If the shell script is modified after registering, the modified shell script must be re-registered with the acquisition unit.

input parameter: key id	Key identifier used to identify the key in the customers key management system.
input parameter: key type	public or private.
output	The requested key must be printed to the standard output. 2048 bit RSA key is currently supported. Keys must be encoded and printed in the following format - private key format - PEM, DER-encoded PKCS8 PrivateKeyInfo RFC 5958 public key format - PEM, DER-encoded X.509 SubjectPublicKeyInfo RFC 5280
exit code	Exit code of zero for success. All other exit values are considered failure.
script permissions	Script must have read and execute permission for the root and "cisys" user.
logs	Script executions are logged. Logs can be found in - /var/log/netapp/cloudinsights/securityadmin/securityad min.log /var/log/netapp/cloudinsights/acq/acq.log

Encrypting a Password for use in API

Option 8 allows you to encrypt a password, which you can then pass to a data collector via API.

Start the SecurityAdmin tool in interactive mode and select option 8: Encrypt Password.

```
securityadmin.sh -i
```

You are prompted to enter the password you want to encrypt. Note that the characters you type are not shown on screen. Re-enter the password when prompted.

Alternatively, if you will use the command in a script, on a command line use *securityadmin.sh* with the "-enc" parameter, passing in your unencrypted password:

```
securityadmin -enc mypassword
```

[root@ci-eng-srivardh-learn bin]# securityadmin.sh -enc mypassword

Please copy paste the encrypted password below:

ciYJAMpdEncBsfc2PiXVBTappugSscDq3XF7Pw7/r5f00JL0mbSel6QA/umLrr8PzBnjcJUHHRwrgf3jFio/2H3GftnqIxSs7ATKiQw5011uvxYiGftkzYaH2BKYHjkIiD

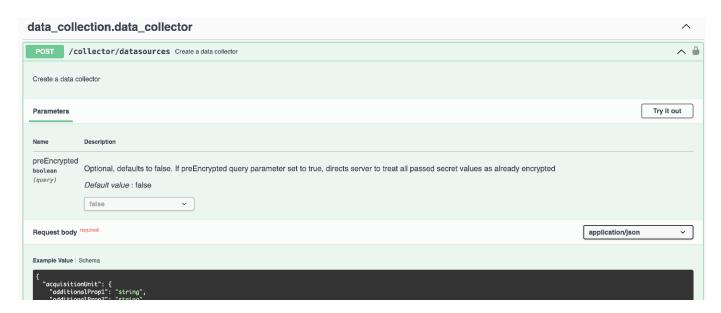
M8BEZZhm7pmTKWWpvAxhJbtjBrwUK2llM1GrnvaFlVVeydvsUMiggOenyJ/wxiko4gddif1Yq6rmia4yzvuYNw6Ppp5k/Pwy+0Hu0voRT+gca1ks80jQToAAO6WSHZfp71

mMokM30af43iV7eJZuMQ5RSq1cjBtYVnTWjp0Rn0g2kBCDf0PpWVrS6EKzh0HKQRRWpBZnQJNPv1bqtP+0pUh5Yd29RGX5Q==

The encrypted password is displayed on screen. Copy the entire string including any leading or trailing symbols.



To send the encrypted password to a data collector, you can use the Data Collection API. The swagger for this API can be found at **Admin > API Access** and click the "API Documentation" link. Select the "Data Collection" API type. Under the *data_collection.data_collector* heading, choose the */collector/datasources* POST API for this example.



If you set the *preEncrypted* option to *True*, any password you pass through the API command will be treated as **already encrypted**; the API will not re-encrypt the password(s). When building your API, simply paste the previously-encrypted password in the appropriate location.

https://<TENANT URL>/rest/v1/collector/datasources?preEncrypted=true

```
"name": "cdot-aaaaa",
 "config": {
  "dsTypeId": "93",
  "vendorModelId": "1",
  "packages": [
    "id": "foundation",
    "displayName": "Inventory",
    "isMandatory": true,
    "attributes": {
     "RELEASESTATUS": "OFFICIAL",
     "enabled": true,
     "ip": "10.62.219.30",
     "user": "admin",
     "password":
"J8bepjwz9oNknfs6mcqbz3zuEThZQp1VyTk+1wE05gWwmmj1u0CB688nfOnB1xnlBVsAWyLmORxFAw
vcDCvGbTragp/+nT0k94L08Z7Q04I5KqhHfTvINGU54S4IVLWiMIFj8kSU4RhMvNNNq5Tarz0gJZhWR+
4RoNF+84R/uFFGwKeblrwfHxWZZMoW7pEJ2kzLFBtBzx2mUvRP0kn6AFbyS4+DM2YTPQkSk3W2Gzc
+nfPDDyH8Tq6AM5WsVCKqnZAa2ZIY1FxMkKT7iFt5oiYnl93ka7OrQlmM9QAyPoyw/JT0nXHDuf683uE
K32yn9CgxNGXy5NcNzRurdFNb5w=="
    "id": "storageperformance",
    "displayName": "Array Performance",
    "isMandatory": false,
    "attributes": {
      "password": "this will not be encrypted on the server side"
 "acquisitionUnit": {
  "id": "1"
}
```

Getting Started

Feature Tutorials

Cloud Insights is loaded with useful features that enable you to quickly and easily find data, troubleshoot issues, and provide insights into your corporate environment. Find data easily with powerful queries, visualize data in dashboards, and send email alerts for data thresholds you set.

Cloud Insights includes a number of video tutorials to help you understand these features and better implement your business insight strategies. Every user who has access to your Cloud Insights environment can take advantage of these tutorials.

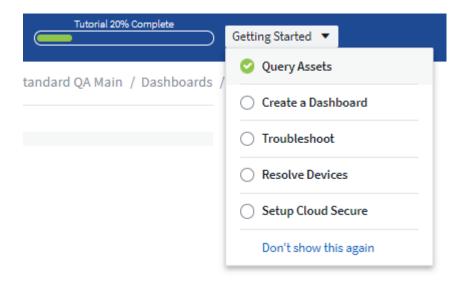
Introduction

Watch a brief tutorial explaining how Cloud Insights works.

► https://docs.netapp.com/us-en/cloudinsights//media/howTo.mp4 (video)

Checklist and Video Tutorials

The **Startup Checklist** displayed on your Cloud Insights site contains a list of several useful tasks and concepts. Selecting an item in the checklist takes you to the appropriate Cloud Insights page for that concept. For example, clicking on the *Create a Dashboard* item opens the Cloud Insights **Dashboards** page.



At the top of the page is a link to a video tutorial showing how to create a dashboard. You can view the video as many times as you like until you click the *Got it! Don't Show This Again* link for that video. The video is available every time you go to the Dashboards page, until you dismiss it.



After watching the video at least once, the *Create a Dashboard* item in the checklist is checked off, indicating that you have completed the tutorial. You can then proceed to the next tutorial.





Dismissing the Checklist

The Startup Checklist is displayed on your site until you click the *Don't Show This Again* link at the bottom of the checklist. Even after dismissing the checklist, the tutorials are still available on each appropriate Cloud Insights page until you dismiss each one from the message header bar.

You can view the tutorials in any order you like, as many times as you like until dismissed.

View the Tutorials

Querying Data

▶ https://docs.netapp.com/us-en/cloudinsights//media/Queries.mp4 (video)

Creating a Dashboard

https://docs.netapp.com/us-en/cloudinsights//media/Dashboards.mp4 (video)

Troubleshooting

https://docs.netapp.com/us-en/cloudinsights//media/Troubleshooting.mp4 (video)

Resolve Devices

► https://docs.netapp.com/us-en/cloudinsights//media/AHR small.mp4 (video)

Collecting Data

Getting started gathering data

After you have signed up for Cloud Insights and log in to your environment for the first time, you will be guided through the following steps in order to begin collecting and managing data.

Data collectors discover information from your data sources, such as storage devices, network switches, and virtual machines. The information gathered is used for analysis, validation, monitoring and troubleshooting.

Cloud Insights has available three types of data collectors:

- Infrastructure (storage devices, network switches, compute infrastructure)
- Operating Systems (such as VMware or Windows)
- · Services (such as Kafka)

Select your first data collector from the supported vendors and models available. You can easily add additional data collectors later.

Install an Acquisition Unit

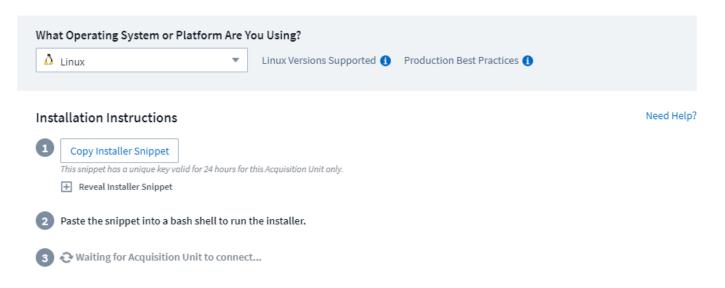
If you selected an *Infrastructure* data collector, an Acquisition Unit is required to inject data into Cloud Insights. You will need to download and install the Acquisition Unit software on a server or VM on the data center from which you will be collecting. A single Acquisition Unit can be used for multiple data collectors.



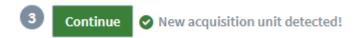
ONTAP Data Management Software

Install Acquisition Unit

Cloud Insights collects device data via one or more Acquisition Units installed on local servers. Each Acquisition Unit can host multiple Data Collectors, which send device metrics to Cloud Insights for analysis.



• Follow the instructions displayed to install your Acquisition Unit. Once the Acquisition Unit software is installed, the Continue button is displayed and you can proceed to the next step.



You may set up additional acquisition units later if needed. For example, you may want different Acquisition Units collecting information from data centers in different regions.

Configure the Data Collector - Infrastructure

For Infrastructure data collectors, you will be asked to fill out the data collector fields presented:

- Give the data collector a unique and meaningful name.
- Enter the credentials (user name and password) to connect to the device, as appropriate.
- Fill in any other mandatory fields in Configuration and Advanced Configuration sections.
- Click Add Collector to save the data collector.

You will be able to configure additional data collectors later.

Configure the Data Collector - Operating Systems and Services

Operating System:

For *Operating System* data collectors, choose a platform (Linux, Windows) to install a Cloud Insights Agent. You must have at least one agent to collect data from Services.

The agent also collects data from the host itself, for use in Cloud Insights. This data is categorized as "Node" data in widgets, etc.

 Open a terminal or command window on the agent host or VM, and paste the displayed command to install the agent. When installation is complete, click Complete Setup.

Services:

For Service data collectors, click on a tile to open the instructions page for that service.

- · Choose a platform and an Agent Access Key.
- If you don't have an agent installed on that platform, follow the instructions to install the agent.
- Click **Continue** to open the data collector instruction page.
- Follow the instructions to configure the data collector.
- When configuration is complete, click Complete Setup.

Add Dashboards

Depending on the type of initial data collector you selected to configure (storage, switch, etc.), one or more relevant dashboards will be imported. For example, if you configured a storage data collector, a set of storage-related dashboards will be imported, and one will be set as your Cloud Insights Home Page. You can change the home page from the **Dashboards > Show All Dashboards** list.

You can import additional dashboards later, or create your own.

That's all there is to it

After you complete the initial setup process, your environment will begin to collect data.

If your initial setup process is interrupted (for example, if you close the browser window), you will need to follow the steps manually:

- · Choose a Data Collector
- · Install an Agent or Acquisition Unit if prompted
- · Configure the Data Collector

Useful definitions

The following definitions may be useful when talking about Cloud Insights data collectors or features:

- Collector life cycle: A collector will belong to one of the following states in its life cycle:
 - Preview: Available in a limited capacity or to a limited audience. Preview features and data collectors
 are expected to become GA following the preview period. Preview periods vary based on audience or
 functionality.
 - GA: A feature or data collector that is Generally Available to all customers, based on Edition or feature set.
 - Deprecated: Applies to data collectors that are, or are expected to become, no longer functionally sustainable. Deprecated data collectors are often replaced with newer, functionally-updated data collectors.
 - Deleted: A data collector that has been removed and is no longer available.
- Acquisition Unit: a computer dedicated to hosting data collectors, typically a Virtual Machine. This computer is typically located in the same data center / VPC as the monitored items.
- Data Source: a module for communicating with a hardware or software stack. It consists of a configuration

Acquisition Unit Requirements

You must install an Acquisition Unit (AU) in order to acquire information from your infrastructure data collectors (storage, VM, port, EC2, etc.). Before you install the Acquisition Unit, you should ensure that your environment meets operating system, CPU, memory, and disk space requirements.

Requirements

Component	Linux Requirement	Windows Requirement
Operating system	A computer running a licensed version of one of the following: * Centos (64-bit): 7.2 through 7.9, 8.1 through 8.4, Stream 8, Stream 9 * Debian (64-bit): 9 and 10 * OpenSUSE Leap 15.1 through 15.5 * Oracle Enterprise Linux (64-bit): 7.5 through 7.9, 8.1 through 8.8 * Red Hat Enterprise Linux (64-bit): 7.2 through 7.9, 8.1 through 8.8, 9.1, 9.2 * Rocky 9.0, 9.1, 9.3	A computer running a licensed version of one of the following: * Microsoft Windows 10 64-bit * Microsoft Windows Server 2012 * Microsoft Windows Server 2012 R2 * Microsoft Windows Server 2016 * Microsoft Windows Server 2019 * Microsoft Windows Server 2022 * Microsoft Windows 11 This computer should be running no other application-level software.
CPU	2 CPU cores	Same

Available disk space	50 GB (100 GB recommended) For Linux, disk space should be allocated in this manner: /opt/netapp 10 GB (20 GB for large environments) /var/log/netapp 40 GB (80 GB for large environments) /tmp at least 1 GB available during installation	50 GB
Network	100 Mbps/1 Gbps Ethernet connection, static IP address, and port 80 or 443 connectivity from Acquisition Unit to *.cloudinsights.netapp.com or your Cloud Insights environment (i.e. https:// <environment_id>.c01.cloudinsights.netapp.com) is required. For requirements between Acquisition Unit and each Data Collector, please refer to instructions for the Data Collector. If your organization requires proxy usage for internet access, you may need to understand your organization's proxy behavior and seek certain exceptions for Cloud Insights to work. For example, does your organization block access by default, and only allow access to specific web sites/domains by exception? If so, you will need to get the following domain added to the exception list: *.cloudinsights.netapp.com For more information, read about Proxies here (Linux) or here (Windows).</environment_id>	Same
Permissions	Sudo permissions on the Acquisition Unit server. /tmp must be mounted with exec capabilities.	Administrator permissions on the Acquisition Unit server
Virus Scan		During installation, you must completely disable all virus scanners. Following installation, the paths used by the Acquisition Unit software must be excluded from virus scanning.

Additional recommendations

 For accurate audit and data reporting, it is strongly recommended to synchronize the time on the Acquisition Unit machine using Network Time Protocol (NTP) or Simple Network Time Protocol (SNTP).

Regarding Sizing

You can get started with a Cloud Insights Acquisition Unit with just 8GB memory and 50GB of disk space, however, for larger environments you should ask yourself the following questions:

Do you expect to:

- Discover more than 2500 virtual machines or 10 large (> 2 node) ONTAP clusters, Symmetrix, or HDS/HPE VSP/XP arrays on this Acquisition Unit?
- Deploy 75 or more total data collectors on this Acquisition Unit?

For each "Yes" answer above, it is recommend to add 8 GB of memory and 50 GB of disk space to the AU. So for example if you answered "Yes" to both, you should deploy a 24GB memory system with 150GB or more of disk space. On Linux, the disk space to be added to the log location.

For additional sizing questions, contact NetApp Support.

Additional Federal Edition requirement

• For Acquisition Unit installations in Cloud Insights Federal Edition clusters, the underlying operating system must have a good source of entropy. On Linux systems this is typically done by installing *rng-tools* or by using hardware random number generation (RNG). It is the customer's responsibility to ensure this requirement is met on the Acquisition Unit machine.

Configuring Acquisition Units

Cloud Insights collects device data using one or more Acquisition Units installed on local servers. Each Acquisition Unit can host multiple Data Collectors, which send device metrics to Cloud Insights for analysis.

This topic describes how to add Acquisition Units and describes additional steps required when your environment uses a proxy.



For accurate audit and data reporting, it is strongly recommended to synchronize the time on the Acquisition Unit machine using **Network Time Protocol (NTP)** or **Simple Network Time Protocol (SNTP)**.

Read about Cloud Insights security here.

Adding a Linux Acquisition Unit

Before you begin

• If your system is using a proxy, you must set the proxy environment variables before the acquisition unit is installed. For more information, see Setting proxy environment variables.

Steps for Linux Acquisition Unit Installation

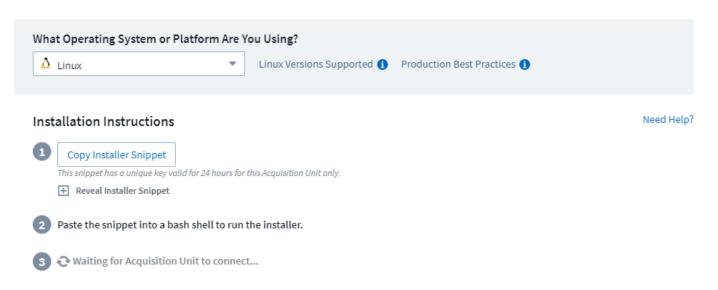
- 1. Log in as Administrator or Account Owner to your Cloud Insights environment.
- 2. Click Observability > Collectors > Acquisition Units > +Acquisition Unit

The system displays the *Install Acquisition Unit* dialog. Choose Linux.



Install Acquisition Unit

ONTAP Data Management Software Cloud Insights collects device data via one or more Acquisition Units installed on local servers. Each Acquisition Unit can host multiple Data Collectors, which send device metrics to Cloud Insights for analysis.



- 1. Verify that the server or VM hosting the Acquisition Unit meets the recommended system requirements.
- 2. Verify that the server is running a supported version of Linux. Click OS Versions Supported (i) for a list of supported versions.
- Copy the Installation command snippet in the dialog into a terminal window on the server or VM that will host the Acquisition unit.
- 4. Paste and execute the command in the Bash shell.

After you finish

- Click **Observability > Collectors > Acquisition units** to check the status of Acquisition Units.
- You can access the Acquisition Unit logs at /var/log/netapp/cloudinsights/acq/acq.log
- Use the following script to control the Acquisition Unit:
 - cloudinsights-service.sh (stop, start, restart, check the status)
- Use the following script to uninstall the Acquisition Unit:
 - · cloudinsights-uninstall.sh

Setting proxy environment variables

For environments that use a proxy, you must set the proxy environment variables before you add the Acquisition Unit. The instructions for configuring the proxy are provided on the *Add Acquisition Unit* dialog.

- 1. Click + in Have a Proxy Server?
- Copy the commands to a text editor and set your proxy variables as needed.

Note: Be aware of restrictions on special characters in proxy username and password fields: '%' and '!' are allowed in the username field. ':', '%', and '!' are allowed in the password field.

- 3. Run the edited command in a terminal using the Bash shell.
- 4. Install the Acquisition Unit software.

Proxy Configuration

The Acquisition Unit uses 2-way/mutual authentication to connect to the Cloud Insights server. The client certificate must be passed to the Cloud Insights server to be authenticated. To accomplish this, the proxy must be set up to forward the https request to the Cloud Insights server without decrypting the data.

The simplest way to do this is to specify wildcard configuration in your proxy/firewall to communicate with Cloud Insights, for example:

*.cloudinsights.netapp.com

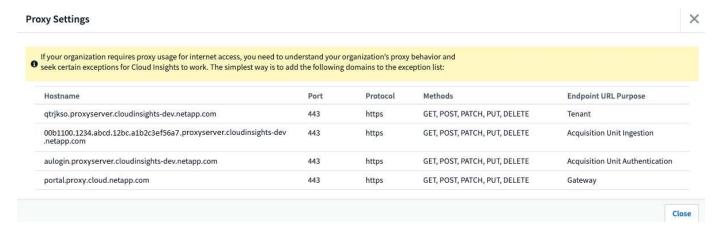


The use of an asterisk (*) for wildcard is common, but your proxy/firewall configuration may use a different format. Check with your proxy documentation to ensure correct wildcard specification in your environment.

More information on proxy configuration can be found in the NetApp Knowledgbase.

Viewing Proxy URLs

You can view your proxy endpoint URLs by clicking the **Proxy Settings** link when choosing a data collector during onboarding, or the link under *Proxy Settings* on the **Help > Support** page. A table like the following is displayed.



If you have Workload Security in your environment, the configured endpoint URLs will also be displayed in this list.

Adding a Windows Acquisition Unit

Steps for Windows Acquisition Unit Installation

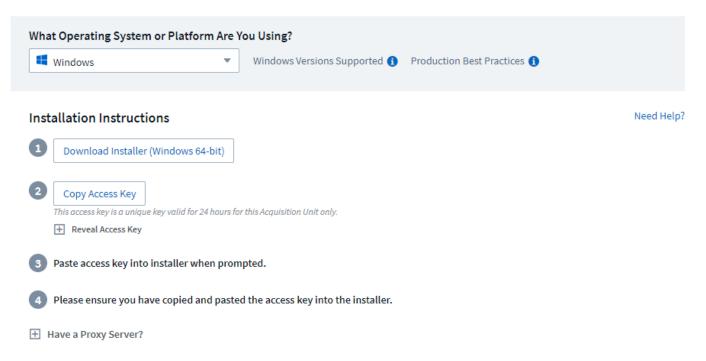
- 1. Log in to the Acquisition Unit server/VM as a user with Administrator permissions.
- On that server, open a browser window and log in to your Cloud Insights environment as Administrator or Account Owner.

3. Click Observability > Collectors > Acquisition Units > +Acquisition Unit .

The system displays the Install Acquisition Unit dialog. Choose Windows.

Install Acquisition Unit

Cloud Insights collects device data via one or more Acquisition Units installed on local servers. Each Acquisition Unit can host multiple Data Collectors, which send device metrics to Cloud Insights for analysis.



- 1. Verify that the server or VM hosting the Acquisition Unit meets the recommended system requirements.
- 2. Verify that the server is running a supported version of Windows. Click OS Versions Supported (i) for a list of supported versions.
- 3. Click the Download Installer (Windows 64-bit) button.
- 4. Copy the Access Key. You will need this during the Installation.
- 5. On the Acquisition Unit server/VM, execute the downloaded installer.
- 6. Paste the Access Key into the installation wizard when prompted.
- 7. During installation, you will be presented with the opportunity to provide your proxy server settings.

After you finish

- Click * > Observability > Collectors > Acquisition units* to check the status of Acquisition Units.
- You can access the Acquisition Unit log in <install dir>\Cloud Insights\Acquisition Unit\log\acq.log
- · Use the following script to stop, start, restart, or check the status of the Acquisition Unit:

cloudinsights-service.sh

Proxy Configuration

The Acquisition Unit uses 2-way/mutual authentication to connect to the Cloud Insights server. The client certificate must be passed to the Cloud Insights server to be authenticated. To accomplish this, the proxy must

be set up to forward the https request to the Cloud Insights server without decrypting the data.

The simplest way to do this is to specify wildcard configuration in your proxy/firewall to communicate with Cloud Insights, for example:

*.cloudinsights.netapp.com

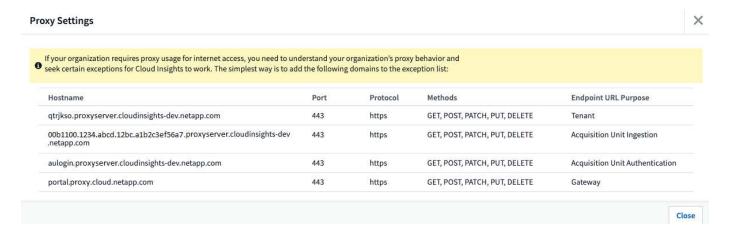


The use of an asterisk (*) for wildcard is common, but your proxy/firewall configuration may use a different format. Check with your proxy documentation to ensure correct wildcard specification in your environment.

More information on proxy configuration can be found in the NetApp Knowledgbase.

Viewing Proxy URLs

You can view your proxy endpoint URLs by clicking the **Proxy Settings** link when choosing a data collector during onboarding, or the link under *Proxy Settings* on the **Help > Support** page. A table like the following is displayed.



If you have Workload Security in your environment, the configured endpoint URLs will also be displayed in this list.

Uninstalling an Acquisition Unit

To uninstall the Acquisition Unit software, do the following:

Windows:

If you are uninstalling a Windows acquisition unit:

- 1. On the Acquisition Unit server/VM, open Control Panel and choose **Uninstall a Program**. Select the Cloud Insights Acquisition Unit program for removal.
- 2. Click Uninstall and follow the prompts.

Linux:

If you are uninstalling a Linux acquisition unit:

1. On the Acquisition Unit server/VM, run the following command:

```
sudo cloudinsights-uninstall.sh -p
```

2. For help with uninstall, run:

```
sudo cloudinsights-uninstall.sh --help
```

Windows and Linux:

After uninstalling the AU:

- 1. In Cloud Insights, go to Observability > Collectors and select the *Acquisition Units tab.
- 2. Click the Options button to the right of the Acquisition Unit you wish to uninstall, and select *Delete*. You can delete an Acquisition Unit only if there are no data collectors assigned to it.



You cannot delete an Acquisition Unit (AU) that has data collectors connected to it. Move all of the AU's data collectors to another AU (edit the collector and simply select a different AU) before deleting the original AU.

An Acquisition unit with a star next to it is being used for device resolution. Before removing this AU, you must select another AU to use for Device Resolution. Hover over a different AU and open the "three dots" menu to select "Use for Device Resolution".



Reinstalling an Acquisition Unit

To re-install an Acquisition Unit on the same server/VM, you must follow these steps:

Before you begin

You must have a temporary Acquisition Unit configured on a separate server/VM before re-installing an Acquisition Unit.

Steps

- 1. Log in to the Acquisition Unit server/VM and uninstall the AU software.
- 2. Log into your Cloud Insights environment and go to **Observability > Collectors**.
- 3. For each data collector, click the Options menu on the right and select *Edit*. Assign the data collector to the temporary Acquisition Unit and click **Save**.

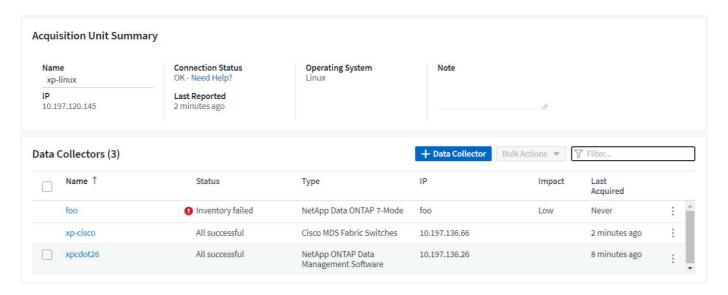
You can also select multiple data collectors of the same type and click the **Bulk Actions** button. Choose *Edit* and assign the data collectors to the temporary Acquisition Unit.

- 4. After all of the data collectors have been moved to the temporary Acquisition Unit, go to **Observability > Collectors** and select the **Acquisition Units** tab.
- 5. Click the Options button to the right of the Acquisition Unit you wish to re-install, and select *Delete*. You can delete an Acquisition Unit only if there are no data collectors assigned to it.
- 6. You can now re-install the Acquisition Unit software on the original server/VM. Click **+Acquisition Unit** and follow the instructions above to install the Acquisition Unit.
- 7. Once the Acquisition Unit has been re-installed, assign your data collectors back to the Acquisition Unit.

Viewing AU Details

The Acquisition Unit (AU) detail page provides useful detail for an AU as well as information to help with troubleshooting. The AU detail page contains the following sections:

- A **summary** section showing the following:
 - Name and IP of the Acquisition Unit
 - · Current connection Status of the AU
 - Last Reported successful data collector poll time
 - The Operating System of the AU machine
 - Any current **Note** for the AU. Use this field to enter a comment for the AU. The field displays the most recently added note.
- A table of the AU's Data Collectors showing, for each data collector:
 - · Name Click this link to drill down into the data collector's detail page with additional information
 - Status Success or error information
 - Type Vendor/model
 - · IP address of the data collector
 - Current Impact level
 - · Last Acquired time when the data collector was last successfully polled



For each data collector, you can click on the "three dots" menu to Clone, Edit, Poll, or Delete the data collector. You can also select multiple data collectors in this list to perform bulk actions on them.

To restart the Acquisition Unit, click the **Restart** button at the top of the page. Drop down this button to attempt to **Restore Connection** to the AU in the event of a connection problem.

Configuring an Agent to Collect Data (Windows/Linux)

Cloud Insights uses Telegraf as its agent for collection of integration data. Telegraf is a plugin-driven server agent that can be used to collect and report metrics, events, and logs. Input plugins are used to collect the desired information into the agent by accessing the system/OS directly, by calling third-party APIs, or by listening to configured streams (i.e. Kafka, statsD, etc). Output plugins are used to send the collected metrics, events, and logs from the agent to Cloud Insights.

The current Telegraf version for Cloud Insights is 1.24.0.

For information on installing on Kubernetes, see the NetApp Kubernetes Monitoring Operator page.



For accurate audit and data reporting, it is strongly recommended to synchronize the time on the Agent machine using **Network Time Protocol (NTP)** or **Simple Network Time Protocol (SNTP)**.



If you want to verify the installation files before instaling the Agent, see the section below on Verifying Checksums.

Installing an Agent

If you are installing a Service data collector and have not yet configured an Agent, you are prompted to first install an Agent for the appropriate Operating System. This topic provides instructions for installing the Telegraf agent on the following Operating Systems:

- Windows
- RHEL and CentOS
- Ubuntu and Debian

To install an agent, regardless of the platform you are using, you must first do the following:

- 1. Log into the host you will use for your agent.
- 2. Log in to your Cloud Insights environment and navigate to Observability > Collectors.
- 3. Click on +Data Collector and choose a data collector to install.
- 4. Choose the appropriate platform for your host (Windows, Linux)
- 5. Follow the remaining steps for each platform.



Once you have installed an agent on a host, you do not need to install an agent again on that host.



Once you have installed an agent on a server/VM, Cloud Insights collects metrics from that system in addition to collecting from any data collectors you configure. These metrics are gathered as "Node" metrics.



If you are using a proxy, read the proxy instructions for your platform before installing the Telegraf agent.

Log Locations

Telegraf log messages are redirected from stdout to the following log files be default:

- RHEL/CentOS: /var/log/telegraf/telegraf.log
- Ubuntu/Debian: /var/log/telegraf/telegraf.log
- Windows: C:\Program Files\telegraf\telegraf.log

Windows

Pre-requisites:

- · PowerShell must be installed
- If you are behind a proxy, you must follow the instructions in the Configuring Proxy Support for Windows section.

Configuring Proxy Support for Windows



If your environment uses a proxy, read this section before you install.



The steps below outline the actions needed to set the https_proxy environment variables. For some proxy environments, users may also need to set the no_proxy environment variable.

For systems residing behind a proxy, perform the following to set the *https_proxy* and/or *http_proxy* environment variable(s) **PRIOR** to installing the Telegraf agent:

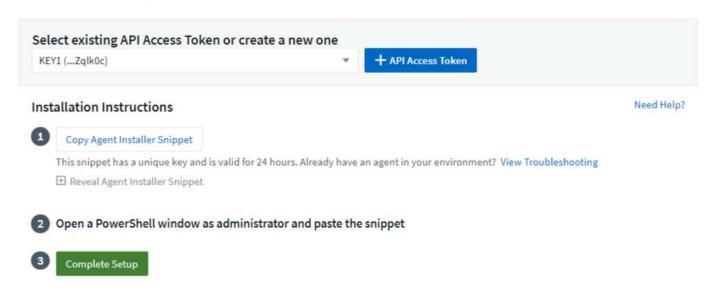
```
[System.Environment]::SetEnvironmentVariable("https_proxy",
"cproxy_server>:cystem.EnvironmentVariableTarget]::Machine)
```

Installing the agent



Install Agent

Quickly setup an agent in your environment and immediately start monitoring data



Steps to install agent on Windows:

- 1. Choose an Agent Access Key.
- 2. Copy the command block from the agent installation dialog. You can click the clipboard icon to quickly copy the command to the clipboard.
- Open a PowerShell window
- 4. Paste the command into the PowerShell window and press Enter.
- 5. The command will download the appropriate agent installer, install it, and set a default configuration. When finished, it will restart the agent service. The command has a unique key and is valid for 24 hours.
- Click Finish or Continue

After the agent is installed, you can use the following commands to start/stop the service:

```
Start-Service telegraf
Stop-Service telegraf
```

Uninstalling the Agent

To uninstall the agent on Windows, do the following in a PowerShell window:

1. Stop and delete the Telegraf service:

```
Stop-Service telegraf
sc.exe delete telegraf
```

Remove the certificate from the trustore:

cd Cert:\CurrentUser\Root
//rm E5FB7B68C08B1CA902708584C274F8EFC7BE8ABC
rm 1A918038E8E127BB5C87A202DF173B97A05B4996

- 3. Delete the C:\Program Files\telegraf folder to remove the binary, logs, and configuration files
- 4. Remove the SYSTEM\CurrentControlSet\Services\EventLog\Application\telegraf key from the registry

Upgrading the Agent

To upgrade the telegraf agent, do the following:

1. Stop and delete the telegraf service:

Stop-Service telegraf sc.exe delete telegraf

- 2. Delete the SYSTEM\CurrentControlSet\Services\EventLog\Application\telegraf key from the registry
- 3. Delete C:\Program Files\telegraf\telegraf.conf
- 4. Delete C:\Program Files\telegraf\telegraf.exe
- 5. Install the new agent.

RHEL and CentOS

Pre-requisites:

- The following commands must be available: curl, sudo, ping, sha256sum, openssl, and dmidecode
- If you are behind a proxy, you must follow the instructions in the Configuring Proxy Support for RHEL/CentOS section.

Configuring Proxy Support for RHEL/CentOS



If your environment uses a proxy, read this section before you install.



The steps below outline the actions needed to set the https_proxy environment variables. For some proxy environments, users may also need to set the no_proxy environment variable.

For systems residing behind a proxy, perform the following steps **PRIOR** to installing the Telegraf agent:

1. Set the https_proxy and/or http_proxy environment variable(s) for the current user:

```
export https_proxy=<proxy_server>:<proxy_port>
```

2. Create /etc/default/telegraf, and insert definitions for the https_proxy and/or http_proxy variable(s):

```
https_proxy=<proxy_server>:<proxy_port>
```

Installing the agent



Install Agent

Quickly setup an agent in your environment and immediately start monitoring data

	elect existing API Access Token or create a new default_ingestion_api_key1 (xEKVyK)	▼	+ API Access Token	Production Best Practices @
Ins	stallation Instructions			Need Help?
1	For environments operating behind a proxy serve Telegraf.	er, follow t	he instructions to conf	igure proxy support to install and run
2	Copy Agent Installer Snippet			
	This snippet has a unique key and is valid for 24 hours. Alr	eady have a	n agent in your environme	nt? View Troubleshooting
	Reveal Agent Installer Snippet			
3	Open a terminal window and paste the snippet in	n a Bash sh	ell (requires curl, sudo	, ping, sha256sum, and dmidecode).

Steps to install agent on RHEL/CentOS:

- 1. Choose an Agent Access Key.
- 2. Copy the command block from the agent installation dialog. You can click the clipboard icon to quickly copy the command to the clipboard.
- 3. Open a Bash window
- 4. Paste the command into the Bash window and press Enter.
- 5. The command will download the appropriate agent installer, install it, and set a default configuration. When finished, it will restart the agent service. The command has a unique key and is valid for 24 hours.
- 6. Click Finish or Continue

After the agent is installed, you can use the following commands to start/stop the service:

If your operating system is using systemd (CentOS 7+ and RHEL 7+):

```
sudo systemctl start telegraf
sudo systemctl stop telegraf
```

If your operating system is not using systemd (CentOS 7+ and RHEL 7+):

```
sudo service telegraf start
sudo service telegraf stop
```

Uninstalling the Agent

To uninstall the agent on RHEL/CentOS, in a Bash terminal, do the following:

1. Stop the Telegraf service:

```
systemctl stop telegraf (If your operating system is using systemd
(CentOS 7+ and RHEL 7+)
/etc/init.d/telegraf stop (for systems without systemd support)
```

2. Remove the Telegraf agent:

```
yum remove telegraf
```

3. Remove any configuration or log files that may be left behind:

```
rm -rf /etc/telegraf*
rm -rf /var/log/telegraf*
```

Upgrading the Agent

To upgrade the telegraf agent, do the following:

1. Stop the telegraf service:

```
systemctl stop telegraf (If your operating system is using systemd
(CentOS 7+ and RHEL 7+)
/etc/init.d/telegraf stop (for systems without systemd support)
```

2. Remove the previous telegraf agent:

```
yum remove telegraf
```

3. Install the new agent.

Ubuntu and Debian

Pre-requisites:

- The following commands must be available: curl, sudo, ping, sha256sum, openssl, and dmidecode
- If you are behind a proxy, you must follow the instructions in the Configuring Proxy Support for Ubuntu/Debian section.

Configuring Proxy Support for Ubuntu/Debian



If your environment uses a proxy, read this section before you install.



The steps below outline the actions needed to set the https_proxy environment variables. For some proxy environments, users may also need to set the no_proxy environment variable.

For systems residing behind a proxy, perform the following steps **PRIOR** to installing the Telegraf agent:

1. Set the https proxy and/or http proxy environment variable(s) for the current user:

```
export https_proxy=<proxy_server>:<proxy_port>
```

2. Create /etc/default/telegraf, and insert definitions for the https_proxy and/or http_proxy variable(s):

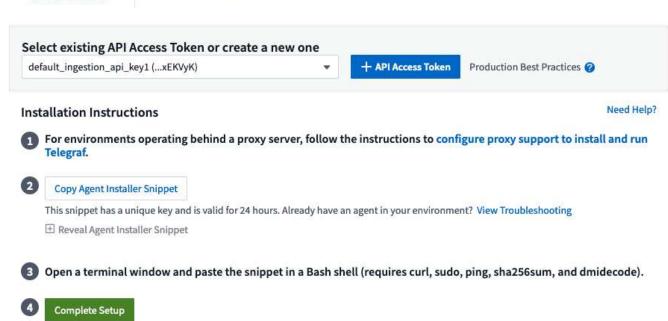
```
https_proxy=<proxy_server>:<proxy_port>
```

Installing the agent



Install Agent

Quickly setup an agent in your environment and immediately start monitoring data



Steps to install agent on Debian or Ubuntu:

- 1. Choose an Agent Access Key.
- 2. Copy the command block from the agent installation dialog. You can click the clipboard icon to quickly copy the command to the clipboard.
- 3. Open a Bash window
- 4. Paste the command into the Bash window and press Enter.
- 5. The command will download the appropriate agent installer, install it, and set a default configuration. When finished, it will restart the agent service. The command has a unique key and is valid for 24 hours.
- 6. Click Finish or Continue

After the agent is installed, you can use the following commands to start/stop the service:

If your operating system is using systemd:

```
sudo systemctl start telegraf
sudo systemctl stop telegraf
```

If your operating system is not using systemd:

```
sudo service telegraf start
sudo service telegraf stop
```

Uninstalling the Agent

To uninstall the agent on Ubuntu/Debian, in a Bash terminal, run the following:

1. Stop the Telegraf service:

```
systemctl stop telegraf (If your operating system is using systemd)
/etc/init.d/telegraf stop (for systems without systemd support)
```

2. Remove the Telegraf agent:

```
dpkg -r telegraf
```

3. Remove any configuration or log files that may be left behind:

```
rm -rf /etc/telegraf*
rm -rf /var/log/telegraf*
```

Upgrading the Agent

To upgrade the telegraf agent, do the following:

1. Stop the telegraf service:

```
systemctl stop telegraf (If your operating system is using systemd)
/etc/init.d/telegraf stop (for systems without systemd support)
```

2. Remove the previous telegraf agent:

```
dpkg -r telegraf
```

3. Install the new agent.

Verifying Checksums

The Cloud Insights agent installer performs integrity checks, but some users may want to perform their own verifications before installing or applying downloaded artifacts. This can be done by downloading the installer and generating a checksum for the downloaded package, then comparing the checksum to the value shown in the install instructions.

Download the installer package without installing

To perform a download-only operation (as opposed to the default download-and-install), users can edit the agent installation command obtained from the UI and remove the trailing "install" option.

Follow these steps:

- 1. Copy the Agent Installer snippet as directed.
- 2. Instead of pasting the snippet into a command window, paste it into a text editor.
- 3. Remove the trailing "--install" (Linux) or "-install" (Windows) from the command.
- 4. Copy the entire command from the text editor.
- 5. Now paste it into your command window (in a working directory) and run it.

Non-Windows (these examples are for Kubernetes; actual script names may vary):

· Download and install (default):

```
installerName=cloudinsights-kubernetes.sh ... && sudo -E -H ./$installerName --download --install
```

· Download-only:

```
installerName=cloudinsights-kubernetes.sh ... && sudo -E -H ./$installerName --download
```

Windows:

· Download and install (default):

```
!$($installerName=".\cloudinsights-windows.ps1") ... -and
$(&$installerName -download -install)
```

· Download-only:

```
!$($installerName=".\cloudinsights-windows.ps1") ... -and
$(&$installerName -download)
```

The download-only command will download all required artifacts from Cloud Insights to the working directory. The artifacts include, but may not be limited to:

- · an installation script
- · an environment file
- YAML files
- a checksum file (ending in sha256.signed or sha256.ps1)

The installation script, environment file, and YAML files can be verified using visual inspection.

Generate checksum value

To generate the checksum value, perform the following command for your appropriate platform:

• RHEL/Ubuntu:

```
sha256sum <package_name>
```

· Windows:

```
Get-FileHash telegraf.zip -Algorithm SHA256 | Format-List
```

Verify checksum

Extract the expected checksum from the checksum file

Non-Windows:

```
openssl smime -verify -in telegraf*.sha256.signed -CAfile netapp_cert.pem -purpose any -nosigs -noverify
```

• Windows:

```
(Get-Content telegraf.zip.sha256.ps1 -First 1).toUpper()
```

Install the downloaded package

Once all of the artifacts have been satisfactorily verified, the agent installation can be initiated by running:

Non-Windows:

```
sudo -E -H ./<installation_script_name> --install
```

Windows:

.\cloudinsights-windows.ps1 -install

Troubleshooting

Some things to try if you encounter problems setting up an agent:

Problem:	Try this:
After configuring a new plugin and restarting Telegraf, Telegraf fails to start up. The logs indicate that an error resembling the following:	The installed Telegraf version is outdated. Follow the steps on this page to Upgrade the Agent for your appropriate platform.
"[telegraf] Error running agent: Error loading config file /etc/telegraf/telegraf.d/cloudinsights-default.conf: plugin outputs.http: line linenumber>: configuration specified the fields ["use_system_proxy"], but they weren't used"	
I ran the installer script on an old installation and now the agent isn't sending data	Uninstall the telegraf agent and then re-run the installation script. Follow the Upgrade the Agent steps on this page for your appropriate platform.
I already installed an agent using Cloud Insights	If you have already installed an agent on your host/VM, you do not need to install the agent again. In this case, simply choose the appropriate Platform and Key in the Agent Installation screen, and click on Continue or Finish .
I already have an agent installed but not by using the Cloud Insights installer	Remove the previous agent and run the Cloud Insights Agent installation, to ensure proper default configuration file settings. When complete, click on Continue or Finish .

Additional information may be found from the Support page or in the Data Collector Support Matrix.

Configuring Data Collectors

You configure Data Collectors in your Cloud Insights environment to collect data from devices in the data center.

Before you begin

- · You must have configured an Acquisition Unit before you can start collecting data.
- You need credentials for the devices from which you are collecting Data.
- Device network addresses, account information, and passwords are required for all devices you are collecting data from.

Steps

1. From the Cloud Insights menu, click Observability > Collectors

The system displays the available Data Collectors arranged by vendor.

2. Click + Collector and select the data collector to configure.

In the dialog box you can configure the data collector and add an Acquisition Unit.

3. Enter a name for the data collector.

Names can contain can contain letters (a-z), numbers (0-9), dashes (-), underscores (_), apostrophes ('), and periods (.).

- 4. Enter the Acquisition Unit to associate with this data collector.
- 5. Enter the required fields in the Configuration screen.
- 6. When prompted to configure notifications, choose to alert by Email, Webhook, or both, and choose the alert types on which to notify (Critical, Warning, Informational, and/or Resolved). You can choose to notify to the Global Monitor Recipient list (configured in **Admin > Notifications**), or specify additional recipients. When ready to continue, click **Complete Setup**.

Customize notifications for this collector ONTAP Default monitors are preconfigured to send email notifications to "Global Monitor Recipient List", you can add additional email addresses for this data collector. Send to Global Monitor Recipient List Gritical, Warning, Informa... Other Email Recipients By Webhook Enable webhook notification to add recipients

When viewing an **ONTAP data collector** landing page, you can modify the notifications by clicking the pencil icon in the "Notifications" field of the data collector summary section.



ONTAP Data Collector notifications take precedence over any specific Monitor notifications that are relevant to the cluster/data collector. The recipient list you set for the Data Collector itself will receive the data collector alerts. If there are no active data collector alerts, then monitor-generated alerts will be sent to specific monitor recipients.



- Click Advanced Configuration to add additional configuration fields. (Not all data collectors require advanced configuration.)
- Click Test Configuration to verify that the data collector is properly configured.
- 3. Click Add Collector to save the configuration and add the data collector to your Cloud Insights tenant.

After adding a new data collector, Cloud Insights initiates three polls:

- 1st inventory poll: immediately
- 1st performance data poll to establish a baseline: immediately after inventory poll
- 2nd performance poll: within 15 seconds after completion of 1st performance poll

Polling then proceeds according to the configured inventory and performance poll intervals.

Determining data collector acquisition status

Because data collectors are the primary source of information for Cloud Insights, it is imperative that you ensure that they remain in a running state.

Data collector status is displayed in the upper right corner of any asset page as the message "Acquired N minutes ago", where N indicates the most recent acquisition time of the asset's data collector(s). The acquisition time/date is also displayed.

Clicking on the message displays a table with data collector name, status, and last successful acquisition time. If you are signed in as an Administrator, clicking on the data collector name link in the table takes you to detail page for that data collector.

Managing configured data collectors

The Installed Data Collectors page provides access to the data collectors that have been configured for Cloud Insights. You can use this page to modify existing data collectors.

Steps

1. In the Cloud Insights menu, click Observability > Collectors

The Available Data Collectors screen is displayed.

2. Click Installed Data Collectors

A list of all of the installed Data Collectors is displayed. The list provides collector name, status, the IP address the collector is accessing, and when data was last acquired from the a device. Action that can be performed on this screen include:

- Control polling
- Change data collector credentials
- Clone data collectors

Controlling Data Collector polling

After making a change to a data collector, you might want it to poll immediately to check your changes, or you might want to postpone the data collection on a data collector for one, three, or five days while you work on a problem.

Steps

- 1. In the Cloud Insights menu, click Observability > Collectors
- 2. Click Installed Data Collectors
- 3. Select the check box to the left of the Data Collector you want to change
- 4. Click **Bulk Actions** and select the polling action you want to take.

Bulk actions can be performed simultaneously on multiple Data Collectors. Select the data collectors, and chose the action to perform from the **Bulk Action** menu.

Editing data collector information

You can edit existing data collector setup information.

To edit a single data collector:

- 1. In the Cloud Insights menu, click **Observability > Collectors** to open the list of installed Data Collectors.
- 2. In the options menu to the right of the data collector you want to modify, click Edit.

The Edit Collector dialog is opened.

3. Enter the changes and click **Test Configuration** to test the new configuration or click **Save** to save the configuration.

You can also edit multiple data collectors:

- 1. Select the check box to the left of each data collector you want to change.
- Click the Bulk Actions button and choose Edit to open the Edit data Collector dialog.
- 3. Modify the fields as above.



The data collectors selected must be the same vendor and model, and reside on the same Acquisition Unit.

When editing multiple data collectors, the Data Collector Name field shows "Mixed" and cannot be edited. Other fields such as user name and password show "Mixed" and can be edited. Fields that share the same value across the selected data collectors show the current values and can be edited.

When editing multiple data collectors, the **Test Configuration** button is not available.

Cloning data collectors

Using the clone facility, you can quickly add a data source that has the same credentials and attributes as another data source. Cloning allows you to easily configure multiple instances of the same device type.

Steps

- 1. In the Cloud Insights menu, click Observability > Collectors.
- 2. Click Installed Data Collectors.
- 3. Click the check box to the left of the data collector you want to copy.
- 4. In the options menu to the right of the selected data collector, click **Clone**.

The Clone Data Collector dialog is displayed.

- 5. Enter new information in the required fields.
- 6. Click Save.

After you finish

The clone operation copies all other attributes and settings to create the new data collector.

Performing bulk actions on data collectors

You can simultaneously edit some information for multiple data collectors. This feature allows you to initiate a poll, postpone polling, and resume polling on multiple data collectors. In addition, you can delete multiple data collectors.

Steps

- 1. In the Cloud Insights menu, click Observability > Collectors
- 2. Click Installed Data Collectors
- 3. Click the check box to the left of the data collectors you want to modify.
- 4. In the options menu to the right, click the option you want to perform.

After you finish

The operation you selected is performed on the data collectors. When you chose to delete data collectors, a dialog is displayed requiring you to conform the action.

Researching a failed data collector

If a data collector has failure message and a High or Medium Impact, you need to research this problem using the data collector summary page with its linked information.

Use the following steps to determine the cause of failed data collectors. Data collector failure messages are displayed on the **Admin** menu and on the **Installed Data Collectors** page.

Steps

- 1. Click Admin > Data Collectors > Installed Data Collectors.
- 2. Click the linked Name of the failing data collector to open the Summary page.
- 3. On the Summary page, check the Comments area to read any notes that might have been left by another engineer who might also be investigating this failure.

- 4. Note any performance messages.
- 5. Move your mouse pointer over the segments of the Event Timeline graph to display additional information.
- 6. Select an error message for a Device and displayed below the Event Timeline and click the Error details icon that displays to the right of the message.
 - The Error details include the text of the error message, most likely causes, information in use, and suggestions of what can be tried to correct the problem.
- 7. In the Devices Reported By This Data Collector area, you might filter the list to display only devices of interest, and you can click the linked **Name** of a device to display the asset page for that device.
- 8. When you return to the data collector summary page, check the **Show Recent Changes** area at the bottom of the page to see if recent changes could have caused the problem.

Importing from the Dashboard Gallery

Cloud Insights provides a number of Recommended Dashboards to provide business insights into your data. Each dashboard contains widgets designed to help answer a particular question or solve a particular problem relevant to the data currently being collected in your environment.

To import a dashboard from the gallery, do the following:

- 1. Select Dashboards > Dashboards
- 2. Click on +From Gallery

A list of **Recommended Dashboards** is displayed. Each dashboard is named with a particular question the dashboard can help you solve. Dashboards are available to help answer questions around different types of objects, including AWS, NetApp, Storage, VMware, and others

3. Select one or more dashboards from the list and click **Add Dashboards**. These dashboards now show in your dashboard list.

In addition to the Recommended Dashboards, you can also choose to import **Additional Dashboards** that are not relevant to your current data. For example, if you have no storage data collectors currently installed but are planning on configuring some in the future, you may still choose to import the storage-relevant dashboards. These dashboards will be available for display but may not show any relevant data until at least one storage data collector is configured.

Importing from the dashboard gallery is available to users with Administrator or Account Owner role.

User Accounts and Roles

Cloud Insights provides up to four user account roles: Account Owner, Administrator, User, and Guest. Each account is assigned specific permission levels as noted in the table below. Users are either invited to Cloud Insights and assigned a specific role, or can sign in via Single Sign-On (SSO) Authorization with a default role. SSO Authorization is available as a feature in Cloud Insights Premium Edition.



User logins in Cloud Insights Federal Edition are limited to configured identity providers (with their specified email domains). When a new user is invited to a Cloud Insights Federal environment, their email address must match the domain configured for that environment.

Permission levels

You use an account that has Administrator privileges to create or modify user accounts. Each user account is assigned a role for each Cloud Insights feature from the following permission levels.

Role	Observability	Workload Security	Reporting	
Account Owner	Administrator functions for Owners can also invite and Authentication and Identity The first Account Owner is	Can modify subscriptions, view billing and usage information, and perform all Administrator functions for Observability, Security, and Reporting. Owners can also invite and manage users, as well as manage SSO Authentication and Identity Federation settings. The first Account Owner is created when you register for Cloud Insights. It is strongly recommended to have at least two Account Owners for each Cloud Insights environment.		
Administrator	Can perform all Observability functions, all user functions, as well as management of data collectors, Observability API tokens, and notifications. An Administrator can also invite other users but can only assign Observability roles.	Can perform all Security functions, including those for Alerts, Forensics, data collectors, automated response policies, and API tokens for Security. An Administrator can also invite other users but can only assign Security roles.	Can perform all User/Author functions including managing Reporting API tokens, as well as all administrative tasks such as configuration of reports, and the shutdown and restart of reporting tasks. An Administrator can also invite other users but can only assign Reporting roles.	
User	Can view and modify dashboards, queries, alerts, annotations, annotation rules, and applications, and manage device resolution.	Can view and manage Alerts and view Forensics. User role can change alert status, add a note, take snapshots manually, and manage restrict user access.	Can perform all Guest/Consumer functions as well as create and manage reports and dashboards.	
Guest	Has read-only access to asset pages, dashboards, alerts, and can view and run queries.	Can view Alerts and Forensics. Guest role cannot change alert status, add a note, take snapshots manually, or restrict user access.	Can view, schedule, and run reports and set personal preferences such as those for languages and time zones. Guests/Consumers cannot create reports or perform administrative tasks.	

Best practice is to limit the number of users with Administrator permissions. The greatest number of accounts should be user or guest accounts.

Cloud Insights Permissions by User Role

The following table shows the Cloud Insights permissions granted to each user role.

Feature	Administrator/ Account Owner	User	Guest
Acquisition Units: Add/Modify/Delete	Υ	N	N
Alerts*: Create/Modify/Delete	Υ	Υ	N
Alerts*: View	Υ	Υ	Υ
Annotation Rules: Create/Run/Modify/Delete	Υ	Υ	N
Annotations: Create/Modify/Assign/Vie w/Remove/Delete	Υ	Υ	N
API Access*: Create/Rename/Disable/R evoke	Υ	N	N
Applications: Create/View/Modify/Delet e	Υ	Υ	N
Asset Pages: Modify	Υ	Υ	N
Asset Pages: View	Υ	Υ	Υ
Audit: View	Υ	N	N
Cloud Cost	Υ	N	N
Security	Υ	N	N
Dashboards: Create/Modify/Delete	Υ	Υ	N
Dashboards: View	Υ	Υ	Υ
Data Collectors: Add/Modify/Poll/Delete	Υ	N	N
Notifications: View	Υ	Υ	Υ
Notifications: Modify	Υ	N	N
Queries: Create/Modify/Delete	Υ	Υ	N
Queries: View/Run	Υ	Υ	Υ
Device Resolution	Υ	Υ	N
Reports*: View/Run	Υ	Υ	Υ

Reports*: Create/Modify/Delete/Sch edule	Υ	Υ	N
Subscription: View/Modify	Υ	N	N
User Management: Invite/Add/Modify/Deactiv ate	Υ	N	N

^{*}Requires Premium Edition

Creating Accounts by Inviting Users

Creating a new user account is achieved through Cloud Central. A user can respond to the invitation sent through email, but if the user does not have an account with Cloud Central, the user needs to sign up with Cloud Central so that they can accept the invitation.

Before you begin

- The user name is the email address of the invitation.
- · Understand the user roles you will be assigning.
- Passwords are defined by the user during the sign up process.

Steps

- 1. Log into Cloud Insights
- 2. In the menu, click **Admin > User Management**

The User Management screen is displayed. The screen contains a list of all of the accounts on the system.

3. Click + User

The Invite User screen is displayed.

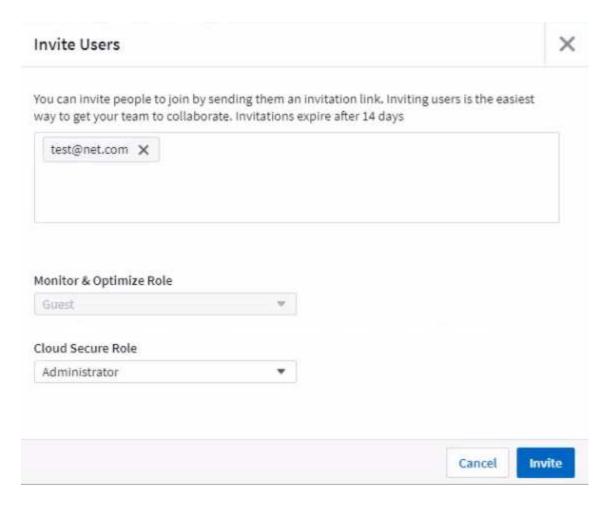
4. Enter an email address or multiple addresses for invitations.

Note: When you enter multiple addresses, they are all created with the same role. You can only set multiple users to the same role.

1. Select the user's role for each feature of Cloud Insights.



The features and roles you can choose from depend on which features you have access to in your particular Administrator role. For example, if you have Admin role only for Reporting, you will be able to assign users to any role in Reporting, but will not be able to assign roles for Observability or Security.



2. Click Invite

The invitation is sent to the user. Users will have 14 days to accept the invitation. Once a user accepts the invitation, they will be taken to the NetApp Cloud Portal, where they will sign up using the email address in the invitation. If they have an existing account for that email address, they can simply sign in and will then be able to access their Cloud Insights environment.

Modifying an existing user's role

To modify an existing user's role, including adding them as a **secondary account owner**, follow these steps.

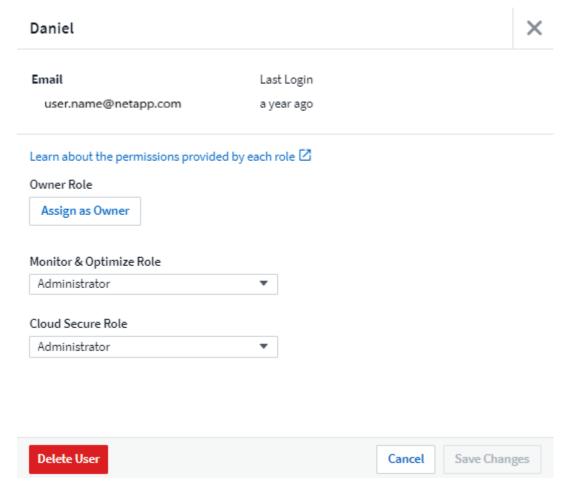
- 1. Click **Admin > User Management**. The screen displays a list of all of the accounts on the system.
- 2. Click the user name of the account you want to change.
- 3. Modify the user's role in each Cloud Insights feature set as needed.
- 4. Click Save Changes.

To assign a Secondary Account Owner

You must be logged in as an account owner for Observability in order to assign the account owner role to another user.

- 1. Click Admin > User Management.
- 2. Click the user name of the account you want to change.

- 3. In the User dialog, click on **Assign as Owner**.
- 4. Save the changes.



You can have as many account owners as you wish, but best practice is to limit the owner role to only select people.

Deleting Users

A user with the Administrator role can delete a user (for example, someone no longer with the company) by clicking on the user's name and clicking *Delete User* in the dialog. The user will be removed from the Cloud Insights environment.

Note that any dashboards, queries, etc. that were created by the user will remain available in the Cloud Insights environment even after the user is removed.

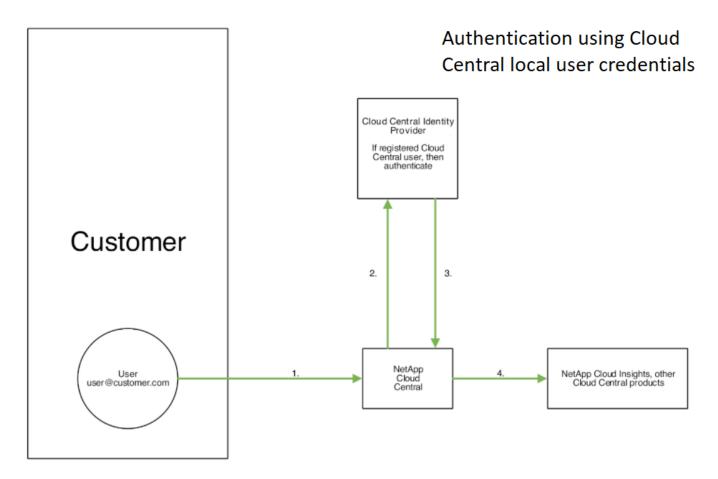
Single Sign-On (SSO) and Identity Federation

Enabling Identity Federation for SSO In Cloud Insights

With Identity Federation:

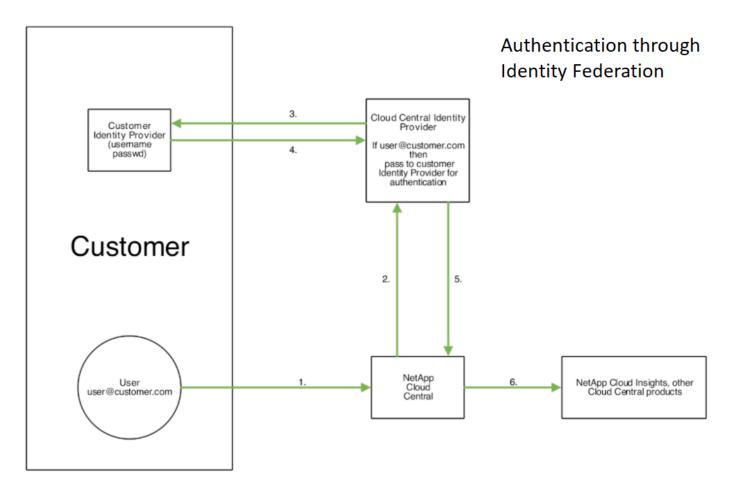
 Authentication is delegated to the customer's identity management system, using the customer's credentials from your corporate directory, and automatization policies such as Multi-Factor Authentication (MFA). • Users log in once to all NetApp Cloud Services (Single Sign On).

User accounts are managed in NetApp Cloud Central for all Cloud Services. By default, authentication is done using Cloud Central local user profile. Below is a simplified overview of that process:



However, some customers would like to use their own identity provider to authenticate their users for Cloud Insights and their other NetApp Cloud Central Services. With Identity Federation, NetApp Cloud Central accounts are authenticated using credentials from your corporate directory.

The following is a simplified example of that process:



In the above diagram, when a user accesses Cloud Insights, that user is directed to the customer's identity management system for authentication. Once the account is authenticated, the user is directed to the Cloud Insights tenant URL.

Cloud Central uses Auth0 to implement Identity Federation and integrate with services like Active Directory Federation Services (ADFS) and Microsoft Azure Active Directory (AD). For more information on Identity Federation setup and configuration, see Cloud Central documentation on Identity Federation.

It is important to understand that changing identity federation in Cloud Central will apply not only to Cloud Insights but to all NetApp Cloud Services. The customer should discuss this change with the NetApp team of each Cloud Central product they own to make sure the configuration they are using will work with Identity Federation or if adjustments need to be made on any accounts. The customer will need to involve their internal SSO team in the change to identity federation as well.

It is also important to realize that once identity federation is enabled, that any changes to the company's identity provider (such moving from SAML to Microsoft AD) will likely require troubleshooting/changes/attention in Cloud Central to update the profiles of the users.

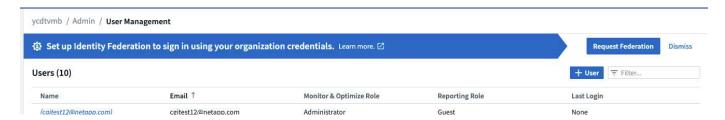
Single Sign-On (SSO) User Auto-Provisioning

In addition to inviting users, administrators can enable **Single Sign-On (SSO) User Auto-Provisioning** access to Cloud Insights for all users in their corporate domain, without having to invite them individually. With SSO enabled, any user with the same domain email address can log into Cloud Insights using their corporate credentials.



SSO User Auto-Provisioning is available in Cloud Insights Premium Edition, and must be configured before it can be enabled for Cloud Insights. SSO User Auto-Provisining configuration includes Identity Federation through NetApp Cloud Central as described in the section above. Federation allows single sign-on users to access your NetApp Cloud Central accounts using credentials from your corporate directory, using open standards such as Security Assertion Markup Language 2.0 (SAML) and OpenID Connect (OIDC).

To configure *SSO User Auto-Provisioning*, on the **Admin > User Management** page, click the **Request Federation** button. Once configured, administrators can then enable SSO user login. When an administrator enables *SSO User Auto-Provisioning*, they choose a default role for all SSO users (such as Guest or User). Users who log in through SSO will have that default role.



Occasionally, an administrator will want to promote a single user out of the default SSO role (for example, to make them an administrator). They can accomplish this on the **Admin > User Management** page by clicking on the right-side menu for the user and selecting *Assign Role*. Users who are assigned an explicit role in this way continue to have access to Cloud Insights even if *SSO User Auto-Provisioning* is subsequently disabled.

If the user no longer requires the elevated role, you can click the menu to *Remove User*. The user will be removed from the list. If *SSO User Auto-Provisioning* is enabled, the user can continue log in to Cloud Insights through SSO, with the default role.

You can choose to hide SSO users by unchecking the Show SSO Users checkbox.

However, do not enable SSO User Auto-Provisioning if either of these are true:

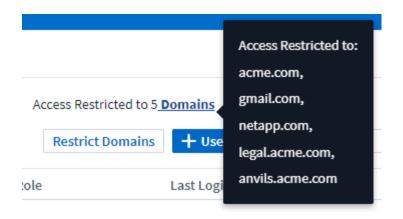
- · Your organization has more than one Cloud Insights tenant
- Your organization does not want any/every user in the federated domain to have some level of automatic access to the Cloud Insights tenant. At this point in time, we do not have the ability to use groups to control role access with this option.

Restricting Access by Domain

Cloud Insights can restrict user access to only the domains you specify. On the **Admin > User Management** page, select "Restrict Domains".

You are presented with these choices:

- No restrictions: Cloud Insights remains accessible to users regardless of their domain.
- Limit access to default domains: default domains are those used by your Cloud Insights environment account owners. These domains are always accessible.
- Limit access to defaults plus domains you specify. List any domains you want to have access to your Cloud Insights environment, in addition to the default domains.



Cloud Insights Data Collector List

Cloud Insights supports a variety of Data Collectors from many vendors and services.

Data Collectors are categorized by these types:

- Infrastructure: Acquired from vendor devices such as storage arrays, switches, hypervisors, or backup devices.
- Service: Acquired from services such as Kubernetes or Docker. Also called *Integration*.

Alphabetical list of Data Collectors supported by Cloud Insights:

Data Collector	Туре
Amazon EC2 and EBS	Infrastructure
AWS S3 as Storage	Infrastrusture
Amazon FSx for NetApp ONTAP	Infrastructure
Apache	Service
Azure NetApp Files	Infrastructure
Azure VMs and VHD	Infrastructure
Brocade Network Advisor (BNA)	Infrastructure
Brocade Fibre Channel Switches	Infrastructure
Brocade FOS REST	Infrastructure
Cisco MDS Fabric Switches	Infrastructure
Consul	Service
Couchbase	Service
CouchDB	Service
Cohesity SmartFiles	Infrastructure
Dell EMC Data Domain	Infrastructure
Dell EMC ECS	Infrastructure
Dell EMC PowerScale (previously Isilon)	Infrastructure
Dell EMC Isilon / PowerScale REST	Infrastructure
Dell EMC PowerStore	Infrastructure
Dell EMC Recoverpoint	Infrastructure
Dell EMC ScaleIO/PowerFlex	Infrastructure
Dell EMC Unity	Infrastructure
Dell EMC Unisphere REST	Infrastructure
Dell EMC VMAX/PowerMax Family of Devices	Infrastructure
Dell EMC VNX Block Storage	Infrastructure

Data Collector	Туре
Dell EMC VNX File	Infrastructure
Dell EMC VNX Unified	Infrastructure
Dell EMC VPLEX	Infrastructure
Dell EMC XtremIO	Infrastructure
Dell XC Series	Infrastructure
Docker	Service
Elasticsearch	Service
Flink	Service
Fujitsu ETERNUS DX	Infrastructure
Google Compute and Storage	Infrastructure
Hadoop	Service
HAProxy	Service
Hitachi Content Platform (HCP)	Infrastructure
Hitachi Vantara Command Suite	Infrastructure
Hitachi Vantara NAS Platform	Infrastructure
Hitachi Ops Center	Infrastructure
HP Enterprise Alletra 6000 (previously Nimble) Storage	Infrastructure
HP Enterprise Alletra 9000 / Primera (previously 3PAR) Storage	Infrastructure
HP Enterprise Command View	Infrastructure
Huawei OceanStor and Dorado Devices	Infrastructure
IBM Cleversafe	Infrastructure
IBM CS Series	Infrastructure
IBM PowerVM	Infrastructure
IBM SAN Volume Controller (SVC)	Infrastructure
IBM System Storage DS8000 Series	Infrastructure
IBM XIV and A9000 Storages	Infrastructure
Infinidat InfiniBox	Infrastructure
Java	Service
Kafka	Service
Kapacitor	Service
Kibana	Service
Kubernetes	Service

Data Collector	Туре
Lenovo HX Series	Infrastructure
Memcached	Service
Microsoft Azure NetApp Files	Infrastructure
Microsoft Hyper-V	Infrastructure
MongoDB	Service
MySQL	Service
NetApp Cloud Volumes ONTAP	Infrastructure
NetApp Cloud Volumes Services for AWS	Infrastructure
NetApp Cloud Connection for ONTAP 9.9+	Infrastructure
NetApp Data ONTAP 7-Mode	Infrastructure
NetApp E-Series	Infrastructure
Amazon FSx for NetApp ONTAP	Infrastructure
NetApp HCI Virtual Center	Infrastructure
NetApp ONTAP Data Management Software	Infrastructure
NetApp ONTAP Select	Infrastructure
NetApp SolidFire All-Flash Array	Infrastructure
NetApp StorageGRID	Infrastructure
Netstat	Service
Nginx	Service
Node	Service
Nutanix NX Series	Infrastructure
OpenStack	Infrastructure
OpenZFS	Service
Oracle ZFS Storage Appliance	Infrastructure
PostgreSQL	Service
Puppet Agent	Service
Pure Storage FlashArray	Infrastructure
Red Hat Virtualization	Infrastructure
Redis	Service
RethinkDB	Service
RHEL & CentOS	Service
Rubrik CDM Storage	Infrastructure
Ubuntu & Debian	Service

Data Collector	Туре
VMware vSphere	Infrastructure
Windows	Service
ZooKeeper	Service

Subscribing to Cloud Insights

Getting started with Cloud Insights is as easy as three simple steps:

- Sign up for an account on NetApp BlueXP to get access to all of NetApp's Cloud offerings.
- Register for a free trial of Cloud Insights to explore the features available.
- Subscribe to Cloud Insights for on-going, uninterrupted access to your data via NetApp Sales direct or AWS Marketplace.

During the registration process, you can choose the global region to host your Cloud Insights environment. For more information, read about Cloud Insights Information and Region.



Unless otherwise noted, the information on this page generally applies to Cloud Insights Commercial Editions. The Federal Edition of Cloud Insights may not contain some of the functionality described on this page.

For a full comparison of the features available in Cloud Insights Basic and Premium Editions, see the Cloud Insights Pricing page.



Inactive Cloud Insights Basic Edition environments are deleted and their resources are reclaimed. An environment is considered inactive if there is no user activity for 30 consecutive days, of if there is no data ingested for 7 consecutive days. Cloud Insights will send a notification and provide a grace period of four days before an environment is deleted.

While using Cloud Insights, if you see a padlock icon it means the feature is not available in your current Edition, or is available in a limited form. Upgrade for full access to the feature.

Trial Edition

When you sign up for Cloud Insights and your environment is active, you enter into a free, 30-day trial of Cloud Insights. During this trial you can explore the features that Cloud Insights has to offer, in your own environment.

At any time during your trial period, you can subscribe to Cloud Insights. Subscribing to Cloud Insights ensures uninterrupted access to your data as well as extended **product support** options.

Cloud Insights displays a banner when your free trial is nearing its end. Within that banner is a *View Subscription* link, which opens the **Admin** → **Subscription** page. Non-Admin users will see the banner but will not be able to go to the Subscription page.



If you need additional time to evaluate Cloud Insights and your trial is set to expire in 4 days or less, you can extend your trial for an additional 30 days. You can extend the trial only once. You cannot extend if your trial has expired.

Trial through AWS Marketplace

You may also sign up for a free trial through the AWS Marketplace. The AWS Marketplace free trial gives you access to Cloud Insights for a trial period of 33 days, and allows up to 499 Managed Units (MUs).

Note: If you configure more than 499 MUs, you will enter "breached" state. While your trial is in breached state, you will lose access to some Cloud Insights functionality until the breach is resolved, either by reducing the

number of MUs configured, or by subscribing to Cloud Insights.

The AWS Marketplace free trial cannot be extended. At any time during your trial, you can downgrade to a Cloud Insights Basic Edition subscription or change to a paid Cloud Insights subscription by visiting the **Admin** → **Subscription** page.

What if My Trial has Expired?

If your free trial has expired and you have not yet subscribed to Cloud Insights, you will have limited functionality until you subscribe.

Module Trials

Coming Soon!



Module Trial is considered Preview functionality and is therefore subject to change.

In addition to your initial trial of Cloud Insights, you may also take advantage of **Module Trials**. For example, if you are already subscribed to Infrastructure Observability but are adding Kubernetes to your environment, you will automatically enter into a 30-day trial of Kubernetes Observability, starting from when you install the NetApp Kubernetes Monitoring Operator. You will only be charged for your Kubernetes Observability managed unit usage at the end of the trial period.



Keep in mind that you will be charged for new managed unit (MU) usage following the trial period, so be sure to plan accordingly. When your module trial is ending, you will be notified if you will need to add more MUs to avoid service interruption.

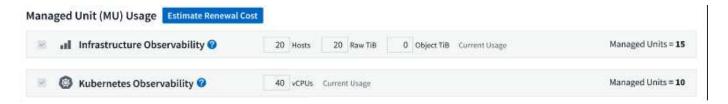
You can monitor your managed unit usage on the **Admin > Subscription** page in the **Usage** tab.



Estimator

During a module trial, you are not changed MU usage for resources consumed for the module, but you can open the **Estimator** (on the *Summary* tab) to see how MUs will be charged following the trial, as well as play with "What if" scenarios with the number of MUs you may need in the future. Reset the numbers by exiting the

Estimator.



Select the checkbox next to a module to add or remove the entire module's MU's from the estimated cost.

The Estimator also allow you to see how the numbers stack up for either an Add On - where you keep your current subscription term and increase the number of managed units licensed - or a Renew option for a the renewal subscription you would purchase when your current subscription term ends.

Note that customers are only eligible for a module trial once per subscription.

Subscription Options

To subscribe, go to **Admin** → **Subscription**. In addition to the **Subscribe** buttons, you will be able to see your installed data collectors and calculate your estimated pricing. For a typical environment, you can click the self-serve AWS Marketplace button. If your environment includes or is expected to include 1,000 or more Managed Units, you are eligible for Volume Pricing.

Pricing

Cloud Insights is priced per **Managed Unit**. Usage of your Managed Units is calculated based on the number of **hosts or virtual machines** and amount of **unformatted capacity** being managed in your infrastructure environment.

- 1 Managed Unit = 2 hosts (any virtual or physical machine)
- 1 Managed Unit = 4 TiB of unformatted capacity of physical or virtual disks
- 1 Managed Unit = 40 TiB of unformatted capacity of select secondary storage: AWS S3, Cohesity SmartFiles, Dell EMC Data Domain, Dell EMC ECS, Hitachi Content Platform, IBM Cleversafe, NetApp StorageGRID, Rubrik.
- 1 Managed Unit = 4 vCPUs of Kuberentes

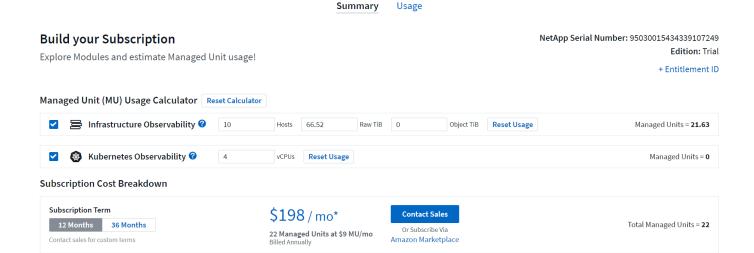
If your environment includes or is expected to include 1,000 or more Managed Units, you are eligible for **Volume Pricing** and will be prompted to Contact NetApp Sales to subscribe. See below for more details.

Estimate Your Subscription Cost

The Subscription Calculators help you estimate your Cloud Insights subscription cost based on the number of Managed Units needed. The current values are pre-populated, and you can adjust those values to assist you with planning for estimated future growth. You can adjust values for Infrastructure, Kubernetes, or both.

Your estimated list price cost will change based on your subscription term.

NOTE: The calculators are for estimation only. Your exact pricing will be set when you subscribe.



How Do I Subscribe?

If your Managed Unit count is less than 1,000, you can subscribe via NetApp Sales, or self-subscribe via AWS Marketplace.

Subscribe through NetApp Sales direct

If your expected Managed Unit count is 1,000 or greater, click on the **Contact Sales** button to subscribe though the NetApp Sales Team.

You must provide your Cloud Insights **Serial Number** to your NetApp sales representative so that your paid subscription can be applied to your Cloud Insights environment. The Serial Number uniquely identifies your Cloud Insights trial environment and can be found on the **Admin > Subscription** page.

Self-Subscribe through AWS Marketplace



You must be an Account Owner or Administrator in order to apply an AWS Marketplace subscription to your existing Cloud Insights trial account. Additionally, you must have an Amazon Web Services (AWS) account.

Clicking on the Amazon Marketplace link opens the AWS Cloud Insights subscription page, where you can complete your subscription. Note that values you entered in the calculator are not populated in the AWS subscription page; you will need to enter the total Managed Units count on this page.

After you have entered the total Managed Units count and chosen either 12-month or 36-month subscription term, click on **Set Up Your Account** to finish the subscription process.

Once the AWS subscription process is complete, you will be taken back to your Cloud Insights environment. Or, if the environment is no longer active (for example, you have logged out), you will be taken to the NetAPp BlueXP sign-in page. When you sign in to Cloud Insights again, your subscription will be active.



After clicking on **Set Up Your account** on the AWS Marketplace page, you must complete the AWS subscription process within one hour. If you do not complete it within one hour, you will need to click on **Set Up Your Account** again to complete the process.

If there is a problem and the subscription process fails to complete correctly, you will still see the "Trial Version"

banner when you log into your environment. In this event, you can go to **Admin > Subscription** and repeat the subscription process.

View Your Subscription Status

Once your subscription is active, you can view your subscription status and Managed Unit usage from the **Admin > Subscription** page.

The Subscription Summary tab displays the following:

- Current Edition
- Subscription Serial Number
- Current MU usage and "what if?" cost estimators
- · Links to modify your subscription
- · Views of your Managed Unit usage

View your Usage Management

The Usage Management tab shows an overview of Managed Unit usage, as well as tabs breaking down Managed Unit consumption by collector or Kubernetes Cluster.



The Unformatted Capacity Managed Unit count reflects a sum of the total raw capacity in the environment and is rounded up to the nearest Managed Unit.



The sum of Managed Units may differ slightly from the Data Collectors count in the summary section. This is because Managed Unit counts are rounded up to the nearest Managed Unit. The sum of these numbers in the Data Collectors list may be slightly higher than the total Managed Units in the status section. The summary section reflects your actual Managed Unit count for your subscription.

In the event that your usage is nearing or exceeding your subscribed amount, you can reduce usage by deleting data collectors or stopping monitoring of Kubernetes Clusters. Delete an item in this list by clicking on the "three dots" menu and selecting *Delete*.

What Happens if I Exceed My Subscribed Usage?

Warnings are displayed when your Managed Unit usage exceeds 80%, 90%, and 100% of your total subscribed amount:

When usage exceeds:	This happens / Recommended action:
80%	An informational banner is displayed. No action is necessary.
90%	A warning banner is displayed. You may want to increase your subscribed Managed Unit count.

100%	An error banner is displayed and you will have limited functionality until you do one of the following: * Remove Data Collectors so that your Managed Unit usage is at or below your subscribed amount * Modify your subscription to increase the subscribed Managed Unit count
------	---

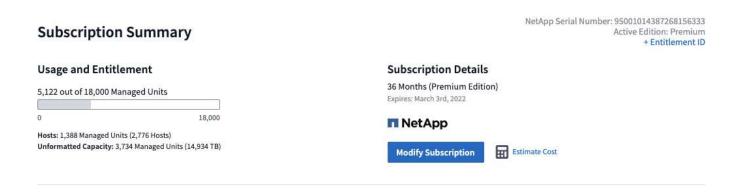
Subscribe Directly and Skip the Trial

You can also subscribe to Cloud Insights directly from the AWS Marketplace, without first creating a trial environment. Once your subscription is complete and your environment is set up, you will immediately be subscribed.

Adding an Entitlement ID

If you own a valid NetApp product that is bundled with Cloud Insights, you can add that product serial number to your existing Cloud Insights subscription. For example, if you have purchased NetApp Astra Control Center, the Astra Control Center license serial number can be used to identify the subscription in Cloud Insights. Cloud Insights refers to this an *Entitlement ID*.

To add an entitlement ID to your Cloud Insights subscription, on the **Admin > Subscription** page, click +*Entitlement ID*.



Automatic Device Resolution

Automatic Device Resolution Overview

You need to identify all of the devices you want to monitor with Cloud Insights. Identification is necessary in order to accurately track performance and inventory in your environment. Typically the majority of devices discovered in your environment are identified through *Automatic Device Resolution*.

After you configure data collectors, devices in your environment including switches, storage arrays, and your virtual infrastructure of hypervisors and VMs are identified. However, this does not normally identify 100% of the devices in your environment.

After data collector type devices have been configured, best practice is to leverage device resolution rules to help identify the remaining unknown devices in your environment. Device resolution can help you resolve unknown devices as the following device types:

- · Physical hosts
- · Storage arrays
- Tapes

Devices remaining as unknown after device resolution are considered generic devices, which you can also show in queries and on dashboards.

The rules created in turn will automatically identify new devices with similar attributes as they are added to your environment. In some cases, device resolution also allows for manual identification bypassing the device resolution rules for undiscovered devices within Cloud Insights.

Incomplete identification of devices can result in issues including:

- Incomplete paths
- · Unidentified multipath connections
- The inability to group applications
- · Inaccurate topology views
- Inaccurate data in the Data warehouse and reporting

The device resolution feature (Manage > Device resolution) includes the following tabs, each of which plays a role in device resolution planning and viewing results:

- Fibre Channel Identify contains a list WWNs and port information of Fibre Channel devices that were not resolved through automatic device resolution. The tab also identifies the percentage of devices that have been identified.
- **IP Address Identify** contains a list of devices accessing CIFS shares and NFS shares that were not identified through automatic device resolution. The tab also identifies the percentage of devices that have been identified.
- Auto resolution rules contains the list of rules that are run when performing Fibre channel device resolution. These are rules you create to resolve unidentified Fibre channel devices.
- Preferences provides configuration options that you use to customize device resolution for your

environment.

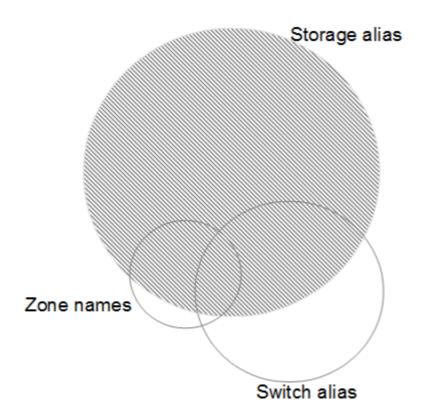
Before You Begin

You need to know how your environment is configured before you define the rules for identifying devices. The more you know about your environment the easier it will be to identify devices.

You need to answer questions similar to the following to help you create accurate rules:

- Does your environment have naming standards for zones or hosts and what percentage of these are accurate?
- Does your environment use a switch alias or storage alias and do they match the host name?
- · How often do naming schemes change in your environment?
- · Have there been any acquisitions or mergers that introduced different naming schemes?

After analyzing your environment, you should be able to identify what naming standards exist that you can expect to reliability encounter. The information you gathered might be represented graphically in a figure similar to the following:



In this example the largest number of devices are reliably represented by storage aliases. Rules that identify hosts using storage aliases should be written first, rules using switch aliases should be written next, and the last rules created should use zone aliases. Due to the overlap of the use of zone aliases and switch aliases, some storage alias rules might identify additional devices, leaving less rules required for zone aliases and switch aliases.

Steps to Identifying devices

Typically, you would use a workflow similar to the following to identify devices in your environment. Identification is an iterative process and might require multiple steps of planning and refining rules.

- Research environment
- Plan rules
- · Create/Revise rules
- · Review results
- · Create additional rules or Manually Identify devices
- Done



If you have unidentified devices (otherwise known as unknown or generic devices) in your environment and you subsequently configure a data source that identifies those devices upon polling, they will no longer be displayed or counted as generic devices.

Related:

Creating Device Resolution Rules Fibre Channel Device Resolution IP Device Resolution Setting Device Resolution Preferences

Device Resolution rules

You create device resolution rules to identify hosts, storage, and tapes that are not automatically identified currently by Cloud Insights. The rules that you create identify devices currently in your environment and also identify similar devices as they are added to your environment.

Creating Device Resolution Rules

When you create rules you start by identifying the source of information that the rule runs against, the method used to extract information, and whether DNS lookup is applied to the results of the rule.

Source that is used to identify the device	* SRM aliases for hosts * Storage alias containing an embedded host or tape name * Switch alias containing an embedded host or tape name * Zone names containing an embedded host name
Method that is used to extract the device name from the source	* As is (extract a name from an SRM) * Delimiters * Regular expressions
DNS lookup	Specifies if you use DNS to verify the host name

You create rules in the Auto Resolution Rules tab. The following steps describe the rule creation process.

Procedure

- 1. Click Manage > Device Resolution
- 2. In the Auto resolution rules tab, click + Host Rule or + Tape Rule.

The **Resolution Rule** screen is displayed.



Click the *View matching criteria* link for help with and examples for creating regular expressions.

3. In the **Type** list select the device you want to identify.

You can select Host or Tape.

4. In the Source list, select the source you want to use to identify the host.

Depending on the source you chose, Cloud Insights displays the following response:

- a. **Zones** lists the zones and WWN that need to be identified by Cloud Insights.
- b. SRM lists the unidentified aliases that need to be identified by Cloud Insights
- c. Storage alias lists storage aliases and WWN that need to be identified by Cloud Insights
- d. Switch alias lists the switch aliases that need to be identified by Cloud Insights
- 5. In the **Method** list select the method you want to employ to identify the host.

Source	Method
SRM	As is, Delimiters, Regular expressions
Storage alias	Delimiters, Regular expressions
Switch alias	Delimiters, Regular expressions
Zones	Delimiters, Regular expressions

Rules using Delimiters require the delimiters and the minimum length of the host name.
 The minimum length of the host name is number of characters that Cloud Insights should use to identify a host. Cloud Insights performs DNS lookups only for host names that are this long or longer.

For rules using Delimiters, the input string is tokenized by the delimiter and a list of host name candidates is created by making several combinations of the adjacent token. The list is then sorted, largest to smallest. For example, for an input sring of *vipsnq03_hba3_emc3_12ep0* the list would result in the following:

- vipsnq03 hba3 emc3 12ep0
- vipsnq03 hba3 emc3
- hba3 emc3 12ep0
- vipsnq03 hba3
- emc3 12ep0
- hba3 emc3
- vipsnq03
- 12ep0

- emc3
- hba3
- Rules using Regular expressions require a regular expression, the format, and cases sensitivity selection.
- 6. Click **Run AR** to run all rules, or click the down-arrow in the button to run the rule you created (and any other rules that have been created since the last full run of AR).

The results of the rule run are displayed in the **FC identify** tab.

Starting an automatic device resolution update

A device resolution update commits manual changes that have been added since the last full automatic device resolution run. Running an update can be used to commit and run only the new manual entries made to the device resolution configuration. No full device resolution run is performed.

Procedure

- 1. Log into the Cloud Insights web UI.
- 2. Click Manage > Device Resolution
- 3. In the **Device Resolution** screen, click the down-arrow in the **Run AR** button.
- Click **Update** to start the update.

Rule-assisted manual identification

This feature is used for special cases where you want to run a specific rule or a list of rules (with or without a one-time reordering) to resolve unknown hosts, storage, and tape devices.

Before you begin

You have a number of devices that have not been identified and you also have multiple rules that successfully identified other devices.



If your source only contains part of a host or device name, use a regular expression rule and format it to add the missing text.

Procedure

- 1. Log into the Cloud Insights web UI.
- 2. Click Manage > Device Resolution
- Click the Fibre Channel Identify tab.

The system displays the devices along with their resolution status.

- 4. Select multiple unidentified devices.
- 5. Click Bulk Actions and select Set host resolution or Set tape resolution.

The system displays the Identify screen which contains a list of all of the rules that successfully identified devices.

6. Change the order of the rules to an order that meets your needs.

The order of the rules are changed in the Identify screen, but are not changed globally.

7. Select the method that that meets your needs.

Cloud Insights executes the host resolution process in the order in which the methods appear, beginning with those at the top.

When rules that apply are encountered, rule names are shown in the rules column and identified as manual.

Related:

Fibre Channel Device Resolution
IP Device Resolution
Setting Device Resolution Preferences

Fibre Channel device resolution

The Fibre Channel Identify screen displays the WWN and WWPN of fibre channel devices whose hosts have not been identified by automatic device resolution. The screen also displays any devices that have been resolved by manual device resolution.

Devices that have been resolved by manual resolution contain a status of *OK* and identify the rule used to identify the device. Missing devices have a status of *Unidentified*. Devices that are specifically excluded from identification have a status of *Excluded*. The total coverage for identification of devices is listed on this page.

You perform bulk actions by selecting multiple devices on the left-hand side of the Fibre Channel Identify screen. Actions can be performed on a single device by hovering over a device and selecting the *Identify* or *Unidentify* buttons on the far right of the list.

The *Total Coverage* link displays a list of the number of devices identified/number of devices available for your configuration:

- · SRM alias
- Storage alias
- · Switch alias
- Zones
- · User defined

Adding a Fibre Channel device manually

You can manually add a fibre channel device to Cloud Insights using the *Manual Add* feature available in the device resolution Fibre Channel Identify tab. This process might be used for pre-identification of a device that is expected to be discovered in the future.

Before you begin

To successfully add a device identification to the system you need to know the WWN or IP address and the device name.

About this task

You can add a Host, Storage, Tape or Unknown fibre channel device manually.

Procedure

1. Log in to the Cloud Insights web UI

- 2. Click Manage > Device Resolution
- 3. Click the Fibre Channel Identify tab.
- 4. Click the Add button.

The Add Device dialog is displayed

5. Enter the WWN or IP address, the device name, and select the device type.

The device you enter is added to the list of devices in the Fibre Channel Identify tab. The Rule is identified as *Manual*.

Importing Fibre Channel device identification from a .CSV file

You can manually import fibre channel device identification into Cloud Insights device resolution using a list of devices in a .CSV file.

1. Before you begin

You must have a correctly formatted .CSV file in order to import device identifications directly into device resolution. The .CSV file for fibre channel devices requires the following information:

WWN	IP	Name	Туре

The data fields must be enclosed in quotes, as shown in the example below.

```
"WWN","IP","Name","Type"
"WWN:2693","ADDRESS2693|IP2693","NAME-2693","HOST"
"WWN:997","ADDRESS997|IP997","NAME-997","HOST"
"WWN:1860","ADDRESS1860|IP1860","NAME-1860","HOST"
```



As a best practice, it is recommended to first export the Fibre Channel Identify information to a .CSV file, make your desired changes in that file, and then import the file back into Fibre Channel Identify. This ensures that the expected columns are present and in the proper order.

To import Fibre Channel Identify information:

- 1. Log into the Cloud Insights web UI.
- 2. Click Manage > Device Resolution
- Select the Fibre Channel Identify tab.
- 4. Click the **Identify > Identify from file** button.
- 5. Navigate to the folder containing your .CSV files for import and select the desired file.

The devices you enter are added to the list of devices in the Fibre Channel Identify tab. The "Rule" is identified as Manual.

Exporting Fibre Channel device identifications to a .CSV file

You can export existing fibre channel device identifications to a .CSV file from the Cloud Insights device resolution feature. You might want to export a device identification so that you can modify it and then import it back into Cloud Insights where it is then used to identify devices that are similar to those originally matching the exported identification.

About this task

This scenario might be used when devices have similar attributes that can be easily edited in the .CSV file and then imported back into the system.

When you export a Fibre Channel device identification to a .CSV file, the file contains the following information in the order shown:

WWN IP Name	Туре
-------------	------

Procedure

- 1. Log into the Cloud Insights web UI.
- 2. Click Manage > Device Resolution
- 3. Select the Fibre Channel Identify tab.
- 4. Select the Fibre Channel device or devices whose identification you want to export.
- 5. Click the **Export** 📆 button.

Select whether to open the .CSV file or save the file.

Related:

IP Device Resolution
Creating Device Resolution Rules
Setting Device Resolution Preferences

IP device resolution

The IP Identify screen displays any iSCSI and CIFS or NFS shares that have been identified by automatic device resolution or by manual device resolution. Unidentified devices are also shown. The screen includes the IP address, Name, Status, iSCSI node, and share name for devices. The percentage of devices that have been successfully identified is also displayed.



Adding IP devices manually

You can manually add an IP device to Cloud Insights using the manual add feature available in the IP Identify screen

Procedure

- 1. Log in to the Cloud insights web UI.
- 2. Click Manage > Device resolution
- 3. Click the IP Address Identify tab.
- 4. Click the Add button.

The Add Device dialog is displayed

5. Enter the address, IP address, and a unique device name.

Result

The device you enter is added to the list of devices in the IP Address Identify tab.

Importing IP device identification from a .CSV file

You can manually import IP device identifications into the Device Resolution feature using a list of device identifications in a .CSV file.

1. Before you begin

You must have a correctly formatted .CSV file in order to import device identifications directly into the Device Resolution feature. The .CSV file for IP devices requires the following information:

Address	IP	Name

The data fields must be enclosed in quotes, as shown in the example below.

```
"Address","IP","Name"
"ADDRESS6447","IP6447","NAME-6447"
"ADDRESS3211","IP3211","NAME-3211"
"ADDRESS593","IP593","NAME-593"
```



As a best practice, it is recommended to first export the IP Address Identify information to a .CSV file, make your desired changes in that file, and then import the file back into IP Address Identify. This ensures that the expected columns are present and in the proper order.

Exporting IP device identification to a .CSV file

You can export existing IP device identifications to a .CSV file from the Cloud Insights device resolution feature. You might want to export a device identification so that you can modify it and then import it back into Cloud Insights where it is then used to identify devices that are similar to those originally matching the exported identification.

About this task

This scenario might be used when devices have similar attributes that can be easily edited in the .CSV file and then imported back into the system.

When you export an IP device identification to a .CSV file, the file contains the following information in the order shown:

Address	IP	Name

Procedure

- 1. Log into the Cloud Insights web UI.
- 2. Click Manage > Device Resolution
- 3. Select the IP Address Identify tab.
- 4. Select the IP device or devices whose identification you want to export.
- 5. Click the **Export** n button.

Select whether to open the .CSV file or save the file.

Related:

Fibre Channel device resolution Creating Device Resolution Rules Setting Device Resolution Preferences

Setting options in the Preferences tab

The device resolution preferences tab lets you create an auto resolution schedule, specify storage and tape venders to include or exclude from identification, and set DNS lookup options.

Auto resolution schedule

An auto resolution schedule can specify when automatic device resolution is run:

Option	Description
Every	Use this option to run automatic device resolution on intervals of days, hours, or minutes.
Every day	Use this option to run automatic device resolution daily at a specific time.
Manually	Use this option to only run automatic device resolution manually.
On every environment change	Use this option to run automatic device resolution whenever there is a change in the environment.

If you specify *Manually*, nightly automatic device resolution is disabled.

DNS processing options

DNS processing options allow you to select the following features:

- When DNS lookup result processing is enabled, you can add a list of DNS names to append to resolved devices.
- You can select Auto resolution of IPs: to enables automatic host resolution for iSCSI initiators and hosts accessing NFS shares by using DNS lookup. If this is not specified, only FC-based resolution is performed.
- You can choose to allow underscores in host names and to use a "connected to" alias instead of the standard port alias in results.

Including or excluding specific storage and tape vendors

You can include or exclude specific storage and tape vendors for automatic resolution. You might want to exclude specific vendors if you know, for example, that a specific host will become a legacy host and should be excluded from your new environment. You can also re-add vendors that you earlier excluded but no longer want excluded.



Device resolution rules for tape only work for WWNs where the Vendor for that WWN is set to *Included as Tape only* in the Vendors preferences.

See also: Regular Expression Examples

Regular expression examples

If you have selected the regular expression approach as your source naming strategy, you can use the regular expression examples as guides for your own expressions used in the Cloud Insights automatic resolution methods.

Formatting regular expressions

When creating regular expressions for Cloud Insights automatic resolution, you can configure output format by entering values in a field named *FORMAT*.

The default setting is \1, which means that a zone name that matches the regular expression is replaced by the contents of the first variable created by the regular expression. In a regular expression, variable values are created by parenthetical statements. If multiple parenthetical statements occur, the variables are referenced numerically, from left to right. The variables can be used in the output format in any order. Constant text can also be inserted in the output, by adding it to the FORMAT field.

For example, you might have the following zone names for this zone naming convention:

```
[Zone number]_[data center]_[hostname]_[device type]_[interface number]
```

- S123 Miami hostname1 filer FC1
- S14_Tampa_hostname2_switch_FC4
- S3991 Boston hostname3 windows2K FC0
- S44 Raleigh hostname4 solaris FC1

And you might want the output to be in the following format:

```
[hostname]-[data center]-[device type]
```

To do this, you need to capture the host name, data center, and device type fields in variables, and use them in the output. The following regular expression would do this:

```
.*?_([a-zA-Z0-9]+)_([a-zA-Z0-9]+)_([a-zA-Z0-9]+)_.*
```

Because there are three sets of parentheses, the variables \1, \2 and \3 would be populated.

You could then use the following format to receive output in your preferred format:

```
\2-\1-\3
```

Your output would be as follows:

```
hostname1-Miami-filer
hostname2-Tampa-switch
hostname3-Boston-windows2K
hostname4-Raleigh-solaris
```

The hyphens between the variables provide an example of constant text that is inserted in the formatted output.

Examples

Example 1 showing zone names

In this example, you use the regular expression to extract a host name from the zone name. You could create a regular expression if you have something similar to the following zone names:

- S0032 myComputer1Name-HBA0
- S0434_myComputer1Name-HBA1
- S0432 myComputer1Name-HBA3

The regular expression that you could use to capture the host name would be:

```
S[0-9]+_([a-zA-Z0-9]*)[_-]HBA[0-9]
```

The outcome is a match of all zones beginning with S that are followed by any combination of digits, followed by an underscore, the alphanumeric hostname (myComputer1Name), an underscore or hyphen, the capital letters HBA, and a single digit (0-9). The hostname alone is stored in the \1 variable.

The regular expression can be broken into its components:

- "S" represents the zone name and begins the expression. This matches only an "S" at the beginning of the zone name.
- The characters [0-9] in brackets indicate that what follows "S" must be a digit between 0 and 9, inclusive.
- The + sign indicates that the occurrence of the information in the preceding brackets has to exist 1 or more times.
- The _ (underscore) means that the digits after S must be followed immediately by only an underscore
 character in the zone name. In this example, the zone naming convention uses the underscore to separate
 the zone name from the host name.
- After the required underscore, the parentheses indicate that the pattern contained within will be stored in the \1 variable.
- The bracketed characters [a-zA-Z0-9] indicate that the characters being matched are all letters (regardless of case) and numbers.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The bracketed characters [_-] (underscore and dash) indicate that the alphanumeric pattern must be followed by an underscore or a dash.
- The letters HBA in the regular expression indicate that this exact sequence of characters must occur in the zone name.
- The final set of bracketed characters [0-9] match a single digit from 0 through 9, inclusive.

Example 2

In this example, skip up to the first underscore "", then match E and everything after that up to the second ", and then skip everything after that.

Zone: Z_E2FHDBS01_E1NETAPP

Hostname: E2FHDBS01

RegExp: .?(E.?).*?

Example 3

The parentheses "()" around the last section in the Regular Expression (below) identifies which part is the hostname. If you wanted VSAN3 to be the host name, it would be: ([a-zA-Z0-9]).*

Zone: A VSAN3 SR48KENT A CX2578 SPA0

Hostname: SR48KENT

RegExp: [a-zA-Z0-9]+ ([a-zA-Z0-9]).*

Example 4 showing a more complicated naming pattern

You could create a regular expression if you have something similar to the following zone names:

- myComputerName123-HBA1 Symm1 FA3
- myComputerName123-HBA2_Symm1_FA5
- myComputerName123-HBA3_Symm1_FA7

The regular expression that you could use to capture these would be:

```
([a-zA-Z0-9]*)_.*
```

The \1 variable would contain only myComputerName123 after being evaluated by this expression.

The regular expression can be broken into its components:

- The parentheses indicate that the pattern contained within will be stored in the \1 variable.
- The bracketed characters [a-zA-Z0-9] mean that any letter (regardless of case) or digit will match.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The _ (underscore) character in the regular expression means that the zone name must have an underscore immediately following the alphanumeric string matched by the preceding brackets.
- The . (period) matches any character (a wildcard).
- The * (asterisk) indicates that the preceding period wildcard may occur 0 or more times.

In other words, the combination .* indicates any character, any number of times.

Example 5 showing zone names without a pattern

You could create a regular expression if you have something similar to the following zone names:

- myComputerName_HBA1_Symm1_FA1
- myComputerName123 HBA1 Symm1 FA1

The regular expression that you could use to capture these would be:

```
(.*?)_.*
```

The \1 variable would contain *myComputerName* (in the first zone name example) or *myComputerName123* (in the second zone name example). This regular expression would thus match everything prior to the first underscore.

The regular expression can be broken into its components:

- The parentheses indicate that the pattern contained within will be stored in the \1 variable.
- The .* (period asterisk) match any character, any number of times.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The ? character makes the match non-greedy. This forces it to stop matching at the first underscore, rather than the last.
- The characters .* match the first underscore found and all characters that follow it.

Example 6 showing computer names with a pattern

You could create a regular expression if you have something similar to the following zone names:

- Storage1 Switch1 myComputerName123A A1 FC1
- Storage2 Switch2 myComputerName123B A2 FC2
- Storage3 Switch3 myComputerName123T A3 FC3

The regular expression that you could use to capture these would be:

Because the zone naming convention has more of a pattern, we could use the above expression, which will match all instances of a hostname (myComputerName in the example) that ends with either an A, a B, or a T, placing that hostname in the \1 variable.

The regular expression can be broken into its components:

- The .* (period asterisk) match any character, any number of times.
- The ? character makes the match non-greedy. This forces it to stop matching at the first underscore, rather than the last.
- The underscore character matches the first underscore in the zone name.
- Thus, the first .*? combination matches the characters Storage1 in the first zone name example.
- The second .*?_ combination behaves like the first, but matches Switch1_ in the first zone name example.
- The parentheses indicate that the pattern contained within will be stored in the \1 variable.
- The bracketed characters [a-zA-Z0-9] mean that any letter (regardless of case) or digit will match.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The bracketed characters in the regular expression [ABT] match a single character in the zone name which must be A, B, or T.
- The _ (underscore) following the parentheses indicates that the [ABT] character match must be followed up an underscore.
- The .* (period asterisk) match any character, any number of times.

The result of this would therefore cause the \1 variable to contain any alphanumeric string which:

- was preceded by some number of alphanumeric characters and two underscores
- was followed by an underscore (and then any number of alphanumeric characters)
- had a final character of A, B or T, prior to the third underscore.

Example 7

Zone: myComputerName123 HBA1 Symm1 FA1

Hostname: myComputerName123

RegExp: ([a-zA-Z0-9]+) .*

Example 8

This example finds everything before the first _.

Zone: MyComputerName_HBA1_Symm1_FA1

MyComputerName123 HBA1 Symm1 FA1

Hostname: MyComputerName

RegExp: (.?)_.

Example 9

This example finds everything after the 1st _ and up to the second _.

Zone: Z MyComputerName StorageName

Hostname: MyComputerName

RegExp: .?(.?).*?

Example 10

This example extracts "MyComputerName123" from the zone examples.

Zone: Storage1_Switch1_MyComputerName123A_A1_FC1

Storage2 Switch2 MyComputerName123B A2 FC2

Storage3_Switch3_MyComputerName123T_A3_FC3

Hostname: MyComputerName123

RegExp: .?.?([a-zA-Z0-9]+)[ABT]_.

Example 11

Zone: Storage1 Switch1 MyComputerName123A A1 FC1

Hostname: MyComputerName123A

RegExp: .?.?([a-zA-z0-9]+).*?

Example 12

The ^ (circumflex or caret) **inside square brackets** negates the expression, for example, [^Ff] means anything except uppercase or lowercase F, and [^a-z] means everything except lowercase a to z, and in the case above, anything except the _. The format statement adds in the "-" to the output host name.

Zone: mhs_apps44_d A 10a0 0429

Hostname: mhs-apps44-d

RegExp: ()_([AB]).*Format in Cloud Insights: $\1-\2$ ([^])_ ()_([^]).*Format in Cloud Insights: $\1-\2-\3$

Example 13

In this example, the storage alias is delimited by "\" and the expression needs to use "\\" to define that there are actually "\" being used in the string, and that those are not part of the expression itself.

Storage Alias: \Hosts\E2DOC01C1\E2DOC01N1

Hostname: E2DOC01N1

RegExp: \\.**?**.?\\(.*?)

Example 14

This example extracts "PD-RV-W-AD-2" from the zone examples.

Zone: PD D-PD-RV-W-AD-2 01

Hostname: PD-RV-W-AD-2

RegExp: -(.*-\d).*

Example 15

The format setting in this case adds the "US-BV-" to the hostname.

Zone: SRV_USBVM11_F1

Hostname: US-BV-M11

RegExp: SRV_USBV([A-Za-z0-9]+)_F[12]

Format: US-BV-\1

Creating Dashboards

Dashboards Overview

Cloud Insights provides users the flexibility to create operational views of infrastructure data, by allowing you to create custom dashboards with a variety of widgets, each of which provides extensive flexibility in displaying and charting your data.



The examples in these sections are for explanation purposes only and do not cover every possible scenario. The concepts and steps herein can be used to create your own dashboards to highlight the data specific to your particular needs.

Creating a Dashboard

You create a new dashboard in one of two places:

- Dashboards > [+New dashboard]
- Dashboards > Show all dashboards > click the [+Dashboard] button

Dashboard Controls

The Dashboard screen has several controls:

- **Time selector**: allows you to view dashboard data for a range of time from the last 15 minutes to the last 30 days, or a custom time range of up to 31 days. You can choose to override this global time range in individual widgets.
- **Edit** button: Selecting this will enable Edit mode, which allows you to make changes to the dashboard. New dashboards open in Edit mode by default.
- Save button: Allows you to save or delete the dashboard.

You can rename the current dashboard by typing a new name before clicking Save.

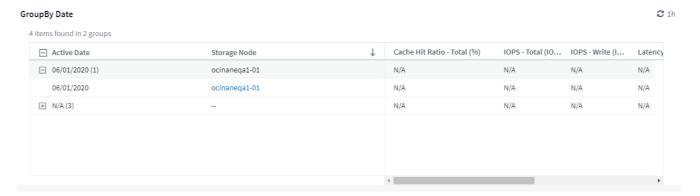
 Add Widget button, which allows you to add any number of tables, charts, or other widgets to the dashboard.

Widgets can be resized and relocated to different positions within the dashboard, to give you the best view of your data according to your current needs.

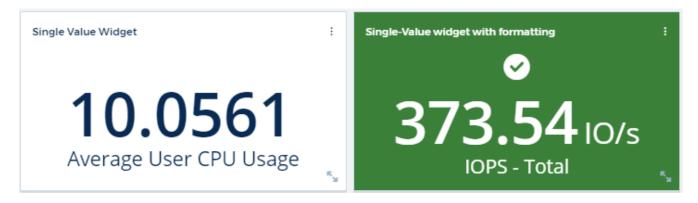
Widget types

You can choose from the following widget types:

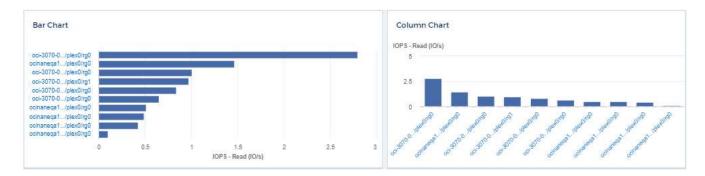
• **Table widget**: A table displaying data according to filters and columns you choose. Table data can be combined in groups that can be collapsed and expanded.



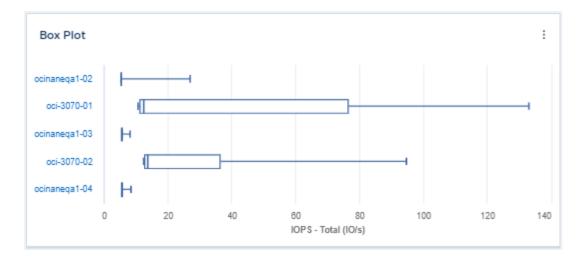
- Line, Spline, Area, Stacked Area Charts: These are time-series chart widgets on which you can display performance and other data over time.
- **Single Value widget**: A widget allowing you to display a single value that can be derived either directly from a counter or calculated using a query or expression. You can define color formatting thresholds to show whether the value is in expected, warning, or critical range.



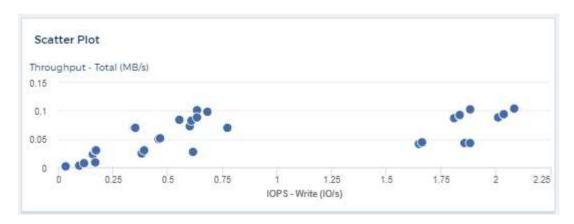
- **Gauge widget**: Displays single-value data in a traditional (solid) gauge or bullet gauge, with colors based on "Warning" or "Critical" values you customize.
- Bar, Column Charts: Displays top or bottom N values, for example, Top 10 storages by capacity or bottom 5 volumes by IOPS.



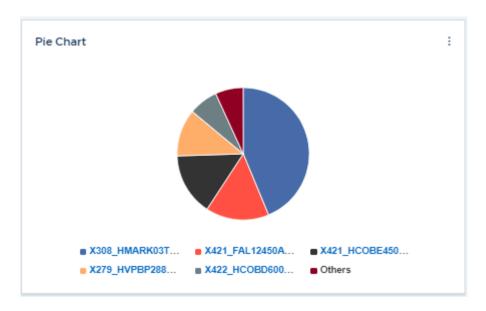
• Box Plot Chart: A plot of the min, max, median, and the range between lower and upper quartile of data in a single chart.



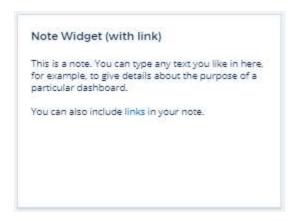
• Scatter Plot Chart: Plots related data as points, for example, IOPS and latency. In this example, you can quickly locate assets with high throughput and low IOPS.



• Pie Chart: a traditional pie chart to display data as a piece of the total.



• Note widget: Up to 1000 characters of free text.



• Alerts Table: Displays up to the last 1,000 alerts.

For more detailed explanations of these and other Dashboard Features, click here.

Setting a Dashboard as your Home Page

You can choose which dashboard to set as your environment's **home page** using either of the following methods:

- Go to **Dashboards** > **Show All Dashboards** to display the list of dashboards in your environment. Click on the options menu to the right of the desired dashboard and select **Set as Home Page**.
- Click on a dashboard from the list to open the dashboard. Click the drop-down menu in the upper corner and select **Set as Home Page**.

Dashboard Features

Dashboards and widgets allow great flexibility in how data is displayed. Here are some concepts to help you get the most from your custom dashboards.

Widget Naming

Widgets are automatically named based on the object, metric, or attribute selected for the first widget query. If you also choose a grouping for the widget, the "Group by" attributes are included in the automatic naming (aggregation method and metric).



Selecting a new object or grouping attribute updates the automatic name.

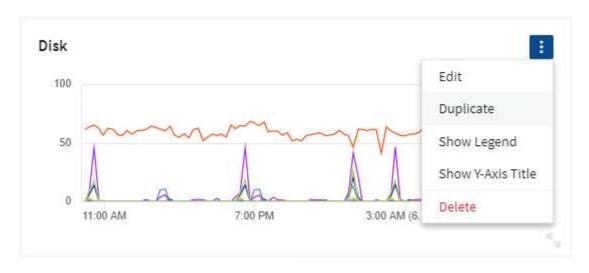
If you do not want to use the automatic widget name, you can simply type a new name.

Widget Placement and Size

All dashboard widgets can be positioned and sized according to your needs for each particular dashboard.

Duplicating a Widget

In dashboard Edit mode, click the menu on the widget and select **Duplicate**. The widget editor is launched, pre-filled with the original widget's configuration and with a "copy" suffix in the widget name. You can easily make any necessary changes and Save the new widget. The widget will be placed at the bottom of your dashboard, and you can position it as needed. Remember to Save your dashboard when all changes are complete.



Displaying Widget Legends

Most widgets on dashboards can be displayed with or without legends. Legends in widgets can be turned on or off on a dashboard by either of the following methods:

 When displaying the dashboard, click the Options button on the widget and select Show Legends in the menu.

As the data displayed in the widget changes, the legend for that widget is updated dynamically.

When legends are displayed, if the landing page of the asset indicated by the legend can be navigated to, the legend will display as a link to that asset page. If the legend displays "all", clicking the link will display a query page corresponding to the first query in the widget.

Transforming Metrics

Cloud Insights provides different **transform** options for certain metrics in widgets (specifically, those metrics called "Custom" or Integration Metrics, such as from Kubernetes, ONTAP Advanced Data, Telegraf plugins, etc.), allowing you to display the data in a number of ways. When adding transformable metrics to a widget, you are presented with a drop-down giving the following transform choices:

None

Data is displayed as is, with no manipulation.

Rate

Current value divided by the time range since the previous observation.

Cumulative

The accumulation of the sum of previous values and the current value.

Delta

The difference between the previous observation value and the current value.

Delta rate

Delta value divided by the time range since the previous observation.

Cumulative Rate

Cumulative value divided by the time range since the previous observation.

Note that transforming metrics does not change the underlying data itself, but only the way that data is displayed.

Dashboard widget queries and filters

Queries

The Query in a dashboard widget is a powerful tool for managing the display of your data. Here are some things to note about widget queries.

Some widgets can have up to five queries. Each query will plot its own set of lines or graphs in the widget. Setting rollup, grouping, top/bottom results, etc. on one query does not affect any other queries for the widget.

You can click on the eye icon to temporarily hide a query. The widget display updates automatically when you hide or show a query. This allows you to check your displayed data for individual queries as you build your widget.

The following widget types can have multiple queries:

- · Area chart
- · Stacked area chart
- · Line chart
- · Spline chart
- · Single value widget

The remaining widget types can have only a single query:

- Table
- · Bar chart
- Box plot
- Scatter plot

Filtering in dashboard widget gueries

Here are some things you can do to get the most out of your filters.

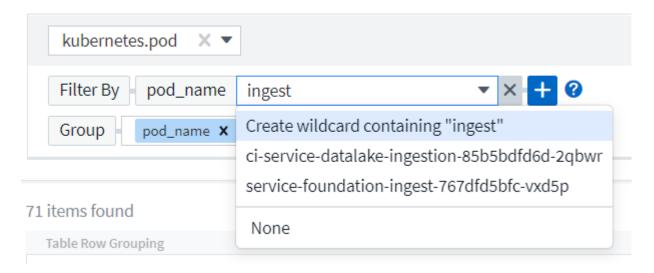
Exact Match Filtering

If you enclose a filter string in double quotes, Insight treats everything between the first and last quote as an exact match. Any special characters or operators inside the quotes will be treated as literals. For example, filtering for "*" will return results that are a literal asterisk; the asterisk will not be treated as a wildcard in this case. The operators AND, OR, and NOT will also be treated as literal strings when enclosed in double quotes.

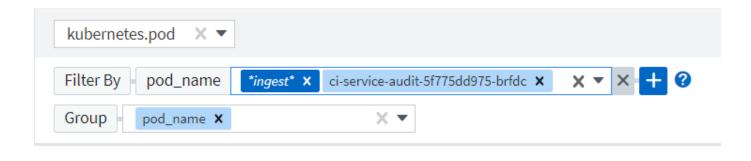
You can use exact match filters to find specific resources, for example hostname. If you want to find only the hostname 'marketing' but exclude 'marketing01', 'marketing-boston', etc., simply enclose the name "marketing" in double quotes.

Wildcards and Expressions

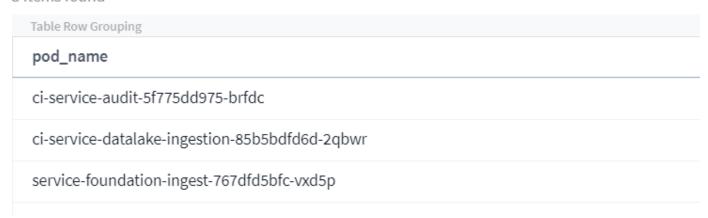
When you are filtering for text or list values in queries or dashboard widgets, as you begin typing you are presented with the option to create a **wildcard filter** based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create **expressions** using NOT or OR, or you can select the "None" option to filter for null values in the field.



Filters based on wildcards or expressions (e.g. NOT, OR, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.



3 items found



Note that Wildcard and Expression filtering works with text or lists but not with numerics, dates or booleans.

Advanced Text Filtering with Contextual Type-Ahead Suggestions

Filtering in widget queries is *contextual*; when you select a filter value or values for a field, the other filters for that query will show values relevant to that filter.

For example, when setting a filter for a specific object *Name*, the field to filter for *Model* will only show values relevant to that object Name.

Contextual filtering also applies to dashboard page variables (text-type attributes or annotations only). When you select a filer value for one variable, any other variables using related objects will only show possible filter values based on the context of those related variables.

Note that only Text filters will show contextual type-ahead suggestions. Date, Enum (list), etc. will not show type-ahead suggestions. That said, you *can* set a filter on an Enum (i.e. list) field and have other text fields be filtered in context. For example, selecting a value in an Enum field like Data Center, then other filters will show only the models/names in that data center), but not vice-versa.

The selected time range will also provide context for the data shown in filters.

Choosing the filter units

As you type a value in a filter field, you can select the units in which to display the values on the chart. For example, you can filter on raw capacity and choose to display in the deafult GiB, or select another format such as TiB. This is useful if you have a number of charts on your dashboard showing values in TiB and you want all your charts to show consistent values.

capacity.raw by Storage A) Query Storage.performance.capacity.raw Filter By capacity.raw 100 X T gibibyte (GiB) - default Storage Гор Group Units Display: Line Chart . d In: bit (b) + Query byte (B) kibibyte (KiB) capacity.raw (TiB) 390.625 mebibyte (MiB) gibibyte (GiB) 341.79688 tahihuta (TiR) 292.96875 None 244.14083 195.3125 146.48438 7:00 PM 7:00 AM (3. Dec) 7:00 PM

Additional Filtering Refinements

The following can be used to further refine your filters.

· An asterisk enables you to search for everything. For example,

```
vol*rhel
```

displays all resources that start with "vol" and end with "rhel".

• The question mark enables you to search for a specific number of characters. For example,

```
BOS-PRD??-S12
```

displays BOS-PRD12-S12, BOS-PRD13-S12, and so on.

• The OR operator enables you to specify multiple entities. For example,

```
FAS2240 OR CX600 OR FAS3270
```

finds multiple storage models.

• The NOT operator allows you to exclude text from the search results. For example,

```
NOT EMC*
```

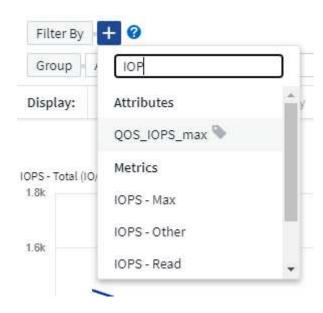
finds everything that does not start with "EMC". You can use

```
NOT *
```

to display fields that contain no value.

Identifying objects returned by queries and filters

The objects returned by queries and filters look similar to those shown in the following illustration. Objects with 'tags' assigned to them are annotations while the objects without tags are performance counters or object attributes.



Grouping and Aggregation

Grouping (Rolling Up)

Data displayed in a widget is grouped (sometimes called rolled-up) from the underlying data points collected during acquisition. For example, if you have a line chart widget showing Storage IOPS over time, you might want to see a separate line for each of your data centers, for a quick comparison. You can choose to group this data in one of several ways:

- Average: displays each line as the average of the underlying data.
- Maximum: displays each line as the maximum of the underlying data.
- Minimum: displays each line as the minimum of the underlying data.
- Sum: displays each line as the sum of the underlying data.
- **Count**: displays a *count* of objects that have reported data within the specified time frame. You can choose the *Entire Time Window* as determined by the dashboard time range (or the widget time range, if set to override the dashboard time), or a *Custom Time Window* that you select.

Steps

To set the grouping method, do the following.

- 1. In your widget's query, choose an asset type and metric (for example, *Storage*) and metric (such as *Performance IOPS Total*).
- 2. For **Group**, choose a roll up method (such as *Average*) and select the attributes or metrics by which to roll up the data (for example, *Data Center*).

The widget updates automatically and shows data for each of your data centers.

You can also choose to group *all* of the underlying data into the chart or table. In this case, you will get a single line for each query in the widget, which will show the average, min, max, sum, or count of the chosen metric or metrics for all of the underlying assets.

Clicking the legend for any widget whose data is grouped by "All" opens a query page showing the results of the first query used in the widget.

If you have set a filter for the query, the data is grouped based on the filtered data.

Note that when you choose to group a widget by any field (for example, *Model*), you will still need to Filter by that field in order to properly display the data for that field on the chart or table.

Aggregating data

You can further align your time-series charts (line, area, etc.) by aggregating data points into minute, hour, or day buckets before that data is subsequently rolled up by attribute (if chosen). You can choose to aggregate data points according to their *Average*, *Maximum*, *Minimum*, *Sum*, or *Count*.

A small interval combined with a long time range may result in an "Aggregation interval resulted in too many data points." warning. You might see this if you have a small interval and increase the dashboard time frame to 7 days. In this case, Insight will temporarily increase the aggregation interval until you select a smaller time frame.

You can also aggregate data in the bar chart widget and single-value widget.

Most asset counters aggregate to *Average* by default. Some counters aggregate to *Max, Min*, or *Sum* by default. For example, port errors aggregate to *Sum* by default, where storage IOPS aggregate to *Average*.

Showing Top/Bottom Results

In a chart widget, you can show either the **Top** or **Bottom** results for rolled up data, and choose the number of results shown from the drop-down list provided. In a table widget, you can sort by any column.

Chart widget top/bottom

In a chart widget, when you choose to rollup data by a specific attribute, you have the option of viewing either the top N or bottom N results. Note that you cannot choose the top or bottom results when you choose to rollup by *all* attributes.

You can choose which results to display by choosing either **Top** or **Bottom** in the query's **Show** field, and selecting a value from the list provided.

Table widget show entries

In a table widget, you can select the number of results shown in the table results. You are not given the option to choose top or bottom results because the table allows you to sort ascending or descending by any column on demand.

You can choose the number of results to show in the table on the dashboard by selecting a value from the query's **Show entries** field.

Grouping in Table Widget

Data in a table widget can be grouped by any available attribute, allowing you to see an overview of your data, and to drill-down into it for more detail. Metrics in the table are rolled up for easy viewing in each collapsed row.

Table widgets allow you to group your data based on the attributes you set. For example, you might want your table to show total storage IOPS grouped by the data centers in which those storages live. Or you might want to display a table of virtual machines grouped according to the hypervisor that hosts them. From the list, you can expand each group to view the assets in that group.

Grouping is only available in the Table widget type.

Grouping example (with rollup explained)

Table widgets allow you to group data for easier display.

In this example, we will create a table widget showing all VMs grouped by Data Center.

Steps

- 1. Create or open a dashboard, and add a **Table** widget.
- 2. Select *Virtual Machine* as the asset type for this widget.
- Click on the Column Selector and choose Hypervisor name and IOPS Total.

Those columns are now displayed in the table.

4. Let's disregard any VM's with no IOPS, and include only VMs that have total IOPS greater than 1. Click the **Filter by [+]** button and select *IOPS - Total*. Click on *Any*, and in the **from** field, type **1**. Leave the **to** field empty. Hit Enter ot click off the filter field to apply the filter.

The table now shows all VMs with Total IOPS greater than or equal to 1. Notice that there is no grouping in the table. All VMs are shown.

5. Click the **Group by [+]** button.

You can group by any attribute or annotation shown. Choose All to display all VMs in a single group.

Any column header for a performance metric displays a "three dot" menu containing a **Roll up** option. The default roll up method is *Average*. This means that the number shown for the group is the average of all the Total IOPS reported for each VM inside the group. You can choose to roll this column up by *Average*, *Sum*, *Minimum* or *Maximum*. Any column that you display that contains performance metrics can be rolled up individually.



6. Click All and select Hypervisor name.

The VM list is now grouped by Hypervisor. You can expand each hypervisor to view the VMs hosted by it.

- 7. Click Save to save the table to the dashboard. You can resize or move the widget as desired.
- 8. Click **Save** to save the dashboard.

Performance data roll up

If you include a column for performance data (for example, *IOPS - Total*) in a table widget, when you choose to group the data you can then choose a roll up method for that column. The default roll up method is to display the average (*avg*) of the underlying data in the group row. You can also choose to display the sum, minimum, or maximum of the data.

Dashboard time range selector

You can select the time range for your dashboard data. Only data relevant to the selected time range will be displayed in widgets on the dashboard. You can select from the following time ranges:

- · Last 15 Minutes
- · Last 30 Minutes
- · Last 60 Minutes
- Last 2 Hours
- · Last 3 Hours (this is the default)
- Last 6 Hours
- · Last 12 Hours
- · Last 24 Hours
- Last 2 Days
- · Last 3 Days
- · Last 7 Days

- Last 30 Days
- · Custom time range

The Custom time range allows you to select up to 31 consecutive days. You can also set the Start Time and End Time of day for this range. The default Start Time is 12:00 AM on the first day selected and the default End Time is 11:59 PM on the last day selected. Clicking **Apply** will apply the custom time range to the dashboard.

Overriding Dashboard Time in Individual widgets

You can override the main dashboard time range setting in individual widgets. These widgets will display data based on their set time frame, not the dashboard time frame.

To override the dashboard time and force a widget to use its own time frame, in the widget's edit mode set the **Override dashboard time** to **On** (check the box), and select a time range for the widget. **Save** the widget to the dashboard.

The widget will display its data according to the time frame set for it, regardless of the time frame you select on the dashboard itself.

The time frame you set for one widget will not affect any other widgets on the dashboard.

Primary and Secondary Axis

Different metrics use different units of measurements for the data they report in a chart. For example, when looking at IOPS, the unit of measurement is the number of I/O operations per second of time (IO/s), while Latency is purely a measure of time (milliseconds, microseconds, seconds, etc.). When charting both metrics on a single line chart using a single set a values for the Y-Axis, the latency numbers (typically a handful of milliseconds) are charted on the same scale with the IOPS (typically numbering in the thousands), and the latency line gets lost at that scale.

But it is possible to chart both sets of data on a single meaningful graph, by setting one unit of measurement on the primary (left-side) Y-axis, and the other unit of measurement on the secondary (right-side) Y-axis. Each metric is charted at its own scale.

Steps

This example illustrates the concept of Primary and Secondary axes in a chart widget.

- 1. Create or open a dashboard. Add a line chart, spline chart, area chart or stacked area chart widget to the dashboard.
- 2. Select an asset type (for example *Storage*) and choose *IOPS Total* for your first metric. Set any filters you like, and choose a roll-up method if desired.

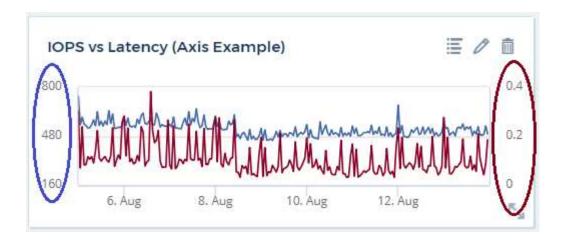
The IOPS line is displayed on the chart, with its scale shown on the left.

3. Click [+Query] to add a second line to the chart. For this line, choose Latency - Total for the metric.

Notice that the line is displayed flat at the bottom of the chart. This is because it is being drawn at the same scale as the IOPS line.

4. In the Latency query, select **Y-Axis: Secondary**.

The Latency line is now drawn at its own scale, which is displayed on the right side of the chart.



Expressions in widgets

In a dashboard, any time series widget (line, spline, area, stacked area) bar chart, column chart, pie chart, or table widget allows you to build expressions from metrics you choose, and show the result of those expressions in a single graph (or column in the case of the table widget). The following examples use expressions to solve specific problems. In the first example, we want to show Read IOPS as a percentage of Total IOPS for all storage assets in our environment. The second example gives visibility into the "system" or "overhead" IOPS that occur in your environment—those IOPS that are not directly from reading or writing data.

You can use variables in expressions (for example, \$Var1 * 100)

Expressions Example: Read IOPS percentage

In this example, we want to show Read IOPS as a percentage of Total IOPS. You can think of this as the following formula:

```
Read Percentage = (Read IOPS / Total IOPS) x 100
```

This data can be shown in a line graph on your dashboard. To do this, follow these steps:

Steps

- 1. Create a new dashboard, or open an existing dashboard in edit mode.
- 2. Add a widget to the dashboard. Choose Area chart.

The widget opens in edit mode. By default, a query is displayed showing *IOPS - Total* for *Storage* assets. If desired, select a different asset type.

3. Click the Convert to Expression link on the right.

The current query is converted to Expression mode. Notice that you cannot change the asset type while in Expression mode. While you are in Expression mode, the link changes to **Revert to Query**. Click this if you wish to switch back to Query mode at any time. Be aware that switching between modes will reset fields to their defaults.

For now, stay in Expression mode.

4. The **IOPS - Total** metric is now in the alphabetic variable field "a". In the "b" variable field, click **Select** and choose **IOPS - Read**.

You can add up to a total of five alphabetic variables for your expression by clicking the + button following the variable fields. For our Read Percentage example, we only need Total IOPS ("a") and Read IOPS ("b").

5. In the **Expression** field, you use the letters corresponding to each variable to build your expression. We know that Read Percentage = (Read IOPS / Total IOPS) x 100, so we would write this expression as:

```
(b / a) * 100
```

- 6. The **Label** field identifies the expression. Change the label to "Read Percentage", or something equally meaningful for you.
- 7. Change the **Units** field to "%" or "Percent".

The chart displays the IOPS Read percentage over time for the chosen storage devices. If desired, you can set a filter, or choose a different rollup method. Be aware that if you select Sum as the rollup method, all percentage values are added together, which potentially may go higher than 100%.

8. Click **Save** to save the chart to your dashboard.

Expressions example: "System" I/O

Example 2: Among the metrics collected from data sources are read, write, and total IOPS. However, the total number of IOPS reported by a data source sometimes includes "system" IOPS, which are those IO operations that are not a direct part of data reading or writing. This system I/O can also be thought of as "overhead" I/O, necessary for proper system operation but not directly related to data operations.

To show these system I/Os, you can subtract read and write IOPS from the total IOPS reported from acquisition. The formula might look like this:

```
System IOPS = Total IOPS - (Read IOPS + Write IOPS)
```

This data can then be shown in a line graph on your dashboard. To do this, follow these steps:

Steps

- 1. Create a new dashboard, or open an existing dashboard in edit mode.
- 2. Add a widget to the dashboard. Choose Line chart.

The widget opens in edit mode. By default, a query is displayed showing *IOPS - Total* for *Storage* assets. If desired, select a different asset type.

3. In the Roll Up field, choose Sum by All.

The Chart displays a line showing the sum of total IOPS.

4. Click the *Duplicate this Query* icon to create a copy of the query.

A duplicate of the query is added below the original.

5. In the second query, click the **Convert to Expression** button.

The current query is converted to Expression mode. Click Revert to Query if you wish to switch back to

Query mode at any time. Be aware that switching between modes will reset fields to their defaults.

For now, stay in Expression mode.

- 6. The IOPS Total metric is now in the alphabetic variable field "a". Click on IOPS Total and change it to IOPS Read.
- 7. In the "b" variable field, click **Select** and choose *IOPS Write*.
- 8. In the **Expression** field, you use the letters corresponding to each variable to build your expression. We would write our expression simply as:

```
a + b
```

In the Display section, choose Area chart for this expression.

The Label field identifies the expression. Change the label to "System IOPS", or something equally meaningful for you.

The chart displays the total IOPS as a line chart, with an area chart showing the combination of read and write IOPS below that. The gap between the two shows the IOPS that are not directly related to data read or write operations. These are your "system" IOPS.

10. Click **Save** to save the chart to your dashboard.

To use a variable in an expression, simply type the variable name, for example, \$var1 * 100. Only numeric variables can be used in expressions.

Expressions in a Table Widget

Table widgets handle expressions a little differently. You can have up to five expressions in a single table widget, each of which is added as a new column to the table. Each expression can include up to five values on which to perform its calculation. You can easily name the column something meaningful.



Variables

Variables allow you to change the data displayed in some or all widgets on a dashboard at once. By setting one or more widgets to use a common variable, changes made in one place cause the data displayed in each widget to update automatically.

Dashboard variables come in several types, can be used across different fields, and must follow rules for naming. These concepts are explained here.

Variable types

A variable can be one the following types:

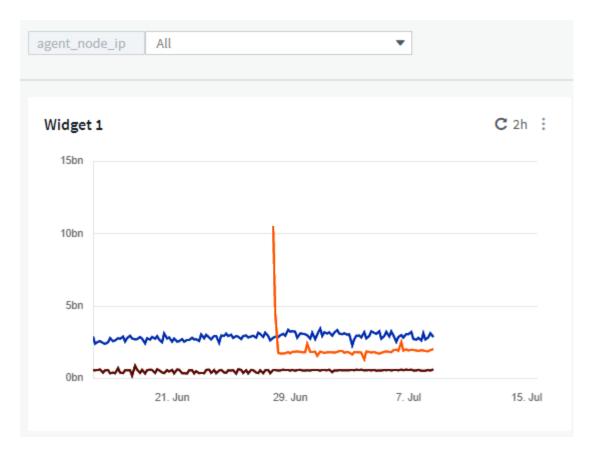
Attribute: Use an object's attributes or metrics to filter

- Annotation: Use a pre-defined Annotation to filter widget data.
- Text: An alphanumeric string.
- Numerical: A number value. Use by itself, or as a "from" or "to" value, depending on your widget field.
- **Boolean**: Use for fields with values of True/False, Yes/No, etc. For the boolean variable, the choices are Yes, No, None, Any.
- Date: A date value. Use as a "from" or "to" value, depending on your widget's configuration.

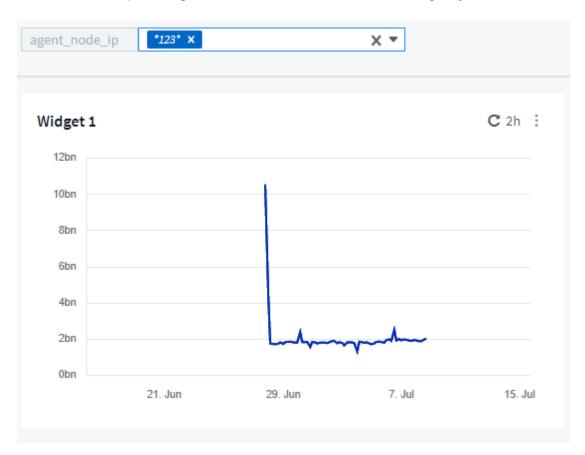


Attribute variables

Selecting an Attribute type variable allows you to filter for widget data containing the specified attribute value or values. The example below shows a line widget displaying free memory trends for Agent nodes. We have created a variable for Agent Node IPs, currently set to show all IPs:

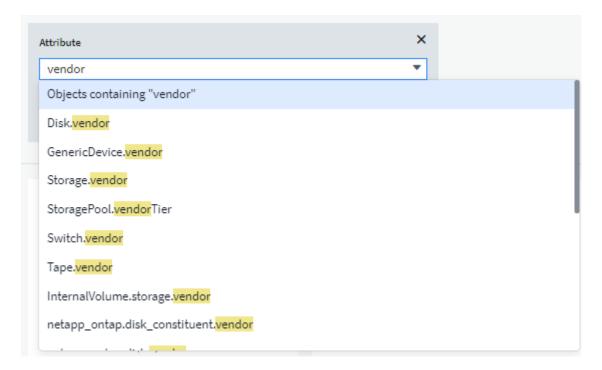


But if you temporarily want to see only nodes on individual subnets in your environment, you can set or change the variable to a specific Agent Node IP or IPs. Here we are viewing only the nodes on the "123" subnet:

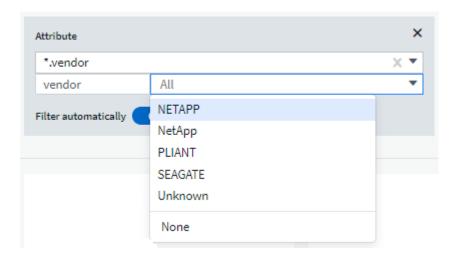


You can also set a variable to filter on all objects with a particular attribute regardless of object type, for

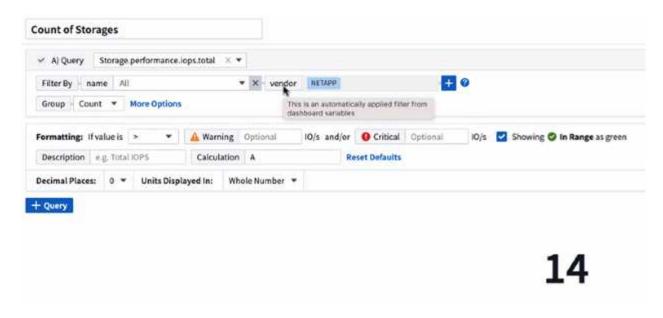
example objects with an attribute of "vendor", by specifying *.vendor in the variable field. You do not need to type the "*."; Cloud Insights will supply this if you select the wildcard option.



When you drop-down the list of choices for the variable value, the results are filtered so show only the available vendors based on the objects on your dashboard.



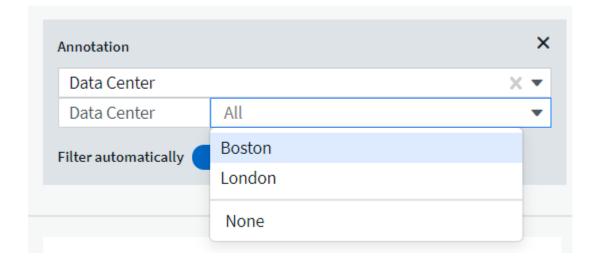
If you edit a widget on your dashboard where the attribute filter is relevant (meaning, the widget's objects contain any *.vendor attribute), it shows you that the attribute filter is automatically applied.



Applying variables is as easy as changing the attribute data of your choice.

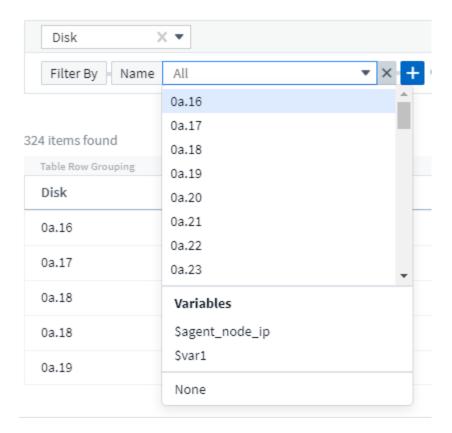
Annotation variables

Choosing an Annotation variable allows you to filter for objects associated with that annotation, for example, those belonging to the same Data Center.



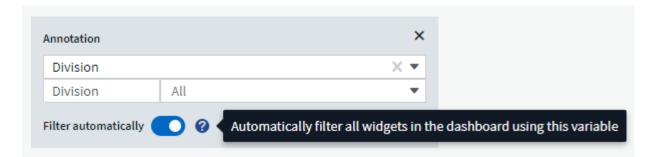
Text, Number, Date, or Boolean variable

You can create generic variables that are not associated with a particular attribute by selecting a variable type of *Text*, *Number*, *Boolean*, or *Date*. Once the variable has been created, you can select it in a widget filter field. When setting a filter in a widget, in addition to specific values that you can select for the filter, any variables that have been created for the dashboard are displayed in the list—these are grouped under the "Variables" section in the drop-down and have names starting with "\$". Choosing a variable in this filter will allow you to search for values that you enter in the variable field on the dashboard itself. Any widgets using that variable in a filter will be updated dynamically.



Variable Filter Scope

When you add an Annotation or Attribute variable to your dashboard, the variable can be applied to *all* widgets on the dashboard, meaning that all widgets on your dashboard will display results filtered according to the value you set in the variable.



Note that only Attribute and Annotation variables can be filtered automatically like this. Non-Annotation or -Attribute variables cannot be automatically filtered. Individual widgets must each be configured to use variables of these types.

To disable automatic filtering so that the variable only applies to the widgets where you have specifically set it, click the "Filter automatically" slider to disable it.

To set a variable in an individual widget, open the widget in edit mode and select the specific annotation or attribute in the *Filter By* field. With an Annotation variable, you can select one or more specific values, or select the Variable name (indicated by the leading "\$") to allow typing in the variable at the dashboard level. The same applies to Attribute variables. Only those widgets for which you set the variable will show the filtered results.

Filtering in variables is *contextual*; when you select a filter value or values for a variable, the other variables on

your page will show only values relevant to that filter.

For example, when setting a variable filter to a specific storage *Model*, any variables set to filter for storage *Name* will only show values relevant to that Model.

To use a variable in an expression, simply type the variable name as part of the expression, for example, *\$var1*100*. Only Numeric variables can be used in expressions. You cannot use numeric Annotation or Attribute variables in expressions.

Filtering in variables is *contextual*; when you select a filter value or values for a variable, the other variables on your page will show only values relevant to that filter.

For example, when setting a variable filter to a specific storage *Model*, any variables set to filter for storage *Name* will only show values relevant to that Model.

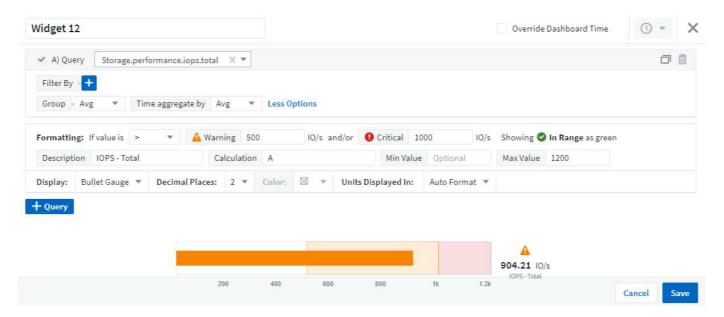
Variable naming

Variables names:

- Must include only the letters a-z, the digits 0-9, period (.), underscore (), and space ().
- Cannot be longer than 20 characters.
- Are case-sensitive: \$CityName and \$cityname are different variables.
- Cannot be the same as an existing variable name.
- · Cannot be empty.

Formatting Gauge Widgets

The Solid and Bullet Gauge widgets allow you to set thresholds for *Warning* and/or *Critical* levels, providing clear representation of the data you specify.



To set formatting for these widgets, follow these steps:

- 1. Choose whether you want to highlight values greater than (>) or less than (<) your thresholds. In this example, we will highlight values greater than (>) the threshold levels.
- 2. Choose a value for the "Warning" threshold. When the widget displays values greater than this level, it displays the gauge in orange.

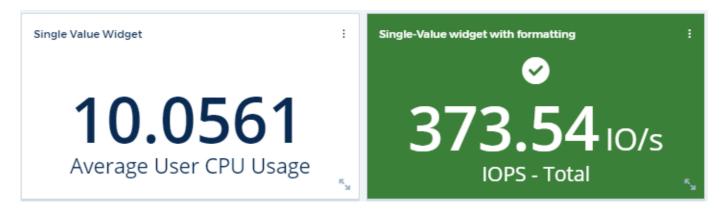
3. Choose a value for the "Critical" threshold. Values greater than this level will cause the gauge to display in red.

You can optionally choose a minimum and maximum value for the gauge. Values below minimum will not display the gauge. Values above maximum will display a full gauge. If you do not choose minimum or maximum values, the widget selects optimal min and max based on the widget's value.



Formatting Single-Value Widget

in the Single-Value widget, in addition to setting Warning (orange) and Critical (red) thresholds, you can choose to have "In Range" values (those below Warning level) shown with either green or white background.



Clicking the link in either a single-value widget or a gauge widget will display a query page corresponding to the first query in the widget.

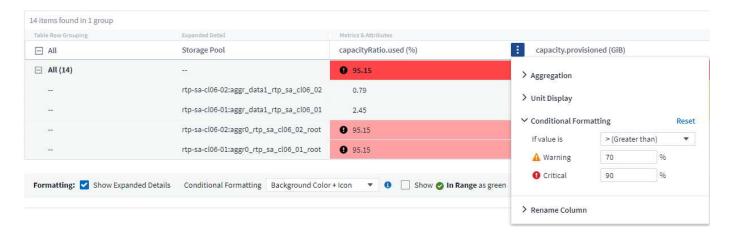
Formatting Table Widgets

Like single-value and gauge widgets, you can set conditional formatting in table widgets, allowing you to highlight data with colors and/or special icons.



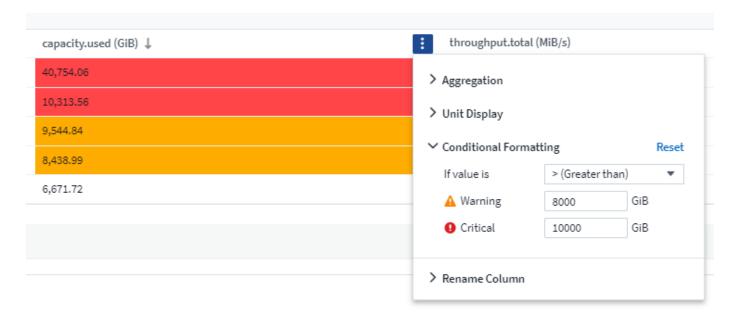
Conditional Formatting is not currently available in Cloud Insights Federal Edition.

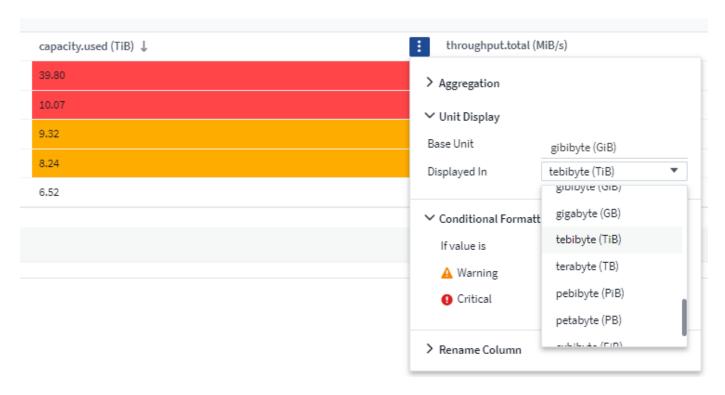
Conditional Formatting allows you to set and highlight Warning-level and Critical-level thresholds in table widgets, bringing instant visibility to outliers and exceptional data points.



Conditional formatting is set separately for each column in a table. For example, you can choose one set of thresholds for a capacity column, and another set for a throughput column.

If you change the Unit Display for a column, the conditional formatting remains and reflects the change in values. The images below show the same conditional formatting even though the display unit is different.



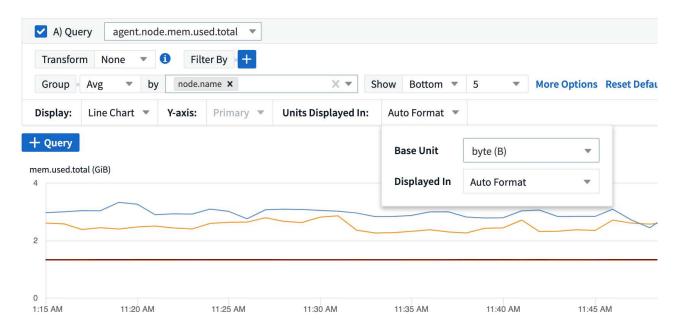


You can choose whether to display condition formatting as color, icons, or both.

Choosing the Unit for Displaying Data

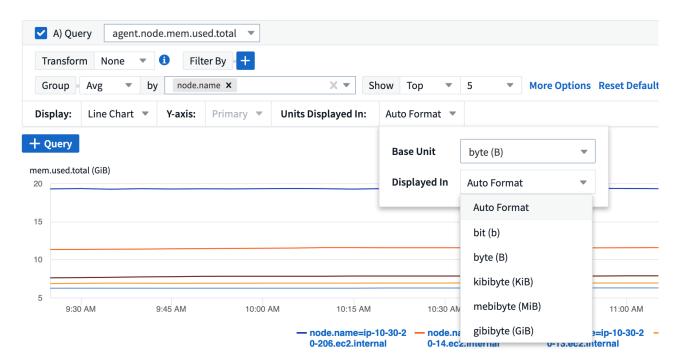
Most widgets on a dashboard allow you to specify the Units in which to display values, for example *Megabytes*, *Thousands*, *Percentage*, *Milliseconds* (*ms*), etc. In many cases, Cloud Insights knows the best format for the data being acquired. In cases where the best format is not known, you can set the format you want.

In the line chart example below, the data selected for the widget is known to be in *bytes* (the base IEC Data unit: see the table below), so the Base Unit is automatically selected as 'byte (B)'. However, the data values are large enough to be presented as gibibytes (GiB), so Cloud Insights by default auto-formats the values as GiB. The Y-axis on the graph shows 'GiB' as the display unit, and all values are displayed in terms of that unit.

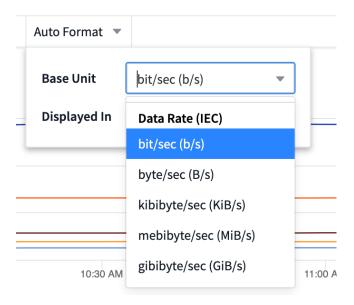


If you want to display the graph in a different unit, you can choose another format in which to display the

values. Since the base unit in this example is *byte*, you can choose from among the supported "byte-based" formats: bit (b), byte (B), kibibyte (KiB), mebibyte (MiB), gibibyte (GiB). The Y-Axis label and values change according to the format you choose.



In cases where the base unit is not known, you can assign a unit from among the available units, or type in your own. Once you assign a base unit, you can then select to display the data in one of the appropriate supported formats.



To clear out your settings and start again, click on Reset Defaults.

A word about Auto-Format

Most metrics are reported by data collectors in the smallest unit, for example as a whole number such as 1,234,567,890 bytes. By default, Cloud Insights will automatically format the value for the most readable display. For example a data value of 1,234,567,890 bytes would be auto formatted to 1.23 *Gibibytes*. You can choose to display it in another format, such as *Mebibytes*. The value will display accordingly.



Cloud Insights uses American English number naming standards. American "billion" is equivalent to "thousand million".

Widgets with multiple queries

If you have a time-series widget (i.e. line, spline, area, stacked area) that has two queries where both are plotted the primary Y-Axis, the base unit is not shown at the top of the Y-Axis. However, if your widget has a query on the primary Y-Axis and a query on the secondary Y-Axis, the base units for each are shown.



If your widget has three or more queries, base units are not shown on the Y-Axis.

Available Units

The following table shows all the available units by category.

Category	Units
Currency	cent dollar
Data(IEC)	bit byte kibibyte mebibyte gibibyte tebibyte pebibyte exbibyte
DataRate(IEC)	bit/sec byte/sec kibibyte/sec mebibyte/sec gibibyte/sec tebibyte/sec pebibyte/sec
Data(Metric)	kilobyte megabyte gigabyte terabyte petabyte exabyte

DataRate(Metric)	kilobyte/sec megabyte/sec gigabyte/sec terabyte/sec petabyte/sec exabyte/sec
IEC	kibi mebi gibi tebi pebi exbi
Decimal	whole number thousand million bilion trillion
Percentage	percentage
Time	nanosecond microsecond millisecond second minute hour
Temperature	celsius fahrenheit
Frequency	hertz kilohertz megahertz gigahertz
CPU	nanocores microcores millicores cores kilocores megacores gigacores teracores petacores exacores
Throughput	I/O ops/sec ops/sec requests/sec reads/sec writes/sec ops/min reads/min writes/min

TV Mode and Auto-Refresh

Data in widgets on Dashboards and Asset Landing Pages auto-refresh according a refresh interval determined by the Dashboard Time Range selected (or widget time range, if set to override the dashboard time). The refresh interval is based on whether the widget is time-series (line, spline, area, stacked area chart) or non-time-series (all other charts).

Dashboard Time Range	Time-Series Refresh Interval	Non-Time-Series Refresh Interval
Last 15 Minutes	10 Seconds	1 Minute
Last 30 Minutes	15 Seconds	1 Minute
Last 60 Minutes	15 Seconds	1 Minute
Last 2 Hours	30 Seconds	5 Minutes
Last 3 Hours	30 Seconds	5 Minutes
Last 6 Hours	1 Minute	5 Minutes
Last 12 Hours	5 Minutes	10 Minutes
Last 24 Hours	5 Minutes	10 Minutes
Last 2 Days	10 Minutes	10 Minutes
Last 3 Days	15 Minutes	15 Minutes
Last 7 Days	1 Hour	1 Hour
Last 30 Days	2 Hours	2 Hours

Each widget displays its auto-refresh interval in the upper-right corner of the widget.

Auto-refresh is not available for Custom dashboard time range.

When combined with **TV Mode**, auto-refresh allows for near-real-time display of data on a dashboard or asset page. TV Mode provides an uncluttered display; the navigation menu is hidden, providing more screen real estate for your data display, as is the Edit button. TV Mode ignores typical Cloud Insights timeouts, leaving the display live until logged out manually or automatically by authorization security protocols.



Because NetApp Cloud Central has its own user login timeout of 7 days, Cloud Insights must log out with that event as well. You can simply log in again and your dashboard will continue to display.

To activate TV Mode, click the
 TV Mode button.

To disable TV Mode, click the **Exit** button in the upper left of the screen.

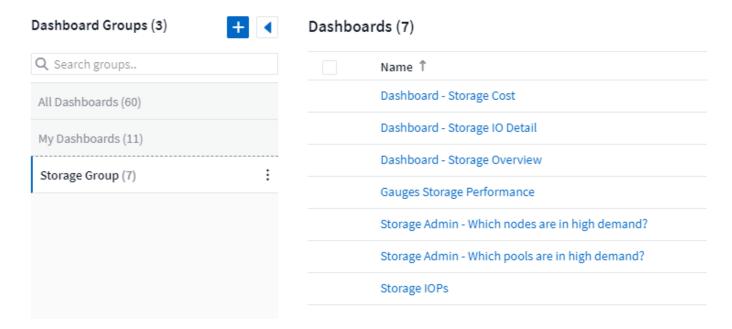


You can temporarily suspend auto-refresh by clicking the Pause button in the upper right corner. While paused, the dashboard time range field will display the paused data's active time range. Your data is still being acquired and updated while auto-refresh is paused. Click the Resume button to continue auto-refreshing of data.



Dashboard Groups

Grouping allows you to view and manage related dashboards. For example, you can have a dashboard group dedicated to the storage in your environment. Dashboard groups are managed on the **Dashboards > Show All Dashboards** page.



Two groups are shown by default:

- All Dashboards lists all the dashboards that have been created, regardless of owner.
- My Dashboards lists only those dashboards created by the current user.

The number of dashboards contained in each group is shown next to the group name.

To create a new group, click the "+" Create New Dashboard Group button. Enter a name for the group and click Create Group. An empty group is created with that name.

To add dashboards to the group, click the *All Dashboards* group to show all dashboards in your environment, of click *My Dashboards* if you only want to see the dashboards you own, and do one of the following:

- To add a single dashboard, click the menu to the right of the dashboard and select Add to Group.
- To add multiple dashboards to a group, select them by clicking the checkbox next to each dashboard, then click the **Bulk Actions** button and select *Add to Group*.

Remove dashboards from the current group in the same manner by selecting *Remove From Group*. You can not remove dashboards from the *All Dashboards* or *My Dashboards* group.



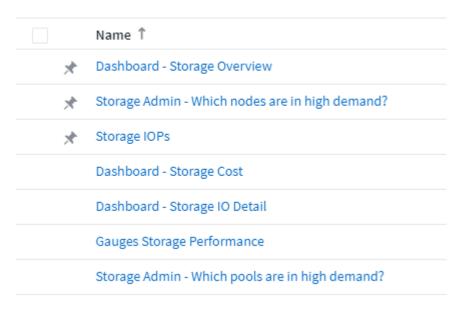
Removing a dashboard from a group does not delete the dashboard from Cloud Insights. To completely remove a dashboard, select the dashboard and click *Delete*. This removes it from any groups to which it belonged and it is no longer available to any user.

Pin your Favorite Dashboards

You can further manage your dashboards by pinning favorite ones to the top of your dashboard list. To pin a dashboard, simply click the thumbtack button displayed when you hover over a dashboard in any list.

Dashboard pin/unpin is an individual user preference and independent of the group (or groups) to which the dashboard belongs.

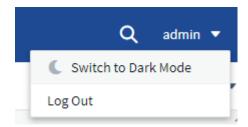
Dashboards (7)



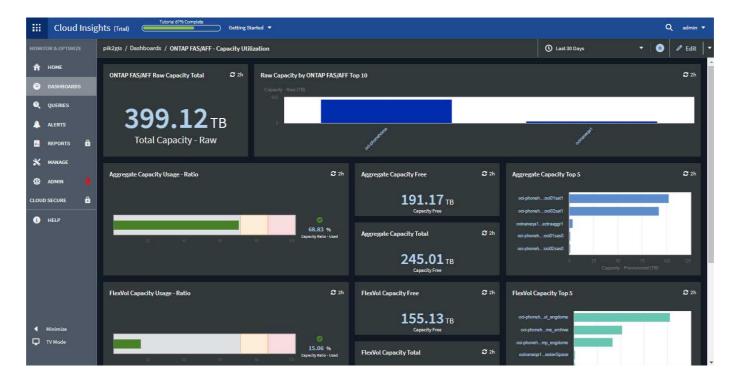
Dark Theme

You can choose to display Cloud Insights using either a light theme (the default), which displays most screens using a light background with dark text, or a dark theme which displays most screens using a dark background with light text.

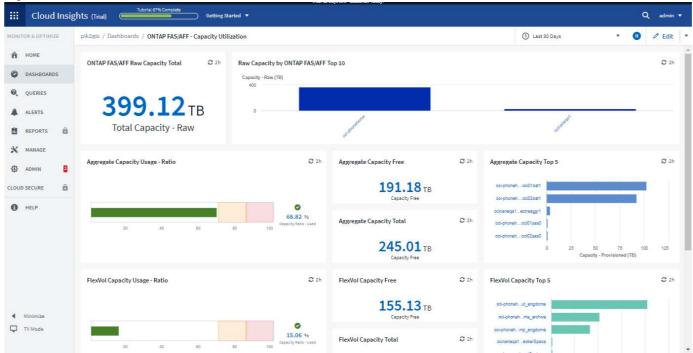
To switch between light and dark themes, click the username button in the upper right corner of the screen and choose the desired theme.



Dark Theme Dashboard view:



Light Theme Dashboard view:



(i)

Some screen areas, such as certain widget charts, still show light backgrounds even while viewed in dark theme.

Line Chart interpolation

Different data collectors often poll their data at different intervals. For example, data collector A may poll every 15 minutes while data collector B polls every five minutes. When a line chart widget (also spline, area, and stacked area charts) is aggregating this data from multiple data collectors into a single line (for example, when the widget is grouping by "all"), and refreshing the line every five minutes, data from collector B may be shown accurately while data from collector A may have gaps, thus affecting the aggregate until collector A polls again.

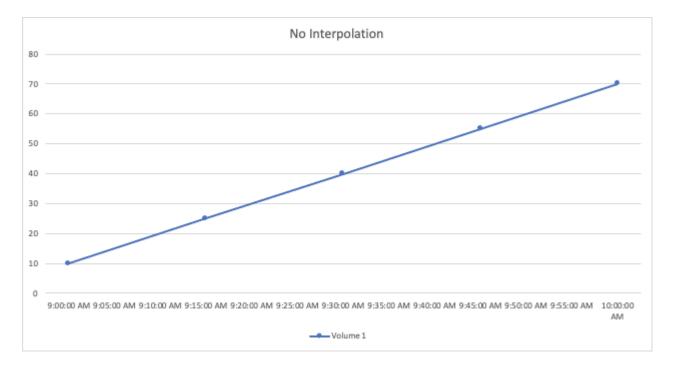
To alleviate this, Cloud Insights interpolates data when aggregating, using the surrounding data points to take a "best guess" at data until data collectors poll again. You can always view each data collector's object data individually by adjusting the widget's grouping.

Interpolation Methods

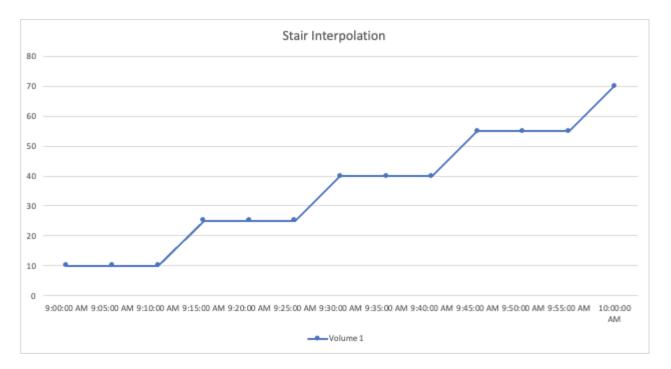
When creating or modifying a line chart (or spline, area, or stacked area chart), you can set the interpolation method to one of three types. In the "Group by" section, choose the desired Interpolation.



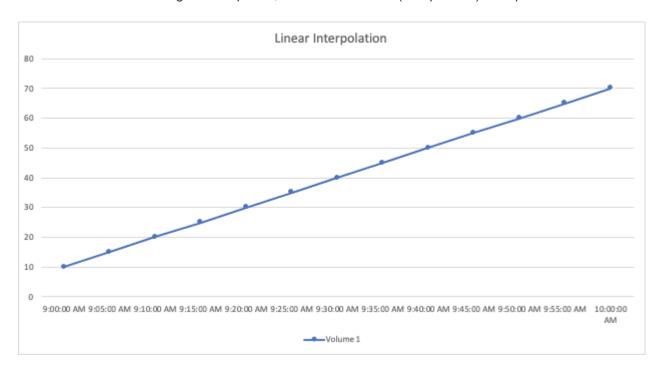
• None: Do nothing, i.e. do not generate points in between.



• **Stair**: A point is generated from the value of previous point. In a straight line, this would display as a typical "stair" layout.



• **Linear**: a point is generated as the value in between connecting the two points. Generates a line that looks like the line connecting the two points, but with additional (interpolated) data points.



Sample Dashboards

Dashboard Example: Virtual Machine Performance

There are many challenges facing IT operations today. Administrators are being asked to do more with less, and having full visibility into your dynamic data centers is a must. In this example, we will show you how to create a dashboard with widgets that give you operational insights into the virtual machine (VM) performance in your environment. By

following this example, and creating widgets to target your own specific needs, you can do things like visualizing backend storage performance compared to frontend virtual machine performance, or viewing VM latency versus I/O demand.

About this task

Here we will create a Virtual Machine Performance dashboard containing the following:

- · a table listing VM names and performance data
- · a chart comparing VM Latency to Storage Latency
- · a chart showing Read, Write and Total IOPS for VMs
- · a chart showing Max Throughput for your VMs

This is just a basic example. You can customize your dashboard to highlight and compare any performance data you choose, in order to target for your own operational best practices.

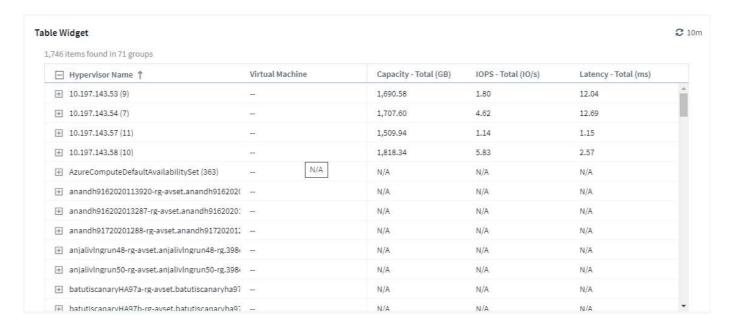
Steps

- 1. Log in to Insight as a user with administrative permissions.
- From the Dashboards menu, select [+New dashboard].

The **New dashboard** page opens.

- At the top of the page, enter a unique name for the dashboard, for example "VM Performance by Application".
- Click Save to save the dashboard with the new name.
- Let's start adding our widgets. If necessary, click the Edit icon to enable Edit mode.
- Click the Add Widget icon and select Table to add a new table widget to the dashboard.

The Edit Widget dialog opens. The default data displayed is for all storages in your environment.



1. We can customize this widget. In the Name field at the top, delete "Widget 1" and enter "Virtual Machine Performance table".

2. Click the asset type drop-down and change Storage to Virtual Machine.

The table data changes to show all virtual machines in your environment.

 Let's add a few columns to the table. Click the Gear icon on the right and select Hypervisor name, IOPS -Total, and Latency - Total. You can also try typing the name into the search to quickly display the desired field.

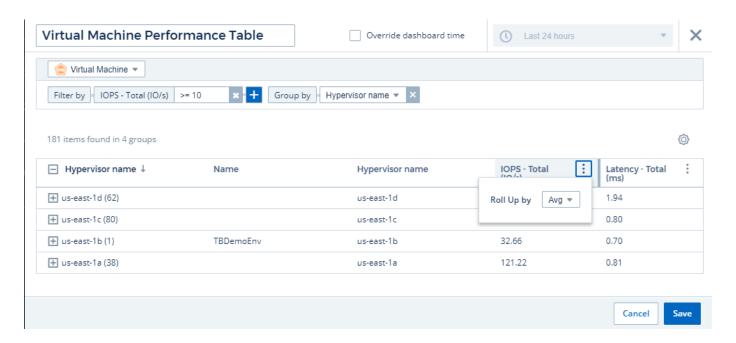
These columns are now displayed in the table. You can sort the table by any of these columns. Note that the columns are displayed in the order in which they were added to the widget.

4. For this exercise we will exclude VMs that are not actively in use, so let's filter out anything with less than 10 total IOPS. Click the [+] button next to **Filter by** and select *IOPS - Total*. Click on **Any** and enter "10" in the **from** field. Leave the **to** field empty. Click outsude the filter field or press Enter to set the filter.

The table now shows only VMs with 10 or more total IOPS.

5. We can further collapse the table by grouping results. Click the [+] button next to **Group by** and select a field to group by, such as *Application* or *Hypervisor name*. Grouping is automatically applied.

The table rows are now grouped according to your setting. You can expand and collapse the groups as needed. Grouped rows show rolled up data for each of the columns. Some columns allow you to choose the roll up method for that column.



1. When you have customized the table widget to your satisfaction, click the [Save] button.

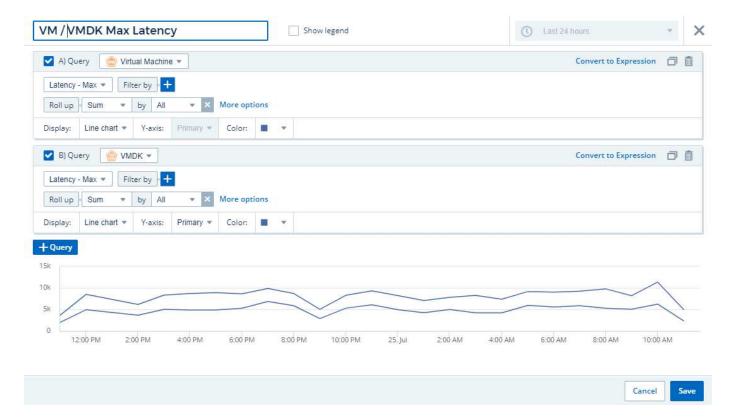
The table widget is saved to the dashboard.

You can resize the widget on the dashboard by dragging the lower-right corner. Make the widget wider to show all the columns clearly. Click **Save** to save the current dashboard.

Next we will add some charts to show our VM Performance. Let's create a line chart comparing VM latency with VMDK latency.

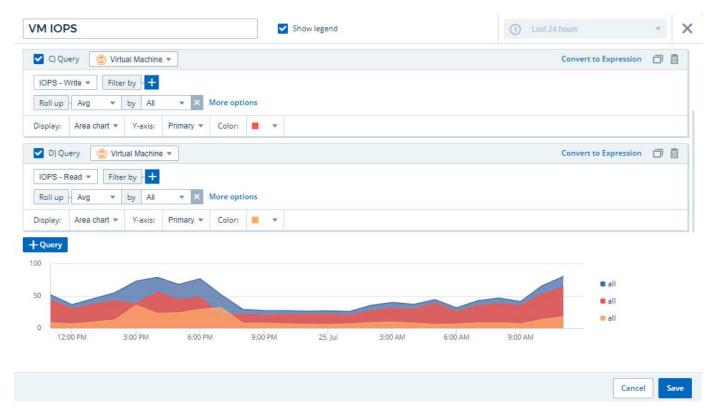
1. If necessary, click the **Edit** icon on the dashboard to enable Edit mode.

- 2. Click the [Add widget] icon and select Line Chart to add a new line chart widget to the dashboard.
- 3. The Edit Widget dialog opens. Name this widget "VM / VMDK Max Latency"
- 4. Select **Virtual Machine** and choose *Latency Max*. Set any filters you wish, or leave **Filter by** empty. For **Roll up**, choose *Sum* by *All*. Display this data as a *Line Chart*, and leave *Y-Axis* as *Primary*.
- 5. Click the **[+Query]** button to add a second data line. For this line, select *VMDK* and *Latency Max*. Set any filters you wish, or leave **Filter by** empty. For **Roll up**, choose *Sum* by *All*. Display this data as a *Line Chart*, and leave *Y-Axis* as *Primary*.
- 6. Click [Save] to add this widget to the dashboard.



Next we will add a chart showing VM Read, Write and Total IOPS in a single chart.

- 1. Click the [Add widget] icon and select Area Chart to add a new area chart widget to the dashboard.
- 2. The Edit Widget dialog opens. Name this widget "VM IOPS"
- 3. Select **Virtual Machine** and choose *IOPS Total*. Set any filters you wish, or leave **Filter by** empty. for **Roll up**, choose *Sum* by *All*. Display this data as an *Area Chart*, and leave *Y-Axis* as *Primary*.
- 4. Click the **[+Query]** button to add a second data line. For this line, select **Virtual Machine** and choose *IOPS Read*.
- Click the [+Query] button to add a third data line. For this line, select Virtual Machine and choose IOPS -Write.
- 6. Click **Show legend** to display a legend for this widget on the dashboard.



1. Click [Save] to add this widget to the dashboard.

Next we will add a chart showing VM Throughput for each Application associated with the VM. We will use the Roll Up feature for this.

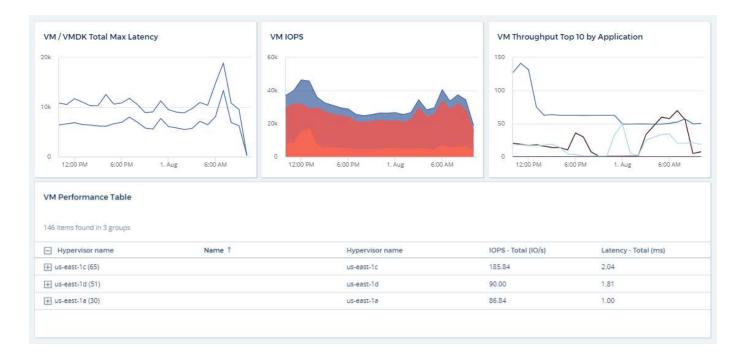
- 1. Click the [Add widget] icon and select Line Chart to add a new line chart widget to the dashboard.
- 2. The Edit Widget dialog opens. Name this widget "VM Throughput by Application"
- 3. Select Virtual Machine and choose Throughput Total. Set any filters you wish, or leave Filter by empty. For Roll up, choose "Max" and select by "Application" or "Name". Show the Top 10 applications. Display this data as a Line Chart, and leave Y-Axis as Primary.
- 4. Click [Save] to add this widget to the dashboard.

You can move widgets on the dashboard by holding down the mouse button anywhere in the top of the widget and dragging it to a new location.

You can resize widgets by dragging the lower-right corner.

Be sure to [Save] the dashboard after you make your changes.

Your final VM Performance Dashboard will look something like this:



Best Practices for Dashboards and Widgets

Tips and tricks to help you get the most out of the powerful features of dashboards and widgets.

Finding the Right Metric

Cloud Insights acquires counters and metrics using names that sometimes differ from data collector to data collector.

When searching for the right metric or counter for your dashboard widget, keep in mind that the metric you want could be under a different name from the one you are thinking of. While drop-down lists in Cloud Insights are usually alphabetical, sometimes a term may not show up in the list where you think it should. For example, terms like "raw capacity" and "used capacity" do not appear together in most lists.

Best practice: Use the search feature in fields such as Filter by or places like the column selector to find what you are looking for. For example, searching for "cap" will show all metrics with "capacity" in their names, no matter where they occur in the list. You can then easily select the metrics you want from that shorter list.

Here are a few alternative phrases you can try when searching for metrics:

When you want to find:	Try also searching for:
CPU	Processor
Capacity	Used capacity Raw capacity Provisioned capacity Storage pools capacity <other asset="" type=""> capacity Written capacity</other>

Disk Speed	Lowest disk speed Least performing disk type
Host	Hypervisor Hosts
Hypervisor	Host Is hypervisor
Microcode	Firmware
Name	Alias Hypervisor name Storage name <other asset="" type=""> name Simple name Resource name Fabric Alias</other>
Read / Write	Partial R/W Pending writes IOPS - Write Written capacity Latency - Read Cache utilization - read
Virtual Machine	VM Is virtual

This is not a comprehensive list. These are examples of possible search terms only.

Finding the Right Assets

The assets you can reference in widget filters and searches vary from asset type to asset type.

In dashboards and asset pages, the asset type around which you are building your widget determines the other asset type counters for which you can filter or add a column. Keep the following in mind when building your widget:

This asset type / counter:	Can be filtered for under these assets:
Virtual Machine	VMDK
Datastore(s)	Internal Volume VMDK Virtual Machine Volume
Hypervisor	Virtual Machine Is hypervisor Host
Host(s)	Internal Volume Volume Cluster Host Virtual Machine

Fabric	Port

This is not a comprehensive list.

Best practice: If you are filtering for a particular asset type that does not appear in the list, try building your query around an alternate asset type.

Scatter Plot Example: Knowing your Axis

Changing the order of counters in a scatter plot widget changes the axes on which the data is displayed.

About this task

This example will create a scatter plot that will allow you to see under-performing VMs that have high latency compared to low IOPS.

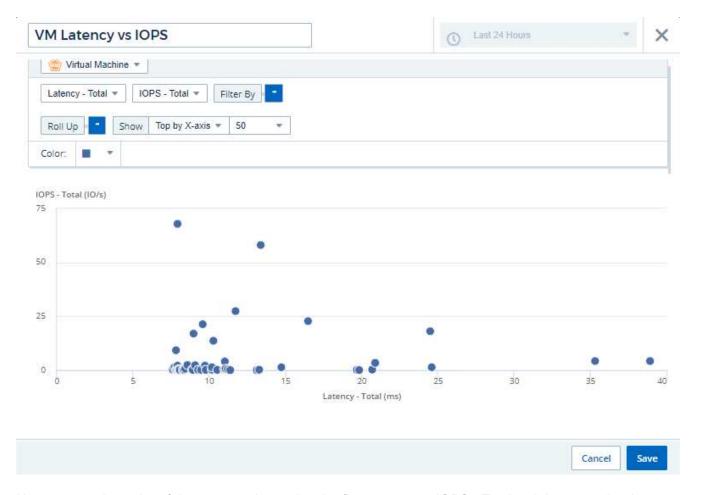
Steps

- 1. Create or open a dashboard in edit mode and add a **Scatter Plot Chart** widget.
- 2. Select an asset type, for example, Virtual Machine.
- 3. Select the first counter you wish to plot. For this example, select Latency Total.

Latency - Total is charted along the X-axis of the chart.

4. Select the second counter you wish to plot. For this example, select IOPS - Total.

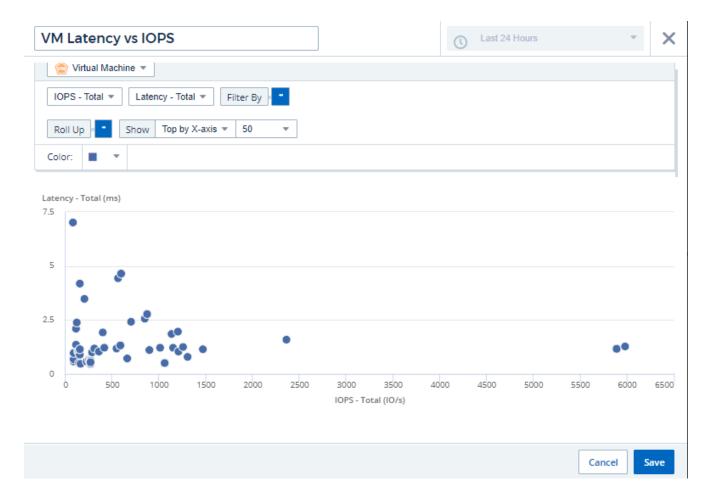
IOPS - Total is charted along the Y-axis in the chart. VMs with higher latency display on the right side of the chart. Only the top 100 highest-latency VMs are displayed, because the **Top by X-axis** setting is current.



5. Now reverse the order of the counters by setting the first counter to *IOPS - Total* and the second to *Latency - Total*.

Latency- Total is now charted along the Y-axis in the chart, and IOPS - Total along the X-axis. VMs with higher IOPS now display on the right side of the chart.

Note that because we haven't changed the **Top by X-Axis** setting, the widget now displays the top 100 highest-IOPS VMs, since this is what is currently plotted along the X-axis.



You can choose for the chart to display the Top N by X-axis, Top N by Y-axis, Bottom N by X-axis, or Bottom N by Y-axis. In our final example, the chart is displaying the Top 100 VMs that have the highest total IOPS. If we change it to **Top by Y-axis**, the chart will once again display the top 100 VMs that have the highest total latency.

Note that in a scatter plot chart, you can click on a point to drill down to the asset page for that resource.

Kubernetes

Kubernetes Cluster Overview

The Cloud Insights Kubernetes Explorer is a powerful tool for displaying the overall health and usage of your Kubernetes clusters and allows you to easily drill down into areas of investigation.

Clicking on **Dashboards > Kubernetes Explorer** opens the Kubernetes Cluster list page. This overview page contains table of the Kubernetes clusters in your environment.



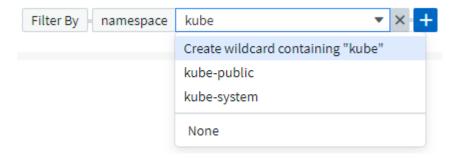
Cluster list

The cluster list displays the following information for each cluster in your environment:

- Cluster Name. Clicking on a cluster name will open the detail page for that cluster.
- Saturation percentages. Overall Saturation is the highest of CPU, Memory, or Storage Saturation.
- Number of Nodes in the cluster. Clicking this number will open the Node list page.
- Number of **Pods** in the cluster. Clicking this number will open the Pod list page.
- Number of Namespaces in the cluster. Clicking this number will open the Namespace list page.
- Number of Workloads in the cluster. Clicking this number will open the Workload list page.

Refining the Filter

When you are filtering, as you begin typing you are presented with the option to create a **wildcard filter** based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create **expressions** using NOT or AND, or you can select the "None" option to filter for null values in the field.



Filters based on wildcards or expressions (e.g. NOT, AND, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.



Kubernetes filters are contextual, meaning for example that if you are on a specific node page, the pod_name filter only lists pods related to that node. Moreover, if you apply a filter for a specific namespace, then the pod_name filter will list only pods on that node *and* in that namespace.

Note that Wildcard and Expression filtering works with text or lists but not with numerics, dates or booleans.

Before Installing or Upgrading the NetApp Kubernetes Monitoring Operator

Read this information before installing of upgrading your NetApp Kubernetes Monitoring Operator

Pre-requisites:

- If you are using a custom or private docker repository, follow the instructions in the Using a custom or private docker repository section
- NetApp Kubernetes Monitoring Operator installation is supported with Kubernetes version 1.20 or greater.
- When Cloud Insights is monitoring the backend storage and Kubernetes is used with the Docker container runtime, Cloud Insights can display pod-to-PV-to-storage mappings and metrics for NFS and iSCSI; other runtimes only show NFS.
- Beginning August 2022, the NetApp Kubernetes Monitoring Operator includes support for Pod Security Policy (PSP). You must upgrade to the latest NetApp Kubernetes Monitoring Operator if your environment uses PSP.
- If you are running on OpenShift 4.6 or higher, you must follow the OpenShift Instructions below in addition to ensuring these pre-requisites are met.
- Monitoring is only installed on Linux nodes
 Cloud Insights supports monitoring of Kubernetes nodes that are running Linux, by specifying a
 Kubernetes node selector that looks for the following Kubernetes labels on these platforms:

Platform	Label
Kubernetes v1.20 and above	Kubernetes.io/os = linux
Rancher + cattle.io as orchestration/Kubernetes platform	cattle.io/os = linux

- The NetApp Kubernetes Monitoring Operator and its dependencies (telegraf, kube-state-metrics, fluentbit, etc.) are not supported on nodes that are running with Arm64 architecture.
- The following commands must be available: curl, kubectl. The docker command is required for an optional
 installation step. For best results, add these commands to the PATH. Note that kubectl needs to be
 configured with access to the following kubernetes objects at a minimum: agents, clusterroles,
 clusterrolebindings, customresourcedefinitions, deployments, namespaces, roles, rolebindings, secrets,
 serviceaccounts, and services. See here for an example .yaml file with these minimum clusterrole
 privileges.
- The host you will use for the NetApp Kubernetes Monitoring Operator installation must have kubectl configured to communicate with the target K8s cluster, and have Internet connectivity to your Cloud Insights environment.

- If you are behind a proxy during installation, or when operating the K8s cluster to be monitored, follow the instructions in the Configuring Proxy Support section.
- The NetApp Kubernetes Monitoring Operator installs its own kube-state-metrics to avoid conflict with any
 other instances.
 - For accurate audit and data reporting, it is strongly recommended to synchronize the time on the Agent machine using Network Time Protocol (NTP) or Simple Network Time Protocol (SNTP).
- If you are re-deploying the Operator (i.e. you are updating or replacing it), there is no need to create a *new* API token; you can re-use the previous token.
- Also note that if you have a recent NetApp Kubernetes Monitoring Operator installed and are using an API
 access token that is renewable, expiring tokens will automatically be replaced by new/refreshed API
 access tokens.
- · Network monitoring:
 - Requires Linux kernel version 4.18.0 and above
 - Photon OS is not supported.

Configuring the Operator

In newer versions of the operator, most commonly modified settings can be configured in the *AgentConfiguration* custom resource. You can edit this resource before deploying the operator by editing the *operator-config.yaml* file. This file includes commented out examples of some settings. See the list of available settings for the most recent version of the operator.

You can also edit this resource after the operator has been deployed using the following command:

```
kubectl -n netapp-monitoring edit AgentConfiguration
```

To determine if your deployed version of the operator supports AgentConfiguration, run the following command:

```
kubectl get crd agentconfigurations.monitoring.netapp.com
```

If you see an "Error from server (NotFound)" message, your operator must be upgraded before you can use the AgentConfiguration.

Important Things to Note Before You Start

If you are running with a proxy, have a custom repository, or are using OpenShift, read the following sections carefully.

Also read about Permissions.

If you are upgrading from a previous installation, read the Upgrading information.

Configuring Proxy Support

There are two places where you may use a proxy in your environment in order to install the NetApp Kubernetes Monitoring Operator. These may be the same or separate proxy systems:

- Proxy needed during execution of the installation code snippet (using "curl") to connect the system where the snippet is executed to your Cloud Insights environment
- Proxy needed by the target Kubernetes cluster to communicate with your Cloud Insights environment

If you use a proxy for either or both of these, to install the NetApp Kubernetes Operating Monitor you must first ensure that your proxy is configured to allow good communication to your Cloud Insights environment. For example, from the servers/VMs from which you wish to install the Operator, you need to be able to access Cloud Insights and be able to download binaries from Cloud Insights.

For the proxy used to install the NetApp Kubernetes Operating Monitor, before installing the Operator, set the http_proxy/https_proxy environment variables. For some proxy environments, you may also need to set the no_proxy environment variable.

To set the variable(s), perform the following steps on your system **before** installing the NetApp Kubernetes Monitoring Operator:

- 1. Set the http_proxy environment variable(s) for the current user:
 - a. If the proxy being setup does not have Authentication (username/password), run the following command:

```
export https_proxy=<proxy_server>:<proxy_port>
```

b. If the proxy being setup does have Authentication (username/password), run this command:

```
export
http_proxy=proxy_username>:cproxy_password>@cproxy_server>:cproxy_po
rt>
```

For the proxy used for your Kubernetes cluster to communicate with your Cloud Insights environment, install the NetApp Kubernetes Monitoring Operator after reading all of these instructions.

Configure the proxy section of AgentConfiguration in operator-config.yaml before deploying the NetApp Kubernetes Monitoring Operator.

```
agent:
 proxy:
   server: <server for proxy>
   port: <port for proxy>
   username: <username for proxy>
   password: <password for proxy>
   # In the noproxy section, enter a comma-separated list of
    # IP addresses and/or resolvable hostnames that should bypass
    # the proxy
   noproxy: <comma separated list>
    isTelegrafProxyEnabled: true
   isFluentbitProxyEnabled: <true or false> # true if Events Log enabled
    isCollectorsProxyEnabled: <true or false> # true if Network
Performance and Map enabled
    isAuProxyEnabled: <true or false> # true if AU enabled
. . .
```

Using a custom or private docker repository

By default, the NetApp Kubernetes Monitoring Operator will pull container images from the Cloud Insights repository. If you have a Kubernetes cluster used as the target for monitoring, and that cluster is configured to only pull container images from a custom or private Docker repository or container registry, you must configure access to the containers needed by the NetApp Kubernetes Monitoring Operator.

Run the "Image Pull Snippet" from the NetApp Monitoring Operator install tile. This command will log into the Cloud Insights repository, pull all image dependencies for the operator, and log out of the Cloud Insights repository. When prompted, enter the provided repository temporary password. This command downloads all images used by the operator, including for optional features. See below for which features these images are used for.

Core Operator Functionality and Kubernetes Monitoring

- netapp-monitoring
- kube-rbac-proxy
- kube-state-metrics
- · telegraf
- · distroless-root-user

Events Log

- · fluent-bit
- · kubernetes-event-exporter

Network Performance and Map

· ci-net-observer

Push the operator docker image to your private/local/enterprise docker repository according to your corporate policies. Ensure that the image tags and directory paths to these images in your repository are consistent with those in the Cloud Insights repository.

Edit the monitoring-operator deployment in operator-deployment.yaml, and modify all image references to use your private Docker repository.

```
image: <docker repo of the enterprise/corp docker repo>/kube-rbac-
proxy:<kube-rbac-proxy version>
image: <docker repo of the enterprise/corp docker repo>/netapp-
monitoring:<version>
```

Edit the AgentConfiguration in operator-config.yaml to reflect the new docker repo location. Create a new imagePullSecret for your private repository, for more details see https://kubernetes.io/docs/tasks/configure-pod-container/pull-image-private-registry/

```
agent:

...

# An optional docker registry where you want docker images to be pulled from as compared to CI's docker registry

# Please see documentation link here: https://docs.netapp.com/us-en/cloudinsights/task_config_telegraf_agent_k8s.html#using-a-custom-or-private-docker-repository

dockerRepo: your.docker.repo/long/path/to/test

# Optional: A docker image pull secret that maybe needed for your private docker registry

dockerImagePullSecret: docker-secret-name
```

OpenShift Instructions

If you are running on OpenShift 4.6 or higher, you must edit the AgentConfiguration in *operator-config.yaml* to enable the *runPrivileged* setting:

```
# Set runPrivileged to true SELinux is enabled on your kubernetes nodes runPrivileged: true
```

Openshift may implement an added level of security that may block access to some Kubernetes components.

Permissions

If the cluster you are monitoring contains Custom Resources which do not have a ClusterRole which aggregates to view, you will need to manually grant the operator access to these resources to monitor them with Event Logs.

- 1. Edit *operator-additional-permissions.yaml* before installing, or after installing edit the resource *ClusterRole/<namespace>-additional-permissions*
- Create a new rule for the desired apiGroups and resources with the verbs ["get", "watch", "list"]. See https://kubernetes.io/docs/reference/access-authn-authz/rbac/
- 3. Apply your changes to the cluster

Tolerations and Taints

The netapp-ci-telegraf-ds, netapp-ci-fluent-bit-ds, and netapp-ci-net-observer-l4-ds DaemonSets must schedule a pod on every node in your cluster in order to correctly collect data on all nodes. The operator has been configured to tolerate some well known **taints**. If you have configured any custom taints on your nodes, thus preventing pods from running on every node, you can create a **toleration** for those taints in the AgentConfiguration. If you have applied custom taints to all nodes in your cluster, you must also add the necessary tolerations to the operator deployment to allow the operator pod to be scheduled and executed.

Learn More about Kubernetes Taints and Tolerations.

Return to the NetApp Kubernetes Monitoring Operator Installation page

Kubernetes Monitoring Operator Installation and Configuration

Cloud Insights offers the **NetApp Kubernetes Monitoring Operator** (NKMO) for Kubernetes collection. When adding a data collector, simply choose the "Kubernetes" tile.



If you have Cloud Insights Federal Edition, your installation and configuration instructions may be different than the instructions on this page. Follow the instructions in Cloud Insights to install the NetApp Kubernetes Monitoring Operator.

Choose a Data Collector to Monitor



The Kubernetes Operator and the data collectors are downloaded from the Cloud Insights Docker Registry. Once installed, the Operator then manages any Operator-compatible collectors deployed in the Kubernetes cluster nodes to acquire data, including managing the life cycle of those collectors. Following this chain, data is acquired from the collectors and sent through to Cloud Insights.

Before installing the NetApp Kubernetes Monitoring Operator



Read the **Before Installing or Upgrading** pre-requisites documentation before installing or upgrading the NetApp Kubernetes Monitoring Operator.

Installing the NetApp Kubernetes Monitoring Operator



Deploy NetApp Monitoring Operator

Provide the Kubernetes cluster name and specify a namespace for deploying the monitoring components.

Quickly install and configure a Kubernetes Operator to send cluster information to Cloud Insights.

Select existing API Access Token or create a new	v one			
KEY2024 (vw6NdM)	•	+ API Access Token	Production Best Practices 🕢	
Installation Instructions				Need Help?
Please review the pre-requisites for installing the NetApp Kube To update an existing operator installation please follow these		itoring Operator.		
Define Kubernetes cluster name and namespace				

Cluster Namespace clustername netapp-monitoring

2 Download the operator YAML files

Execute the following download command in a bash prompt.

Copy Download Command Snippet

Reveal Download Command Snippet

This snippet includes a unique access key that is valid for 24 hours.

3 Optional: Upload the operator images to your private repository

By default, the operator pulls container images from the Cloud Insights repository. To use a private repository, download the required images using the Image Pull command. Then upload them to your private repository maintaining the same tags and directory structure. Finally, update the image paths in operator-deployment and the docker repository settings in operator-config.yaml. For more information review the documentation.

Copy Image Pull Snippet

Reveal Image Pull Snippet

Copy Repository Password

This password is valid for 24 hours.

4 Optional: Review available configuration options

Configure custom options such as proxy and private repository settings. Review the instructions and available options.

5 Deploy the operator (create new or upgrade existing)

Execute the kubectl snippet to apply the following operator YAML files.

- operator-setup.yaml Create the operator's dependencies.
- operator-secrets.yaml Create secrets holding your API key.
- · operator-deployment.yaml, operator-cr.yaml Deploy the NetApp Kubernetes Monitoring Operator.
- operator-config.yaml Apply the configuration settings if not already present.

Copy kubectl Apply Snippet

Reveal kubectl Apply Snipper

After deploying the operator, delete or securely store operator-secrets.yaml.



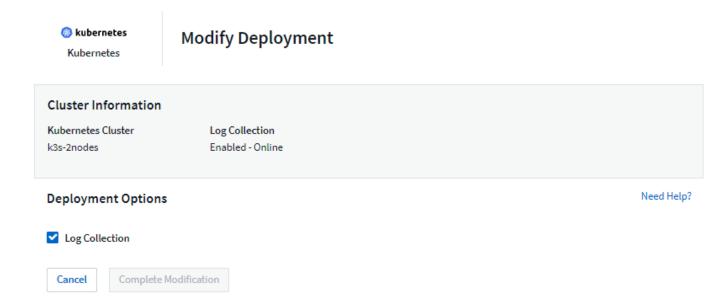
Steps to install NetApp Kubernetes Monitoring Operator agent on Kubernetes:

- 1. Enter a unique cluster name and namespace. If you are upgrading from a previous Kubernetes Operator, use the same cluster name and namespace.
- 2. Once these are entered, you can copy the Download Command snippet to the clipboard.
- 3. Paste the snippet into a *bash* window and execute it. The Operator installation files will be downloaded. Note that the snippet has a unique key and is valid for 24 hours.
- 4. If you have a custom or private repository, copy the optional Image Pull snippet, paste it into a *bash* shell and execute it. Once the images have been pulled, copy them to your private repository. Be sure to maintain the same tags and folder structure. Update the paths in *operator-deployment.yaml* as well as the docker repository settings in *operator-config.yaml*.
- 5. If desired, review available configuration options such as proxy or private repository settings. You can read more about configuration options.
- 6. When you are ready, deploy the Operator by copying the kubectl Apply snippet, downloading it, and executing it.
- 7. The installation proceeds automatically. When it is complete, click the *Next* button.
- 8. When installation is complete, click the *Next* button. Be sure to also delete or securely store the *operator-secrets.yaml* file.

Read more about configuring proxy.

Read more about using a custom/private docker repository.

Kubernetes EMS log collection is enabled by default when installing the NetApp Kubernetes Monitoring Operator. To disable this collection following installation, click the **Modify Deployment** button at the top of the Kubernetes cluster detail page, and un-select "Log collection".



This screen also shows current Log Collection status. Below are the possible states:

- Disabled
- Enabled
- Enabled Installation in progress
- · Enabled Offline
- · Enabled Online
- · Error API Key has insufficient permissions

Upgrading

Upgrading to the latest NetApp Kubernetes Monitoring Operator

Determine whether an AgentConfiguration exists with the existing Operator (if your namespace is not the default *netapp-monitoring*, substitute the appropriate namespace):

 $\verb|kubectl -n| \ \texttt{netapp-monitoring} \ \texttt{get} \ \texttt{agentconfiguration} \ \texttt{netapp-monitoring-configuration} \\$

If an AgentConfiguration exists:

- Install the latest Operator over the existing Operator.
 - Ensure you are pulling the latest container images if you are using a custom repository.

If the AgentConfiguration does not exist:

 Make note of your cluster name as recognized by Cloud Insights (if your namespace is not the default netapp-monitoring, substitute the appropriate namespace):

```
kubectl -n netapp-monitoring get agent -o
jsonpath='{.items[0].spec.cluster-name}'
```

 Create a backup of the existing Operator (if your namespace is not the default netapp-monitoring, substitute the appropriate namespace):

```
kubectl -n netapp-monitoring get agent -o yaml > agent_backup.yaml
```

- · Uninstall the existing Operator.
- · Install the latest Operator.
 - Use the same cluster name.
 - After downloading the latest Operator YAML files, port any customizations found in agent_backup.yaml to the downloaded operator-config.yaml before deploying.
 - Ensure you are pulling the latest container images if you are using a custom repository.

Stopping and Starting the Netapp Kubernetes Monitoring Operator

To stop the Netapp Kubernetes Monitoring Operator:

```
kubectl -n netapp-monitoring scale deploy monitoring-operator --replicas=0
```

To start the Netapp Kubernetes Monitoring Operator:

```
kubectl -n netapp-monitoring scale deploy monitoring-operator --replicas=1
```

Uninstalling

To remove the NetApp Kubernetes Monitoring Operator

Note that the default namespace for the NetApp Kubernetes Monitoring Operator is "netapp-monitoring". If you have set your own namespace, substitute that namespace in these and all subsequent commands and files.

Newer versions of the monitoring operator can be uninstalled with the following commands:

```
kubectl -n <NAMESPACE> delete agent -l installed-by=nkmo-<NAMESPACE>
kubectl -n <NAMESPACE> delete
clusterrole, clusterrolebinding, crd, svc, deploy, role, rolebinding, secret, sa
-l installed-by=nkmo-<NAMESPACE>
```

If the monitoring operator was deployed in its own dedicated namespace, delete the namespace:

```
kubectl delete ns <NAMESPACE>
```

If the first command returns "No resources found", use the following instructions to uninstall older versions of the monitoring operator.

Execute each of the following commands in order. Depending on your current installation, some of these commands may return 'object not found' messages. These messages may be safely ignored.

```
kubectl -n <NAMESPACE> delete agent agent-monitoring-netapp
kubectl delete crd agents.monitoring.netapp.com
kubectl -n <NAMESPACE> delete role agent-leader-election-role
kubectl delete clusterrole agent-manager-role agent-proxy-role agent-
metrics-reader <NAMESPACE>-agent-manager-role <NAMESPACE>-agent-proxy-role
<NAMESPACE>-cluster-role-privileged
kubectl delete clusterrolebinding agent-manager-rolebinding agent-proxy-
rolebinding agent-cluster-admin-rolebinding <NAMESPACE>-agent-manager-
rolebinding <NAMESPACE>-agent-proxy-rolebinding <NAMESPACE>-cluster-role-
binding-privileged
kubectl delete <NAMESPACE>-psp-nkmo
kubectl delete ns <NAMESPACE>
```

If a Security Context Constraint was previously-created:

```
kubectl delete scc telegraf-hostaccess
```

About Kube-state-metrics

The NetApp Kubernetes Monitoring Operator installs kube-state-metrics automatically; no user interaction is needed.

kube-state-metrics Counters

Use the following links to access information for these kube state metrics counters:

- 1. ConfigMap Metrics
- 2. DaemonSet Metrics
- 3. Deployment Metrics
- 4. Ingress Metrics
- Namespace Metrics
- 6. Node Metrics
- 7. Persistent Volume Metrics
- 8. Persistant Volume Claim Metrics
- 9. Pod Metrics

- 10. ReplicaSet metrics
- 11. Secret metrics
- 12. Service metrics
- 13. StatefulSet metrics

```
== Configuring the Operator
```

In newer versions of the operator, most commonly modified settings can be configured in the *AgentConfiguration* custom resource. You can edit this resource before deploying the operator by editing the *operator-config.yaml* file. This file includes commented out examples of some settings. See the list of available settings for the most recent version of the operator.

You can also edit this resource after the operator has been deployed using the following command:

```
kubectl -n netapp-monitoring edit AgentConfiguration
```

To determine if your deployed version of the operator supports AgentConfiguration, run the following command:

```
kubectl get crd agentconfigurations.monitoring.netapp.com
```

If you see an "Error from server (NotFound)" message, your operator must be upgraded before you can use the AgentConfiguration.

Configuring Proxy Support

There are two places where you may use a proxy in your environment in order to install the NetApp Kubernetes Monitoring Operator. These may be the same or separate proxy systems:

- Proxy needed during execution of the installation code snippet (using "curl") to connect the system where the snippet is executed to your Cloud Insights environment
- Proxy needed by the target Kubernetes cluster to communicate with your Cloud Insights environment

If you use a proxy for either or both of these, in order to install the NetApp Kubernetes Operating Monitor you must first ensure that your proxy is configured to allow good communication to your Cloud Insights environment. If you have a proxy and can access Cloud Insights from the server/VM from which you wish to install the Operator, then your proxy is likely configured properly.

For the proxy used to install the NetApp Kubernetes Operating Monitor, before installing the Operator, set the http_proxy/https_proxy environment variables. For some proxy environments, you may also need to set the no_proxy environment variable.

To set the variable(s), perform the following steps on your system **before** installing the NetApp Kubernetes Monitoring Operator:

1. Set the https proxy and/or http proxy environment variable(s) for the current user:

a. If the proxy being setup does not have Authentication (username/password), run the following command:

```
export https_proxy=<proxy_server>:<proxy_port>
```

b. If the proxy being setup does have Authentication (username/password), run this command:

```
export
http_proxy=proxy_username>:cy_password>@expoxy_server>:cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_password>@cy_passwor
```

For the proxy used for your Kubernetes cluster to communicate with your Cloud Insights environment, install the NetApp Kubernetes Monitoring Operator after reading all of these instructions.

Configure the proxy section of AgentConfiguration in operator-config.yaml before deploying the NetApp Kubernetes Monitoring Operator.

```
agent:
 . . .
 proxy:
   server: <server for proxy>
   port: <port for proxy>
   username: <username for proxy>
   password: <password for proxy>
   # In the noproxy section, enter a comma-separated list of
    # IP addresses and/or resolvable hostnames that should bypass
   # the proxy
   noproxy: <comma separated list>
   isTelegrafProxyEnabled: true
   isFluentbitProxyEnabled: <true or false> # true if Events Log enabled
   isCollectorsProxyEnabled: <true or false> # true if Network
Performance and Map enabled
   isAuProxyEnabled: <true or false> # true if AU enabled
. . .
```

Using a custom or private docker repository

By default, the NetApp Kubernetes Monitoring Operator will pull container images from the Cloud Insights repository. If you have a Kubernetes cluster used as the target for monitoring, and that cluster is configured to only pull container images from a custom or private Docker repository or container registry, you must configure access to the containers needed by the NetApp Kubernetes Monitoring Operator.

Run the "Image Pull Snippet" from the NetApp Monitoring Operator install tile. This command will log into the Cloud Insights repository, pull all image dependencies for the operator, and log out of the Cloud Insights repository. When prompted, enter the provided repository temporary password. This command downloads all images used by the operator, including for optional features. See below for which features these images are used for.

Core Operator Functionality and Kubernetes Monitoring

- netapp-monitoring
- · ci-kube-rbac-proxy
- · ci-ksm
- · ci-telegraf
- · distroless-root-user

Events Log

- · ci-fluent-bit
- · ci-kubernetes-event-exporter

Network Performance and Map

· ci-net-observer

Push the operator docker image to your private/local/enterprise docker repository according to your corporate policies. Ensure that the image tags and directory paths to these images in your repository are consistent with those in the Cloud Insights repository.

Edit the monitoring-operator deployment in operator-deployment.yaml, and modify all image references to use your private Docker repository.

```
image: <docker repo of the enterprise/corp docker repo>/kube-rbac-
proxy:<ci-kube-rbac-proxy version>
image: <docker repo of the enterprise/corp docker repo>/netapp-
monitoring:<version>
```

Edit the AgentConfiguration in operator-config.yaml to reflect the new docker repo location. Create a new imagePullSecret for your private repository, for more details see https://kubernetes.io/docs/tasks/configure-pod-container/pull-image-private-registry/

```
agent:
...
# An optional docker registry where you want docker images to be pulled from as compared to CI's docker registry
# Please see documentation link here: https://docs.netapp.com/us-en/cloudinsights/task_config_telegraf_agent_k8s.html#using-a-custom-or-private-docker-repository
dockerRepo: your.docker.repo/long/path/to/test
# Optional: A docker image pull secret that maybe needed for your private docker registry
dockerImagePullSecret: docker-secret-name
```

OpenShift Instructions

If you are running on OpenShift 4.6 or higher, you must edit the AgentConfiguration in *operator-config.yaml* to enable the *runPrivileged* setting:

```
# Set runPrivileged to true SELinux is enabled on your kubernetes nodes runPrivileged: true
```

Openshift may implement an added level of security that may block access to some Kubernetes components.

A Note About Secrets

To remove permission for the NetApp Kubernetes Monitoring Operator to view secrets cluster-wide, delete the following resources from the *operator-setup.yaml* file before installing:

```
ClusterRole/netapp-ci-<namespace>-agent-secret-clusterrole
ClusterRoleBinding/netapp-ci-<namespace>-agent-secret-clusterrolebinding
```

If this is an upgrade, also delete the resources from your cluster:

```
kubectl delete ClusterRole/netapp-ci-<namespace>-agent-secret-clusterrole
kubectl delete ClusterRoleBinding/netapp-ci-<namespace>-agent-secret-
clusterrolebinding
```

If Change Analysis is enabled, modify the *AgentConfiguration* or *operator-config.yaml* to uncomment the change-management section and include *kindsTolgnoreFromWatch: "secrets"* under the change-management section. Note the presence and position of single and double quotes in this line.

```
# change-management:
...
# # A comma separated list of kinds to ignore from watching from the
default set of kinds watched by the collector
# # Each kind will have to be prefixed by its apigroup
# # Example: '"networking.k8s.io.networkpolicies,batch.jobs",
"authorization.k8s.io.subjectaccessreviews"'
kindsToIgnoreFromWatch: '"secrets"'
...
```

Verifying Kubernetes Checksums

The Cloud Insights agent installer performs integrity checks, but some users may want to perform their own verifications before installing or applying downloaded artifacts. To perform a download-only operation (as opposed to the default download-and-install), these users can edit the agent installation command obtained from the UI and remove the trailing "install" option.

Follow these steps:

- 1. Copy the Agent Installer snippet as directed.
- 2. Instead of pasting the snippet into a command window, paste it into a text editor.
- 3. Remove the trailing "--install" from the command.
- 4. Copy the entire command from the text editor.
- 5. Now paste it into your command window (in a working directory) and run it.
 - Download and install (default):

```
installerName=cloudinsights-rhel_centos.sh ... && sudo -E -H
./$installerName --download --install
```

Download-only:

```
installerName=cloudinsights-rhel_centos.sh ... && sudo -E -H
./$installerName --download
```

The download-only command will download all required artifacts from Cloud Insights to the working directory. The artifacts include, but may not be limited to:

- · an installation script
- an environment file
- YAML files
- a signed checksum file (sha256.signed)
- a PEM file (netapp_cert.pem) for signature verification

The installation script, environment file, and YAML files can be verified using visual inspection.

The PEM file can be verified by confirming its fingerprint to be the following:

1A918038E8E127BB5C87A202DF173B97A05B4996

More specifically,

openssl x509 -fingerprint -shal -noout -inform pem -in netapp_cert.pem

The signed checksum file can be verified using the PEM file:

openssl smime -verify -in sha256.signed -CAfile netapp_cert.pem -purpose any

Once all of the artifacts have been satisfactorily verified, the agent installation can be initiated by running:

sudo -E -H ./<installation script name> --install

Troubleshooting

Some things to try if you encounter problems setting up the NetApp Kubernetes Monitoring Operator:

Problem:	Try this:
I do not see a hyperlink/connection between my Kubernetes Persistent Volume and the corresponding back-end storage device. My Kubernetes Persistent Volume is configured using the hostname of the storage server.	Follow the steps to uninstall the existing Telegraf agent, then re-install the latest Telegraf agent. You must be using Telegraf version 2.0 or later, and your Kubernetes cluster storage must be actively monitored by Cloud Insights.

Problem:

I'm seeing messages in the logs resembling the following:

E0901 15:21:39.962145 1 reflector.go:178] k8s.io/kube-state-metrics/internal/store/builder.go:352: Failed to list *v1.MutatingWebhookConfiguration: the server could not find the requested resource E0901 15:21:43.168161 1 reflector.go:178] k8s.io/kube-state-metrics/internal/store/builder.go:352: Failed to list *v1.Lease: the server could not find the requested resource (get leases.coordination.k8s.io) etc.

Try this:

These messages may occur if you are running kubestate-metrics version 2.0.0 or above with Kubernetes versions below 1.20.

To get the Kubernetes version:

kubectl version

To get the kube-state-metrics version:

kubectl get deploy/kube-state-metrics -o jsonpath='{..image}'

To prevent these messages from happening, users can modify their kube-state-metrics deployment to disable the following Leases:

mutatingwebhookconfigurations validatingwebhookconfigurations volumeattachments resources

More specifically, they can use the following CLI argument:

resources=certificatesigningrequests,configmaps,cron jobs,daemonsets,

deployments, endpoints, horizontal podautos calers, ingresses, jobs, limitranges,

namespaces, network policies, nodes, persistent volume claims, persistent volumes,

poddisruptionbudgets,pods,replicasets,replicationcont rollers,resourcequotas,

secrets, services, stateful sets, storage classes

The default resource list is:

"certificatesigningrequests,configmaps,cronjobs,daem onsets,deployments,

endpoints, horizontal podautos calers, ingresses, jobs, lea ses, limitranges,

mutatingwebhookconfigurations,namespaces,network policies,nodes,

persistentvolumeclaims, persistentvolumes, poddisrupti onbudgets, pods, replicasets,

replicationcontrollers,resourcequotas,secrets,services, statefulsets,storageclasses,

validatingwebhookconfigurations,volumeattachments"

Problem: Try this: This is a known issue. Refer to This GitHub article for I see error messages from Telegraf resembling the following, but Telegraf does start up and run: more details. As long as Telegraf is up and running. users can ignore these error messages. Oct 11 14:23:41 ip-172-31-39-47 systemd[1]: Started The plugin-driven server agent for reporting metrics into InfluxDB. Oct 11 14:23:41 ip-172-31-39-47 telegraf[1827]: time="2021-10-11T14:23:41Z" level=error msg="failed to create cache directory. /etc/telegraf/.cache/snowflake, err: mkdir /etc/telegraf/.ca che: permission denied. ignored\n" func="gosnowflake.(*defaultLogger).Errorf" file="log.go:120" Oct 11 14:23:41 ip-172-31-39-47 telegraf[1827]: time="2021-10-11T14:23:41Z" level=error msg="failed to open. Ignored. open /etc/telegraf/.cache/snowflake/ocsp_response_cache.j son: no such file or directorv\n" func="gosnowflake.(*defaultLogger).Errorf" file="log.go:120" Oct 11 14:23:41 ip-172-31-39-47 telegraf[1827]: 2021-10-11T14:23:41Z I! Starting Telegraf 1.19.3 On Kubernetes, my Telegraf pod(s) are reporting the If SELinux is enabled and enforcing, it is likely preventing the Telegraf pod(s) from accessing the following error: "Error in processing mountstats info: failed to open /proc/1/mountstats file on the Kubernetes node. To mountstats file: /hostfs/proc/1/mountstats, error: open overcome this restriction, edit the agent configuration, /hostfs/proc/1/mountstats: permission denied" and enable the runPrivileged setting. For more details, refer to: https://docs.netapp.com/us-en/cloudinsights/ task config telegraf agent k8s.html#openshiftinstructions. On Kubernetes, my Telegraf ReplicaSet pod is The Telegraf ReplicaSet pod is intended to run on a node designated as a master or for etcd. If the reporting the following error: ReplicaSet pod is not running on one of these nodes. [inputs.prometheus] Error in plugin: could not load you will get these errors. Check to see if your master/etcd nodes have taints on them. If they do. kevpair /etc/kubernetes/pki/etcd/server.crt:/etc/kubernetes/pki/ add the necessary tolerations to the Telegraf etcd/server.key: open ReplicaSet, telegraf-rs. /etc/kubernetes/pki/etcd/server.crt: no such file or directory For example, edit the ReplicaSet... kubectl edit rs telegraf-rs ...and add the appropriate tolerations to the spec. Then, restart the ReplicaSet pod.

Problem:	Try this:
I have a PSP/PSA environment. Does this affect my monitoring operator?	If your Kubernetes cluster is running with Pod Security Policy (PSP) or Pod Security Admission (PSA) in place, you must upgrade to the latest NetApp Kubernetes Monitoring Operator. Follow these steps to upgrade to the current NKMO with support for PSP/PSA: 1. Uninstall the previous monitoring operator: kubectl delete agent agent-monitoring-netapp -n netapp-monitoring kubectl delete rs netapp-monitoring kubectl delete crd agents.monitoring.netapp.com kubectl delete clusterrole agent-manager-role agent-proxy-role agent-metrics-reader kubectl delete clusterrolebinding agent-manager-rolebinding agent-proxy-rolebinding agent-cluster-admin-rolebinding 2. Install the latest version of the monitoring operator.
I ran into issues trying to deploy the NKMO, and I have PSP/PSA in use.	1. Edit the agent using the following command: kubectl -n <name-space> edit agent 2. Mark 'security-policy-enabled' as 'false'. This will disable Pod Security Policies and Pod Security Admission and allow the NKMO to deploy. Confirm by using the following commands: kubectl get psp (should show Pod Security Policy removed) kubectl get all -n <namespace> grep -i psp (should show that nothing is found)</namespace></name-space>
"ImagePullBackoff" errors seen	These errors may be seen if you have a custom or private docker repository and have not yet configured the NetApp Kubernetes Monitoring Operator to properly recognize it. Read more about configuring for custom/private repo.

Problem:	Try this:
I am having an issue with my monitoring-operator deployment, and the current documentation does not help me resolve it.	Capture or otherwise note the output from the following commands, and contact the Technical Support team.
	<pre>kubectl -n netapp-monitoring get all kubectl -n netapp-monitoring describe all kubectl -n netapp-monitoring logs <monitoring-operator-pod>all -containers=true kubectl -n netapp-monitoring logs <telegraf-pod>all -containers=true</telegraf-pod></monitoring-operator-pod></pre>
net-observer (Workload Map) pods in NKMO namespace are in CrashLoopBackOff	These pods correspond to Workload Map data collector for Network Observability. Try these: • Check the logs of one of the pods to confirm minimum kernel version. For example: {"ci-tenant-id":"your-tenant-id","collector-cluster":"your-k8s-cluster-name","environment":"prod","level":"error","msg":"faile d in validation. Reason: kernel version 3.10.0 is less than minimum kernel version of 4.18.0","time":"2022-11-09T08:23:08Z"} • Net-observer pods requires the Linux kernel version to be at least 4.18.0. Check the kernel version using the command "uname -r" and ensure they are >= 4.18.0
Pods are running in NKMO namespace (default: netapp-monitoring), but no data is shown in UI for workload map or Kubernetes metrics in Queries	Check the time setting on the nodes of the K8S cluster. For accurate audit and data reporting, it is strongly recommended to synchronize the time on the Agent machine using Network Time Protocol (NTP) or Simple Network Time Protocol (SNTP).
Some of the net-observer pods in NKMO namespace are in Pending state	Net-observer is a DaemonSet and runs a pod in each Node of the k8s cluster. • Note the pod which is in Pending state, and check if it is experiencing a resource issue for CPU or memory. Ensure the required memory and CPU is available in the node.

Problem:	Try this:
I'm seeing the following in my logs immediately after installing the NetApp Kubernetes Monitoring Operator: [inputs.prometheus] Error in plugin: error making HTTP request to http://kube-state-metrics. <namespace>.svc.cluster.local:8080/metrics: Get http://kube-state-metrics.<namespace>.svc.cluster.local:8080/metrics: dial tcp: lookup kube-state-metrics.<namespace>.svc.cluster.local: no such host</namespace></namespace></namespace>	This message is typically only seen when a new operator is installed and the <i>telegraf-rs</i> pod is up before the <i>ksm</i> pod is up. These messages should stop once all pods are running.
I do see not any metrics being collected for the Kubernetes CronJobs that exist in my cluster.	Verify your Kubernetes version (i.e. kubectl version). If it is v1.20.x or below, this is an expected limitation. The kube-state-metrics release deployed with the Netapp Kubernetes Monitoring Operator only supports v1.CronJob. With Kubernetes 1.20.x and below, the CronJob resource is at v1beta.CronJob. As a result, kube-state-metrics cannot find the CronJob resource.
After installing the operator, the telegraf-ds pods enter CrashLoopBackOff and the pod logs indicate "su: Authentication failure".	Edit the telegraf section in AgentConfiguration, and set dockerMetricCollectionEnabled to false. For more details, refer to the operator's configuration options. NOTE: If you are using Cloud Insights Federal Edition, users with restrictions on the use of su will not be able to collect docker metrics because access to the docker socket requires either running the telegraf container as root or using su to add the telegraf user to the docker group. Docker metric collection and the use of su is enabled by default; to disable both, remove the telegraf.docker entry in the AgentConfiguration file: - name: docker run-mode: - DaemonSet substitutions: - key: DOCKER_UNIX_SOCK_PLACEHOLDER value: unix:///run/docker.sock

Problem:	Try this:
I see repeating error messages resembling the following in my Telegraf logs: E! [agent] Error writing to outputs.http: Post "https:// <tenant_url>/rest/v1/lake/ingest/influxdb": context deadline exceeded (Client.Timeout exceeded while awaiting headers)</tenant_url>	Edit the telegraf section in <i>AgentConfiguration</i> , and set <i>dockerMetricCollectionEnabled</i> to false. For more details, refer to the operator's configuration options.
I'm missing involvedobject data for some Event Logs.	Be sure you have followed the steps in the Permissions section above.
Why am I seeing two monitoring operator pods running, one named netapp-ci-monitoring-operator- <pod> and the other named monitoring-operator- <pod>?</pod></pod>	As of October 12, 2023, Cloud Insights has refactored the operator to better serve our users; for those changes to be fully adopted, you must remove the old operator and install the new one.
My kubernetes events unexpectedly stopped reporting to Cloud Insights.	·

Problem:	Try this:
I'm seeing pod(s) deployed by the Netapp Kubernetes Monitoring Operator crash because of insufficient resources.	Refer to the Netapp Kubernetes Monitoring Operator configuration options to increase the CPU and/or memory limits as needed.

Additional information may be found from the Support page or in the Data Collector Support Matrix.

NetApp Kubernetes Monitoring Operator Configuration Options

The NetApp Kubernetes Monitoring Operator installation and configuration can be customized.

The table below lists the possible options for the AgentConfiguration file:

Component	Option	Description
agent		Configuration options that are common to all components that the operator can install. These can be considered as "global" options.
	dockerRepo	A dockerRepo override to pull images from customers private docker repos as compared to Cloud Insights docker repo. Default is cloud insights docker repo
	dockerlmagePullSecret	Optional: A secret for the customers private repo
	clusterName	Free text field that uniquely identifies a cluster across all customers clusters. This should be unique across a cloud insights tenant. Default is what the customer enters in the UI for the "Cluster Name" field
	proxy Format: proxy: server: port: username: password: noProxy: isTelegrafProxyEnabled: isAuProxyEnabled: isFluentbitProxyEnabled: isCollectorProxyEnabled:	Optional to set proxy. This is usually the customer's corporate proxy.
telegraf		Configuration options that can customize the telegraf installation of the Operator
	collectionInterval	Metrics collection interval, in seconds (Max=60s)
	dsCpuLimit	CPU Limit for telegraf ds

Component	Option	Description
	dsMemLimit	Memory limit for telegraf ds
	dsCpuRequest	CPU request for telegraf ds
	dsMemRequest	Memory request for telegraf ds
	rsCpuLimit	CPU Limit for telegraf rs
	rsMemLimit	Memory limit for telegraf rs
	rsCpuRequest	CPU request for telegraf rs
	rsMemRequest	Memory request for telegraf rs
	dockerMountPoint	an override for dockerMountPoint path. This is for non standard docker installations on k8s platforms like cloud foundry
	dockerUnixSocket	an override for dockerUnixSocket path. This is for non standard docker installations on k8s platforms like cloud foundry.
	crioSockPath	an override for crioSockPath path. This is for non standard docker installations on k8s platforms like cloud foundry.
	runPrivileged	Run the telegraf container in privileged mode. Set this to true if SELinux is enabled on your k8s nodes
	batchSize	See Telegraf configuration documentation
	bufferLimit	See Telegraf configuration documentation
	roundInterval	See Telegraf configuration documentation
	collectionJitter	See Telegraf configuration documentation
	precision	See Telegraf configuration documentation
	flushInterval	See Telegraf configuration documentation
	flushJitter	See Telegraf configuration documentation
	outputTimeout	See Telegraf configuration documentation
	dockerMetricCollectionEn abled	Collect Docker metrics. By default, this is set to true and docker metrics will be collected for on-premise, docker-based k8s deployments. To disable docker metric collection, set this to false.
	dsTolerations	telegraf-ds additional tolerations.
	rsTolerations	telegraf-rs additional tolerations.
kube-state-metrics		Configuration options that can customize the kube state metrics installation of the Operator
	cpuLimit	CPU limit for kube-state-metrics deployment
	memLimit	Mem limit for kube-state-metrics deployment
	cpuRequest	CPU request for kube state metrics deployment

Component	Option	Description
	memRequest	Mem request for kube state metrics deployment
	resources	a comma separated list of resources to capture. example: cronjobs,daemonsets,deployments,ingresses,jobs,na mespaces,nodes,persistentvolumeclaims, persistentvolumes,pods,replicasets,resourcequotas,se rvices,statefulsets
	tolerations	kube-state-metrics additional tolerations.
	labels	a comma separated list of resources that kube-state-metrics should capture example: cronjobs=[*],daemonsets=[*],deployments=[*],ingress es=[*],jobs=[*],namespaces=[*],nodes=[*],
		persistentvolumeclaims=[*],persistentvolumes=[*],pod s=[*],replicasets=[*],resourcequotas=[*],services=[*],st atefulsets=[*]
logs		Configuration options that can customize logs collection and installation of the Operator
	readFromHead	true/false, should fluent bit read the log from head
	timeout	timeout, in secs
	dnsMode	TCP/UDP, mode for DNS
	fluent-bit-tolerations	fluent-bit-ds additional tolerations.
	event-exporter-tolerations	event-exporter additional tolerations.
workload-map		Configuration options that can customize the workload map collection and installation of the Operator
	cpuLimit	CPU limit for net observer ds
	memLimit	mem limit for net observer ds
	cpuRequest	CPU request for net observer ds
	memRequest	mem request for net observer ds
	metricAggregationInterval	metric aggregation interval, in seconds
	bpfPollInterval	BPF poll interval, in seconds
	enableDNSLookup	true/false, enable DNS lookup
	I4-tolerations	net-observer-l4-ds additional tolerations.
	runPrivileged	true/false - Set runPrivileged to true if SELinux is enabled on your Kubernetes nodes.
change-management		Configuration options for Kubernetes Change Management and Analysis

Component	Option	Description
	cpuLimit	CPU limit for change-observer-watch-rs
	memLimit	Mem limit for change-observer-watch-rs
	cpuRequest	CPU request for change-observer-watch-rs
	memRequest	mem request for change-observer-watch-rs
	failureDeclarationInterval Mins	Interval in minutes after which a non-successful deployment of a workload will be marked as failed
	deployAggrIntervalSecond s	Frequency at which workload deployment in-progress events are sent
	nonWorkloadAggrInterval Seconds	Frequency at which non-workload deployments are combined and sent
	termsToRedact	A set of regular expressions used in env names and data maps whose value will be redacted Example terms:"pwd", "password", "token", "apikey", "api-key", "jwt"
	additionalKindsToWatch	A comma separated list of additional kinds to watch from the default set of kinds watched by the collector
	kindsTolgnoreFromWatch	A comma separated list of kinds to ignore from watching from the default set of kinds watched by the collector
	logRecordAggrIntervalSec onds	Frequency with which log records are sent to CI from the collector
	watch-tolerations	change-observer-watch-ds additional tolerations. Abbreviated single line format only. Example: '{key: taint1, operator: Exists, effect: NoSchedule},{key: taint2, operator: Exists, effect: NoExecute}'

Sample AgentConfiguration file

Below is a sample AgentConfiguration file.

```
apiVersion: monitoring.netapp.com/vlalpha1
kind: AgentConfiguration
metadata:
   name: netapp-monitoring-configuration
   namespace: "NAMESPACE_PLACEHOLDER"
   labels:
      installed-by: nkmo-NAMESPACE_PLACEHOLDER

spec:
   # # You can modify the following fields to configure the operator.
   # # Optional settings are commented out and include default values for reference
```

```
# # To update them, uncomment the line, change the value, and apply
the updated AgentConfiguration.
 agent:
    # # [Required Field] A uniquely identifiable user-friendly
    # # clusterName must be unique across all clusters in your Cloud
Insights environment.
    clusterName: "CLUSTERNAME PLACEHOLDER"
    # # Proxy settings. The proxy that the operator should use to send
metrics to Cloud Insights.
    # # Please see documentation here: https://docs.netapp.com/us-
en/cloudinsights/task config telegraf agent k8s.html#configuring-proxy-
support
    # proxy:
    # server:
    # port:
    # noproxy:
    # username:
    # password:
    # isTelegrafProxyEnabled:
    # isFluentbitProxyEnabled:
    # isCollectorsProxyEnabled:
    # # [Required Field] By default, the operator uses the CI repository.
    # # To use a private repository, change this field to your repository
name.
    # # Please see documentation here: https://docs.netapp.com/us-
en/cloudinsights/task config telegraf agent k8s.html#using-a-custom-or-
private-docker-repository
    dockerRepo: 'DOCKER REPO PLACEHOLDER'
    # # [Required Field] The name of the imagePullSecret for dockerRepo.
    # # If you are using a private repository, change this field from
'docker' to the name of your secret.
    {{ if not (contains .Values.config.cloudType "aws") }}# {{ end -}}
    dockerImagePullSecret: 'docker'
    # # Allow the operator to automatically rotate its ApiKey before
expiration.
    # tokenRotationEnabled: '{{
.Values.telegraf installer.kubernetes.rs.shim token rotation }}'
    # # Number of days before expiration that the ApiKey should be
rotated. This must be less than the total ApiKey duration.
    # tokenRotationThresholdDays: '{{
.Values.telegraf installer.kubernetes.rs.shim token rotation threshold day
s }}'
```

```
telegraf:
    # # Settings to fine-tune metrics data collection. Telegraf config
names are included in parenthesis.
https://github.com/influxdata/telegraf/blob/master/docs/CONFIGURATION.md#a
    # # The default time telegraf will wait between inputs for all plugins
(interval). Max=60
    # collectionInterval: '{{
.Values.telegraf installer.agent resources.collection interval }}'
    # # Maximum number of records per output that telegraf will write in
one batch (metric batch size).
    # batchSize: '{{
.Values.telegraf installer.agent resources.metric batch size }}'
    # # Maximum number of records per output that telegraf will cache
pending a successful write (metric_buffer_limit).
    # bufferLimit: '{{
.Values.telegraf installer.agent resources.metric buffer limit }}'
    # # Collect metrics on multiples of interval (round interval).
    # roundInterval: '{{
.Values.telegraf installer.agent resources.round interval }}'
    # # Each plugin waits a random amount of time between the scheduled
collection time and that time + collection jitter before collecting inputs
(collection jitter).
    # collectionJitter: '{{
.Values.telegraf installer.agent resources.collection jitter }}'
    # # Collected metrics are rounded to the precision specified. When set
to "0s" precision will be set by the units specified by interval
(precision).
    # precision: '{{ .Values.telegraf installer.agent resources.precision
}}'
    # # Time telegraf will wait between writing outputs (flush interval).
Max=collectionInterval
    # flushInterval: '{{
.Values.telegraf installer.agent resources.flush interval }}'
    # # Each output waits a random amount of time between the scheduled
write time and that time + flush jitter before writing outputs
(flush jitter).
    # flushJitter: '{{
.Values.telegraf installer.agent resources.flush jitter }}'
    # # Timeout for writing to outputs (timeout).
    # outputTimeout: '{{
.Values.telegraf installer.http output plugin.timeout }}'
```

```
# # telegraf-ds CPU/Mem limits and requests.
    # # See https://kubernetes.io/docs/concepts/configuration/manage-
resources-containers/
    dsCpuLimit: '{{
.Values.telegraf installer.telegraf resources.ds cpu limits }}'
    dsMemLimit: '{{
.Values.telegraf installer.telegraf resources.ds mem limits }}'
    dsCpuRequest: '{{
.Values.telegraf installer.telegraf resources.ds cpu request }}'
    dsMemRequest: '{{
.Values.telegraf_installer.telegraf_resources.ds mem request }}'
    # # telegraf-rs CPU/Mem limits and requests.
    rsCpuLimit: '{{
.Values.telegraf installer.telegraf resources.rs cpu limits }}'
    rsMemLimit: '{{
.Values.telegraf installer.telegraf resources.rs mem limits }}'
    rsCpuRequest: '{{
.Values.telegraf installer.telegraf resources.rs cpu request }}'
    rsMemRequest: '{{
.Values.telegraf_installer.telegraf_resources.rs_mem_request }}'
    # # telegraf additional tolerations. Use the following abbreviated
single line format only.
    # # Inspect telegraf-rs/-ds to view tolerations which are always
present.
    # # Example: '{key: taint1, operator: Exists, effect:
NoSchedule}, {key: taint2, operator: Exists, effect: NoExecute}'
    # dsTolerations: ''
    # rsTolerations: ''
    # # Set runPrivileged to true if SELinux is enabled on your Kubernetes
nodes.
    # runPrivileged: 'false'
    # # Collect NFS IO metrics.
    # dsNfsIOEnabled: '{{
.Values.telegraf installer.kubernetes.ds.shim nfs io processing }}'
    # # Collect kubernetes.system container metrics and objects in the
kube-system|cattle-system namespaces for managed kubernetes clusters (EKS,
AKS, GKE, managed Rancher). Set this to true if you want collect these
metrics.
    # managedK8sSystemMetricCollectionEnabled: '{{
.Values.telegraf installer.kubernetes.shim managed k8s system metric colle
ction }}'
```

```
# # Collect kubernetes.pod volume (pod ephemeral storage) metrics.
Set this to true if you want to collect these metrics.
    # podVolumeMetricCollectionEnabled: '{{
.Values.telegraf installer.kubernetes.shim pod volume metric collection
} } '
    # # Declare Rancher cluster as managed. Set this to true if your
Rancher cluster is managed as opposed to on-premise.
    # isManagedRancher: '{{
.Values.telegraf installer.kubernetes.is managed rancher }}'
  # kube-state-metrics:
    # # kube-state-metrics CPU/Mem limits and requests. By default, when
unset, kube-state-metrics has no CPU/Mem limits nor request.
    # cpuLimit:
    # memLimit:
    # cpuRequest:
    # memRequest:
    # # Comma-separated list of metrics to enable.
    # # See metric-allowlist in https://github.com/kubernetes/kube-state-
metrics/blob/main/docs/cli-arguments.md
    # resources:
'cronjobs, daemonsets, deployments, ingresses, jobs, namespaces, nodes, persisten
tvolumeclaims, persistentvolumes, pods, replicasets, resourcequotas, services, s
tatefulsets'
    # # Comma-separated list of Kubernetes label keys that will be used in
the resources' labels metric.
    # # See metric-labels-allowlist in https://github.com/kubernetes/kube-
state-metrics/blob/main/docs/cli-arguments.md
    # labels:
'cronjobs=[*], daemonsets=[*], deployments=[*], ingresses=[*], jobs=[*], namesp
aces=[*],nodes=[*],persistentvolumeclaims=[*],persistentvolumes=[*],pods=[
*], replicasets=[*], resourcequotas=[*], services=[*], statefulsets=[*]'
    # # kube-state-metrics additional tolerations. Use the following
abbreviated single line format only.
    # # No tolerations are applied by default
    # # Example: '{key: taint1, operator: Exists, effect:
NoSchedule}, {key: taint2, operator: Exists, effect: NoExecute}'
    # tolerations: ''
  # # Settings for the Events Log feature.
  # logs:
    # # If Fluent Bit should read new files from the head, not tail.
```

```
# # See Read from Head in
https://docs.fluentbit.io/manual/pipeline/inputs/tail
    # readFromHead: "true"
    # # Network protocol that Fluent Bit should use for DNS: "UDP" or
"TCP".
    # dnsMode: "UDP"
    # # Logs additional tolerations. Use the following abbreviated single
line format only.
    # # Inspect fluent-bit-ds to view tolerations which are always
present. No tolerations are applied by default for event-exporter.
    # # Example: '{key: taint1, operator: Exists, effect:
NoSchedule}, {key: taint2, operator: Exists, effect: NoExecute}'
    # fluent-bit-tolerations: ''
    # event-exporter-tolerations: ''
  # # Settings for the Network Performance and Map feature.
  # workload-map:
    # # net-observer-14-ds CPU/Mem limits and requests.
    # # See https://kubernetes.io/docs/concepts/configuration/manage-
resources-containers/
    # cpuLimit: '500m'
    # memLimit: '500Mi'
    # cpuRequest: '100m'
    # memRequest: '500Mi'
    # # Metric aggregation interval in seconds. Min=30, Max=120
    # metricAggregationInterval: '60'
    # # Interval for bpf polling. Min=3, Max=15
    # bpfPollInterval: '8'
    # # Enable performing reverse DNS lookups on observed IPs.
    # enableDNSLookup: 'true'
    # # net-observer-14-ds additional tolerations. Use the following
abbreviated single line format only.
    # # Inspect net-observer-14-ds to view tolerations which are always
present.
    # # Example: '{key: taint1, operator: Exists, effect:
NoSchedule}, {key: taint2, operator: Exists, effect: NoExecute}'
    # 14-tolerations: ''
    # # Set runPrivileged to true if SELinux is enabled on your Kubernetes
nodes.
```

```
# # Note: In OpenShift environments, this is set to true
automatically.
    # runPrivileged: 'false'
  # change-management:
    # # change-observer-watch-rs CPU/Mem limits and requests.
    # # See https://kubernetes.io/docs/concepts/configuration/manage-
resources-containers/
    # cpuLimit: '500m'
    # memLimit: '500Mi'
    # cpuRequest: '100m'
    # memRequest: '500Mi'
    # # Interval in minutes after which a non-successful deployment of a
workload will be marked as failed
    # failureDeclarationIntervalMins: '30'
    # # Frequency at which workload deployment in-progress events are sent
    # deployAggrIntervalSeconds: '300'
    # # Frequency at which non-workload deployments are combined and sent
    # nonWorkloadAggrIntervalSeconds: '15'
    # # A set of regular expressions used in env names and data maps whose
value will be redacted
    # termsToRedact: '"pwd", "password", "token", "apikey", "api-key",
"jwt"'
    # # A comma separated list of additional kinds to watch from the
default set of kinds watched by the collector
    # # Each kind will have to be prefixed by its apigroup
    # # Example: 'authorization.k8s.io.subjectaccessreviews'
    # additionalKindsToWatch: ''
    # # A comma separated list of kinds to ignore from watching from the
default set of kinds watched by the collector
    # # Each kind will have to be prefixed by its apigroup
    # # Example: 'networking.k8s.io.networkpolicies,batch.jobs'
    # kindsToIgnoreFromWatch: ''
    # # Frequency with which log records are sent to CI from the collector
    # logRecordAggrIntervalSeconds: '20'
    # # change-observer-watch-ds additional tolerations. Use the following
abbreviated single line format only.
    # # Inspect change-observer-watch-ds to view tolerations which are
always present.
```

```
# # Example: '{key: taint1, operator: Exists, effect:
NoSchedule}, {key: taint2, operator: Exists, effect: NoExecute}'
    # watch-tolerations: ''----
[[ID75e726adea6e42936168224f4e5d2b10]]
= Kubernetes Cluster Detail Page
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The Kubernetes cluster detail page displays a detailed overview of your
Kubernetes cluster.
//The detail page is comprised of three distinct but linked landing pages
showing cluster, node, and pod information. The "Resource Usage" section
changes to show the details of the selected item (cluster, node, or pod).
You can see the current page type and name at the top of the screen. The
current page is shown in the following heirarchy: Site Name / Kubernetes
/ Cluster / Node / Pod . You can click any part of this "breadcrumb"
trail to go directly to that specific page.
//image:Kubernetes Breadcrumb.png[Page Breadcrumb]
//== Cluster Overview
//The cluster overview page provides useful information at a glance:
image:Kubernetes Detail Page new.png[Cluster detail page]
== Namespace, Node, and Pod Counts
The counts at the top of the page show you the total number of namespaces,
nodes, and pods in the cluster, as well as the number of popds that are
currently alerting and pending.
//NOTE: It is possible that the three pod sub-counts (healthy, alerting,
```

pending) can add up to more than the displayed total number of pods. This can happen because the _pending_ count includes _all_ pending pods, both unscheduled and scheduled (in other words, unattached and attached to nodes).

== Shared Resources and Saturation

On the top right of the detail page is your cluster saturation as a current percentage as well as a graph showing the recent trend over time. Cluster saturation is the highest of CPU, memory, or storage saturation at each point in time.

Below that, the page shows by default *Shared Resources* usage, with tabs for CPU, Memory, and Storage. Each tab shows the saturation percentage and trend over time, with additional usage details. For storage, the value shown is the greater of backend and filesystem saturation, which are calculated independently.

The devices with the highest usage are shown in a table at the bottom. Click any link to explore these devices.

== Namespaces

The Namespaces tab displays a list of all the namespaces in your Kubernetes environment, showing CPU and Memory usage as well as a count of workloads in each namespace. Click the Name links to explore each namespace.

image:Kubernetes_Namespace_tab_new.png[list of current namespaces in your
K8s environment]

== Workloads

Similarly, the Workloads tab displays a list of the workloads in each namespace, again showing CPU and Memory usage. Clicking the Namespace links drills into each.

image:Kubernetes_Workloads_tab_new.png[list of current namespaces in your
K8s environment]

== The Cluster "Wheel"

image:Kubernetes Wheel Section.png[Cluster Wheel]

The Cluster "Wheel" section provides node and pod health at a glance, which you can drill into for more information. If your cluster contains more nodes than can be displayed in this area of the page, you will be able to turn the wheel using the buttons available.

Alerting pods or nodes are displayed in red. "Warning" areas are displayed in orange. Pods that are unscheduled (that is, unattached) will display in the lower corner of the Cluster "Wheel".

Hovering over a pod (circle) or Node (bar) will extend the view of the node.

image:Kubernetes_Node_Expand.png[Expanded Node]

Clicking on the pod or node in that view will zoom in to the expanded Node view.

image:Kubernetes Critical Pod Zoom.png[Expanded Node View]

From here, you can hover over an element to display details about that element. For example, hovering over the critical pod in this example displays details about that pod.

image:Kubernetes Pod Red.png[Critical Pod Information]

You can view Filesystem, Memory, and CPU information by hovering over the Node elements.

image: Kubernetes Capacity Info.png[Node Capacity and Allocation]

//// == Detail Section

Each page of the Kubernetes Explorer displays a set of usage graphs that may include CPU, Memory, and Storage. Below these graphs are summaries and lists of the top objects in each category, with links to underlying details. For example, _Node_ shows pods and containers, _Pod_ shows PVCs and related objects and containers, etc. The following illustration shows an example of the detailed information on a Node page:

image:Kubernetes Node Resource Usage.png[Resource Usage Example]

```
Resources experiencing alerts will show at the top of the lists. Click on
the affected resource to drill into it for more detail.
////
== A note about the gauges
The Memory and CPU gauges show three colors, since they show used in
relation to both allocatable capacity and total capacity .
////
Keep the following in mind when reading the gauges.
The dark blue band shows the amount used.
* When viewed against the light blue band , the dark blue shows used as
the % of allocatable amount. This matches the "% of allocatable" value
shown (81 in the example below).
* When viewed against the white background , the dark blue shows used as
the % of total capacity. This matches the "% of capacity" value shown (63
in this example).
image:Kubernetes Gauge Explained.png[Gauge Numbers Explained]
//The length of the light blue band against the white background shows the
total allocatable amount vs the total capacity; that figure itself is not
shown, but it's derived using the formula shown in the red text: (capacity
/ allocatable) * 100.
//Much of the time, our own environments show the same number for the 2
percent values in the gauge, and so you don't often see the white band
because the light blue covers it completely (meaning 100% of the total
capacity is allocatable).
////
[[ID3e544f26d29e5f01ad54475e63205d82]]
= Kubernetes Network Performance Monitoring and Map
```

```
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

The Kubernetes Network Performance Monitoring and Map feature simplifies troubleshooting by mapping dependencies between services (also called workloads), and provides real-time visibility into network performance latencies and anomalies to identify performance issues before they affect users.

This capability helps organizations reduce overall costs by analyzing and auditing Kubernetes traffic flows.

Key Features:

- The Workload Map presents Kubernetes workload dependencies and flows and highlights network and performance issues.
- Monitor network traffic between Kubernetes pods, workloads, and nodes; identifies the source of traffic and latency problems.
- Reduce overall costs by analyzing ingress, egress, cross-region, and cross-zone network traffic.

//image:Workload Map Example_withSlideout.png[Workload Map example showing
"Slideout" panel with details]

image:workload-map-animated.gif[Workload Map Example]

== Pre-Requisites

Before you can use the Kubernetes Network Performance Monitoring and Map, you must have configured the

xref:{relative_path}task_config_telegraf_agent_k8s.html[NetApp Kubernetes
Monitoring Operator] to enable this option. During deployment of the
Operator, select the "Network Performance and Map" checkbox to enable. You
can also enable this option by navigating to a Kubernetes landing page and
selecting "Modify Deployment".

image:ServiceMap_NKMO_Deployment_Options.png[selecting the Map option
during NKMO stup]

== Monitors

The Workload Map uses

xref:{relative_path}task_create_monitor.html[monitors] to derive
information. Cloud Insights provides a number of default Kubernetes
Monitors (note that these may be _Paused_ by default. You can _Resume_
(i.e. enable) the monitors you want), or you can create custom monitors
for kubernetes objects, which the Workload Map will also use.

You can create Cloud insights metric alerts on any of the object types below. Make sure the data is grouped by the default object type.

- * kubernetes.workload
- * kubernetes.daemonset
- * kubernetes.deployment
- * kubernetes.cronjob
- * kubernetes.job
- * kubernetes.replicaset
- * kubernetes.statefulset
- * kubernetes.pod
- * kubernetes.network_traffic_14

== The Map

The Map shows services/workloads and their relationships to each other. Arrows show directions of traffic. Hovering over a workload displays summary information for that workload, as you can see in this example:

image:ServiceMap Simple Example.png[Example of a Workload Map workload]

Icons within the circles represent different service types. Note that icons are only visible if the underlying objects have <<workload-labels, labels>>.

image:ServiceMap Icons.png[Service Icons Explained]

The size of each circle indicates node size. Note that these sizes are relative, your browser zoom level or screen size may affect actual circle sizes. In the same way, the traffic line style gives you an at-a-glance view of the connection size; bold solid lines are high traffic, while light dotted lines are lower traffic.

Numbers inside the circles are the number of external connections currently being processed by the service.

image:ServiceMap_Node_and_Connection_Legend.png[legend showing relative
circle (node) and connection sizes]

////

== Details

Hovering over a circle displays a summary of information for that service.

image:Workload_Map_Summary.png[Workload Hover Summary]
///

== Workload Details and Alerts

Circles displayed in color indicate a warning- or critical-level alert for the workload. Hover over the circle for a summary of the issue, or click on the circle to open a slideout panel with more detail.

image:Workload_Map_Slideout_with_Alert.png[Workload Slideout Details With
Alerts]

== Finding and Filtering

As with other Cloud Insights features, you can easily set filters to focus on the specific objects or workload attributes you want.

image:Workload Map Filtering.png[Workload Map filtering]

Likewise, typing a string in the _Find_ field will highlight matching workloads.

image:Workload_Map_Find_Highlighting.png[typing in find box highlights
workloads]

== Workload Labels

Workload labels are necessary if you want the Map to identify the types of workloads displayed (i.e. the circle icons). Labels are derived as follows:

- * Name of the service/application running in generic terms
- * If the source is a pod:
- ** Label is derived from the workload label of the pod
- ** Expected label on the workload: app.kubernetes.io/component
- ** Label name reference:

https://kubernetes.io/docs/concepts/overview/working-with-objects/common-

labels/ ** Recommended labels: *** frontend *** backend *** database *** cache *** queue *** kafka * If the source is external to the kubernetes cluster: ** Cloud Insights will attempt to parse the DNS resolved name to extract the service type. For example, with a DNS resolved name of s3.eu-north-1.amazonaws.com , the resolved name is parsed to get s3 as the service type. == Dive Deep Right-clicking on a workload presents you with additional options to explore further. For example, from here you can zoom in to view the connections for that workload. image: Workload Map Zoom Into Connections.png [Workload Map Right-Click Zoom to show the workload's connections] Or you can open the detail slideout panel to directly view the _Summary_, Network , or Pod & Storage tab. image: Workload Map Detail Network Slideout.png[Detail Slideout Network Tab Example] Finally, selecting Go to Asset Page will open the detailed asset landing

page for the workload.

image:Workload Map Asset Page.png[Workload Asset Page]

```
[[IDf6a45412d8c403ca6ce47ee2a1e88b2e]]
= Kubernetes Change Analytics
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

[.lead]

Kubernetes Change Analytics provide you with an all-in-one view of recent changes to your K8s environment. Alerts and deployment status are at your fingertips. With Change Analytics, you can track every deployment and configuration change, and correlate it with the health and performance of K8s services, infrastructure, and clusters.

Keep the following in mind:

- * In multi-tenant environments, outages may happen because of misconfigured changes. In very dynamic environments, Cloud Insights Change Analytics may not be able to properly track all changes.
- * Change Analytics provides a single pane to view and correlate the health of workloads and configuration changes. This may help in troubleshooting dynamic environments.

To view Kubernetes Change Analytics, navigate to *Kubernetes > Change Analysis*.

image:ChangeAnalytitcs_Main_Screen.png[Kubernetes Change Analytics main screen, showing warning and critical alerts, successful and failed deployments, and the top 3 workloads triggering alerts].

The page automatically refreshes based on the currently-selected Cloud Insights time range. Smaller time ranges mean more frequent screen refreshing.

== Filtering

As with all features of Cloud Insights, filtering the change list is intuitive: at the top of the page, enter or select values for your

Kubernetes Cluster, Namespace, or Workload, or add your own filters by selecting the {+] button.

When you filter down to a specific Cluster, Namespace, and Workload (along with any other filters you set), you are shown a timeline of deployments and alerts for that workload in that namespace on that cluster. Zoom in further by clicking and dragging in the graph to focus on a more specific time range.

image:ChangeAnalytitcs Filtered Timeline.png[Workload Timeline example]

== Quick Status

Below the filtering area are a number of high-level indicators. On the left is the number of alerts (Warning and Critical). This number includes _Active_ as well as _Resolved_ alerts. To see only _Active_ alerts, set a filter for "Status" and choose "Active".

image:ChangeAnalytitcs Alerts.png[Change Analytics Alerts]

Deployment status is also shown here. Again, the default is to show the count of _Started_, _Complete_, and _Failed_ deployments. To see only _Failed_ deployments, set a filter for "Status" and select "Failed".

image:ChangeAnalytitcs Deploys.png[Change Analytics DSeploys]

The top 3 workloads with the most alerts are next. The number in red next to each workload indicates the number of alerts related to that workload. Click the workload link to explore through your Infrastructure (Kubernetes Explorer), Dependencies (Workload Map), or Log Analysis (Event Logs).

image:ChangeAnalytitcs_ExploreWorkloadAlerts.png[Change Analytics Workload
Exploration Options]

== Detail Panel

Selecting a change in the list opens a panel describing the change in more detail. For example, selecting a failed Deploy shows a summary of the Deploy, with start and end times, duration, and where the deploy was triggered, with links to explore those resources. It also displays the reason for the failure, any related changes, and any associated events.

image:ChangeAnalytitcs DeployDetailPanel.png[Deploy Failure Detail Panel]

Selecting an Alert similarly provides details about the alert, including the monitor that triggered the alert as well as a chart showing a visual timeline for the alert.

```
:leveloffset: -1
= Insights
:leveloffset: +1
[[ID332fdfcd2c88695bed2cdeb02c5b3edd]]
= Insights
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Insights allow you to look into things like resource usage and how it
affects other resources, or time-to-full analyses.
A number of Insights are available. Navigate to *Dashboards > Insights* to
start diving in. You can view active Insights (Insights that are
currently occurring) on the main tab, or inactive Insights on the
Inactive Insights tab. Inactive Insights are those that were previously
active but are no longer occurring.
== Insight Types
=== Shared Resources Under Stress
High-impact workloads can reduce the performance of other workloads in a
shared resource. This puts the shared resource under stress. Cloud
Insights provides tools to help you investigate resource saturation and
impact in your environment.
xref:{relative path}insights shared resources under stress.html[Learn
Morel
//NOTE: This is a Preview feature and may change over time as
improvements are made.
xref:{relative path}/concept preview features.html[Learn more] about Cloud
Insights Preview features.
```

=== Kubernetes Namespaces Running Out of Space The Kubernetes Namespaces Running Out of Space Insight gives you a view into workloads on your Kubernetes namespaces that are at risk of running out of space, with an estimate for the number of days remaining before each space becomes full. xref:{relative path}insights k8s namespaces running out of space.html[Lear n Morel //NOTE: This is a Preview feature and may change over time as improvements are made. xref:{relative path}/concept preview features.html[Learn more] about Cloud Insights Preview features. === Reclaim ONTAP Cold Storage The Reclaim ONTAP Cold Storage Insight provides data about cold capacity, potential cost/power savings and recommended action items for volumes on ONTAP systems. xref:{relative path}insights reclaim ontap cold storage.html[Learn More] NOTE: This is a Preview feature and may change over time as improvements are made. xref:{relative path}/concept preview features.html[Learn more] about Cloud Insights Preview features. [[ID1fa1e3a0a4bd412ee2c99aeeb95ba1f8]] = Insights: Shared Resources Under Stress :toc: macro :hardbreaks: :toclevels: 2 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ High-impact workloads can reduce the performance of other workloads in a shared resource. This puts the shared resource under stress. Cloud Insights provides tools to help you investigate resource saturation and impact in your environment. //// NOTE: The Shared Resources Under Stress Insight is considered a

Preview feature and is therefore subject to change.

== Terminology

When talking about workload or resource impact, the following definitions are useful.

//A *Demanding* or *Impactful* resource is one that causes a negative impact on another resource. For example, a volume experiencing very high IOPS may cause increased latency in other volumes (i.e. _Impacted_ or _Degraded_ resources). Demanding and Impacted resources are "peers" that utilize the same shared resource, for example, a storage pool or Volume. Demanding resources are sometimes called greedy resources.

A *Demanding Workload* is a workload that is currently identified as impacting other resources in the shared storage pool. These workloads drive higher IOPS (for example), reducing IOPS in the Impacted Workloads. Demanding workloads are sometimes called high-consuming workloads.

An *Impacted Workload* is a workload that is affected by a high-consuming workload in the shared Storage Pool. These workloads are experiencing reduced IOPS and/or higher latency, caused by the Demanding Workloads.

Note that if Cloud Insights has not discovered the leading compute workload, the volume or internal volume itself will be recognized as the workload. This applies to both demanding and impacted workloads.

Shared Resource Saturation is the ratio of impacting IOPS to baseline .

Baseline is defined as the maximum reported data point for each workload in the hour immediately preceding the detected saturation.

A *Contention* or *Saturation* occurs when IOPS are determined to be affecting other resources or workloads in the shared storage pool.

== Demanding Workloads

To start looking into Demanding and impacted workloads in your shared resources, click on *Dashboards > Insights* and select the *Shared Resources Under Stress* Insight.

//image:Shared_resources_Under_Stress_menu.png[Workloads List]
image:InsightsMenu.png[Insights Menu]

Cloud Insights displays a list of any workloads where a saturation has

been detected. Note that Cloud Insights will show workloads where at least one demanding resource *or* impacted resource has been detected.

Click on a workload to view the details page for it. The top chart shows the activity on the shared resource (for example, a storage pool) on which the contention/saturation is occurring.

//image:Shared_resources_Under_Stress_SharedResource.png[Shared resource
showing contention]

image:ResourceInsightShared.png[Shared resource showing contention]

//image:Insights_Shared_Resource_Contention_Chart.png[Chart showing
activity on the shared resource]

Below that are two charts showing the _demanding_ workloads and the workloads that are impacted by those demanding workloads.

//image:Insights_Demanding_Workload_Chart.png[Demanding workload chart]
//image:Insights_Impacted_Workload_Chart.png[Impacted workload chart]
image:ResourceInsightDemanding.png[Demanding workload chart]
image:ResourceInsightImpacted-a.png[Impacted workload chart]

Below each table is a list of workloads and/or resources affecting or affected by the contention. Clicking on a resource (for example, a VM) opens a detail page for that resource. Clicking on a workload opens a query page showing the pods involved. Note that if the link opens an empty query, it may be because the affected pod is no longer part of the active contention. You can modify the query's time range to view the pod list in greater or more focused time range.

== What do I do to resolve saturation?

There are a number of steps you can take to reduce or eliminate the chance of saturation in your environment. These are shown by expanding the *+Show Recommendations* link on the page. Here are a few things you can try.

* Move high-IOPS consumers

+

Move the "greedy" workloads to less-saturated Storage Pools. It is recommended to assess the tier and capacity of these pools before moving the workloads, to avoid unnecessary costs or additional contentions.

* Implement a quality of service (QoS) policy

+

Implementing a QoS policy per workload to ensure enough free resources available will alleviate saturation on the Storage Pool. This is a long-

term solution. * Add additional resources If the shared resource (for example, Storage Pool) has reached the IOPS saturation point, adding more or faster disks to the pool will ensure enough free resources available to alleviate saturation. Finally, you can click the *Copy Insight Link* to copy the page url to the clipboard, to more easily share with colleagues. [[ID87eb09797f799b1d82d7901d8e469fbb]] = Insights: Kubernetes Namespaces Running out of Space :toc: macro :hardbreaks: :toclevels: 2 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead] Running out of space in your environment is never a good situation. Cloud Insights helps you predict the time you have before Kubernetes persistent volumes become full. //NOTE: This is a Preview feature and may change over time as improvements are made. xref:{relative path}/concept preview features.html[Learn more] about Cloud Insights Preview features. The Kubernetes Namespaces Running Out of Space Insight gives you a view into workloads on your Kubernetes namespaces that are at risk of running out of space, with an estimate for the number of days remaining before each persistent volume becomes full. You can view this Insight by navigating to *Dashboards > Insights*. image:K8sRunningOutOfSpaceWorkloadList.png[List of workloads in K8s namespaces that are at risk of running out of space] Click on a workload to open a detail page for the Insight. On this page you will see a graph showing the workload capacity trends as well as a

table showing the following: * Workload Name * Persistent Volume affected * Predicted Time-to-Full in days * Persistent Volume capacity * Backend Storage Resource affected, with current capacity used out of total capacity. Clicking this link will opoen the detailed landing page for the backend volume. image: K8sRunningOutOfSpaceWorkloadTable.png[Workload table showing detailsl === What can I do if I'm running out of space? On the Insight page, click the *+Show Recommendations* to view possible solutions. The easiest option when running out of space is always to add more capacity, and Cloud Insights shows you the optimal capacity to add to increase time-to-full to a target 60-day prediction. Other recommendations are also shown. image: K8sRunningOutOfSpaceRecommendations.png[Capacity to add to return to 60-day TTF] It is here also that you can copy a convenient link to this Insight, to bookmark the page or to easily share with your team. [[ID4b1b00862510e153bdb91c4c56bbad22]] = Insights: Reclaim ONTAP Cold Storage :toc: macro :hardbreaks: :toclevels: 2 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/

[.lead]

The _Reclaim ONTAP Cold Storage_ Insight provides data about cold capacity, potential cost/power savings and recommended action items for volumes on ONTAP systems.

To view these Insights, navigate to *Dashboards > Insights* and take a look at the _Reclaim ONTAP Cold Storage_ Insight. Note that this Insight will only list affected storages if Cloud Insights has detected cold storage, otherwise you will see an "all clear" message.

Keep in mind that cold data less than 30 days old is not shown.

image:Cold_Data_Insight_List.png[Insight list showing 3 storages with cold
data]

The Insight description gives a quick indication of the amount of data detected as "cold" and which storage that data resides on. The table also provides a count of workloads with cold data.

Selecting an Insight from the list opens a page showing more details, including recommendations to move data to the Cloud or cycle down unised disks, as well as estimated cost and power savings you could potentially realize from implementing those recommendations. The page even provides a handy link to link:https://bluexp.netapp.com/cloud-tiering-service-tco[NetApp's TCO Calculator] so you can experiment with the numbers.

image:Cold Data Power Info.png[Cold Data Power Info]

=== Recommendations

On the Insight page, expand the *Recommendations* to explore the following options:

 * Move unused workloads (zombies) to a lower cost storage tier (HDD) $_{\cdot}$

Utilizing the zombie flag, cold storage and number of days, find the coldest and largest amount of data and move the workload to a lower cost storage tier (such as a storage pool using hard disk storage). A workload is considered a "zombie" when is has not received any significant IO requests for 30 days or more.

* Delete unused workloads

+

Verify which workloads are not in use and consider archiving them or remove them from the storage system.

* Consider NetApp's Fabric Pool Solution

+

NetApp's link:https://docs.netapp.com/us-en/cloud-manager-tiering/concept-cloud-tiering.html#features[Fabric Pool Solution] automatically tiers cold data to low cost cloud storage, thus increasing the efficiency of your

```
performance tier as well as providing remote data protection.
=== Visualize and Explore
The graphs and table provide additional trending information as well as
allow you to drill into the individual workloads.
//image:Cold Data Workload Graph and Table.png[Cold Data Workload Graph an
d Table]
image:Cold Data Storage Trend.png[Cold Data Summary Graphs]
image:Cold Data Workload Table.png[Cold Data Workload Table]
:leveloffset: -1
[[ID0b84cb710989b4c928a01d0cbe64fe1a]]
= ONTAP Essentials
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
ONTAP Essentials is a set of dashboards and workflows that provide
detailed overviews of your ONTAP inventories and workloads. You may see
the following terms used when working in ONTAP Essentials:
* Infrastructure/Inventory: Objects that provide storage/networking
services to user data
* Workloads: Objects that provide interface to users to read/write data.
* Data Protection: Objects that can be protected using NetApp data
protection technologies
For additional terms and definitions related to ONTAP, see the
xref:{relative path}task dc na cdot.html[ONTAP Data Collector]
documentation.
ONTAP Essentials requires at least one working ONTAP data collector with
```

data collected within the last seven days.

== Overview

To begin exploring, select *ONTAP Essentials* from the main Cloud Insights menu.

//image:ONTAP_Essentials_Overview_Sept.png[Overview dashboard for ONTAP
Essentials]

image:OE Overview.png[Overview dashboard for ONTAP Essentials]

The *Overview* dashboard displays useful information like the number of clusters in your environment with their overall capacity and performance percentages. You will also see predictive data regarding the number of expected days until storage capacity or performance capacity runs out of space. Additionally, if any controllers in your infrastructure are running with their CPU at more than 65%--potentially putting your cluster at risk in case of failover--ONTAP Essentials shows those as "Hot" controllers.

Informative graphs give you a look into performance over time as well as breakdowns of capacity usage. Each of these graphs or data points can be used as a starting point for exploration or investigation.

Note: A "days to full" number of "0" (zero) indicates that days to full is estimated at greater than 90 days. In other words, your systems aren't in danger of running out of space any time soon.

== Data Protection

//Select the *Data Protection* page to view SnapMirror relationships. Click through to source or destination volume information, or click the gear icon to add columns for the data you wish to view.

The *Data Protection* page shows the status of volumes protected by *Snapshot copies* or *SnapMirror policies*.

In the _Local Protection Overview_ section, the charts provide the following information for volumes protected by Snapshot copies:

- * The number of volumes protected by Snapshot copies, as well as those not protected.
- * The number of volumes that are using or exceeding the reserve space for the Snapshot copies.
- * The number of volumes in specific ranges of Snapshot copy count (i.e. less than 10 copies, 10 to 200, etc.).

In the _Remote Protection Overview_ section, the charts provide information related to volumes protected by SnapMirror policies:

- * The number of healthy and unhealthy SnapMirror relationships.
- * The number of SnapMirror relationships experiencing recovery point objective (RPO) lag based on the lag status.
- * The number of relationships protected by SnapMirror volume protection types such as Volume SnapMirror, SVMDR relationships, FlexGroup SnapMirror relationships, SnapMirror Business Continuity (SMBC) consistency Group relationships, as well as unprotected volumes.
- * The number of relationships protected by the SnapMirror relationship types such as Asynchronous Mirror, Asynchronous Vault, Asynchronous MirrorVault, StrictSync, and Sync.

image:DataProtectionDashboard_OverviewWidgets_.png[Data Protection
Dashboard widgets showing Local and Remote overviews]

The _Clusters_ grid at the bottom of the page provides details related to the following:

- * Volumes not protected by Snapshots.
- * Volumes breaching snapshot reserve space.
- * Volumes not protected by snapmirror policies and snapmirror relationships experiencing lag.
- * Unhealthy SnapMirror relationships.

//image:ONTAP_Essentials_data_protection.png[SnapMirror Relationships list
page]

//image:OE DataProtection.png[SnapMirror Relationships list page]

image:DataProtectionDashboard_ClusterList.png[Data Protection Cluster List]

== Security

The Security Dashboard gives you an instant view of your current security situation, showing charts for hardware and software volume encryption, anti-ransomware status, and cluster authentication methods. Security criteria is evaluated based on recommendations defined in the link:https://www.netapp.com/pdf.html?item=/media/10674-tr4569.pdf[*NetApp Security Hardening Guide for ONTAP 9*].

Select any of the encryption or anti-ransomware counts to dive into your environment.

```
image:OE SecurityDashboard.png[ONTAP Essentials Security Dashboard]
The ONTAP Essentials Security dashboard monitors your environment to
determine cluster compliance status. Refer to the
link:https://docs.netapp.com/us-en/active-iq-unified-manager/health-
checker/reference cluster compliance categories.html[Cluster Compliance
Categories] to learn more. ONTAP Essentials uses the following monitors to
determine compliance:
|===
|Monitor Name |Attribute Name (Displayed in Cluster Details) |Attribute
Compliant Value
|FIPS Mode Disabled |FIPS mode |Enabled
|Cluster Insecure ciphers for SSH | |Secure SSH Settings | |Yes
|Telnet Protocol Enabled |Telnet |Disabled
|Remote Shell Enabled |Remote Shell |Disabled
|Default Local Admin User Enabled|Default Admin User |Disabled
|MD5 Hashed password |MD5 in use |No
|Cluster Peer Communication Not Encrypted |Cluster Peering
|Encrypted/ No Peer
|AutoSupport HTTPS Transport Disabled |AutoSupport using HTTPS
|No NTP Servers are Configured |Network Time Protocol |Configured
|NTP Server Count is Low | Network Time Protocol | Configured
|Cluster Login Banner Disabled |Login Banner |Enabled
|Log Forwarding Not Encrypted |Log Forwarding Encrypted |Yes
|===
Note that if a monitor above is disabled, the cluster details will show
the value as 'Not checked' for the corresponding security compliance
attribute.
image:OE Cluster Compliance Example.png[Cluster Compliance Status]
For SVMs, the Security dashboard looks at the following monitors:
|===
|Monitor Name | Attribute Name (Displayed in Storage VM Settings)
|Attribute Compliant Value
|Storage VM Insecure ciphers for SSH | |Secure SSH Settings | |Yes
|Storage VM Login banner disabled |Login Banner |Enabled
|Storage VM Audit Log Disabled |Audit Log |Enabled
|===
```

```
//The following link:https://docs.netapp.com/us-en/active-iq-unified-
manager/health-
checker/reference svm compliance categories.html[parameters] also factor
in determining SVM security status:
//image:OE SVM Parameters.png[SVM Parameters for compliance]
In the cluster list, select View Details for each cluster to open a
"slideout" panel showing you the current settings for Cluster, Storage
VM, or Anti-Ransomware.
Cluster details include connection status, certificate information, and
image:OE Cluster Slideout.png[Cluster Detail Slideout Panel]
Storage VM details show audit and SSH information:
image:OE Storage Slideout.png[Storage tab]
Anti-Ransomware details show whether a storage VM is protected by ONTAP's
Anti-Ransomware Protection or Cloud Insights Workload Security. Note that
the ONTAP ARP column displays the current status of ONTAP's on-board Anti-
Ransomware Protection, which is configured on the ONTAP system. Cloud
Insights Workload Security can be enabled by selecting "Protect" in that
column.
image:OE Anti-Ransomware Slideout.png[Anti-Ransomware tab]
== Alerts
Here you can view the Active alerts in your environment and quickly drill
down into potential problems. Select the Resolved tab to view alerts
that have been resolved.
//image:ONTAP Essentials Alerts Menu.png[ONTAP Alerts Menu]
//image:ONTAP Essentials Alerts Page.png[ONTAP Alerts page example showing
active alerts]
image:OE Alerts.png[ONTAP Essentials Alerts List]
== Infrastructure
The ONTAP Essentials *Infrastructure* page gives you a view of cluster
```

health and performance, using pre-built (yet further customizable) queries on all the basic ONTAP objects. Select the object type you wish to explore (cluster, storage pool, etc.) and choose whether to view health or performance information. Set filters to dive deeper into individual systems. image: ONTAP Essentials Health Performance.png[Infrastructure selections for storage pools] Infrastructure page showing cluster health: image: ONTAP Essentials Infrastructure A.png[Infrastructure objects to explore] == Networking ONTAP Essentials Networking gives you views into your FC, NVME FC, Ethernet, and iSCSI infrastructure. On these pages you can explore things like ports in your clusters and their nodes. image:ONTAP Essentials Alerts Menu.png[ONTAP Essentials Networking Menu] image:ONTAP Essentials Alerts Page.png[ONTAP Essentials Networking FC page showing ports into cluster nodes] == Workloads View and explore workloads on LUNs/Volumes, NFS or SMB Shares, or Qtrees in your environment. image:ONTAP Essentials Workloads Menu.png[Workloads Menu] image:ONTAP Essentials Workloads Page.png[Workloads list page] = Working with Queries :leveloffset: +1 [[ID0c53f1ced857101edad3a284f43eda6d]] = Assets used in queries :toc: macro :hardbreaks: :toclevels: 1 :icons: font

```
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

[.lead]

Queries enable you to monitor and troubleshoot your network by searching the assets and metrics in your environment at a granular level based on user-selected criteria (for example, annotations).

Note that annotation rules, which automatically assign annotations to assets, _require_ a query.

You can query the physical or virtual inventory assets (and their associated metrics) in your environment, or the metrics provided with integration such as Kubernetes or ONTAP Advanced Data.

== Inventory Assets

The following asset types can be used in queries, dashboard widgets, and custom asset landing pages. The fields and counters available for filters, expressions, and display will vary among asset types. Not all assets can be used in all widget types.

- * Application
- * Datastore
- * Disk
- * Fabric
- * Generic Device
- * Host
- * Internal Volume
- * iSCSI Session
- * iSCSI Network Portal
- * Path
- * Port
- * Qtree
- * Quota
- * Share
- * Storage
- * Storage Node
- * Storage Pool
- * Storage Virtual Machine (SVM)
- * Switch
- * Tape
- * VMDK
- * Virtual Machine
- * Volume

- * Zone
- * Zone Member
- == Integration Metrics

In addition to querying for inventory assets and their associated performance metrics, you can query for *integration data* metrics as well, such as those generated by Kubernetes or Docker, or provided with ONTAP Advanced Metrics.

image:QueryPageFilter.png[Integration Query Filter Example]

```
[[IDfea8d8052c340d37b5a245397923881b]]
= Creating Queries
:toc: macro
:hardbreaks:
```

:toclevels: 1
:icons: font
:linkattrs:

:relative_path: ./

:imagesdir: {root path}{relative path}./media/

[.lead]

Queries enable you to search the assets in your environment at a granular level, allowing to filter for the data you want and sort the results to your liking.

For example, you can create a query for _volumes_, add a filter to find particular _storages_ associated with the selected volumes, add another filter to find a particular _annotation_ such as "Tier 1" on the selected storages, and finally add another filter to find all storages with _IOPS - Read (IO/s)_ greater than 25. When the results are displayed, you can then sort the columns of information associated with the query in ascending or descending order.

Note: When a new data collector is added which acquires assets, or any annotation or application assignments are made, you can query for those new assets, annotations, or applications only after the queries are indexed. Indexing occurs at a regularly scheduled interval or during certain events such as running annotation rules.

.Creating a Query is very simple:

. Navigate to *Queries > *+New Query*.

. From the 'Select...' list, select the object type you want to query for. You can scroll through the list or you can start typing to more quickly find what you're searching for.

.Scroll list:

image:QueryDrop-DownList.png[Query Drop-Down]

.Type-to-Search:

image:QueryPageFilter.png[Typing to find]

You can add filters to further narrow down your query by clicking the *+* button in the *Filter By* field.

Group rows by object or attribute. When working with integration data (Kubernetes, ONTAP Advanced Metrics, etc.), you can group by multiple attributes, if desired.

image:QueryFilterExample.png[Query Filtering and Grouping]

The query results list shows a number of default columns, depending on the object type searched for. To add, remove, or change the columns, click the gear icon on the right of the table. The available columns variy based on the asset/metric type.

image:QuerySelectColumns.png[Select Columns]

== Choosing Aggregation, Units, Conditional Formatting

=== Aggregation and Units

For "value" columns, you can further refine your query results by choosing how the displayed values are aggregated as well as selecting the units in which those values are displayed. These options are found by selecting the "three dots" menu at the top corner of a column.

image:Query_Page_Aggregation_etc.png[Query Page results showing
Aggregation, Conditional Formatting, Unit Display, and Column Renaming]

==== Units

You can select the units in which to display the values. For example, if the selected column shows raw capacity and the values are shown in GiB, but you prefer to display them as TiB, simply select TiB from the Unit Display drop-down.

==== Aggregation

By the same token, if the values shown are aggregated from the underlying data as "Average", but you prefer to show the sum of all values, select "Sum" from either the _Group by_ drop-down (if you want any grouped values to show the sums) or from the _Time Aggregate By_ drop-down (if you want the row values to show sums of underlying data).

You can choose to aggregate grouped data points by $_{\rm Avg}$, ${\rm Max}$, ${\rm Min}$, or ${\rm Sum}$ _.

You can aggregate individual row data by $_$ Average, Last data point acquired, Maximum, Minimum, or Sum .

==== Conditional Formatting

Conditional Formatting allows you to highlight Warning-level and Critical-level thresholds in the query results list, bringing instant visibility to outliers and exceptional data points.

image:Query_Page_Conditional_Formatting.png[Conditional Formatting on Query Page]

Conditional formatting is set separately for each column. For example, you can choose one set of thresholds for a capacity column, and another set for a throughput column.

==== Rename Column

Renaming a column changes the displayed name on the Query results list. The new column name is also shown in the resulting file if you export the query list to .CSV.

=== Save

After you have configured your query to show you the results you want, you can click the *Save* button to save the query for future use. Give it a meaningful and unique name.

== More on Filtering

=== Wildcards and Expressions

When you are filtering for text or list values in queries or dashboard widgets, as you begin typing you are presented with the option to create a *wildcard filter* based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create

expressions using NOT or OR, or you can select the "None" option to filter for null values in the field. image:Type-Ahead-Example-ingest.png[Wildcard Filter] Filters based on wildcards or expressions (e.g. NOT, OR, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue. image:Type-Ahead-Example-Wildcard-DirectSelect.png[Wildcard Filter Results] Note that Wildcard and Expression filtering works with text or lists but not with numerics, dates or booleans. === Refining Filters You can use the following to refine your filter: |Filter|What it does | Example | Result | * (Asterisk) | enables you to search for everything | vol*rhel | returns all resources that start with "vol" and end with "rhel" | ? (question mark) | enables you to search for a specific number of characters| BOS-PRD??-S12 | returns BOS-PRD** 12 **-S12, BOS-PRD** 23 **-S12, and so on | OR |enables you to specify multiple entities | FAS2240 OR CX600 OR FAS3270 |returns any of FAS2440, CX600, or FAS3270 | NOT |allows you to exclude text from the search results | NOT EMC* |returns everything that does not start with "EMC" | None |searches for NULL values in all fields | None |returns results where the target field is empty | Not * |searches for NULL values in text-only fields | Not * |returns results where the target field is empty |=== If you enclose a filter string in double quotes, Insight treats everything between the first and last quote as an exact match. Any special characters or operators inside the quotes will be treated as literals. For example, filtering for "*" will return results that are a literal asterisk; the asterisk will not be treated as a wildcard in this case. The operators OR and NOT will also be treated as literal strings when enclosed in double

//The operator AND is always trated as a literal string.

quotes.

== What do I do now that I have query results? Querying provides a simple place to add annotations or assign applications to assets. Note that you can only assign applications or annotations to your inventory assets (Disk, Storage, etc.). Integration metrics cannot

take on annotation or application assignments.

To assign an annotation or application to the assets resulting from your query, sinply select the asset(s) using the check box column on the left of the results table, then click the *Bulk Actions* button on the right. Choose the desired action to apply to the selected assets.

image:QueryVolumeBulkActions.png[Query Bulk Actions Example]

== Annotation Rules require query

If you are configuring

xref:{relative_path}task_create_annotation_rules.html[Annotation Rules],
each rule must have an underlying query to work with. But as you've seen
above, queries can be made as broad or as narrow as you need.

```
[[IDe32272ad6699297a7f90f0ee04ac4892]]
```

= Viewing queries

:toc: macro

:hardbreaks:

:toclevels: 1

:icons: font

:linkattrs:

:relative path: ./

:imagesdir: {root path}{relative path}./media/

[.lead]

You can view your queries to monitor your assets and change how your queries display the data related to your assets.

.Steps

- . Log in to your Cloud Insights tenant.
- . Click *Queries* and select *Show all queries*.

You can change how queries display by doing any of the following:

- . You can enter text in the filter box to search to display specific queries.
- . You can change the sort order of the columns in the table of queries to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header.

```
. To resize a column, hover the mouse over the column header until a blue
bar appears. Place the mouse over the bar and drag it right or left.
. To move a column, click on the column header and drag it right or left.
When scrolling through the query results, be aware that the results may
change as Cloud Insights automatically polls your data collectors. This
may result in some items being missing, or some items appearing out of
order depending on how they are sorted.
[[ID21e78ac9b3d9e5e924e4d396f7c471d1]]
= Exporting query results to a .CSV file
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
You can export the results of any query to a .CSV file, which will allow
you to analyze the data or import it into another application.
.Steps
. Log in to Cloud Insights.
. Click *Queries* and select *Show all queries*.
The Queries page is displayed.
. Click a query.
. Click image: ExportButton.png[Export Button] to export the query results
to a .CSV file.
NOTE: Export to .CSV is also available in the "three dots" menu in
dashboard table widgets as well as most landing page tables.
////
. When prompted, do one of the following:
.. Click *Open with* and then *OK* to open the file with Microsoft Excel
and save the file to a specific location.
.. Click *Save file* and then *OK* to save the file to your Downloads
folder.
////
//All of the attributes for the objects in the columns currently selected
```

for display are exported to the file, regardless of whether those attributes are being displayed. The exported data will reflect the current filtering, columns, and column names displayed. Note: When a comma appears in an asset name, the export encloses the name in quotes, preserving the asset name and the proper .csv format. When opening an exported .CSV file with Excel, if you have an object name or other field that is in the format NN:NN (two digits followed by a colon followed by two more digits), Excel will sometimes interpret that name as a Time format, instead of Text format. This can result in Excel displaying incorrect values in those columns. For example, an object named "81:45" would show in Excel as "81:45:00". To work around this, import the .CSV into Excel using the following steps: . Open a new sheet in Excel. . On the "Data" tab, choose "From Text". . Locate the desired .CSV file and click "Import". . In the Import wizard, choose "Delimited" and click Next. . Choose "Comma" for the delimiter and click Next. . Select the desired columns and choose "Text" for the column data format. . Click Finish. Your objects should show in Excel in the proper format.

= Modifying or Deleting a Query
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]]
You can change the criteria that are associated with a query when you want to change the search criteria for the assets that you are querying.
== Modifying a Query

[[ID5e7b91b2ad4dd1782c9a4e7a2be0f69b]]

.Steps

. Click *Queries* and select *Show all queries*. The Queries page is displayed. . Click the query name . To add a criteria to the query, click image: GearIcon.png[Columns] and select a criteria from the list. . To remove a filter from the query, click the ${}^{\star}X^{\star}$ next to the filter to remove. When you have made all necessary changes, do one of the following: * Click the *Save* button to save the query with the name that was used initially. * Click the drop-down next to the *Save* button and select *Save As* to save the query with another name. This does not overwrite the original * Click the drop-down next to the *Save* button and select *Rename* to change the query name that you had used initially. This overwrites the original query. * Click the drop-down next to the *Save* button and select *Discard Changes* to revert the query back to the last saved changes. == Deleting a Query To delete a query, click *Queries* and select *Show all queries*, and do one of the following: . Click on the "three dot" menu to the right of the query and click *Delete*. . Click on the query name and select *Delete* from the *Save* drop-down menu. [[IDbde2c2467525d4876b7fe6a434cede4c]] = Assigning multiple applications to or removing multiple applications from assets :toc: macro :hardbreaks: :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/

[.lead] You can assign multiple xref:{relative_path}task_create_application.html[applications] to or remove multiple applications from assets by using a query instead of having to manually assign or remove them. NOTE: You can use these steps to add or remove xref:{relative_path}task_defining_annotations.html[annotations] in the same way.

.Before you begin

You must have already created a query that finds all the assets that you to edit.

.Steps

. Click *Queries* and select *Show all queries*.

+

The Queries page displays.

. Click the name of the query that finds the assets.

+

The list of assets associated with the query displays.

. Select the desired assets in the list or click the top checkbox to select $\mbox{All}\,.$

+

The image:BulkActions.png[Actions] button displays.

- . To add an application to the selected assets, click image:BulkActions.png[Actions] and select *Add Application*.
- . Select one or more applications.

+

You can select multiple applications for hosts, internal volumes, qtrees, and virtual machines; however, you can select only one application for a volume or a share.

- . Click *Save*.
- . To remove an application assigned to the assets, click image:BulkActions.png[Actions] and select *Remove Application*.
- . Select the application or applications you want to remove.
- . Click *Delete*.

Any new applications you assign override any applications on the asset that were derived from another asset. For example, volumes inherit applications from hosts, and when new applications are assigned to a volume, the new application takes precedence over the derived application.

After you click Save on a bulk add or Remove on a bulk delete action,

Cloud Insights informs you that the action will take some time. You can dismiss this message; the action will continue in the background.

NOTE: For environments with large amounts of related assets, inheritance of application assignments to those assets could take several minutes. Please allow more time for inheritance to occur if you have many related assets.

```
[[IDb0cb004c9043028f353db8247c833d88]]
= Copying table values
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

You can copy values in tables to the clipboard for use in search boxes or other applications.

.About this task

There are two methods you can use to copy values from tables or query results to the clipboard.

.Steps

- . Method 1: Highlight the desired text with the mouse, copy it, and paste it into search fields or other applications.
- . Method 2: For single-value fields, hover over the field and click the clipboard icon image:ClipboardIcon.png[Clipboard] that appears. The value is copied to the clipboard for use in search fields or other applications.

Note that only values that are links to assets can be copied using this method. Only fields that include single values (i.e. non-lists) have the copy icon.

```
[[ID14ed8bbd70e441dd198bf3982b3dce5a]]
= Log Explorer
:toc: macro
```

```
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The Cloud Insights Log Explorer is a powerful tool for querying system
logs. In addition to helping with investigations, you can also save a log
query in a Monitor to provide alerts when those particular log triggers
are activated.
To begin exploring logs, click *Log Queries > +New Log Query*.
//image:LogExplorerMenu.png[Log queries menu, 480]
Select an available log from the list.
//This list may vary based on your current Cloud Insights environment
configuration.
image:LogExplorer 2022.png[Choose your log]
NOTE: The types of logs available for querying may vary based on your
environment. Additional log types may be added over time.
You can set filters to further refine the results of the query. For
example, to find all log messages showing a failure, set a filter for
Messages containing the word "failed".
TIP: You can begin typing the desired text in the filter field; Cloud
Insights will prompt you to create a wildcard search containing the string
as you type.
The results are displayed in a graph showing the number of log instances
in each time period shown. Below the graph are the log entries temselves.
The graph and the entries refresh automatically based on the selected time
range.
image:LogExplorer QueryForFailed.png[Query example showing filter]
== Filtering
=== Include / Exclude
When filtering the logs, you can choose to *include* (i.e. "Filter to") or
*exclude* the strings you type. Excluded strings are displayed in the
```

completed filter as "NOT <string>".

image:Log_Advanced_Query_Filter_Exclude.png[Log Filter Showing Exclude
option]

Filters based on wildcards or expressions (e.g. NOT, OR, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.

NOTE: At any point, you can click on _Create a Log Monitor_ to create a new Monitor based on the current filter.

=== Advanced Filtering

When you are filtering for text or list values in queries or dashboard widgets, as you begin typing you are presented with the option to create a *wildcard filter* based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create expressions using NOT, AND, or OR, or you can select the "None" option to filter for null values.

NOTE: Be sure to Save your query early and often as you build your filtering. Advanced Querying is "free-form" string entry, and parsing mistakes may occur as you build.

Take a look at this screen image showing filtered results for an advanced query of the _logs.kubernetes.event_ log. There is a lot going on in this page, which is explained below the image:

image:Log_Advanced_Query_ScreenExplained.png[The Advanced Query Screen
Explained]

- 1. This advanced query string filters for the following:
- * Filter for log entries with a _reason_ that includes the word "failed", but not anything with the specific reason of "FailedMount".
- * Include any of those entries that also include a _metadata.namespace_ including the word "monitoring", but exclude the specific namespaces of "cm-monitoring" or "eg-monitoring".

+

Note that in the case above, since both "cm-monitoring" and "eg-monitoring" contain a dash ("-"), the strings must be included in double-quotes or a parsing error will be displayed. Strings that do not include dashes, spaces, etc. do not need to be enclosed in quotes. If in doubt, try putting the string in quotes.

- 2. The results of the current filter, including any "Filter By" values AND the Advanced Query filter, are displayed in the results list. The list can be sorted by any displayed columns. To display additional columns, select the "gear" icon.
- 3. The graph has been zoomed in to show only log results that occurred within a specific time frame. The time range shown here reflects the current zoom level. Select the _Reset Zoom_ button to set the zoom level back to the current Cloud Insights time range.
- 4. The chart results have been Grouped By the _source_ field. The chart shows results in each column grouped into colors. Hovering over a column in the chart will display some details about the specific entries.

image:Log Advanced Query Group Detail.png[Grouped Hover Details]

==== Refining Filters

You can use the following to refine your filter:

|===

|Filter|What it does

- | * (Asterisk) | enables you to search for everything
- | ? (question mark) | enables you to search for a specific number of characters
- | OR |enables you to specify multiple entities
- | NOT |allows you to exclude text from the search results
- | None |searches for NULL values in all fields
- | Not * |searches for NULL values in _text-only_ fields

|===

If you enclose a filter string in double quotes, Insight treats everything between the first and last quote as an exact match. Any special characters or operators inside the quotes will be treated as literals. For example, filtering for "*" will return results that are a literal asterisk; the asterisk will not be treated as a wildcard in this case. The operators OR and NOT will also be treated as literal strings when enclosed in double quotes.

You can combine a simple filter with an advanced query filter; the resulting filter is an "AND" of the two.

=== The Chart Legend

The _Legend_ below the chart has a few surprises as well. For each result (based on the current filter) shown in the Legend, you have an option to display only results for that line (Add Filter), or to display any results

NOT for that line (Add Exclude Filter). The chart and the Log Entries list update to show results based on your selection. To remove this filtering, open the Legend again and select the [X] to clear the Legend-based filter. image:Log Advanced Query Legend.png[Advanced Query Legen showing "Add Filter" for the first result] //// == The Log Graph The graph shows the number of log entries, grouped into buckets, which are based on the selected dashboard time range. The buckets for each time range are as follows: |=== |Dashboard Time Range|Bucket size |Last 15 Minutes|10 Seconds |Last 30 Minutes|15 Seconds |Last 60 Minutes|30 Seconds |Last 2 Hours|1 Minute |Last 3 Hours|5 Minutes |Last 6 Hours|5 Minutes |Last 12 Hours|10 Minutes |Last 24 Hours|15 Minutes |Last 2 Days|30 Minutes |Last 3 Days|45 Minutes |Last 7 Days|2 Hours |Last 30 Days|1 Day |=== //To zoom in the graph, simply drag the sliders from either side. To pan the zoomed area, click and hold in the white area and move left or right. Click Reset Zoom to reset the zoom level. //image:LogExplorer Zoom.png[Zoom in by dragging in the sides of the graph] //image:LogExplorer Zoom 2.png[Zoom in by dragging in the sides of the graph] Note that when zooming the graph or scrolling the table, dashboard autorefresh will pause and the time range will show the frozen time. To resume refresh, click the Resume button image:ResumeButton.png[]. This will also reset the zoom level. ////

```
== Log Details
Clicking anywhere in a log entry in the list will open a detail pane for
that entry. Here you can explore more information about the event.
Click on "Add Filter" to add the selected field to the current filter. The
log entry list will update based on the new filter.
image:LogExplorer DetailPane.png[Log Entry Detail Pane]
== Troubleshooting
Here you will find suggestions for troubleshooting problems with Log
Oueries.
|===
|*Problem:* | *Try this:*
|I don't see "debug" messages in my log query
|Debug log messaging is not collected. To capture messages you want,
change the relevant message severity to _informational, error, alert,
emergency, or _notice_ level.
|===
:leveloffset: -1
= Working with Annotations
:leveloffset: +1
[[IDbfbdf2696413e9f03c2c72188df922e9]]
= Defining annotations
:toc: macro
:hardbreaks:
:toclevels: 2
```

```
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

When customizing Cloud Insights to track data for your corporate requirements, you can define specialized notes, called annotations, and assign them to your assets.

You can assign annotations to assets with information such as asset end of life, data center, building location, storage tier, or volume service level.

Using annotations to help monitor your environment includes the following high-level tasks:

- * Creating or editing definitions for all annotation types.
- * Displaying asset pages and associating each asset with one or more annotations.

+

For example, if an asset is being leased and the lease expires within two months, you might want to apply an end-of-life annotation to the asset. This helps prevent others from using that asset for an extended time.

- * Creating rules to automatically apply annotations to multiple assets of the same type.
- * Filter assets by their annotations.
- == Default annotation types

Cloud Insights provides some default annotation types. These annotations can be used to filter or group data.

You can associate assets with default annotation types such as the following:

- * Asset life cycle, such as birthday, sunset, or end of life
- * Location information about a device, such as data center, building, or floor
- * Classification of assets, such as by quality (tiers), by connected devices (switch level), or by service level
- * Status, such as hot (high utilization)

The following table lists the Cloud Insights-provided annotation types.

[cols=3*, Optiosn="header", cols="30,53, 16"]

```
|Annotation types
|Description
|Type
|Alias|User-friendly name for a resource|Text
//|Birthday|Date device was/will be brought online|Date
//|Building|Physical location of assets|List
//|City|Municipality location of assets|List
|Compute Resource Group|Group assignment used by the Host and VM
Filesystems data collector|List
//|Continent|Geographic location of assets|List
//|Country|National location of assets|List
|Data Center|Physical location|List
//|Direct Attached|Indicates (Yes or No) if a storage resource is
connected directly to hosts|Boolean
//|End of Life |Date when a device will be taken offline|Date
//|Fabric Alias|User-friendly name for a fabric|Text
//|Floor|Location of a device on a floor of a building (hosts, storage
arrays, switches, and tapes) | List
|Hot|Devices under heavy use on a regular basis or at the threshold of
capacity|Boolean
|Note|Comments associated with a resource|Test
//|Rack|Rack in which the resource resides|List
//|SAN|Logical partition of the network for hosts, storage arrays, tapes,
switches, and applications.|List
|Service Level|A set of supported service levels that you can assign to
resources. Provides an ordered options list for internal volumes, gtree,
and volumes. Edit service levels to set performance policies for different
levels.|List
//|State/Province|State or province where the resource is located.|List
|Sunset|Threshold set after which no new allocations can be made to that
device. Useful for planned migrations and other pending network
changes. | Date
|Switch Level|Predefined options for setting up categories for switches.
Typically, these designations remain for the life of the device, although
you can edit them. Available only for switches. |List
|Tier|Can be used to define different levels of service within your
environment. Tiers can define the type of level, such as speed needed (for
example, gold or silver). This feature is available only on internal
volumes, gtrees, storage arrays, storage pools, and volumes. | List
|Violation Severity|Rank (for example, major) of a violation (for example,
missing host ports or missing redundancy), in a hierarchy of highest to
lowest importance.|List
NOTE: Alias, Data Center, Hot, Service Level, Sunset, Switch Level,
```

and Violation Severity are system-level annotations, which you cannot delete or rename; you can change only their assigned values.

== Creating custom annotations

Using annotations, you can add custom business-specific data that matches your business needs to assets. While Cloud Insights provides a set of default annotations, you might find that you want to view data in other ways. The data in custom annotations supplements device data already collected, such as storage manufacturer, number volumes, and performance statistics. The data you add using annotations is not discovered by Cloud Insights.

.Steps

. In the Cloud Insights menu, click *Manage > Annotations*.

+

The Annotations page displays the list of annotations.

- . Click *+Add*
- . Enter a *Name* and *Description* of the annotation.

+

You can enter up to 255 characters in these fields.

. Click *Type* and then select one of the following options that represents the type of data allowed in this annotation:

.Annotation types

Boolean:: Creates a drop-down list with the choices of yes and no. For example, the "Direct Attached" annotation is Boolean.

Date:: This creates a field that holds a date. For example, if the annotation will be a date, select this.

List:: Creates either of the following:

* A drop-down fixed list

+

When others are assigning this annotation type on a device, they cannot add more values to the list.

* A drop-down flexible list

+

If you select the Add new values on the fly option when you create this list, when others are assigning this annotation type on a device, they can add more values to the list.

Number:: Creates a field where the user assigning the annotation can enter a number. For example, if the annotation type is "Floor", the user could select the Value Type of "number" and enter the floor number.

Text:: Creates a field that allows free-form text. For example, you might

enter "Language" as the annotation type, select "Text" as the value type, and enter a language as a value. NOTE: After you set the type and save your changes, you cannot change the type of the annotation. If you need to change the type, you have to delete the annotation and create a new one. . If you select List as the annotation type, do the following: .. Select *Add new values on the fly* if you want the ability to add more values to the annotation when on an asset page, which creates a flexible list. For example, suppose you are on an asset page and the asset has the City annotation with the values Detroit, Tampa, and Boston. If you selected the *Add new values on the fly* option, you can add additional values to City like San Francisco and Chicago directly on the asset page instead of having to go to the Annotations page to add them. If you do not choose this option, you cannot add new annotation values when applying the annotation; this creates a fixed list. .. Enter a value and description in *Value* and *Description* fields. .. Click *+Add+* to add additional values. .. Click the Trash icon to delete a value. . Click *Save* Your annotations appear in the list on the Annotations page. .After you finish In the UI, the annotation is available immediately for use. [[ID7dfa666d767426ba6c05aeb98f7d54a8]] = Using annotations :toc: macro :hardbreaks: AA :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/

[.lead]

You create annotations and assign them to assets you monitor. Annotations are notes that provide information about an asset, such as physical location, end of life, storage tier, or volume service levels.

== Defining annotations

Using annotations, you can add custom business-specific data that matches your business needs to assets. While Cloud Insights provides a set of default annotations, such as asset life cycle (birthday or end of life), building or data center location, and tier, you might find that you want to view data in other ways.

The data in custom annotations supplements device data already collected, such as switch manufacturer, number of ports, and performance statistics. The data you add using annotations is not discovered by Cloud Insights.

.Before you begin

- * List any industry terminology to which environment data must be associated.
- * List corporate terminology to which environment data must be associated.
- * Identify any default annotation types that you might be able to use.
- * Identify which custom annotations you need to create. You need to create the annotation before it can be assigned to an asset.

Use the following steps to create an annotation.

.Steps

- . In the Cloud Insights menu, click *Manage > Annotations*
- . Click *+ Annotation* to create a new annotation.
- . Enter a Name, Description, and type for the new annotation.

+

For example, enter the following to create a text annotation that defines the physical location of an asset in Data Center 4:

+

- * Enter a name for the annotation, such as "Location"
- * Enter a description of what the annotation is describing, such as
- "Physical location is Data Center 4"
- * Enter the 'type' of annotation it is, such as "Text".
- == Manually assigning annotations to assets

Assigning annotations to assets helps you sort, group, and report on assets in ways that are relevant to your business. Although you can assign

annotations to assets of a particular type automatically using annotation rules, you can assign annotations to an individual asset by using its asset page.

- .Before you begin
- * You must have created the annotation you want to assign.
- .Steps
- . Log in to your Cloud Insights environment.
- . Locate the asset to which you want to apply the annotation.
- ** You can locate assets by querying, choosing from a dashoard widget, or search. When you have located the asset you want, click the link to open the asset's landing page.
- . On the asset page, in the User Data section, click *+ Annotation*.
- . The Add Annotation dialog box displays.
- . Select an annotation from the list.
- . Click Value and do either of the following, depending on type of annotation you selected:
- ** If the annotation type is list, date, or Boolean, select a value from the list.
- ** If the annotation type is text, type a value.
- . Click *Save*.

If you want to change the value of the annotation after you assign it, click the annotation field and select a different value.

If the annotation is of list type for which the _Add new values on the fly_ option is selected, you can type a new value in addition to selecting an existing value.

== Assigning annotations using annotation rules

To automatically assign annotations to assets based on criteria that you define, you configure annotation rules. Cloud Insights assigns the annotations to assets based on these rules. Cloud Insights also provides two default annotation rules, which you can modify to suit your needs or remove if you do not want to use them.

////

=== Default storage annotation rules

To expedite the assignment of storage annotations to your resources, Cloud Insights includes 21 default annotation rules, which associate a tier level with a storage tier model. All of your storage resources are automatically associated with a tier upon acquisition of the assets in your environment.

The default annotation rules apply tier annotations in the following way:

Tier 1, storage quality tier::

The Tier 1 annotation is applied to the following vendors and their specified families: EMC (Symmetrix), HDS (HDS9500V, HDS9900, HDS9900V, R600, R700, USP r, USP V), IBM (DS8000), NetApp (FAS6000 or FAS6200), and Violin (Memory).

Tier 2, storage quality tier::

The Tier 2 annotation is applied to the following vendors and their specified families: HP (3PAR StoreServ or EVA), EMC (CLARiiON), HDS (AMS or D800), IBM (XIV), and NetApp (FAS3000, FAS3100, and FAS3200).

You can edit the default settings of these rules to match your tier requirements, or you can remove them if you do not need them. ////

=== Creating annotation rules

As an alternative to manually applying annotations to individual assets, you can automatically apply annotations to multiple assets using annotation rules. Annotations set manually on an individual asset pages take precedence over rule-based annotations when Insight evaluates the annotation rules.

.Before you begin

You must have created a query for the annotation rule.

.About this task

Although you can edit the annotation types while you are creating the rules, you should have defined the types ahead of time.

.Steps

. Click *Manage > Annotation rules*

+

The Annotation Rules page displays the list of existing annotation rules.

- . Click *+ Add*.
- . Do the following:
- .. In the *Name* box, enter a unique name that describes the rule.

+

This name will appear in the Annotation Rules page.

- .. Click *Query* and select the query that is used to apply the annotation to assets.
- .. Click *Annotation* and select the annotation you want to apply.

```
.. Click *Value* and select a value for the annotation.
For example, if you choose Birthday as the annotation, you specify a date
for the value.
.. Click *Save*
.. Click *Run all rules* if you want to run all the rules immediately;
otherwise, the rules are run at a regularly scheduled interval.
[[IDe8a6ca007914e5f22b359eb147b2d278]]
= Creating annotation rules
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
You can use annotation rules to automatically apply annotations to
multiple assets based on criteria that you define. Cloud Insights assigns
the annotations to assets based on these rules. Annotations set manually
on an individual asset pages take precedence over rule-based annotations
when Cloud Insight evaluates the annotation rules.
.Before you begin
You must have created a query for the annotation rule.
.Steps
. In the Cloud Insights menu click *Manage > Annotation rules*.
. Click *+ Rule* to add a new annotation rule.
The Add Rule dialog is displayed.
. Do the following:
.. In the *Name* box, enter a unique name that describes the rule.
The name appears in the Annotation Rules page.
.. Click *Query* and select the query that Cloud Insights uses to identify
the assets the annotation applies to.
.. Click *Annotation* and select the annotation you want to apply.
.. Click *Value* and select a value for the annotation.
For example, if you choose Birthday as the annotation, you specify a date
```

for the value.

- .. Click *Save*
- .. Click *Run all rules* if you want to run all the rules immediately; otherwise, the rules are run at a regularly scheduled interval.

+

NOTE: In a large Cloud Insights environment, you may notice that running annotation rules seems to take a while to complete. This is because the indexer runs first and must complete prior to running the rules. The indexer is what gives Cloud Insights the ability to search or filter for new or updated objects and counters in your data. The rules engine waits until the indexer completes its update before applying the rules.

== Modifying annotation rules

You can modify an annotation rule to change the rule's name, its annotation, the annotation's value, or the query associated with the rule.

.Steps

. In the Cloud Insights menu, Click *Manage > Annotation rules*.

+

The Annotation Rules page displays the list of existing annotation rules.

. Locate the Annotation Rule you want to modify.

+

You can filter the annotation rules by entering a value in the filter box or click a page number to browse through the annotation rules by page.

- . Click the menu icon for the rule that you want to modify.
- . Click *Edit*

+

The Edit Rule dialog is displayed.

- . Modify the annotation rule's name, annotation, value, or query.
- == Changing the Order of Rules

Annotation rules are processed from the top of the rules list to the bottom. To change the order in which a rule is processed, do the following:

.Steps

- . Click on the menu icon for the rule you want to move.
- . Click *Move Up* or *Move Down* as needed until the rule appears in the location you want.

Note that when running multiple rules that update the same annotation on an asset, the first rule (as run from the top down) applies the annotation and updates the asset, then the second rule applies but doesn't change any annotation that was already set by the previous rule.

```
== Deleting annotation rules
You might want to delete annotation rules that are no longer used.
.Steps
. In the Cloud Insights menu, Click *Manage > Annotation rules*.
The Annotation Rules page displays the list of existing annotation rules.
. Locate the Annotation Rule you want to delete.
You can filter the annotation rules by entering a value in the filter box
or click a page number to browse through the annotation rules by page.
. Click the menu icon for the rule that you want to delete.
. Click *Delete*
A confirmation message is displayed, prompting whether you want to delete
the rule.
. Click *OK*
[[IDe514772f06adbd0f4a0d253c559d5a99]]
= Importing Annotations
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights includes an API for importing annotations or applications
from a CSV file, and assigning them to objects you specify.
NOTE: The Cloud Insights API is available in *Cloud Insights Premium
Edition*.
//The export and import functions are supported only between servers that
are running the same version of OnCommand Insight.
== Importing
The *Admin > API Access* links contain
```

```
xref: {relative path}API Overview.html[documentation] for the
*Assets/Import* API. This documentation contains information on the .CSV
file format.
image:api assets import.png[Import API]
== .CSV File Format
The general format of the CSV file is as follows. The first line of the
file defines the import fields and specifies the order of the fields. This
is followed by separate lines for each annotation or application. You do
not need to define every field. However, the subsequent annotation lines
must follow the same order as the definition line.
[Object Type] , [Object Name or ID] , Annotation Type [, Annotation
Type, ...] [, Application] [, Tenant] [, Line Of Business] [,
Business Unit] [, Project]
See the API Documentation for examples of .CSV files.
////
<Object Type Value 1>, <Object Name or Key 1>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line Of Business>] [, <Business Unit>] [, <Project>]
<Object Type Value 2>, <Object Name or Key 2>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line Of Business>] [, <Business Unit>] [, <Project>]
<Object Type Value 3>, <Object Name or Key 3>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line Of Business>] [, <Business Unit>] [, <Project>]
. . .
<Object Type Value N>, <Object Name or Key N>, <Annotation Value> [,
<Annotation Value> ...] [, <Application>] [, <Tenant>] [,
<Line Of Business>] [, <Business Unit>] [, <Project>]
////
You can import and assign annotations from a .CSV file from within the API
swagger itself. Simply choose the file to use and click the Execute
button:
image:api assets import assign.png[Import and Assign]
== Import Behavior
```

During the import operation, data is added, merged, or replaced, depending on the objects and object types that are being imported. While importing, keep in mind the following behaviors.

- * Adds an annotation or application if none exists with the same name in the target system.
- * Merges an annotation if the annotation type is a list, and an annotation with the same name exists in the target system.
- * Replaces an annotation if the annotation type is anything other than a list, and an annotation with the same name exists in the target system.

Note: If an annotation with the same name but with a different type exists in the target system, the import fails. If objects depend on the failed annotation, those objects may show incorrect or unwanted information. You must check all annotation dependencies after the import operation is complete.

- * If an annotation value is empty then that annotation is removed from the object. Inherited annotations are not affected.
- * Date type annotation values must be passed in as unix time in milliseconds.
- * When annotating volumes or internal volumes, the object name is a combination of storage name and volume name using the "\->" separator. For example: _<Storage Name>\-><Volume Name>_
- * If an object name contains a comma, the whole name must be in double quotes. For example: _"NetApp1, NetApp2"\->023F_
- * When attaching annotating to storages, switches, and ports, the 'Application' column will be ignored.
- * Tenant, Line_Of_Business, Business_Unit, and/or Project makes a business entity. As with all business entities, any of the values can be empty.

The following object types can be annotated.

|=== |OBJECT TYPE |NAME OR KEY |Host|id\-><id> or <Name> or <IP> |VM|id\-><id> or <Name> |StoragePool|id\-><id> or <Storage Name>\-><Storage Pool Name> |InternalVolume|id\-><id> or <Storage Name>\-><Internal Volume Name> |Volume|id\-><id> or <Storage Name>\-><Volume Name> |Storage|id\-><id> or <Name> or <IP> |Switch|id\-><id> or <Name> or <IP> |Port|id\-><id> or <Storage Name>\-><Internal Volume Name>\-><Qtree Name> |All or <Name> or <IP> |All or <Name> or <IIP> |All or <III or <I

|===

//• TBD: ONLY FOR OCI? The user can import a csv file and execute the API for the annotation assignment -

////

* Annotation Rules

+

Adds an annotation rule if no annotation rule with the same name exists in the target system.

Replaces an annotation rule if an annotation rule with the same name exists in the target system.

Note: Annotation rules are dependent on both queries and annotations. You must check all the annotation rules for accuracy after the import operation is complete.

////

////

Policies

Adds a policy if no policy with the same name exists in the target system. Replaces a policy if a policy with the same name exists in the target system.

Note: Policies may be out of order after the import operation is complete. You must check the policy order after the import.

Policies that are dependent on annotations may fail if the annotations are incorrect. You must check all the annotation dependencies after the import.

Queries

Adds a query if no query with the same name exists in the target system. Replaces a query if a query with the same name exists in the target system, even if the resource type of the query is different.

Note: If the resource type of a query is different, after the import, any dashboard widgets that use that query may display unwanted or incorrect results. You must check all query-based widgets for accuracy after the import.

Queries that are dependent on annotations may fail if the annotations are incorrect. You must check all the annotation dependencies after the import.

Dashboards

Adds a dashboard if no dashboard with the same name exists in the target system.

Replaces a dashboard if a dashboard with the same name exists in the

target system, even if the resource type of the query is different.

Note: You must check all query-based widgets in dashboards for accuracy after the import.

If the source server has multiple dashboards with the same name, they are all exported. However, only the first one will be imported to the target server. To avoid errors during import, you should ensure that your dashboards have unique names before exporting them.

////

:leveloffset: -1

= Working with Applications

:leveloffset: +1

[[IDaadceffc6e72831fcc38a065335771fa]]

= Tracking asset usage by application

[.lead]

Understanding the applications used in your company's environment helps you to keep track of asset usage and cost.

Before you can track data associated with the applications running in your environment, you must first define those applications and associate them with the appropriate assets. You can associate applications with the following assets: hosts, virtual machines, volumes, internal volumes, qtrees, shares, and hypervisors.

This topic provides an example of tracking the usage of virtual machines that the Marketing Team uses for its Exchange email.

You might want to create a table similar to the following to identify applications used in your environment and note the group or business unit using each applications.

[cols=5*, options="header]

|Tenant|Line of Business|Business Unit|Project|Applications |NetApp|Data Storage|Legal|Patents|Oracle Identity Manager, Oracle On Demand, PatentWiz

|NetApp|Data Storage|Marketing|Sales Events|Exchange, Oracle Shared DataBase, BlastOff Event Planner

|===

The table shows that that Marketing Team uses the Exchange application. We want to track their virtual machine utilization for Exchange, so that we can predict when we will need to add more storage. We can associate the Exchange application with all of Marketing's virtual machines: . Create an application named Exchange . Go to *Queries > +New Query* to create a new query for virtual machines (or select an existing VM query, if applicable). Assuming the Marketing team's VMs all have a name containing the string "*mkt*", create your query to filter VM name for "mkt". . Select the VMs. . Associate the VMs with the Exchange application using *Bulk Actions > Add Applications*. . Select the desired application and click *Save*. . When finished, *Save* the query. [[ID0c7d50d5843dcda73e3411ac7fc182e6]] = Creating Applications :toc: macro :hardbreaks: :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead] To track data associated with specific applications running in your environment, you can define the applications in Cloud Insights. .Before you begin If you want to associate the application with a business entity, you must create the business entity before you define the application. .About this task Cloud Insights allows you to track data from assets associated with applications for things like usage or cost reporting. .Steps

. In the Cloud Insights menu, click *Manage > Applications*.

+

The Add Application dialog box displays.

- . Enter a unique name for the application.
- . Select a priority for the application.
- . Click *Save*.

After defining an application, it can be assigned to assets.

== Assigning applications to assets

This procedure assigns the application to a host as an example. You can assign host, virtual machine, volume, or internal volumes to an application.

.Steps

- . Locate the asset to which you want to assign to the application:
- . Click *Queries > +New Query* and search for Host.
- . Click the check box on the left of the Host you want to associate with the application.
- . Click *Bulk Actions > Add Application*.
- . Select the Application you are assigning the asset to.

Any new applications you assign override any applications on the asset that were derived from another asset. For example, volumes inherit applications from hosts, and when new applications are assigned to a volume, the new application takes precedence over the derived application.

NOTE: For environments with large amounts of related assets, inheritance of application assignments to those assets could take several minutes. Please allow more time for inheritance to occur if you have many related assets.

.After you finish

After assigning the host to the application you can assign the remaining assets to the application. To access the landing page for the application, click *Manage > Application* and select the application you created.

:leveloffset: -1

```
= Monitors and Alerts
:leveloffset: +1
[[ID0316d9dda54967bdc76c3867fce774f8]]
= Alerting with Monitors
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
You create monitors to set thresholds that trigger alerts to notify you
about issues related to the resources in your network. For example, you
can create a monitor to alert for node write latency for any of a
multitude of protocols.
TIP: Monitors and Alerting is available in all Cloud Insights Editions,
however, Basic Edition is subject to the following:
* You may only have up to five custom monitors active at a time. Any
monitors beyond five will be created in or moved to Paused state.
* VMDK, Virtual Machine, Host, and DataStore metrics monitors are not
supported. If you have monitors created for these metrics, they will be
paused and cannot be resumed when downgrading to Basic Edition.
toc::[]
//When the monitored threshold and conditions are reached or exceeded,
Cloud Insights creates an alert. A Monitor can have a Warning threshold,
a Critical threshold, or both. Log Monitors can also have an
Informational alert level.
Monitors allow you to set thresholds on metrics generated by
"infrastructure" objects such as storage, VM, EC2, and ports, as well as
for "integration" data such as those collected for Kubernetes, ONTAP
advanced metrics, and Telegraf plugins. These _metric_ monitors alert you
when warning-level or critical-level thresholds are crossed.
You can also create monitors to trigger warning-, critical-, or
informational-level alerts when specified log events are detected.
Cloud Insights provides a number of
```

xref:{relative path}task system monitors.html[System-Defined Monitors] as well, based on your environment. == Security Best Practice Cloud Insights alerts are designed to highlight data points and trends in your environment, and Cloud Insights allows you to enter any valid email address as an alert recipient. If you are working in a secure environment, be especially mindful of who is receiving the notification or otherwise has access to the alert. == Metric or Log Monitor? . From the Cloud Insights menu, click *Alerts > Manage Monitors* The Monitors list page is displayed, showing currently configured monitors. . To modify an existing monitor, click the monitor name in the list. . To add a monitor, Click *+ Monitor*. image:Monitor log or metric.png[Choose system or log monitor] When you add a new monitor, you are prompted to create a Metric Monitor or a Log Monitor.

- * _Metric_ monitors alert on infrastructure- or performance-related triggers
- * Log monitors alert on log-related activity

After you choose your monitor type, the Monitor Configuration dialog is displayed. Configuration varies depending on which type of monitor you are creating.

=== Metric Monitor

. In the drop-down, search for and choose an object type and metric to monitor.

You can set filters to narrow down which object attributes or metrics to monitor.

//image:select metric to monitor.png[Select Metric]

image:MonitorMetricFilter.png[Metrics Filtering]

//When working with integration data (Kubernetes, ONTAP Advanced Data, etc.), metric filtering works against the data samples themselves, not the objects as with infrastructure data (storage, VMs, ports, etc.).

When working with integration data (Kubernetes, ONTAP Advanced Data, etc.), metric filtering removes the individual/unmatched data points from the plotted data series, unlike infrastructure data (storage, VM, ports etc.) where filters work on the aggregated value of the data series and potentially remove the entire object from the chart.

//image:IntegrationMetricFilterExample.png[Integration Metric Filtering]

TIP: To create a multi-condition monitor (e.g., IOPS > X and latency > Y), define the first condition as a threshold and the second condition as a filter.

==== Define the Conditions of the Monitor.

- . After choosing the object and metric to monitor, set the Warning-level and/or Critical-level thresholds.
- . For the _Warning_ level, enter 200 for our example. The dashed line indicating this Warning level displays in the example graph.
- . For the _Critical_ level, enter 400. The dashed line indicating this Critical level displays in the example graph.

The graph displays historical data. The Warning and Critical level lines on the graph are a visual representation of the Monitor, so you can easily see when the Monitor might trigger an alert in each case.

. For the occurence interval, choose <code>_Continuously_</code> for a period of <code>_15</code> <code>Minutes</code> .

+

You can choose to trigger an alert the moment a threshold is breached, or wait until the threshold has been in continuous breach for a period of time. In our example, we do not want to be alerted every time the Total IOPS peaks above the Warning or Critical level, but only when a monitored object continuously exceeds one of these levels for at least 15 minutes.

//image:define_monitor_conditions.png[Define Conditions]
image:Monitor_metric_conditions.png[Define the monitor's conditions]

=== Log Monitor

When creating a *Log monitor*, first choose which log to monitor from the available log list. You can then filter based on the available attributes as above. You can also choose one or more "Group By" attributes.

NOTE: The Log Monitor filter cannot be empty.

//image:Monitor_log_monitor_filter.png[choose which log to monitor, and
set a filter]

image:Monitor_Group_By_Example.png[Choose log to monitor, set a filter, and select group by method, if desired]

```
////
///==== Define the alert behavior
```

Choose how you want to alert when a log alert is triggered. You can set the monitor to alert with _Warning_, _Critical_, or _Informational_ severity, based on the filter conditions you set above.

image:Monitor_log_alert_behavior.png[define the log behavior to monitor]
////

==== Define the alert Behavior

You can create the monitor to alert with a severity level of _Critical_, _Warning_, or _Informational_, when the conditions you defined above occur once (i.e. immediately), or wait to alert until the conditions occur 2 times or more.

==== Define the alert resolution behavior

You can choose how a log monitor alert is resolved. You are presented with three choices:

- * Resolve instantly
- * Purge after the data retention period (please refer to the Editions Page for details). Note that the Monitor has no resolution condition by definition, so an Alert will stay _active_ and suppress all subsequent alerts with matching _group_by_ generated by this monitor, until the data retention period has passed.
- * Resolve based on log entry: Resolve alert when the log line is discovered as outlined in the following definition, or purge after the

data retention period. image:LogMonitorAlertResolution.png[Alert Resolution Options] //// * *Resolve instantly*: The alert is immediately resolved with no further action needed * *Resolve based on time*: The alert is resolved after the specified time has passed * *Resolve based on log entry*: The alert is resolved when a subsequent log activity has occurred. For example, when an object is logged as "available". image:Monitor log monitor resolution.png[Alert Resolution] //// === Select notification type and recipients In the _Set up team notification(s)_ section, you can choose whether to alert your team via email or Webhook. image:Webhook Choose Monitor Notification.png[Choose alerting method] *Alerting via Email:* Specify the email recipients for alert notifications. If desired, you can choose different recipients for warning or critical alerts. image:email monitor alerts.png[Email Alert Recipients] *Alerting via Webhook:* Specify the webhook(s) for alert notifications. If desired, you can choose different webhooks for warning or critical alerts. image:Webhook Monitor Notifications.png[Webhook Alerting] NOTE: ONTAP Data Collector notifications take precedence over any specific Monitor notifications that are relevant to the cluster/data collector. The recipient list you set for the Data Collector itself will receive the data collector alerts. If there are no active data collector alerts, then monitor-generated alerts will be sent to specific monitor recipients. === Setting Corrective Actions or Additional Information

You can add an optional description as well as additional insights and/or corrective actions by filling in the *Add an Alert Description* section. The description can be up to 1024 characters and will be sent with the alert. The insights/corrective action field can be up to 67,000 characters and will be displayed in the summary section of the alert landing page.

In these fields you can provide notes, links, or steps to take to correct or otherwise address the alert.

image:Monitors_Alert_Description.png[Alert Corrective Actions and
Description]

=== Save your Monitor

. If desired, you can add a description of the monitor.

+

. Give the Monitor a meaningful name and click *Save*.

+

Your new monitor is added to the list of active Monitors.

== Monitor List

The Monitor page lists the currently configured monitors, showing the following:

- * Monitor Name
- * Status
- * Object/metric being monitored
- * Conditions of the Monitor

You can choose to temporarily pause monitoring of an object type by clicking the menu to the right of the monitor and selecting *Pause*. When you are ready to resume monitoring, click *Resume*.

You can copy a monitor by selecting *Duplicate* from the menu. You can then modify the new monitor and change the object/metric, filter, conditions, email recipients, etc.

If a monitor is no longer needed, you can delete it by selecting *Delete* from the menu.

== Monitor Groups

Grouping allows you to view and manage related monitors. For example, you can have a monitor group dedicated to the storage in your environment, or monitors relevant to a certain recipient list.

image:Monitors GroupList.png[Monitor Grouping]

The following monitor groups are shown. The number of monitors contained in a group is shown next to the group name.

- * *All Monitors* lists all monitors.
- * *Custom Monitors* lists all user-created monitors.
- * *Suspended Monitors* will list any system monitors that have been suspended by Cloud Insights.
- * Cloud Insights will also show a number of *System Monitor Groups*, which will list one or more groups of

xref:{relative_path}task_system_monitors.html[system-defined monitors],
including ONTAP Infrastructure and Workload monitors.

//* *Data Collection* monitors will alert on Data Collector or Acquisition
Unit issues.

NOTE: Custom monitors can be paused, resumed, deleted, or moved to another group. System-defined monitors can be paused and resumed but can not be deleted or moved.

=== Suspended Monitors

This group will only be shown if Cloud Insights has suspended one or more monitors. A monitor may be suspended if it is generating excessive or continuous alerts. If the monitor is a custom monitor, modify the conditions to prevent the continuous alerting, and then resume the monitor. The monitor will be removed from the Suspended Monitors group when the issue causing the suspension is resolved.

```
////
=== Data Collection Monitors
This group will show monitors
////
```

=== System-Defined Monitors

These groups will show monitors provided by Cloud Insights, as long as your environment contains the devices and/or log availability required by the monitors.

System-Defined monitors cannot be modified, moved to another group, or deleted. However, you can duplicate a system monitor and modify or move the duplicate.

System monitors may include monitors for ONTAP Infrastructure (storage, volume, etc.) or Workloads (i.e. log monitors), or other groups. NetApp is constantly evaluating customer need and product functionality, and will update or add to system monitors and groups as needed.

=== Custom Monitor Groups

You can create your own groups to contain monitors based on your needs. For example, you may want a group for all of your storage-related monitors.

To create a new custom monitor group, click the *"+" Create New Monitor Group* button. Enter a name for the group and click *Create Group*. An empty group is created with that name.

To add monitors to the group, go to the _All Monitors_ group (recommended) and do one of the following:

- * To add a single monitor, click the menu to the right of the monitor and select Add to Group . Choose the group to which to add the monitor.
- * Click on the monitor name to open the monitor's edit view, and select a group in the _Associate to a monitor group_ section.

image:Monitors AssociateToGroup.png[Associate to group]

//* To add multiple monitors to a group, select them by clicking the
checkbox next to each monitor, then click the *Bulk Actions* button and
select _Move to Group_.

Remove monitors by clicking on a group and selecting _Remove from Group_ from the menu. You can not remove monitors from the _All Monitors_ or _Custom Monitors_ group. To delete a monitor from these groups, you must delete the monitor itself.

//To remove a monitor from a group while editing the monitor, in the Associate with a group section, click the *X* next to the group name.

NOTE: Removing a monitor from a group does not delete the monitor from Cloud Insights. To completely remove a monitor, select the monitor and click _Delete_. This also removes it from the group to which it belonged and it is no longer available to any user.

You can also move a monitor to a different group in the same manner, selecting Move to Group . //NOTE: Each monitor can belong to only a single group at any given time (in addition to belonging to "All Monitors" and "Custom Monitors"). To pause or resume all monitors in a group at once, select the menu for the group and click _Pause_ or _Resume_. Use the same menu to rename or delete a group. Deleting a group does not delete the monitors from Cloud Insights; they are still available in All Monitors . image:Monitors PauseGroup.png[Pause a group] //// === Creating a Monitor In the example below, we will create a Monitor to give a Warning alert when Volume Node NFS Write Latency reaches or exceeds 200ms, and a Critical alert when it reaches or exceeds 400ms. We only want to be alerted when either threshold is exceeded for at least 15 continuous minutes. ==== Requirements * Cloud Insights must be configured to collect integration data, and that data is being collected. //// //// ==== Create the Monitor . From the Cloud Insights menu, click *Alerts > Manage Monitors* The Monitors list page is displayed, showing currently configured monitors. . To add a monitor, Click *+ Monitor*. To modify an existing monitor, click the monitor name in the list. The Monitor Configuration dialog is displayed. . In the drop-down, search for and choose an object type and metric to monitor, for example netapp ontap volume node nfs write latency .

You can set filters to narrow down which object attributes or metrics to monitor.

//image:select_metric_to_monitor.png[Select Metric]

image:MonitorMetricFilter.png[Metrics Filtering]

//When working with integration data (Kubernetes, ONTAP Advanced Data, etc.), metric filtering works against the data samples themselves, not the objects as with infrastructure data (storage, VMs, ports, etc.).

When working with integration data (Kubernetes, ONTAP Advanced Data, etc.), metric filtering removes the individual/unmatched data points from the plotted data series, unlike infrastructure data (storage, VM, ports etc.) where filters work on the aggregated value of the data series and potentially remove the entire object from the chart.

//image:IntegrationMetricFilterExample.png[Integration Metric Filtering]

TIP: To create a multi-condition monitor (e.g., IOPS > X and latency > Y), define the first condition as a threshold and the second condition as a filter.

==== Define the Conditions of the Monitor.

- . After choosing the object and metric to monitor, set the Warning-level and/or Critical-level thresholds.
- . For the _Warning_ level, enter 200. The dashed line indicating this Warning level displays in the example graph.
- . For the _Critical_ level, enter 400. The dashed line indicating this Critical level displays in the example graph.

The graph displays historical data. The Warning and Critical level lines on the graph are a visual representation of the Monitor, so you can easily see when the Monitor might trigger an alert in each case.

. For the occurence interval, choose <code>_Continuously_</code> for a period of <code>_15</code> <code>Minutes</code> .

+

You can choose to trigger an alert the moment a threshold is breached, or wait until the threshold has been in continuous breach for a period of time. In our example, we do not want to be alerted every time the Total IOPS peaks above the Warning or Critical level, but only when a monitored object continuously exceeds one of these levels for at least 15 minutes.

+

image:define monitor conditions.png[Define Conditions]

==== Refining the Filter

When you are filtering, as you begin typing you are presented with the option to create a *wildcard filter* based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create *expressions* using NOT or OR, or you can select the "None" option to filter for null values in the field.

image:Type-Ahead_Monitor_1.png[Wildcard Filter]

Filters based on wildcards or expressions (e.g. NOT, OR, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.

image:Type-Ahead-Example-Wildcard-DirectSelect.png[Wildcard Filter
Results]

Note that Wildcard and Expression filtering works with text or lists but not with numerics, dates or booleans.

=== Setting Corrective Actions or Additional Information

You can add an optional description as well as additional insights and/or corrective actions by filling in the *Add an Alert Description* section. The description can be up to 1024 characters and will be sent with the alert. The insights/corrective action field can be up to 67,000 characters and will be displayed in the summary section of the alert landing page.

In these fields you can provide notes, links, or steps to take to correct or otherwise address the alert.

image:Monitors_Alert_Description.png[Alert Corrective Actions and Description]

=== Select notification type and recipients

In the _Set up team notification(s)_ section, you can choose whether to alert your team via email or Webhook.

```
image:Webhook Choose Monitor Notification.png[Choose alerting method]
*Alerting via Email:*
Specify the email recipients for alert notifications. If desired, you can
choose different recipients for warning or critical alerts.
image:email monitor alerts.png[Email Alert Recipients]
*Alerting via Webhook:*
Specify the webhook(s) for alert notifications. If desired, you can choose
different webhooks for warning or critical alerts.
image:Webhook Monitor Notifications.png[Webhook Alerting]
==== Warning vs. Critical vs. Resolved alerting
Whether a monitor sends a Warning, Critical, or Resolved alert
notification depends on which threshold is crossed:
* Crossing from non-triggered to WARNING - Send Warning Alert
* Crossing from non-triggered to CRITICAL - Send Critical Alert
* Crossing from WARNING to CRITICAL - Send Critical Alert
* Crossing from CRITICAL to WARNING - Send Warning Alert
* Crossing from WARNING to Non-Triggered - Send RESOLVED Alert
* Crossing from CRITICAL to Non-Triggered - Send RESOLVED Alert
=== Save your Monitor
. If desired, you can add a description of the monitor.
. Give the Monitor a meaningful name and click *Save*.
Your new monitor is added to the list of active Monitors.
=== Monitor List
The Monitor page lists the currently configured monitors, showing the
following:
* Monitor Name
* Status
* Object/metric being monitored
* Conditions of the Monitor
```

You can view any active alerts associated with a monitor by clicking rthe "bell" icon next to the Monitor name.

image:ViewActiveAlerts.png[Icon showing active alerts for a monitor]

You can choose to temporarily suspend monitoring of an object type by clicking the menu to the right of the monitor and selecting *Pause*. When you are ready to resume monitoring, click *Resume*.

You can copy a monitor by selecting *Duplicate* from the menu. You can then modify the new monitor and change the object/metric, filter, conditions, email recipients, etc.

If a monitor is no longer needed, you can delete it by selecting *Delete* from the menu.

== Monitor Groups

Grouping allows you to view and manage related monitors. For example, you can have a monitor group dedicated to the storage in your environment, or monitors relevant to a certain recipient list.

image:Monitors GroupList.png[Monitor Grouping]

////

Two groups are shown by default:

- * *All Monitors* lists all monitors.
- * *Custom Monitors* lists only user-created monitors.
- * *Suspended Monitors* will be shown only if a monitor has been suspended by the system.

////

////

The number of monitors contained in a group is shown next to the group name.

To create a new group, click the *"+" Create New Monitor Group* button. Enter a name for the group and click *Create Group*. An empty group is created with that name.

To add monitors to the group, go to the _All Monitors_ group (recommended) and do one of the following:

- * To add a single monitor, click the menu to the right of the monitor and select Add to Group . Choose the group to which to add the monitor.
- * Click on the monitor name to open the monitor's edit view, and select a group in the _Associate to a monitor group_ section.

image:Monitors_AssociateToGroup.png[Associate to group]

//* To add multiple monitors to a group, select them by clicking the
checkbox next to each monitor, then click the *Bulk Actions* button and
select _Move to Group_.

Remove monitors by clicking on a group and selecting _Remove from Group_ from the menu. You can not remove monitors from the _All Monitors_ or _Custom Monitors_ group. To delete a monitor from these groups, you must delete the monitor itself.

NOTE: Removing a monitor from a group does not delete the monitor from Cloud Insights. To completely remove a monitor, select the monitor and click _Delete_. This also removes it from the group to which it belonged and it is no longer available to any user.

You can also move a monitor to a different group in the same manner, selecting _Move to Group_.

NOTE: Each monitor can belong to only a single group at any given time.

To pause or resume all monitors in a group at once, select the menu for the group and click \mbox{Pause} or \mbox{Resume} .

Use the same menu to rename or delete a group. Deleting a group does not delete the monitors from Cloud Insights; they are still available in $_$ All Monitors .

image:Monitors_PauseGroup.png[Pause a group]
////

== System-Defined Monitors

Cloud Insights includes a number of system-defined monitors for both metrics and logs. The system monitors available are dependent on the data collectors present in your environment. Because of that, the monitors available in Cloud Insights may change as data collectors are added or their configurations changed.

View the xref:{relative path}task system monitors.html[System-Defined

```
Monitors] page for descriptions of monitors included with Cloud Insights.
=== More Information
* xref:{relative path}task view and manage alerts.html[Viewing and
Dismissing Alerts]
[[ID745c0eb118e066ed3de06e47d2d49bfe]]
= Viewing and Managing Alerts from Monitors
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights displays alerts when
xref:{relative path}task create monitor.html[monitored thresholds] are
exceeded.
TIP: Monitors and Alerting is available in Cloud Insights Standard Edition
and higher.
== Viewing and Managing Alerts
To view and manage alerts, do the following.
. Navigate to the *Alerts > All Alerts* page.
. A list of up to the most recent 1,000 alerts is displayed. You can sort
this list on any field by clicking the column header for the field. The
list displays the following information. Note that not all of these
columns are displayed by default. You can select columns to display by
clicking on the "gear" icon image:gear.png[gear icon]:
** *Alert ID*: System-generated unique alert ID
** *Triggered Time*: The time at which the relevant Monitor triggered the
```

alert

- ** *Current Severity* (Active alerts tab): The current severity of the active alert
- ** *Top Severity* (Resolved alerts tab); The maximum severity of the alert before it was resolved
- ** *Monitor*: The monitor configured to trigger the alert
- ** *Triggered On*: The object on which the monitored threshold was breached
- ** *Status*: Current alert status, New or In Process
- ** *Active Status*: Active or Resolved
- ** *Condition*: The threshold condition that triggered the alert
- ** *Metric*: The object's metric on which the monitored threshold was breached
- ** *Monitor Status*: Current status of the monitor that triggered the alert
- ** *Has Corrective Action*: The alert has suggested corrective actions. Open the alert page to view these.

You can manage an alert by clicking the menu to the right of the alert and choosing one of the following:

- * *In Process* to indicate that the alert is under investigation or otherwise needs to be kept open
- * *Dismiss* to remove the alert from the list of active alerts.

You can manage multiple alerts by selecting the checkbox to the left of each Alert and clicking _Change Selected Alerts Status_.

Clicking on an Alert ID opens the Alert Detail Page.

== Alert Detail Page

The Alert Detail Page provides additional detail about the alert, including a _Summary_, an _Expert View_ showing graphs related to the object's data, any _Related Assets_, and _Comments_ entered by alert investigators.

image:alert_detail_page.png[Alert Detail Page]

== Alerts When Data Is Missing

In a realtime system such as Cloud Insights, to trigger the analysis of a Monitor to decide if an Alert should be generated, we rely on one of two things:

- * the next datapoint to arrive
- * a timer to fire when there is no datapoint and you have waited long

enough

As is the case with slow data arrival—or no data arrival—the timer mechanism needs to take over as the data arrival rate is insufficient to trigger alerts in "real time." So the question typically becomes "How long do I wait before I close the analysis window and look at what I have?" If you wait too long then you are not generating the alerts fast enough to be useful.

If you have a Monitor with a 30-minute window that notices that a condition is violated by the last data point before a long-term loss-of-data, an Alert will be generated because the Monitor received no other information to use to confirm a recovery of the metric or notice that the condition persisted.

== "Permanently Active" Alerts

It is possible to configure a monitor in such a way for the condition to *always* exist on the monitored object--for example, IOPS > 1 or latency > 0. These are often created as 'test' monitors and then forgotten. Such monitors create alerts that stay permanently open on the constituent objects, which can cause system stress and stability issues over time.

To prevent this, Cloud Insights will automatically close any "permanently active" alert after 7 days. Note that the underlying monitor conditions may (probably will) continue to exist, causing a new alert to be issued almost immediately, but this closing of "always active" alerts alleviates some of the system stress that can otherwise occur.

```
[[ID3a7b28b8f4de7780e8f60fa5f08babc6]]
= Configuring Email Notifications
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

You can configure an email list for subscription-related notifications, as well as a global email list of recipients for notification of performance policy threshold violations.

//You can also configure a global email list of recipients for monitor-

related notifications. To configure notification email recipient settings, go to the *Admin > Notifications* page and select the Email tab. [.thumb] image:Notifications email list.png[Email Notifications] == Subscription Notification Recipients To configure recipients for subscription-related event notifications, go to the "Subscription Notification Recipients" section. You can choose to have email notifications sent for subscription-related events to any or all of the following recipients: * All Account Owners * All Monitor & Optimize Administrators * Additional Email Addresses that you specify The following are examples of the types of notifications that might be sent, and user actions you can take. |=== |*Notification:*|*User Action:* |Trial or subscription has been updated|Review subscription details on the xref:{relative path}concept subscribing to cloud insights.html[Subscriptio |Subscription will expire in 90 days Subscription will expire in 30 days | No action needed if "Auto Renewal" is enabled Contact link:https://www.netapp.com/us/forms/sales-inquiry/cloud-insightssales-inquiries.aspx[NetApp sales] to renew the subscription |Trial ends in 2 days|Renew trial from the xref:{relative path}concept subscribing to cloud insights.html[Subscriptio n] page. You can renew a trial one time. Contact link:https://www.netapp.com/us/forms/sales-inquiry/cloud-insightssales-inquiries.aspx[NetApp sales] to purchase a subscription |Trial or subscription has expired Account will stop collecting data in 48 hours Account will be deleted after 48 hours | Contact link:https://www.netapp.com/us/forms/sales-inquiry/cloud-insights-salesinquiries.aspx[NetApp sales] to purchase a subscription |=== == Global Recipient List for Alerts

Email notifications of alerts are sent to the alert recipient list for every action on the alert. You can choose to send alert notifications to a global recipient list.

To configure global alert recipients, choose the desired recipients in the *Global Monitor Notification Recipients* section.

You can always override the global recipients list for an individual monitor when creating or modifying the monitor.

NOTE: ONTAP Data Collector notifications take precedence over any specific Monitor notifications that are relevant to the cluster/data collector. The recipient list you set for the Data Collector itself will receive the data collector alerts. If there are no active data collector alerts, then monitor-generated alerts will be sent to specific monitor recipients.

//image:MonitorTeamNotifications.png[Override Global Notifications]

== Editing Notifications for ONTAP

You can modify notifications for ONTAP clusters by selecting _Edit Notifications from the upper-right drop-down on a Storage landing page.

image:EditONTAPNotifications.png[ONTAP Edit Notifications]

From here, you can set notifications for Critical, Warning, Informational, and/or Resolved alerts. Each scenario can notify the Global Recipient list or other recipients you choose.

image:EditONTAPNotifications_MultipleScenarios.png[ONTAP Notifications for different scenarios]

```
[[ID7b60d43812e045bbfcf7405e2f5d4d94]]
= System Monitors
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

Cloud Insights includes a number of system-defined monitors for both metrics and logs. The system monitors available are dependent on the data collectors present in your environment. Because of that, the monitors available in Cloud Insights may change as data collectors are added or their configurations changed.

NOTE: Many System Monitors are in _Paused_ state by default. You can enable a system monitor by selecting the _Resume_ option for the monitor. Ensure that _Advanced Counter Data Collection_ and _Enable ONTAP EMS log collection_ are enabled in the Data Collector. These options can be found in the ONTAP Data Collector under _Advanced Configuration_: image:Enable_Log_Monitor_Collection.png[Enabling Advanced Counter and EMS Log collection for ONTAP]

```
[#top]
toc::[]

////
== Create the Monitor

. From the Cloud Insights menu, click *Alerts > Manage Monitors*
+
The Monitors list page is displayed, showing currently configured monitors.

. To modify an existing monitor, click the monitor name in the list.

. To add a monitor, Click *+ Monitor*.
+
image:Monitor_log_or_metric.png[Choose system or log monitor]
+
```

When you add a new monitor, you are prompted to create a Metric Monitor or a Log Monitor.

- * _Metric_ monitors alert on infrastructure- or performance-related triggers
- * $_$ Log $_$ monitors alert on log-related activity

After you choose your monitor type, the Monitor Configuration dialog is displayed.

==== Metric Monitor

. In the drop-down, search for and choose an object type and metric to monitor.

You can set filters to narrow down which object attributes or metrics to monitor.

//image:select metric to monitor.png[Select Metric]

image:MonitorMetricFilter.png[Metrics Filtering]

//When working with integration data (Kubernetes, ONTAP Advanced Data, etc.), metric filtering works against the data samples themselves, not the objects as with infrastructure data (storage, VMs, ports, etc.).

When working with integration data (Kubernetes, ONTAP Advanced Data, etc.), metric filtering removes the individual/unmatched data points from the plotted data series, unlike infrastructure data (storage, VM, ports etc.) where filters work on the aggregated value of the data series and potentially remove the entire object from the chart.

//image:IntegrationMetricFilterExample.png[Integration Metric Filtering]

TIP: To create a multi-condition monitor (e.g., IOPS > X and latency > Y), define the first condition as a threshold and the second condition as a filter.

==== Define the Conditions of the Monitor.

- . After choosing the object and metric to monitor, set the Warning-level and/or Critical-level thresholds.
- . For the _Warning_ level, enter 200 for our example. The dashed line indicating this Warning level displays in the example graph.
- . For the _Critical_ level, enter 400. The dashed line indicating this Critical level displays in the example graph.

+

The graph displays historical data. The Warning and Critical level lines on the graph are a visual representation of the Monitor, so you can easily see when the Monitor might trigger an alert in each case.

. For the occurence interval, choose <code>_Continuously_</code> for a period of <code>_15</code> <code>Minutes_.</code>

+

You can choose to trigger an alert the moment a threshold is breached, or

wait until the threshold has been in continuous breach for a period of time. In our example, we do not want to be alerted every time the Total IOPS peaks above the Warning or Critical level, but only when a monitored object continuously exceeds one of these levels for at least 15 minutes.

//image:define_monitor_conditions.png[Define Conditions]
image:Monitor metric conditions.png[Define the monitor's conditions]

=== Log Monitor

When creating a *Log monitor*, first choose which log to monitor from the available log list. You can then filter based on the available attributes as above.

For example, you might choose to filter for "object.store.unavailable" message type in the logs.netapp.ems source:

NOTE: The Log Monitor filter cannot be empty.

image:Monitor_log_monitor_filter.png[choose which log to monitor, and set
a filter]

==== Define the alert behavior

Choose how you want to alert when a log alert is triggered. You can set the monitor to alert with _Warning_, _Critical_, or _Informational_ severity, based on the filter conditions you set above.

image:Monitor log alert behavior.png[define the log behavior to monitor]

==== Define the alert resolution behavior

You can choose how an log monitor alert is resolved. You are presented with three choices:

- * *Resolve instantly*: The alert is immediately resolved with no further action needed
- * *Resolve based on time*: The alert is resolved after the specified time has passed
- * *Resolve based on log entry*: The alert is resolved when a subsequent log activity has occurred. For example, when an object is logged as "available".

```
image:Monitor log monitor resolution.png[Alert Resolution]
==== Select notification type and recipients
In the Set up team notification(s) section, you can choose whether to
alert your team via email or Webhook.
image:Webhook Choose Monitor Notification.png[Choose alerting method]
*Alerting via Email:*
Specify the email recipients for alert notifications. If desired, you can
choose different recipients for warning or critical alerts.
image:email monitor alerts.png[Email Alert Recipients]
*Alerting via Webhook:*
Specify the webhook(s) for alert notifications. If desired, you can choose
different webhooks for warning or critical alerts.
image:Webhook Monitor Notifications.png[Webhook Alerting]
==== Setting Corrective Actions or Additional Information
You can add an optional description as well as additional insights and/or
corrective actions by filling in the *Add an Alert Description* section.
The description can be up to 1024 characters and will be sent with the
alert. The insights/corrective action field can be up to 67,000 characters
and will be displayed in the summary section of the alert landing page.
In these fields you can provide notes, links, or steps to take to correct
or otherwise address the alert.
image: Monitors Alert Description.png [Alert Corrective Actions and
Description]
==== Save your Monitor
. If desired, you can add a description of the monitor.
. Give the Monitor a meaningful name and click *Save*.
```

```
Your new monitor is added to the list of active Monitors.
==== Monitor List
The Monitor page lists the currently configured monitors, showing the
following:
* Monitor Name
* Status
* Object/metric being monitored
* Conditions of the Monitor
You can choose to temporarily pause monitoring of an object type by
clicking the menu to the right of the monitor and selecting *Pause*. When
you are ready to resume monitoring, click *Resume*.
You can copy a monitor by selecting *Duplicate* from the menu. You can
then modify the new monitor and change the object/metric, filter,
conditions, email recipients, etc.
If a monitor is no longer needed, you can delete it by selecting *Delete*
from the menu.
////
////
== Monitor Groups
Grouping allows you to view and manage related monitors. For example, you
can have a monitor group dedicated to the storage in your environment, or
monitors relevant to a certain recipient list.
image:Monitors GroupList.png[Monitor Grouping]
////
////
Two groups are shown by default:
* *All Monitors* lists all monitors.
* *Custom Monitors* lists only user-created monitors.
////
////
The number of monitors contained in a group is shown next to the group
name.
```

NOTE: Custom monitors can be paused, resumed, deleted, or moved to another group. System-defined monitors can be paused and resumed but can not be deleted or moved.

=== Custom Monitor Groups

To create a new custom monitor group, click the *"+" Create New Monitor Group* button. Enter a name for the group and click *Create Group*. An empty group is created with that name.

To add monitors to the group, go to the _All Monitors_ group (recommended) and do one of the following:

- * To add a single monitor, click the menu to the right of the monitor and select Move to Group . Choose the group to which to add the monitor.
- * To add multiple monitors, select the ones you want to move, and click _Bulk Actions_. Choose _Move to group_. You cannot move system-defined monitors.
- * Click on the monitor name to open the monitor's edit view, and select a group in the _Associate to a monitor group_ section.

image:Monitors AssociateToGroup.png[Associate to group]

//* To add multiple monitors to a group, select them by clicking the checkbox next to each monitor, then click the *Bulk Actions* button and select Move to Group .

Remove monitors by clicking on a group and selecting _Remove from Group_ from the menu. You can not remove monitors from the _All Monitors_ or _Custom Monitors_ group, or from any of the system-defined monitor groups. To delete a monitor from these groups, you must delete the monitor itself.

//To remove a monitor from a group while editing the monitor, in the Associate with a group section, click the *X* next to the group name.

NOTE: Removing a monitor from a group does not delete the monitor from Cloud Insights. To completely remove a monitor, select the monitor and click _Delete_. This also removes it from any group to which it belonged and it is no longer available to any user.

You can also move a monitor to a different group in the same manner, selecting _Move to Group_. You cannot move monitors from the system-defined monitor groups; to move a monitor from one of these groups, you must first _Duplicate_ the monitor, and then move the duplicated monitor to the desired group.

To pause or resume all monitors in a group at once, select the menu for the group and click \mbox{Pause} or \mbox{Resume} .

Use the same menu to rename or delete a group. Deleting a group does not delete the monitors from Cloud Insights; they are still available in _All Monitors_. You cannot delete system-defined monitor groups.

image:Monitors_PauseGroup.png[Pause a group]
////

== Monitor Descriptions

System-defined monitors are comprised of pre-defined metrics and conditions, as well as default descriptions and corrective actions, which can not be modified. You _can_ modify the notification recipient list for system-defined monitors. To view the metrics, conditions, description and corrective actions, or to modify the recipient list, open a system-defined monitor group and click the monitor name in the list.

System-defined monitor groups cannot be modified or removed.

The following system-defined monitors are available, in the noted groups.

- * *ONTAP Infrastructure* includes monitors for infrastructure-related issues in ONTAP clusters.
- * *ONTAP Workload Examples* includes monitors for workload-related issues.
- * Monitors in both group default to Paused state.

Below are the system monitors currently included with Cloud Insights:

=== Metric Monitors

|===

|Monitor Name|Severity|Monitor Description|Corrective Action |Fiber Channel Port Utilization High|CRITICAL|Fiber Channel Protocol ports are used to receive and transfer the SAN traffic between the customer host system and the ONTAP LUNs. If the port utilization is high, then it will become a bottleneck and it will ultimately affect the performance of sensitive of Fiber Channel Protocol workloads...A warning alert indicates that planned action should be taken to balance network traffic...A critical alert indicates that service disruption is imminent and emergency measures should be taken to balance network traffic to ensure service continuity.|If critical threshold is breached, consider immediate actions to minimize service disruption:

- 1. Move workloads to another lower utilized FCP port.
- 2. Limit the traffic of certain LUNs only to essential work, either via QoS policies in ONTAP or host-side configuration to lighten the utilization of the FCP ports....

If warning threshold is breached, plan to take the following actions:

- 1. Configure more FCP ports to handle the data traffic so that the port utilization gets distributed among more ports.
- 2. Move workloads to another lower utilized FCP port.
- 3. Limit the traffic of certain LUNs only to essential work, either via QoS policies in ONTAP or host-side configuration to lighten the utilization of the FCP ports.

|Lun Latency High|CRITICAL|LUNs are objects that serve the I/O traffic often driven by performance sensitive applications such as databases. High LUN latencies means that the applications themselves might suffer and be unable to accomplish their tasks...A warning alert indicates that planned action should be taken to move the LUN to appropriate Node or Aggregate...A critical alert indicates that service disruption is imminent and emergency measures should be taken to ensure service continuity. Following are expected latencies based on media type - SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds, and SATA HDD 17-20 milliseconds|If critical threshold is breached, consider following actions to minimize service disruption:

If the LUN or its volume has a QoS policy associated with it, then evaluate its threshold limits and validate if they are causing the LUN workload to get throttled....

If warning threshold is breached, plan to take the following actions:

- 1. If aggregate is also experiencing high utilization, move the LUN to another aggregate.
- 2. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node.
- 3. If the LUN or its volume has a QoS policy associated with it, evaluate its threshold limits and validate if they are causing the LUN workload to get throttled.

|Network Port Utilization High |CRITICAL|Network ports are used to receive and transfer the NFS, CIFS, and iSCSI protocol traffic between the customer host systems and the ONTAP volumes. If the port utilization is high, then it becomes a bottleneck and it will ultimately affect the performance of NFS, CIFS and iSCSI workloads...A warning alert indicates that planned action should be taken to balance network traffic...A critical alert indicates that service disruption is imminent and emergency measures should be taken to balance network traffic to ensure service continuity.|If critical threshold is breached, consider following immediate actions to minimize service disruption:

1. Limit the traffic of certain volumes only to essential work, either via QoS policies in ONTAP or host-side analysis to decrease the utilization of

the network ports.

- 2. Configure one or more volumes to use another lower utilized network port...
- If warning threshold is breached, consider the following immediate actions:
- 1. Configure more network ports to handle the data traffic so that the port utilization gets distributed among more ports.
- 2. Configure one or more volumes to use another lower utilized network port.

| NVMe Namespace Latency High | CRITICAL | NVMe Namespaces are objects that serve the I/O traffic that is driven by performance sensitive applications such as databases. High NVMe Namespaces latency means that the applications themselves may suffer and be unable to accomplish their tasks...A warning alert indicates that planned action should be taken to move the LUN to appropriate Node or Aggregate...A critical alert indicates that service disruption is imminent and emergency measures should be taken to ensure service continuity. | If critical threshold is breached, consider immediate actions to minimize service disruption:

If the NVMe namespace or its volume has a QoS policy assigned to them, then evaluate its limit thresholds in case they are causing the NVMe namespace workload to get throttled....

If warning threshold is breached, consider to take the following actions:

- 1. If aggregate is also experiencing high utilization, move the LUN to another aggregate.
- 2. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node.
- 3. If the NVMe namespace or its volume has a QoS policy assigned to them, evaluate its limit thresholds in case they are causing the NVMe namespace workload to get throttled.

|QTree Capacity Full|CRITICAL|A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a default space quota or a quota defined by a quota policy to limit amount of data stored in the tree within the volume capacity...A warning alert indicates that planned action should be taken to increase the space...A critical alert indicates that service disruption is imminent and emergency measures should be taken to free up space to ensure service continuity.|If critical threshold is breached, consider immediate actions to minimize service disruption:

- 1. Increase the space of the qtree in order to accommodate the growth.
- 2. Delete unwanted data to free up space....
- If warning threshold is breached, plan to take the following immediate actions:
- 1. Increase the space of the qtree in order to accommodate the growth.
- 2. Delete unwanted data to free up space.

|QTree Capacity Hard Limit|CRITICAL|A qtree is a logically defined file system that can exist as a special subdirectory of the root directory

within a volume. Each qtree has a space quota measured in KBytes that is used to store data in order to control the growth of user data in volume and not exceed its total capacity....A qtree maintains a soft storage capacity quota that provides alert to the user proactively before reaching the total capacity quota limit in the qtree and being unable to store data anymore. Monitoring the amount of data stored within a qtree ensures that the user receives uninterrupted data service. |If critical threshold is breached, consider following immediate actions to minimize service disruption:

- 1. Increase the tree space quota in order to accommodate the growth
- 2. Instruct the user to delete unwanted data in the tree to free up space |QTree Capacity Soft Limit|WARNING|A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a space quota measured in KBytes that it can use to store data in order to control the growth of user data in volume and not exceed its total capacity...A qtree maintains a soft storage capacity quota that provides alert to the user proactively before reaching the total capacity quota limit in the qtree and being unable to store data anymore. Monitoring the amount of data stored within a qtree ensures that the user receives uninterrupted data service.|If warning threshold is breached, consider the following immediate actions:
- 1. Increase the tree space quota to accommodate the growth.
- 2. Instruct the user to delete unwanted data in the tree to free up space. |QTree Files Hard Limit|CRITICAL|A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a quota of the number of files that it can contain to maintain a manageable file system size within the volume...A qtree maintains a hard file number quota beyond which new files in the tree are denied. Monitoring the number of files within a qtree ensures that the user receives uninterrupted data service.|If critical threshold is breached, consider immediate actions to minimize service disruption:
- 1. Increase the file count quota for the gtree.
- 2. Delete unwanted files from the qtree file system.

|QTree Files Soft Limit|WARNING|A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a quota of the number of files that it can contain in order to maintain a manageable file system size within the volume...A qtree maintains a soft file number quota to provide alert to the user proactively before reaching the limit of files in the qtree and being unable to store any additional files. Monitoring the number of files within a qtree ensures that the user receives uninterrupted data service.|If warning threshold is breached, plan to take the following immediate actions:

- 1. Increase the file count quota for the qtree.
- 2. Delete unwanted files from the qtree file system .

|Snapshot Reserve Space Full|CRITICAL|Storage capacity of a volume is

necessary to store application and customer data. A portion of that space, called snapshot reserved space, is used to store snapshots which allow data to be protected locally. The more new and updated data stored in the ONTAP volume the more snapshot capacity is used and less snapshot storage capacity is available for future new or updated data. If the snapshot data capacity within a volume reaches the total snapshot reserve space, it might lead to the customer being unable to store new snapshot data and reduction in the level of protection for the data in the volume.

Monitoring the volume used snapshot capacity ensures data services continuity. If critical threshold is breached, consider immediate actions to minimize service disruption:

- 1. Configure snapshots to use data space in the volume when the snapshot reserve is full.
- 2. Delete some older unwanted snapshots to free up space....
- If warning threshold is breached, plan to take the following immediate actions:
- 1. Increase the snapshot reserve space within the volume to accommodate the growth.
- 2. Configure snapshots to use data space in the volume when the snapshot reserve is full.

|Storage Capacity Limit|CRITICAL|When a storage pool (aggregate) is filling up, I/O operations slow down and finally stop resulting in storage outage incident. A warning alert indicates that planned action should be taken soon to restore minimum free space. A critical alert indicates that service disruption is imminent and emergency measures should be taken to free up space to ensure service continuity.|If critical threshold is breached, immediately consider the following actions to minimize service disruption:

- 1. Delete Snapshots on non-critical volumes.
- 2. Delete Volumes or LUNs that are non-essential workloads and that may be restored from off storage copies......If warning threshold is breached, plan the following immediate actions:
- 1. Move one or more volumes to a different storage location.
- 2. Add more storage capacity.
- 3. Change storage efficiency settings or tier inactive data to cloud storage.

|Storage Performance Limit|CRITICAL|When a storage system reaches its performance limit, operations slow down, latency goes up and workloads and applications may start failing. ONTAP evaluates the storage pool utilization for workloads and estimates what percent of performance has been consumed...A warning alert indicates that planned action should be taken to reduce storage pool load to ensure that there will be enough storage pool performance left to service workload peaks...A critical alert indicates that a performance brownout is imminent and emergency measures should be taken to reduce storage pool load to ensure service continuity.|If critical threshold is breached, consider following

immediate actions to minimize service disruption:

- 1. Suspend scheduled tasks such as Snapshots or SnapMirror replication.
- 2. Idle non-essential workloads....
- If warning threshold is breached, take the following actions immediately:
- 1. Move one or more workloads to a different storage location.
- 3. Change workload characteristics (block size, application caching).

 |User Quota Capacity Hard Limit|CRITICAL|ONTAP recognizes the users of
 Unix or Windows systems who have the rights to access volumes, files or
 directories within a volume. As a result, ONTAP allows the customers to
 configure storage capacity for their users or groups of users of their
 Linux or Windows systems. The user or group policy quota limits the amount
 of space the user can utilize for their own data...A hard limit of this
 quota allows notification of the user when the amount of capacity used
 within the volume is right before reaching the total capacity quota.

 Monitoring the amount of data stored within a user or group quota ensures
 that the user receives uninterrupted data service.|If critical threshold
 is breached, consider following immediate actions to minimize service
 disruption:
- 1. Increase the space of the user or group quota in order to accommodate the growth.
- 2. Instruct the user or group to delete unwanted data to free up space. |User Quota Capacity Soft Limit|WARNING|ONTAP recognizes the users of Unix or Windows systems that have the rights to access volumes, files or directories within a volume. As a result, ONTAP allows the customers to configure storage capacity for their users or groups of users of their Linux or Windows systems. The user or group policy quota limits the amount of space the user can utilize for their own data...A soft limit of this quota allows proactive notification to the user when the amount of capacity used within the volume is reaching the total capacity quota. Monitoring the amount of data stored within a user or group quota ensures that the user receives uninterrupted data service.|If warning threshold is breached, plan to take the following immediate actions:
- 1. Increase the space of the user or group quota in order to accommodate the growth.
- 2. Delete unwanted data to free up space.
- |Volume Capacity Full|CRITICAL|Storage capacity of a volume is necessary to store application and customer data. The more data stored in the ONTAP volume the less storage availability for future data. If the data storage capacity within a volume reaches the total storage capacity may lead to the customer being unable to store data due to lack of storage capacity. Monitoring the volume used storage capacity ensures data services continuity.|If critical threshold is breached, consider following immediate actions to minimize service disruption:
- 1. Increase the space of the volume to accommodate the growth.

- 2. Delete unwanted data to free up space.
- 3. If snapshot copies occupy more space than the snapshot reserve, delete old Snapshots or enable Volume Snapshot Autodelete....If warning threshold is breached, plan to take the following immediate actions:
- 1. Increase the space of the volume in order to accommodate the growth
- 2. If snapshot copies occupy more space than the snapshot reserve, delete old Snapshots or enabling Volume Snapshot Autodelete......

|Volume Inodes Limit|CRITICAL|Volumes that store files use index nodes (inode) to store file metadata. When a volume exhausts its inode allocation, no more files can be added to it...A warning alert indicates that planned action should be taken to increase the number of available inodes...A critical alert indicates that file limit exhaustion is imminent and emergency measures should be taken to free up inodes to ensure service continuity.|If critical threshold is breached, consider following immediate actions to minimize service disruption:

- 1. Increase the inodes value for the volume. If the inodes value is already at the max value, then split the volume into two or more volumes because the file system has grown beyond the maximum size.
- 2. Use FlexGroup as it helps to accommodate large file systems....

 If warning threshold is breached, plan to take the following immediate actions:
- 1. Increase the inodes value for the volume. If the inodes value is already at the max, then split the volume into two or more volumes because the file system has grown beyond the maximum size.
- 2. Use FlexGroup as it helps to accommodate large file systems |Volume Latency High|CRITICAL|Volumes are objects that serve the I/O traffic often driven by performance sensitive applications including devOps applications, home directories, and databases. High volume latencies means that the applications themselves may suffer and be unable to accomplish their tasks. Monitoring volume latencies is critical to maintain application consistent performance. The following are expected latencies based on media type SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.|If critical threshold is breached, consider following immediate actions to minimize service disruption:

If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled.... If warning threshold is breached, consider the following immediate actions:

- 1. If aggregate is also experiencing high utilization, move the volume to another aggregate.
- 2. If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled.
- 3. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node.

|Monitor Name|Severity|Monitor Description|Corrective Action |Node High Latency|WARNING / CRITICAL|Node latency has reached the levels where it might affect the performance of the applications on the node. Lower node latency ensures consistent performance of the applications. The expected latencies based on media type are: SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.|If critical threshold is breached, then immediate actions should be taken to minimize service disruption:

- 1. Suspend scheduled tasks, Snapshots or SnapMirror replication
- 2. Lower the demand of lower priority workloads via QoS limits
- 3. Inactivate non-essential workloads

Consider immediate actions when warning threshold is breached:

- 1. Move one or more workloads to a different storage location
- 2. Lower the demand of lower priority workloads via QoS limits
- 3. Add more storage nodes (AFF) or disk shelves (FAS) and redistribute workloads
- 4. Change workload characteristics (block size, application caching etc) |Node Performance Limit|WARNING / CRITICAL|Node performance utilization has reached the levels where it might affect the performance of the IOs and the applications supported by the node. Low node performance utilization ensures consistent performance of the applications.|Immediate actions should be taken to minimize service disruption if critical threshold is breached:
- 1. Suspend scheduled tasks, Snapshots or SnapMirror replication
- 2. Lower the demand of lower priority workloads via QoS limits
- 3. Inactivate non-essential workloads

Consider the following actions if warning threshold is breached:

- 1. Move one or more workloads to a different storage location
- 2. Lower the demand of lower priority workloads via QoS limits
- 3. Add more storage nodes (AFF) or disk shelves (FAS) and redistribute workloads ${}^{\prime}$
- 4. Change workload characteristics (block size, application caching etc) |Storage VM High Latency|WARNING / CRITICAL|Storage VM (SVM) latency has reached the levels where it might affect the performance of the applications on the storage VM. Lower storage VM latency ensures consistent performance of the applications. The expected latencies based on media type are: SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.|If critical threshold is breached, then immediately evaluate the threshold limits for volumes of the storage VM with a QoS policy assigned, to verify whether they are causing the volume workloads to get throttled

Consider following immediate actions when warning threshold is breached:

1. If aggregate is also experiencing high utilization, move some volumes

of the storage VM to another aggregate.

- 2. For volumes of the storage VM with a QoS policy assigned, evaluate the threshold limits if they are causing the volume workloads to get throttled
- 3. If the node is experiencing high utilization, move some volumes of the storage VM to another node or reduce the total workload of the node |User Quota Files Hard Limit|CRITICAL|The number of files created within the volume has reached the critical limit and additional files cannot be created. Monitoring the number of files stored ensures that the user receives uninterrupted data service.|Immediate actions are required to minimize service disruption if critical threshold is breached...Consider taking following actions:
- 1. Increase the file count quota for the specific user
- 2. Delete unwanted files to reduce the pressure on the files quota for the specific user

|User Quota Files Soft Limit|WARNING|The number of files created within the volume has reached the threshold limit of the quota and is near to the critical limit. You cannot create additional files if quota reaches the critical limit. Monitoring the number of files stored by a user ensures that the user receives uninterrupted data service.|Consider immediate actions if warning threshold is breached:

- 1. Increase the file count quota for the specific user quota
- 2. Delete unwanted files to reduce the pressure on the files quota for the specific user

|Volume Cache Miss Ratio|WARNING / CRITICAL|Volume Cache Miss Ratio is the percentage of read requests from the client applications that are returned from the disk instead of being returned from the cache. This means that the volume has reached the set threshold.|If critical threshold is breached, then immediate actions should be taken to minimize service disruption:

- 1. Move some workloads off of the node of the volume to reduce the IO load
- 2. If not already on the node of the volume, increase the WAFL cache by purchasing and adding a Flash Cache
- 3. Lower the demand of lower priority workloads on the same node via QoS limits

Consider immediate actions when warning threshold is breached:

- 1. Move some workloads off of the node of the volume to reduce the IO load
- 2. If not already on the node of the volume, increase the WAFL cache by purchasing and adding a Flash Cache
- 3. Lower the demand of lower priority workloads on the same node via QoS limits
- 4. Change workload characteristics (block size, application caching etc) |Volume Qtree Quota Overcommit|WARNING / CRITICAL|Volume Qtree Quota Overcommit specifies the percentage at which a volume is considered to be overcommitted by the qtree quotas. The set threshold for the qtree quota is reached for the volume. Monitoring the volume qtree quota overcommit

ensures that the user receives uninterrupted data service. If critical threshold is breached, then immediate actions should be taken to minimize service disruption:

- 1. Increase the space of the volume
- 2. Delete unwanted data

When warning threshold is breached, then consider increasing the space of the volume.

|===

<<top, Back to Top>>

=== Log Monitors

|===

|Monitor Name|Severity|Description|Corrective Action

|AWS Credentials Not Initialized|INFO|This event occurs when a module attempts to access Amazon Web Services (AWS) Identity and Access Management (IAM) role-based credentials from the cloud credentials thread before they are initialized. |Wait for the cloud credentials thread, as well as the system, to complete initialization.

|Cloud Tier Unreachable|CRITICAL|A storage node cannot connect to Cloud Tier object store API. Some data will be inaccessible.|If you use on-premises products, perform the following corrective actions: ...Verify that your intercluster LIF is online and functional by using the "network interface show" command....Check the network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF....Ensure the following:...The configuration of your object store has not changed....The login and connectivity information is still valid....Contact NetApp technical support if the issue persists.

If you use Cloud Volumes ONTAP, perform the following corrective actions: ...Ensure that the configuration of your object store has not changed... Ensure that the login and connectivity information is still valid....Contact NetApp technical support if the issue persists.

|Disk Out of Service|INFO|This event occurs when a disk is removed from service because it has been marked failed, is being sanitized, or has entered the Maintenance Center.|None.

|FlexGroup Constituent Full|CRITICAL|A constituent within a FlexGroup volume is full, which might cause a potential disruption of service. You can still create or expand files on the FlexGroup volume. However, none of the files that are stored on the constituent can be modified. As a result, you might see random out-of-space errors when you try to perform write operations on the FlexGroup volume.|It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X"

command....Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent. |Flexgroup Constituent Nearly Full|WARNING|A constituent within a FlexGroup volume is nearly out of space, which might cause a potential disruption of service. Files can be created and expanded. However, if the constituent runs out of space, you might not be able to append to or modify the files on the constituent. |It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X" command....Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent. |FlexGroup Constituent Nearly Out of Inodes|WARNING|A constituent within a FlexGroup volume is almost out of inodes, which might cause a potential disruption of service. The constituent receives lesser create requests than average. This might impact the overall performance of the FlexGroup volume, because the requests are routed to constituents with more inodes. | It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X" command....Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent.

|FlexGroup Constituent Out of Inodes|CRITICAL|A constituent of a FlexGroup volume has run out of inodes, which might cause a potential disruption of service. You cannot create new files on this constituent. This might lead to an overall imbalanced distribution of content across the FlexGroup volume.|It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X" command....Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent.

|LUN Offline|INFO|This event occurs when a LUN is brought offline manually. |Bring the LUN back online.

|Main Unit Fan Failed|WARNING|One or more main unit fans have failed. The system remains operational....However, if the condition persists for too long, the overtemperature might trigger an automatic shutdown.|Reseat the failed fans. If the error persists, replace them.

|Main Unit Fan in Warning State|INFO|This event occurs when one or more main unit fans are in a warning state.|Replace the indicated fans to avoid overheating.

| NVRAM Battery Low | WARNING | The NVRAM battery capacity is critically low. There might be a potential data loss if the battery runs out of power....Your system generates and transmits an AutoSupport or "call home" message to NetApp technical support and the configured destinations if it is configured to do so. The successful delivery of an AutoSupport message significantly improves problem determination and resolution. | Perform the following corrective actions:...View the battery's current status, capacity, and charging state by using the "system node environment sensors show" command....If the battery was replaced recently or the system was non-operational for an extended period of time, monitor the battery to verify

that it is charging properly....Contact NetApp technical support if the battery runtime continues to decrease below critical levels, and the storage system shuts down automatically.

|Service Processor Not Configured|WARNING|This event occurs on a weekly basis, to remind you to configure the Service Processor (SP). The SP is a physical device that is incorporated into your system to provide remote access and remote management capabilities. You should configure the SP to use its full functionality. |Perform the following corrective actions:...Configure the SP by using the "system service-processor network modify" command....Optionally, obtain the MAC address of the SP by using the "system service-processor network show" command....Verify the SP network configuration by using the "system service-processor network show" command....Verify that the SP can send an AutoSupport email by using the "system service-processor autosupport invoke" command.

NOTE: AutoSupport email hosts and recipients should be configured in ONTAP before you issue this command.

|Service Processor Offline|CRITICAL|ONTAP is no longer receiving heartbeats from the Service Processor (SP), even though all the SP recovery actions have been taken. ONTAP cannot monitor the health of the hardware without the SP....The system will shut down to prevent hardware damage and data loss. Set up a panic alert to be notified immediately if the SP goes offline. |Power-cycle the system by performing the following actions:...Pull the controller out from the chassis...Push the controller back in....Turn the controller back on....If the problem persists, replace the controller module.

|Shelf Fans Failed|CRITICAL|The indicated cooling fan or fan module of the shelf has failed. The disks in the shelf might not receive enough cooling airflow, which might result in disk failure.|Perform the following corrective actions:...Verify that the fan module is fully seated and secured.

NOTE: The fan is integrated into the power supply module in some disk shelves...If the issue persists, replace the fan module...If the issue still persists, contact NetApp technical support for assistance.

|System Cannot Operate Due to Main Unit Fan Failure |CRITICAL|One or more main unit fans have failed, disrupting system operation. This might lead to a potential data loss. |Replace the failed fans.

|Unassigned Disks|INFO|System has unassigned disks - capacity is being wasted and your system may have some misconfiguration or partial configuration change applied.|Perform the following corrective actions:...Determine which disks are unassigned by using the "disk show -n" command...Assign the disks to a system by using the "disk assign" command.

|Antivirus Server Busy|WARNING|The antivirus server is too busy to accept any new scan requests.|If this message occurs frequently, ensure that there are enough antivirus servers to handle the virus scan load generated by the SVM.

|AWS Credentials for IAM Role Expired|CRITICAL|Cloud Volume ONTAP has become inaccessible. The Identity and Access Management (IAM) role-based credentials have expired. The credentials are acquired from the Amazon Web Services (AWS) metadata server using the IAM role, and are used to sign API requests to Amazon Simple Storage Service (Amazon S3). | Perform the following:...Log in to the AWS EC2 Management Console...Navigate to the Instances page...Find the instance for the Cloud Volumes ONTAP deployment and check its health....Verify that the AWS IAM role associated with the instance is valid and has been granted proper privileges to the instance. |AWS Credentials for IAM Role Not Found|CRITICAL|The cloud credentials thread cannot acquire the Amazon Web Services (AWS) Identity and Access Management (IAM) role-based credentials from the AWS metadata server. The credentials are used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible...|Perform the following:...Log in to the AWS EC2 Management Console...Navigate to the Instances page....Find the instance for the Cloud Volumes ONTAP deployment and check its health...Verify that the AWS IAM role associated with the instance is valid and has been granted proper privileges to the instance. |AWS Credentials for IAM Role Not Valid|CRITICAL|The Identity and Access Management (IAM) role-based credentials are not valid. The credentials are acquired from the Amazon Web Services (AWS) metadata server using the IAM role, and are used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible. | Perform the following:...Log in to the AWS EC2 Management Console...Navigate to the Instances page...Find the instance for the Cloud Volumes ONTAP deployment and check its health...Verify that the AWS IAM role associated with the instance is valid and has been granted proper privileges to the instance. | AWS IAM Role Not Found | CRITICAL | The Identity and Access Management (IAM) roles thread cannot find an Amazon Web Services (AWS) IAM role on the AWS metadata server. The IAM role is required to acquire role-based credentials used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible...|Perform the following: ... Log in to the AWS EC2 Management Console. ... Navigate to the Instances page...Find the instance for the Cloud Volumes ONTAP deployment and check its health....Verify that the AWS IAM role associated with the instance is valid.

|AWS IAM Role Not Valid|CRITICAL|The Amazon Web Services (AWS) Identity and Access Management (IAM) role on the AWS metadata server is not valid. The Cloud Volume ONTAP has become inaccessible...|Perform the following:...Log in to the AWS EC2 Management Console...Navigate to the Instances page...Find the instance for the Cloud Volumes ONTAP deployment and check its health...Verify that the AWS IAM role associated with the instance is valid and has been granted proper privileges to the instance. |AWS Metadata Server Connection Fail|CRITICAL|The Identity and Access Management (IAM) roles thread cannot establish a communication link with the Amazon Web Services (AWS) metadata server. Communication should be

established to acquire the necessary AWS IAM role-based credentials used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible....|Perform the following:...Log in to the AWS EC2 Management Console....Navigate to the Instances page....Find the instance for the Cloud Volumes ONTAP deployment and check its health.... |FabricPool Space Usage Limit Nearly Reached|WARNING|The total clusterwide FabricPool space usage of object stores from capacity-licensed providers has nearly reached the licensed limit.|Perform the following corrective actions:...Check the percentage of the licensed capacity used by each FabricPool storage tier by using the "storage aggregate object-store show-space" command...Delete Snapshot copies from volumes with the tiering policy "snapshot" or "backup" by using the "volume snapshot delete" command to clear up space...Install a new license on the cluster to increase the licensed capacity.

|FabricPool Space Usage Limit Reached|CRITICAL|The total cluster-wide FabricPool space usage of object stores from capacity-licensed providers has reached the license limit.|Perform the following corrective actions:...Check the percentage of the licensed capacity used by each FabricPool storage tier by using the "storage aggregate object-store show-space" command...Delete Snapshot copies from volumes with the tiering policy "snapshot" or "backup" by using the "volume snapshot delete" command to clear up space...Install a new license on the cluster to increase the licensed capacity.

|Giveback of Aggregate Failed|CRITICAL|This event occurs during the migration of an aggregate as part of a storage failover (SFO) giveback, when the destination node cannot reach the object stores. |Perform the following corrective actions:...Verify that your intercluster LIF is online and functional by using the "network interface show" command....Check network connectivity to the object store server by using the "'ping" command over the destination node intercluster LIF. ...Verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the "aggregate object-store config show" command....Alternatively, you can override the error by specifying false for the "require-partner-waiting" parameter of the giveback command....Contact NetApp technical support for more information or assistance.

| HA Interconnect Down | WARNING | The high-availability (HA) interconnect is down. Risk of service outage when failover is not available. | Corrective actions depend on the number and type of HA interconnect links supported by the platform, as well as the reason why the interconnect is down. ... If the links are down: ... Verify that both controllers in the HA pair are operational. ... For externally connected links, make sure that the interconnect cables are connected properly and that the small form-factor pluggables (SFPs), if applicable, are seated properly on both controllers. ... For internally connected links, disable and re-enable the links, one after the other, by using the "ic link off" and "ic link on"

commands. ...If links are disabled, enable the links by using the "ic link on" command. ...If a peer is not connected, disable and re-enable the links, one after the other, by using the "ic link off" and "ic link on" commands....Contact NetApp technical support if the issue persists.

|Max Sessions Per User Exceeded | WARNING

|You have exceeded the maximum number of sessions allowed per user over a TCP connection. Any request to establish a session will be denied until some sessions are released. ...|Perform the following corrective actions: ...Inspect all the applications that run on the client, and terminate any that are not operating properly...Reboot the client...Check if the issue is caused by a new or existing application:...If the application is new, set a higher threshold for the client by using the "cifs option modify -max -opens-same-file-per-tree" command.

In some cases, clients operate as expected, but require a higher threshold. You should have advanced privilege to set a higher threshold for the client. ...If the issue is caused by an existing application, there might be an issue with the client. Contact NetApp technical support for more information or assistance.

|Max Times Open Per File Exceeded|WARNING|You have exceeded the maximum number of times that you can open the file over a TCP connection. Any request to open this file will be denied until you close some open instances of the file. This typically indicates abnormal application behavior...|Perform the following corrective actions:...Inspect the applications that run on the client using this TCP connection.

The client might be operating incorrectly because of the application running on it...Reboot the client...Check if the issue is caused by a new or existing application:...If the application is new, set a higher threshold for the client by using the "cifs option modify -max-opens-same-file-per -tree" command.

In some cases, clients operate as expected, but require a higher threshold. You should have advanced privilege to set a higher threshold for the client. ...If the issue is caused by an existing application, there might be an issue with the client. Contact NetApp technical support for more information or assistance.

|NetBIOS Name Conflict|CRITICAL

|The NetBIOS Name Service has received a negative response to a name registration request, from a remote machine. This is typically caused by a conflict in the NetBIOS name or an alias. As a result, clients might not be able to access data or connect to the right data-serving node in the cluster.|Perform any one of the following corrective actions:...If there is a conflict in the NetBIOS name or an alias, perform one of the following:...Delete the duplicate NetBIOS alias by using the "vserver cifs delete -aliases alias -vserver vserver" command....Rename a NetBIOS alias by deleting the duplicate name and adding an alias with a new name by using the "vserver cifs create -aliases alias -vserver vserver" command. ...If there are no aliases configured and there is a conflict in the NetBIOS

name, then rename the CIFS server by using the "vserver cifs delete -vserver vserver" and "vserver cifs create -cifs-server netbiosname" commands.

NOTE: Deleting a CIFS server can make data inaccessible. ...Remove NetBIOS name or rename the NetBIOS on the remote machine.

|NFSv4 Store Pool Exhausted|CRITICAL|A NFSv4 store pool has been exhausted.|If the NFS server is unresponsive for more than 10 minutes after this event, contact NetApp technical support.

|No Registered Scan Engine|CRITICAL|The antivirus connector notified ONTAP that it does not have a registered scan engine. This might cause data unavailability if the "scan-mandatory" option is enabled. |Perform the following corrective actions:...Ensure that the scan engine software installed on the antivirus server is compatible with ONTAP....Ensure that scan engine software is running and configured to connect to the antivirus connector over local loopback.

|No Vscan Connection|CRITICAL|ONTAP has no Vscan connection to service virus scan requests. This might cause data unavailability if the "scanmandatory" option is enabled.|Ensure that the scanner pool is properly configured and the antivirus servers are active and connected to ONTAP.
|Node Root Volume Space Low|CRITICAL|The system has detected that the root volume is dangerously low on space. The node is not fully operational.

Data LIFs might have failed over within the cluster, because of which NFS and CIFS access is limited on the node. Administrative capability is limited to local recovery procedures for the node to clear up space on the root volume.|Perform the following corrective actions:...Clear up space on the root volume by deleting old Snapshot copies, deleting files you no longer need from the /mroot directory, or expanding the root volume capacity...Reboot the controller....Contact NetApp technical support for more information or assistance.

|Nonexistent Admin Share|CRITICAL|Vscan issue: a client has attempted to connect to a nonexistent ONTAP_ADMIN\$ share. |Ensure that Vscan is enabled for the mentioned SVM ID. Enabling Vscan on a SVM causes the ONTAP_ADMIN\$ share to be created for the SVM automatically.

| NVMe Namespace Out of Space | CRITICAL | An NVMe namespace has been brought offline because of a write failure caused by lack of space. | Add space to the volume, and then bring the NVMe namespace online by using the "vserver nvme namespace modify" command.

|NVMe-oF Grace Period Active|WARNING|This event occurs on a daily basis when the NVMe over Fabrics (NVMe-oF) protocol is in use and the grace period of the license is active. The NVMe-oF functionality requires a license after the license grace period expires. NVMe-oF functionality is disabled when the license grace period is over. |Contact your sales representative to obtain an NVMe-oF license, and add it to the cluster, or remove all instances of NVMe-oF configuration from the cluster. |NVMe-oF Grace Period Expired|WARNING|The NVMe over Fabrics (NVMe-oF)

license grace period is over and the NVMe-oF functionality is

disabled. | Contact your sales representative to obtain an NVMe-oF license, and add it to the cluster.

|NVMe-oF Grace Period Start|WARNING|The NVMe over Fabrics (NVMe-oF) configuration was detected during the upgrade to ONTAP 9.5 software. NVMe-oF functionality requires a license after the license grace period expires.|Contact your sales representative to obtain an NVMe-oF license, and add it to the cluster.

|Object Store Host Unresolvable|CRITICAL|The object store server host name cannot be resolved to an IP address. The object store client cannot communicate with the object-store server without resolving to an IP address. As a result, data might be inaccessible. |Check the DNS configuration to verify that the host name is configured correctly with an IP address.

|Object Store Intercluster LIF Down|CRITICAL|The object-store client cannot find an operational LIF to communicate with the object store server. The node will not allow object store client traffic until the intercluster LIF is operational. As a result, data might be inaccessible. |Perform the following corrective actions:...Check the intercluster LIF status by using the "network interface show -role intercluster" command...Verify that the intercluster LIF is configured correctly and operational...If an intercluster LIF is not configured, add it by using the "network interface create -role intercluster" command.

|Object Store Signature Mismatch|CRITICAL|The request signature sent to the object store server does not match the signature calculated by the client. As a result, data might be inaccessible. |Verify that the secret access key is configured correctly. If it is configured correctly, contact NetApp technical support for assistance.

|READDIR Timeout|CRITICAL|A READDIR file operation has exceeded the timeout that it is allowed to run in WAFL. This can be because of very large or sparse directories. Corrective action is recommended. |Perform the following corrective actions:...Find information specific to recent directories that have had READDIR file operations expire by using the following 'diag' privilege nodeshell CLI command:

wafl readdir notice show....Check if directories are indicated as sparse or not:...If a directory is indicated as sparse, it is recommended that you copy the contents of the directory to a new directory to remove the sparseness of the directory file. ...If a directory is not indicated as sparse and the directory is large, it is recommended that you reduce the size of the directory file by reducing the number of file entries in the directory.

|Relocation of Aggregate Failed|CRITICAL|This event occurs during the relocation of an aggregate, when the destination node cannot reach the object stores. |Perform the following corrective actions:...Verify that your intercluster LIF is online and functional by using the "network interface show" command....Check network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF.

...Verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the "aggregate object-store config show" command....Alternatively, you can override the error by using the "override-destination-checks" parameter of the relocation command....Contact NetApp technical support for more information or assistance.

|Shadow Copy Failed|CRITICAL|A Volume Shadow Copy Service (VSS), a
Microsoft Server backup and restore service operation, has failed.|Check
the following using the information provided in the event message:...Is
shadow copy configuration enabled?...Are the appropriate licenses installed?
...On which shares is the shadow copy operation performed?...Is the share name
correct?...Does the share path exist?...What are the states of the shadow copy
set and its shadow copies?

|Storage Switch Power Supplies Failed | WARNING | There is a missing power supply in the cluster switch. Redundancy is reduced, risk of outage with any further power failures. | Perform the following corrective actions: ... Ensure that the power supply mains, which supplies power to the cluster switch, is turned on...Ensure that the power cord is connected to the power supply....Contact NetApp technical support if the issue persists. |Too Many CIFS Authentication|WARNING|Many authentication negotiations have occurred simultaneously. There are 256 incomplete new session requests from this client.|Investigate why the client has created 256 or more new connection requests. You might have to contact the vendor of the client or of the application to determine why the error occurred. |Unauthorized User Access to Admin Share|WARNING|A client has attempted to connect to the privileged ONTAP ADMIN\$ share even though their logged-in user is not an allowed user. | Perform the following corrective actions: ... Ensure that the mentioned username and IP address is configured in one of the active Vscan scanner pools....Check the scanner pool configuration that is currently active by using the "vserver vscan scanner pool show-active" command.

|Virus Detected|WARNING|A Vscan server has reported an error to the storage system. This typically indicates that a virus has been found. However, other errors on the Vscan server can cause this event....Client access to the file is denied. The Vscan server might, depending on its settings and configuration, clean the file, quarantine it, or delete it.|Check the log of the Vscan server reported in the "syslog" event to see if it was able to successfully clean, quarantine, or delete the infected file. If it was not able to do so, a system administrator might have to manually delete the file.

|Volume Offline|INFO|This message indicates that a volume is made offline.|Bring the volume back online.

|Volume Restricted|INFO|This event indicates that a flexible volume is made restricted.|Bring the volume back online.

|Storage VM Stop Succeeded|INFO|This message occurs when a 'vserver stop'

operation succeeds. | Use 'vserver start' command to start the data access on a storage VM.

|Node Panic|WARNING|This event is issued when a panic occurs|Contact NetApp customer support.|1 day

|===

<<top, Back to Top>>

=== Anti-Ransomware Log Monitors

|===

|Monitor Name|Severity|Description|Corrective Action |Storage VM Anti-ransomware Monitoring Disabled|WARNING|The antiransomware monitoring for the storage VM is disabled. Enable antiransomware to protect the storage VM.|None

|Storage VM Anti-ransomware Monitoring Enabled (Learning Mode)|INFO|The anti-ransomware monitoring for the storage VM is enabled in learning mode.|None

|Volume Anti-ransomware Monitoring Enabled|INFO|The anti-ransomware monitoring for the volume is enabled.|None

|Volume Anti-ransomware Monitoring Disabled|WARNING|The anti-ransomware monitoring for the volume is disabled. Enable anti-ransomware to protect the volume.|None

|Volume Anti-ransomware Monitoring Enabled (Learning Mode) |INFO|The anti-ransomware monitoring for the volume is enabled in learning mode.|None |Volume Anti-ransomware Monitoring Paused (Learning Mode) |WARNING|The anti-ransomware monitoring for the volume is paused in learning mode.|None |Volume Anti-ransomware Monitoring Paused|WARNING|The anti-ransomware monitoring for the volume is paused.|None

|Volume Anti-ransomware Monitoring Disabling|WARNING|The anti-ransomware monitoring for the volume is disabling.|None

|Ransomware Activity Detected|CRITICAL|To protect the data from the detected ransomware, a Snapshot copy has been taken that can be used to restore original data.

Your system generates and transmits an AutoSupport or "call home" message to NetApp technical support and any configured destinations. AutoSupport message improves problem determination and resolution. |Refer to the "FINAL-DOCUMENT-NAME" to take remedial measures for ransomware activity.

|===

<<top, Back to Top>>

=== FSx for NetApp ONTAP Monitors

|===

|Monitor Name|Thresholds|Monitor Description|Corrective Action |FSx Volume Capacity is Full|Warning @ > 85 %...Critical @ > 95 %|Storage capacity of a volume is necessary to store application and customer data. The more data stored in the ONTAP volume the less storage availability for future data. If the data storage capacity within a volume reaches the total storage capacity may lead to the customer being unable to store data due to lack of storage capacity. Monitoring the volume used storage capacity ensures data services continuity. | Immediate actions are required to minimize service disruption if critical threshold is breached:...1. Consider deleting data that is not needed anymore to free up space |FSx Volume High Latency|Warning @ > 1000 µs...Critical @ > 2000 µs|Volumes are objects that serve the IO traffic often driven by performance sensitive applications including devOps applications, home directories, and databases. High volume latencies means that the applications themselves may suffer and be unable to accomplish their tasks. Monitoring volume latencies is critical to maintain application consistent performance. | Immediate actions are required to minimize service disruption if critical threshold is breached:...1. If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled.....Plan to take the following actions soon if warning threshold is breached: ... If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled ... 2. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node.

|FSx Volume Inodes Limit|Warning @ > 85 %...Critical @ > 95 %|Volumes that store files use index nodes (inode) to store file metadata. When a volume exhausts its inode allocation no more files can be added to it. A warning alert indicates that planned action should be taken to increase the number of available inodes. A critical alert indicates that file limit exhaustion is imminent and emergency measures should be taken to free up inodes to ensure service continuity|Immediate actions are required to minimize service disruption if critical threshold is breached:...1. Consider increasing the inodes value for the volume. If the inodes value is already at the max, then consider splitting the volume into two or more volumes because the file system has grown beyond the maximum size.....Plan to take the following actions soon if warning threshold is breached:...1. Consider increasing the inodes value for the volume. If the inodes value is already

at the max, then consider splitting the volume into two or more volumes because the file system has grown beyond the maximum size |FSx Volume Qtree Quota Overcommit|Warning @ > 95 %...Critical @ > 100 %|Volume Qtree Quota Overcommit specifies the percentage at which a volume is considered to be overcommitted by the qtree quotas. The set threshold for the qtree quota is reached for the volume. Monitoring the volume qtree quota overcommit ensures that the user receives uninterrupted data service.|If critical threshold is breached, then immediate actions should be taken to minimize service disruption:

1. Delete unwanted data...When warning threshold is breached, then consider increasing the space of the volume.

|FSx Snapshot Reserve Space is Full|Warning @ > 90 %...Critical @ > 95 %|Storage capacity of a volume is necessary to store application and customer data. A portion of that space, called snapshot reserved space, is used to store snapshots which allow data to be protected locally. The more new and updated data stored in the ONTAP volume the more snapshot capacity is used and less snapshot storage capacity will be available for future new or updated data. If the snapshot data capacity within a volume reaches the total snapshot reserve space it may lead to the customer being unable to store new snapshot data and reduction in the level of protection for the data in the volume. Monitoring the volume used snapshot capacity ensures data services continuity. | Immediate actions are required to minimize service disruption if critical threshold is breached:...1. Consider configuring snapshots to use data space in the volume when the snapshot reserve is full...2. Consider deleting some older snapshots that may not be needed anymore to free up space.....Plan to take the following actions soon if warning threshold is breached:...1. Consider increasing the snapshot reserve space within the volume to accommodate the growth...2. Consider configuring snapshots to use data space in the volume when the snapshot reserve is full

|FSx Volume Cache Miss Ratio|Warning @ > 95 %...Critical @ > 100 %|Volume Cache Miss Ratio is the percentage of read requests from the client applications that are returned from the disk instead of being returned from the cache. This means that the volume has reached the set threshold.|If critical threshold is breached, then immediate actions should be taken to minimize service disruption:

- 1. Move some workloads off of the node of the volume to reduce the ${\tt IO}$ load
- 2. Lower the demand of lower priority workloads on the same node via QoS limits...Consider immediate actions when warning threshold is breached:
- 1. Move some workloads off of the node of the volume to reduce the IO load
- 2. Lower the demand of lower priority workloads on the same node via QoS limits
- 3. Change workload characteristics (block size, application caching etc)

|===

<<top, Back to Top>> === K8s Monitors |=== |Monitor Name|Description|Corrective Actions|Severity/Threshold |Persistent Volume Latency High |High persistent volume latencies means that the applications themselves may suffer and be unable to accomplish their tasks. Monitoring persistent volume latencies is critical to maintain application consistent performance. The following are expected latencies based on media type -SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20milliseconds. | **Immediate Actions ** If critical threshold is breached, consider immediate actions to minimize service disruption: If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled. **Actions To Do Soon** If warning threshold is breached, plan the following immediate actions: 1. If storage pool is also experiencing high utilization, move the volume to another storage pool. 2. If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled. 3. If the controller is also experiencing high utilization, move the volume to another controller or reduce the total workload of the controller. |Warning @ > 6,000 μ s Critical @ > 12,000 μ s |Cluster Memory Saturation High |Cluster allocatable memory saturation is high. Cluster CPU saturation is calculated as the sum of memory usage divided by the sum of allocatable memory across all K8s nodes. |Add nodes. Fix any unscheduled nodes. Right-size pods to free up memory on nodes. |Warning @ > 80 % Critical @ > 90 % IPOD Attach Failed |This alert occurs when a volume attachment with POD is failed. |Warning

| High Retransmit Rate

| High TCP Retransmit Rate

|Check for Network congestion - Identify workloads that consume a lot of network bandwidth.

Check for high Pod CPU utilization.

Check hardware network performance.

|Warning @ > 10 %

Critical @ > 25 %

|Node File System Capacity High

|Node File System Capacity High

|- Increase the size of the node disks to ensure that there is sufficient room for the application files.

- Decrease application file usage.

|Warning @ > 80 %

Critical @ > 90 %

|Workload Network Jitter High

| High TCP Jitter (high latency/response time variations)

|Check for Network congestion. Identify workloads that consume a lot of network bandwidth.

Check for high Pod CPU utilization.

Check hardware network performance

|Warning @ > 30 ms

Critical @ > 50 ms

|Persistent Volume Throughput

IMBPS thresholds on persistent volumes can be used to alert an administrator when persistent volumes exceed predefined performance expectations, potentially impacting other persistent volumes. Activating this monitor will generate alerts appropriate for the typical throughput profile of persistent volumes on SSDs. This monitor will cover all persistent volumes in your environment. The warning and critical threshold values can be adjusted based on your monitoring goals by duplicating this monitor and setting thresholds appropriate for your storage class. A duplicated monitor can be further targeted to a subset of the persistent volumes in your environment.

| **Immediate Actions**

If critical threshold is breached, plan immediate actions to minimize service disruption:

- 1. Introduce QoS MBPS limits for the volume.
- 2. Review the application driving the workload on the volume for anomalies.
- **Actions To Do Soon**

If warning threshold is breached, plan to take the following immediate

```
actions:
1. Introduce QoS MBPS limits for the volume.
2. Review the application driving the workload on the volume for
anomalies.
|Warning @ > 10,000 MB/s
Critical @ > 15,000 MB/s
|Container at Risk of Going OOM Killed
|The container's memory limits are set too low. The container is at risk
of eviction (Out of Memory Kill).
|Increase container memory limits.
Warning @ > 95 %
|Workload Down
|Workload has no healthy pods.
|Critical @ < 1
|Persistent Volume Claim Failed Binding
|This alert occurs when a binding is failed on a PVC.
|Warning
|ResourceQuota Mem Limits About to Exceed
|Memory limits for Namespace are about to exceed ResourceQuota
|Warning @ > 80 %
 Critical @ > 90 %
|ResourceQuota Mem Requests About to Exceed
|Memory requests for Namespace are about to exceed ResourceQuota
|Warning @ > 80 %
 Critical @ > 90 %
|Node Creation Failed
|The node could not be scheduled because of a configuration error.
|Check the Kubernetes event log for the cause of the configuration
failure.
|Critical
|Persistent Volume Reclamation Failed
|The volume has failed its automatic reclamation.
|Warning @ > 0 B
```

```
|Container CPU Throttling
|The container's CPU Limits are set too low. Container processes are
slowed.
|Increase container CPU limits.
|Warning @ > 95 %
Critical @ > 98 %
|Service Load Balancer Failed to Delete
|Warning
|Persistent Volume IOPS
| IOPS thresholds on persistent volumes can be used to alert an
administrator when persistent volumes exceed predefined performance
expectations. Activating this monitor will generate alerts appropriate for
the typical IOPS profile of persistence volumes. This monitor will cover
all persistent volumes in your environment. The warning and critical
threshold values can be adjusted based on your monitoring goals by
duplicating this monitor and setting thresholds appropriate for your
workload.
| **Immediate Actions **
If critical threshold is breached, plan Immediate actions to minimize
service disruption :
1. Introduce QoS IOPS limits for the volume.
2. Review the application driving the workload on the volume for
anomalies.
**Actions To Do Soon**
If warning threshold is breached, plan the following immediate actions:
1. Introduce QoS IOPS limits for the volume.
2. Review the application driving the workload on the volume for
anomalies.
|Warning @ > 20,000 IO/s
Critical @ > 25,000 IO/s
|Service Load Balancer Failed to Update
|Warning
| POD Failed Mount
|This alert occurs when a mount is failed on a POD.
|Warning
```

```
|Node PID Pressure
|Available process identifiers on the (Linux) node has fallen below an
eviction threshold.
|Find and fix pods that generate many processes and starve the node of
available process IDs.
Set up PodPidsLimit to protect your node against pods or containers that
spawn too many processes.
|Critical @ > 0
| Pod Image Pull Failure
|Kubernetes failed to pull the pod container image.
|- Make sure the pod's image is spelled correctly in the pod
configuration.
- Check image tag exists in your registry.
- Verify the credentials for the image registry.
- Check for registry connectivity issues.
- Verify you are not hitting the rate limits imposed by public registry
providers.
|Warning
|Job Running Too Long
|Job is running for too long
|Warning @ > 1 hr
Critical @ > 5 hr
|Node Memory High
|Node memory usage is high
|Add nodes.
Fix any unscheduled nodes.
Right-size pods to free up memory on nodes.
|Warning @ > 85 %
Critical @ > 90 %
|ResourceQuota CPU Limits About to Exceed
|CPU limits for Namespace are about to exceed ResourceQuota
|Warning @ > 80 %
Critical @ > 90 %
| Pod Crash Loop Backoff
|Pod has crashed and attempted to restart multiple times.
|Critical @ > 3
|Node CPU High
```

```
|Node CPU usage is high.
|Add nodes.
Fix any unscheduled nodes.
Right-size pods to free up CPU on nodes.
|Warning @ > 80 %
Critical @ > 90 %
|Workload Network Latency RTT High
| High TCP RTT (Round Trip Time) latency
|Check for Network congestion | Identify workloads that consume a lot of
network bandwidth.
Check for high Pod CPU utilization.
Check hardware network performance.
|Warning @ > 150 ms
Critical @ > 300 ms
|Job Failed
|Job did not complete successfully due to a node crash or reboot, resource
exhaustion, job timeout, or pod scheduling failure.
|Check the Kubernetes event logs for failure causes.
|Warning @ > 1
|Persistent Volume Full in a Few Days
|Persistent Volume will run out of space in a few days
|-Increase the volume size to ensure that there is sufficient room for the
application files.
-Reduce the amount of data stored in applications.
|Warning @ < 8 day
Critical @ < 3 day
|Node Memory Pressure
|Node is running out of memory. Available memory has met eviction
threshold.
|Add nodes.
Fix any unscheduled nodes.
Right-size pods to free up memory on nodes.
|Critical @ > 0
|Node Unready
|Node has been unready for 5 minutes
|Verify the node have enough CPU, memory, and disk resources.
Check node network connectivity.
Check the Kubernetes event logs for failure causes.
|Critical @ < 1
|Persistent Volume Capacity High
```

```
|Persistent Volume backend used capacity is high.
|- Increase the volume size to ensure that there is sufficient room for
the application files.
- Reduce the amount of data stored in applications.
|Warning @ > 80 %
Critical @ > 90 %
|Service Load Balancer Failed to Create
|Service Load Balancer Create Failed
|Critical
|Workload Replica Mismatch
|Some pods are currently not available for a Deployment or DaemonSet.
|Warning @ > 1
|ResourceQuota CPU Requests About to Exceed
|CPU requests for Namespace are about to exceed ResourceQuota
|Warning @ > 80 %
Critical @ > 90 %
|High Retransmit Rate
| High TCP Retransmit Rate
|Check for Network congestion - Identify workloads that consume a lot of
network bandwidth.
Check for high Pod CPU utilization.
Check hardware network performance.
|Warning @ > 10 %
Critical @ > 25 %
|Node Disk Pressure
|Available disk space and inodes on either the node's root filesystem or
image filesystem has satisfied an eviction threshold.
|- Increase the size of the node disks to ensure that there is sufficient
room for the application files.
- Decrease application file usage.
|Critical @ > 0
|Cluster CPU Saturation High
|Cluster allocatable CPU saturation is high.
Cluster CPU saturation is calculated as the sum of CPU usage divided by
the sum allocatable CPU across all K8s nodes.
|Add nodes.
Fix any unscheduled nodes.
```

```
Right-size pods to free up CPU on nodes.
|Warning @ > 80 %
Critical @ > 90 %
|===
////
//Original:
|===
|Monitor Name|Severity|Monitor Description
| POD Created | Informational | This alert occurs when a POD is created.
| POD Deleted | Informational | This alert occurs when a POD is deleted.
|Daemonset Created|Informational|This alert occurs when a Daemonset is
created.
|Daemonset Deleted|Informational|This alert occurs when a Daemonset is
deleted.
|Replicaset Created|Informational|This alert occurs when a Replicaset is
created.
|Replicaset Deleted|Informational|This alert occurs when a Replicaset is
deleted.
|Deployment Created|Informational|This alert occurs when a Deployment is
created.
| POD Failed | WARNING | This alert occurs when a POD is failed.
| POD Attach Failed | WARNING | This alert occurs when a volume attachment with
POD is failed.
|Persistent Volume Claim Failed Binding|WARNING|This alert occurs when a
binding is failed on a PVC.
| POD Failed Mount | WARNING | This alert occurs when a mount is failed on a
POD.
|===
////
<<top, Back to Top>>
=== Change Log Monitors
|===
|Monitor Name|Severity|Monitor Description
|Internal Volume Discovered|Informational|This message occurs when an
Internal Volume is discovered.
|Internal Volume Modified|Informational|This message occurs when an
```

Internal Volume is modified.

|Storage Node Discovered|Informational|This message occurs when an Storage Node is discovered.

|Storage Node Removed|Informational|This message occurs when an Storage Node is removed.

|Storage Pool Discovered|Informational|This message occurs when an Storage Pool is discovered.

|Storage Virtual Machine Discovered|Informational|This message occurs when an Storage Virtual Machine is discovered.

 $|Storage\ Virtual\ Machine\ Modified | Informational | This\ message\ occurs\ when an\ Storage\ Virtual\ Machine\ is\ modified.$

|===

<<top, Back to Top>>

=== Data Collection Monitors

|===

Monitor Name | Description | Corrective Action

|Acquisition Unit Shutdown|Cloud Insights Acquisition Units periodically restart as part of upgrades to introduce new features. This happens once a month or less in a typical environment. A Warning Alert that an Acquisition Unit has shutdown should be followed soon after by a Resolution noting that the newly-restarted Acquisition Unit has completed a registration with Cloud Insights. Typically this shutdown-to-registration cycle takes 5 to 15 minutes. |If the alert occurs frequently or lasts longer than 15 minutes, check on the operation of the system hosting the Acquisition Unit, the network, and any proxy connecting the AU to the Internet.

|Collector Failed|The poll of a data collector encountered an unexpected failure situation.|Visit the data collector page in Cloud Insights to learn more about the situation.

|Collector Warning|This Alert typically can arise because of an erroneous configuration of the data collector or of the target system. Revisit the configurations to prevent future Alerts. It can also be due to a retrieval of less-than-complete data where the data collector gathered all the data that it could. This can happen when situations change during data collection (e.g., a virtual machine present at the beginning of data collection is deleted during data collection and before its data is captured).|Check the configuration of the data collector or target system.

Note that the monitor for Collector Warning can send more alerts than other monitor types, so it is recommended to set no alert recipients unless you are troubleshooting.

|===

<<top,Back to Top>>
=== Security Monitors
|===

|Monitor Name|Threshold|Monitor Description|Corrective Action |AutoSupport HTTPS transport disabled|Warning @ < 1|AutoSupport supports HTTPS, HTTP, and SMTP for transport protocols. Because of the sensitive nature of AutoSupport messages, NetApp strongly recommends using HTTPS as the default transport protocol for sending AutoSupport messages to NetApp support.|To set HTTPS as the transport protocol for AutoSupport messages, run the following ONTAP command:...system node autosupport modify -transport https

|Cluster Insecure ciphers for SSH|Warning @ < 1|Indicates that SSH is using insecure ciphers, for example ciphers beginning with *cbc.|To remove the CBC ciphers, run the following ONTAP command:...security ssh remove -vserver <admin vserver> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc

|Cluster Login Banner Disabled|Warning @ < 1|Indicates that the Login banner is disabled for users accessing the ONTAP system. Displaying a login banner is helpful for establishing expectations for access and use of the system.|To configure the login banner for a cluster, run the following ONTAP command:...security login banner modify -vserver <admin svm> -message "Access restricted to authorized users"

|Cluster Peer Communication Not Encrypted|Warning @ < 1|When replicating data for disaster recovery, caching, or backup, you must protect that data during transport over the wire from one ONTAP cluster to another.

Encryption must be configured on both the source and destination clusters.|To enable encryption on cluster peer relationships that were created prior to ONTAP 9.6, the source and destination cluster must be upgraded to 9.6. Then use the "cluster peer modify" command to change both the source and destination cluster peers to use Cluster Peering Encryption....See the NetApp Security Hardening Guide for ONTAP 9 for details.

|Default Local Admin User Enabled|Warning @ > 0|NetApp recommends locking (disabling) any unneeded Default Admin User (built-in) accounts with the lock command. They are primarily default accounts for which passwords were never updated or changed.|To lock the built-in "admin" account, run the following ONTAP command:...security login lock -username admin |FIPS Mode Disabled|Warning @ < 1|When FIPS 140-2 compliance is enabled, TLSv1 and SSLv3 are disabled, and only TLSv1.1 and TLSv1.2 remain enabled. ONTAP prevents you from enabling TLSv1 and SSLv3 when FIPS 140-2

compliance is enabled. | To enable FIPS 140-2 compliance on a cluster, run the following ONTAP command in advanced privilege mode: ... security config modify -interface SSL -is-fips-enabled true

|Log Forwarding Not Encrypted|Warning @ < 1|Offloading of syslog information is necessary for limiting the scope or footprint of a breach to a single system or solution. Therefore, NetApp recommends securely offloading syslog information to a secure storage or retention location.|Once a log forwarding destination is created, its protocol cannot be changed. To change to an encrypted protocol, delete and recreate the log forwarding destination using the following ONTAP command:...cluster log-forwarding create -destination <destination ip> -protocol tcp-encrypted

|MD5 Hashed password|Warning @ > 0|NetApp strongly recommends to use the more secure SHA-512 hash function for ONTAP user account passwords.

Accounts using the less secure MD5 hash function should migrate to the SHA-512 hash function.|NetApp strongly recommends user accounts migrate to the more secure SHA-512 solution by having users change their passwords...to lock accounts with passwords that use the MD5 hash function, run the following ONTAP command:...security login lock -vserver * -username * -hash-function md5

|No NTP servers are configured|Warning @ < 1|Indicates that the cluster has no configured NTP servers. For redundancy and optimum service, NetApp recommends that you associate at least three NTP servers with the cluster.|To associate an NTP server with the cluster, run the following ONTAP command:

cluster time-service ntp server create -server <ntp server host name or ip
address>

|NTP server count is low|Warning @ < 3|Indicates that the cluster has less than 3 configured NTP servers. For redundancy and optimum service, NetApp recommends that you associate at least three NTP servers with the cluster.|To associate an NTP server with the cluster, run the following ONTAP command:...cluster time-service ntp server create -server <ntp server host name or ip address>

|Remote Shell Enabled|Warning @ > 0|Remote Shell is not a secure method for establishing command-line access to the ONTAP solution. Remote Shell should be disabled for secure remote access.|NetApp recommends Secure Shell (SSH) for secure remote access....To disable Remote shell on a cluster, run the following ONTAP command in advanced privilege mode:...security protocol modify -application rsh- enabled false |Storage VM Audit Log Disabled|Warning @ < 1|Indicates that Audit logging is disabled for SVM.|To configure the Audit log for a vserver, run the following ONTAP command:...vserver audit enable -vserver <svm> |Storage VM Insecure ciphers for SSH|Warning @ < 1|Indicates that SSH is using insecure ciphers, for example ciphers beginning with *cbc.|To remove the CBC ciphers, run the following ONTAP command:...security ssh remove

-vserver <vserver> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc |Storage VM Login banner disabled|Warning @ < 1|Indicates that the Login banner is disabled for users accessing SVMs on the system. Displaying a login banner is helpful for establishing expectations for access and use of the system.|To configure the login banner for a cluster, run the following ONTAP command:...security login banner modify -vserver <svm> -message "Access restricted to authorized users"

|Telnet Protocol Enabled|Warning @ > 0|Telnet is not a secure method for establishing command-line access to the ONTAP solution. Telnet should be disabled for secure remote access.|NetApp recommends Secure Shell (SSH) for secure remote access.

To disable Telnet on a cluster, run the following ONTAP command in advanced privilege mode:...security protocol modify -application telnet -enabled false

|===

<<top, Back to Top>>

=== Data Protection Monitors

|===

|Monitor Name|Thresholds|Monitor Description|Corrective Action
|Insufficient Space for Lun Snapshot Copy|(Filter contains_luns = Yes)
Warning @ > 95 %...Critical @ > 100 %|Storage capacity of a volume is
necessary to store application and customer data. A portion of that space,
called snapshot reserved space, is used to store snapshots which allow
data to be protected locally. The more new and updated data stored in the
ONTAP volume the more snapshot capacity is used and less snapshot storage
capacity will be available for future new or updated data. If the snapshot
data capacity within a volume reaches the total snapshot reserve space it
may lead to the customer being unable to store new snapshot data and
reduction in the level of protection for the data in the LUNs in the
volume. Monitoring the volume used snapshot capacity ensures data services
continuity.| **Immediate Actions**

If critical threshold is breached, consider immediate actions to minimize service disruption:

- 1. Configure snapshots to use data space in the volume when the snapshot reserve is full.
- 2. Delete some older unwanted snapshots to free up space.

Actions To Do Soon

If warning threshold is breached, plan to take the following immediate

actions:

- 1. Increase the snapshot reserve space within the volume to accommodate the growth.
- 2. Configure snapshots to use data space in the volume when the snapshot reserve is full.

|SnapMirror Relationship Lag|Warning @ > 150%...Critical @ > 300%|SnapMirror relationship lag is the difference between the snapshot timestamp and the time on the destination system. The lag_time_percent is the ratio of lag time to the SnapMirror Policy's schedule interval. If the lag time equals the schedule interval, the lag_time_percent will be 100%. If the SnapMirror policy does not have a schedule, lag_time_percent will not be calculated.|Monitor SnapMirror status using the "snapmirror show" command. Check the SnapMirror transfer history using the "snapmirror show-history" command

|===

<<top, Back to Top>>

=== Cloud Volume (CVO) Monitors

|===

|Monitor Name|CI Severity|Monitor Description|Corrective Action |CVO Disk Out of Service|INFO|This event occurs when a disk is removed from service because it has been marked failed, is being sanitized, or has entered the Maintenance Center.|None

|CVO Giveback of Storage Pool Failed|CRITICAL|This event occurs during the migration of an aggregate as part of a storage failover (SFO) giveback, when the destination node cannot reach the object stores.|Perform the following corrective actions:

Verify that your intercluster LIF is online and functional by using the "network interface show" command.

Check network connectivity to the object store server by using the "'ping" command over the destination node intercluster LIF.

Verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the "aggregate object-store config show" command.

Alternatively, you can override the error by specifying false for the "require-partner-waiting" parameter of the giveback command.

Contact NetApp technical support for more information or assistance. |CVO HA Interconnect Down|WARNING|The high-availability (HA) interconnect is down. Risk of service outage when failover is not available. |Corrective actions depend on the number and type of HA interconnect links supported by the platform, as well as the reason why the interconnect is down.

If the links are down:

Verify that both controllers in the HA pair are operational.

For externally connected links, make sure that the interconnect cables are connected properly and that the small form-factor pluggables (SFPs), if applicable, are seated properly on both controllers.

For internally connected links, disable and re-enable the links, one after the other, by using the "ic link off" and "ic link on" commands.

If links are disabled, enable the links by using the "ic link on" command.

If a peer is not connected, disable and re-enable the links, one after the other, by using the "ic link off" and "ic link on" commands.

Contact NetApp technical support if the issue persists. | CVO Max Sessions Per User Exceeded | WARNING

|You have exceeded the maximum number of sessions allowed per user over a TCP connection. Any request to establish a session will be denied until some sessions are released.

|Perform the following corrective actions:

Inspect all the applications that run on the client, and terminate any that are not operating properly.

Reboot the client.

Check if the issue is caused by a new or existing application:

If the application is new, set a higher threshold for the client by using the "cifs option modify -max-opens-same-file-per-tree" command. In some cases, clients operate as expected, but require a higher threshold. You should have advanced privilege to set a higher threshold for the client.

If the issue is caused by an existing application, there might be an issue with the client. Contact NetApp technical support for more information or assistance.

|CVO NetBIOS Name Conflict|CRITICAL

|The NetBIOS Name Service has received a negative response to a name registration request, from a remote machine. This is typically caused by a conflict in the NetBIOS name or an alias. As a result, clients might not be able to access data or connect to the right data-serving node in the cluster. |Perform any one of the following corrective actions:

If there is a conflict in the NetBIOS name or an alias, perform one of the following:

Delete the duplicate NetBIOS alias by using the "vserver cifs delete -aliases alias -vserver vserver" command.

Rename a NetBIOS alias by deleting the duplicate name and adding an alias with a new name by using the "vserver cifs create -aliases alias -vserver vserver" command.

If there are no aliases configured and there is a conflict in the NetBIOS name, then rename the CIFS server by using the "vserver cifs delete -vserver vserver" and "vserver cifs create -cifs-server netbiosname" commands.

NOTE: Deleting a CIFS server can make data inaccessible.

Remove NetBIOS name or rename the NetBIOS on the remote machine. |CVO NFSv4 Store Pool Exhausted|CRITICAL|A NFSv4 store pool has been exhausted.|If the NFS server is unresponsive for more than 10 minutes after this event, contact NetApp technical support.

| CVO Node Panic|WARNING|This event is issued when a panic occurs|Contact NetApp customer support.

|CVO Node Root Volume Space Low|CRITICAL|The system has detected that the root volume is dangerously low on space. The node is not fully operational. Data LIFs might have failed over within the cluster, because of which NFS and CIFS access is limited on the node. Administrative capability is limited to local recovery procedures for the node to clear up space on the root volume.|Perform the following corrective actions:

Clear up space on the root volume by deleting old Snapshot copies, deleting files you no longer need from the /mroot directory, or expanding the root volume capacity.

Reboot the controller.

Contact NetApp technical support for more information or assistance. |CVO Nonexistent Admin Share|CRITICAL|Vscan issue: a client has attempted to connect to a nonexistent ONTAP_ADMIN\$ share. |Ensure that Vscan is

enabled for the mentioned SVM ID. Enabling Vscan on a SVM causes the ONTAP ADMIN\$ share to be created for the SVM automatically.

|CVO Object Store Host Unresolvable|CRITICAL|The object store server host name cannot be resolved to an IP address. The object store client cannot communicate with the object-store server without resolving to an IP address. As a result, data might be inaccessible. |Check the DNS configuration to verify that the host name is configured correctly with an IP address.

|CVO Object Store Intercluster LIF Down|CRITICAL|The object-store client cannot find an operational LIF to communicate with the object store server. The node will not allow object store client traffic until the intercluster LIF is operational. As a result, data might be inaccessible. |Perform the following corrective actions:

Check the intercluster LIF status by using the "network interface show -role intercluster" command.

Verify that the intercluster LIF is configured correctly and operational.

If an intercluster LIF is not configured, add it by using the "network interface create -role intercluster" command.

|CVO Object Store Signature Mismatch|CRITICAL|The request signature sent to the object store server does not match the signature calculated by the client. As a result, data might be inaccessible. |Verify that the secret access key is configured correctly. If it is configured correctly, contact NetApp technical support for assistance.

|CVO QoS Monitor Memory Maxed Out|CRITICAL|The QoS subsystem's dynamic memory has reached its limit for the current platform hardware. Some QoS features might operate in a limited capacity.|Delete some active workloads or streams to free up memory. Use the "statistics show -object workload -counter ops" command to determine which workloads are active. Active workloads show non-zero ops. Then use the "workload delete <workload_name>" command multiple times to remove specific workloads.
Alternatively, use the "stream delete -workload <workload name> *" command to delete the associated streams from the active workload.

|CVO READDIR Timeout|CRITICAL|A READDIR file operation has exceeded the

| CVO READDIR Timeout | CRITICAL | A READDIR file operation has exceeded the timeout that it is allowed to run in WAFL. This can be because of very large or sparse directories. Corrective action is recommended. | Perform the following corrective actions:

Find information specific to recent directories that have had READDIR file operations expire by using the following 'diag' privilege nodeshell CLI command:

wafl readdir notice show.

Check if directories are indicated as sparse or not:

If a directory is indicated as sparse, it is recommended that you copy the contents of the directory to a new directory to remove the sparseness of the directory file.

If a directory is not indicated as sparse and the directory is large, it is recommended that you reduce the size of the directory file by reducing the number of file entries in the directory.

|CVO Relocation of Storage Pool Failed|CRITICAL|This event occurs during the relocation of an aggregate, when the destination node cannot reach the object stores. |Perform the following corrective actions:

Verify that your intercluster LIF is online and functional by using the "network interface show" command.

Check network connectivity to the object store server by using the "'ping" command over the destination node intercluster LIF.

Verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the "aggregate object-store config show" command.

Alternatively, you can override the error by using the "override-destination-checks" parameter of the relocation command.

Contact NetApp technical support for more information or assistance. |CVO Shadow Copy Failed|CRITICAL|A Volume Shadow Copy Service (VSS), a Microsoft Server backup and restore service operation, has failed.|Check the following using the information provided in the event message:

Is shadow copy configuration enabled?

Are the appropriate licenses installed?

On which shares is the shadow copy operation performed?

Is the share name correct?

Does the share path exist?

What are the states of the shadow copy set and its shadow copies? |CVO Storage VM Stop Succeeded|INFO|This message occurs when a 'vserver stop' operation succeeds.|Use 'vserver start' command to start the data access on a storage VM.

|CVO Too Many CIFS Authentication|WARNING|Many authentication negotiations have occurred simultaneously. There are 256 incomplete new session requests from this client.|Investigate why the client has created 256 or more new connection requests. You might have to contact the vendor of the

client or of the application to determine why the error occurred. |CVO Unassigned Disks|INFO|System has unassigned disks - capacity is being wasted and your system may have some misconfiguration or partial configuration change applied.|Perform the following corrective actions:

Determine which disks are unassigned by using the "disk show -n" command.

Assign the disks to a system by using the "disk assign" command. |CVO Unauthorized User Access to Admin Share|WARNING|A client has attempted to connect to the privileged ONTAP_ADMIN\$ share even though their logged-in user is not an allowed user.|Perform the following corrective actions:

Ensure that the mentioned username and IP address is configured in one of the active Vscan scanner pools.

Check the scanner pool configuration that is currently active by using the "vserver vscan scanner pool show-active" command.

|CVO Virus Detected|WARNING|A Vscan server has reported an error to the storage system. This typically indicates that a virus has been found. However, other errors on the Vscan server can cause this event.

Client access to the file is denied. The Vscan server might, depending on its settings and configuration, clean the file, quarantine it, or delete it. |Check the log of the Vscan server reported in the "syslog" event to see if it was able to successfully clean, quarantine, or delete the infected file. If it was not able to do so, a system administrator might have to manually delete the file.

|CVO Volume Offline|INFO|This message indicates that a volume is made offline.|Bring the volume back online.

|CVO Volume Restricted|INFO|This event indicates that a flexible volume is made restricted.|Bring the volume back online.

|===

<<top, Back to Top>>

=== SnapMirror for Business Continuity (SMBC) Mediator Log Monitors

|===

|Monitor Name|Severity|Monitor Description|Corrective Action

|ONTAP Mediator Added|INFO|This message occurs when ONTAP Mediator is added successfully on a cluster.|None

|ONTAP Mediator Not Accessible|CRITICAL|This message occurs when either the ONTAP Mediator is repurposed or the Mediator package is no longer installed on the Mediator server. As a result, SnapMirror failover is not possible.|Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

|ONTAP Mediator Removed|INFO|This message occurs when ONTAP Mediator is removed successfully from a cluster.|None

|ONTAP Mediator Unreachable|WARNING|This message occurs when the ONTAP Mediator is unreachable on a cluster. As a result, SnapMirror failover is not possible.|Check the network connectivity to the ONTAP Mediator by using the "network ping" and "network traceroute" commands. If the issue persists, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

|SMBC CA Certificate Expired|CRITICAL|This message occurs when the ONTAP Mediator certificate authority (CA) certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.|Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new CA certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

|SMBC CA Certificate Expiring|WARNING|This message occurs when the ONTAP Mediator certificate authority (CA) certificate is due to expire within the next 30 days.|Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new CA certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

|SMBC Client Certificate Expired|CRITICAL|This message occurs when the ONTAP Mediator client certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.|Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

|SMBC Client Certificate Expiring|WARNING|This message occurs when the ONTAP Mediator client certificate is due to expire within the next 30 days.|Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

|SMBC Relationship Out of Sync

Note: UM doesn't have this one|CRITICAL|This message occurs when a SnapMirror for Business Continuity (SMBC) relationship changes status from "in-sync" to "out-of-sync". Due to this RPO=0 data protection will be disrupted.|Check the network connection between the source and destination

volumes. Monitor the SMBC relationship status by using the "snapmirror show" command on the destination, and by using the "snapmirror list-destinations" command on the source. Auto-resync will attempt to bring the relationship back to "in-sync" status. If the resync fails, verify that all the nodes in the cluster are in quorum and are healthy.

|SMBC Server Certificate Expired|CRITICAL|This message occurs when the ONTAP Mediator server certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.|Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new server certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

|SMBC Server Certificate Expiring|WARNING|This message occurs when the ONTAP Mediator server certificate is due to expire within the next 30 days.|Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new server certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

|===

<<top, Back to Top>>

=== Additional Power, Heartbeat, and Miscellaneous System Monitors

|===

|Monitor Name|Severity|Monitor Description|Corrective Action

|Disk Shelf Power Supply Discovered|INFORMATIONAL|This message occurs when a power supply unit is added to the disk shelf.|NONE |Disk Shelves Power Supply Removed|INFORMATIONAL|This message occurs when a power supply unit is removed from the disk shelf.|NONE |MetroCluster Automatic Unplanned Switchover Disabled|CRITICAL|This message occurs when automatic unplanned switchover capability is disabled.|Run the "metrocluster modify -node-name <nodename> -automatic -switchover-onfailure true" command for each node in the cluster to enable automatic switchover.

|MetroCluster Storage Bridge Unreachable|CRITICAL|The storage bridge is not reachable over the management network|1) If the bridge is monitored by SNMP, verify that the node management LIF is up by using the "network interface show" command. Verify that the bridge is alive by using the "network ping" command.

2) If the bridge is monitored in-band, check the fabric cabling to the bridge, and then verify that the bridge is powered up.

|MetroCluster Bridge Temperature Abnormal - Below Critical|CRITICAL|The

sensor on the Fibre Channel bridge is reporting a temperature that is below the critical threshold. |1) Check the operational status of the fans on the storage bridge.

2) Verify that the bridge is operating under recommended temperature conditions.

|MetroCluster Bridge Temperature Abnormal - Above Critical|CRITICAL|The sensor on the Fibre Channel bridge is reporting a temperature that is above the critical threshold.|1) Check the operational status of the chassis temperature sensor on the storage bridge using the command "storage bridge show -cooling".

2) Verify that the storage bridge is operating under recommended temperature conditions.

|MetroCluster Aggregate Left Behind|WARNING|The aggregate was left behind during switchback.|1) Check the aggregate state by using the command "aggr show".

2) If the aggregate is online, return it to its original owner by using the command "metrocluster switchback".

|All Links Between Metrocluster Partners Down|CRITICAL|RDMA interconnect adapters and intercluster LIFs have broken connections to the peered cluster or the peered cluster is down.|1) Ensure that the intercluster LIFs are up and running. Repair the intercluster LIFs if they are down.

- 2) Verify that the peered cluster is up and running by using the "cluster peer ping" command. See the MetroCluster Disaster Recovery Guide if the peered cluster is down.
- 3) For fabric MetroCluster, verify that the back-end fabric ISLs are up and running. Repair the back-end fabric ISLs if they are down.
- 4) For non-fabric MetroCluster configurations, verify that the cabling is correct between the RDMA interconnect adapters. Reconfigure the cabling if the links are down.

|MetroCluster Partners Not Reachable Over Peering Network|CRITICAL|The connectivity to the peer cluster is broken.|1) Ensure that the port is connected to the correct network/switch.

- 2) Ensure that the intercluster LIF is connected with the peered cluster.
- 3) Ensure that the peered cluster is up and running by using the command "cluster peer ping". Refer to the MetroCluster Disaster Recovery Guide if the peered cluster is down.

|MetroCluster Inter Switch All Links Down|CRITICAL|All Inter-Switch Links (ISLs) on the storage switch are down.|1) Repair the back-end fabric ISLs on the storage switch.

- 2) Ensure that the partner switch is up and its ISLs are operational.
- 3) Ensure that intermediate equipment, such as xWDM devices, are operational.

|MetroCluster Node To Storage Stack SAS Link Down|WARNING|The SAS adapter or its attached cable might be at fault.|1. Verify that the SAS adapter is online and running.

2. Verify that the physical cable connection is secure and operating, and

replace the cable if necessary.

3. If the SAS adapter is connected to disk shelves, ensure IOMs and disks are properly seated.

|MetroClusterFC Initiator Links Down|CRITICAL|The FC initiator adapter is at fault.|1. Ensure that the FC initiator link has not been tampered with.

- 2. Verify the operational status of the FC initiator adapter by using the command "system node run -node local -command storage show adapter".
- |FC-VI Interconnect Link Down|CRITICAL|The physical link on the FC-VI port is offline.|1. Ensure that the FC-VI link has not been tampered with.
- 2. Verify that the physical status of the FC-VI adapter is "Up" by using the command "metrocluster interconnect adapter show".
- 3. If the configuration includes fabric switches, ensure that they are properly cabled and configured.

|MetroCluster Spare Disks Left Behind|WARNING|The spare disk was left behind during switchback.|If the disk is not failed, return it to its original owner by using the command "metrocluster switchback".

|MetroCluster Storage Bridge Port Down|CRITICAL|The port on the storage bridge is offline.|1) Check the operational status of the ports on the storage bridge by using the command "storage bridge show -ports".

- 2) Verify logical and physical connectivity to the port.
- |MetroCluster Storage Switch Fans Failed|CRITICAL|The fan on the storage switch failed.|1) Ensure that the fans in the switch are operating correctly by using the command "storage switch show -cooling".
- 2) Ensure that the fan FRUs are properly inserted and operational. |MetroCluster Storage Switch Unreachable|CRITICAL|The storage switch is not reachable over the management network.|1) Ensure that the node management LIF is up by using the command "network interface show".
- 2) Ensure that the switch is alive by using the command "network ping".
- 3) Ensure that the switch is reachable over SNMP by checking its SNMP settings after logging into the switch.

|MetroCluster Switch Power Supplies Failed|CRITICAL|A power supply unit on the storage switch is not operational.|1) Check the error details by using the command "storage switch show -error -switch-name <swtich name>".

- 2) Identify the faulty power supply unit by using the command "storage switch show -power -switch-name <switch name>".
- 3) Ensure that the power supply unitis properly inserted into the chassis of the storage switch and fully operational.

|MetroCluster Switch Temperature Sensors Failed|CRITICAL|The sensor on the Fibre Channel switch failed.|1) Check the operational status of the temperature sensors on the storage switch by using the command "storage switch show -cooling".

2) Verify that the switch is operating under recommended temperature conditions.

|MetroCluster Switch Temperature Abnormal|CRITICAL|The temperature sensor on the Fibre Channel switch reported abnormal temperature.|1) Check the operational status of the temperature sensors on the storage switch by

using the command "storage switch show -cooling".

2) Verify that the switch is operating under recommended temperature conditions.

|Service Processor Heartbeat Missed |INFORMATIONAL|This message occurs when ONTAP does not receive an expected "heartbeat" signal from the Service Processor (SP). Along with this message, log files from SP will be sent out for debugging. ONTAP will reset the SP to attempt to restore communication. The SP will be unavailable for up to two minutes while it reboots.|Contact NetApp technical support.

|Service Processor Heartbeat Stopped|WARNING|This message occurs when ONTAP is no longer receiving heartbeats from the Service Processor (SP). Depending on the hardware design, the system may continue to serve data or may determine to shut down to prevent data loss or hardware damage. The system continues to serve data, but because the SP might not be working, the system cannot send notifications of down appliances, boot errors, or Open Firmware (OFW) Power-On Self-Test (POST) errors. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution. |If the system has shut down, attempt a hard power cycle: Pull the controller out from the chassis, push it back in then power on the system. Contact NetApp technical support if the problem persists after the power cycle, or for any other condition that may warrant attention.

```
<<top,Back to Top>>

== More Information

//* xref:{relative_path}concept_notifications_email.html[Email Alerting]
for Monitors

* xref:{relative_path}task_view_and_manage_alerts.html[Viewing and Dismissing Alerts]

:leveloffset: -1
```

[[IDca0b23332a47f2d9fdb8f4b6aa1b32d9]]

```
= Cloud Insights API
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The Cloud Insights API enables NetApp customers and independent software
vendors (ISVs) to integrate Cloud Insights with other applications, such
as CMDB's or other ticketing systems.
//NOTE: This documentation is considered *Preview Documentation* and is
therefore subject to change.
//NOTE: The Cloud Insights API is available in *Cloud Insights Premium
Edition*.
Note that Cloud Insights APIs are available based on your current Edition:
[cols="<,^s,^s,^s"]
|===
|API Type|Basic|Standard|Premium
|Acquisition
Unit|image:SmallCheckMark.png[]|image:SmallCheckMark.png[]|image:SmallChec
kMark.png[]
| Data
Collection|image:SmallCheckMark.png[]|image:SmallCheckMark.png[]|image:Sma
llCheckMark.png[]
|Alerts| |image:SmallCheckMark.png[]|image:SmallCheckMark.png[]
|Assets| |image:SmallCheckMark.png[]|image:SmallCheckMark.png[]
|Data Ingestion| |image:SmallCheckMark.png[]|image:SmallCheckMark.png[]
|Log Ingestion| |image:SmallCheckMark.png[]|image:SmallCheckMark.png[]
|===
Additionally, your Cloud Insights link:https://docs.netapp.com/us-
en/cloudinsights/concept user roles.html#permission-levels[feature set
role] will determine which APIs you can access. User and Guest roles have
fewer privileges than Administrator role. For example, if you have
Administrator role in Monitor and Optimize, but User role in Reporting,
you can manage all API types except Data Warehouse.
```

- == Requirements for API Access
- * An API Access Token model is used to grant access.
- * API Token management is performed by Cloud Insights users with the Administrator role.
- == API Documentation (Swagger)

The latest API information is found by logging in to Cloud Insights and navigating to *Admin > API Access*. Click the *API Documentation* link.

```
//image:API_Types_Example.png[API Types]
//image:API_Documentation.png[API Documentation]
image:API_Swagger_Types.png[API types]
```

The API Documentation is Swagger-based, which provides a brief description and usage information for the API, and allows you to try it out in your environment. Depending on your user role and/or Cloud Insights edition, the API types available to you may vary.

image:API Swagger Example.png[API Swagger Example]

== API Access Tokens

Before using the Cloud Insights API, you must create one or more *API Access Tokens*. Access tokens are used for specified API types, and can grant read and/or write permissions. You can also set the expiration for each access token. All APIs under the specified types are valid for the access token. Each token is void of a username or password.

To create an Access Token:

- * Click *Admin > API Access*
- * Click *+API Access Token*
 - ** Enter Token Name
 - ** Select API Types
 - ** Specify the permissions granted for this API access
 - ** Specify Token Expiration

NOTE: Your token will only be available for copying to the clipboard and saving during the creation process. Tokens can not be retrieved after they are created, so it is highly recommended to copy the token and save it in a secure location. You will be prompted to click the *Copy API Access Token* button before you can close the token creation screen.

You can disable, enable, and revoke tokens. Tokens that are disabled can be enabled.

Tokens grant general purpose access to APIs from a customer perspective; managing access to APIs in the scope of their own tenant. Customer administrators may grant and revoke these tokens without direct involvement from Cloud Insights back end personnel.

The application receives an Access Token after a user successfully authenticates and authorizes access, then passes the Access Token as a credential when it calls the target API. The passed token informs the API that the bearer of the token has been authorized to access the API and perform specific actions specified by the scope that was granted during authorization.

The HTTP header where the Access Token is passed is *X-CloudInsights-ApiKey:*.

For example, use the following to retrieve storages assets:

curl https://<tenant_host_name>/rest/v1/assets/storages -H 'XCloudInsights-ApiKey:<API_Access_Token>'

Where _<API_Access_Token>_ is the token you saved during API access creation.

See the swagger pages for examples specific to the API you wish to use.

== API Type

The Cloud Insights API is category-based, and currently contains the following types:

- * ASSETS type contains asset, query, and search APIs.
- ** Assets: Enumerate top-level objects and retrieve a specific object or an object hierarchy.
- ** Query: Retrieve and manage Cloud Insights queries.
- ** Import: Import annotations or applications and assign them to objects
- ** Search: Locate a specific object without knowing the object's unique ID or full name.

- * DATA COLLECTION type is used to retrieve and manage data collectors.
- * DATA INGESTION type is used to retrieve and manage ingestion data and custom metrics, such as from Telegraf agents
- * LOG INGESTION is used to retrieve and manage log data

//* DATA WAREHOUSE type is available for environments with Cloud Insights Reporting, and is used to manage data used with Reporting

Additional types and/or APIs may become available over time. You can find the latest API information in the xref:{relative_path}#api-documentation-swagger[API Swagger documentation].

Note that the API types to which a user has access depend also on the xref:{relative_path}concept_user_roles.html[User Role] they have in each Cloud Insights feature set (Monitoring, Workload Security, Reporting).

```
////
== API Commands
```

Each of the REST API commands comprises the API's URL, an HTTP action, URL parameters, and an expected API responses.

The Cloud Insights APIs can be generally separated into the following sections:

////

////

admin::

Provides an entry point to the system for administrative operations in Cloud Insights. Allowed roles are Any, User, and Administrator and includes the following:

- ** admin/acquisitionUnits
- ** admin/certificates
- ** admin/datasources
- ** admin/licenses
- ** admin/patches
- ** admin/users

////

////

=== ASSETS type

Assets::

Lets you enumerate top-level objects and retrieve a specific object or an

object hierarchy from a system when you know the object unique ID or full name. ** /assets/<asset type> ** /assets/annotations ** /assets/applications ** /assets/businessEntities ** /assets/dataStores ** /assets/devices ** /assets/disks ** /assets/fabrics ** /assets/fileSystems ** /assets/genericDevices ** /assets/hosts ** /assets/iSCSINetworkPortal ** /assets/iSCSISessions ** /assets/internalVolumes ** /assets/paths ** /assets/ports ** /assets/qtrees ** /assets/quotas ** /assets/risks ** /assets/shares ** /assets/storageNodes ** /assets/storeagePools ** /assets/storageVirtualMachines ** /assets/storages ** /assets/switches ** /assets/tapes ** /assets/virtualMachines ** /assets/vmdks ** /assets/volumes ** /assets/zones ** /assets/zoneMembers Queries:: Retrieve and manage queries. ** /queries ** /query (to run an ad-hoc query) login:: Perform user session management.

```
Locate a specific object without knowing the object's unique ID or full
name.
** /search
=== DATA COLLECTION type
Data Collectors::
Retrieve and manage data collectors
** /collector/datasourceTypes
** /collector/datasources
////
== Inventory Traversal
This section describes how to traverse a hierarchy of Cloud Insights
objects.
=== Top Level Objects
Individual objects are identified in requests by unique URL (called "self"
in JSON) and require knowledge of object type and internal id. For some of
the top level objects (Hosts, Storages, and so on), REST API provides
access to the complete collection.
The general format of an API URL is:
https://<tenant>/rest/v1/<type>/<object>
For example, to retrieve all storages from a tenant named
mysite.c01.cloudinsights.netapp.com , the request URL is:
https://mysite.c01.cloudinsights.netapp.com/rest/v1/assets/storages
=== Children and Related Objects
Top level objects, such as Storage, can be used to traverse to other
children and related objects. For example, to retrieve all disks for a
specific storage, concatenate the storage "self" URL with "/disks", for
example:
// https://<tenant> + "/rest/v1/assets/storages/4537" + "/disks"
 https://<tenant>/rest/v1/assets/storages/4537/disks
```

== Expands

Many API commands support the *expand* parameter, which provides additional details about the object or URLs for related objects.

The one common expand parameter is _expands_. The response contains a list of all available specific expands for the object.

For example, when you request the following:

https://<tenant>/rest/v1/assets/storages/2782?expand= expands

The API returns all available expands for the object as follows:

//comment here

// https://<tenant>/rest/v1/assets/storages/4537/disks

//lots of text here in the original as a picture.

image:expands.gif[expands example]

Each expand contains data, a URL, or both. The expand parameter supports multiple and nested attributes, for example:

https://<tenant>/rest/v1/assets/storages/2782?expand=performance,storageResources.storage

Expand allows you to bring in a lot of related data in one response. NetApp advises that you do not request too much information at one time; this can cause performance degradation.

To discourage this, requests for top-level collections cannot be expanded. For example, you cannot request expand data for all storage objects at once. Clients are required to retrieve the list of objects and then choose specific objects to expand.

== Performance Data

Performance data is gathered across many devices as separate samples. Every hour (the default), Cloud Insights aggregates and summarizes performance samples.

The API allows access to both the samples and the summarized data. For an object with performance data, a performance summary is available as

expand=performance. Performance history time series are available through nested _expand=performance.history_.

Examples of Performance Data objects include:

- * StoragePerformance
- * StoragePoolPerformance
- * PortPerformance
- * DiskPerformance

A Performance Metric has a description and type and contains a collection of performance summaries. For example, Latency, Traffic, and Rate.

A Performance Summary has a description, unit, sample start time, sample end time, and a collection of summarized values (current, min, max, avg, etc.) calculated from a single performance counter over a time range (1 hour, 24 hours, 3 days, and so on).

```
// == Performance Data
```

//The Performance data JSON...

image:API_Performance.png[API Performance Example]

The resulting Performance Data dictionary has the following keys:

- * "self" is the object's unique URL
- * "history" is the list of pairs of timestamp and map of counters values
- * Every other dictionary key ("diskThroughput" and so on) is the name of a performance metric.

Each performance data object type has a unique set of performance metrics. For example, the Virtual Machine performance object supports "diskThroughput" as a performance metric. Each supported performance metric is of a certain "performanceCategory" presented in the metric dictionary. Cloud Insights supports several performance metric type listed later in this document. Each performance metric dictionary will also have the "description" field that is a human-readable description of this performance metric and a set of performance summary counter entries.

The Performance Summary counter is the summarization of performance counters. It presents typical aggregated values like min, max, and avg for a counter and also the latest observed value, time range for summarized data, unit type for counter and thresholds for data. Only thresholds are optional; the rest of attributes are mandatory.

Performance summaries are available for these types of counters:

```
* Read - Summary for read operations
* Write - Summary for write operations
* Total - Summary for all operations. It may be higher than the simple sum
of read and write; it may include other operations.
* Total Max - Summary for all operations. This is the maximum total value
in the specified time range.
== Object Performance Metrics
The API can return detailed metrics for objects in your environment, for
example:
* Storage Performance Metrics such as IOPS (Number of input/output
requests per second), Latency, or Throughput.
////
[cols=2*,options="header",cols="25,75"]
| Performance Metric
| Description
| IOPS | Number of input/output requests per second
|Latency|Average time, in milliseconds, it takes to complete an IO request
|Cache hit ratio|Percentage of requests satisfied from cache
|Utilization|Percent of theoretical maximum for average utilization of the
resource in specified time frame. The range is 0 to 100.
|Throughput|Number of bytes transferred in MB/second
|IO density|Number of IO per second per terabyte of used capacity. IO
density is used only for storage, volume and internal volume assets.
|===
The following table describes performance metrics available for different
type of objects:
[cols=7*, options="header"]
|===
||IOPS|Latency|Throughput|Cache Hit|Utilization|IO Density
|Disk|X||X||X|
|Storage Pool|X||X||X||
Internal Volume|X|X|X||X|
Volume | X | X | X | X | X | X |
StorageNode|X|X|X|X|X||
Storage | X | X | X | | | X |
Data Store | X | X | X | | | |
VMDK |X|X|X|||
VM | X | X | X | | X | |
```

Host |X|X|X||X|||=== //// * Switch Performance Metrics, such as Traffic Utilization, BB Credit Zero data, or Port Errors. See the xref:{relative path}#api-documentation-swagger[API Swagger documentation] for information on metrics for each object type. //// The following table describes performance metrics available for switches: [cols=5*, options="header"] |Type|REST Name|Unit|Range|Description |Traffic Utilization|trafficUtilization.rx|%|0-100|Receive traffic utilization. Calculated as 100 * total bytes received / receive capacity during the sampling period. |Traffic Utilization|trafficUtilization.tx|%|0-100|Transmit traffic utilization. Calculated as 100 * total bytes transmitted / transmit capacity during the sampling period |Traffic Utilization|trafficUtilization.totaL|%|0-100|Total traffic utilization. Calculated as 100 * total bytes received and transmitted / port capacity during the sampling period. |Traffic Utilization|trafficUtilization.rxMax|%|0-100 |Maximum of trafficUtilization.tx during the sampling period. |Traffic Rate|trafficRate.rx|MiB/s|≥0|Traffic received during the sampling |Traffic Rate|trafficRate.tx|MiB/s|≥0|Traffic transmitted during the sampling period. |Traffic Frame|trafficFrameRate.rx|frames/s|≥0|Number of FC frames received per second during the sampling period. |Traffic Frame|trafficFrameRate.tx|frames/s|≥0|Number of FC frames transmitted per second during the sampling period. |Traffic Frame|trafficFrameRate.total|frames|s≥0|Number of FC frames transmitted and received per second during the sampling period. |Traffic Frame|trafficFrameSizeAvg.rx|Bytes/frame|0-2,148|Average length of received FC frames during the sampling period. |Zero BB Credit|bbCreditZero.rx|none (count)|≥0|Number of times the

receive buffer-to-buffer credit count transitioned to zero during the

sampling period. It represents the number of times the attached port had to stop transmitting because this port was out of credits to provide.

|Zero BB Credit |bbCreditZero.tx|none (count)|\geq0|Number of times the transmit buffer-to-buffer credit count transitioned to zero during the sampling period.

|Zero BB Credit|bbCreditZero.total|none (counnt)|≥0|Number of times the transmit and receive buffer-to-buffer credit counts transitioned to zero during the sampling period.

|Zero BB Credit|bbCreditZeroMs.tx|ms|≥0|Time in milliseconds during which the transmit buffer-to-buffer credit count was zero during the sample period.

|Port Errors|portErrors.timeoutDiscardTx|none (count)| \geq 0|Number of receive link resets during the sample period. Represents the number of link resets issued by the attached port to this port.

|Port Errors|portErrors.linkResetRx|none(count)|≥0|Number of receive link resets during the sample period. Represents the number of link resets issued by the attached port to this port.

|Port Errors|portErrors.syncLoss|none (count)|≥0|Number of loss of synchronization failures during the sample period.

|Port Errors|portErrors.signalLoss|none (count)|≥0|Number of signal losses during the sample period.

|Port Errors|portErrors.class3Discard|none (count)|≥0|Number of class 3 FC frames discarded during the sample period.

|Port Errors|portErrors.frameTooLong|none(count)| \geq 0|Number of FC frames discarded by this port during the sample period because their length exceeded the agreed to maximum limit.

|Port Errors|portErrors.frameTooShort|none(count)|≥0|Number of FC frames discarded by this port during the sample period because their actual length was less than the length given in the frame header.

|Port Errors|portErrors.linkFailure|none (count)|≥0|Number of link failures detected by this port during the sample period.

|Port Errors|portErrors.crc|none (count)|≥0|Number of frames with invalid CRCs detected by this port during the sample period.

|Port Errors|portErrors.encIn (Brocade only)|none (count)|≥0|The number of 8b/10b encoding errors that have occurred inside frame boundaries. This counter is generally a zero value, although occasional errors may occur on a normal link and give a nonzero result.

|Port Errors|portErrors.encOut (Brocade only)|none (count)|≥0|The number of 8b/10b encoding errors that have occurred outside frames boundaries. This counter may become a nonzero value during link initialization but indicates a problem if it increments faster than the link-bit error rate allows (approximately once every 20 minutes for 1 Gb/s). This is usually caused by corrupted primitive sequences.

|Port Errors|portErrors.total|none (count)|≥0|Total number of errors detected by this port during the sample period. Is equal to the sum of all error counters defined above (with portErrors. prefix in REST name). |=== //// == Performance History Data History data is presented in performance data as a list of timestamp and counter maps pairs. //graphic History counters are named based on the performance metric object name. For example, the virtual machine performance object supports "diskThroughput" so the history map will contain keys named "diskThroughput.read", "diskThroughput.write" and "diskThroughput.total". NOTE: Timestamp is in UNIX time format. The following is an example of a performance data JSON for a disk: //Graphic image:DiskPerformanceExample.png[Disk Performance JSON] == Objects with Capacity Attributes Objects with capacity attributes use basic data types and the CapacityItem for representation. === CapacityItem CapacityItem is a single logical unit of capacity. It has "value" and "highThreshold" in units defined by its parent object. It also supports an optional breakdown map that explains how the capacity value is constructed. For example, the total capacity of a 100 TB storagePool would be a CapacityItem with a value of 100. The breakdown may show 60 TB allocated for "data" and 40 TB for "snapshots". Note:: "highThreshold" represents system defined thresholds for the corresponding metrics, which a client can use to generate alerts or visual cues on values that are out of acceptable configured ranges.

//Graphic

The following shows the capacity for StoragePools with multiple capacity counters:

//Graphic

image:StoragePoolCapacity.png[Storage Pool Capacity Example]

== Using Search to Look Up Objects

The search API is a simple entry point to the system. The only input parameter to the API is a free-form string and the resulting JSON contains a categorized list of results. Types are different asset types from the Inventory, such as storages, hosts, dataStores, and so on. Each type would contain a list of objects of the type that match the search criteria.

Cloud Insights is an extensible (wide open) solution that allows integrations with third party orchestration, business management, change control and ticketing systems as well as custom CMDB integrations.

Cloud Insight's RESTful API is a primary point of integration that allows simple and effective movement of data as well as allows users to gain seamless access to their data.

== Disabling or Revoking an API Token

To temporarily disable an API token, on the API token list page, click the "three dots" menu for the API, and select _Disable_. You can Re-enable the token at any time using the same menu and selecting Enable.

To permanently remove an API token, from the menu, select "Revoke". You cannot re-enable a revoked token; you must create a new token.

image:API_Disable_Token.png[Disable or Revoke and API token]

== Rotating Expired API Access Tokens

API access tokens have an expiration date. When an API access token expires, users need to generate a new token (of type _Data Ingestion_ with Read/Write permissions) and reconfigure Telegraf to use the newlygenerated token instead of the expired token. The steps below detail how to do this.

==== Kubernetes

Note that these commands are using the default namespace "netappmonitoring". If you have set your own namespace, substitute that namespace in these and all subsequent commands and files. Note: If you have the latest NetApp Kubernetes Monitoring Operator installed and using an API access token that is renewable, expiring tokens will automatically be replaced by new/refreshed API access tokens. There is no need to perform the manual steps listed below. * Edit the the NetApp Kubernetes Monitoring Operator. kubectl -n netapp-monitoring edit agent agent-monitoring-netapp * Modify the spec.output-sink.api-key value, replacing the old API token with the new API token. spec: output-sink: - api-key: <NEW API TOKEN> ==== RHEL/CentOS and Debian/Ubuntu * Edit the Telegraf configuration files, and replace all instances of the old API token with the new API token. sudo sed -i.bkup 's/<OLD API TOKEN>/<NEW API TOKEN>/g' /etc/telegraf/telegraf.d/*.conf * Restart Telegraf. sudo systemctl restart telegraf ==== Windows * For each Telegraf configuration file in C:\Program Files\telegraf\telegraf.d , replace all instances of the old API token with the new API token. cp <plugin>.conf <plugin>.conf.bkup (Get-Content <plugin>.conf).Replace('<OLD API TOKEN>', '<NEW API TOKEN>') | Set-Content <plugin>.conf

```
* Restart Telegraf.
 Stop-Service telegraf
 Start-Service telegraf
[[ID54ecf4813aa1e06e4f4d5a3a90ea16dd]]
= Notification using Webhooks
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Webhooks allow users to send alert notifications to various applications
using a customized webhook channel.
Many commercial applications support webhooks as a standard input
interface, for example: Slack, PagerDuty, Teams, and Discord all support
webhooks. By supporting a generic, customizable webhook channel, Cloud
Insights can support many of these delivery channels. Information on
webhooks can be found on these application websites. For example, Slack
provides link:https://api.slack.com/messaging/webhooks[this useful guide].
You can create multiple webhook channels, each channel targeted for a
different purpose; separate applications, different recipients, etc.
The webhook channel instance is comprised of the following elements:
|===
|Name|Unique name
|URL|Webhook target URL, including the http:// or https:// prefix
along with the url params
|Method |GET, POST - Default is POST
|Custom Header|Specify any custom header lines here
|Message Body|Put the body of your message here
|Default Alert Parameters|Lists the default parameters for the webhook
|Custom Parameters and Secrets|Custom parameters and secrets allow you to
add unique parameters and secure elements such as passwords
|===
```

== Creating a Webhook

To create a Cloud Insights webhook, go to *Admin > Notifications* and select the *Webhooks* tab.

The following image shows an example webhook configured for Slack:

image:Webhook Example Slack.png[Webhook Example]

Enter appropriate information for each of the fields, and click "Save" when complete.

You can also click the "Test Webhook" button to test the connection. Note that this will send the "Message Body" (without substitutions) to the defined URL according to the selected Method.

Cloud Insights webhooks comprise a number of default parameters.

Additionally, you can create your own custom parameters or secrets.

image:Webhook_Default_Parameters.png[Cloud Insights Default Webhook
Parameters]

=== Parameters: What are they and how do I use them?

Alert Parameters are dynamic values populated per alert. For example, the $_{\%}TriggeredOn\%$ parameter will be replaced with the object on which the alert was triggered.

Note that in this section, substitutions are _not_ performed when clicking the "Test Webhook" button; the button sends a payload that shows the _%%_ substitutions but does not replace them with data.

=== Custom Parameters and Secrets

In this section you can add any custom parameters and/or secrets you wish. For security reasons, if a secret is defined only the webhook creator can modify this webhook channel. It is read-only for others. You can use secrets in URL/Headers as %%<secret name>%% .

```
=== Webhooks List Page
On the Webhooks list page, displayed are the Name, Created By, Created On,
Status, Secure, and Last Reported fields.
== Choosing Webhook Notification in a Monitor
To choose the webhook notification in a
xref:{relative path}task create monitor.html#creating-a-monitor[monitor],
go to *Alerts > Manage Monitors* and select the desired monitor, or add a
new monitor. In the Set up team notifications section, choose Webhook
as the delivery method. Select the alert levels (Critical, Warning,
Resolved), then choose the desired webhook.
image:Webhook Monitor Notify.png[Webhook Monitor Notification]
//To be published after Feb 5:
//Select the alert levels (Critical, Warning, Resolved), then choose the
desired webhook(s). You can choose multiple webhooks for each alert, and
you can choose the same webhook for different alerts.
//image: Webhook Monitor Notifications.png [Webhook Monitor Notifications]
== Webhook Examples:
Webhooks for xref:{relative path}task webhook example slack.html[Slack]
Webhooks for
xref:{relative path}task webhook example pagerduty.html[PagerDuty]
Webhooks for xref:{relative path}task webhook example teams.html[Teams]
Webhooks for
xref:{relative path}task webhook example discord.html[Discord]
= Monitoring your Environment
```

```
:leveloffset: +1
[[ID80c4ff78cefbd7d2c73d58ddc89b9d9e]]
= Auditing
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
To identify changes both expected (for tracking) or unexpected (for
troubleshooting), you can view an audit trail of the Cloud Insights system
events and user activities.
== Viewing Audited Events
To View the Audit page, click *Admin > Audit* in the menu. The Audit page
is displayed, providing the following details for each audit entry:
* *Time* - Date and time of the event or activity
* *User* - The User who initiated the activity
* *Role* - The user's role in Cloud Insights (guest, user, administrator)
* *IP* - The IP address associated with the event
* *Action* - Type of activity, for example Login, Create, Update
* *Category* - The category of activity
* *Details* - Details of the activity
//When there is a user activity that affects a resource, such as a data
collector or an application, the details include a link to the resource's
landing page.
//*Note* When a data collector is deleted, the user activity details
related to the data collector no longer contain a link to the data
collector's landing page.
== Displaying audit entries
There are a number of different ways to view audit entries:
* You can display audit entries by choosing a particular time period (1
hour, 24 hours, 3 days, etc.).
```

```
* You can change the sort order of entries to either ascending (up arrow)
or descending (down arrow) by clicking the arrow in the column header.
By default, the table displays the entries in descending time order.
* You can use the filter fields to show only the entries you want in the
table. Click the [+] button to add additional filters.
image:Audit Filters.png[Audit Filters]
=== More on Filtering
You can use any of the following to refine your filter:
|Filter|What it does | Example | Result
| * (Asterisk) | enables you to search for everything | vol*rhel | returns
all resources that start with "vol" and end with "rhel"
| ? (question mark) | enables you to search for a specific number of
characters| BOS-PRD??-S12 | returns BOS-PRD** 12 **-S12, BOS-
PRD** 23 **-S12, and so on
OR | enables you to specify multiple entities | FAS2240 OR CX600 OR
FAS3270 | returns any of FAS2440, CX600, or FAS3270
| NOT |allows you to exclude text from the search results | NOT EMC*
|returns everything that does not start with "EMC"
| None |searches for blank/NULL/None in any field where selected |
None | returns results where the target field is not empty
| Not * | as with None above, but you can also use this form to search
for NULL values in _text-only_ fields | Not * |returns results where the
target field is not empty.
| "" |searches for an exact match| "NetApp*" | returns results containing
the exact literal string NetApp*
|===
If you enclose a filter string in double quotes, Insight treats everything
between the first and last quote as an exact match. Any special characters
or operators inside the quotes will be treated as literals. For example,
filtering for "*" will return results that are a literal asterisk; the
asterisk will not be treated as a wildcard in this case. The operators OR
and NOT will also be treated as literal strings when enclosed in double
quotes.
== Audited Events and Actions
The events and actions audited by Cloud insights can be categorized in the
```

following broad areas:

```
* *User account*: Log in, log out, role change, etc.
Example: User *Tony Lavoie* logged in from *10.1.120.15*, user agent
*Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like
Gecko) Chrome/85.0.4183.121 Safari/537.36*, login method(s) *Cloud Central
Portal Login *
* *Acquisition Unit*: create, delete, etc.
Example: Acquisition unit *AU-Boston-1* removed .
* *Data Collector*: add, remove, modify, postpone/resume, change
acquisition unit, start/stop, etc.
+
Example: Datasource *FlexPod Lab* removed, vendor *NetApp*, model *ONTAP
Data Management Software*, ip *192.168.106.5 *.
* *Application*: add, assign to object, remove, etc.
Example: Internal Volume *ocisedev:tlappSVM01:tlappFlexVol01* added to
application *Test App_*.
* *Annotation*: add, assign, remove, annotation rule actions, annotation
value changes, etc.
Example: Annotation value *Boston* added to annotation type
*SalesOffice *.
* *Query*: add, remove, etc.
Example: Query *TL Sales Query* is added .
* *Monitor*: add, remove, etc.
Example: Monitor Aggr Size - CI Alerts Notifications Dev updated
* *Notification*: change email, etc.
Example: Recipient ci-alerts-notifications-dl created
== Exporting Audit Events
You can export the results of your Audit display to a .CSV file, which
will allow you to analyze the data or import it into another application.
.Steps
```

- . On the Audit page, set the desired time range and any filters you want. Cloud Insights will export only the Audit entries that match the filtering and time range you have set.
- . Click the _Export_ button image:ExportButton.png[Export Button] in the upper right of the table.

The displayed Audit events will be exported to a .CSV file, up to a maximum of 10,000 rows.

== Retention of Audit Data

The amount of time Cloud Insights retains Audit data is based on your Edition:

- * Basic Edition: Audit data is retained for 30 days
- * Standard and Premium Editions: Audit data is retained for 1 year plus 1 day

Audit entries older than the retention time are automatically purged. No user interaction is needed.

== Troubleshooting

[.lead]

Here you will find suggestions for troubleshooting problems with Audit.

|===

|*Problem:* |*Try this:*

|I see Audit messages telling me that a monitor has been exported. |Export of a custom monitor configuration is typically used by NetApp engineers during development and testing of new features. If you did not expect to see this message, please consider exploring the actions of the user named in the audited action or contact support.

|===

:leveloffset: -1

= Asset Page Information

```
:leveloffset: +1
[[IDa19085ab1c700496a456832ec222bca4]]
= Asset Page Overview
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
Asset pages summarize the current status of an asset and contain links to
additional information about the asset and its related assets.
== Types of Asset Pages
Cloud Insights provides asset pages for the following assets:
* Virtual machine
* Storage Virtual Machine (SVM)
* Volume
* Internal volume
* Host (including Hypervisor)
* Storage pool
* Storage
* Datastore
* Application
* Storage node
* Qtree
* Disk
* VMDK
* Port
* Switch
* Fabric
//* Object storage (for example, Atmos, Centera, Amazon S3)
//* Zone
//Mapping and Masking information can be viewed in tables on Zone, Volume,
VM, and Host/Hypervisor asset pages.
//Note: Summary information is available for object storage assets;
however, you can only access this information from the Data sources detail
```

page.

== Changing the Time Range of Displayed Data

By default, an asset page displays the last 24 hours of data; however, you can change the segment of data displayed by selecting another fixed time range or a custom range of time to view less or more data.

You can change the time segment of displayed data by using an option that is located on every asset page, regardless of asset type. To change the time range, click the displayed time range in the top bar and choose from among the following time segments:

- * Last 15 Minutes
- * Last 30 Minutes
- * Last 60 Minutes
- * Last 2 Hours
- * Last 3 Hours (this is the default)
- * Last 6 Hours
- * Last 12 Hours
- * Last 24 Hours
- * Last 2 Days
- * Last 3 Days
- * Last 7 Days
- * Last 30 Days
- * Custom time range

The Custom time range allows you to select up to 31 consecutive days. You can also set the Start Time and End Time of day for this range. The default Start Time is 12:00 AM on the first day selected and the default End Time is 11:59 PM on the last day selected. Clicking Apply will apply the custom time range to the asset page.

The information in an asset page summary section, as well as in any tables or custom widgets on the page, refreshes automatically based on the selected time range. The current refresh rate is displayed in the upperright corner of the Summary section as well as on any relevant tables or widgets on the page.

== Add Custom Widgets

You can add your own widgets to any asset page. Widgets you add will appear on asset pages for all objects of that type. For example, adding a custom widget to a storage asset page will display that widget on asset pages for all storage assets.

```
[[ID4ad21205efa62d96ca33ced0d7b7734a]]
= Filtering for Objects In-Context
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

[.lead]

When configuring a widget on an asset's landing page, you can set _in-context_ filters to show only objects directly related to the current asset. By default, when you add a widget, _all_ objects of the selected type in your environment are displayed. In-context filters allow you to display only the data relevant to your current asset.

On most asset landing pages, widgets allow you to filter for objects related to the current asset. In filter drop-downs, object types that display a link icon image:LinkIcon.png[Link Icon] can be filtered incontext to the current asset.

For example, on a Storage asset page, you can add a Bar Chart widget to show the top IOPS on internal volumes only on that storage. By default, when you add a widget, _all_ internal volumes in your environment are displayed.

To show only internal volumes on the current storage asset, do the following:

.Steps

- . Open an asset page for any *Storage* asset.
- . Click *Edit* to open the asset page in Edit mode.
- . Click *Add Widget* and select Bar Chart .
- . Select *Internal Volume* for the object type to display on the bar chart. Notice that the internal volume object type has a link icon image:LinkIcon.png[Link Icon] beside it. The "linked" icon is enabled by default.

+

image:LinkingObjects.png[bar chart volumes]

- . Choose _IOPS Total_ and set any additional filters you like.
- . Collapse the *Roll Up* field by clicking the [X] beside it. The *Show* field is displayed.
- . Choose to show Top 10.
- . Save the widget.

+

The bar chart shows only the internal volumes that reside on the current storage asset.

The widget will be displayed on the asset pages for all storage objects. When the in-context link is enabled in the widget, the bar chart shows data for internal volumes related only to the currently-displayed storage asset.

To unlink the object data, edit the widget and click the link icon image:LinkIconEnabled.png[Link Icon Enabled] next to the object type. The link becomes disabled image:LinkIconDisabled.png[Link Icon Disabled] and the chart displays data for _all_ objects in your environment.

You can also use

xref:{relative_path}concept_dashboard_features.html#variables[*special
variables in widgets*] to display asset-related information on landing
pages.

```
[[IDeb483595aefc478d54f03bb88374d43e]]
= Asset Page Summary section
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

The Summary section of an asset page displays general information about an asset, including whether any metrics or performance policies are cause for concern. Potential problem areas are indicated by a red circle.

The information in the summary section, as well as in any tables or custom widgets on the asset page, refreshes automatically based on the selected time range. You can see the current refresh rate in the upper-right corner of the Summary section, the tables, and any custom widgets.

image:Summary Section Example.png[]

Note: The information displayed in the Summary section varies, depending on the type of asset you are viewing.

//For example, if your Storage Pool is experiencing an active
link:https://docs.netapp.com/us-

en/cloudinsights/insights_overview.html[_Shared Resource Under Stress_] Insight, the Summary section will include a link to that Insight.

You can click any of the asset links to view their asset pages. For example, if you are viewing a storage node, you can click a link to view the asset page of the storage it is associated with.

You can view the metrics associated with the asset. A red circle next to a metric indicates that you might need to diagnose and resolve potential problems.

NOTE: You may notice that volume capacity might show greater than 100% on some storage assets. This is due to metadata related to the capacity of the volume being part of the consumed capacity data reported by the asset.

If applicable, you can click an alert link to view the alert and monitor associated with the asset.

== Topology

On certain asset pages, the summary section contains a link to view the topology of the asset and its connections.

Topology is available for the following asset types:

- * Application
- * Disk
- * Fabric
- * Host
- * Internal Volume
- * Port
- * Switch
- * Virtual Machine
- * VMDK
- * Volume

image:TopologyExample.png[]

[[IDf1733ae91408ca23ab93d3624a49f2eb]]

= Expert View

```
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The Expert View section of an asset page enables you to view a performance
sample for the base asset based on any number of applicable metrics in
context with a chosen time period in the performance chart and any assets
related to it. The data in the charts refreshes automatically as data
collectors poll and updated data is acquired.
== Using the Expert View section
The following is an example of the Expert View section in a storage asset
page:
//image:ExpertViewExample1.png[Expert View Example]
image:Expert View 2021.png[Expert View Example]
You can select the metrics you want to view in the performance chart for
the time period selected. Click on the Display Metrics drop-down and
choose from the metrics listed.
The *Resources* section shows the name of the base asset and the color
representing the base asset in the performance chart. If the *Top
Correlated* section does not contain an asset you want to view in the
performance chart, you can use the *Search Assets* box in the *Additional
Resources* section to locate the asset and add it to the performance
chart. As you add resources, they appear in the Additional resources
section.
Also shown in the Resources section, when applicable, are any assets
related to the base asset in the following categories:
* Top correlated
Shows the assets that have a high correlation (percentage) with one or
more performance metrics to the base asset.
* Top contributors
Shows the assets that contribute (percentage) to the base asset.
```

* Workload Contentions Shows the assets that impact or are impacted by other shared resources, such as hosts, networks, and storage. These are sometimes called greedy and degraded resources. //// * Degraded Shows the assets that are depleted of system resources due to this asset. //// == Alerts in Expert View Alerts are also displayed in the Expert View section of an asset landing page, showing the time and duration of the alert as well as the monitor condition that triggered it. image:Alerts In Expert View.png[Alerts in Expert View] == Expert View metric definitions The Expert View section of an asset page displays several metrics based on the time period selected for the asset. Each metric is displayed in its own performance chart. You can add or remove metrics and related assets from the charts depending on what data you want to see. The metrics you can choose will vary depending on asset type. |=== | *Metric* | *Description* | BB credit zero Rx, Tx | Number of times the receive/transmit buffer-tobuffer credit count transitioned to zero during the sampling period. This metric represents the number of times the attached port had to stop transmitting because this port was out of credits to provide. | BB credit zero duration Tx | Time in milliseconds during which the transmit BB credit was zero during the sampling interval. | Cache hit ratio (Total, Read, Write) % | Percentage of requests that result in cache hits. The higher the number of hits versus accesses to the volume, the better is the performance. This column is empty for storage arrays that do not collect cache hit information.

| Cache utilization (Total) % | Total percentage of cache requests that

| Class 3 discards | Count of Fibre Channel Class 3 data transport

result in cache hits

discards.

399

- | CPU utilization (Total) % | Amount of actively used CPU resources, as a percentage of total available (over all virtual CPUs).
- | CRC error | Number of frames with invalid cyclic redundancy checks (CRCs) detected by the port during the sampling period
- | Frame rate | Transmit frame rate in frames per second (FPS)
- \mid Frame size average (Rx, Tx) \mid Ratio of traffic to frame size. This metric enables you to identify whether there are any overhead frames in the fabric.
- | Frame size too long | Count of Fibre Channel data transmission frames that are too long.
- | Frame size too short | Count of Fibre Channel data transmission frames that are too short.
- | I/O density (Total, Read, Write) | Number of IOPS divided by used capacity (as acquired from the most recent inventory poll of the data source) for the Volume, Internal Volume or Storage element. Measured in number of I/O operations per second per TB.
- | IOPS (Total, Read, Write) | Number of read/write I/O service requests passing through the I/O channel or a portion of that channel per unit of time (measured in I/O per sec)
- | IP throughput (Total, Read, Write) | Total: Aggregated rate at which IP data was transmitted and received in megabytes per second.
- | Read: IP Throughput (Receive): | Average rate at which IP data was received in megabytes per second.
- | Write: IP Throughput (Transmit): | Average rate at which IP data was transmitted in megabytes per second.
- | Latency (Total, Read, Write) | Latency (R&W): Rate at which data is read or written to the virtual machines in a fixed amount of time. The value is measured in megabytes per second.
- | Latency: | Average response time from the virtual machines in a data store.
- | Top Latency: | The highest response time from the virtual machines in a data store.
- | Link failure | Number of link failures detected by the port during the sampling period.
- \mid Link reset Rx, Tx \mid Number of receive or transmit link resets during the sampling period. This metric represents the number of link resets that were issued by the attached port to this port.
- \mid Memory utilization (Total) % \mid Threshold for the memory used by the host.
- | Partial R/W (Total) % | Total number of times that a read/write operation crosses a stripe boundary on any disk module in a RAID 5, RAID 1/0, or RAID 0 LUN Generally, stripe crossings are not beneficial, because each one requires an additional I/O. A low percentage indicates an efficient stripe element size and is an indication of improper alignment of a volume (or a NetApp LUN). For CLARiiON, this value is the number of stripe crossings divided by the total number of IOPS.

- | Port errors | Report of port errors over the sampling period/given time span.
- | Signal loss count | Number of signal loss errors. If a signal loss error occurs, there is no electrical connection, and a physical problem exists.
- | Swap rate (Total Rate, In rate, Out rate) | Rate at which memory is swapped in, out, or both from disk to active memory during the sampling period. This counter applies to virtual machines.
- | Sync loss count | Number of synchronization loss errors. If a synchronization loss error occurs, the hardware cannot make sense of the traffic or lock onto it. All the equipment might not be using the same data rate, or the optics or physical connections might be of poor quality. The port must resynchronize after each such error, which impacts system performance. Measured in KB/sec.
- | Throughput (Total, Read, Write) | Rate at which data is being transmitted, received, or both in a fixed amount of time in response to I/O service requests (measured in MB per sec).
- \mid Timeout discard frames Tx \mid Count of discarded transmit frames caused by timeout.
- | Traffic rate (Total, Read, Write) | Traffic transmitted, received, or both received during the sampling period, in mebibytes per second.
- | Traffic utilization (Total, Read, Write) | Ratio of traffic received/transmitted/total to receive/transmit/total capacity, during the sampling period.
- | Utilization (Total, Read, Write) % | Percentage of available bandwidth used for transmission (Tx) and reception (Rx).
- \mid Write pending (Total) \mid Number of write I/O service requests that are pending.

|===

== Using the Expert View section

The Expert view section enables you to view performance charts for an asset based on any number of applicable metrics during a chosen time period, and to add related assets to compare and contrast asset and related asset performance over different time periods.

.Steps

- . Locate an asset page by doing either of the following:
- * Search for and select a specific asset.
- * Select an asset from a dashboard widget.
- * Query for a set of assets and select one from the results list.

The asset page displays. By default, the performance chart shows two

metrics for time period selected for the asset page. For example, for a storage, the performance chart shows latency and total IOPS by default. The Resources section displays the resource name and an Additional resources section, which enables you to search for assets. Depending on the asset, you might also see assets in the Top correlated, Top contributor, Greedy, and Degraded sections. If there are no assets relevant to these sections, they are not displayed.

. You can add a performance chart for a metric by clicking *Display Metrics* and selecting the metrics you want displayed.

+

A separate chart is displayed for each metric selected. The chart displays the data for the selected time period. You can change the time period by clicking on another time period in the top right corner of the asset page, or by zooming in on any chart.

+

Click on *Display Metrics* to de-select any chart. The performance chart for the metric is removed from Expert View.

. You can position your cursor over the chart and change the metric data that displays for that chart by clicking any of the following, depending on the asset:

+

* Read, Write, or Total

+

* Tx, Rx, or Total

+

Total is the default.

+

You can drag your cursor over the data points in the chart to see how the value of the metric changes over the time period selected.

. In the Resources section, you can add any related assets to the performance charts:

+

* You can select a related asset in the *Top Correlated*, *Top Contributors*, *Greedy*, and *Degraded* sections to add data from that asset to the performance chart for each selected metric.

+

After you select the asset, a color block appears next to the asset to denote the color of its data points in the chart.

. Click on *Hide Resources* to hide the additional resources pane. Click on *Resources* to show the pane.

+

* For any asset shown, you can click the asset name to display its asset page, or you can click the percentage that the asset correlates or contributes to the base asset to view more information about the asset's

relation to the base asset. For example, clicking the linked percentage next to a top correlated asset displays an informational message comparing the type of correlation that asset has with the base asset. * If the Top correlated section does not contain an asset you want to display in a performance chart for comparison purposes, you can use the Search assets box in the Additional resources section to locate other assets. After you select an asset, it displays in the additional resources section. When you no longer want to view information about the asset, click image:TrashCanIcon.png[Delete]. [[ID8fd4c7b6c09ae7f12b4c0a968fd7c1aa]] = User Data Section :toc: macro :hardbreaks: :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead] The User Data section of an asset page displays and enables you to change any user-defined data such as applications and annotations. == Using the User Data section to assign or modify applications You can assign applications running in your environment to certain assets (host, virtual machines, volumes, internal volumes, qtrees, and hypervisors). The User Data section enables you to add, change, or remove the applications assigned to an asset. For all of these asset types except for volumes, you can assign more than one application. .Steps . Locate an asset page by doing any of the following: .. Query for a list of assets and then select one from the list. .. On a Dashboard, locate an asset name and click it. .. Perform a search and choose an asset from the results.

The asset page displays. The User Data section of the page shows

currently-assigned applications or annotations.

To change the application assigned, or to assign an application or additional applications, drop down the *Application* list and select the application(s) you want to assign to the asset. You can type to search for an application, or select one from the list.

// If you choose an application that is associated with a business entity, the business entity is automatically assigned to the asset. In this case, when you place your cursor over the business entity name, the word derived displays. If you want to maintain the entity for only the asset and not the associated application, you can manually override the assignment of the application.

To remove an application, drop down the application list and un-check the application.

== Using the User Data section to assign or modify annotations

When customizing Cloud Insights to track data for your corporate requirements, you can define specialized notes called annotations, and assign them to your assets. The User Data section of an asset page displays annotations assigned to an asset and also enables you to change the annotations assigned to that asset.

.Steps

- . To add an annotation to the asset, in the User Data section of the asset page, click $^*+$ Annotation $^*.$
- . Select an annotation from the list.
- .. If the annotation type is list, date, or Boolean, select a value from the list.
- .. If the annotation type is text, type a value.
- . Click Save.

The annotation is assigned to the asset. You can later filter assets by annotation using a query.

If you want to change the value of the annotation after you assign it, drop down the annotation list and enter a different value.

If the annotation is of list type for which the _Add new values on the fly_ option is selected, you can type to add a new value in addition to selecting an existing value.

```
[[ID6187304bd1cd0feba7f4a45e93539b15]]
= Asset Page Related Alerts section
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
You can use the Related Alerts section of an asset page to see any alerts
that occur in your environment as a result of a monitor assigned to an
asset. Monitors generate alerts based on conditions you set, and enable
you to identify the implication and analyze the impact and root cause of
the problem in a manner that enables rapid and effective correction.
The following example shows a typical Related Alerts section that displays
on an asset page:
image:Alerts on Landing Page.png[Related Alerts Table]
The Related Alerts section enables you to view and manage the alerts that
occur in your network as the result of monitor conditions assigned to an
asset.
.Steps
* Locate an asset page by doing any of the following:
** Type the name of the asset in the Search area, and then select the
asset from the list.
** In a dashboard widget, click on the name of an asset.
** Query for a set of assets and select on from the results list.
The asset page displays. The Related Alerts section displays the time the
alert was triggered as well as current status of the alert and the monitor
that triggered it. You can click the Alert ID to open the landing page for
the alert for further investigation.
////
```

You can perform any of the following optional tasks:

- * Use the filter box to show only specific alerts.
- * Change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header.
- * Click the asset name in any description to display its asset page; a red circle indicates issues that need further investigation.
- * You can click the performance policy, which displays the Edit Policy dialog box, to review the performance policy and make changes to the policy if necessary.

If you determine the issue is no longer a cause for concern, click the "three dots" menu on the right and select "Dismiss Violation" to remove a violation from the list.

[[IDa8c44a8cecdff13cb638ca44ac762e4d]]

= Storage Virtualization

:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:

:toc: macro

:relative path: ./

:imagesdir: {root path}{relative path}./media/

[.lead]

Cloud Insights can differentiate between a storage array having local storage or virtualization of other storage arrays. This gives you the ability to relate cost and distinguish performance from the front-end all the way to the back-end of your infrastructure.

=== Virtualization in a Table Widget

One of the easiest ways to begin looking at your storage virtualization is to create a dashboard table widget showing Virtualized type. When building the query for the widget, simply add "virtualizedType" to your grouping or filter.

image:StorageVirtualization_TableWidgetSettings.png[Basic query to show storage virtualizedTypes]

The resulting table widget shows you the _Standard_, _Backend_, and Virtual storages in your environment.

image:StorageVirtualization_TableWidgetShowingVirtualizedTypes.png[Table
showing virtualized types]

=== Landing Pages show Virtualized information

On a storage, volume, internal volume, or disk landing page, you can see relevant virtualization information. For example, looking at the storage landing page below, you can see that this is a Virtual storage, and which backend storage system applies. Any relevant tables on landing pages will also show virtualization information as applicable.

image:StorageVirtualization_StorageSummary.png[Storage Landing Page
showing Virtual and backed storage information]

//In the _Volumes_ table on that landing page, you can also see virtualization information, and selecting one of those volumes to display its landing page will also display virtualization information in the Summary and relevant tables, including a table showing the Internal Volumes as well as a table listing Virtual Volume Relations, if any.

=== Existing landing pages and dashboards

Be aware that if you currently have customized landing pages or dashboards in your environment, these will not automatically show all virtualization information by default. However, you can _Revert to Default_ any customized dashboard or landing page (you will have to re-implement your customizations), or modify the relevant widgets to include the desired virtualization attributes or metrics.

Revert to Default is available in the upper-right corner of a custom dashboard or landing page screen.

image:RevertToDefault.png[Revert to Default button]

[[IDd95f3ff39196ae8deeb85f406310744f]]

= Hints and Tips to Search for Assets and Alerts

:toc: macro
:hardbreaks:
:toclevels: 1

```
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Multiple search techniques can be used to search for data or objects in
your monitored environment.
* *Wildcard search*
You can perform multiple character wildcard search using the * character.
For example, applic*n would return application .
* *Phrases used in search*
A phrase is a group of words surrounded by double quotation marks; for
example, "VNX LUN 5". You can use double quotes to search for documents
that contain spaces in their names or attributes.
* *Boolean Operators*
Using Boolean operators OR, AND, and NOT, you can combine multiple terms
to form a more complex query.
+
OR
The OR operator is the default conjunction operator.
If there is no Boolean operator between two terms, the OR operator is
The OR operator links two terms and finds a matching document if either of
the terms exists in a document.
For example, storage OR netapp searches for documents that contain
either _storage_ or _netapp_.
High scores are given to documents that match most of the terms.
AND
You can use the AND operator to find documents in which both the search
terms exist in a single document. For example, _storage AND netapp_
searches for documents that contain both storage and netapp .
```

You can use the symbol *&&* instead of the word AND. +

NOT +

When you use the NOT operator, all the documents that contain the term after NOT are excluded from the search results. For example, _storage NOT netapp_ searches for documents that contains only _storage_ and not _netapp_.

+

You can use the symbol *!* instead of the word NOT.

////

== Prefix and suffix search

As soon as you start typing a search string, the search engine does a prefix and suffix search to find the best match.

Exact matches are given a higher score than a prefix or suffix match. The score is calculated based on the distance of the search term from the actual search result. For example, we have three storages: "aurora", "aurora1", and "aurora11". Searching for "aur" will return all three storages. However, the search result for "aurora" will have the highest score because it has the closest distance to the prefix search string.

The search engine also searches for terms in reverse order, which allows you to perform a suffix search. For example, when you type "345" in the search box, the search engine searches for "345".

////

Search is case-insensitive.

== Search using indexed terms
Searches that match more of the indexed terms result in higher scores.

The search string is split into separate search terms by space. For example, the search string "storage aurora netapp" is split into three keywords: "storage", "aurora", and "netapp". The search is performed using all three terms. The documents that match most of these terms will have the highest score. The more information you provide, the better are the search results. For example, you can search for a storage by its name and model.

The UI displays the search results across categories, with the three top results per category. If you did not find an object that you were expecting, you can include more terms in the search string to improve the search results.

```
The following table provides a list of indexed terms that can be added to
the search string.
|===
|Category |Indexed terms
|Storage | "storage"
name
vendor
model
|StoragePool |"storagepool"
name
name of the storage
IP addresses of the storage
serial number of the storage
storage vendor
storage model
names for all associated internal volumes
names for all associated disks
|Internal Volume |"internal volume"
name
name of the storage
IP addresses of the storage
serial number of the storage
storage vendor
storage model
name of the storage pool
names of all associated shares
names of all associated applications
//and business entities
|Volume | "volume"
name
label
names of all internal volumes
name of the storage pool
name of the storage
IP addresses of the storage
serial number of the storage
storage vendor
storage model
|Storage Node | "storagenode"
name
```

```
name of the storage
IP addresses of the storage
serialnumber of the storage
storage vendor
storage model
|Host |"host"
name
IP addresses
names of all associated applications
//and business entities
|Datastore | "datastore"
name
virtual center IP
names of all volumes
names of all internal volumes
|Virtual Machines | "virtualmachine"
name
DNS name
IP addresses
name of the host
IP addresses of the host
names of all datastores
names of all associated applications
//and business entities
|Switches (regular and NPV) | "switch"
IP address
wwn
name
serial number
model
domain ID
name of the fabric
wwn of the fabric
|Application | "application"
name
tenant
line of business
business unit
project
|Tape |"tape"
```

```
IP address
name
serial number
vendor
|Port |"port"
wwn
name
|Fabric |"fabric"
wwn
name
|Storage Virtual Machine (SVM)|"storagevirtualmachine"
name
UUID
|===
:leveloffset: -1
= Reporting
:leveloffset: +1
[[ID4fd3f893bbc52bc194e70a687b5b820c]]
= Cloud Insights Reporting Overview
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights reporting is a business intelligence tool that enables you
to view pre-defined reports or create custom reports.
NOTE: The Reporting feature is available in Cloud Insights
xref:{relative path}concept subscribing to cloud insights.html[Premium
Edition].
```

With Cloud Insights reporting you can perform the following tasks:

- * Run a pre-defined report
- * Create a custom report
- * Customize a report's format and delivery method
- * Schedule reports to run automatically
- * Email reports
- * Use colors to represent thresholds on data

Cloud Insights Reporting can generate custom reports for areas like chargeback, consumption analysis, and forecasting, and can help answer questions such as the following:

- * What inventory do I have?
- * Where is my inventory?
- * Who is using our assets?
- * What is the chargeback for allocated storage for a business unit?
- * How long until I need to acquire additional storage capacity?
- * Are business units aligned along the proper storage tiers?
- * How is storage allocation changing over a month, quarter, or year?
- == Accessing Cloud Insights Reporting

You can access Cloud Insights Reporting by clicking the *Reports* link in the menu.

//image:ReportsMenu.png[Reports Menu Link]

You will be taken to the Reporting interface. Cloud Insights uses IBM Cognos Analytics for its reporting engine.

== What is ETL?

When working with Reporting, you will hear the terms "Data Warehouse" and "ETL". ETL stands for "Extract, Transform, and Load". The ETL process retrieves data collected in Cloud Insights, and transforms the data into a format for use in Reporting. "Data Warehouse" refers to the collected data available for Reporting.

The ETL process includes these individual processes:

- * *Extract*: Takes data from Cloud Insights.
- * *Transform*: Applies business logic rules or functions to the data as it is extracted from Cloud Insights.
- $\mbox{\ensuremath{\star}}$ *Load*: Saves the transformed data into the data warehouse for use in Reporting.

```
[[ID2bc95ddd7e1306e9b54874b4c82bee7d]]
= Cloud Insights Reporting User Roles
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
If you have Cloud Insights Premium Edition with Reporting, every Cloud
Insights user in your environment also has a Single Sign-On (SSO) login to
the Reporting application (i.e. Cognos). Simply click the *Reports* link
in the menu and you will automatically be logged in to Reporting.
Your user role in Cloud Insights determines your Reporting user role:
|===
|Cloud Insights Role|Reporting Role|Reporting Permissions
|Guest|Consumer|Can view, schedule, and run reports and set personal
preferences such as those for languages and time zones. Consumers cannot
create reports or perform administrative tasks.
|User|Author|Can perform all Consumer functions as well as create and
manage reports and dashboards.
|Administrator|Administrator|Can perform all Author functions as well as
all administrative tasks such as configuration of reports and the shutdown
and restart of reporting tasks.
|===
The following table shows the functions available to each Reporting role.
|===
          |Consumer |Author |Administrator
|Feature
|View reports in the Team Content tab | Yes | Yes
                                                      |Yes
|Run reports | Yes
                       |Yes
|Schedule reports | Yes | Yes
|Upload external files |No |Yes |Yes
|Create Jobs| No|Yes|Yes
//|Create Users|No|No|Yes
```

|Create stories |No |Yes |Yes | Create reports |No |Yes |Yes | Create Packages and Data Modules |No |Yes|Yes | Perform administrative tasks |No |No |Yes |
|Add/Edit HTML Item |No |No |Yes | Run report with HTML Item |Yes |Yes |Yes |Add/Edit Custom SQL |No |No |Yes | Run reports with Custom SQL |Yes |Yes |Yes |Yes |

== Setting Reporting (Cognos) email preferences

NOTE: If you change your user email preferences within Cloud Insights Reporting (i.e. the Cognos application), those preferences are active _only for the current session_. Logging out of Cognos and back in again will reset your email preferences.

== Important note for existing customers

If you are new to Cloud Insights with Reporting, welcome! There is nothing more you need to do to begin enjoying Reporting.

If you are a current Premium Edition customer, SSO is not automatically enabled for your environment. When you enable SSO, the administrator user for the reporting portal (Cognos) ceases to exist. This means that any reports that are in the _My Content_ folder are removed and must be reinstalled or re-created in _Team Content_. Additionally, scheduled reports will need to be configured once SSO is enabled.

=== What steps should I take to prepare my existing environment for enabling SSO?

To ensure your reports are retained, migrate all reports from _My Content_ to _Team Content_ using the following steps. You must do this prior to enabling SSO in your environment:

. Navigate to *Menu > Content*

image:Reporting Menu.png[Cognos upper-left Menu]

- . Create a new folder in * Team Content *
- .. If multiple users have been created, please create a separate folder for each user to avoid overwriting reports with duplicate names

- . Navigate to My Content
- . Select all of the reports you wish to retain.
- . In the upper right corner of the menu, select "Copy or move"
- . Navigate to the newly created folder in _Team Content
- . Paste the reports to the newly created folder using the "Copy to" or "Move to" buttons $\ \ \,$
- . Once SSO is enabled for Cognos, log into Cloud Insights with the email address used to create your account.
- . Navigate to the _Team Content_ folder within Cognos, and Copy or Move the previously saved reports back to My Content .

```
[[ID83057b363dd656776131892a9aa84ab4]]
= Predefined Reports Made Easy
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

Cloud Insights Reporting includes predefined reports that address a number of common reporting requirements, providing critical insight that stakeholders need to make informed decisions about their storage infrastructure.

NOTE: The Reporting feature is available in Cloud Insights xref:{relative_path}concept_subscribing_to_cloud_insights.html[Premium Edition].

You can generate pre-defined reports from the Cloud Insights Reporting Portal, email them to other users, and even modify them. Several reports

enable you to filter by device, business entity, or tier. The reporting tools use IBM Cognos as a foundation and give you many data presentation options.

The pre-defined reports show your inventory, storage capacity, chargeback, performance, storage efficiency, and cloud cost data. You can modify these pre-defined reports and save your modifications.

You can generate reports in various formats, including HTML, PDF, CSV, XML, and Excel.

////

Cloud Insights accommodates multiple tenancy in reporting by enabling you to associate users with business units. With this feature, administrators can separate data or reports according to the attributes of a user or his/her affiliation.

////

== Navigating to Pre-defined Reports

When you open the Reporting Portal, the _Team Content_ folder is the starting point for you to select the type of information that you require in the Cloud Insights reports.

- . In the left navigation pane, select *Content > Team Content*.
- . Select *Reports* to access the pre-defined reports.

image:Reporting_Menu.png[Reporting menu]

image:Reporting_Team_Content.png[Team Content showing Reports highlighted]

== Using predefined reports to answer common questions

The following predefined reports are available in *Team content > Reports*.

=== Application Service Level Capacity and Performance
The Application Service Level Capacity and Performance report provides a
high level overview of your applications. You can use this information for
capacity planning or for a migration plan.

=== Chargeback

The Chargeback report provides storage capacity chargeback and accountability information by hosts, application, and business entities, and includes both current and historical data.

To prevent double counting do not include ESX servers, only monitor the VMs.

//An updated version of this report is available at the NetApp Storage Automation Store.

=== Data Sources

The Data Sources report shows all the data sources that are installed on your site, the status of the data source (success/failure), and status messages. The report provides information about where to start troubleshooting data sources. Failed data sources impact the accuracy of reporting and the general usability of the product.

=== ESX vs VM Performance

The ESX vs VM Performance report provides a comparison of ESX servers and VMs, showing average and peak IOPs, throughput, and latency and utilizations for ESX servers and VMs. To prevent double counting, exclude the ESX servers; only include the VMs.

An updated version of this report is available at the NetApp Storage Automation Store.

=== Fabric Summary

The Fabric Summary report identifies switches and switch information, including port counts, firmware versions, and license status. The report does not include NPV switch ports.

=== Host HBAs

The Host HBAs report provides an overview of the hosts in the environment and provides the vendor, model, and firmware version of HBAs, and the firmware level of the switches to which they are connected. This report can be used to analyze firmware compatibility when planning a firmware upgrade for a switch or an HBA.

=== Host Service Level Capacity and Performance
The Host Service Level Capacity and Performance report provides an overview of storage utilization by host for block only applications.

=== Host Summary

The Host Summary report provides an overview of storage utilization by each selected host with information for Fibre Channel and iSCSI hosts. The report enables you to compare ports and paths, the Fibre Channel and ISCSI capacity, and violation counts.

=== License Details

The License Details report shows the entitled quantity of resources you are licensed for across all sites with active licenses. The report also

shows a summation of actual quantity across all the sites with active licenses. The summation may include overlaps of storage arrays managed by multiple servers.

=== Mapped but not Masked Volumes

The Mapped but not Masked Volumes report lists the volumes whose logical unit number (LUN) has been mapped for use by a particular host, but is not masked to that host. In some cases these could be decommissioned LUNs that have been unmasked. Unmasked volumes can be accessed by any host, making them vulnerable to data corruption.

=== NetApp Capacity and Performance

The NetApp Capacity and Performance report provides global data for allocated, utilized, and committed capacity with trending and performance data for NetApp capacity.

=== Scorecard

The Scorecard report provides a summary and general status of all assets acquired by Cloud Insights. Status is indicated with green, yellow, and red flags:

- * Green indicates normal condition
- * Yellow indicates a potential issue in the environment
- * Red indicates an issue that requires attention

All of the fields in the report are described in the Data Dictionary provided with the report.

=== Storage Summary

The Storage Summary report provides a global summary of used and unused capacity data for raw, allocated, storage pools, and volumes. This report provides an overview of all of the storage discovered.

 $//{\rm A}$ newer version of this report is available at the NetApp Storage Automation Store.

=== VM Capacity and Performance

Describes the virtual machine (VM) environment and its capacity usage. VM tools must be enabled to view some data, such as when VMs were powered down.

=== VM Paths

The VM Paths report provides data store capacity data and performance metrics for which virtual machine is running on which host, which hosts are accessing which shared volumes, what the active access path is, and what comprises capacity allocation and usage.

=== HDS Capacity by Thin Pool The HDS Capacity by Thin Pool report shows the amount of usable capacity on a storage pool that is thin provisioned. === NetApp Capacity by Aggregate The NetApp Capacity by Aggregate report shows raw total, total, used, available, and committed space of aggregates. === Symmetrix Capacity by Thick Array The Symmetrix Capacity by Thick Array report shows raw capacity, useable capacity, free capacity, mapped, masked, and total free capacity. === Symmetrix Capacity by Thin Pool The Symmetrix Capacity by Thin Pool report shows raw capacity, useable capacity, used capacity, free capacity, used percentage, subscribed capacity, and subscription rate. === XIV Capacity by Array The XIV Capacity by Array report shows used and unused capacity for the array. === XIV Capacity by Pool The XIV Capacity by Pool report shows used and unused capacity for storage pools. [[ID139fef8627ad58f0ca431a31226acf92]] = Storage Manager Dashboard :toc: macro :hardbreaks: :toclevels: 2 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead] The Storage Manager Dashboard provides you with a centralized visualization that enables you to compare and contrast resource usage over time against the acceptable ranges and previous days of activity. Showing only the key performance metrics for your storage services, you can make decisions about how to maintain your data centers. NOTE: The Reporting feature is available in Cloud Insights xref:{relative path}concept subscribing to cloud insights.html[Premium

Edition].

== Summary

Selecting *Storage Manager Dashboard* from Team Content gives you several reports that provide information on of your traffic and storage.

image:Reporting_Storage_Manager_Dashboard_Choices.png[Storage Manager
Dashboard Options]

For an at-a-glance view, the *Storage Manager Report* comprises seven components that contain contextual information on many aspects of your storage environment. You can drill down on the aspects of your storage services to perform an in-depth of analysis of a section that interests you most.

image:Reporting-SMD.png[Storage Manager Dashboard]

This component shows the used versus usable storage capacity, total switch ports versus the number of switch ports connected, and total connected switch port utilization versus the total bandwidth, and how each of these trend over time. You can view the actual utilization compared against the low, mid, and high ranges, which enables you to compare and contrast usage between projections and your desired actuals, based on a target. For capacity and switch ports, you can configure this target. The forecast is based on an extrapolation of the current growth rate and the date you set. When the forecasted used capacity, which is based on future usage projection date, exceeds the target, an alert (solid red circle) appears next to Capacity.

=== Storage Tiers Capacity

This component shows the tier capacity used versus the capacity allocated to the tier, which indicates how the used capacity increases or decreases over a 12-month period and how many months are remaining to full capacity. Capacity usage is shown with values provided for actual usage, the usage forecast, and a target for capacity, which you can configure. When the forecasted used capacity, which is based on future usage projection date, exceeds the target capacity, an alert (solid red circle) appears next to a tier.

You can click any tier to display the Storage Pools Capacity and Performance Details report, which shows free versus used capacities, number of days to full, and performance (IOPS and Response Time) details for all the pools in the selected tier. You can also click any storage or storage pool name in this report to display the asset page summarizing the

current state of that resource.

=== Daily Storage Traffic

This component shows how the environment is performing, if there is any large growth, changes, or potential issues compared to the previous six months. It also shows the average traffic versus the traffic for the previous seven days, and for the previous day. You can visualize any abnormalities in the way the infrastructure is performing because it provides information that highlights both cyclical (previous seven days) and seasonal variations (previous six months).

You can click the title (Daily Storage Traffic) to display the Storage Traffic Details report, which shows the heat map of the hourly storage traffic for the previous day for each storage system. Click any storage name in this report to display the asset page summarizing the current state of that resource.

=== Data Centers Time to Full

This component shows all the data centers versus all of the tiers and how much capacity remains in each data center for each tier of storage based on forecasted growth rates. Tier capacity level is shown in blue; the darker the color, the lesser time the tier at the location has left before it is full.

You can click a section of a tier to display the Storage Pools Days to Full Details report, which shows total capacity, free capacity, and number of days to full for all the pools in the selected tier and the data center. Click any storage or storage pool name in this report to display the asset page summarizing the current state of that resource.

=== Top 10 Applications

This component shows the top 10 applications based on the used capacity. Regardless of how the tier organizes the data, this area displays the current used capacity and share of the infrastructure. You can visualize the range of user experience for the previous seven days to see if consumers experience acceptable (or, more importantly, unacceptable) response times.

This area also shows trending, which indicates if the applications meet their performance service level objectives (SLO). You can view the previous week's minimum response time, the first quartile, the third quartile, and the maximum response time, with a median shown against an acceptable SLO, which you can configure. When the median response time for any application is out of the acceptable SLO range, an alert (solid red circle) appears next to the application. You can click an application to display the asset page summarizing the current state of that resource.

=== Storage Tiers Daily Performance

This component shows a summary of the tier's performance for response time and IOPS for the previous seven days. This performance is compared against a SLO, which you can configure, enabling you to see if there is opportunity to consolidate tiers, realign workloads delivered from those tiers, or identify issues with particular tiers. When median response time or median IOPS is out of the acceptable SLO range, an alert (solid red circle) appears next to a tier.

You can click a tier name to display the Storage Pools Capacity and Performance Details report, which shows free versus used capacities, number of days to full, and performance (IOPS and response time) details for all the pools in the selected tier. Click any storage or storage pool in this report to display the asset page summarizing the current state of that resource.

=== Orphaned Capacity

This component shows the total orphaned capacity and orphaned capacity by tier, comparing it against acceptable ranges for total usable capacity and showing the actual capacity that is orphaned. Orphaned capacity is defined by configuration and by performance. Storage orphaned by configuration describes a situation in which there is storage allocated to a host. However, the configuration has not been performed properly and the host cannot access the storage. Orphaned by performance is when the storage is correctly configured to be accessed by a host. However, there has been no storage traffic.

The horizontal stacked bar shows the acceptable ranges. The darker the gray, the more unacceptable the situation is. The actual situation is shown with the narrow bronze bar that shows the actual capacity that is orphaned.

You can click a tier to display the Orphaned Storage Details report, which shows all the volumes identified as orphaned by configuration and performance for the selected tier. Click any storage, storage pool, or volume in this report to display the asset page summarizing the current state of that resource.

[[IDa8ec453cc78b114e61ab5cff36a3e9ec]]
= Creating a Report (Example)
:toc: macro

:hardbreaks:
:toclevels: 2
:icons: font

```
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
Use the steps in this example to generate a simple report on physical
capacity of storage and storage pools in a number of data centers.
.Steps
. Navigate to *Menu > Content > Team Content > Reports*
. In the upper-right of the screen, select *[New +]*
. Select *Report*
image:Reporting New Report.png[Creating a New Report]
. On the *Templates* tab, select Blank
The Source and Data tabs is displayed
. Open *Select a source +*
. Under *Team content*, open *Packages*
A list of available packages is displayed.
. Choose *Storage and Storage Pool Capacity*
image: Reporting Select Source For Report.png[Selecting a source for the
report]
+
. Select *Open*
The available styles for your report are displayed.
. Select *List*
Add appropriate names for List and Query
. Select *OK*
. Expand Physical Capacity
. Expand to the lowest level of Data Center
. Drag Data Center to the Reporting palate.
. Expand Capacity (MB)
. Drag _Capacity (MB) _ to the Reporting palate.
. Drag Used Capacity (MB) to the Reporting palate.
. Run the report by selecting an output type from the *Run* menu.
image:Reporting Running A Report.png[Selecting a Report Output]
.Result
```

```
A report similar to the following is created:
image:Reporting-Example1.png[Report Example]
[[IDdc4ac107645563ff5c3bb56d257d3274]]
= Managing Reports
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
You can customize a report's output format and delivery, set report
properties or schedules, and email reports.
NOTE: The Reporting feature is available in Cloud Insights
xref:{relative path}concept subscribing to cloud insights.html[Premium
Edition].
== Customizing a report's output format and delivery
You can customize the format and delivery method of reports.
. In the Cloud Insights Reporting Portal, Go to *Menu > Content > My
Content/Team Content*. Mouse over the report you want to customize and
open the "three dots" menu.
//image:ReportCustomizationMenu.png[Report Customization Menu]
image:Reporting Output and Delivery.png[Report Output and Delivery]
. Click *Properties > Schedule*
//image:ReportSchedule.png[Scheduling a Report]
. You can set the following options:
** *Schedule* when you want reports to run.
** Choose *Options* for report format and delivery (Save, Print, Email)
and Languages for the report.
. Click *Save* to produce the report using the selections you made.
```

== Copying a report to the clipboard

Use this process to copy a report to the clipboard.

- . Select a report to copy from (*Menu > Content > My Content or Team Content*)
- . Choose $_{\tt Edit}$ report $_$ from the report's drop-down menu

+

image:Reporting_Edit_Report.png[Editing a Report]

+

- . In the upper-right of the screen, open the "three dots" menu next to "Properties".
- . Select *Copy Report to Clipboard*.

+

image:Reporting_Copy_To_Clipboard.png[Copying a report to the clipboard]

== Opening reports from the clipboard

You can open a report specification that was previously copied to the clipboard.

About this task

Start by creating a new report or opening an existing report that you wish to replace with the copied report. The steps below are for a new report.

- . Select *Menu > +New > Report* and create a blank report.
- . In the upper-right of the screen, open the "three dots" menu next to "Properties".
- . Select *Open Report from Clipboard*.

image:Reporting_Open_From_Clipboard.png[Opening a report from the clipboard]

- . Paste the copied code into the window and select *OK*.
- . Select the floppy disk icon to save the report.
- . Choose where to save the report (_My Content_, _Team Content_, or create a new folder).
- . Give the new report a meaningful name and select *Save*.
- == Editing an existing report

Be aware that editing files in their default location runs a risk of those reports being overwritten upon the next report catalog refresh. It is recommended to save the edited report under a new name or store it in a

non-default location. == Troubleshooting Here you will find suggestions for troubleshooting problems with Reporting. |=== |*Problem:* |*Try this:* |When scheduling a report to be sent via email, the name of the user logged in is pre-populated to the email's "To" field. However, the name is in the form of "firstname lastname" (first name, space, last name). Since this is not a valid email address, the email will fail to send when the scheduled report is run. |When scheduling the report to be sent via email, clear the pre-populated name and enter a valid, properly-formatted email address in the "To" field. |=== [[ID49e684accaa92d0946f9b49132e46a9b]] = Creating Custom Reports :toc: macro :hardbreaks: :toclevels: 2 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead] You can use the report authoring tools to create custom reports. After creating reports, you can save them and run them on a regular schedule. The results of reports can be automatically sent by email to yourself and others. NOTE: The Reporting feature is available in Cloud Insights xref:{relative path}concept subscribing to cloud insights.html[Premium Edition]. The examples in this section show the following process, which can be used for any of the Cloud Insights Reporting data models:

- * Identifying a question to be answered with a report
- * Determining the data needed to support the results
- * Selecting data elements for the report

Before designing your custom report, you need to complete some prerequisite tasks. If you do not complete these, reports could be inaccurate or incomplete.

For example, if you do not finish the device identification process, your capacity reports will not be accurate. Or, if you do not finish setting annotations (such as tiers, business units, and data centers), your custom reports might not accurately report data across your domain or might show "N/A" for some data points.

Before you design your reports, complete the following tasks:

- * Configure all
- xref:{relative_path}task_configure_data_collectors.html[data collectors]
 properly.
- * Enter annotations (such as tiers, data centers, and business units) on devices and resources in your environment. It is beneficial to have annotations stable before generating reports, because Cloud Insights Reporting collects historical information.
- //* Configure OnCommand Insight Data Warehouse to accept the data from the OnCommand Insight server in the Extract, Transform, and Load (ETL) process.
- == Report Creation Process

The process of creating custom (also called "ad hoc") reports involves several tasks:

- * Plan the results of your report.
- * Identify data to support your results.
- * Select the data model (for example, Chargeback data model, Inventory data model, and so on) that contains the data.
- * Select data elements for the report.
- * Optionally format, sort, and filter report results.
- === Planning the Results of Your Custom Report

Before you open the report authoring tools, you might want to plan the results you want from the report. With report authoring tools, you can create reports easily and might not need a great deal of planning; however, it is a good idea to get a sense from the report requestor about the report requirements.

- * Identify the exact question you want to answer. For example:
- ** How much capacity do I have left?
- ** What are the chargeback costs per business unit?
- ** What is the capacity by tier to ensure that business units are aligned at the proper tier of storage?
- ** How can I forecast power and cooling requirements? (Add customized metadata by adding annotations to resources.)
- * Identify the data elements that you need to support the answer.
- * Identify the relationships between data that you want to see in the answer. Do not include illogical relationships in your question, for example, "I want to see the ports that relate to capacity."
- * Identify any calculations needed on data.
- * Determine what types of filtering are needed to limit the results.
- * Determine if you need to use current or historical data.
- * Determine if you need to set access privileges on reports to limit the data to specific audiences.
- * Identify how the report will be distributed. For example, should it be emailed on a set schedule or included in the Team content folder area?
- * Determine who will maintain the report. This might affect the complexity of the design.
- * Create a mockup of the report.
- === Tips for designing reports
 Several tips might be helpful when you are designing reports.
- * Determine whether you need to use current or historical data.

+

Most reports only need to report on the latest data available in the ${\tt Cloud}$ ${\tt Insights.}$

- * Cloud Insights Reporting provides historical information on capacity and performance, but not on inventory.
- * Everybody sees all data; however, you might need to limit data to specific audiences.

+

To segment the information for different users, you can create reports and set access permissions on them.

== Reporting data models

Cloud Insights includes several data models from which you can either select predefined reports or create your own custom report.

Each data model contains a simple data mart and an advanced data mart:

* The simple data mart provides quick access to the most commonly used

data elements and includes only the last snapshot of Data Warehouse data; it does not include historical data.

- * The advanced data mart provides all values and details available from the simple data mart and includes access to historical data values.
- === Capacity data models

Enables you to answer questions about storage capacity, file system utilization, internal volume capacity, port capacity, qtree capacity, and virtual machine (VM) capacity. The Capacity data model is a container for several capacity data models. You can create reports answering various types of questions using this data model:

==== Storage and Storage Pool Capacity data model

Enables you to answer questions about storage capacity resource planning, including storage and storage pools, and includes both physical and virtual storage pool data. This simple data model can help you answer questions related to capacity on the floor and the capacity usage of storage pools by tier and data center over time.

If you are new to capacity reporting, you should start with this data model because it is a simpler, targeted data model. You can answer questions similar to the following using this data model:

- * What is the projected date for reaching the capacity threshold of 80% of my physical storage?
- * What is the physical storage capacity on an array for a given tier?
- * What is my storage capacity by manufacturer and family as well as by data center?
- * What is the storage utilization trend on an array for all of the tiers?
- * What are my top 10 storage systems with the highest utilization?
- * What is the storage utilization trend of the storage pools?
- * How much capacity is already allocated?
- * What capacity is available for allocation?

==== File System Utilization data model

This data model provides visibility about capacity utilization by hosts at the file system level. Administrators can determine allocated and used capacity per file system, determine the type of file system, and identify trending statistics by file system type. You can answer the following questions using this data model:

- * What is the size of the file system?
- * Where is the data kept and how is it accessed, for example, local or SAN?

* What are the historical trends for the file system capacity? Then, based on this, what can we anticipate for future needs?

==== Internal Volume Capacity data model

Enables you to answer questions about internal volume used capacity, allocated capacity, and capacity usage over time:

- * Which internal volumes have a utilization higher than a predefined threshold?
- * Which internal volumes are in danger of running out of capacity based on a trend?
- 8 What is the used capacity versus the allocated capacity on our internal volumes?

==== Port Capacity data model

Enables you to answer questions about switch port connectivity, port status, and port speed over time. You can answer questions similar the following to help you plan for purchases of new switches:

How can I create a port consumption forecast that predicts resource (port) availability (according to data center, switch vendor and port speed)?

- * Which ports are likely to run out of capacity, providing data speed, data center, vendor and number of Host and storage ports?
- * What are the switch port capacity trends over time?
- * What are the port speeds?
- * What type of port capacity is needed and which organization is about to run out of a certain port type or vendor?
- * What is the optimal time to purchase that capacity and make it available?

==== Qtree Capacity data model

Enables you to trend qtree utilization (with data such as used versus allocated capacity) over time. You can view the information by different dimensions—for example, by business entity, application, tier, and service level. You can answer the following questions using this data model:

- * What is the used capacity for qtrees versus the limits set per application or business entity?
- * What are the trends of our used and free capacity so that we can do capacity planning?
- * Which business entities are using the most capacity?
- * Which applications consume the most capacity?

==== VM Capacity data model

Enables you to report your virtual environment and its capacity usage. This data model lets you report on changes in capacity usage over time for VMs and data stores. The data model also provides thin provisioning and virtual machine chargeback data.

- * How can I determine capacity chargeback based on capacity provisioned to VMs and data stores?
- * What capacity is not used by VMs and which portion of unused is free, orphaned, or other?
- * What do we need to purchase based on consumption trends?
- * What are my storage efficiency savings achieved by using storage thin provisioning and deduplication technologies?

Capacities in the VM Capacity data model are taken from virtual disks (VMDKs). This means that the provisioned size of a VM using the VM Capacity data model is the size of its virtual disks. This is different from the provisioned capacity in the Virtual Machines view in Cloud Insights, which shows the provisioned size for the VM itself.

==== Volume Capacity data model

Enables you to analyze all aspects of the volumes in your environment and organize data by vendor, model, tier, service level, and data center.

You can view the capacity related to orphaned volumes, unused volumes, and protection volumes (used for replication). You can also see different volume technologies (iSCSI or FC), and compare virtual volumes to non-virtual volumes for array virtualization issues.

You can answer questions similar to the following with this data model:

- * Which volumes have a utilization higher than a predefined threshold?
- * What is the trend in my data center for orphan volume capacity?
- * How much of my data center capacity is virtualized or thin provisioned?
- * How much of my data center capacity must be reserved for replication?

=== Chargeback data model

Enables you to answer questions about used capacity and allocated capacity on storage resources (volumes, internal volumes, and qtrees). This data model provides storage capacity chargeback and accountability information by hosts, application, and business entities, and includes both current and historical data. Report data can be categorized by service level and storage tier.

You can use this data model to generate chargeback reports by finding the

amount of capacity that is used by a business entity. This data model enables you to create unified reporting of multiple protocols (including NAS, SAN, FC, and iSCSI).

- * For storage without internal volumes, chargeback reports show chargeback by volumes.
- * For storage with internal volumes:
- ** If business entities are assigned to volumes, chargeback reports show chargeback by volumes.
- ** If business entities are not assigned to volumes but assigned to qtrees, chargeback reports show chargeback by qtrees.
- ** If business entities are not assigned to volumes and not assigned to gtrees, chargeback reports show the internal volume.
- ** The decision whether to show chargeback by volume, qtree or internal volume is made per each internal volume, so it is possible for different internal volumes in the same storage pool to show chargeback at different levels.

Capacity facts are purged after a default time interval. For details, see Data Warehouse processes.

Reports using the Chargeback data model might display different values than reports using the Storage Capacity data model.

- * For storage arrays that are not NetApp storage systems, the data from both data models is the same.
- * For NetApp and Celerra storage systems, the Chargeback data model uses a single layer (of volumes, internal volumes, or qtrees) to base its charges, while the Storage Capacity data model uses multiple layers (of volumes and internal volumes) to base its charges.

=== Inventory data model

Enables you to answer questions about inventory resources including hosts, storage systems, switches, disks, tapes, qtrees, quotas, virtual machines and servers, and generic devices. The Inventory data model includes several submarts that enable you to view information about replications, FC paths, iSCSI paths, NFS paths, and violations. The Inventory data model does not include historical data. Questions you can answer with this data

- * What assets do I have and where are they?
- * Who is using the assets?
- * What types of devices do I have and what are components of those devices?
- * How many hosts per OS do I have and how many ports exist on those hosts?

- * What storage arrays per vendor exist in each data center?
- * How many switches per vendor do I have in each data center?
- * How many ports are not licensed?
- * What vendor tapes are we using and how many ports exist on each tape?re all the generic devices identified before we begin working on reports?
- * What are the paths between hosts and storage volumes or tapes?
- * What are the paths between generic devices and storage volumes or tapes?
- * How many violations of each type do I have per data center?
- * For each replicated volume, what are the source and target volumes?
- * Do I have any firmware incompatibilities or port speed mismatches between Fibre Channel host HBAs and switches?

=== Performance data model

Enables you to answer questions about performance for volumes, application volumes, internal volumes, switches, applications, VMs, VMDKs, ESX versus VM, hosts, and application nodes. Many of these report _Hourly_ data, _Daily_ data, or both. Using this data model, you can create reports that answer several types of performance management questions:

- * What volumes or internal volumes have not been used or accessed during a specific period?
- * Can we pinpoint any potential misconfiguration for storage for an application (unused)?
- * What was the overall access behavior pattern for an application?
- * Are tiered volumes assigned appropriately for a given application?
- * Could we use cheaper storage for an application currently running without impact to application performance?
- * What are the applications that are producing more accesses to currently configured storage?

When you use the switch performance tables, you can obtain the following information:

- * Is my host traffic through connected ports balanced?
- * Which switches or ports are exhibiting a high number of errors?
- * What are the most used switches based on port performance?
- * What are the underutilized switches based on port performance?
- * What is the host trending throughput based on port performance?
- * What is the performance utilization for last X days for one specified host, storage system, tape, or switch?

- * Which devices are producing traffic on a specific switch (for example, which devices are responsible for use of a highly utilized switch)?
- * What is the throughput for a specific business unit in our environment?

When you use the disk performance tables, you can obtain the following information:

- * What is the throughput for a specified storage pool based on disk performance data?
- * What is the highest used storage pool?
- * What is the average disk utilization for a specific storage?
- * What is the trend of usage for a storage system or storage pool based on disk performance data?
- * What is the disk usage trending for a specific storage pool?

When you use VM and VMDK performance tables, you can obtain the following information:

- * Is my virtual environment performing optimally?
- * Which VMDKs are reporting the highest workloads?
- * How can I use the performance reported from VMDs mapped to different datastores to make decisions about re-tiering.

The Performance data model includes information that helps you determine the appropriateness of tiers, storage misconfigurations for applications, and last access times of volumes and internal volumes. This data model provides data such as response times, IOPs, throughput, number of writes pending, and accessed status.

=== Storage Efficiency data model

Enables you to track the storage efficiency score and potential over time. This data model stores measurements of not only the provisioned capacity, but also the amount that is used or consumed (the physical measurement). For example, when thin provisioning is enabled, Cloud Insights indicates how much capacity is taken from the device. You can also use this model to determine efficiency when deduplication is enabled. You can answer various questions using the Storage Efficiency data mart:

* What is our storage efficiency savings as a result of implementing thin

provisioning and deduplication technologies?

- * What are the storage savings across data centers?
- * Based on historical capacity trends, when do we need to purchase additional storage?
- * What would be the capacity gain if we enabled technologies such as thin provisioning and deduplication?
- * Regarding storage capacity, am I at risk now?
- === Data model fact and dimension tables

Each data model includes both fact and dimension tables.

- * Fact tables: Contain data that is measured, for example, quantity, raw and usable capacity. Contain foreign keys to dimension tables.
- * Dimension tables: Contain descriptive information about facts, for example, data center and business units. A dimension is a structure, often composed of hierarchies, that categorizes data. Dimensional attributes help describe the dimensional values.

Using different or multiple dimension attributes (seen as columns in the reports), you construct reports that access data for each dimension described in the data model.

=== Colors used in data model elements

Colors on data model elements have different indications.

- * Yellow assets: Represent measurements.
- * Non-yellow assets: Represent attributes. These values do not aggregate.
- === Using multiple data models in one report

Typically, you use one data model per report. However, you can write a report that combines data from multiple data models.

To write a report that combines data from multiple data models, choose one of the data models to use as the base, then write SQL queries to access the data from the additional data marts. You can use the SQL Join feature to combine the data from the different queries into a single query that you can use to write the report.

For example, say you want the current capacity for each storage array and you want to capture custom annotations on the arrays. You could create the report using the Storage Capacity data model. You could use the elements from the Current Capacity and dimension tables and add a separate SQL query to access the annotations information in the Inventory data model. Finally, you could combine the data by linking the Inventory storage data to the Storage Dimension table using the storage name and the join criteria.

```
[[ID0247ebfe26efee39e3b8b193a40b9980]]
= Access the Reporting Database via API
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights' powerful API allows users to query the Cloud Insights
Reporting database directly, without going through the Cognos Reporting
environment.
NOTE: This documentation refers to the Cloud Insights Reporting feature,
which is available in the
xref:{relative path}/concept subscribing to cloud insights.html#editions[P
remium Edition].
== Odata
The Cloud Insights Reporting API follows the
link:https://www.odata.org/[OData v4] (Open Data Protocol) standard for
its querying of the Reporting database.
For more information or to learn more, check out
link:https://www.odata.org/getting-started/basic-tutorial/[this tutorial]
on OData.
All requests will start with the url _\https://<Cloud Insights
URL>/rest/v1/dwh-management/odata
== Generating an APIKey
```

Read more about xref:{relative path}API Overview.html[Cloud Insights APIs]. To generate an API key, do the following: * Log into your Cloud Insights environment and select *Admin > API Access*. * Click "+ API Access Token". * Enter a Name & Description. * For type, choose Data Warehouse . * Set Permissions as Read/Write. * Set a desires Expiration date. * Click "Save", then *copy the key and save it* somewhere safe. You will not be able to access the full key later. APIkeys are good for Sync or Async . //<TBD More information> == Direct query of tables With the API Key in place, direct queries of the Reporting database are now possible. Long URLs may be simplified to \https://.../odata/ for display purposes rather than the full \https://<Cloud Insights URL>/rest/v1/dwh-management/odata/ Try simple queries like * \https://<Cloud Insights URL>/rest/v1/dwh-management/odata/dwh custom * \https://<Cloud Insights URL>/rest/v1/dwh-management/odata/dwh inventory * \https://<Cloud Insights URL>/rest/v1/dwhmanagement/odata/dwh inventory/storage * \https://<Cloud Insights URL>/rest/v1/dwhmanagement/odata/dwh inventory/disk * \https://.../odata/dwh custom/custom queries == REST API Examples The URL for all calls is \https://<Cloud Insights URL>/rest/v1/dwhmanagement/odata . * GET /{schema}/** - Retrieves data from the Reporting Database. Format: \https://<Cloud Insights URL>/rest/v1/dwh-

```
management/odata/<schema name>/<query>
Example:
https://<domain>/rest/v1/dwh-
management/odata/dwh inventory/fabric?$count=true&$orderby=name
Result:
 {
    "@odata.context": "$metadata#fabric",
    "@odata.count": 2,
    "value": [
        {
            "id": 851,
            "identifier": "10:00:50:EB:1A:40:3B:44",
            "wwn": "10:00:50:EB:1A:40:3B:44",
            "name": "10:00:50:EB:1A:40:3B:44",
            "vsanEnabled": "0",
            "vsanId": null,
            "zoningEnabled": "0",
            "url": "https://<domain>/web/#/assets/fabrics/941716"
        },
            "id": 852,
            "identifier": "10:00:50:EB:1A:40:44:0C",
            "wwn": "10:00:50:EB:1A:40:44:0C",
            "name": "10:00:50:EB:1A:40:44:0C",
            "vsanEnabled": "0",
            "vsanId": null,
            "zoningEnabled": "0",
            "url": "https://<domain>/web/#/assets/fabrics/941836"
        }
     1
}
// Example: get all Storage Pools of storage with id 287 with expanded
Internal Volume data: https://<Cloud Insights URL>/rest/v1/dwh-
management/odata/dwh inventory/storage(287)/storage pool?$expand=internal
volume
////
* POST /{schema}/** - Write data and create queries in dwh custom schema
of Data Warehouse database through ODATA protocol, requires ADMIN role
```

```
Format: https://<Cloud Insights URL>/rest/v1/dwh-
management/odata/<schema name>/. The body contains the record
in JSON format
Example: add a new record to the storage table: https://<Cloud Insights
URL>/rest/v1/dwh-management/odata/dwh custom/storage , Request body:
{"storageId": 123, "storageName": "storage123"}
Creating queries: POST https://<Cloud Insights URL>/rest/v1/dwh-
management/odata/dwh custom/custom queries -d '{"queryName":
"<query name>", "querySql": "<query sql>"}'
 * PATCH /{schema}/** - Modify data and modify queries in dwh custom
schema of Data Warehouse database through ODATA protocol, requires ADMIN
role
Format: https://<Cloud Insights URL>/rest/v1/dwh-
management/odata/<schema name>/('<record id>'). The body
contains the record in JSON format
Example: modify a record in the storage table: https://<Cloud Insights
URL>/rest/v1/dwh-management/odata/dwh custom/storage('123'), Request
body: {"storageId": 123, "storageName": "storage123"}
Modifying queries: PATCH https://<Cloud Insights URL>/rest/v1/dwh-
management/odata/dwh custom/custom queries('queryName') -d '{"queryName":
"<query name>", "querySql": "<query sql>"}'
* DELETE /{schema}/** - Delete data and delete queries in dwh custom
schema of Data Warehouse database through ODATA protocol, requires ADMIN
role
Format: https://<Cloud Insights URL>/rest/v1/dwh-
management/odata/<schema name>/('<record id>')
Example: delete a record from the storage table: https://<Cloud Insights
URL>/rest/v1/dwh-management/odata/dwh custom/storage('123')
Deleting queries: DELETE https://<Cloud Insights URL>/rest/v1/dwh-
management/odata/dwh custom/custom queries('queryName')
////
```

```
== Helpful Hints
Keep the following in mind when working with Reporting API queries.
* The query payload must be a valid JSON string
* The query payload must be contained in a single line
* Double quotes must be escaped, i.e. \"
* Tabs are supported as \t
* Avoid comments
* Lower-case table names are supported
Additionally:
* 2 Headers are required:
** Name "X-CloudInsights-ApiKey"
** Attribute Value "<apikey>"
Your API key will be specific to your Cloud Insights environment.
[[ID10dcbd24e8982bd95a4e5869840ba490]]
= How historical data is retained for Reporting
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights retains historical data for use in Reporting based on the
data marts and granularity of the data, as shown in the following table.
|===
|Data mart | Measured object | Granularity | Retention period
|Performance marts | Volumes and internal volumes | Hourly | 14 days
|Performance marts | Volumes and internal volumes | Daily | 13 months
|Performance marts | Application | Hourly | 13 months
|Performance marts | Host | Hourly | 13 months
|Performance marts | Switch performance for port | Hourly | 35 days
|Performance marts | Switch performance for host, storage, and tape |
Hourly | 13 months
|Performance marts | Storage node | Hourly | 14 days
```

```
|Performance marts | Storage node | Daily | 13 months
|Performance marts | VM performance | Hourly | 14 days
|Performance marts | VM performance | Daily | 13 months
|Performance marts | Hypervisor performance | Hourly | 35 days
|Performance marts | Hypervisor performance | Daily | 13 months
|Performance marts | VMDK performance | Hourly | 35 days
|Performance marts | VMDK performance | Daily | 13 months
|Performance marts | Disk performance | Hourly | 14 days
|Performance marts | Disk performance | Daily | 13 months
|Capacity marts | All (except individual volumes) | Daily | 13 months
|Capacity marts | All (except individual volumes) | Monthly representative
| 14 months and beyond
|Inventory marts | Individual volumes | Current state | 1 day (or until
|===
[[ID09bdedcdcce86d2f420f85fb726a2748]]
= Cloud Insights Reporting Schema Diagrams
:toc: macro
:hardbreaks:
:toclevekls: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This document provides the schema diagrams for the Reporting Database. You
can also download a file containing the
xref:{relative path}ci reporting database schema.pdf[schema tables].
NOTE: The Reporting feature is available in Cloud Insights
xref:{relative path}concept subscribing to cloud insights.html[Premium
Edition].
== Inventory Datamart
The following images describe the inventory datamart.
```

```
=== Annotations
image:annotations.png[]
=== Applications
image:apps annot.png[]
=== Kubernetes Metrics
image:k8s_schema.jpg[Kubernetes]
==== Kubernetes Cluster Metrics Fact
image:k8s_cluster_metrics_fact.jpg[Kubernetes Cluster Metrics Fact]
==== Kubernetes Namespace Metrics Fact
image:k8s namespace metrics fact.jpg[Kubernetes Namespace Metrics Fact]
==== Kubernetes Node Metrics Fact
image:k8s node metrics fact.jpg[Kubernetes Node Metrics Fact]
==== Kubernetes PVC Metrics Fact
image:k8s pvc metrics fact.jpg[Kubernetes PVC Metrics Fact]
==== Kubernetes Workload Metrics Fact
image:k8s workload metrics fact.jpg[Kubernetes Workload Metrics Fact]
=== NAS
image:nas.png[]
=== Paths and Violations
image:logical.png[]
=== Port Connectivity
```

```
image:connectivity.png[]
=== SAN Fabric
image:fabric.png[]
=== Storage
image:storage.png[]
=== Storage Node
image:storage_node.png[]
=== VM
image:vm.png[]
1 1 1
== Capacity Datamart
The following images describe the capacity datamart.
=== Chargeback
image:Chargeback_Fact.png[]
=== Disk Group Capacity
image:Disk_Group_Capacity.png[]
=== File System Utilization
image:fs util.png[]
=== Internal Volume Capacity
image:Internal_Volume_Capacity_Fact.png[]
```

```
=== Kubernetes PV Capacity
image:k8s pvc capacity fact.jpg[]
=== Port Capacity
image:ports.png[]
=== Qtree Capacity
image:Qtree_Capacity_Fact.png[]
=== Storage Capacity Efficiency
image:efficiency.png[]
=== Storage and Storage Pool Capacity
image:Storage_and_Storage_Pool_Capacity_Fact.png[]
=== Storage Node Capacity
image:Storage Node Capacity Fact.jpg[]
=== VM Capacity
image:VM_Capacity_Fact.png[]
=== Volume Capacity
image:Volume Capacity.png[]
```

```
1 1 1
== Performance Datamart
The following images describe the performance datamart.
=== Application Volume Hourly Performance
image:application performance fact.jpg[]
=== Disk Daily Performance
image:disk_daily_performance_fact.png[]
=== Disk Hourly Performance
image:disk_hourly_performance_fact.png[]
=== Host Hourly Performance
image:host performance fact.jpg[]
=== Internal Volume Hourly Performance
image:internal volume performance fact.jpg[]
=== Internal Volume Daily Performance
image:internal volume daily performance fact.jpg[]
=== Qtree Daily Performance
image:QtreeDailyPerformanceFact.png[]
=== Storage Node Daily Performance
```

```
image:storage node daily performance fact.jpg[]
=== Storage Node Hourly Performance
image:storage node hourly performance fact.jpg[]
=== Switch Hourly Performance for Host
image:switch performance for host hourly fact.png[]
=== Switch Hourly Performance for Port
image:switch performance for port hourly fact.png[]
=== Switch Hourly Performance for Storage
image:switch performance for storage hourly fact.png[]
=== Switch Hourly Performance for Tape
image:switch performance for tape hourly fact.png[]
=== VM Performance
image:vm hourly performance fact.png[]
=== VM Daily Performance for Host
image:vm daily performance fact.png[]
=== VM Hourly Performance for Host
image:vm_hourly_performance_fact.png[]
=== VM Daily Performance for Host
image:vm daily performance fact.png[]
=== VM Hourly Performance for Host
```

```
image:vm_hourly_performance_fact.png[]
=== VMDK Daily Performance
image:vmdk_daily_performance_fact.png[]
=== VMDK Hourly Performance
image:vmdk hourly performance_fact.png[]
=== Volume Hourly Performance
image:volume performance fact.jpg[]
=== Volume Daily Performance
image:volume daily_performance_fact.jpg[]
// == Kubernetes Capacity
// Moved to Capacity section above
////
=== Kubernetes PV Capacity
image:k8s pvc capacity fact.jpg[]
////
// Moved to Inventory section above:
////
=== Kubernetes Metrics
image:k8s schema.jpg[Kubernetes]
==== Kubernetes Cluster Metrics Fact
image:k8s cluster metrics fact.jpg[Kubernetes Cluster Metrics Fact]
==== Kubernetes Namespace Metrics Fact
```

```
image:k8s namespace metrics fact.jpg[Kubernetes Namespace Metrics Fact]
==== Kubernetes Node Metrics Fact
image:k8s node metrics fact.jpg[Kubernetes Node Metrics Fact]
==== Kubernetes PVC Metrics Fact
image:k8s pvc metrics fact.jpg[Kubernetes PVC Metrics Fact]
==== Kubernetes Workload Metrics Fact
image:k8s workload metrics fact.jpg[Kubernetes Workload Metrics Fact]
////
[[ID63181e2196487d92430e5ab3388a329a]]
= Cloud Insights Schemas for Reporting
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
These schema tables and diagrams are provided here as a reference for
Cloud Insights Reporting.
link:https://docs.netapp.com/us-
en/cloudinsights/ci_reporting_database_schema.pdf[*Schema Tables*] in .PDF
format. Click the link to open, or right-click and choose Save as... to
download.
xref:{relative path}reporting schema diagrams.html[*Schema Diagrams*]
NOTE: The Reporting feature is available in Cloud Insights
xref:{relative path}concept subscribing to cloud insights.html[Premium
Edition].
```

```
:leveloffset: -1
= Workload Security
:leveloffset: +1
[[ID4fec82d76406ae3f9c5e47b88a02a048]]
= About Storage Workload Security
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights Storage Workload Security (formerly Cloud Secure) helps
protect your data with actionable intelligence on insider threats. It
provides centralized visibility and control of all corporate data access
across hybrid cloud environments to ensure security and compliance goals
are met.
NOTE: Workload Security is not available in Cloud Insights Federal
Edition.
== Visibility
Gain centralized visibility and control of user access to your critical
corporate data stored on-premise or in the cloud.
Replace tools and manual processes that fail to provide timely and
accurate visibility into data access and control. Workload Security
uniquely operates on both cloud and on-premise storage systems to give you
```

real-time alerts of malicious user behavior.

== Protection

Protect organizational data from being misused by malicious or compromised users through advanced machine learning and anomaly detection.

Alerts you to any abnormal data access through advanced machine learning and anomaly detection of user behavior.

== Compliance

//not complete? 4/17

Ensure corporate compliance by auditing user data access to your critical corporate data stored on-premise or in the cloud.

```
= Getting Started
:leveloffset: +1
[[ID0658a047603dcd22f5526018db7050a0]]
= Getting Started with Workload Security
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
There are configuration tasks that need to be completed before you can
start using Workload Security to monitor user activity.
NOTE: Workload Security is not available in Cloud Insights Federal
Edition.
//not complete? 4/17
The Workload Security system uses an agent to collect access data from
storage systems and user information from Directory Services servers.
```

You need to configure the following before you can start collecting data:

```
[cols="2*"]
|===
|Task|Related information
| Configure an Agent a|
xref:{relative path}concept cs agent requirements.html[Agent Requirements]
xref:{relative path}task cs add agent.html[Add Agent]
link:https://netapp.hubs.vidyard.com/watch/Lce7EaGg7NZfvCUw4Jwy5P?[*Video*
: Agent Deployment]
|Configure a User Directory
Connector|link:task config_user_dir_connect.html[Add User Directory
Connectorl
link:https://netapp.hubs.vidyard.com/watch/NEmbmYrFjCHvPps7QMy8me?[*Video*
: Active Directory Connection]
|Configure data collectors | Click *Workload Security > Collectors*
Click the data collector you want to configure.
See the Data Collector Vendor Reference section of the documentation.
link:https://netapp.hubs.vidyard.com/watch/YSQrcYA7DKXbj1UGeLYnSF?[*Video*
: ONTAP SVM Connection]
|Create Users Accounts|
xref:{relative path}concept user roles.html[Manage User Accounts]
|Troubleshooting|link:https://netapp.hubs.vidyard.com/watch/Fs8N2w9wBtsFGr
hRH9X85U?[*Video*: Troubleshooting]
|===
Workload Security can integrate with other tools as well. For example,
link:http://docs.netapp.com/us-
en/cloudinsights/CloudInsights CloudSecure Splunk integration guide.pdf[se
e this guide] on integration with Splunk.
[[ID43c57ad6c72a800a8c122ccf2649a664]]
= Workload Security Agent Requirements
```

```
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
You must xref:{relative path}task cs add agent.html[install an Agent] in
order to acquire information from your data collectors. Before you install
the Agent, you should ensure that your environment meets operating system,
CPU, memory, and disk space requirements.
NOTE: Storage Workload Security is not available in Cloud Insights Federal
Edition.
[cols=2*,options="header",cols="36,60"]
|Component|Linux Requirement
|Operating system|A computer running a licensed version of one of the
following:
Red Hat Enterprise Linux 7.x, 8.x 64-bit, SELinux
CentOS 7.x 64-bit, SELinux
CentOS 8 Stream, SELinux
Ubuntu 20 through 22 64-bit
Rocky 8.x 64-bit, Rocky 9.x 64-bit, SELinux
SUSE Linux Enterprise Server 15 SP3, SUSE Linux Enterprise Server 15 SP4,
SELinux on SUSE 15 SP3
This computer should be running no other application-level software. A
dedicated server is recommended.
|Commands|'unzip' is required for installation. Additionally, the 'sudo su
-' command is required for installation, running scripts, and uninstall.
//|Docker | The Docker CE package must be installed on the VM hosting the
agent.
//The agent systems should always have the Docker CE package installed.
Users should not install the Docker-client-xx or Docker-common-xx native
RHEL Docker packages since these do not support the 'docker run' CLI
format that Workload Security supports.
//|Java |OpenJDK Java is required.
|CPU |4 CPU cores
|Memory | 16 GB RAM
```

```
|Available disk space | Disk space should be allocated in this manner:
//50 GB available for the root partition
/opt/netapp 35 GB (minimum)
Note: It is recommended to allocate a little extra disk space to allow for
the creation of the filesystem. Ensure that there is at least 35 GB free
space in the filesystem.
If /opt is a mounted folder from a NAS storage, make sure that local users
have access to this folder. Agent or Data collector may fail to install if
local users do not have permission to this folder. see the
xref:{relative path}task cs add agent.html#troubleshooting-agent-
errors[troubleshooting] section for more details.
|Network|100 Mbps to 1 Gbps Ethernet connection, static IP address, IP
connectivity to all devices, and a required port to the Workload Security
instance (80 or 443).
|===
Please note: The Workload Security agent can be installed in the same
machine as a Cloud Insights acquisition unit and/or agent. However, it is
a best practice to install these in separate machines. In the event that
these are installed on the same machine, please allocate disk space as
shown below:
|===
|Available disk space | 50-55 GB
For Linux, disk space should be allocated in this manner:
/opt/netapp 25-30 GB
/var/log/netapp 25 GB
|===
== Additional recommendations
* It is strongly recommended to synchronize the time on both the ONTAP
system and the Agent machine using *Network Time Protocol (NTP)* or
*Simple Network Time Protocol (SNTP)*.
////
Removed from Table:
|Agent outbound URLs (port 433)|
```

```
\https://<site name>.cs01.cloudinsights.netapp.com
//You can get the site ID from the product URL. For example:
https://*ab1234*.cs01.cloudinsights.netapp.com
You can use a broader range to specify the tenant ID:
\https://*.cs01.cloudinsights.netapp.com/
\https://gateway.c01.cloudinsights.netapp.com
\https://agentlogin.cs01.cloudinsights.netapp.com
////
== Cloud Network Access Rules
For *US-based* Workload Security environments:
[cols=5*, options="header"]
|===
|Protocol|Port|Source |Destination |Description
|TCP|443|Workload Security Agent|<site name>.cs01.cloudinsights.netapp.com
<site name>.c01.cloudinsights.netapp.com
<site name>.c02.cloudinsights.netapp.com|Access to Cloud Insights
|TCP|443|Workload Security Agent|gateway.c01.cloudinsights.netapp.com
agentlogin.cs01.cloudinsights.netapp.com|Access to authentication services
For *Europe-based* Workload Security environments:
[cols=5*, options="header"]
|===
|Protocol|Port|Source |Destination |Description
|TCP|443|Workload Security Agent|<site name>.cs01-eu-
1.cloudinsights.netapp.com
<site name>.c01-eu-1.cloudinsights.netapp.com
<site name>.c02-eu-1.cloudinsights.netapp.com
|Access to Cloud Insights
|TCP|443|Workload Security Agent|gateway.c01.cloudinsights.netapp.com
agentlogin.cs01-eu-1.cloudinsights.netapp.com
|Access to authentication services
|===
For *APAC-based* Workload Security environments:
[cols=5*, options="header"]
|===
```

```
|Protocol|Port|Source | Destination | Description
|TCP|443|Workload Security Agent|<site name>.cs01-ap-
1.cloudinsights.netapp.com
<site name>.c01-ap-1.cloudinsights.netapp.com
<site name>.c02-ap-1.cloudinsights.netapp.com
|Access to Cloud Insights
|TCP|443|Workload Security Agent|gateway.c01.cloudinsights.netapp.com
agentlogin.cs01-ap-1.cloudinsights.netapp.com
|Access to authentication services
|===
== In-network rules
//Note that when adding link:task add collector svm.html#permissions-
when-adding-via-cluster-management-ip[csuser] , that user requires SSH
access to the ONTAP management LIF.
[cols=5*, options="header"]
|===
|Protocol|Port|Source |Destination |Description
|TCP|389(LDAP)
636 (LDAPs / start-tls) |Workload Security Agent|LDAP Server URL|Connect
to LDAP
|TCP|443|Workload Security Agent|Cluster or SVM Management IP Address
(depending on SVM collector configuration) | API communication with ONTAP
|TCP|35000 - 55000|SVM data LIF IP Addresses|Workload Security
Agent|Communication from ONTAP to the Workload Security Agent for Fpolicy
events. These ports must be opened towards the Workload Security Agent in
order for ONTAP to send events to it, including any firewall on the
Workload Security Agent itself (if present).
|TCP|7|Workload Security Agent|SVM data LIF IP Addresses|Echo from Agent
to SVM Data LIFs
|SSH |22|Workload Security Agent| Cluster management |Needed for CIFS/SMB
user blocking.
|===
//link:task cs add agent.html[Configure an Agent]
// Supported browsers a
// * Internet Explorer 11
// * Firefox ESR 60
// * Chrome latest nightly (73.0)6
== System Sizing
```

```
See the xref:{relative path}concept cs event rate checker.html[Event Rate
Checker] documentation for information about sizing.
[[IDdb6b887d512c0b3ed6add7fd84ae290e]]
= Workload Security Agent Installation
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Workload Security (formerly Cloud Secure) collects user activity data
using one or more agents. Agents connect to devices in your environment
and collect data that is sent to the Workload Security SaaS layer for
analysis. See xref:{relative path}concept cs agent requirements.html[Agent
Requirements] to configure an agent VM.
NOTE: Workload Security is not available in Cloud Insights Federal
Edition.
== Before You Begin
* The sudo privilege is required for installation, running scripts, and
uninstall.
* While installing the agent, a local user cssys and a local group
cssys are created on the machine. If permission settings do not allow
creation of a local user, and instead require Active Directory, a user
with the username cssys must be created in the Active Directory server.
* You can read about Cloud Insights security
xref:{relative path}security overview.html[here].
== Steps to Install Agent
. Log in as Administrator or Account Owner to your Workload Security
environment.
. Select *Collectors > Agents > +Agent*
The system displays the Add an Agent page:
image::Add-agent-1.png[]
```

```
. Verify that the agent server meets the minimum system requirements.
. To verify that the agent server is running a supported version of Linux,
click Versions Supported (i) .
. If your network is using proxy server, please set the proxy server
details by following the instructions in the Proxy section.
image:CloudSecureAgentWithProxy Instructions.png[Agent Install with Proxy
Notel
//image::add-agent-2.png[]
. Click the Copy to Clipboard icon to copy the installation command.
. Run the installation command in a terminal window.
. The system displays the following message when the installation
completes successfully:
image::new-agent-detect.png[]
//cd /var NEW
//Grep /var/
.After You Finish
//. Verify that the agent is installed using the following command:
//`sudo grep -irn register agent.log`
. You need to configure a
xref:{relative path}task config user dir connect.html[User Directory
Collector ].
. You need to configure one or more Data Collectors.
////
== Files Created During Installation
* Installation directory:
/opt/netapp/cloudsecure/agent
* Installation logs:
/var/log/netapp/cloudsecure/install
/opt/netapp/cloud-secure/logs
* Agent Logs:
```

```
* You can use the following command to verify the agent installed
correctly:
`sudo grep -irn register /opt/netapp/cloudsecure/agent/logs/agent.log`
//* Use the following script to control the Workload Security service:
//`sudo cloud-secure-agent-service.sh --help`
* Use the following script to uninstall the agent:
`sudo cloudsecure-agent-uninstall.sh`
////
== Network Configuration
Run the following commands on the local system to open ports that will be
used by Workload Security. If there is a security concern regarding the
port range, you can use a lesser port range, for example 35000:35100 .
Each SVM uses two ports.
.Steps
. `sudo firewall-cmd --permanent --zone=public --add-port=35000-55000/tcp`
. `sudo firewall-cmd --reload`
Follow the next steps according to your platform:
*CentOS 7.x / RHEL 7.x*:
. `sudo iptables-save | grep 35000`
Sample output:
-A IN public allow -p tcp -m tcp --dport 35000:55000 -m conntrack
-ctstate NEW, UNTRACKED -j ACCEPT
*CentOS 8.x / RHEL 8.x*:
. `sudo firewall-cmd --zone=public --list-ports | grep 35000` (for CentOS
8)
Sample output:
35000-55000/tcp
== Troubleshooting Agent Errors
```

Known problems and their resolutions are described in the following table.

[cols=2*, options="header", cols"30,70"]

|===

|Problem: | Resolution:

|Agent installation fails to create the

/opt/netapp/cloudsecure/agent/logs/agent.log folder and the install.log file provides no relevant information.|This error occurs during bootstrapping of the agent. The error is not logged in log files because it occurs before logger is initialized.

The error is redirected to standard output, and is visible in the service log using the `journalctl -u cloudsecure-agent.service` command. This command can be used for troubleshooting the issue further.

|Agent installation fails with 'This linux distribution is not supported. Exiting the installation'.|This error appears when you attempt to install the Agent on an unsupported system. See

xref:{relative_path}concept_cs_agent_requirements.html[Agent
Requirements].

|Agent Installation failed with the error:

"-bash: unzip: command not found"

|Install unzip and then run the installation command again. If Yum is installed on the machine, try "yum install unzip" to install unzip software.

After that, re-copy the command from the Agent installation UI and paste it in the CLI to execute the installation again.

|Agent was installed and was running. However agent has stopped suddenly. |SSH to the Agent machine. Check the status of the agent service via `sudo systemctl status cloudsecure-agent.service`.

- 1. Check if the logs shows a message "Failed to start Workload Security daemon service".
- 2. Check if cssys user exists in the Agent machine or not. Execute the following commands one by one with root permission and check if the cssys user and group exists.

`sudo id cssys`

`sudo groups cssys`

- 3. If none exists, then a centralized monitoring policy may have deleted the cssys user.
- 4. Create cssys user and group manually by executing the following commands.
- `sudo useradd cssys`
- `sudo groupadd cssys`

- 5. Restart the agent service after that by executing the following command:
- `sudo systemctl restart cloudsecure-agent.service`
- 6. If it is still not running, please check the other troubleshooting options.

|Unable to add more than 50 Data collectors to an Agent.

|Only 50 Data collectors can be added to an Agent. This can be a combination of all the collector types, for example, Active Directory, SVM and other collectors.

|UI shows Agent is in NOT CONNECTED state.

|Steps to restart the Agent.

- 1. SSH to the Agent machine.
- 2. Restart the agent service after that by executing the following command:
- `sudo systemctl restart cloudsecure-agent.service`
- 3. Check the status of the agent service via `sudo systemctl status cloudsecure-agent.service`.
- 4. Agent should go to CONNECTED state.

|Agent VM is behind Zscaler proxy and the agent installation is failing. Because of Zscaler proxy's SSL inspection, the Workload Security certificates are presented as it is signed by Zscaler CA so the agent is not trusting the communication.

|Disable SSL inspection in the Zscaler proxy for the

*.cloudinsights.netapp.com url. If Zscaler does SSL inspection and replaces the certificates, Workload Security will not work.

|While installing the agent, the installation hangs after unzipping. | "chmod 755 -Rf" command is failing.

The command fails when the agent installation command is being run by a non-root sudo user that has files in the working directory, belonging to another user, and permissions of those files cannot be changed. Because of the failing chmod command, the rest of the installation does not execute.

- 1. Create a new directory named "cloudsecure".
- 2. Go to that directory.
- 3. Copy and paste the full "token=...../cloudsecure-agent-install.sh" installation command and press enter.
- 4. Installation should be able to proceed.

|If the Agent is still not able to connect to Saas, please open a case with NetApp Support. Provide the Cloud Insights serial number to open a case, and attach logs to the case as noted.

|To attach logs to the case:

- 1. Execute the following script with root permission and share the output file (cloudsecure-agent-symptoms.zip).
- a. /opt/netapp/cloudsecure/agent/bin/cloudsecure-agent-symptom-collector.sh
- 2. Execute the following commands one by one with root permission and share the output.
 - a. id cssys
 - b. groups cssys
 - c. cat /etc/os-release

|The cloudsecure-agent-symptom-collector.sh script fails with the following error.

 $[root@machine \ tmp] \# \ /opt/netapp/cloudsecure/agent/bin/cloudsecure-agent-symptom-collector.sh$

Collecting service log

Collecting application logs

Collecting agent configurations

Taking service status snapshot

Taking agent directory structure snapshot

.....

 $/ {\tt opt/netapp/cloudsecure/agent/bin/cloudsecure-agent-symptom-collector.sh:} \\$

line 52: zip: command not found

ERROR: Failed to create /tmp/cloudsecure-agent-symptoms.zip

|Zip tool is not installed..

Install the zip tool by running the command "yum install zip".

Then run the cloudsecure-agent-symptom-collector.sh again.

|Agent installation Fails with useradd: cannot create directory /home/cssys

|This error can occur if user's login directory cannot be created under /home, due to lack of permissions.

The workaround would be to create cssys user and add its login directory manually using the following command:

_sudo useradd user_name -m -d HOME_DIR_

-m : Create the user's home directory if it does not exist.

-d : The new user is created using HOME_DIR as the value for the user's login directory.

For instance, _sudo useradd cssys -m -d /cssys_, adds a user _cssys_ and creates its login directory under root.

|Agent is not running after installation. Systemctl status cloudsecure-agent.service shows the following: [root@demo ~] # systemctl status cloudsecure-agent.service agent.service - Workload Security Agent Daemon Service Loaded: loaded (/usr/lib/systemd/system/cloudsecure-agent.service; enabled; vendor preset: disabled) Active: activating (auto-restart) (Result: exit-code) since Tue 2021-08-03 21:12:26 PDT; 2s ago Process: 25889 ExecStart=/bin/bash /opt/netapp/cloudsecure/agent/bin/cloudsecure-agent (code=exited status=126) Main PID: 25889 (code=exited, status=126), Aug 03 21:12:26 demo systemd[1]: cloudsecure-agent.service: main process exited, code=exited, status=126/n/a Aug 03 21:12:26 demo systemd[1]: Unit cloudsecure-agent.service entered failed state. Aug 03 21:12:26 demo systemd[1]: cloudsecure-agent.service failed. |This can be failing because cssys user may not have permission to install. If /opt/netapp is an NFS mount and if cssys user does not have access to this folder, installation will fail. cssys is a local user created by the Workload Security installer that may not have permission to access the mounted share. You can check this by attempting to access /opt/netapp/cloudsecure/agent/bin/cloudsecure-agent using cssys user. If it returns "Permission denied", installation permission is not present. Instead of a mounted folder, install on a directory local to the machine. |Agent was initially connected via a proxy server and the proxy was set

|Agent was initially connected via a proxy server and the proxy was set during Agent installation. Now the proxy server has changed. How can the Agent's proxy configuration be changed?

|You can edit the agent.properties to add the proxy details. Follow these steps:

- 1. Change to the folder containing the properties file:
- cd /opt/netapp/cloudsecure/conf
- 2. Using your favorite text editor, open the _agent.properties_ file for editing.

```
3. Add or modify the following lines:
AGENT PROXY HOST=scspa1950329001.vm.netapp.com
AGENT PROXY PORT=80
AGENT PROXY USER=pxuser
AGENT PROXY PASSWORD=pass1234
4. Save the file.
5. Restart the agent:
sudo systemctl restart cloudsecure-agent.service
|===
//image:CloudSecureAgentInstallationCommand.png[]
[[ID5f837baf01267791533025dfc8510af7]]
= Deleting a Workload Security Agent
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
When you delete a Workload Security Agent, all the data collectors
associated with the Agent must be deleted first.
== Deleting an Agent
[IMPORTANT]
Deleting an Agent deletes all of the Data Collectors associated with the
Agent. If you plan to configure the data collectors with a different agent
you should create a backup of the Data Collector configurations before you
```

```
delete the Agent.
.Before you begin
. Make sure all the data collectors associated with the agent are deleted
from the Workload Security portal.
Note: Ignore this step if all the associated collectors are in STOPPED
state.
.Steps to delete an Agent:
. SSH into the agent VM and execute the following command. When prompted,
enter "y" to continue.
 sudo /opt/netapp/cloudsecure/agent/install/cloudsecure-agent-uninstall.sh
Uninstall CloudSecure Agent? [y|N]:
. Click *Workload Security > Collectors > Agents*
The system displays the list of configured Agents.
. Click the options menu for the Agent you are deleting.
. Click *Delete*.
The system displays the *Delete Agent* page.
. Click *Delete* to confirm the deletion.
[[IDb5c235ff634549f83549d1c9e930c0d5]]
= Configuring an Active Directory (AD) User Directory Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Workload Security can be configured to collect user attributes from Active
Directory servers.
```

```
.Before you begin
* You must be a Cloud Insights Administrator or Account Owner to perform
* You must have the IP address of the server hosting the Active Directory
server.
* An Agent must be configured before you configure a User Directory
connector.
.Steps to Configure a User Directory Collector
. In the Workload Security menu, click:
*Collectors > User Directory Collectors > + User Directory Collector* and
select *Active Directory*
The system displays the Add User Directory screen.
Configure the User Directory Collector by entering the required data in
the following tables:
[cols=2*, cols"30,70"]
[Options=header]
|===
|Name|Description
|Name |Unique name for the user directory. For example GlobalADCollector
|Agent|Select a configured agent from the list
|Server IP/Domain Name|IP address or Fully-Qualified Domain Name (FQDN) of
server hosting the active directory
|Forest Name|Forest level of the directory structure.
Forest name allows both of the following formats:
x.y.z => direct domain name as you have it on your SVM. [Example:
hq.companyname.com]
DC=x,DC=y,DC=z => Relative distinguished names [Example: DC=hq,DC=
companyname, DC=com]
Or you can specify as the following:
OU=engineering, DC=hq, DC= companyname, DC=com [to filter by specific OU
engineering]
CN=username, OU=engineering, DC=companyname, DC=netapp, DC=com [to get
only specific user with <username> from OU <engineering>]
CN=Acrobat Users, CN=Users, DC=hq, DC=companyname, DC=com, O=
```

```
companyname, L=Boston, S=MA, C=US [to get all Acrobat Users within the Users
in that organization]
Trusted Active Directory domains are also supported.
|Bind DN|User permitted to search the directory. For example:
username@companyname.com or username@domainname.com
In addition, Domain Read Only permission is required.
User must be a member of the Security group Read-only Domain
Controllers .
|BIND password|Directory server password (i.e. password for username used
in Bind DN)
|Protocol|Idap, Idaps, Idap-start-tls
|Ports|Select port
|===
////
Add to table once link is provided:
For more details about forest names, please refer to this
xref:{relative path}///
Enter the following Directory Server required attributes if the default
attribute names have been modified in Active Directory. Most often these
attributes names are not modified in Active Directory, in which case you
can simply proceed with the default attribute name.
[cols=2*, cols"50,50"]
[Options=header]
|Attributes |Attribute name in Directory Server
|Display Name|name
|SID|objectsid
|User Name|sAMAccountName
|===
Click Include Optional Attributes to add any of the following attributes:
[cols=2*, cols"50,50"]
[Options=header]
|Attributes |Attribute Name in Directory Server
|Email Address|mail
|Telephone Number|telephonenumber
|Role|title
|Country|co
|State|state
```

```
|Department|department
|Photo|thumbnailphoto
|ManagerDN|manager
|Groups|memberOf
|===
//Removed based on review comments
//Enter the following user search parameters in the Advanced Configuration
attributes table:
//[cols=2*, cols"50,50"]
//[Options=header]
//|===
//|*Base DN*|*Query*
//|Attributes //|(&(objectCategory=person)(objectClass=user))
//|Email Address|mail
//|Phone|telephoneNumber
//|Country|Country
//|State|state
//|Department|department
//|Photo|thumbnailPhoto
//
//|===
== Testing Your User Directory Collector Configuration
You can validate LDAP User Permissions and Attribute Definitions using the
following procedures:
* Use the following command to validate Workload Security LDAP user
permission:
`ldapsearch -o ldif-wrap=no -LLL -x -b "dc=netapp,dc=com" -h 10.235.40.29
-p 389 -D \Administrator@netapp.com -W`
* Use AD Explorer to navigate an AD database, view object properties and
attributes, view permissions, view an object's schema, execute
sophisticated searches that you can save and re-execute.
** Install link:https://docs.microsoft.com/en-
us/sysinternals/downloads/adexplorer[AD Explorer] on any windows machine
which can connect to the AD Server.
** Connect to the AD server using the username/password of the AD
directory server.
```

image:cs ADExample.png[AD Connection]

== Troubleshooting User Directory Collector Configuration Errors

The following table describes known problems and resolutions that can occur during collector configuration:

[cols=2*, cols"50,50"]

[options="header"]

|===

|Problem: | Resolution:

|Adding a User Directory connector results in the 'Error' state. Error says, "Invalid credentials provided for LDAP server".

|Incorrect username or password provided. Edit and provide the correct user name and password.

|Adding a User Directory connector results in the 'Error' state. Error says, "Failed to get the object corresponding to DN=DC=hq,DC=domainname,DC=com provided as forest name."

|Incorrect forest name provided. Edit and provide the correct forest name.

|The optional attributes of domain user are not appearing in the Workload Security User Profile page.

|This is likely due to a mismatch between the names of optional attributes added in CloudSecure and the actual attribute names in Active Directory. Edit and provide the correct optional attribute name(s).

|Data collector in error state with "Failed to retrieve LDAP users. Reason for failure: Cannot connect on the server, the connection is null" |Restart the collector by clicking on the Restart button.

|Adding a User Directory connector results in the 'Error' state. |Ensure you have provided valid values for the required fields (Server, forest-name, bind-DN, bind-Password).

Ensure bind-DN input is always provided as

'Administrator@<domain_forest_name>' or as a user account with domain admin privileges.

|Adding a User Directory connector results in the 'RETRYING' state. Shows error "Unable to define state of the collector, reason Tcp command [Connect(localhost:35012, None, List(), Some(, seconds), true)] failed because of java.net.ConnectionException:Connection refused."

|Incorrect IP or FQDN provided for the AD Server. Edit and provide the correct IP address or FQDN.

|Adding a User Directory connector results in the 'Error' state. Error

says, "Failed to establish LDAP connection".

|Incorrect IP or FQDN provided for the AD Server. Edit and provide the correct IP address or FQDN.

|Adding a User Directory connector results in the 'Error' state. Error says, "Failed to load the settings. Reason: Datasource configuration has an error. Specific reason: /connector/conf/application.conf: 70: ldap.ldap-port has type STRING rather than NUMBER" |Incorrect value for Port provided. Try using the default port values or the correct port number for the AD server.

|I started with the mandatory attributes, and it worked. After adding the optional ones, the optional attributes data is not getting fetched from AD.

|This is likely due to a mismatch between the optional attributes added in CloudSecure and the actual attribute names in Active Directory. Edit and provide the correct mandatory or optional attribute name.

|After restarting the collector, when will the AD sync happen? |AD sync will happen immediately after the collector restarts. It will take approximately 15 minutes to fetch user data of approximately 300K users, and is refreshed every 12 hours automatically.

|User Data is synced from AD to CloudSecure. When will the data be deleted?

|User data is retained for 13months in case of no refresh. If the tenant is deleted then the data will be deleted.

|User Directory connector results in the 'Error' state. "Connector is in error state. Service name: usersLdap. Reason for failure: Failed to retrieve LDAP users. Reason for failure: 80090308: LdapErr: DSID-0C090453, comment: AcceptSecurityContext error, data 52e, v3839" |Incorrect forest name provided. See above on how to provide the correct forest name.

|Telephone number is not getting populated in the user profile page. |This is most likely due to an attribute mapping problem with the Active Directory.

- 1. Edit the particular Active Directory collector which is fetching the user's information from Active Directory.
- 2. Notice under optional attributes, there is a field name "Telephone Number" mapped to Active Directory attribute 'telephonenumber'.
- 4. Now, please use the Active Directory Explorer tool as described above to browse the Active Directory and see the correct attribute name.
- 3. Make sure that in Active Directory there is an attribute named

'telephonenumber' which has indeed the telephone number of the user.

- 5. Let us say in Active Directory it has been modified to 'phonenumber'.
- 6. Then Edit the CloudSecure User Directory collector. In optional attribute section, replace 'telephonenumber' with 'phonenumber'.
- 7. Save the Active Directory collector, the collector will restart and get the telephone number of the user and display the same in the user profile page.

|If encryption certificate (SSL) is enabled on the Active Directory (AD) Server, the Workload Security User Directory Collector can not connect to the AD Server.

|Disable AD Server encryption before Configuring a User Directory Collector.

Once the user detail is fetched it will be there for 13 months. If the AD server gets disconnected after fetching the user details, the newly added users in AD won't get fetched. To fetch again, the user directory collector needs to be connected to AD.

|Data from Active Directory is present in CloudInsights Security.

Want to delete all the user information from CloudInsights.

|It is not possible to ONLY delete Active Directory user information from CloudInsights Security. In order to delete the user, the complete tenant needs to be deleted.

|===

Directory servers.

.Before you begin

```
[[ID263a58ff7727aa1194c02e437f331456]]
= Configuring an LDAP Directory Server Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
You configure Workload Security to collect user attributes from LDAP
```

* You must be a Cloud Insights Administrator or Account Owner to perform this task.

```
* You must have the IP address of the server hosting the LDAP Directory
server.
* An Agent must be configured before you configure an LDAP Directory
connector.
.Steps to Configure a User Directory Collector
. In the Workload Security menu, click:
*Collectors > User Directory Collectors > + User Directory Collector* and
select *LDAP Directory Server*
The system displays the Add User Directory screen.
Configure the User Directory Collector by entering the required data in
the following tables:
[cols=2*, cols"30,70"]
[Options=header]
|===
|Name|Description
|Name |Unique name for the user directory. For example
GlobalLDAPCollector
|Agent|Select a configured agent from the list
|Server IP/Domain Name|IP address or Fully-Qualified Domain Name (FQDN) of
server hosting the LDAP Directory Server
|Search Base|Search Base of the LDAP server
Search Base allows both of the following formats:
x.y.z => direct domain name as you have it on your SVM. [Example:
hq.companyname.com]
DC=x,DC=y,DC=z => Relative distinguished names [Example: DC=hq,DC=
companyname, DC=com]
Or you can specify as the following:
OU=engineering, DC=hq, DC= companyname, DC=com [to filter by specific OU
engineering]
CN=username, OU=engineering, DC=companyname, DC=netapp, DC=com [to get
only specific user with <username> from OU <engineering>]
CN=Acrobat Users, CN=Users, DC=hq, DC=companyname, DC=com, O=
companyname, L=Boston, S=MA, C=US [to get all Acrobat Users within the Users
in that organization]
```

```
|Bind DN|User permitted to search the directory. For example:
uid=ldapuser, cn=users, cn=accounts, dc=domain, dc=companyname, dc=com
uid=john,cn=users,cn=accounts,dc=dorp,dc=company,dc=com for a user
john@dorp.company.com.
dorp.company.com
|--accounts
        |--users
              |--john
              |--anna
|BIND password|Directory server password (i.e. password for username used
in Bind DN)
|Protocol|Idap, ldaps, ldap-start-tls
|Ports|Select port
|===
////
Add to table once link is provided:
For more details about forest names, please refer to this
xref:{relative path}///
Enter the following Directory Server required attributes if the default
attribute names have been modified in LDAP Directory Server. Most often
these attributes names are not modified in LDAP Directory Server, in
which case you can simply proceed with the default attribute name.
[cols=2*, cols"50,50"]
[Options=header]
|===
|Attributes |Attribute name in Directory Server
|Display Name|name
|UNIXID|uidnumber
|User Name|uid
|===
Click Include Optional Attributes to add any of the following attributes:
[cols=2*, cols"50,50"]
[Options=header]
|===
|Attributes |Attribute Name in Directory Server
|Email Address|mail
|Telephone Number|telephonenumber
|Role|title
```

```
|Country|co
|State|state
|Department|departmentnumber
|Photo|photo
|ManagerDN|manager
|Groups|memberOf
|===
//Removed based on review comments
//Enter the following user search parameters in the Advanced Configuration
attributes table:
//[cols=2*, cols"50,50"]
//[Options=header]
//|===
//|*Base DN*|*Query*
//|Attributes //|(&(objectCategory=person)(objectClass=user))
//|Email Address|mail
//|Phone|telephoneNumber
//|Country|Country
//|State|state
//|Department|department
//|Photo|thumbnailPhoto
//
//|===
== Testing Your User Directory Collector Configuration
You can validate LDAP User Permissions and Attribute Definitions using the
following procedures:
* Use the following command to validate Workload Security LDAP user
permission:
ldapsearch -D "uid=john ,cn=users,cn=accounts,dc=dorp,dc=company,dc=com"
-W -x -LLL -o ldif-wrap=no -b "cn=accounts,dc=dorp,dc=company,dc=com" -H
ldap://vmwipaapp08.dorp.company.com
* Use LDAP Explorer to navigate an LDAP database, view object properties
and attributes, view permissions, view an object's schema, execute
sophisticated searches that you can save and re-execute.
** Install LDAP Explorer (http://ldaptool.sourceforge.net/) or Java LDAP
Explorer (http://jxplorer.org/) on any windows machine which can connect
to the LDAP Server.
```

** Connect to the LDAP server using the username/password of the LDAP directory server.

image:CloudSecure LDAPDialog.png[LDAP Connection]

== Troubleshooting LDAP Directory Collector Configuration Errors

The following table describes known problems and resolutions that can occur during collector configuration:

[cols=2*, cols"50,50"]

[options="header"]

|===

|Problem: | Resolution:

|Adding an LDAP Directory connector results in the 'Error' state. Error says, "Invalid credentials provided for LDAP server".

|Incorrect Bind DN or Bind Password or Search Base provided. Edit and provide the correct information.

|Adding an LDAP Directory connector results in the 'Error' state. Error says, "Failed to get the object corresponding to

DN=DC=hq, DC=domainname, DC=com provided as forest name."

|Incorrect Search Base provided. Edit and provide the correct forest name.

|The optional attributes of domain user are not appearing in the Workload Security User Profile page.

|This is likely due to a mismatch between the names of optional attributes added in CloudSecure and the actual attribute names in Active Directory. Fields are case sensitive. Edit and provide the correct optional attribute name(s).

|Data collector in error state with "Failed to retrieve LDAP users. Reason for failure: Cannot connect on the server, the connection is null" |Restart the collector by clicking on the _Restart_ button.

|Adding an LDAP Directory connector results in the 'Error' state. |Ensure you have provided valid values for the required fields (Server, forest-name, bind-DN, bind-Password).

Ensure bind-DN input is always provided as uid=ldapuser, cn=users, cn=accounts, dc=domain, dc=companyname, dc=com.

|Adding an LDAP Directory connector results in the 'RETRYING' state. Shows error "Failed to determine the health of the collector hence retrying again"

|Ensure correct Server IP and Search Base is provided

////

|While adding LDAP directory the following error is shown:
"Failed to determine the health of the collector within 2 retries, try
restarting the collector again(Error Code: AGENT008)"
|Ensure correct Server IP and Search Base is provided

|Adding an LDAP Directory connector results in the 'RETRYING' state. Shows error "Unable to define state of the collector, reason Tcp command [Connect(localhost:35012, None, List(), Some(, seconds), true)] failed because of java.net.ConnectionException:Connection refused."
|Incorrect IP or FQDN provided for the AD Server. Edit and provide the correct IP address or FQDN.

, , , ,

|Adding an LDAP Directory connector results in the 'Error' state. Error says, "Failed to establish LDAP connection".

|Incorrect IP or FQDN provided for the LDAP Server. Edit and provide the correct IP address or FQDN.

Or

Incorrect value for Port provided. Try using the default port values or the correct port number for the LDAP server.

|Adding an LDAP Directory connector results in the 'Error' state. Error says, "Failed to load the settings. Reason: Datasource configuration has an error. Specific reason: /connector/conf/application.conf: 70: ldap.ldap-port has type STRING rather than NUMBER" |Incorrect value for Port provided. Try using the default port values or the correct port number for the AD server.

 \mid I started with the mandatory attributes, and it worked. After adding the optional ones, the optional attributes data is not getting fetched from AD.

|This is likely due to a mismatch between the optional attributes added in CloudSecure and the actual attribute names in Active Directory. Edit and provide the correct mandatory or optional attribute name.

|After restarting the collector, when will the LDAP sync happen? |LDAP sync will happen immediately after the collector restarts. It will take approximately 15 minutes to fetch user data of approximately 300K users, and is refreshed every 12 hours automatically.

|User Data is synced from LDAP to CloudSecure. When will the data be deleted?

|User data is retained for 13months in case of no refresh. If the tenant is deleted then the data will be deleted.

|LDAP Directory connector results in the 'Error' state. "Connector is in error state. Service name: usersLdap. Reason for failure: Failed to retrieve LDAP users. Reason for failure: 80090308: LdapErr: DSID-0C090453, comment: AcceptSecurityContext error, data 52e, v3839" |Incorrect forest name provided. See above on how to provide the correct forest name.

|Telephone number is not getting populated in the user profile page. |This is most likely due to an attribute mapping problem with the Active Directory.

- 1. Edit the particular Active Directory collector which is fetching the user's information from Active Directory.
- 2. Notice under optional attributes, there is a field name "Telephone Number" mapped to Active Directory attribute 'telephonenumber'.
- 4. Now, please use the Active Directory Explorer tool as described above to browse the LDAP Directory server and see the correct attribute name.
- 3. Make sure that in LDAP Directory there is an attribute named 'telephonenumber' which has indeed the telephone number of the user.
- 5. Let us say in LDAP Directory it has been modified to 'phonenumber'.
- 6. Then Edit the CloudSecure User Directory collector. In optional attribute section, replace 'telephonenumber' with 'phonenumber'.
- 7. Save the Active Directory collector, the collector will restart and get the telephone number of the user and display the same in the user profile page.

|If encryption certificate (SSL) is enabled on the Active Directory (AD) Server, the Workload Security User Directory Collector can not connect to the AD Server.

|Disable AD Server encryption before Configuring a User Directory Collector.

Once the user detail is fetched it will be there for 13 months. If the AD server gets disconnected after fetching the user details, the newly added users in AD won't get fetched. To fetch again the user directory collector needs to be connected to AD.

|===

[[ID8aff8313ebb28727869d3d7d9122ac36]]

= Configuring the ONTAP SVM Data Collector

:toc: macro
:hardbreaks:
:toclevels: 1

```
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Workload Security uses data collectors to collect file and user access
data from devices.
== Before you begin
* This data collector is supported with the following:
** Data ONTAP 9.2 and later versions. For best performance, use a Data
ONTAP version greater than 9.13.1.
** SMB protocol version 3.1 and earlier.
//Note that Workload Security does not work with SMB configurations that
use Flexcache. In systems using Flexcache, starting with ONTAP 9.7,
Fpolicy is supported only in an NFS environment.
** NFS protocol version 4.0 and earlier
** Flexgroup is supported from ONTAP 9.4 and later versions
** ONTAP Select is supported
* Only data type SVMs are supported. SVMs with infinite volumes are not
supported.
* SVM has several sub-types. Of these, only _default_, _sync_source_, and
sync destination are supported.
* An Agent xref:{relative path}task cs add agent.html[must be configured]
before you can configure data collectors.
* Make sure that you have a properly configured User Directory Connector,
otherwise events will show encoded user names and not the actual name of
the user (as stored in Active Directory) in the "Activity Forensics" page.
* For optimal performance, you should configure the FPolicy server to be
on the same subnet as the storage system.
//* You need the SVM management IP address or the cluster IP, and username
/ password for login.
* You must add an SVM using one of the following two methods:
** By Using Cluster IP, SVM name, and Cluster Management Username and
Password. * This is the recommended method. *
*** SVM name must be exactly as is shown in ONTAP and is case-sensitive.
** By Using SVM Vserver Management IP, Username, and Password
** If you are not able or not willing to use the full Administrator
```

```
Cluster/SVM Management Username and Password, you can create a custom user
with lesser privileges as mentioned in the xref:{relative path}#a-note-
about-permissions["A note about permissions"] section below. This custom
user can be created for either SVM or Cluster access.
*** o You can also use an AD user with a role that has at least the
permissions of csrole as mentioned in "A note about permissions" section
below. Also refer to the link:https://docs.netapp.com/ontap-
9/index.jsp?topic=%2Fcom.netapp.doc.pow-adm-auth-rbac%2FGUID-0DB65B04-
71DB-43F4-9A0F-850C93C4896C.html[ONTAP documentation].
* Ensure the correct applications are set for the SVM by executing the
following command:
clustershell::> security login show -vserver <vservername> -user-or-group
-name <username>
Example output:
 image:cs svm sample output.png[SVM Command Output Example]
////
security login show -vserver symname
    Vserver: svmname
    Authentication Acct Is-Nsswitch
    User/Group Name Application Method Role Name Locked Group
    vsadmin http password vsadmin yes no
    vsadmin ontapi password vsadmin yes no
    vsadmin ssh password vsadmin yes no
    3 entries were displayed.
////
* Ensure that the SVM has a CIFS server configured:
clustershell::> `vserver cifs show`
The system returns the Vserver name, CIFS server name and additional
fields.
* Set a password for the SVM vsadmin user. If using custom user or cluster
admin user, skip this step.
 clustershell::> `security login password -username vsadmin -vserver
svmname`
* Unlock the SVM vsadmin user for external access. If using custom user or
cluster admin user, skip this step.
clustershell::> `security login unlock -username vsadmin -vserver
svmname`
```

* Ensure the firewall-policy of the data LIF is set to 'mgmt' (not 'data'). Skip this step if using a dedicated management lif to add the SVM.

clustershell::> `network interface modify -lif <SVM_data_LIF_name>
-firewall-policy mgmt`

 * When a firewall is enabled, you must have an exception defined to allow TCP traffic for the port using the Data ONTAP Data Collector.

See xref:{relative_path}concept_cs_agent_requirements.html[Agent requirements] for configuration information. This applies to on-premise Agents and Agents installed in the Cloud.

* When an Agent is installed in an AWS EC2 instance to monitor a Cloud ONTAP SVM, the Agent and Storage must be in the same VPC. If they are in separate VPCs, there must be a valid route between the VPC's.

== Prerequisites for User Access Blocking

Keep the following in mind for
xref:{relative_path}/cloudinsights/cs_restrict_user_access.html[User
Access Blocking]:

Cluster level credentials are needed for this feature to work.

If you are using cluster administration credentials, no new permissions are needed.

If you are using a custom user (for example, _csuser_) with permissions given to the user, then follow the steps below to give permissions to Workload Security to block user.

For csuser with cluster credentials, do the following from the ONTAP command line:

security login role create -role csrole -cmddirname "vserver export-policy rule" -access all

security login role create -role csrole -cmddirname set -access all security login role create -role csrole -cmddirname "vserver cifs session" -access all

security login role create -role csrole -cmddirname "vserver services access-check authentication translate" -access all

security login role create -role csrole -cmddirname "vserver name-mapping" -access all

== A Note About Permissions === Permissions when adding via *Cluster Management IP*: If you cannot use the Cluster management administrator user to allow Workload Security to access the ONTAP SVM data collector, you can create a new user named "csuser" with the roles as shown in the commands below. Use the username "csuser" and password for "csuser" when configuring the Workload Security data collector to use Cluster Management IP. To create the new user, log in to ONTAP with the Cluster management Administrator username/password, and execute the following commands on the ONTAP server: security login role create -role csrole -cmddirname DEFAULT -access readonly security login role create -role csrole -cmddirname "vserver fpolicy" -access all security login role create -role csrole -cmddirname "volume snapshot" -access all -query "-snapshot cloudsecure *" security login role create -role csrole -cmddirname "event catalog" -access all security login role create -role csrole -cmddirname "event filter" -access all security login role create -role csrole -cmddirname "event notification destination" -access all security login role create -role csrole -cmddirname "event notification" -access all security login role create -role csrole -cmddirname "security certificate" -access all security login create -user-or-group-name csuser -application ontapi -authmethod password -role csrole security login create -user-or-group-name csuser -application ssh -authmethod password -role csrole Permissions for ONTAP ARP Integration: security login rest-role create -role arwrole -api /api/storage/volumes -access readonly -vserver <cluster name> security login rest-role create -api /api/security/anti-ransomware -access readonly -role arwrole -vserver <cluster_name> security login create -user-or-group-name csuser -application http -authmethod password -role arwrole

Permissions for ONTAP Access Denied:

security login rest-role create -role csrestrole -api
/api/protocols/fpolicy -access all -vserver <cluster_name>
security login create -user-or-group-name csuser -application http
-authmethod password -role csrestrole

Note: If one rest-role is already added--either _arwrole_ or _csrestrole_--there is no need to add a second rest-role. You can simply add the API permissions as in the example below.

Example: _csrestrole_ is already present, so we only have to enable antiransomware protection and give API permissions to the existing csrestrole:

security login rest-role create -role csrestrole -api
/api/storage/volumes -access readonly -vserver <cluster_name>
security login rest-role create -api /api/security/anti-ransomware
-access readonly -role arwrole -vserver <cluster name>

=== Permissions when adding via *Vserver Management IP*:

If you cannot use the Cluster management administrator user to allow Workload Security to access the ONTAP SVM data collector, you can create a new user named "csuser" with the roles as shown in the commands below. Use the username "csuser" and password for "csuser" when configuring the Workload Security data collector to use Vserver Management IP.

//If you cannot use the "vsadmin" user, since "vsadmin" has all the privileges, create a new user named "csuser" with the following roles as is shown in the command below. Use the username "csuser" and password for "csuser" for adding the Vserver via Vserver Mgmt IP in the ONTAP DataSource Addition UI.

To create the new user, log in to ONTAP with the Cluster management Administrator username/password, and execute the following commands on the ONTAP server. For ease, copy these commands to a text editor and replace the <vservername> with your Vserver name before and executing these commands on ONTAP:

security login role create -vserver <vservername> -role csrole
-cmddirname DEFAULT -access none

security login role create -vserver <vservername> -role csrole
-cmddirname "network interface" -access readonly

security login role create -vserver <vservername> -role csrole -cmddirname version -access readonly security login role create -vserver <vservername> -role csrole -cmddirname volume -access readonly security login role create -vserver <vservername> -role csrole -cmddirname vserver -access readonly security login role create -vserver <vservername> -role csrole -cmddirname "vserver fpolicy" -access all security login role create -vserver <vservername> -role csrole -cmddirname "volume snapshot" -access all security login create -user-or-group-name csuser -application ontapi -authmethod password -role csrole -vserver <vservername> Permissions for ONTAP Access Denied: security login rest-role create -role csrestrole -api /api/protocols/fpolicy -access all -vserver <svm name> security login create -user-or-group-name csuser -application http -authmethod password -role csrestrole -vserver <svm name> === Permissions for ONTAP Autonomous Ransomware Protection If you are using cluster administration credentials, no new permissions are needed. If you are using a custom user (for example, csuser) with permissions given to the user, then follow the steps below to give permissions to Workload Security to collect ARP related information from ONTAP. For csuser with cluster credentials, do the following from the ONTAP command line: security login rest-role create -role arwrole -api /api/storage/volumes -access readonly -vserver <cluster name> security login rest-role create -api /api/security/anti-ransomware -access readonly -role arwrole -vserver <cluster name> security login create -user-or-group-name csuser -application http -authmethod password -role arwrole For more information, read about xref:{relative path}concept cs integration with ontap arp.html[Integration with ONTAP Autonomous Ransomware Protection]

=== Permissions for ONTAP Access Denied

If the Data Collector is added using cluster administration credentials, no new permissions are needed.

If the Collector is added using a custom user (for example, _csuser_) with permissions given to the user, follow the steps below to give Workload Security the necessary permission to register for Access Denied events with ONTAP.

For csuser with _cluster_ credentials, execute the following commands from the ONTAP command line. Note that _csrestrole_ is custom role and _csuser_ is ontap custom user.

security login rest-role create -role csrestrole -api
/api/protocols/fpolicy -access all -vserver <cluster_name>
security login create -user-or-group-name csuser -application http
-authmethod password -role csrestrole

For csuser with _SVM_ credentials, execute the following commands from the ONTAP command line:

security login rest-role create -role csrestrole -api
/api/protocols/fpolicy -access all -vserver <svm_name>
security login create -user-or-group-name csuser -application http
-authmethod password -role csrestrole -vserver <svm_name>

For more information, read about xref:{relative_path}concept_ws_integration_with_ontap_access_denied.html[I ntegration with ONTAP Access Denied]

- == Configure the data collector
- .Steps for Configuration
- . Log in as Administrator or Account Owner to your Cloud Insights environment.

. Click *Workload Security > Collectors > +Data Collectors* The system displays the available Data Collectors. . Hover over the *NetApp SVM tile and click *+Monitor*. The system displays the ONTAP SVM configuration page. Enter the required data for each field. [caption=] .Configuration [cols=2*, cols"50,50"] [Options=header] |=== |Field|Description |Name |Unique name for the Data Collector |Agent|Select a configured agent from the list. |Connect via Management IP for: |Select either Cluster IP or SVM Management |Cluster / SVM Management IP Address|The IP address for the cluster or the SVM, depending on your selection above. |SVM Name|The Name of the SVM (this field is required when connecting via Cluster IP) |Username|User name to access the SVM/Cluster When adding via Cluster IP the options are: 1. Cluster-admin 2. 'csuser' 3. AD-user having similar role as csuser.

When adding via SVM IP the options are:

- 4. vsadmin
- 5. 'csuser'
- 6. AD-username having similar role as csuser.

|Password|Password for the above user name

|Filter Shares/Volumes|Choose whether to include or exclude Shares / Volumes from event collection

|Enter complete share names to exclude/include|Comma-separated list of shares to exclude or include (as appropriate) from event collection |Enter complete volume names to exclude/include|Comma-separated list of volumes to exclude or include (as appropriate) from event collection |Monitor Folder Access|When checked, enables events for folder access monitoring. Note that folder create/rename and delete will be monitored even without this option selected. Enabling this will increase the number of events monitored.

|Set ONTAP Send Buffer size|Sets the ONTAP Fpolicy send buffer size. If an

ONTAP version prior to 9.8p7 is used and performance issue is seen, then the ONTAP send buffer size can be altered to get improved ONTAP performance. Contact NetApp Support if you do not see this option and wish to explore it.

|===

.After you finish

//* Click *Test Configuration* to check the status of the collector you
configured.

* In the Installed Data Collectors page, use the options menu on the right of each collector to edit the data collector. You can restart the data collector or edit data collector configuration attributes.

== Recommended Configuration for Metro Cluster

The following is recommended for Metro Cluster:

- 1. Connect two data collectors, one to the source SVM and another to the destination SVM.
- 2. The data collectors should be connected by _Cluster IP_.
- 3. At any moment of time, one data collector should be in running, another will be in error.

The current 'running' SVM's data collector will show as _Running_. The current 'stopped' SVM's data collector will show as _Error .

- 4. Whenever there is a switchover, the state of the data collector will change from 'running' to 'error' and vice versa.
- 5. It will take up to two minutes for the data collector to move from Error state to Running state.

== Service Policy

If using service policy from ONTAP version 9.9.1, in order to connect to the Data Source Collector, the _data-fpolicy-client_ service is required along with the data service _data-nfs_, and/or _data-cifs_.

Example:

Testcluster-1::*> net int service-policy create -policy only_data_fpolicy -allowed-addresses 0.0.0.0/0 -vserver aniket_svm -services data-cifs,data-nfs,data,-core,data-fpolicy-client (network interface service-policy create)

In versions of ONTAP prior to 9.9.1, _data-fpolicy-client_ need not be set.

== Play-Pause Data Collector

2 new operations are now shown on kebab menu of collector (PAUSE and \mbox{RESUME}).

If the Data Collector is in _Running_ state, you can Pause collection. Open the "three dots" menu for the collector and select PAUSE. While the collector is paused, no data is gathered from ONTAP, and no data is sent from the collector to ONTAP. This means no Fpolicy events will flow from ONTAP to the data collector, and from there to Cloud Insights.

Note that if any new volumes, etc. are created on ONTAP while the collector is Paused, Workload Security won't gather the data and those volumes, etc. will not be reflected in dashboards or tables.

Keep the following in mind:

- * Snapshot purge won't happen as per the settings configured on a paused collector.
- * EMS events (like ONTAP ARP) won't be processed on a paused collector. This means if ONTAP identifies a ransomware attack, Cloud Insights Workload Security won't be able to acquire that event.
- * Health notifications emails will NOT be sent for a paused collector.
- * Manual or Automatic actions (such as Snapshot or User Blocking) will not be supported on a paused collector.
- * On agent or collector upgrades, agent VM restarts/reboots, or agent service restart, a paused collector will remain in Paused state.
- * If the data collector is in _Error_ state, the collector cannot be changed to _Paused_ state. The Pause button will be enabled only if the state of the collector is _Running_.
- * If the agent is disconnected, the collector cannot be changed to _Paused_ state. The collector will go into _Stopped_ state and the Pause button will be disabled.

== Troubleshooting

Known problems and their resolutions are described in the following table.

In the case of an error, click on _more detail_ in the _Status_ column for detail about the error.

```
image:CS_Data_Collector_Error.png[]
```

```
[cols=2*, options="header", cols"30,70"]
```

|===

|Problem: | Resolution:

|Data Collector runs for some time and stops after a random time, failing with: "Error message: Connector is in error state. Service name: audit. Reason for failure: External fpolicy server overloaded."

|The event rate from ONTAP was much higher than what the Agent box can handle. Hence the connection got terminated.

Check the peak traffic in CloudSecure when the disconnection happened. This you can check from the *CloudSecure > Activity Forensics > All Activity* page.

If the peak aggregated traffic is higher than what the Agent Box can handle, then please refer to the Event Rate Checker page on how to size for Collector deployment in an Agent Box.

If the Agent was installed in the Agent box prior to 4 March 2021, run the following commands in the Agent box:

```
echo 'net.core.rmem_max=8388608' >> /etc/sysctl.conf
echo 'net.ipv4.tcp_rmem = 4096 2097152 8388608' >> /etc/sysctl.conf
sysctl -p
```

Restart the collector from the UI after resizing.

|Collector reports Error Message: "No local IP address found on the connector that can reach the data interfaces of the SVM".
|This is most likely due to a networking issue on the ONTAP side. Please follow these steps:

- 1. Ensure that there are no firewalls on the SVM data lif or the management lif which are blocking the connection from the SVM.
- 2. When adding an SVM via a cluster management IP, please ensure that the

data lif and management lif of the SVM are pingable from the Agent VM. In case of issues, check the gateway, netmask and routes for the lif.

You can also try logging in to the cluster via ssh using the cluster management IP, and ping the Agent IP. Make sure that the agent IP is pingable:

network ping -vserver <vserver name> -destination <Agent IP> -lif <Lif Name> -show-detail

If not pingable, make sure the network settings in ONTAP are correct, so that the Agent machine is pingable.

- 3. If you have tried connecting via Cluster IP and it is not working, try connecting directly via SVM IP. Please see above for the steps to connect via SVM IP.
- 4. While adding the collector via SVM IP and vsadmin credentials, check if the SVM Lif has Data plus Mgmt role enabled. In this case ping to the SVM Lif will work, however SSH to the SVM Lif will not work.

If yes, create an SVM Mgmt Only Lif and try connecting via this SVM management only Lif.

- 5. If it is still not working, create a new SVM Lif and try connecting through that Lif. Make sure that the subnet mask is correctly set.
- 6. Advanced Debugging:
- a) Start a packet trace in ONTAP.
- b) Try to connect a data collector to the SVM from CloudSecure UI.
- c) Wait till the error appears. Stop the packet trace in ONTAP.
- d) Open the packet trace from ONTAP. It is available at this location

\https://<cluster mgmt ip>/spi/<clustername>/etc/log/packet traces/

- e) Make sure there is a SYN from ONTAP to the Agent box.
- f) If there is no SYN from ONTAP then it is an issue with firewall in ONTAP.
- g) Open the firewall in ONTAP, so that ONTAP is able to connect the agent box.
- 7. If it is still not working, please consult the networking team to make sure that no external firewall is blocking the connection from ONTAP to the Agent box.
- 8. Verify that port 7 is open.
- 9. If none of the above solves the issue, open a case with

link:http://docs.netapp.com/usen/cloudinsights/concept_requesting_support.html[Netapp Support] for
further assistance.

|Message: "Failed to determine ONTAP type for [hostname: <IP Address>. Reason: Connection error to Storage System <IP Address>: Host is unreachable (Host unreachable)"

- \mid 1. Verify that the correct SVM IP Management address or Cluster Management IP has been provided.
- 2. SSH to the SVM or the Cluster to which you are intending to connect. Once you are connected ensure that the SVM or the Cluster name is correct.

|Error Message: "Connector is in error state. Service.name: audit. Reason for failure: External fpolicy server terminated."

- |1. It is most likely that a firewall is blocking the necessary ports in the agent machine. Verify the port range 35000-55000/tcp is opened for the agent machine to connect from the SVM. Also ensure that there are no firewalls enabled from the ONTAP side blocking communication to the agent machine.
- 2. Type the following command in the Agent box and ensure that the port range is open.

sudo iptables-save \| grep 3500*

Sample output should look like:

- _-A IN_public_allow -p tcp -m tcp --dport 35000 -m conntrack -ctstate NEW -j ACCEPT_
- 3. Login to SVM, enter the following commands and check that no firewall is set to block the communication with ONTAP.

system services firewall show
system services firewall policy show

link:https://docs.netapp.com/ontap-

9/index.jsp?topic=%2Fcom.netapp.doc.dot-cm-nmg%2FGUID-969851BB-4302-4645-8DAC-1B059D81C5B2.html[Check firewall commands] on the ONTAP side.

4. SSH to the SVM/Cluster which you want to monitor. Ping the Agent box from the SVM data lif (with CIFS, NFS protocols support) and ensure that ping is working:

_network ping -vserver <vserver name> -destination <Agent IP> -lif <Lif</pre>

Name> -show-detail

If not pingable, make sure the network settings in ONTAP are correct, so that the Agent machine is pingable.

5.If a single SVM is added twice added to a tenant via 2 data collectors, then this error will be shown. Delete one of the data collectors thru the UI. Then restart the other data collector thru the UI. Then the data collector will show "RUNNING" status and will start receiving events from SVM.

Basically, in a tenant, 1 SVM should be added only once, via 1 data collector. 1 SVM should not added twice via 2 data collectors.

- 6. In instances where the same SVM was added in two different Workload Security environments (tenants), the last one will always succeed. The second collector will configure fpolicy with its own IP address and kick out the first one. So the collector in the first one will stop receiving events and its "audit" service will enter into error state.

 To prevent this, configure each SVM on a single environment.
- 7. This error may also occur if service policies are not configured correctly. With ONTAP 9.8 or later, in order to connect to the Data Source Collector, the data-fpolicy-client service is required along with the data service data-nfs, and/or data-cifs. Additionally, the data-fpolicy-client service must be associated with the data lif(s) for the monitored SVM.

|No events seen in activity page.

- |1. Check if ONTAP collector is in "RUNNING" state. If yes, then ensure that some cifs events are being generated on the cifs client VMs by opening some files.
- 2. If no activities are seen, please login to the SVM and enter the following command.

<SVM>event log show -source fpolicy Please ensure that there are no errors related to fpolicy.

3. If no activities are seen, please login to the SVM. Enter the following command

<SVM>fpolicy show

Check if the fpolicy policy named with prefix "cloudsecure_" has been set and status is "on". If not set, then most likely the Agent is unable to execute the commands in the SVM. Please ensure all the prerequisites as described in the beginning of the page have been followed.

|SVM Data Collector is in error state and Errror message is "Agent failed

to connect to the collector"

- |1. Most likely the Agent is overloaded and is unable to connect to the Data Source collectors.
- 2. Check how many Data Source collectors are connected to the Agent.
- 3. Also check the data flow rate in the "All Activity" page in the UI.
- 4. If the number of activities per second is significantly high, install another Agent and move some of the Data Source Collectors to the new Agent.

|SVM Data Collector shows error message as "fpolicy.server.connectError: Node failed to establish a connection with the FPolicy server "12.195.15.146" (reason: "Select Timed out")"

|Firewall is enabled in SVM/Cluster. So fpolicy engine is unable to connect to fpolicy server.

CLIs in ONTAP which can be used to get more information are:

event log show -source fpolicy which shows the error event log show -source fpolicy -fields event, action, description which shows more details.

link:https://docs.netapp.com/ontap-

9/index.jsp?topic=%2Fcom.netapp.doc.dot-cm-nmg%2FGUID-969851BB-4302-4645-8DAC-1B059D81C5B2.html[Check firewall commands] on the ONTAP side.

|Error Message: "Connector is in error state. Service name:audit. Reason for failure: No valid data interface (role: data, data protocols: NFS or CIFS or both, status: up) found on the SVM."

|Ensure there is an operational interface (having role as data and data protocol as CIFS/NFS.

 \mid The data collector goes into Error state and then goes into RUNNING state after some time, then back to Error again. This cycle repeats.

|This typically happens in the following scenario:

- 1. There are multiple data collectors added.
- 2. The data collectors which show this kind of behavior will have 1 SVM added to these data collectors. Meaning 2 or more data collectors are connected to 1 SVM.
- 3. Ensure 1 data collector connects to only 1 SVM.
- 4. Delete the other data collectors which are connected to the same SVM.

|Connector is in error state. Service name: audit. Reason for failure: Failed to configure (policy on SVM symname. Reason: Invalid value specified for 'shares-to-include' element within 'fpolicy.policy.scopemodify: "Federal'

|The share names need to be given without any quotes. Edit the ONTAP SVM

DSC configuration to correct the share names. Include and exclude shares is not intended for a long list of share names. Use filtering by volume instead if you have a large number of shares to include or exclude. |There are existing fpolicies in the Cluster which are unused. What should be done with those prior to installation of Workload Security? |It is recommended to delete all existing unused fpolicy settings even if they are in disconnected state. Workload Security will create fpolicy with the prefix "cloudsecure ". All other unused fpolicy configurations can be deleted. CLI command to show fpolicy list: fpolicy show Steps to delete fpolicy configurations: fpolicy disable -vserver <svmname> -policy-name <policy name> fpolicy policy scope delete -vserver <svmname> -policy-name <policy name> fpolicy policy delete -vserver <svmname> -policy-name <policy name> fpolicy policy event delete -vserver <svmname> -event-name <event list> fpolicy policy external-engine delete -vserver <svmname> -engine-name <engine name> |After enabling Workload Security, ONTAP performance is impacted: Latency becomes sporadically high, IOPs become sporadically low. |While using ONTAP with Workload Security sometimes latency issues can be seen in ONTAP. There are a number of possible reasons for this as noted in the following: link:https://mysupport.netapp.com/site/bugsonline/product/ONTAP/BURT/1372994[1372994], https://mysupport.netapp.com/site/bugsonline/product/ONTAP/BURT/1415152[1415152], https://mysupport.netapp.com/site/bugsonline/product/ONTAP/BURT/1438207[1438207], https://mysupport.netapp.com/site/bugsonline/product/ONTAP/BURT/1479704[1479704], https://mysupport.netapp.com/site/bugsonline/product/ONTAP/BURT/1354659[1354659]. All of these issues are fixed in ONTAP 9.13.1 and later; it is strongly recommended to use one of these later versions. |Data collector is in error, shows this error message. "Error: Connector is in error state. Service name: audit. Reason for

failure: Failed to configure policy on SVM svm test. Reason: Missing value for zapi field: events. " |Start with a new SVM with only NFS service configured. Add an ONTAP SVM data collector in Workload Security. CIFS is configured as an allowed protocol for the SVM while adding the ONTAP SVM Data Collector in Workload Security. Wait until the Data collector in Workload Security shows an error. Since the CIFS server is NOT configured on the SVM, this error as shown in the left is shown by Workload Security. Edit the ONTAP SVM data collector and un-check CIFs as allowed protocol. Save the data collector. It will start running with only NFS protocol enabled. |Data Collector shows the error message: "Error: Failed to determine the health of the collector within 2 retries, try restarting the collector again (Error Code: AGENT008)". |1. On the Data Collectors page, scroll to the right of the data collector giving the error and click on the 3 dots menu. Select Edit . Enter the password of the data collector again. Save the data collector by pressing on the Save button. Data Collector will restart and the error should be resolved.

2. The Agent machine may not enough CPU or RAM headroom, that is why the

Please check the number of Data Collectors which are added to the Agent in the machine.

If it is more than 20, please increase the CPU and RAM capacity of the Agent machine.

Once the CPU and RAM is increased, the DSCs will get into Initializing and then to Running state automatically.

Look into the sizing guide on link:https://docs.netapp.com/us-en/cloudinsights/concept cs event rate checker.html[this page].

|===

DSCs are failing.

If you are still experiencing problems, reach out to the support links mentioned in the *Help > Support* page.

[[ID7a8d1548286601e3093c9c9eded74cf6]]

= Configuring the Cloud Volumes ONTAP and Amazon FSx for NetApp ONTAP collector

:toc: macro
:hardbreaks:
:toclevels: 1

```
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root path} {relative path} ./media/
```

[.lead]

Workload Security uses data collectors to collect file and user access data from devices.

== Cloud Volumes ONTAP Storage Configuration

See the OnCommand Cloud Volumes ONTAP Documentation to configure a single-node / HA AWS instance to host the Workload Security Agent: https://docs.netapp.com/us-en/cloud-manager-cloud-volumes-ontap/index.html

After the configuration is complete, follow the steps to setup your SVM: https://docs.netapp.com/us-en/cloudinsights/task add collector svm.html

- == Supported Platforms
- * Cloud Volumes ONTAP, supported in all the available cloud service providers wherever available. For example: Amazon, Azure, Google Cloud.
- * ONTAP Amazon FSx
- == Agent Machine Configuration

The agent machine must be configured in the respective subnets of the cloud Service providers. Read more about network access in the [Agent Requirements].

Below are the steps for Agent installation in AWS. Equivalent steps, as applicable to the cloud service provider, can be followed in Azure or Google Cloud for the installation.

In AWS, use the following steps to configure the machine to be used as a Workload Security Agent:

Use the following steps to configure the machine to be used as a Workload Security Agent:

.Steps

. Log in to the AWS console and navigate to EC2-Instances page and select Launch instance .

```
. Select a RHEL or CentOS AMI with the appropriate version as mentioned in
this page:
https://docs.netapp.com/us-
en/cloudinsights/concept cs agent requirements.html
. Select the VPC and Subnet that the Cloud ONTAP instance resides in.
. Select t2.xlarge (4 vcpus and 16 GB RAM) as allocated resources.
.. Create the EC2 instance.
. Install the required Linux packages using the YUM package manager:
.. Install wget and unzip native Linux packages.
////
.. Install selinux (dependency package for the docker-ce):
wget http://mirror.centos.org/centos/7/extras/x86_64/Packages/container-
selinux-2.68-1.el7.noarch.rpm
yum install -y container-selinux-2.68-1.el7.noarch.rpm
. Install the docker-ce (not the native docker) package. You must use a
version higher than 17.03:
https://download.docker.com/linux/centos/7/x86 64/stable/Packages/
. SSH to the Redhat EC2 VM:
 ssh -i "your new pem.pem" <ec2 hostname or IP>
 sudo su -
. Perform a docker login after installing the required AWS CLI package:
 curl "https://s3.amazonaws.com/aws-cli/awscli-bundle.zip" -o "awscli-
bundle.zip"
 unzip awscli-bundle.zip
 sudo ./awscli-bundle/install -i /usr/local/aws -b /usr/local/bin/aws
 /usr/local/bin/aws --version
 aws configure --profile collector_readonly
 aws ecr get-login --no-include-email --region us-east-1 --profile
collector readonly
docker login -u AWS -p <token generated above> <ECR hostname>
. Use the following command to verify the steps completed successfully and
the cs-ontap-dsc image can be successfully pulled:
```

```
docker pull 376015418222.dkr.ecr.us-east-1.amazonaws.com/cs-ontap-
dsc:1.25.0
////
== Install the Workload Security Agent
. Log in as Administrator or Account Owner to your Cloud Insights
environment.
. Navigate to Workload Security *Collectors* and click the *Agents* tab.
. Click *+Agent* and specify RHEL as the target platform.
. Copy the Agent Installation command.
. Paste the Agent Installation command into the RHEL EC2 instance you are
logged in to.
This installs the Workload Security agent, providing all of the
xref:{relative path}concept_cs_agent_requirements.html[Agent
Prerequisites] are met.
For detailed steps please refer to this xref:{relative path}
https://docs.netapp.com/us-en/cloudinsights/task cs add agent.html#steps-
to-install-agent
////
== Add a NetApp ONTAP data collector
. Click *Observability > Collectors > Data Collectors > +Data Collector*
and specify the NetApp ONTAP Cloud Volumes data collector. Enter the
required information in the fields.
[caption=]
.Configuration
[cols=2*, cols"50,50"]
[Options=header]
|===
|Field|Description
|Name |Unique name for the Data Collector
|Agent|Select a configured agent from the list or click *Add Agent* to
configure an Agent. See
xref:{relative path}concept cs agent requirements.html[Agent requirements]
and xref: {relative path} task cs add agent.html [Agent Installation] for
configuration information.
|SVM Management IP Address|Management IP Address
|Username|User name to access the SVM
```

```
|Password|SVM Password
|Enter complete share names to exclude|Comma-separated list of shares to
exclude from event collection
|Enter complete volume names to exclude|Comma-separated list of volumes to
exclude from event collection
|===
.. Click *Add Collector*
. Verify the Agent Server is running using the
`docker ps` command and a `docker logs <docker image id>` file.
All of the data collector's service status should be in the 'running'
state.
// .. Identify an NFS client (in the same VPC subnet as the Agent and
Cloud ONTAP)
//.. Install the nfs-utils package in this VPC Subnet:
//+
//'yum install -y nfs-utils'
//.. NFS mount the volume / qtree container that was created in the SVM.
////
== Troubleshooting
Known problems and their resolutions are described in the following table.
|===
|Problem |Resolution
|"Workload Security: Failed to determine ONTAP type for Amazon FxSN data
collector" error is shown by the Data Collector.
Customer is unable to add new Amazon FSxN data collector into Workload
Security. Connection to FSxN cluster on port 443 from the agent is timing
out. Firewall and AWS security groups have the required rules enabled to
allow communication. An agent is already deployed and is in the same AWS
account as well. This same agent is used to connect and monitor the
remaining NetApp devices (and all of them are working).
|Solve this issue by adding fsxadmin LIF network segment to agent's
security rule.
Allowed all ports if you are not sure about the ports.
|===
```

```
[[IDc848d791112b9127eadbfa0939f1099d]]
= User Management
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Workload Security user accounts are managed through Cloud Insights.
Cloud Insights provides four user account levels: Account Owner,
Administrator, User, and Guest. Each account is assigned specific
permission levels. A User account that has Administrator privileges can
create or modify users, and assign each user one of the following Workload
Security roles:
|===
        |Workload Security Access
|Role
|Administrator
|Can perform all Workload Security functions, including those for Alerts,
Forensics, data collectors, automated response policies, and APIs for
Workload Security.
An Administrator can also invite other users but can only assign Workload
Security roles.
lUser
|Can view and manage Alerts and view Forensics. User role can change alert
status, add a note, take snapshots manually, and restrict user access.
|Guest
|Can view Alerts and Forensics. Guest role cannot change alert status, add
a note, take snapshots manually, or restrict user access.
|===
```

```
.Steps
. Log into Workload Security
. In the menu, click *Admin > User Management*
You will be forwarded to Cloud Insights's User Management page.
. Select the desired role for each user.
While adding a new user, simply select the desired role (usually User or
Guest).
More information on User accounts and roles can be found in the Cloud
Insights link:https://docs.netapp.com/us-
en/cloudinsights/concept user roles.html[User Role] documentation.
[[ID0acdeeac44f83b6c74cfa9027f41b028]]
= SVM Event Rate Checker (Agent Sizing Guide)
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The Event Rate Checker is used to check the NFS/SMB combined event rate in
the SVM before installing an ONTAP SVM data collector, to see how many
SVMs one Agent machine will be able to monitor. Use the Event Rate Checker
as a sizing guide to help plan your security environment.
An Agent can support up to a maximum of 50 data collectors.
== Requirements:
* Cluster IP
* Cluster admin username and password
NOTE: When running this script no ONTAP SVM Data Collector should be
running for the SVM for which event rate is being determined.
Steps:
```

. Install the Agent by following the instructions in CloudSecure. . Once the agent is installed, run the server data rate checker.sh script as a sudo user: /opt/netapp/cloudsecure/agent/install/svm event rate checker.sh . This script requires sshpass to be installed in the linux machine. There are two ways to install it: .. Run the following command: linux_prompt> yum install sshpass .. If that does not work, then download sshpass to the linux machine from the web and run the following command: linux prompt> rpm -i sshpass Provide the correct values when prompted. See below for an example. The script will take approximately 5 minutes to run. After the run is complete, the script will print the event rate from the SVM. You can check Event rate per SVM in the console output:

"Svm svm_rate is generating 100 events/sec".

//This will show the rate of generation of Events for a SVM.

Each Ontap SVM Data Collector can be associated with a single SVM, which means each data collector will be able to receive the number of events which a single SVM generates.

Keep the following in mind:

A) Use this table as a general sizing guide. You can increase the number of cores and/or memory to increase the number of data collectors supported, up to a maximum of 50 data collectors:

|===

|Agent Machine Configuration |Number of SVM Data Collectors |Max event Rate which the Agent Machine can handle

|4 core, 16GB | 10 data collectors | 20K events/sec |4 core, 32GB | 20 data collectors | 20K events/sec |===

- B) To calculate your total events, add the Events generated for all SVMs for that agent.
- C) If the script is not run during peak hours or if peak traffic is difficult to predict, then keep an event rate buffer of 30%.
- ${\sf B}$ + C Should be less than A, otherwise the Agent machine will fail to monitor.

In other words, the number of data collectors which can be added to a single agent machine should comply to the formula below:

Sum of all Event rate of all Data Source Collectors + Buffer Event rate of 30% < 20000 events/second

See the xref:{relative_path}concept_cs_agent_requirements.html[Agent Requirements] page for additional pre-requisites and requirements.

== Example

Let us say we have three SVMS generating event rates of 100, 200, and 300 events per second, respectively.

We apply the formula:

(100+200+300) + [(100+200+300)*30%] = 600+180 = 780events/sec 780 events/second is < 20000 events/second, so the 3 SVMs can be monitored via one agent box.

Console output is available in the Agent machine in the file name __fpolicy_stat_<SVM Name>.log__ in the present working directory.

The script may give erroneous results in the following cases:

- * Incorrect credentials, IP, or SVM name are provided.
- * An already-existing fpolicy with same name, sequence number, etc. will give error.
- * The script is stopped abruptly while running.

An example script run is shown below:

[root@ci-cs-data agent]#
/opt/netapp/cloudsecure/agent/install/svm event rate checker.sh

```
Enter the cluster ip: 10.192.139.166
 Enter the username to SSH: admin
 Enter the password:
 Getting event rate for NFS and SMB events.
 Available SVMs in the Cluster
 QA SVM
 Stage SVM
 Qa-fas8020
 Oa-fas8020-01
 Qa-fas8020-02
 audit svm
 svm rate
 vs new
 vs new2
 _____
 Enter [1/5] SVM name to check (press enter to skip): svm rate
 Enter [2/5] SVM name to check (press enter to skip): audit svm
 Enter [3/5] SVM name to check (press enter to skip):
 Enter [4/5] SVM name to check (press enter to skip):
 Enter [5/5] SVM name to check (press enter to skip):
 Running check for svm svm rate...
 Running check for svm audit svm...
 Waiting 5 minutes for stat collection
 Stopping sample svm rate sample
 Stopping sample audit svm sample
 fpolicy stats of svm svm rate is saved in fpolicy stat svm rate.log
 Svm svm rate is generating 100 SMB events/sec and 100 NFS events/sec
 Overall svm svm rate is generating 200 events/sec
 fpolicy stats of svm audit svm is saved in fpolicy stat audit svm.log
 Svm audit svm is generating 200 SMB events/sec and 100 NFS events/sec
 Overall svm audit svm is generating 300 events/sec
 [root@ci-cs-data agent]#
== Troubleshooting
|===
|Question|Answer
|If I run this script on an SVM that is already configured for Workload
Security, does it just use the existing fpolicy config on the SVM or does
```

it setup a temporary one and run the process? |The Event Rate Checker can run fine even for an SVM already configured for Workload Security. There should be no impact. |Can I increase the number of SVMs on which the script can be run? |Yes. Simply edit the script and change the max number of SVMs from 5 to any desirable number. |If I increase the number of SVMs, will it increase the time of running of the script? |No. The script will run for a max of 5 minutes, even if the number of SVMs is increased. |Can I increase the number of SVMs on which the script can be run? |Yes. You need to edit the script and change the max number of SVMs from 5 to any desirable number. | If I increase the number of SVMs, will it increase the time of running of |No. The script will run for a max of 5mins, even if the number of SVMs are increased. |What happens if I run the Event Rate Checker with an existing agent? |Running the Event Rate Checker against an already-existing agent may cause an increase in latency on the SVM. This increase will be temporary in nature while the Event rate Checker is running. |=== :leveloffset: -1 [[IDecc629975b176d4b874a1ad1505c92d7]] = Alerts :hardbreaks: :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path} {relative path}./media [.lead] The Workload Security Alerts page shows a timeline of recent attacks and/or warnings and allows you to view details for each issue. NOTE: Workload Security is not available in Cloud Insights Federal

Edition. image:CloudSecureAlertsListPage.png[Alerts list] //// The Alerts page shows all alerts generated by Workload Security. Use this page to identify recent alerts and the users generating those alerts. You can also access all alerts that have been raised with the ability to drill down into individual alerts. //// //// == History History shows the number of alerts that have been raised over the last seven days. Hovering over the severity of the alerts displays the number, severity, and occurrence date for each alert type. == Notable Users * Shows a list of the users that have generated the highest number of alerts. * Shows the type of alerts generated. * Shows the total number of alerts generated for each user. //// == Alert The Alert list displays a graph showing the total number of Potential Attacks and/or Warnings that have been raised in the selected time range, followed by a list of the attacks and/or warnings that occurred in that time range. You can change the time range by adjusting the start time and end time sliders in the graph. The following is displayed for each alert: *Potential Attacks:* * The Potential Attack type (for example, Ransomware or Sabotage) * The date and time the potential attack was Detected

```
* The Status of the alert:
** *New*: This is the default for new alerts.
** *In Progress*: The alert is under investigation by a team member or
** *Resolved*: The alert has been marked as resolved by a team member.
** *Dismissed*: The alert has been dismissed as false positive or expected
behavior.
An administrator can change the status of the alert and add a note to
assist with investigation.
image:CloudSecureChangeAlertStatus.png[Change Alert Status]
* The User whose behavior triggered the alert
* Evidence of the attack (for example, a large number of files was
encrypted)
* The _Action Taken (for example, a snapshot was taken)
*Warnings:*
* The _Abnormal Behavior_ that triggered the warning
* The date and time the behavior was Detected
* The Status of the alert (New, In progress, etc.)
* The User whose behavior triggered the alert
* A description of the Change (for example, an abnormal increase in file
access)
* The Action Taken
== Filter Options
You can filter Alerts by the following:
* The Status of the alert
* Specific text in the Note
* The type of Attacks/Warnings
* The User whose actions triggered the alert/warning
```

== The Alert Details page

You can click an alert link on the Alerts list page to open a detail page for the alert. Alert details may vary according to the type of attack or alert. For example, a Ransomware Attack detail page may show the following information:

=== Summary section:

- * Attack type (Ransomware, Sabotage) and Alert ID (assigned by Workload Security)
- * Date and Time the attack was detected
- * Action Taken (for example, an automatic snapshot was taken. Time of snapshot is shown immediately below the summary section))
- * Status (New, In Progress, etc.)

=== Attack Results section:

- * Counts of Affected Volumes and Files
- * An accompanying summary of the detection
- * A graph showing file activity during the attack

=== Related Users section:

This section shows details about the user involved in the potential attack, including a graph of Top Activity for the user.

Alerts page (this example shows a potential ransomware attack): image:RansomwareAlertExample.png[Ransomware Alert Example]

Detail page (this example shows a potential ransomware attack): image:RansomwareDetailPageExample.png[Ransomware Detail Page Example]

== Take a Snapshot Action

Workload Security protects your data by automatically taking a snapshot when malicious activity is detected, ensuring that your data is safely backed up.

You can define

xref:{relative_path}cs_automated_response_policies.html[automated response
policies] that take a snapshot when ransomware attack or other abnormal
user activity is detected.

You can also take a snapshot manually from the alert page.

Automatic Snapshot taken:

image:AlertActionsAutomaticExample.png[Alert Action Screen, 1000]

Manual Snapshot:

image:AlertActionsExample.png[Alert Action Screen,1000]

== Alert Notifications

Email notifications of alerts are sent to an alert recipient list for every action on the alert. To configure alert recipients, click on *Admin > Notifications* and enter an email addresses for each recipient.

== Retention Policy

Alerts and Warnings are retained for 13 months. Alerts and Warnings older than 13 months will be deleted.

If the Workload Security environment is deleted, all data associated with the environment is also deleted.

== Troubleshooting

|===

|Problem: |Try This:

|There is a situation where, ONTAP takes hourly snapshots per day. Will Workload Security (WS) snapshots affect it? Will WS snapshot take the hourly snapshot place? Will the default hourly snapshot get stopped? |Workload Security snapshots will not affect the hourly snapshots. WS snapshots will not take the hourly snapshot space and that should continue as before. The default hourly snapshot will not get stopped.

|What will happen if the maximum snapshot count is reached in ONTAP? |If the maximum Snapshot count is reached, subsequent Snapshot taking will fail and Workload Security will show an error message noting that Snapshot is full.

User needs to define Snapshot policies to delete the oldest snapshots, otherwise snapshots will not be taken.

In ONTAP 9.3 and earlier, a volume can contain up to 255 Snapshot copies. In ONTAP 9.4 and later, a volume can contain up to 1023 Snapshot copies.

See the ONTAP Documentation for information on link:https://docs.netapp.com/ontap9/index.jsp?topic=%2Fcom.netapp.doc.dot-cm-cmpr960%2Fvolume__snapshot__autodelete__modify.html[setting Snapshot deletion policy].

|Workload Security is unable to take snapshots at all. |Make sure that the role being used to create snapshots has xref:{relative path} https://docs.netapp.com/usen/cloudinsights/task add collector svm.html#a-note-aboutpermissions[proper rights assigned]. Make sure csrole is created with proper access rights for taking snapshots: security login role create -vserver <vservername> -role csrole -cmddirname "volume snapshot" -access all |Snapshots are failing for older alerts on SVMs which were removed from Workload Security and subsequently added back again. For new alerts which occur after SVM is added again, snapshots are taken. |This is a rare scenario. In the event you experience this, log in to ONTAP and take the snapshots manually for the older alerts. |In the _Alert Details_ page, the message "Last attempt failed" error is seen below the Take Snapshot button. Hovering over the error displays "Invoke API command has timed out for the data collector with id". |This can happen when a data collector is added to Workload Security via SVM Management IP, if the LIF of the SVM is in disabled state in ONTAP. Enable the particular LIF in ONTAP and trigger Take Snapshot manually from Workload Security. The Snapshot action will then succeed. |=== = Forensics :leveloffset: +1 [[ID71413cd0f902458b4a03f0e3742bab78]] = Forensics - All Activity :toc: macro :hardbreaks: :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/

[.lead]

The All Activity page helps you understand the actions performed on entities in the Workload Security environment.

== Examining All Activity Data

Click *Forensics > Activity Forensics* and click the *All Activity* tab to access the All Activity page.

This page provides an overview of activities in your environment, highlighting the following information:

* A graph showing _Activity History_ (accessed per minute/per 5 minutes/per 10 minutes based on selected global time range)

You can zoom the graph by dragging out a rectangle in the graph. The entire page will be loaded to display the zoomed time range. When zoomed in, a button is displayed that lets the user zoom out.

- * A chart of _Activity Types_. To obtain activity history data by activity type, click on corresponding x-axis label link.
- * A chart of Activity on _Entity Types_. To obtain activity history data by entity type, click on corresponding x-axis label link.
- * A list of the All Activity data

The _*All Activity*_ table shows the following information. Note that not all of these columns are displayed by default. You can select columns to display by clicking on the "gear" icon image:GearIcon.png[gear icon].

- * The *time* an entity was accessed including the year, month, day, and time of the last access.
- * The *user* that accessed the entity with a link to the xref:{relative path}forensic user overview.html[User information].

//Above should be new user profile?

- * The *activity* the user performed. Supported types are:
- ** *Change Group Ownership* Group Ownership is of file or folder is changed. For more details about group ownership please see link:https://docs.microsoft.com/en-us/previous-versions/orphantopics/ws.11/dn789205(v=ws.11)?redirectedfrom=MSDN[this link.]
- ** *Change Owner* Ownership of file or folder is changed to another user.

- ** *Change Permission* File or folder permission is changed.
- ** *Create* Create file or folder.
- ** *Delete* Delete file or folder. If a folder is deleted, _delete_ events are obtained for all the files in that folder and subfolders.
- ** *Read* File is read.
- ** *Read Metadata* Only on enabling folder monitoring option. Will be generated on opening a folder on Windows or Running "ls" inside a folder in Linux.
- ** *Rename* Rename file or folder.
- ** *Write* Data is written to a file.
- ** *Write Metadata* File metadata is written, for example, permission changed.
- ** *Other Change* Any other event which are not described above. All unmapped events are mapped to "Other Change" activity type. Applicable to files and folders.
- * The *Path* to the entity with a link to the xref:{relative_path}forensic_entity_detail.html[Entity Detail Data]
- * The *Entity Type*, including entity (i.e. file) extension (.doc, .docx, .tmp, etc.)
- * The *Device* where the entities reside
- * The *Protocol* used to fetch events.
- * The *Original Path* used for rename events when the original file was renamed. This column is not visible in the table by default. Use the column selector to add this column to the table.
- * The *Volume* where the entities reside. This column is not visible in the table by default. Use the column selector to add this column to the table.
- //* The *Source IP* address from which the activity was performed.

```
== Filtering Forensic Activity History Data
There are two methods you can use to filter data.
  Hover over the field in the table and click the filter icon that
appears. The value is added to the appropriate filters in the top _Filter
By list.
   Filter data by typing in the Filter By field:
+
Select the appropriate filter from the top 'Filter By' widget by clicking
the *[+]* button:
image:Forensic Activity Filter.png[Entity Filer, width=500]
Enter the search text
Press Enter or click outside of the filter box to apply the filter.
You can filter Forensic Activity data by the following fields:
* The *Activity* type.
////
** Change Group Ownership
** Change Owner
** Change Permission
** Copy
** Create
** Delete
** Move
** Read
** Read Metadata
** Rename
** Write
** Write Metadata
** Other Change
////
* *Source IP* from which the entity was accessed. You must provide a valid
source IP address in double quotes, for example "10.1.1.1.". Incomplete
IPs such as "10.1.1.*", "10.1.*.*", etc. will not work.
* *Protocol* to fetch protocol-specific activities.
//* *Noise Reduction* to filter activities on temporary files which are
```

generated as part of the normal operating process. If noise reduction is enabled, temporary files of extension .tmp, .ldb, .laccdb, .\$db etc. are filtered.

- * *Username* of the user performing the activity. You need to provide the exact Username to filter. Search with partial username, or partial username prefixed or suffixed with '*' will not work.
- * *Noise Reduction* to filter files which are created in the last 2 hours by the user. It is also used to filter temporary files (for example, .tmp files) accessed by the user.

The following fields are subject to special filtering rules:

- * *Entity Type*, using entity (file) extension
- * *Path* of the entity
- * *User* performing the activity
- * *Device* (SVM) where entities reside
- * *Volume* where entities reside
- * The *Original Path* used for rename events when the original file was renamed.

The preceding fields are subject to the following when filtering:

- * Exact value should be within quotes: Example: "searchtext"
- * Wildcard strings must contain no quotes: Example: searchtext,
- *searchtext*, will filter for any strings containing 'searchtext'.
- * String with a prefix, Example: searchtext*, will search any strings which start with 'searchtext'.
- == Sorting Forensic Activity History Data

You can sort activity history data by _Time, User, Source IP, Activity, Path_ and _Entity Type_. By default, the table is sorted by descending _Time_ order, meaning the latest data will be displayed first. Sorting is disabled for _Device_ and _Protocol_ fields.

== Exporting All Activity

You can export the activity history to a .CSV file by clicking the _Export_ button above the Activity History table. Note that only the top 100,000 records are exported. Depending on the amount of data, it may take from a few seconds to several minutes for the export to finish.

```
== Column Selection for All Activity
The All activity table shows select columns by default. To add, remove,
or change the columns, click the gear icon on the right of the table and
select from the list of available columns.
image:CloudSecure ActivitySelection.png[Activity Selector, width=30%]
== Activity History Retention
Activity history is retained for 13 months for active Workload Security
environments.
== Applicability of Filters in Forensics Page
|===
|Filter | What it does | Example | Applicable in Which Filters? | Not
applicable for which filters | Result
|* (Asterisk) | enables you to search for everything | Auto*03172022
|User, PATH, Entity Type, Device Type, Volume, Original Path
||returns all resources that start with "Auto" and end with "03172022"
|? (question mark) |enables you to search for a specific number of
characters | AutoSabotageUser1 03172022? | User, Entity Type, Device,
Volume | |returns AutoSabotageUser1 03172022A,
AutoSabotageUser1 03172022AB, AutoSabotageUser1 031720225, and so on
|OR |enables you to specify multiple entities
|AutoSabotageUser1 03172022 OR AutoRansomUser4 03162022
    |User, Domain, Username, PATH, Entity Type, Device, Original Path
||returns any of AutoSabotageUser1 03172022 OR AutoRansomUser4 03162022
       |allows you to exclude text from the search results |NOT
AutoRansomUser4 03162022
    |User, Domain, Username, PATH, Entity Type, Original PATH, Volume
|Device | returns everything that does not start
with"AutoRansomUser4 03162022"
|None | searches for NULL values in all fields | None | Domain |
|returns results where the target field is empty
|===
== Path / Original path Search
```

```
Search results with and without / will be different
|===
|/AutoDir1/AutoFile |Works
|AutoDir1/AutoFile |Doesn't work
|/AutoDir1/AutoFile (Dir1) |Dir1 Partial substring doesn't work
|"/AutoDir1/AutoFile03242022" |Exact search works
|Auto*03242022 |Doesn't work
|AutoSabotageUser1 03172022? |Doesn't work
|/AutoDir1/AutoFile03242022 OR /AutoDir1/AutoFile03242022 |Works
|NOT /AutoDir1/AutoFile03242022 |Works
|NOT /AutoDir1 |Works
|NOT /AutoFile03242022 |Doesn't work
|* |Shows all the entries
|===
== Troubleshooting
|===
|Problem|Try This
|In the "All Activities" table, under the 'User' column, the user name is
shown as:
"ldap:HQ.COMPANYNAME.COM:S-1-5-21-3577637-1906459482-1437260136-1831817"
"ldap:default:80038003"
|Possible reasons could be:
1. No User Directory Collectors have been configured yet. To add one, go
to *Workload Security > Collectors > User Directory Collectors* and click
on *+User Directory Collector*. Choose Active Directory or LDAP
Directory Server .
2. A User Directory Collector has been configured, however it has stopped
or is in error state. Please go to *Collectors > User Directory
Collectors* and check the status. Refer to the
link:http://docs.netapp.com/us-
en/cloudinsights/task config user dir connect.html#troubleshooting-user-
directory-collector-configuration-errors[User Directory Collector
troubleshooting] section of the documentation for troubleshooting tips.
After configuring properly, the name will get automatically resolved
within 24 hours.
If it still does not get resolved, check if you have added the correct
```

User Data Collector. Make sure that the user is indeed part of the added Active Directory/LDAP Directory Server.

|Some NFS events are not seen in UI.

|Check the following:

- 1. A user directory collector for AD server with POSIX attributes set should be running with the unixid attribute enabled from UI.
- 2. Any user doing NFS access should be seen when searched in the user page from ${\tt UI}$
- 3. Raw events (Events for whom the user is not yet discovered) are not supported for NFS $\,$
- 4. Anonymous access to the NFS export will not be monitored.
- 5. Make sure NFS version used in lesser than NFS4.1.

|After typing some letters containing a wildcard character like asterisk (*) in the filters on the Forensics _All Activity_ or _Entities_ pages, the pages load very slowly.

|An asterisk ($\$) in the search string searches for everything. However, leading wildcard strings like _*<searchTerm>_ or _*<searchTerm>*_ will result in a slow query.

To get better performance, use prefix strings instead, in the format _<searchTerm>*_ (in other words, append the asterisk (*) _after_ a search term).

Example: use the string _testvolume*_, rather than _*testvolume_ or _*test*volume_.

Use a prefix-based search to see all activities underneath a given folder recursively (Hierarchical search). e.g. _/path1/path2/path3_ or _"/path1/path2/path3"_ will list all the activities recursively under _/path1/path2/path3_.

Alternatively use the "Add To Filter" option under the All Activity tab.

|I| am encountering a "Request failed with status code 500/503" error when using a Path filter.

|Try using a smaller date range for filtering records.

|===

[[ID4c29e0ffac528ad2d9e73065921c3667]]

= Forensic Entities Page

```
:toc: macro
:hardbreaks:
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[Lead]
The Forensics Entities page provides detailed information about entity
activity in your environment.
== Examining Entity Information
Click *Forensics > Activity Forensics* and click the Entities tab to
access the Entities page.
This page provides an overview of entity activity in your environment,
highlighting the following information:
* A graph showing Unique Entities accessed per minute
* A chart of Entity Types Accessed
* A breakdown of the Common Paths
* A list of the Top 50 Entities out of the total number of entities
image:CS-Entities-Page.png[Entities Page]
Clicking on an entity in the list opens an overview page for the entity,
showing a profile of the entity with details like name, type, device name,
most accessed location IP, and path, as well as the entity behavior such
as the user, IP, and time the entity was last accessed.
image:CS-entity-detail-page.png[Entity Overview Page]
[[ID7a15117d7c1691539dfc241d15546d0e]]
= Forensic User Overview
:hardbreaks:
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root_path}{relative_path}./media/
```

[lead]

Information for each user is provided in the User Overview. Use these views to understand user characteristics, associated entities, and recent activities.

== User Profile

User Profile information includes contact information and location of the user. The profile provides the following information:

- * Name of the user
- * Email address of the user
- * User's Manager
- * Phone contact for the user
- * Location of the user

== User Behavior

The user behavior information identifies recent activities and operations performed by the user. This information includes:

- * Recent activity
- ** Last access location
- ** Activity graph
- ** Alerts
- //** Entities accessed
- * Operations for the last seven days
- ** Number of operations
- //** Number of read operations
- //** Number of times meta data was accessed
- == Refresh Interval

The User list is refreshed every 12 hours.

== Retention Policy

If not refreshed again, the User list is retained for 13 months. After 13 months, the data will be deleted.

If your Workload Security environment is deleted, all data associated with the environment is deleted.

```
:leveloffset: -1
[[ID0f827154d0e47b78b03f3d989303fa61]]
= Automated Response Policies
:hardbreaks:
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path} {relative path}./media
[.lead]
Response Policies trigger actions such as taking a snapshot or restricting
user access in the event of an attack or abnormal user behavior.
NOTE: Workload Security is not available in Cloud Insights Federal
Edition.
You can set policies on specific devices or all devices. To set a response
policy, select *Admin > Automated Response Policies* and click the
appropriate *+Policy* button. You can create policies for Attacks or for
Warnings.
image:Automated Response Screenshot.png[Create Attack Policy]
You must save the policy with a unique name.
To disable an automated response action (for example, Take Snapshot),
simply un-check the action and save the policy.
When an alert is triggered against the specified devices (or all devices,
if selected), the automated response policy takes a snapshot of your data.
You can see snapshot status on the
xref:{relative path}cs alert data.html#the-alert-details-page[Alert detail
page].
See the xref:{relative path}cs restrict user access.html[Restrict User
Access] page for more details on restricting user access by IP.
You can modify or pause an Automated Response Policy by choosing the
option in the policy's drop-down menu.
Workload Security will automatically delete snapshots once per day based
on the Snapshot Purge settings.
//image:AutomatedResponsePolicyList.png[Automated Response Policy Pause]
//image:CloudSecure AutomatedResponsePolicies WithSnapshotPurge.png[Automa
```

```
ted Response Policy Settings]
image:CloudSecure SnapshotPurgeSettings.png[Snapshot Purge Settings]
[[ID7058ce0c7d60255fac67b1b363b4d767]]
= Allowed File Types Policies
:hardbreaks:
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media
[.lead]
If a ransomware attack is detected for a known file extension, and alerts
are being generated on the Alerts screen, then that file extension can be
added to an allowed file types list to prevent unnecessary alerting.
Navigate to *Workload Security > Policies* and go to the Allowed File
Type Policies tab.
image:WS Allowed File Type Policies.png[Allowed File Types Policies]
Once added to the allowed file types list, no ransomware attack alert
will be generated for that allowed file type. Note that the Allowed File
Types policy is only applicable for ransomware detection.
For example, if a file named test.txt is renamed to test.txt.abc and
Workload Security is detecting a ransomware attack because of the .abc
extension, the _.abc_ extension can be added to the _allowed file types_
list. After being added to the list, ransomware attacks will no longer be
generated against files with the .abc extension.
Allowed File Types can be exact matches (e.g., ".abc") or expressions
(e.g., ".*type", ".type*", or "*type*"). Expressions of types ".a*c",
".p*f" are not supported.
[[IDc7756021f9616e005480f10f5a67d12a]]
= Integration with ONTAP Autonomous Ransomware Protection
:toc: macro
:hardbreaks:
```

```
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root path}{relative path}./media/
```

[.lead]

The ONTAP Autonomous Ransomware Protection (ARP) feature uses workload analysis in NAS (NFS and SMB) environments to proactively detect and warn about abnormal in-file activity that might indicate a ransomware attack.

Additional details and license requirements about ARP can be found link:https://docs.netapp.com/us-en/ontap/anti-ransomware/index.html[here].

Workload Security integrates with ONTAP to receive ARP events and provide an additional analytics and automatic responses layer.

Workload Security receives the ARP events from ONTAP and takes the following actions:

- . Correlates volume encryption events with user activity to identify who is causing the damage.
- . Implements automatic response policies (if defined)
- . Provides forensics capabilities:
- ** Allow customers to conduct data breach investigations.
- ** Identify what files were affected, helping to recover faster and conduct data breach investigations.

== Prerequisites

- . Minimum ONTAP version: 9.11.1
- . ARP enabled volumes. Details on enabling ARP can be found link:https://docs.netapp.com/us-en/ontap/anti-ransomware/enable-task.html[here]. ARP must be enabled via OnCommand System Manager. Workload Security cannot enable ARP.
- . Workload Security collector should be added via cluster IP.
- . Cluster level credentials are needed for this feature to work. In other words, cluster level credentials must be used when adding the SVM.
- == User permissions required

If you are using cluster administration credentials, no new permissions are needed.

If you are using a custom user (for example, _csuser_) with permissions given to the user, then follow the steps below to give permissions to Workload Security to collect ARP related information from ONTAP.

```
For csuser with cluster credentials, do the following from the ONTAP
command line:
 security login rest-role create -role arwrole -api /api/storage/volumes
-access readonly -vserver <cluster name>
 security login rest-role create -api /api/security/anti-ransomware
-access readonly -role arwrole -vserver <cluster name>
 security login create -user-or-group-name csuser -application http
-authmethod password -role arwrole
Read more about configuring other
xref:{relative path}task add collector svm.html[ONTAP permissions].
== Sample Alert
A sample alert generated due to ARP event is shown below:
//image:CS ONTAP ARP EXAMPLE.png[ONTAP ARP Example Screen]
image:CS Ransomware Example 1.png[Ransomware Alert top section]
image:CS Ransomware Example 2.png[Ransomware Alert top section]
image:CS Ransomware Example 3.png[Ransomware Alert top section]
A high confidence banner indicates the attack has shown ransomware
behavior along with file encryption activities.
The encrypted files graph indicates the timestamp at which the volume
encryption activity was detected by the ARP solution.
== Limitations
In the case where an SVM is not monitored by Workload Security, but there
are ARP events generated by ONTAP, the events will still be received and
displayed by Workload Security. However, Forensic information related to
the alert, as well as user mapping, will not be captured or shown.
== Troubleshooting
Known problems and their resolutions are described in the following table.
[cols=2*, options="header", cols"30,70"]
|===
|Problem: | Resolution:
|Email alerts are received 24 hrs after an attack is detected. In the UI,
```

the alerts are shown 24 hrs before that when the emails are received by Cloud Insights Workload Security.

|When ONTAP sends the _Ransomware Detected_ Event to Cloud Insights
Workload Security (i.e. Workload Security), the email is sent. The Event
contains a list of attacks and its timestamps. The Workload Security UI
displays the alert timestamp of the first file attacked. ONTAP sends the
Ransomware Detected Event to Cloud Insights when a certain number of
files are encoded.

Therefore, there may be a difference between the time the alert is displayed in the UI and the time the email is sent.

|===

```
= Integration with ONTAP Access Denied
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
```

[[IDf8cef30d801f8ba172c92daa7a719c97]]

:relative path: ./

:imagesdir: {root path}{relative path}./media/

[.lead]

The ONTAP Access Denied feature uses workload analysis in NAS environments (NFS and SMB) to proactively detect and warn about failed file operations (i.e., a user trying to perform an operation for which they do not have permission). These failed file operation notifications—especially in cases of security—related failures—will further help in blocking insider attacks at early stages.

Cloud Insights Workload Security integrates with ONTAP to receive Access Denied events and provide an additional analytic and automatic response layer.

Prerequisites

- * Minimum ONTAP version: 9.13.0.
- * A Workload Security administrator must enable the Access Denied feature while adding a new collector or editing existing collector, by selecting

the Monitor Access Denied Events checkbox under Advanced Configuration.

image:WS_Access_Denied_Enable_in_Collector.png[Enable Access Denied in the
ONTAP collector Advanced Configuration]

== User permissions required

If the Data Collector is added using cluster administration credentials, no new permissions are needed.

If the Collector is added using a custom user (for example, _csuser_) with permissions given to the user, follow the steps below to give Workload Security the necessary permission to register for Access Denied events with ONTAP.

For csuser with _cluster_ credentials, execute the following commands from the ONTAP command line. Note that _csrestrole_ is custom role and _csuser_ is ontap custom user.

security login rest-role create -role csrestrole -api
/api/protocols/fpolicy -access all -vserver <cluster_name>
security login create -user-or-group-name csuser -application http
-authmethod password -role csrestrole

For csuser with _SVM_ credentials, execute the following commands from the ONTAP command line:

security login rest-role create -role csrestrole -api
/api/protocols/fpolicy -access all -vserver <svm_name>
security login create -user-or-group-name csuser -application http
-authmethod password -role csrestrole -vserver <svm name>

Read more about configuring other xref:{relative_path}task_add_collector_svm.html[ONTAP permissions].

== Access Denied events

Once events have been acquired from the ONTAP system, the Workload Security Forensics page will show Access Denied events. In addition to the

information displayed, you can view the missing user permissions for a particular operation by adding the _Desired Activity_ column to the table from the gear icon.

image:WS_Access_Denied_Example_Event_1.png[Example Access Denied Event]

```
[[IDebac69ac74b78527be3bc7fc4f9b5829]]
= Blocking User Access
:hardbreaks:
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media
```

[.lead]

Once an attack is detected, Workload Security can stop the attack by blocking user access to the file system. Access can be blocked automatically, using Automated Response Policies or manually from the alert or user details pages.

NOTE: Workload Security is not available in Cloud Insights Federal Edition.

When blocking user access, you should define a blocking time period. After the selected time period ends, user access is automatically restored. Access blocking is supported for both SMB and NFS protocols.

User is directly blocked for SMB and IP address of the host machines causing the attack will be blocked for NFS. Those machine IP addresses will be blocked from accessing any of the Storage Virtual Machines (SVMs) monitored by Workload Security.

For example, let's say Workload Security manages 10 SVMs and the Automatic Response Policy is configured for four of those SVMs. If the attack originates in one of the four SVMs, the user's access will be blocked in all 10 SVMs. A Snapshot is still taken on the originating SVM.

If there are four SVMs with one SVM configured for SMB, one configured for

NFS, and the remaining two configured for both NFS and SMB, all the SVMs will be blocked if the attack originates in any of the four SVMs.

== Prerequisites for User Access Blocking

Cluster level credentials are needed for this feature to work.

If you are using cluster administration credentials, no new permissions are needed.

If you are using a custom user (for example, _csuser_) with permissions given to the user, then follow the steps below to give permissions to Workload Security to block user.

For csuser with cluster credentials, do the following from the ONTAP command line:

security login role create -role csrole -cmddirname "vserver export-policy rule" -access all

security login role create -role csrole -cmddirname set -access all security login role create -role csrole -cmddirname "vserver cifs session" -access all

security login role create -role csrole -cmddirname "vserver services access-check authentication translate" -access all

security login role create -role csrole -cmddirname "vserver name-mapping" -access all

Be sure to review the Permissions section of the xref:{relative_path}/cloudinsights/task_add_collector_svm.html[Configuring the ONTAP SVM Data Collector] page as well.

- == How to enable the feature?
- * In Workload Security, navigate to *Workload Security > Policies > Automated Response Policies*. Choose *+Attack Policy*.
- * Select (check) Block User File Access .
- == How to set up Automatic user access blocking?
- * Create a new Attack Policy or edit an existing Attack policy.
- * Select the SVMs on which the attack policy should be monitored.
- * Click on the checkbox "Block User File Access". The feature will be enabled when this is selected.
- * Under "Time Period" select the time until which the blocking should be applied.
- * To test automatic user blocking,, you can simulate an attack via a xref:{relative_path}concept_cs_attack_simulator.html[simulated script].

- == How to know if there are blocked users in the system?
- * In the alert lists page, a banner on the top of screen will be displayed in case any user is blocked.
- * Clicking on the banner will take you to the "Users" page, where the list of blocked users can be seen.
- * In the "Users" page, there in a column named "User/IP Access". In that column, the current state of user blocking will be displayed.
- == Restrict and manage user access manually
- * You can go to the alert details or user details screen and then manually block or restore a user from those screens.
- == User Access Limitation History

In the alert details and user details page, in the user panel, you can view an audit of the user's access limitation history: Time, Action (Block, Unblock), duration, action taken by, manual/automatic, and affected IPs for NFS.

== How to disable the feature?

At any time, you can disable the feature. If there are restricted users in the system, you must restore their access first.

- * In Workload Security, navigate to *Workload Security > Policies > Automated Response Policies*. Choose *+Attack Policy*.
- * De-select (uncheck) Block User File Access .

The feature will be hidden from all pages.

== Manually Restore IPs for NFS

Use the following steps to manually restore any IPs from ONTAP if your Workload Security trial expires, or if the agent/collector is down.

. List all export policies on an SVM.

	contrail-qa-	-fas8020::> expoi	rt-policy	rule sho	ow -vserver <svm name=""></svm>	
		Policy	Rule	Access	Client	RO
	Vserver	Name	Index	Protocol	l Match	Rule
-						
	svm0	default	1	nfs3,	cloudsecure_rule,	never

			nfs4, cifs	10.11.12.13	
svm1	default	4	cifs,	0.0.0.0/0	any
			nfs		
svm2	test	1	nfs3,	cloudsecure_rule,	never
			nfs4,	10.11.12.13	
			cifs		
svm3	test	3	cifs,	0.0.0.0/0	any
			nfs,		
			flexcac	he	

⁴ entries were displayed.

. Delete the rules across all policies on the SVM which have "cloudsecure_rule" as Client Match by specifying its respective RuleIndex. Workload Security rule will usually be at 1.

contrail-qa-fas8020::*> export-policy rule delete -vserver <svm name>
-policyname * -ruleindex 1

. Ensure Workload Security rule is deleted (optional step to confirm).

contrail-qa-fas8020::*> export-policy rule show -vserver <svm name> Policy Rule Access Client RO Index Protocol Match Vserver Name Rule svm0 default 4 cifs, 0.0.0.0/0 any nfs 3 cifs, 0.0.0.0/0 svm2 test any nfs, flexcache

== Manually Restore Users for SMB

Use the following steps to manually restore any users from ONTAP if your Workload Security trial expires, or if the agent/collector is down.

You can get the list of users blocked in Workload Security from the users list page.

² entries were displayed.

- 1. Login to the ONTAP cluster (where you want to unblock users) with cluster admin credentials. (For Amazon FSx, login with FSx credentials).
- 2. Run the following command to list all users blocked by Workload Security for SMB in all SVMs:

vserver name-mapping show -direction win-unix -replacement " "

Vserver: <vservername>
Direction: win-unix

Position Hostname IP Address/Mask

1 - Pattern: CSLAB\\US040

Replacement:

2 - Pattern: CSLAB\\US030

Replacement:

2 entries were displayed.

In the above output, 2 users were blocked (US030, US040) with domain CSLAB.

3. Once we identify the position from the above output, run the following command to unblock the user:

vserver name-mapping delete -direction win-unix -position <position>

4. Confirm the users are unblocked by running the command:

vserver name-mapping show -direction win-unix -replacement " "

No entries should be displayed for the users previously blocked.

== Troubleshooting

|===

|Problem|Try This

|Some of the users are not getting restricted, though there is an attack.
|1. Make sure that the Data Collector and Agent for the SVMs are in
Running state. Workload Security won't be able to send commands if the
Data Collector and Agent are stopped.

2. This is because the user may have accessed the storage from a machine

with a new IP which has not been used before.

Restricting happens via IP address of the host through which the user is accessing the storage. Check in the UI (Alert Details > Access Limitation History for This User > Affected IPs) for the list of IP addresses which are restricted. If the user is accessing storage from a host which has an IP different from the restricted IPs, then the user will still be able to access the storage through the non-restricted IP. If the user is trying to access from the hosts whose IPs are restricted, then the storage won't be accessible.

|Manually clicking on Restrict Access gives "IP addresses of this user have already been restricted".

|The IP to be restricted is already being restricted from another user.

|Policy could not be modified. Reason: not authorized for that command. |Check if using csuser, permissions are given to the user as mentioned above.

|User (IP Address) blocking for NFS works, but for SMB / CIFS, I see an error message: "SID to DomainName transformation failed. Reason timeout: socket is not established"

|This can happen is _csuser_ does not have permission to perform ssh.

(Ensure connection at cluster level, then ensure user can perform ssh).

csuser role requires these permissions.

https://docs.netapp.com/usen/cloudinsights/cs_restrict_user_access.html#prerequisites-for-useraccess-blocking

For _csuser_ with cluster credentials, do the following from the ONTAP command line:

security login role create -role csrole -cmddirname "vserver export-policy rule" -access all

security login role create -role csrole -cmddirname set -access all security login role create -role csrole -cmddirname "vserver cifs session" -access all

security login role create -role csrole -cmddirname "vserver services access-check authentication translate" -access all

security login role create -role csrole -cmddirname "vserver name-mapping" -access all

If _csuser_ is not used and if admin user at cluster level is used, make sure that the admin user has ssh permission to ONTAP.

```
|===
[[ID707da7a26baf174db4885e609a9d402e]]
= Workload Security: Simulating an Attack
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
You can use the instructions on this page to simulate an attack for
testing or demonstrating Workload Security using the included Ransomware
Simulation script.
NOTE: Workload Security is not available in Cloud Insights Federal
Edition.
== Things to note before you begin
* The ransomware simulation script works on Linux only.
* The script is provided with the Workload Security agent installation
files. It is available on any machine that has a Workload Security agent
installed.
* You can run the script on the Workload Security agent machine itself;
there is no need to prepare another Linux machine. However, if you prefer
to run the script on another system, simply copy the script and run it
there.
== Have at least 1,000 sample files
This script should run on an SVM with a folder that has files to encrypt.
We recommend having at least 1,000 files within that folder and any sub-
folders. The files must not be empty.
Do not create the files and encrypt them using the same user. Workload
Security considers this a low-risk activity and will therefore not
generate an alert (i.e. the same user modifies files he/she/they just
created).
```

See below for instructions to xref:{relative_path}#create-files-programmatically[programmatically create non-empty files].

- == Guidelines before you run the simulator:
- . Make sure encrypted files are not empty.
- . Make sure you encrypt > 50 files. A small number of files will be ignored.
- . Do not run an attack with the same user multiple times. After a few times, Workload Security will learn this user behavior and assume it is the user's normal behavior.
- . Do not encrypt files the same user has just created. Changing a file that was just created by a user is not considered a risky activity. Instead, use files created by another user OR wait for a few hours between creating the files and encrypting them.

== Prepare the system

First, mount the target volume to machine. You can mount either an NFS mount or CIFs export.

To mount NFS export in Linux:

```
mount -t nfs -o vers=4.0 10.193.177.158:/svmvol1 /mntpt mount -t nfs -o vers=4.0 Vserver data IP>:/nfsvol /destinationlinuxfolder
```

Do not mount NFS version 4.1; it is not supported by Fpolicy.

To mount CIFs in Linux:

mount -t cifs //10.193.77.91/sharedfolderincluster /root/destinationfolder/ -o username=raisa

Next, set up a Data Collector:

- . Configure the Workload Security agent if not already done.
- . Configure SVM data collector if not already done.
- == Run the Ransomware Simulator script
- . Log in (ssh) to the Workload Security agent machine.
- . Navigate to: /opt/netapp/cloudsecure/agent/install
- . Call the simulator script without parameters to see usage:

```
/opt/netapp/cloudsecure/agent/install
 # ./ransomware simulator.sh
 Error: Invalid directory provided.
 Usage: ./ransomware_simulator.sh [-e] [-d] [-i <input_directory>]
        -e to encrypt files (default)
        -d to restore files
        -i <input directory> - Files under the directory to be encrypted
 Encrypt command example: ./ransomware simulator.sh -e -i
/mnt/audit/reports/
 Decrypt command example: ./ransomware simulator.sh -d -i
/mnt/audit/reports/
== Encrypt your test files
To encrypt the files, run the following command:
 # ./ransomware simulator.sh -e -i /root/for/
Encryption key is saved in /opt/netapp/cloudsecure/cloudsecure-agent-
1.251.0/install/encryption-key,
 which can be used for restoring the files.
Encrypted /root/for/File000.txt
Encrypted /root/for/File001.txt
Encrypted /root/for/File002.txt
== Restore files
To decrypt, run the following command:
 [root@scspa2527575001 install]# ./ransomware simulator.sh -d -i
/root/for/
 File /root/for/File000.txt is restored.
 File /root/for/File001.txt is restored.
File /root/for/File002.txt is restored.
== Run the script multiple times
After generating a ransomware attack for a user, switch to another user in
order to generate an additional attack.
Workload Security learns user behavior and will not alert on repeated
ransomware attacks within a short duration for the same user.
```

== Create files programmatically Before creating the files, you must first stop or pause the data collector processing. Perform the steps below before you add the data collector to the Agent. If you have already added the data collector, just edit the data collector, enter an invalid password, and save it. This will temporarily put the data collector in error state. NOTE: Be sure you note the original password! NOTE: The recommended option is to xref:{relative path}task add collector svm.html#play-pause-datacollector[pause the collector] before creating your files.] Before running the simulation, you must first add files to be encrypted. You can either manually copy the files to be encrypted into the target folder, or use a script (see the example below) to programmatically create the files. Whichever method you use, copy at least 1,000 files. If you choose to programmatically create the files, do the following: . Log into the Agent box. . Mount an NFS export from the SVM of the filer to the Agent machine. Cd to that folder. . In that folder create a file named createfiles.sh . Copy the following lines to that file. for i in {000..1000} echo hello > "File\${i}.txt" done echo 3 > /proc/sys/vm/drop caches; sync . Save the file. . Ensure execute permission on the file: chmod 777 ./createfiles.sh . Execute the script: ./createfiles.sh 1000 files will be created in the current folder. . Re-enable the data collector

If you disabled the data collector in step 1, edit the data collector, enter the correct password, and save. Make sure that the data collector is back in running state. . If you paused the collector before following these steps, be sure to xref:{relative path}task add collector svm.html#play-pause-datacollector[resume the collector]. [[IDe2ca62e81e2006ec9e7b65ea578ad8ed]] = Configuring Email Notifications for Alerts, Warnings, and Agent/Data Source Collector health :toc: macro :hardbreaks: :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead] To configure Workload Security alert recipients, click on *Admin > Notifications* and enter an email addresses in the appropriate section(s)

for each recipient.

NOTE: Workload Security is not available in Cloud Insights Federal Edition.

== Potential Attack Alerts and Warnings

To send Potential Attack alert notifications, enter the recipients' email addresses in the Send Potential Attack Alerts section. Email notifications are sent to the alert recipient list for every action on the alert.

To send _Warning _ notifications, enter the recipients' email addresses in the Send Warning Alerts section.

== Agent and Data Collector Health monitoring

You can monitor the health of Agents and Data Sources through

notifications.

In order to receive notifications in the event that an Agent or Data Source collector is not functioning, enter the email addresses of the recipients in the Data Collection Health Alerts section.

Keep the following in mind:

- * Health alerts will be sent only after the agent/collector stops reporting for at least one hour.
- * Only one email notification is sent to the intended recipients in a given 24 hour period, even If the Agent or Data collector is disconnected for a longer duration.
- * In case of an Agent failure, one alert will be sent (not one per collector). The email will include a list of all impacted SVMs.
- * Active directory collection failure is reported as a warning; it does not impact Ransomware detection.
- * The Getting Started setup list now includes a new _Configure email notifications phase.
- == Receiving Agent And Data Collector Upgrade Notifications
- * Enter the email ID(s) in the "Data Collection Health Alerts".
- * The "Enable upgrade notifications" check box becomes enabled.
- * Agent and Data Collector upgrade email notifications are sent to the email IDs one day in advance of the planned upgrade.

== Troubleshooting

|===

|*Problem:* | *Try this:*

|Email IDs are present in the "Data Collector Health Alerts", however I am not receiving notifications.

|Notification emails are sent from the NetApp Cloud Insights domain, i.e from _accounts@service.cloudinsights.netapp.com_. Some companies block incoming emails if they are from an external domain. Ensure that external notifications from NetApp Cloud Insights domains are whitelisted.

|===

[[ID60bf036011375aee49006d9fcfd603be]]

= Workload Security API

:hardbreaks:

:icons: font

:linkattrs:

:relative path: ./

:imagesdir: {root path}{relative path}./media/

[.lead]

The Workload Security API enables NetApp customers and independent software vendors (ISVs) to integrate Workload Security with other applications, such as CMDB's or other ticketing systems.

NOTE: Workload Security is not available in Cloud Insights Federal Edition.

Requirements for API Access:

- * An API Access Token model is used to grant access.
- * API Token management is performed by Workload Security users with the Administrator role.

== API Documentation (Swagger)

The latest API information is found by logging in to Workload Security and navigating to *Admin > API Access*. Click the *API Documentation* link. The API Documentation is Swagger-based, which provides a brief description and usage information for the API and allows you to try it out in your environment.

== API Access Tokens

Before using the Workload Security API, you must create one or more *API Access Tokens*. Access tokens grant read permissions. You can also set the expiration for each access token.

To create an Access Token:

- * Click *Admin > API Access*
- * Click *+API Access Token*
- * Enter *Token Name*
- * Specify *Token Expiration*

NOTE: Your token will only be available for copying to the clipboard and saving during the creation process. Tokens can not be retrieved after they are created, so it is highly recommended to copy the token and save it in a secure location. You will be prompted to click the Copy API Access Token button before you can close the token creation screen.

You can disable, enable, and revoke tokens. Tokens that are disabled can be enabled.

Tokens grant general purpose access to APIs from a customer perspective, managing access to APIs in the scope of their own environment. The application receives an Access Token after a user successfully authenticates and authorizes access, then passes the Access Token as a credential when it calls the target API. The passed token informs the API that the bearer of the token has been authorized to access the API and perform specific actions based on the scope that was granted during authorization. The HTTP header where the Access Token is passed is *X-CloudInsights-ApiKey:* For example, use the following to retrieve storages assets: curl https://<tenant host name>/rest/v1/cloudsecure/activities -H 'X-CloudInsights-ApiKey: <API Access Token>' Where <API Access Token> is the token you saved during API access key creation. Detailed information can be found in the API Documentation link under *Admin > API Access*. :leveloffset: -1 [[ID5aef73fba8fa9bd6dd27216fd763b4a7]] = Active IQ :toc: macro :hardbreaks: :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead]

NetApp link:https://www.netapp.com/us/products/data-infrastructure-management/active-iq.aspx[Active IQ] provides a series of visualizations, analytics, and other support-related services to NetApp customers for their hardware / software systems. The data reported by Active IQ can enhance troubleshooting of system problems and also provide insight into

optimization and predictive analysis related to your devices.

NOTE: ActiveIQ is not available in Cloud Insights Federal Edition.

Cloud Insights collects the *Risks* for any NetApp Clustered Data ONTAP storage system that is monitored and reported by Active IQ. Risks reported for the storage systems are collected automatically by Cloud Insights as part of its data collection from those devices. You must add the appropriate data collector to Cloud Insights to collect Active IQ risk information.

Cloud Insights will not show risk data for ONTAP systems that are not monitored and reported by Active IQ.

The risks reported are shown in Cloud Insights on the _storage_ and _storage node_ asset landing pages, in the "Risks" table. The table shows Risk Detail, Category of risk, and Potential Impact of the risk, and also provides a link to the Active IQ page summarizing all risks for the storage node (NetApp Support account sign-in required).

image:AIQ Risks_Table_Example.png[Active IQ Risks Table]

A count of reported risks is also shown in the landing page's Summary widget, with a link to the appropriate Active IQ page. On a _storage_ landing page, the count is a sum of risks from all underlying storage nodes.

image:AIQ Summary Example.png[Storage Page Summary]

== Opening the Active IQ page

When clicking on the link to an Active IQ page, if you are not currently signed in to your Active IQ account, you must perform the following steps to view the Active IQ page for the storage node.

- . In the Cloud Insights Summary widget or Risks table, click the "Active IQ' link.
- . Sign in to your NetApp Support account. You are taken directly to the storage node page in Active IQ.
- == Querying for Risks

In Cloud Insights, you can add the *monitoring.count* column to a storage or storage node query. If the returned result includes Active IQ-Monitored storage systems, the monitoring.count column will display the number of

```
risks for the storage system or node.
== Dashboards
You can build widgets (e.g. pie chart, table widget, bar, column, scatter
plot, and single value widgets) in order to visualize object risks for
storage and storage nodes for NetApp Clustered Data ONTAP systems
monitored by Active IQ. "Object Risks" can be selected as a column or
metric in these widgets where Storage or Storage Node is the object of
focus.
//Additionally, you can filter on "Object Risks" in widgets or queries.
image:ObjectRiskWidgets.png[Object Risks in Widgets]
= Troubleshooting
:leveloffset: +1
[[ID543ac597a1eeec30e38b297f8ecbd1bc]]
= Troubleshooting General Cloud Insights Problems
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Here you will find suggestions for troubleshooting Cloud insights.
See also
xref:{relative path}task troubleshooting linux acquisition unit problems.h
tml[Troubleshooting Linux Acquisition Unit Problems] and
xref:{relative path}task troubleshooting windows acquisition unit problems
.html[Troubleshooting Windows Acquisition Unit Problems].
```

== Login issues

|===

|*Problem:* | *Try this:*

|Cloud Insights logs out every 5 minutes |Enable third-party acceptance for the necessary NetApp and authO cookies.

Example:

In Chrome, enter "chrome://settings/cookies" in the browser URL.

Select the "Allow all cookies" option.

OR

Select "Block third-party cookies" and add exceptions for [*.]auth0.com and [*.]netapp.com.

Note: Make sure to select the "Including third-party cookies on this site" option when creating an exception.

|I have a Cloud Central account but am unable to login to Cloud Central.
|Open a ticket from https://mysupport.netapp.com/site/help. Select
category "cloud.netapp.com > Account/Login issues" or "cloud.netapp.com >
Federation issues". This is specifically for Cloud Central issues or
questions.

For all other Cloud Insights technical support issues, contact xref:{relative path}concept requesting support.html[NetApp support].

|I got invited to Cloud Insights but I get a "not authorized" message. |Verify that you have signed up for a Cloud Central account, or that your organization uses SSO login with Cloud Central.

Verify your Cloud Central profile email address matches email address shown in your Cloud Insights welcome email. If the email does not match, request a new invitation with the correct email address.

|I logged out from Cloud Central or Cloud Secure and was automatically logged out from Cloud Insights. |Single Sign-On (SSO) across NetApp Cloud logs out all Cloud Insights, Cloud Secure, and Reporting sessions. If you have access to multiple Cloud Insights accounts, logging out from any one logs out all active sessions. Log back in to access your account.

|I was automatically logged out after several days. |NetApp Cloud accounts require reauthentication every few days (current Cloud Central setting is 7 days). Log back in to access your account.

|I receive an error message "no longer authorized to login".

|Contact your account administrator to verify access to Cloud Insights. Verify your Cloud Central profile email address matches email address shown in your Cloud Insights welcome email

|Other login errors

|Try incognito mode in Chrome, or clear browser history, cookies, and cache.

Try with a different browser profile (i.e. Chrome - add Person).

|===

== Other Issues

|===

|*Question:* | *Answer:*

|My Qtree hard quotas are showing correctly in queries, but my soft quotas are showing as the total capacity of the volume. Is that correct?
|Only hard quotas--either manually set or set through Trident--will show as the set quotas; if no hard quotas are specified, the Qtree capacity will be the internal volume capacity.

|I have both a soft and a hard quota manually set in the same Qtree, but the total capacity showing is the hard quota; is that correct?
|Yes, if a hard quota is specified, that will be shown as the total capacity.

|When entering a Cognos report schedule time, sometimes I end up with an extra "m" in the schedule time. For example, if I enter the time as "02:15 PM", it may add an extra character: "02:15 PMM" (or PMm). When I click outside, it changes it to "2:15 AM".

I am able to save the report, but when I re-open the saved report, the schedule time appears as AM (i.e. morning), regardless of whether I entered AM or PM in the schedule time.

|Re-enter the schedule time, being careful not to enter the full "AM" or "PM" characters; it is enough to type "A" for "AM or "P" for "PM". If you are not shown the extra character, the schedule time will be set correctly.

|===

== Resources

Additional troubleshooting tips may be found in the link:https://kb.netapp.com/Advice and Troubleshooting/Cloud Services/Cloud Insights[NetApp Knowledgebase] (support sign-in required). Additional support information may be found from the Cloud Insights xref:{relative path}concept requesting support.html[Support] page. If you have an active Cloud Insights subscription you can use these support options: link:https://www.netapp.com/us/contact-us/support.aspx[Phone] link:https://mysupport.netapp.com/site/cases/mine/create?serialNumber=9500 1014387268156333[Support Ticket] For more information, see the https://docs.netapp.com/usen/cloudinsights/concept requesting support.html[Cloud Insights Support Documentation]. [[IDfdf11c7d87710a5e7a48c4a549f17c21]] = Troubleshooting Acquisition Unit Problems on Linux :toc: macro :hardbreaks: :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead] Here you will find suggestions for troubleshooting problems with Acquisition Units on a Linux server. |=== |*Problem:* | *Try this:* |AU status on the *Observability > Collectors* page in the *Acquisition Units* tab displays "Certificate Expired" or "Certificate Revoked" . |Click on the menu to the right of the AU and select *Restore Connection*. Follow the instructions to restore your Acquisition Unit: 1. Stop the Acquisition Unit (AU) service. You can click the Copy Stop Command button to quickly copy the command to the clipboard, then paste this command into a command prompt on the acquisition unit machine.

- 2. Create a file named "token" in the /var/lib/netapp/cloudinsights/acq/conf folder on the AU.
- 3. Click the _Copy Token_ button, and paste this token into the file you created.
- 4. Restart the AU service. Click the _Copy Restart Command_ button, and paste the command into a command prompt on the AU.

|Permission denied when starting the Acquisition Unit Server Service|When the AU is installed on SELINUX, SE should be set to _permissive_ mode. _Enforcing_ mode is not supported. After setting SELINUX to permissive mode, restart the AU service.

link:https://kb.netapp.com/Advice_and_Troubleshooting/Cloud_Services/Cloud
_Insights/Permission_denied_when_starting_the_Cloud_Insight_Acquisition_Un
it Server Service[Learn more].

|Server Requirements not met | Ensure that your Acquisition Unit server or VM meets

xref:{relative_path}concept_acquisition_unit_requirements.html[requirement
s]

|Network Requirements not met |Ensure that your Acquisition Unit server/VM can access your Cloud Insights environment (<environmentname>.c01.cloudinsights.netapp.com) through SSL connection over port 443.
Try the following commands:

```
_ping <environment-name>.c01.cloudinsights.netapp.com_
_traceroute <environment-name>.c01.cloudinsights.netapp.com_
_curl \https://<environment-name>.c01.cloudinsights.netapp.com_
_wget \https://<environment-name>.c01.cloudinsights.netapp.com_
```

|Proxy Server not configured properly | Verify your proxy settings, and uninstall/re-install the Acquisition Unit software if necessary to enter the correct proxy settings.

- 1. Try "curl". Refer to "man curl" information/documentation regarding proxies: --preproxy, --proxy-* (that's a wildcard "*" because curl supports many proxy settings).
- 2. Try "wget". Check documentation for proxy options.

|Acquisition unit installation failed in Cloud insights with credential errors while starting acquisition service (and visible in the acq.log).|This can be caused by the inclusion of special characters in the proxy credentials. Uninstall the AU (sudo cloudinsights-uninstall.sh)

and reinstall without using special characters.

|Linux: missing library / file not found| Ensure that your Linux
Acquisition Unit server/VM has all necessary libraries. For example, you
must have the _unzip_ library installed on the server. To install the
unzip library, run the command _*sudo yum install unzip*_ before running
the Acquisition Unit install script

|Permission issues| Be sure you are logged in as a user with _sudo_ permissions

|Acquisition Not Running:

| Gather the acq.log from /opt/netapp/cloudinsights/acq/logs (Linux)
Restart the Acquisition Service: sudo cloudinsights-service.sh restart
acquisition

|Data Collection Issues:

|Send an Error Report from the Data Collector landing page by clicking the "Send Error Report" button

|Status: Heartbeat Failed

|The Acquisition Unit (AU) sends a heartbeat to Cloud Insights every 60 seconds to renew its lease. If the heartbeat call fails due to network issue or unresponsive Cloud Insights, the AU's lease time isn't updated. When the AU's lease time expires, Cloud Insights shows a status of "Heartbeat Failed".

Troubleshoot steps:

Check the network connection between the Acquisition Unit sever and CloudInsights.

Check whether the Acquisition Unit service is running. If the service is not running, start the service.

Check the Acquisition Unit log (/var/log/netapp/cloudinsights/acq/acq.log) to see whether there are any errors.

|I'm seeing a "Heartbeat Error: message

|This error can occur if there is a network interruption that causes communication between the Acquisition Unit and the Cloud Insights environment to be interrupted for more than one minute. Verify the connection between the AU and Cloud Insights is stable and active.

|===

////

Moving Data Collectors to Different Acquisition Units:

• What is the proper process for a customer to do this? They have to retype their passwords or go into the security admin tool to transfer the keys?

////

== Considerations about Proxies and Firewalls

If your organization requires proxy usage for internet access, you may need to understand your organization's proxy behavior and seek certain exceptions for Cloud Insights to work. Keep the following in mind:

* First, does your organization block access by default, and only allow access to specific web sites/domains by exception? If so, you will need to get the following domain added to the exception list:

+

*.cloudinsights.netapp.com

+

Your Cloud Insights Acquisition Unit, as well as your interactions in a web browser with Cloud Insights, will all go to hosts with that domain name.

* Second, some proxies attempt to perform TLS/SSL inspection by impersonating Cloud Insights web sites with digital certificates not generated from NetApp. The Cloud Insights Acquisition Unit's security model is fundamentally incompatible with these technologies. You would also need the above domain name excepted from this functionality in order for the Cloud Insights Acquisition Unit to successfully login to Cloud Insights and facilitate data discovery.

In case where the proxy is set up for traffic inspection, the Cloud Insights environment must be added to an exception list in the proxy configuration. The format and setup of this exception list varies according to your proxy environment and tools, but in general you must add the URLs of the Cloud Insights servers to this exception list in order to allow the AU to properly communicate with those servers.

The simplest way to do this is to add the Cloud Insights domain itself to the exception list:

*.cloudinsights.netapp.com

In the case where the proxy is not set up for traffic inspection, an exception list may or may not be required. If you are unsure whether you need to add Cloud Insights to an exception list, or if you experience difficulties installing or running Cloud Insights due to proxy and/or firewall configuration, talk to your proxy administration team to set up

```
the proxy's handling of SSL interception.
=== Viewing Proxy endpoints
You can view your proxy endpoints by clicking the *Proxy Settings* link
when choosing a data collector during onboarding, or the link under Proxy
Settings on the *Help > Support* page. A table like the following is
displayed. If you have Workload Security in your environment, the
configured endpoint URLs will also be displayed in this list.
image:ProxyEndpoints NewTable.png[Proxy Endpoints Table]
== Resources
Additional troubleshooting tips may be found in the
link:https://kb.netapp.com/Advice and Troubleshooting/Cloud Services/Cloud
Insights[NetApp Knowledgebase] (support sign-in required).
Additional support information may be found from the Cloud Insights
xref:{relative path}concept requesting support.html[Support] page.
[[ID075995600dd2e8a0033f536133202f10]]
= Troubleshooting Acquisition Unit Problems on Windows
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Here you will find suggestions for troubleshooting problems with
Acquisition Units on a Windows server.
|===
|*Problem:* | *Try this:*
|AU status on the *Observability > Collectors* page in the *Acquisition
Units* tab displays "Certificate Expired" or "Certificate Revoked" .
|Click on the menu to the right of the AU and select *Restore Connection*.
Follow the instructions to restore your Acquisition Unit:
1. Stop the Acquisition Unit (AU) service. You can click the Copy Stop
Command button to quickly copy the command to the clipboard, then paste
```

this command into a command prompt on the acquisition unit machine.

- 2. Create a file named "token" in the _c:\Program Files\Cloud Insights\Acquisition Unit\conf\ folder on the AU.
- 3. Click the _Copy Token_ button, and paste this token into the file you created.
- 4. Restart the AU service. Click the _Copy Restart Command_ button, and paste the command into a command prompt on the AU.

|Server Requirements not met | Ensure that your Acquisition Unit server or VM meets

xref:{relative_path}concept_acquisition_unit_requirements.html[requirement
s]

|Network Requirements not met |Ensure that your Acquisition Unit server/VM can access your Cloud Insights environment (<environment-name>.c01.cloudinsights.netapp.com) through SSL connection over port 443. Try the following commands:

```
_ping <environment-name>.c01.cloudinsights.netapp.com_
_traceroute <environment-name>.c01.cloudinsights.netapp.com_
_curl \https://<environment-name>.c01.cloudinsights.netapp.com_
_wget \https://<environment-name>.c01.cloudinsights.netapp.com_
```

|Proxy Server not configured properly | Verify your proxy settings, and uninstall/re-install the Acquisition Unit software if necessary to enter the correct proxy settings.

- 1. Try "curl". Refer to "man curl" information/documentation regarding proxies: --preproxy, --proxy-* (that's a wildcard "*" because curl supports many proxy settings).
- 2. Try "wget". Check documentation for proxy options.

|Acquisition unit installation failed in Cloud insights with credential errors while starting acquisition service (and visible in the acq.log).|This can be caused by the inclusion of special characters in the proxy credentials. Uninstall the AU (_sudo cloudinsights-uninstall.sh_) and reinstall without using special characters.

|Permission issues| Be sure you are logged in as a user with administrator permissions

|Acquisition Not Running | You can find information in the acq.log in the _<install directory>\Cloud

Insights\Acquisition Unit\log_ folder.

Restart the Acquisition via Windows Services

|Data Collection Issues

|Send an Error Report from the Data Collector landing page by clicking the "Send Error Report" button

|Status: Heartbeat Failed

|The Acquisition Unit (AU) sends a heartbeat to Cloud Insights every 60 seconds to renew its lease. If the heartbeat call fails due to network issue or unresponsive Cloud Insights, the AU's lease time isn't updated. When the AU's lease time expires, Cloud Insights shows a status of "Heartbeat Failed".

Troubleshoot steps:

- * Check the network connection between the Acquisition Unit sever and CloudInsights.
- * Check whether the Acquisition Unit service is running. If the service is not running, start the service.
- * Check the Acquisition Unit log (<Install dir>:\Program Files\Cloud Insights\Acquisition Unit\log\acq.log) to see whether there are any errors.

|I'm seeing a "Heartbeat Error: message

|This error can occur if there is a network interruption that causes communication between the Acquisition Unit and the Cloud Insights environment to be interrupted for more than one minute. Verify the connection between the AU and Cloud Insights is stable and active.

|===

== Considerations about Proxies and Firewalls

If your organization requires proxy usage for internet access, you may need to understand your organization's proxy behavior and seek certain exceptions for Cloud Insights to work. Keep the following in mind:

* First, does your organization block access by default, and only allow access to specific web sites/domains by exception? If so, you will need to add the following domain to your exception list:

+

*.cloudinsights.netapp.com

+

Your Cloud Insights Acquisition Unit, as well as your interactions in a web browser with Cloud Insights, will all go to hosts with that domain

name. * Second, some proxies attempt to perform TLS/SSL inspection by impersonating Cloud Insights web sites with digital certificates not generated from NetApp. The Cloud Insights Acquisition Unit's security model is fundamentally incompatible with these technologies. You would also need the above domain name excepted from this functionality in order for the Cloud Insights Acquisition Unit to successfully login to Cloud Insights and facilitate data discovery. === Viewing Proxy endpoints You can view your proxy endpoints by clicking the *Proxy Settings* link when choosing a data collector during onboarding, or the link under Proxy Settings on the *Help > Support* page. A table like the following is displayed. If you have Workload Security in your environment, the configured endpoint URLs will also be displayed in this list. image:ProxyEndpoints NewTable.png[Proxy Endpoints Table] == Resources Additional troubleshooting tips may be found in the link:https://kb.netapp.com/Advice and Troubleshooting/Cloud Services/Cloud Insights[NetApp Knowledgebase] (support sign-in required). Additional support information may be found from the Cloud Insights xref:{relative path}concept requesting support.html[Support] page. [[IDfe412a5bb3a89a23a12c6d07cdfd1a28]] = Researching a failed data collector :toc: macro :hardbreaks: :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/

[.lead]

If a data collector has failure message and a High or Medium Impact, you need to research this problem using the data collector summary page with its linked information.

Use the following steps to determine the cause of failed data collectors. Data collector failure messages are displayed on the *Admin* menu and on the *Installed Data Collectors* page.

.Steps

- . Click *Admin* > *Data Collectors* > *Installed Data Collectors*.
- . Click the linked Name of the failing data collector to open the Summary page.
- . On the Summary page, check the Comments area to read any notes that might have been left by another engineer who might also be investigating this failure.
- . Note any performance messages.
- //. If there is a patch being applied to this data collector, click link to check the patch page to see if that has caused the problem.
- . Move your mouse pointer over the segments of the Event Timeline graph to display additional information.
- . Select an error message for a Device and displayed below the Event Timeline and click the Error details icon that displays to the right of the message.

+

The Error details include the text of the error message, most likely causes, information in use, and suggestions of what can be tried to correct the problem.

- . In the Devices Reported By This Data Collector area, you might filter the list to display only devices of interest, and you can click the linked *Name* of a device to display the asset page for that device.
- . When you return to the data collector summary page, check the *Show Recent Changes* area at the bottom of the page to see if recent changes could have caused the problem.

:leveloffset: -1

[[ID04af7f0af35c7de9051bea59999ac866]]
= Cloud Insights Data Collector Support Matrix
:toc: macro
:toc-title: Cloud Insights Data Collectors
:hardbreaks:
:toclevels: 1
:icons: font

```
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The Data Collector Support Matrix provides reference for Data Collectors
supported by Cloud Insights, including vendor and model information.
[#top]
toc::[]
== HP Enterprise 3PAR / Alletra 9000 / Primera StoreServ Storage
:description: Support Matrix Asciidoc for HP Enterprise 3PAR / Alletra
9000 / Primera StoreServ Storage
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|HPE Alletra 9080
HPE 3PAR 20450
HPE 3PAR 20800
HPE 3PAR 20850
HPE 3PAR 20850 R2
HPE 3PAR 7200c
HPE 3PAR 7400
HPE 3PAR 7440c
HPE 3PAR 7450c
HPE 3PAR 8200
HPE 3PAR 8400
HPE 3PAR 8440
HPE 3PAR 8450
HPE 3PAR 9450
HPE 3PAR A630
HPE 3PAR A650
HPE 3PAR A670
HP 3PAR 20800
HP 3PAR 7200
HP 3PAR 7200c
HP 3PAR 7400
HP 3PAR 7400c
HP 3PAR 7450c
HP 3PAR 8200
HP 3PAR 8400
InServ F400
```

```
InServ T400
InServ T800
InServ V400
|3.1.1 (MU1)
3.1.2 (MU3)
3.1.3 (MU1)
3.1.3 (MU2)
3.1.3 (MU3)
3.2.1 (MU3)
3.2.1 (MU5)
3.2.2
3.2.2 (MU2)
3.2.2 (MU4)
3.2.2 (MU6)
3.3.1 (MU1)
3.3.1 (MU2)
3.3.1 (MU5)
3.3.2
3.3.2 (MU1)
4.4.1 Release Type: Standard Support Release
4.5.11 Release Type: Extended Support Release
4.5.3 Release Type: Extended Support Release
4.5.7 Release Type: Extended Support Release
9.5.8 Release Type: Extended Support Release
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.117+|foundation .3+|Device Group|Name|Implemented|SSH|
|Storage Management Id|Implemented|SSH|
|Type|Gap|SSH|
.14+|Disk|Capacity (GB)|Implemented|SSH|use capacity
|Disk Id|Implemented|SSH|Uniquely identifies this disk in the array
|Location|Gap|SSH|Where this disk is physically located in the array
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Role|Implemented|SSH|
|Role Enum|Implemented|SSH|enum for disk role
|Serial Number|Implemented|SSH|
|Status|Implemented|SSH|
|Status Enum|Implemented|SSH|enum for disk status
|Type|Gap|SSH|
|Type Enum|Implemented|SSH|enum for disk type
```

```
|Vendor|Implemented|SSH|
|Vendor Id|Implemented|SSH|
.4+|ISCSI Network Portal|IP|Implemented|SSH|
|Listening Port|Implemented|SSH|
|Nic|Implemented|SSH|
|OID|Implemented|SSH|
.3+|ISCSI Network Portal Group|OID|Implemented|SSH|
|Portal Group Name|Implemented|SSH|
|Portal Group Tag|Implemented|SSH|
.3+|ISCSI Node|Node Name|Implemented|SSH|
|OID|Implemented|SSH|
|Type|Gap|SSH|
.8+|ISCSI Session|OID|Implemented|SSH|
|Initiator OID|Implemented|SSH|
|Portal Group OID|Implemented|SSH|
|Target Session Id|Implemented|SSH|
|Number Of Connections|Implemented|SSH|
|Max Connections|Implemented|SSH|
|Initiator Ips|Implemented|SSH|
|Security|Implemented|SSH|
.5+|Info|Api Name|Implemented|SSH|
|Api Version|Implemented|SSH|
|DataSource Name|Implemented|SSH|Info
|Date|Implemented|SSH|
|Originator ID|Implemented|SSH|
.12+|Storage|Display IP|Implemented|SSH|
|Failed Raw Capacity|Implemented|SSH|Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented|SSH|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Microcode Version|Implemented|SSH|
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|SSH|
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.8+|Storage Node|Memory Size|Gap|SSH|device memory in MB
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Processors Count|Implemented|SSH|device CPU
|State|Implemented|SSH|free text describing the device state
```

```
|UUID|Implemented|SSH|
|Up Time|Implemented|SSH|time in milliseconds
|Version|Implemented|SSH|software version
.24+|Storage Pool|Auto Tiering|Implemented|SSH|indicates if this
storagepool is participating in auto tiering with other pools
|Compression Enabled|Implemented|SSH|Is compression enabled on the storage
pool
|Compression Savings|Implemented|SSH|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|SSH|capacity allocated for data
|Data Used Capacity|Implemented|SSH|
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|SSH|ratio of dedupe savings in percentage
|Include In Dwh Capacity|Implemented|SSH|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|SSH|
|Other Allocated Capacity|Gap|SSH|Capacity allocated for other (not data
and not snapshot)
|Other UsedCapacity (MB)|Implemented|SSH|Any capacity other than data and
snapshot
|Physical Disk Capacity (MB)|Implemented|SSH|used as raw capacity for
storage pool
|Raid Group|Implemented|SSH|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Redundancy|Implemented|SSH|Redundancy level
|Snapshot Allocated Capacity|Gap|SSH|Allocated capacity of snapshots in MB
|Snapshot Used Capacity|Implemented|SSH|
|Storage Pool Id|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Type|Gap|SSH|
|Vendor Tier|Implemented|SSH|Vendor Specific Tier Name
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.7+|Storage Synchronization|Mode|Implemented|SSH|
|Mode Enum|Implemented|SSH|
|Source Volume|Implemented|SSH|
|State|Implemented|SSH|free text describing the device state
|State Enum|Implemented|SSH|
|Target Volume|Implemented|SSH|
|Technology|Implemented|SSH|technology which causes storage efficiency
.12+|Volume|AutoTier Policy Identifier|Implemented|SSH|Dynamic Tier Policy
```

```
identifier
|Auto Tiering|Implemented|SSH|indicates if this storagepool is
participating in auto tiering with other pools
|Capacity|Implemented|SSH|Snapshot Used capacity in MB
|Name|Implemented|SSH|
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|SSH|Redundancy level
|Storage Pool Id|Implemented|SSH|
|Thin Provisioned|Implemented|SSH|
|Type|Gap|SSH|
|UUID|Implemented|SSH|
|Used Capacity|Implemented|SSH|
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.4+|Volume Map|LUN|Implemented|SSH|Name of the backend lun
|Protocol Controller|Implemented|SSH|
|Storage Port|Implemented|SSH|
|Type|Gap|SSH|
.4+|Volume Mask|Initiator|Implemented|SSH|
|Protocol Controller|Implemented|SSH|
|Storage Port|Implemented|SSH|
|Type|Gap|SSH|
.2+|Volume Ref|Name|Implemented|SSH|
|Storage Ip|Implemented|SSH|
.4+|WWN Alias|Host Aliases|Implemented|SSH|
|Object Type|Implemented|SSH|
|Source|Implemented|SSH|
|WWN|Implemented|SSH|
.74+|performance .6+|Disk|IOps Read|Implemented|SMI-S|Number of read IOps
on the disk
| IOPs Total | Implemented | SMI-S |
| IOPs Write | Implemented | SMI-S |
|Throughput Read|Implemented|SMI-S|
|Throughput Total|Implemented|SMI-S|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|SMI-S|
.19+|Storage|Cache Hit Ratio Read|Implemented|SMI-S|
|Cache Hit Ratio Total|Implemented|SMI-S|
|Cache Hit Ratio Write|Implemented|SMI-S|
|Failed Raw Capacity|Implemented|SMI-S|
|Raw Capacity|Implemented|SMI-S|
|Spare Raw Capacity|Implemented|SMI-S|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|SMI-S|
| IOPs other | Implemented | SMI-S |
|IOps Read|Implemented|SMI-S|Number of read IOps on the disk
```

```
| IOPs Total | Implemented | SMI-S |
| IOPs Write | Implemented | SMI-S |
|Latency Read|Implemented|SMI-S|
|Latency Total|Implemented|SMI-S|
|Latency Write|Implemented|SMI-S|
|Partial Blocked Ratio|Implemented|SMI-S|
|Throughput Read|Implemented|SMI-S|
|Throughput Total|Implemented|SMI-S|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|SMI-S|
|Write Pending|Implemented|SMI-S|total write pending
.11+|Storage Node|Cache Hit Ratio Total|Implemented|SMI-S|
|IOps Read|Implemented|SMI-S|Number of read IOps on the disk
|IOPs Total|Implemented|SMI-S|
| IOPs Write | Implemented | SMI-S |
|Latency Read|Implemented|SMI-S|
|Latency Total|Implemented|SMI-S|
|Latency Write|Implemented|SMI-S|
|Throughput Read|Implemented|SMI-S|
|Throughput Total|Implemented|SMI-S|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|SMI-S|
|Utilization Total|Implemented|SMI-S|
.19+|StoragePool Disk|Capacity Provisioned|Implemented|SMI-S|
|Raw Capacity|Implemented|SMI-S|
|Total Capacity|Implemented|SMI-S|
|Used Capacity|Implemented|SMI-S|
|Over Commit Capacity Ratio|Implemented|SMI-S|Reported as a time series
|Capacity Used Ratio|Implemented|SMI-S|
|Total Data Capacity|Implemented|SMI-S|
|Data Used Capacity|Implemented|SMI-S|
| IOps Read | Implemented | SMI - S | Number of read IOps on the disk
| IOPs Total | Implemented | SMI-S |
|IOPs Write|Implemented|SMI-S|
|Other Total Capacity|Implemented|SMI-S|
|Other Used Capacity|Implemented|SMI-S|
|Snapshot Reserved Capacity|Implemented|SMI-S|
|Snapshot Used Capacity|Implemented|SMI-S|
|Snapshot Used Capacity Ratio|Implemented|SMI-S| Reported as a time series
|Throughput Read|Implemented|SMI-S|
|Throughput Total|Implemented|SMI-S|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|SMI-S|
.19+|Volume|Cache Hit Ratio Read|Implemented|SMI-S|
|Cache Hit Ratio Total|Implemented|SMI-S|
|Cache Hit Ratio Write|Implemented|SMI-S|
```

```
|Raw Capacity|Implemented|SMI-S|
|Total Capacity|Implemented|SMI-S|
|Used Capacity|Implemented|SMI-S|
|Capacity Used Ratio|Implemented|SMI-S|
|CapacityRatio Written|Implemented|SMI-S|
| IOps Read | Implemented | SMI-S | Number of read IOps on the disk
|IOPs Total|Implemented|SMI-S|
| IOPs Write | Implemented | SMI-S |
|Latency Read|Implemented|SMI-S|
|Latency Total|Implemented|SMI-S|
|Latency Write|Implemented|SMI-S|
|Partial Blocked Ratio|Implemented|SMI-S|
|Throughput Read|Implemented|SMI-S|
|Throughput Total|Implemented|SMI-S|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|SMI-S|
|Write Pending|Implemented|SMI-S|total write pending
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|3Par SMI-S
|SMI-S
|HTTP/HTTPS
15988/5989
|true
|true
|true
|true
|3Par CLI
ISSH
|SSH
122
|true
|false
|true
|true
```

```
|===
<<top, Back to Top>>
== Amazon AWS EC2
:description: Support Matrix Asciidoc for Amazon AWS EC2
Models and versions supported by this data collector:
|===
<.<|API versions
12014-10-01
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.56+|foundation .7+|Data Store|Capacity|Implemented|HTTPS|Snapshot Used
capacity in MB
|MOID|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
| Provisioned Capacity | Implemented | HTTPS |
|Virtual Center Ip|Implemented|HTTPS|
|subscription Id|Implemented|HTTPS|
.6+|Server|Cluster|Implemented|HTTPS|Cluster name
|DataCenter Name|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
|MOID|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Virtual Center Ip|Implemented|HTTPS|
.8+|Virtual Disk|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|DataStore OID|Implemented|HTTPS|
| Is Chargeable | Implemented | HTTPS |
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
| Is Snapshot | Implemented | HTTPS |
|subscription Id|Implemented|HTTPS|
.20+|VirtualMachine|Dns Name|Implemented|HTTPS|
|Guest State|Implemented|HTTPS|
|DataStore OID|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
```

```
|IPs|Implemented|HTTPS|
|MOID|Implemented|HTTPS|
|Memory|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|OS|Implemented|HTTPS|
|Power State|Implemented|HTTPS|
|State Change Time|Implemented|HTTPS|
|Processors|Implemented|HTTPS|
|Provisioned Capacity|Implemented|HTTPS|
|Instance Type|Implemented|HTTPS|
|Launch Time|Implemented|HTTPS|
|LifeCycle|Implemented|HTTPS|
|public Ips|Implemented|HTTPS|
|Security Groups|Implemented|HTTPS|
|subscription Id|Implemented|HTTPS|
.3+|VirtualMachine Disk|OID|Implemented|HTTPS|
|VirtualDisk OID|Implemented|HTTPS|
|VirtualMachine OID|Implemented|HTTPS|
.5+|Host|Host OS|Implemented|HTTPS|
|IPs|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.7+|Info|Api Description|Implemented|HTTPS|
|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.30+|performance .3+|Data Store|Capacity Provisioned|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio|Implemented|HTTPS|Reported as a time series
.10+|Virtual Disk|Total Capacity|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.17+|vm|Total Capacity|Implemented|HTTPS|
```

```
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
|Total CPU Utilization|Implemented|HTTPS|
| IOps Read | Implemented | HTTPS | Number of read IOps on the disk
|diskIops.total|Implemented|HTTPS|
|Disk IOPs write|Implemented|HTTPS|
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Disk Throughput Read|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|total disk throughput read
|Disk Throughput Write|Implemented|HTTPS|
| IP Throughput Read | Implemented | HTTPS |
|Throughput total|Implemented|HTTPS|IP throughput total
|ipThroughput.write|Implemented|HTTPS|
|Total Memory Utilization|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|EC2 API
|HTTPS
|HTTPS
|443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Amazon AWS S3
:description: Support Matrix Asciidoc for Amazon AWS S3
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
```

```
IS3
12010-08-01
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.40+|foundation .7+|Info|Api Description|Implemented|HTTPS|
|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.10+|Internal Volume|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on
the storage pool
|Internal Volume Id|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
.3+|QTree|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Type|Gap|HTTPS|
.10+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
```

```
.10+|Storage Pool|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ
to cotnrol which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.4+|performance .4+|Internal Volume|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
|Total Objects|Implemented|HTTPS|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|S3 API
| HTTPS
|HTTPS
|443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Microsoft Azure NetApp Files
:description: Support Matrix Asciidoc for Microsoft Azure NetApp Files
```

```
Models and versions supported by this data collector:
|===
<.<|API versions <.<|Models
|2019-06-01
|Azure NetApp Files
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.69+|foundation .5+|File Share|Is InternalVolume|Implemented|HTTPS|whether
the file share represents an internal volume (netapp volume) or is it a
qtree/folder within the internal volume
| Is Shared | Implemented | HTTPS | whether this fileShare has any shares
associated with it
|Name|Implemented|HTTPS|
|Path|Implemented|HTTPS|path of the fileShare
|Qtree Id|Implemented|HTTPS|unique id of the qtree
.4+|Info|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
.18+|Internal Volume|Data Allocated Capacity|Gap|HTTPS|capacity allocated
for data
|Data Used Capacity|Implemented|HTTPS|
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Internal Volume Id|Implemented|HTTPS|
|Last Snapshot Time|Implemented|HTTPS|time of last snapshot
|Name|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Snapshot Count|Implemented|HTTPS|Number of snapshots on the internal
volumes
|Snapshot Used Capacity|Implemented|HTTPS|
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Total Used Capacity (MB)|Implemented|HTTPS|place holder for the used
capacity as read from the device
```

```
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
.6+|QTree|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Quota HardCapacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk
space, allowed for the quota target
|Security Style|Implemented|HTTPS|Security style of the directory: unix,
ntfs, or mixed
|Status|Implemented|HTTPS|
|Type|Gap|HTTPS|
.6+|Quota|Hard Capacity Limit (MB)|Implemented|HTTPS|max amount of disk
space, allowed for the quota target (Hard limit)
|Internal Volume Id|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Quota Id|Implemented|HTTPS|unique id of the quota
|Type|Gap|HTTPS|
|Used Capacity|Implemented|HTTPS|
.3+|Share|IP Interfaces|Implemented|HTTPS|comma separated list of IP
addresses on which this share is exposed
|Name|Implemented|HTTPS|
|Protocol|Implemented|HTTPS|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|HTTPS|
|Permission|Implemented|HTTPS|Permissions for this particular share
.11+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.14+|Storage Pool|Data Allocated Capacity|Gap|HTTPS|capacity allocated for
data
|Data Used Capacity|Implemented|HTTPS|
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
```

```
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.23+|performance .17+|Internal Volume|Latency Total|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
|Latency Read|Implemented||
|IOPs other|Implemented||
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Write|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
| IOPs Total | Implemented | |
|Latency Write|Implemented||
|Total Capacity|Implemented||
|Used Capacity|Implemented||
|Snapshot Used Capacity Ratio|Implemented|| Reported as a time series
|Capacity Used Ratio|Implemented||
|Total Data Capacity|Implemented||
|Data Used Capacity|Implemented||
|Snapshot Used Capacity|Implemented||
.6+|StoragePool Disk|IOps Read|Implemented||Number of read IOps on the
disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
```

```
ports)
|Azure Netapp Files REST API
|HTTPS
|HTTPS
1443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Brocade Fibre Channel Switches
:description: Support Matrix Asciidoc for Brocade Fibre Channel Switches
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
|178.0
183.0
Brocade 200E
Brocade 300E
Brocade 3900
Brocade 4024 Embedded
Brocade 48000
Brocade 5000
Brocade 5100
Brocade 5300
Brocade 5480 Embedded
Brocade 6505
Brocade 6510
Brocade 6520
Brocade 6548
Brocade 7800
Brocade 7840
Brocade DCX
Brocade DCX-4S Backbone
Brocade DCX8510-4
Brocade DCX8510-8
Brocade G610
Brocade G620
```

```
Brocade G630
Brocade G720
Brocade M5424 Embedded
Brocade X6-4
Brocade X6-8
Brocade X7-4
Brocade X7-8
|v5.3.2c
v6.2.1b
v6.2.2g
v6.3.2
v6.4.1a
v6.4.2
v6.4.2a
v7.0.0
v7.0.1b
v7.1.0c
v7.3.0c
v7.3.1d
v7.4.1d
v7.4.1f
v7.4.2a
v7.4.2c
v7.4.2d
v7.4.2g
v7.4.2g_cvr_824494_01
v7.4.2h
v7.4.2j1
v8.0.2a
v8.0.2c
v8.0.2d
v8.1.2g
v8.1.2j
v8.1.2k
v8.2.0
v8.2.0b
v8.2.1c
v8.2.1d
v8.2.2a
v8.2.2b
v8.2.2c
v8.2.2d
v8.2.2d4
v8.2.3
v8.2.3a
v8.2.3a1
```

```
v8.2.3b
v8.2.3c
v8.2.3c1
v9.0.0b
v9.0.1a
v9.0.1b4
v9.0.1c
v9.0.1d
v9.0.1e
v9.0.1e1
v9.1.0b
v9.1.1
v9.1.1 01
v9.1.1b
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.75+|foundation .4+|FC Name Server Entry|FC ID|Implemented|SSH|
|Nx Port WWN|Implemented|SSH|
|Physica lPort WWN|Implemented|SSH|
|Switch Port WWN|Implemented|SSH|
.4+|Fabric|Name|Implemented|Manual Entry|
|VSAN Enabled|Implemented|SSH|
|VSANId|Implemented|SSH|
|WWN|Implemented|SSH|
.2+|IVR Physical Fabric|IVR Chassis WWNs|Implemented|SSH|Comma seperated
list of IVR enabled chassis WWNs
|Lowest IVRChassis WWN|Implemented|SSH|identifier of the IVR fabric
.4+|Info|DataSource Name|Implemented|SSH|Info
|Date|Implemented|SSH|
|Originator ID|Implemented|SSH|
|Originator Key|Implemented|SSH|
.13+|Logical Switch|Chassis WWN|Implemented|SSH|
|Domain Id|Implemented|SSH|
|Firmware Version|Implemented|SSH|
|IP|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Model|Implemented|SSH|
|Name|Implemented|Manual Entry|
|Serial Number|Implemented|SSH|
|Switch Role|Implemented|SSH|
|Switch State|Implemented|SSH|
```

```
|Switch Status|Implemented|SSH|
|Type|Gap|SSH|
|WWN|Implemented|SSH|
.16+|Port|Blade|Implemented|SSH|
| FC4 Protocol | Implemented | SSH |
|GBIC Type|Implemented|SSH|
|Generated|Implemented|SSH|
|Name|Implemented|Manual Entry|
|Node WWN|Implemented|SSH|Mandatory to report with PortId if WWN is not
present
|Port ID|Implemented|SSH|
|Port Number|Implemented|SSH|
|Port Speed|Implemented|SSH|
|Port State|Implemented|SSH|
|Port Status|Implemented|SSH|
|Port Type|Implemented|SSH|
|Raw Port Status|Implemented|SSH|
|Raw Speed GigaBits|Implemented|SSH|
|Unknown Connectivity|Implemented|SSH|
|WWN|Implemented|SSH|
.14+|Switch|Domain Id|Implemented|SSH|
|Firmware Version|Implemented|SSH|
|IP|Implemented|SSH|
|Manage URL|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Model|Implemented|SSH|
|Name|Implemented|Manual Entry|
|Serial Number|Implemented|SSH|
|Switch Role|Implemented|SSH|
|Switch State|Implemented|SSH|
|Switch Status|Implemented|SSH|
|Type|Gap|SSH|
|VSAN Enabled|Implemented|SSH|
|WWN|Implemented|SSH|
.7+|Unknown|Driver|Implemented|SSH|
|Firmware|Implemented|SSH|
|Generated|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Model|Implemented|SSH|
|Name|Implemented|Manual Entry|
|WWN|Implemented|SSH|
.4+|WWN Alias|Host Aliases|Implemented|SSH|
|Object Type|Implemented|SSH|
|Source|Implemented|SSH|
|WWN|Implemented|SSH|
.1+|Zone|Zone Name|Implemented|SSH|
```

```
.2+|Zone Member|Type|Gap|SSH|
|WWN|Implemented|SSH|
.4+|Zoning Capabilities|Active Configuration|Implemented|SSH|
|Configuration Name|Implemented|SSH|
|Default Zoning Behavior|Implemented|SSH|
|WWN|Implemented|SSH|
.28+|performance .28+|port|BB Credit|Implemented|SNMP|
|bbCreditZero.total|Implemented|SNMP|
|BB Credit|Implemented|SNMP|
|bbCreditZeroMs|Implemented|SNMP|
|portErrors.class3Discard|Implemented|SNMP|
|portErrors.crc|Implemented|SNMP|
|Port Error|Implemented|SNMP|
|portErrors.encOut|Implemented|SNMP|
|Port Error|Implemented|SNMP|Port errors due to long frame
|Port Error|Implemented|SNMP|Port errors due to short frame
|portErrors.linkFailure|Implemented|SNMP|Port Errors link failure
|portErrors.linkResetRx|Implemented|SNMP|
|Port Error|Implemented|SNMP|Port Error due to link reset
|Port Error|Implemented|SNMP|Port errors signal loss
|Port Error|Implemented|SNMP|Port error sync loss
|Port Error|Implemented|SNMP|port errors timeout discard
|Port Error|Implemented|SNMP|Total port errors
|Traffic Frame Rate|Implemented|SNMP|
|Total Traffic Frame Rate|Implemented|SNMP|
|Traffic Frame Rate|Implemented|SNMP|
|Average Frame Size|Implemented|SNMP|Average Frame size of traffic
|TX Frames|Implemented|SNMP|traffic average frame size
|Traffic Rate|Implemented|SNMP|
|Total Traffic Rate|Implemented|SNMP|
|Traffic Rate|Implemented|SNMP|
|Traffic Utilization|Implemented|SNMP|
|Traffic Utilization|Implemented|SNMP|Total traffic utilization
|Traffic Utilization|Implemented|SNMP|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Brocade SNMP
| SNMP
```

```
|SNMPv1, SNMPv2, SNMPv3
1161
|true
|true
|true
|true
|Brocade SSH
|SSH
|SSH
122
|false
|false
|true
|true
|Data source wizard configuration
|Manual Entry
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Brocade Network Advisor HTTP
:description: Support Matrix Asciidoc for Brocade Network Advisor HTTP
Models and versions supported by this data collector:
<.<|API versions <.<|Models <.<|Firmware versions
|14.4.1
14.4.3
14.4.4
14.4.5
|Brocade 5300
Brocade 6510
Brocade 6520
```

```
Brocade 6548
Brocade DCX 8510-8
Brocade G620
DS-6620B
EMC Connectrix ED-DCX8510-8B
lv7.2.1a
v7.3.1a
v7.4.1b
v7.4.2d
v8.2.3b
v8.2.3c
v9.0.1a
v9.0.1b
v9.0.1e1
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.58+|foundation .4+|FC Name Server Entry|FC ID|Implemented|HTTP/S|
|Nx Port WWN|Implemented|HTTP/S|
|Physica lPort WWN|Implemented|HTTP/S|
|Switch Port WWN|Implemented|HTTP/S|
.4+|Fabric|Name|Implemented|HTTP/S|
|VSAN Enabled|Implemented|HTTP/S|
|VSANId|Implemented|HTTP/S|
|WWN|Implemented|HTTP/S|
.7+|Info|Api Description|Implemented|HTTP/S|
|Api Name|Implemented|HTTP/S|
|Api Version|Implemented|HTTP/S|
|DataSource Name|Implemented|HTTP/S|Info
|Date|Implemented|HTTP/S|
|Originator ID|Implemented|HTTP/S|
|Originator Key|Implemented|HTTP/S|
.15+|Port|Blade|Implemented|HTTP/S|
|FC4 Protocol|Implemented|HTTP/S|
|GBIC Type|Implemented|HTTP/S|
|Generated|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Port ID|Implemented|HTTP/S|
|Port Number|Implemented|HTTP/S|
|Port Speed|Implemented|HTTP/S|
|Port State|Implemented|HTTP/S|
|Port Status|Implemented|HTTP/S|
```

```
|Port Type|Implemented|HTTP/S|
|Raw Port Status|Implemented|HTTP/S|
|Raw Speed GigaBits|Implemented|HTTP/S|
|Unknown Connectivity|Implemented|HTTP/S|
|WWN|Implemented|HTTP/S|
.13+|Switch|Domain Id|Implemented|HTTP/S|
|Firmware Version|Implemented|HTTP/S|
|IP|Implemented|HTTP/S|
|Manage URL|Implemented|HTTP/S|
|Manufacturer|Implemented|HTTP/S|
|Model|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Serial Number|Implemented|HTTP/S|
|Switch Role|Implemented|HTTP/S|
|Switch State|Implemented|HTTP/S|
|Switch Status|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
|WWN|Implemented|HTTP/S|
.5+|Unknown|Driver|Implemented|HTTP/S|
|Firmware|Implemented|HTTP/S|
|Manufacturer|Implemented|HTTP/S|
|Model|Implemented|HTTP/S|
|WWN|Implemented|HTTP/S|
.4+|WWN Alias|Host Aliases|Implemented|HTTP/S|
|Object Type|Implemented|HTTP/S|
|Source|Implemented|HTTP/S|
|WWN|Implemented|HTTP/S|
.1+|Zone|Zone Name|Implemented|HTTP/S|
.2+|Zone Member|Type|Gap|HTTP/S|
|WWN|Implemented|HTTP/S|
.3+|Zoning Capabilities|Active Configuration|Implemented|HTTP/S|
|Configuration Name|Implemented|HTTP/S|
|WWN|Implemented|HTTP/S|
.18+|performance .18+|port|bbCreditZero.total|Implemented|HTTP/S|
|BB Credit|Implemented|HTTP/S|
|bbCreditZeroMs|Implemented|HTTP/S|
|portErrors.class3Discard|Implemented|HTTP/S|
|portErrors.crc|Implemented|HTTP/S|
|Port Error|Implemented|HTTP/S|
|Port Error|Implemented|HTTP/S|Port errors due to short frame
|portErrors.linkFailure|Implemented|HTTP/S|Port Errors link failure
|Port Error|Implemented|HTTP/S|Port errors signal loss
|Port Error|Implemented|HTTP/S|Port error sync loss
|Port Error|Implemented|HTTP/S|port errors timeout discard
|Port Error|Implemented|HTTP/S|Total port errors
|Traffic Rate|Implemented|HTTP/S|
```

```
|Total Traffic Rate|Implemented|HTTP/S|
|Traffic Rate|Implemented|HTTP/S|
|Traffic Utilization|Implemented|HTTP/S|
|Traffic Utilization|Implemented|HTTP/S|Total traffic utilization
|Traffic Utilization|Implemented|HTTP/S|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Brocade Network Advisor REST API
|HTTP/HTTPS
|HTTP/HTTPS
180/443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Brocade FOS REST
:description: Support Matrix Asciidoc for Brocade FOS REST
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|Brocade 6505
Brocade G720
Brocade X6-8
|v8.2.3c
v8.2.3c1
v9.0.1e1
v9.1.1b
|===
Products supported by this data collector:
```

```
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.72+|foundation .4+|FC Name Server Entry|FC ID|Implemented|HTTPS|
|Nx Port WWN|Implemented|HTTPS|
|Physica lPort WWN|Implemented|HTTPS|
|Switch Port WWN|Implemented|HTTPS|
.4+|Fabric|Name|Implemented|HTTPS|
| VSAN Enabled | Implemented | HTTPS |
|VSANId|Implemented|HTTPS|
|WWN|Implemented|HTTPS|
.7+|Info|Api Description|Implemented|HTTPS|
|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.12+|Logical Switch|WWN|Implemented|HTTPS|
|IP|Implemented|HTTPS|
|Firmware Version|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Switch Role|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Serial Number|Implemented|HTTPS|
|Switch State|Implemented|HTTPS|
|Domain Id|Implemented|HTTPS|
|Chassis WWN|Implemented|HTTPS|
.15+|Port|Blade|Implemented|HTTPS|
|Generated|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Node WWN|Implemented|HTTPS|Mandatory to report with PortId if WWN is not
present
|Port ID|Implemented|HTTPS|
|Port Number|Implemented|HTTPS|
|Port Speed|Implemented|HTTPS|
|Port State|Implemented|HTTPS|
|Port Status|Implemented|HTTPS|
|Port Type|Implemented|HTTPS|
|Raw Port Status|Implemented|HTTPS|
|Raw Speed GigaBits|Implemented|HTTPS|
|Unknown Connectivity|Implemented|HTTPS|
|WWN|Implemented|HTTPS|
```

```
|Description|Implemented|HTTPS|
.14+|Switch|Domain Id|Implemented|HTTPS|
|Firmware Version|Implemented|HTTPS|
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Serial Number|Implemented|HTTPS|
|Switch Role|Implemented|HTTPS|
|Switch State|Implemented|HTTPS|
|Switch Status|Implemented|HTTPS|
|Type|Gap|HTTPS|
|VSAN Enabled|Implemented|HTTPS|
|WWN|Implemented|HTTPS|
.5+|Unknown|Driver|Implemented|HTTPS|
|Firmware|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|WWN|Implemented|HTTPS|
.4+|WWN Alias|Host Aliases|Implemented|HTTPS|
|Object Type|Implemented|HTTPS|
|Source|Implemented|HTTPS|
|WWN|Implemented|HTTPS|
.1+|Zone|Zone Name|Implemented|HTTPS|
.2+|Zone Member|Type|Gap|HTTPS|
|WWN|Implemented|HTTPS|
.4+|Zoning Capabilities|Active Configuration|Implemented|HTTPS|
|Configuration Name|Implemented|HTTPS|
|Default Zoning Behavior|Implemented|HTTPS|
|WWN|Implemented|HTTPS|
.27+|performance .27+|port|BB Credit|Implemented|HTTPS|
|bbCreditZero.total|Implemented|HTTPS|
|BB Credit|Implemented|HTTPS|
|bbCreditZeroMs|Implemented|HTTPS|
|portErrors.class3Discard|Implemented|HTTPS|
|portErrors.crc|Implemented|HTTPS|
|Port Error|Implemented|HTTPS|
|portErrors.encOut|Implemented|HTTPS|
|Port Error|Implemented|HTTPS|Port errors due to long frame
|Port Error|Implemented|HTTPS|Port errors due to short frame
|portErrors.linkFailure|Implemented|HTTPS|Port Errors link failure
|portErrors.linkResetRx|Implemented|HTTPS|
|Port Error|Implemented|HTTPS|Port Error due to link reset
|Port Error|Implemented|HTTPS|Port errors signal loss
|Port Error|Implemented|HTTPS|Port error sync loss
```

```
|Port Error|Implemented|HTTPS|Total port errors
|Traffic Frame Rate|Implemented|HTTPS|
|Total Traffic Frame Rate|Implemented|HTTPS|
|Traffic Frame Rate|Implemented|HTTPS|
|Average Frame Size|Implemented|HTTPS|Average Frame size of traffic
|TX Frames|Implemented|HTTPS|traffic average frame size
|Traffic Rate|Implemented|HTTPS|
|Total Traffic Rate|Implemented|HTTPS|
|Traffic Rate|Implemented|HTTPS|
|Traffic Utilization|Implemented|HTTPS|
|Traffic Utilization|Implemented|HTTPS|Total traffic utilization
|Traffic Utilization|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Brocade FOS REST API
IHTTPS
|443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Cisco MDS & Nexus Fabric Switches
:description: Support Matrix Asciidoc for Cisco MDS & Nexus Fabric
Switches
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
18978-E04
CN1610
```

```
DS-C9124-2-K9
DS-C9124-K9
DS-C9132T-K9
DS-C9134-K9
DS-C9148-16P-K9
DS-C9148-32P-K9
DS-C9148-48P-K9
DS-C9148S-K9
DS-C9148T-K9
DS-C9222I-K9
DS-C9250I-K9
DS-C9396S-K9
DS-C9396T-K9
DS-C9506
DS-C9509
DS-C9513
DS-C9706
DS-C9710
DS-C9718
DS-HP-8GFC-K9
DS-HP-FC-K9
N5K-C5548UP
N5K-C5596UP
N5K-C56128P
N5K-C5696Q
UCS-FI-6248UP
UCS-FI-6296UP
UCS-FI-6332
UCS-FI-6332-16UP
UCS-FI-6454
|3.3(1c)|
4.1(3a)
5.0(1a)
5.0(3)N2(3.11e)
5.0(3)N2(3.23o)
5.0(3)N2(4.01d)
5.0(3)N2(4.04e)
5.0(3)N2(4.13e)
5.0(3)N2(4.13i)
5.0(3) N2(4.21e)
5.0(3)N2(4.21j)
5.0(3)N2(4.21k)
5.0(3)N2(4.22c)
5.0(8)
5.2(2d)
5.2(3)N2(2.28g)
```

```
5.2(6a)
5.2(8)
5.2(8b)
5.2(8c)
5.2(8d)
5.2(8f)
5.2(8g)
5.2(8h)
5.2(8i)
6.2(1)
6.2(11)
6.2(11b)
6.2(11c)
6.2(11e)
6.2(13)
6.2(13a)
6.2(15)
6.2(17)
6.2(19)
6.2(21)
6.2(23)
6.2 (25)
6.2(27)
6.2(29)
6.2(31)
6.2(33)
6.2(5)
6.2(5a)
6.2(7)
6.2(9)
6.2(9a)
6.2(9c)
7.3(0)D1(1)
7.3(0)DY(1)
7.3(1)DY(1)
7.3(1)N1(1)
7.3(13)N1(1)
7.3(6)N1(1)
7.3(8)N1(1)
8.1(1)
8.1(1a)
8.1(1b)
8.2(1)
8.2(2)
8.3(1)
8.3(2)
```

```
8.4(1)
8.4(1a)
8.4(2)
8.4(2a)
8.4(2b)
8.4(2c)
8.4(2d)
8.4(2e)
8.4(2f)
8.5(1)
9.2(1)
9.2(1a)
9.2(2)
9.3(2)
9.3(2a)
9.3(5) I42(2a)
9.3(5) I42(2c)
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.69+|foundation .4+|FC Name Server Entry|FC ID|Implemented|SNMP|
|Nx Port WWN|Implemented|SNMP|
|Physica lPort WWN|Implemented|SNMP|
|Switch Port WWN|Implemented|SNMP|
.4+|Fabric|Name|Implemented|SNMP|
|VSAN Enabled|Implemented|SNMP|
|VSANId|Implemented|SNMP|
|WWN|Implemented|SNMP|
.2+|IVR Physical Fabric|IVR Chassis WWNs|Implemented|SNMP|Comma seperated
list of IVR enabled chassis WWNs
|Lowest IVRChassis WWN|Implemented|SNMP|identifier of the IVR fabric
.4+|Info|DataSource Name|Implemented|SNMP|Info
|Date|Implemented|SNMP|
|Originator ID|Implemented|SNMP|
|Originator Key|Implemented|SNMP|
.9+|Logical Switch|Chassis WWN|Implemented|SNMP|
|Domain Id|Implemented|SNMP|
|DomainId Type|Implemented|SNMP|
|IP|Implemented|SNMP|
|Manufacturer|Implemented|SNMP|
|Priority|Implemented|SNMP|
|Switch Role|Implemented|SNMP|
```

```
|Type|Gap|SNMP|
|WWN|Implemented|SNMP|
.14+|Port|Blade|Implemented|SNMP|
|GBIC Type|Implemented|SNMP|
|Generated|Implemented|SNMP|
|Name|Implemented|SNMP|
|Port ID|Implemented|SNMP|
|Port Number|Implemented|SNMP|
|Port Speed|Implemented|SNMP|
|Port State|Implemented|SNMP|
|Port Status|Implemented|SNMP|
|Port Type|Implemented|SNMP|
|Raw Port Status|Implemented|SNMP|
|Raw Speed GigaBits|Implemented|SNMP|
|Unknown Connectivity|Implemented|SNMP|
|WWN|Implemented|SNMP|
.12+|Switch|Firmware Version|Implemented|SNMP|
|IP|Implemented|SNMP|
|Manage URL|Implemented|SNMP|
|Manufacturer|Implemented|SNMP|
|Model|Implemented|SNMP|
|Name|Implemented|SNMP|
|SANRoute Enabled|Implemented|SNMP|Indicates if this chassis is enabled
for SAN routing (IVR, etc...)
|Serial Number|Implemented|SNMP|
|Switch Status|Implemented|SNMP|
|Type|Gap|SNMP|
|VSAN Enabled|Implemented|SNMP|
|WWN|Implemented|SNMP|
.7+|Unknown|Driver|Implemented|SNMP|
|Firmware|Implemented|SNMP|
|Generated|Implemented|SNMP|
|Manufacturer|Implemented|SNMP|
|Model|Implemented|SNMP|
|Name|Implemented|SNMP|
|WWN|Implemented|SNMP|
.4+|WWN Alias|Host Aliases|Implemented|SNMP|
|Object Type|Implemented|SNMP|
|Source|Implemented|SNMP|
|WWN|Implemented|SNMP|
.2+|Zone|Zone Name|Implemented|SNMP|
|Zone Type|Implemented|SNMP|
.2+|Zone Member|Type|Gap|SNMP|
|WWN|Implemented|SNMP|
.5+|Zoning Capabilities|Active Configuration|Implemented|SNMP|
|Configuration Name|Implemented|SNMP|
```

```
|Default Zoning Behavior|Implemented|SNMP|
|Merge Control|Implemented|SNMP|
|WWN|Implemented|SNMP|
.26+|performance .26+|port|BB Credit|Implemented|SNMP|
|bbCreditZero.total|Implemented|SNMP|
|BB Credit|Implemented|SNMP|
|bbCreditZeroMs|Implemented|SNMP|
|portErrors.class3Discard|Implemented|SNMP|
|portErrors.crc|Implemented|SNMP|
|Port Error|Implemented|SNMP|Port errors due to long frame
|Port Error|Implemented|SNMP|Port errors due to short frame
|portErrors.linkFailure|Implemented|SNMP|Port Errors link failure
|portErrors.linkResetRx|Implemented|SNMP|
|Port Error|Implemented|SNMP|Port Error due to link reset
|Port Error|Implemented|SNMP|Port errors signal loss
|Port Error|Implemented|SNMP|Port error sync loss
|Port Error|Implemented|SNMP|port errors timeout discard
|Port Error|Implemented|SNMP|Total port errors
|Traffic Frame Rate|Implemented|SNMP|
|Total Traffic Frame Rate|Implemented|SNMP|
|Traffic Frame Rate|Implemented|SNMP|
|Average Frame Size|Implemented|SNMP|Average Frame size of traffic
|TX Frames|Implemented|SNMP|traffic average frame size
|Traffic Rate|Implemented|SNMP|
|Total Traffic Rate|Implemented|SNMP|
|Traffic Rate|Implemented|SNMP|
|Traffic Utilization|Implemented|SNMP|
|Traffic Utilization|Implemented|SNMP|Total traffic utilization
|Traffic Utilization|Implemented|SNMP|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Cisco SNMP
| SNMP
|SNMPv1 (Inventory only), SNMPv2, SNMPv3
|161
|true
|true
```

```
Itrue
|true
|===
<<top, Back to Top>>
== Cohesity
:description: Support Matrix Asciidoc for Cohesity
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
IC2500
C2505
C4000 Compute Node
C4600
C5036
C5066
C6025
C6035
C6055
PXG1
UCS-C240M5H10
|6.5.1f release-20210913 13f6a4bf
6.5.1f u1 release-20211027 9e4e40cb
6.6.0d u6 release-20221204 c03629f0
6.8.1 release-20220807 6c9115ef
6.8.1 u1 release-20221022 6f58ed2a
6.8.1 u2 release-20230412 5ced2ed3
6.8.1 u3 release-20230509 1e641b74
7.0 ul release-20230222 8995f044
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.66+|foundation .3+|Disk|Capacity (GB)|Implemented||use capacity
|Disk Id|Implemented||Uniquely identifies this disk in the array
|Name|Implemented||
.5+|File Share|Is InternalVolume|Implemented||whether the file share
represents an internal volume (netapp volume) or is it a qtree/folder
within the internal volume
```

```
|Is Shared|Implemented||whether this fileShare has any shares associated
with it
|Name|Implemented||
|Path|Implemented||path of the fileShare
|Qtree Id|Implemented||unique id of the qtree
.5+|Info|Api Name|Implemented||
|DataSource Name|Implemented||Info
|Date|Implemented||
|Originator ID|Implemented||
|Originator Key|Implemented||
.13+|Internal Volume|Compression Enabled|Implemented||Is compression
enabled on the storage pool
|Dedupe Enabled|Implemented||Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented||ratio of dedupe savings in percentage
|Internal Volume Id|Implemented||
|Name|Implemented||
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
|Storage Pool Id|Implemented||
|Thin Provisioned|Implemented||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented||
|Total Used Capacity|Implemented||Total capacity in MB
|Total Used Capacity (MB)|Implemented||place holder for the used capacity
as read from the device
|Type|Gap||
.3+|QTree|Name|Implemented||
|Qtree Id|Implemented||unique id of the qtree
|Type|Gap||
.3+|Share|IP Interfaces|Implemented||comma separated list of IP addresses
on which this share is exposed
|Name|Implemented||
|Protocol|Implemented||enum for share protocol
.13+|Storage|Display IP|Implemented||
|Failed Raw Capacity|Implemented||Raw capacity of failed disks (sum of
all disks that are failed)
|Family|Implemented||The storage Family could be Clariion, Symmetrix, et
al
|IP|Implemented||
|Manage URL|Implemented||
|Manufacturer|Implemented||
|Microcode Version|Implemented||
|Model|Implemented||
|Name|Implemented||
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
```

```
the array)
|Serial Number|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|Virtual|Implemented||Is this a storage virtualization device?
.5+|Storage Node|Model|Implemented||
|Name|Implemented||
|Serial Number|Implemented||
|UUID|Implemented||
|Version|Implemented||software version
.16+|Storage Pool|Compression Enabled|Implemented||Is compression enabled
on the storage pool
|Dedupe Enabled|Implemented||Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented||ratio of dedupe savings in percentage
|Include In Dwh Capacity|Implemented||A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented||
|Physical Disk Capacity (MB)|Implemented||used as raw capacity for
storage pool
|Raid Group|Implemented||indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
|Status|Implemented||
|Storage Pool Id|Implemented||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented||
|Total Used Capacity|Implemented||Total capacity in MB
|Type|Gap||
|Virtual|Implemented||Is this a storage virtualization device?
|Encrypted|Implemented||
.26+|performance .7+|Disk|IOps Read|Implemented||Number of read IOps on
the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Utilization Total|Implemented||
.3+|Internal Volume|Total Capacity|Implemented||
|Capacity Used Ratio|Implemented||
|Used Capacity|Implemented||
.0+|Qtree.10+|Storage|Latency Total|Implemented||
|IOps Read|Implemented||Number of read IOps on the disk
```

```
|Latency Read|Implemented||
| IOPs Write | Implemented | |
|Throughput Read|Implemented||
|Throughput Write|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|IOPs Total|Implemented||
|Latency Write|Implemented||
|Utilization Total|Implemented||
.6+|StoragePool Disk|IOps Read|Implemented||Number of read IOps on the
disk
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Write|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|IOPs Total|Implemented||
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Cohesity REST API
|HTTPS
|HTTPS
|443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== EMC Celerra (SSH)
:description: Support Matrix Asciidoc for EMC Celerra (SSH)
Models and versions supported by this data collector:
|===
```

```
<.<|Models <.<|Firmware versions
|NS-480FC
NSX
VG8
VNX5200
VNX5300
VNX5400
VNX5600
VNX7600
15.5.38-1
6.0.65 - 2
7.1.76 - 4
7.1.79-8
7.1.83-2
8.1.21-266
8.1.21-303
8.1.9-155
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.85+|foundation .6+|File Share|Is InternalVolume|Implemented|SSH|whether
the file share represents an internal volume (netapp volume) or is it a
qtree/folder within the internal volume
| Is Shared | Implemented | SSH | whether this fileShare has any shares
associated with it
|Name|Implemented|SSH|
|Path|Implemented|SSH|path of the fileShare
|Qtree Id|Implemented|SSH|unique id of the qtree
|Status|Implemented|SSH|
.6+|Info|Api Name|Implemented|SSH|
|Api Version|Implemented|SSH|
|DataSource Name|Implemented|SSH|Info
|Date|Implemented|SSH|
|Originator ID|Implemented|SSH|
|Originator Key|Implemented|SSH|
.21+|Internal Volume|Data Allocated Capacity|Gap|SSH|capacity allocated
for data
|Data Used Capacity|Implemented|SSH|
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|SSH|ratio of dedupe savings in percentage
|GuidKey 1|Implemented|SSH|GuidKey1 is implicit for all objects whose GUID
```

key has not changed since OCI version 7.3.5. |GuidKey 2|Implemented|SSH|GuidKey2 is implicit for all objects whose GUID key has not changed since OCI version 7.3.5. |Internal Volume Id|Implemented|SSH| |Last Snapshot Time|Implemented|SSH|time of last snapshot |Name|Implemented|SSH| |Other Allocated Capacity|Gap|SSH|Capacity allocated for other (not data and not snapshot) |Other UsedCapacity (MB)|Implemented|SSH|Any capacity other than data and snapshot |Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity to raw capacity |Snapshot Count|Implemented|SSH|Number of snapshots on the internal volumes |Storage Pool Id|Implemented|SSH| |Thin Provisioned|Implemented|SSH| |Thin Provisioning Supported|Implemented|SSH|Whether this internal volume supports thin provisioning for the volume layer on top of it |Total Allocated Capacity|Implemented|SSH| |Total Used Capacity|Implemented|SSH|Total capacity in MB |Total Used Capacity (MB)|Implemented|SSH|place holder for the used capacity as read from the device |Type|Gap|SSH| |Virtual Storage|Implemented|SSH|Owning virtual storage (vfiler) .8+|QTree|GuidKey 1|Implemented|SSH|GuidKey1 is implicit for all objects whose GUID key has not changed since OCI version 7.3.5. |GuidKey 2|Implemented|SSH|GuidKey2 is implicit for all objects whose GUID key has not changed since OCI version 7.3.5. |Name|Implemented|SSH| |Qtree Id|Implemented|SSH|unique id of the qtree |Quota HardCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk space, allowed for the quota target |Quota SoftCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk space, allowed for the quota target |Quota UsedCapacity|Implemented|SSH|Space in MB currently used |Type|Gap|SSH| .12+|Quota|GuidKey 1|Implemented|SSH|GuidKey1 is implicit for all objects whose GUID key has not changed since OCI version 7.3.5. |GuidKey 2|Implemented|SSH|GuidKey2 is implicit for all objects whose GUID key has not changed since OCI version 7.3.5. | Hard Capacity Limit (MB) | Implemented | SSH | max amount of disk space, allowed for the quota target (Hard limit) |Hard File Limit|Implemented|SSH|max number of files allowed for the quota target |Internal Volume Id|Implemented|SSH| |Qtree Id|Implemented|SSH|unique id of the qtree

```
|Quota Id|Implemented|SSH|unique id of the quota
|Soft Capacity Limit (MB)|Implemented|SSH|Maximum amount of disk space,
allowed for the quota target
|Soft File Limit|Implemented|SSH|Max number of files allowed for the quota
target
|Type|Gap|SSH|
|Used Capacity|Implemented|SSH|
|Used Files|Implemented|SSH|Number of files currently used
.3+|Share|IP Interfaces|Implemented|SSH|comma separated list of IP
addresses on which this share is exposed
|Name|Implemented|SSH|
|Protocol|Implemented|SSH|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|SSH|
|Permission|Implemented|SSH|Permissions for this particular share
.12+|Storage|Cpu Count|Implemented|SSH|Cpu Count of the storage
|Display IP|Implemented|SSH|
|Failed Raw Capacity|Implemented|SSH|Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented|SSH|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Microcode Version|Implemented|SSH|
|Model|Implemented|SSH|
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|SSH|
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.15+|Storage Pool|Data Allocated Capacity|Gap|SSH|capacity allocated for
data
|Data Used Capacity|Implemented|SSH|
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Include In Dwh Capacity|Implemented|SSH|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|SSH|
|Raid Group|Implemented|SSH|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Snapshot Allocated Capacity|Gap|SSH|Allocated capacity of snapshots in MB
|Snapshot Used Capacity|Implemented|SSH|
|Storage Pool Id|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
```

```
|Total Allocated Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Type|Gap|SSH|
|Virtual|Implemented|SSH|Is this a storage virtualization device?
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Celerra CLI
|SSH
|SSH
|true
|false
|true
|true
|===
<<top, Back to Top>>
== EMC CLARiiON (NaviCLI)
:description: Support Matrix Asciidoc for EMC CLARiiON (NaviCLI)
Models and versions supported by this data collector:
<.<|API versions <.<|Models <.<|Firmware versions
16.23
6.26
6.28
7.30
7.32
7.33
|AX4-5F8
CX3-20f
CX3-40f
CX4-480
VNX5100
```

```
VNX5200
VNX5300
VNX5400
VNX5500
VNX5600
VNX5700
VNX5800
VNX7600
VNX8000
104.28.000.5.710
04.30.000.5.525
05.32.000.5.218
05.32.000.5.219
05.32.000.5.221
05.32.000.5.225
05.32.000.5.249
05.33.000.5.074
05.33.009.5.155
05.33.009.5.184
05.33.009.5.186
05.33.009.5.218
05.33.009.5.231
05.33.009.5.236
05.33.009.5.238
05.33.009.6.305
05.33.021.5.256
05.33.021.5.266
2.23.50.5.710
3.26.20.5.011
3.26.40.5.029
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.101+|foundation .14+|Disk|Capacity (GB)|Implemented|CLI|use capacity
|Disk Id|Implemented|CLI|Uniquely identifies this disk in the array
|Group|Implemented|CLI|
|Location|Gap|CLI|Where this disk is physically located in the array
|Model|Implemented|CLI|
|Name|Implemented|CLI|
|Role|Implemented|CLI|
|Role Enum|Implemented|CLI|enum for disk role
|Serial Number|Implemented|CLI|
```

```
|Status|Implemented|CLI|
|Status Enum|Implemented|CLI|enum for disk status
|Type|Gap|CLI|
|Type Enum|Implemented|CLI|enum for disk type
|Vendor|Implemented|CLI|
.7+|Info|Api Name|Implemented|CLI|
|Api Version|Implemented|CLI|
|Client Api Name|Implemented|CLI|
|Client Api Version|Implemented|CLI|
|DataSource Name|Implemented|CLI|Info
|Date|Implemented|CLI|
|Originator ID|Implemented|CLI|
.14+|Storage|Display IP|Implemented|CLI|
|Failed Raw Capacity|Implemented|CLI|Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented|CLI|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|CLI|
|Manage URL|Implemented|CLI|
|Manufacturer|Implemented|CLI|
|Microcode Version|Implemented|CLI|
|Model|Implemented|CLI|
|Name|Implemented|CLI|
|Total Raw Capacity|Implemented|CLI|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|CLI|
|Spare Raw Capacity|Implemented|CLI|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|CLI|Specified if the storage supports
active-active configurations
|Virtual|Implemented|CLI|Is this a storage virtualization device?
.4+|Storage Node|Name|Implemented|CLI|
|Serial Number|Implemented|CLI|
|UUID|Implemented|CLI|
|ManagementIp Addresses|Implemented|CLI|
.18+|Storage Pool|Dedupe Enabled|Implemented|CLI|Is dedupe enabled on the
storage pool
|Include In Dwh Capacity|Implemented|CLI|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|CLI|
|Other Allocated Capacity|Gap|CLI|Capacity allocated for other (not data
and not snapshot)
|Other UsedCapacity (MB)|Implemented|CLI|Any capacity other than data and
snapshot
|Physical Disk Capacity (MB)|Implemented|CLI|used as raw capacity for
storage pool
```

```
|Raid Group|Implemented|CLI|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|CLI|ratio to convert from usable capacity
to raw capacity
|Redundancy|Implemented|CLI|Redundancy level
|Snapshot Allocated Capacity|Gap|CLI|Allocated capacity of snapshots in MB
|Snapshot Used Capacity|Implemented|CLI|
|Status|Implemented|CLI|
|Storage Pool Id|Implemented|CLI|
|Thin Provisioning Supported|Implemented|CLI|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|CLI|
|Total Used Capacity|Implemented|CLI|Total capacity in MB
|Type|Gap|CLI|
|Virtual|Implemented|CLI|Is this a storage virtualization device?
.7+|Storage Synchronization|Mode|Implemented|CLI|
|Mode Enum|Implemented|CLI|
|Source Volume|Implemented|CLI|
|State|Implemented|CLI|free text describing the device state
|State Enum|Implemented|CLI|
|Target Volume|Implemented|CLI|
|Technology|Implemented|CLI|technology which causes storage efficiency
.17+|Volume|AutoTier Policy Identifier|Implemented|CLI|Dynamic Tier Policy
identifier
|Auto Tiering|Implemented|CLI|indicates if this storagepool is
participating in auto tiering with other pools
|Capacity|Implemented|CLI|Snapshot Used capacity in MB
|DiskGroup|Implemented|CLI|Disk Group Type
|Disk Type|Not Available|CLI|
|Junction Path|Implemented|CLI|
|Meta|Implemented|CLI|Flag saying whether this volume is a meta volume
with memeber or not. Meta volumes will have DiskGroup empty!
|Name|Implemented|CLI|
|Total Raw Capacity|Implemented|CLI|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|CLI|Redundancy level
|Replica Source|Implemented|CLI|
|Replica Target|Implemented|CLI|
|Storage Pool Id|Implemented|CLI|
|Thin Provisioned|Implemented|CLI|
|Type|Gap|CLI|
|UUID|Implemented|CLI|
|Used Capacity|Implemented|CLI|
.4+|Volume Map|LUN|Implemented|CLI|Name of the backend lun
|Protocol Controller|Implemented|CLI|
```

```
|Storage Port|Implemented|CLI|
|Type|Gap|CLI|
.4+|Volume Mask|Initiator|Implemented|CLI|
|Protocol Controller|Implemented|CLI|
|Storage Port|Implemented|CLI|
|Type|Gap|CLI|
.7+|Volume Member|Capacity|Implemented|CLI|Snapshot Used capacity in MB
|Name|Implemented|CLI|
|Rank|Implemented|CLI|
|Total Raw Capacity|Implemented|CLI|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|CLI|Redundancy level
|Storage Pool Id|Implemented|CLI|
|Used Capacity|Implemented|CLI|
.5+|WWN Alias|Host Aliases|Implemented|CLI|
|IP|Implemented|CLI|
|Object Type|Implemented|CLI|
|Source|Implemented|CLI|
|WWN|Implemented|CLI|
.66+|performance .9+|Disk|IOps Read|Implemented|CLI|Number of read IOps on
the disk
| IOPs Total | Implemented | CLI |
| IOPs Write | Implemented | CLI |
|Throughput Read|Implemented|CLI|
|Throughput Total|Implemented|CLI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|CLI|
|Read Utilization|Implemented|CLI|
|Utilization Total|Implemented|CLI|
|Utilization Write|Implemented|CLI|
.16+|Storage|Cache Hit Ratio Read|Implemented|CLI|
|Cache Hit Ratio Total|Implemented|CLI|
|Cache Hit Ratio Write|Implemented|CLI|
|Failed Raw Capacity|Implemented|CLI|
|Raw Capacity|Implemented|CLI|
|Spare Raw Capacity|Implemented|CLI|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|CLI|
|IOPs other|Implemented|CLI|
|IOps Read|Implemented|CLI|Number of read IOps on the disk
| IOPs Total | Implemented | CLI |
|IOPs Write|Implemented|CLI|
|Latency Total|Implemented|CLI|
|Partial Blocked Ratio|Implemented|CLI|
|Throughput Read|Implemented|CLI|
|Throughput Total|Implemented|CLI|Average disk total rate (read and write
```

```
across all disks) in MB/s
|Throughput Write|Implemented|CLI|
.4+|Storage Node|IOps Read|Implemented|CLI|Number of read IOps on the disk
|IOPs Total|Implemented|CLI|
| IOPs Write | Implemented | CLI |
|Utilization Total|Implemented|CLI|
.20+|StoragePool Disk|Capacity Provisioned|Implemented|CLI|
|Raw Capacity|Implemented|CLI|
|Total Capacity|Implemented|CLI|
|Used Capacity|Implemented|CLI|
|Over Commit Capacity Ratio|Implemented|CLI|Reported as a time series
|Capacity Used Ratio|Implemented|CLI|
|IOps Read|Implemented|CLI|Number of read IOps on the disk
| IOPs Total | Implemented | CLI |
| IOPs Write | Implemented | CLI |
|Other Total Capacity|Implemented|CLI|
|Other Used Capacity|Implemented|CLI|
|Snapshot Reserved Capacity|Implemented|CLI|
|Snapshot Used Capacity|Implemented|CLI|
|Snapshot Used Capacity Ratio|Implemented|CLI| Reported as a time series
|Throughput Read|Implemented|CLI|
|Throughput Total|Implemented|CLI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|CLI|
|Read Utilization|Implemented|CLI|
|Utilization Total|Implemented|CLI|
|Utilization Write|Implemented|CLI|
.17+|Volume|Cache Hit Ratio Read|Implemented|CLI|
|Cache Hit Ratio Total|Implemented|CLI|
|Cache Hit Ratio Write|Implemented|CLI|
|Raw Capacity|Implemented|CLI|
|Total Capacity|Implemented|CLI|
|Used Capacity|Implemented|CLI|
|Capacity Used Ratio|Implemented|CLI|
|IOps Read|Implemented|CLI|Number of read IOps on the disk
| IOPs Total | Implemented | CLI |
| IOPs Write | Implemented | CLI |
|Latency Read|Implemented|CLI|
|Latency Total|Implemented|CLI|
|Latency Write|Implemented|CLI|
|Partial Blocked Ratio|Implemented|CLI|
|Throughput Read|Implemented|CLI|
|Throughput Total|Implemented|CLI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|CLI|
```

```
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Navi CLI
|CLI
|6389,2162,2163,443(HTTPS)/80(HTTP)
|true
|true
|true
|false
|===
<<top, Back to Top>>
== EMC Data Domain (SSH)
:description: Support Matrix Asciidoc for EMC Data Domain (SSH)
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
|DD VE
DD2200
DD2500
DD3300
DD4200
DD6300
DD6800
DD6900
DD7200
DD9300
DD9400
DD9500
DD9800
DD990
DD9900
16.1.2.051-633576
```

```
6.1.2.20-606786
6.1.2.50-632120
6.2.0.30-629757
6.2.0.35-635767
6.2.1.30-663869
6.2.1.40-671977
6.2.1.60-686365
7.10.0.0-1017741
7.10.1.0-1042928
7.2.0.30-663847
7.2.0.50-671975
7.2.0.60-682124
7.2.0.70-686759
7.2.0.90-692270
7.6.0.20-689174
7.6.0.30-690691
7.7.0.7-1007134
7.7.1.10-1011247
7.7.2.011-1011427
7.7.2.10-1011249
7.7.3.0-1011963
7.7.4.0-1017976
7.7.5.1-1040473
7.7.5.11-1046187
7.8.0.0-1008134
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.85+|foundation .14+|Disk|Capacity (GB)|Implemented|SSH|use capacity
|Disk Id|Implemented|SSH|Uniquely identifies this disk in the array
|Group|Implemented|SSH|
|Location|Gap|SSH|Where this disk is physically located in the array
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Role|Implemented|SSH|
|Role Enum|Implemented|SSH|enum for disk role
|Serial Number|Implemented|SSH|
|Speed|Implemented|SSH|Speed of disk (RPM)
|Status|Implemented|SSH|
|Status Enum|Implemented|SSH|enum for disk status
|Type|Gap|SSH|
|Type Enum|Implemented|SSH|enum for disk type
```

```
.5+|File Share|Is InternalVolume|Implemented|SSH|whether the file share
represents an internal volume (netapp volume) or is it a qtree/folder
within the internal volume
|Is Shared|Implemented|SSH|whether this fileShare has any shares
associated with it
|Name|Implemented|SSH|
|Path|Implemented|SSH|path of the fileShare
|Qtree Id|Implemented|SSH|unique id of the qtree
.3+|Info|DataSource Name|Implemented|SSH|Info
|Date|Implemented|SSH|
|Originator ID|Implemented|SSH|
.16+|Internal Volume|Data Allocated Capacity|Gap|SSH|capacity allocated
|Data Used Capacity|Implemented|SSH|
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|SSH|ratio of dedupe savings in percentage
|Internal Volume Id|Implemented|SSH|
|Name|Implemented|SSH|
|Other Allocated Capacity|Gap|SSH|Capacity allocated for other (not data
and not snapshot)
|Other UsedCapacity (MB)|Implemented|SSH|Any capacity other than data and
snapshot
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Storage Pool Id|Implemented|SSH|
|Thin Provisioned|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Total Used Capacity (MB)|Implemented|SSH|place holder for the used
capacity as read from the device
|Type|Gap|SSH|
.5+|QTree|Name|Implemented|SSH|
|Qtree Id||Implemented||SSH||unique id of the gtree
|Quota HardCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk
space, allowed for the quota target
|Quota SoftCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk
space, allowed for the quota target
|Type|Gap|SSH|
.7+|Quota|Hard Capacity Limit (MB)|Implemented|SSH|max amount of disk
space, allowed for the quota target (Hard limit)
|Internal Volume Id|Implemented|SSH|
|Qtree Id|Implemented|SSH|unique id of the qtree
|Quota Id|Implemented|SSH|unique id of the quota
|Soft Capacity Limit (MB)|Implemented|SSH|Maximum amount of disk space,
```

```
allowed for the quota target
|Type|Gap|SSH|
|Used Capacity|Implemented|SSH|
.3+|Share|IP Interfaces|Implemented|SSH|comma separated list of IP
addresses on which this share is exposed
|Name|Implemented|SSH|
|Protocol|Implemented|SSH|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|SSH|
|Permission|Implemented|SSH|Permissions for this particular share
.13+|Storage|Cpu Count|Implemented|SSH|Cpu Count of the storage
|Display IP|Implemented|SSH|
|Failed Raw Capacity|Implemented|SSH|Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented|SSH|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Microcode Version|Implemented|SSH|
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|SSH|
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.17+|Storage Pool|Data Allocated Capacity|Gap|SSH|capacity allocated for
data
|Data Used Capacity|Implemented|SSH|
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|SSH|ratio of dedupe savings in percentage
|Include In Dwh Capacity|Implemented|SSH|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|SSH|
|Other Allocated Capacity|Gap|SSH|Capacity allocated for other (not data
and not snapshot)
|Other UsedCapacity (MB)|Implemented|SSH|Any capacity other than data and
snapshot
|Physical Disk Capacity (MB)|Implemented|SSH|used as raw capacity for
storage pool
|Raid Group|Implemented|SSH|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Storage Pool Id|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
```

```
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Type|Gap|SSH|
|Virtual|Implemented|SSH|Is this a storage virtualization device?
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Data Domain CLI
|SSH
|SSH
122
|true
|true
|true
|true
|===
<<top, Back to Top>>
== EMC ECS
:description: Support Matrix Asciidoc for EMC ECS
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
IECS
|3.6.1.1
3.6.1.3
3.6.2.1
3.6.2.4
3.7.0.0
3.7.0.3
3.7.0.4
3.7.0.5
3.8.0.1
```

```
3.8.0.2
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.62+|foundation .12+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
.5+|Disk Group|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|DiskGroup Id|Implemented|HTTPS|unique id of the disk group
|Name|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.3+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
.11+|Internal Volume|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on
the storage pool
|Internal Volume Id|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Total Used Capacity (MB)|Implemented|HTTPS|place holder for the used
capacity as read from the device
|Type|Gap|HTTPS|
.3+|QTree|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Type|Gap|HTTPS|
```

```
.11+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.6+|Storage Node|Name|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
|Version|Implemented|HTTPS|software version
|Node Capacity Utilization Total in MB|Implemented|HTTPS|
|Node Capacity Utilization Usable in MB|Implemented|HTTPS|
|Node Capacity Utilization Used in MB|Implemented|HTTPS|
.11+|Storage Pool|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ
to cotnrol which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.21+|performance .4+|Internal Volume|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
|Total Objects|Implemented|HTTPS|
.4+|Storage|IOPs Total|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
```

```
.7+|Storage Node|IOPs Total|Implemented|HTTPS|
| Node Capacity Utilization Total | Implemented | HTTPS |
|Node Capacity Utilization Usable|Implemented|HTTPS|
|Node Capacity Utilization Used|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.6+|StoragePool Disk|Capacity Provisioned|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio|Implemented|HTTPS|Reported as a time series
|Capacity Used Ratio|Implemented|HTTPS|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|EMC ECS REST API
HTTPS
IHTTPS
1443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Dell EMC Isilon & PowerScale Rest
:description: Support Matrix Asciidoc for Dell EMC Isilon & PowerScale
Rest
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
```

```
|A200
A2000
A300
A3000
F200
F600
F800
F900
H400
H500
NL410
S210
X210
X400
X410
9.1.0.11
9.1.0.6
9.2.1.10
9.2.1.11
9.2.1.12
9.2.1.16
9.2.1.19
9.2.1.21
9.2.1.6
9.2.1.7
9.2.1.9
9.4.0.11
9.4.0.12
9.4.0.13
9.4.0.14
9.4.0.5
9.4.0.7
9.5.0.3
v8.0.0.4
v8.0.0.6
v8.0.0.7
v8.1.2.0
v8.2.2.0
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.115+|foundation .16+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
```

```
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Group|Implemented|HTTPS|
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type|Gap|HTTPS|
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
|Vendor Id|Implemented|HTTPS|
.10+|Disk Group|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|DiskGroup Id|Implemented|HTTPS|unique id of the disk group
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Redundancy|Implemented|HTTPS|Redundancy level
|Status|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Vendor DiskGroup Type|Implemented|HTTPS|vendor's designation of the disk
group type
|Vendor Tier|Implemented|HTTPS|Vendor Specific Tier Name
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.5+|File Share|Is InternalVolume|Implemented|HTTPS|whether the file share
represents an internal volume (netapp volume) or is it a qtree/folder
within the internal volume
| Is Shared | Implemented | HTTPS | whether this fileShare has any shares
associated with it
|Name|Implemented|HTTPS|
|Path|Implemented|HTTPS|path of the fileShare
|Qtree Id|Implemented|HTTPS|unique id of the qtree
.3+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
.16+|Internal Volume|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|HTTPS|capacity allocated for data
|Data Used Capacity|Implemented|HTTPS|
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Internal Volume Id|Implemented|HTTPS|
```

```
|Name|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Snapshot Allocated Capacity|Gap|HTTPS|Allocated capacity of snapshots in
|Snapshot Used Capacity|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
.6+|QTree|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the gtree
|Quota HardCapacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk
space, allowed for the quota target
|Quota SoftCapacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk
space, allowed for the quota target
|Quota UsedCapacity|Implemented|HTTPS|Space in MB currently used
|Type|Gap|HTTPS|
.12+|Quota|Hard Capacity Limit (MB)|Implemented|HTTPS|max amount of disk
space, allowed for the quota target (Hard limit)
|Hard File Limit|Implemented|HTTPS|max number of files allowed for the
quota target
|Internal Volume Id|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Quota Id|Implemented|HTTPS|unique id of the quota
|Soft Capacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk space,
allowed for the quota target
|Soft File Limit|Implemented|HTTPS|Max number of files allowed for the
quota target
|Threshold (MB)|Implemented|HTTPS|Disk space threshold, for the quota
target
|Type|Gap|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Used Files|Implemented|HTTPS|Number of files currently used
|User/Group Target|Implemented|HTTPS|user/group target this quota refers
to
.4+|Share|Description|Implemented|HTTPS|
| IP Interfaces | Implemented | HTTPS | comma separated list of IP addresses on
which this share is exposed
|Name|Implemented|HTTPS|
|Protocol|Implemented|HTTPS|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|HTTPS|
|Permission|Implemented|HTTPS|Permissions for this particular share
```

```
.14+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|UUID|Implemented|HTTPS|
.8+|Storage Node|Memory Size|Gap|HTTPS|device memory in MB
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Serial Number|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
|Version|Implemented|HTTPS|software version
|ManagementIp Addresses|Implemented|HTTPS|
.19+|Storage Pool|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|HTTPS|capacity allocated for data
|Data Used Capacity|Implemented|HTTPS|
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Other Allocated Capacity|Gap|HTTPS|Capacity allocated for other (not data
and not snapshot)
|Other UsedCapacity (MB)|Implemented|HTTPS|Any capacity other than data
and snapshot
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Snapshot Allocated Capacity|Gap|HTTPS|Allocated capacity of snapshots in
```

```
|Snapshot Used Capacity|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.66+|performance .6+|Disk|IOps Read|Implemented|HTTPS|Number of read IOps
on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.12+|Internal Volume|Total Capacity|Implemented|HTTPS|
|Total Data Capacity|Implemented|HTTPS|
|Data Used Capacity|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Snapshot Reserved Capacity|Implemented|HTTPS|
|Snapshot Used Capacity|Implemented|HTTPS|
|Snapshot Used Capacity Ratio|Implemented|HTTPS| Reported as a time series
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.4+|Qtree|Quota Hard Limit|Implemented|HTTPS|Capacity quota hard limit
|Used Capacity|Implemented|HTTPS|
|Total File Count|Implemented|HTTPS|
|Quota Capacity|Implemented|HTTPS|Physical used quota capacity
.12+|Storage|Failed Raw Capacity|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|HTTPS|
|Failed Disks|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Total|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
```

```
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.13+|Storage Node|Failed Disks|Implemented|HTTPS|
| IOPs Read | Implemented | HTTPS | Number of read IOPs on file system
| IOPs Write | Implemented | HTTPS | IOPs write of filesystem
|File Throughput Read|Implemented|HTTPS|
|FileSystem Throughput|Implemented|HTTPS|fileSystem Throughput write
| IOps Read | Implemented | HTTPS | Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Total|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
.19+|StoragePool Disk|Capacity Provisioned|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio|Implemented|HTTPS|Reported as a time series
|Capacity Used Ratio|Implemented|HTTPS|
|Total Data Capacity|Implemented|HTTPS|
|Data Used Capacity|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Other Total Capacity|Implemented|HTTPS|
|Other Used Capacity|Implemented|HTTPS|
|Snapshot Reserved Capacity|Implemented|HTTPS|
|Snapshot Used Capacity|Implemented|HTTPS|
|Snapshot Used Capacity Ratio|Implemented|HTTPS| Reported as a time series
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
```

```
|EMC Isilon & PowerScale REST API
|HTTPS
1
|443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Dell EMC Isilon / PowerScale (CLI)
:description: Support Matrix Asciidoc for Dell EMC Isilon / PowerScale
(CLI)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
A200
A2000
A300
F200
F800
F900
H400
H500
H600
H700
NL400
NL410
S210
X200
X210
X400
X410
|9.1.0.10
9.1.0.12
9.1.0.16
9.1.0.18
9.1.0.19
9.1.0.7
```

```
9.2.1.11
9.2.1.13
9.2.1.15
9.2.1.22
9.2.1.7
9.2.1.9
9.3.0.3
9.4.0.0
9.4.0.10
9.4.0.12
9.4.0.13
9.4.0.14
9.4.0.6
9.4.0.7
v7.1.1.8
v7.2.0.5
v7.2.1.3
v7.2.1.6
v8.0.0.4
v8.0.0.6
v8.0.0.7
v8.0.1.1
v8.1.2.0
v8.2.2.0
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.105+|foundation .16+|Disk|Capacity (GB)|Implemented|SSH|use capacity
|Disk Id|Implemented|SSH|Uniquely identifies this disk in the array
|Group|Implemented|SSH|
|Location|Gap|SSH|Where this disk is physically located in the array
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Role|Implemented|SSH|
|Role Enum|Implemented|SSH|enum for disk role
|Serial Number|Implemented|SSH|
|Speed|Implemented|SSH|Speed of disk (RPM)
|Status|Implemented|SSH|
|Status Enum|Implemented|SSH|enum for disk status
|Type|Gap|SSH|
|Type Enum|Implemented|SSH|enum for disk type
|Vendor|Implemented|SSH|
```

```
|Vendor Id|Implemented|SSH|
.10+|Disk Group|Capacity|Implemented|SSH|Snapshot Used capacity in MB
|DiskGroup Id|Implemented|SSH|unique id of the disk group
|Name|Implemented|SSH|
|Physical Disk Capacity (MB)|Implemented|SSH|used as raw capacity for
storage pool
|Redundancy|Implemented|SSH|Redundancy level
|Status|Implemented|SSH|
|Used Capacity|Implemented|SSH|
|Vendor DiskGroup Type|Implemented|SSH|vendor's designation of the disk
group type
|Vendor Tier|Implemented|SSH|Vendor Specific Tier Name
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.5+|File Share|Is InternalVolume|Implemented|SSH|whether the file share
represents an internal volume (netapp volume) or is it a gtree/folder
within the internal volume
|Is Shared|Implemented|SSH|whether this fileShare has any shares
associated with it
|Name|Implemented|SSH|
|Path|Implemented|SSH|path of the fileShare
|Qtree Id|Implemented|SSH|unique id of the qtree
.3+|Info|DataSource Name|Implemented|SSH|Info
|Date|Implemented|SSH|
|Originator ID|Implemented|SSH|
.14+|Internal Volume|Data Allocated Capacity|Gap|SSH|capacity allocated
for data
|Data Used Capacity|Implemented|SSH|
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Internal Volume Id|Implemented|SSH|
|Name|Implemented|SSH|
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Snapshot Allocated Capacity|Gap|SSH|Allocated capacity of snapshots in MB
|Snapshot Used Capacity|Implemented|SSH|
|Storage Pool Id|Implemented|SSH|
|Thin Provisioned|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Type|Gap|SSH|
.6+|QTree|Name|Implemented|SSH|
|Qtree Id|Implemented|SSH|unique id of the qtree
|Quota HardCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk
space, allowed for the quota target
|Quota SoftCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk
```

```
space, allowed for the quota target
|Quota UsedCapacity|Implemented|SSH|Space in MB currently used
|Type|Gap|SSH|
.12+|Quota|Hard Capacity Limit (MB)|Implemented|SSH|max amount of disk
space, allowed for the quota target (Hard limit)
|Hard File Limit|Implemented|SSH|max number of files allowed for the quota
target
|Internal Volume Id|Implemented|SSH|
|Qtree Id|Implemented|SSH|unique id of the qtree
|Quota Id|Implemented|SSH|unique id of the quota
|Soft Capacity Limit (MB)|Implemented|SSH|Maximum amount of disk space,
allowed for the quota target
|Soft File Limit|Implemented|SSH|Max number of files allowed for the quota
target
|Threshold (MB)|Implemented|SSH|Disk space threshold, for the quota target
|Type|Gap|SSH|
|Used Capacity|Implemented|SSH|
|Used Files|Implemented|SSH|Number of files currently used
|User/Group Target|Implemented|SSH|user/group target this quota refers to
.4+|Share|Description|Implemented|SSH|
|IP Interfaces|Implemented|SSH|comma separated list of IP addresses on
which this share is exposed
|Name|Implemented|SSH|
|Protocol|Implemented|SSH|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|SSH|
|Permission|Implemented|SSH|Permissions for this particular share
.12+|Storage|Display IP|Implemented|SSH|
|Failed Raw Capacity|Implemented|SSH|Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented|SSH|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Microcode Version|Implemented|SSH|
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|SSH|
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.6+|Storage Node|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Serial Number|Implemented|SSH|
|State|Implemented|SSH|free text describing the device state
```

```
|UUID|Implemented|SSH|
|ManagementIp Addresses|Implemented|SSH|
.15+|Storage Pool|Data Allocated Capacity|Gap|SSH|capacity allocated for
data
|Data Used Capacity|Implemented|SSH|
|Include In Dwh Capacity|Implemented|SSH|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|SSH|
|Physical Disk Capacity (MB)|Implemented|SSH|used as raw capacity for
storage pool
|Raid Group|Implemented|SSH|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Snapshot Allocated Capacity|Gap|SSH|Allocated capacity of snapshots in MB
|Snapshot Used Capacity|Implemented|SSH|
|Storage Pool Id|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Type|Gap|SSH|
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.46+|performance .6+|Internal Volume|Total Capacity|Implemented|SSH|
|Total Data Capacity|Implemented|SSH|
|Data Used Capacity|Implemented|SSH|
|Snapshot Reserved Capacity|Implemented|SSH|
|Snapshot Used Capacity|Implemented|SSH|
|Snapshot Used Capacity Ratio|Implemented|SSH| Reported as a time series
.5+|Qtree|Quota Hard Limit|Implemented|SSH|Capacity quota hard limit
|Quota Soft Limit|Implemented|SSH|Capacity Quota soft Limit
|Used Capacity|Implemented|SSH|
|Total File Count|Implemented|SSH|
|Quota Capacity|Implemented|SSH|Physical used quota capacity
.12+|Storage|Failed Raw Capacity|Implemented|SSH|
|Raw Capacity|Implemented|SSH|
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|SSH|
|IOPs other|Implemented|SSH|
|IOps Read|Implemented|SSH|Number of read IOps on the disk
| IOPs Total | Implemented | SSH |
| IOPs Write | Implemented | SSH |
|Latency Total|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
```

```
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
.12+|Storage Node|IOPs Read|Implemented|SSH|Number of read IOPs on file
system
| IOPs Write | Implemented | SSH | IOPs write of filesystem
|File Throughput Read|Implemented|SSH|
|FileSystem Throughput|Implemented|SSH|fileSystem Throughput write
| IOps Read | Implemented | SSH | Number of read IOps on the disk
|IOPs Total|Implemented|SSH|
|IOPs Write|Implemented|SSH|
|Latency Total|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
|Utilization Total|Implemented|SSH|
.11+|StoragePool Disk|Capacity Provisioned|Implemented|SSH|
|Raw Capacity|Implemented|SSH|
|Total Capacity|Implemented|SSH|
|Used Capacity|Implemented|SSH|
|Over Commit Capacity Ratio|Implemented|SSH|Reported as a time series
|Capacity Used Ratio|Implemented|SSH|
|Total Data Capacity|Implemented|SSH|
|Data Used Capacity|Implemented|SSH|
|Snapshot Reserved Capacity|Implemented|SSH|
|Snapshot Used Capacity|Implemented|SSH|
|Snapshot Used Capacity Ratio|Implemented|SSH| Reported as a time series
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Isilon SSH
|SSH
ISSH
122
|true
Ifalse
|true
|true
```

```
|===
<<top, Back to Top>>
== EMC PowerStore REST
:description: Support Matrix Asciidoc for EMC PowerStore REST
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|PowerStore 1000T
PowerStore 1200T
PowerStore 3000T
PowerStore 3200T
PowerStore 5000T
PowerStore 5000X
PowerStore 9000T
PowerStore 9200T
12.0.1.3
2.1.1.0
2.1.1.1
3.0.0.1
3.2.0.0
3.2.0.1
3.2.1.0
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.110+|foundation .8+|Disk|Capacity (GB)|Implemented||use capacity
|Disk Id|Implemented||Uniquely identifies this disk in the array
|Name|Implemented||
|Speed|Implemented||Speed of disk (RPM)
|Status|Implemented||
|Type|Gap||
|Type Enum|Implemented||enum for disk type
|Vendor|Implemented||
.5+|File Share|Is InternalVolume|Implemented||whether the file share
represents an internal volume (netapp volume) or is it a qtree/folder
within the internal volume
|Is Shared|Implemented||whether this fileShare has any shares associated
```

```
with it
|Name|Implemented||
|Path|Implemented||path of the fileShare
|Qtree Id|Implemented||unique id of the qtree
.4+|ISCSI Network Portal|IP|Implemented||
|Listening Port|Implemented||
|Nic|Implemented||
|OID|Implemented||
.3+|ISCSI Network Portal Group|OID|Implemented||
|Portal Group Name|Implemented||
|Portal Group Tag|Implemented||
.4+|ISCSI Node|Host Aliases|Implemented||
|Node Name|Implemented||
|OID|Implemented||
|Type|Gap||
.7+|ISCSI Session|OID|Implemented||
|Initiator OID|Implemented||
|Portal Group OID|Implemented||
|Number Of Connections|Implemented||
|Max Connections|Implemented||
|Initiator Ips|Implemented||
|Security|Implemented||
.5+|Info|Api Name|Implemented||
|DataSource Name|Implemented||Info
|Date|Implemented||
|Originator ID|Implemented||
|Originator Key|Implemented||
.12+|Internal Volume|Dedupe Enabled|Implemented||Is dedupe enabled on the
storage pool
|Internal Volume Id|Implemented||
|Name|Implemented||
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
|Storage Pool Id|Implemented||
|Thin Provisioned|Implemented||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented||
|Total Used Capacity|Implemented||Total capacity in MB
|Total Used Capacity (MB)|Implemented||place holder for the used capacity
as read from the device
|Type|Gap||
|Virtual Storage|Implemented||Owning virtual storage (vfiler)
.3+|QTree|Name|Implemented||
|Qtree Id|Implemented||unique id of the qtree
|Type|Gap||
```

```
.3+|Share|IP Interfaces|Implemented||comma separated list of IP addresses
on which this share is exposed
|Name|Implemented||
|Protocol|Implemented||enum for share protocol
.2+|Share Initiator|Initiator|Implemented||
|Permission|Implemented||Permissions for this particular share
.14+|Storage|Display IP|Implemented||
|Failed Raw Capacity|Implemented||Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented||The storage Family could be Clariion, Symmetrix, et
al
|IP|Implemented||
|Manage URL|Implemented||
|Manufacturer|Implemented||
|Microcode Version|Implemented||
|Model|Implemented||
|Name|Implemented||
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Serial Number|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|SupportActive Active|Implemented||Specified if the storage supports
active-active configurations
|Virtual|Implemented||Is this a storage virtualization device?
.6+|Storage Node|Model|Implemented||
|Name|Implemented||
|Partner Node UUID|Implemented||HA pair's UUID
|UUID|Implemented||
|Version|Implemented||software version
|Parent Serial Number|Implemented||
.12+|Storage Pool|Compression Savings|Implemented||ratio of compression
savings in percentage
|Include In Dwh Capacity|Implemented||A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented||
|Physical Disk Capacity (MB)|Implemented||used as raw capacity for
storage pool
|Raid Group|Implemented||indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
|Storage Pool Id|Implemented||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented||
```

```
|Total Used Capacity|Implemented||Total capacity in MB
|Type|Gap||
|Virtual|Implemented||Is this a storage virtualization device?
.10+|Volume|Capacity|Implemented||Snapshot Used capacity in MB
|Junction Path|Implemented||
|Name|Implemented||
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Storage Pool Id|Implemented||
|Thin Provisioned|Implemented||
|Type|Gap||
|UUID|Implemented||
|Used Capacity|Implemented||
|QoS - Policy|Implemented||
.5+|Volume Map|LUN|Implemented||Name of the backend lun
|Masking Required|Implemented||
|Protocol Controller|Implemented||
|Storage Port|Implemented||
|Type|Gap||
.3+|Volume Mask|Initiator|Implemented||
|Protocol Controller|Implemented||
|Type|Gap||
.4+|WWN Alias|Host Aliases|Implemented||
|Object Type|Implemented||
|Source|Implemented||
|WWN|Implemented||
.54+|performance .10+|Internal Volume|IOPs other|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.0+|Qtree.14+|Storage|Failed Raw Capacity|Implemented||
|Raw Capacity|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|StoragePools Capacity|Implemented||
|IOPs other|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
| IOPs Total | Implemented | |
|IOPs Write|Implemented||
```

```
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.10+|Storage Node|IOps Read|Implemented||Number of read IOps on the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Utilization Total|Implemented||
.7+|StoragePool Disk|Total Compression Savings|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
| IOPs Total | Implemented | |
| IOPs Write | Implemented | |
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.13+|Volume|Raw Capacity|Implemented||
|Total Capacity|Implemented||
|Used Capacity|Implemented||
|Capacity Used Ratio|Implemented||
| IOps Read | Implemented | | Number of read IOps on the disk
| IOPs Total | Implemented | |
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
```

```
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|EMC PowerStore REST API
IHTTPS
|HTTPS
1443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== EMC RecoverPoint (HTTP)
:description: Support Matrix Asciidoc for EMC RecoverPoint (HTTP)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|RecoverPoint
|5.1.P1(c.175)
5.1.SP4.P1(h.89)
5.1.SP4.P2(h.101)
5.1.SP4.P3(h.109)
5.1.SP4.P4(h.97)
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.34+|foundation .4+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.13+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
```

```
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.8+|Storage Node|Memory Size|Gap|HTTPS|device memory in MB
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Processors Count|Implemented|HTTPS|device CPU
|Serial Number|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
|Version|Implemented|HTTPS|software version
.9+|Storage Synchronization|Mode|Implemented|HTTPS|
|Mode Enum|Implemented|HTTPS|
|Source Storage|Implemented|HTTPS|
|Source Volume|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|State Enum|Implemented|HTTPS|
|Target Storage|Implemented|HTTPS|
|Target Volume|Implemented|HTTPS|
|Technology|Implemented|HTTPS|technology which causes storage efficiency
changed
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|RecoverPoint REST API
|HTTPS
| HTTPS
|443
```

```
Itrue
Itrue
|true
|true
|===
<<top, Back to Top>>
== EMC ScaleIO & PowerFlex REST
:description: Support Matrix Asciidoc for EMC ScaleIO & PowerFlex REST
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
|ScaleI0
|R2 6.11000.113
R2 6.11000.115
R3 0.1400.101
R3 5.1200.104
R3 6.500.113
R3 6.700.103
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.51+|foundation .8+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Name|Implemented|HTTPS|
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
.8+|Info|Api Version|Implemented|HTTPS|
|Client Api Description|Implemented|HTTPS|
|Client Api Name|Implemented|HTTPS|
|Client Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
```

```
.13+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
|IP|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|HTTPS|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.2+|Storage Node|Name|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
.12+|Storage Pool|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ
to cotnrol which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.8+|Volume|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Junction Path|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
|Host IPs|Implemented|HTTPS|
```

```
.39+|performance .6+|Disk|IOps Read|Implemented||Number of read IOps on
the disk
| IOPs Write | Implemented | |
|Throughput Read|Implemented||
|Throughput Write|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|IOPs Total|Implemented||
.10+|Storage|Failed Raw Capacity|Implemented||
|Raw Capacity|Implemented||
|Spare Raw Capacity|Implemented||Raw capacity of spare disks (sum of all
disks that are spare)
|StoragePools Capacity|Implemented||
| IOps Read | Implemented | | Number of read IOps on the disk
| IOPs Total | Implemented | |
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.9+|Storage Node|IOps Read|Implemented||Number of read IOps on the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.6+|StoragePool Disk|IOps Read|Implemented||Number of read IOps on the
disk
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Write|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|IOPs Total|Implemented||
.8+|Volume|Raw Capacity|Implemented||
|Total Capacity|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
|IOPs Total|Implemented||
| IOPs Write | Implemented | |
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
```

```
|Throughput Write|Implemented||
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|EMC ScaleIO & PowerFlex REST API
|HTTPS
|HTTPS
|443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== EMC Symmetrix CLI
:description: Support Matrix Asciidoc for EMC Symmetrix CLI
Models and versions supported by this data collector:
<.<|API versions <.<|Models <.<|Firmware versions
V10.0.0.0
V10.0.1.0
V7.6.2.67
V8.3.0.22
V8.3.0.6
V8.4.0.7
V8.4.0.9
V9.1.0.18
V9.1.0.5
V9.1.0.6
V9.2.0.0
V9.2.1.0
V9.2.1.1
```

```
V9.2.1.2
V9.2.2.0
V9.2.3.0
V9.2.3.1
V9.2.3.4
V9.2.3.5
V9.2.3.6
V9.2.4.1
V9.2.4.2
IDMX3-24
DMX4-24
PMax2000
PowerMax 2000
PowerMax 8000
VMAX-1
VMAX100K
VMAX10K
VMAX200K
VMAX250F
VMAX400K
VMAX40K
VMAX450F
VMAX850F
VMAX950F
|5773.198.142(168D0000) build 5
5876.272.177(16F40000) build 39
5876.286.194(16F40000) build 115
5876.309.196(16F40000) build 162
5977.1131.1131(17590000) build 551
5977.1151.1151(17590000) build 45
5977.1151.1151(17590000) build 59
5977.1151.1151(17590000) build 60
5977.1151.1151(17590000) build 9
5978.479.479(175A0000) build 195
5978.711.711(175A0000) build 113
5978.711.711(175A0000) build 139
5978.711.711(175A0000) build 149
5978.711.711(175A0000) build 194
5978.711.711(175A0000) build 196
5978.711.711(175A0000) build 220
5978.711.711(175A0000) build 239
5978.711.711(175A0000) build 252
5978.711.711(175A0000) build 267
5978.711.711(175A0000) build 278
5978.711.711(175A0000) build 287
5978.711.711(175A0000) build 335
```

```
5978.711.711(175A0000) build 365
5978.711.711(175A0000) build 366
5978.711.711(175A0000) build 388
5978.711.711(175A0000) build 416
5978.711.711(175A0000) build 436
5978.711.711(175A0000) build 438
5978.711.711(175A0000) build 448
5978.711.711(175A0000) build 461
5978.711.711(175A0000) build 480
5978.711.711(175A0000) build 484
5978.711.711(175A0000) build 502
5978.711.711(175A0000) build 529
5978.711.711(175A0000) build 8
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.153+|foundation .3+|Device Group|Name|Implemented||
|Storage Management Id|Implemented||
|Type|Gap||
.12+|Disk|Capacity (GB)|Implemented||use capacity
|Disk Id|Implemented||Uniquely identifies this disk in the array
|Group|Implemented||
|Model|Implemented||
|Name|Implemented||
|Role|Implemented||
|Role Enum|Implemented||enum for disk role
|Serial Number|Implemented||
|Status|Implemented||
|Status Enum|Implemented||enum for disk status
|Type Enum|Implemented||enum for disk type
|Vendor|Implemented||
.8+|Disk Group|Capacity|Implemented||Snapshot Used capacity in MB
|DiskGroup Id|Implemented||unique id of the disk group
|Name|Implemented||
|Physical Disk Capacity (MB)|Implemented||used as raw capacity for
storage pool
|Redundancy|Implemented||Redundancy level
|Used Capacity|Implemented||
|Vendor DiskGroup Type|Implemented||vendor's designation of the disk group
type
|Virtual|Implemented||Is this a storage virtualization device?
.4+|ISCSI Network Portal|IP|Implemented||
```

```
|Listening Port|Implemented||
|Nic|Implemented||
|OID|Implemented||
.3+|ISCSI Network Portal Group|OID|Implemented||
|Portal Group Name|Implemented||
|Portal Group Tag|Implemented||
.3+|ISCSI Node|Node Name|Implemented||
|OID|Implemented||
|Type|Gap||
.2+|ISCSI Node Map|OID|Implemented||
|Portal Group OID|Implemented||
.7+|ISCSI Session|Initiator Ips|Implemented||
|Initiator OID|Implemented||
|Max Connections|Implemented||
|Number Of Connections|Implemented||
|OID|Implemented||
|Portal Group OID|Implemented||
|Security|Implemented||
.10+|Info|Api Description|Implemented||
|Api Name|Implemented||
|Api Version|Implemented||
|Client Api Description|Implemented||
|Client Api Name|Implemented||
|Client Api Version|Implemented||
|DataSource Name|Implemented||Info
|Date|Implemented||
|Originator ID|Implemented||
|Originator Key|Implemented||
.5+|Network Pipe|Auto Recover|Implemented||
|Bidirectional|Implemented||
|Operational Status|Implemented||
|Source Id|Implemented||
|Target Id|Implemented||
.1+|Network Pipe Port Wwn|WWN|Implemented||
.3+|Protocol EndPoint|ID|Implemented||
|Name|Implemented||
|Storage Ip|Implemented||
.12+|Storage|Display IP|Implemented||
|Failed Raw Capacity|Implemented||Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented||The storage Family could be Clariion, Symmetrix, et
|IP|Implemented||
|Manufacturer|Implemented||
|Microcode Version|Implemented||
|Model|Implemented||
```

```
|Name|Implemented||
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Serial Number|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|Virtual|Implemented||Is this a storage virtualization device?
.2+|Storage Node|Name|Implemented||
|UUID|Implemented||
.23+|Storage Pool|Auto Tiering|Implemented||indicates if this storagepool
is participating in auto tiering with other pools
|Compression Enabled|Implemented||Is compression enabled on the storage
pool
|Compression Savings|Implemented||ratio of compression savings in
percentage
|Data Allocated Capacity|Gap||capacity allocated for data
|Data Used Capacity|Implemented||
|Dedupe Enabled|Implemented||Is dedupe enabled on the storage pool
|Include In Dwh Capacity|Implemented||A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented||
|Other UsedCapacity (MB)|Implemented||Any capacity other than data and
snapshot
|Physical Disk Capacity (MB)|Implemented||used as raw capacity for
storage pool
|Raid Group|Implemented||indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
|Redundancy|Implemented||Redundancy level
|Snapshot Used Capacity|Implemented||
|Soft Limit (MB)|Implemented||logical volume size that is defined during
volume creation or resizing operations
|Storage Pool Id|Implemented||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented||
|Total Used Capacity|Implemented||Total capacity in MB
|Type|Gap||
|Vendor Tier|Implemented||Vendor Specific Tier Name
|Virtual|Implemented||Is this a storage virtualization device?
|Effective Used Capacity Percent|Implemented||
.9+|Storage Synchronization|Mode|Implemented||
|Mode Enum|Implemented||
|Source Storage|Implemented||
|Source Volume|Implemented||
```

```
|State|Implemented||free text describing the device state
|State Enum|Implemented||
|Target Storage|Implemented||
|Target Volume|Implemented||
|Technology|Implemented||technology which causes storage efficiency
.21+|Volume|AutoTier Policy Identifier|Implemented||Dynamic Tier Policy
identifier
|Auto Tiering|Implemented||indicates if this storagepool is participating
in auto tiering with other pools
|Capacity|Implemented||Snapshot Used capacity in MB
|Disk Size|Implemented||comma seperated list of disk sizes (GB)
|Disk Type|Not Available||
|Mainframe|Implemented||indicates if this volume is a Mainframe Volume
|Meta|Implemented||Flag saying whether this volume is a meta volume with
memeber or not. Meta volumes will have DiskGroup empty!
|Name|Implemented||
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Redundancy|Implemented||Redundancy level
|Replica Source|Implemented||
|Replica Target|Implemented||
|Snapshot|Implemented||
|Storage Pool Id|Implemented||
|Thin Provisioned|Implemented||
|Type|Gap||
|UUID|Implemented||
|Used Capacity|Implemented||
|Virtual|Implemented||Is this a storage virtualization device?
|Written Capacity|Implemented||Total capacity written to this volume by a
Host in MB
|storage Groups|Implemented||
.5+|Volume Map|LUN|Implemented||Name of the backend lun
|Protocol Controller|Implemented||
|Storage Port|Implemented||
|TID|Implemented||
|Type|Gap||
.4+|Volume Mask|Initiator|Implemented||
|Protocol Controller|Implemented||
|Storage Port|Implemented||
|Type|Gap||
.10+|Volume Member|Auto Tiering|Implemented||indicates if this storagepool
is participating in auto tiering with other pools
|Capacity|Implemented||Snapshot Used capacity in MB
|Cylinders|Implemented||
|Name|Implemented||
```

```
|Rank|Implemented||
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Redundancy|Implemented||Redundancy level
|Storage Pool Id|Implemented||
|UUID|Implemented||
|Used Capacity|Implemented||
.2+|Volume Ref|Name|Implemented||
|Storage Ip|Implemented||
.4+|WWN Alias|Host Aliases|Implemented||
|Object Type|Implemented||
|Source|Implemented||
|WWN|Implemented||
.69+|performance .6+|Disk|IOps Read|Implemented||Number of read IOps on
the disk
| IOPs Total | Implemented | |
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.16+|Storage|Cache Hit Ratio Read|Implemented||
|Cache Hit Ratio Total|Implemented||
|Cache Hit Ratio Write|Implemented||
|Cache Utilization Total|Implemented||
|Failed Raw Capacity|Implemented||
|Raw Capacity|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|StoragePools Capacity|Implemented||
| IOPs other | Implemented | |
|IOps Read|Implemented||Number of read IOps on the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Write Pending|Implemented||total write pending
.11+|Storage Node|Cache Hit Ratio Total|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
```

```
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Utilization Total|Implemented||
.17+|StoragePool Disk|Capacity Provisioned|Implemented||
|Raw Capacity|Implemented||
|Total Capacity|Implemented||
|Used Capacity|Implemented||
|Over Commit Capacity Ratio|Implemented||Reported as a time series
|Capacity Used Ratio|Implemented||
|Total Data Capacity|Implemented||
|Data Used Capacity|Implemented||
|IOps Read|Implemented||Number of read IOps on the disk
| IOPs Total | Implemented | |
| IOPs Write | Implemented | |
|Other Used Capacity|Implemented||
|Snapshot Used Capacity|Implemented||
|Snapshot Used Capacity Ratio|Implemented|| Reported as a time series
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.19+|Volume|Cache Hit Ratio Read|Implemented||
|Cache Hit Ratio Total|Implemented||
|Cache Hit Ratio Write|Implemented||
|Raw Capacity|Implemented||
|Total Capacity|Implemented||
|Used Capacity|Implemented||
|Written Capacity|Implemented||
|Capacity Used Ratio|Implemented||
|CapacityRatio Written|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
| IOPs Total | Implemented | |
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Write Pending|Implemented||total write pending
|===
```

```
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|symcli
|CLI
12707
|true
|true
|true
|true
|Symmetrix SMI-S
|SMI-S
|HTTP/HTTPS
15988/5989
|true
|false
|false
|true
|===
<<top, Back to Top>>
== Dell Unisphere REST
:description: Support Matrix Asciidoc for Dell Unisphere REST
Models and versions supported by this data collector:
|===
<.<|API versions <.<|Models <.<|Firmware versions
|V10.0.0.5
V10.0.1.3
V9.2.1.6
V9.2.3.20
V9.2.3.22
V9.2.3.4
V9.2.4.1
|PowerMax 2000
```

```
PowerMax 2500
PowerMax 8000
VMAX250F
VMAX950F
|5978.479.479 build 350
5978.711.711 build 252
5978.711.711 build 278 build 278
5978.711.711 build 287
5978.711.711 build 329 build 329
5978.711.711 build 365
5978.711.711 build 365 build 365
5978.711.711 build 376
5978.711.711 build 388 build 388
5978.711.711 build 416
5978.711.711 build 435
5978.711.711 build 448
5978.711.711 build 461 build 461
5978.711.711 build 481 build 481
5978.711.711 build 484
5978.711.711 build 484 build 484
5978.711.711 build 502
6079.125.0 build 53 build 53
6079.175.0 build 0 build 0
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.80+|foundation .9+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
.10+|Info|Api Description|Implemented|HTTPS|
|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|Client Api Description|Implemented|HTTPS|
|Client Api Name|Implemented|HTTPS|
|Client Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
```

```
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.12+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.2+|Storage Node|Name|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
.13+|Storage Pool|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ
to cotnrol which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|Effective Used Capacity Percent|Implemented|HTTPS|
.9+|Storage Synchronization|Mode|Implemented|HTTPS|
|Mode Enum|Implemented|HTTPS|
|Source Storage|Implemented|HTTPS|
|Source Volume|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|State Enum|Implemented|HTTPS|
|Target Storage|Implemented|HTTPS|
```

```
|Target Volume|Implemented|HTTPS|
|Technology|Implemented|HTTPS|technology which causes storage efficiency
changed
.13+|Volume|AutoTier Policy Identifier|Implemented|HTTPS|Dynamic Tier
Policy identifier
|Auto Tiering|Implemented|HTTPS|indicates if this storagepool is
participating in auto tiering with other pools
|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Disk Type|Not Available|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|HTTPS|Redundancy level
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|storage Groups|Implemented|HTTPS|
.4+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|WWN Alias|Host Aliases|Implemented|HTTPS|
|Object Type|Implemented|HTTPS|
|Source|Implemented|HTTPS|
|WWN|Implemented|HTTPS|
.48+|performance .15+|Storage|Latency Total|Implemented|HTTPS|
|Latency Read|Implemented|HTTPS|
| IOPs other | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Raw Capacity|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
| IOPs Total | Implemented | HTTPS |
|Latency Write|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
|Write Pending|Implemented|HTTPS|total write pending
|Cache Hit Ratio Read|Implemented|HTTPS|
|Cache Hit Ratio Total|Implemented|HTTPS|
|Cache Hit Ratio Write|Implemented|HTTPS|
|Throughput Write|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
```

```
write across all disks) in MB/s
.0+|Storage Node.19+|StoragePool Disk|Capacity
Provisioned | Implemented | HTTPS |
|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio|Implemented|HTTPS|Reported as a time series
|Capacity Used Ratio|Implemented|HTTPS|
|Total Compression Savings|Implemented|HTTPS|
|Compression Savings Space|Implemented|HTTPS|
|Total Data Capacity|Implemented|HTTPS|
|Data Used Capacity|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Snapshot Reserved Capacity|Implemented|HTTPS|
|Snapshot Used Capacity Ratio|Implemented|HTTPS| Reported as a time series
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
.14+|Volume|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
|CapacityRatio Written|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
```

```
|Dell Unisphere API
|HTTPS
| HTTPS
1443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== EMC VNX (SSH)
:description: Support Matrix Asciidoc for EMC VNX (SSH)
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
| VNX5300
VNX5400
VNX5700
VNX5800
|05.33.009.5.231
7.1.76-4
7.1.80 - 3
8.1.9-232
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.135+|foundation .13+|Disk|Disk Id|Implemented|SSH|Uniquely identifies
this disk in the array
|Name|Implemented|SSH|
|Capacity (GB) | Implemented | SSH | use capacity
|Location|Gap|SSH|Where this disk is physically located in the array
|Role Enum|Implemented|SSH|enum for disk role
|Role|Implemented|SSH|
|Status|Implemented|SSH|
|Status Enum|Implemented|SSH|enum for disk status
```

```
|Serial Number|Implemented|SSH|
|Vendor|Implemented|SSH|
|Model|Implemented|SSH|
|Type|Gap|SSH|
|Type Enum|Implemented|SSH|enum for disk type
.6+|File Share|Is InternalVolume|Implemented|SSH|whether the file share
represents an internal volume (netapp volume) or is it a gtree/folder
within the internal volume
|Is Shared|Implemented|SSH|whether this fileShare has any shares
associated with it
|Name|Implemented|SSH|
|Path|Implemented|SSH|path of the fileShare
|Qtree Id|Implemented|SSH|unique id of the qtree
|Status|Implemented|SSH|
.8+|Info|Api Name|Implemented|SSH|
|Api Version|Implemented|SSH|
|Client Api Name|Implemented|SSH|
|Client Api Version|Implemented|SSH|
|DataSource Name|Implemented|SSH|Info
|Date|Implemented|SSH|
|Originator ID|Implemented|SSH|
|Originator Key|Implemented|SSH|
.21+|Internal Volume|Data Allocated Capacity|Gap|SSH|capacity allocated
for data
|Data Used Capacity|Implemented|SSH|
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|SSH|ratio of dedupe savings in percentage
|GuidKey 1|Implemented|SSH|GuidKey1 is implicit for all objects whose GUID
key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|SSH|GuidKey2 is implicit for all objects whose GUID
key has not changed since OCI version 7.3.5.
|Internal Volume Id|Implemented|SSH|
|Last Snapshot Time|Implemented|SSH|time of last snapshot
|Name|Implemented|SSH|
|Other Allocated Capacity|Gap|SSH|Capacity allocated for other (not data
and not snapshot)
|Other UsedCapacity (MB)|Implemented|SSH|Any capacity other than data and
snapshot
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Snapshot Count|Implemented|SSH|Number of snapshots on the internal
|Storage Pool Id|Implemented|SSH|
|Thin Provisioned|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
```

```
|Total Allocated Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Total Used Capacity (MB)|Implemented|SSH|place holder for the used
capacity as read from the device
|Type|Gap|SSH|
|Virtual Storage|Implemented|SSH|Owning virtual storage (vfiler)
.8+|QTree|GuidKey 1|Implemented|SSH|GuidKey1 is implicit for all objects
whose GUID key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|SSH|GuidKey2 is implicit for all objects whose GUID
key has not changed since OCI version 7.3.5.
|Name|Implemented|SSH|
|Qtree Id|Implemented|SSH|unique id of the qtree
|Quota HardCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk
space, allowed for the quota target
|Quota SoftCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk
space, allowed for the quota target
|Quota UsedCapacity|Implemented|SSH|Space in MB currently used
|Type|Gap|SSH|
.11+|Quota|Quota Id|Implemented|SSH|unique id of the quota
|Type|Gap|SSH|
|Internal Volume Id|Implemented|SSH|
|Qtree Id|Implemented|SSH|unique id of the qtree
|Soft File Limit|Implemented|SSH|Max number of files allowed for the quota
target
| Hard Capacity Limit (MB) | Implemented | SSH | max amount of disk space,
allowed for the quota target (Hard limit)
|Soft Capacity Limit (MB)|Implemented|SSH|Maximum amount of disk space,
allowed for the quota target
|Used Files|Implemented|SSH|Number of files currently used
|Used Capacity|Implemented|SSH|
|GuidKey 1|Implemented|SSH|GuidKey1 is implicit for all objects whose GUID
key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|SSH|GuidKey2 is implicit for all objects whose GUID
key has not changed since OCI version 7.3.5.
.3+|Share|IP Interfaces|Implemented|SSH|comma separated list of IP
addresses on which this share is exposed
|Name|Implemented|SSH|
|Protocol|Implemented|SSH|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|SSH|
|Permission|Implemented|SSH|Permissions for this particular share
.15+|Storage|Cpu Count|Implemented|SSH|Cpu Count of the storage
|Display IP|Implemented|SSH|
|Failed Raw Capacity|Implemented|SSH|Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented|SSH|The storage Family could be Clariion, Symmetrix,
et al
```

```
|IP|Implemented|SSH|
|Manage URL|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Microcode Version|Implemented|SSH|
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|SSH|
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|SSH|Specified if the storage supports
active-active configurations
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.4+|Storage Node|UUID|Implemented|SSH|
|Name|Implemented|SSH|
|Serial Number|Implemented|SSH|
|ManagementIp Addresses|Implemented|SSH|
.18+|Storage Pool|Data Allocated Capacity|Gap|SSH|capacity allocated for
data
|Data Used Capacity|Implemented|SSH|
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Include In Dwh Capacity|Implemented|SSH|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|SSH|
|Physical Disk Capacity (MB)|Implemented|SSH|used as raw capacity for
storage pool
|Raid Group|Implemented|SSH|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Redundancy|Implemented|SSH|Redundancy level
|Snapshot Allocated Capacity|Gap|SSH|Allocated capacity of snapshots in MB
|Snapshot Used Capacity|Implemented|SSH|
|Status|Implemented|SSH|
|Storage Pool Id|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Type|Gap|SSH|
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.13+|Volume|Name|Implemented|SSH|
|Junction Path|Implemented|SSH|
|Storage Pool Id|Implemented|SSH|
|Auto Tiering|Implemented|SSH|indicates if this storagepool is
```

```
participating in auto tiering with other pools
|AutoTier Policy Identifier|Implemented|SSH|Dynamic Tier Policy identifier
|UUID|Implemented|SSH|
|Type|Gap|SSH|
|Thin Provisioned|Implemented|SSH|
|Capacity|Implemented|SSH|Snapshot Used capacity in MB
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Used Capacity|Implemented|SSH|
|Redundancy|Implemented|SSH|Redundancy level
|Disk Type|Not Available|SSH|
.4+|Volume Map|LUN|Implemented|SSH|Name of the backend lun
|Storage Port|Implemented|SSH|
|Protocol Controller|Implemented|SSH|
|Type|Gap|SSH|
.4+|Volume Mask|Storage Port|Implemented|SSH|
|Initiator|Implemented|SSH|
|Protocol Controller|Implemented|SSH|
|Type|Gap|SSH|
.5+|WWN Alias|Source|Implemented|SSH|
|Host Aliases|Implemented|SSH|
|WWN|Implemented|SSH|
|Object Type|Implemented|SSH|
|IP|Implemented|SSH|
.27+|performance .9+|Disk|IOps Read|Implemented|SSH|Number of read IOps on
the disk
| IOPs Total | Implemented | SSH |
|IOPs Write|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
|Read Utilization|Implemented|SSH|
|Utilization Total|Implemented|SSH|
|Utilization Write|Implemented|SSH|
.14+|Storage|Failed Raw Capacity|Implemented|SSH|
|Raw Capacity|Implemented|SSH|
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|SSH|
| IOPs other | Implemented | SSH |
| IOps Read | Implemented | SSH | Number of read IOps on the disk
| IOPs Total | Implemented | SSH |
|IOPs Write|Implemented|SSH|
|Latency Read|Implemented|SSH|
|Latency Total|Implemented|SSH|
```

```
|Latency Write|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
.4+|Storage Node|IOps Read|Implemented|SSH|Number of read IOps on the disk
|IOPs Total|Implemented|SSH|
|IOPs Write|Implemented|SSH|
|Utilization Total|Implemented|SSH|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|VNX SSH and CLI
|SSH
|SSH
122
|true
|false
|true
|true
|===
<<top, Back to Top>>
== EMC VNXe & Unity Unisphere (CLI)
:description: Support Matrix Asciidoc for EMC VNXe & Unity Unisphere (CLI)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|Unity 300
Unity 300F
Unity 350F
Unity 380
Unity 380F
Unity 400
```

```
Unity 400F
Unity 450F
Unity 480F
Unity 500
Unity 550F
Unity 600
Unity 600F
Unity 650F
Unity 680F
Unity 880
VNXe3200
|3.1.17.10223906
3.1.17.10229825
4.1.2.9257522
4.2.1.9535982
4.2.3.9670635
4.5.1.0.5.001
5.0.2.0.5.009
5.0.6.0.5.008
5.0.8.0.5.007
5.1.2.0.5.007
5.1.3.0.5.003
5.2.1.0.5.013
5.2.2.0.5.004
5.2.2.0.6.201
5.3.0.0.5.120
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.124+|foundation .15+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Group|Implemented|HTTPS|
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type|Gap|HTTPS|
```

```
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
.6+|Disk Group|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|DiskGroup Id|Implemented|HTTPS|unique id of the disk group
|Disk Type|Not Available|HTTPS|
|Name|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.5+|File Share|Is InternalVolume|Implemented|HTTPS|whether the file share
represents an internal volume (netapp volume) or is it a qtree/folder
within the internal volume
|Is Shared|Implemented|HTTPS|whether this fileShare has any shares
associated with it
|Name|Implemented|HTTPS|
|Path|Implemented|HTTPS|path of the fileShare
|Qtree Id|Implemented|HTTPS|unique id of the qtree
.4+|ISCSI Network Portal|IP|Implemented|HTTPS|
|Listening Port|Implemented|HTTPS|
|Nic|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.3+|ISCSI Network Portal Group|OID|Implemented|HTTPS|
|Portal Group Name|Implemented|HTTPS|
|Portal Group Tag|Implemented|HTTPS|
.3+|ISCSI Node|Node Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
.3+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
.17+|Internal Volume|Data Allocated Capacity|Gap|HTTPS|capacity allocated
for data
|Data Used Capacity|Implemented|HTTPS|
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|Internal Volume Id|Implemented|HTTPS|
|Last Snapshot Time|Implemented|HTTPS|time of last snapshot
|Name|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Snapshot Count|Implemented|HTTPS|Number of snapshots on the internal
volumes
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
```

```
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Total Used Capacity (MB)|Implemented|HTTPS|place holder for the used
capacity as read from the device
|Type|Gap|HTTPS|
.3+|QTree|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Type|Gap|HTTPS|
.4+|Share|Description|Implemented|HTTPS|
| IP Interfaces | Implemented | HTTPS | comma separated list of IP addresses on
which this share is exposed
|Name|Implemented|HTTPS|
|Protocol|Implemented|HTTPS|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|HTTPS|
|Permission|Implemented|HTTPS|Permissions for this particular share
.12+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.6+|Storage Node|Memory Size|Gap|HTTPS|device memory in MB
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Serial Number|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
.16+|Storage Pool|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
```

```
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
|Soft Limit (MB)|Implemented|HTTPS|logical volume size that is defined
during volume creation or resizing operations
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.9+|Storage Synchronization|Mode|Implemented|HTTPS|
|Mode Enum|Implemented|HTTPS|
|Source Storage|Implemented|HTTPS|
|Source Volume|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|State Enum|Implemented|HTTPS|
|Target Storage|Implemented|HTTPS|
|Target Volume|Implemented|HTTPS|
|Technology|Implemented|HTTPS|technology which causes storage efficiency
changed
.8+|Volume|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Junction Path|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
.4+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.45+|performance .6+|Disk|IOps Read|Implemented|HTTPS|Number of read IOps
on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
```

```
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.7+|Internal Volume|Total Capacity|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.0+|Qtree.4+|Storage|Failed Raw Capacity|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|HTTPS|
.4+|Storage Node|IOps Read|Implemented|HTTPS|Number of read IOps on the
disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Utilization Total|Implemented|HTTPS|
.13+|StoragePool Disk|Capacity Provisioned|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Capacity Soft Limit|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio | Implemented | HTTPS | Reported as a time series
|Capacity Used Ratio|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.11+|Volume|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
| IOps Read | Implemented | HTTPS | Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Total|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
```

```
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
| VNXe & Unisphere CLI
|HTTPS
|HTTPS
1443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== EMC VPLEX
:description: Support Matrix Asciidoc for EMC VPLEX
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
| VPLEX
|5.4.1.00.00.07
5.4.1.01.00.05
6.2.0.03.00.02
6.2.0.04.00.07
6.2.0.05.00.11
6.2.0.07.00.02
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
```

```
.61+|foundation .4+|Info|DataSource Name|Implemented|HTTP/S|Info
|Date|Implemented|HTTP/S|
|Originator ID|Implemented|HTTP/S|
|Originator Key|Implemented|HTTP/S|
.13+|Storage|Display IP|Implemented|HTTP/S|
|Failed Raw Capacity|Implemented|HTTP/S|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTP/S|The storage Family could be Clariion,
Symmetrix, et al
|IP|Implemented|HTTP/S|
|Manage URL|Implemented|HTTP/S|
|Manufacturer|Implemented|HTTP/S|
|Microcode Version|Implemented|HTTP/S|
|Model|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Total Raw Capacity|Implemented|HTTP/S|Total raw capacity (sum of all
disks on the array)
|Serial Number|Implemented|HTTP/S|
|Spare Raw Capacity|Implemented|HTTP/S|Raw capapcity of spare disks (sum
of all disks that are spare)
|Virtual|Implemented|HTTP/S|Is this a storage virtualization device?
.4+|Storage Node|Name|Implemented|HTTP/S|
|Serial Number|Implemented|HTTP/S|
|State|Implemented|HTTP/S|free text describing the device state
|UUID|Implemented|HTTP/S|
.15+|Storage Pool|Dedupe Enabled|Implemented|HTTP/S|Is dedupe enabled on
the storage pool
|Name|Implemented|HTTP/S|
|Other Allocated Capacity|Gap|HTTP/S|Capacity allocated for other (not
data and not snapshot)
|Other UsedCapacity (MB)|Implemented|HTTP/S|Any capacity other than data
and snapshot
|Physical Disk Capacity (MB)|Implemented|HTTP/S|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTP/S|indicates whether this storagePool is a
raid group
|Raw to Usable Ratio|Implemented|HTTP/S|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTP/S|Redundancy level
|Status|Implemented|HTTP/S|
|Storage Pool Id|Implemented|HTTP/S|
|Thin Provisioning Supported|Implemented|HTTP/S|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTP/S|
|Total Used Capacity|Implemented|HTTP/S|Total capacity in MB
|Type|Gap|HTTP/S|
```

```
|Virtual|Implemented|HTTP/S|Is this a storage virtualization device?
.9+|Storage Synchronization|Mode|Implemented|HTTP/S|
|Mode Enum|Implemented|HTTP/S|
|Source Storage|Implemented|HTTP/S|
|Source Volume|Implemented|HTTP/S|
|State|Implemented|HTTP/S|free text describing the device state
|State Enum|Implemented|HTTP/S|
|Target Storage|Implemented|HTTP/S|
|Target Volume|Implemented|HTTP/S|
|Technology|Implemented|HTTP/S|technology which causes storage efficiency
changed
.8+|Volume|Capacity|Implemented|HTTP/S|Snapshot Used capacity in MB
|Name|Implemented|HTTP/S|
|Total Raw Capacity|Implemented|HTTP/S|Total raw capacity (sum of all
disks on the array)
|Redundancy|Implemented|HTTP/S|Redundancy level
|Storage Pool Id|Implemented|HTTP/S|
|Thin Provisioned|Implemented|HTTP/S|
|UUID|Implemented|HTTP/S|
|Virtual|Implemented|HTTP/S|Is this a storage virtualization device?
.4+|Volume Map|LUN|Implemented|HTTP/S|Name of the backend lun
|Protocol Controller|Implemented|HTTP/S|
|Storage Port|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.4+|Volume Mask|Initiator|Implemented|HTTP/S|
|Protocol Controller|Implemented|HTTP/S|
|Storage Port|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.34+|performance .7+|Storage|IOPs Total|Implemented|SSH|
|Latency Read|Implemented|SSH|
|Latency Total|Implemented|SSH|
|Latency Write|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
.11+|Storage Node|Cache Hit Ratio Total|Implemented|SSH|
|IOps Read|Implemented|SSH|Number of read IOps on the disk
|IOPs Total|Implemented|SSH|
|IOPs Write|Implemented|SSH|
|Latency Read|Implemented|SSH|
|Latency Total|Implemented|SSH|
|Latency Write|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
```

```
|Throughput Write|Implemented|SSH|
|Utilization Total|Implemented|SSH|
.7+|StoragePool Disk|Capacity Provisioned|Implemented|SSH|
|Total Capacity|Implemented|SSH|
|Used Capacity|Implemented|SSH|
|Over Commit Capacity Ratio|Implemented|SSH|Reported as a time series
|Capacity Used Ratio|Implemented|SSH|
|Other Total Capacity|Implemented|SSH|
|Other Used Capacity|Implemented|SSH|
.9+|Volume|Raw Capacity|Implemented|SSH|
|Total Capacity|Implemented|SSH|
| IOPs Total | Implemented | SSH |
|Latency Read|Implemented|SSH|
|Latency Total|Implemented|SSH|
|Latency Write|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|EMC VPLEX CLI
|SSH
|SSH
122
|true
|true
|true
|true
|EMC VPLEX API
|HTTP/HTTPS
|HTTP/HTTPS
180/443
|true
|true
```

```
|true
|true
|===
<<top, Back to Top>>
== EMC XtremIO (HTTP)
:description: Support Matrix Asciidoc for EMC XtremIO (HTTP)
Models and versions supported by this data collector:
<.<|API versions <.<|Models <.<|Firmware versions
|6.2.1
6.2.2
6.3.1
6.3.2
6.3.3
6.4.0
|1 Bricks & 125TB
1 Bricks & 24TB
1 Bricks & 26TB
1 Bricks & 31TB
1 Bricks & 62TB
1 Bricks & 8TB
1X10TB
1X20TB
1X40TB
2 Bricks & 52TB
2 Bricks & 62TB
2 Bricks & 76TB
2 Bricks & 83TB
2X10TB
2X20TB
2X40TB
3 Bricks & 251TB
3 Bricks & 283TB
4 Bricks & 125TB
4 Bricks & 503TB
4 Bricks & 628TB
4 Bricks & 754TB
4X20TB
4X40TB
6X20TB
|4.0.27-1
```

```
4.0.31-11
6.1.0-99 X2
6.3.3-8 X2
6.4.0-22 X2
6.4.0-36 hotfix 2 X2
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.92+|foundation .18+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Group|Implemented|HTTPS|
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Potential Transfer Rate|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Seek Time|Implemented|HTTPS|
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type|Gap|HTTPS|
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
|Vendor Id|Implemented|HTTPS|
.8+|Info|Api Version|Implemented|HTTPS|
|Client Api Description|Implemented|HTTPS|
|Client Api Name|Implemented|HTTPS|
|Client Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.15+|Storage|Cpu Count|Implemented|HTTPS|Cpu Count of the storage
|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
```

```
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|HTTPS|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.6+|Storage Node|Name|Implemented|HTTPS|
|Processors Count|Implemented|HTTPS|device CPU
|Serial Number|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
|ManagementIp Addresses|Implemented|HTTPS|
.19+|Storage Pool|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|HTTPS|capacity allocated for data
|Data Used Capacity|Implemented|HTTPS|
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.7+|Storage Synchronization|Mode|Implemented|HTTPS|
|Mode Enum|Implemented|HTTPS|
```

```
|Source Volume|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|State Enum|Implemented|HTTPS|
|Target Volume|Implemented|HTTPS|
|Technology|Implemented|HTTPS|technology which causes storage efficiency
changed
.13+|Volume|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Disk Size|Implemented|HTTPS|comma seperated list of disk sizes (GB)
|Disk Speed|Implemented|HTTPS|comma seperated list of disk speeds (rpm)
|Disk Type|Not Available|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|HTTPS|Redundancy level
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.3+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Protocol Controller|Implemented|HTTPS|
|Type|Gap|HTTPS|
.3+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Type|Gap|HTTPS|
.41+|performance .10+|Storage|IOps Read|Implemented|HTTPS|Number of read
IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Partial Blocked Ratio|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.9+|Storage Node|IOps Read|Implemented|HTTPS|Number of read IOps on the
disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
```

```
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.8+|StoragePool Disk|Capacity Provisioned|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio|Implemented|HTTPS|Reported as a time series
|Capacity Used Ratio|Implemented|HTTPS|
|Total Data Capacity|Implemented|HTTPS|
|Data Used Capacity|Implemented|HTTPS|
.14+|Volume|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
| IOps Read | Implemented | HTTPS | Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Partial Blocked Ratio|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|EMC XtremIO REST API
|HTTPS
|HTTPS
1443
|true
|true
|true
|true
```

```
|===
<<top, Back to Top>>
== NetApp E-Series
:description: Support Matrix Asciidoc for NetApp E-Series
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
|2600
2660
2680
2702
2704
2800B
2804
2806
3000
5480
5486
5488
5504
5564
5600
5700
5700B
6000
08.40.60.01
8.10.14.0
8.20.11.0
8.20.27.0
8.20.30.0
8.20.5.0
8.20.8.0
8.25.14.0
8.25.6.0
8.30.1.0
8.40.0.1
8.40.0.3
8.40.20.0
8.40.30.3
8.40.40.0
8.40.50.0
8.40.60.1
```

```
8.40.60.2
8.40.60.3
8.42.20.0
8.50.0.3
8.50.0.4
8.51.0.0
8.52.0.0
8.52.0.1
8.53.0.1
8.53.0.4
8.62.0.0
8.62.0.2
8.63.0.2
8.70.0.3
8.71.2.0
8.71.3.0
8.72.0.0
8.72.1.0
8.72.2.0
8.73.0.0
8.74.0.0
8.74.1.0
8.74.2.0
8.74.3.0
8.75.0.0
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.94+|foundation .16+|Disk|Capacity (GB)|Implemented|RMI|use capacity
|Disk Id|Implemented|RMI|Uniquely identifies this disk in the array
|Group|Implemented|RMI|
|Location|Gap|RMI|Where this disk is physically located in the array
|Model|Implemented|RMI|
|Name|Implemented|RMI|
|Role|Implemented|RMI|
|Role Enum|Implemented|RMI|enum for disk role
|Serial Number|Implemented|RMI|
|Speed|Implemented|RMI|Speed of disk (RPM)
|Status|Implemented|RMI|
|Status Enum|Implemented|RMI|enum for disk status
|Type|Gap|RMI|
|Type Enum|Implemented|RMI|enum for disk type
```

```
|Vendor|Implemented|RMI|
|Vendor Id|Implemented|RMI|
.4+|ISCSI Network Portal|IP|Implemented|RMI|
|Listening Port|Implemented|RMI|
|Nic|Implemented|RMI|
|OID|Implemented|RMI|
.3+|ISCSI Network Portal Group|OID|Implemented|RMI|
|Portal Group Name|Implemented|RMI|
|Portal Group Tag|Implemented|RMI|
.4+|ISCSI Node|Host Aliases|Implemented|RMI|
|Node Name|Implemented|RMI|
|OID|Implemented|RMI|
|Type|Gap|RMI|
.8+|ISCSI Session|Initiator Ips|Implemented|RMI|
|Initiator OID|Implemented|RMI|
|Max Connections|Implemented|RMI|
|Number Of Connections|Implemented|RMI|
|OID|Implemented|RMI|
|Portal Group OID|Implemented|RMI|
|Security|Implemented|RMI|
|Target Session Id|Implemented|RMI|
.3+|Info|DataSource Name|Implemented|RMI|Info
|Date|Implemented|RMI|
|Originator ID|Implemented|RMI|
.12+|Storage|Display IP|Implemented|RMI|
|Failed Raw Capacity|Implemented|RMI|Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented|RMI|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|RMI|
|Manufacturer|Implemented|RMI|
|Microcode Version|Implemented|RMI|
|Model|Implemented|RMI|
|Name|Implemented|RMI|
|Total Raw Capacity|Implemented|RMI|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|RMI|
|Spare Raw Capacity|Implemented|RMI|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|RMI|Is this a storage virtualization device?
.8+|Storage Node|Cache Size|Implemented|RMI|device cache size in MB
|Memory Size|Gap|RMI|device memory in MB
|Model|Implemented|RMI|
|Name|Implemented|RMI|
|Serial Number|Implemented|RMI|
|State|Implemented|RMI|free text describing the device state
```

```
|UUID|Implemented|RMI|
|Up Time|Implemented|RMI|time in milliseconds
.18+|Storage Pool|Data Allocated Capacity|Gap|RMI|capacity allocated for
data
|Data Used Capacity|Implemented|RMI|
|Dedupe Enabled|Implemented|RMI|Is dedupe enabled on the storage pool
|Include In Dwh Capacity|Implemented|RMI|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|RMI|
|Other Allocated Capacity|Gap|RMI|Capacity allocated for other (not data
and not snapshot)
|Other UsedCapacity (MB)|Implemented|RMI|Any capacity other than data and
snapshot
|Physical Disk Capacity (MB)|Implemented|RMI|used as raw capacity for
storage pool
|Raid Group|Implemented|RMI|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|RMI|ratio to convert from usable capacity
to raw capacity
|Redundancy|Implemented|RMI|Redundancy level
|Status|Implemented|RMI|
|Storage Pool Id|Implemented|RMI|
|Thin Provisioning Supported|Implemented|RMI|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|RMI|
|Total Used Capacity|Implemented|RMI|Total capacity in MB
|Type|Gap|RMI|
|Virtual|Implemented|RMI|Is this a storage virtualization device?
.12+|Volume|Capacity|Implemented|RMI|Snapshot Used capacity in MB
|Disk Type|Not Available|RMI|
|Name|Implemented|RMI|
|Total Raw Capacity|Implemented|RMI|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|RMI|Redundancy level
|Storage Pool Id|Implemented|RMI|
|Thin Provisioned|Implemented|RMI|
|Type|Gap|RMI|
|UUID|Implemented|RMI|
|Used Capacity|Implemented|RMI|
|Virtual|Implemented|RMI|Is this a storage virtualization device?
|Written Capacity|Implemented|RMI|Total capacity written to this volume by
a Host in MB
.3+|Volume Map|LUN|Implemented|RMI|Name of the backend lun
|Storage Port|Implemented|RMI|
|Type|Gap|RMI|
.3+|Volume Mask|Initiator|Implemented|RMI|
```

```
|Storage Port|Implemented|RMI|
|Type|Gap|RMI|
.73+|performance .9+|Disk|IOps Read|Implemented|RMI|Number of read IOps on
the disk
| IOPs Total | Implemented | RMI |
| IOPs Write | Implemented | RMI |
|Throughput Read|Implemented|RMI|
|Throughput Total|Implemented|RMI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|RMI|
|Read Utilization|Implemented|RMI|
|Utilization Total|Implemented|RMI|
|Utilization Write|Implemented|RMI|
.17+|Storage|Cache Hit Ratio Read|Implemented|RMI|
|Cache Hit Ratio Total|Implemented|RMI|
|Cache Hit Ratio Write|Implemented|RMI|
|Failed Raw Capacity|Implemented|RMI|
|Raw Capacity|Implemented|RMI|
|Spare Raw Capacity|Implemented|RMI|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|RMI|
| IOPs other | Implemented | RMI |
| IOps Read | Implemented | RMI | Number of read IOps on the disk
| IOPs Total | Implemented | RMI |
|IOPs Write|Implemented|RMI|
|Latency Read|Implemented|RMI|
|Latency Total|Implemented|RMI|
|Latency Write|Implemented|RMI|
|Throughput Read|Implemented|RMI|
|Throughput Total|Implemented|RMI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|RMI|
.10+|Storage Node|Cache Hit Ratio Total|Implemented|RMI|
| IOps Read | Implemented | RMI | Number of read IOps on the disk
| IOPs Total | Implemented | RMI |
|IOPs Write|Implemented|RMI|
|Latency Read|Implemented|RMI|
|Latency Total|Implemented|RMI|
|Latency Write|Implemented|RMI|
|Throughput Read|Implemented|RMI|
|Throughput Total|Implemented|RMI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|RMI|
.19+|StoragePool Disk|Capacity Provisioned|Implemented|RMI|
|Raw Capacity|Implemented|RMI|
|Total Capacity|Implemented|RMI|
```

```
| Used Capacity | Implemented | RMI |
|Over Commit Capacity Ratio|Implemented|RMI|Reported as a time series
|Capacity Used Ratio|Implemented|RMI|
|Total Data Capacity|Implemented|RMI|
|Data Used Capacity|Implemented|RMI|
|IOps Read|Implemented|RMI|Number of read IOps on the disk
| IOPs Total | Implemented | RMI |
|IOPs Write|Implemented|RMI|
|Other Total Capacity|Implemented|RMI|
|Other Used Capacity|Implemented|RMI|
|Throughput Read|Implemented|RMI|
|Throughput Total|Implemented|RMI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|RMI|
|Read Utilization|Implemented|RMI|
|Utilization Total|Implemented|RMI|
|Utilization Write|Implemented|RMI|
.18+|Volume|Cache Hit Ratio Read|Implemented|RMI|
|Cache Hit Ratio Total|Implemented|RMI|
|Cache Hit Ratio Write|Implemented|RMI|
|Raw Capacity|Implemented|RMI|
|Total Capacity|Implemented|RMI|
|Used Capacity|Implemented|RMI|
|Written Capacity|Implemented|RMI|
|Capacity Used Ratio|Implemented|RMI|
|CapacityRatio Written|Implemented|RMI|
| IOps Read | Implemented | RMI | Number of read IOps on the disk
|IOPs Total|Implemented|RMI|
|IOPs Write|Implemented|RMI|
|Latency Read|Implemented|RMI|
|Latency Total|Implemented|RMI|
|Latency Write|Implemented|RMI|
|Throughput Read|Implemented|RMI|
|Throughput Total|Implemented|RMI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|RMI|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
```

```
|SANtricity API
IRMI
| TCP
|true
|true
|false
|false
|===
<<top, Back to Top>>
== Google Cloud Compute
:description: Support Matrix Asciidoc for Google Cloud Compute
Models and versions supported by this data collector:
<.<|API versions
|v1
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.48+|foundation .7+|Data Store|Capacity|Implemented|HTTPS|Snapshot Used
capacity in MB
|MOID|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Provisioned Capacity|Implemented|HTTPS|
|Virtual Center Ip|Implemented|HTTPS|
|subscription Id|Implemented|HTTPS|
.4+|Server|DataCenter Name|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Virtual Center Ip|Implemented|HTTPS|
.5+|Virtual Disk|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
|subscription Id|Implemented|HTTPS|
```

```
.17+|VirtualMachine|Guest State|Implemented|HTTPS|
|DataStore OID|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
|IPs|Implemented|HTTPS|
|MOID|Implemented|HTTPS|
|Memory|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|OS|Implemented|HTTPS|
|Power State|Implemented|HTTPS|
|State Change Time|Implemented|HTTPS|
|Processors|Implemented|HTTPS|
|Provisioned Capacity|Implemented|HTTPS|
|Instance Type|Implemented|HTTPS|
|Launch Time|Implemented|HTTPS|
|public Ips|Implemented|HTTPS|
|subscription Id|Implemented|HTTPS|
.3+|VirtualMachine Disk|OID|Implemented|HTTPS|
|VirtualDisk OID|Implemented|HTTPS|
|VirtualMachine OID|Implemented|HTTPS|
.5+|Host|Host OS|Implemented|HTTPS|
|IPs|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.7+|Info|Api Description|Implemented|HTTPS|
|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.27+|performance .4+|Data Store|Capacity Provisioned|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio|Implemented|HTTPS|Reported as a time series
|Capacity Used Ratio|Implemented|HTTPS|
.6+|Virtual Disk|IOps Read|Implemented|HTTPS|Number of read IOps on the
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.17+|vm|Total Capacity|Implemented|HTTPS|
|Total CPU Utilization|Implemented|HTTPS|
```

```
| IOps Read | Implemented | HTTPS | Number of read IOps on the disk
|diskIops.total|Implemented|HTTPS|
|Disk IOPs write|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Disk Throughput Read|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|total disk throughput read
|Disk Throughput Write|Implemented|HTTPS|
| IP Throughput Read | Implemented | HTTPS |
|Throughput total|Implemented|HTTPS|IP throughput total
|ipThroughput.write|Implemented|HTTPS|
|Total Memory Utilization|Implemented|HTTPS|
|swapRate.inRate|Implemented|HTTPS|
|Swap Rate|Implemented|HTTPS|
|Total Swap Rate|Implemented|HTTPS|
|Schedule wait time|Implemented|HTTPS|Waiting to be scheduled time in
percent
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Google Compute Platform REST API
|HTTPS
|443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== HDS HCP (HTTPS)
:description: Support Matrix Asciidoc for HDS HCP (HTTPS)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
```

```
| Hitachi Content Platform
19.3.7.2
9.5.0.121
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.41+|foundation .3+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
.11+|Internal Volume|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on
the storage pool
|Internal Volume Id|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Total Used Capacity (MB)|Implemented|HTTPS|place holder for the used
capacity as read from the device
|Type|Gap|HTTPS|
.3+|QTree|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Type|Gap|HTTPS|
.10+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.2+|Storage Node|Name|Implemented|HTTPS|
```

```
|UUID|Implemented|HTTPS|
.12+|Storage Pool|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ
to cotnrol which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Soft Limit (MB)|Implemented|HTTPS|logical volume size that is defined
during volume creation or resizing operations
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.19+|performance .1+|Internal Volume|Total Objects|Implemented||
.7+|Storage|Failed Raw Capacity|Implemented||
|Raw Capacity|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|StoragePools Capacity|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.4+|Storage Node|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Utilization Total|Implemented||
.7+|StoragePool Disk|Total Capacity|Implemented||
|Capacity Used Ratio|Implemented||
|Capacity Provisioned|Implemented||
|Used Capacity|Implemented||
|Raw Capacity|Implemented||
|Capacity Soft Limit|Implemented||
|Over Commit Capacity Ratio|Implemented||Reported as a time series
|===
Management APIs used by this data collector:
|===
```

```
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
| HDS HCP REST API
|HTTPS
|HTTPS
19090
|true
|true
|true
|true
|===
<<top, Back to Top>>
== HiCommand Device Manager
:description: Support Matrix Asciidoc for HiCommand Device Manager
Models and versions supported by this data collector:
|===
<.<|API versions <.<|Models <.<|Firmware versions
|7.6.1
8.7.7
8.8.1
8.8.3
8.8.5
|DF850MH
DF850S
008MH
HM850
P9500
RAID700
RAID800
VSP5000
XP24000
XP7
|0983/A-H
0988/H-S
DKC:60-08-22
DKC:60-08-65
DKC:70-06-46
```

```
DKC:70-06-67-00/00
DKC:80-06-80
DKC:80-06-82-00/00
DKC:80-06-86-00/00
DKC:80-06-87
DKC:80-06-88-00/00
DKC:80-06-91
DKC:80-06-91-00/00
DKC:80-06-93-00/00
DKC:83-05-45-40/00
DKC:83-05-45-60/00
DKC:83-05-46-60/00
DKC:83-05-47-60/00
DKC:83-05-48-40/00
DKC:83-05-48-60/00
DKC:88-08-08-60/00
DKC:88-08-09-60/00
DKC:90-08-81-00/00
DKC:90-08-83-00/01
SVP:60-08-21/00
SVP:60-08-54/00
SVP:70-06-32/00
SVP:70-06-51/00
SVP:80-06-76/02
SVP:80-06-78/00
SVP:80-06-81/00
SVP:80-06-82/00
SVP:80-06-83/00
SVP:80-06-86/00
SVP:80-06-88/00
SVP:83-05-49-40/00
SVP:83-05-49-60/00
SVP:83-05-50-60/00
SVP:83-05-51-60/00
SVP:83-05-52-40/00
SVP:83-05-52-60/00
SVP:88-08-10-60/00
SVP:88-08-11-60/00
SVP:90-08-81/00
SVP:90-08-83/00
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
```

```
.113+|foundation .14+|Disk|Capacity (GB)|Implemented|HDS XML API|use
capacity
|Disk Id|Implemented|HDS XML API|Uniquely identifies this disk in the
|Group|Implemented|HDS XML API|
|Location|Gap|HDS XML API|Where this disk is physically located in the
array
|Model|Implemented|HDS XML API|
|Name|Implemented|HDS XML API|
|Role|Implemented|HDS XML API|
|Role Enum|Implemented|HDS XML API|enum for disk role
|Serial Number|Implemented|HDS XML API|
|Speed|Implemented|HDS XML API|Speed of disk (RPM)
|Type|Gap|HDS XML API|
|Type Enum|Implemented|HDS XML API|enum for disk type
|Vendor|Implemented|HDS XML API|
|Vendor Id|Implemented|HDS XML API|
.11+|Disk Group|Capacity|Implemented|HDS XML API|Snapshot Used capacity in
|DiskGroup Id|Implemented|HDS XML API|unique id of the disk group
|Disk Type|Not Available|HDS XML API|
|Name|Implemented|HDS XML API|
|Physical Disk Capacity (MB)|Implemented|HDS XML API|used as raw capacity
for storage pool
|Redundancy|Implemented|HDS XML API|Redundancy level
|Status|Implemented|HDS XML API|
|Used Capacity|Implemented|HDS XML API|
|Vendor DiskGroup Type|Implemented|HDS XML API|vendor's designation of the
disk group type
|Vendor Tier|Implemented|HDS XML API|Vendor Specific Tier Name
|Virtual|Implemented|HDS XML API|Is this a storage virtualization device?
.6+|Info|Api Name|Implemented|HDS XML API|
|Api Version|Implemented|HDS XML API|
|DataSource Name|Implemented|HDS XML API|Info
|Date|Implemented|HDS XML API|
|Originator ID|Implemented|HDS XML API|
|Originator Key|Implemented|HDS XML API|
.3+|Network Pipe|Source Id|Implemented|HDS XML API|
|Target Id|Implemented|HDS XML API|
|Bidirectional|Implemented|HDS XML API|
.3+|Protocol EndPoint|ID|Implemented|HDS XML API|
|Name|Implemented|HDS XML API|
|Storage Ip|Implemented|HDS XML API|
.13+|Storage|Display IP|Implemented|HDS XML API|
|Failed Raw Capacity|Implemented|HDS XML API|Raw capapcity of failed disks
```

```
(sum of all disks that are failed)
|Family|Implemented|HDS XML API|The storage Family could be Clariion,
Symmetrix, et al
|IP|Implemented|HDS XML API|
|Manage URL|Implemented|HDS XML API|
|Manufacturer|Implemented|HDS XML API|
|Microcode Version|Implemented|HDS XML API|
|Model|Implemented|HDS XML API|
|Name|Implemented|HDS XML API|
|Total Raw Capacity|Implemented|HDS XML API|Total raw capacity (sum of all
disks on the array)
|Serial Number|Implemented|HDS XML API|
|Spare Raw Capacity|Implemented|HDS XML API|Raw capapcity of spare disks
(sum of all disks that are spare)
|Virtual|Implemented|HDS XML API|Is this a storage virtualization device?
.2+|Storage Node|Name|Implemented|HDS XML API|
|UUID|Implemented|HDS XML API|
.18+|Storage Pool|Auto Tiering|Implemented|HDS XML API|indicates if this
storagepool is participating in auto tiering with other pools
|Compression Enabled|Implemented|HDS XML API|Is compression enabled on the
storage pool
|Compression Savings|Implemented|HDS XML API|ratio of compression savings
in percentage
|Dedupe Enabled|Implemented|HDS XML API|Is dedupe enabled on the storage
pool
|Include In Dwh Capacity|Implemented|HDS XML API|A way from ACQ to cotnrol
which stroage pools are interesting in DWH Capacity
|Name|Implemented|HDS XML API|
|Physical Disk Capacity (MB)|Implemented|HDS XML API|used as raw capacity
for storage pool
|Raid Group|Implemented|HDS XML API|indicates whether this storagePool is
a raid group
|Raw to Usable Ratio|Implemented|HDS XML API|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HDS XML API|Redundancy level
|Soft Limit (MB)|Implemented|HDS XML API|logical volume size that is
defined during volume creation or resizing operations
|Status|Implemented|HDS XML API|
|Storage Pool Id|Implemented|HDS XML API|
|Thin Provisioning Supported|Implemented|HDS XML API|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HDS XML API|
|Total Used Capacity|Implemented|HDS XML API|Total capacity in MB
|Type|Gap|HDS XML API|
|Virtual|Implemented|HDS XML API|Is this a storage virtualization device?
.9+|Storage Synchronization|Mode|Implemented|HDS XML API|
```

```
|Mode Enum|Implemented|HDS XML API|
|Source Storage|Implemented|HDS XML API|
|Source Volume|Implemented|HDS XML API|
|State|Implemented|HDS XML API|free text describing the device state
|State Enum|Implemented|HDS XML API|
|Target Storage|Implemented|HDS XML API|
|Target Volume|Implemented|HDS XML API|
|Technology|Implemented|HDS XML API|technology which causes storage
efficiency changed
.16+|Volume|AutoTier Policy Identifier|Implemented|HDS XML API|Dynamic
Tier Policy identifier
|Auto Tiering|Implemented|HDS XML API|indicates if this storagepool is
participating in auto tiering with other pools
|Capacity|Implemented|HDS XML API|Snapshot Used capacity in MB
|Junction Path|Implemented|HDS XML API|
|Mainframe|Implemented|HDS XML API|indicates if this volume is a Mainframe
Volume
|Meta|Implemented|HDS XML API|Flag saying whether this volume is a meta
volume with memeber or not. Meta volumes will have DiskGroup empty!
|Name|Implemented|HDS XML API|
|Total Raw Capacity|Implemented|HDS XML API|Total raw capacity (sum of all
disks on the array)
|Redundancy|Implemented|HDS XML API|Redundancy level
|Replica Source|Implemented|HDS XML API|
|Replica Target|Implemented|HDS XML API|
|Storage Pool Id|Implemented|HDS XML API|
|Thin Provisioned|Implemented|HDS XML API|
|Type|Gap|HDS XML API|
|Used Capacity|Implemented|HDS XML API|
|Virtual|Implemented|HDS XML API|Is this a storage virtualization device?
.4+|Volume Map|LUN|Implemented|HDS XML API|Name of the backend lun
|Masking Required|Implemented|HDS XML API|
|Protocol Controller|Implemented|HDS XML API|
|Storage Port|Implemented|HDS XML API|
.3+|Volume Mask|Initiator|Implemented|HDS XML API|
|Protocol Controller|Implemented|HDS XML API|
|Storage Port|Implemented|HDS XML API|
.7+|Volume Member|Name|Implemented|HDS XML API|
|Storage Pool Id|Implemented|HDS XML API|
|Rank|Implemented|HDS XML API|
|Cylinders|Implemented|HDS XML API|
|Capacity|Implemented|HDS XML API|Snapshot Used capacity in MB
|Total Raw Capacity|Implemented|HDS XML API|Total raw capacity (sum of all
disks on the array)
|Used Capacity|Implemented|HDS XML API|
.4+|WWN Alias|Host Aliases|Implemented|HDS XML API|
```

```
|Object Type|Implemented|HDS XML API|
|Source|Implemented|HDS XML API|
|WWN|Implemented|HDS XML API|
.47+|performance .7+|Disk|IOps Read|Implemented|Export/CLI|Number of read
IOps on the disk
| IOPs Total | Implemented | Export / CLI |
| IOPs Write | Implemented | Export / CLI |
|Throughput Read|Implemented|Export/CLI|
|Throughput Total|Implemented|Export/CLI|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|Export/CLI|
|Utilization Total|Implemented|Export/CLI|
.19+|Storage|Cache Hit Ratio Read|Implemented|Export/CLI|
|Cache Hit Ratio Total|Implemented|Export/CLI|
|Cache Hit Ratio Write|Implemented|Export/CLI|
|Cache Utilization Total|Implemented|Export/CLI|
|Failed Raw Capacity|Implemented|Export/CLI|
|Raw Capacity|Implemented|Export/CLI|
|Spare Raw Capacity|Implemented|Export/CLI|Raw capapcity of spare disks
(sum of all disks that are spare)
|StoragePools Capacity|Implemented|Export/CLI|
| IOPs other | Implemented | Export / CLI |
|IOps Read|Implemented|Export/CLI|Number of read IOps on the disk
| IOPs Total | Implemented | Export / CLI |
| IOPs Write | Implemented | Export / CLI |
|Latency Read|Implemented|Export/CLI|
|Latency Total|Implemented|Export/CLI|
|Latency Write|Implemented|Export/CLI|
|Throughput Read|Implemented|Export/CLI|
|Throughput Total|Implemented|Export/CLI|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|Export/CLI|
|Write Pending|Implemented|Export/CLI|total write pending
.2+|Storage Node|Throughput Total|Implemented|Export/CLI|Average disk
total rate (read and write across all disks) in MB/s
| IOPs Total | Implemented | Export / CLI |
.7+|StoragePool Disk|Total Capacity|Implemented|Export/CLI|
|Capacity Used Ratio|Implemented|Export/CLI|
|Capacity Provisioned|Implemented|Export/CLI|
|Used Capacity|Implemented|Export/CLI|
|Raw Capacity|Implemented|Export/CLI|
|Capacity Soft Limit|Implemented|Export/CLI|
|Over Commit Capacity Ratio|Implemented|Export/CLI|Reported as a time
series
.12+|Volume|Latency Total|Implemented|Export/CLI|
| IOps Read | Implemented | Export / CLI | Number of read IOps on the disk
```

```
|Latency Read|Implemented|Export/CLI|
|Cache Hit Ratio Read|Implemented|Export/CLI|
| IOPs Write | Implemented | Export / CLI |
|Cache Hit Ratio Total|Implemented|Export/CLI|
|Cache Hit Ratio Write|Implemented|Export/CLI|
|Throughput Read|Implemented|Export/CLI|
|Throughput Write|Implemented|Export/CLI|
|Throughput Total|Implemented|Export/CLI|Average disk total rate (read and
write across all disks) in MB/s
| IOPs Total | Implemented | Export / CLI |
|Latency Write|Implemented|Export/CLI|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Export utility (USPV) / SNM CLI (AMS)
|Export/CLI
|false
|false
|false
|false
|HiCommand Device Manager XML API
|HDS XML API
|HTTP/HTTPS
12001
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Hitachi Ops Center
```

```
:description: Support Matrix Asciidoc for Hitachi Ops Center
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
|VSP 5100
VSP 5500
VSP F1500
VSP F600
VSP G800
|80-06-92-00/00:01-65-03/05
83-05-46-60/00:01-65-03/05
83-05-47-40/00:01-65-03/05
83-05-48-40/00:01-65-03/05
90-08-81-00/00:01-65-03/05
90-08-82-00/00:01-65-03/05
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.108+|foundation .15+|Disk|Capacity (GB)|Implemented||use capacity
|Disk Id|Implemented||Uniquely identifies this disk in the array
|Group|Implemented||
|Location|Gap||Where this disk is physically located in the array
|Model|Implemented||
|Name|Implemented||
|Potential Transfer Rate|Implemented||
|Role|Implemented||
|Role Enum|Implemented||enum for disk role
|Seek Time|Implemented||
|Serial Number|Implemented||
|Speed|Implemented||Speed of disk (RPM)
|Status Enum|Implemented||enum for disk status
|Type Enum|Implemented||enum for disk type
|Vendor|Implemented||
.10+|Disk Group|Capacity|Implemented||Snapshot Used capacity in MB
|DiskGroup Id|Implemented||unique id of the disk group
|Disk Type|Not Available||
|Name|Implemented||
|Physical Disk Capacity (MB)|Implemented||used as raw capacity for
storage pool
|Redundancy|Implemented||Redundancy level
```

```
|Status|Implemented||
|Used Capacity|Implemented||
|Vendor DiskGroup Type|Implemented||vendor's designation of the disk group
type
|Virtual|Implemented||Is this a storage virtualization device?
.4+|ISCSI Network Portal|OID|Implemented||
|IP|Implemented||
|Nic|Implemented||
|Listening Port|Implemented||
.3+|ISCSI Network Portal Group|OID|Implemented||
|Portal Group Name|Implemented||
|Portal Group Tag|Implemented||
.3+|ISCSI Node|OID|Implemented||
|Node Name|Implemented||
|Type|Gap||
.2+|ISCSI Node Map|OID|Implemented||
|Portal Group OID|Implemented||
.7+|ISCSI Session|OID|Implemented||
|Initiator OID|Implemented||
|Portal Group OID|Implemented||
|Number Of Connections|Implemented||
|Max Connections|Implemented||
|Initiator Ips|Implemented||
|Security|Implemented||
.4+|Info|DataSource Name|Implemented||Info
|Date|Implemented||
|Originator ID|Implemented||
|Originator Key|Implemented||
.14+|Storage|Display IP|Implemented||
|Failed Raw Capacity|Implemented||Raw capapcity of failed disks (sum of
all disks that are failed)
|Family|Implemented||The storage Family could be Clariion, Symmetrix, et
al
|IP|Implemented||
|Manage URL|Implemented||
|Manufacturer|Implemented||
|Microcode Version|Implemented||
|Model|Implemented||
|Name|Implemented||
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Serial Number|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|SupportActive Active|Implemented||Specified if the storage supports
active-active configurations
```

```
|Virtual|Implemented||Is this a storage virtualization device?
.2+|Storage Node|Name|Implemented||
|UUID|Implemented||
.16+|Storage Pool|Dedupe Enabled|Implemented||Is dedupe enabled on the
storage pool
|Include In Dwh Capacity|Implemented||A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented||
|Physical Disk Capacity (MB)|Implemented||used as raw capacity for
storage pool
|Raid Group|Implemented||indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
|Redundancy|Implemented||Redundancy level
|Soft Limit (MB)|Implemented||logical volume size that is defined during
volume creation or resizing operations
|Status|Implemented||
|Storage Pool Id|Implemented||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented||
|Total Used Capacity|Implemented||Total capacity in MB
|Type|Gap||
|Virtual|Implemented||Is this a storage virtualization device?
|Encrypted|Implemented||
.9+|Storage Synchronization|Source Volume|Implemented||
|Target Volume|Implemented||
|Mode|Implemented||
|Mode Enum|Implemented||
|State|Implemented||free text describing the device state
|State Enum|Implemented||
|Source Storage|Implemented||
|Target Storage|Implemented||
|Technology|Implemented||technology which causes storage efficiency
changed
.10+|Volume|Capacity|Implemented||Snapshot Used capacity in MB
|Junction Path|Implemented||
|Name|Implemented||
|Protection Type|Implemented||
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Storage Pool Id|Implemented||
|Thin Provisioned|Implemented||
|Type|Gap||
|Used Capacity|Implemented||
```

```
|Compression Enabled|Implemented||
.5+|Volume Map|LUN|Implemented||Name of the backend lun
|Masking Required|Implemented||
|Protocol Controller|Implemented||
|Storage Port|Implemented||
|Type|Gap||
.4+|Volume Mask|Initiator|Implemented||
|Protocol Controller|Implemented||
|Storage Port|Implemented||
|Type|Gap||
.28+|performance .6+|Disk|IOps Read|Implemented||Number of read IOps on
the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.13+|Storage|Failed Raw Capacity|Implemented||
|Raw Capacity|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|StoragePools Capacity|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.2+|Storage Node|Throughput Total|Implemented||Average disk total rate
(read and write across all disks) in MB/s
|IOPs Total|Implemented||
.7+|StoragePool Disk|Total Capacity|Implemented||
|Capacity Used Ratio|Implemented||
|Capacity Provisioned|Implemented||
|Used Capacity|Implemented||
|Raw Capacity|Implemented||
|Capacity Soft Limit|Implemented||
|Over Commit Capacity Ratio|Implemented||Reported as a time series
|===
```

```
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Hitachi Ops Center REST API
| HTTPS
| HTTPS
|443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== HDS HNAS (CLI)
:description: Support Matrix Asciidoc for HDS HNAS (CLI)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
IG600
G800
HNAS 4080
HNAS 4100
N800
113.9.6918.05
14.5.7413.01
14.6.7520.04
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.87+|foundation .6+|Disk Group|DiskGroup Id|Implemented|SSH|unique id of
the disk group
|Name|Implemented|SSH|
```

```
|Virtual|Implemented|SSH|Is this a storage virtualization device?
|Vendor Tier|Implemented|SSH|Vendor Specific Tier Name
|Capacity|Implemented|SSH|Snapshot Used capacity in MB
|Used Capacity|Implemented|SSH|
.5+|File Share|Name|Implemented|SSH|
|Path|Implemented|SSH|path of the fileShare
|Qtree Id|Implemented|SSH|unique id of the qtree
|Is InternalVolume|Implemented|SSH|whether the file share represents an
internal volume (netapp volume) or is it a qtree/folder within the
internal volume
|Is Shared|Implemented|SSH|whether this fileShare has any shares
associated with it
.4+|Info|DataSource Name|Implemented|SSH|Info
|Originator ID|Implemented|SSH|
|Date|Implemented|SSH|
|Originator Key|Implemented|SSH|
.16+|Internal Volume|Internal Volume Id|Implemented|SSH|
|Name|Implemented|SSH|
|Storage Pool Id|Implemented|SSH|
|Type|Gap|SSH|
|Thin Provisioned|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Status|Implemented|SSH|
|Virtual Storage|Implemented|SSH|Owning virtual storage (vfiler)
|Snapshot Used Capacity|Implemented|SSH|
|Data Used Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Total Used Capacity (MB)|Implemented|SSH|place holder for the used
capacity as read from the device
|Total Allocated Capacity|Implemented|SSH|
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Dedupe Savings|Implemented|SSH|ratio of dedupe savings in percentage
.6+|QTree|Qtree Id|Implemented|SSH|unique id of the qtree
|Name|Implemented|SSH|
|Type|Gap|SSH|
|Quota HardCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk
space, allowed for the quota target
|Quota UsedCapacity|Implemented|SSH|Space in MB currently used
|Quota SoftCapacity Limit (MB)|Implemented|SSH|Maximum amount of disk
space, allowed for the quota target
.10+|Quota|Quota Id|Implemented|SSH|unique id of the quota
|Type|Gap|SSH|
|Internal Volume Id|Implemented|SSH|
```

```
|Qtree Id|Implemented|SSH|unique id of the qtree
|Hard File Limit|Implemented|SSH|max number of files allowed for the quota
target
|Soft File Limit|Implemented|SSH|Max number of files allowed for the quota
| Hard Capacity Limit (MB) | Implemented | SSH | max amount of disk space,
allowed for the quota target (Hard limit)
|Used Files|Implemented|SSH|Number of files currently used
|Used Capacity|Implemented|SSH|
|Soft Capacity Limit (MB)|Implemented|SSH|Maximum amount of disk space,
allowed for the quota target
.4+|Share|Name|Implemented|SSH|
|Protocol|Implemented|SSH|enum for share protocol
| IP Interfaces | Implemented | SSH | comma separated list of IP addresses on
which this share is exposed
|Description|Implemented|SSH|
.2+|Share Initiator|Initiator|Implemented|SSH|
|Permission|Implemented|SSH|Permissions for this particular share
.14+|Storage|IP|Implemented|SSH|
|Display IP|Implemented|SSH|
|Name|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Model|Implemented|SSH|
|Family|Implemented|SSH|The storage Family could be Clariion, Symmetrix,
et al
|Serial Number|Implemented|SSH|
|Microcode Version|Implemented|SSH|
|Virtual|Implemented|SSH|Is this a storage virtualization device?
|SupportActive Active|Implemented|SSH|Specified if the storage supports
active-active configurations
|Cpu Count|Implemented|SSH|Cpu Count of the storage
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|Failed Raw Capacity|Implemented|SSH|Raw capapcity of failed disks (sum of
all disks that are failed)
.7+|Storage Node|UUID|Implemented|SSH|
|Name|Implemented|SSH|
|Version|Implemented|SSH|software version
|Serial Number|Implemented|SSH|
|State|Implemented|SSH|free text describing the device state
|Memory Size|Gap|SSH|device memory in MB
|Processors Count|Implemented|SSH|device CPU
.13+|Storage Pool|Storage Pool Id|Implemented|SSH|
|Name|Implemented|SSH|
```

```
|Type|Gap|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Include In Dwh Capacity|Implemented|SSH|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Virtual|Implemented|SSH|Is this a storage virtualization device?
|Raid Group|Implemented|SSH|indicates whether this storagePool is a raid
group
|Snapshot Used Capacity|Implemented|SSH|
|Data Used Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Total Allocated Capacity|Implemented|SSH|
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
| HDS HNAS CLI
|SSH
|SSH
122
|true
|true
|true
|true
|===
<<top, Back to Top>>
== HPE Nimble / Alletra 6000 Storage
:description: Support Matrix Asciidoc for HPE Nimble / Alletra 6000
Storage
Models and versions supported by this data collector:
<.<|API versions <.<|Models <.<|Firmware versions
```

```
lv1
16030
AF1000
AF20Q
AF3000
AF40
AF5000
CS1000
CS300
CS3000
CS500
CS5000
HF20
HF20H
HF40
HF60
|5.0.10.0-742719-opt
5.0.7.0-604814-opt
5.0.8.0-677726-opt
5.2.1.1000-1017822-opt
5.2.1.400-796142-opt
5.2.1.600-841103-opt
5.2.1.700-882343-opt
5.2.1.800-930936-opt
5.2.1.900-1003439-opt
6.0.0.300-956221-opt
6.0.0.400-991061-opt
6.1.1.200-1020304-opt
6.1.1.300-1028597-opt
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.104+|foundation .17+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Group|Implemented|HTTPS|
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Potential Transfer Rate|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Seek Time|Implemented|HTTPS|
```

```
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type|Gap|HTTPS|
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
.4+|ISCSI Network Portal|IP|Implemented|HTTPS|
|Listening Port|Implemented|HTTPS|
|Nic|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.3+|ISCSI Network Portal Group|OID|Implemented|HTTPS|
|Portal Group Name|Implemented|HTTPS|
|Portal Group Tag|Implemented|HTTPS|
.4+|ISCSI Node|Host Aliases|Implemented|HTTPS|
|Node Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
.7+|ISCSI Session|Initiator Ips|Implemented|HTTPS|
|Initiator OID|Implemented|HTTPS|
|Max Connections|Implemented|HTTPS|
|Number Of Connections|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
|Security|Implemented|HTTPS|
.6+|Info|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.14+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
```

```
all disks that are spare)
|SupportActive Active|Implemented|HTTPS|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.6+|Storage Node|Name|Implemented|HTTPS|
|Partner Node UUID|Implemented|HTTPS|HA pair's UUID
|Serial Number|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
|Parent Serial Number|Implemented|HTTPS|
.18+|Storage Pool|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|HTTPS|capacity allocated for data
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
|Reserved Capacity|Implemented|HTTPS|Reserved Capacity in MB
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.12+|Volume|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|HTTPS|Redundancy level
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
```

```
|Compression Enabled|Implemented|HTTPS|
|Encrypted|Implemented|HTTPS|
.5+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Masking Required|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|WWN Alias|Host Aliases|Implemented|HTTPS|
|Object Type|Implemented|HTTPS|
|Source|Implemented|HTTPS|
|WWN|Implemented|HTTPS|
.35+|performance .14+|Storage|Failed Raw Capacity|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|HTTPS|
| IOPs other | Implemented | HTTPS |
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.6+|StoragePool Disk|IOps Read|Implemented|HTTPS|Number of read IOps on
the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.15+|Volume|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
|Total Compression Savings|Implemented|HTTPS|
|Compression Savings Space|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
```

```
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
| HP Nimble REST API
|HTTPS
| HTTPS
|5392
|true
|false
Itrue
|true
|===
<<top, Back to Top>>
== Huawei OceanStor (REST/HTTPS)
:description: Support Matrix Asciidoc for Huawei OceanStor (REST/HTTPS)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
15300 V5
5500 V3
5500 V5
5800 V3
Dorado 5000 V6 SAS
Dorado 6000 V3
```

```
Dorado 6000 V6 NVMe
IV300R001C01
V300R002C10
V300R006C20
V300R006C50
V500R007C10
V500R007C30
V600R003C00
V600R005C03
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.126+|foundation .14+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type|Gap|HTTPS|
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
.5+|File Share|Name|Implemented|HTTPS|
|Path|Implemented|HTTPS|path of the fileShare
|Qtree Id|Implemented|HTTPS|unique id of the qtree
| Is InternalVolume | Implemented | HTTPS | whether the file share represents an
internal volume (netapp volume) or is it a qtree/folder within the
internal volume
|Is Shared|Implemented|HTTPS|whether this fileShare has any shares
associated with it
.4+|ISCSI Network Portal|OID|Implemented|HTTPS|
|IP|Implemented|HTTPS|
|Listening Port|Implemented|HTTPS|
|Nic|Implemented|HTTPS|
.3+|ISCSI Network Portal Group|OID|Implemented|HTTPS|
|Portal Group Name|Implemented|HTTPS|
|Portal Group Tag|Implemented|HTTPS|
.3+|ISCSI Node|OID|Implemented|HTTPS|
```

```
|Node Name|Implemented|HTTPS|
|Type|Gap|HTTPS|
.2+|ISCSI Node Map|OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
.7+|ISCSI Session|OID|Implemented|HTTPS|
|Initiator OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
|Number Of Connections|Implemented|HTTPS|
|Max Connections|Implemented|HTTPS|
|Initiator Ips|Implemented|HTTPS|
|Security|Implemented|HTTPS|
.4+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.13+|Internal Volume|Internal Volume Id|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|Snapshot Used Capacity|Implemented|HTTPS|
|Data Used Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Total Allocated Capacity|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
.7+|QTree|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Name|Implemented|HTTPS|
|Status|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Quota HardCapacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk
space, allowed for the quota target
|Quota SoftCapacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk
space, allowed for the quota target
|Quota UsedCapacity|Implemented|HTTPS|Space in MB currently used
.11+|Quota|Quota Id|Implemented|HTTPS|unique id of the quota
|Type|Gap|HTTPS|
|Internal Volume Id|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Hard File Limit|Implemented|HTTPS|max number of files allowed for the
quota target
|Soft File Limit|Implemented|HTTPS|Max number of files allowed for the
```

```
quota target
| Hard Capacity Limit (MB) | Implemented | HTTPS | max amount of disk space,
allowed for the quota target (Hard limit)
|Soft Capacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk space,
allowed for the quota target
|Used Files|Implemented|HTTPS|Number of files currently used
|Used Capacity|Implemented|HTTPS|
|User/Group Target|Implemented|HTTPS|user/group target this quota refers
.4+|Share|Name|Implemented|HTTPS|
|Protocol|Implemented|HTTPS|enum for share protocol
|Description|Implemented|HTTPS|
| IP Interfaces | Implemented | HTTPS | comma separated list of IP addresses on
which this share is exposed
.2+|Share Initiator|Initiator|Implemented|HTTPS|
|Permission|Implemented|HTTPS|Permissions for this particular share
.14+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|HTTPS|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.3+|Storage Node|Name|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
|Version|Implemented|HTTPS|software version
.12+|Storage Pool|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ
to cotnrol which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
```

```
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.10+|Volume|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Junction Path|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|HTTPS|Redundancy level
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.4+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.33+|performance .7+|Disk|IOps Read|Implemented|HTTPS|Number of read IOps
on the disk
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Write|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
| IOPs Total | Implemented | HTTPS |
|Utilization Total|Implemented|HTTPS|
.10+|Storage Node|Latency Total|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
|Latency Read|Implemented|HTTPS|
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Write|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
| IOPs Total | Implemented | HTTPS |
```

```
|Latency Write|Implemented|HTTPS|
| Utilization Total | Implemented | HTTPS |
.16+|Volume|Cache Hit Ratio Read|Implemented|HTTPS|
|Cache Hit Ratio Total|Implemented|HTTPS|
|Cache Hit Ratio Write|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
| Huawei OceanStor REST API
IHTTPS
|HTTPS
18088
|true
|true
|true
|true
| Huawei OceanStor Performance REST API
| HTTPS
|HTTPS
18088
|true
|false
```

```
Itrue
Itrue
|===
<<top, Back to Top>>
== IBM Cleversafe
:description: Support Matrix Asciidoc for IBM Cleversafe
|===
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.57+|foundation .10+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
.3+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
.14+|Internal Volume|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on
the storage pool
|Internal Volume Id|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Protection Type|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
```

```
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Comment|Gap|HTTPS|state: free text comment describing the svm
.3+|QTree|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Type|Gap|HTTPS|
.10+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
|IP|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.6+|Storage Node|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Serial Number|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
|Version|Implemented|HTTPS|software version
|ManagementIp Addresses|Implemented|HTTPS|
.11+|Storage Pool|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ
to cotnrol which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|===
```

```
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
| IBM Cleversafe REST API
|HTTPS
|HTTPS
1443
ltrue
|true
|true
Itrue
|===
<<top, Back to Top>>
== IBM DS 8K (DSCLI)
:description: Support Matrix Asciidoc for IBM DS 8K (DSCLI)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
12107-951
2107-961
2107-985
2107-996
17.6.31.4250
7.7.51.1400
7.8.57.18
7.9.21.91
7.9.32.126
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.67+|foundation .14+|Disk|Capacity (GB)|Implemented|DSNI|use capacity
```

```
|Disk Id|Implemented|DSNI|Uniquely identifies this disk in the array
|Group|Implemented|DSNI|
|Location|Gap|DSNI|Where this disk is physically located in the array
|Model|Implemented|DSNI|
|Name|Implemented|Manual Entry|
|Role|Implemented|DSNI|
|Role Enum|Implemented|DSNI|enum for disk role
|Speed|Implemented|DSNI|Speed of disk (RPM)
|Status|Implemented|DSNI|
|Type|Gap|DSNI|
|Type Enum|Implemented|DSNI|enum for disk type
|Vendor|Implemented|DSNI|
|Vendor Id|Implemented|DSNI|
.4+|Info|DataSource Name|Implemented|DSNI|Info
|Date|Implemented|DSNI|
|Originator ID|Implemented|DSNI|
|Originator Key|Implemented|DSNI|
.13+|Storage|Display IP|Implemented|DSNI|
|Failed Raw Capacity|Implemented|DSNI|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|DSNI|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|DSNI|
|Manage URL|Implemented|DSNI|
|Manufacturer|Implemented|DSNI|
|Microcode Version|Implemented|DSNI|
|Model|Implemented|DSNI|
|Name|Implemented|Manual Entry|
|Total Raw Capacity|Implemented|DSNI|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|DSNI|
|Spare Raw Capacity|Implemented|DSNI|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|DSNI|Is this a storage virtualization device?
.16+|Storage Pool|Data Allocated Capacity|Gap|DSNI|capacity allocated for
data
|Data Used Capacity|Implemented|DSNI|
|Dedupe Enabled|Implemented|DSNI|Is dedupe enabled on the storage pool
|Include In Dwh Capacity|Implemented|DSNI|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|Manual Entry|
|Physical Disk Capacity (MB)|Implemented|DSNI|used as raw capacity for
storage pool
|Raid Group|Implemented|DSNI|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|DSNI|ratio to convert from usable
```

```
capacity to raw capacity
|Redundancy|Implemented|DSNI|Redundancy level
|Status|Implemented|DSNI|
|Storage Pool Id|Implemented|DSNI|
|Thin Provisioning Supported|Implemented|DSNI|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|DSNI|
|Total Used Capacity|Implemented|DSNI|Total capacity in MB
|Type|Gap|DSNI|
|Virtual|Implemented|DSNI|Is this a storage virtualization device?
.9+|Volume|Capacity|Implemented|DSNI|Snapshot Used capacity in MB
|Disk Type|Not Available|DSNI|
|Junction Path|Implemented|DSNI|
|Name|Implemented|Manual Entry|
|Total Raw Capacity|Implemented|DSNI|Total raw capacity (sum of all disks
on the array)
|Storage Pool Id|Implemented|DSNI|
|Thin Provisioned|Implemented|DSNI|
|Type|Gap|DSNI|
|Used Capacity|Implemented|DSNI|
.3+|Volume Map|LUN|Implemented|DSNI|Name of the backend lun
|Protocol Controller|Implemented|DSNI|
|Storage Port|Implemented|DSNI|
.3+|Volume Mask|Initiator|Implemented|DSNI|
|Protocol Controller|Implemented|DSNI|
|Storage Port|Implemented|DSNI|
.5+|WWN Alias|Host Aliases|Implemented|DSNI|
|Host OS|Implemented|DSNI|
|Object Type|Implemented|DSNI|
|Source|Implemented|DSNI|
|WWN|Implemented|DSNI|
.53+|performance .9+|Disk|Utilization Write|Implemented|DSNI|
| IOps Read | Implemented | DSNI | Number of read IOps on the disk
|Read Utilization|Implemented|DSNI|
| IOPs Write | Implemented | DSNI |
|Throughput Read|Implemented|DSNI|
|Throughput Write|Implemented|DSNI|
|Throughput Total|Implemented|DSNI|Average disk total rate (read and write
across all disks) in MB/s
| IOPs Total | Implemented | DSNI |
|Utilization Total|Implemented|DSNI|
.13+|Storage|Latency Total|Implemented|DSNI|
|Latency Read|Implemented|DSNI|
| IOPs other | Implemented | DSNI |
| IOPs Write | Implemented | DSNI |
|Throughput Read|Implemented|DSNI|
```

```
| IOPs Total | Implemented | DSNI |
|Latency Write|Implemented|DSNI|
| IOps Read | Implemented | DSNI | Number of read IOps on the disk
|Cache Hit Ratio Read|Implemented|DSNI|
|Cache Hit Ratio Total|Implemented|DSNI|
|Cache Hit Ratio Write|Implemented|DSNI|
|Throughput Write|Implemented|DSNI|
|Throughput Total|Implemented|DSNI|Average disk total rate (read and write
across all disks) in MB/s
.17+|StoragePool Disk|Capacity Provisioned|Implemented|DSNI|
|Raw Capacity|Implemented|DSNI|
|Total Capacity|Implemented|DSNI|
|Used Capacity|Implemented|DSNI|
|Over Commit Capacity Ratio|Implemented|DSNI|Reported as a time series
|Capacity Used Ratio|Implemented|DSNI|
|Total Data Capacity|Implemented|DSNI|
|Data Used Capacity|Implemented|DSNI|
|IOps Read|Implemented|DSNI|Number of read IOps on the disk
| IOPs Total | Implemented | DSNI |
| IOPs Write | Implemented | DSNI |
|Throughput Read|Implemented|DSNI|
|Throughput Total|Implemented|DSNI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|DSNI|
|Read Utilization|Implemented|DSNI|
|Utilization Total|Implemented|DSNI|
|Utilization Write|Implemented|DSNI|
.14+|Volume|Cache Hit Ratio Read|Implemented|DSNI|
|Cache Hit Ratio Total|Implemented|DSNI|
|Cache Hit Ratio Write|Implemented|DSNI|
|Raw Capacity|Implemented|DSNI|
|Total Capacity|Implemented|DSNI|
| IOps Read | Implemented | DSNI | Number of read IOps on the disk
| IOPs Total | Implemented | DSNI |
| IOPs Write | Implemented | DSNI |
|Latency Read|Implemented|DSNI|
|Latency Total|Implemented|DSNI|
|Latency Write|Implemented|DSNI|
|Throughput Read|Implemented|DSNI|
|Throughput Total|Implemented|DSNI|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|DSNI|
|===
Management APIs used by this data collector:
```

```
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Data source wizard configuration
|Manual Entry
|true
|true
|true
|true
| IBM DS CLI
| DSNI
|DSNI
|true
|true
|true
|true
|===
<<top, Back to Top>>
== IBM PowerVM (SSH)
:description: Support Matrix Asciidoc for IBM PowerVM (SSH)
|===
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.54+|foundation .6+|Data Store|Capacity|Implemented|SSH|Snapshot Used
capacity in MB
|MOID|Implemented|SSH|
```

```
|Name|Implemented|SSH|
|OID|Implemented|SSH|
|Used Capacity|Implemented|SSH|
|Virtual Center Ip|Implemented|SSH|
.5+|LUN|Disk Name|Implemented|SSH|
|DataStore OID|Implemented|SSH|
|Host OID|Implemented|SSH|
|Number|Implemented|SSH|
|OID|Implemented|SSH|
.6+|Path|Active|Implemented|SSH|
|Lun OID|Implemented|SSH|
|Host Port WWPN|Implemented|SSH|
|OID|Implemented|SSH|
|Storage Port WWPN|Implemented|SSH|
|Type|Gap|SSH|
.4+|Server|Host OID|Implemented|SSH|
|MOID|Implemented|SSH|
|OID|Implemented|SSH|
|Virtual Center Ip|Implemented|SSH|
.6+|Virtual Disk|Capacity|Implemented|SSH|Snapshot Used capacity in MB
|DataStore OID|Implemented|SSH|
|Lun OID|Implemented|SSH|
|Name|Implemented|SSH|
|OID|Implemented|SSH|
|Type|Gap|SSH|
.12+|VirtualMachine|Dns Name|Implemented|SSH|
|Guest State|Implemented|SSH|
|Host OID|Implemented|SSH|
|IPs|Implemented|SSH|
|MOID|Implemented|SSH|
|Memory|Implemented|SSH|
|Name|Implemented|SSH|
|OID|Implemented|SSH|
|OS|Implemented|SSH|
|Power State|Implemented|SSH|
|State Change Time|Implemented|SSH|
|Processors|Implemented|SSH|
.3+|VirtualMachine Disk|OID|Implemented|SSH|
|VirtualDisk OID|Implemented|SSH|
|VirtualMachine OID|Implemented|SSH|
.9+|Host|Host Cpu Count|Implemented|SSH|
|Host Installed Memory|Implemented|SSH|
|Host Model|Implemented|SSH|
|NIC count|Implemented|SSH|
|IPs|Implemented|SSH|
|Manufacturer|Implemented|SSH|
```

```
|Name|Implemented|SSH|
|OID|Implemented|SSH|
|Platform Type|Implemented|SSH|
.3+|Info|DataSource Name|Implemented|SSH|Info
|Date|Implemented|SSH|
|Originator ID|Implemented|SSH|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
| IBM Hardware Management Console SSH access
|SSH
|SSH
122
|true
Ifalse
|true
|true
|===
<<top, Back to Top>>
== IBM SVC (CLI)
:description: Support Matrix Asciidoc for IBM SVC (CLI)
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
|2072-12F
2072-12G
2072-2N4
2072-324
2072-3H4
2072-3N4
2076-124
2076-12F
2076-224
```

```
2076-24F
2076-24G
2076-624
2076-724
2076-824
2076-AF6
2076-AFF
2077-24F
2077-424
2078-12F
2078-224
2078-24C
2078-24F
2078-324
2078-424
2078-4H4
2078-92G
2078-AF3
4657-924
4662-12G
4662-6H2
4666-AH8
9843-AE2
9843-AE3
9846-AG8
9848-AE2
9848-AF7
9848-AF8
9848-AG8
SVC
11.5.2.7
1.6.1.2
1.6.1.4
1.6.1.5
7.5.0.11
7.5.0.12
7.7.1.8
7.8.1.14
7.8.1.6
7.8.1.8
8.2.1.10
8.2.1.11
8.2.1.14
8.2.1.9
8.3.1.1
8.3.1.2
```

```
8.3.1.5
8.3.1.6
8.3.1.7
8.3.1.9
8.4.0.10
8.4.0.11
8.4.0.6
8.4.0.7
8.4.0.8
8.4.0.9
8.5.0.5
8.5.0.6
8.5.0.7
8.5.0.8
8.5.0.9
8.5.2.2
8.5.3.1
8.5.4.0
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.94+|foundation .15+|Disk|Capacity (GB)|Implemented|SSH|use capacity
|Disk Id|Implemented|SSH|Uniquely identifies this disk in the array
|Group|Implemented|SSH|
|Location|Gap|SSH|Where this disk is physically located in the array
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Role|Implemented|SSH|
|Role Enum|Implemented|SSH|enum for disk role
|Serial Number|Implemented|SSH|
|Speed|Implemented|SSH|Speed of disk (RPM)
|Status|Implemented|SSH|
|Type|Gap|SSH|
|Type Enum|Implemented|SSH|enum for disk type
|Vendor|Implemented|SSH|
|Vendor Id|Implemented|SSH|
.3+|Info|DataSource Name|Implemented|SSH|Info
|Date|Implemented|SSH|
|Originator ID|Implemented|SSH|
.13+|Storage|Display IP|Implemented|SSH|
|Failed Raw Capacity|Implemented|SSH|Raw capapcity of failed disks (sum of
all disks that are failed)
```

```
|Family|Implemented|SSH|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|SSH|
|Manage URL|Implemented|SSH|
|Manufacturer|Implemented|SSH|
|Microcode Version|Implemented|SSH|
|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|SSH|
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|SSH|Is this a storage virtualization device?
.7+|Storage Node|Model|Implemented|SSH|
|Name|Implemented|SSH|
|Partner Node UUID|Implemented|SSH|HA pair's UUID
|Serial Number|Implemented|SSH|
|State|Implemented|SSH|free text describing the device state
|UUID|Implemented|SSH|
|Parent Serial Number|Implemented|SSH|
.18+|Storage Pool|Compression Enabled|Implemented|SSH|Is compression
enabled on the storage pool
|Compression Savings|Implemented|SSH|ratio of compression savings in
percentage
|Dedupe Enabled|Implemented|SSH|Is dedupe enabled on the storage pool
|Include In Dwh Capacity|Implemented|SSH|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|SSH|
|Physical Disk Capacity (MB)|Implemented|SSH|used as raw capacity for
storage pool
|Raid Group|Implemented|SSH|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|SSH|ratio to convert from usable capacity
to raw capacity
|Redundancy|Implemented|SSH|Redundancy level
|Status|Implemented|SSH|
|Storage Pool Id|Implemented|SSH|
|Thin Provisioning Supported|Implemented|SSH|Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|SSH|
|Total Used Capacity|Implemented|SSH|Total capacity in MB
|Type|Gap|SSH|
|Uses FlashPools|Implemented|SSH|indicate if this storagepool uses Flash
Pools(NetApp specific)
|Virtual|Implemented|SSH|Is this a storage virtualization device?
```

```
|Encrypted|Implemented|SSH|
.9+|Storage Synchronization|Mode|Implemented|SSH|
|Mode Enum|Implemented|SSH|
|Source Storage|Implemented|SSH|
|Source Volume|Implemented|SSH|
|State|Implemented|SSH|free text describing the device state
|State Enum|Implemented|SSH|
|Target Storage|Implemented|SSH|
|Target Volume|Implemented|SSH|
|Technology|Implemented|SSH|technology which causes storage efficiency
changed
.18+|Volume|AutoTier Policy Identifier|Implemented|SSH|Dynamic Tier Policy
|Auto Tiering|Implemented|SSH|indicates if this storagepool is
participating in auto tiering with other pools
|Capacity|Implemented|SSH|Snapshot Used capacity in MB
|DiskGroup|Implemented|SSH|Disk Group Type
|Head|Implemented|SSH|Specify the head (in netapp) for this volume
|Junction Path|Implemented|SSH|
|Name|Implemented|SSH|
|Protection Type|Implemented|SSH|
|Total Raw Capacity|Implemented|SSH|Total raw capacity (sum of all disks
on the array)
|Storage Pool Id|Implemented|SSH|
|Thin Provisioned|Implemented|SSH|
|Type|Gap|SSH|
|UUID|Implemented|SSH|
|Used Capacity|Implemented|SSH|
|Virtual|Implemented|SSH|Is this a storage virtualization device?
|Written Capacity|Implemented|SSH|Total capacity written to this volume by
a Host in MB
|Compression Enabled|Implemented|SSH|
|Encrypted|Implemented|SSH|
.3+|Volume Map|LUN|Implemented|SSH|Name of the backend lun
|Protocol Controller|Implemented|SSH|
|Storage Port|Implemented|SSH|
.4+|Volume Mask|Initiator|Implemented|SSH|
|Protocol Controller|Implemented|SSH|
|Storage Port|Implemented|SSH|
|Type|Gap|SSH|
.4+|WWN Alias|Host Aliases|Implemented|SSH|
|Object Type|Implemented|SSH|
|Source|Implemented|SSH|
|WWN|Implemented|SSH|
.70+|performance .9+|Disk|IOps Read|Implemented|SSH|Number of read IOps on
the disk
```

```
| IOPs Total | Implemented | SSH |
| IOPs Write | Implemented | SSH |
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
|Read Utilization|Implemented|SSH|
|Utilization Total|Implemented|SSH|
|Utilization Write|Implemented|SSH|
.17+|Storage|Cache Hit Ratio Read|Implemented|SSH|
|Cache Hit Ratio Total|Implemented|SSH|
|Cache Hit Ratio Write|Implemented|SSH|
|Failed Raw Capacity|Implemented|SSH|
|Raw Capacity|Implemented|SSH|
|Spare Raw Capacity|Implemented|SSH|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|SSH|
|IOPs other|Implemented|SSH|
|IOps Read|Implemented|SSH|Number of read IOps on the disk
| IOPs Total | Implemented | SSH |
| IOPs Write | Implemented | SSH |
|Latency Read|Implemented|SSH|
|Latency Total|Implemented|SSH|
|Latency Write|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
.11+|Storage Node|Cache Hit Ratio Total|Implemented|SSH|
|IOps Read|Implemented|SSH|Number of read IOps on the disk
| IOPs Total | Implemented | SSH |
| IOPs Write | Implemented | SSH |
|Latency Read|Implemented|SSH|
|Latency Total|Implemented|SSH|
|Latency Write|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
|Utilization Total|Implemented|SSH|
.15+|StoragePool Disk|Capacity Provisioned|Implemented|SSH|
|Raw Capacity|Implemented|SSH|
|Total Capacity|Implemented|SSH|
|Used Capacity|Implemented|SSH|
|Over Commit Capacity Ratio|Implemented|SSH|Reported as a time series
|Capacity Used Ratio|Implemented|SSH|
```

```
|IOps Read|Implemented|SSH|Number of read IOps on the disk
| IOPs Total | Implemented | SSH |
|IOPs Write|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
|Read Utilization|Implemented|SSH|
|Utilization Total|Implemented|SSH|
|Utilization Write|Implemented|SSH|
.18+|Volume|Cache Hit Ratio Read|Implemented|SSH|
|Cache Hit Ratio Total|Implemented|SSH|
|Cache Hit Ratio Write|Implemented|SSH|
|Raw Capacity|Implemented|SSH|
|Total Capacity|Implemented|SSH|
|Used Capacity|Implemented|SSH|
|Written Capacity|Implemented|SSH|
|Capacity Used Ratio|Implemented|SSH|
|CapacityRatio Written|Implemented|SSH|
|IOps Read|Implemented|SSH|Number of read IOps on the disk
| IOPs Total | Implemented | SSH |
| IOPs Write | Implemented | SSH |
|Latency Read|Implemented|SSH|
|Latency Total|Implemented|SSH|
|Latency Write|Implemented|SSH|
|Throughput Read|Implemented|SSH|
|Throughput Total|Implemented|SSH|Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented|SSH|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
| IBM SVC CLI
|SSH
|SSH
122
|true
|false
```

```
Itrue
Itrue
|===
<<top, Back to Top>>
== IBM XIV & A9000 (XIVCLI)
:description: Support Matrix Asciidoc for IBM XIV & A9000 (XIVCLI)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
1415
A14
|10.2.4.e
12.3.2.c
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.94+|foundation .13+|Disk|Capacity (GB)|Implemented|XIV CLI|use capacity
|Disk Id|Implemented|XIV CLI|Uniquely identifies this disk in the array
|Model|Implemented|XIV CLI|
|Name|Implemented|XIV CLI|
|Role|Implemented|XIV CLI|
|Role Enum|Implemented|XIV CLI|enum for disk role
|Serial Number|Implemented|XIV CLI|
|Speed|Implemented|XIV CLI|Speed of disk (RPM)
|Status|Implemented|XIV CLI|
|Status Enum|Implemented|XIV CLI|enum for disk status
|Type|Gap|XIV CLI|
|Type Enum|Implemented|XIV CLI|enum for disk type
|Vendor|Implemented|XIV CLI|
.6+|Info|Client Api Name|Implemented|XIV CLI|
|Client Api Version|Implemented|XIV CLI|
|DataSource Name|Implemented|XIV CLI|Info
|Date|Implemented|XIV CLI|
|Originator ID|Implemented|XIV CLI|
|Originator Key|Implemented|XIV CLI|
.16+|Internal Volume|Internal Volume Id|Implemented|XIV CLI|
|Name|Implemented|XIV CLI|
```

```
|Storage Pool Id|Implemented|XIV CLI|
|Type|Gap|XIV CLI|
|Thin Provisioned|Implemented|XIV CLI|
|Thin Provisioning Supported|Implemented|XIV CLI|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Dedupe Enabled|Implemented|XIV CLI|Is dedupe enabled on the storage pool
|Snapshot Used Capacity|Implemented|XIV CLI|
|Snapshot Allocated Capacity|Gap|XIV CLI|Allocated capacity of snapshots
|Data Used Capacity|Implemented|XIV CLI|
|Data Allocated Capacity|Gap|XIV CLI|capacity allocated for data
|Total Used Capacity|Implemented|XIV CLI|Total capacity in MB
|Total Used Capacity (MB)|Implemented|XIV CLI|place holder for the used
capacity as read from the device
|Total Allocated Capacity|Implemented|XIV CLI|
|Other Allocated Capacity|Gap|XIV CLI|Capacity allocated for other (not
data and not snapshot)
|Raw to Usable Ratio|Implemented|XIV CLI|ratio to convert from usable
capacity to raw capacity
.3+|QTree|Qtree Id|Implemented|XIV CLI|unique id of the qtree
|Name|Implemented|XIV CLI|
|Type|Gap|XIV CLI|
.12+|Storage|Display IP|Implemented|XIV CLI|
|Failed Raw Capacity|Implemented|XIV CLI|Raw capapcity of failed disks
(sum of all disks that are failed)
|Family|Implemented|XIV CLI|The storage Family could be Clariion,
Symmetrix, et al
|IP|Implemented|XIV CLI|
|Manufacturer|Implemented|XIV CLI|
|Microcode Version|Implemented|XIV CLI|
|Model|Implemented|XIV CLI|
|Name|Implemented|XIV CLI|
|Total Raw Capacity|Implemented|XIV CLI|Total raw capacity (sum of all
disks on the array)
|Serial Number|Implemented|XIV CLI|
|Spare Raw Capacity|Implemented|XIV CLI|Raw capapcity of spare disks (sum
of all disks that are spare)
|Virtual|Implemented|XIV CLI|Is this a storage virtualization device?
.23+|Storage Pool|Auto Tiering|Implemented|XIV CLI|indicates if this
storagepool is participating in auto tiering with other pools
|Compression Enabled|Implemented|XIV CLI|Is compression enabled on the
storage pool
|Compression Savings|Implemented|XIV CLI|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|XIV CLI|capacity allocated for data
|Data Used Capacity|Implemented|XIV CLI|
```

```
|Dedupe Enabled|Implemented|XIV CLI|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|XIV CLI|ratio of dedupe savings in percentage
| Include In Dwh Capacity | Implemented | XIV CLI | A way from ACQ to cotnrol
which stroage pools are interesting in DWH Capacity
|Name|Implemented|XIV CLI|
|Other UsedCapacity (MB)|Implemented|XIV CLI|Any capacity other than data
and snapshot
|Physical Disk Capacity (MB)|Implemented|XIV CLI|used as raw capacity for
storage pool
|Raid Group|Implemented|XIV CLI|indicates whether this storagePool is a
raid group
|Raw to Usable Ratio|Implemented|XIV CLI|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|XIV CLI|Redundancy level
|Snapshot Allocated Capacity|Gap|XIV CLI|Allocated capacity of snapshots
in MB
|Snapshot Used Capacity|Implemented|XIV CLI|
|Soft Limit (MB)|Implemented|XIV CLI|logical volume size that is defined
during volume creation or resizing operations
|Storage Pool Id|Implemented|XIV CLI|
|Thin Provisioning Supported|Implemented|XIV CLI|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|XIV CLI|
|Total Used Capacity|Implemented|XIV CLI|Total capacity in MB
|Type|Gap|XIV CLI|
|Virtual|Implemented|XIV CLI|Is this a storage virtualization device?
.12+|Volume|Capacity|Implemented|XIV CLI|Snapshot Used capacity in MB
|DiskGroup|Implemented|XIV CLI|Disk Group Type
|Disk Type|Not Available|XIV CLI|
|Name|Implemented|XIV CLI|
|Qtree Id|Implemented|XIV CLI|unique id of the qtree
|Total Raw Capacity|Implemented|XIV CLI|Total raw capacity (sum of all
disks on the array)
|Redundancy|Implemented|XIV CLI|Redundancy level
|Storage Pool Id|Implemented|XIV CLI|
|Thin Provisioned|Implemented|XIV CLI|
|Type|Gap|XIV CLI|
|Used Capacity|Implemented|XIV CLI|
|Compression Enabled|Implemented|XIV CLI|
.2+|Volume Map|LUN|Implemented|XIV CLI|Name of the backend lun
|Protocol Controller|Implemented|XIV CLI|
.2+|Volume Mask|Initiator|Implemented|XIV CLI|
|Protocol Controller|Implemented|XIV CLI|
.5+|WWN Alias|Host Aliases|Implemented|XIV CLI|
|Host OS|Implemented|XIV CLI|
|Object Type|Implemented|XIV CLI|
```

```
|Source|Implemented|XIV CLI|
|WWN|Implemented|XIV CLI|
.33+|performance .13+|Storage|Latency Total|Implemented|DSNI|
|Latency Read|Implemented|DSNI|
| IOPs other | Implemented | DSNI |
| IOPs Write | Implemented | DSNI |
|Throughput Read|Implemented|DSNI|
| IOPs Total | Implemented | DSNI |
|Latency Write|Implemented|DSNI|
| IOps Read | Implemented | DSNI | Number of read IOps on the disk
|Cache Hit Ratio Read|Implemented|DSNI|
|Cache Hit Ratio Total|Implemented|DSNI|
|Cache Hit Ratio Write|Implemented|DSNI|
|Throughput Write|Implemented|DSNI|
|Throughput Total|Implemented|DSNI|Average disk total rate (read and write
across all disks) in MB/s
.6+|StoragePool Disk|Total Capacity|Implemented|DSNI|
|Capacity Used Ratio|Implemented|DSNI|
|Capacity Provisioned|Implemented|DSNI|
|Used Capacity|Implemented|DSNI|
|Raw Capacity|Implemented|DSNI|
|Over Commit Capacity Ratio|Implemented|DSNI|Reported as a time series
.14+|Volume|Latency Total|Implemented|DSNI|
|Latency Read|Implemented|DSNI|
| IOPs Write | Implemented | DSNI |
|Compression Savings Space|Implemented|DSNI|
|Throughput Read|Implemented|DSNI|
| IOPs Total | Implemented | DSNI |
|Latency Write|Implemented|DSNI|
|IOps Read|Implemented|DSNI|Number of read IOps on the disk
|Cache Hit Ratio Read|Implemented|DSNI|
|Total Compression Savings|Implemented|DSNI|
|Cache Hit Ratio Total|Implemented|DSNI|
|Cache Hit Ratio Write|Implemented|DSNI|
|Throughput Write|Implemented|DSNI|
|Throughput Total|Implemented|DSNI|Average disk total rate (read and write
across all disks) in MB/s
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
```

```
| IBM DS CLI
|DSNI
|DSNI
|true
|true
|true
|true
| IBM XIV CLI
|XIV CLI
|TCP
17778
|true
Ifalse
|true
|false
|===
<<top, Back to Top>>
== Infinidat Infinibox (HTTP)
:description: Support Matrix Asciidoc for Infinidat Infinibox (HTTP)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|F6230
F6240
F6303
F6304
16.0.31.0
7.0.14.20
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.99+|foundation .11+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
```

```
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Serial Number|Implemented|HTTPS|
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Vendor|Implemented|HTTPS|
.5+|File Share|Is InternalVolume|Implemented|HTTPS|whether the file share
represents an internal volume (netapp volume) or is it a qtree/folder
within the internal volume
|Is Shared|Implemented|HTTPS|whether this fileShare has any shares
associated with it
|Name|Implemented|HTTPS|
|Path|Implemented|HTTPS|path of the fileShare
|Qtree Id|Implemented|HTTPS|unique id of the qtree
.4+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.11+|Internal Volume|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on
the storage pool
|Internal Volume Id|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Other Allocated Capacity|Gap|HTTPS|Capacity allocated for other (not data
and not snapshot)
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
.3+|QTree|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Type|Gap|HTTPS|
.3+|Share|IP Interfaces|Implemented|HTTPS|comma separated list of IP
addresses on which this share is exposed
|Name|Implemented|HTTPS|
|Protocol|Implemented|HTTPS|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|HTTPS|
|Permission|Implemented|HTTPS|Permissions for this particular share
```

```
.14+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|HTTPS|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.4+|Storage Node|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
.17+|Storage Pool|Data Allocated Capacity|Gap|HTTPS|capacity allocated for
data
|Data Used Capacity|Implemented|HTTPS|
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
|Reserved Capacity|Implemented|HTTPS|Reserved Capacity in MB
|Soft Limit (MB)|Implemented|HTTPS|logical volume size that is defined
during volume creation or resizing operations
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
```

```
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.13+|Volume|Name|Implemented|HTTPS|
|Junction Path|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|Thin Provisioned|Implemented|HTTPS|
|Replica Source|Implemented|HTTPS|
|Replica Target|Implemented|HTTPS|
|Snapshot|Implemented|HTTPS|
|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Used Capacity|Implemented|HTTPS|
.4+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Protocol Controller|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Storage Port|Implemented|HTTPS|
.4+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Storage Port|Implemented|HTTPS|
.4+|WWN Alias|Source|Implemented|HTTPS|
|Host Aliases|Implemented|HTTPS|
|WWN|Implemented|HTTPS|
|Object Type|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Infinidat REST API
HTTPS
| HTTPS
|443
|true
|true
|true
|true
```

```
|===
<<top, Back to Top>>
== Microsoft Azure Compute
:description: Support Matrix Asciidoc for Microsoft Azure Compute
Models and versions supported by this data collector:
|===
<.<|API versions
12018-06-01
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.54+|foundation .7+|Data Store|Capacity|Implemented|HTTPS|Snapshot Used
capacity in MB
|MOID|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
| Provisioned Capacity | Implemented | HTTPS |
|Virtual Center Ip|Implemented|HTTPS|
|subscription Id|Implemented|HTTPS|
.6+|Server|Cluster|Implemented|HTTPS|Cluster name
|DataCenter Name|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
|MOID|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Virtual Center Ip|Implemented|HTTPS|
.8+|Virtual Disk|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|DataStore OID|Implemented|HTTPS|
|Lun OID|Implemented|HTTPS|
| Is Chargeable | Implemented | HTTPS |
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
|subscription Id|Implemented|HTTPS|
.18+|VirtualMachine|Guest State|Implemented|HTTPS|
|DataStore OID|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
|IPs|Implemented|HTTPS|
```

```
|MOID|Implemented|HTTPS|
|Memory|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|OS|Implemented|HTTPS|
|Power State|Implemented|HTTPS|
|State Change Time|Implemented|HTTPS|
|Processors|Implemented|HTTPS|
|Provisioned Capacity|Implemented|HTTPS|
|Instance Type|Implemented|HTTPS|
|Launch Time|Implemented|HTTPS|
|public Ips|Implemented|HTTPS|
|Security Groups|Implemented|HTTPS|
|subscription Id|Implemented|HTTPS|
.3+|VirtualMachine Disk|OID|Implemented|HTTPS|
|VirtualDisk OID|Implemented|HTTPS|
|VirtualMachine OID|Implemented|HTTPS|
.5+|Host|Host OS|Implemented|HTTPS|
|IPs|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.7+|Info|Api Description|Implemented|HTTPS|
|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.20+|performance .3+|Data Store|Capacity Provisioned|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio|Implemented|HTTPS|Reported as a time series
.7+|Virtual Disk|Total Capacity|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.10+|vm|Total CPU Utilization|Implemented|HTTPS|
| IOps Read | Implemented | HTTPS | Number of read IOps on the disk
|diskIops.total|Implemented|HTTPS|
|Disk IOPs write|Implemented|HTTPS|
|Disk Throughput Read|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|total disk throughput read
```

```
|Disk Throughput Write|Implemented|HTTPS|
| IP Throughput Read | Implemented | HTTPS |
|Throughput total|Implemented|HTTPS|IP throughput total
|ipThroughput.write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Microsoft Azure Compute REST API
|HTTPS
|HTTPS
| 443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Microsoft Hyper-V
:description: Support Matrix Asciidoc for Microsoft Hyper-V
|===
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.82+|foundation .7+|Data Store|Capacity|Implemented|WMI|Snapshot Used
capacity in MB
|MOID|Implemented|WMI|
|Name|Implemented|WMI|
|OID|Implemented|WMI|
```

```
| Provisioned Capacity | Implemented | WMI |
|Used Capacity|Implemented|WMI|
|Virtual Center Ip|Implemented|WMI|
.3+|NasShare DataStore|DataStore OID|Implemented|WMI|
|Nas Share OID|Implemented|WMI|
|OID|Implemented|WMI|
.4+|NasShare Host|Nas Share OID|Implemented|WMI|
|Host OID|Implemented|WMI|
|OID|Implemented|WMI|
|Read Only|Implemented|WMI|
.6+|LUN|Disk Name|Implemented|WMI|
|DataStore OID|Implemented|WMI|
|Host OID|Implemented|WMI|
|Number|Implemented|WMI|
|OID|Implemented|WMI|
|TID|Implemented|WMI|
.6+|NAS Share|Capacity|Implemented|WMI|Allocated capacity in MB
|Filer Ip|Implemented|WMI|
|Filer Name|Implemented|WMI|
|OID|Implemented|WMI|
|Share Path|Implemented|WMI| For a HvNasShare to be matched to a Share
|Type|Gap|WMI|
.6+|Path|Active|Implemented|WMI|
|Lun OID|Implemented|WMI|
|Host Port WWPN|Implemented|WMI|
|OID|Implemented|WMI|
|Storage Port WWPN|Implemented|WMI|
|Type|Gap|WMI|
.6+|Server|Cluster|Implemented|WMI|Cluster name
|DataCenter Name|Implemented|WMI|
|Host OID|Implemented|WMI|
|MOID|Implemented|WMI|
|OID|Implemented|WMI|
|Virtual Center Ip|Implemented|WMI|
.7+|Virtual Disk|Capacity|Implemented|WMI|Snapshot Used capacity in MB
|DataStore OID|Implemented|WMI|
|Name|Implemented|WMI|
|OID|Implemented|WMI|
|Type|Gap|WMI|
|Used Capacity|Implemented|WMI|used capacity for reporting (MB)
|Used Capacity|Implemented|WMI|
.15+|VirtualMachine|Dns Name|Implemented|WMI|
|Guest State|Implemented|WMI|
|DataStore OID|Implemented|WMI|
|Host OID|Implemented|WMI|
|IPs|Implemented|WMI|
```

```
|MOID|Implemented|WMI|
|Memory|Implemented|WMI|
|Name|Implemented|WMI|
|OID|Implemented|WMI|
|OS|Implemented|WMI|
| Power State | Implemented | WMI |
|State Change Time|Implemented|WMI|
|Processors|Implemented|WMI|
|Provisioned Capacity|Implemented|WMI|
|Used Capacity|Implemented|WMI|
.3+|VirtualMachine Disk|OID|Implemented|WMI|
|VirtualDisk OID|Implemented|WMI|
|VirtualMachine OID|Implemented|WMI|
.12+|Host|Host Cpu Count|Implemented|WMI|
|Host Cpu Speed|Implemented|WMI|
|Host Domain|Implemented|WMI|
|Host Installed Memory|Implemented|WMI|
|Host Model|Implemented|WMI|
|NIC count|Implemented|WMI|
|NIC speed|Implemented|WMI|
|IPs|Implemented|WMI|
|Manufacturer|Implemented|WMI|
|Name|Implemented|WMI|
|OID|Implemented|WMI|
|Platform Type|Implemented|WMI|
.4+|ISCSI Node|Host Aliases|Implemented|WMI|
|Node Name|Implemented|WMI|
|OID|Implemented|WMI|
|Type|Gap|WMI|
.3+|Info|DataSource Name|Implemented|WMI|Info
|Date|Implemented|WMI|
|Originator ID|Implemented|WMI|
.43+|performance .5+|Data Store|Capacity Provisioned|Implemented|WS-
Management |
|Total Capacity|Implemented|WS-Management|
|Used Capacity|Implemented|WS-Management|
|Over Commit Capacity Ratio | Implemented | WS-Management | Reported as a time
|Capacity Used Ratio|Implemented|WS-Management|
.14+|Host|Total CPU Utilization|Implemented|WS-Management|
| IOps Read | Implemented | WS-Management | Number of read IOps on the disk
|disklops.total|Implemented|WS-Management|
|Disk IOPs write|Implemented|WS-Management|
|Latency Read|Implemented|WS-Management|
|Latency Total|Implemented|WS-Management|
|Latency Write|Implemented|WS-Management|
```

```
|Disk Throughput Read|Implemented|WS-Management|
|Throughput Read|Implemented|WS-Management|total disk throughput read
|Disk Throughput Write|Implemented|WS-Management|
| IP Throughput Read | Implemented | WS-Management |
|Throughput total|Implemented|WS-Management|IP throughput total
|ipThroughput.write|Implemented|WS-Management|
|Total Memory Utilization|Implemented|WS-Management|
.10+|Virtual Disk|Total Capacity|Implemented|WS-Management|
| Used Capacity | Implemented | WS-Management |
|Capacity Used Ratio|Implemented|WS-Management|
| IOps Read | Implemented | WS-Management | Number of read IOps on the disk
| IOPs Total | Implemented | WS-Management |
| IOPs Write | Implemented | WS-Management |
|Latency Total|Implemented|WS-Management|
|Throughput Read|Implemented|WS-Management|
|Throughput Total|Implemented|WS-Management|Average disk total rate (read
and write across all disks) in MB/s
|Throughput Write|Implemented|WS-Management|
.14+|vm|Total Capacity|Implemented|WS-Management|
| Used Capacity | Implemented | WS-Management |
|Capacity Used Ratio|Implemented|WS-Management|
|Total CPU Utilization|Implemented|WS-Management|
|IOps Read|Implemented|WS-Management|Number of read IOps on the disk
|diskIops.total|Implemented|WS-Management|
|Disk IOPs write|Implemented|WS-Management|
|Latency Total|Implemented|WS-Management|
|Disk Throughput Read|Implemented|WS-Management|
|Throughput Read|Implemented|WS-Management|total disk throughput read
|Disk Throughput Write|Implemented|WS-Management|
| IP Throughput Read | Implemented | WS-Management |
|Throughput total|Implemented|WS-Management|IP throughput total
|ipThroughput.write|Implemented|WS-Management|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|PowerShell
|WS-Management
|HTTP
|5985
```

```
|true
|false
|false
|true
|WMI
|WMI
|WMI
|135
|true
|false
|true
|true
|===
<<top, Back to Top>>
== NetApp 7 Mode
:description: Support Matrix Asciidoc for NetApp 7 Mode
Models and versions supported by this data collector:
|===
<.<|API versions <.<|Models <.<|Firmware versions
|1.12
1.14
1.17
1.19
1.20
1.21
|FAS2040
FAS2050
FAS2220
FAS2240-2
FAS2240-4
FAS2520
FAS2554
FAS3140
FAS3160
FAS3210
FAS3220
FAS3240
FAS3250
```

```
FAS3270
FAS6240
FAS6290
FAS8020
FAS8040
FAS8060
FAS8080
N6070
N6240
V3240
17.3.6
8.1.1 7-Mode
8.1.3P2 7-Mode
8.1.4P1 7-Mode
8.1.4P10 7-Mode
8.1.4P9D18 7-Mode
8.2.1 7-Mode
8.2.2 7-Mode
8.2.3 7-Mode
8.2.3P2 7-Mode
8.2.3P3 7-Mode
8.2.4 7-Mode
8.2.4P2 7-Mode
8.2.4P4 7-Mode
8.2.4P5 7-Mode
8.2.4P6 7-Mode
8.2.5 7-Mode
8.2.5P1 7-Mode
8.2.5P2 7-Mode
8.2.5P4 7-Mode
8.2.5P5 7-Mode
8.2P3 7-Mode
8.2P4 7-Mode
Data ONTAP Release 7.3.3
Data ONTAP Release 7.3.4
Data ONTAP Release 8.2.5 7-Mode
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.180+|foundation .4+|Capability|Active|Implemented||
|Licensed|Implemented||
|Type|Gap||
```

```
|Used|Implemented||
.15+|Disk|Capacity (GB)|Implemented||use capacity
|Disk Id|Implemented||Uniquely identifies this disk in the array
|Group|Implemented||
|Location|Gap||Where this disk is physically located in the array
|Model|Implemented||
|Name|Implemented||
|Role|Implemented||
|Role Enum|Implemented||enum for disk role
|Serial Number|Implemented||
|Speed|Implemented||Speed of disk (RPM)
|Status|Implemented||
|Status Enum|Implemented||enum for disk status
|Type|Gap||
|Type Enum|Implemented||enum for disk type
|Vendor|Implemented||
.8+|File Share|Is InternalVolume|Implemented||whether the file share
represents an internal volume (netapp volume) or is it a qtree/folder
within the internal volume
| Is Oplocks Enabled | Implemented | | whether opportunistic locks are enabled
on the FileShare
|Is Shared|Implemented||whether this fileShare has any shares associated
with it
|Name|Implemented||
|Path|Implemented||path of the fileShare
|Qtree Id|Implemented||unique id of the gtree
|Security Type|Implemented||
|Status|Implemented||
.4+|ISCSI Network Portal|IP|Implemented||
|Listening Port|Implemented||
|Nic|Implemented||
|OID|Implemented||
.3+|ISCSI Network Portal Group|OID|Implemented||
|Portal Group Name|Implemented||
|Portal Group Tag|Implemented||
.4+|ISCSI Node|Host Aliases|Implemented||
|Node Name|Implemented||
|OID|Implemented||
|Type|Gap||
.2+|ISCSI Node Map|OID|Implemented||
|Portal Group OID|Implemented||
.9+|ISCSI Session|Initiator Ips|Implemented||
|Initiator OID|Implemented||
|Initiator Session Id|Implemented||
|Max Connections|Implemented||
|Number Of Connections|Implemented||
```

```
|OID|Implemented||
|Portal Group OID|Implemented||
|Security|Implemented||
|Target Session Id|Implemented||
.6+|Info|Api Name|Implemented||
|Api Version|Implemented||
|DataSource Name|Implemented||Info
|Date|Implemented||
|Originator ID|Implemented||
|Originator Key|Implemented||
.27+|Internal Volume|Clone Sourc|Implemented||name of the source internal
volume
|Compression Enabled|Implemented||Is compression enabled on the storage
|Compression Savings|Implemented||ratio of compression savings in
percentage
|Data Allocated Capacity|Gap||capacity allocated for data
|Data Used Capacity|Implemented||
|Dedupe Enabled|Implemented||Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented||ratio of dedupe savings in percentage
|Internal Volume Id|Implemented||
|Last Known Access Time|Implemented||Last know access to the volume
|Last Snapshot Time|Implemented||time of last snapshot
|Name|Implemented||
|Protection Type|Implemented||
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
|Snapshot Allocated Capacity|Gap||Allocated capacity of snapshots in MB
|Snapshot Count|Implemented||Number of snapshots on the internal volumes
|Snapshot Used Capacity|Implemented||
|Space Guarantee|Implemented||Space Guatantee policy (file, volume or
none)
|Status|Implemented||
|Storage Pool Id|Implemented||
|Thin Provisioned|Implemented||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented||
|Total Clone Saved Capacity in MB|Implemented||
|Total Used Capacity|Implemented||Total capacity in MB
|Total Used Capacity (MB)|Implemented||place holder for the used capacity
as read from the device
|Type|Gap||
|Virtual Storage|Implemented||Owning virtual storage (vfiler)
.3+|Storage Synchronization|Source Internal Volume|Implemented||
|Target Internal Volume|Implemented||
```

```
|Technology|Implemented||technology which causes storage efficiency
changed
.9+|QTree|Name|Implemented||
|Oplocks|Implemented||Indicates whether opportunistic locks are enabled on
|Qtree Id|Implemented||unique id of the gtree
|Quota HardCapacity Limit (MB)|Implemented||Maximum amount of disk space,
allowed for the quota target
|Quota SoftCapacity Limit (MB)|Implemented||Maximum amount of disk space,
allowed for the quota target
|Quota UsedCapacity|Implemented||Space in MB currently used
|Security Style|Implemented||Security style of the directory: unix, ntfs,
or mixed
|Status|Implemented||
|Type|Gap||
.3+|Qtree Storage Synchronization|Source Qtree Id|Implemented||
|Target Qtree Id|Implemented||
|Technology|Implemented||technology which causes storage efficiency
changed
.12+|Quota|Hard Capacity Limit (MB)|Implemented||max amount of disk space,
allowed for the quota target (Hard limit)
|Hard File Limit|Implemented||max number of files allowed for the quota
target
|Internal Volume Id|Implemented||
|Qtree Id|Implemented||unique id of the qtree
|Quota Id|Implemented||unique id of the quota
|Soft Capacity Limit (MB)|Implemented||Maximum amount of disk space,
allowed for the quota target
|Soft File Limit|Implemented||Max number of files allowed for the quota
|Threshold (MB)|Implemented||Disk space threshold, for the quota target
|Type|Gap||
|Used Capacity|Implemented||
|Used Files|Implemented||Number of files currently used
|User/Group Target|Implemented||user/group target this quota refers to
.4+|Share|Description|Implemented||
| IP Interfaces | Implemented | | comma separated list of IP addresses on which
this share is exposed
|Name|Implemented||
|Protocol|Implemented||enum for share protocol
.2+|Share Initiator|Initiator|Implemented||
|Permission|Implemented||Permissions for this particular share
.15+|Storage|Cpu Count|Implemented||Cpu Count of the storage
|Display IP|Implemented||
|Failed Raw Capacity|Implemented||Raw capapcity of failed disks (sum of
all disks that are failed)
```

```
|Family|Implemented||The storage Family could be Clariion, Symmetrix, et
al
|IP|Implemented||
|Manage URL|Implemented||
|Manufacturer|Implemented||
|Memory|Implemented||
|Microcode Version|Implemented||
|Model|Implemented||
|Name|Implemented||
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Serial Number|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|Virtual|Implemented||Is this a storage virtualization device?
.10+|Storage Node|Cache Size|Implemented||device cache size in MB
|Memory Size|Gap||device memory in MB
|Model|Implemented||
|Name|Implemented||
|Processors Count|Implemented||device CPU
|Partner Node UUID|Implemented||HA pair's UUID
|Serial Number|Implemented||
|UUID|Implemented||
|Version|Implemented||software version
|ManagementIp Addresses|Implemented||
.18+|Storage Pool|Data Allocated Capacity|Gap||capacity allocated for data
|Data Used Capacity|Implemented||
|Include In Dwh Capacity|Implemented||A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented||
|Physical Disk Capacity (MB)|Implemented||used as raw capacity for
storage pool
|Raid Group|Implemented||indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
|Redundancy|Implemented||Redundancy level
|Snapshot Allocated Capacity|Gap||Allocated capacity of snapshots in MB
|Snapshot Used Capacity|Implemented||
|Status|Implemented||
|Storage Pool Id|Implemented||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented||
|Total Used Capacity|Implemented||Total capacity in MB
|Type|Gap||
```

```
|Uses FlashPools|Implemented||indicate if this storagepool uses Flash
Pools(NetApp specific)
|Virtual|Implemented||Is this a storage virtualization device?
.15+|Volume|Capacity|Implemented||Snapshot Used capacity in MB
|DiskGroup|Implemented||Disk Group Type
|Disk Size|Implemented||comma seperated list of disk sizes (GB)
|Disk Speed|Implemented||comma seperated list of disk speeds (rpm)
|Disk Type|Not Available||
|Junction Path|Implemented||
|Last Known Access Time|Implemented||Last know access to the volume
|Name|Implemented||
|Protection Type|Implemented||
|Qtree Id|Implemented||unique id of the qtree
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Storage Pool Id|Implemented||
|Thin Provisioned|Implemented||
|Type|Gap||
|Used Capacity|Implemented||
.3+|Volume Map|LUN|Implemented||Name of the backend lun
|Protocol Controller|Implemented||
|Type|Gap||
.4+|Volume Mask|Initiator|Implemented||
|Protocol Controller|Implemented||
|Storage Port|Implemented||
|Type|Gap||
.94+|performance .9+|Disk|IOps Read|Implemented||Number of read IOps on
the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Read Utilization|Implemented||
|Utilization Total|Implemented||
|Utilization Write|Implemented||
.13+|Internal Volume|IO Density Read|Implemented||
| IO Density Total | Implemented | |
|Write IO Density|Implemented||
| IOPs other | Implemented | |
| IOps Read | Implemented | Number of read IOps on the disk
| IOPs Total | Implemented | |
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
```

```
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.5+|Qtree|Quota Hard Limit|Implemented||Capacity quota hard limit
|Quota Soft Limit|Implemented||Capacity Quota soft Limit
|Used Capacity|Implemented||
|Total File Count|Implemented||
| IOPs Total | Implemented | |
.17+|Storage|Failed Raw Capacity|Implemented||
|Raw Capacity|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|StoragePools Capacity|Implemented||
| IO Density Read | Implemented | |
| IO Density Total | Implemented | |
|Write IO Density|Implemented||
|IOPs other|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
|IOPs Total|Implemented||
| IOPs Write | Implemented | |
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.12+|Storage Node|Cache Hit Ratio Total|Implemented||
|Total Disk Read Replaced|Implemented||
|IOps Read|Implemented||Number of read IOps on the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Utilization Total|Implemented||
.20+|StoragePool Disk|Capacity Provisioned|Implemented||
|Raw Capacity|Implemented||
|Total Capacity|Implemented||
|Used Capacity|Implemented||
```

```
|Over Commit Capacity Ratio|Implemented||Reported as a time series
|Capacity Used Ratio|Implemented||
|Total Data Capacity|Implemented||
|Data Used Capacity|Implemented||
|IOps Read|Implemented||Number of read IOps on the disk
| IOPs Total | Implemented | |
|IOPs Write|Implemented||
|Snapshot Reserved Capacity|Implemented||
|Snapshot Used Capacity|Implemented||
|Snapshot Used Capacity Ratio|Implemented|| Reported as a time series
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Read Utilization|Implemented||
|Utilization Total|Implemented||
|Utilization Write|Implemented||
.18+|Volume|Raw Capacity|Implemented||
|Total Capacity|Implemented||
|Used Capacity|Implemented||
|Capacity Used Ratio|Implemented||
| IO Density Read | Implemented | |
| IO Density Total | Implemented | |
|Write IO Density|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
| IOPs Total | Implemented | |
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Partial Blocked Ratio|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Write Pending|Implemented||total write pending
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
```

```
|NetApp 7 Mode ZAPI
IZAPI
|ZAPI
|true
|true
|true
|true
|===
<<top, Back to Top>>
== NetApp Cloud Volumes Service
:description: Support Matrix Asciidoc for NetApp Cloud Volumes Service
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|AWS Cloud Volumes
|v1
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.54+|foundation .4+|Capability|Type|Gap||
|Active|Implemented||
|Licensed|Implemented||
|Used|Implemented||
.3+|Info|DataSource Name|Implemented||Info
|Originator ID|Implemented||
|Date|Implemented||
.17+|Internal Volume|Internal Volume Id|Implemented||
|Name|Implemented||
|Storage Pool Id|Implemented||
|Type|Gap||
|Thin Provisioned|Implemented||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Dedupe Enabled|Implemented||Is dedupe enabled on the storage pool
|Snapshot Count|Implemented||Number of snapshots on the internal volumes
```

```
|Status|Implemented||
|UUID|Implemented||
|Snapshot Used Capacity|Implemented||
|Data Used Capacity|Implemented||
|Data Allocated Capacity|Gap||capacity allocated for data
|Total Used Capacity|Implemented||Total capacity in MB
|Total Used Capacity (MB)|Implemented||place holder for the used capacity
as read from the device
|Total Allocated Capacity|Implemented||
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
.5+|QTree|Qtree Id|Implemented||unique id of the qtree
|Name|Implemented||
|Security Style|Implemented||Security style of the directory: unix, ntfs,
or mixed
|Status|Implemented||
|Type|Gap||
.11+|Storage|IP|Implemented||
|Display IP|Implemented||
|Name|Implemented||
|Manufacturer|Implemented||
|Model|Implemented||
|Family|Implemented||The storage Family could be Clariion, Symmetrix, et
|Microcode Version|Implemented||
|Virtual|Implemented||Is this a storage virtualization device?
|Total Raw Capacity|Implemented||Total raw capacity (sum of all disks on
the array)
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|Failed Raw Capacity|Implemented||Raw capapcity of failed disks (sum of
all disks that are failed)
.14+|Storage Pool|Storage Pool Id|Implemented||
|Name|Implemented||
|Type|Gap||
|Thin Provisioning Supported|Implemented||Whether this internal volume
supports thin provisioning for the volume layer on top of it
|Include In Dwh Capacity|Implemented||A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Virtual|Implemented||Is this a storage virtualization device?
|Raid Group|Implemented||indicates whether this storagePool is a raid
group
|Snapshot Used Capacity|Implemented||
|Data Used Capacity|Implemented||
|Data Allocated Capacity|Gap||capacity allocated for data
|Total Used Capacity|Implemented||Total capacity in MB
```

```
|Total Allocated Capacity|Implemented||
|Physical Disk Capacity (MB)|Implemented||used as raw capacity for
storage pool
|Raw to Usable Ratio|Implemented||ratio to convert from usable capacity to
raw capacity
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Cloud Volumes Service REST API
IHTTPS
|HTTPS
1443
|true
Itrue
|true
|true
|===
<<top, Back to Top>>
== Amazon FSx for NetApp ONTAP
:description: Support Matrix Asciidoc for Amazon FSx for NetApp ONTAP
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|FSx for ONTAP
|Data ONTAP
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.200+|foundation .4+|Capability|Active|Implemented|HTTPS|
|Licensed|Implemented|HTTPS|
```

```
|Type|Gap|HTTPS|
|Used|Implemented|HTTPS|
.8+|File Share|Is InternalVolume|Implemented|HTTPS|whether the file share
represents an internal volume (netapp volume) or is it a qtree/folder
within the internal volume
| Is Oplocks Enabled | Implemented | HTTPS | whether opportunistic locks are
enabled on the FileShare
|Is Shared|Implemented|HTTPS|whether this fileShare has any shares
associated with it
|Name|Implemented|HTTPS|
|Path|Implemented|HTTPS|path of the fileShare
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Security Type|Implemented|HTTPS|
|Status|Implemented|HTTPS|
.4+|ISCSI Network Portal|IP|Implemented|HTTPS|
|Listening Port|Implemented|HTTPS|
|Nic|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.3+|ISCSI Network Portal Group|OID|Implemented|HTTPS|
|Portal Group Name|Implemented|HTTPS|
|Portal Group Tag|Implemented|HTTPS|
.4+|ISCSI Node|Host Aliases|Implemented|HTTPS|
|Node Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
.2+|ISCSI Node Map|OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
.6+|ISCSI Security Info|Auth Type|Implemented|HTTPS|
|Inbound Keyword|Implemented|HTTPS|
|Initiator Node Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Outbound Keyword|Implemented|HTTPS|
|Target Node Name|Implemented|HTTPS|
.9+|ISCSI Session|Initiator Ips|Implemented|HTTPS|
|Initiator OID|Implemented|HTTPS|
|Initiator Session Id|Implemented|HTTPS|
|Max Connections|Implemented|HTTPS|
|Number Of Connections|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
|Security|Implemented|HTTPS|
|Target Session Id|Implemented|HTTPS|
.6+|Info|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
```

```
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.36+|Internal Volume|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|HTTPS|capacity allocated for data
|Data Used Capacity|Implemented|HTTPS|
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|GuidKey 1|Implemented|HTTPS|GuidKey1 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|HTTPS|GuidKey2 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|Internal Volume Id|Implemented|HTTPS|
|Last Known Access Time|Implemented|HTTPS|Last know access to the volume
|Name|Implemented|HTTPS|
|Protection Type|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Snapshot Allocated Capacity|Gap|HTTPS|Allocated capacity of snapshots in
|Snapshot Count|Implemented|HTTPS|Number of snapshots on the internal
volumes
|Snapshot Used Capacity|Implemented|HTTPS|
|Space Guarantee|Implemented|HTTPS|Space Guatantee policy (file, volume
or none)
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Clone Saved Capacity in MB|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Total Used Capacity (MB)|Implemented|HTTPS|place holder for the used
capacity as read from the device
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Virtual Storage|Implemented|HTTPS|Owning virtual storage (vfiler)
|Comment|Gap|HTTPS|state: free text comment describing the svm
|Group Id|Implemented|HTTPS|
|Group Name|Implemented|HTTPS|
|Encrypted|Implemented|HTTPS|
||Implemented|HTTPS|
||Implemented|HTTPS|
```

```
||Implemented|HTTPS|
|Tiering Minimum Cooling Days|Implemented|HTTPS|
.3+|Storage Synchronization|Source Internal Volume|Implemented|HTTPS|
|Target Internal Volume|Implemented|HTTPS|
|Technology|Implemented|HTTPS|technology which causes storage efficiency
changed
.12+|QTree|GuidKey 1|Implemented|HTTPS|GuidKey1 is implicit for all
objects whose GUID key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|HTTPS|GuidKey2 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|GuidKey 3|Implemented|HTTPS|GuidKey3 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|Name|Implemented|HTTPS|
|Oplocks|Implemented|HTTPS|Indicates whether opportunistic locks are
enabled on the Qtree
|Qtree Id|Implemented|HTTPS|unique id of the gtree
|Quota HardCapacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk
space, allowed for the quota target
|Quota SoftCapacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk
space, allowed for the quota target
|Quota UsedCapacity|Implemented|HTTPS|Space in MB currently used
|Security Style|Implemented|HTTPS|Security style of the directory: unix,
ntfs, or mixed
|Status|Implemented|HTTPS|
|Type|Gap|HTTPS|
.15+|Quota|GuidKey 1|Implemented|HTTPS|GuidKey1 is implicit for all
objects whose GUID key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|HTTPS|GuidKey2 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|GuidKey 3|Implemented|HTTPS|GuidKey3 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
| Hard Capacity Limit (MB) | Implemented | HTTPS | max amount of disk space,
allowed for the quota target (Hard limit)
|Hard File Limit|Implemented|HTTPS|max number of files allowed for the
quota target
|Internal Volume Id|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the gtree
|Quota Id|Implemented|HTTPS|unique id of the quota
|Soft Capacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk space,
allowed for the quota target
|Soft File Limit|Implemented|HTTPS|Max number of files allowed for the
quota target
|Threshold (MB)|Implemented|HTTPS|Disk space threshold, for the quota
target
|Type|Gap|HTTPS|
|Used Capacity|Implemented|HTTPS|
```

```
|Used Files|Implemented|HTTPS|Number of files currently used
|User/Group Target|Implemented|HTTPS|user/group target this quota refers
to
.4+|Share|Description|Implemented|HTTPS|
|IP Interfaces|Implemented|HTTPS|comma separated list of IP addresses on
which this share is exposed
|Name|Implemented|HTTPS|
|Protocol|Implemented|HTTPS|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|HTTPS|
|Permission|Implemented|HTTPS|Permissions for this particular share
.18+|Storage|Cluster|Not Available|HTTPS|Whether this storage is a cluster
|Cpu Count|Implemented|HTTPS|Cpu Count of the storage
|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Memory|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|HTTPS|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|UUID|Implemented|HTTPS|
.6+|Storage Node|Cache Size|Implemented|HTTPS|device cache size in MB
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
|Version|Implemented|HTTPS|software version
.23+|Storage Pool|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|HTTPS|capacity allocated for data
|Data Used Capacity|Implemented|HTTPS|
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
```

```
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
| Include In Dwh Capacity | Implemented | HTTPS | A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
|Snapshot Allocated Capacity|Gap|HTTPS|Allocated capacity of snapshots in
|Snapshot Used Capacity|Implemented|HTTPS|
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Uses FlashPools|Implemented|HTTPS|indicate if this storagepool uses Flash
Pools(NetApp specific)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|UUID|Implemented|HTTPS|
.11+|Storage VirtualMachine|Allocated capacity
(MB) | Implemented | HTTPS | Allocated capacity of snapshots
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|Guid Key|Implemented|HTTPS|Globally Unique Key of the storage virtual
machine
| IP Space | Implemented | HTTPS |
|Name|Implemented|HTTPS|
|Protocols|Implemented|HTTPS|csv of svm enabled protocols (fcp, iscsi,
nfs, cifs...)
|State|Implemented|HTTPS|free text describing the device state
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
.1+|VirtualMachine Internal Volume Ref|Internal Volume
Id|Implemented|HTTPS|
.1+|VirtualMachine StoragePool lRef|Storage Pool Id|Implemented|HTTPS|
.1+|VirtualMachine Volume Ref|Volume Name|Implemented|HTTPS|
.13+|Volume|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|DiskGroup|Implemented|HTTPS|Disk Group Type
```

```
|Junction Path|Implemented|HTTPS|
|Last Known Access Time|Implemented|HTTPS|Last know access to the volume
|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Encrypted|Implemented|HTTPS|
.4+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|performance .0+|Qtree.1+|Storage|Failed Disks|Implemented|HTTPS|
.3+|Storage Node|Cache Hit Ratio Total|Implemented|HTTPS|
|Total Disk Read Replaced|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|NetApp ONTAP API
|HTTP/HTTPS
|HTTP/HTTPS
180/443
Itrue
|true
|true
|true
|===
```

```
<<top, Back to Top>>
== NetApp Clustered Data ONTAP 8.1.1+
:description: Support Matrix Asciidoc for NetApp Clustered Data ONTAP
8.1.1+
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|AFF-A150
AFF-A200
AFF-A220
AFF-A250
AFF-A300
AFF-A320
AFF-A400
AFF-A700
AFF-A700s
AFF-A800
AFF-A900
AFF-C190
AFF-C250
AFF-C400
AFF-C800
AFF8020
AFF8040
AFF8060
AFF8080
CDvM100
CDvM200
DM5000H
FAS2240-2
FAS2240-4
FAS2520
FAS2552
FAS2554
FAS2620
FAS2650
FAS2720
FAS2750
FAS3220
FAS3250
FAS3270
FAS500f
FAS6210
```

FAS6220 FAS8020 FAS8040 FAS8060 FAS8080 FAS8200 FAS8300 FAS8700 FAS9000 FAS9500 FASDvM300 SIMBOX V6240 |8.2.3P5 8.3.0 8.3.1 8.3.1P2 8.3.2 8.3.2P12 8.3.2P2 8.3.2P5 9.0.1 9.1.0 9.1.0P1 9.1.0P10 9.1.0P11 9.1.0P12 9.1.0P14 9.1.0P15 9.1.0P17 9.1.0P19 9.1.0P20 9.1.0P5 9.1.0P7 9.1.0P8 9.10.0 9.10.1 9.10.1P1 9.10.1P10 9.10.1P11 9.10.1P12 9.10.1P13 9.10.1P2 9.10.1P3 9.10.1P4 9.10.1P5

```
9.10.1P6
9.10.1P7
9.10.1P8
9.10.1P9
9.11.0P1
9.11.1
9.11.1P1
9.11.1P10
9.11.1P2
9.11.1P3
9.11.1P4
9.11.1P5
9.11.1P6
9.11.1P7
9.11.1P8
9.11.1P9
9.11.1X12
9.11.1X26
9.12.1
9.12.1P1
9.12.1P2
9.12.1P3
9.12.1P4
9.13.0
9.13.0P1
9.13.0P2
9.13.0P3
9.13.1
9.13.1X19
9.13.1X25
9.13.1X33
9.13.1X34
9.13.1X35
9.13.1X36
9.3.0P1
9.3.0P10
9.3.0P12
9.3.0P13
9.3.0P14
9.3.0P15
9.3.0P18
9.3.0P19
9.3.0P2
9.3.0P20
9.3.0P21
9.3.0P4
```

- 9.3.0P6 9.3.0P7 9.3.0P8 9.3.0P9
 - 9.4.0
 - 9.4.0P2
 - 9.4.0P3
 - 9.4.0P4
 - 9.4.0P6
 - 9.5.0
 - 9.5.0P1
 - 9.5.0P10
 - 9.5.0P12
 - 9.5.0P13
 - 9.5.0P14
 - 9.5.0P15
 - 9.5.0P16
 - 9.5.0P17
 - 9.5.0P18
 - 9.5.0P19
 - 9.5.0P2
 - 9.5.0P3
 - 9.5.0P4
 - 9.5.0P5
 - 9.5.0P6
 - 9.5.0P7
 - 9.5.0P8
 - 9.5.0P9
 - 9.6.0
 - 9.6.0P11
 - 9.6.0P14
 - 9.6.0P15
 - 9.6.0P18
 - 9.6.0P2
 - 9.6.0P3
 - 9.6.0P4
 - 9.6.0P5
 - 9.6.0X11
 - 9.7.0
 - 9.7.0P1
 - 9.7.0P10
 - 9.7.0P11
 - 9.7.0P12
 - 9.7.0P13
 - 9.7.0P14
 - 9.7.0P15

9.7.0P16 9.7.0P17 9.7.0P18 9.7.0P19 9.7.0P2 9.7.0P20 9.7.0P21 9.7.0P22 9.7.0P3 9.7.0P4 9.7.0P5 9.7.0P6 9.7.0P7 9.7.0P8 9.7.0P9 9.7.0X17 9.8.0 9.8.0P1 9.8.0P10 9.8.0P11 9.8.0P11D2 9.8.0P12 9.8.0P13 9.8.0P14 9.8.0P15 9.8.0P16 9.8.0P17 9.8.0P18 9.8.0P19 9.8.0P2 9.8.0P3 9.8.0P4 9.8.0P5 9.8.0P6 9.8.0P7 9.8.0P8 9.8.0P9 9.9.0 9.9.1 9.9.1P1 9.9.1P10 9.9.1P11 9.9.1P12 9.9.1P13 9.9.1P14 9.9.1P15

```
9.9.1P16
9.9.1P2
9.9.1P3
9.9.1P4
9.9.1P5
9.9.1P6
9.9.1P7
9.9.1P8
9.9.1P9
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.258+|foundation .4+|Capability|Active|Implemented|HTTPS|
|Licensed|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Used|Implemented|HTTPS|
.15+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Group|Implemented|HTTPS|
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type|Gap|HTTPS|
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
.5+|Disk Group|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|DiskGroup Id|Implemented|HTTPS|unique id of the disk group
|Name|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.8+|File Share|Is InternalVolume|Implemented|HTTPS|whether the file share
represents an internal volume (netapp volume) or is it a qtree/folder
within the internal volume
|Is Oplocks Enabled|Implemented|HTTPS|whether opportunistic locks are
enabled on the FileShare
|Is Shared|Implemented|HTTPS|whether this fileShare has any shares
```

```
associated with it
|Name|Implemented|HTTPS|
|Path|Implemented|HTTPS|path of the fileShare
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Security Type|Implemented|HTTPS|
|Status|Implemented|HTTPS|
.4+|ISCSI Network Portal|IP|Implemented|HTTPS|
|Listening Port|Implemented|HTTPS|
|Nic|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.3+|ISCSI Network Portal Group|OID|Implemented|HTTPS|
|Portal Group Name|Implemented|HTTPS|
|Portal Group Tag|Implemented|HTTPS|
.4+|ISCSI Node|Host Aliases|Implemented|HTTPS|
|Node Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
.2+|ISCSI Node Map|OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
.6+|ISCSI Security Info|Auth Type|Implemented|HTTPS|
|Inbound Keyword|Implemented|HTTPS|
| Initiator Node Name | Implemented | HTTPS |
|OID|Implemented|HTTPS|
|Outbound Keyword|Implemented|HTTPS|
|Target Node Name|Implemented|HTTPS|
.9+|ISCSI Session|Initiator Ips|Implemented|HTTPS|
|Initiator OID|Implemented|HTTPS|
|Initiator Session Id|Implemented|HTTPS|
|Max Connections|Implemented|HTTPS|
|Number Of Connections|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
|Security|Implemented|HTTPS|
|Target Session Id|Implemented|HTTPS|
.6+|Info|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.43+|Internal Volume|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|HTTPS|capacity allocated for data
|Data Used Capacity|Implemented|HTTPS|
```

```
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|FlashPool Eligibility|Implemented|HTTPS|Whether the internal volume can
participate in hybrid caching
|GuidKey 1|Implemented|HTTPS|GuidKey1 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|HTTPS|GuidKey2 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|Internal Volume Id|Implemented|HTTPS|
|Last Known Access Time|Implemented|HTTPS|Last know access to the volume
|Name|Implemented|HTTPS|
|Other Allocated Capacity|Gap|HTTPS|Capacity allocated for other (not data
and not snapshot)
|Protection Type|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Snapshot Allocated Capacity|Gap|HTTPS|Allocated capacity of snapshots in
MB
|Snapshot Count|Implemented|HTTPS|Number of snapshots on the internal
|Snapshot Used Capacity|Implemented|HTTPS|
|Space Guarantee|Implemented|HTTPS|Space Guatantee policy (file, volume
or none)
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Clone Saved Capacity in MB|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Total Used Capacity (MB)|Implemented|HTTPS|place holder for the used
capacity as read from the device
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Virtual Storage|Implemented|HTTPS|Owning virtual storage (vfiler)
|Adaptive Qos Policy|Implemented|HTTPS|
|Comment|Gap|HTTPS|state: free text comment describing the svm
|Group Id|Implemented|HTTPS|
|Group Name|Implemented|HTTPS|
|Encrypted|Implemented|HTTPS|
||Implemented|HTTPS|
||Implemented|HTTPS|
||Implemented|HTTPS|
|Qos Limit IOPS|Implemented|HTTPS|
|Qos Limit MBPS|Implemented|HTTPS|
```

```
|Qos Limit Raw|Implemented|HTTPS|
|QoS - Policy|Implemented|HTTPS|
|Tiering Minimum Cooling Days|Implemented|HTTPS|
.3+|Storage Synchronization|Source Internal Volume|Implemented|HTTPS|
|Target Internal Volume|Implemented|HTTPS|
|Technology|Implemented|HTTPS|technology which causes storage efficiency
changed
.12+|QTree|GuidKey 1|Implemented|HTTPS|GuidKey1 is implicit for all
objects whose GUID key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|HTTPS|GuidKey2 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|GuidKey 3|Implemented|HTTPS|GuidKey3 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|Name|Implemented|HTTPS|
|Oplocks|Implemented|HTTPS|Indicates whether opportunistic locks are
enabled on the Qtree
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Quota HardCapacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk
space, allowed for the quota target
|Quota SoftCapacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk
space, allowed for the quota target
|Quota UsedCapacity|Implemented|HTTPS|Space in MB currently used
|Security Style|Implemented|HTTPS|Security style of the directory: unix,
ntfs, or mixed
|Status|Implemented|HTTPS|
|Type|Gap|HTTPS|
.15+|Quota|GuidKey 1|Implemented|HTTPS|GuidKey1 is implicit for all
objects whose GUID key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|HTTPS|GuidKey2 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|GuidKey 3|Implemented|HTTPS|GuidKey3 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
| Hard Capacity Limit (MB) | Implemented | HTTPS | max amount of disk space,
allowed for the quota target (Hard limit)
|Hard File Limit|Implemented|HTTPS|max number of files allowed for the
quota target
|Internal Volume Id|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Quota Id|Implemented|HTTPS|unique id of the quota
|Soft Capacity Limit (MB)|Implemented|HTTPS|Maximum amount of disk space,
allowed for the quota target
|Soft File Limit|Implemented|HTTPS|Max number of files allowed for the
quota target
|Threshold (MB)|Implemented|HTTPS|Disk space threshold, for the quota
target
|Type|Gap|HTTPS|
```

```
|Used Capacity|Implemented|HTTPS|
|Used Files|Implemented|HTTPS|Number of files currently used
|User/Group Target|Implemented|HTTPS|user/group target this quota refers
to
.4+|Share|Description|Implemented|HTTPS|
| IP Interfaces | Implemented | HTTPS | comma separated list of IP addresses on
which this share is exposed
|Name|Implemented|HTTPS|
|Protocol|Implemented|HTTPS|enum for share protocol
.2+|Share Initiator|Initiator|Implemented|HTTPS|
|Permission|Implemented|HTTPS|Permissions for this particular share
.18+|Storage|Cluster|Not Available|HTTPS|Whether this storage is a cluster
|Cpu Count|Implemented|HTTPS|Cpu Count of the storage
|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Memory|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|HTTPS|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|UUID|Implemented|HTTPS|
.12+|Storage Node|Cache Size|Implemented|HTTPS|device cache size in MB
|Memory Size|Gap|HTTPS|device memory in MB
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Processors Count|Implemented|HTTPS|device CPU
|Partner Node UUID|Implemented|HTTPS|HA pair's UUID
|Serial Number|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
|Up Time|Implemented|HTTPS|time in milliseconds
|Version|Implemented|HTTPS|software version
|ManagementIp Addresses|Implemented|HTTPS|
```

```
.38+|Storage Pool|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|HTTPS|capacity allocated for data
|Data Used Capacity|Implemented|HTTPS|
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
|Reserved Capacity|Implemented|HTTPS|Reserved Capacity in MB
|Snapshot Allocated Capacity|Gap|HTTPS|Allocated capacity of snapshots in
MB
|Snapshot Used Capacity|Implemented|HTTPS|
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Uses FlashPools|Implemented|HTTPS|indicate if this storagepool uses Flash
Pools(NetApp specific)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|CompactionSavings Enabled|Implemented|HTTPS|
|Encrypted|Implemented|HTTPS|
|License Used Percent|Implemented|HTTPS|
||Implemented|HTTPS|
```

```
|UUID|Implemented|HTTPS|
.13+|Storage VirtualMachine|Allocated capacity
(MB) | Implemented | HTTPS | Allocated capacity of snapshots
|Comment|Implemented|HTTPS|
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|Guid Key|Implemented|HTTPS|Globally Unique Key of the storage virtual
machine
| IP Space | Implemented | HTTPS |
|Internal Volume Limit|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Protocols|Implemented|HTTPS|csv of svm enabled protocols (fcp, iscsi,
nfs, cifs...)
|State|Implemented|HTTPS|free text describing the device state
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
.1+|VirtualMachine Internal Volume Ref|Internal Volume
Id|Implemented|HTTPS|
.1+|VirtualMachine StoragePool lRef|Storage Pool Id|Implemented|HTTPS|
.1+|VirtualMachine Volume Ref|Volume Name|Implemented|HTTPS|
.21+|Volume|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|DiskGroup|Implemented|HTTPS|Disk Group Type
|Disk Size|Implemented|HTTPS|comma seperated list of disk sizes (GB)
|Disk Speed|Implemented|HTTPS|comma seperated list of disk speeds (rpm)
|Disk Type|Not Available|HTTPS|
|Junction Path|Implemented|HTTPS|
|Last Known Access Time|Implemented|HTTPS|Last know access to the volume
|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|Encrypted|Implemented|HTTPS|
|Qos Limit IOPS|Implemented|HTTPS|
|Qos Limit MBPS|Implemented|HTTPS|
|Qos Limit Raw|Implemented|HTTPS|
|QoS - Policy|Implemented|HTTPS|
.4+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Protocol Controller|Implemented|HTTPS|
```

```
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.142+|performance .9+|Disk|IOps Read|Implemented|HTTPS|Number of read IOps
on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|Read Utilization|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
|Utilization Write|Implemented|HTTPS|
.31+|Internal Volume|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
|Clone Saved Capacity Total|Implemented|HTTPS|
|Total Compression Savings|Implemented|HTTPS|
|Compression Savings Space|Implemented|HTTPS|
|Total Data Capacity|Implemented|HTTPS|
|Data Used Capacity|Implemented|HTTPS|
|deduplicationSavingsPercent.total|Implemented|HTTPS|
|Deduplication Savings Space|Implemented|HTTPS|
|File System Capacity Logical Used|Implemented|HTTPS|
|File System Capacity Physical Available|Implemented|HTTPS|
|File System Capacity Physical Used|Implemented|HTTPS|
| IO Density Read | Implemented | HTTPS |
| IO Density Total | Implemented | HTTPS |
|Write IO Density|Implemented|HTTPS|
| IOPs other | Implemented | HTTPS |
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Other Total Capacity|Implemented|HTTPS|
|QOS Wait|Implemented|HTTPS|Total qos wait time
|Snapshot Reserved Capacity|Implemented|HTTPS|
|Snapshot Used Capacity|Implemented|HTTPS|
|Snapshot Used Capacity Ratio|Implemented|HTTPS| Reported as a time series
|Throughput Read|Implemented|HTTPS|
```

```
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.5+|Qtree|Quota Hard Limit|Implemented|HTTPS|Capacity quota hard limit
|Quota Soft Limit|Implemented|HTTPS|Capacity Quota soft Limit
|Used Capacity|Implemented|HTTPS|
|Total File Count|Implemented|HTTPS|
| IOPs Total | Implemented | HTTPS |
.18+|Storage|Failed Raw Capacity|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|HTTPS|
|Failed Disks|Implemented|HTTPS|
| IO Density Read | Implemented | HTTPS |
| IO Density Total | Implemented | HTTPS |
|Write IO Density|Implemented|HTTPS|
| IOPs other | Implemented | HTTPS |
| IOps Read | Implemented | HTTPS | Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.16+|Storage Node|Cache Hit Ratio Total|Implemented|HTTPS|
|Total Disk Read Replaced|Implemented|HTTPS|
|Failed Disks|Implemented|HTTPS|
|Failed Fans|Implemented|HTTPS|
|Failed Power Units|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|Optimal Utilization|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
.32+|StoragePool Disk|Capacity Provisioned|Implemented|HTTPS|
```

```
|Raw Capacity|Implemented|HTTPS|
|Reserved Capacity|Implemented|HTTPS|Fractional reserved (included in data
allocated capacity)
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio|Implemented|HTTPS|Reported as a time series
|Capacity Used Ratio|Implemented|HTTPS|
|compactionSavingsPercent.total|Implemented|HTTPS|
|Total Compaction Space|Implemented|HTTPS|
|Total Compression Savings|Implemented|HTTPS|
|Compression Savings Space|Implemented|HTTPS|
|Total Data Capacity|Implemented|HTTPS|
|Data Used Capacity|Implemented|HTTPS|
|deduplicationSavingsPercent.total|Implemented|HTTPS|
|Deduplication Savings Space|Implemented|HTTPS|
|Failed Disks|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Object Store Space Physical Space Used|Implemented|HTTPS|
|Object Store Space Referenced Space|Implemented|HTTPS|
|Object Store Space SIS Space Saved|Implemented|HTTPS|
|Object Store Space Used|Implemented|HTTPS|
|Snapshot Reserved Capacity|Implemented|HTTPS|
|Snapshot Used Capacity|Implemented|HTTPS|
|Snapshot Used Capacity Ratio|Implemented|HTTPS| Reported as a time series
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|Read Utilization|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
|Utilization Write|Implemented|HTTPS|
.14+|Virtual Machine|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
| IO Density Read | Implemented | HTTPS |
| IO Density Total | Implemented | HTTPS |
|Write IO Density|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
```

```
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.17+|Volume|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
| IO Density Read | Implemented | HTTPS |
| IO Density Total | Implemented | HTTPS |
|Write IO Density|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Partial Blocked Ratio|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|NetApp ONTAP API
|HTTP/HTTPS
|HTTP/HTTPS
180/443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== NetApp SolidFire 8.1+
:description: Support Matrix Asciidoc for NetApp SolidFire 8.1+
```

```
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
| H410S-2
H610S-2
 H610S-4
 SF19210
 SF2405
 SF38410
 SF4805
 SF9605
 SF9608
FCN001
H300S
H410S-0
H410S-1
H410S-2
H500S
H610S-1
H610S-2
H610S-4
H610S2
SF19210
SF38410
SF4805
SF9605
111.1.0.72
11.5.0.63
11.7.0.76
11.8.0.23
12.0.0.333
12.2.0.777
12.3.0.958
12.3.1.103
12.3.1.165
12.3.2.3
12.5.0.897
12.7.0.380
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
```

```
.117+|foundation .18+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Group|Implemented|HTTPS|
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Potential Transfer Rate|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Seek Time|Implemented|HTTPS|
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type|Gap|HTTPS|
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
|Encrypted|Implemented|HTTPS|
.4+|ISCSI Network Portal|IP|Implemented|HTTPS|
|Listening Port|Implemented|HTTPS|
|Nic|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.3+|ISCSI Network Portal Group|OID|Implemented|HTTPS|
|Portal Group Name|Implemented|HTTPS|
|Portal Group Tag|Implemented|HTTPS|
.4+|ISCSI Node|Host Aliases|Implemented|HTTPS|
|Node Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
.2+|ISCSI Node Map|OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
.7+|ISCSI Session|Initiator Ips|Implemented|HTTPS|
|Initiator OID|Implemented|HTTPS|
|Max Connections|Implemented|HTTPS|
|Number Of Connections|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
|Security|Implemented|HTTPS|
.4+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.15+|Storage|Cpu Count|Implemented|HTTPS|Cpu Count of the storage
|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
```

```
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|HTTPS|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.8+|Storage Node|Memory Size|Gap|HTTPS|device memory in MB
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Processors Count|Implemented|HTTPS|device CPU
|Serial Number|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
|Version|Implemented|HTTPS|software version
|ManagementIp Addresses|Implemented|HTTPS|
.22+|Storage Pool|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Data Allocated Capacity|Gap|HTTPS|capacity allocated for data
|Data Used Capacity|Implemented|HTTPS|
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTPS|ratio of dedupe savings in percentage
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Other Allocated Capacity|Gap|HTTPS|Capacity allocated for other (not data
and not snapshot)
|Other UsedCapacity (MB)|Implemented|HTTPS|Any capacity other than data
and snapshot
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
```

```
|Status|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|Encrypted|Implemented|HTTPS|
.21+|Volume|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Junction Path|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|HTTPS|Redundancy level
|Replica Source|Implemented|HTTPS|
|Replica Target|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|Account Identifier|Implemented|HTTPS|
|Account Name|Implemented|HTTPS|
|Compression Enabled|Implemented|HTTPS|
|Encrypted|Implemented|HTTPS|
|Qos Burst IOPS|Implemented|HTTPS|
|Qos Limit IOPS|Implemented|HTTPS|
|qos Min IOPS|Implemented|HTTPS|
|QoS - Policy|Implemented|HTTPS|
.5+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Masking Required|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.62+|performance .6+|Disk|IOps Read|Implemented|HTTPS|Number of read IOps
on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
```

```
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.16+|Storage|Failed Raw Capacity|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|HTTPS|
| IOPs other | Implemented | HTTPS |
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Partial Blocked Ratio|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
.10+|Storage Node|IOps Read|Implemented|HTTPS|Number of read IOps on the
disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
.14+|StoragePool Disk|Capacity Provisioned|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Over Commit Capacity Ratio|Implemented|HTTPS|Reported as a time series
|Capacity Used Ratio|Implemented|HTTPS|
|Total Compression Savings|Implemented|HTTPS|
|Compression Savings Space|Implemented|HTTPS|
|Total Data Capacity|Implemented|HTTPS|
|Data Used Capacity|Implemented|HTTPS|
|deduplicationSavingsPercent.total|Implemented|HTTPS|
|Deduplication Savings Space|Implemented|HTTPS|
|Other Total Capacity|Implemented|HTTPS|
|Other Used Capacity|Implemented|HTTPS|
```

```
.16+|Volume|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
|Total Compression Savings|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Partial Blocked Ratio|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|SolidFire REST API
|HTTPS
HTTPS
1443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== NetApp StorageGRID (HTTPS)
:description: Support Matrix Asciidoc for NetApp StorageGRID (HTTPS)
Models and versions supported by this data collector:
|===
```

```
<.<|API versions <.<|Models <.<|Firmware versions
13.0
3.2
3.3
3.4
3.5
|Webscale
|11.2.0
11.4.0
11.4.0.3
11.4.0.4
11.5.0.1
11.5.0.11
11.5.0.2
11.5.0.3
11.5.0.6
11.5.0.7
11.5.0.8
11.5.0.9
11.6.0
11.6.0.1
11.6.0.10
11.6.0.2
11.6.0.4
11.6.0.5
11.6.0.7
11.6.0.8
11.6.0.9
11.7.0
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.60+|foundation .7+|Info|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
|Version|Implemented|HTTPS|software version
.13+|Internal Volume|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on
the storage pool
```

```
|GuidKey 1|Implemented|HTTPS|GuidKey1 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|HTTPS|GuidKey2 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|Internal Volume Id|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Total Used Capacity (MB)|Implemented|HTTPS|place holder for the used
capacity as read from the device
|Type|Gap|HTTPS|
.5+|QTree|GuidKey 1|Implemented|HTTPS|GuidKey1 is implicit for all objects
whose GUID key has not changed since OCI version 7.3.5.
|GuidKey 2|Implemented|HTTPS|GuidKey2 is implicit for all objects whose
GUID key has not changed since OCI version 7.3.5.
|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Type|Gap|HTTPS|
.13+|Storage|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.11+|Storage Node|Cache Size|Implemented|HTTPS|device cache size in MB
|Memory Size|Gap|HTTPS|device memory in MB
|Name|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
|ManagementIp Addresses|Implemented|HTTPS|
```

```
|Node Capacity Utilization Allowed Meta Data in MB|Implemented|HTTPS|
| Node Capacity Utilization Total in MB | Implemented | HTTPS |
|Node Capacity Utilization Usable in MB|Implemented|HTTPS|
|Node Capacity Utilization Used in MB|Implemented|HTTPS|
|Node Capacity Utilization Used Meta Data in MB|Implemented|HTTPS|
|Site Name|Implemented|HTTPS|
.11+|Storage Pool|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ
to cotnrol which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.22+|performance .4+|Internal Volume|Total Capacity|Implemented||
|Used Capacity|Implemented||
|Capacity Used Ratio|Implemented||
|Total Objects|Implemented||
.0+|Qtree.4+|Storage|Failed Raw Capacity|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|StoragePools Capacity|Implemented||
|Raw Capacity|Implemented||
.8+|Storage Node|Node Capacity Utilization Allowed Meta Data|Implemented||
|Node Capacity Utilization Total|Implemented||
|Node Capacity Utilization Usable|Implemented||
|Node Capacity Utilization Used|Implemented||
|Node Capacity Utilization Used Meta Data|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.6+|StoragePool Disk|Capacity Provisioned|Implemented||
|Raw Capacity|Implemented||
|Total Capacity|Implemented||
|Used Capacity|Implemented||
|Over Commit Capacity Ratio|Implemented||Reported as a time series
|Capacity Used Ratio|Implemented||
```

```
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|StorageGrid REST API
|HTTPS
| HTTPS
1443
Itrue
|false
|true
|true
|===
<<top, Back to Top>>
== Nutanix Storage (REST)
:description: Support Matrix Asciidoc for Nutanix Storage (REST)
Models and versions supported by this data collector:
|===
<.<|Models <.<|Firmware versions
| HPE DL360-8 G10
NX-3060-G6
 NX-3170-G6
NX-8035-G6
 NX-8150-G7
HPE DL360-8 G10
HPE DL380-12 G10
NX-3060-G5
NX-3170-G7
NX-5155-G6
NX-8035-G6
NX-8035-G7
NX-8150-G7
NX-8150-G8
16.5.1.6
```

```
6.5.2
6.5.2.5
6.5.2.6
6.5.2.7
6.5.3
6.5.3.1
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.121+|foundation .14+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Location|Gap|HTTPS|Where this disk is physically located in the array
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Serial Number|Implemented|HTTPS|
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type|Gap|HTTPS|
|Type Enum|Implemented|HTTPS|enum for disk type
|Vendor|Implemented|HTTPS|
.5+|File Share|Is InternalVolume|Implemented|HTTPS|whether the file share
represents an internal volume (netapp volume) or is it a gtree/folder
within the internal volume
| Is Shared | Implemented | HTTPS | whether this fileShare has any shares
associated with it
|Name|Implemented|HTTPS|
|Path|Implemented|HTTPS|path of the fileShare
|Qtree Id|Implemented|HTTPS|unique id of the qtree
.4+|ISCSI Network Portal|IP|Implemented|HTTPS|
|Listening Port|Implemented|HTTPS|
|Nic|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.3+|ISCSI Network Portal Group|OID|Implemented|HTTPS|
|Portal Group Name|Implemented|HTTPS|
|Portal Group Tag|Implemented|HTTPS|
.3+|ISCSI Node|Node Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
.2+|ISCSI Node Map|OID|Implemented|HTTPS|
```

```
|Portal Group OID|Implemented|HTTPS|
.7+|ISCSI Session|Initiator Ips|Implemented|HTTPS|
|Initiator OID|Implemented|HTTPS|
|Max Connections|Implemented|HTTPS|
|Number Of Connections|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Portal Group OID|Implemented|HTTPS|
|Security|Implemented|HTTPS|
.5+|Info|Api Name|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.14+|Internal Volume|Compression Enabled|Implemented|HTTPS|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTPS|ratio of compression savings in
percentage
|Dedupe Enabled|Implemented|HTTPS|Is dedupe enabled on the storage pool
|Internal Volume Id|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Total Used Capacity (MB)|Implemented|HTTPS|place holder for the used
capacity as read from the device
|Type|Gap|HTTPS|
|UUID|Implemented|HTTPS|
.3+|Storage Synchronization|Source Internal Volume|Implemented|HTTPS|
|Target Internal Volume|Implemented|HTTPS|
|Technology|Implemented|HTTPS|technology which causes storage efficiency
changed
.3+|QTree|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Type|Gap|HTTPS|
.3+|Share|IP Interfaces|Implemented|HTTPS|comma separated list of IP
addresses on which this share is exposed
|Name|Implemented|HTTPS|
|Protocol|Implemented|HTTPS|enum for share protocol
.1+|Share Initiator|Initiator|Implemented|HTTPS|
.15+|Storage|Cluster|Not Available|HTTPS|Whether this storage is a cluster
|Display IP|Implemented|HTTPS|
```

```
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Serial Number|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|SupportActive Active|Implemented|HTTPS|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.7+|Storage Node|Memory Size|Gap|HTTPS|device memory in MB
|Model|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Processors Count|Implemented|HTTPS|device CPU
|Serial Number|Implemented|HTTPS|
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
.15+|Storage Pool|Auto Tiering|Implemented|HTTPS|indicates if this
storagepool is participating in auto tiering with other pools
|Data Used Capacity|Implemented|HTTPS|
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTPS|
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTPS|Redundancy level
|Reserved Capacity|Implemented|HTTPS|Reserved Capacity in MB
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
```

```
.9+|Volume|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|Junction Path|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|Qtree Id|Implemented|HTTPS|unique id of the qtree
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Redundancy|Implemented|HTTPS|Redundancy level
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioned|Implemented|HTTPS|
|UUID|Implemented|HTTPS|
.4+|Volume Map|LUN|Implemented|HTTPS|Name of the backend lun
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|Volume Mask|Initiator|Implemented|HTTPS|
|Protocol Controller|Implemented|HTTPS|
|Storage Port|Implemented|HTTPS|
|Type|Gap|HTTPS|
.54+|performance .6+|Disk|IOps Read|Implemented|HTTPS|Number of read IOps
on the disk
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Write|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
| IOPs Total | Implemented | HTTPS |
.10+|Internal Volume|IOPs other|Implemented|HTTPS|
| IOps Read | Implemented | HTTPS | Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.0+|Qtree.14+|Storage|Failed Raw Capacity|Implemented|HTTPS|
|Raw Capacity|Implemented|HTTPS|
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|StoragePools Capacity|Implemented|HTTPS|
| IOPs other | Implemented | HTTPS |
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
```

```
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.9+|Storage Node|Latency Total|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
|Latency Read|Implemented|HTTPS|
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Write|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
| IOPs Total | Implemented | HTTPS |
|Latency Write|Implemented|HTTPS|
.6+|StoragePool Disk|IOps Read|Implemented|HTTPS|Number of read IOps on
the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
.9+|Volume|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Latency Read|Implemented|HTTPS|
|Latency Total|Implemented|HTTPS|
|Latency Write|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Nutanix REST API
```

```
| HTTPS
IHTTPS
1443
Itrue
|true
|true
|true
|===
<<top, Back to Top>>
== OpenStack (REST API / SSH)
:description: Support Matrix Asciidoc for OpenStack (REST API / SSH)
|===
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.76+|foundation .7+|Data Store|Capacity|Implemented|HTTPS|Snapshot Used
capacity in MB
|MOID|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
| Provisioned Capacity | Implemented | HTTPS |
|Used Capacity|Implemented|HTTPS|
|Virtual Center Ip|Implemented|HTTPS|
.3+|NasShare DataStore|DataStore OID|Implemented|HTTPS|
|Nas Share OID|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.4+|NasShare Host|Nas Share OID|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Read Only|Implemented|HTTPS|
.6+|LUN|Disk Name|Implemented|HTTPS|
|DataStore OID|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
|Number|Implemented|HTTPS|
```

```
|OID|Implemented|HTTPS|
|TID|Implemented|HTTPS|
.6+|NAS Share|Capacity|Implemented|HTTPS|Allocated capacity in MB
|Filer Ip|Implemented|HTTPS|
|Filer Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Share Path|Implemented|HTTPS| For a HvNasShare to be matched to a Share
|Type|Gap|HTTPS|
.6+|Path|Active|Implemented|HTTPS|
|Lun OID|Implemented|HTTPS|
|Host Port WWPN|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Storage Port WWPN|Implemented|HTTPS|
|Type|Gap|HTTPS|
.6+|Server|Cluster|Implemented|HTTPS|Cluster name
|DataCenter Name|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
|MOID|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Virtual Center Ip|Implemented|HTTPS|
.7+|Virtual Disk|Capacity|Implemented|HTTPS|Snapshot Used capacity in MB
|DataStore OID|Implemented|HTTPS|
|Lun OID|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
|Used Capacity|Implemented|HTTPS|
.13+|VirtualMachine|Guest State|Implemented|HTTPS|
|DataStore OID|Implemented|HTTPS|
|Host OID|Implemented|HTTPS|
|IPs|Implemented|HTTPS|
|MOID|Implemented|HTTPS|
|Memory|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Power State|Implemented|HTTPS|
|State Change Time|Implemented|HTTPS|
|Processors|Implemented|HTTPS|
|Provisioned Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
.3+|VirtualMachine Disk|OID|Implemented|HTTPS|
|VirtualDisk OID|Implemented|HTTPS|
|VirtualMachine OID|Implemented|HTTPS|
.7+|Host|Host Cpu Count|Implemented|HTTPS|
|Host Domain|Implemented|HTTPS|
|Host Installed Memory|Implemented|HTTPS|
```

```
|Host OS|Implemented|HTTPS|
|IPs|Implemented|HTTPS|
|Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
.4+|ISCSI Node|Host Aliases|Implemented|HTTPS|
|Node Name|Implemented|HTTPS|
|OID|Implemented|HTTPS|
|Type|Gap|HTTPS|
.4+|Info|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.10+|performance .5+|Data Store|Total Capacity|Implemented||
|Capacity Used Ratio|Implemented||
|Capacity Provisioned|Implemented||
|Used Capacity|Implemented||
|Over Commit Capacity Ratio|Implemented||Reported as a time series
.2+|Host|Total CPU Utilization|Implemented||
|Total Memory Utilization|Implemented||
.3+|Virtual Disk|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|OpenStack REST API
|HTTPS
|HTTPS
1443
|true
|false
ltrue
|true
|OpenStack SSH
|SSH
|SSH
122
```

```
|true
|false
|true
|true
|===
<<top, Back to Top>>
== Oracle ZFS (HTTPS)
:description: Support Matrix Asciidoc for Oracle ZFS (HTTPS)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|Sun ZFS Storage 7330
Sun ZFS Storage 7335
Sun ZFS Storage 7350
Sun ZFS Storage 7370
Sun ZFS Storage 7420
Sun ZFS Storage 7430
Sun ZFS Storage 7450
1-1.1
1-1.2
1-1.3
1-1.34
1-1.4
2013.06.05.6.12
2013.06.05.6.15
2013.06.05.7.21
2013.06.05.7.24
2013.06.05.7.25
2013.06.05.7.26
2013.06.05.8.0
2013.06.05.8.26
2013.06.05.8.29
2013.06.05.8.35
2013.06.05.8.37
2013.06.05.8.47
2013.06.05.8.50
2013.06.05.8.53
2013.06.05.8.6
2013.06.05.8.7
```

```
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.114+|foundation .14+|Disk|Capacity (GB)|Implemented|HTTP/S|use capacity
|Disk Id|Implemented|HTTP/S|Uniquely identifies this disk in the array
|Location|Gap|HTTP/S|Where this disk is physically located in the array
|Model|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Role|Implemented|HTTP/S|
|Role Enum|Implemented|HTTP/S|enum for disk role
|Serial Number|Implemented|HTTP/S|
|Speed|Implemented|HTTP/S|Speed of disk (RPM)
|Status|Implemented|HTTP/S|
|Status Enum|Implemented|HTTP/S|enum for disk status
|Type|Gap|HTTP/S|
|Type Enum|Implemented|HTTP/S|enum for disk type
|Vendor|Implemented|HTTP/S|
.4+|Info|DataSource Name|Implemented|HTTP/S|Info
|Date|Implemented|HTTP/S|
|Originator ID|Implemented|HTTP/S|
|Originator Key|Implemented|HTTP/S|
.18+|Internal Volume|Compression Enabled|Implemented|HTTP/S|Is compression
enabled on the storage pool
|Compression Savings|Implemented|HTTP/S|ratio of compression savings in
percentage
|Data Used Capacity|Implemented|HTTP/S|
|Dedupe Enabled|Implemented|HTTP/S|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTP/S|ratio of dedupe savings in percentage
|Internal Volume Id|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Other UsedCapacity (MB)|Implemented|HTTP/S|Any capacity other than data
and snapshot
|Protection Type|Implemented|HTTP/S|
|Raw to Usable Ratio|Implemented|HTTP/S|ratio to convert from usable
capacity to raw capacity
|Snapshot Used Capacity|Implemented|HTTP/S|
|Storage Pool Id|Implemented|HTTP/S|
|Thin Provisioned|Implemented|HTTP/S|
|Thin Provisioning Supported|Implemented|HTTP/S|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTP/S|
|Total Used Capacity|Implemented|HTTP/S|Total capacity in MB
|Total Used Capacity (MB)|Implemented|HTTP/S|place holder for the used
```

```
capacity as read from the device
|Type|Gap|HTTP/S|
.4+|QTree|Name|Implemented|HTTP/S|
|Qtree Id|Implemented|HTTP/S|unique id of the qtree
|Quota UsedCapacity|Implemented|HTTP/S|Space in MB currently used
|Type|Gap|HTTP/S|
.7+|Quota|Internal Volume Id|Implemented|HTTP/S|
|Qtree Id|Implemented|HTTP/S|unique id of the qtree
|Quota Id|Implemented|HTTP/S|unique id of the quota
|Soft Capacity Limit (MB)|Implemented|HTTP/S|Maximum amount of disk space,
allowed for the quota target
|Type|Gap|HTTP/S|
|Used Capacity|Implemented|HTTP/S|
|User/Group Target|Implemented|HTTP/S|user/group target this quota refers
.13+|Storage|Display IP|Implemented|HTTP/S|
|Failed Raw Capacity|Implemented|HTTP/S|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTP/S|The storage Family could be Clariion,
Symmetrix, et al
|IP|Implemented|HTTP/S|
|Manufacturer|Implemented|HTTP/S|
|Microcode Version|Implemented|HTTP/S|
|Model|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Total Raw Capacity|Implemented|HTTP/S|Total raw capacity (sum of all
disks on the array)
|Serial Number|Implemented|HTTP/S|
|Spare Raw Capacity|Implemented|HTTP/S|Raw capapcity of spare disks (sum
of all disks that are spare)
|SupportActive Active|Implemented|HTTP/S|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTP/S|Is this a storage virtualization device?
.8+|Storage Node|Memory Size|Gap|HTTP/S|device memory in MB
|Model|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Processors Count|Implemented|HTTP/S|device CPU
|Serial Number|Implemented|HTTP/S|
|State|Implemented|HTTP/S|free text describing the device state
|UUID|Implemented|HTTP/S|
|Version|Implemented|HTTP/S|software version
.22+|Storage Pool|Auto Tiering|Implemented|HTTP/S|indicates if this
storagepool is participating in auto tiering with other pools
|Compression Enabled|Implemented|HTTP/S|Is compression enabled on the
storage pool
|Compression Savings|Implemented|HTTP/S|ratio of compression savings in
```

```
percentage
|Data Allocated Capacity|Gap|HTTP/S|capacity allocated for data
|Data Used Capacity|Implemented|HTTP/S|
|Dedupe Enabled|Implemented|HTTP/S|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTP/S|ratio of dedupe savings in percentage
| Include In Dwh Capacity | Implemented | HTTP/S | A way from ACQ to cotnrol
which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTP/S|
|Other UsedCapacity (MB)|Implemented|HTTP/S|Any capacity other than data
and snapshot
|Physical Disk Capacity (MB)|Implemented|HTTP/S|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTP/S|indicates whether this storagePool is a
raid group
|Raw to Usable Ratio|Implemented|HTTP/S|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTP/S|Redundancy level
|Snapshot Used Capacity|Implemented|HTTP/S|
|Status|Implemented|HTTP/S|
|Storage Pool Id|Implemented|HTTP/S|
|Thin Provisioning Supported|Implemented|HTTP/S|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTP/S|
|Total Used Capacity|Implemented|HTTP/S|Total capacity in MB
|Type|Gap|HTTP/S|
|Virtual|Implemented|HTTP/S|Is this a storage virtualization device?
.15+|Volume|Name|Implemented|HTTP/S|
|Junction Path|Implemented|HTTP/S|
|Storage Pool Id|Implemented|HTTP/S|
|Auto Tiering|Implemented|HTTP/S|indicates if this storagepool is
participating in auto tiering with other pools
|UUID|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
|Disk Type|Not Available|HTTP/S|
|Virtual|Implemented|HTTP/S|Is this a storage virtualization device?
|Thin Provisioned|Implemented|HTTP/S|
|Disk Speed|Implemented|HTTP/S|comma seperated list of disk speeds (rpm)
|Disk Size|Implemented|HTTP/S|comma seperated list of disk sizes (GB)
|Capacity|Implemented|HTTP/S|Snapshot Used capacity in MB
|Total Raw Capacity|Implemented|HTTP/S|Total raw capacity (sum of all
disks on the array)
|Used Capacity|Implemented|HTTP/S|
|Redundancy|Implemented|HTTP/S|Redundancy level
.5+|Volume Map|LUN|Implemented|HTTP/S|Name of the backend lun
|Storage Port|Implemented|HTTP/S|
|Masking Required|Implemented|HTTP/S|
```

```
|Protocol Controller|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.4+|Volume Mask|Storage Port|Implemented|HTTP/S|
|Initiator|Implemented|HTTP/S|
|Protocol Controller|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.41+|performance .1+|Internal Volume|IOPs Total|Implemented||
.2+|Qtree|Used Capacity|Implemented||
|Total File Count|Implemented||
.8+|Storage|Cache Hit Ratio Total|Implemented||
|IOPs other|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
| IOPs Total | Implemented | |
|IOPs Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
.8+|Storage Node|Cache Hit Ratio Total|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
|IOPs Total|Implemented||
| IOPs Write | Implemented | |
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|Utilization Total|Implemented||
.10+|Storage Node Data|key|Implemented||
|Server ID|Implemented||
|Throughput Read|Implemented||
|Throughput Write|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|IOps Read|Implemented||Number of read IOps on the disk
|IOPs Write|Implemented||
| IOPs Total | Implemented | |
|Cache Hit Ratio Total|Implemented||
|Utilization Total|Implemented||
.12+|StoragePool Disk|IOPs Total|Implemented||
|Total Capacity|Implemented||
|Capacity Used Ratio|Implemented||
|Total Data Capacity|Implemented||
|Capacity Provisioned|Implemented||
|Data Used Capacity|Implemented||
|Used Capacity|Implemented||
|Other Used Capacity|Implemented||
```

```
|Raw Capacity|Implemented||
|Over Commit Capacity Ratio|Implemented||Reported as a time series
|Snapshot Used Capacity|Implemented||
|Snapshot Used Capacity Ratio|Implemented|| Reported as a time series
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Oracle ZFS REST API
|HTTP/HTTPS
|HTTP/HTTPS
1215
|true
|true
|true
|true
|===
<<top, Back to Top>>
== Pure Storage FlashArray (HTTP)
:description: Support Matrix Asciidoc for Pure Storage FlashArray (HTTP)
Models and versions supported by this data collector:
<.<|Models <.<|Firmware versions
|DFSC1
FA-420
FA-450
FA-C40R3
FA-C60
FA-C60R3
FA-X10R2
FA-X10R3
FA-X20R2
FA-X20R3
FA-X50R2
```

FA-X50R3 FA-X70R2 FA-X70R3 FA-X90R2 FA-X90R3 FA-XL130 FA-XL170 FA-m10r2 FA-m20 FA-m20r2 FA-m50 FA-m50r2 FA-m70 FA-m70r2 FA-x70 |4.8.8 5.3.14 5.3.15 5.3.17 5.3.18 5.3.20 5.3.21 5.3.6 5.3.8 6.1.10 6.1.11 6.1.13 6.1.14 6.1.15 6.1.17 6.1.18 6.1.19 6.1.21 6.1.22 6.1.23 6.1.5 6.2.13 6.2.7 6.2.9 6.3.10 6.3.11 6.3.12 6.3.2 6.3.5 6.3.6 6.3.7

```
6.3.9
6.4.3
6.4.4
6.4.5
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.98+|foundation .13+|Disk|Capacity (GB)|Implemented|HTTP/S|use capacity
|Disk Id|Implemented|HTTP/S|Uniquely identifies this disk in the array
|Name|Implemented|HTTP/S|
|Potential Transfer Rate|Implemented|HTTP/S|
|Role|Implemented|HTTP/S|
|Role Enum|Implemented|HTTP/S|enum for disk role
|Seek Time|Implemented|HTTP/S|
|Speed|Implemented|HTTP/S|Speed of disk (RPM)
|Status|Implemented|HTTP/S|
|Status Enum|Implemented|HTTP/S|enum for disk status
|Type|Gap|HTTP/S|
|Type Enum|Implemented|HTTP/S|enum for disk type
|Vendor|Implemented|HTTP/S|
.4+|ISCSI Network Portal|IP|Implemented|HTTP/S|
|Listening Port|Implemented|HTTP/S|
|Nic|Implemented|HTTP/S|
|OID|Implemented|HTTP/S|
.3+|ISCSI Network Portal Group|OID|Implemented|HTTP/S|
|Portal Group Name|Implemented|HTTP/S|
|Portal Group Tag|Implemented|HTTP/S|
.3+|ISCSI Node|Node Name|Implemented|HTTP/S|
|OID|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.2+|ISCSI Node Map|OID|Implemented|HTTP/S|
|Portal Group OID|Implemented|HTTP/S|
.7+|ISCSI Session|Initiator Ips|Implemented|HTTP/S|
|Initiator OID|Implemented|HTTP/S|
|Max Connections|Implemented|HTTP/S|
|Number Of Connections|Implemented|HTTP/S|
|OID|Implemented|HTTP/S|
|Portal Group OID|Implemented|HTTP/S|
|Security|Implemented|HTTP/S|
.4+|Info|DataSource Name|Implemented|HTTP/S|Info
|Date|Implemented|HTTP/S|
|Originator ID|Implemented|HTTP/S|
```

```
|Originator Key|Implemented|HTTP/S|
.14+|Storage|Display IP|Implemented|HTTP/S|
|Failed Raw Capacity|Implemented|HTTP/S|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTP/S|The storage Family could be Clariion,
Symmetrix, et al
|IP|Implemented|HTTP/S|
|Manage URL|Implemented|HTTP/S|
|Manufacturer|Implemented|HTTP/S|
|Microcode Version|Implemented|HTTP/S|
|Model|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Total Raw Capacity|Implemented|HTTP/S|Total raw capacity (sum of all
disks on the array)
|Serial Number|Implemented|HTTP/S|
|Spare Raw Capacity|Implemented|HTTP/S|Raw capapcity of spare disks (sum
of all disks that are spare)
|SupportActive Active|Implemented|HTTP/S|Specified if the storage supports
active-active configurations
|Virtual|Implemented|HTTP/S|Is this a storage virtualization device?
.6+|Storage Node|Memory Size|Gap|HTTP/S|device memory in MB
|Model|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|State|Implemented|HTTP/S|free text describing the device state
|UUID|Implemented|HTTP/S|
|Version|Implemented|HTTP/S|software version
.18+|Storage Pool|Data Allocated Capacity|Gap|HTTP/S|capacity allocated
for data
|Data Used Capacity|Implemented|HTTP/S|
|Dedupe Enabled|Implemented|HTTP/S|Is dedupe enabled on the storage pool
|Dedupe Savings|Implemented|HTTP/S|ratio of dedupe savings in percentage
|Include In Dwh Capacity|Implemented|HTTP/S|A way from ACQ to cotnrol
which stroage pools are interesting in DWH Capacity
|Name|Implemented|HTTP/S|
|Other UsedCapacity (MB)|Implemented|HTTP/S|Any capacity other than data
and snapshot
|Physical Disk Capacity (MB)|Implemented|HTTP/S|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTP/S|indicates whether this storagePool is a
raid group
|Raw to Usable Ratio|Implemented|HTTP/S|ratio to convert from usable
capacity to raw capacity
|Redundancy|Implemented|HTTP/S|Redundancy level
|Snapshot Used Capacity|Implemented|HTTP/S|
|Storage Pool Id|Implemented|HTTP/S|
|Thin Provisioning Supported|Implemented|HTTP/S|Whether this internal
```

```
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTP/S|
|Total Used Capacity|Implemented|HTTP/S|Total capacity in MB
|Type|Gap|HTTP/S|
|Virtual|Implemented|HTTP/S|Is this a storage virtualization device?
.12+|Volume|Capacity|Implemented|HTTP/S|Snapshot Used capacity in MB
|Disk Size|Implemented|HTTP/S|comma seperated list of disk sizes (GB)
|Disk Speed|Implemented|HTTP/S|comma seperated list of disk speeds (rpm)
|Disk Type|Not Available|HTTP/S|
|Name|Implemented|HTTP/S|
|Total Raw Capacity|Implemented|HTTP/S|Total raw capacity (sum of all
disks on the array)
|Redundancy|Implemented|HTTP/S|Redundancy level
|Storage Pool Id|Implemented|HTTP/S|
|Thin Provisioned|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
|Used Capacity|Implemented|HTTP/S|
|Virtual|Implemented|HTTP/S|Is this a storage virtualization device?
.4+|Volume Map|LUN|Implemented|HTTP/S|Name of the backend lun
|Protocol Controller|Implemented|HTTP/S|
|Storage Port|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.4+|Volume Mask|Initiator|Implemented|HTTP/S|
|Protocol Controller|Implemented|HTTP/S|
|Storage Port|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.4+|WWN Alias|Host Aliases|Implemented|HTTP/S|
|Object Type|Implemented|HTTP/S|
|Source|Implemented|HTTP/S|
|WWN|Implemented|HTTP/S|
.38+|performance .14+|Storage|Failed Raw Capacity|Implemented||
|Raw Capacity|Implemented||
|Spare Raw Capacity|Implemented||Raw capapcity of spare disks (sum of all
disks that are spare)
|StoragePools Capacity|Implemented||
|IOPs other|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
```

```
.11+|StoragePool Disk|Capacity Provisioned|Implemented||
|Raw Capacity|Implemented||
|Total Capacity|Implemented||
|Used Capacity|Implemented||
|Over Commit Capacity Ratio|Implemented||Reported as a time series
|Capacity Used Ratio|Implemented||
|Total Data Capacity|Implemented||
|Data Used Capacity|Implemented||
|Other Used Capacity|Implemented||
|Snapshot Used Capacity|Implemented||
|Snapshot Used Capacity Ratio|Implemented|| Reported as a time series
.13+|Volume|Raw Capacity|Implemented||
|Total Capacity|Implemented||
|Used Capacity|Implemented||
|Capacity Used Ratio|Implemented||
| IOps Read | Implemented | Number of read IOps on the disk
|IOPs Total|Implemented||
|IOPs Write|Implemented||
|Latency Read|Implemented||
|Latency Total|Implemented||
|Latency Write|Implemented||
|Throughput Read|Implemented||
|Throughput Total|Implemented||Average disk total rate (read and write
across all disks) in MB/s
|Throughput Write|Implemented||
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
| Pure Storage REST API
|HTTP/HTTPS
|HTTP/HTTPS
180/443
|true
|true
|true
|true
|===
```

```
<<top, Back to Top>>
== Red Hat RHV (REST)
:description: Support Matrix Asciidoc for Red Hat RHV (REST)
|===
|===
Products supported by this data collector:
|===
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.74+|foundation .7+|Data Store|OID|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Capacity|Implemented|HTTP/S|Snapshot Used capacity in MB
|Used Capacity|Implemented|HTTP/S|
|Provisioned Capacity|Implemented|HTTP/S|
|Virtual Center Ip|Implemented|HTTP/S|
|MOID|Implemented|HTTP/S|
.3+|NasShare DataStore|DataStore OID|Implemented|HTTP/S|
|Nas Share OID|Implemented|HTTP/S|
|OID|Implemented|HTTP/S|
.4+|NasShare Host|Host OID|Implemented|HTTP/S|
|Nas Share OID|Implemented|HTTP/S|
|Read Only|Implemented|HTTP/S|
|OID|Implemented|HTTP/S|
.4+|LUN|OID|Implemented|HTTP/S|
|Number|Implemented|HTTP/S|
|DataStore OID|Implemented|HTTP/S|
|Host OID|Implemented|HTTP/S|
.6+|NAS Share|Filer Name|Implemented|HTTP/S|
|Filer Ip|Implemented|HTTP/S|
|Share Path|Implemented|HTTP/S| For a HvNasShare to be matched to a Share
|Capacity|Implemented|HTTP/S|Allocated capacity in MB
|Type|Gap|HTTP/S|
|OID|Implemented|HTTP/S|
.4+|Path|OID|Implemented|HTTP/S|
|Lun OID|Implemented|HTTP/S|
|Active|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.6+|Server|OID|Implemented|HTTP/S|
```

```
|Virtual Center Ip|Implemented|HTTP/S|
|Cluster|Implemented|HTTP/S|Cluster name
|DataCenter Name|Implemented|HTTP/S|
|Host OID|Implemented|HTTP/S|
|MOID|Implemented|HTTP/S|
.6+|Virtual Disk|OID|Implemented|HTTP/S|
|DataStore OID|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|Capacity|Implemented|HTTP/S|Snapshot Used capacity in MB
|Used Capacity|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.14+|VirtualMachine|OID|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|OS|Implemented|HTTP/S|
|Processors|Implemented|HTTP/S|
|Memory|Implemented|HTTP/S|
|DataStore OID|Implemented|HTTP/S|
|Power State|Implemented|HTTP/S|
|State Change Time|Implemented|HTTP/S|
|Host OID|Implemented|HTTP/S|
|IPs|Implemented|HTTP/S|
|Guest State|Implemented|HTTP/S|
|Used Capacity|Implemented|HTTP/S|
|Provisioned Capacity|Implemented|HTTP/S|
|MOID|Implemented|HTTP/S|
.3+|VirtualMachine Disk|OID|Implemented|HTTP/S|
|VirtualMachine OID|Implemented|HTTP/S|
|VirtualDisk OID|Implemented|HTTP/S|
.11+|Host|OID|Implemented|HTTP/S|
|Name|Implemented|HTTP/S|
|IPs|Implemented|HTTP/S|
|Platform Type|Implemented|HTTP/S|
| Host Installed Memory | Implemented | HTTP / S |
|Manufacturer|Implemented|HTTP/S|
|Host Model|Implemented|HTTP/S|
|Host Cpu Count|Implemented|HTTP/S|
|Host Cpu Speed|Implemented|HTTP/S|
|NIC count|Implemented|HTTP/S|
|NIC speed|Implemented|HTTP/S|
.3+|ISCSI Node|OID|Implemented|HTTP/S|
|Node Name|Implemented|HTTP/S|
|Type|Gap|HTTP/S|
.3+|Info|DataSource Name|Implemented|HTTP/S|Info
|Originator ID|Implemented|HTTP/S|
|Date|Implemented|HTTP/S|
```

```
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Red Hat RHEV REST API
|HTTPS
| HTTPS
1443
|true
Itrue
|true
|true
|===
<<top, Back to Top>>
== Rubrik Storage
:description: Support Matrix Asciidoc for Rubrik Storage
Models and versions supported by this data collector:
|===
<.<|API versions <.<|Firmware versions
|v5.3
|5.3.3-p1-19391
6.0.3-p3-13584
7.0.2-p4-15876
7.0.3-p1-15949
8.0.3-p2-22743
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.57+|foundation .12+|Disk|Capacity (GB)|Implemented|HTTPS|use capacity
|Disk Id|Implemented|HTTPS|Uniquely identifies this disk in the array
|Location|Gap|HTTPS|Where this disk is physically located in the array
```

```
|Name|Implemented|HTTPS|
|Role|Implemented|HTTPS|
|Role Enum|Implemented|HTTPS|enum for disk role
|Speed|Implemented|HTTPS|Speed of disk (RPM)
|Status|Implemented|HTTPS|
|Status Enum|Implemented|HTTPS|enum for disk status
|Type|Gap|HTTPS|
|Type Enum|Implemented|HTTPS|enum for disk type
|Encrypted|Implemented|HTTPS|
.7+|Info|Api Description|Implemented|HTTPS|
|Api Name|Implemented|HTTPS|
|Api Version|Implemented|HTTPS|
|DataSource Name|Implemented|HTTPS|Info
|Date|Implemented|HTTPS|
|Originator ID|Implemented|HTTPS|
|Originator Key|Implemented|HTTPS|
.13+|Storage|Cluster|Not Available|HTTPS|Whether this storage is a cluster
|Cpu Count|Implemented|HTTPS|Cpu Count of the storage
|Display IP|Implemented|HTTPS|
|Failed Raw Capacity|Implemented|HTTPS|Raw capapcity of failed disks (sum
of all disks that are failed)
|Family|Implemented|HTTPS|The storage Family could be Clariion, Symmetrix,
et al
|IP|Implemented|HTTPS|
|Manage URL|Implemented|HTTPS|
|Manufacturer|Implemented|HTTPS|
|Memory|Implemented|HTTPS|
|Microcode Version|Implemented|HTTPS|
|Total Raw Capacity|Implemented|HTTPS|Total raw capacity (sum of all disks
on the array)
|Spare Raw Capacity|Implemented|HTTPS|Raw capapcity of spare disks (sum of
all disks that are spare)
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
.10+|Storage Node|Memory Size|Gap|HTTPS|device memory in MB
|Name|Implemented|HTTPS|
|Processors Count|Implemented|HTTPS|device CPU
|State|Implemented|HTTPS|free text describing the device state
|UUID|Implemented|HTTPS|
|ManagementIp Addresses|Implemented|HTTPS|
|Node Capacity Utilization Total in MB|Implemented|HTTPS|
|Node Capacity Utilization Usable in MB|Implemented|HTTPS|
|Node Capacity Utilization Used in MB|Implemented|HTTPS|
|Site Name|Implemented|HTTPS|
.15+|Storage Pool|Data Used Capacity|Implemented|HTTPS|
|Include In Dwh Capacity|Implemented|HTTPS|A way from ACQ to cotnrol which
stroage pools are interesting in DWH Capacity
```

```
|Name|Implemented|HTTPS|
|Other UsedCapacity (MB)|Implemented|HTTPS|Any capacity other than data
and snapshot
|Physical Disk Capacity (MB)|Implemented|HTTPS|used as raw capacity for
storage pool
|Raid Group|Implemented|HTTPS|indicates whether this storagePool is a raid
group
|Raw to Usable Ratio|Implemented|HTTPS|ratio to convert from usable
capacity to raw capacity
|Snapshot Used Capacity|Implemented|HTTPS|
|Storage Pool Id|Implemented|HTTPS|
|Thin Provisioning Supported|Implemented|HTTPS|Whether this internal
volume supports thin provisioning for the volume layer on top of it
|Total Allocated Capacity|Implemented|HTTPS|
|Total Used Capacity|Implemented|HTTPS|Total capacity in MB
|Type|Gap|HTTPS|
|Virtual|Implemented|HTTPS|Is this a storage virtualization device?
|Effective Used Capacity Percent|Implemented|HTTPS|
.30+|performance .7+|Storage|IOps Read|Implemented|HTTPS|Number of read
IOps on the disk
| IOPs other | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
|Throughput Read|Implemented|HTTPS|
|Throughput Write|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
| IOPs Total | Implemented | HTTPS |
.10+|Storage Node|IOps Read|Implemented|HTTPS|Number of read IOps on the
disk
| IOPs Total | Implemented | HTTPS |
| IOPs Write | Implemented | HTTPS |
| Node Capacity Utilization Total | Implemented | HTTPS |
|Node Capacity Utilization Usable|Implemented|HTTPS|
|Node Capacity Utilization Used|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|Utilization Total|Implemented|HTTPS|
.13+|StoragePool Disk|Raw Capacity|Implemented|HTTPS|
|Total Capacity|Implemented|HTTPS|
|Used Capacity|Implemented|HTTPS|
|Capacity Used Ratio|Implemented|HTTPS|
|Data Used Capacity|Implemented|HTTPS|
|IOps Read|Implemented|HTTPS|Number of read IOps on the disk
| IOPs Total | Implemented | HTTPS |
```

```
| IOPs Write | Implemented | HTTPS |
|Other Used Capacity|Implemented|HTTPS|
|Snapshot Used Capacity|Implemented|HTTPS|
|Throughput Read|Implemented|HTTPS|
|Throughput Total|Implemented|HTTPS|Average disk total rate (read and
write across all disks) in MB/s
|Throughput Write|Implemented|HTTPS|
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|Rubrik Storage REST API
|HTTPS
|HTTPS
1443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== NetApp HCI Virtual Center
:description: Support Matrix Asciidoc for NetApp HCI Virtual Center
Models and versions supported by this data collector:
|===
<.<|API versions
|VMware vCenter Server 6.7.0 build-10244857
VMware vCenter Server 6.7.0 build-14368073
VMware vCenter Server 7.0.3 build-19234570
VMware vCenter Server 7.0.3 build-20150588
VMware vCenter Server 7.0.3 build-20395099
VMware vCenter Server 7.0.3 build-20990077
VMware vCenter Server 7.0.3 build-21477706
VMware vCenter Server 7.0.3 build-21784236
```

```
VMware vCenter Server 8.0.1 build-21815093
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.91+|foundation .8+|Data Store|Capacity|Implemented|Web Services|Snapshot
Used capacity in MB
|MOID|Implemented|Web Services|
|Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|Provisioned Capacity|Implemented|Web Services|
|Used Capacity|Implemented|Web Services|
|Virtual Center Ip|Implemented|Web Services|
|Type|Implemented|Web Services|
.3+|NasShare DataStore|DataStore OID|Implemented|Web Services|
|Nas Share OID|Implemented|Web Services|
|OID|Implemented|Web Services|
.4+|NasShare Host|Nas Share OID|Implemented|Web Services|
|Host OID|Implemented|Web Services|
|OID|Implemented|Web Services|
|Read Only|Implemented|Web Services|
.8+|LUN|Disk Name|Implemented|Web Services|
|DataStore OID|Implemented|Web Services|
| Host OID | Implemented | Web Services |
|Number|Implemented|Web Services|
|OID|Implemented|Web Services|
|Policy|Implemented|Web Services|
|TID|Implemented|Web Services|
|Volume Uuid|Implemented|Web Services|
.6+|NAS Share|Capacity|Implemented|Web Services|Allocated capacity in MB
|Filer Ip|Implemented|Web Services|
|Filer Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|Share Path|Implemented|Web Services| For a HvNasShare to be matched to a
Share
|Type|Gap|Web Services|
.6+|Path|Active|Implemented|Web Services|
|Lun OID|Implemented|Web Services|
| Host Port WWPN | Implemented | Web Services |
|OID|Implemented|Web Services|
|Storage Port WWPN|Implemented|Web Services|
|Type|Gap|Web Services|
.6+|Server|Cluster|Implemented|Web Services|Cluster name
```

```
|DataCenter Name|Implemented|Web Services|
| Host OID | Implemented | Web Services |
|MOID|Implemented|Web Services|
|OID|Implemented|Web Services|
|Virtual Center Ip|Implemented|Web Services|
.8+|Virtual Disk|Capacity|Implemented|Web Services|Snapshot Used capacity
in MB
|DataStore OID|Implemented|Web Services|
|Lun OID|Implemented|Web Services|
|Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|Type|Gap|Web Services|
|Used Capacity|Implemented|Web Services|used capacity for reporting (MB)
|Used Capacity|Implemented|Web Services|
.15+|VirtualMachine|Dns Name|Implemented|Web Services|
|Guest State|Implemented|Web Services|
|DataStore OID|Implemented|Web Services|
|Host OID|Implemented|Web Services|
|IPs|Implemented|Web Services|
|MOID|Implemented|Web Services|
|Memory|Implemented|Web Services|
|Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|OS|Implemented|Web Services|
|Power State|Implemented|Web Services|
|State Change Time|Implemented|Web Services|
|Processors|Implemented|Web Services|
|Provisioned Capacity|Implemented|Web Services|
|Used Capacity|Implemented|Web Services|
.3+|VirtualMachine Disk|OID|Implemented|Web Services|
|VirtualDisk OID|Implemented|Web Services|
|VirtualMachine OID|Implemented|Web Services|
.12+|Host|Host Cpu Count|Implemented|Web Services|
|Host Cpu Speed|Implemented|Web Services|
|Host Domain|Implemented|Web Services|
| Host Installed Memory | Implemented | Web Services |
|Host Model|Implemented|Web Services|
|NIC count|Implemented|Web Services|
|NIC speed|Implemented|Web Services|
|IPs|Implemented|Web Services|
|Manufacturer|Implemented|Web Services|
|Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|Platform Type|Implemented|Web Services|
.4+|ISCSI Node|Host Aliases|Implemented|Web Services|
|Node Name|Implemented|Web Services|
```

```
|OID|Implemented|Web Services|
|Type|Gap|Web Services|
.8+|Info|Api Description|Implemented|Web Services|
|Api Name|Implemented|Web Services|
|Api Version|Implemented|Web Services|
|Client Api Name|Implemented|Web Services|
|Client Api Version|Implemented|Web Services|
|DataSource Name|Implemented|Web Services|Info
|Date|Implemented|Web Services|
|Originator ID|Implemented|Web Services|
.53+|performance .14+|Data Store|Capacity Provisioned|Implemented|Web
Services|
|Total Capacity|Implemented|Web Services|
|Used Capacity|Implemented|Web Services|
|Over Commit Capacity Ratio | Implemented | Web Services | Reported as a time
series
|Capacity Used Ratio|Implemented|Web Services|
|IOps Read|Implemented|Web Services|Number of read IOps on the disk
| IOPs Total | Implemented | Web Services |
| IOPs Write | Implemented | Web Services |
|Latency Read|Implemented|Web Services|
|Latency Total|Implemented|Web Services|
|Latency Write|Implemented|Web Services|
|Throughput Read|Implemented|Web Services|
|Throughput Total|Implemented|Web Services|Average disk total rate (read
and write across all disks) in MB/s
|Throughput Write|Implemented|Web Services|
.17+|Host|Total CPU Utilization|Implemented|Web Services|
|IOps Read|Implemented|Web Services|Number of read IOps on the disk
|diskIops.total|Implemented|Web Services|
|Disk IOPs write|Implemented|Web Services|
|Latency Read|Implemented|Web Services|
|Latency Total|Implemented|Web Services|
|Latency Write|Implemented|Web Services|
|Disk Throughput Read|Implemented|Web Services|
|Throughput Read|Implemented|Web Services|total disk throughput read
|Disk Throughput Write|Implemented|Web Services|
| IP Throughput Read | Implemented | Web Services |
|Throughput total|Implemented|Web Services|IP throughput total
|ipThroughput.write|Implemented|Web Services|
|Total Memory Utilization|Implemented|Web Services|
|swapRate.inRate|Implemented|Web Services|
|Swap Rate|Implemented|Web Services|
|Total Swap Rate|Implemented|Web Services|
.22+|vm|cpuCoSchedulingDelayTimePercent.total|Implemented|Web Services|
|cpuDemandToEntitlementPercent.total|Implemented|Web Services|
```

```
|Idle CPU Time|Implemented|Web Services|idle time in percent
|CPU Wait Time|Implemented|Web Services|total cpu wait time in percent
|Total CPU Utilization|Implemented|Web Services|
| IOps Read | Implemented | Web Services | Number of read IOps on the disk
|diskIops.total|Implemented|Web Services|
|Disk IOPs write|Implemented|Web Services|
|Latency Read|Implemented|Web Services|
|Latency Total|Implemented|Web Services|
|Latency Write|Implemented|Web Services|
|Disk Throughput Read|Implemented|Web Services|
|Throughput Read|Implemented|Web Services|total disk throughput read
|Disk Throughput Write|Implemented|Web Services|
| IP Throughput Read | Implemented | Web Services |
|Throughput total|Implemented|Web Services|IP throughput total
|ipThroughput.write|Implemented|Web Services|
|Total Memory Utilization|Implemented|Web Services|
|swapRate.inRate|Implemented|Web Services|
|Swap Rate|Implemented|Web Services|
|Total Swap Rate|Implemented|Web Services|
|Schedule wait time|Implemented|Web Services|Waiting to be scheduled time
in percent
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|VMware REST API
|Web Services
|HTTP/HTTPS
180/443
|true
|true
|true
Itrue
|===
<<top, Back to Top>>
== VMware Cloud on AWS
```

```
:description: Support Matrix Asciidoc for VMware Cloud on AWS
Models and versions supported by this data collector:
|===
<.<|API versions
|VMware vCenter Server 7.0.3 build-20532039
VMware vCenter Server 7.0.3 build-20870699
VMware vCenter Server 8.0.0 build-21709157
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.59+|foundation .8+|Data Store|OID|Implemented|Web Services|
|Name|Implemented|Web Services|
|Capacity|Implemented|Web Services|Snapshot Used capacity in MB
| Used Capacity | Implemented | Web Services |
|Provisioned Capacity|Implemented|Web Services|
|Virtual Center Ip|Implemented|Web Services|
|MOID|Implemented|Web Services|
|Type|Implemented|Web Services|
.6+|Server|OID|Implemented|Web Services|
|Virtual Center Ip|Implemented|Web Services|
|Cluster|Implemented|Web Services|Cluster name
|DataCenter Name|Implemented|Web Services|
|Host OID|Implemented|Web Services|
|MOID|Implemented|Web Services|
.7+|Virtual Disk|OID|Implemented|Web Services|
|DataStore OID|Implemented|Web Services|
|Name|Implemented|Web Services|
|Capacity|Implemented|Web Services|Snapshot Used capacity in MB
|Used Capacity|Implemented|Web Services|
|Used Capacity|Implemented|Web Services|used capacity for reporting (MB)
|Type|Gap|Web Services|
.15+|VirtualMachine|OID|Implemented|Web Services|
|Name|Implemented|Web Services|
|Dns Name|Implemented|Web Services|
|OS|Implemented|Web Services|
|Processors|Implemented|Web Services|
|Memory|Implemented|Web Services|
|DataStore OID|Implemented|Web Services|
|Power State|Implemented|Web Services|
|State Change Time|Implemented|Web Services|
```

```
| Host OID | Implemented | Web Services |
|IPs|Implemented|Web Services|
|Guest State|Implemented|Web Services|
|Used Capacity|Implemented|Web Services|
|Provisioned Capacity|Implemented|Web Services|
|MOID|Implemented|Web Services|
.3+|VirtualMachine Disk|OID|Implemented|Web Services|
|VirtualMachine OID|Implemented|Web Services|
|VirtualDisk OID|Implemented|Web Services|
.12+|Host|OID|Implemented|Web Services|
|Name|Implemented|Web Services|
|IPs|Implemented|Web Services|
|Host Domain|Implemented|Web Services|
|Platform Type|Implemented|Web Services|
| Host Installed Memory | Implemented | Web Services |
|Manufacturer|Implemented|Web Services|
| Host Model | Implemented | Web Services |
|Host Cpu Count|Implemented|Web Services|
| Host Cpu Speed | Implemented | Web Services |
|NIC count|Implemented|Web Services|
|NIC speed|Implemented|Web Services|
.8+|Info|DataSource Name|Implemented|Web Services|Info
|Originator ID|Implemented|Web Services|
|Date|Implemented|Web Services|
|Api Name|Implemented|Web Services|
|Api Version|Implemented|Web Services|
|Api Description|Implemented|Web Services|
|Client Api Name|Implemented|Web Services|
|Client Api Version|Implemented|Web Services|
.62+|performance .14+|Data Store|Latency Total|Implemented|Web Services|
|IOps Read|Implemented|Web Services|Number of read IOps on the disk
|Latency Read|Implemented|Web Services|
| IOPs Write | Implemented | Web Services |
|Throughput Read|Implemented|Web Services|
|Throughput Write|Implemented|Web Services|
|Throughput Total|Implemented|Web Services|Average disk total rate (read
and write across all disks) in MB/s
| IOPs Total | Implemented | Web Services |
|Latency Write|Implemented|Web Services|
|Total Capacity|Implemented|Web Services|
| Used Capacity | Implemented | Web Services |
|Capacity Used Ratio|Implemented|Web Services|
|Capacity Provisioned|Implemented|Web Services|
|Over Commit Capacity Ratio | Implemented | Web Services | Reported as a time
.17+|Host|Latency Write|Implemented|Web Services|
```

```
|Throughput total|Implemented|Web Services|IP throughput total
|Swap Rate|Implemented|Web Services|
|Disk IOPs write|Implemented|Web Services|
| IOps Read | Implemented | Web Services | Number of read IOps on the disk
|Total Memory Utilization|Implemented|Web Services|
|Latency Read|Implemented|Web Services|
|Latency Total|Implemented|Web Services|
|ipThroughput.write|Implemented|Web Services|
| IP Throughput Read | Implemented | Web Services |
|Disk Throughput Read|Implemented|Web Services|
|swapRate.inRate|Implemented|Web Services|
|Disk Throughput Write|Implemented|Web Services|
|Total CPU Utilization|Implemented|Web Services|
|diskIops.total|Implemented|Web Services|
|Total Swap Rate|Implemented|Web Services|
|Throughput Read|Implemented|Web Services|total disk throughput read
.9+|Virtual Disk|Latency Total|Implemented|Web Services|
| IOps Read | Implemented | Web Services | Number of read IOps on the disk
|Latency Read|Implemented|Web Services|
| IOPs Write | Implemented | Web Services |
|Throughput Read|Implemented|Web Services|
|Throughput Write|Implemented|Web Services|
|Throughput Total|Implemented|Web Services|Average disk total rate (read
and write across all disks) in MB/s
| IOPs Total | Implemented | Web Services |
|Latency Write|Implemented|Web Services|
.22+|vm|Latency Write|Implemented|Web Services|
|Throughput total|Implemented|Web Services|IP throughput total
|Swap Rate|Implemented|Web Services|
|Disk IOPs write|Implemented|Web Services|
|cpuCoSchedulingDelayTimePercent.total|Implemented|Web Services|
|IOps Read|Implemented|Web Services|Number of read IOps on the disk
|Total Memory Utilization|Implemented|Web Services|
|Idle CPU Time|Implemented|Web Services|idle time in percent
|Latency Read|Implemented|Web Services|
|cpuDemandToEntitlementPercent.total|Implemented|Web Services|
|Latency Total|Implemented|Web Services|
|ipThroughput.write|Implemented|Web Services|
| IP Throughput Read | Implemented | Web Services |
|Disk Throughput Read|Implemented|Web Services|
|CPU Wait Time|Implemented|Web Services|total cpu wait time in percent
|swapRate.inRate|Implemented|Web Services|
|Disk Throughput Write|Implemented|Web Services|
|Total CPU Utilization|Implemented|Web Services|
|Schedule wait time|Implemented|Web Services|Waiting to be scheduled time
in percent
```

```
|diskIops.total|Implemented|Web Services|
|Total Swap Rate|Implemented|Web Services|
|Throughput Read|Implemented|Web Services|total disk throughput read
|===
Management APIs used by this data collector:
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|VMware REST API
|Web Services
|HTTP/HTTPS
180/443
|true
|true
|true
|true
|===
<<top, Back to Top>>
== VMware vSphere (Web Services)
:description: Support Matrix Asciidoc for VMware vSphere (Web Services)
Models and versions supported by this data collector:
|===
<.<|API versions
|VMware ESXi 6.0.0 build-10719132
VMware ESXi 6.0.0 build-2494585
VMware ESXi 6.0.0 build-5572656
VMware ESXi 6.0.0 build-9313334
VMware ESXi 6.5.0 build-14990892
VMware ESXi 6.5.0 build-5969303
VMware ESXi 7.0.0 build-15843807
VMware ESXi 7.0.3 build-20036589
VMware ESXi 7.0.3 build-20328353
VMware ESXi 7.0.3 build-20842708
VMware vCenter Server 5.0.0 build-3073236
VMware vCenter Server 5.0.0 build-455964
```

```
VMware vCenter Server 5.0.0 build-623373
VMware vCenter Server 5.1.0 build-3814779
VMware vCenter Server 5.5.0 build-1750787
VMware vCenter Server 5.5.0 build-2442329
VMware vCenter Server 5.5.0 build-3000241
VMware vCenter Server 5.5.0 build-3252642
VMware vCenter Server 5.5.0 build-3721164
VMware vCenter Server 5.5.0 build-4180647
VMware vCenter Server 5.5.0 build-6516310
VMware vCenter Server 5.5.0 build-9911218
VMware vCenter Server 6.0.0 build-13638472
VMware vCenter Server 6.0.0 build-14510545
VMware vCenter Server 6.0.0 build-2776511
VMware vCenter Server 6.0.0 build-3634793
VMware vCenter Server 6.0.0 build-3634794
VMware vCenter Server 6.0.0 build-5960847
VMware vCenter Server 6.0.0 build-7924803
VMware vCenter Server 6.0.0 build-8803875
VMware vCenter Server 6.0.0 build-9313458
VMware vCenter Server 6.5.0 build-10964411
VMware vCenter Server 6.5.0 build-15679215
VMware vCenter Server 6.5.0 build-17590285
VMware vCenter Server 6.5.0 build-17994927
VMware vCenter Server 6.5.0 build-18499837
VMware vCenter Server 6.5.0 build-18711281
VMware vCenter Server 6.5.0 build-19261680
VMware vCenter Server 6.5.0 build-20510539
VMware vCenter Server 6.5.0 build-7119157
VMware vCenter Server 6.7.0 build-10244857
VMware vCenter Server 6.7.0 build-11727113
VMware vCenter Server 6.7.0 build-13007421
VMware vCenter Server 6.7.0 build-13639324
VMware vCenter Server 6.7.0 build-14368073
VMware vCenter Server 6.7.0 build-15129973
VMware vCenter Server 6.7.0 build-15679289
VMware vCenter Server 6.7.0 build-17137327
VMware vCenter Server 6.7.0 build-18010599
VMware vCenter Server 6.7.0 build-18485185
VMware vCenter Server 6.7.0 build-18831049
VMware vCenter Server 6.7.0 build-19299595
VMware vCenter Server 6.7.0 build-19832247
VMware vCenter Server 6.7.0 build-19832280
VMware vCenter Server 6.7.0 build-20504301
VMware vCenter Server 6.7.0 build-20504362
VMware vCenter Server 6.7.0 build-8170161
VMware vCenter Server 6.7.0 build-9433931
```

```
VMware vCenter Server 7.0.0 build-16620013
VMware vCenter Server 7.0.0 build-16749670
VMware vCenter Server 7.0.1 build-17491160
VMware vCenter Server 7.0.2 build-17694817
VMware vCenter Server 7.0.2 build-17958471
VMware vCenter Server 7.0.2 build-18356314
VMware vCenter Server 7.0.2 build-18455184
VMware vCenter Server 7.0.3 build-18700403
VMware vCenter Server 7.0.3 build-18778458
VMware vCenter Server 7.0.3 build-19234570
VMware vCenter Server 7.0.3 build-19480866
VMware vCenter Server 7.0.3 build-19717403
VMware vCenter Server 7.0.3 build-20051473
VMware vCenter Server 7.0.3 build-20150588
VMware vCenter Server 7.0.3 build-20395099
VMware vCenter Server 7.0.3 build-20845200
VMware vCenter Server 7.0.3 build-20990077
VMware vCenter Server 7.0.3 build-21290409
VMware vCenter Server 7.0.3 build-21477706
VMware vCenter Server 7.0.3 build-21784236
VMware vCenter Server 7.0.3 build-21958406
VMware vCenter Server 8.0.0 build-20920323
VMware vCenter Server 8.0.0 build-21216066
VMware vCenter Server 8.0.0 build-21457384
VMware vCenter Server 8.0.1 build-21560480
VMware vCenter Server 8.0.1 build-21815093
VMware vCenter Server 8.0.1 build-21860503
|===
Products supported by this data collector:
^|Product ^|Category ^|Feature/Attribute ^|Status ^|Protocol Used
^|Additional Information
.91+|foundation .8+|Data Store|Capacity|Implemented|Web Services|Snapshot
Used capacity in MB
|MOID|Implemented|Web Services|
|Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|Provisioned Capacity|Implemented|Web Services|
| Used Capacity | Implemented | Web Services |
|Virtual Center Ip|Implemented|Web Services|
|Type|Implemented|Web Services|
.3+|NasShare DataStore|DataStore OID|Implemented|Web Services|
|Nas Share OID|Implemented|Web Services|
|OID|Implemented|Web Services|
```

```
.4+|NasShare Host|Nas Share OID|Implemented|Web Services|
| Host OID | Implemented | Web Services |
|OID|Implemented|Web Services|
|Read Only|Implemented|Web Services|
.8+|LUN|Disk Name|Implemented|Web Services|
|DataStore OID|Implemented|Web Services|
|Host OID|Implemented|Web Services|
|Number|Implemented|Web Services|
|OID|Implemented|Web Services|
|Policy|Implemented|Web Services|
|TID|Implemented|Web Services|
|Volume Uuid|Implemented|Web Services|
.6+|NAS Share|Capacity|Implemented|Web Services|Allocated capacity in MB
|Filer Ip|Implemented|Web Services|
|Filer Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|Share Path|Implemented|Web Services| For a HvNasShare to be matched to a
Share
|Type|Gap|Web Services|
.6+|Path|Active|Implemented|Web Services|
|Lun OID|Implemented|Web Services|
| Host Port WWPN | Implemented | Web Services |
|OID|Implemented|Web Services|
|Storage Port WWPN|Implemented|Web Services|
|Type|Gap|Web Services|
.6+|Server|Cluster|Implemented|Web Services|Cluster name
|DataCenter Name|Implemented|Web Services|
|Host OID|Implemented|Web Services|
|MOID|Implemented|Web Services|
|OID|Implemented|Web Services|
|Virtual Center Ip|Implemented|Web Services|
.8+|Virtual Disk|Capacity|Implemented|Web Services|Snapshot Used capacity
in MB
|DataStore OID|Implemented|Web Services|
|Lun OID|Implemented|Web Services|
|Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|Type|Gap|Web Services|
|Used Capacity|Implemented|Web Services|used capacity for reporting (MB)
|Used Capacity|Implemented|Web Services|
.15+|VirtualMachine|Dns Name|Implemented|Web Services|
|Guest State|Implemented|Web Services|
|DataStore OID|Implemented|Web Services|
|Host OID|Implemented|Web Services|
|IPs|Implemented|Web Services|
|MOID|Implemented|Web Services|
```

```
|Memory|Implemented|Web Services|
|Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|OS|Implemented|Web Services|
|Power State|Implemented|Web Services|
|State Change Time|Implemented|Web Services|
|Processors|Implemented|Web Services|
|Provisioned Capacity|Implemented|Web Services|
|Used Capacity|Implemented|Web Services|
.3+|VirtualMachine Disk|OID|Implemented|Web Services|
|VirtualDisk OID|Implemented|Web Services|
|VirtualMachine OID|Implemented|Web Services|
.12+|Host|Host Cpu Count|Implemented|Web Services|
| Host Cpu Speed | Implemented | Web Services |
|Host Domain|Implemented|Web Services|
| Host Installed Memory | Implemented | Web Services |
| Host Model | Implemented | Web Services |
|NIC count|Implemented|Web Services|
|NIC speed|Implemented|Web Services|
|IPs|Implemented|Web Services|
|Manufacturer|Implemented|Web Services|
|Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|Platform Type|Implemented|Web Services|
.4+|ISCSI Node|Host Aliases|Implemented|Web Services|
|Node Name|Implemented|Web Services|
|OID|Implemented|Web Services|
|Type|Gap|Web Services|
.8+|Info|Api Description|Implemented|Web Services|
|Api Name|Implemented|Web Services|
|Api Version|Implemented|Web Services|
|Client Api Name|Implemented|Web Services|
|Client Api Version|Implemented|Web Services|
|DataSource Name|Implemented|Web Services|Info
|Date|Implemented|Web Services|
|Originator ID|Implemented|Web Services|
.68+|performance .14+|Data Store|Capacity Provisioned|Implemented|Web
Services|
|Total Capacity|Implemented|Web Services|
|Used Capacity|Implemented|Web Services|
|Over Commit Capacity Ratio | Implemented | Web Services | Reported as a time
series
|Capacity Used Ratio|Implemented|Web Services|
|IOps Read|Implemented|Web Services|Number of read IOps on the disk
| IOPs Total | Implemented | Web Services |
| IOPs Write | Implemented | Web Services |
```

```
|Latency Read|Implemented|Web Services|
|Latency Total|Implemented|Web Services|
|Latency Write|Implemented|Web Services|
|Throughput Read|Implemented|Web Services|
|Throughput Total|Implemented|Web Services|Average disk total rate (read
and write across all disks) in MB/s
|Throughput Write|Implemented|Web Services|
.17+|Host|Total CPU Utilization|Implemented|Web Services|
|IOps Read|Implemented|Web Services|Number of read IOps on the disk
|diskIops.total|Implemented|Web Services|
|Disk IOPs write|Implemented|Web Services|
|Latency Read|Implemented|Web Services|
|Latency Total|Implemented|Web Services|
|Latency Write|Implemented|Web Services|
|Disk Throughput Read|Implemented|Web Services|
|Throughput Read|Implemented|Web Services|total disk throughput read
|Disk Throughput Write|Implemented|Web Services|
| IP Throughput Read | Implemented | Web Services |
|Throughput total|Implemented|Web Services|IP throughput total
|ipThroughput.write|Implemented|Web Services|
|Total Memory Utilization|Implemented|Web Services|
|swapRate.inRate|Implemented|Web Services|
|Swap Rate|Implemented|Web Services|
|Total Swap Rate|Implemented|Web Services|
.12+|Virtual Disk|Total Capacity|Implemented|Web Services|
|Used Capacity|Implemented|Web Services|
|Capacity Used Ratio|Implemented|Web Services|
|IOps Read|Implemented|Web Services|Number of read IOps on the disk
| IOPs Total | Implemented | Web Services |
| IOPs Write | Implemented | Web Services |
|Latency Read|Implemented|Web Services|
|Latency Total|Implemented|Web Services|
|Latency Write|Implemented|Web Services|
|Throughput Read|Implemented|Web Services|
|Throughput Total|Implemented|Web Services|Average disk total rate (read
and write across all disks) in MB/s
|Throughput Write|Implemented|Web Services|
.25+|vm|Total Capacity|Implemented|Web Services|
|Used Capacity|Implemented|Web Services|
|Capacity Used Ratio|Implemented|Web Services|
|cpuCoSchedulingDelayTimePercent.total|Implemented|Web Services|
|cpuDemandToEntitlementPercent.total|Implemented|Web Services|
|Idle CPU Time|Implemented|Web Services|idle time in percent
|CPU Wait Time|Implemented|Web Services|total cpu wait time in percent
|Total CPU Utilization|Implemented|Web Services|
| IOps Read | Implemented | Web Services | Number of read IOps on the disk
```

```
|diskIops.total|Implemented|Web Services|
|Disk IOPs write|Implemented|Web Services|
|Latency Read|Implemented|Web Services|
|Latency Total|Implemented|Web Services|
|Latency Write|Implemented|Web Services|
|Disk Throughput Read|Implemented|Web Services|
|Throughput Read|Implemented|Web Services|total disk throughput read
|Disk Throughput Write|Implemented|Web Services|
| IP Throughput Read | Implemented | Web Services |
|Throughput total|Implemented|Web Services|IP throughput total
|ipThroughput.write|Implemented|Web Services|
|Total Memory Utilization|Implemented|Web Services|
|swapRate.inRate|Implemented|Web Services|
|Swap Rate|Implemented|Web Services|
|Total Swap Rate|Implemented|Web Services|
|Schedule wait time|Implemented|Web Services|Waiting to be scheduled time
in percent
|===
Management APIs used by this data collector:
|===
^|API ^|Protocol Used ^|Transport layer protocol used ^|Incoming ports
used ^|Outgoing ports used ^|Supports authentication ^|Requires only
'Read-only' credentials ^|Supports Encryption ^|Firewall friendly (static
ports)
|VMware REST API
|Web Services
|HTTP/HTTPS
180/443
|true
|true
|true
|true
|===
<<top, Back to Top>>
= Reference & Support
:leveloffset: +1
```

```
[[ID8ef22cc0df5070a4c8a59bcf2480e66a]]
= Requesting Support
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
toc::[]
You can access support options in Cloud Insights by clicking on *Help >
Support*. The support options available to you depend on your Cloud
Insights Edition.
NOTE: The Live Chat support option is not available in Cloud Insights
Federal Edition.
//image:SupportPageExample.png[Support Page]
//image:SupportPageExample-NA.png[Support page]
image:SupportPageWithLearningCenter.png[Support page]
== Activating support entitlement
//Once you have access to Cloud Insights Service shortly after subscribing
in the AWS marketplace, it is strongly recommended that you activate
support entitlement. Activating support entitlement enables you to access
technical support over online chat, web ticketing system, and phone. The
default support level is self-service until registration is completed.
Cloud Insights offers self-service and email support when running in trial
mode. Once you have subscribed to the service, it is strongly recommended
that you activate support entitlement. Activating support entitlement
enables you to access technical support over the online chat, the web
ticketing system, and the phone. The default support mode is self-service
until registration is completed. See xref:{relative path}#obtaining-
support-information[details] below.
During the initial subscription process, your Cloud Insights instance will
generate a 20-digit NetApp serial number starting with "950". This NetApp
serial number represents the Cloud Insights subscription associated with
your account. You must register the NetApp serial number to activate
```

support entitlement. We offer two options for support registration:

- . User with pre-existing NetApp Support Site (NSS) SSO account (e.g. current NetApp customer)
- . New NetApp customer with no pre-existing NetApp Support Site (NSS) SSO account
- === Option 1: Steps for a user with a pre-existing NetApp Support Site (NSS) SSO account

.Steps

- . Navigate to the NetApp registration website https://register.netapp.com
- . Select "I am already registered as NetApp Customer" and choose _Cloud Insights_ as the Product Line. Select your Billing Provider (NetApp or AWS) and provide your Serial Number and your NetApp Subscription Name or AWS Customer ID by referring to the "Help > Support" menu within the Cloud Insights user interface:

+

image:SupportPage SN Section-NA.png[SN Screen]

. Complete the Existing Customer Registration form and click *Submit*.

+

image:ExistingCustomerRegExample.png[existing customer form]

- . If no errors occur, user will be directed to a "Registration Submitted Successfully" page. The email address associated with the NSS SSO username used for registration will receive an email within a couple minutes stating "your product is now eligible for support".
- . This is a onetime registration for the Cloud Insights NetApp serial number.
- === Option 2: Steps for a new NetApp customer with no pre-existing NetApp Support Site (NSS) SSO account

.Steps

- . Navigate to the NetApp registration website https://register.netapp.com
- . Select "I am not a registered NetApp Customer" and complete the required information in example form below:

image:NewCustomerRegExample.png[new customer form]

. Select Cloud Insights as the Product Line. Select your Billing

Provider (NetApp or AWS) and provide your Serial Number and your NetApp Subscription Name or AWS Customer ID by referring to the "Help > Support" menu within the Cloud Insights user interface:

+

image:SupportPage_SN_Section-NA.png[SN_Screen]

- . If no errors occur, user will be directed to a "Registration Submitted Successfully" page. The email address associated with the NSS SSO username used for registration will receive an email within a few hours stating "your product is now eligible for support".
- . As a new NetApp customer, you will also need to create a NetApp Support Site (NSS) user account for future registrations and access to support portal for technical support chat and web ticketing. This link is located at https://mysupport.netapp.com/eservice/public/now.do. You can provide your newly registered Cloud Insights serial number to expedite the process.
- . This is a one-time registration for the Cloud Insights NetApp serial number.
- == Obtaining Support Information

NetApp provides support for Cloud Insights in a variety of ways. Extensive free self-support options are available 24x7, such as knowledgebase (KB) articles or the NetApp community. For users who are subscribed to any of the Cloud Insights Editions (Basic*, Standard, Premium), technical support is available via phone or web ticketing. A NetApp Support Site (NSS) SSO account is required for web ticket along with case management.

*Support is available with Basic Edition as long as all your NetApp storage systems are covered at least at the Premium Support level.

=== Self-Service Support:

These support options are available in Trial mode and are available for free 24×7 :

*

link:https://mysupport.netapp.com/site/search?q=cloud%20insights&offset=0 &searchType=Manual&autocorrect=true&origin=CI_Suppport_KB&filter=%28conten t_type%3D%3D%22knowledgebase%22;product%3D%3D%22Cloud%20Insights%22%29[Knowledgebase]

+

Clicking the links in this section takes you to the NetApp Knowledgebase, where you can search through relevant articles, how-to's, and more.

////

link:https://kb.netapp.com/app/browse/a_status/published/channelRecordID/H OW_TO/currentSelectedID/RN_PRODUCT_473/isProductSelected/true/isRecommenda tionAllowed/true/pageSize/10/productRecordID/RN_PRODUCT_473/sortColumn/pub lishDate/sortDirection/DESC/truncate/200/type/browse[How-to's], link:https://kb.netapp.com/app/browse/a_status/published/channelRecordID/F AQ/currentSelectedID/RN_PRODUCT_473/isProductSelected/true/isRecommendatio nAllowed/true/pageSize/10/productRecordID/RN_PRODUCT_473/sortColumn/publis hDate/sortDirection/DESC/truncate/200/type/browse[FAQ's], or link:https://kb.netapp.com/app/browse/a_status/published/channelRecordID/B REAK_FIX/currentSelectedID/RN_PRODUCT_473/isProductSelected/true/isRecomme ndationAllowed/true/pageSize/10/productRecordID/RN_PRODUCT_473/sortColumn/publishDate/sortDirection/DESC/truncate/200/type/browse[Break Fix] information related to Cloud Insights.

////

* *link:https://docs.netapp.com/us-en/cloudinsights/[Documentation]*

+

Clicking on the Documentation link takes you to this documentation center.

*

link:https://mysupport.netapp.com/site/search?q=cloud%20insights&offset=0 &searchType=Manual&autocorrect=true&origin=CI_Support_Community&filter=%28 content_type%3D%3D%22community%22;product%3D%3D%22Cloud%20Insights%22%29[Community]

+

Clicking on the community link takes you to the NetApp Cloud Insights community, where you can connect with peers and experts.

There is also a link to provide xref:{relative_path}mailto:ng-cloudinsights-customerfeedback@netapp.com[*Feedback*] to help us improve Cloud Insights.

=== Subscription Support

In addition to the self-support options above, if you have a Cloud Insights subscription or paid support for monitored NetApp products or services, you can work with a NetApp Support Engineer to resolve your problem.

NOTE: You must register in order to <<Activating support entitlement and accessing support, activate support>> for NetApp Cloud products. To register, go to NetApp's link:https://register.netapp.com[Cloud Data Services Support Registration].

It is highly recommended that you check the box to allow a NetApp Support Engineer access to your Cloud Insights environment during your support session. This will allow the engineer to troubleshoot the problem and help you resolve it quickly. When your issue is resolved or your support session has ended, you can un-check the box.

You can request support by any of the following methods. You must have an active Cloud Insights subscription to use these support options:

* link:https://www.netapp.com/us/contact-us/support.aspx[*Phone*]

link:https://mysupport.netapp.com/portal?_nfpb=true&_st=initialPage=true&_
pageLabel=submitcase[*Support Ticket*]

* *Chat* - You will be connected with NetApp support personnel for assistance (weekdays only). Chat is available in the *Help > Live Chat* menu option in the upper right of any Cloud Insights screen.

You can also request sales support by clicking on the link:https://www.netapp.com/us/forms/sales-inquiry/cloud-insights-sales-inquiries.aspx[*Contact Sales*] link.

Your Cloud Insights serial number is visible within the service from the *Help > Support* menu. If you are experiencing issues accessing the service and have registered a serial number with NetApp previously, you can also view your list of Cloud Insights serial numbers from the NetApp Support Site as follows:

- * Login to mysupport.netapp.com
- * From the Products > My Products menu tab, use Product Family "SaaS Cloud Insights" to locate all your registered serial numbers:

image:Support View SN.png[View Support SN]

== Cloud Insights Data Collector Support Matrix

You can view or download information and details about supported Data Collectors in the

xref:{relative_path}reference_data_collector_support_matrix.html[*Cloud
Insights Data Collector Support Matrix*, role="external"].

=== Learning Center

Regardless of your subscription, *Help > Support* links to several NetApp University course offerings to help you get the most out of Cloud Insights. Check them out!

```
= Data Collector Reference - Infrastructure
:leveloffset: +1
[[IDf10b5f428939dec54ffa2eb33acbca1b]]
= Vendor-Specific Reference
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The topics in this section provide vendor-specific reference information.
In most cases, configuring a data collector is straightforward. In some
cases, you may need additional information or commands to properly
configure the data collector.
Click on a *vendor* in the menu to the left to see information for their
data collectors.
[[ID6f55bb51576e8727d36368ffebb19f53]]
= Configuring the Amazon EC2 data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Amazon EC2 data collector to acquire inventory and
performance data from EC2 instances.
== Requirements
In order to collect data from Amazon EC2 devices, you must have the
```

following information:

- * You must have one of the following:
- ** The *IAM Role* for your Amazon EC2 cloud account, if using IAM Role Authentication. IAM Role only applies if your acquisition unit is installed on an AWS instance.
- ** The *IAM Access Key* ID and Secret Access Key for your Amazon EC2 cloud account, if using IAM Access Key authentication.
- * You must have the "list organization" privilege
- * Port 443 HTTPS
- * EC2 Instances can be reported as a Virtual Machine, or (less naturally) a Host. EBS Volumes can be reported as both a VirtualDisk used by the VM, as well as a DataStore providing the Capacity for the VirtualDisk.

Access keys consist of an access key ID (for example, AKIAIOSFODNN7EXAMPLE) and a secret access key (for example, wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY). You use access keys to sign programmatic requests that you make to EC2 if you use the Amazon EC2 SDKs, REST, or Query API operations. These keys are provided with your contract from Amazon.

== Configuration

Enter data into the data collector fields according to the table below:

```
[cols=2*, options="header", cols"50,50"]
|===
```

|Field | Description

|AWS Region|Choose AWS region

|IAM Role|For use only when acquired on an AU in AWS. See below for more information on xref:{relative_path}task_dc_amazon_ec2.html#iam-roles[IAM Roles].

|AWS IAM Access Key ID|Enter AWS IAM Access Key ID. Required if you do not use IAM Role.

|AWS IAM Secret Access Key|Enter AWS IAM Secret Access Key. Required if you do not use IAM Role.

|I understand AWS bills me for API requests|Check this to verify your understanding that AWS bills you for API requests made by Cloud Insights polling.

|===

== Advanced Configuration

```
[cols=2*, options="header", cols"50,50"]
|===
```

|Field | Description

|Include Extra Regions| Specify additional regions to include in polling. |Cross Account Role|Role for accessing resources in different AWS accounts.

|Inventory Poll Interval (min)|The default is 60

|Choose 'Exclude' or 'Include' to Apply to Filter VMs by Tags|Specify whether to include or exclude VM's by Tags when collecting data. If 'Include' is selected, the Tag Key field can not be empty.

|Tag Keys and Values on which to Filter VMs|Click *+ Filter Tag* to choose which VMs (and associated disks) to include/exclude by filtering for keys and values that match keys and values of tags on the VM. Tag Key is required, Tag Value is optional. When Tag Value is empty, the VM is filtered as long as it matches the Tag Key.

//|HTTP connection and socket timeout (sec)|The default is 300 //|Include AWS tags|Check to enable support for AWS tags in Cloud Insights annotations.

|Performance Poll Interval (sec)|The default is 1800

|CloudWatch Agent Metrics Namespace

|Namespace in EC2/EBS from which to collect data. Note that if the names of the default metrics in this namespace are changed, Cloud Insights may not be able to collect that renamed data. It is recommended to leave the default metric names.

|===

== IAM Access Key

Access keys are long-term credentials for an IAM user or the AWS account root user. Access keys are used to sign programmatic requests to the AWS CLI or AWS API (directly or using the AWS SDK).

Access keys consist of two parts: an access key ID and a secret access key. When you use _IAM Access Key_ authentication (as opposed to _IAM Role_ authentication), you must use both the access key ID and secret access key together for authentication of requests. For more information, see the Amazon documentation on

link:https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_acces
s-keys.html[Access Keys].

== IAM Role

When using _IAM Role_ authentication (as opposed to IAM Access Key authentication), you must ensure that the role you create or specify has the appropriate permissions needed to access your resources.

For example, if you create an IAM role named _InstanceEc2ReadOnly_, you must set up the policy to grant EC2 read-only list access permission to all EC2 resources for this IAM role. Additionally, you must grant STS (Security Token Service) access so that this role is allowed to assume roles cross accounts.

After you create an IAM role, you can attach it when you create a new EC2 instance or any existing EC2 instance.

After you attach the IAM role _InstanceEc2ReadOnly_ to an EC2 instance, you will be able to retrieve the temporary credential through instance metadata by IAM role name and use it to access AWS resources by any application running on this EC2 instance.

For more information see the Amazon documentaiton on link:https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles.html[IAM Roles].

Note: IAM role can be used only when the Acquisition Unit is running in an AWS instance.

== Mapping Amazon tags to Cloud Insights annotations

The Amazon EC2 data collector includes an option that allows you to populate Cloud Insights annotations with tags configured on EC2. The annotations must be named exactly as the EC2 tags. Cloud Insights will always populate same-named text-type annotations, and will make a "best attempt" to populate annotations of other types (number, boolean, etc). If your annotation is of a different type and the data collector fails to populate it, it may be necessary to remove the annotation and re-create it as a text type.

Note that AWS is case-sensitive, while Cloud Insights is case-insensitive. So if you create an annotation named "OWNER" in Cloud Insights, and tags named "OWNER", "Owner", and "owner" in EC2, all of the EC2 variations of "owner" will map to Cloud Insight's "OWNER" annotation.

////

.Related Information

* https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html[Managing Access Keys for IAM Users^]
////

== Include Extra Regions

In the AWS Data Collector *Advanced Configuration* section, you can set

the *Include extra regions* field to include additional regions, separated by comma or semi-colon. By default, this field is set to *_us-.*_*, which collects on all US AWS regions. To collect on _all_ regions, set this field to * .* *.

If the *Include extra regions* field is empty, the data collector will collect on assets specified in the *AWS Region* field as specified in the *Configuration* section.

== Collecting from AWS Child Accounts

Cloud Insights supports collection of child accounts for AWS within a single AWS data collector. Configuration for this collection is performed in the AWS environment:

- * You must configure each child account to have an AWS Role that allows the main account ID to access EC2 details from the children account.
- * Each child account must have the role name configured as the same string.
- * Enter this role name string into the Cloud Insights AWS Data Collector *Advanced Configuration* section, in the *Cross account role* field.

Best Practice: It is highly recommended to assign the AWS predefined _AmazonEC2ReadOnlyAccess_ policy to the EC2 main account. Also, the user configured in the data source should have at least the predefined AWSOrganizationsReadOnlyAccess policy assigned, in order to query AWS.

Please see the following for information on configuring your environment to allow Cloud Insights to collect from AWS child accounts:

link:https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_crossaccount-with-roles.html[Tutorial: Delegate Access Across AWS Accounts
Using IAM Roles]

link:https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_commonscenarios_aws-accounts.html[AWS Setup: Providing Access to an IAM User in
Another AWS Account That You Own]

link:https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_foruser.html[Creating a Role to Delegate Permissions to an IAM User]

== Troubleshooting

Additional information on this Data Collector may be found from the xref:{relative_path}concept_requesting_support.html[Support] page or in the xref:{relative path}reference data collector support matrix.html[Data

```
Collector Support Matrix].
[[ID2c387bf0f3b5f0edaac716826249fea9]]
= Amazon FSx for NetApp ONTAP data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector acquires inventory and performance data from Amazon
FSx for NetApp ONTAP. This data collector will be made available
incrementally throughout the Cloud Insights service regions. Contact your
sales person if you do not see the Icon for this collector in your Cloud
Insights Environment.
NOTE: This Cloud Insights collector requires an ONTAP user with a
Filesystem-Scoped role. Please review the AWS
link:https://docs.aws.amazon.com/fsx/latest/ONTAPGuide/roles-and-
users.html[Roles and Rules] documentation for available options. At this
time AWS supports only one kind of user role with Filesystem Scope, which
is fsxadmin . This is the appropriate role to be used for the Cloud
Insights collector. The user should also have all three of these
applications assigned to it: http, ontapi, ssh.
== Terminology
Cloud Insights acquires inventory and performance data from the FSx-NetApp
data collector. For each asset type acquired, the most common terminology
used for the asset is shown. When viewing or troubleshooting this data
collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Cluster|Storage
|LUN|Volume
|Volume|Internal Volume
|===
```

== FSx-NetApp Terminology

The following terms apply to objects or references that you might find on FSx-NetApp storage asset landing pages. Many of these terms apply to other data collectors as well.

=== Storage

- * Model A comma-delimited list of the unique, discrete model names within this cluster.
- * Vendor AWS
- * Serial number The array serial number.
- * IP generally will be the IP(s) or hostname(s) as configured in the data source.
- //* Microcode version firmware.
- * Raw Capacity base 2 summation of all the SSD storage assigned to the FSx filesystem.
- * Latency a representation of what the host facing workloads are experiencing, across both reads and writes. Ideally, Cloud Insights is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, Cloud Insights is generally performing an IOPs-weighted calculation derived from the individual internal volumes' statistics.
- * Throughput aggregated from internal volumes.

Management - this may contain a hyperlink for the management interface of the device. Created programmatically by the Cloud Insights data source as part of inventory reporting.

=== Storage Pool

- * Storage what storage array this pool lives on. Mandatory.
- * Type a descriptive value from a list of an enumerated list of possibilities. Most commonly will be "Aggregate" or "RAID Group"".
- //* Node if this storage array's architecture is such that pools belong
 to a specific storage node, its name will be seen here as a hyperlink to
 its own landing page.
- * Capacity the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these.
- * IOPS the sum IOPs of all the volumes allocated on this storage pool.
- * Throughput the sum throughput of all the volumes allocated on this storage pool.

== Requirements

The following are requirements to configure and use this data collector:

```
* You must have access to an account with the "fsxadmin" role, with three
applications assigned to it - ssh, ontapi, http
* Account details include username and password.
* Port requirements: 443
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|NetApp Management IP | IP address or fully-qualified domain name of the
NetApp cluster
|User Name |User name for NetApp cluster
|Password |Password for NetApp cluster
|===
== Advanced Metrics
This data collector collects the following advanced metrics from the FSx
for NetApp ONTAP storage:
* fpolicy
* nfsv3
* nfsv3:node
* nfsv4
* nfsv4 1
* nfsv4 1:node
* nfsv4:node
* policy group
* qtree
* volume
* workload volume
Note that FSx CLI and API commands retrieve some capacity values that
Cloud Insights ZAPI doesn't collect, so certain capacity values (such as
those for storage pools) may be different in Cloud Insights than they are
on the FSx itself.
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
```

|===

|Problem: |Try this:

|Receive 401 HTTP response or 13003 ZAPI error code and ZAPI returns "Insufficient privileges" or "not authorized for this command" | Check username and password, and user privileges/permissions.

|ZAPI returns "cluster role is not cluster_mgmt LIF"|AU needs to talk to cluster management IP. Check the IP and change to a different IP if necessary

|ZAPI command fails after retry| AU has communication problem with the cluster. Check network, port number, and IP address. User should also try to run a command from command line from the AU machine.

|AU failed to connect to ZAPI via HTTP| Check whether ZAPI port accepts plaintext. If AU tries to send plaintext to an SSL socket, the communication fails.

|Communication fails with SSLException|AU is attempting to send SSL to a plaintext port on a filer. Check whether the ZAPI port accepts SSL, or use a different port.

|Additional Connection errors:

ZAPI response has error code 13001, "database is not open"

ZAPI error code is 60 and response contains "API did not finish on time"

ZAPI response contains "initialize session() returned NULL environment"

ZAPI error code is 14007 and response contains "Node is not healthy"

| Check network, port number, and IP address. User should also try to run a command from command line from the AU machine.

|===

Additional information may be found from the xref:{relative_path}concept_requesting_support.html[Support] page or in the xref:{relative_path}reference_data_collector_support_matrix.html[Data Collector Support Matrix].

```
[[ID9914c95b38c03dab796103c5bd0cd831]]
= Configuring the Azure compute data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Azure compute data collector to acquire inventory
and performance data from Azure compute instances.
== Requirements
You need the following information to configure this data collector.
* Port requirement: 443 HTTPS
* Azure OAuth 2.0 Redirect URI (login.microsoftonline.com)
* Azure Management Rest IP (management.azure.com)
* Azure Resource Manager IP (management.core.windows.net)
* Azure Service Principal Application (Client) ID (Reader role required)
* Azure service principal authentication key (user password)
* You need to set up an Azure account for Cloud Insights discovery.
Once the account is properly configured and you register the application
in Azure, you will have the credentials required to discover the Azure
instance with Cloud Insights. The following link describes how to set up
the account for discovery.
https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-
create-service-principal-portal
== Configuration
Enter data into the data collector fields according to the table below:
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|Azure Service Principal Application (Client) ID (Reader role
required) | Sign-in ID to Azure. Requires Reader Role access.
|Azure tenant ID|Microsoft tenant ID
|Azure Service Principal Authentication Key|Login authentication key
|I understand Microsoft bills me for API requests|Check this to verify
```

```
your understanding that Microsoft bills you for API requests made by
Insight polling.
|===
== Advanced Configuration
[cols=2*, options="header", cols"50,50"]
|Field | Description
|Inventory Poll Interval (min)|The default is 60
//|HTTP connection and socket timeout (sec)|The default is 300
|Choose 'Exclude' or 'Include' to Apply to Filter VMs by Tags|Specify
whether to include or exclude VM's by Tags when collecting data. If
'Include' is selected, the Tag Key field can not be empty.
|Tag Keys and Values on which to Filter VMs|Click *+ Filter Tag* to choose
which VMs (and associated disks) to include/exclude by filtering for keys
and values that match keys and values of tags on the VM. Tag Key is
required, Tag Value is optional. When Tag Value is empty, the VM is
filtered as long as it matches the Tag Key.
|Performance Poll Interval (sec)|The default is 300
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
= Broadcom
:leveloffset: +1
[[ID834b8f3fdd1104b5b55315e04ea96a29]]
= Brocade Network Advisor data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
```

```
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Brocade Network Advisor data collector to acquire
inventory and performance data from Brocade switches.
== Terminology
Cloud Insights acquires the following inventory information from the
Brocade Network Advisor data collector. For each asset type acquired by
Cloud Insights, the most common terminology used for this asset is shown.
When viewing or troubleshooting this data collector, keep the following
terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Switch|Switch
|Port|Port
|Virtual Fabric, Physical Fabric|Fabric
|Logical Switch|Logical Switch
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
The following are required to configure this data collector:
* The Cloud Insights Acquisition Unit will initate connections to TCP port
443 on the BNA server. BNA server must be running version 14.2.1 or
higher.
* Brocade Network Advisor Server IP address
* User name and password to an administrator account
* Port requirement: HTTP/HTTPS 443
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Brocade Network Advisor Server IP|IP address of the Network Advisor
Server
```

```
|User Name|User name for the switch
|User Name|Administrator user name
|Password|Administrator password
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Connection Type|HTTPS (default port 443) or HTTP (default port 80)
|Override Connection Port | If blank, use the default port in the
Connection Type field, otherwise enter the connection port to use
|Password|Password for the switch
|Inventory poll interval (min) |The default is 40
//|Connection timeout (sec)|The default is 60
|Report Access Gateway|Check to include devices in Access Gateway mode
|Performance Poll Interval (sec)|The default is 1800
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|Receive a message that more than 1 node is logged into the Access Gateway
port, or data collector fails to discover Access Gateway device. | Check
that the NPV device is operating correctly and that all connected WWNs are
expected. Do not directly acquire the NPV device. Instead, acquisition of
the core fabric switch will collect the NPV device data.
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID9d772d4dfe46012ec4c47182fe5a2a78]]
= Brocade FC Switch data collector
:toc: macro
```

```
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Brocade FC Switch (SSH) data source to discover
inventory for Brocade or rebranded switch devices running Factored
Operating System (FOS) firmware 4.2 and later. Devices in both FC switch
and Access Gateway modes are supported.
== Terminology
Cloud Insights acquires the following inventory information from the
Brocade FC Switch data collector. For each asset type acquired by Cloud
Insights, the most common terminology used for this asset is shown. When
viewing or troubleshooting this data collector, keep the following
terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Switch|Switch
|Port|Port
|Virtual Fabric, Physical Fabric|Fabric
|Zone|Zone
|Logical Switch|Logical Switch
|Virtual Volume|Volume
|LSAN Zone|IVR Zone
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* The Cloud Insights Acquisition Unit (AU) will initiate connections to
TCP Port 22 on Brocade switches to collect inventory data. The AU will
also initiate connections to UDP port 161 for collection of performance
* There must be IP connectivity to all switches in the fabric. If you
```

select the Discover all switches in the fabric check box, Cloud Insights

identifies all the switches in the fabric; however, it needs IP connectivity to these additional switches to discover them. * The same account is needed globally across all switches in the fabric. You can use PuTTY (open source terminal emulator) to confirm access. * Ports 161 and 162 must be open to all switches in the fabric for SNMP performance polling. * SNMP read-only Community String == Configuration [cols=2*, options="header", cols"50,50"] |Field|Description |Switch IP|IP address or fully-qualified domain name of the EFC Server |User Name|User name for the switch |Password|Password for the switch |SNMP | SNMP version |SNMP Community String|SNMP read-only community string used to access the switch |SNMP User Name|SNMP user name |SNMP Password|SNMP password |=== == Advanced configuration [cols=2*, options="header", cols"50,50"] |Field|Description |Fabric name|Fabric name to be reported by the data collector. Leave blank to report the fabric name as WWN. |Inventory Poll Interval (min) | Interval between inventory polls. The default is 15. |Excluded Devices|Comma-separated list of device IDs to exclude from //|Timeout (sec)|Connection timeout. The default is 30. //|Banner Wait Timeout (sec)|SSHAdmin Domains Active Select if using Admin Domainsbanner wait timeout. The default is 5 seconds. |Admin Domains Active|Select if using Admin Domains |Retrieve MPR Data|Select to acquire routing data from your multiprotocol router. |Enable Trapping|Select to enable acquisition upon receiving an SNMP trap from the device. If you select enable trapping, you must also activate SNMP. |Minimum Time Between Traps (sec)|Minimum time between acquisition attempts triggered by traps. The default is 10. |Discover all switches in the fabric|Select to discover all switches in

the fabric |Choose Favoring HBA vs. Zone Aliases|Choose whether to favor HBA or zone aliases |Performance Poll Interval (sec)|Interval between performance polls. The | SNMP Auth Protocol | SNMP authentication protocol (SNMP v3 only) |SNMP Privacy Password|SNMP privacy password (SNMP v3 only) |SNMP Retries|Number of SNMP retry attempts //|SNMP Timeout (ms)|SNMP timeout. The default is 5000. |=== == Troubleshooting Some things to try if you encounter problems with this data collector: === Inventory [cols=2*, options="header", cols"50,50"] |=== |Problem:|Try this: |The inventory acquisition of the Brocade datasource fails with the error: <date> <time> ERROR [com.onaro.sanscreen.acquisition.framework.datasource.BaseDataSource] Error 2 out of 2: <datasource name> [Internal error] - Unable to generate the model for device <IP>. Error detecting prompt ([Device name <name>]: Unable to generate the model for device <IP>. Error detecting prompt) |The issue may be caused when the Brocade switch takes too long to return with a prompt, exceeding the default timeout of 5 seconds. In the data collector's Advanced Configuration settings in Cloud Insights, try increasing the SSH Banner Wait Timeout (sec) to a higher value. |Error: "Cloud Insights received Invalid Chassis Role" |Check that the user configured in this data source has been granted the chassis role permission. |Error: "Mismatched Chassis IP Address" |Change the data source configuration to use chassis IP address. |Receive a message that more than 1 node is logged into the Access Gateway port | Check that the NPV device is operating correctly and that all connected WWNs are expected. Do not directly acquire the NPV device. Instead, acquisition of the core fabric switch will collect the NPV device data. |Performance collection fails with "Timed out during sending SNMP

```
|Depending on query variables and switch configuration, some queries may
exceed the default timeout.
link:https://kb.netapp.com/Cloud/BlueXP/Cloud Insights/Cloud Insight Broca
de data source fails performance collection with a timeout due to default
SNMP configuration[Learn More].
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[IDde9a712b1706d0b5a1dfde05dc14d2c0]]
= Brocade FOS REST Data Collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Brocade FOS REST collector to discover inventory
and performance for Brocade switch devices running FabricOS (FOS) firmware
8.2 and later.
By default, this collector will attempt to discover all the FOS devices
that are a part of all the fabrics the switch is part of.
== Terminology
Cloud Insights acquires the following inventory information from the
Brocade FOS REST data collector. For each asset type acquired by Cloud
Insights, the most common terminology used for this asset is shown. When
viewing or troubleshooting this data collector, keep the following
terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term|Cloud Insights Term
```

```
|Switch|Switch
|Port|Port
|Virtual Fabric, Physical Fabric|Fabric
|Zone|Zone
|Logical Switch|Logical Switch
|Virtual Volume|Volume
|LSAN Zone|IVR Zone
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* There must be TCP connectivity to all switches in the fabric. This data
collector type will seamlessly try both HTTP and HTTPS for each device in
the fabric. If you select the Discover all switches in the fabric check
box, Cloud Insights identifies all the switches in the fabric; however, it
needs TCP connectivity to these additional switches to discover them.
* The same account is needed globally across all switches in the fabric.
You can use the device's Web interface to confirm access.
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|Switch IP|IP address or fully-qualified domain name of the FOS switch
|User Name|User name for the switch
|Password|Password for the switch
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Excluded Devices|Comma-separated list of device IPv4 addresses to exclude
from polling.
|Inventory Poll Interval (min) | Interval between inventory polls. The
default is 60.
|Discover all switches in the fabric|Select to discover all switches in
|Choose Favoring HBA vs. Zone Aliases|Choose whether to favor HBA or zone
aliases.
```

|Connection type| HTTP or HTTPS.|Note that this setting only changes which protocol CI attempts to use per device first - CI will attempt the opposite protocol automatically if the default fails |Override TCP Port|Specify a port if not using the default. |Performance Poll Interval (sec)|Interval between performance polls. The default is 300.

|===

== Troubleshooting

Some things to try if you encounter problems with this data collector:

=== Inventory

[cols=2*, options="header", cols"50,50"]
|===

|Problem:|Try this:

|The Test feature warns me that a protocol is inaccessible

|A given Brocade FOS 8.2+ device will only want to speak on HTTP or HTTPS - if a switch has a digital certificate installed, the switch will throw HTTP errors if one attempts to communicate to it with unencrypted HTTP versus HTTPS. The test feature attempts communication with both HTTP and HTTPS - if the Test tells you that one protocol passes, you can safely save the collector and not worry that the other protocol was unsuccessful - the collector will attempt both protocols during collection, and only fail if neither works.

|Error: "Cloud Insights received Invalid Chassis Role" |Check that the user configured in this data source has been granted the chassis role permission.

|Error: "Mismatched Chassis IP Address" |Change the data source configuration to use chassis IP address.

|Receive a message that more than 1 node is logged into the Access Gateway port| Check that the NPV device is operating correctly and that all connected WWNs are expected. Do not directly acquire the NPV device. Instead, acquisition of the core fabric switch will collect the NPV device data.

|Performance collection fails with "Timed out during sending SNMP request".

|Depending on query variables and switch configuration, some queries may exceed the default timeout.

link:https://kb.netapp.com/Cloud/BlueXP/Cloud Insights/Cloud Insight Broca

```
de data source fails performance collection with a timeout due to default
SNMP configuration[Learn More].
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
:leveloffset: -1
[[ID99d561bfb646b6d65907bbdecf5d263c]]
= Cisco MDS Fabric Switches data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Cisco MDS Fabric Switches data collector to
discover inventory for Cisco MDS Fabric Switches as well as a variety of
Cisco Nexus FCoE switches on which the FC service is enabled.
Additionally, you can discover many models of Cisco devices running in NPV
mode with this data collector.
== Terminology
Cloud Insights acquires the following inventory information from the Cisco
FC Switch data collector. For each asset type acquired by Cloud Insights,
the most common terminology used for this asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
```

```
|Switch|Switch
|Port|Port
|VSAN|Fabric
|Zone|Zone
|Logical Switch|Logical Switch
|Name Server Entry|Name Server Entry
|Inter-VSAN Routing (IVR) Zone|IVR Zone
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* An IP address of one switch in the fabric or individual switches
* Chassis discovery, to enable fabric discovery
* If using SNMP V2, read-only community string
* Port 161 is used to access the device
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Cisco Switch IP|IP address or fully-qualified domain name of the switch
|SNMP Version|Select V1, V2, or V3. V2 or later is required for
performance acquisition.
|SNMP Community String|SNMP read-only community string used to access the
switch (not applicable for SNMP v3)
|User Name|User name for the switch (SNMP v3 only)
|Password|Password used for the switch (SNMPv3 only)
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls (default
40 minutes)
|SNMP Auth Protocol|SNMP authentication protocol (SNMPv3 only)
| SNMP Privacy Protocol | SNMP privacy protocol (SNMPv3 only)
|SNMP Privacy Password|SNMP Privacy Password
|SNMP Retries|Number of SNMP retry attempts
|SNMP Timeout (ms)|SNMP timeout (default 5000 ms)
|Enable Trapping|Select to enable trapping. If you enable trapping, you
```

must also activate SNMP notifications. |Minimum Time Between Traps (sec)|Minimum time between acquisition attempts triggered by traps (default 10 seconds) |Discover All Fabric Switches|Select to discover all switches in the |Excluded Devices|Comma-separated list of device IPs to exclude from polling |Included Devices|Comma-separated list of device IPs to include in polling |Check Device Type|Select to accept only those devices that explicitly advertise themselves as Cisco devices |First Alias Type|Provide a first preference for resolution of the alias. Choose from the following: *Device Alais* This is a user-friendly name for a port WWN (pWWN) that can be used in all configuration commands, as required. All switches in the Cisco MDS 9000 Family support Distributed Device Alias Services (device aliases). *None* Do not report any alias. *Port Description* A description to help identify the port in a list of ports. *Zone Alias (all) * A user-friendly name for a port that can be used only for the active configuration. This is the default. |Second Alias Type|Provide a second preference for resolution of the alias |Third Alias Type|Provide a third preference for resolution of the alias |Enable SANTap Proxy Mode Support|Select if your Cisco switch is using SANTap in proxy mode. If you are using EMC RecoverPoint, then you are probably using SANTap. |Performance Poll Interval (sec)|Interval between performance polls (default 300 seconds) |=== == Troubleshooting Some things to try if you encounter problems with this data collector: === Inventory [cols=2*, options="header", cols"50,50"] |=== |Problem:|Try this: |Error: Failed to discover chassis - no switches have been discovered | • Ping the device with the IP configured

- Login to the device using Cisco Device Manager GUI
- Login to the device using CLI
- Try to run SNMP walk

|Error: Device is not a Cisco MDS switch

- | Make sure the data source IP configured for the device is correct
- Login to the device using Cisco Device Manager GUI
- Login to the device using CLI

|Error: Cloud Insights is not able to obtain the switch's WWN. |This may not be a FC or FCoE switch, and as such may not be supported. Make sure the IP/FQDN configured in the datasource is truly a FC/FCoE switch.

|Error: Found more than one nodes logged into NPV switch port

|Disable direct acquisition of the NPV switch

|Error: Could not connect to the switch

- | Make sure the device is UP
- Check the IP address and listening port
- Ping the device
- Login to the device using Cisco Device Manager GUI
- Login to the device using CLI
- Run SNMP walk

```
|===
```

```
=== Performance
```

[cols=2*, options="header", cols"50,50"]

| ===

|Problem:|Try this:

|Error: Performance acquisition not supported by SNMP v1

- | Edit Data Source and disable Switch Performance
- Modify Data Source and switch configuration to use SNMP v2 or higher \mid ===

Additional information may be found from the

xref:{relative_path}concept_requesting_support.html[Support] page or in
the xref:{relative_path}reference_data_collector_support_matrix.html[Data
Collector Support Matrix].

```
[[IDb620847bdb9ce38d237939a7913e8edd]]
```

= Cohesity SmartFiles data collector

:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font

:linkattrs:

:relative path: ./

```
:imagesdir: {root path}{relative path}./media/
[.lead]
This REST API-based collector will acquire a Cohesity cluster, discovering
the "Views" (as Cloud Insights Internal Volumes), the various nodes, as
well as collecting performance metrics.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Cohesity Cluster IP|IP address of the Cohesity cluster
|User Name|User name for the Cohesity cluster
|Password|Password used for the Cohesity cluster
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|TCP Port|Port used for TCP communication with the Cohesity cluster
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 60 minutes.
|Performance Poll Interval (min)|Interval between performance polls. The
default is 900 seconds.
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
= Dell
:leveloffset: +1
[[IDbfb5b8f844eb9d360c742bb44887aabd]]
= Dell EMC XC Series data collector
:toc: macro
:hardbreaks:
```

```
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to discover inventory and
performance information for the Dell EMC XC Series storage arrays.
////
== Terminology
Cloud Insights acquires the following inventory information from this data
collector. For each asset type acquired by Cloud Insights, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Disk|Disk
|Disk Folder|Disk Group
|Storage Center|Storage
|Controller|Storage Node
|Storage Type|Storage Pool
|Volume|Volume
|Fiber Channel I/O Port|Port
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* Administrator credentials for the Dell EMC XC Enterprise Manager server
* IP address of the XC Enterprise Manager server
////
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|Prism External IP Address|IP address of the XC server
```

```
|User Name|User name for the XC server
|Password|Password used for the XC server
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|TCP Port|Port used for TCP communication with the XC server
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 60 minutes.
|Performance Poll Interval (min)|Interval between performance polls. The
default is 300 seconds.
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
:leveloffset: -1
= Dell EMC
:leveloffset: +1
[[IDfe579555d0b33858ad5d417d8168738a]]
= DELL EMC Data Domain data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector gathers inventory and performance information from
DELL EMC Data Domain deduplication storage systems. To configure this data
collector, there are specific configuration instructions and usage
```

```
recommendations you must follow.
== Terminology
Cloud Insights acquires the following inventory information from the Data
Domain data collector. For each asset type acquired by Cloud Insights, the
most common terminology used for this asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term|Cloud Insights Term
|Disk|Disk
|Array|Storage
|FC Port|Port
|File System|Internal Volume
|Quota|Quota
|NFS and CIFS share|FileShare
Note: These are common terminology mappings only and might not represent
every case for this data colletor.
== Requirements
You need the following information to configure this data collector:
* IP address of the Data Domain device
* Read-only user name and password to the Data Domain storage
* SSH port 22
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
| IP address | The IP address or fully-qualified domain name of the Data
Domain storage array
|User name|The user name for the Data Domain storage array
|Password|The password for the Data Domain storage array
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
```

```
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 20.
//|SSH Process Wait Timeout (sec)|SSH process timeout. The default is 180.
|SSH Port|SSH service port
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID9e4a5deeefbc9e5d53956e2821afc1e7]]
= Configuring the EMC ECS data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector acquires inventory and performance data from EMC ECS
storage systems. For configuration, the data collector requires an IP
address or hostname of the ECS cluster and a username and password.
NOTE: Dell EMC ECS is metered at a different Raw TB to Managed Unit rate.
Every 40 TB of unformatted ECS capacity is charged as 1
xref:{relative path}concept subscribing to cloud insights.html#pricing[Man
aged Unit (MU)].
== Terminology
Cloud Insights acquires the following inventory information from the ECS
data collector. For each asset type acquired, the most common terminology
used for this asset is shown. When viewing or troubleshooting this data
collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
```

```
|Vendor/Model Term | Cloud Insights Term
|Cluster|Storage
|Tenant|Storage Pool
|Bucket|Internal Volume
|Disk|Disk
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* An IP address or hostname of the ECS cluster
* A username and password for the ECS system
* Port 4443 (HTTPS). Requires outbound connectivity to TCP port 4443 on
the ECS system.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|ECS Host|IP address or fully-qualified domain name of the ECS system
|ECS Host Port | Port used for communication with ECS Host
|ECS User ID|User ID for ECS
|Password|Password used for ECS
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|Inventory Poll Interval (min)|The default is 360 minutes.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
==== Inventory
[cols=2*, options="header", cols"50,50"]
```

```
|===
|Problem:|Try this:
|Error: User authentication failed.
|Make sure your credentials for this device are correct.
|===
==== Performance
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|Error: No enough data collected.
| Check collection timestamp in log file and modify polling interval
accordingly
• Wait for longer time
|Error: Performance polling interval is too big.
|Check collection timestamp in log file ${logfile} and modify polling
interval accordingly
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[IDbba2e5f5dfde9f33c740bd279bfb711a]]
= Dell EMC PowerScale data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Dell EMC PowerScale (previously Isilon) SSH data
collector to acquire inventory and performance data from PowerScale scale-
out NAS storage.
```

```
== Terminology
Cloud Insights acquires the following inventory information from this data
collector. For each asset type acquired by Cloud Insights, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Drive|Disk
|Cluster|Storage
|Node|Storage Node
|File System|Internal Volume
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
You need the following information to configure this data collector:
* Administrator permissions to the PowerScale storage
* IP address of the PowerScale cluster
* SSH access to port 22
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
| IP address | The IP address or fully-qualified domain name of the
PowerScale cluster
|User Name|User name for the PowerScale cluster
|Password|Password used for the PowerScale cluster
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min) | Interval between inventory polls. The
default is 20.
|Performance Poll Interval (sec)|Interval between performance polls. The
```

```
default is 300.
//|SSH Process Wait Timeout|SSH process timeout period. The default is 60.
|SSH Port|SSH service port. The default is 22.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem: |Try this:
|"Invalid login credentials" with error messages "Commands not enabled for
role-based administration require root user access"
|* Verify that the user has permissions to run the following commands on
the device:
 > isi version osrelease
 > isi status -q
 > isi status -n
 > isi devices -d %s
  > isi license
* Verify credentials used in the wizard are matching device credentials
|"Internal Error" with error messages "Command <Your command> run failed
with permission: <Your current permission>. Sudo command run permission
|Verify that the user has sudo permissions to run the following command on
the device
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID28f30f0d1b6e5b77b425e29d1f672437]]
= Dell EMC Isilon / PowerScale REST data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
```

```
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Dell EMC Isilon / PowerScale REST data collector
to acquire inventory and performance data from Dell EMC Isilon or
PowerScale storage. This collector supports arrays running OneFS 8.0.0+.
== Terminology
Cloud Insights acquires the following inventory information from this data
collector. For each asset type acquired by Cloud Insights, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term|Cloud Insights Term
|Drive|Disk
|Cluster|Storage
|Node|Storage Node
|OneFS File System|Internal Volume
|OneFS File System|Storage Pool
|Qtree|Qtree
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
You need the following information to configure this data collector:
* A user account and password. This account does NOT need to be
admin/root, but you MUST grant a substantial number of read only
privileges to your service account - see table below
* IP address / Fully Qualified Domain Name of the Dell EMC Isilon /
PowerScale cluster
* HTTPS access to port 8080
* Isilon / PowerScale cluster running OneFS 8.0.0 or higher
[cols=3*, options="header", cols"33,33,33"]
|===
|Privilege Name|Description|r(read) or rw (read+write)
| ISI PRIV LOGIN PAPI | Platform API | r
| ISI PRIV SYS TIME | Time | r
|ISI PRIV AUTH|Auth|r
```

```
| ISI PRIV ROLE | Privilege | r
|ISI PRIV DEVICES|Devices|r
| ISI PRIV EVENT| Event| r
|ISI PRIV HDFS|HDFS|r
| ISI PRIV NDMP|NDMP|r
| ISI PRIV NETWORK | Network | r
|ISI PRIV NFS|NFS|r
| ISI PRIV PAPI CONFIG | Configure Platform API | r
| ISI PRIV QUOTA | Quota | r
| ISI PRIV SMARTPOOLS | SmartPools | r
|ISI PRIV SMB|SMB|r
| ISI PRIV STATISTICS | Statistics | r
|ISI PRIV SWIFT|Swift|r
| ISI PRIV JOB ENGINE | Job Engine | r
|===
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|Isilon IP address|The IP address or fully-qualified domain name of the
Isilon storage
|User Name|User name for the Isilon
|Password|Password used for the Isilon
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|HTTPS Port|The default is 8080.
|Inventory Poll Interval (min) | Interval between inventory polls. The
default is 20.
|Performance Poll Interval (sec)|Interval between performance polls. The
default is 300.
//|SSH Process Wait Timeout|SSH process timeout period. The default is 60.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
```

```
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|"Invalid login credentials" with error messages "Commands not enabled for
role-based administration require root user access"
|* Verify that the user has permissions to run the following commands on
the device:
 > isi version osrelease
 > isi status -q
 > isi status -n
 > isi devices -d %s
 > isi license
* Verify credentials used in the wizard are matching device credentials
|"Internal Error" with error messages "Command <Your command> run failed
with permission: <Your current permission>. Sudo command run permission
|Verify that the user has sudo permissions to run the following command on
the device
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID6d3827c2b20cbb0c1b01455c6ea00c3a]]
= Dell EMC PowerStore data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The EMC PowerStore data collector gathers inventory information from EMC
PowerStore storage. For configuration, the data collector requires the IP
address of the storage processors and a read-only user name and password.
The EMC PowerStore data collector gathers the volume-to-volume replication
relationships that PowerStore coordinates across other storage arrays.
Cloud Insights shows a storage array for each PowerStore cluster, and
collects inventory data for nodes and storage ports on that cluster. No
```

```
storage pool or volume data is collected.
== Terminology
Cloud Insights acquires the following inventory information from this data
collector. For each asset type acquired by Cloud Insights, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term|Cloud Insights Term
|host|host
|host volume mapping|host volume mapping
|hardware (it has Drives under "extra details" object): Drives|Disk
|Appliance|StoragePool
|Cluster|Storage Array
|Node|StorageNode
|fc port|Port
|volume|Volume
|InternalVolume|file system
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
The following information is required to configure this data collector:
* IP address or fully-qualified domain name of storage processor
* Read-only user name and password
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|PowerStore gateway(s)|IP addresses or fully-qualified domain names of
PowerStore storage
|User Name|User name for PowerStore
|Password|Password used for PowerStore
|===
== Advanced configuration
```

```
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|HTTPS Port|Default is 443
|Inventory Poll Interval (minutes)|Interval between inventory polls. The
default is 60 minutes.
|===
Cloud Insight's PowerStore performance collection makes use of
PowerStore's 5-minute granularity source data. As such, Cloud Insights
polls for that data every five minutes, and this is not configurable.
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID7e6c65221f5ed468368dbeaf6447b33b]]
= Dell EMC RecoverPoint data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The EMC RecoverPoint data collector's primary use case is to discover
volume-to-volume replication relationships that the RecoverPoint storage
appliance facilitates. This collector will also discover the Recoverpoint
appliance itself. Please note that Dell/EMC sells a VMware backup solution
for VMs--"RecoverPoint for VMs"--which is not supported by this collector
```

For configuration, the data collector requires the IP address of the storage processors and a read-only user name and password.

The EMC RecoverPoint data collector gathers the volume-to-volume replication relationships that RecoverPoint coordinates across other storage arrays. Cloud Insights shows a storage array for each RecoverPoint

```
cluster, and collects inventory data for nodes and storage ports on that
cluster. No storage pool or volume data is collected.
== Requirements
The following information is required to configure this data collector:
* IP address or fully-qualified domain name of storage processor
* Read-only user name and password
* REST API access via port 443
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|Address of RecoverPoint|IP address or fully-qualified domain name of
RecoverPoint cluster
|User Name|User name for the RecoverPoint cluster
|Password|Password used for the RecoverPoint cluster
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|TCP Port|TCP Port used to connect to Recoverpoint cluster
|Inventory Poll Interval (minutes) |Interval between inventory polls. The
default is 20 minutes.
|Excluded Clusters|Comma-separated list of cluster IDs or names to exclude
when polling.
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
```

```
[[ID290615de6b4d768fae49082af2144c0e]]
= DELL EMC ScaleIO / PowerFlex data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The ScaleIO/PowerFlex data collector collects inventory information from
ScaleIO & PowerFlex storage. For configuration, this data collector
requires the ScaleIO/PowerFlex gateway address and an admin user name and
password.
== Terminology
Cloud Insights acquires the following inventory information from the
ScaleIO/PowerFlex data collector. For each asset type acquired by Cloud
Insights, the most common terminology used for this asset is shown. When
viewing or troubleshooting this data collector, keep the following
terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|MDM (Meta Data Manager) Cluster|Storage
|SDS (ScaleIO/PowerFlex Data Server)|Storage Node
|Storage Pool|Storage Pool
|Volume|Volume
|Device|Disk
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* Read-only access to the Admin user account
* Port requirement: HTTPS Port 443
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
```

```
|Field|Description
|ScaleIO/PowerFlex Gateway(s)|IP addresses or FQDNs of ScaleIO/PowerFlex
gateways, separated by comma (,) or semicolon (;)
|User Name|Admin user name used to log in to the ScaleIO/PowerFlex device
|Password|Password used to log in to the ScaleIO/PowerFlex device
== Advanced configuration
Click the Inventory check box to enable inventory collection.
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|HTTPS port|443
|Inventory poll interval (min)|The default is 60.
|Connection Timeout (sec)|The default is 60.
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[IDf9a5a99904924aa497c9dc1452edc334]]
= Configuring the EMC Unity data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The DELL EMC Unity (formerly VNXe) data collector provides inventory
support for VNXe unified storage arrays. Cloud Insights currently supports
iSCSI and NAS protocols.
== Requirements
* The Unity data collector is CLI based; you must install the Unisphere
```

for Unity CLI, (uemcli.exe) onto the acquisition unit where your VNXe data collector resides.

- * uemcli.exe uses HTTPS as the transport protocol, so the acquisition unit will need to be able to initiate HTTPS connections to the Unity.
- * IP address or fully-qualified domain name of the Unity device
- * You must have at least a read-only user for use by the data collector.

//* IP address of the managing Solutions enabler server.

- * HTTPS on Port 443 is required
- * The EMC Unity data collector provides NAS and iSCSI support for inventory; fibre channel volumes will be discovered, but Cloud Insights does not report on FC mapping, masking, or storage ports.

== Terminology

Cloud Insights acquires the following inventory information from the Unity data collector. For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

```
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term|Cloud Insights Term
|Disk|Disk
|Storage Array|Storage
|Processor|Storage Node
|Storage Pool|Storage Pool
|General iSCSI Block info, VMware VMFS|Share
|Replication Remote System|Synchronization
|iSCSI Node|iSCSI Target Node
|iSCSI Initiator|iSCSI Target Initiator
Note: These are common terminology mappings only and might not represent
every case for this data source.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Unity Storage|IP address or fully-qualified domain name of the Unity
device
|User Name |User name for the Unity device
|Password |Password for the Unity device
|Full Path to the Executable UEMCLI|Full path to the folder containing the
```

```
uemcli.exe executable
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min) |Interval between inventory polls. The
default is 40 minutes
|Unity CLI Port | Port used for the Unity CLI
//|Inventory External Process Timeout (sec) |The default is 1800.
|Performance poll interval (sec)|The default is 300.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|"Failed to execute external utility" with error messages "Failed to find
Unisphere executable uemcli"
|* Verify correct IP address, username, and password
* Confirm that Unisphere CLI is installed on the Cloud Insights
Acquisition Unit
* Confirm that Unisphere CLI installation directory is correct in the
datasource configuration
* Confirm that the IP of the VNXe is correct in the configuration of the
datasource. From the Cloud Insights Acquisition Unit, open a CMD and
change to to the configured installation directory: ${INSTALLDIR. Try to
make a connection with the VNXe device by typing: uemcli -d <Your IP> -u
<Your ID> /sys/general show
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[IDe0c9b3293bba10b8a2d196ed90897693]]
= Dell EMC VMAX and PowerMax Family of Devices data collector
```

```
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Cloud Insights discovers EMC VMAX and PowerMax storage arrays by using
```

Cloud Insights discovers EMC VMAX and PowerMax storage arrays by using Solutions Enabler symcli commands in conjunction with an existing Solutions Enabler server in your environment. The existing Solutions Enabler server has connectivity to the VMAX/PowerMax storage array through access to gatekeeper volumes.

== Requirements

Before configuring this data collector, you should ensure that Cloud Insights has TCP connectivity to port 2707 on the existing Solutions Enabler server. Cloud Insights discovers all the Symmetrix arrays that are "Local" to this server, as seen in "symcfg list" output from that server.

- * The EMC Solutions Enabler (CLI) with SMI-S provider application must be installed on the Acquisition Unit server and the version must match or be earlier than the version running on the Solutions Enabler Server.
- * A properly configured {installdir}\EMC\SYMAPI\config\netcnfg file is required. This file defines service names for Solutions Enabler servers, as well as the access method (SECURE / NOSECURE /ANY).
- * If you require read/write latency at the storage node level, the SMI-S Provider must communicate with a running instance of the UNISPHERE for VMAX application.
- * IP address of the managing Solutions Enabler server
- * Administrator permissions on the Solutions Enabler (SE) Server
- * Read-only user name and password to the SE software
- * The UNISPHERE for VMAX application must be running and collecting statistics for the EMC VMAX and PowerMax sstorage arrays that are managed by the SMI-S Provider installation
- * Access validation for performance: In a web browser on your Acquisition Unit, go to _\https://<SMI-S Hostname or IP>:5989/ecomconfig_ where "SMI-S Hostname or IP" is the IP address or hostname of your SMI-S server. This URL is for an administrative portal for the EMC SMI-S (aka "ECOM") service you will receive a login popup.
- * Permissions must be declared in the Solutions Enabler server's daemon configuration file, usually found here: _/var/symapi/config/daemon_users_

```
Here is an example file with the proper cisys permissions.
root@cernciaukc101:/root
14:11:25 # tail /var/symapi/config/daemon users
 ###
       Refer to the storrdfd(3) man page for additional details.
 ###
 ###
        As noted above, only authorized users can perform stordaemon
control
       operations (e.g., shutdown).
 ###
 ######
# smith
               storrdfd
cisys storapid <all>
== Terminology
Cloud Insights acquires the following inventory information from the EMC
VMAX/PowerMax data source. For each asset type acquired, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Disk|Disk
|Disk Group|Disk Group
|Storage|Array Storage
|Director|Storage Node
|Device Pool, Storage Resource Pool (SRP)|Storage Pool
|Device TDev|Volume
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Configuration
*Note:* If SMI-S user authentication is not enabled, the default values in
the Cloud Insights data collector are ignored.
[cols=2*, options="header", cols"50,50"]
|Field|Description
```

```
|Service Name|Service name as specified in netcnfg file
|Full path to CLI|Full path to the folder containing the Symmetrix CLI
|SMI-S Host IP Address| IP address of the SMI-S host
|===
== Advanced Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 40 minutes.
|Choose 'Exclude' or 'Include' to specify a list|Specify whether to
include or exclude the array list below when collecting data.
|Inventory Filter Device List|Comma-separated list of device IDs to
include or exclude
|Connection Caching|Choose connection caching method:
* LOCAL means that the Cloud Insights Acquisition service is running on
the Solutions Enabler server, which has Fibre Channel connectivity to the
Symmetrix arrays you seek to discover, and has access to gatekeeper
volumes. This might be seen in some Remote Acquisition Unit (RAU)
configurations.
* REMOTE CACHED is the default and should be used in most cases. This uses
the NETCNFG file settings to connect using IP to the Solutions Enabler
server, which must have Fibre Channel connectivity to the Symmetrix arrays
you seek to discover, and has access to Gatekeeper volumes.
* In the event that REMOTE CACHED options make CLI commands fail, use the
REMOTE option. Keep in mind that it will slow down the acquisition process
(possibly to hours or even days in extreme cases). The NETCNFG file
settings are still used for an IP connection to the Solutions Enabler
server that has Fibre Channel connectivity to the Symmetrix arrays being
discovered.
*Note: * This setting does not change Cloud Insights behavior with respect
to the arrays listed as REMOTE by the "symcfg list" output. Cloud Insights
gathers data only on devices shown as LOCAL by this command.
|SMI-S Protocol|Protocol used to connect to the SMI-S provider. Also
displays the default port used.
|Override SMIS-Port|If blank, use the default port in the Connection Type
field, otherwise enter the connection port to use
//|CLI Timeout (sec)|CLI process timeout (default 7200 seconds)
//|SMI-S Host IP|IP address of the SMI-S Provider Host
//|SMI-S Port|Port used by SMI-S Provider Host
```

//|SMI-S Namespace|Interoperability namespace that the SMI-S provider is configured to use |SMI-S User Name|User name for the SMI-S Provider Host |SMI-S Password|User name for the SMI-S Provider Host |Performance Polling Interval (sec) |Interval between performance polls (default 1000 seconds) |hoose 'Exclude' or 'Include' to specify a list| Specify whether to include or exclude the array list below when collecting performance data |Performance Filter Device List|Comma-separated list of device IDs to include or exclude //|RPO Polling Interval (sec)|Interval between RPO polls (default 300 seconds) |=== == Troubleshooting Some things to try if you encounter problems with this data collector: [cols=2*, options="header", cols"50,50"] |=== |Problem:|Try this: |Error: The feature being requested is not currently licensed |Install the SYMAPI server license. |Error: No devices were found |Make sure Symmetrix devices are configured to be managed by the the Solutions Enabler server: - Run symcfq list -v to see the list of configured Symmetrix devices. |Error: A requested network service was not found in the service file |Make sure the Solutions Enabler Service Name is defined the netcnfg file for Solutions Enabler. This file is usually located under SYMAPI\config\ in the Solutions Enabler client installation. |Error: The remote client/server handshake failed |Check the most recent storsrvd.log* files on the Solutions Enabler host we are trying to discover. |Error: Common name in client certificate not valid |Edit the hosts file on the Solutions Enabler server so that the Acquisition Unit's hostname resolves to the IP address as reported in the storsrvd.log on the Solutions Enabler server. |Error: The function could not obtain memory |Make sure there is enough free memory available in the system to execute Solutions Enabler |Error: Solutions Enabler was unable to serve all data required. |Investigate the health status and load profile of Solutions Enabler |Error: • The "symcfg list -tdev" CLI command may return incorrect data when collected with Solutions Enabler 7.x from a Solutions Enabler server 8.x.

• The "symcfg list -srp" CLI command may return incorrect data when collected with Solutions Enabler 8.1.0 or earlier from a Solutions Enabler server 8.3 or later.

|Be sure you are using the same Solutions Enabler major release |I'm seeing data collection errors with the message: "unknown code" |You may see this message if permissions are not declared in the Solutions Enabler server's daemon configuration file (see the <<requirements, Requirements>> above.) This assumes your SE client version matches your SE server version.

This error may also occur if the _cisys_ user (which executes Solutions Enabler commands) has not been configured with the necessary daemon permissions in the /var/symapi/config/daemon users configuration file.

To fix this, edit the /var/symapi/config/daemon_users file and make sure the cisys user has <all> permission specified for the storapid daemon.

Example:

```
14:11:25 # tail /var/symapi/config/daemon_users
...
cisys storapid <all>
```

Additional information may be found from the xref:{relative_path}concept_requesting_support.html[Support] page or in the xref:{relative_path}reference_data_collector_support_matrix.html[Data Collector Support Matrix].

```
[[ID38dle362aecffc2236f0d782a408c534]]
= Dell EMC VNX Block Storage (NaviCLI) data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
Cloud Insights uses the Dell EMC VNX Block Storage (NaviSec) data
collector (formerly CLARiiON) to acquire inventory and performance data.
```

== Terminology Cloud Insights acquires the following inventory information from the EMC VNX Block Storage data collector. For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind: [cols=2*, options="header", cols"50,50"] |=== |Vendor/Model Term|Cloud Insights Term |Disk|Disk |Storage|Storage |Storage Processor|Storage Node |This Pool, RAID Group|Storage Pool |LUN|Volume |=== Note: These are common terminology mappings only and might not represent every case for this data source. == Requirements The following requirements must be met in order to collect data: * An IP address of each VNX block storage processor * Read-only Navisphere username and password to the VNX block storage arrays * NaviSecCli must be installed on the Cloud Insights AU * Access validation: Run NaviSecCLI from the Cloud Insights AU to each array using the username and password. * Port requirements: 80, 443 * NaviSecCLI version should correspond with the newest FLARE code on your * For performance, statistics logging must be enabled. == NaviSphere command line interface syntax naviseccli.exe -h <IP address> -user <user> -password <password> -scope <scope,use 0 for global scope> -port <use 443 by default> command == Configuration [cols=2*, options="header", cols"50,50"]

|VNX Block Storage IP Address|IP address or fully-qualified domain name of

|Field |Description

```
the VNX block storage
|User Name |Name used to log into the VNX block storage device.
|Password|Password used to log into the VNX block storage device.
|CLI Path to naviseccli.exe|Full path to the folder containing the
naviseccli.exe executable
|===
== Advanced Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field |Description
|Inventory Poll Interval (min)|Interval between inventory polls. Default
is 40 minutes.
//|Use Secure Client |Select to use the navseccl secure client (i)
|Scope|The secure client scope. The default is Global.
//|VNX block storage CLI Port|Port used for VNX block storage CLI. The
default is 443.
//|Inventory External Process Timeout (sec)|External process timeout. The
default is 1800 seconds.
|Performance Poll Interval (sec)|Interval between performance polls. The
default is 300 seconds.
//|Performance External process timeout (sec)|External process timeout.
The default is 1800 seconds.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|Problem:|Try this:
|Error:
• Agent Not Running
  Failed to find naviseccli
• Failed to execute any command
| Confirm that NaviSphere CLI is installed on the Cloud Insight
Acquisition Unit
   You have not selected the "Use secure client" option in the data
collector configuration wizard and do not have a non-secure version of
Naviphere CLI installed.
• Confirm that NaviSphere CLI installation directory is correct in the
```

data collector configuration

- Confirm that the IP of the VNX block storage is correct in the data collector configuration:
- From the Cloud Insights Acquisition Unit:
 - Open a CMD.
 - Change the directory to the configured installation directory
- Try to make a connection with the VNX block storage device by typing "navicli -h {ip} getagent" (replace the {ip} with the actual IP) | Error: 4.29 emc235848 emc241018 getall Failed to parse host alias info | This is likely caused by a FLARE 29 corruption issue of the host initiator database on the array itself. See EMC knowledge base articles: emc235848, emc241018. You can also check

https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb58128 |Error: Unable to retrieve Meta LUNs. Error Executing java -jar navicli.jar

- | Modify the data collector configuration to use the secure client (recommended)
- Install navicli.jar in the CLI path to navicli.exe OR naviseccli.exe
- Note: navicli.jar is deprecated as of EMC Navisphere version 6.26
- The navicli.jar may be available on http://powerlink.emc.com

|Error: Storage Pools not reporting disks on Service Processor at configured IP address

|Configure the data collector with both Service Processor IPs, separated by a comma

|Error: Revision mismatch error

- |• This is usually caused by updating the firmware on the VNX block storage device, but not updating the installation of NaviCLI.exe. This also might be caused by having different devices with different firmwares, but only one CLI installed (with a different firmware version).
- Verify that the device and the host are both running identical versions of the software:
- From the Cloud Insights Acquisition Unit, open a command line window
 - Change the directory to the configured installation directory
- Make a connection with the CLARiiON device by typing "navicli -h \${ip} getagent"
- Look for the version number on the first couple of lines. Example: "Agent Rev: 6.16.2 (0.1)"
- Look for and compare the version on the first line. Example: "Navisphere CLI Revision 6.07.00.04.07"

|Error: Unsupported Configuration - No Fibre Channel Ports |The device is not configured with any Fibre Channel ports. Currently, only FC configurations are supported. Verify this version/firmware is supported.

|===

Additional information may be found from the

```
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative_path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID0082ee8aa2954403272e9f7a829e815c]]
= DELL EMC VNX File (formerly Celerra Unified Storage System) data
collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector acquires inventory information from the VNX File
Storage System. For configuration, this data collector requires the IP
address of the storage processors and a read-only user name and password.
== Terminology
Cloud Insights acquires the following inventory information from the VNX
File data collector. For each asset type acquired by Cloud Insights, the
most common terminology used for this asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term|Cloud Insights Term
|Celerra Network Server/Celerra Storage Pool|Storage Pool
|File System|Internal Volume
|Data Mover|Controller
|File System mounted on a data mover|File Share
|CIFS and NFS Exports|Share
|Disk Volume|Backend LUN
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
```

```
You need the following to configure this data collector:
* The IP address of the storage processor
* Read-only user name and password
* SSH port 22
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|VNX File IP Address|IP address or fully-qualified domain name of the VNX
File device
|User Name|Name used to log in to the VNX File device
|Password|Password used to log in to the VNX File device
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (minutes)|Interval between inventory polls. The
default is 20 minutes.
//|SSH Process Wait Timeout (sec)|SSH process timeout. The default is 600
seconds.
//|Number of Retries|Number of inventory retry attempts. The default is 2.
//|SSH Banner Wait Timeout (sec)|SSH banner wait timeout. The default is
20 seconds.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|Error: Unable to proceed while DART update in progress
| Possible solution: Pause the data collector and wait for the DART upgrade
to complete before attempting another acquisition request.
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
```

```
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID8d4dd5fd608263a9839d4dc9a15d6688]]
= Configuring the Dell EMC VNX Unified data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
For configuration, the Dell EMC VNX Unified (SSH) data collector requires
the IP address of the Control Station and a read-only username and
password.
== Terminology
Cloud Insights acquires the following inventory information from this data
collector. For each asset type acquired by Cloud Insights, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Disk|Disk
|Disk Folder|Disk Group
|File system|Internal Volume
|Storage|Storage
|Storage Processor|Storage Node
|Storage Pool, RAID Group|Storage Pool
|LUN|Volume
|Data Mover|Controller
|File System mounted on a data mover|File Share
|CIFS and NFS Exports|Share
|Disk Volume|Backend LUN
|===
== Requirements
```

```
You need the following to configure the VNX (SSH) data collector:
* VNX IP address & Credentials to the Celerra Control Station.
* Read-only username and password.
* The data collector is able to run NaviCLI/NaviSecCLI commands against
the backend array utilizing the DART OS NAS heads
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|VNX IP Address|IP address or fully-qualified domain name of the VNX
Control Station
|User Name |User name for the VNX Control Station
|Password |Password for the VNX Control Station
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 40 minutes.
//|VNX SSH Process Wait Timeout|VNX SSH process timeout (default 600
seconds)
//|Celerra Command Retry Attempts|Number of Celerra command retry attempts
//|CLARiiON External Process Timeout for Inventory (sec)| CLARiiON
external process timeout for inventory. The default is 1800 seconds.)
|Performance Poll Interval (sec).|Interval between performance polls. The
default is 300 seconds.
//|CLARiiON External Process Timeout for Performance (sec).|CLARiiON
external process timeout for performance. The default is 1800 seconds.
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
```

```
[[IDb1117375addbb8bf3b774369f365545f]]
= Configuring the EMC VPLEX data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector acquires inventory and performance data from EMC VPLEX
storage systems. For configuration, the data collector requires an IP
address of the VPLEX server and an administrative level domain account.
NOTE: Cloud Insights' performance collection from Vplex clusters requires
that the performance archive service be operational, in order to populate
the .CSV files and logs that Cloud Insights retrieves via SCP-based file
copies. NetApp has observed that many Vplex firmware upgrade/management
station updates will leave this functionality non-operational. Customers
planning such upgrades may want to proactively ask Dell/EMC if their
planned upgrade will leave this functionality inoperable, and if so, how
can they re-enable it to minimize gaps in performance visibility? Cloud
Insight's Vplex performance code will assess on each poll whether all the
expected files exist, and if they are being properly updated; if they are
missing or stale, Cloud Insights will log performance collection failures.
== Terminology
Cloud Insightst acquires the following inventory information from the
VPLEX data collector. For each asset type acquired, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Cluster|Storage
|Engine|Storage Node
|Device, System Extent|Backend Storage Pool
```

|Virtual Volume|Volume

```
|Front-End Port, Back-End Port|Port
|Distributed Device|Storage Synchronization
|Storage View|Volume Map, Volume Mask
|Storage Volume|Backend LUN
|ITLs|Backend Path
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* An IP address of the VPLEX Management Console
* Administrative level domain account for the VPLEX server
* Port 443 (HTTPS). Requires outbound connectivity to TCP port 443 on the
VPLEX management station.
* For performance, read-only username and password for ssh/scp access.
* For performance, port 22 is required.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
| IP address of VPLEX Management Console | IP address or fully-qualified
domain name of the VPLEX Management Console
|User Name|User name for VPLEX CLI
|Password|Password used for VPLEX CLI
| Performance Remote IP Address | Performance Remote IP address of the VPLEX
Management Console
|Performance Remote User Name|Performance Remote user name of VPLEX
Management Console
|Performance Remote Password|Performance Remote Password of VPLEX
Management Console
| ====
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Communication Port|Port used for VPLEX CLI. The default is 443.
|Inventory Poll Interval (min)|The default is 20 minutes.
//|Connection timeout (sec)|The default is 60 seconds.
```

|Number of connection retries|The default is 3. |Performance Poll Interval (sec)|Interval between performance polls. The default is 600 seconds. //|Performance SSH Process Wait Timeout (sec).| SSH process timeout. The default is 600 seconds. //|SSH Banner Wait Timeout (sec).|The default is 20 seconds. |Number of Retries|The default is 2. |=== == Troubleshooting Some things to try if you encounter problems with this data collector: === Inventory [cols=2*, options="header", cols"50,50"] |=== |Problem: |Try this: |Error: User authentication failed. |Make sure your credentials for this device are correct. |=== === Performance [cols=2*, options="header", cols"50,50"] |=== |Problem:|Try this: |Error: VPLEX performance for version below 5.3 is not supported. |Upgrade VPLEX to 5.3 or above |Error: No enough data collected. | Check collection timestamp in log file and modify polling interval accordingly • Wait for longer time |Error: Perpetual Log files not being updated. |Please contact EMC support to enable updating the perpetual log files |Error: Performance polling interval is too big. |Check collection timestamp in log file \${logfile} and modify polling interval accordingly |Error: Performance Remote IP address of VPLEX Management Console is not configured. |Edit the data source to set Performance Remote IP address of VPLEX Management Console. |Error: No performance data reported from director | Check that the system performance monitors are running correctly · Please contact EMC support to enable updating the system performance monitor log files

```
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID2e1376a470dfdbdf1fda21f911dec5e9]]
= Dell EMC XtremeIO data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The EMC XtremIO data collector acquires inventory and performance data
from the EMC XtremIO storage system.
== Requirements
To configure the EMC XtremIO (HTTP) data collector, you must have:
* The XtremIO Management Server (XMS) Host address
* An account with administrator privileges
* Access to port 443 (HTTPS)
== Terminology
Cloud Insights acquires the following inventory information from the EMC
XtremIO data collector. For each asset type acquired by Cloud Insights,
the most common terminology used for this asset is shown. When viewing or
troubleshooting this data source, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Disk (SSD)|Disk
|Cluster|Storage
|Controller|Storage Node
|Volume|Volume
|LUN Map|Volume Map
```

```
|Target FC Initiator|Volume Mask
|===
Note: These are common terminology mappings only and might not represent
every case for this data source.
== Requirements
* The XtremIO Management Server (XMS) Host IP address
* Administrator user name and password for the XtremIO
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|XMS Host|IP address or fully-qualified domain name of the XtremIO
Management Server
|User name|User name for the XtremIO Management Server
|Password|Password for the XtremIO Management Server
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|TCP port|TCP Port used to connect to XTremIO Management Server. The
default is 443.
|Inventory poll interval (min)|Interval between inventory polls. The
default is 60 minutes.
//|Connection timeout (sec)|Connection timeout The default is 60 seconds.
|Performance poll interval (sec)|Interval between performance polls. The
default is 300 seconds.
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
```

```
:leveloffset: -1
[[ID394489cdae380e5e9493d6bdd0840e4a]]
= Fujitsu Eternus data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
The Fujitsu Eternus data collector acquires inventory data using
administration-level access to the storage system.
== Terminology
Cloud Insights acquires the following inventory information from the
Fujitsu Eternus storage. For each asset type acquired by Cloud Insights,
the most common terminology used for this asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term|Cloud Insights Term
|Disk|Disk
|Storage|Storage
|Thin Pool, Flexible Tier Pool, Raid Group|Storage Pool
|Standard Volume, Snap Data Volume (SDV), Snap Data Pool Volume (SDPV),
Thin Provisioning Volume (TPV), Flexible Tier Volume (FTV), Wide Striping
Volume (WSV) | Volume
|Channel adapter|Controller
|===
Note: These are common terminology mappings only and might not represent
every case for this data collectior.
== Requirements
The following are required to configure this data collector:
* An IP address of the Eternus storage, which cannot be comma delimited
* SSH Administration-level user name and password
* Port 22
```

```
* Ensure that the page scroll is disabled (clienv-show-more-scroll
disable)
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
| IP Address of Eternus Storage | IP address of the Eternus storage
|User Name|User name for Eternus storage
|Password|Password for the Eternus storage
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|The default is 20 minutes.
//|SSH Process Wait Timeout (sec)|SSH process timeout. The default is
600 seconds.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|"Error retrieving data" with error messages "Error Finding Prompt CLI" or
"Error finding prompt at the end of shell results"
|Likely caused by: Storage system has page scrolling enabled.
Possible solution:
* Try to disable page scrolling by running the following command:
set clienv-show-more -scroll disable
|"Connecting error" with error messages "Failed to instantiate an SSH
connection to storage" or "Failed to instantiate a connection to
VirtualCenter"
|Likely causes:
* Incorrect credentials.
* Incorrect IP address.
* Network problem.
* Storage may be down or unresponsive.
```

```
Possible solutions:
* Verify credentials and IP address entered.
* Try to communicate with storage using SSH Client.
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[IDa111551265852bdb8c31e849685b5f56]]
= NetApp Google Compute data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector supports inventory and performance collection from
Google Compute cloud platform configurations. This collector will seek to
discover all the Compute resources within all the Projects within one
Google organization. If you have multiple Google organizations you want to
discover with Cloud Insights, you will want to deploy one Cloud Insights
collector per organization.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Organization ID|The organization ID you want to discover with this
collector. This field is required if your service account is able to see
more than one organization
|Choose 'Exclude' or 'Include' to filter GCP Projects by IDs| If you want
to limit what projects' resources are brought into Cloud Insights.
|Project IDs |The list of Project IDs that you want to filter in, or out
from discovery, depending on the value of the "Choose 'Exclude"...."
value. The default list is empty
|Client ID |Client ID for the Google Cloud Platform configuration
|Copy and paste the contents of your Google Credential File here|Copy your
```

```
Google credentials for the Cloud Platform account to this field
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min) |Default is 60 minutes
|Choose 'Exclude' or 'Include' to Apply to Filter VMs by Labels|Specify
whether to include or exclude VM's by Labels when collecting data. If
'Include' is selected, the Label Key field can not be empty.
|Label Keys and Values on which to Filter VMs|Click *+ Filter Label* to
choose which VMs (and associated disks) to include/exclude by filtering
for keys and values that match keys and values of labels on the VM. Label
Key is required, Label Value is optional. When Label Value is empty, the
VM is filtered as long as it matches the Label Key.
|Performance Poll Interval (sec)|Default is 1800 seconds
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
= HP Enterprise
:leveloffset: +1
[[ID8fd6376d4fada350187f8cd06733511a]]
= HP Enterprise Alletra 9000 / Primera Storage data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
```

```
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the HP Enterprise Alletra 9000 / HP Enterprise Primera
(previously 3PAR) data collector to discover inventory and performance.
== Terminology
Cloud Insights acquires the following inventory information from this data
collector. For each asset type acquired by Cloud Insights, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Field|Description
|Physical Disk|Disk
|Storage System|Storage
|Controller Node|Storage Node
|Common Provisioning Group|Storage Pool
|Virtual Volume|Volume
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
The following are required to configure this data colletor:
* IP address or FQDN of the InServ cluster
* For inventory, read-only user name and password to the StoreServ Server
* For performance, read-write user name and password to the StoreServ
Server
* Port requirements: 22 (inventory collection), 5988 or 5989 (performance
collection) [Note: Performance is supported for StoreServ OS 3.x+]
* For performance collection confirm that SMI-S is enabled by logging into
the array via SSH.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Storage IP address|Storage IP address or fully-qualified domain name of
the StoreServ cluster
```

```
|User Name|User name for the StoreServ Server
|Password|Password used for the StoreServ Server
//|SMI-S Host IP address|IP address of the SMI-S Provider Host
|SMI-S User Name|User name for the SMI-S Provider Host
|SMI-S Password|Password used for the SMI-S Provider Host
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 40 minutes.
//|Excluded Devices|Comma-separated list of device IPs to exclude
//|SSH Process Wait Timeout (sec)|SSH process timeout. The default is 60
seconds.
//|Number of SSH Retries|Number of SSH retry attempts
//|SSH Banner Wait Timeout (sec)|SSH banner wait timeout. The default is
|SMI-S Connectivity|Protocol used to connect to the SMI-S provider
|Override SMI-S Default Port|If blank, use the default port from SMI-S
Connectivity, otherwise enter the connection port to use
//|SMI-S Password|Password used for the SMI-S Provider Host
//|SMI-S namespace|SMI-S namespace. The default path is root/PG InterOp.
|Performance Poll Interval (sec)|Interval between performance polls. The
default is 300 seconds.
//|Number of SMI-S Connection Retries|Number of SMI-S connection retry
attempts
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
////
error: "Cache server is waiting for the system manager"
Customer can take action. What can customer do about this scenario?
////
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|"showsys" command doesn't return any result.
|Run "showsys" and "showversion -a" from the command line and check if the
```

```
version is supported by the array.
|===
=== Performance
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
| Failed to connect or login. Provider initialization failed.
|An all-numeric array name can cause problems with SMI-S server. Try
changing the array name.
|SMI-S user configured does not have any domain
|Grant appropriate domain privileges to the configured SMI-S user
|Cloud Insights states that it cannot connect/login to SMI-S service.
|Confirm there is no firewall between the CI AU and the array that would
block the CI AU from making TCP connections to 5988 or 5989.
Once that is done, and if you have confirmed there is no firewall, you
should SSH to the array, and use the "showcim" command to confirm.
Verify that:
* Service is enabled
* HTTPS is enabled
* HTTPS port should be 5989
If those all are so, you can try to "stopcim" and then a "startcim" to
restart the CIM (i.e. SMI-S service).
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID2a209f90c9deb914343890a6cef04307]]
= HP Enterprise Command View data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
```

```
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
The HP Enterprise Command View Advanced Edition data collector supports
discovering XP and P9500 arrays via Command View Advanced Edition (CVAE)
server. Cloud Insights communicates with CVAE using the standard Command
View API to collect inventory and performance data.
== Terminology
Cloud Insights acquires the following inventory information from the HP
Enterprise Command View data collector. For each asset type acquired by
Cloud Insights, the most common terminology used for this asset is shown.
When viewing or troubleshooting this data collector, keep the following
terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term|Cloud Insights Term
|PDEV|Disk
|Journal Pool|Disk Group
|Storage Array|Storage
|Port Controller|Storage Node
|Array Group, DP Pool|Storage Pool
|Logical Unit, LDEV|Volume
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Inventory requirements
You must have the following in order to collect inventory data:
* IP address of the CVAE server
* Read-only user name and password for the CVAE software and peer
privileges
* Port requirement: 2001
== Performance requirements
The following requirements must be met in order to collect performance
data:
* HDS USP, USP V, and VSP performance
```

- ** Performance Monitor must be licensed.
- ** Monitoring switch must be enabled.
- ** The Export Tool (Export.exe) must be copied to the Cloud Insights AU and extracted to a location. On CI Linux AUs, ensure "cisys" has read and execute permissions.
- ** The Export Tool version must match the microcode version of the target array.
- * AMS performance:
- ** Performance Monitor must be licensed.
- ** The Storage Navigator Modular 2 (SNM2) CLI utility be installed on the Cloud Insights AU.
- * Network requirements
- ** The Export Tools are Java based, and use RMI to speak to the array. These tools may not be firewall-friendly as they may dynamically negotiate source and destination TCP ports on each invocation. Also, different model array's Export Tools may behave differently across the network consult HPE for your model's requirements
- == Configuration

[cols=2*, options="header", cols"50,50"]

|Field|Description

|Command View Server |IP address or fully-qualified domain name of the Command View server

|User Name |User name for the Command View server.

|Password|Password used for the Command View server.

|Devices - VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages |Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires:

- * Array's IP: IP address of the storage
- * User Name: User name for the storage
- * Password: Password for the storage
- * Folder Containing Export Utility JAR Files

|SNM2Devices - WMS/SMS/AMS Storages|Device list for WMS/SMS/AMS storages. Each storage requires:

- * Array's IP: IP address of the storage
- * Storage Navigator CLI Path: SNM2 CLI path
- * Account Authentication Valid: Select to choose valid account authentication
- * User Name: User name for the storage

* Password: Password for the storage |Choose Tuning Manager for Performance |Override other performance options |Tuning Manager Host|IP address or fully-qualified domain name of tuning manager |Tuning Manager Port|Port used for Tuning Manager |Tuning Manager Username|User name for Tuning Manager |Tuning Manager Password|Password for Tuning Manager Note: In HDS USP, USP V, and VSP, any disk can belong to more than one array group. == Advanced configuration |=== |Field|Description |Command View Server Port |Port used for the Command View Server |HTTPs Enabled|Select to enable HTTPs |Inventory Poll Interval (min) | Interval between inventory polls. The default is 40. |Choose 'Exclude' or 'Include' to specify a list|Specify whether to include or exclude the array list below when collecting data. |Exclude or Include Devices|Comma-separated list of device ID's or array names to include or exclude |Query Host Manager|Select to query host manager //|HTTP Timeout (sec)|HTTP connection timeout The default is 60. |Performance Polling Interval (sec) | Interval between performance polls. The default is 300. //|Export timeout in seconds|Export utility timeout. The default is 300. |=== == Troubleshooting Some things to try if you encounter problems with this data collector: === Inventory [cols=2*, options="header", cols"50,50"] |Problem:|Try this: |Error: User does not have enough permission |Use a different user account that has more privilege or increase the privilege of user account configured in the data collector |Error: Storages list is empty. Either devices are not configured or the user does not have enough permission |* Use DeviceManager to check if the devices are configured. * Use a different user account that has more privilege, or increase the

```
privilege of the user account
|Error: HDS storage array was not refreshed for some days
|Investigate why this array is not being refreshed in HP CommandView AE.
|===
=== Performance
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|Error:
```

- * Error executing export utility
- * Error executing external command
- $\mid *$ Confirm that Export Utility is installed on the Cloud Insights Acquisition Unit
- * Confirm that Export Utility location is correct in the data collector configuration
- * Confirm that the IP of the USP/R600 array is correct in the configuration of the data collector
- \star Confirm that the User name and password are correct in the configuration of the data collector
- * Confirm that Export Utility version is compatible with storage array micro code version
- * From the Cloud Insights Acquisition Unit, open a CMD prompt and do the following:
- Change the directory to the configured installation directory
- Try to make a connection with the configured storage array by executing batch file runWin.bat

|Error: Export tool login failed for target IP

- |* Confirm that username/password is correct
- * Create a user ID mainly for this HDS data collector
- * Confirm that no other data collectors are configured to acquire this array

|Error: Export tools logged "Unable to get time range for monitoring".

- |* Confirm performance monitoring is enabled on the array.
- * Try invoking the export tools outside of Cloud Insights to confirm the problem lies outside of Cloud Insights.

|Error:

- * Configuration error: Storage Array not supported by Export Utility
- * Configuration error: Storage Array not supported by Storage Navigator Modular CLI
- |* Configure only supported storage arrays.
- * Use "Filter Device List" to exclude unsupported storage arrays.

|Error:

- * Error executing external command
- * Configuration error: Storage Array not reported by Inventory
- * Configuration error:export folder does not contains jar files

- |* Check Export utility location.
- * Check if Storage Array in question is configured in Command View server
- * Set Performance poll interval as multiple of 60 seconds.

|Error:

- * Error Storage navigator CLI
- * Error executing auperform command
- * Error executing external command
- |* Confirm that Storage Navigator Modular CLI is installed on the Cloud Insights Acquisition Unit
- * Confirm that Storage Navigator Modular CLI location is correct in the data collector configuration
- * Confirm that the IP of the WMS/SMS/SMS array is correct in the configuration of the data collector
- * Confirm that Storage Navigator Modular CLI version is compatible with micro code version of storage array configured in the data collector
- * From the Cloud Insights Acquisition Unit, open a CMD prompt and do the following:
- Change the directory to the configured installation directory
- Try to make a connection with the configured storage array by executing following command "auunitref.exe"

|Error: Configuration error: Storage Array not reported by Inventory |Check if Storage Array in question is configured in Command View server | Error:

- * No Array is registered with the Storage Navigator Modular 2 CLI
- * Array is not registered with the Storage Navigator Modular 2 CLI
- * Configuration error: Storage Array not registered with StorageNavigator Modular CLI
- |* Open Command prompt and change directory to the configured path
- * Run the command "set=STONAVM HOME=."
- * Run the command "auunitref"
- * Confirm that the command output contains details of the array with IP
- * If the output does not contain the array details then register the array with Storage Navigator CLI:
 - Open Command prompt and change directory to the configured path
 - Run the command "set=STONAVM HOME=."
- Run command "auunitaddauto -ip ${ip}''$. Replace ${ip}$ with real IP ===

Additional information may be found from the xref:{relative_path}concept_requesting_support.html[Support] page or in the xref:{relative_path}reference_data_collector_support_matrix.html[Data Collector Support Matrix].

[[IDa1c3c01da6ffbbb44ff71aa9832142a3]]

```
= HPE Alletra 6000 data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The HP Enterprise Alletra 6000 (previously Nimble) data collector supports
inventory and performance data for Alletra 6000 storage arrays.
== Terminology
Cloud Insights acquires the following inventory information from this
collector. For each asset type acquired by Cloud Insights, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Array|Storage
|Disk|Disk
|Volume|Volume
|Pool|Storage Pool
|Initiator|Storage Host Alias
|Controller|Storage Node
|Fibre Channel Interface|Controller
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
You must have the following in order to collect inventory and
configuration data from the storage array:
* The array must be installed and configured, and reachable from the
client through its fully qualified domain name (FQDN) or array management
IP address.
* The array must be running NimbleOS 2.3.x or later.
* You must have a valid user name and password to the array with at least
"Operator" level role. The "Guest" role does not have sufficient access to
```

```
understand initiator configurations.
* Port 5392 must be open on the array.
You must have the following in order to collect performance data from the
storage array:
* The array must be running NimbleOS 4.0.0 or later
* The array must have volumes configured. The only performance API
NimbleOS has is for volumes, and any statistics Cloud Insights reports are
derived from the statistics on volumes
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|Array Management IP Address|Fully qualified domain name (FQDN) or array
management IP address.
|User Name|User name for the array
|Password|Password for the array
|===
== Advanced configuration
|===
|Field|Description
|Port|Port used by Nimble REST API. The default is 5392.
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 60 minutes.
|===
Note: The default performance poll interval is 300 seconds and can not be
changed. This is the only interval supported by HPE Alletra 6000.
:leveloffset: -1
= Hitachi Data Systems
:leveloffset: +1
[[IDe0b5755cbc12b6116f116f82ee4abe9b]]
```

```
= Hitachi Vantara Command Suite data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The Hitachi Vantara Command Suite data collector supports the HiCommand
Device Manager server. Cloud Insights communicates with the HiCommand
Device Manager server using the standard HiCommand API.
== Terminology
Cloud Insights acquires the following inventory information from the
Hitachi Vantara Command Suite data collector. For each asset type acquired
by Cloud Insights, the most common terminology used for this asset is
shown. When viewing or troubleshooting this data collector, keep the
following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|PDEV|Disk
|Journal Pool|Disk Group
|Storage Array|Storage
|Port Controller|Storage Node
|Array Group, HDS Pool|Storage Pool
|Logical Unit, LDEV|Volume
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
=== Storage
The following terms apply to objects or references that you might find on
HDS storage asset landing pages. Many of these terms apply to other data
collectors as well.
* Name - comes directly from HDS HiCommand Device Manager's "name"
attribute via the GetStorageArray XML API call
* Model - comes directly from HDS HiCommand Device Manager's "arrayType"
attribute via the GetStorageArray XML API call
```

- * Vendor HDS
- * Family comes directly from HDS HiCommand Device Manager's "arrayFamily" attribute via the GetStorageArray XML API call
- \star IP this is the management IP address of the array, not an exhaustive list of all IP addresses on the array
- * Raw Capacity a base2 value representing the sum of the total capacity of all disks in this system, regardless of disk role.

=== Storage Pool

The following terms apply to objects or references that you might find on HDS storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- * Type: The value here will be one of:
- ** RESERVED if this pool is dedicated for purposes other than data volumes, i.e, journaling, snapshots
- ** Thin Provisioning if this is a HDP pool
- ** Raid Group you will not likely see these for a few reasons:

Cloud Insights takes a strong stance to avoid double counting capacity at all costs. On HDS, one typically needs to build Raid Groups from disks, create pool volumes on those Raid Groups, and construct pools (often HDP, but could be special purpose) from those pool volumes. If Cloud Insights reported both the underlying Raid Groups as is, as well as the Pools, the sum of their raw capacity would vastly exceed the sum of the disks.

Instead, Cloud Insights' HDS Command Suite data collector arbitrarily shrinks the size of Raid Groups by the capacity of pool volumes. This may result in Cloud Insights not reporting the Raid Group at all. Additionally, any resulting Raid Groups are flagged in a way such that they are not visible in the Cloud Insights WebUI, but they do flow into the Cloud Insights Data Warehouse (DWH). The purpose of these decisions is to avoid UI clutter for things that most users do not care about — if your

HDS array has Raid Groups with 50MB free, you probably cannot use that

* Node - N/A, as HDS pools are not tied to any one specific node

free space for any meaningful outcome.

- * Redundancy the RAID level of the pool. Possibly multiple values for a HDP pool comprised of multiple RAID types
- * Capacity % the percent used of the pool for data usage, with the used GB and total logical GB size of the pool
- * Over-committed Capacity a derived value, stating "the logical capacity of this pool is oversubscribed by this percentage by virtue of the sum of the logical volumes exceeding the logical capacity of the pool by this percentage"

- * Snapshot shows the capacity reserved for snapshot usage on this pool
- === Storage Node

The following terms apply to objects or references that you might find on HDS storage node asset landing pages. Many of these terms apply to other data collectors as well.

- * Name The name of the Front-end director (FED) or Channel Adapter on monolithic arrays, or the name of the controller on a modular array. A given HDS array will have 2 or more Storage Nodes
- * Volumes The Volume table will show any volume mapped to any port owned by this storage node
- == Inventory Requirements

You must have the following in order to collect inventory data:

- * IP address of the HiCommand Device Manager server
- * Read-only user name and password for the HiCommand Device Manager software and peer privileges
- * Port requirements: 2001 (http) or 2443 (https)
- * Log into HiCommand Device Manager software using username and password
- * Verify access to HiCommand Device Manager

http://<HiCommand_Device_Manager_IP>:2001/service/StorageManager

== Performance requirements

The following requirements must be met in order to collect performance data:

- * HDS USP, USP V, and VSP performance
- ** Performance Monitor must be licensed.
- ** Monitoring switch must be enabled.
- ** The Export Tool (Export.exe) must be copied to the Cloud Insights AU.
- ** The Export Tool version must match the microcode version of the target array.
- * AMS performance:
- ** NetApp strongly recommends creating a dedicated service account on AMS arrays for Cloud Insights to use to retrieve performance data. Storage Navigator only allows a user account one concurrent login to the array. Having Cloud Insights use the same user account as management scripts or HiCommand may result in Cloud Insights, management scripts, or HiCommand being unable to communicate to the array due to the one concurrent user account login limit

- ** Performance Monitor must be licensed.
- ** The Storage Navigator Modular 2 (SNM2) CLI utility needs to be installed on the Cloud Insights ${\tt AU}$.

== Configuration

[cols=2*, options="header", cols"50,50"]

|Field|Description

 $| \mbox{ HiCommand Server } | \mbox{ IP address or fully-qualified domain name of the } \\ \mbox{ HiCommand Device Manager server} \\$

|User Name |User name for the HiCommand Device Manager server. |Password|Password used for the HiCommand Device Manager server. |Devices - VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages |Device list for VSP G1000 (R800), VSP (R700), HUS VM (HM700) and USP storages. Each storage requires:

- * Array's IP: IP address of the storage
- * User Name: User name for the storage
- * Password: Password for the storage
- * Folder Containing Export Utility JAR Files

|SNM2Devices - WMS/SMS/AMS Storages|Device list for WMS/SMS/AMS storages. Each storage requires:

- * Array's IP: IP address of the storage
- * Storage Navigator CLI Path: SNM2 CLI path
- * Account Authentication Valid: Select to choose valid account authentication
- * User Name: User name for the storage
- * Password: Password for the storage

|Choose Tuning Manager for Performance|Override other performance options |Tuning Manager Host|IP address or fully-qualified domain name of tuning manager

Note: In HDS USP, USP V, and VSP, any disk can belong to more than one array group.

== Advanced configuration

|===

```
|Field|Description
|Connection Type|HTTPS or HTTP, also displays the default port
| HiCommand Server Port | Port used for the HiCommand Device Manager
//|HTTPs Enabled|Select to enable HTTPs
|Inventory Poll Interval (min) | Interval between inventory polls. The
default is 40.
|Choose 'Exclude' or 'Include' to specify a list|Specify whether to
include or exclude the array list below when collecting data.
|Filter device List|Comma-separated list of device serial numbers to
include or exclude
//|Query Host Manager|Select to query host manager
//|HTTP Timeout (sec)|HTTP connection timeout The default is 60.
|Performance Poll Interval (sec)|Interval between performance polls. The
default is 300.
|Export timeout in seconds|Export utility timeout. The default is 300.
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|Error: User does not have enough permission
|Use a different user account that has more privilege or increase the
privilege of user account configured in the data collector
|Error: Storages list is empty. Either devices are not configured or the
user does not have enough permission
|* Use DeviceManager to check if the devices are configured.
* Use a different user account that has more privilege, or increase the
privilege of the user account
|Error: HDS storage array was not refreshed for some days
|Investigate why this array is not being refreshed in HDS HiCommand.
|===
=== Performance
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|Error:
* Error executing export utility
* Error executing external command
|* Confirm that Export Utility is installed on the Cloud Insights
```

Acquisition Unit

- * Confirm that Export Utility location is correct in the data collector configuration
- * Confirm that the IP of the USP/R600 array is correct in the configuration of the data collector $\,$
- * Confirm that the User name and password are correct in the configuration of the data collector
- * Confirm that Export Utility version is compatible with storage array micro code version
- * From the Cloud Insights Acquisition Unit, open a CMD prompt and do the following:
- Change the directory to the configured installation directory
- Try to make a connection with the configured storage array by executing batch file runWin.bat

|Error: Export tool login failed for target IP

- |* Confirm that username/password is correct
- * Create a user ID mainly for this HDS data collector
- * Confirm that no other data collectors are configured to acquire this array

|Error: Export tools logged "Unable to get time range for monitoring".

- |* Confirm performance monitoring is enabled on the array.
- * Try invoking the export tools outside of Cloud Insights to confirm the problem lies outside of Cloud Insights.

|Error:

- * Configuration error: Storage Array not supported by Export Utility
- * Configuration error: Storage Array not supported by Storage Navigator Modular CLI
- |* Configure only supported storage arrays.
- * Use "Filter Device List" to exclude unsupported storage arrays.

|Error:

- * Error executing external command
- * Configuration error: Storage Array not reported by Inventory
- * Configuration error:export folder does not contains jar files
- |* Check Export utility location.
- * Check if Storage Array in question is configured in HiCommand server
- * Set Performance poll interval as multiple of 60 seconds.

|Error:

- * Error Storage navigator CLI
- * Error executing auperform command
- * Error executing external command
- |* Confirm that Storage Navigator Modular CLI is installed on the Cloud Insights Acquisition Unit
- * Confirm that Storage Navigator Modular CLI location is correct in the data collector configuration
- * Confirm that the IP of the WMS/SMS/SMS array is correct in the configuration of the data collector

- * Confirm that Storage Navigator Modular CLI version is compatible with micro code version of storage array configured in the data collector
- * From the Cloud Insights Acquisition Unit, open a CMD prompt and do the following:
- Change the directory to the configured installation directory
- Try to make a connection with the configured storage array by executing following command "auunitref.exe"

|Error: Configuration error: Storage Array not reported by Inventory |Check if Storage Array in question is configured in HiCommand server |Error:

- * No Array is registered with the Storage Navigator Modular 2 CLI
- * Array is not registered with the Storage Navigator Modular 2 CLI
- * Configuration error: Storage Array not registered with StorageNavigator Modular CLI
- |* Open Command prompt and change directory to the configured path
- * Run the command "set=STONAVM HOME=."
- * Run the command "auunitref"
- * Confirm that the command output contains details of the array with IP
- * If the output does not contain the array details then register the array with Storage Navigator CLI:
 - Open Command prompt and change directory to the configured path
 - Run the command "set=STONAVM HOME=."
- Run command "auunitaddauto -ip \${ip}". Replace \${ip} with real IP

Additional information may be found from the xref:{relative_path}concept_requesting_support.html[Support] page or in the xref:{relative_path}reference_data_collector_support_matrix.html[Data Collector Support Matrix].

```
[[ID93da16b799a62ad6fdbfaa67f352ba33]]
= Configuring the Hitachi Vantara NAS data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

The Hitachi Vantara NAS data collector is an inventory and configuration data collector that supports discovery of HDS NAS clusters. Cloud

Insights supports discovering NFS and CIFS shares, file systems (Internal Volumes), and spans (Storage Pools). == Terminology Cloud Insights acquires the following inventory information from the HNAS data collector. For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind: [cols=2*, options="header", cols"50,50"] |Vendor/Model Term|Cloud Insights Term |Tier|Disk Group |Cluster|Storage |Node|Storage Node |Span|Storage Pool |System Drive|Backend Lun |Files System|Internal Volume |=== Note: These are common terminology mappings only and might not represent every case for this data collector. == Requirements * Device IP address * Port 22, SSH protocol * Username and password - privilege level: Supervisor * Note: This data collector is SSH based, so the AU that hosts it must be able to initiate SSH sessions to TCP 22 on the HNAS itself, or the Systems Management Unit (SMU) that the cluster is connected to. == Configuration [cols=2*, options="header", cols"50,50"] |=== |Field|Description | HNAS Host|IP address or fully-qualified domain name of HNAS Management Host |User Name|User name for HNAS CLI |Password|Password used for HNAS CLI

|===

== Advanced configuration

```
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 30 minutes.
//|SSH Banner Wait Timeout (sec)|SSH banner wait timeout. The default is
15 seconds.
//|SSH Command Timeout (sec)|SSH command timeout. The default is 30
seconds.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|Problem:|Try this:
|"Error connecting" with error messages "Error setting up shell
channel: "or "Error opening shell channel"
|Likely caused by network connectivity issues or SSH is misconfigured.
Confirm connection with alternate SSH client
|"Timeout" or "Error retrieving data" with error messages "Command: XXX
has timed out."
|* Try the command with alternate SSH client
* Increase timeout
|"Error connecting " or "Invalid login credentials" with error messages
"Could not communicate with the device:"
| * Check IP address
* Check user name and password
* Confirm connection with alternate SSH client
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[IDec1c0ccdb532f25cebc30e9acebbb2b9]]
= Hitachi Ops Center data collector
:toc: macro
:hardbreaks:
```

```
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector uses Hitachi Ops Center's integrated suite of
applications to access inventory and performance data of multiple storage
devices. For inventory and capacity discovery, your Ops Center
installation must include both the "Common Services" and "Administrator"
components. For performance collection, you must additionally have
"Analyzer" deployed.
== Terminology
Cloud Insights acquires the following inventory information from this data
collector. For each asset type acquired by Cloud Insights, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term|Cloud Insights Term
|Storage Systems|Storage
|Volume|Volume
|Parity Groups|Storage Pool(RAID), Disk Groups
|Disk|Disk
|Storage Pool|Storage Pool(Thin, SNAP)
|External Parity Groups|Storage Pool(Backend), Disk Groups
|Port|Storage Node \rightarrow Controller Node \rightarrow Port
|Host Groups|Volume Mapping and Masking
|Volume Pairs|Storage Synchronization
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Inventory Requirements
You must have the following in order to collect inventory data:
* IP address or hostname of the Ops Center server hosting the "Common
Services" component
* Root/sysadmin user account and password that exist on all servers
hosting Ops Center components. HDS did not implement REST API support for
```

```
usage by LDAP/SSO users until Ops Center 10.8+
== Performance requirements
The following requirements must be met in order to collect performance
data:
The HDS Ops Center "Analyzer" module must be installed
Storage arrays must be feeding the Ops Center "Analyzer" module
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
| Hitachi Ops Center IP Address | IP address or fully-qualified domain name
of the Ops Center server hosting the "Common Services" component
|User Name |User name for the Ops Center server.
|Password|Password used for the Ops Center server.
|===
== Advanced configuration
|===
|Field|Description
|Connection Type|HTTPS (port 443) is the default
|Override TCP Port |Specify the port to use if not the default
|Inventory Poll Interval (min) | Interval between inventory polls. The
default is 40.
|Choose 'Exclude' or 'Include' to specify a list|Specify whether to
include or exclude the array list below when collecting data.
|Filter device List|Comma-separated list of device serial numbers to
include or exclude
|Performance Poll Interval (sec)|Interval between performance polls. The
default is 300.
|===
////
== Troubleshooting
Some things to try if you encounter problems with this data collector:
==== Inventory
[cols=2*, options="header", cols"50,50"]
|===
```

```
|Problem:|Try this:
|Error: Error message seen showing a IP/hostname and port number the
collector is not set to use
|This will be indicative that the collector is able to speak to Common
Services, learns where Administrator and Analyzer are,
| and is subsequently having difficulty to speak to them. If a HTTP 40x
error is observed, this likely means you are attempting to use a non
root/sysadmin account
| If a HTTP 5xx error is observed, that is likely a problem with the Ops
Center module in question
|===
////
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
:leveloffset: -1
[[ID2b21d6f120be407b177b385e9f54f4d5]]
= Infinidat InfiniBox data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The Infinidat InfiniBox (HTTP) data collector is used to collect inventory
information from the Infinidat InfiniBox storage system.
== Terminology
Cloud Insights acquires the following inventory information from the
Infinidat InfiniBox data collector. For each asset type acquired by Cloud
Insights, the most common terminology used for this asset is shown. When
viewing or troubleshooting this data collector, keep the following
terminology in mind:
```

```
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Storage Pool|Storage Pool
|Node|Controller
|Filesystem|Internal Volume
|Filesystem|File Share
|Filesystem Exports|Share
|===
== Requirements
The following are requirements when configuring this data collector.
* IP address or FQDN of InfiniBox management Node
* Admin userid and password
* Port 443 via REST API
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|InfiniBox Host|IP address or fully-qualified domain name of the InfiniBox
Management Node
|User Name|User name for InfiniBox Management Node
|Password|Password for the InfiniBox Management Node
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|TCP Port|TCP Port used to connect to InfiniBox Server. The default is
443.
|Inventory Poll Interval | Interval between inventory polls. The default is
60 minutes.
//|Connection Timeout|Connection timeout. The default is 60 seconds.
|===
== Troubleshooting
```

```
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID4ac369f4967fa1f162b16d107ce6156e]]
= Huawei OceanStor data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Huawei OceanStor (REST/HTTPS) data collector to
discover inventory and performance for Huawei OceanStor and OceanStor
Dorado storage.
== Terminology
Cloud Insights acquires the following inventory and performance
information from the Huawei OceanStor. For each asset type acquired by
Cloud Insights, the most common terminology used for this asset is shown.
When viewing or troubleshooting this data collector, keep the following
terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Storage Pool|Storage Pool
|File System|Internal Volume
|Controller|Storage Node
|FC Port (Mapped)|Volume Map
|Host FC Initiator (Mapped)|Volume Mask
|NFS/CIFS Share|Share
|iSCSI Link Target|iSCSI Target Node
|iSCSI Link Initiator|iSCSI Initiator Node
|Disk|Disk
|LUN|Volume
|===
== Requirements
```

```
The following requirements are required to configure this data collector:
* Device IP address
* Credentials to access OceanStor device manager
* Port 8088 must be available
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|OceanStor Host IP Address|IP address or fully-qualified domain name of
the OceanStor Device Manager
|User Name|Name used to log into the OceanStor Device Manager
|Password|Password used to log into the OceanStor Device Manager
== Advanced Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|TCP Port|TCP Port used to connect to OceanStor Device Manager. The
default is 8088.
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 60 minutes.
//|Connection Timeout (sec)|Connection timeout. The default is 60 seconds.
|Performance poll interval (sec).| The default is 300 seconds.
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
= IBM
:leveloffset: +1
```

```
[[ID8d7e8cabfa8abb2411ed8ee468219d9a]]
= IBM Cleversafe data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to discover inventory and
performance data for IBM Cleversafe storage systems.
NOTE: IBM Cleversafe is metered at a different Raw TB to Managed Unit
rate. Every 40 TB of unformatted IBM Cleversafe capacity is charged as 1
xref:{relative path}concept subscribing to cloud insights.html#pricing[Man
aged Unit (MU)].
== Terminology
Cloud Insights acquires the following inventory information from the IBM
Cleversafe data collector. For each asset type acquired by Cloud Insights,
the most common terminology used for this asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term |Cloud Insights Term
|Storage Pool|Storage Pool
|Container|Internal Volume
|Container|File Share
|NFS Share|Share
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* The external data services IP address for the cluster
* Administrator user name and password
* Port 9440
== Configuration
```

```
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|Manager IP or host name|IP address or hostname of management node
|User name|Username for the user account with super user or system
administrator role
|Password|Password for the user account with super user or system
administrator role
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|Inventory poll interval (min)|Interval between inventory polls.
|HTTP Connection Timeout (sec) | HTTP timeout in seconds.
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID42024e87118c3f81742507638b4096f3]]
= IBM CS data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to discover inventory and
performance data for IBM CS storage systems.
== Terminology
Cloud Insights acquires the following inventory information from the IBM
```

CS data collector. For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind: [cols=2*, options="header", cols"50,50"] |=== |Vendor/Model Term |Cloud Insights Term |Storage Pool|Storage Pool |Container|Internal Volume |Container|File Share |NFS Share|Share |=== Note: These are common terminology mappings only and might not represent every case for this data collector. == Requirements * The external data services IP address for the cluster * Administrator user name and password * Port 9440 == Configuration [cols=2*, options="header", cols"50,50"] |Field | Description |Prism External IP Address|The external data services IP address for the |User name|User name for the Admin account |Password|Password for the Admin account |=== == Advanced configuration [cols=2*, options="header", cols"50,50"] |Field | Description |TCP port|TCP Port used to connect to the IBM CS array. The default is |Inventory poll interval (min)|Interval between inventory polls. The default is 60 minutes. //|Connection timeout (sec)|Connection timeout The default is 60 seconds. |Performance poll interval(sec)|Interval between performance polls. The default is 300 seconds.

|=== == Troubleshooting Additional information on this Data Collector may be found from the xref:{relative path}concept requesting support.html[Support] page or in the xref:{relative path}reference data collector support matrix.html[Data Collector Support Matrix]. [[ID09e887e2e21cc0c2b0ac1cc9914ce141]] = IBM System Storage DS8000 Series data collector :toc: macro :hardbreaks: :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead] The IBM DS (CLI) data collector supports inventory and performance data acquisition for DS6xxx and DS8xxx devices. DS3xxx, DS4xxx, and DS5xxx devices are supported by the xref:{relative path}task dc na eseries.html[NetApp E-Series data collector]. You should refer to the Cloud Insights support matrix for supported models and firmware versions. == Terminology Cloud Insights acquires the following inventory information from the IBM DS data collector. For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind: [cols=2*, options="header", cols"50,50"]

```
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Disk Drive Module|Disk
|Storage Image|Storage
|Extent Pool|Storage Node
```

```
|Fixed Block Volume|Volume
| Host FC Initiator (Mapped) | Volume Mask
|===
Note: These are common terminology mappings only and might not represent
every case for this data collecor.
== Requirements
You need the following to configure this data collector:
* IP address of each DS array
* Read-only username and password on each DS array
* Third-party software installed on the Cloud Insights AU: IBM dscli
* Access validation: Run dscli commands using the username and password
* Port requirements: 80, 443, & 1750
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|DS Storage|IP address or fully-qualified domain name of the DS device
|User Name |User name for the DS CLI
|Password |Password for the DS CLI
| dscli executable path | Full path to the dscli executable
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min) |Interval between inventory polls (min).
The default is 40.
|Storage Display Name | Name of the IBM DS storage array
|Inventory Exclude Devices|Comma-separated list of device serial numbers
to exclude from inventory collection
|Performance Poll Interval (sec)|The default is 300.
|Performance Filter Type|Include: Data collected only from devices on
list. Exclude: No data from these devices is collected
|Performance Filter Device List|Comma-separated list of device IDs to
include or exclude from performance collection
| ====
== Troubleshooting
```

```
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|Problem:|Try this:
|Error containing: CMUC00192E, CMUC00191E or CMUC00190E
|* Verify credentials and IP address entered.
* Try to communicate with the array through web management console
https://${ip}:8452/DS8000/Console. Replace the ${ip} with data collector
configured IP.
|Error:
* Cannot run program
* Error executing command
|* From Cloud Insights Acquisition Unit Open a CMD
* Open CLI.CFG file in CLI's home dir/lib and check property JAVA INSTALL,
edit the value to match your environment
* Display Java version installed on this machine, typing: "java -version"
* Ping the IP address of the IBM Storage device specified in CLI command
issued.
* If all the above worked fine then manually run a CLI command
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID9e8d865a8947a0ec27b93dca173699ea]]
= Configuring the IBM PowerVM data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The IBM PowerVM (SSH) data collector is used to collect information about
virtual partitions running on IBM POWER hardware instances managed by a
```

hardware management console (HMC).

== Terminology

Cloud Insights acquires inventory information from the virtual partitions running on IBM POWER hardware instances. For each asset type acquired, the most common terminology used for the asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

```
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term | Cloud Insights Term
|hdisk |Virtual Disk
|Managed System|Host
|LPAR, VIO Server|Virtual Machine
|Volume Group|Data Store
|Physical Volume|LUN
|===
```

Note: These are common terminology mappings only and might not represent every case for this data collector.

== Requirements

The following requirements must be met to configure and use this data collector:

- * IP address of the Hardware Management Console (HMC)
- * User name and password that provide access to Hardware Management Console (HMC) through SSH
- * Port requirement SSH-22
- * View permission on all management systems and logical partition security domains

T m1

The user must also have View permission on HMC configurations and the ability to collect VPD information for the HMC console security grouping. The user must also be allowed Virtual IO Server Command access under the Logical Partition security grouping. It is a best practice to start from a role of an operator and then remove all roles. Read-only users on the HMC do not have privileges to run proxied commands on AIX hosts.

* IBM best practice is to have the devices monitored by two or more HMCs. Be aware that this may cause OnCommand Insight to report duplicated devices, therefore it is highly recommended to add redundant devices to the "Exclude Devices" list in the Advanced Configuration for this data

```
collector.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
| Hardware Management Console (HMC) IP Address | IP address or fully-
qualified domain name of the PowerVM Hardware Management Console
|HMC User|User name for the Hardware Management Console
|Password|Password used for the Hardware Management Console
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 20 minutes.
|SSH Port|Port used for SSH to the PowerVM
|Password|Password used for the Hardware Management Console
//|SSH Process Wait Timeout (sec)|SSH process timeout. The default is 600
seconds.
|Number of Retries|Number of inventory retry attempts
|Exclude Devices|Comma-separated list of device IDs or display names to
exclude
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID5a330644772d4419d9d88d91b30e4315]]
= Configuring the IBM SAN Volume Controller data collector
:toc: macro
:hardbreaks:
:toclevels: 2
```

```
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path} {relative_path}./media/
```

[.lead]

The IBM SAN Volume Controller (SVC) data collector collects inventory and performance data using SSH, supporting a variety of devices that run the SVC operating system.

The list of supported devices includes models such as the SVC, the v7000, the v5000, and the v3700. Refer to the Cloud Insights support matrix for supported models and firmware versions.

== Terminology

Cloud Insights acquires the following inventory information from the IBM SVC data collector. For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

```
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term | Cloud Insights Term
|Drive|Disk
|Cluster|Storage
|Node|Storage Node
|Mdisk Group|Storage Pool
|Vdisk|Volume
|Mdisk|Backend LUNs and paths
|===
```

Note: These are common terminology mappings only and might not represent every case for this data collector.

- === Inventory Requirements
- * IP address of each SVC cluster
- * Port 22 available
- * Read-only user name and password
- === Performance Requirements
- * SVC Console, which is mandatory for any SVC cluster and required for the SVC discovery foundation package.
- * Credentials will require administrative access level only for copying performance files from cluster nodes to the config node.

```
* Enable data collection by connecting to the SVC cluster by SSH and
running: svctask startstats -interval 1
+
Note: Alternatively, enable data collection using the SVC management user
interface.
////
* IP address of each SVC cluster
* Port 22 available
* Public and private key pair that you either generate with Cloud Insights
or reuse a keypair already in use on your SVC
If you are reusing an existing keypair, you must convert them from Putty
format to OpenSSH format.
* Public key installed on the SVC cluster
* Private key needs to be identified in the Acquisition Unit.
* Access validation: Open ssh session to the SVC cluster using the private
key
//Note: No third-party software needs to be installed.
== Performance Requirements
* SVC Console, which is mandatory for any SVC cluster and required for the
SVC discovery foundation package.
* Administrative access level required only for copying performance files
from cluster nodes to the config node.
Note: Because this access level is not required for the SVC foundation
discovery package, the SVC foundation user might not work successfully.
//* A private and public SSH key must be generated for this user, and the
private key stored so that it is accessible from the Acquisition Unit. If
the SVC foundation user has the proper permissions, then the same user and
key works. The same SSH key can be used for inventory and performance
* Enable data collection by connecting to the SVC cluster by SSH and
running: svctask startstats -interval 1
Note: Alternatively, enable data collection using the SVC management user
interface.
* Port Requirement: 22
////
```

```
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Cluster IP Addresses | IP addresses or fully-qualified domain names of the
SVC storage
//|'Password' or 'OpenSSH Key File'|Credential type used to connect to the
device via SSH
|Inventory User Name|User name for the SVC CLI
|Inventory Password | Password for the SVC CLI
//|Full Path to Inventory Private Key|Full path to the Inventory private
key file
//|Performance User Name|User name for the SVC CLI for performance
collection
collection
//|Full Path to Performance Private Key|Full path to the Performance
private key file
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 40 minutes.
//|Exclude Devices|Comma-separated list of device IDs to exclude from
inventory collection
//|SSH Process Wait Timeout (sec)|SSH process timeout. The default is 200
seconds.
|Performance Poll Interval (sec)|Interval between performance polls. The
default is 300 seconds.
//|Performance Exclude Devices|Comma-separated list of device IDs to
exclude from performance collection
//|Performance SSH Process Wait Timeout (sec)|SSH process timeout. The
default is 200 seconds.
collection
|To clean up dumped stats files|Select this checkbox to clean up dumped
stats files
|===
== Troubleshooting
```

```
Some things to try if you encounter problems with this data collector:
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|Error: "The command cannot be initiated because it was not run on the
configuration node."
|The command must be executed on the configuration node.
|===
Some things to try if you encounter problems with this data collector:
[cols=2*, options="header", cols"50,50"]
|===
|Problem: |Try this:
|Error: "The command cannot be initiated because it was not run on the
configuration node."
|The command must be executed on the configuration node.
|===
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID711dfdaa052c8ff4613c30e05448b416]]
= Configuring the IBM XIV/A9000 data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
IBM XIV and A9000 (CLI) data collector uses the XIV command-line interface
to collect inventory data while performance collection is accomplished by
making SMI-S calls to the XIV/A9000 array, which runs a SMI-S provider on
port 7778.
== Terminology
[cols=2*, options="header", cols"50,50"]
```

```
|Vendor/Model Term | Cloud Insights Term
|Disk|Disk
|Storage System|Storage
|Storage Pool|Storage Pool
|Volume|Volume
|===
== Requirements
The following requirements must be met to configure and use this data
collector:
* Port requirement: TCP port 7778
* Read-only user name and password
* The XIV CLI must be installed on the AU
== Performance requirements
The following are requirements for performance collection:
* SMI-S Agent 1.4 or higher
* SMI-S compatible CIMService running on array. Most XIV arrays have a
CIMServer installed by default.
* User login must be provided for the CIMServer. The login must have full
read access to the array configuration and properties.
* SMI-S namespace. Default is root/ibm. This is configurable in the
CIMServer.
* Port Requirements: 5988 for HTTP, 5989 for HTTPS.
* Refer to the following link on how to create an account for SMI-S
performance collection:
http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp?topic=%2F
com.ibm.tpc V41.doc%2Ffqz0 t adding cim agent.html
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|XIV IP address|IP address or fully-qualified domain name of the XIV
storage
|User Name |User name for the XIV storage
```

```
|Password|Password for the XIV storage
|Full Path to XIV CLI Directory|Full path to the folder containing the XIV
CLI
|SMI-S Host IP Address|IP address of the SMI-S host
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 40 minutes.
//|CLI Process Wait Timeout (sec)|CLI process timeout. The default is
7200000 ms.
//|SMI-S Host IP|IP address of the SMI-S Provider Host
//|SMI-S Port|Port used by SMI-S Provider Host
|SMI-S Protocol|Protocol used to connect to the SMI-S provider. Also
displays the default port.
|Override SMI-S Port|If blank, use the default port in the Connection Type
field, otherwise enter the connection port to use
//|SMI-S Namespace|SMI-S namespace
|Username|User name for the SMI-S Provider Host
|Password|Password for the SMI-S Provider Host
|Performance Poll Interval (sec)|Interval between performance polls. The
default is 300 seconds.
//|Number of SMI-S Connection Retries|Number of SMI-S connection retry
attempts.
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
:leveloffset: -1
[[IDd6eeb0721ffe4460a0ee68a55a674a51]]
= Lenovo data collector
```

```
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Lenovo data collector to discover inventory and
performance data for Lenovo HX storage systems.
////
== Terminology
Cloud Insights acquires the following inventory information from the
Lenovo data collector. For each asset type acquired by Cloud Insights, the
most common terminology used for this asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Storage Pool|Storage Pool
|Nutanix Container|Internal Volume
|Nutanix Container|File Share
|NFS Share|Share
|===
////
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* Prism External IP Address
* Administrator user name and password
* TCP Port requirement: 9440
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|Prism External IP Address|The external data services IP address for the
```

```
cluster
|User name|User name for the Admin account
|Password|Password for the Admin account
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|TCP port|TCP Port used to connect to array. The default is 9440.
|Inventory poll interval (min)|Interval between inventory polls. The
default is 60 minutes.
//|Connection timeout (sec)|Connection timeout The default is 60 seconds.
|Performance poll interval (sec)|Interval between performance polls. The
default is 300 seconds.
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
= Microsoft
:leveloffset: +1
[[IDda03135a518b1e2121a57a766c1e1a10]]
= Configuring the Azure NetApp Files data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Azure NetApp Files data collector to acquire
inventory and performance data.
```

```
== Requirements
You need the following information to configure this data collector.
* Port requirement: 443 HTTPS
* Azure Management Rest IP (management.azure.com)
* Azure service principal client ID (user account)
* Azure service principal authentication key (user password)
* You need to set up an Azure account for Cloud Insights discovery.
Once the account is properly configured and you register the application
in Azure, you will have the credentials required to discover the Azure
instance with Cloud Insights. The following link describes how to set up
the account for discovery:
https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-
create-service-principal-portal
== Configuration
Enter data into the data collector fields according to the table below:
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|Azure Service Principal Client ID|Sign-in ID to Azure
|Azure Tenant ID|Azure Tenant ID
|Azure Service Principal Authentication Key|Login authentication key
|I understand Microsoft bills me for API requests|Check this to verify
your understanding that Microsoft bills you for API requests made by
Insight polling.
|===
== Advanced Configuration
[cols=2*, options="header", cols"50,50"]
|Field | Description
|Inventory Poll Interval (min)|The default is 60
//|HTTP connection and socket timeout (sec)|The default is 300
//|Choose 'Exclude' or 'Include' to Apply to Filter VMs by Tags|Specify
whether to include or exclude VM's by Tags when collecting data. If
'Include' is selected, the Tag Key field can not be empty.
//|Tag Keys and Values on which to Filter VMs|Click *+ Filter Tag* to
```

choose which VMs (and associated disks) to include/exclude by filtering for keys and values that match keys and values of tags on the VM. Tag Key is required, Tag Value is optional. When Tag Value is empty, the VM is filtered as long as it matches the Tag Key.

//|Performance Poll Interval (sec)|The default is 300
|===

== Troubleshooting

- * The credentials used by your ANF data collector must not have access to any Azure subscriptions that contain ANF volumes.
- * If Reader access causes performance collection to fail, try granting contributor access on a resource group level.

Additional information on this Data Collector may be found from the xref:{relative_path}concept_requesting_support.html[Support] page or in the xref:{relative_path}reference_data_collector_support_matrix.html[Data Collector Support Matrix].

```
[[IDc7514bd3b84716786ca613375b144b50]]
= Microsoft Hyper-V data collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

The Microsoft Hyper-V data collector acquires inventory and performance data from the virtualized server computing environment. This data collector can discover a standalone Hyper-V host, or an entire cluster - create one collector per standalone host or cluster.

== Terminology

Cloud Insights acquires the following inventory information from the Microsoft Hyper-V (WMI). For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

```
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term|Cloud Insights Term
|Virtual Hard Disk|Virtual Disk
|Host|Host
|Virtual Machine|Virtual Machine
|Cluster Shared Volumes (CSV), Partition Volume|Data Store
|Internet SCSI Device, Multi Path SCSI LUN|LUN
|Fiber Channel Port|Port
|===
```

Note: These are common terminology mappings only and might not represent every case for this data collector.

== Requirements

The following are required to configure this data collector:

- * The Hyper-V requires port 5985 opened for data collection and remote access/management.
- * IP address or FQDN of cluster or standalone hypervisor. Using the floating cluster hostname or IP is likely the most reliable approach versus pointing the collector at just one specific node in a cluster.
- * Administrative-level user account that works on all the hypervisors in the cluster.
- * WinRM needs to be enabled and listening on all hypervisors
- * Port requirements: Port 135 via WMI & Dynamic TCP ports assigned 1024-65535 for Windows 2003 and older and 49152-65535 for Windows 2008.
- * DNS resolution must succeed, even if the data collector is pointed at only an IP address
- * Each Hyper-V hypervisor must have "Resource Metering" turned on for every VM, on every host. This allows each hypervisor to have more data available for Cloud Insights on each guest. If this is not set, fewer performance metrics are acquired for each guest. More information on Resource metering can be found in the Microsoft documentation:

link:https://docs.microsoft.com/en-us/previous-versions/windows/itpro/windows-server-2012-R2-and-2012/hh831661(v=ws.11)[Hyper-V Resource
Metering Overview]

+

link:https://docs.microsoft.com/en-us/powershell/module/hyper-v/enablevmresourcemetering?view=win10-ps[Enable-VMResourceMetering]

NOTE: The Hyper-V data collector requires a Windows Acquisition Unit.

////

```
Best Practice: It is highly recommended for each Hyper-V hypervisor to
have "Resource Metering" turned on for every VM, on every host. This
allows each hypervisor to have more data available for Cloud Insights on
each guest. If this is not set, fewer performance metrics are acquired for
each quest. More information on Resource metering can be found in the
link:https://docs.microsoft.com/en-us/previous-versions/windows/it-
pro/windows-server-2012-R2-and-2012/hh831661(v=ws.11)[Microsoft
documentation].
////
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Cluster IP address or floating cluster FQDN|The IP address or fully-
qualified domain name for the the cluster, or a standalone, non-clustered
hypervisor
|User Name|Administrator user name for the hypervisor
|Password|Password for the hypervisor
|DNS domain suffix|The hostname suffix that combines with the simple
hostname to render the FQDN of a hypervisor
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min)|The default is 20 minutes.
//|Connection Timeout (ms)|The default is 60000 ms.
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
:leveloffset: -1
```

```
= NetApp
:leveloffset: +1
[[ID59533b9a61330dbafaece5779d6daff3]]
= NetApp Cloud Connection for ONTAP 9.9+ data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector creates a cloud connection to support data collection
from ONTAP 9.9+ CVO, AFF, and FAS systems.
NOTE: This data collector is no longer available to install in Cloud
Insights as of April 4, 2023, and will be removed from all Cloud Insights
installations in July 2023. For information on transitioning to AU-based
data collection, see the
link:https://kb.netapp.com/Advice and Troubleshooting/Cloud Services/Cloud
Insights/How to transition from NetApp Cloud Connection to AU based data
collector[Knowledgebase^].
////
NOTE: This data collector is
xref:{relative path}task getting started with cloud insights.html#useful-
definitions[deprecated] as of January 1, 2023, and is no longer available
as of April 2023. For information on transitioning to AU-based data
collection, see the
link:https://kb.netapp.com/Advice and Troubleshooting/Cloud Services/Cloud
Insights/How to transition from NetApp Cloud Connection to AU based data
collector[Knowledgebase^].
////
////
== Configuration
Cloud Insights collects data from ONTAP 9.9+ using a *cloud connection*,
eliminating the need to install an external acquisition unit, simplifying
troubleshooting, maintenance, and initial deployment. Configuration of the
cloud connection for the ONTAP 9.9+ data collector requires you to copy a
*Pairing Code* to the ONTAP System Manager, which will then establish a
```

```
connection to your Cloud Insights environment. After the connection is
established, the data collected is the same as it would be if it was
collected through an acquisition unit.
This data collector supports ONTAP 9.9+ CVO, AFF, and FAS systems.
image:Cloud Agent DC.png[Cloud Agent Data Collector Configuration]
Follow these steps to configure the connection:
* Generate a unique token which will be used to establish the connection
to the ONTAP system.
* Copy the Pairing Code, which includes the token. You can view the
pairing code by clicking on [+] Reveal Code Snippet .
Once you copy the pairing code, the data collector configuration screen
will reveal a step 6, prompting you to wait for the connection to be
established. Nothing more needs to be done on this screen until the
connection is established.
image:Cloud Agent Step Waiting.png[Waiting for connection]
* In a new browser tab, log into the ONTAP System Manager and navigate to
Cluster > Settings > Cloud Connections .
* Click Add Cloud Connection and paste the pairing code.
* Return to the Cloud Insights browser tab and wait for the connection to
be established. Once it is established, a Complete button is revealed.
* Click Complete .
// The Cloud Connection data collector acquires EMS (Event Monitoring
System) logs from ONTAP.
== Troubleshooting
Some things to try if you encounter problems with this data collector:
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|I'm seeing the following error while trying to connect to Azure CVO: "The
certificate signing request to broker/manager CA service was not
```

```
completed."
|Verify that your Cloud manager proxy settings are set to the Cloud
Manager private IP. Cloud Manager installation may set a different proxy.
Once the proxy is set to the correct IP and you reference the proxy in the
Cloud Connector dialog, the connection to Cloud Insights should connect
successfully.
|===
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
////
[[IDd518645fd62916ca57bd1b8c03c2dd99]]
= NetApp Cloud Volumes ONTAP data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector supports inventory collection from Cloud Volumes ONTAP
configurations.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|NetApp Management IP Address | IP address for Cloud Volumens ONTAP
|User Name | User name for Cloud Volumes ONTAP
|Password| Password for the above user
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
```

```
|Field|Description
|Connection Type | HTTPS recommended. Also shows default port.
|Override Communication Port | Port to use if not default.
|Inventory Poll Interval (min) |Default is 60 minutes.
| Inventory Concurrent Thread Count | Number of concurrent threads.
|Force TLS for HTTPS |Force TLS over HTTPS
|Automatically Lookup Netgroups|Automatically Lookup Netgroups
|Netgroup Expansion |Select Shell or File
|HTTP read timeout seconds |Default is 30 seconds
|Force responses as UTF-8 |Force responses as UTF-8
|Performance Poll Interval (min) | Default is 900 seconds.
| Performance Concurrent Thread Count | Number of concurrent threads.
|Advanced Counter Data Collection | Check this to have Cloud Insights
collect the advanced metrics from the list below.
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[IDdb838af8d4c0192eb4f802447f1370a0]]
= NetApp Cloud Volumes Services for AWS data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector supports inventory collection from NetApp Cloud
Volumes Services for AWS configurations.
== Configuration
[cols=2*, options="header", cols"50,50"]
```

```
|===
|Field|Description
|Cloud Volumes Region | Region of the NetApp Cloud Volumes Services for AWS
|API Key |Cloud Volumes API key
|Secret Key |Cloud Volumes secret key
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min) |Default is 60 minutes
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
[cols=2*, options="header", cols"50,50"]
|Problem:|Try this:
|I received an error similar to this one:
'Failed to execute request: Connect to <AWS region endpoint>:8080 [<AWS
region endpoint>/AWS region endpoint IP>] failed: connect timed out: GET
https://<AWS Region Endpoint FQDN>:8080/v1/Storage/IPRanges HTTP/1.1'
|The xref:{relative path}task configure acquisition unit.html#proxy-
configuration-2[proxy] used by Cloud Insights to communicate with the
Acquisition Unit does not communicate between Cloud Insights and the Data
Collector itself. Here are a few things you can try:
Ensure that the acquisition unit is able to resolve the fqdn and reach the
required port.
Confirm that a proxy is not required to reach the specified endpoint in
the error message.
Curl can be used to test the communication between the acquisition unit
and the endpoint. Make sure that you are *not* using a Proxy for this
test.
Example:
root@acquisitionunit# curl -s -H accept:application/json -H "Content-type:
application/json" -H api-key: <api key used in the data collector
credentials -H secret-key: < secret key used in the data collector
credentials> -X GET https://<AWS Regional</pre>
```

```
Endpoint>:8080/v1/Storage/IPRanges
See this
link:https://kb.netapp.com/Advice and Troubleshooting/Cloud Services/Cloud
Insights/Cloud Insights fails discovery for Cloud Volumes Service for AWS
[NetApp KB article].
|===
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID905962d0c629ffcada0f0982e4559b16]]
= NetApp ONTAP Data Management Software data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
This data collector acquires inventory and performance data from storage
systems running ONTAP using read-only API calls from an ONTAP account.
This data collector also creates a record in the cluster application
registry to accelerate support.
== Terminology
Cloud Insights acquires inventory and performance data from the ONTAP data
collector. For each asset type acquired, the most common terminology used
for the asset is shown. When viewing or troubleshooting this data
collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Disk|Disk
|Raid Group|Disk Group
|Cluster|Storage
|Node|Storage Node
|Aggregate|Storage Pool
```

|LUN|Volume |Volume|Internal Volume |===

== ONTAP Data Management Terminology

The following terms apply to objects or references that you might find on ONTAP Data Management storage asset landing pages. Many of these terms apply to other data collectors as well.

=== Storage

- * Model A comma-delimited list of the unique, discrete node model names within this cluster. If all the nodes in the clusters are the same model type, just one model name will appear.
- * Vendor same Vendor name you would see if you were configuring a new data source.
- * Serial number The array serial number. On cluster architecture storage systems like ONTAP Data Management, this serial number may be less useful than the individual "Storage Nodes" serial numbers.
- * IP generally will be the IP(s) or hostname(s) as configured in the data source.
- * Microcode version firmware.
- * Raw Capacity base 2 summation of all the physical disks in the system, regardless of their role.
- * Latency a representation of what the host facing workloads are experiencing, across both reads and writes. Ideally, Cloud Insights is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, Cloud Insights is generally performing an IOPs-weighted calculation derived from the individual internal volumes' statistics.
- * Throughput aggregated from internal volumes.

Management - this may contain a hyperlink for the management interface of the device. Created programmatically by the Cloud Insights data source as part of inventory reporting.

=== Storage Pool

- * Storage what storage array this pool lives on. Mandatory.
- * Type a descriptive value from a list of an enumerated list of possibilities. Most commonly will be "Aggregate" or "RAID Group"".
- * Node if this storage array's architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page.
- * Uses Flash Pool Yes/No value does this SATA/SAS based pool have SSDs used for caching acceleration?

- * Redundancy RAID level or protection scheme. RAID_DP is dual parity, RAID TP is triple parity.
- * Capacity the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these.
- * Over-committed capacity If by using efficiency technologies you have allocated a sum total of volume or internal volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- * Snapshot snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots. ONTAP in MetroCluster configurations are likely to exhibit this, while other ONTAP configurations are less so.
- * Utilization a percentage value showing the highest disk busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host driven workloads. Also, many arrays' replication implementations may drive disk utilization while not showing as internal volume or volume workload.
- * IOPS the sum IOPs of all the disks contributing capacity to this storage pool.

Throughput - the sum throughput of all the disks contributing capacity to this storage pool.

=== Storage Node

- * Storage what storage array this node is part of. Mandatory.
- * HA Partner on platforms where a node will fail over to one and only one other node, it will generally be seen here.
- * State health of the node. Only available when the array is healthy enough to be inventoried by a data source.
- * Model model name of the node.
- * Version version name of the device.
- * Serial number The node serial number.
- * Memory base 2 memory if available.
- * Utilization On ONTAP, this is a controller stress index from a proprietary algorithm. With every performance poll, a number between 0 and 100% will be reported that is the higher of either WAFL disk contention, or average CPU utilization. If you observe sustained values > 50%, that is indicative of undersizing potentially a controller/node not large enough or not enough spinning disks to absorb the write workload.
- * IOPS Derived directly from ONTAP ZAPI calls on the node object.
- * Latency Derived directly from ONTAP ZAPI calls on the node object.
- * Throughput Derived directly from ONTAP ZAPI calls on the node object.
- * Processors CPU count.

```
== Requirements
The following are requirements to configure and use this data collector:
* You must have access to an Administrator account configured for read-
only API calls.
* Account details include username and password.
* Port requirements: 80 or 443
* Account permissions:
** Read only role name to ontapi application to the default Vserver
** You may require additional optional write permissions. See the Note
About Permissions below.
* ONTAP License requirements:
** FCP license and mapped/masked volumes required for fibre-channel
discovery
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|NetApp Management IP | IP address or fully-qualified domain name of the
NetApp cluster
|User Name |User name for NetApp cluster
|Password |Password for NetApp cluster
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Connection type|Choose HTTP (default port 80) or HTTPS (default port
443). The default is HTTPS
|Override Communication Port|Specify a different port if you do not want
to use the default
|Inventory Poll Interval (min) |Default is 60 minutes.
//|Inventory concurrent thread count|Thread count for parallel foundation
|For TLS for HTTPS|Only allow TLS as protocol when using HTTPS
|Automatically Lookup Netgroups|Enable the automatic netgroup lookups for
export policy rules
|Netgroup Expansion|Netgroup Expansion Strategy. Choose file or shell .
The default is _shell_.
```

|HTTP read timeout seconds|Default is 30

|Force responses as UTF-8|Forces data collector code to interpret responses from the CLI as being in UTF-8

//|Foundation model writing timeout|Max time before aborting when writing the model to disk. -1 Disables the feature. 0 is for testing only.
|Performance Poll Interval (sec)|Default is 900 seconds.

//|Performance Concurrent thread count|Thread count for parallel performance queries

//|Performance model writing timeout|Max time before aborting when writing the model to disk. -1 Disables the feature. 0 is for testing only.

|Advanced Counter Data Collection|Enable ONTAP integration. Select this to include ONTAP Advanced Counter data in polls. Choose the desired counters from the list.

//|<TBD: New Micro Poll> | Default is 60 seconds

|Cluster Switch Metrics|Allow Cloud Insights to collect cluster switch data. Note that in addition to enabling this on the Cloud Insights side, you must also configure the ONTAP system to provide link:https://docs.netapp.com/us-en/ontap-cli-98/system-switch-ethernet-create.html[switch information], and ensure the correct <<a-note-about-permissions, permissions>> are set, in order to allow the switch data to be sent to Cloud Insights. See "A Note About Permissions" below.

|===

== ONTAP Power Metrics

Several ONTAP models provide power metrics for Cloud Insights that can be used for monitoring or alerting. The lists of supported and unsupported models below are not comprehensive but should provide some guidance; in general, if a model is in the same family as one on the list, the support should be the same.

Supported Models:

A200

A220

A250

A300

A320

A400

A700

A700s

A800 A900 C190 FAS2240-4 FAS2552 FAS2650 FAS2720 FAS2750 FAS8200 FAS8300 FAS8700 FAS8700

Unsupported Models:

FAS2620

FAS3250

FAS3270

FAS500f

FAS6280

FAS/AFF 8020

FAS/AFF 8040

FAS/AFF 8060

FAS/AFF 8080

== A Note About Permissions

Since a number of Cloud Insights' ONTAP dashboards rely on advanced ONTAP counters, you must enable *Advanced Counter Data Collection* in the data collector Advanced Configuration section.

You should also ensure that write permission to the ONTAP API is enabled. This typically requires an account at the cluster level with the necessary permissions.

To create a local account for Cloud Insights at the cluster level, log in to ONTAP with the Cluster management Administrator username/password, and execute the following commands on the ONTAP server:

- . Before you begin, you must be signed in to ONTAP with an _Administrator_ account, and _diagnostic-level commands_ must be enabled.
- . Create a read-only role using the following commands.

security login role create -role ci readonly -cmddirname DEFAULT -access

readonly security login role create -role ci_readonly -cmddirname security -access readonly security login role create -role ci readonly -access all -cmddirname {cluster application-record create} . Create the read-only user using the following command. Once you have executed the create command, you will be prompted to enter a password for this user. security login create -username ci user -application ontapi -authentication-method password -role ci_readonly If AD/LDAP account is used, the command should be security login create -user-or-group-name DOMAIN\aduser/adgroup -application ontapi -authentication-method domain -role ci readonly If you are collecting cluster switch data: security login rest-role create -role ci readonly -api /api/network/ethernet -access readonly The resulting role and user login will look something like the following. Your actual output may vary: Role Command/ Access Vserver Name Directory Query Level cluster1 ci readonly DEFAULT read only cluster1 ci readonly security readonly cluster1::security login> show Vserver: cluster1 Authentication Acct UserName Application Method Role Name Locked _____ ----ci_user ontapi password ci_readonly no NOTE: If ONTAP access control is not set correctly, then Cloud Insights REST calls may fail, resulting in gaps in data for the device. For

NOTE: If ONTAP access control is not set correctly, then Cloud Insights REST calls may fail, resulting in gaps in data for the device. For example, if you have enabled it on the Cloud Insights collector but have not configured the permissions on the ONTAP, acquisition will fail. Additionally, if the role is previously defined on the ONTAP and you are adding the Rest API abilities, ensure that http is added to the role.

== Troubleshooting Some things to try if you encounter problems with this data collector: === Inventory

[cols=2*, options="header", cols"50,50"]
|===

|Problem:|Try this:

|Receive 401 HTTP response or 13003 ZAPI error code and ZAPI returns "Insufficient privileges" or "not authorized for this command" | Check username and password, and user privileges/permissions.

|Cluster version is < 8.1| Cluster minimum supported version is 8.1. Upgrade to minimum supported version.

|ZAPI returns "cluster role is not cluster_mgmt LIF"|AU needs to talk to cluster management IP. Check the IP and change to a different IP if necessary

|Error: "7 Mode filers are not supported"| This can happen if you use this data collector to discover 7 mode filer. Change IP to point to cdot cluster instead.

|ZAPI command fails after retry| AU has communication problem with the cluster. Check network, port number, and IP address. User should also try to run a command from command line from the AU machine.

|AU failed to connect to ZAPI via HTTP| Check whether ZAPI port accepts plaintext. If AU tries to send plaintext to an SSL socket, the communication fails.

|Communication fails with SSLException|AU is attempting to send SSL to a plaintext port on a filer. Check whether the ZAPI port accepts SSL, or use a different port.

|Additional Connection errors:

ZAPI response has error code 13001, "database is not open"

ZAPI error code is 60 and response contains "API did not finish on time"

ZAPI response contains "initialize session() returned NULL environment"

ZAPI error code is 14007 and response contains "Node is not healthy"

```
|Check network, port number, and IP address. User should also try to run a
command from command line from the AU machine.
|===
=== Performance
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|"Failed to collect performance from ZAPI" error|This is usually due to
perf stat not running. Try the following command on each node:
> system node systemshell -node * -command "spmctl -h cmd -stop; spmctl
-h cmd -exec"
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID9eeaf6444336c6a509e848c7c60981a3]]
= NetApp Data ONTAP operating in 7-Mode data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
For storage systems using Data ONTAP software operating in 7-Mode, you use
the 7-mode data collector, which uses the CLI to obtain capacity and
performance data.
== Terminology
Cloud Insights acquires the following inventory information from the
NetApp 7-mode data collector. For each asset type acquired, the most
common terminology used for this asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
```

```
NOTE: This data collector is
xref:{relative path}task getting started with cloud insights.html#useful-
definitions[deprecated].
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Disk|Disk
|Raid Group|Disk Group
|Filer|Storage
|Filer|Storage Node
|Aggregate|Storage Pool
|LUN|Volume
|Volume|Internal Volume
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
You need the following to configure and use this data collector:
* IP addresses of the FAS storage controller and partner.
* Port 443
* A custom admin level username and password for controller and partner
controller with the following role capabilities for 7-Mode:
** "api-*": Use this to allow OnCommand Insight to execute all NetApp
storage API commands.
** "login-http-admin": Use this to allow OnCommand Insight to connect to
the NetApp storage via HTTP.
** "security-api-vfiler": Use this to allow OnCommand Insight to execute
NetApp storage API commands to retrieve vFiler unit information.
** "cli-options": Use this to read storage system options.
** "cli-lun": Access these commands for managing LUNs. Displays the status
(LUN path, size, online/offline state, and shared state) of the given LUN
or class of LUNs.
** "cli-df": Use this to display free disk space.
** "cli-ifconfig": Use this to display interfaces and IP addresses.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Address of storage system|IP address or fully-qualified domain name for
```

```
the NetApp storage system
|User Name|User name for the NetApp storage system
|Password|Password for the NetApp storage system
|Address of HA Partner in Cluster|IP address or fully-qualified domain
name for the HA Partner
|User Name of HA Partner in Cluster|User name for the HA partner
|Password of HA Partner Filer in Cluster|Password for the HA Partner
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 20 minutes.
|Connection Type | HTTPS or HTTP, also displays the default port
|Override Connection Port | If blank, use the default port in the
Connection Type field, otherwise enter the connection port to use
|Performance Poll Interval (sec)|Interval between performance polls. The
default is 300 seconds.
|===
== Storage systems connection
As an alternative to using the default administrative user for this data
collector, you can configure a user with administrative rights directly on
the NetApp storage systems so that this data collector can acquire data
from NetApp storage systems.
Connecting to NetApp storage systems requires that the user, who is
specified when acquiring the main pfiler (on which the storage system
exist), meet the following conditions:
* The user must be on vfiler0 (root filer/pfiler).
Storage systems are acquired when acquiring the main pfiler.
* The following commands define the user role capabilities:
** "api-*": Use this to allow Cloud Insights to execute all NetApp storage
API commands.
This command is required to use the ZAPI.
** "login-http-admin": Use this to allow Cloud Insights to connect to the
NetApp storage via HTTP. This command is required to use the ZAPI.
```

- ** "security-api-vfiler": Use this to allow Cloud Insights to execute NetApp storage API commands to retrieve vFiler unit information.
- ** "cli-options": For "options" command and used for partner IP and enabled licenses.
- ** "cli-lun": Access these command for managing LUNs. Displays the status (LUN path, size, online/offline state, and shared state) of the given LUN or class of LUNs.
- ** "cli-df": For "df -s", "df -r", "df -A -r" commands and used to display free space.
- ** "cli-ifconfig": For "ifconfig -a" command and used for getting filer IP address.
- ** "cli-rdfile": For "rdfile /etc/netgroup" command and used for getting netgroups.
- ** "cli-date": For "date" command and used to get full date for getting Snapshot copies.
- ** "cli-snap": For "snap list" command and used for getting Snapshot copies.

If cli-date or cli-snap permissions are not provided, acquisition can finish, but Snapshot copies are not reported.

To acquire a 7-Mode data source successfully and generate no warnings on the storage system, you should use one of the following command strings to define your user roles. The second string listed here is a streamlined version of the first:

- * login-http-admin,api-*,security-api-vfile,cli-rdfile,cli-options,cli-df,cli-lun,cli-ifconfig,cli-date,cli-snap,
- * login-http-admin,api-* ,security-api-vfile,cli-
- == Troubleshooting

Some things to try if you encounter problems with this data collector:

=== Inventory

[cols=2*, options="header", cols"50,50"]

|Problem:|Try this:

|Receive 401 HTTP response or 13003 ZAPI error code and ZAPI returns "Insufficient privileges" or "not authorized for this command" | Check username and password, and user privileges/permissions.

|"Failed to execute command" error|Check whether the user has the following permission on the device:

- api-*
- cli-date
- cli-df
- cli-ifconfig
- cli-lun
- cli-operations
- cli-rdfile
- cli-snap
- login-http-admin
- security-api-vfiler

Also check if the ONTAP version is supported by Cloud Insights and verify if the credentials used match device credentials

|Cluster version is < 8.1| Cluster minimum supported version is 8.1. Upgrade to minimum supported version.

|ZAPI returns "cluster role is not cluster_mgmt LIF"|AU needs to talk to cluster management IP. Check the IP and change to a different IP if necessary

|Error: "7 Mode filers are not supported"| This can happen if you use this data collector to discover 7 mode filer. Change IP to point to cdot filer instead.

|ZAPI command fails after retry| AU has communication problem with the cluster. Check network, port number, and IP address. User should also try to run a command from command line from the AU machine.

|AU| failed to connect to ZAPI|Check IP/port connectivity and assert ZAPI configuration.

|AU failed to connect to ZAPI via HTTP| Check whether ZAPI port accepts plaintext. If AU tries to send plaintext to an SSL socket, the communication fails.

|Communication fails with SSLException|AU is attempting to send SSL to a plaintext port on a filer. Check whether the ZAPI port accepts SSL, or use a different port.

|Additional Connection errors:

ZAPI response has error code 13001, "database is not open"

ZAPI error code is 60 and response contains "API did not finish on time"

ZAPI response contains "initialize session() returned NULL environment"

ZAPI error code is 14007 and response contains "Node is not healthy"

|Check network, port number, and IP address. User should also try to run a command from command line from the AU machine.

|Socket timeout error with ZAPI|Check filer connectivity and/or increase timeout.

```
|"C Mode clusters are not supported by the 7 Mode data source" error|Check
IP and change the IP to a 7 Mode cluster.
|"Failed to connect to vFiler" error| Check that the acquiring user
capabilities include the following at a minimum:
security-api-vfiler
login-http-admin
Confirm that filer is running minimum ONTAPI version 1.7.
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID5112e5b1f7c56a48eb408892a304da3b]]
= NetApp E-Series data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The NetApp E-Series data collector gathers inventory and performance data.
The collector supports firmware 7.x+ using the same configurations and
reporting the same data.
== Terminology
Cloud insight acquires the following inventory information from the NetApp
E-Series data collector. For each asset type acquired, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Disk|Disk
```

|Volume Group|Disk Group |Storage Array|Storage |Controller|Storage Node |Volume Group|Storage Pool |Volume|Volume |===

Note: These are common terminology mappings only and might not represent every case for this data collector.

== E-Series Terminology (Landing Page)

The following terms apply to objects or references that you might find on NetApp E-Series asset landing pages. Many of these terms apply to other data collectors as well.

=== Storage

- * Model model name of the device.
- * Vendor same Vendor name you would see if you were configuring a new datasource
- * Serial number The array serial number. On cluster architecture storage systems like NetApp Clustered Data Ontap, this serial number may be less useful than the individual "Storage Nodes" serial numbers

// EMC Isilon, IBM SVC

- * IP generally will be the IP(s) or hostname(s) as configured in the data source
- * Microcode version firmware
- * Raw Capacity base 2 summation of all the physical disks in the system, regardless of their role
- * Latency a representation of what the host facing workloads are experiencing, across both reads and writes. Ideally, Cloud Insights is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, Cloud Insights is generally performing an IOPs-weighted calculation derived from the individual volumes' statistics.
- * Throughput the array's total host facing throughput. Ideally sourced directly from the array, if unavailable, Cloud Insights is summing the volumes' throughput to derive this value
- * Management this may contain a hyperlink for the management interface of the device. Created programmatically by the Cloud Insights datasource as part of inventory reporting

=== Storage Pool

* Storage - what storage array this pool lives on. Mandatory

- * Type a descriptive value from a list of an enumerated list of possibilities. Most commonly will be "Thin Provisioning" or "RAID Group"
- * Node if this storage array's architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page
- * Uses Flash Pool Yes/No value
- * Redundancy RAID level or protection scheme. E-Series reports "RAID 7" for DDP pools
- * Capacity the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these. These value both include E-Series "preservation" capacity, resulting both in numbers and the percentage being higher than what the E-Series own user interface may show
- * Over-committed capacity If via efficiency technologies you have allocated a sum total of volume or internal volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- * Snapshot snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots
- * Utilization a percentage value showing the highest disk busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host driven workloads. Also, many arrays' replication implementations may drive disk utilization while not showing as volume workload.
- * IOPS the sum IOPs of all the disks contributing capacity to this storage pool. If disk IOPs is not available on a given platform, this value will be sourced from the sum of volume IOPs for all the volumes sitting on this storage pool
- * Throughput the sum throughput of all the disks contributing capacity to this storage pool. If disk throughput is not available on a given platform, this value will be sourced from the sum of volume throughout for all the volumes sitting on this storage pool

=== Storage Node

- * Storage what storage array this node is part of. Mandatory
- * HA Partner on platforms where a node will fail over to one and only one other node, it will generally be seen here
- * State health of the node. Only available when the array is healthy enough to be inventoried by a data source
- * Model model name of the node
- * Version version name of the device.

```
* Serial number - The node serial number
* Memory - base 2 memory if available
* Utilization - Generally a CPU utilization number, or in the case of
NetApp Ontap, a controller stress index. Utilization is not currently
available for NetApp E-Series
* IOPS - a number representing the host driven IOPs on this controller.
Ideally sourced directly from the array, if unavailable, it will be
calculated by summing all the IOPs for volumes that belong exclusively to
// Available for NetApp E-Series with OCI 7.3.10, or by installing Data
Source Service Pack 8 or higher for OCI 7.3.1+
* Latency - a number representing the typical host latency or response
time on this controller. Ideally sourced directly from the array, if
unavailable, it will be calculated by performing an IOPs weighted
calculation from volumes that belong exclusively to this node.
// Available for NetApp E-Series with OCI 7.3.10, or by installing Data
Source Service Pack 8 or higher for OCI 7.3.1+
* Throughput - a number representing the host driven throughput on this
controller. Ideally sourced directly from the array, if unavailable, it
will be calculated by summing all the throughput for volumes that belong
exclusively to this node.
// Available for NetApp E-Series with OCI 7.3.10, or by installing Data
Source Service Pack 8 or higher for OCI 7.3.1+
* Processors - CPU count
== Requirements
* The IP address of each controller on the array
* Port requirement 2463
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Comma-separated list of Array SANtricity Controller IPs| IP addresses
and/or fully-qualified domain names for the array controllers
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min) |Default is 30 minutes
```

```
|Performance Poll Interval up to 3600 seconds|Default is 300 seconds
|===
== Troubleshooting
Additional information on this data collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[IDeaf72c390e9af9d0dcd8278f4263c064]]
= Configuring the NetApp HCI Management server data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The NetApp HCI Management server data collector collects NetApp HCI Host
information and requires read-only privileges on all objects within the
Management server.
This data collector acquires from the *NetApp HCI Management server only*.
To collect data from the storage system, you must also configure the
xref:{relative path}task dc na solidfire.html[NetApp SolidFire] data
collector.
== Terminology
Cloud Insights acquires the following inventory information from this data
collector. For each asset type acquired, the most common terminology used
for the asset is shown. When viewing or troubleshooting this data
collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Virtual disk|Disk
|Host|Host
|Virtual machine|Virtual machine
```

```
|Data store|Data store
|LUN|Volume
|Fibre channel port|Port
|===
These are common terminology mappings only and might not represent every
case for this data collector.
== Requirements
The following information is required to configure this data collector:
* IP address of the NetApp HCI Management server
* Read-only username and password for the NetApp HCI Management server
* Read only privileges on all objects in the NetApp HCI Management server.
* SDK access on the NetApp HCI Management server - normally already set
//* 3rd party software installed on NetApp HCI Management server / RAU:
none
* Port requirements: http-80 https-443
* Validate access:
** Log into the NetApp HCI Management server using above username and
password
** Verify SDK enabled: telnet <vc ip> 443
== Setup and connection
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|Name|Unique name for the data collector
|Acquisition unit|Name of acquisition unit
|===
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|NetApp HCI Storage Cluster MVIP |Management Virtual IP Address
|SolidFire Management Node (mNode)|Management Node IP Address
|User name |User name used to access the NetApp HCI Management server
|Password|Password used to access the NetApp HCI Management server
| VCenter User Name | User name for VCenter
|VCenter Password|Password for VCenter
|===
```

```
== Advanced configuration
In the advanced configuration screen, check the *VM Performance* box to
collect performance data. Inventory collection is enabled by default.
The following fields can be configured:
[cols=2*, cols"50,50"]
|===
|Field|Description
|Inventory poll interval (min) | Deafult is 20
//|Connection Timeout (ms)|Default is 60000
|Filter VMs by|Select CLUSTER, DATACENTER, or ESX HOST
|Choose 'Exclude' or 'Include' to Specify a List|Specify Whether to
Include or Exclude VMs
|Filter Device List |List of VMs to filter (comma separated, or semicolon
separated if comma is used in the value) for for Filtering by ESX HOST,
CLUSTER, and DATACENTER Only
//or you can choose to filter by TAG
//|Number of retries | Default is 3
//|HCI Management port| Default is 8443
//|Tag Keys and Values on which to Filter VMs|Click *+ Filter Tag* to
choose which VMs (and associated disks) to include/exclude by filtering
for keys and values that match keys and values of tags on the VM. Tag Key
is required, Tag Value is optional. When Tag Value is empty, the VM is
filtered as long as it matches the Tag Key.
//Tag filtering is only available in VSphere 6.0 Beta or later.
|Performance poll interval (sec)|Default is 300
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|Error: Include list to filter VMs cannot be empty
|If Include List is selected, please list valid DataCenter, Cluster, or
Host names to filter VMs
|Error: Failed to instantiate a connection to VirtualCenter at IP
| Possible solutions:
```

- * Verify credentials and IP address entered.
- * Try to communicate with Virtual Center using Infrastructure Client.
- * Try to communicate with Virtual Center using Managed Object Browser (e.g MOB).

|Error: VirtualCenter at IP has non-conform certificate that JVM requires |Possible solutions:

- * Recommended: Re-generate certificate for Virtual Center by using stronger (e.g. 1024-bit) RSA key.
- * Not Recommended: Modify the JVM java.security configuration to leverage the constraint jdk.certpath.disabledAlgorithms to allow 512-bit RSA key. See JDK 7 update 40 release notes at

"http://www.oracle.com/technetwork/java/javase/7u40-relnotes-2004172.html"

Additional information may be found from the xref:{relative_path}concept_requesting_support.html[Support] page or in the xref:{relative_path}reference_data_collector_support_matrix.html[Data Collector Support Matrix].

```
[[ID1faa28ec2492156f31faf6588887020e]]
= NetApp SolidFire All-Flash Array data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

The NetApp SolidFire All-Flash Array data collector supports inventory and performance collection from both iSCSI and Fibre Channel SolidFire configurations.

The SolidFire data collector utilizes the SolidFire REST API. The acquisition unit where the data collector resides needs to be able to initiate HTTPS connections to TCP port 443 on the SolidFire cluster management IP address. The data collector needs credentials capable of making REST API queries on the SolidFire cluster.

== Terminology

Cloud Insights acquires the following inventory information from the NetApp SolidFire All-Flash Array data collector. For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind: [cols=2*, options="header", cols"50,50"] |Vendor/Model Term | Cloud Insights Term |Drive|Disk |Cluster|Storage |Node|Storage Node |Volume|Volume |Fibre channel port|Port |Volume Access Group, LUN Assignment| Volume Map |iSCSI Session|Volume Mask |=== Note: These are common terminology mappings only and might not represent every case for this data collector. == Requirements The following are requirements for configuring this data collector: * Management Virtual IP Address * Read-only username and credentials * Port 443 == Configuration [cols=2*, options="header", cols"50,50"] |=== |Field|Description |Management Virtual IP Address (MVIP) |Management Virtual IP address of the SolidFire Cluster |User Name |Name used to log into the SolidFire cluster |Password |Password used to log into the SolidFire cluster |=== == Advanced configuration [cols=2*, options="header", cols"50,50"] |=== |Field|Description |Connection Type |Choose connection type

```
|Communication Port |Port used for NetApp API |Inventory Poll Interval (min) |Default is 20 minutes |Performance Poll Interval (sec) |Default is 300 seconds |===
```

== Troubleshooting

When SolidFire reports an error it is displayed in Cloud Insights as follows:

_An error message was received from a SolidFire device while trying to retrieve data. The call was <method> (<parameterString>). The error message from the device was (check the device manual): <message>

Where:

- The <method> is an HTTP method, such as GET or PUT.
- The <parameterString> is a comma separated list of parameters that were included in the REST call.
- The <message> is whatever the device returned as the error message.

Additional information on this Data Collector may be found from the xref:{relative_path}concept_requesting_support.html[Support] page or in the xref:{relative_path}reference_data_collector_support_matrix.html[Data Collector Support Matrix].

```
[[IDcab887c9c5a2bdla3a4bda0dl810708c]]
= NetApp StorageGRID data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/

[.lead]
The NetApp StorageGRID data collector supports inventory and performance collection from StorageGRID configurations.
NOTE: StorageGRID is metered at a different Raw TB to Managed Unit rate.
Every 40 TB of unformatted StorageGRID capacity is charged as 1
```

xref:{relative path}concept subscribing to cloud insights.html#pricing[Man

```
aged Unit (MU)].
== Terminology
Cloud Insights acquires the following inventory information from the
NetApp StorageGRID collector. For each asset type acquired, the most
common terminology used for this asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|StorageGRID|Storage
| Node | Node
|Tenant|Storage Pool
|Bucket|Internal Volume
|===
== Requirements
The following are requirements for configuring this data source:
* StorageGRID Host IP Address
* A username and password for a user that has had the Metric Query and
Tenant Access roles assigned
* Port 443
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|StorageGRID Host IP Address | Management Virtual IP address of the
StorageGRID appliance
|User Name |Name used to log into the StorageGRID appliance
|Password |Password used to log into the StorageGRID appliance
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Inventory Poll Interval (min) |Default is 60 minutes
|performance Poll Interval (sec)|Default is 900 seconds
|===
```

```
== Single Sign-On (SSO)
The link:https://docs.netapp.com/sqws-112/index.jsp[StorageGRID] firmware
versions have corresponding API versions; 3.0 API and newer versions
support single sign-on (SSO) login.
|===
|Firmware version |API version |Support single sign on (SSO)
|11.1 |2 |No
|11.2 |3.0 |Yes
|11.5 |3.3 |Yes
|===
== Troubleshooting
Additional information on this Data Collector may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
:leveloffset: -1
[[ID54dbb9cccd5aac9b28794886e5bfb5dc]]
= Nutanix NX data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Nutanix data collector to discover inventory and
performance data for Nutanix NX storage systems.
== Terminology
Cloud Insights acquires the following inventory information from the
Nutanix data collector. For each asset type acquired by Cloud Insights,
the most common terminology used for this asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
```

```
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term | Cloud Insights Term
|Storage Pool|Storage Pool
|Nutanix Container|Internal Volume
|Nutanix Container|File Share
|NFS Share|Share
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* The external data services IP address for the cluster
* Read-only user name and password, unless volume groups are in use, in
which case, Admin user name and password are required
* Port requirement: HTTPS 443
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|Prism External IP Address|The external data services IP address for the
|User name|User name for the Admin account
|Password|Password for the Admin account
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|TCP port|TCP Port used to connect to Nutanix array. The default is 9440.
|Inventory poll interval (min)|Interval between inventory polls. The
default is 60 minutes.
//|Connection timeout (sec)|Connection timeout The default is 60 seconds.
|Performance poll interval(sec)|Interval between performance polls. The
default is 300 seconds.
|===
== Troubleshooting
```

Additional information on this Data Collector may be found from the xref:{relative path}concept requesting support.html[Support] page or in the xref:{relative path}reference data collector support matrix.html[Data Collector Support Matrix]. [[IDf05ba047ac3c677536da9cbc1771c12e]] = OpenStack data collector :toc: macro :hardbreaks: :toclevels: 2 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root_path}{relative_path}./media/ [.lead] The OpenStack (REST API / KVM) data collector acquires inventory data for all OpenStack instances, and optionally, VM performance data. == Requirements * IP address of the OpenStack controller * OpenStack admin role credential and sudo access to the Linux KVM hypervisor. If you are not using the admin account or admin equivalent privileges, you will need to use trial and error to identify the default policies to relax for your data collector userid. * The OpenStack Gnocchi module must be installed and configured for performance collection. Configuring Gnocchi is done by editing the nova.conf file for each hypervisor and then restarting the Nova Compute service on each hypervisor. The option name changes for different releases of OpenStack: ** Icehouse ** Juno ** Kilo ** Liberty ** Mitaka ** Newton ** Ocata * For CPU stats, "compute monitors=ComputeDriverCPUMonitor" needs to be turned on in /etc/nova/nova.conf on compute nodes. * Port requirements: ** 5000 for http and 13000 for https, for the Keystone service ** 22 for KVM SSH

```
** 8774 for Nova Compute Service
** 8776 for Cinder Block Service
** 8777 for Gnocchi Performance Service
** 9292 for Glance Image Service
*Note* The port binds to the specific service, and the service may run on
the controller or another host in larger environments.
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
| OpenStack Controller IP Address|IP address or fully-qualified domain
name of the OpenStack Controller
|OpenStack Administrator|User name for an OpenStack Admin
|OpenStack Password|Password used for the OpenStack Admin
|OpenStack Administrator Tenant|OpenStack Administrator Tenant name
|KVM Sudo User|KVM Sudo User name
|Choose 'Password' or 'OpenSSH Key File' to specify credential
type | Credential type used to connect to the device via SSH
|Full Path to Inventory Private Key|Full Path to Inventory Private Key
| KVM Sudo Password | KVM Sudo Password
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field|Description
|Enable hypervisor inventory discovery through SSH|Check this to enable
hypervisor inventory discovery through SSH
|OpenStack Admin URL port|OpenStack Admin URL port
|Use HTTPS|Check to use secure HTTP
//|HTTP Connection Timeout (sec)|Timeout for HTTP connection. The default
is 300 seconds.
|SSH Port|Port used for SSH
//|SSH Process Wait Timeout (sec)|SSH process timeout. The default is 30
seconds.
|SSH Process Retries|Number of inventory retry attempts
|Inventory Poll Interval (min)|Interval between inventory polls. The
default is 20 minutes.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
```

```
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|"Configuration error" with error messages start with "Policy doesn't
allow" or "You are not authorized"
| * Check ip address
* Check User name and password
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID1a10d0991a059eb3d193e5768d4dc9f4]]
= Oracle ZFS Storage Appliance data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Oracle ZFS Storage Appliance data collector to
gather inventory and performance data.
== Terminology
Cloud Insights acquires inventory information with the Oracle ZFS data
collector. For each asset type acquired by Cloud Insights, the most common
terminology used for this asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term | Cloud Insights Term
|Disk (SSD)|Disk
|Cluster|Storage
|Controller|Storage Node
|LUN|Volume
```

```
|LUN Map|Volume Map
|Initiator, Target|Volume Mask
|Share|Internal Volume
|===
Note: These are common terminology mappings only and might not represent
every case for this data source.
== Requirements
* Host names for the ZFS Controller-1 and the ZFS Controller-2
* Administrator user name and password
* Port requirement: 215 HTTP/HTTPS
== Required Performance metrics
Oracle ZFS appliances give storage administators large amounts of
flexibility to capture performance statistics. Cloud Insights expects you
to have each controller in a high availability pair configured to
capture the following metrics:
* smb2.ops[share]
* nfs3.ops[share]
* nfs4.ops[share]
* nfs4-1.ops[share]
Failure to have the controller capture any or all of these will likely
result in Cloud Insights not having, or underreporting, the workload on
the "Internal Volumes".
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field | Description
|ZFS Controller-1 Hostname|Host name for storage controller 1
|ZFS Controller-2 Hostname|Host name for storage controller 2
|User name|User name for the storage system administrator user account
|Password|Password for the administrator user account
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|Field | Description
|Connection Type | HTTPS or HTTP, also displays the default port
```

```
|Override Connection Port | If blank, use the default port in the
Connection Type field, otherwise enter the connection port to use
|Inventory poll interval|The default is 60 seconds
|Performance Poll Interval (sec)|The default is 300.
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|"Invalid login credentials"
|validate Zfs user account and password
|"Configuration error" with error message "REST Service is disabled"
|Verify REST service is enabled on this device.
|"Configuration error " with error message "User unauthorized for command"
|Likely due to certain roles (for example, 'advanced analytics') are not
included for the configured user <userName>.
Possible Solution:
* Correct the Analytics (statistic) scope for the user ${user} with the
read only role:
- From the Configuration -> Users screen, put your mouse over the role and
double click to allow editing
- Select "Analytics" from the Scope drop down menu. A list of the
possible properties appears.
   Click the top most check box and it will select all three properties.
  Click the Add button on the right side.
- Click the Apply button at the top right of the pop-up window. The pop-
up window will close.
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID22d6e8f56f59716d0d25739a47b7264f]]
= Pure Storage FlashArray data collector
:toc: macro
:hardbreaks:
```

```
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Pure Storage FlashArray data collector to gather
inventory and performance data.
== Terminology
For each asset type acquired by Cloud Insights, the most common
terminology used for the asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term | Cloud Insights Term
|Drive (SSD)|Disk
|Array|Storage
|Controller|Storage Node
|Volume|Volume
|LUN Map|Volume Map
|Initiator, Target|Volume Mask
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* Storage system IP address
* User name and password for the Administrator account of the Pure storage
system.
* Port requirement: HTTP/HTTPS 80/443
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field | Description
|FlashArray Host IP Address|IP address of the storage system
|User name |User name with admin privileges
|Password for the admin privileged account|Password
|===
```

```
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|Connection type|Choose HTTP or HTTPS. Also displays the default port.
|Override TCP port|If blank, use the default port in the Connection Type
field, otherwise enter the connection port to use
|Inventory poll interval (min)|The default is 60 minutes
//|Connection Timeout (sec)|The default is 60
|Performance Poll Interval (sec)|The default is 300
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|Problem:|Try this:
|"Invalid login credentials" with error messages "Policy doesn't allow" or
"You are not authorized"
|Validate Pure user account and password via Pure http interface
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[ID9bbc5fb799637b6d7ef819e74ea672d9]]
= Red Hat Virtualization data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses the Red Hat Virtualization data collector to gather
inventory data from virtualized Linux and Microsoft Windows workloads.
```

```
== Terminology
For each asset type acquired by Cloud Insights, the most common
terminology used for the asset is shown. When viewing or troubleshooting
this data collector, keep the following terminology in mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Disk|Virtual Disk
|Host|Host
|Virtual Machine|Virtual Machine
|Storage Domain|Data Store
|Logical Unit|LUN
|===
Note: These are common terminology mappings only and might not represent
every case for this data collector.
== Requirements
* IP address of the RHEV server over port 443 via REST API
* Read-only username and password
* RHEV Version 3.0+
== Configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|RHEV Server IP Address|IP address of the storage system
|User name |User name with admin privileges
|Password for the admin privileged account|Password
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|HTTPS Communication Port|Port used for HTTPS communication to RHEV
|Inventory poll interval (min)|The default is 20 minutes.
//|Connection Timeout (sec)|The default is 60.
|===
== Troubleshooting
```

Additional information on this Data Collector may be found from the xref:{relative_path}concept_requesting_support.html[Support] page or in the xref:{relative_path}reference_data_collector_support_matrix.html[Data Collector Support Matrix].

```
[[ID3c54d197c1ebcf19953fcb63af17f421]]
= Rubrik CDM Data Collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

Cloud Insights uses the Rubrik data collector to acquire inventory and performance data from Rubrik storage appliances.

== Terminology

Cloud Insights acquires the following inventory information from the Rubrik data collector. For each asset type acquired by Cloud Insights, the most common terminology used for this asset is shown. When viewing or troubleshooting this data collector, keep the following terminology in mind:

```
[cols=2*, options="header", cols"50,50"]
|===
|Vendor/Model Term | Cloud Insights Term
|Cluster|Storage, Storage Pool
|Node|Storage Node
|Disk|Disk
|===
```

Note: These are common terminology mappings only and might not represent every case for this data source.

```
== Requirements
The following are required to configure this data collector:
* The Cloud Insights Acquisition Unit will initiate connections to TCP
port 443 to Rubrik cluster. One collector per cluster.
* Rubrik cluster IP address.
* User name and password to the cluster.
* Port requirement: HTTPS 443
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field | Description
|IP|IP address of the Rubrik cluster
|User name|User name for the cluster
|Password|Password for the cluster
|===
== Advanced configuration
[cols=2*, options="header", cols"50,50"]
|Inventory poll interval (min)|The default is 60
|Performance Poll Interval (sec)|The default is 300
|===
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|Problem:|Try this:
|I received a message that more than one storage is created.
|Check that the cluster is configured correctly, and the collector is
pointing to a single cluster.
|I received a warning that disk API returned more data
|Check with support to get additional data.
Additional information may be found from the
```

```
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative_path}reference data collector support matrix.html[Data
Collector Support Matrix].
[[IDda57dfa15d453b4a43c9b4c9b517129c]]
= Configuring the VMware VSphere data collector
:toc: macro
:hardbreaks:
:toclevels: 2
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
The data collector for VMware vSphere collects ESX Host information and
requires read-only privileges on all objects within the Virtual Center.
== Terminology
Cloud Insights acquires the following inventory information from the
VMware vSphere data collector. For each asset type acquired, the most
common terminology used for the asset is shown. When viewing or
troubleshooting this data collector, keep the following terminology in
mind:
[cols=2*, options="header", cols"50,50"]
|Vendor/Model Term | Cloud Insights Term
|Virtual disk|Disk
|Host|Host
|Virtual machine|Virtual machine
|Data store|Data store
|LUN|Volume
|Fibre channel port|Port
These are common terminology mappings only and might not represent every
case for this data collector.
== Requirements
The following information is required to configure this data collector:
* IP address of the Virtual Center server
* Read-only username and password in Virtual Center
```

```
* We require read only privileges on all objects within Virtual Center.
* SDK access on the Virtual Center server - normally already setup.
* Port requirements: http-80 https-443
* Validate access:
** Log into Virtual Center Client using above username and password
** Verify SDK enabled: telnet <vc ip> 443
== Setup and connection
[cols=2*, options="header", cols"50,50"]
|===
|Field | Description
|Name|Unique name for the data collector
|Acquisition unit|Name of acquisition unit
|===
== Configuration
[cols=2*, options="header", cols"50,50"]
|Field|Description
|Virtual center IP Address | IP address of the Virtual Center
|User name |User name used to access the Virtual Center
|Password|Password used to access the Virtual Center
|===
== Advanced configuration
In the advanced configuration screen, check the *VM Performance* box to
collect performance data. Inventory collection is enabled by default.
The following fields can be configured:
[cols=2*, cols"50,50"]
|===
|Field|Description
|Inventory poll interval (min) | Default is 20
//|Connection Timeout (ms)|Default is 60000
|Filter VMs |Select CLUSTER, DATACENTER, or ESX HOST
//or you can choose to filter by TAG
|Choose 'Exclude' or 'Include' to Specify a List|Create a filter list
(CLUSTER, DATACENTER, and/or ESX HOST)
|Number of retries | Default is 3
|Communication port| Default is 443
//|Tag Keys and Values on which to Filter VMs|Click *+ Filter Tag* to
```

choose which VMs (and associated disks) to include/exclude by filtering for keys and values that match keys and values of tags on the VM. Tag Key is required, Tag Value is optional. When Tag Value is empty, the VM is filtered as long as it matches the Tag Key.

//Tag filtering is only available in VSphere 6.0 Beta or later.

|Filter Device List...|This list must consist of exact string matches - if you intend to filter by ESX_HOST, you must build a comma delimited list of the exact "names" of your ESX hosts as reported in both Cloud Insights and vSphere. These "names" may be either IP addresses, simple hostnames, or fully qualified domain names (FQDNs) - this is determined by how these hosts were named when they were originally added to vSphere.

When filtering by CLUSTER, use the Cloud Insights-style cluster names as reported by CI on hypervisors - Cloud Insights prepends the vSphere cluster name with the vSphere datacenter name and a forward slash - "DC1/clusterA" is the cluster name Cloud Insights would report on a hypervisor in clusterA within data center DC1.

|Performance poll interval (sec)|Default is 300 |===

== Mapping VMware tags to Cloud Insights annotations

The VMware data collector allows you to populate Cloud Insights annotations with tags configured on VMware. The annotations must be named exactly as the VMware tags. Cloud Insights will always populate same-named text-type annotations, and will make a "best attempt" to populate annotations of other types (number, boolean, etc). If your annotation is of a different type and the data collector fails to populate it, it may be necessary to remove the annotation and re-create it as a text type.

Note that VMware tags may be case-sensitive, while Cloud Insights tags are case-insensitive. So if you create an annotation named "OWNER" in Cloud Insights, and tags named "OWNER", "Owner", and "owner" in VMware, all of those variations of "owner" would map to Cloud Insight's "OWNER" annotation.

Keep the following in mind:

- * Cloud Insights currently only auto-publishes support information for NetApp devices.
- * Since this support information is held in annotation form, you can query it or use it in dashboards.
- * If a user overwrites or empties the annotation value, the value is autofilled again when Cloud Insights updates annotations, which it does

```
once a day.
== Troubleshooting
Some things to try if you encounter problems with this data collector:
=== Inventory
[cols=2*, options="header", cols"50,50"]
|===
|Problem:|Try this:
|Error: Include list to filter VMs cannot be empty
|If Include List is selected, please list valid DataCenter, Cluster, or
Host names to filter VMs
|Error: Failed to instantiate a connection to VirtualCenter at IP
| Possible solutions:
* Verify credentials and IP address entered.
* Try to communicate with Virtual Center using VMware Infrastructure
Client.
* Try to communicate with Virtual Center using Managed Object Browser (e.g
|Error: VirtualCenter at IP has non-conform certificate that JVM requires
|Possible solutions:
* Recommended: Re-generate certificate for Virtual Center by using
stronger (e.g. 1024-bit) RSA key.
* Not Recommended: Modify the JVM java.security configuration to leverage
the constraint jdk.certpath.disabledAlgorithms to allow 512-bit RSA key.
See JDK 7 update 40 release notes at
"http://www.oracle.com/technetwork/java/javase/7u40-relnotes-2004172.html"
|===
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page or in
the xref:{relative path}reference data collector support matrix.html[Data
Collector Support Matrix].
:leveloffset: -1
```

```
= Data Collector Reference - Services
:leveloffset: +1
[[ID2927ab2bae3f5ad3be64398a6cce61d7]]
= Node Data Collection
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights gathers metrics from the node on which you install an
agent.
== Installation
. From *Observability > Collectors*, choose an operating system/platform.
Note that installing any integration data collector (Kubernetes, Docker,
Apache, etc.) will also configure node data collection.
. Follow the instructions to configure the agent. The instructions vary
depending on the type of Operating System or Platform you are using to
collect data.
== Objects and Counters
The following objects and their counters are collected as Node metrics:
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
|Node Filesystem
|Node UUID
Device
Path
Type
|Node IP
Node Name
Node OS
```

Mode |Free Inodes Free Inodes Total Inodes Used Total Used Total Used |Node Disk |Node UUID Disk |Node IP Node Name Node OS |IO Time Total IOPS In Progress Read Bytes (per sec) Read Time Total Reads (per sec) Weighted IO Time Total Write Bytes (per sec) Write Time Total Writes (per sec) Current Disk Queue Length Write Time Read Time IO Time |Node CPU |Node UUID CPU |Node IP Node Name Node OS |System CPU Usage User CPU Usage Idle CPU Usage Processor CPU Usage

Interrupt CPU Usage DPC CPU Usage Node |Node UUID |Node IP Node Name Node OS |Kernel Boot Time Kernel Context Switches (per sec) Kernel Entropy Available Kernel Interrupts (per sec) Kernel Processes Forked (per sec) Memory Active Memory Available Total Memory Available Memory Buffered Memory Cached Memory Commit Limit Memory Committed As Memory Dirty Memory Free Memory High Free Memory High Total Memory Huge Page Size Memory Huge Pages Free Memory Huge Pages Total Memory Low Free Memory Low Total Memory Mapped Memory Page Tables Memory Shared Memory Slab Memory Swap Cached Memory Swap Free Memory Swap Total Memory Total Memory Used Total Memory Used Memory Vmalloc Chunk Memory Vmalloc Total Memory Vmalloc Used Memory Wired

Memory Writeback Total Memory Writeback Tmp Memory Cache Faults Memory Demand Zero Faults Memory Page Faults Memory Pages Memory Nonpaged Memory Paged Memory Cache Core Memory Standby Cache Normal Memory Standby Cache Reserve Memory Transition Faults Processes Blocked Processes Dead Processes Idle Processes Paging Processes Running Processes Sleeping Processes Stopped Processes Total Processes Total Threads Processes Unknown Processes Zombies Processor Queue Length Swap Free Swap Total Swap Used Total Swap Used Swap In Swap Out System Uptime System Num CPU System Num Users System Calls |Node Network |Network Interface Node UUID |Node Name Node IP Node OS |Bytes Received Bytes Sent

```
Packets Outboud Discarded
Packets Outboud Errors
Packets Received Discarded
Packets Received Errors
Packets Received
Packets Sent
|===
== Setup
Setup and Troubleshooting information can be found on the
xref:{relative_path}task_config_telegraf_agent.html[Configuring an Agent]
page.
[[IDe84898c1fe7d47ef5bd856be5e6c7dc7]]
= ActiveMQ Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from ActiveMQ.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
ActiveMQ.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
```

group data collectors, for example, by OS/Platform. . Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data. image:ActiveMQDCConfigWindows.png[ActiveMQ configuration] == Setup Information may be found in the http://activemq.apache.org/gettingstarted.html[ActiveMQ documentation] == Objects and Counters The following objects and their counters are collected: [cols="<.<,<.<,<.<"] |=== |Object:|Identifiers:|Attributes: |Datapoints: |ActiveMQ Queue |Namespace Queue Port Server |Node Name Node IP Node UUID |Consumer Count Dequeue Count Enqueue Count Queue Size |ActiveMQ Subscriber |Client ID Connection ID Port Server Namespace | Is Active Destination Node Name

```
Node IP
Node UUID
Node OS
Selector
Subscription
| Dequeue Count
Dispatched Count
Dispatched Queue Size
Enqueue Count
Pending Queue Size
|ActiveMQ Topic
|Topic
Port
Server
Namespace
|Node Name
Node IP
Node UUID
Node OS
|Consumer Count
Dequeue Count
Enqueue Count
Size
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[ID857fe755c17bc31bc55ddee376107680]]
= Apache Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

This data collector allows collection of data from Apache servers in your environment.

.Pre-requisites

- * You must have your Apache HTTP Server set up and properly running
- * You must have sudo or administrator permissions on your agent host/VM
- * Typically, the Apache <code>_mod_status_</code> module is configured to expose a page at the <code>'/server-status?auto'</code> location of the Apache server. The

ExtendedStatus option must be enabled in order to collect all available fields. For information about how to configure your server, see the Apache module documentation:

https://httpd.apache.org/docs/2.4/mod/mod status.html#enable

== Installation

. From *Observability > Collectors*, click *+Data Collector*. Choose Apache.

+

Select the Operating System or Platform on which the Telegraf agent is installed.

- . If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click _Show Instructions_ to expand the
- xref:{relative_path}task_config_telegraf_agent.html[Agent installation]
 instructions.
- . Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the *+ Agent Access Key* button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- . Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

image:ApacheDCConfigLinux.png[Apache configuration]

== Setup

Telegraf's plugin for Apache's HTTP Server relies on the 'mod_status' module to be enabled. When this is enabled, Apache's HTTP Server will expose an HTML endpoint that can be viewed on your browser or scraped for extraction of status of all Apache's HTTP Server configuration.

=== Compatibility:

```
Configuration was developed against Apache's HTTP Server version 2.4.38.

=== Enabling mod_status:
Enabling and exposing the 'mod_status' modules involves two steps:

* Enabling module

* Exposing stats from module

=== Enabling module:
The loading of modules is controlled by the config file under
'/usr/local/apache/conf/httpd.conf'. Edit the config file and uncomment
the following lines:

LoadModule status_module modules/mod_status.so
```

Include conf/extra/httpd-info.conf

```
=== Exposing stats from module:
```

<Location /server-status>

The exposing of 'mod_status' is controlled by the config file under '/usr/local/apache2/conf/extra/httpd-info.conf'. Make sure you have the following in that configuration file (at least, other directives will be there):

Allow server status reports generated by mod_status, with the URL of http://servername/server-status

```
SetHandler server-status
</Location>

#
# ExtendedStatus controls whether Apache will generate "full" status
# information (ExtendedStatus On) or just basic information (ExtendedStatus
# Off) when the "server-status" handler is called. The default is Off.
#
ExtendedStatus On
```

```
For detailed instructions on the 'mod_status' module, see the link:https://httpd.apache.org/docs/2.4/mod/mod_status.html#enable[Apache documentation] == Objects and Counters
```

```
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|Object:|Identifiers:|Attributes: |Datapoints:
|Apache
|Namespace
Server
|Node IP
Node Name
Port
Parent Server Config Generation
Parent Server MPM Generation
Server Uptime
Is Stopping
|Busy Workers
Bytes per Request
Bytes per Second
CPU Children System
CPU Children User
CPU Load
CPU System
CPU User
Asynchronous Connections Closing
Asynchronous Connections Keep Alive
Asynchronous Connections Writing
Connections Total
Duration per Request
Idle Workers
Load Average (last 1m)
Load Average (last 15m)
Load Average (last 5m)
Processes
Requests per Second
Total Accesses
Total Duration
Total KBytes
Scoreboard Closing
Scoreboard DNS Lookups
Scoreboard Finishing
Scoreboard Idle Cleanup
```

```
Scoreboard Keep Alive
Scoreboard Logging
Scoreboard Open
Scoreboard Reading
Scoreboard Sending
Scoreboard Starting
Scoreboard Waiting
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[ID720d49953357b2bf3f836b70064a163a]]
= Consul Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from Consul.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Consul.
If you haven't configured an Agent for collection, you are prompted to
xref:{relative path}task config telegraf agent.html[install an agent] in
your environment.
If you have an agent already configured, select the appropriate Operating
System or Platform and click *Continue*.
. Follow the instructions in the Consul Configuration screen to configure
the data collector. The instructions vary depending on the type of
Operating System or Platform you are using to collect data.
//image:ConsulDCConfigWindows.png[Consul configuration]
```

```
== Setup
Information may be found in the
link:https://www.consul.io/docs/index.html[Consul documentation].
== Objects and Counters for consul
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
|Consul
|Namespace
Check ID
Service Node
|Node IP
Node OS
Node UUID
Node Name
Service Name
Check Name
Service ID
Status
|Critical
Passing
Warning
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[ID5bd25d9ed944ea316f61e423b818ad9b]]
= Couchbase Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
```

```
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from Couchbase.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Couchbase.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:CouchbaseDCConfigWindows.png[Couchbase configuration]
== Setup
Information may be found in the
link:https://docs.couchbase.com/home/index.html[Couchbase documentation].
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|Object:|Identifiers:|Attributes: |Datapoints:
|Couchbase Node
```

```
|Namespace
Cluster
Couchbase Node Hostname
|Node Name
Node IP
|Memory Free
Memory Total
|Couchbase Bucket
|Namespace
Bucket
Cluster
|Node Name
Node IP
|Data Used
Data Fetches
Disk Used
Item Count
Memory Used
Operations Per Second
Quota Used
|===
== Troubleshooting
Additional information may be found from the
xref:{relative_path}concept_requesting_support.html[Support] page.
[[IDd09fbdab98bf79ff55e3a2af1a26ddf6]]
= CouchDB Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
```

```
Cloud Insights uses this data collector to gather metrics from CouchDB.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
CouchDB.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:CouchDBDCConfigLinux.png[CouchDB configuration]
== Setup
Information may be found in the
link:http://docs.couchdb.org/en/stable/[CouchDB documentation].
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
| CouchDB
|Namespace
Server
|Node Name
Node IP
```

```
|Authentication Cache Hits
Authentication Cache Miss
Database Reads
Database Writes
Databases Open
Open OS Files
Max Request Time
Min Request Time
Httpd Request Methods Copy
Httpd Request Methods Delete
Httpd Request Methods Get
Httpd Request Methods Head
Httpd Request Methods Post
Httpd Request Methods Put
Status Codes 200
Status Codes 201
Status Codes 202
Status Codes 301
Status Codes 304
Status Codes 400
Status Codes 401
Status Codes 403
Status Codes 404
Status Codes 405
Status Codes 409
Status Codes 412
Status Codes 500
|===
== Troubleshooting
Additional information may be found from the
xref:{relative_path}concept requesting support.html[Support] page.
[[ID00534640191ba34dba8046b70c9fdcac]]
= Docker Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
```

```
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from Docker.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Docker.
If you haven't configured an Agent for collection, you are prompted to
xref: {relative path}task config telegraf agent.html[install an agent] in
your environment.
If you have an agent already configured, select the appropriate Operating
System or Platform and click *Continue*.
. Follow the instructions in the Docker Configuration screen to configure
the data collector. The instructions vary depending on the type of
Operating System or Platform you are using to collect data.
//The example below shows the instructions for Linux:
image:DockerDCConfigLinux.png[Docker configuration]
== Setup
The Telegraf input plugin for Docker collects metrics through a specified
UNIX socket or a TCP endpoint.
=== Compatibility
Configuration was developed against Docker version 1.12.6.
=== Setting Up
=== Accessing Docker through a UNIX socket
If the Telegraf agent is running on baremetal, add the telegraf Unix user
to the docker Unix group by running the following:
sudo usermod -aG docker telegraf
If the Telegraf agent is running within a Kubernetes pod, expose the
Docker Unix socket by mapping the socket into the pod as a volume and then
mounting that volume to /var/run/docker.sock. For example, add the
following to the PodSpec:
```

```
volumes:
- name: docker-sock
hostPath:
path: /var/run/docker.sock
type: File
Then, add the following to the Container:
____
volumeMounts:
- name: docker-sock
mountPath: /var/run/docker.sock
Note that the Cloud Insights installer provided for the Kubernetes
platform takes care of this mapping automatically.
=== Access Docker through a TCP endpoint
By default, Docker uses port 2375 for unencrypted access and port 2376 for
encrypted access.
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
|Docker Engine
|Namespace
Docker Engine
|Node Name
Node IP
Node UUID
Node OS
Kubernetes Cluster
Docker Version
Unit
```

|Memory Containers Containers Paused Containers Running Containers Stopped CPUs Go Routines Images Listener Events Used File Descriptors Data Available Data Total Data Used Metadata Available Metadata Total Metadata Used Pool Blocksize |Docker Container |Namespace Container Name Docker Engine |Kubernetes Container Hash Kubernetes Container Ports Kubernetes Container Restart Count Kubernetes Container Termination Message Path Kubernetes Container Termination Message Policy Kubernetes Pod Termination Grace Period Container Image Container Status Container Version Node Name Kubernetes Container Log Path Kubernetes Container Name Kubernetes Docker Type Kubernetes Pod Name Kubernetes Pod Namespace Kubernetes Pod UID Kubernetes Sandbox ID Node IP Node UUID Docker Version Kubernetes IO Config Seen

Kubernetes IO Config Source OpenShift IO SCC Kubernetes Description Kubernetes Display Name OpenShift Tags Kompose Service Pod Template Hash Controller Revision Hash Pod Template Generation License Schema Build Date Schema License Schema Name Schema URL Schema VCS URL Schema Vendor Schema Version Schema Schema Version Maintainer Customer Pod Kubernetes StatefulSet Pod Name Tenant Webconsole Architecture Authoritative Source URL Build Date RH Build Host RH Component Distribution Scope Install Release Run Summary Uninstall VCS Ref VCS Type Vendor Version Health Status Container ID |Memory Active Anonymous Memory Active File Memory Cache Memory Hierarchical Limit Memory Inactive Anonymous Memory Inactive File

Memory Limit Memory Mapped File Memory Max Usage Memory Page Fault Memory Page Major Fault Memory Paged In Memory Paged Out Memory Resident Set Size Memory Resident Set Size Huge Memory Total Active Anonymous Memory Total Active File Memory Total Cache Memory Total Inactive Anonymous Memory Total Inactive File Memory Total Mapped File Memory Total Page Fault Memory Total Page Major Fault Memory Total Paged In Memory Total Paged Out Memory Total Resident Set Size Memory Total Resident Set Size Huge Memory Total Unevictable Memory Unevictable Memory Usage Memory Usage Percent Exit Code OOM Killed PID Started At Failing Streak |Docker Container Block IO |Namespace Container Name Device Docker Engine |Kubernetes Container Hash Kubernetes Container Ports Kubernetes Container Restart Count Kubernetes Container Termination Message Path Kubernetes Container Termination Message Policy Kubernetes Pod Termination Grace Period Container Image

Container Status Container Version Node Name Kubernetes Container Log Path Kubernetes Container Name Kubernetes Docker Type Kubernetes Pod Name Kubernetes Pod Namespace Kubernetes Pod UID Kubernetes Sandbox ID Node IP Node UUID Docker Version Kubernetes Config Seen Kubernetes Config Source OpenShift SCC Kubernetes Description Kubernetes Display Name OpenShift Tags Schema Schema Version Pod Template Hash Controller Revision Hash Pod Template Generation Kompose Service Schema Build Date Schema License Schema Name Schema Vendor Customer Pod Kubernetes StatefulSet Pod Name Tenant Webconsole Build Date License Vendor Architecture Authoritative Source URL RH Build Host RH Component Distribution Scope Install Maintainer

Release Run Summary Uninstall

VCS Ref VCS Type Version Schema URL Schema VCS URL Schema Version Container ID | IO Service Bytes Recursive Async IO Service Bytes Recursive Read IO Service Bytes Recursive Sync IO Service Bytes Recursive Total IO Service Bytes Recursive Write IO Serviced Recursive Async IO Serviced Recursive Read IO Serviced Recursive Sync IO Serviced Recursive Total IO Serviced Recursive Write |Docker Container Network |Namespace Container Name Network Docker Engine |Container Image Container Status Container Version Node Name Node IP Node UUID Node OS K8s Cluster Docker Version Container ID |RX Dropped RX Bytes RX Errors RX Packets TX Dropped TX Bytes TX Errors TX Packets

| Docker Container CPU |Namespace Container Name CPU Docker Engine |Kubernetes Container Hash Kubernetes Container Ports Kubernetes Container Restart Count Kubernetes Container Termination Message Path Kubernetes Container Termination Message Policy Kubernetes Pod Termination Grace Period Kubernetes Config Seen Kubernetes Config Source OpenShift SCC Container Image Container Status Container Version Node Name Kubernetes Container Log Path Kubernetes Container name Kubernetes Docker Type Kubernetes Pod Name Kubernetes Pod Namespace Kubernetes Pod UID Kubernetes Sandbox ID Node IP Node UUID Node OS Kubernetes Cluster Docker Version Kubernetes Description Kubernetes Display Name OpenShift Tags Schema Version Pod Template Hash Controller Revision Hash Pod Template Generation Kompose Service Schema Build Date Schema License Schema Name Schema Vendor Customer Pod

```
Kubernetes StatefulSet Pod Name
Tenant
Webconsole
Build Date
License
Vendor
Architecture
Authoritative Source URL
RH Build Host
RH Component
Distribution Scope
Install
Maintainer
Release
Run
Summary
Uninstall
VCS Ref
VCS Type
Version
Schema URL
Schema VCS URL
Schema Version
Container ID
|Throttling Periods
Throttling Throttled Periods
Throttling Throttled Time
Usage In Kernel Mode
Usage In User Mode
Usage Percent
Usage System
Usage Total
|===
== Troubleshooting
[cols=2*, options="header", cols"50,50"]
|Problem:|Try this:
|I do not see my Docker metrics in Cloud Insights after following the
instructions on the configuration page.
|Check the Telegraf agent logs to see if it reports the following error:
```

E! Error in plugin [inputs.docker]: Got permission denied while trying to connect to the Docker daemon socket If it does, take the necessary steps to provide the Telegraf agent access to the Docker Unix socket as specified above. |=== Additional information may be found from the xref:{relative path}concept requesting support.html[Support] page. [[ID69d5405fbb149eb19c40d7c245cde45a]] = Elasticsearch Data Collector :toc: macro :hardbreaks: :toclevels: 1 :icons: font :linkattrs: :relative path: ./ :imagesdir: {root path}{relative path}./media/ [.lead] Cloud Insights uses this data collector to gather metrics from Elasticsearch. . From *Observability > Collectors*, click *+Data Collector*. Choose Elasticsearch. Select the Operating System or Platform on which the Telegraf agent is installed. . If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click Show Instructions to expand the xref:{relative path}task config telegraf agent.html[Agent installation] instructions. . Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the *+ Agent Access Key* button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform. . Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform

```
you are using to collect data.
image:ElasticsearchDCConfigLinux.png[Elasticsearch configuration]
== Setup
Information may be found in the
link:https://www.elastic.co/guide/index.html[Elasticsearch documentation].
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
|Elasticsearch Cluster
|Namespace
Cluster
|Node IP
Node Name
Cluster Status
|Master Node Count
Total Node Count
Filesystem Data Available (bytes)
Filesystem Data Free (bytes)
Filesystem Data Total (bytes)
JVM Threads
OS Allocated Proccessors
OS Available Processors
OS Mem Free (bytes)
OS Mem Free
OS Mem Total (bytes)
OS Mem Used (bytes)
OS Mem Used
Process CPU
Indices Completion Size (bytes)
Indices Count
Indices Docs Count
Indices Docs Deleted
Indices Field Data Evictions
Indices Field Data Memory Size (bytes)
Indices Query Cache Count
```

```
Indices Cache Size
Indices Segments Count
Indices Segments Doc Values Memory (bytes)
Indices Shards Index Primaries Avg
Indices Shards Index Primaries Max
Indices Shards Index Primaries Min
Indices Shards Index Replication Avg
Indices Shards Index Replication Max
Indices Shards Index Replication Min
Indices Shards Avg
Indices Shards Max
Indices Shards Primaries
Indices Shards Replication
Indices Shards Total
Indices Store Size (bytes)
|Elasticsearch Node
|Namespace
Cluster
ES Node ID
ES Node IP
ES Node
|Zone ID
|Machine Learning Enabled
Machine Learning Memory
Machine Learning Max Open Jobs
X-Pack Installed
Breakers Accounting Estimated Size (bytes)
Breakers Accounting Limit Size (bytes)
Breakers Accounting Overhead
Breakers Accounting Tripped
Breakers Field Data Estimated Size (bytes)
Breakers Field Data Limit Size (bytes)
Breakers Field Data Overhead
Breakers Field Data Tripped
Breakers In-Flight Sstimated Size (bytes)
Breakers In-Flight Limit Size (bytes)
Breakers In-Flight Overhead
Breakers In-Flight Tripped
Breakers Parent Estimated Size (bytes)
Breakers Parent Limit Size (bytes)
Breakers Parent Overhead
Breakers Parent Tripped
```

Breakers Request Estimated Size (bytes) Breakers Request Limit Size (bytes) Breakers Request Overhead Breakers Request Tripped Filesystem Data Available (bytes) Filesystem Data Free (bytes) Filesystem Data Total (bytes) Filesystem IO Stats Devices Ops Filesystem IO Stats Devices Read (kb) Filesystem IO Stats Devices Read Ops Filesystem IO Stats Devices Erite (kb) Filesystem IO Stats Devices Write Ops Filesystem IO Stats Total Ops Filesystem IO Stats Total Read (kb) Filesystem IO Stats Read Ops Filesystem IO Stats Total Write (kb) Filesystem IO Stats Write Ops Filesystem Least Usage Estimate Available (bytes) Filesystem Least Usage Estimate Total (bytes) Filesystem Least Usage Used Disk Filesystem Most Usage Estimate Available (bytes) Filesystem Most Usage Estimate Total (bytes) Filesystem Most Usage Used Disk Filesystem Total Available (bytes) Filesystem Total Free (bytes) Filesystem Total (bytes) Indices Completion Size (bytes) Indices Docs Count Indices Docs Deleted Indices Field Data Evictions Indices Field Data Memory Size (bytes) Indices Flush Periodic Indices Flush Total Indices Flush Total Time Indices Get Current Indices Get Exists Time Indices Get Exists Total Indices Get Total Indices Indexing Delete Total Indices Indexing Index Total Indices Indexing Noop Update Total Indices Indexing Throttle Time HTTP Current Open HTTP Total Opened JVM Buffer Pool Direct Count JVM Classes Current Loaded Count

```
JVM GC Collectors Old Collection Count
JVM Mem Heap Committed (bytes)
OS CPU Load Average 15m
OS CPU
OS Mem Free (bytes)
OS Swap Free (bytes)
Process CPU
Process CPU Total
Process Max File Descriptors
Process Mem Total Virtual (bytes)
Thread Pool Analyze Active
Thread Pool Analyze Completed
Thread Pool Analyze Largest
Thread Pool Analyze Queue
Thread Pool Analyze Rejected
Thread Pool Analyze Threads
Thread Pool Fetch Shard Started Active
Thread Pool Fetch Shard Started Completed
Thread Pool Fetch Shard Started Largest
Thread Pool Fetch Shard Started Queue
Thread Pool Fetch Shard Started Rejected
Thread Pool Fetch Shard Started Shreads
Thread Pool Fetch Shard Store Active
Thread Pool Fetch Shard Store Completed
Transport RX (per sec)
Transport RX Bytes (per sec)
Transport Server Open
Transport TX (per sec)
Transport TX Bytes (per sec)
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[ID153ce3ee29f9380043c8501881dc2d6f]]
= Flink Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
```

```
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
Cloud Insights uses this data collector to gather metrics from Flink.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Flink.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:FlinkDCConfigWindows.png[Flink configuration]
== Setup
A full Flink deployment involves the following components:
JobManager: The Flink primary system. Coordinates a series of
TaskManagers. In a High Availability setup, system will have more than one
JobManager.
TaskManager: This is where Flink operators are executed.
The Flink plugin is based on the telegraf's Jolokia plugin. As such as a
requirement to gather info from all Flink components, JMX needs to be
configured and exposed via Jolokia on all components.
=== Compatibility
Configuration was developed against Flink version 1.7.0.
```

=== Setting Up ==== Jolokia Agent Jar For all individual components, a version the Jolokia agent jar file must be downloaded. The version tested against was link:https://jolokia.org/download.html[Jolokia agent 1.6.0]. Instructions below assume that downloaded jar file (jolokia-jvm-1.6.0agent.jar) is placed under location '/opt/flink/lib/'. ==== JobManager To configure JobManager to expose the Jolokia API, you can setup the following environment variable on your nodes then restart the JobManager: export FLINK ENV JAVA OPTS="-javaagent:/opt/flink/lib/jolokia-jvm-1.6.0agent.jar=port=8778,host=0.0.0.0" You can choose a different port for Jolokia (8778). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. ==== TaskManager To configure TaskManager(s) to expose the Jolokia API, you can setup the following environment variable on your nodes then restart the TaskManager: export FLINK ENV JAVA OPTS="-javaagent:/opt/flink/lib/jolokia-jvm-1.6.0agent.jar=port=8778,host=0.0.0.0" You can choose a different port for Jolokia (8778). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. == Objects and Counters The following objects and their counters are collected: [cols="<.<,<.<,<.<"] |Object:|Identifiers:|Attributes: |Datapoints: |Flink Task Manager |Cluster Namespace Server

```
|Node Name
Task Manager ID
Node IP
|Network Available Memory Segments
Network Total Memory Segments
Garbage Collection PS MarkSweep Count
Garbage Collection PS MarkSweep Time
Garbage Collection PS Scavenge Count
Garbage Collection PS Scavenge Time
Heap Memory Committed
Heap Memory Init
Heap Memory Max
Heap Memory Used
Thread Count Daemon
Thread Count Peak
Thread Count
Thread Count Total Started
|Flink Job
|Cluster
Namespace
server
Job ID
|Node Name
Job Name
Node IP
Last Checkpoint External Path
Restarting Time
IDowntime
Full Restarts
Last Checkpoint Alignment Buffered
Last Checkpoint Duration
Last Checkpoint Size
Number of Completed Checkpoints
Number of Failed Checkpoints
Number of in Progress Checkpoints
Number of Checkpoints
Uptime
|Flink Job Manager
```

|Cluster Namespace Server |Node Name Node IP |Garbage Collection PS MarkSweep Count Garbage Collection PS MarkSweep Time Garbage Collection PS Scavenge Count Garbage Collection PS Scavenge Time Heap Memory Committed Heap Memory Init Heap Memory Max Heap Memory Used Number Registered Task Managers Number Running Jobs Task Slots Available Task Slots Total Thread Count Daemon Thread Count Peak Thread Count Thread Count Total Started |Flink Task |Cluster Namespace Job ID Task ID Server Node Name Job Name Sub Task Index Task Attempt ID Task Attempt Number Task Name Task Manager ID Node IP Current Input Watermark |Buffers In Pool Usage Buffers In Queue Length Buffers Out Pool Usage Buffers Out Queue Length

Number Buffers In Local Number Bufffers In Local Per Second Count Number Buffers in Local Per Second Rate Number Buffers In Remote Number Buffers In Remote Per Second Count Number Buffers In Remote Per Second Rate Number Buffers Out Number Buffers Out Per Second Count Number Buffers Out Per Second Rate Number Bytes In Local Number Bytes In Local Per Second Count Number Bytes In Local Per Second Rate Number Bytes In Remote Number Bytes In Remote Per Second Count Number Bytes In Remote Per Second Rate Number Bytes Out Number Bytes Out Per Second Count Number Bytes Out Per Second Rate Number Records In Number Records In Per Second Count Number Records In Per Second Rate Number Records Out Number Records Out Per Second Count Number Records Out Per Second Rate |Flink Task Operator |Cluster Namespace Job ID Operator ID Task ID Server Node Name Job Name Operator Name Sub Task Index Task Attempt ID Task Attempt Number Task Name Task Manager ID Node IP |Current Input Watermark Current Output Watermark

Number Records In Number Records In Per Second Count Number Records In Per Second Rate Number Records Out Number Records Out Per Second Count Number Records Out Per Second Rate Number Late Records Dropped Assigned Partitions Bytes Consumed Rate Commit Latency Avg Commit Latency Max Commit Rate Commits Failed Commits Succeeded Connection Close Rate Connection Count Connection Creation Rate Count Fetch Latency Avg Fetch Latency Max Fetch Rate Fetch Size Avg Fetch Size Max Fetch Throttle Time Avq Fetch Throttle Time Max Heartbeat Rate Incoming Byte Rate IO Ratio IO Time Avg (ns) IO Wait Ratio IO Wait Time Avg (ns) Join Rate Join Time Avg Last Heartbeat Ago Network IO Rate Outgoing Byte Rate Records Consumed Rate Records Lag Max Records per Request Avg Request Rate Request Size Avg Request Size Max Response Rate Select Rate Sync Rate Sync Time Avg

```
Heartbeat Response Time Max
Join Time Max
Sync Time Max
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[IDebe40bb015da730311f9b41c8487a867]]
= Hadoop Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from Hadoop.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Hadoop.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
```

. Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data. image:HadoopDCConfigLinux-1.png[Hadoop configuration] image:HadoopDCConfigLinux-2.png[Hadoop configuration] == Setup A full Hadoop deployment involves the following components: * NameNode: The Hadoop Distributed File System (HDFS) primary system. Coordinates a series of DataNodes. * Secondary NameNode: a warm failover for the main NameNode. In Hadoop the promotion to NameNode does not occur automatically. Secondary NameNode gathers information from NameNode to be ready to be promoted when needed. * DataNode: Actual owner for data. * ResourceManager: The compute primary system (Yarn). Coordinates a series of NodeManagers. * NodeManager: The resource for compute. Actual location for running of applications. * JobHistoryServer: Responsible for servicing all job history related requests. The Hadoop plugin is based on the telegraf's Jolokia plugin. As such as a requirement to gather info from all Hadoop components, JMX needs to be configured and exposed via Jolokia on all components. === Compatibility Configuration was developed against Hadoop version 2.9.2. === Setting Up ==== Jolokia Agent Jar For all individual components, a version the Jolokia agent jar file must be downloaded. The version tested against was link:https://jolokia.org/download.html[Jolokia agent 1.6.0].

Instructions below assume that downloaded jar file (jolokia-jvm-1.6.0-agent.jar) is placed under location '/opt/hadoop/lib/'.

==== NameNode

To configure NameNode to expose the Jolokia API, you can setup the following in <HADOOP HOME>/etc/hadoop/hadoop-env.sh:

export HADOOP_NAMENODE_OPTS="\$HADOOP_NAMENODE_OPTS -javaagent:/opt/hadoop/lib/jolokia

- -jvm-1.6.0-agent.jar=port=7800,host=0.0.0.0 -Dcom.sun.management.jmxremote
- -Dcom.sun.management.jmxremote.port=8000 -Dcom.sun.management.jmxremote.ssl=false
- -Dcom.sun.management.jmxremote.password.file=\$HADOOP_HOME/conf/jmxremote.password" You can choose a different port for JMX (8000 above) and Jolokia (7800). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. You can use the option '-Dcom.sun.management.jmxremote.authenticate=false' if you don't

want to authenticate. Use at your own risk.

==== Secondary NameNode

To configure the Secondary NameNode to expose the Jolokia API, you can setup the following in <HADOOP HOME>/etc/hadoop/hadoop-env.sh:

export HADOOP_SECONDARYNAMENODE_OPTS="\$HADOOP_SECONDARYNAMENODE_OPTS

- -javaagent:/opt/hadoop/lib/jolokia-jvm-1.6.0-agent.jar=port=7802,host=0.0.0.0
- -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=8002
- -Dcom.sun.management.jmxremote.ssl=false
- -Dcom.sun.management.jmxremote.password.file=\$HADOOP_HOME/conf/jmxremote.password" You can choose a different port for JMX (8002 above) and Jolokia (7802). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. You can use the option '-Dcom.sun.management.jmxremote.authenticate=false' if you don't want to authenticate. Use at your own risk.

==== DataNode

To configure the DataNodes to expose the Jolokia API, you can setup the following in <hadoop_HOME>/etc/hadoop/hadoop-env.sh:

export HADOOP DATANODE OPTS="\$HADOOP DATANODE OPTS -javaagent:/opt/hadoop/lib/jolokia-jvm

- -1.6.0-agent.jar=port=7801,host=0.0.0.0 -Dcom.sun.management.jmxremote
- -Dcom.sun.management.jmxremote.port=8001 -Dcom.sun.management.jmxremote.ssl=false
- -Dcom.sun.management.jmxremote.password.file=\$HADOOP_HOME/conf/jmxremote.password" You can choose a different port for JMX (8001 above) and Jolokia (7801). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. You can use the option '-Dcom.sun.management.jmxremote.authenticate=false' if you don't want to authenticate. Use at your own risk.

==== ResourceManager

To configure the ResourceManager to expose the Jolokia API, you can setup the following in <HADOOP HOME>/etc/hadoop/hadoop-env.sh:

export YARN RESOURCEMANAGER OPTS="\$YARN RESOURCEMANAGER OPTS

- -javaagent:/opt/hadoop/lib/jolokia-jvm-1.6.0-agent.jar=port=7803,host=0.0.0.0
- -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=8003
- -Dcom.sun.management.jmxremote.ssl=false
- -Dcom.sun.management.jmxremote.password.file=\$HADOOP_HOME/conf/jmxremote.password"

You can choose a different port for JMX (8003 above) and Jolokia (7803). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. You can use the option '-Dcom.sun.management.jmxremote.authenticate=false' if you don't want to authenticate. Use at your own risk.

==== NodeManager

To configure the NodeManagers to expose the Jolokia API, you can setup the following in <HADOOP HOME>/etc/hadoop/hadoop-env.sh:

export YARN_NODEMANAGER_OPTS="\$YARN_NODEMANAGER_OPTS -javaagent:/opt/hadoop/lib/jolokia -jvm-1.6.0-agent.jar=port=7804,host=0.0.0.0 -Dcom.sun.management.jmxremote

- -Dcom.sun.management.jmxremote.port=8004 -Dcom.sun.management.jmxremote.ssl=false
- -Dcom.sun.management.jmxremote.password.file=\$HADOOP_HOME/conf/jmxremote.password" You can choose a different port for JMX (8004 above) and Jolokia (7804). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. You can use the option '-Dcom.sun.management.jmxremote.authenticate=false' if you don't want to authenticate. Use at your own risk.

==== JobHistoryServer

To configure the JobHistoryServer to expose the Jolokia API, you can setup the following in <HADOOP HOME>/etc/hadoop/hadoop-env.sh:

export HADOOP_JOB_HISTORYSERVER_OPTS="\$HADOOP_JOB_HISTORYSERVER_OPTS

- -javaagent:/opt/hadoop/lib/jolokia-jvm-1.6.0-agent.jar=port=7805,host=0.0.0.0
- -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=8005
- -Dcom.sun.management.jmxremote.password.file=\$HADOOP_HOME/conf/jmxremote.password" You can choose a different port for JMX (8005 above) and Jolokia (7805). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. You can use the option '-Dcom.sun.management.jmxremote.authenticate=false' if you don't want to authenticate. Use at your own risk.

```
== Objects and Counters

The following objects and their counters are collected:

[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:

|Hadoop Secondary NameNode

|Cluster
|Namespace
|Server
|Node Name
|Node IP
|Compile Info
|Version
|GC Count
```

```
GC Copies Count
GC Marks Sweep Compact Count
GC Number Info Threshold Exceeded
GC Number Warning Threshold Exceeded
GC Time
GC Copy Time
GC Marks Sweep Compact Time
GC Total Extra Sleep Time
Logs Error Count
Logs Fatal Count
Logs Info Count
Logs Warn Count
Memory Heap Committed
Memory Heap Max
Memory Heap Used
Memory Max
Memory Non Heap Committed
Memory Non Heap Max
Memory Non Heap Used
Threads Blocked
Threads New
Threads Runnable
Threads Terminated
Threads Timed Waiting
Threads Waiting
| Hadoop NodeManager
|Cluster
Namespace
Server
|Node Name
Node IP
|Containers Allocated
Memory Allocate
Memory Allocated Oportunistic
Virtual Cores Allocated Oportunistic
Virtual Cores Allocated
Memory Available
Virtual Cores Available
Directories Bad Local
Directories Bad Log
Cache Size Before Clean
Container Launch Duration Avg Time
```

Container Launch Duration Number Of Operations Containers Completed Containers Failed Containers Initing Containers Killed Containers Launched Containers Reiniting ContaIners Rolled Back on Failure Containers Running Disk Utilization Good Local Directories Disk Utilization Good Log Directories Bytes Deleted Private Bytes Deleted Public Containers Running Opportunistic Bytes Deleted Total Shuffle Connections Shuffle Output Bytes Shuffle Outputs Failed Shuffle Outputs Ok GC Count GC Copies Count GC Marks Sweep Compact Count GC Number Info Threshold Exceeded GC Number Warning Threshold Exceeded GC Time GC Copy Time GC Marks Sweep Compact Time GC Total Extra Sleep Time Logs Error Count Logs Fatal Count Logs Info Count Logs Warn Count Memory Heap Committed Memory Heap Max Memory Heap Used Memory Max Memory Non Heap Committed Memory Non Heap Max Memory Non Heap Used Threads Blocked Threads New Threads Runnable Threads Terminated Threads Timed Waiting Threads Waiting

| Hadoop Resource Manager |Cluster Namespace Server |Node Name Node IP |ApplicationMaster Launch Delay Avg ApplicationMaster Launch Delay Number ApplicationMaster Register Delay Avg ApplicationMaster Register Delay Number NodeManager Active Number NodeManager Decomissioned Number NodeManager Decomissioning Number NodeManager Lost Number NodeManager Rebooted Number NodeManager Shutdown Number NodeManager Healthy Number NodeManager Memory Limit NodeManager Virtual Cores Limit Used Capacity Active Applications Active Users Aggregate Containers Allocated Aggregate Containers Preempted Aggregate Containers Released Aggregate Memory Seconds Preempted Aggregate Node Local Containers Allocated Aggregate Off Switch Containers Allocated Aggregate Ack Local Containers Allocated Aggregate Virtual Cores Seconds Preempted Containers Allocated Memory Allocated Virtual Cores Allocated Application Attempt First Container Allocation Delay Avg Time Application Attempt First Container Allocation Delay Number Applications Completed Applications Failed Applications Killed Applications Pending Applications Running Applications Submitted Memory Available Virtual Cores Available

Containers Pending Memory Pending Virtual Cores Pending Containers Reserved Memory Reserved Virtual Cores Reserved Memory ApplicationMaster Used Virtual Cores ApplicationMaster Used Capacity Used GC Count GC Copies Count GC Marks Sweep Compact Count GC Number Info Threshold Exceeded GC Number Warning Threshold Exceeded GC Time GC Copy Time GC Marks Sweep Compact Time GC Total Extra Sleep Time Logs Error Count Logs Fatal Count Logs Info Count Logs Warn Count Memory Heap Committed Memory Heap Max Memory Heap Used Memory Max Memory Non Heap Committed Memory Non Heap Max Memory Non Heap Used Threads Blocked Threads New Threads Runnable Threads Terminated Threads Timed Waiting Threads Waiting | Hadoop DataNode |Cluster Namespace Server |Node Name Node IP Cluster ID Version

|Transceiver Count Transmits in Progress Cache Capacity Cache Used Capacity DFS Used Estimated Capacity Lost Total Last Volume Failure Rate Blocks Number Cached Blocks Number Failed to Cache Blocks Number Failed to Uncache Volumes Number Failed Capacity Remaining GC Count GC Copies Count GC Marks Sweep Compact Count GC Number Info Threshold Exceeded GC Number Warning Threshold Exceeded GC Time GC Copy Time GC Marks Sweep Compact Time GC Total Extra Sleep Time Logs Error Count Logs Fatal Count Logs Info Count Logs Warn Count Memory Heap Committed Memory Heap Max Memory Heap Used Memory Max Memory Non Heap Committed Memory Non Heap Max Memory Non Heap Used Threads Blocked Threads New Threads Runnable Threads Terminated Threads Timed Waiting Threads Waiting | Hadoop NameNode |Cluster Namespace Server

|Node Name Node IP Transaction ID Last Written Time Since Last Loaded Edits HA State File System State Block Pool ID Cluster ID Compile Info Distinct Version Count Version |Block Capacity Blocks Total Capacity Total Capacity Used Capacity Used Non DFS Blocks Corrupt Estimated Capacity Lost Total Blocks Excess Heartbeats Expired Files Total File System Lock Queue Length Blocks Missing Blocks Missing Replication with Factor One Clients Active Data Nodes Dead Data Nodes Decommissioning Dead Data Nodes Decommissioning Live Data Nodes Decomissioning Encryption Zones Number Data Nodes Entering Maintenance Files Under Construction Data Nodes Dead in Maintenance Data Nodes Live in Maintenance Data Nodes Live Storages Stale Replication Pending Timeouts Data Node Message Pending Blocks Pending Deletion Blocks Pending Replication Blocks Misreplicated Postponed Blocks Scheduled Replication Snapshots Snapshottable Directories

Data Nodes Stale Files Total Load Total Sync Count Total Transactions Since Last Checkpoint Transactions Since Last Log Roll Blocks Underreplicated Volume Failures Total Sync Times Total Objects Max Operations Block Add Operations Allow Snapshots Operations Block Batched Operations Block Queued Operations Block Received and Deleted Operations Report Avg Time Operations Block Report Number Cache Report Avg Time Cache Report Number Operations Create File Operations Create Snapshots Operations Create SymLink Operations Delete File Operations Delete Snapshot Operations Disallow Snapshot Operations File In/Out Files Appended Files Created Files Deleted Files Listing Files Renamed Files Truncated File System Load Time Operations Generate EDEK Avg Time Operations Generate EDEK Operations Get Additional Data Node Blocks Get Locations Get Edit Avg Time Get Edit Number Get Image Avg Time Get Image Number Operations Get Link Target Operations Get Listing Operations List Snapshottable Dir Replication Not Scheduled Number Put Image Avg Time

Put Image Number Operations Rename Snapshots Resource Check Time Avg Time Resource Check Time Number Safe Mode Time Operations Snapshot Diff Report Operations Storage Block Report Replication Successful Sync Avg Time Operations Sync Number Replication Timeout Operations Total Transaction Avg Time Transaction Batchd In Sync Transaction Number EDEK Warmup Time Avg EDEK Warmup Number Block Pool Used Space Cache Capacity Cache Used Capacity Free Block Pool Used Percent Percent Remaining Percent Used Threads GC Count GC Copies Count GC Marks Sweep Compact Count GC Number Info Threshold Exceeded GC Number Warning Threshold Exceeded GC Time GC Copy Time GC Marks Sweep Compact Time GC Total Extra Sleep Time Logs Error Count Logs Fatal Count Logs Info Count Logs Warn Count Memory Heap Committed Memory Heap Max Memory Heap Used Memory Max Memory Non Heap Committed Memory Non Heap Max Memory Non Heap Used Threads Blocked

```
Threads New
Threads Runnable
Threads Terminated
Threads Timed Waiting
Threads Waiting
| Hadoop JobHistoryServer
|Cluster
Namespace
Server
|Node Name
Node IP
|GC Count
GC Copies Count
GC Marks Sweep Compact Count
GC Number Info Threshold Exceeded
GC Number Warning Threshold Exceeded
GC Time
GC Copy Time
GC Marks Sweep Compact Time
GC Total Extra Sleep Time
Logs Error Count
Logs Fatal Count
Logs Info Count
Logs Warn Count
Memory Heap Committed
Memory Heap Max
Memory Heap Used
Memory Max
Memory Non Heap Committed
Memory Non Heap Max
Memory Non Heap Used
Threads Blocked
Threads New
Threads Runnable
Threads Terminated
Threads Timed Waiting
Threads Waiting
|===
== Troubleshooting
```

```
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[ID4a79e1763cf52d863c5a1d60747427ce]]
= HAProxy Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from HAProxy.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
HAProxy.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:HAProxyDCConfigLinux.png[HAProxy configuration]
== Setup
Telegraf's plugin for HAProxy relies on HAProxy Stats enablement. This is
a configuration built into HAProxy but it is not enabled out of the box.
```

When enabled, HAProxy will expose an HTML endpoint that can be viewed on your browser or scraped for extraction of status of all HAProxy configurations.

=== Compatibility:

Configuration was developed against HAProxy version 1.9.4.

=== Setting Up:

To enable stats, edit your haproxy configuration file and add the the following lines after the 'defaults' section, using your own user/password and/or haproxy URL:

stats enable stats auth myuser:mypassword stats uri /haproxy?stats

The following is a simplified example configuration file with stats enabled:

global daemon maxconn 256

defaults
mode http
stats enable
stats uri /haproxy?stats
stats auth myuser:mypassword
timeout connect 5000ms
timeout client 50000ms
timeout server 50000ms

frontend http-in bind *:80 default backend servers

frontend http-in9080 bind *:9080 default backend servers 2

backend servers server server1 10.128.0.55:8080 check ssl verify none server server2 10.128.0.56:8080 check ssl verify none

backend servers_2 server server3 10.128.0.57:8080 check ssl verify none server server4 10.128.0.58:8080 check ssl verify none

For complete and up to date instructions, see the link:https://cbonte.github.io/haproxy-dconv/1.8/configuration.html#4stats%20enable[HAProxy documentation]. == Objects and Counters The following objects and their counters are collected: [cols="<.<,<.<,<.<"] |=== |Object:|Identifiers:|Attributes: |Datapoints: | HAProxy Frontend |Namespace Address Proxy |Node IP Node Name Proxy ID Mode Process id Sessions Rate Limit Server id Sessions Limit Status |Bytes In Bytes Out Cache Hits Cache Lookups Compression Bytes Bypassed Compression Bytes In Compression Bytes Out Compression Responses Connection Rate Connection Rate Max Connections Total Requests Denied by Connection Rule Requests Denied by Security Concerns Responses Denied by Security Concerns Requests Denied by Session Rule Requests Errors Responses 1xx

Responses 2xx Responses 3xx Responses 4xx Responses 5xx Responses Other Requests Intercepted Sessions Rate Sessions Rate Max Requests Rate Requests Rate Max Requests Total Sessions Sessions Max Sessions Total Requests Rewrites |HAProxy Server |Namespace Address Proxy Server |Node IP Node Name Check Time to Finish Check Fall Configuration Check Health Value Check Rise Configuration Check Status Proxy ID Last Change Time Last Session Time Mode Process id Server id Status Weight |Active Servers Backup Servers Bytes In Bytes Out Check Downs Check Fails Client Aborts

Connections Connection Average Time Downtime Total Denied Responses Connection Errors Response Errors Responses 1xx Responses 2xx Responses 3xx Responses 4xx Responses 5xx Responses Other Server Selected Total Queue Current Queue Max Queue Average Time Sessions per Second Sessions per Second Max Connection Reuse Response Time Average Sessions Sessions Max Server Transfer Aborts Sessions Total Sessions Total Time Average Requests Redispatches Requests Retries Requests Rewrites | HAProxy Backend |Namespace Address Proxy |Node IP Node Name Proxy ID Last Change Time Last Session Time Mode Process id Server id Sessions Limit Status Weight

|Active Servers Backup Servers Bytes In Bytes Out Cache Hits Cache Lookups Check Downs Client Aborts Compression Bytes Bypassed Compression Bytes In Compression Bytes Out Compression Responses Connections Connection Average Time Downtime Total Requests Denied by Security Concerns Responses Denied by Security Concerns Connection Errors Response Errors Responses 1xx Responses 2xx Responses 3xx Responses 4xx Responses 5xx Responses Other Server Selected Total Queue Current Queue Max Queue Average Time Sessions per Second Sessions per Second Max Requests Total Connection Reuse Response Time Average Sessions Sessions Max Server Transfer Aborts Sessions Total Sessions Total Time Average Requests Redispatches Requests Retries Requests Rewrites |=== == Troubleshooting

```
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[IDc3d4c1df6da4575c11be4f68327e8f59]]
= JVM Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from JVM.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose JVM.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:JVMDCConfigLinux.png[JVM configuration]
== Setup
Information may be found in
link:https://docs.oracle.com/javase/specs/jvms/se12/html/index.html[JVM
```

```
documentation].
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|Object:|Identifiers:|Attributes: |Datapoints:
| JVM
|Namespace
JVM
|OS Architecture
OS Name
OS Version
Runtime Specification
Runtime Specification Vendor
Runtime Specification Version
Uptime
Runtime VM Name
Runtime VM Vendor
Runtime VM Version
Node Name
Node IP
|Class Loaded
Class Loaded Total
Class Unloaded
Memory Heap Committed
Memory Heap Init
Memory Heap Used Max
Memory Heap Used
Memory Non Heap Committed
Memory Non Heap Init
Memory Non Heap Max
Memory Non Heap Used
Memory Objects Pending Finalization
OS Processors Available
OS Committed Virtual Memory Size
OS Free Physical Memory Size
OS Free Swap Space Size
OS Max File Descriptor Count
OS Open File Descriptors Count
```

```
OS Processor CPU Load
OS Processor CPU Time
OS System CPU Load
OS System Load Average
OS Total Physical Memory Size
OS Total Swap Space Size
Thread Daemon Count
Thread Peak Count
Thread Count
Thread Total Started Count
Garbage Collector Copy Collection Count
Garbage Collector Copy Collection Time
Garbage Collector Mark-sweep Collection Count
Garbage Collector Mark-sweep Collection Time
Garbage Collector G1 Old Generation Collection Count
Garbage Collector G1 Old Generation Collection Time
Garbage Collector G1 Young Generation Collection Count
Garbage Collector G1 Young Generation Collection Time
Garbage Collector Concurrent Mark-sweep Collection Count
Garbage Collector Concurrent Mark-sweep Collection Time
Garbage Collector Parallel Collection Count
Garbage Collector Parallel Collection Time
Garbage Collector Parallel Scavenge Mark-sweep Collection Count
Garbage Collector Parallel Scavenge Mark-sweep Collection Time
Garbage Collector Parallel Scavenge Collection Count
Garbage Collector Parallel Scavenge Collection Time
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[ID0472d9d496fee3715d6a96496db536be]]
= Kafka Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
```

[.lead]

Cloud Insights uses this data collector to gather metrics from Kafka.

== Installation

. From *Observability > Collectors*, click *+Data Collector*. Choose Kafka.

+

Select the Operating System or Platform on which the Telegraf agent is installed.

- . If you haven't already installed an Agent for collection, or you wish to install an Agent for a different Operating System or Platform, click _Show Instructions_ to expand the
- xref:{relative_path}task_config_telegraf_agent.html[Agent installation]
 instructions.
- . Select the Agent Access Key for use with this data collector. You can add a new Agent Access Key by clicking the *+ Agent Access Key* button. Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform.
- . Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data.

image:KafkaDCConfigWindows.png[Kafka configuration]

== Setup

The Kafka plugin is based on the telegraf's Jolokia plugin. As such as a requirement to gather info from all Kafka brokers, JMX needs to be configured and exposed via Jolokia on all components.

=== Compatibility

Configuration was developed against Kafka version 0.11.0.2.

=== Setting up

All the instructions below assume your install location for kafka is '/opt/kafka'. You can adapt instructions below to reflect your install location.

==== Jolokia Agent Jar

A version the Jolokia agent jar file must be link:https://jolokia.org/download.html[downloaded]. The version tested against was Jolokia agent 1.6.0.

Instructions below assume that the downloaded jar file (jolokia-jvm-1.6.0agent.jar) is placed under the location '/opt/kafka/libs/'.

==== Kafka Brokers
To configure Kafka Brokers to expose the Jolokia API, you can add the
following in <KAFKA_HOME>/bin/kafka-server-start.sh, just before the
'kafka-run-class.sh' call:

export JMX_PORT=9999
export RMI_HOSTNAME=hostname -I
export KAFKA_JMX_OPTS="-javaagent:/opt/kafka/libs/jolokia-jvm-1.6.0-agent.jar=port=8778,host=0.0.0.0
-Dcom.sun.management.jmxremote.password.file=/opt/kafka/config/jmxremote.password
-Dcom.sun.management.jmxremote.ssl=false -Djava.rmi.server.hostname=\$RMI_HOSTNAME
-Dcom.sun.management.jmxremote.rmi.port=\$JMX_PORT"

Note that example above is using 'hostname -I' to setup the 'RMI_HOSTNAME'

environment variable. In multiple IP machines, this will need to be tweaked to gather the IP you care about for RMI connections. You can choose a different port for JMX (9999 above) and Jolokia (8778). If you have an internal IP to lock Jolokia onto you can replace the "catch all" 0.0.0.0 by your own IP. Notice this IP needs to be accessible from the telegraf plugin. You can use the option '-Dcom.sun.management.jmxremote.authenticate=false' if you don't want to authenticate. Use at your own risk. == Objects and Counters The following objects and their counters are collected: [cols="<.<,<.<,<.<"] |=== |Object:|Identifiers:|Attributes: |Datapoints: |Kafka Broker |Cluster Namespace Broker |Node Name Node IP |Replica Manager Fetcher Max Lag Zookeeper Client Connections Zookeeper Client Connections (15m rate)

```
Zookeeper Client Connections (5m rate)
Zookeeper Client Connections (mean rate)
Zookeeper Client Connections (1m rate)
Replica Manager Partition Count
Thread Count Daemon
Thread Count Peak
Thread Count Current
Thread Count Total Started
Offline Partitions
Produce Requests Total Time (50th Percentile)
Produce Requests Total Time (75th Percentile)
Produce Requests Total Time (95th Percentile)
Produce Requests Total Time (98 Percentile)
Produce Requests Total Time (999th Percentile)
Produce Requests Total Time (99th Percentile)
Produce Requests Total Time
Produce Requests Total Time Max
Produce Requests Total Time Mean
Produce Requests Total Time Min
Produce Requests Total Time Stddev
Replica Manager ISR Shrinks
Replica Manager ISR Shrinks (15m rate)
Replica Manager ISR Shrinks (5m rate)
Replica Manager ISR Shrinks (mean rate)
Replica Manager ISR Shrinks (1m rate)
Request Handler Avg Idle
Request Handler Avg Idle (15m rate)
Request Handler Avg Idle (5m rate)
Request Handler Avg Idle (mean rate)
Request Handler Avg Idle (1m rate)
Garbage Collection G1 Old Generation Count
Garbage Collection G1 Old Generation Time
Garbage Collection G1 Young Generation Count
Garbage Collection G1 Young Generation Time
Zookeeper Read Only Connects
Zookeeper Read Only Connects (15m rate)
Zookeeper Read Only Connects (5m rate)
Zookeeper Read Only Connects (mean rate)
Zookeeper Read Only Connects (1m rate)
Network Processor Avg Idle
Requests Fetch Follower Total Time (50th percentile)
Requests Fetch Follower Total Time (75th percentile)
Requests Fetch Follower Total Time (95th percentile)
Requests Fetch Follower Total Time (98th percentile)
Requests Fetch Follower Total Time (999th percentile)
Requests Fetch Follower Total Time (99th percentile)
```

```
Requests Fetch Follower Total Time
Requests Fetch Follower Total Time Max
Requests Fetch Follower Total Time Mean
Requests Fetch Follower Total Time Min
Requests Fetch Follower Total Time Stddev
Requests Waiting in Produce Purgatory
Network Requests Fetch Consumer
Network Requests Fetch Consumer (5m rate)
Network Requests Fetch Consumer (15m rate)
Network Requests Fetch Consumer (mean rate)
Network Requests Fetch Consumer (1m rate)
Unclean Leader Elections
Unclean Leader Elections (15m rate)
Unclean Leader Elections (5m rate)
Unclean Leader Elections (mean rate)
Unclean Leader Elections (1m rate)
Active Controllers
Heap Memory Committed
Heap Memory Init
Heap Memory Max
Heap Memory Used
Zookeeper Session Expires
Zookeeper Session Expires (15m rate)
Zookeeper Session Expires (5m rate)
Zookeeper Session Expires (mean rate)
Zookeeper Session Expires (1m rate)
Zookeeper Authentication Failures
Zookeeper Authentication Failures (15m rate)
Zookeeper Authentication Failures (5m rate)
Zookeeper Authentication Failures (mean rate)
Zookeeper Authentication Failures (1m rate)
Leader Election Time (50th percentile)
Leader Election Time (75th percentile)
Leader Election Time (95th percentile)
Leader Election Time (98th percentile)
Leader Election Time (999th percentile)
Leader Election Time (99th percentile)
Leader Election Count
Leader Election Time (15m rate)
Leader Election Time (5m rate)
Leader Election Time Max
Leader Election Time Mean
Leader Election Time (mean rate)
Leader Election Time Min
Leader Election Time (1m rate)
Leader Election Time (stddev)
```

```
Network Requests Fetch Follower
Network Requests Fetch Follower (15m rate)
Network Requests Fetch Follower (5m rate)
Network Requests Fetch Follower (mean rate)
Network Requests Fetch Follower (1m rate)
Broker Topic Messages
Broker Topic Messages (15m rate)
Broker Topic Messages (5m rate)
Broker Topic Messages (mean rate)
Broker Topic Messages (1m rate)
Broker Topic Bytes In
Broker Topic Bytes In (15m rate)
Broker Topic Bytes In (5m rate)
Broker Topic Bytes In (mean rate)
Broker Topic Bytes In (1m rate)
Zookeeper Disconnects Count
Zookeeper Disconnects (15m rate)
Zookeeper Disconnects (5m rate)
Zookeeper Disconnects (mean rate)
Zookeeper Disconnects (1m rate)
Network Requests Fetch Consumer Total Time (50th percentile)
Network Requests Fetch Consumer Total Time (75th percentile)
Network Requests Fetch Consumer Total Time (95th percentile)
Network Requests Fetch Consumer Total Time (98th percentile)
Network Requests Fetch Consumer Total Time (999th percentile)
Network Requests Fetch Consumer Total Time (99th percentile)
Network Requests Fetch Consumer Total Time
Network Requests Fetch Consumer Total Time Max
Network Requests Fetch Consumer Total Time Mean
Network Requests Fetch Consumer Total Time Min
Network Requests Fetch Consumer Total Time Stddev
LeaderCount
Requests Waiting in Fetch Purgatory
Broker Topic Bytes Out
Broker Topic Bytes Out (15m rate)
Broker Topic Bytes Out (5m rate)
Broker Topic Bytes Out (mean rate)
Broker Topic Bytes Out (1m rate)
Zookeeper Authentications
Zookeeper Authentications (15m rate)
Zookeeper Authentications (5m rate)
Zookeeper Authentications (mean rate)
Zookeeper Authentications (1m rate)
Requests Produce Count
Requests Produce (15m rate)
Requests Produce (5m rate)
```

```
Requests Produce (mean rate)
Requests Produce (1m rate)
Replica Manager ISR Expands
Replica Manager ISR Expands (15m rate)
Replica Manager ISR Expands (5m rate)
Replica Manager ISR Expands (mean rate)
Replica Manager ISR Expands (1m rate)
Replica Manager Under Replicated Partitions
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[ID933594a6de1603ff8a26d2d21e4d0128]]
= Kibana Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from Kibana.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Kibana.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
```

Best practice: Use a different Agent Access Key only when you want to group data collectors, for example, by OS/Platform. . Follow the configuration steps to configure the data collector. The instructions vary depending on the type of Operating System or Platform you are using to collect data. image:KibanaDCConfigLinux.png[Kibana configuration] == Setup Information may be found in the link:https://www.elastic.co/guide/index.html[Kibana documentation]. == Objects and Counters The following objects and their counters are collected: [cols="<.<,<.<,<.<"] |=== |Object:|Identifiers:|Attributes: |Datapoints: |Kibana |Namespace Address |Node IP Node Name Version Status |Concurrent Connections Heap Max Heap Used Requests per Second Response Time Average Response Time Max Uptime |=== == Troubleshooting Additional information may be found from the xref:{relative path}concept requesting support.html[Support] page.

```
[[IDf703e903ebdcf0ce4f9fd4c95ef8af49]]
= Memcached Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from Memcached.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Memcached.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:MemcachedDCConfigWindows.png[Memcached configuration]
== Setup
Information may be found in the
link:https://github.com/memcached/memcached/wiki[Memcached wiki].
== Objects and Counters
The following objects and their counters are collected:
```

```
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
| Memcached
|Namespace
Server
|Node IP
Node Name
|Accepting Connections
Handled Authentication Requests
Failed Authentications
Bytes Used
Bytes Read (per sec)
Bytes Written (per sec)
CAS Badval
CAS Hits
CAS Misses
Flush Regs (per sec)
Get Reqs (per sec)
Set Reqs (per sec)
Touch Regs (per sec)
Connection Yields (per sec)
Connection Structures
Open Connections
Current Stored Items
Decr Requests Hits (per sec)
Decr Requests Misses (per sec)
Delete Requests Hits (per sec)
Delete Requests Misses (per sec)
Items Evicted
Valid Evictions
Expired Items
Get Hits (per sec)
Get Misses (per sec)
Used Hash Bytes
Hash Is Expanding
Hash Power Level
Incr Requests Hits (per sec)
Incr Requests Misses (per sec)
Server Max Bytes
Listen Disabled Num
```

```
Reclaimed
Worker Threads Count
Total Opened Connections
Total Items Stored
Touch Hits
Touch Misses
Server Uptime
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[IDdfb5ca2eddc053341b0300bf95f52889]]
= MongoDB Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from MongoDB.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
MongoDB.
+
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
```

```
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:MongoDBDCConfigLinux.png[MongoDB configuration]
== Setup
Information may be found in the link:https://docs.mongodb.com/[MongoDB
documentation].
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|Object:|Identifiers:|Attributes: |Datapoints:
|MongoDB
|Namespace
Hostname
|MongoDB Database
|Namespace
Hostname
Database name
|===
== Troubleshooting
Information may be found from the
xref:{relative_path}concept_requesting_support.html[Support] page.
[[ID620cabc86342f799681831c173d09775]]
```

```
= MySQL Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from MySQL.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
MySQL.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative_path}task_config_telegraf_agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:MySQLDCConfigWindows.png[MySQL Configuration]
== Setup
Information may be found in the link:https://dev.mysql.com/doc/[MySQL
documentation].
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
```

```
|Object:|Identifiers:|Attributes: |Datapoints:
|MySQL
|Namespace
MySQL Server
|Node IP
Node Name
|Aborted Clients (per sec)
Aborted Connects (per sec)
RX Bytes (per sec)
TX Bytes (per sec)
Commands Admin (per sec)
Commands Alter Event
Commands Alter Function
Commands Alter Instance
Commands Alter Procedure
Commands Alter Server
Commands Alter Table
Commands Alter Tablespace
Commands Alter User
Commands Analyze
Commands Assign To Keycache
Commands Begin
Commands Binlog
Commands Call Procedure
Commands Change DB
Commands Change Master
Commands Change Repl Filter
Commands Check
Commands Checksum
Commands Commit
Commands Create DB
Commands Create Event
Commands Create Function
Commands Create Index
Commands Create Procedure
Commands Create Server
Commands Create Table
Commands Create Trigger
Commands Create UDF
Commands Create User
Commands Create View
```

```
Commands Dealloc SQL
Connection Errors Accept
Created Tmp Disk Tables
Delayed Errors
Flush Commands
Handler Commit
Innodb Buffer Pool Bytes Data
Key Blocks Not Flushed
Key Read Requests
Key Write Requests
Key Writes
Max Execution Time Exceeded
Max Used Connections
Open Files
Performance Schema Accounts Lost
Prepared Stmt Count
Qcache Free Blocks
Queries
Questions
Select Full Join
Select Full Range Join
Select Range Check
Select Scan
Table Locks Immediate
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[IDf5dc21fc8190cba9efe13d65f2cf9f2b]]
= Netstat Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
```

```
Cloud Insights uses this data collector to gather Netstat metrics.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Netstat.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:NetstatDCConfigWindows.png[Windows Netstat Configuration]
== Setup
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
| Netstat
|Node UUID
|Node IP
Node Name
|===
== Troubleshooting
```

```
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[IDfa6a0e9840003ffb792d321d7fc438c3]]
= Nginx Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from Nginx.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Nginx.
+
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:NginxDCConfigLinux-1.png[Linux Nginx Configuration]
image:NginxDCConfigLinux-2.png[Linux Nginx Configuration]
== Setup
```

```
Nginx metric collection requires that Nginx
link:http://nginx.org/en/docs/http/ngx_http_stub_status_module.html[http_s
tub status module] be enabled.
Additional information may be found in the
link:http://nginx.org/en/docs/[Nginx documentation].
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
|Nginx
|Namespace
Server
|Node IP
Node Name
Port
|Accepts
Active
Handled
Reading
Requests
Waiting
Writing
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[IDa6e4923c3419a22f0ba85c774ad8728b]]
= PostgreSQL Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
```

```
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from PostgreSQL.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
PostgreSQL.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:PostgreSQLDCConfigLinux.png[PostgreSQL configuration]
== Setup
Information may be found in the
link:https://www.postgresql.org/docs/[PostgreSQL documentation].
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
|PostgreSQL Server
|Namespace
```

Database Server |Node Name Node IP |Buffers Allocated Buffers Backend Buffers Backend File Sync Buffers Checkpoint Buffers Clean Checkpoints Sync Time Checkpoints Write Time Checkpoints Requests Checkpoints Timed Max Written Clean |PostgreSQL Database |Namespace Database Server |Database OID Node Name Node IP |Blocks Read Time Blocks Write Time Blocks Hits Blocks Reads Conflicts Deadlocks Client Number Temp Files Bytes Temp Files Number Rows Deleted Rows Fetched Rows Inserted Rows Returned Rows Updated Transactions Committed Transactions Rollbacked |===

```
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[IDde41269bc3f652ea999dfc31460c97dc]]
= Puppet Agent Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from Puppet
Agent.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Puppet.
+
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:PuppetDCConfigWindows.png[Puppet configuration]
```

```
== Setup
Information may be found in the https://puppet.com/docs[Puppet
documentation]
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|Object:|Identifiers:|Attributes: |Datapoints:
| Puppet Agent
|Namespace
Node UUID
|Node Name
Location
Node IP
Version Configstring
Version Puppet
|Changes Total
Events Failure
Events Success
Events Total
Resources Changed
Resources Failed
Resources Failed To Restart
Resources Outofsync
Resources Restarted
Resources Scheduled
Resources Skipped
Resources Total
Time Anchor
Time Configretrieval
Time Cron
Time Exec
Time File
Time Filebucket
Time Lastrun
Time Package
Time Schedule
Time Service
```

```
Time Sshauthorizedkey
Time Total
Time User
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
[[IDbee01232e480e390a4c87d0a969ff80c]]
= Redis Data Collector
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root path}{relative path}./media/
[.lead]
Cloud Insights uses this data collector to gather metrics from Redis.
Redis is an open source, in-memory data structure store used as a
database, cache, and message broker, supporting the following data
structures: strings, hashes, lists, sets, and more.
== Installation
. From *Observability > Collectors*, click *+Data Collector*. Choose
Redis.
Select the Operating System or Platform on which the Telegraf agent is
installed.
. If you haven't already installed an Agent for collection, or you wish to
install an Agent for a different Operating System or Platform, click Show
Instructions to expand the
xref:{relative path}task config telegraf agent.html[Agent installation]
instructions.
. Select the Agent Access Key for use with this data collector. You can
add a new Agent Access Key by clicking the *+ Agent Access Key* button.
Best practice: Use a different Agent Access Key only when you want to
group data collectors, for example, by OS/Platform.
```

```
. Follow the configuration steps to configure the data collector. The
instructions vary depending on the type of Operating System or Platform
you are using to collect data.
image:RedisDCConfigWindows.png[Redis Data Collector Configuration]
== Setup
Information may be found in the link:https://redis.io/documentation[Redis
documentation].
== Objects and Counters
The following objects and their counters are collected:
[cols="<.<,<.<,<.<"]
|===
|Object:|Identifiers:|Attributes: |Datapoints:
|Redis
|Namespace
Server
|===
== Troubleshooting
Additional information may be found from the
xref:{relative path}concept requesting support.html[Support] page.
:leveloffset: -1
[[ID60a4d0044acf46e6ac514daf98f5abcb]]
= Object Icon Reference
:toc: macro
:hardbreaks:
:toclevels: 1
:icons: font
:linkattrs:
:relative_path: ./
:imagesdir: {root_path}{relative_path}./media/
```

```
[.lead]
A quick reference for object icons used in Cloud Insights.
== Infrastructure Icons:
image:Icon Glossary.png[Infrastructure Icon Reference]
== Kubernetes Icons:
image:K8sIconsWithLabels.png[Kubernetes Icon Reference, width=180]
== Kubernetes Network Performance Monitoring and Map Icons:
image: ServiceMap Icons.png [Kubernetes Network Performance Monitoring and
Map, width=640]
:leveloffset: -1
[[ID870611dcda2ee7c98d95c23f5bb785f7]]
= Legal notices
:toc: macro
:hardbreaks:
:icons: font
:linkattrs:
:relative path: ./
:imagesdir: {root_path}{relative_path}./media/
[.lead]
Legal notices provide access to copyright statements, trademarks, patents,
and more.
== Copyright
http://www.netapp.com/us/legal/copyright.aspx[^]
== Trademarks
NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks
page are trademarks of NetApp, Inc. Other company and product names may be
trademarks of their respective owners.
http://www.netapp.com/us/legal/netapptmlist.aspx[^]
== Patents
A current list of NetApp owned patents can be found at:
```

https://www.netapp.com/us/media/patents-page.pdf[^] == Privacy policy https://www.netapp.com/us/legal/privacypolicy/index.aspx[^] == Open source Notice files provide information about third-party copyright and licenses used in NetApp software. //* xref:{relative path}media/Notice Cloud Insights-2021.05.18.pdf[Notice for Cloud Insights] xref:{relative path}media/Notice Cloud Insights-2023-04.pdf[Notice for Cloud Insights] xref:{relative path}media/Notice Cloud Secure-2022-12-14.pdf[Notice for Workload Security (formerly Cloud Secure)] :leveloffset: -1 <<< *Copyright information* Copyright © 2024 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 link:http://www.netapp.com/TM[http://www.netapp.com/TM^] are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.