



Cloud Insights documentation

Cloud Insights

NetApp
April 23, 2021

This PDF was generated from <https://docs.netapp.com/us-en/cloudinsights/index.html> on April 23, 2021.
Always check docs.netapp.com for the latest.

Table of Contents

Cloud Insights documentation	1
Getting Started	1
What's New with Cloud Insights	2
Cloud Insights Onboarding	35
Creating your NetApp Cloud Central account	35
Starting your Cloud Insights free trial	35
Sign in and go	35
Logging Out	36
Security	37
Cloud Insights Security	37
Information and Region	39
Getting Started	42
Feature Tutorials	42
Collecting Data	43
Importing from the Dashboard Gallery	74
User accounts	74
Switching to other Cloud Central products	77
Cloud Insights Data Collector List	79
Subscribing to Cloud Insights	82
Editions	82
Trial Version	83
Subscription Options	83
How Do I Subscribe?	85
Subscription Mode	85
Subscribe Directly and Skip the Trial	87
Automatic Device Resolution	88
Automatic Device Resolution Overview	88
Device Resolution rules	90
Fibre Channel device resolution	93
IP device resolution	95
Setting options in the Preferences tab	97
Regular expression examples	98
Creating Dashboards	105
Dashboards Overview	105
Dashboard Features	108
Sample Dashboards	130
Best Practices for Dashboards and Widgets	135
Kubernetes Explorer	140
Kubernetes Cluster Overview	140
Kubernetes Cluster Detail Page	143
Working with Queries	150
Assets used in queries	150
Creating Queries	151

Viewing queries	154
Exporting query results to a .CSV file	155
Modifying or Deleting a Query	156
Assigning multiple applications to or removing multiple applications from assets	157
Copying table values	157
Working with Annotations	159
Defining annotations	159
Using annotations	161
Creating annotation rules	164
Importing Annotations	165
Working with Applications	168
Tracking asset usage by application	168
Creating Applications	168
Monitors and Alerts	170
Alerting with Monitors	170
Viewing and Managing Alerts from Monitors	176
Configuring Email Notifications	177
Cloud Insights API	180
Requirements for API Access	180
API Documentation (Swagger)	180
API Access Tokens	181
API Categories	182
Inventory Traversal	182
Expands	183
Performance Data	185
Object Performance Metrics	187
Performance History Data	187
Objects with Capacity Attributes	188
Using Search to Look Up Objects	189
Performance Policies and Alerts	190
Alerting with Monitors	190
Viewing and Managing Alerts from Monitors	196
Configuring Email Notifications	197
Notification using Webhooks	200
Webhook Examples:	203
Monitoring your Environment	204
Auditing	204
Asset Page Information	206
Asset Page Overview	206
Filtering for Objects In-Context	207
Asset Page Summary section	208
Expert View	210
User Data Section	215
Asset Page Violations section	216
Hints and Tips to Search for Assets	217

Reporting	221
Cloud Insights Reporting Overview	221
Cloud Insights Reporting User Roles	221
Predefined Reports Made Easy	223
Storage Manager Dashboard	226
Creating a Report (Example)	229
Managing Reports	230
Creating Custom Reports	233
How historical data is retained for Reporting	240
Cloud Insights Data Warehouse Schema Diagrams	241
Cloud Insights Schemas for Reporting	281
Cloud Secure	282
About Cloud Secure	282
Getting Started	282
Alerts	309
Forensics	315
Automated Response Policies	321
Configuring Alert Email Notifications	323
Cloud Secure API	323
Active IQ	325
Troubleshooting	327
Troubleshooting General Cloud Insights Problems	327
Troubleshooting Acquisition Unit Problems on Linux	328

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.

Cloud Insights helps you:

- **Reduce mean time to resolution by as much as 90%**

Stop log hunting for days and failing to manually correlate infrastructure: use our dynamic topology and correlation analysis to pinpoint the problem area immediately.

- **Reduce cloud infrastructure costs by an average of 33%**

Remove inefficiencies by identifying abandoned and unused resources and right-size workloads to their optimized performance and cost tiers.

- **Prevent as much as 80% of cloud issues from impacting end users**

End searching through vast amounts of data to find the relevant piece, by using advanced analytics and machine learning to identify issues before they become critical outages.

Getting Started

- I'm a New Customer: [Onboarding](#)
- I'm signed up: Now what do I do?
[Acquiring Data](#)
[Setting up users](#)
- How do I get the most out of Cloud Insights?
[Preparing Assets: Annotating](#)
[Finding the Assets You Want: Querying](#)
[Seeing the Data You want: Dashboards](#)
[Alerting: Performance Policies and Thresholds](#)
- I've tried out Cloud Insights and now 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 with Cloud Insights.

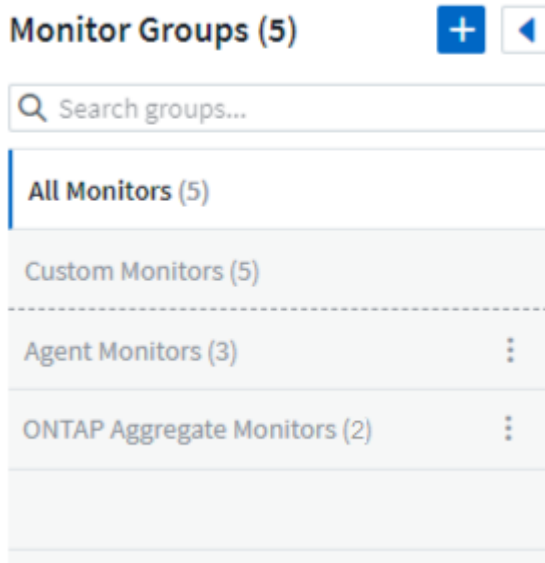
April 2021

The Cloud Insights team has developed a wealth of new content in March. Here are some of the highlights we are super excited to show:

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.

3 Set up team notification(s) (alert your team via email, or Webhook)

By Webhook	Notify team on	Use Webhook(s)
	<div>Critical, Warning</div>	<div>PagerDuty Trigger x</div>
	Notify team on	Use Webhook(s)
	<div>Resolved</div>	<div>PagerDuty Resolve x</div>

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.



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
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

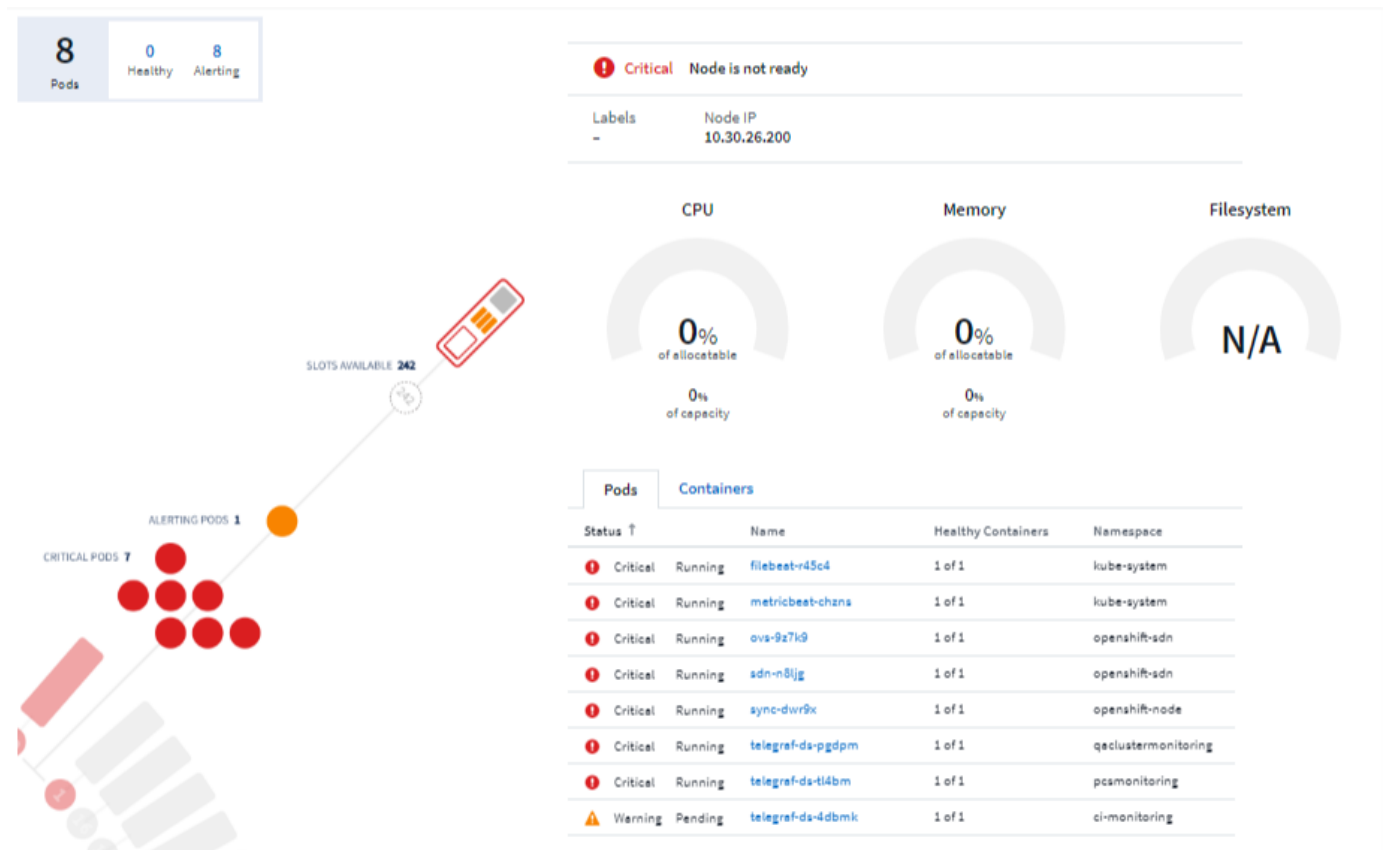
Previous Metric Name	New Metric Name
netapp_ontap.lun.write_data	netapp_ontap.lun.write_throughput
netapp_ontap.nic_common.rx_bytes	netapp_ontap.nic_common.rx_throughput
netapp_ontap.nic_common.tx_bytes	netapp_ontap.nic_common.tx_throughput
netapp_ontap.path.read_data	netapp_ontap.path.read_throughput
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
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

Previous Metric Name	New Metric Name
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.

Metrics & Attributes	
Metric Name	
qa-ots-cl01	<div> <div> <div>▼</div> <div>Rename Column</div> <div>Reset</div> </div> <div> <div>Metric Name</div> </div> </div>
ngslabc90	
kuat	
hkdemo-cluster	

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.

agent.node_fs

Filter By +

Group agent_node_name agent_node_os

3 items found

Table Row Grouping		Metrics & Attributes	
agent_node_name	agent_node_os	free	inodes_used
WIN2K12R2IMAGE	Microsoft Windows	70,594,338,816.00	0.00
WIN2K19IMAGE	Microsoft Windows	72,546,041,856.00	0.00
ci-qa-chunge-qaau	Red Hat Enterprise Linux Server	169,010,801,322.67	21,844.00

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.

Table Widget

agent.node

Filter By + Group agent_node_name

8 items found

Table Row Grouping		Metrics & Attributes	
agent_node_name ↑		mem.used (GiB)	
ci-qa-avinashp-k8-bakra-1		12.41	
ci-qa-avinashp-k8-bakra-2		9.31	
ci-qa-avinashp-k8-bakra-3		4.46	
ci-qa-avinashp-k8-bakra-4		1.15	
ci-qa-avinashp-k8swheel-1		15.23	

> Aggregation

< Unit Display

Base Unit byte (B)

Displayed In gibibyte (GiB)

Cancel Save

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

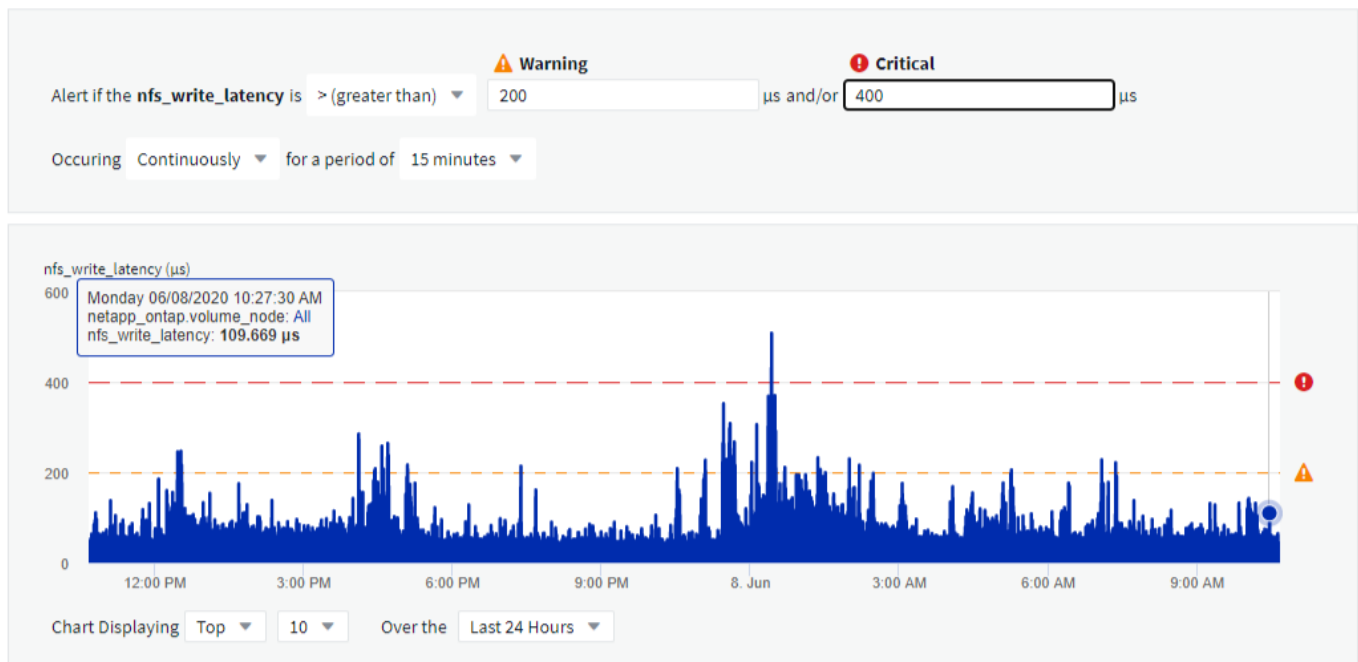
Cloud Insights has added a **Data Ingestion** API category, giving you greater control over custom data and agents. Detailed documentation for this and other API Categories can be found in Cloud Insights by navigating to **Admin > API Access** and clicking the *API Documentation* link. You can also attach 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:



POTENTIAL ATTACK: AL_307
Ransomware Attack

Detected
4 days ago
Jul 26, 2020 3:38 AM

Action Taken
Snapshots Taken

Status
In Progress

Last snapshots taken by
Amit Schwartz
Jul 30, 2020 2:54 PM

How To:
[Restore Entities](#)

[Re-Take Snapshots](#)

Total Attack Results


1 Affected Volumes | 0 Deleted Files | 5148 Encrypted Files

5148 Files have been copied, deleted, and potentially encrypted by 1 user account.

This is potentially a sign of ransomware attack.
The extension "crypt" was added to each file.

Encrypted Files

Activity per minute



Related Users



Ewen Hall
Developer
Engineering

5148
Encrypted Files

Detected
4 days ago
Jul 26, 2020 3:38 AM

Action Taken
Snapshots Taken



Manual Snapshot:

Cloud Insights

Abhi Basu Thakur

MONITOR & OPTIMIZE
Alerts / Nabilah Howell had an abnormal change in activity rate
Jul 23, 2020 - Jul 26, 2020
1:44 AM 1:44 AM

CLOUD SECURE
ALERTS
FORENSICS
ADMIN
HELP

Alert Detail

WARNING: AL_306
Nabilah Howell had an abnormal change in activity rate.

Detected
5 days ago
Jul 25, 2020 1:44 PM

Action Taken
None

Status
New

Recommendation: Setup an Automated Response Policy
An Automated Response Policy will trigger measures to contain the damage automatically when a future attack is detected. Try it now.

Take Snapshots
How To:
Restore Entities

Nabilah Howell's Activity Rate Change

Typical	Alert	
122.8 Activities Per Minute	210 Activities Per Minute	↑ 71%

Nabilah Howell's activity rate grew 71% over their typical average.

Activity Rate

Activity per 5 minutes

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

Object Type	Counter
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:



POTENTIAL ATTACK: AL_305
Ransomware Attack

Detected
5 days ago
Jul 11, 2020 4:02 AM

Action Taken
None

Status
New

Total Attack Results

1	0	4173
Affected Volumes	Deleted Files	Encrypted Files

4173 Files have been copied, deleted, and potentially encrypted by 1 user account.

This is potentially a sign of ransomware attack.
The extension "crypt" was added to each file.

Encrypted Files

Activity per minute



Related Users



Kristjan Egilsson
Accountant
Finance

4173
Encrypted Files

Detected
5 days ago
Jul 11, 2020 4:02 AM

Action Taken
None



Username
us035

Email
Egilsson@netapp.com

Phone
387224312607


Department
Finance

Manager
Lyndsey Maddox

Top Activity Types

Activity per minute
Last access location: 10.197.144.115

[View Activity Detail](#)



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.



GroupBy All

☐ Override Dashboard Time

index_0.index_0 

Filter By  Group

agent_version  

1 item found

Table Row Grouping

Metrics & Attributes

agent_version	value	consumer	protocol_name	level0	level1
Java/1.8.0_242	1,649.80	CloudInsights	GENERATED	simulated	N/A

- 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 with Integration Data (shown in Edit mode):

Table Widget - Integration Data Example Override Dashboard Time Last 7 Days

index_0.index_0

Filter By + Group agent_version name protocol_name

500 items found

Table Row Grouping			Metrics & Attributes				
agent_version	name	protocol_name	value	consumer	protocol_name	level0	level1
Java/1.8.0_242	simulated.shinchaku-client-1010.counter.2...	GENERATED	1,597.16	CloudInsights	GENERATED	simulated	shinchaku-
Java/1.8.0_242	simulated.shinchaku-client-1008.counter.1...	GENERATED	1,604.92	CloudInsights	GENERATED	simulated	shinchaku-
Java/1.8.0_242	simulated.shinchaku-client-1015.counter.1...	GENERATED	1,684.82	CloudInsights	GENERATED	simulated	shinchaku-
Java/1.8.0_242	simulated.shinchaku-client-1008.counter.0...	GENERATED	1,677.15	CloudInsights	GENERATED	simulated	shinchaku-

Cancel Save

- 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):

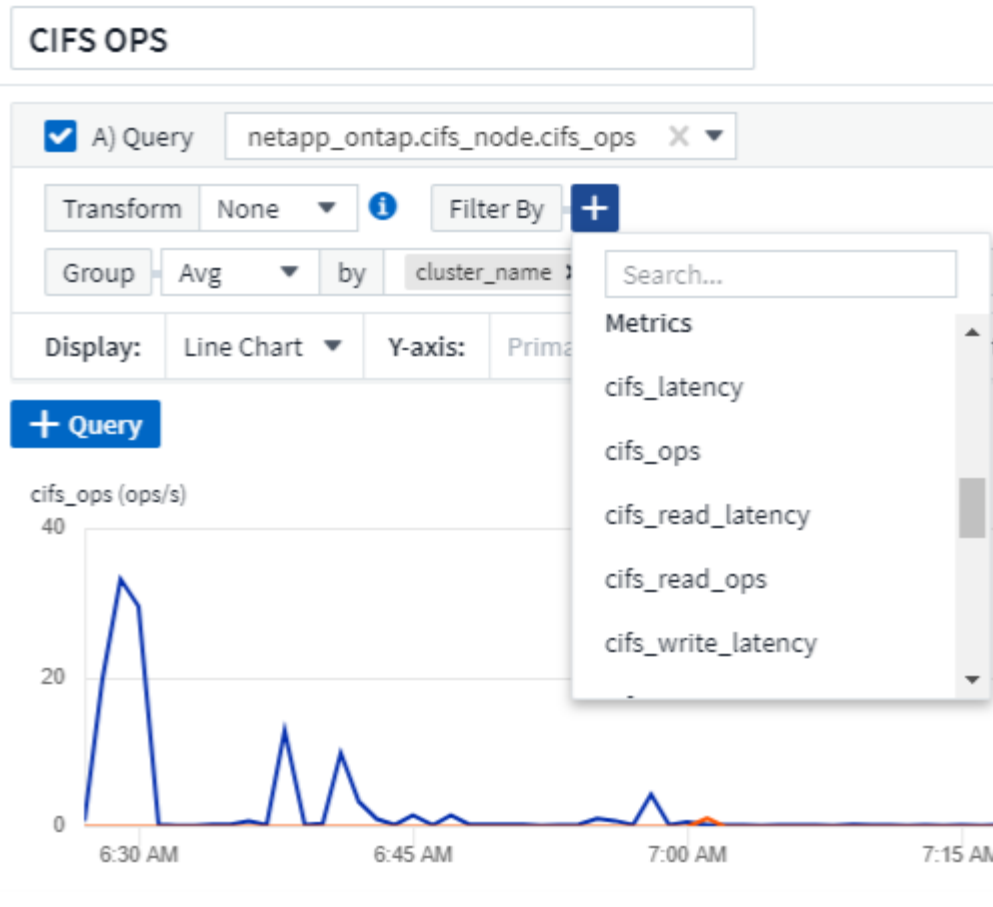
GroupBy Date 1h

4 items found in 2 groups


Active Date	Storage Node	Cache Hit Ratio - Total (%)	IOPS - Total (IO...	IOPS - Write (I...	Latency
06/01/2020 (1)	ocinaneqa1-01	N/A	N/A	N/A	N/A
06/01/2020	ocinaneqa1-01	N/A	N/A	N/A	N/A
N/A (3)	--	N/A	N/A	N/A	N/A

Metrics Filtering

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



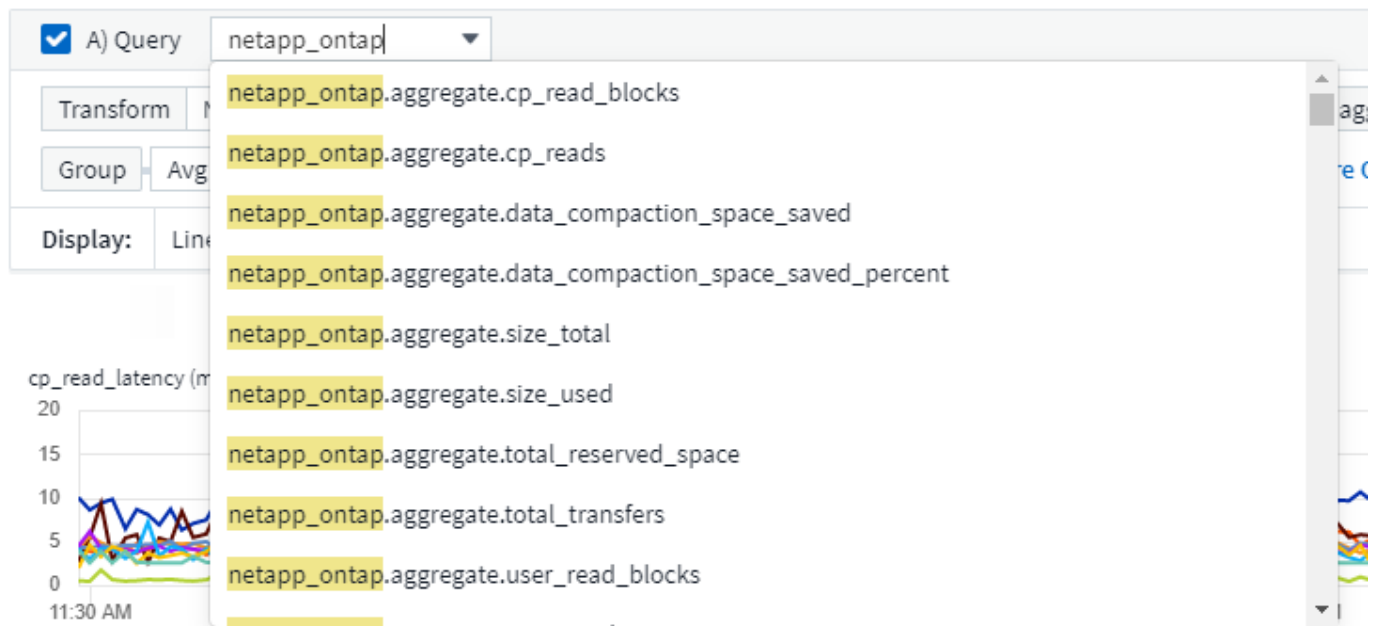
When working with integration data (Kubernetes, ONTAP Advanced Data, etc.), metric filtering removes the individual/unmatched data points from the plotted data series, unlike infrastructure data (storage, VM, ports etc.) where filters work on the aggregated value of the data series and potentially remove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



ONTAP Advanced Counter Data

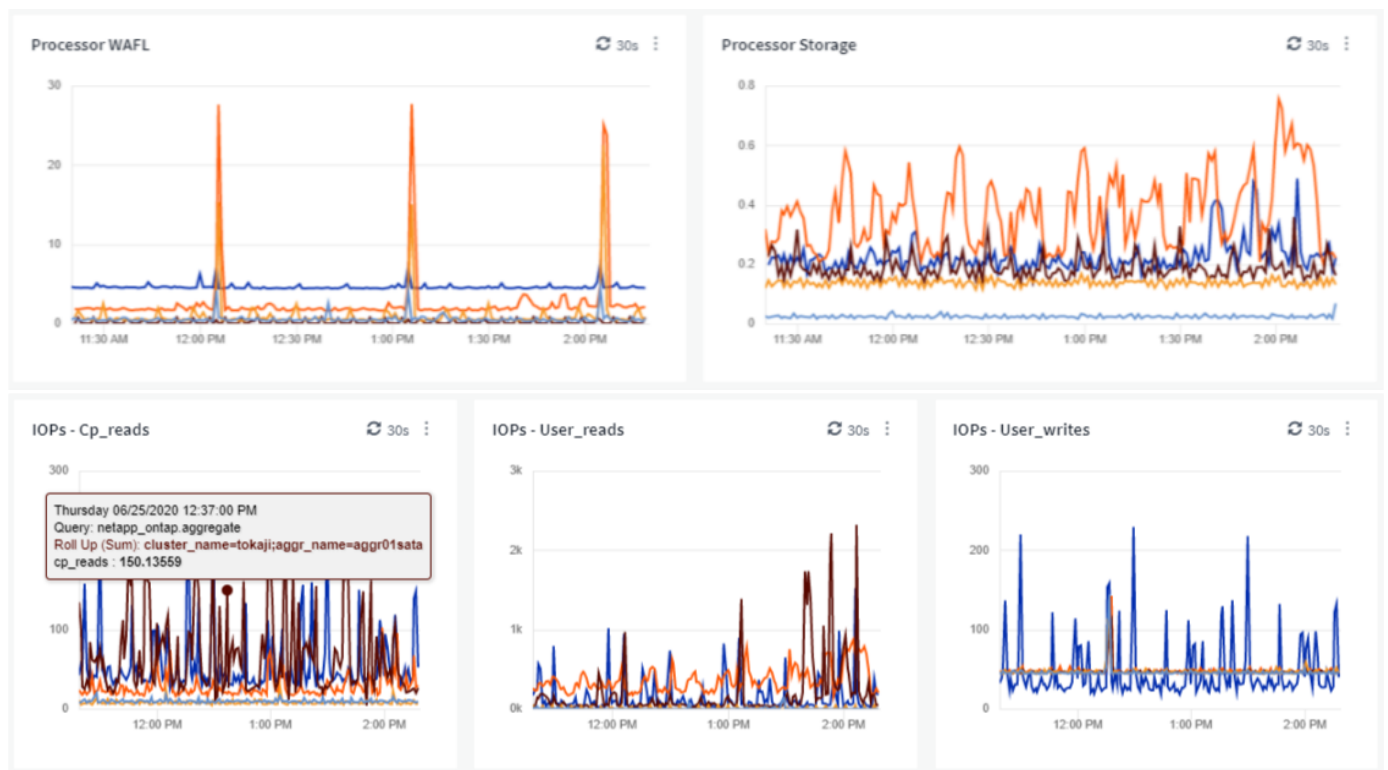
Cloud Insights takes advantage of NetApp's ONTAP-specific **Advanced Counter Data**, which provides a host of counters and metrics collected from ONTAP devices. ONTAP Advanced Counter Data is available to all NetApp ONTAP customers. These metrics enable customized and wide-ranging visualization in Cloud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

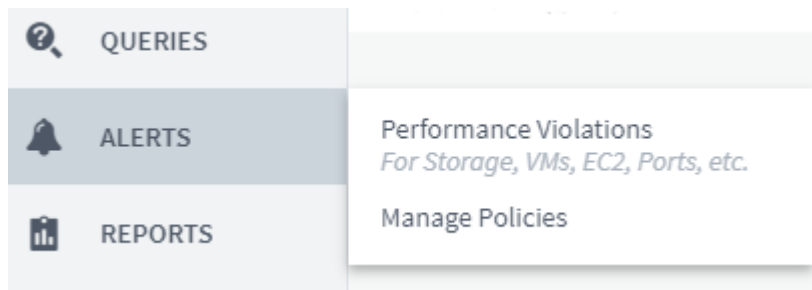
Learn More

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 bug 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

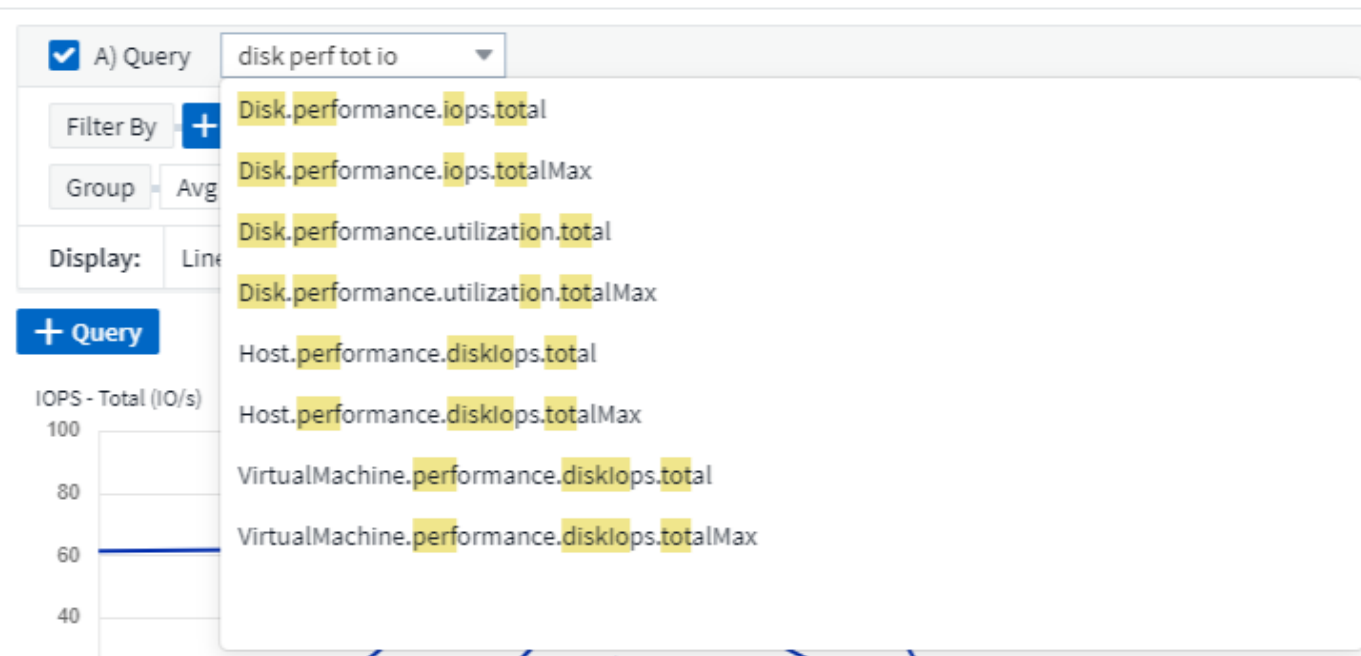
Reporting a data collector error is easier with the *Send Error Report* button on the data collector page. Clicking the button sends basic information about the error to NetApp and prompts investigation into the problem. Once pressed, Cloud Insights acknowledges that NetApp has been notified, and the 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.



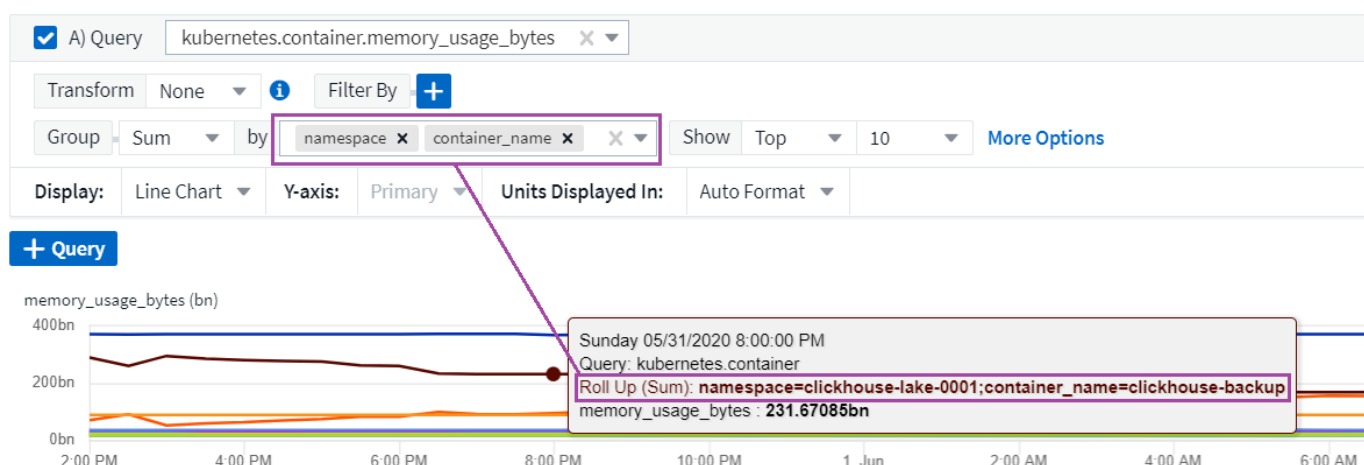
Widget Improvements

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:

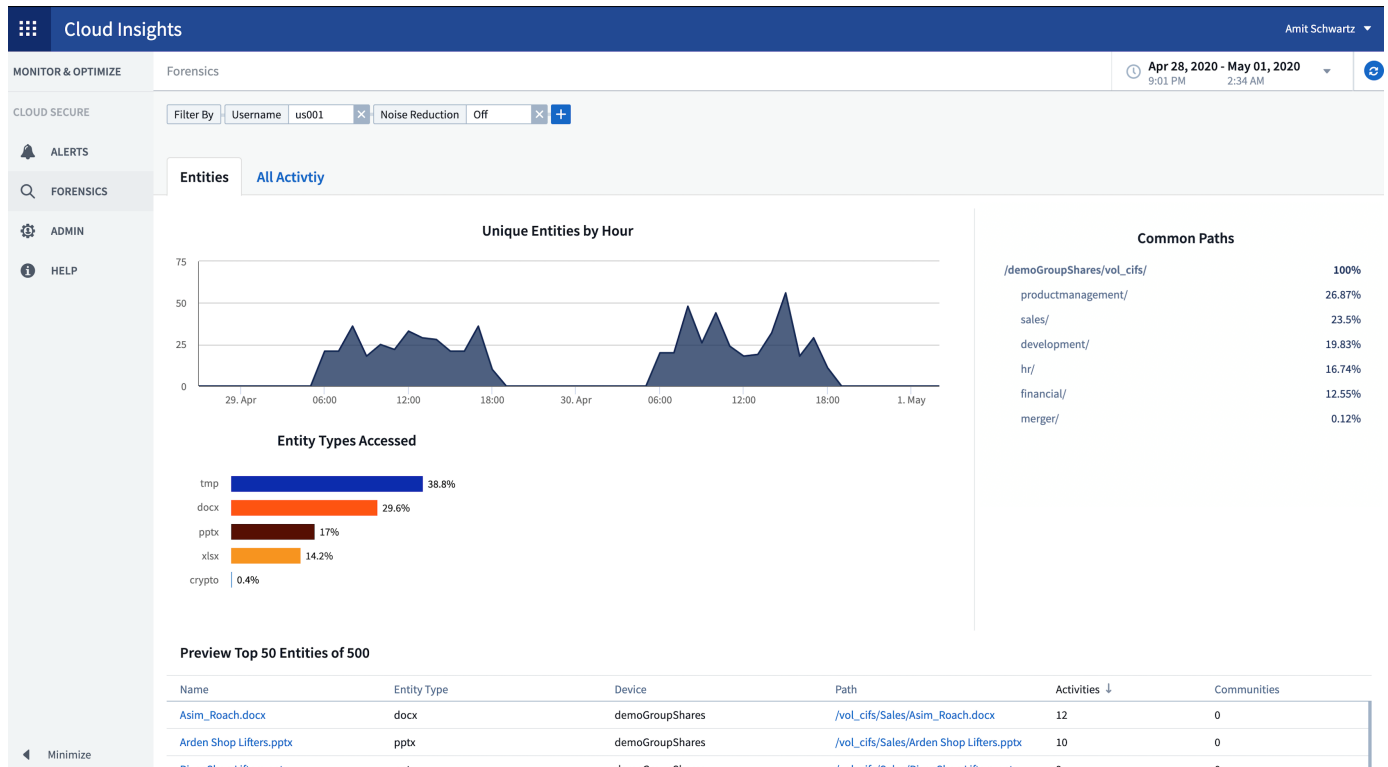
- Cloud Insights Consumers: can run and view reports
- Cloud Insights Authors: can perform the Consumer functions as well as create and manage reports and dashboards
- Cloud Insights Administrators: can perform the Author functions as well as all administrative tasks

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 Groups (3)



All Dashboards (60)

My Dashboards (11)

Storage Group (7) ⋮

Dashboards (7)



Name ↑

[Dashboard - Storage Cost](#)

[Dashboard - Storage IO Detail](#)

[Dashboard - Storage Overview](#)

[Gauges Storage Performance](#)

[Storage Admin - Which nodes are in high demand?](#)

[Storage Admin - Which pools are in high demand?](#)

[Storage IOPs](#)

Dashboard Pinning

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

Dashboards (7)



Name ↑



[Dashboard - Storage Overview](#)



[Storage Admin - Which nodes are in high demand?](#)



[Storage IOPs](#)

[Dashboard - Storage Cost](#)

[Dashboard - Storage IO Detail](#)

[Gauges Storage Performance](#)

[Storage Admin - Which pools are in high demand?](#)

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

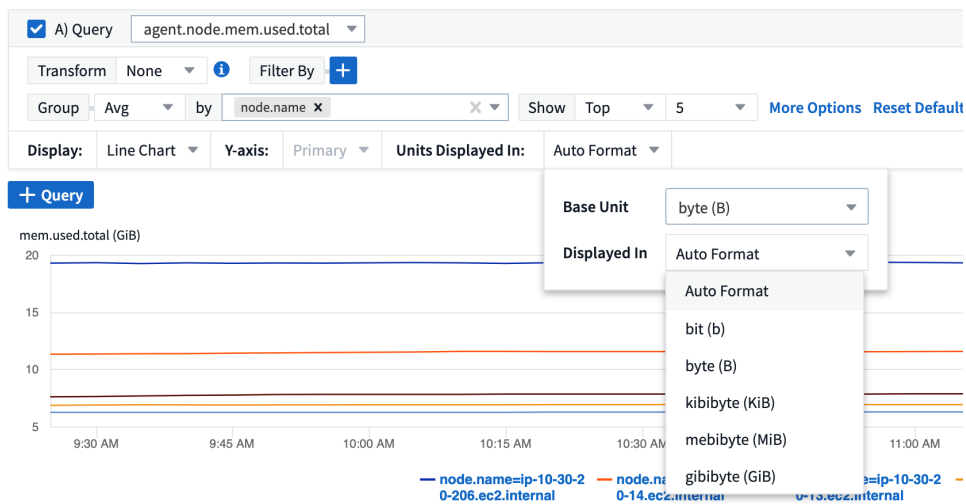
Optimized [time aggregation](#) intervals in time-series widgets (Line, Spline, Area, and Stacked Area charts) are more frequent for 3-hour and 24-hour dashboard/widget time ranges, allowing for faster charting of data.

- 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.

ASSETS.import

PUT /assets/import Import assets from a CSV file.

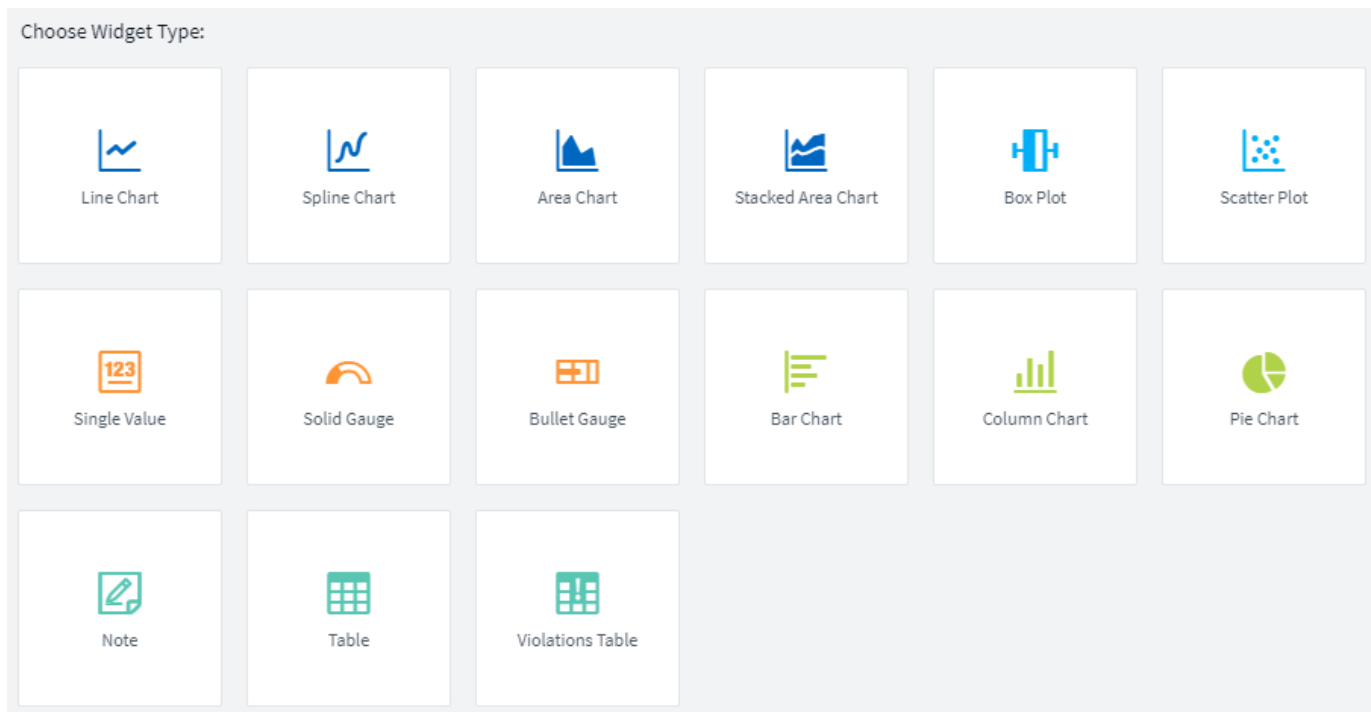
Import annotations and applications from the given CSV file. The format of the CSV file is following:


```

Project]
, <Annotation Type> [, <Annotation Type> ...] [, Application] [, Tenant] [, Line_Of_Business] [, Business_Unit] [,
<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>]
      
```

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.



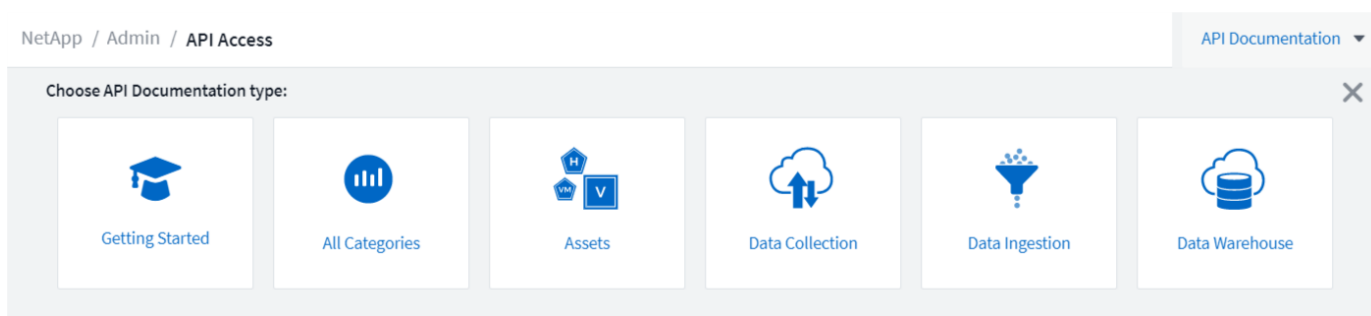
February 2020

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 Access**, 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.

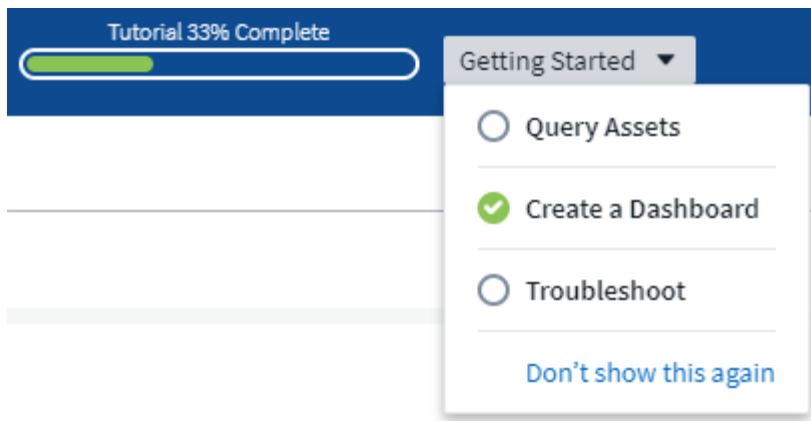
January 2020

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 errors for the data collector.

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 HCI data collector acquires from the HCI 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 [pre-requisites](#) 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 **Admin > Data 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 Cloud Central** portal. If you already have a NetApp Cloud Central login, you can start a free trial of Cloud Insights with a few quick steps.

Creating your NetApp Cloud Central account

To sign up for access to NetApp's cloud services, go to [NetApp Cloud Central](#) and click **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 click **Sign Up**.

You will then be taken to NetApp's cloud offerings page.

Select Cloud Insights.

What if I already have a NetApp Cloud login?

If you already have a NetApp Cloud Central account, simply choose **Log In** on the [NetApp Cloud Central](#) 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 then create your company's environment.

Once the creation of your environment is complete, you can use your Cloud Central credentials to log in and start your free, 30-day trial of Cloud Insights. During this trial you can explore all the features that Cloud Insights Standard Edition 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 tenant has been created, at any time you can simply log in to the NetApp Cloud 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-name>.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 Cloud Central, they will be prompted to log in.



New users must still sign up for access to Cloud Central 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 Cloud Central sign in screen.



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

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.

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 Cloud Secure 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 AES-256, 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.

- Cloud Secure User Directory Data

In cases where the Cloud Secure 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 Cloud Secure User Directory data collector, not to data about the users of Cloud Insights/Cloud Secure 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 ActiveIQ. However, depending on a customer's naming conventions, data for shares, volumes, VMs, qtrees, applications, etc. may contain personally identifiable information.

If Cloud Secure 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 Cloud Secure 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
- Cloud Secure Active Directory information
- Cloud Secure 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.

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](#)

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.

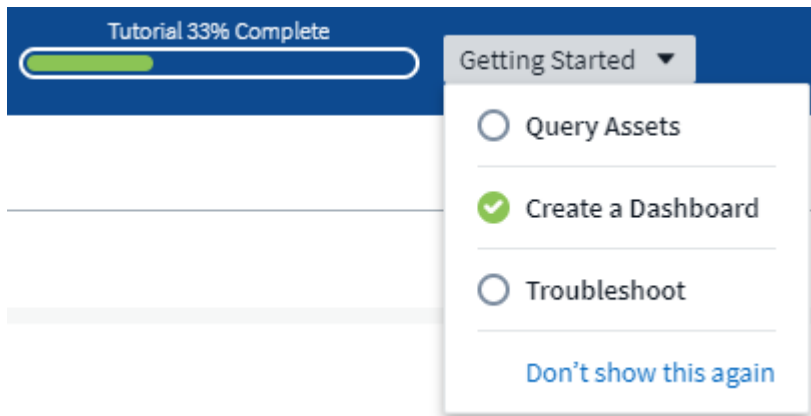
Introduction

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.



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

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.

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)

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 utilizes three types of data collectors:

- Operating Systems
- Services
- Infrastructure

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 to collect data for Cloud Insights. A single Acquisition Unit can be used for multiple data collectors.

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.

What Operating System or Platform Are You Using?

 Linux

Linux Versions Supported 

Production Best Practices 

Installation Instructions

[Need Help?](#)

1 [Copy Installer Snippet](#)

This snippet has a unique key valid for 24 hours for this Acquisition Unit only.

 [Reveal Installer Snippet](#)

2 Paste the snippet into a bash shell to run the installer.

3 Waiting for Acquisition Unit to connect...

- 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.

3 [Continue](#) **New acquisition unit detected!**

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 user name and password to connect to the device, as appropriate.
- Fill in any other mandatory fields in *Configuration* and *Advanced Configuration* sections.
- Click **Test Connection** to test the connection to the device.
- 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 (MacOS, 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

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	<p>A computer running a licensed version of one of the following:</p> <ul style="list-style-type: none"> * Centos (64-bit): 7.2 through 7.9, 8.1 through 8.2 * Debian (64-bit): 9 * Oracle Enterprise Linux (64-bit): 7.5 through 7.9, 8.1 through 8.2 * Red Hat Enterprise Linux (64-bit): 7.2 through 7.9, 8.1 through 8.2 * Ubuntu Server: 18.04 LTS <p>This computer should be running no other application-level software. A dedicated server is recommended.</p>	<p>A computer running a licensed version of one of the following:</p> <ul style="list-style-type: none"> * Microsoft Windows 10 64-bit * Microsoft Windows Server 2012 * Microsoft Windows Server 2012 R2 * Microsoft Windows Server 2016 * Microsoft Windows Server 2019 <p>This computer should be running no other application-level software. A dedicated server is recommended.</p>
CPU	2 CPU cores	Same
Memory	8 GB RAM	Same
Available disk space	<p>50 GB</p> <p>For Linux, disk space should be allocated in this manner:</p> <p>/opt/netapp 25 GB</p> <p>/var/log/netapp 25 GB</p>	50 GB
Network	100 Mbps / 1 Gbps Ethernet connection, static IP address, IP connectivity to all FC devices, and a required port to the Cloud Insights instance (80 or 443).	Same
Permissions	<p>Sudo permissions on the Acquisition Unit server.</p> <p>/tmp must be mounted with exec capabilities.</p>	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, 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.

For additional sizing questions, contact NetApp Support.

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)**.

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 **Admin > Data Collectors > Acquisition Units > +Acquisition Unit**

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

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.

What Operating System or Platform Are You Using?

 Linux

Linux Versions Supported ⓘ Production Best Practices ⓘ

Installation Instructions

[Need Help?](#)

1 [Copy Installer Snippet](#)

This snippet has a unique key valid for 24 hours for this Acquisition Unit only.

[+ Reveal Installer Snippet](#)

2 Paste the snippet into a bash shell to run the installer.

3 Waiting for Acquisition Unit to connect...

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.
3. 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 **Admin > Data 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?*
2. 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.

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.
2. On that server, open a browser window and log in to your Cloud Insights environment as Administrator or Account Owner.
3. Click **Admin > Data 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.

What Operating System or Platform Are You Using?



Windows Versions Supported ⓘ Production Best Practices ⓘ

Installation Instructions

[Need Help?](#)

1 [Download Installer \(Windows 64-bit\)](#)

2 [Copy Access Key](#)

This access key is a unique key valid for 24 hours for this Acquisition Unit only.

[+ Reveal Access Key](#)

3 Paste access key into installer when prompted.

4 Please ensure you have copied and pasted the access key into the installer.

[+ Have a Proxy Server?](#)

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 **Admin > Data 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.

Uninstalling an Acquisition Unit

To uninstall the Acquisition Unit software, do the following:

Windows:

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:

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
```

Both:

1. After uninstalling the AU software, go to **Admin > Data 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.

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 **Admin > Data 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 **Admin > Data 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

Acquisition Unit Summary

Name xp-linux	Connection Status OK - Need Help?	Operating System Linux	Note
IP 10.197.120.145	Last Reported 2 minutes ago		

Data Collectors (3)

[+ Data Collector](#)
[Bulk Actions](#)

<input type="checkbox"/>	Name ↑	Status	Type	IP	Impact	Last Acquired	
	foo	Inventory failed	NetApp Data ONTAP 7-Mode	foo	Low	Never	⋮
	xp-cisco	All successful	Cisco MDS Fabric Switches	10.197.136.66		2 minutes ago	⋮
<input type="checkbox"/>	xpcdot26	All successful	NetApp ONTAP Data Management Software	10.197.136.26		8 minutes ago	⋮

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

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.



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)**.

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](#)
- [macOS](#)
- [Kubernetes](#)

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 site and go to **Admin > Data Collectors**.
 3. Click on **+Data Collector** and choose a data collector to install.
1. Choose the appropriate platform for your host (Windows, Linux, macOS, etc.)
 2. 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.

Windows



Install Agent

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

Select existing API Access Token or create a new one

KEY1 (...Zqlk0c)

+ API Access Token

Installation Instructions

[Need Help?](#)

1

Copy Agent Installer Snippet

This snippet has a unique key and is valid for 24 hours. Already have an agent in your environment? [View Troubleshooting](#)

⊞ Reveal Agent Installer Snippet

2

Open a PowerShell window as administrator and paste the snippet

3

Complete Setup

Pre-requisites:

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

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.
3. 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.
6. 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
```

Configuring Proxy Support for Windows

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",  
"<proxy_server>:<proxy_port>",  
[System.EnvironmentVariableTarget]::Machine)
```

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
```

2. Delete the *C:\Program Files\telegraf* folder to remove the binary, logs, and configuration files
3. 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



Install Agent

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

Select existing API Access Token or create a new one

KEY1 (...Zq|k0c)

+ API Access Token

Installation Instructions

[Need Help?](#)

1

Copy Agent Installer Snippet

This snippet has a unique key and is valid for 24 hours. Already have an agent in your environment? [View Troubleshooting](#)

📄 Reveal Agent Installer Snippet

2

Open a terminal window and paste the snippet in a Bash shell (requires curl, sudo, ping, and dmidecode)

3

Complete Setup

Pre-requisites:

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

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
```

Configuring Proxy Support for RHEL/CentOS

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>
```

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



Install Agent

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

Select existing API Access Token or create a new one

KEY1 (...ZqIk0c)

+ API Access Token

Installation Instructions

[Need Help?](#)

1

Copy Agent Installer Snippet

This snippet has a unique key and is valid for 24 hours. Already have an agent in your environment? [View Troubleshooting](#)

⊞ Reveal Agent Installer Snippet

2

Open a terminal window and paste the snippet in a Bash shell (requires curl, sudo, ping, and dmidecode)

3

Complete Setup

Pre-requisites:

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

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
```

Configuring Proxy Support for Ubuntu/Debian

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>
```

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.](#)

macOS



Install Agent

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

Select existing API Access Token or create a new one

KEY1 (...Zqlk0c)

+ API Access Token

Installation Instructions

[Need Help?](#)

1

Copy Agent Installer Snippet

This snippet has a unique key and is valid for 24 hours. Already have an agent in your environment? [View Troubleshooting](#)
⊞ Reveal Agent Installer Snippet

2

Open a terminal window and paste the snippet in a Bash shell (requires sudo and curl)

3

Complete Setup

Pre-requisites:

- The "curl" command must be available
- If you are behind a proxy, follow the instructions in the **Configuring Proxy Support for macOS** section.

Steps to install agent on macOS:

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. If you previously installed a Telegraf agent using Homebrew, you will be prompted to uninstall it. Once the previously installed Telegraf agent is uninstalled, re-run the command in step 5 above.
7. Click **Finish** or **Continue**

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

```
sudo launchctl start telegraf
sudo launchctl stop telegraf
```

Configuring Proxy Support for macOS

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

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

AFTER installing the Telegraf agent, add and set the appropriate *https_proxy* and/or *http_proxy* variable(s) in */Applications/telegraf.app/Contents/telegraf.plist*:

```
...
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN"
"http://www.apple.com/DTDs/PropertyList-1.0.dtd">
<plist version="1.0">
<dict>
  <key>EnvironmentVariables</key>
  <dict>
    <key>https_proxy</key>
    <string><proxy_server>:<proxy_port></string>
  </dict>
  <key>Program</key>
  <string>/Applications/telegraf.app/Contents/MacOS/telegraf</string>
  <key>Label</key>
  <string>telegraf</string>
  <key>ProgramArguments</key>
  <array>
    <string>/Applications/telegraf.app/Contents/MacOS/telegraf</string>
    <string>--config</string>
    <string>/usr/local/etc/telegraf.conf</string>
    <string>--config-directory</string>
    <string>/usr/local/etc/telegraf.d</string>
  </array>
  <key>RunAtLoad</key>
  <true/>
</dict>
</plist>
...
```

Then, restart Telegraf after loading the above changes:

```
sudo launchctl stop telegraf
sudo launchctl unload -w /Library/LaunchDaemons/telegraf.plist
sudo launchctl load -w /Library/LaunchDaemons/telegraf.plist
sudo launchctl start telegraf
```

Uninstalling the Agent

To uninstall the agent on macOS, in a Bash terminal, run the following:

1. Stop the Telegraf service:

```
sudo launchctl stop telegraf
```

2. Uninstall the telegraf agent:

```
cp /Applications/telegraf.app/scripts/uninstall /tmp
sudo /tmp/uninstall
```

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

```
rm -rf /usr/local/etc/telegraf*
rm -rf /usr/local/var/log/telegraf.*
```

Upgrading the Agent

To upgrade the telegraf agent, do the following:

1. Stop the telegraf service:

```
sudo launchctl stop telegraf
```

2. Uninstall the previous telegraf agent:

```
cp /Applications/telegraf.app/scripts/uninstall /tmp
sudo /tmp/uninstall
```

3. [Install the new agent.](#)

Kubernetes

Install Agent

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

Select existing API Access Token or create a new one

default_ingestion_api_key1 (...Y6G511) ▼

+ API Access Token

Production Best Practices ⓘ

Installation Instructions

[Need Help?](#)

1 **kube-state-metrics** must be installed and running. Note that some variants of Kubernetes may require additional **security considerations**. For environments operating behind a proxy server, follow the instructions to [configure proxy support to install and run Telegraf](#).

2 [Copy Agent Installer Snippet](#)

This snippet has a unique key and is valid for 24 hours. Already have an agent in your environment? [View Troubleshooting](#)

[+ Reveal Agent Installer Snippet](#)

3 **Open a terminal window and paste the snippet in a Bash shell on the target Kubernetes cluster (requires curl, sudo and kubectl).**

4 [Complete Setup](#)

Pre-requisites:

- The following commands must be available: curl, sudo, kubectl

For best results, add these commands to the PATH.

- kube-state-metrics must be installed. See below for more information.
- If you are behind a proxy, follow the instructions in the **Configuring Proxy Support for Kubernetes** section.
- If you are running a kubernetes variant that requires security context constraints, follow the instructions in the **Configuring the Agent to Collect Data from Kubernetes** section.

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 v1.14 and above	Kubernetes.io/os = linux
Kubernetes v1.13 and below	beta.kubernetes.io/os = linux
Rancher + cattle.io as orchestration/Kubernetes platform	cattle.io/os = linux

Installation

Steps to install agent on Kubernetes:

1. Choose an Agent Access Key.

2. Click the **Copy Agent Installer Snippet** button in the installation dialog. You can optionally click the **+Reveal Agent Installer Snippet** button if you want to view the command block.
3. Open a Bash window.
4. Paste the command into the Bash window.
5. Optionally, you can override the namespace or provide the cluster name as part of the install command by modifying the command block to add one or both of the following before the final `./$installerName`
 - `CLUSTER_NAME=<Cluster Name>`
 - `NAMESPACE=<Namespace>`

Scroll through the following example to see this in place in the command block:

```
installerName=cloudinsights-kubernetes.sh && token=<token> &&  
key=c642e336-91f4-4c6f-8086-72faabd6aff6 &&  
domain=tenant1.testk8.cloudinsights-test.netapp.com && curl -k -X GET  
-H "Authorization: Bearer $token" -H "X-CloudInsights-APIKey-Id:  
$key" -o $installerName  
https://$domain/rest/v1/lake/telegraf/platforms/installer?platform=ku  
bernetes && chmod +x $installerName && sudo --preserve-env JWT=$token  
DOMAIN_NAME=$domain API_KEY_ID=$key CLUSTER_NAME=TEST_CLUSTER  
NAMESPACE=NEW-NAMESPACE ./$installerName
```



CLUSTER_NAME is the name of the Kubernetes cluster from Cloud Insights collects metrics, while *NAMESPACE* is the namespace to which the Telegraf agent will be deployed. The specified namespace will be created if it does not already exist.

1. When ready, execute the command block.
2. The command will download the appropriate agent installer, install it, and set a default configuration. If you have not explicitly set the *namespace*, you will be prompted to enter it. When finished, the script will restart the agent service. The command has a unique key and is valid for 24 hours.
3. When finished, click **Complete Setup**.

After the agent is installed, generate the Telegraf DaemonSet YAML and ReplicaSet YAML using the following commands. Note that these commands are using the default namespace "ci-monitoring".

If you have set your own namespace, substitute that namespace in these and all subsequent commands and files:

```
kubectl --namespace ci-monitoring get ds telegraf-ds -o yaml >  
/tmp/telegraf-ds.yaml  
kubectl --namespace ci-monitoring get rs telegraf-rs -o yaml >  
/tmp/telegraf-rs.yaml
```

You can use the following commands to stop and start the Telegraf service:

```
kubectl --namespace ci-monitoring delete ds telegraf-ds
kubectl --namespace ci-monitoring delete ds telegraf-rs
```

```
kubectl --namespace ci-monitoring apply -f /tmp/telegraf-ds.yaml
kubectl --namespace ci-monitoring apply -f /tmp/telegraf-rs.yaml
```

Configuring Proxy Support for Kubernetes

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

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

AFTER installing the Telegraf agent, add and set the appropriate *https_proxy* and/or *http_proxy* environment variable(s) to the *telegraf-ds* daemonset and *telegraf-rs* replicaset.

```
kubectl edit ds telegraf-ds
```

```
...
  env:
  - name: https_proxy
    value: <proxy_server>:<proxy_port>
  - name: HOSTIP
    valueFrom:
      fieldRef:
        apiVersion: v1
        fieldPath: status.hostIP
...
```

```
kubectl edit rs telegraf-rs
```

```
...
  env:
    - name: https_proxy
      value: <proxy_server>:<proxy_port>
    - name: HOSTIP
      valueFrom:
        fieldRef:
          apiVersion: v1
          fieldPath: status.hostIP
...

```

Then, restart Telegraf:

```
kubectl delete pod telegraf-ds-*
kubectl delete pod telegraf-rs-*

```

Configuring the Agent to Collect Data from Kubernetes

For Kubernetes environments, Cloud Insights deploys the Telegraf agent as a DaemonSet and a ReplicaSet. The pods in which the agents run need to have access to the following:

- hostPath
- configMap
- secrets

These Kubernetes objects are automatically created as part of the Kubernetes agent install command provided in the Cloud Insights UI. Some variants of Kubernetes, such as OpenShift, implement an added level of security that may block access to these components. The *SecurityContextConstraint* is not created as part of the Kubernetes agent install command provided in the Cloud Insights UI, and must be created manually. Once created, restart the Telegraf pod(s).

```

apiVersion: v1
kind: SecurityContextConstraints
metadata:
  name: telegraf-hostaccess
  creationTimestamp:
  annotations:
    kubernetes.io/description: telegraf-hostaccess allows hostpath
volume mounts for restricted SAs.
  labels:
    app: ci-telegraf
priority: 10
allowPrivilegedContainer: false
defaultAddCapabilities: []
requiredDropCapabilities: []
allowedCapabilities: []
allowedFlexVolumes: []
allowHostDirVolumePlugin: true
volumes:
- hostPath
- configMap
- secret
allowHostNetwork: false
allowHostPorts: false
allowHostPID: false
allowHostIPC: false
seLinuxContext:
  type: MustRunAs
runAsUser:
  type: RunAsAny
supplementalGroups:
  type: RunAsAny
fsGroup:
  type: RunAsAny
readOnlyRootFilesystem: false
users:
- system:serviceaccount:ci-monitoring:telegraf-user
groups: []

```

Installing the kube-state-metrics server

When you install the kube-state-metrics server you can enable collection of the following Kubernetes objects: StatefulSet, DaemonSet, Deployment, PV, PVC, ReplicaSet, Service, Namespace, Secret, ConfigMap, Pod Volume, and Ingress.



It is strongly recommended to use kube-state-metrics version 2.0 or later in order to take advantage of the full feature set including the ability to link Kubernetes persistent volumes (PVs) to backend storage devices.

Use the following steps to install the kube-state-metrics server:

Steps

1. Create a temporary folder (for example, `/tmp/kube-state-yaml-files/`) and copy the .yaml files from <https://github.com/kubernetes/kube-state-metrics/tree/master/examples/standard> to this folder.
2. Run the following command to apply the .yaml files needed for installing kube-state-metrics:

```
kubectl apply -f /tmp/kube-state-yaml-files/
```

kube-state-metrics Counters

Use the following links to access information for the kube state metrics counters:

1. [ConfigMap Metrics](#)
2. [DaemonSet Metrics](#)
3. [Deployment Metrics](#)
4. [Ingress Metrics](#)
5. [Namespace Metrics](#)
6. [Node Metrics](#)
7. [Persistent Volume Metrics](#)
8. [Persistent Volume Claim Metrics](#)
9. [Pod Metrics](#)
10. [ReplicaSet metrics](#)
11. [Secret metrics](#)
12. [Service metrics](#)
13. [StatefulSet metrics](#)

Uninstalling the Agent

To uninstall the agent on Kubernetes, do the following:

1. If the monitoring namespace is being used solely for Telegraf:

```
kubectl delete ns ci-monitoring
```

If the monitoring namespace is being used for other purposes in addition to Telegraf:

1. Stop and delete the Telegraf service:

```
kubectl --namespace ci-monitoring delete ds telegraf-ds
kubectl --namespace ci-monitoring delete rs telegraf-rs
```

2. Delete the Telegraf ConfigMap and ServiceAccount:

```
kubectl --namespace ci-monitoring delete cm telegraf-conf
kubectl --namespace ci-monitoring delete cm telegraf-conf-rs
kubectl --namespace ci-monitoring delete sa telegraf-user
```

3. Delete the Telegraf ClusterRole and ClusterRolebinding:

```
kubectl --namespace ci-monitoring delete clusterrole endpoint-access
kubectl --namespace ci-monitoring delete clusterrolebinding endpoint-access
```

Upgrading the Agent

To upgrade the telegraf agent, do the following:

1. Remove the current the telegraf deployments:

```
kubectl --namespace ci-monitoring delete ds telegraf-ds
kubectl --namespace ci-monitoring delete rs telegraf-rs
```

2. Back up the existing configurations:

```
kubectl --namespace ci-monitoring get cm telegraf-conf -o yaml >
/tmp/telegraf-conf.yaml
kubectl --namespace ci-monitoring get cm telegraf-conf-rs -o yaml >
/tmp/telegraf-conf-rs.yaml
```

3. [Install the new agent.](#)

Troubleshooting Agent Installation

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

Problem:	Try this:
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 .
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.

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 **Admin > Data Collectors**

The system displays the available Data Collectors arranged by vendor.

2. Click **+ Collector** on the required vendor 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 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. Click **Advanced Configuration** to add additional configuration fields. (Not all data collectors require advanced configuration.)
7. Click **Test Configuration** to verify that the data collector is properly configured.

8. 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 **Admin > Data 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 **Admin > Data 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 **Admin > Data 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.
2. 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 **Admin > Data 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 **Admin > Data 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 **Dashboard > Show all 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.

User accounts

Cloud Insights provides four user accounts: Account Owner, Administrator, User, and Guest. Each account is assigned specific permission levels. Users are either [invited](#) to Cloud Insights and assigned a specific role, or can sign in via [Single Sign-On \(SSO\)](#) with a default role. SSO is available as a feature in Cloud Insights Premium Edition.

Permission levels

You use an account that has Administrator privileges to create or modify user accounts. Each user account is assigned one of the following permission levels.

- **Guest** can view asset pages, dashboards, and queries, and run queries.
- **User** can perform all guest-level privileges as well as create, modify, or delete dashboards, queries, annotations, annotation rules, and applications.
- **Administrator** and **Account Owner** can perform all functions, as well as create, modify and delete policies, import dashboards, and manage all users and data collectors.

The Account Owner is created when you register for Cloud Insights.

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

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	Y	N	N
Alerts*/Policies: Create/Modify/Delete	Y	Y	N
Alerts*/Policies: View	Y	Y	Y
Annotation Rules: Create/Run/Modify/Delete	Y	Y	N
Annotations: Create/Modify/Assign/View/Remove/Delete	Y	Y	N
API Access*: Create/Rename/Disable/Revoke	Y	N	N
Applications: Create/View/Modify/Delete	Y	Y	N
Asset Pages: Modify	Y	Y	N
Asset Pages: View	Y	Y	Y
Audit: View	Y	N	N
Cloud Cost	Y	N	N
Cloud Secure*	Y	N	N
Dashboards: Create/Modify/Delete	Y	Y	N
Dashboards: View	Y	Y	Y
Data Collectors: Add/Modify/Poll/Delete	Y	N	N
Notifications: View/Modify	Y	N	N
Queries: Create/Modify/Delete	Y	Y	N
Queries: View/Run	Y	Y	Y
Reports*: View/Run	Y	Y	Y
Reports*: Create/Modify/Delete/Schedule	Y	Y	N
Subscription: View/Modify	Y	N	N

User Management: Invite/Add/Modify/Deactivate	Y	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.

5. Enter the user's e-mail address.
6. Select the user role.
7. 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.

Single Sign-On (SSO) Accounts

In addition to inviting users, 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 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, using open standards such as Security Assertion Markup Language 2.0 (SAML) and OpenID Connect (OIDC).

To configure SSO, on the **Admin > User Management** page, click the **Configure SSO** button. Once configured, administrators can then enable SSO user login. When an administrator enables SSO, 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.

The screenshot shows the 'Cloud Insights' interface with the 'Admin > User Management' page. A blue banner at the top states 'Single Sign-on (SSO) now available! Allow user access to Cloud Insights through corporate credentials.' with a 'Configure SSO' button and a 'Dismiss' button. Below the banner, there is a '+ User' button and a 'Filter...' input field. A table lists two users:

Name	Email ↑	Role	Last Login
DOW	user11@netappci1.onmicrosoft.com	Administrator	4 minutes ago
user1	user1@netappci1.onmicrosoft.com	Account Owner	2 hours ago

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 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 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.

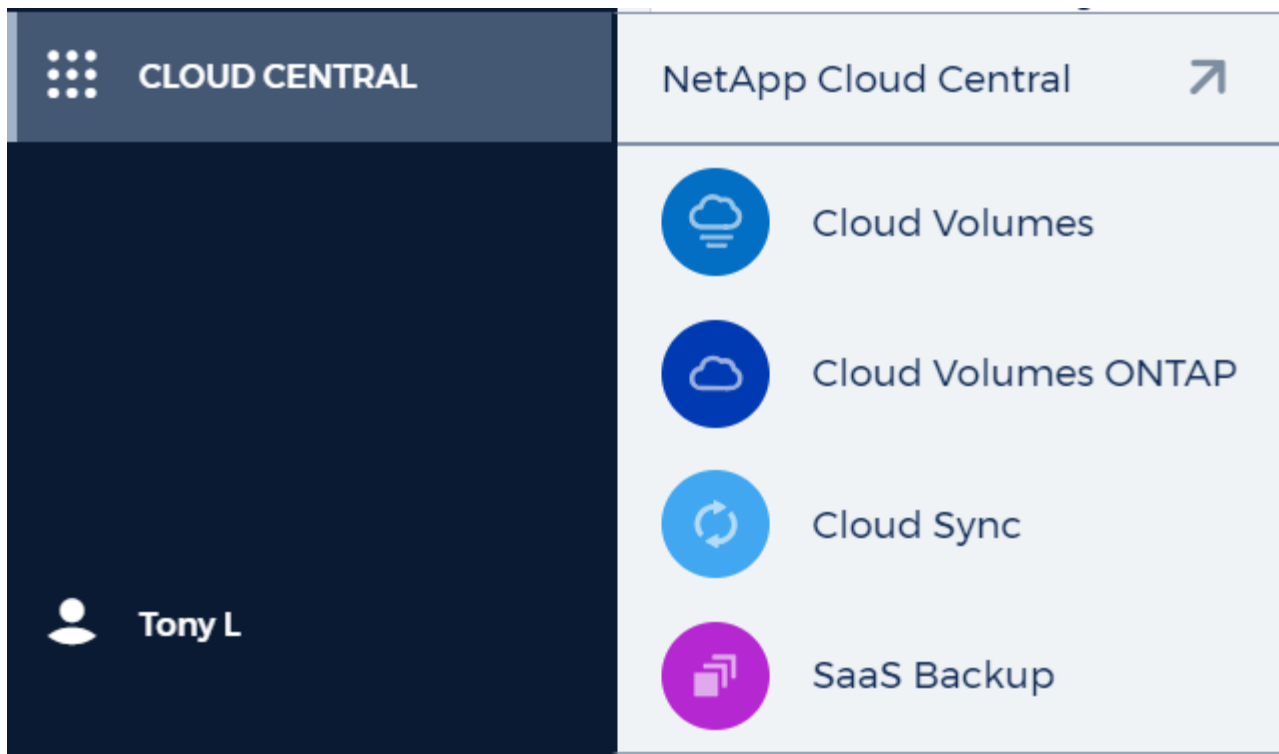
The screenshot shows the 'Cloud Insights' interface with the 'Admin > User Management' page. A toggle switch labeled 'Enable SSO' is turned on. Below the toggle, there is a 'Show SSO Users' checkbox which is checked. A table lists three users:

Name	Email ↑	Role
user1 user1	user1@netappci1.onmicrosoft.com	Account Owner
user2 user2	user2@netappci1.onmicrosoft.com	SSO (Guest)
user3 user3	user3@netappci1.onmicrosoft.com	SSO (Guest)

A 'Single Sign-On' dialog box is open, showing the 'Default role for SSO users' dropdown set to 'Guest'.

Switching to other Cloud Central products

You can switch to other NetApp cloud products by clicking on the **Cloud Central** menu and selecting from the NetApp Cloud products listed. You can also choose to go to the Cloud Central portal itself.



You will be automatically authenticated and signed in to the selected cloud product.

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	Type
ActiveMQ	Service
Amazon EC2 and EBS	Infrastructure
Apache	Service
Azure NetApp Files	Infrastructure
Azure VMs and VHD	Infrastructure
Brocade Network Advisor (BNA)	Infrastructure
Brocade Fibre Channel Switches	Infrastructure
Cisco MDS Fabric Switches	Infrastructure
Consul	Service
Couchbase	Service
CouchDB	Service
Dell EMC Data Domain	Infrastructure
Dell EMC ECS	Infrastructure
Dell EMC Isilon	Infrastructure
Dell EMC PowerStore	Infrastructure
Dell EMC Recoverpoint	Infrastructure
Dell EMC ScaleIO	Infrastructure
Dell EMC Unity	Infrastructure
Dell EMC VMAX/PowerMax Family of Devices	Infrastructure
Dell EMC VNX Block Storage	Infrastructure
Dell EMC VNX File	Infrastructure
Dell EMC VNX Unified	Infrastructure
Dell EMC VPLEX	Infrastructure
Dell EMC XtremIO	Infrastructure
Dell XC Series	Infrastructure

Data Collector	Type
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
HPE Nimble Storage	Infrastructure
HP Enterprise 3PAR StoreServ 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
Lenovo HX Series	Infrastructure
macOS	Service
Memcached	Service
Microsoft Azure NetApp Files	Infrastructure
Microsoft Hyper-V	Infrastructure
MongoDB	Service

Data Collector	Type
MySQL	Service
NetApp Cloud Volumes ONTAP	Infrastructure
NetApp Cloud Volumes Services for AWS	Infrastructure
NetApp Data ONTAP 7-Mode	Infrastructure
NetApp E-Series	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
Ubuntu & Debian	Service
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 Cloud Central](#) 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](#).

Editions

The features and functionality available to you in Cloud Insights depend on the Edition to which you subscribe. The Editions available are explained here.

- **Basic Edition** is free and available to existing NetApp customers with an active NetApp support account.



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 90 consecutive days.



The features, data retention times, and objects or metrics collected in Cloud Insights Basic Edition are subject to change with or without notice.

- **Standard Edition** is available via subscription and offers all the features of Basic Edition plus more.
- **Premium Edition** offers additional benefits such as Business Intelligence and Reporting, as well as Cloud Secure Auditing and Threat Detection.


Key Features

Here are the key features available in Basic, Standard, and Premium Edition:

Key Feature	Basic Edition	Standard Edition	Premium Edition
Data Retention	7 Days	90 Days	13 Months
Infrastructure & Storage Metrics	NetApp Only	Multi-Vendor	Multi-Vendor
Customizable Dashboards	✓	✓	✓
Forum, Documentation, and Training Videos	✓	✓	✓
Live Chat and Technical Support	-	✓	✓
VM Metrics	-	✓	✓
Cloud Metrics	-	✓	✓

Service Metrics	-	✓	✓
Monitors and Alerting	-	✓	✓
API Access	-	-	✓
Single Sign-On (SSO)	-	-	✓
Cloud Secure User Data Access Auditing	-	-	✓
Cloud Secure Insider Threat Detection with AI/ML	-	-	✓
Business Intelligence and Reporting*	-	-	✓

*Available for environments of 500 managed units and larger

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 Version

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 all the features that Cloud Insights Standard Edition 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.

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.

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 **Subscribe Now** button. If your environment includes or is expected to include 1,000 or more Managed Units, you are eligible for Volume Pricing.

1. Estimate Your Managed Unit (MU) Usage

Hosts	Unformatted Capacity (TB)	Subscription Term
<input type="text" value="0"/> = 0 Managed Units	<input type="text" value="28"/> = 7 Managed Units	<input type="button" value="12 Months"/> <input type="button" value="36 Months"/>
Based on Current Usage		Contact sales for custom terms

2. Choose an Edition That's Right for Your Business

	Basic Free Available Only for NetApp Customers Downgrade	Standard \$42 / mo* 7 Managed Units at \$6 MU/mo Billed Annually Contact Sales Or Subscribe Via Amazon Marketplace	Premium \$63 / mo* 7 Managed Units at \$9 MU/mo Billed Annually Contact Sales Or Subscribe Via Amazon Marketplace
Key Features			
Data Retention	7 Days	90 Days	13 Months
Infrastructure and Storage Metrics	NetApp Only	Multi-Vendor	Multi-Vendor
Customizable Dashboards	✓	✓	✓
Real-time Dashboards	✓	✓	✓
Forum, Documentation and Training Videos	✓	✓	✓
Live Chat and Technical Support	—	✓	✓
VM Metrics	—	✓	✓
Cloud Metrics	—	✓	✓
Service Metrics	—	✓	✓
Monitors and Alerting	—	✓	✓
API Access	—	—	✓
Single Sign-On (SSO)	—	—	✓
Cloud Secure User Data Access Auditing	—	—	✓
Cloud Secure Insider Threat Detection with AI/ML	—	—	✓
Business Intelligence and Reporting	—	—	Requires at least 500 MUs

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 TB of unformatted capacity of physical or virtual disks

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 Calculator gives you an estimated list-price monthly Cloud Insights cost based on the number of hosts and amount of unformatted capacity being reported by your data collectors. The current values are

pre-populated in the *Hosts* and *Unformatted Capacity* fields. You can enter different values to assist you with planning for estimated future growth.

Your estimated list price cost will change based on your subscription term.



The calculator is 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 through 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 **Subscribe Now** button 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 Cloud Central 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.

Subscription Mode


Once your subscription is active, you can view your subscription status and Managed Unit usage on the **Admin > Subscription** page.

Subscription Summary

NetApp Serial Number : 1234568796059595050
Active Edition: Standard

Usage and Entitlement

38 out of 50 Managed Units



Hosts: 13 Managed Units (13 Hosts)

Unformatted Capacity: 25 Managed Units (50 TB)

Subscription Details

Multiple Active Subscriptions

Total 50 Managed Units with some valid through March 20, 2021

[View/Modify Managed Units](#)

The Subscription status page displays the following:

- Current subscription or active Edition
- Details about your subscription(s)
- Current Managed Unit usage, including breakdown counts for hosts and capacity



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.

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: <ul style="list-style-type: none">* Modify your subscription to increase the subscribed Managed Unit count* Remove Data Collectors so that your Managed Unit usage is at or below your subscribed amount

Installed Data Collectors

Click on the **View Data Collectors** button to expand the list of installed Data Collectors.

Installed Data Collectors (9)					
		Bulk Actions ▼		Filter...	
<input type="checkbox"/>	Name	Type	Total Managed Units ↓	Host Managed Units	Storage Managed Units
	aws	Amazon EC2	734	733	1
	aws2	Amazon EC2	732	731	1
	vm1	VMware vSphere	116	116	
	vm_2_	VMware vSphere	21	21	
	vm_2	VMware vSphere	21	21	

The Data Collectors section shows the Data Collectors installed in your environment and the breakdown of Managed Units for each.



The sum of Managed Units may differ slightly from the Data Collectors count in the status 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 Status 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 delete data collectors in this list by clicking on the "three dots" menu and selecting **Delete**.

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.

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 reliably 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	<ul style="list-style-type: none">* 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	<ul style="list-style-type: none">* 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 string 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.
4. 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**
3. Click the **Fibre Channel Identify** tab.

The system displays the identified and unidentified devices.

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*. 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	Type
-----	----	------	------

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**
3. 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	Type
-----	----	------	------

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.

+ Add							Total coverage
							20% (2/10)
IP identify (10)							Identify Unidentify <input type="text" value="filter..."/>
<input type="checkbox"/>	Address	IP	Name	Status	iSCSI node	Share name	
<input type="checkbox"/>	1.1.1.1	1.1.1.1	LA3-CNS-SQL-06A	OK		/vol/ServerLogs_STG/	
<input type="checkbox"/>	0.0.0.0/0					/vol/ServerLogs_STG/	
<input type="checkbox"/>	10.56.100.18				iqn.1991-05.com.microsoft:la3-cns-sql-06b.cns.comcastnets.com		
<input type="checkbox"/>	10.56.100.19				iqn.1991-05.com.microsoft:jec20643597717.tfyd.com	/vol/wc_sc_libraries_prod/libraries_qtree/	
<input type="checkbox"/>	100.54.18.100	100.54.18.100	ushapl000961b	OK			

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**  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 vendors 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.

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:

```
. *? _ . *? _ ( [a-zA-Z0-9] * [ABT] ) _ . *
```

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

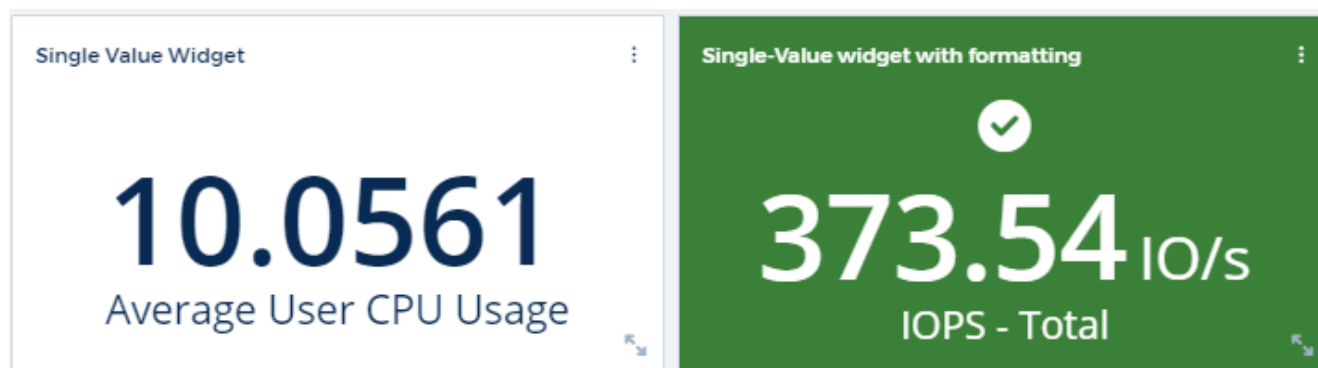
You can choose from the following widget types:

- **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.


4 items found in 2 groups

Active Date	Storage Node	Cache Hit Ratio - Total (%)	IOPS - Total (IO...	IOPS - Write (L...	Latency
06/01/2020 (1)	ocinaneqa1-01	N/A	N/A	N/A	N/A
06/01/2020	ocinaneqa1-01	N/A	N/A	N/A	N/A
N/A (3)	--	N/A	N/A	N/A	N/A


- **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.



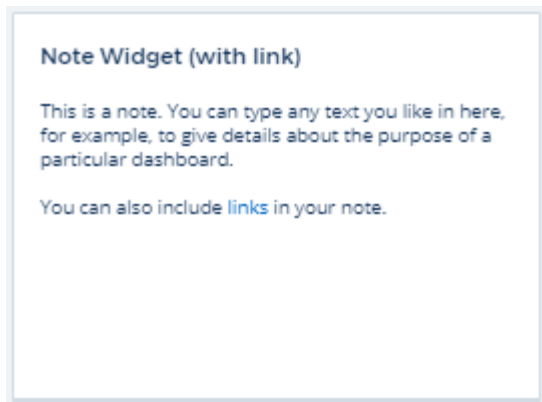
- **Scatter Plot Chart:** Plots related data as points, for example, IOPS and latency. In this example, you can quickly locate assets with high throughput and low IOPS.



- **Pie Chart:** a traditional pie chart to display data as a piece of the total.



- **Note widget:** Up to 1000 characters of free text.



- **Violations Table:** Displays up to the last 1,000 performance policy violations.

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 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.

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 queries

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.

Advanced Filtering

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*
```

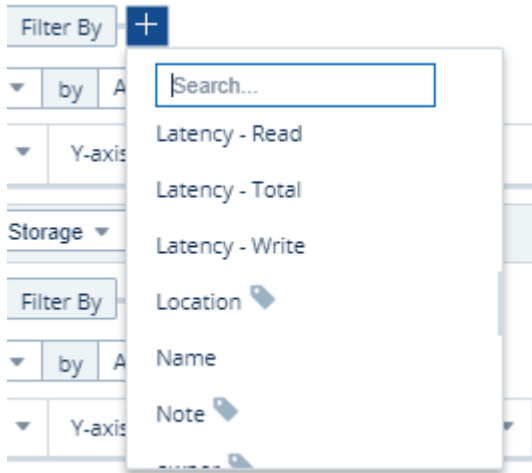
finds everything that does not start with "EMC". You can use

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:

- **Avg**: displays each line as the *average* of the underlying data.
- **Max**: displays each line as the *maximum* of the underlying data.
- **Min**: 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 *Avg*) 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 *Avg*, *Max*, *Min*, or *Sum*, or by the *Last* data point collected during the chosen interval. To choose an aggregation method, click on **More options** in the widget's query section.

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 *Avg* 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 *Avg*.

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.
3. 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 or 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 *Avg*. 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 *Avg*, *Sum*, *Min* or *Max*. 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 of 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) allows you to build expressions from metrics you choose, and show the result of those expressions in a single graph. 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.

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:

$$\text{Read Percentage} = (\text{Read IOPS} / \text{Total IOPS}) \times 100$$

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.

You can also use expressions in Line chart, Spline chart, or Stacked Area chart widgets.

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.

9. 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.

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.

Before you begin

The example below requires the **City** annotation (also called City attribute) to be set on multiple storage assets. For best results, set different cities on different storages. See the [Annotations](#) topics for more information on using annotations.

About this task

Variables provide a quick and simple way of filtering the data shown in some or all of the widgets on a custom dashboard. The following steps will guide you to creating widgets that use variables, and show you how to use them on your dashboard.

Steps

1. Click on **Dashboards > +New Dashboard**.
2. Before adding widgets, you must define the variables we will use to filter the dashboard data. Click on the **Add Variable** button.

The list of attributes is displayed.

3. Let's say we want to set the dashboard to filter based on City. Select the *City* attribute from the list.

The \$city variable field is created and added to the dashboard. Variables used by the dashboard are displayed above any widgets.

4. Next, we must tell our widgets to use this variable. The simplest way to illustrate this is to add a table widget showing the *City* column. Click on the **Add Widget** button and select the *Table* widget.
5. First, add the *City* column to the table by selecting it from the "gear" button.

City is a list-type attribute, so it contains a list of previously-defined choices. You may also choose text, boolean, or date-type attributes.

6. Next, click the **Filter by +** button and choose *City*.
7. Click *Any* to view the possible filter choices for *City*. Notice that the list now includes "**\$city**" at the top, in addition to any previously-available choices. Select *\$city* to use this dashboard variable.

The *\$city* choice only appears here if it was defined previously on the main dashboard page. If the variable was not previously defined, only the existing choices for the filter will be shown. Only variables that are applicable to the selected attribute type will be displayed in the drop-down for that filter.

8. **Save** the widget.
9. On the dashboard page, click on *Any* next to the *\$city* variable, and select the city or cities you want to see.

Your table widget updates to show only the cities you selected. You can change the values in the *\$city* variable at will, and all widgets on your dashboard that are set to use the *\$city* variable will refresh automatically to show only data for the values you selected.

Be sure to **Save** your dashboard when you have it configured as you want it.

More on dashboard variables

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:

- **Text:** Alphanumeric string. This is the default variable type.
- **Numerical:** a number or range of numbers.
- **Boolean:** Use for fields with values of True/False, Yes/No, 0/1, etc. For the boolean variable, the choices are Yes, No, None, Any.
- **Date:** A date or range of dates.

"Generic" variables

You can set a generic or universal variable by clicking the **Add Variable** button and selecting one of the types listed above. These types are always shown at the top of the drop-down list. The variable is given a default name, for example "\$var1", and is not tied to a specific annotation or attribute.

Configuring a generic variable allows you to use that variable in widgets to filter for any field of that type. For example, if you have a table widget showing Name, Alias, and Vendor (which are all text-type attributes), and "\$var1" is a text-type variable, you can set filters for each of those fields in the widget to use the \$var1 variable. You can set other widgets to use \$var1 for those or any text fields.

On your dashboard page, setting \$var1 to a value (for example "NetApp") will filter all of those fields in all widgets that are set to use that variable. In this way, you can update multiple widgets at once to highlight

dashboard data you choose at will.

Because generic variables can be used for any field of that type, you can change the name of a generic variable without changing its functionality.

Note: All variables are treated as "generic" variables, even those you create for a specific attribute, because all configured variables of a type are shown when you set a filter for any attributes or annotations of that type. However, best practice is to create a generic variable when you will use it to filter for a value across multiple fields, as in the Name/Alias/Vendor example above.

Variable naming

Variables names:

- Must always be prefixed with a "\$". This is added automatically when you configure a variable.
- Cannot contain any special characters; only the letters a-z and the digits 0-9 are allowed.
- Cannot be longer than 20 characters, including the "\$" symbol.
- Are case-sensitive: \$CityName and \$cityname are different variables.
- Cannot be the same as an existing variable name.
- Cannot be only the "\$" symbol.

Widgets that use variables

Variables can be used with the following widgets:

- Area Chart
- Bar Chart
- Box Plot Chart
- Line Chart
- Scatter Plot Chart
- Single Value Widget
- Spline Chart
- Stacked Area Chart
- Table Widget
- Pie Chart

Understanding "\$this" variables

Special variables on an asset's landing page allow you to easily showcase additional information that is directly related to the current asset. These special variables have names beginning with '\$this'.

1. About this task

To use the "\$this" variables in widgets on your asset's landing page, follow the steps below. For this example, we will add a **table widget**.



"\$this" variables are only valid for an asset's landing page. They are not available for other dashboards. The available "\$this" variables varies according to asset type.

Steps

1. Navigate to the landing page for an asset of your choosing. For this example, let's choose a Virtual Machine (VM) asset page. Query or search for a VM and click on the link to go to that VM's asset page.

The asset page for the VM opens.

2. Click **Edit** to switch to edit mode, and click the **Add Widget** button. Choose the **Table** widget.

The Table widget opens for editing. By default, all storages are shown in the table.

3. We want to show all virtual machines. Click on the asset selector and change *Storage* to *Virtual Machine*.

All virtual machines are now shown in the table.

4. Click on the gear button and add the *Hypervisor Name* column to the table.

The hypervisor name is shown for each VM in the table.

5. We only care about the hypervisor that hosts the current VM. Click on the **Filter by** field's **+** button and select *Hypervisor Name*.

6. Click on *Any* and select the ***\$this.host.name*** variable. Press Enter or click off the field to apply the filter.

The table now shows all the VM's hosted by the current VM's hypervisor.

7. Click **Save** to save the widget.

8. Click **Save** to save the asset page.

Result

The table that you created for this VM asset page will be displayed for any VM asset page you display. The use of the *\$this.host.name* variable in the widget means that only the VM's owned by the *current assets's* hypervisor will be displayed in the table.

You can also apply **in-context filters** to asset page widgets to accomplish a similar result.

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.

Widget 12

☐ Override Dashboard Time

⌚

✕

✓ A) Query

Storage.performance.iops.total ✕

⌵

Filter By

+

Group

Avg

Time aggregate by

Avg

Less Options

Formatting: If value is

>

Warning

500

IO/s and/or

Critical

1000

IO/s

Showing

🟢 In Range as green

Description

IOPS - Total

Calculation

A

Min Value

Optional

Max Value

1200

Display:

Bullet Gauge

Decimal Places:

2

Color:

⊗

Units Displayed In:

Auto Format

+ Query

200

400

600

800

1k

1.2k

904.21 IO/s

IOPS - Total


Cancel


Save

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.

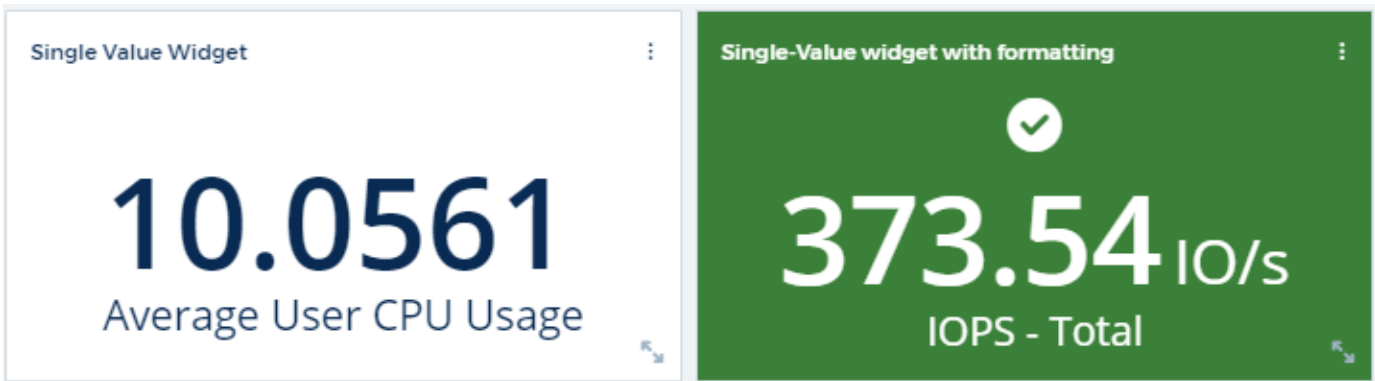
You can optionally choose a minimum and maximum value for the gauge. Values below minimum will not display the gauge. Values above maximum will display a full gauge. If you do not choose minimum or maximum values, the widget selects optimal min and max based on the widget's value.





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.


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.

Auto Format ▼

Base Unit

bit/sec (b/s) ▼

Displayed In

Data Rate (IEC)

bit/sec (b/s)

byte/sec (B/s)

kibibyte/sec (KiB/s)

mebibyte/sec (MiB/s)

gibibyte/sec (GiB/s)

10:30 AM

11:00 A

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 billion 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.

Dashboard Groups (3)



All Dashboards (60)

My Dashboards (11)

Storage Group (7)



Dashboards (7)



Name ↑

Dashboard - Storage Cost

Dashboard - Storage IO Detail

Dashboard - Storage Overview

Gauges Storage Performance

Storage Admin - Which nodes are in high demand?

Storage Admin - Which pools are in high demand?

Storage IOPs

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, or 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)

<input type="checkbox"/>	Name ↑
	Dashboard - Storage Overview
	Storage Admin - Which nodes are in high demand?
	Storage IOPs
	Dashboard - Storage Cost
	Dashboard - Storage IO Detail
	Gauges Storage Performance
	Storage Admin - Which pools are in high demand?

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.
2. From the **Dashboards** menu, select **[+New dashboard]**.

The **New dashboard** page opens.

3. At the top of the page, enter a unique name for the dashboard, for example "VM Performance by Application".
4. Click **Save** to save the dashboard with the new name.
5. Let's start adding our widgets. If necessary, click the **Edit** icon to enable Edit mode.
6. 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.

Table Widget

10m

1,746 items found in 71 groups

Hypervisor Name ↑	Virtual Machine	Capacity - Total (GB)	IOPS - Total (IO/s)	Latency - Total (ms)
10.197.143.53 (9)	--	1,690.58	1.80	12.04
10.197.143.54 (7)	--	1,707.60	4.62	12.69
10.197.143.57 (11)	--	1,509.94	1.14	1.15
10.197.143.58 (10)	--	1,818.34	5.83	2.57
AzureComputeDefaultAvailabilitySet (363)	--	N/A	N/A	N/A
anandh9162020113920-rg-avset.anandh91620201	--	N/A	N/A	N/A
anandh916202013287-rg-avset.anandh91620201	--	N/A	N/A	N/A
anandh91720201288-rg-avset.anandh91720201	--	N/A	N/A	N/A
anjalivIngrun48-rg-avset.anjalivIngrun48-rg.398	--	N/A	N/A	N/A
anjalivIngrun50-rg-avset.anjalivIngrun50-rg.398	--	N/A	N/A	N/A
batutiscanaryHA97a-rg-avset.batutiscanaryha97	--	N/A	N/A	N/A
batutiscanaryHA97b-rg-avset.batutiscanaryha97	--	N/A	N/A	N/A

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.

3. 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 outside 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.

Virtual Machine Performance Table

☐ Override dashboard time

Last 24 hours

×

Virtual Machine

Filter by

IOPS - Total (IO/s)

>= 10

×

+

Group by

Hypervisor name

×

181 items found in 4 groups

Hypervisor name ↓	Name	Hypervisor name	IOPS - Total	Latency - Total (ms)
+	us-east-1d (62)	us-east-1d		1.94
+	us-east-1c (80)	us-east-1c		0.80
+	us-east-1b (1)	TBDemoEnv	32.66	0.70
+	us-east-1a (38)	us-east-1a	121.22	0.81

Cancel

Save

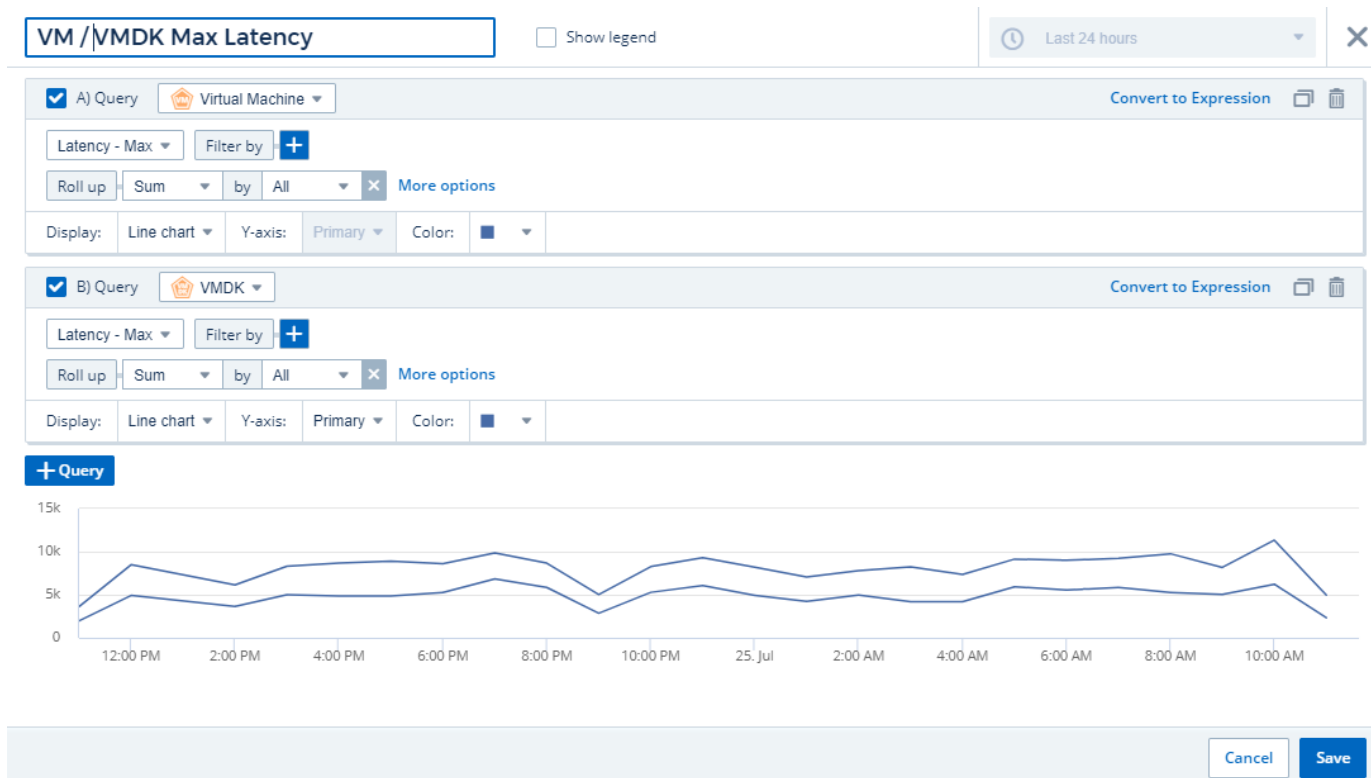
- When you have customized the table widget to your satisfaction, click the **[Save]** button.

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.

- If necessary, click the **Edit** icon on the dashboard to enable Edit mode.
- Click the **[Add widget]** icon and select *Line Chart* to add a new line chart widget to the dashboard.
- The **Edit Widget** dialog opens. Name this widget "VM / VMDK Max Latency"
- 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*.
- 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*.
- 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*.
5. 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

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
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.

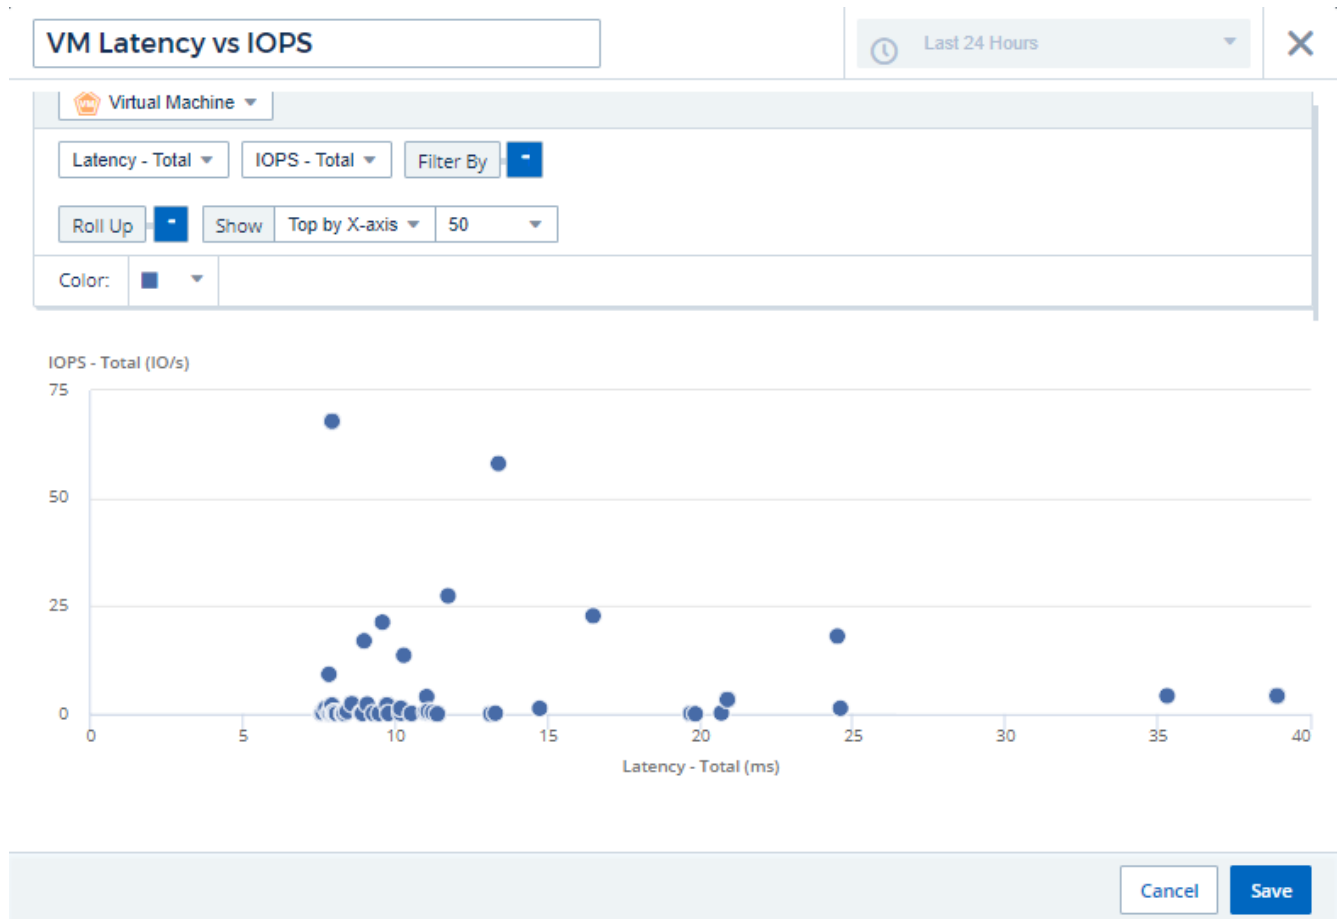
Steps

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 Explorer

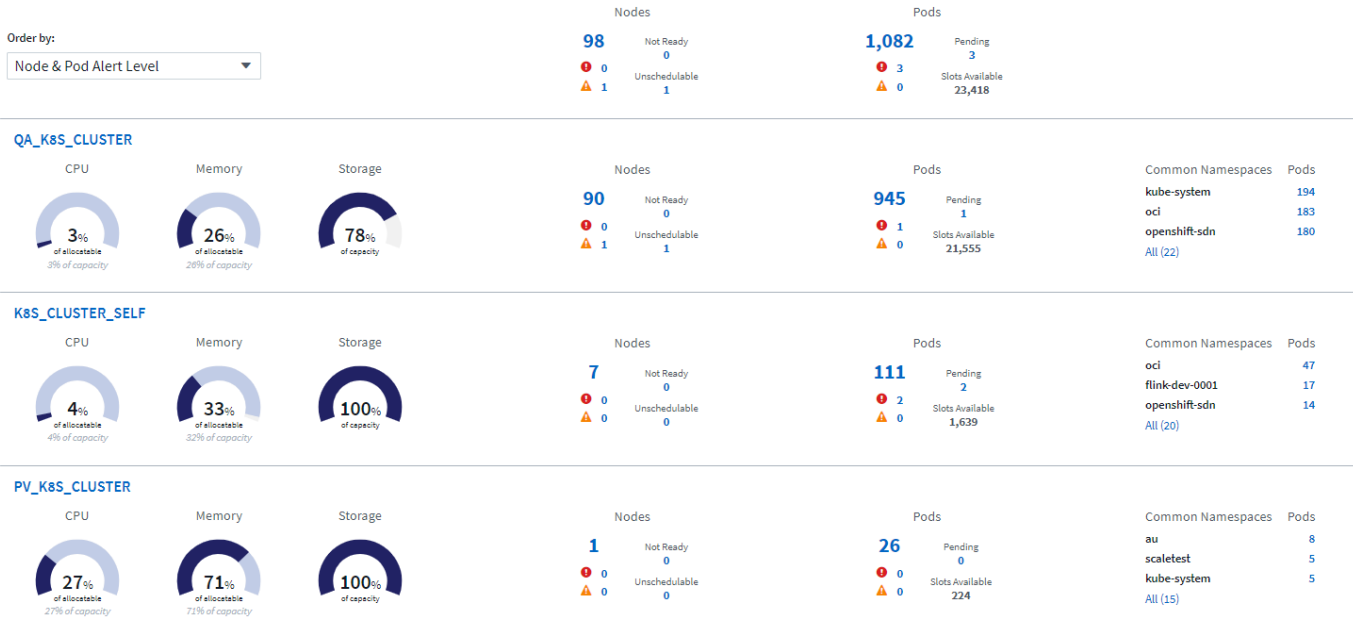
Kubernetes Cluster Overview

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 overview page. This overview page contains a wide variety of at-a-glance information.

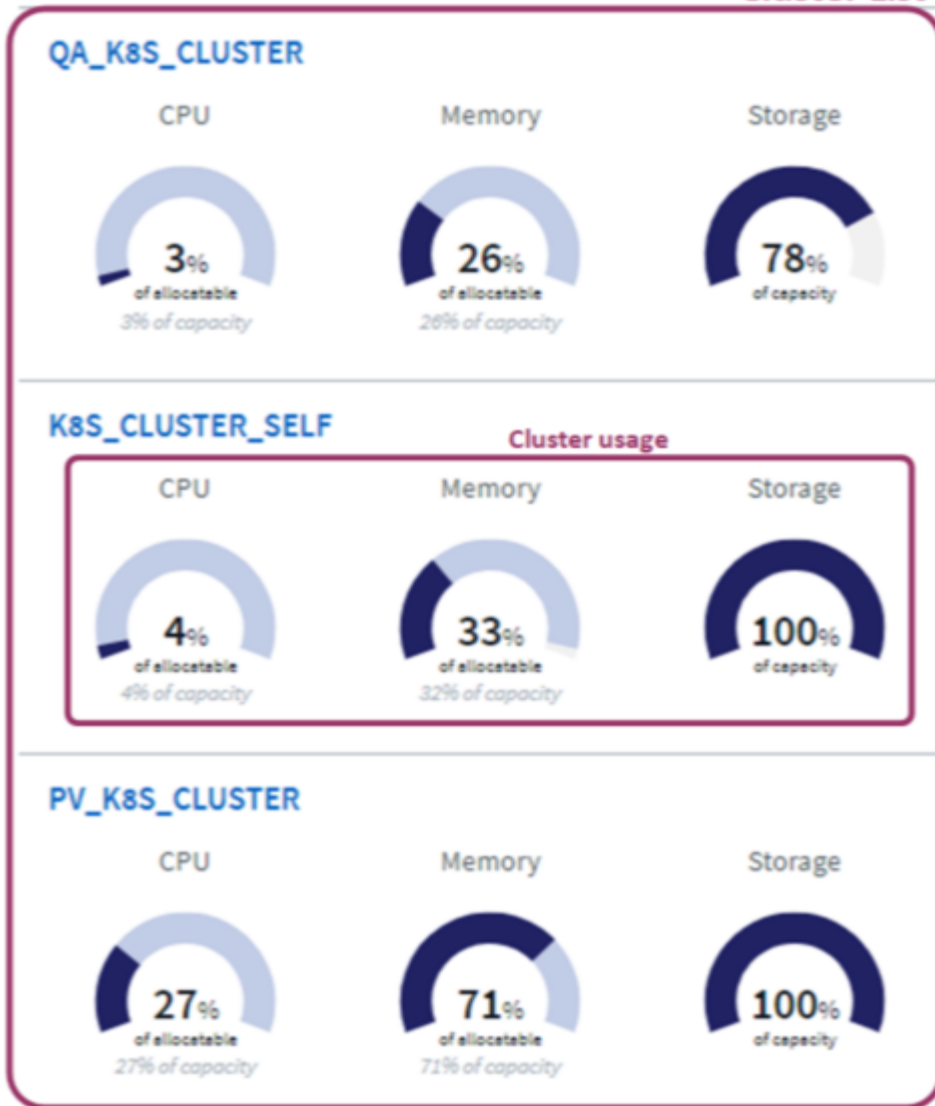


Numbers displayed in blue in the various sections of this and subsequent Kubernetes Explorer pages (for example, node/pod status, namespace counts, etc.) are links to related query pages that show the details behind those numbers.



Cluster Usage overview

Cluster List



The cluster list displays the following usage information for each cluster in your environment:

- CPU: percentage of total CPU capacity in use
- Memory: percentage of total memory used
- Storage: percentage of total storage in use

You can sort the cluster list by any of the following factors:

- Node & Pod Alert Level (default)
- Cluster Name
- Number of Nodes
- Most Utilized by Compute
- Least Utilized by Compute
- Most Utilized by Storage
- Least Utilized by Storage

Clicking on a Cluster Name will open the [detail page](#) for that cluster

Node and Pod Status



Namespaces

To the right of the screen is a list of the top three namespaces utilized in each cluster. Click the **All** link to see all namespaces for the cluster.

Top Namespaces

Common Namespaces	Pods
-------------------	------

kube-system	194
-------------	-----

oci	183
-----	-----

openshift-sdn	180
---------------	-----

All (22)	Click to see all namespaces
--------------------------	---

Common Namespaces	Pods
-------------------	------

oci	47
-----	----

flink-dev-0001	17
----------------	----

openshift-sdn	14
---------------	----

All (20)	
--------------------------	--

Common Namespaces	Pods
-------------------	------

au	8
----	---

scaletest	5
-----------	---

kube-system	5
-------------	---

All (15)	
--------------------------	--

Kubernetes Cluster Detail Page


The Kubernetes cluster detail page displays detailed information about 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.

My_Cl_Site / Kubernetes / * My_Cluster / ip-10-30-12-200 / ds-4dbmk


Cluster Overview

The cluster overview page provides useful information at a glance:

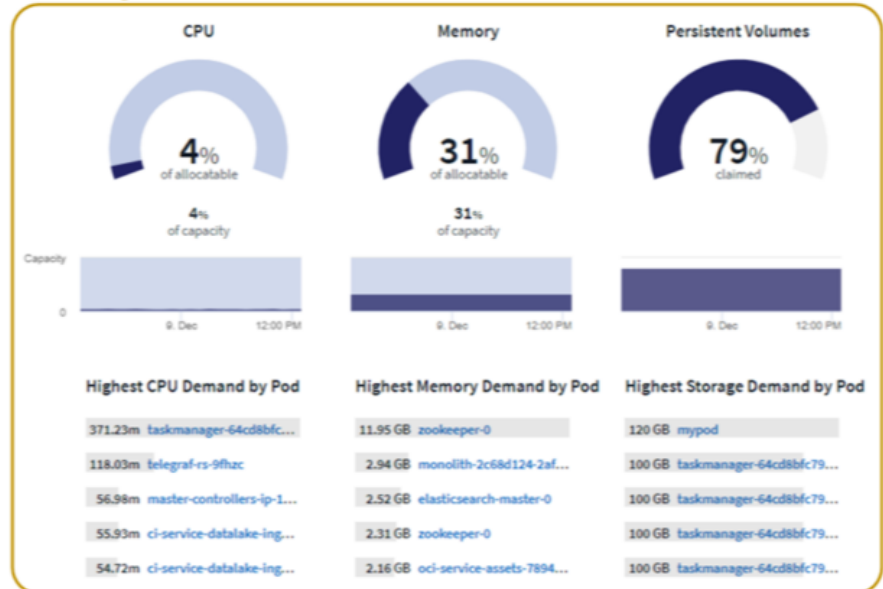


Order by: Node & Pod Alert Level

Cluster Health at a Glance




Resource Usage



Node and Pod Counts

The Node/Pod counts at the top of the page show you the total number of nodes and pods in the cluster, as well as a breakdown of how many pods are currently alerting or pending.



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).

The 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.




Clicking on the pod or node in that view will zoom in to the 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.






You can view Filesystem, Memory, and CPU information by hovering over the Node elements.



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:

Labels	Node IP
-	10.30.23.207



Pods		Containers			
Status ↑		Name	Healthy Containers	Namespace	
❗ Critical	Pending	demo-pod2	1 of 2	k8wheel	
● Healthy	Running	ci-exclusive-node-scheduler-6dc4dd96-s6h9t	2 of 2	kafka-lake-0001	
● Healthy	Running	ci-service-apikey-7676fd5f7d-ptmh9	1 of 1	oci	
● Healthy	Running	ci-service-notifications-7f594c4bbd-4p7hz	1 of 1	oci	
● Healthy	Running	ci-service-webui-rest-5d454c8648-98llk	1 of 1	oci	
● Healthy	Succeeded	job-odata-2c68d124-2af5-4b6b-864f-f04c04e77de5-75fnf	1 of 1	oci	

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*. Storage gauges only show two colors because they only show a single pair of values: *used* in relation to *total*.


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).



Working with Queries

Assets used in queries

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

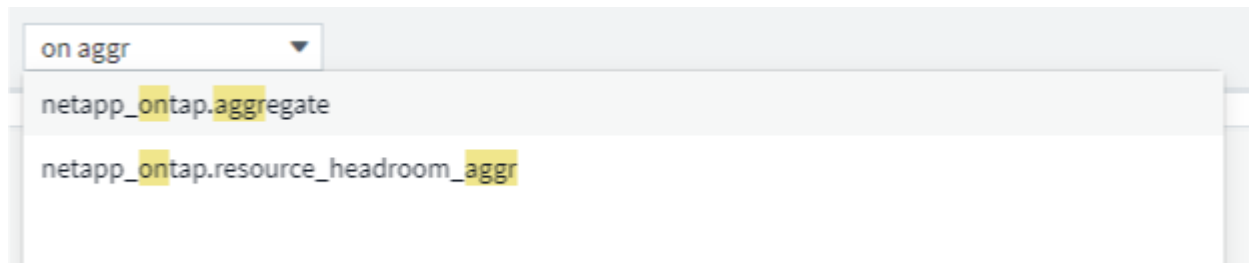
The following asset types can be used in queries, dashboard widgets, and custom asset landing pages. The fields and counters available for filters, expressions, and display will vary among asset types. Not all assets can be used in all widget types.

- 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.



Creating Queries

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.

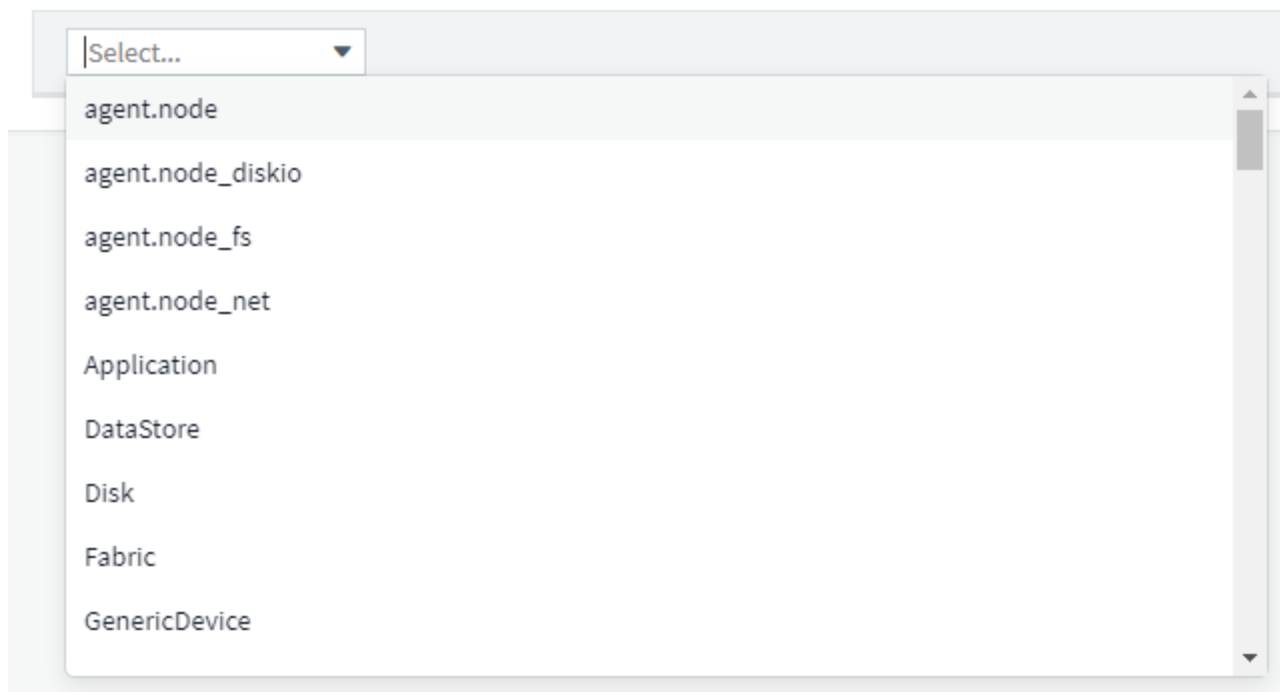
For example, you can create a query for *volumes*, add a filter to find particular *storages* associated with the selected volumes, add another filter to find a particular *annotation* such as "Tier 1" on the selected storages, and finally add another filter to find all storages with *IOPS - Read (IO/s)* 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:

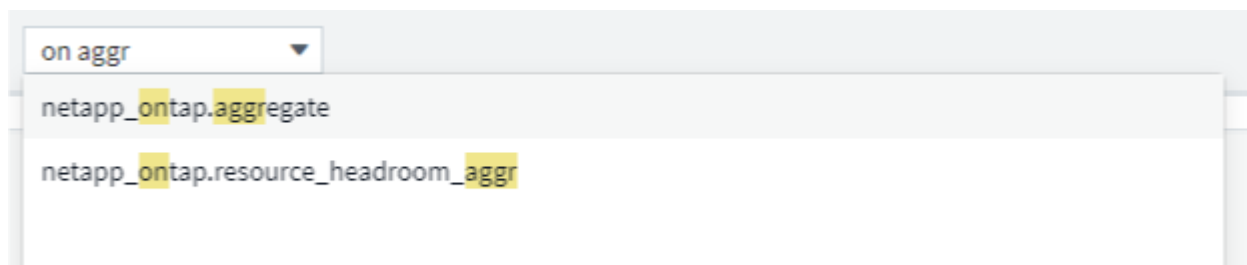
1. Navigate to **Queries > *+New Query**.
2. From the 'Select...' list, select the object type you want to query for. You can scroll through the list or you can start typing to more quickly find what you're searching for.

Scroll list:



A screenshot of a web interface showing a dropdown menu. The dropdown is open, displaying a list of metrics. The first item is highlighted. The list includes:

- agent.node
- agent.node_diskio
- agent.node_fs
- agent.node_net
- Application
- DataStore
- Disk
- Fabric
- GenericDevice

Type-to-Search:

A screenshot of a web interface showing a 'Type-to-Search' dropdown menu. The dropdown is open, displaying a list of search results. The first item is highlighted. The list includes:

- netapp_ontap.aggregate
- netapp_ontap.resource_headroom_aggr

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.

netapp_ontap.aggregate X ▼

Filter By cluster_name ci- X +

Group aggr_name X X ▼

5 items found

Table Row Grouping	Metrics & Attributes	
aggr_name	cp_read_blocks	cluster_name ↓
oci02sat0	0.59	oci-phonehome
oci02sat1	0.15	oci-phonehome
oci02sat2	212.64	oci-phonehome
oci01sat0	0.39	oci-phonehome
oci01sat1	48.89	oci-phonehome

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 vary based on the asset/metric type.

netapp_ontap.aggregate X ▼

Filter By +

Group aggr_name X X ▼

14 items found

Table Row Grouping	Metrics & Attributes	
aggr_name	cp_read_blocks	agent_version ↑
aggr0_optimus_02	1.72	Apache-HttpClie
aggr1_optimus_02	408.84	Apache-HttpClie
ocinaneqa1_04_aggr0	6.19	Apache-HttpClie
ocinaneqa1_03_aggr0	6.48	Apache-HttpClie
oci02sat0	1.04	Apache-HttpClie

- ☐ Show Selected Only
- ☒ agent_version
- ☐ aggr_name
- ☐ cluster_location
- ☒ cluster_name
- ☐ cluster_serial_number
- ☐ cluster_version

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

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 NULL values in all fields	None	returns results where the target field is empty
Not *	searches for NULL values in <i>text-only</i> 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 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, simply 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.

Volume X

Filter By Name Any X +

Query Results (5) | 2 Selected

<input type="checkbox"/>	Name ↑	Storage Pools	Capacity - Raw (GB)	Mapped Ports
<input type="checkbox"/>	DmoESX_optimus:mc_Dm...	optimus-02:aggr1_optimu...	N/A	
<input checked="" type="checkbox"/>	DmoSAN_optimus:hoffma...	optimus-02:aggr1_optimu...	N/A	
<input checked="" type="checkbox"/>	DmoSAN_optimus:mc_D...	optimus-02:aggr1_optimu...	N/A	US:windows_zu08
<input type="checkbox"/>	oci-3070-01:/vol/vfiler_lun...	oci-3070-01:aggr5	N/A	OS:windows
<input type="checkbox"/>	spectrav1:sjimmylscsi:/v...	ocinaneqa1-01:spectraaggr1	N/A	OS:linux

Bulk Actions

- Add Annotation
- Remove Annotation
- Add Application
- Remove Application

Annotation Rules require query

If you are configuring [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.

Viewing queries

You can view your queries to monitor your assets and change how your queries display the data related to your assets.

Steps

1. Log in to your Cloud Insights tenant.
2. Click **Queries** and select **Show all queries**.
You can change how queries display by doing any of the following:
3. You can enter text in the filter box to search to display specific queries.
4. 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.
5. 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.
6. 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.


Exporting query results to a .CSV file

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

1. Log in to your Cloud Insights tenant.
2. Click **Queries** and select **Show all queries**.

The Queries page is displayed.

3. Click a query.
4. Click  to export the query results to a .CSV file.
5. When prompted, do one of the following:
 - a. Click **Open with** and then **OK** to open the file with Microsoft Excel and save the file to a specific location.
 - b. 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.

When exporting query results, be aware that all rows in the results table will be exported, not just those selected or displayed on the screen, up to a maximum of 10,000 rows.

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:

1. Open a new sheet in Excel.
2. On the "Data" tab, choose "From Text".
3. Locate the desired .CSV file and click "Import".
4. In the Import wizard, choose "Delimited" and click Next.
5. Choose "Comma" for the delimiter and click Next.
6. Select the desired columns and choose "Text" for the column data format.
7. Click Finish.

Your objects should show in Excel in the proper format.

Modifying or Deleting a Query


Modifying a Query

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.

Steps

1. Click **Queries** and select **Show all queries**.

The Queries page is displayed.

2. Click the query name
3. To add a criteria to the query, click  and select a criteria from the list.
4. To remove a filter from the query, click the **X** 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 query.
- 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:

1. Click on the "three dot" menu to the right of the query and click **Delete**.
2. Click on the query name and select **Delete** from the **Save** drop-down menu.

Assigning multiple applications to or removing multiple applications from assets

You can assign multiple applications to or remove multiple applications from assets by using a query instead of having to manually assign or remove them.



You can use these steps to add or remove 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

1. Click **Queries** and select **Show all queries**.


The Queries page displays.

2. Click the name of the query that finds the assets.

The list of assets associated with the query displays.

3. Select the desired assets in the list or click the top checkbox to select All.

The  button displays.

4. To add an application to the selected assets, click  and select **Add Application**.

5. 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.

6. Click **Save**.

7. To remove an application assigned to the assets, click  and select **Remove Application**.

8. Select the application or applications you want to remove.

9. 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.


Copying table values

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

1. Method 1: Highlight the desired text with the mouse, copy it, and paste it into search fields or other applications.
2. Method 2: For single-value fields, hover over the field and click the clipboard icon  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.

Working with Annotations

Defining 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.

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.

Annotation types	Description	Type
Alias	User-friendly name for a resource	Text
Compute Resource Group	Group assignment used by the Host and VM Filesystems data collector	List
Data Center	Physical location	List
Hot	Devices under heavy use on a regular basis or at the threshold of capacity	Boolean
Note	Comments associated with a resource	Text
Service Level	A set of supported service levels that you can assign to resources. Provides an ordered options list for internal volumes, qtrees, and volumes. Edit service levels to set performance policies for different levels.	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, qtrees, 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



Alias, Data Center, Hot, Service Level, Sunset, Switch Level, Tier, 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

1. In the Cloud Insights menu, click **Manage > Annotations**.

The Annotations page displays the list of annotations.

2. Click **+Add**
3. Enter a **Name** and **Description** of the annotation.

You can enter up to 255 characters in these fields.

4. 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.



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.

1. If you select List as the annotation type, do the following:
 - a. 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.
 - b. Enter a value and description in **Value** and **Description** fields.
 - c. Click **Add** to add additional values.
 - d. Click the Trash icon to delete a value.
2. 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.

Using annotations

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

1. In the Cloud Insights menu, click **Manage > Annotations**
2. Click **+ Annotation** to create a new annotation.
3. 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

1. Log in to your Cloud Insights environment.
2. Locate the asset to which you want to apply the annotation.
 - You can locate assets by querying, choosing from a dashboard widget, or search. When you have located the asset you want, click the link to open the asset's landing page.
3. On the asset page, in the User Data section, click **+ Annotation**.
4. The Add Annotation dialog box displays.
5. Select an annotation from the list.

6. 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.

7. 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.

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

1. Click **Manage > Annotation rules**

The Annotation Rules page displays the list of existing annotation rules.

2. Click **+ Add**.

3. Do the following:

- a. In the **Name** box, enter a unique name that describes the rule.

This name will appear in the Annotation Rules page.

- b. Click **Query** and select the query that is used to apply the annotation to assets.

- c. Click **Annotation** and select the annotation you want to apply.

- d. 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.

- e. Click **Save**

- f. Click **Run all rules** if you want to run all the rules immediately; otherwise, the rules are run at a regularly scheduled interval.

Creating annotation rules

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

1. In the Cloud Insights menu click **Manage > Annotation rules**.
2. Click **+ Rule** to add a new annotation rule.

The Add Rule dialog is displayed.

3. Do the following:
 - a. In the **Name** box, enter a unique name that describes the rule.

The name appears in the Annotation Rules page.
 - b. Click **Query** and select the query that Cloud Insights uses to identify the assets the annotation applies to.
 - c. Click **Annotation** and select the annotation you want to apply.
 - d. 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.

- e. Click **Save**
- f. Click **Run all rules** if you want to run all the rules immediately; otherwise, the rules are run at a regularly scheduled interval.



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

1. In the Cloud Insights menu, Click **Manage > Annotation rules**.

The Annotation Rules page displays the list of existing annotation rules.

2. 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.

3. Click the menu icon for the rule that you want to modify.
4. Click **Edit**

The Edit Rule dialog is displayed.

5. 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

1. Click on the menu icon for the rule you want to move.
2. Click **Move Up** or **Move Down** as needed until the rule appears in the location you want.

Deleting annotation rules

You might want to delete annotation rules that are no longer used.

Steps

1. In the Cloud Insights menu, Click **Manage > Annotation rules**.

The Annotation Rules page displays the list of existing annotation rules.

2. 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.

3. Click the menu icon for the rule that you want to delete.
4. Click **Delete**

A confirmation message is displayed, prompting whether you want to delete the rule.

5. Click **OK**

Importing Annotations

Cloud Insights includes an API for importing annotations or applications from a CSV file, and assigning them to objects you specify.



The Cloud Insights API is available in **Cloud Insights Premium Edition**.

Importing

The **Admin > API Access** links contain [documentation](#) for the **Assets/Import** API. This documentation contains information on the .CSV file format.

ASSETS.import

PUT /assets/import Import assets from a CSV file.

Import annotations and applications from the given CSV file. The format of the CSV file is following:

```
[Project]
[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]]
```

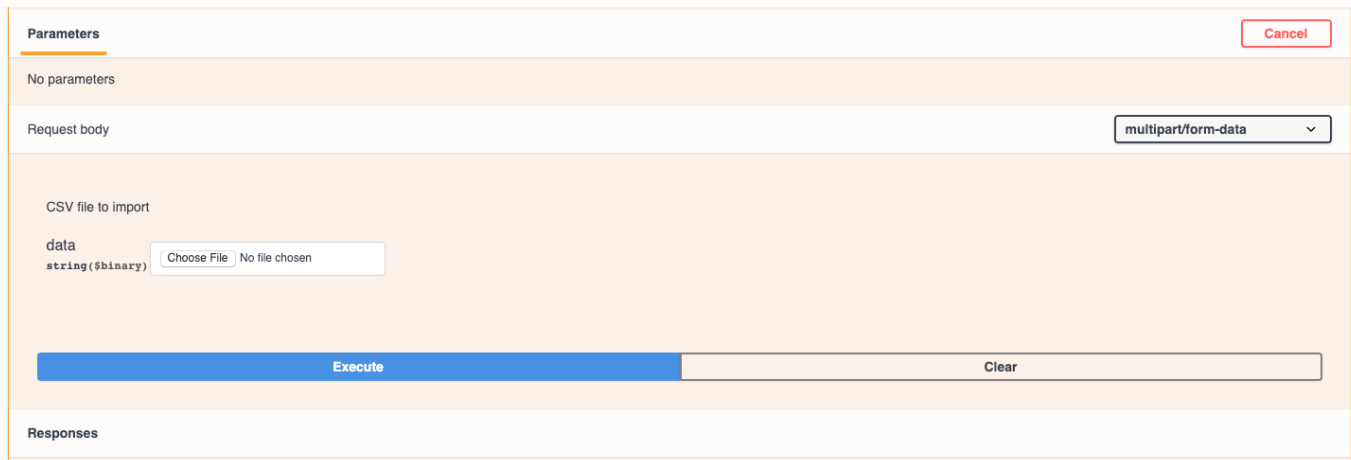
.CSV File Format

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.

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:



The image shows the Swagger UI for the `ASSETS.import` endpoint. It features a `Parameters` tab with a `Cancel` button. Below the parameters, there is a `Request body` section with a dropdown menu set to `multipart/form-data`. The request body is a `CSV file to import` with a `data` field of type `string($binary)`. A file selection interface shows `Choose File` and `No file chosen`. At the bottom, there are `Execute` and `Clear` buttons. The `Responses` section is empty.

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 <WWN>
Qtree	id-><id> or <Storage Name>-><Internal Volume Name>-><Qtree Name>
Share	id-><id> or <Storage Name>-><Internal Volume Name>-><Share Name>-><Protocol>[-><Qtree Name (optional in case of default Qtree)>]

Working with Applications

Tracking asset usage by application

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.

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:

1. Create an application named *Exchange*
2. 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".

3. Select the VMs.
4. Associate the VMs with the *Exchange* application using **Bulk Actions > Add Applications**.
5. Select the desired application and click **Save**.
6. When finished, **Save** the query.

Creating Applications

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.

- a. Enter a unique name for the application.
- b. Select a priority for the application.
- c. Click **Save**.

After you finish

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

1. Locate the asset to which you want to assign to the application:
2. Click **Queries > +New Query** and search for Host.
3. Click the check box on the left of the Host you want to associate with the application.
4. Click **Bulk Actions > Add Application**.
5. Select the Application you are assigning the asset to.

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.

Monitors and Alerts

Alerting with Monitors

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.



Monitors and Alerting is available in Cloud Insights Standard Edition and higher.

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.

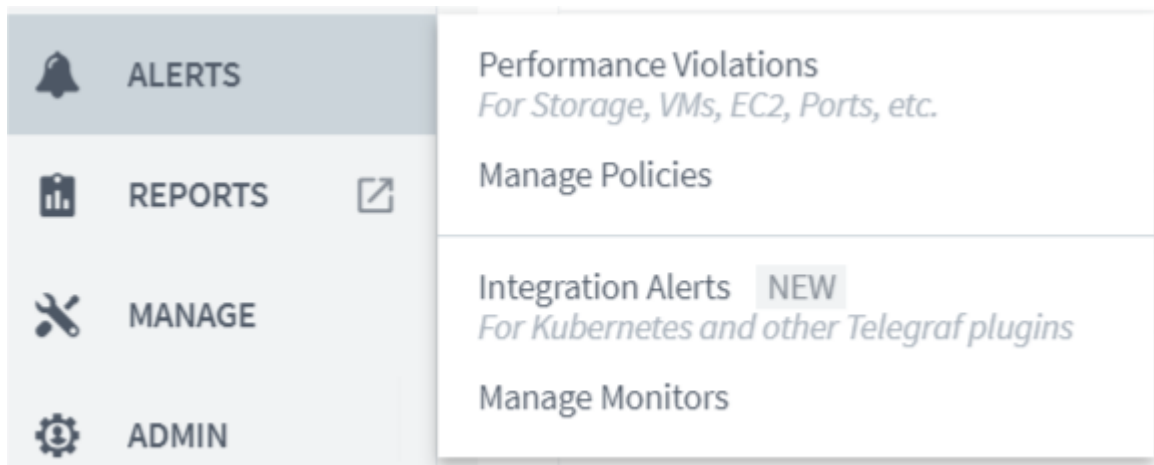
Monitor or Performance Policy?

What's the difference between a **Performance Policy** and a **Monitor**?

Policies allow you to set thresholds on "infrastructure" objects such as storage, VM, EC2, and ports. These policies trigger violations when thresholds are met or exceeded. Each violation can be investigated for troubleshooting. Policies are described in detail elsewhere in this [documentation](#).

Monitors provide similar functionality for "integration" data such as those collected for Kubernetes, ONTAP advanced metrics, and Telegraf plugins, and alert when thresholds are crossed. With Monitors, you can set thresholds for Warning- or Critical-level alerts, or both.

Policies and Monitors are available under the **Alerts** menu.



Emails can be sent when a policy or monitor is triggered.

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

1. From the Cloud Insights menu, click **Alerts > Manage Monitors**

The Monitors list page is displayed, showing currently configured monitors.


2. To add a monitor, Click **+ Monitor**. To modify an existing monitor, click the monitor name in the list.

The Monitor Configuration dialog is displayed.

3. 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.

1 Select a metric to monitor



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.



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.


1. After choosing the object and metric to monitor, set the Warning-level and/or Critical-level thresholds.
2. For the *Warning* level, enter 200. The dashed line indicating this Warning level displays in the example graph.
3. 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.

4. For the occurrence interval, choose *Continuously* for a period of *15 Minutes*.

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.


2 Define the monitor's conditions (set at least one threshold condition)



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.

3 Set up team notification(s) (alert your team via email, or Webhook)



Alerting via Email:

Specify the email recipients for alert notifications. If desired, you can choose different recipients for warning or critical alerts.

3 Set up team notification(s)

Email

Notify team on

Critical, Resolved

☒ Critical

☐ Warning

☒ Resolved

Add Recipients (Required)

user_1@email.com X

user_2@email.com X

Email

Notify team on

Warning

Add Recipients (Required)

user_3@email.com X

Alerting via Webhook:

Specify the webhook(s) for alert notifications. If desired, you can choose different webhooks for warning or critical alerts.

3 Set up team notification(s) (alert your team via email, or Webhook)

By Webhook

Slack

Notify team on

Critical

Use Webhook(s)

Slack X Teams X X

Notify team on

Resolved

Use Webhook(s)

Slack X Teams X X

Notify team on

Warning

Use Webhook(s)

Slack X Teams X X



Webhooks is considered a Preview feature and is therefore subject to change.

Save your Monitor

1. If desired, you can add a description of the monitor.
2. 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

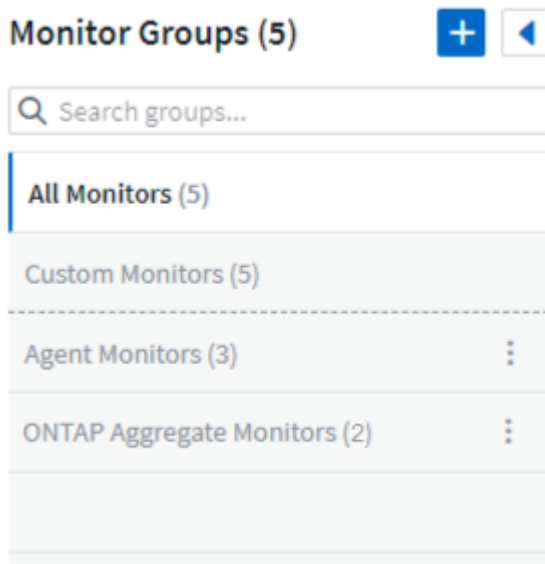
You can choose to temporarily suspend monitoring of an object type by clicking the menu to the right of the monitor and selecting **Pause**. When you are ready to resume monitoring, click **Resume**.

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.



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.

5 Associate to a monitor group (optional)

ONTAP Monitors

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.



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*.



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 *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*.

Monitor Groups (3)



Search: Agent Monitors

- All Monitors (4)
- Custom Monitors (4)
- Agent Monitors (3) 

Pause

Resume

Rename

Delete

More Information

- [Viewing and Dismissing Alerts](#)

Viewing and Managing Alerts from Monitors


Cloud Insights displays alerts when [monitored thresholds](#) are exceeded.



Monitors and Alerting is available in Cloud Insights Standard Edition and higher.

Viewing and Managing Alerts

To view and manage alerts, do the following.

1. Navigate to the **Alerts > Integration Alerts** page.
2. 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  :
 - **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

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.

Alert Summary

Monitor:

Volume Total Data

Triggered On:

cluster_name: tawny
aggr_name: Multiple_Values

Duration / Time Triggered:

1d 6h / Jun 9, 2020 2:22 AM

Top Severity:

❗ Critical

Metric:

① netapp_ontap.workload_volume.total_data

Condition:

Average total_data is > (greater than) 0m and/or 0m all the time in 2-hour window.

Filters Applied:

cluster_name: Any


Status:

New

Expert View

Display Metrics ▾

total_data (m)



Related Alerts

1 item found

Alert ID	Active Status	Triggered Time ↓	Top Severity	Monitor	Triggered On	Status
AL-46769	Resolved	a day ago Jun 9, 2020 2:22 AM	❗ Critical	Volume Total Data	cluster_name: tawny aggr_name: Multiple_Values	New

Comments

There are no comments yet on this alert.

[+ Comment](#)

Configuring Email Notifications

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.

To configure notification email recipient settings, go to the **Admin > Notifications** page.

Subscription Notification Recipients

Subscription Notification Recipients

Send subscription related notifications to the following:

☒ All Account Owners

☒ All Administrators

☒ Additional Email Addresses

Enter email addresses separated by commas.

Save

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 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 Subscription page
Subscription will expire in 90 days Subscription will expire in 30 days	No action needed if "Auto Renewal" is enabled Contact NetApp sales to renew the subscription
Trial ends in 2 days	Renew trial from the Subscription page. You can renew a trial one time. Contact 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 NetApp sales to purchase a subscription

Global Recipient List for Performance Policy Notifications

Global Performance Policy Recipients

Default email recipients for Performance Policy related notifications:

Recipients

Enter email addresses separated by commas.

Email Signature

Email signature added to messages sent by Cloud Insights

Save

To add recipients to the global performance policy notification email list, go to the "Global Performance Policy Recipients" section and enter email addresses separated by commas. Emails sent as alerts from performance policy threshold violations will be sent to all recipients on the list.

If you make a mistake, you can click on [x] to remove a recipient from the list.

You can also add an optional signature block that will be attached to the email notifications sent.



You can override the global list for a specific policy when you configure that policy.

Cloud Insights API

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.



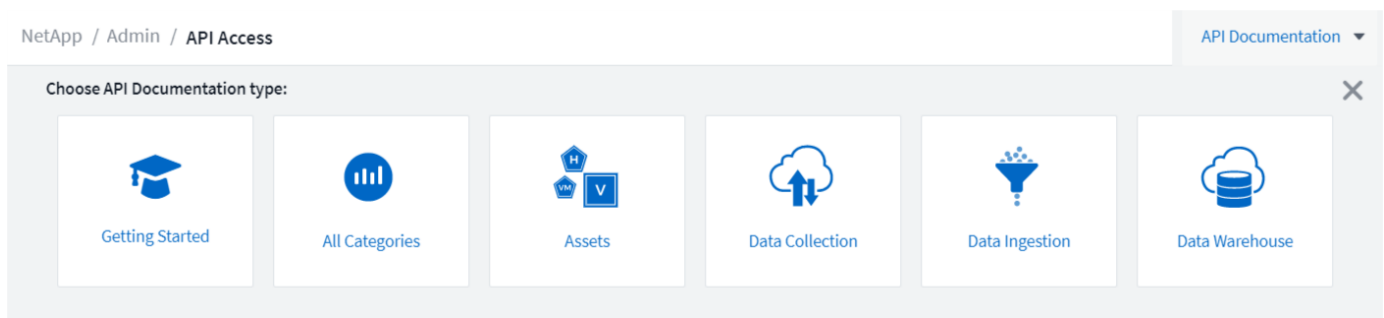
The Cloud Insights API is available in **Cloud Insights Premium Edition**.

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.



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.

POST

/assets/annotations Create annotation definition



Parameters

Try it out

No parameters

Request body

application/json ▼

Request body should include required name, type, optional description and enumValues (if enum type). Enums should contain name and label. Example:

```
{
  "name": "StorageLocation",
  "type": "FIXED_ENUM",
  "description": "Storage Location",
  "enumValues": [
    {
      "name": "PT_LISBON",
      "label": "Lisbon (Portugal)"
    },
    {
      "name": "US_WALTHAM",
      "label": "Waltham (USA)"
    }
  ]
}
```

Example Value | Schema

{}

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 categories, and can grant read and/or write permissions. You can also set the expiration for each access token. All APIs under the specified categories 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 Categories
 - Specify the permissions granted for this API access
 - Specify Token Expiration



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 keys 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 'X-CloudInsights-ApiKey: <API_Access_Token>'
```

Where *<API_Access_Token>* is the token you saved during API access key creation.

API Categories

The Cloud Insights API is category-based, and currently contains the following categories:

- **ASSETS** category 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** category is used to retrieve and manage data collectors.
- **DATA INGESTION** category is used to retrieve and manage ingestion data and custom metrics, such as from Telegraf agents

Additional categories and/or APIs may become available over time. You can find the latest API information in the [API Swagger documentation](#).

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/<category>/<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
```

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:

```

{
  "id": "1247936",
  "self": "/rest/v1/assets/storages/1247936",
  "name": "amsprdclu01",
  "simpleName": "amsprdclu01",
  "naturalKey": "5DF483F0-1729-11DC-9A79-123478563412",
  "ip": "10.64.0.132",
  "serialNumber": "1-80-000011",
  "model": "FAS3270,FAS6290",
  "vendor": "NetApp",
  "microcodeVersion": "8.1.3 clustered Data ONTAP",
  "capacity": {
    "description": "Storage Capacity",
    "unitType": "MB",
    "total": {
      "value": 8.23185105E8
    },
    "storagePools": {
      "value": 5.43220974E8
    }
  },
  "isActive": true,
  "createTime": "2013-05-07T16:52:21-0700",
  "family": "FAS3200,FAS6200",
  "managementUrl": null,
  "virtualizedType": "STANDARD",
  "protocols": [
    "NAS",
    "NFS",
    "CIFS",
    "FC",
    "ISCSI"
  ],
  "_expands": {
    "performance": {
      "url": "/rest/v1/assets/storages/1247936/performance",
      "name": "Performance Data"
    },
    "storageNodes": {
      "url": "/rest/v1/assets/storages/1247936/storageNodes",
      "name": "Storage Storage Nodes"
    },
    "storagePools": {
      "url": "/rest/v1/assets/storages/1247936/storagePools",
      "name": "Storage Storage Pools"
    },
    "storageResources": {
      "url": "/rest/v1/assets/storages/1247936/storageResources",
      "name": "Storage Storage Resources"
    },
    "internalVolumes": {
      "url": "/rest/v1/assets/storages/1247936/internalVolumes",
      "name": "Storage Internal Volumes"
    },
    "volumes": {
      "url": "/rest/v1/assets/storages/1247936/volumes",
      "name": "Storage Volumes"
    },
    "disks": {
      "url": "/rest/v1/assets/storages/1247936/disks",
      "name": "Disks"
    },
    "datasources": {
      "url": "/rest/v1/assets/storages/1247936/datasources",
      "name": "Storage Datasources"
    },
    "ports": {
      "url": "/rest/v1/assets/storages/1247936/ports",
      "name": "Storage Ports"
    },
    "annotations": {
      "url": "/rest/v1/assets/storages/1247936/annotations",
      "name": "Storage Annotations"
    },
    "qtrees": {
      "url": "/rest/v1/assets/storages/1247936/qtrees",
      "name": "Qtrees"
    }
  },
  "-----"

```

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,storageRe  
sources.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 category 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).

<https://tenant.cloudinsights.netapp.com/rest/v1/assets/storages/1/performance?expand=history>

Details

Response body

```
{
  "self": "/rest/v1/assets/storages/1/performance",
  "cacheHitRatio": {
    "read": {
      "description": "Cache Hit Ratio - Read",
      "unitType": "%",
      "start": null,
      "end": null,
      "current": null,
      "min": null,
      "max": null,
      "avg": null,
      "sum": null,
      "isDownsampled": false
    },
    "write": {
      "description": "Cache Hit Ratio - Write",
      "unitType": "%",
      "start": null,
      "end": null,
      "current": null,
      "min": null,
      "max": null,
      "avg": null,
      "sum": null,
      "isDownsampled": false
    }
  }
}
```

Self

Performance Metric

Response body

```
}
},
"history": [
  [
    1578418848140,
    {
      "latency.total": 1.30578,
      "latency.read": 3.64681,
      "ioDensity.read": 9.62065,
      "iops.write": 686.35502,
      "ioDensity.total": 31.36259,
      "capacity.raw": 80024.92772,
      "throughput.read": 7.32371,
      "iops.total": 1488.7974,
      "latency.write": 0.39495,
      "ioDensity.write": 14.45856,
      "iops.read": 456.69703,
      "capacity.storagePools": 56058.1041,
      "throughput.write": 14.59581,
      "throughput.total": 21.91953
    }
  ],
  [
    1578419748198,
    {

```

History

Timestamp

Counter Values

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 categories 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.
- Switch Performance Metrics, such as Traffic Utilization, BB Credit Zero data, or Port Errors.

See the [API Swagger documentation](#) for information on metrics for each object type.

Performance History Data

History data is presented in performance data as a list of timestamp and counter maps pairs.

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”.



Timestamp is in UNIX time format.

The following is an example of a performance data JSON for a disk:

```

"performance": {
  "self": "/rest/v1/assets/disks/4013931/performance",
  "iops": {
    "performanceCategory": "IOPS",
    "description": "Disk IOPS",
    "read": {
      "description": "Disk Read Iops",
      "unitType": "IO/s",
      "start": 1399305599999,
      "end": 1402604368055,
      "current": 1,
      "min": 0,
      "max": 6,
      "avg": 0.5532
    },
    [...]
  },
  "total": {
    "description": "Disk Total Throughput",
    "unitType": "MB/s",
    "start": 1399305599999,
    "end": 1402604368055,
    "current": 0,
    "min": 0,
    "max": 2,
    "avg": 0.1702
  }
},
"history":
[
  [
    1399300412690,
    {
      "utilization.total": 12,
      "iops.total": 26,
      "iops.write": 22,
      "iops.read": 4,
      "throughput.read": 0,
      "utilization.read": 2.12,
      "throughput.total": 5,
      "utilization.write": 10.24,
      "throughput.write": 5
    }
  ]
]

```

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.

The following shows the capacity for `StoragePools` with multiple capacity counters:

StoragePoolCapacity

```
Model properties:
{
  description: string
  unitType: 'MB' or 'GB' or 'TB' or 'KiB' or 'MiB' or 'TiB'
  total: CapacityItem
  used: CapacityItem
  provisioned: CapacityItem
  reservedCapacity: CapacityItem
  softLimit: Double
  rawToUsableRatio: Double
  isDedupeEnabled: boolean
  dedupeSavings: NumericValueWithUnit
  isCompressionEnabled: boolean
  compressionSavings: NumericValueWithUnit
  isThinProvisioningSupported: boolean
}
```

close

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. Categories are different asset types from the Inventory, such as storages, hosts, dataStores, and so on. Each category would contain a list of objects of the category 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.

Performance Policies and Alerts

Alerting with Monitors

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.



Monitors and Alerting is available in Cloud Insights Standard Edition and higher.

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.


Monitor or Performance Policy?

What's the difference between a **Performance Policy** and a **Monitor**?

Policies allow you to set thresholds on "infrastructure" objects such as storage, VM, EC2, and ports. These policies trigger violations when thresholds are met or exceeded. Each violation can be investigated for troubleshooting. Policies are described in detail elsewhere in this [documentation](#).

Monitors provide similar functionality for "integration" data such as those collected for Kubernetes, ONTAP advanced metrics, and Telegraf plugins, and alert when thresholds are crossed. With Monitors, you can set thresholds for Warning- or Critical-level alerts, or both.

Policies and Monitors are available under the **Alerts** menu.



Emails can be sent when a policy or monitor is triggered.

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

1. From the Cloud Insights menu, click **Alerts > Manage Monitors**

The Monitors list page is displayed, showing currently configured monitors.

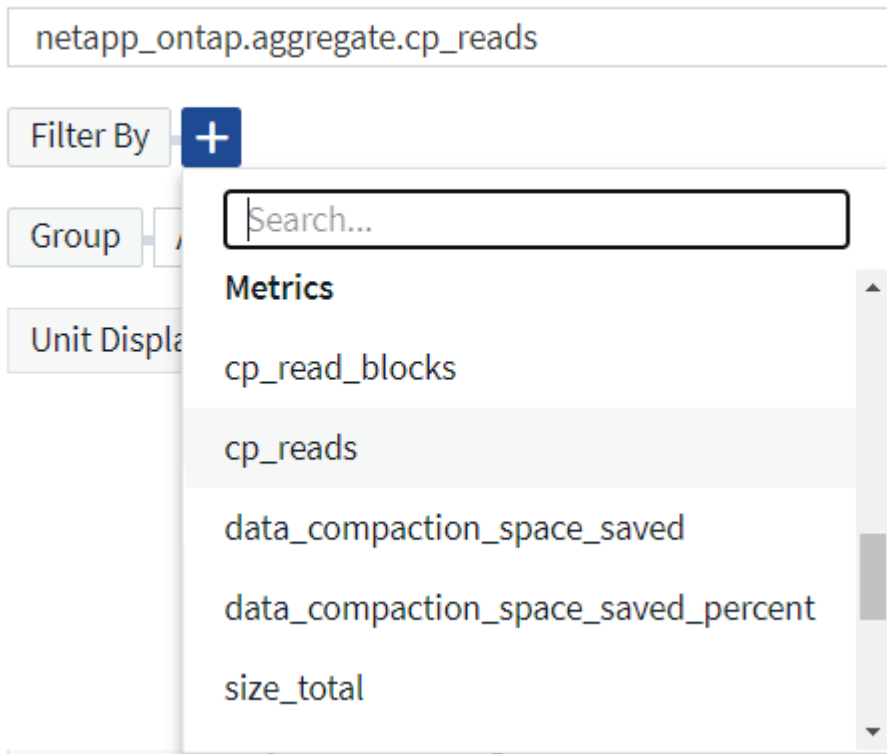
2. To add a monitor, Click **+ Monitor**. To modify an existing monitor, click the monitor name in the list.

The Monitor Configuration dialog is displayed.

3. 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.

1 Select a metric to monitor



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.



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.


1. After choosing the object and metric to monitor, set the Warning-level and/or Critical-level thresholds.
2. For the *Warning* level, enter 200. The dashed line indicating this Warning level displays in the example graph.
3. 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.

4. For the occurrence interval, choose *Continuously* for a period of *15 Minutes*.

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.


2 Define the monitor's conditions (set at least one threshold condition)



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.

3 Set up team notification(s) (alert your team via email, or Webhook)



Alerting via Email:

Specify the email recipients for alert notifications. If desired, you can choose different recipients for warning or critical alerts.

3 Set up team notification(s)

Email

Notify team on

Critical, Resolved

☒ Critical

☐ Warning

☒ Resolved

Add Recipients (Required)

user_1@email.com X

user_2@email.com X

Email

Notify team on

Warning

Add Recipients (Required)

user_3@email.com X

Alerting via Webhook:

Specify the webhook(s) for alert notifications. If desired, you can choose different webhooks for warning or critical alerts.

3 Set up team notification(s) (alert your team via email, or Webhook)

By Webhook

Slack

Notify team on

Critical

Use Webhook(s)

Slack X Teams X

Notify team on

Resolved

Use Webhook(s)

Slack X Teams X

Notify team on

Warning

Use Webhook(s)

Slack X Teams X



Webhooks is considered a Preview feature and is therefore subject to change.

Save your Monitor

1. If desired, you can add a description of the monitor.
2. 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 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.



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.

5 Associate to a monitor group (optional)

ONTAP Monitors

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.



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*.



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 *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*.

Monitor Groups (3)



Agent Monitors

All Monitors (4)

Custom Monitors (4)

Agent Monitors (3)



Pause

Resume

Rename

Delete

More Information

- [Viewing and Dismissing Alerts](#)

Viewing and Managing Alerts from Monitors


Cloud Insights displays alerts when [monitored thresholds](#) are exceeded.



Monitors and Alerting is available in Cloud Insights Standard Edition and higher.

Viewing and Managing Alerts

To view and manage alerts, do the following.

1. Navigate to the **Alerts > Integration Alerts** page.
2. 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  :
 - **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

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.

Alert Summary

Monitor:

Volume Total Data

Triggered On:

cluster_name: tawny
aggr_name: Multiple_Values

Duration / Time Triggered:

1d 6h / Jun 9, 2020 2:22 AM

Top Severity:

❗ Critical

Metric:

① netapp_ontap.workload_volume.total_data

Condition:

Average total_data is > (greater than) 0m and/or 0m all the time in 2-hour window.

Filters Applied:


cluster_name: Any

Status:

New

Expert View

Display Metrics ▾



Related Alerts

1 item found

Alert ID	Active Status	Triggered Time ↓	Top Severity	Monitor	Triggered On	Status
AL-46769	Resolved	a day ago Jun 9, 2020 2:22 AM	❗ Critical	Volume Total Data	cluster_name: tawny aggr_name: Multiple_Values	New

Comments

There are no comments yet on this alert.

[+ Comment](#)

Configuring Email Notifications

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.

To configure notification email recipient settings, go to the **Admin > Notifications** page.

Subscription Notification Recipients

Subscription Notification Recipients

Send subscription related notifications to the following:

☒ All Account Owners

☒ All Administrators

☒ Additional Email Addresses

Enter email addresses separated by commas.

Save

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 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 Subscription page
Subscription will expire in 90 days Subscription will expire in 30 days	No action needed if "Auto Renewal" is enabled Contact NetApp sales to renew the subscription
Trial ends in 2 days	Renew trial from the Subscription page. You can renew a trial one time. Contact 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 NetApp sales to purchase a subscription

Global Recipient List for Performance Policy Notifications

Global Performance Policy Recipients

Default email recipients for Performance Policy related notifications:

Recipients

Enter email addresses separated by commas.

Email Signature

Email signature added to messages sent by Cloud Insights

Save

To add recipients to the global performance policy notification email list, go to the "Global Performance Policy Recipients" section and enter email addresses separated by commas. Emails sent as alerts from performance policy threshold violations will be sent to all recipients on the list.

If you make a mistake, you can click on [x] to remove a recipient from the list.

You can also add an optional signature block that will be attached to the email notifications sent.



You can override the global list for a specific policy when you configure that policy.

Notification using Webhooks

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 [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 <i>http://</i> or <i>https://</i> prefix
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:

Edit a Webhook

Name

Slack Test

Template Type

Slack

URL

https://hooks.slack.com/services/T12345678/B01234567H/B01234567kT89mFrtvpyA

Method

POST

Custom Header

Content-Type: application/json
Accept: application/json

Message Body

```
{
  "blocks": [
    {
      "type": "section",
      "text": {
        "type": "mrkdwn",
        "text": "**Cloud Insights Alert - %alertId%*  
Severity - %severity%*"
      }
    }
  ],
}
```

Cancel

Test Webhook

Save Webhook

Cloud Insights webhooks comprise a number of default parameters. Additionally, you can create your own custom parameters or secrets.


Default Alert Parameters

Name	Description
%%alertDescription%%	Alert description
%%alertId%%	Alert ID
%%alertRelativeUrl%%	Relative URL to the Alert page. To build alert link use <code>https://%%cloudInsightsHostName%%%%alertRelativeUrl%%</code>
%%metricName%%	Monitored metric
%%monitorName%%	Monitor name
%%objectType%%	Monitored object type
%%severity%%	Alert severity level
%%alertCondition%%	Alert condition
%%triggerTime%%	Alert trigger time in GMT ("Tue, 27 Oct 2020 01:20:30 GMT")
%%triggerTimeEpoch%%	Alert trigger time in Epoch format (milliseconds)
%%triggeredOn%%	Triggered On (key:value pairs separated by commas)
%%value%%	Metric value that triggered the alert
%%cloudInsightsLogoUrl%%	Cloud Insights logo URL
%%cloudInsightsHostname%%	Cloud Insights Hostname (concatenate with relative URL to build alert link)

Custom Parameters and Secrets

Name	Value	Description
------	-------	-------------

No Data Available

 Parameter

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.

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.

Choosing Webhook Notification in a Monitor

To choose the webhook notification in a [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.

3 Set up team notification(s) (alert your team via email, or Webhook)

By Webhook

Notify team on

Critical, Warning, Resolved

Use Webhook

Please Select

Search...

ci-alerts-notifications-dev

ci-alerts-notifications-qa

Webhook Examples:

Webhooks for [Slack](#)

Webhooks for [PagerDuty](#)

Webhooks for [Teams](#)

Webhooks for [Discord](#)

Monitoring your Environment

Auditing

To identify unexpected changes, you can view an audit trail of the Cloud Insights system and its user activities.

Cloud Insights generates audit entries for changes that impact the system or users. These changes include:

- Logging in
- Authorizing or unauthorizing a path
- Updating an authorized path
- Setting global policies or thresholds
- Adding or removing a data collector
- Starting or stopping a data collector
- Updating data collector properties
- Adding, editing, or deleting a task
- Removing an application group
- Identifying or changing the identification for a device

Use the following steps to access the Audit system:

Steps

1. In the Cloud Insights menu, click **Admin > Audit**

The Audit page is displayed, providing the following details for each audit entry:

- **Time** - Date and time that the changes were made
- **User** - User account's role, (guest, user, or administrator)
- **IP** - IP address associated with the audit entry
- **Action** - Type of activity in the audit entry
- **Details** - Details of the audit entry

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, 3 hours, 24 hours, 3 days, and 7 days), with Cloud Insights showing a maximum number of 1000 violations for the selected time period.
- You 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.

By default, the table displays the entries in descending time order.

- You can use the filter box to show only the entries you want in the table.

Asset Page Information

Asset Page Overview

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

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 1 Hour
- Last 3 Hours (this is the default)
- Last 24 Hours
- 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.

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.

Filtering for Objects In-Context


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.

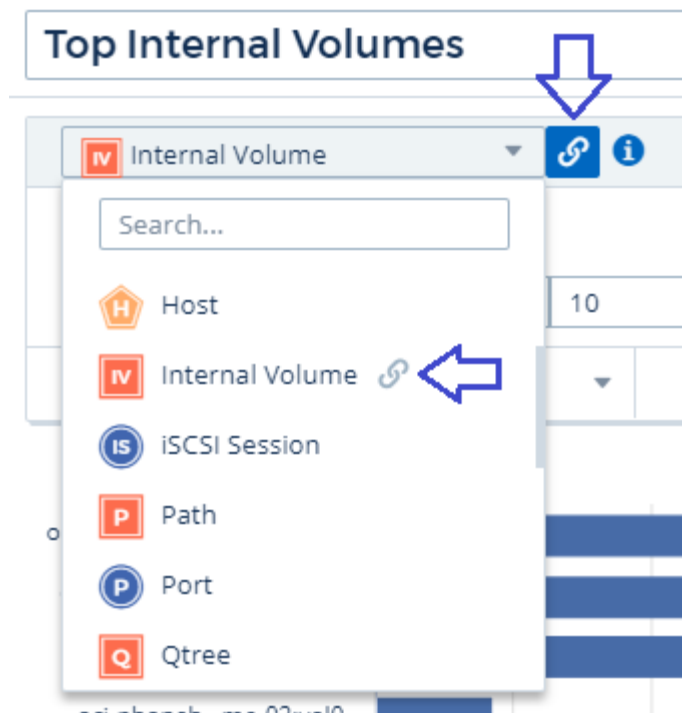
On most asset landing pages, widgets allow you to filter for objects related to the current asset. In filter drop-downs, object types that display a link icon  can be filtered in-context to the current asset.

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



1. Open an asset page for any **Storage** asset.
2. Click **Edit** to open the asset page in Edit mode.
3. Click **Add Widget** and select *Bar Chart*.
4. Select **Internal Volume** for the object type to display on the bar chart. Notice that the internal volume object type has a link icon  beside it. The "linked" icon is enabled by default.



5. Choose *IOPS - Total* and set any additional filters you like.
6. Collapse the **Roll Up** field by clicking the [X] beside it. The **Show** field is displayed.
7. Choose to show Top 10.
8. 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  next to the object type. The link becomes disabled  and the chart displays data for *all* objects in your environment.

You can also use [special variables in widgets](#) to display asset-related information on landing pages.

Asset Page Summary section

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 next to the metric or performance policy.

Storage

Model:
FAS3070

Vendor:
NetApp

Family:
FAS3000

Serial Number:
1082247,30012349

IP:
10.197.138.10,10.197.138.11

Microcode Version:
8.1.4P10 7-Mode

Raw Capacity:
83,211.0 GB

Latency - Total:
0.02 ms

IOPS - Total:

287.35 IO/s

135 'Storage' violations with 'IOPS - Total' > 2.00 IO/s

0.02 MB/s

Management:

FC Fabrics Connected:
0

Performance Policies:

Storage

Note: The information displayed in the Summary section varies, depending on the type of asset you are viewing.

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.



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 a performance policy link to view the performance policy or policies associated with the asset.

If a red circle appears next to a performance policy, this indicates an asset has crossed the performance policy's defined threshold. You can examine the performance policy to further diagnose the issue.

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

Internal Volume

Storage: [barbados1,barbados2](#)

Storage Pool: [barbados1.aggr1](#)

Status: Online

Type: FlexVol

UUID:

SVM/vFiler: vfiler0

Capacity - Total: 1.0 GB

Capacity - Used: 0.0 GB

Snapshot: <0.1 GB

Latency - Total: 0.02 ms

Storage Pool Utilization: 0.68 %

IOPS - Total: 0.13 IO/s

Datastore:

Deduplication Savings: 0.0 %

Thin Provisioned: No

Replication Source(s):

Performance Policies: [Find High Latency FlexVols](#)

[View Topology](#)

Topology

ocise-esx-1431... → NAS → barbados1,bar...


[Close](#)

Expert View

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.

Using the Expert View section

The following is an example of the Expert View section in a storage asset page:



You can select the metrics you want to view in the performance chart for the time period selected.

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.
- Greedy

Shows the assets that take away system resources from the asset through sharing the same resources, such as hosts, networks, and storage.
- Degraded

Shows the assets that are depleted of system resources due to this asset.

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-to-buffer 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 result in cache hits
Class 3 discards	Count of Fibre Channel Class 3 data transport discards.
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)
Frame size average (Rx, Tx)	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.
Link reset Rx, Tx	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.
Memory utilization (Total) %	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).
Timeout discard frames - Tx	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).
Write pending (Total)	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

1. 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.

2. 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.

3. 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.

4. 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.

5. 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 .

User Data Section

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

1. Locate an asset page by doing any of the following:
 - a. Query for a list of assets and then select one from the list.
 - b. On a Dashboard, locate an asset name and click it.
 - c. 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.

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

1. To add an annotation to the asset, in the User Data section of the asset page, click **+Annotation**.
2. Select an annotation from the list.
3. Click Value and do either of the following, depending on type of annotation you selected:
 - a. If the annotation type is list, date, or Boolean, select a value from the list.
 - b. If the annotation type is text, type a value.
4. 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.

Asset Page Violations section

You can use the Violations section of an asset page to see the violations, if any, that occur in your environment as a result of a performance policy assigned to an asset. Performance policies monitor thresholds and enable you to detect a violation of a threshold immediately, 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 Violations section that displays on an asset page for a hypervisor:

Violations		
243 items found		<input type="text" value="Filter..."/>
Time ↓	Description	
08/27/2018 8:59:27 PM	 10.197.144.53 violation with 'Memory Utilization - Total' > 5.00 % (value of 70.19 %)	⋮
08/27/2018 8:59:27 PM	 10.197.144.53 violation with 'CPU Utilization - Total' > 5.00 % (value of 27.59 %)	⋮
08/27/2018 1:45:04 PM	 10.197.144.53 violation with 'CPU Utilization - Total' > 5.00 % (value of 25.89 %)	⋮
08/27/2018 1:45:04 PM	 10.197.144.53 violation with 'Memory Utilization - Total' > 5.00 % (value of 68.61 %)	⋮
08/13/2018 12:55:04 PM	 10.197.144.53 violation with 'Memory Utilization - Total' > 5.00 % (value of 50.17 %)	⋮
08/13/2018 12:55:04 PM	 10.197.144.53 violation with 'CPU Utilization - Total' > 5.00 % (value of 16.66 %)	⋮
08/13/2018 11:30:44 AM	 10.197.144.53 violation with 'Memory Utilization - Total' > 5.00 % (value of 68.05 %)	⋮

The Violations section enables you to view and manage any of the violations that occur in your network as the result of a performance policy 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 one from the results list.

The asset page displays. The Violations section displays the time the violation occurred and a description of

the threshold that was crossed, along with a hyperlink to the asset on which the violation occurred (for example **2 violations for ds-30 with Latency - Total > 50**).

You can perform any of the following optional tasks:

- Use the filter box to show only specific violations.
- 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.

Hints and Tips to Search for Assets

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 used.

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.

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	"internalvolume" 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

Category	Indexed terms
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
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
Switches (regular and NPV)	"switch" IP address wwn name serial number model domain ID name of the fabric wwn of the fabric

Category	Indexed terms
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

Reporting

Cloud Insights Reporting Overview

Cloud Insights reporting is a business intelligence tool that enables you to view pre-defined reports or create custom reports.



The Reporting feature is available in Cloud Insights [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.

You will be taken to the Reporting interface. Cloud Insights uses IBM Cognos Analytics for its reporting engine. Log in using your Reporting credentials.

Cloud Insights 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). 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.

Feature	Consumer	Author	Administrator
View reports in the Team Content tab	Yes	Yes	Yes
Run reports	Yes	Yes	Yes
Schedule reports	Yes	Yes	Yes
Upload external files	No	Yes	Yes
Create Jobs	No	Yes	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

Setting Reporting (Cognos) email preferences



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:

1. Create a new folder in *Team Content*
 - a. If multiple users have been created, please create a separate folder for each user to avoid overwriting reports with duplicate names
2. Navigate to *My Content*
3. Select all of the reports you wish to retain.
4. In the upper right corner of the menu, select "Copy or move"
5. Navigate to the newly created folder in *Team Content*
6. Paste the reports to the newly created folder using the "Copy to" or "Move to" buttons
7. Once SSO is enabled for Cognos, log into Cloud Insights with the email address used to create your account.
8. Navigate to the *Team Content* folder within Cognos, and Copy or Move the previously saved reports back to *My Content*.

Predefined Reports Made Easy

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.



The Reporting feature is available in Cloud Insights [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.

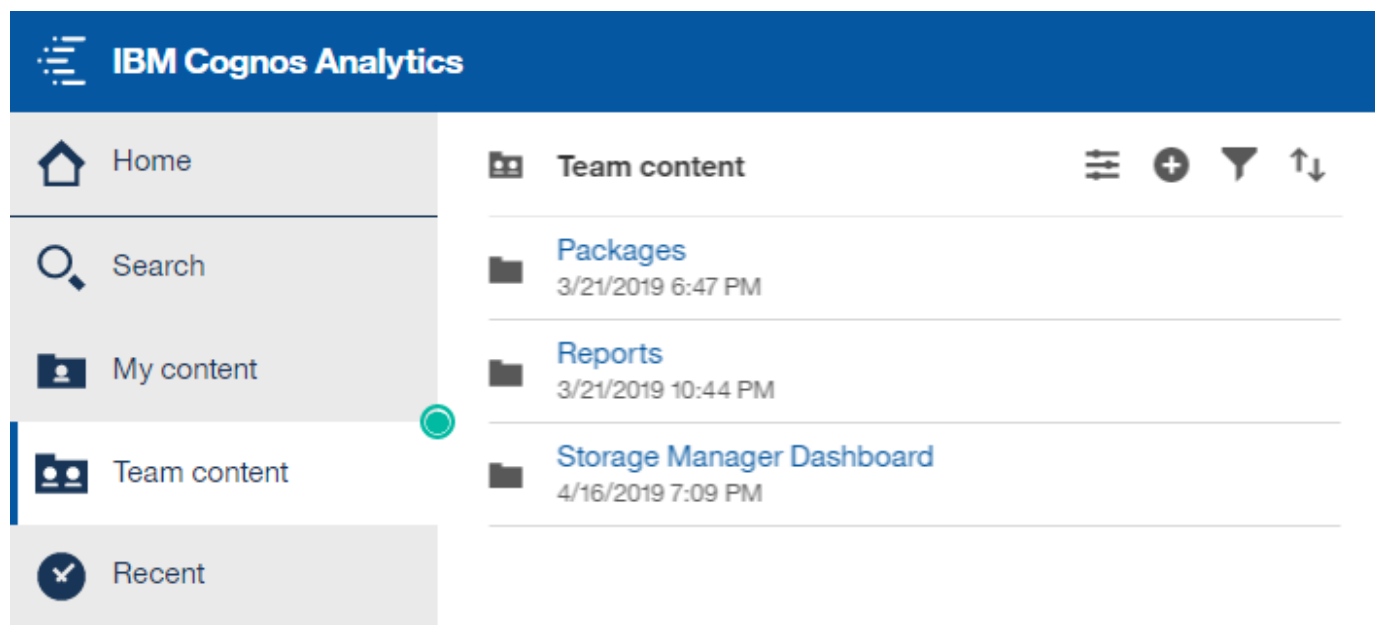
Cloud Insights accommodates multiple tenancy in reporting by enabling you to associate users with business units. With this feature, administrators can separate data or reports according to the attributes of a user or his/her affiliation.

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.

1. In the left navigation pane, click **Team Content** and select the information category that you want to use.

2. Click **Reports** to access the pre-defined reports.



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.

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.

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.

Storage Manager Dashboard

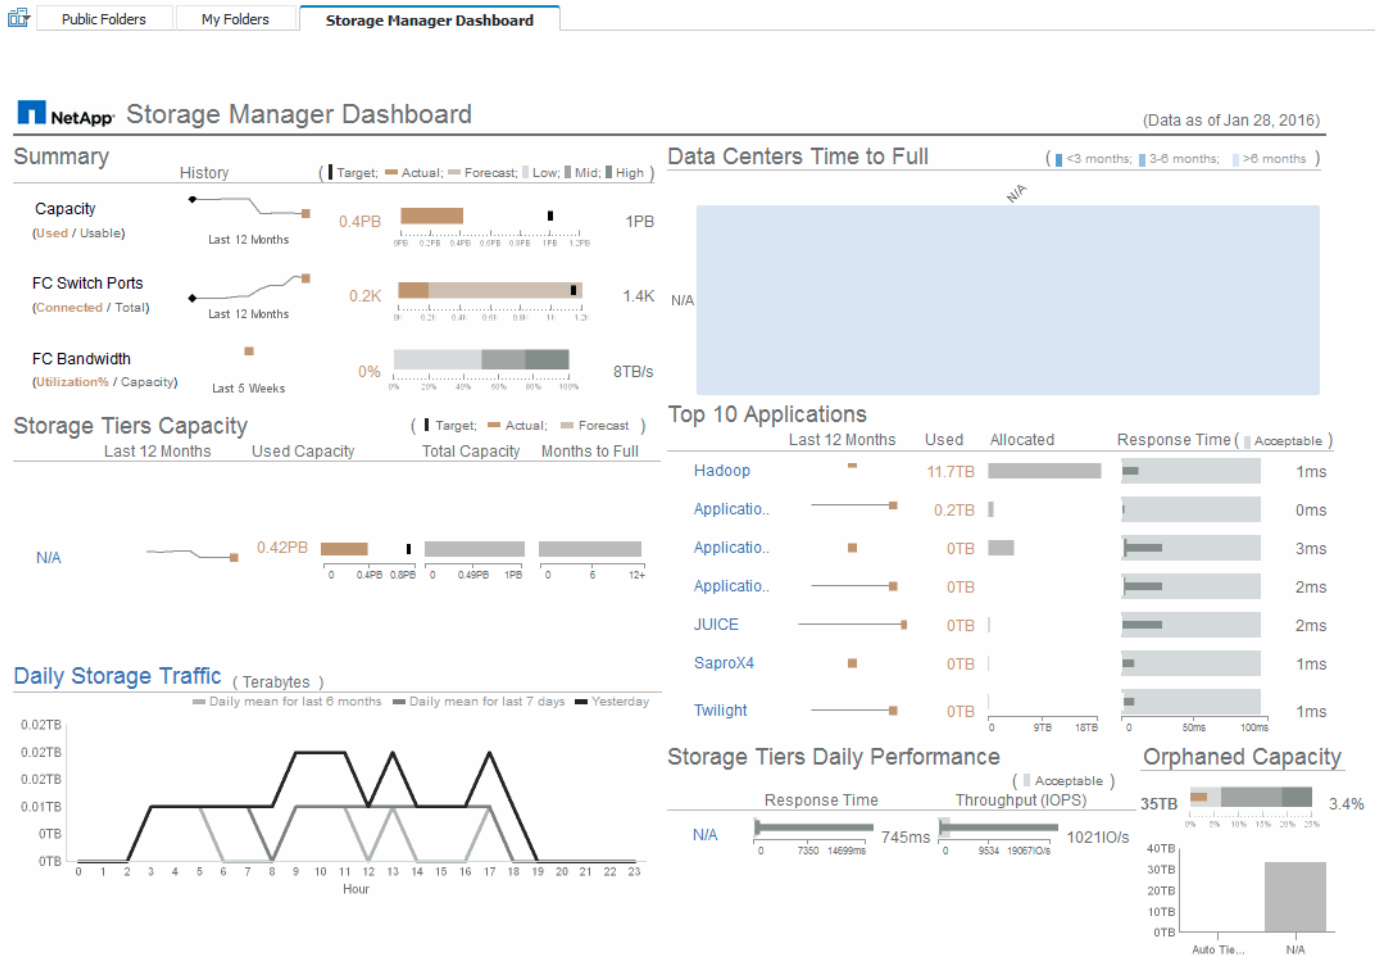
The Storage Manager Dashboard provides you with a centralized visualization that enables you to compare and contrast resource usage over time against the acceptable ranges and previous days of activity. Showing only the key performance metrics for your storage services, you can make decisions about how to maintain your data centers.



The Reporting feature is available in Cloud Insights [Premium Edition](#).

The dashboard comprises seven components that contain contextual information on certain 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.

Summary



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.

Creating a Report (Example)

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

1. In the toolbar, click **[+]**
2. Click **Report**
3. Click **Templates > Blank**
4. Click **Themes > Cool Blue > OK**

The Source and Data tabs is displayed

5. Click **Source > [+]**
6. In the **Open file** dialog, click **Team content > Packages**

A list of available packages is displayed.


7. Click **Storage and Storage Pool Capacity > Open**
8. Click **[+]**

The available styles for your report are displayed.

9. Click **List**












Add appropriate names for List and Query

10. Click **OK**
11. Expand *Physical Capacity*
12. Expand to the lowest level of *Data Center*
13. Drag *Data Center* to the Reporting palate.
14. Expand *Capacity (MB)*

15. Drag *Capacity (MB)* to the Reporting palate.
16. Drag *Used Capacity (MB)* to the Reporting palate.
17. Run the report by clicking  and selecting an output type.

Result

A report similar to the following is created:

	Data Center	Capacity (MB)	Used Capacity (MB)
	Asia	122,070,096.00	45,708,105.00
	BLR	100,709,506.00	54,982,204.00
	Boulder	22,883,450.00	12,011,075.00
	DC01	1,707,024,715.00	1,407,609,686.00
	DC02	732,370,688.00	732,370,688.00
	DC03	314,598,162.00	65,448,975.00
	DC04	573,573,884.00	282,645,615.00
	DC05	89,245,458.00	62,145,011.00
	DC06	19,455,433,799.00	11,283,487,744.00
	DC08	100,709,506.00	44,950,171.00
	DC10	112,916,718.00	43,346,818.00
	DC14	23,565,735,054.00	17,357,431,924.00
	DC56	137,549,084.00	10,657,793.00
	Europe	743,942,208.00	240,369,325.00
	HIO	9,823,036,853.00	4,216,750,338.00
	London	0.00	0.00
	N/A	9,049,939,023.00	5,887,911,992.00
	RTP	12,386,326,262.00	5,638,948,477.00
	SAC	9,269,642,330.00	6,197,549,437.00
	 Top  Page up  Page down  Bottom		

Managing Reports

You can customize a report's output format and delivery, set report properties or schedules, and email reports.

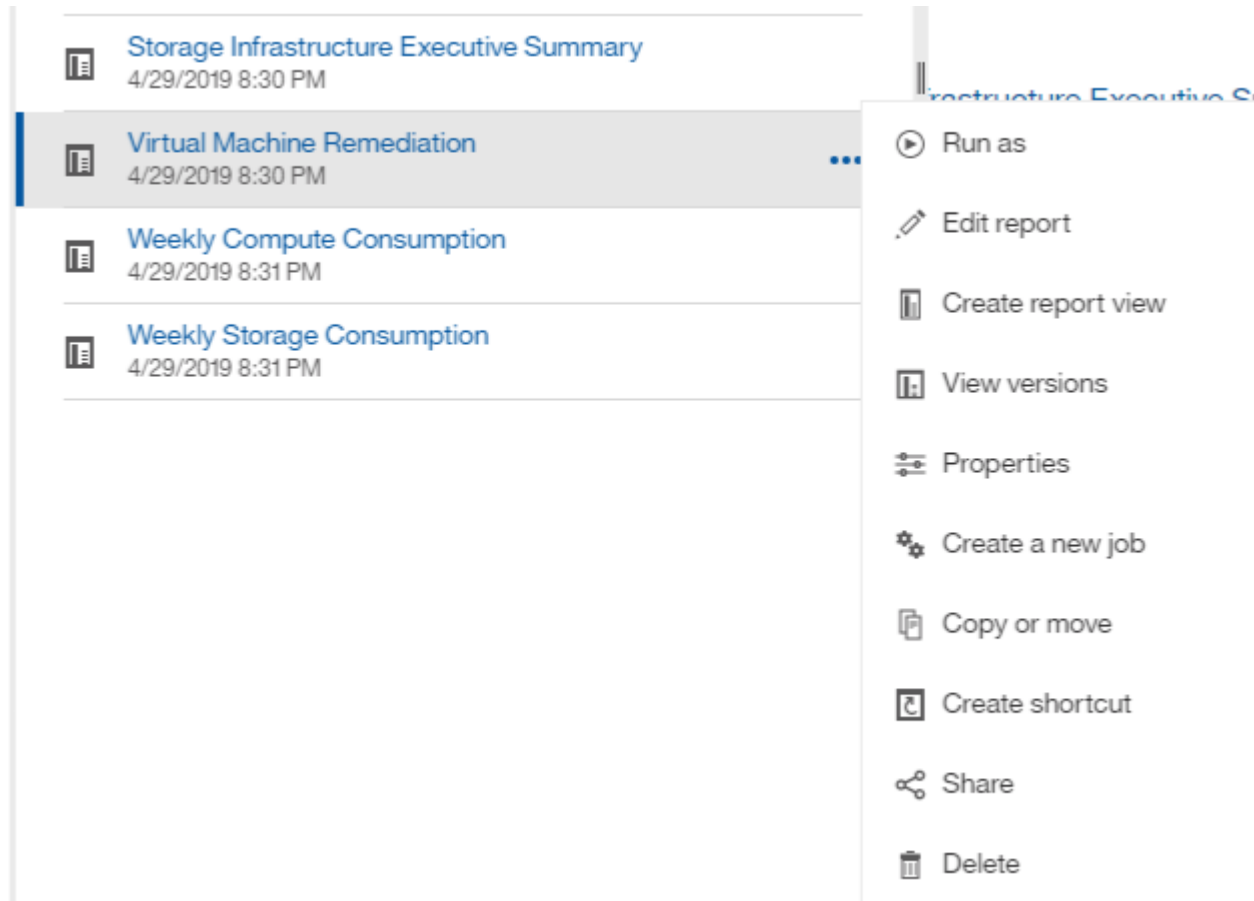


The Reporting feature is available in Cloud Insights [Premium Edition](#).

Customizing a report's output format and delivery

You can customize the format and delivery method of reports.

1. In the Cloud Insights Reporting Portal, mouse over the report you want to customize and click [...].






1. Click **Properties > Schedule**

[< Back](#)

Create schedule

Period

Start	 2018-04-06	 1:49 PM
End	 2018-07-06	 1:49 PM
<input type="checkbox"/> No end date		

Run every week(s)

On day(s)

M	T	W	T	F	S	S
---	---	---	---	----------	---	---

☐ Daily time interval

Options

Format  HTML >

Delivery  Save >



Prompts Set values >

Languages English (United States) >

1. You can set the following options:
 - Schedule when you want reports to run.
 - Format the report output.
 - Delivery: print, save, or email the report.
 - Languages define the language the report is delivered in.
2. Click **Create** 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.

1. In the toolbar, click **[+]**
2. Click **Report**
3. Click the **Pages** icon . The Report icon  **Report** is displayed



4. Left click the Report icon. Report options are displayed.
5. Click **Copy Report to Clipboard**.

Opening reports (xml) from the clipboard

You can open a report specification that was previously copied to the clipboard.

About this task

You enter the Reporting user interface by creating a new report or opening an existing report

1. In the toolbar, click **[+]**
2. Click **Report**
3. Click the **Pages** icon . The Report icon  **Report** is displayed
4. Left click the Report icon. Report options are displayed.
5. Click **Open report from clipboard**.

Creating Custom Reports

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.



The Reporting feature is available in Cloud Insights [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 [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.

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 Cloud 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 model

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 qtrees, 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. 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.

How historical data is retained for Reporting

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	35 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 next ETL)

Cloud Insights Data Warehouse Schema Diagrams

This document provides the schema diagrams for the Data Warehouse Database. You can also download a file containing the [schema tables](#).

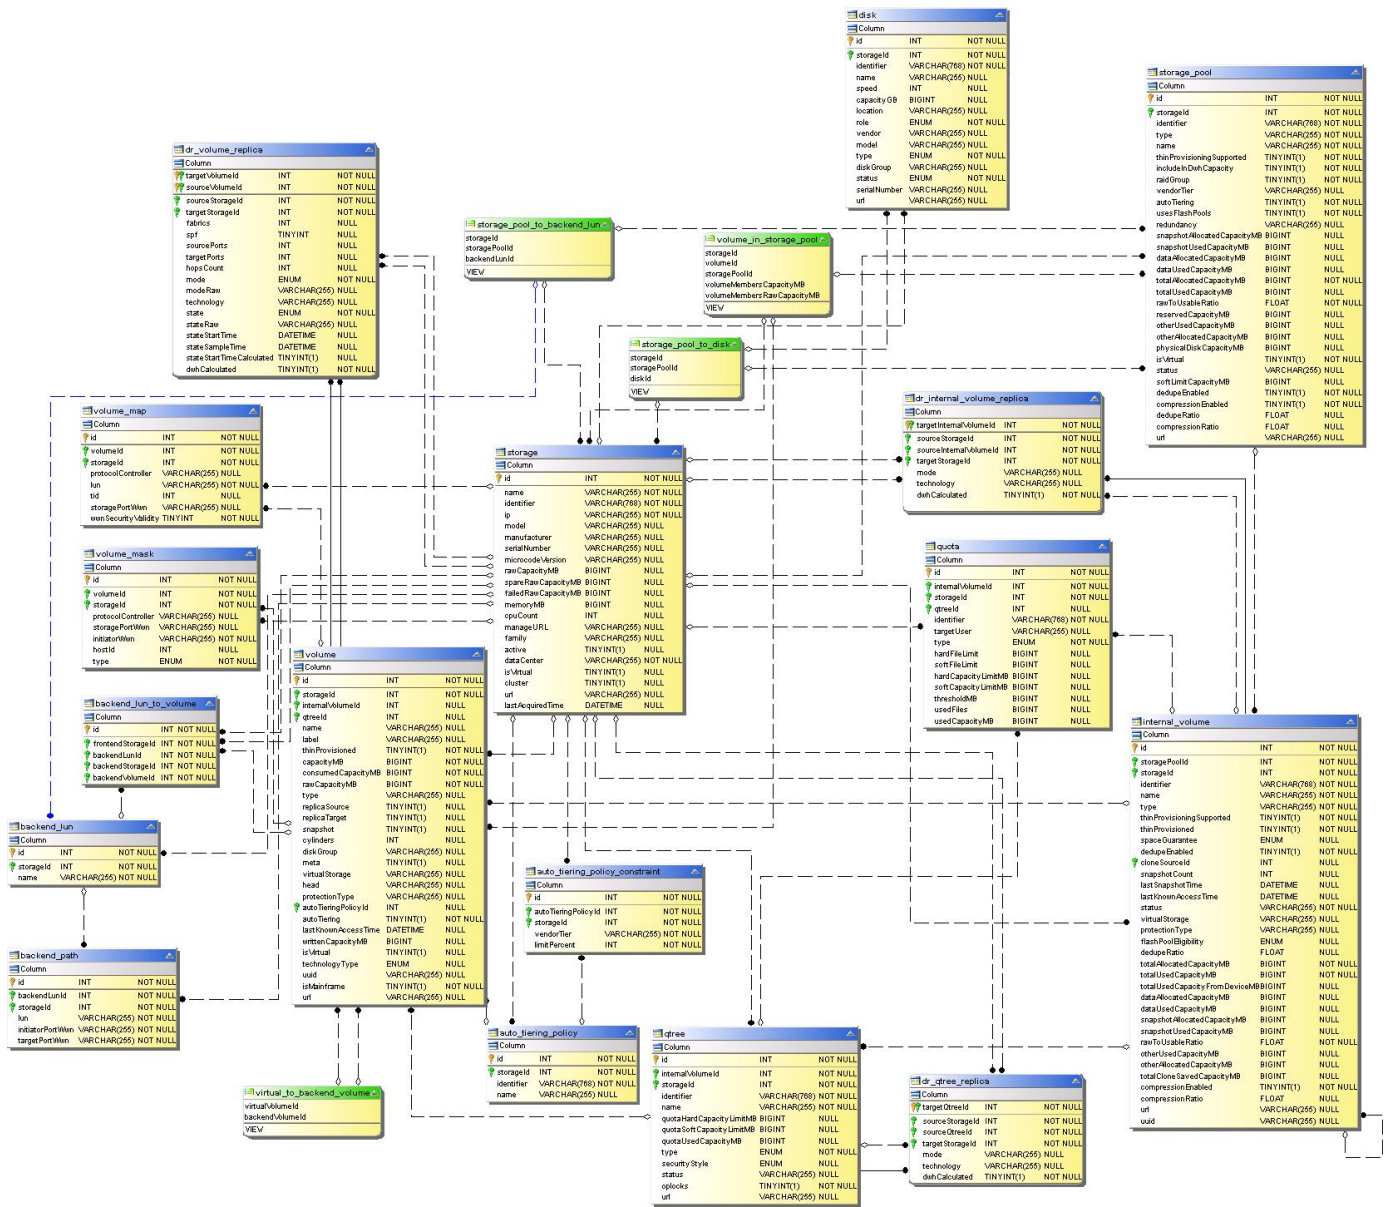


The Reporting feature is available in Cloud Insights [Premium Edition](#).

Inventory Datamart

The following images describe the inventory datamart.

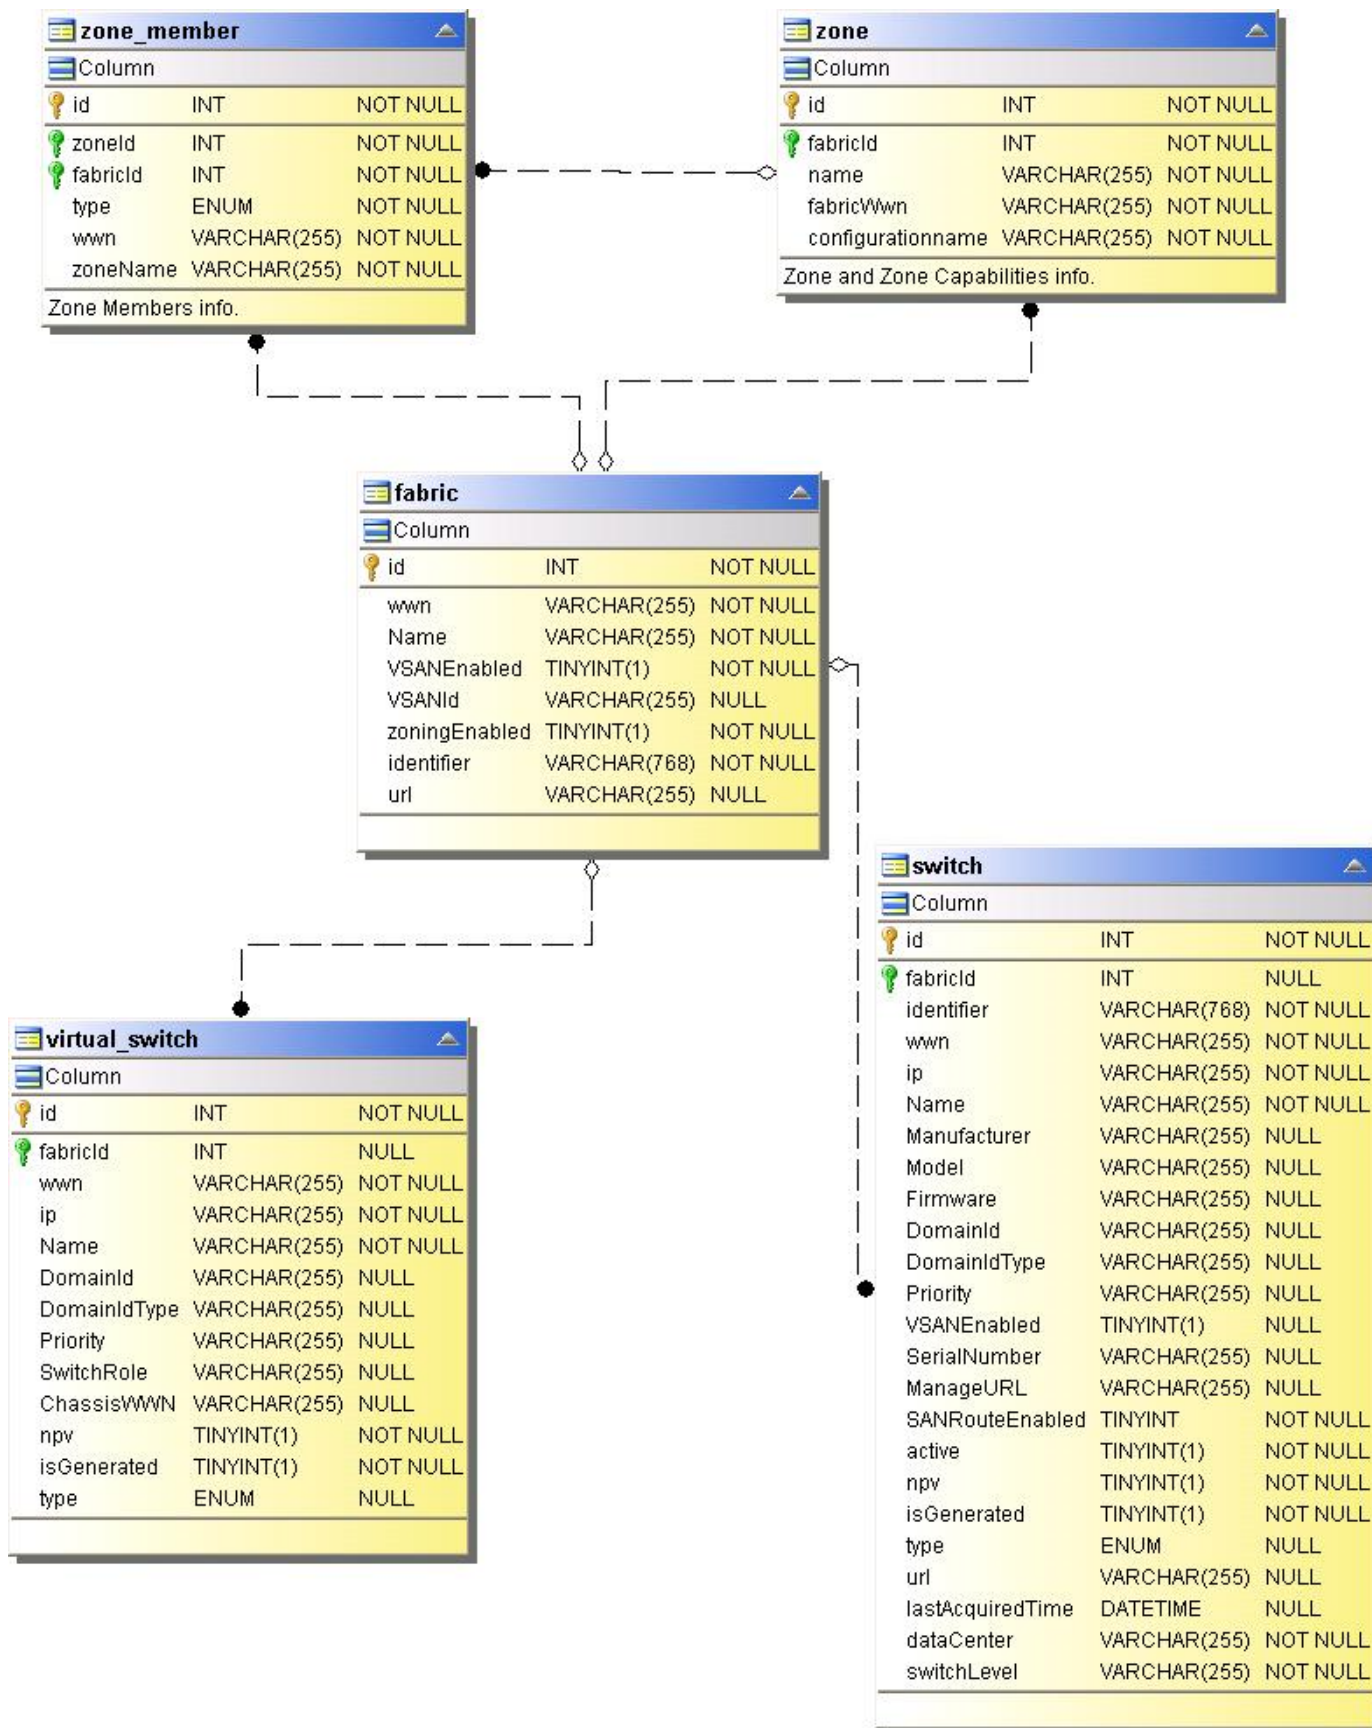
Storage



Storage Node



SAN Fabric



Port Connectivity

host	
Column	
id	INT
name	VARCHAR(255)
identifier	VARCHAR(768)
ip	VARCHAR(255)
os	VARCHAR(255)
model	VARCHAR(255)
manufacturer	VARCHAR(255)
installedMemoryMB	VARCHAR(255)
hostFsFreeGB	VARCHAR(255)
hostFsTotalGB	VARCHAR(255)
hostFsUsedGB	VARCHAR(255)
cpuCount	VARCHAR(255)
cpuSpeed	VARCHAR(255)
nicCount	VARCHAR(255)
nicSpeed	VARCHAR(255)
url	VARCHAR(255)
active	TINYINT(1)
dataCenter	VARCHAR(255)

host_adapter	
Column	
id	INT
hostid	INT
wwn	VARCHAR(255)
model	VARCHAR(255)
manufacturer	VARCHAR(255)
driver	VARCHAR(255)
firmware	VARCHAR(255)

switch	
Column	
id	INT
fabricid	INT
identifier	VARCHAR(768)
wwn	VARCHAR(255)
ip	VARCHAR(255)
name	VARCHAR(255)
manufacturer	VARCHAR(255)
model	VARCHAR(255)
firmware	VARCHAR(255)
domainid	VARCHAR(255)
domainidType	VARCHAR(255)
priority	VARCHAR(255)
vsanEnabled	TINYINT(1)
serialNumber	VARCHAR(255)
manageURL	VARCHAR(255)
sanRouteEnabled	TINYINT(1)
npv	TINYINT(1)
type	ENUM
url	VARCHAR(255)
lastAcquiredTime	DATETIME
active	TINYINT(1)
dataCenter	VARCHAR(255)
switchLevel	VARCHAR(255)
isGenerated	TINYINT(1)

storage	
Column	
id	INT
name	VARCHAR(255)
identifier	VARCHAR(768)
ip	VARCHAR(255)
model	VARCHAR(255)
manufacturer	VARCHAR(255)
serialNumber	VARCHAR(255)
microcodeVersion	VARCHAR(255)
rawCapacityMB	BIGINT
spareRawCapacityMB	BIGINT
failedRawCapacityMB	BIGINT
memoryMB	BIGINT
cpuCount	INT
manageURL	VARCHAR(255)
family	VARCHAR(255)
cluster	TINYINT(1)
url	VARCHAR(255)
lastAcquiredTime	DATETIME
active	TINYINT(1)
dataCenter	VARCHAR(255)
isVirtual	TINYINT(1)

storage_controller	
Column	
id	INT
storageid	INT
wwn	VARCHAR(255)
model	VARCHAR(255)
manufacturer	VARCHAR(255)
driver	VARCHAR(255)
firmware	VARCHAR(255)
numberOfPorts	VARCHAR(255)

switch_port	
Column	
id	INT
switchid	INT
fabricid	INT
virtualSwitchid	INT
wwn	VARCHAR(255)
status	VARCHAR(100)
rawPortStatus	VARCHAR(255)
type	VARCHAR(255)
portPhysicalState	VARCHAR(255)
number	BIGINT
blade	BIGINT
portid	VARCHAR(255)
name	VARCHAR(255)
speed	VARCHAR(12)
fc4Protocol	VARCHAR(255)
classOfService	VARCHAR(255)
gbicType	VARCHAR(255)
url	VARCHAR(255)
active	TINYINT(1)
isGenerated	TINYINT(1)

tape	
Column	
id	INT
name	VARCHAR(255)
identifier	VARCHAR(768)
ip	VARCHAR(255)
manufacturer	VARCHAR(255)
serialNumber	VARCHAR(255)
active	TINYINT(1)

tape_controller	
Column	
id	INT
tapeid	INT
wwn	VARCHAR(255)
model	VARCHAR(255)
manufacturer	VARCHAR(255)
driver	VARCHAR(255)
firmware	VARCHAR(255)
numberOfPorts	VARCHAR(255)

virtual_switch	
Column	
id	INT
fabricid	INT
wwn	VARCHAR(255)
ip	VARCHAR(255)
name	VARCHAR(255)
domainid	VARCHAR(255)
domainidType	VARCHAR(255)
priority	VARCHAR(255)
switchRole	VARCHAR(255)
chassisWwn	VARCHAR(255)
npv	TINYINT(1)
generated	TINYINT(1)
type	ENUM
isGenerated	TINYINT(1)

generic_device	
Column	
id	INT
wwn	VARCHAR(255)
identifier	VARCHAR(768)
manufacturer	VARCHAR(255)
model	VARCHAR(255)
firmware	VARCHAR(255)
driver	VARCHAR(255)
serialNumber	VARCHAR(255)

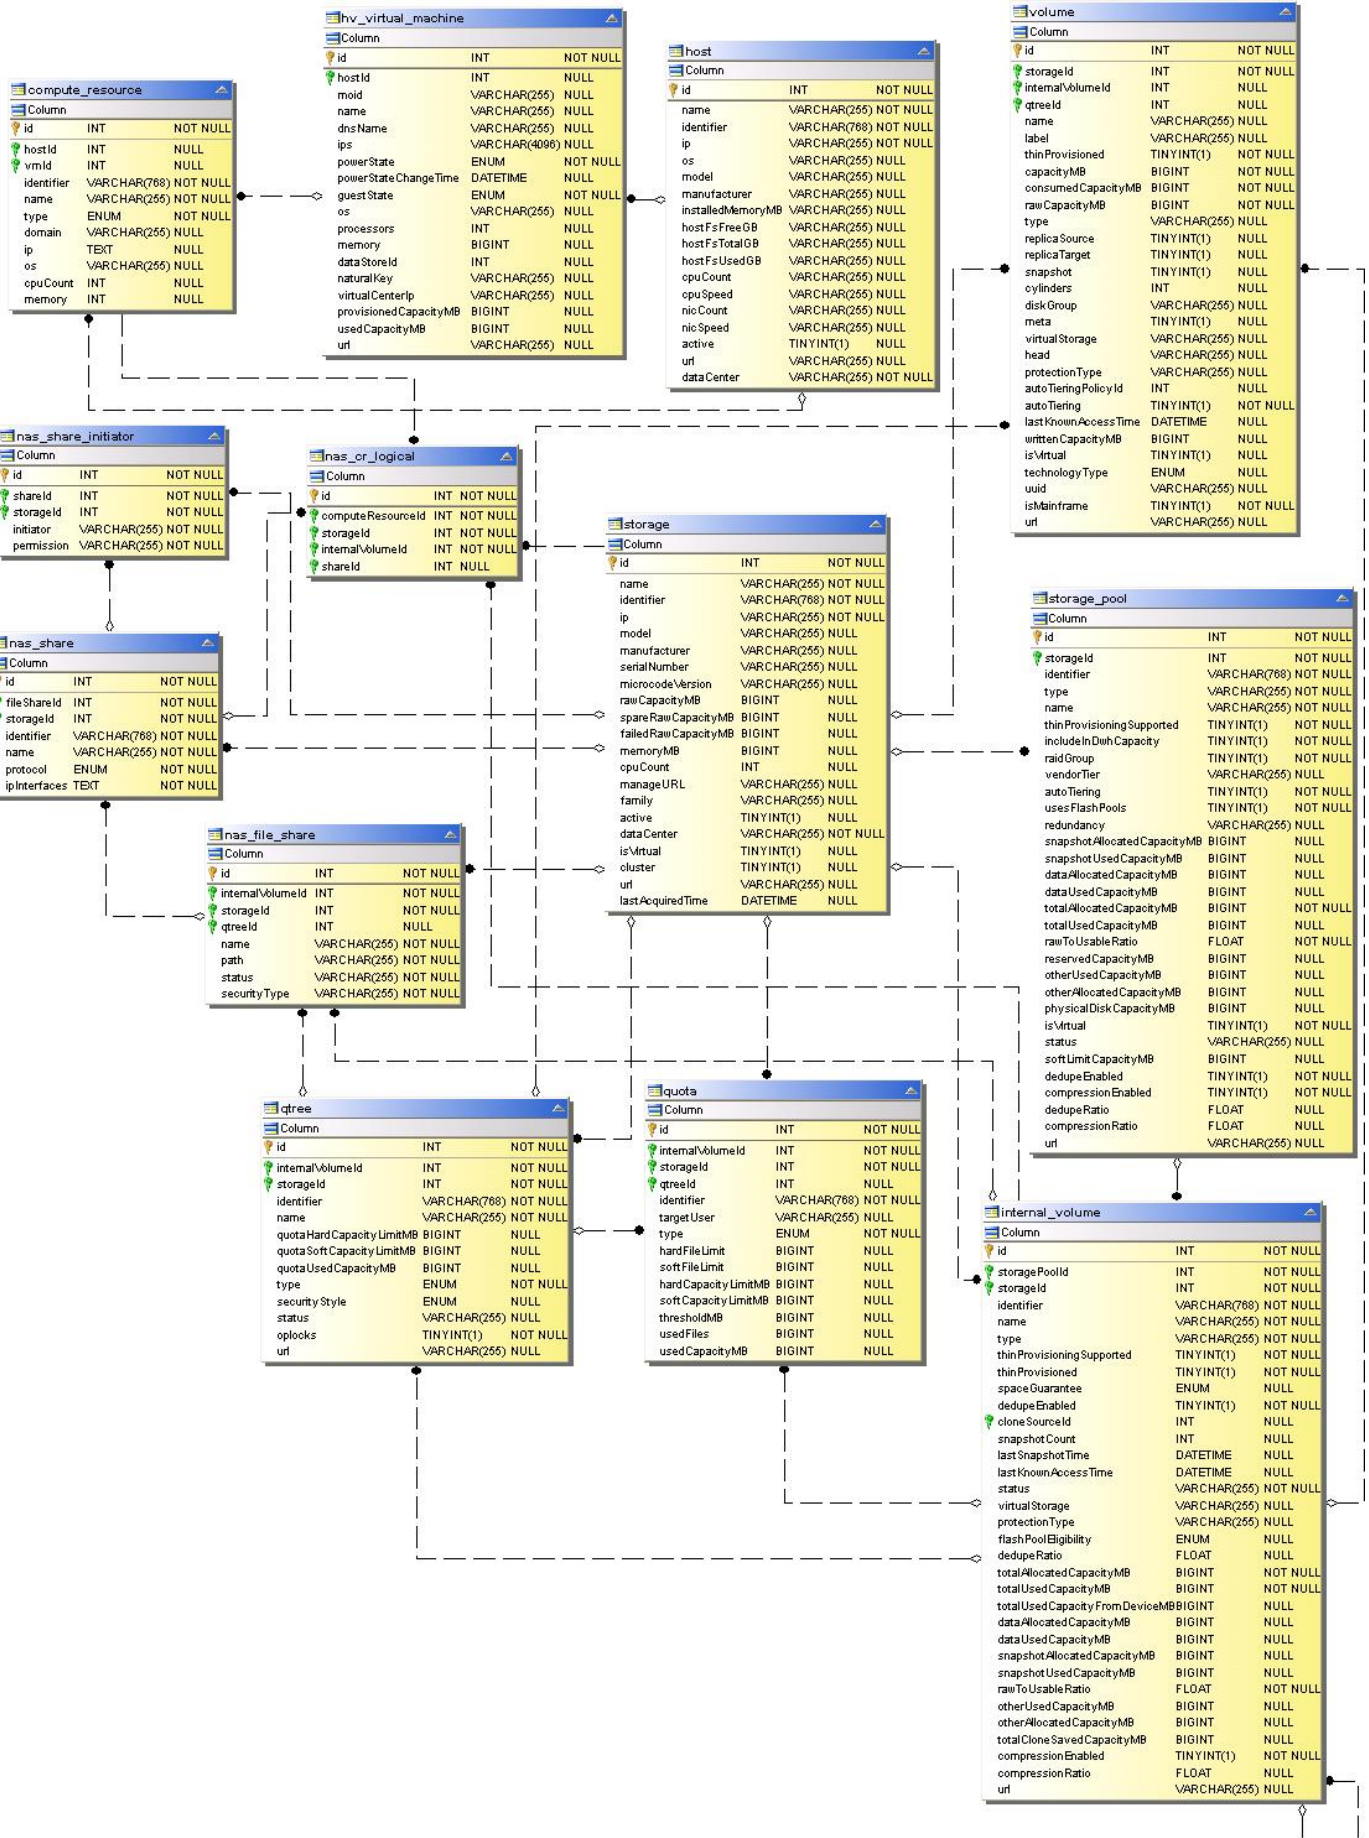
tape_port	
Column	
id	INT
controllerid	INT
tapeid	INT
wwn	VARCHAR(255)
nodeWwn	VARCHAR(255)
portid	VARCHAR(255)
name	VARCHAR(255)
speed	VARCHAR(12)
controller	VARCHAR(255)
url	VARCHAR(255)
active	TINYINT(1)

port_connectivity	
Column	
id	INT
portid	INT
type	ENUM
wwn	VARCHAR(255)
connectedid	INT
connectedType	ENUM
connectedWwn	VARCHAR(255)

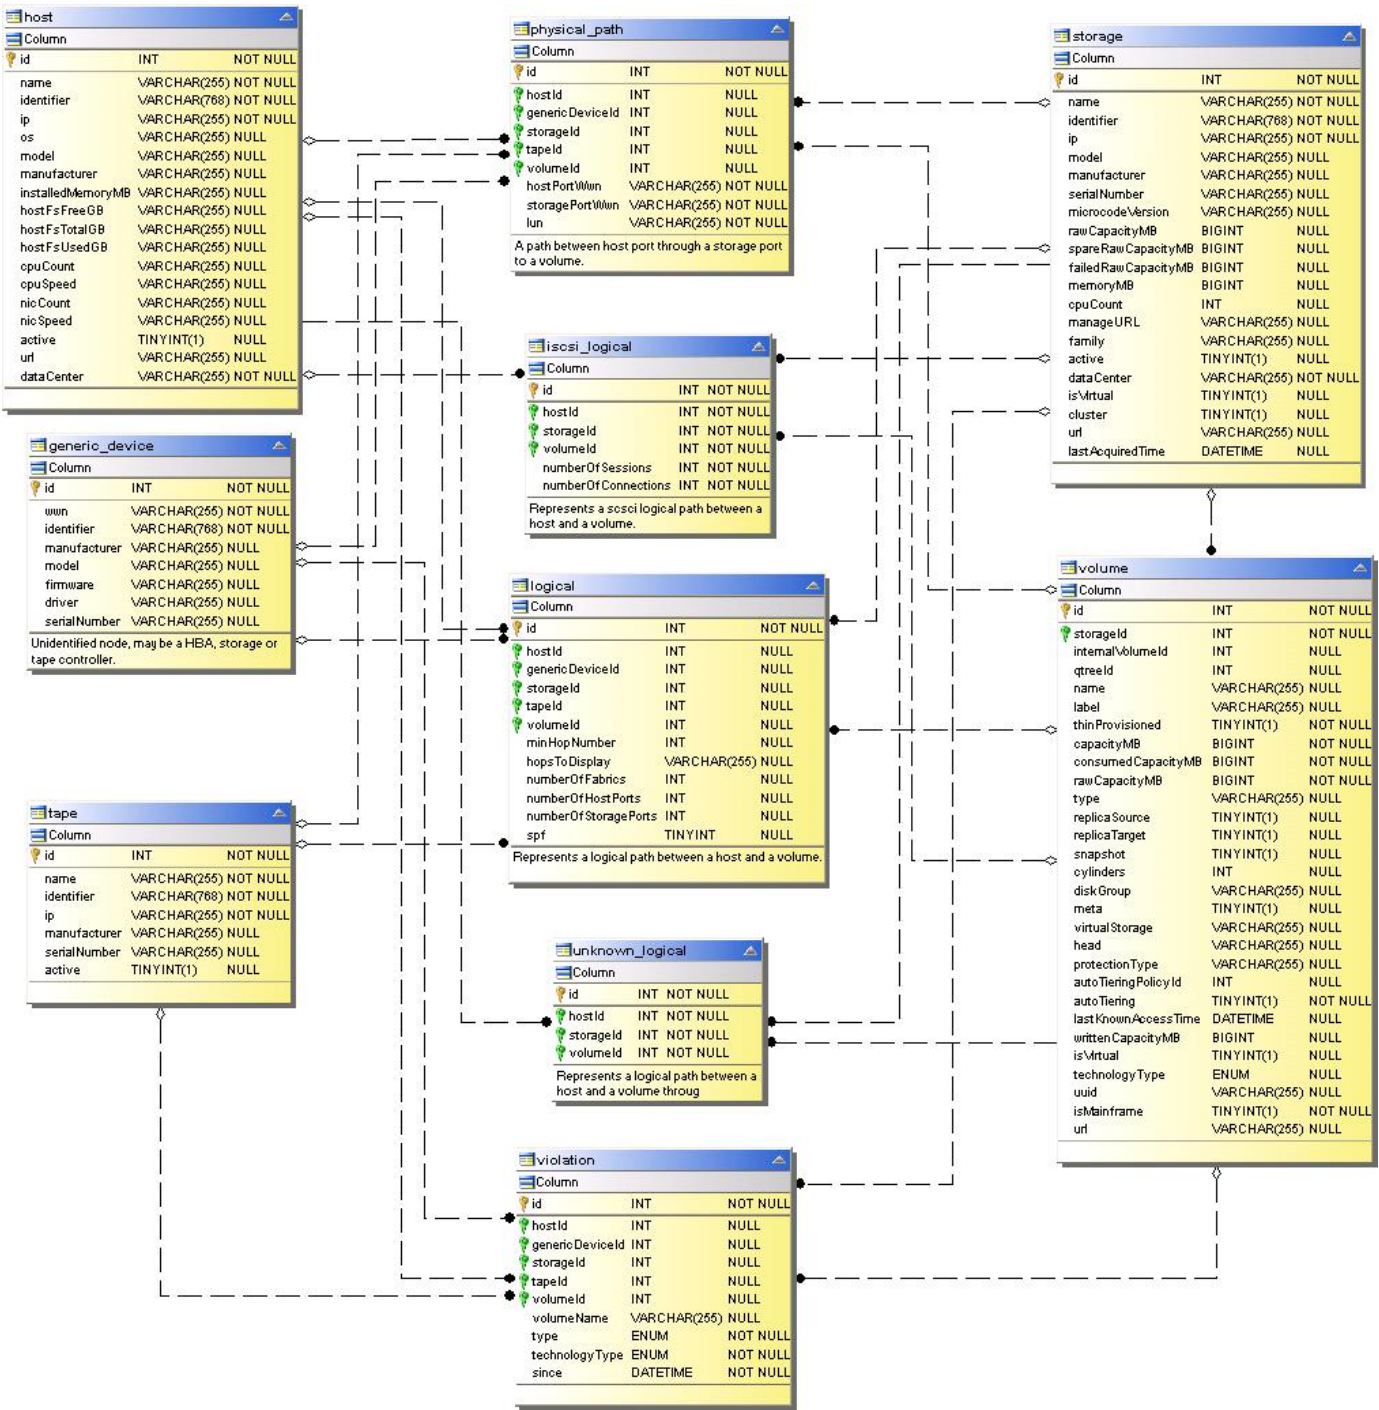
generic_device_port	
Column	
id	INT
genericDeviceid	INT
wwn	VARCHAR(255)
number	BIGINT
portid	VARCHAR(255)
name	VARCHAR(255)
speed	VARCHAR(12)
url	VARCHAR(255)
active	TINYINT(1)

fc_name_server_entry	
Column	
id	INT
portid	INT
type	ENUM
wwn	VARCHAR(255)
connectedSwitchPortid	INT
connectedSwitchPortWwn	VARCHAR(255)
physicalPortWwn	VARCHAR(255)
fcid	VARCHAR(255)






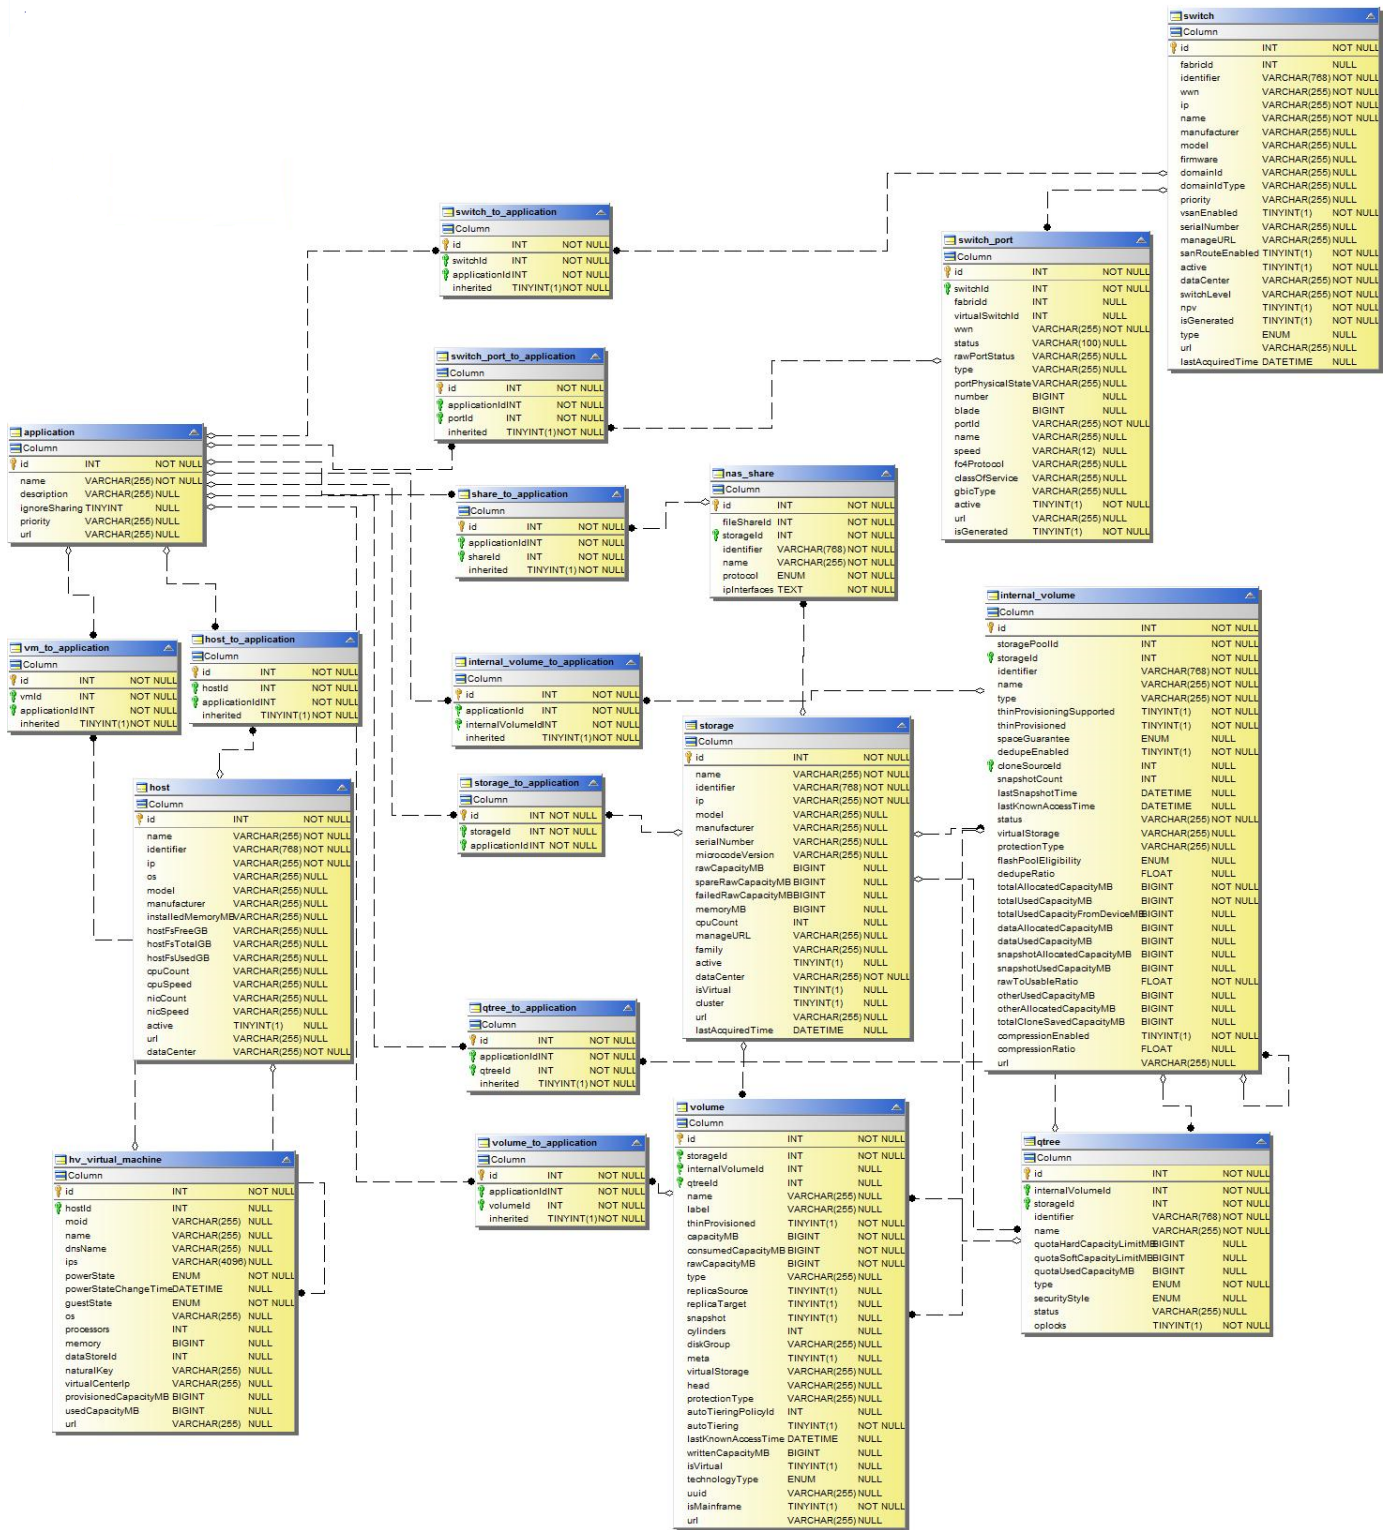
Paths and Violations



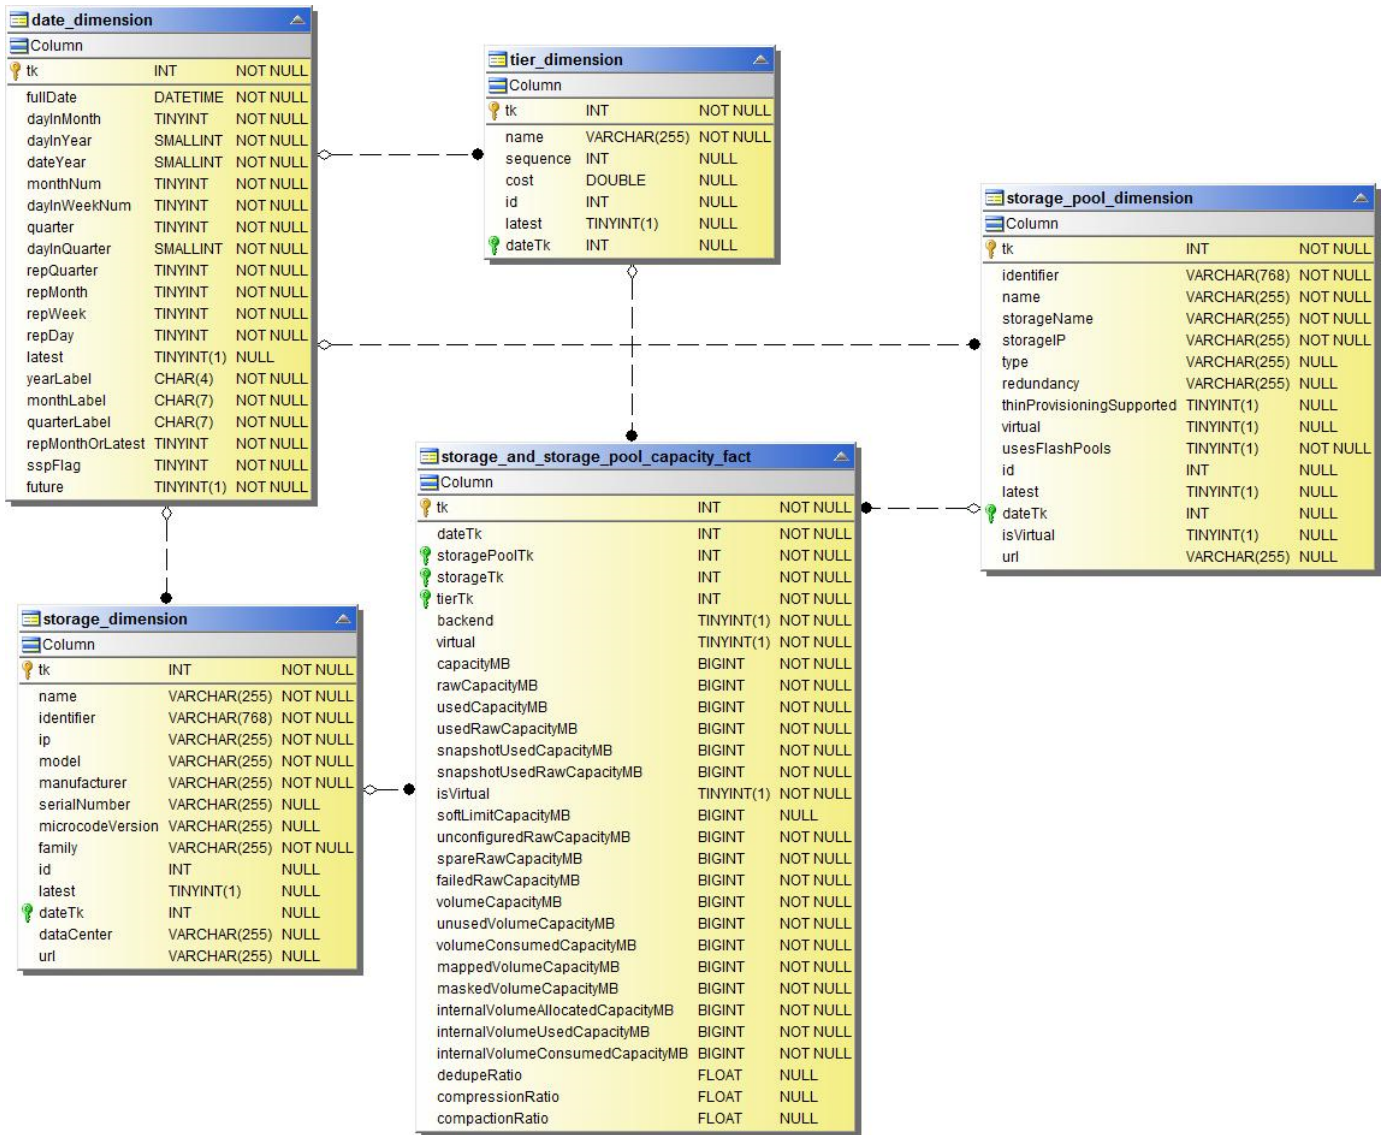
Annotations



Applications



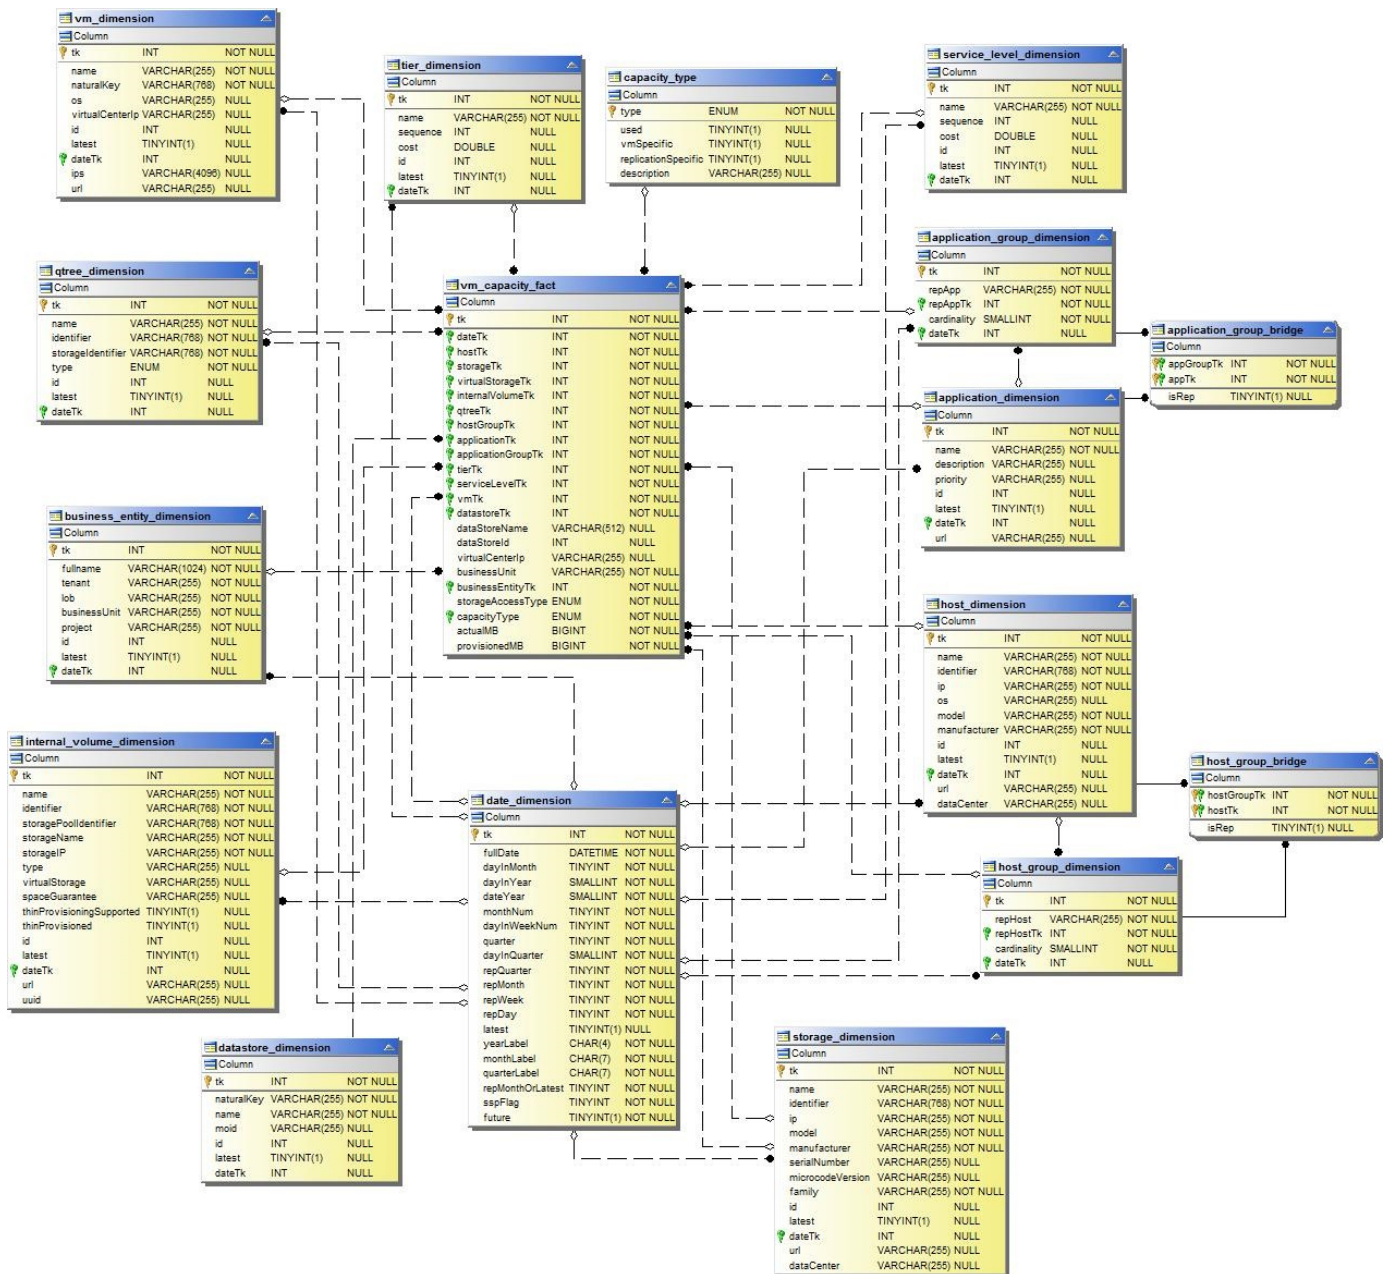
Storage and Storage Pool Capacity



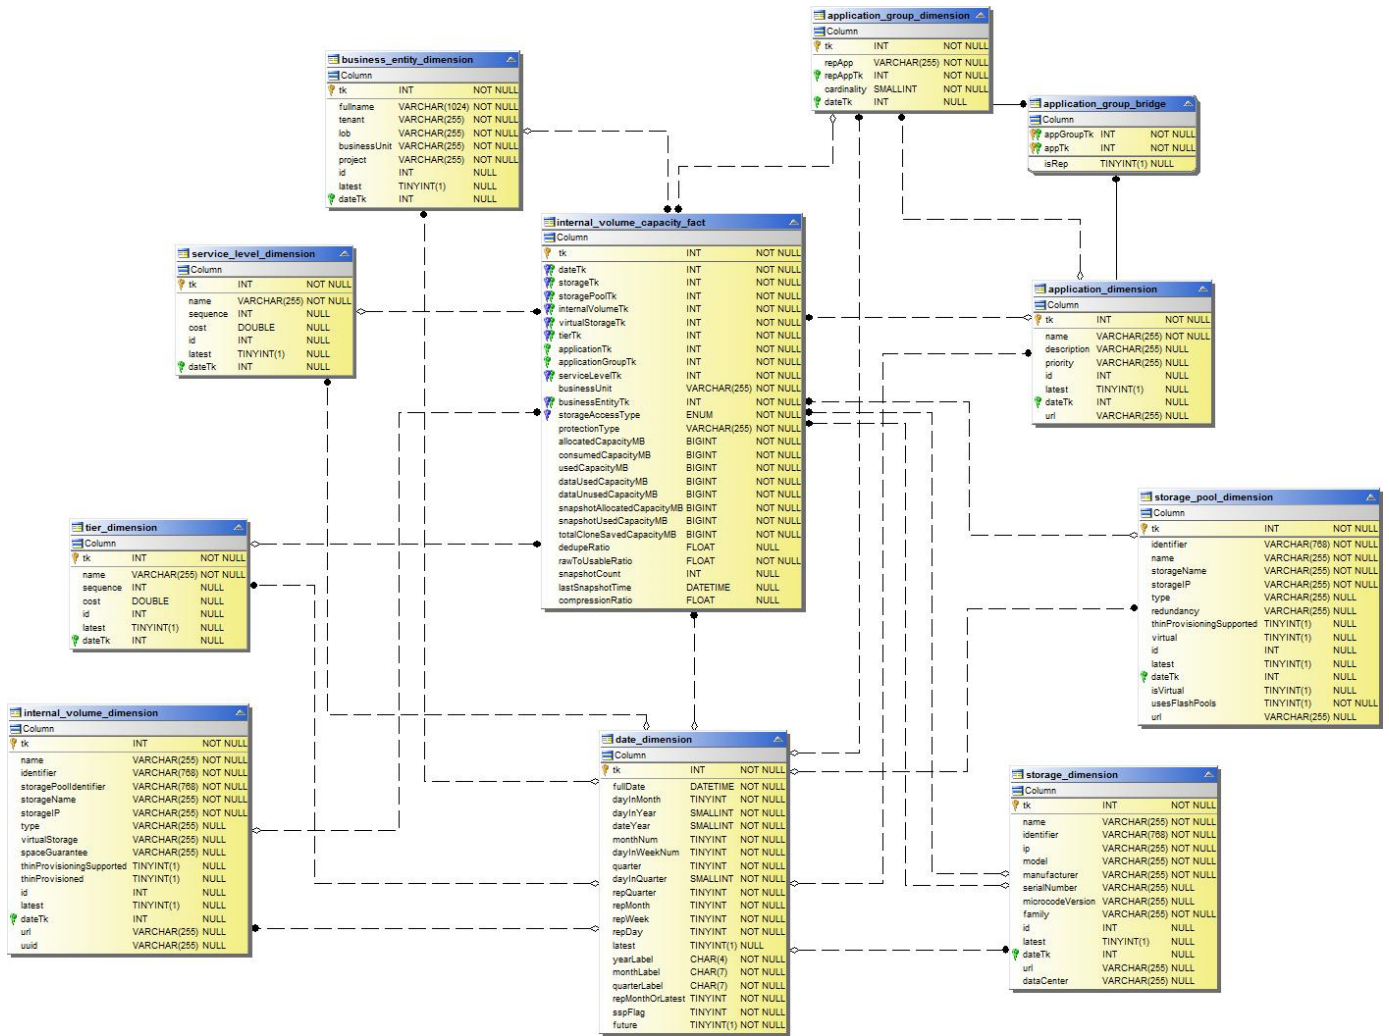
Volume Capacity

[Volume Capacity]

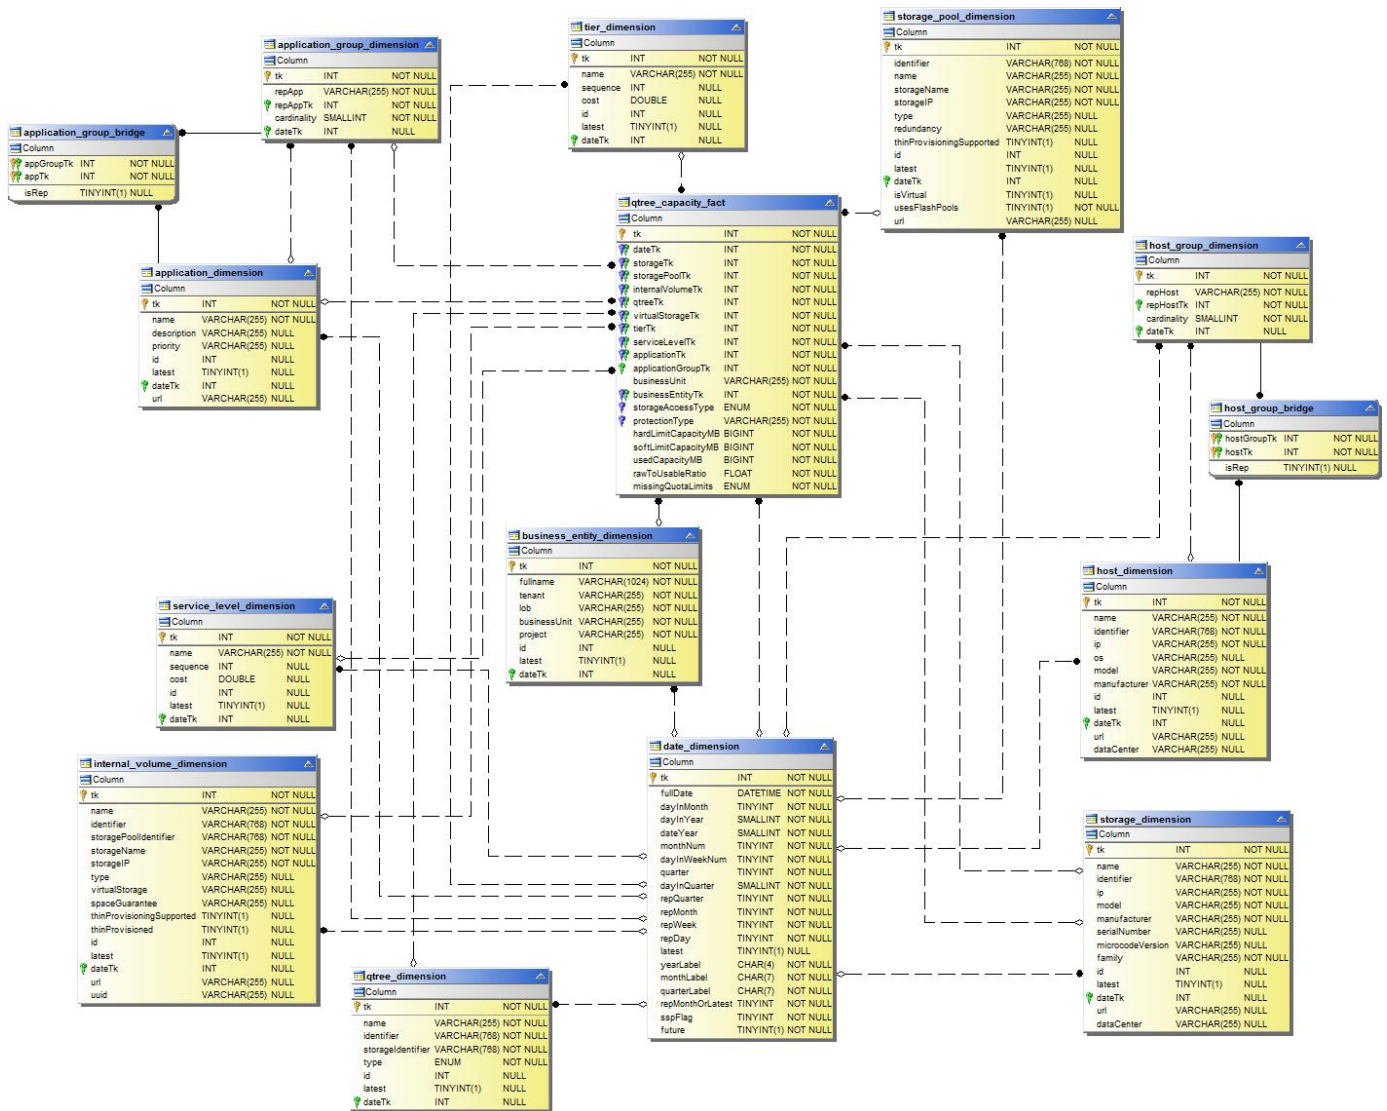
VM Capacity



Internal Volume Capacity



Qtree Capacity



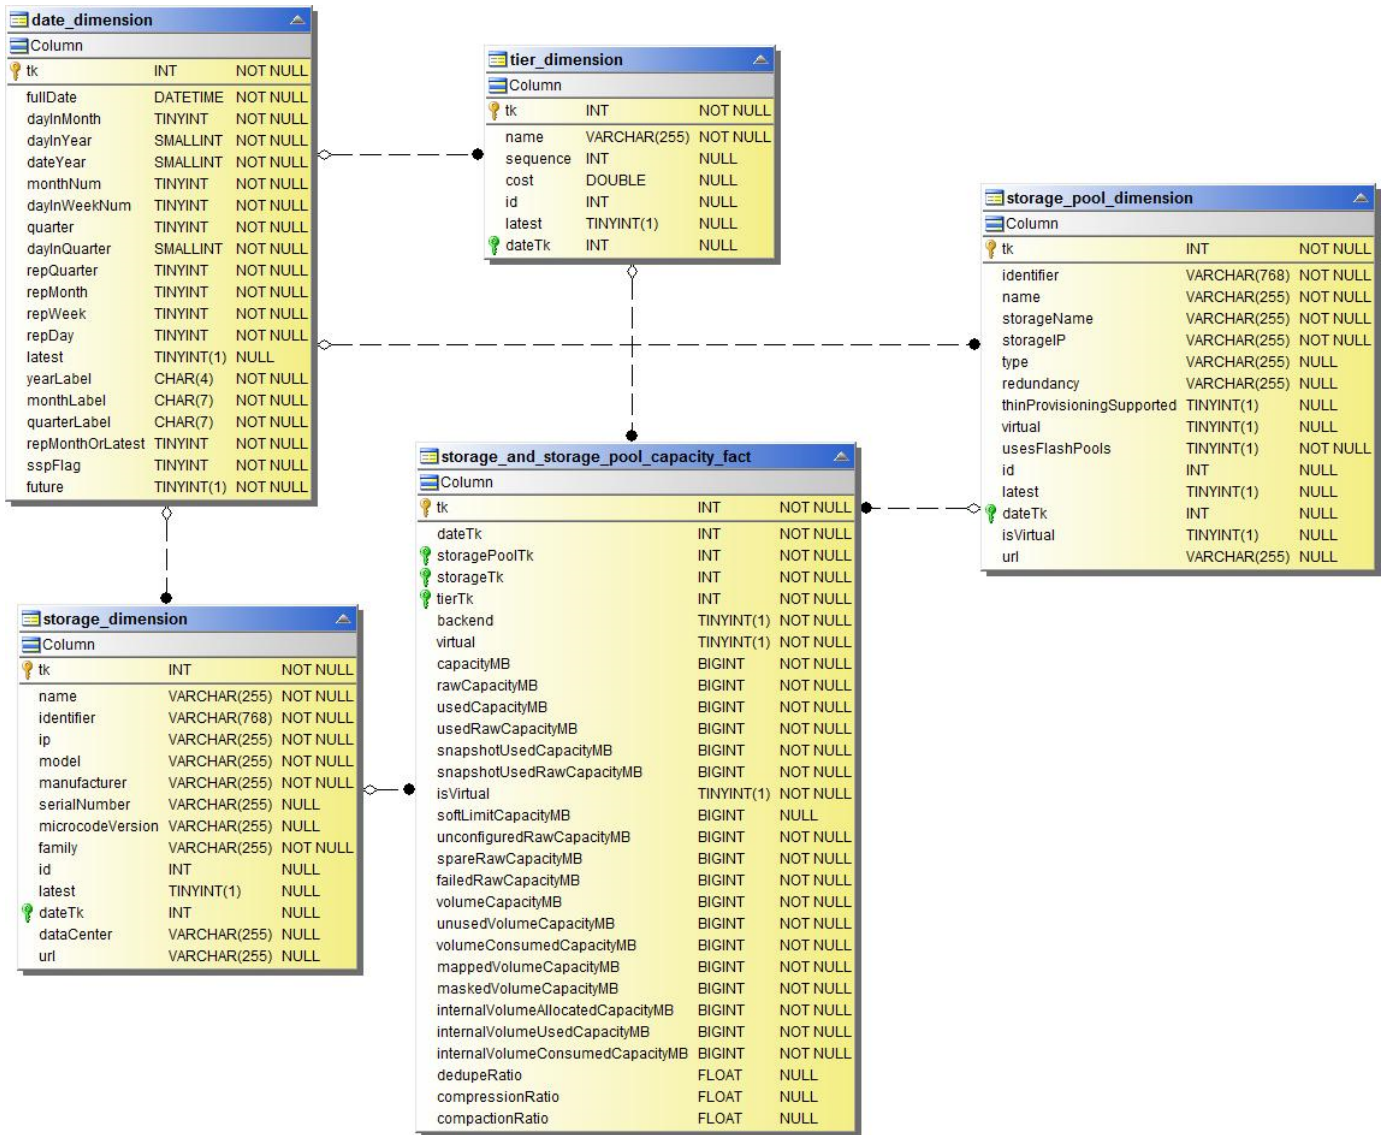
Storage Capacity Efficiency

efficiency_fact		
Column		
tk	INT	NOT NULL
dateTk	INT	NOT NULL
storageTk	INT	NOT NULL
rawCapacityMB	BIGINT	NOT NULL
backendCapacityMB	BIGINT	NOT NULL
storageTechnology	VARCHAR(255)	NULL
gainMB	BIGINT	NOT NULL
lossMB	BIGINT	NOT NULL
potentialGainMB	BIGINT	NOT NULL
potentialLossMB	BIGINT	NOT NULL

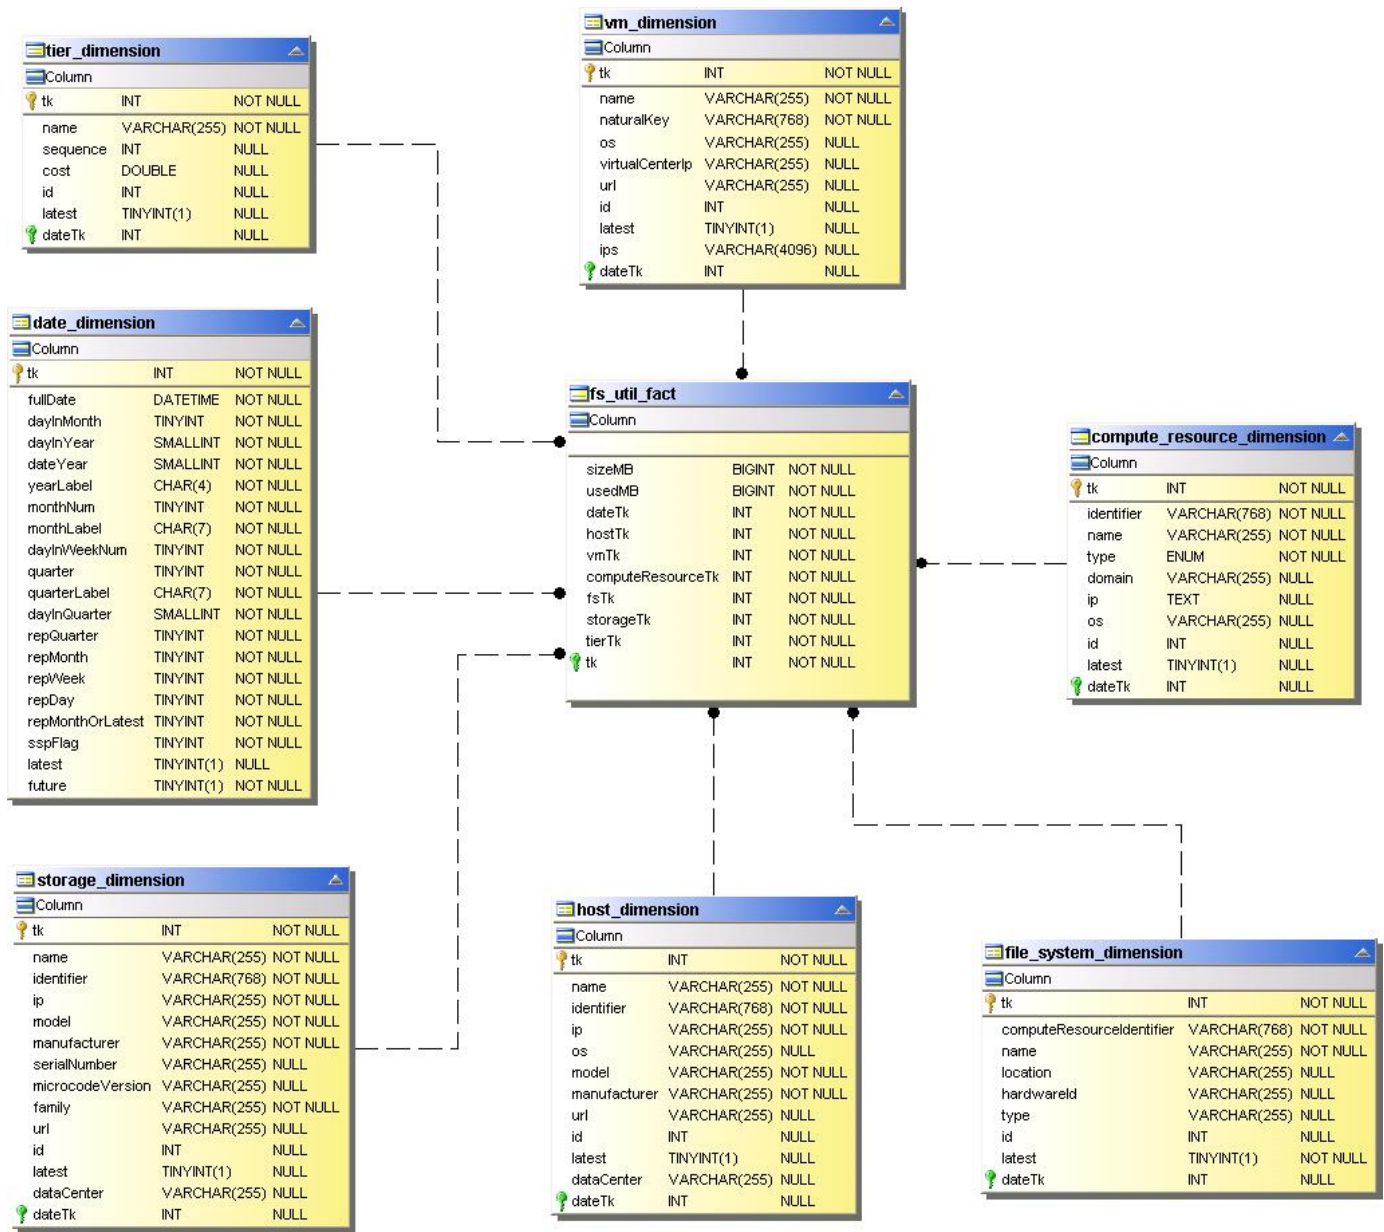
date_dimension		
Column		
tk	INT	NOT NULL
fullDate	DATETIME	NOT NULL
dayInMonth	TINYINT	NOT NULL
dayInYear	SMALLINT	NOT NULL
dateYear	SMALLINT	NOT NULL
monthNum	TINYINT	NOT NULL
dayInWeekNum	TINYINT	NOT NULL
quarter	TINYINT	NOT NULL
dayInQuarter	SMALLINT	NOT NULL
repQuarter	TINYINT	NOT NULL
repMonth	TINYINT	NOT NULL
repWeek	TINYINT	NOT NULL
repDay	TINYINT	NOT NULL
latest	TINYINT(1)	NULL
yearLabel	CHAR(4)	NOT NULL
monthLabel	CHAR(7)	NOT NULL
quarterLabel	CHAR(7)	NOT NULL
repMonthOrLatest	TINYINT	NOT NULL
sspFlag	TINYINT	NOT NULL
future	TINYINT(1)	NOT NULL

storage_dimension		
Column		
tk	INT	NOT NULL
name	VARCHAR(255)	NOT NULL
identifier	VARCHAR(768)	NOT NULL
ip	VARCHAR(255)	NOT NULL
model	VARCHAR(255)	NOT NULL
manufacturer	VARCHAR(255)	NOT NULL
serialNumber	VARCHAR(255)	NULL
microcodeVersion	VARCHAR(255)	NULL
family	VARCHAR(255)	NOT NULL
id	INT	NULL
latest	TINYINT(1)	NULL
dateTk	INT	NULL
url	VARCHAR(255)	NULL
dataCenter	VARCHAR(255)	NULL

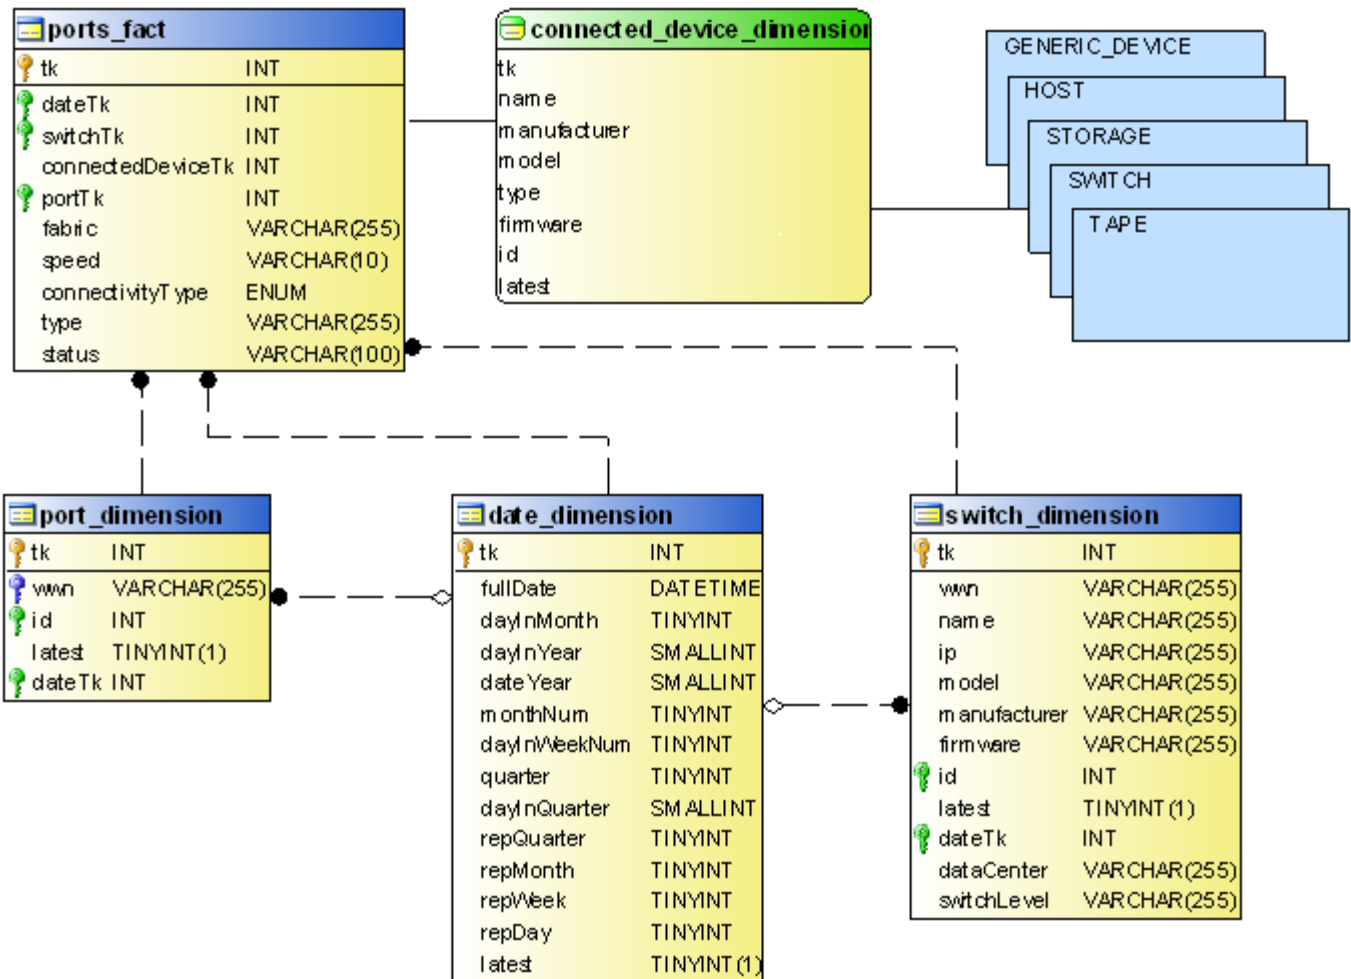
Storage and Storage Pool Capacity



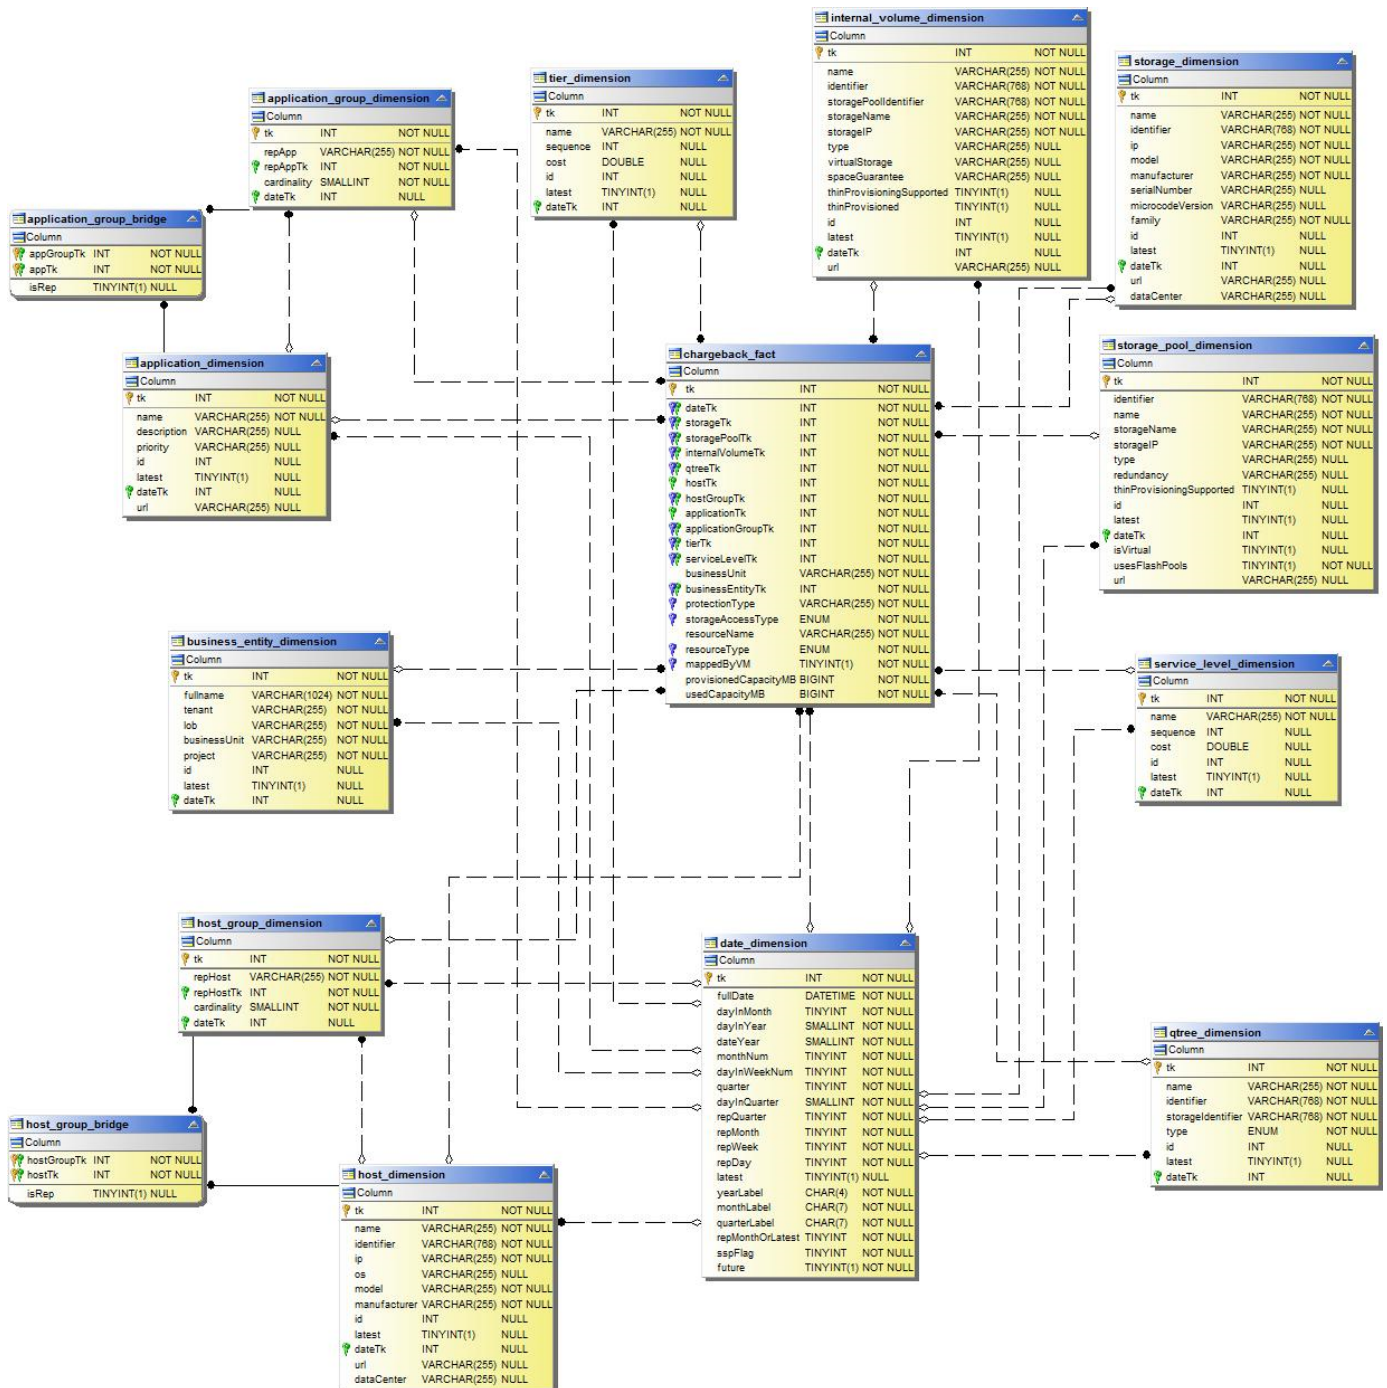
File System Utilization



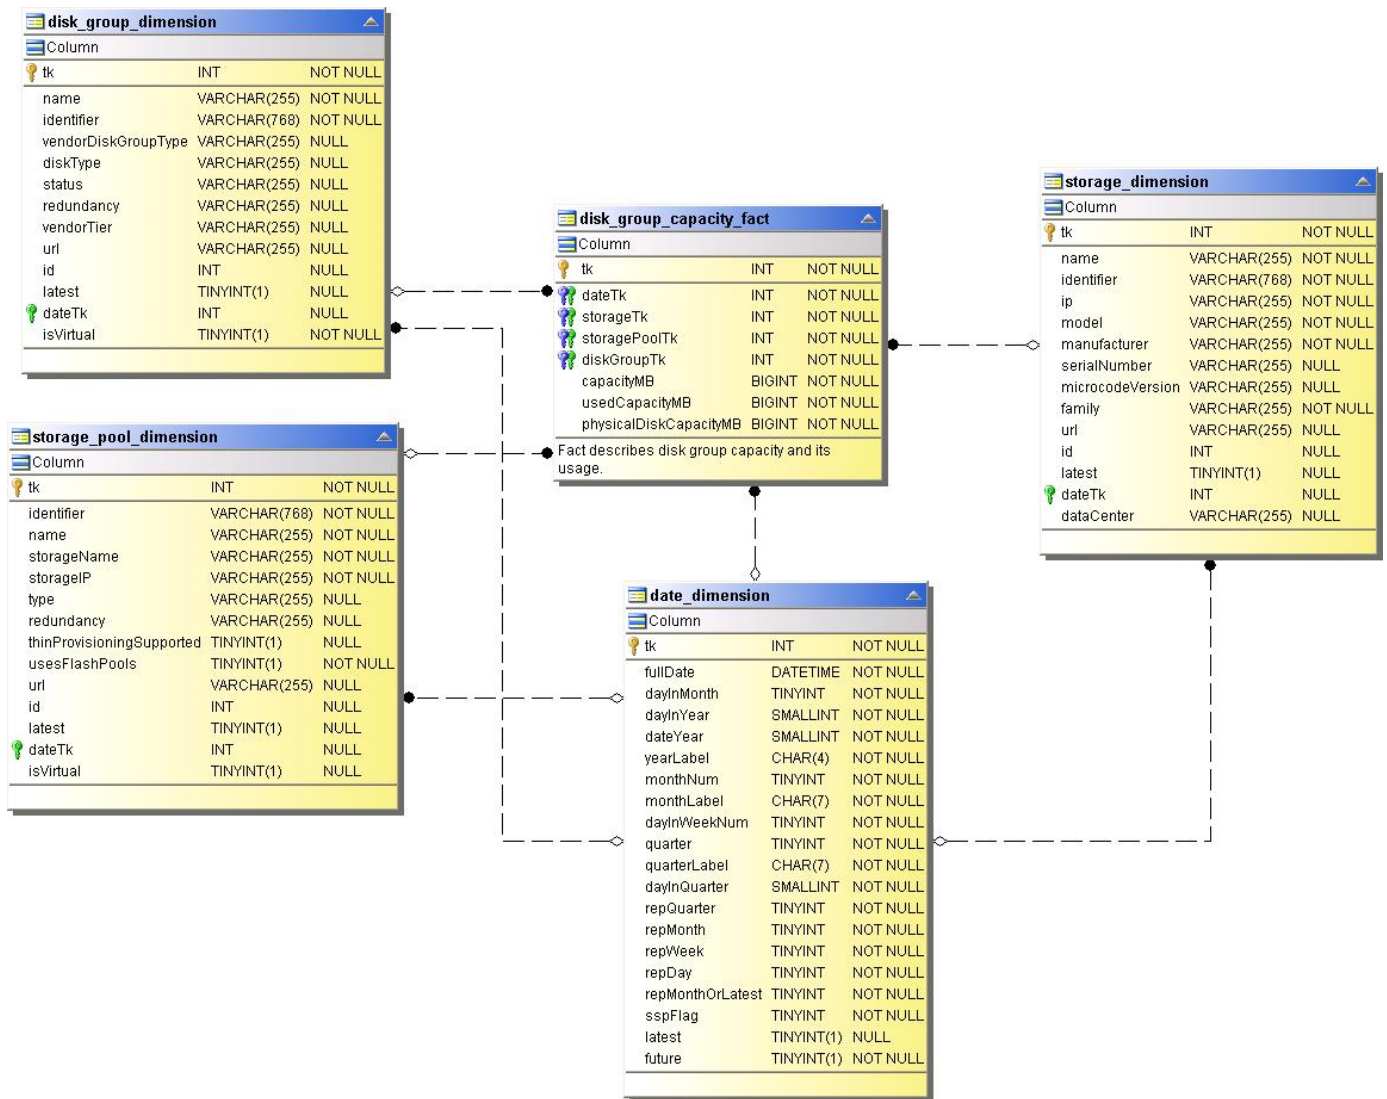
Port Capacity



Chargeback




Disk Group Capacity



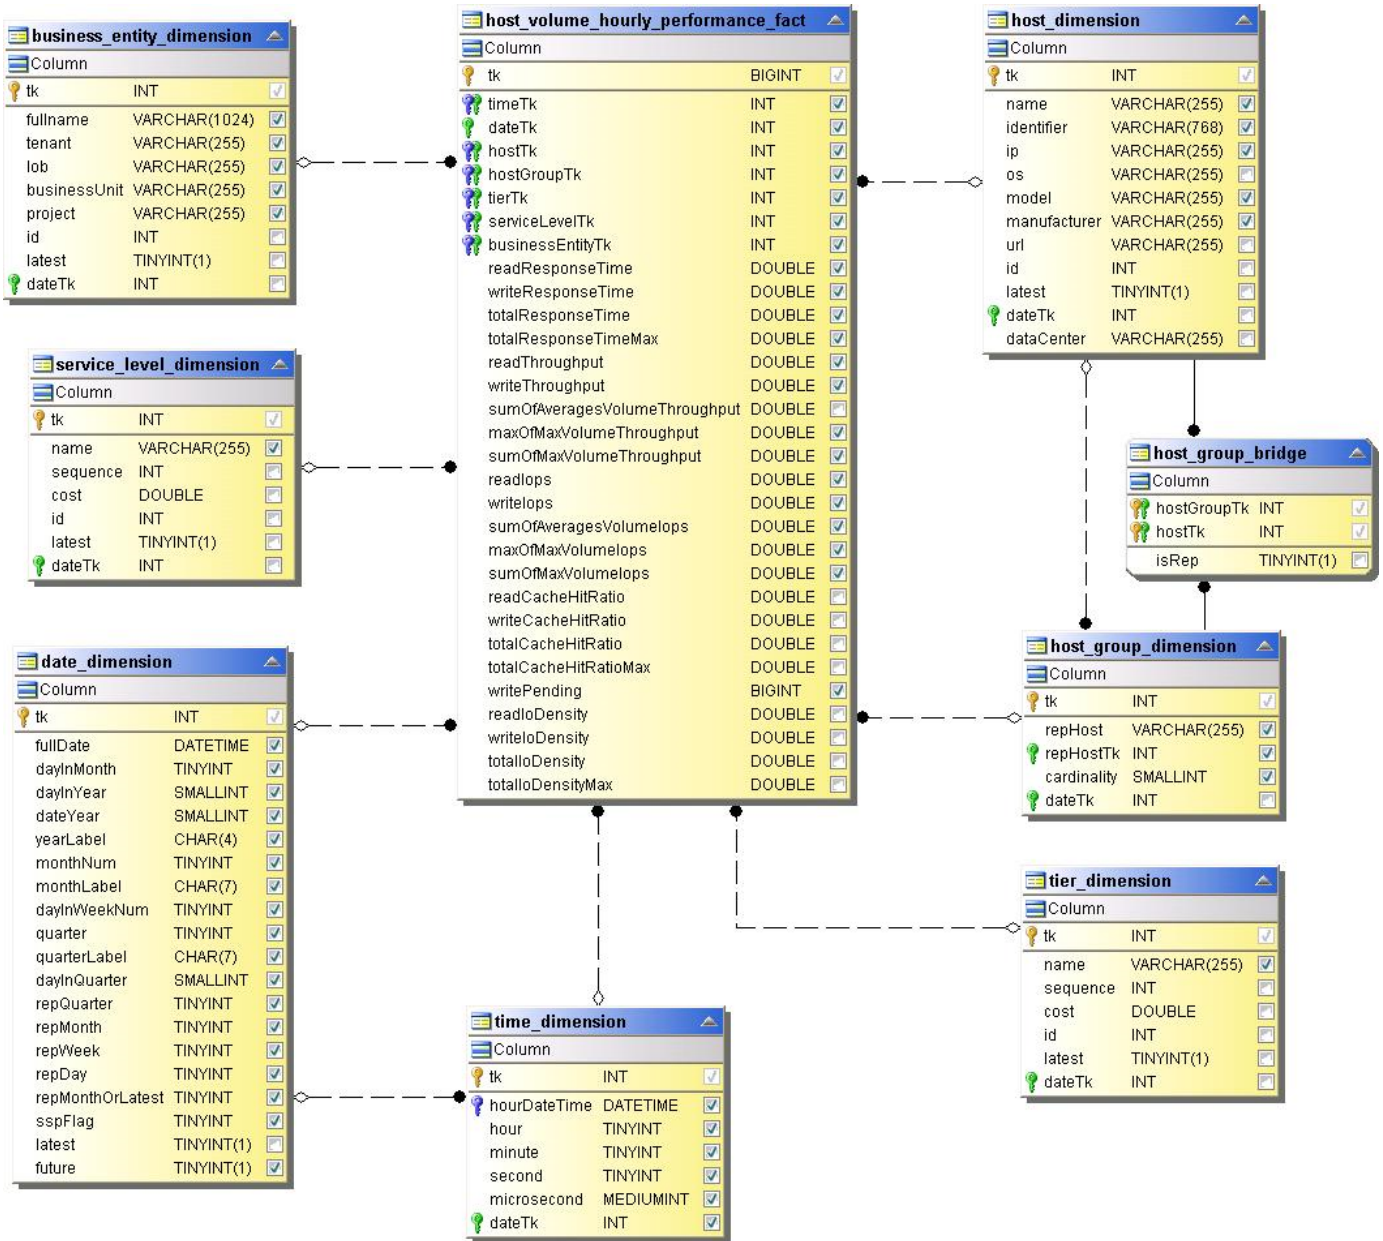
Performance Datamart

The following images describe the performance datamart.


Application Volume Hourly Performance




Internal Volume Hourly Performance




Internal Volume Hourly Performance




Internal Volume Daily Performance



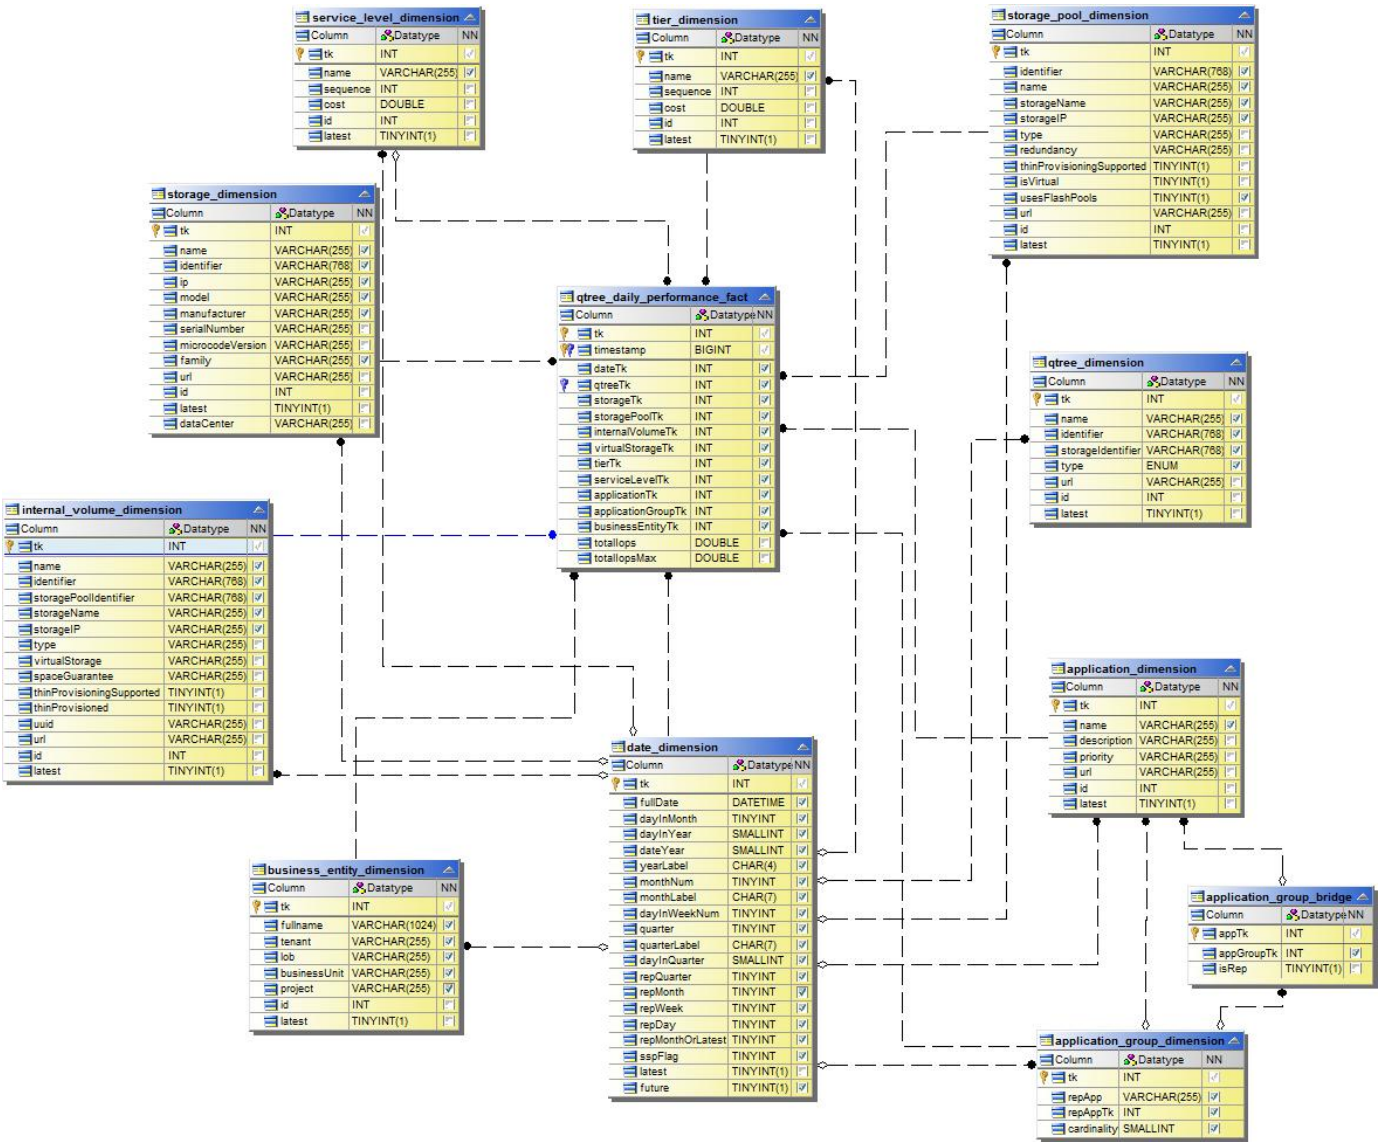
Volume Hourly Performance



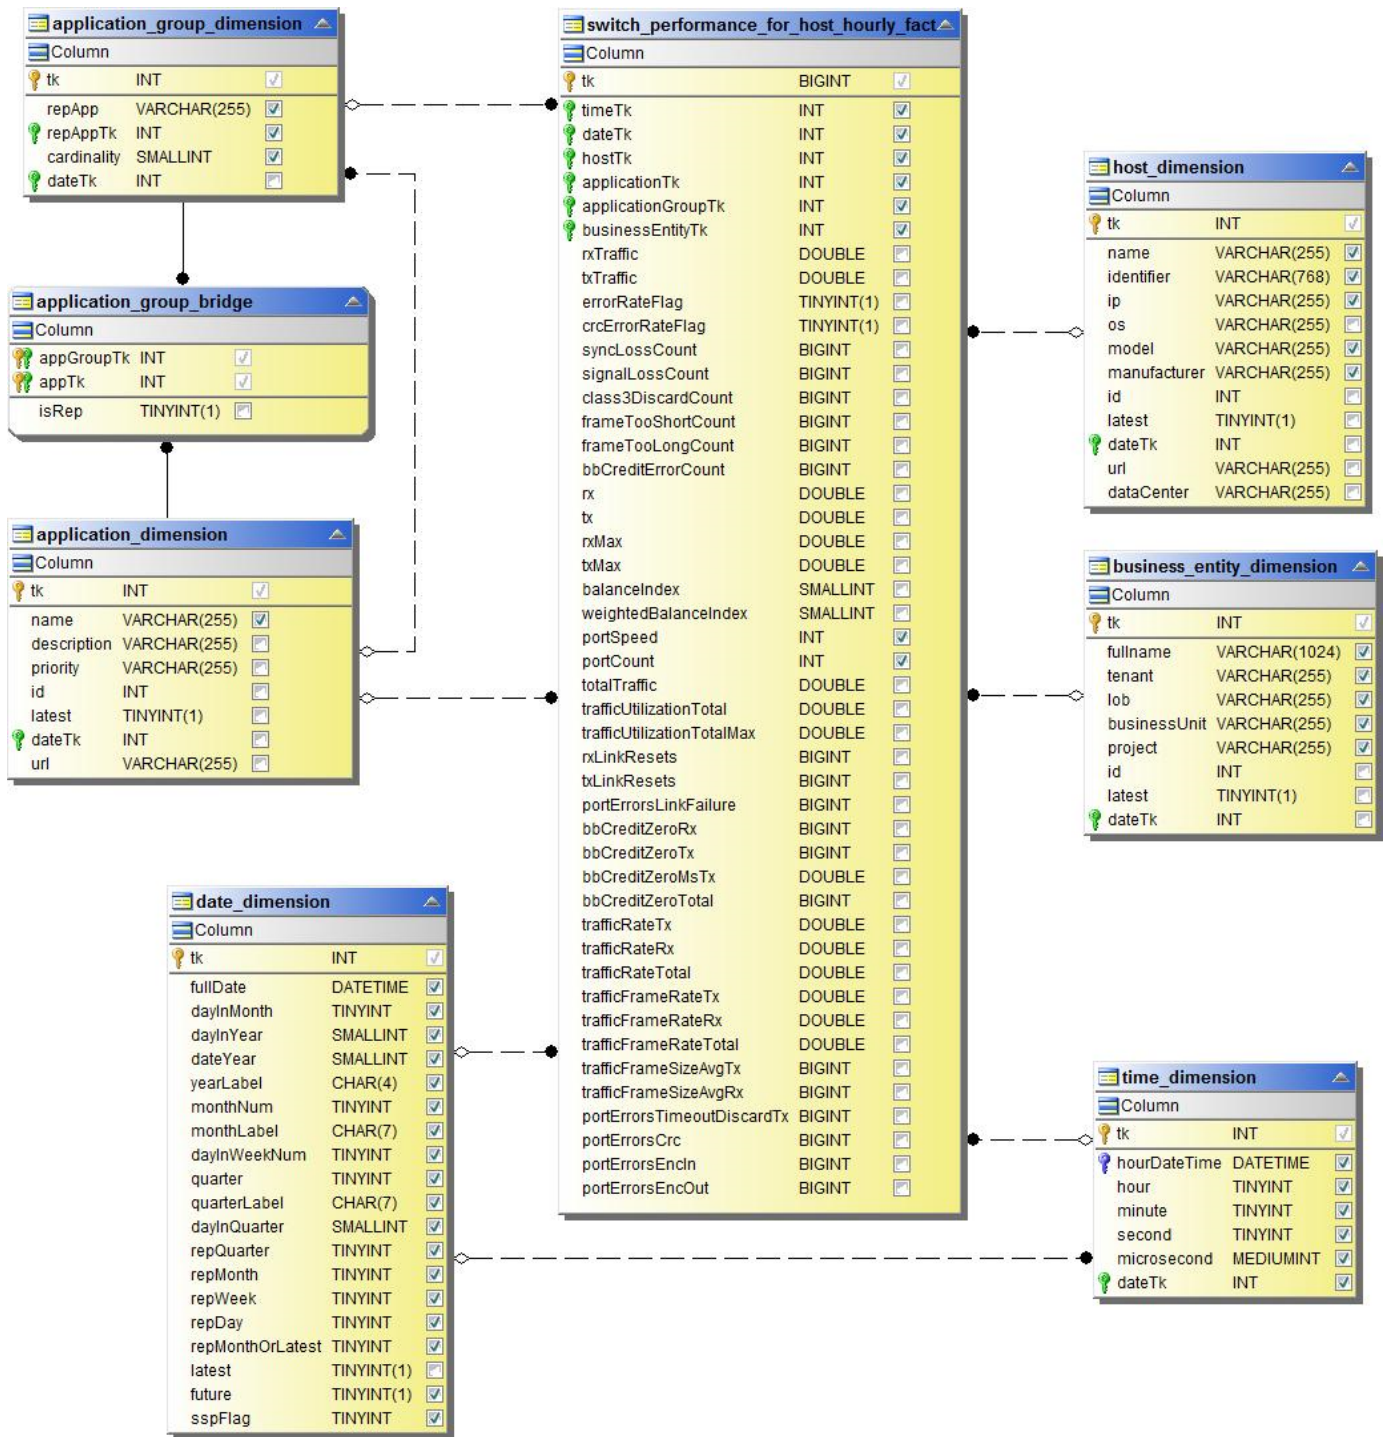
Volume Daily Performance



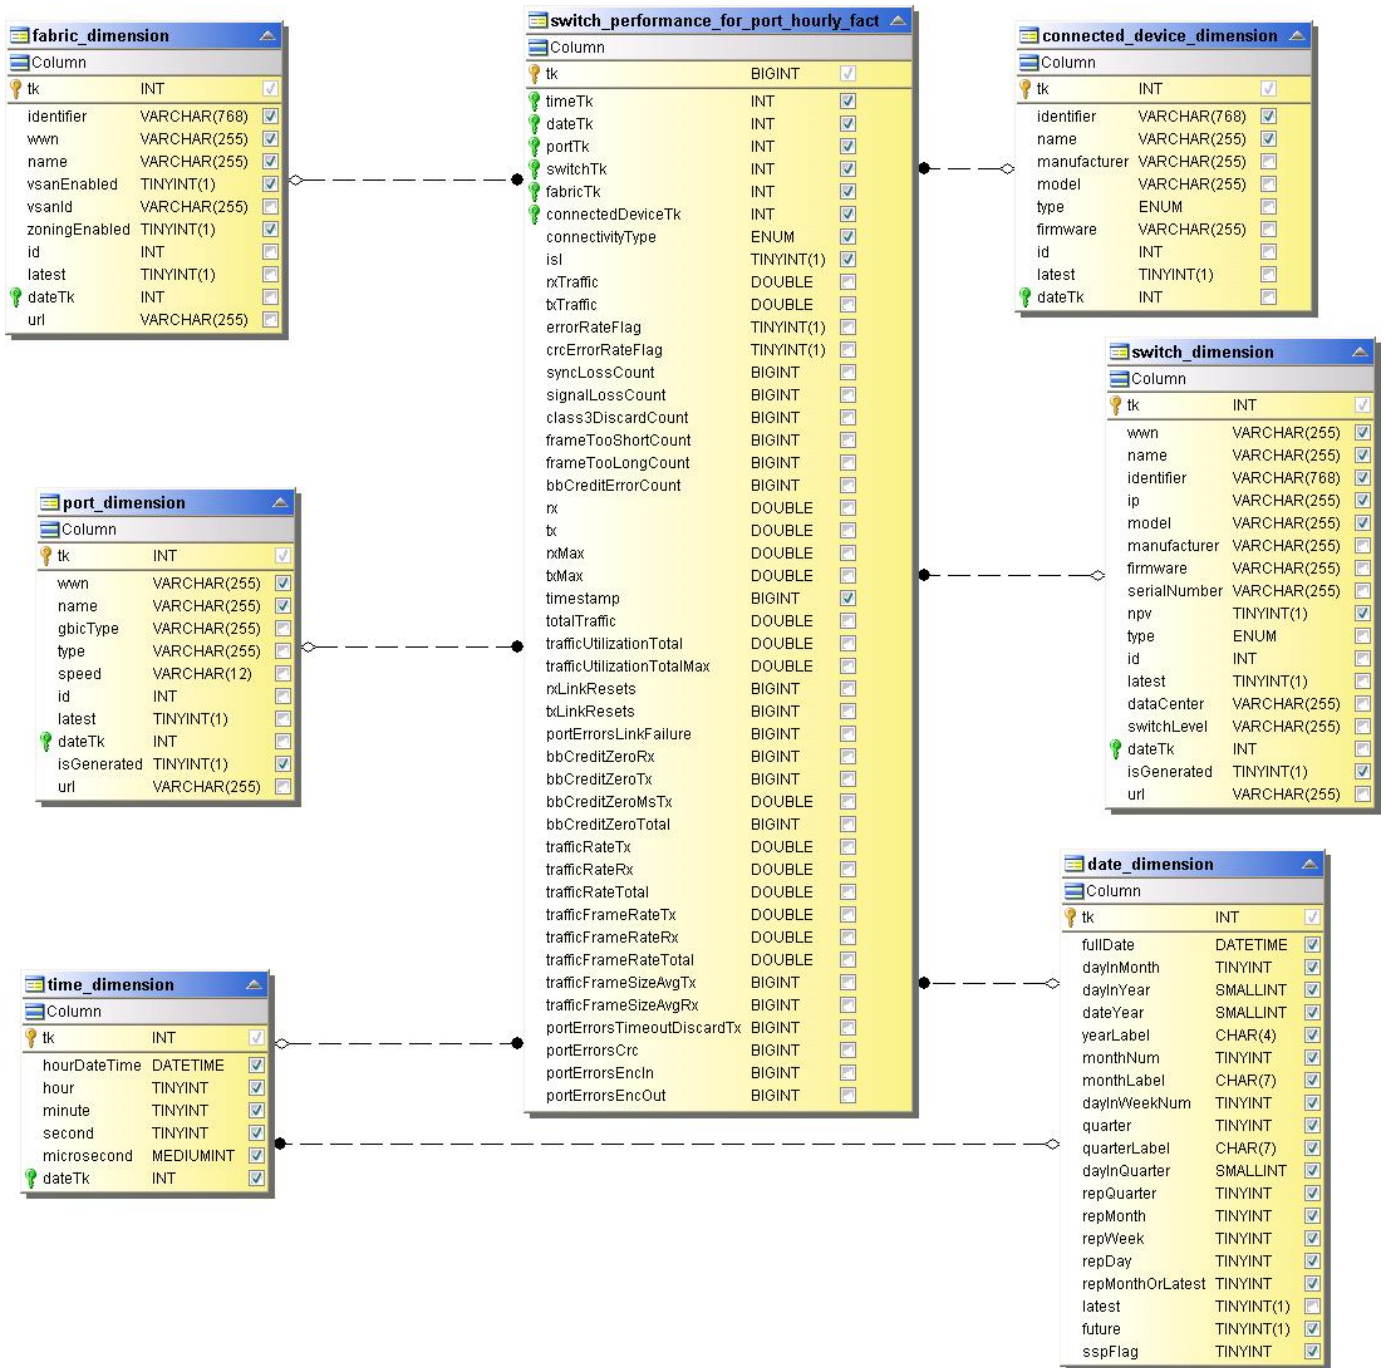
Qtree Daily Performance



Switch Hourly Performance for Host



Switch Hourly Performance for Port



Switch Hourly Performance for Storage

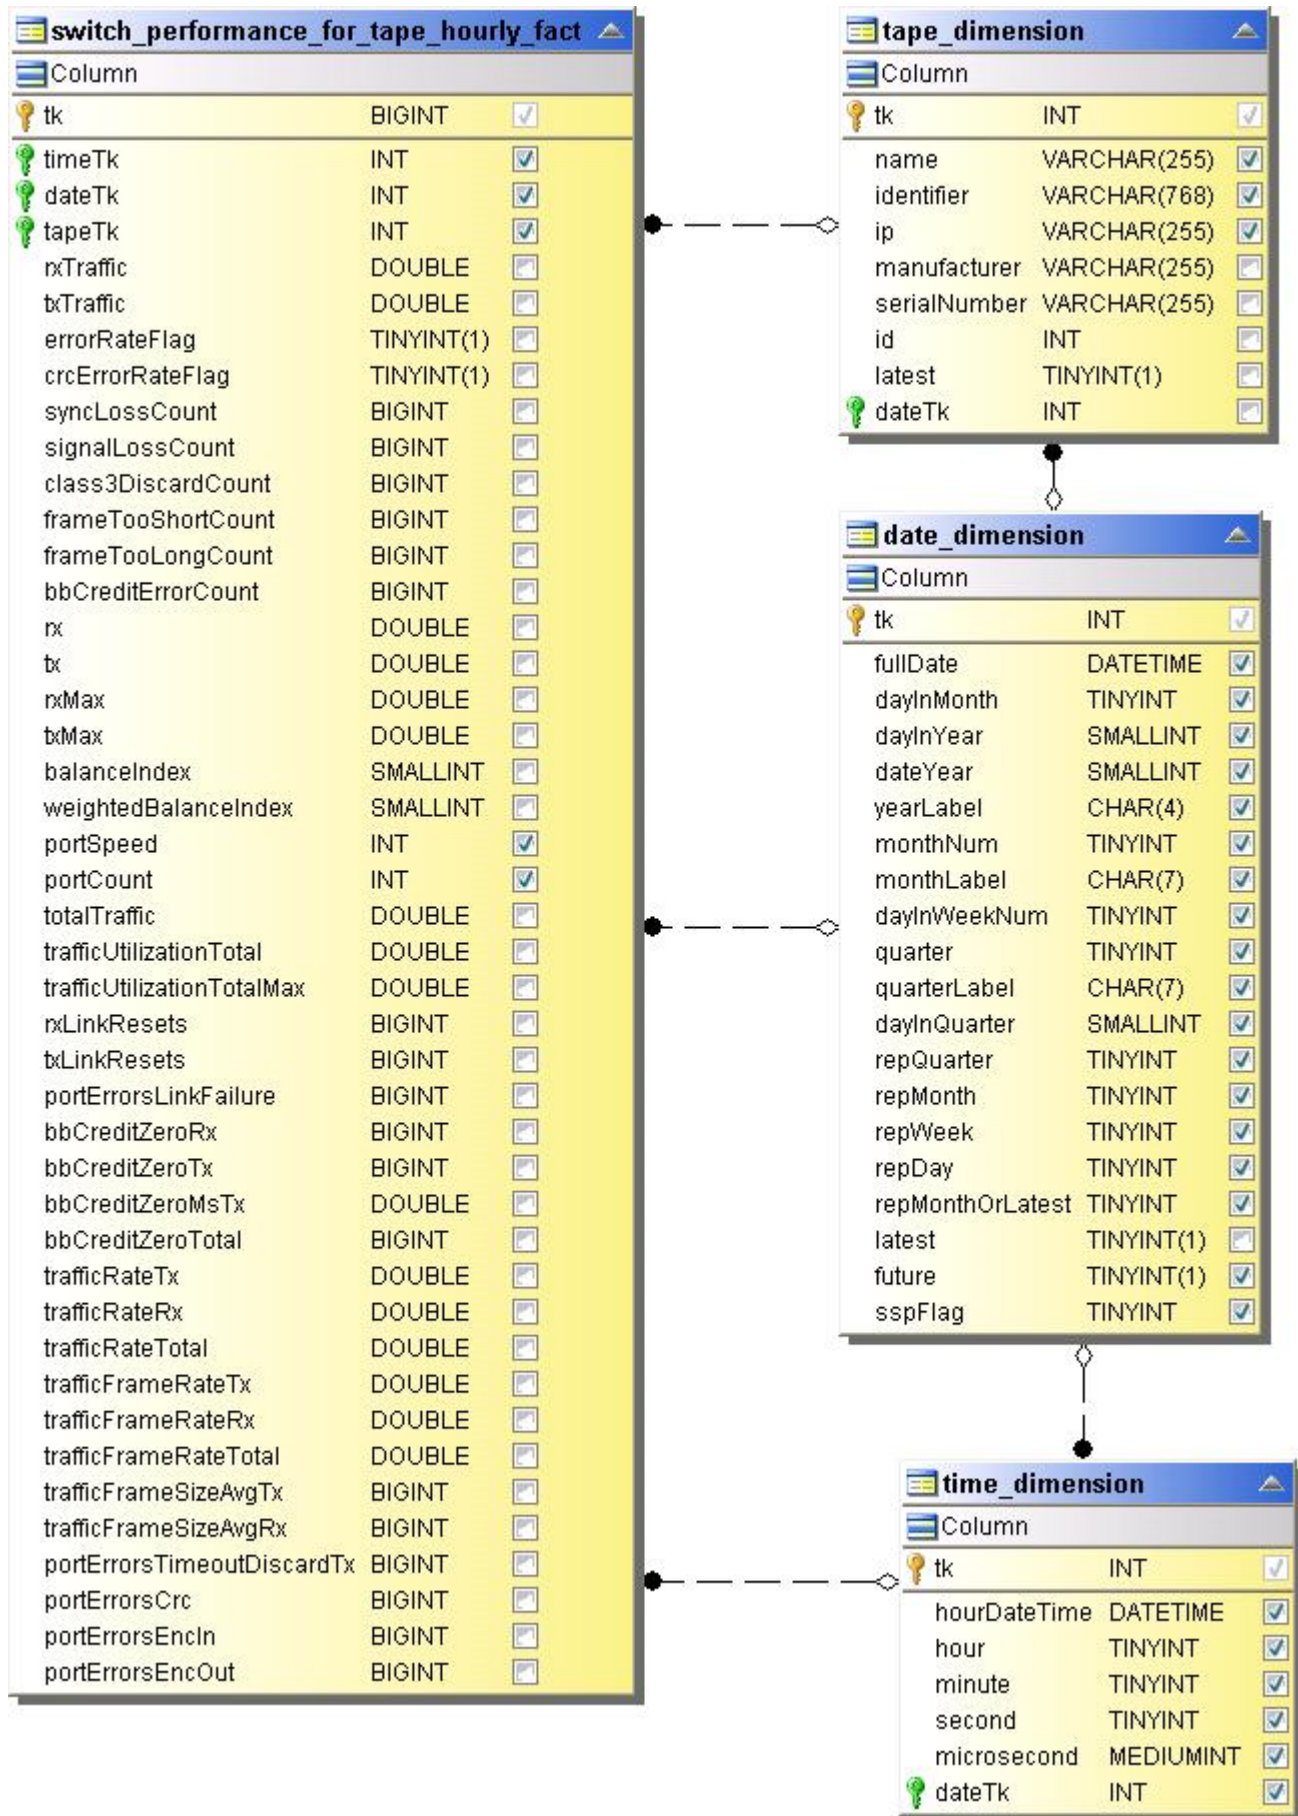
switch_performance_for_storage_hourly_fact		
Column		
tk	BIGINT	<input checked="" type="checkbox"/>
timeTk	INT	<input checked="" type="checkbox"/>
dateTk	INT	<input checked="" type="checkbox"/>
storageTk	INT	<input checked="" type="checkbox"/>
rxTraffic	DOUBLE	<input type="checkbox"/>
txTraffic	DOUBLE	<input type="checkbox"/>
errorRateFlag	TINYINT(1)	<input type="checkbox"/>
crcErrorRateFlag	TINYINT(1)	<input type="checkbox"/>
syncLossCount	BIGINT	<input type="checkbox"/>
signalLossCount	BIGINT	<input type="checkbox"/>
class3DiscardCount	BIGINT	<input type="checkbox"/>
frameTooShortCount	BIGINT	<input type="checkbox"/>
frameTooLongCount	BIGINT	<input type="checkbox"/>
bbCreditErrorCount	BIGINT	<input type="checkbox"/>
rx	DOUBLE	<input type="checkbox"/>
tx	DOUBLE	<input type="checkbox"/>
rxMax	DOUBLE	<input type="checkbox"/>
txMax	DOUBLE	<input type="checkbox"/>
balanceIndex	SMALLINT	<input type="checkbox"/>
weightedBalanceIndex	SMALLINT	<input type="checkbox"/>
portSpeed	INT	<input checked="" type="checkbox"/>
portCount	INT	<input checked="" type="checkbox"/>
totalTraffic	DOUBLE	<input type="checkbox"/>
trafficUtilizationTotal	DOUBLE	<input type="checkbox"/>
trafficUtilizationTotalMax	DOUBLE	<input type="checkbox"/>
rxLinkResets	BIGINT	<input type="checkbox"/>
txLinkResets	BIGINT	<input type="checkbox"/>
portErrorsLinkFailure	BIGINT	<input type="checkbox"/>
bbCreditZeroRx	BIGINT	<input type="checkbox"/>
bbCreditZeroTx	BIGINT	<input type="checkbox"/>
bbCreditZeroMsTx	DOUBLE	<input type="checkbox"/>
bbCreditZeroTotal	BIGINT	<input type="checkbox"/>
trafficRateTx	DOUBLE	<input type="checkbox"/>
trafficRateRx	DOUBLE	<input type="checkbox"/>
trafficRateTotal	DOUBLE	<input type="checkbox"/>
trafficFrameRateTx	DOUBLE	<input type="checkbox"/>
trafficFrameRateRx	DOUBLE	<input type="checkbox"/>
trafficFrameRateTotal	DOUBLE	<input type="checkbox"/>
trafficFrameSizeAvgTx	BIGINT	<input type="checkbox"/>
trafficFrameSizeAvgRx	BIGINT	<input type="checkbox"/>
portErrorsTimeoutDiscardTx	BIGINT	<input type="checkbox"/>
portErrorsCrc	BIGINT	<input type="checkbox"/>
portErrorsEncln	BIGINT	<input type="checkbox"/>
portErrorsEncOut	BIGINT	<input type="checkbox"/>

storage_dimension		
Column		
tk	INT	<input checked="" type="checkbox"/>
name	VARCHAR(255)	<input checked="" type="checkbox"/>
identifier	VARCHAR(768)	<input checked="" type="checkbox"/>
ip	VARCHAR(255)	<input checked="" type="checkbox"/>
model	VARCHAR(255)	<input checked="" type="checkbox"/>
manufacturer	VARCHAR(255)	<input checked="" type="checkbox"/>
serialNumber	VARCHAR(255)	<input type="checkbox"/>
microcodeVersion	VARCHAR(255)	<input type="checkbox"/>
family	VARCHAR(255)	<input checked="" type="checkbox"/>
id	INT	<input type="checkbox"/>
latest	TINYINT(1)	<input type="checkbox"/>
dateTk	INT	<input checked="" type="checkbox"/>
dataCenter	VARCHAR(255)	<input type="checkbox"/>
url	VARCHAR(255)	<input type="checkbox"/>

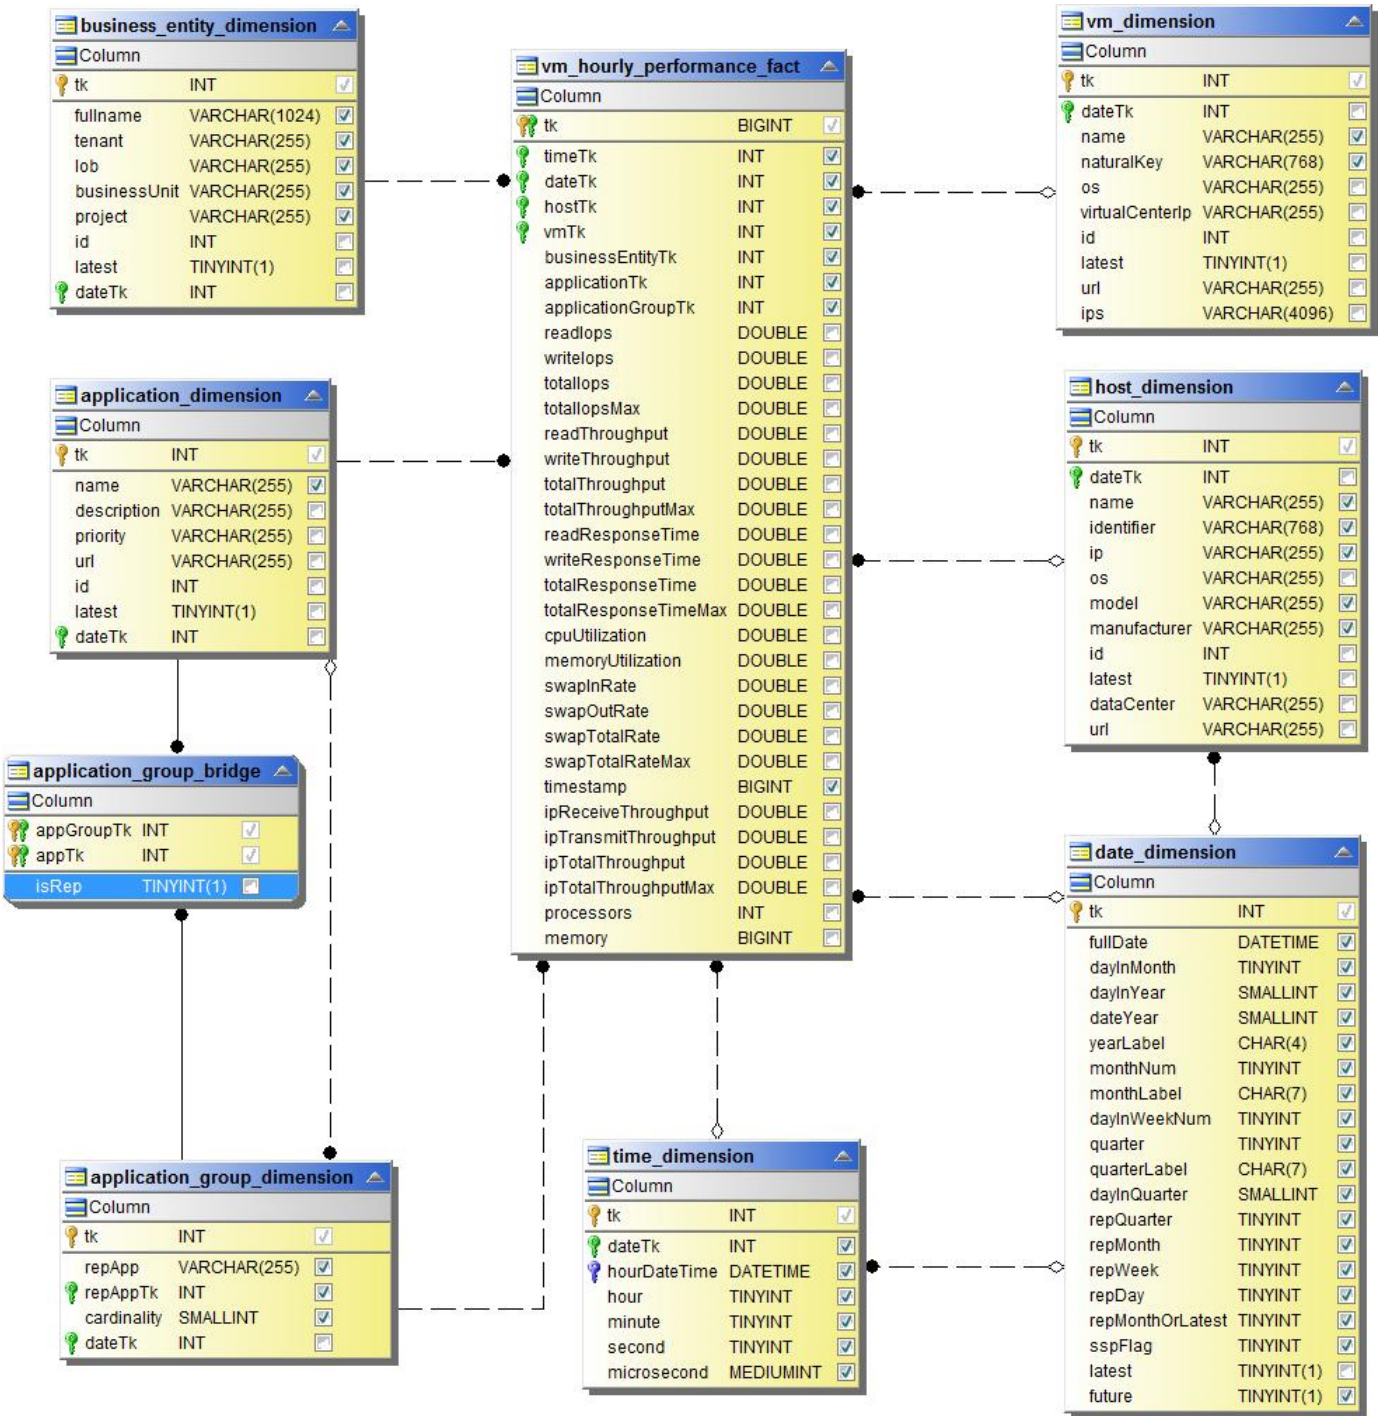
date_dimension		
Column		
tk	INT	<input checked="" type="checkbox"/>
fullDate	DATETIME	<input checked="" type="checkbox"/>
dayInMonth	TINYINT	<input checked="" type="checkbox"/>
dayInYear	SMALLINT	<input checked="" type="checkbox"/>
dateYear	SMALLINT	<input checked="" type="checkbox"/>
yearLabel	CHAR(4)	<input checked="" type="checkbox"/>
monthNum	TINYINT	<input checked="" type="checkbox"/>
monthLabel	CHAR(7)	<input checked="" type="checkbox"/>
dayInWeekNum	TINYINT	<input checked="" type="checkbox"/>
quarter	TINYINT	<input checked="" type="checkbox"/>
quarterLabel	CHAR(7)	<input checked="" type="checkbox"/>
dayInQuarter	SMALLINT	<input checked="" type="checkbox"/>
repQuarter	TINYINT	<input checked="" type="checkbox"/>
repMonth	TINYINT	<input checked="" type="checkbox"/>
repWeek	TINYINT	<input checked="" type="checkbox"/>
repDay	TINYINT	<input checked="" type="checkbox"/>
repMonthOrLatest	TINYINT	<input checked="" type="checkbox"/>
latest	TINYINT(1)	<input type="checkbox"/>
future	TINYINT(1)	<input checked="" type="checkbox"/>
sspFlag	TINYINT	<input checked="" type="checkbox"/>

time_dimension		
Column		
tk	INT	<input checked="" type="checkbox"/>
hourDateTime	DATETIME	<input checked="" type="checkbox"/>
hour	TINYINT	<input checked="" type="checkbox"/>
minute	TINYINT	<input checked="" type="checkbox"/>
second	TINYINT	<input checked="" type="checkbox"/>
microsecond	MEDIUMINT	<input checked="" type="checkbox"/>
dateTk	INT	<input checked="" type="checkbox"/>

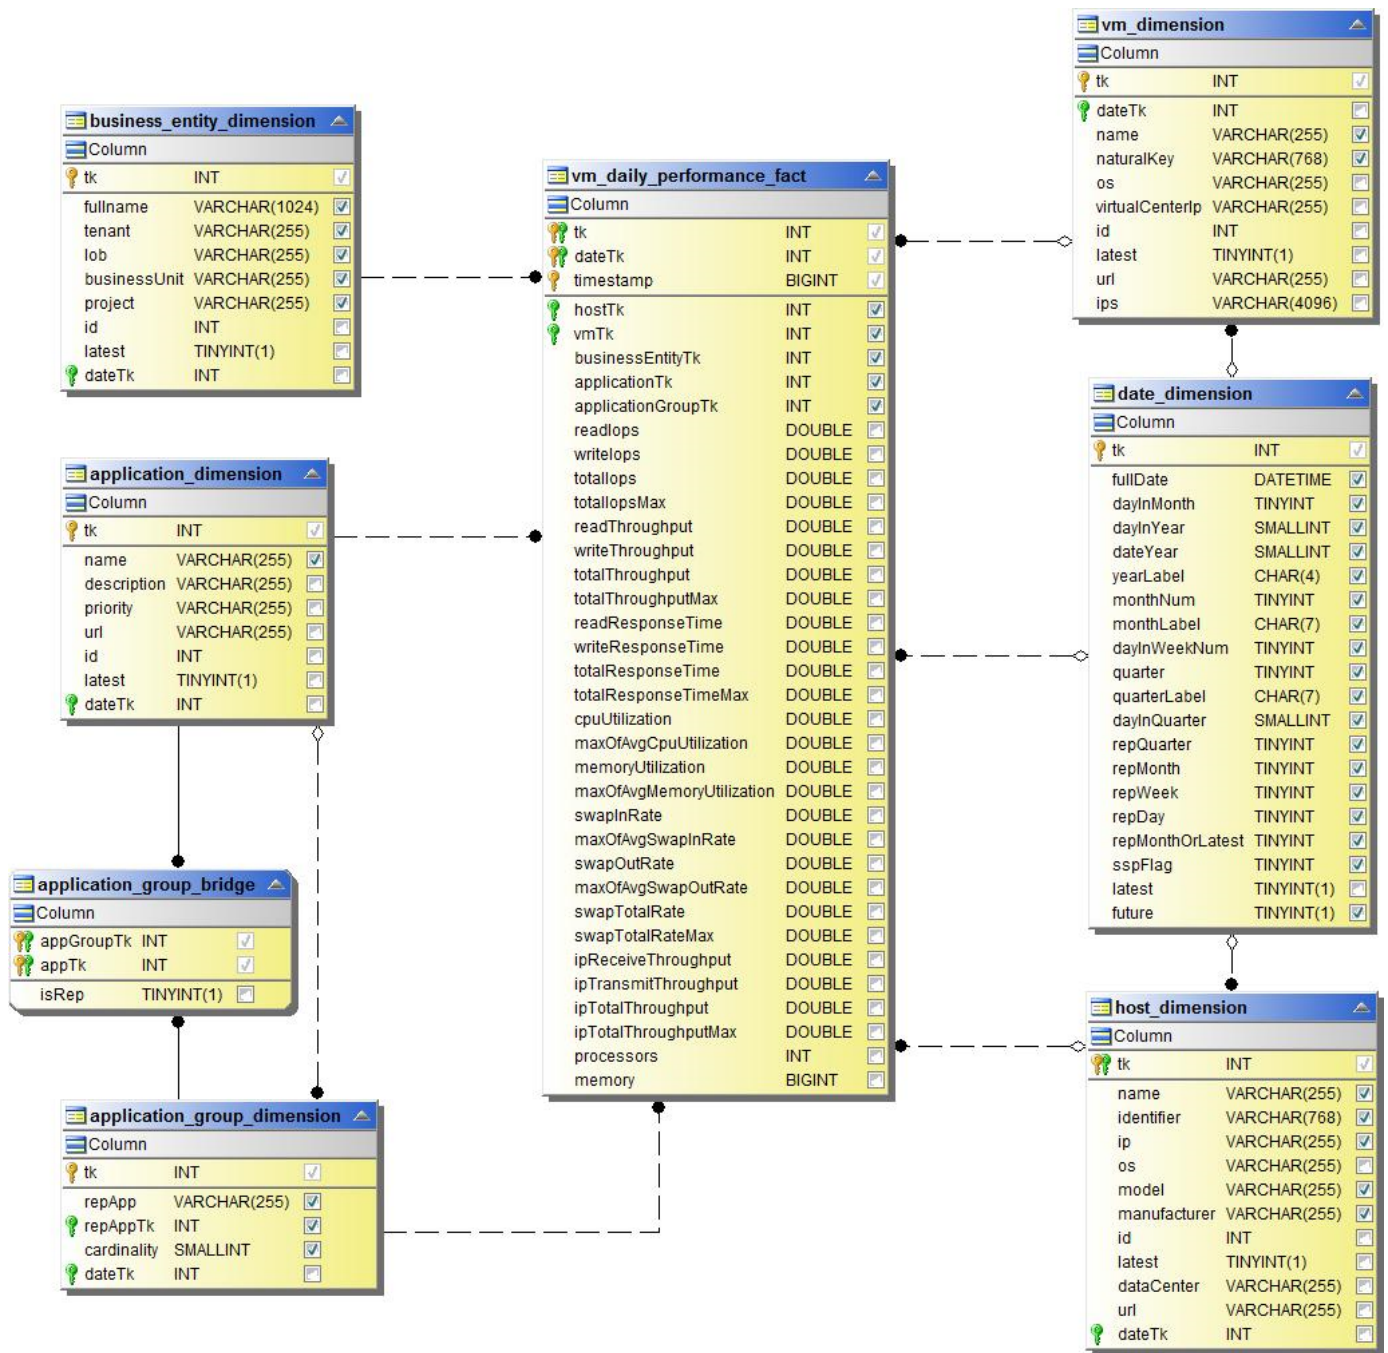
Switch Hourly Performance for Tape



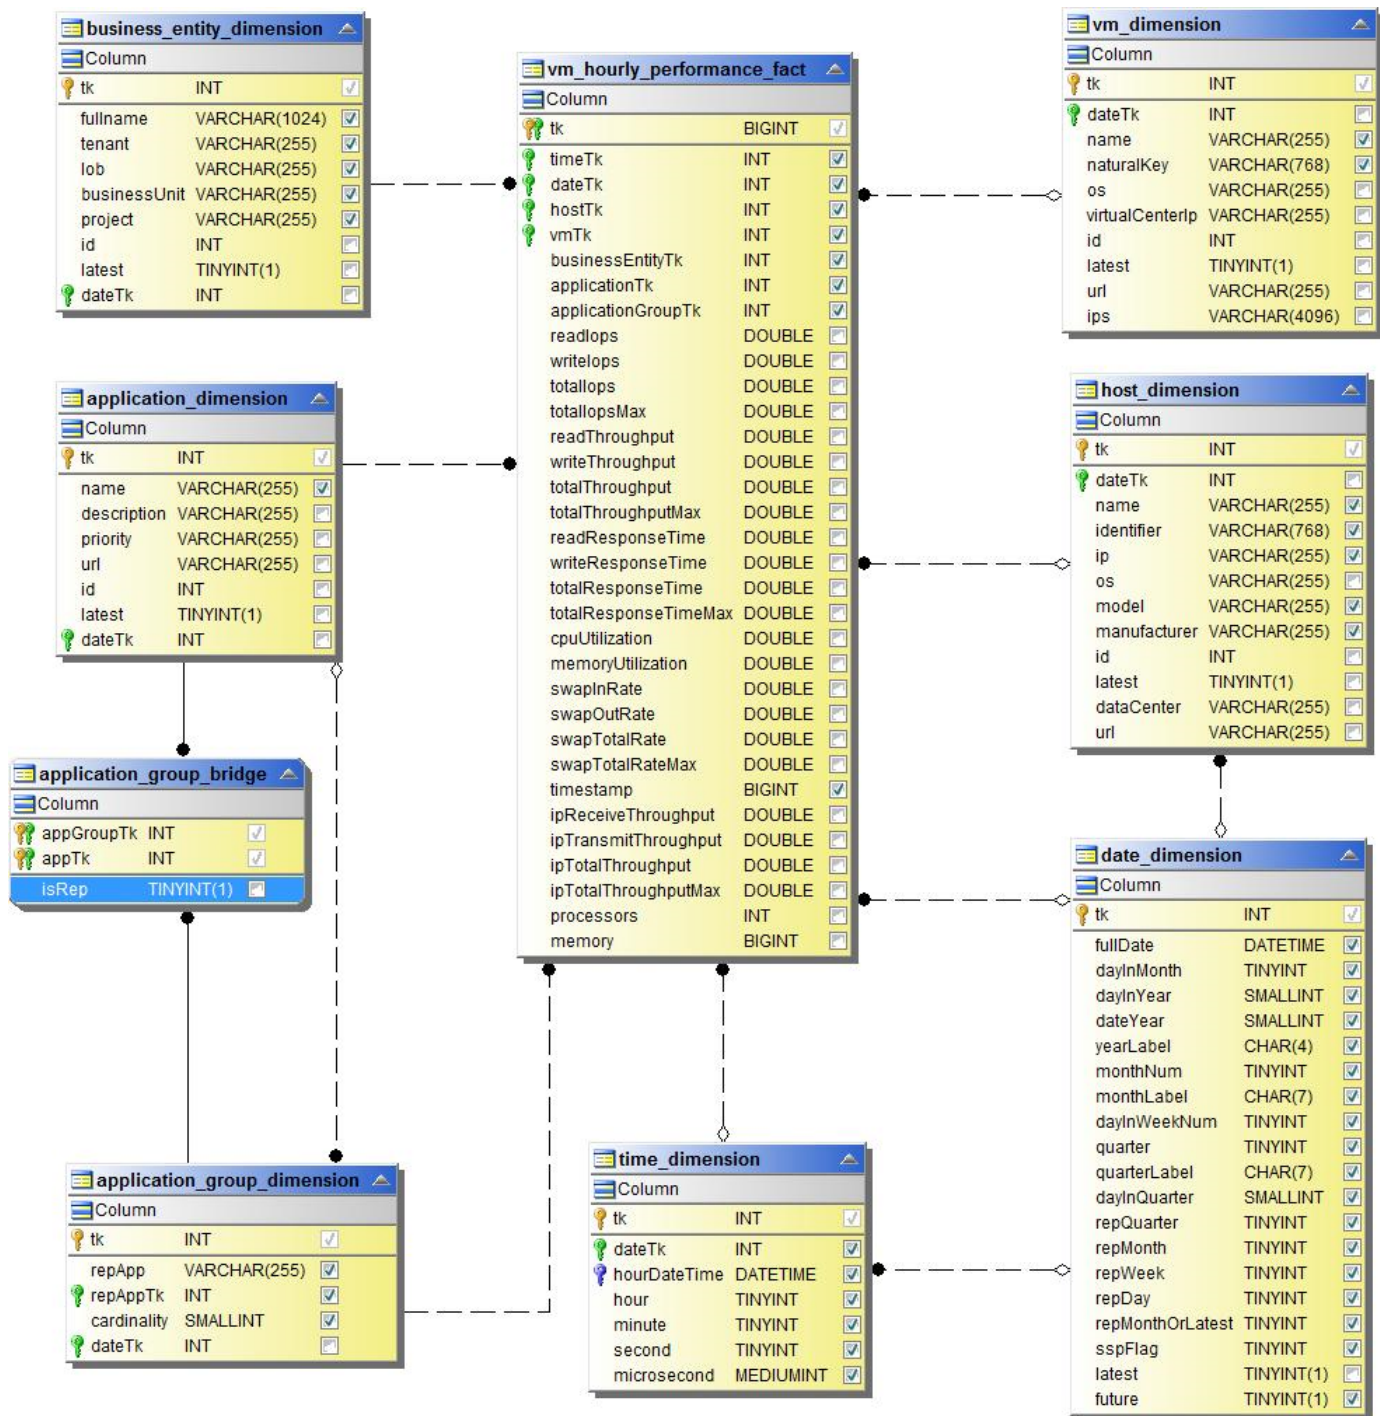
VM Performance



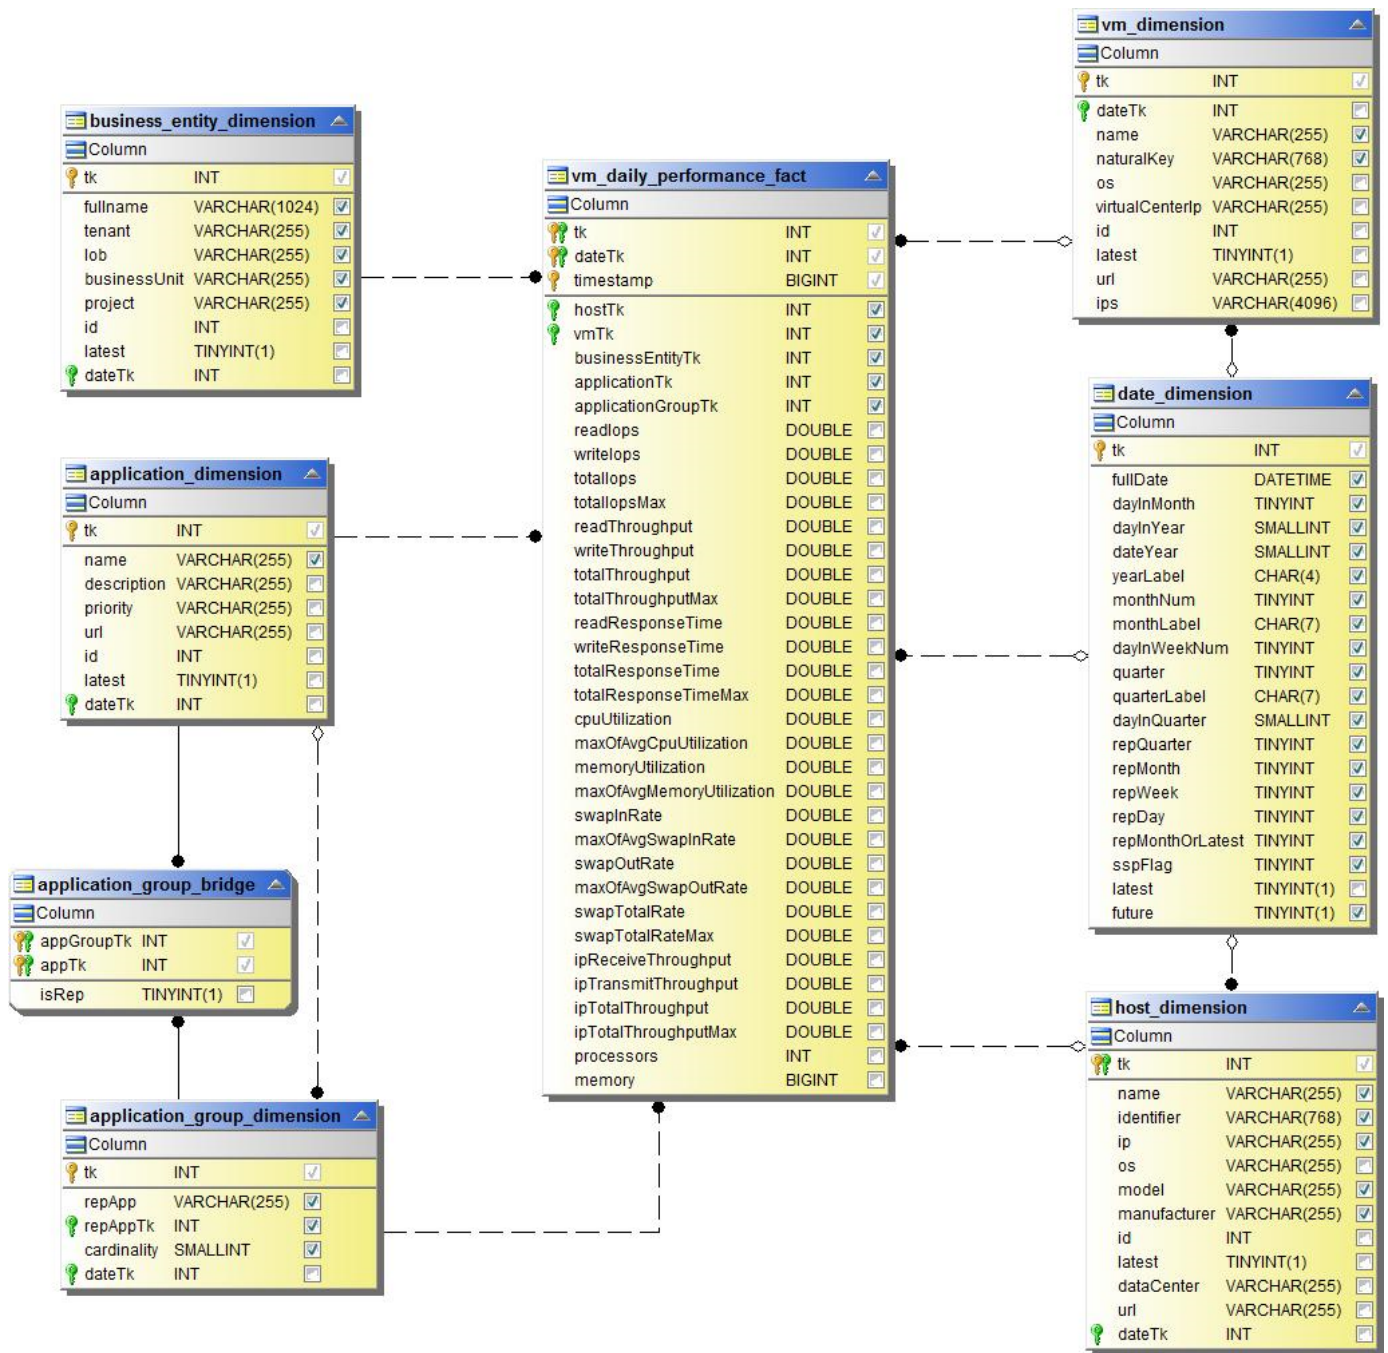
VM Daily Performance for Host



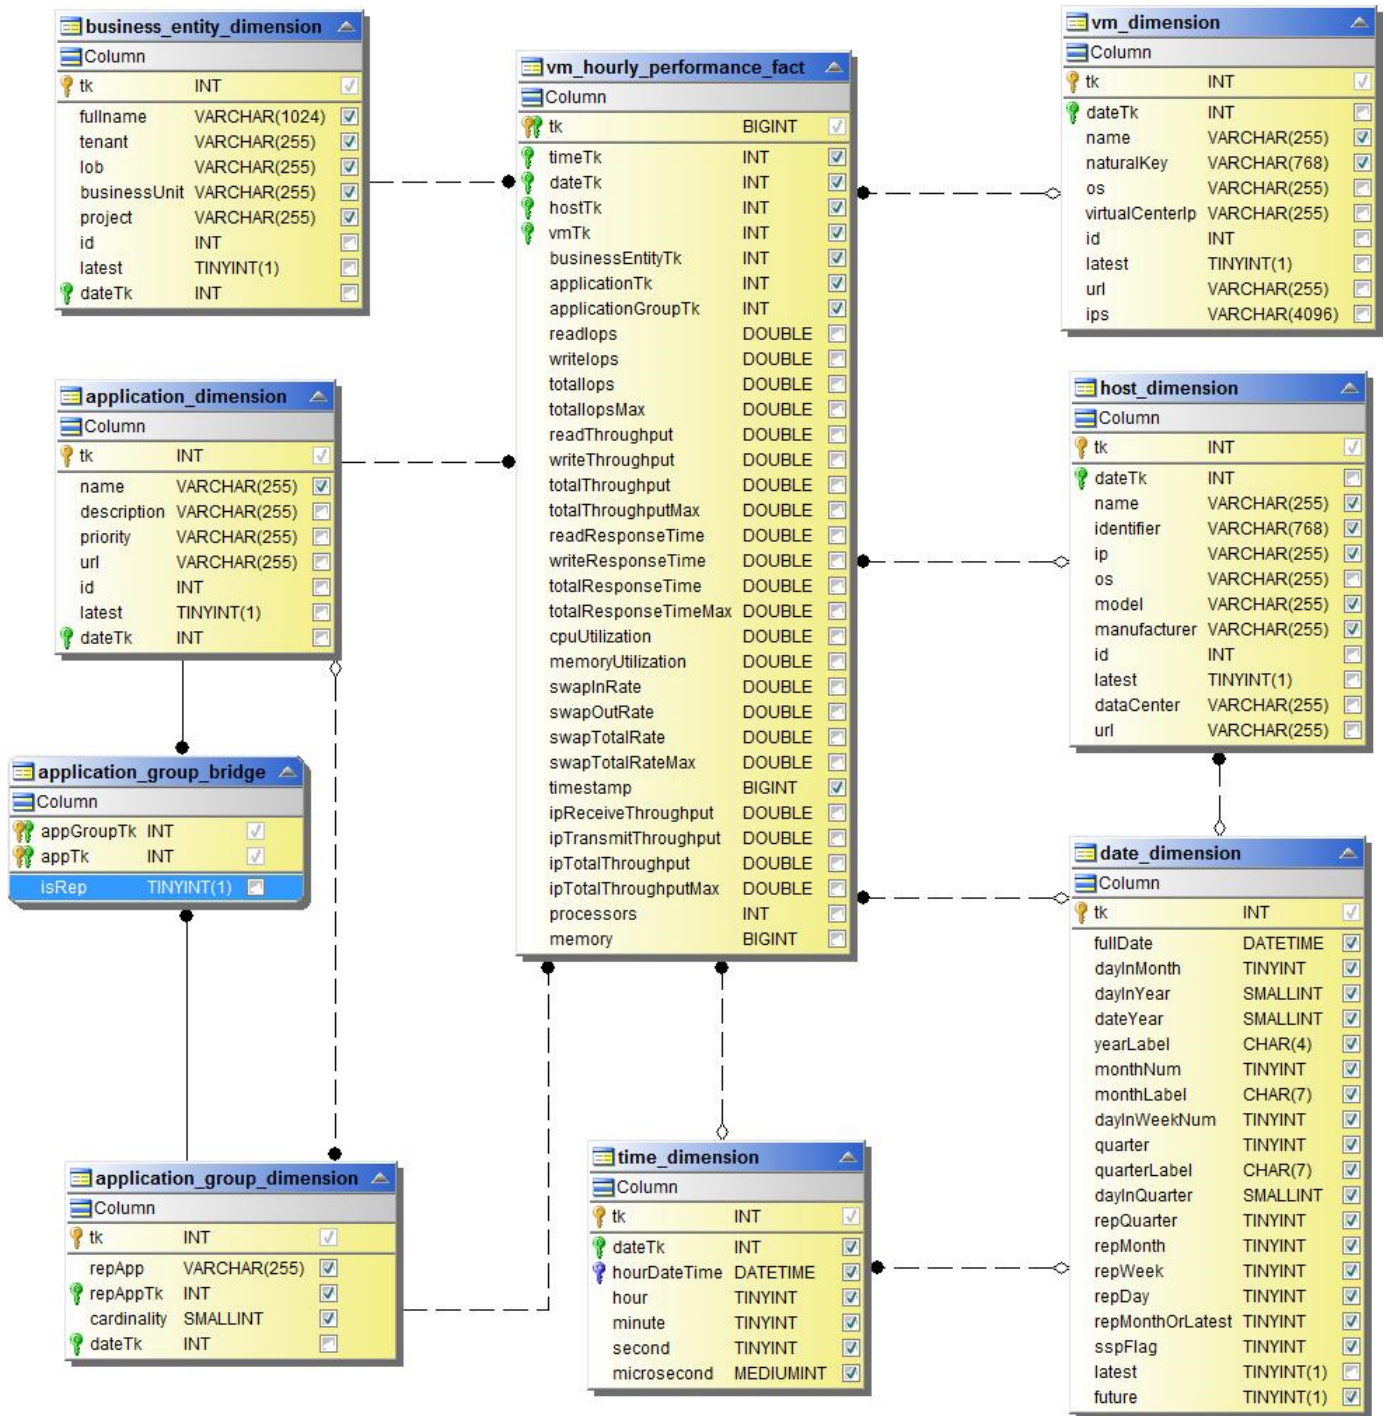
VM Hourly Performance for Host



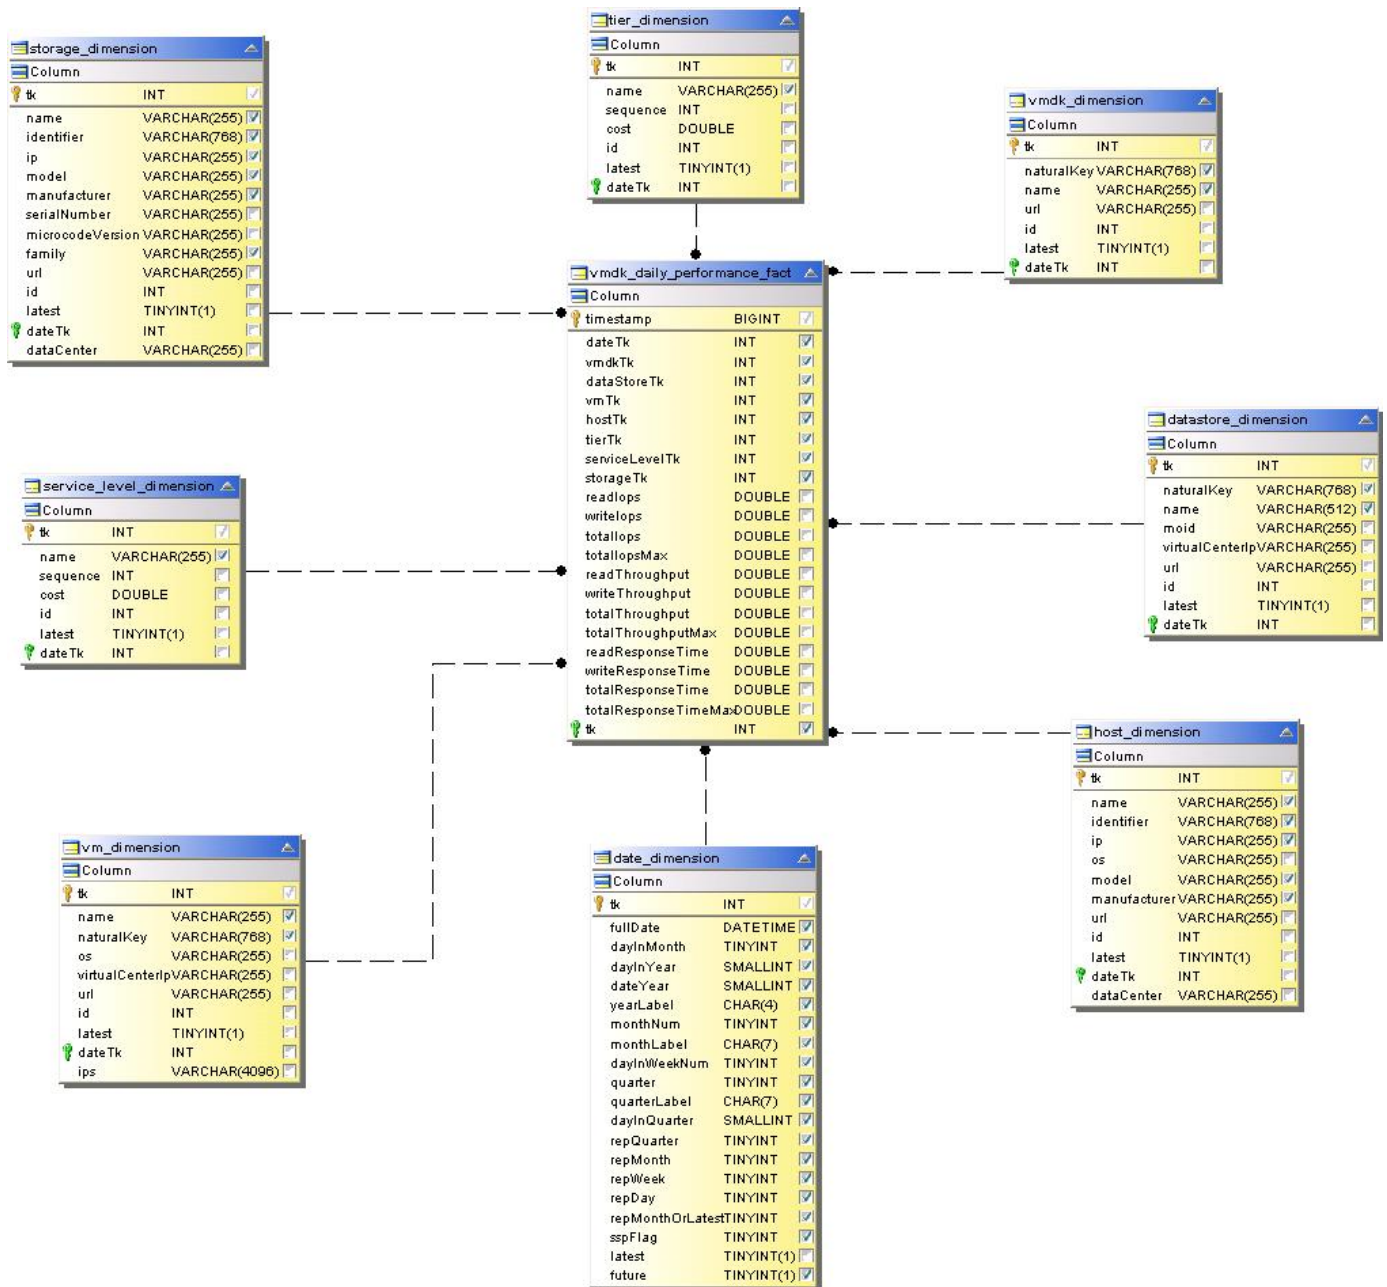
VM Daily Performance for Host




VM Hourly Performance for Host




VMDK Daily Performance




VMDK Hourly Performance



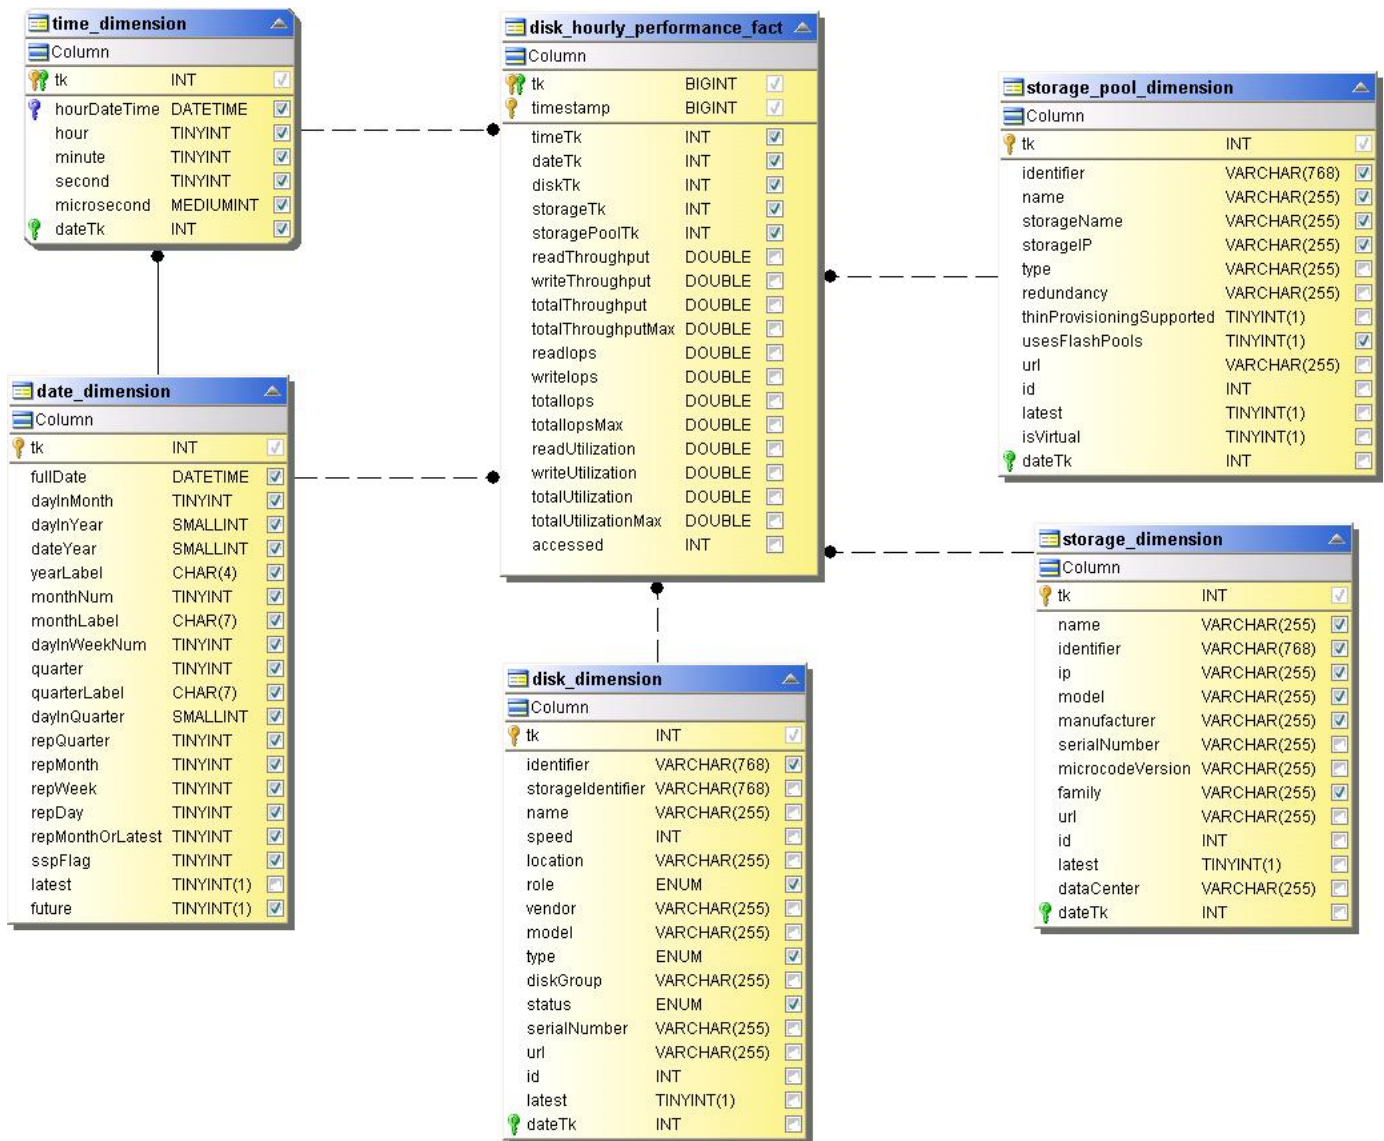
Storage Node Hourly Performance



Disk Daily Performance



Disk Hourly Performance



Cloud Insights Schemas for Reporting

These schema tables and diagrams are provided here as a reference for Cloud Insights Reporting.

[Schema Tables](#) in .PDF format. Click the link to open, or right-click and choose **Save as...** to download.

Schema Diagrams



The Reporting feature is available in Cloud Insights [Premium Edition](#).

Cloud Secure

About Cloud Secure

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.

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. Cloud Secure 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

Ensure corporate compliance by auditing user data access to your critical corporate data stored on-premise or in the cloud.

Getting Started

Getting Started with Cloud Secure

There are configuration tasks that need to be completed before you can start using Cloud Secure to monitor user activity.

The Cloud Secure 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:

Task	Related information
Configure an Agent	Agent Requirements Add Agent Video: Agent Deployment

Configure a User Directory Connector	Add User Directory Connector Video: Active Directory Connection
Configure data collectors	Click Admin > Data Collectors Click the data collector you want to configure. See the Data Collector Vendor Reference section of the documentation. Video: ONTAP SVM Connection
Create Users Accounts	Manage User Accounts
Troubleshooting	Video: Troubleshooting

Agent Requirements

You must [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.

Component	Linux Requirement
Operating system	<p>A computer running a licensed version of one of the following:</p> <ul style="list-style-type: none"> Red Hat Enterprise Linux 7.2 64-bit Red Hat Enterprise Linux 7.2 64-bit KVM Red Hat Enterprise Linux 7.5 64-bit Red Hat Enterprise Linux 7.5 64-bit KVM Red Hat Enterprise Linux 7.8 64-bit Red Hat Enterprise Linux 7.8 64-bit KVM CentOS 7.2 64-bit CentOS 7.2 64-bit KVM CentOS 7.5 64-bit CentOS 7.5 64-bit KVM CentOS 7.8 64-bit CentOS 7.8 64-bit KVM <p>This computer should be running no other application-level software. A dedicated server is recommended.</p>
Commands	The 'sudo su –' command is required for installation, running scripts, and uninstall.
CPU	4 CPU cores
Memory	16 GB RAM
Available disk space	Disk space should be allocated in this manner: /opt/netapp 25 GB
Network	100 Mbps to 1 Gbps Ethernet connection, static IP address, IP connectivity to all devices, and a required port to the Cloud Secure instance (80 or 443).

Please note: Cloud Insights agent and Cloud Secure agent can be installed in the same machine. However, it is a best practice to install them in separate machines. In the event that both agents are installed on the same machine, please allocate disk space as shown below:

Available disk space	50 GB For Linux, disk space should be allocated in this manner: /opt/netapp 25 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)**.

Cloud Network Access Rules

For **US-based** Cloud Secure environments:

Protocol	Port	Destination	Direction	Description
TCP	443	<site_name>.cs01.cloudinsights.netapp.com <site_name>.c01.cloudinsights.netapp.com <site_name>.c02.cloudinsights.netapp.com	Outbound	Access to Cloud Insights
TCP	443	gateway.c01.cloudinsights.netapp.com agentlogin.cs01.cloudinsights.netapp.com	Outbound	Access to authentication services

For **Europe-based** Cloud Secure environments:

Protocol	Port	Destination	Direction	Description
TCP	443	<site_name>.cs01-eu-1.cloudinsights.netapp.com <site_name>.c01-eu-1.cloudinsights.netapp.com <site_name>.c02-eu-1.cloudinsights.netapp.com	Outbound	Access to Cloud Insights

Protocol	Port	Destination	Direction	Description
TCP	443	gateway.c01.cloudinsights.netapp.com agentlogin.cs01-eu-1.cloudinsights.netapp.com	Outbound	Access to authentication services

In-network rules

Protocol	Port	Destination	Direction	Description
TCP	389(LDAP) 636 (LDAPs / start-tls)	LDAP Server URL	Outbound	Connect to LDAP
TCP	443	SVM Management IP Address	Outbound	API communication with ONTAP
TCP	35000 - 55000	SVM data LIF IP Addresses	Inbound/Outbound	Communication with ONTAP for Fpolicy events

Related:

See the [Event Rate Checker](#) documentation for information about sizing.

Cloud Secure Agent Installation

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 Cloud Secure SaaS layer for analysis. See [Agent Requirements](#) to configure an agent VM.

Before You Begin

- The sudo privilege is required for installation, running scripts, and uninstall.

Steps to Install Agent

1. Log in as Administrator or Account Owner to your Cloud Secure environment.
2. Click **Admin > Data Collectors > Agents > +Agent**

The system displays the Add an Agent page:

Add an Agent



Cloud Secure collects device and user data using one or more Agents installed on local servers. Each Agent can host multiple Data Collectors, which send data to Cloud Secure for analysis.

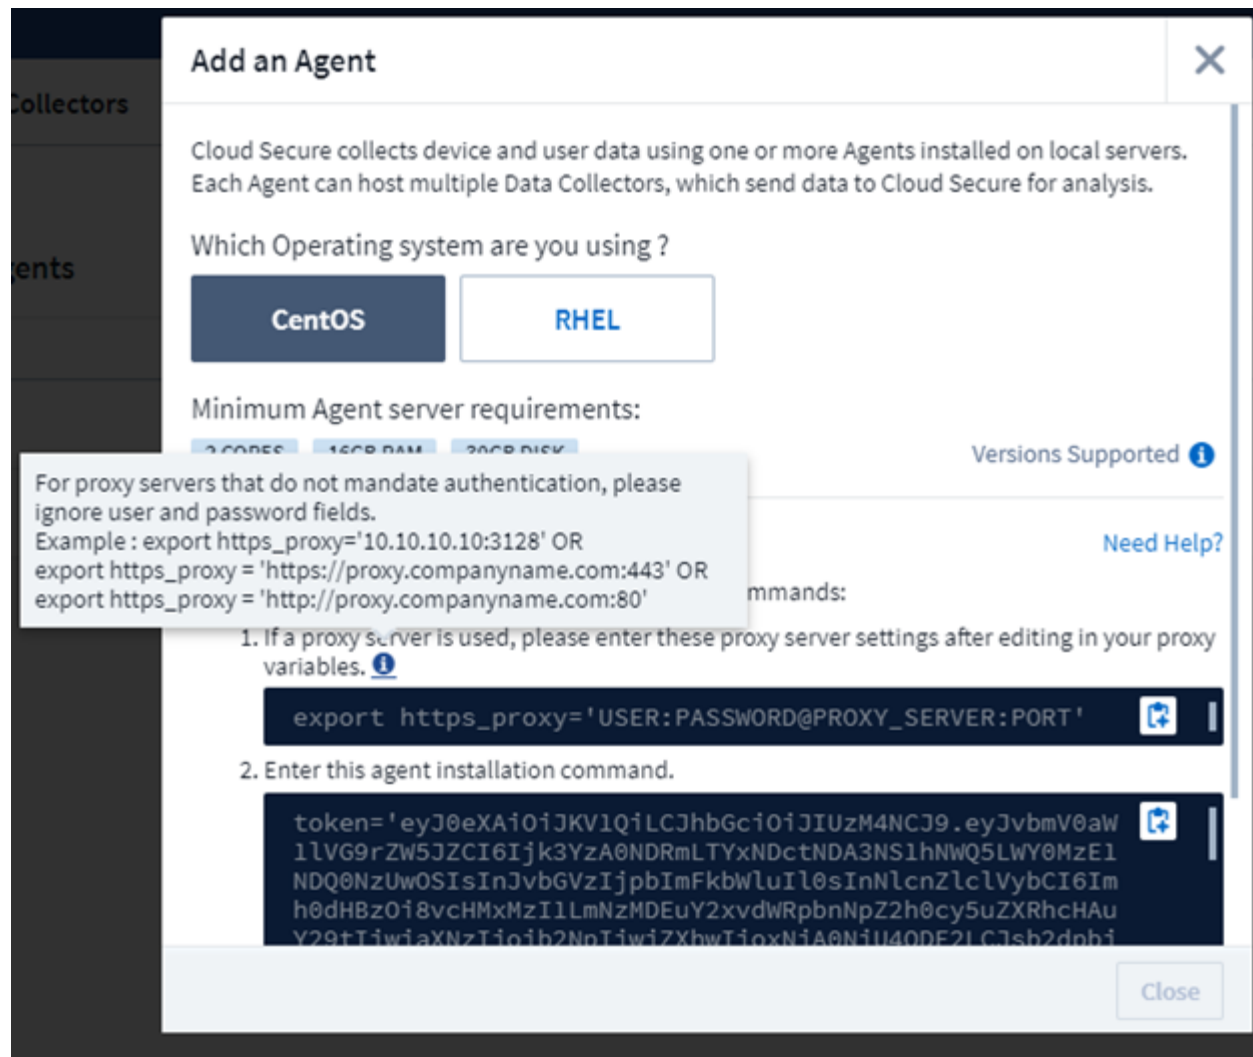
Which Operating system are you using ?

CentOS


RHEL

Close

3. Select the operating system on which you are installing the agent.
4. Verify that the agent server meets the minimum system requirements.
5. To verify that the agent server is running a supported version of Linux, click *Versions Supported (i)*.
6. If your network is using proxy server, please set the proxy server details by following the instructions in the Proxy section.



7. Click the Copy to Clipboard icon to copy the installation command.
8. Run the installation command in a terminal window.
9. The system displays the following message when the installation completes successfully:



After You Finish

1. You need to configure a [User Directory Collector](#) .
2. You need to configure one or more Data Collectors.

Network Configuration

Run the following commands on the local system to open ports that will be used by Cloud Secure.

Steps

1. `sudo firewall-cmd --permanent --zone=public --add-port=35000-55000/tcp`
2. `sudo firewall-cmd --reload`

3. `sudo iptables-save | grep 35000`

sample output:

```
-A IN_public_allow -p tcp -m tcp --dport 35000 -m conntrack -ctstate NEW -j ACCEPT
```

Troubleshooting Agent Errors

Known problems and their resolutions are described in the following table.

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 <code>journalctl -u cloudsecure-agent.service</code> command. This command can be used for troubleshooting the issue further.
Agent installation fails with 'This linux distribution is not supported. Exiting the installation'.	The supported platforms for Cloud Secure 1.0.0 are RHEL 7.x / CentOS 7.x. Ensure that you are not installing the agent on a RHEL 6.x or CentOS 6.x system.
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.

Problem:	Resolution:
Agent was installed and was running. However agent has stopped suddenly.	<p>SSH to the Agent machine. Check the status of the agent service via <code>sudo systemctl status cloudsecure-agent.service</code>.</p> <ol style="list-style-type: none"> 1. Check if the logs shows a message“Failed to start Cloud Secure 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. <pre>sudo id cssys sudo groups cssys</pre> <ol style="list-style-type: none"> 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. <pre>sudo useradd cssys sudo groupadd cssys</pre> <ol style="list-style-type: none"> 5. Restart the agent service after that by executing the following command: <pre>sudo systemctl restart cloudsecure-agent.service</pre> <ol style="list-style-type: none"> 6. If it is still not running, please check the other troubleshooting options.
Unable to add more than 10 Data collectors to an Agent.	Only 10 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.	<p>Steps to restart the Agent.</p> <ol style="list-style-type: none"> 1. SSH to the Agent machine. 2. Restart the agent service after that by executing the following command: <pre>sudo systemctl restart cloudsecure-agent.service</pre> <ol style="list-style-type: none"> 3. Check the status of the agent service via <code>sudo systemctl status cloudsecure-agent.service</code>. 4. Agent should go to CONNECTED state.

Deleting a Cloud Secure Agent

When you delete a Cloud Secure Agent, all the data collectors associated with the Agent must be deleted first.

Deleting an Agent



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

1. Make sure all the data collectors associated with the agent are deleted from the Cloud Secure portal.

Note: Ignore this step if all the associated collectors are in STOPPED state.

Steps to delete an Agent:

1. 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]:
```

2. Click **Admin > Data Collectors > Agents**

The system displays the list of configured Agents.

3. Click the options menu for the Agent you are deleting.
4. Click **Delete**.

The system displays the **Delete Agent** page.

5. Click **Delete** to confirm the deletion.

Configuring an Active Directory (AD) User Directory Collector

Cloud Secure 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 this task.
- 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

1. In the Cloud Secure menu, click:
Admin > Data 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:

Name	Description
Name	Unique name for the user directory. For example <i>GlobalADCollector</i>
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	<p>Forest level of the directory structure. Forest name allows both of the following formats:</p> <p><i>x.y.z</i> ⇒ direct domain name as you have it on your SVM. [Example: <i>hq.companyname.com</i>]</p> <p><i>DC=x,DC=y,DC=z</i> ⇒ Relative distinguished names [Example: <i>DC=hq,DC= companyname,DC=com</i>]</p> <p>Or you can specify as the following:</p> <p><i>OU=engineering,DC=hq,DC= companyname,DC=com</i> [to filter by specific OU engineering]</p> <p><i>CN=username,OU=engineering,DC=companyname,DC=netapp, DC=com</i> [to get only specific user with <username> from OU <engineering>]</p> <p><i>CN=Acrobat Users,CN=Users,DC=hq,DC=companyname,DC=com ,O= companyname,L=Boston,S=MA,C=US</i> [to get all Acrobat Users within the Users in that organization]</p>
Bind DN	User permitted to search the directory. For example: <i>username@companyname.com</i> or <i>username@domainname.com</i>
BIND password	Directory server password (i.e. password for username used in Bind DN)
Protocol	ldap, 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:///

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.

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:

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

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 Cloud Secure 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.

Configuration

Configuration | **Server** | **Connection** | Option | SSL/TLS

User DN: ☐ Anonymous login

Password: ☒ Store password

Use SSL port: ☐ Yes ☒ No

Use TLS: ☐ Yes ☒ No (TLS is only used on non SSL ports)

Base DN:

Troubleshooting LDAP Directory Collector Configuration Errors

The following table describes known problems and resolutions that can occur during collector configuration:

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 Cloud Secure 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 <i>Restart</i> button.

Problem:	Resolution:
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.
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.

Problem:	Resolution:
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.	<p>This is most likely due to an attribute mapping problem with the Active Directory.</p> <ol style="list-style-type: none"> 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 Cloud Secure User Directory Collector can not connect to the AD Server.	<p>Disable AD Server encryption before Configuring a User Directory Collector.</p> <p>Once the user detail is fetched it will be there for 13 months.</p> <p>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.</p>

Configuring the ONTAP SVM Data Collector

Cloud Secure 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
 - SMB protocol version 3.1 and earlier
 - NFS protocol version 4.0 and earlier
- Only data type SVMs are supported. SVMs with infinite/flexgroup volumes are not supported

- SVM has several sub-types. Of these, only *default* and *sync_source* are supported.
- An Agent [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 must add an SVM using one of the following two methods:
 - By Using Cluster IP, SVM name, and Cluster Management Username and Password
 - 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 [“A note about permissions”](#) section below. This custom user can be created for either SVM or Cluster access.
 - 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 [ONTAP documentation](#).
- Ensure the correct applications are set for the SVM by executing the following command:

```
clustershell::> security login show -vserver <vservename> -user-or
-group-name <username>
```

Example output:

```
Vserver: svmname
```

User/Group Name	Application	Authentication Method	Role Name	Acct Locked	Second Authentication Method
vsadmin	http	password	vsadmin	no	none
vsadmin	ontapi	password	vsadmin	no	none
vsadmin	ssh	password	vsadmin	no	none

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 [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.

A Note About Permissions

Permissions when adding via Cluster Management IP:

If you cannot use the Cluster management administrator user to allow Cloud Secure 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 Cloud Secure 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 none
security login role create -role csrole -cmddirname "network interface"
-access readonly
security login role create -role csrole -cmddirname version -access
readonly
security login role create -role csrole -cmddirname volume -access
readonly
security login role create -role csrole -cmddirname vserver -access
readonly
security login role create -role csrole -cmddirname "vserver fpolicy"
-access all
security login role create -role csrole -cmddirname "volume snapshot"
-access all
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 when adding via Vserver Management IP:

If you cannot use the Cluster management administrator user to allow Cloud Secure 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 Cloud Secure data collector to use Vserver 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. For ease, copy these commands to a text editor and replace the <vservename> with your Vserver name before and executing these commands on ONTAP:

```

security login role create -vserver <vservname> -role csrole -cmddirname
DEFAULT -access none
security login role create -vserver <vservname> -role csrole -cmddirname
"network interface" -access readonly
security login role create -vserver <vservname> -role csrole -cmddirname
version -access readonly
security login role create -vserver <vservname> -role csrole -cmddirname
volume -access readonly
security login role create -vserver <vservname> -role csrole -cmddirname
vserver -access readonly
security login role create -vserver <vservname> -role csrole -cmddirname
"vserver fpolicy" -access all
security login role create -vserver <vservname> -role csrole -cmddirname
"volume snapshot" -access all
security login create -user-or-group-name csuser -application ontapi
-authmethod password -role csrole -vserver <vservname>

```

Configure the data collector

Steps for Configuration

1. Log in as Administrator or Account Owner to your Cloud Insights environment.
2. Click **Admin > Data Collectors > +Data Collectors**

The system displays the available Data Collectors.

3. Hover over the **NetApp SVM** tile and click ***+Monitor**.

The system displays the ONTAP SVM configuration page. Enter the required data for each field.

Configuration

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 IP
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.

After you finish


- 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.

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.

Installed Data Collectors

Name	Status	Type	Agent
9.8_vs1	 Error more detail	ONTAP SVM	agent-11

Problem:	Resolution:
<p>Error message: "Connection to the FPolicy server <IP> is broken. (reason: "FPolicy server is removed from external engine.")")"</p>	<p>SVM is unable to reach the Fpolicy Server.</p> <ol style="list-style-type: none"> 1. Make sure there is route available from SVM to the Fpolicy Server/Agent machine IP. Login to the cluster/SVM and ping the Fpolicy Server IP address using the following command: <pre>net ping -lif <data_lif> -destination <agent IP> -vserver <svmname> -show-detail</pre> <ol style="list-style-type: none"> 2. In instances where the same SVM was added in two different Cloud Secure 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.
<p>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."</p>	<p>The event rate from ONTAP was much higher than what the Agent box can handle. Hence the connection got terminated.</p> <p>Check the peak traffic in CloudSecure when the disconnection happened. This you can check from the CloudSecure > Activity Forensics > All Activity page.</p> <p>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.</p> <p>If the Agent was installed in the Agent box prior to 4 March 2021, run the following commands in the Agent box:</p> <pre>echo 'net.core.rmem_max=8388608' >> /etc/sysctl.conf echo 'net.ipv4.tcp_rmem = 4096 2097152 8388608' >> /etc/sysctl.conf sysctl -p</pre> <p>Restart the collector from the UI after resizing.</p>

Problem:	Resolution:
<p>Collector reports Error Message: “No local IP address found on the connector that can reach the data interfaces of the SVM”.</p>	<p>This is most likely due to a networking issue on the ONTAP side. Please follow these steps:</p> <ol style="list-style-type: none"> 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. <p>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:</p> <pre>network ping -vserver <vserver name> -destination <Agent IP> -lif <Lif Name> -show-detail</pre> <p>If not pingable, make sure the network settings in ONTAP are correct, so that the Agent machine is pingable.</p> <ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> 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 <pre>https://<cluster_mgmt_ip>/spi/<clustername>/etc/log/packet_traces/</pre> <ol style="list-style-type: none"> 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

Problem:	Resolution:
<p>Message: "Failed to determine ONTAP type for [hostname: <IP Address>. Reason: Connection error to Storage System <IP Address>: Host is unreachable (Host unreachable)"</p>	<ol style="list-style-type: none"> 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.
<p>Error Message: "Connector is in error state. Service.name: audit. Reason for failure: External fpolicy server terminated."</p>	<ol style="list-style-type: none"> 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. <p><i>sudo iptables-save grep 3500*</i></p> <p>Sample output should look like:</p> <p><i>-A IN_public_allow -p tcp -m tcp --dport 35000 -m conntrack -ctstate NEW -j ACCEPT</i></p> 3. Login to SVM, enter the following commands and check that no firewall is set to block the communication with ONTAP. <p><i>system services firewall show</i> <i>system services firewall policy show</i></p> <p>Check firewall commands on the ONTAP side.</p> 4. SSH to the SVM/Cluster which you want to monitor. Ping the Agent box from the SVM management lif (with CIFS, NFS protocols support) and ensure that ping is working: <p><i>network ping -vserver <vserver name> -destination <Agent IP> -lif <Lif Name> -show-detail</i></p> <p>If not pingable, make sure the network settings in ONTAP are correct, so that the Agent machine is pingable.</p>

Problem:	Resolution:
No events seen in activity page.	<p>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.</p> <p>2. If no activities are seen, please login to the SVM and enter the following command. <code><SVM>event log show -source fpolicy</code> Please ensure that there are no errors related to fpolicy.</p> <p>3. If no activities are seen, please login to the SVM. Enter the following command <code><SVM>fpolicy show</code> Please check if the fpolicy policy named with prefix “metadata_service” 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.</p>
SVM Data Collector is in error state and Error message is “Agent failed to connect to the collector”	<p>1. Most likely the Agent is overloaded and is unable to connect to the Data Source collectors.</p> <p>2. Check how many Data Source collectors are connected to the Agent.</p> <p>3. Also check the data flow rate in the “All Activity” page in the UI.</p> <p>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.</p>
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")"	<p>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:</p> <p>event log show -source fpolicy which shows the error event log show -source fpolicy -fields event,action,description which shows more details.</p> <p>Check firewall commands on the ONTAP side.</p>
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).

Problem:	Resolution:
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 svmname. Reason: Invalid value specified for 'shares-to-include' element within 'fpolicy.policy.scope-modify: "Federal"	The share names need to be given without any quotes. Edit the ONTAP SVM DSC configuration to correct the share names. <i>Include and exclude shares</i> 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.

If you are still experiencing problems, reach out to the support links mentioned in the **Help > Support** page.

Configuring the Cloud Volumes ONTAP Data Collector

Cloud Secure uses data collectors to collect file and user access data from devices.

Cloud Volumes ONTAP Storage Configuration

See the OnCommand Cloud Manager Documentation to configure a single-node / HA AWS instance to host the Cloud Secure Agent:

<https://docs.netapp.com/us-en/occm/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

Agent Machine Configuration

Use the following steps to configure the machine to be used as a Cloud Secure Agent:

Steps

1. Log in to the AWS console and navigate to EC2-Instances page and select *Launch instance*.
2. 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
3. Select the VPC and Subnet that the Cloud ONTAP instance resides in.
4. Select *t2.xlarge* (4 vcpus and 16 GB RAM) as allocated resources.
 - a. Create the EC2 instance.
5. Install the required Linux packages using the YUM package manager:
 - a. Install *wget* and *unzip* native Linux packages.

Install the Cloud Secure Agent

1. Log in as Administrator or Account Owner to your Cloud Insights environment.
2. Navigate to Cloud Secure **Admin > Data Collectors** and click the **Agents** tab.
3. Click **+Agent** and specify RHEL as the target platform.
4. Copy the Agent Installation command.
5. Paste the Agent Installation command into the RHEL EC2 instance you are logged in to.
This installs the Cloud Secure agent, providing all of the [Agent Prerequisites](#) are met.

For detailed steps please refer to this xref:

https://docs.netapp.com/us-en/cloudinsights/task_cs_add_agent.html#steps-to-install-agent

User Management

Cloud Secure user accounts are managed through Cloud Insights.

Cloud Insights provides four user accounts: Account Owner, Administrator, User, and Guest. Each account is assigned specific permission levels. A User account that has Administrator privileges can access Cloud Secure and can create or modify users.

Steps

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

You will be forwarded to Cloud Insights's User Management page.

More information on User accounts and roles can be found in the Cloud Insights [User Role](#) documentation.

SVM Event Rate Checker

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.

Requirements:

- Cluster IP
- Cluster admin username and password



When running this script no ONTAP SVM Data Collector should be running for the SVM for which event rate is being determined.

Steps:

1. Install the Agent by following the instructions in CloudSecure.
2. 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
```

3. Provide the correct values when prompted. See below for an example.
4. The script will take approximately 5 minutes to run.
5. 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".
```

1. 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) A single Agent machine can handle upto 7000 events per second and maximum of 10 data collectors.

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%.

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% < 7000 events/second
```

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 < 7000 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  
Qa-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 shails3 is generating 200 SMB events/sec and 100 NFS events/sec  
Overall svm shails3 is generating 300 events/sec
```

```
[root@ci-cs-data agent]#
```

Troubleshooting

Question: If I run this script on an SVM that is already configured for Cloud Secure, does it just use the existing fpolicy config on the SVM or does it setup a temporary one and run the process?
Answer: The Event Rate Checker can run fine even for an SVM already configured for Cloud Secure. There should be no impact.
Question: Can I increase the number of SVMs on which the script can be run?
Answer: Yes. Simply edit the script and change the max number of SVMs from 5 to any desirable number.
Question: If I increase the number of SVMs, will it increase the time of running of the script?
Answer: No. The script will run for a max of 5 minutes, even if the number of SVMs is increased.

Alerts

The Cloud Secure Alerts page shows a timeline of recent attacks and/or warnings and allows you to view details for each issue.

Alerts

Last 3 Days

Filter By

Status

New

Date	Potential Attacks	Warning
24 Jul	1	0
26 Jul	1	0
30 Jul	1	0
1 Aug	1	0
3 Aug	4	0
4 Aug	3	0
5 Aug	5	0
6 Aug	1	0
7 Aug	6	1
8 Aug	3	0
9 Aug	1	0

Potential Attacks

(3)

Potential Attacks	Detected ↓	Status	User	Evidence	Action Taken
Ransomware Attack	5 hours ago Aug 10, 2020 4:38 AM	New	Iris McIntosh	> 700 Files Encrypted	Snapshots Taken
Ransomware Attack	a day ago Aug 9, 2020 3:51 AM	New	Christy Santos	> 500 Files Encrypted	Snapshots Taken
Ransomware Attack	2 days ago Aug 8, 2020 4:29 AM	New	Safwan Langley	> 700 Files Encrypted	Snapshots Taken

Warnings

(7)

Abnormal Behaviour	Detected ↓	Status	User	Change	Action Taken
User Activity Rate	2 days ago Aug 8, 2020 7:49 PM	New	Iris McIntosh	↑ 192.46%	None
User Activity Rate	2 days ago Aug 8, 2020 7:32 PM	New	Jenny Bryan	↑ 73.64%	None
User Activity Rate	3 days ago Aug 7, 2020 8:07 PM	New	Szymon Owen	↑ 189.88%	None

Alert

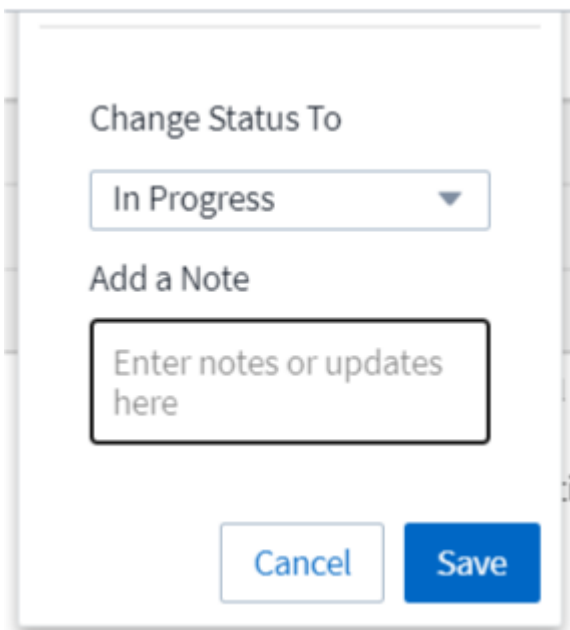
The Alert list displays a graph showing the total number of Potential Attacks and/or Warnings that have been raised in the selected time range, followed by a list of the attacks and/or warnings that occurred in that time range. You can change the time range by adjusting the start time and end time sliders in the graph.

The following is displayed for each alert:

Potential Attacks:

- The *Potential Attack* type (for example, Ransomware)
- The date and time the potential attack was *Detected*
- The *Status* of the alert:
 - New (this is the default for new alerts)
 - In Progress
 - Resolved
 - Dismissed

An administrator can change the status of the alert and add a note to assist with investigation.



The image shows a modal dialog box with a light gray border. At the top, it says "Change Status To". Below this is a dropdown menu with "In Progress" selected and a downward arrow. Underneath is a section titled "Add a Note" followed by a text input area with the placeholder text "Enter notes or updates here". At the bottom of the dialog are two buttons: "Cancel" in a light blue box and "Save" in a dark blue box.

- 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 (this is the default for new alerts)

- In Progress
- Resolved
- Dismissed

An administrator can change the status of the alert and add a note to assist with investigation.

- 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 (in this example, Ransomware) and Alert ID (assigned by Cloud Secure)
- 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 showing potential ransomware attack:



Filter By



Potential Attacks (1)

Potential Attacks	Detected ↓	Status	User	Evidence	Action Taken
Ransomware Attack	5 days ago Jul 11, 2020 4:02 AM	New	Kristjan Egilsson	> 700 Files Encrypted	None

Warnings (0)

Abnormal Behaviour	Detected ↓	Status	User	Change	Action Taken
No Data Available					

Detail page for potential ransomware attack:



POTENTIAL ATTACK: AL_305
Ransomware Attack

Detected
5 days ago
Jul 11, 2020 4:02 AM

Action Taken
None

Status
New

Total Attack Results

1	0	4173
Affected Volumes	Deleted Files	Encrypted Files

4173 Files have been copied, deleted, and potentially encrypted by 1 user account.

This is potentially a sign of ransomware attack.
The extension "crypt" was added to each file.

Encrypted Files

Activity per minute



Related Users



Kristjan Egilsson
Accountant
Finance

4173
Encrypted Files

Detected
5 days ago
Jul 11, 2020 4:02 AM

Action Taken
None



Username
us035

Email
Egilsson@netapp.com

Phone
387224312607


Department
Finance

Manager
Lyndsey Maddox

Top Activity Types

Activity per minute
Last access location: 10.197.144.115

[View Activity Detail](#)



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:

Potential Attack Detail / Ransomware Attack

POTENTIAL ATTACK: AL_307
Ransomware Attack

Detected
4 days ago
Jul 26, 2020 3:38 AM

Action Taken
Snapshots Taken

Status
In Progress

Last snapshots taken by
Amit Schwartz
Jul 30, 2020 2:54 PM

How To:
[Restore Entities](#)

[Re-Take Snapshots](#)

Total Attack Results

1
Affected Volumes

0
Deleted Files

5148
Encrypted Files

5148 Files have been copied, deleted, and potentially encrypted by 1 user account.

This is potentially a sign of ransomware attack.
The extension "crypt" was added to each file.

Encrypted Files
Activity per minute

Related Users

Ewen Hall
Developer
Engineering

5148
Encrypted Files

Detected
4 days ago
Jul 26, 2020 3:38 AM

Action Taken
Snapshots Taken

Manual Snapshot:

Cloud Insights

Abhi Basu Thakur

MONITOR & OPTIMIZE

Alerts / **Nabilah Howell** had an abnormal change in activity rate

Jul 23, 2020 - Jul 26, 2020
1:44 AM 1:44 AM

CLOUD SECURE

ALERTS

FORENSICS

ADMIN

HELP

WARNING: AL_306
Nabilah Howell had an abnormal change in activity rate.

Detected
5 days ago
Jul 25, 2020 1:44 PM

Action Taken
None

Status
New

Recommendation: Setup an Automated Response Policy
An Automated Response Policy will trigger measures to contain the damage automatically when a future attack is detected. Try it now.

[Take Snapshots](#)

How To:
[Restore Entities](#)

Nabilah Howell's Activity Rate Change

Typical
122.8
Activities Per Minute

Alert
210
Activities Per Minute

↑ 71%

Nabilah Howell's activity rate grew 71% over their typical average.

Activity Rate
Activity per 5 minutes

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 Cloud Secure environment is deleted, all data associated with the environment is also deleted.

Troubleshooting

Problem:	Try This:
For snapshots taken by Cloud Secure (CS), is there a purging/archiving period for CS snapshots?	No. There is no purging/archiving period set for CS snapshots. The user needs to define purging policy for CS snapshots. Please refer to the ONTAP documentation on how to setup the policies.
There is a situation where, ONTAP takes hourly snapshots per day. Will Cloud Secure (CS) snapshots affect it? Will CS snapshot take the hourly snapshot place? Will the default hourly snapshot get stopped?	Cloud Secure snapshots will not affect the hourly snapshots. CS 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?	<p>If the maximum Snapshot count is reached, subsequent Snapshot taking will fail and Cloud Secure will show an error message noting that Snapshot is full.</p> <p>User needs to define Snapshot policies to delete the oldest snapshots, otherwise snapshots will not be taken.</p> <p>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.</p> <p>See the ONTAP Documentation for information on setting Snapshot deletion policy.</p>
Cloud Secure is unable to take snapshots at all.	<p>Make sure that the role being used to create snapshots has xref: proper rights assigned.</p> <p>Make sure <i>csrole</i> is created with proper access rights for taking snapshots:</p> <pre>security login role create -vserver <vservename> -role csrole -cmddirname "volume snapshot" -access all</pre>
Snapshots are failing for older alerts on SVMs which were removed from Cloud Secure 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 <i>Alert Details</i> page, the message “Last attempt failed” error is seen below the <i>Take Snapshot</i> button. Hovering over the error displays “Invoke API command has timed out for the data collector with id”.	<p>This can happen when a data collector is added to Cloud Secure via SVM Management IP, if the LIF of the SVM is in <i>disabled</i> state in ONTAP.</p> <p>Enable the particular LIF in ONTAP and trigger <i>Take Snapshot manually</i> from Cloud Secure. The Snapshot action will then succeed.</p>

Forensics

Forensics - All Activity

The All Activity page helps you understand the actions performed on entities in the Cloud Secure 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  .

- 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 [User information](#).
- The **activity** the user performed. Supported types are:
 - **Change Group Ownership** - Group Ownership of file or folder is changed. For more details about group ownership please see [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 [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.

Filtering Forensic Activity History Data

There are two methods you can use to filter data.

1. 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.
2. Filter data by typing in the *Filter By* field:

Select the appropriate filter from the top 'Filter By' widget by clicking the **[+]** button:



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.
- **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.

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 10,000 records are exported.

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.



Activity History Retention

Activity history is retained for 13 months for active Cloud Secure environments.

Troubleshooting

Problem	Try This
---------	----------

<p>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” or “ldap:default:80038003”</p>	<p>Possible reasons could be:</p> <ol style="list-style-type: none"> 1. No User Directory Collectors have been configured yet. To add one, go to Admin > Data Collectors > User Directory Collectors and click on +User Directory Collector. Choose <i>Active Directory</i> or <i>LDAP Directory Server</i>. 2. A User Directory Collector has been configured, however it has stopped or is in error state. Please go to Admin > Data Collectors > User Directory Collectors and check the status. Refer to the 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.
--	---

Forensic Entities Page


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



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.



Forensic User Overview

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
- Operations for the last seven days
 - Number of operations

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 Cloud Secure environment is deleted, all data associated with the environment is deleted.

Automated Response Policies

Response Policies trigger actions such as taking a snapshot in the event of an attack or abnormal user behavior.

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.

Add Attack Policy

X

Policy Name*

Unique New Policy Name

For Ransomware Attacks

Currently Cloud Secure discovers and tracks possible Ransomware attacks.
Coming Soon: Tracking for additional attack types, including Identity Theft, Sabotage, and Snooping.

On Device

All Devices

+ Another Device

Actions

☒ Take Snapshot

Cancel

Save

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 [Alert detail page](#).

You can modify or pause an Automated Response Policy by choosing the option in the policy's drop-down menu.

Admin / Automated Response Policies

Attack Policies

+ Attack Policy

Filter...

Name	Alert Type	Device	Status	
RansomwareAttack	Ransomware Attack	svm_ci svm_ci2 demoGroupShares2	Active	<div>Edit Pause Delete</div>

Warning Policies

Configuring Alert Email Notifications

Email notifications of alerts are sent to the 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.

Admin / Notifications

Global Recipient List

Send Policy Alerts to the following email addresses

admin@company.com X

alert_list@company.com X

Save

Cloud Secure API

The Cloud Secure API enables NetApp customers and independent software vendors (ISVs) to integrate Cloud Secure with other applications, such as CMDB's or other ticketing systems.

Requirements for API Access:

- An API Access Token model is used to grant access.
- API Token management is performed by Cloud Secure users with the Administrator role.

API Documentation (Swagger)

The latest API information is found by logging in to Cloud Secure 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 Cloud Secure 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**



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**.







Active IQ

NetApp [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.



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).

Risks				
108 items found Filter...				
Object ↑	Risk Detail	Category	Potential Impact	Source
 tawny01	The following certificates have expired or are expiring within 30 days: Expired: 53CF9553, 53C504D4, 53D671B4, Expiring within 30 days: None	System Configuration	Clients may not be able to connect to the cluster over secure (SSL based) protocols.	 Active IQ ↗
 tawny01	None of the NIS servers configured for SVM(s) tawny_svm_oci_markc can be contacted.	CIFS Protocol	Potential CIFS and NFS outages may occur.	 Active IQ ↗
 tawny01	ONTAP version 8.3.2 has entered the Self-Service Support period.	ONTAP	Self-Service Support is the time period where NetApp does not provide support for a version of a software product, but related documentation is still available on the NetApp Support Site.	 Active IQ ↗

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.

Storage Summary		
Model: FAS6210	Microcode Version: 8.3.2 clustered Data ONTAP	Management: HTTPS://10.197.143.25:443
Vendor: NetApp	Raw Capacity: 80,024.3 GB	FC Fabrics Connected: 0
Family: FAS6200	Latency - Total: 0.77 ms	Performance Policies:
Serial Number: 1-80-000013	IOPS - Total: 1,819.19 IO/s	Risks:  108 risks detected by  Active IQ ↗
IP: 10.197.143.25	Throughput - Total: 41.69 MB/s	

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.

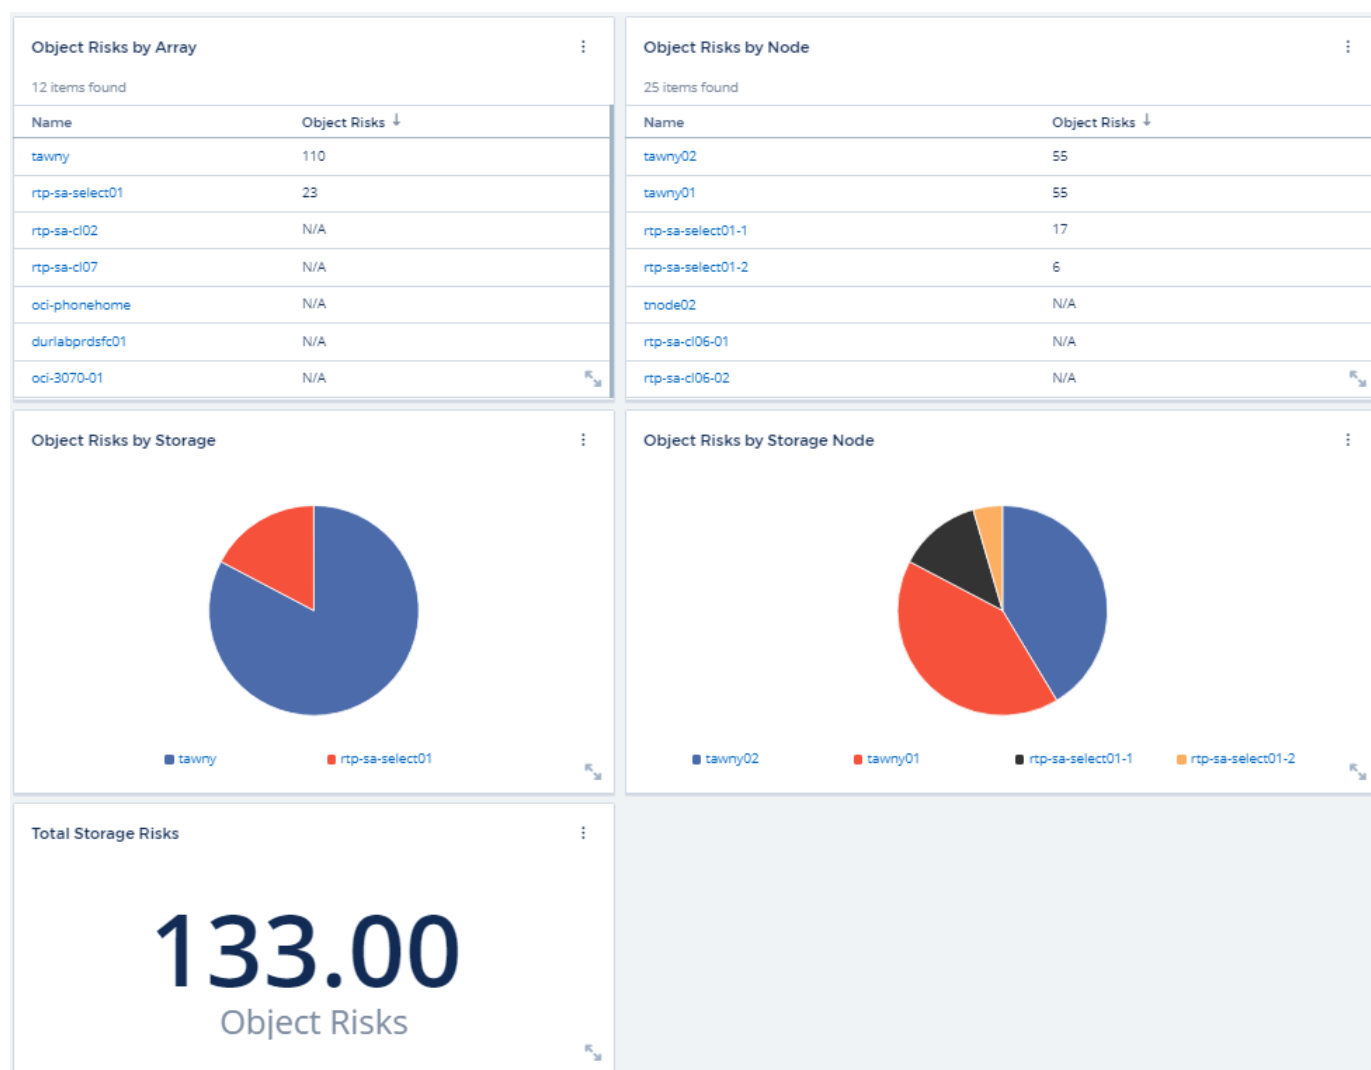
1. In the Cloud Insights Summary widget or Risks table, click the "Active IQ" link.
2. Sign in to your NetApp Support account. You are taken directly to the storage node page in Active IQ.

Querying for Risks

In Cloud Insights, you can add the **Object Risks** column to a storage or storage node query. If the returned result includes Active IQ-Monitored storage systems, the Object Risks 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.



Troubleshooting

Troubleshooting General Cloud Insights Problems

Here you will find suggestions for troubleshooting Cloud insights.

See also [Troubleshooting Linux Acquisition Unit Problems](#) and [Troubleshooting Windows Acquisition Unit Problems](#).

Problem:	Try this:
Cloud Insights logs out every 5 minutes	Enable 3rd party cookie acceptance from <code>[*.]auth0.com</code> in your browser settings. For example in Chrome's incognito mode, the default browser settings block third-party cookies. Try the following: Enter "chrome://settings/cookies" in browser URL. Select "Allow all cookies" option.
I have a Cloud Central account but am unable to login to Cloud Central.	Contact NetApp Cloud Central 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	Clear browser history, cookies, and cache. Try with a different browser profile (i.e. Chrome - add Person).

If you have an active Cloud Insights subscription you can use these support options:

[Phone](#)

For more information, see the [Cloud Insights Support Documentation](#).

Troubleshooting Acquisition Unit Problems on Linux

Here you will find suggestions for troubleshooting problems with Acquisition Units on a Linux server.

Problem:	Try this:
AU status on the Admin > Data Collectors page in the Acquisition Units tab displays "Certificate Expired" or "Certificate Revoked" .	<p>Click on the menu to the right of the AU and select Restore Connection. Follow the instructions to restore your Acquisition Unit:</p> <ol style="list-style-type: none">1. Stop the Acquisition Unit (AU) service. You can click the <i>Copy Stop Command</i> 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 <code>/var/lib/netapp/cloudinsights/acq/conf</code> folder on the AU.3. Click the <i>Copy Token</i> button, and paste this token into the file you created.4. Restart the AU service. Click the <i>Copy Restart Command</i> button, and paste the command into a command prompt on the AU.
Installation fails on SELinux	When the AU is installed on SELinux, SE should be either disabled or set to permissive mode. Once the installation is complete, enforcing mode can be enabled.
Server Requirements not met	Ensure that your Acquisition Unit server or VM meets requirements
Network Requirements not met	<p>Ensure that your Acquisition Unit server/VM can access your Cloud Insights environment (<code><environment-name>.c01.cloudinsights.netapp.com</code>) through SSL connection over port 443. Try the following commands:</p> <pre>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</pre>

Proxy Server not configured properly	<p>Verify your proxy settings, and uninstall/re-install the Acquisition Unit software if necessary to enter the correct proxy settings.</p> <p>1. Try "curl". Refer to "man curl" information/documentation regarding proxies: --preproxy, --proxy-* (that's a wildcard "*" because curl supports many proxy settings).</p> <p>2. Try "wget". Check documentation for proxy options.</p>
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 (<i>sudo cloudinsights-uninstall.sh</i>) 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 <i>unzip</i> library installed on the server. To install the <i>unzip</i> library, run the command <i>*sudo yum install unzip*</i> before running the Acquisition Unit install script
Permission issues	Be sure you are logged in as a user with <i>sudo</i> permissions
Acquisition Not Running:	<p>Gather the acq.log from /opt/netapp/cloudinsights/acq/logs (Linux)</p> <p>Restart the Acquisition Service: <i>sudo cloudinsights-service.sh restart acquisition</i></p>
Data Collection Issues:	Send an Error Report from the Data Collector landing page by clicking the "Send Error Report" button
Status: Heartbeat Failed	<p>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".</p> <p>Troubleshoot steps:</p> <p>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.</p>

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.

Related Information

- [Managing Access Keys for IAM Users](#)

Terminology

Cloud Insights acquires the following inventory information from the Dell Xc Series 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:

Field	Description
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 XC Enterprise Manager server
- IP address of the XC Enterprise Manager server

Some things to try if you encounter problems with this data collector:

Inventory

Problem:	Try this:
Error: User authentication failed.	Make sure your credentials for this device are correct.

Performance

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.	<ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• Check that the system performance monitors are running correctly• Please contact EMC support to enable updating the system performance monitor log files

error: "Cache server is waiting for the system manager"
Customer can take action. What can customer do about this scenario?

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 AU.

Troubleshooting

Some things to try if you encounter problems with this data collector:

Inventory

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	<ul style="list-style-type: none">* 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

Problem:	Try this:
Error: <ul style="list-style-type: none">* Error executing export utility* Error executing external command	<ul style="list-style-type: none">* 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:<ul style="list-style-type: none">- 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	<ul style="list-style-type: none">* 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

Problem:	Try this:
Error: Export tools logged "Unable to get time range for monitoring".	<ul style="list-style-type: none"> * 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	<ul style="list-style-type: none"> * 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	<ul style="list-style-type: none"> * 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	<ul style="list-style-type: none"> * 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: <ul style="list-style-type: none"> - 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	<ul style="list-style-type: none"> * 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: <ul style="list-style-type: none"> - 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

- 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

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.

- 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

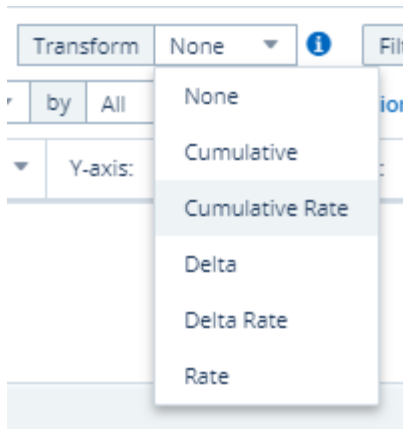
Best Practice: It is highly recommended for each Hyper-V hypervisor to have “Resource Metering” turned on for every VM, on every host. This allows each hypervisor to have more data available for Cloud Insights on each guest. If this is not set, fewer performance metrics are acquired for each guest. More information on Resource metering can be found in the [Microsoft documentation](#).

Categorizing Telegraf Data

Because of the varied nature of the ingested data, Cloud Insights makes its best attempt to categorize the data in meaningful ways. To do this, it examines the meta-data (data about the data) that is gathered.

- For meta-data that it knows, Cloud Insights attaches a meaningful name and categorizes the metric or counter as "Stable". Stable metrics will have names that use "Title Case", with individual words separated by space (), such as *IOs In Progress*.
- For data that it might know but isn't 100% certain, it keeps the source-determined name and assigns it to the "Guess" category. Guess metrics will appear with lower-case names, with individual words separated by an underscore (_), such as *virtual_machine_io*. "Guess" data cannot be used in filters or grouping.

All metrics gathered by Telegraf data collectors can be used in dashboard widgets. In Time-Series widgets (Line, Spline, Area, Stacked Area), you are given a choice of how you want to **Transform** the values: None, Sum, Delta, etc.. These [display options](#) are explained below.



Telegraf Data Display Options