



Observability

Data Infrastructure Insights

NetApp
August 28, 2025

This PDF was generated from https://docs.netapp.com/us-en/data-infrastructure-insights/concept_dashboards_overview.html on August 28, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Observability	1
Creating Dashboards	1
Dashboards Overview	1
Dashboard Features	4
Sample Dashboards	36
Best Practices for Dashboards and Widgets	41
Working with Queries	45
Assets used in queries	45
Creating Queries	47
Viewing queries	52
Exporting query results to a .CSV file	53
Modifying or Deleting a Query	54
Assigning multiple applications to or removing multiple applications from assets	54
Copying table values	55
Log Explorer	56
Identifying inactive devices	61
Insights	62
Insights	62
Insights: Shared Resources Under Stress	62
Insights: Kubernetes Namespaces Running out of Space	65
Insights: Reclaim ONTAP Cold Storage	66
Monitors and Alerts	69
Alerting with Monitors	69
Viewing and Managing Alerts from Monitors	79
Configuring Email Notifications	82
Anomaly Detection Monitors	85
System Monitors	88
Configuring Email Notifications	159
Webhook Notifications	162
Working with Annotations	176
Defining annotations	176
Using annotations	178
Creating annotation rules	181
Importing Annotations	182
Working with Applications	184
Tracking asset usage by application	185
Creating Applications	185
Automatic Device Resolution	186
Automatic Device Resolution Overview	186
Device Resolution rules	189
Fibre Channel device resolution	191
IP device resolution	194
Setting options in the Preferences tab	196

Regular expression examples	196
Asset Page Information	203
Asset Page Overview	203
Filtering for Objects In-Context	208
Storage Virtualization	209
Hints and Tips to Search for Assets and Alerts	211
Analyzing Data	214
SAN Analyzer Overview	214
VM Analyzer Overview	217
Monitor Infrastructure Health	219
Reporting	220
Data Infrastructure Insights Reporting Overview	220
Data Infrastructure Insights Reporting User Roles	221
Predefined Reports Made Easy	223
Storage Manager Dashboard	227
Creating a Report (Example)	230
Managing Reports	232
Creating Custom Reports	235
Access the Reporting Database via API	242
How historical data is retained for Reporting	246
Data Infrastructure Insights Reporting Schema Diagrams	247
Data Infrastructure Insights Schemas for Reporting	299

Observability

Creating Dashboards

Dashboards Overview

Data Infrastructure 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.
- **Save button:** Allows you to save or delete the dashboard.

You can rename the current dashboard by clicking **Rename** from the Save menu.

- **+ 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.

- **+ Add Variable** button, which allows you to use variables to actively filter the dashboard data.

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.



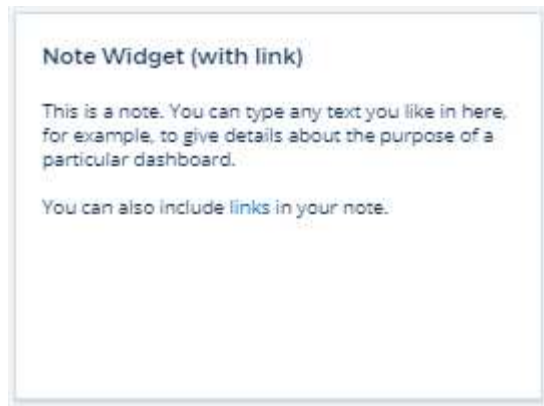
- **Scatter Plot Chart:** Plots related data as points, for example, IOPS and latency. In this example, you can quickly locate assets with high throughput and low IOPS.




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



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



- **Time Bar Chart:** Displays log or metric data over time.



- **Alerts Table:** Displays up to the last 1,000 alerts.

For more detailed explanations of these and other Dashboard Features, [click here](#).

Setting a Dashboard as your Home Page

You can choose which dashboard to set as your tenant's **home page** using either of the following methods:

- Go to **Dashboards > Show All Dashboards** to display the list of dashboards on your tenant. 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.

toc:[]

Widget Naming

Widgets are automatically named based on the object, metric, or attribute selected for the first widget query. If you also choose a grouping for the widget, the "Group by" attributes are included in the automatic naming (aggregation method and metric).

C
B
A

☒ A) Query
 Chart Type: Bar Chart
 Chart Color: ■
 Decimal Places: 2
 Convert to Expression

Object: agent.node
 Metric: cpu.time_active
 Display Unit: cpu.time_active (None)

Display: Last 24 Hours
 Aggregated by: Last
 Save
Reset

Filter by Attribute: +

Filter by Metric: +

Group by: agent_node_ip
 aggregated by: Maximum
 Apply f(x): Rank
 Top: 10

A
C

Selecting a new object or grouping attribute updates the automatic name.

If you do not want to use the automatic widget name, you can simply type a new name.

Widget Placement and Size

All dashboard widgets can be positioned and sized according to your needs for each particular dashboard.

Duplicating a Widget

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

Displaying Widget Legends

Most widgets on dashboards can be displayed with or without legends. Legends in widgets can be turned on or off on a dashboard by either of the following methods:

- When displaying the dashboard, click the **Options** button on the widget and select **Show Legends** in the menu.

As the data displayed in the widget changes, the legend for that widget is updated dynamically.

When legends are displayed, if the landing page of the asset indicated by the legend can be navigated to, the legend will display as a link to that asset page. If the legend displays "all", clicking the link will display a query page corresponding to the first query in the widget.

Transforming Metrics

Data Infrastructure Insights provides different **transform** options for certain metrics in widgets (specifically, those metrics called "Custom" or Integration Metrics, such as from Kubernetes, ONTAP Advanced Data, Telegraf plugins, etc.), allowing you to display the data in a number of ways. When adding transformable metrics to a widget, you are presented with a drop-down giving the following transform choices:

None:

Data is displayed as is, with no manipulation.

Rate:

Current value divided by the time range since the previous observation.

Cumulative:

The accumulation of the sum of previous values and the current value.

Delta:

The difference between the previous observation value and the current value.

Delta rate:

Delta value divided by the time range since the previous observation.

Cumulative Rate:

Cumulative value divided by the time range since the previous observation.

Note that transforming metrics does not change the underlying data itself, but only the way that data is displayed.

Dashboard widget queries and filters

Queries

The Query in a dashboard widget is a powerful tool for managing the display of your data. Here are some things to note about widget queries.

Some widgets can have up to five queries. Each query will plot its own set of lines or graphs in the widget. Setting rollup, grouping, top/bottom results, etc. on one query does not affect any other queries for the widget.

You can click on the eye icon to temporarily hide a query. The widget display updates automatically when you hide or show a query. This allows you to check your displayed data for individual queries as you build your widget.

The following widget types can have multiple queries:

- Area chart
- Stacked area chart
- Line chart
- Spline chart
- Single value widget

The remaining widget types can have only a single query:

- Table
- Bar chart
- Box plot
- Scatter plot

Filtering in dashboard widget 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.

Wildcards and Expressions

When you are filtering for text or list values in queries or dashboard widgets, as you begin typing you are presented with the option to create a **wildcard filter** based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create **expressions** using NOT or OR, or you can select the "None" option to filter for null values in the field.



Filters based on wildcards or expressions (e.g. NOT, OR, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.

kubernetes.pod

Filter By
pod_name
ingest
ci-service-audit-5f775dd975-brfdc
+
?

Group
pod_name

3 items found

pod_name
ci-service-audit-5f775dd975-brfdc
ci-service-datalake-ingestion-85b5bdfd6d-2qbwr
service-foundation-ingest-767dfd5bfc-vxd5p

Note that Wildcard and Expression filtering works with text or lists but not with numerics, dates or booleans.

Advanced Text Filtering with Contextual Type-Ahead Suggestions

Filtering in widget queries is *contextual*; when you select a filter value or values for a field, the other filters for that query will show values relevant to that filter.

For example, when setting a filter for a specific object *Name*, the field to filter for *Model* will only show values relevant to that object Name.

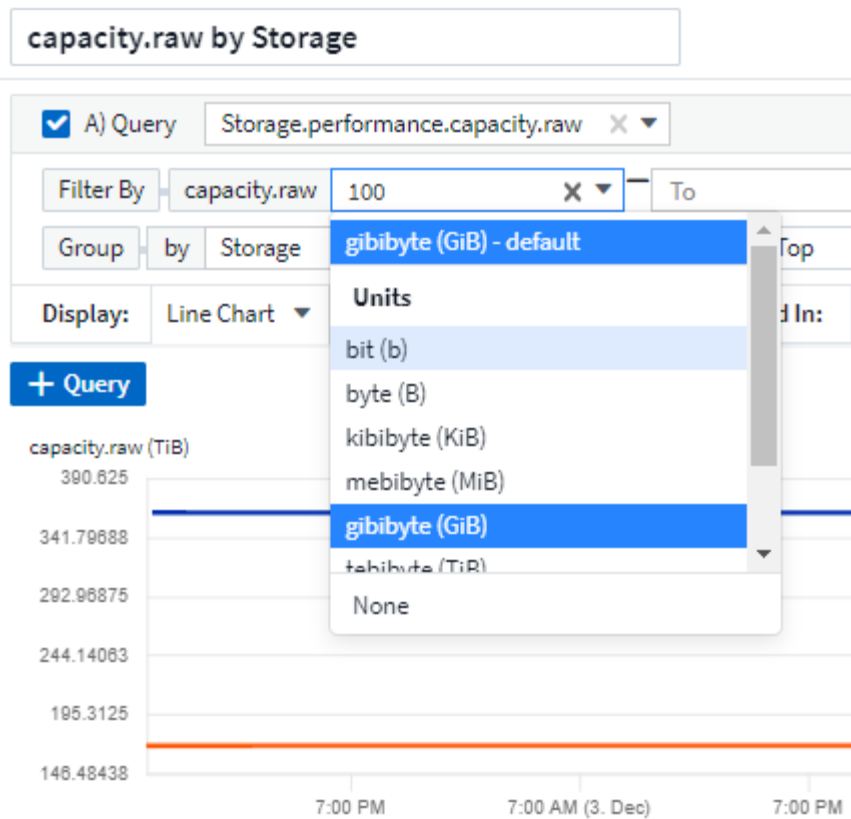
Contextual filtering also applies to dashboard page variables (text-type attributes or annotations only). When you select a filter value for one variable, any other variables using related objects will only show possible filter values based on the context of those related variables.

Note that only Text filters will show contextual type-ahead suggestions. Date, Enum (list), etc. will not show type-ahead suggestions. That said, you *can* set a filter on an Enum (i.e. list) field and have other text fields be filtered in context. For example, selecting a value in an Enum field like Data Center, then other filters will show only the models/names in that data center), but not vice-versa.

The selected time range will also provide context for the data shown in filters.

Choosing the filter units

As you type a value in a filter field, you can select the units in which to display the values on the chart. For example, you can filter on raw capacity and choose to display in the default GiB, or select another format such as TiB. This is useful if you have a number of charts on your dashboard showing values in TiB and you want all your charts to show consistent values.



Additional Filtering Refinements

The following can be used to further refine your filters.

- An asterisk enables you to search for everything. For example,

```
vol*rhel
```

displays all resources that start with "vol" and end with "rhel".

- The question mark enables you to search for a specific number of characters. For example,

```
BOS-PRD??-S12
```

displays *BOS-PRD12-S12*, *BOS-PRD13-S12*, and so on.

- The OR operator enables you to specify multiple entities. For example,

```
FAS2240 OR CX600 OR FAS3270
```

finds multiple storage models.

- The NOT operator allows you to exclude text from the search results. For example,

NOT EMC*


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

NOT *

to display fields that contain no value.

Identifying objects returned by queries and filters

The objects returned by queries and filters look similar to those shown in the following illustration. Objects with 'tags' assigned to them are annotations while the objects without tags are performance counters or object attributes.



Grouping and Aggregation

Grouping (Rolling Up)

Data displayed in a widget is grouped (sometimes called rolled-up) from the underlying data points collected during acquisition. For example, if you have a line chart widget showing Storage IOPS over time, you might want to see a separate line for each of your data centers, for a quick comparison. You can choose to group this data in one of several ways:

- **Average:** displays each line as the *average* of the underlying data.
- **Maximum:** displays each line as the *maximum* of the underlying data.
- **Minimum:** displays each line as the *minimum* of the underlying data.
- **Sum:** displays each line as the *sum* of the underlying data.
- **Count:** displays a *count* of objects that have reported data within the specified time frame. You can choose the *Entire Time Window* as determined by the dashboard time range.

Steps

To set the grouping method, do the following.

1. In your widget's query, choose an asset type and metric (for example, *Storage*) and metric (such as *Performance IOPS Total*).
2. For **Group**, choose a roll up method (such as *Average*) and select the attributes or metrics by which to roll up the data (for example, *Data Center*).

The widget updates automatically and shows data for each of your data centers.

You can also choose to group *all* of the underlying data into the chart or table. In this case, you will get a single line for each query in the widget, which will show the average, min, max, sum, or count of the chosen metric or metrics for all of the underlying assets.

Clicking the legend for any widget whose data is grouped by "All" opens a query page showing the results of the first query used in the widget.

If you have set a filter for the query, the data is grouped based on the filtered data.

Note that when you choose to group a widget by any field (for example, *Model*), you will still need to Filter by that field in order to properly display the data for that field on the chart or table.

Aggregating data

You can further align your time-series charts (line, area, etc.) by aggregating data points into minute, hour, or day buckets before that data is subsequently rolled up by attribute (if chosen). You can choose to aggregate data points according to their *Average*, *Maximum*, *Minimum*, *Sum*, or *Count*.

A small interval combined with a long time range may result in an "Aggregation interval resulted in too many data points." warning. You might see this if you have a small interval and increase the dashboard time frame to 7 days. In this case, Insight will temporarily increase the aggregation interval until you select a smaller time frame.

You can also aggregate data in the bar chart widget and single-value widget.

Most asset counters aggregate to *Average* by default. Some counters aggregate to *Max*, *Min*, or *Sum* by default. For example, port errors aggregate to *Sum* by default, where storage IOPS aggregate to *Average*.

Showing Top/Bottom Results

In a chart widget, you can show either the **Top** or **Bottom** results for rolled up data, and choose the number of results shown from the drop-down list provided. In a table widget, you can sort by any column.

Chart widget top/bottom

In a chart widget, when you choose to rollup data by a specific attribute, you have the option of viewing either the top N or bottom N results. Note that you cannot choose the top or bottom results when you choose to rollup by *all* attributes.

You can choose which results to display by choosing either **Top** or **Bottom** in the query's **Show** field, and selecting a value from the list provided.

Table widget show entries

In a table widget, you can select the number of results shown in the table results. You are not given the option to choose top or bottom results because the table allows you to sort ascending or descending by any column

on demand.

You can choose the number of results to show in the table on the dashboard by selecting a value from the query's **Show entries** field.

Grouping in Table Widget

Data in a table widget can be grouped by any available attribute, allowing you to see an overview of your data, and to drill-down into it for more detail. Metrics in the table are rolled up for easy viewing in each collapsed row.

Table widgets allow you to group your data based on the attributes you set. For example, you might want your table to show total storage IOPS grouped by the data centers in which those storages live. Or you might want to display a table of virtual machines grouped according to the hypervisor that hosts them. From the list, you can expand each group to view the assets in that group.

Grouping is only available in the Table widget type.

Grouping example (with rollup explained)

Table widgets allow you to group data for easier display.

In this example, we will create a table widget showing all VMs grouped by Data Center.

Steps

1. Create or open a dashboard, and add a **Table** widget.
2. Select *Virtual Machine* as the asset type for this widget.
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 *Average*. This means that the number shown for the group is the average of all the Total IOPS reported for each VM inside the group. You can choose to roll this column up by *Average*, *Sum*, *Minimum* or *Maximum*. Any column that you display that contains performance metrics can be rolled up individually.



6. Click *All* and select *Hypervisor name*.

The VM list is now grouped by Hypervisor. You can expand each hypervisor to view the VMs hosted by it.

7. Click **Save** to save the table to the dashboard. You can resize or move the widget as desired.

8. Click **Save** to save the dashboard.

Performance data roll up

If you include a column for performance data (for example, *IOPS - Total*) in a table widget, when you choose to group the data you can then choose a roll up method for that column. The default roll up method is to display the average (*avg*) of the underlying data in the group row. You can also choose to display the sum, minimum, or maximum of the data.

Dashboard time range selector

You can select the time range for your dashboard data. Only data relevant to the selected time range will be displayed in widgets on the dashboard. You can select from the following time ranges:

- Last 15 Minutes
- Last 30 Minutes
- Last 60 Minutes
- Last 2 Hours
- Last 3 Hours (this is the default)
- Last 6 Hours
- Last 12 Hours
- Last 24 Hours
- Last 2 Days
- Last 3 Days
- Last 7 Days

- Last 30 Days
- Custom time range

The Custom time range allows you to select up to 31 consecutive days. You can also set the Start Time and End Time of day for this range. The default Start Time is 12:00 AM on the first day selected and the default End Time is 11:59 PM on the last day selected. Clicking **Apply** will apply the custom time range to the dashboard.

Zooming in to a time range

While viewing a time-series widget (Line, Spline, Area, Stacked Area)—or a graph on a landing page—you can drag the mouse over the graph to zoom in. In the upper right of the screen you can then lock that time range so that graphs on other pages reflect data for that locked time range. To unlock, select a different time range from the list.


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 choose the desired time range, and 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) bar chart, column chart, pie chart, or table widget allows you to build expressions from metrics you choose, and show the result of those expressions in a single graph (or column in the case of the [table widget](#)). The following examples use expressions to solve specific problems. In the first example, we want to show Read IOPS as a percentage of Total IOPS for all storage assets on your tenant. The second example gives visibility into the "system" or "overhead" IOPS that occur on your tenant—those IOPS that are not directly from reading or writing data.

You can use variables in expressions (for example, $\$Var1 * 100$)

Expressions Example: Read IOPS percentage

In this example, we want to show Read IOPS as a percentage of Total IOPS. You can think of this as the following formula:

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

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 $\text{Read Percentage} = (\text{Read IOPS} / \text{Total IOPS}) \times 100$, so we would write this expression as:

$$(b / a) * 100$$

6. The **Label** field identifies the expression. Change the label to "Read Percentage", or something equally meaningful for you.

7. Change the **Units** field to "%" or "Percent".

The chart displays the IOPS Read percentage over time for the chosen storage devices. If desired, you can set a filter, or choose a different rollup method. Be aware that if you select Sum as the rollup method, all percentage values are added together, which potentially may go higher than 100%.

8. Click **Save** to save the chart to your dashboard.

Expressions example: "System" I/O

Example 2: Among the metrics collected from data sources are read, write, and total IOPS. However, the total number of IOPS reported by a data source sometimes includes "system" IOPS, which are those IO operations that are not a direct part of data reading or writing. This system I/O can also be thought of as "overhead" I/O, necessary for proper system operation but not directly related to data operations.

To show these system I/Os, you can subtract read and write IOPS from the total IOPS reported from acquisition. The formula might look like this:

$$\text{System IOPS} = \text{Total IOPS} - (\text{Read IOPS} + \text{Write IOPS})$$

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.

To use a variable in an expression, simply type the variable name, for example, *\$var1 * 100*. Only numeric variables can be used in expressions.

Expressions in a Table Widget

Table widgets handle expressions a little differently. You can have up to five expressions in a single table widget, each of which is added as a new column to the table. Each expression can include up to five values on which to perform its calculation. You can easily name the column something meaningful.




Variables

Variables allow you to change the data displayed in some or all widgets on a dashboard at once. By setting one or more widgets to use a common variable, changes made in one place cause the data displayed in each widget to update automatically.

Variable types


A variable can be one the following types:

- **Attribute:** Use an object's attributes or metrics to filter
- **Annotation:** Use a pre-defined [Annotation](#) to filter widget data.
- **Text:** An alphanumeric string.
- **Numerical:** A number value. Use by itself, or as a "from" or "to" value, depending on your widget field.
- **Boolean:** Use for fields with values of True/False, Yes/No, etc. For the boolean variable, the choices are Yes, No, None, Any.
- **Date:** A date value. Use as a "from" or "to" value, depending on your widget's configuration.




Attribute variables

Selecting an Attribute type variable allows you to filter for widget data containing the specified attribute value or values. The example below shows a line widget displaying free memory trends for Agent nodes. We have created a variable for Agent Node IPs, currently set to show all IPs:

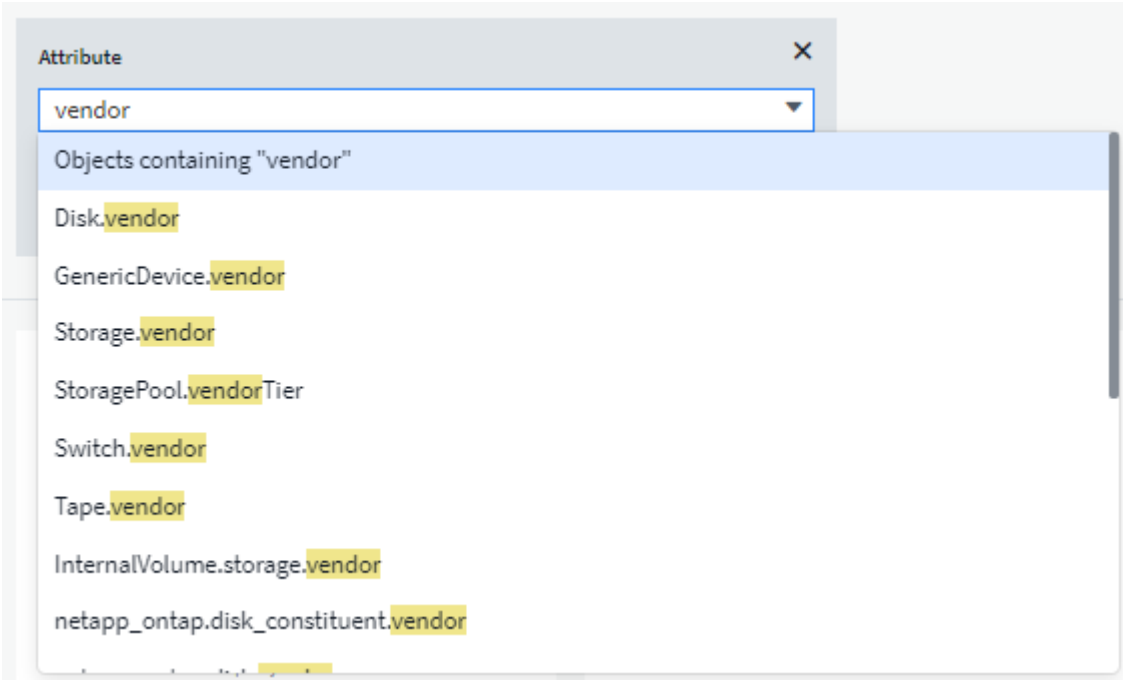


But if you temporarily want to see only nodes on individual subnets on your tenant, you can set or change the variable to a specific Agent Node IP or IPs. Here we are viewing only the nodes on the "123" subnet:

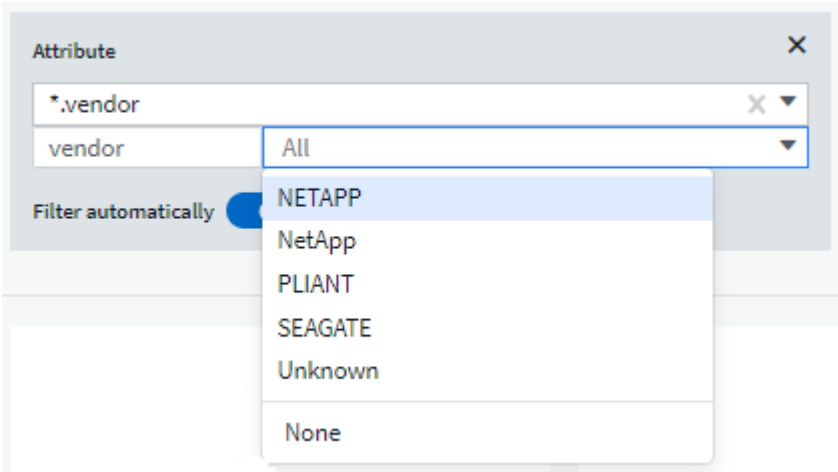


You can also set a variable to filter on *all* objects with a particular attribute regardless of object type, for

example objects with an attribute of "vendor", by specifying `*.vendor` in the variable field. You do not need to type the `"*."`; Data Infrastructure Insights will supply this if you select the wildcard option.



When you drop-down the list of choices for the variable value, the results are filtered so show only the available vendors based on the objects on your dashboard.



If you edit a widget on your dashboard where the attribute filter is relevant (meaning, the widget's objects contain any `*.vendor attribute`), it shows you that the attribute filter is automatically applied.

14

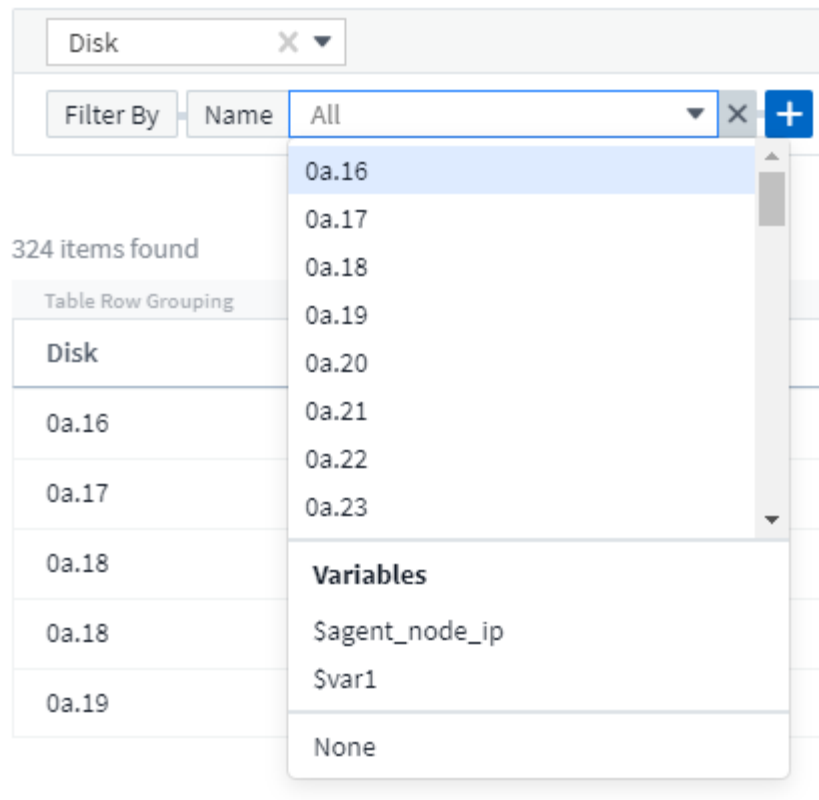
Applying variables is as easy as changing the attribute data of your choice.

Annotation variables

Choosing an Annotation variable allows you to filter for objects associated with that annotation, for example, those belonging to the same Data Center.


Text, Number, Date, or Boolean variable

You can create generic variables that are not associated with a particular attribute by selecting a variable type of *Text*, *Number*, *Boolean*, or *Date*. Once the variable has been created, you can select it in a widget filter field. When setting a filter in a widget, in addition to specific values that you can select for the filter, any variables that have been created for the dashboard are displayed in the list—these are grouped under the "Variables" section in the drop-down and have names starting with "\$". Choosing a variable in this filter will allow you to search for values that you enter in the variable field on the dashboard itself. Any widgets using that variable in a filter will be updated dynamically.



Variable Filter Scope

When you add an Annotation or Attribute variable to your dashboard, the variable can be applied to *all* widgets on the dashboard, meaning that all widgets on your dashboard will display results filtered according to the value you set in the variable.



Note that only Attribute and Annotation variables can be filtered automatically like this. Non-Annotation or -Attribute variables cannot be automatically filtered. Individual widgets must each be configured to use variables of these types.

To disable automatic filtering so that the variable only applies to the widgets where you have specifically set it, click the "Filter automatically" slider to disable it.

To set a variable in an individual widget, open the widget in edit mode and select the specific annotation or attribute in the *Filter By* field. With an Annotation variable, you can select one or more specific values, or select the Variable name (indicated by the leading "\$") to allow typing in the variable at the dashboard level. The same applies to Attribute variables. Only those widgets for which you set the variable will show the filtered results.

Filtering in variables is *contextual*; when you select a filter value or values for a variable, the other variables on

your page will show only values relevant to that filter.

For example, when setting a variable filter to a specific storage *Model*, any variables set to filter for storage *Name* will only show values relevant to that *Model*.

To use a variable in an expression, simply type the variable name as part of the expression, for example, *\$var1 * 100*. Only Numeric variables can be used in expressions. You cannot use numeric Annotation or Attribute variables in expressions.

Filtering in variables is *contextual*; when you select a filter value or values for a variable, the other variables on your page will show only values relevant to that filter.

For example, when setting a variable filter to a specific storage *Model*, any variables set to filter for storage *Name* will only show values relevant to that *Model*.

Variable naming

Variables names:

- Must include only the letters a-z, the digits 0-9, period (.), underscore (_), and space ().
- Cannot be longer than 20 characters.
- Are case-sensitive: *\$CityName* and *\$cityname* are different variables.
- Cannot be the same as an existing variable name.
- Cannot be empty.

Formatting Gauge Widgets

The Solid and Bullet Gauge widgets allow you to set thresholds for *Warning* and/or *Critical* levels, providing clear representation of the data you specify.

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: *Blue* Units Displayed In: *Auto Format*

+ Query

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.

Formatting Table Widgets

Like single-value and gauge widgets, you can set conditional formatting in table widgets, allowing you to

highlight data with colors and/or special icons.

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

14 items found in 1 group

Table Row Grouping	Expanded Detail	Metrics & Attributes
All	Storage Pool	capacityRatio.used (%)
All (14)	--	95.15
--	rtp-sa-cl06-02:aggr_data1_rtp_sa_cl06_02	0.79
--	rtp-sa-cl06-01:aggr_data1_rtp_sa_cl06_01	2.45
--	rtp-sa-cl06-02:aggr0_rtp_sa_cl06_02_root	95.15
--	rtp-sa-cl06-01:aggr0_rtp_sa_cl06_01_root	95.15

Formatting: ☒ Show Expanded Details Conditional Formatting: Background Color + Icon ☐ Show In Range as green

capacity.provisioned (GiB)

> Aggregation

> Unit Display

> Conditional Formatting Reset

If value is > (Greater than) >Warning 70 %>Critical 90 %

> Rename Column

Conditional formatting is set separately for each column in a table. For example, you can choose one set of thresholds for a capacity column, and another set for a throughput column.

If you change the Unit Display for a column, the conditional formatting remains and reflects the change in values. The images below show the same conditional formatting even though the display unit is different.

capacity.used (GiB) ↓

40,754.06
10,313.56
9,544.84
8,438.99
6,671.72

throughput.total (MiB/s)

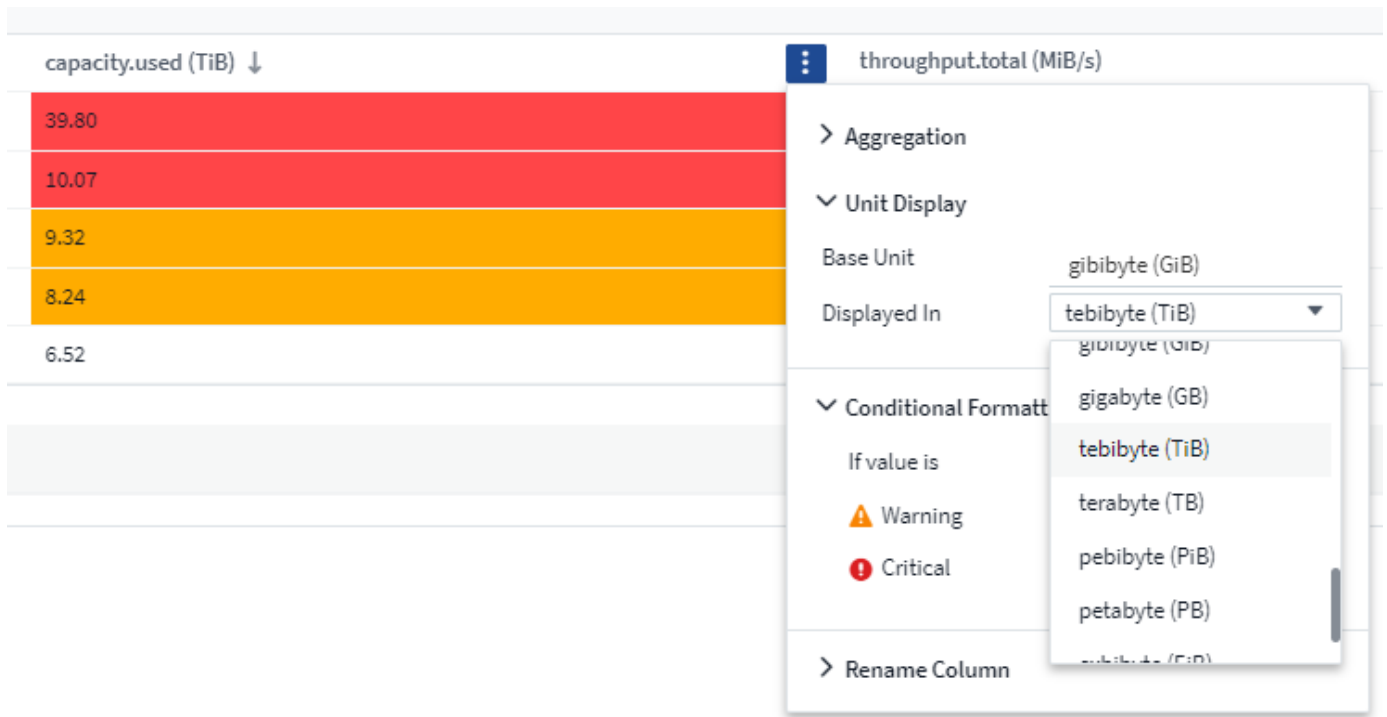
> Aggregation

> Unit Display

> Conditional Formatting Reset

If value is > (Greater than) >Warning 8000 GiB>Critical 10000 GiB

> Rename Column

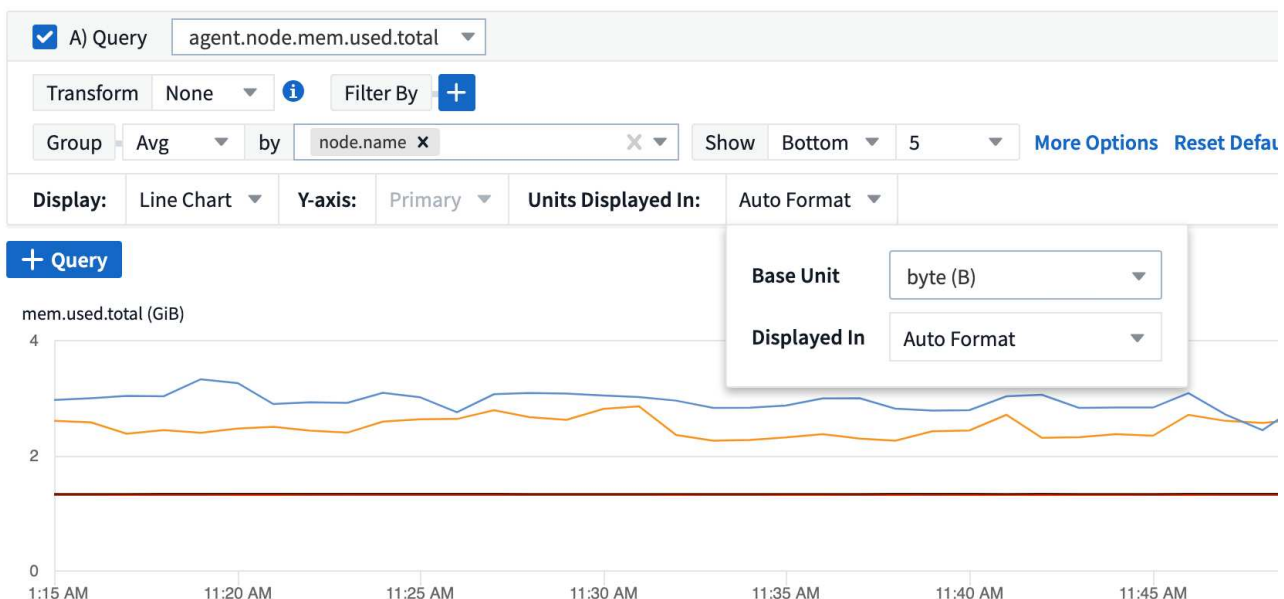


You can choose whether to display condition formatting as color, icons, or both.

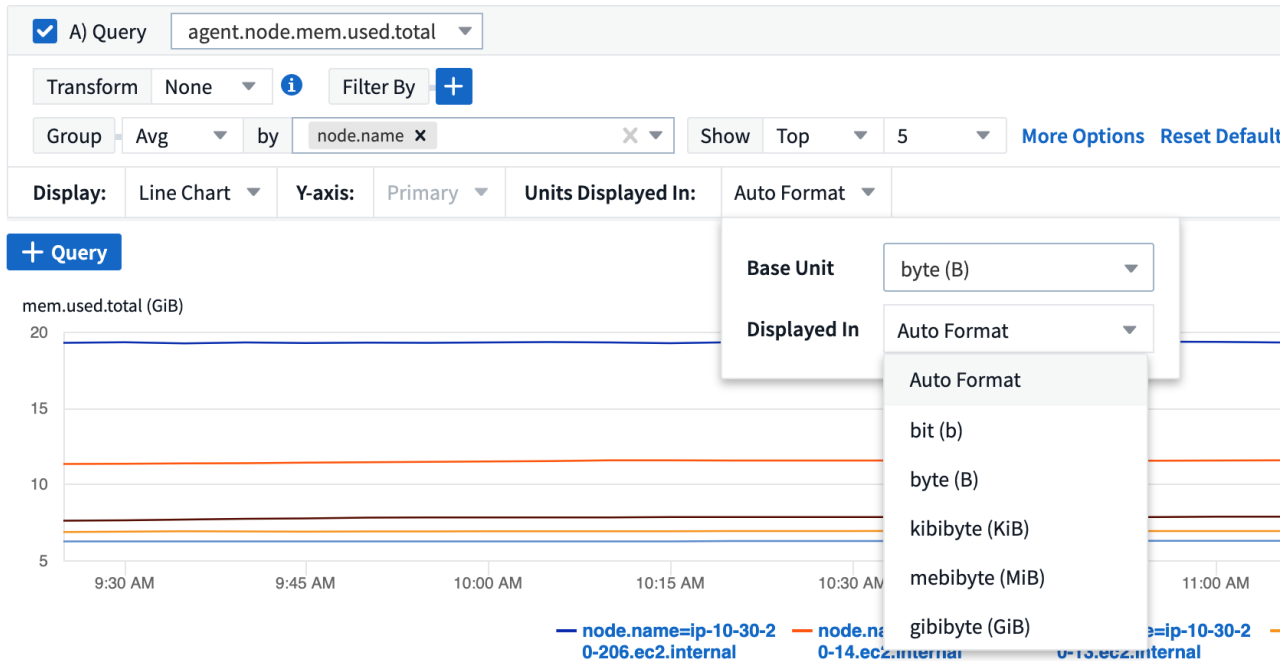
Choosing the Unit for Displaying Data

Most widgets on a dashboard allow you to specify the Units in which to display values, for example *Megabytes*, *Thousands*, *Percentage*, *Milliseconds (ms)*, etc. In many cases, Data Infrastructure 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.

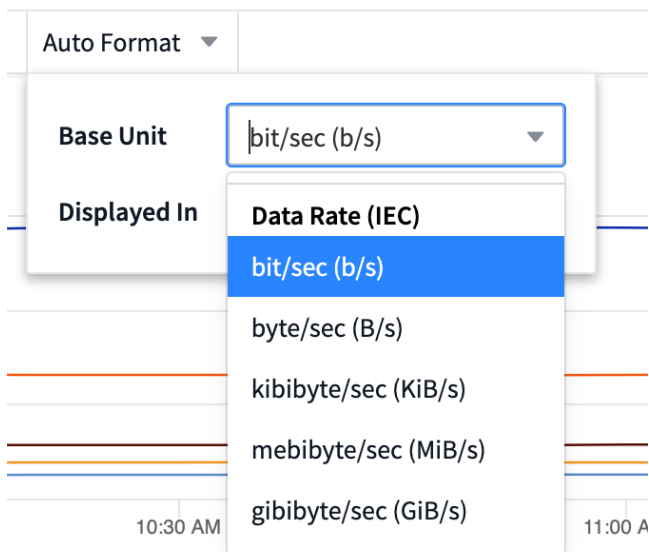
In the line chart example below, the data selected for the widget is known to be in *bytes* (the base IEC Data unit: see the table below), so the Base Unit is automatically selected as 'byte (B)'. However, the data values are large enough to be presented as gibibytes (GiB), so Data Infrastructure Insights by default auto-formats the values as GiB. The Y-axis on the graph shows 'GiB' as the display unit, and all values are displayed in terms of that unit.



If you want to display the graph in a different unit, you can choose another format in which to display the values. Since the base unit in this example is *byte*, you can choose from among the supported "byte-based" formats: bit (b), byte (B), kibibyte (KiB), mebibyte (MiB), gibibyte (GiB). The Y-Axis label and values change according to the format you choose.



In cases where the base unit is not known, you can assign a unit from among the [available units](#), or type in your own. Once you assign a base unit, you can then select to display the data in one of the appropriate supported formats.



To clear out your settings and start again, click on **Reset Defaults**.

A word about Auto-Format

Most metrics are reported by data collectors in the smallest unit, for example as a whole number such as 1,234,567,890 bytes. By default, Data Infrastructure 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.



Data Infrastructure 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. 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 Data Infrastructure Insights timeouts, leaving the display live until logged out manually or automatically by authorization security protocols.



Because NetApp BlueXP has its own user login timeout of 7 days, Data Infrastructure 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 on your tenant. 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)

<input type="checkbox"/>	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 on your tenant, 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 Data Infrastructure 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?


Dark Theme

You can choose to display Data Infrastructure Insights using either a light theme (the default), which displays most screens using a light background with dark text, or a dark theme which displays most screens using a dark background with light text.


To switch between light and dark themes, click the username button in the upper right corner of the screen and choose the desired theme.



Dark Theme Dashboard view:



Light Theme Dashboard view:



Some screen areas, such as certain widget charts, still show light backgrounds even while viewed in dark theme.

Line Chart interpolation

Different data collectors often poll their data at different intervals. For example, data collector A may poll every 15 minutes while data collector B polls every five minutes. When a line chart widget (also spline, area, and stacked area charts) is aggregating this data from multiple data collectors into a single line (for example, when the widget is grouping by "all"), and refreshing the line every five minutes, data from collector B may be shown accurately while data from collector A may have gaps, thus affecting the aggregate until collector A polls again.

To alleviate this, Data Infrastructure Insights interpolates data when aggregating, using the surrounding data points to take a "best guess" at data until data collectors poll again. You can always view each data collector's object data individually by adjusting the widget's grouping.

Interpolation Methods

When creating or modifying a line chart (or spline, area, or stacked area chart), you can set the interpolation method to one of three types. In the "Group by" section, choose the desired Interpolation.

Group by

All

▼

aggregated by

Average

▼

Apply f(x)

Interpolation

Linear


▼

None


Linear

Stair


- **None:** Do nothing, i.e. do not generate points in between.



- **Stair:** A point is generated from the value of previous point. In a straight line, this would display as a typical "stair" layout.




- **Linear:** a point is generated as the value in between connecting the two points. Generates a line that looks like the line connecting the two points, but with additional (interpolated) data points.



Anomaly Bounds in Line Widgets

When including a Line or Spline chart widget on a dashboard or landing page, you may choose to view the chart in context of the **expected bounds** for the data. You can think of this as looking for anomalies in the patterns of your data.

DII uses seasonal data (hourly or daily) to set upper and lower bounds on where it *expects* the data to fall at a given time. If the data spikes above or falls below those expected bounds, the chart will highlight that as an anomaly.



To view anomaly bounds, edit the widget and choose *Show Anomaly Bounds*. You may choose from among two detection algorithms:

- **Adaptive Detector** adapts to changes quickly, making it helpful for detailed investigations.
- **Smooth Detector** minimizes noise and false positives, filtering out short-term fluctuations while still detecting significant shifts.

Additionally, you may choose to show either *Hourly* or *Daily* seasonality, as well as set the sensitivity of detection. *High* sensitivity detects more boundary crossing, *Low* sensitivity detects less.

Keep in mind that you may only view expected bounds when the chart is set to display a single line. If your Group By settings or filters show multiple lines, or if you have set multiple queries for the widget, the option to show expected bounds will be disabled.

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 on your tenant. 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 on your tenant.

Table Widget

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	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.anandh917202012	--	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 on your tenant.

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

X

Virtual Machine

Filter by IOPS - Total (IO/s) >= 10 + Group by Hypervisor name X

181 items found in 4 groups

Hypervisor name ↓	Name	Hypervisor name	IOPS - Total (IO/s)	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	us-east-1b	32.66	0.70
+ us-east-1a (38)		us-east-1a	121.22	0.81

Roll Up by Avg

Cancel Save

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


The table widget is saved to the dashboard.

You can resize the widget on the dashboard by dragging the lower-right corner. Make the widget wider to show all the columns clearly. Click **Save** to save the current dashboard.

Next we will add some charts to show our VM Performance. Let's create a line chart comparing VM latency with VMDK latency.


1. If necessary, click the **Edit** icon on the dashboard to enable Edit mode.

2. Click the **[Add widget]** icon and select *Line Chart* to add a new line chart widget to the dashboard.
3. The **Edit Widget** dialog opens. Name this widget "VM / VMDK Max Latency"
4. Select **Virtual Machine** and choose *Latency - Max*. Set any filters you wish, or leave **Filter by** empty. For **Roll up**, choose *Sum* by *All*. Display this data as a *Line Chart*, and leave *Y-Axis* as *Primary*.
5. Click the **[+Query]** button to add a second data line. For this line, select *VMDK* and *Latency - Max*. Set any filters you wish, or leave **Filter by** empty. For **Roll up**, choose *Sum* by *All*. Display this data as a *Line Chart*, and leave *Y-Axis* as *Primary*.
6. Click **[Save]** to add this widget to the dashboard.



Next we will add a chart showing VM Read, Write and Total IOPS in a single chart.

1. Click the **[Add widget]** icon and select *Area Chart* to add a new area chart widget to the dashboard.
2. The Edit Widget dialog opens. Name this widget "VM IOPS"
3. Select **Virtual Machine** and choose *IOPS - Total*. Set any filters you wish, or leave **Filter by** empty. for **Roll up**, choose *Sum* by *All*. Display this data as an *Area Chart*, and leave *Y-Axis* as *Primary*.
4. Click the **[+Query]** button to add a second data line. For this line, select **Virtual Machine** and choose *IOPS - Read*.
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

Data Infrastructure 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 Data Infrastructure 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.



- Now reverse the order of the counters by setting the first counter to *IOPS - Total* and the second to *Latency - Total*.

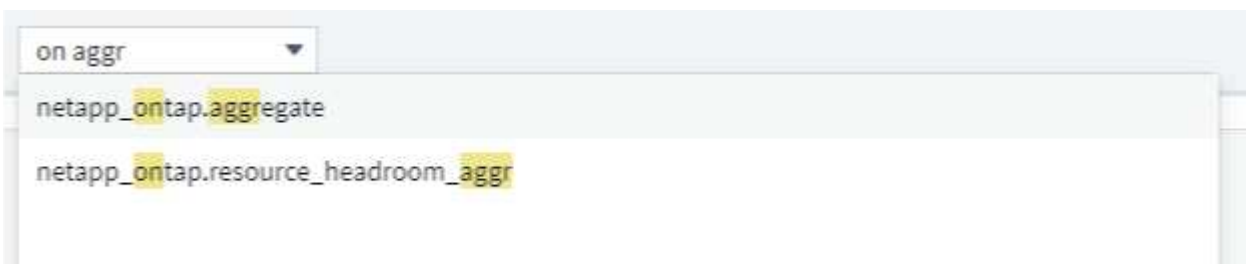
Latency - Total is now charted along the Y-axis in the chart, and *IOPS - Total* along the X-axis. VMs with higher IOPS now display on the right side of the chart.

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.

- 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 on your tenant at a granular level, allowing to filter for the data you want and sort the results to your liking.

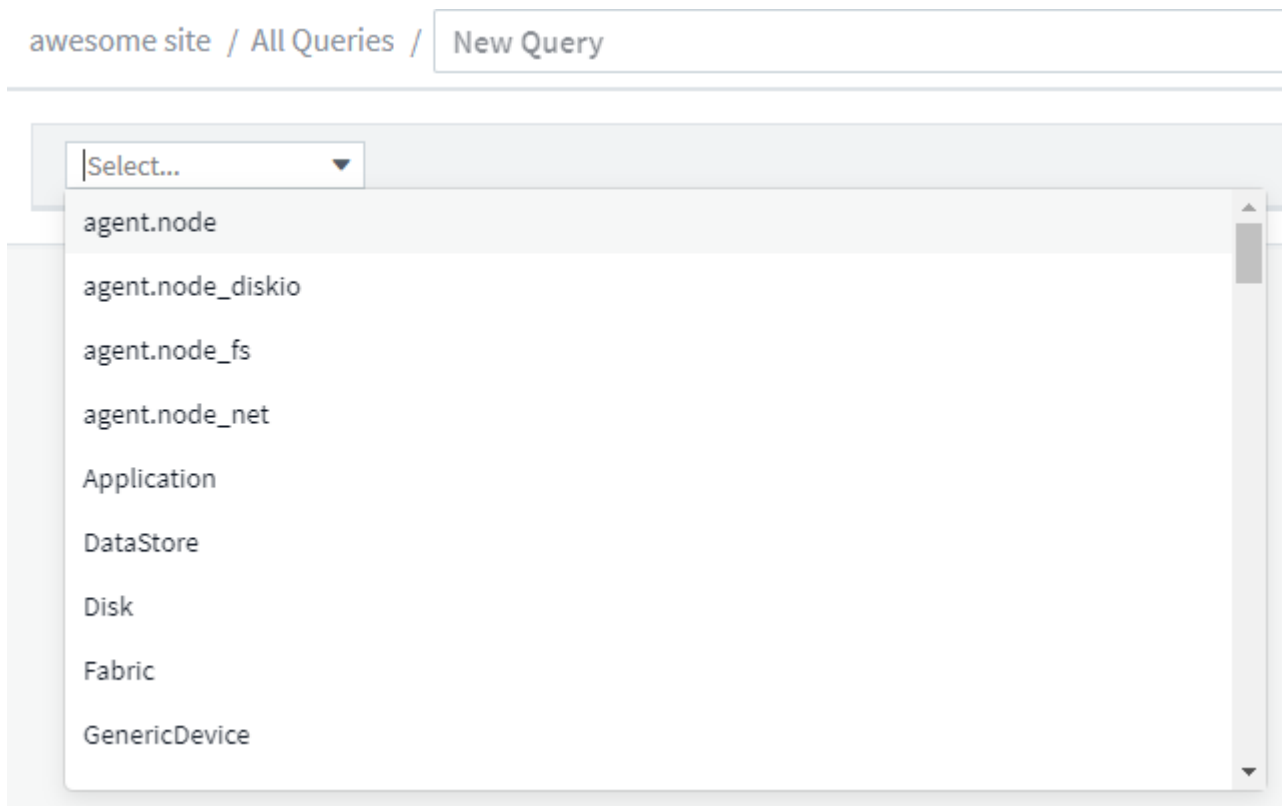
For example, you can create a query for *volumes*, add a filter to find particular *storages* associated with the selected volumes, add another filter to find a particular *annotation* such as "Tier 1" on the selected storages, and finally add another filter to find all storages with *IOPS - Read (IO/s)*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



Type-to-Search:

on aggr

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

Filter By
cluster_name
ci-

Group
aggr_name

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

Filter By

Group
aggr_name

14 items found

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

Search...
☐ Show Selected Only
☒ agent_version
☐ aggr_name
☐ cluster_location
☒ cluster_name
☐ cluster_serial_number
☐ cluster_version

See it in Action

[Explore and analyze with queries in Data Infrastructure Insights \(Video\)](#)

Choosing Aggregation, Units, Conditional Formatting

Aggregation and Units

For "value" columns, you can further refine your query results by choosing how the displayed values are aggregated as well as selecting the units in which those values are displayed. These options are found by selecting the "three dots" menu at the top corner of a column.

Units

You can select the units in which to display the values. For example, if the selected column shows raw capacity and the values are shown in GiB, but you prefer to display them as TiB, simply select TiB from the Unit Display drop-down.

Aggregation

By the same token, if the values shown are aggregated from the underlying data as "Average", but you prefer to show the sum of all values, select "Sum" from either the *Group by* drop-down (if you want any grouped values to show the sums) or from the *Time Aggregate By* drop-down (if you want the row values to show sums of underlying data).

You can choose to aggregate grouped data points by *Avg*, *Max*, *Min*, or *Sum*.

You can aggregate individual row data by *Average*, *Last data point acquired*, *Maximum*, *Minimum*, or *Sum*.

Conditional Formatting

Conditional Formatting allows you to highlight Warning-level and Critical-level thresholds in the query results list, bringing instant visibility to outliers and exceptional data points.

143 items found

Table Row Grouping	Metrics & Attributes
agent.node_diskio ↑	io_time (sec)
nvme0n1	20,604.96
nvme0n1	29,184.97
nvme0n1	4,642.68
nvme0n1	31,918.99
nvme0n1	29,258.26
nvme0n1	18,022.16
nvme0n1	28,483.30
nvme0n1	69,835.02
nvme0n1	15,952.78

> Aggregation

> Unit Display

Conditional Formatting [Reset](#)

If value is > (Greater than)

Warning 10000 sec

Critical 20000 sec

> Rename Column

Conditional formatting is set separately for each column. For example, you can choose one set of thresholds for a capacity column, and another set for a throughput column.

Rename Column

Renaming a column changes the displayed name on the Query results list. The new column name is also

shown in the resulting file if you export the query list to .CSV.

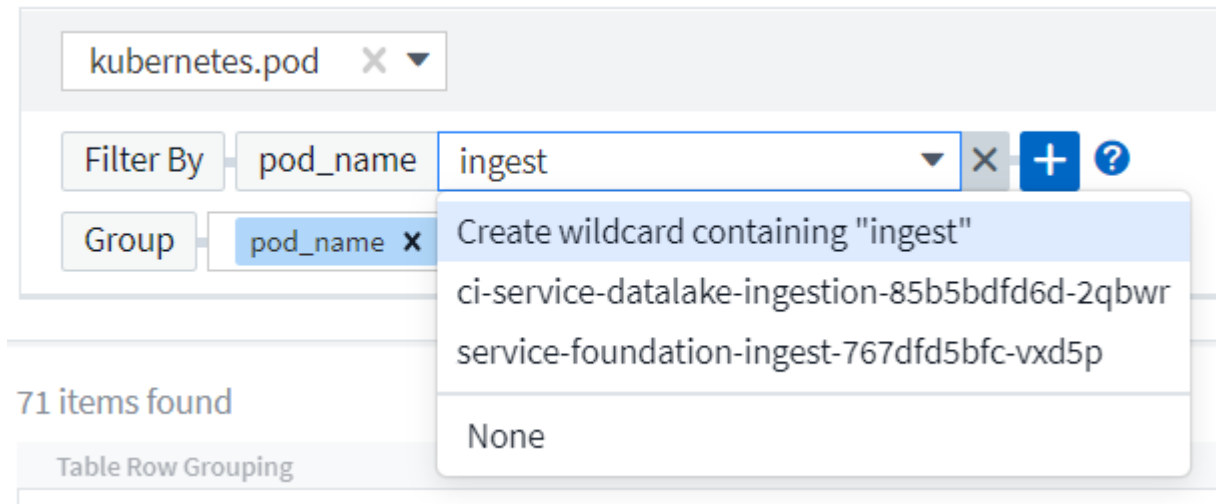
Save

After you have configured your query to show you the results you want, you can click the **Save** button to save the query for future use. Give it a meaningful and unique name.

More on Filtering

Wildcards and Expressions

When you are filtering for text or list values in queries or dashboard widgets, as you begin typing you are presented with the option to create a **wildcard filter** based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create **expressions** using NOT or OR, or you can select the "None" option to filter for null values in the field.



Filters based on wildcards or expressions (e.g. NOT, OR, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.

kubernetes.pod X ▼

Filter By

pod_name

ingest X

ci-service-audit-5f775dd975-brfdc X

X ▼

X

+

?

Group

pod_name X

X ▼

3 items found

Table Row Grouping
pod_name
ci-service-audit-5f775dd975-brfdc
ci-service-datalake-ingestion-85b5bdfd6d-2qbwr
service-foundation-ingest-767dfd5bfc-vxd5p

Note that Wildcard and Expression filtering works with text or lists but not with numerics, dates or booleans.

Refining Filters

You can use the following to refine your filter:

Filter	What it does	Example	Result
* (Asterisk)	enables you to search for everything	vol*rhel	returns all resources that start with "vol" and end with "rhel"
? (question mark)	enables you to search for a specific number of characters	BOS-PRD??-S12	returns BOS-PRD 12 -S12, BOS-PRD 23 -S12, and so on
OR	enables you to specify multiple entities	FAS2240 OR CX600 OR FAS3270	returns any of FAS2440, CX600, or FAS3270
NOT	allows you to exclude text from the search results	NOT EMC*	returns everything that does not start with "EMC"
None	searches for NULL values in all fields	None	returns results where the target field is empty
Not *	searches for NULL values in <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.

The screenshot shows a web interface for querying assets. At the top, there's a search bar with 'Volume' and a filter section with 'Filter By' and 'Name' tabs. Below this is a table titled 'Query Results (5) | 2 Selected'. The table has columns: Name, Storage Pools, Capacity - Raw (GB), and Mapped Ports. Two rows are selected, indicated by blue checkmarks in the first column. A 'Bulk Actions' dropdown menu is open on the right, showing options: Add Annotation, Remove Annotation, Add Application, and Remove Application. The table data is as follows:

Name	Storage Pools	Capacity - Raw (GB)	Mapped Ports
DmoESX_optimus:mc_Dm...	optimus-02:aggr1_optimu...	N/A	
DmoSAN_optimus:hoffma...	optimus-02:aggr1_optimu...	N/A	
DmoSAN_optimus:mc_D...	optimus-02:aggr1_optimu...	N/A	OS:windows_zu08
oci-3070-01:/vol/vfiler_lun...	oci-3070-01:aggr5	N/A	OS:windows
spectrav1:sjimmyscsi:/v...	ocinaneqa1-01:spectraaggr1	N/A	OS:linux

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 Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure Insights.
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.

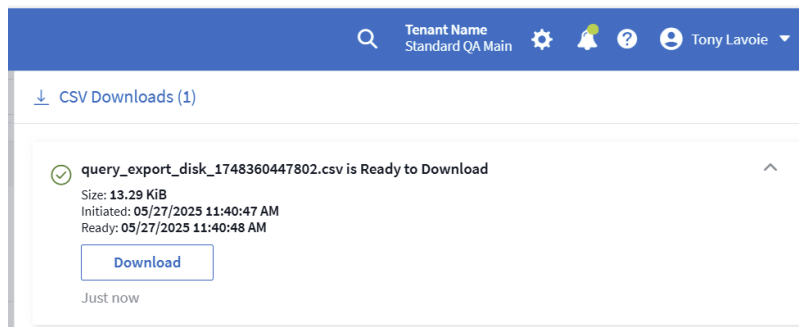


Export to .CSV is also available in the "three dots" menu in dashboard table widgets as well as most landing page tables.

Asynchronous export

Exporting data to .CSV can take anywhere from a few seconds to a number of hours, depending on the amount of data to export. Data Infrastructure Insights exports that data asynchronously, so you can continue working while the .CSV is being compiled.

View and download your .CSV exports by selecting the "Bell" icon in the upper-right toolbar.



The exported data will reflect the current filtering, columns, and column names displayed.

Commas in asset names

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.

Time format, or not time 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

You can change the criteria that are associated with a query when you want to change the search criteria for the assets that you are querying.

Modifying a Query

Steps

1. Click **Explore** and select **All Metric Queries**.

The Queries page is displayed.

2. Click the query name
3. To add a criteria to the query, click Columns icon and select a metric or attribute from the list.

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 **Explore** and select **Metric 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 Bulk Actions drop-down displays.

4. To add an application to the selected assets, click Bulk Actions 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 Bulk Actions 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.

After you click *Save* on a bulk add or *Remove* on a bulk delete action, Data Infrastructure Insights informs you that the action will take some time. You can dismiss this message; the action will continue in the background.



For environments with large amounts of related assets, inheritance of application assignments to those assets could take several minutes. Please allow more time for inheritance to occur if you have many related assets.

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.

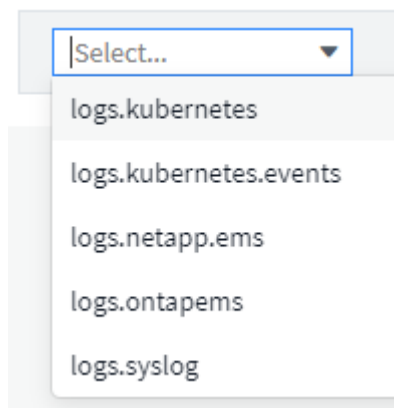
Note that only values that are links to assets can be copied using this method. Only fields that include single values (i.e. non-lists) have the copy icon.

Log Explorer

The Data Infrastructure Insights Log Explorer is a powerful tool for querying system logs. In addition to helping with investigations, you can also save a log query in a Monitor to provide alerts when those particular log triggers are activated.

To begin exploring logs, click **Log Queries > +New Log Query**.

Select an available log from the list.



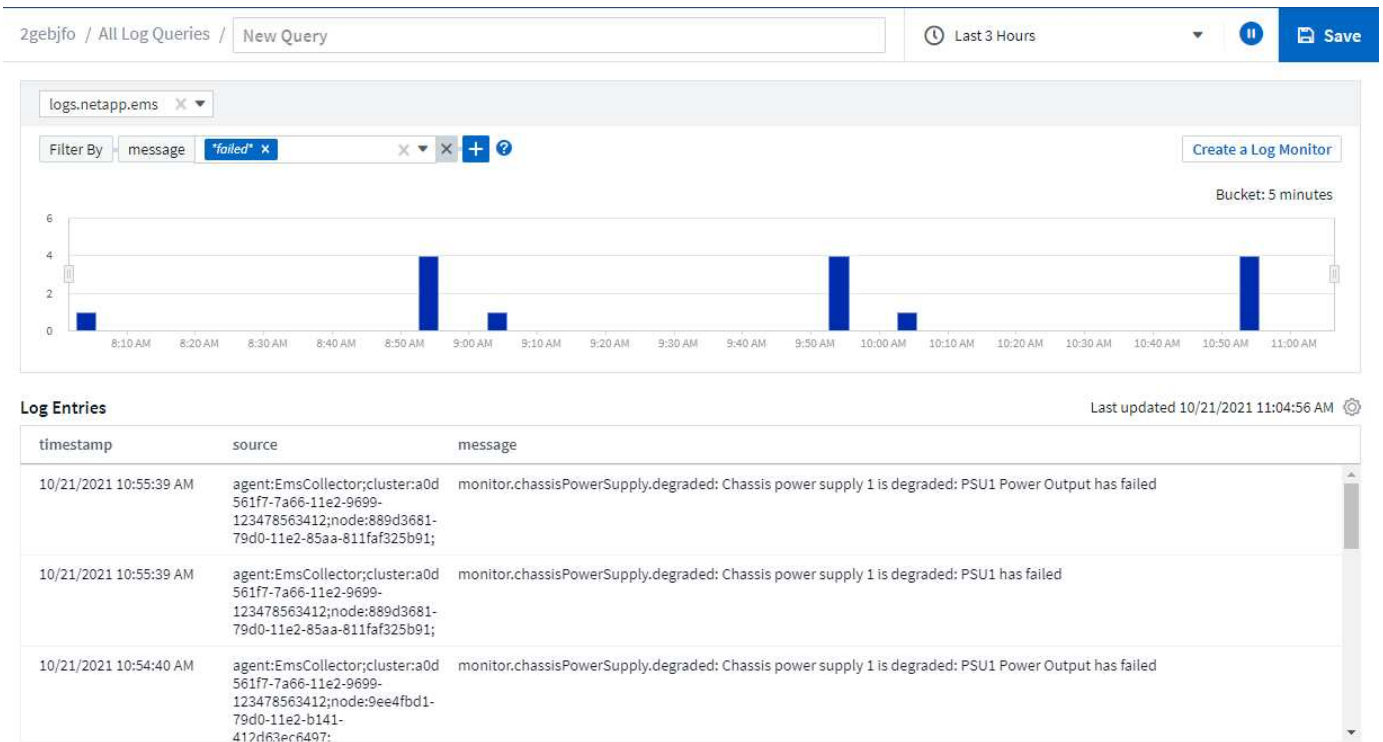
The types of logs available for querying may vary based on your environment. Additional log types may be added over time.

You can set filters to further refine the results of the query. For example, to find all log messages showing a failure, set a filter for *Messages* containing the word "failed".



You can begin typing the desired text in the filter field; Data Infrastructure Insights will prompt you to create a wildcard search containing the string as you type.

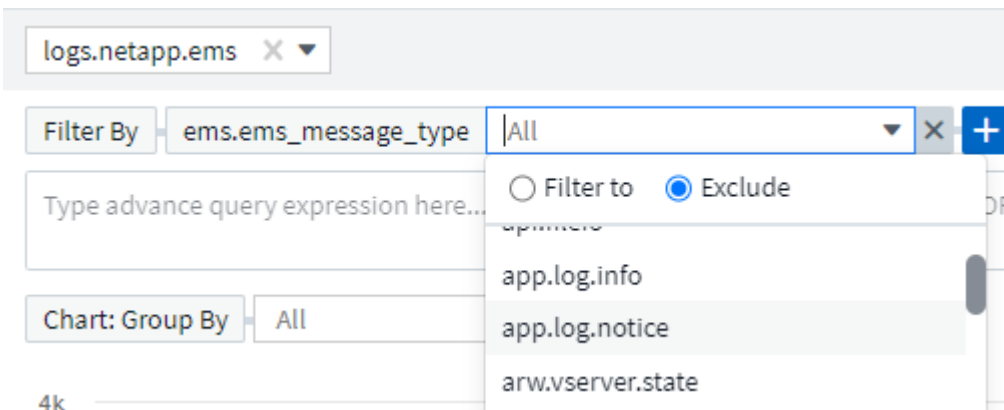
The results are displayed in a graph showing the number of log instances in each time period shown. Below the graph are the log entries themselves. The graph and the entries refresh automatically based on the selected time range.



Filtering

Include / Exclude

When filtering the logs, you can choose to **include** (i.e. "Filter to") or **exclude** the strings you type. Excluded strings are displayed in the completed filter as "NOT <string>".



Filters based on wildcards or expressions (e.g. NOT, OR, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.



At any point, you can click on *Create a Log Monitor* to create a new Monitor based on the current filter.

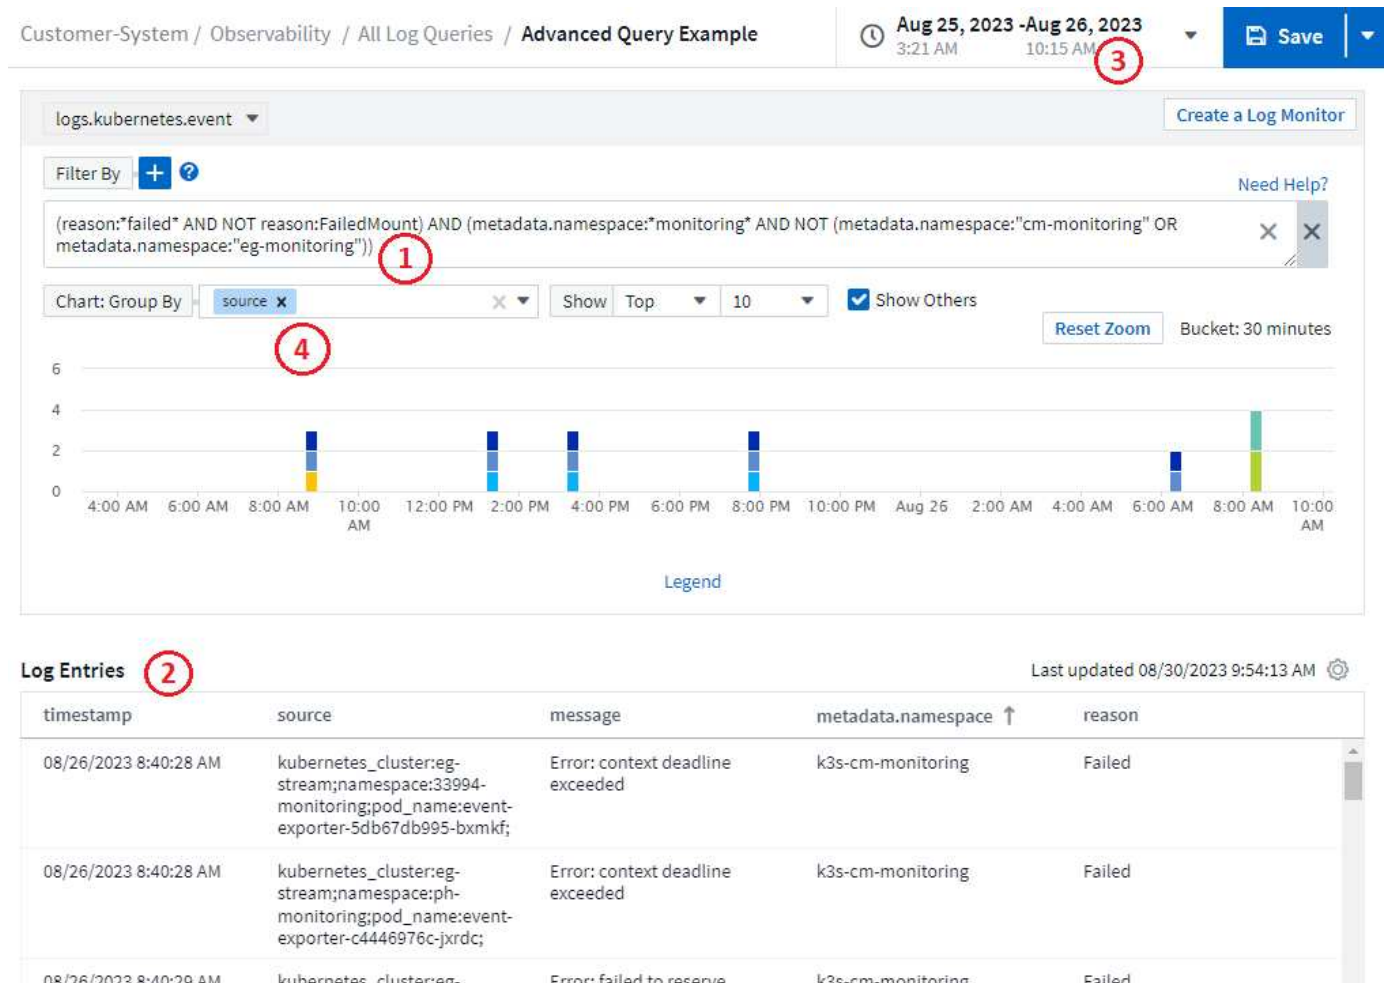
Advanced Filtering

When you are filtering for text or list values in queries or dashboard widgets, as you begin typing you are presented with the option to create a **wildcard filter** based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create expressions using NOT, AND, or OR, or you can select the "None" option to filter for null values.



Be sure to Save your query early and often as you build your filtering. Advanced Querying is "free-form" string entry, and parsing mistakes may occur as you build.

Take a look at this screen image showing filtered results for an advanced query of the *logs.kubernetes.event* log. There is a lot going on in this page, which is explained below the image:



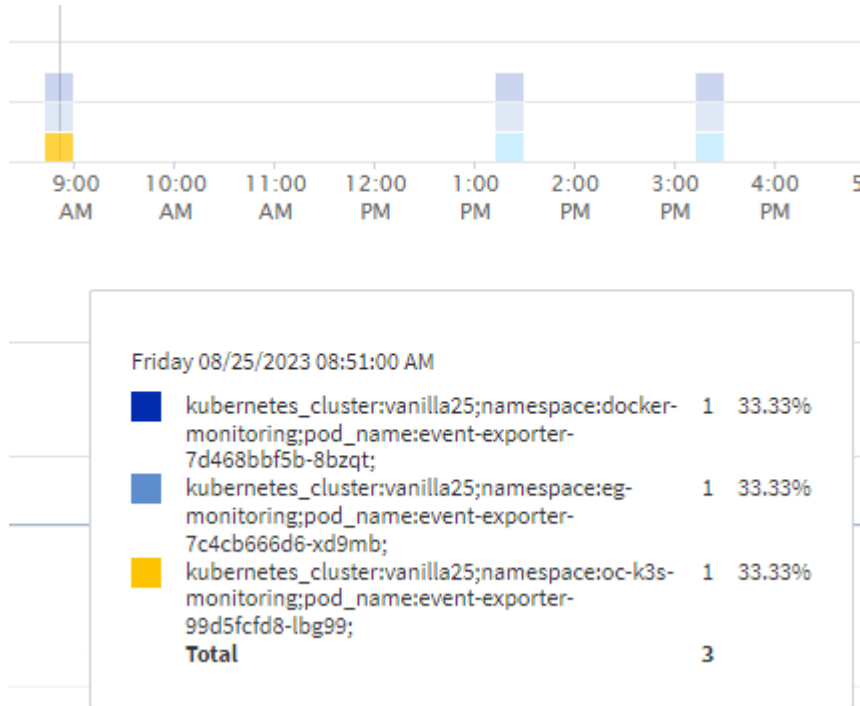
1. This advanced query string filters for the following:

- Filter for log entries with a *reason* that includes the word "failed", but not anything with the specific reason of "FailedMount".
- Include any of those entries that also include a *metadata.namespace* including the word "monitoring", but exclude the specific namespaces of "cm-monitoring" or "eg-monitoring".

Note that in the case above, since both "cm-monitoring" and "eg-monitoring" contain a dash ("-"), the strings must be included in double-quotes or a parsing error will be displayed. Strings that do not include dashes, spaces, etc. do not need to be enclosed in quotes. If in doubt, try putting the string in quotes.

- The results of the current filter, including any "Filter By" values AND the Advanced Query filter, are displayed in the results list. The list can be sorted by any displayed columns. To display additional columns, select the "gear" icon.
- The graph has been zoomed in to show only log results that occurred within a specific time frame. The time range shown here reflects the current zoom level. Select the *Reset Zoom* button to set the zoom level back to the current Data Infrastructure Insights time range.

4. The chart results have been Grouped By the *source* field. The chart shows results in each column grouped into colors. Hovering over a column in the chart will display some details about the specific entries.



Refining Filters

You can use the following to refine your filter:

Filter	What it does
* (Asterisk)	enables you to search for everything
? (question mark)	enables you to search for a specific number of characters
OR	enables you to specify multiple entities
NOT	allows you to exclude text from the search results
None	searches for NULL values in all fields
Not *	searches for NULL values in <i>text-only</i> fields

If you enclose a filter string in double quotes, Insight treats everything between the first and last quote as an exact match. Any special characters or operators inside the quotes will be treated as literals. For example, filtering for "*" will return results that are a literal asterisk; the asterisk will not be treated as a wildcard in this case. The operators OR and NOT will also be treated as literal strings when enclosed in double quotes.

You can combine a simple filter with an advanced query filter; the resulting filter is an "AND" of the two.

The Chart Legend

The *Legend* below the chart has a few surprises as well. For each result (based on the current filter) shown in the Legend, you have an option to display only results for that line (Add Filter), or to display any results NOT for that line (Add Exclude Filter). The chart and the Log Entries list update to show results based on your

selection. To remove this filtering, open the Legend again and select the [X] to clear the Legend-based filter.

Legend

<div><div></div><div>kubernetes_cluster:vanila25;namespace:docker-monitoring;pod_name:event-exporter-7d468bbf5b-8bzqt;</div></div>	<div><div></div><div></div></div>	527.78%
<div><div></div><div>kubernetes_cluster:vanila25;namespace:eg-monitoring;pod_name:event-exporter-7c4cb666d6-xd9mb;</div></div>	<div><div></div><div></div></div>	527.78%
<div><div></div><div>kubernetes_cluster:vanila25;namespace:oc-k3s-monitoring;pod_name:event-exporter-</div></div>	<div><div></div><div></div></div>	316.67%

Log Details

Clicking anywhere in a log entry in the list will open a detail pane for that entry. Here you can explore more information about the event.

Click on "Add Filter" to add the selected field to the current filter. The log entry list will update based on the new filter.

Note that some fields cannot be added as filters; in those cases, the *Add Filter* icon is not available.

Log Details



timestamp

09/20/2021 9:03:36 PM

message

2021-09-20T15:33:36Z E! [processors.execd] stderr: "Total time to process mountstats file: /hostfs/proc/1/mountstats, was: 0s"

id: 227814532095936770

node_name: ci-auto-dsacq-insights-1.cloudinsights-dev.netapp.com

Add Filter



source: telegraf-ds-dfcc5

type: logs.kubernetes

kubernetes

kubernetes.annotations.openshift.io_scc: telegraf-hostaccess

kubernetes.container_hash: ci-registry.nane.openenglab.netapp.com:8077/telegraf@sha256:00b45a7cc0761c

Troubleshooting

Here you will find suggestions for troubleshooting problems with Log Queries.

Problem:	Try this:
I don't see "debug" messages in my log query	Debug log messaging is not collected. To capture messages you want, change the relevant message severity to <i>informational</i> , <i>error</i> , <i>alert</i> , <i>emergency</i> , or <i>notice</i> level.

Identifying inactive devices

Identifying the assets you have and who's using them is critical to "right-sizing" and freeing up unused infrastructure. You can easily reallocate or decommission underused resources and avoid unnecessary purchases.

Use the following steps to identify inactive assets.

Steps

- Navigate to **Observability > Explore** → **+New Metric Query**.
- Select *Storage* from the drop down.

- Click the gear and add *isActive* as a column.

Rows showing a Check are active. "X" indicates inactive devices.

To remove inactive devices, simply select the devices to remove and in the **Bulk Actions** drop-down, select *Delete Inactive Devices*.

Insights

Insights

Insights allow you to look into things like resource usage and how it affects other resources, or time-to-full analyses.

A number of Insights are available. Navigate to **Dashboards > Insights** to start diving in. You can view active Insights (Insights that are currently occurring) on the main tab, or inactive Insights on the *Inactive Insights* tab. Inactive Insights are those that were previously active but are no longer occurring.

Insight Types

Shared Resources Under Stress

High-impact workloads can reduce the performance of other workloads in a shared resource. This puts the shared resource under stress. Data Infrastructure Insights provides tools to help you investigate resource saturation and impact on your tenant. [Learn More](#)

Kubernetes Namespaces Running Out of Space

The Kubernetes Namespaces Running Out of Space Insight gives you a view into workloads on your Kubernetes namespaces that are at risk of running out of space, with an estimate for the number of days remaining before each space becomes full. [Learn More](#)

Reclaim ONTAP Cold Storage

The *Reclaim ONTAP Cold Storage* Insight provides data about cold capacity, potential cost/power savings and recommended action items for volumes on ONTAP systems. [Learn More](#)



This is a *Preview* feature and may change over time as improvements are made. [Learn more](#) about Data Infrastructure Insights Preview features.

Insights: Shared Resources Under Stress

High-impact workloads can reduce the performance of other workloads in a shared resource. This puts the shared resource under stress. Data Infrastructure Insights provides tools to help you investigate resource saturation and impact on your tenant.

Terminology

When talking about workload or resource impact, the following definitions are useful.

A **Demanding Workload** is a workload that is currently identified as impacting other resources in the shared

storage pool. These workloads drive higher IOPS (for example), reducing IOPS in the Impacted Workloads. Demanding workloads are sometimes called *high-consuming workloads*.

An **Impacted Workload** is a workload that is affected by a high-consuming workload in the shared Storage Pool. These workloads are experiencing reduced IOPS and/or higher latency, caused by the Demanding Workloads.

Note that if Data Infrastructure Insights has not discovered the leading compute workload, the volume or internal volume itself will be recognized as the workload. This applies to both demanding and impacted workloads.

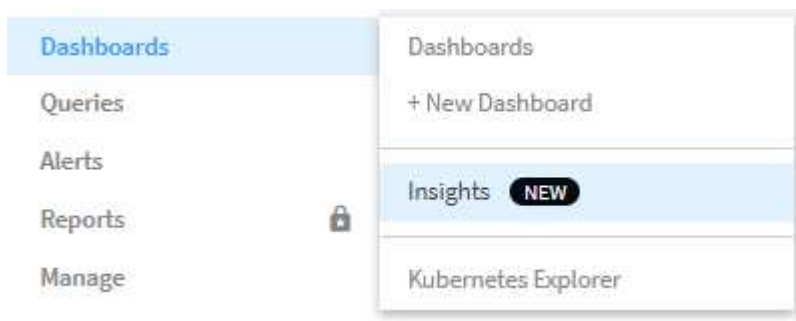
Shared Resource Saturation is the ratio of impacting IOPS to *baseline*.

Baseline is defined as the maximum reported data point for each workload in the hour immediately preceding the detected saturation.

A **Contention** or **Saturation** occurs when IOPS are determined to be affecting other resources or workloads in the shared storage pool.


Demanding Workloads

To start looking into Demanding and impacted workloads in your shared resources, click on **Dashboards > Insights** and select the **Shared Resources Under Stress** Insight.



Data Infrastructure Insights displays a list of any workloads where a saturation has been detected. Note that Data Infrastructure Insights will show workloads where at least one *demanding resource* **or** *impacted resource* has been detected.

Click on a workload to view the details page for it. The top chart shows the activity on the shared resource (for example, a storage pool) on which the contention/saturation is occurring.




Below that are two charts showing the *demanding* workloads and the workloads that are *impacted* by those demanding workloads.

Demanding Workloads (1) ⓘ

Potentially impacted the shared resource and other related workloads

Contributing IOPS ▾

lops.total (IOPS)




Workload	Current Contributing IOPS (IOPS) ↓	Change Since Detection (IOPS)
 internal-volume-331	500.00	+190.00

Impacted Workloads (1) ⓘ

Impacted by changed workloads on the shared resource

Latency ▾

latency.total (ms)



Workload	Current Latency (ms) ↓	Change Since Detection (ms)
 internal-volume-332	200.00	+110.00

Below each table is a list of workloads and/or resources affecting or affected by the contention. Clicking on a resource (for example, a VM) opens a detail page for that resource. Clicking on a workload opens a query page showing the pods involved. Note that if the link opens an empty query, it may be because the affected pod is no longer part of the active contention. You can modify the query's time range to view the pod list in greater or more focused time range.

What do I do to resolve saturation?

There are a number of steps you can take to reduce or eliminate the chance of saturation on your tenant. These are shown by expanding the **+Show Recommendations** link on the page. Here are a few things you can try.

- Move high-IOPS consumers

Move the "greedy" workloads to less-saturated Storage Pools. It is recommended to assess the tier and capacity of these pools before moving the workloads, to avoid unnecessary costs or additional contentions.

- Implement a quality of service (QoS) policy

Implementing a QoS policy per workload to ensure enough free resources available will alleviate saturation on the Storage Pool. This is a long-term solution.

- Add additional resources

If the shared resource (for example, Storage Pool) has reached the IOPS saturation point, adding more or faster disks to the pool will ensure enough free resources available to alleviate saturation.

Finally, you can click the **Copy Insight Link** to copy the page url to the clipboard, to more easily share with colleagues.

Insights: Kubernetes Namespaces Running out of Space

Running out of space on your tenant is never a good situation. Data Infrastructure Insights helps you predict the time you have before Kubernetes persistent volumes become full.

The *Kubernetes Namespaces Running Out of Space* Insight gives you a view into workloads on your Kubernetes namespaces that are at risk of running out of space, with an estimate for the number of days remaining before each persistent volume becomes full.

You can view this Insight by navigating to **Dashboards > Insights**.

Kubernetes Namespaces Running Out of Space (3)

Description	Estimated Days to Full	Workloads at Risk	Detected ↓
1 workload at risk on es	35	1	2 days ago
1 workload at risk on manager	24	1	2 days ago
2 workloads at risk on cloudinsights	1	2	2 days ago

Click on a workload to open a detail page for the Insight. On this page you will see a graph showing the workload capacity trends as well as a table showing the following:

- Workload Name
- Persistent Volume affected
- Predicted Time-to-Full in days
- Persistent Volume capacity
- Backend Storage Resource affected, with current capacity used out of total capacity. Clicking this link will open the detailed landing page for the backend volume.

Workloads at risk (2)

Workloads	Persistent Volume (pvClaim)	Time to Full (Days) ↓	Persistent Volume Capacity (GiB)	Backend Storage Resource (Capacity Used)
multi (1)	pv1 (pvc1)	1	4.00	internal-volume-601 60.00% (3.00/5.00 GiB)
taskmanager (1)	pv1 (pvc1)	1	4.00	internal-volume-601 60.00% (3.00/5.00 GiB)

What can I do if I'm running out of space?


On the Insight page, click the **+Show Recommendations** to view possible solutions. The easiest option when running out of space is always to add more capacity, and Data Infrastructure Insights shows you the optimal capacity to add to increase time-to-full to a target 60-day prediction. Other recommendations are also shown.

Show Recommendations

- 1 Get time to full back up to 60 days by adding more capacity to backend resources
Add to the following resources to bring time-to-full up to ideal capacity.

Backend Resource ↓	Current Capacity (time to full)	Recommended Capacity to Add		Ideal Capacity (time to full)
internal-volume-601	2.00 GiB 1 Days	+	518.79 GiB	= 520.79 GiB 60 Days

- 2 Use NetApp Astra Trident with your K8s to automatically grow capacity
Astra Trident can keep your capacity lean without risk of running out of space.

[Learn more about](#)  **Astra Trident**

 [Copy Insight Link](#)

It is here also that you can copy a convenient link to this Insight, to bookmark the page or to easily share with your team.

Insights: Reclaim ONTAP Cold Storage

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

To view these Insights, navigate to **Dashboards > Insights** and take a look at the *Reclaim ONTAP Cold Storage* Insight. Note that this Insight will only list affected storages if Data Infrastructure Insights has detected cold storage, otherwise you will see an "all clear" message.

Keep in mind that cold data less than 30 days old is not shown.

Reclaim ONTAP Cold Storage (3)

Description	Cold data storage(TiB)	Workloads with cold data	Detected ↓
0.30 TiB of cold data on storage rtp-sa-cl04	0.30	45	an hour ago
1.22 TiB of cold data on storage umeng-aff300-01-02	1.22	84	16 days ago
11.62 TiB of cold data on storage rtp-sa-cl01	11.62	171	16 days ago

The Insight description gives a quick indication of the amount of data detected as "cold" and which storage that data resides on. The table also provides a count of workloads with cold data.

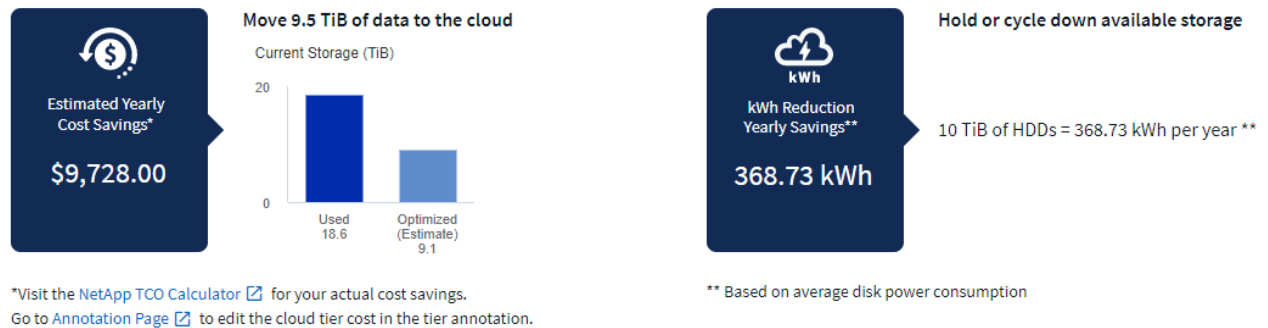
Selecting an Insight from the list opens a page showing more details, including recommendations to move data to the Cloud or cycle down unused disks, as well as estimated cost and power savings you could potentially realize from implementing those recommendations. The page even provides a handy link to [NetApp's TCO Calculator](#) so you can experiment with the numbers.



150 Workloads on storage **rtp-sa-cl01** contains a total of 9.5 TiB of cold data.

You could lower costs 9.3% a year and reduce your carbon footprint by moving cold storage to the cloud.

Detected: 2 months ago, 9:21 AM
(ACTIVE)
May 19, 2023 10:05AM



Recommendations

On the Insight page, expand the **Recommendations** to explore the following options:

- Move unused workloads (zombies) to a lower cost storage tier (HDD)

Utilizing the zombie flag, cold storage and number of days, find the coldest and largest amount of data and move the workload to a lower cost storage tier (such as a storage pool using hard disk storage). A workload is considered a "zombie" when it has not received any significant IO requests for 30 days or more.

- Delete unused workloads

Verify which workloads are not in use and consider archiving them or remove them from the storage system.

- Consider NetApp's Fabric Pool Solution


NetApp's [Fabric Pool Solution](#) automatically tiers cold data to low cost cloud storage, thus increasing the efficiency of your performance tier as well as providing remote data protection.

Visualize and Explore


The graphs and table provide additional trending information as well as allow you to drill into the individual workloads.

Cluster Cold Storage Trend [Show Details](#)

Cold Data (TiB)



Cold Storage by Days Cold (TiB)



Workloads with cold data (150) [View all workloads](#)

Filter...

Workloads	# Days cold	↑	Total Size (GiB)	Cold Data Size (GiB)	Percent Cold (%)	Is Zombie	i Disk Type
SelectPool	31		8,192.00	1,714.21	20.93	N A	SAS
nj_UCS_VMw_Infrastructure	31		5,120.00	934.74	18.26	N A	SAS
Oracle_SAP_DS_220	31		2,048.00	861.97	42.09	N A	SSD
rtp_sa_workspace	31		13,000.00	741.32	5.70	N A	SAS
vc220_migrate	31		4,311.58	685.30	15.89	N A	SAS
H01_shared	31		998.25	646.55	64.77	N A	SSD
ProdSelectPool	31		8,192.00	555.30	6.78	N A	SAS
vcenter_migrate	31		6,144.00	475.99	7.75	N A	SAS
rtp_sa_mgmt_apps	31		4,096.00	449.26	10.97	N A	SAS
SOFTWARE	31		600.00	365.54	60.92	N A	SAS
DP_Migrate	31		7,168.00	347.20	4.84	N A	SAS

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 all Data Infrastructure Insights Editions, however, Basic Edition is subject to the following:

- * You may only have up to five custom monitors active at a time. Any monitors beyond five will be created in or moved to *Paused* state.
- * VMDK, Virtual Machine, Host, and DataStore metrics monitors are not supported. If you have monitors created for these metrics, they will be paused and cannot be resumed when downgrading to Basic Edition.

Monitors allow you to set thresholds on metrics generated by "infrastructure" objects such as storage, VM, EC2, and ports, as well as for "integration" data such as those collected for Kubernetes, ONTAP advanced metrics, and Telegraf plugins. These *metric* monitors alert you when warning-level or critical-level thresholds are crossed.

You can also create monitors to trigger warning-, critical-, or informational-level alerts when specified *log events* are detected.

Data Infrastructure Insights provides a number of [System-Defined Monitors](#) as well, based on your environment.

Security Best Practice

Data Infrastructure Insights alerts are designed to highlight data points and trends on your tenant, and Data Infrastructure Insights allows you to enter any valid email address as an alert recipient. If you are working in a secure environment, be especially mindful of who is receiving the notification or otherwise has access to the alert.

Metric or Log Monitor?

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

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

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



When you add a new monitor, you are prompted to create a Metric Monitor or a Log Monitor.

- *Metric* monitors alert on infrastructure- or performance-related triggers
- *Log* monitors alert on log-related activity


After you choose your monitor type, the Monitor Configuration dialog is displayed. Configuration varies depending on which type of monitor you are creating.

Metric Monitor

1. In the drop-down, search for and choose an object type and metric to monitor.

You can set filters to narrow down which object attributes or metrics to monitor.

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 for our example. 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.



Define the alert resolution behavior

You can choose how a metric monitor alert is resolved. You are presented with two choices:

- **Resolve when the metric returns to the acceptable range.**
- **Resolve when the metric is within the acceptable range for a specified amount of time, from 1 minute to 7 days.**

Log Monitor

When creating a **Log monitor**, first choose which log to monitor from the available log list. You can then filter based on the available attributes as above. You can also choose one or more "Group By" attributes.



The Log Monitor filter cannot be empty.

Define the alert Behavior

You can create the monitor to alert with a severity level of *Critical*, *Warning*, or *Informational*, when the conditions you defined above occur once (i.e. immediately), or wait to alert until the conditions occur 2 times or more.

Define the alert resolution behavior

You can choose how a log monitor alert is resolved. You are presented with three choices:

- **Resolve instantly:** The alert is immediately resolved with no further action needed
- **Resolve based on time:** The alert is resolved after the specified time has passed
- **Resolve based on log entry:** The alert is resolved when a subsequent log activity has occurred. For example, when an object is logged as "available".

- ☐ Resolve instantly
- ☐ Resolve based on time
- ☒ Resolve based on log entry

Log Source logs.netapp.ems ▼

Filter By ems.ems_message_type "object.store.available" x x ▼ x +

Anomaly Detection Monitor

1. In the drop-down, search for and choose an object type and metric to monitor.

You can set filters to narrow down which object attributes or metrics to monitor.

1 Select a metric anomaly to monitor

Object Storage x ▼ Metric iops.total x ▼

Filter by Attribute + ?

Filter by Metric + ?

Group by Storage ▼

Unit Displayed In Whole Number ▼

Define the Conditions of the Monitor.

1. After choosing the object and metric to monitor, you set the conditions under which an anomaly is detected.
 - Choose whether to detect an anomaly when the chosen metric **spikes above** the predicted bounds, **drops below** those bounds, or **spikes above or drops below** the bounds.
 - Set the **sensitivity** of detection. **Low** (fewer anomalies are detected), **Medium**, or **High** (more anomalies are detected).
 - Set the alerts to be wither **Warning** or **Critical**.
 - If desired, you can choose to reduce noise, ignoring anomalies when the chosen metric is below a threshold that you set.

2 Define the monitor's conditions

Trigger alert when **performance.iops.total** Spikes above ▼ the predicted bounds.

Set sensitivity: Low (detect fewer anomalies) ▼

Alert severity: Critical ▼

To reduce noise, ignore anomalies when **performance.iops.total** is below Optional IO/s

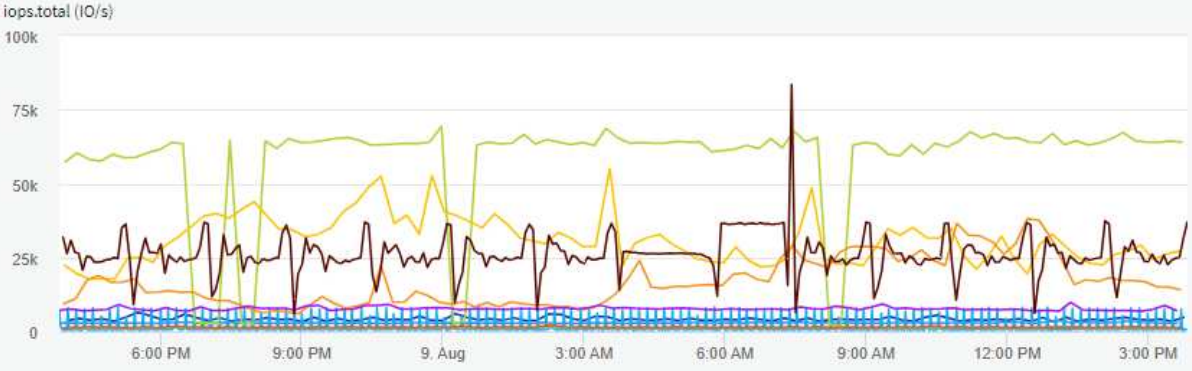


Chart Displaying Top ▼ 10 ▼ Over the Last 24 Hours ▼

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)



Add Delivery Method ▼

- Email
- 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 <div>Critical, Resolved</div> <div><input checked="" type="checkbox"/> Critical</div> <div><input type="checkbox"/> Warning</div> <div><input checked="" type="checkbox"/> Resolved</div>	Add Recipients (Required) <div>user_1@email.com X</div> <div>user_2@email.com X</div>
 Email	Notify team on <div>Warning</div>	Add Recipients (Required) <div>user_3@email.com X</div>

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	Notify team on	Slack	Use Webhook(s)
	<div>Critical</div>		<div>Slack X Teams X</div>
	Notify team on		Use Webhook(s)
	<div>Resolved</div>		<div>Slack X Teams X</div>
	Notify team on		Use Webhook(s)
	<div>Warning</div>		<div>Slack X Teams X</div>



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

Setting Corrective Actions or Additional Information

You can add an optional description as well as additional insights and/or corrective actions by filling in the **Add an Alert Description** section. The description can be up to 1024 characters and will be sent with the alert. The insights/corrective action field can be up to 67,000 characters and will be displayed in the summary section of the alert landing page.

In these fields you can provide notes, links, or steps to take to correct or otherwise address the alert.

You can add any object attribute (for example, storage name) as a parameter to an alert description. For example, you can set parameters for volume name and storage name in a description like: "High Latency for Volume: `%%relatedObject.volume.name%%`, Storage: `%%relatedObject.storage.name%%`".

4 Add an alert description (optional)

Add a description	<input type="text" value="Enter a description that will be sent with this alert (1024 character limit)"/>
Add insights and corrective actions	<input type="text" value="Enter a url or details about the suggested actions to fix the issue raised by the alert"/>

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

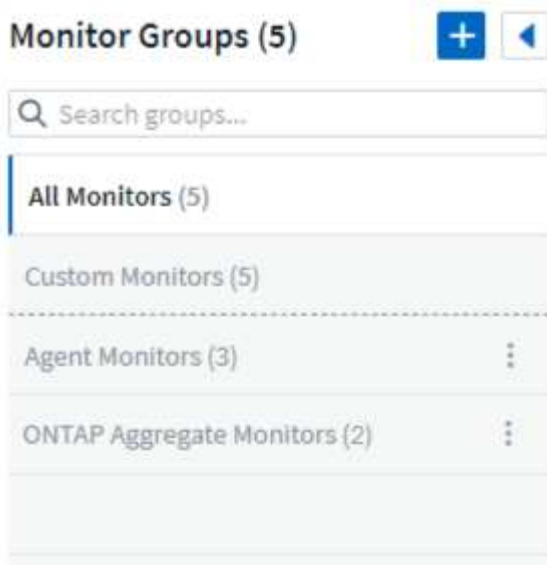
You can choose to temporarily pause monitoring of an object type by clicking the menu to the right of the monitor and selecting **Pause**. When you are ready to resume monitoring, click **Resume**.

You can copy a monitor by selecting **Duplicate** from the menu. You can then modify the new monitor and change the object/metric, filter, conditions, email recipients, etc.

If a monitor is no longer needed, you can delete it by selecting **Delete** from the menu.

Monitor Groups

Grouping allows you to view and manage related monitors. For example, you can have a monitor group dedicated to the storage on your tenant, or monitors relevant to a certain recipient list.



The following monitor groups are shown. The number of monitors contained in a group is shown next to the group name.

- **All Monitors** lists all monitors.
- **Custom Monitors** lists all user-created monitors.
- **Suspended Monitors** will list any system monitors that have been suspended by Data Infrastructure Insights.
- Data Infrastructure Insights will also show a number of **System Monitor Groups**, which will list one or more groups of [system-defined monitors](#), including ONTAP Infrastructure and Workload monitors.



Custom monitors can be paused, resumed, deleted, or moved to another group. System-defined monitors can be paused and resumed but can not be deleted or moved.

Suspended Monitors

This group will only be shown if Data Infrastructure Insights has suspended one or more monitors. A monitor may be suspended if it is generating excessive or continuous alerts. If the monitor is a custom monitor, modify the conditions to prevent the continuous alerting, and then resume the monitor. The monitor will be removed from the Suspended Monitors group when the issue causing the suspension is resolved.

System-Defined Monitors

These groups will show monitors provided by Data Infrastructure Insights, as long as your environment contains the devices and/or log availability required by the monitors.

System-Defined monitors cannot be modified, moved to another group, or deleted. However, you can duplicate a system monitor and modify or move the duplicate.

System monitors may include monitors for ONTAP Infrastructure (storage, volume, etc.) or Workloads (i.e. log monitors), or other groups. NetApp is constantly evaluating customer need and product functionality, and will update or add to system monitors and groups as needed.

Custom Monitor Groups

You can create your own groups to contain monitors based on your needs. For example, you may want a

group for all of your storage-related monitors.

To create a new custom monitor group, click the **"+" Create New Monitor Group** button. Enter a name for the group and click **Create Group**. An empty group is created with that name.

To add monitors to the group, go to the *All Monitors* group (recommended) and do one of the following:

- To add a single monitor, click the menu to the right of the monitor and select *Add to Group*. Choose the group to which to add the monitor.
- Click on the monitor name to open the monitor's edit view, and select a group in the *Associate to a monitor group* section.

5 Associate to a monitor group (optional)

A screenshot of a web interface showing a dropdown menu. The dropdown is open, displaying the text 'ONTAP Monitors' and a small downward-pointing arrow on the right side of the menu box.

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 Data Infrastructure 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*.

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 Data Infrastructure Insights; they are still available in *All Monitors*.



System-Defined Monitors

Data Infrastructure Insights includes a number of system-defined monitors for both metrics and logs. The system monitors available are dependent on the data collectors present on your tenant. Because of that, the monitors available in Data Infrastructure Insights may change as data collectors are added or their configurations changed.

View the [System-Defined Monitors](#) page for descriptions of monitors included with Data Infrastructure Insights.

More Information

- [Viewing and Dismissing Alerts](#)

Viewing and Managing Alerts from Monitors

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



Monitors and Alerting is available in Data Infrastructure Insights Standard Edition and higher.

Viewing and Managing Alerts

To view and manage alerts, do the following.

1. Navigate to the **Alerts > All 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
- **Has Corrective Action**: The alert has suggested corrective actions. Open the alert page to view these.

You can manage an alert by clicking the menu to the right of the alert and choosing one of the following:

- **In Process** to indicate that the alert is under investigation or otherwise needs to be kept open
- **Dismiss** to remove the alert from the list of active alerts.

You can manage multiple alerts by selecting the checkbox to the left of each Alert and clicking *Change Selected Alerts Status*.

Clicking on an Alert ID opens the Alert Detail Page.

Alert Detail Panel

Select any alert row to open the alert's detail panel. The alert detail panel 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.

Metric Alert

Jun 3, 2025
9:29 AM - 10:47 AM



Critical Alert AL-14930837 ACTIVE [Collapse Details](#)

Triggered On

Storage:
 CI-GDL1-Ontap-fas8080

Details

Top Severity: Critical
Condition: **Average iops.total** is > (greater than) 1,700 IO/s and/or 2,000 IO/s all the time in 15-minute window.

Monitor

altimeout

Attributes

Filters Applied: N/A

Description

No Description Provided

Resolution conditions

Resolve when metric is within acceptable range for 10 mins

Status

New


Time

Triggered time: Jun 3, 2025 10:44 AM Duration: 17m (Active)

Alert Summary

[Alert Attributes](#)

Jun 03, 2025 09:29 AM - 10:47 AM [Settings](#)



Close

Alerts When Data Is Missing

In a realtime system such as Data Infrastructure Insights, to trigger the analysis of a Monitor to decide if an Alert should be generated, we rely on one of two things:

- the next datapoint to arrive
- a timer to fire when there is no datapoint and you have waited long enough

As is the case with slow data arrival—or no data arrival—the timer mechanism needs to take over as the data arrival rate is insufficient to trigger alerts in "real time." So the question typically becomes "How long do I wait before I close the analysis window and look at what I have?" If you wait too long then you are not generating the alerts fast enough to be useful.

If you have a Monitor with a 30-minute window that notices that a condition is violated by the last data point before a long-term loss-of-data, an Alert will be generated because the Monitor received no other information

to use to confirm a recovery of the metric or notice that the condition persisted.

"Permanently Active" Alerts

It is possible to configure a monitor in such a way for the condition to **always** exist on the monitored object—for example, IOPS > 1 or latency > 0. These are often created as 'test' monitors and then forgotten. Such monitors create alerts that stay permanently open on the constituent objects, which can cause system stress and stability issues over time.

To prevent this, Data Infrastructure Insights will automatically close any "permanently active" alert after 7 days. Note that the underlying monitor conditions may (probably will) continue to exist, causing a new alert to be issued almost immediately, but this closing of "always active" alerts alleviates some of the system stress that can otherwise occur.

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 and select the *Email* tab.

Subscription Notification Recipients

Send subscription related notifications to the following:

- ☒ All Account Owners
- ☒ All Monitor & Optimize Administrators
- ☒ Additional Email Addresses

name@email.com X

Save

Global Monitor Notification Recipients

Default email recipients for monitor related notifications:

- ☐ All Account Owners
- ☒ All Monitor & Optimize Administrators
- ☐ Additional Email Addresses

Save

Subscription Notification Recipients

To configure recipients for subscription-related event notifications, go to the "Subscription Notification Recipients" section.

You can choose to have email notifications sent for subscription-related events to any or all of the following recipients:

- All Account Owners
- All *Monitor & Optimize* Administrators

- Additional Email Addresses that you specify

The following are examples of the types of notifications that might be sent, and user actions you can take.

Notification:	User Action:
Trial or subscription has been updated	Review subscription details on the 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



To ensure your recipients receive notifications from Data Infrastructure Insights, add the following email addresses to any "allow" lists:

- accounts@service.cloudinsights.netapp.com
- DoNotReply@cloudinsights.netapp.com

Global Recipient List for Alerts

Email notifications of alerts are sent to the alert recipient list for every action on the alert. You can choose to send alert notifications to a global recipient list.

To configure global alert recipients, choose the desired recipients in the **Global Monitor Notification Recipients** section.

You can always override the global recipients list for an individual monitor when creating or modifying the monitor.



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

Editing Notifications for ONTAP

You can modify notifications for ONTAP clusters by selecting *Edit Notifications* from the upper-right drop-down on a Storage landing page.

Edit

Poll Again

Postpone 3 Days

Postpone 7 Days

Postpone 30 Days

Edit Notifications

Delete

From here, you can set notifications for Critical, Warning, Informational, and/or Resolved alerts. Each scenario can notify the Global Recipient list or other recipients you choose.

Edit Notifications

☒ By Email

Notify team on

Critical, Warn... ▾

Send to

☐ Global Monitor Recipient List

☒ Other Email Recipients

email@email.one ✕

email2@email2.two ✕ |

Notify team on

Resolved ▾

Send to

☒ Global Monitor Recipient List

☐ Other Email Recipients


☐ By Webhook

Enable webhook notification to add recipients

Anomaly Detection Monitors

Anomaly Detection provides insight into unexpected changes in the patterns of data on your tenant. An anomaly occurs when the pattern of an object's behavior changes, for example, if an object experiences a certain level of latency at a certain time on Wednesdays, but latency spikes above that level at that time on the subsequent Wednesday, that spike would be considered an anomaly. Data Infrastructure Insights allows the creation of monitors to alert when anomalies such as this occur.

Anomaly detection is suitable for object metrics that exhibit a recurring, predictable pattern. When these object metrics spike above or drop below their expected levels, Data Infrastructure Insights can generate an alert to prompt investigation.



What is Anomaly Detection?

An anomaly occurs when the mean value of a metric is a number of standard deviations away from the weighted mean of that metric for the previous few weeks, with recent weeks having more weight than previous weeks. Data Infrastructure Insights provides the ability to monitor data and alert when anomalies are detected. You have a choice to set the "sensitivity" levels of detection. For example, a higher sensitivity would be when the mean value is fewer standard deviations from the mean, thus causing more alerts to be generated. Conversely, lower sensitivity = more standard deviations from mean = fewer alerts.

Anomaly Detection monitoring differs from Threshold Monitoring.

- **Threshold-based monitoring** works when you have pre-defined thresholds for specific metrics. In other words, when you have a clear understanding of what is expected (i.e. within a normal range).

Metric Monitor

Set the high and low parameters that will trigger an alert if exceeded



Use when you know the upper and lower operating range

- **Anomaly Detection monitoring** uses machine learning algorithms to identify outliers that deviate from the norm, for when the definition of "normal" is not clear.

Anomaly Detection Monitor

Detect and be alerted to abnormal performance changes



Use when you want to trigger alerts against performance spikes and drops

When would I need Anomaly Detection?

Anomaly Detection monitoring can provide helpful alerts for many situations, including the following:


- When the definition of *normal* is unclear. For example, SAN error rates may be expected in varying amounts depending on port. Alerting on one error is noisy and unnecessary, but a sudden or significant increase could indicate a widespread issue.
- Where there are changes over time. Workloads that exhibit seasonality (i.e. they are busy or quiet at certain times). This could include unexpected quiet periods that may indicate a batch stall.
- Working with large amounts of data where manually defining and adjusting thresholds is impractical. For example, a tenant with a large numbers of hosts and/or volumes with varying workloads. Each may have different SLAs, so understanding the ones that exceed the norm is important.

Creating an Anomaly Detection Monitor

To alert on anomalies, create a monitor by navigating to **Observability > Alerts > +Monitor**. Select *Anomaly Detection Monitor* as the monitor type.

Metric Monitor

Set the high and low parameters that will trigger an alert if exceeded



Use when you know the upper and lower operating range

Log Monitor


Monitor logs and configure alerts



Use when you want to trigger alerts in response to log activity

Anomaly Detection Monitor

Detect and be alerted to abnormal performance changes



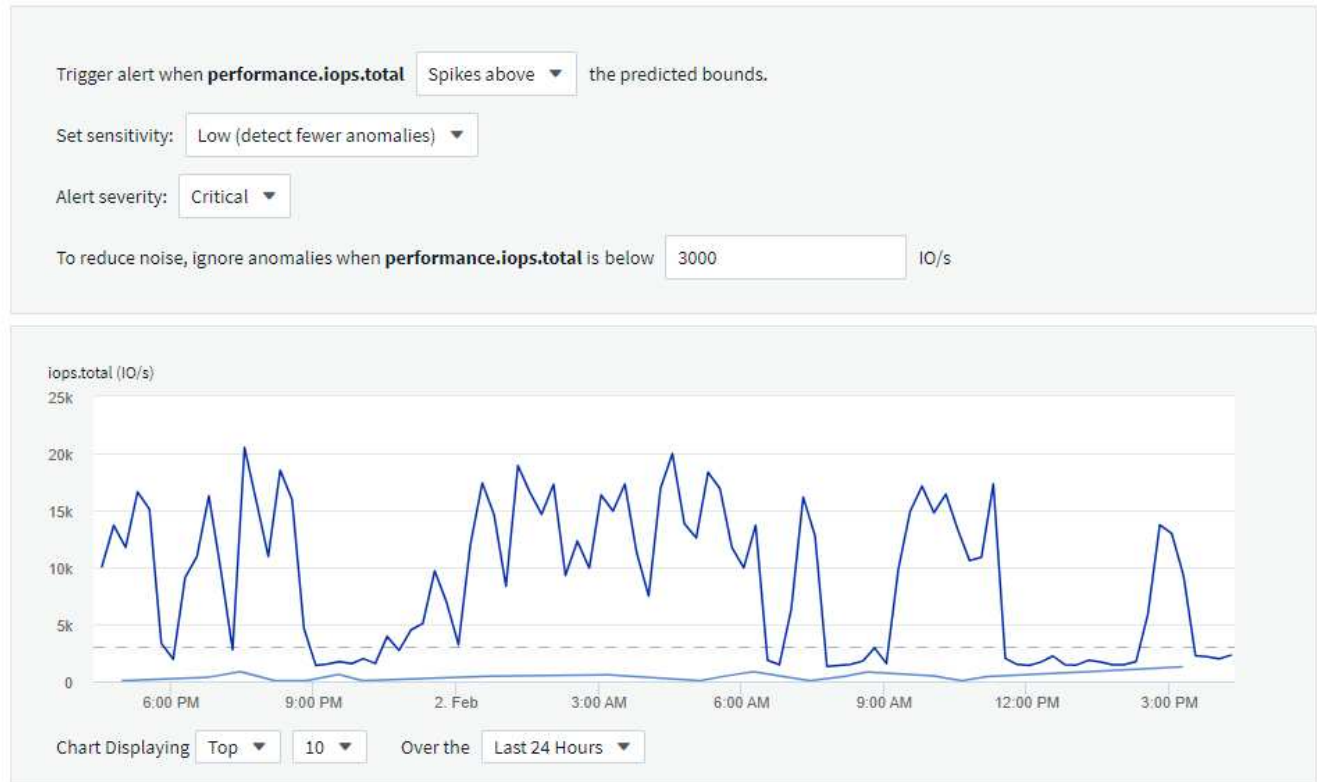
Use when you want to trigger alerts against performance spikes and drops

Choose the object and metric you want to monitor. You can set filters and grouping as with other types of monitors.

Next, set the conditions for the monitor.

- Trigger an alert when the selected metric either *Spikes above* the predicted bounds, *Drops below* those bounds, or both.
- Set sensitivity to *Medium*, *Low* (fewer anomalies are detected), or *High* (more anomalies are detected).
- Determine whether the alert level is *Critical* or *Warning*.
- Optionally, set a value below which anomalies are *ignored*. This can help reduce noise. This value is shown as a dashed line on the sample graph.

2 Define the monitor's conditions



Finally, you can configure a delivery method for the alerts (email, webhook, or both), give the monitor an optional description or corrective actions, and add the monitor to a custom group, if desired.

Save the monitor with a meaningful name, and you're done.


Upon creation, the monitor analyzes data from the previous week to establish an initial baseline. Anomaly detection becomes more accurate as time passes and more history occurs.



When a monitor is created, DII looks at any existing data for the week prior for significant data spikes or drops; these are considered anomalies. During the first week after monitor creation (the "learning" phase), there is a chance for increased "noise" in alerts. To mitigate this noise, only spikes or drops lasting longer than 30 minutes are considered anomalies and generate alerts. In the subsequent week as more data is analyzed the noise will typically reduce and a significant spike or drop lasting any period of time will be considered an anomaly..

Viewing the Anomalies

On an alert landing page, alerts triggered when anomalies are detected will show a highlighted band in the chart, from the time when the metric spiked outside the predicted bounds to when it moved back inside those bounds.



While viewing an anomaly chart on an alert landing page, you can choose the following options:

- Weekly Trend: compare values to the same time, same day on previous weeks, for up to 5 previous weeks.
- Full Anomaly Bounds: by default, the graph focuses on the metric value so you can better analyze the metric behavior. Select to show full anomaly bounds (maximum value, etc.)

You can also view objects that contributed to the anomaly by selecting those in the landing page's expert view. The chart will show the behavior of the selected objects.



System Monitors

Data Infrastructure Insights includes a number of system-defined monitors for both metrics and logs. The system monitors available are dependent on the data collectors present on your tenant. Because of that, the monitors available in Data Infrastructure Insights may change as data collectors are added or their configurations changed.



Many System Monitors are in *Paused* state by default. You can enable a system monitor by selecting the *Resume* option for the monitor. Ensure that *Advanced Counter Data Collection* and *Enable ONTAP EMS log collection* are enabled in the Data Collector. These options can be found in the ONTAP Data Collector under *Advanced Configuration*:

- ☒ Enable ONTAP EMS log collection
- ☒ Opt in for Advanced Counter Data Collection rollout.

toc:[]

Monitor Descriptions

System-defined monitors are comprised of pre-defined metrics and conditions, as well as default descriptions and corrective actions, which can not be modified. You *can* modify the notification recipient list for system-defined monitors. To view the metrics, conditions, description and corrective actions, or to modify the recipient list, open a system-defined monitor group and click the monitor name in the list.

System-defined monitor groups cannot be modified or removed.

The following system-defined monitors are available, in the noted groups.

- **ONTAP Infrastructure** includes monitors for infrastructure-related issues in ONTAP clusters.
- **ONTAP Workload Examples** includes monitors for workload-related issues.
- Monitors in both group default to *Paused* state.

Below are the system monitors currently included with Data Infrastructure Insights:

Metric Monitors

Monitor Name	Severity	Monitor Description	Corrective Action
Fiber Channel Port Utilization High	CRITICAL	Fiber Channel Protocol ports are used to receive and transfer the SAN traffic between the customer host system and the ONTAP LUNs. If the port utilization is high, then it will become a bottleneck and it will ultimately affect the performance of sensitive of Fiber Channel Protocol workloads....A warning alert indicates that planned action should be taken to balance network traffic....A critical alert indicates that service disruption is imminent and emergency measures should be taken to balance network traffic to ensure service continuity.	If critical threshold is breached, consider immediate actions to minimize service disruption: <ol style="list-style-type: none">1. Move workloads to another lower utilized FCP port.2. Limit the traffic of certain LUNs only to essential work, either via QoS policies in ONTAP or host-side configuration to lighten the utilization of the FCP ports.... If warning threshold is breached, plan to take the following actions: <ol style="list-style-type: none">1. Configure more FCP ports to handle the data traffic so that the port utilization gets distributed among more ports.2. Move workloads to another lower utilized FCP port.3. Limit the traffic of certain LUNs only to essential work, either via QoS policies in ONTAP or host-side configuration to lighten the utilization of the FCP ports.

Lun Latency High	CRITICAL	<p>LUNs are objects that serve the I/O traffic often driven by performance sensitive applications such as databases. High LUN latencies means that the applications themselves might suffer and be unable to accomplish their tasks....A warning alert indicates that planned action should be taken to move the LUN to appropriate Node or Aggregate....A critical alert indicates that service disruption is imminent and emergency measures should be taken to ensure service continuity. Following are expected latencies based on media type - SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds, and SATA HDD 17-20 milliseconds</p>	<p>If critical threshold is breached, consider following actions to minimize service disruption: If the LUN or its volume has a QoS policy associated with it, then evaluate its threshold limits and validate if they are causing the LUN workload to get throttled.... If warning threshold is breached, plan to take the following actions:</p> <ol style="list-style-type: none"> 1. If aggregate is also experiencing high utilization, move the LUN to another aggregate. 2. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node. 3. If the LUN or its volume has a QoS policy associated with it, evaluate its threshold limits and validate if they are causing the LUN workload to get throttled.
------------------	----------	--	---

Network Port Utilization High	CRITICAL	<p>Network ports are used to receive and transfer the NFS, CIFS, and iSCSI protocol traffic between the customer host systems and the ONTAP volumes. If the port utilization is high, then it becomes a bottleneck and it will ultimately affect the performance of NFS, CIFS and iSCSI workloads....A warning alert indicates that planned action should be taken to balance network traffic....A critical alert indicates that service disruption is imminent and emergency measures should be taken to balance network traffic to ensure service continuity.</p>	<p>If critical threshold is breached, consider following immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Limit the traffic of certain volumes only to essential work, either via QoS policies in ONTAP or host-side analysis to decrease the utilization of the network ports. 2. Configure one or more volumes to use another lower utilized network port.... <p>If warning threshold is breached, consider the following immediate actions:</p> <ol style="list-style-type: none"> 1. Configure more network ports to handle the data traffic so that the port utilization gets distributed among more ports. 2. Configure one or more volumes to use another lower utilized network port.
-------------------------------	----------	---	--

NVMe Namespace Latency High	CRITICAL	<p>NVMe Namespaces are objects that serve the I/O traffic that is driven by performance sensitive applications such as databases. High NVMe Namespaces latency means that the applications themselves may suffer and be unable to accomplish their tasks....A warning alert indicates that planned action should be taken to move the LUN to appropriate Node or Aggregate....A critical alert indicates that service disruption is imminent and emergency measures should be taken to ensure service continuity.</p>	<p>If critical threshold is breached, consider immediate actions to minimize service disruption: If the NVMe namespace or its volume has a QoS policy assigned to them, then evaluate its limit thresholds in case they are causing the NVMe namespace workload to get throttled.... If warning threshold is breached, consider to take the following actions:</p> <ol style="list-style-type: none"> 1. If aggregate is also experiencing high utilization, move the LUN to another aggregate. 2. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node. 3. If the NVMe namespace or its volume has a QoS policy assigned to them, evaluate its limit thresholds in case they are causing the NVMe namespace workload to get throttled.
-----------------------------	----------	---	---

QTree Capacity Full	CRITICAL	<p>A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a default space quota or a quota defined by a quota policy to limit amount of data stored in the tree within the volume capacity....A warning alert indicates that planned action should be taken to increase the space....A critical alert indicates that service disruption is imminent and emergency measures should be taken to free up space to ensure service continuity.</p>	<p>If critical threshold is breached, consider immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Increase the space of the qtree in order to accommodate the growth. 2. Delete unwanted data to free up space.... <p>If warning threshold is breached, plan to take the following immediate actions:</p> <ol style="list-style-type: none"> 1. Increase the space of the qtree in order to accommodate the growth. 2. Delete unwanted data to free up space.
QTree Capacity Hard Limit	CRITICAL	<p>A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a space quota measured in KBytes that is used to store data in order to control the growth of user data in volume and not exceed its total capacity....A qtree maintains a soft storage capacity quota that provides alert to the user proactively before reaching the total capacity quota limit in the qtree and being unable to store data anymore. Monitoring the amount of data stored within a qtree ensures that the user receives uninterrupted data service.</p>	<p>If critical threshold is breached, consider following immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Increase the tree space quota in order to accommodate the growth 2. Instruct the user to delete unwanted data in the tree to free up space

QTree Capacity Soft Limit	WARNING	<p>A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a space quota measured in KBytes that it can use to store data in order to control the growth of user data in volume and not exceed its total capacity....A qtree maintains a soft storage capacity quota that provides alert to the user proactively before reaching the total capacity quota limit in the qtree and being unable to store data anymore. Monitoring the amount of data stored within a qtree ensures that the user receives uninterrupted data service.</p>	<p>If warning threshold is breached, consider the following immediate actions:</p> <ol style="list-style-type: none"> 1. Increase the tree space quota to accommodate the growth. 2. Instruct the user to delete unwanted data in the tree to free up space.
QTree Files Hard Limit	CRITICAL	<p>A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a quota of the number of files that it can contain to maintain a manageable file system size within the volume....A qtree maintains a hard file number quota beyond which new files in the tree are denied. Monitoring the number of files within a qtree ensures that the user receives uninterrupted data service.</p>	<p>If critical threshold is breached, consider immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Increase the file count quota for the qtree. 2. Delete unwanted files from the qtree file system.

QTree Files Soft Limit	WARNING	<p>A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a quota of the number of files that it can contain in order to maintain a manageable file system size within the volume....A qtree maintains a soft file number quota to provide alert to the user proactively before reaching the limit of files in the qtree and being unable to store any additional files. Monitoring the number of files within a qtree ensures that the user receives uninterrupted data service.</p>	<p>If warning threshold is breached, plan to take the following immediate actions:</p> <ol style="list-style-type: none"> 1. Increase the file count quota for the qtree. 2. Delete unwanted files from the qtree file system.
Snapshot Reserve Space Full	CRITICAL	<p>Storage capacity of a volume is necessary to store application and customer data. A portion of that space, called snapshot reserved space, is used to store snapshots which allow data to be protected locally. The more new and updated data stored in the ONTAP volume the more snapshot capacity is used and less snapshot storage capacity is available for future new or updated data. If the snapshot data capacity within a volume reaches the total snapshot reserve space, it might lead to the customer being unable to store new snapshot data and reduction in the level of protection for the data in the volume. Monitoring the volume used snapshot capacity ensures data services continuity.</p>	<p>If critical threshold is breached, consider immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Configure snapshots to use data space in the volume when the snapshot reserve is full. 2. Delete some older unwanted snapshots to free up space.... <p>If warning threshold is breached, plan to take the following immediate actions:</p> <ol style="list-style-type: none"> 1. Increase the snapshot reserve space within the volume to accommodate the growth. 2. Configure snapshots to use data space in the volume when the snapshot reserve is full.

Storage Capacity Limit	CRITICAL	When a storage pool (aggregate) is filling up, I/O operations slow down and finally stop resulting in storage outage incident. A warning alert indicates that planned action should be taken soon to restore minimum free space. A critical alert indicates that service disruption is imminent and emergency measures should be taken to free up space to ensure service continuity.	<p>If critical threshold is breached, immediately consider the following actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Delete Snapshots on non-critical volumes. 2. Delete Volumes or LUNs that are non-essential workloads and that may be restored from off storage copies.....If warning threshold is breached, plan the following immediate actions: 1. Move one or more volumes to a different storage location. 2. Add more storage capacity. 3. Change storage efficiency settings or tier inactive data to cloud storage.
Storage Performance Limit	CRITICAL	When a storage system reaches its performance limit, operations slow down, latency goes up and workloads and applications may start failing. ONTAP evaluates the storage pool utilization for workloads and estimates what percent of performance has been consumed....A warning alert indicates that planned action should be taken to reduce storage pool load to ensure that there will be enough storage pool performance left to service workload peaks....A critical alert indicates that a performance brownout is imminent and emergency measures should be taken to reduce storage pool load to ensure service continuity.	<p>If critical threshold is breached, consider following immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Suspend scheduled tasks such as Snapshots or SnapMirror replication. 2. Idle non-essential workloads.... <p>If warning threshold is breached, take the following actions immediately:</p> <ol style="list-style-type: none"> 1. Move one or more workloads to a different storage location. 2. Add more storage nodes (AFF) or disk shelves(FAS) and redistribute workloads 3. Change workload characteristics(block size, application caching).

User Quota Capacity Hard Limit	CRITICAL	<p>ONTAP recognizes the users of Unix or Windows systems who have the rights to access volumes, files or directories within a volume. As a result, ONTAP allows the customers to configure storage capacity for their users or groups of users of their Linux or Windows systems. The user or group policy quota limits the amount of space the user can utilize for their own data....A hard limit of this quota allows notification of the user when the amount of capacity used within the volume is right before reaching the total capacity quota. Monitoring the amount of data stored within a user or group quota ensures that the user receives uninterrupted data service.</p>	<p>If critical threshold is breached, consider following immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Increase the space of the user or group quota in order to accommodate the growth. 2. Instruct the user or group to delete unwanted data to free up space.
--------------------------------	----------	--	--

User Quota Capacity Soft Limit	WARNING	<p>ONTAP recognizes the users of Unix or Windows systems that have the rights to access volumes, files or directories within a volume. As a result, ONTAP allows the customers to configure storage capacity for their users or groups of users of their Linux or Windows systems. The user or group policy quota limits the amount of space the user can utilize for their own data....A soft limit of this quota allows proactive notification to the user when the amount of capacity used within the volume is reaching the total capacity quota. Monitoring the amount of data stored within a user or group quota ensures that the user receives uninterrupted data service.</p>	<p>If warning threshold is breached, plan to take the following immediate actions:</p> <ol style="list-style-type: none"> 1. Increase the space of the user or group quota in order to accommodate the growth. 2. Delete unwanted data to free up space.
--------------------------------	---------	--	--

Volume Capacity Full	CRITICAL	<p>Storage capacity of a volume is necessary to store application and customer data. The more data stored in the ONTAP volume the less storage availability for future data. If the data storage capacity within a volume reaches the total storage capacity may lead to the customer being unable to store data due to lack of storage capacity. Monitoring the volume used storage capacity ensures data services continuity.</p>	<p>If critical threshold is breached, consider following immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Increase the space of the volume to accommodate the growth. 2. Delete unwanted data to free up space. 3. If snapshot copies occupy more space than the snapshot reserve, delete old Snapshots or enable Volume Snapshot Autodelete....If warning threshold is breached, plan to take the following immediate actions: <ol style="list-style-type: none"> 1. Increase the space of the volume in order to accommodate the growth 2. If snapshot copies occupy more space than the snapshot reserve, delete old Snapshots or enabling Volume Snapshot Autodelete.....
----------------------	----------	---	---

Volume Inodes Limit	CRITICAL	<p>Volumes that store files use index nodes (inode) to store file metadata. When a volume exhausts its inode allocation, no more files can be added to it....A warning alert indicates that planned action should be taken to increase the number of available inodes....A critical alert indicates that file limit exhaustion is imminent and emergency measures should be taken to free up inodes to ensure service continuity.</p>	<p>If critical threshold is breached, consider following immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Increase the inodes value for the volume. If the inodes value is already at the max value, then split the volume into two or more volumes because the file system has grown beyond the maximum size. 2. Use FlexGroup as it helps to accommodate large file systems.... <p>If warning threshold is breached, plan to take the following immediate actions:</p> <ol style="list-style-type: none"> 1. Increase the inodes value for the volume. If the inodes value is already at the max, then split the volume into two or more volumes because the file system has grown beyond the maximum size. 2. Use FlexGroup as it helps to accommodate large file systems
---------------------	----------	---	---

Volume Latency High	CRITICAL	<p>Volumes are objects that serve the I/O traffic often driven by performance sensitive applications including devOps applications, home directories, and databases. High volume latencies means that the applications themselves may suffer and be unable to accomplish their tasks. Monitoring volume latencies is critical to maintain application consistent performance. The following are expected latencies based on media type - SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.</p>	<p>If critical threshold is breached, consider following immediate actions to minimize service disruption: If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled....</p> <p>If warning threshold is breached, consider the following immediate actions:</p> <ol style="list-style-type: none"> 1. If aggregate is also experiencing high utilization, move the volume to another aggregate. 2. If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled. 3. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node.
Monitor Name	Severity	Monitor Description	Corrective Action

Node High Latency	WARNING / CRITICAL	<p>Node latency has reached the levels where it might affect the performance of the applications on the node. Lower node latency ensures consistent performance of the applications. The expected latencies based on media type are: SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.</p>	<p>If critical threshold is breached, then immediate actions should be taken to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Suspend scheduled tasks, Snapshots or SnapMirror replication 2. Lower the demand of lower priority workloads via QoS limits 3. Inactivate non-essential workloads <p>Consider immediate actions when warning threshold is breached:</p> <ol style="list-style-type: none"> 1. Move one or more workloads to a different storage location 2. Lower the demand of lower priority workloads via QoS limits 3. Add more storage nodes (AFF) or disk shelves (FAS) and redistribute workloads 4. Change workload characteristics (block size, application caching etc)
-------------------	--------------------	---	---

Node Performance Limit	WARNING / CRITICAL	<p>Node performance utilization has reached the levels where it might affect the performance of the IOs and the applications supported by the node. Low node performance utilization ensures consistent performance of the applications.</p>	<p>Immediate actions should be taken to minimize service disruption if critical threshold is breached:</p> <ol style="list-style-type: none"> 1. Suspend scheduled tasks, Snapshots or SnapMirror replication 2. Lower the demand of lower priority workloads via QoS limits 3. Inactivate non-essential workloads <p>Consider the following actions if warning threshold is breached:</p> <ol style="list-style-type: none"> 1. Move one or more workloads to a different storage location 2. Lower the demand of lower priority workloads via QoS limits 3. Add more storage nodes (AFF) or disk shelves (FAS) and redistribute workloads 4. Change workload characteristics (block size, application caching etc)
------------------------	--------------------	--	---

Storage VM High Latency	WARNING / CRITICAL	Storage VM (SVM) latency has reached the levels where it might affect the performance of the applications on the storage VM. Lower storage VM latency ensures consistent performance of the applications. The expected latencies based on media type are: SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.	<p>If critical threshold is breached, then immediately evaluate the threshold limits for volumes of the storage VM with a QoS policy assigned, to verify whether they are causing the volume workloads to get throttled</p> <p>Consider following immediate actions when warning threshold is breached:</p> <ol style="list-style-type: none"> 1. If aggregate is also experiencing high utilization, move some volumes of the storage VM to another aggregate. 2. For volumes of the storage VM with a QoS policy assigned, evaluate the threshold limits if they are causing the volume workloads to get throttled 3. If the node is experiencing high utilization, move some volumes of the storage VM to another node or reduce the total workload of the node
User Quota Files Hard Limit	CRITICAL	The number of files created within the volume has reached the critical limit and additional files cannot be created. Monitoring the number of files stored ensures that the user receives uninterrupted data service.	<p>Immediate actions are required to minimize service disruption if critical threshold is breached....Consider taking following actions:</p> <ol style="list-style-type: none"> 1. Increase the file count quota for the specific user 2. Delete unwanted files to reduce the pressure on the files quota for the specific user

User Quota Files Soft Limit	WARNING	The number of files created within the volume has reached the threshold limit of the quota and is near to the critical limit. You cannot create additional files if quota reaches the critical limit. Monitoring the number of files stored by a user ensures that the user receives uninterrupted data service.	Consider immediate actions if warning threshold is breached: <ol style="list-style-type: none"> 1. Increase the file count quota for the specific user quota 2. Delete unwanted files to reduce the pressure on the files quota for the specific user
Volume Cache Miss Ratio	WARNING / CRITICAL	Volume Cache Miss Ratio is the percentage of read requests from the client applications that are returned from the disk instead of being returned from the cache. This means that the volume has reached the set threshold.	<p>If critical threshold is breached, then immediate actions should be taken to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Move some workloads off of the node of the volume to reduce the IO load 2. If not already on the node of the volume, increase the WAFL cache by purchasing and adding a Flash Cache 3. Lower the demand of lower priority workloads on the same node via QoS limits <p>Consider immediate actions when warning threshold is breached:</p> <ol style="list-style-type: none"> 1. Move some workloads off of the node of the volume to reduce the IO load 2. If not already on the node of the volume, increase the WAFL cache by purchasing and adding a Flash Cache 3. Lower the demand of lower priority workloads on the same node via QoS limits 4. Change workload characteristics (block size, application caching etc)

Volume Qtree Quota Overcommit	WARNING / CRITICAL	Volume Qtree Quota Overcommit specifies the percentage at which a volume is considered to be overcommitted by the qtree quotas. The set threshold for the qtree quota is reached for the volume. Monitoring the volume qtree quota overcommit ensures that the user receives uninterrupted data service.	<p>If critical threshold is breached, then immediate actions should be taken to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Increase the space of the volume 2. Delete unwanted data <p>When warning threshold is breached, then consider increasing the space of the volume.</p>
-------------------------------	--------------------	--	--

[Back to Top](#)

Log Monitors

Monitor Name	Severity	Description	Corrective Action
AWS Credentials Not Initialized	INFO	This event occurs when a module attempts to access Amazon Web Services (AWS) Identity and Access Management (IAM) role-based credentials from the cloud credentials thread before they are initialized.	Wait for the cloud credentials thread, as well as the system, to complete initialization.

Cloud Tier Unreachable	CRITICAL	A storage node cannot connect to Cloud Tier object store API. Some data will be inaccessible.	<p>If you use on-premises products, perform the following corrective actions: ...Verify that your intercluster LIF is online and functional by using the "network interface show" command....Check the network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF....Ensure the following:...The configuration of your object store has not changed....The login and connectivity information is still valid....Contact NetApp technical support if the issue persists.</p> <p>If you use Cloud Volumes ONTAP, perform the following corrective actions: ...Ensure that the configuration of your object store has not changed.... Ensure that the login and connectivity information is still valid....Contact NetApp technical support if the issue persists.</p>
Disk Out of Service	INFO	This event occurs when a disk is removed from service because it has been marked failed, is being sanitized, or has entered the Maintenance Center.	None.

FlexGroup Constituent Full	CRITICAL	A constituent within a FlexGroup volume is full, which might cause a potential disruption of service. You can still create or expand files on the FlexGroup volume. However, none of the files that are stored on the constituent can be modified. As a result, you might see random out-of-space errors when you try to perform write operations on the FlexGroup volume.	It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X" command....Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent.
Flexgroup Constituent Nearly Full	WARNING	A constituent within a FlexGroup volume is nearly out of space, which might cause a potential disruption of service. Files can be created and expanded. However, if the constituent runs out of space, you might not be able to append to or modify the files on the constituent.	It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X" command....Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent.
FlexGroup Constituent Nearly Out of Inodes	WARNING	A constituent within a FlexGroup volume is almost out of inodes, which might cause a potential disruption of service. The constituent receives lesser create requests than average. This might impact the overall performance of the FlexGroup volume, because the requests are routed to constituents with more inodes.	It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X" command....Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent.

FlexGroup Constituent Out of Inodes	CRITICAL	A constituent of a FlexGroup volume has run out of inodes, which might cause a potential disruption of service. You cannot create new files on this constituent. This might lead to an overall imbalanced distribution of content across the FlexGroup volume.	It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X" command....Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent.
LUN Offline	INFO	This event occurs when a LUN is brought offline manually.	Bring the LUN back online.
Main Unit Fan Failed	WARNING	One or more main unit fans have failed. The system remains operational....However, if the condition persists for too long, the overtemperature might trigger an automatic shutdown.	Reseat the failed fans. If the error persists, replace them.
Main Unit Fan in Warning State	INFO	This event occurs when one or more main unit fans are in a warning state.	Replace the indicated fans to avoid overheating.
NVRAM Battery Low	WARNING	The NVRAM battery capacity is critically low. There might be a potential data loss if the battery runs out of power....Your system generates and transmits an AutoSupport or "call home" message to NetApp technical support and the configured destinations if it is configured to do so. The successful delivery of an AutoSupport message significantly improves problem determination and resolution.	Perform the following corrective actions:...View the battery's current status, capacity, and charging state by using the "system node environment sensors show" command....If the battery was replaced recently or the system was non-operational for an extended period of time, monitor the battery to verify that it is charging properly....Contact NetApp technical support if the battery runtime continues to decrease below critical levels, and the storage system shuts down automatically.

Service Processor Not Configured	WARNING	This event occurs on a weekly basis, to remind you to configure the Service Processor (SP). The SP is a physical device that is incorporated into your system to provide remote access and remote management capabilities. You should configure the SP to use its full functionality.	Perform the following corrective actions:...Configure the SP by using the "system service-processor network modify" command.... Optionally, obtain the MAC address of the SP by using the "system service-processor network show" command.... Verify the SP network configuration by using the "system service-processor network show" command.... Verify that the SP can send an AutoSupport email by using the "system service-processor autosupport invoke" command. NOTE: AutoSupport email hosts and recipients should be configured in ONTAP before you issue this command.
Service Processor Offline	CRITICAL	ONTAP is no longer receiving heartbeats from the Service Processor (SP), even though all the SP recovery actions have been taken. ONTAP cannot monitor the health of the hardware without the SP.... The system will shut down to prevent hardware damage and data loss. Set up a panic alert to be notified immediately if the SP goes offline.	Power-cycle the system by performing the following actions:... Pull the controller out from the chassis.... Push the controller back in.... Turn the controller back on.... If the problem persists, replace the controller module.

Shelf Fans Failed	CRITICAL	The indicated cooling fan or fan module of the shelf has failed. The disks in the shelf might not receive enough cooling airflow, which might result in disk failure.	Perform the following corrective actions:...Verify that the fan module is fully seated and secured. NOTE: The fan is integrated into the power supply module in some disk shelves....If the issue persists, replace the fan module....If the issue still persists, contact NetApp technical support for assistance.
System Cannot Operate Due to Main Unit Fan Failure	CRITICAL	One or more main unit fans have failed, disrupting system operation. This might lead to a potential data loss.	Replace the failed fans.
Unassigned Disks	INFO	System has unassigned disks - capacity is being wasted and your system may have some misconfiguration or partial configuration change applied.	Perform the following corrective actions:...Determine which disks are unassigned by using the "disk show -n" command....Assign the disks to a system by using the "disk assign" command.
Antivirus Server Busy	WARNING	The antivirus server is too busy to accept any new scan requests.	If this message occurs frequently, ensure that there are enough antivirus servers to handle the virus scan load generated by the SVM.
AWS Credentials for IAM Role Expired	CRITICAL	Cloud Volume ONTAP has become inaccessible. The Identity and Access Management (IAM) role-based credentials have expired. The credentials are acquired from the Amazon Web Services (AWS) metadata server using the IAM role, and are used to sign API requests to Amazon Simple Storage Service (Amazon S3).	Perform the following:...Log in to the AWS EC2 Management Console....Navigate to the Instances page....Find the instance for the Cloud Volumes ONTAP deployment and check its health....Verify that the AWS IAM role associated with the instance is valid and has been granted proper privileges to the instance.

AWS Credentials for IAM Role Not Found	CRITICAL	The cloud credentials thread cannot acquire the Amazon Web Services (AWS) Identity and Access Management (IAM) role-based credentials from the AWS metadata server. The credentials are used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible....	Perform the following:...Log in to the AWS EC2 Management Console....Navigate to the Instances page....Find the instance for the Cloud Volumes ONTAP deployment and check its health....Verify that the AWS IAM role associated with the instance is valid and has been granted proper privileges to the instance.
AWS Credentials for IAM Role Not Valid	CRITICAL	The Identity and Access Management (IAM) role-based credentials are not valid. The credentials are acquired from the Amazon Web Services (AWS) metadata server using the IAM role, and are used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible.	Perform the following:...Log in to the AWS EC2 Management Console....Navigate to the Instances page....Find the instance for the Cloud Volumes ONTAP deployment and check its health....Verify that the AWS IAM role associated with the instance is valid and has been granted proper privileges to the instance.
AWS IAM Role Not Found	CRITICAL	The Identity and Access Management (IAM) roles thread cannot find an Amazon Web Services (AWS) IAM role on the AWS metadata server. The IAM role is required to acquire role-based credentials used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible....	Perform the following:...Log in to the AWS EC2 Management Console....Navigate to the Instances page....Find the instance for the Cloud Volumes ONTAP deployment and check its health....Verify that the AWS IAM role associated with the instance is valid.

AWS IAM Role Not Valid	CRITICAL	The Amazon Web Services (AWS) Identity and Access Management (IAM) role on the AWS metadata server is not valid. The Cloud Volume ONTAP has become inaccessible....	Perform the following:...Log in to the AWS EC2 Management Console....Navigate to the Instances page....Find the instance for the Cloud Volumes ONTAP deployment and check its health....Verify that the AWS IAM role associated with the instance is valid and has been granted proper privileges to the instance.
AWS Metadata Server Connection Fail	CRITICAL	The Identity and Access Management (IAM) roles thread cannot establish a communication link with the Amazon Web Services (AWS) metadata server. Communication should be established to acquire the necessary AWS IAM role-based credentials used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible....	Perform the following:...Log in to the AWS EC2 Management Console....Navigate to the Instances page....Find the instance for the Cloud Volumes ONTAP deployment and check its health....
FabricPool Space Usage Limit Nearly Reached	WARNING	The total cluster-wide FabricPool space usage of object stores from capacity-licensed providers has nearly reached the licensed limit.	Perform the following corrective actions:...Check the percentage of the licensed capacity used by each FabricPool storage tier by using the "storage aggregate object-store show-space" command....Delete Snapshot copies from volumes with the tiering policy "snapshot" or "backup" by using the "volume snapshot delete" command to clear up space....Install a new license on the cluster to increase the licensed capacity.

FabricPool Space Usage Limit Reached	CRITICAL	The total cluster-wide FabricPool space usage of object stores from capacity-licensed providers has reached the license limit.	Perform the following corrective actions:...Check the percentage of the licensed capacity used by each FabricPool storage tier by using the "storage aggregate object-store show-space" command....Delete Snapshot copies from volumes with the tiering policy "snapshot" or "backup" by using the "volume snapshot delete" command to clear up space....Install a new license on the cluster to increase the licensed capacity.
Giveback of Aggregate Failed	CRITICAL	This event occurs during the migration of an aggregate as part of a storage failover (SFO) giveback, when the destination node cannot reach the object stores.	Perform the following corrective actions:...Verify that your intercluster LIF is online and functional by using the "network interface show" command....Check network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF. ...Verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the "aggregate object-store config show" command....Alternatively, you can override the error by specifying false for the "require-partner-waiting" parameter of the giveback command....Contact NetApp technical support for more information or assistance.

HA Interconnect Down	WARNING	The high-availability (HA) interconnect is down. Risk of service outage when failover is not available.	<p>Corrective actions depend on the number and type of HA interconnect links supported by the platform, as well as the reason why the interconnect is down.</p> <p>...If the links are down:...Verify that both controllers in the HA pair are operational....For externally connected links, make sure that the interconnect cables are connected properly and that the small form-factor pluggables (SFPs), if applicable, are seated properly on both controllers....For internally connected links, disable and re-enable the links, one after the other, by using the "ic link off" and "ic link on" commands.</p> <p>...If links are disabled, enable the links by using the "ic link on" command.</p> <p>...If a peer is not connected, disable and re-enable the links, one after the other, by using the "ic link off" and "ic link on" commands....Contact NetApp technical support if the issue persists.</p>
----------------------	---------	---	--

Max Sessions Per User Exceeded	WARNING	<p>You have exceeded the maximum number of sessions allowed per user over a TCP connection. Any request to establish a session will be denied until some sessions are released. ...</p>	<p>Perform the following corrective actions:</p> <p>...Inspect all the applications that run on the client, and terminate any that are not operating properly...Reboot the client...Check if the issue is caused by a new or existing application:...If the application is new, set a higher threshold for the client by using the "cifs option modify -max-opens -same-file-per-tree" command.</p> <p>In some cases, clients operate as expected, but require a higher threshold. You should have advanced privilege to set a higher threshold for the client. ...If the issue is caused by an existing application, there might be an issue with the client. Contact NetApp technical support for more information or assistance.</p>
--------------------------------	---------	---	---

Max Times Open Per File Exceeded	WARNING	<p>You have exceeded the maximum number of times that you can open the file over a TCP connection. Any request to open this file will be denied until you close some open instances of the file. This typically indicates abnormal application behavior....</p>	<p>Perform the following corrective actions:...Inspect the applications that run on the client using this TCP connection.</p> <p>The client might be operating incorrectly because of the application running on it....Reboot the client....Check if the issue is caused by a new or existing application:...If the application is new, set a higher threshold for the client by using the "cifs option modify -max-opens -same-file-per-tree" command.</p> <p>In some cases, clients operate as expected, but require a higher threshold. You should have advanced privilege to set a higher threshold for the client. ...If the issue is caused by an existing application, there might be an issue with the client. Contact NetApp technical support for more information or assistance.</p>
----------------------------------	---------	---	---

NetBIOS Name Conflict	CRITICAL	<p>The NetBIOS Name Service has received a negative response to a name registration request, from a remote machine. This is typically caused by a conflict in the NetBIOS name or an alias. As a result, clients might not be able to access data or connect to the right data-serving node in the cluster.</p>	<p>Perform any one of the following corrective actions:...</p> <p>If there is a conflict in the NetBIOS name or an alias, perform one of the following:...</p> <p>Delete the duplicate NetBIOS alias by using the "vserver cifs delete -aliases alias -vserver vserver" command....</p> <p>Rename a NetBIOS alias by deleting the duplicate name and adding an alias with a new name by using the "vserver cifs create -aliases alias -vserver vserver" command. ...</p> <p>If there are no aliases configured and there is a conflict in the NetBIOS name, then rename the CIFS server by using the "vserver cifs delete -vserver vserver" and "vserver cifs create -cifs -server netbiosname" commands.</p> <p>NOTE: Deleting a CIFS server can make data inaccessible. ...</p> <p>Remove NetBIOS name or rename the NetBIOS on the remote machine.</p>
NFSv4 Store Pool Exhausted	CRITICAL	A NFSv4 store pool has been exhausted.	<p>If the NFS server is unresponsive for more than 10 minutes after this event, contact NetApp technical support.</p>
No Registered Scan Engine	CRITICAL	<p>The antivirus connector notified ONTAP that it does not have a registered scan engine. This might cause data unavailability if the "scan-mandatory" option is enabled.</p>	<p>Perform the following corrective actions:...</p> <p>Ensure that the scan engine software installed on the antivirus server is compatible with ONTAP....</p> <p>Ensure that scan engine software is running and configured to connect to the antivirus connector over local loopback.</p>

No Vscan Connection	CRITICAL	ONTAP has no Vscan connection to service virus scan requests. This might cause data unavailability if the "scan-mandatory" option is enabled.	Ensure that the scanner pool is properly configured and the antivirus servers are active and connected to ONTAP.
Node Root Volume Space Low	CRITICAL	The system has detected that the root volume is dangerously low on space. The node is not fully operational. Data LIFs might have failed over within the cluster, because of which NFS and CIFS access is limited on the node. Administrative capability is limited to local recovery procedures for the node to clear up space on the root volume.	Perform the following corrective actions:...Clear up space on the root volume by deleting old Snapshot copies, deleting files you no longer need from the /mroot directory, or expanding the root volume capacity....Reboot the controller....Contact NetApp technical support for more information or assistance.
Nonexistent Admin Share	CRITICAL	Vscan issue: a client has attempted to connect to a nonexistent ONTAP_ADMIN\$ share.	Ensure that Vscan is enabled for the mentioned SVM ID. Enabling Vscan on a SVM causes the ONTAP_ADMIN\$ share to be created for the SVM automatically.
NVMe Namespace Out of Space	CRITICAL	An NVMe namespace has been brought offline because of a write failure caused by lack of space.	Add space to the volume, and then bring the NVMe namespace online by using the "vserver nvme namespace modify" command.
NVMe-oF Grace Period Active	WARNING	This event occurs on a daily basis when the NVMe over Fabrics (NVMe-oF) protocol is in use and the grace period of the license is active. The NVMe-oF functionality requires a license after the license grace period expires. NVMe-oF functionality is disabled when the license grace period is over.	Contact your sales representative to obtain an NVMe-oF license, and add it to the cluster, or remove all instances of NVMe-oF configuration from the cluster.

NVMe-oF Grace Period Expired	WARNING	The NVMe over Fabrics (NVMe-oF) license grace period is over and the NVMe-oF functionality is disabled.	Contact your sales representative to obtain an NVMe-oF license, and add it to the cluster.
NVMe-oF Grace Period Start	WARNING	The NVMe over Fabrics (NVMe-oF) configuration was detected during the upgrade to ONTAP 9.5 software. NVMe-oF functionality requires a license after the license grace period expires.	Contact your sales representative to obtain an NVMe-oF license, and add it to the cluster.
Object Store Host Unresolvable	CRITICAL	The object store server host name cannot be resolved to an IP address. The object store client cannot communicate with the object-store server without resolving to an IP address. As a result, data might be inaccessible.	Check the DNS configuration to verify that the host name is configured correctly with an IP address.
Object Store Intercluster LIF Down	CRITICAL	The object-store client cannot find an operational LIF to communicate with the object store server. The node will not allow object store client traffic until the intercluster LIF is operational. As a result, data might be inaccessible.	Perform the following corrective actions:...Check the intercluster LIF status by using the "network interface show -role intercluster" command....Verify that the intercluster LIF is configured correctly and operational....If an intercluster LIF is not configured, add it by using the "network interface create -role intercluster" command.
Object Store Signature Mismatch	CRITICAL	The request signature sent to the object store server does not match the signature calculated by the client. As a result, data might be inaccessible.	Verify that the secret access key is configured correctly. If it is configured correctly, contact NetApp technical support for assistance.

READDIR Timeout	CRITICAL	<p>A READDIR file operation has exceeded the timeout that it is allowed to run in WAFL. This can be because of very large or sparse directories. Corrective action is recommended.</p>	<p>Perform the following corrective actions:...</p> <p>Find information specific to recent directories that have had READDIR file operations expire by using the following 'diag' privilege nodeshell CLI command:</p> <pre>wafl readdir notice show....</pre> <p>Check if directories are indicated as sparse or not:...</p> <p>If a directory is indicated as sparse, it is recommended that you copy the contents of the directory to a new directory to remove the sparseness of the directory file. ...</p> <p>If a directory is not indicated as sparse and the directory is large, it is recommended that you reduce the size of the directory file by reducing the number of file entries in the directory.</p>
-----------------	----------	--	---

Relocation of Aggregate Failed	CRITICAL	This event occurs during the relocation of an aggregate, when the destination node cannot reach the object stores.	Perform the following corrective actions:...Verify that your intercluster LIF is online and functional by using the "network interface show" command....Check network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF. ...Verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the "aggregate object-store config show" command....Alternatively, you can override the error by using the "override-destination-checks" parameter of the relocation command....Contact NetApp technical support for more information or assistance.
Shadow Copy Failed	CRITICAL	A Volume Shadow Copy Service (VSS), a Microsoft Server backup and restore service operation, has failed.	Check the following using the information provided in the event message:...Is shadow copy configuration enabled?...Are the appropriate licenses installed? ...On which shares is the shadow copy operation performed?...Is the share name correct?...Does the share path exist?...What are the states of the shadow copy set and its shadow copies?

Storage Switch Power Supplies Failed	WARNING	There is a missing power supply in the cluster switch. Redundancy is reduced, risk of outage with any further power failures.	Perform the following corrective actions:...Ensure that the power supply mains, which supplies power to the cluster switch, is turned on....Ensure that the power cord is connected to the power supply....Contact NetApp technical support if the issue persists.
Too Many CIFS Authentication	WARNING	Many authentication negotiations have occurred simultaneously. There are 256 incomplete new session requests from this client.	Investigate why the client has created 256 or more new connection requests. You might have to contact the vendor of the client or of the application to determine why the error occurred.
Unauthorized User Access to Admin Share	WARNING	A client has attempted to connect to the privileged ONTAP_ADMIN\$ share even though their logged-in user is not an allowed user.	Perform the following corrective actions:...Ensure that the mentioned username and IP address is configured in one of the active Vscan scanner pools....Check the scanner pool configuration that is currently active by using the "vserver vscan scanner pool show-active" command.
Virus Detected	WARNING	A Vscan server has reported an error to the storage system. This typically indicates that a virus has been found. However, other errors on the Vscan server can cause this event....Client access to the file is denied. The Vscan server might, depending on its settings and configuration, clean the file, quarantine it, or delete it.	Check the log of the Vscan server reported in the "syslog" event to see if it was able to successfully clean, quarantine, or delete the infected file. If it was not able to do so, a system administrator might have to manually delete the file.
Volume Offline	INFO	This message indicates that a volume is made offline.	Bring the volume back online.

Volume Restricted	INFO	This event indicates that a flexible volume is made restricted.	Bring the volume back online.
Storage VM Stop Succeeded	INFO	This message occurs when a 'vserver stop' operation succeeds.	Use 'vserver start' command to start the data access on a storage VM.
Node Panic	WARNING	This event is issued when a panic occurs	Contact NetApp customer support.

[Back to Top](#)

Anti-Ransomware Log Monitors

Monitor Name	Severity	Description	Corrective Action
Storage VM Anti-ransomware Monitoring Disabled	WARNING	The anti-ransomware monitoring for the storage VM is disabled. Enable anti-ransomware to protect the storage VM.	None
Storage VM Anti-ransomware Monitoring Enabled (Learning Mode)	INFO	The anti-ransomware monitoring for the storage VM is enabled in learning mode.	None
Volume Anti-ransomware Monitoring Enabled	INFO	The anti-ransomware monitoring for the volume is enabled.	None
Volume Anti-ransomware Monitoring Disabled	WARNING	The anti-ransomware monitoring for the volume is disabled. Enable anti-ransomware to protect the volume.	None
Volume Anti-ransomware Monitoring Enabled (Learning Mode)	INFO	The anti-ransomware monitoring for the volume is enabled in learning mode.	None
Volume Anti-ransomware Monitoring Paused (Learning Mode)	WARNING	The anti-ransomware monitoring for the volume is paused in learning mode.	None
Volume Anti-ransomware Monitoring Paused	WARNING	The anti-ransomware monitoring for the volume is paused.	None
Volume Anti-ransomware Monitoring Disabling	WARNING	The anti-ransomware monitoring for the volume is disabling.	None

Ransomware Activity Detected	CRITICAL	<p>To protect the data from the detected ransomware, a Snapshot copy has been taken that can be used to restore original data.</p> <p>Your system generates and transmits an AutoSupport or "call home" message to NetApp technical support and any configured destinations. AutoSupport message improves problem determination and resolution.</p>	Refer to the "FINAL-DOCUMENT-NAME" to take remedial measures for ransomware activity.
------------------------------	----------	---	---

[Back to Top](#)

FSx for NetApp ONTAP Monitors

Monitor Name	Thresholds	Monitor Description	Corrective Action
FSx Volume Capacity is Full	Warning @ > 85 %...Critical @ > 95 %	<p>Storage capacity of a volume is necessary to store application and customer data. The more data stored in the ONTAP volume the less storage availability for future data. If the data storage capacity within a volume reaches the total storage capacity may lead to the customer being unable to store data due to lack of storage capacity.</p> <p>Monitoring the volume used storage capacity ensures data services continuity.</p>	<p>Immediate actions are required to minimize service disruption if critical threshold is breached:...1. Consider deleting data that is not needed anymore to free up space</p>

FSx Volume High Latency	Warning @ > 1000 µs...Critical @ > 2000 µs	Volumes are objects that serve the IO traffic often driven by performance sensitive applications including devOps applications, home directories, and databases. High volume latencies means that the applications themselves may suffer and be unable to accomplish their tasks. Monitoring volume latencies is critical to maintain application consistent performance.	Immediate actions are required to minimize service disruption if critical threshold is breached:... 1. If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled.....Plan to take the following actions soon if warning threshold is breached:...1. If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled....2. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node.
FSx Volume Inodes Limit	Warning @ > 85 %...Critical @ > 95 %	Volumes that store files use index nodes (inode) to store file metadata. When a volume exhausts its inode allocation no more files can be added to it. A warning alert indicates that planned action should be taken to increase the number of available inodes. A critical alert indicates that file limit exhaustion is imminent and emergency measures should be taken to free up inodes to ensure service continuity	Immediate actions are required to minimize service disruption if critical threshold is breached:... 1. Consider increasing the inodes value for the volume. If the inodes value is already at the max, then consider splitting the volume into two or more volumes because the file system has grown beyond the maximum size.....Plan to take the following actions soon if warning threshold is breached:...1. Consider increasing the inodes value for the volume. If the inodes value is already at the max, then consider splitting the volume into two or more volumes because the file system has grown beyond the maximum size

FSx Volume Qtree Quota Overcommit	Warning @ > 95 %...Critical @ > 100 %	Volume Qtree Quota Overcommit specifies the percentage at which a volume is considered to be overcommitted by the qtree quotas. The set threshold for the qtree quota is reached for the volume. Monitoring the volume qtree quota overcommit ensures that the user receives uninterrupted data service.	If critical threshold is breached, then immediate actions should be taken to minimize service disruption: 1. Delete unwanted data...When warning threshold is breached, then consider increasing the space of the volume.
FSx Snapshot Reserve Space is Full	Warning @ > 90 %...Critical @ > 95 %	Storage capacity of a volume is necessary to store application and customer data. A portion of that space, called snapshot reserved space, is used to store snapshots which allow data to be protected locally. The more new and updated data stored in the ONTAP volume the more snapshot capacity is used and less snapshot storage capacity will be available for future new or updated data. If the snapshot data capacity within a volume reaches the total snapshot reserve space it may lead to the customer being unable to store new snapshot data and reduction in the level of protection for the data in the volume. Monitoring the volume used snapshot capacity ensures data services continuity.	Immediate actions are required to minimize service disruption if critical threshold is breached:...1. Consider configuring snapshots to use data space in the volume when the snapshot reserve is full...2. Consider deleting some older snapshots that may not be needed anymore to free up space.....Plan to take the following actions soon if warning threshold is breached:...1. Consider increasing the snapshot reserve space within the volume to accommodate the growth...2. Consider configuring snapshots to use data space in the volume when the snapshot reserve is full

FSx Volume Cache Miss Ratio	Warning @ > 95 %...Critical @ > 100 %	Volume Cache Miss Ratio is the percentage of read requests from the client applications that are returned from the disk instead of being returned from the cache. This means that the volume has reached the set threshold.	<p>If critical threshold is breached, then immediate actions should be taken to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Move some workloads off of the node of the volume to reduce the IO load 2. Lower the demand of lower priority workloads on the same node via QoS limits...Consider immediate actions when warning threshold is breached: <ol style="list-style-type: none"> 1. Move some workloads off of the node of the volume to reduce the IO load 2. Lower the demand of lower priority workloads on the same node via QoS limits 3. Change workload characteristics (block size, application caching etc)
-----------------------------	---------------------------------------	---	--

[Back to Top](#)

K8s Monitors


Monitor Name	Description	Corrective Actions	Severity/Threshold
--------------	-------------	--------------------	--------------------

Persistent Volume Latency High	<p>High persistent volume latencies means that the applications themselves may suffer and be unable to accomplish their tasks. Monitoring persistent volume latencies is critical to maintain application consistent performance. The following are expected latencies based on media type - SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.</p>	<p>Immediate Actions If critical threshold is breached, consider immediate actions to minimize service disruption: If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled.</p> <p>Actions To Do Soon If warning threshold is breached, plan the following immediate actions: 1. If storage pool is also experiencing high utilization, move the volume to another storage pool. 2. If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled. 3. If the controller is also experiencing high utilization, move the volume to another controller or reduce the total workload of the controller.</p>	<p>Warning @ > 6,000 μs Critical @ > 12,000 μs</p>
Cluster Memory Saturation High	<p>Cluster allocatable memory saturation is high. Cluster CPU saturation is calculated as the sum of memory usage divided by the sum of allocatable memory across all K8s nodes.</p>	<p>Add nodes. Fix any unscheduled nodes. Right-size pods to free up memory on nodes.</p>	<p>Warning @ > 80 % Critical @ > 90 %</p>
POD Attach Failed	<p>This alert occurs when a volume attachment with POD is failed.</p>		<p>Warning</p>

High Retransmit Rate	High TCP Retransmit Rate	Check for Network congestion - Identify workloads that consume a lot of network bandwidth. Check for high Pod CPU utilization. Check hardware network performance.	Warning @ > 10 % Critical @ > 25 %
Node File System Capacity High	Node File System Capacity High	- Increase the size of the node disks to ensure that there is sufficient room for the application files. - Decrease application file usage.	Warning @ > 80 % Critical @ > 90 %
Workload Network Jitter High	High TCP Jitter (high latency/response time variations)	Check for Network congestion. Identify workloads that consume a lot of network bandwidth. Check for high Pod CPU utilization. Check hardware network performance	Warning @ > 30 ms Critical @ > 50 ms
Persistent Volume Throughput	MBPS thresholds on persistent volumes can be used to alert an administrator when persistent volumes exceed predefined performance expectations, potentially impacting other persistent volumes. Activating this monitor will generate alerts appropriate for the typical throughput profile of persistent volumes on SSDs. This monitor will cover all persistent volumes on your tenant. The warning and critical threshold values can be adjusted based on your monitoring goals by duplicating this monitor and setting thresholds appropriate for your storage class. A duplicated monitor can be further targeted to a subset of the persistent volumes on your tenant.	Immediate Actions If critical threshold is breached, plan immediate actions to minimize service disruption: 1. Introduce QoS MBPS limits for the volume. 2. Review the application driving the workload on the volume for anomalies. Actions To Do Soon If warning threshold is breached, plan to take the following immediate actions: 1. Introduce QoS MBPS limits for the volume. 2. Review the application driving the workload on the volume for anomalies.	Warning @ > 10,000 MB/s Critical @ > 15,000 MB/s

Container at Risk of Going OOM Killed	The container's memory limits are set too low. The container is at risk of eviction (Out of Memory Kill).	Increase container memory limits.	Warning @ > 95 %
Workload Down	Workload has no healthy pods.		Critical @ < 1
Persistent Volume Claim Failed Binding	This alert occurs when a binding is failed on a PVC.		Warning
ResourceQuota Mem Limits About to Exceed	Memory limits for Namespace are about to exceed ResourceQuota		Warning @ > 80 % Critical @ > 90 %
ResourceQuota Mem Requests About to Exceed	Memory requests for Namespace are about to exceed ResourceQuota		Warning @ > 80 % Critical @ > 90 %
Node Creation Failed	The node could not be scheduled because of a configuration error.	Check the Kubernetes event log for the cause of the configuration failure.	Critical
Persistent Volume Reclamation Failed	The volume has failed its automatic reclamation.		Warning @ > 0 B
Container CPU Throttling	The container's CPU Limits are set too low. Container processes are slowed.	Increase container CPU limits.	Warning @ > 95 % Critical @ > 98 %
Service Load Balancer Failed to Delete			Warning
Persistent Volume IOPS	IOPS thresholds on persistent volumes can be used to alert an administrator when persistent volumes exceed predefined performance expectations. Activating this monitor will generate alerts appropriate for the typical IOPS profile of persistence volumes. This monitor will cover all persistent volumes on your tenant. The warning and critical threshold values can be adjusted based on your monitoring goals by duplicating this monitor and setting thresholds appropriate for your workload.	Immediate Actions If critical threshold is breached, plan Immediate actions to minimize service disruption : 1. Introduce QoS IOPS limits for the volume. 2. Review the application driving the workload on the volume for anomalies. Actions To Do Soon If warning threshold is breached, plan the following immediate actions: 1. Introduce QoS IOPS limits for the volume. 2. Review the application driving the workload on the volume for anomalies.	Warning @ > 20,000 IO/s Critical @ > 25,000 IO/s

Service Load Balancer Failed to Update			Warning
POD Failed Mount	This alert occurs when a mount is failed on a POD.		Warning
Node PID Pressure	Available process identifiers on the (Linux) node has fallen below an eviction threshold.	Find and fix pods that generate many processes and starve the node of available process IDs. Set up PodPidsLimit to protect your node against pods or containers that spawn too many processes.	Critical @ > 0
Pod Image Pull Failure	Kubernetes failed to pull the pod container image.	<ul style="list-style-type: none"> - Make sure the pod's image is spelled correctly in the pod configuration. - Check image tag exists in your registry. - Verify the credentials for the image registry. - Check for registry connectivity issues. - Verify you are not hitting the rate limits imposed by public registry providers. 	Warning
Job Running Too Long	Job is running for too long		Warning @ > 1 hr Critical @ > 5 hr
Node Memory High	Node memory usage is high	Add nodes. Fix any unscheduled nodes. Right-size pods to free up memory on nodes.	Warning @ > 85 % Critical @ > 90 %
ResourceQuota CPU Limits About to Exceed	CPU limits for Namespace are about to exceed ResourceQuota		Warning @ > 80 % Critical @ > 90 %
Pod Crash Loop Backoff	Pod has crashed and attempted to restart multiple times.		Critical @ > 3
Node CPU High	Node CPU usage is high.	Add nodes. Fix any unscheduled nodes. Right-size pods to free up CPU on nodes.	Warning @ > 80 % Critical @ > 90 %

Workload Network Latency RTT High	High TCP RTT (Round Trip Time) latency	Check for Network congestion  Identify workloads that consume a lot of network bandwidth. Check for high Pod CPU utilization. Check hardware network performance.	Warning @ > 150 ms Critical @ > 300 ms
Job Failed	Job did not complete successfully due to a node crash or reboot, resource exhaustion, job timeout, or pod scheduling failure.	Check the Kubernetes event logs for failure causes.	Warning @ > 1
Persistent Volume Full in a Few Days	Persistent Volume will run out of space in a few days	-Increase the volume size to ensure that there is sufficient room for the application files. -Reduce the amount of data stored in applications.	Warning @ < 8 day Critical @ < 3 day
Node Memory Pressure	Node is running out of memory. Available memory has met eviction threshold.	Add nodes. Fix any unscheduled nodes. Right-size pods to free up memory on nodes.	Critical @ > 0
Node Unready	Node has been unready for 5 minutes	Verify the node have enough CPU, memory, and disk resources. Check node network connectivity. Check the Kubernetes event logs for failure causes.	Critical @ < 1
Persistent Volume Capacity High	Persistent Volume backend used capacity is high.	- Increase the volume size to ensure that there is sufficient room for the application files. - Reduce the amount of data stored in applications.	Warning @ > 80 % Critical @ > 90 %
Service Load Balancer Failed to Create	Service Load Balancer Create Failed		Critical
Workload Replica Mismatch	Some pods are currently not available for a Deployment or DaemonSet.		Warning @ > 1
ResourceQuota CPU Requests About to Exceed	CPU requests for Namespace are about to exceed ResourceQuota		Warning @ > 80 % Critical @ > 90 %

High Retransmit Rate	High TCP Retransmit Rate	Check for Network congestion - Identify workloads that consume a lot of network bandwidth. Check for high Pod CPU utilization. Check hardware network performance.	Warning @ > 10 % Critical @ > 25 %
Node Disk Pressure	Available disk space and inodes on either the node's root filesystem or image filesystem has satisfied an eviction threshold.	- Increase the size of the node disks to ensure that there is sufficient room for the application files. - Decrease application file usage.	Critical @ > 0
Cluster CPU Saturation High	Cluster allocatable CPU saturation is high. Cluster CPU saturation is calculated as the sum of CPU usage divided by the sum allocatable CPU across all K8s nodes.	Add nodes. Fix any unscheduled nodes. Right-size pods to free up CPU on nodes.	Warning @ > 80 % Critical @ > 90 %

[Back to Top](#)

Change Log Monitors

Monitor Name	Severity	Monitor Description
Internal Volume Discovered	Informational	This message occurs when an Internal Volume is discovered.
Internal Volume Modified	Informational	This message occurs when an Internal Volume is modified.
Storage Node Discovered	Informational	This message occurs when an Storage Node is discovered.
Storage Node Removed	Informational	This message occurs when an Storage Node is removed.
Storage Pool Discovered	Informational	This message occurs when an Storage Pool is discovered.
Storage Virtual Machine Discovered	Informational	This message occurs when an Storage Virtual Machine is discovered.
Storage Virtual Machine Modified	Informational	This message occurs when an Storage Virtual Machine is modified.

[Back to Top](#)

Data Collection Monitors

Monitor Name	Description	Corrective Action
Acquisition Unit Shutdown	Data Infrastructure Insights Acquisition Units periodically restart as part of upgrades to introduce new features. This happens once a month or less in a typical environment. A Warning Alert that an Acquisition Unit has shutdown should be followed soon after by a Resolution noting that the newly-restarted Acquisition Unit has completed a registration with Data Infrastructure Insights. Typically this shutdown-to-registration cycle takes 5 to 15 minutes.	If the alert occurs frequently or lasts longer than 15 minutes, check on the operation of the system hosting the Acquisition Unit, the network, and any proxy connecting the AU to the Internet.
Collector Failed	The poll of a data collector encountered an unexpected failure situation.	Visit the data collector page in Data Infrastructure Insights to learn more about the situation.
Collector Warning	This Alert typically can arise because of an erroneous configuration of the data collector or of the target system. Revisit the configurations to prevent future Alerts. It can also be due to a retrieval of less-than-complete data where the data collector gathered all the data that it could. This can happen when situations change during data collection (e.g., a virtual machine present at the beginning of data collection is deleted during data collection and before its data is captured).	<p>Check the configuration of the data collector or target system.</p> <p>Note that the monitor for Collector Warning can send more alerts than other monitor types, so it is recommended to set no alert recipients unless you are troubleshooting.</p>

[Back to Top](#)

Security Monitors

Monitor Name	Threshold	Monitor Description	Corrective Action
--------------	-----------	---------------------	-------------------

AutoSupport HTTPS transport disabled	Warning @ < 1	AutoSupport supports HTTPS, HTTP, and SMTP for transport protocols. Because of the sensitive nature of AutoSupport messages, NetApp strongly recommends using HTTPS as the default transport protocol for sending AutoSupport messages to NetApp support.	To set HTTPS as the transport protocol for AutoSupport messages, run the following ONTAP command:...system node autosupport modify -transport https
Cluster Insecure ciphers for SSH	Warning @ < 1	Indicates that SSH is using insecure ciphers, for example ciphers beginning with *cbc.	To remove the CBC ciphers, run the following ONTAP command:...security ssh remove -vserver <admin vserver> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc
Cluster Login Banner Disabled	Warning @ < 1	Indicates that the Login banner is disabled for users accessing the ONTAP system. Displaying a login banner is helpful for establishing expectations for access and use of the system.	To configure the login banner for a cluster, run the following ONTAP command:...security login banner modify -vserver <admin svm> -message "Access restricted to authorized users"
Cluster Peer Communication Not Encrypted	Warning @ < 1	When replicating data for disaster recovery, caching, or backup, you must protect that data during transport over the wire from one ONTAP cluster to another. Encryption must be configured on both the source and destination clusters.	To enable encryption on cluster peer relationships that were created prior to ONTAP 9.6, the source and destination cluster must be upgraded to 9.6. Then use the "cluster peer modify" command to change both the source and destination cluster peers to use Cluster Peering Encryption....See the NetApp Security Hardening Guide for ONTAP 9 for details.

Default Local Admin User Enabled	Warning @ > 0	NetApp recommends locking (disabling) any unneeded Default Admin User (built-in) accounts with the lock command. They are primarily default accounts for which passwords were never updated or changed.	To lock the built-in "admin" account, run the following ONTAP command:...security login lock -username admin
FIPS Mode Disabled	Warning @ < 1	When FIPS 140-2 compliance is enabled, TLSv1 and SSLv3 are disabled, and only TLSv1.1 and TLSv1.2 remain enabled. ONTAP prevents you from enabling TLSv1 and SSLv3 when FIPS 140-2 compliance is enabled.	To enable FIPS 140-2 compliance on a cluster, run the following ONTAP command in advanced privilege mode:...security config modify -interface SSL -is-fips-enabled true
Log Forwarding Not Encrypted	Warning @ < 1	Offloading of syslog information is necessary for limiting the scope or footprint of a breach to a single system or solution. Therefore, NetApp recommends securely offloading syslog information to a secure storage or retention location.	Once a log forwarding destination is created, its protocol cannot be changed. To change to an encrypted protocol, delete and recreate the log forwarding destination using the following ONTAP command:...cluster log-forwarding create -destination <destination ip> -protocol tcp-encrypted
MD5 Hashed password	Warning @ > 0	NetApp strongly recommends to use the more secure SHA-512 hash function for ONTAP user account passwords. Accounts using the less secure MD5 hash function should migrate to the SHA-512 hash function.	NetApp strongly recommends user accounts migrate to the more secure SHA-512 solution by having users change their passwords....to lock accounts with passwords that use the MD5 hash function, run the following ONTAP command:...security login lock -vserver * -username * -hash-function md5

No NTP servers are configured	Warning @ < 1	Indicates that the cluster has no configured NTP servers. For redundancy and optimum service, NetApp recommends that you associate at least three NTP servers with the cluster.	To associate an NTP server with the cluster, run the following ONTAP command: cluster time-service ntp server create -server <ntp server host name or ip address>
NTP server count is low	Warning @ < 3	Indicates that the cluster has less than 3 configured NTP servers. For redundancy and optimum service, NetApp recommends that you associate at least three NTP servers with the cluster.	To associate an NTP server with the cluster, run the following ONTAP command:...cluster time-service ntp server create -server <ntp server host name or ip address>
Remote Shell Enabled	Warning @ > 0	Remote Shell is not a secure method for establishing command-line access to the ONTAP solution. Remote Shell should be disabled for secure remote access.	NetApp recommends Secure Shell (SSH) for secure remote access....To disable Remote shell on a cluster, run the following ONTAP command in advanced privilege mode:...security protocol modify -application rsh- enabled false
Storage VM Audit Log Disabled	Warning @ < 1	Indicates that Audit logging is disabled for SVM.	To configure the Audit log for a vserver, run the following ONTAP command:...vserver audit enable -vserver <svm>
Storage VM Insecure ciphers for SSH	Warning @ < 1	Indicates that SSH is using insecure ciphers, for example ciphers beginning with *cbc.	To remove the CBC ciphers, run the following ONTAP command:...security ssh remove -vserver <vserver> -ciphers aes256-cbc,aes192-cbc,aes128-cbc,3des-cbc
Storage VM Login banner disabled	Warning @ < 1	Indicates that the Login banner is disabled for users accessing SVMs on the system. Displaying a login banner is helpful for establishing expectations for access and use of the system.	To configure the login banner for a cluster, run the following ONTAP command:...security login banner modify -vserver <svm> -message "Access restricted to authorized users"

Telnet Protocol Enabled	Warning @ > 0	Telnet is not a secure method for establishing command-line access to the ONTAP solution. Telnet should be disabled for secure remote access.	NetApp recommends Secure Shell (SSH) for secure remote access. To disable Telnet on a cluster, run the following ONTAP command in advanced privilege mode: ...security protocol modify -application telnet -enabled false
-------------------------	---------------	---	---

[Back to Top](#)

Data Protection Monitors

Monitor Name	Thresholds	Monitor Description	Corrective Action
Insufficient Space for Lun Snapshot Copy	(Filter contains _luns = Yes) Warning @ > 95 %...Critical @ > 100 %	Storage capacity of a volume is necessary to store application and customer data. A portion of that space, called snapshot reserved space, is used to store snapshots which allow data to be protected locally. The more new and updated data stored in the ONTAP volume the more snapshot capacity is used and less snapshot storage capacity will be available for future new or updated data. If the snapshot data capacity within a volume reaches the total snapshot reserve space it may lead to the customer being unable to store new snapshot data and reduction in the level of protection for the data in the LUNs in the volume. Monitoring the volume used snapshot capacity ensures data services continuity.	<p>Immediate Actions</p> <p>If critical threshold is breached, consider immediate actions to minimize service disruption:</p> <ol style="list-style-type: none"> 1. Configure snapshots to use data space in the volume when the snapshot reserve is full. 2. Delete some older unwanted snapshots to free up space. <p>Actions To Do Soon</p> <p>If warning threshold is breached, plan to take the following immediate actions:</p> <ol style="list-style-type: none"> 1. Increase the snapshot reserve space within the volume to accommodate the growth. 2. Configure snapshots to use data space in the volume when the snapshot reserve is full.

SnapMirror Relationship Lag	Warning @ > 150%...Critical @ > 300%	SnapMirror relationship lag is the difference between the snapshot timestamp and the time on the destination system. The lag_time_percent is the ratio of lag time to the SnapMirror Policy's schedule interval. If the lag time equals the schedule interval, the lag_time_percent will be 100%. If the SnapMirror policy does not have a schedule, lag_time_percent will not be calculated.	Monitor SnapMirror status using the "snapmirror show" command. Check the SnapMirror transfer history using the "snapmirror show-history" command
-----------------------------	--------------------------------------	---	--

[Back to Top](#)

Cloud Volume (CVO) Monitors

Monitor Name	CI Severity	Monitor Description	Corrective Action
CVO Disk Out of Service	INFO	This event occurs when a disk is removed from service because it has been marked failed, is being sanitized, or has entered the Maintenance Center.	None

CVO Giveback of Storage Pool Failed	CRITICAL	<p>This event occurs during the migration of an aggregate as part of a storage failover (SFO) giveback, when the destination node cannot reach the object stores.</p>	<p>Perform the following corrective actions:</p> <p>Verify that your intercluster LIF is online and functional by using the "network interface show" command.</p> <p>Check network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF.</p> <p>Verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the "aggregate object-store config show" command.</p> <p>Alternatively, you can override the error by specifying false for the "require-partner-waiting" parameter of the giveback command.</p> <p>Contact NetApp technical support for more information or assistance.</p>
-------------------------------------	----------	---	---

CVO HA Interconnect Down	WARNING	The high-availability (HA) interconnect is down. Risk of service outage when failover is not available.	<p>Corrective actions depend on the number and type of HA interconnect links supported by the platform, as well as the reason why the interconnect is down.</p> <p>If the links are down:</p> <p>Verify that both controllers in the HA pair are operational.</p> <p>For externally connected links, make sure that the interconnect cables are connected properly and that the small form-factor pluggables (SFPs), if applicable, are seated properly on both controllers.</p> <p>For internally connected links, disable and re-enable the links, one after the other, by using the "ic link off" and "ic link on" commands.</p> <p>If links are disabled, enable the links by using the "ic link on" command.</p> <p>If a peer is not connected, disable and re-enable the links, one after the other, by using the "ic link off" and "ic link on" commands.</p> <p>Contact NetApp technical support if the issue persists.</p>
--------------------------	---------	---	---

CVO Max Sessions Per User Exceeded	WARNING	<p>You have exceeded the maximum number of sessions allowed per user over a TCP connection. Any request to establish a session will be denied until some sessions are released.</p>	<p>Perform the following corrective actions:</p> <p>Inspect all the applications that run on the client, and terminate any that are not operating properly.</p> <p>Reboot the client.</p> <p>Check if the issue is caused by a new or existing application:</p> <p>If the application is new, set a higher threshold for the client by using the "cifs option modify -max-opens -same-file-per-tree" command.</p> <p>In some cases, clients operate as expected, but require a higher threshold. You should have advanced privilege to set a higher threshold for the client.</p> <p>If the issue is caused by an existing application, there might be an issue with the client. Contact NetApp technical support for more information or assistance.</p>
------------------------------------	---------	---	---

CVO NetBIOS Name Conflict	CRITICAL	<p>The NetBIOS Name Service has received a negative response to a name registration request, from a remote machine. This is typically caused by a conflict in the NetBIOS name or an alias. As a result, clients might not be able to access data or connect to the right data-serving node in the cluster.</p>	<p>Perform any one of the following corrective actions:</p> <p>If there is a conflict in the NetBIOS name or an alias, perform one of the following:</p> <p>Delete the duplicate NetBIOS alias by using the "vserver cifs delete -aliases alias -vserver vserver" command.</p> <p>Rename a NetBIOS alias by deleting the duplicate name and adding an alias with a new name by using the "vserver cifs create -aliases alias -vserver vserver" command.</p> <p>If there are no aliases configured and there is a conflict in the NetBIOS name, then rename the CIFS server by using the "vserver cifs delete -vserver vserver" and "vserver cifs create -cifs -server netbiosname" commands.</p> <p>NOTE: Deleting a CIFS server can make data inaccessible.</p> <p>Remove NetBIOS name or rename the NetBIOS on the remote machine.</p>
CVO NFSv4 Store Pool Exhausted	CRITICAL	A NFSv4 store pool has been exhausted.	If the NFS server is unresponsive for more than 10 minutes after this event, contact NetApp technical support.
CVO Node Panic	WARNING	This event is issued when a panic occurs	Contact NetApp customer support.

CVO Node Root Volume Space Low	CRITICAL	<p>The system has detected that the root volume is dangerously low on space. The node is not fully operational. Data LIFs might have failed over within the cluster, because of which NFS and CIFS access is limited on the node.</p> <p>Administrative capability is limited to local recovery procedures for the node to clear up space on the root volume.</p>	<p>Perform the following corrective actions:</p> <p>Clear up space on the root volume by deleting old Snapshot copies, deleting files you no longer need from the /mroot directory, or expanding the root volume capacity.</p> <p>Reboot the controller.</p> <p>Contact NetApp technical support for more information or assistance.</p>
CVO Nonexistent Admin Share	CRITICAL	<p>Vscan issue: a client has attempted to connect to a nonexistent ONTAP_ADMIN\$ share.</p>	<p>Ensure that Vscan is enabled for the mentioned SVM ID. Enabling Vscan on a SVM causes the ONTAP_ADMIN\$ share to be created for the SVM automatically.</p>
CVO Object Store Host Unresolvable	CRITICAL	<p>The object store server host name cannot be resolved to an IP address. The object store client cannot communicate with the object-store server without resolving to an IP address. As a result, data might be inaccessible.</p>	<p>Check the DNS configuration to verify that the host name is configured correctly with an IP address.</p>
CVO Object Store Intercluster LIF Down	CRITICAL	<p>The object-store client cannot find an operational LIF to communicate with the object store server. The node will not allow object store client traffic until the intercluster LIF is operational. As a result, data might be inaccessible.</p>	<p>Perform the following corrective actions:</p> <p>Check the intercluster LIF status by using the "network interface show -role intercluster" command.</p> <p>Verify that the intercluster LIF is configured correctly and operational.</p> <p>If an intercluster LIF is not configured, add it by using the "network interface create -role intercluster" command.</p>

CVO Object Store Signature Mismatch	CRITICAL	The request signature sent to the object store server does not match the signature calculated by the client. As a result, data might be inaccessible.	Verify that the secret access key is configured correctly. If it is configured correctly, contact NetApp technical support for assistance.
CVO QoS Monitor Memory Maxed Out	CRITICAL	The QoS subsystem's dynamic memory has reached its limit for the current platform hardware. Some QoS features might operate in a limited capacity.	Delete some active workloads or streams to free up memory. Use the "statistics show -object workload -counter ops" command to determine which workloads are active. Active workloads show non-zero ops. Then use the "workload delete <workload_name>" command multiple times to remove specific workloads. Alternatively, use the "stream delete -workload <workload name> *" command to delete the associated streams from the active workload.

CVO READDIR Timeout	CRITICAL	<p>A READDIR file operation has exceeded the timeout that it is allowed to run in WAFL. This can be because of very large or sparse directories. Corrective action is recommended.</p>	<p>Perform the following corrective actions:</p> <p>Find information specific to recent directories that have had READDIR file operations expire by using the following 'diag' privilege nodeshell CLI command: wafl readdir notice show.</p> <p>Check if directories are indicated as sparse or not:</p> <p>If a directory is indicated as sparse, it is recommended that you copy the contents of the directory to a new directory to remove the sparseness of the directory file.</p> <p>If a directory is not indicated as sparse and the directory is large, it is recommended that you reduce the size of the directory file by reducing the number of file entries in the directory.</p>
---------------------	----------	--	---

CVO Relocation of Storage Pool Failed	CRITICAL	This event occurs during the relocation of an aggregate, when the destination node cannot reach the object stores.	<p>Perform the following corrective actions:</p> <p>Verify that your intercluster LIF is online and functional by using the "network interface show" command.</p> <p>Check network connectivity to the object store server by using the "ping" command over the destination node intercluster LIF.</p> <p>Verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the "aggregate object-store config show" command.</p> <p>Alternatively, you can override the error by using the "override-destination-checks" parameter of the relocation command.</p> <p>Contact NetApp technical support for more information or assistance.</p>
---------------------------------------	----------	--	--

CVO Shadow Copy Failed	CRITICAL	A Volume Shadow Copy Service (VSS), a Microsoft Server backup and restore service operation, has failed.	<p>Check the following using the information provided in the event message:</p> <p>Is shadow copy configuration enabled?</p> <p>Are the appropriate licenses installed?</p> <p>On which shares is the shadow copy operation performed?</p> <p>Is the share name correct?</p> <p>Does the share path exist?</p> <p>What are the states of the shadow copy set and its shadow copies?</p>
CVO Storage VM Stop Succeeded	INFO	This message occurs when a 'vserver stop' operation succeeds.	Use 'vserver start' command to start the data access on a storage VM.
CVO Too Many CIFS Authentication	WARNING	Many authentication negotiations have occurred simultaneously. There are 256 incomplete new session requests from this client.	Investigate why the client has created 256 or more new connection requests. You might have to contact the vendor of the client or of the application to determine why the error occurred.
CVO Unassigned Disks	INFO	System has unassigned disks - capacity is being wasted and your system may have some misconfiguration or partial configuration change applied.	<p>Perform the following corrective actions:</p> <p>Determine which disks are unassigned by using the "disk show -n" command.</p> <p>Assign the disks to a system by using the "disk assign" command.</p>

CVO Unauthorized User Access to Admin Share	WARNING	A client has attempted to connect to the privileged ONTAP_ADMIN\$ share even though their logged-in user is not an allowed user.	<p>Perform the following corrective actions:</p> <p>Ensure that the mentioned username and IP address is configured in one of the active Vscan scanner pools.</p> <p>Check the scanner pool configuration that is currently active by using the "vserver vscan scanner pool show-active" command.</p>
CVO Virus Detected	WARNING	<p>A Vscan server has reported an error to the storage system. This typically indicates that a virus has been found. However, other errors on the Vscan server can cause this event.</p> <p>Client access to the file is denied. The Vscan server might, depending on its settings and configuration, clean the file, quarantine it, or delete it.</p>	Check the log of the Vscan server reported in the "syslog" event to see if it was able to successfully clean, quarantine, or delete the infected file. If it was not able to do so, a system administrator might have to manually delete the file.
CVO Volume Offline	INFO	This message indicates that a volume is made offline.	Bring the volume back online.
CVO Volume Restricted	INFO	This event indicates that a flexible volume is made restricted.	Bring the volume back online.

[Back to Top](#)

SnapMirror for Business Continuity (SMBC) Mediator Log Monitors

Monitor Name	Severity	Monitor Description	Corrective Action
ONTAP Mediator Added	INFO	This message occurs when ONTAP Mediator is added successfully on a cluster.	None

ONTAP Mediator Not Accessible	CRITICAL	This message occurs when either the ONTAP Mediator is repurposed or the Mediator package is no longer installed on the Mediator server. As a result, SnapMirror failover is not possible.	Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
ONTAP Mediator Removed	INFO	This message occurs when ONTAP Mediator is removed successfully from a cluster.	None
ONTAP Mediator Unreachable	WARNING	This message occurs when the ONTAP Mediator is unreachable on a cluster. As a result, SnapMirror failover is not possible.	Check the network connectivity to the ONTAP Mediator by using the "network ping" and "network traceroute" commands. If the issue persists, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
SMBC CA Certificate Expired	CRITICAL	This message occurs when the ONTAP Mediator certificate authority (CA) certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.	Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new CA certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

SMBC CA Certificate Expiring	WARNING	This message occurs when the ONTAP Mediator certificate authority (CA) certificate is due to expire within the next 30 days.	Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new CA certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
SMBC Client Certificate Expired	CRITICAL	This message occurs when the ONTAP Mediator client certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.	Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
SMBC Client Certificate Expiring	WARNING	This message occurs when the ONTAP Mediator client certificate is due to expire within the next 30 days.	Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
SMBC Relationship Out of Sync Note: UM doesn't have this one	CRITICAL	This message occurs when a SnapMirror for Business Continuity (SMBC) relationship changes status from "in-sync" to "out-of-sync". Due to this RPO=0 data protection will be disrupted.	Check the network connection between the source and destination volumes. Monitor the SMBC relationship status by using the "snapmirror show" command on the destination, and by using the "snapmirror list-destinations" command on the source. Auto-resync will attempt to bring the relationship back to "in-sync" status. If the resync fails, verify that all the nodes in the cluster are in quorum and are healthy.

SMBC Server Certificate Expired	CRITICAL	This message occurs when the ONTAP Mediator server certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.	Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new server certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
SMBC Server Certificate Expiring	WARNING	This message occurs when the ONTAP Mediator server certificate is due to expire within the next 30 days.	Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new server certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

[Back to Top](#)

Additional Power, Heartbeat, and Miscellaneous System Monitors

Monitor Name	Severity	Monitor Description	Corrective Action
Disk Shelf Power Supply Discovered	INFORMATIONAL	This message occurs when a power supply unit is added to the disk shelf.	NONE
Disk Shelves Power Supply Removed	INFORMATIONAL	This message occurs when a power supply unit is removed from the disk shelf.	NONE
MetroCluster Automatic Unplanned Switchover Disabled	CRITICAL	This message occurs when automatic unplanned switchover capability is disabled.	Run the "metrocluster modify -node-name <nodename> -automatic -switchover-onfailure true" command for each node in the cluster to enable automatic switchover.

Monitor Name	Severity	Monitor Description	Corrective Action
MetroCluster Storage Bridge Unreachable	CRITICAL	The storage bridge is not reachable over the management network	1) If the bridge is monitored by SNMP, verify that the node management LIF is up by using the "network interface show" command. Verify that the bridge is alive by using the "network ping" command. 2) If the bridge is monitored in-band, check the fabric cabling to the bridge, and then verify that the bridge is powered up.
MetroCluster Bridge Temperature Abnormal - Below Critical	CRITICAL	The sensor on the Fibre Channel bridge is reporting a temperature that is below the critical threshold.	1) Check the operational status of the fans on the storage bridge. 2) Verify that the bridge is operating under recommended temperature conditions.
MetroCluster Bridge Temperature Abnormal - Above Critical	CRITICAL	The sensor on the Fibre Channel bridge is reporting a temperature that is above the critical threshold.	1) Check the operational status of the chassis temperature sensor on the storage bridge using the command "storage bridge show -cooling". 2) Verify that the storage bridge is operating under recommended temperature conditions.
MetroCluster Aggregate Left Behind	WARNING	The aggregate was left behind during switchback.	1) Check the aggregate state by using the command "aggr show". 2) If the aggregate is online, return it to its original owner by using the command "metrocluster switchback".

Monitor Name	Severity	Monitor Description	Corrective Action
All Links Between Metrocluster Partners Down	CRITICAL	RDMA interconnect adapters and intercluster LIFs have broken connections to the peered cluster or the peered cluster is down.	<ol style="list-style-type: none"> 1) Ensure that the intercluster LIFs are up and running. Repair the intercluster LIFs if they are down. 2) Verify that the peered cluster is up and running by using the "cluster peer ping" command. See the MetroCluster Disaster Recovery Guide if the peered cluster is down. 3) For fabric MetroCluster, verify that the back-end fabric ISLs are up and running. Repair the back-end fabric ISLs if they are down. 4) For non-fabric MetroCluster configurations, verify that the cabling is correct between the RDMA interconnect adapters. Reconfigure the cabling if the links are down.
MetroCluster Partners Not Reachable Over Peering Network	CRITICAL	The connectivity to the peer cluster is broken.	<ol style="list-style-type: none"> 1) Ensure that the port is connected to the correct network/switch. 2) Ensure that the intercluster LIF is connected with the peered cluster. 3) Ensure that the peered cluster is up and running by using the command "cluster peer ping". Refer to the MetroCluster Disaster Recovery Guide if the peered cluster is down.
MetroCluster Inter Switch All Links Down	CRITICAL	All Inter-Switch Links (ISLs) on the storage switch are down.	<ol style="list-style-type: none"> 1) Repair the back-end fabric ISLs on the storage switch. 2) Ensure that the partner switch is up and its ISLs are operational. 3) Ensure that intermediate equipment, such as xWDM devices, are operational.

Monitor Name	Severity	Monitor Description	Corrective Action
MetroCluster Node To Storage Stack SAS Link Down	WARNING	The SAS adapter or its attached cable might be at fault.	<ol style="list-style-type: none"> 1. Verify that the SAS adapter is online and running. 2. Verify that the physical cable connection is secure and operating, and replace the cable if necessary. 3. If the SAS adapter is connected to disk shelves, ensure IOMs and disks are properly seated.
MetroClusterFC Initiator Links Down	CRITICAL	The FC initiator adapter is at fault.	<ol style="list-style-type: none"> 1. Ensure that the FC initiator link has not been tampered with. 2. Verify the operational status of the FC initiator adapter by using the command "system node run -node local -command storage show adapter".
FC-VI Interconnect Link Down	CRITICAL	The physical link on the FC-VI port is offline.	<ol style="list-style-type: none"> 1. Ensure that the FC-VI link has not been tampered with. 2. Verify that the physical status of the FC-VI adapter is "Up" by using the command "metrocluster interconnect adapter show". 3. If the configuration includes fabric switches, ensure that they are properly cabled and configured.
MetroCluster Spare Disks Left Behind	WARNING	The spare disk was left behind during switchback.	If the disk is not failed, return it to its original owner by using the command "metrocluster switchback".
MetroCluster Storage Bridge Port Down	CRITICAL	The port on the storage bridge is offline.	<ol style="list-style-type: none"> 1) Check the operational status of the ports on the storage bridge by using the command "storage bridge show -ports". 2) Verify logical and physical connectivity to the port.

Monitor Name	Severity	Monitor Description	Corrective Action
MetroCluster Storage Switch Fans Failed	CRITICAL	The fan on the storage switch failed.	1) Ensure that the fans in the switch are operating correctly by using the command "storage switch show -cooling". 2) Ensure that the fan FRUs are properly inserted and operational.
MetroCluster Storage Switch Unreachable	CRITICAL	The storage switch is not reachable over the management network.	1) Ensure that the node management LIF is up by using the command "network interface show". 2) Ensure that the switch is alive by using the command "network ping". 3) Ensure that the switch is reachable over SNMP by checking its SNMP settings after logging into the switch.
MetroCluster Switch Power Supplies Failed	CRITICAL	A power supply unit on the storage switch is not operational.	1) Check the error details by using the command "storage switch show -error -switch-name <switch name>". 2) Identify the faulty power supply unit by using the command "storage switch show -power -switch -name <switch name>". 3) Ensure that the power supply unit is properly inserted into the chassis of the storage switch and fully operational.
MetroCluster Switch Temperature Sensors Failed	CRITICAL	The sensor on the Fibre Channel switch failed.	1) Check the operational status of the temperature sensors on the storage switch by using the command "storage switch show -cooling". 2) Verify that the switch is operating under recommended temperature conditions.

Monitor Name	Severity	Monitor Description	Corrective Action
MetroCluster Switch Temperature Abnormal	CRITICAL	The temperature sensor on the Fibre Channel switch reported abnormal temperature.	1) Check the operational status of the temperature sensors on the storage switch by using the command "storage switch show -cooling". 2) Verify that the switch is operating under recommended temperature conditions.
Service Processor Heartbeat Missed	INFORMATIONAL	This message occurs when ONTAP does not receive an expected "heartbeat" signal from the Service Processor (SP). Along with this message, log files from SP will be sent out for debugging. ONTAP will reset the SP to attempt to restore communication. The SP will be unavailable for up to two minutes while it reboots.	Contact NetApp technical support.

Monitor Name	Severity	Monitor Description	Corrective Action
Service Processor Heartbeat Stopped	WARNING	This message occurs when ONTAP is no longer receiving heartbeats from the Service Processor (SP). Depending on the hardware design, the system may continue to serve data or may determine to shut down to prevent data loss or hardware damage. The system continues to serve data, but because the SP might not be working, the system cannot send notifications of down appliances, boot errors, or Open Firmware (OFW) Power-On Self-Test (POST) errors. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.	If the system has shut down, attempt a hard power cycle: Pull the controller out from the chassis, push it back in then power on the system. Contact NetApp technical support if the problem persists after the power cycle, or for any other condition that may warrant attention.

[Back to Top](#)

More Information

- [Viewing and Dismissing Alerts](#)

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 and select the *Email* tab.

Subscription Notification Recipients

Send subscription related notifications to the following:

- ☒ All Account Owners
- ☒ All Monitor & Optimize Administrators
- ☒ Additional Email Addresses

name@email.com X

Save

Global Monitor Notification Recipients

Default email recipients for monitor related notifications:

- ☐ All Account Owners
- ☒ All Monitor & Optimize Administrators
- ☐ Additional Email Addresses

Save

Subscription Notification Recipients

To configure recipients for subscription-related event notifications, go to the "Subscription Notification Recipients" section.

You can choose to have email notifications sent for subscription-related events to any or all of the following recipients:

- All Account Owners
- All *Monitor & Optimize* Administrators
- Additional Email Addresses that you specify

The following are examples of the types of notifications that might be sent, and user actions you can take.

Notification:	User Action:
Trial or subscription has been updated	Review subscription details on the 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



To ensure your recipients receive notifications from Data Infrastructure Insights, add the following email addresses to any "allow" lists:

- accounts@service.cloudinsights.netapp.com
- DoNotReply@cloudinsights.netapp.com

Global Recipient List for Alerts

Email notifications of alerts are sent to the alert recipient list for every action on the alert. You can choose to send alert notifications to a global recipient list.

To configure global alert recipients, choose the desired recipients in the **Global Monitor Notification Recipients** section.


You can always override the global recipients list for an individual monitor when creating or modifying the monitor.



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

Editing Notifications for ONTAP

You can modify notifications for ONTAP clusters by selecting *Edit Notifications* from the upper-right drop-down on a Storage landing page.



From here, you can set notifications for Critical, Warning, Informational, and/or Resolved alerts. Each scenario can notify the Global Recipient list or other recipients you choose.

☒ By Email

Notify team on

Critical, Warn... ▼

Send to

- ☐ Global Monitor Recipient List
- ☒ Other Email Recipients



email@email.one ✕

email2@email2.two ✕ |

Notify team on

Resolved ▼

Send to

- ☒ Global Monitor Recipient List
- ☐ Other Email Recipients

☐ By Webhook

Enable webhook notification to add recipients

Webhook Notifications

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, Data Infrastructure 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 along with the url params
Method	GET, POST - Default is POST
Custom Header	Specify any custom header lines here
Message Body	Put the body of your message here
Default Alert Parameters	Lists the default parameters for the webhook
Custom Parameters and Secrets	Custom parameters and secrets allow you to add unique parameters and secure elements such as passwords

Creating a Webhook

To create a Data Infrastructure 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/<token>

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%%*"
      }
    }
  ],
  "r"
```

Cancel

Test Webhook

Save Webhook

Enter appropriate information for each of the fields, and click "Save" when complete.

You can also click the "Test Webhook" button to test the connection. Note that this will send the "Message Body" (without substitutions) to the defined URL according to the selected Method.

Data Infrastructure 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.

You can add any object attribute (for example, storage name) as a parameter to a webhook. For example, you can set parameters for volume name and storage name in a webhook description like: "High Latency for Volume: `%%relatedObject.volume.name%%`, Storage: `%%relatedObject.storage.name%%`".

Note that in this section, substitutions are *not* performed when clicking the "Test Webhook" button; the button

sends a payload that shows the %% substitutions but does not replace them with data.

Custom Parameters and Secrets

In this section you can add any custom parameters and/or secrets you wish. For security reasons, if a secret is defined only the webhook creator can modify this webhook channel. It is read-only for others. You can use secrets in URL/Headers as %%<secret_name>%%.

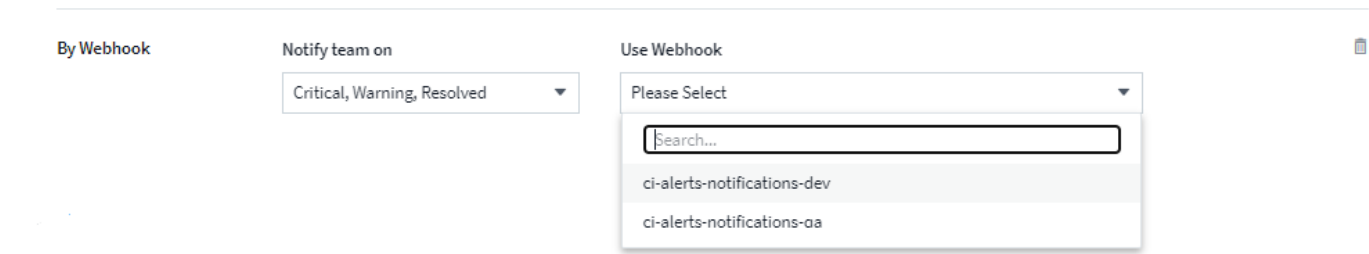
Webhooks List Page

On the Webhooks list page, displayed are the Name, Created By, Created On, Status, Secure, and Last Reported fields.

Choosing Webhook Notification in a Monitor

To choose the webhook notification in a [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)



Webhook Examples:

Webhooks for [Slack](#)
Webhooks for [PagerDuty](#)
Webhooks for [Teams](#)
Webhooks for [Discord](#)

Webhook Example for Discord

Webhooks allow users to send alert notifications to various applications using a customized webhook channel. This page provides an example for setting up webhooks for Discord.



This page refers to third-party instructions, which could be subject to change. Refer to the [Discord documentation](#) for the most up-to-date information.

Discord Setup:

- In Discord, select the Server, under Text Channels, select Edit Channel (gear icon)
- Select **Integrations > View Webhooks** and click **New Webhook**
- Copy the Webhook URL. You will need to paste this into the Data Infrastructure Insights webhook configuration.

Create Data Infrastructure Insights Webhook:

1. In Data Infrastructure Insights, navigate to **Admin > Notifications** and select the **Webhooks** tab. Click **+Webhook** to create a new webhook.
2. Give the webhook a meaningful Name, such as "Discord".
3. In the *Template Type* drop-down, select **Discord**.
4. Paste the URL from above into the *URL* field.

Edit a Webhook

Name

Discord Webhook

Template Type

Discord ▼

URL

https://discord.com/api/webhooks/<token string>

Method

POST ▼

Custom Header

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

Message Body

```
{
  "content": null,
  "embeds": [
    {
      "title": "%%severity%% | %%alertId%% | %%triggeredOn%%",
      "description": "%%monitorName%%",
      "url": "https://%%cloudInsightsHostname%%/%%alertRelativeUrl%%",
      "color": 3244733,
      "fields": [
        {
          "name": "%%metricName%%"
```

Cancel

Test Webhook

Save Webhook



In order to test the webhook, temporarily replace the url value in the message body with any valid URL (such as <https://netapp.com>) then click the *Test Webhook* button. Be sure to set the message body back once the test completes.

Notifications via Webhook

To notify on events via webhook, in Data Infrastructure Insights navigate to **Alerts > Monitors** and click **+Monitor** to create a new [monitor](#).

- Select a metric and define the monitor's conditions.
- Under *_Set up team notification(s)*, choose the **Webhook** Delivery Method.
- Choose the "Discord" webhook for the desired events (Critical, Warning, Resolved)

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

The screenshot shows a configuration interface for team notifications. It has three main sections: 'By Webhook', 'Notify team on', and 'Use Webhook(s)'. The 'By Webhook' section is active. The 'Notify team on' section has a dropdown menu showing 'Critical, Warning, Resolved'. The 'Use Webhook(s)' section has a button labeled 'Discord' and a close icon (X).

Webhook Example for PagerDuty

Webhooks allow users to send alert notifications to various applications using a customized webhook channel. This page provides an example for setting up webhooks for PagerDuty.



This page refers to third-party instructions, which could be subject to change. Refer to the [PagerDuty documentation](#) for the most up-to-date information.

PagerDuty Setup:

1. In PagerDuty, navigate to **Services > Service Directory** and click on the **+New Service** button
2. Enter in a *Name* and select *Use our API directly*. Click on *Add Service*.

Add a Service

A service may represent an application, component or team you wish to open incidents against.

General Settings


Name


Description

Integration Settings

Connect with one of PagerDuty's supported integrations, or create a custom integration through email or API. Alerts from a service from a supported integration or through the Events V2 API.

You can add more than one integration to a service, for example, one for monitoring alerts and one for [change events](#).

Integration Type 

☐ Select a tool 


PagerDuty integrates with hundreds of tools, including monitoring tools, ticketing systems, code repositories, and deploy pipelines. This may involve configuration steps in the tool you are integrating with PagerDuty.

☐ Integrate via email

If your monitoring tool can send email, it can integrate with PagerDuty using a custom email address.

☒ Use our API directly

If you're writing your own integration, use our Events API. More information is in our developer documentation.



☐ Don't use an integration

If you only want incidents to be manually created. You can always add additional integrations later.

3. Click on the *Integrations* tab to see the **Integration Key**. You will need this key when you create the Data Infrastructure Insights webhook below.

1. Go to **Incidents** or **Services** to view Alerts.


PagerDuty [Incidents](#) [Services](#) [People](#) [Analytics](#) [Status](#)


Incidents on All Teams

Your open incidents: 8 triggered, 2 acknowledged

All open incidents: 8 triggered, 2 acknowledged

1 acknowledged 20 triggered 47 resolved 10 search

Go to incident at:  All Teams

Open Triggered Acknowledged Resolved Any Status  Assigned to me All

<input type="checkbox"/>	Status	Urgency	Title	Created	Service	Assigned To
<input type="checkbox"/>	Triggered	High	incident1 / AL18 / aggregate_name_team02xall id: incident181x123 (triggered)	at 5:45 PM	Test3	Edwin Chung
<input type="checkbox"/>	Triggered	High	incident2 / AL20 / aggregate_name_team02xall id: incident201x123 (triggered)	at 5:45 PM	Test3	Edwin Chung
<input type="checkbox"/>	Triggered	High	incident3 / AL19 / aggregate_name_team02xall id: incident191x123 (triggered)	at 5:45 PM	Test3	Edwin Chung
<input type="checkbox"/>	Triggered	High	incident4 / AL17 / aggregate_name_team02xall id: incident171x123 (triggered)	at 5:45 PM	Test3	Edwin Chung
<input type="checkbox"/>	Triggered	High	incident5 / AL16 / aggregate_name_team02xall id: incident161x123 (triggered)	at 5:22 PM	Test3	Edwin Chung
<input type="checkbox"/>	Triggered	High	incident6 / AL15 / aggregate_name_team02xall id: incident151x123 (triggered)	at 5:17 PM	Alerts	Edwin Chung


Create Data Infrastructure Insights Webhook:

1. In Data Infrastructure Insights, navigate to **Admin > Notifications** and select the **Webhooks** tab. Click **+Webhook** to create a new webhook.
 2. Give the webhook a meaningful Name, such as "PagerDuty Trigger". You will use this webhook for critical- and warning-level events.
 3. In the *Template Type* drop-down, select **PagerDuty**.
1. Create a custom parameter secret named *routingKey* and set the value to the PagerDuty *Integration Key* value from above.

Custom Parameters and Secrets

Name	Value ↑	Description
%%routingKey%%	*****	⋮

+ Parameter

Name 	Value
<input type="text" value="routingKey"/>	<input type="text" value="*****"/>
Type	Description
<input type="text" value="Secret"/>	<div></div>

Cancel

Save Parameter

Repeat these steps to create a "PagerDuty Resolve" webhook for resolved events.

PagerDuty to Data Infrastructure Insights Field Mapping

The following table and image show the mapping of fields between PagerDuty and Data Infrastructure Insights:

PagerDuty	Data Infrastructure Insights
Alert Key	Alert ID
Source	Triggered On
Component	Metric Name
Group	Object Type
Class	Monitor Name

Message Body

```
{
  "dedup_key": "%%alertId%%",
  "event_action": "trigger",
  "links": [
    {
      "href": "https://%%cloudInsightsHostname%%/%%alertRelativeUrl%%",
      "text": "'%%metricName%%' value of %%value%% (%%alertCondition%%) for %%triggeredOn%%"
    }
  ],
  "payload": {
    "class": "%%monitorName%%",
    "component": "%%metricName%%",
    "group": "%%objectType%%",
    "severity": "critical",
    "source": "%%triggeredOn%%",
    "summary": "%%severity%% | %%alertId%% | %%triggeredOn%%"
  },
  "routing_key": "%%routingKey%%"
}
```

Notifications via Webhook

To notify on events via webhook, in Data Infrastructure Insights navigate to **Alerts > Monitors** and click **+Monitor** to create a new [monitor](#).

- Select a metric and define the monitor's conditions.
- Under **_Set up team notification(s)**, choose the **Webhook** Delivery Method.
- Choose the "PagerDuty Trigger" webhook for Critical- and Warning-level events.
- Choose the "PagerDuty Resolve" for resolved events.

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

By Webhook

Notify team on

Critical, Warning

Use Webhook(s)

PagerDuty Trigger x

Notify team on

Resolved

Use Webhook(s)

PagerDuty Resolve x



Setting separate notifications for trigger events versus resolved events is a best practice, since PagerDuty handles trigger events differently than resolved events.

Webhook Example for Slack

Webhooks allow users to send alert notifications to various applications using a customized webhook channel. This page provides an example for setting up webhooks for Slack.



This page refers to third-party instructions, which could be subject to change. Refer to the [Slack documentation](#) for the most up-to-date information.

Slack Example:

- Go to <https://api.slack.com/apps> and Create a new App. Give it a meaningful name and select the Slack Workspace.

Create a Slack App

App Name

e.g. Super Service

Don't worry; you'll be able to change this later.

Development Slack Workspace

Development Slack Workspace

Your app belongs to this workspace—leaving this workspace will remove your ability to manage this app. Unfortunately, this can't be changed later.

By creating a Web API Application, you agree to the [Slack API Terms of Service](#).

Cancel>Create App

- Go to Incoming Webhooks, click on *Activate Incoming Webhooks*, Request to *Add New Webhook*, and select the Channel on which to Post.
- Copy the Webhook URL. You will need to paste this into the Data Infrastructure Insights webhook configuration.

Create Data Infrastructure Insights Webhook:

1. In Data Infrastructure Insights, navigate to **Admin > Notifications** and select the **Webhooks** tab. Click **+Webhook** to create a new webhook.
2. Give the webhook a meaningful Name, such as "Slack Webhook".
3. In the *Template Type* drop-down, select **Slack**.

4. Paste the URL from above into the *URL* field.

Edit a Webhook

Name

Slack

Template Type

Slack

URL

https://hooks.slack.com/services/<token string>

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

Notifications via Webhook

To notify on events via webhook, in Data Infrastructure Insights navigate to **Alerts > Monitors** and click **+Monitor** to create a new [monitor](#).

- Select a metric and define the monitor's conditions.
- Under *_Set up team notification(s)*, choose the **Webhook** Delivery Method.
- Choose the "Slack" webhook for the desired events (Critical, Warning, Resolved)

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

By Webhook	Notify team on	Use Webhook(s)
	<div>Critical, Warning, Resolved ▼</div>	<div>Slack x ▼</div>

More information:

- To modify message format and layout, see <https://api.slack.com/messaging/composing>
- Error handling: https://api.slack.com/messaging/webhooks#handling_errors

Webhook Example for Microsoft Teams

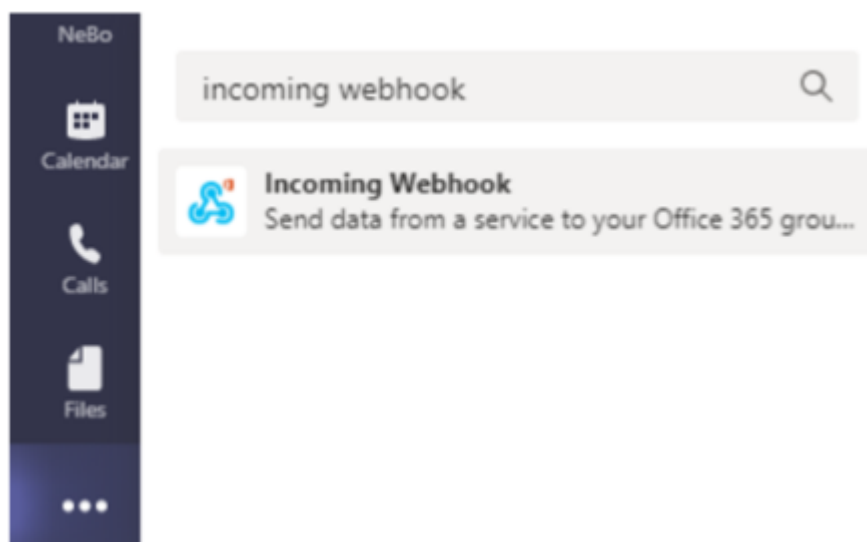
Webhooks allow users to send alert notifications to various applications using a customized webhook channel. This page provides an example for setting up webhooks for Teams.



This page refers to third-party instructions, which could be subject to change. Refer to the [Teams documentation](#) for the most up-to-date information.

Teams Setup:

1. In Teams, select the kebab, and search for Incoming Webhook.



2. Select **Add to a Team > Select a Team > Setup a Connector**.
3. Copy the Webhook URL. You will need to paste this into the Data Infrastructure Insights webhook configuration.

Create Data Infrastructure Insights Webhook:

1. In Data Infrastructure Insights, navigate to **Admin > Notifications** and select the **Webhooks** tab. Click **+Webhook** to create a new webhook.
2. Give the webhook a meaningful Name, such as "Teams Webhook".

3. In the *Template Type* drop-down, select **Teams**.

Edit a Webhook

Name

Teams Webhook

Template Type

Teams

URL

https://netapp.webhook.office.com/webhookb2/<token string>

Method

POST

Custom Header

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

Message Body

```
{
  "@type": "MessageCard",
  "@context": "http://schema.org/extensions",
  "themeColor": "0076D7",
  "summary": "Cloud Insights Alert",
  "sections": [
    {
      "activityTitle": "%%severity%% | %%alertId%% | %%triggeredOn%%",
      "activitySubtitle": "%%triggerTime%%",
      "markdown": false,
      "facts": [
```

Cancel

Test Webhook

Save Webhook

1. Paste the URL from above into the *URL* field.

Notifications via Webhook

To notify on events via webhook, in Data Infrastructure Insights navigate to **Alerts > Monitors** and click **+Monitor** to create a new [monitor](#).

- Select a metric and define the monitor's conditions.
- Under *_Set up team notification(s)*, choose the **Webhook** Delivery Method.
- Choose the "Teams" webhook for the desired events (Critical, Warning, Resolved)

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

By Webhook

Notify team on

Use Webhook(s)

Critical, Warning, Resolved

Teams - Edwin x

Working with Annotations

Defining annotations

When customizing Data Infrastructure 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

Data Infrastructure 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 Data Infrastructure 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, qtree, 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 Data Infrastructure 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 Data Infrastructure Insights.

Steps

1. In the Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure Insights menu, click **Observability > Enrich > 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 Data Infrastructure 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. Data Infrastructure Insights assigns the annotations to assets based on these rules. Data Infrastructure 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. Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure 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.

Note that when running multiple rules that update the same annotation on an asset, the first rule (as run from the top down) applies the annotation and updates the asset, then the second rule applies but doesn't change any annotation that was already set by the previous rule.

Deleting annotation rules

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

Steps

1. In the Data Infrastructure 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

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



The Data Infrastructure Insights API is available in **Data Infrastructure 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:

Parameters Cancel

No parameters

Request body multipart/form-data

CSV file to import

data
string(\$binary) Choose File No file chosen

Execute Clear

Responses

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 on your tenant, 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 on your tenant 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 on your tenant, you can define the applications in Data Infrastructure Insights.

About this task

Data Infrastructure Insights allows you to track data from assets associated with applications for things like usage or cost reporting.

Steps

1. In the Data Infrastructure Insights menu, click **Observability > Enrich > Applications**. Select

The Add Application dialog box displays.

2. Enter a unique name for the application.
3. Select a priority for the application.
4. Click **Save**.

After defining an application, it can be assigned to assets.

Assigning applications to assets

This procedure assigns the application to a host as an example. You can assign host, virtual machine, volume, or internal volumes to an application.

Steps

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.

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.



For environments with large amounts of related assets, inheritance of application assignments to those assets could take several minutes. Please allow more time for inheritance to occur if you have many related assets.

After you finish

After assigning the host to the application you can assign the remaining assets to the application. To access the landing page for the application, click **Manage > Application** and select the application you created.

Automatic Device Resolution

Automatic Device Resolution Overview

You need to identify all of the devices you want to monitor with Data Infrastructure Insights. Identification is necessary in order to accurately track performance and inventory on your tenant. Typically the majority of devices discovered on your tenant are identified through *Automatic Device Resolution*.

After you configure data collectors, devices on your tenant including switches, storage arrays, and your virtual infrastructure of hypervisors and VMs are identified. However, this does not normally identify 100% of the devices on your tenant.

After data collector type devices have been configured, best practice is to leverage device resolution rules to help identify the remaining unknown devices on your tenant. 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 Data Infrastructure 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 on your tenant?
- Have there been any acquisitions or mergers that introduced different naming schemes?

After analyzing your environment, you should be able to identify what naming standards exist that you can

expect to reliability encounter. The information you gathered might be represented graphically in a figure similar to the following:



In this example the largest number of devices are reliably represented by storage aliases. Rules that identify hosts using storage aliases should be written first, rules using switch aliases should be written next, and the last rules created should use zone aliases. Due to the overlap of the use of zone aliases and switch aliases, some storage alias rules might identify additional devices, leaving less rules required for zone aliases and switch aliases.

Steps to Identifying devices

Typically, you would use a workflow similar to the following to identify devices on your tenant. 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) on your tenant 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](#)

Device Resolution rules

You create device resolution rules to identify hosts, storage, and tapes that are not automatically identified currently by Data Infrastructure 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, Data Infrastructure Insights displays the following response:

- a. **Zones** lists the zones and WWN that need to be identified by Data Infrastructure Insights.
- b. **SRM** lists the unidentified aliases that need to be identified by Data Infrastructure Insights
- c. **Storage alias** lists storage aliases and WWN that need to be identified by Data Infrastructure Insights

d. **Switch alias** lists the switch aliases that need to be identified by Data Infrastructure 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 Data Infrastructure Insights should use to identify a host. Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure Insights web UI.
2. Click **Manage > Device Resolution**
3. Click the **Fibre Channel Identify** tab.

The system displays the devices along with their resolution status.

4. Select multiple unidentified devices.
5. Click **Bulk Actions** and select **Set host resolution** or **Set tape resolution**.

The system displays the Identify screen which contains a list of all of the rules that successfully identified devices.

6. Change the order of the rules to an order that meets your needs.

The order of the rules are changed in the Identify screen, but are not changed globally.

7. Select the method that that meets your needs.

Data Infrastructure Insights executes the host resolution process in the order in which the methods appear, beginning with those at the top.

When rules that apply are encountered, rule names are shown in the rules column and identified as manual.

Related:

[Fibre Channel Device Resolution](#)

[IP Device Resolution](#)

[Setting Device Resolution Preferences](#)

Fibre Channel device resolution

The Fibre Channel Identify screen displays the WWN and WWPN of fibre channel devices whose hosts have not been identified by automatic device resolution. The screen also displays any devices that have been resolved by manual device resolution.

Devices that have been resolved by manual resolution contain a status of *OK* and identify the rule used to

identify the device. Missing devices have a status of *Unidentified*. Devices that are specifically excluded from identification have a status of *Excluded*. The total coverage for identification of devices is listed on this page.

You perform bulk actions by selecting multiple devices on the left-hand side of the Fibre Channel Identify screen. Actions can be performed on a single device by hovering over a device and selecting the *Identify* or *Unidentify* buttons on the far right of the list.

The *Total Coverage* link displays a list of the number of devices identified/number of devices available for your configuration:

- SRM alias
- Storage alias
- Switch alias
- Zones
- User defined

Adding a Fibre Channel device manually

You can manually add a fibre channel device to Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure Insights device resolution feature. You might want to export a device identification so that you can modify it and then import it back into Data Infrastructure 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 Data Infrastructure 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:

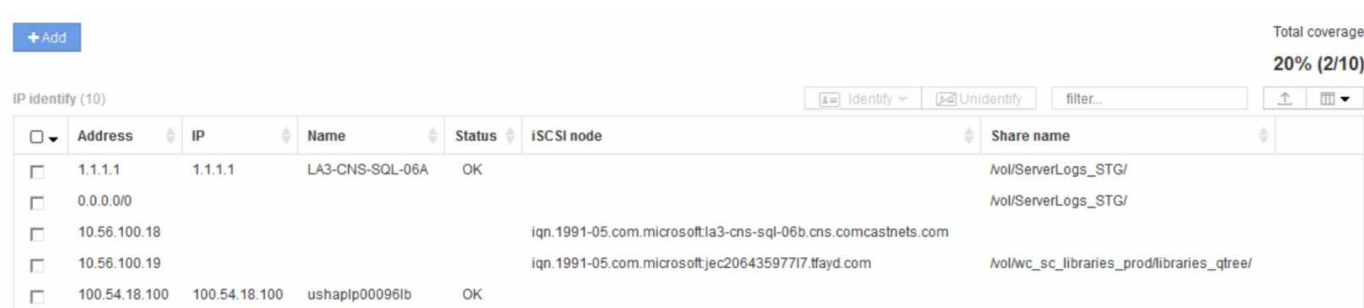
[IP Device Resolution](#)

[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.



The screenshot shows the 'IP identify (10)' screen. At the top left is a '+ Add' button. At the top right, it says 'Total coverage 20% (2/10)'. Below this is a toolbar with 'Identify' and 'Unidentify' buttons, a 'filter...' input, and icons for sorting and filtering. The main table has columns: Address, IP, Name, Status, iSCSI node, and Share name. There are 10 rows of data, with the first row highlighted.

<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 Data Infrastructure Insights using the manual add feature available in the IP Identify screen.

Procedure

1. Log in to the Data Infrastructure 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 Data Infrastructure Insights device resolution feature. You might want to export a device identification so that you can modify it and then import it back into Data Infrastructure 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 Data Infrastructure 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 Data Infrastructure Insights automatic resolution methods.

Formatting regular expressions

When creating regular expressions for Data Infrastructure Insights automatic resolution, you can configure output format by entering values in a field named *FORMAT*.

The default setting is \1, which means that a zone name that matches the regular expression is replaced by the contents of the first variable created by the regular expression. In a regular expression, variable values are created by parenthetical statements. If multiple parenthetical statements occur, the variables are referenced numerically, from left to right. The variables can be used in the output format in any order. Constant text can also be inserted in the output, by adding it to the *FORMAT* field.

For example, you might have the following zone names for this zone naming convention:

```
[Zone number]_[data center]_[hostname]_[device type]_[interface number]
```

- S123_Miami_hostname1_filer_FC1
- S14_Tampa_hostname2_switch_FC4
- S3991_Boston_hostname3_windows2K_FC0
- S44_Raleigh_hostname4_solaris_FC1

And you might want the output to be in the following format:

```
[hostname]-[data center]-[device type]
```

To do this, you need to capture the host name, data center, and device type fields in variables, and use them in the output. The following regular expression would do this:

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

Because there are three sets of parentheses, the variables \1, \2 and \3 would be populated.

You could then use the following format to receive output in your preferred format:

```
\2-\1-\3
```

Your output would be as follows:

```
hostname1-Miami-filer
hostname2-Tampa-switch
hostname3-Boston-windows2K
hostname4-Raleigh-solaris
```

The hyphens between the variables provide an example of constant text that is inserted in the formatted output.

Examples

Example 1 showing zone names

In this example, you use the regular expression to extract a host name from the zone name. You could create a regular expression if you have something similar to the following zone names:

- S0032_myComputer1Name-HBA0
- S0434_myComputer1Name-HBA1
- S0432_myComputer1Name-HBA3

The regular expression that you could use to capture the host name would be:

```
S[0-9]+_([a-zA-Z0-9]*)[_-]HBA[0-9]
```

The outcome is a match of all zones beginning with S that are followed by any combination of digits , followed by an underscore, the alphanumeric hostname (myComputer1Name), an underscore or hyphen, the capital letters HBA, and a single digit (0-9). The hostname alone is stored in the `\1` variable.

The regular expression can be broken into its components:

- "S" represents the zone name and begins the expression. This matches only an "S" at the beginning of the zone name.
- The characters [0-9] in brackets indicate that what follows "S" must be a digit between 0 and 9, inclusive.
- The + sign indicates that the occurrence of the information in the preceding brackets has to exist 1 or more times.
- The _ (underscore) means that the digits after S must be followed immediately by only an underscore character in the zone name. In this example, the zone naming convention uses the underscore to separate the zone name from the host name.
- After the required underscore, the parentheses indicate that the pattern contained within will be stored in the `\1` variable.
- The bracketed characters [a-zA-Z0-9] indicate that the characters being matched are all letters (regardless of case) and numbers.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The bracketed characters [_-] (underscore and dash) indicate that the alphanumeric pattern must be followed by an underscore or a dash.
- The letters HBA in the regular expression indicate that this exact sequence of characters must occur in the zone name.
- The final set of bracketed characters [0-9] match a single digit from 0 through 9, inclusive.

Example 2

In this example, skip up to the first underscore "", *then match E and everything after that up to the second ""*, and then skip everything after that.

Zone: Z_E2FHDBS01_E1NETAPP

Hostname: E2FHDBS01

RegExp: .?(E.?).*?

Example 3

The parentheses "()" around the last section in the Regular Expression (below) identifies which part is the hostname. If you wanted VSAN3 to be the host name, it would be: `_[a-zA-Z0-9].*`

Zone: A_VSAN3_SR48KENT_A_CX2578_SPA0

Hostname: SR48KENT

RegExp: `_[a-zA-Z0-9]+_([a-zA-Z0-9]).*`

Example 4 showing a more complicated naming pattern

You could create a regular expression if you have something similar to the following zone names:

- myComputerName123-HBA1_Symm1_FA3
- myComputerName123-HBA2_Symm1_FA5
- myComputerName123-HBA3_Symm1_FA7

The regular expression that you could use to capture these would be:

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

The `\1` variable would contain only *myComputerName123* after being evaluated by this expression.

The regular expression can be broken into its components:

- The parentheses indicate that the pattern contained within will be stored in the `\1` variable.
- The bracketed characters `[a-zA-Z0-9]` mean that any letter (regardless of case) or digit will match.
- The `*` (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The `_` (underscore) character in the regular expression means that the zone name must have an underscore immediately following the alphanumeric string matched by the preceding brackets.
- The `.` (period) matches any character (a wildcard).
- The `*` (asterisk) indicates that the preceding period wildcard may occur 0 or more times.

In other words, the combination `.*` indicates any character, any number of times.

Example 5 showing zone names without a pattern

You could create a regular expression if you have something similar to the following zone names:

- myComputerName_HBA1_Symm1_FA1
- myComputerName123_HBA1_Symm1_FA1

The regular expression that you could use to capture these would be:

```
(.*? )_.*
```

The \1 variable would contain *myComputerName* (in the first zone name example) or *myComputerName123* (in the second zone name example). This regular expression would thus match everything prior to the first underscore.

The regular expression can be broken into its components:

- The parentheses indicate that the pattern contained within will be stored in the \1 variable.
- The .* (period asterisk) match any character, any number of times.
- The * (asterisk) following the brackets indicates that the bracketed characters occur 0 or more times.
- The ? character makes the match non-greedy. This forces it to stop matching at the first underscore, rather than the last.
- The characters _.* match the first underscore found and all characters that follow it.

Example 6 showing computer names with a pattern

You could create a regular expression if you have something similar to the following zone names:

- Storage1_Switch1_myComputerName123A_A1_FC1
- Storage2_Switch2_myComputerName123B_A2_FC2
- Storage3_Switch3_myComputerName123T_A3_FC3

The regular expression that you could use to capture these would be:

```
.*?_.*?_([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 Data Infrastructure Insights: \1-\2 ([^_])_()
()_([AB]).*Format in Data Infrastructure 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]

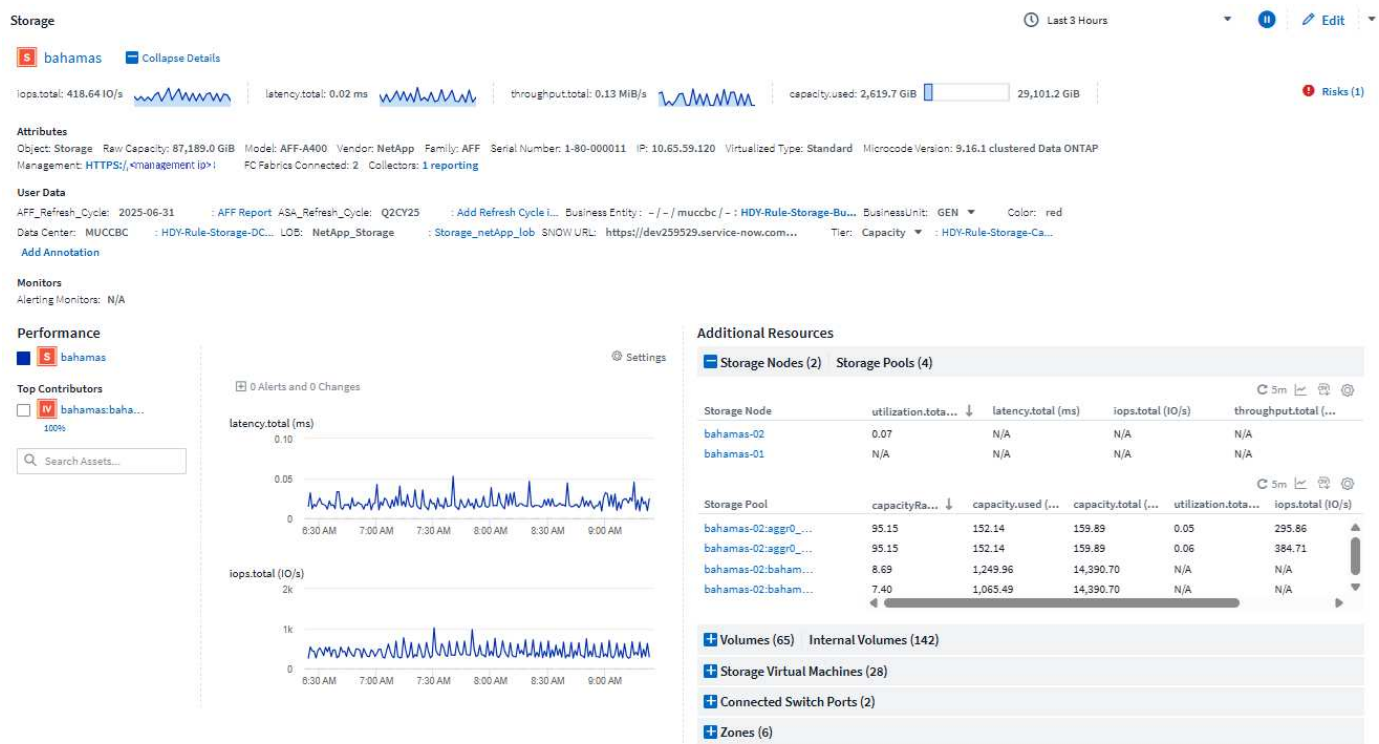
Format: US-BV-\\1

Asset Page Information

Asset Page Overview

Asset landing pages summarize the current status of an asset and contain links to additional information about the asset and its related assets.

Landing pages give you a single-page view of the object, with Summary, Performance, and Related Resource information.



Summary Section

At the top of a landing page is the expandible Summary section, which includes several sparkline graphs displaying recent data trends for things like throughput or latency, as well as object information and attributes, and any monitors that may be alerting for the object.

The Summary section also displays and enables you to change annotations assigned to the asset.

Performance Section

The Performance Section displays performance data for the object. Select *Settings* to add additional charts to the display like Throughput or Capacity, or select correlated or contributing resources to chart their data alongside the object's. Devices that may potentially be causing contention will also be listed in the Performance section. The data in the charts refreshes automatically as data collectors poll and updated data is acquired.

You can select the [metrics](#) you want to view in the performance chart for the time period selected. Click on the *Settings* drop-down and choose from the metrics listed.

In addition to performance data, any alerts that are or were active within the selected page time range will also be shown.

Performance

 **jamaica**
2 Alerts

Top Correlated


☐  **jamaica-02**
83%

Top Contributors


☐  **jamaica:DmoES...**
51%

☐  **DmoESX_jamaic...**
49%

Top Changes

☐  **svm_pdiddens_...**
3 Changes

Search Assets...



You can select from among the following devices that may be listed in the Performance section:

- 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.

- Top Changes

Assets related to recent changes.

- Workload Contentions

Shows the assets that impact or are impacted by other shared resources, such as hosts, networks, and storage. These are sometimes called *greedy* and *degraded* resources.

Additional Resources Section

The Additional Resources section displays tables of data for resources related to the current object type. You can expand and collapse these tables in order to focus on specific resources. Select the gear icon to temporarily show additional metrics or attributes in a table.

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.

Custom widgets are placed at the bottom of a landing page, below the Performance and Resource sections.

Types of Asset Pages

Data Infrastructure 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
- Host
- Zone

Changing the Time Range of Displayed Data

By default, an asset page displays the last 3 hours of data; however, you can change the time segment of displayed data by using an option that is located on every asset page, regardless of asset type. To change the time range, click the displayed time range in the top bar and choose from among the following time segments:

- Last 15 Minutes
- Last 30 Minutes
- Last 60 Minutes
- Last 2 Hours
- Last 3 Hours (this is the default)
- Last 6 Hours
- Last 12 Hours
- Last 24 Hours
- Last 2 Days
- Last 3 Days
- Last 7 Days
- Last 14 Days

- Last 30 Days
- Custom time range

The Custom time range allows you to select up to 31 consecutive days. You can also set the Start Time and End Time of day for this range. The default Start Time is 12:00 AM on the first day selected and the default End Time is 11:59 PM on the last day selected. Clicking Apply will apply the custom time range to the asset page.

The information on the page refreshes automatically based on the selected time range. The current refresh rate is displayed in the upper-right corner of the Summary section as well as on any relevant tables or widgets on the page.

Performance metric definitions

The Performance section can display 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 from 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.

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 on your tenant 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 on your tenant 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 on your tenant.

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

Storage Virtualization

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

Virtualization in a Table Widget

One of the easiest ways to begin looking at your storage virtualization is to create a dashboard table widget showing Virtualized type. When building the query for the widget, simply add "virtualizedType" to your grouping or filter.

Storage X ▼

Display

Last 3 Hours (Dashboard Time) ▼

☐ Override Dashboard Time

Filter by Attribute +

Filter by Metric +

Group by virtualizedType X ▼

The resulting table widget shows you the *Standard*, *Backend*, and *Virtual* storages on your tenant.

Storage by virtualizedType

50 items found in 4 groups

virtualizedType ↑	Storage
Backend (5)	--
Backend	Sym-Perf
Backend	Sym-000050074300343
Backend	CX600_26_CK00351029326
Backend	VNX8000_46_CK00351029346
Backend	Sym-000050074300324
Standard (36)	--
Virtual (8)	--

Landing Pages show Virtualized information

On a storage, volume, internal volume, or disk landing page, you can see relevant virtualization information. For example, looking at the storage landing page below, you can see that this is a Virtual storage, and which backend storage system applies. Any relevant tables on landing pages will also show virtualization information as applicable.

Storage Summary

Model:
V-Series

Vendor:
NetApp

Family:
V-Series

Serial Number:
1306894

IP:
192.168.7.41

Virtualized Type:
Virtual

Backend Storage:
[Sym-000050074300343](#)

Microcode Version:
8.0.2 7-Mode

Raw Capacity:
0.0 GiB

Latency - Total:
N/A

IOPS - Total:
N/A

Throughput - Total:
N/A

Management:


FC Fabrics Connected:
7

Alert Monitors:

Existing landing pages and dashboards

Be aware that if you currently have customized landing pages or dashboards on your tenant, these will not automatically show all virtualization information by default. However, you can *Revert to Default* any customized dashboard or landing page (you will have to re-implement your customizations), or modify the relevant widgets to include the desired virtualization attributes or metrics.

Revert to Default is available in the upper-right corner of a custom dashboard or landing page screen.



Hints and Tips to Search for Assets and Alerts

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

Category	Indexed terms
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
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

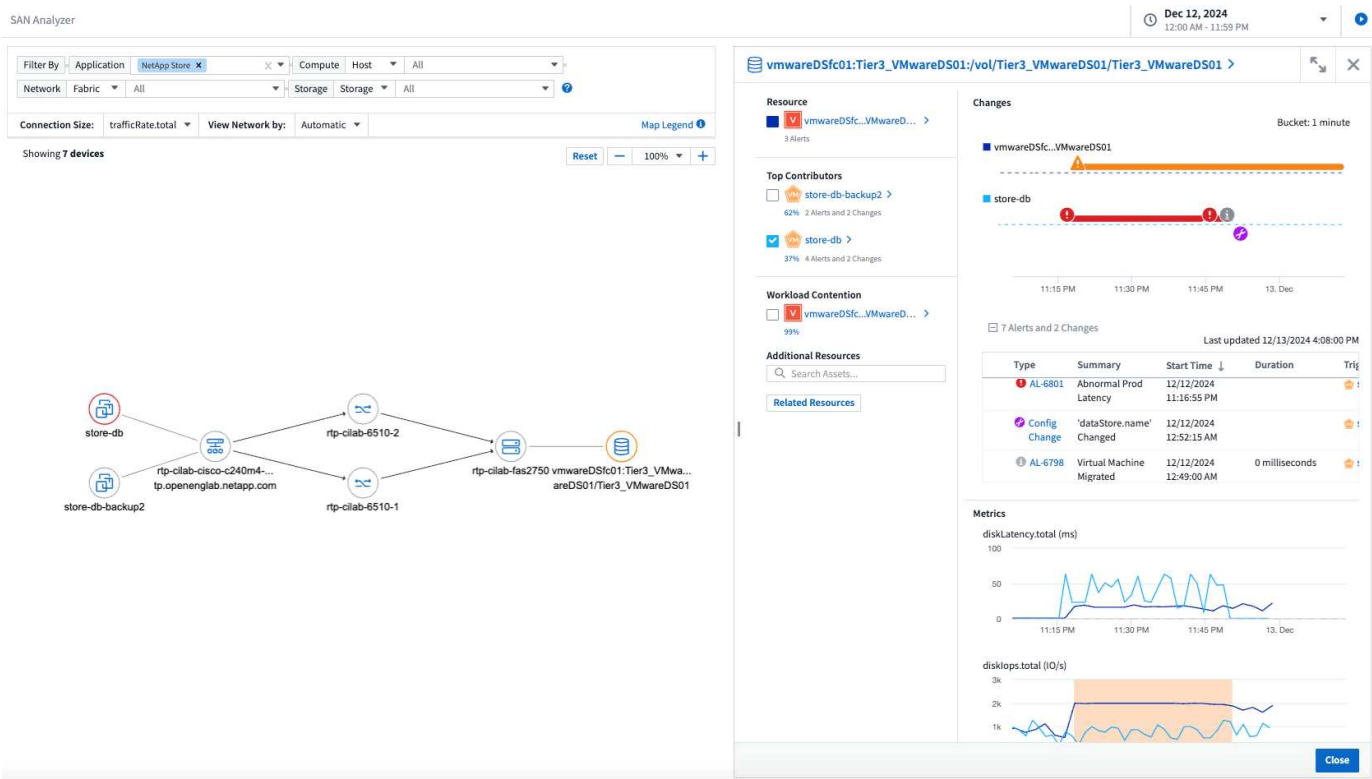
Category	Indexed terms
Virtual Machines	"virtualmachine" name DNS name IP addresses name of the host IP addresses of the host names of all datastores names of all associated applications
Switches (regular and NPV)	"switch" IP address wwn name serial number model domain ID name of the fabric wwn of the fabric
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

Analyzing Data

SAN Analyzer Overview

SAN plays a crucial role in handling vital workloads, but its complexity can result in significant outages and customer disruptions. With DII's **SAN Analyzer**, managing SAN becomes simpler and more efficient. This powerful tool offers end-to-end visibility, mapping dependencies from VM/Host to network to LUN and storage.

By providing an interactive topology map, SAN Analyzer enables you to pinpoint issues, understand changes, and enhance comprehension of data flow. Streamline SAN management in complex IT environments with SAN Analyzer and increase your visibility into block workloads.



Explore connections among your assets

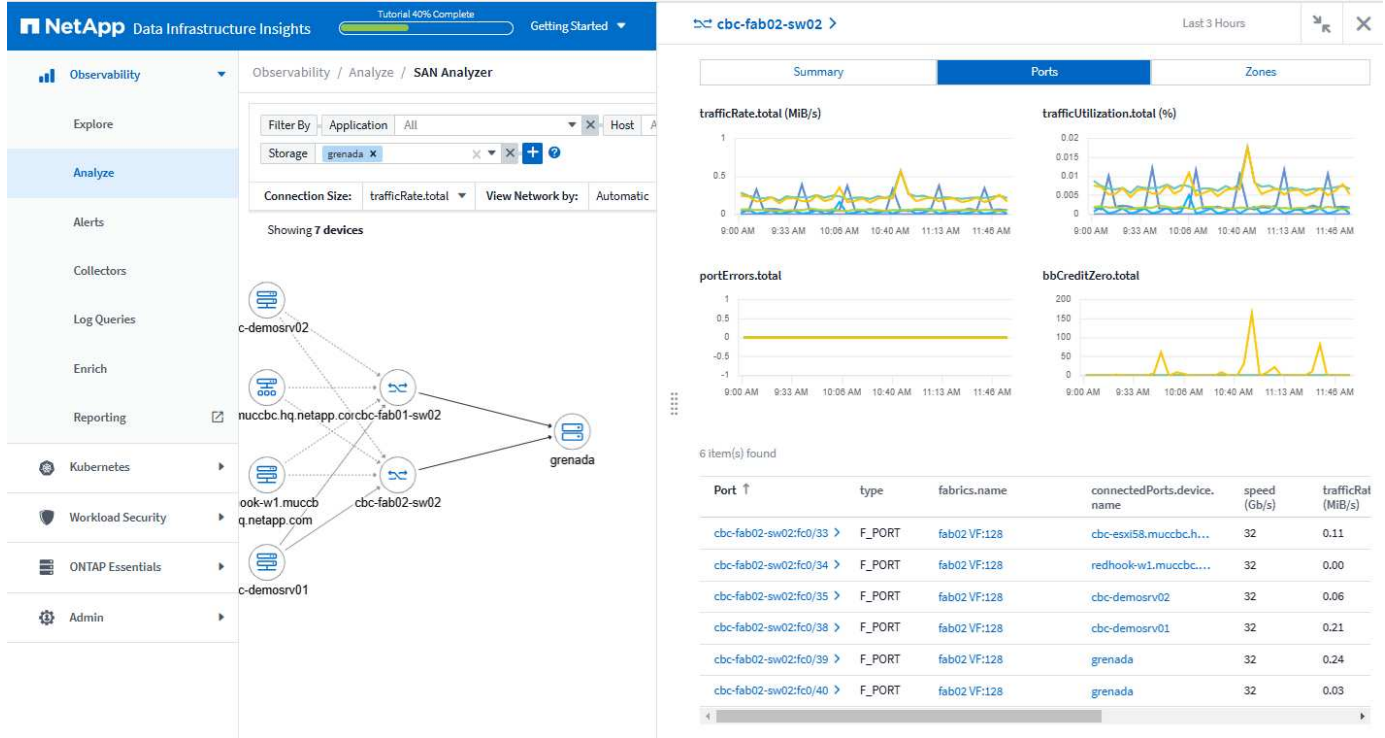
Select **Observability > Analyze > SAN Analyzer** to view the SAN Analyzer. Set a filter for Application, Host, Fabric, and/or Storage. The map for the objects is displayed, showing connected objects. Hover over an object to view traffic metrics for those connections.



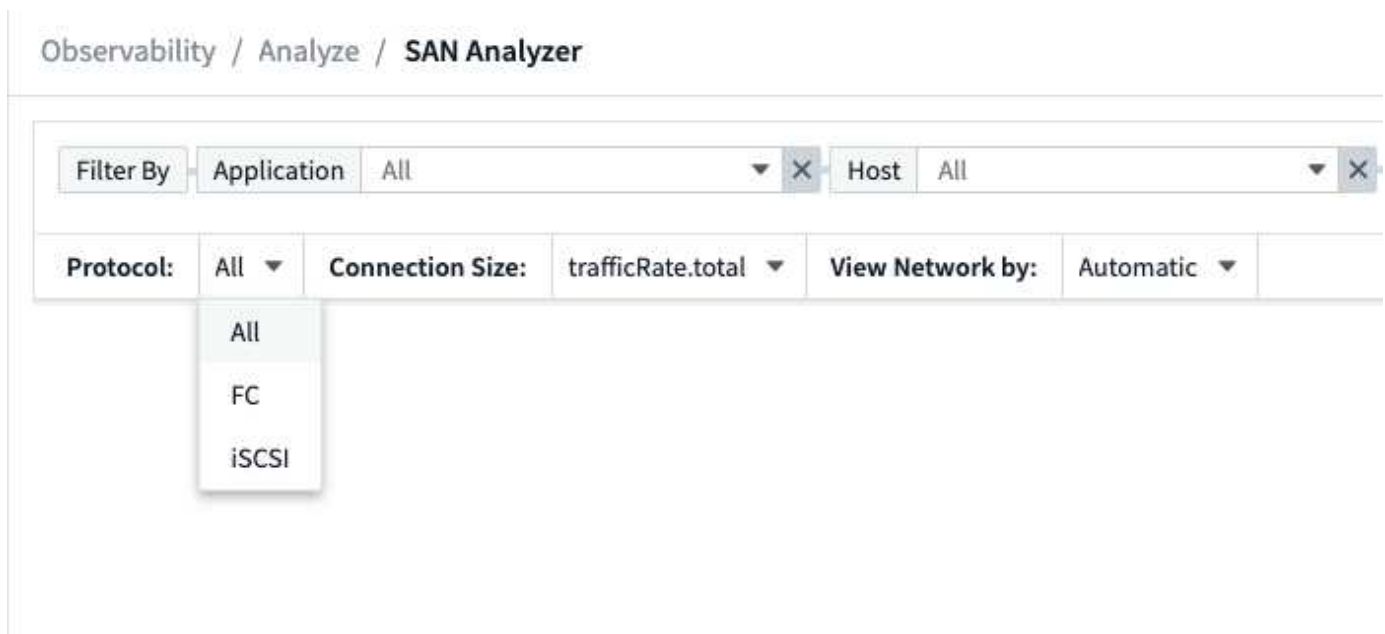
Most SAN Analyzer filters (including those you may add) are contextual; when you select an object in one of these filters, the choices presented in the other filter drop-downs are refreshed in context with the selected object or objects. The only exceptions to this are Application, Port, and Switch; these filters are not contextual.

Clicking on an object or group opens a slideout panel providing additional details about the object and its connections. The slideout panel displays a summary, which provides details about the selected object (for example, IP, Hypervisor, Connected Fabrics, etc., depending on the type of object), and charts showing metrics for the object such as latency or IOPS, and changes and alerts if relevant. You can select to display metrics for

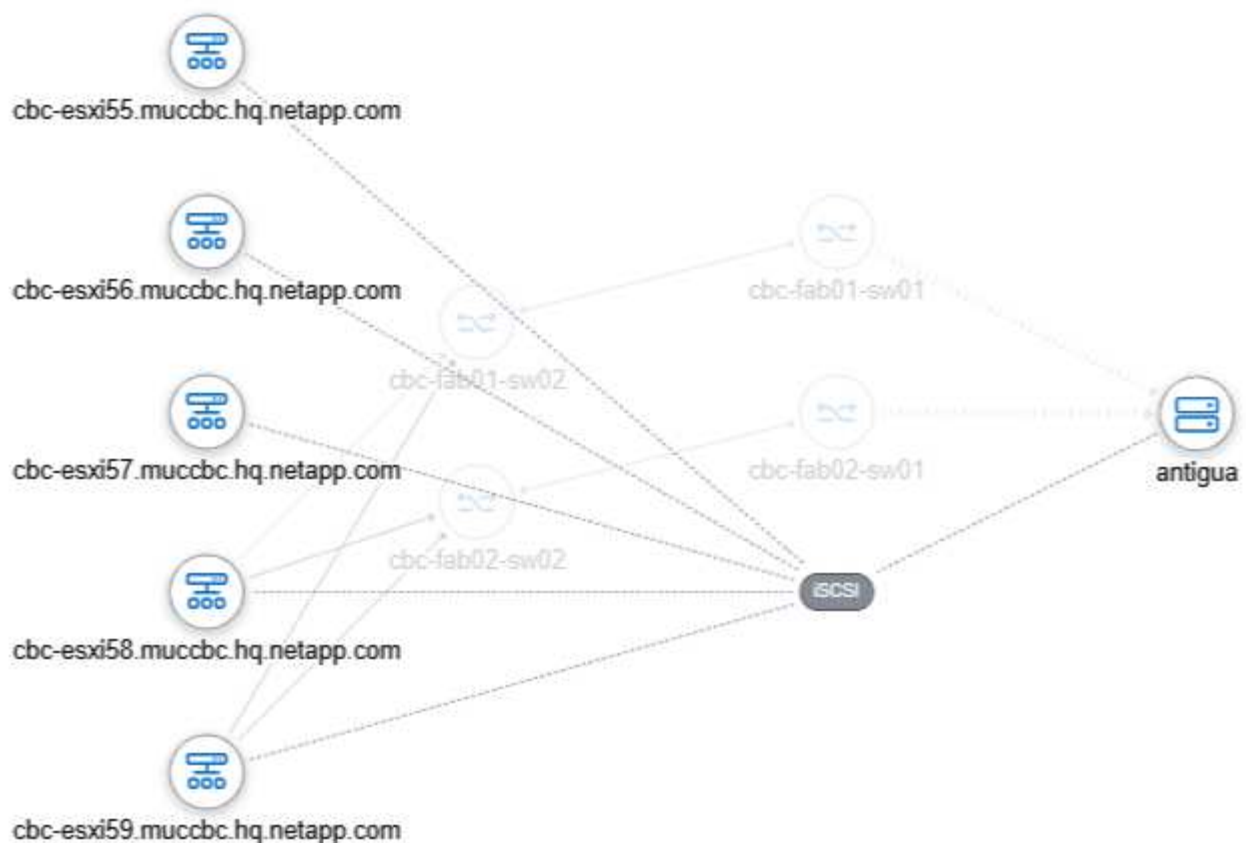
In addition to the Summary tab, the slideout panel displays tabs for things like Port details or Zone information, as applicable to the chosen object.



If your environment has different protocols, you can filter by iSCSI or FC:



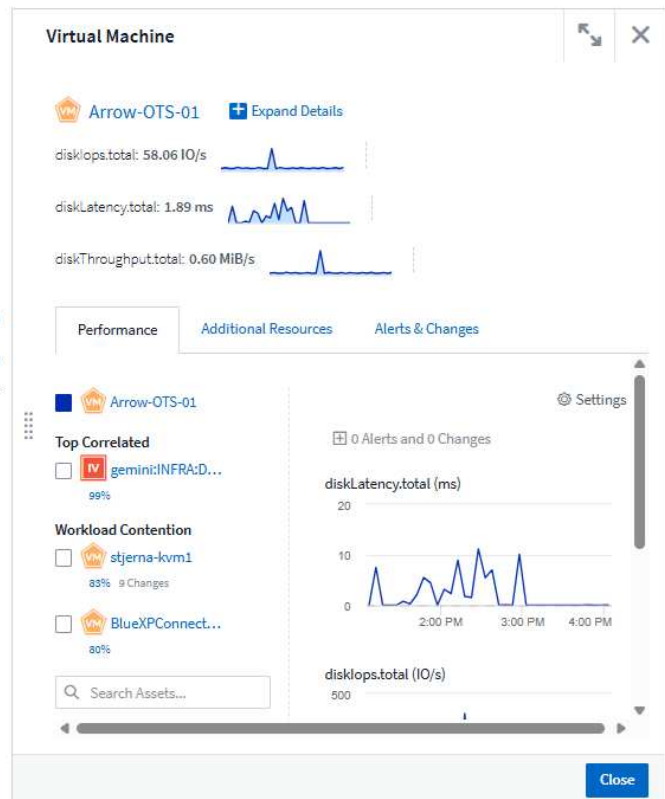
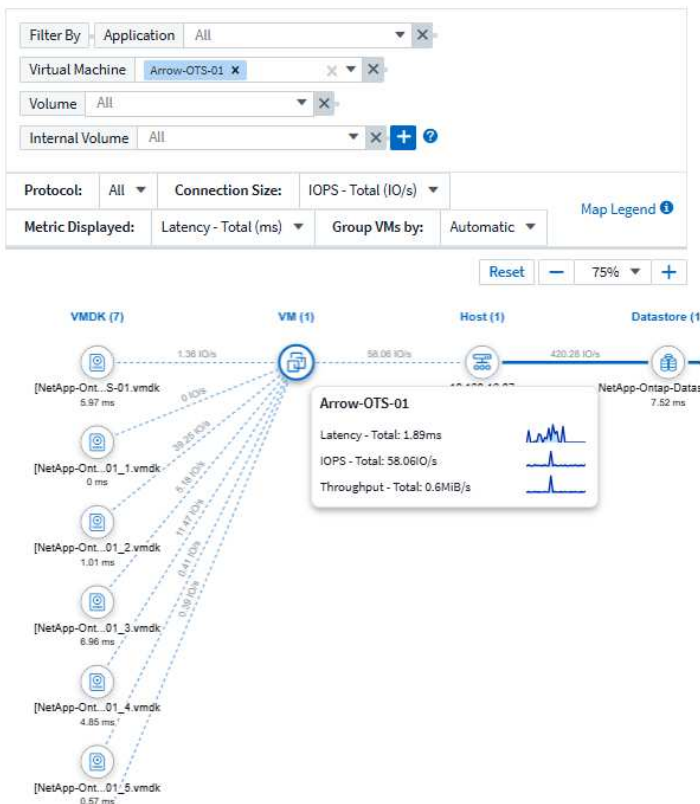
If your environment includes iSCSI devices, hovering over the *iSCSI* object highlights the connections related to those relevant iSCSI devices.



VM Analyzer Overview

With DII's **VM Analyzer**, managing your virtual assets becomes simpler and more efficient. This powerful tool offers end-to-end visibility, mapping dependencies from VMDK/VM to Host to Datastore to Internal Volume/Volume to Storage.

By providing an interactive topology map, VM Analyzer enables you to pinpoint issues, understand changes, and enhance comprehension of data flow. Streamline VM management and increase your visibility into virtual workloads.



Explore connections among your assets

Select **Observability > Analyze > VM Analyzer** to view the VM Analyzer. Set a filter for Application, Virtual Machine, Volume, Internal Volume, or add your own filters. The map for the objects is displayed, showing connected objects. Hover over an object to view traffic metrics for those connections.



Most VM Analyzer filters (including those you may add) are contextual; when you select an object in one of these filters, the choices presented in the other filter drop-downs are refreshed in context with the selected object or objects.

Clicking on an object or group opens a slideout panel providing additional details about the object and its connections. The slideout panel displays a summary, which provides details about the selected object (for example, throughput or utilization, depending on the type of object), and charts showing metrics for the object such as latency or IOPS. Additional tabs allow you to explore related additional resources or changes and alerts. You can select to display metrics for Top Correlated or contending objects on the charts as well, if desired.

See it in Action

[Simplified troubleshooting with VM Analyzer \(Video\)](#)

Monitor Infrastructure Health

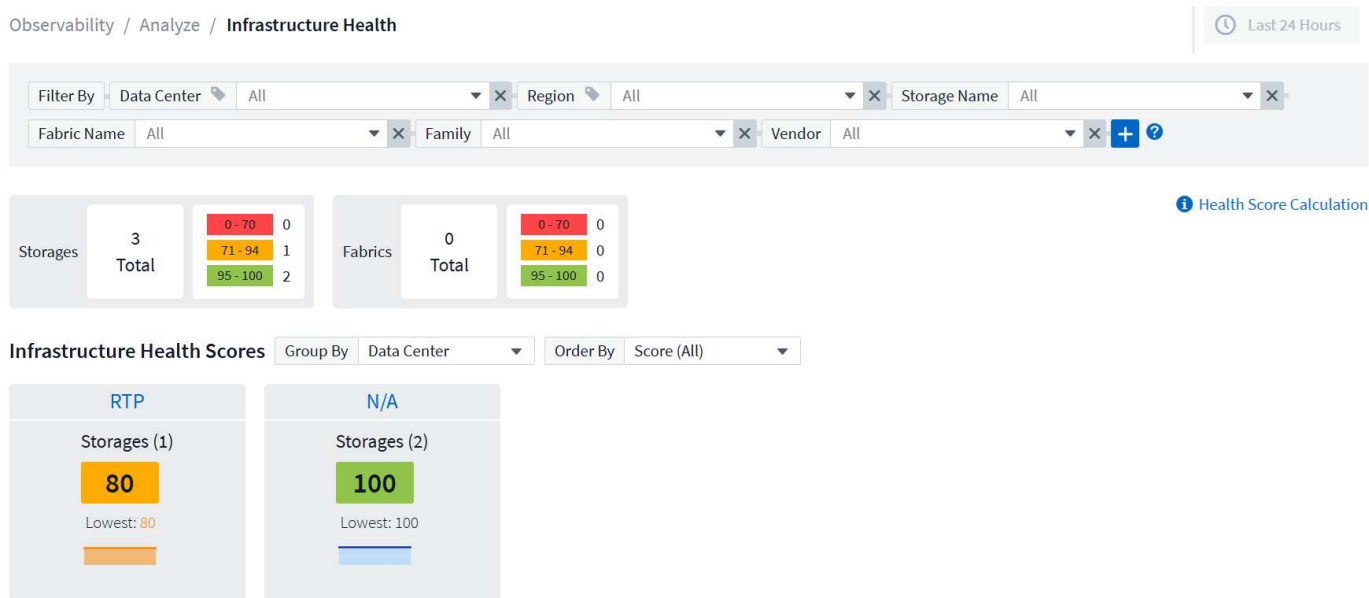
Data Infrastructure Insights provides comprehensive infrastructure health monitoring that tracks the performance, capacity, configuration, and component status of your storage environment. Health scores are calculated based on monitor alerts across these categories, giving you a unified view of system health and enabling proactive issue resolution.

The Infrastructure Health dashboard



Monitoring Infrastructure Health is a [Preview](#) feature and is subject to change.

Navigate to **Observability > Analyze** and select **Infrastructure Health**. The dashboard provides an overview of your system health, based on monitor alert categories and scores as explained below. Set filters at the top to narrow down the focus of your investigation.



By default, health scores are grouped by data center; you can select the grouping that works best for your session.

Configure Monitors to use for infrastructure health

Health scores are driven by alerts that are configured for inclusion in system health calculations.

When creating a monitor for an infrastructure object, you can choose whether to include alerts from the monitor in the calculations. At the bottom of the screen, expand the Advanced Configuration and select to *Include in Infrastructure Health Calculation*. Select a category to which to apply the calculation for the monitor:

- **Component Health** - fan failure, service processor offline, etc.
- **Performance Health** - high storage node utilization, abnormal spike in node latency, etc.

- **Capacity Health** - storage Pool capacity approaching full, insufficient space for LUN snapshot, etc.
- **Configuration Health** - cloud tier unreachable, SnapMirror relationship out of sync, etc.

☐ Advanced Configuration

Associate to an Infrastructure Health Category (optional)

☒ Include in Infrastructure Health Calculation

Select a Health Category

- Capacity
- Components
- Configuration
- Performance

Health scores explained

Scores are presented on a scale of 0 to 100, with 100 being at full health. Monitored infrastructure objects currently or recently experiencing issues will lower this score according to the following weighted averages:

- Components, Performance, or Capacity: 30% each
- Configuration: 10%

Health scores are impacted by alerts generated by the monitors you configured to include in infrastructure health calculations in the following ways:

- Critical alerts drop the health score by the *full* category weight
- Warning alerts drop the score by *half* the category weight.

If any categories are not reporting, the weighted average will adjust accordingly.

For example: 1 critical alert on Components (-30) and 1 warning alert on Performance (50% of 30 = -15) yield a health score of 55 (100 minus 45).

When alerts are resolved, these health score reductions gradually fade, and the score fully recovers within 2 hours.

Reporting

Data Infrastructure Insights Reporting Overview

Data Infrastructure 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 Data Infrastructure Insights [Premium Edition](#). Availability of the reporting feature is subject to a minimum footprint requirement. [Contact your NetApp sales representative](#) for more information.

With Data Infrastructure 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

Data Infrastructure 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 Data Infrastructure Insights Reporting

You can access Data Infrastructure Insights Reporting by clicking the **Reports** link in the menu.

You will be taken to the Reporting interface. Data Infrastructure Insights uses IBM Cognos Analytics for its reporting engine.

What is ETL?

When working with Reporting, you will hear the terms "Data Warehouse" and "ETL". ETL stands for "Extract, Transform, and Load". The ETL process retrieves data collected in Data Infrastructure Insights, and transforms the data into a format for use in Reporting. "Data Warehouse" refers to the collected data available for Reporting.

The ETL process includes these individual processes:

- **Extract:** Takes data from Data Infrastructure Insights.
- **Transform:** Applies business logic rules or functions to the data as it is extracted from Data Infrastructure Insights.
- **Load:** Saves the transformed data into the data warehouse for use in Reporting.

Data Infrastructure Insights Reporting User Roles

If you have Data Infrastructure Insights Premium Edition with Reporting, every Data Infrastructure Insights user on your tenant 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 Data Infrastructure Insights determines your Reporting user role:

Data Infrastructure 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
Add/Edit HTML Item	No	No	Yes
Run report with HTML Item	Yes	Yes	Yes
Add/Edit Custom SQL	No	No	Yes
Run reports with Custom SQL	Yes	Yes	Yes

Setting Reporting (Cognos) email preferences




If you change your user email preferences within Data Infrastructure 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.

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 on your tenant:

1. Navigate to **Menu > Content**



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 Data Infrastructure 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

Data Infrastructure 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 Data Infrastructure Insights [Premium Edition](#).

You can generate pre-defined reports from the Data Infrastructure 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.

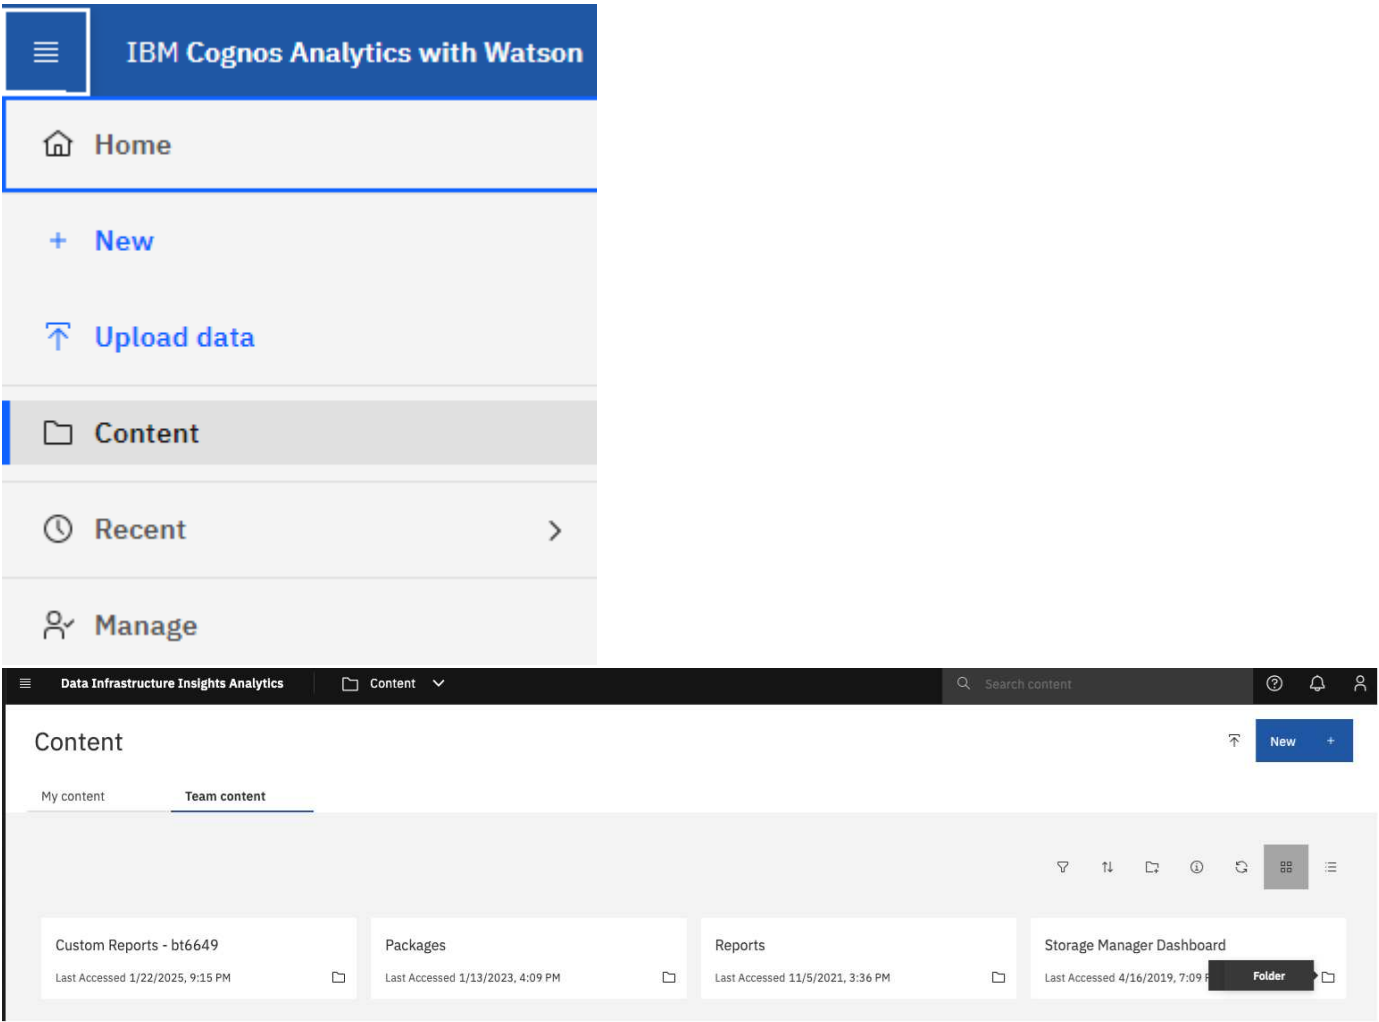
The pre-defined reports show your inventory, storage capacity, chargeback, performance, storage efficiency, and cloud cost data. You can modify these pre-defined reports and save your modifications.

You can generate reports in various formats, including HTML, PDF, CSV, XML, and Excel.

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 Data Infrastructure Insights reports.

- 1. In the left navigation pane, select **Content > Team Content**.
- 2. Select **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 Data Infrastructure 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

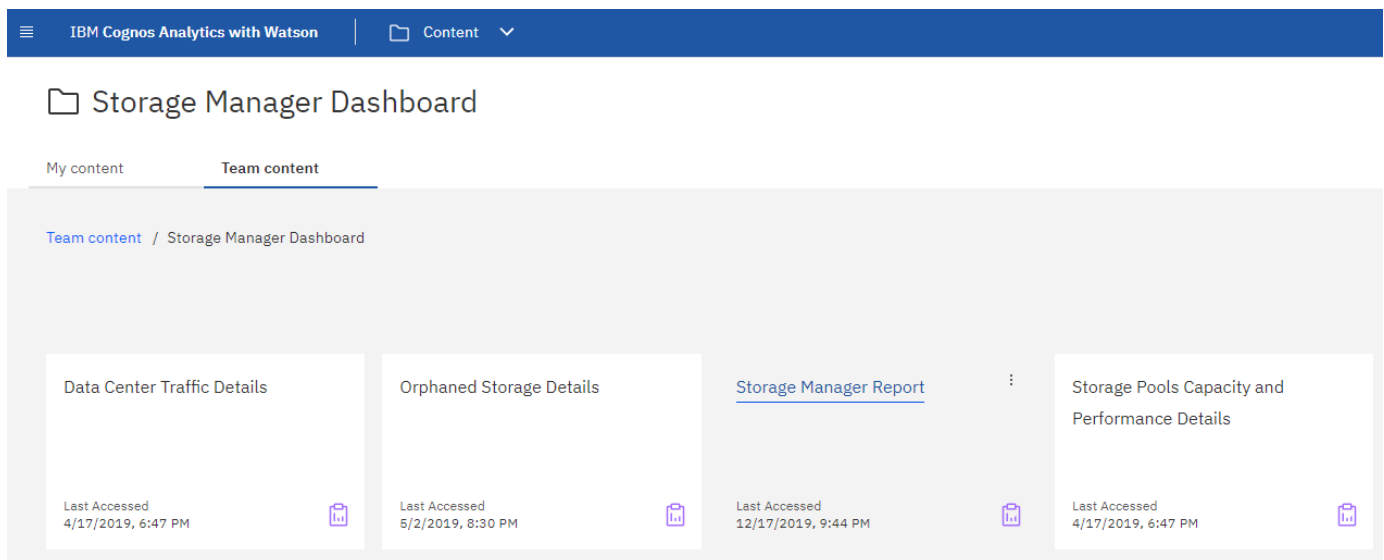
The Storage Manager Dashboard provides you with a centralized visualization that enables you to compare and contrast resource usage over time against the acceptable ranges and previous days of activity. Showing only the key performance metrics for your storage services, you can make decisions about how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



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

Summary

Selecting **Storage Manager Dashboard** from Team Content gives you several reports that provide information on of your traffic and storage.




For an at-a-glance view, the **Storage Manager Report** comprises seven components that contain contextual information on many aspects of your storage environment. You can drill down on the aspects of your storage services to perform an in-depth of analysis of a section that interests you most.

NetApp Storage Manager Dashboard

(Data as of Jan 28, 2016)

Summary

History (Target: Actual: Forecast: Low: Mid: High)


**Data Centers Time to Full**

(<3 months; 3-6 months; >6 months)

**Storage Tiers Capacity**


(Target: Actual: Forecast)

Last 12 Months Used Capacity Total Capacity Months to Full


**Daily Storage Traffic**

(Terabytes)


Daily mean for last 6 months Daily mean for last 7 days Yesterday

**Top 10 Applications**


Last 12 Months Used Allocated Response Time (Acceptable)

**Storage Tiers Daily Performance**

(Acceptable)

**Orphaned Capacity**

35TB 3.4%



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.

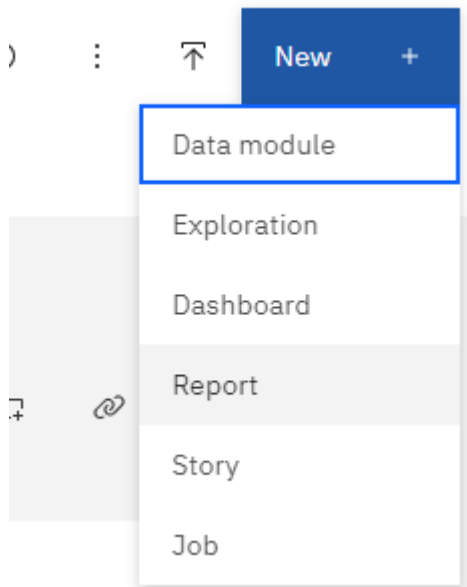
You can click a tier to display the Orphaned Storage Details report, which shows all the volumes identified as orphaned by configuration and performance for the selected tier. Click any storage, storage pool, or volume in this report to display the asset page summarizing the current state of that resource.

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. Navigate to **Menu > Content > Team Content > Reports**
2. In the upper-right of the screen, select **[New +]**
3. Select **Report**



4. On the **Templates** tab, select *Blank*
The Source and Data tabs is displayed
5. Open **Select a source +**
6. Under **Team content**, open **Packages**
A list of available packages is displayed.
7. Choose **Storage and Storage Pool Capacity**



8. Select **Open**

The available styles for your report are displayed.

9. Select **List**

Add appropriate names for List and Query

10. Select **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 selecting an output type from the **Run** menu.



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 Data Infrastructure Insights [Premium Edition](#).

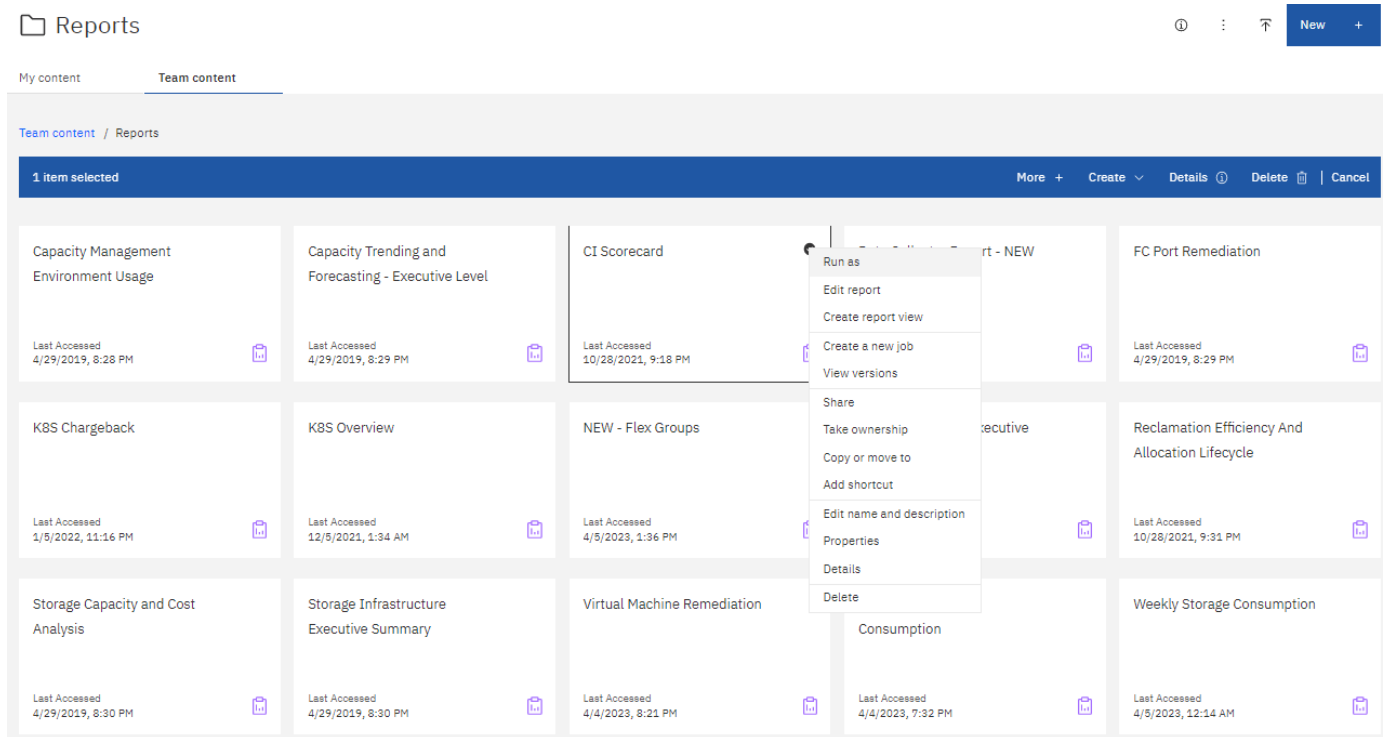


Before making changes to reporting permissions or security, you must copy "My Content" reports to the "Team Content" folder to ensure Reports are saved.

Customizing a report's output format and delivery

You can customize the format and delivery method of reports.

1. In the Data Infrastructure Insights Reporting Portal, Go to **Menu > Content > My Content/Team Content**. Mouse over the report you want to customize and open the "three dots" menu.

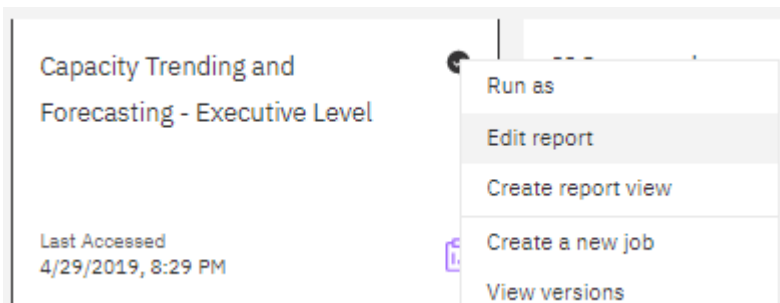


1. Click **Properties > Schedule**
2. You can set the following options:
 - **Schedule** when you want reports to run.
 - Choose **Options** for report format and delivery (Save, Print, Email) and Languages for the report.
3. Click **Save** to produce the report using the selections you made.


Copying a report to the clipboard

Use this process to copy a report to the clipboard.

1. Select a report to copy from (**Menu > Content > My Content or Team Content**)
2. Choose *Edit report* from the report's drop-down menu



3. In the upper-right of the screen, open the "three dots" menu next to "Properties".
4. Select **Copy Report to Clipboard**.




Opening reports from the clipboard

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

About this task

Start by creating a new report or opening an existing report that you wish to replace with the copied report. The steps below are for a new report.

1. Select **Menu > +New > Report** and create a blank report.
2. In the upper-right of the screen, open the "three dots" menu next to "Properties".
3. Select **Open Report from Clipboard**.



1. Paste the copied code into the window and select **OK**.

2. Select the floppy disk icon to save the report.
3. Choose where to save the report (*My Content*, *Team Content*, or create a new folder).
4. Give the new report a meaningful name and select **Save**.

Editing an existing report

Be aware that editing files in their default location runs a risk of those reports being overwritten upon the next report catalog refresh. It is recommended to save the edited report under a new name or store it in a non-default location.

Troubleshooting

Here you will find suggestions for troubleshooting problems with Reporting.

Problem:	Try this:
When scheduling a report to be sent via email, the name of the user logged in is pre-populated to the email's "To" field. However, the name is in the form of "firstname lastname" (first name, space, last name). Since this is not a valid email address, the email will fail to send when the scheduled report is run.	When scheduling the report to be sent via email, clear the pre-populated name and enter a valid, properly-formatted email address in the "To" field.
My scheduled report sends out via email, but the report can not be accessed if the origination is from the "My Content" folder.	In order to avoid this, the report or report-view must be saved to the "Team Content > Custom Reports - xxxxxx" folder, and the schedule created from that saved version. The "Custom Reports - xxxxxx" folder is visible to all users on the tenant.
when saving a Job, the folder may show "Team Content" with the list of content from "Custom Reports - xxxxxx", however you can not save the Job here because Cognos thinks this is the "Team Content" folder where you don't have access to write.	The work around is to create a new folder with a unique name (i.e. "NewFolder") and save there, or to save to "My Content" and then copy/move to "Custom Reports - xxxxxx".

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 Data Infrastructure Insights [Premium Edition](#).

The examples in this section show the following process, which can be used for any of the Data Infrastructure 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 on your tenant. It is beneficial to have annotations stable before generating reports, because Data Infrastructure 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 Data Infrastructure Insights.

- Data Infrastructure 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

Data Infrastructure Insights includes several data models from which you can either select predefined reports or create your own custom report.

Each data model contains a simple data mart and an advanced data mart:

- The simple data mart provides quick access to the most commonly used data elements and includes only the last snapshot of Data Warehouse data; it does not include historical data.
- The advanced data mart provides all values and details available from the simple data mart and includes access to historical data values.

Capacity data models

Enables you to answer questions about storage capacity, file system utilization, internal volume capacity, port capacity, qtree capacity, and virtual machine (VM) capacity. The Capacity data model is a container for several capacity data models. You can create reports answering various types of questions using this data model:

Storage and Storage Pool Capacity data model

Enables you to answer questions about storage capacity resource planning, including storage and storage pools, and includes both physical and virtual storage pool data. This simple data model can help you answer questions related to capacity on the floor and the capacity usage of storage pools by tier and data center over time.

If you are new to capacity reporting, you should start with this data model because it is a simpler, targeted data model. You can answer questions similar to the following using this data model:

- What is the projected date for reaching the capacity threshold of 80% of my physical storage?
- What is the physical storage capacity on an array for a given tier?
- What is my storage capacity by manufacturer and family as well as by data center?
- What is the storage utilization trend on an array for all of the tiers?
- What are my top 10 storage systems with the highest utilization?
- What is the storage utilization trend of the storage pools?
- How much capacity is already allocated?
- What capacity is available for allocation?

File System Utilization data model

This data model provides visibility about capacity utilization by hosts at the file system level. Administrators can determine allocated and used capacity per file system, determine the type of file system, and identify trending statistics by file system type. You can answer the following questions using this data model:

- What is the size of the file system?
- Where is the data kept and how is it accessed, for example, local or SAN?
- What are the historical trends for the file system capacity? Then, based on this, what can we anticipate for future needs?

Internal Volume Capacity data model

Enables you to answer questions about internal volume used capacity, allocated capacity, and capacity usage over time:

- Which internal volumes have a utilization higher than a predefined threshold?
- Which internal volumes are in danger of running out of capacity based on a trend?
- 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 Data Infrastructure Insights, which shows the provisioned size for the VM itself.

Volume Capacity data model

Enables you to analyze all aspects of the volumes on your tenant 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. Many of these report *Hourly* data, *Daily* data, or both. Using this data model, you can create reports that answer several types of performance management questions:

- What volumes or internal volumes have not been used or accessed during a specific period?
- Can we pinpoint any potential misconfiguration for storage for an application (unused)?
- What was the overall access behavior pattern for an application?
- Are tiered volumes assigned appropriately for a given application?
- Could we use cheaper storage for an application currently running without impact to application performance?

- What are the applications that are producing more accesses to currently configured storage?

When you use the switch performance tables, you can obtain the following information:

- Is my host traffic through connected ports balanced?
- Which switches or ports are exhibiting a high number of errors?
- What are the most used switches based on port performance?
- What are the underutilized switches based on port performance?
- What is the host trending throughput based on port performance?
- What is the performance utilization for last X days for one specified host, storage system, tape, or switch?
- Which devices are producing traffic on a specific switch (for example, which devices are responsible for use of a highly utilized switch)?
- What is the throughput for a specific business unit in our environment?

When you use the disk performance tables, you can obtain the following information:

- What is the throughput for a specified storage pool based on disk performance data?
- What is the highest used storage pool?
- What is the average disk utilization for a specific storage?
- What is the trend of usage for a storage system or storage pool based on disk performance data?
- What is the disk usage trending for a specific storage pool?

When you use VM and VMDK performance tables, you can obtain the following information:

- Is my virtual environment performing optimally?
- Which VMDKs are reporting the highest workloads?
- How can I use the performance reported from VMDs mapped to different datastores to make decisions about re-tiering.

The Performance data model includes information that helps you determine the appropriateness of tiers, storage misconfigurations for applications, and last access times of volumes and internal volumes. This data model provides data such as response times, IOPs, throughput, number of writes pending, and accessed status.

Storage Efficiency data model

Enables you to track the storage efficiency score and potential over time. This data model stores measurements of not only the provisioned capacity, but also the amount that is used or consumed (the physical measurement). For example, when thin provisioning is enabled, Data Infrastructure 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.

Access the Reporting Database via API

Data Infrastructure Insights' powerful API allows users to query the Data Infrastructure Insights Reporting database directly, without going through the Cognos Reporting environment.



This documentation refers to the Data Infrastructure Insights Reporting feature, which is available in Data Infrastructure Insights Premium Edition.

Odata

The Data Infrastructure Insights Reporting API follows the [OData v4](#) (Open Data Protocol) standard for its querying of the Reporting database.

For more information or to learn more, check out [this tutorial](#) on OData.

All requests will start with the url *https://<Data Infrastructure Insights URL>/rest/v1/dwh-management/odata*

Generating an APIKey

Read more about [Data Infrastructure Insights APIs](#).

To generate an API key, do the following:

- Log into your Data Infrastructure Insights environment and select **Admin > API Access**.
- Click “+ API Access Token”.
- Enter a Name & Description.
- For type, choose *Data Warehouse*.
- Set Permissions as Read/Write.
- Set a desired Expiration date.
- Click “Save”, then **copy the key and save it** somewhere safe. You will not be able to access the full key later.

APIkeys are good for [Sync or Async](#).

Direct query of tables

With the API Key in place, direct queries of the Reporting database are now possible. Long URLs may be simplified to `https://.../odata/` for display purposes rather than the full `https://<Data Infrastructure Insights URL>/rest/v1/dwh-management/odata/`

Try simple queries like

- `https://<Data Infrastructure Insights URL>/rest/v1/dwh-management/odata/dwh_custom`
- `https://<Data Infrastructure Insights URL>/rest/v1/dwh-management/odata/dwh_inventory`
- `https://<Data Infrastructure Insights URL>/rest/v1/dwh-management/odata/dwh_inventory/storage`
- `https://<Data Infrastructure Insights URL>/rest/v1/dwh-management/odata/dwh_inventory/disk`
- `https://.../odata/dwh_custom/custom_queries`

REST API Examples

The URL for all calls is `https://<Data Infrastructure Insights URL>/rest/v1/dwh-management/odata`.

- GET `/[{schema}]/**` - Retrieves data from the Reporting Database.

Format: `https://<Data Infrastructure Insights URL>/rest/v1/dwh-management/odata/<schema_name>/<query>`

Example:

```
https://<domain>/rest/v1/dwh-  
management/odata/dwh_inventory/fabric?$count=true&$orderby=name
```

Result:

```
{
  "@odata.context": "$metadata#fabric",
  "@odata.count": 2,
  "value": [
    {
      "id": 851,
      "identifier": "10:00:50:EB:1A:40:3B:44",
      "wwn": "10:00:50:EB:1A:40:3B:44",
      "name": "10:00:50:EB:1A:40:3B:44",
      "vsanEnabled": "0",
      "vsanId": null,
      "zoningEnabled": "0",
      "url": "https://<domain>/web/#/assets/fabrics/941716"
    },
    {
      "id": 852,
      "identifier": "10:00:50:EB:1A:40:44:0C",
      "wwn": "10:00:50:EB:1A:40:44:0C",
      "name": "10:00:50:EB:1A:40:44:0C",
      "vsanEnabled": "0",
      "vsanId": null,
      "zoningEnabled": "0",
      "url": "https://<domain>/web/#/assets/fabrics/941836"
    }
  ]
}
```

Helpful Hints

Keep the following in mind when working with Reporting API queries.

- The query payload must be a valid JSON string
- The query payload must be contained in a single line
- Double quotes must be escaped, i.e. \"
- Tabs are supported as \t
- Avoid comments
- Lower-case table names are supported

Additionally:

- 2 Headers are required:
 - Name "X-CloudInsights-ApiKey"
 - Attribute Value "<apikey>"

Your API key will be specific to your Data Infrastructure Insights environment.

Synchronous or Asynchronous?

By default, an API command will operate in *synchronous* mode, meaning that you send the request and the response is returned immediately. However, at times a query may take a long time to execute, which could lead to the request timing out. To get around this, you can execute a request *asynchronously*. In asynchronous mode, the request will return a URL through which the execution can be monitored. The URL will return the result when it is ready.

To execute a query in async mode, add the header **Prefer: respond-async** to the request. Upon successful execution, the response will contain the following headers:

```
Status Code: 202 (which means ACCEPTED)
preference-applied: respond-async
location: https://<Data Infrastructure Insights URL>/rest/v1/dwh-
management/odata/dwh_custom/asyncStatus/<token>
```

Querying the location URL will return the same headers if the response is not ready yet, or will return with status 200 if the response is ready. The response content will be of type text and contains the http status of the original query and some metadata, followed by the results of the original query.

```
HTTP/1.1 200 OK
OData-Version: 4.0
Content-Type: application/json;odata.metadata=minimal
oDataResponseSizeCounted: true

{ <JSON_RESPONSE> }
```

To see a list of all async queries and which of them are ready, use the following command:

```
GET https://<Data Infrastructure Insights URL>/rest/v1/dwh-
management/odata/dwh_custom/asyncList
```

The response has the following format:

```

{
  "queries" : [
    {
      "Query": "https://<Data Infrastructure Insights
URL>/rest/v1/dwh-
management/odata/dwh_custom/heavy_left_join3?$count=true",
      "Location": "https://<Data Infrastructure Insights
URL>/rest/v1/dwh-management/odata/dwh_custom/asyncStatus/<token>",
      "Finished": false
    }
  ]
}

```

How historical data is retained for Reporting

Data Infrastructure Insights retains historical data for use in Reporting based on the data marts and granularity of the data, as shown in the following table.

Data mart	Measured object	Granularity	Retention period
Performance marts	Volumes and internal volumes	Hourly	14 days
Performance marts	Volumes and internal volumes	Daily	13 months
Performance marts	Application	Hourly	13 months
Performance marts	Host	Hourly	13 months
Performance marts	Switch performance for port	Hourly	35 days
Performance marts	Switch performance for host, storage, and tape	Hourly	13 months
Performance marts	Storage node	Hourly	14 days
Performance marts	Storage node	Daily	13 months
Performance marts	VM performance	Hourly	14 days
Performance marts	VM performance	Daily	13 months
Performance marts	Hypervisor performance	Hourly	35 days
Performance marts	Hypervisor performance	Daily	13 months
Performance marts	VMDK performance	Hourly	35 days
Performance marts	VMDK performance	Daily	13 months
Performance marts	Disk performance	Hourly	14 days
Performance marts	Disk performance	Daily	13 months

Capacity marts	All (except individual volumes)	Daily	13 months
Capacity marts	All (except individual volumes)	Monthly representative	14 months and beyond
Inventory marts	Individual volumes	Current state	1 day (or until next ETL)

Data Infrastructure Insights Reporting Schema Diagrams

This document provides schema diagrams for the Reporting Database.

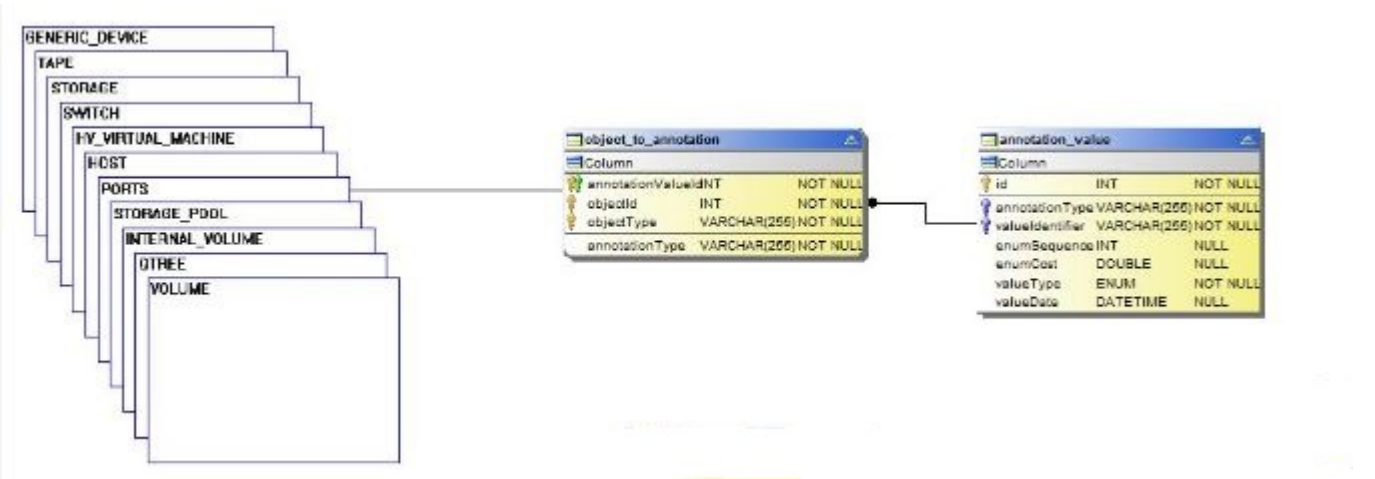


The Reporting feature is available in Data Infrastructure Insights [Premium Edition](#).

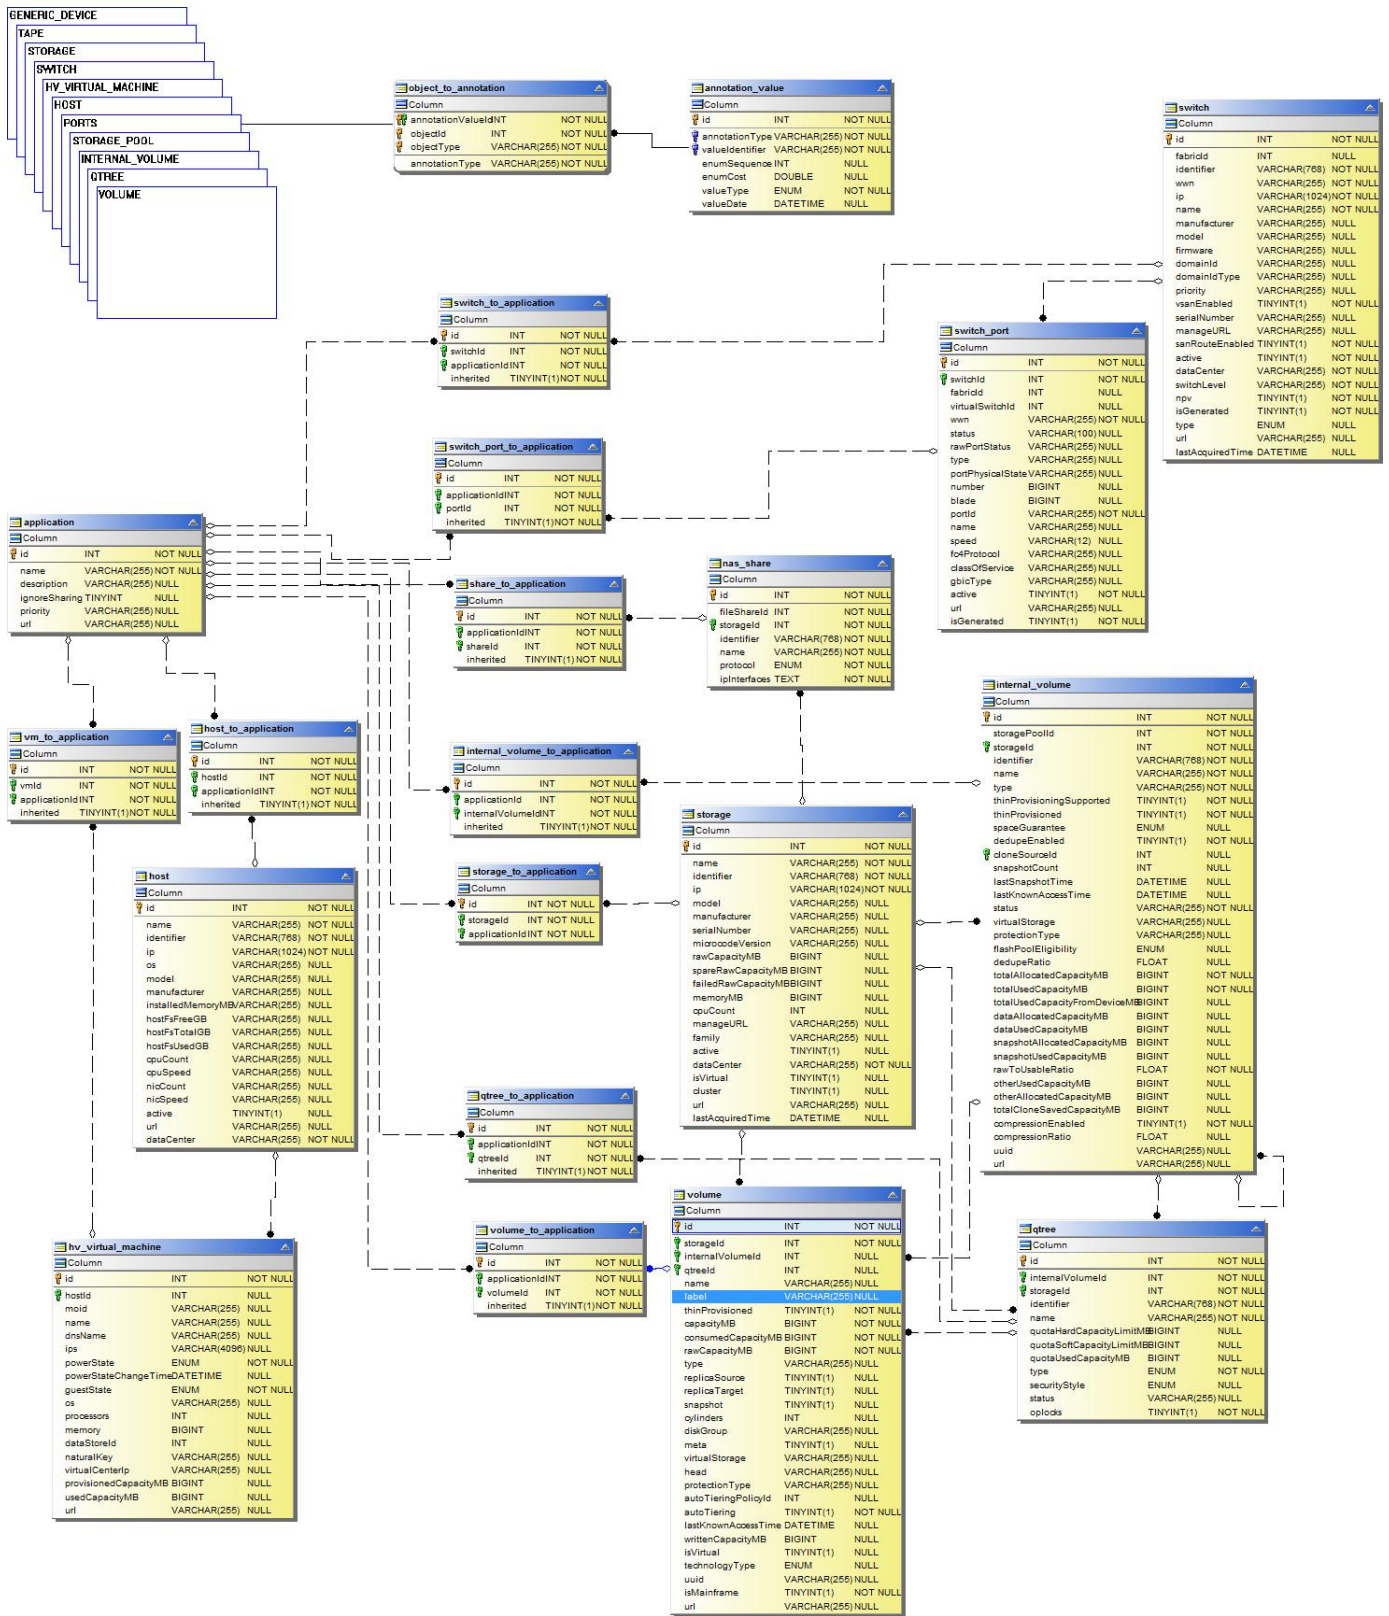
Inventory Datamart

The following images describe the inventory datamart.


Annotations




Applications




Kubernetes Metrics



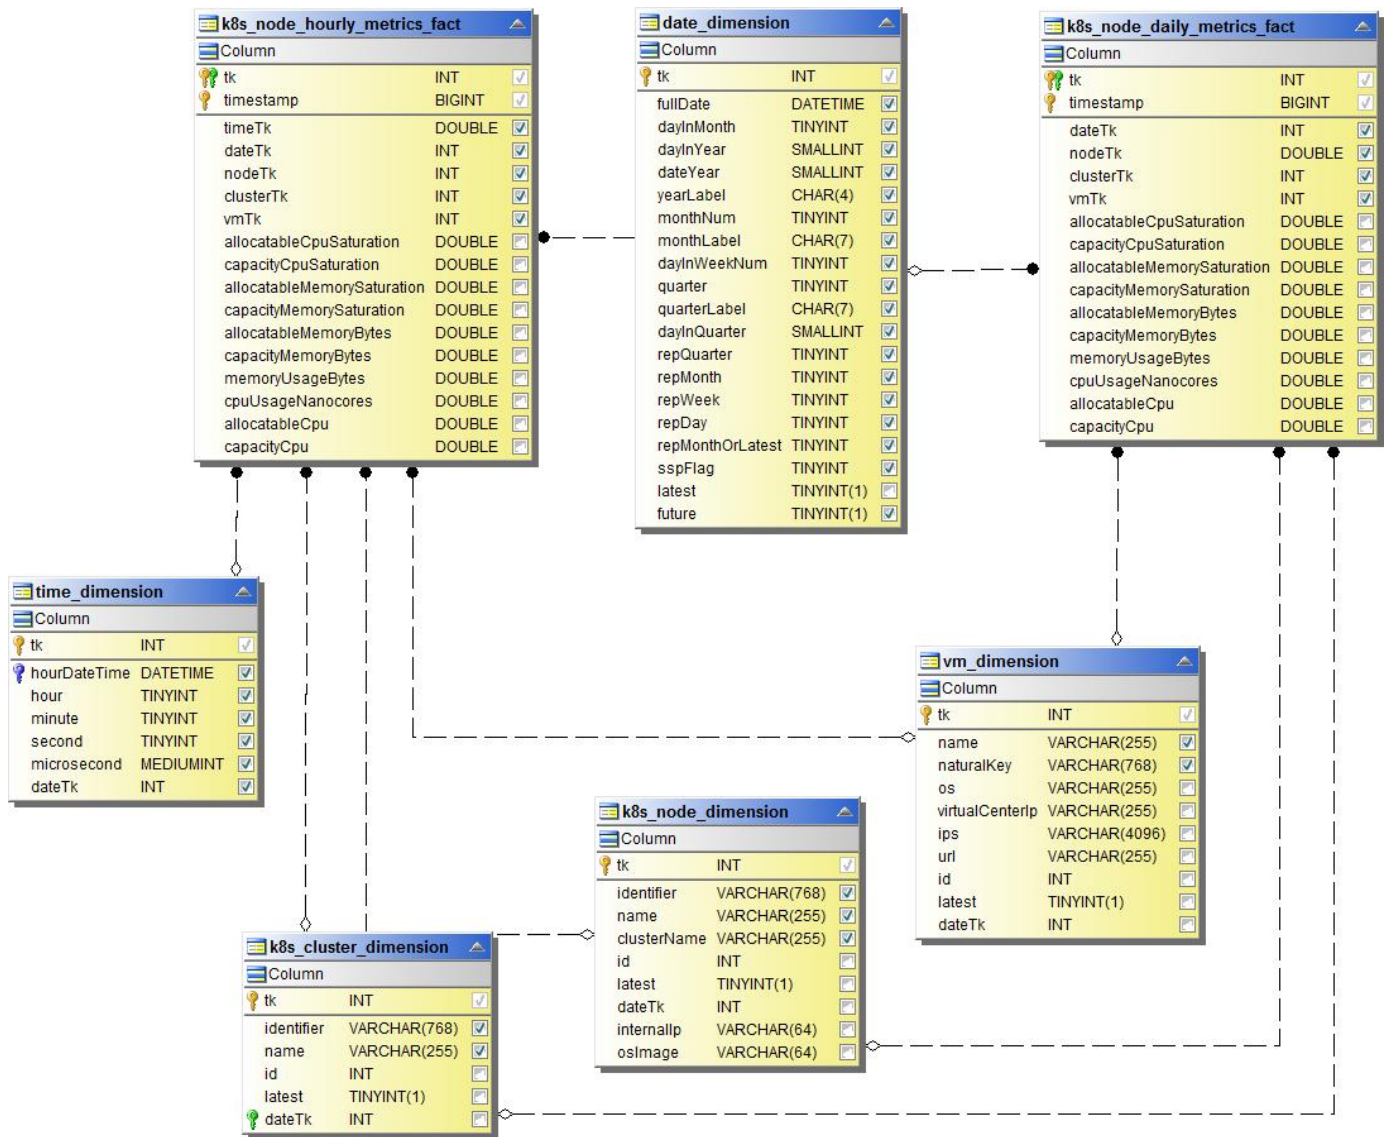
Kubernetes Cluster Metrics Fact




Kubernetes Namespace Metrics Fact




Kubernetes Node Metrics Fact




Kubernetes PVC Metrics Fact



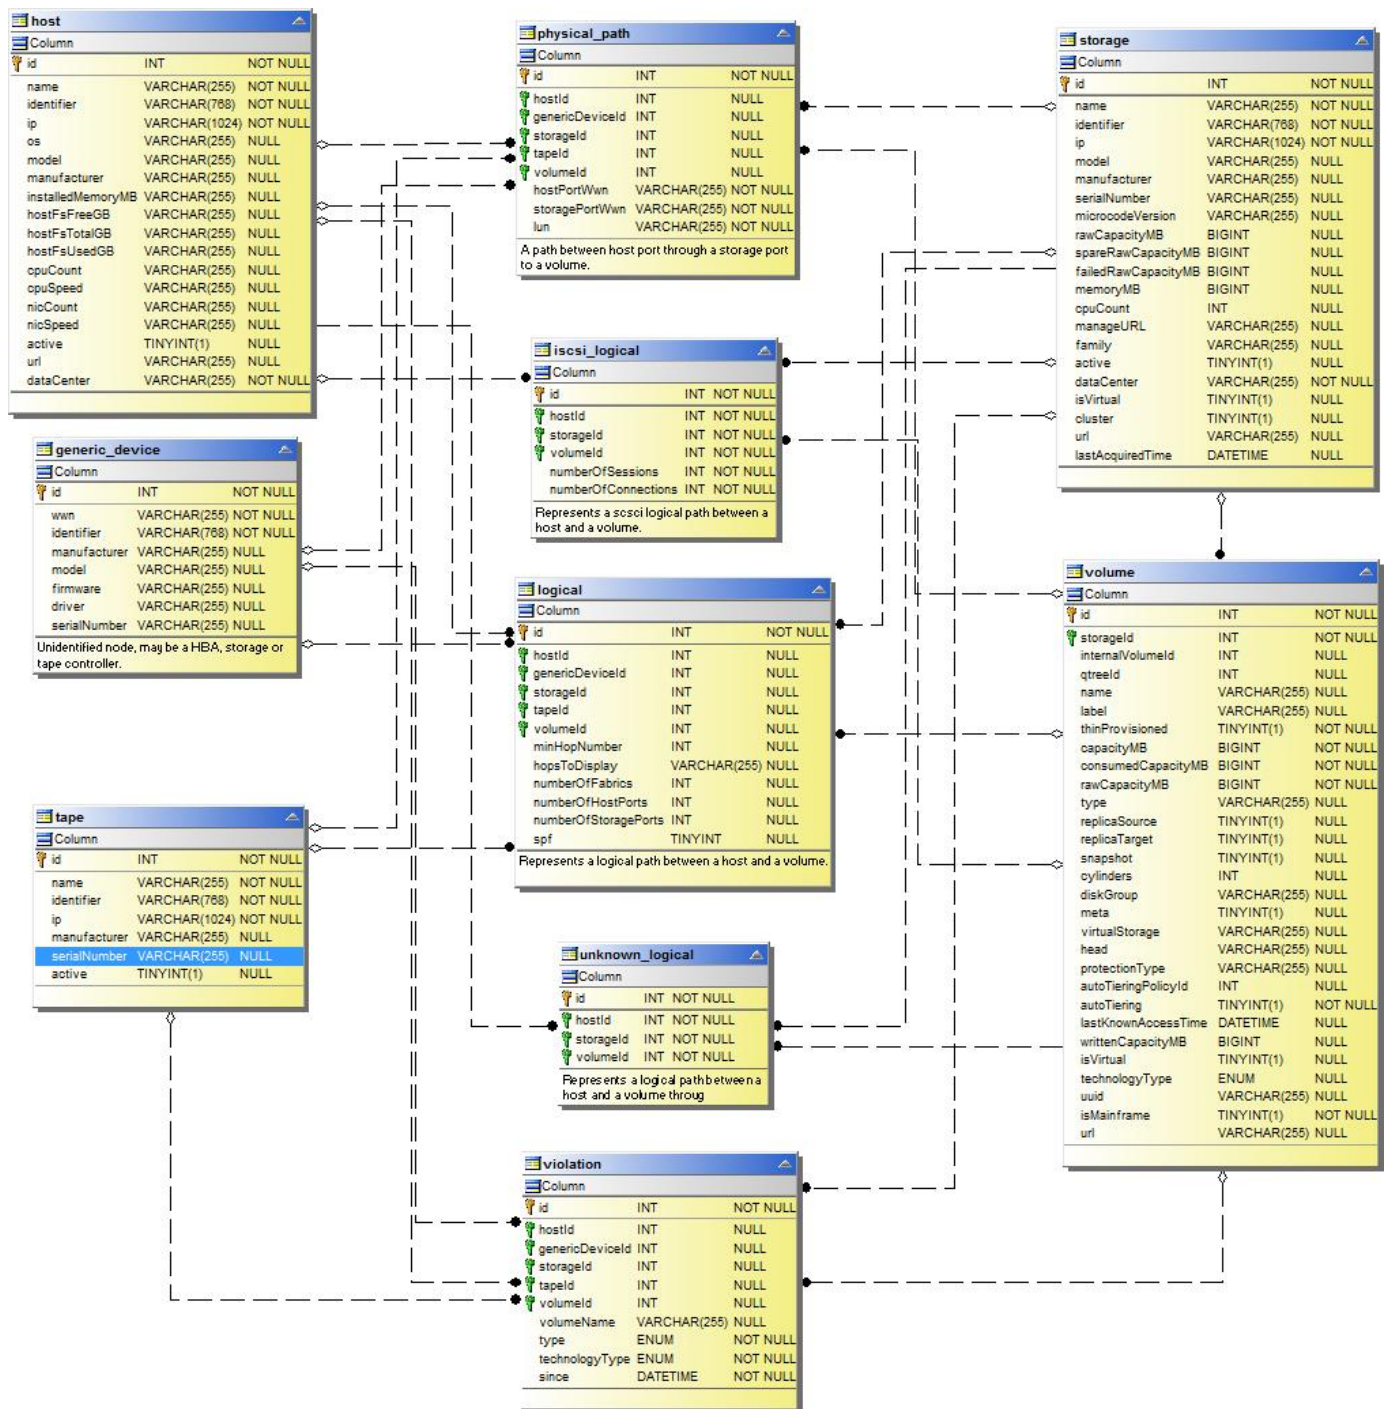
Kubernetes Workload Metrics Fact




NAS




Paths and Violations



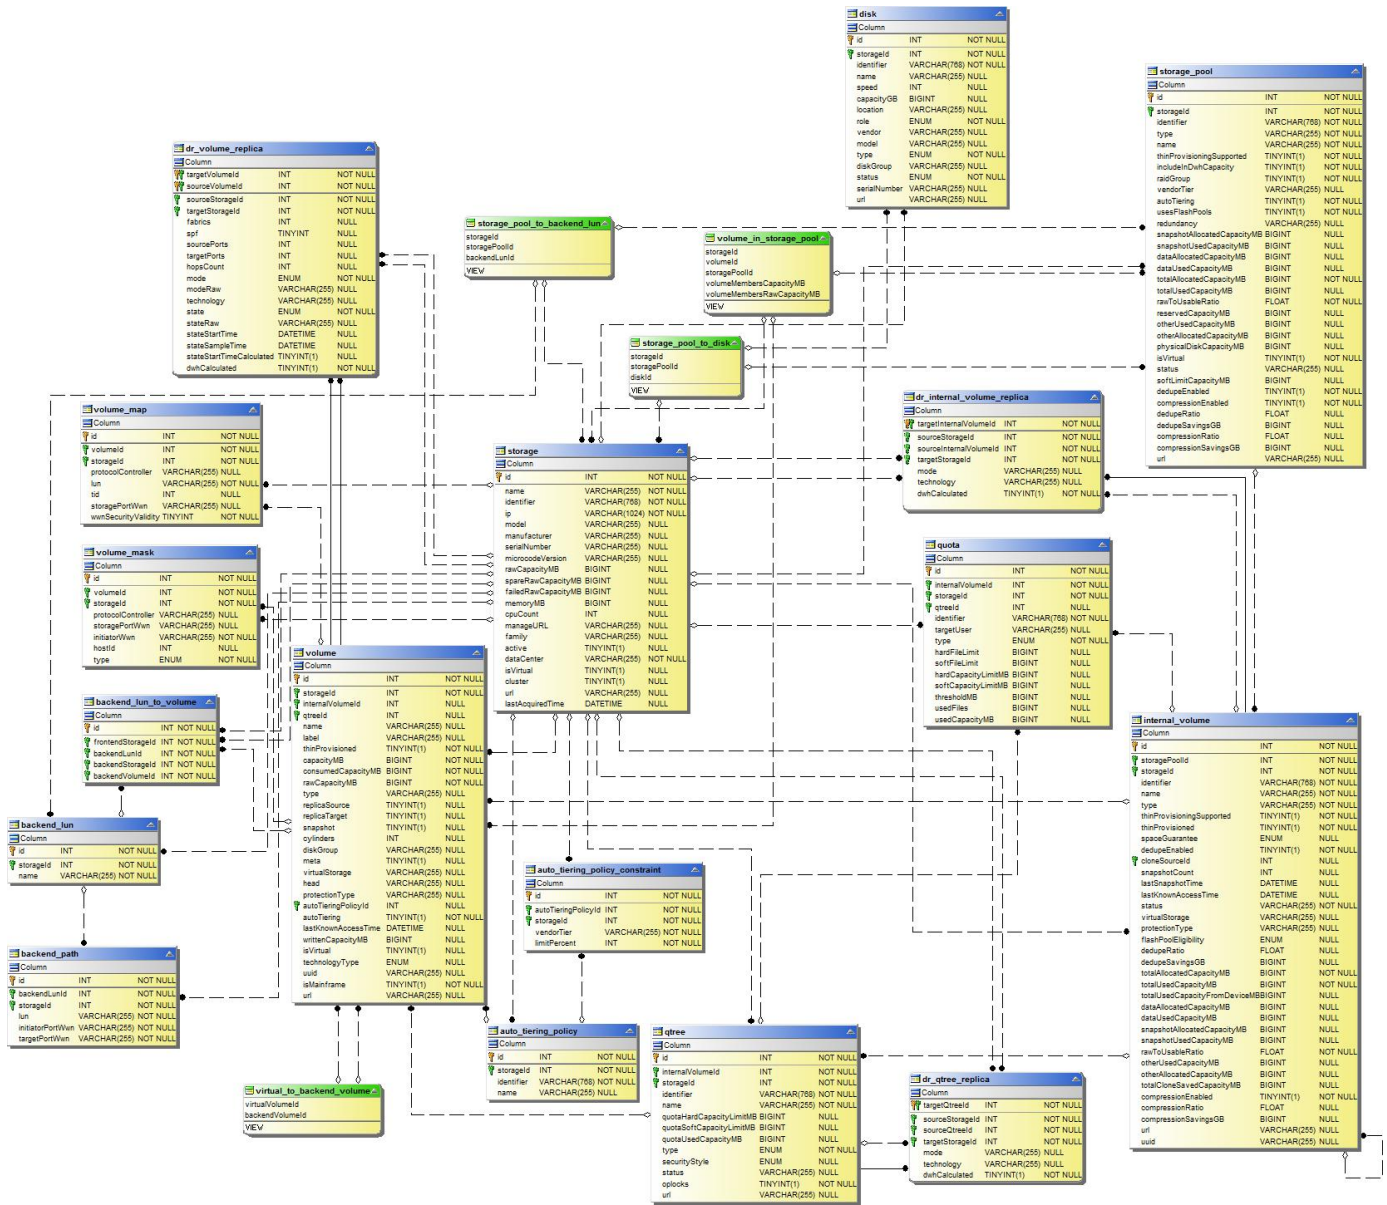
Port Connectivity



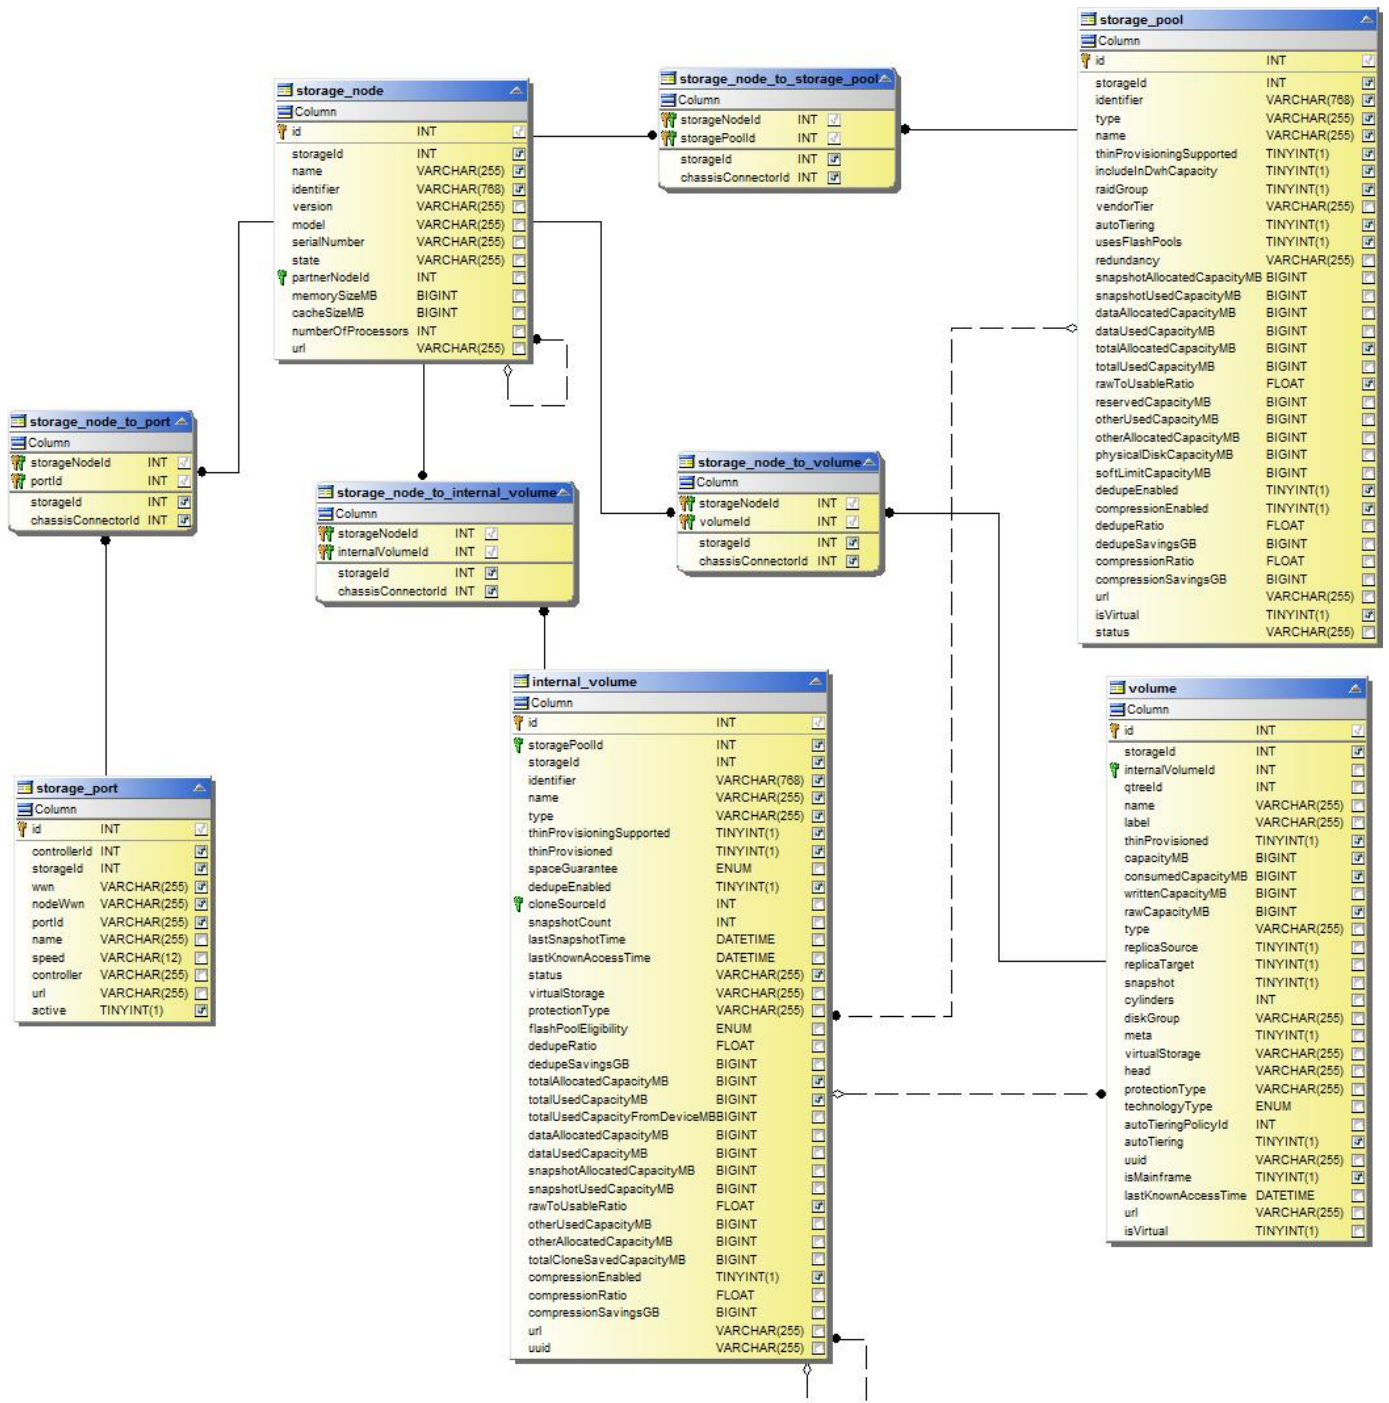
SAN Fabric




Storage



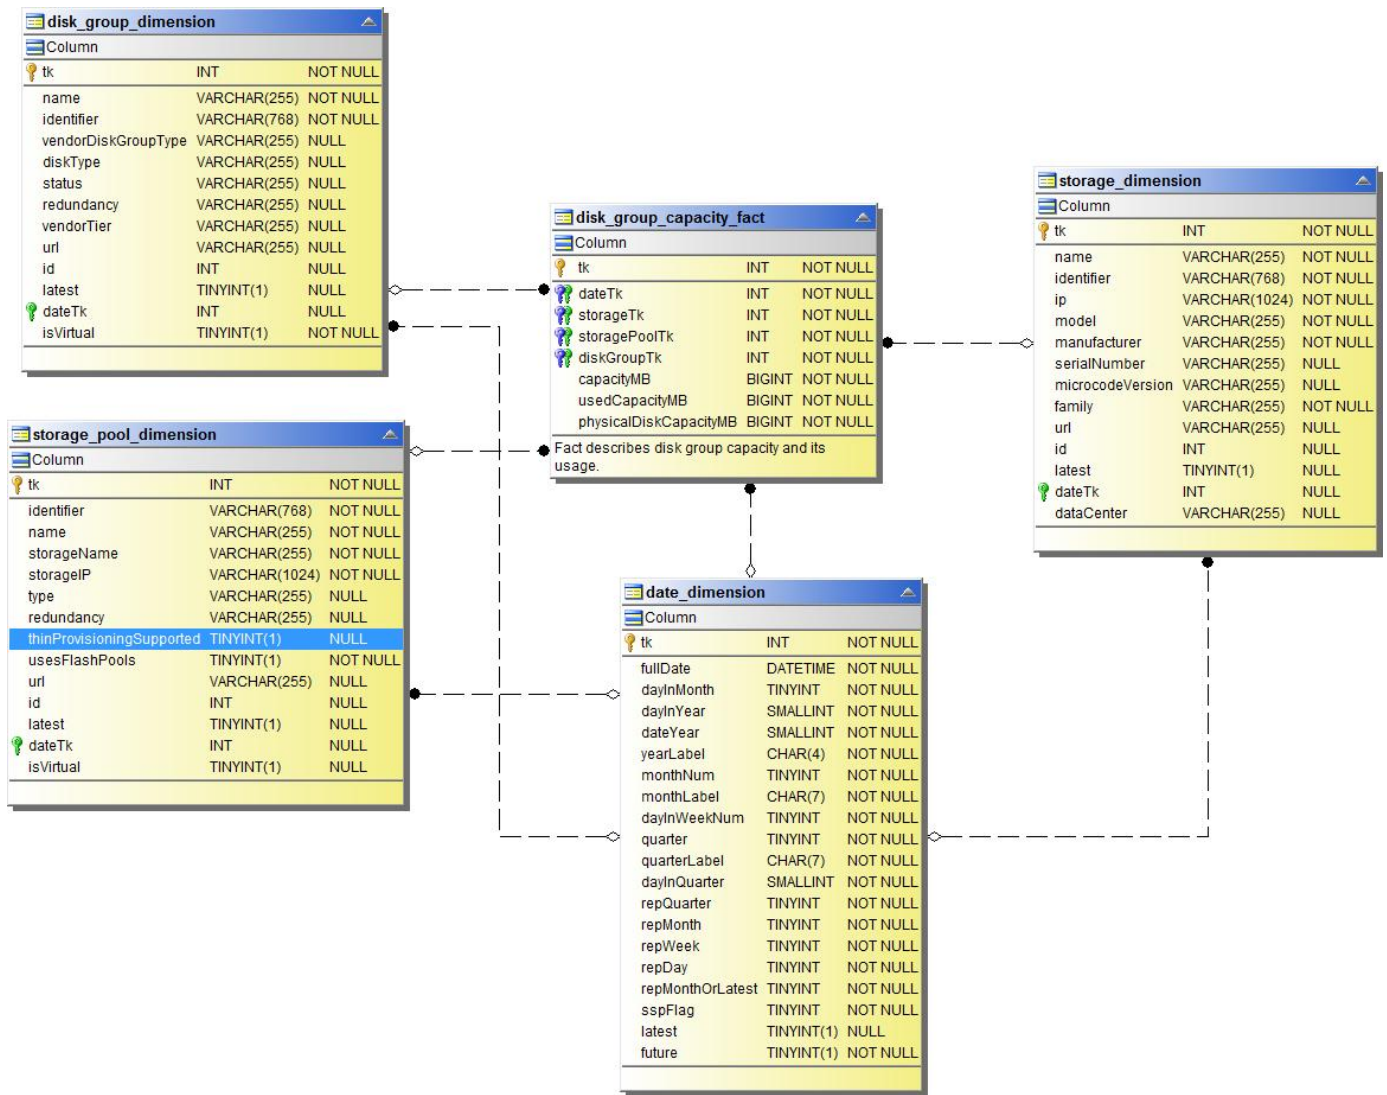
Storage Node



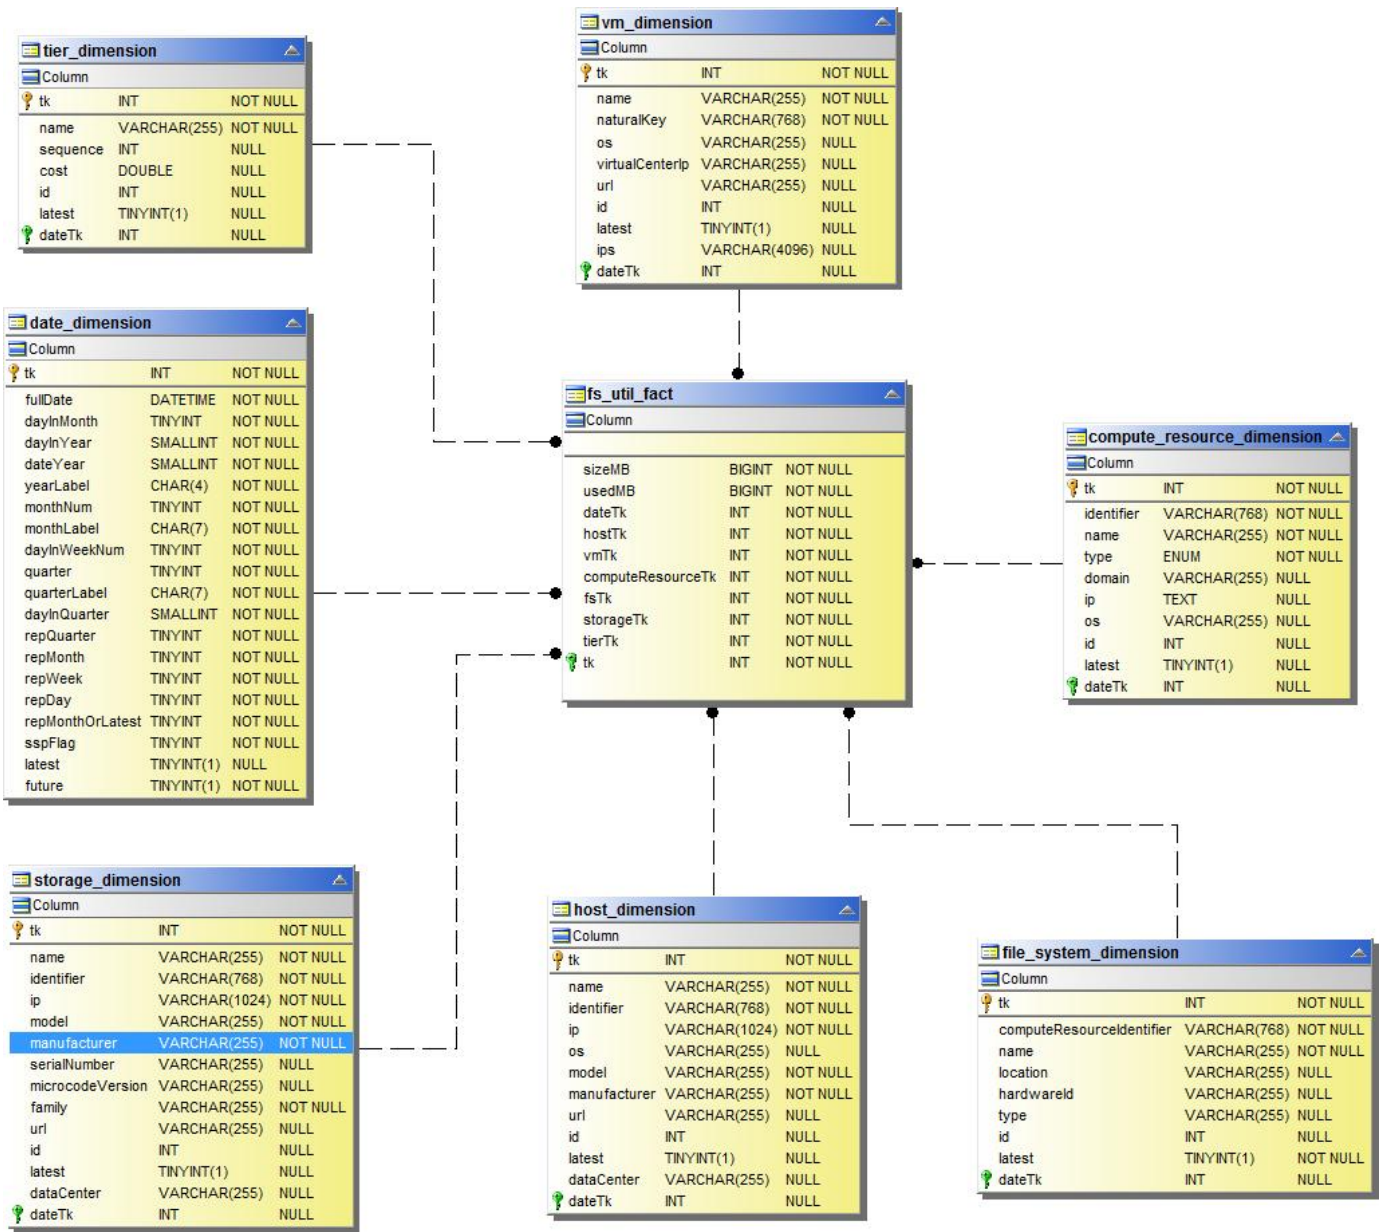
VM




Disk Group Capacity




File System Utilization



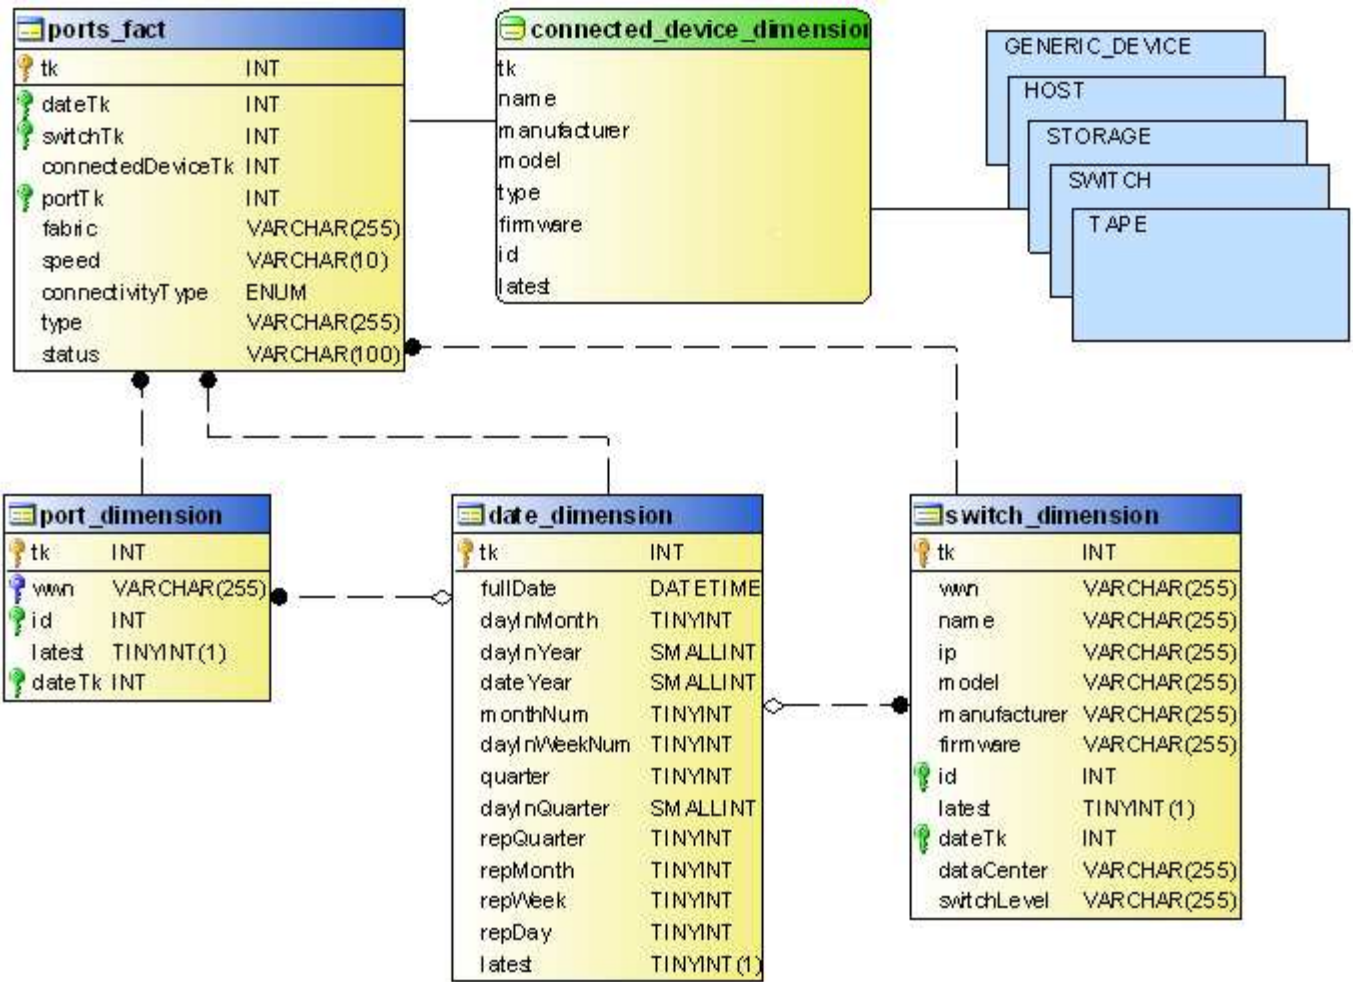
Internal Volume Capacity




Kubernetes PV Capacity



Port 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(1024)	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	

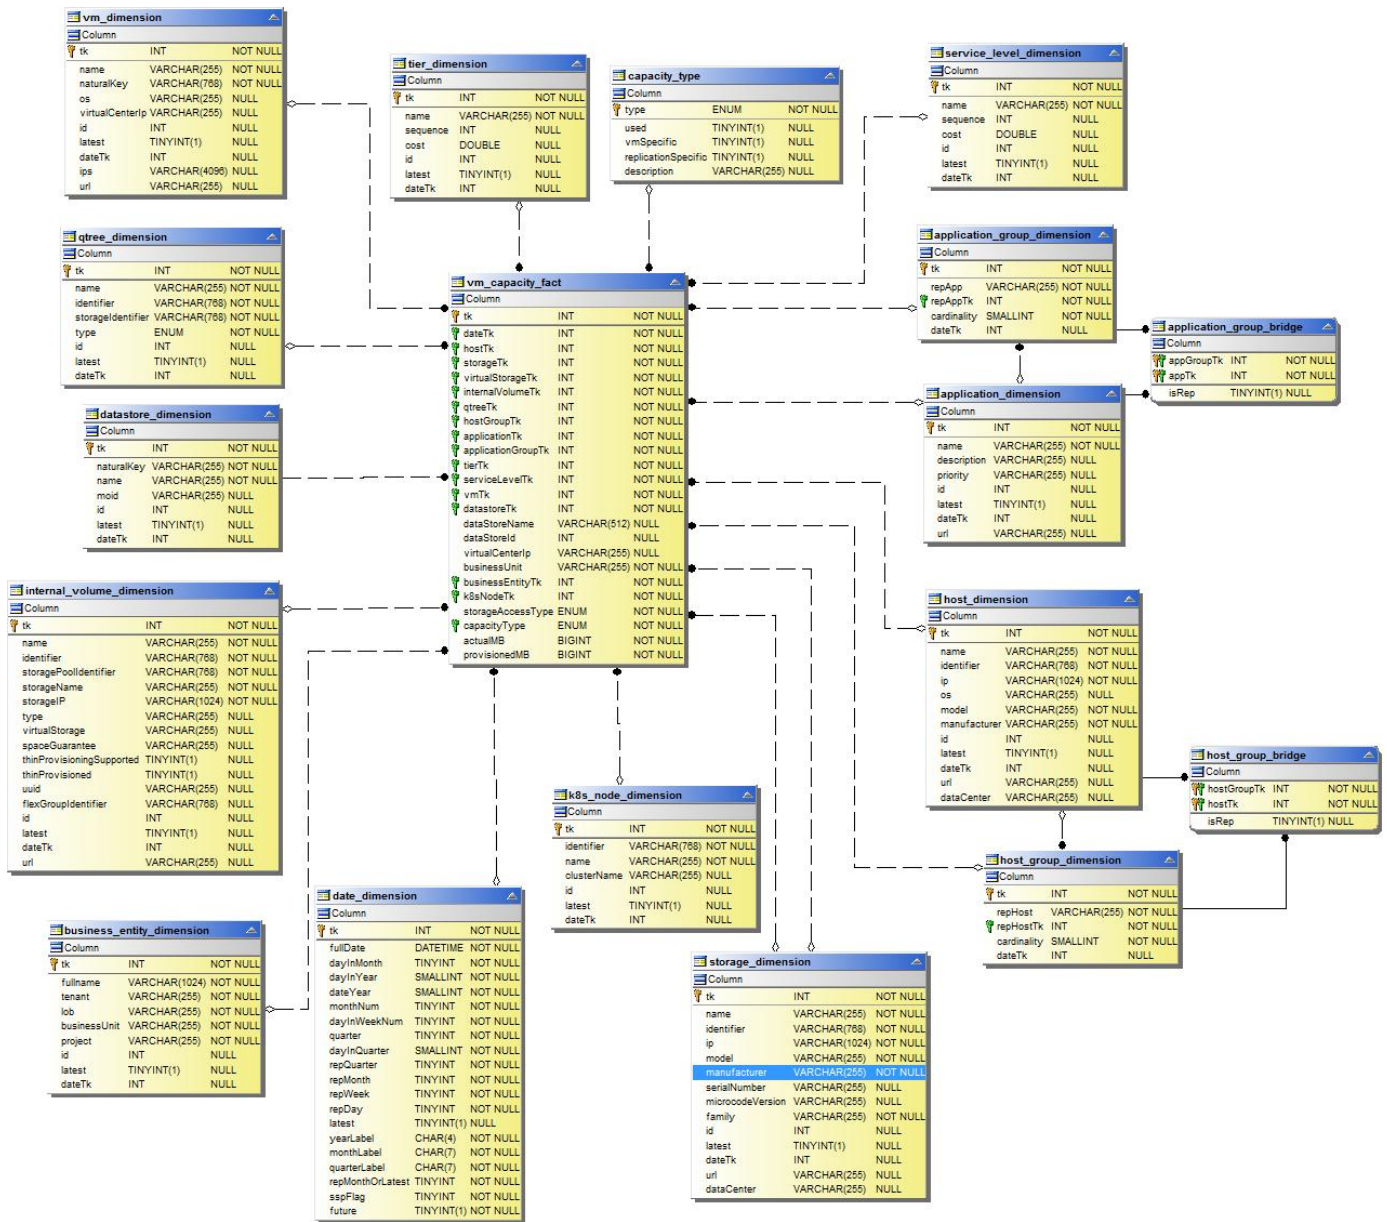
Storage and Storage Pool Capacity



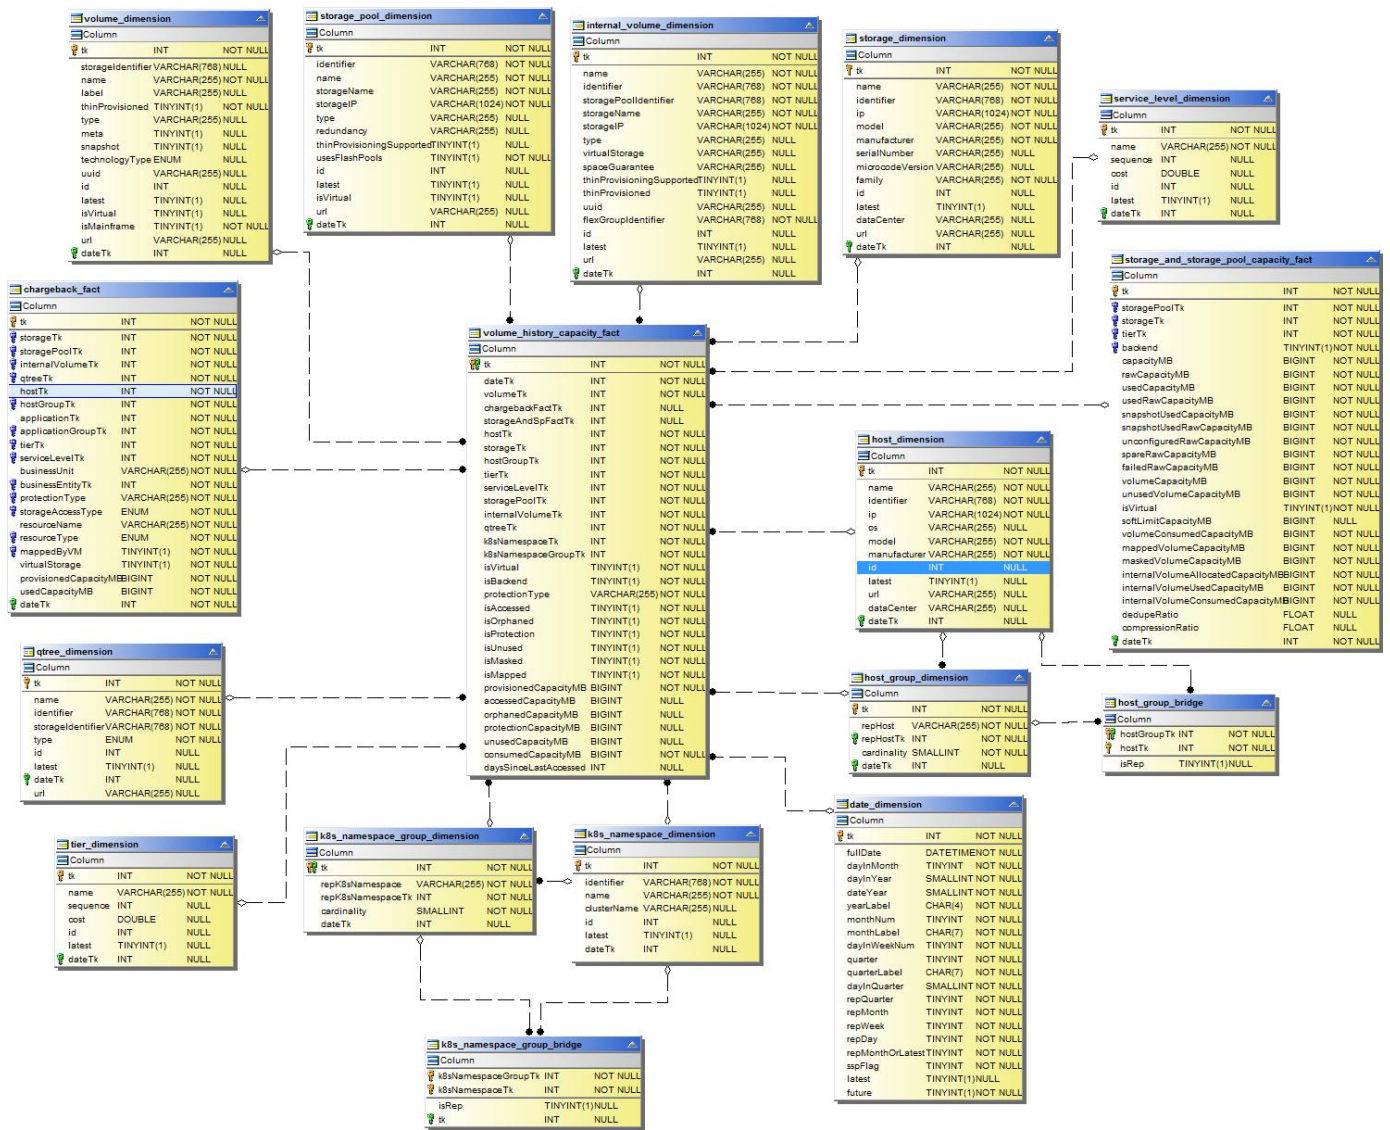
Storage Node Capacity



VM Capacity




Volume Capacity



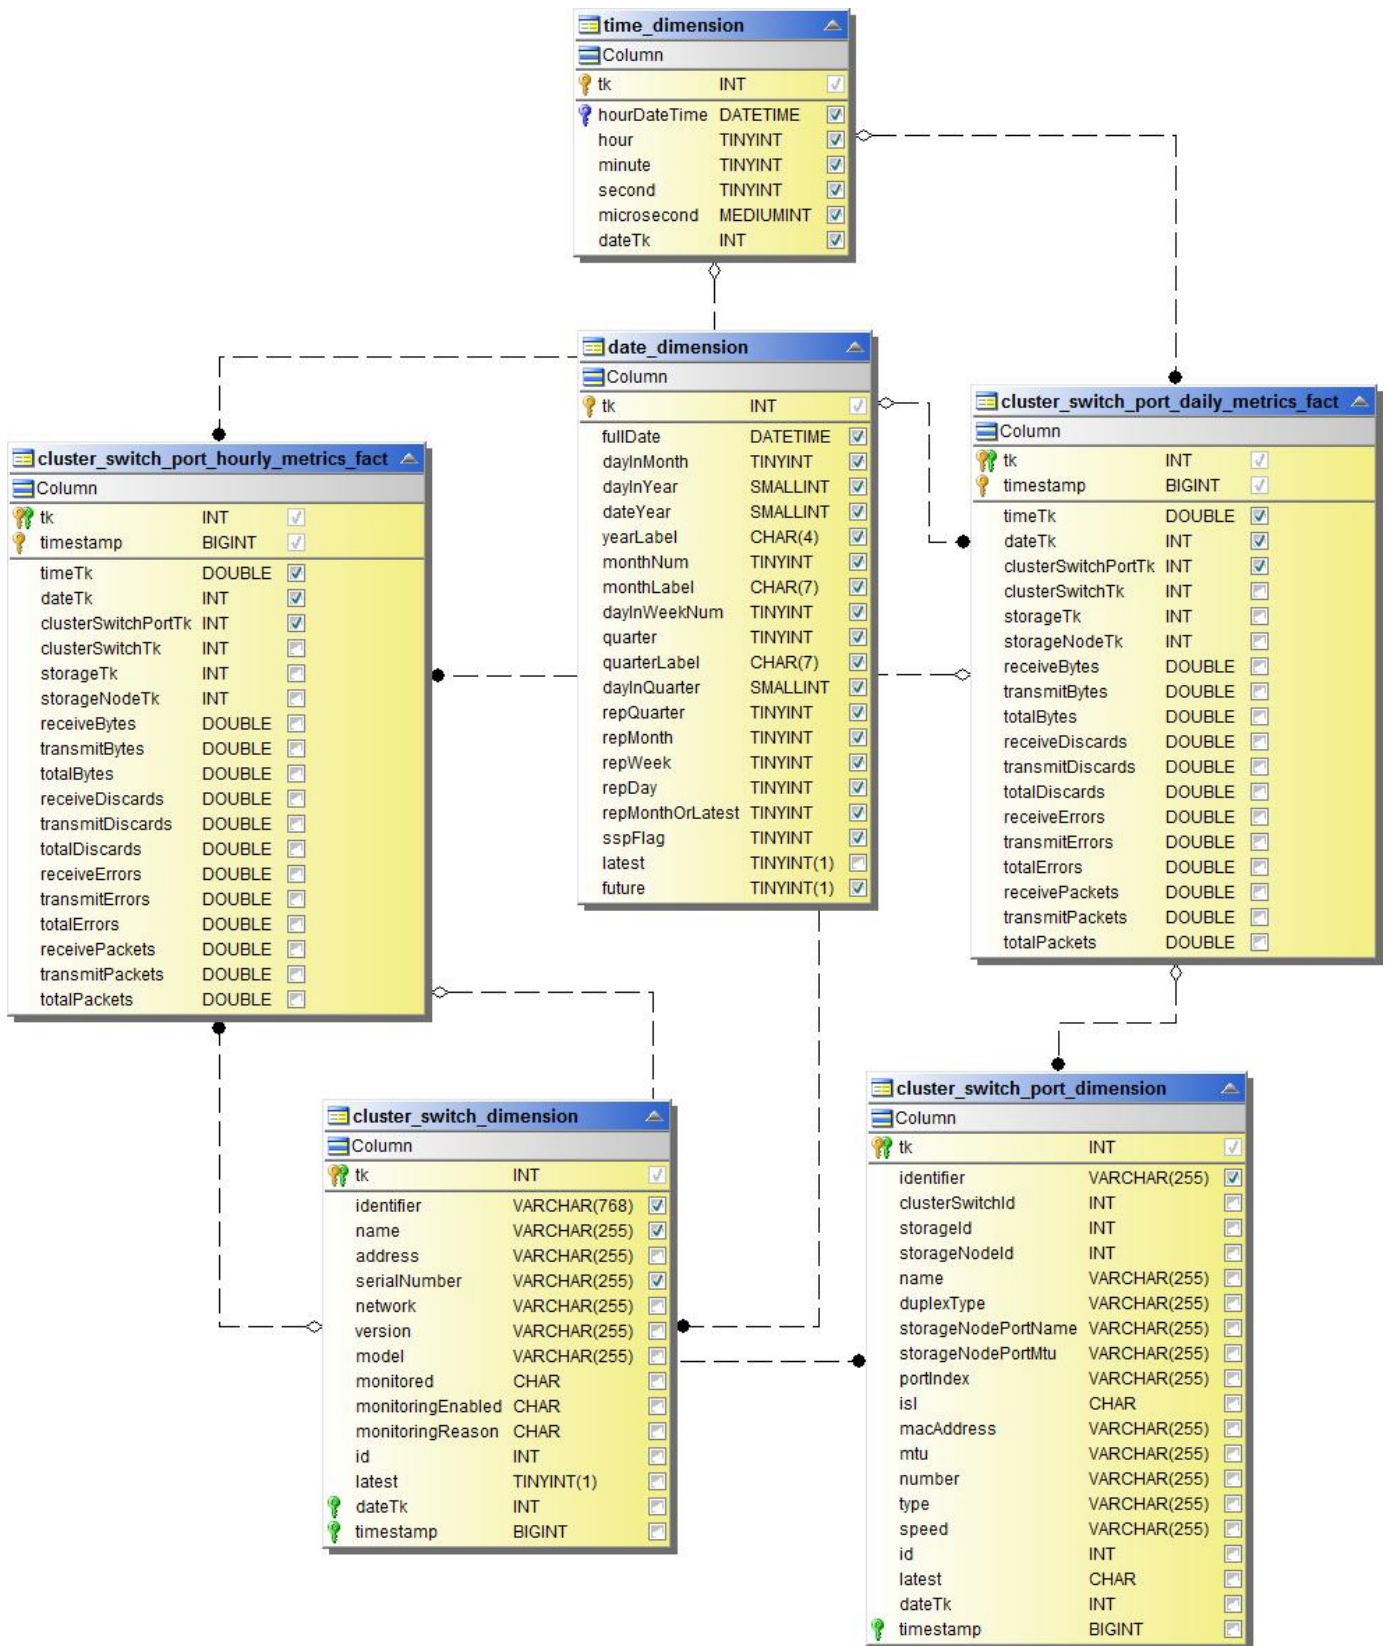
Performance Datamart

The following images describe the performance datamart.


Application Volume Hourly Performance



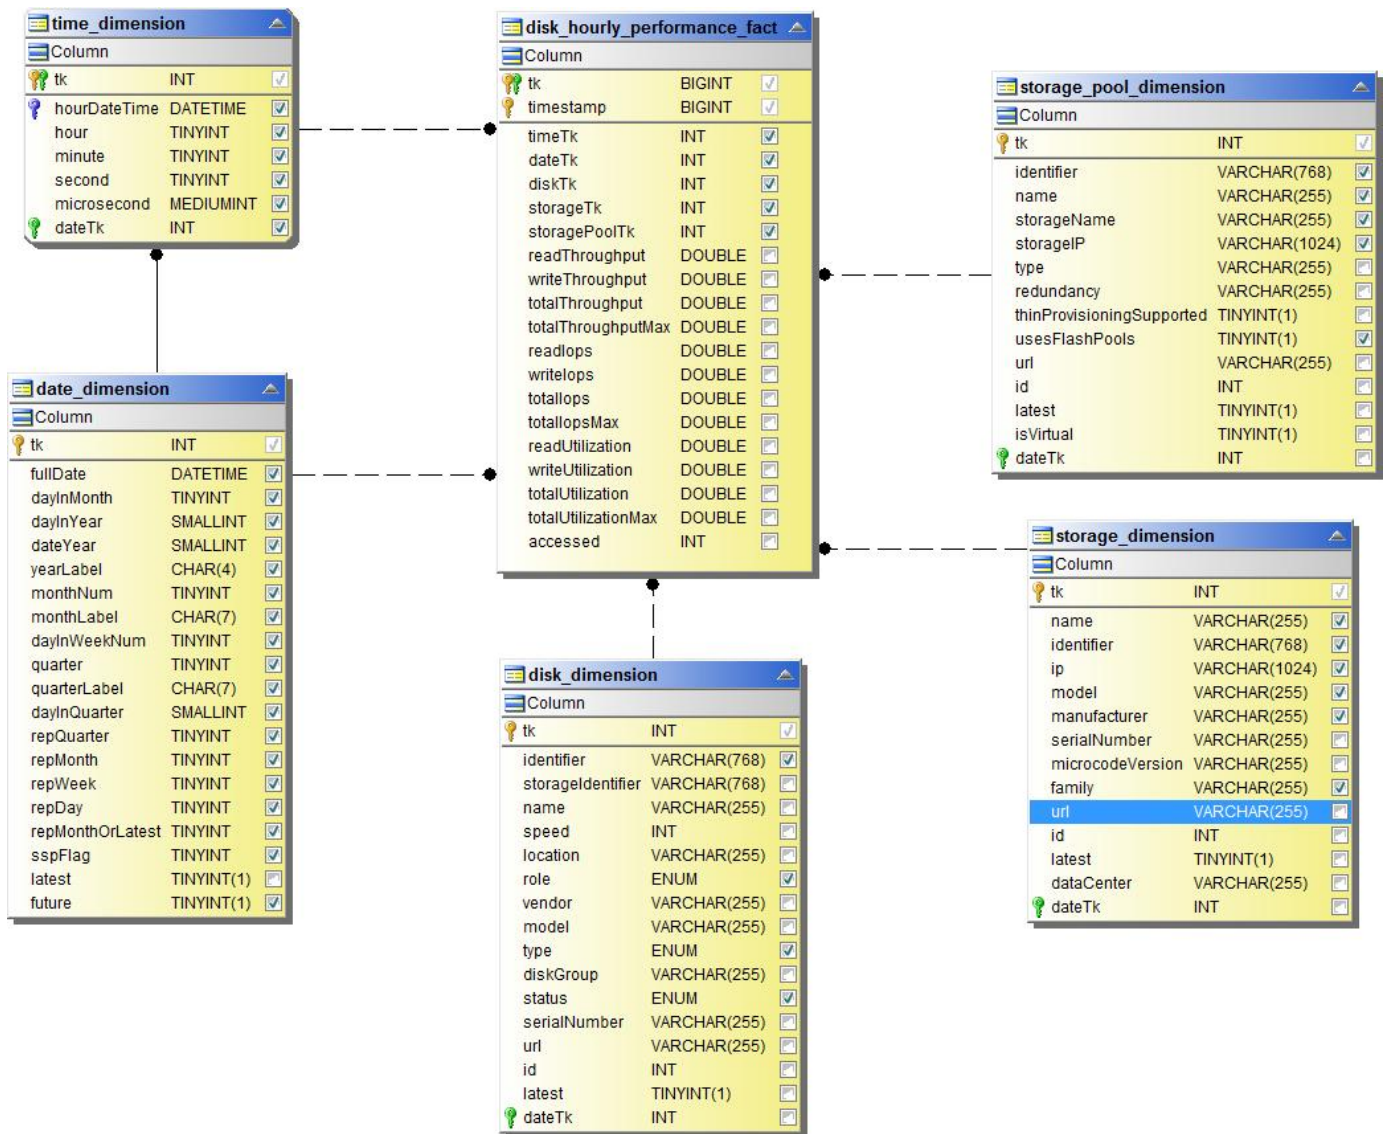
Cluster Switch Performance




Disk Daily Performance



Disk Hourly Performance



Host Hourly Performance



Host VM Daily Performance


Column		
tk	INT	<input checked="" type="checkbox"/>
dateTk	INT	<input checked="" type="checkbox"/>
hostTk	INT	<input checked="" type="checkbox"/>
readIops	DOUBLE	<input type="checkbox"/>
writeIops	DOUBLE	<input type="checkbox"/>
totalIops	DOUBLE	<input type="checkbox"/>
totalIopsMax	DOUBLE	<input type="checkbox"/>
readThroughput	DOUBLE	<input type="checkbox"/>
writeThroughput	DOUBLE	<input type="checkbox"/>
totalThroughput	DOUBLE	<input type="checkbox"/>
totalThroughputMax	DOUBLE	<input type="checkbox"/>
readResponseTime	DOUBLE	<input type="checkbox"/>
writeResponseTime	DOUBLE	<input type="checkbox"/>
totalResponseTime	DOUBLE	<input type="checkbox"/>
totalResponseTimeMax	DOUBLE	<input type="checkbox"/>
cpuUtilization	DOUBLE	<input type="checkbox"/>
maxOfAvgCpuUtilization	DOUBLE	<input type="checkbox"/>
memoryUtilization	DOUBLE	<input type="checkbox"/>
maxOfAvgMemoryUtilization	DOUBLE	<input type="checkbox"/>
swapInRate	DOUBLE	<input type="checkbox"/>
maxOfAvgSwapInRate	DOUBLE	<input type="checkbox"/>
swapOutRate	DOUBLE	<input type="checkbox"/>
maxOfAvgSWapOutRate	DOUBLE	<input type="checkbox"/>
swapTotalRate	DOUBLE	<input type="checkbox"/>
swapTotalRateMax	DOUBLE	<input type="checkbox"/>
timestamp	BIGINT	<input checked="" type="checkbox"/>
ipReceiveThroughput	DOUBLE	<input type="checkbox"/>
ipTransmitThroughput	DOUBLE	<input type="checkbox"/>
ipTotalThroughput	DOUBLE	<input type="checkbox"/>
ipTotalThroughputMax	DOUBLE	<input type="checkbox"/>

The performance daily data for host vm performance.


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"/>
sspFlag	TINYINT	<input checked="" type="checkbox"/>
latest	TINYINT(1)	<input type="checkbox"/>
future	TINYINT(1)	<input checked="" type="checkbox"/>

Column		
tk	INT	<input checked="" type="checkbox"/>
dateTk	INT	<input type="checkbox"/>
name	VARCHAR(255)	<input checked="" type="checkbox"/>
identifier	VARCHAR(768)	<input checked="" type="checkbox"/>
ip	VARCHAR(1024)	<input checked="" type="checkbox"/>
os	VARCHAR(255)	<input type="checkbox"/>
model	VARCHAR(255)	<input checked="" type="checkbox"/>
manufacturer	VARCHAR(255)	<input checked="" type="checkbox"/>
id	INT	<input type="checkbox"/>
latest	TINYINT(1)	<input type="checkbox"/>
dataCenter	VARCHAR(255)	<input type="checkbox"/>
url	VARCHAR(255)	<input type="checkbox"/>


Host VM Hourly Performance




Internal Volume Hourly Performance




Internal Volume Daily Performance




Qtree Daily Performance




Qtree Hourly performance



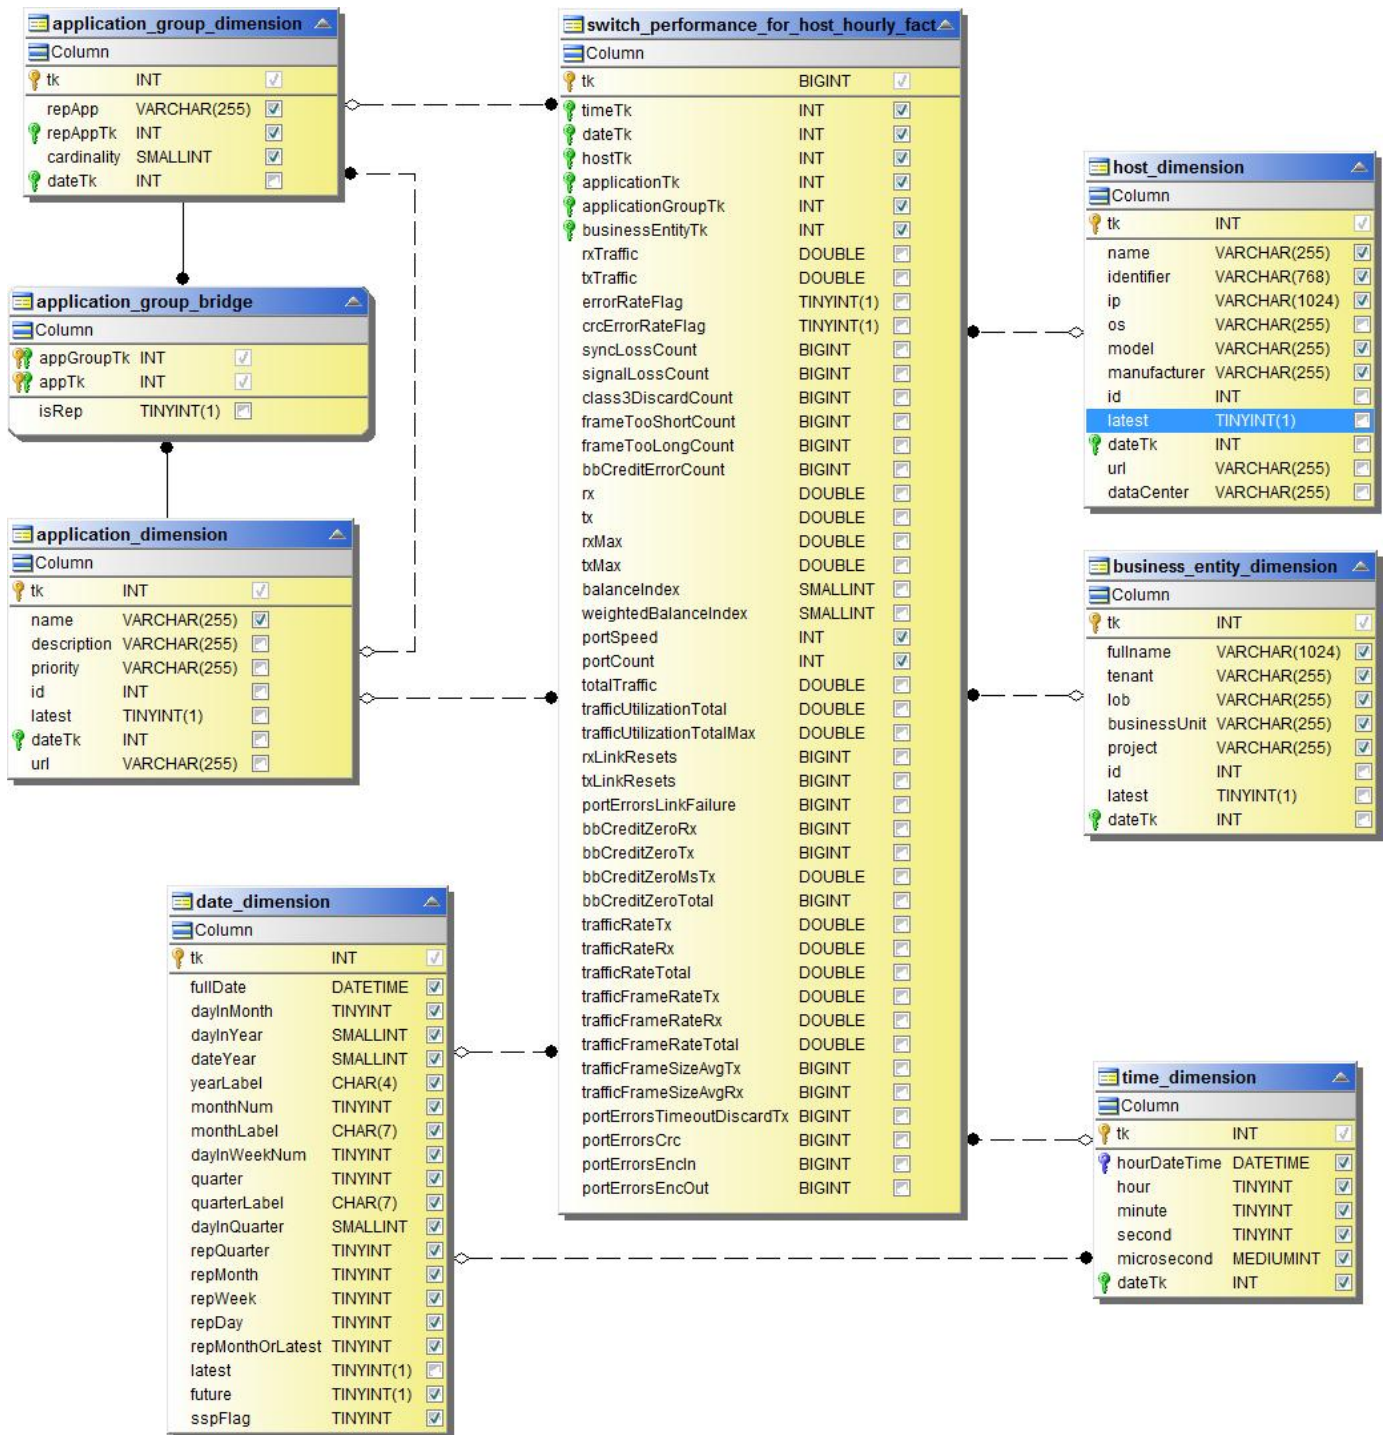
Storage Node Daily Performance



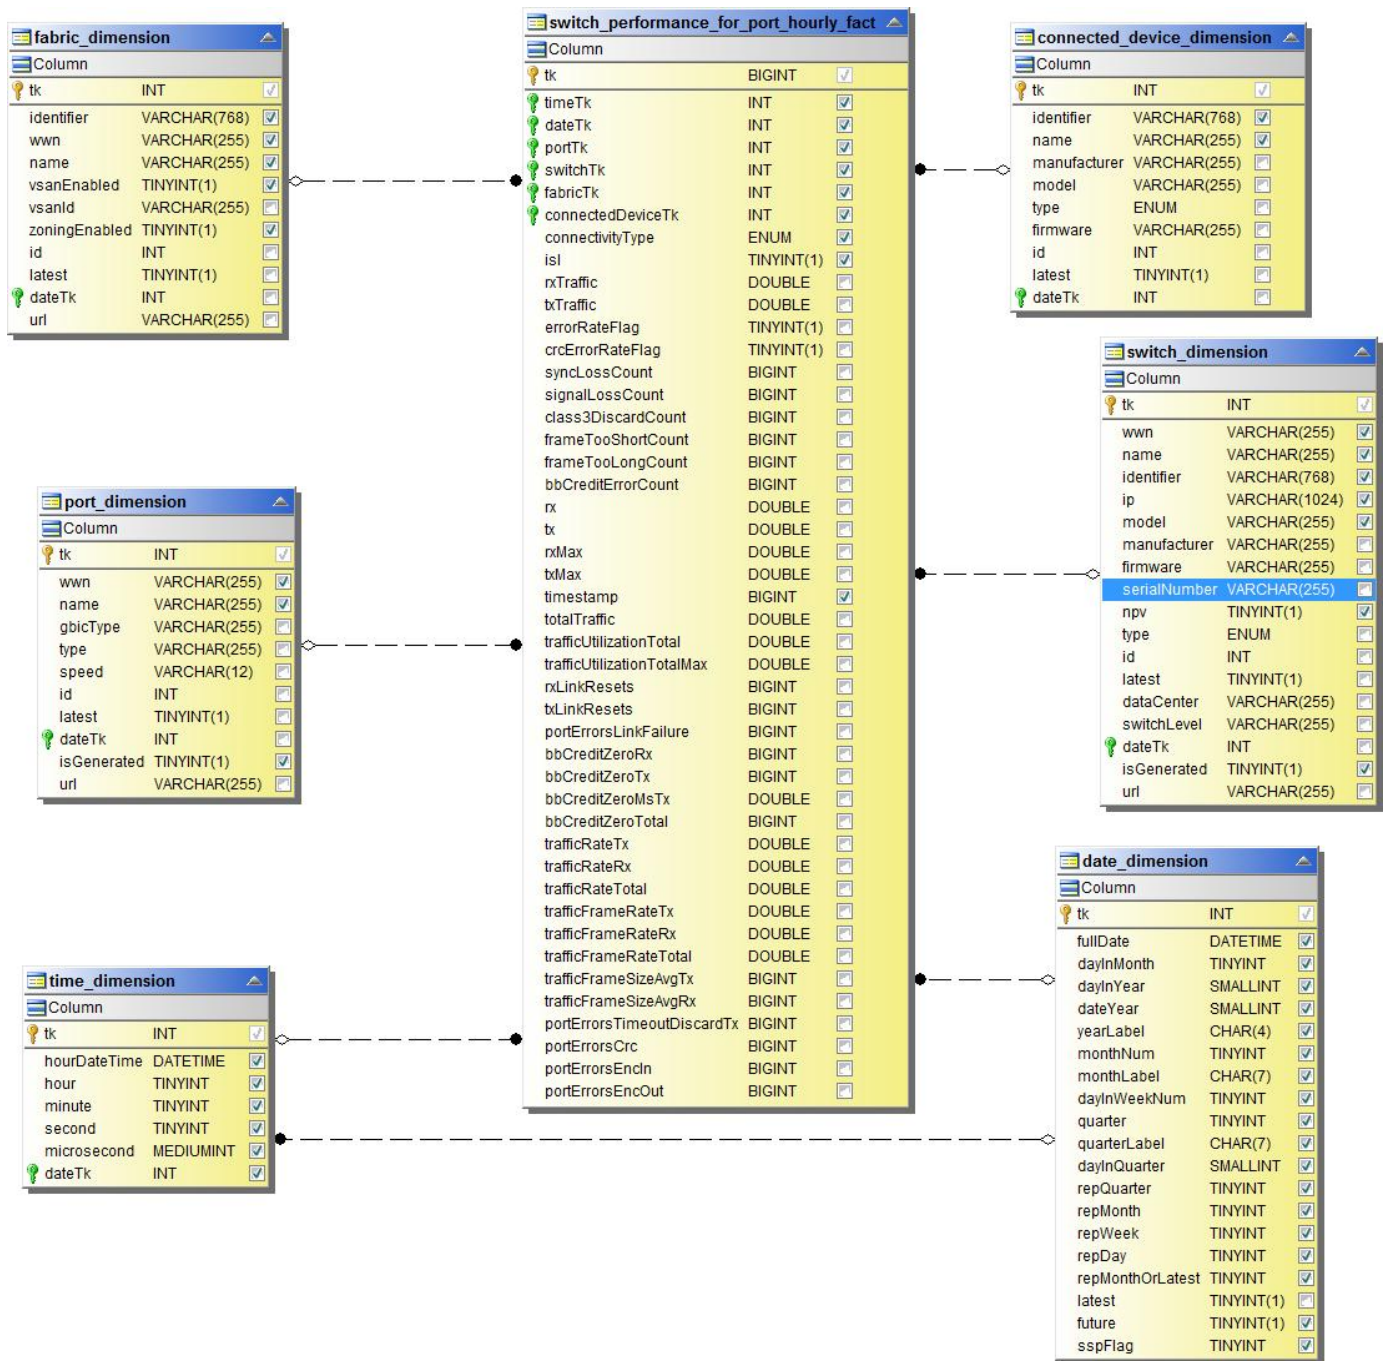
Storage Node Hourly Performance




Switch Hourly Performance for Host



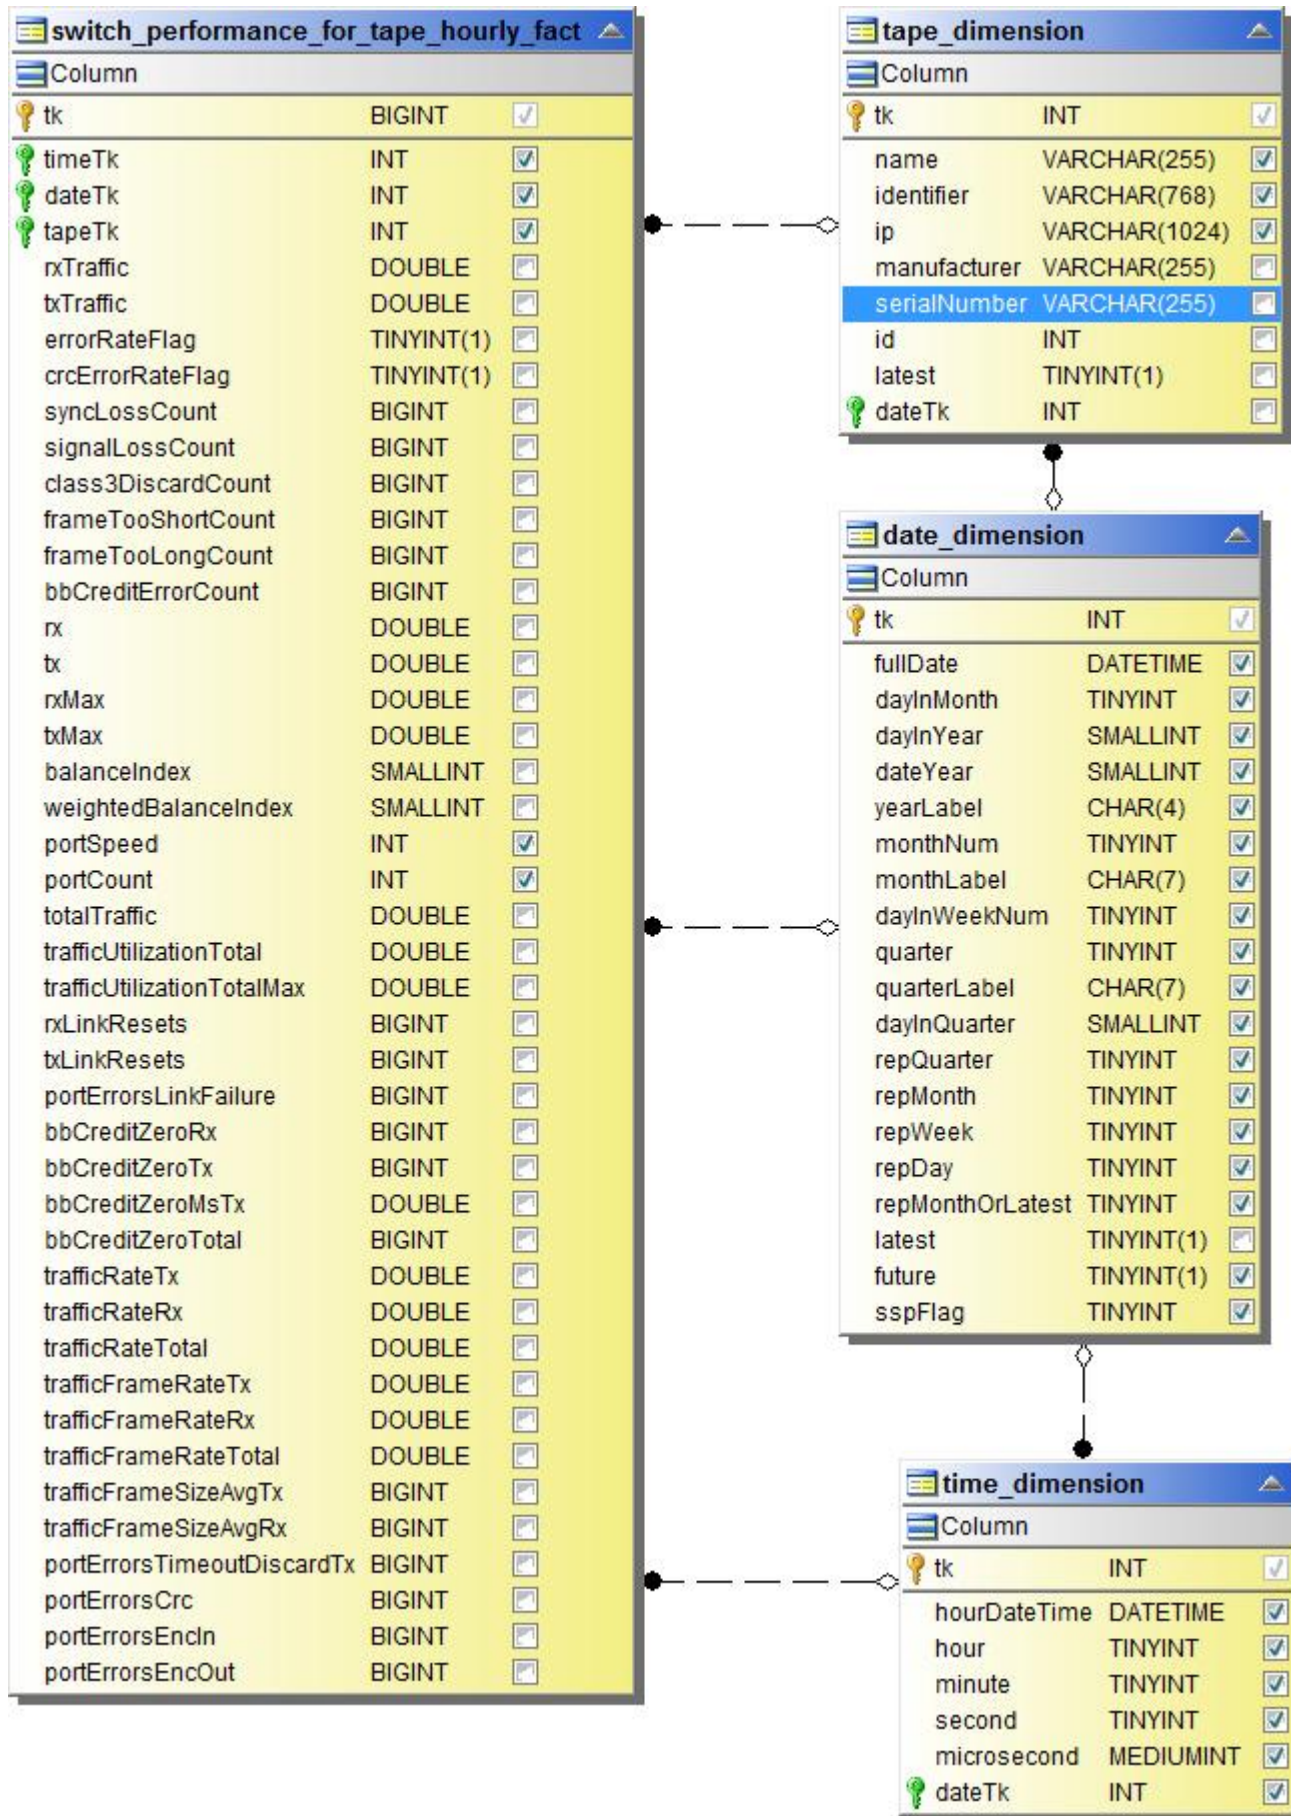
Switch Hourly Performance for Port




Switch Hourly Performance for Storage




Switch Hourly Performance for Tape




VM Performance




VM Daily Performance for Host



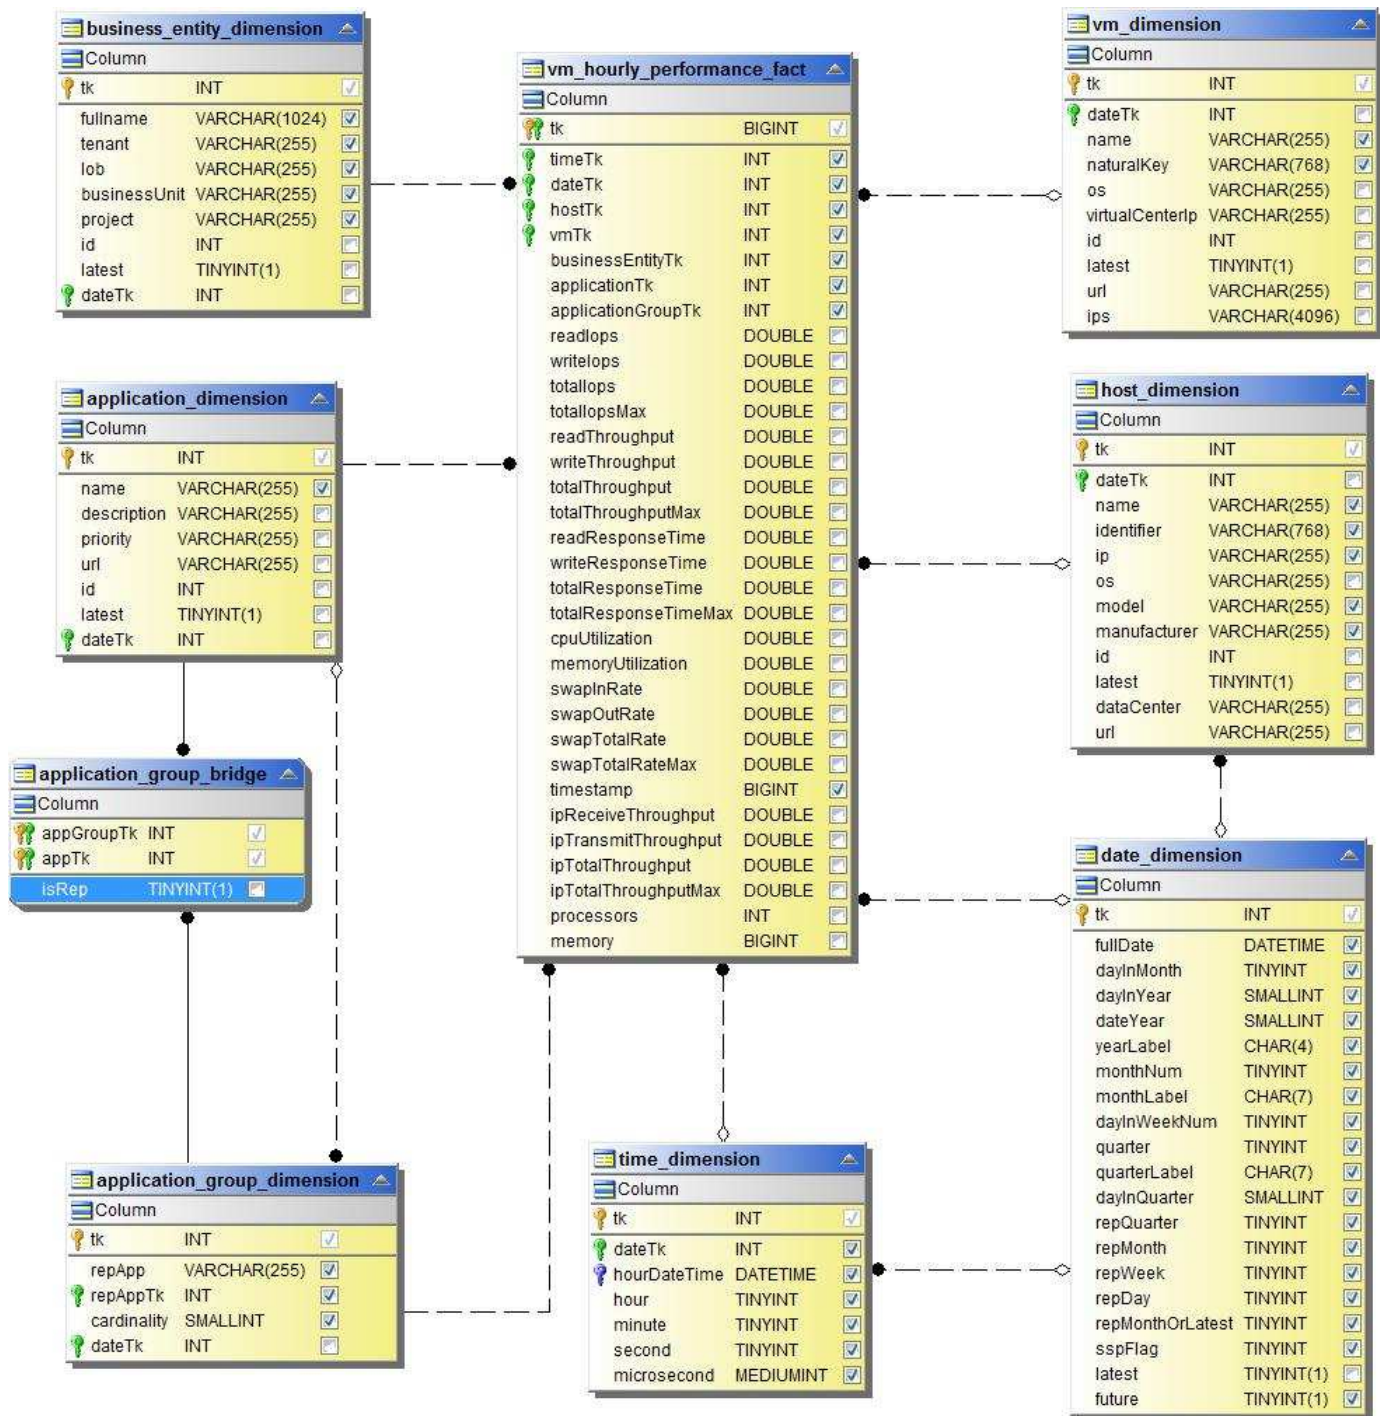
VM Hourly Performance for Host




VM Daily Performance for Host




VM Hourly Performance for Host




VMDK Daily Performance




VMDK Hourly Performance



Volume Hourly Performance



Volume Daily Performance



Data Infrastructure Insights Schemas for Reporting

These schema tables and diagrams are provided here as a reference for Data Infrastructure 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 Data Infrastructure Insights [Premium Edition](#).

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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