



How-To guides

OnCommand Insight

NetApp
April 01, 2024

This PDF was generated from <https://docs.netapp.com/us-en/oncommand-insight/howto/setting-up-email-for-notifications.html> on April 01, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- How-To guides 1
 - Getting Started with Insight 1
 - Creating custom dashboards 14
 - Creating performance policies 48
 - Troubleshooting Fibre Channel BB credit 0 errors 52
 - Analyzing your infrastructure 57
 - Introduction to minimizing risk in thin provisioning 62
 - Collecting Host and VM file system utilization data 68
 - Configuring your system to report chargeback data 72
 - Ensuring IO density reports describe only internal data volumes 78
 - Collecting integration data 80
 - Analyzing an application performance problem 89
 - Collecting and reporting AWS billing data 97
 - Integrating with ServiceNow 100

How-To guides

Getting Started with Insight

After OnCommand Insight has been installed and properly licensed, there are a number of tasks you should do to begin preparing your environment to showcase the data that is important to you.

Some of the tasks performed in a typical environment include the following:

1. **Annotating your assets** to prepare them for querying and reporting. Useful initial annotations typically include data Center, Tier, and Service Level.
2. **Creating queries** to show important data and help troubleshooting
3. **Assigning applications and business entities** to assets
4. **Creating performance policies and alerts** for violations against those policies
5. **Creating custom dashboards** to highlight data according to need or user role

Setting up notifications

You can configure Insight to send notifications on trigger events such as performance policy, global path, or capacity violations via email, SNMP, or Syslog. You can also configure Insight to send email notifications on system-level events such as data source errors or acquisition unit failures.

These are basic instructions. For more detail on notifications, see Configuration and administration > Insight configuration and administration > Setting up Insight.

Setting up email for notifications

Insight can send email notifications on trigger events, such as performance policy violations.

About this task

Follow these basic steps to configure email notifications:

Steps

1. Click **Admin > Notifications** and go to the **Email** section.
2. In the **Server** box, enter the name of your SMTP server. You can enter either a fully-qualified domain name or an IP address.
3. Enter your SMTP user name and (if required by your SMTP server) password.
4. In the **Sender email** box, enter the sender email account that will be identified as the sender on the notifications.

This account must be a valid email account within your organization.

5. In the **Email signature** box, enter any text that you want to be inserted in every email that is sent.

6. In the **Recipients** box, click **+** to enter an email address, and click **OK**.
7. Click **Save**.

To edit or remove an email address, or to send a test email, select the address and click the appropriate button that appears.

Note that you can configure Insight to send email notifications for specified performance policy violations to specific individuals or groups. For example, you might send cloud asset violations to one group, and physical host events to another. Go to **Manage > Performance policies** to configure individual policy notifications.

Setting up Syslog for logging

Insight can send syslog events for capacity or path violations and performance alerts.

About this task

Follow these basic steps to configure syslog notification in Insight:

Steps

1. Click **Admin > Notifications** and go to the **Syslog** section.
2. Place a check in the **Syslog enabled** checkbox..
3. In the **Server** field, enter the IP address of the log server.
4. In the **Facility** field, select the facility level that corresponds to the type of program that is logging the message.
5. Click **Save**.

Setting up SNMP for notifications

Insight can send SNMP notifications on trigger events, such as violations or when data source thresholds are exceeded.

About this task

Follow these basic steps to configure SNMP in Insight:

Steps

1. Click **Admin > Notifications** and go to the **SNMP** section.
2. Click **Actions** and select **Add trap source**.
3. In the **Add SNMP trap recipients** dialog box, enter the **IP** address and **Port** to which you want SNMP trap messages sent. For **Community String**, use "public" for SNMP trap messages.
4. Click **Save**.

Preparing assets: Annotating

Annotating allows you to associate specific tags or labels to the assets you choose, which aids in managing and reporting on those assets.

Creating annotations for your enterprise

This guide describes how to create and customize annotations for your environment that can be used for querying, filtering, alert notifications and reporting.

An annotation is a note or tag that you associate with specific assets in your environment. OnCommand Insight provides several annotations that you can configure for your assets as needed, or you can create your own custom annotations based on your business needs.

The examples that follow are those that are typically configured first in new customer environments, to serve as a baseline for additional actions. Your own annotation needs may vary, but the steps described herein can be used as a guide to configuring any annotations you may need on the assets you desire.

This guide is based on the following assumptions:

- You have OnCommand Insight Server installed and properly licensed.
- You want to explore best practices, not every available option.
- You understand that these are examples only and that your specific needs may vary.

This guide walks you through modifying existing annotations as well as creating custom ones

In our example environment, we wish to be able to list assets according to Data Center, Tier, Service Level and Environment.

Configuring Data Center annotations

The Data Center annotation is typically used to associate a storage array, switch, or physical host asset with a data center location. You may choose to associate the Data Center annotation with other assets in your environment as well.

Steps:

- Log in to Insight as a user with administrative permissions.
- Select **Manage > Annotations**.
- Choose the **Data Center** annotation and click on the **Edit** icon.
- Click **+Add** and add the Name and Description of your first data center to the annotation list.
- Do the same for your other data centers.
- When finished, click **Save**.

Example Data Center annotations:

Name	Description
DC1_SVL	Sunnyvale Bldg 1
DC2_SVLb3	SVL Bldg3 ENG
DC3_NY	New York

DC4_London	London
...	

Insight comes with several out-of-the-box annotation types that allow users to define or modify values to fit their needs. These default annotation types will always be available to the Insight web UI as well as reporting. Newly-created custom annotations are visible in the Insight web UI but require additional measures to make them available to reporting. For information on including custom annotations in reports, contact your NetApp Customer Support representative.



Some users may be inclined to use the Country annotation to set asset locations, as opposed to, or in conjunction with, the Data Center annotation. However, you should bear in mind that the Country annotation is treated as a custom annotation type in the Insight data warehouse and therefore may not show up in reporting at the same granularity as Data Center.

Configure Tier annotations

The Tier annotation is used to associate assets with their respective Tiers, for use in cost accounting, for example. Insight comes with a number of default Tier annotations; you may modify these according to your tiering naming conventions or create your own Tiers as needed.

When setting Tier annotations, keep the following in mind:

- Cost is per gigabyte.
- Tier 1, 2, 3 are default tiers configured at a storage array level, by disk type. However, many customers will have multiple disk types within an array, or across arrays of the same type.
- Best practice is to create Tier annotations based on disk type and/or disk speed. This is a typical Tier methodology; your own needs may vary.

Steps:

- Choose the **Tier** annotation and click on the **Edit** icon.
- If desired, click **+Add** and add the Name and Description of your first tier to the annotation list.
- Do the same for your other tiers.
- When finished, click **Save**.

Example Tier annotations:

Name	Description	Cost per Gb
Auto Tier	Automatic Storage Tiering Tier	0.5
Tier 1 SSD	All Flash Array	0.5
Tier 2 SAS	SAS	0.25

Tier 3 SATA	SATA	0.1
...		

Configure Service Level annotations

The Service Level annotation is used to associate assets with their respective service levels.

Service Level annotations are typically only set in customer environments that use auto-tiering. In the Insight data warehouse, Tier is preferred. However, best practice is to use Service Level when you want to detail Provisioned Cost vs. Customer Cost. When both are present in Data Warehouse, Service Level will supersede Tier.

Steps:

- Choose the **Service Level** annotation and click on the **Edit** icon.
- Click **+Add** and add the Name and Description of your first service level to the annotation list.
- Do the same for your other service levels.
- When finished, click **Save**.

Example Service Level annotations:

Name	Description	Cost per Gb
Service Level 1	FAS controllers with FC or SAS, local & remote mirror and tape	0.93
Service Level 2	FAS controllers with FC or SAS, local & remote mirrors	0.85
Service Level 3	FAS controllers with SATA and local mirror	0.48
...		

Configure custom Environment annotations

The Environment annotation is a custom annotation for associating assets with their respective environmental location or use, for example, Lab, R&D, Production, etc. By creating the Environment annotation and setting it onto these assets, you can easily find, filter and report on your Lab assets separately from your Production assets, for example.

Steps:

- Select **Manage > Annotations**.
- Click the **+Add** button at the top of the page.

- For **Name**, enter '**Environment**'.
- For **Description**, enter '**Asset environment type**'.
- For **Type**, select **List**. New fields are displayed for you to create your list.
- For now, leave **Add new assets on the fly** unchecked. You will check this on if you want to be able to add new Environments to the list of choices at the same time you are associating them with assets.
- Enter the Name and Description of your first environment.
- Click **+Add** and do the same for your other environments.
- When finished, click **Save**.

Example Environment annotations:

Name	Description
Lab	Lab
Dev	Development
Prd	Production
...	

Finding assets: Querying

You can easily find and display assets in your environment using powerful queries.

Using queries to annotate your assets

Now that you have created your initial annotations, let's take a look at how to associate those annotations with specific assets.


In the examples that follow, we will apply these annotations to specific assets. For example, we will create a query to list all the storage arrays that reside in a specific data center, and mark those with the appropriate annotation. Then we'll do the same for assets belonging to a specific tier and service level.

Querying and annotating data centers

You use queries to associate your annotations with the appropriate assets in your environment. In this example, we will associate Data Center annotations with selected assets.

During data source acquisition, Insight gathers (among other information) the names of each asset it discovers. For this example, we will assume that all of your storage arrays have been named according to the data center in which they reside, such as "SVL_NN_<label>" for arrays residing in Sunnyvale. Insight queries make annotating these assets simple.

- Log in to Insight as a user with administrative permissions
- Select **Queries** > **+New Query**

- Drop down the **Search for...** field and select **Storage**. A list of all of your storage arrays is displayed.
- In the **Name** filter field, type “SVL” and click the  button (or press Enter). The query results list is now updated to show only those arrays that contain the string “SVL”.
- When filtering, you can use any of the following characters alone or combined to refine your search in any text box on the Query page:
 - An asterisk enables you to search for everything. For example, “vol*rhel” displays assets that start with “vol” and end with “rhel”.
 - The question mark enables you to search for a specific number of characters. For example, filtering for “SVL-PRD??-S12” displays SVL-PRD12-S12, SVL-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.
- Select the storage arrays you wish to associate with this data center. When all of your desired arrays are selected, click the **Actions** button and select **Edit annotation**.
- In the **Add Annotation** dialog, select the **Data Center** annotation.
- Choose the desired **Value**, for example, “DC1_SVL”.
- Click **Save**.
- If the Data Center column is not visible on the Query results page, select it by dropping down the **Columns** button and choosing **Data Center**.
- If desired, you can save the query for future use by clicking the **Save** button in the upper right corner of the Query page and giving it a unique and explicit name. For example, “Storage Arrays - SVL data center”.

If you wish to associate the “SVL” annotation with other assets, create a new query and follow these steps for each asset type you desire.

Repeat these steps for assets in each of your data centers.

Querying and annotating tiers


You use queries to associate your annotations with the appropriate assets in your environment. Here, we will associate those tiers with the appropriate assets.

Earlier, we set up annotations for your Tiers. For this example, we will associate tiers with storage pools, and will assume your tier annotations are configured as follows:

Value	Description	Cost per Gb
Tier 1 SSD	All Flash Array	0.5
Tier 2 SAS	SAS	0.25
Tier 3 SATA	SATA	0.1

Let’s search for all SSD disks in your environment, and associate the “Tier 1 SSD” annotation with them.

- Log in to Insight as a user with administrative permissions
- Select **Queries > +New Query**

- Drop down the **Search for...** field and select **Storage Pool**. A list of all of your storage pools is displayed.
- The **Name** field may not be helpful this time, so let's use another field. Click the **More** drop-down and select "Least performing disk type". This field lists the disk types that we are interested in. Enter "SSD" in the field and click the  button. The query results list shows only your SSD storage pools.
- You may filter further by clicking the **More** drop-down and selecting additional fields.
- Select the storage pools you wish to associate with this tier. When all of your desired storage pools are selected, click the **Actions** button and select **Edit annotation**.
- In the **Add Annotation** dialog, select the **Tier** annotation.
- Choose the desired **Value** from the list. For this example, choose "Tier 1 SSD".
- Click **Save**.
- If the Tier column is not visible on the Query results page, select it by dropping down the **Columns** button and choosing **Tier**. You should see the appropriate annotation now associated with your assets.
- Save the query by clicking the **Save** button in the upper right corner of the Query page and giving it a unique and explicit name. For example, "Storage Pools - Tier 1 SSD".

If you wish to associate the "Tier 1 SSD" annotation with other assets, create a new query and follow these steps for each asset type you desire.

Repeat these steps for the assets in each of your remaining tiers.

Service Level and Environment annotations

Add Service Level and Environment annotations to the appropriate assets using the steps and concepts you've learned.

To add Service Level and Environment annotations to the appropriate assets in your environment, follow the steps noted above, choosing the desired assets and appropriate Service Level or Environment annotations. You can have multiple annotations associated with the same assets, and in fact this practice will allow you greater flexibility in managing your environment through Insight.

Now that you have created queries to annotate your assets, you can use these annotations in many different ways, such as:

- Performance policies to alert you when events occur on desired assets
- Custom dashboards and widgets to monitor activity
- Reporting

Your corporate structure: setting up business entities and applications

Understanding the elements of your corporate structure helps you to keep track of asset usage and report on costs.

Configuring business entities for your company

Understanding the business elements of your corporate structure helps you to keep track of asset usage and report on costs. Here we will configure your company business entities.

About this task

OnCommand Insight allows you to define business entities in a hierarchy that includes up to four levels of granularity.

- **Tenant**

Primarily used by service providers to associate resources with a customer. Tenant level is needed if your company is an ISP and you want to track customer usage of resources.

- **Line of Business (LOB)**

A line of business or product line within a company, for example, Data Storage. Line of Business is needed in the hierarchy if the data for different product lines needs to be tracked.

- **Business Unit**

Represents a traditional business unit such as Legal or Marketing. Business Unit is required if you need to track data for different departments. This level of the hierarchy is often valuable in separating a resource that one department uses that other departments do not.

- **Project**

Often used to identify a specific project within a business unit for which you want capacity chargeback. For example, “Patents” might be a project name for the Legal business unit and “Sales Events” might be a project name for the Marketing business unit. Note that level names may include spaces.

An example of a business entity hierarchy would be:



Best practice: Create a table with each row showing one full business entity in your hierarchy:

Tenant	Line of Business	Business Unit	Project
NetApp Inc.	Data Storage	Legal	Patents
NetApp Inc.	Data storage	Marketing	Sales Events

N/A	N/A	Safety and Security	N/A
...			



You are not required to use all of the levels in the design of your corporate hierarchy. You can choose “N/A” for levels that you do not use.

To create a business entity hierarchy in Insight:

Steps

1. Log in to Insight as a user with administrative permissions.
2. Select **Manage > Business entities**.
3. Click the **+Add** button
4. Click in the **Tenant** box and type your tenant name.

If you have already entered tenants for your environment, a list of existing tenants will appear from which you may choose. You may also choose N/A if tenant does not apply for this business entity.

5. Repeat for **Line of Business**, **Business Unit**, and **Project**.
6. Click **Save**

After you finish

Best Practices:

- Map out your business hierarchy in a table, and check that the names in the hierarchy will be self-explanatory in Insight views and reports.
- Create your business entities in Insight before creating applications.
- Identify and list all applications that will be associated with each business entity.

Configuring applications for your company

Understanding the applications used in your company’s environment helps you to keep track of asset usage and report on costs. Here we will configure your company’s applications and associate them with the appropriate assets.

About this task

In the *Configuring business entities for your company* section, we created some business entities, and recommended that you list out all the applications you associate with each business entity. OnCommand Insight allows us to then track data associated with those applications for things like usage or cost reporting.

Before you can track data associated with the applications running in your environment, you must first define those applications and associate them with the appropriate assets. You can associate applications with the following assets: hosts, virtual machines, volumes, internal volumes, qtrees, shares, and hypervisors.

In this walkthrough, we want to track the usage of virtual machines that the Marketing Team uses for its Exchange email. You will remember the following table we created while defining our business entities. Let’s

add a column to this worksheet listing the applications used by each business entity. (This table is a worksheet example only. You will not see an “Applications” column in the business entities table in Insight.)

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
N/A	N/A	Safety and Security	N/A	N/A
...				

Creating applications in Insight:

Steps

1. Log in to Insight as a user with administrative permissions.
2. Select **Manage > Applications**
3. Click the **+Add** button
4. Enter the name of the application (for our example, enter “Exchange”)
5. Select a priority for the application
6. If you wish to associate the application to a business entity, select one from the **Business Entity** drop-down. Otherwise, you can leave this as “None”.
7. If you want to ensure each host has access to the same volumes in a cluster, ensure that the **Validate volume sharing** box is checked. For example, hosts in high-availability clusters often need to be masked to the same volumes to allow for failover; however, hosts in unrelated applications usually have no need to access the same physical volumes. Additionally, regulatory policies might require you to explicitly disallow unrelated applications from accessing the same physical volumes for security reasons. If you do not use volume sharing, clear the **Validate volume sharing** box. This requires the Assure license.
8. Click Save.
9. Repeat for all other Applications in your environment.

After you finish

We see that the Marketing Team uses the Exchange application. We want to see their virtual machine utilization for Exchange, in order to predict when we will need to add more storage. Let’s associate the Exchange application with all of Marketing’s VM’s. The easiest way to accomplish this is through a query.

By following these steps, you can associate each of your applications with their appropriate assets.

Associating Applications to assets:

Now that you have created your applications (and tied them to business entities, as desired), we can now

associate those applications with assets in your environment. In this example, we will associate the Exchange application with a number of virtual machines in your company. The easiest way to do this is what a query.

1. Select **Queries** > **+New query**.
2. In the **Select Resource Type** drop-down, choose *Virtual Machine*
3. We will assume the Marketing team names their assets with the string “_mktg_”. In the Name filter box, enter “_mktg_”(without quotes) and click the apply (checkmark) button.
4. The list of all VM’s with the “_mktg_”string is shown.
5. If desired, click the **More** drop-down and add additional filters.
6. Select the VM’s used for Exchange by clicking in the checkbox beside each desired VM Name, or select all VMs by clicking the checkbox at the top of the column.
7. When the desired VM’s are selected, click the **Actions** button and choose **Add Application**.
8. In the Assign Application dialog, click the **Application** drop-down and select “Exchange”.
9. Click **Save**.
10. Repeat as necessary to associate the Exchange application with other assets (hosts, volumes, etc.)

Creating performance policies for alerting

Performance Policies allow you to monitor and send alerts when specific conditions are met.

About this task

Now that we have annotated our assets, let’s create a Performance Policy that we can use to alert us when latency is higher than 2ms in any storage array at our Sunnyvale (DC1_SVL) data center. When these conditions occur, we will send an email to selected recipients.

Steps

1. Select **Manage** > **Performance policies**.

The Performance policies page opens. There are several default policies already set, which you can modify to suit your needs. We will create a new policy, however.

2. Click the **+Add** button.

The **Add Policy** dialog opens.

3. In the **Policy name** field, enter “SVL Data Center Latency policy”.

You must use a name that is different from all the other policy names for the object. For example, you cannot have two policies named "Latency" for an internal volume; however, you can have a "Latency" policy for an internal volume and another "Latency" policy for a different volume. The best practice is to always use a unique name for any policy, regardless of the object type.

4. For **Apply to objects of type**, select **Storage**.
5. In the **With annotation** field, select **Data Center** is “DC1_SVL” (or choose the name of your desired data center here).
6. Apply after a window of **First occurrence** .

The First occurrence option triggers an alert when a threshold is exceeded on the first sample of data. All other options trigger an alert when the threshold is crossed once and is continuously crossed for at least the specified amount of time.

7. From the **With severity** list, select **Warning**.
8. Under **Email recipients**, click to override the global recipient list. Click **+** to add the email address of your first desired alert recipient, and click **OK**. Repeat for any additional desired email recipients.
9. Leave the default choice to Create alert if **any** of the following are true. This will send an alert if any one of the set thresholds is met. You can also choose to alert only if **all** of the set thresholds are met.
10. To set your first threshold, select **Latency - Total** in the drop-down and set it to greater than 2 ms.
11. If desired, add additional thresholds to alert on by clicking the **Add threshold** button. When the policy is customized the way you want it, click **Save**.
12. You can also choose to **Stop processing further policies if alert is generated**. This will halt additional policy alerting if this policy's conditions are met.
13. You may add as many new policies as you like, setting alerts for other recipients based on different conditions, according to business need. Any policies configured without specific recipients will send alerts to the global recipient list set in the **Admin > Notifications** page

After you finish

Each new policy is automatically activated when it is saved, and recipients will start receiving alerts when the policy's conditions are met (known as a *violation*). You can also monitor these violations in the **Dashboards > Violations Dashboard**.

Highlighting data using dashboards

Now that you have your assets annotated and have configured performance policies to alert for violations, you can create dashboards to highlight specific data that you want to target.

About this task

In this example we will provide a high-level view of dashboard creation by creating a dashboard with a single widget highlighting VM Performance data. You can add as many widgets as you need on a single dashboard, and you can create as many dashboards as you need. Widgets can be resized and moved as desired.

More information on Dashboards and Widgets can be found throughout the OnCommand Insight documentation.

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. Best practice: Name and save your dashboard as soon as you create it. Click the **Save** button and enter a unique name for the dashboard in the **Name** field. For example "VM Performance Dashboard". Click **Save**.
4. If necessary, slide the **Edit** switch to "On" to enable Edit mode. This allows you to begin adding widgets to your dashboard.

5. Click the **+Widgets** button and select **Table** to add a new table widget to the dashboard.

The Edit Widget dialog opens.

6. In the Name field, delete "Widget 1" and enter "Virtual Machine Performance table".
7. Click the asset type drop-down and change **Storage** to **Virtual Machine**.

The table data changes to show all virtual machines in your environment.

8. To add additional columns to the table, click the "Columns"  button and select the desired columns, for example *Data Center*, *Storage name*, and *Tier*. You can sort the table by any of these columns.
9. You can set filters as needed to highlight the data that is important to you for this dashboard, for example, you might choose to only show only virtual machines with "Tier 1 - SSD" annotation. Click the " + " button next to **Filter by** and select *Tier*. Click on **Any** and enter "Tier 1 - SSD". Click the  button to save the filter.

The table now shows only virtual machines in the "SSD" tier.

10. You can group results by clicking the " + " button next to **Group by** and selecting a field to group by, such as *Data Center*. Grouping is automatically applied to the table.
11. When you have customized the widget to your satisfaction, click the **Save** button.

The table widget is saved to the dashboard.

12. You can resize the widget on the dashboard by dragging the lower-right corner.
13. To add more widgets, click the **+Widget** button. Each widget is added to the dashboard when it is saved.
14. When you have made all the changes desired, be sure to click **Save** to save the dashboard.
15. You can create additional dashboards to highlight different data.

Creating custom dashboards

OnCommand Insight 7.3 includes enhanced custom dashboard capabilities to give users an operational view of the data that is important to them and provide a single-stop view of that data.

OnCommand Insight provides users the flexibility to create operational views infrastructure data across IT platforms, by allowing you to create custom dashboards with a variety of widgets, each of which provides extensive flexibility in displaying and charting your data. In this How-To, we will create an example dashboard to highlight VM Performance.

This How-To should serve as an example only and does not cover every possible scenario. The concepts and steps herein can be used to create your own custom dashboards to highlight the data specific to your particular needs.

Overview

You create a custom dashboard by either of the following methods:

- **Dashboards > +New dashboard**

- **Dashboards > Show all dashboards** and click **+Add**

The New Dashboard screen has several controls:

- **Time selector**: allows you to view dashboard data for a range of time from 3 hours up to 90 days by use of the custom date range selector. You can choose to override this global time range in individual widgets.
- **Edit** button: Selecting “On” will enable Edit mode, which allows you to make changes to the dashboard. New dashboards open in Edit mode by default.
- **Save** button: Allows you to save, rename or delete the dashboard.
- **Variable** button: Variables can be added to dashboards. Changing the variable updates all of your widgets at once. For more information on variables, see [Custom Dashboard concepts](#)
- **Widget** button, which allows you to add any number of tables, charts, or other widgets to the dashboard.

Widgets can be resized and relocated to different positions within the dashboard, to give you the best view of your data according to your current needs.

Widget types

You can choose from the following widget types:

Table widget: A table displaying data according to filters and columns you choose. Table data can be combined in groups that can be collapsed and expanded.

Line, Spline, Area, Stacked Area charts: These are time-series chart widgets on which you can display performance and other data over time.

Single value widget: A widget allowing you to display a single value that can be derived either directly from a counter or calculated using a query or expression. For example, you can display the sum of total IOPS for all storage in your environment as a single value at the top of your dashboard.

Bar chart: A chart to display top or bottom 5, 10, 20, or 50 values.

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 would quickly see assets with high latency and low IOPS.

Additionally, there are a number of legacy widgets that you can choose. In the **Widgets** drop-down, select **Show more...** to see these widgets.

Custom Dashboard concepts

Custom 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. Each concept is explained in greater detail in the following sections.

Variables

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

Multiple queries and/or expressions

Each time series widget (line, spline, area, or stacked area charts) can have up to five queries and/or expressions to determine what data to display, allowing you to compare different sets of data on a single chart. For example, you can have a line chart showing IOPS for both Storage and VM's, or a single chart comparing Throughput and Latency for all Storage Pools.

Rollup and Grouping

Data displayed in each widget is rolled up from the data points collected. You can choose to roll up this data in one of several ways:

- Avg: rolls up data as the average of the underlying data
- Max: rolls up data to the maximum of the underlying data
- Min: rolls up data to the minimum of the underlying data
- Sum: rolls up data as the sum of the underlying data

By default, all of the underlying data is rolled up into a single entry (All) on the chart or table. You can choose to roll up data for a specific attribute instead, such as Data Center or Tier, to distribute the underlying data into desired groups. Your widget will display data for only the attributes you select.

You can group data in a table widget according to the attribute you choose. For example, you might choose to group your table by Data Center. Groups can be expanded or collapsed at will. Performance data in a table is rolled up in the group header according to the rollup method you set in the widget (average, max, min, or sum).

Table widgets can be sorted by any column, and columns can be moved or resized as needed.

Top / Bottom

Use this to limit the result set in chart widgets, to select whether to display the top N results in your widget, or the bottom N results. You can choose this option when data is either not rolled up or is rolled up by a specific attribute.

Override dashboard time

By default, most widgets you add to a dashboard show data according to the dashboard's time range setting (3h, 24h, 3d, 7d, 30d or custom range). However, you can override this time setting in individual widgets to force them to show data in a specific time context, regardless of the dashboard's time setting.

These concepts are explained in greater detail in the following section.

Dashboard variables

Dashboard variables allow you to filter data across multiple widgets on a dashboard quickly and easily.

Before you begin

This example requires the **City** annotation (also called City attribute) to be set on multiple storage assets.

For best results, set different cities on different storages.

About this task

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

Steps

1. Log in to Insight as a user with administrative permissions
2. Click on **Dashboards > +New Dashboard**.
3. Before adding widgets, we want to define the variables we will use to filter the dashboard data. Click on the **Variable** button.

The list of attributes is displayed.

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

The \$city variable field is created and added to the dashboard.


5. Next, we must tell our widgets to use this variable. The simplest way to illustrate this is to add a table widget showing the City column. Click on the **Widget** button and select the **Table** widget.

6. First, add the City field to the table by selecting it from the column picker  button.

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

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

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



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

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

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

More on dashboard variables

Dashboard variables come in several types, can be used across different fields, and must follow rules for naming. These concepts are explained here.

Variable types

A variable can be one the following types:

Text: Alphanumeric string. This is the default variable type.

Numerical: a number or range of numbers.

Boolean: Use for fields with values of True/False, Yes/No, 0/1, etc. For the boolean variable, the choices are *Yes*, *No*, *None*, *Any*.

Date: A date or range of dates.

“Generic” variables

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

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

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

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



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

Variable naming

Variables names:

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

Widgets that use variables

Variables can be used with the following widgets:

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

Displaying widget legends

Widgets in dashboards can be displayed with or without legends.

Legends in widgets can be turned on or off on a dashboard by either of two methods:

1. When creating or editing the widget itself, check the Legends checkbox and save the widget.
2. With the dashboard in Edit mode, click the Options button on the widget and check the Legends checkbox in the menu.

As you edit and change the data displayed in the widget, 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.

Dashboard widget queries and filters

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

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

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

The following widget types can have multiple queries:

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

The remaining widget types can have only a single query:

- Table
- Bar chart
- Box plot
- Scatter plot

Filtering in dashboard queries

You can filter using any of the following to refine your search in any **text field** in the query:

- 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.

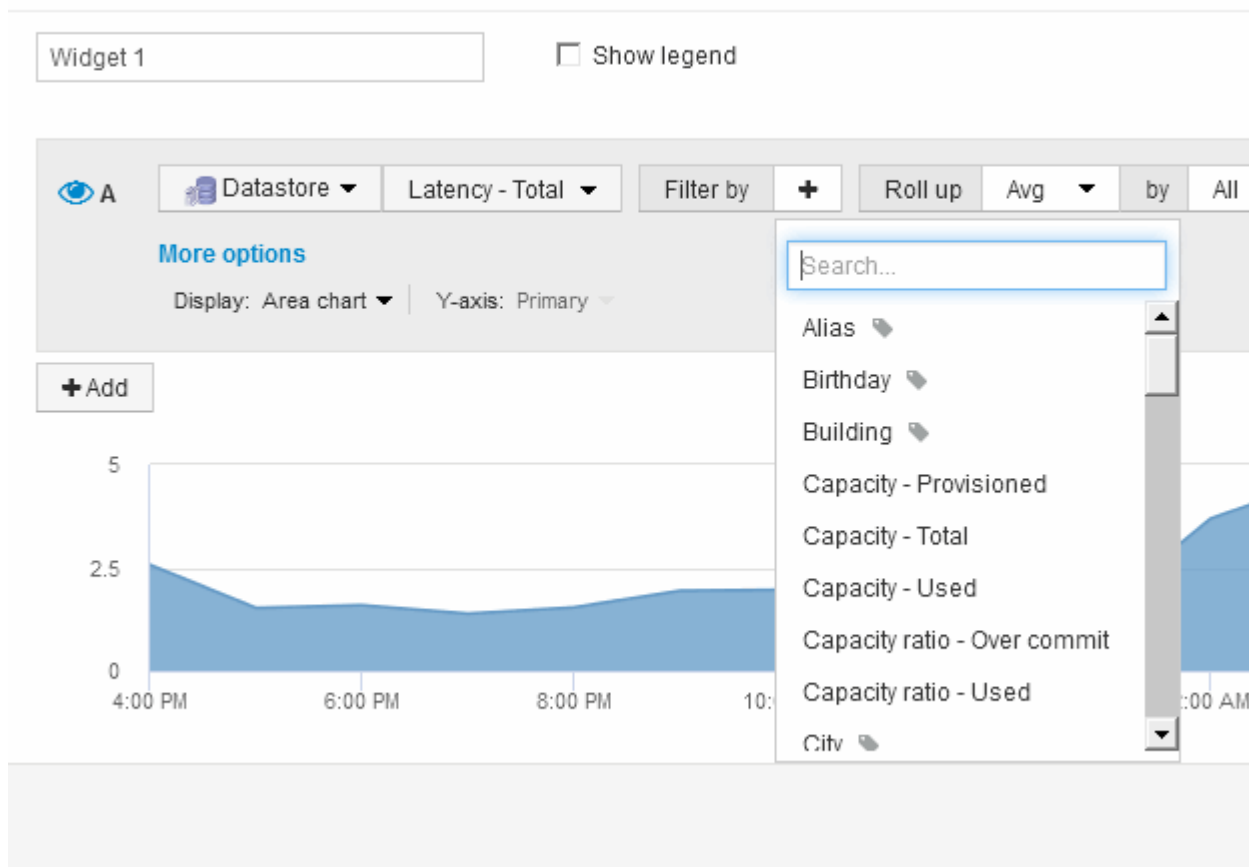
- The OR operator enables you to specify multiple entities. For example, `FAS2240 OR CX600 OR FAS3270` finds multiple storage models.
- The NOT operator allows you to exclude text from the search results. For example, `NOT EMC*` finds everything that does not start with “EMC”. You can use `NOT *` to display fields that contain null values.

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.

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.

Edit widget



Roll up and Aggregation

Data displayed in dashboard widgets is rolled up from acquired data points, allowing flexibility and conciseness in your dashboards.

Data displayed in each widget is 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 roll up this data in one of several ways:

- **Avg**: displays each line as the *average* of the underlying data.
- **Max**: displays each line as the *maximum* of the underlying data.
- **Min**: displays each line as the *minimum* of the underlying data.
- **Sum**: displays each line as the *sum* of the underlying data.

To do this, in your widget's query, first choose an asset type (for example, *Storage*) and metric (such as *IOPS - Total*). For **Roll up**, choose a rollup method (such as *Avg*) and select an attribute or annotation by which to roll up the data (for example, *Data Center*). The widget updates automatically and shows a line for each of your data centers.

You can also choose to roll up *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 or sum of the chosen metric for all of the underlying assets.

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

Note that when you choose to roll up 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 properly on the chart or table.

Aggregating data: You can further align your time-series charts (line, area, etc.) by aggregating data points into minute, hour, or day buckets before that data is subsequently rolled up by attribute (if chosen). You can choose to aggregate data points according to their Avg, Max, Min, or Sum, or by the Last data point collected during the chosen interval. To choose an aggregation method, click on **More options** in the widget's query section.

The minimum allowed interval is ten minutes. A small interval combined with a long time range may result in a "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 to 1 hour until you select a smaller time frame.

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

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

Showing top/bottom results in dashboard widgets

In a chart widget on a custom dashboard, you can show either the Top or Bottom results for rolled up data, and choose the number of results shown. In a table widget, you can select the number of rows displayed and 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 can choose from 5, 10,

20, or 50 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.

Note that the more results you choose to display, the taller your widget will be when you save it to the dashboard. You will not be able to resize the widget smaller than the number of rows displayed.

Grouping in table widgets

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.

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* of the underlying data in the group row. You can also choose to display the *sum*, *minimum*, or *maximum* of the data..


Grouping example (with rollup explained)

Table widgets allow you to group data for easier display.

About this task

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. Click the check button 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.

Because **All** is selected as the grouping method by default, all VMs are moved into a single group named “All”.

- Above the *IOPS - Total* column is now a **Roll up** option. The default roll up method is *Avg*. This means that the number shown for the group is the average of all the Total IOPS reported for each VM inside the group. You can choose to roll this column up by *Avg*, *Sum*, *Min* or *Max*. Each column that you display that contains performance metrics can be rolled up individually.
- 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.

Table - Grouping Example

Override dashboard time: OFF

Virtual Machine Filter by: IOPS - Total (IOPS) >= 5 Group By: Hypervisor name

Show entries: 5

Hypervisor name		Name	Hypervisor name	IOPS - Total (IOPS)
hv-72-001.nane.neta... (3)			hv-72-001.nane.neta...	8.88
hv-72-002.nane.neta... (4)			hv-72-002.nane.neta...	12.34
hv-72-002.nane.neta...		vsa-5-vc	hv-72-002.nane.neta...	14.77
hv-72-002.nane.neta...		ns5	hv-72-002.nane.neta...	7.01
hv-72-002.nane.neta...		ns6	hv-72-002.nane.neta...	6.94

37 items found in 35 groups

Cancel Save

- Click **Save** to save the table to the dashboard. You can resize the widget.
- Click **Save** to save the dashboard.

Overriding dashboard time for individual widgets

You can override the main dashboard time frame 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 a its own time frame, in the widget's edit mode set the **Override dashboard time** to **On**, and select a time frame for the widget. **Save** the widget to the dashboard.

The widget will display its data according to the time frame set, 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 explained

The secondary axis makes it easier to view data from two different sets of values that use different units of measurement.

About this task

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

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.

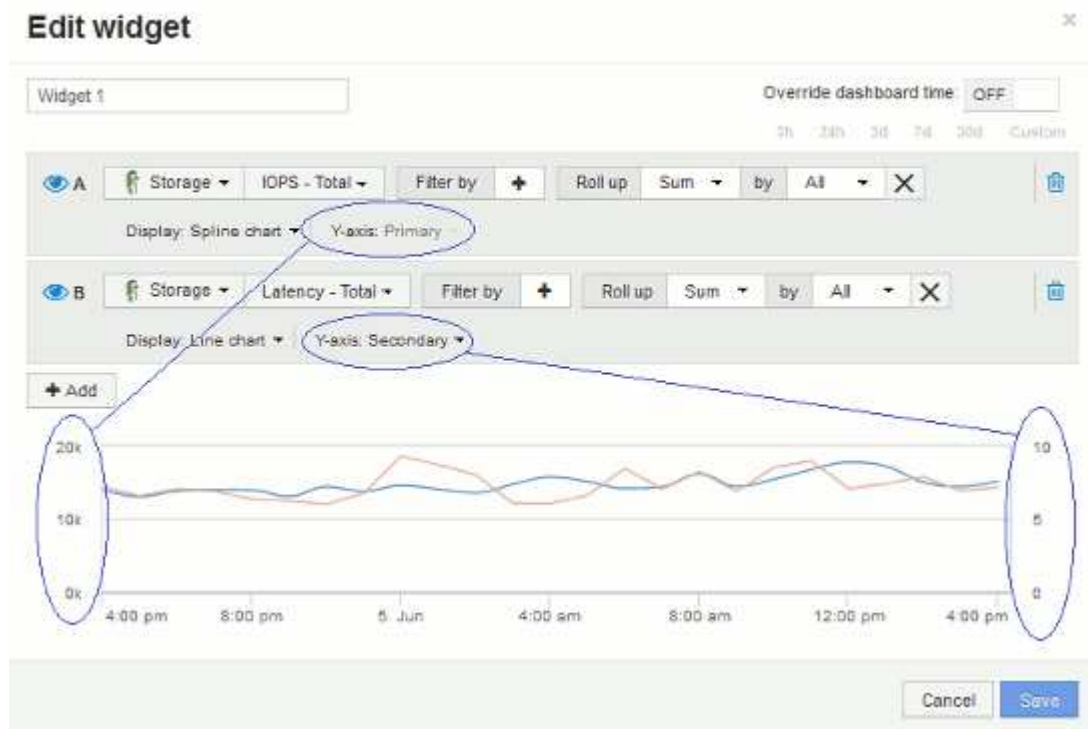
The IOPS line is displayed on the chart, with its scale shown on the left.

3. Click **+Add** 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 dashboard widgets

Expressions in time series widgets allow you to display data based on calculations with metrics of your choosing.

In a dashboard, any time series widget (line, spline, area, stacked area) allows you to build expressions from metrics you choose, and show the result of those expressions in a single graph. The following examples use expressions to solve specific problems. In the first example, we want to show Read IOPS as a percentage of Total IOPS for all storage assets in our environment. The second example gives us visibility into the "system" or "overhead" IOPS that occur in our environment—those IOPS that are not from reading or writing data.

Expressions example: Read IOPS percentage

Using expressions, you can view metrics by alternate means, such as percentage of total.

About this task

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

- Read Percentage = (Read IOPS / Total IOPS) x 100

This data can be shown in a line graph on your dashboard. To do this, follow these steps:

Steps

1. Create a new dashboard, or open an existing dashboard in **edit mode**.
2. Add a widget to the dashboard. Choose **Area chart**.

The widget opens in edit mode. By default, a query is displayed showing **IOPS - Total** for **Storage** assets. If desired, select a different asset type.

3. Click the **Convert to Expression** button.

The current query is converted to Expression mode. Notice that you cannot change the asset type while in Expression mode. Notice that while you are in Expression mode, the button changes to **Revert to Query**. Click this if you wish to switch back to Query mode at any time. Be aware that switching between modes will reset fields to their defaults.

For now, stay in **Expression** mode.

4. The **IOPS - Total** metric is now in the alphabetic variable field "a". In the "b" variable field, click **Select** and choose **IOPS - Read**.

You can add up to a total of five alphabetic variables for your expression by clicking the **+** button following the variable fields. For our Read Percentage example, we only need Total IOPS ("a") and Read IOPS ("b").

5. In the **Expression** field, you use the letters corresponding to each variable to build your expression. We know that *Read Percentage = (Read IOPS / Total IOPS) x 100*, so we would write this expression as: (b / a) * 100
6. The **Label** field identifies the expression. Change the label to "Read Percentage", or something equally meaningful for you.
7. Change the **Units** field to "%" or "Percent".

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

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

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

Expressions example: "System" I/O

Expressions give you the freedom to chart data that can be calculated from other metrics.

About this task

Example 2: OnCommand Insight acquires many metrics from data sources. Among them are read, write, and total IOPS. However, the total number of IOPS reported by acquisition 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. Click the button to create a copy of the query.

A duplicate of the query is added below the original.

4. 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.

5. The **IOPS - Total** metric is now in the alphabetic variable field "a". Click on **IOPS - Total** and change it to **IOPS - Read**.
6. In the "b" variable field, click **Select** and choose **IOPS - Write**.
7. 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.
8. 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.

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

Custom Dashboard: Virtual Machine Performance

OnCommand Insight's custom dashboards and widgets help provide operational views into inventory and performance trends.

About this task

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 custom dashboard with widgets that give you operational insights into the virtual machine performance in your environment. By following this example, and creating widgets to target your own specific needs, you will be able to visualize backend storage performance compared to frontend virtual machine (VM) performance, or view VM latency versus I/O demand.

Custom dashboards allow you to prioritize efforts and identify resource availability. You can respond to the ebb and flow of workloads and minimize the time to detect and remediate emerging issues. Custom dashboards allow you the flexibility to create prioritized views into business-critical infrastructure, and are useful for identifying performance availability across multi-vendor technologies.

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 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. Let's give our dashboard a meaningful name. Click **Save**. In the **Name** field, 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, slide the **Edit** switch to "On" to enable Edit mode.
6. Click the **Widget** button and select **Table widget** to add a new table widget to the dashboard.

The Edit Widget dialog opens. The default name is "Widget 1" and the default data displayed is for all storages in your environment.

Edit widget ✕

Widget 5 Override dashboard time: ☐ OFF

3h 24h 3d 7d 30d Custom

Storage Filter by + Group By +

Show entries: 5


Name	Vendor
3070-a,3070-b	NetApp
APM000934007420000	EMC
Ds4800	NetApp
FNIM00142500950	EMC
Storage Center 6145...	Dell

18 items found


Cancel Save

7. We can customize this widget. In the Name field, delete "Widget 1" and enter "Virtual Machine Performance table".
8. Click the asset type drop-down and change **Storage** to **Virtual Machine**.

The table data changes to show all virtual machines in your environment. For now, the table only shows the VM names. Let's add a few columns to the table.

9. Click the *Columns*  button and select *Data Center*, *Storage name*, and *IOPS - Total*. You can also try typing the name into the search to quickly display the desired field(s).


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.

10. 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 (IO/s)*. Click on **Any** and enter "10" in the **from** field. Leave the **to** field empty. Click the  button to save the filter.

The table now shows only VMs with 10 or more total IOPS.

11. 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 Cluster. 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.



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

The table widget is saved to the dashboard.

13. 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.

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

15. If necessary, slide the **Edit** switch to "On" to enable Edit mode.


16. Click the **Widget** button and select **Line Chart** to add a new line chart widget to the dashboard.

The Edit Widget dialog opens. Click the **Name** field and name this widget "VM vs Storage Max Latency"

17. 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**.

18. Click the **+Add** button to add a second data line. For this line, select **Storage** 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**.

Edit widget



19. Click **Save** to add this widget to the dashboard.
20. Next we will add a chart showing VM Read, Write and Total IOPS in a single chart.
21. Click the **Widget** button and select **Area Chart** to add a new area chart widget to the dashboard.

The Edit Widget dialog opens. Click the **Name** field and name this widget "VM IOPS"

22. 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 a**Area Chart**, and leave Y-Axis as **Primary**.
23. Click the +Add button to add a second data line. For this line, select **Virtual Machine** and choose **IOPS - Read**. Leave Y-Axis as **Primary**.
24. Click the +Add button to add a third data line. For this line, select **Virtual Machine** and choose **IOPS - Write**. Leave Y-Axis as **Primary**.

Edit widget



25. Click **Save** to add this widget to the dashboard.


26. 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.

27. Click the **Widget** button and select **Line Chart** to add a new line chart widget to the dashboard.

The Edit Widget dialog opens. Click the **Name** field and name this widget “VM Throughput by Application”

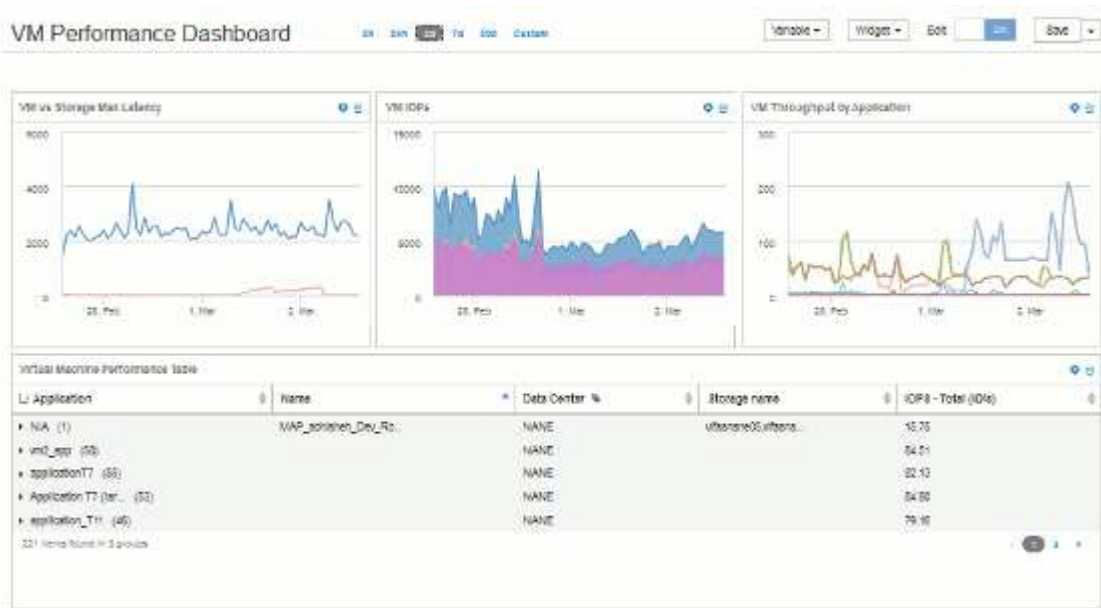
28. Select **Virtual Machine** and choose **Throughput - Total**. Set any filters you wish, or leave **Filter by** empty. For **Roll up**, choose “Max” and select by “Application” or “Name”. Show the **Top 10** applications. Display this data as a **Line Chart**, and leave Y-Axis as **Primary**.

Edit widget



29. Click **Save** to add this widget to the dashboard.
30. You can move widgets by holding down the mouse button anywhere in the top of the widget and dragging 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 like this:



Example storage node utilization dashboard with variables

Create a custom dashboard for Storage Analysis which has variables for storage, storage pool, node, tier, utilization and latency.

Before you begin

Familiarity with dashboards in Insight is recommended but not required.

About this task

The following procedure will create a custom Storage Analysis Overview dashboard which uses variables for storage, storage pool, node, tier, utilization and latency. Variables in the example below will be used to filter the displayed assets or metrics across one or many widgets available on the dashboard. The widgets that use these variables as filters will be updated with filtered content on-demand according to the values entered in the variable fields on the dashboard, allowing you to quickly filter multiple charts and graphs to drill down to a specific area of interest.

By following the steps in this example, you will create a dashboard like the following. You can change these widgets, or add any number of additional widgets, to highlight any data you choose.



Steps

1. Create a new dashboard, and name it "Analysis: Storage Overview", or something equally descriptive.
2. Click on the **Variable** drop-down and select **Text** variable type. By default, the variable is named `$var1`. Click on `$var1` to edit the name, and change it to `$storage`, then click the check mark to save the variable. Repeat to create text variables for `$node`, `$pool`, and `$volume`.
3. Repeat the above process to create **Number**-type variables named `$utilization` and `$latency`.
4. Click on the **Variable** drop-down and search for the `Tier` annotation. Select it to create a variable named `$tier`.


You can add variables at any time, however it is easier to create them up front and therefore make them available to all widgets as you create them.

5. Add a widget by clicking on the **Widget** drop-down and selecting either a **line chart** or **area chart** widget. Name the widget “Node Utilization”. Click on the **Storage** asset type and change it to **Storage Node**. Select **Utilization - Total** for the chart data.
6. Click on the **Filter by +** button to add a filter. Search for and select **Storage**, then click on **Any** and select the *\$storage* variable.
7. Click the **+** button to add another filter for **Name**. Set the variable to *\$node*.

Different variables can be assigned to the annotation name filter. Use the name/variable pair at the lowest level depending on the object in the widget. For example:

- You can assign the *\$node* variable to the **Name** filter for a Node-focused widget.
 - You can assign the *\$pool* variable to the **Name** filter for a Pool-focused widget.
8. Click the **+** button to add another filter for **Utilization - Total (%)**. Set the variable to *>= \$utilization*.
 9. Click the **X** after the **Roll up** field to collapse the field.
 10. Select **Show Top 5** and click **Save** to save the widget and return to your Dashboard.

Your widget should look something like this:



11. Add another line or area chart widget to your dashboard. Select **Storage Node** as the asset type and **Latency - Total** as the metric to chart.
12. Click on the **Filter by +** button to add filters for **Storage: \$storage** and **Name: \$node**.
13. Add a filter for **Latency - Total** and select the *\$latency* variable.
14. Name the widget “Node Latency” and save it.
15. You can add supporting tables to show more details for the charts you created, for example, Max or Avg node utilization. Add a **Table widget** to the dashboard and select **Storage Node** as the asset type, and create filters for **Storage: \$storage**, **Name: \$node**, and **Utilization - Total: \$utilization**.
16. Add columns to the table for **Utilization - Max**, **Utilization - Total**, or any other desired columns.
17. Name the widget “Node Peak and Avg Utilization” and save it.

Edit widget

Node Peak and Avg Utilization

Override dashboard time:

OFF

3h
24h
3d
7d
30d
Custom

Storage Node

Filter by

Storage: \$storage

Name: \$node

Utilization - Total (%): >= \$utilization

+

Group by

+

Name	Utilization - Max (%)	Utilization - Total (%)
3070-a	76.79	21.57
3070-b	76.79	21.57
vfasnane01	54.83	18.55
vfasnane02	32.50	6.06
aurora3	29.27	12.88

53 items found

Cancel

Save

Node Peak and Avg Utilization

Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Storage Node

Filter by: Storage: \$storage × Name: \$node × Utilization - Total (%) >= \$utilization × +

Group by: +

Name	Utilization - Max (%)	Utilization - Total (%)
3070-a	76.79	21.57
3070-b	76.79	21.57
vifasname01	54.83	18.55
vifasname02	32.50	6.06
aurora3	29.27	12.88

53. *Hearts Found*

Cancel Save

- Repeat the steps to create a table for Node Latency, showing **Latency - Max**, **Latency - Total**, or other columns as desired.
- To round out your dashboard, you might add additional chart and table widgets for some or all of the following:

Chart	Table
Storage Pool Utilization	Storage Pool Peak and Avg Utilization
Storage Pool Throughput	Storage Pool Peak and Avg Throughput
Volume Latency	Volume Peak and Avg Latency
Volume IOPS	Volume Peak and Avg IOPS


1. You can move and resize the widgets in whatever placement you want on your dashboard. When finished, be sure to **Save** the dashboard.

Your final dashboard will look something like this:



- You can use the variables to focus on specific assets in your dashboard. As you enter values into the variable fields, your widgets update automatically to reflect those variables. For example, by entering "15" in the \$utilization variable field, the widgets using that variable update to show only assets with total utilization $\geq 15\%$.

Node utilization widget showing top 5 of all nodes:



Node utilization widget showing nodes with 15% or greater utilization:



3. Keep in mind the following when creating your widgets:

- The \$tier variable will only impact resources that are annotated with the **Tier** annotation.
- Not all filters will impact all widgets, depending on whether the widget is designed to accept the variable(s) specified.
- Number variables are applied as “greater than or equal to” the value specified. Note that any variable can be used as a filter on any widget at any level in a storage hierarchy, as long as the variable is valid for the asset against which the widget is running. As you move down from a Node level to Storage Pool to a Volume widget, more variables are present for assignment as filters. For example, at a Storage Node level widget, the *Storage* and *Name* variables can be assigned as filters. At a Storage Pool level, *Storage*, *Nodes*, *Storage Pools* and *Name* are all available. Assign your variables as appropriate and use the \$name variable at the lowest level in the stack. Doing this will allow your \$name variable to filter on the actual name of the asset against which the widget is running.


Node dashboard example widget settings

Widget settings for the node dashboard with variables example.

Following are the settings for each of the widgets in the storage node dashboard example.

Node utilization:

Edit widget



Edit widget

Node Peak and Avg Utilization Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Storage Node

Filter by: **Storage: \$storage** **Name: \$node** **Utilization - Total (%): >= \$utilization**

Group by: **+**


Name	Utilization - Max (%)	Utilization - Total (%)
3070-a	76.79	21.57
3070-b	76.79	21.57
vifasane01	54.83	18.55
vifasane02	32.50	6.06
aurora3	29.27	12.88

53 items found

Cancel **Save**

Node latency:

Edit widget



Edit widget

Node Peak and Avg Latency

Override dashboard time: OFF

3h 3m 3d 7d 30d Custom

Storage Node

Filter by: Storage: \$storage Name: \$node Latency - Total (ms) >= \$latency

Group by: +


Name	Latency - Max (ms)	Latency - Total (ms)
vfasname04	9.05	7.70
vfasname05	2.25	0.41
vfasname02	1.62	0.90
vfasname01	1.42	1.03
vfasname06	0.97	0.64

8 items found

Cancel Save

Storage pool utilization:

Edit widget



Edit widget

Storage Pool Peak and Avg Utilization

Override dashboard time:

3h 24h 3d 7d 30d Custom

Storage Pool

Filter by: Storage: \$storage x Nodes: \$node x Name: \$pool x

Utilization - Total (%) >= Utilization x Tier: \$tier x +

Group by: +

Table view


Name	Utilization - Max (%)	Utilization - Total (%)
vfasname01:aggr1	15.85	8.52
vfasname01:vfasna...	16.19	4.71
vfasname02:aggr2	9.28	3.65
vfasname02:vfasna...	4.96	1.63
vfasname03:aggr3	1.04	0.68

14 items found

Cancel Save

Storage pool throughput:

Edit widget



Edit widget

Storage Pool Peak and Avg Throughput

Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Storage Pool

Filter by: Storage: \$storage Nodes: \$node Name: \$pool

Tier: \$tier

Group by: +


Name	Throughput - Max (MB/s)	Throughput - Total (MB/s)
vfasname01:aggr1	181.17	143.62
vfasname06:aggr1	127.19	26.75
vfasname05:aggr1	89.83	18.20
vfasname02:aggr2	24.57	9.70
vfasname05:aggr_opm1	14.61	4.75

14 items found

Cancel Save

Volume latency:

Edit widget



Edit widget

Volume Peak and Avg Latency

Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Volume

Filter by

Storage: \$storage

Nodes: \$node

Storage pools: \$pool

Name: \$volume

Latency - Total (ms): >= Latency

Tier: \$tier

Group by: +

Name	Latency - Max (ms)	Latency - Total (ms)
vfasname05/vol/bo...	0.00	0.00
vfasname05/vol/bo...	0.19	0.06
vfasname05/vol/bo...	0.00	0.00
vfasname05/vol/bo...	0.00	0.00
vfasname05/vol/bo...	0.00	0.00


51 items found

Cancel

Save

Volume IOPS:

Edit widget



Edit widget

Volume Peak and Avg IOPS

Override dashboard time: OFF

3h 24h 3d 7d 30d Custom

Volume

Filter by: Storage: Sstorage x Nodes: Snode x Storage pools: Spool x

Name: Svvolume x Tier: Sstier x +

Group by +

Name	IOPS - Max (IO/s)	IOPS - Total (IO/s)
vfasname05:/vol/vl...	1,889.31	198.97
vfasname05:/vol/vl...	50.03	19.18
vfasname05:/vol/bo...	1.51	1.11
vfasname05:/vol/bo...	0.00	0.00
vfasname06:/vol/bo...	0.00	0.00

31 items found

Cancel Save

Best Practices for Dashboards and Widgets

Tips and tricks to help you get the most out of the powerful features of dashboards and widgets.

Best Practice: finding the right metric

OnCommand Insight acquires counters and metrics using names that sometimes differ from data source to data source.

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 OnCommand Insight 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 it occurs. You can then easily select the metrics you want from that short 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 capacityRaw capacity Provisioned capacity Storage pools capacity <other asset type> capacity Written capacity
Disk Speed	Lowest disk speedLeast performing disk type
Host	HypervisorHosts
Hypervisor	HostIs hypervisor
Microcode	Firmware
Name	AliasHypervisor 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	VMIs virtual

This is not a comprehensive list. These are examples of possible search terms only.

Best Practice: finding the right assets

The Insight assets you can reference in widget filters and searches vary from asset type to asset type.

In dashboards, 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.

Scatterplot Example: knowing your axis

Changing the order of counters in a scatterplot widget changes the axes on which the data is displayed.

About this task

This example will create a scatter plot that will allow you to see under-performing VMs that have high latency compared to low IOPS.


Steps

1. Create or open a dashboard in edit mode and add a **Scatter Plot Chart** widget.
2. Select an asset type, for example, **Virtual Machine**.
3. Select the first counter you wish to plot. For this example, select **Latency - Total**.

Latency - Total is charted along the X-axis of the chart.

4. Select the second counter you wish to plot. For this example, select **IOPS - Total**.


IOPS - Total is charted along the Y-axis in the chart. VMs with higher latency display on the right side of the chart. Only the top 100 highest-latency VMs are displayed, because the **Top by X-axis** setting is current.



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

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

Note that because we haven't changed the **Top by X-Axis** setting, the widget now displays the top 100 highest-IOPS VMs, since this is what is currently plotted along the X-axis.



- You can choose for the chart to display the Top N by X-axis, Top N by Y-axis, Bottom N by X-axis, or Bottom N by Y-axis. In our final example, the chart is displaying the Top 100 VMs that have the highest *total IOPS*. If we change it to Top by Y-axis, the chart will once again display the top 100 VMs that have the highest *total latency*.

Note that in a scatterplot chart, you can click on a point to open the asset page for that resource.

Creating performance policies

You create performance policies to set thresholds that trigger alerts to notify you about issues related to the resources in your network. For example, you can create a performance policy to alert you when the total utilization for storage pools is greater than 60%.

Steps

- Open OnCommand Insight in your browser.
- Select **Manage > Performance Policies**.

The Performance Policies page is

Performance Policies

[Add new policy](#)

Database policies

Policy Name	Severity	Annotations	Time Window	Thresholds
Latency	Warning		First occurrence	'Latency - Total' > 200 ms
Database_0	Warning		First occurrence	IOPS - Total > 0 I/Os or 'Latency - Total' > 0 ms

Showing 1 of 2 entries

Internal volume policies

Policy Name	Severity	Annotations	Time Window	Thresholds
Atmos Service Level	Critical	Service_Level = Atmos	First occurrence	'Latency - Total' > 100 ms or IOPS - Total > 100 I/Os or 'Throughput - Total' > 200 MB/s
Global	Critical		First occurrence	'Latency - Total' > 200 ms or IOPS - Total > 1 I/Os or 'Throughput - Total' > 300 MB/s

Showing 1 of 2 entries

Storage policies

Policy Name	Severity	Annotations	Time Window	Thresholds
Storage_Storage	Warning		First occurrence	IOPS - Read > 10 I/Os
Storage_0	Warning		First occurrence	'Throughput - Total' > 0 MB/s or IOPS - Total > 0 I/Os

Showing 1 of 2 entries

displayed.

Policies are organized by object, and are evaluated in the order in which they appear in the list for that object.

3. Click **Add new policy**.

The Add Policy dialog box is displayed.

4. In the **Policy name** field, enter a name for the policy.

You must use a name that is different from all the other policy names for the object. For example, you cannot have two policies named “Latency” for an internal volume; however, you can have a “Latency” policy for an internal volume and another “Latency” policy for a different volume. The best practice is to always use a unique name for any policy, regardless of the object type.

5. From the **Apply to objects of type** list, select the type of object to which the policy applies.

6. From the **With annotation** list, select an annotation type, if applicable, and enter a value for the annotation in the **Value** box to apply the policy only to objects that have this particular annotation set.

7. If you selected **Port** as the object type, from the **Connected to** list, select what the port is connected to.

8. From the **Apply after a window of** list, select when an alert is raised to indicate a threshold violation.

The First occurrence option triggers an alert when a threshold is exceeded on the first sample of data. All other options trigger an alert when the threshold is crossed once and is continuously crossed for at least the specified amount of time.

9. From the **With severity** list, select the severity for the violation.

10. By default, email alerts on policy violations will be sent to the recipients in the global email list. You can override these settings so that alerts for a particular policy are sent to specific recipients.

- Click the link to open the recipients list, then click the **+** button to add recipients. Violation alerts for that policy will be sent to all recipients in the list.

11. Click the **any** link in the **Create alert if any of the following are true** section to control how alerts are triggered:

- **any**

This is the default setting, which creates alerts when any of the thresholds related to a policy are crossed.

- **all**

This setting creates an alert when all of the thresholds for a policy are crossed. When you select **all**, the first threshold that you create for a performance policy is referred to as the primary rule. You must ensure that the primary rule threshold is the violation that you are most concerned about for the performance policy.

12. In the **Create alert if** section, select a performance counter and an operator, and then enter a value to create a threshold.
13. Click **Add threshold** to add more thresholds.
14. To remove a threshold, click the trash can icon.
15. Select the **Stop processing further policies if alert is generated** check box if you want the policy to stop processing when an alert occurs.

For example, if you have four policies for datastores, and the second policy is configured to stop processing when an alert occurs, the third and fourth policies are not processed while a violation of the second policy is active.

16. Click **Save**.

The Performance Policies page displays, and the performance policy appears in the list of policies for the object type.

Configuring performance and assure violation notifications

OnCommand Insight supports notifications for performance and assure violations. By default, Insight does not send notifications for these violations; you must configure Insight to send email, to send syslog messages to the syslog server, or to send SNMP notifications when a violation occurs.

Before you begin

You must have configured email, syslog, and SNMP sending methods for violations.

Steps

1. Click **Admin > Notifications**.
2. Click **Events**.
3. In the **Performance Violations events** or **Assure Violations events** section, click the list for the notification method (**Email**, **Syslog**, or **SNMP**) you want, and select the severity level (**Warning and above** or **Critical**) for the violation.
4. Click **Save**.



Monitoring the violations in your network

When Insight generates violations due to the thresholds set in performance policies, you can view them using the Violations Dashboard. The dashboard lists all the violations that occur in your network and enables you to locate and address issues.





Steps


1. Open OnCommand Insight in your browser.
2. On the Insight toolbar, click **Dashboards** and select **Violations Dashboard**.

The Violations Dashboard displays.

3. You can use the **Violations By Policies** pie chart in the following ways:
 - You can position your cursor over any slice of a chart to display the percentage of the total violations that occurred for a particular policy or metric.
 - You can click a slice of a chart to “enlarge” it, which enables you to emphasize and study more carefully that slice by moving it away from the rest of the chart.
 - You can click the  icon in the upper-right corner to display the pie chart in full screen mode, and click  again to minimize the pie chart. A pie chart can contain a maximum of five slices; thus, if you have six policies that generate violations, Insight combines the fifth and sixth slices into an “Others” slice. Insight assigns the most violations to the first slice, the second most violations to the second slice, and so on.
4. You can use the **Violations History** chart in the following ways:
 - You can position your cursor over the chart to display the total number of violations that occurred at a particular time and the number that occurred out of the total for each specified metric.
 - You can click a legend label to remove the data associated with the legend from the chart.

Click on the legend to display the data again.

 - You can click the  icon in the upper-right corner to display the chart in full screen mode, and click  again to minimize the pie chart.
5. You can use the **Violations Table** in the following ways:
 - You can click the  icon in the upper-right corner to display the table in full screen mode, and click  again to minimize the pie chart.

If your window size is too small, then the Violations Table displays only three columns; however, when you click , additional columns (up to seven) display.


- You can display violations for a particular time period (**1h, 3h, 24h, 3d, 7d, and 30d**), with Insight showing a maximum number of 1000 violations for the selected time period.
- You can use the **filter** box to show only the violations you want.
- You can change the sort order of the columns in a table to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header; to return to the default sort order, click any other column header.

By default, the table displays the violations in descending order.

- You can click a violation in the ID column to display the asset page for the duration of the violation.

- You can click the resource links (for example, storage pool and storage volume) in the Description column to display the asset pages associated with those resources.
- You can click the performance policy link in the Policy column to display the Edit Policy dialog box.

You might want to adjust the thresholds for a policy if you feel it generates too few or too many violations.

- You can click a page number to browse through data by page if there is more data than fits on a single page.
- You can click  to dismiss the violation.

Troubleshooting Fibre Channel BB credit 0 errors

Fibre Channel uses buffer-to-buffer credits (BB credits) to control transmission flow. The credit value is decremented when a frame is sent from a port and the credit value is replenished when the port receives a response. If the BB credits in the port are not replenished, the transmission flow can be impacted. Ports need memory, or buffers, to temporarily store frames until they are assembled in sequence, and delivered. The number of buffers is the number of frames a port can store and is called a Buffer Credit.

As the available credits for a given port approach zero, an error warns that the port will stop receiving transmissions when zero is reached and will not resume until the BB credits are replenished.

Insight performance policies allow you to set thresholds on the following port metrics.

BB credit zero - Rx
Number of times the receive buffer-to-buffer credit count transitioned to zero during the sampling period
BB credit zero - Tx
Number of times the transmit buffer-to-buffer credit count transitioned to zero during the sampling period
BB credit zero - Total
Number of times this port had to stop transmitting because the attached port was out of credits to provide
BB credit zero duration - Tx
Time in milliseconds during which the Tx BB credit was zero during the sampling interval

BB Credit errors might be caused by some of the following scenarios:

- If a given implementation has a high percentage of FC frames of sizes significantly less than the maximum size, then more BB_Credits might be required.
- Workload changes to your environment that could be impacting ports or devices that are connected to them, such as storage nodes.

You can use the fabric, switch, and port asset pages to monitor your Fibre Channel environment. Port asset pages present summary information about the resource, its topology (the device and its connections), performance charts, and tables of associated resources. When troubleshooting Fibre Channel issues the performance chart for each port asset is helpful because it shows traffic for the selected top contributor port. Port asset pages also show buffer-to-buffer credit metrics and port errors in this chart, with Insight displaying a separate performance chart for each metric.

Creating performance policies and thresholds for ports

You can create performance policies with thresholds for metrics associated with a port. By default, performance policies apply to all devices of the specified type when you create them. You can create an annotation to include only a specific device or a set of devices in the performance policy. For simplicity, an annotation is not used in this procedure.

Before you begin

If you want to use an annotation with this performance policy, you must create the annotation before you create the performance policy.

Steps

1. From the Insight toolbar, click **Manage > Performance policies**

The existing policies are displayed. If a policy exists for switch ports, you can edit the existing policy, adding the new policies and thresholds.

2. Edit an existing port policy or create a new port policy
 - Click the pencil icon at the far right of the existing policy. Add the thresholds described in steps “d” and “e”.
 - Click **+Add** to add a new policy
 - a. Add a “Policy Name”: Slow Drain Device
 - b. Select port as object type
 - c. Enter First occurrence for “Apply after window” of
 - d. Enter threshold: BB credit zero - Rx > 1,000,000
 - e. Enter threshold: BB credit zero - Tx > 1,000,000
 - f. Click “Stop processing further policies if alert is generated”
 - g. Click “Save”

The policy you create monitors the thresholds you set over a period of 24 hours. If the threshold is exceeded a violation is reported.

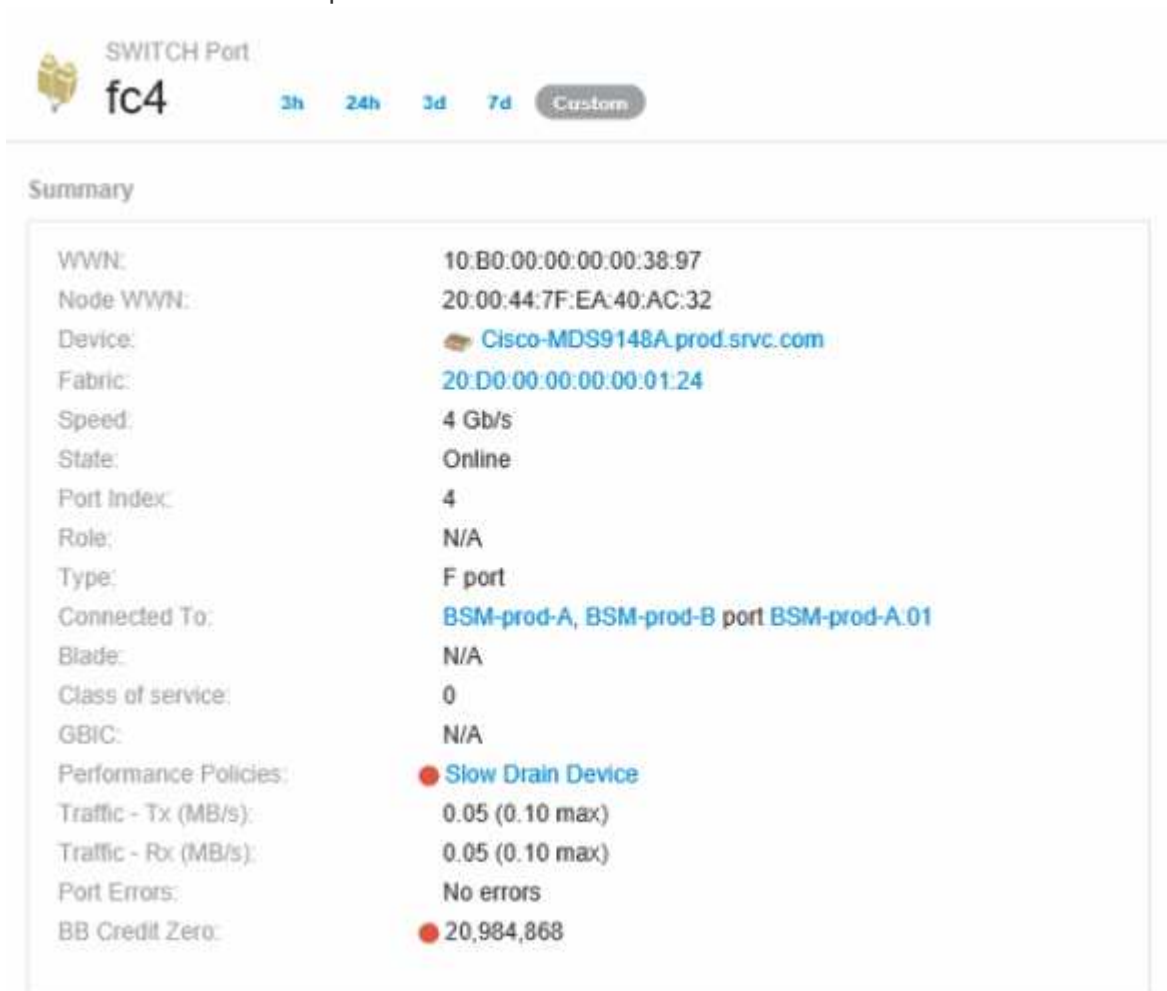
3. Click **Dashboards > Violations Dashboard**

The system displays all of the violations that have occurred on the system. Search or sort the violations to view the “Slow Drain Device” violations. The Violations Dashboard shows all of the ports that experienced BB Credit 0 errors exceeding the thresholds that were set in the performance policy. Each switch port identified in the violations dashboard is a highlighted link to the Port Landing Page.


4. Click on a highlighted port link to display the Port Landing Page.

The port landing page is displayed and includes information useful for BB Credit 0 troubleshooting:

- Devices the port is connected to
- Identification of the port reporting the violation, which is a Fiber Channel Switch port.
- The port's speed
- The associated node and port name



5. Scroll down to view the port metrics. Click **Select metrics to show > BB credit zero** to display the BB credit graph.



6. Click **Top correlated**

The top correlated resource analysis shows the connected controller node the port is servicing as the resource most correlated with performance. This step compares the IOPS metrics of the port activity with the overall node activity. The display shows the Tx and Rx BB Credit Zero metrics and the controller node's IOPS. The display shows the following:

- The controller IOs are highly correlated with the port traffic
- The performance policy is violated when the port is transmitting IO to the server.
- Given that our port performance violation is occurring in conjunction with a high IOPS load on the storage controller it is likely that the violation is due to the workload on the storage node.



7. Return to the Port Landing Page and access the storage controller node's landing page to analyze the workload metrics.

The Node shows a utilization violation and metrics show high "cache reads replaced" that correlates to buffer-to-buffer zero credit states.

Storage: BSM-prod-A, BSM-prod-B
 HA partner: BSM-prod-B
 State: N/A
 Model: FAS6070
 Version: 8.0.5 7-Mode
 Serial number: 700001181351
 Memory: 98,304 MB
 Utilization: ● 21.26% (94.56% max)
 IOPS: 232.73 IO/s (1,153.00 IO/s max)
 Latency: 7.07 ms (15.00 ms max)
 Throughput: 22.44 MB/s (106.00 MB/s max)
 Processors: 12
 Performance Policies: ● Node Utilization
 Node Read Latency

8. From the Node landing page, you can compare the BB credit zeroes by selecting the port from the correlated resources list, and select utilization data, including Cache utilization data, for our node from the metrics menu.



This data makes it clear that the Cache hit Ratio is inversely correlated to our other metrics. Instead of being able to respond to the server load from cache, the storage node is experiencing high cache reads replaced. It is likely that having to retrieve most of the data from disk rather than cache is causing the delay in the port's transmission of data to the server. The cause of the performance problem appears likely to be a workload generated change in IO behavior, and that the node cache, and its configuration, are the cause. The problem might be solved by increasing the node's cache size or change the caching algorithm's behavior.

Analyzing your infrastructure

The procedures in this topic are ones that you might use to perform an analysis of parts of the infrastructure in your environment. The steps, views, and data you gather in this exercise use virtual computing objects as an example. Analysis of other assets in your environment will follow similar steps using relevant counters for each specific asset. The intent of this exercise is to familiarize you with the variety of options Insight offers to monitor and understand the characteristics of the assets in your data center.

About this task

Some of the actions you can take to analyze the state of your infrastructure might include the following:

- Observe an object's behavior over time
- Compare an object's metrics against the metrics of the top 10 like objects
- Compare numbers for objects
- Compare the top 10 objects against the average
- Compare metrics A vs. B for many objects to show categories and anomalies
- Compare a range of objects against other objects
- Use an expression to display metrics not available in the web UI

You can create all of these views of objects in your infrastructure in a dashboard using widgets for each analysis you perform. The dashboards can be saved to provide quick access to current data on your infrastructure.

Observe an object's behavior over time

You can observe the behavior of a single object to determine if the object is operating within expected operational levels.

Steps

1. Use a query to identify the VM that will be the subject of analysis: **Query > + New query > Virtual machine > "name"**

Leaving the name field blank returns all VMs. Select the VM that you want to use in this exercise. You can select it by scrolling through the list of VMs.

2. Create a new dashboard for the information you want to collect. From the toolbar, click **Dashboards > +New Dashboard**.

3. In the new Dashboard, select **Variable > Text**.
 - a. Add the VM name from your query as the `$var1` value.
 - b. Click the check box.

The variable is used to easily swap between different sets of objects you want to analyze. In other steps of your analysis, you may reuse this variable for additional analysis against the single VM initially chosen. Variables become more useful when identifying multiple objects.

4. Add a line chart widget to the new dashboard: **Widget > Line chart**.
 - a. Change the default asset type to virtual machine: Click **Virtual machine > Latency-Total**.
 - b. Click **Filter by > Name > \$var1**.
 - c. Change the time period on the dashboard: **Override dashboard time > On > 7 days**.


You can change the duration of the display using any of the preset selections or by specifying a custom time range.

+ The dashboard displays the VM's **IOPS-Total** for the period of time you specify.

5. Assign a name to the widget and save the widget.

Results

Your widget should contain data similar to the following:



The VM shows a period of abnormally high latency for a short period of time in the 7 days that are displayed.

Compare objects with the top 10 latency total to the average latency for all like objects

You might want to compare the VMs with the top 10 latency total to the average latency total to identify any that are extremely out of the average range. This information could help in decisions to balance workloads on VMs.

Steps

1. Add a widget with a stacked area chart to the to the new dashboard: **Widget > Stacked Area Chart**

- a. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency total**

The widget displays the Latency Total, for all VMs, for 24 hours in a stacked area chart.

- b. Create a second display in this widget that shows Latency Total averaged for all VMs: **Widget > Line chart**

- c. Change the default device to Virtual machine: Click **Virtual machine > Latency-total**

The widget displays the Latency Total for the default 24 hour period of time using a line chart.

- d. Click **X** on the **Roll up** bar and select **Show > Top > 10**

The system displays the Top 10 VMs based on Latency Total.

2. To compare the average Latency Total for all VMs to the Top 10 IOPS total use the following steps:

- a. Click **+Add**

- b. Change the default device to Virtual machine: Click **Storage > Virtual machine > IOPS total**

- c. Click **X** on the **Roll up** bar and select **Show > Top > 10**

The system displays the 10 objects with high latency and shows the average latency in a line chart.

+ image::.../media/analytics-top10-avg.gif[]

+ The average latency is 1.6 ms, while in the top ten, the are VMs experiencing latency of over 200 ms.

Compare one object's latency total to the latency total of the top 10 objects

The following steps compare a single VM's Latency Total to the VMs reporting the Top 10 Latency Total in the entire virtual infrastructure.

Steps

1. Add a widget with a line chart to the to the new dashboard: **Widget > Line Chart**

- a. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency-total**

The widget displays the total Latency, for all VMs, for the default 24 hours in an area chart.


- b. Create a second display in this widget that shows Latency Total averaged for all VMs: **Widget > Line chart**

c. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency-Total**

The widget displays the Latency total for the default 24 hour period of time using a line chart.

d. Click **X** on the **Roll up** bar and select **Show > Top > 10**

The system displays the Top 10 VMs based on Latency - Total.



2. Add the VM that you want to compare to the Top 10:

a. Click **+Add**

b. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency total**

c. Click **Filter by > Name > \$var1**

3. Click **Show legend**

Results

A legend identifies each of the VMs under analysis. You can easily identify VM_Exchange_1 and determine if it is experiencing latency similar to the top ten VMs in the environment.

Compare metrics-A against metrics-B to show categories and anomalies


You can use a scatter plot to show two sets of data for each object. For example, you can specify IOPS Read and Latency Total to be displayed for each object. Using this chart you can identify the object you consider troublesome based on both the IOPS and the Latency combined.

Steps

1. Add a widget with a scatter plot chart to the to the new dashboard: **Widget > Scatter Plot Chart**

2. Change the default device to Virtual machine: Click **Storage > Virtual machine > Latency total > IOPS Read**

The system displays a scatter plot similar to the following:



Use an expression to identify alternate metrics


You can use expressions to view metrics not provided by the web UI, such as the IOPS that are system overhead generated.

About this task

You might want use an expression to show total IOPS generated by non-read or non-write operations, such as overhead operations for an internal volume.

Steps

1. Add a widget to the dashboard. Choose **Area chart**.
2. Change the default device to Internal volume: Click **Storage > Internal volume > IOPS Write**
3. Click the **Convert to Expression** button.
4. The **IOPS - Write** metric is now in the alphabetic variable field “a”.
5. In the “b” variable field, click **Select** and choose **IOPS - Read**.
6. In the **Expression** field, type **a + b**. In the **Display** section, choose **Area chart** for the expression.
7. In the **Filter by** field, enter the name of the internal volume you are analyzing.
8. The **Label** field identifies the expression. Change the label to something meaningful like “R + W IOPS”.
9. Click **+Add** to add a line for total IOPS to the widget.
10. Change the default device to Internal volume: Click **Storage > Internal volume > IOPS Total**
11. In the **Filter by** field, enter the name of the internal volume you analyzing.



The chart displays the total IOPS as a line, with the chart showing the combination of read and write IOPS in blue. The gap between 9:30 and 9:45 shows non-read and non-write IO (overhead) operations.

Introduction to minimizing risk in thin provisioning

In today's hybrid IT data centers, administrators are pressured to stretch resource utilization beyond physical bounds by employing capacity efficiency technologies such as thin provisioning to control over allocation and leverage what was once unavailable capacities.

OnCommand Insight provides near real time capacity usage and utilization details historically across multiple thin provisioned layers within the IT service stack. Failing to properly manage oversubscription risk could result in untimely downtime to the business.

Monitoring the storage pool

Each storage pool landing page provides over-subscription ratios, identifies correlated resources, LUN and disk utilization, as well as policy breaches and violations that have occurred with the storage pool.

Use the storage pool landing page to identify any potential problems with the physical assets supporting your virtual infrastructure. You can track capacity and capacity ratios trending over 30 days or use a custom time frame. Pay attention to data in the following sections to monitor the status of the storage pool.

- **Summary**

Use this section to understand:

- Storage pool capacity information including physical capacity and the overcommitted capacity.
- Whether the aggregate is oversubscribed, and by how much.
- Any policy violations that have occurred.

- **Storage resources and Disks sections**

The storage resources section shows the LUN utilization.

The disks section shows the individual disks that make up the storage pool.

- **Resources**

Use this section to understand the VMDKs to LUNs correlation and understand the storage to VM application path.

- **Violations section**

The violations section identifies any breaches to performance policies that have been set for the storage pool.

Monitoring the Datastores

The Datastore landing page identifies over-subscription ratios, LUN and disk utilization, correlated resources, and shows policy breeches and violations that have occurred with the Datastore.

Use this landing page to identify problems with your virtual infrastructure. You can track capacity and capacity ratio trending to anticipate changes in your capacity.

- **Summary**

Use this section to understand:

- Datastore capacity information including physical capacity and the overcommitted capacity.
- The percentage of overcommitted capacity.
- Metrics for latency, IOPS, and throughput.

- **VMDKs**

The VMDKs section shows virtual disk capacity and performance.

- **Storage resources**

This section shows the capacity used and the performance metrics for the internal volume correlated to the Datastore.

- **Resources**

Use this section to understand the VMDKs to LUNs correlation, and understand the storage to VM application path.

- **Violations section**

The violations section identifies any breaches to performance policies that have been set for the Datastore.

Create dashboards to monitor thin provisioned environments

OnCommand Insight's flexible dashboard widget design and display charting options allow deep analysis into capacity usage and utilization, strategic information for minimizing risks in thin provisioned data center infrastructures.

You can create dashboards that provide access to Datastore and Storage pool information that you want to monitor.

Using dashboards to access Datastore information

You might want to create dashboards that provide quick access to the data you want to monitor in your virtual infrastructure. A dashboard could include widgets similar to the following to identify the top 10 Datastores based on their overcommitted % and a widget showing the capacity data for Datastores. The dashboards use variables to highlight Datastores that are overcommitted by more than 150% and Datastores that have exceeded more than 80% used capacity.

New dashboard

3h 24h 3d 7d 30d Custom



Overcommit Subscription %

Name	Capacity - Total (GB)	Capacity - Used (GB)	Capacity - Provisioned (GB)	Capacity ratio - Over commit (%)	Capacity ratio - Used (%)
dc14_1407_...1_prd_cl03	5,008.00	4,091.04	12,876.38	257.12	81.69
dc14_1407_...2_prd_cl03	6,936.69	5,872.31	14,633.80	210.96	84.66
dc14_1407_...3_prd_cl03	9,437.03	7,951.36	17,639.86	186.92	84.26
dc14_1407_...4_prd_cl03	7,911.09	6,627.00	17,891.24	226.15	83.77


4 items found

Additional widgets that could be used to monitor your thin-provisioned environment could include some of the following information:

- VMDK capacities correlated to Datastores
- VM capacities
- Data store capacity used trending

Using dashboards to access Storage pool information

A dashboard could include widgets similar to the following, identifying the amount of physical storage capacity used, or identifying the overcommitted capacity for a Storage Pool.



Using performance policies to reduce risk in thin provisioning

You should create performance policies to raise alerts when thresholds in your virtual infrastructure have been breached. The alerts allow you to respond to changes in your environment before they cause interruptions or outages in operations.

Policies that help in monitoring the virtual infrastructure include the following:

- **Datastore**

You could use the following policies on the Datastore:

- Capacity ratio - Overcommit
- Capacity ratio - Used
- Capacity - Used
- Capacity - Total

- **Storage pool**

The following policies can protect against storage related capacity outages in thin provisioned environments:

- Capacity provisioned
- Capacity used

- Capacity ratio - Overcommit
- Capacity ratio - Used

You can expand from these policies to monitor capacity in the virtual infrastructure, including:

- Internal volumes
- LUNs
- Disks
- VMDKs
- VMs

You can configure policies using annotations. You assign the same annotation to the specific assets that support an application. For example, you can assign annotations to the Datastores and the Storage pools of a thin provisioned application. You might have annotations named Production for the production environment, Development for the development environment, and so on. You can change the thresholds and criticality of warnings depending on the type of application the assets are supporting. For example, a breach of a threshold for a production application's DataStore might raise a *critical warning*, while the same breach for a development environment might only raise a *warning*. Incorporating annotations within defined policies can help to further reduce unwanted alerting noise for non-critical assets.

Creating performance policies for Storage Pools

You can create performance policies that trigger alerts to notify you when thresholds for Storage Pool assets have been exceeded.

Before you begin

This procedure assumes that you have thin provisioned the storage pool.

About this task

You want to create policies that monitor and report changes in a storage pool that could contribute to outages. For the thin provisioned physical storage pool, you want to monitor the physical capacity and monitor the Overcommit Ratio.

Steps

1. Open OnCommand Insight in your browser.
2. Select **Manage > Performance Policies**

The Performance Policies page is displayed. Policies are organized by object, and are evaluated in the order in which they appear in the list. If notifications are enabled (**Admin > Notifications**), you can configure Insight to send email when performance policies are breached.

3. Click **+Add** to create a new policy.
4. In **Policy Name** enter a policy name for the Storage Pool.
5. In **Apply to objects of type** select Storage Pool.
6. In **Apply after window of** enter First occurrence.
7. In **With severity** enter Critical

8. Configure the Email recipients that you want notified when thresholds are breached.

By default, email alerts on policy violations are sent to the recipients in the global email list. You can override these settings so that alerts for a particular policy are sent to specific recipients.

Click the link to open the recipients list, then click the + button to add recipients. Violation alerts for this policy will be sent to all recipients in the list.

9. In **Create alert if any of the following are true** enter Capacity ratio - Used > 85%

Results

This configuration results in the system sending a critical warning message when more than 85% of the physical capacity of the storage pool is used. Using 100% of the physical memory will result in application failure.

Create additional Storage Pool policies

About this task

Create an additional “Capacity ratio - Used” policy that raises a warning message when the Storage Pool capacity used exceeds 75%. If notifications are enabled (**Admin > Notifications**), you can configure Insight to send email when performance policies are breached.

Creating performance policies for Datastores

You can create performance policies with thresholds for metrics associated with the datastores that correlate to the storage pools you are monitoring. By default, performance policies apply to all devices of the specified type when you create them. You can create an annotation to include only a specific device or a set of devices in the performance policy.

Before you begin

When using an annotation in a performance policy, the annotation must exist before the policy is created.

About this task

You create a performance policy that provides notification when one or more Datastores you are monitoring exceeds a threshold you set. Your system might already contain a global policy that meets your needs or a policy using annotations might also work if you annotate your Datastores.

Steps

1. From the Insight toolbar, select **Manage > Performance Policies**

The performance policies page is displayed. Review any existing performance policies to identify existing policies that address the metrics for thresholds you want to monitor.

2. Click **+Add** to add a new policy

3. Add a “Policy Name”

You must use a name that is different from all the other policy names for the object. For example, you

cannot have two policies named "Latency" for an internal volume; however, you can have a "Latency" policy for an internal volume and another "Latency" policy for a data store. The best practice is to always use a unique name for any policy, regardless of the object type.

4. Select "Datastore" as the Object Type
5. Click "First Occurrence"

The First occurrence option triggers an alert when a threshold is exceeded on the first sample of data. All other options trigger an alert when the threshold is crossed once and is continuously crossed for at least the specified amount of time.

6. Click "Warning"
7. For "Create alert", select **Capacity ratio - Over commit** and set the value to **> 150**

You might want to create additional capacity related alerts, such as **Capacity total** and **Capacity used**.

Collecting Host and VM file system utilization data

The Host and VM File Systems data source, combined with the Host Utilization license, enables reporting and chargeback at the file system level for known Hosts and VMs.

OnCommand Insight collects data from storage devices, most of which report their volumes as block devices. This allows Insight to report on utilization at the storage level, but not at the file system level. Storage arrays typically know which blocks have been written to, but not which blocks have been freed.

Client hosts and VMs implement file systems (ntfs, ext*...) on top of these block devices. Most file systems keep a table of contents containing directory and file metadata. When files are deleted, their entries are simply removed from the table of contents. Blocks consumed by those files are now eligible for re-use by the file system, but the storage array doesn't know this. In order for Insight to report on filesystem usage, it must be collected from the client host or VM point of view for accurate chargeback.

Insight allows this level of file system utilization data collection through the **NetApp Host and VM File System** data source, in combination with the **Host Utilization** license. VM's must be annotated with the appropriate **Compute Resource Group** name, and associated storage arrays must be annotated with appropriate **Tier** annotations with proper costs for accurate cost reporting.



The Host Utilization License is resource-based, as opposed to capacity-based as other Insight licenses.

Configure Insight for file system collection

To configure Insight for collection of file system utilization data, you must install the Host Utilization Pack license and configure the NetApp Host and VM File Systems data source.

Before you begin

If you haven't already, install the Host Utilization Pack license. You can check for the license in the **Admin > Setup** page, on the **Licenses** tab.

The Host and VM File Systems data source only reports file system utilization and file system metadata for

known **Compute Resources** (hosts and VMs) currently being collected or discovered in Insight:

- Virtual Machines are collected by hypervisor data sources such as Hyper-V and VMware.
- Hosts are discovered via device resolution.

The proper Tier annotations must be present on the appropriate storage resources.

The following connected block storage devices are supported:

- NetApp Clustered Data OnTap (cDOT)
- NetApp 7-Mode
- Clariion
- Windows: VMWare virtual disks (VMDKs) for FC, iSCSI
- Linux: VMWare VMDKs (iSCSI and FC not supported)

A **Compute Resource Group** is an annotation that allows grouping of hosts and/or virtual machines that share a common administrative credential.

Steps

1. First, annotate the hosts and/or virtual machines to be included in your **Compute Resource Group**. Go to **Queries > +New query** and search for *Virtual Machine* assets.

You will need to repeat these steps for *Host* assets.

2. Click on the column selector on the right of the table and select the **Compute Resource Group** column to display it in the query results table.
3. Select the virtual machines you wish to add to the desired compute resource group. You can use a filter to search for specific assets.
4. Click on the **Actions** button and choose **Edit annotation**.
5. Select the *Compute Resource Group* annotation, then choose the desired resource group name in the *Value* field.

The resource group annotation is added to the selected VMs. The resource group name must match the name you will configure in the Host and VM File Systems data source later.

6. To configure the Host and VM File Systems data source for a compute resource group, click on **Admin > Data sources** and **Add** the *NetApp Host and VM File Systems* data source.

Settings

*Name

Vendor

Model

Where to run

What to collect

Clustered Data ONTAP 8.1.1+

Clustered Data ONTAP 8.1.1+ (Unified Manager 6.0+)

Data ONTAP 7-Mode

E-Series (Firmware 6.x)

E-Series (Firmware 7.x+)

Host and VM File Systems

SolidFire 8.1+

StorageGrid

Configuration

Advanced configuration

Test

Cancel Save

7. In the **Configuration** section, enter a **User Name** and **Password** for an operating system user with appropriate rights to retrieve file system data. For Windows operating system users, this must include the domain prefix if your Windows environment uses it.

Note that an Insight Acquisition Unit (AU) installed on Linux can report on Linux compute resources, while an AU installed on Windows can talk to either Linux or Windows compute resources.

8. Enter the name of the **Compute Resource Group** for the assets from which you will want to collect file system utilization data. This name must match the resource group name you used to annotate the assets above.

If you leave the Compute Resource Group field empty, the data source will collect data for hosts or VMs that have no Compute Resource Group annotation.

9. In the **Advanced Configuration** section, enter the desired polling interval for this data source. The default of 6 hours is usually adequate.
10. It is recommended to **Test** the data source connection before saving it. A successful connection result will also show you how many compute resource targets are contained in the group.
11. Click **Save**. The Host and VM File Systems data source will begin collecting data on its next poll.
12. Once file system data is being collected, you can view it on the host's or VM's asset page, in the File System widget:

File Systems

Name	Capacity (Used / Total GB)	Type	Storage Resource
/	9.15% (11.0 / 120.0)	xfs	vifasnane:...vm_oci_
/boot	23.79% (0.1 / 0.5)	xfs	vifasnane:...vm_oci_
/dev/dm-1	7.8	swap	vifasnane:...vm_oci_

Showing 1 to 3 of 3 entries

- Repeat these steps for each Compute Resource Group you will have. Each compute resource group must be associated with its own Host and VM File Systems data source.

Note that file system information will be collected for hosts and VM's that are already being acquired by any traditional VMware or Hyper-V data sources in your environment.

File system chargeback and reporting


Chargeback for file systems is always performed from the storage perspective. Storage arrays associated with virtual machines annotated for a particular compute resource group will be included in chargeback reports for that resource group.

Before you begin

Any virtual machines which you wish to include in file system utilization chargeback must be annotated with the appropriate compute resource group name. Storage arrays associated with those virtual machines must be annotated with the appropriate Tier annotations. ETL to data warehouse must have occurred after these annotations are in place.

Steps

- Open a browser to your Reporting server, usually <https://<host>:9300/p2pd> or <http://<host>:9300/bi> (7.3.3 or later) and log in.
- Choose the **File System Utilization** package and create a new report.



- Drag and drop items from your data mart(s) to build your report.

The example below is a very simple report. You can create complex reports built around your specific business needs.

Name	Type	Allocated Capacity GB	Used Capacity GB	Tier Name	Cost	Storage Name
/	xfs	119.96	9.96	N/A		vifasnane05,vifasnane06
/	xfs	5,492.53	799.63	Tier 1	100	vifasnane
/boot	xfs	0.48	0.17	N/A		vifasnane05,vifasnane06
/boot	xfs	8.72	2.41	Tier 1	100	vifasnane
/dev/dm-1	swap	7.81	0.00	N/A		vifasnane05,vifasnane06
/dev/dm-1	swap	140.61	0.78	Tier 1	100	vifasnane
C:\	NTFS	948.27	331.98	Tier 1	100	vifasnane
PHYSICALDRIVE0: System Reserved	NTFS	1.70	1.41	Tier 1	100	vifasnane

Configuring your system to report chargeback data

Chargeback reports provide storage capacity chargeback and accountability information by hosts, application, and business entities, and include both current and historical data.

This guide describes how to configure Insight to generate a chargeback report providing accountability for service level costs and storage usage cost. The intent of the guide is to provide the steps required to create a simple chargeback report, and familiarize Insight users with the options available when configuring chargeback in their unique environment.

For each application, the example report identifies the resources provisioned and the cost of the resources. The output for the report is created by defining the following data in Insight

- Storage tiers
- Cost associated with each storage tier
- Provisioned storage capacity
- Service levels
- Cost per service level

The following sections describe the steps required to configure this data so that it can be accessed by Insight Reporting.

Defining annotations for use with chargeback

When customizing OnCommand Insight to track data for your corporate requirements, you can define specialized annotations needed to provide a complete picture of your data: for example, an annotation can define an asset's end of life, or the data center the asset resides in, or a storage tier defining the cost per GB of the storage.

About this task

The chargeback report example in this guide provides data for Service level and for Tier level. You must create annotations for each Service level and Tier level and then define costs for the Service levels and Tier levels.

Steps

1. Log in to the Insight web UI
2. Click **Manage > Annotations**

The annotations page displays.

3. Position your cursor over the Service Level, or Tier annotation, and click .

The Edit Annotation dialog box displays.

4. Click **ADD** to add new Tiers and cost.

In the report example, the Tier and Service level names use the precious metal analogy of Gold, Silver, and Bronze. You can use any naming convention chosen by your organization, for example, Tier 1, Level 2, Supreme.

5. Enter the values for the Gold-Fast, Gold, Silver, and Bronze Tiers and the costs associated with each.

The values you enter define the cost per-GB for the storage that is used by applications. The service level cost can be the cost of providing the service or the actual price to service the consumer. These costs will be reported in the Chargeback report.

6. Click **Save** when you finish.

Defining applications for use with chargeback

If you want to track cost data associated with specific applications running in your environment, you first need to define the applications.

Before you begin

If you want to associate application with a business entity, you must have already created the business entity.



This example does not associate any application with business entities.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage > Application**

After you define an application, the Applications page displays the application's name, its priority, and, if applicable, the business entity associated with the application.

3. Click **Add**

The Add Application dialog box displays.

4. Enter a unique name for the application in the Name box. Enter the Applications identified in the Report: African Tours, APAC Commercial Sales, and so on.
5. Click **Priority** and select the priority (critical, high, medium, or low) for the application in your environment.
6. If you plan to use this application with a business entity, click **Business entity** and select the entity from

the list.

7. You will not use volume sharing, click to clear the **Validate** volume sharing box.
8. Click **Save**.

The applications appears in the Applications page. If you click the application's name, Insight displays the asset page for the application. After defining an application, you can go to an asset page for host, virtual machine, volume, internal volume, or hypervisor to assign an application to an asset.

Assigning applications to assets

After defining your applications, you need to associate the applications with specific assets. You can use a simple ad hoc method to apply an applications to an asset. Users looking to apply applications in bulk should use a query method to identify the assets they want to assign to an application.

Assigning applications to assets using an ad hoc method


You assign an application to an asset so that you can identify the resources of the asset that the application uses. If an asset has a cost assigned to it you can identify the cost that is incurred by the application and if the resource is measured by size, you can determine if the resource will need to be replenished.


About this task

Use the following method to assign applications to assets.

Steps

1. Log in to the OnCommand Insight web UI.
2. Locate the asset (host, virtual machine, volume, or internal volume) to which you want to apply the application by doing either of the following:

Option	Description
Navigate to the list of assets	Click Dashboard > Assets Dashboard and select the asset.
Search for the asset	Click  on the toolbar to display the Search assets box, type the name of the asset, and then select the asset from the list.


3. In the **User Data** section of the asset page, position your cursor over the name of the application currently assigned to the asset (if there is no application assigned, **None** is displayed) and then click  (Edit application).

The list of available applications for the selected asset is displayed. The applications that are currently associated with the asset are preceded by a check mark.

4. You can type in the Search box to filter the application names, or you can scroll down the list.

5. Select the applications you want to associate with the asset.

You can assign multiple applications to host, virtual machine, and internal volume; however, you can only assign one application to a volume.

6. Click  to assign the selected application or applications to the asset.

The application names appear in the User Data section; if the application is associated with a business entity, the name of the business entity appears in this section also.

Assigning applications to an asset using a query

You assign an application to an asset so that you can identify the resources of the asset that the application uses. If an asset has a cost assigned to it you can identify the cost that is incurred by the application and if the resource is measured by size, you can determine if the resource will need to be replenished.

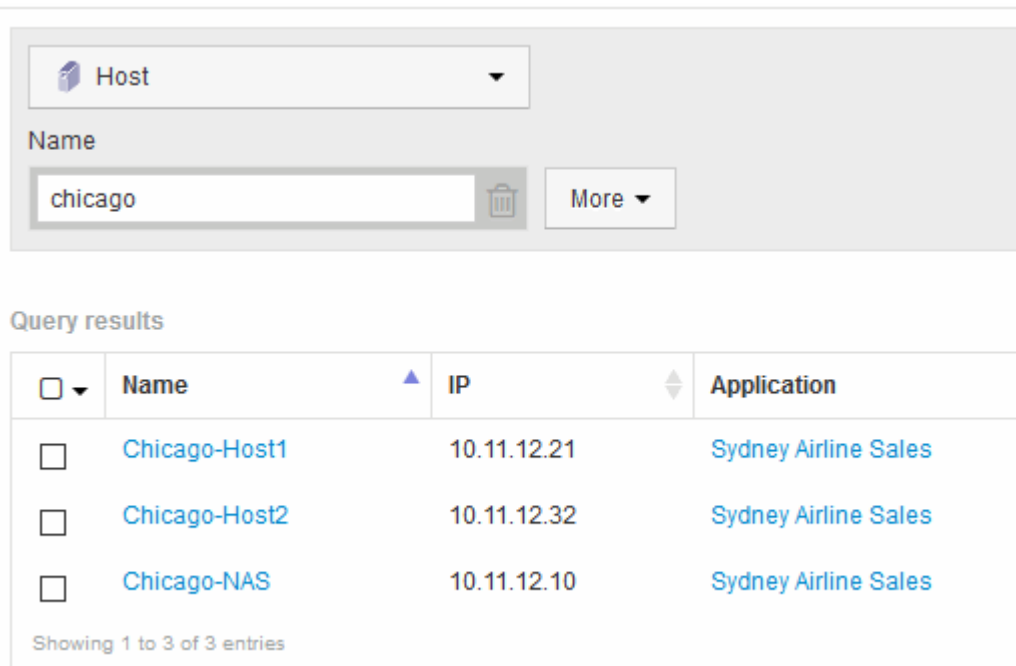
About this task

You can simplify the task of assigning multiple assets to an application by using a query.

Steps

1. Create a new query to identify the assets you want to assign an application to. For example, if you want to assign it to a host with a specific name that relates to a geographic location, Click **Queries** > **+New Query**
2. Click **Host**
3. In the **Name** field, enter Chicago

The system displays all of the Hosts with **Chicago** as part of their name.




The screenshot shows a web interface for selecting hosts. At the top, there is a dropdown menu with 'Host' selected. Below it is a 'Name' field containing the text 'chicago'. To the right of the field is a trash icon and a 'More' button with a dropdown arrow. Below the 'Name' field is a section titled 'Query results' which contains a table with three columns: 'Name', 'IP', and 'Application'. The table lists three hosts: 'Chicago-Host1', 'Chicago-Host2', and 'Chicago-NAS', all with IP addresses and 'Sydney Airline Sales' as the application. At the bottom of the table, it says 'Showing 1 to 3 of 3 entries'.

<input type="checkbox"/>	Name	IP	Application
<input type="checkbox"/>	Chicago-Host1	10.11.12.21	Sydney Airline Sales
<input type="checkbox"/>	Chicago-Host2	10.11.12.32	Sydney Airline Sales
<input type="checkbox"/>	Chicago-NAS	10.11.12.10	Sydney Airline Sales


Showing 1 to 3 of 3 entries

4. Select one or more of the Hosts identified by your query.

5. Click **Actions > Add Application**



The system displays the Assign application dialog.

6. Select the Application you want to assign to the Host and click 
7. Click **Save**

The application name appears in the User Data section.

Creating a simple Chargeback report

Chargeback reports allow administrators and managers to evaluate capacity usage by application, business entity, service level, and tier. The chargeback reports include capacity accountability, historical capacity accountability, and trending data. The data for these reports are built and scheduled from the OnCommand Insight Data Warehouse.

Before you begin

To create the sample report your system must be configured to report costs for storage tiers. The following tasks need to be completed:

- Define annotations for tiers.

- Assign costs to annotations.
- Define the applications you want to track data for.
- Assign applications to assets.

About this task

This example uses the Cognos Workspace Advanced reporting tool to create the Chargeback report. With Workspace Advanced, you create reports by dragging and dropping data elements into a report pallet.

Steps

1. In the OnCommand Insight web UI, click the reporting icon.
2. Log in to the Reporting Portal.
3. In the IBM Cognos Connection toolbar, click **Launch > Cognos Worksapce Advanced**

The Workspace Advanced package screen opens.

4. Click **Packages > Chargeback**

The IBM Workspace Advanace screen is displayed.

5. Click **New**
6. In the **New** report dialog Click **List** to specify a list report.

The report palette is displayed and the Chargeback “Simple data mart” and “Advanced data mart” are displayed under the Source heading.

7. Click the arrows next to each data mart to expand them.

The full contents of the data marts are displayed.

8. Drag “Application” from the “Simple Data Mart” into the far left column of the report palette.

When you drag an item into the palette the column shrinks and is highlighted. Dropping the application data into the highlighted columns results in all of the applications being listed correctly in the column.

9. Drag “Tier” from the “Simple Data Mart” into the next column of the report palette.

The storage tier associated with each application is added to the palette.

10. Drag “Tier Cost” from the “Simple Data Mart” into the next column of the report palette.
11. Drag “Provisioned capacity” from the “Simple Data Mart” into the next column of the report palette.
12. Hold down the **Ctrl** key and select the “Tier cost” and “Provisioned capacity” columns in the pallet.
13. Right-click the mouse in either of the selected columns.
14. Click **Calculate > Tier cost * Provisioned capacity DB**

A new column is added to the pallet with the title “Tier Cost * Provision Capacity GB”.

15. Right click the **Tier Cost * Provision Capacity GB** column.
16. Click **Style > Data Type**

17. Click **Format type > Currency**

18. Click **OK**

The column data is now formatted as US currency.

19. Right click “Tier Cost * Provision Capacity GB” and select **Edit Data Item Label**

20. Replace the Name field with “Provisioned Capacity Cost”

21. To run the report, click **Run > Run report - HTML**

A report similar to the following is displayed.

Application	Service Level	Service Level Cost	Tier	Tier Cost	Provisioned Capacity GB	Provisioned Capacity Cost
APAC Commercial Sales	Gold-Fast	12	Gold-Fast	12	674.04	\$8,088.42
APAC Commercial Sales	Silver	10	Silver	7	1,903.83	\$13,326.82
APAC Cruises	Gold-Fast	12	Gold-Fast	12	730.20	\$8,762.44
African Tours	Gold	12	Gold	10	4,856.12	\$48,561.16
African Tours	Silver	10	Silver	7	1,480.85	\$10,365.93
CRM	Bronze	3	Bronze	3	5,689.08	\$17,067.23
Caribbean	Gold	12	Gold	10	4,590.41	\$45,904.08
Commercial Applications	Bronze	3	Bronze	3	14,312.88	\$42,938.64
Commercial Applications	Gold-Fast	12	Gold-Fast	12	40,308.42	\$483,701.05
Commercial Environments	Bronze	3	Bronze	3	16,812.27	\$50,436.81
Commercial Environments	Gold	12	Gold	10	9,313.51	\$93,135.13
Commercial Environments	Silver	10	Silver	7	1,480.79	\$10,365.54
Concur	Gold	12	Gold	10	247.39	\$2,473.91
Concur	Gold-Fast	12	Gold-Fast	12	575.17	\$6,902.09
Consumer Feedback	Gold	12	Gold	10	1,335.89	\$13,358.94

Ensuring IO density reports describe only internal data volumes

In NetApp storage systems the root aggregate contains the root volume. The root volume contains special directories and configuration files for managing and controlling the storage system. The management and control operations might result a large amount of activity in the root aggregate. When you query the Insight system for the top 10 internal volumes with the highest IO density, your results might include NetApp root aggregates as members of the top 10.

When monitoring your environment it is more important to determine which internal data volumes are producing high I/O density numbers. In order to accurately identify only the data volumes, you need to isolate the NetApp internal volumes from queries you use to monitor I/O density.

This guide describes how to easily identify the NetApp root aggregates, isolate them from the results of internal volume queries, and create rules that exclude any new NetApp root aggregates as they are added to the system. The following Insight features are used to insure that your I/O density reports are derived from internal data volumes.

- A query is created to identify all of the NetApp root aggregates that are monitored by Insight.

- An annotation is assigned to each of the NetApp root aggregates.
- An annotation rule is created to exclude the NetApp aggregates

Creating a query to identify NetApp root aggregates in your environment

Queries provide search at a granular level, based on user-selected criteria. Using a query allows you to search for Internal volumes in your environment that contain the NetApp root aggregate.

Steps

1. In the OnCommand Insight web UI, create a query to identify NetApp root aggregates in your environment: **Queries > New Query > Select Resource Type**
2. Click **Storage Pool**
3. Enter the name for the root aggregate

This example uses “aggr0” for the name. When creating an aggregate, only the following requirements for the name must be followed:

- It must begin with either a letter or an underscore (_).
 - It can contain only letters, digits, and underscores.
 - It can be 250 characters or less. In most cases the aggregate is name aggr0, aggr_0, or something similar. It might require an iterative process to identify all of the NetApp root aggregates in your environment.
4. Click **Save** and enter a name for the new query.

As mentioned before, this might be an iterative process and require multiple queries to identify all of the NetApp root aggregates.

Create an annotation for the root volumes returned by your queries

Annotations are specialized notes that you assign to your assets, allowing you to filter assets by their annotations. The annotation you create will be used to identify the NetApp root aggregates in your environment and ensure that they are not included in a specific report.

Before you begin

You must have identified all of the root aggregates you want to exclude from the “High I/O Density” report.

Steps

1. Create an annotation to associate all of the NetApp root aggregates you identified with queries: **Manage > Annotations**
2. Click **Add**
 - a. Enter the name for the annotation: **RootAggr**
 - b. Enter a description of the annotation: **Remove root aggregate from "High I/O Density" report**

c. Enter the type of annotation: **Boolean**

3. Click **Save**

Create an annotation rule to automate excluding specific aggregates from your I/O density report

As an alternative to manually applying annotations to individual assets, you can automatically apply annotations to multiple assets using annotation rules. Annotation rules are based on queries you create and when run on the system they add new assets to existing sets of assets. When these sets of assets are excluded from a report the new assets are automatically excluded too.

Before you begin

You must have created and saved a query that identifies the NetApp root aggregates you identified in your environment.

Steps

1. Log in to the OnCommand Insight web UI.
2. Click **Manage > Annotation rules**
3. Click **Add**

The Add Rule dialog box displays.

4. Do the following:
 - a. In the Name box, enter a unique name that describes the rule: "RootAggrExclude"
 - b. Click Query and select the Query that Insight should use to apply the annotation rule to: "Aggregate0"
 - c. Click Annotation and select: "Root agg exclude"
 - d. Click Value and enter True

Collecting integration data

You can import integration data into your OnCommand insight system. Data can be imported using collectd, open source software that runs as a daemon to collect performance data, or by using the integration SNMP data source which allows you to collect generic SNMP data.

Data flow for integration data

The following applies to the total amount of integration data that is allowed to be presented to the OnCommand Insight server:

- A queue of 100 calls is maintained.

When a client waits in the queue for more than one minute, a timeout error occurs.

- The recommended ingestion rate for integration data is once per minute, per client.

- There is a limit of 300 integration object types allowed.

Accessing collectd software and documentation

You can access the output writer plugin software and documentation for collectd at NetApp's GitHub site:
https://github.com/NetApp/OCI_collectd

Backup and restore of integration data

Backup and restore of integration data is modeled after OnCommand Insight performance data backup and restore policies. When a backup is configured for performance data, integration data is also included in the backup. As with performance backup, the most recent seven days of integration data is included in the backup. Any integration data that is present in a backup is restored in a restore operation.

Licenses

A Perform license is required for integration data to be reported. If a Perform license is not present an error occurs with a message "Perform license required to report integration data".

Collecting SNMP Integration data

The integration SNMP data source allows you to collect generic SNMP data in OnCommand Insight.

Integration packs

The SNMP Integration data source uses an "Integration Pack" to define what integration values are collected, and what SNMP objects provide those values.

An Integration Pack consists of:

- A JSON configuration file (integration.json) defining integration payload contents in terms of SNMP objects of a specific device type (switch, router, and so on).
- A list of MIB files that the integration pack depends on.

An integration pack can define multiple data types. For example, when integrating an RHEL host, a data type can be defined for the general system information such as uptime, number of users, and number of running processes, a second data type can be defined for data on memory and file system usage. In general, each data type must be "flat" and cannot contain nested data.

A single integration pack should not define more than 24 data types. Insight limits the amount of integration data that is collected. Attempting to ingest more than 24 reports over a period of one minute results in a rate error.


The names for integration types must adhere to the following rules:

- The name cannot start with the following characters: `_`, `-`, or `+`
- The name cannot contain the following characters: `#`, `\`, `/`, `*`, `?`, `"`, `<`, `>`, `|`, `'`, ```,
- Cannot be longer than 100 UTF-8 encoded bytes
- Cannot be named `.` or `..`

Integration file format

An integration pack is a ZIP file that contains a JSON configuration file (integration.json) defining the integration payload contents in terms of SNMP objects. It also contains a MIBS folder that contains all of the MIB files and their MIB dependencies.

The integration.json file must exist at the top level of the ZIP file and the MIB files must exist in the "resources/mibs" subdirectory within the ZIP. The ZIP file may also contain files, such as a "readme.txt", if desired. An example of integration ZIP structure is:



Importing SNMP integration packs

You import SNMP integration packs into OnCommand Insight using the web UI. Integration packs are identified by the "integrationPacName" value defined in the integration.json configuration file contained in the ZIP file.

Before you begin

You must have created a properly formatted ZIP file that contains the integration pack you want to import to the OnCommand Insight server.

About this task



Use the following steps to import SNMP integration packs to the Insight server.

Steps

1. Click **Admin > Setup > SNMP Integration**

The system displays the Import SNMP package screen:

Import SNMP package

 Select file	No file selected	 Import
Warning: This will overwrite any conflicting package from existing database.		

2. Click **Select file** to select the local file containing the SNMP package.

The file you select is displayed in the File box.



Any existing integration pack with the same name is overwritten.

3. Click **Import**

The file is imported to the Insight server.

Creating an SNMP integration data source

The Integration SNMP Data Source provides common SNMP configuration properties similar to other SNMP based data sources included with the OnCommand Insight data sources for Brocade and Cisco.

Before you begin

In order to successfully use the Integration SNMP data source to collect, the following must be true:

- You must have already imported an integration pack you will use for this SNMP data source.
- All target devices share the same credentials.
- All target devices implement the SNMP Objects referenced by the configured integration pack.

About this task

To create an SNMP Integration data source, choose vendor "Integration" and model "SNMP" in the data source creation wizard.

Steps

1. In the OnCommand Insight web UI, click **Admin > Data Sources**
2. Click **+Add**
3. Enter a name for the Data Source
4. For Vendor, select **Integration**
5. For Model, select **SNMP**

Add data source

Settings

*Name

Vendor

Integration

Model

SNMP

Where to run

local

What to collect

☒ Integration (BETA)

Configure

Configuration

Advanced configuration

Test

Cancel

Save

- For What to collect, check **Integration**

This is the only package on this data source and is checked by default:

- Click **Configuration**
- Enter the IP addresses for the systems from which you will collect SNMP data
- Select an imported SNMP Integration Pack
- Set the integration poll interval
- Select the SNMP version
- Enter the SNMP community string

For SNMP V1 and V2.

- Add the user name and password for systems you will be collecting data from.

For SNMP V3.

- Click **Advanced Configuration**

The Advanced Configuration default settings are displayed. Make any changes to these settings that are required.

Integration.json file information

The integration.json file identifies the payload .

The following illustration provides a color-coded representation of a simple integration.json file. The accompanying table identifies the function of the objects in the file.

```
{
  "integrationPackName": "WindowsSnmp",
  "description": "Generic integration for mibs supported by the default
SNMP Agent for Windows 2012, including HOST-RESOURCES",
  "acquisitionType": "SNMP",
  "integrationTypes": [
    {
      "integrationType": "snmp_win2012_host",
      "name": {
        "mibModuleName": "RFC1213-MIB",
        "objectName": "sysName"
      },
      "identifiers": {
        "hostname": {
          "mibModuleName": "RFC1213-MIB",
        }
      },
      "attributes": {
        "description": {
          "mibModuleName": "RFC1213-MIB",
          "objectName": "sysDescr"
        },
        "snmp_sys_obj_id": {
          "mibModuleName": "RFC1213-MIB",
          "objectName": "sysObjectID"
        }
      },
      "dataPoints": {
        "uptime": {
          "num": {
            "mibModuleName": "RFC1213-MIB",
            "objectName": "sysUpTime"
          }
        }
      }
    }
  ]
}
```

Blue	Reserved
Red	User customizable strings and IDs
Green	MIB names
Purple	MIB object
Black	JSON structure

About integration.json files

Each field has the following characteristics:

- The "identifiers" section forms a unique compound key to create a new "object" in Insight
- The "attributes" provide supporting meta-data about the object.

In both of these cases, only the value of the latest report for that object (identified by the identifiers) is preserved.

- The "dataPoints" are time-series data and must be numeric values. Insight keeps each and every value reported here for 90 days (by default) and links them time-series to the object identified.

Numeric Expressions

By default, all value expressions are reported as strings in the integration payload. "identifiers" and "attributes" may only define string values. "dataPoints" may define string or numeric values. Numeric values are defined using one of the following modifier keys:

- num - the total number of bytes received since the counter was last initialized
- delta - the number of bytes received during the poll interval
- rate - the average receive rate during the poll interval in bytes per second

An average receive rate during the poll interval in megabytes per second can be accomplished using a combination of rate and math operations

Math operations

The `integration.json` file supports the following math operations: add, subtract, multiply, divide. The following example shows multiplication, division, and sum operations in a JSON file.


```

"network_utilization":
{
  "mult": [
    {
      "div": [
        {
          "sum": [
            "rate": {
              "mibModuleName": "IF-MIB",
              "objectName": "ifHCOutOctets",
              "comment": "bytes per second out"
            },
            "rate": {
              "mibModuleName": "IF-MIB",
              "objectName": "ifHCInOctets",
              "comment": "bytes per second in"
            }
          ]
        },
        {
          "num": {
            "mibModuleName": "IF-MIB",
            "objectName": "ifSpeed",
            "comment": "1,000,000 bits per second"
          }
        }
      ]
    },
    {
      "const": 0.0008,
      "comment": "normalize to ratio of bits and convert to percent:
8 * 100 / 1,000,000 = 0.0008"
    }
  ]
}

```

Keywords

An integration pack keyword, string, is implemented to force OCTET STRINGS or proprietary types derived from OCTET STRING that would normally be rendered in hexadecimal format to instead be rendered as ASCII characters.

Often OCTET STRINGS contain binary data, for example MAC addresses and WWNs:

```

"interface_mac": {
  "mibModuleName": "IF-MIB",
  "objectName": "ifPhysAddress"
}

```

ifPhysAddress is type PhysAddress, which is just an OCTET STRING:

```

PhysAddress ::= TEXTUAL-CONVENTION
    DISPLAY-HINT "1x:"
    STATUS      current
    DESCRIPTION
        "Represents media- or physical-level
addresses."
    SYNTAX      OCTET STRING

```

When ifPhysAddress is rendered as hex by default, the result is:

```
"interface_mac": "00:50:56:A2:07:E7"
```

However if you have an OCTET STRING or proprietary type derived from OCTET STRING that you want to interpret as ASCII, you can use the "string" keyword:

```

"string_test_1": {
    "string": {
        "mibModuleName":      "IF-MIB",
        "objectName":         "ifPhysAddress"
    }
},

"string_test_2": {
    "string": [
        {
            "mibModuleName":      "IF-MIB",
            "objectName":         "ifPhysAddress"
        },
        {
            "const": "JSD"
        },
        {
            "mibModuleName":      "IF-MIB",
            "objectName":         "ifPhysAddress"
        }
    ]
}

```

The keyword follows the existing string concatenation rules, inserting a single space between terms in the following example:

```
"string_test_1": "PVçç",  
  "string_test_2": "PVçç JSD PVçç"
```

The "string" keyword acts on a single term or a list of terms, but not nested expressions. Nested expressions are only supported for dataPoint expressions. Attempting to use a "string" expression in a dataPoint expression will result in an error similar to the following:

```
java.lang.IllegalArgumentException: Integration pack 'GenericSwitch32' index 'snmp_generic_interface_32'  
section 'dataPoints' key 'string_test_3' unsupported JSON numeric expression  
'{"string":{"mibModuleName":"IF-MIB","objectName":"ifPhysAddress"}}'
```

Some derived OCTET STRING types such as DisplayString, SnmpAdminString have hard-coded precedence over the "string" keyword. This is because SnmpAdminString is specifically UTF-8 encoded, and we want to handle it correctly, whereas the "string" keyword forces the default string representation returned by the snmp_framework, which assumes single byte ascii code points per character.

Analyzing an application performance problem

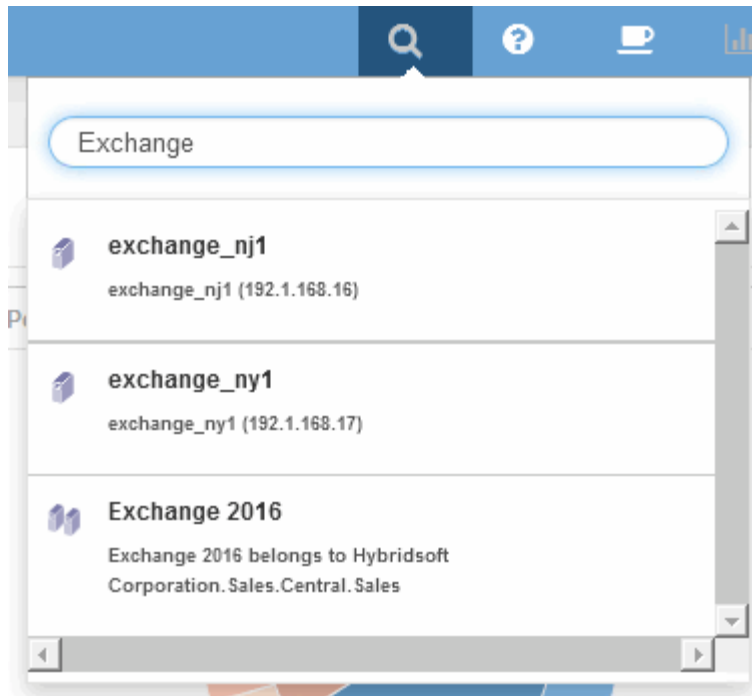
This document describes steps you might take to address reports of performance problems for an application that are impacting users or administrators. For example, users are complaining that their Exchange application is experiencing periods of slowness throughout the day.

About this task

In OnCommand Insight, an application is a configured entity. You assign a name and business entity to the application and you assign compute and storage resources to the application. This allows a better end-to-end view of infrastructure health and more pro-active management of infrastructure asset management.

Steps

1. To begin investigating the issue, use the Insight toolbar to perform a global search for the Exchange application.



When performing a search, you can add an object descriptor before the object name to refine the search results.

- When you select "Exchange 2016" from the search results, the system displays the Application landing page.

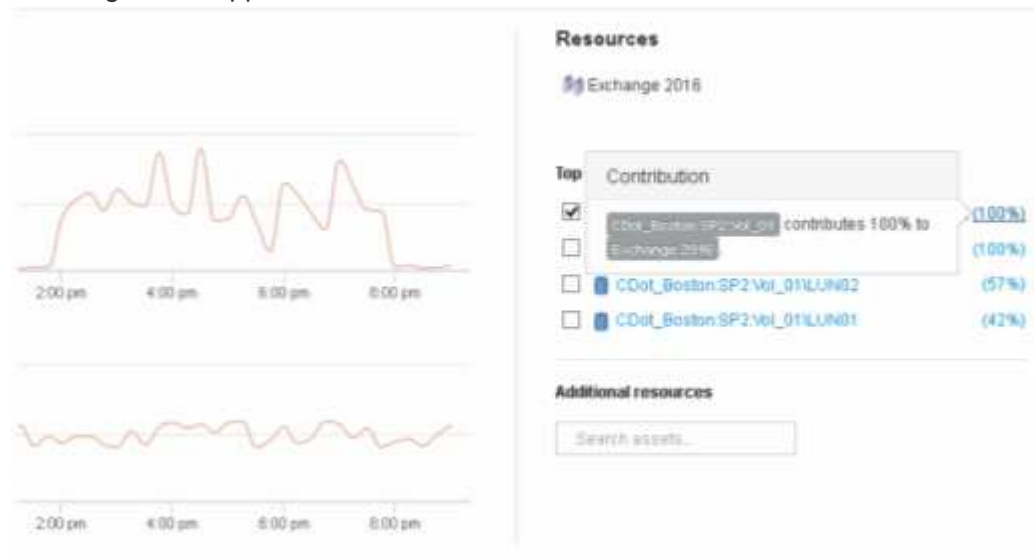


In the Application landing page, the following information is of interest:

- In the 24-hour time period selected, an increase in latency is shown on the right of the latency graph.
- During the period of increased latency there is no significant change in the level of IOPS. It appears the latency increase is not caused by a heavier application usage. We are not really seeing a high IOPS

demand on the storage that could account for the latency spike. The increase in latency could be due to an external factor.

- On the right of the charts in the Top contributors section, click on the 100% for the selected internal volume (CDot_Boston:SP2:Vol_01). The system shows this resource is contributing 100% to the Exchange 2016 application.



- Click on the navigation link for this internal volume (CDot_Boston:SP2:Vol_01) to access the internal volume landing page. Analysis of the internal volume might provide information pertaining to the latency spike.

Examining the internal volume



In the Internal Volume landing page, you see:

- The performance charts for the internal volume match what was previously seen for the application performance for both latency and IOPS.
- In the Resources section, where the correlated assets are displayed, a “Greedy” resource is identified (CDot_Boston:SP1:Vol_01).

A greedy resource is identified by insight correlation analytics. Greedy/degraded resources are “peers” that utilize the same shared resource. The greedy resource has IOPS or utilization rates that negatively impact the degraded resource’s IOPS or latency.

Greedy and Degraded resources can be identified on Virtual Machine, Volume, and Internal Volume landing pages. A maximum of two greedy resources will be displayed on each landing page.

Selecting the correlation ranking (%) provides the Greedy resource analysis findings. For example, clicking a greedy percentage value identifies the operation on an asset that impacts the operation on the Degraded asset, similar to what is shown in the following example.

Resources

- CDot_Boston:SP1:Vol... (98%)

Top correlated

- VM_Exchange_1 (98%)
- CDot_Boston_N1 (85%)

Greedy

- CDot_Boston:SP1:Vol... (98%)

Resources

- hionpcmsac...4_prd_cl05

Greedy

IOPS of CDot_Boston:SP1:Vol... impacts Latency of CDot_Boston:SP1:Vol... by 98%. (98%)

When a degraded resource is identified, you can select the degraded (%) score to identify the operation and the resource that is impacting the degraded resource.

Resources

- CDot_Boston:SP2:Vol... (98%)

Top correlated

- VM_Cs_travBook (99%)
- CDot_Boston:SP1 (56%)

Degraded

- CDot_Boston:SP2:Vol... (98%)

Additional resources

Search assets...

Resources

- hionpcmsac...p13_splunk

Top correlated

- hionpcmsac...01n01b:...sac...01n01b_ex...

Degraded

- hionpcmsac...01:svmn...170_vmdk04_p... (88%)
- hionpcmsac...01:svmn...180_vmdk04_p... (40%)


Degraded

IOPS of hionpcmsac...p13_splunk impacts Latency of hionpcmsac...4_prd_cl03 by 88%. (88%) (40%)

Examining the greedy resource

Clicking on the internal volume identified as the greedy resource opens the landing page for the volume CDot_Boston:SP1:Vol_01.

Note in the summary details this internal volume is a resource for a different application (Travel Booking) and although contained in a different storage pool is on the same node as the internal volume for Exchange 2016 (CDot_Boston_N1)



The landing page shows:

- The internal volume associated with a Travel Booking application.
- A new storage pool is identified in the correlated resources.
- The original internal volume you were examining (CDot_Boston:SP2:Vol_01) is identified as “Degraded”.
- In the performance graph, the application has a steady latency profile and does have an IOPS spike roughly at the same time we see the latency spike on the Exchange application.

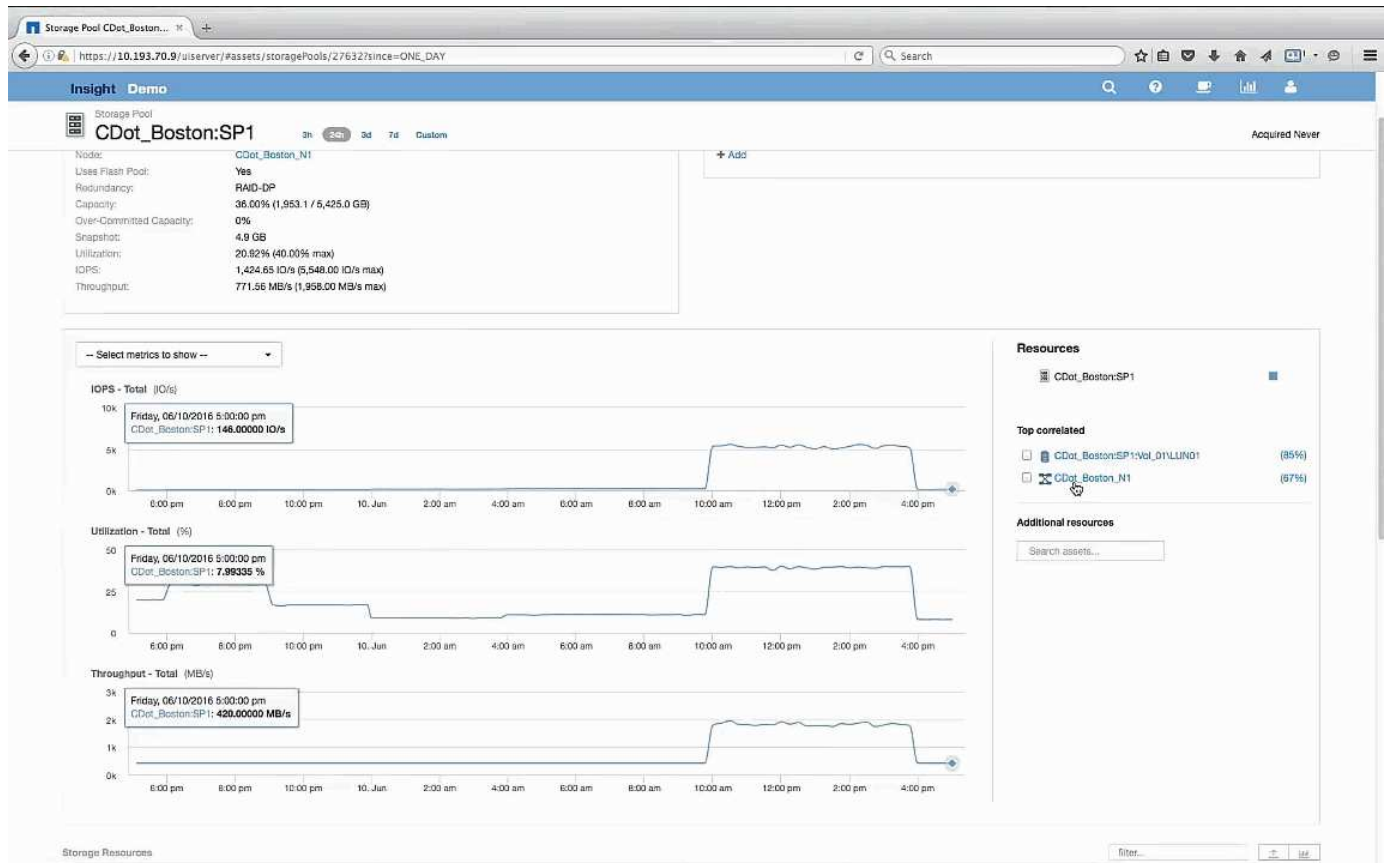
This might indicate that the latency spike on the Exchange application is likely caused by the IOPS spike on this volume.

To the right of the charts in the Resource section notice the correlated Degraded resource which is the Exchange 2016 internal volume (CDot_Boston:SP2:Vol_01). Click on the check box to include the degraded internal volume in the in the performance graphs. Aligning the two performance graphs shows that the latency and IOPS spikes occur at nearly the exact same time. This tells us that we want to get a better understanding of the Travel Booking application. We need to understand why the application is experiencing such a prolonged IOPS spike.

Examining the Storage pool associated with the Travel Booking application might identify why the application is experiencing the IOPS spike. Click CDot_Boston:SP1 to view the Storage Pool landing page.


Examine the storage pool

Examining the storage pool landing page shows the same IOPS spike seen in its correlated assets. In the Resources section you can see that this storage pool landing page links to the volume of the travel application. Click on the volume to open the volume landing page.



Examining the volume

The volume landing page shows the same familiar IOPS spike seen in its correlated assets.




In the resources section the VM for the Travel Booking application is identified. Click on the VM link to view the VM landing page.

Examining the VM

In the VM landing page, select additional metrics to display and include CPU utilization and Memory utilization. The graphs for CPU and Memory utilization show that both are operating at nearly 100% of their capacity. This tells us that the problem with the Exchange server is not a storage problem, but instead is the result of the high VM CPU and memory utilization and the consequential memory swapping of I/O to disk.



To solve this problem, you can look for additional similar resources. Enter “Node” in the Additional resources input dialog to show metrics for assets similar to the Exchange VM. The comparison can help identify a node that might be a better fit for hosting the workload should a change be necessary.



Collecting and reporting AWS billing data

The Amazon AWS Cloud Cost data source imports billing data generated by Amazon into Insight as integration data, making it available to the data warehouse for reporting.

There are three parts to making cloud billing data available to Insight:

Verifying your AWS account information.

Configuring the AWS Cloud Cost data source in Insight to collect the data.

Sending the data to Data Warehouse via ETL for use in reports.

Preparing AWS for Insight data collection

Your AWS account must be properly configured to allow Insight to collect cloud cost data.

About this task

The following steps are done through your AWS account. See the Amazon documentation for more information: <http://docs.aws.amazon.com>. If you are unfamiliar with setting up an AWS cloud account, contact your cloud provider for assistance.



These steps are provided here as a courtesy and are believed correct as of the time of publication. NetApp makes no guarantee of the correctness of these steps. Contact your cloud provider or AWS account holder for information or assistance on configuring your AWS account.

Best practice: Insight recommends that you create a primary IAM user on the same account that owns the S3 bucket where the billing reports are uploaded, and use this user to configure and collect AWS billing data.

To configure your AWS account to allow Insight to collect data, perform the following steps:

Steps

1. Log in to your AWS account as an Identity Access Management (IAM) user. For proper collection, log in to the primary IAM account, as opposed to a group IAM account.
2. Go to **Amazon S3** to create your bucket. Enter a unique bucket name and verify the correct Region.
3. Turn on your Amazon Cost and Usage Report. See <https://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/billing-reports-gettingstarted-turnonreports.html> for information.
 - a. Go to the **AWS Billing and Cost Management Dashboard** and choose **Reports**.
 - b. Click on **Create report** and enter in the Report Name. For **Time unit**, choose Daily. Check the box to include **Resource IDs**, and click **Next**.
 - c. Click on the **Sample Policy** link in the Select delivery options page. Copy the Sample Policy text in the box to the clipboard. Click **Close**.
 - d. Go back to the S3 Bucket that was created, click on the **Permissions** tab and select the **Bucket Policy** button.
 - e. Paste the text from the Sample Policy, and replace <bucketname> with your actual bucket name in each line that contains the following: "Resource": "arn:aws:s3:: <bucketname>". **Save** the policy.
 - f. Go back to your Create Report screen, enter in your S3 Bucket and click the **Verify** button. Click **Next**.
 - g. Verify your information and click **Review and Complete**.
4. You must grant permissions in order for Insight to collect data from AWS. The following link provides details on how to grant permissions to **List All Buckets** (Step 4.1) and set permissions on the objects in the folder (Step 5.2): <https://docs.aws.amazon.com/AmazonS3/latest/dev/walkthrough1.html>.

5. In the IAM console, go to **Policies** and click **Create policy**.
6. Enter a name in the **Policy Name** field, and click **Create policy** at the bottom.
7. In the IAM console, select your user, then select **Add Inline Policy** at the bottom of the screen.
8. Click on **Choose a service** and select S3.
9. Go to the **JSON** tab. Copy the JSON sample text from step 5.1.2.g of the AWS walkthrough into the JSON box.
10. Replace the *companybucket* and *Development* fields in the JSON with your S3 information.
11. Click **Review Policy** to review your policy settings.

Configuring the AWS Cloud Cost data source

You configure the AWS Cloud Cost data source as you would for any Insight data source.

Before you begin

You must have your Amazon AWS account already set up and prepared for Insight data collection, and have the following pieces of information at hand.

- Report Name
- S3 Bucket Name
- AWS Region where your S3 bucket resides.
- Report path prefix

About this task

Once your AWS account is ready and has the proper permissions set, you are ready to configure OnCommand Insight to collect billing report data.



You will need to add a separate AWS Cloud Cost data source for each billable user/account from which you wish to retrieve billing data.

Steps

1. Log in to OnCommand Insight as an administrator.
2. Click on **Admin > Data sources** to open the Insight Data Source page.
3. Click **+Add** to add a new data source. Choose **Amazon** and select **AWS Cloud Cost**.
4. In the **Configuration** section, fill in the *Report name*, *S3 Bucket name*, *S3 Region* (must be the region where your S3 bucket resides), *Report path prefix*, *AWS IAM Access Key ID*, and *AWS IAM Secret Access Key*. If you are unsure of any of these, check with your cloud provider or AWS account holder.
5. Click the checkbox to verify your understanding that AWS will bill you for API requests and data transfers made by the Insight data source.
6. In **Advanced Configuration**, enter the HTTP connection and socket timeout. The default is 300 seconds.
7. Click **Save**.

Processing AWS Cloud Cost data in Insight

Insight collects data from your AWS billing report once a month for the previous month, and reflects the finalized cloud cost for that month.

After you set up your AWS Cloud Cost data source(s), if you already had billing reports generated to S3, you will get up to three months of past data immediately after the first data source poll.

Insight collects AWS “final” data once a month. This collection occurs a few days after the close of the previous month, allowing AWS time to finalize the actual data.

AWS billing data is sent to Insight's Data Warehouse for use in reporting.

Keep in mind that each data source must be configured for a single billable account/user.

Reporting on Cloud Cost data in Insight

Cloud cost monthly data collected in Insight is sent to the data warehouse and is available in the Cloud Cost datamart for use in reports.

Before you begin

You must have data sources configured to collect cloud cost data from AWS. Each billable user/account must have a separate data source.

Allow Insight at least 36 hours to begin collecting data.


Allow ETL to run at least once after that time, to send the data to the data warehouse.

About this task

After your data has been collected and sent to the data warehouse, you can view it in any of several pre-configured reports, or create custom reports. Insight stores the data in its own Cloud Cost datamart.

To view your cloud cost data in one of the pre-configured reports:

Steps

1. Open Insight Reporting by one of these methods:
 - Click on the Reporting Portal icon  in the Insight server web UI or in the Data Warehouse UI.
 - Launch Reporting directly by entering the following URL: https://<dw_server_name>:9300/p2pd/servlet/dispatch or https://<dw_server_name>:9300/bi (7.3.3 and later)
2. Once you are logged in to Reporting, click on **Public Folders** and select **Cloud Cost**.
3. You can view your AWS billing data in the available reports located in the **Cloud Cost** folder, or create your own custom report using the **Cloud Cost datamart** available from the **Packages** folder.

Integrating with ServiceNow

OnCommand Insight integrates with ServiceNow management software to provide greater value than the products have separately.

Using a Python script, Insight can integrate data with ServiceNow, synchronizing the following information:

- Storage asset data for ServiceNow servers
- Host and VM URLs for ServiceNow servers
- Relationships between Hosts/VMs and Storage

Preparation and prerequisites for Service Now integration

The necessary preparations and prerequisites must be satisfied for ServiceNow, Insight, and the Python middleware connector prior to integration.

Recommended workflow

The following workflow is strongly recommended when integrating ServiceNow with Insight:

1. Deploy the Python middleware connector in your development instance first.
2. Once you have confirmed all faults have been identified and corrected in your development instance, deploy the connector in your test/stage instance.
3. Once you have confirmed correct operation in your staging instance, deploy the connector in your production instance.

If problems are found during any of these stages, follow your rollback steps and disable the connector, then troubleshoot the problem and re-deploy.

General prerequisites:

- You can use either a standalone host or VM (recommended) or the Insight server host/VM to host the python middleware connector.
- It is highly recommended to backup the production Insight server and deploy it on a development instance.
- ServiceNow must be accurately discovering servers in the CMDB.
- Insight must be accurately discovering your storage and compute environments.
- Port 443 and 80 to the Insight Server and ServiceNow Instance.

ServiceNow prerequisites:

- It is highly recommended to use a development/test instance.
- Permission to load ServiceNow update sets.
- Permission to create users.
- ServiceNow version Jakarta or later

Insight prerequisites:

- It is highly recommended to use a development/test instance.
- Permission to create users (Admin permissions).
- Insight version 7.3.1 or later is supported, but to get the most out of Insight, use the latest version.

Python middleware connector prerequisites:

- Python version 3.6 or greater installed.
- When installing Python, check the box to enable all users. This sets Python for standard application install locations.
- When installing Python, check the box to enable the installer to update the path. Otherwise, you will have to update the path manually.
- Download the Python **pysnow** and **requests** libraries.

Downloading the ServiceNow Python connector

You must download the Python connector for ServiceNow integration and extract it to a location of your choosing.

Steps

1. Download the **ServiceNow Integration connector** from the [NetApp Storefront](#).
2. Extract the .zip file to a folder, for example `c:\OCI2SNOW`.

The integration connector script is named `oci_snow_sync.pyz`.

Configuring ServiceNow for integration

Integrating ServiceNow with Insight requires several setup tasks.

About this task

The following tasks must be performed when integrating ServiceNow with Insight:

On the ServiceNow side:

- Elevate Role
- Install Update Sets
- Set up users

On the Insight side:

- Add the ServiceNow user

On the Python connector side:

- Install Python
- Install additional libraries
- Initialize the connector
- Edit the config.ini file
- Test the connector
- Synchronize the connector
- Schedule daily task execution

Each of these is explained in greater detail in the following sections.

Elevate role

You must elevate your ServiceNow role to `security_admin` before you can integrate with Insight.

Steps

1. Log into your ServiceNow instance with administrator permissions.
2. Under the **System Administrator** drop-down, choose **Elevate Roles** and elevate your role to `security_admin`. Click OK.

Install update set

As part of the integration between ServiceNow and OnCommand Insight you must install an Update Set, which loads pre-configured data into ServiceNow in order to provide the connector with specific fields and tables for extracting and loading data.

Steps

1. Navigate to the remote update sets table in ServiceNow by searching for “Retrieved update sets”.
2. Click on **Import Update Set from XML**.
3. The update set is in the Python connector .zip file previously downloaded to your local drive (in our example, the `C:\OCI2SNOW` folder) in the `\update_sets` sub-folder. Click on **Choose File** and select the .xml file in this folder. Click **Upload**.
4. Once the Update Set is loaded, open it and click on **Preview Update Set**.

If errors are detected, you must correct them before you can commit the Update Set.

5. If there are no errors, click **Commit Update Set**.

Once the Update Set has been committed it will show on the **System Update Sets > Update Sources** page.

ServiceNow integration - Set up user

You must set up a ServiceNow user for Insight to connect with and synchronize data.

About this task

Steps

1. Create a services account in ServiceNow. Login to ServiceNow and navigate to **system security > users and groups > users**. Click on **New**.
2. Enter a user name. In this example, we will use “OCI2SNOW” as our integration user. Enter a password for this user.



In this How-to we use a services account user named “OCI2SNOW” across the documentation. You may use a different services account, but be sure it is consistent across your environment.

3. Right-click on the menu bar and click **Save**. This will allow you to stay on this user in order to add roles.
4. Click **Edit** and add the following roles to this user:
 - asset
 - import_transformer
 - rest_service
5. Click **Save**.
6. This same user must be added to OnCommand Insight. Log in to Insight as a user with Administrator permissions.
7. Navigate to **Admin > Setup** and click on the **Users** tab.
8. Click the **Actions** button and select **Add user**.
9. For name, enter “OCI2SNOW”. If you used a different user name above, enter that name here. Enter the same password you used for the ServiceNow user above. You may leave the email field blank.
10. Assign this user the **User** role. Click **Save**.

Install Python and libraries

Python can be installed on the Insight server or on a standalone host or VM.

Steps

1. On your VM or host, download Python 3.6 or later.
2. Choose custom installation and choose the following options. These are either necessary for proper connector script operation or are highly recommended.
 - Install launcher for all users
 - Add Python to the PATH
 - Install pip (which allows Python to install other packages)
 - Install tk/tcl and IDLE
 - Install the Python test suite
 - Install py launcher for all users
 - Associate files with Python
 - Create shortcuts for installed applications
 - Add python to environment variables
 - Precompile standard library
3. After Python is installed, install the “requests” and “psnow” Python libraries. Run the following command:

```
python -m pip install requests pysnow
```

NOTE: This command might fail when you are operating in a proxy environment. To work around this issue, you need to manually download each one of the Python Libraries and run the install requests one by one and in the correct order.

The command will install several files.

4. Verify the Python libraries are installed correctly. Start Python using one of the following methods:
 - Open a cmd prompt and type `python`
 - On Windows, open **Start** and choose **Python > python-<version>.exe**
5. At the Python prompt, type `modules`

Python will ask you to wait a moment while it gathers a list of modules, which it will then display.

Setup Python middleware

Now that Python and the necessary libraries are installed, you can configure the middleware connector to communicate with OnCommand Insight and ServiceNow.

Steps

1. On the host or VM where you downloaded the connector software, open a cmd window as administrator and change to the `\OCI2SNOW\` folder.
2. You must initialize the script to generate an empty **config.ini** file. Run the following command:
`oci_snow_sync.pyz init`
3. Open the **config.ini** file in a text editor and make the following changes in the [OCI] section:
 - Set **url** to `<a href="https://<name.domain>" class="bare">https://<name.domain>` or `<a href="https://<ip" class="bare">https://<ip` address for the Insight instance.
 - Set **user** and **password** to the Insight user created, for example, OCI2SNOW.
 - Set **include_off_vms** to **false**
4. In the [SNOW] section, make the following changes:
 - Set **Instance** to the FQDN or ip address for your ServiceNow instance
 - Set **User** and **Password** to the ServiceNow service account user, for example, the OCI2SNOW.
 - Under **Field for the OCI URL**, set the **url** field to "u_oci_url". This field is created as part of the connector OCI update set. You can change this in the customer environment, but if you do so, you need to modify it here and in ServiceNow. Best practice is to leave this field as is.
 - Set the **filter_status** field to "Installed, In Stock". If you have a status that is different, you must set that status here in order to get all the records to match with Insight records prior to upload of new records. In most cases this field should remain unchanged.
 - Set **stale_status** to "Retired".
5. The [Proxy] section is only required if you use a proxy server. If you need to use this section, ensure the following settings:
 - `;https = <a href="http://<host>:<port>" class="bare">http://<host>:<port>`;
 - `;http = <a href="http://<host>:<port>" class="bare">http://<host>:<port>`;
 - `;include_oci = True`
 - `;include_snow = True`
6. Edit the [Log] section only if you need deeper debug information.

7. To test the connector, open a cmd prompt as administrator and change to the \OCI2SNOW folder. Run the following command: `oci_snow_sync.pyz test`

Details can be seen in the `logs\` folder.

Syncing the connector

Once ServiceNow, Insight and the connector are properly configured, you can synchronize the connector.

Steps

1. Open a cmd prompt and change to the \OCI2SNOW folder.
2. Run the following command twice. The first sync updates the items, the second sync updates the relationships: `oci_snow_sync.pyz sync`
3. Verify that the Storage Server table in your ServiceNow instance is populated. Open a storage server and verify that resources related to that storage are listed.

Scheduling synchronization to occur daily

You can use the Windows task scheduler to automatically sync the ServiceNow connector.

About this task

Automatic synchronization ensures that Insight data is regularly moved to ServiceNow. You can use any method for scheduling. The following steps use the Windows task scheduler to accomplish automatic syncing.

Steps

1. On the Windows screen, click **Start** and enter **run > task scheduler**.
2. Click **Create Basic Task...**
3. Enter a meaningful name, such as "OCI2SNOW Connector Sync". Enter a description of the task. Click **Next**.
4. Select to run the task **Daily**. Click **Next**.
5. Choose a time of day to run the task. Click **Next**.
6. For Action, select **Start a program**. Click **Next**.
7. In the **Program/script** field, enter `C:\OCI2SNOW\oci_snow_sync_pyz`. In the **Arguments** field, enter `sync`. In the **Start in** field, enter `C:\OCI2SNOW`. Click **next**.
8. Review the Summary details, and click **Finish**.

The synchronization is now scheduled to run daily.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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