Monitor, Alerts, and Notifications

Data Infrastructure Insights

NetApp
October 22, 2026

This PDF was generated from https://docs.netapp.com/us-en/data-infrastructure-insights/task_create_monitor.html on October 22, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Monitor, Alerts, and Notifications . . . 1
- Alerting with Monitors . . . 1
  - Security Best Practice . . . 1
  - Metric or Log Monitor? . . . 1
  - Monitor List . . . 9
  - Monitor Groups . . . 9
  - System-Defined Monitors . . . 12
- Viewing and Managing Alerts from Monitors . . . 12
  - Viewing and Managing Alerts . . . 12
  - Alert Detail Panel . . . 13
  - Alerts When Data Is Missing . . . 14
  - "Permanently Active" Alerts . . . 15
  - Maintenance Windows . . . 15
- Configuring Email Notifications . . . 16
  - Subscription Notification Recipients . . . 16
  - Global Recipient List for Alerts . . . 17
  - Editing Notifications for ONTAP . . . 17
- Anomaly Detection Monitors . . . 19
  - What is Anomaly Detection? . . . 19
  - When would I need Anomaly Detection? . . . 20
  - Creating an Anomaly Detection Monitor . . . 20
  - Viewing the Anomalies . . . 21
- System Monitors . . . 22
  - Monitor Descriptions . . . 23
  - More Information . . . 91
- Webhook Notifications . . . 91
  - Notification using Webhooks . . . 92
  - Webhook Example for Discord . . . 95
  - Webhook Example for PagerDuty . . . 97
  - Webhook Example for Slack . . . 100
  - Webhook Example for Microsoft Teams . . . 102
- Configuring a Login Message . . . 104

Monitor, Alerts, and Notifications

Alerting with Monitors

Configure monitors to track performance thresholds, log events, and anomalies across your infrastructure resources. Create custom alerts for metrics like node write latency, storage capacity, or application performance, and receive notifications when these conditions are met.

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

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

Data Infrastructure Insights provides a number of System-Defined Monitors as well, based on your environment.

Security Best Practice

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

Metric or Log Monitor?

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

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

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

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

- *Metric* monitors alert on infrastructure- or performance-related triggers
- *Log* monitors alert on log-related activity
- *Anomaly Detection* monitors alert on changes in the patterns of your data
- *Attribute* monitors alert on changes to an object's attributes, such as moving an object to a new data center

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

Metric Monitor

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

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

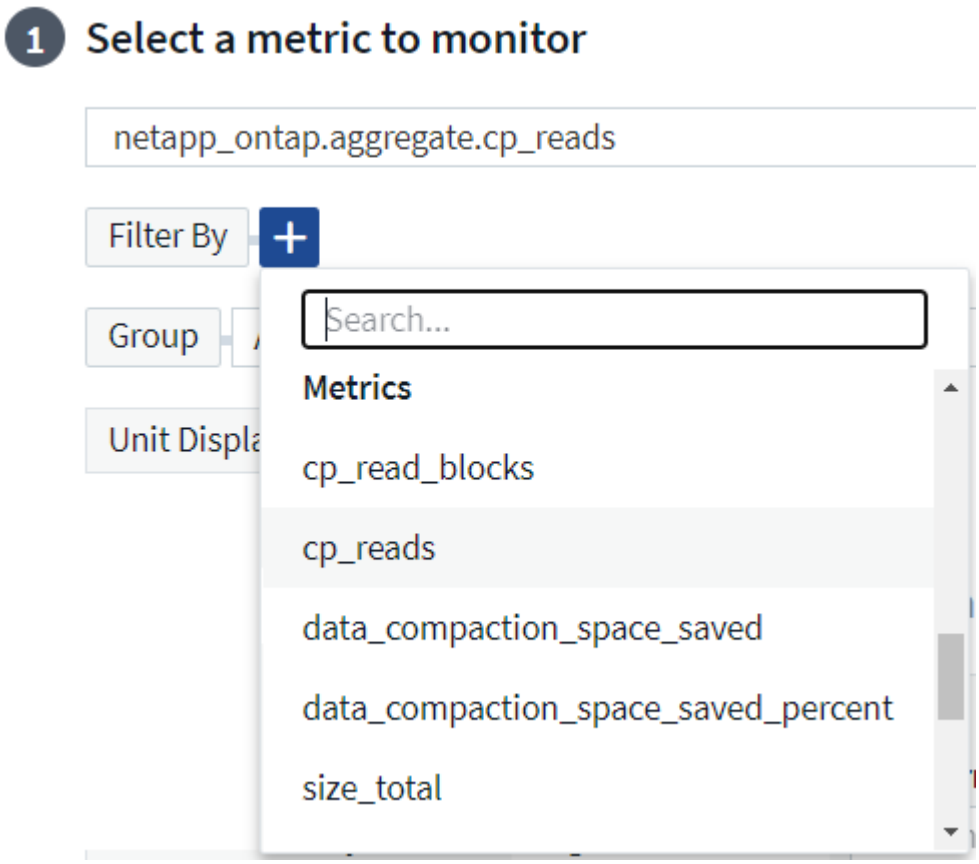

When working with integration data (Kubernetes, ONTAP Advanced Data, etc.), metric filtering removes the individual/unmatched data points from the plotted data series, unlike infrastructure data (storage, VM, ports etc.) where filters work on the aggregated value of the data series and potentially remove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

Metric monitors apply to Inventory objects such as storage, switch, host, vm, etc., as well as integration metrics such as ONTAP Advanced or Kubernetes data. When monitoring inventory objects, note that you cannot select a "Group By" method. However, grouping *is* allowed when monitoring integration data.

Multi-Condition monitors

You may choose to further refine your metric monitor by adding a second condition. Simply expand the "+Add Secondary Metric Condition" prompt and configure the additional condition.

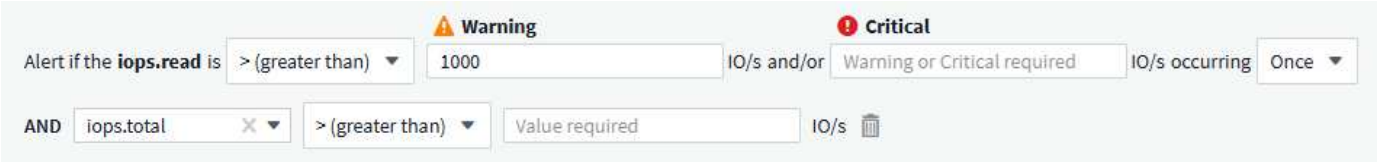

The monitor will alert if both conditions are met.

Note that you can only "AND" a second condition; you cannot choose to alert on one condition OR the other.

Define the Conditions of the Monitor.

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

the example graph.

3. For the *Critical* level, enter 400. The dashed line indicating this Critical level displays in the example graph.

   The graph displays historical data. The Warning and Critical level lines on the graph are a visual representation of the Monitor, so you can easily see when the Monitor might trigger an alert in each case.

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

   You can choose to trigger an alert the moment a threshold is breached, or wait until the threshold has been in continuous breach for a period of time. In our example, we do not want to be alerted every time the Total IOPS peaks above the Warning or Critical level, but only when a monitored object continuously exceeds one of these levels for at least 15 minutes.

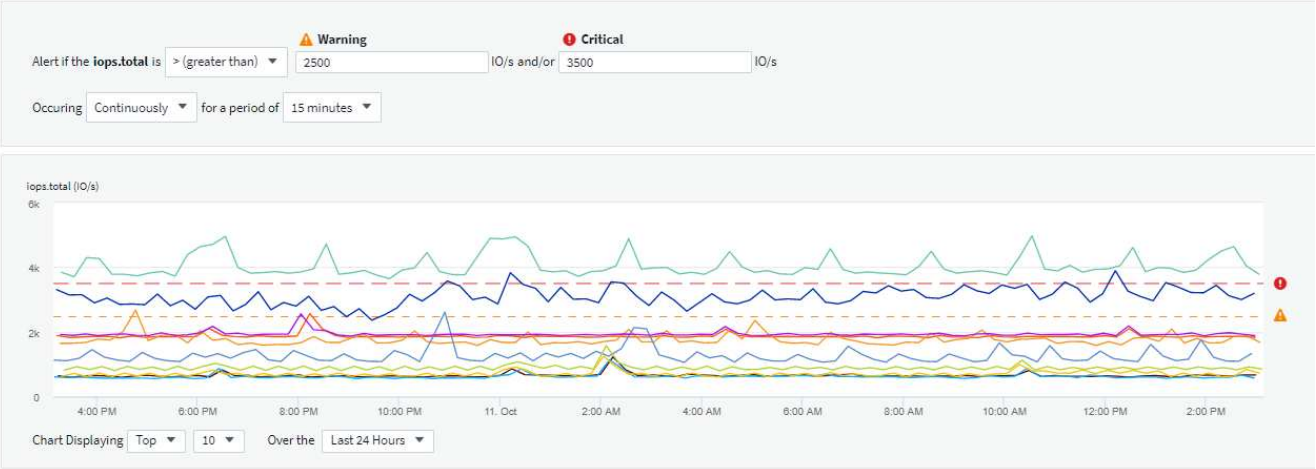

Define the alert resolution behavior

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

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

Log Monitor

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

> The Log Monitor filter cannot be empty.

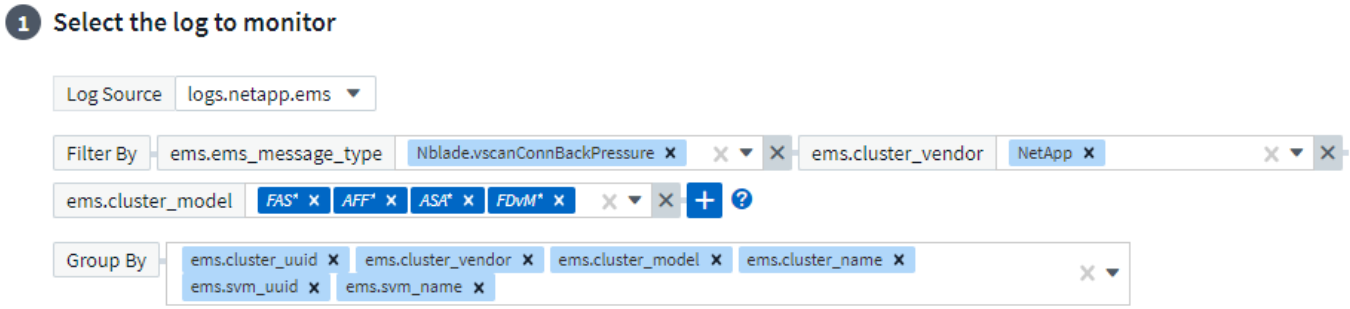

Define the alert Behavior

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

Define the alert resolution behavior

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

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

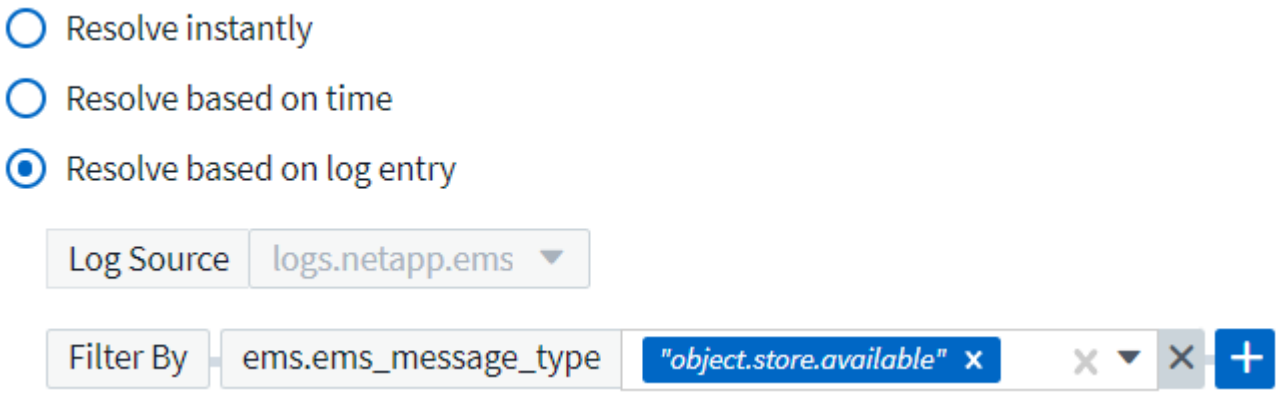

Anomaly Detection Monitor

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

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

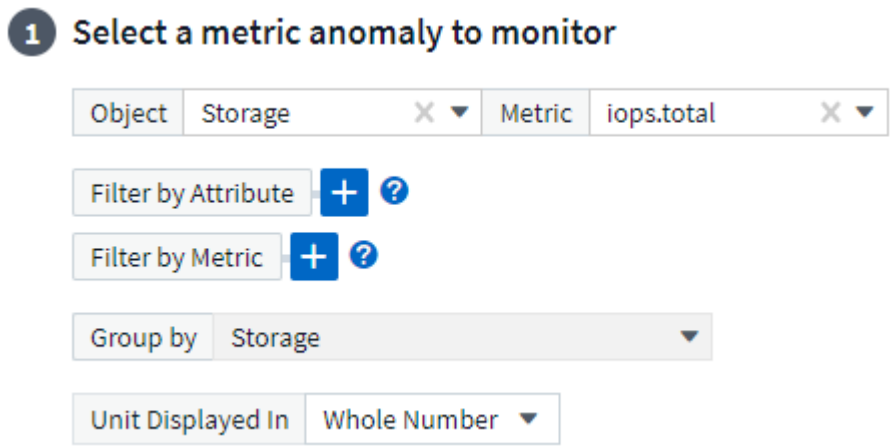

Define the Conditions of the Monitor.

1. After choosing the object and metric to monitor, yous et the conditions under which an anomaly is detected.
    - Choose whether to detect an anomaly when the chosen metric **spikes above** the predicted bounds,

drops below those bounds, or **spikes above or drops below** the bounds.

- Set the **sensitivity** of detection. **Low** (fewer anomalies are detcted), **Medium**, or **High** (more anomalies are detected).
- Set the alerts to be wither **Warning** or **Critical**.
- If desired, you can choose to reduce noise, ignoring anomalies when the chosen metric is below a threshold that you set.

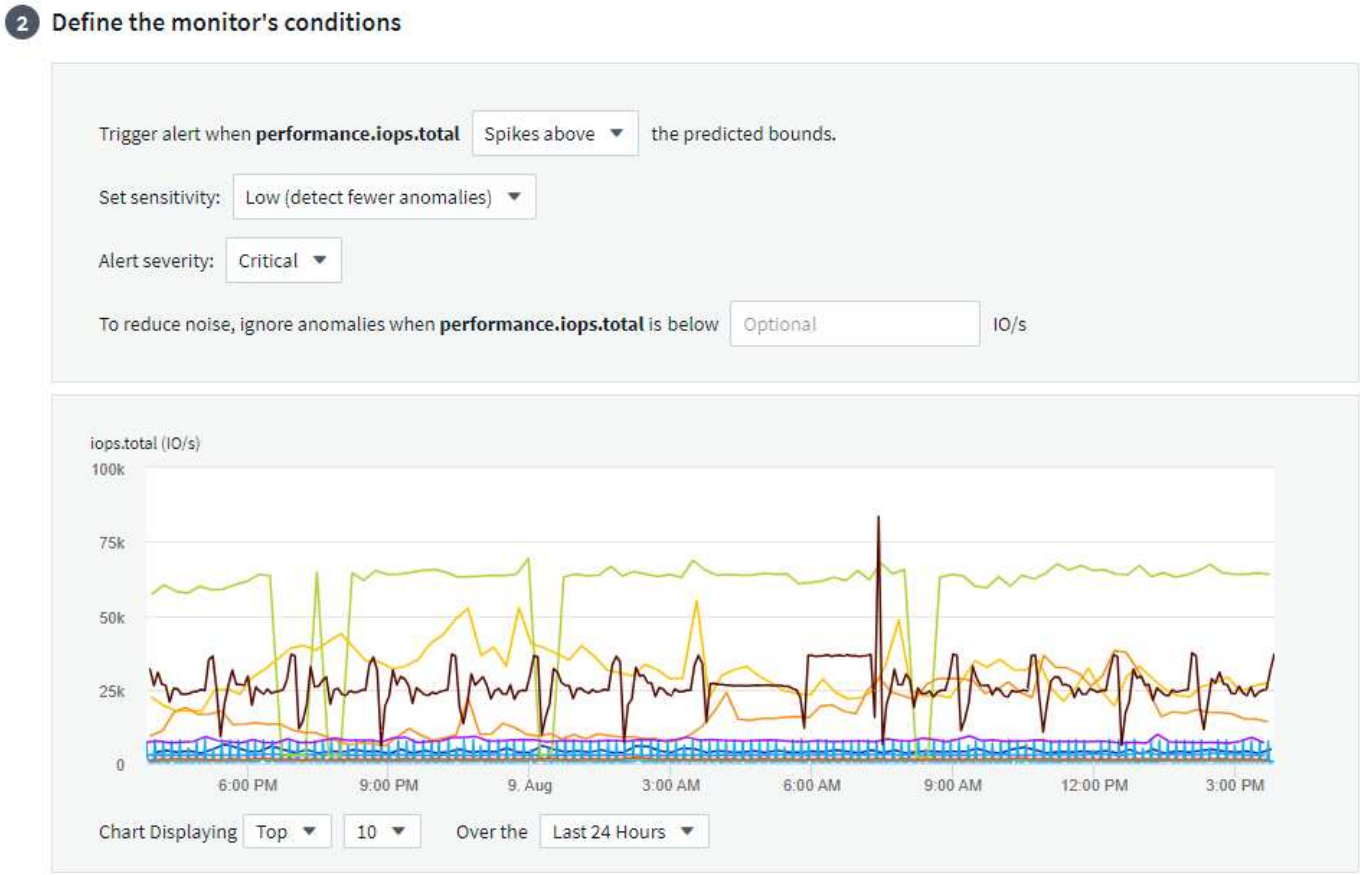

Attribute Monitor

Attribute Monitors automatically alert your team when assets move between data centers, a storage pool is not encrypted, or any attribute modifications occur. This helps to maintain configuration baselines and helps ensure nothing slips through the cracks during new deployments or changes.

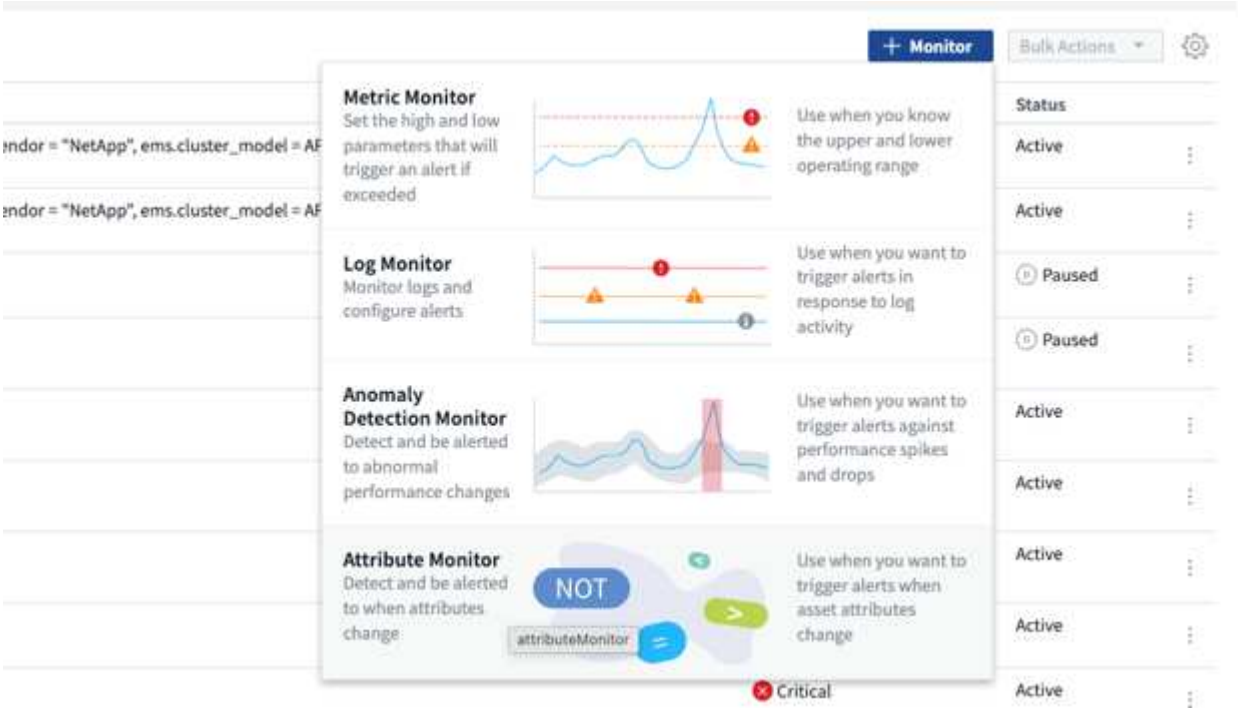

Simply choose the object type and select which attributes you want to monitor for changes.

Select notification type and recipients

In the *Set up team notification(s)* section, you can choose whether to alert your team via email or Webhook.

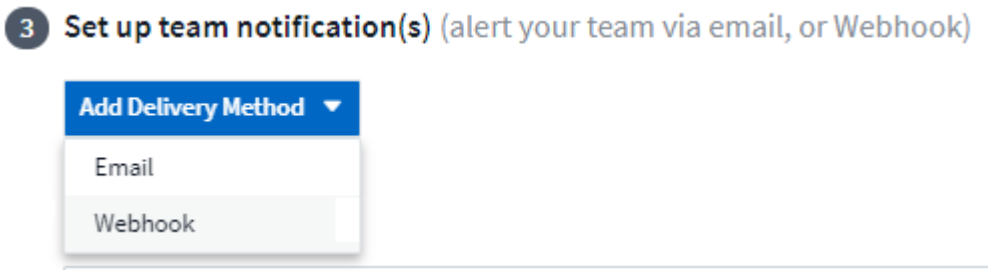

Alerting via Email:

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

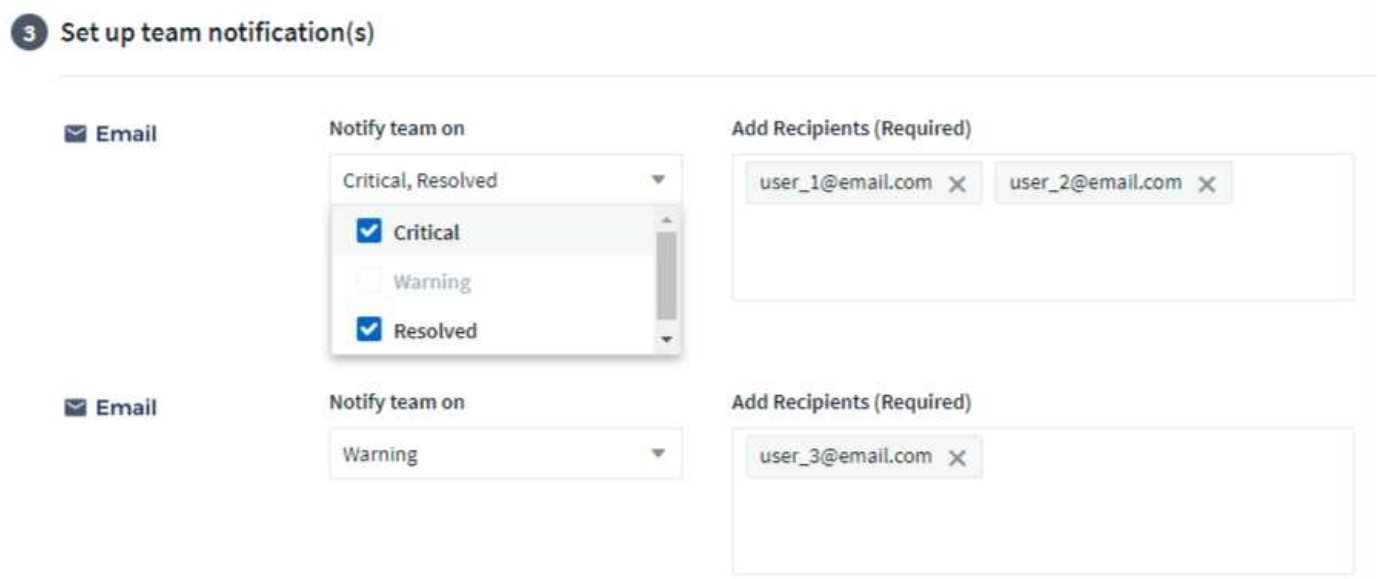

Alerting via Webhook:

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

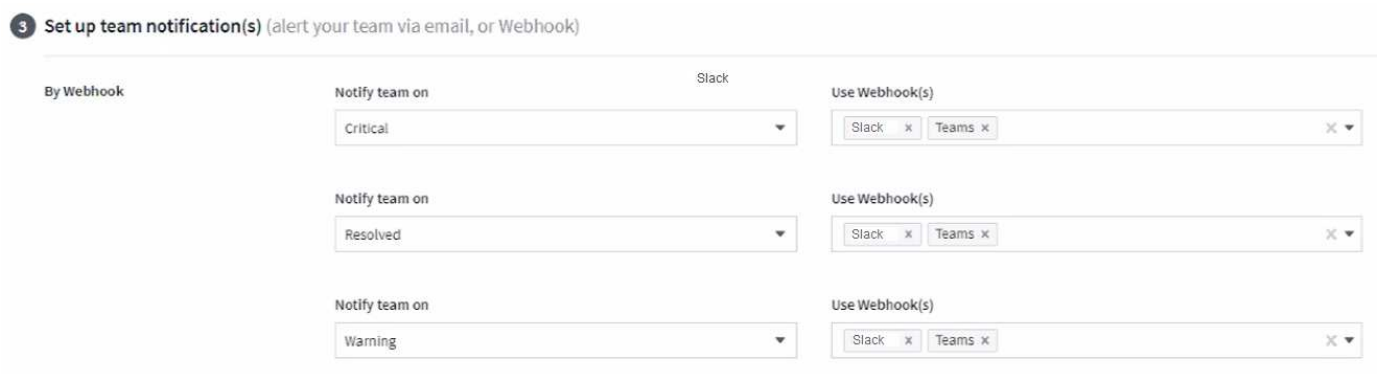

> ONTAP Data Collector notifications take precedence over any specific Monitor notifications that are relevant to the cluster/data collector. The recipient list you set for the Data Collector itself will receive the data collector alerts. If there are no active data collector alerts, then monitor-generated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

Setting Corrective Actions or Additional Information

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

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

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

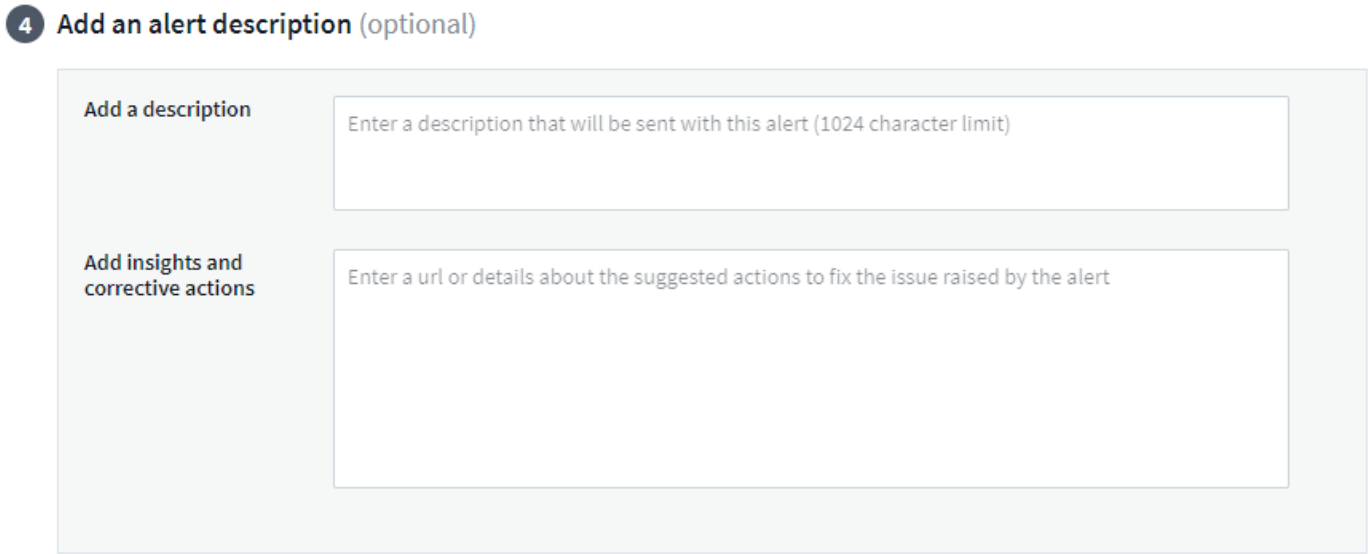

Save your Monitor

1. If desired, you can add a description of the monitor.
2. Give the Monitor a meaningful name and click **Save**.

   Your new monitor is added to the list of active Monitors.

Monitor List

The Monitor page lists the currently configured monitors, showing the following:

- Monitor Name
- Status
- Object/metric being monitored
- Conditions of the Monitor

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

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

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

Monitor Groups

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

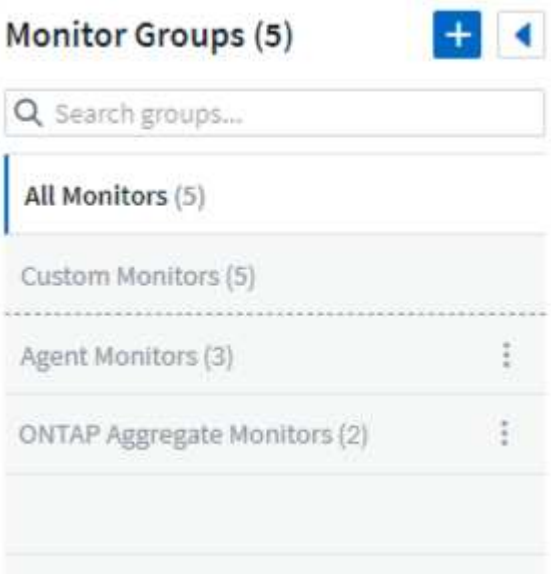

The following monitor groups are shown. The number of monitors contained in a group is shown next to the group name.

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

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

Suspended Monitors

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

System-Defined Monitors

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

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

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

Custom Monitor Groups

You can create your own groups to contain monitors based on your needs. For example, you may want a group for all of your storage-related monitors.

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

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

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

Remove monitors by clicking on a group and selecting *Remove from Group* from the menu. You can not remove monitors from the *All Monitors* or *Custom Monitors* group. To delete a monitor from these groups, you must delete the monitor itself.

> Removing a monitor from a group does not delete the monitor from Data Infrastructure Insights. To completely remove a monitor, select the monitor and click *Delete*. This also removes it from the group to which it belonged and it is no longer available to any user.

You can also move a monitor to a different group in the same manner, selecting *Move to Group*.

To pause or resume all monitors in a group at once, select the menu for the group and click *Pause* or *Resume*.

Use the same menu to rename or delete a group. Deleting a group does not delete the monitors from Data Infrastructure Insights; they are still available in *All Monitors*.

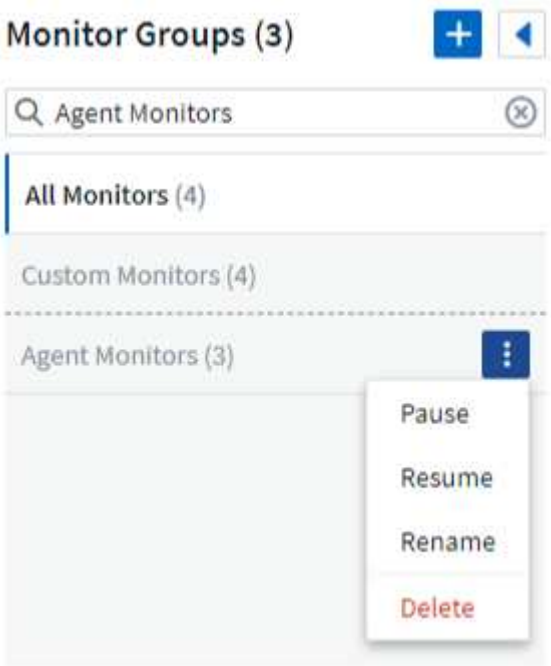

System-Defined Monitors

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

View the System-Defined Monitors page for descriptions of monitors included with Data Infrastructure Insights.

More Information

- Viewing and Dismissing Alerts

Viewing and Managing Alerts from Monitors

Data Infrastructure Insights displays alerts when monitored thresholds are exceeded.

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

Viewing and Managing Alerts

To view and manage alerts, do the following.

1. Navigate to the **Alerts > All Alerts** page.
2. A list of up to the most recent 1,000 alerts is displayed. You can sort this list on any field by clicking the column header for the field. The list displays the following information. Note that not all of these columns are displayed by default. You can select columns to display by clicking on the "gear" icon:
    - **Alert ID**: System-generated unique alert ID
    - **Triggered Time**: The time at which the relevant Monitor triggered the alert
    - **Current Severity** (Active alerts tab): The current severity of the active alert

- **Top Severity** (Resolved alerts tab); The maximum severity of the alert before it was resolved
- **Monitor**: The monitor configured to trigger the alert
- **Triggered On**: The object on which the monitored threshold was breached
- **Status**: Current alert status, *New* or *In Process*
- **Active Status**: *Active* or *Resolved*
- **Condition**: The threshold condition that triggered the alert
- **Metric**: The object's metric on which the monitored threshold was breached
- **Monitor Status**: Current status of the monitor that triggered the alert
- **Has Corrective Action**: The alert has suggested corrective actions. Open the alert page to view these.

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

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

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

Clicking on an Alert ID opens the Alert Detail Page.

Alert Detail Panel

Select any alert row to open the alert's detail panel. The alert detail panel provides additional detail about the alert, including a *Summary*, a *Performance* section showing graphs related to the object's data, any *Related Assets*, and *Comments* entered by alert investigators.

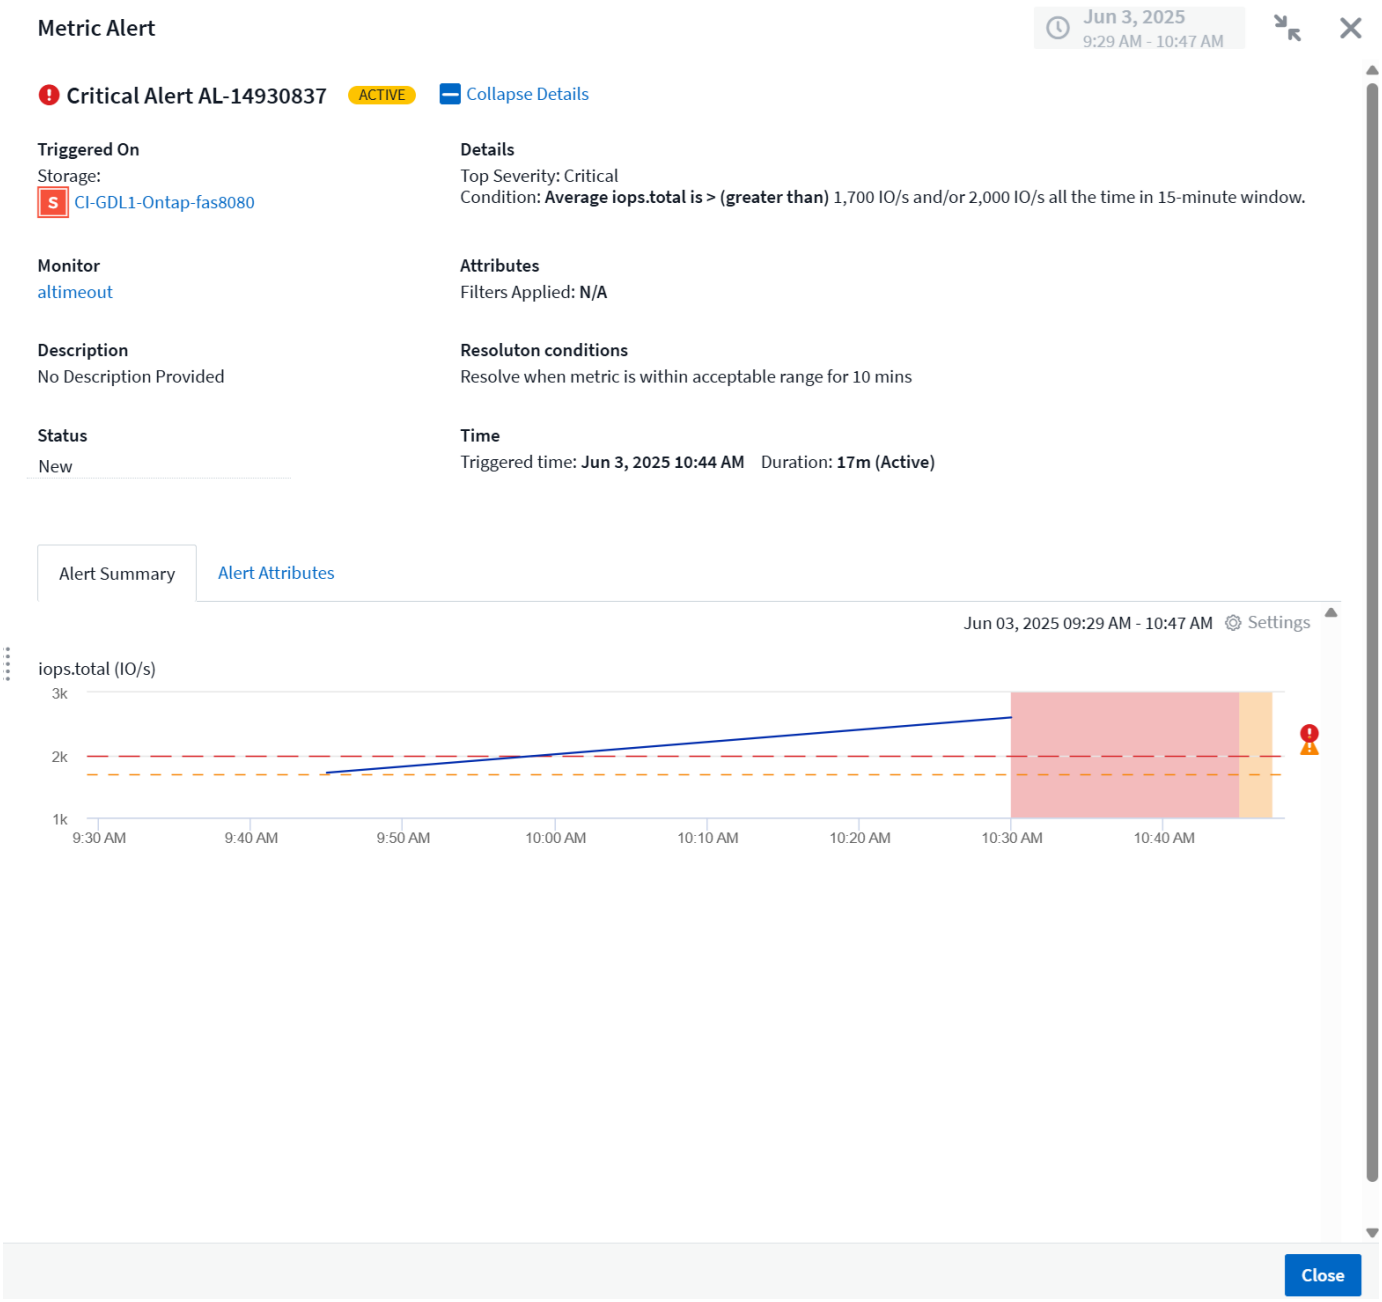

Alerts When Data Is Missing

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

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

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

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

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

"Permanently Active" Alerts

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

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

Maintenance Windows

Maintenance Windows allow you to suppress alert notifications during scheduled maintenance periods, helping you avoid unnecessary disruptions.

With Maintenance Windows, you can schedule specific maintenance periods during which alert notifications are suppressed, for objects and metrics you choose. For example, you can suppress alert notifications triggered by specific storage systems, when those storage systems are in a planned upgrade period.

You can choose whether to suppress both alerts and notifications during the window, or suppress notifications only but continue to detect alerts.

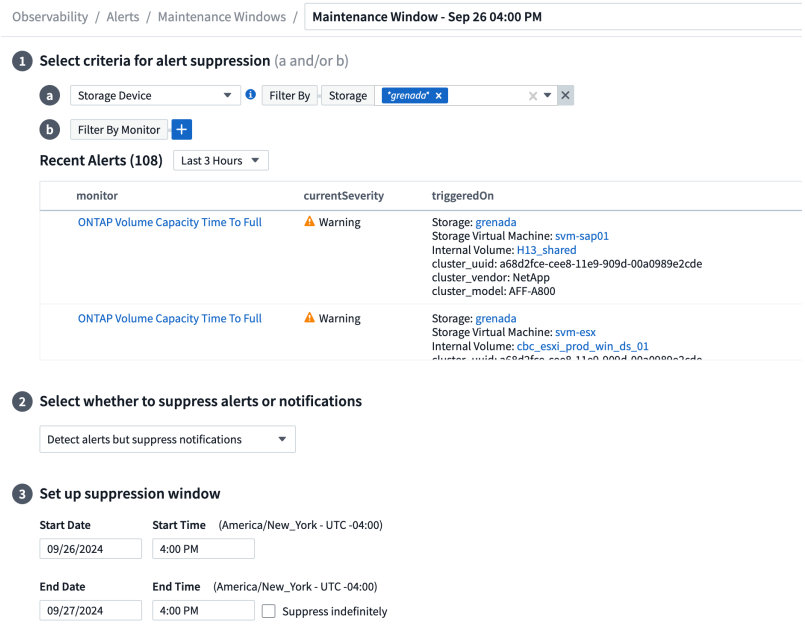

Creating a Maintenance Window

To create a Maintenance Window:

1. Navigate to **Observability > Alerts > Maintenance Windows** and click **+ Maintenance Window**.
2. Enter a **Name** for the maintenance window. For best results, make the name meaningful and specific.
3. Select the **storage device** or **object attribute**.

4. Filter for the **monitor or monitors**. Wildcard are accepted.
5. Select the **Start Time** and **End Time** for the maintenance period.
6. Choose whether to suppress **alerts and notifications**, or **notifications only** (alerts are still detected during the window).
7. Click **Save** to create the maintenance window.

Configuring Email Notifications

You can configure an email list for subscription-related notifications, as well as a global email list of recipients for notification of performance policy threshold violations.

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

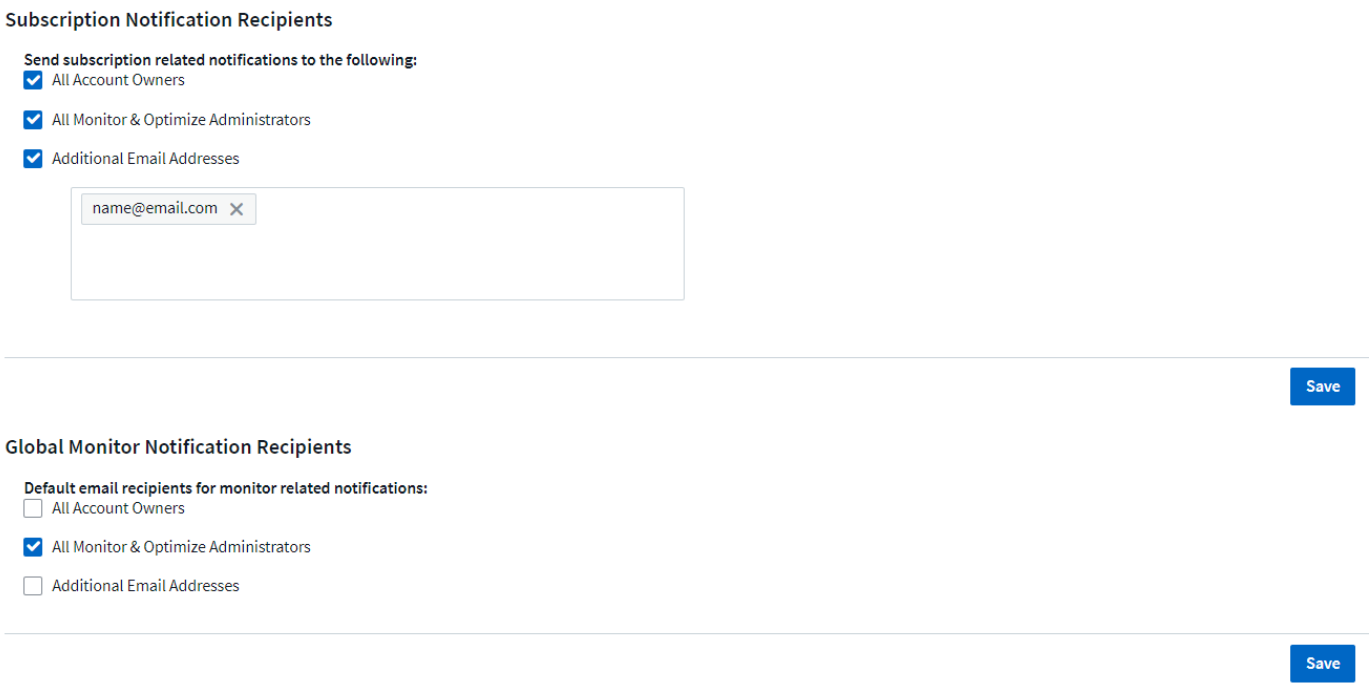

Subscription Notification Recipients

To configure recipients for subscription-related event notifications, go to the "Subscription Notification Recipients" section.
You can choose to have email notifications sent for subscription-related events to any or all of the following recipients:

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

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

Notification:	User Action:

Trial or subscription has been updated	Review subscription details on the Subscription page
Subscription will expire in 90 days Subscription will expire in 30 days	No action needed if “Auto Renewal” is enabled Contact NetApp sales to renew the subscription
Trial ends in 2 days	Renew trial from the Subscription page. You can renew a trial one time. Contact NetApp sales to purchase a subscription
Trial or subscription has expired Account will stop collecting data in 48 hours Account will be deleted after 48 hours	Contact NetApp sales to purchase a subscription

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

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

Global Recipient List for Alerts

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

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

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

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

Editing Notifications for ONTAP

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

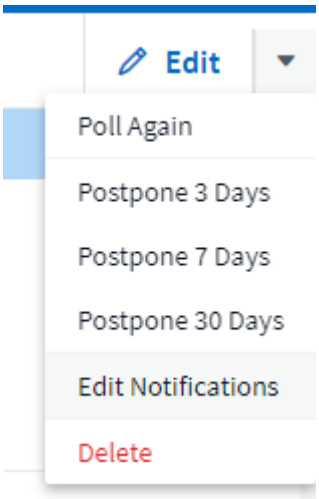

From here, you can set notifications for Critical, Warning, Informational, and/or Resolved alerts. Each scenario can notify the Global Recipient list or other recipients you choose.

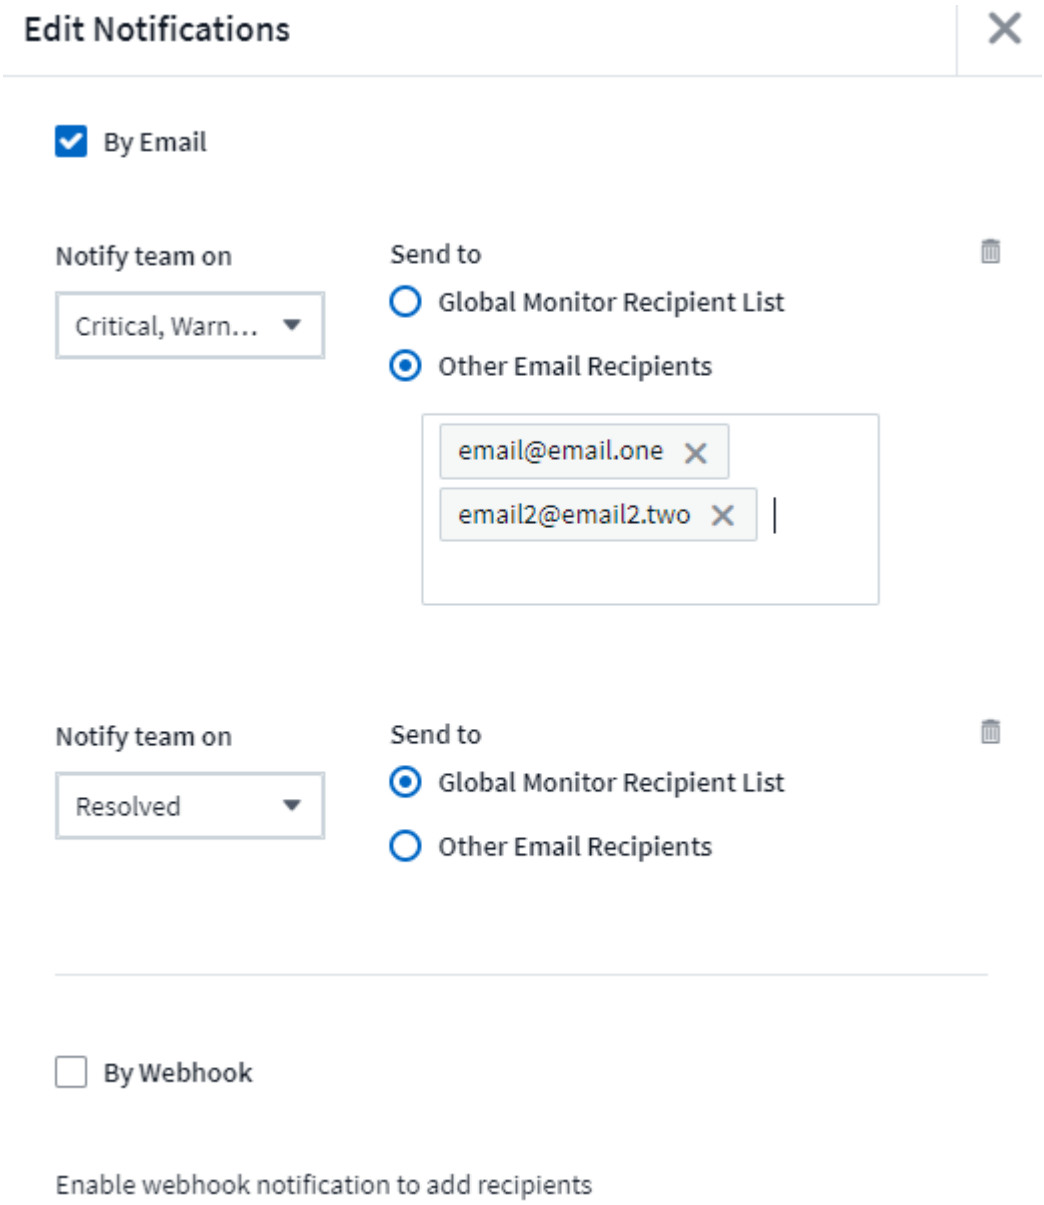

Anomaly Detection Monitors

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

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

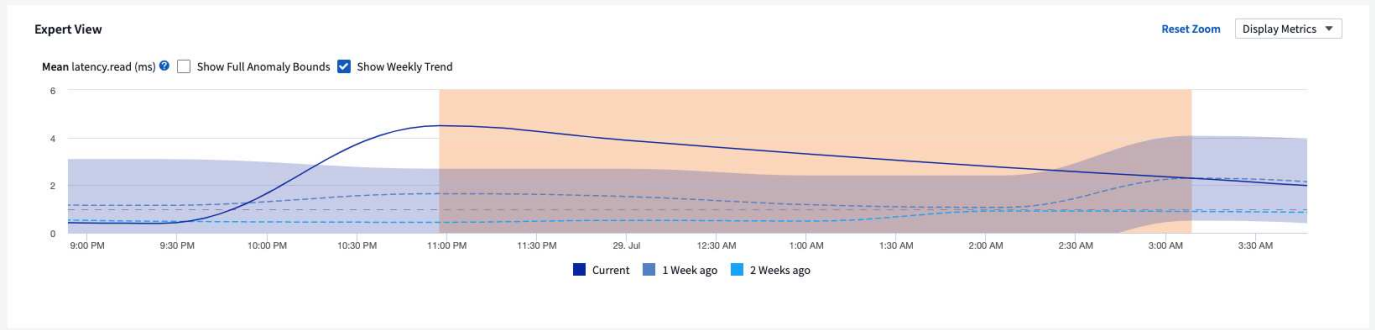

What is Anomaly Detection?

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

Anomaly Detection monitoring differs from Threshold Monitoring.

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

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

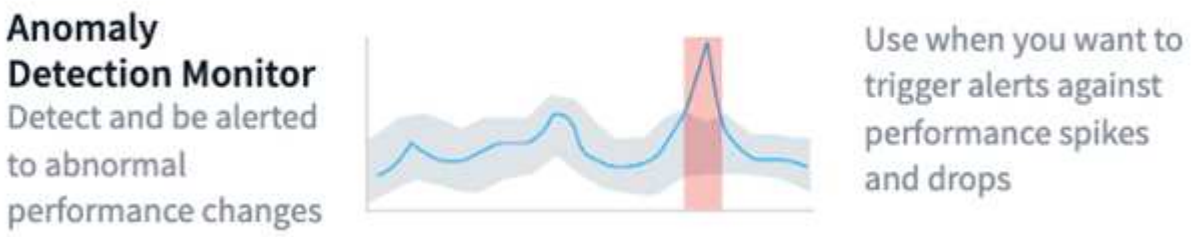

When would I need Anomaly Detection?

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

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

Creating an Anomaly Detection Monitor

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

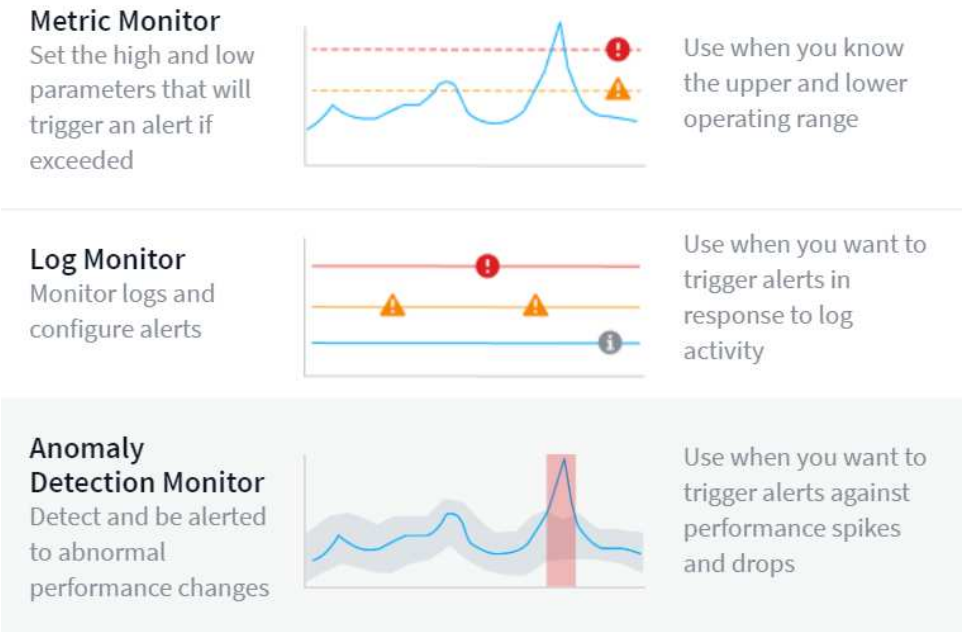

Choose the object and metric you want to monitor. You can set filters and grouping as with other types of monitors.

Next, set the conditions for the monitor.

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

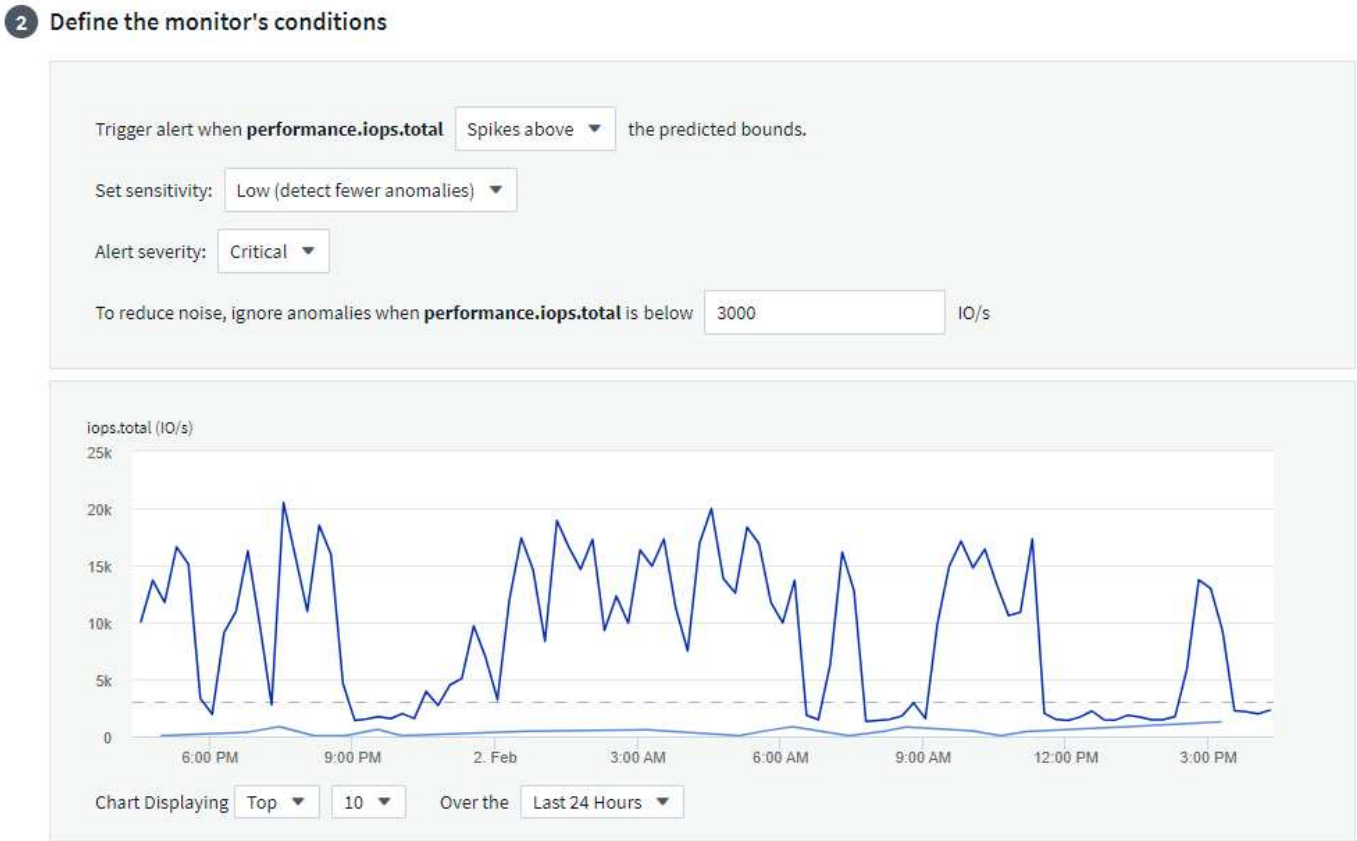

Finally, you can configure a delivery method for the alerts (email, webhook, or both), give the monitor an optional description or corrective actions, and add the monitor to a custom group, if desired.

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

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

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

Viewing the Anomalies

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

bounds.

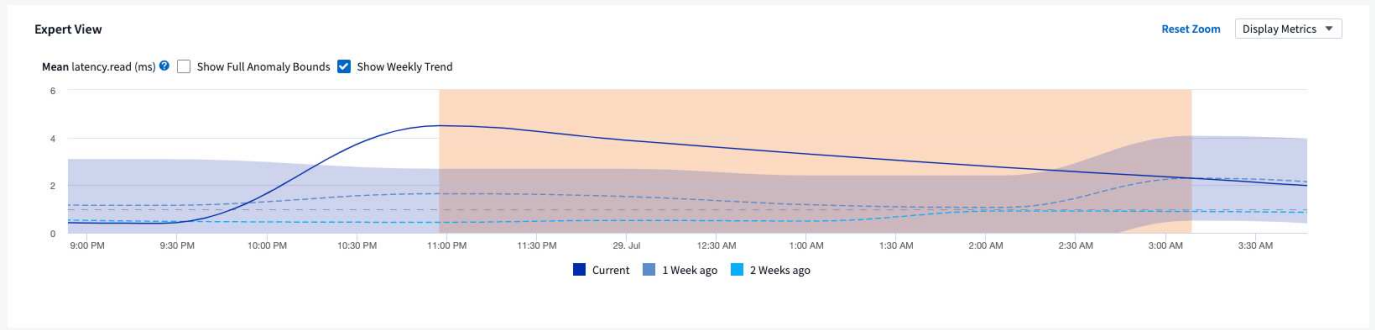

While viewing an anomaly chart on an alert landing page, you can choose the following options:

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

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

Internal Volumes	Value at Time of Detection (ms)	Change Over Detection Window (ms)
vol_mysql_logs_1	176.92	+99.82
vol_mysql_logs_2	159.32	-3.24
vol_mysql_logs_3	158.21	-2.43
vol_mysql_logs_4	127.93	+0.25
vol_mysql_logs_5	126.40	+0.17
vol_mysql_logs_6	126.11	+0.20

System Monitors

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

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

Monitor Descriptions

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

System-defined monitor groups cannot be modified or removed.

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

- **ONTAP Infrastructure** includes monitors for infrastructure-related issues in ONTAP clusters.
- **ONTAP Workload Examples** includes monitors for workload-related issues.
- **CVO Infrastructure** and **CVO Workload Examples** include corresponding monitors for Cloud Volumes ONTAP.
- Default state varies by monitor. Refer to the tables below and review a monitor before enabling notifications.

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

Metric Monitors

Monitor Name	Severity	Monitor Description	Corrective Action

Fiber Channel Port Utilization High	CRITICAL	Fiber Channel Protocol ports are used to receive and transfer the SAN traffic between the customer host system and the ONTAP LUNs. If the port utilization is high, then it will become a bottleneck and it will ultimately affect the performance of sensitive of Fiber Channel Protocol workloads.…A warning alert indicates that planned action should be taken to balance network traffic.…A critical alert indicates that service disruption is imminent and emergency measures should be taken to balance network traffic to ensure service continuity.	If critical threshold is breached, consider immediate actions to minimize service disruption: 1. Move workloads to another lower utilized FCP port. 2. Limit the traffic of certain LUNs only to essential work, either via QoS policies in ONTAP or host-side configuration to lighten the utilization of the FCP ports.… If warning threshold is breached, plan to take the following actions: 1. Configure more FCP ports to handle the data traffic so that the port utilization gets distributed among more ports. 2. Move workloads to another lower utilized FCP port. 3. Limit the traffic of certain LUNs only to essential work, either via QoS policies in ONTAP or host-side configuration to lighten the utilization of the FCP ports.

Lun Latency High	CRITICAL	LUNs are objects that serve the I/O traffic often driven by performance sensitive applications such as databases. High LUN latencies means that the applications themselves might suffer and be unable to accomplish their tasks.…A warning alert indicates that planned action should be taken to move the LUN to appropriate Node or Aggregate.…A critical alert indicates that service disruption is imminent and emergency measures should be taken to ensure service continuity. Following are expected latencies based on media type - SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds, and SATA HDD 17-20 milliseconds	If critical threshold is breached, consider following actions to minimize service disruption: If the LUN or its volume has a QoS policy associated with it, then evaluate its threshold limits and validate if they are causing the LUN workload to get throttled.… If warning threshold is breached, plan to take the following actions: 1. If aggregate is also experiencing high utilization, move the LUN to another aggregate. 2. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node. 3. If the LUN or its volume has a QoS policy associated with it, evaluate its threshold limits and validate if they are causing the LUN workload to get throttled.

Network Port Utilization High	CRITICAL	Network ports are used to receive and transfer the NFS, CIFS, and iSCSI protocol traffic between the customer host systems and the ONTAP volumes. If the port utilization is high, then it becomes a bottleneck and it will ultimately affect the performance of NFS, CIFS and iSCSI workloads.…A warning alert indicates that planned action should be taken to balance network traffic.…A critical alert indicates that service disruption is imminent and emergency measures should be taken to balance network traffic to ensure service continuity.	If critical threshold is breached, consider following immediate actions to minimize service disruption: 1. Limit the traffic of certain volumes only to essential work, either via QoS policies in ONTAP or host-side analysis to decrease the utilization of the network ports. 2. Configure one or more volumes to use another lower utilized network port.… If warning threshold is breached, consider the following immediate actions: 1. Configure more network ports to handle the data traffic so that the port utilization gets distributed among more ports. 2. Configure one or more volumes to use another lower utilized network port.

NVMe Namespace Latency High	CRITICAL	NVMe Namespaces are objects that serve the I/O traffic that is driven by performance sensitive applications such as databases. High NVMe Namespaces latency means that the applications themselves may suffer and be unable to accomplish their tasks.…A warning alert indicates that planned action should be taken to move the LUN to appropriate Node or Aggregate.…A critical alert indicates that service disruption is imminent and emergency measures should be taken to ensure service continuity.	If critical threshold is breached, consider immediate actions to minimize service disruption: If the NVMe namespace or its volume has a QoS policy assigned to them, then evaluate its limit thresholds in case they are causing the NVMe namespace workload to get throttled.… If warning threshold is breached, consider to take the following actions: 1. If aggregate is also experiencing high utilization, move the LUN to another aggregate. 2. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node. 3. If the NVMe namespace or its volume has a QoS policy assigned to them, evaluate its limit thresholds in case they are causing the NVMe namespace workload to get throttled.

QTree Capacity Full	CRITICAL	A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a default space quota or a quota defined by a quota policy to limit amount of data stored in the tree within the volume capacity.…A warning alert indicates that planned action should be taken to increase the space.…A critical alert indicates that service disruption is imminent and emergency measures should be taken to free up space to ensure service continuity.	If critical threshold is breached, consider immediate actions to minimize service disruption: 1. Increase the space of the qtree in order to accommodate the growth. 2. Delete unwanted data to free up space.… If warning threshold is breached, plan to take the following immediate actions: 1. Increase the space of the qtree in order to accommodate the growth. 2. Delete unwanted data to free up space.
QTree Capacity Hard Limit	CRITICAL	A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a space quota measured in KBytes that is used to store data in order to control the growth of user data in volume and not exceed its total capacity.…A qtree maintains a soft storage capacity quota that provides alert to the user proactively before reaching the total capacity quota limit in the qtree and being unable to store data anymore. Monitoring the amount of data stored within a qtree ensures that the user receives uninterrupted data service.	If critical threshold is breached, consider following immediate actions to minimize service disruption: 1. Increase the tree space quota in order to accommodate the growth 2. Instruct the user to delete unwanted data in the tree to free up space

QTree Capacity Soft Limit	WARNING	A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a space quota measured in KBytes that it can use to store data in order to control the growth of user data in volume and not exceed its total capacity.…A qtree maintains a soft storage capacity quota that provides alert to the user proactively before reaching the total capacity quota limit in the qtree and being unable to store data anymore. Monitoring the amount of data stored within a qtree ensures that the user receives uninterrupted data service.	If warning threshold is breached, consider the following immediate actions: 1. Increase the tree space quota to accommodate the growth. 2. Instruct the user to delete unwanted data in the tree to free up space.
QTree Files Hard Limit	CRITICAL	A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a quota of the number of files that it can contain to maintain a manageable file system size within the volume.…A qtree maintains a hard file number quota beyond which new files in the tree are denied. Monitoring the number of files within a qtree ensures that the user receives uninterrupted data service.	If critical threshold is breached, consider immediate actions to minimize service disruption: 1. Increase the file count quota for the qtree. 2. Delete unwanted files from the qtree file system.

QTree Files Soft Limit	WARNING	A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a volume. Each qtree has a quota of the number of files that it can contain in order to maintain a manageable file system size within the volume.…A qtree maintains a soft file number quota to provide alert to the user proactively before reaching the limit of files in the qtree and being unable to store any additional files. Monitoring the number of files within a qtree ensures that the user receives uninterrupted data service.	If warning threshold is breached, plan to take the following immediate actions: 1. Increase the file count quota for the qtree. 2. Delete unwanted files from the qtree file system.
Snapshot Reserve Space Full	CRITICAL	Storage capacity of a volume is necessary to store application and customer data. A portion of that space, called snapshot reserved space, is used to store snapshots which allow data to be protected locally. The more new and updated data stored in the ONTAP volume the more snapshot capacity is used and less snapshot storage capacity is available for future new or updated data. If the snapshot data capacity within a volume reaches the total snapshot reserve space, it might lead to the customer being unable to store new snapshot data and reduction in the level of protection for the data in the volume. Monitoring the volume used snapshot capacity ensures data services continuity.	If critical threshold is breached, consider immediate actions to minimize service disruption: 1. Configure snapshots to use data space in the volume when the snapshot reserve is full. 2. Delete some older unwanted snapshots to free up space.… If warning threshold is breached, plan to take the following immediate actions: 1. Increase the snapshot reserve space within the volume to accommodate the growth. 2. Configure snapshots to use data space in the volume when the snapshot reserve is full.

Storage Capacity Limit	CRITICAL	When a storage pool (aggregate) is filling up, I/O operations slow down and finally stop resulting in storage outage incident. A warning alert indicates that planned action should be taken soon to restore minimum free space. A critical alert indicates that service disruption is imminent and emergency measures should be taken to free up space to ensure service continuity.	If critical threshold is breached, immediately consider the following actions to minimize service disruption: 1. Delete Snapshots on non-critical volumes. 2. Delete Volumes or LUNs that are non-essential workloads and that may be restored from off storage copies.…….If warning threshold is breached, plan the following immediate actions: 1. Move one or more volumes to a different storage location. 2. Add more storage capacity. 3. Change storage efficiency settings or tier inactive data to cloud storage.
Storage Performance Limit	CRITICAL	When a storage system reaches its performance limit, operations slow down, latency goes up and workloads and applications may start failing. ONTAP evaluates the storage pool utilization for workloads and estimates what percent of performance has been consumed.…A warning alert indicates that planned action should be taken to reduce storage pool load to ensure that there will be enough storage pool performance left to service workload peaks.…A critical alert indicates that a performance brownout is imminent and emergency measures should be taken to reduce storage pool load to ensure service continuity.	If critical threshold is breached, consider following immediate actions to minimize service disruption: 1. Suspend scheduled tasks such as Snapshots or SnapMirror replication. 2. Idle non-essential workloads.… If warning threshold is breached, take the following actions immediately: 1. Move one or more workloads to a different storage location. 2. Add more storage nodes (AFF) or disk shelves(FAS) and redistribute workloads 3. Change workload characteristics(block size, application caching).

User Quota Capacity Hard Limit	CRITICAL	ONTAP recognizes the users of Unix or Windows systems who have the rights to access volumes, files or directories within a volume. As a result, ONTAP allows the customers to configure storage capacity for their users or groups of users of their Linux or Windows systems. The user or group policy quota limits the amount of space the user can utilize for their own data.…A hard limit of this quota allows notification of the user when the amount of capacity used within the volume is right before reaching the total capacity quota. Monitoring the amount of data stored within a user or group quota ensures that the user receives uninterrupted data service.	If critical threshold is breached, consider following immediate actions to minimize service disruption: 1. Increase the space of the user or group quota in order to accommodate the growth. 2. Instruct the user or group to delete unwanted data to free up space.

ONTAP Volume Capacity Full	CRITICAL	Storage capacity of a volume is necessary to store application and customer data. The more data stored in the ONTAP volume the less storage availability for future data. If the data storage capacity within a volume reaches the total storage capacity may lead to the customer being unable to store data due to lack of storage capacity. Monitoring the volume used storage capacity ensures data services continuity.	If critical threshold is breached, consider following immediate actions to minimize service disruption: 1. Increase the space of the volume to accommodate the growth. 2. Delete unwanted data to free up space. 3. If snapshot copies occupy more space than the snapshot reserve, delete old Snapshots or enable Volume Snapshot Autodelete.…If warning threshold is breached, plan to take the following immediate actions: 1. Increase the space of the volume in order to accommodate the growth 2. If snapshot copies occupy more space than the snapshot reserve, delete old Snapshots or enabling Volume Snapshot Autodelete.……

Volume Inodes Limit	CRITICAL	Volumes that store files use index nodes (inode) to store file metadata. When a volume exhausts its inode allocation, no more files can be added to it.…A warning alert indicates that planned action should be taken to increase the number of available inodes.…A critical alert indicates that file limit exhaustion is imminent and emergency measures should be taken to free up inodes to ensure service continuity.	If critical threshold is breached, consider following immediate actions to minimize service disruption: 1. Increase the inodes value for the volume. If the inodes value is already at the max value, then split the volume into two or more volumes because the file system has grown beyond the maximum size. 2. Use FlexGroup as it helps to accommodate large file systems.… If warning threshold is breached, plan to take the following immediate actions: 1. Increase the inodes value for the volume. If the inodes value is already at the max, then split the volume into two or more volumes because the file system has grown beyond the maximum size. 2. Use FlexGroup as it helps to accommodate large file systems

Volume Latency High	CRITICAL	Volumes are objects that serve the I/O traffic often driven by performance sensitive applications including devOps applications, home directories, and databases. High volume latencies means that the applications themselves may suffer and be unable to accomplish their tasks. Monitoring volume latencies is critical to maintain application consistent performance. The following are expected latencies based on media type - SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.	If critical threshold is breached, consider following immediate actions to minimize service disruption: If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled.… If warning threshold is breached, consider the following immediate actions: 1. If aggregate is also experiencing high utilization, move the volume to another aggregate. 2. If the volume has a QoS policy assigned to it, evaluate its limit thresholds in case they are causing the volume workload to get throttled. 3. If the node is also experiencing high utilization, move the volume to another node or reduce the total workload of the node.
Monitor Name	Severity	Monitor Description	Corrective Action

Node High Latency	WARNING / CRITICAL	Node latency has reached the levels where it might affect the performance of the applications on the node. Lower node latency ensures consistent performance of the applications. The expected latencies based on media type are: SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.	If critical threshold is breached, then immediate actions should be taken to minimize service disruption: 1. Suspend scheduled tasks, Snapshots or SnapMirror replication 2. Lower the demand of lower priority workloads via QoS limits 3. Inactivate non-essential workloads Consider immediate actions when warning threshold is breached: 1. Move one or more workloads to a different storage location 2. Lower the demand of lower priority workloads via QoS limits 3. Add more storage nodes (AFF) or disk shelves (FAS) and redistribute workloads 4. Change workload characteristics (block size, application caching etc)

Node Performance Limit	WARNING / CRITICAL	Node performance utilization has reached the levels where it might affect the performance of the IOs and the applications supported by the node. Low node performance utilization ensures consistent performance of the applications.	Immediate actions should be taken to minimize service disruption if critical threshold is breached: 1. Suspend scheduled tasks, Snapshots or SnapMirror replication 2. Lower the demand of lower priority workloads via QoS limits 3. Inactivate non-essential workloads Consider the following actions if warning threshold is breached: 1. Move one or more workloads to a different storage location 2. Lower the demand of lower priority workloads via QoS limits 3. Add more storage nodes (AFF) or disk shelves (FAS)and redistribute workloads 4. Change workload characteristics (block size, application caching etc)

Storage VM High Latency	WARNING / CRITICAL	Storage VM (SVM) latency has reached the levels where it might affect the performance of the applications on the storage VM. Lower storage VM latency ensures consistent performance of the applications. The expected latencies based on media type are: SSD up to 1-2 milliseconds; SAS up to 8-10 milliseconds and SATA HDD 17-20 milliseconds.	If critical threshold is breached, then immediately evaluate the threshold limits for volumes of the storage VM with a QoS policy assigned, to verify whether they are causing the volume workloads to get throttled Consider following immediate actions when warning threshold is breached: 1. If aggregate is also experiencing high utilization, move some volumes of the storage VM to another aggregate. 2. For volumes of the storage VM with a QoS policy assigned, evaluate the threshold limits if they are causing the volume workloads to get throttled 3. If the node is experiencing high utilization, move some volumes of the storage VM to another node or reduce the total workload of the node
User Quota Files Hard Limit	CRITICAL	The number of files created within the volume has reached the critical limit and additional files cannot be created. Monitoring the number of files stored ensures that the user receives uninterrupted data service.	Immediate actions are required to minimize service disruption if critical threshold is breached….Consider taking following actions: 1. Increase the file count quota for the specific user 2. Delete unwanted files to reduce the pressure on the files quota for the specific user

User Quota Files Soft Limit	WARNING	The number of files created within the volume has reached the threshold limit of the quota and is near to the critical limit. You cannot create additional files if quota reaches the critical limit. Monitoring the number of files stored by a user ensures that the user receives uninterrupted data service.	Consider immediate actions if warning threshold is breached: 1. Increase the file count quota for the specific user quota 2. Delete unwanted files to reduce the pressure on the files quota for the specific user
Volume Cache Miss Ratio	WARNING / CRITICAL	Volume Cache Miss Ratio is the percentage of read requests from the client applications that are returned from the disk instead of being returned from the cache. This means that the volume has reached the set threshold.	If critical threshold is breached, then immediate actions should be taken to minimize service disruption: 1. Move some workloads off of the node of the volume to reduce the IO load 2. If not already on the node of the volume, increase the WAFL cache by purchasing and adding a Flash Cache 3. Lower the demand of lower priority workloads on the same node via QoS limits Consider immediate actions when warning threshold is breached: 1. Move some workloads off of the node of the volume to reduce the IO load 2. If not already on the node of the volume, increase the WAFL cache by purchasing and adding a Flash Cache 3. Lower the demand of lower priority workloads on the same node via QoS limits 4. Change workload characteristics (block size, application caching etc)

Volume Qtree Quota Overcommit	WARNING / CRITICAL	Volume Qtree Quota Overcommit specifies the percentage at which a volume is considered to be overcommitted by the qtree quotas. The set threshold for the qtree quota is reached for the volume. Monitoring the volume qtree quota overcommit ensures that the user receives uninterrupted data service.	If critical threshold is breached, then immediate actions should be taken to minimize service disruption: 1. Increase the space of the volume 2. Delete unwanted data When warning threshold is breached, then consider increasing the space of the volume.

Log Monitors

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

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

FlexGroup Constituent Full	CRITICAL	A constituent within a FlexGroup volume is full, which might cause a potential disruption of service. You can still create or expand files on the FlexGroup volume. However, none of the files that are stored on the constituent can be modified. As a result, you might see random out-of-space errors when you try to perform write operations on the FlexGroup volume.	It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X" command.…Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent.
FlexGroup Constituent Out of Inodes	CRITICAL	A constituent of a FlexGroup volume has run out of inodes, which might cause a potential disruption of service. You cannot create new files on this constituent. This might lead to an overall imbalanced distribution of content across the FlexGroup volume.	It is recommended that you add capacity to the FlexGroup volume by using the "volume modify -files +X" command.…Alternatively, delete files from the FlexGroup volume. However, it is difficult to determine which files have landed on the constituent.
Main Unit Fan Failed	WARNING	One or more main unit fans have failed. The system remains operational.…However, if the condition persists for too long, the overtemperature might trigger an automatic shutdown.	Reseat the failed fans. If the error persists, replace them.

NVRAM Battery Low	WARNING	The NVRAM battery capacity is critically low. There might be a potential data loss if the battery runs out of power.…Your system generates and transmits an AutoSupport or "call home" message to NetApp technical support and the configured destinations if it is configured to do so. The successful delivery of an AutoSupport message significantly improves problem determination and resolution.	Perform the following corrective actions:…View the battery's current status, capacity, and charging state by using the "system node environment sensors show" command.…If the battery was replaced recently or the system was non-operational for an extended period of time, monitor the battery to verify that it is charging properly.…Contact NetApp technical support if the battery runtime continues to decrease below critical levels, and the storage system shuts down automatically.
Service Processor Not Configured	WARNING	This event occurs on a weekly basis, to remind you to configure the Service Processor (SP). The SP is a physical device that is incorporated into your system to provide remote access and remote management capabilities. You should configure the SP to use its full functionality.	Perform the following corrective actions:…Configure the SP by using the "system service-processor network modify" command.…Optionally, obtain the MAC address of the SP by using the "system service-processor network show" command.…Verify the SP network configuration by using the "system service-processor network show" command.…Verify that the SP can send an AutoSupport email by using the "system service-processor autosupport invoke" command. NOTE: AutoSupport email hosts and recipients should be configured in ONTAP before you issue this command.

Service Processor Offline	CRITICAL	ONTAP is no longer receiving heartbeats from the Service Processor (SP), even though all the SP recovery actions have been taken. ONTAP cannot monitor the health of the hardware without the SP.…The system will shut down to prevent hardware damage and data loss. Set up a panic alert to be notified immediately if the SP goes offline.	Power-cycle the system by performing the following actions:…Pull the controller out from the chassis.…Push the controller back in.…Turn the controller back on.…If the problem persists, replace the controller module.
Shelf Fan Failed	CRITICAL	The indicated cooling fan or fan module of the shelf has failed. The disks in the shelf might not receive enough cooling airflow, which might result in disk failure.	Perform the following corrective actions:…Verify that the fan module is fully seated and secured. NOTE: The fan is integrated into the power supply module in some disk shelves.…If the issue persists, replace the fan module.…If the issue still persists, contact NetApp technical support for assistance.
System Cannot Operate Due to Main Unit Fan Failure	CRITICAL	One or more main unit fans have failed, disrupting system operation. This might lead to a potential data loss.	Replace the failed fans.
Unassigned Disks	INFO	System has unassigned disks - capacity is being wasted and your system may have some misconfiguration or partial configuration change applied.	Perform the following corrective actions:…Determine which disks are unassigned by using the "disk show -n" command.…Assign the disks to a system by using the "disk assign" command.
Antivirus Server Busy	WARNING	The antivirus server is too busy to accept any new scan requests.	If this message occurs frequently, ensure that there are enough antivirus servers to handle the virus scan load generated by the SVM.

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

AWS IAM Role Not Found	CRITICAL	The Identity and Access Management (IAM) roles thread cannot find an Amazon Web Services (AWS) IAM role on the AWS metadata server. The IAM role is required to acquire role-based credentials used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible.…	Perform the following:…Log in to the AWS EC2 Management Console.…Navigate to the Instances page.…Find the instance for the Cloud Volumes ONTAP deployment and check its health.…Verify that the AWS IAM role associated with the instance is valid.
AWS IAM Role Not Valid	CRITICAL	The Amazon Web Services (AWS) Identity and Access Management (IAM) role on the AWS metadata server is not valid. The Cloud Volume ONTAP has become inaccessible.…	Perform the following:…Log in to the AWS EC2 Management Console.…Navigate to the Instances page.…Find the instance for the Cloud Volumes ONTAP deployment and check its health.…Verify that the AWS IAM role associated with the instance is valid and has been granted proper privileges to the instance.
AWS Metadata Server Connection Fail	CRITICAL	The Identity and Access Management (IAM) roles thread cannot establish a communication link with the Amazon Web Services (AWS) metadata server. Communication should be established to acquire the necessary AWS IAM role-based credentials used to sign API requests to Amazon Simple Storage Service (Amazon S3). Cloud Volume ONTAP has become inaccessible.…	Perform the following:…Log in to the AWS EC2 Management Console.…Navigate to the Instances page.…Find the instance for the Cloud Volumes ONTAP deployment and check its health.…

FabricPool Space Usage Limit Nearly Reached	WARNING	The total cluster-wide FabricPool space usage of object stores from capacity-licensed providers has nearly reached the licensed limit.	Perform the following corrective actions:…Check the percentage of the licensed capacity used by each FabricPool storage tier by using the "storage aggregate object-store show-space" command.…Delete Snapshot copies from volumes with the tiering policy "snapshot" or "backup" by using the "volume snapshot delete" command to clear up space.…Install a new license on the cluster to increase the licensed capacity.
FabricPool Space Usage Limit Reached	CRITICAL	The total cluster-wide FabricPool space usage of object stores from capacity-licensed providers has reached the license limit.	Perform the following corrective actions:…Check the percentage of the licensed capacity used by each FabricPool storage tier by using the "storage aggregate object-store show-space" command.…Delete Snapshot copies from volumes with the tiering policy "snapshot" or "backup" by using the "volume snapshot delete" command to clear up space.…Install a new license on the cluster to increase the licensed capacity.

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

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

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

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

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

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

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

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

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

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

Anti-Ransomware Log Monitors

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

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

FSx for NetApp ONTAP Monitors

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

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

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

FSx Volume Cache Miss Ratio	Warning @ > 95 %…Critical @ > 100 %	Volume Cache Miss Ratio is the percentage of read requests from the client applications that are returned from the disk instead of being returned from the cache. This means that the volume has reached the set threshold.	If critical threshold is breached, then immediate actions should be taken to minimize service disruption: 1. Move some workloads off of the node of the volume to reduce the IO load 2. Lower the demand of lower priority workloads on the same node via QoS limits…Consider immediate actions when warning threshold is breached: 1. Move some workloads off of the node of the volume to reduce the IO load 2. Lower the demand of lower priority workloads on the same node via QoS limits 3. Change workload characteristics (block size, application caching etc)

K8s Monitors

Monitor Name	Description	Corrective Actions	Severity/Threshold

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

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

Container at Risk of Going OOM Killed	The container's memory limits are set too low. The container is at risk of eviction (Out of Memory Kill).	Increase container memory limits.	Warning @ > 95 %
Workload Down	Workload has no healthy pods.		Critical @ < 1
Persistent Volume Claim Failed Binding	This alert occurs when a binding is failed on a PVC.		Warning
ResourceQuota Mem Limits About to Exceed	Memory limits for Namespace are about to exceed ResourceQuota		Warning @ > 80 % Critical @ > 90 %
ResourceQuota Mem Requests About to Exceed	Memory requests for Namespace are about to exceed ResourceQuota		Warning @ > 80 % Critical @ > 90 %
Node Creation Failed	The node could not be scheduled because of a configuration error.	Check the Kubernetes event log for the cause of the configuration failure.	Critical
Kubernetes Persistent Volume Reclamation Failed	The volume has failed its automatic reclamation.		Warning @ > 0 B
Container CPU Throttling	The container's CPU Limits are set too low. Container processes are slowed.	Increase container CPU limits.	Warning @ > 95 % Critical @ > 98 %
Service Load Balancer Failed to Delete			Warning

Persistent Volume IOPS	IOPS thresholds on persistent volumes can be used to alert an administrator when persistent volumes exceed predefined performance expectations. Activating this monitor will generate alerts appropriate for the typical IOPS profile of persistence volumes. This monitor will cover all persistent volumes on your tenant. The warning and critical threshold values can be adjusted based on your monitoring goals by duplicating this monitor and setting thresholds appropriate for your workload.	**Immediate Actions** If critical threshold is breached, plan Immediate actions to minimize service disruption : 1. Introduce QoS IOPS limits for the volume. 2. Review the application driving the workload on the volume for anomalies. **Actions To Do Soon** If warning threshold is breached, plan the following immediate actions: 1. Introduce QoS IOPS limits for the volume. 2. Review the application driving the workload on the volume for anomalies.	Warning @ > 20,000 IO/s Critical @ > 25,000 IO/s
Service Load Balancer Failed to Update			Warning
POD Failed Mount	This alert occurs when a mount is failed on a POD.		Warning
Node PID Pressure	Available process identifiers on the (Linux) node has fallen below an eviction threshold.	Find and fix pods that generate many processes and starve the node of available process IDs. Set up PodPidsLimit to protect your node against pods or containers that spawn too many processes.	Critical @ > 0
Pod Image Pull Failure	Kubernetes failed to pull the pod container image.	- Make sure the pod's image is spelled correctly in the pod configuration. - Check image tag exists in your registry. - Verify the credentials for the image registry. - Check for registry connectivity issues. - Verify you are not hitting the rate limits imposed by public registry providers.	Warning
Job Running Too Long	Job is running for too long		Warning @ > 1 hr Critical @ > 5 hr

Node Memory High	Node memory usage is high	Add nodes. Fix any unscheduled nodes. Right-size pods to free up memory on nodes.	Warning @ > 85 % Critical @ > 90 %
ResourceQuota CPU Limits About to Exceed	CPU limits for Namespace are about to exceed ResourceQuota		Warning @ > 80 % Critical @ > 90 %
Kubernetes Pod Crash Loop Backoff	Pod has crashed and attempted to restart multiple times.		Critical @ > 3
Node CPU High	Node CPU usage is high.	Add nodes. Fix any unscheduled nodes. Right-size pods to free up CPU on nodes.	Warning @ > 80 % Critical @ > 90 %
Workload Network Latency RTT High	High TCP RTT (Round Trip Time) latency	Check for Network congestion ▒ Identify workloads that consume a lot of network bandwidth. Check for high Pod CPU utilization. Check hardware network performance.	Warning @ > 150 ms Critical @ > 300 ms
Kubernetes Job Failed	Job did not complete successfully due to a node crash or reboot, resource exhaustion, job timeout, or pod scheduling failure.	Check the Kubernetes event logs for failure causes.	Warning @ > 1
Persistent Volume Full in a Few Days	Persistent Volume will run out of space in a few days	-Increase the volume size to ensure that there is sufficient room for the application files. -Reduce the amount of data stored in applications.	Warning @ < 8 day Critical @ < 3 day
Node Memory Pressure	Node is running out of memory. Available memory has met eviction threshold.	Add nodes. Fix any unscheduled nodes. Right-size pods to free up memory on nodes.	Critical @ > 0
Node Unready	Node has been unready for 5 minutes	Verify the node have enough CPU, memory, and disk resources. Check node network connectivity. Check the Kubernetes event logs for failure causes.	Critical @ < 1

Persistent Volume Capacity High	Persistent Volume backend used capacity is high.	- Increase the volume size to ensure that there is sufficient room for the application files. - Reduce the amount of data stored in applications.	Warning @ > 80 % Critical @ > 90 %
Service Load Balancer Failed to Create	Service Load Balancer Create Failed		Critical
Workload Replica Mismatch	Some pods are currently not available for a Deployment or DaemonSet.		Warning @ > 1
ResourceQuota CPU Requests About to Exceed	CPU requests for Namespace are about to exceed ResourceQuota		Warning @ > 80 % Critical @ > 90 %
High Retransmit Rate	High TCP Retransmit Rate	Check for Network congestion - Identify workloads that consume a lot of network bandwidth. Check for high Pod CPU utilization. Check hardware network performance.	Warning @ > 10 % Critical @ > 25 %
Node Disk Pressure	Available disk space and inodes on either the node's root filesystem or image filesystem has satisfied an eviction threshold.	- Increase the size of the node disks to ensure that there is sufficient room for the application files. - Decrease application file usage.	Critical @ > 0
Cluster CPU Saturation High	Cluster allocatable CPU saturation is high. Cluster CPU saturation is calculated as the sum of CPU usage divided by the sum allocatable CPU across all K8s nodes.	Add nodes. Fix any unscheduled nodes. Right-size pods to free up CPU on nodes.	Warning @ > 80 % Critical @ > 90 %

Change Log Monitors

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

Node Removed	Informational	This message occurs when an Storage Node is removed.
Storage Pool Discovered	Informational	This message occurs when an Storage Pool is discovered.
Storage Virtual Machine Discovered	Informational	This message occurs when an Storage Virtual Machine is discovered.
Storage Virtual Machine Modified	Informational	This message occurs when an Storage Virtual Machine is modified.

Data Collection Monitors

Monitor Name	Description	Corrective Action
Collector Failed	The poll of a data collector encountered an unexpected failure situation.	Visit the data collector page in Data Infrastructure Insights to learn more about the situation.
Collector Warning	This Alert typically can arise because of an erroneous configuration of the data collector or of the target system. Revisit the configurations to prevent future Alerts. It can also be due to a retrieval of less-than-complete data where the data collector gathered all the data that it could. This can happen when situations change during data collection (e.g., a virtual machine present at the beginning of data collection is deleted during data collection and before its data is captured).	Check the configuration of the data collector or target system. Note that the monitor for Collector Warning can send more alerts than other monitor types, so it is recommended to set no alert recipients unless you are troubleshooting.

Attribute Monitors

Attribute monitors evaluate configuration attributes reported for an object. Unlike metric monitors, they do not compare a time-series metric with a numeric threshold. Each condition below generates a *Warning* alert when the reported attribute matches a noncompliant or unavailable state.

The previous standalone ONTAP Essentials security metric monitors have been replaced by the consolidated **ONTAP Cluster Misconfigured** and **ONTAP Storage Misconfigured** attribute monitors. The corresponding CVO monitors apply the same checks to Cloud Volumes ONTAP.

ONTAP Attribute Monitors

Monitor Name	Group	Default State	Conditions	Corrective Action
ONTAP Cluster Misconfigured	ONTAP Infrastructure	Active	Alerts when one or more of these 12 security conditions is detected: * AutoSupport HTTPS is not configured * Web HTTP is enabled * SSH is not secured * The login banner is disabled * Cluster peer communication is not encrypted * The certificate issuer is self-signed * The default local administrator is enabled * FIPS is disabled * Log forwarding is not encrypted * MD5 password hashing is enabled * Remote Shell (RSH) is enabled * Telnet is enabled	Review the affected attributes in the alert. Disable insecure protocols and password hashing, secure SSH and peer communication, enable the required security controls, and replace self-signed certificates according to your organization's security policy.
ONTAP Storage Misconfigured	ONTAP Infrastructure	Active	Alerts when an ONTAP Storage VM has audit logging disabled, SSH is not secured, or its login banner is disabled.	Enable Storage VM audit logging, secure SSH, and configure an appropriate login banner.
ONTAP Cluster Lacks Spare Disks	ONTAP Infrastructure	Active	Alerts when the cluster reports that spare disks are low.	Review spare disks by node and add compatible spare capacity to maintain data protection.
ONTAP Storage Service Unavailable	ONTAP Workload Examples	Active	Alerts when an enabled FCP or iSCSI service is unavailable on an ONTAP Storage VM.	Verify Storage VM, protocol, and data LIF status. Review recent ONTAP health and event messages, and restore the affected SAN service.

Monitor Name	Group	Default State	Conditions	Corrective Action
ONTAP Node Misconfigured	ONTAP Workload Examples	Paused	Alerts when Cisco Discovery Protocol (CDP) is disabled on an ONTAP node.	Verify the node's CDP setting and network port status, then enable CDP where it is required for network discovery and topology monitoring.

Cloud Volumes ONTAP Attribute Monitors

Monitor Name	Group	Default State	Conditions	Corrective Action
CVO Cluster Misconfigured	CVO Infrastructure	Active	Applies the cluster security configuration checks described for **ONTAP Cluster Misconfigured** to Cloud Volumes ONTAP.	Review the affected attributes and correct the reported insecure or noncompliant cluster settings.
CVO Storage Misconfigured	CVO Infrastructure	Active	Alerts when a Cloud Volumes ONTAP Storage VM has audit logging disabled, SSH is not secured, or its login banner is disabled.	Enable Storage VM audit logging, secure SSH, and configure an appropriate login banner.
CVO Node Misconfigured	CVO Workload Examples	Paused	Alerts when CDP is disabled on a Cloud Volumes ONTAP node.	Verify the node's CDP and network settings, then enable CDP where required.

NTP Configuration Monitors

The following NTP monitors remain available as standalone system-defined monitors for ONTAP and Cloud Volumes ONTAP.

Monitor Name	Platform	Threshold	Monitor Description	Corrective Action
No NTP Servers are Configured	ONTAP	Warning @ < 1	The cluster has no configured NTP servers. For redundancy and optimum service, NetApp recommends associating at least three NTP servers with the cluster.	Associate an NTP server with the cluster by using `cluster time-service ntp server create -server <ntp_server>`.

Monitor Name	Platform	Threshold	Monitor Description	Corrective Action
NTP Server Count is Low	ONTAP	Warning @ < 3	The cluster has fewer than three configured NTP servers.	Associate additional NTP servers with the cluster by using `cluster time-service ntp server create -server <ntp_server>`.
CVO No NTP Servers are Configured	Cloud Volumes ONTAP	Warning @ < 1	The cluster has no configured NTP servers.	Associate an NTP server with the cluster by using `cluster time-service ntp server create -server <ntp_server>`.
CVO NTP Server Count is Low	Cloud Volumes ONTAP	Warning @ < 3	The cluster has fewer than three configured NTP servers.	Associate additional NTP servers with the cluster by using `cluster time-service ntp server create -server <ntp_server>`.

Data Protection Monitors

Monitor Name	Thresholds	Monitor Description	Corrective Action

Insufficient Space for Lun Snapshot Copy	(Filter contains_luns = Yes) Warning @ > 95 %…Critical @ > 100 %	Storage capacity of a volume is necessary to store application and customer data. A portion of that space, called snapshot reserved space, is used to store snapshots which allow data to be protected locally. The more new and updated data stored in the ONTAP volume the more snapshot capacity is used and less snapshot storage capacity will be available for future new or updated data. If the snapshot data capacity within a volume reaches the total snapshot reserve space it may lead to the customer being unable to store new snapshot data and reduction in the level of protection for the data in the LUNs in the volume. Monitoring the volume used snapshot capacity ensures data services continuity.	**Immediate Actions** If critical threshold is breached, consider immediate actions to minimize service disruption: 1. Configure snapshots to use data space in the volume when the snapshot reserve is full. 2. Delete some older unwanted snapshots to free up space. **Actions To Do Soon** If warning threshold is breached, plan to take the following immediate actions: 1. Increase the snapshot reserve space within the volume to accommodate the growth. 2. Configure snapshots to use data space in the volume when the snapshot reserve is full.
SnapMirror Relationship Lag	Warning @ > 150%…Critical @ > 300%	SnapMirror relationship lag is the difference between the snapshot timestamp and the time on the destination system. The lag_time_percent is the ratio of lag time to the SnapMirror Policy's schedule interval. If the lag time equals the schedule interval, the lag_time_percent will be 100%. If the SnapMirror policy does not have a schedule, lag_time_percent will not be calculated.	Monitor SnapMirror status using the "snapmirror show" command. Check the SnapMirror transfer history using the "snapmirror show-history" command

Cloud Volume (CVO) Monitors

Monitor Name	CI Severity	Monitor Description	Corrective Action

CVO Disk Out of Service	INFO	This event occurs when a disk is removed from service because it has been marked failed, is being sanitized, or has entered the Maintenance Center.	None
CVO Giveback of Storage Pool Failed	CRITICAL	This event occurs during the migration of an aggregate as part of a storage failover (SFO) giveback, when the destination node cannot reach the object stores.	Perform the following corrective actions: Verify that your intercluster LIF is online and functional by using the "network interface show" command. Check network connectivity to the object store server by using the"'ping" command over the destination node intercluster LIF. Verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the "aggregate object-store config show" command. Alternatively, you can override the error by specifying false for the "require-partner-waiting" parameter of the giveback command. Contact NetApp technical support for more information or assistance.

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

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

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

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

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

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

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

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

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

SnapMirror for Business Continuity (SMBC) Mediator Log Monitors

Monitor Name	Severity	Monitor Description	Corrective Action
ONTAP Mediator Added	INFO	This message occurs when ONTAP Mediator is added successfully on a cluster.	None
ONTAP Mediator Not Accessible	CRITICAL	This message occurs when either the ONTAP Mediator is repurposed or the Mediator package is no longer installed on the Mediator server. As a result, SnapMirror failover is not possible.	Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
ONTAP Mediator Removed	INFO	This message occurs when ONTAP Mediator is removed successfully from a cluster.	None

ONTAP Mediator Unreachable	WARNING	This message occurs when the ONTAP Mediator is unreachable on a cluster. As a result, SnapMirror failover is not possible.	Check the network connectivity to the ONTAP Mediator by using the "network ping" and "network traceroute" commands. If the issue persists, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
SMBC CA Certificate Expired	CRITICAL	This message occurs when the ONTAP Mediator certificate authority (CA) certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.	Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new CA certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
SMBC CA Certificate Expiring	WARNING	This message occurs when the ONTAP Mediator certificate authority (CA) certificate is due to expire within the next 30 days.	Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new CA certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.
SMBC Client Certificate Expired	CRITICAL	This message occurs when the ONTAP Mediator client certificate has expired. As a result, all further communication to the ONTAP Mediator will not be possible.	Remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

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

SMBC Server Certificate Expiring	WARNING	This message occurs when the ONTAP Mediator server certificate is due to expire within the next 30 days.	Before this certificate expires, remove the configuration of the current ONTAP Mediator by using the "snapmirror mediator remove" command. Update a new server certificate on the ONTAP Mediator server. Reconfigure access to the ONTAP Mediator by using the "snapmirror mediator add" command.

Additional Power, Heartbeat, and Miscellaneous System Monitors

Monitor Name	Severity	Monitor Description	Corrective Action
Disk Shelf Power Supply Discovered	INFORMATIONAL	This message occurs when a power supply unit is added to the disk shelf.	NONE
Disk Shelves Power Supply Removed	INFORMATIONAL	This message occurs when a power supply unit is removed from the disk shelf.	NONE
MetroCluster Automatic Unplanned Switchover Disabled	CRITICAL	This message occurs when automatic unplanned switchover capability is disabled.	Run the "metrocluster modify -node-name <nodename> -automatic -switchover-onfailure true" command for each node in the cluster to enable automatic switchover.
MetroCluster Storage Bridge Unreachable	CRITICAL	The storage bridge is not reachable over the management network	1) If the bridge is monitored by SNMP, verify that the node management LIF is up by using the "network interface show" command. Verify that the bridge is alive by using the "network ping" command. 2) If the bridge is monitored in-band, check the fabric cabling to the bridge, and then verify that the bridge is powered up.

Monitor Name	Severity	Monitor Description	Corrective Action
MetroCluster Bridge Temperature Abnormal - Below Critical	CRITICAL	The sensor on the Fibre Channel bridge is reporting a temperature that is below the critical threshold.	1) Check the operational status of the fans on the storage bridge. 2) Verify that the bridge is operating under recommended temperature conditions.
MetroCluster Bridge Temperature Abnormal - Above Critical	CRITICAL	The sensor on the Fibre Channel bridge is reporting a temperature that is above the critical threshold.	1) Check the operational status of the chassis temperature sensor on the storage bridge using the command "storage bridge show -cooling". 2) Verify that the storage bridge is operating under recommended temperature conditions.
MetroCluster Aggregate Left Behind	WARNING	The aggregate was left behind during switchback.	1) Check the aggregate state by using the command "aggr show". 2) If the aggregate is online, return it to its original owner by using the command "metrocluster switchback".

Monitor Name	Severity	Monitor Description	Corrective Action
All Links Between Metrocluster Partners Down	CRITICAL	RDMA interconnect adapters and intercluster LIFs have broken connections to the peered cluster or the peered cluster is down.	1) Ensure that the intercluster LIFs are up and running. Repair the intercluster LIFs if they are down. 2) Verify that the peered cluster is up and running by using the "cluster peer ping" command. See the MetroCluster Disaster Recovery Guide if the peered cluster is down. 3) For fabric MetroCluster, verify that the back-end fabric ISLs are up and running. Repair the back-end fabric ISLs if they are down. 4) For non-fabric MetroCluster configurations, verify that the cabling is correct between the RDMA interconnect adapters. Reconfigure the cabling if the links are down.
MetroCluster Partners Not Reachable Over Peering Network	CRITICAL	The connectivity to the peer cluster is broken.	1) Ensure that the port is connected to the correct network/switch. 2) Ensure that the intercluster LIF is connected with the peered cluster. 3) Ensure that the peered cluster is up and running by using the command "cluster peer ping". Refer to the MetroCluster Disaster Recovery Guide if the peered cluster is down.
MetroCluster Inter-Switch All Links Down	CRITICAL	All Inter-Switch Links (ISLs) on the storage switch are down.	1) Repair the back-end fabric ISLs on the storage switch. 2) Ensure that the partner switch is up and its ISLs are operational. 3) Ensure that intermediate equipment, such as xWDM devices, are operational.

Monitor Name	Severity	Monitor Description	Corrective Action
MetroCluster Node To Storage Stack SAS Link Down	WARNING	The SAS adapter or its attached cable might be at fault.	1. Verify that the SAS adapter is online and running. 2. Verify that the physical cable connection is secure and operating, and replace the cable if necessary. 3. If the SAS adapter is connected to disk shelves, ensure IOMs and disks are properly seated.
MetroCluster FC Initiator Link Down	CRITICAL	The FC initiator adapter is at fault.	1. Ensure that the FC initiator link has not been tampered with. 2. Verify the operational status of the FC initiator adapter by using the command "system node run -node local -command storage show adapter".
MetroCluster FC-VI Interconnect Link Down	CRITICAL	The physical link on the FC-VI port is offline.	1. Ensure that the FC-VI link has not been tampered with. 2. Verify that the physical status of the FC-VI adapter is "Up" by using the command "metrocluster interconnect adapter show". 3. If the configuration includes fabric switches, ensure that they are properly cabled and configured.
MetroCluster Spare Disks Left Behind	WARNING	The spare disk was left behind during switchback.	If the disk is not failed, return it to its original owner by using the command "metrocluster switchback".
MetroCluster Storage Bridge Port Down	CRITICAL	The port on the storage bridge is offline.	1) Check the operational status of the ports on the storage bridge by using the command "storage bridge show -ports". 2) Verify logical and physical connectivity to the port.

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

Monitor Name	Severity	Monitor Description	Corrective Action
MetroCluster Switch Temperature Abnormal	CRITICAL	The temperature sensor on the Fibre Channel switch reported abnormal temperature.	1) Check the operational status of the temperature sensors on the storage switch by using the command "storage switch show -cooling". 2) Verify that the switch is operating under recommended temperature conditions.
Service Processor Heartbeat Stopped	WARNING	This message occurs when ONTAP is no longer receiving heartbeats from the Service Processor (SP). Depending on the hardware design, the system may continue to serve data or may determine to shut down to prevent data loss or hardware damage. The system continues to serve data, but because the SP might not be working, the system cannot send notifications of down appliances, boot errors, or Open Firmware (OFW) Power-On Self-Test (POST) errors. If your system is configured to do so, it generates and transmits an AutoSupport (or 'call home') message to NetApp technical support and to the configured destinations. Successful delivery of an AutoSupport message significantly improves problem determination and resolution.	If the system has shut down, attempt a hard power cycle: Pull the controller out from the chassis, push it back in then power on the system. Contact NetApp technical support if the problem persists after the power cycle, or for any other condition that may warrant attention.

More Information

- Viewing and Dismissing Alerts

Webhook Notifications

Notification using Webhooks

Webhooks allow users to send alert notifications to various applications using a customized webhook channel.

Many commercial applications support webhooks as a standard input interface, for example: Slack, PagerDuty, Teams, and Discord all support webhooks. By supporting a generic, customizable webhook channel, Data Infrastructure Insights can support many of these delivery channels. Information on webhooks can be found on these application websites. For example, Slack provides this useful guide.

You can create multiple webhook channels, each channel targeted for a different purpose; separate applications, different recipients, etc.

The webhook channel instance is comprised of the following elements:

Name	Unique name
URL	Webhook target URL, including the *http://* or *https://* prefix along with the url params
Method	GET, POST - Default is POST
Custom Header	Specify any custom header lines here
Message Body	Put the body of your message here
Default Alert Parameters	Lists the default parameters for the webhook
Custom Parameters and Secrets	Custom parameters and secrets allow you to add unique parameters and secure elements such as passwords

Creating a Webhook

To create a Data Infrastructure Insights webhook, go to **Admin > Notifications** and select the **Webhooks** tab.

The following image shows an example webhook configured for Slack:

Edit a Webhook

Name

Slack Test

Template Type

Slack

URL

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

Method

POST

Custom Header

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

Message Body

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

Cancel | Test Webhook | Save Webhook

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

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

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

Default Alert Parameters

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

Custom Parameters and Secrets

Name	Value	Description
No Data Available		

+ Parameter

Parameters: What are they and how do I use them?

Alert Parameters are dynamic values populated per alert. For example, the *%%TriggeredOn%%* parameter will be replaced with the object on which the alert was triggered.

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

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

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

Custom Parameters and Secrets

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

Webhooks List Page

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

Choosing Webhook Notification in a Monitor

To choose the webhook notification in a monitor, go to **Alerts > Manage Monitors** and select the desired monitor, or add a new monitor. In the *Set up team notifications* section, choose *Webhook* as the delivery method. Select the alert levels (Critical, Warning, Resolved), then choose the desired webhook.

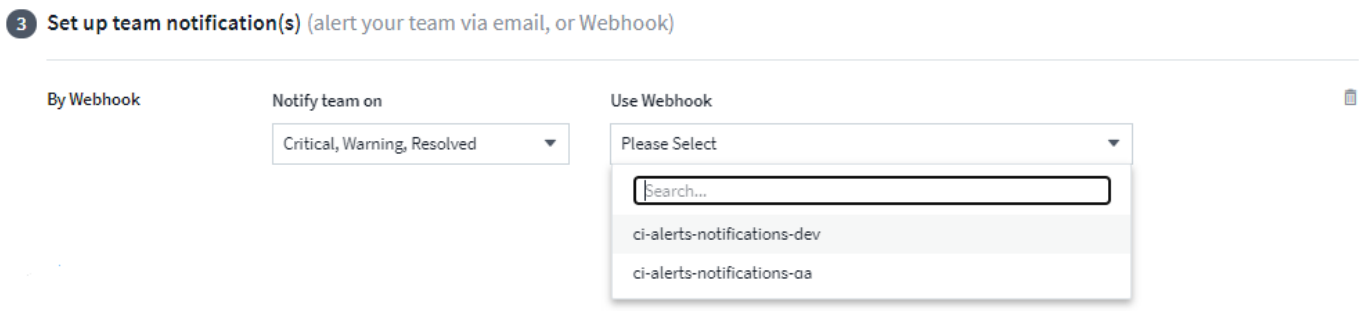

Webhook Examples:

Webhooks for Slack
Webhooks for PagerDuty
Webhooks for Teams
Webhooks for Discord

Webhook Example for Discord

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

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

Discord Setup:

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

configuration.

Create Data Infrastructure Insights Webhook:

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

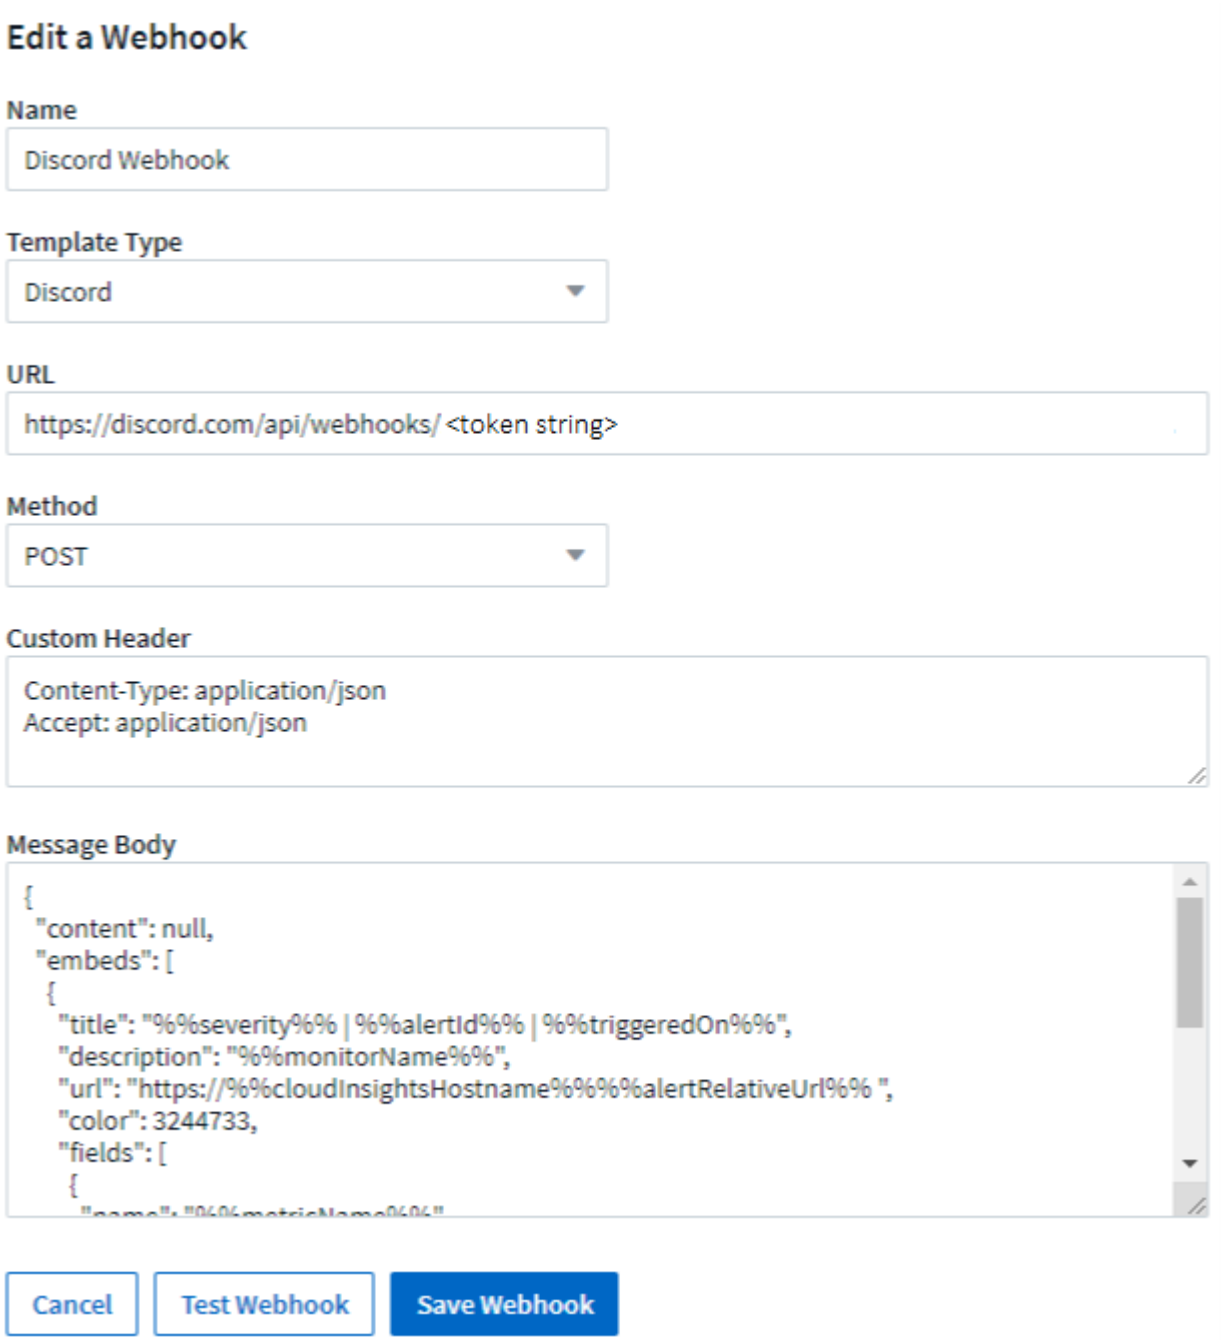

In order to test the webhook, temporarily replace the url value in the message body with any valid URL (such as *https://netapp.com*) then click the *Test Webhook* button. Be sure to set the

message body back once the test completes.

Notifications via Webhook

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

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

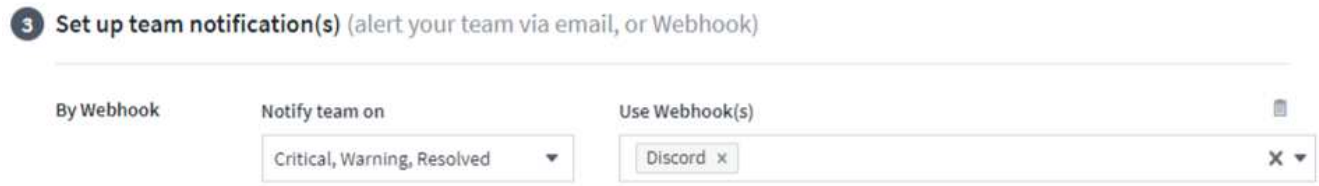

Webhook Example for PagerDuty

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

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

PagerDuty Setup:

1. In PagerDuty, navigate to **Services > Service Directory** and click on the **+New Service** button
2. Enter in a *Name* and a meaningful *description*. Click *Next*.

   [PagerDuty screen]

3. On the next screen, select an existing escalation policy, or generate a new one. Click *Next*.

   [Assigning an Escalation Policy]

4. Select AIOps features as desired. Click *Next*.

   [aiops service selection screen]

5. Select the Integration(s) you wish to use to feed alerts to PagerDuty. For DII, use the Events API V2 Integration.

   [PagerDuty integrations selection screen]

On the Integration configuration screen, make note of the **Integration Key**. You will need this key when you create the Data Infrastructure Insights webhook below.

+

[Copying the Integration Key]

Create Data Infrastructure Insights Webhook:

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

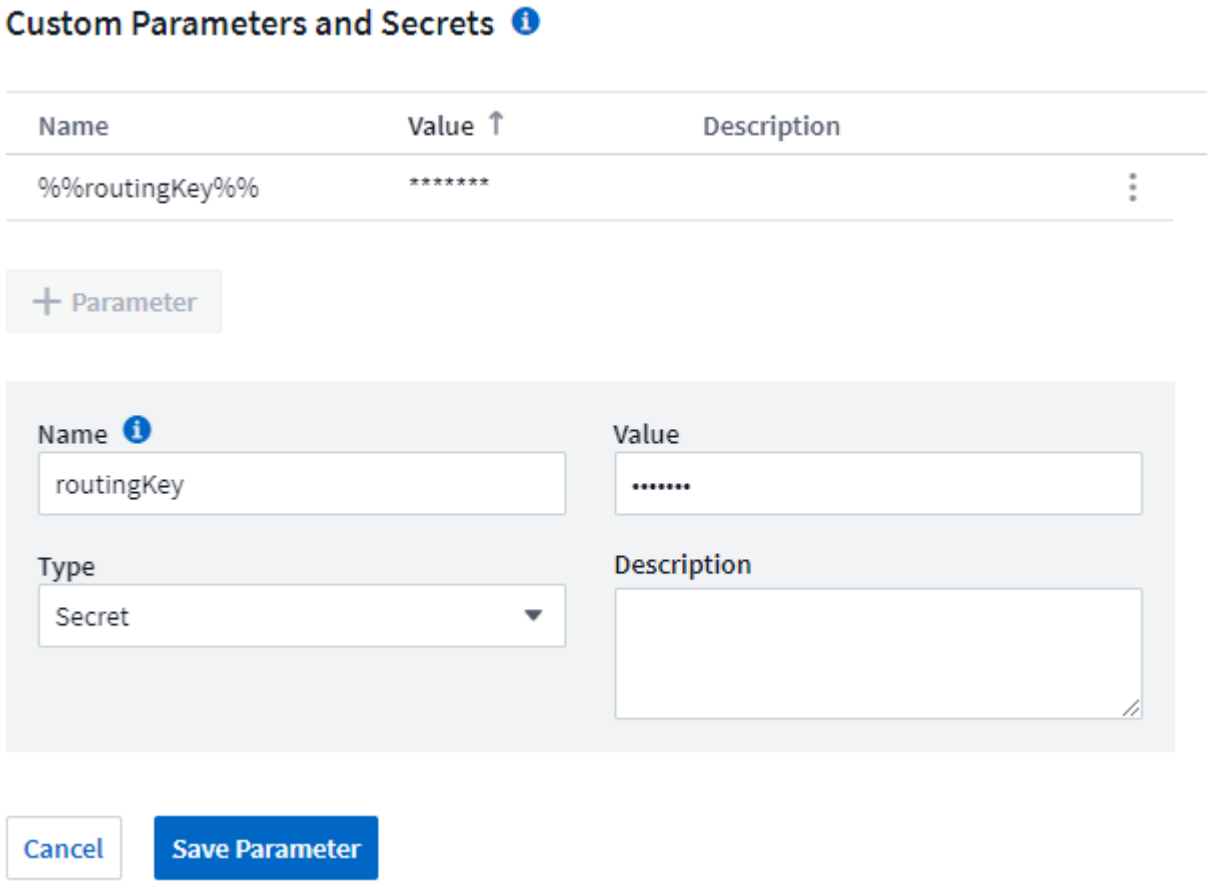

Repeat these steps to create a "PagerDuty Resolve" webhook for resolved events.

PagerDuty to Data Infrastructure Insights Field Mapping

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

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

PagerDuty	Data Infrastructure Insights
Class	Monitor Name

Message Body

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

Notifications via Webhook

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

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

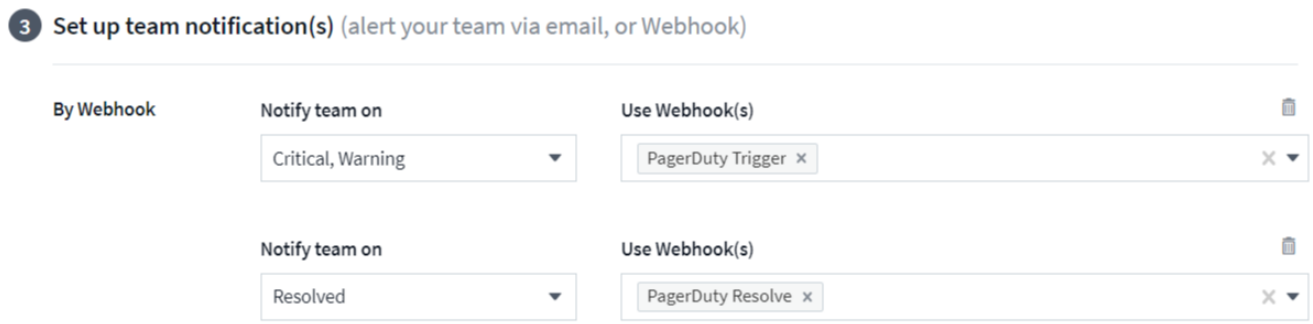

Setting separate notifications for trigger events versus resolved events is a best practice, since

> PagerDuty handles trigger events differently than resolved events.

Webhook Example for Slack

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

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

Slack Example:

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

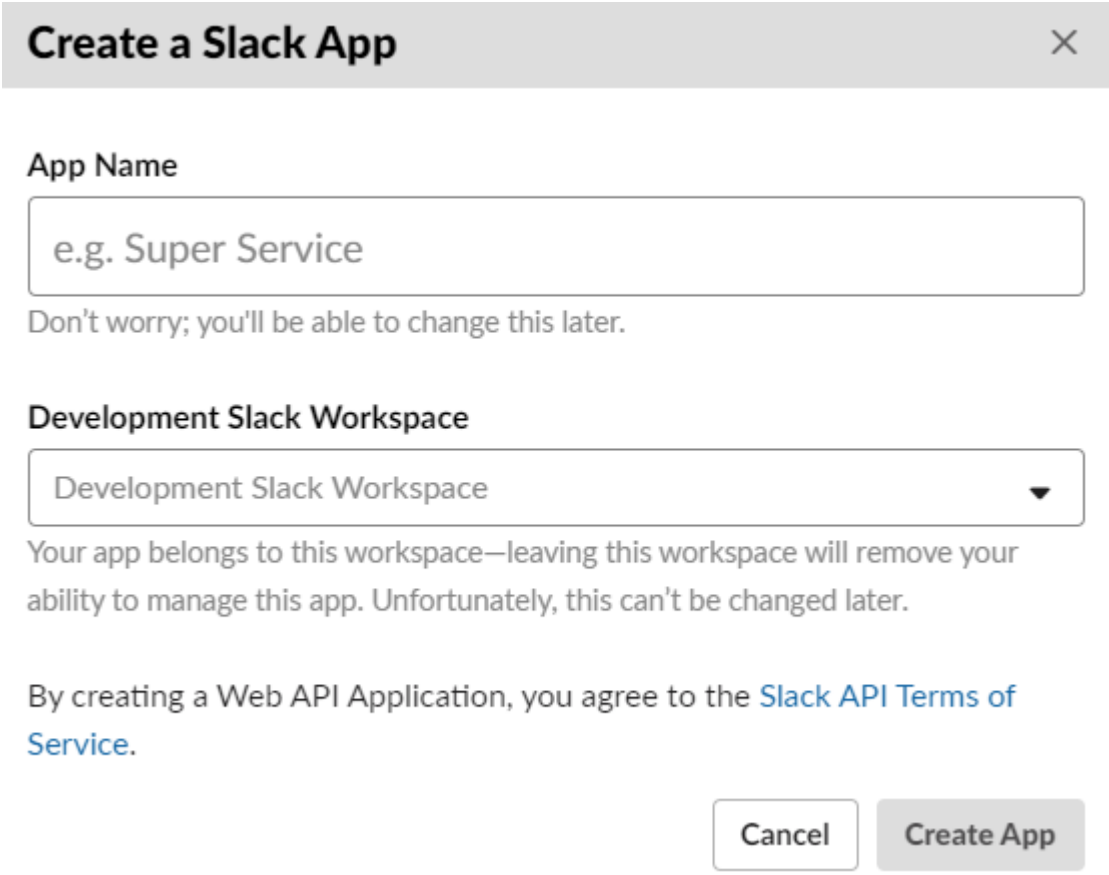

- Go to Incoming Webhooks, click on *Activate Incoming Webhooks*, Request to *Add New Webhook*, and select the Channel on which to Post.
- Copy the Webhook URL. You will need to paste this into the Data Infrastructure Insights webhook configuration.

Create Data Infrastructure Insights Webhook:

1. In Data Infrastructure Insights, navigate to **Admin > Notifications** and select the **Webhooks** tab. Click **+Webhook** to create a new webhook.

2. Give the webhook a meaningful Name, such as "Slack Webhook".
3. In the *Template Type* drop-down, select **Slack**.
4. Paste the URL from above into the *URL* field.

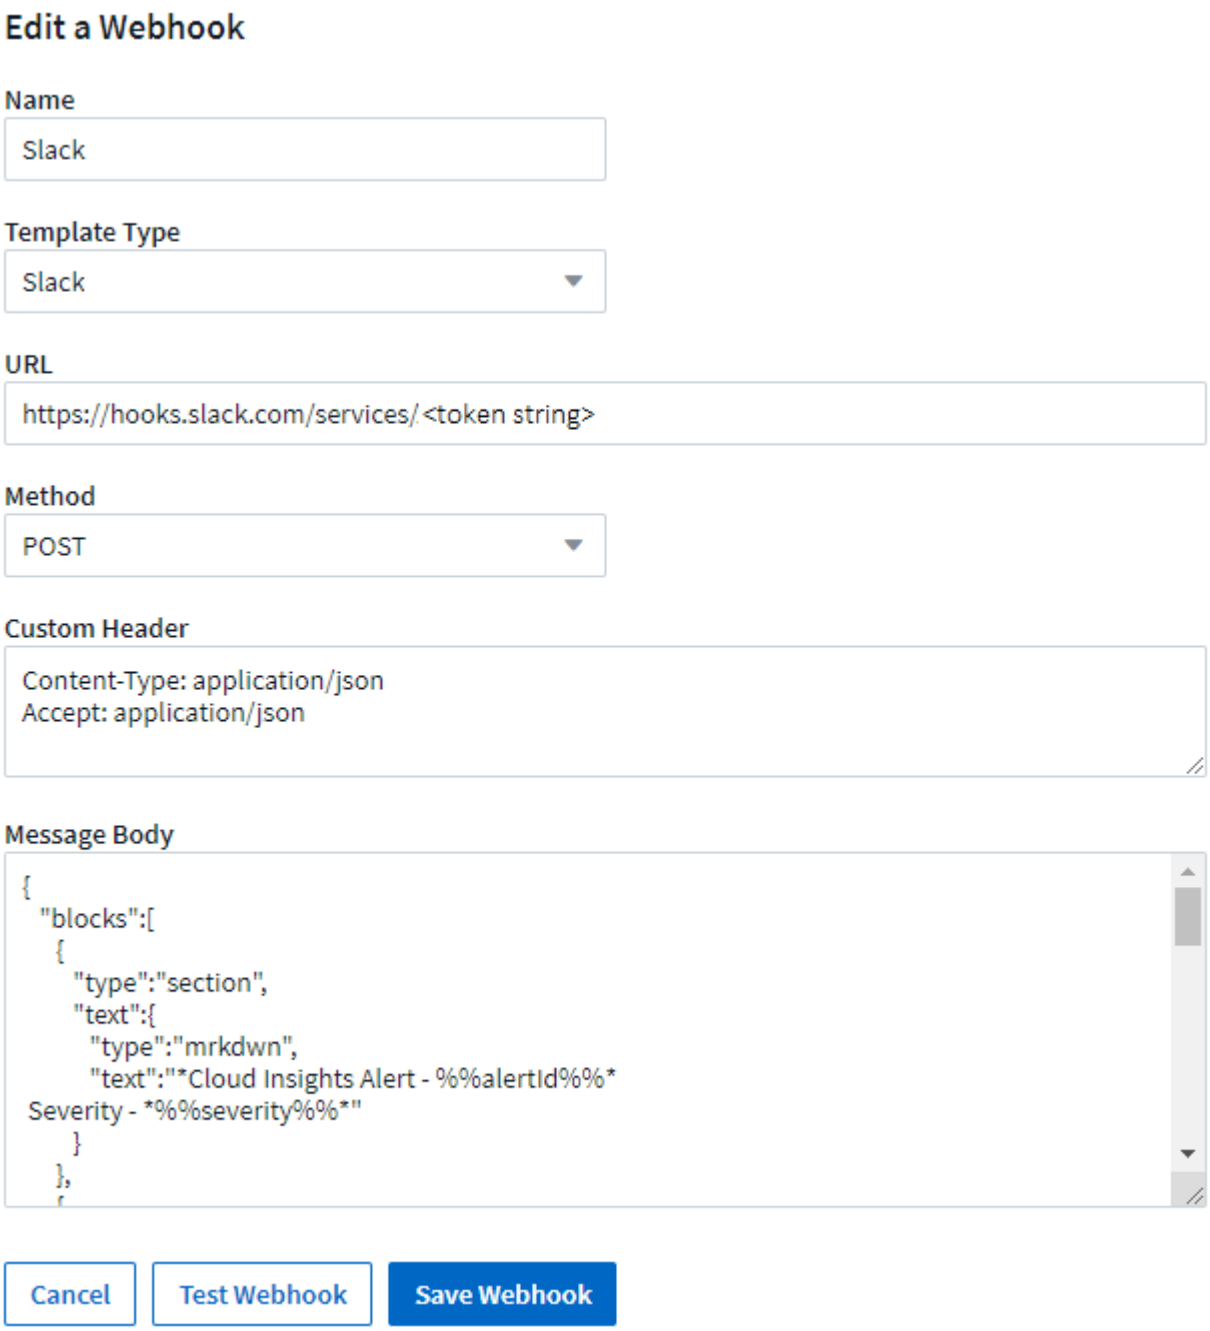

Notifications via Webhook

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

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

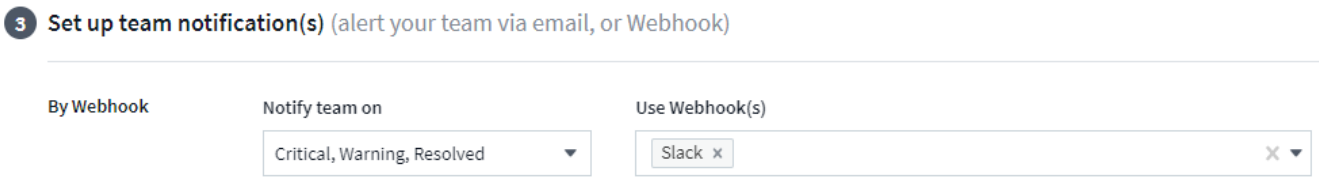

More information:

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

Webhook Example for Microsoft Teams

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

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

Teams Setup:

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

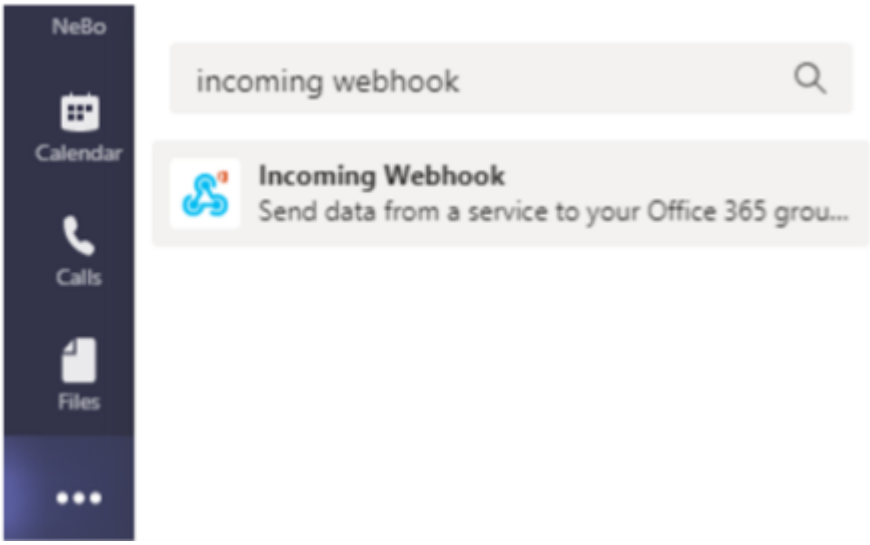

2. Select **Add to a Team > Select a Team > Setup a Connector**.
3. Copy the Webhook URL. You will need to paste this into the Data Infrastructure Insights webhook configuration.

Create Data Infrastructure Insights Webhook:

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

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

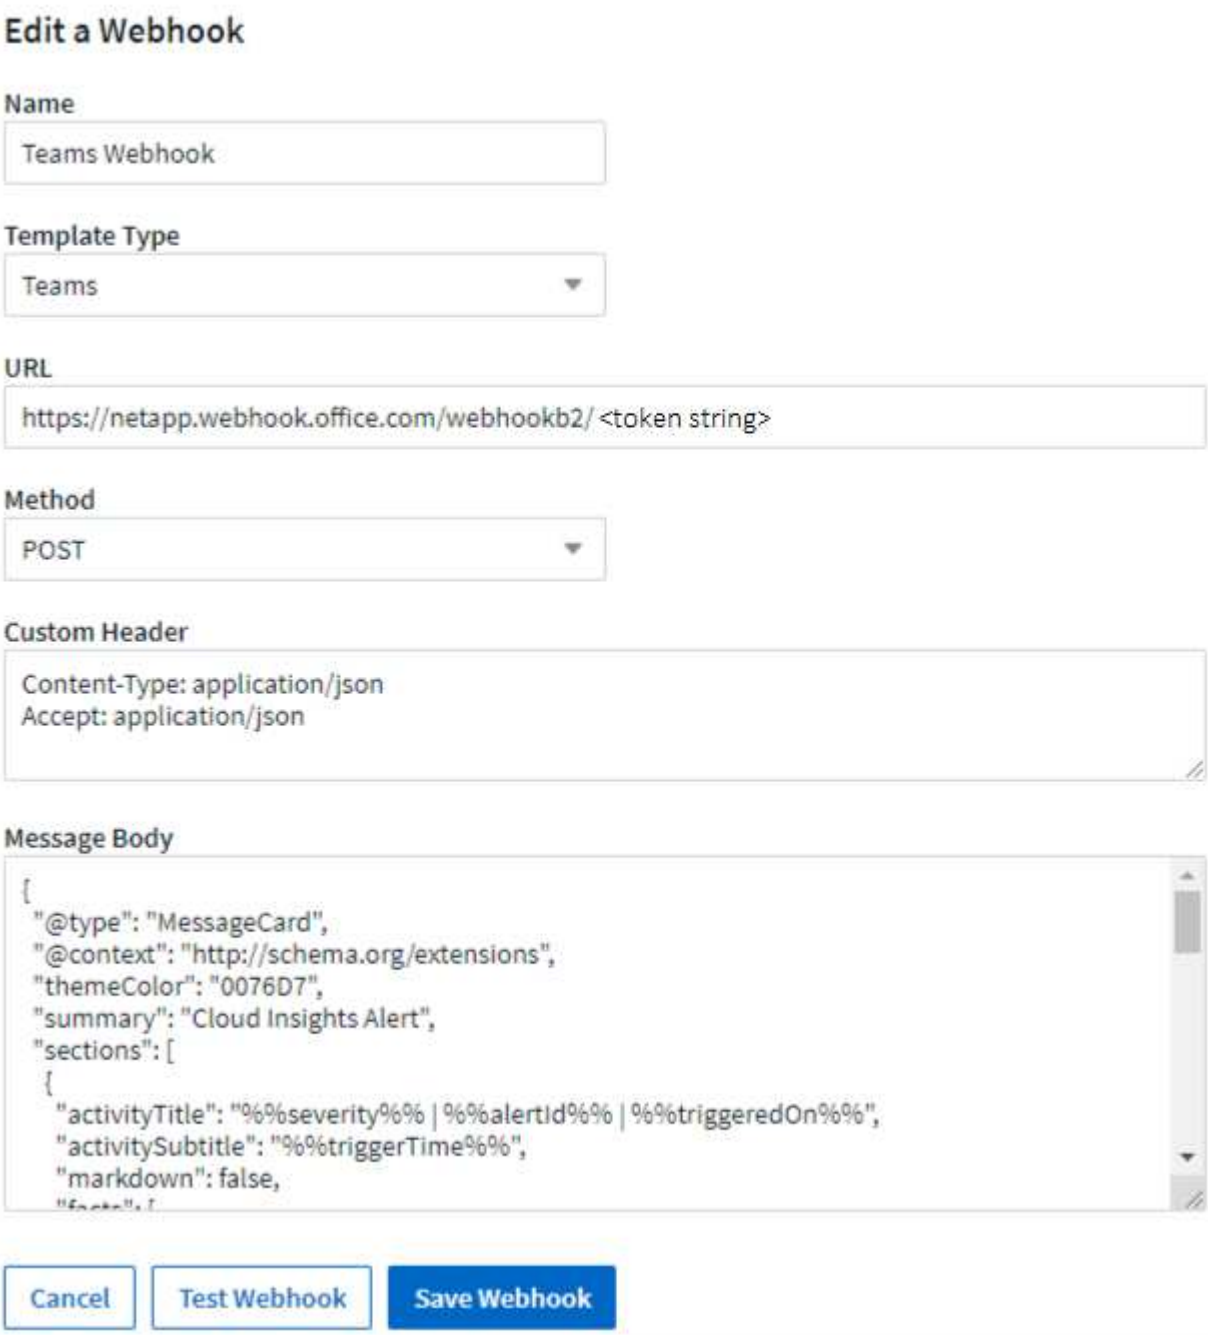

1. Paste the URL from above into the *URL* field.

Notifications via Webhook

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

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

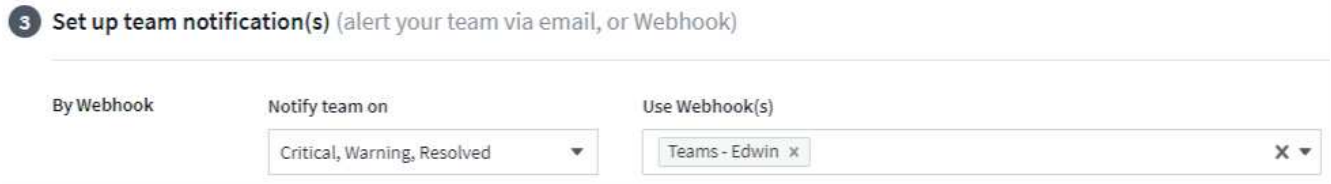

Configuring a Login Message

You can configure a message to display as a pop-up for all users when they log in. You choose to display only for the next login, or for all subsequent logins.

An administrator can set the message in Admin > Notifications, on the "Login Banner" tab.

The message supports basic Markdown, so you can apply formatting such as bold, italics, tables, etc.

Choose whether the users sees the message only on their next login, or on every subsequent login.

Note that the user *must* acknowledge the login message in order to dismiss it. Once dismissed, any additional pop-ups (What’s New updates, Terms-of-Service updates, etc.) will display.

Copyright information

Copyright © 2026 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.