



Event, performance, and health monitoring

ONTAP 9

NetApp
April 16, 2024

Table of Contents

- Event, performance, and health monitoring 1
 - Monitor cluster performance with System Manager 1
 - Monitor and manage cluster performance using the CLI 11
 - Monitor cluster performance with Unified Manager 47
 - Monitor cluster performance with Cloud Insights 47
- Audit logging 48
- AutoSupport 53
- Health monitoring 82
- File System Analytics 95
- EMS configuration 109

Event, performance, and health monitoring

Monitor cluster performance with System Manager

Monitor cluster performance using System Manager

The topics in this section show you how to manage cluster health and performance with System Manager in ONTAP 9.7 and later releases.

You can monitor cluster performance by viewing information about your system on the System Manager Dashboard. The Dashboard displays information about important alerts and notifications, the efficiency and capacity of storage tiers and volumes, the nodes that are available in a cluster, the status of the nodes in an HA pair, the most active applications and objects, and the performance metrics of a cluster or a node.

The Dashboard lets you determine the following information:

- **Health:** How healthy is the cluster?
- **Capacity:** What capacity is available on the cluster?
- **Performance:** How well is the cluster performing, based on latency, IOPS, and throughput?
- **Network:** How is the network configured with hosts and storage objects, such as ports, interfaces, and storage VMs?

In the Health and Capacity overviews, you can click [→](#) to view additional information and perform tasks.

In the Performance overview, you can view metrics based on the hour, the day, the week, the month, or the year.

In the Network overview, the number of each object in the network is displayed (for example, "8 NVMe/FC ports"). You can click on the numbers to view details about each network object.

View performance on cluster dashboard

Use the dashboard to make informed decisions about workloads you might want to add or move. You can also look at peak usage times to plan for potential changes.

The performance values refresh every 3 seconds and the performance graph refreshes every 15 seconds.

Steps

1. Click **Dashboard**.
2. Under **Performance**, select the interval.

Identify hot volumes and other objects

Accelerate your cluster performance by identifying the frequently accessed volumes (hot volumes) and data (hot objects).



Beginning in ONTAP 9.10.1, you can use the Activity Tracking feature in File System Analytics to monitor hot objects in a volume.


Steps

1. Click **Storage > Volumes**.
2. Filter the IOPS, latency, and throughput columns to view the frequently accessed volumes and data.

Modify QoS

Beginning with ONTAP 9.8, when you provision storage, [Quality of Service \(QoS\)](#) is enabled by default. You can disable QoS or choose a custom QoS policy during the provisioning process. You can also modify QoS after your storage has been provisioned.

Steps

1. In System Manager, select **Storage** then **Volumes**.
2. Next to the volume for which you want to modify QoS, select  then **Edit**.

Monitor risks

Beginning with ONTAP 9.10.0, you can use System Manager to monitor the risks reported by Active IQ Digital Advisor. Beginning with ONTAP 9.10.1, you can use System Manager to also acknowledge the risks.

NetApp Active IQ Digital Advisor reports opportunities to reduce risk and improve the performance and efficiency of your storage environment. With System Manager, you can learn about risks reported by Active IQ and receive actionable intelligence that helps you administer storage and achieve higher availability, improved security, and better storage performance.

Link to your Active IQ account

To receive information about risks from Active IQ, you should first link to your Active IQ account from System Manager.

Steps

1. In System Manager, click **Cluster > Settings**.
2. Under **Active IQ Registration**, click **Register**.
3. Enter your credentials for Active IQ.
4. After your credentials are authenticated, click **Confirm to link Active IQ with System Manager**.

View the number of risks

Beginning with ONTAP 9.10.0, you can view from the dashboard in System Manager the number of risks reported by Active IQ.

Before you begin

You must establish a connection from System Manager to your Active IQ account. Refer to [Link to your Active IQ account](#).

Steps

1. In System Manager, click **Dashboard**.
2. In the **Health** section, view the number of reported risks.



You can view more detailed information about each risk by clicking the message showing the number of risks. See [View details of risks](#).

View details of risks

Beginning with ONTAP 9.10.0, you can view from System Manager how the risks reported by Active IQ are categorized by impact areas. You can also view detailed information about each reported risk, its potential impact on your system, and corrective actions you can take.

Before you begin

You must establish a connection from System Manager to your Active IQ account. Refer to [Link to your Active IQ account](#).

Steps

1. Click **Events > All Events**.
2. In the **Overview** section, under **Active IQ Suggestions**, view the number of risks in each impact area category. The risk categories include:
 - Performance & efficiency
 - Availability & protection
 - Capacity
 - Configuration
 - Security
3. Click on the **Active IQ Suggestions** tab to view information about each risk, including the following:
 - Level of impact to your System
 - Category of the risk
 - Nodes that are affected
 - Type of mitigation needed
 - Corrective actions you can take

Acknowledge risks

Beginning with ONTAP 9.10.1, you can use System Manager to acknowledge any of the open risks.

Steps

1. In System Manager, display the list of risks by performing the procedure in [View details of risks](#).
2. Click on the risk name of an open risk that you want to acknowledge.
3. Enter information into the following fields:
 - Reminder (date)
 - Justification
 - Comments
4. Click **Acknowledge**.



After you acknowledge a risk, it takes a few minutes for the change to be reflected in the list of Active IQ suggestions.

Unacknowledge risks

Beginning with ONTAP 9.10.1, you can use System Manager to unacknowledge any risk that was previously acknowledged.

Steps

1. In System Manager, display the list of risks by performing the procedure in [View details of risks](#).
2. Click on the risk name of an acknowledged risk that you want to unacknowledge.
3. Enter information into the following fields:
 - Justification
 - Comments
4. Click **Unacknowledge**.



After you unacknowledge a risk, it takes a few minutes for the change to be reflected in the list of Active IQ suggestions.

System Manager insights

Beginning with ONTAP 9.11.1, System Manager displays *insights* that help you optimize the performance and security of your system.



To view, customize, and respond to insights, refer to [Gain insights to help optimize your system](#)

Capacity insights

System Manager can display the following insights in response to capacity conditions in your system:

Insight	Severity	Condition	Fixes
Local tiers are lacking space	Remediate risks	One or more local tiers are more than 95% full and quickly growing. Existing workloads might be unable to grow, or in extreme cases, existing workloads might run out of space and fail.	Recommended fix: Perform one of following options. <ul style="list-style-type: none">• Clear the volume recovery queue.• Enable thin provisioning on thick provisioned volumes to free up trapped storage.• Move volumes to another local tier.• Delete unneeded Snapshot copies.• Delete unneeded directories or files in the volumes.• Enable Fabric Pool to tier the data to the cloud.

Applications are lacking space	Needs attention	One or more volumes are more than 95% full, but they do not have autogrow enabled.	<p>Recommended: Enable autogrow up to 150% of current capacity.</p> <p>Other options:</p> <ul style="list-style-type: none"> • Reclaim space by deleting Snapshot copies. • Resize the volumes. • Delete directories or files.
FlexGroup volume's capacity is imbalanced	Optimize storage	The size of the constituent volumes of one or more FlexGroup volumes has grown unevenly over time, leading to an imbalance in capacity usage. If the constituent volumes become full, write failures could occur.	Recommended: Rebalance the FlexGroup volumes.
Storage VMs are running out of capacity	Optimize storage	One or more storage VMs are near their maximum capacity. You will not be able to provision more space for new or existing volumes if the storage VMs reach maximum capacity.	Recommended: If possible, increase the maximum capacity limit of the storage VM.

Security insights

System Manager can display the following insights in response to conditions that might jeopardize the security of your data or your system.

Insight	Severity	Condition	Fixes
Volumes are still in anti-ransomware learning mode	Needs attention	One or more volumes have been in the anti-ransomware learning mode for 90 days.	Recommended: Enable the anti-ransomware active mode for those volumes.

Automatic deletion of Snapshot copies is enabled on volumes	Needs attention	Snapshot auto-deletion is enabled on one or more volumes.	Recommended: Disable the automatic deletion of Snapshot copies. Otherwise, in case of a ransomware attack, data recovery for these volumes might not be possible.
Volumes don't have Snapshot policies	Needs attention	One or more volumes don't have an adequate Snapshot policy attached to them.	Recommended: Attach a Snapshot policy to volumes that don't have one. Otherwise, in case of a ransomware attack, data recovery for these volumes might not be possible.
Native FPolicy is not configured	Best practice	Native FPolicy is not configured on one or more NAS storage VMs.	Recommended: IMPORTANT: Blocking extensions might lead to unexpected results. Beginning in 9.11.1, you can enable native FPolicy for storage VMs, which blocks over 3000 file extensions known to be used for ransomware attacks. Configure native FPolicy in NAS storage VMs to control the file extensions that are allowed or not allowed to be written on volumes in your environment.
Telnet is enabled	Best practice	Secure Shell (SSH) should be used for secure remote access.	Recommended: Disable Telnet and use SSH for secure remote access.
Too few NTP servers are configured	Best practice	The number of servers configured for NTP is less than 3.	Recommended: Associate at least three NTP servers with the cluster. Otherwise, problems can occur with the synchronization of the cluster time.
Remote Shell (RSH) is enabled	Best practice	Secure Shell (SSH) should be used for secure remote access.	Recommended: Disable RSH and use SSH for secure remote access.
Login banner isn't configured	Best practice	Login messages are not configured either for the cluster, for the storage VM, or for both.	Recommended: Setup the login banners for the cluster and the storage VM and enable their use.
AutoSupport is using a nonsecure protocol	Best practice	AutoSupport is not configured to communicate via HTTPS.	Recommended: It is strongly recommended to use HTTPS as the default transport protocol to send AutoSupport messages to technical support.

Default admin user is not locked	Best practice	Nobody has logged in using a default administrative account (admin or diag), and these accounts are not locked.	Recommended: Lock default administrative accounts when they are not being used.
Secure Shell (SSH) is using nonsecure ciphers	Best practice	The current configuration uses nonsecure CBC ciphers.	Recommended: You should allow only secure ciphers on your web server to protect secure communication with your visitors. Remove ciphers that have names containing "cbc", such as "ais128-cbc", "aes192-cbc", "aes256-cbc", and "3des-cbc".
Global FIPS 140-2 compliance is disabled	Best practice	Global FIPS 140-2 compliance is disabled on the cluster.	Recommended: For security reasons, you should enable Global FIPS 140-2 compliant cryptography to ensure ONTAP can safely communicate with external clients or server clients.
Volumes aren't being monitored for ransomware attacks	Needs attention	Anti-ransomware is disabled on one or more volumes.	Recommended: Enable anti-ransomware on the volumes. Otherwise, you might not notice when volumes are being threatened or under attack.
Storage VMs aren't configured for anti-ransomware	Best practice	One or more storage VMs aren't configured for anti-ransomware protection.	Recommended: Enable anti-ransomware on the storage VMs. Otherwise, you might not notice when storage VMs are being threatened or under attack.

Configuration insights

System Manager can display the following insights in response to concerns about the configuration of your system.

Insight	Severity	Condition	Fixes
Cluster isn't configured for notifications	Best practice	Email, webhooks, or an SNMP trap host is not configured to let you receive notifications about problems with the cluster.	Recommended: Configure notifications for the cluster.

Cluster isn't configured for automatic updates.	Best practice	The cluster hasn't been configured to receive automatic updates for the latest disk qualification package, disk firmware, shelf firmware, and SP/BMC firmware files when they are available.	Recommended: Enable this feature.
Cluster firmware isn't up-to-date	Best practice	Your system doesn't have the latest update to the firmware which could have improvements, security patches, or new features that help secure the cluster for better performance.	Recommended: Update the ONTAP firmware.

Gain insights to help optimize your system

With System Manager, you can view insights that help you optimize your system.

About this task

Beginning with ONTAP 9.11.0, you can view insights in System Manager that help you optimize the capacity and security compliance of your system.

Beginning with ONTAP 9.11.1, you can view additional insights that help you optimize the capacity, security compliance, and configuration of your system.



Blocking extensions might lead to unexpected results. Beginning with ONTAP 9.11.1, you can enable native FPolicy for storage VMs using System Manager. You might receive a System Manager Insight message recommending that you [configure native FPolicy](#) for a storage VM.

With FPolicy Native Mode, you can allow or disallow specific file extensions. System Manager recommends over 3000 disallowed file extensions that have been used in past ransomware attacks. Some of these extensions might be used by legitimate files in your environment and blocking them might lead to unexpected issues.

Therefore, it is strongly advised that you modify the list of extensions to meet the needs of your environment. Refer to [How to remove a file extension from a native FPolicy configuration created by System Manager using System Manager to recreate the policy](#).

To learn more about native FPolicy, see [FPolicy configuration types](#).

Based on best practices, these insights are displayed on one page from which you can initiate immediate

actions to optimize your system. For more detail about each insight, see [System Manager insights](#).





View optimization insights

Steps

1. In System Manager, click **Insights** in the left-hand navigation column.

The **Insights** page shows groups of insights. Each group of insights might contain one or more insights. The following groups are displayed:

- Needs your attention
 - Remediate risks
 - Optimize your storage
2. (Optional) Filter the insights that are displayed by clicking these buttons in the upper-right corner of the page:

-  Displays the security-related insights.
-  Displays the capacity-related insights.
-  Displays the configuration-related insights.
-  Displays all of the insights.

Respond to insights to optimize your system

In System Manager, you can respond to insights by either dismissing them, exploring different ways to remediate the problems, or initiating the process to fix the problems.

Steps

1. In System Manager, click **Insights** in the left-hand navigation column.
2. Hover over an insight to reveal the buttons to perform the following actions:
 - **Dismiss:** Remove the insight from the view. To “undismiss” the insight, refer to [Customize the settings for insights](#).
 - **Explore:** Find out various ways to remediate the problem mentioned in the insight. This button appears only if there is more than one method of remediation.
 - **Fix:** Initiate the process of remediating the problem mentioned in the insight. You will be asked to confirm whether you want to take the action needed to apply the fix.




Some of these actions can be initiated from other pages in System Manager, but the **Insights** page helps you streamline your day-to-day tasks by allowing you to initiate these action from this one page.

Customize the settings for insights

You can customize which insights you will be notified about in System Manager.

Steps

1. In System Manager, click **Insights** in the left-hand navigation column.
2. In the upper-right corner of the page, click , then select **Settings**.
3. On the **Settings** page, ensure there is a check in the check boxes next to the insights you want to be notified about. If you previously dismissed an insight, you can “undismiss” it by ensuring a check is in its check box.
4. Click **Save**.

Export the insights as a PDF file

You can export all applicable insights as a PDF file.

Steps

1. In System Manager, click **Insights** in the left-hand navigation column.
2. In the upper-right corner of the page, click , then select **Export**.

Configure native FPolicy

Beginning with ONTAP 9.11.1, when you receive a System Manager Insight that suggests implementing native FPolicy, you can configure it on your storage VMs and volumes.

Before you begin

When you access System Manager Insights, under **Apply best practices**, you might receive a message saying that native FPolicy is not configured.

To learn more about FPolicy configuration types, refer to [FPolicy configuration types](#).

Steps

1. In System Manager, click **Insights** in the left-hand navigation column.
2. Under **Apply best practices**, locate **Native FPolicy is not configured**.
3. Read the following message before taking action:



Blocking extensions might lead to unexpected results. Beginning with ONTAP 9.11.1, you can enable native FPolicy for storage VMs using System Manager. With FPolicy Native Mode, you can allow or disallow specific file extensions. System Manager recommends over 3000 disallowed file extensions that have been used in past ransomware attacks. Some of these extensions might be used by legitimate files in your environment and blocking them might lead to unexpected issues.

Therefore, it is strongly advised that you modify the list of extensions to meet the needs of your environment. Refer to [How to remove a file extension from a native FPolicy configuration created by System Manager using System Manager to recreate the policy](#).

4. Click **Fix**.
5. Select the storage VMs to which you want to apply the native FPolicy.
6. For each storage VM, select the volumes that will receive the native FPolicy.
7. Click **Configure**.

Monitor and manage cluster performance using the CLI

Performance monitoring and management overview

You can set up basic performance monitoring and management tasks and identify and resolve common performance issues.

You can use these procedures to monitor and manage cluster performance if the following assumptions apply to your situation:

- You want to use best practices, not explore every available option.
- You want to display system status and alerts, monitor cluster performance, and perform root-cause analysis by using Active IQ Unified Manager (formerly OnCommand Unified Manager), in addition to the ONTAP command-line interface.
- You are using the ONTAP command-line interface to configure storage quality of service (QoS).

QoS is also available in System Manager, NSLM, WFA, VSC (VMware Plug-in), and APIs.

- You want to install Unified Manager by using a virtual appliance, instead of a Linux or Windows-based installation.
- You're willing to use a static configuration rather than DHCP to install the software.
- You can access ONTAP commands at the advanced privilege level.
- You are a cluster administrator with the "admin" role.

Related information


If these assumptions are not correct for your situation, you should see the following resources:

- [Active IQ Unified Manager 9.8 Installation](#)
- [System administration](#)

Monitor performance

Performance monitoring and maintenance workflow overview

Monitoring and maintaining cluster performance involves installing Active IQ Unified Manager software, setting up basic monitoring tasks, identifying performance issues, and making adjustments as needed.



Verify that your VMware environment is supported

To successfully install Active IQ Unified Manager, you must verify that your VMware environment meets the necessary requirements.

Steps

1. Verify that your VMware infrastructure meets the sizing requirements for the installation of Unified Manager.
2. Go to the [Interoperability Matrix](#) to verify that you have a supported combination of the following components:
 - ONTAP version
 - ESXi operating system version
 - VMware vCenter Server version
 - VMware Tools version
 - Browser type and version



The [Interoperability Matrix](#) lists the supported configurations for Unified Manager.

3. Click the configuration name for the selected configuration.

Details for that configuration are displayed in the Configuration Details window.

4. Review the information in the following tabs:

- Notes

Lists important alerts and information that are specific to your configuration.

- Policies and Guidelines

Provides general guidelines for all configurations.

Active IQ Unified Manager worksheet

Before you install, configure, and connect Active IQ Unified Manager, you should have specific information about your environment readily available. You can record the information in the worksheet.

Unified Manager installation information

Virtual machine on which software is deployed	Your value
ESXi server IP address	
Host fully qualified domain name	
Host IP address	
Network mask	
Gateway IP address	
Primary DNS address	
Secondary DNS address	
Search domains	
Maintenance user name	
Maintenance user password	


Unified Manager configuration information

Setting	Your value
Maintenance user email address	
NTP server	

SMTP server host name or IP address	
SMTP user name	
SMTP password	
SMTP default port	25 (Default value)
Email from which alert notifications are sent	
LDAP bind distinguished name	
LDAP bind password	
Active Directory administrator name	
Active Directory password	
Authentication server base distinguished name	
Authentication server host name or IP address	

Cluster information

Capture the following information for each cluster on Unified Manager.

Cluster 1 of N	Your value
Host name or cluster-management IP address	
ONTAP administrator user name  The administrator must have been assigned the "admin" role.	
ONTAP administrator password	
Protocol (HTTP or HTTPS)	

Related information

[Administrator authentication and RBAC](#)

Install Active IQ Unified Manager

Download and deploy Active IQ Unified Manager

To install the software, you must download the virtual appliance (VA) installation file and then use a VMware vSphere Client to deploy the file to a VMware ESXi server. The VA is available in an OVA file.

Steps

1. Go to the **NetApp Support Site Software Download** page and locate Active IQ Unified Manager.

<https://mysupport.netapp.com/products/index.html>

2. Select **VMware vSphere** in the **Select Platform** drop-down menu and click **Go!**
3. Save the “OVA” file to a local or network location that is accessible to your VMware vSphere Client.
4. In VMware vSphere Client, click **File > Deploy OVF Template**.
5. Locate the “OVA” file and use the wizard to deploy the virtual appliance on the ESXi server.

You can use the **Properties** tab in the wizard to enter your static configuration information.

6. Power on the VM.
7. Click the **Console** tab to view the initial boot process.
8. Follow the prompt to install VMware Tools on the VM.
9. Configure the time zone.
10. Enter a maintenance user name and password.
11. Go to the URL displayed by the VM console.

Configure initial Active IQ Unified Manager settings

The Active IQ Unified Manager Initial Setup dialog box appears when you first access the web UI, which enables you to configure some initial settings and to add clusters.

Steps

1. Accept the default AutoSupport enabled setting.
2. Enter the NTP server details, the maintenance user email address, the SMTP server host name, and additional SMTP options, and then click **Save**.

After you finish

When the initial setup is complete, the Cluster Data Sources page is displayed where you can add the cluster details.

Specify the clusters to be monitored

You must add a cluster to an Active IQ Unified Manager server to monitor the cluster, view the cluster discovery status, and monitor its performance.

What you'll need

- You must have the following information:
 - Host name or cluster-management IP address

The host name is the fully qualified domain name (FQDN) or short name that Unified Manager uses to connect to the cluster. This host name must resolve to the cluster-management IP address.

The cluster-management IP address must be the cluster-management LIF of the administrative storage virtual machine (SVM). If you use a node-management LIF, the operation fails.

- ONTAP administrator user name and password
- Type of protocol (HTTP or HTTPS) that can be configured on the cluster and the port number of the cluster
- You must have the Application Administrator or Storage Administrator role.
- The ONTAP administrator must have the ONTAPI and SSH administrator roles.
- The Unified Manager FQDN must be able to ping ONTAP.

You can verify this by using the ONTAP command `ping -node node_name -destination Unified_Manager_FQDN`.

About this task

For a MetroCluster configuration, you must add both the local and remote clusters, and the clusters must be configured correctly.

Steps

1. Click **Configuration > Cluster Data Sources**.
2. From the Clusters page, click **Add**.
3. In the **Add Cluster** dialog box, specify the required values, such as the host name or IP address (IPv4 or IPv6) of the cluster, user name, password, protocol for communication, and port number.

By default, the HTTPS protocol is selected.

You can change the cluster-management IP address from IPv6 to IPv4 or from IPv4 to IPv6. The new IP address is reflected in the cluster grid and the cluster configuration page after the next monitoring cycle finishes.

4. Click **Add**.
5. If HTTPS is selected, perform the following steps:
 - a. In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information about the cluster.
 - b. Click **Yes**.

Unified Manager checks the certificate only when the cluster is initially added, but does not check it for each API call to ONTAP.

If the certificate has expired, you cannot add the cluster. You must renew the SSL certificate and then add the cluster.

6. **Optional:** View the cluster discovery status:
 - a. Review the cluster discovery status from the **Cluster Setup** page.

The cluster is added to the Unified Manager database after the default monitoring interval of approximately 15 minutes.

Set up basic monitoring tasks

Perform daily monitoring

You can perform daily monitoring to ensure that you do not have any immediate performance issues that require attention.

Steps

1. From the Active IQ Unified Manager UI, go to the **Event Inventory** page to view all current and obsolete events.
2. From the **View** option, select `Active Performance Events` and determine what action is required.

Use weekly and monthly performance trends to identify performance issues

Identifying performance trends can assist you in identifying whether the cluster is being overused or underused by analyzing volume latency. You can use similar steps to identify CPU, network, or other system bottlenecks.

Steps

1. Locate the volume that you suspect is being underused or overused.
2. On the **Volume Details** tab, click **30 d** to display the historical data.
3. In the "Break down data by" drop-down menu, select **Latency**, and then click **Submit**.
4. Deselect **Aggregate** in the cluster components comparison chart, and then compare the cluster latency with the volume latency chart.
5. Select **Aggregate** and deselect all other components in the cluster components comparison chart, and then compare the aggregate latency with the volume latency chart.
6. Compare the reads/writes latency chart to the volume latency chart.
7. Determine whether client application loads have caused a workload contention and rebalance workloads as needed.
8. Determine whether the aggregate is overused and causing contention and rebalance workloads as needed.

Use performance thresholds to generate event notifications

Events are notifications that the Active IQ Unified Manager generates automatically when a predefined condition occurs, or when a performance counter value crosses a threshold. Events help you identify performance issues in the clusters you are monitoring. You can configure alerts to send email notification automatically when events of certain severity types occur.

Set performance thresholds

You can set performance thresholds to monitor critical performance issues. User-defined thresholds trigger a warning or a critical event notification when the system approaches or exceeds the defined threshold.

Steps

1. Create the Warning and Critical event thresholds:
 - a. Select **Configuration > Performance Thresholds**.
 - b. Click **Create**.
 - c. Select the object type and specify a name and description of the policy.
 - d. Select the object counter condition and specify the limit values that define Warning and Critical events.
 - e. Select the duration of time that the limit values must be breached for an event to be sent, and then click **Save**.
2. Assign the threshold policy to the storage object.
 - a. Go to the Inventory page for the same cluster object type that you previously selected and choose the **Performance** from the View option.
 - b. Select the object to which you want to assign the threshold policy, and then click **Assign Threshold Policy**.
 - c. Select the policy you previously created, and then click **Assign Policy**.

Example

You can set user-defined thresholds to learn about critical performance issues. For example, if you have a Microsoft Exchange Server and you know that it crashes if volume latency exceeds 20 milliseconds, you can set a warning threshold at 12 milliseconds and a critical threshold at 15 milliseconds. With this threshold setting, you can receive notifications when the volume latency exceeds the limit.

Object Counter Condition* Average Latency ms/op Warning 12 ms/op Critical 15 ms/op

Add alerts

You can configure alerts to notify you when a particular event is generated. You can configure alerts for a single resource, for a group of resources, or for events of a particular severity type. You can specify the frequency with which you want to be notified and associate a script to the alert.

What you'll need

- You must have configured notification settings such as the user email address, SMTP server, and SNMP trap host to enable the Active IQ Unified Manager server to use these settings to send notifications to users when an event is generated.
- You must know the resources and events for which you want to trigger the alert, and the user names or email addresses of the users that you want to notify.
- If you want to have a script execute based on the event, you must have added the script to Unified Manager by using the Scripts page.
- You must have the Application Administrator or Storage Administrator role.

About this task

You can create an alert directly from the Event details page after receiving an event in addition to creating an alert from the Alert Setup page, as described here.

Steps

1. In the left navigation pane, click **Storage Management > Alert Setup**.

2. In the **Alert Setup** page, click **Add**.
3. In the **Add Alert** dialog box, click **Name**, and enter a name and description for the alert.
4. Click **Resources**, and select the resources to be included in or excluded from the alert.

You can set a filter by specifying a text string in the **Name contains** field to select a group of resources. Based on the text string that you specify, the list of available resources displays only those resources that match the filter rule. The text string that you specify is case-sensitive.

If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes precedence over the include rule, and the alert is not generated for events related to the excluded resource.

5. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.



To select more than one event, press the Ctrl key while you make your selections.

6. Click **Actions**, and select the users that you want to notify, choose the notification frequency, choose whether an SNMP trap will be sent to the trap receiver, and assign a script to be executed when an alert is generated.



If you modify the email address that is specified for the user and reopen the alert for editing, the Name field appears blank because the modified email address is no longer mapped to the user that was previously selected. Also, if you modified the email address of the selected user from the Users page, the modified email address is not updated for the selected user.

You can also choose to notify users through SNMP traps.

7. Click **Save**.

Example of adding an alert

This example shows how to create an alert that meets the following requirements:

- Alert name: HealthTest
- Resources: includes all volumes whose name contains "abc" and excludes all volumes whose name contains "xyz"
- Events: includes all critical health events
- Actions: includes "sample@domain.com", a "Test" script, and the user has to be notified every 15 minutes

Perform the following steps in the Add Alert dialog box:

1. Click **Name**, and enter HealthTest in the **Alert Name** field.
2. Click **Resources**, and in the Include tab, select **Volumes** from the drop-down list.
 - a. Enter abc in the **Name contains** field to display the volumes whose name contains "abc".
 - b. Select <<All Volumes whose name contains 'abc'>> from the Available Resources area, and move it to the Selected Resources area.
 - c. Click **Exclude**, and enter xyz in the **Name contains** field, and then click **Add**.
3. Click **Events**, and select **Critical** from the Event Severity field.

4. Select **All Critical Events** from the Matching Events area, and move it to the Selected Events area.
5. Click **Actions**, and enter `sample@domain.com` in the Alert these users field.
6. Select **Remind every 15 minutes** to notify the user every 15 minutes.

You can configure an alert to repeatedly send notifications to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.

7. In the Select Script to Execute menu, select **Test** script.
8. Click **Save**.

Configure alert settings

You can specify which events from Active IQ Unified Manager trigger alerts, the email recipients for those alerts, and the frequency for the alerts.

What you'll need

You must have the Application Administrator role.

About this task

You can configure unique alert settings for the following types of performance events:

- Critical events triggered by breaches of user-defined thresholds
- Warning events triggered by breaches of user-defined thresholds, system-defined thresholds, or dynamic thresholds

By default, email alerts are sent to Unified Manager admin users for all new events. You can have email alerts sent to other users by adding those users' email addresses.



To disable alerts from being sent for certain types of events, you must clear all of the check boxes in an event category. This action does not stop events from appearing in the user interface.

Steps

1. In the left navigation pane, select **Storage Management > Alert Setup**.

The Alert Setup page is displayed.

2. Click **Add** and configure the appropriate settings for each of the event types.

To have email alerts sent to multiple users, enter a comma between each email address.

3. Click **Save**.

Identify performance issues in Active IQ Unified Manager

If a performance event occurs, you can locate the source of the issue within Active IQ Unified Manager and use other tools to fix it. You might receive an email notification of an event or notice the event during daily monitoring.

Steps

1. Click the link in the email notification, which takes you directly to the storage object having a performance event.

If you...	Then...
Receive an email notification of an event	Click the link to go directly to the event details page.
Notice the event while analyzing the Event Inventory page	Select the event to go directly to the event details page.

2. If the event has crossed a system-defined threshold, follow the suggested actions in the UI to troubleshoot the issue.
3. If the event has crossed a user-defined threshold, analyze the event to determine if you need to take action.
4. If the issue persists, check the following settings:
 - Protocol settings on the storage system
 - Network settings on any Ethernet or fabric switches
 - Network settings on the storage system
 - Disk layout and aggregate metrics on the storage system
5. If the issue persists, contact technical support for assistance.

Use Active IQ Digital Advisor to view system performance

For any ONTAP system that sends AutoSupport telemetry to NetApp, you can view extensive performance and capacity data. Active IQ shows system performance over a longer period than you can see in System Manager.

You can view graphs of CPU utilization, latency, IOPS, IOPS by protocol, and network throughput. You can also download this data in .csv format for analysis in other tools.

In addition to this performance data, Active IQ can show you storage efficiency by workload and compare that efficiency to the expected efficiency for that type of workload. You can view capacity trends and see an estimate of how much additional storage you might need to add in a given time frame.



- Storage Efficiency is available at the customer, cluster, and node level on the left-hand-side of the main dashboard.
- Performance is available at the cluster and node level on the left-hand-side of the main dashboard.


Related information

- [Active IQ Digital Advisor documentation](#)
- [Active IQ Digital Advisor video playlist](#)
- [Active IQ Web Portal](#)

Manage performance issues

Performance management workflow

Once you have identified a performance issue, you can conduct some basic diagnostic checks of your infrastructure to rule out obvious configuration errors. If those don't pinpoint the problem, you can start looking at workload management issues.



Perform basic infrastructure checks

Check protocol settings on the storage system

Check the NFS TCP maximum transfer size

For NFS, you can check whether the TCP maximum transfer size for reads and writes might be causing a performance issue. If you think the size is slowing performance, you can increase it.

What you'll need

- You must have cluster administrator privileges to perform this task.
- You must use advanced privilege level commands for this task.

Steps

1. Change to the advanced privilege level:

```
set -privilege advanced
```

2. Check the TCP maximum transfer size:

```
vserver nfs show -vserver vserver_name -instance
```

3. If the TCP maximum transfer size is too small, increase the size:

```
vserver nfs modify -vserver vserver_name -tcp-max-xfer-size integer
```

4. Return to the administrative privilege level:

```
set -privilege admin
```

Example

The following example changes the TCP maximum transfer size of SVM1 to 1048576:

```
cluster1::*> vserver nfs modify -vserver SVM1 -tcp-max-xfer-size 1048576
```

Check the iSCSI TCP read/write size

For iSCSI, you can check the TCP read/write size to determine if the size setting is creating a performance issue. If the size is the source of an issue, you can correct it.

What you'll need

Advanced privilege level commands are required for this task.

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Check the TCP window size setting:

```
vserver iscsi show -vserver,er vserver_name -instance
```

3. Modify the TCP window size setting:

```
vserver iscsi modify -vserver vserver_name -tcp-window-size integer
```

4. Return to administrative privilege:

```
set -privilege admin
```

Example

The following example changes the TCP window size of SVM1 to 131,400 bytes:

```
cluster1::*> vserver iscsi modify -vserver vs1 -tcp-window-size 131400
```

Check the CIFS multiplex settings

If slow CIFS network performance causes a performance issue, you can modify the multiplex settings to improve and correct it.

Steps

1. Check the CIFS multiplex setting:

```
vserver cifs options show -vserver -vserver_name -instance
```

2. Modify the CIFS multiplex setting:

```
vserver cifs options modify -vserver -vserver_name -max-mpx integer
```

Example

The following example changes the maximum multiplex count on SVM1 to 255:

```
cluster1:::> vserver cifs options modify -vserver SVM1 -max-mpx 255
```

Check the FC adapter port speed

The adapter target port speed should match the speed of the device to which it connects, to optimize performance. If the port is set to autonegotiation, it can take longer to reconnect after a takeover and giveback or other interruption.

What you'll need

All LIFs that use this adapter as their home port must be offline.

Steps

1. Take the adapter offline:

```
network fcp adapter modify -node nodename -adapter adapter -state down
```

2. Check the maximum speed of the port adapter:

```
fcp adapter show -instance
```

3. Change the port speed, if necessary:

```
network fcp adapter modify -node nodename -adapter adapter -speed  
{1|2|4|8|10|16|auto}
```

4. Bring the adapter online:

```
network fcp adapter modify -node nodename -adapter adapter -state up
```

5. Bring all the LIFs on the adapter online:

```
network interface modify -vserver * -lif * { -home-node node1 -home-port e0c }  
-status-admin up
```

Example

The following example changes the port speed of adapter 0d on node1 to 2 Gbps:

```
cluster1::> network fcp adapter modify -node node1 -adapter 0d -speed 2
```

Check the network settings on the data switches

Although you must maintain the same MTU settings on your clients, servers and storage systems (that is, network endpoints), intermediate network devices such as NICs and switches should be set to their maximum MTU values to ensure that performance is not impacted.

For best performance, all components in the network must be able to forward jumbo frames (9000 bytes IP, 9022 bytes including Ethernet). Data switches should be set to at least 9022 bytes, but a typical value of 9216 is possible with most switches.

Procedure

For data switches, check that the MTU size is set to 9022 or higher.

For more information, see the switch vendor documentation.

Check the MTU network setting on the storage system

You can change the network settings on the storage system if they are not the same as on the client or other network endpoints. Whereas the management network MTU setting is set to 1500, the data network MTU size should be 9000.

About this task

All ports within a broadcast-domain have the same MTU size, with the exception of the e0M port handling management traffic. If the port is part of a broadcast-domain, use the `broadcast-domain modify` command to change the MTU for all ports within the modified broadcast-domain.

Note that intermediate network devices such as NICs and data switches can be set to higher MTU sizes than network endpoints. For more information, see [Check the network settings on the data switches](#).

Steps

1. Check the MTU port setting on the storage system:

```
network port show -instance
```

2. Change the MTU on the broadcast domain used by the ports:

```
network port broadcast-domain modify -ipspace ipspace -broadcast-domain  
broadcast_domain -mtu new_mtu
```

Example

The following example changes the MTU port setting to 9000:

```
network port broadcast-domain modify -ipspace Cluster -broadcast-domain  
Cluster -mtu 9000
```

Check disk throughput and latency

You can check the disk throughput and latency metrics for cluster nodes to assist you in troubleshooting.

About this task

Advanced privilege level commands are required for this task.

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Check the disk throughput and latency metrics:

```
statistics disk show -sort-key latency
```

Example

The following example displays the totals in each user read or write operation for node2 on cluster1:

```

::*> statistics disk show -sort-key latency
cluster1 : 8/24/2015 12:44:15

```

Disk	Node	Busy (%)	Total Ops	Read Ops	Write Ops	Read (Bps)	Write (Bps)	*Latency (us)
1.10.20	node2	4	5	3	2	95232	367616	23806
1.10.8	node2	4	5	3	2	138240	386048	22113
1.10.6	node2	3	4	2	2	48128	371712	19113
1.10.19	node2	4	6	3	2	102400	443392	19106
1.10.11	node2	4	4	2	2	122880	408576	17713

Check throughput and latency between nodes

You can use the `network test-path` command to identify network bottlenecks, or to prequalify network paths between nodes. You can run the command between intercluster nodes or intracluster nodes.

What you'll need

- You must be a cluster administrator to perform this task.
- Advanced privilege level commands are required for this task.
- For an intercluster path, the source and destination clusters must be peered.

About this task

Occasionally, network performance between nodes may not meet expectations for your path configuration. A 1 Gbps transmission rate for the kind of large data transfers seen in SnapMirror replication operations, for example, would not be consistent with a 10 GbE link between the source and destination clusters.

You can use the `network test-path` command to measure throughput and latency between nodes. You can run the command between intercluster nodes or intracluster nodes.



The test saturates the network path with data, so you should run the command when the system is not busy and when network traffic between nodes is not excessive. The test times out after ten seconds. The command can be run only between ONTAP 9 nodes.

The `session-type` option identifies the type of operation you are running over the network path—for example, "AsyncMirrorRemote" for SnapMirror replication to a remote destination. The type dictates the amount of data used in the test. The following table defines the session types:

Session Type	Description
AsyncMirrorLocal	Settings used by SnapMirror between nodes in the same cluster

AsyncMirrorRemote	Settings used by SnapMirror between nodes in different clusters (default type)
RemoteDataTransfer	Settings used by ONTAP for remote data access between nodes in the same cluster (for example, an NFS request to a node for a file stored in a volume on a different node)

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Measure throughput and latency between nodes:

```
network test-path -source-node source_nodename |local -destination-cluster destination_clustername -destination-node destination_nodename -session-type Default|AsyncMirrorLocal|AsyncMirrorRemote|SyncMirrorRemote|RemoteDataTransfer
```

The source node must be in the local cluster. The destination node can be in the local cluster or in a peered cluster. A value of "local" for `-source-node` specifies the node on which you are running the command.

The following command measures throughput and latency for SnapMirror-type replication operations between `node1` on the local cluster and `node3` on `cluster2`:

```
cluster1::> network test-path -source-node node1 -destination-cluster cluster2 -destination-node node3 -session-type AsyncMirrorRemote
Test Duration:          10.88 secs
Send Throughput:       18.23 MB/sec
Receive Throughput:    18.23 MB/sec
MB sent:                198.31
MB received:           198.31
Avg latency in ms:     2301.47
Min latency in ms:     61.14
Max latency in ms:     3056.86
```

3. Return to administrative privilege:

```
set -privilege admin
```

After you finish

If performance does not meet expectations for the path configuration, you should check node performance statistics, use available tools to isolate the problem in the network, check switch settings, and so forth.

Manage workloads

Identify remaining performance capacity

Performance capacity, or *headroom*, measures how much work you can place on a node or an aggregate before performance of workloads on the resource begins to be affected by latency. Knowing the available performance capacity on the cluster helps you provision and balance workloads.

What you'll need

Advanced privilege level commands are required for this task.

About this task

You can use the following values for the `-object` option to collect and display headroom statistics:

- For CPUs, `resource_headroom_cpu`.
- For aggregates, `resource_headroom_aggr`.

You can also complete this task using System Manager and Active IQ Unified Manager.

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Start real-time headroom statistics collection:

```
statistics start -object resource_headroom_cpu|aggr
```

For complete command syntax, see the man page.

3. Display real-time headroom statistics information:

```
statistics show -object resource_headroom_cpu|aggr
```

For complete command syntax, see the man page.

4. Return to administrative privilege:

```
set -privilege admin
```

Example

The following example displays the average hourly headroom statistics for cluster nodes.

You can compute the available performance capacity for a node by subtracting the `current_utilization` counter from the `optimal_point_utilization` counter. In this example, the utilization capacity for `CPU_sti2520-213` is -14% (72%-86%), which suggests that the CPU has been overutilized on average for the past hour.

You could have specified `ewma_daily`, `ewma_weekly`, or `ewma_monthly` to get the same information averaged over longer periods of time.

```
sti2520-2131454963690::*> statistics show -object resource_headroom_cpu
-raw -counter ewma_hourly
(statistics show)
```

```
Object: resource_headroom_cpu
Instance: CPU_sti2520-213
Start-time: 2/9/2016 16:06:27
End-time: 2/9/2016 16:06:27
Scope: sti2520-213
```

Counter	Value
-----	-----
ewma_hourly	-
current_ops	4376
current_latency	37719
current_utilization	86
optimal_point_ops	2573
optimal_point_latency	3589
optimal_point_utilization	72
optimal_point_confidence_factor	1

```
Object: resource_headroom_cpu
Instance: CPU_sti2520-214
Start-time: 2/9/2016 16:06:27
End-time: 2/9/2016 16:06:27
Scope: sti2520-214
```

Counter	Value
-----	-----
ewma_hourly	-
current_ops	0
current_latency	0
current_utilization	0
optimal_point_ops	0
optimal_point_latency	0
optimal_point_utilization	71
optimal_point_confidence_factor	1

```
2 entries were displayed.
```

Identify high-traffic clients or files

You can use ONTAP Active Objects technology to identify clients or files that are responsible for a disproportionately large amount of cluster traffic. Once you have identified these "top" clients or files, you can rebalance cluster workloads or take other steps to resolve the issue.

What you'll need

You must be a cluster administrator to perform this task.

Steps

1. View the top clients accessing the cluster:

```
statistics top client show -node node_name -sort-key sort_column -interval  
seconds_between_updates -iterations iterations -max number_of_instances
```

For complete command syntax, see the man page.

The following command displays the top clients accessing cluster1:

```
cluster1::> statistics top client show  
  
cluster1 : 3/23/2016 17:59:10  
  
                *Total  
      Client Vserver          Node Protocol    Ops  
-----
```

172.17.180.170	vs4	siderop1-vs4	nfs	668
172.17.180.169	vs3	siderop1-vs3	nfs	337
172.17.180.171	vs3	siderop1-vs3	nfs	142
172.17.180.170	vs3	siderop1-vs3	nfs	137
172.17.180.123	vs3	siderop1-vs3	nfs	137
172.17.180.171	vs4	siderop1-vs4	nfs	95
172.17.180.169	vs4	siderop1-vs4	nfs	92
172.17.180.123	vs4	siderop1-vs4	nfs	92
172.17.180.153	vs3	siderop1-vs3	nfs	0

2. View the top files accessed on the cluster:

```
statistics top file show -node node_name -sort-key sort_column -interval  
seconds_between_updates -iterations iterations -max number_of_instances
```

For complete command syntax, see the man page.

The following command displays the top files accessed on cluster1:

```
cluster1::> statistics top file show
```

```
cluster1 : 3/23/2016 17:59:10
```

```

                                *Total
      File Volume Vserver      Node      Ops
-----
/vol/vol1/vm170-read.dat  vol1      vs4 siderop1-vs4      22
/vol/vol1/vm69-write.dat  vol1      vs3 siderop1-vs3       6
  /vol/vol2/vm171.dat     vol2      vs3 siderop1-vs3       2
  /vol/vol2/vm169.dat     vol2      vs3 siderop1-vs3       2
  /vol/vol2/p123.dat      vol2      vs4 siderop1-vs4       2
  /vol/vol2/p123.dat      vol2      vs3 siderop1-vs3       2
/vol/vol1/vm171.dat      vol1      vs4 siderop1-vs4       2
/vol/vol1/vm169.dat      vol1      vs4 siderop1-vs4       2
/vol/vol1/vm169.dat      vol1      vs4 siderop1-vs3       2
  /vol/vol1/p123.dat      vol1      vs4 siderop1-vs4       2
```

Guarantee throughput with QoS

Guarantee throughput with QoS overview

You can use storage quality of service (QoS) to guarantee that performance of critical workloads is not degraded by competing workloads. You can set a throughput *ceiling* on a competing workload to limit its impact on system resources, or set a throughput *floor* for a critical workload, ensuring that it meets minimum throughput targets, regardless of demand by competing workloads. You can even set a ceiling and floor for the same workload.


About throughput ceilings (QoS Max)

A throughput ceiling limits throughput for a workload to a maximum number of IOPS or MBps, or IOPS and MBps. In the figure below, the throughput ceiling for workload 2 ensures that it does not "bully" workloads 1 and 3.

A *policy group* defines the throughput ceiling for one or more workloads. A workload represents the I/O operations for a *storage object*: a volume, file, qtree or LUN, or all the volumes, files, qtrees, or LUNs in an SVM. You can specify the ceiling when you create the policy group, or you can wait until after you monitor workloads to specify it.



Throughput to workloads might exceed the specified ceiling by up to 10%, especially if a workload experiences rapid changes in throughput. The ceiling might be exceeded by up to 50% to handle bursts. Bursts occur on single nodes when tokens accumulate up to 150%



About throughput floors (QoS Min)

A throughput floor guarantees that throughput for a workload does not fall below a minimum number of IOPS or MBps, or IOPS and MBps. In the figure below, the throughput floors for workload 1 and workload 3 ensure that they meet minimum throughput targets, regardless of demand by workload 2.



As the examples suggest, a throughput ceiling throttles throughput directly. A throughput floor throttles throughput indirectly, by giving priority to the workloads for which the floor has been set.

You can specify the floor when you create the policy group, or you can wait until after you monitor workloads to specify it.


Beginning in ONTAP 9.13.1, you can set throughput floors at the SVM scope with [Adaptive policy group templates](#). In releases of ONTAP before 9.13.1, a policy group that defines a throughput floor cannot be applied to an SVM.



In releases before ONTAP 9.7, throughput floors are guaranteed when there is sufficient performance capacity available.

In ONTAP 9.7 and later, throughput floors can be guaranteed even when there is insufficient performance capacity available. This new floor behavior is called floors v2. To meet the guarantees, floors v2 can result in higher latency on workloads without a throughput floor or on work that exceeds the floor settings. Floors v2 applies to both QoS and adaptive QoS.

The option of enabling/disabling the new behavior of floors v2 is available in ONTAP 9.7P6 and later. A workload might fall below the specified floor during critical operations like `volume move trigger-cutover`. Even when sufficient capacity is available and critical operations are not taking place, throughput to a workload might fall below the specified floor by up to 5%. If floors are overprovisioned and there is no performance capacity, some workloads might fall below the specified floor.



About shared and non-shared QoS policy groups

Beginning with ONTAP 9.4, you can use a *non-shared* QoS policy group to specify that the defined throughput ceiling or floor applies to each member workload individually. Behavior of *shared* policy groups depends on the policy type:

- For throughput ceilings, the total throughput for the workloads assigned to the shared policy group cannot exceed the specified ceiling.
- For throughput floors, the shared policy group can be applied to a single workload only.

About adaptive QoS

Ordinarily, the value of the policy group you assign to a storage object is fixed. You need to change the value manually when the size of the storage object changes. An increase in the amount of space used on a volume, for example, usually requires a corresponding increase in the throughput ceiling specified for the volume.

Adaptive QoS automatically scales the policy group value to workload size, maintaining the ratio of IOPS to TB|GBs as the size of the workload changes. That is a significant advantage when you are managing hundreds or thousands of workloads in a large deployment.

You typically use adaptive QoS to adjust throughput ceilings, but you can also use it to manage throughput floors (when workload size increases). Workload size is expressed as either the allocated space for the storage object or the space used by the storage object.



Used space is available for throughput floors in ONTAP 9.5 and later. It is not supported for throughput floors in ONTAP 9.4 and earlier.

- An *allocated space* policy maintains the IOPS/TB|GB ratio according to the nominal size of the storage object. If the ratio is 100 IOPS/GB, a 150 GB volume will have a throughput ceiling of 15,000 IOPS for as long as the volume remains that size. If the volume is resized to 300 GB, adaptive QoS adjusts the throughput ceiling to 30,000 IOPS.
- A *used space* policy (the default) maintains the IOPS/TB|GB ratio according to the amount of actual data stored before storage efficiencies. If the ratio is 100 IOPS/GB, a 150 GB volume that has 100 GB of data stored would have a throughput ceiling of 10,000 IOPS. As the amount of used space changes, adaptive QoS adjusts the throughput ceiling according to the ratio.

Beginning with ONTAP 9.5, you can specify an I/O block size for your application that enables a throughput limit to be expressed in both IOPS and MBps. The MBps limit is calculated from the block size multiplied by the

IOPS limit. For example, an I/O block size of 32K for an IOPS limit of 6144IOPS/TB yields an MBps limit of 192MBps.

You can expect the following behavior for both throughput ceilings and floors:

- When a workload is assigned to an adaptive QoS policy group, the ceiling or floor is updated immediately.
- When a workload in an adaptive QoS policy group is resized, the ceiling or floor is updated in approximately five minutes.

Throughput must increase by at least 10 IOPS before updates take place.

Adaptive QoS policy groups are always non-shared: the defined throughput ceiling or floor applies to each member workload individually.

Beginning with ONTAP 9.6, throughput floors are supported on ONTAP Select premium with SSD.

Adaptive policy group template

Beginning in ONTAP 9.13.1, you can set an adaptive QoS template on an SVM. Adaptive policy group templates enable you to set throughput floors and ceilings for all volumes in an SVM.

Adaptive policy group templates can only be set after the SVM has been created. Use the `vserver modify` command with the `-qos-adaptive-policy-group-template` parameter to set the policy.

When you set an an adaptive policy group template, volumes created or migrated after setting the policy automatically inherit the policy. Any volumes existing on the SVM are not impacted when you assign the policy template. If you disable the policy on the SVM, any volume subsequently migrated to or created on the SVM will not receive the policy. Disabling the adaptive policy group template does not impact volumes that inherited the policy template as they retain the policy template.

For more information, see [Set an adaptive policy group template](#).

General support

The following table shows the differences in support for throughput ceilings, throughput floors, and adaptive QoS.

Resource or feature	Throughput ceiling	Throughput floor	Throughput floor v2	Adaptive QoS
ONTAP 9 version	All	9.2 and later	9.7 and later	9.3 and later
Platforms	All	<ul style="list-style-type: none">• AFF• C190 *• ONTAP Select premium with SSD *	<ul style="list-style-type: none">• AFF• C190• ONTAP Select premium with SSD	All
Protocols	All	All	All	All

Resource or feature	Throughput ceiling	Throughput floor	Throughput floor v2	Adaptive QoS
FabricPool	Yes	Yes, if the tiering policy is set to "none" and no blocks are in the cloud.	Yes, if the tiering policy is set to "none" and no blocks are in the cloud.	No
SnapMirror Synchronous	Yes	No	No	Yes

C190 and ONTAP Select support started with the ONTAP 9.6 release.

Supported workloads for throughput ceilings

The following table shows workload support for throughput ceilings by ONTAP 9 version. Root volumes, load-sharing mirrors, and data protection mirrors are not supported.

Workload support - ceiling	ONTAP 9.0	ONTAP 9.1	ONTAP 9.2	ONTAP 9.3	ONTAP 9.4 - 9.7	ONTAP 9.8 and later
Volume	yes	yes	yes	yes	yes	yes
File	yes	yes	yes	yes	yes	yes
LUN	yes	yes	yes	yes	yes	yes
SVM	yes	yes	yes	yes	yes	yes
FlexGroup volume	no	no	no	yes	yes	yes
qtrees*	no	no	no	no	no	yes
Multiple workloads per policy group	yes	yes	yes	yes	yes	yes
Non-shared policy groups	no	no	no	no	yes	yes

Beginning with ONTAP 9.8, NFS access is supported in qtrees in FlexVol and FlexGroup volumes with NFS enabled. Beginning with ONTAP 9.9.1, SMB access is also supported in qtrees in FlexVol and FlexGroup volumes with SMB enabled.

Supported workloads for throughput floors

The following table shows workload support for throughput floors by ONTAP 9 version. Root volumes, load-sharing mirrors, and data protection mirrors are not supported.

Workload support - floor	ONTAP 9.2	ONTAP 9.3	ONTAP 9.4 - 9.7	ONTAP 9.8 - 9.13.0	ONTAP 9.13.1 and later
Volume	yes	yes	yes	yes	yes
File	no	yes	yes	yes	yes
LUN	yes	yes	yes	yes	yes
SVM	no	no	no	no	yes
FlexGroup volume	no	no	yes	yes	yes
qtrees *	no	no	no	yes	yes
Multiple workloads per policy group	no	no	yes	yes	yes
Non-shared policy groups	no	no	yes	yes	yes

*Beginning with ONTAP 9.8, NFS access is supported in qtrees in FlexVol and FlexGroup volumes with NFS enabled. Beginning with ONTAP 9.9.1, SMB access is also supported in qtrees in FlexVol and FlexGroup volumes with SMB enabled.

Supported workloads for adaptive QoS

The following table shows workload support for adaptive QoS by ONTAP 9 version. Root volumes, load-sharing mirrors, and data protection mirrors are not supported.

Workload support - adaptive QoS	ONTAP 9.3	ONTAP 9.4 - 9.13.0	ONTAP 9.13.1 and later
Volume	yes	yes	yes
File	no	yes	yes
LUN	no	yes	yes
SVM	no	no	yes
FlexGroup volume	no	yes	yes
Multiple workloads per policy group	yes	yes	yes
Non-shared policy groups	yes	yes	yes

Maximum number of workloads and policy groups

The following table shows the maximum number of workloads and policy groups by ONTAP 9 version.

Workload support	ONTAP 9.3 and earlier	ONTAP 9.4 and later
Maximum workloads per cluster	12,000	40,000
Maximum workloads per node	12,000	40,000
Maximum policy groups	12,000	12,000

Enable or disable throughput floors v2

You can enable or disable throughput floors v2 on AFF. The default is enabled. With floors v2 enabled, throughput floors can be met when controllers are heavily used at the expense of higher latency on other workloads. Floors v2 applies to both QoS and Adaptive QoS.

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Enter one of the following commands:

If you want to...	Use this command:
Disable floors v2	<code>qos settings throughput-floors-v2 -enable false</code>
Enable floors v2	<code>qos settings throughput-floors-v2 -enable true</code>



To disable throughput floors v2 in an MetroCluster cluster, you must run the

```
qos settings throughput-floors-v2 -enable false
```

command on both the source and destination clusters.

```
cluster1::*> qos settings throughput-floors-v2 -enable false
```

Storage QoS workflow

If you already know the performance requirements for the workloads you want to manage with QoS, you can specify the throughput limit when you create the policy group. Otherwise, you can wait until after you monitor the workloads to specify the limit.

Set a throughput ceiling with QoS

You can use the `max-throughput` field for a policy group to define a throughput ceiling for storage object workloads (QoS Max). You can apply the policy group when you create or modify the storage object.

What you'll need

- You must be a cluster administrator to create a policy group.
- You must be a cluster administrator to apply a policy group to an SVM.

About this task

- Beginning with ONTAP 9.4, you can use a *non-shared* QoS policy group to specify that the defined throughput ceiling applies to each member workload individually. Otherwise, the policy group is *shared*: the total throughput for the workloads assigned to the policy group cannot exceed the specified ceiling.

Set `-is-shared=false` for the `qos policy-group create` command to specify a non-shared policygroup.

- You can specify the throughput limit for the ceiling in IOPS, MB/s, or IOPS, MB/s. If you specify both IOPS and MB/s, whichever limit is reached first is enforced.



If you set a ceiling and a floor for the same workload, you can specify the throughput limit for the ceiling in IOPS only.

- A storage object that is subject to a QoS limit must be contained by the SVM to which the policy group belongs. Multiple policy groups can belong to the same SVM.
- You cannot assign a storage object to a policy group if its containing object or its child objects belong to the policy group.
- It is a QoS best practice to apply a policy group to the same type of storage objects.

Steps

1. Create a policy group:

```
qos policy-group create -policy-group policy_group -vserver SVM -max  
-throughput number_of_iops|Mb/S|iops,Mb/S -is-shared true|false
```

For complete command syntax, see the man page. You can use the `qos policy-group modify` command to adjust throughput ceilings.

The following command creates the shared policy group `pg-vs1` with a maximum throughput of 5,000 IOPS:

```
cluster1::> qos policy-group create -policy-group pg-vs1 -vserver vs1  
-max-throughput 5000iops -is-shared true
```

The following command creates the non-shared policy group `pg-vs3` with a maximum throughput of 100 IOPS and 400 Kb/S:

```
cluster1::> qos policy-group create -policy-group pg-vs3 -vserver vs3
-max-throughput 100iops,400KB/s -is-shared false
```

The following command creates the non-shared policy group `pg-vs4` without a throughput limit:

```
cluster1::> qos policy-group create -policy-group pg-vs4 -vserver vs4
-is-shared false
```

2. Apply a policy group to an SVM, file, volume, or LUN:

```
storage_object create -vserver SVM -qos-policy-group policy_group
```

For complete command syntax, see the man pages. You can use the `storage_object modify` command to apply a different policy group to the storage object.

The following command applies policy group `pg-vs1` to SVM `vs1`:

```
cluster1::> vserver create -vserver vs1 -qos-policy-group pg-vs1
```

The following commands apply policy group `pg-app` to the volumes `app1` and `app2`:

```
cluster1::> volume create -vserver vs2 -volume app1 -aggregate aggr1
-qos-policy-group pg-app
```

```
cluster1::> volume create -vserver vs2 -volume app2 -aggregate aggr1
-qos-policy-group pg-app
```

3. Monitor policy group performance:

```
qos statistics performance show
```

For complete command syntax, see the man page.



Monitor performance from the cluster. Do not use a tool on the host to monitor performance.

The following command shows policy group performance:

```
cluster1::> qos statistics performance show
Policy Group          IOPS      Throughput  Latency
-----
-total-              12316      47.76MB/s  1264.00us
pg_vs1                5008      19.56MB/s   2.45ms
_System-Best-Effort    62        13.36KB/s   4.13ms
_System-Background    30         0KB/s       0ms
```

4. Monitor workload performance:

```
qos statistics workload performance show
```

For complete command syntax, see the man page.



Monitor performance from the cluster. Do not use a tool on the host to monitor performance.

The following command shows workload performance:

```
cluster1::> qos statistics workload performance show
Workload             ID      IOPS      Throughput  Latency
-----
-total-              -      12320      47.84MB/s  1215.00us
app1-wid7967         7967    7219      28.20MB/s  319.00us
vs1-wid12279         12279   5026      19.63MB/s   2.52ms
_USERSPACE_APPS      14      55        10.92KB/s  236.00us
_Scan_Backgro...     5688    20         0KB/s       0ms
```



You can use the `qos statistics workload latency show` command to view detailed latency statistics for QoS workloads.

Set a throughput floor with QoS

You can use the `min-throughput` field for a policy group to define a throughput floor for storage object workloads (QoS Min). You can apply the policy group when you create or modify the storage object. Beginning with ONTAP 9.8, you can specify the throughput floor in IOPS or MBps, or IOPS and MBps.

Before you begin

- You must be running ONTAP 9.2 or later. Throughput floors are available beginning with ONTAP 9.2.
- You must be a cluster administrator to create a policy group.
- Beginning in ONTAP 9.13.1, you can enforce throughput floors at the SVM level using an [adaptive policy group template](#). You cannot set an adaptive policy group template on an SVM with a QoS policy group.

About this task

- Beginning with ONTAP 9.4, you can use a *non-shared* QoS policy group to specify that the defined throughput floor be applied to each member workload individually. This is the only condition in which a policy group for a throughput floor can be applied to multiple workloads.

Set `-is-shared=false` for the `qos policy-group create` command to specify a non-shared policy group.

- Throughput to a workload might fall below the specified floor if there is insufficient performance capacity (headroom) on the node or aggregate.
- A storage object that is subject to a QoS limit must be contained by the SVM to which the policy group belongs. Multiple policy groups can belong to the same SVM.
- It is a QoS best practice to apply a policy group to the same type of storage objects.
- A policy group that defines a throughput floor cannot be applied to an SVM.

Steps

1. Check for adequate performance capacity on the node or aggregate, as described in [Identifying remaining performance capacity](#).
2. Create a policy group:

```
qos policy-group create -policy group policy_group -vserver SVM -min
-throughput qos_target -is-shared true|false
```

For complete command syntax, see the man page for your ONTAP release. You can use the `qos policy-group modify` command to adjust throughput floors.

The following command creates the shared policy group `pg-vs2` with a minimum throughput of 1,000 IOPS:

```
cluster1::> qos policy-group create -policy group pg-vs2 -vserver vs2
-min-throughput 1000iops -is-shared true
```

The following command creates the non-shared policy group `pg-vs4` without a throughput limit:

```
cluster1::> qos policy-group create -policy group pg-vs4 -vserver vs4
-is-shared false
```

3. Apply a policy group to a volume or LUN:

```
storage_object create -vserver SVM -qos-policy-group policy_group
```

For complete command syntax, see the man pages. You can use the `_storage_object_modify` command to apply a different policy group to the storage object.

The following command applies policy group `pg-app2` to the volume `app2`:

```
cluster1::> volume create -vserver vs2 -volume app2 -aggregate aggr1
-qos-policy-group pg-app2
```

4. Monitor policy group performance:

```
qos statistics performance show
```

For complete command syntax, see the man page.



Monitor performance from the cluster. Do not use a tool on the host to monitor performance.

The following command shows policy group performance:

```
cluster1::> qos statistics performance show
Policy Group           IOPS           Throughput     Latency
-----
-total-                12316          47.76MB/s     1264.00us
pg_app2                7216           28.19MB/s     420.00us
_System-Best-Effort    62             13.36KB/s     4.13ms
_System-Background    30             0KB/s         0ms
```

5. Monitor workload performance:

```
qos statistics workload performance show
```

For complete command syntax, see the man page.



Monitor performance from the cluster. Do not use a tool on the host to monitor performance.

The following command shows workload performance:

```
cluster1::> qos statistics workload performance show
Workload              ID           IOPS           Throughput     Latency
-----
-total-               -            12320          47.84MB/s     1215.00us
app2-wid7967          7967         7219           28.20MB/s     319.00us
vs1-wid12279          12279        5026           19.63MB/s     2.52ms
_USERSPACE_APPS       14           55             10.92KB/s     236.00us
_Scan_Backgro...     5688         20             0KB/s         0ms
```



You can use the `qos statistics workload latency show` command to view detailed latency statistics for QoS workloads.

Use adaptive QoS policy groups

You can use an *adaptive QoS* policy group to automatically scale a throughput ceiling or floor to volume size, maintaining the ratio of IOPS to TBs|GBs as the size of the volume changes. That is a significant advantage when you are managing hundreds or thousands of workloads in a large deployment.

Before you begin

- You must be running ONTAP 9.3 or later. Adaptive QoS policy groups are available beginning with ONTAP 9.3.
- You must be a cluster administrator to create a policy group.

About this task

A storage object can be a member of an adaptive policy group or a non-adaptive policy group, but not both. The SVM of the storage object and the policy must be the same. The storage object must be online.

Adaptive QoS policy groups are always non-shared: the defined throughput ceiling or floor applies to each member workload individually.

The ratio of throughput limits to storage object size is determined by the interaction of the following fields:

- `expected-iops` is the minimum expected IOPS per allocated TB|GB.



`expected-iops` is guaranteed on AFF platforms only. `expected-iops` is guaranteed for FabricPool only if the tiering policy is set to "none" and no blocks are in the cloud. `expected-iops` is guaranteed for volumes that are not in a SnapMirror Synchronous relationship.

- `peak-iops` is the maximum possible IOPS per allocated or used TB|GB.
- `expected-iops-allocation` specifies whether allocated space (the default) or used space is used for `expected-iops`.



`expected-iops-allocation` is available in ONTAP 9.5 and later. It is not supported in ONTAP 9.4 and earlier.

- `peak-iops-allocation` specifies whether allocated space or used space (the default) is used for `peak-iops`.
- `absolute-min-iops` is the absolute minimum number of IOPS. You can use this field with very small storage objects. It overrides both `peak-iops` and/or `expected-iops` when `absolute-min-iops` is greater than the calculated `expected-iops`.

For example, if you set `expected-iops` to 1,000 IOPS/TB, and the volume size is less than 1 GB, the calculated `expected-iops` will be a fractional IOP. The calculated `peak-iops` will be an even smaller fraction. You can avoid this by setting `absolute-min-iops` to a realistic value.

- `block-size` specifies the application I/O block size. The default is 32K. Valid values are 8K, 16K, 32K, 64K, ANY. ANY means that the block size is not enforced.

Three default adaptive QoS policy groups are available, as shown in the following table. You can apply these policy groups directly to a volume.

Default policy group	Expected IOPS/TB	Peak IOPS/TB	Absolute Min IOPS
extreme	6,144	12,288	1000
performance	2,048	4,096	500
value	128	512	75

You cannot assign a storage object to a policy group if its containing object or its child objects belong to a policy group. The following table lists the restrictions.

If you assign the...	Then you cannot assign...
SVM to a policy group	Any storage objects contained by the SVM to a policy group
Volume to a policy group	The volume's containing SVM or any child LUNs to a policy group
LUN to a policy group	The LUN's containing volume or SVM to a policy group
File to a policy group	The file's containing volume or SVM to a policy group

Steps

1. Create an adaptive QoS policy group:

```
qos adaptive-policy-group create -policy group policy_group -vserver SVM
-expected-iops number_of_iops/TB|GB -peak-iops number_of_iops/TB|GB -expected
-iops-allocation-space|used-space -peak-iops-allocation allocated-space|used-
space -absolute-min-iops number_of_iops -block-size 8K|16K|32K|64K|ANY
```

For complete command syntax, see the man page.



`-expected-iops-allocation` and `-block-size` is available in ONTAP 9.5 and later. These options are not supported in ONTAP 9.4 and earlier.

The following command creates adaptive QoS policy group `adpg-app1` with `-expected-iops` set to 300 IOPS/TB, `-peak-iops` set to 1,000 IOPS/TB, `-peak-iops-allocation` set to `used-space`, and `-absolute-min-iops` set to 50 IOPS:

```
cluster1::> qos adaptive-policy-group create -policy group adpg-app1
-vserver vs2 -expected-iops 300iops/tb -peak-iops 1000iops/TB -peak-iops
-allocation used-space -absolute-min-iops 50iops
```

2. Apply an adaptive QoS policy group to a volume:

```
volume create -vserver SVM -volume volume -aggregate aggregate -size number_of
TB|GB -qos-adaptive-policy-group policy_group
```

For complete command syntax, see the man pages.

The following command applies adaptive QoS policy group `adpg-app1` to volume `app1`:

```
cluster1::> volume create -vserver vs1 -volume app1 -aggregate aggr1
-size 2TB -qos-adaptive-policy-group adpg-app1
```

The following commands apply the default adaptive QoS policy group `extreme` to the new volume `app4` and to the existing volume `app5`. The throughput ceiling defined for the policy group applies to volumes `app4` and `app5` individually:

```
cluster1::> volume create -vserver vs4 -volume app4 -aggregate aggr4
-size 2TB -qos-adaptive-policy-group extreme
```

```
cluster1::> volume modify -vserver vs5 -volume app5 -qos-adaptive-policy
-group extreme
```

Set an adaptive policy group template

Beginning in ONTAP 9.13.1, you can enforce throughput floors and ceilings at the SVM level using an adaptive policy group template.

About this task

- The adaptive policy group template is a default policy `apg1`. The policy can be modified at any time. It can only be set with the CLI or ONTAP REST API and can only be applied to existing SVMs.
- The adaptive policy group template only impacts volumes created on or migrated to the SVM after you set the policy. Existing volumes on the SVM retain their existing status.

If you disable the adaptive policy group template, volumes on the SVM retain their existing policies. Only volumes subsequently created on or migrated to the SVM will be impacted by the disablement.

- You cannot set an adaptive policy group template on an SVM with a QoS policy group.
- Adaptive policy group templates are designed for AFF platforms. An adaptive policy group template can be set on other platforms, but the policy may not enforce a minimum throughput. Similarly, you can add an adaptive policy group template to an SVM in a FabricPool aggregate or in an aggregate that does not support a minimum throughput, however the throughput floor will not be enforced.
- If the SVM is in a MetroCluster configuration or an SnapMirror relationship, the adaptive policy group template will be enforced on the mirrored SVM.

Steps

1. Modify the SVM to apply the adaptive policy group template:

```
vserver modify -qos-adaptive-policy-group-template apg1
```


2. Confirm the policy was set:

```
vserver show -fields qos-adaptive-policy-group
```

Monitor cluster performance with Unified Manager

With Active IQ Unified Manager, you can maximize availability and maintain control of your NetApp AFF and FAS storage infrastructure for improved scalability, supportability, performance, and security.

Active IQ Unified Manager continuously monitors system health and send alerts, so your organization can free up IT staff resources. You can instantly view storage status from a single dashboard and quickly address issues through recommended actions.

Data management is simplified because you can discover, monitor, and receive notifications to proactively manage storage and quickly resolve issues. Admin efficiency is improved because you can monitor petabytes of data from a single dashboard and manage your data at scale.

With Active IQ Unified Manager, you can keep pace with fluctuating business demands, optimizing performance using performance data and advanced analytics. The reporting capabilities allow you to access standard reports or create custom operational reports to meet the specific needs of your business.

Related links:

- [Learn more about Active IQ Unified Manager](#)
- [Get started with Active IQ Unified Manager for VMware](#)
- [Get started with Active IQ Unified Manager for Linux](#)
- [Get started with Active IQ Unified Manager for Windows](#)

Monitor cluster performance with Cloud Insights

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

Cloud Insights comes in two editions

Cloud Insights Basic Edition is designed specifically to monitor and optimize your NetApp Data Fabric assets. It provides advanced analytics for the connections between all NetApp resources including HCI and All Flash FAS (AFF) within the environment free of charge.

Cloud Insights Standard Edition focuses not only on NetApp Data Fabric-enabled infrastructure components, but also on multi-vendor and multi-cloud environments. With its enriched capabilities, you can access support for over 100 services and resources.

In today's world, with resources in play from your on-premises data centers to multiple public clouds, it's crucial to have the complete picture from the application itself to the backend disk of the storage array. The additional support for application monitoring (like Kafka, MongoDB, and Nginx) gives you the information and knowledge you need to operate at the optimal level of utilization as well as with the perfect risk buffer.

Both editions (Basic and Standard) can integrate with NetApp Active IQ Unified Manager. Customers who use

Active IQ Unified Manager can see join information inside the Cloud Insights user interface. Notifications posted on Active IQ Unified Manager are not overlooked and can be correlated to events in Cloud Insights. In other words, you get the best of both worlds.

Monitor, troubleshoot, and optimize all your resources

Cloud Insights helps you significantly reduce the time to resolve issues and prevent them from impacting end users. It also helps you reduce cloud infrastructure costs. Your exposure to insider threats is reduced by protecting your data with actionable intelligence.

Cloud Insights gives you visibility to your entire hybrid infrastructure in one place—from the public cloud to your data center. You can instantly create relevant dashboards that can be customized to your specific needs. You can also create targeted and conditional alerts that are specific and relevant to your organization's needs.

Advanced anomaly detection helps you proactively fix issues before they arise. You can view resource contention and degradation automatically to quickly restore impacted workloads. Troubleshooting goes more quickly with the automatically built hierarchy of relationships between the different components in your stack.

You can identify unused or abandoned resources across your environment, which helps you discover opportunities to right-size the infrastructure and optimize your entire spend.

Cloud Insights visualizes your system topology to gain an understanding of your Kubernetes architecture. You can monitor the health of your Kubernetes clusters, including which nodes are in trouble, and zoom in when you see a problem.

Cloud Insights helps you protect organizational data from being misused by malicious or compromised users through advanced machine learning and anomaly detection that gives you actionable intelligence on insider threats.

Cloud Insights helps you to visualize Kubernetes metrics so you can fully understand the relations between your pods, nodes, and clusters. You're able to assess the health of a cluster or a working pod, as well as the load it is currently processing—enabling you to take command of your K8S cluster and to control both the health and the cost of your deployment.

Related links

- [Learn more about Cloud Insights](#)
- [Get started with Cloud Insights](#)

Audit logging

How ONTAP implements audit logging

Management activities recorded in the audit log are included in standard AutoSupport reports, and certain logging activities are included in EMS messages. You can also forward the audit log to destinations that you specify, and you can display audit log files by using the CLI or a web browser.

Beginning with ONTAP 9.11.1, you can display audit log contents using System Manager.

Beginning with ONTAP 9.12.1, ONTAP provides tampering alerts for audit logs. ONTAP runs a daily background job to check for tampering of audit.log files and sends an EMS alert if it finds any log files that have been changed or tampered with.

ONTAP logs management activities that are performed on the cluster, for example, what request was issued, the user who triggered the request, the user's access method, and the time of the request.

The management activities can be one of the following types:

- SET requests, which typically apply to non-display commands or operations
 - These requests are issued when you run a `create`, `modify`, or `delete` command, for instance.
 - Set requests are logged by default.
- GET requests, which retrieve information and display it in the management interface
 - These requests are issued when you run a `show` command, for instance.
 - GET requests are not logged by default, but you can control whether GET requests sent from the ONTAP CLI (`-cliget`), from the ONTAP API (`-ontapget`), or from the REST API (`-httpget`) are logged in the file.

ONTAP records management activities in the `/mroot/etc/log/mlog/audit.log` file of a node.

Commands from the three shells for CLI commands—the clustershell, the nodeshell, and the non-interactive systemshell (interactive systemshell commands are not logged)—as well as API commands are logged here. Audit logs include timestamps to show whether all nodes in a cluster are time synchronized.

The `audit.log` file is sent by the AutoSupport tool to the specified recipients. You can also forward the content securely to external destinations that you specify; for example, a Splunk or a syslog server.

The `audit.log` file is rotated daily. The rotation also occurs when it reaches 100 MB in size, and the previous 48 copies are preserved (with a maximum total of 49 files). When the audit file performs its daily rotation, no EMS message is generated. If the audit file rotates because its file size limit is exceeded, an EMS message is generated.

Changes to audit logging in ONTAP 9

Beginning with ONTAP 9, the `command-history.log` file is replaced by `audit.log`, and the `mgwd.log` file no longer contains audit information. If you are upgrading to ONTAP 9, you should review any scripts or tools that refer to the legacy files and their contents.

After upgrade to ONTAP 9, existing `command-history.log` files are preserved. They are rotated out (deleted) as new `audit.log` files are rotated in (created).

Tools and scripts that check the `command-history.log` file might continue to work, because a soft link from `command-history.log` to `audit.log` is created at upgrade. However, tools and scripts that check the `mgwd.log` file will fail, because that file no longer contains audit information.

In addition, audit logs in ONTAP 9 and later no longer include the following entries because they are not considered useful and cause unnecessary logging activity:

- Internal commands run by ONTAP (that is, where `username=root`)
- Command aliases (separately from the command they point to)

Beginning with ONTAP 9, you can transmit the audit logs securely to external destinations using the TCP and TLS protocols.

Display audit log contents

You can display the contents of the cluster's `/mroot/etc/log/mlog/audit.log` files by using the ONTAP CLI, System Manager, or a web browser.

The cluster's log file entries include the following:

Time

The log entry timestamp.

Application

The application used to connect to the cluster. Examples of possible values are `internal`, `console`, `ssh`, `http`, `ontapi`, `snmp`, `rsh`, `telnet`, and `service-processor`.

User

The username of the remote user.

State

The current state of the audit request, which could be `success`, `pending`, or `error`.

Message

An optional field that might contain error or additional information about the status of a command.

Session ID

The session ID on which the request is received. Each SSH *session* is assigned a session ID, while each HTTP, ONTAPI, or SNMP *request* is assigned a unique session ID.

Storage VM

The SVM through which the user connected.

Scope

Displays `svm` when the request is on a data storage VM; otherwise displays `cluster`.

Command ID

The ID for each command received on a CLI session. This enables you to correlate a request and response. ZAPI, HTTP, and SNMP requests do not have command IDs.

You can display the cluster's log entries from the ONTAP CLI, from a web browser, and beginning with ONTAP 9.11.1, from System Manager.

System Manager

- To display the inventory, select **Events & Jobs > Audit Logs**. Each column has controls to filter, sort, search, show, and inventory categories. The inventory details can be downloaded as an Excel workbook.
- To set filters, click the **Filter** button on the upper right side, then select the desired fields. You can also view all the commands executed in the session in which a failure occurred by clicking on the Session ID link.

CLI

To display audit entries merged from multiple nodes in the cluster, enter:

```
security audit log show [parameters]
```

You can use the `security audit log show` command to display audit entries for individual nodes or merged from multiple nodes in the cluster. You can also display the content of the `/mroot/etc/log/mlog` directory on a single node by using a web browser. See the man page for details.

Web browser

You can display the content of the `/mroot/etc/log/mlog` directory on a single node by using a web browser. [Learn about how to access a node's log, core dump, and MIB files by using a web browser.](#)

Manage audit GET request settings

While SET requests are logged by default, GET requests are not. However, you can control whether GET requests sent from ONTAP HTML (`-httpget`), the ONTAP CLI (`-cliget`), or from the ONTAP APIs (`-ontapiget`) are logged in the file.

You can modify audit logging settings from the ONTAP CLI, and beginning with ONTAP 9.11.1, from System Manager.

System Manager

1. Select **Events & Jobs > Audit Logs**.
2. Click  in the upper-right corner, then choose the requests to add or remove.

CLI

- To specify that GET requests from the ONTAP CLI or APIs should be recorded in the audit log (the `audit.log` file), in addition to default set requests, enter:

```
security audit modify [-cliget {on|off}][-httpget {on|off}][-ontapiget {on|off}]
```
- To display the current settings, enter:

```
security audit show
```

See the man pages for details.

Manage audit log destinations

You can forward the audit log to a maximum of 10 destinations. For example, you can forward the log to a Splunk or syslog server for monitoring, analysis, or backup purposes.

About this task

To configure forwarding, you must provide the IP address of the syslog or Splunk host, its port number, a transmission protocol, and the syslog facility to use for the forwarded logs. [Learn about syslog facilities.](#)

You can select one of the following transmission values:

UDP Unencrypted

User Datagram Protocol with no security (default)

TCP Unencrypted

Transmission Control Protocol with no security




TCP Encrypted

Transmission Control Protocol with Transport Layer Security (TLS)

A **Verify server** option is available when the TCP Encrypted protocol is selected.

You can forward audit logs from the ONTAP CLI, and beginning with ONTAP 9.11.1, from System Manager.

System Manager

- To display audit log destinations, select **Cluster >Settings**. A count of log destinations is shown in the **Notification Management** tile. Click  to show details.
- To add, modify, or delete audit log destinations, select **Events & Jobs > Audit Logs**, then click **Manage Audit Destinations** in the upper right of the screen. Click  **Add**, or click  in the **Host Address** column to edit or delete entries.

CLI

1. For each destination that you want to forward the audit log to, specify the destination IP address or host name and any security options.

```
cluster1::> cluster log-forwarding create -destination
192.168.123.96
-port 514 -facility user

cluster1::> cluster log-forwarding create -destination
192.168.123.98
-port 514 -protocol tcp-encrypted -facility user
```

- If the `cluster log-forwarding create` command cannot ping the destination host to verify connectivity, the command fails with an error. Although not recommended, using the `-force` parameter with the command bypasses the connectivity verification.
 - When you set the `-verify-server` parameter to `true`, the identity of the log forwarding destination is verified by validating its certificate. You can set the value to `true` only when you select the `tcp-encrypted` value in the `-protocol` field.
2. Verify that the destination records are correct by using the `cluster log-forwarding show` command.

```
cluster1::> cluster log-forwarding show

Destination Host          Port    Protocol          Verify Syslog
-----
192.168.123.96           514    udp-unencrypted  false  user
192.168.123.98           514    tcp-encrypted   true   user
2 entries were displayed.
```

See the man pages for details.

AutoSupport

Manage AutoSupport settings with System Manager

You can use System Manager to manage the settings for your AutoSupport account.

You can perform the following procedures:

View AutoSupport settings

You can use System Manager to view the settings for your AutoSupport account.

Steps

1. In System Manager, click **Cluster > Settings**.

In the **AutoSupport** section, the following information is displayed:

- Status
- Transport protocol
- Proxy server
- From email address


2. In the **AutoSupport** section, select , then select **More Options**.

Additional information is displayed about the AutoSupport connection and email settings. Also, the transfer history of messages is listed.

Generate and send AutoSupport data

In System Manager, you can initiate the generation of AutoSupport messages and choose from which cluster node or nodes the data is collected.

Steps

1. In System Manager, select **Cluster > Settings**.
2. In the **AutoSupport** section, select , then select **Generate and Send**.
3. Enter a subject.
4. Select the check box under **Collect Data From** to specify the nodes from which to collect the data.

Test the connection to AutoSupport

From System Manager, you can send a test message to verify the connection to AutoSupport.

Steps

1. In System Manager, click **Cluster > Settings**.
2. In the **AutoSupport** section, select , then select **Test Connectivity**.
3. Enter a subject for the message.

Enable or disable AutoSupport

AutoSupport delivers proven business benefits to NetApp customers, including proactive identification of possible configuration issues and accelerated support case resolution. AutoSupport is enabled by default in new systems. If required, you can use System Manager to disable the ability of AutoSupport to monitor the



health of your storage system and send you notification messages. You can enable AutoSupport again after it has been disabled.

About this task

Before you disable AutoSupport, you should be aware that you are turning off the NetApp call-home system and you'll lose the following benefits:

- **Health Monitoring:** AutoSupport monitors the health of your storage system and sends notifications to technical support and your internal support organization.
- **Automation:** AutoSupport automates the reporting of support cases. Most support cases are opened automatically before customers realize there's a problem.
- **Faster resolution:** Systems sending AutoSupport data have their support cases resolved in half of the time compared to cases for systems that not sending AutoSupport data.
- **Faster upgrades:** AutoSupport powers customer self-service workflows, such as version upgrades, add-ons, renewals, and firmware update automation in System Manager.
- **More functions:** Certain functions in other tools work only when AutoSupport is enabled, for example, some workflows in BlueXP.

Steps

1. Select **Cluster > Settings**.
2. In the **AutoSupport** section, select , then select **Disable**.
3. If you want to enable AutoSupport again, in the **AutoSupport** section, select , then select **Enable**.

Suppress the generation of support cases


Beginning with ONTAP 9.10.1, you can use System Manager to send a request to AutoSupport to suppress the generation of support cases.

About this task

To suppress the generation of support cases, you specify the nodes and number of hours for which you want the suppression to occur.

Suppressing support cases can be especially helpful if you do not want AutoSupport to create automated cases while you are performing maintenance on your systems.


Steps

1. Select **Cluster > Settings**.
2. In the **AutoSupport** section, select , then select **Suppress Support Case Generation**.
3. Enter the number of hours that you want the suppression to occur.
4. Select the nodes for which you want the suppression to occur.

Resume the generation of support cases

Beginning with ONTAP 9.10.1, you can use System Manager to resume the generation of support cases from AutoSupport if it has been suppressed.

Steps



1. Select **Cluster > Settings**.
2. In the **AutoSupport** section, select , then select **Resume Support Case Generation**.

3. Select the nodes for which you want the generation to resume.

Edit AutoSupport settings

You can use System Manager to modify the connection and email settings for your AutoSupport account.

Steps

1. Select **Cluster > Settings**.
2. In the **AutoSupport** section, select , then select **More Options**.
3. In the **Connections** section or the **Email** section, select  **Edit** to modify the settings for either section.

Manage AutoSupport with the CLI

Manage AutoSupport overview

AutoSupport is a mechanism that proactively monitors the health of your system and automatically sends messages to NetApp technical support, your internal support organization, and a support partner. Although AutoSupport messages to technical support are enabled by default, you must set the correct options and have a valid mail host to have messages sent to your internal support organization.

Only the cluster administrator can perform AutoSupport management. The storage virtual machine (SVM) administrator has no access to AutoSupport.

AutoSupport is enabled by default when you configure your storage system for the first time. AutoSupport begins sending messages to technical support 24 hours after AutoSupport is enabled. You can shorten the 24-hour period by upgrading or reverting the system, modifying the AutoSupport configuration, or changing the system time to be something other than a 24-hour period.



You can disable AutoSupport at any time, but you should leave it enabled. Enabling AutoSupport can significantly help speed problem determination and resolution should a problem occur on your storage system. By default, the system collects AutoSupport information and stores it locally, even if you disable AutoSupport.

For more information about AutoSupport, see the NetApp Support Site.

Related information

- [NetApp Support](#)
- [Learn more about the AutoSupport commands in the ONTAP CLI](#)

Use AutoSupport and Active IQ Digital Advisor

The AutoSupport component of ONTAP collects telemetry and sends it for analysis. Active IQ Digital Advisor analyzes the data from AutoSupport and provides proactive care and optimization. Using artificial intelligence, Active IQ can identify potential problems and help you resolve them before they impact your business.

Active IQ enables you to optimize your data infrastructure across your global hybrid cloud by delivering actionable predictive analytics and proactive support through a cloud-based portal and mobile app. Data-driven insights and recommendations from Active IQ are available to all NetApp customers with an active

SupportEdge contract (features vary by product and support tier).

Here are some things you can do with Active IQ:

- Plan upgrades. Active IQ identifies issues in your environment that can be resolved by upgrading to a newer version of ONTAP and the Upgrade Advisor component helps you plan for a successful upgrade.
- View system wellness. Your Active IQ dashboard reports any issues with wellness and helps you correct those issues. Monitor system capacity to make sure you never run out of storage space. View support cases for your system.
- Manage performance. Active IQ shows system performance over a longer period than you can see in System Manager. Identify configuration and system issues that are impacting your performance.
- Maximize efficiency. View storage efficiency metrics and identify ways to store more data in less space.
- View inventory and configuration. Active IQ displays complete inventory and software and hardware configuration information. See when service contracts are expiring and renew them to ensure you remain supported.

Related information

[NetApp Documentation: Active IQ Digital Advisor](#)

[Launch Active IQ](#)

[SupportEdge Services](#)

When and where AutoSupport messages are sent

AutoSupport sends messages to different recipients, depending on the type of message. Learning when and where AutoSupport sends messages can help you understand messages that you receive through email or view on the Active IQ (formerly known as My AutoSupport) web site.

Unless specified otherwise, settings in the following tables are parameters of the `system node autosupport modify` command.

Event-triggered messages

When events occur on the system that require corrective action, AutoSupport automatically sends an event-triggered message.

When the message is sent	Where the message is sent
AutoSupport responds to a trigger event in the EMS	Addresses specified in <code>-to</code> and <code>-noteto</code> . (Only critical, service-affecting events are sent.) Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to <code>enable</code>

Scheduled messages

AutoSupport automatically sends several messages on a regular schedule.

When the message is sent	Where the message is sent
Daily (by default, sent between 12:00 a.m. and 1:00 a.m. as a log message)	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to <code>enable</code>
Daily (by default, sent between 12:00 a.m. and 1:00 a.m. as a performance message), if the <code>-perf</code> parameter is set to <code>true</code>	Addresses specified in <code>-partner-address`</code> Technical support, if <code>-support</code> is set to <code>enable</code>
Weekly (by default, sent Sunday between 12:00 a.m. and 1:00 a.m.)	Addresses specified in <code>-partner-address</code> Technical support, if <code>-support</code> is set to <code>enable</code>

Manually triggered messages

You can manually initiate or resend an AutoSupport message.

When the message is sent	Where the message is sent
You manually initiate a message using the <code>system node autosupport invoke</code> command	If a URI is specified using the <code>-uri</code> parameter in the <code>system node autosupport invoke</code> command, the message is sent to that URI. If <code>-uri</code> is omitted, the message is sent to the addresses specified in <code>-to</code> and <code>-partner-address</code> . The message is also sent to technical support if <code>-support</code> is set to <code>enable</code> .
You manually initiate a message using the <code>system node autosupport invoke-core-upload</code> command	If a URI is specified using the <code>-uri</code> parameter in the <code>system node autosupport invoke-core-upload</code> command, the message is sent to that URI, and the core dump file is uploaded to the URI. If <code>-uri</code> is omitted in the <code>system node autosupport invoke-core-upload</code> command, the message is sent to technical support, and the core dump file is uploaded to the technical support site. Both scenarios require that <code>-support</code> is set to <code>enable</code> and <code>-transport</code> is set to <code>https</code> or <code>http</code> . Due to the large size of core dump files, the message is not sent to the addresses specified in the <code>-to</code> and <code>-partner-addresses</code> parameters.

When the message is sent	Where the message is sent
You manually initiate a message using the <code>system node autosupport invoke-performance-archive</code> command	<p>If a URI is specified using the <code>-uri</code> parameter in the <code>system node autosupport invoke-performance-archive</code> command, the message is sent to that URI, and the performance archive file is uploaded to the URI.</p> <p>If <code>-uri</code> is omitted in the <code>system node autosupport invoke-performance-archive</code>, the message is sent to technical support, and the performance archive file is uploaded to the technical support site.</p> <p>Both scenarios require that <code>-support</code> is set to <code>enable</code> and <code>-transport</code> is set to <code>https</code> or <code>http</code>.</p> <p>Due to the large size of performance archive files, the message is not sent to the addresses specified in the <code>-to</code> and <code>-partner-addresses</code> parameters.</p>
You manually resend a past message using the <code>system node autosupport history retransmit</code> command	Only to the URI that you specify in the <code>-uri</code> parameter of the <code>system node autosupport history retransmit</code> command

Messages triggered by technical support

Technical support can request messages from AutoSupport using the AutoSupport OnDemand feature.

When the message is sent	Where the message is sent
When AutoSupport obtains delivery instructions to generate new AutoSupport messages	<p>Addresses specified in <code>-partner-address</code></p> <p>Technical support, if <code>-support</code> is set to <code>enable</code> and <code>-transport</code> is set to <code>https</code></p>
When AutoSupport obtains delivery instructions to resend past AutoSupport messages	Technical support, if <code>-support</code> is set to <code>enable</code> and <code>-transport</code> is set to <code>https</code>
When AutoSupport obtains delivery instructions to generate new AutoSupport messages that upload core dump or performance archive files	Technical support, if <code>-support</code> is set to <code>enable</code> and <code>-transport</code> is set to <code>https</code> . The core dump or performance archive file is uploaded to the technical support site.

How AutoSupport creates and sends event-triggered messages

AutoSupport creates event-triggered AutoSupport messages when the EMS processes a trigger event. An event-triggered AutoSupport message alerts recipients to problems that require corrective action and contains only information that is relevant to the problem. You

can customize what content to include and who receives the messages.

AutoSupport uses the following process to create and send event-triggered AutoSupport messages:

1. When the EMS processes a trigger event, EMS sends AutoSupport a request.

A trigger event is an EMS event with an AutoSupport destination and a name that begins with a `callhome.` prefix.

2. AutoSupport creates an event-triggered AutoSupport message.

AutoSupport collects basic and troubleshooting information from subsystems that are associated with the trigger to create a message that includes only information that is relevant to the trigger event.

A default set of subsystems is associated with each trigger. However, you can choose to associate additional subsystems with a trigger by using the `system node autosupport trigger modify` command.

3. AutoSupport sends the event-triggered AutoSupport message to the recipients defined by the `system node autosupport modify` command with the `-to`, `-noteto`, `-partner-address`, and `-support` parameters.

You can enable and disable delivery of AutoSupport messages for specific triggers by using the `system node autosupport trigger modify` command with the `-to` and `-noteto` parameters.

Example of data sent for a specific event

The `storage shelf PSU failed` EMS event triggers a message that contains basic data from the Mandatory, Log Files, Storage, RAID, HA, Platform, and Networking subsystems and troubleshooting data from the Mandatory, Log Files, and Storage subsystems.

You decide that you want to include data about NFS in any AutoSupport messages sent in response to a future `storage shelf PSU failed` event. You enter the following command to enable troubleshooting-level data for NFS for the `callhome.shlf.ps.fault` event:

```
cluster1::\>
system node autosupport trigger modify -node node1 -autosupport
-message shlf.ps.fault -troubleshooting-additional nfs
```

Note that the `callhome.` prefix is dropped from the `callhome.shlf.ps.fault` event when you use the `system node autosupport trigger` commands, or when referenced by AutoSupport and EMS events in the CLI.

Types of AutoSupport messages and their content

AutoSupport messages contain status information about supported subsystems. Learning what AutoSupport messages contain can help you interpret or respond to messages that you receive in email or view on the Active IQ (formerly known as My AutoSupport) web site.

Type of message	Type of data the message contains
Event-triggered	Files containing context-sensitive data about the specific subsystem where the event occurred
Daily	Log files
Performance	Performance data sampled during the previous 24 hours
Weekly	Configuration and status data
Triggered by the <code>system node autosupport invoke</code> command	<p>Depends on the value specified in the <code>-type</code> parameter:</p> <ul style="list-style-type: none"> • <code>test</code> sends a user-triggered message with some basic data. <p>This message also triggers an automated email response from technical support to any specified email addresses, using the <code>-to</code> option, so that you can confirm that AutoSupport messages are being received.</p> <ul style="list-style-type: none"> • <code>performance</code> sends performance data. • <code>all</code> sends a user-triggered message with a complete set of data similar to the weekly message, including troubleshooting data from each subsystem. <p>Technical support typically requests this message.</p>
Triggered by the <code>system node autosupport invoke-core-upload</code> command	Core dump files for a node
Triggered by the <code>system node autosupport invoke-performance-archive</code> command	Performance archive files for a specified period of time

Type of message	Type of data the message contains
Triggered by AutoSupport OnDemand	<p>AutoSupport OnDemand can request new messages or past messages:</p> <ul style="list-style-type: none"> • New messages, depending on the type of AutoSupport collection, can be <code>test</code>, <code>all</code>, or <code>performance</code>. • Past messages depend on the type of message that is resent. <p>AutoSupport OnDemand can request new messages that upload the following files to the NetApp Support Site at mysupport.netapp.com:</p> <ul style="list-style-type: none"> • Core dump • Performance archive

What AutoSupport subsystems are

Each subsystem provides basic and troubleshooting information that AutoSupport uses for its messages. Each subsystem is also associated with trigger events that allow AutoSupport to collect from subsystems only information that is relevant to the trigger event.

AutoSupport collects context-sensitive content. You can view information about subsystems and trigger events by using the `system node autosupport trigger show` command.

AutoSupport size and time budgets

AutoSupport collects information, organized by subsystem, and enforces a size and time budget on content for each subsystem. As storage systems grow, AutoSupport budgets provide control over the AutoSupport payload, which in turn provides scalable delivery of AutoSupport data.

AutoSupport stops collecting information and truncates the AutoSupport content if the subsystem content exceeds its size or time budget. If the content cannot be truncated easily (for example, binary files), AutoSupport omits the content.

You should modify the default size and time budgets only if asked to do so by NetApp Support. You can also review the default size and time budgets of the subsystems by using the `autosupport manifest show` command.

Files sent in event-triggered AutoSupport messages

Event-triggered AutoSupport messages only contain basic and troubleshooting information from subsystems that are associated with the event that caused AutoSupport to generate the message. The specific data helps NetApp support and support partners troubleshoot the problem.

AutoSupport uses the following criteria to control content in event-triggered AutoSupport messages:

- Which subsystems are included

Data is grouped into subsystems, including common subsystems, such as Log Files, and specific subsystems, such as RAID. Each event triggers a message that contains only the data from specific subsystems.

- The detail level of each included subsystem

Data for each included subsystem is provided at a basic or troubleshooting level.

You can view all possible events and determine which subsystems are included in messages about each event using the `system node autosupport trigger show` command with the `-instance` parameter.

In addition to the subsystems that are included by default for each event, you can add additional subsystems at either a basic or a troubleshooting level using the `system node autosupport trigger modify` command.

Log files sent in AutoSupport messages

AutoSupport messages can contain several key log files that enable technical support staff to review recent system activity.

All types of AutoSupport messages might include the following log files when the Log Files subsystem is enabled:

Log file	Amount of data included from the file
<ul style="list-style-type: none">• Log files from the <code>/mroot/etc/log/mlog/</code> directory• The MESSAGES log file	<p>Only new lines added to the logs since the last AutoSupport message up to a specified maximum. This ensures that AutoSupport messages have unique, relevant—not overlapping—data.</p> <p>(Log files from partners are the exception; for partners, the maximum allowed data is included.)</p>
<ul style="list-style-type: none">• Log files from the <code>/mroot/etc/log/shelflog/</code> directory• Log files from the <code>/mroot/etc/log/acp/</code> directory• Event Management System (EMS) log data	<p>The most recent lines of data up to a specified maximum.</p>

The content of AutoSupport messages can change between releases of ONTAP.

Files sent in weekly AutoSupport messages

Weekly AutoSupport messages contain additional configuration and status data that is useful to track changes in your system over time.

The following information is sent in weekly AutoSupport messages:

- Basic information about every subsystem
- Contents of selected `/mroot/etc` directory files
- Log files
- Output of commands that provide system information
- Additional information, including replicated database (RDB) information, service statistics, and more

How AutoSupport OnDemand obtains delivery instructions from technical support

AutoSupport OnDemand periodically communicates with technical support to obtain delivery instructions for sending, resending, and declining AutoSupport messages as well as uploading large files to the NetApp support site. AutoSupport OnDemand enables AutoSupport messages to be sent on-demand instead of waiting for the weekly AutoSupport job to run.

AutoSupport OnDemand consists of the following components:

- AutoSupport OnDemand client that runs on each node
- AutoSupport OnDemand service that resides in technical support

The AutoSupport OnDemand client periodically polls the AutoSupport OnDemand service to obtain delivery instructions from technical support. For example, technical support can use the AutoSupport OnDemand service to request that a new AutoSupport message be generated. When the AutoSupport OnDemand client polls the AutoSupport OnDemand service, the client obtains the delivery instructions and sends the new AutoSupport message on-demand as requested.

AutoSupport OnDemand is enabled by default. However, AutoSupport OnDemand relies on some AutoSupport settings to continue communicating with technical support. AutoSupport OnDemand automatically communicates with technical support when the following requirements are met:

- AutoSupport is enabled.
- AutoSupport is configured to send messages to technical support.
- AutoSupport is configured to use the HTTPS transport protocol.


The AutoSupport OnDemand client sends HTTPS requests to the same technical support location to which AutoSupport messages are sent. The AutoSupport OnDemand client does not accept incoming connections.



AutoSupport OnDemand uses the “autosupport” user account to communicate with technical support. ONTAP prevents you from deleting this account.

If you want to disable AutoSupport OnDemand, but keep AutoSupport enabled, use the command: `system node autosupport modify -ondemand-state disable`.

The following illustration shows how AutoSupport OnDemand sends HTTPS requests to technical support to obtain delivery instructions.



The delivery instructions can include requests for AutoSupport to do the following:

- Generate new AutoSupport messages.

Technical support might request new AutoSupport messages to help triage issues.

- Generate new AutoSupport messages that upload core dump files or performance archive files to the NetApp support site.

Technical support might request core dump or performance archive files to help triage issues.

- Retransmit previously generated AutoSupport messages.

This request automatically happens if a message was not received due to a delivery failure.

- Disable delivery of AutoSupport messages for specific trigger events.

Technical support might disable delivery of data that is not used.

Structure of AutoSupport messages sent by email

When an AutoSupport message is sent by email, the message has a standard subject, a brief body, and a large attachment in 7z file format that contains the data.



If AutoSupport is configured to hide private data, certain information, such as the hostname, is omitted or masked in the header, subject, body, and attachments.

Subject

The subject line of messages sent by the AutoSupport mechanism contains a text string that identifies the reason for the notification. The format of the subject line is as follows:

HA Group Notification from *System_Name* (*Message*) *Severity*

- *System_Name* is either the hostname or the system ID, depending on the AutoSupport configuration

Body

The body of the AutoSupport message contains the following information:

- Date and timestamp of the message
- Version of ONTAP on the node that generated the message

- System ID, serial number, and hostname of the node that generated the message
- AutoSupport sequence number
- SNMP contact name and location, if specified
- System ID and hostname of the HA partner node

Attached files

The key information in an AutoSupport message is contained in files that are compressed into a 7z file called `body.7z` and attached to the message.

The files contained in the attachment are specific to the type of AutoSupport message.

AutoSupport severity types

AutoSupport messages have severity types that help you understand the purpose of each message—for example, to draw immediate attention to an emergency problem, or only to provide information.

Messages have one of the following severities:

- **Alert:** Alert messages indicate that a next-higher level event might occur if you do not take some action.
You must take an action against alert messages within 24 hours.
- **Emergency:** Emergency messages are displayed when a disruption has occurred.
You must take an action against emergency messages immediately.
- **Error:** Error conditions indicate what might happen if you ignore.
- **Notice:** Normal but significant condition.
- **Info:** Informational message provides details about the issue, which you can ignore.
- **Debug:** Debug-level messages provide instructions you should perform.

If your internal support organization receives AutoSupport messages through email, the severity appears in the subject line of the email message.

Requirements for using AutoSupport

You must use HTTPS with TLSv1.2 or secure SMTP for delivery of AutoSupport messages to provide the best security and to support all of the latest AutoSupport features. AutoSupport messages delivered with any other protocol will be rejected.

Supported protocols

All of these protocols run on IPv4 or IPv6, based on the address family to which the name resolves.

Protocol and port	Description
HTTPS on port 443	<p>This is the default protocol. You should use this whenever possible.</p> <p>This protocol supports AutoSupport OnDemand and uploads of large files.</p> <p>The certificate from the remote server is validated against the root certificate, unless you disable validation.</p> <p>The delivery uses an HTTPS PUT request. With PUT, if the request fails during transmission, the request restarts where it stopped. If the server receiving the request does not support PUT, the delivery uses an HTTPS POST request.</p>
HTTP on port 80	<p>This protocol is preferred over SMTP.</p> <p>This protocol supports uploads of large files, but not AutoSupport OnDemand.</p> <p>The delivery uses an HTTPS PUT request. With PUT, if the request fails during transmission, the request restarts where it stopped. If the server receiving the request does not support PUT, the delivery uses an HTTPS POST request.</p>
SMTP on port 25 or another port	<p>You should use this protocol only if the network connection does not allow HTTPS.</p> <p>The default port value is 25, but you can configure AutoSupport to use a different port.</p> <p>Keep the following limitations in mind when using SMTP:</p> <ul style="list-style-type: none"> • AutoSupport OnDemand and uploads of large files are not supported. • Data is not encrypted. <p>SMTP sends data in clear text, making text in the AutoSupport message easy to intercept and read.</p> <ul style="list-style-type: none"> • Limitations on message length and line length can be introduced.

If you configure AutoSupport with specific email addresses for your internal support organization, or a support partner organization, those messages are always sent by SMTP.

For example, if you use the recommended protocol to send messages to technical support and you also want to send messages to your internal support organization, your messages will be transported using both HTTPS

and SMTP, respectively.

AutoSupport limits the maximum file size for each protocol. The default setting for HTTP and HTTPS transfers is 25 MB. The default setting for SMTP transfers is 5 MB. If the size of the AutoSupport message exceeds the configured limit, AutoSupport delivers as much of the message as possible. You can edit the maximum size by modifying AutoSupport configuration. See the `system node autosupport modify` man page for more information.



AutoSupport automatically overrides the maximum file size limit for the HTTPS and HTTP protocols when you generate and send AutoSupport messages that upload core dump or performance archive files to the NetApp support site or a specified URI. The automatic override applies only when you upload files by using the `system node autosupport invoke-core-upload` or the `system node autosupport invoke-performance-archive` commands.

Configuration requirements

Depending on your network configuration, the HTTPS protocol may require additional configuration of a proxy URL. If HTTPS to send AutoSupport messages to technical support and you have a proxy, you must identify the URL for that proxy. If the proxy uses a port other than the default port, which is 3128, you can specify the port for that proxy. You can also specify a user name and password for proxy authentication.

If you use SMTP to send AutoSupport messages either to your internal support organization or to technical support, you must configure an external mail server. The storage system does not function as a mail server; it requires an external mail server at your site to send mail. The mail server must be a host that listens on the SMTP port (25) or another port, and it must be configured to send and receive 8-bit Multipurpose Internet Mail Extensions (MIME) encoding. Example mail hosts include a UNIX host running an SMTP server such as the sendmail program and a Windows server running the Microsoft Exchange server. You can have one or more mail hosts.

Set up AutoSupport

You can control whether and how AutoSupport information is sent to technical support and your internal support organization, and then test that the configuration is correct.

About this task

In ONTAP 9.5 and later releases, you can enable AutoSupport and modify its configuration on all nodes of the cluster simultaneously. When a new node joins the cluster, the node inherits the AutoSupport cluster configuration automatically. You do not have to update the configuration on each node separately.



Beginning with ONTAP 9.5, the scope of the `system node autosupport modify` command is cluster-wide. The AutoSupport configuration is modified on all nodes in the cluster, even when the `-node` option is specified. The option is ignored, but it has been retained for CLI backward compatibility.

In ONTAP 9.4 and earlier releases, the scope of the `system node autosupport modify` command is specific to the node. The AutoSupport configuration should be modified on each node in your cluster.

By default, AutoSupport is enabled on each node to send messages to technical support by using the HTTPS transport protocol.

You must use HTTPS with TLSv1.2 or secure SMTP for delivery of AutoSupport messages to provide the best

security and to support all of the latest AutoSupport features.

Steps

1. Ensure that AutoSupport is enabled:

```
system node autosupport modify -state enable
```

2. If you want technical support to receive AutoSupport messages, use the following command:

```
system node autosupport modify -support enable
```

You must enable this option if you want to enable AutoSupport to work with AutoSupport OnDemand or if you want to upload large files, such as core dump and performance archive files, to technical support or a specified URL.

3. If technical support is enabled to receive AutoSupport messages, specify which transport protocol to use for the messages.

You can choose from the following options:

If you want to...	Then set the following parameters of the <code>system node autosupport modify</code> command...
Use the default HTTPS protocol	<ol style="list-style-type: none">a. Set <code>-transport</code> to <code>https</code>.b. If you use a proxy, set <code>-proxy-url</code> to the URL of your proxy. This configuration supports communication with AutoSupport OnDemand and uploads of large files.
Use SMTP	Set <code>-transport</code> to <code>smtp</code> . This configuration does not support AutoSupport OnDemand or uploads of large files.

4. If you want your internal support organization or a support partner to receive AutoSupport messages, perform the following actions:

- a. Identify the recipients in your organization by setting the following parameters of the `system node autosupport modify` command:

Set this parameter...	To this...
<code>-to</code>	Up to five comma-separated individual email addresses or distribution lists in your internal support organization that will receive key AutoSupport messages

<code>-noteto</code>	Up to five comma-separated individual email addresses or distribution lists in your internal support organization that will receive a shortened version of key AutoSupport messages designed for cell phones and other mobile devices
<code>-partner-address</code>	Up to five comma-separated individual email addresses or distribution lists in your support partner organization that will receive all AutoSupport messages

- b. Check that addresses are correctly configured by listing the destinations using the `system node autosupport destinations show` command.
5. If you are sending messages to your internal support organization or you chose SMTP transport for messages to technical support, configure SMTP by setting the following parameters of the `system node autosupport modify` command:
 - Set `-mail-hosts` to one or more mail hosts, separated by commas.

You can set a maximum of five.

You can configure a port value for each mail host by specifying a colon and port number after the mail host name: for example, `mymailhost.example.com:5678`, where 5678 is the port for the mail host.
 - Set `-from` to the email address that sends the AutoSupport message.
 6. Configure DNS.
 7. Optionally, add command options if you want to change specific settings:

If you want to do this...	Then set the following parameters of the <code>system node autosupport modify</code> command...
Hide private data by removing, masking, or encoding sensitive data in the messages	Set <code>-remove-private-data</code> to <code>true</code> . If you change from <code>false</code> to <code>true</code> , all AutoSupport history and all associated files are deleted.
Stop sending performance data in periodic AutoSupport messages	Set <code>-perf</code> to <code>false</code> .

8. Check the overall configuration by using the `system node autosupport show` command with the `-node` parameter.
9. Verify the AutoSupport operation by using the `system node autosupport check show` command.

If any problems are reported, use the `system node autosupport check show-details` command to view more information.
10. Test that AutoSupport messages are being sent and received:
 - a. Use the `system node autosupport invoke` command with the `-type` parameter set to `test`.


```
cluster1::> system node autosupport invoke -type test -node node1
```

- b. Confirm that NetApp is receiving your AutoSupport messages:

```
system node autosupport history show -node local
```

The status of the latest outgoing AutoSupport message should eventually change to `sent-successful` for all appropriate protocol destinations.

- c. Optionally, confirm that the AutoSupport message is being sent to your internal support organization or to your support partner by checking the email of any address that you configured for the `-to`, `-noteto`, or `-partner-address` parameters of the `system node autosupport modify` command.

Upload core dump files

When a core dump file is saved, an event message is generated. If the AutoSupport service is enabled and configured to send messages to NetApp support, an AutoSupport message is transmitted, and an automated email acknowledgement is sent to you.

What you'll need

- You must have set up AutoSupport with the following settings:
 - AutoSupport is enabled on the node.
 - AutoSupport is configured to send messages to technical support.
 - AutoSupport is configured to use the HTTP or HTTPS transport protocol.

The SMTP transport protocol is not supported when sending messages that include large files, such as core dump files.

About this task

You can also upload the core dump file through the AutoSupport service over HTTPS by using the `system node autosupport invoke-core-upload` command, if requested by NetApp support.

How to upload a file to NetApp

Steps

1. View the core dump files for a node by using the `system node coredump show` command.

In the following example, core dump files are displayed for the local node:

```
cluster1::> system node coredump show -node local
Node:Type Core Name Saved Panic Time
-----
node:kernel
core.4073000068.2013-09-11.15_05_01.nz true 9/11/2013 15:05:01
```

2. Generate an AutoSupport message and upload a core dump file by using the `system node autosupport invoke-core-upload` command.

In the following example, an AutoSupport message is generated and sent to the default location, which is technical support, and the core dump file is uploaded to the default location, which is the NetApp support site:

```
cluster1::> system node autosupport invoke-core-upload -core-filename
core.4073000068.2013-09-11.15_05_01.nz -node local
```

In the following example, an AutoSupport message is generated and sent to the location specified in the URI, and the core dump file is uploaded to the URI:

```
cluster1::> system node autosupport invoke-core-upload -uri
https://files.company.com -core-filename
core.4073000068.2013-09-11.15_05_01.nz -node local
```

Upload performance archive files

You can generate and send an AutoSupport message that contains a performance archive. By default, NetApp technical support receives the AutoSupport message, and the performance archive is uploaded to the NetApp support site. You can specify an alternate destination for the message and upload.

What you'll need

- You must have set up AutoSupport with the following settings:
 - AutoSupport is enabled on the node.
 - AutoSupport is configured to send messages to technical support.
 - AutoSupport is configured to use the HTTP or HTTPS transport protocol.

The SMTP transport protocol is not supported when sending messages that include large files, such as performance archive files.

About this task

You must specify a start date for the performance archive data that you want to upload. Most storage systems retain performance archives for two weeks, enabling you to specify a start date up to two weeks ago. For example, if today is January 15, you can specify a start date of January 2.

Step

1. Generate an AutoSupport message and upload the performance archive file by using the `system node autosupport invoke-performance-archive` command.

In the following example, 4 hours of performance archive files from January 12, 2015 are added to an AutoSupport message and uploaded to the default location, which is the NetApp support site:

```
cluster1::> system node autosupport invoke-performance-archive -node
local -start-date 1/12/2015 13:42:09 -duration 4h
```

In the following example, 4 hours of performance archive files from January 12, 2015 are added to an AutoSupport message and uploaded to the location specified by the URI:

```
cluster1::> system node autosupport invoke-performance-archive -node
local -start-date 1/12/2015 13:42:09 -duration 4h -uri
https://files.company.com
```

Get AutoSupport message descriptions

The descriptions of the AutoSupport messages that you receive are available through the ONTAP Syslog Translator.

Steps

1. Go to the [Syslog Translator](#).
2. In the **Release** field, enter the the version of ONTAP you are using. In the **Search String** field, enter "callhome". Select **Translate**.
3. The Syslog Translator will alphabetically list all events that match the message string you entered.

Commands for managing AutoSupport

You use the `system node autosupport` commands to change or view AutoSupport configuration, display information about previous AutoSupport messages, and send, resend or cancel an AutoSupport message.

Configure AutoSupport

If you want to...	Use this command...
Control whether any AutoSupport messages are sent	<code>system node autosupport modify with the -state parameter</code>
Control whether AutoSupport messages are sent to technical support	<code>system node autosupport modify with the -support parameter</code>
Set up AutoSupport or modify the configuration of AutoSupport	<code>system node autosupport modify</code>
Enable and disable AutoSupport messages to your internal support organization for individual trigger events, and specify additional subsystem reports to include in messages sent in response to individual trigger events	<code>system node autosupport trigger modify</code>



Display information about the AutoSupport configuration

If you want to...	Use this command...
Display the AutoSupport configuration	<code>system node autosupport show</code> with the <code>-node</code> parameter
View a summary of all addresses and URLs that receive AutoSupport messages	<code>system node autosupport destinations show</code>
Display which AutoSupport messages are sent to your internal support organization for individual trigger events	<code>system node autosupport trigger show</code>
Display status of AutoSupport configuration as well as delivery to various destinations	<code>system node autosupport check show</code>
Display detailed status of AutoSupport configuration as well as delivery to various destinations	<code>system node autosupport check show-details</code>

Display information about past AutoSupport messages

If you want to...	Use this command...
Display information about one or more of the 50 most recent AutoSupport messages	<code>system node autosupport history show</code>
Display information about recent AutoSupport messages generated to upload core dump or performance archive files to the technical support site or a specified URI	<code>system node autosupport history show-upload-details</code>
View the information in the AutoSupport messages including the name and size of each file collected for the message along with any errors	<code>system node autosupport manifest show</code>

Send, resend, or cancel AutoSupport messages

If you want to...	Use this command...
<p>Retransmit a locally stored AutoSupport message, identified by its AutoSupport sequence number</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>If you retransmit an AutoSupport message, and if support already received that message, the support system will not create a duplicate case. If, on the other hand, support did not receive that message, then the AutoSupport system will analyze the message and create a case, if necessary.</p> </div>	<pre>system node autosupport history retransmit</pre>
<p>Generate and send an AutoSupport message—for example, for testing purposes</p>	<pre>system node autosupport invoke</pre> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>Use the <code>-force</code> parameter to send a message even if AutoSupport is disabled. Use the <code>-uri</code> parameter to send the message to the destination you specify instead of the configured destination.</p> </div>
<p>Cancel an AutoSupport message</p>	<pre>system node autosupport history cancel</pre>

Related information

[ONTAP 9 Commands](#)

Information included in the AutoSupport manifest

The AutoSupport manifest provides you with a detailed view of the files collected for each AutoSupport message. The AutoSupport manifest also includes information about collection errors when AutoSupport cannot collect the files it needs.

The AutoSupport manifest includes the following information:

- Sequence number of the AutoSupport message
- Which files AutoSupport included in the AutoSupport message
- Size of each file, in bytes
- Status of the AutoSupport manifest collection
- Error description, if AutoSupport failed to collect one or more files

You can view the AutoSupport manifest by using the `system node autosupport manifest show` command.

The AutoSupport manifest is included with every AutoSupport message and presented in XML format, which means that you can either use a generic XML viewer to read it or view it using the Active IQ (formerly known as

My AutoSupport) portal.

AutoSupport case suppression during scheduled maintenance windows

AutoSupport case suppression enables you to stop unnecessary cases from being created by AutoSupport messages that are triggered during scheduled maintenance windows.

To suppress AutoSupport cases, you must manually invoke an AutoSupport message with a specially formatted text string: `MAINT=xh`. `x` is the duration of the maintenance window in units of hours.

Related information

[How to suppress automatic case creation during scheduled maintenance windows](#)

Troubleshoot AutoSupport when messages are not received

If the system does not send the AutoSupport message, you can determine whether that is because AutoSupport cannot generate the message or cannot deliver the message.

Steps

1. Check delivery status of the messages by using the `system node autosupport history show` command.
2. Read the status.

This status	Means
initializing	The collection process is starting. If this state is temporary, all is well. However, if this state persists, there is an issue.
collection-failed	AutoSupport cannot create the AutoSupport content in the spool directory. You can view what AutoSupport is trying to collect by entering the <code>system node autosupport history show -detail</code> command.
collection-in-progress	AutoSupport is collecting AutoSupport content. You can view what AutoSupport is collecting by entering the <code>system node autosupport manifest show</code> command.
queued	AutoSupport messages are queued for delivery, but not yet delivered.
transmitting	AutoSupport is currently delivering messages.
sent-successful	AutoSupport successfully delivered the message. You can find out where AutoSupport delivered the message by entering the <code>system node autosupport history show -delivery</code> command.
ignore	AutoSupport has no destinations for the message. You can view the delivery details by entering the <code>system node autosupport history show -delivery</code> command.

This status	Means
re-queued	AutoSupport tried to deliver messages, but the attempt failed. As a result, AutoSupport placed the messages back in the delivery queue for another attempt. You can view the error by entering the <code>system node autosupport history show</code> command.
transmission-failed	AutoSupport failed to deliver the message the specified number of times and stopped trying to deliver the message. You can view the error by entering the <code>system node autosupport history show</code> command.
ondemand-ignore	The AutoSupport message was processed successfully, but the AutoSupport OnDemand service chose to ignore it.

3. Perform one of the following actions:

For this status	Do this
initializing or collection-failed	Contact NetApp Support, because AutoSupport cannot generate the message. Mention the following Knowledge Base article: AutoSupport is failing to deliver: status is stuck in initializing
ignore, re-queued, or transmission failed	Check that destinations are correctly configured for SMTP, HTTP, or HTTPS because AutoSupport cannot deliver the message.

Troubleshoot AutoSupport message delivery over HTTP or HTTPS

If the system does not send the expected AutoSupport message and you are using HTTP or HTTPS, or the Automatic Update feature is not working, you can check a number of settings to resolve the problem.

What you'll need

You should have confirmed basic network connectivity and DNS lookup:

- Your node management LIF must be up for operational and administrative status.
- You must be able to ping a functioning host on the same subnet from the cluster management LIF (not a LIF on any of the nodes).
- You must be able to ping a functioning host outside the subnet from the cluster management LIF.
- You must be able to ping a functioning host outside the subnet from the cluster management LIF using the name of the host (not the IP address).

About this task

These steps are for cases when you have determined that AutoSupport can generate the message, but cannot deliver the message over HTTP or HTTPS.

If you encounter errors or cannot complete a step in this procedure, determine and address the problem before proceeding to the next step.

Steps

1. Display the detailed status of the AutoSupport subsystem:

```
system node autosupport check show-details
```

This includes verifying connectivity to AutoSupport destinations by sending test messages and providing a list of possible errors in your AutoSupport configuration settings.

2. Verify the status of the node management LIF:

```
network interface show -home-node local -role node-mgmt -fields  
vserver,lif,status-oper,status-admin,address,role
```

The `status-oper` and `status-admin` fields should return “up”.

3. Record the SVM name, the LIF name, and the LIF IP address for later use.
4. Ensure that DNS is enabled and configured correctly:

```
vserver services name-service dns show
```

5. Address any errors returned by the AutoSupport message:

```
system node autosupport history show -node * -fields node,seq-  
num,destination,last-update,status,error
```

For assistance troubleshooting any returned errors, see the [ONTAP AutoSupport \(Transport HTTPS and HTTP\) Resolution Guide](#).

6. Confirm that the cluster can access both the servers it needs and the Internet successfully:

- a. `network traceroute -lif node-management_LIF -destination DNS server`

- b. `network traceroute -lif node_management_LIF -destination support.netapp.com`



The address `support.netapp.com` itself does not respond to ping/traceroute, but the per-hop information is valuable.

- c. `system node autosupport show -fields proxy-url`

- d. `network traceroute -node node_management_LIF -destination proxy_url`

If any of these routes are not functioning, try the same route from a functioning host on the same subnet as the cluster, using the “traceroute” or “tracert” utility found on most third-party network clients. This assists you in determining whether the issue is in your network configuration or your cluster configuration.

7. If you are using HTTPS for your AutoSupport transport protocol, ensure that HTTPS traffic can exit your network:
 - a. Configure a web client on the same subnet as the cluster management LIF.

Ensure that all configuration parameters are the same values as for the AutoSupport configuration, including using the same proxy server, user name, password, and port.

- b. Access `https://support.netapp.com` with the web client.

The access should be successful. If not, ensure that all firewalls are configured correctly to allow HTTPS and DNS traffic, and that the proxy server is configured correctly. For more information on configuring static name resolution for `support.netapp.com`, see the Knowledge Base article [How would a HOST entry be added in ONTAP for support.netapp.com?](#)

8. Beginning with ONTAP 9.10.1, if you enabled the Automatic Update feature, ensure you have HTTPS connectivity to the following additional URLs:
 - `https://support-sg-emea.netapp.com`
 - `https://support-sg-naeast.netapp.com`
 - `https://support-sg-nawest.netapp.com`

Troubleshoot AutoSupport message delivery over SMTP

If the system cannot deliver AutoSupport messages over SMTP, you can check a number of settings to resolve the problem.

What you'll need

You should have confirmed basic network connectivity and DNS lookup:

- Your node management LIF must be up for operational and administrative status.
- You must be able to ping a functioning host on the same subnet from the cluster management LIF (not a LIF on any of the nodes).
- You must be able to ping a functioning host outside the subnet from the cluster management LIF.
- You must be able to ping a functioning host outside the subnet from the cluster management LIF using the name of the host (not the IP address).

About this task

These steps are for cases when you have determined that AutoSupport can generate the message, but cannot deliver the message over SMTP.

If you encounter errors or cannot complete a step in this procedure, determine and address the problem before proceeding to the next step.

All commands are entered at the ONTAP command-line interface, unless otherwise specified.

Steps

1. Verify the status of the node management LIF:

```
network interface show -home-node local -role node-mgmt -fields  
vserver,lif,status-oper,status-admin,address,role
```

The `status-oper` and `status-admin` fields should return up.

2. Record the SVM name, the LIF name, and the LIF IP address for later use.
3. Ensure that DNS is enabled and configured correctly:

```
vserver services name-service dns show
```

4. Display all of the servers configured to be used by AutoSupport:

```
system node autosupport show -fields mail-hosts
```

Record all server names displayed.

5. For each server displayed by the previous step, and `support.netapp.com`, ensure that the server or URL can be reached by the node:

```
network traceroute -node local -destination server_name
```

If any of these routes is not functioning, try the same route from a functioning host on the same subnet as the cluster, using the “traceroute” or “tracert” utility found on most third-party network clients. This assists you in determining whether the issue is in your network configuration or your cluster configuration.

6. Log in to the host designated as the mail host, and ensure that it can serve SMTP requests:

```
netstat -aAn|grep 25
```

25 is the listener SMTP port number.

A message similar to the following text is displayed:

```
ff64878c tcp          0          0 *.25      *.*      LISTEN.
```

7. From some other host, open a Telnet session with the SMTP port of the mail host:

```
telnet mailhost 25
```

A message similar to the following text is displayed:

```
220 filer.yourco.com Sendmail 4.1/SMI-4.1 ready at Thu, 30 Nov 2014  
10:49:04 PST
```

8. At the telnet prompt, ensure that a message can be relayed from your mail host:

```
HELO domain_name
```

```
MAIL FROM: your_email_address
```

```
RCPT TO: autosupport@netapp.com
```

`domain_name` is the domain name of your network.

If an error is returned saying that relaying is denied, relaying is not enabled on the mail host. Contact your system administrator.

9. At the telnet prompt, send a test message:

```
DATA
```

SUBJECT: TESTING
THIS IS A TEST

.



Ensure that you enter the last period (.) on a line by itself. The period indicates to the mail host that the message is complete.

If an error is returned, your mail host is not configured correctly. Contact your system administrator.

10. From the ONTAP command-line interface, send an AutoSupport test message to a trusted email address that you have access to:

```
system node autosupport invoke -node local -type test
```

11. Find the sequence number of the attempt:

```
system node autosupport history show -node local -destination smtp
```

Find the sequence number for your attempt based on the timestamp. It is probably the most recent attempt.

12. Display the error for your test message attempt:

```
system node autosupport history show -node local -seq-num seq_num -fields error
```

If the error displayed is `Login denied`, your SMTP server is not accepting send requests from the cluster management LIF. If you do not want to change to using HTTPS as your transport protocol, contact your site network administrator to configure the SMTP gateways to address this issue.

If this test succeeds but the same message sent to `mailto:autosupport@netapp.com` does not, ensure that SMTP relay is enabled on all of your SMTP mail hosts, or use HTTPS as a transport protocol.

If even the message to the locally administered email account does not succeed, confirm that your SMTP servers are configured to forward attachments with both of these characteristics:

- The “7z” suffix
- The “application/x-7x-compressed” MIME type.

Troubleshoot the AutoSupport subsystem

The `system node check show` commands can be used to verify and troubleshoot any issues related to the AutoSupport configuration and delivery.

Step

1. Use the following commands to display the status of the AutoSupport subsystem.

Use this command...	To do this...
<code>system node autosupport check show</code>	Display overall status of the AutoSupport subsystem, such as the status of AutoSupport HTTP or HTTPS destination, AutoSupport SMTP destinations, AutoSupport OnDemand Server, and AutoSupport configuration
<code>system node autosupport check show-details</code>	Display detailed status of the AutoSupport subsystem, such as detailed descriptions of errors and the corrective actions

Health monitoring

Monitor the health of your system overview

Health monitors proactively monitor certain critical conditions in your cluster and raise alerts if they detect a fault or risk. If there are active alerts, the system health status reports a degraded status for the cluster. The alerts include the information that you need to respond to degraded system health.

If the status is degraded, you can view details about the problem, including the probable cause and recommended recovery actions. After you resolve the problem, the system health status automatically returns to OK.

The system health status reflects multiple separate health monitors. A degraded status in an individual health monitor causes a degraded status for the overall system health.

For details on how ONTAP supports cluster switches for system health monitoring in your cluster, you can refer to the *Hardware Universe*.

[Supported switches in the Hardware Universe](#)

For details on the causes of Cluster Switch Health Monitor (CSHM) AutoSupport messages, and the necessary actions required to resolve these alerts, you can refer to the Knowledgebase article.

[AutoSupport Message: Health Monitor Process CSHM](#)

How health monitoring works

Individual health monitors have a set of policies that trigger alerts when certain conditions occur. Understanding how health monitoring works can help you respond to problems and control future alerts.

Health monitoring consists of the following components:

- Individual health monitors for specific subsystems, each of which has its own health status

For example, the Storage subsystem has a node connectivity health monitor.

- An overall system health monitor that consolidates the health status of the individual health monitors

A degraded status in any single subsystem results in a degraded status for the entire system. If no subsystems have alerts, the overall system status is OK.

Each health monitor is made up of the following key elements:

- Alerts that the health monitor can potentially raise

Each alert has a definition, which includes details such as the severity of the alert and its probable cause.

- Health policies that identify when each alert is triggered

Each health policy has a rule expression, which is the exact condition or change that triggers the alert.

A health monitor continuously monitors and validates the resources in its subsystem for condition or state changes. When a condition or state change matches a rule expression in a health policy, the health monitor raises an alert. An alert causes the subsystem's health status and the overall system health status to become degraded.

Ways to respond to system health alerts

When a system health alert occurs, you can acknowledge it, learn more about it, repair the underlying condition, and prevent it from occurring again.

When a health monitor raises an alert, you can respond in any of the following ways:

- Get information about the alert, which includes the affected resource, alert severity, probable cause, possible effect, and corrective actions.
- Get detailed information about the alert, such as the time when the alert was raised and whether anyone else has acknowledged the alert already.
- Get health-related information about the state of the affected resource or subsystem, such as a specific shelf or disk.
- Acknowledge the alert to indicate that someone is working on the problem, and identify yourself as the "Acknowledger."
- Resolve the problem by taking the corrective actions provided in the alert, such as fixing cabling to resolve a connectivity problem.
- Delete the alert, if the system did not automatically clear it.
- Suppress an alert to prevent it from affecting the health status of a subsystem.

Suppressing is useful when you understand a problem. After you suppress an alert, it can still occur, but the subsystem health displays as "ok-with-suppressed." when the suppressed alert occurs.

System health alert customization

You can control which alerts a health monitor generates by enabling and disabling the system health policies that define when alerts are triggered. This enables you to customize the health monitoring system for your particular environment.

You can learn the name of a policy either by displaying detailed information about a generated alert or by displaying policy definitions for a specific health monitor, node, or alert ID.

Disabling health policies is different from suppressing alerts. When you suppress an alert, it does not affect the subsystem's health status, but the alert can still occur.

If you disable a policy, the condition or state that is defined in its policy rule expression no longer triggers an alert.

Example of an alert that you want to disable

For example, suppose an alert occurs that is not useful to you. You use the `system health alert show -instance` command to obtain the Policy ID for the alert. You use the policy ID in the `system health policy definition show` command to view information about the policy. After reviewing the rule expression and other information about the policy, you decide to disable the policy. You use the `system health policy definition modify` command to disable the policy.

How health alerts trigger AutoSupport messages and events

System health alerts trigger AutoSupport messages and events in the Event Management System (EMS), enabling you to monitor the health of the system using AutoSupport messages and the EMS in addition to using the health monitoring system directly.

Your system sends an AutoSupport message within five minutes of an alert. The AutoSupport message includes all alerts generated since the previous AutoSupport message, except for alerts that duplicate an alert for the same resource and probable cause within the previous week.


Some alerts do not trigger AutoSupport messages. An alert does not trigger an AutoSupport message if its health policy disables the sending of AutoSupport messages. For example, a health policy might disable AutoSupport messages by default because AutoSupport already generates a message when the problem occurs. You can configure policies to not trigger AutoSupport messages by using the `system health policy definition modify` command.

You can view a list of all of the alert-triggered AutoSupport messages sent in the previous week using the `system health autosupport trigger history show` command.

Alerts also trigger the generation of events to the EMS. An event is generated each time an alert is created and each time an alert is cleared.

Available cluster health monitors

There are several health monitors that monitor different parts of a cluster. Health monitors help you to recover from errors within ONTAP systems by detecting events, sending alerts to you, and deleting events as they clear.

Health monitor name (identifier)	Subsystem name (identifier)	Purpose
Cluster switch(cluster-switch)	Switch (Switch-Health)	<p>Monitors cluster network switches and management network switches for temperature, utilization, interface configuration, redundancy (cluster network switches only), and fan and power supply operation. The cluster switch health monitor communicates with switches through SNMP. SNMPv2c is the default setting.</p> <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;">  <p>Beginning with ONTAP 9.2, this monitor can detect and report when a cluster switch has rebooted since the last polling period.</p> </div>
MetroCluster Fabric	Switch	Monitors the MetroCluster configuration back-end fabric topology and detects misconfigurations such as incorrect cabling and zoning, and ISL failures.
MetroCluster Health	Interconnect, RAID, and storage	Monitors FC-VI adapters, FC initiator adapters, left-behind aggregates and disks, and inter-cluster ports
Node connectivity(node-connect)	CIFS nondisruptive operations (CIFS-NDO)	Monitors SMB connections for nondisruptive operations to Hyper-V applications.
	Storage (SAS-connect)	Monitors shelves, disks, and adapters at the node level for appropriate paths and connections.
System	not applicable	Aggregates information from other health monitors.
System connectivity (system-connect)	Storage (SAS-connect)	Monitors shelves at the cluster level for appropriate paths to two HA clustered nodes.

Receive system health alerts automatically

You can manually view system health alerts by using the `system health alert show` command. However, you should subscribe to specific Event Management System (EMS) messages to automatically receive notifications when a health monitor generates an alert.

About this task

The following procedure shows you how to set up notifications for all `hm.alert.raised` messages and all `hm.alert.cleared` messages.

All `hm.alert.raised` messages and all `hm.alert.cleared` messages include an SNMP trap. The names of the SNMP traps are `HealthMonitorAlertRaised` and `HealthMonitorAlertCleared`. For information about SNMP traps, see the *Network Management Guide*.

Steps

1. Use the `event destination create` command to define the destination to which you want to send the EMS messages.

```
cluster1::> event destination create -name health_alerts -mail
admin@example.com
```

2. Use the `event route add-destinations` command to route the `hm.alert.raised` message and the `hm.alert.cleared` message to a destination.

```
cluster1::> event route add-destinations -messagename hm.alert*
-destinations health_alerts
```

Related information

[Network management](#)

Respond to degraded system health

When your system's health status is degraded, you can show alerts, read about the probable cause and corrective actions, show information about the degraded subsystem, and resolve the problem. Suppressed alerts are also shown so that you can modify them and see whether they have been acknowledged.

About this task

You can discover that an alert was generated by viewing an AutoSupport message or an EMS event, or by using the `system health` commands.

Steps

1. Use the `system health alert show` command to view the alerts that are compromising the system's health.
2. Read the alert's probable cause, possible effect, and corrective actions to determine whether you can resolve the problem or need more information.

3. If you need more information, use the `system health alert show -instance` command to view additional information available for the alert.
4. Use the `system health alert modify` command with the `-acknowledge` parameter to indicate that you are working on a specific alert.
5. Take corrective action to resolve the problem as described by the `Corrective Actions` field in the alert.

The corrective actions might include rebooting the system.

When the problem is resolved, the alert is automatically cleared. If the subsystem has no other alerts, the health of the subsystem changes to `OK`. If the health of all subsystems is `OK`, the overall system health status changes to `OK`.

6. Use the `system health status show` command to confirm that the system health status is `OK`.

If the system health status is not `OK`, repeat this procedure.

Example of responding to degraded system health

By reviewing a specific example of degraded system health caused by a shelf that lacks two paths to a node, you can see what the CLI displays when you respond to an alert.

After starting ONTAP, you check the system health and you discover that the status is degraded:

```
cluster1::>system health status show
Status
-----
degraded
```

You show alerts to find out where the problem is, and see that shelf 2 does not have two paths to node1:

```
cluster1::>system health alert show
      Node: node1
      Resource: Shelf ID 2
      Severity: Major
      Indication Time: Mon Nov 10 16:48:12 2013
      Probable Cause: Disk shelf 2 does not have two paths to controller
                      node1.
      Possible Effect: Access to disk shelf 2 via controller node1 will be
                      lost with a single hardware component failure (e.g.
                      cable, HBA, or IOM failure).
      Corrective Actions: 1. Halt controller node1 and all controllers attached
                          to disk shelf 2.
                          2. Connect disk shelf 2 to controller node1 via two
                          paths following the rules in the Universal SAS and ACP Cabling Guide.
                          3. Reboot the halted controllers.
                          4. Contact support personnel if the alert persists.
```

You display details about the alert to get more information, including the alert ID:

```

cluster1::>system health alert show -monitor node-connect -alert-id
DualPathToDiskShelf_Alert -instance
    Node: node1
    Monitor: node-connect
    Alert ID: DualPathToDiskShelf_Alert
Alerting Resource: 50:05:0c:c1:02:00:0f:02
    Subsystem: SAS-connect
    Indication Time: Mon Mar 21 10:26:38 2011
Perceived Severity: Major
    Probable Cause: Connection_establishment_error
    Description: Disk shelf 2 does not have two paths to controller
node1.
    Corrective Actions: 1. Halt controller node1 and all controllers
attached to disk shelf 2.
                        2. Connect disk shelf 2 to controller node1 via
two paths following the rules in the Universal SAS and ACP Cabling Guide.
                        3. Reboot the halted controllers.
                        4. Contact support personnel if the alert
persists.
    Possible Effect: Access to disk shelf 2 via controller node1 will
be lost with a single
hardware component failure (e.g. cable, HBA, or IOM failure).
    Acknowledge: false
    Suppress: false
    Policy: DualPathToDiskShelf_Policy
Acknowledger: -
Suppressor: -
Additional Information: Shelf uuid: 50:05:0c:c1:02:00:0f:02
                        Shelf id: 2
                        Shelf Name: 4d.shelf2
                        Number of Paths: 1
                        Number of Disks: 6
                        Adapter connected to IOMA:
                        Adapter connected to IOMB: 4d
Alerting Resource Name: Shelf ID 2

```

You acknowledge the alert to indicate that you are working on it.

```

cluster1::>system health alert modify -node node1 -alert-id
DualPathToDiskShelf_Alert -acknowledge true

```

You fix the cabling between shelf 2 and node1, and then reboot the system. Then you check system health again, and see that the status is OK:

```
cluster1::>system health status show
Status
-----
OK
```

Configure discovery of cluster and management network switches

The cluster switch health monitor automatically attempts to discover your cluster and management network switches using the Cisco Discovery Protocol (CDP). You must configure the health monitor if it cannot automatically discover a switch or if you do not want to use CDP for automatic discovery.

About this task

The `system cluster-switch show` command lists the switches that the health monitor discovered. If you do not see a switch that you expected to see in that list, then the health monitor cannot automatically discover it.

Steps

1. If you want to use CDP for automatic discovery, do the following:
 - a. Ensure that the Cisco Discovery Protocol (CDP) is enabled on your switches.

Refer to your switch documentation for instructions.
 - b. Run the following command on each node in the cluster to verify whether CDP is enabled or disabled:

```
run -node node_name -command options cdpd.enable
```

If CDP is enabled, go to step d. If CDP is disabled, go to step c.

- c. Run the following command to enable CDP:

```
run -node node_name -command options cdpd.enable on
```

Wait five minutes before you go to the next step.

- d. Use the `system cluster-switch show` command to verify whether ONTAP can now automatically discover the switches.

2. If the health monitor cannot automatically discover a switch, use the `system cluster-switch create` command to configure discovery of the switch:

```
cluster1::> system cluster-switch create -device switch1 -address
192.0.2.250 -snmp-version SNMPv2c -community cshml! -model NX5020 -type
cluster-network
```

Wait five minutes before you go to the next step.

3. Use the `system cluster-switch show` command to verify that ONTAP can discover the switch for

which you added information.

After you finish

Verify that the health monitor can monitor your switches.

Verify the monitoring of cluster and management network switches

The cluster switch health monitor automatically attempts to monitor the switches that it discovers; however, monitoring might not happen automatically if the switches are not configured correctly. You should verify that the health monitor is properly configured to monitor your switches.

Steps

1. To identify the switches that the cluster switch health monitor discovered, enter the following command:

ONTAP 9.8 and later

```
system switch ethernet show
```

ONTAP 9.7 and earlier

```
system cluster-switch show
```

If the `Model` column displays the value `OTHER`, then ONTAP cannot monitor the switch. ONTAP sets the value to `OTHER` if a switch that it automatically discovers is not supported for health monitoring.



If a switch does not display in the command output, you must configure discovery of the switch.

2. Upgrade to the latest supported switch software and reference the configuration file (RCF) from the NetApp Support Site.

[NetApp Support Downloads page](#)

The community string in the switch's RCF must match the community string that the health monitor is configured to use. By default, the health monitor uses the community string `cshml!`.



At this time, the health monitor only supports SNMPv2.

If you need to change information about a switch that the cluster monitors, you can modify the community string that the health monitor uses by using the following command:

ONTAP 9.8 and later

```
system switch ethernet modify
```

ONTAP 9.7 and earlier

```
system cluster-switch modify
```

3. Verify that the switch's management port is connected to the management network.

This connection is required to perform SNMP queries.

Commands for monitoring the health of your system

You can use the `system health` commands to display information about the health of system resources, to respond to alerts, and to configure future alerts. Using the CLI commands enables you to view in-depth information about how health monitoring is configured. The man pages for the commands contain more information.

Display the status of system health

If you want to...	Use this command...
Display the health status of the system, which reflects the overall status of individual health monitors	<code>system health status show</code>
Display the health status of subsystems for which health monitoring is available	<code>system health subsystem show</code>

Display the status of node connectivity

If you want to...	Use this command...
Display details about connectivity from the node to the storage shelf, including port information, HBA port speed, I/O throughput, and the rate of I/O operations per second	<code>storage shelf show -connectivity</code> Use the <code>-instance</code> parameter to display detailed information about each shelf.
Display information about drives and array LUNs, including the usable space, shelf and bay numbers, and owning node name	<code>storage disk show</code> Use the <code>-instance</code> parameter to display detailed information about each drive.
Display detailed information about storage shelf ports, including port type, speed, and status	<code>storage port show</code> Use the <code>-instance</code> parameter to display detailed information about each adapter.

Manage the discovery of cluster, storage, and management network switches

If you want to...	Use this command.. (ONTAP 9.8 and later)	Use this command.. (ONTAP 9.7 and earlier)
Display the switches that the cluster monitors	<code>system switch ethernet show</code>	<code>system cluster-switch show</code>

If you want to...	Use this command.. (ONTAP 9.8 and later)	Use this command.. (ONTAP 9.7 and earlier)
<p>Display the switches that the cluster currently monitors, including switches that you deleted (shown in the Reason column in the command output), and configuration information that you need for network access to the cluster and management network switches.</p> <p>This command is available at the advanced privilege level.</p>	<pre>system switch ethernet show-all</pre>	<pre>system cluster-switch show-all</pre>
<p>Configure discovery of an undiscovered switch</p>	<pre>system switch ethernet create</pre>	<pre>system cluster-switch create</pre>
<p>Modify information about a switch that the cluster monitors (for example, device name, IP address, SNMP version, and community string)</p>	<pre>system switch ethernet modify</pre>	<pre>system cluster-switch modify</pre>
<p>Disable monitoring of a switch</p>	<pre>system switch ethernet modify -disable-monitoring</pre>	<pre>system cluster-switch modify -disable-monitoring</pre>
<p>Disable discovery and monitoring of a switch and delete switch configuration information</p>	<pre>system switch ethernet delete</pre>	<pre>system cluster-switch delete</pre>
<p>Permanently remove the switch configuration information which is stored in the database (doing so reenables automatic discovery of the switch)</p>	<pre>system switch ethernet delete -force</pre>	<pre>system cluster-switch delete -force</pre>
<p>Enable automatic logging to send with AutoSupport messages.</p>	<pre>system switch ethernet log</pre>	<pre>system cluster-switch log</pre>

Respond to generated alerts



If you want to...	Use this command...
<p>Display information about generated alerts, such as the resource and node where the alert was triggered, and the alert's severity and probable cause</p>	<pre>system health alert show</pre>

If you want to...	Use this command...
Display information about each generated alert	<code>system health alert show -instance</code>
Indicate that someone is working on an alert	<code>system health alert modify</code>
Acknowledge an alert	<code>system health alert modify -acknowledge</code>
Suppress a subsequent alert so that it does not affect the health status of a subsystem	<code>system health alert modify -suppress</code>
Delete an alert that was not automatically cleared	<code>system health alert delete</code>
Display information about the AutoSupport messages that alerts triggered within the last week, for example, to determine whether an alert triggered an AutoSupport message	<code>system health autosupport trigger history show</code>

Configure future alerts

If you want to...	Use this command...
Enable or disable the policy that controls whether a specific resource state raises a specific alert	<code>system health policy definition modify</code>

Display information about how health monitoring is configured

If you want to...	Use this command...
Display information about health monitors, such as their nodes, names, subsystems, and status	<code>system health config show</code> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>Use the <code>-instance</code> parameter to display detailed information about each health monitor.</p> </div>
Display information about the alerts that a health monitor can potentially generate	<code>system health alert definition show</code> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>Use the <code>-instance</code> parameter to display detailed information about each alert definition.</p> </div>

If you want to...	Use this command...
Display information about health monitor policies, which determine when alerts are raised	<pre>system health policy definition show</pre> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <p>Use the <code>-instance</code> parameter to display detailed information about each policy. Use other parameters to filter the list of alerts—for example, by policy status (enabled or not), health monitor, alert, and so on.</p> </div>

Display environmental information

Sensors help you monitor the environmental components of your system. The information you can display about environmental sensors include their type, name, state, value, and threshold warnings.

Step

1. To display information about environmental sensors, use the `system node environment sensors show` command.

File System Analytics

File System Analytics overview

File System Analytics (FSA) was first introduced in ONTAP 9.8 to provide real-time visibility into file usage and storage capacity trends inside ONTAP FlexGroup or FlexVol volumes. This native capability eliminates the need for external tools and provides key insights into how your storage is utilized and whether there are opportunities to optimize the storage for your business needs.


With FSA, you have visibility at all levels of a volume's file system hierarchy in NAS. For example, you can gain usage and capacity insights at the Storage VM (SVM), volume, directory, and file levels. You can use FSA to answer questions like:

- What is filling up my storage, and are there any large files I can move to another storage location?
- Which are my most active volumes, directories, and files? Is my storage performance optimized for the needs of my users?
- How much data was added in the last month?
- Who are my most active or least active storage users?
- How much inactive or dormant data is on my primary storage? Can I move that data to a lower cost cold tier?
- Will my planned quality-of-service changes negatively impact access to critical, frequently accessed files?

File System Analytics is integrated into ONTAP System Manager. Views within System Manager provide:

- Real-time visibility for effective data management and operation

- Real-time data collection and aggregation
- Subdirectory and file sizes and counts, together with associated performance profiles
- File age histograms for modify and access histories



Supported volume types

File System Analytics is designed to provide visibility on volumes with active NAS data, with the exception of FlexCache caches and SnapMirror destination volumes.

File System Analytics feature availability

Each ONTAP release expands the scope of File System Analytics.

	ONTAP 9.14.1	ONTAP 9.13.1	ONTAP 9.12.1	ONTAP 9.11.1	ONTAP 9.10.1	ONTAP 9.9.1	ONTAP 9.8
Visualization in System Manager	✓	✓	✓	✓	✓	✓	✓
Capacity analytics	✓	✓	✓	✓	✓	✓	✓
Inactive data information	✓	✓	✓	✓	✓	✓	✓

	ONTAP 9.14.1	ONTAP 9.13.1	ONTAP 9.12.1	ONTAP 9.11.1	ONTAP 9.10.1	ONTAP 9.9.1	ONTAP 9.8
Support for volumes transitioned from Data ONTAP 7-Mode	✓	✓	✓	✓	✓	✓	
Ability to customize inactive period in System Manager	✓	✓	✓	✓	✓	✓	
Volume-level Activity Tracking	✓	✓	✓	✓	✓		
Download Activity Tracking data to CSV	✓	✓	✓	✓	✓		
SVM-level Activity Tracking	✓	✓	✓	✓			
Timeline	✓	✓	✓	✓			
Usage Analytics	✓	✓	✓				
Option to enable File System Analytics by default	✓	✓					
Initialization scan progress monitor	✓						

Learn more about File System Analytics

ONTAP File System Analytics

Daniel Tennant
Director of Software Engineering
December 13, 2020

© 2020 NetApp, Inc. All rights reserved. — NETAPP CONFIDENTIAL —






Further Reading

- [TR 4687: Best-practice guidelines for ONTAP File System Analytics](#)
- [Knowledge Base: High or fluctuating latency after turning on NetApp ONTAP File System Analytics](#)

Enable File System Analytics

To collect and display usage data such as capacity analytics, you need to enable File System Analytics on a volume.

About this task

- Beginning with ONTAP 9.8, you can enable File System Analytics on a new or existing volume. If you upgrade a system to ONTAP 9.8 or later, ensure that all upgrade processes have completed before you enable File System Analytics.
- Depending on the size and contents of the volume, enabling analytics may take time while ONTAP processes existing data in the volume. System Manager displays progress and presents analytics data when complete. If you need more precise information about initialization progress, you can use the ONTAP CLI command `volume analytics show`.

Beginning in ONTAP 9.14.1, ONTAP provides progress tracking for the initialization scan in addition to notifications about throttling events that affect the scan progress.

For additional considerations related to the initialization scan, see [Scan considerations](#).

Steps

You can enable File System Analytics with ONTAP System Manager or the CLI.

System Manager

In ONTAP 9.8 and 9.9.1	Beginning in ONTAP 9.10.1
<ol style="list-style-type: none">1. Select Storage > Volumes.2. Select the desired volume, then select Explorer.3. Select Enable Analytics or Disable Analytics.	<ol style="list-style-type: none">1. Select Storage > Volumes.2. Select the desired volume. From the individual volume menu, select File System > Explorer.3. Select Enable Analytics or Disable Analytics.

CLI

Enable File System Analytics with the CLI

1. Run the following command:

```
volume analytics on -vserver svm_name -volume volume_name [-foreground {true|false}]
```

By default, the command runs in the foreground; ONTAP displays progress and presents analytics data when complete. If you need more precise information, you can run the command in the background by using the `-foreground false` option and then use the `volume analytics show` command to display initialization progress in the CLI.
2. After successfully enabling File System Analytics, use System Manager or the ONTAP REST API to display the analytic data.


Modify default File System Analytics settings

Beginning in ONTAP 9.13.1, you can modify SVM or clusters settings to enable File System Analytics by default on new volumes.

System Manager

If you are using System Manager, you can modify the storage VM or cluster settings to enable capacity analytics and Activity Tracking at volume creation by default. Default enablement only applies to volumes created after you modify the settings, not existing volumes.

Modify File System Analytics settings on a cluster

1. In System Manager, navigate to **Cluster settings**.
2. In **Cluster settings**, review the File System Settings tab. To modify the settings, select the  icon.
3. In the **Activity Tracking** field, enter the names of the SVMs to enable Activity Tracking for by default. Leaving the field blank will leave Activity Tracking disabled on all SVMs.

Uncheck the **Enable on new storage VMs** box to disable Activity Tracking by default on new storage VMs.

4. In the **Analytics** field, enter the names of the storage VMs you want capacity analytics enabled for by default. Leaving the field blank will leave capacity analytics disabled on all SVMs.

Uncheck the **Enable on new storage VMs** box to disable capacity analytics by default on new storage VMs.

5. Select **Save**.

Modify File System Analytics settings on an SVM

1. Select the SVM you want to modify then **Storage VM settings**.
2. In the **File System Analytics** card, use the toggles to enable or disable Activity Tracking and capacity analytics for all new volumes on the storage VM.

CLI

You can configure the storage VM to enable File System Analytics by default on new volumes using the ONTAP CLI.

Enable File System Analytics by default on an SVM

1. Modify the SVM to enable capacity analytics and Activity Tracking by default on all newly created volumes:

```
vserver modify -vserver svm_name -auto-enable-activity-tracking true -auto-enable-analytics true
```

View file system activity

After File System Analytics (FSA) is enabled, you can view the root directory contents of a selected volume sorted by the space used in each subtree.

Select any file system object to browse the file system and to display detailed information about each object in a directory. Information about directories can also be displayed graphically. Over time, historical data is displayed for each subtree. Space used is not sorted if there are more than 3000 directories.

Explorer

The File System Analytics **Explorer** screen consists of three areas:

- Tree view of directories and subdirectories; expandable list showing name, size, modify history, and access history.
- Files; showing name, size, and accessed time for the object selected in the directory list.
- Active and inactive data comparison for the object selected in the directory list.

Beginning with ONTAP 9.9.1, you can customize the range to be reported. The default value is one year. Based on these customizations, you can take corrective actions, such as moving volumes and modifying the tiering policy.

Accessed time is shown by default. However, if the volume default has been altered from the CLI (by setting the `-atime-update` option to `false` with the `volume modify` command), then only last modified time is shown. For example:

- The tree view will not display the **access history**.
- The files view will be altered.
- The active/inactive data view will be based on modified time (`mtime`).

Using these displays, you can examine the following:

- File system locations consuming the most space
- Detailed information about a directory tree, including file and subdirectory count within directories and subdirectories
- File system locations that contain old data (for example, scratch, temp, or log trees)

Keep the following points in mind when interpreting FSA output:

- FSA show where and when your data is in use, not how much data is being processed. For example, large space consumption by recently accessed or modified files does not necessarily indicate high system processing loads.
- The way that the **Volume Explorer** tab calculates space consumption for FSA might differ from other tools. In particular, there could be significant differences compared to the consumption reported in the **Volume Overview** if the volume has storage efficiency features enabled. This is because the **Volume Explorer** tab does not include efficiency savings.
- Due to space limitations in the directory display, it is not possible to view a directory depth greater than 8 levels in the *List View*. To view directories more than 8 levels deep, you must switch to *Graphical View*, locate the desired directory, then switch back to *List View*. This will allow additional screen space in the display.

Steps

1. View the root directory contents of a selected volume:

In ONTAP 9.8 and 9.9.1	Beginning in ONTAP 9.10.1
Click Storage > Volumes , select the desired volume, then click Explorer .	Select Storage > Volumes , select the desired volume. From the individual volume menu, select File System > Explorer .

Enable Activity Tracking

Beginning with ONTAP 9.10.1, File System Analytics includes an Activity Tracking feature that allows you to identify hot objects and download the data as a CSV file. Beginning with ONTAP 9.11.1, Activity Tracking is expanded to the SVM scope. Also beginning in ONTAP 9.11.1, System Manager features a timeline for Activity Tracking, allowing you to look through up to five minutes of Activity Tracking data.

Activity Tracking enables monitoring in four categories:

- Directories
- Files
- Clients
- Users

For each category monitored, Activity Tracking will display read IOPs, write IOPs, read throughputs, and write throughputs. Queries on Activity Tracking refresh every 10 to 15 seconds pertaining to hot spots seen in the system over the previous five-second interval.

Activity tracking information is approximate, and the accuracy of the data depends on the distribution of the incoming I/O traffic.

When viewing Activity Tracking in System Manager at the volume level, only the menu of the expanded volume will actively refresh. If the view of any volumes are collapsed, they will not refresh until the volume display is expanded. You can stop the refreshes with the **Pause Refresh** button. Activity data can be downloaded in a CSV format that will display all the point-in-time data captured for the selected volume.

With the timeline feature available beginning in ONTAP 9.11.1, you can keep a record of hotspot activity on a volume or SVM, continuously updating approximately every five seconds and retaining the previous five minutes of data. Timeline data is only retained for fields that are visible area of the page. If you collapse a tracking category or scroll so the timeline is out of view, the timeline will stop collecting data. By default, timelines are disabled and will automatically be disabled when you navigate away from the Activity tab.

Enable Activity Tracking for a single volume

You can enable Activity Tracking with ONTAP System Manager or the CLI.

About this task

If you use RBAC with the ONTAP REST API or System Manager, you will need to create custom roles to manage access to Activity Tracking. See [Role-based access control](#) for this process.

System Manager

Steps

1. Select **Storage > Volumes**. Select the desired volume. From the individual volume menu, select File System and then select the Activity tab.
2. Ensure **Activity Tracking** is turned on to view individual reports on top directories, files, clients, and users.
3. To analyze data in greater depth without refreshes, select **Pause Refresh**. You can download the data to have a CSV record of the report as well.

CLI

Steps

1. Enable Activity Tracking:

```
volume activity-tracking on -vserver svm_name -volume volume_name
```

2. Check if the Activity Tracking state for a volume is on or off with the command:

```
volume activity-tracking show -vserver svm_name -volume volume_name -state
```

3. Once enabled, use ONTAP System Manager or the ONTAP REST API to display Activity Tracking data.

Enable Activity Tracking for multiple volumes

You can enable Activity Tracking for multiple volumes with System Manager or the CLI.

About this task

If you use RBAC with the ONTAP REST API or System Manager, you will need to create custom roles to manage access to Activity Tracking. See [Role-based access control](#) for this process.

System Manager

Enable for specific volumes

1. Select **Storage > Volumes**. Select the desired volume. From the individual volume menu, select File System and then select the Activity tab.
2. Select the volumes that you want to enable Activity Tracking on. At the top of the volume list, select the **More Options** button. Select **Enable Activity Tracking**.
3. To view Activity Tracking at the SVM level, select the specific SVM you would like to view from **Storage > Volumes**. Navigate to the File System tab then Activity and you will see data for the volumes that have Activity Tracking enabled.

Enable for all volumes

1. Select **Storage > Volumes**. Select an SVM from the menu.
2. Navigate to the **File System** tab, choose the **More** tab to enable Activity Tracking on all volumes in the SVM.

CLI

Beginning in ONTAP 9.13.1, you can enable Activity Tracking for multiple volumes using the ONTAP CLI.

Steps

1. Enable Activity Tracking:

```
volume activity-tracking on -vserver svm_name -volume [*|!volume_names]
```

Use *** to enable Activity Tracking for all volumes on the specified storage VM.

Use *!* followed by volume names to enable Activity Tracking for all volumes on the SVM except the named volumes.

2. Confirm the operation succeeded:

```
volume show -fields activity-tracking-state
```

3. Once enabled, use ONTAP System Manager or the ONTAP REST API to display Activity Tracking data.

Enable usage analytics

Beginning in ONTAP 9.12.1, you can enable usage analytics to see which directories within a volume are using the most space. You can view the total number of directories in a volume or the total number of files in a volume. Reporting is limited to the 25 directories that use the most space.

Analytics for large directories refresh every 15 minutes. You can monitor the most recent refresh by checking the Last Refreshed timestamp at the top of the page. You can also click the Download button to download data to an Excel workbook. The download operation runs in the background and presents the most recently reported information for the selected volume. If the scan returns without any results, ensure the volume is online. Events such as SnapRestore will cause File System Analytics to rebuild its list of large directories.

Steps

1. Select **Storage > Volumes**. Select the desired volume.
2. From the individual volume menu, select **File System**. Then select the **Usage** tab.
3. Toggle the **Analytics** switch to enable usage analytics.
4. System Manager will display a bar graph identifying the directories with the largest size in descending order.



ONTAP might display partial data or no data at all while the list of top directories is being collected. The progress of the scan can be in the **Usage** tab that displays during the scan.

To gain more insights into a specific directory, you can [view activity on a file system](#).

Take corrective action based on analytics

Beginning with ONTAP 9.9.1, you can take corrective actions based on current data and desired outcomes directly from the File System Analytics displays.

Delete directories and files

In the Explorer display, you can select directories or individual files to delete. Directories are deleted with low-latency fast directory delete functionality. (Fast directory delete is also available beginning in ONTAP 9.9.1 without analytics enabled.)

Steps

1. Click **Storage > Volumes**, then click **Explorer**.

When you hover over a file or folder, the option to delete appears. You can only delete one object at a time.



When directories and files are deleted, the new storage capacity values are not displayed immediately.

Assign media cost in storage tiers to compare costs of inactive data storage locations

Media cost is a value that you assign based on your evaluation of storage costs, represented as your choice of currency per GB. When set, System Manager uses the assigned media cost to project estimated savings when you move volumes.

The media cost you set is not persistent; it can only be set for a single browser session.

Steps

1. Click **Storage > Tiers**, then click **Set Media Cost** in the desired local tier (aggregate) tiles.

Be sure to select active and inactive tiers to enable comparison.

2. Enter a currency type and amount.


When you enter or change the media cost, the change is made in all media types.

Move volumes to reduce storage costs

Based on analytics displays and media cost comparisons, you can move volumes to less expensive storage in local tiers.

Only one volume at a time can be compared and moved.

Steps

1. After enabling media cost display, click **Storage > Tiers**, then click **Volumes**.
2. To compare destination options for a volume, click  for the volume, then click **Move**.
3. In the **Select Destination Local Tier** display, select destination tiers to display the estimated cost difference.
4. After comparing options, select the desired tier and click **Move**.

Role-based access control with File System Analytics

Beginning in ONTAP 9.12.1, ONTAP includes a predefined role-based access control (RBAC) role called `admin-no-fsa`. The `admin-no-fsa` role grants administrator-level privileges but prevents the user from performing operations related to the `files` endpoint (i.e. File System Analytics) in the ONTAP CLI, REST API, and in System Manager.



For more information on the `admin-no-fsa` role, refer to [Predefined roles for cluster administrators](#).

If you are using a version of ONTAP released prior to ONTAP 9.12.1, you will need to create a dedicated role to control access to File System Analytics. In versions of ONTAP prior to ONTAP 9.12.1, you must configure RBAC permissions through the ONTAP CLI or ONTAP REST API.

System Manager

Beginning in ONTAP 9.12.1, you can configure RBAC permissions for File System Analytics using System Manager.

Steps

1. Select **Cluster > Settings**. Under **Security**, navigate to **Users and Roles** and select .
2. Under **Roles**, select  **Add**.
3. Provide a name for the role. Under Role Attributes, configure the access or restrictions for the user role by providing the appropriate [API endpoints](#). See the table below for primary paths and secondary paths to configure File System Analytics access or restrictions.

Restriction	Primary Path	Secondary Path
Activity Tracking on volumes	/api/storage/volumes	<ul style="list-style-type: none">• /:uuid/top-metrics/directories• /:uuid/top-metrics/files• /:uuid/top-metrics/clients• /:uuid/top-metrics/users
Activity Tracking on SVMs	/api/svm/svms	<ul style="list-style-type: none">• /:uuid/top-metrics/directories• /:uuid/top-metrics/files• /:uuid/top-metrics/clients• /:uuid/top-metrics/users
All File System Analytics operations	/api/storage/volumes	/:uuid/files

You can use `/*` instead of an UUID to set the policy for all volumes or SVMs at the endpoint.

Choose the access privileges for each endpoint.

4. Select **Save**.
5. To assign the role to a user or users, see [Control administrator access](#).

CLI

If you are using a version of ONTAP released prior to ONTAP 9.12.1, use the ONTAP CLI to create a custom-role.

Steps

1. Create a default role to have access to all features.

This needs to be done before creating the restrictive role to ensure the role is only restrictive on the Activity Tracking:

```
security login role create -cmddirname DEFAULT -access all -role storageAdmin
```

2. Create the restrictive role:

```
security login role create -cmddirname "volume file show-disk-usage" -access none -role storageAdmin
```

3. Authorize roles to access the SVM's web services:

- `rest` for REST API calls
- `security` for password protection
- `sysmgr` for System Manager access

```
vserver services web access create -vserver svm-name -name_ -name rest -role storageAdmin
```

```
vserver services web access create -vserver svm-name -name security -role storageAdmin
```

```
vserver services web access create -vserver svm-name -name sysmgr -role storageAdmin
```

4. Create a user.

You must issue a distinct create command for each application you would like to apply to the user. Calling create multiple times on the same user simply applies all the applications to that one user and does not create a new user each time. The `http` parameter for application type applies for the ONTAP REST API and System Manager.

```
security login create -user-or-group-name storageUser -authentication -method password -application http -role storageAdmin
```

5. With the new user credentials, you can now log in to System Manager or use the ONTAP REST API to access File Systems Analytics data.

More information

- [Predefined roles for cluster administrators](#)
- [Control administrator access with System Manager](#)
- [Learn more about RBAC roles and the ONTAP REST API](#)

Considerations for File System Analytics

You should be aware of certain usage limits and potential performance impacts associated with implementing File System Analytics.

SVM-protected relationships

If you have enabled File System Analytics on volumes whose containing SVM is in a protection relationship, the analytics data is not replicated to the destination SVM. If the source SVM must be resynchronized in a recovery operation, you must manually reenables analytics on desired volumes after recovery.

Performance considerations

In some cases, enabling File System Analytics could negatively impact performance during the initial metadata collection. This is most typically seen on systems that are at maximum utilization. To avoid enabling analytics on such systems, you can use ONTAP System Manager performance monitoring tools.

If you experience a notable increase in latency, refer to the Knowledge Base article [High or fluctuating latency after turning on NetApp ONTAP File System Analytics](#).

Scan considerations

When you enable capacity analytics, ONTAP conducts an initialization scan for capacity analytics. The scan accesses metadata for all files in volumes for which capacity analytics is enabled. No file data is read during the scan. Beginning in ONTAP 9.14.1, you can track the progress of the scan with the REST API, in the **Explorer** tab of System Manager, or with the `volume analytics show` CLI command. If there is a throttling event, ONTAP provides a notification.

After the scan completes, File System Analytics is continuously updated in real time as the filesystem changes without the need to run the scan again.

The time required for the scan is proportional to the number of directories and files on the volume. Because the scan collects metadata, file size does not impact the scan time.

For more information about the initialization scan, see [TR-4867: Best practice guidelines for File System Analytics](#).

Best practices

You should start the scan on volumes that do not share aggregates. You can see which aggregates are currently hosting which volumes using the command:

```
volume show -volume comma-separated-list_of_volumes -fields aggr-list
```

While the scan runs, volumes continue to serve client traffic. It's recommended you start the scan during periods where you anticipate lower client traffic.

If client traffic increases, it will consume system resources and cause the scan to take longer.

Beginning in ONTAP 9.12.1, you can pause data collection in System Manager and with the ONTAP CLI.

- If you are using the ONTAP CLI:
 - You can pause data collection with the command: `volume analytics initialization pause -vserver svm_name -volume volume_name`
 - Once client traffic has slowed, you can resume data collection with the command: `volume analytics initialization resume -vserver svm_name -volume volume_name`
- If you are using System Manager, in the **Explorer** view of the volume menu, you use the **Pause Data Collection** and **Resume Data Collection** buttons to manage the scan.

EMS configuration

EMS configuration overview

You can configure ONTAP 9 to send important EMS (Event Management System) event notifications directly to an email address, syslog server, Simple Management Network Protocol (SNMP) trap host, or webhook application so that you are immediately notified of system issues that require prompt attention.

Because important event notifications are not enabled by default, you need to configure the EMS to send notifications to either an email address, a syslog server, an SNMP trap host, or webhook application.

Review release-specific versions of the [ONTAP 9 EMS Reference](#).

If your EMS event mapping uses deprecated ONTAP command sets (such as event destination, event route), it's recommended that you update your mapping. [Learn how to update your EMS mapping from deprecated ONTAP commands](#).

Configure EMS event notifications and filters with System Manager

You can use System Manager to configure how the Event Management System (EMS) delivers event notifications so that you can be notified of system issues that require your prompt attention.

ONTAP version	With System Manager, you can...
ONTAP 9.12.1 and later	Specify Transport Layer Security (TLS) protocol when sending events to remote syslog servers.
ONTAP 9.10.1 and later	Configure email addresses, syslog servers, and webhook applications, as well as SNMP trap hosts.
ONTAP 9.7 to 9.10.0	Configure only SNMP trap hosts. You can configure other EMS destination with the ONTAP CLI. See EMS configuration overview .

You can perform the following procedures:

- [Add an EMS event notification destination](#)
- [Create a new EMS event notification filter](#)
- [Edit an EMS event notification destination](#)
- [Edit an EMS event notification filter](#)
- [Delete an EMS event notification destination](#)
- [Delete an EMS event notification filter](#)

Related information


- [ONTAP EMS Reference](#)
- [Using the CLI to configure SNMP trap hosts to receive event notifications](#)

Add an EMS event notification destination

You can use System Manager to specify to where you want EMS messages sent.

Beginning with ONTAP 9.12.1, EMS events can be sent to a designated port on a remote syslog server via the Transport Layer Security (TLS) protocol. For details, see the [event notification destination create man page](#).

Steps

1. Click **Cluster > Settings**.
2. In the **Notifications Management** section, click , then click **View Event Destinations**.
3. On the **Notification Management** page, select the **Events Destinations** tab.
4. Click **+ Add**.
5. Specify a name, an EMS destination type, and filters.



If needed, you can add a new filter. Click **Add a New Event Filter**.

6. Depending on the EMS destination type you selected, specify the following:


To configure...	Specify or select...
SNMP traphost	<ul style="list-style-type: none">• Traphost name
Email (Beginning with 9.10.1)	<ul style="list-style-type: none">• Destination email address• Mail server• From email address
Syslog server (Beginning with 9.10.1)	<ul style="list-style-type: none">• Host name or IP address of the server• Syslog port (beginning with 9.12.1)• Syslog transport (beginning with 9.12.1) <p>Selecting TCP Encrypted enables the Transport Layer Security (TLS) protocol. If no value is entered for Syslog port, a default is used based on the Syslog transport selection.</p>
Webhook (Beginning with 9.10.1)	<ul style="list-style-type: none">• Webhook URL• Client authentication (select this option to specify a client certificate)

Create a new EMS event notification filter

Beginning with ONTAP 9.10.1, you can use System Manager to define new customized filters that specify the rules for handling EMS notifications.

Steps

1. Click **Cluster > Settings**.



2. In the **Notifications Management** section, click , then click **View Event Destinations**.
3. On the **Notification Management** page, select the **Event Filters** tab.
4. Click **+ Add**.
5. Specify a name, and select whether you want to copy rules from an existing event filter or add new rules.
6. Depending on your choice, perform the following steps:

If you choose....	Then, perform these steps...
Copy rules from existing event filter	<ol style="list-style-type: none"> 1. Select an existing event filter. 2. Modify the existing rules. 3. Add other rules, if needed, by clicking + Add.
Add new rules	Specify the type, name pattern, severities, and SNMP trap type for each new rule.

Edit an EMS event notification destination

Beginning with ONTAP 9.10.1, you can use System Manager to change the event notification destination information.

Steps

1. Click **Cluster > Settings**.
2. In the **Notifications Management** section, click , then click **View Event Destinations**.
3. On the **Notifications Management** page, select the **Events Destinations** tab.
4. Next to the name of the event destination, click , then click **Edit**.
5. Modify the event destination information, then click **Save**.



Edit an EMS event notification filter

Beginning with ONTAP 9.10.1, you can use System Manager to modify customized filters to change how event notifications are handled.



You cannot modify system-defined filters.

Steps

1. Click **Cluster > Settings**.
2. In the **Notifications Management** section, click , then click **View Event Destinations**.
3. On the **Notification Management** page, select the **Event Filters** tab.
4. Next to the name of the event filter, click , then click **Edit**.
5. Modify the event filter information, then click **Save**.



Delete an EMS event notification destination

Beginning with ONTAP 9.10.1, you can use System Manager to delete an EMS event notification destination.



You cannot delete SNMP destinations.

Steps

1. Click **Cluster > Settings**.
2. In the **Notifications Management** section, click , then click **View Event Destinations**.
3. On the **Notification Management** page, select the **Events Destinations** tab.
4. Next to the name of the event destination, click , then click **Delete**.



Delete an EMS event notification filter

Beginning with ONTAP 9.10.1, you can use System Manager to delete customized filters.



You cannot delete system-defined filters.

Steps

1. Click **Cluster > Settings**.
2. In the **Notifications Management** section, click , then click **View Event Destinations**.
3. On the **Notification Management** page, select the **Event Filters** tab.
4. Next to the name of the event filter, click , then click **Delete**.

Configure EMS event notifications with the CLI

EMS configuration workflow


You must configure important EMS event notifications to be sent either as email, forwarded to a syslog server, forwarded to an SNMP traphost, or forwarded to a webhook application. This helps you to avoid system disruptions by taking corrective actions in a timely manner.

About this task

If your environment already contains a syslog server for aggregating the logged events from other systems, such as servers and applications, then it is easier to use that syslog server also for important event notifications from storage systems.

If your environment does not already contain a syslog server, then it is easier to use email for important event notifications.

If you already forward event notifications to an SNMP traphost, then you might want to monitor that traphost for important events.



Choices

- Set EMS to send event notifications.

If you want...	Refer to this...
The EMS to send important event notifications to an email address	Configure important EMS events to send email notifications
The EMS to forward important event notifications to a syslog server	Configure important EMS events to forward notifications to a syslog server
If you want the EMS to forward event notifications to an SNMP traphost	Configure SNMP traphosts to receive event notifications
If you want the EMS to forward event notifications to a webhook application	Configure important EMS events to forward notifications to a webhook application

Configure important EMS events to send email notifications

To receive email notifications of the most important events, you must configure the EMS to send email messages for events that signal important activity.

What you'll need

DNS must be configured on the cluster to resolve the email addresses.

About this task

You can perform this task any time the cluster is running by entering the commands on the ONTAP command line.

Steps

1. Configure the event SMTP mail server settings:

```
event config modify -mail-server mailhost.your_domain -mail-from
cluster_admin@your_domain
```

2. Create an email destination for event notifications:

```
event notification destination create -name storage-admins -email
your_email@your_domain
```

3. Configure the important events to send email notifications:

```
event notification create -filter-name important-events -destinations storage-
admins
```

Configuring important EMS events to forward notifications to a syslog server

To log notifications of the most severe events on a syslog server, you must configure the EMS to forward notifications for events that signal important activity.

What you'll need

DNS must be configured on the cluster to resolve the syslog server name.

About this task

If your environment does not already contain a syslog server for event notifications, you must first create one. If your environment already contains a syslog server for logging events from other systems, then you might want to use that one for important event notifications.

You can perform this task any time the cluster is running by entering the commands on the ONTAP CLI.

Beginning with ONTAP 9.12.1, EMS events can be sent to a designated port on a remote syslog server via the Transport Layer Security (TLS) protocol. Two new parameters are available:

tcp-encrypted

When `tcp-encrypted` is specified for the `syslog-transport`, ONTAP verifies the identity of the destination host by validating its certificate. The default value is `udp-unencrypted`.

syslog-port

The default value `syslog-port` parameter depends on the setting for the `syslog-transport` parameter. If `syslog-transport` is set to `tcp-encrypted`, `syslog-port` has the default value 6514.

For details, see the `event notification destination create` man page.

Steps

1. Create a syslog server destination for important events:

```
event notification destination create -name syslog-ems -syslog syslog-server-address -syslog-transport {udp-unencrypted|tcp-unencrypted|tcp-encrypted}
```

Beginning with ONTAP 9.12.1, the following values can be specified for `syslog-transport`:

- `udp-unencrypted` - User Datagram Protocol with no security
- `tcp-unencrypted` - Transmission Control Protocol with no security
- `tcp-encrypted` - Transmission Control Protocol with Transport Layer Security (TLS)

The default protocol is `udp-unencrypted`.

2. Configure the important events to forward notifications to the syslog server:

```
event notification create -filter-name important-events -destinations syslog-ems
```

Configure SNMP traphosts to receive event notifications

To receive event notifications on an SNMP traphost, you must configure a traphost.

What you'll need

- SNMP and SNMP traps must be enabled on the cluster.



SNMP and SNMP traps are enabled by default.

- DNS must be configured on the cluster to resolve the traphost names.

About this task

If you do not already have an SNMP traphost configured to receive event notifications (SNMP traps), you must add one.

You can perform this task any time the cluster is running by entering the commands on the ONTAP command line.

Step

1. If your environment does not already have an SNMP traphost configured to receive event notifications, add one:

```
system snmp traphost add -peer-address snmp_traphost_name
```

All event notifications that are supported by SNMP by default are forwarded to the SNMP traphost.

Configure important EMS events to forward notifications to a webhook application

You can configure ONTAP to forward important event notifications to a webhook application. The configuration steps needed depend on the level of security you choose.

Prepare to configure EMS event forwarding

There are several concepts and requirements you should consider before configuring ONTAP to forward event notifications to a webhook application.

Webhook application

You need a webhook application capable of receiving the ONTAP event notifications. A webhook is a user-defined callback routine that extends the capability of the remote application or server where it runs. Webhooks are called or activated by the client (in this case ONTAP) by sending an HTTP request to the destination URL. Specifically, ONTAP sends an HTTP POST request to the server hosting the webhook application along with the event notification details formatted in XML.

Security options

There are several security options available depending on how the Transport Layer Security (TLS) protocol is used. The option you choose determines the required ONTAP configuration.



TLS is a cryptographic protocol that is widely used on the internet. It provides privacy as well as data integrity and authentication using one or more public key certificates. The certificates are issued by trusted certificate authorities.

HTTP

You can use HTTP to transport the event notifications. With this configuration, the connection is not secure. The identities of the ONTAP client and webhook application are not verified. Further, the network traffic is not encrypted or protected. See [Configure a webhook destination to use HTTP](#) for the configuration details.

HTTPS

For additional security, you can install a certificate at the server hosting the webhook routine. The HTTPS protocol is used by ONTAP to verify the identity of the webhook application server as well as by both parties to ensure the privacy and integrity of the network traffic. See [Configure a webhook destination to use HTTPS](#) for the configuration details.

HTTPS with mutual authentication

You can further enhance the HTTPS security by installing a client certificate at the ONTAP system issuing the webhook requests. In addition to ONTAP verifying the identity of the webhook application server and protecting the network traffic, the webhook application verifies the identity of the ONTAP client. This two-way peer authentication is known as *Mutual TLS*. See [Configure a webhook destination to use HTTPS with mutual authentication](#) for the configuration details.

Related information

- [The Transport Layer Security \(TLS\) Protocol Version 1.3](#)

Configure a webhook destination to use HTTP

You can configure ONTAP to forward event notifications to a webhook application using HTTP. This is the least secure option but the simplest to set up.

Steps

1. Create a new destination `restapi-ems` to receive the events:

```
event notification destination create -name restapi-ems -rest-api-url
http://<webhook-application>
```

In the above command, you must use the **HTTP** scheme for the destination.

2. Create a notification linking the `important-events` filter with the `restapi-ems` destination:

```
event notification create -filter-name important-events -destinations restapi-ems
```

Configure a webhook destination to use HTTPS

You can configure ONTAP to forward event notifications to a webhook application using HTTPS. ONTAP uses the server certificate to confirm the identity of the webhook application as well as secure the network traffic.

Before you begin

- Generate a private key and certificate for the webhook application server
- Have the root certificate available to install in ONTAP

Steps

1. Install the appropriate server private key and certificates at the server hosting your webhook application. The specific configuration steps are dependent on the server.
2. Install the server root certificate in ONTAP:

```
security certificate install -type server-ca
```

The command will ask for the certificate.

3. Create the `restapi-ems` destination to receive the events:

```
event notification destination create -name restapi-ems -rest-api-url https://<webhook-application>
```

In the above command, you must use the **HTTPS** scheme for the destination.

4. Create the notification that links the `important-events` filter with the new `restapi-ems` destination:

```
event notification create -filter-name important-events -destinations restapi-ems
```

Configure a webhook destination to use HTTPS with mutual authentication

You can configure ONTAP to forward event notifications to a webhook application using HTTPS with mutual authentication. With this configuration there are two certificates. ONTAP uses the server certificate to confirm the identity of the webhook application and secure the network traffic. In addition, the application hosting the webhook uses the client certificate to confirm the identity of the ONTAP client.

Before you begin

You must do the following before configuring ONTAP:

- Generate a private key and certificate for the webhook application server
- Have the root certificate available to install in ONTAP
- Generate a private key and certificate for the ONTAP client

Steps

1. Perform the first two steps in the task [Configure a webhook destination to use HTTPS](#) to install the server certificate so that ONTAP can verify the identity of the server.

2. Install the appropriate root and intermediate certificates at the webhook application to validate the client certificate.

3. Install the client certificate in ONTAP:

```
security certificate install -type client
```

The command will ask for the private key and certificate.

4. Create the `restapi-ems` destination to receive the events:

```
event notification destination create -name restapi-ems -rest-api-url  
https://<webhook-application> -certificate-authority <issuer of the client  
certificate> -certificate-serial <serial of the client certificate>
```

In the above command, you must use the **HTTPS** scheme for destination.

5. Create the notification that links the `important-events` filter with the new `restapi-ems` destination:

```
event notification create -filter-name important-events -destinations restapi-  
ems
```

Update deprecated EMS event mapping

EMS event mapping models

Prior to ONTAP 9.0, EMS events could only be mapped to event destinations based on event name pattern matching. The ONTAP command sets (`event destination`, `event route`) that use this model continue to be available in the latest versions of ONTAP, but they have been deprecated starting with ONTAP 9.0.

Beginning with ONTAP 9.0, the best practice for ONTAP EMS event destination mapping is to use the more scalable event filter model in which pattern matching is done on multiple fields, using the `event filter`, `event notification`, and `event notification destination` command sets.

If your EMS mapping is configured using the deprecated commands, you should update your mapping to use the `event filter`, `event notification`, and `event notification destination` command sets.

There are two types of event destinations:

1. **System-generated destinations:** There are five system-generated event destinations (created by default)

- `allevents`
- `asup`
- `criticals`
- `pager`
- `traphost`

Some of the system-generated destinations are for special purpose. For example, the `asup` destination routes `callhome.*` events to the AutoSupport module in ONTAP to generate AutoSupport messages.

2. **User-created destinations:** These are manually created using the event destination create command.

```
cluster-1::event*> destination show
```

Name	Mail Dest.	SNMP Dest.	Syslog Dest.	Hide
allevents	-	-	-	
false				
asup	-	-	-	
false				
criticals	-	-	-	
false				
pager	-	-	-	
false				
traphost	-	-	-	
false				

5 entries were displayed.

+

```
cluster-1::event*> destination create -name test -mail test@xyz.com
```

This command is deprecated. Use the "event filter", "event notification destination" and "event notification" commands, instead.

+

```
cluster-1::event*> destination show
```

+

Name	Mail Dest.	SNMP Dest.	Syslog Dest.	Hide
allevents	-	-	-	
false				
asup	-	-	-	
false				
criticals	-	-	-	
false				
pager	-	-	-	
false				
test	test@xyz.com	-	-	
false				
traphost	-	-	-	
false				

6 entries were displayed.

In the deprecated model, EMS events are individually mapped to a destination using the `event route add-destinations` command.

```
cluster-1::event*> route add-destinations -message-name raid.aggr.*
-destinations test
This command is deprecated. Use the "event filter", "event notification
destination" and "event notification" commands, instead.
4 entries were acted on.

cluster-1::event*> route show -message-name raid.aggr.*
```

Time	Message	Severity	Destinations	Freq	Threshd
	raid.aggr.autoGrow.abort	NOTICE	test	0	0
	raid.aggr.autoGrow.success	NOTICE	test	0	0
	raid.aggr.lock.conflict	INFORMATIONAL	test	0	0
	raid.aggr.log.CP.count	DEBUG	test	0	0

4 entries were displayed.


The new, more scalable EMS event notifications mechanism is based on event filters and event notification destinations. Refer to the following KB article for detailed information on the new event notification mechanism:

- [Overview of Event Management System for ONTAP 9](#)

Legacy routing based model



Event notification based model



Update EMS event mapping from deprecated ONTAP commands

If your EMS event mapping is currently configured using the deprecated ONTAP command sets (`event destination`, `event route`), you should follow this procedure to update your mapping to use the `event filter`, `event notification`, and `event notification destination` command sets.

Steps

1. List all the event destinations in the system using the `event destination show` command.

```
cluster-1::event*> destination show
```

```
Hide
Name          Mail Dest.      SNMP Dest.      Syslog Dest.
Params
-----
-----
allevents     -              -              -
false
asup          -              -              -
false
criticals     -              -              -
false
pager         -              -              -
false
test          test@xyz.com    -              -
false
traphost      -              -              -
false
6 entries were displayed.
```

2. For each destination, list the events being mapped to it using the `event route show -destinations <destination name>` command.

```
cluster-1::event*> route show -destinations test
```

```
Time
Message          Severity      Destinations    Freq
Threshd          Threshd
-----
-----
raid.aggr.autoGrow.abort      NOTICE      test            0      0
raid.aggr.autoGrow.success    NOTICE      test            0      0
raid.aggr.lock.conflict       INFORMATIONAL test            0      0
raid.aggr.log.CP.count        DEBUG        test            0      0
4 entries were displayed.
```

3. Create a corresponding `event filter` which includes all these subsets of events. For example, if you want to include only the `raid.aggr.*` events, use a wildcard for the `message-name` parameter when creating the filter. You can also create filters for single events.



You can create up to 50 event filters.

```

cluster-1::event*> filter create -filter-name test_events

cluster-1::event*> filter rule add -filter-name test_events -type
include -message-name raid.aggr.*

cluster-1::event*> filter show -filter-name test_events
Filter Name Rule      Rule      Message Name      SNMP Trap Type
Severity
      Position Type
-----
test_events
      1      include  raid.aggr.*      *      *
      2      exclude *      *      *
2 entries were displayed.

```

4. Create an event notification destination for each of the event destination endpoints (i.e., SMTP/SNMP/syslog)

```

cluster-1::event*> notification destination create -name dest1 -email
test@xyz.com

cluster-1::event*> notification destination show
Name      Type      Destination
-----
dest1      email      test@xyz.com (via "localhost" from
"admin@localhost", configured in "event config")
snmp-traphost  snmp      - (from "system snmp traphost")
2 entries were displayed.

```

5. Create an event notification by mapping the event filter to the event notification destination.

```

cluster-1::event*> notification create -filter-name asup_events
-destinations dest1

cluster-1::event*> notification show
ID  Filter Name      Destinations
----
1   default-trap-events  snmp-traphost
2   asup_events         dest1
2 entries were displayed.

```

6. Repeat steps 1-5 for each event destination that has an event route mapping.



Events routed to SNMP destinations should be mapped to the `snmp-traphost` event notification destination. The SNMP traphost destination uses the system configured SNMP traphost.

```
cluster-1::event*> system snmp traphost add 10.234.166.135

cluster-1::event*> system snmp traphost show
      scspr2410142014.gdl.englab.netapp.com
(scspr2410142014.gdl.englab.netapp.com) <10.234.166.135>   Community:
public

cluster-1::event*> notification destination show -name snmp-traphost

      Destination Name: snmp-traphost
      Type of Destination: snmp
      Destination: 10.234.166.135 (from "system snmp
traphost")
      Server CA Certificates Present?: -
      Client Certificate Issuing CA: -
Client Certificate Serial Number: -
      Client Certificate Valid?: -
```

Copyright information

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

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

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

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

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

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

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

Trademark information

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