Monitor storage health

NetApp Console local deployment

NetApp
August 10, 2026

Table of Contents

Monitor storage health . . . 1
- Storage monitoring in NetApp Console local deployment . . . 1
  - Monitor fleet-wide health with dashboards . . . 1
  - Investigate performance issues . . . 1
  - Configure alerts and notifications . . . 1
  - Remediate issues . . . 1
- Monitor with dashboards . . . 2
  - View Home page metrics in NetApp Console local deployment . . . 2
  - Use custom dashboards . . . 3
  - Manage custom dashboards . . . 11
- Monitor fleets using inventory in NetApp Console local deployment . . . 13
  - Access Inventory . . . 13
  - Inventory views . . . 13
  - Investigate storage resources . . . 13
  - Provision storage resources . . . 14
  - Inventory and alerts . . . 14
  - Advanced management tasks . . . 14
- Remediate alerts . . . 14
  - Learn about alerts in NetApp Console local deployment . . . 14
  - Monitor and investigate alerts in NetApp Console local deployment . . . 16
  - Configure notification rules for alerts in NetApp Console local deployment . . . 20
  - Remediate storage class drift in NetApp Console local deployment . . . 23
  - Alert remediation actions in NetApp Console local deployment . . . 24
- Investigate performance issues . . . 26
  - Learn about NetApp Console local deployment Workload Analyzer . . . 26
  - Investigate high latency on a volume in NetApp Console local deployment . . . 27
  - Investigate high throughput on a volume in NetApp Console local deployment . . . 29
  - Investigate capacity growth on a volume in NetApp Console local deployment . . . 30
  - Workload Analyzer metrics in NetApp Console local deployment . . . 31

Monitor storage health

Storage monitoring in NetApp Console local deployment

NetApp Console local deployment helps you monitor storage across fleets with dashboards, alerts, deep-dive analysis, and remediation so you can spot and resolve issues before they affect workloads.

Monitor fleet-wide health with dashboards

Start with dashboards for a quick view of storage health and performance. Use the Home page for at-a-glance metrics, open the dashboards table for predefined and custom dashboards, and move into alerts, investigation, and remediation when you need more detail.

- View Home page metrics in NetApp Console local deployment
- Keep dashboards current in NetApp Console local deployment
- Learn about custom dashboards in NetApp Console local deployment

Investigate performance issues

Use the Workload Analyzer to investigate performance issues on individual volumes. It correlates latency, throughput, IOPS, and configuration changes so you can pinpoint root causes.

- Learn about the NetApp Console local deployment Workload Analyzer
- Investigate high latency on a volume in NetApp Console local deployment
- Investigate high throughput demand on a volume in NetApp Console local deployment
- Learn about Workload Analyzer metrics in NetApp Console local deployment

Configure alerts and notifications

Configure alert policies and notification channels so the right people hear about storage conditions that require attention. For example, route a capacity warning to the storage team and a protection alert to the backup admin who can act on it.

Configure notifications and alert policies in NetApp Console local deployment.
View storage alerts in NetApp Console local deployment.

Remediate issues

When you detect an issue, remediate it directly from NetApp Console local deployment:

- **Automated remediation** — NetApp Console local deployment applies corrective actions automatically based on predefined rules.
- **Guided remediation** — Follow step-by-step recommendations to resolve issues with full control over each action.

Monitor with dashboards

View Home page metrics in NetApp Console local deployment

Use the Home page dashboard in NetApp Console local deployment as your default starting point for monitoring storage health and performance. It provides high-level metrics for fleet health, alerts, security, capacity usage, latency, throughput, and IOPS.

The Home page dashboard is automatically scoped to only the fleets and systems you are authorized to view. You can also add a custom dashboard to the Home page for role-specific monitoring.

Use the cards as a layered triage workflow. Start with **Fleet health** and **Alerts** to find risk hotspots. Use **System health status** and **Recommended remediations** to identify impacted resources. Use **Capacity usage**, **Latency**, **Throughput**, and **IOPS** to investigate root causes.

Fleet health

The **Fleet health** card shows a high-level status summary for the top five fleets, including system counts, used capacity, security compliance, and critical alerts. Use this card to identify which fleets need immediate attention. Select the fleet name to view a fleet-only dashboard for the selected fleet.

Recommended remediations

The **Recommended remediations** card lists active issues and suggested actions. The card prioritizes remediations based on capacity pressure, offline resources, and protection-related risks.

Alerts

The **Alerts** card summarizes alert volume in the selected time range and groups it by severity. Use this card to understand current incident load and to spot recent changes in critical, warning, or informational alerts.

Security

The **Security** card summarizes compliance and protection metrics, such as system compliance and ransomware-related controls. Use this view to monitor security trends across your resources.

Capacity usage

The **Capacity usage** card displays projected and historical usage trends across selected fleets or systems. Use this card to identify growth patterns and plan for additional capacity.

Latency

The **Latency** card tracks response-time behavior over time across the selected systems. Use this trend to detect emerging performance delays and compare relative latency across resources.

Throughput

The **Throughput** card shows data transfer rates over time for selected systems. Use this card to understand workload intensity and monitor throughput demand changes.

IOPS

The **IOPS** card shows input/output operation rates over time. Use this card to compare activity levels between systems and identify sustained or intermittent I/O demand.

Dashboard card (example metric card)

The lower **Dashboard** area can contain a user-selected custom dashboard such as average latency for a selected production scope. Use these cards for focused monitoring by team, workload, or goal.

System health status

The **System health status** card lists individual systems with their current health state. Use this card to quickly identify unhealthy systems and correlate status changes with alerts and performance signals.

Related information

- Learn about custom dashboards in NetApp Console local deployment
- Learn about custom dashboard cards and metrics in NetApp Console local deployment

Use custom dashboards

Learn about custom dashboards in NetApp Console local deployment

Use custom dashboards in NetApp Console local deployment to create focused monitoring views for specific teams, fleets, workloads, and operational goals. Custom dashboards complement the Home page dashboard by adding targeted visibility to broad, always-available monitoring.

How dashboard types work together

Home page dashboard
: Ready-to-use monitoring panes for capacity, alerts, performance capacity, licenses, and data protection status. Views are preconfigured and update automatically.
View Home page metrics in NetApp Console local deployment.

Custom dashboards
: Role-specific monitoring views built from predefined cards, including selectable scope, metrics, target resources, and time windows.

Use both together:
* Start on the Home page for daily baseline monitoring.
* Use custom dashboards for focused investigation by team, fleet, workload, or operational question.

You can add one custom dashboard to the Home page at a time.

Understand custom dashboards

A custom dashboard is a saved collection of cards that you arrange on a grid. Each card is based on a predefined card template with a metric and chart type. When you add a card, you configure the scope, target objects, aggregation, and time window.

Custom dashboards and cards are private to you. Each card shows data only for the storage resources you

are authorized to view.

You can remove the custom dashboard from your Home page and add a different one as monitoring priorities change.

Monitor what matters

Card templates are organized by metric type so you can focus a dashboard on a specific operational area. For example, use capacity cards to watch a volume that is growing faster than expected, or use alert cards to track repeated failures on a busy cluster:

Performance
: Latency, IOPS, throughput, utilization, and performance capacity across the fleet, systems, SVMs, volumes, and LUNs.

Capacity
: Used capacity, free capacity, capacity used percentage, and snapshot capacity across the fleet, systems, SVMs, volumes, LUNs, and local tiers.

Health
: Health status, degraded or critical object counts, failed disks, network interface errors, and hot objects.

Alerts
: Critical alert counts, alerts by severity, alerts by impact area, recent alerts, and alert trends over time.

For a complete catalog of predefined cards and metrics, see Learn about custom dashboard cards and metrics in NetApp Console local deployment.

Resource types

You can scope cards to any level of the ONTAP storage hierarchy: fleet, cluster, system (node), SVM, volume, LUN, and local tier (aggregate). The available resource types depend on the card template you select.

Plan dashboard views

Organize custom dashboards around storage administration goals such as workload performance triage, capacity planning, health risk detection, and alert prioritization.

Related information

- Get started with dashboards in NetApp Console local deployment
- View Home page metrics in NetApp Console local deployment
- Learn about custom dashboard cards and metrics in NetApp Console local deployment

Get started with dashboards in NetApp Console local deployment

In NetApp Console local deployment, follow this workflow to monitor from broad to focused views: review the Home page dashboard, review predefined dashboards, create custom dashboards, and keep dashboards current.

Required access roles

All roles. Dashboards show only the resources you are authorized to view. If you do not see a resource in the list of available resources, you do not have permission to view it.

1 Check baseline health on the Home page

Use the Home page dashboard to review organization-level capacity, alerts, performance, and data protection status.

- View Home page metrics in NetApp Console local deployment
- Learn about storage monitoring in NetApp Console local deployment

2 Review other predefined dashboards

Open **Storage > Dashboards** to review predefined dashboards for additional role-specific views.

- Manage dashboards in NetApp Console local deployment

3 Build custom dashboards for focused monitoring

Create custom dashboards when you need targeted views for selected resources, metrics, and time windows.

- Create a custom dashboard
- Customize dashboard cards in NetApp Console local deployment

4 Keep dashboards current

Update dashboards as priorities change by editing, duplicating, or deleting custom dashboards.

- Manage dashboards in NetApp Console local deployment
- Learn about custom dashboard cards and metrics in NetApp Console local deployment

Related information

- Learn about custom dashboards in NetApp Console local deployment
- View Home page metrics in NetApp Console local deployment

Build dashboards for monitoring goals in NetApp Console local deployment

Create a custom dashboard in NetApp Console local deployment, add cards, and save a monitoring view for daily operations.

If you only need organization-level monitoring that is ready to use, start with predefined Home page dashboards. Create a custom dashboard when you need a role-specific view, resource-level targeting, or a tailored card layout.

Console local deployment dashboard guidelines

- Only the user who creates the dashboard can view it. You cannot share dashboards with other users, even when you add them to the Console Home page.
- You can only view resources that you have permission to view. If you do not see a resource in the list of available resources, you do not have permission to view it.
- Before you begin, identify the metric types and resources you want the dashboard to track. Learn about

custom dashboard cards and metrics in NetApp Console local deployment

Build a dashboard for your monitoring goals

Create a dashboard when you need one place to monitor the metrics that matter to your role, such as fleet health, performance hotspots, and alert trends.

Steps

1. Select **Storage > Dashboards**.
2. Select **Add Dashboard**.

   Console local deployment creates a new dashboard with a default name and no description.

3. Select **Add Card**.
4. Select the **Resource type** for the card, such as fleet, system, or volume and then choose resources from the list of available resources.
5. Select **Next** to continue to choose a metric and card type.
6. Choose a metric to display. You can filter the available metrics by type: the **Metric Type**: **Performance**, **Capacity**, **Health**, or **Alerts** or search for a specific metric by name. The available metrics depend on the resource type you selected.

   Learn about the available metrics.

7. Select a card to include in the dashboard and select **Next**.

   You can only select one card at a time. The card preview shows live data for the selected metric and resource type.

8. Name your card. The name must be unique within the dashboard.
9. Review the card preview, enter a card name and description, and then select **Add**.
10. Repeat these steps for additional cards.
11. Select the pencil icon to edit the dashboard name and description to describe its purpose.
12. Select the action menu to the right of the Add card button and select **Save dashboard**.

    The saved dashboard is available under **Storage > Dashboards**.

13. Optionally, add this dashboard to your Home page by selecting **Add to Home page** from the action menu. The dashboard is added to the Home page. Only you can view custom dashboards that you create. It is not shared with others.

Related information

- View Home page metrics in NetApp Console local deployment
- Learn about custom dashboards in NetApp Console local deployment
- Customize dashboard cards in NetApp Console local deployment
- Manage dashboards in NetApp Console local deployment

Customize cards in NetApp Console local deployment

In NetApp Console local deployment, customize dashboard cards when you need to

focus operators on the signals that matter most, such as SLA risk, capacity pressure, or recurring alert noise. Use this topic to shape dashboard views around the operational decisions your team needs to make.

Before you begin

- Create or open a dashboard where you want to place cards.
- Review available templates and metrics in Learn about custom dashboard cards and metrics in NetApp Console local deployment.

Add cards while refining an existing dashboard

Use this option when you are refining an existing dashboard and need to quickly add cards that answer a specific operational question.

Steps

1. Open the dashboard you want to add a card to.
2. Select **Add Card**.
3. Select the **Resource type** for the card, such as fleet, system, or volume and then choose resources from the list of available resources.
4. Select **Next** to continue to choose a metric and card type.
5. Choose a metric to display. You can filter the available metrics by type: the **Metric Type**: **Performance**, **Capacity**, **Health**, or **Alerts** or search for a specific metric by name. The available metrics depend on the resource type you selected.

   Learn about the available metrics.

6. Select a card to include in the dashboard and select **Next**.

   You can only select one card at a time. The card preview shows live data for the selected metric and resource type.

7. Name your card. The name must be unique within the dashboard.
8. Review the card preview, enter a card name and description, and then select **Add**.
9. Repeat these steps for additional cards.
10. Save the dashboard.

Related information

- View Home page metrics in NetApp Console local deployment
- Learn about custom dashboards in NetApp Console local deployment
- Create a custom dashboard
- Manage custom dashboards
- Learn about custom dashboard cards and metrics in NetApp Console local deployment

Predefined dashboards and custom dashboard cards reference in NetApp Console local deployment

In NetApp Console local deployment, predefined dashboards and custom dashboard card

templates provide monitoring coverage for ONTAP performance, capacity, health, and alert operations.

Predefined dashboards

The following predefined dashboards are available out of the box.

Name	Description
Volumes	Volume performance, capacity, QoS, FabricPool, growth rate, and forecast metrics.
Time Till Full	ONTAP time-till-full predictions for clusters, volumes, and aggregates.
Security	Security compliance, encryption, autonomous ransomware protection, and authentication overview.
SVMs	ONTAP SVM performance, capacity, volume, LIF, protocol (CIFS, FCP, iSCSI, NFS, NVMe/FC), QoS, and latency heatmap monitoring.
Power	ONTAP cluster power consumption, temperature, and fan speed monitoring for nodes, shelves, and aggregates.
Nodes	ONTAP node performance, CPU, network, and backend monitoring.
Network	Network performance metrics including Ethernet, FibreChannel, NVMe/FC, Link Aggregation Groups, and Routes.
LUNs	ONTAP LUN performance, capacity, and efficiency monitoring.
Health	ONTAP cluster health monitoring including errors, warnings, EMS alerts, HA issues, node/disk/shelf health, volumes, licenses, and network ports.
Aggregates	ONTAP aggregate capacity, efficiency ratios, and growth rate monitoring.

Dashboard-level time windows and aggregation

Available time windows are 30 minutes, 1 hour, 24 hours, 48 hours, 7 days, and 30 days. Aggregation options vary by template and include views such as fleet-wide rollup or top-N style results. Time window and aggregation are configured at the dashboard level and apply to all cards in the dashboard.

Dashboard card templates and metrics

Cards are the building blocks of custom dashboards. Each card uses a predefined template with a metric and chart type, then maps that view to specific ONTAP objects and monitoring goals.

Card templates are predefined and not user-created.

Storage operations metric index (at a glance)

Metric types align to common storage administration outcomes, including performance bottlenecks, capacity pressure, health risk, and alert volume.

Metric type	Supported metrics	Storage admin use case	Detailed templates
Performance	Average latency, peak latency, IOPS, throughput (MB/s), utilization, performance capacity	Identify bottlenecks, sustained pressure, and performance headroom	Go to templates
Capacity	Used capacity, free capacity, capacity used (%), snapshot capacity used	Track growth, identify low-free-capacity areas, and monitor snapshot impact	Go to templates
Health	Health status, degraded/critical objects, failed disks, network interface errors, hot objects (by latency)	Detect health regressions and prioritize operational triage	Go to templates
Alerts	Alerts (Critical), alerts by severity, alerts by impact area, recent alerts, alert trend	Monitor active incidents and trend alert volume over time	Go to templates

Supported ONTAP object hierarchy (resource types)

Card data can be represented at multiple ONTAP object levels, from fleet-level visibility to object-level troubleshooting.

Resource type settings map to one or more of the following resource levels:

- Fleet
- Cluster
- System (node)
- SVM
- Volume
- LUN
- Local tier (aggregate)

The available resource levels depend on the template and metric you select.

Chart types for operational analysis

Chart type affects how quickly you can interpret trends, comparisons, distributions, and point-in-time values.

Chart type	Best for
Line chart	Time-series trends, such as latency, IOPS, throughput, utilization, and alert trends over time.
Bar chart	Categorical comparison and distribution, such as alerts by severity or impact area.
Pie or doughnut	Proportional distribution, such as capacity used percentage.
Table	Detailed, multi-column data, such as used capacity, health status, and recent alerts.
Single number (stat)	A single key value at a glance, such as critical alert counts or failed disks.

Performance templates for workload triage

Performance templates focus on latency, throughput, and utilization signals that indicate workload pressure.

Template	Chart type	Description
Average Latency	Line	Average latency across the selected targets over the chosen time window.
Peak Latency	Line	Maximum latency observed across the selected targets over the chosen time window.
IOPS	Line	Sum of I/O operations per second across the selected targets.
Throughput (MB/s)	Line	Sum of data throughput across the selected targets.
Utilization	Line	Average CPU or disk utilization across the selected targets.
Performance Capacity	Table	Headroom analysis comparing current utilization to the optimal point.

Capacity templates for growth and headroom

Capacity templates focus on consumed and available space to support growth planning and risk detection.

Template	Chart type	Description
Used Capacity	Table	Sum of used capacity across the selected targets.
Free Capacity	Table	Sum of free or available capacity across the selected targets.
Capacity Used (%)	Pie or doughnut	Average used-to-total ratio across the selected targets.
Snapshot Capacity Used	Table	Sum of snapshot space consumed across the selected targets.

Health templates for risk detection

Health templates focus on object status and failure indicators to support operational triage.

Template	Chart type	Description
Health Status	Table	Latest health state for each selected target object.
Degraded/Critical Objects	Single number	Count of objects whose health state is not OK.
Failed Disks	Single number	Sum of failed disks across the selected systems.
Network Interface Errors	Pie or doughnut	Sum of network interface error counters across the selected targets.
Hot Objects (by Latency)	Table	Resources ranked by peak latency, in descending order.

Alerts templates for incident prioritization

Alerts templates focus on incident volume, severity, and impact area to support monitoring and prioritization.

Template	Chart type	Description
Alerts (Critical)	Single number	Count of critical-severity alerts within the time window.
Alerts by Severity	Bar	Alerts grouped by severity level.
Alerts by Impact Area	Bar	Alerts grouped by impact area, such as capacity, performance, or security.
Recent Alerts	Table	Most recent alerts, sorted by time in descending order.
Alert Trend	Line	Alert count per time bucket over the chosen time window.

Related information

- View Home page metrics in NetApp Console local deployment
- Learn about custom dashboards in NetApp Console local deployment
- Create a custom dashboard
- Customize dashboard cards in NetApp Console local deployment
- Manage custom dashboards

Manage custom dashboards

Manage dashboards in NetApp Console local deployment

Manage custom dashboards in NetApp Console local deployment to keep views aligned with operations. You can edit, duplicate, and delete dashboards as teams, fleets, and priorities change.

This task applies only to custom dashboards under **Storage > Dashboards**. Predefined Home page dashboards cannot be edited in the same way.

View dashboards

View custom and predefined dashboards to find the dashboard you want to open, edit, duplicate, or delete.

You can duplicate a predefined dashboard to create a custom version that you can edit.

Steps

1. Select **Storage > Dashboards**.
2. Review the dashboard table to find the dashboard you want to open.
3. Select the dashboard name to open it and view its metrics.
4. If needed, use the available actions to edit, duplicate, or delete custom dashboards.

Update a dashboard as priorities change

Edit a dashboard when your operational priorities change and you need to adjust which cards and metrics are visible.

Steps

1. Select **Storage > Dashboards**.

2. Open the dashboard you want to edit.
3. Modify an existing card on the dashboard by selecting Edit from the actions menu of the card.
4. Add a new card to the dashboard by selecting **Add Card**.

   Customize cards in NetApp Console local deployment.

5. Select **Save changes** to save the dashboard.
   - You can also choose to save your changes as a new dashboard by selecting **Save as new dashboard** from the action menu. This option is useful when you want to keep the original dashboard for other teams or fleets while creating a new version for your current monitoring needs.
   - You can also choose to discard your changes by selecting **Discard changes** from the action menu. This option is useful when you want to revert to the last saved version of the dashboard.

Reuse a custom dashboard for another resource or fleet

Duplicate a dashboard when you want to reuse a proven layout for another team, fleet, or monitoring scenario without rebuilding it from scratch.

You can duplicate a predefined dashboard to create a custom version that you can edit.

Steps

1. Select **Storage > Dashboards**.
2. For the dashboard you want to duplicate, select **Duplicate** from the actions menu.

   Console local deployment creates a copy that includes all cards from the original dashboard.

3. Provide a name and description for the duplicated dashboard.
4. Select the duplicate dashboard you just created. Change the metrics, resources, and card types to suit your new monitoring scenario.

   Customize cards in NetApp Console local deployment.

Remove dashboards you no longer use

Delete a dashboard when it no longer supports current operations and you want to keep your dashboard list focused on active use cases.

Steps

1. Select **Storage > Dashboards**.
2. For the dashboard you want to remove, select **Delete**.
3. Confirm the deletion.

Related information

- View Home page metrics in NetApp Console local deployment
- Get started with dashboards in NetApp Console local deployment
- Learn about custom dashboards in NetApp Console local deployment
- Create a custom dashboard

- Customize dashboard cards in NetApp Console local deployment

Monitor fleets using inventory in NetApp Console local deployment

In NetApp Console local deployment, use inventory to monitor fleet health and investigate storage resources across your environment. Inventory provides capacity, security, and performance views that help you identify issues, understand resource utilization, and track managed storage systems.

Inventory is primarily an informational and diagnostic workspace. From inventory, you can review systems, volumes, and other storage objects across your fleets and perform common actions such as provisioning storage resources. For advanced storage management operations, NetApp Console local deployment redirects you to System Manager.

Access Inventory

You can access **Inventory** from several areas of Console local deployment, including:

- The Home page
- Fleet overview pages
- Fleet health cards and related inventory views

Depending on where you enter **Inventory**, the displayed scope can be at the organization level or fleet level.

Inventory views

Inventory provides multiple views that help you analyze different aspects of your storage environment.

Capacity
: View capacity utilization and identify systems, volumes, and other storage objects that consume storage resources.

Security
: Review security-related information and compliance status across managed storage resources.

Performance
: Analyze performance-related metrics and identify resources that might require further investigation.

The information displayed varies depending on the selected view and object type.

Investigate storage resources

Use **Inventory** to:

- Monitor fleet health
- Review storage capacity usage
- Track managed storage systems
- Identify volumes and other objects that consume capacity

- Investigate potential security or performance issues
- Locate resources that require administrative attention

Inventory helps you move from high-level visibility into your environment to more detailed investigation of specific storage resources.

Provision storage resources

Inventory includes workflows that allow you to provision storage resources such as volumes and LUNs.

When provisioning resources, you select an existing managed storage system and provide the configuration settings required for the workload.

Inventory and alerts

Alerts are generated for specific storage objects and conditions within your environment.

When you select an alert, Console local deployment might direct you to System Manager for additional investigation or management actions. **Inventory** complements alerts by providing broader visibility into the overall state of your storage environment and managed resources.

Advanced management tasks

Inventory helps you identify resources that require action, but some advanced storage management operations are performed in System Manager.

Examples include:

- Advanced workload management
- Resource relocation and optimization tasks
- Detailed system administration activities

Use **Inventory** to identify and investigate issues, and then use System Manager when deeper management capabilities are required.

Remediate alerts

Learn about alerts in NetApp Console local deployment

NetApp Console local deployment generates alerts that identify health issues on your on-premises ONTAP clusters and for NetApp Backup and Recovery operations.

Console local deployment consolidates alerts from ONTAP clusters and Backup and Recovery operations into a single view. The Alerts page includes a dashboard, an alerts list, and a systems view so you can review alerts across the entire organization or scope them to a specific fleet or system. Each alert identifies the affected system, its severity, and its impact area.

Use alerts in Console local deployment to:

- Detect capacity, connectivity, data protection, performance, and security issues.
- Prioritize alerts by severity and impact area.

- Open an alert in System Manager or Workload Analyzer, then run a guided or automated Fix-It action when the page offers one.
- Route notifications through email to users with specified access roles, or send alert data to an external system through a webhook.
- Export the alerts list to a CSV file for offline analysis.

Alert severity and impact areas

Each alert includes a severity and an impact area:

- **Severity** indicates urgency, from highest to lowest: Critical, Major, Minor, Warning, and Info.
- **Impact area** identifies the affected domain: Capacity, Connectivity, Data protection, Performance, and Security.

Alert sources and scope

- **ONTAP alerts** originate from the ONTAP Event Management System (EMS) on the on-premises clusters that Console local deployment has discovered. Learn more about the ONTAP EMS and available alerts.
- **NetApp Backup and Recovery alerts** originate from Backup and Recovery actions. Learn more about NetApp Backup and Recovery alerts.
- Your permissions determine which fleets you can view. Scope the view to the entire organization, a specific fleet, or an individual system. The **Systems health** tab shows per-system status and the **Alerts** tab shows the current alert list for the selected scope.
- The CSV export reflects the current scope, search text, and filters.

Alert notification rules

Create notification rules to determine which alerts generate notifications and who receives them. A notification rule has the following characteristics:

- It matches alerts on the service (such as ONTAP management or Backup and Recovery), severity, impact area, and target system.
- For email delivery, it sends notifications to users with the specified access roles. For webhook delivery, it sends alert data to the configured endpoint—access to that data is controlled by the receiving application, not by Console local deployment.
- It applies to specific systems, not to fleets.
- It can be muted during a planned maintenance window to suppress expected alerts.

Console local deployment includes a default notification rule that is active immediately after installation. The default rule monitors all services and systems for Critical-severity alerts and delivers notifications to the Console Notification Center only—no email or webhook delivery is configured until you set it up. You can view and adjust this rule, and create custom rules to match your environment.

Info-level alerts are visible in the Alerts page but are not delivered by notification unless you explicitly create a rule that includes the Info severity level.

Related information

- Monitor and investigate alerts
- Create and manage alert notification rules

- Learn about ONTAP alerts in NetApp Console local deployment

Monitor and investigate alerts in NetApp Console local deployment

Monitor and investigate alerts in NetApp Console local deployment to keep your storage healthy and minimize disruption.

View active alerts across your entire organization or a specific fleet, prioritize them by severity and impact area, and investigate root causes so you can resolve issues before they escalate. When an alert includes a recommended action or Fix It option, you can move from investigation directly into remediation.

Required access role

All roles except Federation admin and Federation viewer. Users with the Federation admin or Federation viewer role cannot view alerts. Learn about access roles.

For ONTAP alerts, you can access tools such as System Manager and Workload Analyzer directly from the Alerts page to investigate and resolve issues.

For external reporting, you can export the alerts list to a CSV file for offline analysis.

You can also set up notification rules to receive alerts by email or webhook. Learn how to set up alert notifications.

Available alerts

NetApp Console local deployment supports alerts for ONTAP and NetApp Backup and Recovery.

- Learn about ONTAP alerts in NetApp Console local deployment
- Learn more about NetApp Backup and Recovery alerts.

View the alerts dashboard

Use the alerts dashboard to get a quick view of current alert activity for the scope you selected. The dashboard summarizes active alerts by severity and impact area, and it includes a chart that shows alert distribution over the past 24 hours so you can spot trends quickly.

You can filter the dashboard to focus on critical alerts or alerts for a specific impact area.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
   - **All** (default) to see alerts across all fleets you have access to
   - A specific fleet to see alerts only for that fleet
3. Filter the dashboard according to the alerts you need to view, including:
   - Severity (Critical, Warning, or Informational)
   - Critical alerts grouped by impact area
   - Impact area (Security, Protection, Availability, Performance, Capacity, or Configuration)

View all alerts

Use the Alerts tab to view the full list of active alerts for the scope you've selected. The list updates to reflect the current organization or fleet scope, and you can search the list for specific alerts by name or description.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
   - **All** (default) to see alerts across all fleets you have access to
   - A specific fleet to see alerts only for that fleet
3. Select the **Alerts** tab to view the full list of active alerts for the selected scope.
4. Optionally, search the list for a specific alert by name or description.

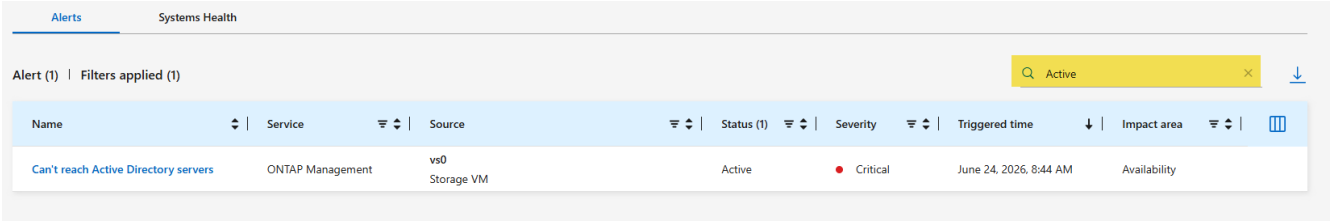

View alerts by system

You can view alerts for a specific system within a fleet or your organization.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
   - **All** (default) to see alerts across all fleets you have access to
   - A specific fleet to see alerts only for that fleet
3. Select the **Systems health** tab to view a summary of the alerts associated with the systems within the scope you've selected.
4. Select **View alerts** on the row of the respective system to view the full list of active alerts associated with that system.

Investigate an alert

You can investigate an alert to see what triggered it and who received the notification.

When investigating an ONTAP alert, you can also go directly to the System Manager interface for the system that generated the alert to view additional details and take action.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
   - **All** (default) to see alerts across all fleets you have access to
   - A specific fleet to see alerts only for that fleet

3. In either tab, select the name of the alert you want to investigate to view its details.

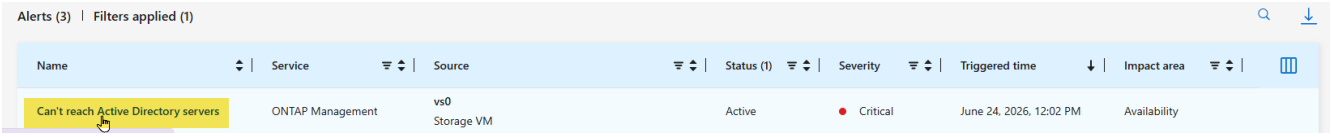

4. If applicable, select the VM or cluster name to open the System Manager interface for the system that generated the alert.

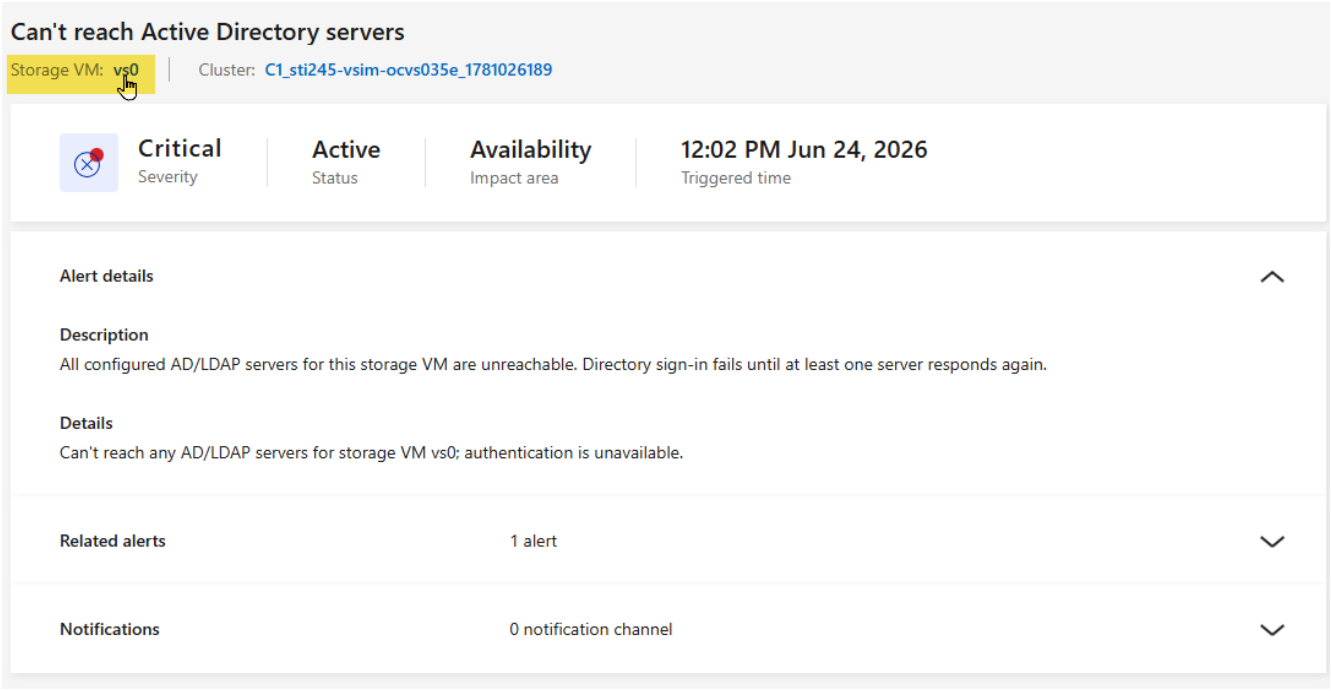

Remediate an alert

When an alert includes a recommended action or Fix It option, use it to move from investigation into remediation without leaving the alert details.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
   - **All** (default) to see alerts across all fleets you have access to
   - A specific fleet to see alerts only for that fleet
3. Select the alert you want to remediate.
4. Review the alert details, then select the recommended action or **Fix It** option if one is available.
5. Follow the guided steps to apply the remediation and then return to the alert details to confirm the alert state.

If the action resolves the issue, the alert no longer appears as active for that scope. If the alert remains active, review the alert details again and continue the investigation.

Analyze an alert

You can analyze an ONTAP alert in Workload Analyzer to understand what triggered it.

When investigating an ONTAP alert, you can also go directly to the System Manager interface for the system that generated the alert to view additional details and take action.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
   - **All** (default) to see alerts across all fleets you have access to
   - A specific fleet to see alerts only for that fleet
3. Select the **Systems health** tab to view the full list of active alerts for the selected scope.
4. In either tab, select the name of the alert you want to investigate to view its details.

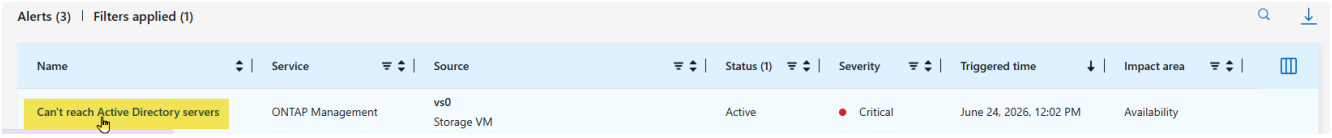

5. If applicable, select **Analyze** to open the Workload Analyzer interface for the system that generated the alert.

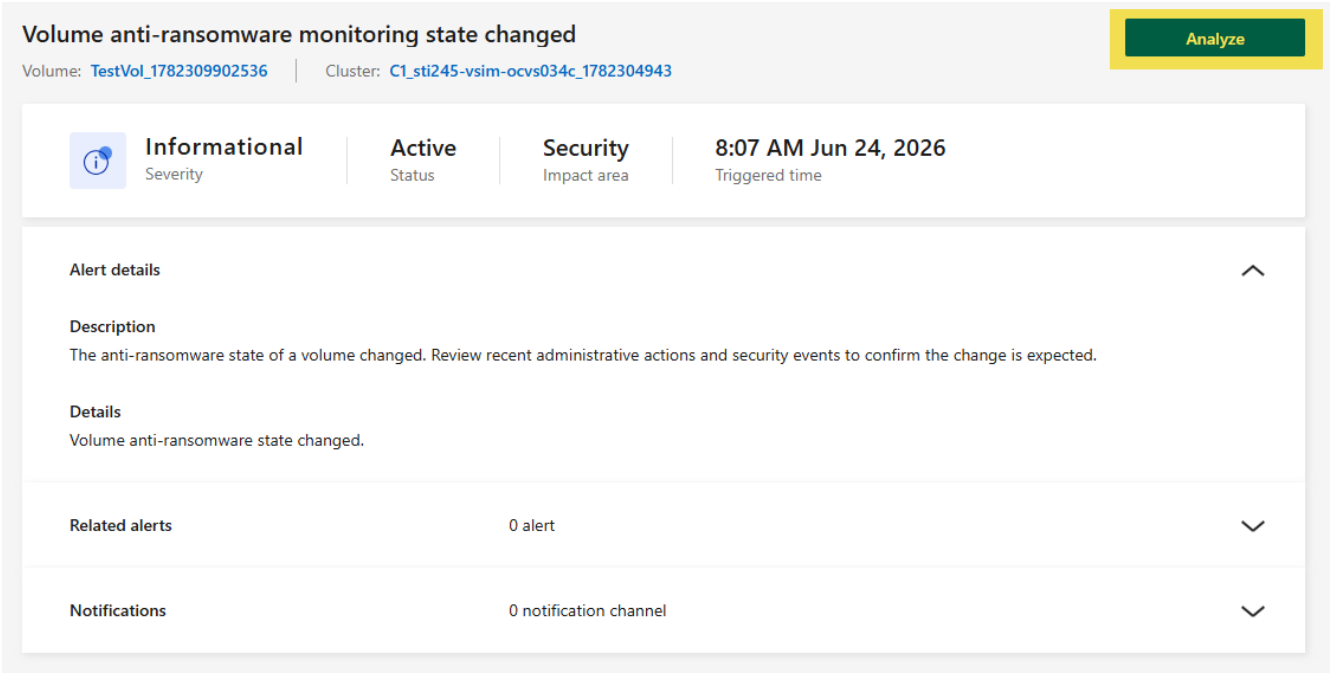

6. Enter the time range that covers the period when the alert was triggered, then select **Analyze** to view the analysis results. Learn about Workload Analyzer.

Export alerts to CSV

Export the alerts list in the NetApp Console local deployment to a CSV file for offline analysis or reporting. The export reflects the current scope (organization or fleet), search text, and filters.

Steps

1. Select **Health > Alerts** and then select the **Alerts** tab.
2. Set the scope (organization or fleet) and apply any filters or search text you want included in the export.
3. Select the download icon to export the alerts list to a CSV file.

Configure notification rules for alerts in NetApp Console local deployment

Configure notification rules in NetApp Console local deployment to control which alerts raise notifications, who receives them, and how they are delivered.

Required access roles

Super admin, Organization admin, Storage admin, Operations support analyst, or any of the admin roles for NetApp Backup and Recovery. Learn about access roles.

Use notification rules to route the right alerts to the right people through the channels that fit your environment, while reducing noise from alerts that aren't relevant to you. Notification rules complement the Alerts page: you investigate or remediate the alert in the Alerts workflow, then use rules to control how similar alerts are delivered in the future.

A notification rule matches alerts based on conditions you define—such as service, severity, and impact area—and then delivers a notification through the channels you select. The Console local deployment includes default notification rules that you can view and adjust, and you can create your own custom rules. You can also mute rules to temporarily suppress notifications during planned events like maintenance windows.

- Learn more about the ONTAP EMS and available alerts.

Before you begin

- To deliver notifications by email, your Organization admin must configure an email server in the Console local deployment. To deliver notifications to an external system, your Organization admin must configure a webhook. For the steps, see Configure notification delivery.
- The Console local deployment Notification Center displays notifications by default, so no additional setup is required for in-console notifications.

View notification rules

The notification rules page is the central place to review and manage every rule that applies to your organization. Each rule shows its key attributes so that you can quickly assess and adjust your alerting behavior.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification settings**.
3. Review the rules in the list. Each rule shows its active status, name, the services and severity levels it monitors, the roles authorized to receive notifications, the enabled delivery channels, when it was last modified, and any scheduled mute period.
4. To find a specific rule, use the filter and search controls to filter by active status, service, severity, role, channel, or mute status.

Create a notification rule

Create a notification rule to define the conditions that trigger a notification and how that notification is delivered.

> You can't set notification rules for fleets. You can set them only for specific systems. In large environments with many systems, consider creating broader rules by severity rather than duplicating rules per system—for example, one Critical rule that covers all systems acts as a catch-all, with additional targeted rules for systems that need different routing or recipients.

Example notification rule for Critical alerts across all systems

Suppose you want to ensure no critical alert goes unnoticed, regardless of which system it originates from. You can create a broad rule that routes all critical alerts to the Console Notification Center for storage admins.

In this example, you create a rule with the following settings:

- **Services**: All
- **Severity**: Critical
- **IAM role**: Storage admin
- **System**: (Leave this blank to apply to all systems)
- **Delivery method**: Console Notification Center

With this rule in place, storage admins see a notification in the Console for every critical alert across all systems. This rule acts as a baseline that you can supplement with more targeted rules.

Example of a targeted notification rule for Storage admin capacity alerts

Suppose your storage admins need to respond immediately to critical capacity issues on a specific production system, but you don't want lower-severity alerts to flood their inboxes. You can create a targeted rule that emails storage admins only when a critical capacity alert fires on that system.

In this example, you create a rule with the following settings:

- **Services**: ONTAP management
- **Severity**: Critical
- **Impact area**: Capacity
- **IAM role**: Storage admin
- **System**: <system_name>
- **Delivery method**: Email

With this rule in place, Console local deployment emails all storage admins for the specified system when a critical capacity alert occurs. Alerts with a lower severity, such as warning or informational, don't generate an email, so the team can focus on the issues that require urgent attention.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification settings**, and then select **Add rule**.
3. Define the conditions that the rule matches:
    - **Rule name**: enter a concise, descriptive name. This field is required.
    - **Rule description**: optionally, enter additional context about the rule's purpose.

- **Services**: select one or more ONTAP or platform services that the rule applies to.
- **Roles**: select the access roles that users must have to receive notifications when the rule is triggered.
- **Severity levels**: select which severity levels the rule monitors, such as Critical, Warning, Error, or Information.
- **Impact area**: select the operational areas to match, such as Capacity or Performance.
- **Alert name**: select the specific alert types that trigger the rule.
- **System**: select the system context that the rule applies to.

4. Select the delivery channels for the notification:
   - **Email**: sends alert emails to users who have the selected roles. Requires a configured email server.
   - **Webhook**: sends alert data to an external system. Requires selecting a configured webhook integration.
   - **Console Notification Center**: displays real-time alerts for users who have the selected roles.
5. Optionally, to suppress notifications from this rule during a planned period, such as a maintenance window, set a **Mute notifications** schedule, and choose a recurrence if you want the mute period to repeat. You can add multiple mute windows per rule, which is useful for recurring maintenance cycles such as weekly patching.
6. Select **Create**.

Enable or disable a notification rule

Enable or disable individual notification rules to control which conditions raise notifications. Disable rules that aren't relevant to your environment to reduce noise, and enable rules to start receiving their notifications again.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification settings**.
3. Locate the rule you want to change.
4. Toggle the rule's **Active** switch on or off. The change takes effect immediately.

Mute a notification rule

Mute a notification rule to suppress alert delivery during a planned event, such as a maintenance window, without disabling the rule. Alerts continue to appear in the Alerts page during the mute period—only the email, webhook, and in-console notifications are suppressed. When the mute window expires, notifications resume automatically.

You can add multiple mute windows to a single rule and configure each one to recur on a schedule.

Examples of mute windows

- **One-off maintenance window**: You're running a firmware upgrade on a storage system Saturday night and want to suppress the expected alerts during that window. Set a mute window for Saturday 10:00 PM to Sunday 2:00 AM with no recurrence. When the window expires, notifications resume automatically.
- **Recurring weekly patching window**: Your team patches systems every Sunday from midnight to 4:00 AM. Add a mute window with weekly recurrence so notifications are suppressed automatically each week without manual intervention. If a second system has its own patching schedule at a different time, add a second mute window to the same rule with its own start time and recurrence.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification rules**.
3. In the rules list, locate the rule you want to mute and select the action menu for that rule.
4. Select **Edit**.
5. In the **Mute notifications** section, set the start and end date and time for the mute window.
6. Optionally, select a recurrence to repeat the window on a schedule.
7. To add additional mute windows, select **Add mute window** and repeat the previous steps.
8. Select **Save**.

Delete a notification rule

Delete a notification rule if you no longer need it. Deleting a rule permanently removes it, so you can't recover it.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification settings**.
3. In the rules list, locate the rule you want to delete and select the action menu for that rule.
4. Select **Delete** from the action menu.

Related information

- Configure notification delivery
- Learn about Console alerts
- View alerts

Remediate storage class drift in NetApp Console local deployment

In NetApp Console local deployment, find drift-related alerts, analyze the drift findings, and remediate workloads that no longer match their storage class intent.

Required roles

Storage admin

You can only view and remediate workloads in fleets where you have permission.

Remediate drift

When a workload no longer matches its storage class intent, NetApp Console local deployment raises an alert. Use the alert to analyze the drift and apply the recommended remediation.

You can apply some remediations directly from the alert. For other issues, update the workload configuration or assign a different storage class to restore compliance. The remediation workflow shows the expected impact before you apply changes.

Steps

1. Select **Storage > Storage classes**.

   A summary of storage class drift appears in the **Drifts by storage class** area. The complete list appears in the table.

2. In the **Drifts** column, select **View** for the relevant storage class.

   The **Alerts > Overview** page opens with filters applied.

3. Select an alert to open the alert details page.
4. Select **Analyze**.
5. Review the findings in **Workload analyzer**.

   For configuration drift findings, compare the intended and actual settings to determine what changed.

6. Select the recommended remediation action, such as **Fix it**.
7. Review the proposed changes and apply the remediation.
8. Return to **Storage > Storage classes** and verify that the workload no longer appears as drifted.

   Compliance evaluations can take several minutes to reflect recent configuration changes. If the workload is still listed as drifted, return to the alert details page and select **Analyze** to review any remaining findings.

Related information

- Learn about storage class drift and compliance in NetApp Console local deployment
- Learn about storage alerts
- View alerts

Alert remediation actions in NetApp Console local deployment

Use this reference to understand the remediation actions available from **Fix it** in NetApp Console local deployment. For each action, the table describes what the action does and the related alert. Review the action details in the confirmation dialog before you run it.

Remediation actions

Remediation action	Alert name	Alert description	Impact area	Remediation summary
Enable automatic volume growth	Volume full	The volume is running low on space and is nearing full capacity.	Capacity	Automatically increases a volume's size (up to a configured limit) to reduce the risk of running out of space.
Resize volume	Volume full	The volume is running low on space and is nearing full capacity.	Capacity	Increases a volume's size to provide more space immediately.

Remediation action	Alert name	Alert description	Impact area	Remediation summary
Bring volume online	Volume offline	The volume is offline and not serving data.	Availability	Brings an offline volume online so it can resume serving data.
Bring aggregate online	Aggregate offline	The aggregate is not online and volumes hosted on it may be unavailable.	Availability	Brings an offline aggregate online to restore access to the volumes it hosts.
Bring LUN online	LUN offline	The LUN is offline and host access is unavailable.	Availability	Brings an offline LUN online so hosts can resume access and I/O.
Unrestrict volume	Volume restricted	The volume is in a restricted state and may not serve I/O normally.	Availability	Returns a restricted volume to an online state so clients can access it again.
Bring NVMe namespace online	NVMe namespace is offline	The NVMe namespace is offline and host access is unavailable.	Availability	Brings an offline NVMe namespace online to restore host access.
Bring intercluster LIF up	Intercluster LIF down	An intercluster LIF is down and cluster-to-cluster communication may be impacted.	Connectivity	Brings an intercluster LIF up to restore cluster-to-cluster connectivity used for replication.
Re-enable MetroCluster AUSO	MetroCluster automatic unplanned switchover disabled	Automatic unplanned switchover is disabled on this cluster.	Availability	Re-enables automatic unplanned switchover (AUSO) so MetroCluster protection works as intended.
Configure Service Processor network	Service Processor is not configured	The Service Processor (SP) network is not configured.	Manageability	Configures the Service Processor (SP) network settings to enable out-of-band management access.
Assign unassigned disks	Unassigned disks detected	Unassigned disks are present and available capacity may be unused. Assign the disks to a node so they can be used for storage.	Availability	Assigns currently unassigned disks to a node so they can be used for storage.
Root volume space recovery	Node root volume space low	Available space on the node root volume is low and may affect normal system operation.	System health	Frees space on the node root volume by deleting the largest snapshot or the largest core dump file. Deletions are permanent.

Investigate performance issues

Learn about NetApp Console local deployment Workload Analyzer

Use NetApp Console local deployment Workload Analyzer to view a single volume's behavior over time. You can investigate performance degradation, capacity trends, and resource contention issues from the volume itself, from an alert, or from inventory.

How Workload Analyzer identifies root causes

Use the Workload Analyzer to correlate metrics from a single volume across six sections so you can identify root causes quickly. Select a volume and a time range to view events, performance, and capacity together within the same window. This view can help you determine whether storage is the likely source of an application issue or whether another layer, such as networking, is causing the problem.

The analyzer is most useful when you already know which volume is affected. If you open it from an alert or from volume inventory, Console local deployment opens the analyzer with the volume selected so you can focus on the analysis window and the charts.

The analyzer tracks one volume at a time, so use it to investigate a specific issue rather than compare multiple volumes.

Review the capacity section as part of this assessment. When an ONTAP system is more than 85% full, the high utilization can itself cause performance issues, so low available capacity can explain latency that the performance charts alone don't account for.

After you identify a root cause, you can act directly from the analyzer. When automated options are available, Console local deployment provides Fix-It recommendations with guided steps or one-click remediation.

Workload charts for LUNs don't provide the same level of detail as the charts for volumes, so you'll see fewer statistics when you analyze a LUN.

Access Workload Analyzer

You can open the Workload Analyzer in several ways:

- From the **Health** menu, select **Workload Analyzer**, then search for and select the volume you want to investigate in the **Volume** field.
- From the **Storage > Fleets > Inventory** page, navigate to a volume you want to analyze and select **Analyze** from the action menu.
- From the **Alerts > Overview** page, select an alert for the volume you want to investigate, then select **Analyze** from the alert details.

You can also analyze LUNs, but they show fewer statistics than volumes.

Analyze latency, throughput, and storage capacity

Use these six sections to investigate a volume:

Volume selection

Choose the volume and time range to investigate so all charts and events align to the same workload and analysis window.

Event timeline

View current and historical configuration changes, warning alerts, and critical alerts for the selected volume during the analysis window. Use this section to correlate "what changed" with changes in performance or capacity behavior.

Latency analysis

View volume latency over time, either as a total or broken down by delay center—network, data processing, aggregate operations, cluster interconnect, and others. If the volume uses a QoS policy, the analyzer displays the policy's maximum and minimum latency limits as reference lines so you can see how close the workload is to a throttling threshold. Toggle the forecast to project whether latency is trending toward a warning or critical threshold.

Throughput analysis

View IOPS and MB/s over time, with an optional read/write/other breakdown. If the volume uses a QoS policy, the analyzer displays IOPS and throughput limit lines. Toggle the forecast to project whether demand is trending toward QoS limits.

Resource utilization

View how busy the nodes and aggregates serving the volume are during the analysis window. The analyzer plots each resource as an independent line rather than stacking values, so you can identify whether a specific node or aggregate is a contention source. Toggle the forecast to project whether any resource is trending toward a high-utilization threshold.

Capacity analysis

View how physical space and available space in the volume have changed over time. Hover to see a storage efficiency breakdown including deduplication, compression, and compaction savings. Toggle the forecast to project when the volume will reach warning or critical capacity thresholds. Switch to the snapshot view to see how snapshot space is tracking against the configured snapshot reserve.

Forecast capacity and performance trends

Use the forecast toggle in each analysis section to extend the chart timeline beyond the current time and project future values based on the observed trend. A dashed line distinguishes forecast values from actual data. The analyzer also displays an insight message when the forecast indicates that a threshold breach is likely, along with an estimated time to breach.

Related information

Investigate high latency on a volume.
Investigate high throughput demand on a volume.
Investigate capacity growth on a volume.
Learn about Workload Analyzer metrics in NetApp Console local deployment.

Investigate high latency on a volume in NetApp Console local deployment

Use Workload Analyzer in NetApp Console local deployment to find the source of a volume's slow response times.

If you open the analyzer from an alert or from volume inventory, the volume is already selected and you can focus on the time range and chart breakdowns.

When application performance degrades, identify which part of the I/O path is causing the slowdown. The latency analysis section breaks down response time by delay center so you can distinguish between network issues, data processing overhead, aggregate contention, and other contributors.

Steps

1. Open Workload Analyzer using one of the following methods:
   - From the **Health** menu, select **Workload Analyzer**, then search for and select the volume you want to investigate in the **Volume** field.
   - From the **Storage > Fleets > Inventory** page, navigate to a volume you want to analyze and select **Analyze** from the action menu.
   - From the **Alerts > Overview** page, select a capacity-related alert for the volume you want to investigate, then select **Analyze** from the alert details.
2. Set the **Time Range** to cover the period when the performance issue occurred, then select **Analyze**.
3. Review the **Event timeline** section to see configuration changes and alerts that line up with the latency increase.

   The analyzer plots configuration change events (wrench icon) and alert events (warning and critical icons) along the same time axis as the latency chart, making it easy to see whether a change preceded the degradation.

4. In the **Latency** section, open the **View** dropdown and select **Breakdown by delay center**. The chart shows the contribution of each delay center to the total latency over time. Delay centers include:
   - Read latency
   - Write latency
   - Other latency
5. Hover over a point on the chart where latency is highest to see each delay center's value and its percentage of total latency.

   The largest delay center usually points to the resource under contention. When a shared resource cannot keep up with demand, the volumes that use it wait longer for I/O. Reduce the load on that resource, such as by moving workloads or adjusting a QoS limit, to lower latency.

6. Identify the dominant contributor and use the value to determine next steps:
   - High **Read latency** may indicate disk contention. Check the **Resource Utilization** section to confirm aggregate busy percentage.
   - High **Write latency** may indicate NIC saturation or network path issues outside the cluster.
   - High **Other latency** may indicate that inline efficiency settings are consuming more CPU than the workload can tolerate. Consider adjusting efficiency settings or QoS.
   - High **Cloud Latency** may indicate the workload is frequently accessing cold data on the capacity tier. Consider adjusting the tiering policy.
7. If the delay centers don't account for the slowdown, check the **Capacity Analysis** section. When the ONTAP system is more than 85% full, the high utilization can cause performance issues regardless of the delay center breakdown.
8. If latency increased immediately after a change event, use the event details and the related charts together to validate the likely cause:
   - For QoS changes, compare the latency line to the QoS limit lines and review the throughput charts to see whether demand is meeting a policy ceiling.
   - For aggregate or volume moves, compare the latency spike to the node and aggregate utilization values shown for the same timeframe.
   - For tiering policy changes, review the cloud latency contribution to determine whether cold-data access

increased after the change.

After you finish

If a configuration or placement issue caused the problem and Console local deployment can resolve it, the alert on the volume may offer a Fix-It option.

Learn about Workload Analyzer metrics in NetApp Console local deployment.

Investigate high throughput on a volume in NetApp Console local deployment

Use Workload Analyzer in NetApp Console local deployment to examine a volume's IOPS and throughput demand over time.

When a volume's workload is consuming more IOPS or throughput than expected, identify what's driving the demand. The throughput analysis section shows IOPS and MB/s with an optional read, write, and other breakdown so you can identify which type of I/O is driving high demand and whether that demand is trending toward a QoS limit.

If you open the analyzer from an alert or from volume inventory, the volume is already selected and you can focus on the time range and throughput charts.

Steps

1. Open Workload Analyzer using one of the following methods:
   - From the **Health** menu, select **Workload Analyzer**, then search for and select the volume you want to investigate in the **Volume** field.
   - From the **Storage > Fleets > Inventory** page, navigate to a volume you want to analyze and select **Analyze** from the action menu.
   - From the **Alerts > Overview** page, select an alert for the volume you want to investigate, then select **Analyze** from the alert details.
2. Set the **Time Range** to cover the period of high demand, then select **Analyze**.
3. Scroll to the **Throughput Analysis** section.
4. Open the **View** dropdown and select **Breakdown (RWO)**.

   The charts split into read, write, and other components for both IOPS and MB/s. This lets you determine whether read-heavy access patterns, write-heavy patterns, or a combination is driving the demand.

5. Use the component picker to enable or disable the metrics you want to examine:
   - **Read IOPS** and **Write IOPS** — operations per second by I/O direction
   - **Read MB/s** and **Write MB/s** — throughput in megabytes per second by I/O direction
   - **Other IOPS** and **Other MB/s** — metadata and other non-data operations
   - **QoS IOPS Limit** and **QoS MB/s Limit** — policy ceilings that appear only when the volume uses a QoS policy
   - **Cloud Throughput** — MB/s written to and read from the cloud, shown only for workloads that tier capacity to the cloud through FabricPool
6. Hover over a peak in the chart to see the exact value and percentage breakdown for each component at that point in time.

   The hover tooltip also shows the average I/O size (throughput divided by IOPS) for each category, which

can indicate whether the workload is performing large sequential I/O or small random I/O.

7. Enable **Show Forecast** to project whether current throughput trends will reach the QoS IOPS or MB/s limit.

   If the forecast line is approaching a QoS limit line, ONTAP might throttle the workload soon.

8. If you want to see aggregate demand without the breakdown, open the **View** dropdown and select **Totals**.

   The totals view shows a single IOPS line and a single MB/s line, which is useful for a quick overall demand summary.

After you finish

Use the throughput patterns you observe to decide what to do next:

- If read IOPS are very high relative to write IOPS, consider whether read-ahead settings or caching can reduce the load on the volume.
- If the workload is approaching or has reached the QoS IOPS or MB/s limit, use the QoS limit lines and related metric definitions to confirm throttling and decide whether you need to adjust the limit.
- If throughput is high and latency is also elevated, check the **Resource Utilization** section to determine whether the underlying node or aggregate is under contention. Compare throughput peaks, latency spikes, and resource utilization over the same period.
- If the analyzer shows rapid capacity or snapshot growth while demand is increasing, review the capacity and snapshot metrics to understand how that growth might affect the volume.

Investigate high latency on a volume.
Learn about Workload Analyzer metrics in NetApp Console local deployment.

Investigate capacity growth on a volume in NetApp Console local deployment

Use Workload Analyzer in NetApp Console local deployment to investigate why a volume is running out of space.

When a volume is filling up or snapshot copies are consuming more space than expected, examine capacity and snapshot trends over time. The capacity analysis section shows how physical and available space change, breaks down storage efficiency savings, and projects when the volume will reach a capacity threshold.

If you open the analyzer from an alert or from volume inventory, the volume is already selected and you can focus on capacity trends and the forecast window.

Steps

1. Open Workload Analyzer using one of the following methods:
   - From the **Health** menu, select **Workload Analyzer**, then search for and select the volume you want to investigate in the **Volume** field.
   - From the **Storage > Fleets > Inventory** page, navigate to a volume you want to analyze and select **Analyze** from the action menu.
   - From the **Alerts > Overview** page, select an alert for the volume you want to investigate, then select **Analyze** from the alert details.
2. Set the **Time Range** to cover the period of capacity growth, then select **Analyze**.
3. Scroll to the **Capacity Analysis** section.

4. Select **Capacity view** from the **View** dropdown.

   The chart shows physical capacity as the bottom stacked area and available capacity as the top stacked area. Together they equal the total volume size, so you can see how quickly available space is decreasing.

5. Hover over a point on the chart to see the storage efficiency breakdown, including deduplication, compression, and compaction savings, and the overall storage efficiency ratio.

   A declining efficiency ratio while physical capacity rises can indicate that newly written data isn't deduplicating or compressing as well as the existing data.

6. To investigate snapshot consumption, select **Snapshot View** from the **View** dropdown.

   The chart shows the snapshot reserve as a horizontal reference line and snapshot used capacity as an area below it. Hover to see the snapshot space used and its percentage of the reserve, the total snapshot count, and the age of the oldest snapshot.

   When snapshot consumption exceeds the reserve, snapshots begin consuming space from the volume's data area, which reduces the space available for new writes.

7. If the volume tiers data to the cloud, select **Cloud Tier View** from the **View** dropdown to compare how much capacity is on the local performance tier versus the cloud capacity tier.
8. Enable **Show Forecast** to project when the volume will reach a warning or critical capacity threshold.

   The capacity forecast requires at least 10 days of historical data. When fewer than 30 days of capacity remain before the volume is full, the analyzer identifies the storage as almost full. The forecast displays an insight message with an estimated time to breach.

After you finish

Use the capacity trends you observe to decide what to do next:

- If available capacity is trending toward full, consider increasing the volume size, enabling autogrow, or moving data off the volume.
- If snapshot used capacity is exceeding the reserve, review your snapshot policy and retention to reclaim space.
- If the storage efficiency ratio is dropping, review whether the workload's data type or inline efficiency settings are reducing savings. Use Learn about Workload Analyzer metrics in NetApp Console local deployment for detailed descriptions of capacity, snapshot, and efficiency metrics.

Workload Analyzer metrics in NetApp Console local deployment

In NetApp Console local deployment, the Workload Analyzer shows metrics across six sections. Each metric description explains how the analyzer renders the metric on the chart and what it indicates about volume behavior.

QoS reference lines in latency charts

If the selected volume uses a QoS policy, you can use two additional reference lines in the metric picker:

Reference line	Description
QoS Limit Max	The maximum latency ceiling defined by the QoS policy. When the latency line approaches or touches this line, ONTAP is throttling the workload, and the QoS limit—rather than a physical resource—is the dominant source of latency.
QoS Limit Min	The guaranteed minimum latency floor defined by the QoS policy. This line indicates the response time floor enforced for the workload. When other workloads use QoS minimums to guarantee their throughput, ONTAP can throttle this workload, which then shows higher latency.

These horizontal lines do not appear in the delay center breakdown on hover. They serve as reference thresholds, not as contributors to latency.

Latency breakdown by I/O type

Use the I/O type breakdown in the **Latency Analysis** section after you select **Breakdown by delay center** from the View dropdown. The chart breaks total latency into three categories based on the direction of the I/O operation. Use these metrics to determine whether a performance issue is affecting reads, writes, or metadata operations.

Latency type	Description	Common causes of elevated values
Read latency	Average time to complete a client read request on the volume, from when the request is received until data is returned. Read requests must traverse the full I/O path from the network protocol layer through the data processing stack to the aggregate where the data resides.	Aggregate or disk contention; cache misses that promote reads to physical disk; cold data retrieved from a cloud capacity tier via FabricPool; cross-node reads when data resides on a different node than the one handling the request.
Write latency	Average time to acknowledge a client write request on the volume. ONTAP acknowledges a write once it is committed to the NVRAM write log; the data is later flushed to the aggregate.	NIC saturation or high round-trip latency between the client and the storage node; synchronous SnapMirror waiting for the destination cluster to confirm receipt before the write can be acknowledged; NVRAM pressure under heavy write loads; CPU overhead from inline deduplication, compression, or compaction running in the write I/O path.
Other latency	Average time to complete non-read, non-write operations on the volume, including file opens, attribute lookups, directory traversals, file creates, and locks. Elevated other latency can affect application responsiveness even when read and write latency appear normal.	CPU pressure from inline efficiency operations competing with metadata processing; high namespace activity from workloads that generate large numbers of file creates, deletes, or directory scans; QoS throttling that constrains total IOPS, leaving fewer IOPS available for metadata operations; volume activation overhead when many volumes are concurrently active.

Throughput components

Use the throughput components in the **Throughput Analysis** section after you select **Breakdown (RWO)** from

the View dropdown. The analyzer shows both IOPS and MB/s simultaneously.

Component	Description	Rendered as
Read IOPS	Read operations per second.	Stacked area on the IOPS chart.
Write IOPS	Write operations per second.	Stacked area on the IOPS chart.
Other IOPS	Metadata and other non-data operations per second.	Stacked area on the IOPS chart.
Read MB/s	Read throughput in megabytes per second.	Stacked area on the MB/s chart.
Write MB/s	Write throughput in megabytes per second.	Stacked area on the MB/s chart.
Cloud Throughput	Throughput in megabytes per second between the volume and the cloud capacity tier. This metric appears only for workloads that tier capacity to the cloud through FabricPool.	Stacked area on the MB/s chart.

The Read, Write, and Other components sum to the total IOPS or MB/s value. Hover over the chart to see each component's value, its percentage of the total, and the average I/O size (MB/s ÷ IOPS) per category.

Resource utilization metrics

Use the **Resource Utilization** section to review resource utilization metrics. The analyzer shows each resource serving the volume as an independent line. Resources do not stack.

Node metrics

Metric	Description
Node utilization	Composite utilization percentage for the node. Hover to see the CPU utilization, network utilization, and available headroom for each node.
CPU utilization	Percentage of node CPU capacity in use. The hover tooltip shows this value.
Network utilization	Percentage of node network capacity in use. The hover tooltip shows this value.
Headroom	Remaining node capacity available for additional workload. The hover tooltip shows this value.

Aggregate metrics

Metric	Description
Aggregate utilization	Disk busy percentage for the aggregate. Hover to see the disk busy value, number of volumes on the aggregate, and available headroom.
Disk busy	Percentage of time the aggregate's disks actively service I/O. The hover tooltip shows this value.
Volume count	Number of volumes currently hosted on the aggregate. The hover tooltip shows this value.
Headroom	Remaining aggregate capacity available for additional workload. The hover tooltip shows this value.

When a node or aggregate utilization line crosses the high-utilization warning threshold, which the chart marks with a dashed reference line, other workloads on that resource may be competing with the selected volume for I/O capacity.

Capacity metrics

Capacity View

Use the **Capacity Analysis** section to review Capacity View metrics after you select **Capacity View** from the View dropdown.

Metric	Description
Physical Capacity	Space consumed on disk after storage efficiency operations. This is the bottom stacked area in the chart. Physical + Available always equals the total volume size.
Available Capacity	Free space remaining in the volume. This is the top stacked area in the chart. It decreases as Physical Capacity grows.
Storage Efficiency Ratio	Overall efficiency ratio (logical data ÷ physical data). The hover tooltip shows this value. It reflects combined savings from deduplication, compression, and compaction.
Deduplication savings	Space saved by removing duplicate data blocks. The hover tooltip shows this value.
Compression savings	Space saved by compressing data inline. The hover tooltip shows this value.
Compaction savings	Space saved by packing multiple small blocks into a single physical block. The hover tooltip shows this value.

Snapshot View

Use Snapshot View when you want to review snapshot-specific metrics.

Metric	Description	Rendered as
Available snapshot	Pre-allocated space reserved for snapshots, configured as a percentage of the volume size.	Horizontal reference line.
Used snapshot	Actual space consumed by snapshot copies.	Area chart below the reserve line.

Hover over the chart to see the snapshot reserve size, snapshot space used and its percentage of the reserve, total snapshot count, and the age of the oldest snapshot. When snapshot consumption exceeds the reserve, snapshots begin consuming space from the volume’s data area.

Forecast behavior

Use the **Show Forecast** toggle in any analysis section when you want to project future values beyond the current time based on the observed trend. The following behaviors apply across all sections:

Aspect	Behavior
Projection window	Projects forward based on the selected time range. For a 2-hour view, the projection window is approximately 1 hour. For a 7-day view, the projection window extends several days.
Visual style	The analyzer renders forecast values as a dashed line. A vertical separator marks the boundary between actual data and projected values.
Threshold alerts	The analyzer displays an insight message when the forecast indicates that the volume will breach a warning or critical threshold, along with an estimated time to breach.
Trend basis	The projection uses the rate of change from the most recent portion of the selected time range. High volatility in recent data reduces forecast confidence.
Minimum data required	Capacity forecasting requires at least 10 days of historical data. When fewer than 30 days of capacity remain before the volume is full, the analyzer identifies the storage as almost full.

Related information

- Learn about NetApp Console local deployment Workload Analyzer
- Investigate high latency on a volume
- Investigate high throughput demand on a volume
- Investigate capacity growth on a volume

Copyright information

Trademark information