



Investigate performance issues

NetApp Console local deployment

NetApp
August 10, 2026

Table of Contents

- Investigate performance issues 1
 - Learn about NetApp Console local deployment Workload Analyzer 1
 - How Workload Analyzer identifies root causes 1
 - Access Workload Analyzer..... 1
 - Analyze latency, throughput, and storage capacity 1
 - Forecast capacity and performance trends 2
 - Investigate high latency on a volume in NetApp Console local deployment 2
 - Investigate high throughput on a volume in NetApp Console local deployment 4
 - Investigate capacity growth on a volume in NetApp Console local deployment 5
 - Workload Analyzer metrics in NetApp Console local deployment 7
 - Throughput components 8
 - Resource utilization metrics 8
 - Capacity metrics..... 9
 - Forecast behavior 10
 - Related information 10

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.

Latency type	Description	Common causes of elevated values
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.

Metric	Description
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

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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