



Resolving performance events

Active IQ Unified Manager 9.14

NetApp

November 11, 2024

Table of Contents

- Resolving performance events 1
 - Confirming that the latency is within the expected range 1
 - Review the impact of configuration changes on workload performance 1
 - Options for improving workload performance from the client-side 1
 - Check for client or network issues 2
 - Verify whether other volumes in the QoS policy group have unusually high activity 2
 - Move logical interfaces (LIFs) 3
 - Run storage efficiency operations at less busy times 3
 - Add disks and reallocate data 4
 - How enabling Flash Cache on a node can improve workload performance 4
 - How enabling Flash Pool on a storage aggregate can improve workload performance 5
 - MetroCluster configuration health check 5
 - MetroCluster configuration verification 6
 - Moving workloads to a different aggregate 6
 - Moving workloads to a different node 7
 - Moving workloads to an aggregate on a different node 9
 - Moving workloads to a node in a different HA pair 10
 - Moving workloads to another node in a different HA pair 12
 - Use QoS policy settings to prioritize the work on this node 14
 - Remove inactive volumes and LUNs 14
 - Add disks and perform aggregate layout reconstruction 15

Resolving performance events

You can use the suggested actions to try and resolve performance events on your own. The first three suggestions are always displayed, and the actions under the fourth suggestion are specific to the type of event displayed.

The **Help me do this** links provide additional information for each suggested action, including instructions for performing a specific action. Some of the actions may involve using Unified Manager, ONTAP System Manager, OnCommand Workflow Automation, ONTAP CLI commands, or a combination of these tools.

Confirming that the latency is within the expected range

When a cluster component is in contention, volume workloads that use it might have decreased response time (latency). You can review the latency of each victim workload on the component in contention to confirm that its actual latency is within its expected range. You can also click a volume name to view the historical data for the volume.

If the performance event is in the obsolete state, the latency of each victim involved in the event might have returned within its expected range.

Review the impact of configuration changes on workload performance

Configuration changes on the cluster, such as a failed disk, HA failover, or a moved volume, could negatively impact volume performance and cause increased latency.

In Unified Manager, you can review the Workload Analysis page to see when a recent configuration change occurred and compare it to the operations and latency (response time) to see whether there was a change in activity for the selected volume workload.

The performance pages of Unified Manager can only detect a limited number of change events. The health pages provide alerts for other events caused by configuration changes. You can search for the volume in Unified Manager to see the event history.

Options for improving workload performance from the client-side

You can check your client workloads, such as applications or databases, that are sending I/O to volumes involved in a performance event to determine if a client-side change might correct the event.

When the clients that are connected to volumes on a cluster increase their I/O requests, the cluster must work harder to meet the demand. If you know which clients have a high number of I/O requests to a particular volume on the cluster, you can improve cluster performance by adjusting the number of clients accessing the volume or decreasing the amount of I/O to the volume. You can also apply or increase a limit on the QoS policy group of which the volume is a member.

You can investigate clients and their applications to determine whether the clients are sending more I/O than

usual, which might be causing contention on a cluster component. On the Event details page, the System Diagnosis section displays the top volume workloads using the component in contention. If you know which client is accessing a particular volume, you can go to the client to determine whether the client hardware or an application is not operating as expected or is doing more work than usual.

In a MetroCluster configuration, write requests to a volume on a local cluster are mirrored to a volume on the remote cluster. Keeping the source volume on the local cluster in sync with the destination volume on the remote cluster can also increase the demand of both clusters in the MetroCluster configuration. By reducing write requests to these mirrored volumes, the clusters can perform fewer sync operations, which reduces the performance impact on other workloads.

Check for client or network issues

When the clients that are connected to volumes on a cluster increase their I/O requests, the cluster must work harder to meet the demand. The increased demand on the cluster can put a component in contention, increase the latency of workloads that use it, and trigger an event in Unified Manager.

On the Event details page, the System Diagnosis section displays the top volume workloads using the component in contention. If you know which client is accessing a particular volume, you can go to the client to determine whether the client hardware or an application is not operating as expected or is doing more work than usual. You might need to contact your client administrator or application vendor for assistance.

You can check your network infrastructure to determine whether there are hardware issues, bottlenecks, or competing workloads that might have caused I/O requests between the cluster and connected clients to perform slower than expected. You might need to contact your network administrator for assistance.

Verify whether other volumes in the QoS policy group have unusually high activity

You can review the workloads in the Quality of Service (QoS) policy group with the highest change in activity to determine whether more than one workload caused the event. You can also see whether other workloads are still exceeding the set throughput limit or whether they are back within their expected range of activity.

On the Event details page, in the System Diagnosis section, you can sort the workloads by peak deviation in activity to display the workloads with the highest change in activity at the top of the table. These workloads might be the “bullies” whose activity exceeded the set limit and might have caused the event.

You can navigate to the Workload Analysis page for each volume workload to review its IOPS activity. If the workload has periods of very high operations activity, it might have contributed to the event. You can change the policy group settings for the workload or move the workload to a different policy group.

You can use ONTAP System Manager or the ONTAP CLI commands to manage policy groups, as follows:

- Create a policy group.
- Add or remove workloads in a policy group.
- Move a workload between policy groups.
- Change the throughput limit of a policy group.

Move logical interfaces (LIFs)

Moving logical interfaces (LIFs) to a less busy port can help improve load balancing, assist with maintenance operations and performance tuning, and reduce indirect access.

Indirect access can reduce system efficiency. It occurs when a volume workload is using different nodes for network processing and data processing. To reduce indirect access, you can rearrange LIFs, which involves moving LIFs to use the same node for network processing and data processing. You can configure load balancing to have ONTAP automatically move busy LIFs to a different port or you can move a LIF manually.

Benefits	Considerations
<ul style="list-style-type: none">• Improve load balancing.• Reduce indirect access.	 <p>When moving a LIF connected to CIFS shares, clients accessing the CIFS shares are disconnected. Any read or write requests to the CIFS shares are disrupted.</p>

You use the ONTAP commands to configure load balancing. For more information, see the ONTAP networking documentation.

You use ONTAP System Manager and the ONTAP CLI commands to move LIFs manually.

Run storage efficiency operations at less busy times

You can modify the policy or schedule that handles storage efficiency operations to run when the impacted volume workloads are less busy.

Storage efficiency operations can use a high amount of cluster CPU resources and become a bully to the volumes on which the operations are being run. If the victim volumes have high activity at the same time when the storage efficiency operations are run, their latency can increase and trigger an event.

On the Event details page, the System Diagnosis section displays workloads in the QoS policy group by peak deviation in activity to identify the bully workloads. If you see “storage efficiency” displayed near the top of the table, these operations are bullying the victim workloads. By modifying the efficiency policy or schedule to run when these workloads are less busy, you can prevent the storage efficiency operations from causing contention on a cluster.

You can use ONTAP System Manager to manage efficiency policies. You can use the ONTAP commands to manage efficiency policies and schedules.

What storage efficiency is

Storage efficiency enables you to store the maximum amount of data for the lowest cost and accommodates rapid data growth while consuming less space. NetApp strategy for storage efficiency is based on the built-in foundation of storage virtualization and unified storage provided by its core ONTAP operating system and Write Anywhere File Layout (WAFL) file system.

Storage efficiency includes using technologies such as thin provisioning, Snapshot copy, deduplication, data

compression, FlexClone, thin replication with SnapVault and volume SnapMirror, RAID-DP, Flash Cache, Flash Pool aggregate, and FabricPool-enabled aggregates which help to increase storage utilization and decrease storage costs.

The unified storage architecture allows you to efficiently consolidate a storage area network (SAN), network-attached storage (NAS), and secondary storage on a single platform.

High-density disk drives, such as serial advanced technology attachment (SATA) drives configured within Flash Pool aggregate or with Flash Cache and RAID-DP technology, increase efficiency without affecting performance and resiliency.

A FabricPool-enabled aggregate includes an all SSD aggregate or HDD aggregate (starting with ONTAP 9.8) as the local performance tier and an object store that you specify as the cloud tier. Configuring FabricPool helps you manage which storage tier (the local tier or the cloud tier) data should be stored based on whether the data is frequently accessed.

Technologies such as thin provisioning, Snapshot copy, deduplication, data compression, thin replication with SnapVault and volume SnapMirror, and FlexClone offer better savings. You can use these technologies individually or together to achieve maximum storage efficiency.

Add disks and reallocate data

You can add disks to an aggregate to increase the storage capacity and the performance of that aggregate. After adding the disks, you will see an improvement in read performance only after reallocating the data across the disks you added.

You can use these instructions when Unified Manager has received aggregate events triggered by dynamic thresholds or by system-defined performance thresholds:

- When you have received a dynamic threshold event, on the Event details page, the cluster component icon that represents the aggregate in contention is highlighted red.

Beneath the icon, in parentheses, is the name of the aggregate, which identifies the aggregate to which you can add disks.

- When you have received a system-defined threshold event, on the Event details page, the event description text lists the name of the aggregate that is having the problem.

You can add disks and reallocate data on this aggregate.

The disks you add to the aggregate must already exist in the cluster. If the cluster does not have extra disks available, you might need to contact your administrator or purchase more disks. You can use ONTAP System Manager or the ONTAP commands to add disks to an aggregate.



You should reallocate data when using HDD and Flash Pool aggregates only. Do not reallocate data on SSD or FabricPool aggregates.

How enabling Flash Cache on a node can improve workload performance

You can improve workload performance by enabling Flash Cache™ intelligent data

caching on each node in the cluster.

A Flash Cache module, or Performance Acceleration Module PCIe-based memory module, optimizes the performance of random read-intensive workloads by functioning as an intelligent external read cache. This hardware works in tandem with the WAFL External Cache software component of ONTAP.

In Unified Manager, on the Event details page, the cluster component icon that represents the aggregate in contention is highlighted red. Beneath the icon, in parentheses, is the name of the aggregate, which identifies the aggregate. You can enable Flash Cache on the node on which the aggregate resides.

You can use ONTAP System Manager or the ONTAP commands to see whether Flash Cache is installed or enabled, and to enable it if not already enabled. The following command indicates whether Flash Cache is enabled on a specific node: **cluster::> run local options flexscale.enable**

For more information about Flash Cache and the requirements for using it, see the following technical report:

[Technical Report 3832: Flash Cache Best Practices Guide](#)

How enabling Flash Pool on a storage aggregate can improve workload performance

You can improve workload performance by enabling the Flash Pool feature on an aggregate. A Flash Pool is an aggregate that incorporates both HDDs and SSDs. The HDDs are used for primary storage and the SSDs provide a high-performance read and write cache to boost aggregate performance.

In Unified Manager, the Event details page displays the name of the aggregate in contention. You can use ONTAP System Manager or the ONTAP commands to see whether Flash Pool is enabled for an aggregate. If you have SSDs installed, you can use the command-line interface to enable it. If you have SSDs installed, you can run the following command on the aggregate to see whether Flash Pool is enabled: **cluster::> storage aggregate show -aggregate aggr_name -field hybrid-enabled**

In this command, *aggr_name* is the name of the aggregate, such as the aggregate in contention.

For more information about Flash Pool and the requirements for using it, see the *Clustered Data ONTAP Physical Storage Management Guide*.

MetroCluster configuration health check

You can use Unified Manager to review the health of the clusters in a MetroCluster configuration over IP or FC. The health status and events help you determine whether there are hardware or software issues that might be impacting the performance of your workloads.

If you configure Unified Manager to send email alerts, you can check your email for any health issues on the local or remote cluster that might have contributed to a performance event. In the Unified Manager GUI, you can select **Event Management** to see a list current events and then use the filters to display MetroCluster configuration events only.

For more information, see [Checking the health of clusters in a MetroCluster configuration](#)

MetroCluster configuration verification

You can prevent performance issues for mirrored workloads in a MetroCluster over FC and IP configurations by ensuring that the MetroCluster configuration is set up correctly. You can also improve workload performance by changing the configuration or upgrading software or hardware components.

Refer to [MetroCluster Documentation](#) for the instructions of setting up the clusters in the MetroCluster configuration, including the Fibre Channel (FC) switches, cables, and inter-switch links (ISLs). It also helps you configure the MetroCluster software so that the local and remote clusters can communicate with mirror volume data. For information specific to your MetroCluster over IP setup, see [Install a MetroCluster IP configuration](#).

You can compare your MetroCluster configuration to the requirements in [MetroCluster Documentation](#) to determine whether changing or upgrading components in your MetroCluster configuration might improve workload performance. This comparison can help you answer the following questions:

- Are the controllers appropriate for your workloads?
- Do you need to upgrade your ISL bundles to a larger bandwidth to handle more throughput?
- Can you adjust the buffer-to-buffer credits (BBC) on your switches to increase the bandwidth?
- If your workloads have high write throughput to solid state drive (SSD) storage, do you need to upgrade your FC-to-SAS bridges to accommodate the throughput?

Related information

- For information about replacing or upgrading MetroCluster components, see the [MetroCluster Documentation](#).
- For information about upgrading controllers see [Upgrading controllers in a MetroCluster FC configuration using switchover and switchback](#) and [Upgrading controllers in a MetroCluster IP configuration using switchover and switchback](#)

Moving workloads to a different aggregate

You can use Unified Manager to help identify an aggregate that is less busy than the aggregate where your workloads currently reside, and then you can move selected volumes or LUNs to that aggregate. Moving high performing workloads to a less busy aggregate, or an aggregate with flash storage enabled, allows the workload to perform more efficiently.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have recorded the name of the aggregate that is currently having a performance issue.
- You must have recorded the date and time at which the aggregate received the event.
- Unified Manager must have collected and analyzed a month or more of performance data.

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized aggregate:

- The aggregates on the same cluster that are less utilized
- The highest-performing volumes on the current aggregate

Steps

1. Identify the aggregate in the cluster that is the least utilized:

- a. From the **Event** details page, click the name of the cluster on which the aggregate resides.

The cluster details are displayed in the Performance/Cluster Landing page.

- b. On the **Summary** page, click **Aggregates** from the **Managed Objects** pane.

The list of aggregates on this cluster are displayed.

- c. Click the **Utilization** column to sort the aggregates by least utilized.

You can also identify those aggregates that have the greatest **Free Capacity**. This provides a list of potential aggregates to which you might want to move workloads.

- d. Write down the name of the aggregate to which you want to move the workloads.

2. Identify the high-performing volumes from the aggregate that received the event:

- a. Click the aggregate that is having the performance issue.

The aggregate details are displayed in the Performance/Aggregate Explorer page.

- b. From the **Time Range** selector, select **Last 30 Days**, and then click **Apply Range**.

This enables you to view a longer performance history period than the default 72 hours. You want to move a volume that is using a lot of resources on a consistent basis, not just over the past 72 hours.

- c. From the **View and Compare** control, select **Volumes on this Aggregate**.

A list of FlexVol volumes and FlexGroup constituent volumes on this aggregate are displayed.

- d. Sort the volumes by highest MB/s, and then by highest IOPS, to see the highest performing volumes.

- e. Write down the names of the volumes that you want to move to a different aggregate.

3. Move the high-performing volumes to the aggregate you identified as having low utilization.

You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, check to see whether you are receiving the same type of events from this node or aggregate.

Moving workloads to a different node

You can use Unified Manager to help identify an aggregate on a different node that is less busy than the node on which your workloads are currently running, and then you can move selected volumes to that aggregate. Moving high-performing workloads to an aggregate on a less busy node allows the workloads on both nodes to perform more efficiently.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have recorded the name of the node that is currently having a performance issue.
- You must have recorded the date and time at which the node received the performance event.
- Unified Manager must have collected and analyzed performance data for a month or longer.

This procedure helps you to identify the following resources so that you can move high-performing workloads to a lower utilized node:

- The nodes on the same cluster that have the greatest free performance capacity
- The aggregates on the new node that have the greatest free performance capacity
- The highest-performing volumes on the current node

Steps

1. Identify a node in the cluster that has the greatest free performance capacity:
 - a. On the **Event Details** page, click the name of the cluster on which the node resides.

The cluster details are displayed in the Performance/Cluster Landing page.
 - b. On the **Summary** tab, click **Nodes** from the **Managed Objects** pane.

The list of nodes on this cluster are displayed.
 - c. Click the **Performance Capacity Used** column to sort the nodes by least percentage used.

This provides a list of potential nodes to which you might want to move workloads.
 - d. Write down the name of the node to which you want to move the workloads.
2. Identify an aggregate on the new node that is the least utilized:
 - a. In the left navigation pane, click **Storage > Aggregates** and select **Performance > All Aggregates** from the View menu.

The Performance: All Aggregates view is displayed.
 - b. Click **Filtering**, select **Node** from the left drop-down menu, type the name of the node in the text field, and then click **Apply Filter**.

The Performance: All Aggregates view is redisplayed with the list of aggregates that are available on this node.
 - c. Click the **Performance Capacity Used** column to sort the aggregates by least used.

This provides a list of potential aggregates to which you might want to move workloads.
 - d. Write down the name of the aggregate to which you want to move the workloads.
3. Identify the high-performing workloads from the node that received the event:
 - a. Return to the **Event Details** page for the event.
 - b. In the **Affected Volumes** field, click the link for the number of volumes.

The Performance: All Volumes view is displayed with a filtered list of the volumes on that node.

- c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

This provides a list of potential volumes that you may want to move.

- d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.
4. Move the volumes to the aggregates that you identified as having greatest free performance capacity on the new node.

You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, you can check whether you are receiving the same type of events from this node or aggregate.

Moving workloads to an aggregate on a different node

You can use Unified Manager to help identify an aggregate on a different node that is less busy than the node where your workloads are currently running, and then you can move selected volumes to that aggregate. Moving high-performing workloads to an aggregate on a less busy node allows workloads on both nodes to perform more efficiently.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have recorded the name of the node that is currently having a performance issue.
- You must have recorded the date and time at which the node received the performance event.
- Unified Manager must have collected and analyzed a month or more of performance data.

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized node:

- The nodes on the same cluster that are less utilized
- The aggregates on the new node that are the least utilized
- The highest-performing volumes on the current node

Steps

1. Identify a node in the cluster that is the least utilized:
 - a. From the **Event** details page, click the name of the cluster on which the node resides.

The cluster details are displayed in the Performance/Cluster Landing page.
 - b. On the **Summary** page, click **Nodes** from the **Managed Objects** pane.

The list of nodes on this cluster are displayed.
 - c. Click the **Utilization** column to sort the nodes by least utilized.

You can also identify those nodes that have the greatest **Free Capacity**. This provides a list of potential nodes to which you might want to move workloads.

d. Write down the name of the node to which you want to move the workloads.

2. Identify an aggregate on the new node that is the least utilized:

a. In the left navigation pane, click **Storage > Aggregates** and select **Performance > All Aggregates** from the View menu.

The Performance: All Aggregates view is displayed.

b. Click **Filtering**, select **Node** from the left drop-down menu, type the name of the node in the text field, and then click **Apply Filter**.

The Performance: All Aggregates view is redisplayed with the list of aggregates that are available on this node.

c. Click the **Utilization** column to sort the aggregates by least utilized.

You can also identify those aggregates that have the greatest **Free Capacity**. This provides a list of potential aggregates to which you might want to move workloads.

d. Write down the name of the aggregate to which you want to move the workloads.

3. Identify the high-performing workloads from the node that received the event:

a. Return to the **Event** details page for the event.

b. In the **Affected Volumes** field, click the link for the number of volumes.

The Performance: All Volumes view is displayed with a filtered list of the volumes on that node.

c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

This provides a list of potential volumes that you may want to move.

d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.

4. Move the volumes to the aggregates you identified as having low utilization on the new node.

You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, check to see whether you are receiving the same type of events from this node or aggregate.

Moving workloads to a node in a different HA pair

You can use Unified Manager to help identify an aggregate on a node in a different high-availability (HA) pair that has more free performance capacity than the HA pair where your workloads are currently running. Then you can move selected volumes to aggregates on the new HA pair.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- Your cluster must consist of a minimum of two HA pairs

You cannot use this remediation process if you have only one HA pair in your cluster.

- You must have recorded the names of the two nodes in the HA pair that are currently having a performance issue.
- You must have recorded the date and time at which the nodes received the performance event.
- Unified Manager must have collected and analyzed performance data for a month or longer.

Moving high-performing workloads to an aggregate on a node with more free performance capacity allows workloads on both nodes to perform more efficiently. This procedure helps you to identify the following resources so that you can move high-performing workloads to a node that has more free performance capacity on a different HA pair:

- The nodes in a different HA pair on the same cluster that have the greatest free performance capacity
- The aggregates on the new nodes that have the greatest free performance capacity
- The highest-performing volumes on the current nodes

Steps

1. Identify the nodes that are part of a different HA pair on the same cluster:

- a. On the **Event Details** page, click the name of the cluster on which the nodes reside.

The cluster details are displayed in the Performance/Cluster Landing page.

- b. On the **Summary** page, click **Nodes** from the **Managed Objects** pane.

The list of nodes on this cluster is displayed in the Performance: All Nodes view.

- c. Write down the names of the nodes that are in different HA pairs from the HA pair that is currently having a performance issue.

2. Identify a node in the new HA pair that has the greatest free performance capacity:

- a. On the **Performance: All Nodes** view, click the **Performance Capacity Used** column to sort the nodes by least percentage used.

This provides a list of potential nodes to which you might want to move workloads.

- b. Write down the name of the node on a different HA pair to which you want to move the workloads.

3. Identify an aggregate on the new node that has the greatest free performance capacity:

- a. On the **Performance: All Nodes** view, click the node.

The node details are displayed in the Performance/Node Explorer page.

- b. In the **View and Compare** menu, select **Aggregates on this Node**.

The aggregates on this node are displayed in the grid.

- c. Click the **Performance Capacity Used** column to sort the aggregates by least used.

This provides a list of potential aggregates to which you might want to move workloads.

- d. Write down the name of the aggregate to which you want to move the workloads.
4. Identify the high-performing workloads from the nodes that received the event:
 - a. Return to the **Event** details page for the event.
 - b. In the **Affected Volumes** field, click the link for the number of volumes for the first node.

The Performance: All Volumes view is displayed with a filtered list of the volumes on that node.

- c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

This provides a list of potential volumes that you might want to move.

- d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.
 - e. Perform steps 4c and 4d for the second node that was part of this event to identify possible volumes that you want to move from that node as well.
5. Move the volumes to the aggregates that you identified as having greatest free performance capacity on the new node.

You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, you can check whether you are receiving the same type of events from this node or aggregate.

Moving workloads to another node in a different HA pair

You can use Unified Manager to help identify an aggregate on a node in a different HA pair that is less busy than the HA pair where your workloads are currently running. Then you can move selected volumes to aggregates on the new HA pair. Moving high-performing workloads to an aggregate on a less busy node allows workloads on both nodes to perform more efficiently.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- Your cluster must consist of a minimum of two HA pairs; you cannot use this remediation process if you have only one HA pair in your cluster.
- You must have recorded the names of the two nodes in the HA pair that are currently having the performance issue.
- You must have recorded the date and time at which the nodes received the performance event.
- Unified Manager must have collected and analyzed a month or more of performance data.

These steps help you identify the following resources so that you can move high-performing workloads to a lower utilized node on a different HA pair:

- The nodes in a different HA pair on the same cluster that are less utilized
- The aggregates on the new nodes that are the least utilized

- The highest-performing volumes on the current nodes

Steps

1. Identify the nodes that are part of a different HA pair on the same cluster:

- a. In the left navigation pane, click **Storage > Clusters** and select **Performance > All Clusters** from the View menu.

The Performance: All Clusters view is displayed.

- b. Click the number in the **Node Count** field for the current cluster.

The Performance: All Nodes view is displayed.

- c. Write down the names of the nodes that are in different HA pairs from the HA pair that is currently having the performance issue.

2. Identify a node in the new HA pair that is the least utilized:

- a. Click the **Utilization** column to sort the nodes by least utilized.

You can also identify those nodes that have the greatest **Free Capacity**. This provides a list of potential nodes to which you might want to move workloads.

- b. Write down the name of the node to which you want to move the workloads.

3. Identify an aggregate on the new node that is the least utilized:

- a. In the left navigation pane, click **Storage > Aggregates** and select **Performance > All Aggregates** from the View menu.

The Performance: All Aggregates view is displayed.

- b. Click **Filtering**, select **Node** from the left drop-down menu, type the name of the node in the text field, and then click **Apply Filter**.

The Performance: All Aggregates view is redisplayed with the list of aggregates that are available on this node.

- c. Click the **Utilization** column to sort the aggregates by least utilized.

You can also identify those aggregates that have the greatest **Free Capacity**. This provides a list of potential aggregates to which you might want to move workloads.

- d. Write down the name of the aggregate to which you want to move the workloads.

4. Identify the high-performing workloads from the nodes that received the event:

- a. Return to the **Event** details page for the event.

- b. In the **Affected Volumes** field, click the link for the number of volumes for the first node.

The Performance: All Volumes view is displayed with a filtered list of the volumes on that node.

- c. Click the **Total Capacity** column to sort the volumes by the largest allocated space.

This provides a list of potential volumes that you might want to move.

- d. Write down the names of the volumes that you want to move, and the names of the current aggregates on which they reside.

- e. Perform steps 4c and 4d for the second node that was part of this event to identify possible volumes that you want to move from that node as well.
5. Move the volumes to the aggregates you identified as having low utilization on the new node.

You can perform the move operation by using ONTAP System Manager, OnCommand Workflow Automation, ONTAP commands, or a combination of these tools.

After a few days, check to see whether you are receiving the same type of events from this node or aggregate.

Use QoS policy settings to prioritize the work on this node

You can set a limit on a QoS policy group to control the I/O per second (IOPS) or MBps throughput limit for the workloads it contains. If workloads are in a policy group with no set limit, such as the default policy group, or the set limit does not meet your needs, you can increase the set limit or move the workloads to a new or existing policy group that has the desired limit.

If a performance event on a node is caused by workloads overusing the node resources, the event description on the Event details page displays a link to the list of volumes involved. In the Performance/Volumes page, you can sort the affected volumes by IOPS and MBps to see which workloads have the highest usage that might have contributed to the event.

By assigning the volumes that are overusing the node resources to a more restrictive policy group setting, the policy group throttles the workloads to restrict their activity, which can reduce the use of the resources on that node.

You can use ONTAP System Manager or the ONTAP commands to manage policy groups, including the following tasks:

- Creating a policy group
- Adding or removing workloads in a policy group
- Moving a workload between policy groups
- Changing the throughput limit of a policy group

Remove inactive volumes and LUNs

When aggregate free space has been identified as an issue, you can search for unused volumes and LUNs and delete them from the aggregate. This can help to alleviate the low disk space issue.

If a performance event on an aggregate is caused by low disk space, there are a few ways you can determine which volumes and LUNs are no longer being used.

To identify unused volumes:

- On the Event details page, the **Affected Objects Count** field provides a link that displays the list of affected volumes.

Click the link to display the volumes on the Performance: All Volumes view. From there you can sort the affected volumes by **IOPS** to see which volumes have not been active.

To identify unused LUNs:

1. From the Event details page, write down the name of the aggregate on which the event occurred.
2. In the left navigation pane, click **Storage > LUNs** and select **Performance > All LUNs** from the View menu.
3. Click **Filtering**, select **Aggregate** from the left drop-down menu, type the name of the aggregate in the text field, and then click **Apply Filter**.
4. Sort the resulting list of affected LUNs by **IOPS** to view the LUNs that are not active.

After you have identified the unused volumes and LUNs, you can use ONTAP System Manager or the ONTAP commands to delete those objects.

Add disks and perform aggregate layout reconstruction

You can add disks to an aggregate to increase the storage capacity and the performance of that aggregate. After adding the disks, you only see an improvement in performance after reconstructing the aggregate.

When you receive a system-defined threshold event on the Event details page, the event description text lists the name of the aggregate that is having the problem. You can add disks and reconstruct data on this aggregate.

The disks you add to the aggregate must already exist in the cluster. If the cluster does not have extra disks available, you might need to contact your administrator or purchase more disks. You can use ONTAP System Manager or the ONTAP commands to add disks to an aggregate.

[Technical Report 3838: Storage Subsystem Configuration Guide](#)

Copyright information

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

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

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

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

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

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

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

Trademark information

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