



Collecting data and monitoring workload performance

Active IQ Unified Manager 9.10

NetApp
June 11, 2024

Table of Contents

- Collecting data and monitoring workload performance 1
 - Types of workloads monitored by Unified Manager 1
 - Workload performance measurement values 2
 - What the expected range of performance is 4
 - How the latency forecast is used in performance analysis 5
 - How Unified Manager uses workload latency to identify performance issues 6
 - How cluster operations can affect workload latency 7
 - Performance monitoring of MetroCluster configurations 7
 - What performance events are 10

Collecting data and monitoring workload performance

Unified Manager collects and analyzes workload activity every 5 minutes to identify performance events, and it detects configuration changes every 15 minutes. It retains a maximum of 30 days of 5-minute historical performance and event data, and it uses this data to forecast the expected latency range for all monitored workloads.

Unified Manager must collect a minimum of 3 days of workload activity before it can begin its analysis and before the latency forecast for I/O response time can be displayed on the Workload Analysis page and in the Event details page. While this activity is being collected, the latency forecast does not display all changes occurring from workload activity. After collecting 3 days of activity, Unified Manager adjusts the latency forecast every 24 hours at 12:00 a.m., to reflect workload activity changes and establish a more accurate dynamic performance threshold.

During the first 4 days that Unified Manager is monitoring a workload, if more than 24 hours have passed since the last data collection, the latency charts will not display the latency forecast for that workload. Events detected prior to the last collection are still available.



Daylight savings time (DST) changes the system time, which alters the latency forecast of performance statistics for monitored workloads. Unified Manager immediately begins to correct the latency forecast, which takes approximately 15 days to complete. During this time you can continue to use Unified Manager, but, since Unified Manager uses the latency forecast to detect dynamic events, some events might not be accurate. Events detected prior to the time change are not affected.

Types of workloads monitored by Unified Manager

You can use Unified Manager to monitor the performance of two types of workloads: user-defined and system-defined.

• **User-defined workloads**

The I/O throughput from applications to the cluster. These are processes involved in read and write requests. A volume, LUN, NFS share, SMB/CIFS share, and a workload is a user-defined workload.



Unified Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

If one or more of the following is true for a workload, it cannot be monitored by Unified Manager:

- It is a data protection (DP) copy in read-only mode. (DP volumes are monitored for user-generated traffic.)
- It is an offline data clone.
- It is a mirrored volume in a MetroCluster configuration.

• **System-defined workloads**

The internal processes involved with storage efficiency, data replication, and system health, including:


- Storage efficiency, such as deduplication
- Disk health, which includes RAID reconstruct, disk scrubbing, and so on
- Data replication, such as SnapMirror copies
- Management activities
- File system health, which includes various WAFL activities
- File system scanners, such as WAFL scan
- Copy offload, such as offloaded storage efficiency operations from VMware hosts
- System health, such as volume moves, data compression, and so on
- Unmonitored volumes

Performance data for system-defined workloads is displayed in the GUI only when the cluster component used by these workloads is in contention. For example, you cannot search for the name of a system-defined workload to view its performance data in the GUI.



Workload performance measurement values

Unified Manager measures the performance of workloads on a cluster based on historical and expected statistical values, which form the latency forecast of values for the workloads. It compares the actual workload statistical values to the latency forecast to determine when workload performance is too high or too low. A workload that is not performing as expected triggers a dynamic performance event to notify you.

In the following illustration, the actual value, in red, represents the actual performance statistics in the time frame. The actual value has crossed the performance threshold, which is the upper bounds of the latency forecast. The peak is the highest actual value in the time frame. The deviation measures the change between the expected values (the forecast) and the actual values, while the peak deviation indicates the largest change between the expected values and the actual values.



The following table lists the workload performance measurement values.

Measurement	Description
Activity	<p>The percentage of the QoS limit used by the workloads in the policy group.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <p>If Unified Manager detects a change to a policy group, such as adding or removing a volume or changing the QoS limit, the actual and expected values might exceed 100% of the set limit. If a value exceeds 100% of the set limit it is displayed as >100%. If a value is less than 1% of the set limit it is displayed as <1%.</p> </div>
Actual	<p>The measured performance value at a specific time for a given workload.</p>
Deviation	<p>The change between the expected values and the actual values. It is the ratio of the actual value minus the expected value to the upper value of the expected range minus the expected value.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <p>A negative deviation value indicates that workload performance is lower than expected, while a positive deviation value indicates that workload performance is higher than expected.</p> </div>
Expected	<p>The expected values are based on the analysis of historical performance data for a given workload. Unified Manager analyzes these statistical values to determine the expected range (latency forecast) of values.</p>
Latency Forecast (Expected Range)	<p>The latency forecast is a prediction of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Unified Manager triggers a dynamic performance event.</p>
Peak	<p>The maximum value measured over a period of time.</p>
Peak Deviation	<p>The maximum deviation value measured over a period of time.</p>

Measurement	Description
Queue Depth	The number of pending I/O requests that are waiting at the interconnect component.
Utilization	For the network processing, data processing, and aggregate components, the percentage of busy time to complete workload operations over a period of time. For example, the percentage of time for the network processing or data processing components to process an I/O request or for an aggregate to fulfill a read or write request.
Write Throughput	The amount of write throughput, in Megabytes per second (MB/s), from workloads on a local cluster to the partner cluster in a MetroCluster configuration.

What the expected range of performance is

The latency forecast is a prediction of what the upper and lower performance values are expected to be at a specific time. For the workload latency, the upper values form the performance threshold. When the actual value crosses the performance threshold, Unified Manager triggers a dynamic performance event.

For example, during regular business hours between 9:00 a.m. and 5:00 p.m., most employees might check their email between 9:00 a.m. and 10:30 a.m. The increased demand on the email servers means an increase in workload activity on the back-end storage during this time. Employees might notice slow response time from their email clients.

During the lunch hour between 12:00 p.m. and 1:00 p.m. and at the end of the work day after 5:00 p.m., most employees are likely away from their computers. The demand on the email servers typically decreases, also decreasing the demand on back-end storage. Alternatively, there could be scheduled workload operations, such as storage backups or virus scanning, that start after 5:00 p.m. and increase activity on the back-end storage.

Over several days, the increase and decrease in workload activity determines the expected range (latency forecast) of activity, with upper and lower boundaries for a workload. When the actual workload activity for an object is outside the upper or lower boundaries, and remains outside the boundaries for a period of time, this might indicate that the object is being overused or underused.

How the latency forecast is formed

Unified Manager must collect a minimum of 3 days of workload activity before it can begin its analysis and before the latency forecast for I/O response time can be displayed in the GUI. The minimum required data collection does not account for all changes occurring from workload activity. After collecting the first 3 days of activity, Unified Manager adjusts the latency forecast every 24 hours at 12:00 a.m. to reflect workload activity changes and establish a more accurate dynamic performance threshold.



Daylight savings time (DST) changes the system time, which alters the latency forecast of performance statistics for monitored workloads. Unified Manager immediately begins to correct the latency forecast, which takes approximately 15 days to complete. During this time you can continue to use Unified Manager, but, since Unified Manager uses the latency forecast to detect dynamic events, some events might not be accurate. Events detected prior to the time change are not affected.

How the latency forecast is used in performance analysis

Unified Manager uses the latency forecast to represent the typical I/O latency (response time) activity for your monitored workloads. It alerts you when the actual latency for a workload is above the upper bounds of the latency forecast, which triggers a dynamic performance event, so that you can analyze the performance issue and take corrective action for resolving it.

The latency forecast sets the performance baseline for the workload. Over time, Unified Manager learns from past performance measurements to forecast the expected performance and activity levels for the workload. The upper boundary of the expected range establishes the dynamic performance threshold. Unified Manager uses the baseline to determine when the actual latency is above or below a threshold, or outside the bounds of their expected range. The comparison between the actual values and the expected values creates a performance profile for the workload.

When the actual latency for a workload exceeds the dynamic performance threshold, due to contention on a cluster component, the latency is high and the workload performs more slowly than expected. The performance of other workloads that share the same cluster components might also be slower than expected.

Unified Manager analyzes the threshold crossing event and determines whether the activity is a performance event. If the high workload activity remains consistent for a long period of time, such as several hours, Unified Manager considers the activity to be normal and dynamically adjusts the latency forecast to form the new dynamic performance threshold.

Some workloads might have consistently low activity, where the latency forecast for latency does not have a high rate of change over time. To minimize the number of events during analysis of performance events, Unified Manager triggers an event only for low-activity volumes whose operations and latencies are much higher than expected.



In this example, the latency for a volume has a latency forecast, in gray, of 3.5 milliseconds per operation

(ms/op) at its lowest and 5.5 ms/op at its highest. If the actual latency, in blue, suddenly increases to 10 ms/op, due to an intermittent spike in network traffic or contention on a cluster component, it is then above the latency forecast and has exceeded the dynamic performance threshold.

When network traffic has decreased, or the cluster component is no longer in contention, the latency returns within the latency forecast. If the latency remains at or above 10 ms/op for a long period of time, you might need to take corrective action to resolve the event.

How Unified Manager uses workload latency to identify performance issues

The workload latency (response time) is the time it takes for a volume on a cluster to respond to I/O requests from client applications. Unified Manager uses the latency to detect and alert you to performance events.


A high latency means that requests from applications to a volume on a cluster are taking longer than usual. The cause of the high latency could be on the cluster itself, due to contention on one or more cluster components. High latency could also be caused by issues outside of the cluster, such as network bottlenecks, issues with the client hosting the applications, or issues with the applications themselves.



Unified Manager only monitors the workload activity on the cluster. It does not monitor the applications, the clients, or the paths between the applications and the cluster.

Operations on the cluster, such as making backups or running deduplication, that increase their demand of cluster components shared by other workloads can also contribute to high latency. If the actual latency exceeds the dynamic performance threshold of the expected range (latency forecast), Unified Manager analyzes the event to determine whether it is a performance event that you might need to resolve. The latency is measured in milliseconds per operation (ms/op).

On the Latency Total chart in the Workload Analysis page, you can view an analysis of the latency statistics to see how the activity of individual processes, such as read and write requests, compares to the overall latency statistics. The comparison helps you determine which operations have the highest activity or whether specific operations have abnormal activity that is impacting the latency for a volume. When analyzing performance events, you can use the latency statistics to determine whether an event was caused by an issue on the cluster. You can also identify the specific workload activities or cluster components that are involved in the event.



This example shows the Latency chart . The actual response time (latency) activity is a blue line and the latency forecast (expected range) is green.




There can be gaps in the blue line if Unified Manager was unable to gather data. This can occur because the cluster or volume was unreachable, Unified Manager was turned off during that time, or the collection was taking longer than the 5 minute collection period.

How cluster operations can affect workload latency

Operations (IOPS) represent the activity of all user-defined and system-defined workloads on a cluster. The IOPS statistics help you determine whether cluster processes, such as making backups or running deduplication, are impacting workload latency (response time) or might have caused, or contributed to, a performance event.

When analyzing performance events, you can use the IOPS statistics to determine whether a performance event was caused by an issue on the cluster. You can identify the specific workload activities that might have been the main contributors to the performance event. IOPS are measured in operations per second (ops/sec).



This example shows the IOPS chart. The actual operations statistics is a blue line and the IOPS forecast of operations statistics is green.



In some cases where a cluster is overloaded, Unified Manager might display the message `Data collection is taking too long on Cluster cluster_name`. This means that not enough statistics have been collected for Unified Manager to analyze. You need to reduce the resources the cluster is using so that statistics can be collected.

Performance monitoring of MetroCluster configurations

Unified Manager enables you to monitor the write throughput between clusters in a MetroCluster configuration to identify workloads with a high amount of write throughput. If these high-performing workloads are causing other volumes on the local cluster to have high I/O response times, Unified Manager triggers performance events to notify you.

When a local cluster in a MetroCluster configuration mirrors its data to its partner cluster, the data is written to NVRAM and then transferred over the interswitch links (ISLs) to the remote aggregates. Unified Manager

analyzes the NVRAM to identify the workloads whose high write throughput is overutilizing the NVRAM, placing the NVRAM in contention.

Workloads whose deviation in response time has exceeded the performance threshold are called *victims* and workloads whose deviation in write throughput to the NVRAM is higher than usual, causing the contention, are called *bullies*. Because only the write requests are mirrored to the partner cluster, Unified Manager does not analyze read throughput.


Unified Manager treats the clusters in a MetroCluster configuration as individual clusters. It does not distinguish between clusters that are partners or correlate the write throughput from each cluster.

Volume behavior during switchover and switchback


Events that trigger a switchover or switchback cause active volumes to be moved from one cluster to the other cluster in the disaster recovery group. The volumes on the cluster that were active and serving data to clients are stopped, and the volumes on the other cluster are activated and start serving data. Unified Manager monitors only those volumes that are active and running.

Because volumes are moved from one cluster to another, it is recommended that you monitor both clusters. A single instance of Unified Manager can monitor both clusters in a MetroCluster configuration, but sometimes the distance between the two locations necessitates using two Unified Manager instances to monitor both clusters. The following figure shows a single instance of Unified Manager:


Normal operation



Cluster B fails --- switchover to Cluster A



Cluster B is repaired --- switchback to Cluster B



□ = active and monitored

■ = inactive and not monitored

The volumes with p in their names indicate the primary volumes, and the volumes with b in their names are mirrored backup volumes that are created by SnapMirror.

During normal operation:

- Cluster A has two active volumes: Vol1_p and Vol2_p.
- Cluster B has two active volumes: Vol3_p and Vol4_p.
- Cluster A has two inactive volumes: Vol3_b and Vol4_b.
- Cluster B has two inactive volumes: Vol1_b and Vol2_b.

Information pertaining to each of the active volumes (statistics, events, and so on) is collected by Unified Manager. Vol1_p and Vol2_p statistics are collected by Cluster A, and Vol3_p and Vol4_p statistics are collected by Cluster B.

After a catastrophic failure causes a switchover of active volumes from Cluster B to Cluster A:

- Cluster A has four active volumes: Vol1_p, Vol2_p, Vol3_b, and Vol4_b.


- Cluster B has four inactive volumes: Vol3p, Vol4p, Vol1b, and Vol2b.

As during normal operation, information pertaining to each of the active volumes is collected by Unified Manager. But in this case, Vol1p and Vol2p statistics are collected by Cluster A, and Vol3b and Vol4b statistics are also collected by Cluster A.

Note that Vol3p and Vol3b are not the same volumes, because they are on different clusters. The information in Unified Manager for Vol3p is not the same as Vol3b:

- During switchover to Cluster A, Vol3p statistics and events are not visible.
- On the very first switchover, Vol3b looks like a new volume with no historical information.

When Cluster B is repaired and a switchback is performed, Vol3p is active again on Cluster B, with the historical statistics and a gap of statistics for the period during the switchover. Vol3b is not viewable from Cluster A until another switchover occurs:



- MetroCluster volumes that are inactive, for example, Vol3b on Cluster A after switchback, are identified with the message “This volume was deleted”. The volume is not actually deleted, but it is not currently being monitored by Unified Manager because it is not the active volume.
- If a single Unified Manager is monitoring both clusters in a MetroCluster configuration, volume search returns information for whichever volume is active at that time. For example, a search for “Vol3” would return statistics and events for Vol3b on Cluster A if a switchover has occurred and Vol3 has become active on Cluster A.

What performance events are

Performance events are incidents related to workload performance on a cluster. They help you identify workloads with slow response times. Together with health events that occurred at the same time, you can determine the issues that might have caused, or contributed to, the slow response times.

When Unified Manager detects multiple occurrences of the same event condition for the same cluster component, it treats all occurrences as a single event, not as separate events.

Performance event analysis and notification

Performance events notify you about I/O performance issues on a workload caused by contention on a cluster component. Unified Manager analyzes the event to identify all workloads involved, the component in contention, and whether the event is still an issue that you might need to resolve.

Unified Manager monitors the I/O latency (response time) and IOPS (operations) for volumes on a cluster. When other workloads overuse a cluster component, for example, the component is in contention and cannot perform at an optimal level to meet workload demands. The performance of other workloads that are using the same component might be impacted, causing their latencies to increase. If the latency crosses the dynamic performance threshold, Unified Manager triggers a performance event to notify you.

Event analysis

Unified Manager performs the following analyses, using the previous 15 days of performance statistics, to identify the victim workloads, bully workloads, and the cluster component involved in an event:

- Identifies victim workloads whose latency has crossed the dynamic performance threshold, which is the upper boundary of the latency forecast:
 - For volumes on HDD or Flash Pool hybrid aggregates (local tier), events are triggered only when the latency is greater than 5 milliseconds (ms) and the IOPS are more than 10 operations per second (ops/sec).
 - For volumes on all-SSD aggregates or FabricPool aggregates (cloud tier), events are triggered only when the latency is greater than 1 ms and the IOPS are more than 100 ops/sec.
- Identifies the cluster component in contention.



If the latency of victim workloads at the cluster interconnect is greater than 1 ms, Unified Manager treats this as significant and triggers an event for the cluster interconnect.

- Identifies the bully workloads that are overusing the cluster component and causing it to be in contention.
- Ranks the workloads involved, based on their deviation in utilization or activity of a cluster component, to determine which bullies have the highest change in usage of the cluster component and which victims are the most impacted.

An event might occur for only a brief moment and then correct itself after the component it is using is no longer in contention. A continuous event is one that reoccurs for the same cluster component within a five-minute interval and remains in the active state. For continuous events, Unified Manager triggers an alert after detecting the same event during two consecutive analysis intervals.

When an event is resolved, it remains available in Unified Manager as part of the record of past performance issues for a volume. Each event has a unique ID that identifies the event type and the volumes, cluster, and cluster components involved.



A single volume can be involved in more than one event at the same time.

Event state

Events can be in one of the following states:

- **Active**

Indicates that the performance event is currently active (new or acknowledged). The issue causing the event has not corrected itself or has not been resolved. The performance counter for the storage object remains above the performance threshold.

- **Obsolete**

Indicates that the event is no longer active. The issue causing the event has corrected itself or has been resolved. The performance counter for the storage object is no longer above the performance threshold.

Event notification

The events are displayed on the Dashboard page and on many other pages in the user interface, and alerts for those events are sent to specified email addresses. You can view detailed analysis information about an event and get suggestions for resolving it on the Event details page and on the Workload Analysis page.

Event interaction

On the Event details page and on the Workload Analysis page, you can interact with events in the following ways:

- Moving the mouse over an event displays a message that shows the date and time when the event was detected.

If there are multiple events for the same time period, the message shows the number of events.

- Clicking a single event displays a dialog box that shows more detailed information about the event, including the cluster components that are involved.

The component in contention is circled and highlighted red. You can click **View full analysis** to view the full analysis on the Event details page. If there are multiple events for the same time period, the dialog box shows details about the three most recent events. You can click an event to view the event analysis on the Event details page.

How Unified Manager determines the performance impact for an event

Unified Manager uses the deviation in activity, utilization, write throughput, cluster component usage, or I/O latency (response time) for a workload to determine the level of impact to workload performance. This information determines the role of each workload in the event and how they are ranked on the Event details page.

Unified Manager compares the last analyzed values for a workload to the expected range (latency forecast) of values. The difference between the values last analyzed and the expected range of values identifies the workloads whose performance was most impacted by the event.

For example, suppose a cluster contains two workloads: Workload A and Workload B. The latency forecast for Workload A is 5-10 milliseconds per operation (ms/op) and its actual latency is usually around 7 ms/op. The latency forecast for Workload B is 10-20 ms/op and its actual latency is usually around 15 ms/op. Both

workloads are well within their latency forecast. Due to contention on the cluster, the latency of both workloads increases to 40 ms/op, crossing the dynamic performance threshold, which is the upper bounds of the latency forecast, and triggering events. The deviation in latency, from the expected values to the values above the performance threshold, for Workload A is around 33 ms/op, and the deviation for Workload B is around 25 ms/op. The latency of both workloads spike to 40 ms/op, but Workload A had the bigger performance impact because it had the higher latency deviation at 33 ms/op.

On the Event details page, in the System Diagnosis section, you can sort workloads by their deviation in activity, utilization, or throughput for a cluster component. You can also sort workloads by latency. When you select a sort option, Unified Manager analyzes the deviation in activity, utilization, throughput, or latency since the event was detected from the expected values to determine the workload sort order. For the latency, the red dots (●) indicate a performance threshold crossing by a victim workload, and the subsequent impact to the latency. Each red dot indicates a higher level of deviation in latency, which helps you identify the victim workloads whose latency was impacted the most by an event.

Cluster components and why they can be in contention

You can identify cluster performance issues when a cluster component goes into contention. The performance of workloads that use the component slow down and their response time (latency) for client requests increases, which triggers an event in Unified Manager.

A component that is in contention cannot perform at an optimal level. Its performance has declined, and the performance of other cluster components and workloads, called *victims*, might have increased latency. To bring a component out of contention, you must reduce its workload or increase its ability to handle more work, so that the performance can return to normal levels. Because Unified Manager collects and analyzes workload performance in five-minute intervals, it detects only when a cluster component is consistently overused. Transient spikes of overusage that last for only a short duration within the five-minute interval are not detected.

For example, a storage aggregate might be under contention because one or more workloads on it are competing for their I/O requests to be fulfilled. Other workloads on the aggregate can be impacted, causing their performance to decrease. To reduce the amount of activity on the aggregate, there are different steps you can take, such as moving one or more workloads to a less busy aggregate or node, to lessen the overall workload demand on the current aggregate. For a QoS policy group, you can adjust the throughput limit, or move workloads to a different policy group, so that the workloads are no longer being throttled.

Unified Manager monitors the following cluster components to alert you when they are in contention:

- **Network**

Represents the wait time of I/O requests by the external networking protocols on the cluster. The wait time is time spent waiting for “transfer ready” transactions to finish before the cluster can respond to an I/O request. If the network component is in contention, it means high wait time at the protocol layer is impacting the latency of one or more workloads.

- **Network Processing**

Represents the software component in the cluster involved with I/O processing between the protocol layer and the cluster. The node handling network processing might have changed since the event was detected. If the network processing component is in contention, it means high utilization at the network processing node is impacting the latency of one or more workloads.

When using an All SAN Array cluster in an active-active configuration, the network processing latency value is displayed for both nodes so you can verify the nodes are sharing the load equally.

- **QoS Limit Max**

Represents the throughput maximum (peak) setting of the storage Quality of Service (QoS) policy group assigned to the workload. If the policy group component is in contention, it means all workloads in the policy group are being throttled by the set throughput limit, which is impacting the latency of one or more of those workloads.

- **QoS Limit Min**

Represents the latency to a workload that is being caused by QoS throughput minimum (expected) setting assigned to other workloads. If the QoS minimum set on certain workloads use the majority of the bandwidth to guarantee the promised throughput, other workloads will be throttled and see more latency.

- **Cluster Interconnect**

Represents the cables and adapters with which clustered nodes are physically connected. If the cluster interconnect component is in contention, it means high wait time for I/O requests at the cluster interconnect is impacting the latency of one or more workloads.

- **Data Processing**

Represents the software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains the workload. The node handling data processing might have changed since the event was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the latency of one or more workloads.

- **Volume Activation**

Represents the process that tracks the usage of all active volumes. In large environments where more than 1000 volumes are active, this process tracks how many critical volumes need to access resources through the node at the same time. When the number of concurrent active volumes exceeds the recommended maximum threshold, some of the non-critical volumes will experience latency as identified here.

- **MetroCluster Resources**

Represents the MetroCluster resources, including NVRAM and interswitch links (ISLs), used to mirror data between clusters in a MetroCluster configuration. If the MetroCluster component is in contention, it means high write throughput from workloads on the local cluster or a link health issue is impacting the latency of one or more workloads on the local cluster. If the cluster is not in a MetroCluster configuration, this icon is not displayed.

- **Aggregate or SSD Aggregate Ops**

Represents the storage aggregate on which the workloads are running. If the aggregate component is in contention, it means high utilization on the aggregate is impacting the latency of one or more workloads. An aggregate consists of all HDDs, or a mix of HDDs and SSDs (a Flash Pool aggregate), or a mix of HDDs and a cloud tier (a FabricPool aggregate). An “SSD Aggregate” consists of all SSDs (an all-flash aggregate), or a mix of SSDs and a cloud tier (a FabricPool aggregate).

- **Cloud Latency**

Represents the software component in the cluster involved with I/O processing between the cluster and the cloud tier on which user data is stored. If the cloud latency component is in contention, it means that a large amount of reads from volumes that are hosted on the cloud tier are impacting the latency of one or more workloads.

- **Sync SnapMirror**

Represents the software component in the cluster involved with replicating user data from the primary volume to the secondary volume in a SnapMirror Synchronous relationship. If the sync SnapMirror component is in contention, it means that the activity from SnapMirror Synchronous operations are impacting the latency of one or more workloads.

Roles of workloads involved in a performance event

Unified Manager uses roles to identify the involvement of a workload in a performance event. The roles include victims, bullies, and sharks. A user-defined workload can be a victim, bully, and shark at the same time.

Role	Description
Victim	A user-defined workload whose performance has decreased due to other workloads, called bullies, that are over-using a cluster component. Only user-defined workloads are identified as victims. Unified Manager identifies victim workloads based on their deviation in latency, where the actual latency, during an event, has greatly increased from its latency forecast (expected range).
Bully	A user-defined or system-defined workload whose over-use of a cluster component has caused the performance of other workloads, called victims, to decrease. Unified Manager identifies bully workloads based on their deviation in usage of a cluster component, where the actual usage, during an event, has greatly increased from its expected range of usage.
Shark	A user-defined workload with the highest usage of a cluster component compared to all workloads involved in an event. Unified Manager identifies shark workloads based on their usage of a cluster component during an event.

Workloads on a cluster can share many of the cluster components, such as aggregates and the CPU for network and data processing. When a workload, such as a volume, increases its usage of a cluster component to the point that the component cannot efficiently meet workload demands, the component is in contention. The workload that is over-using a cluster component is a bully. The other workloads that share those components, and whose performance is impacted by the bully, are the victims. Activity from system-defined workloads, such as deduplication or Snapshot copies, can also escalate into “bullying”.

When Unified Manager detects an event, it identifies all workloads and cluster components involved, including the bully workloads that caused the event, the cluster component that is in contention, and the victim workloads whose performance has decreased due to the increased activity of bully workloads.



If Unified Manager cannot identify the bully workloads, it only alerts on the victim workloads and the cluster component involved.

Unified Manager can identify workloads that are victims of bully workloads, and also identify when those same workloads become bully workloads. A workload can be a bully to itself. For example, a high-performing workload that is being throttled by a policy group limit causes all workloads in the policy group to be throttled, including itself. A workload that is a bully or a victim in an ongoing performance event might change its role or no longer be a participant in the event.

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.