Cisco and Hitachi Adaptive Solutions with SAN Analytics

Best Practices Guide

MK-SL-244-01
September 2022

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@HitachiVantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision history

Changes	Date
Updated image quality	September 22, 2022
Initial release	August 1, 2022

Best Practices Guide

This guide documents the best practices of using SAN Analytics software to identify, resolve, and troubleshoot performance degradation of a Cisco UCS fabric backed by a Hitachi Virtual Storage Platform. The solution uses the Fibre Channel-Non-Volatile Memory Express (NVMe) protocol to back a VMware 7.0U3 virtualized environment. This document does not cover configuration.

This guide is written for professional services staff such as storage administrators, VMware administrators, sales engineers, field consultants, and validated Hitachi and Cisco resale partners. Readers of this document must have knowledge of RAID systems and functionality, VMware ESXi and vCenter environments, and converged infrastructure.

Note: Testing of these procedures was in a lab environment. Many factors impact production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Introduction

Fibre Channel is a high-speed data transfer protocol that provides in-order lossless delivery of raw block data. It is designed to provide connections between endpoint servers and storage devices. This technology supports point-to-point devices through a common fabric known as a Storage Area Network (SAN). Fibre Channel SANs are typically deployed for low latency applications that are best suited to block-based storage, such as databases used for high-speed online transactional processing (OLTP) and those found in banking, online ticketing, and virtual environments.

Because Fibre Channel SANs are lossless, end users often find themselves in situations of SAN congestion caused by overutilization of an edge link or a slow drain device. Although both scenarios have the same results, the reasons are different. Congestion caused by overutilization of an edge-link happens when the switch receives more traffic than can be sent on the link. In contrast, congestion caused by a slow drain happens when a device cannot process the frames as fast as its ingress rate because of issues such as high CPU usage or even a software defect. In turn, this device, called a slow drain device, applies backpressure by slowing down Receiver Ready (R_RDY) signals.

SAN configuration

This is a best-practice datacenter architecture built by Hitachi Vantara and Cisco Systems to meet your enterprise needs using virtual server workloads. It uses a Hitachi Virtual Storage Platform (VSP) storage system to connect the Cisco MDS Multilayer switches that control Fibre Channel/SAN communication to the Cisco UCS Fabric Interconnects and Cisco UCS chassis.

Northbound Ethernet/LAN networking is enabled through the Cisco Nexus 9000 family of switches. For information about SAN connectivity and network connectivity see the Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Deployment Guide.

The following figure shows the validated architecture for Cisco and Hitachi Adaptive Solutions for Converged Infrastructure. Red lines represent Fabric A connections, blue lines represent Fabric B connections, and the rest are port channel connections.

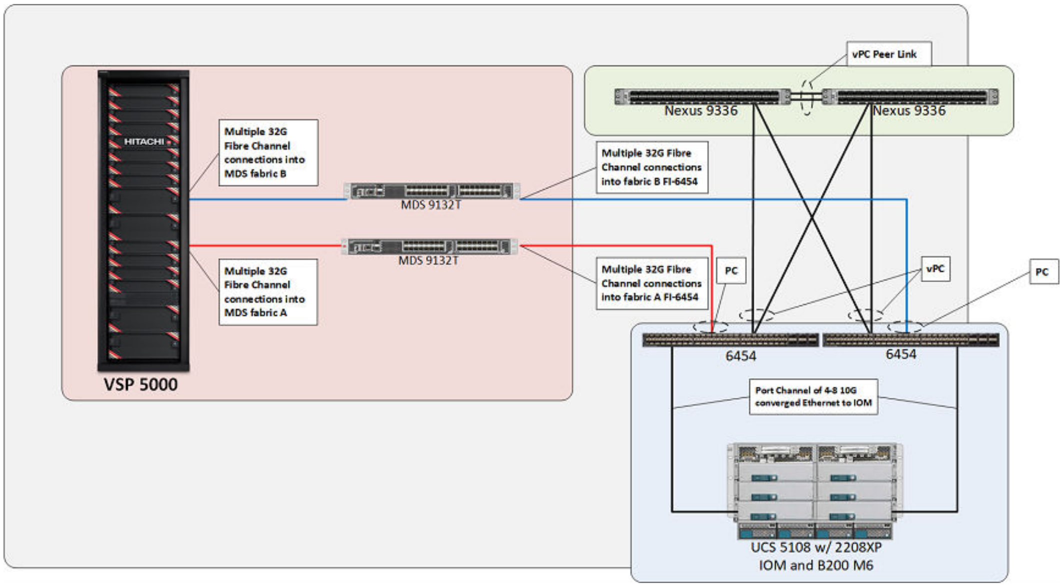
vPC Peer Link
Nexus 9336
Nexus 9336
HITACHI
Multiple 32G Fibre Channel connections into MDS fabric B
MDS 9132T
Multiple 32G Fibre Channel connections into fabric B FI-6454
MDS 9132T
Multiple 32G Fibre Channel connections into MDS fabric A
Multiple 32G Fibre Channel connections into fabric A FI-6454
PC
vPC
PC
6454
6454
VSP 5000
Port Channel of 4-8 10G converged Ethernet to IOM
UCS 5108 w/ 2208XP IOM and B200 M6

Solution components

The following tables list the tested solution components.

Table 1 Hardware components

Component	Version
Hitachi Virtual Storage Platform VSP 5600H	90-07-01-00/00
Cisco MDS 9132T Fibre Channel switch	MDS NX-OS 9.2(2)
Cisco Nexus 9332-FX2 switch	NX-OS 7.0(3)I7(9)
Cisco Fabric Interconnect 6454	4.2(1i)
Cisco Unified Computing System B200 M6 Blade Servers	4.2(1i)
Cisco Unified Computing System 2208XP IOM	4.2(1i)

Table 2 Software components

Component	Version
VMware vCenter Standalone (VCSA) 7.0 U3	7.0.3, 19234570
VMware ESXi 7.0 U3 Cisco Custom Image	7.0.3, 19193900
VMware ESXi 7.0U3 nenic	1.0.42.0
VMware ESXi 7.0U3 nfnic	5.0.0.34
Hitachi Ops Center Analyzer	10.8.1
Cisco Nexus Dashboard Fabric Controller	12.0.1a

Overutilization

Overutilization was purposely created within the Cisco UCS environment backed by Hitachi VSP storage to explore the software capabilities of Hitachi Ops Center Analyzer and Cisco SAN Analytics in conjunction with Cisco Nexus Dashboard Fabric Controller (NDFC). With overutilization, the data transfer rate (Tx Data rate) from the storage system is faster than the host port speed. Tx Data rate is the transmitted data rate respective to the port.

A Tx Data rate mismatch can be caused by an 8G link or port-channel that is connected to a server and a storage system that is backed by a 32G link which causes a speed mismatch where the incoming data is more than what can be sent by the 8G link. This requires the switch to buffer the data until it can be processed, which causes backpressure to occur directly on the Cisco MDS switch, where R_RDY signals are sent slowly because of a lack of free receiver buffers.

The host that is connected to an over-utilized link might not be impacted, but other servers on the fabric are impacted because of backpressure congestion. The following figure represents the Cisco UCS during overutilization conditions. The red arrows depicted on the storage ports represent SAN congestion.

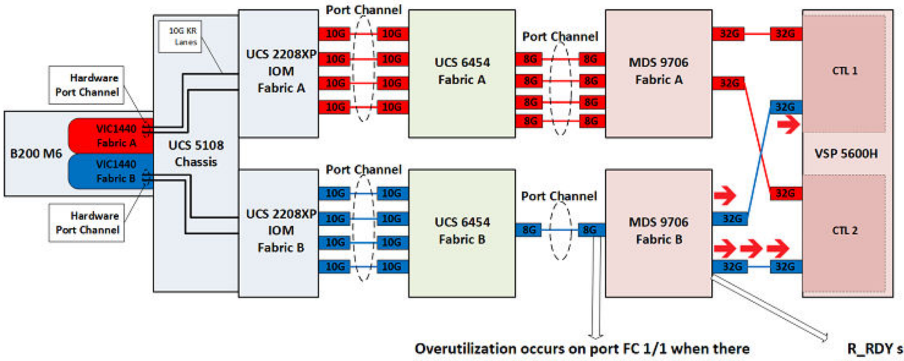

Overutilization occurs on port FC 1/1 when there is a speed mismatch between host and storage ports. In this case host ports on fabric B are running at 8G while storage ports are at 32G.

R_RDY signal causes back pressure on target ports. This also causes RxB2Bto0 counter to increase on MDS.

Hitachi Ops Center Analyzer

Hitachi Ops Center Analyzer allows end-to-end metric collection from compute hosts, Fibre Channel switches, and backend storage systems. The best practices in this guide cover configuration and installation instructions as well as how to use Hitachi Ops Center Analyzer features to detect, troubleshoot, and resolve performance degradation within a Cisco UCS SAN fabric backed by Hitachi VSP storage. The following figure shows Ops Center Analyzer capabilities.

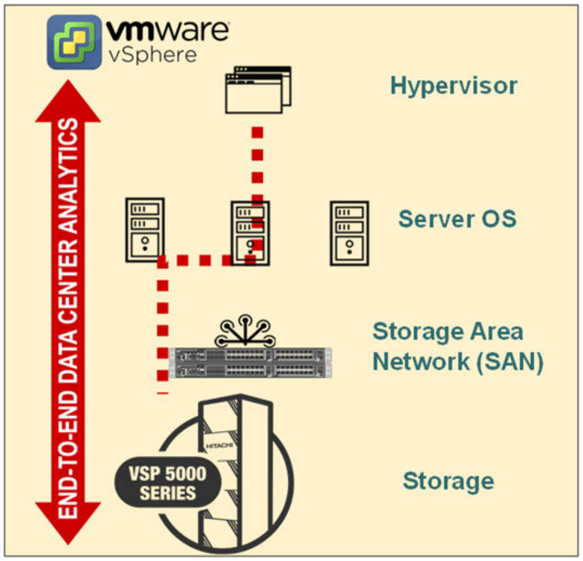

Cisco SAN Analytics and Cisco Nexus Dashboard Fabric Controller

The Cisco SAN Analytics solution offers end-to-end visibility into Fibre Channel block storage traffic. The solution is natively available on the storage area network because of the integrated-by-design architecture with the Cisco MDS 9000 switch family.

Cisco SAN Analytics delivers deep visibility into I/O traffic between the compute and the storage infrastructure. This information is in addition to the already-available visibility obtained from individual ports, switches, servers, virtual machines, and storage systems. Cisco MDS switches export all the metrics to Cisco Nexus Dashboard Fabric Controller (NDFC) using streaming telemetry. In turn, NDFC automatically calculates performance baselines and categorizes the devices based on their deviations. It also provides alerting using the Anomaly Detection feature.

Cisco SAN Analytics and NDFC can be used to detect, troubleshoot, and resolve performance degradation within a Cisco UCS SAN fabric backed by Hitachi VSP storage. The following figure shows the NDFC capabilities.

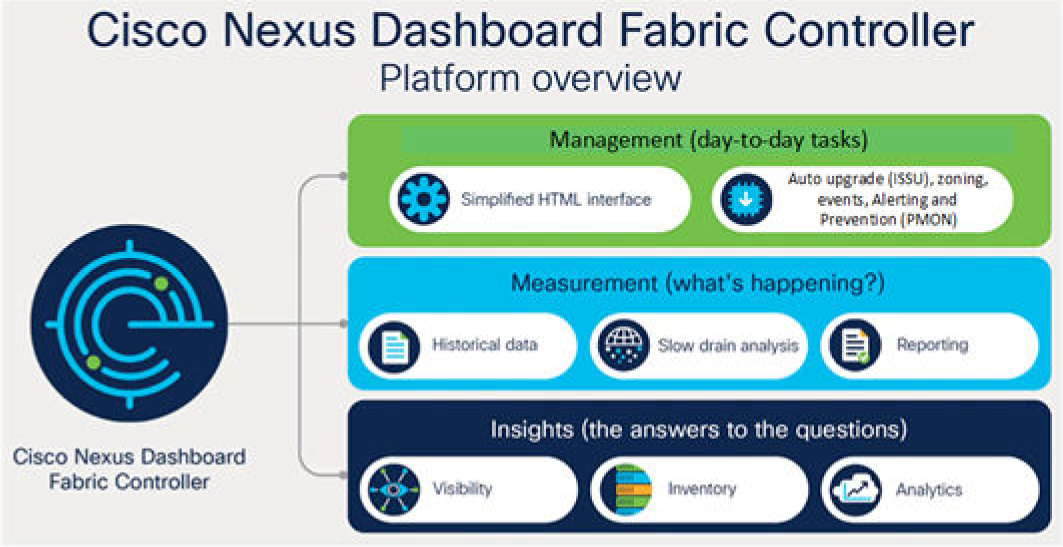
Cisco Nexus Dashboard Fabric Controller
Platform overview

Management (day-to-day tasks)
Simplified HTML interface
Auto upgrade (ISSU), zoning, events, Alerting and Prevention (PMON)
Measurement (what's happening?)
Historical data
Slow drain analysis
Reporting
Insights (the answers to the questions)
Visibility
Inventory
Analytics
Cisco Nexus Dashboard Fabric Controller

Cisco SAN Analytics and NDFC best practices

This section describes how to identify, troubleshoot, and resolve performance and congestion issues using Cisco SAN Analytics and NDFC.

Cisco MDS

This section describes Cisco MDS NX-OS commands.

Tx and Rx detection

Using the Cisco NX-OS CLI on Cisco MDS switches, administrators can run simple commands to obtain insight into counters and credits available for ports being used as targets within the fabric. In the following examples, Fibre Channel ports 1/19 and 1/20 on Fabric B are the target ports. The following command displays the ports that are used through counters as well as the remaining B2B credits:

```
show interface counters brief
```

This command helps administrators understand link utilization. The following figure shows that fc1/1, fc1/19, and fc1/20 are experiencing high I/O rate and frame count. The port-channel between Cisco MDS and Fabric Interconnect is on fc1/1. The output rate on fc1/1 is approximately 800 Mbps, which is the maximum data-rate of an 8G Fibre Channel port. This condition indicates congestion because of the overutilization of fc1/1.

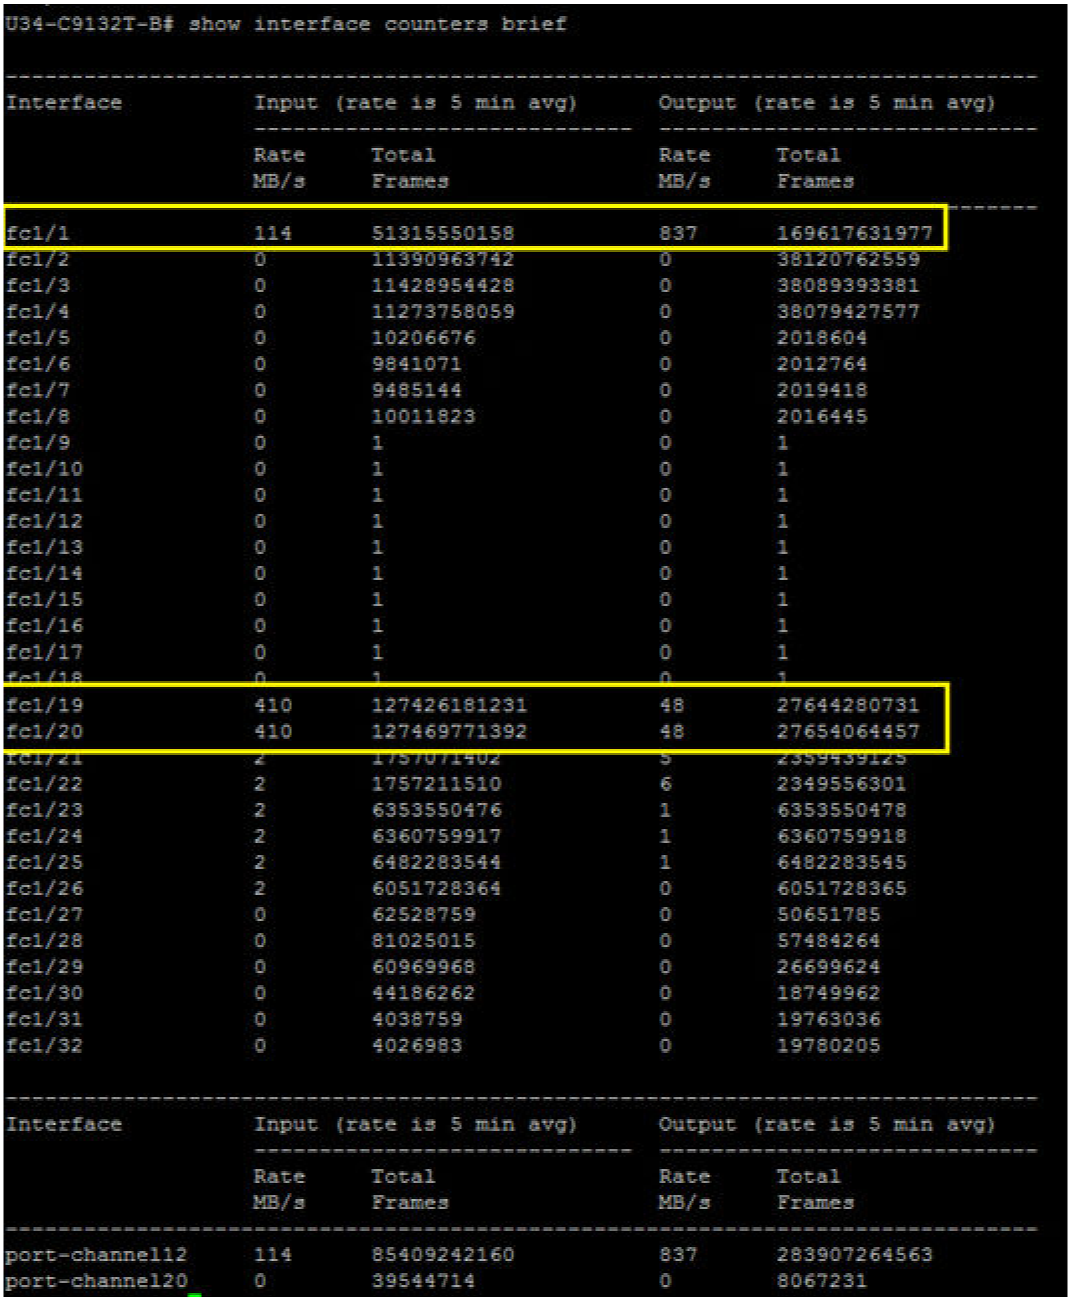

```
U34-C9132T-B# show interface counters brief

-------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   -----------------------------  -----------------------------
                   Rate     Total                 Rate     Total
                   MB/s     Frames                MB/s     Frames
-------------------------------------------------------------------------------
fc1/1              114      51315550158           837      169617631977
fc1/2              0        11390963742           0        38120762559
fc1/3              0        11428954428           0        38089393381
fc1/4              0        11273758059           0        38079427577
fc1/5              0        10206676              0        2018604
fc1/6              0        9841071               0        2012764
fc1/7              0        9485144               0        2019418
fc1/8              0        10011823              0        2016445
fc1/9              0        1                     0        1
fc1/10             0        1                     0        1
fc1/11             0        1                     0        1
fc1/12             0        1                     0        1
fc1/13             0        1                     0        1
fc1/14             0        1                     0        1
fc1/15             0        1                     0        1
fc1/16             0        1                     0        1
fc1/17             0        1                     0        1
fc1/18             0        1                     0        1
fc1/19             410      127426181231          48       27644280731
fc1/20             410      127469771392          48       27654064457
fc1/21             2        1757071402            5        2359439125
fc1/22             2        1757211510            6        2349556301
fc1/23             2        6353550476            1        6353550478
fc1/24             2        6360759917            1        6360759918
fc1/25             2        6482283544            1        6482283545
fc1/26             2        6051728364            0        6051728365
fc1/27             0        62528759              0        50651785
fc1/28             0        81025015              0        57484264
fc1/29             0        60969968              0        26699624
fc1/30             0        44186262              0        18749962
fc1/31             0        4038759               0        19763036
fc1/32             0        4026983               0        19780205

-------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   -----------------------------  -----------------------------
                   Rate     Total                 Rate     Total
                   MB/s     Frames                MB/s     Frames
-------------------------------------------------------------------------------
port-channel12     114      85409242160           837      283907264563
port-channel20     0        39544714              0        8067231
```

RxB2Bto0 counters

Use the following command to obtain additional interface details:

```
show interface <fc-port> counters
```

This command helps administrators understand how many buffer credits remain on a per-port basis. A receive B2B credits value of 0 indicates that the port is under congestion conditions. The receive B2B credits remaining shown in the following figure is an instantaneous value. Therefore, the quickly incrementing value of a receive B2B credit that transitions to zero is an important metric.

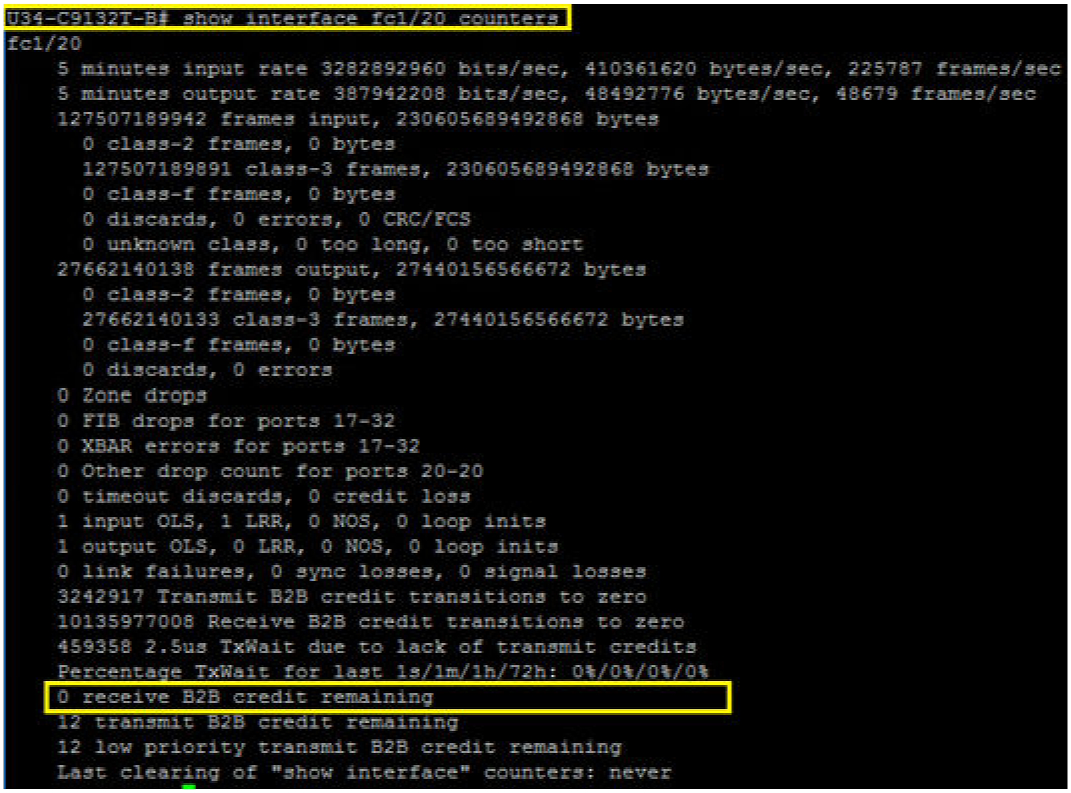

```
U34-C9132T-B# show interface fc1/20 counters
fc1/20
    5 minutes input rate 3282892960 bits/sec, 410361620 bytes/sec, 225787 frames/sec
    5 minutes output rate 387942208 bits/sec, 48492776 bytes/sec, 48679 frames/sec
    127507189942 frames input, 230605689492868 bytes
      0 class-2 frames, 0 bytes
      127507189891 class-3 frames, 230605689492868 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors, 0 CRC/FCS
      0 unknown class, 0 too long, 0 too short
    27662140138 frames output, 27440156566672 bytes
      0 class-2 frames, 0 bytes
      27662140133 class-3 frames, 27440156566672 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 errors
    0 Zone drops
    0 FIB drops for ports 17-32
    0 XBAR errors for ports 17-32
    0 Other drop count for ports 20-20
    0 timeout discards, 0 credit loss
    1 input OLS, 1 LRR, 0 NOS, 0 loop inits
    1 output OLS, 0 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
    3242917 Transmit B2B credit transitions to zero
    10135977008 Receive B2B credit transitions to zero
    459358 2.5us TxWait due to lack of transmit credits
    Percentage TxWait for last 1s/1m/1h/72h: 0%/0%/0%/0%
    0 receive B2B credit remaining
    12 transmit B2B credit remaining
    12 low priority transmit B2B credit remaining
    Last clearing of "show interface" counters: never
```

Port Monitor logging

Port Monitor is a feature of Cisco MDS 9000 switches that monitors data-plane metrics at a low granularity (such as 1 second) and takes automatic actions, such as generating alerts, shutting down a port, isolating a port, or enabling Dynamic Ingress Rate Limiting.

A Port Monitor policy must be configured using the NX-OS CLI or NDFC. When configured thresholds are exceeded, alerts can be sent to remote syslog servers or SNMP trap receivers. The switch stores these alerts in the form of logs. The following output shows the thresholds of various metrics set using Port Monitor.

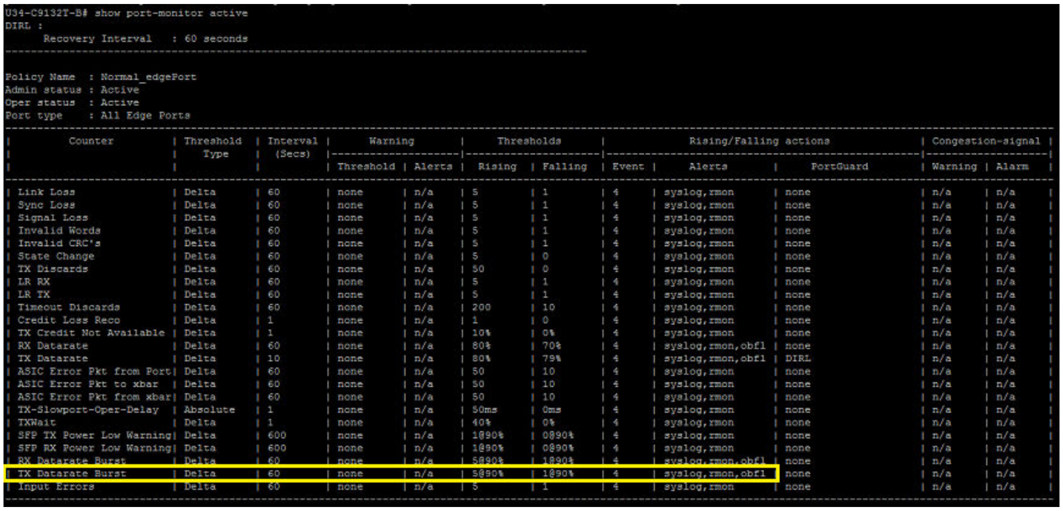

```
U34-C9132T-B# show port-monitor active
DIRL :
      Recovery Interval    : 60 seconds
------------------------------------------------------------------------------------------

Policy Name  : Normal_edgePort
Admin status : Active
Oper status  : Active
Port type    : All Edge Ports
-----------------------------------------------------------------------------------------------------------------------------------------------------------------
|         Counter          | Threshold  | Interval |       Warning        |       Thresholds       |                Rising/Falling actions                 | Congestion-signal |
|                          |   Type     |  (Secs)  |----------------------|------------------------|-------------------------------------------------------|-------------------|
|                          |            |          | Threshold | Alerts   |  Rising    | Falling   | Event |      Alerts          |       PortGuard        | Warning | Alarm   |
-----------------------------------------------------------------------------------------------------------------------------------------------------------------
| Link Loss                | Delta      | 60       | none      | n/a      | 5          | 1         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| Sync Loss                | Delta      | 60       | none      | n/a      | 5          | 1         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| Signal Loss              | Delta      | 60       | none      | n/a      | 5          | 1         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| Invalid Words            | Delta      | 60       | none      | n/a      | 5          | 1         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| Invalid CRC's            | Delta      | 60       | none      | n/a      | 5          | 1         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| State Change             | Delta      | 60       | none      | n/a      | 5          | 0         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| TX Discards              | Delta      | 60       | none      | n/a      | 50         | 0         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| LR RX                    | Delta      | 60       | none      | n/a      | 5          | 1         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| LR TX                    | Delta      | 60       | none      | n/a      | 5          | 1         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| Timeout Discards         | Delta      | 60       | none      | n/a      | 200        | 10        | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| Credit Loss Reco         | Delta      | 1        | none      | n/a      | 1          | 0         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| TX Credit Not Available  | Delta      | 1        | none      | n/a      | 10%        | 0%        | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| RX Datarate              | Delta      | 60       | none      | n/a      | 80%        | 70%       | 4     | syslog,rmon,obfl     | none                   | n/a     | n/a     |
| TX Datarate              | Delta      | 10       | none      | n/a      | 80%        | 79%       | 4     | syslog,rmon,obfl     | DIRL                   | n/a     | n/a     |
| ASIC Error Pkt from Port | Delta      | 60       | none      | n/a      | 50         | 10        | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| ASIC Error Pkt to xbar   | Delta      | 60       | none      | n/a      | 50         | 10        | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| ASIC Error Pkt from xbar | Delta      | 60       | none      | n/a      | 50         | 10        | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| TX-Slowport-Oper-Delay   | Absolute   | 1        | none      | n/a      | 50ms       | 0ms       | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| TXWait                   | Delta      | 1        | none      | n/a      | 40%        | 0%        | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| SFP TX Power Low Warning | Delta      | 600      | none      | n/a      | 1@90%      | 0@90%     | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| SFP RX Power Low Warning | Delta      | 600      | none      | n/a      | 1@90%      | 0@90%     | 4     | syslog,rmon          | none                   | n/a     | n/a     |
| RX Datarate Burst        | Delta      | 60       | none      | n/a      | 5@90%      | 1@90%     | 4     | syslog,rmon,obfl     | none                   | n/a     | n/a     |
| TX Datarate Burst        | Delta      | 60       | none      | n/a      | 5@90%      | 1@90%     | 4     | syslog,rmon,obfl     | none                   | n/a     | n/a     |
| Input Errors             | Delta      | 60       | none      | n/a      | 5          | 1         | 4     | syslog,rmon          | none                   | n/a     | n/a     |
-----------------------------------------------------------------------------------------------------------------------------------------------------------------
```

Run the following command to view logs based on port monitor thresholds:

```
show log | grep fc<interface number>
```

The output of this command shows alerts that have been triggered according to port monitoring policies. The following figure represents fc1/1 showing a Tx Datarate Burst reaching a rising threshold of 5 times at 90% utilization at a polling interval of 60 seconds. The Tx Datarate Burst identifies an interface that faces the maximum data transmission for which the port is capable.

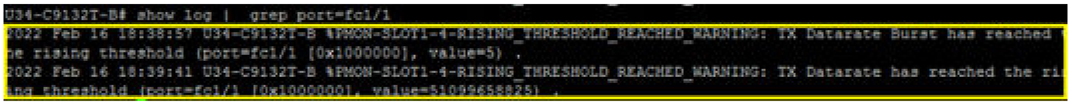

```
U34-C9132T-B# show log |  grep port=fc1/1
2022 Feb 16 18:38:57 U34-C9132T-B %PMON-SLOT1-4-RISING_THRESHOLD_REACHED_WARNING: TX Datarate Burst has reached
he rising threshold (port=fc1/1 [0x1000000], value=5) .
2022 Feb 16 18:39:41 U34-C9132T-B %PMON-SLOT1-4-RISING_THRESHOLD_REACHED_WARNING: TX Datarate has reached the ri
ing threshold (port=fc1/1 [0x1000000], value=51099658825) .
```

Overutilization congestion prevention using Cisco MDS Dynamic Ingress Rate Limiting

Cisco MDS Dynamic Ingress Rate Limiting (DIRL) identifies and resolves SAN congestion. DIRL must be enabled from Port Monitor to detect any symptoms of egress congestion on the switch ports. Then DIRL limits ingress data to prevent congestion in the egress direction. DIRL dynamically adapts the ingress traffic rate until the egress congestion is gone. By limiting ingress frames, DIRL also slows down the data-requesting frames (read I/O command) to the storage system.

To view DIRL settings on a fabric using MDS NX-OS, run the following command:

```
show port-monitor active
```

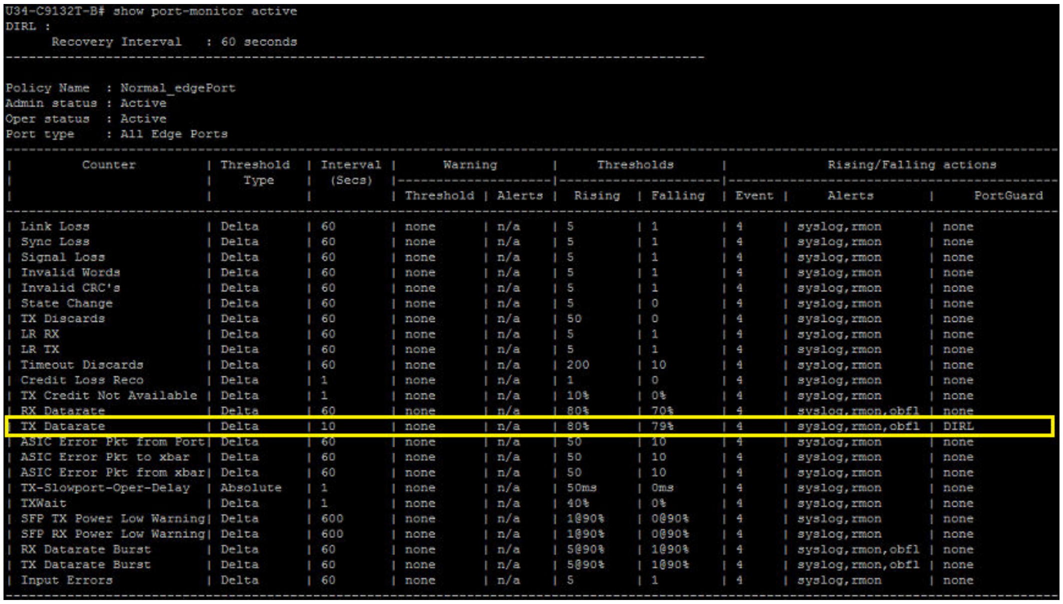

```
U34-C9132T-B# show port-monitor active
DIRL :
      Recovery Interval   : 60 seconds
-------------------------------------------------------------------------------------------

Policy Name  : Normal_edgePort
Admin status : Active
Oper status  : Active
Port type    : All Edge Ports
--------------------------------------------------------------------------------------------------------------------------------------------
|         Counter          | Threshold  | Interval |       Warning       |      Thresholds      |            Rising/Falling actions
|                          |    Type    |  (Secs)  |---------------------|----------------------|-------------------------------------------
|                          |            |          | Threshold | Alerts  |  Rising   | Falling  | Event |      Alerts        |     PortGuard
--------------------------------------------------------------------------------------------------------------------------------------------
| Link Loss                | Delta      | 60       | none      | n/a     | 5         | 1        | 4     | syslog,rmon        | none
| Sync Loss                | Delta      | 60       | none      | n/a     | 5         | 1        | 4     | syslog,rmon        | none
| Signal Loss              | Delta      | 60       | none      | n/a     | 5         | 1        | 4     | syslog,rmon        | none
| Invalid Words            | Delta      | 60       | none      | n/a     | 5         | 1        | 4     | syslog,rmon        | none
| Invalid CRC's            | Delta      | 60       | none      | n/a     | 5         | 1        | 4     | syslog,rmon        | none
| State Change             | Delta      | 60       | none      | n/a     | 5         | 0        | 4     | syslog,rmon        | none
| TX Discards              | Delta      | 60       | none      | n/a     | 50        | 0        | 4     | syslog,rmon        | none
| LR RX                    | Delta      | 60       | none      | n/a     | 5         | 1        | 4     | syslog,rmon        | none
| LR TX                    | Delta      | 60       | none      | n/a     | 5         | 1        | 4     | syslog,rmon        | none
| Timeout Discards         | Delta      | 60       | none      | n/a     | 200       | 10       | 4     | syslog,rmon        | none
| Credit Loss Reco         | Delta      | 1        | none      | n/a     | 1         | 0        | 4     | syslog,rmon        | none
| TX Credit Not Available  | Delta      | 1        | none      | n/a     | 10%       | 0%       | 4     | syslog,rmon        | none
| RX Datarate              | Delta      | 60       | none      | n/a     | 80%       | 70%      | 4     | syslog,rmon,obfl   | none
| TX Datarate              | Delta      | 10       | none      | n/a     | 80%       | 79%      | 4     | syslog,rmon,obfl   | DIRL
| ASIC Error Pkt from Port | Delta      | 60       | none      | n/a     | 50        | 10       | 4     | syslog,rmon        | none
| ASIC Error Pkt to xbar   | Delta      | 60       | none      | n/a     | 50        | 10       | 4     | syslog,rmon        | none
| ASIC Error Pkt from xbar | Delta      | 60       | none      | n/a     | 50        | 10       | 4     | syslog,rmon        | none
| TX-Slowport-Oper-Delay   | Absolute   | 1        | none      | n/a     | 50ms      | 0ms      | 4     | syslog,rmon        | none
| TXWait                   | Delta      | 1        | none      | n/a     | 40%       | 0%       | 4     | syslog,rmon        | none
| SFP TX Power Low Warning | Delta      | 600      | none      | n/a     | 1@90%     | 0@90%    | 4     | syslog,rmon        | none
| SFP RX Power Low Warning | Delta      | 600      | none      | n/a     | 1@90%     | 0@90%    | 4     | syslog,rmon        | none
| RX Datarate Burst        | Delta      | 60       | none      | n/a     | 5@90%     | 1@90%    | 4     | syslog,rmon,obfl   | none
| TX Datarate Burst        | Delta      | 60       | none      | n/a     | 5@90%     | 1@90%    | 4     | syslog,rmon,obfl   | none
| Input Errors             | Delta      | 60       | none      | n/a     | 5         | 1        | 4     | syslog,rmon        | none
--------------------------------------------------------------------------------------------------------------------------------------------
```

This command output shows that the Tx Data rate port guard DIRL has been set to prevent congestion among both slow and fast devices within the same fabric. Additionally, to view alerts for DIRL prevention, you can use the following command:

```
show logging last 20
```

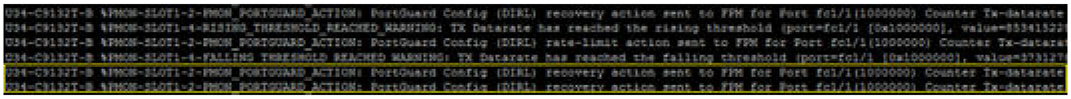

```
U34-C9132T-B %PMON-SLOT1-2-PMON_PORTGUARD_ACTION: PortGuard Config (DIRL) recovery action sent to FPM for Port fc1/1(1000000) Counter Tx-datarate
U34-C9132T-B %PMON-SLOT1-4-RISING_THRESHOLD_REACHED_WARNING: TX Datarate has reached the rising threshold (port=fc1/1 [0x1000000], value=85341522
U34-C9132T-B %PMON-SLOT1-2-PMON_PORTGUARD_ACTION: PortGuard Config (DIRL) rate-limit action sent to FPM for Port fc1/1(1000000) Counter Tx-datarat
U34-C9132T-B %PMON-SLOT1-4-FALLING_THRESHOLD_REACHED_WARNING: TX Datarate has reached the falling threshold (port=fc1/1 [0x1000000], value=373127
U34-C9132T-B %PMON-SLOT1-2-PMON_PORTGUARD_ACTION: PortGuard Config (DIRL) recovery action sent to FPM for Port fc1/1(1000000) Counter Tx-datarate
U34-C9132T-B %PMON-SLOT1-2-PMON_PORTGUARD_ACTION: PortGuard Config (DIRL) recovery action sent to FPM for Port fc1/1(1000000) Counter Tx-datarate
```

This command output shows DIRL action that prevents fc1/1 TxData rate so that overutilization does not occur.

Per-flow traffic utilization on a switch port in real-time

Real-time performance statistics are a useful tool in dynamic troubleshooting and fault isolation within the fabric. For any port on a switch, you can monitor several statistics with granular intervals that can consist of the top 10 metrics for initiator, target, and LUN (ITL) flows and initiator, target, and Namespace (ITN) flows. SAN Analytics must be enabled for this feature.

Run the following command to display per-flow traffic utilization with the highest throughput.

Note: Additional metrics aside from throughput, such an I/O per second (IOPS) and Exchange Completion Time (ECT), are also available.

```
show analytics -top -key THPUT
```

The output of this command shows the top 10 ITL flows with the highest throughput in correlation with the initiator, target, and VSP LUN. This command updates the display every 3-4 seconds. It finds the top 10 flows among tens of thousands of flows that might be active on the switch.

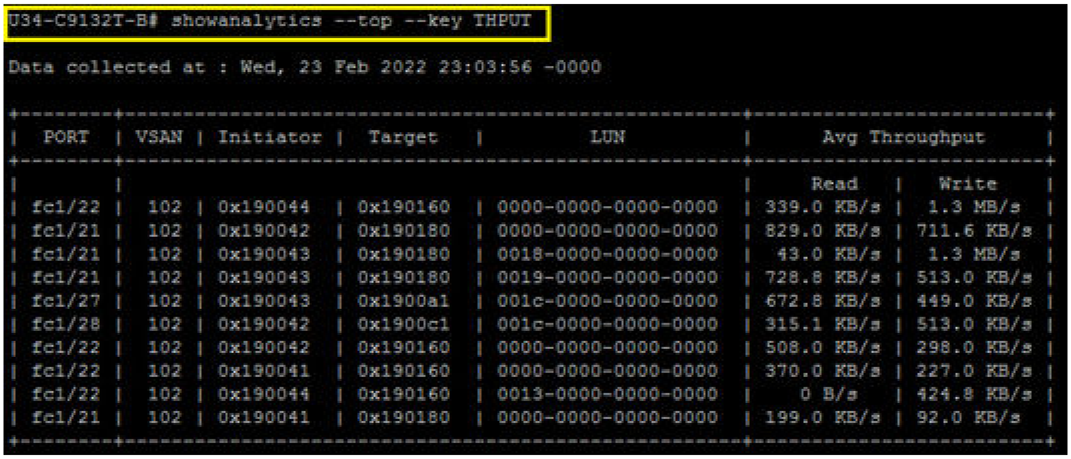

```
U34-C9132T-B# showanalytics --top --key THPUT

Data collected at : Wed, 23 Feb 2022 23:03:56 -0000

+--------+----------------------------------------------------------+-----------------------------+
|  PORT  | VSAN | Initiator |  Target   |          LUN              |       Avg Throughput        |
+--------+----------------------------------------------------------+-----------------------------+
|        |                                                          |    Read     |     Write     |
| fc1/22 |  102 | 0x190044  | 0x190160  | 0000-0000-0000-0000       | 339.0 KB/s |   1.3 MB/s    |
| fc1/21 |  102 | 0x190042  | 0x190180  | 0000-0000-0000-0000       | 829.0 KB/s | 711.6 KB/s    |
| fc1/21 |  102 | 0x190043  | 0x190180  | 0018-0000-0000-0000       |  43.0 KB/s |   1.3 MB/s    |
| fc1/21 |  102 | 0x190043  | 0x190180  | 0019-0000-0000-0000       | 728.8 KB/s | 513.0 KB/s    |
| fc1/27 |  102 | 0x190043  | 0x1900a1  | 001c-0000-0000-0000       | 672.8 KB/s | 449.0 KB/s    |
| fc1/28 |  102 | 0x190042  | 0x1900c1  | 001c-0000-0000-0000       | 315.1 KB/s | 513.0 KB/s    |
| fc1/22 |  102 | 0x190042  | 0x190160  | 0000-0000-0000-0000       | 508.0 KB/s | 298.0 KB/s    |
| fc1/22 |  102 | 0x190041  | 0x190160  | 0000-0000-0000-0000       | 370.0 KB/s | 227.0 KB/s    |
| fc1/22 |  102 | 0x190044  | 0x190160  | 0013-0000-0000-0000       |    0 B/s   | 424.8 KB/s    |
| fc1/21 |  102 | 0x190041  | 0x190180  | 0000-0000-0000-0000       | 199.0 KB/s | 92.0 KB/s     |
+--------+----------------------------------------------------------+-----------------------------+
```

Use the following command output to retrieve additional performance information to investigate SAN performance statistics of the initiator, target, and LUN. The following command drills down into initiator 0x190043, target 0x190160 and LUN 18:

```
showanalytics --initiator 0x190043 --target 0x190160 --lun 0018-0000-0000-0000 --info
--target-itl
```

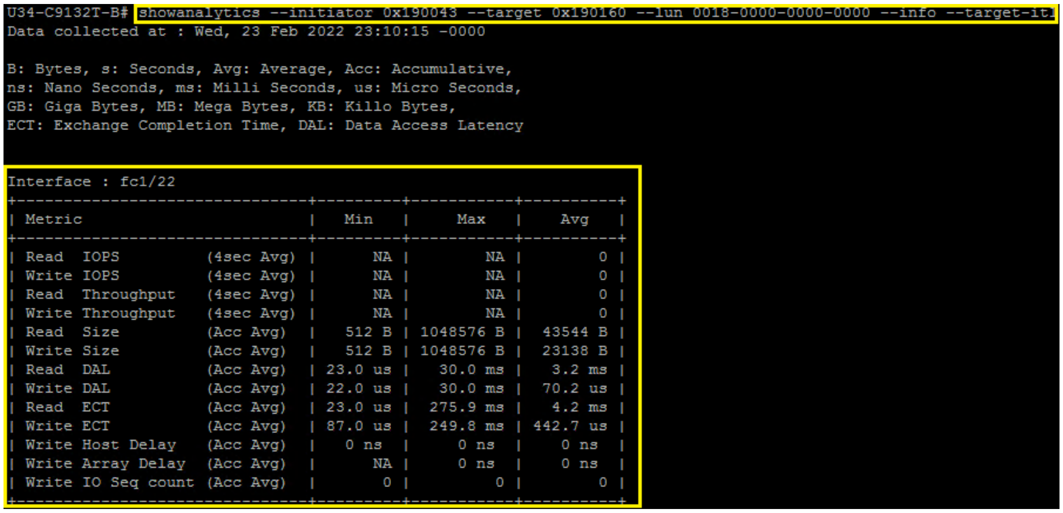

```
U34-C9132T-B# showanalytics --initiator 0x190043 --target 0x190160 --lun 0018-0000-0000-0000 --info --target-itl
Data collected at : Wed, 23 Feb 2022 23:10:15 -0000

B: Bytes, s: Seconds, Avg: Average, Acc: Accumulative,
ns: Nano Seconds, ms: Milli Seconds, us: Micro Seconds,
GB: Giga Bytes, MB: Mega Bytes, KB: Killo Bytes,
ECT: Exchange Completion Time, DAL: Data Access Latency

Interface : fc1/22
+--------------------------------+---------+-----------+----------+
| Metric                         |   Min   |    Max    |   Avg    |
+--------------------------------+---------+-----------+----------+
| Read  IOPS         (4sec Avg) |      NA |        NA |        0 |
| Write IOPS         (4sec Avg) |      NA |        NA |        0 |
| Read  Throughput   (4sec Avg) |      NA |        NA |        0 |
| Write Throughput   (4sec Avg) |      NA |        NA |        0 |
| Read  Size         (Acc Avg)  |   512 B | 1048576 B |  43544 B |
| Write Size         (Acc Avg)  |   512 B | 1048576 B |  23138 B |
| Read  DAL          (Acc Avg)  | 23.0 us |   30.0 ms |   3.2 ms |
| Write DAL          (Acc Avg)  | 22.0 us |   30.0 ms |  70.2 us |
| Read  ECT          (Acc Avg)  | 23.0 us |  275.9 ms |   4.2 ms |
| Write ECT          (Acc Avg)  | 87.0 us |  249.8 ms | 442.7 us |
| Write Host Delay   (Acc Avg)  |   0 ns  |     0 ns  |    0 ns  |
| Write Array Delay  (Acc Avg)  |      NA |     0 ns  |    0 ns  |
| Write IO Seq count (Acc Avg)  |       0 |         0 |        0 |
+--------------------------------+---------+-----------+----------+
```

Slowest and busiest switch ports in real-time

The slowest and busiest correlation is provided by the SAN Analytics feature. Exchange completion time (ECT) helps administrators understand which ports are the slowest in a fabric. Increased ECT values show performance degradation in the SAN fabric, especially if the backing storage system is using fast NVMe drive sets in conjunction with FC-NVMe.

Run the following command to gain insight into switch ports with the highest ECT.

```
showanalytics --top --key ECT
```

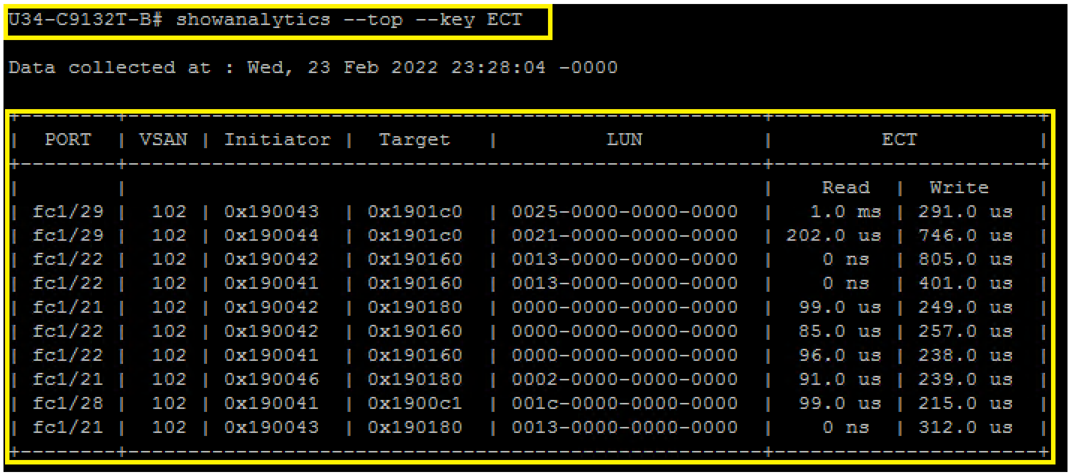

```
U34-C9132T-B# showanalytics --top --key ECT

Data collected at : Wed, 23 Feb 2022 23:28:04 -0000

+--------+---------------------------------------------------------+----------------------+
|  PORT  | VSAN | Initiator |  Target   |          LUN           |         ECT          |
+--------+---------------------------------------------------------+----------------------+
|        |                                                         |   Read   |   Write   |
| fc1/29 |  102 | 0x190043  | 0x1901c0  | 0025-0000-0000-0000    |  1.0 ms  | 291.0 us  |
| fc1/29 |  102 | 0x190044  | 0x1901c0  | 0021-0000-0000-0000    | 202.0 us | 746.0 us  |
| fc1/22 |  102 | 0x190042  | 0x190160  | 0013-0000-0000-0000    |   0 ns   | 805.0 us  |
| fc1/22 |  102 | 0x190041  | 0x190160  | 0013-0000-0000-0000    |   0 ns   | 401.0 us  |
| fc1/21 |  102 | 0x190042  | 0x190180  | 0000-0000-0000-0000    |  99.0 us | 249.0 us  |
| fc1/22 |  102 | 0x190042  | 0x190160  | 0000-0000-0000-0000    |  85.0 us | 257.0 us  |
| fc1/22 |  102 | 0x190041  | 0x190160  | 0000-0000-0000-0000    |  96.0 us | 238.0 us  |
| fc1/21 |  102 | 0x190046  | 0x190180  | 0002-0000-0000-0000    |  91.0 us | 239.0 us  |
| fc1/28 |  102 | 0x190041  | 0x1900c1  | 001c-0000-0000-0000    |  99.0 us | 215.0 us  |
| fc1/21 |  102 | 0x190043  | 0x190180  | 0013-0000-0000-0000    |   0 ns   | 312.0 us  |
+--------+---------------------------------------------------------+----------------------+
```

This output provides granularity in the form of milliseconds (ms), microseconds (us), and nanoseconds (ns).

Cisco Nexus Dashboard Fabric Controller

This section describes Cisco Nexus Dashboard Fabric Controller (NDFC) capabilities.

Tx and Rx detection

From NDFC, administrators can view the performance utilization of ports. Administrators can select the arrow icon next specific ports and view a historical diagram on the performance of these ports.

Use the following procedure to view the Rx and Tx utilization for the port channel supporting Fabric B.

To access the vantage point, perform the following steps:

Procedure

1. Select **Fabrics** from the navigation tree.
2. Click on the fabric that is experiencing congestion.

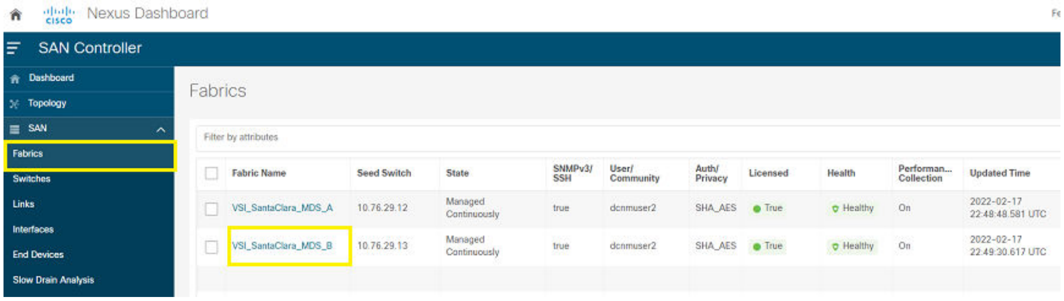

□	Fabric Name	Seed Switch	State	SNMPv3/ SSH	User/ Community	Auth/ Privacy	Licensed	Health	Performan... Collection	Updated Time
□	VSI_SantaClara_MDS_A	10.76.29.12	Managed Continuously	true	dcnmuser2	SHA_AES	● True	Healthy	On	2022-02-17 22:48:48.581 UTC
□	VSI_SantaClara_MDS_B	10.76.29.13	Managed Continuously	true	dcnmuser2	SHA_AES	● True	Healthy	On	2022-02-17 22:49:30.617 UTC

3. In the fabric pop-up window, select the highlighted arrow.

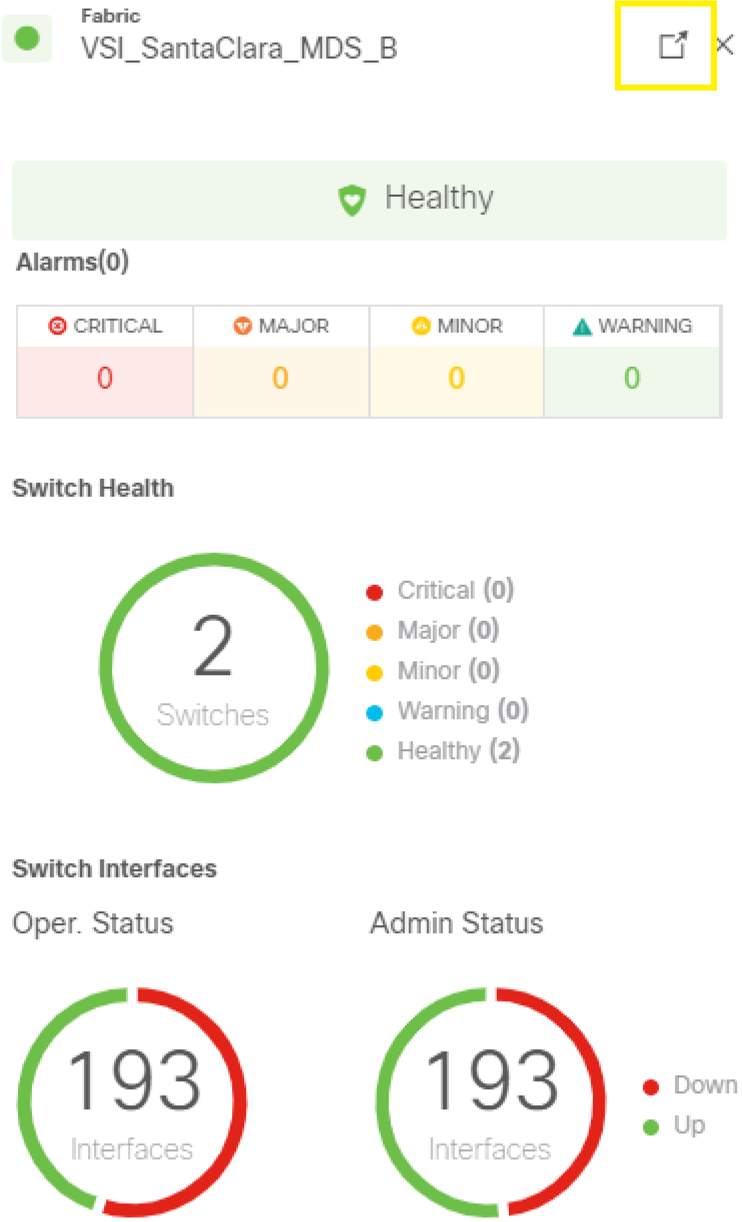

CRITICAL	MAJOR	MINOR	WARNING
0	0	0	0

4. Select **Interfaces** from the menu.
5. Apply filters to list the suspected port.
6. Select the arrow icon next to a specific port based on the performance graph.

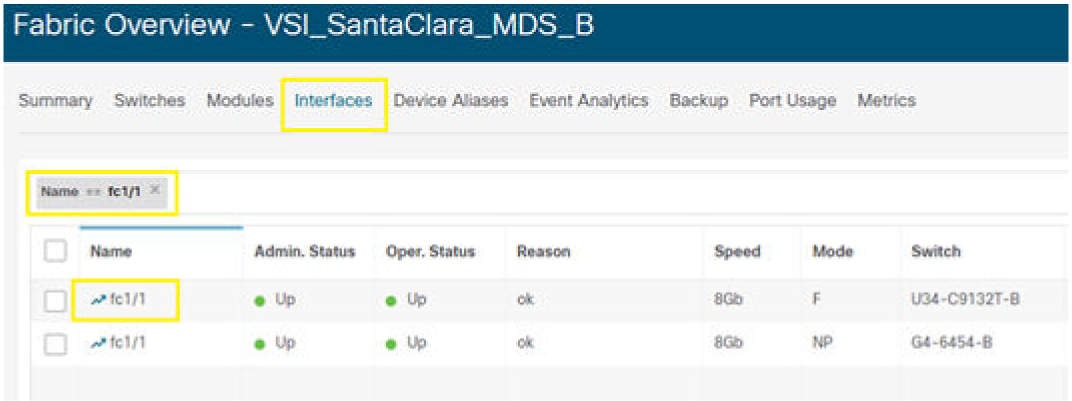

Fabric Overview - VSI_SantaClara_MDS_B

Name == fc1/1 ×

□	Name	Admin. Status	Oper. Status	Reason	Speed	Mode	Switch
□	fc1/1	● Up	● Up	ok	8Gb	F	U34-C9132T-B
□	fc1/1	● Up	● Up	ok	8Gb	NP	G4-6454-B

You will be presented with a performance graph showing Tx and Rx utilization.

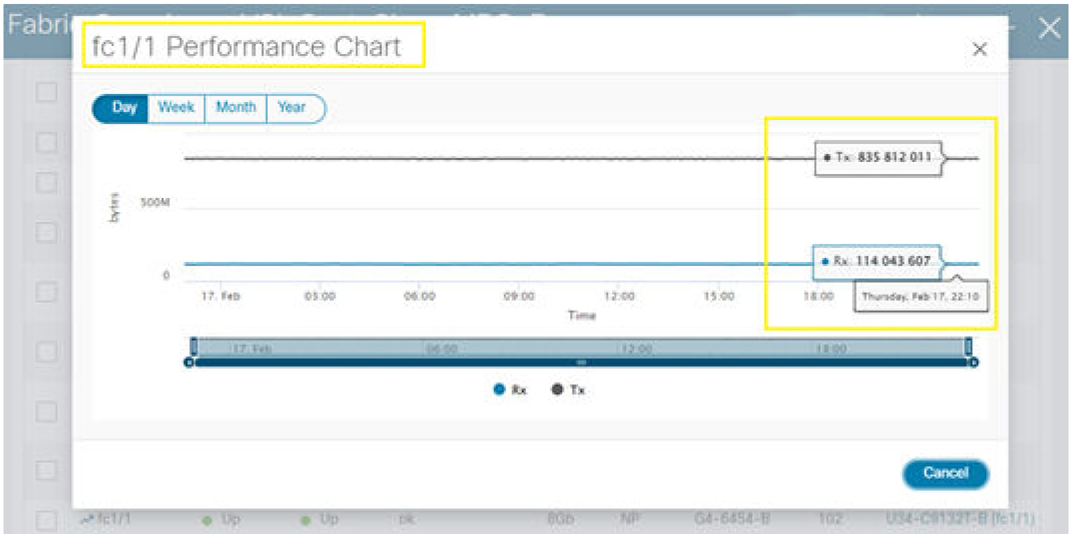

fc1/1 Performance Chart

RxB2Bto0 counters

Additionally, administrators can view native Cisco MDS counters from NDFC to provide insight into port congestion. Administrators must enable slow drain analysis on the fabrics to use this feature. Slow drain analysis must run 24/7 on each fabric.

To view RxB2Bto0 counters, perform the following steps:

Procedure

1. Select **Slow Drain Analysis** from the navigation tree.
2. Select the fabric suspected of congestion.

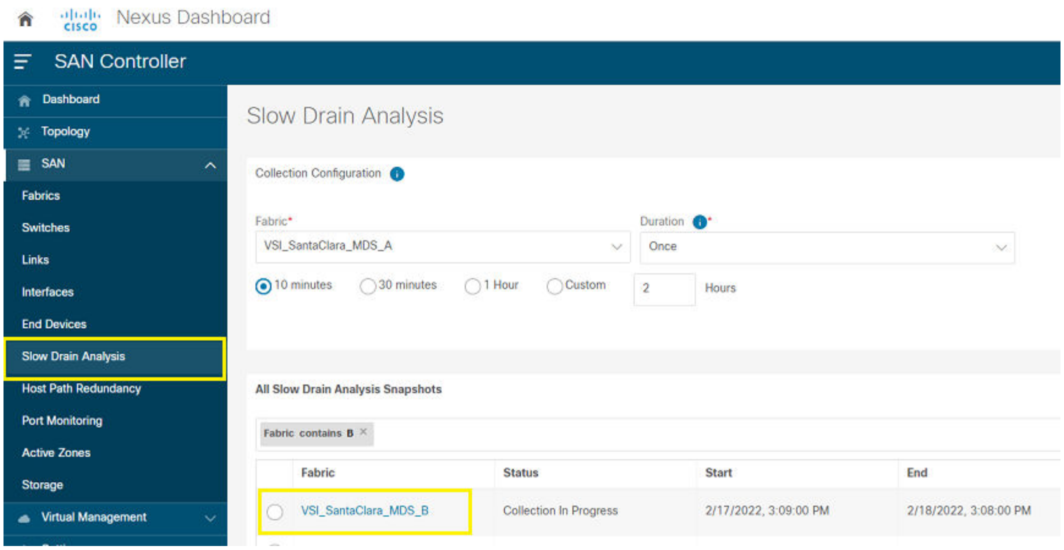

Slow Drain Analysis

Collection Configuration

Fabric*: VSI_SantaClara_MDS_A

Duration*: Once

● 10 minutes ○ 30 minutes ○ 1 Hour ○ Custom 2 Hours

All Slow Drain Analysis Snapshots

Fabric contains B ×

	Fabric	Status	Start	End
○	VSI_SantaClara_MDS_B	Collection In Progress	2/17/2022, 3:09:00 PM	2/18/2022, 3:08:00 PM

The following figure shows a representation of Fibre Channel ports along with the counters. In this example, RxB2Bto0 has a high value indicating ingress congestion on ports FC1/19 and FC 1/20.

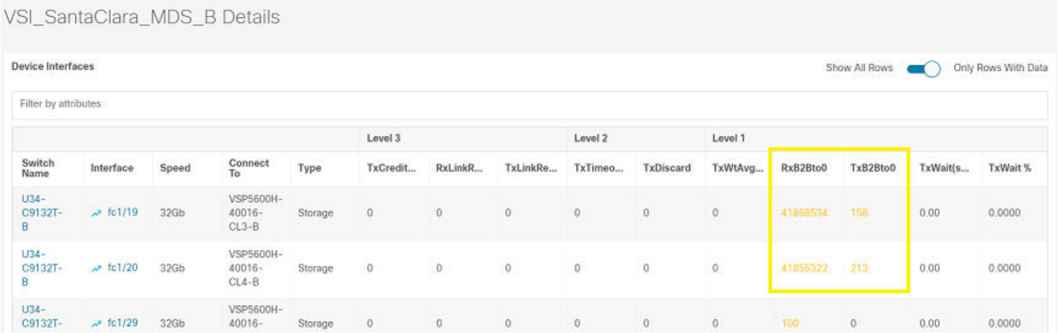
VSI_SantaClara_MDS_B Details

Device Interfaces

					Level 3			Level 2		Level 1				
Switch Name	Interface	Speed	Connect To	Type	TxCredit...	RxLinkR...	TxLinkRe...	TxTimeo...	TxDiscard	TxWtAvg...	RxB2Bto0	TxB2Bto0	TxWait(s...	TxWait %
U34-C9132T-B	fc1/19	32Gb	VSP5600H-40016-CL3-B	Storage	0	0	0	0	0	0	41868534	158	0.00	0.0000
U34-C9132T-B	fc1/20	32Gb	VSP5600H-40016-CL4-B	Storage	0	0	0	0	0	0	41855322	213	0.00	0.0000
U34-C9132T-	fc1/29	32Gb	VSP5600H-40016-	Storage	0	0	0	0	0	0	100	0	0.00	0.0000

Port Monitor logging

After you apply port monitoring polices to the fabric, you can natively view logs from NDFC. The following example shows that port fc 1/1 has set off a port alarm because of a high Tx data rate.

To view logs, perform the following steps:

Procedure

1. From the navigation tree expand **Operations**, and then select **Event Analytics**.
2. Click **Events**.
3. Filter based on the fabric.
4. View the logs.

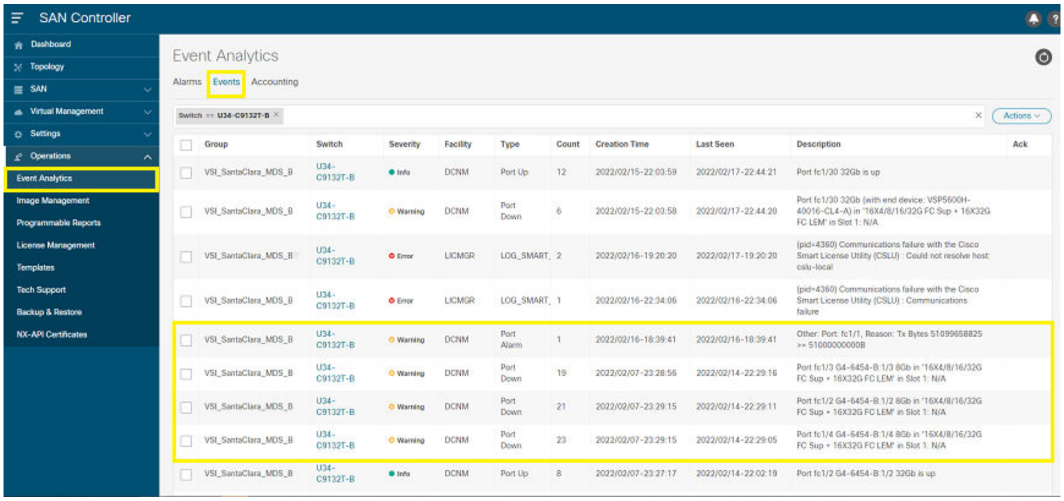

	Group	Switch	Severity	Facility	Type	Count	Creation Time	Last Seen	Description	Ack
□	VSI_SantaClara_MDS_B	U34-C9132T-B	● Info	DCNM	Port Up	12	2022/02/15-22:03:59	2022/02/17-22:44:21	Port fc1/30 32Gb is up	
□	VSI_SantaClara_MDS_B	U34-C9132T-B	● Warning	DCNM	Port Down	6	2022/02/15-22:03:58	2022/02/17-22:44:20	Port fc1/30 32Gb (with end device: VSP5600H-40016-CL4-A) in '16X4/8/16/32G FC Sup + 16X32G FC LEM' in Slot 1: N/A	
□	VSI_SantaClara_MDS_B	U34-C9132T-B	● Error	LICMGR	LOG_SMART_	2	2022/02/16-19:20:20	2022/02/17-19:20:20	(pid=4360) Communications failure with the Cisco Smart License Utility (CSLU) : Could not resolve host: cslu-local	
□	VSI_SantaClara_MDS_B	U34-C9132T-B	● Error	LICMGR	LOG_SMART_	1	2022/02/16-22:34:06	2022/02/16-22:34:06	(pid=4360) Communications failure with the Cisco Smart License Utility (CSLU) : Communications failure	
□	VSI_SantaClara_MDS_B	U34-C9132T-B	● Warning	DCNM	Port Alarm	1	2022/02/16-18:39:41	2022/02/16-18:39:41	Other: Port: fc1/1, Reason: Tx Bytes 51099658825 >= 51000000008	
□	VSI_SantaClara_MDS_B	U34-C9132T-B	● Warning	DCNM	Port Down	19	2022/02/07-23:28:56	2022/02/14-22:29:16	Port fc1/3 G4-6454-B:1/3 8Gb in '16X4/8/16/32G FC Sup + 16X32G FC LEM' in Slot 1: N/A	
□	VSI_SantaClara_MDS_B	U34-C9132T-B	● Warning	DCNM	Port Down	21	2022/02/07-23:29:15	2022/02/14-22:29:11	Port fc1/2 G4-6454-B:1/2 8Gb in '16X4/8/16/32G FC Sup + 16X32G FC LEM' in Slot 1: N/A	
□	VSI_SantaClara_MDS_B	U34-C9132T-B	● Warning	DCNM	Port Down	23	2022/02/07-23:29:15	2022/02/14-22:29:05	Port fc1/4 G4-6454-B:1/4 8Gb in '16X4/8/16/32G FC Sup + 16X32G FC LEM' in Slot 1: N/A	
□	VSI_SantaClara_MDS_B	U34-C9132T-B	● Info	DCNM	Port Up	8	2022/02/07-23:27:17	2022/02/14-22:02:19	Port fc1/2 G4-6454-B:1/2 32Gb is up	

Cisco Nexus Dashboard Fabric Controller

Cisco Nexus Dashboard Fabric Controller (NDFC) provides real time insight into the top 10 busiest and slowest host systems. Select the attribute drop-down on the respective dashboard to view details about storage systems backing the fabric IOPs, throughput, and ECT. To enable this dashboard, SAN Insights (Cisco SAN Analytics) must be configured on the fabric for data collection to occur.

To access the NFDC dashboard, perform the following steps:

Procedure

1. From the navigation tree, select **Dashboard**.
2. Click **SAN Insights**.
3. Select a protocol.

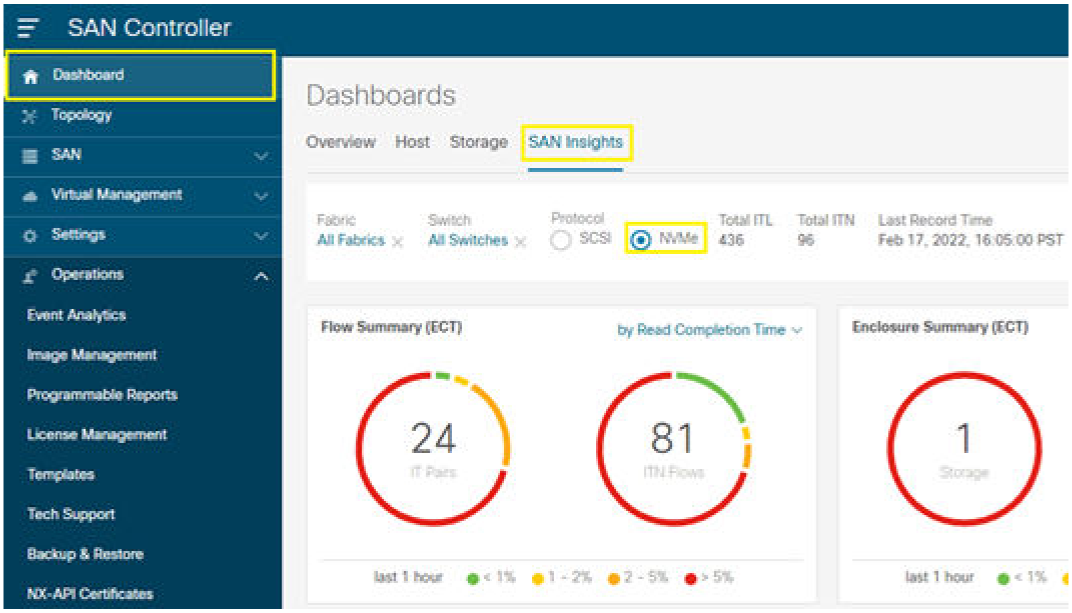

Fabric: All Fabrics
Switch: All Switches
Protocol: ○ SCSI ● NVMe
Total ITL: 436
Total ITN: 96
Last Record Time: Feb 17, 2022, 16:05:00 PST

You will be presented with multiple dashboards that provide insight into the top 10 busiest hosts and storage systems within the fabric.

4. Select other metrics such as IOPs, throughput, or ECT from the <*name of the menu*> menu.

   The following example shows two NVMe hosts with the highest utilization within the environment that helps identify the root cause of congestion caused by overutilization.

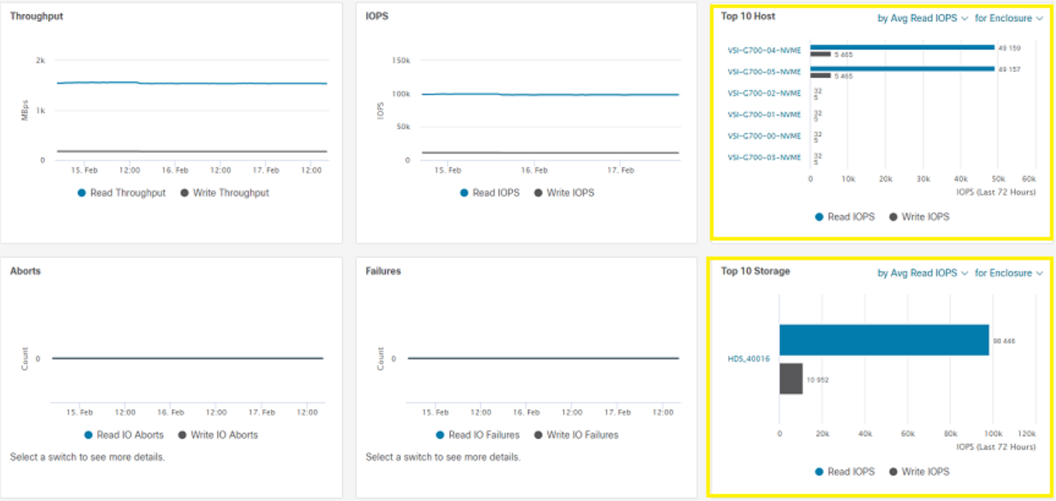

Visibility into UCS Blade, vNIC, and vHBA traffic

After a fabric has been discovered with UCSM credentials and SAN Insights has been enabled on the fabric, you can select the primary fabric interconnect and gain performance and configuration insight into vHBAs, vNIC, and blade performance from NDFC.

To gain insight into these resources, perform the following steps:

Procedure

1. From the navigation tree, select **Switches**.
2. Select the primary Fabric Interconnect.
3. In the Fabric window, select the highlighted arrow.

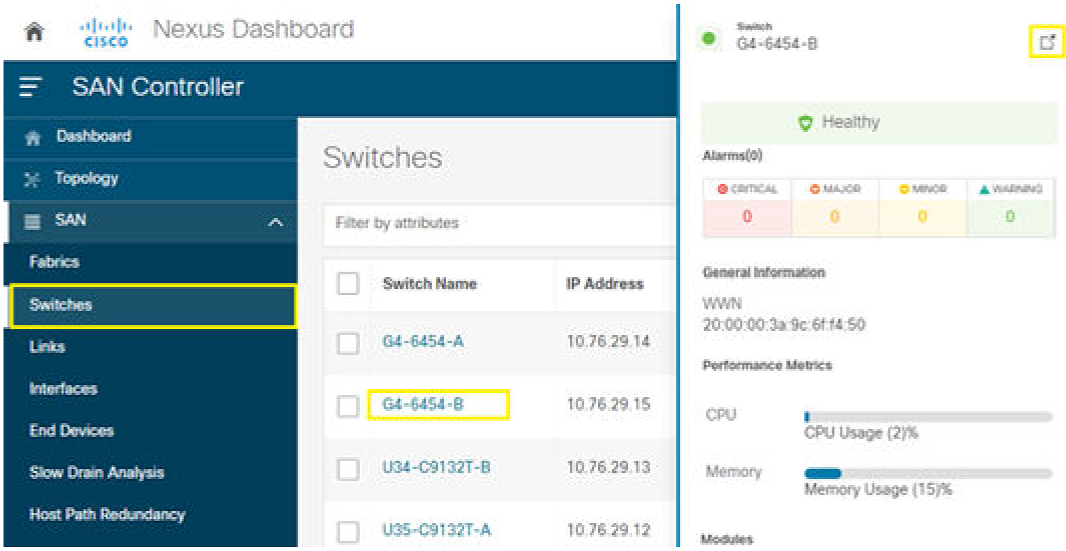

Switch
G4-6454-B

CRITICAL	MAJOR	MINOR	WARNING
0	0	0	0

WWN
20:00:00:3a:9c:6f:f4:50

CPU
CPU Usage (2)%

Memory
Memory Usage (15)%

□	Switch Name	IP Address
□	G4-6454-A	10.76.29.14
□	G4-6454-B	10.76.29.15
□	U34-C9132T-B	10.76.29.13
□	U35-C9132T-A	10.76.29.12

When selecting the primary Fabric Interconnect, you will see tabs at the top of the screen for Blades, vNIC, and vHBA. vHBA is selected in the following example.

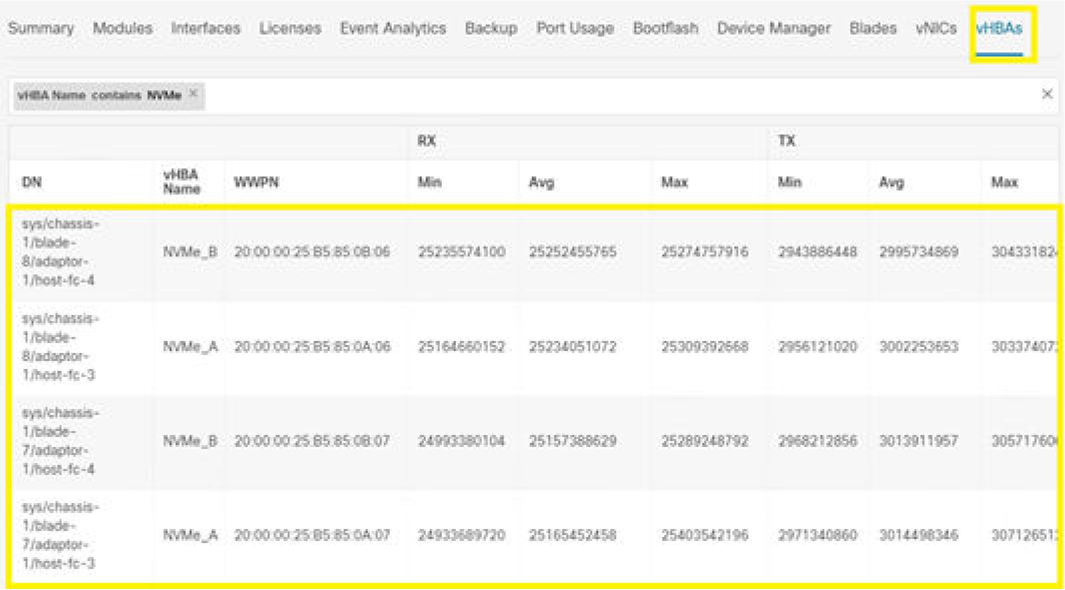

			RX			TX		
DN	vHBA Name	WWPN	Min	Avg	Max	Min	Avg	Max
sys/chassis-1/blade-8/adaptor-1/host-fc-4	NVMe_B	20:00:00:25:B5:85:0B:06	25235574100	25252455765	25274757916	2943886448	2995734869	30433182
sys/chassis-1/blade-8/adaptor-1/host-fc-3	NVMe_A	20:00:00:25:B5:85:0A:06	25164660152	25234051072	25309392668	2956121020	3002253653	30337407
sys/chassis-1/blade-7/adaptor-1/host-fc-4	NVMe_B	20:00:00:25:B5:85:0B:07	24993380104	25157388629	25289248792	2968212856	3013911957	30571760
sys/chassis-1/blade-7/adaptor-1/host-fc-3	NVMe_A	20:00:00:25:B5:85:0A:07	24933689720	25165452458	25403542196	2971340860	3014498346	30712651

Multipath verification to identify the cause of overutilization

NDFC provides administrators insight into SAN performance based on pathing, which correlates host profiles to storage system ports. Cisco SAN Analytics must be enabled on the fabric to use this feature.

To view storage performance on each path and pinpoint the cause of overutilization, perform the following steps:

Procedure

1. From the navigation tree, select **Dashboard**.
2. Click **SAN Insights**.
3. Select **Monitor Metrics**.

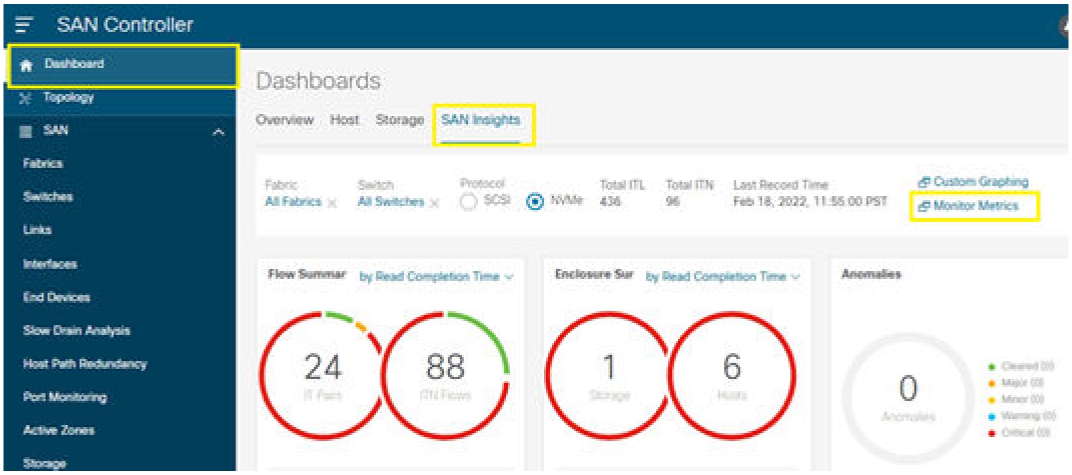

Fabric: All Fabrics ×
Switch: All Switches ×
Protocol: ○ SCSI ● NVMe
Total ITL: 436
Total ITN: 96
Last Record Time: Feb 18, 2022, 11:55:00 PST

4. In the upper-left corner verify that **Viewing NVMe metrics** is visible. Select the appropriate vantage point based on host, storage, or pathing.

The following figure shows that hosts VSI-G700-04-NVME and VSI-G700-05-NVME have higher than average ECT, IOPs, and throughput, which indicates overutilization that stems from these hosts.

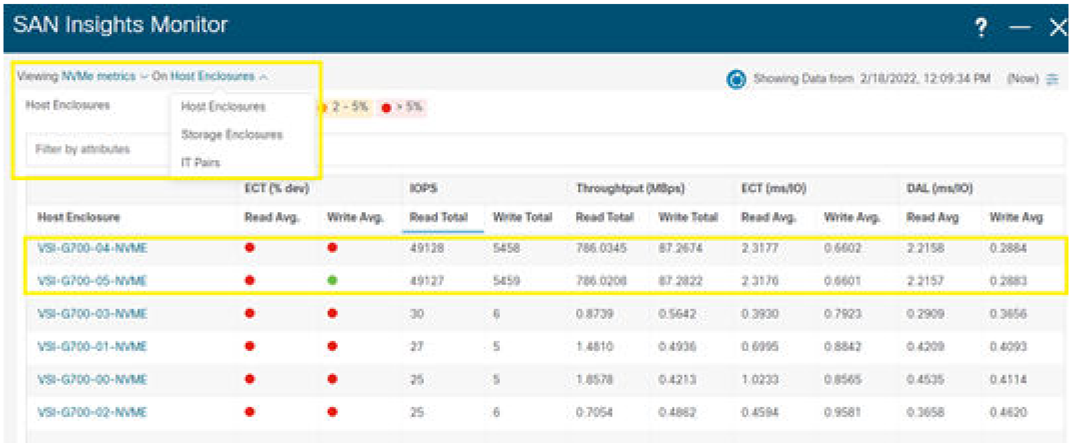

	ECT (% dev)		IOPS		Throughput (MBps)		ECT (ms/IO)		DAL (ms/IO)	
Host Enclosure	Read Avg.	Write Avg.	Read Total	Write Total	Read Total	Write Total	Read Avg.	Write Avg.	Read Avg	Write Avg
VSI-G700-04-NVME	●	●	49128	5458	786.0345	87.2674	2.3177	0.6602	2.2158	0.2884
VSI-G700-05-NVME	●	●	49127	5459	786.0208	87.2822	2.3176	0.6601	2.2157	0.2883
VSI-G700-03-NVME	●	●	30	6	0.8739	0.5642	0.3930	0.7923	0.2909	0.3656
VSI-G700-01-NVME	●	●	27	5	1.4810	0.4936	0.6995	0.8842	0.4209	0.4093
VSI-G700-00-NVME	●	●	25	5	1.8578	0.4213	1.0233	0.8565	0.4535	0.4114
VSI-G700-02-NVME	●	●	25	6	0.7054	0.4862	0.4594	0.9581	0.3658	0.4620

Alternatively, viewing the IT Pairs shows that Hitachi VSP ports CL3-B and CL4-B used by host VSI-G700-04 are showing higher than average utilization and ECT, which is the potential root cause of fabric congestion.

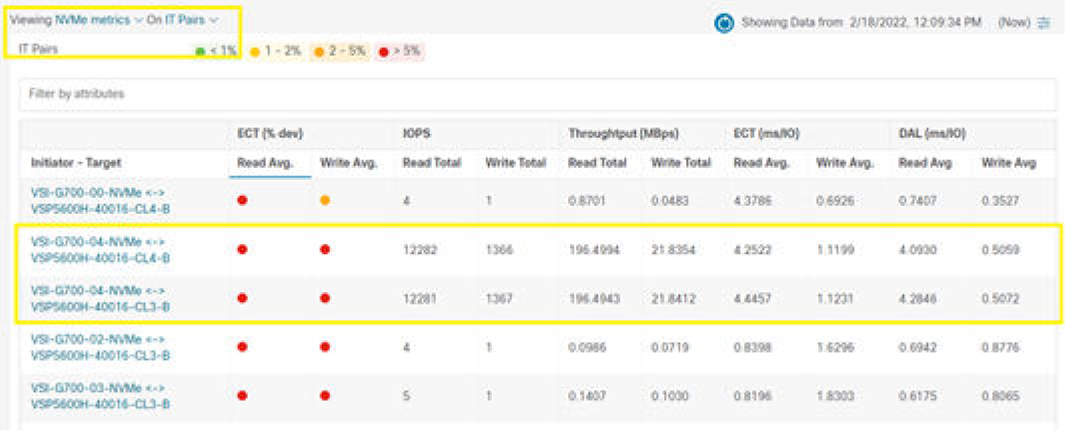
Viewing NVMe metrics ⌄ On IT Pairs ⌄
IT Pairs
● < 1%
● 1 - 2%
● 2 - 5%
● > 5%
Showing Data from 2/18/2022, 12:09:34 PM (Now)
Filter by attributes

	ECT (% dev)		IOPS		Throughput (MBps)		ECT (ms/IO)		DAL (ms/IO)	
Initiator - Target	Read Avg.	Write Avg.	Read Total	Write Total	Read Total	Write Total	Read Avg.	Write Avg.	Read Avg	Write Avg
VSI-G700-00-NVMe <-> VSP5600H-40016-CL4-B	●	●	4	1	0.8701	0.0483	4.3786	0.6926	0.7407	0.3527
VSI-G700-04-NVMe <-> VSP5600H-40016-CL4-B	●	●	12282	1366	196.4994	21.8354	4.2522	1.1199	4.0930	0.5059
VSI-G700-04-NVMe <-> VSP5600H-40016-CL3-B	●	●	12281	1367	196.4943	21.8412	4.4457	1.1231	4.2846	0.5072
VSI-G700-02-NVMe <-> VSP5600H-40016-CL3-B	●	●	4	1	0.0986	0.0719	0.8398	1.6296	0.6942	0.8776
VSI-G700-03-NVMe <-> VSP5600H-40016-CL3-B	●	●	5	1	0.1407	0.1030	0.8196	1.8303	0.6175	0.8065

Performance investigation of initiator, target, and LUN or namespace

SAN Insights (Cisco SAN Analytics) must be enabled to use this feature. From the NDFC dashboard, you can access detailed statistics of hosts that show high utilization. Viewing host information allows administrators to correlate VSP target ports to host initiator ports along with target LUNs. In the case of FC-NVMe, the respective namespace ID is shown.

The following is an example of performance investigation using SAN Analytics in conjunction with the NDFC dashboard to correlate initiator, target, and LUN or Namespace information.

Procedure

1. From NDFC, select **Dashboard** from the navigation tree.
2. Select the **SAN Insights** tab.
3. Select NVMe for the protocol.

4. From the top 10 hosts dashboard, select a host that shows high utilization to investigate. Metrics that can sort top 10 hosts include Read/Write IOPS, Read/Write Throughput, Read/Write ECT, and Read/Write DAL.

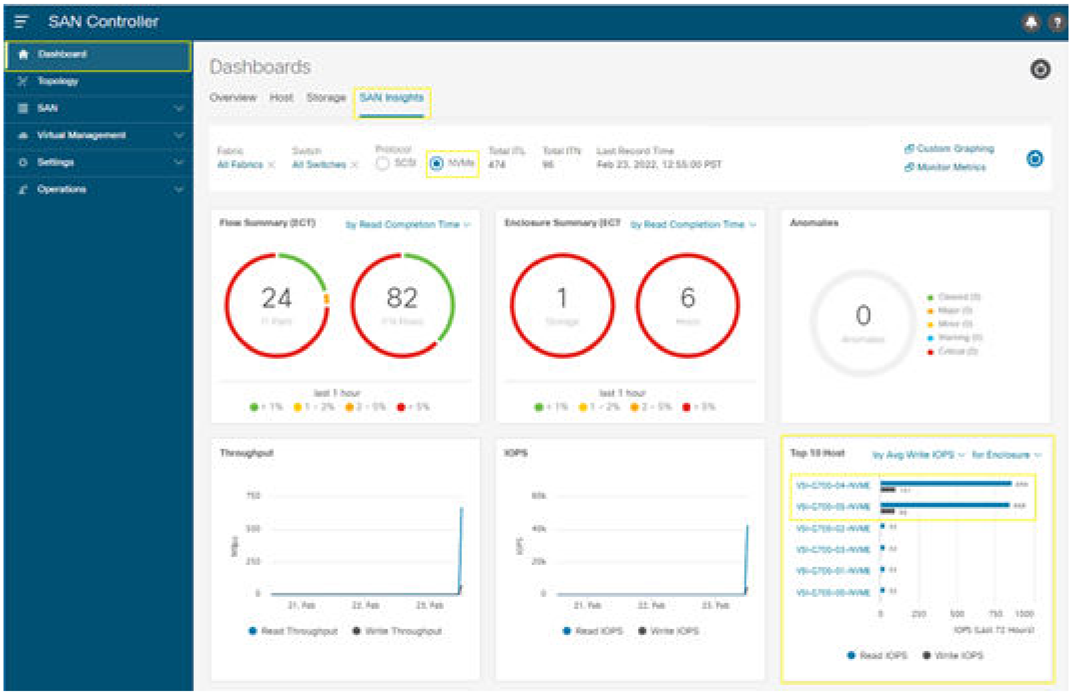

In this example, host VSI-G700-4-NVMe is selected, and its graph metrics page is presented. This provides insight to valuable metrics such as read and write, IOPS, throughput, ECT, DAL, and I/O.

5. Select the **Table** tab to correlate the physical to virtual infrastructure as well as performance trends.

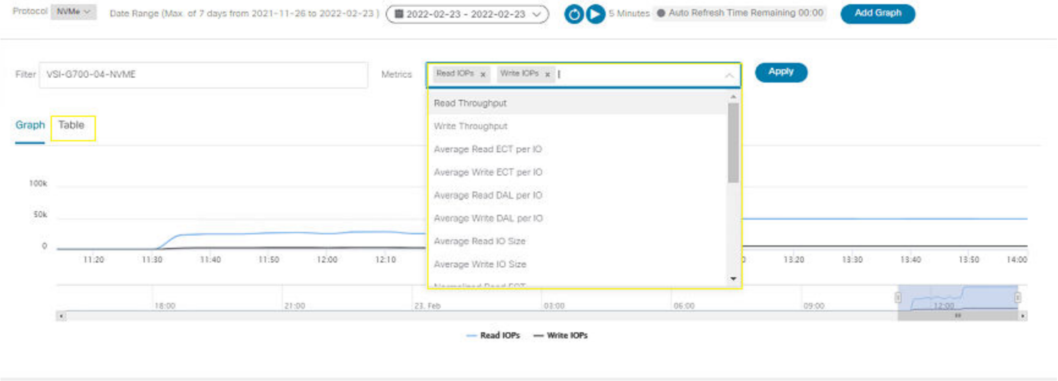

Protocol NVMe

Date Range (Max. of 7 days from 2021-11-26 to 2022-02-23) 2022-02-23 - 2022-02-23

5 Minutes ● Auto Refresh Time Remaining 00:00

Filter VSI-G700-04-NVME

Metrics Read IOPs x Write IOPs x

6. From this vantage point, administrators can investigate host performance trends and correlate VSP target ports, NVMe namespace IDs, and Cisco MDS ports used in the SAN fabric on a per initiator basis. From the **Metrics** menu, administrators can modify the metrics.

   The following example shows VSI-G700-04-NVME, which is using VSP 5600 storage port CL3-B. It is backed by namespace ID 4 and shows above average ECT per I/O compared to its respective counterparts, which indicates that there is performance degradation and further investigation is needed for these resources to pinpoint the cause of congestion within fabric B.

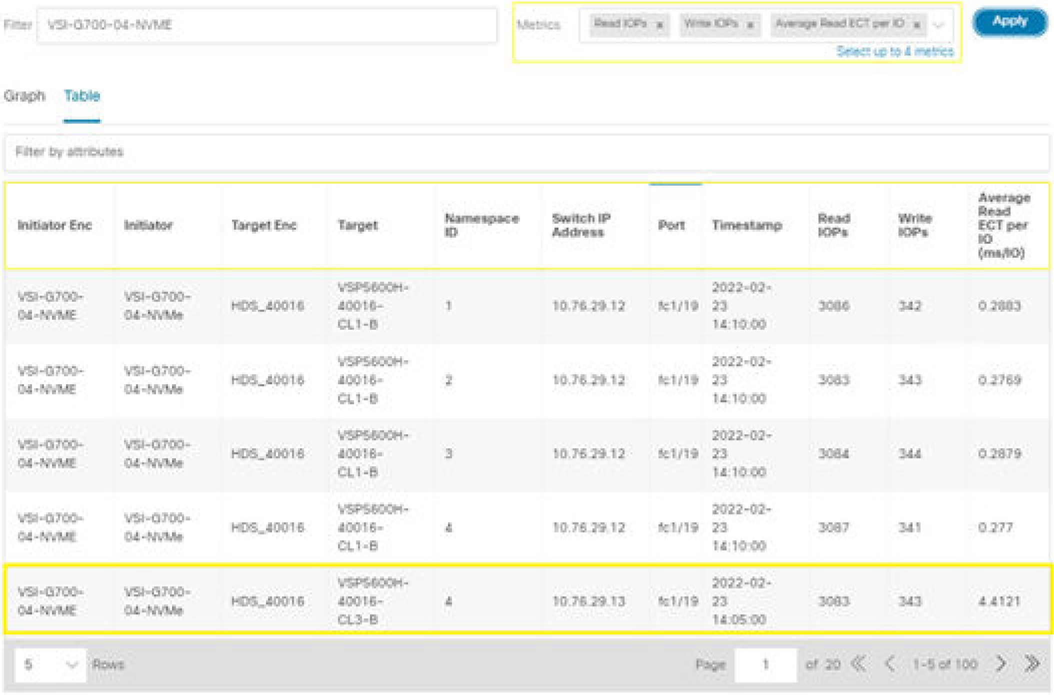

Initiator Enc	Initiator	Target Enc	Target	Namespace ID	Switch IP Address	Port	Timestamp	Read IOPs	Write IOPs	Average Read ECT per IO (ms/IO)
VSI-G700-04-NVME	VSI-G700-04-NVMe	HDS_40016	VSP5600H-40016-CL1-B	1	10.76.29.12	fc1/19	2022-02-23 14:10:00	3086	342	0.2883
VSI-G700-04-NVME	VSI-G700-04-NVMe	HDS_40016	VSP5600H-40016-CL1-B	2	10.76.29.12	fc1/19	2022-02-23 14:10:00	3083	343	0.2769
VSI-G700-04-NVME	VSI-G700-04-NVMe	HDS_40016	VSP5600H-40016-CL1-B	3	10.76.29.12	fc1/19	2022-02-23 14:10:00	3084	344	0.2879
VSI-G700-04-NVME	VSI-G700-04-NVMe	HDS_40016	VSP5600H-40016-CL1-B	4	10.76.29.12	fc1/19	2022-02-23 14:10:00	3087	341	0.277
VSI-G700-04-NVME	VSI-G700-04-NVMe	HDS_40016	VSP5600H-40016-CL3-B	4	10.76.29.13	fc1/19	2022-02-23 14:05:00	3083	343	4.4121

Hitachi Ops Center Analyzer best practices

This section describes how to identify, troubleshoot, and resolve performance congestion using the Ops Center Analyzer.

Ops Center Analyzer dashboard

Ops Center Analyzer dashboards are visual representations of the performance metrics of your infrastructure resources. The consolidated view helps administrators quickly interpret performance metrics and identify performance problems. The consolidated dashboard view allows for the unified management of the server, storage, and network infrastructure resources. You can ensure the health of your datacenter by proactively monitoring consumer groups, storage components, volumes, VMs, servers, and network devices.

Analyzer allows grouping of custom resources known as consumers so administrators can easily distinguish resources based on customer, region, or usage. The advanced visual analytics aid in visualizing the performance data in easy-to-use graphs and charts. These visual cues allow for intuitive performance management.

The following example shows an Analyzer dashboard indicating performance issues related to VMs and block-based storage LDEVs as well as a consumer resource.

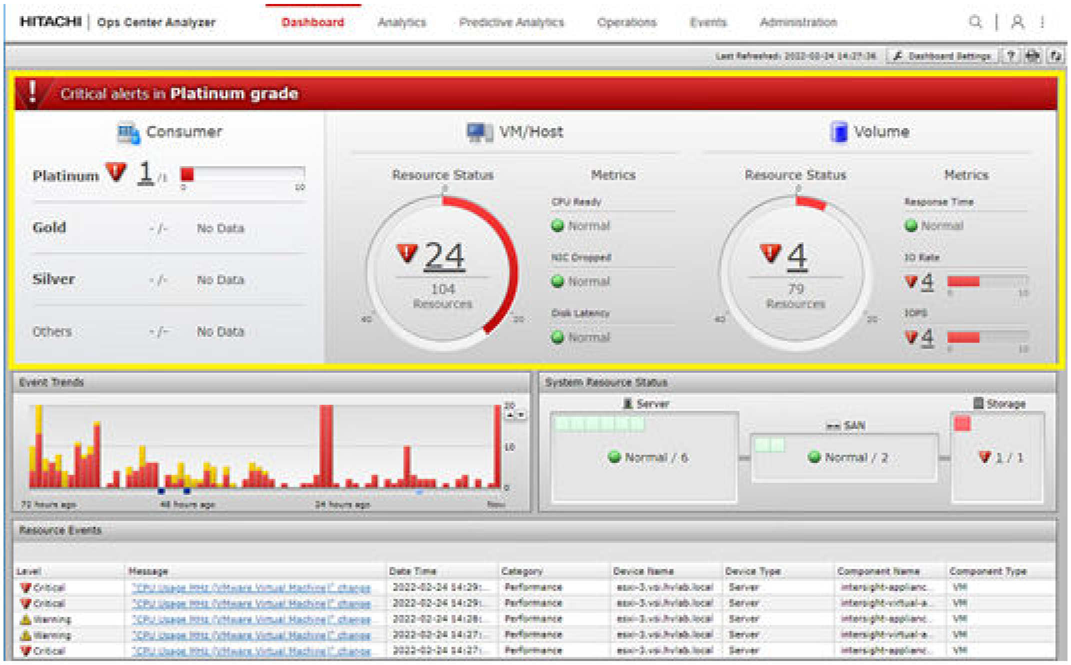

Level	Message	Date Time	Category	Device Name	Device Type	Component Name	Component Type
Critical	"CPU Usage MHz (VMware Virtual Machine)" change...	2022-02-24 14:29:...	Performance	esxi-3.vsi.hvlab.local	Server	intersight-applianc...	VM
Critical	"CPU Usage MHz (VMware Virtual Machine)" change...	2022-02-24 14:29:...	Performance	esxi-3.vsi.hvlab.local	Server	intersight-virtual-a...	VM
Warning	"CPU Usage MHz (VMware Virtual Machine)" change...	2022-02-24 14:28:...	Performance	esxi-3.vsi.hvlab.local	Server	intersight-applianc...	VM
Warning	"CPU Usage MHz (VMware Virtual Machine)" change...	2022-02-24 14:27:...	Performance	esxi-3.vsi.hvlab.local	Server	intersight-virtual-a...	VM
Critical	"CPU Usage MHz (VMware Virtual Machine)" change...	2022-02-24 14:27:...	Performance	esxi-3.vsi.hvlab.local	Server	intersight-applianc...	VM

End-2-End (E2E) data analysis

Performance analysis starts with understanding whether any E2E data path resources are overloaded. By using E2E data analysis in Ops Center Analyzer, administrators can identify the root cause of performance degradation within VMs, hosts, SAN fabric, and VSP storage systems by viewing a visual representation of the end-to-end data path within the datacenter. From the E2E view, administrators are presented with easy to identify markers indicating exactly which resources are causing issues.

To access the E2E view from the dashboard, perform the following steps:

Procedure

1. Log in to Hitachi Ops Center Analyzer.
2. From the dashboard, select the impacted resource based on consumer, VM/host, or volume.

   In this example Platinum consumer is selected, and it shows that consumer resource Cisco UCS consumer has performance degradation because of exceeded VM and Volume thresholds.
3. Click **Show volume E2E View**.

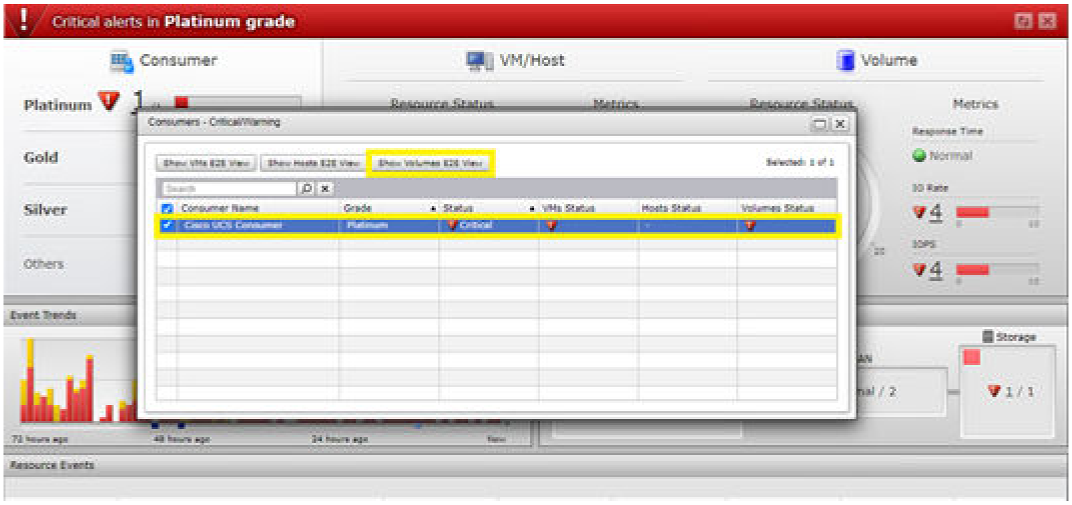

This vantage point shows the visual representation of VSP storage volumes and their association with the SAN fabric. LDEV IDs 0E, 0F, 10, and 11 have performance degradation alerts as well as the Cisco MDS 9132T SAN fabric switch that they use. The VSP storage resources that support these volumes such as ports, DKC processors, cache, VSP pools, and parity groups are also highlighted. This visual relation provides critical information for troubleshooting the SAN fabric with all end devices shown in correlation.

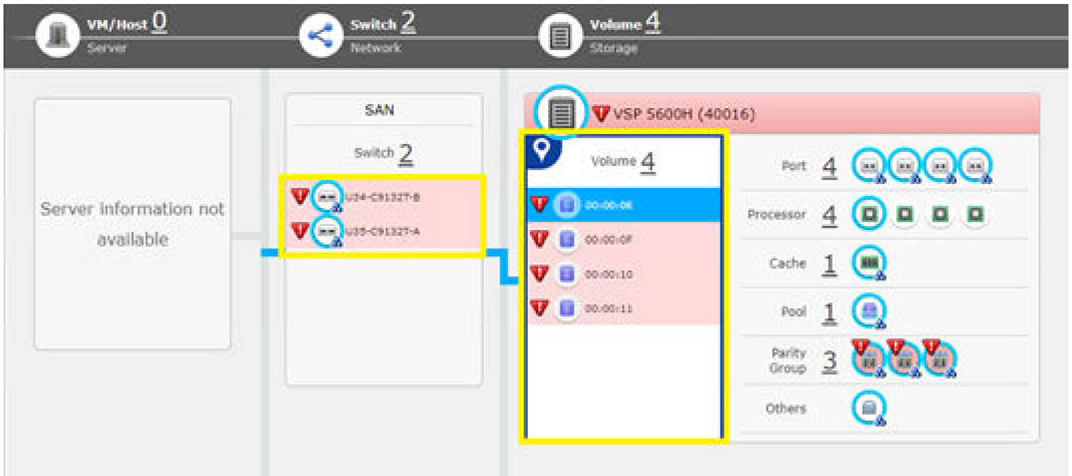

Note: Analyzer 10.8.1 does not visually associate VM/hosts to utilized NVMe resources; FC-SCSI resources will be available. FC-NVMe VM/host visualization will be included in future updates of Ops Center Analyzer. At the time of writing this paper, host association for FC-NVMe devices is a manual process from the **Show Details** view, which is covered in the following section.

Show Detail view

Select a resource, right-click on the object, and then select **Show Detail** to view the resource performance summary report. The resource performance summary report opens in a new window.

In this example, to gain additional insight into storage configuration and performance, right-click on LDEV 0E and select Show Detail. LDEV 0E is an NVMe device using the FC-NVMe protocol. This best practice also applies to VMs, hosts, and SAN switches.

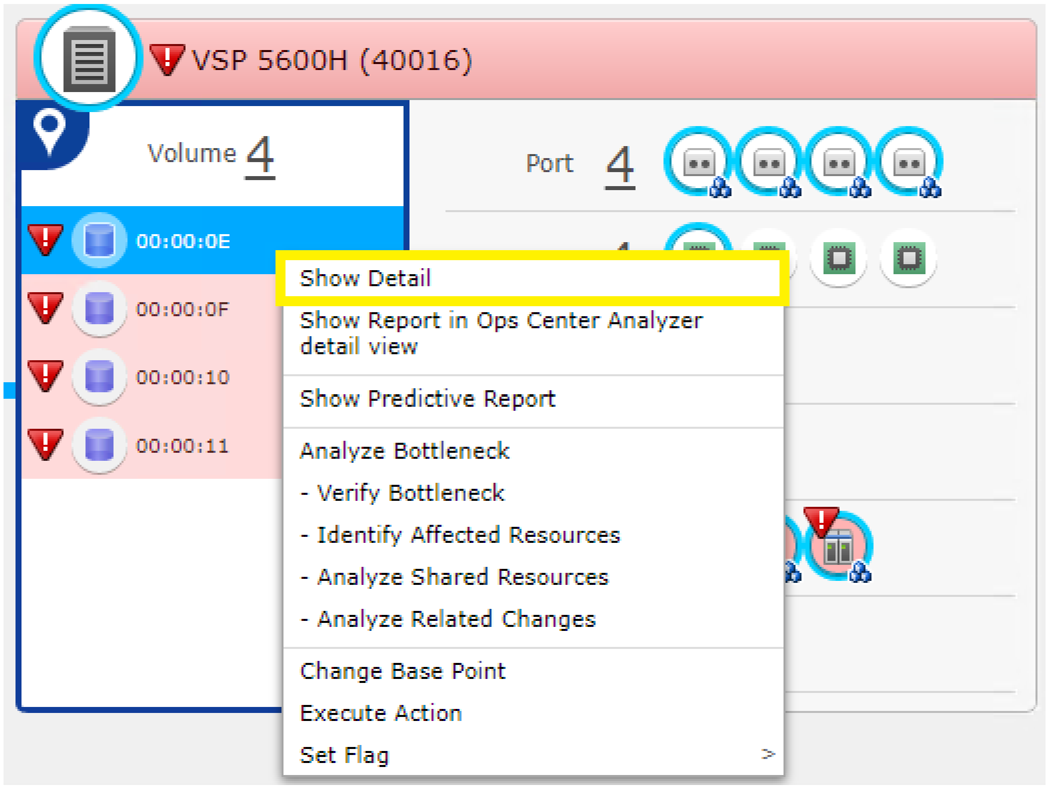

Basic information

The following example shows information such as VSP Volume Label, NVM Subsystem ID/ Nickname, NVM Namespace ID, LDEV and capacity status. It also shows information about which end hosts use these NVMe resources from the NVM Host NQN. Additionally, from Show Detail, the Trend, Events, and Change History tabs provide additional information about the cause of performance degradation on the fabric.

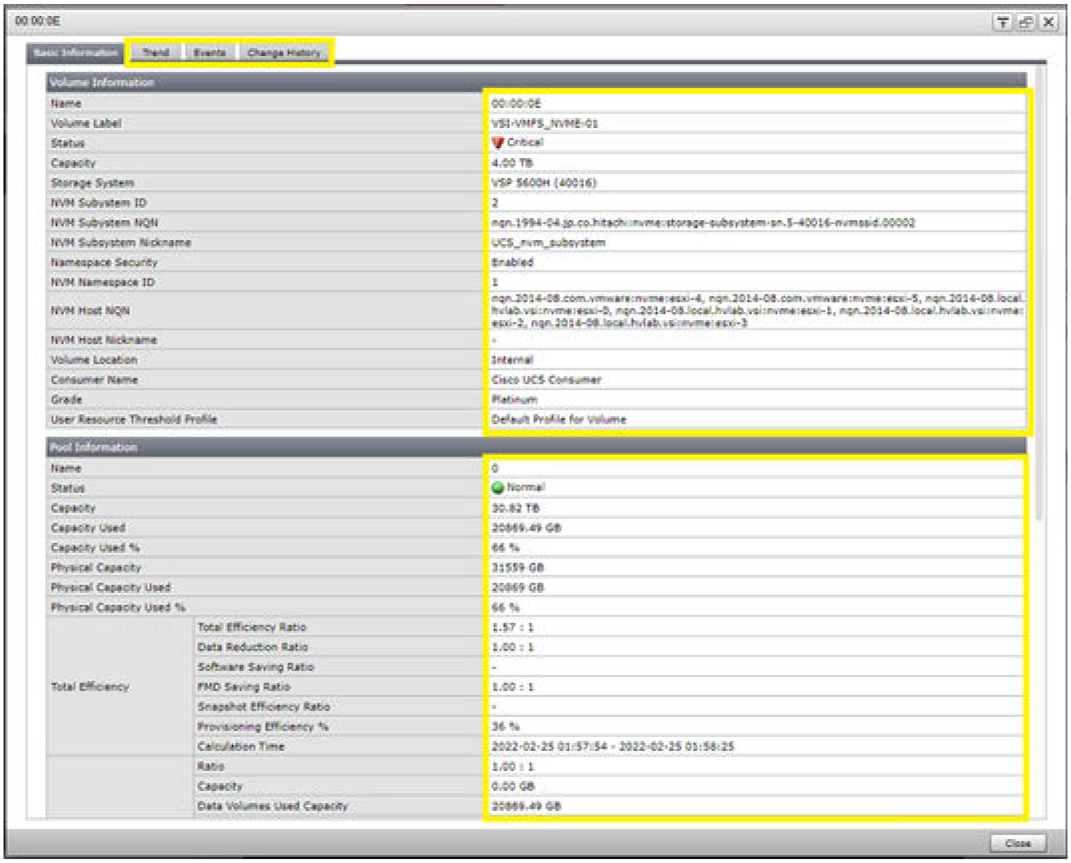

Volume Information

Name	00:00:0E
Volume Label	VSI-VMFS_NVME-01
Status	Critical
Capacity	4.00 TB
Storage System	VSP 5600H (40016)
NVM Subsystem ID	2
NVM Subsystem NQN	nqn.1994-04.jp.co.hitachi:nvme:storage-subsystem-sn.5-40016-nvmssid.00002
NVM Subsystem Nickname	UCS_nvm_subsystem
Namespace Security	Enabled
NVM Namespace ID	1
NVM Host NQN	nqn.2014-08.com.vmware:nvme:esxi-4, nqn.2014-08.com.vmware:nvme:esxi-5, nqn.2014-08.local.hvlab.vsi:nvme:esxi-0, nqn.2014-08.local.hvlab.vsi:nvme:esxi-1, nqn.2014-08.local.hvlab.vsi:nvme:esxi-2, nqn.2014-08.local.hvlab.vsi:nvme:esxi-3
NVM Host Nickname	-
Volume Location	Internal
Consumer Name	Cisco UCS Consumer
Grade	Platinum
User Resource Threshold Profile	Default Profile for Volume

Pool Information

Name		0
Status		Normal
Capacity		30.82 TB
Capacity Used		20869.49 GB
Capacity Used %		66 %
Physical Capacity		31559 GB
Physical Capacity Used		20869 GB
Physical Capacity Used %		66 %
Total Efficiency	Total Efficiency Ratio	1.57 : 1
	Data Reduction Ratio	1.00 : 1
	Software Saving Ratio	-
	FMD Saving Ratio	1.00 : 1
	Snapshot Efficiency Ratio	-
	Provisioning Efficiency %	36 %
	Calculation Time	2022-02-25 01:57:54 - 2022-02-25 01:58:25
	Ratio	1.00 : 1
	Capacity	0.00 GB
	Data Volumes Used Capacity	20869.49 GB

Trend

Select the Trends tab to understand the critical performance degradation of an LDEV. Based on configured resource profiles, administrators can set static or dynamic thresholds to receive alerts after the resources pass the specified thresholds.

The following example shows that LDEV 0E has critical alerts for high IOPs and Transfer Rate utilization. Select the resource and Analyzer will provide a historical graph of the selected resource performance ranging from the last hour to the last 14 days.

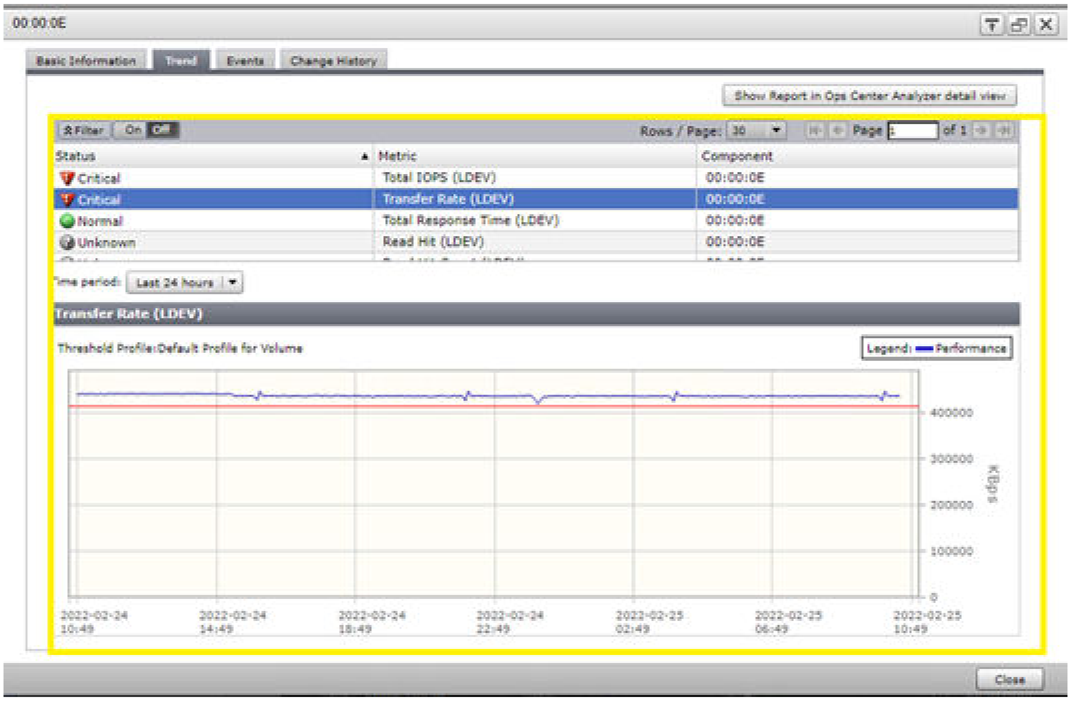

Status	Metric	Component
Critical	Total IOPS (LDEV)	00:00:0E
Critical	Transfer Rate (LDEV)	00:00:0E
Normal	Total Response Time (LDEV)	00:00:0E
Unknown	Read Hit (LDEV)	00:00:0E

Time period: Last 24 hours

Transfer Rate (LDEV)

Threshold Profile:Default Profile for Volume

Events

Select the Events tab to access a historical log of any performance events related to the selected resource.

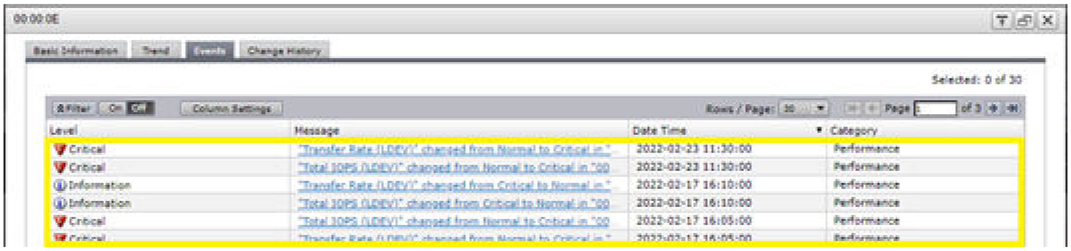

Level	Message	Date Time	Category
Critical	"Transfer Rate (LDEV)" changed from Normal to Critical in "	2022-02-23 11:30:00	Performance
Critical	"Total IOPS (LDEV)" changed from Normal to Critical in "00	2022-02-23 11:30:00	Performance
Information	"Transfer Rate (LDEV)" changed from Critical to Normal in "	2022-02-17 16:10:00	Performance
Information	"Total IOPS (LDEV)" changed from Critical to Normal in "00	2022-02-17 16:10:00	Performance
Critical	"Total IOPS (LDEV)" changed from Normal to Critical in "00	2022-02-17 16:05:00	Performance
Critical	"Transfer Rate (LDEV)" changed from Normal to Critical in "	2022-02-17 16:05:00	Performance

Change History

The Change History tab provides a chronological timeline of any changes to resources including storage, SAN switches, VMs, and hosts. In this example, the Change History tab is selected on the supporting SAN switch, which shows that the firmware was upgraded on this resource recently.

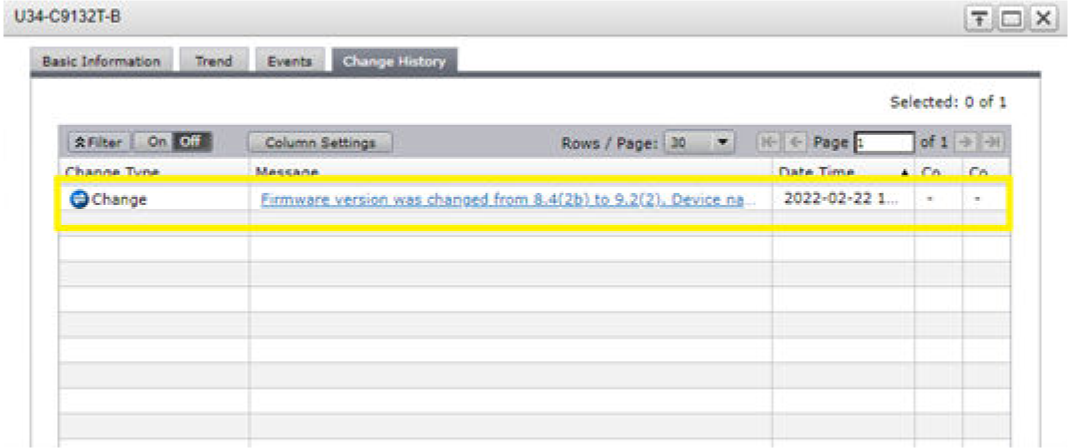

U34-C9132T-B

Basic Information | Trend | Events | Change History

Selected: 0 of 1

Filter On Off | Column Settings | Rows / Page: 30 | Page 1 of 1

Change Type	Message	Date Time	Co...	Co...
Change	Firmware version was changed from 8.4(2b) to 9.2(2). Device na...	2022-02-22 1...	-	-

Analyze Bottleneck

Ops Center Analyzer offers multiple troubleshooting tools for isolating a bottleneck candidate and identifying the root cause. Access the tools shown in the following figure for further analysis by right-clicking on an object and selecting Analyze Bottleneck.

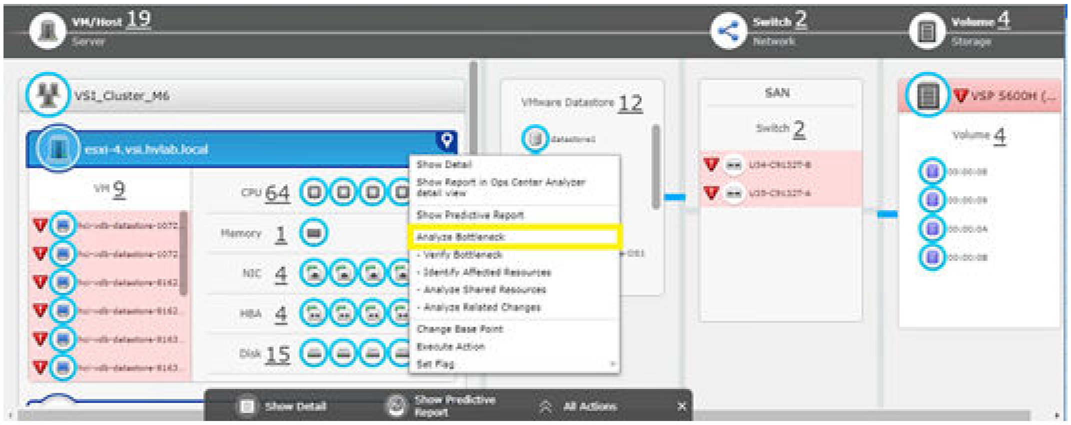

Verify Bottleneck

Use Verify Bottleneck at the initial stage of analysis to compare performance charts of the base point of analysis with the bottlenecked candidate. The following example shows a VM with high VM CPU usage, which you can compare to the physical hypervisor. Administrators can change viewed metrics based on the selected resource from the Metric menu. In the following example, you can conclude that the VM selected shows above average utilization compared to the host.

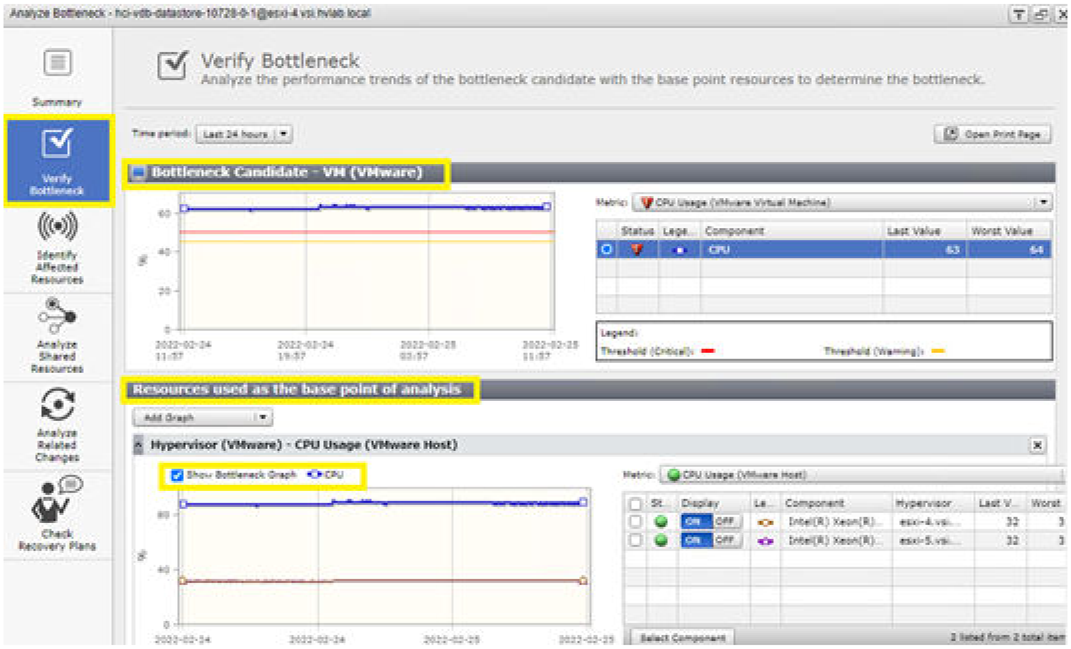

Analyze Bottleneck - hci-vdb-datastore-10728-0-1@esxi-4.vsi.hvlab.local

Verify Bottleneck

Analyze the performance trends of the bottleneck candidate with the base point resources to determine the bottleneck.

Time period: Last 24 hours

Bottleneck Candidate - VM (VMware)

Metric: CPU Usage (VMware Virtual Machine)

	Status	Lege...	Component	Last Value	Worst Value
○	🟥	●	CPU	63	64

Legend:
Threshold (Critical): 🟥 Threshold (Warning): 🟨

Resources used as the base point of analysis

Hypervisor (VMware) - CPU Usage (VMware Host)

☑ Show Bottleneck Graph ● CPU

Metric: CPU Usage (VMware Host)

□	St...	Display	Le...	Component	Hypervisor	Last V...	Worst
□	●	ON OFF	●	Intel(R) Xeon(R)...	esxi-4.vsi...	32	3
□	●	ON OFF	●	Intel(R) Xeon(R)...	esxi-5.vsi...	32	3

Identify Affected Resources

Use Identify Affected Resources to display the user resources that rely on the bottlenecked resource. In the following example, the Cisco MDS 9132T SAN switch supporting fabric B is selected. Viewing Identify Affected Resources, shows that VSP NVMe LDEVs 0E, 0F,10, and 11 rely on this Cisco MDS SAN switch.

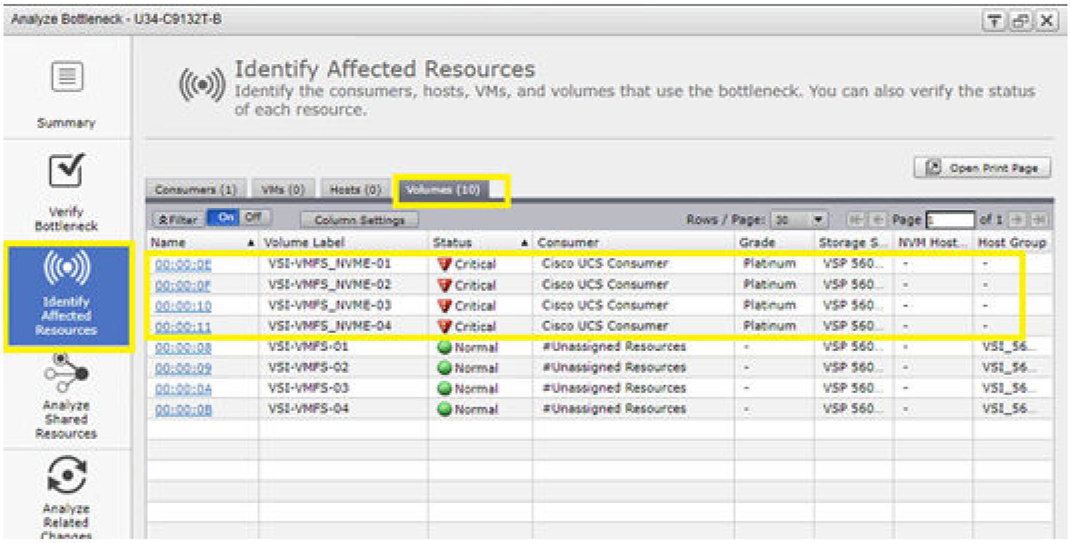

Analyze Bottleneck - U34-C9132T-B

Identify Affected Resources

Identify the consumers, hosts, VMs, and volumes that use the bottleneck. You can also verify the status of each resource.

Rows / Page: 30 Page 1 of 1

Name	Volume Label	Status	Consumer	Grade	Storage S...	NVM Host...	Host Group
00:00:0E	VSI-VMFS_NVME-01	Critical	Cisco UCS Consumer	Platinum	VSP 560...	-	-
00:00:0F	VSI-VMFS_NVME-02	Critical	Cisco UCS Consumer	Platinum	VSP 560...	-	-
00:00:10	VSI-VMFS_NVME-03	Critical	Cisco UCS Consumer	Platinum	VSP 560...	-	-
00:00:11	VSI-VMFS_NVME-04	Critical	Cisco UCS Consumer	Platinum	VSP 560...	-	-
00:00:08	VSI-VMFS-01	Normal	#Unassigned Resources	-	VSP 560...	-	VSI_56...
00:00:09	VSI-VMFS-02	Normal	#Unassigned Resources	-	VSP 560...	-	VSI_56...
00:00:0A	VSI-VMFS-03	Normal	#Unassigned Resources	-	VSP 560...	-	VSI_56...
00:00:0B	VSI-VMFS-04	Normal	#Unassigned Resources	-	VSP 560...	-	VSI_56...

Analyze Shared Resources

Use Analyze Shared Resources if you suspect that the root cause of the problem is resource contention, a noisy neighbor that disrupts the balance of resource usage. Compare performance charts of the bottleneck candidate to the resources using the bottleneck. After comparing performance across several resources with Analyze Shared Resources, isolate the actual bottleneck. In the following example, the hypervisor is selected as the bottleneck candidate, which you can compare to other VMs that use this candidate to see if there is performance degradation.

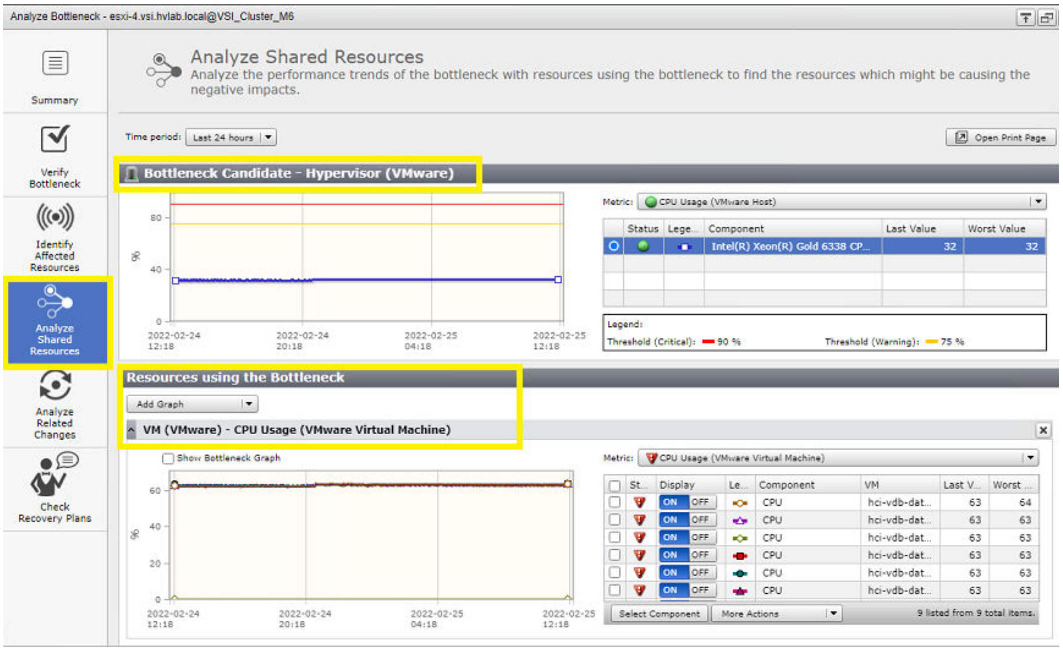

Analyze Bottleneck - esxi-4.vsi.hvlab.local@VSI_Cluster_M6

Analyze Shared Resources

Analyze the performance trends of the bottleneck with resources using the bottleneck to find the resources which might be causing the negative impacts.

Time period: Last 24 hours

Bottleneck Candidate - Hypervisor (VMware)

Metric: CPU Usage (VMware Host)

	Status	Lege...	Component	Last Value	Worst Value
○	●		Intel(R) Xeon(R) Gold 6338 CP...	32	32

Legend:
Threshold (Critical): 🟥 90 %   Threshold (Warning): 🟨 75 %

Resources using the Bottleneck

VM (VMware) - CPU Usage (VMware Virtual Machine)

Metric: CPU Usage (VMware Virtual Machine)

□	St...	Display	Le...	Component	VM	Last V...	Worst ...
□		ON OFF		CPU	hci-vdb-dat...	63	64
□		ON OFF		CPU	hci-vdb-dat...	63	63
□		ON OFF		CPU	hci-vdb-dat...	63	63
□		ON OFF		CPU	hci-vdb-dat...	63	63
□		ON OFF		CPU	hci-vdb-dat...	63	63
□		ON OFF		CPU	hci-vdb-dat...	63	63

9 listed from 9 total items.

Analyze Related Changes

Using Analyze Shared Resources does not reveal the actual bottleneck (noisy neighbor), or that the root cause of the problem is a recent configuration change. Use **Analyze Related Changes** to compare performance charts with configuration events. The bar graph portion of the chart represents the configuration changes made at a specific time. You can click on a bar to list the changes.

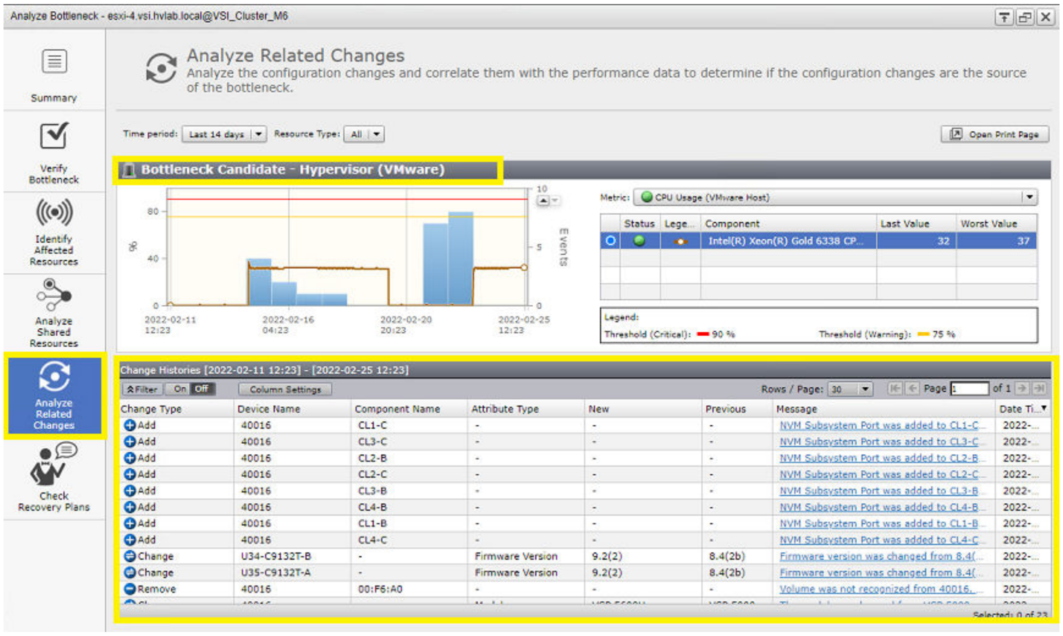

Analyze Related Changes

Analyze the configuration changes and correlate them with the performance data to determine if the configuration changes are the source of the bottleneck.

Time period: Last 14 days   Resource Type: All

Bottleneck Candidate - Hypervisor (VMware)

Metric: CPU Usage (VMware Host)

	Status	Lege...	Component	Last Value	Worst Value
○	●		Intel(R) Xeon(R) Gold 6338 CP...	32	37

Legend:

Threshold (Critical): 🟥 90 %   Threshold (Warning): 🟨 75 %

Change Histories [2022-02-11 12:23] - [2022-02-25 12:23]

Rows / Page: 30   Page 1 of 1

Change Type	Device Name	Component Name	Attribute Type	New	Previous	Message	Date Ti...
Add	40016	CL1-C	-	-	-	NVM Subsystem Port was added to CL1-C...	2022-...
Add	40016	CL3-C	-	-	-	NVM Subsystem Port was added to CL3-C...	2022-...
Add	40016	CL2-B	-	-	-	NVM Subsystem Port was added to CL2-B...	2022-...
Add	40016	CL2-C	-	-	-	NVM Subsystem Port was added to CL2-C...	2022-...
Add	40016	CL3-B	-	-	-	NVM Subsystem Port was added to CL3-B...	2022-...
Add	40016	CL4-B	-	-	-	NVM Subsystem Port was added to CL4-B...	2022-...
Add	40016	CL1-B	-	-	-	NVM Subsystem Port was added to CL1-B...	2022-...
Add	40016	CL4-C	-	-	-	NVM Subsystem Port was added to CL4-C...	2022-...
Change	U34-C9132T-B	-	Firmware Version	9.2(2)	8.4(2b)	Firmware version was changed from 8.4(...	2022-...
Change	U35-C9132T-A	-	Firmware Version	9.2(2)	8.4(2b)	Firmware version was changed from 8.4(...	2022-...
Remove	40016	00:F6:A0	-	-	-	Volume was not recognized from 40016....	2022-...

Selected: 0 of 23

Sparkline View

Use the Sparkline View to analyze the health and performance of the resources in the datacenter monitoring environment. The Sparkline View displays performance reports for multiple nodes in the same pane for a quick comparison between different nodes. You can display detailed performance metrics for each node and find a correlation with other nodes. In the following example, a VSP 5600H is selected as the sparkline candidate, where you can compare VSP resources such as port, processor, cache, pool, and parity group performance in a single pane.

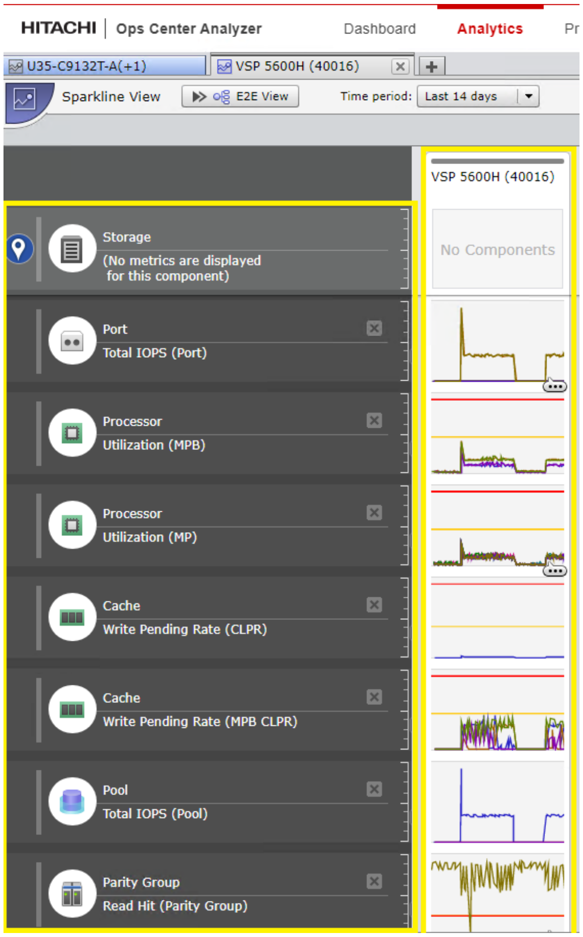

Time period: Last 14 days

VSP 5600H (40016)

No Components

Storage
(No metrics are displayed for this component)

Port
Total IOPS (Port)

Processor
Utilization (MPB)

Processor
Utilization (MP)

Cache
Write Pending Rate (CLPR)

Cache
Write Pending Rate (MPB CLPR)

Pool
Total IOPS (Pool)

Parity Group
Read Hit (Parity Group)

Monitor settings

Define the monitoring conditions for detecting deterioration in the service performance of virtual machines and volumes using dynamic thresholds or static thresholds. You can create rules and conditions to automate resource assignment to monitoring profiles. Using these rules, the newly discovered user resources are automatically assigned to the user resource threshold profiles. When you do not create monitoring threshold profiles or define assignment rules, the newly discovered resources are automatically registered to the default threshold profiles.

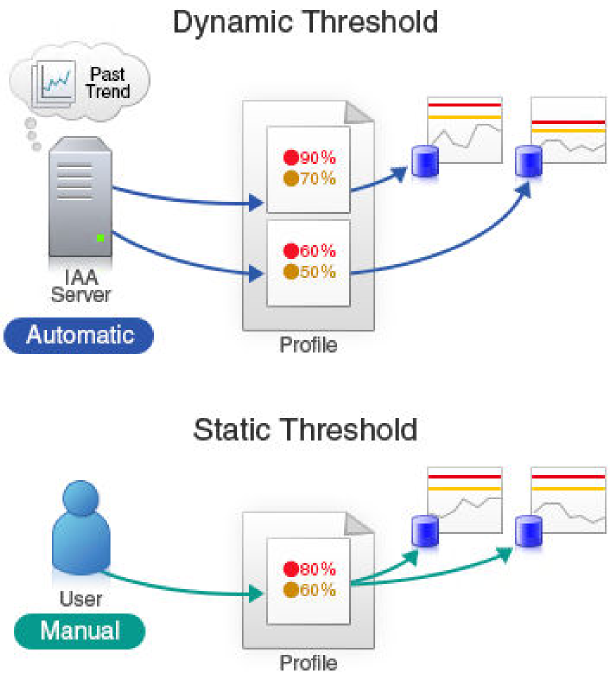

Ops Center Analyzer notifications

Setting notifications is an easy way to stay informed on the status of infrastructure resources and events. Monitoring resources is both an active and passive activity for IT administrators. You can use Ops Center Analyzer to configure email notifications that provide detailed information about issues with resource management. If multiple administrators rely on the Ops Center Analyzer service, you can create different profiles to deliver different types of information based on the profile settings. In the following example profile, a notification alert provides critical, warning, and information alerts based on performance and event action categories.

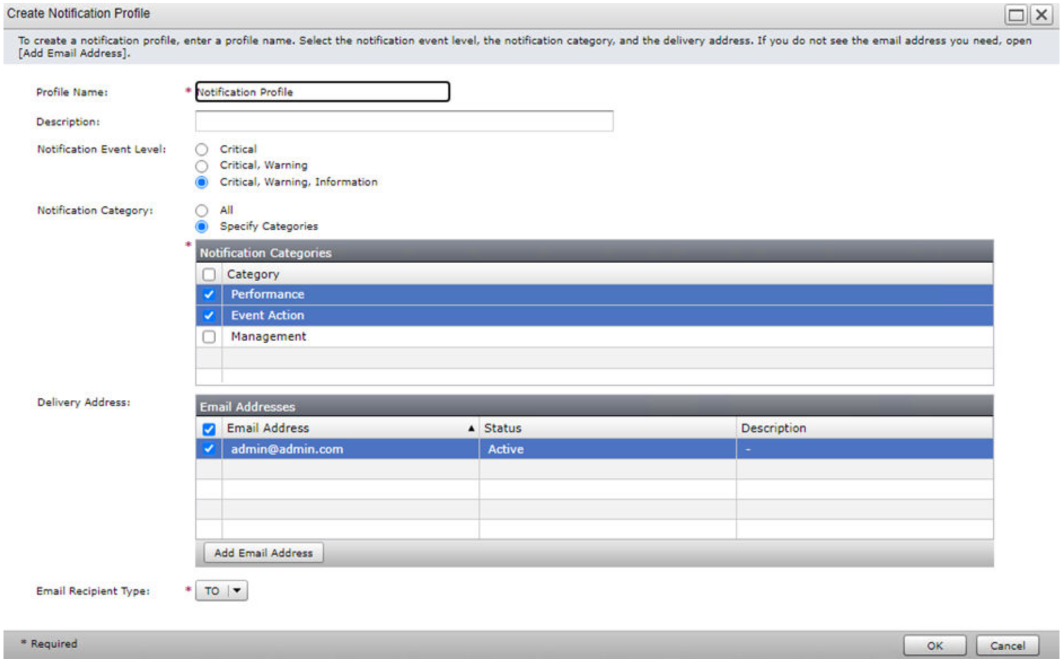

Create Notification Profile

To create a notification profile, enter a profile name. Select the notification event level, the notification category, and the delivery address. If you do not see the email address you need, open [Add Email Address].

Profile Name: * Notification Profile

Description:

Notification Event Level:
○ Critical
○ Critical, Warning
● Critical, Warning, Information

Notification Category:
○ All
● Specify Categories

* Notification Categories

□	Category
☑	Performance
☑	Event Action
□	Management

Delivery Address:

Email Addresses

☑	Email Address ▲	Status	Description
☑	admin@admin.com	Active	-

Add Email Address

Email Recipient Type: * TO ▾

* Required

OK Cancel

Hitachi Ops Center Automator integration

Ops Center Analyzer supports integration with Ops Center Automator. This support provides direct access to the service templates in Ops Center Automator from the Execute Action window in the Ops Center Analyzer UI. When administrators notice a performance problem in a shared infrastructure, they can run the appropriate action or service template to resolve it. This allows administrators to have auto-remediation capabilities in the event of a degradation of performance or capacity.

View the following video for more information on Ops Center Analyzer integration with Ops Center Automator.

Hitachi Ops Center detail view server

Ops Center Analyzer detail view server provides historical report analysis across the entire datacenter infrastructure, enabling creating advanced monitoring reports, and performing additional troubleshooting and diagnostics. Unlike the Analyzer server, the detail view server provides additional performance outside of the 14-day window. The following section describes the capabilities of the Analyzer detail view server for troubleshooting performance congestion. See the Appendix (on page 40) for a list of Cisco SAN performance metrics that can be collected from detail view server.

> **Note:** You can access Analyzer detail view reports any time directly from the E2E view by clicking any resource icon and selecting Show Report in Analyzer detail view.

Reports

After logging in to Ops Center Detail View Server, administrators are presented with the resource tree where resources can be selected to view the latest available default reports in the Performance view. From the resource tree, selections can be made based on Hypervisor, SAN switch, or VSP storage system.

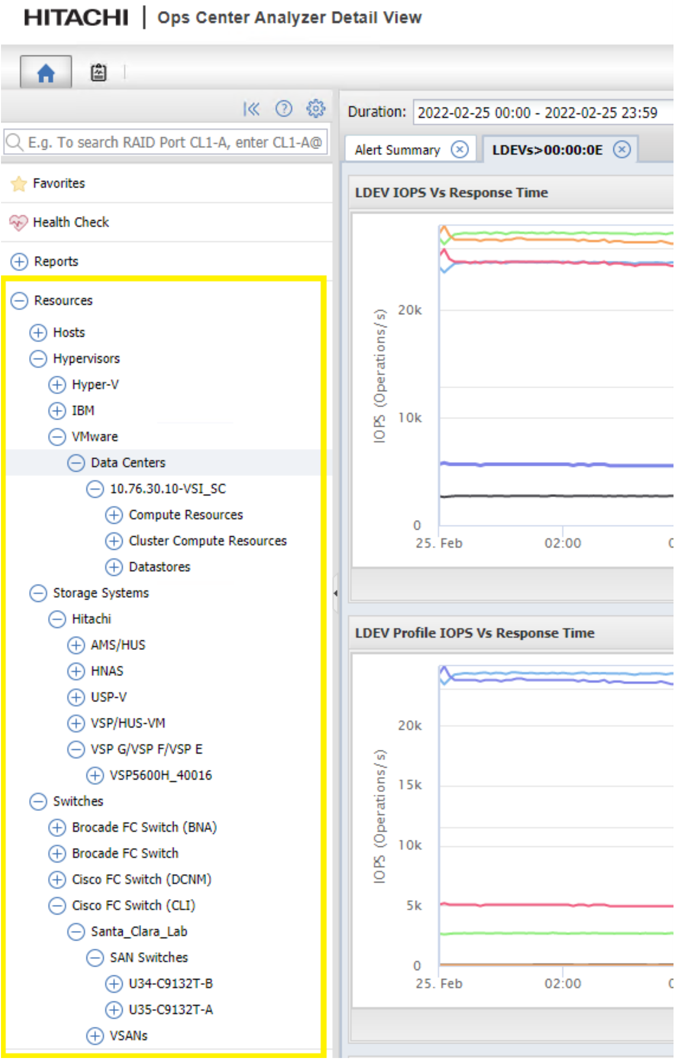

After a resource has been selected, administrators can extract reports in multiple formats including PNG, JPEG, PDF, SVG, and CSV from the 3 ellipses icon. Along with extracting reports, administrators can do performance comparisons using the Compare With function built into every report. The following example shows NVMe LDEV 0E IOPs performance monitored over two custom time durations. You can conclude that over this date range, resources that rely on this LDEV have been impacted by a configuration change or performance issue that is causing higher than normal LDEV response times.

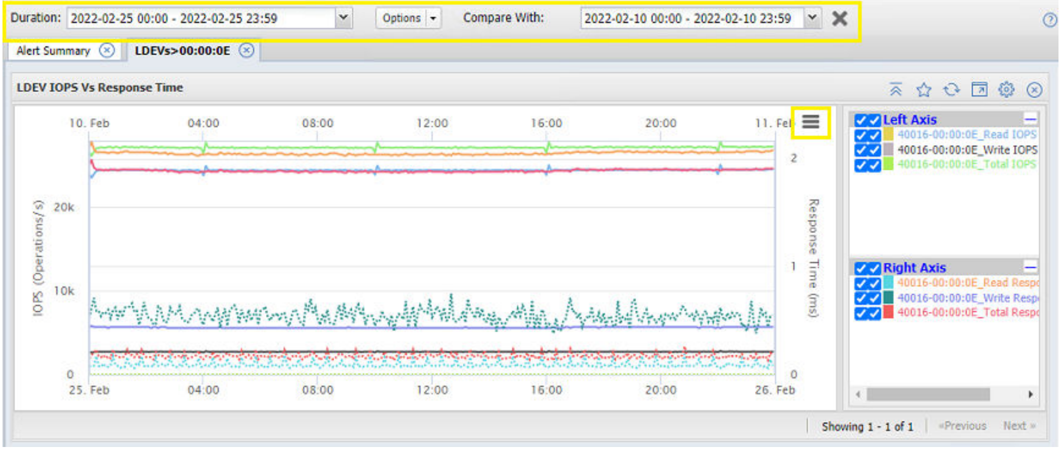
Duration: 2022-02-25 00:00 - 2022-02-25 23:59
Compare With: 2022-02-10 00:00 - 2022-02-10 23:59

LDEV IOPS Vs Response Time

Left Axis
- ☑ 40016-00:00:0E_Read IOPS
- ☑ 40016-00:00:0E_Write IOPS
- ☑ 40016-00:00:0E_Total IOPS

Right Axis
- ☑ 40016-00:00:0E_Read Respo
- ☑ 40016-00:00:0E_Write Resp
- ☑ 40016-00:00:0E_Total Respo

Showing 1 - 1 of 1

Custom reports

From the Ops Center detail view server, administrators can create their own custom reports that might not originally be available from a selected resource. In this example, a custom report is created for the fabric B SAN switch that is under congestion conditions. From the detailed view server, select the report name, as well as X and Y axis values, which in this case is Cisco switch port RX B2B Credit Transitions to 0 and Port Channel Rx Utilization.

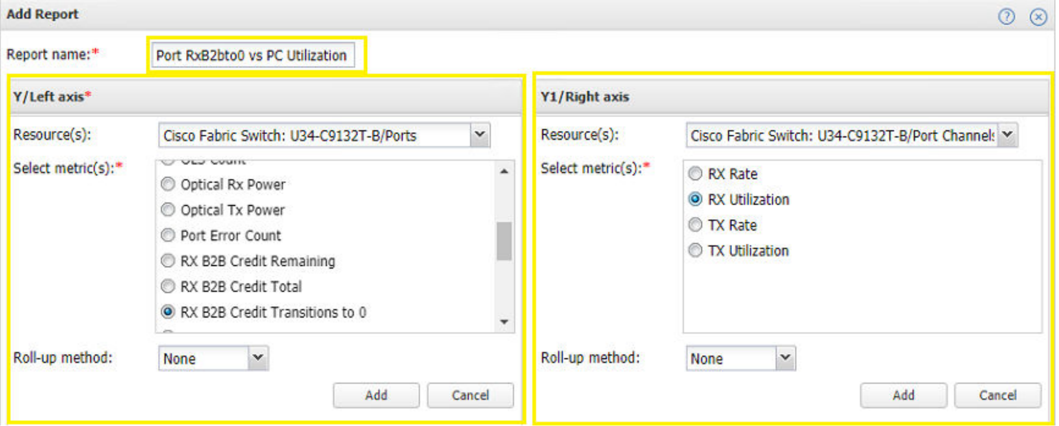
Add Report

Report name:* Port RxB2bto0 vs PC Utilization

Y/Left axis*

Resource(s): Cisco Fabric Switch: U34-C9132T-B/Ports

Select metric(s):*
- ○ Optical Rx Power
- ○ Optical Tx Power
- ○ Port Error Count
- ○ RX B2B Credit Remaining
- ○ RX B2B Credit Total
- ● RX B2B Credit Transitions to 0

Roll-up method: None

Y1/Right axis

Resource(s): Cisco Fabric Switch: U34-C9132T-B/Port Channel:

Select metric(s):*
- ○ RX Rate
- ● RX Utilization
- ○ TX Rate
- ○ TX Utilization

Roll-up method: None

Viewing the custom report output, administrators can narrow down ports that are experiencing congestion conditions. Here we can see Fibre Channel ports fc1/19 and fc1/20 are showing extremely high RXB2Bto0 transitions. We can also understand that port-channel 12, which uses fc1/1 that supports host connectivity of these storage ports to our hosts, is showing extremely high utilization narrowing the source of congestion.

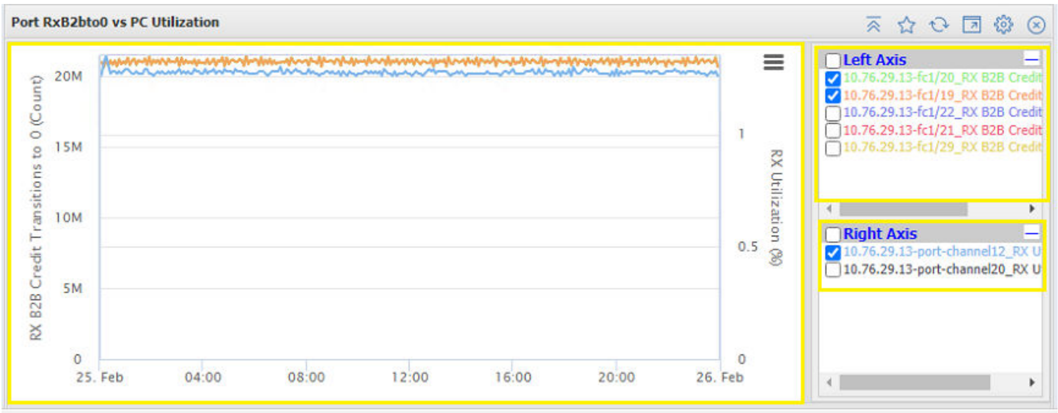
Port RxB2bto0 vs PC Utilization
RX B2B Credit Transitions to 0 (Count)
20M
15M
10M
5M
0
25. Feb
04:00
08:00
12:00
16:00
20:00
26. Feb
1
0.5
0
RX Utilization (%)
□ Left Axis
☑ 10.76.29.13-fc1/20_RX B2B Credit
☑ 10.76.29.13-fc1/19_RX B2B Credit
□ 10.76.29.13-fc1/22_RX B2B Credit
□ 10.76.29.13-fc1/21_RX B2B Credit
□ 10.76.29.13-fc1/29_RX B2B Credit
□ Right Axis
☑ 10.76.29.13-port-channel12_RX U
□ 10.76.29.13-port-channel20_RX U

Custom MQL reports for Cisco UCS

The Ops Center Analyzer detail view query language is a regex-based query language used to retrieve and filter data stored in the Analyzer detail view database. MQL allows complex analysis on this data in real-time with constant runtime. MQL syntax makes it possible to traverse relations, identify patterns in data, and provides a mechanism to establish correlations. Administrators can use the following queries in Ops Center Analyzer detail view server for Cisco UCS-based deployments.

Cisco Port and Hitachi storage port performance:

```
fabCiscoSwitchPort/*switchPort/*sp/raidPort[@totalIOPS rx b .*]
```

In this example, we can correlate VSP storage port performance for utilized Cisco host ports. We filtered based on Fabric B, and we can see that fc1/19 using VSP 5600H port CL3-B is showing high IOP utilization. This query can help narrow down performance issues from the VSP system perspective.

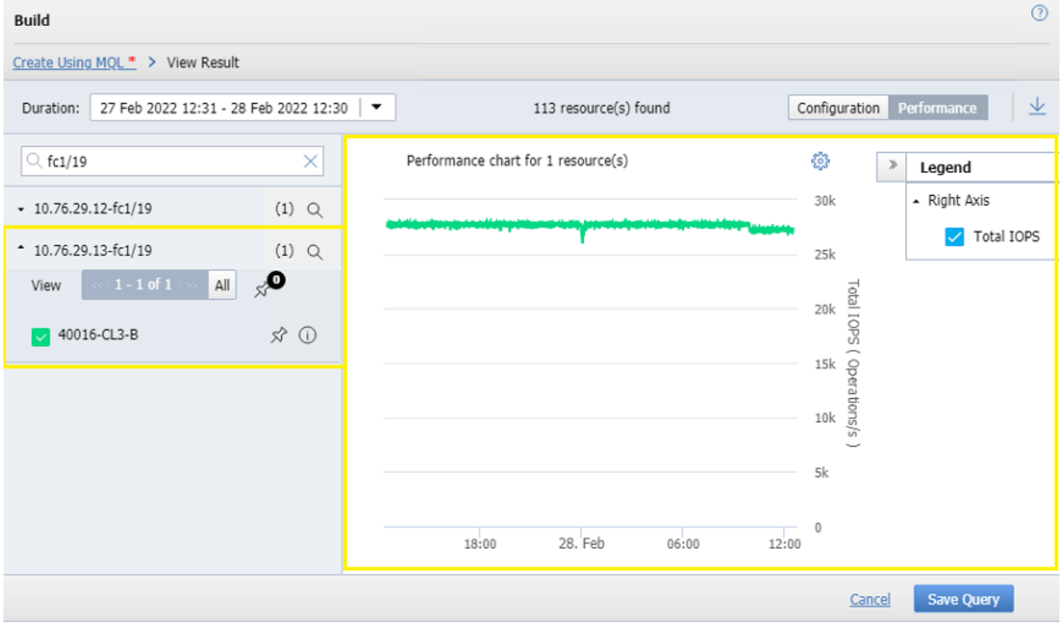

Build

Create Using MQL * > View Result

Duration: 27 Feb 2022 12:31 - 28 Feb 2022 12:30

113 resource(s) found

fc1/19

10.76.29.12-fc1/19 (1)

10.76.29.13-fc1/19 (1)

View 1 - 1 of 1 All

☑ 40016-CL3-B

Performance chart for 1 resource(s)

Legend

Right Axis

☑ Total IOPS

Cisco port channel and VMware host:

```
fabCiscoPortChannel/*portChannel/*hbaPort/*hba/*vmhba/h[@diskWrite rx b .*]
```

This example correlates VMware host performance per host vHBA to the Cisco Port Channel supporting the hosts. This query can be used to narrow down performance congestion on a Cisco UCS environment running VMware.

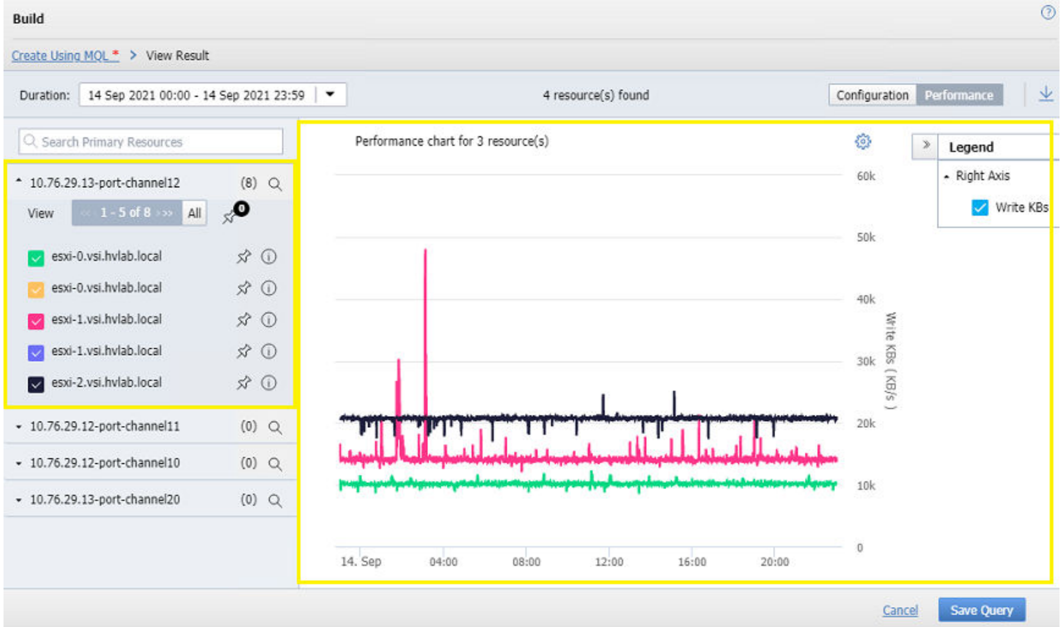

Build

Create Using MQL * > View Result

Duration: 14 Sep 2021 00:00 - 14 Sep 2021 23:59

4 resource(s) found

Search Primary Resources

10.76.29.13-port-channel12 (8)

View 1 - 5 of 8 All

☑ esxi-0.vsi.hvlab.local

☑ esxi-0.vsi.hvlab.local

☑ esxi-1.vsi.hvlab.local

☑ esxi-1.vsi.hvlab.local

☑ esxi-2.vsi.hvlab.local

10.76.29.12-port-channel11 (0)

10.76.29.12-port-channel10 (0)

10.76.29.13-port-channel20 (0)

Performance chart for 3 resource(s)

Legend

Right Axis

☑ Write KBs

Alerting

Ops Center Analyzer detail view classifies metrics into two types: configuration and performance. Configuration metrics change infrequently, and include resource status, capacity, while performance metrics include IOPS and response time.

Administrators can set various alert conditions for configuration and performance metrics. When these conditions are met, a notification can be sent by email, SNMP, or a Syslog server that aids in prevention and identification of performance degradation. In the following example, a custom alert for Cisco MDS 9132T, which supports fabric B under congestion conditions, specifies an alert to be shown when RxB2Bto0 credits on selected switch ports exceeds a limit of 2000.

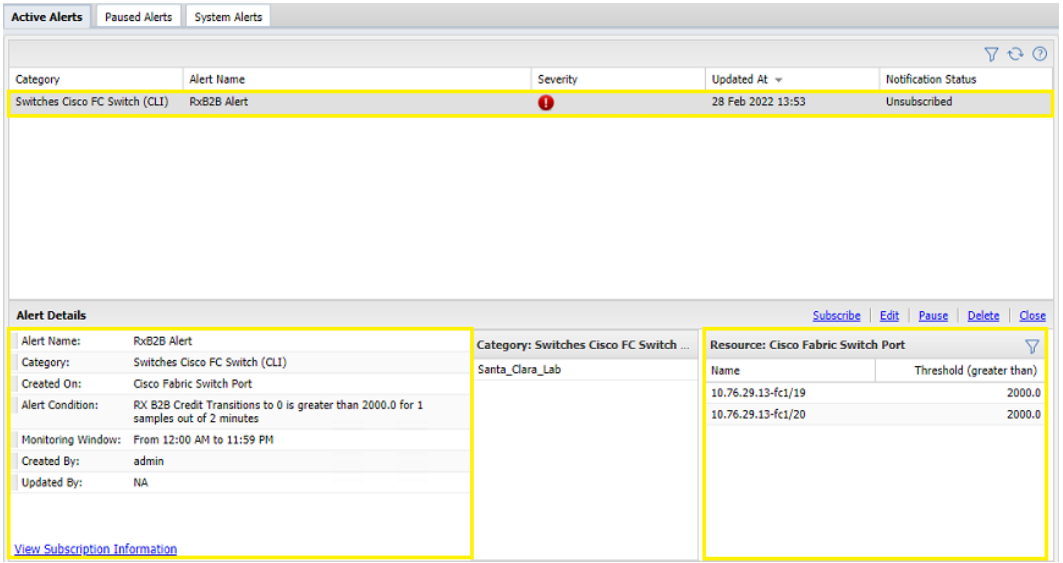

Category	Alert Name	Severity	Updated At	Notification Status
Switches Cisco FC Switch (CLI)	RxB2B Alert	!	28 Feb 2022 13:53	Unsubscribed

Alert Details

Alert Name:	RxB2B Alert
Category:	Switches Cisco FC Switch (CLI)
Created On:	Cisco Fabric Switch Port
Alert Condition:	RX B2B Credit Transitions to 0 is greater than 2000.0 for 1 samples out of 2 minutes
Monitoring Window:	From 12:00 AM to 11:59 PM
Created By:	admin
Updated By:	NA

Category: Switches Cisco FC Switch ...

Santa_Clara_Lab

Resource: Cisco Fabric Switch Port

Name	Threshold (greater than)
10.76.29.13-fc1/19	2000.0
10.76.29.13-fc1/20	2000.0

Appendix

This appendix lists the metrics that the Hitachi Ops Center Analyzer detail view server can collect from Cisco SAN switches onboarded using the CLI as the data collection method.

Attribute Name	Aggregation Operation	Unit	Data Granularity
CPU Utilization	Sum of user CPU and kernalCpu	%	5 Minutes
Memory Utilization	Divide memTotal by memUsed and then convert to %	%	5 Minutes
User CPU Utilization	Direct	%	5 Minutes
Kernal CPU Utilization	Direct	%	5 Minutes
Idle CPU Utilization	Direct	%	5 Minutes
Memory Total In GB	Convert KB to GB	GB/Sec	5 Minutes
Memory Used In GB	Convert KB to GB	GB/Sec	5 Minutes
TX Rate	Divide txBytes by 1024	KB/Sec	5 Minutes
RX Rate	Divide rxBytes by 1024	KB/Sec	5 Minutes
TX Utilization	Divide speed by txBytes and then convert to GB and then to %	%	5 Minutes
RX Utilization	Divide speed by rxBytes and then convert to GB and then to %	%	5 Minutes
Discarded Frames Count	Sum of rxDiscardCount and txDiscardCount	Number	5 Minutes
Invalid CRC Count	Direct	Number	5 Minutes
NOS Count	Sum of outputNOS and inputNOS	Number	5 Minutes
OLS Count	Sum of outputOLS and inputOLS	Number	5 Minutes
LRR Count	Sum of outputLRR and inputLRR	Number	5 Minutes
Port Error Count	Sum of rxErrorCount and txErrorCount	Number	5 Minutes

Attribute Name	Aggregation Operation	Unit	Data Granularity
Temperature	Direct	Centigrade	5 Minutes
Voltage	Direct	V	5 Minutes
Current	Direct	mA	5 Minutes
Optical Tx Power	Direct	dBm	5 Minutes
Optical Rx Power	Direct	dBm	5 Minutes
RX B2B Credit Transitions to 0	DELTA between previous and current data points	Count	5 Minutes
TX B2B Credit Transitions to 0	DELTA between previous and current data points	Count	5 Minutes
RX B2B Credit Remaining	Direct	Count	5 Minutes
TX B2B Credit Remaining	Direct	Count	5 Minutes
RX B2B Credit Total	Direct	Count	5 Minutes
TX B2B Credit Total	Direct	Count	5 Minutes
TX Rate	Divide txBytes by 1024	KB/Sec	5 Minutes
RX Rate	Divide rxBytes by 1024	KB/Sec	5 Minutes
TX Utilization	Divide speed by txBytes and then convert to GB and then to %	%	5 Minutes
RX Utilization	Divide speed by rxBytes and then convert to GB and then to %	%	5 Minutes

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact