



Washington Systems Center - Storage

Accelerate with IBM Storage: Cisco / IBM c-type SAN Analytics

Art Scrimo
Product Manager Cisco
January 2020



Accelerate with IBM Storage Webinars

The Free IBM Storage Technical Webinar Series Continues in 2020...

Washington Systems Center – Storage experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

Located on the Accelerate with IBM Storage Site: <https://www.ibm.com/support/pages/node/1125513>

Also, check out the WSC YouTube Channel here:
https://www.youtube.com/channel/UCNuks0go01_ZrVVF1jgOD6Q



2020 Upcoming Webinars:

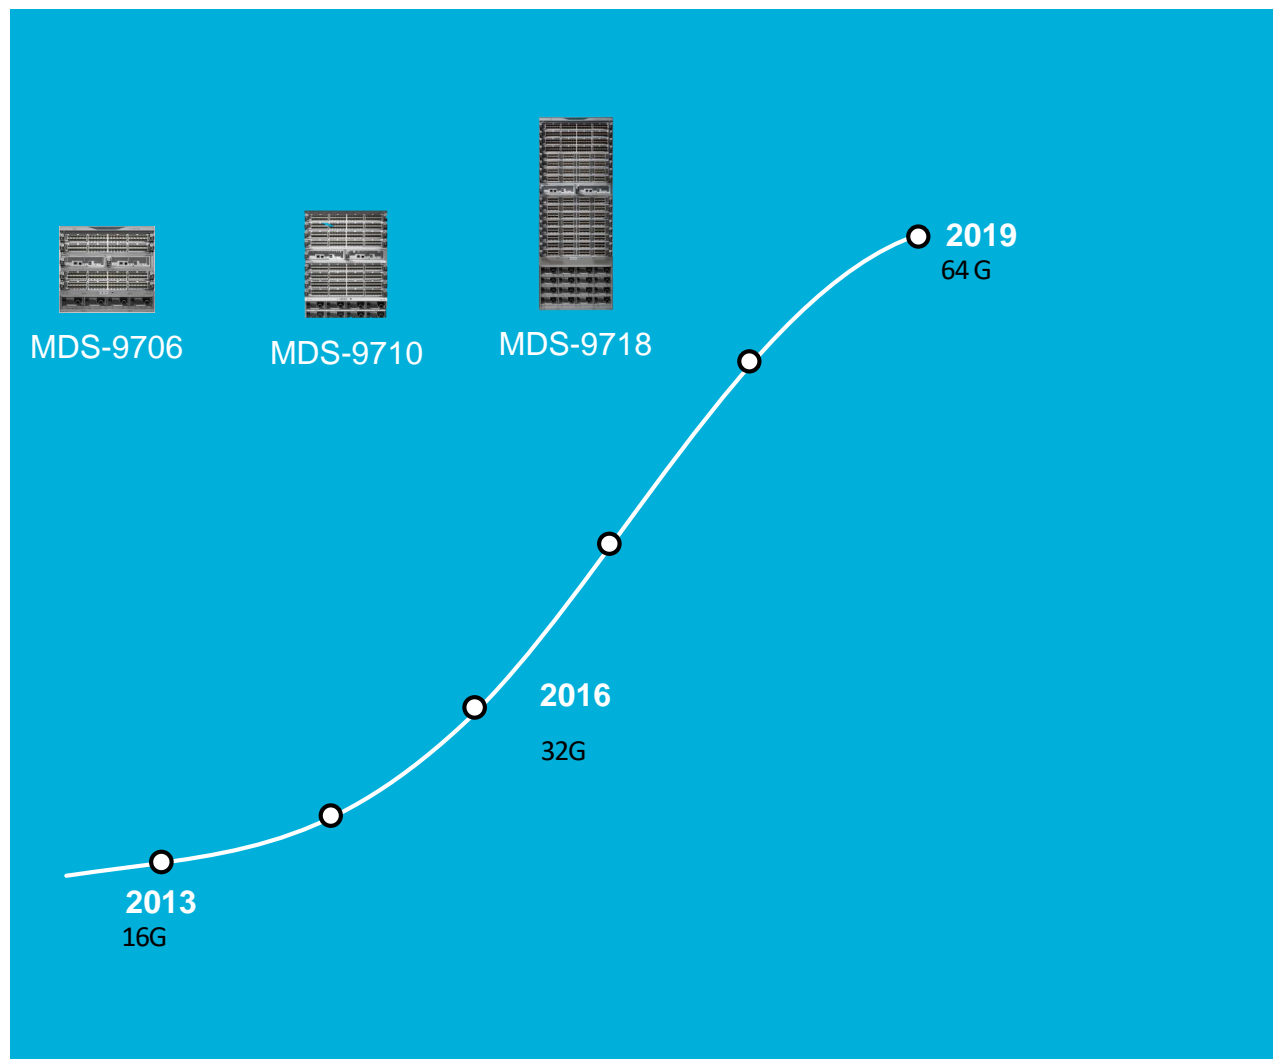
January 14 – Brocade / IBM b-type SAN Modernization

Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=ed703072fdd739b7e1384ed84869aed8d>

January 21 - Cisco / IBM c-type SAN Analytics

Register Here: <https://ibm.webex.com/ibm/onstage/g.php?MTID=eabc0e050b2ff8bdcea4e012c30dfbf71>

Built to last



Investment protection for NVMe/FC and all flash arrays

- 64G ready director
- No forklift upgrade
- Built for the most demanding storage environments



Gain actionable insights

- Industries first NVMe/FC analytics
- FC-SCSI and NVMe/FC support
- Built for customer choice and flexibility



Reduce operational complexity

- Extending devops support for IT automation: ANSIBLE
- Reduce OPEX; simple integration
- Built for advanced SAN automation

Built to last

The good...

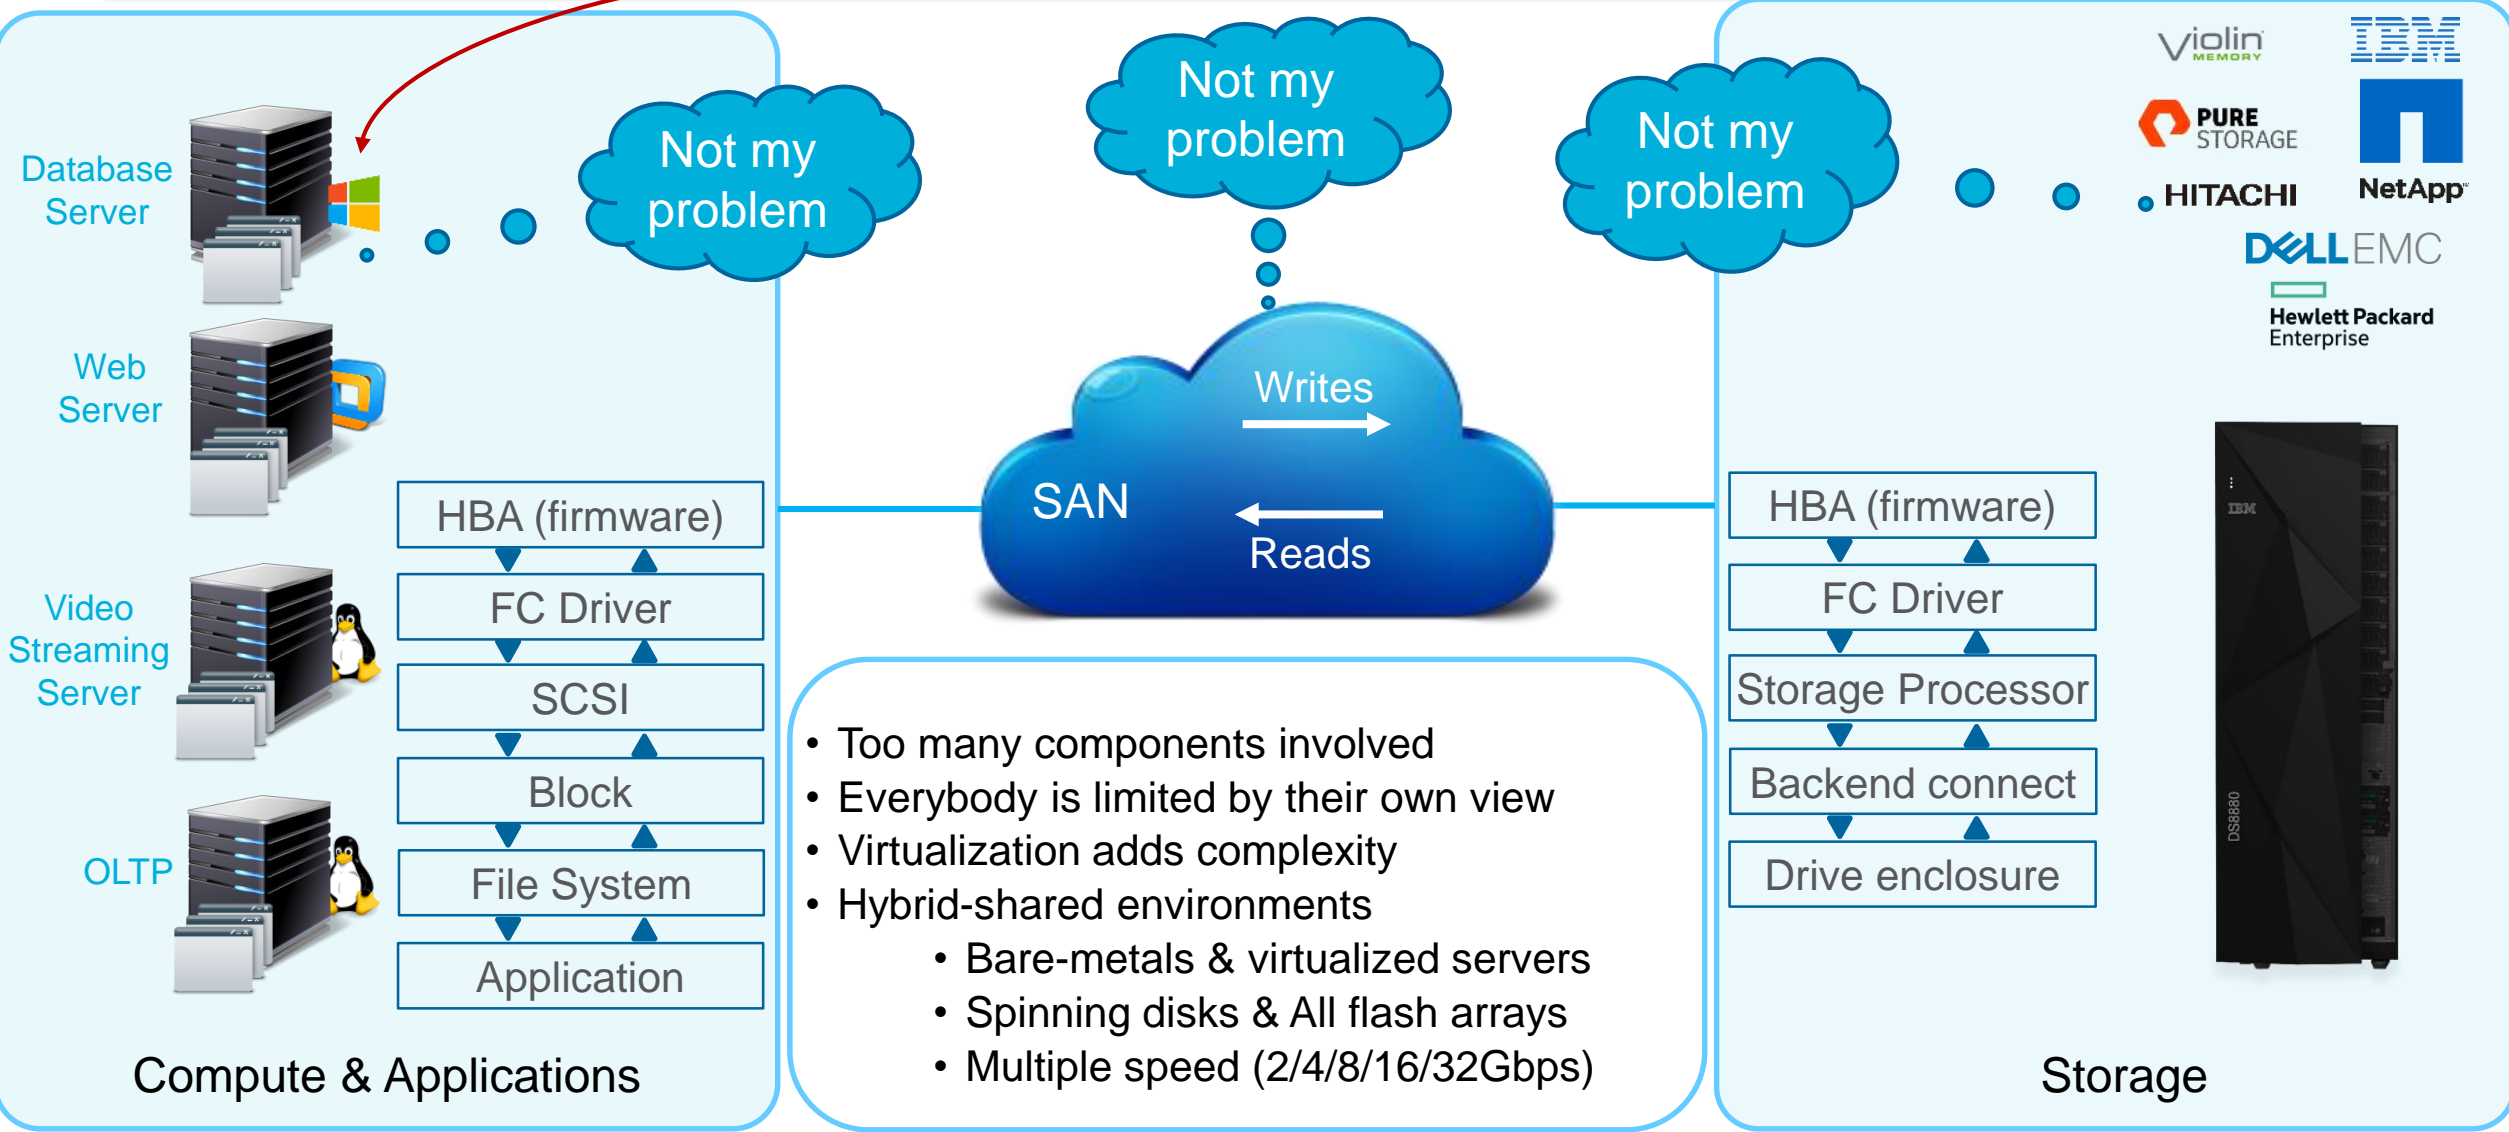
- All flash arrays are great for performance, smaller footprint, less power, higher density
- NVMe is great for performance and will have an impact on future deployments
 - Pods and application specific deployments to start

and the bad...

- Are your hosts ready?
 - Typically we see slow drain caused by slower host attached devices
- Not many NVMe supported devices today
 - Can you plan around the future deployments today?

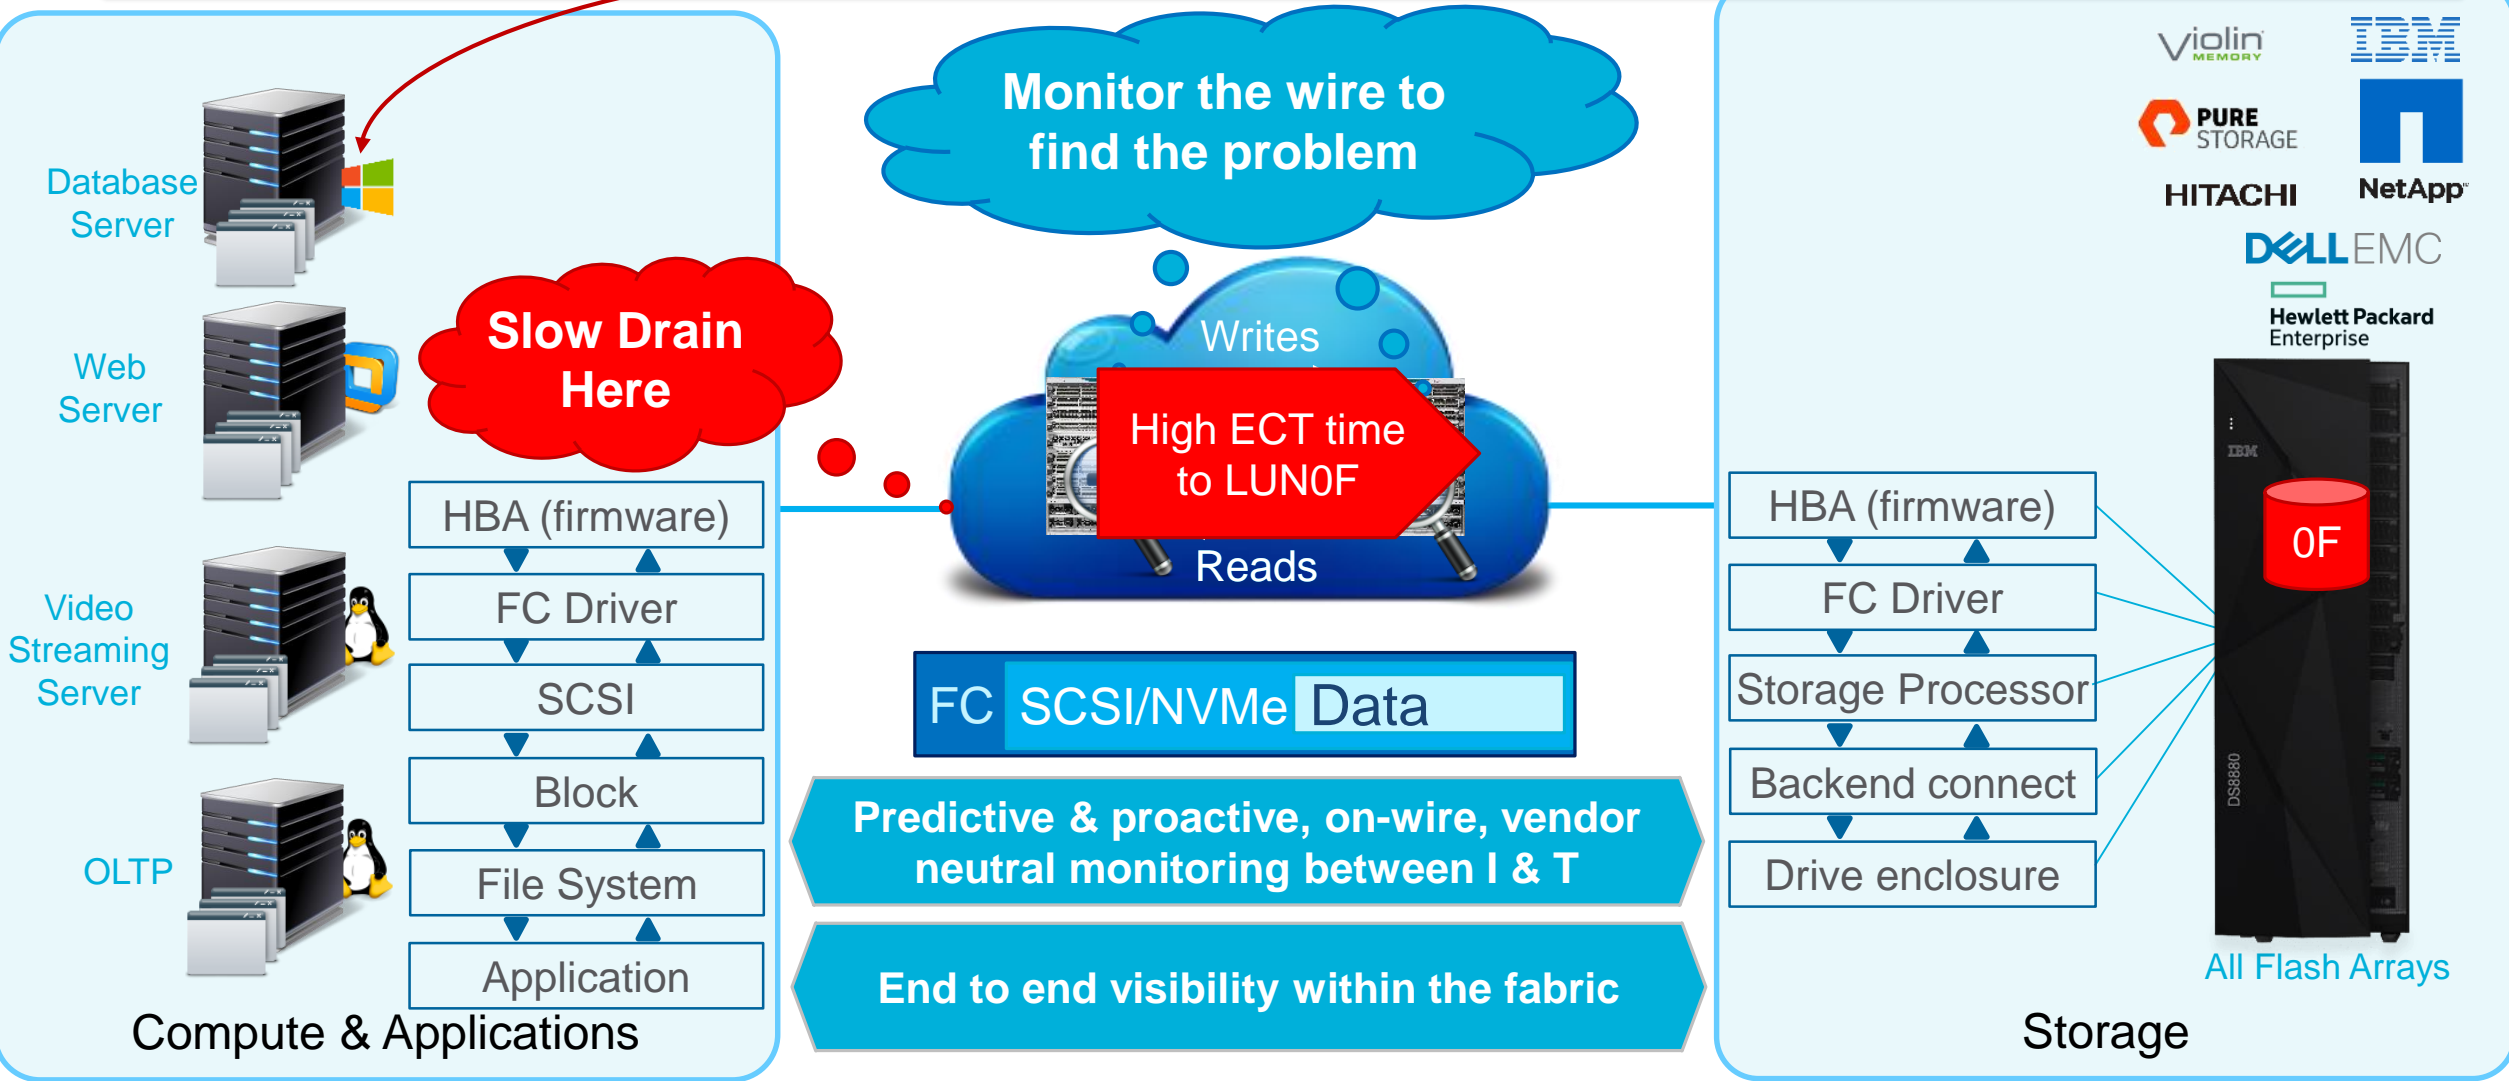
The current state of storage I/O problems

Application issues identified



The current state of storage I/O problems


Application issues identified



IBM 32GB Portfolio with SAN Telemetry Capabilities

Cisco MDS 9700

(IBM SAN768C-6,
SAN384C-6, SAN912C-6)



MDS 9700 32GB Line Card



MDS 9132T 32GB 1U Switch

(IBM SAN32C-6)




MDS 9148T 32GB 1U Switch

(IBM SAN48C-6)



MDS 9396T 32GB 2U Switch

(IBM SAN96C-6)



- Hardware Integrated Telemetry
- End-to-end metrics
- * These features are available for all Cisco 32Gbps products, not available on the 16Gbps products



Cisco Telemetry - Architecture

Data Collection

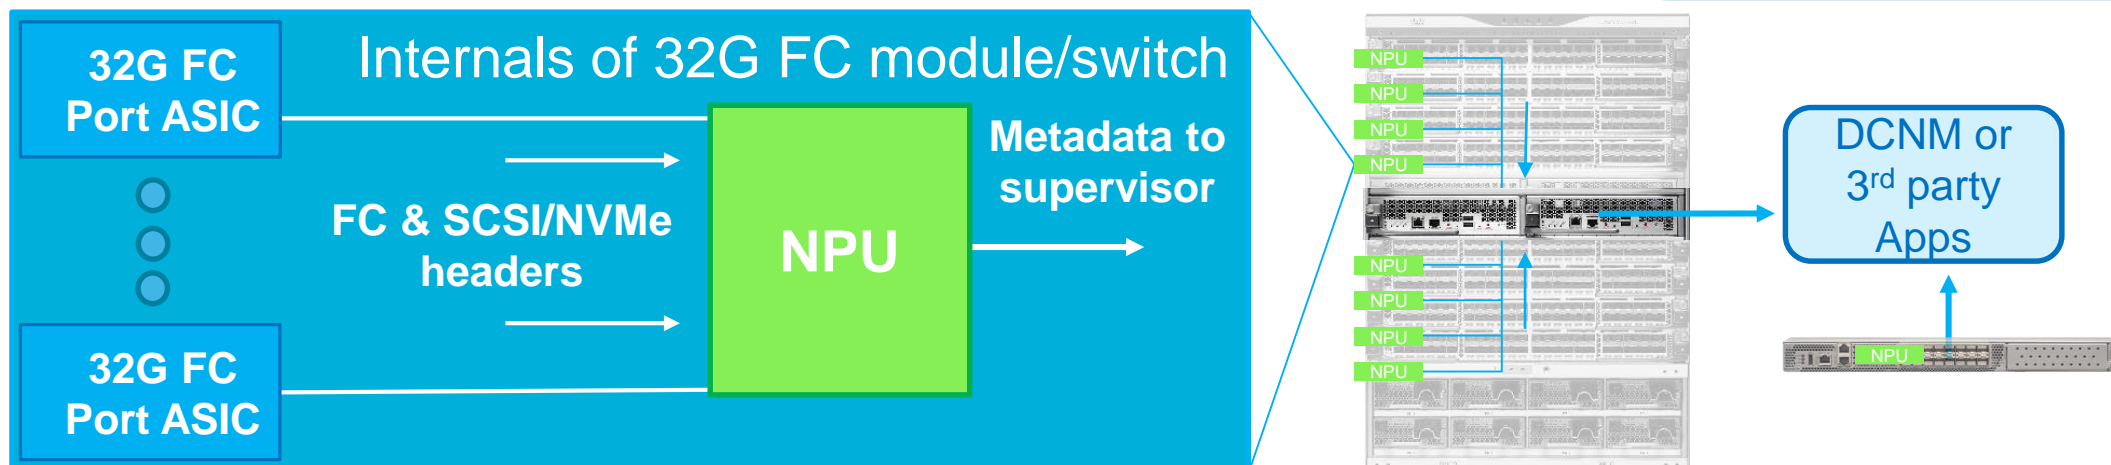
- Inbuilt data collection capability into 32G FC port-ASIC
- No impact to data traffic
- Collects only FC & SCSI/NVMe headers, not data

Data Processing

- On-board Network Processing Unit (NPU) on 32G FC module
- Receives headers of frames from port-ASIC
- Extracts metadata from headers
- Stores meta-data in multiple views


Data Analytics and Visualization

- DCNM receives meta-data processed by NPU
- Stores data for long duration
- Analyzes & provides end-to-end visibility, trending, variations, etc
- Correlation with host & target data




60+ metrics per flow

Write CMD



Read CMD




- Exchange Completion Time (ECT)
- Data Access Latency (DAL)
- Outstanding IO (Queue Depth)
- IOPS, Detailed flow level counters
- Per frame timeout drop frames
- Failed FC exchanges, IO retransmissions
- SCSI/NVMe Error conditions (Aborts, Rejects, ect)
- IO block size and other flow level stats

All metrics at granularity of SID – DID – LUN

32GB Line Card Deployment models


Location of 32G FC module

Storage Edge




- Closest to storage
- 32G FC module in data path for all flows
- Best for existing MDS 9700 deployment

Core Deployment



- High capacity core with 32G FC speed
- 32G FC module in data path for all flows
- Best for existing MDS 9700 deployment

Full Fabric




- Best model for end-to-end visibility and correlation
- 32G FC module in data path for all flows
- Best for upgrading from MDS 9500

32G FC module must be in frame switching path


Fixed Switched Deployment models

16Gbps Fixed Switch Host Edge



- Closest to storage
- 32G FC module in data path for all flows
- Best for existing MDS deployment

32Gbps Fixed Switch Host Edge



- Closest to the host edge
- 32G FC switches in data path for all flows
- Best for new edge refresh

32G FC product must be in frame switching path

Metric visualization flexibility

On-Switch CLI

- Flexible SQL-like CLIs to pull raw data from 32G FC module
- Output in key-value format (similar to dictionaries or JSON)
- Unmatched flexibility for advanced user
 - Basic trending & correlation Optional standard CLIs
- Output in nice-to-read tabular format, just like any other NX-OS show command output

Remote RESTful access

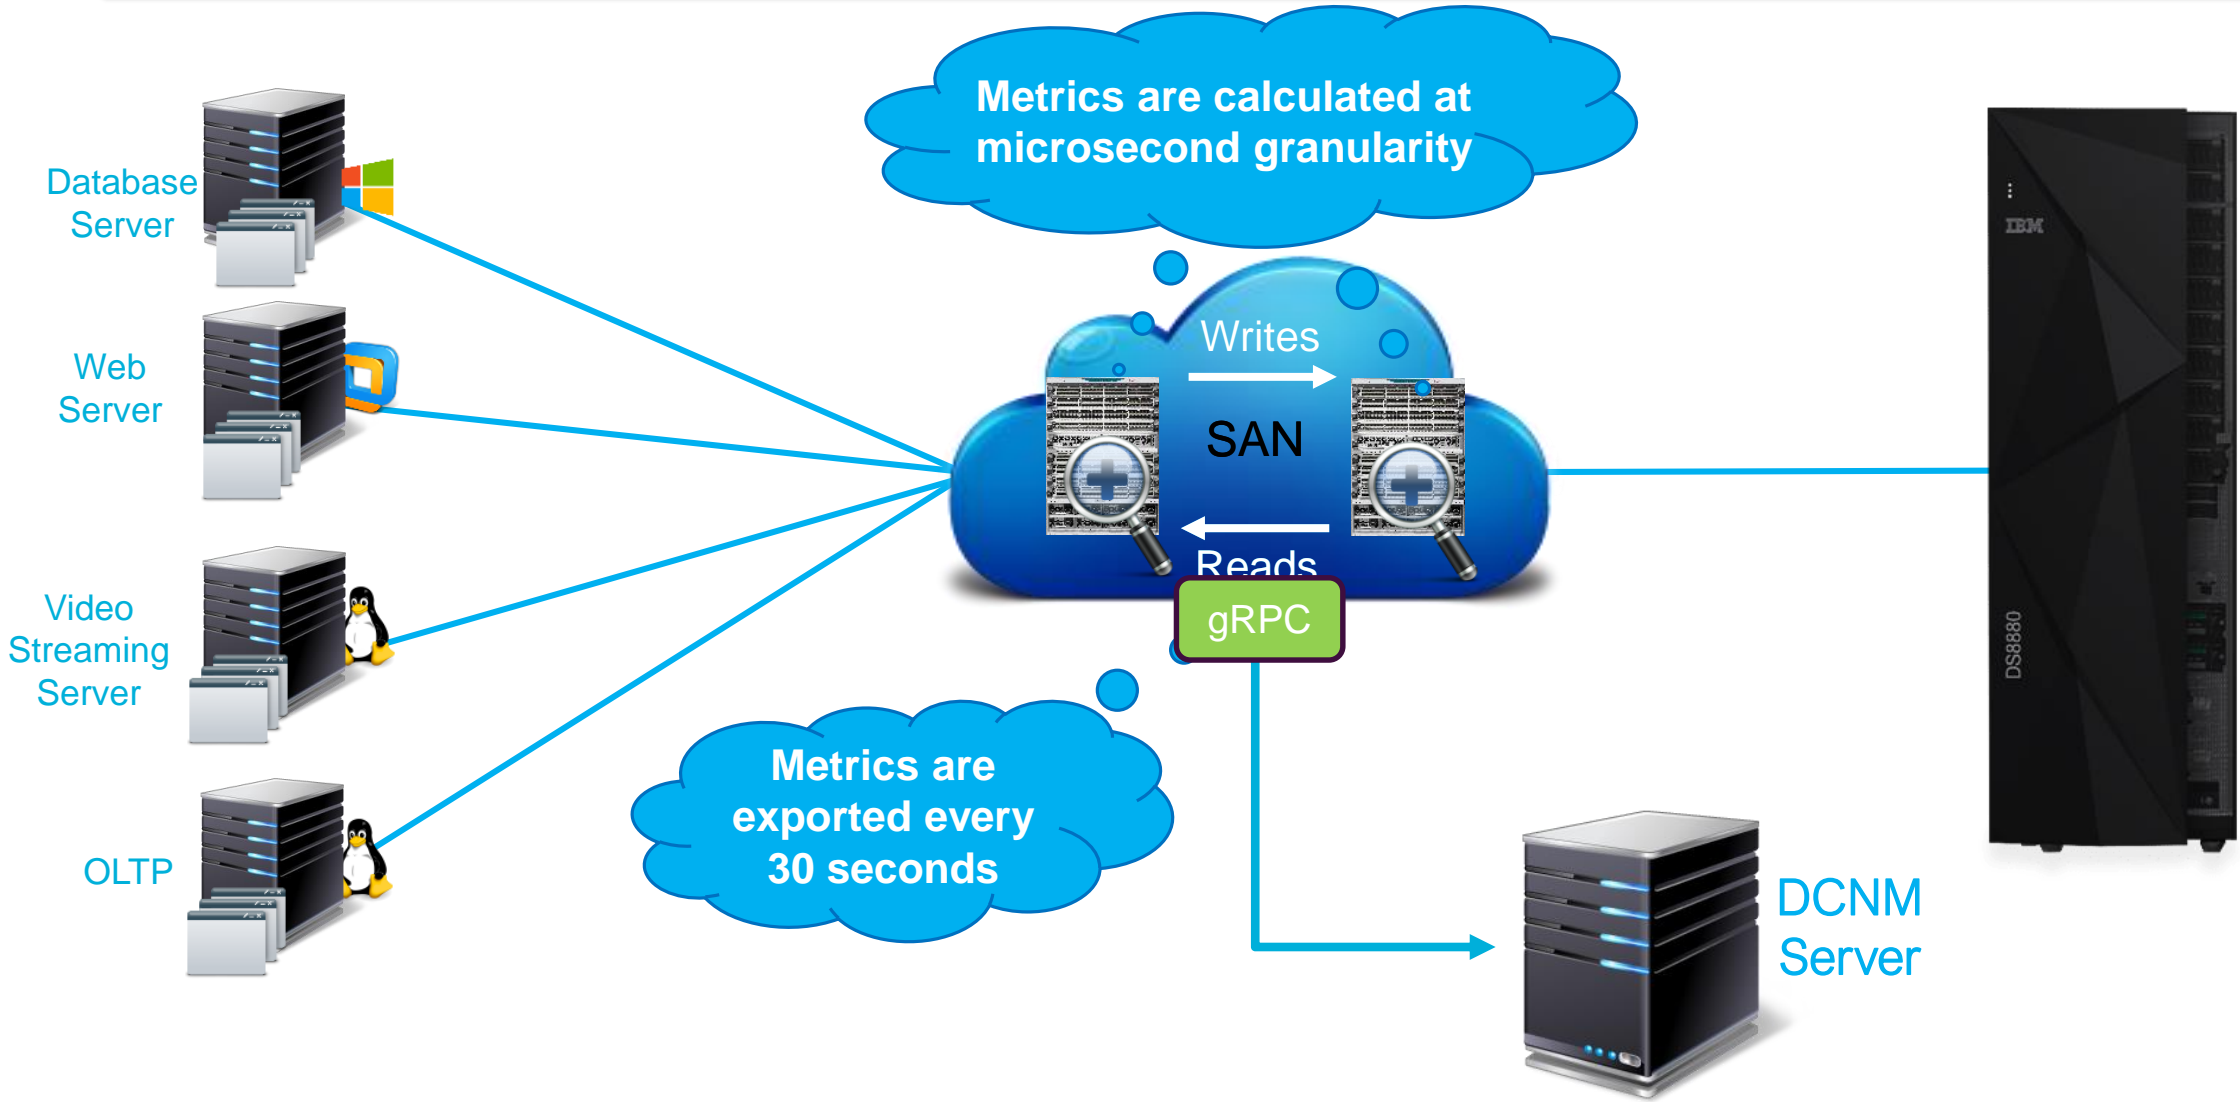
- Capability to query switch remotely
- Response in JSON format
- Response output depends on the SQL-like input query under the request
- Same as the existing NX-API model
- Works best to track specific metrics
- Extremely flexible

Streaming Telemetry

- MDS streams out the metrics to remote apps regularly
- Data transport and encoding in industry leading open format
- Optimized performance for continuous data (instead of Pull mode) export at high frequency.

DCNM and 3rd Party Apps

SAN Analytics data delivery



DCNM 11.3(1)

Features – SAN Insights – Use
cases

Show me which hosts are talking to which storage connections (Different from zoning)

The screenshot shows the Cisco Data Center Network Manager interface. The 'Monitor / SAN / SAN Insights' section is active, with 'Host Enclosures' selected. A timeline at the top shows data from 72 hours ago to the present. The 'Host Enclosures' table on the left lists 19 hosts, with 'UCSB_RHEL_17439' selected. The 'Initiator Target Pairs' table on the right shows connections between the selected host and various storage targets. A network diagram below the table visualizes these connections. At the bottom, performance metrics for a specific flow are displayed.

Host Enclosure	Status
172.25.174.138	Critical
UCSB_RHEL_17439	Warning
Win_17416	Healthy
HOST_Emulex_476e41	Healthy
RHEL_17480	Healthy
WIN_174186	Healthy
RHEL17438	Healthy
SUSE	Healthy
172.25.174.33	Healthy
172.25.174.36	Healthy
WIN_174182	Healthy
172.25.174.135	Healthy
ESX_174184	Healthy
172.25.174.152	Healthy
RHEL_17414	Healthy
RHEL_17417	Healthy
UCSB_WIN_17440	Healthy
WIN_174233	Healthy
172.25.174.137	Healthy

Source PWWN	SID	Destination PWWN	DID	Fabric	Status
20:00:06:25:b5:00:00:2f	700002	50:05:07:60:5e:bf:e9:62	5801a0	Fabric_N5596UP-17486	Critical
20:00:06:25:b5:00:00:3f	700001	50:05:07:60:5e:bf:e9:62	5801a0	Fabric_N5596UP-17486	Warning
20:00:06:25:b5:00:00:4f	ef0042	50:05:07:60:5e:bf:e9:41	d80240	Fabric_N5596UP-17486	Warning
20:00:06:25:b5:00:00:5f	ef0041	50:05:07:60:5e:bf:e9:41	d80240	Fabric_N5596UP-17486	Warning

Name	1-Hour Average	Baseline
Normalized Read ECT	0.1218 ms/KB	0.1232 ms/KB
Normalized Write ECT	0.0000 ms/KB	0.0000 ms/KB
Normalized Read DAL	0.1211 ms/KB	
Normalized Write DAL	0.0000 ms/KB	
Read Active I/O	0.00	
Write Active I/O	0.00	

Name	Value
In Errors	
Out Errors	
In Discards	
Out Discards	
Tx	
Rx	



- Select Host Enclosure on the left side
- See all connections to storage from each host
- Check health of each active traffic flow across all fabrics

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Show me which storage devices are talking to which hosts (Different from zoning)

The screenshot shows the Cisco Data Center Network Manager interface. The 'Storage Enclosures' tab is selected, and a list of storage enclosures is shown on the left. The 'Initiator Target Pairs' table is displayed in the center, showing connections between storage enclosures and hosts. A blue arrow points to the 'Status' column of the Initiator Target Pairs table, highlighting a red status icon. Below the table, there is a network diagram and performance metrics for a specific flow.

Storage Enclosures

Storage Enclosure	Status
<input checked="" type="radio"/> IBM_F840	Warning
<input type="radio"/> ONTAP:1e8b0c2b-f019-11e4-a817-00a0985e37e2	Healthy
<input type="radio"/> NETAPP_8060	Healthy
<input type="radio"/> HDS_93040500_HUS150	Healthy
<input type="radio"/> HDS_210083_HUSVM	Healthy
<input type="radio"/> HDS_410081_G600	Healthy
<input type="radio"/> StorageSystem:VSP.53100	Healthy
<input type="radio"/> IBM_XIV	Healthy

Initiator Target Pairs

Source PWWN	SID	Destination PWWN	DID	Fabric	Status
<input checked="" type="radio"/> 10:00:00:90:fa:73:b0:7d	d80160	50:05:07:60:5e:bf:e9:41	d80240	Fabric_N5596UP-17486	Critical
<input type="radio"/> 20:00:06:25:b5:00:00:2e	700004	50:05:07:60:5e:bf:e9:62	5801a0	Fabric_N5596UP-17486	Healthy
<input type="radio"/> 20:00:06:25:b5:00:00:2f	700002	50:05:07:60:5e:bf:e9:62	5801a0	Fabric_N5596UP-17486	Warning
<input type="radio"/> 20:00:06:25:b5:00:00:3e	700003	50:05:07:60:5e:bf:e9:62	5801a0	Fabric_N5596UP-17486	Healthy

Flow: 10:00:00:90:fa:73:b0:7d to 50:05:07:60:5e:bf:e9:41

Name	1-Hour Average	Baseline
Normalized Read ECT	0.7710 ms/KB	0.6775 ms/KB
Normalized Write ECT	0.0000 ms/KB	0.0000 ms/KB
Normalized Read DAL	0.7700 ms/KB	
Normalized Write DAL	0.0000 ms/KB	
Read Active I/O	0.00	
Write Active I/O	0.00	

Interface

Name	Value
In Errors	
Out Errors	
In Discards	
Out Discards	
Tx	
Rx	



- Select Storage Enclosure on the left side
- See all connections to Hosts from each Storage Device
- Check health of each active traffic flow across all fabrics

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Show me the poorest performing flows in my environment

Flows Selected 1 / Total 12

Initiator - Target	Status
10:00:00:00:c9:ef:42:f3 <-> 50:05:07:60:5e:bf:e9:41	Critical
20:00:78:da:6e:a0:01:e0 <-> 50:05:07:60:5e:bf:e9:...	Critical
10:00:00:90:fa:47:6e:55 <-> 50:06:0e:80:10:1a:fa:44	Healthy
10:00:00:90:fa:47:6e:55 <-> 50:06:0e:80:12:27:61:...	Healthy
10:00:00:90:fa:47:6e:55 <-> 50:06:0e:80:13:27:63:...	Healthy
10:00:00:90:fa:47:6e:55 <-> 50:01:73:80:64:5d:01:...	Healthy
10:00:00:90:fa:47:6e:55 <-> 50:05:07:60:5e:bf:e9:41	Healthy
10:00:00:90:fa:02:4c:61 <-> 50:06:0e:80:10:1a:fa:44	Healthy
10:00:00:90:fa:02:4c:61 <-> 50:06:0e:80:12:27:61:13	Healthy
10:00:00:90:fa:02:4c:61 <-> 50:06:0e:80:13:27:63:00	Healthy
21:00:00:0e:1e:18:51 <-> 50:05:07:60:5e:bf:e9:...	Healthy
10:00:00:90:fa:73:b0:7d <-> 50:05:07:60:5e:bf:e9:41	Healthy

Initiator Target Pairs Selected 1 / Total 1

Source PWWN	SID	Destination PWWN	DID	Fabric	Status
10:00:00:00:c9:ef:42:f3	7f00ac	50:05:07:60:5e:bf:e9:41	d80240	Fabric_N5596UP-17486	Critical

Flow: 10:00:00:00:c9:ef:42:f3 to 50:05:07:60:5e:bf:e9:41

Name	1-Hour Average	Baseline
Normalized Read ECT	7.3291 ms/KB	0.5810 ms/KB
Normalized Write ECT	0.0000 ms/KB	0.0000 ms/KB
Normalized Read DAL	7.3286 ms/KB	
Normalized Write DAL	0.0000 ms/KB	
Read Active I/O	0.00	
Write Active I/O	0.00	

Interface

Name	Value
In Errors	
Out Errors	
In Discards	
Out Discards	
Tx	
Rx	



- Select Flows up at the top
- Select the Slowest Flow from the left side (top most is worst)
- Check health of each active traffic flow across all fabrics

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Find trends and patterns using slider in Monitor/SAN Insights

Flows Selected 1 / Total 12

Initiator - Target	Status
10:00:00:00:c9:ef:42:f3 <-> 50:05:07:60:5e:bf:e9:41	Critical
20:00:78:da:6e:a0:01:e0 <-> 50:05:07:60:5e:bf:e9:...	Critical
10:00:00:90:fa:47:6e:55 <-> 50:06:0e:80:10:1a:fa:44	Healthy
10:00:00:90:fa:47:6e:55 <-> 50:06:0e:80:12:27:61:...	Healthy
10:00:00:90:fa:47:6e:55 <-> 50:06:0e:80:13:27:63:...	Healthy
10:00:00:90:fa:47:6e:55 <-> 50:01:73:80:64:5d:01:...	Healthy
10:00:00:90:fa:47:6e:55 <-> 50:05:07:60:5e:bf:e9:41	Healthy
10:00:00:90:fa:02:4c:61 <-> 50:06:0e:80:10:1a:fa:44	Healthy
10:00:00:90:fa:02:4c:61 <-> 50:06:0e:80:12:27:61:13	Healthy
10:00:00:90:fa:02:4c:61 <-> 50:06:0e:80:13:27:63:00	Healthy
21:00:00:0e:1e:18:51 <-> 50:05:07:60:5e:bf:e9:...	Healthy
10:00:00:90:fa:73:b0:7d <-> 50:05:07:60:5e:bf:e9:41	Healthy

Initiator Target Pairs Selected 1 / Total 1

Source PWWN	SID	Destination	Fabric	Status
10:00:00:00:c9:ef:42:f3	7f00ac	50:05:07:60:5e:bf:e9:41	d80240 Fabric_N5596UP-17486	Critical

Flow: 10:00:00:00:c9:ef:42:f3 to 50:05:07:60:5e:bf:e9:41

Name	1-Hour Average	Baseline
Normalized Read ECT	7.3291 ms/KB	0.5810 ms/KB
Normalized Write ECT	0.0000 ms/KB	0.0000 ms/KB
Normalized Read DAL	7.3286 ms/KB	
Normalized Write DAL	0.0000 ms/KB	
Read Active I/O	0.00	
Write Active I/O	0.00	

Interface

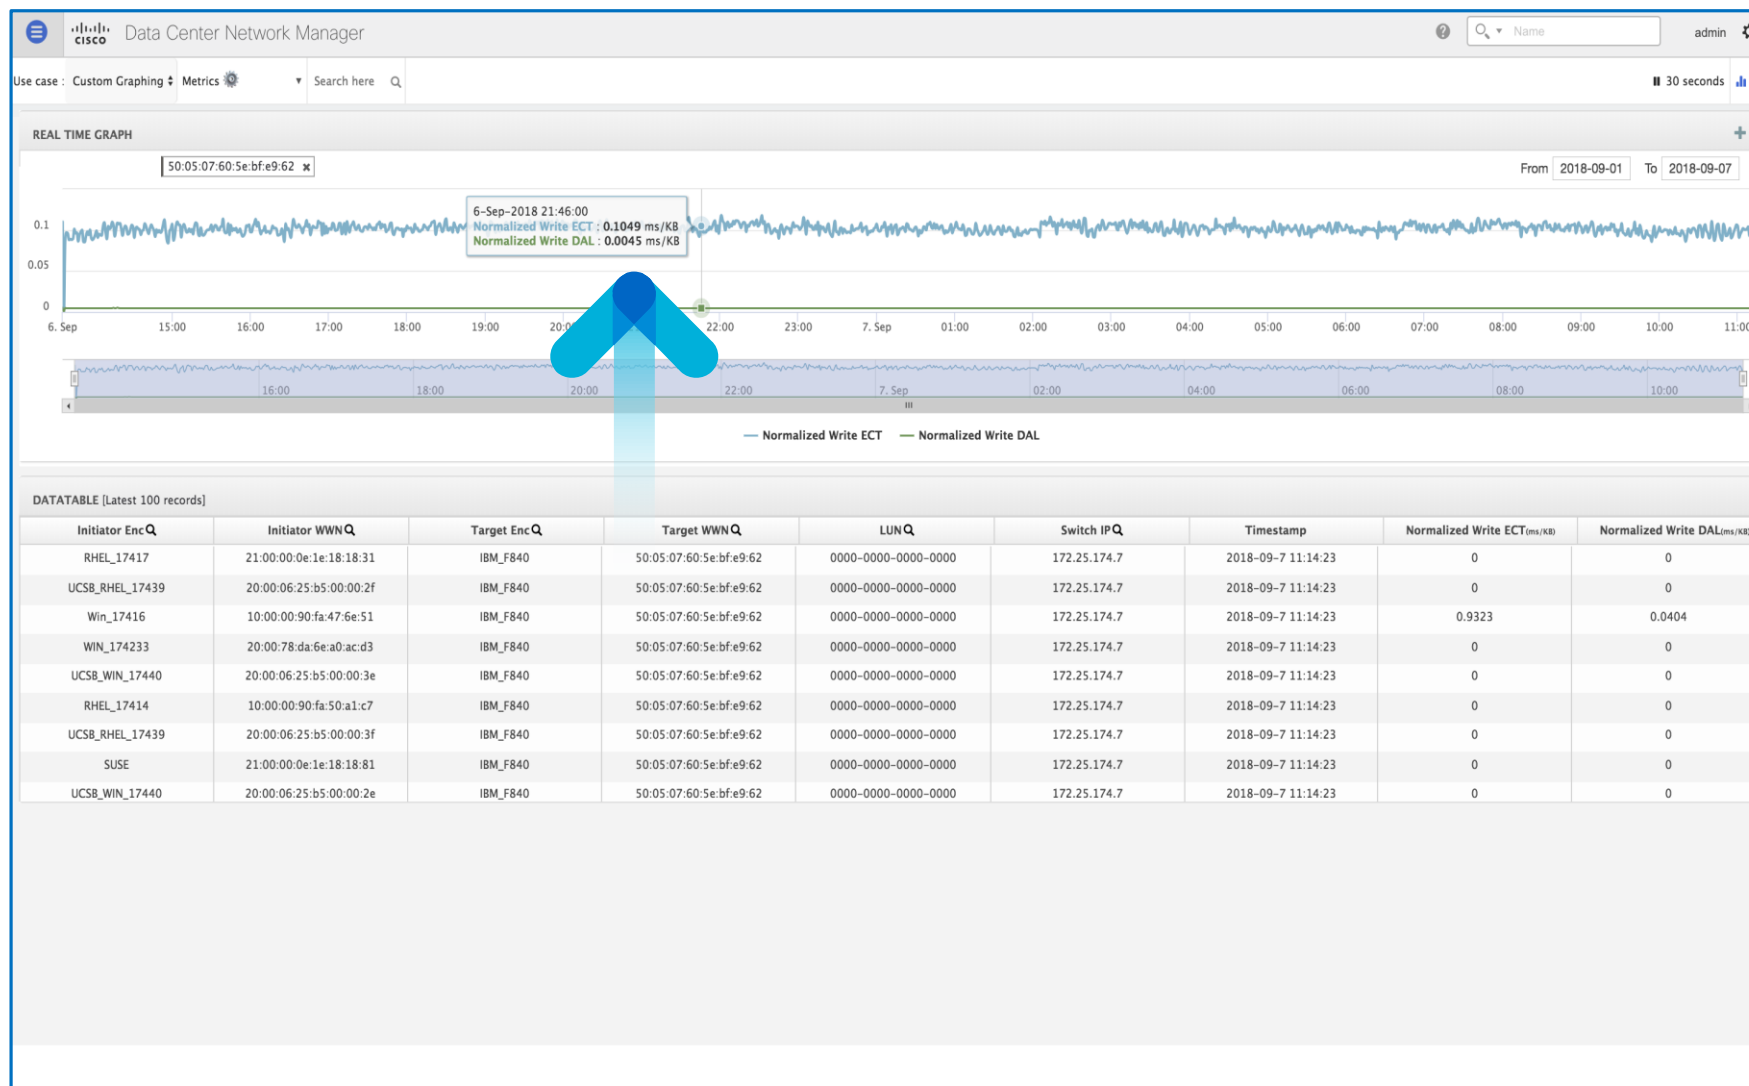
Name	Value
In Errors	
Out Errors	
In Discards	
Out Discards	
Tx	
Rx	



- User slider to go back in time and find patterns
- One hour increments for every dash
- Watch how things change over time with each click into the past
- Use custom graphing to dive down deeper to see trends

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

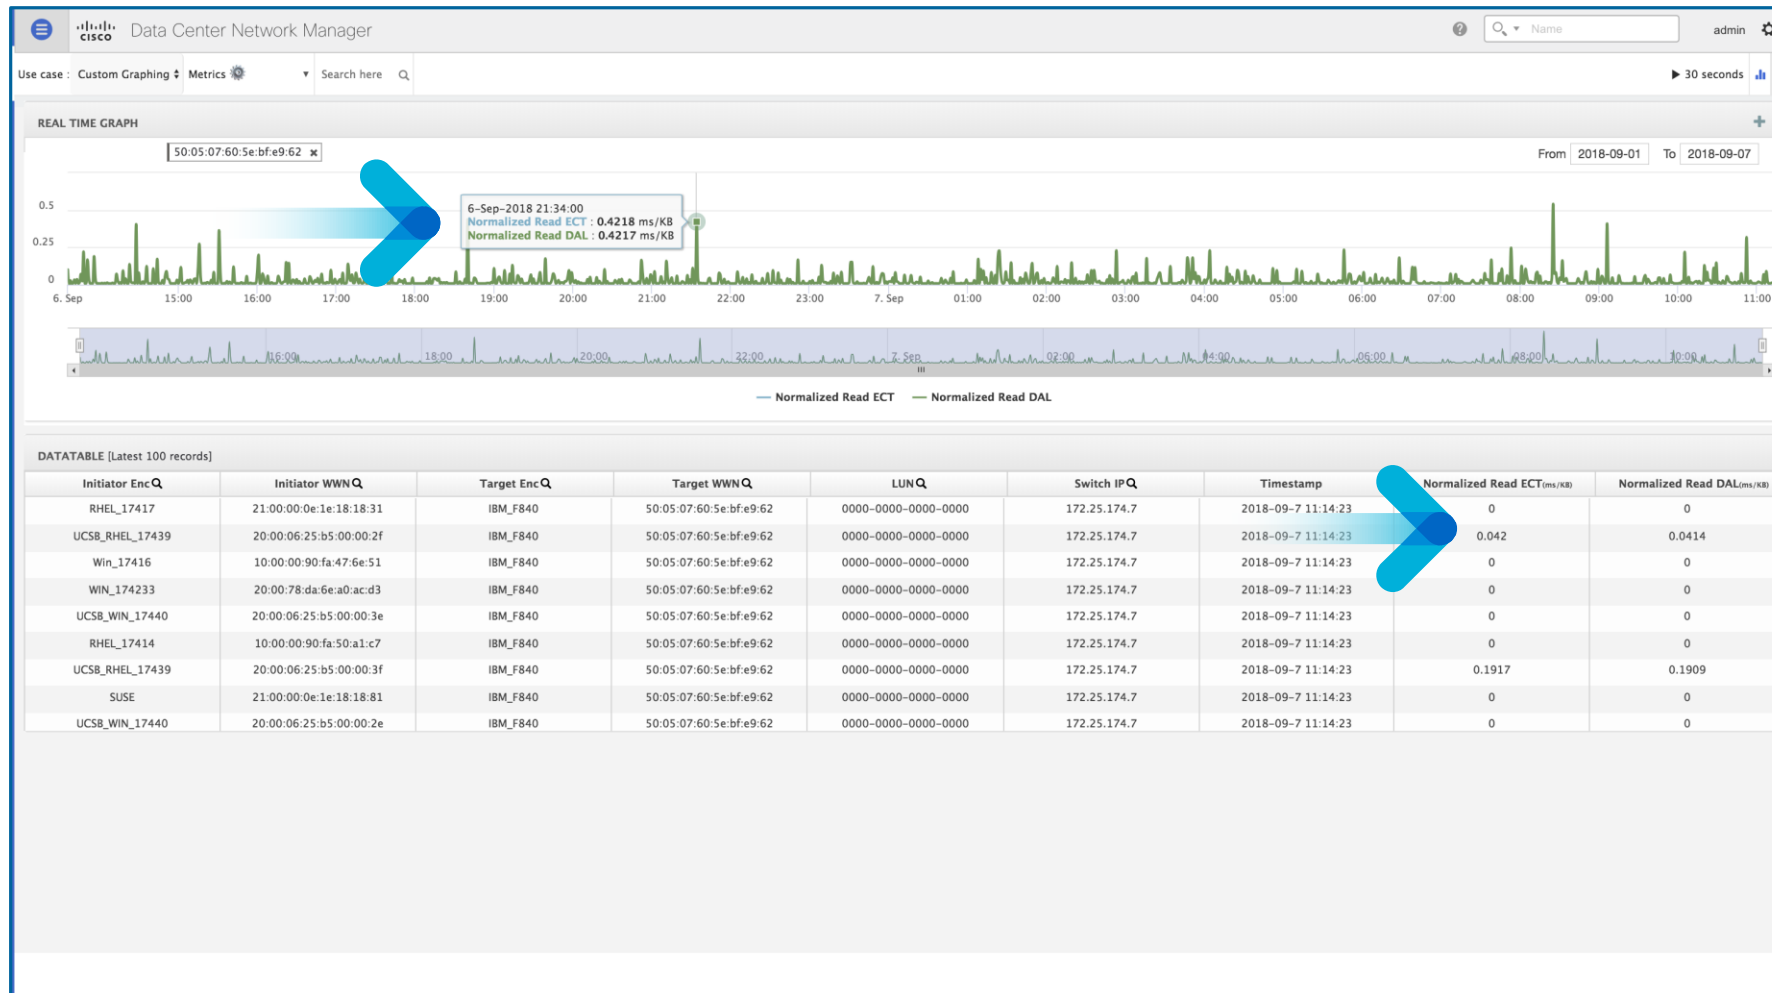
Is it a slow SAN or slow storage device?



- Use Custom Graphing to search by LUN ID, Host or Storage port
- Compare DAL and ECT
- DAL high storage Device causing slowness; ECT High SAN causing slowness
- **Based on this graph all issues are pointing to a SAN related slowness on this storage port**

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Is my Storage Device slow to respond to read or write requests?



- Use Custom Graphing
- Input WWPN or host name of storage device
- Select Read or Write DAL and ECT
- If they track on a similar trend than the storage device is impacting ECT
- **Based on this graph all issues are pointing to a storage device causing slow response times**

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Show me how busy my host is

The screenshot displays the Cisco Data Center Network Manager interface. At the top, it shows 'Monitor / SAN / SAN Insights' with tabs for Host Enclosures, Storage Enclosures, and Flows. A legend indicates activity levels: 0-1% (green), 1-15% (yellow), 15-30% (orange), and >30% (red). A time slider is set to 'now'.

Host Enclosures (Selected 1 / Total 4):

Host Enclosure	Status
Win_17416	Healthy
172.25.174.137	Healthy
172.25.174.138	Healthy
172.25.174.152	Healthy

Initiator Target Pairs (Selected 1 / Total 8):

Source PWWN	SID	Destination PWWN	DID	Fabric	Status
10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:10:1a:fa:44	d80040	Fabric_N5596UP-17486	Healthy
10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:12:27:61:13	d80220	Fabric_N5596UP-17486	Healthy
10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:13:27:63:00	d80200	Fabric_N5596UP-17486	Healthy

A topology map below shows the host (Win_17416) connected to various storage nodes (MDS9710-174141, MDS9710-17480, N5596UP-17486, N5596UP-17480) and other components like HDS_210083_HUB/IN, HDS_410081_0800, and IBM_SVC. A legend indicates Critical (red), Warning (yellow), Healthy (green), and Unknown (grey).

Flow: 10:00:00:90:fa:47:6e:55 to 50:06:0e:80:10:1a:fa:44

Name	1-Hour Average	Baseline
Read Active I/O	0.00	
Write Active I/O	0.14	
Read IO Aborts	0.00	
Write IO Aborts	0.00	
Read IO Failure	0.00	
Write IO Failure	0.00	

Interface: MDS9710-174141 fc7/9

Name	Value
In Errors	0
Out Errors	0
In Discards	0
Out Discards	0
Tx	1.6750 MB/s
Rx	10.6625 MB/s



- Use Host Enclosure view
- Find Host in list
- Select Interface in Topology Map
- Check each path under Initiator Target Pairs
- See both Flow based perf metrics and interface based metrics
- Errors, Aborts and other error related metrics provided

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Show me how busy my storage device is

☰ **csico** Data Center Network Manager ? admin ⚙️

🏠 Monitor / SAN / SAN Insights ○ Host Enclosures ● Storage Enclosures ○ Flows ● 0 - 1% ● 1 - 15% ● 15 - 30% ● > 30% Time Chosen now

72-hours ago 66 60 54 48-hours ago 42 36 30 24-hours ago 18 12 6 now

Storage Enclosures Selected 1 / Total 5

Storage Enclosure	Status
<input checked="" type="radio"/> HDS_210083_HUSVM	●
<input type="radio"/> HDS_410081_G600	●
<input type="radio"/> HDS_93040500_HUS150	●
<input type="radio"/> IBM_F840	●
<input type="radio"/> IBM_XIV	●


Initiator Target Pairs Selected 1 / Total 2

Source PWWN	SID	Destination PWWN	DID	Fabric	Status
<input checked="" type="radio"/> 10:00:00:90:fa:02:4c:61	d801c1	50:06:0e:80:13:27:63:00	d80200	Fabric_N5596UP-17486	●
<input type="radio"/> 10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:13:27:63:00	d80200	Fabric_N5596UP-17486	●

Flow: 10:00:00:90:fa:02:4c:61 to 50:06:0e:80:13:27:63:00 Double-click to collapse or restore the panel 747 fc17

Name	1-Hour Average	Baseline
<input checked="" type="checkbox"/> Read Active I/O	0.00	
<input checked="" type="checkbox"/> Write Active I/O	0.13	
<input checked="" type="checkbox"/> Read IO Aborts	0.00	
<input checked="" type="checkbox"/> Write IO Aborts	0.00	
<input checked="" type="checkbox"/> Read IO Failure	0.00	
<input checked="" type="checkbox"/> Write IO Failure	0.00	

Name	Value
<input checked="" type="checkbox"/> In Errors	0
<input checked="" type="checkbox"/> Out Errors	0
<input checked="" type="checkbox"/> In Discards	0
<input checked="" type="checkbox"/> Out Discards	0
<input checked="" type="checkbox"/> Tx	2.2750 B/s
<input checked="" type="checkbox"/> Rx	0.4500 B/s



- Use Storage Enclosure view
- Find Storage Device in list
- Select Interface in Topology Map
- Check each path under Initiator Target Pairs
- See both Flow based perf metrics and interface based metrics
- Errors, Aborts and other error related metrics provided

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Are there errors impacting my hosts?

The screenshot shows the Cisco Data Center Network Manager interface. The top navigation bar includes 'Data Center Network Manager' and a search bar. Below the navigation, there are tabs for 'Monitor / SAN / SAN Insights' and radio buttons for 'Host Enclosures', 'Storage Enclosures', and 'Flows'. A legend indicates error rates: 0-1% (green), 1-15% (yellow), 15-30% (orange), and >30% (red). A time slider is set to 'now'.

Host Enclosures (Selected 1 / Total 4):

Host Enclosure	Status
Win_17416	Healthy
172.25.174.137	Healthy
172.25.174.138	Healthy
172.25.174.152	Healthy

Initiator Target Pairs (Selected 1 / Total 8):


Source PWWN	SID	Destination PWWN	DID	Fabric	Status
10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:10:1a:fa:44	d80040	Fabric_N5596UP-17486	Healthy
10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:12:27:61:13	d80220	Fabric_N5596UP-17486	Healthy
10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:13:27:63:00	d80200	Fabric_N5596UP-17486	Healthy

Flow: 10:00:00:90:fa:47:6e:55 to 50:06:0e:80:10:1a:fa:44

Name	1-Hour Average	Baseline
Read Active I/O	0.00	
Write Active I/O	0.14	
Read IO Aborts	0.00	
Write IO Aborts	0.00	
Read IO Failure	0.00	
Write IO Failure	0.00	

Interface: MDS9710-174141 fc7/9

Name	Value
In Errors	0
Out Errors	0
In Discards	0
Out Discards	0
Tx	1.6750 MB/s
Rx	10.6625 MB/s



- Use Host Enclosure view
- Find Host in list
- Select Interface in Topology Map
- Check each path under Initiator Target Pairs
- Errors, Aborts and other error related metrics provided

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Are there errors impacting my storage device?

Storage Enclosures Selected 1 / Total 5

Storage Enclosure	Status
<input checked="" type="radio"/> HDS_210083_HUSVM	●
<input type="radio"/> HDS_410081_G600	●
<input type="radio"/> HDS_93040500_HUS150	●
<input type="radio"/> IBM_F840	●
<input type="radio"/> IBM_XIV	●

Initiator Target Pairs Selected 1 / Total 2

Source PWWN	SID	Destination PWWN	DID	Fabric	Status
<input checked="" type="radio"/> 10:00:00:90:fa:02:4c:61	d801c1	50:06:0e:80:13:27:63:00	d80200	Fabric_N5596UP-17486	●
<input type="radio"/> 10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:13:27:63:00	d80200	Fabric_N5596UP-17486	●

Flow: 10:00:00:90:fa:02:4c:61 to 50:06:0e:80:13:27:63:00

Name	1-Hour Average	Baseline
Read Active I/O	0.00	
Write Active I/O	0.13	
Read IO Aborts	0.00	
Write IO Aborts	0.00	
Read IO Failure	0.00	
Write IO Failure	0.00	

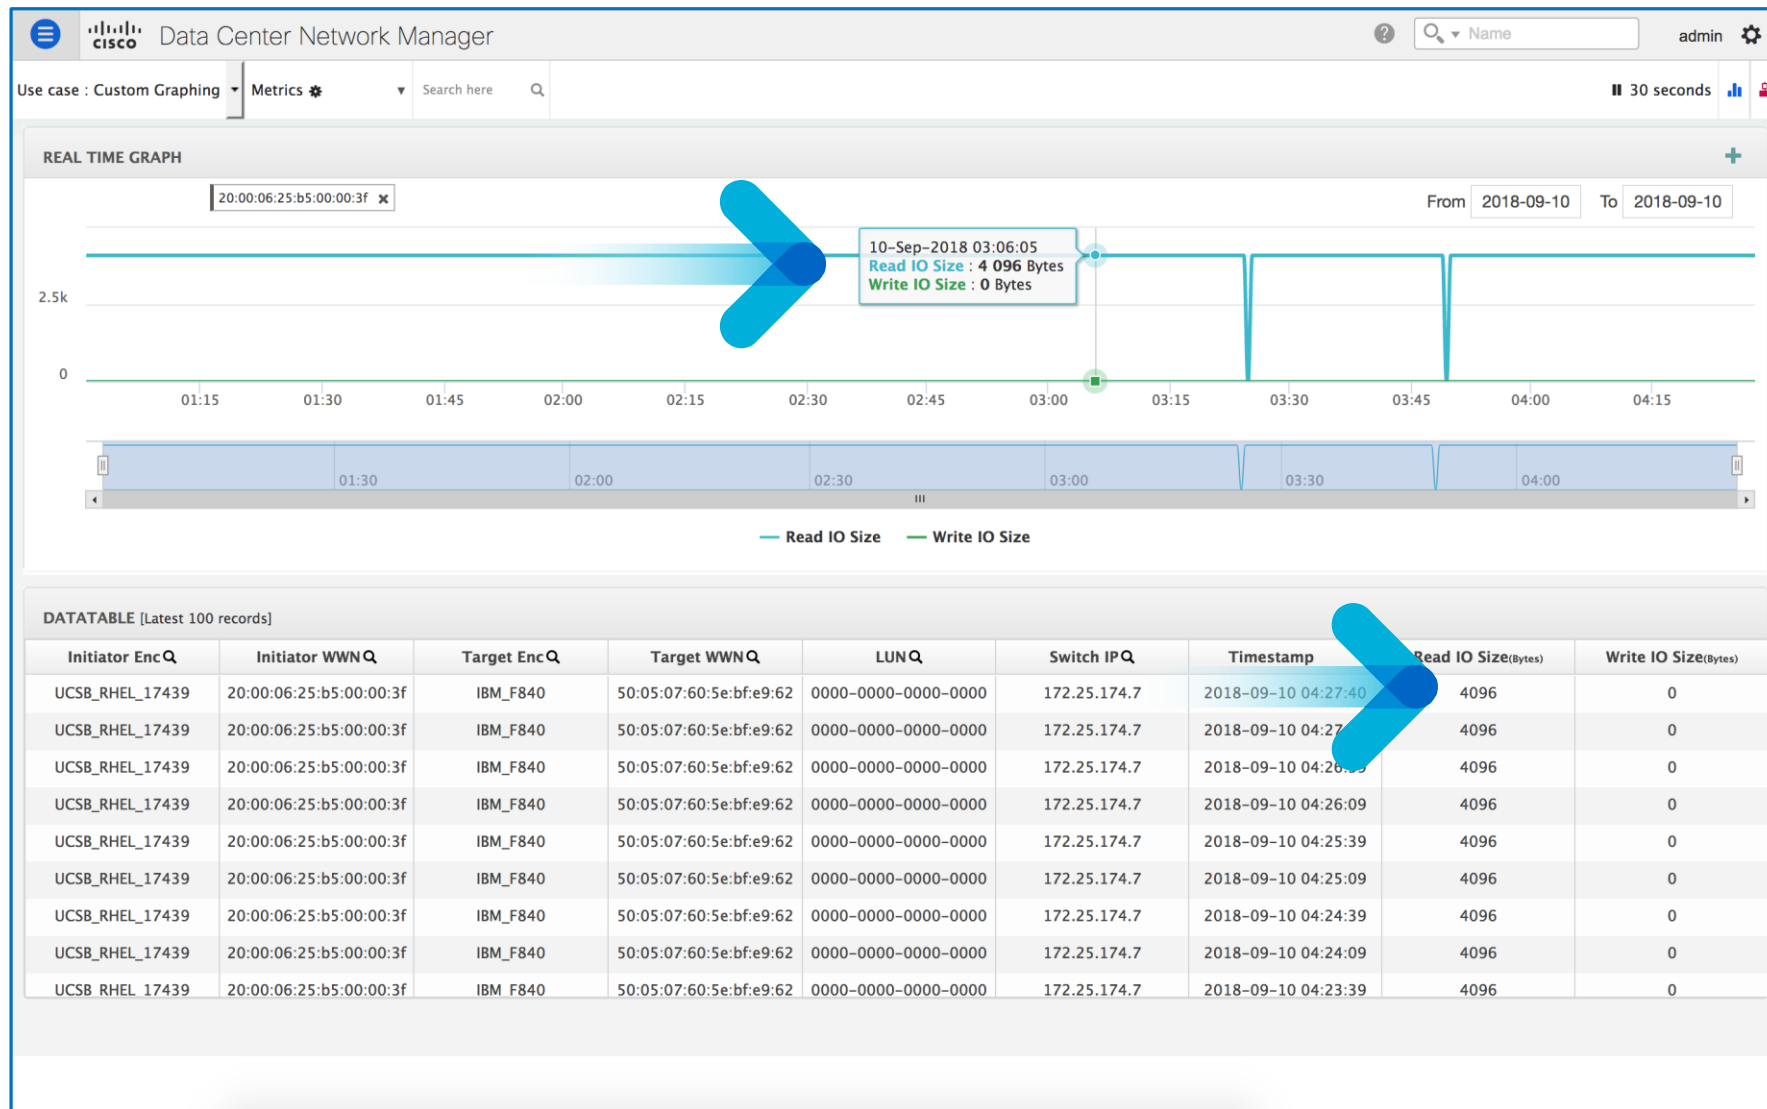
Name	Value
In Errors	0
Out Errors	0
In Discards	0
Out Discards	0
Tx	2.2750 B/s
Rx	0.4500 B/s



- Use Storage Enclosure view
- Find Storage Device in list
- Select Interface in Topology Map
- Check each path under Initiator Target Pairs
- Errors, Aborts and other error related metrics provided

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

What is my current I/O Block Size for my host?



- Use Custom Graphing
- Input WWPN or host name of host
- Select Read or Write IO Size in Metrics
- See block size over time

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

What is the Queue Depth for my host?

The screenshot shows the Cisco Data Center Network Manager interface. The top navigation bar includes 'Monitor / SAN / SAN Insights' and a legend for queue depth percentages: 0-1% (green), 1-15% (yellow), 15-30% (orange), and >30% (red). A time slider is set to 'now'. On the left, under 'Host Enclosures', the host 'Win_17416' is selected. The 'Initiator Target Pairs' table shows three entries with a 'Status' of 'Healthy' (green dot). Below this is a network diagram showing connections between various storage enclosures. At the bottom, two tables provide performance metrics for a specific flow and interface.

Host Enclosure	Status
Win_17416	Healthy
172.25.174.137	Healthy
172.25.174.138	Healthy
172.25.174.152	Healthy

Source PWWN	SID	Destination PWWN	DID	Fabric	Status
10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:10:1a:fa:44	d80040	Fabric_N5596UP-17486	Healthy
10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:12:27:61:13	d80220	Fabric_N5596UP-17486	Healthy
10:00:00:90:fa:47:6e:55	94010f	50:06:0e:80:13:27:63:00	d80200	Fabric_N5596UP-17486	Healthy

Name	1-Hour Average	Baseline
Read Active I/O	0.00	
Write Active I/O	0.14	
Read IO Aborts	0.00	
Write IO Aborts	0.00	
Read IO Failure	0.00	
Write IO Failure	0.00	

Name	Value
In Errors	0
Out Errors	0
In Discards	0
Out Discards	0
Tx	1.6750 MB/s
Rx	10.6625 MB/s



- Use Monitor/SAN Insights
- Select a Host on the left
- Scroll down and look at Read Active I/O and Write Active I/O
- number of outstanding IO's that are yet to be completed in real-time

For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Can I see NVMe Analytics data?


The screenshot shows the Cisco Data Center Network Manager interface. A blue arrow points to the 'NVMe' radio button in the top navigation bar. Below, the 'Host Enclosures' table shows 'NVME_SB_P0' with a red status indicator. The 'Initiator Target Pairs' table lists several connections, with the first one selected. A network diagram shows the connection between 'NVME_SB_P0' and 'MDS9396T-17445'. Below the diagram, two tables provide detailed performance metrics for the selected flow.

Host Enclosure	Read (% dev) Avg.	Write (% dev) Avg.
<input checked="" type="radio"/> NVME_SB_P0	●	●
<input type="radio"/> NVME_SCALE_INIT_0	●	●
<input type="radio"/> NVME_SCALE_INIT_3	●	●
<input type="radio"/> NVME_SCALE_INIT_2	●	●
<input type="radio"/> NVME_SCALE_INIT_4	●	●
<input type="radio"/> NVME_SCALE_INIT_1	●	●

Source PWWN	SID	Destination PWWN	DID	Fabric	Read (% dev) Avg.	Write (% dev) Avg.
<input checked="" type="radio"/> 20:00:00:11:0d:e5:fa:14	1d00d5	20:01:00:11:0d:e5:fb:13	1d00e3	Fabric_N93180FX...	●	●
<input type="radio"/> 20:00:00:11:0d:e5:fa:13	1d00d4	20:01:00:11:0d:e5:fb:12	1d00e2	Fabric_N93180FX...	●	●
<input type="radio"/> 20:00:00:11:0d:e5:fa:12	1d00d3	20:01:00:11:0d:e5:fb:11	1d00e1	Fabric_N93180FX...	●	●
<input type="radio"/> 20:00:00:11:0d:e5:fa:14	1d00d5	20:09:00:a0:98:9f:90:71	1d01e1	Fabric_N93180FX...	●	●
<input type="radio"/> 20:00:00:11:0d:e5:fa:11	1d00d2	20:01:00:11:0d:e5:fb:10	1d00e0	Fabric_N93180FX...	●	●

Name	1-Hour Average	Baseline
Average Read ECT Deviation	27.7098 %	
Average Write ECT Deviation	19.4052 %	
Average Read ECT	0.0754 ms/IO	
Average Write ECT	0.1165 ms/IO	
Average Read DAL	0.0737 ms/IO	
Average Write DAL	0.0879 ms/IO	
Read Active I/O	1	
Write Active I/O	1	
Read IO Aborts	0	
Write IO Aborts	0	
Read IO Failure	0	
Write IO Failure	0	
Read IOPs	784	

Name	Value (1-Hour)
In Errors	
Out Errors	
In Discards	
Out Discards	
Tx	
Rx	

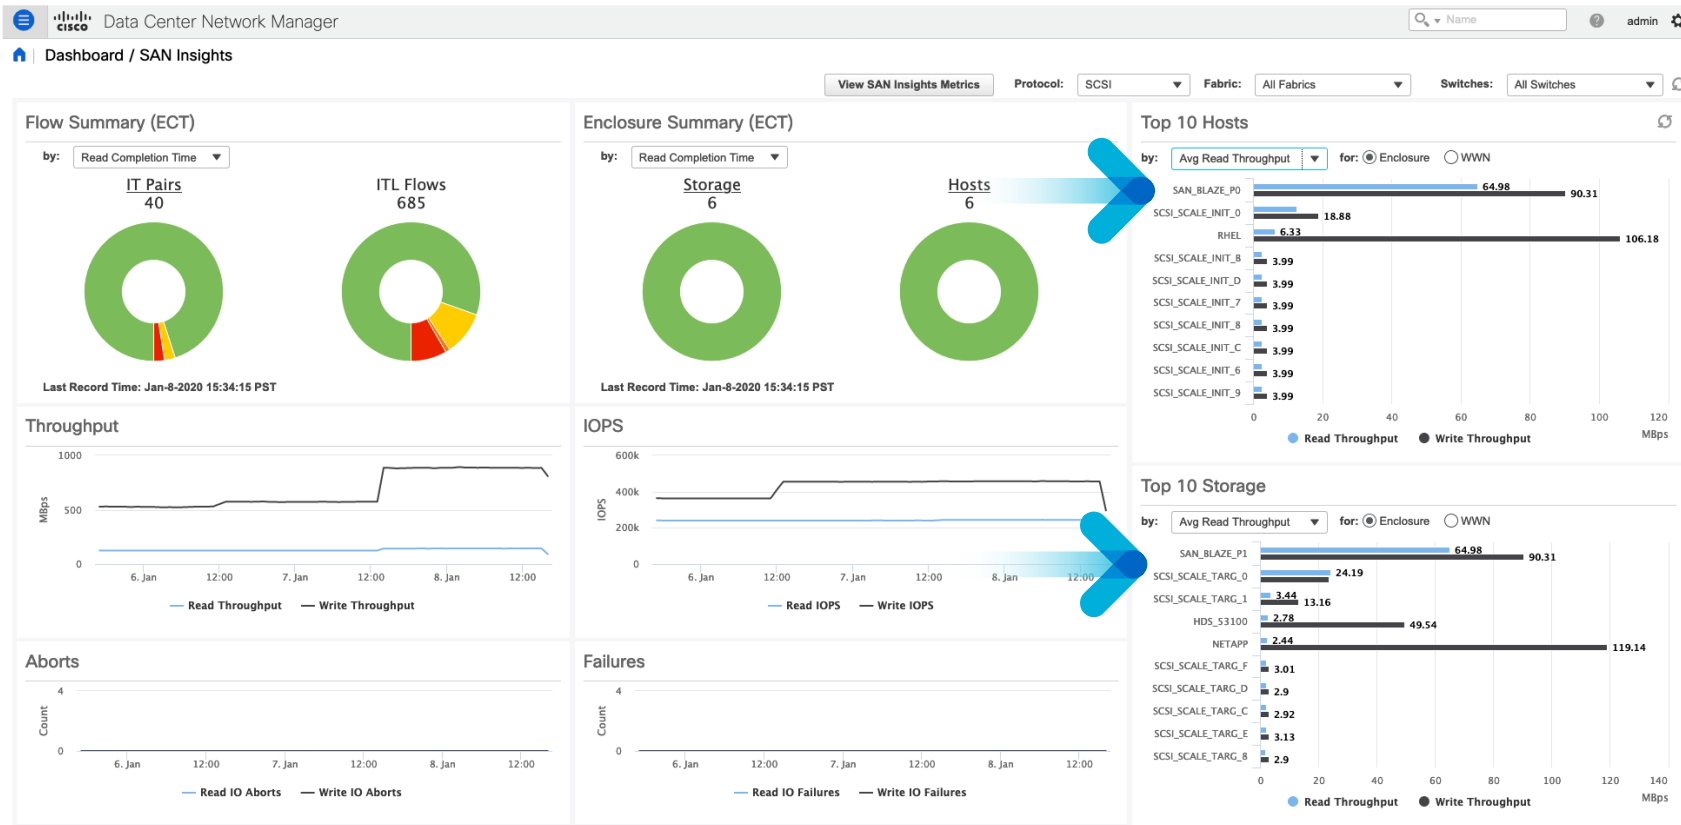


- Use Monitor/SAN Insights
- Select NVMe Radio button
- See all NVMe Flows



For more information please check <https://www.youtube.com/watch?v=mIWq7Oko0ml&t=>

Show me what is using the most throughput?



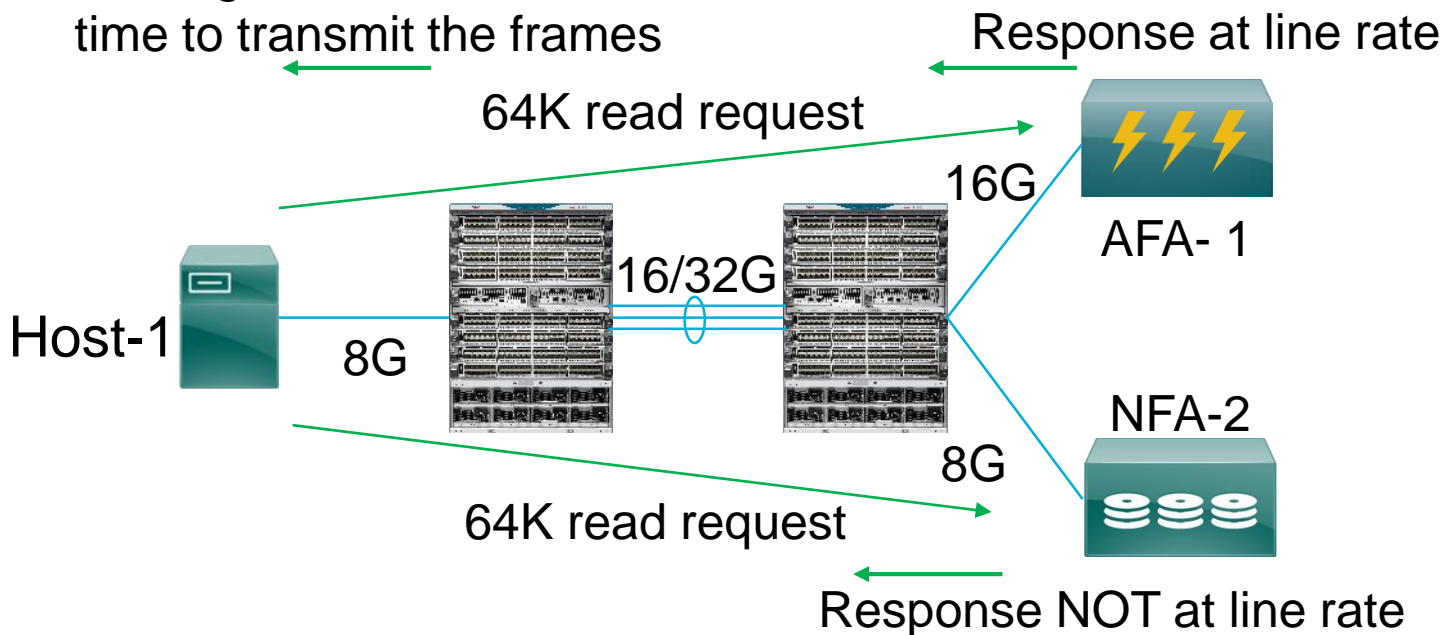
- Use Dashboard/SAN Insights
- Select Read or Write Throughput for Hosts
- Select Read or Write Throughput for Storage




For more information please check <https://www.youtube.com/watch?v=mIWq70ko0ml&t=>

Host Oversubscription in SAN

Host-edge switch needs more time to transmit the frames




Time to transmit 32 full size FC frames



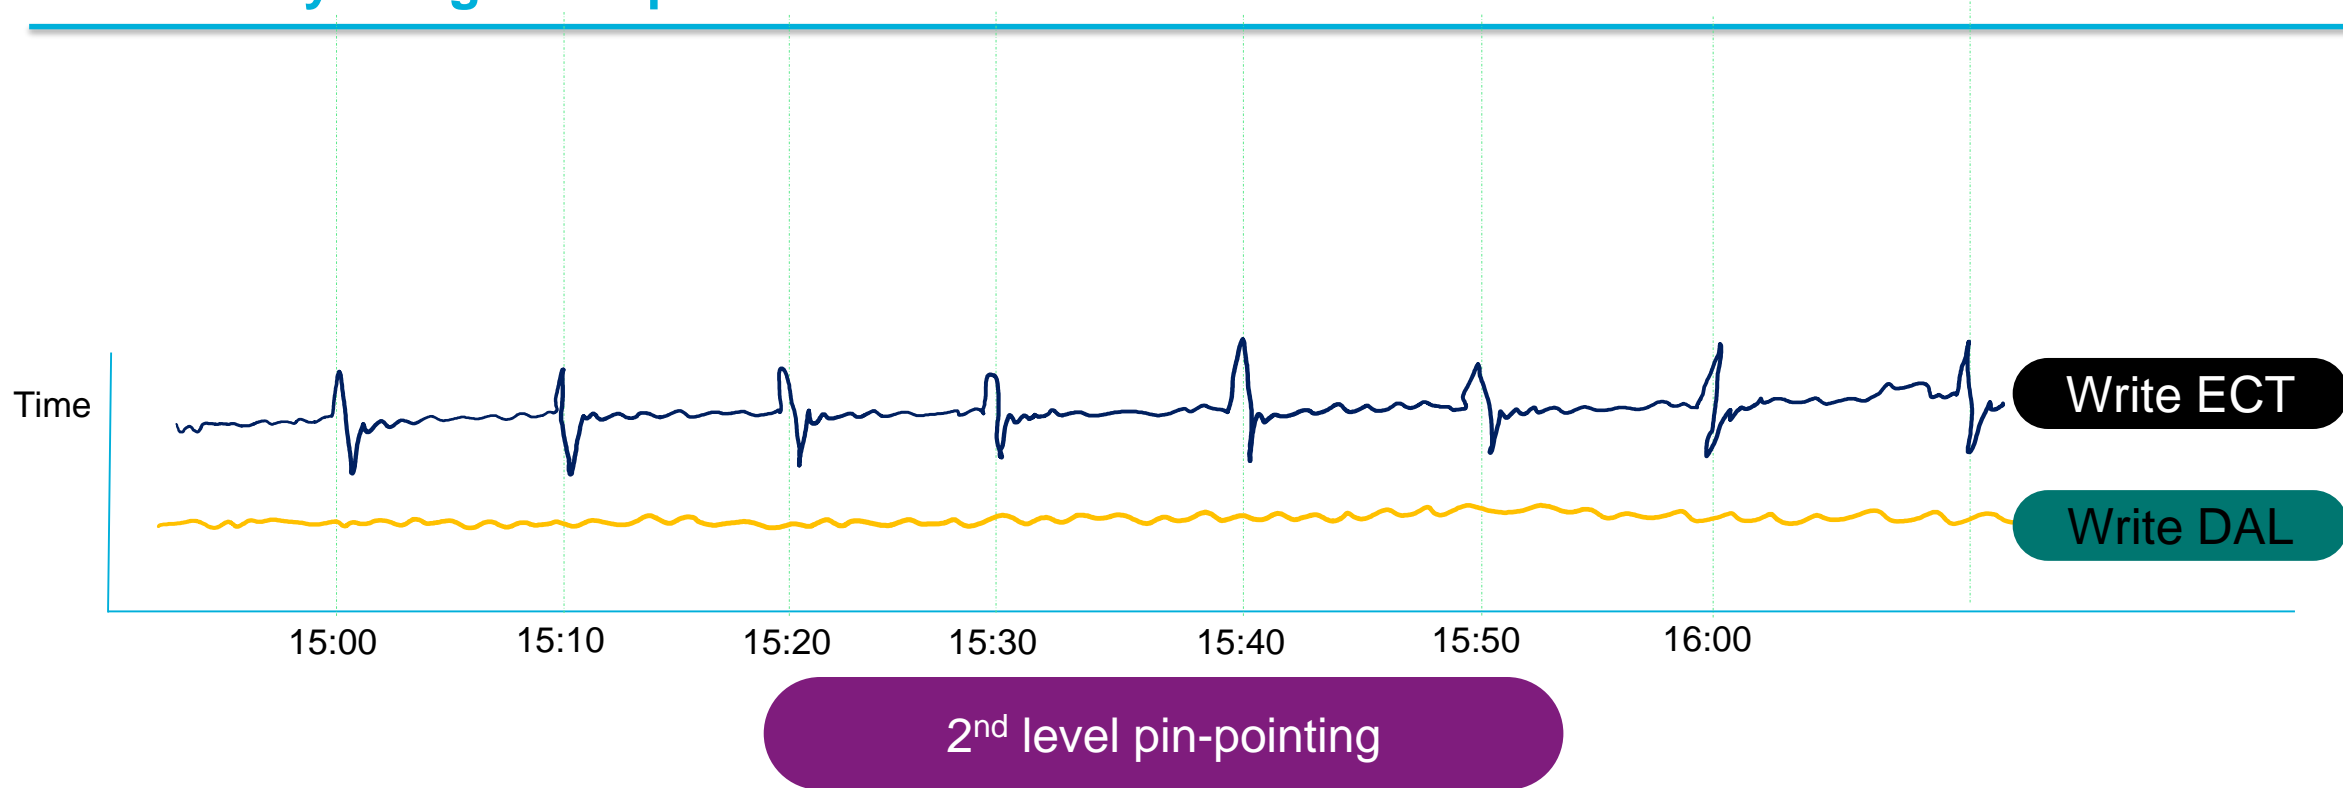
- AFA are extremely fast
- Frames are transmitted at line rate (of the directly connected interface)
- Backpressure is created if downstream links have less bandwidth
- Faster responses from AFA may not leave enough BW on host-edge for responses from NFA

Case study: Large European bank




- Write ECT spike followed by dip
 - May be the cause of application performance issues
- Frequency – every 10 minutes

Case study: Large European bank



- Write ECT spike followed by dip. Frequency – every 10 minutes
- DAL is stable (no change)
 - Not a storage array issue

Case study: Large European bank



3rd level pin-pointing


Write ECT spike followed by dip. Frequency – every 10 minutes

DAL is stable (no change). Not a storage array issue

No changes in Read ECT and DAL. No fabric congestion observed.

No indication of fabric delay. Indication of delay within host.

Case study: Large European bank



3rd level pin-pointing

Write ECT spike followed by dip. Frequency – every 10 minutes

DAL is stable (no change). Not a storage array issue

No changes in Read ECT and DAL. Not a fabric issue

Delay within host → Resulted in detection of an unpatched Oracle app on host

DCNM 11.3(1)

Features – SAN Insights -
Licensing

DCNM SAN Insights and SAN Telemetry Streaming Licensing

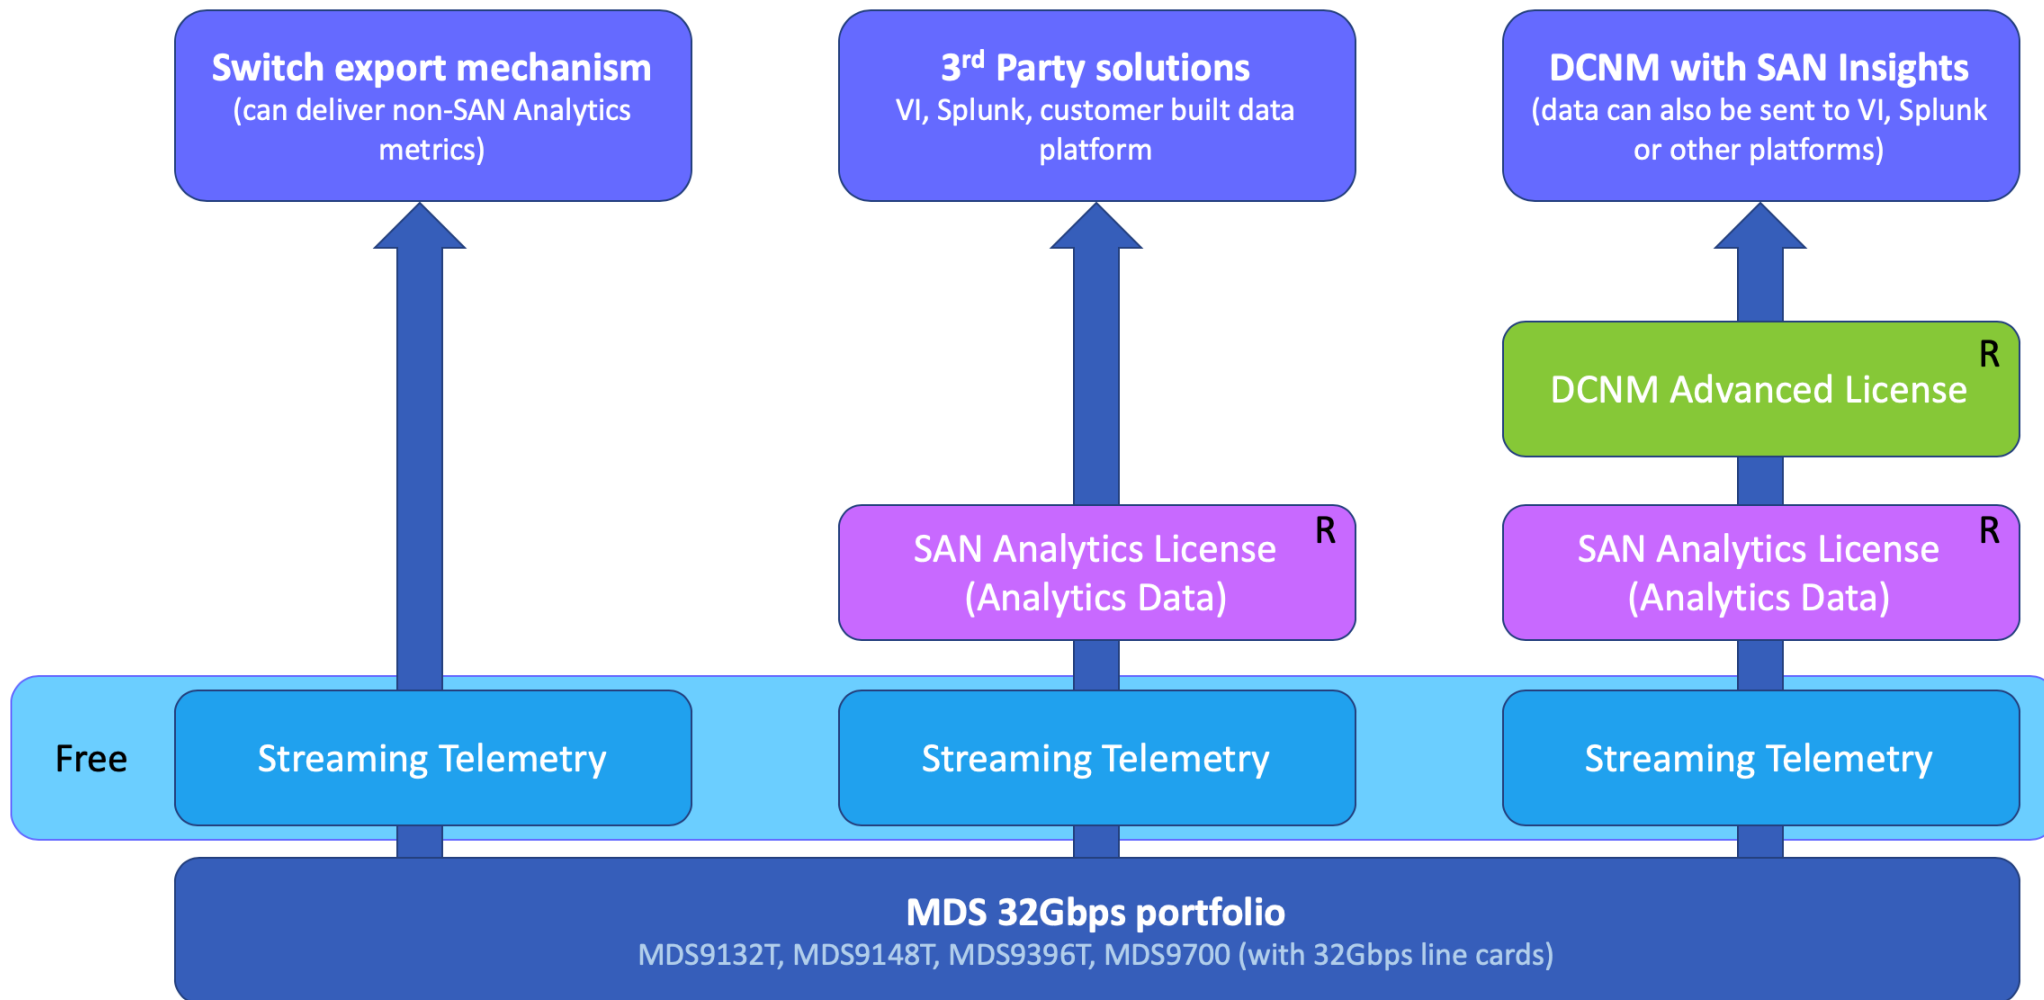
SAN Insights – This is a switch-based license that enables the switch to provide both on-board and outbound access to the telemetry data provided by the 32GB module with the added functions of DCNM visualization. This is provided by the enablement of the SAN_ANALYTICS_PKG on each switch.

SAN_ANALYTICS_PKG is a switch-based license (tied to the serial number of the switch). A single license allows you to enable ITL metric collection on all the ports on the licensed switch.

SAN_ANALYTICS_PKG is a time-based license which is offered in 3 years and 5 years term. At the end of the term of the license the switch will provide a message 30, 60 and 90 days before the expiration date. DCNM will also show the expiration of the license in the license views. You need to buy a similar license at the expiry of the term to continue using the functionality.

Since only the ITL flow metric collection is licensed but export is not, SAN_ANALYTICS_PKG is enough to export metrics to DCNM.

DCNM SAN Insights and SAN Telemetry Streaming Licensing



R= Required License

DCNM 11.3(1)

Features – SAN Insights -
Summary

Session summary

- Hardware and software-based solution
- Deploy anywhere in the fabric as long as 32Gbps line card or 32Gbps switch is in the path
- Uses industry standard form factors for compression and transport
- Supports CLI, pull and push methods for Analytics data delivery
- 60+ metrics
- DCNM is the destination for SAN Analytics data and is contained within a feature called SAN Insights
- SAN Analytics is licensed per switch (telemetry is free)
- Supports 20K flows per DCNM server
- Must be at NX-OS 8.3.1 and DCNM 11.2.1 minimum



Washington Systems Center - Storage



Thank You!