

Cisco Tech Day

Denver

March 3, 2026

Storage Networking



Cisco Tech Day
Denver



Todd Philippi
DC Solutions Engineer

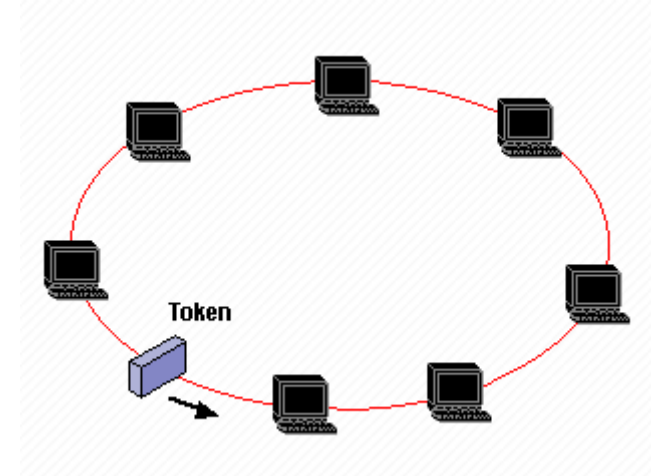
Agenda

- 01 **Ethernet Always Wins**
- 02 **NVME/TCP**
- 03 **SAN congestion solutions**
- 04 **Demo**

Ethernet Never Loses!

Ethernet Never Loses!

Token Ring?



1993 Cisco Acquires Crescendo Networks.


100 Mbps (Fast Ethernet) adapters \$119

IBM 16Mb LAN streamer adapter = \$1,455

Ethernet Never Loses!

Token Ring?

TDM/Voice/PBX Networks?



Cisco Acquires Celsius systems - Oct, 1998

Ethernet Never Loses!

Token Ring?

TDM Networks?

Physical Security/Video Surveillance?



March 2007 Cisco acquires SyPixx

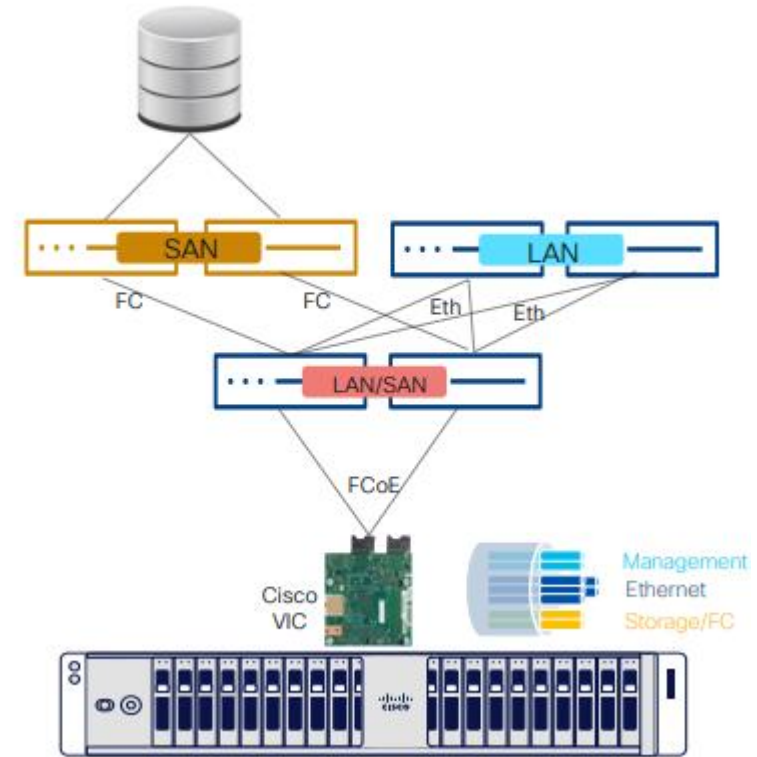
Ethernet Never Loses!

Token Ring?

TDM Networks?

Physical Security/Video
Surveillance?

FCOE



Cisco acquires Nuova Systems - May 2008 (MPLS)
FCOE and UCS


2008 – I learned FC

Brocade merges with McData.

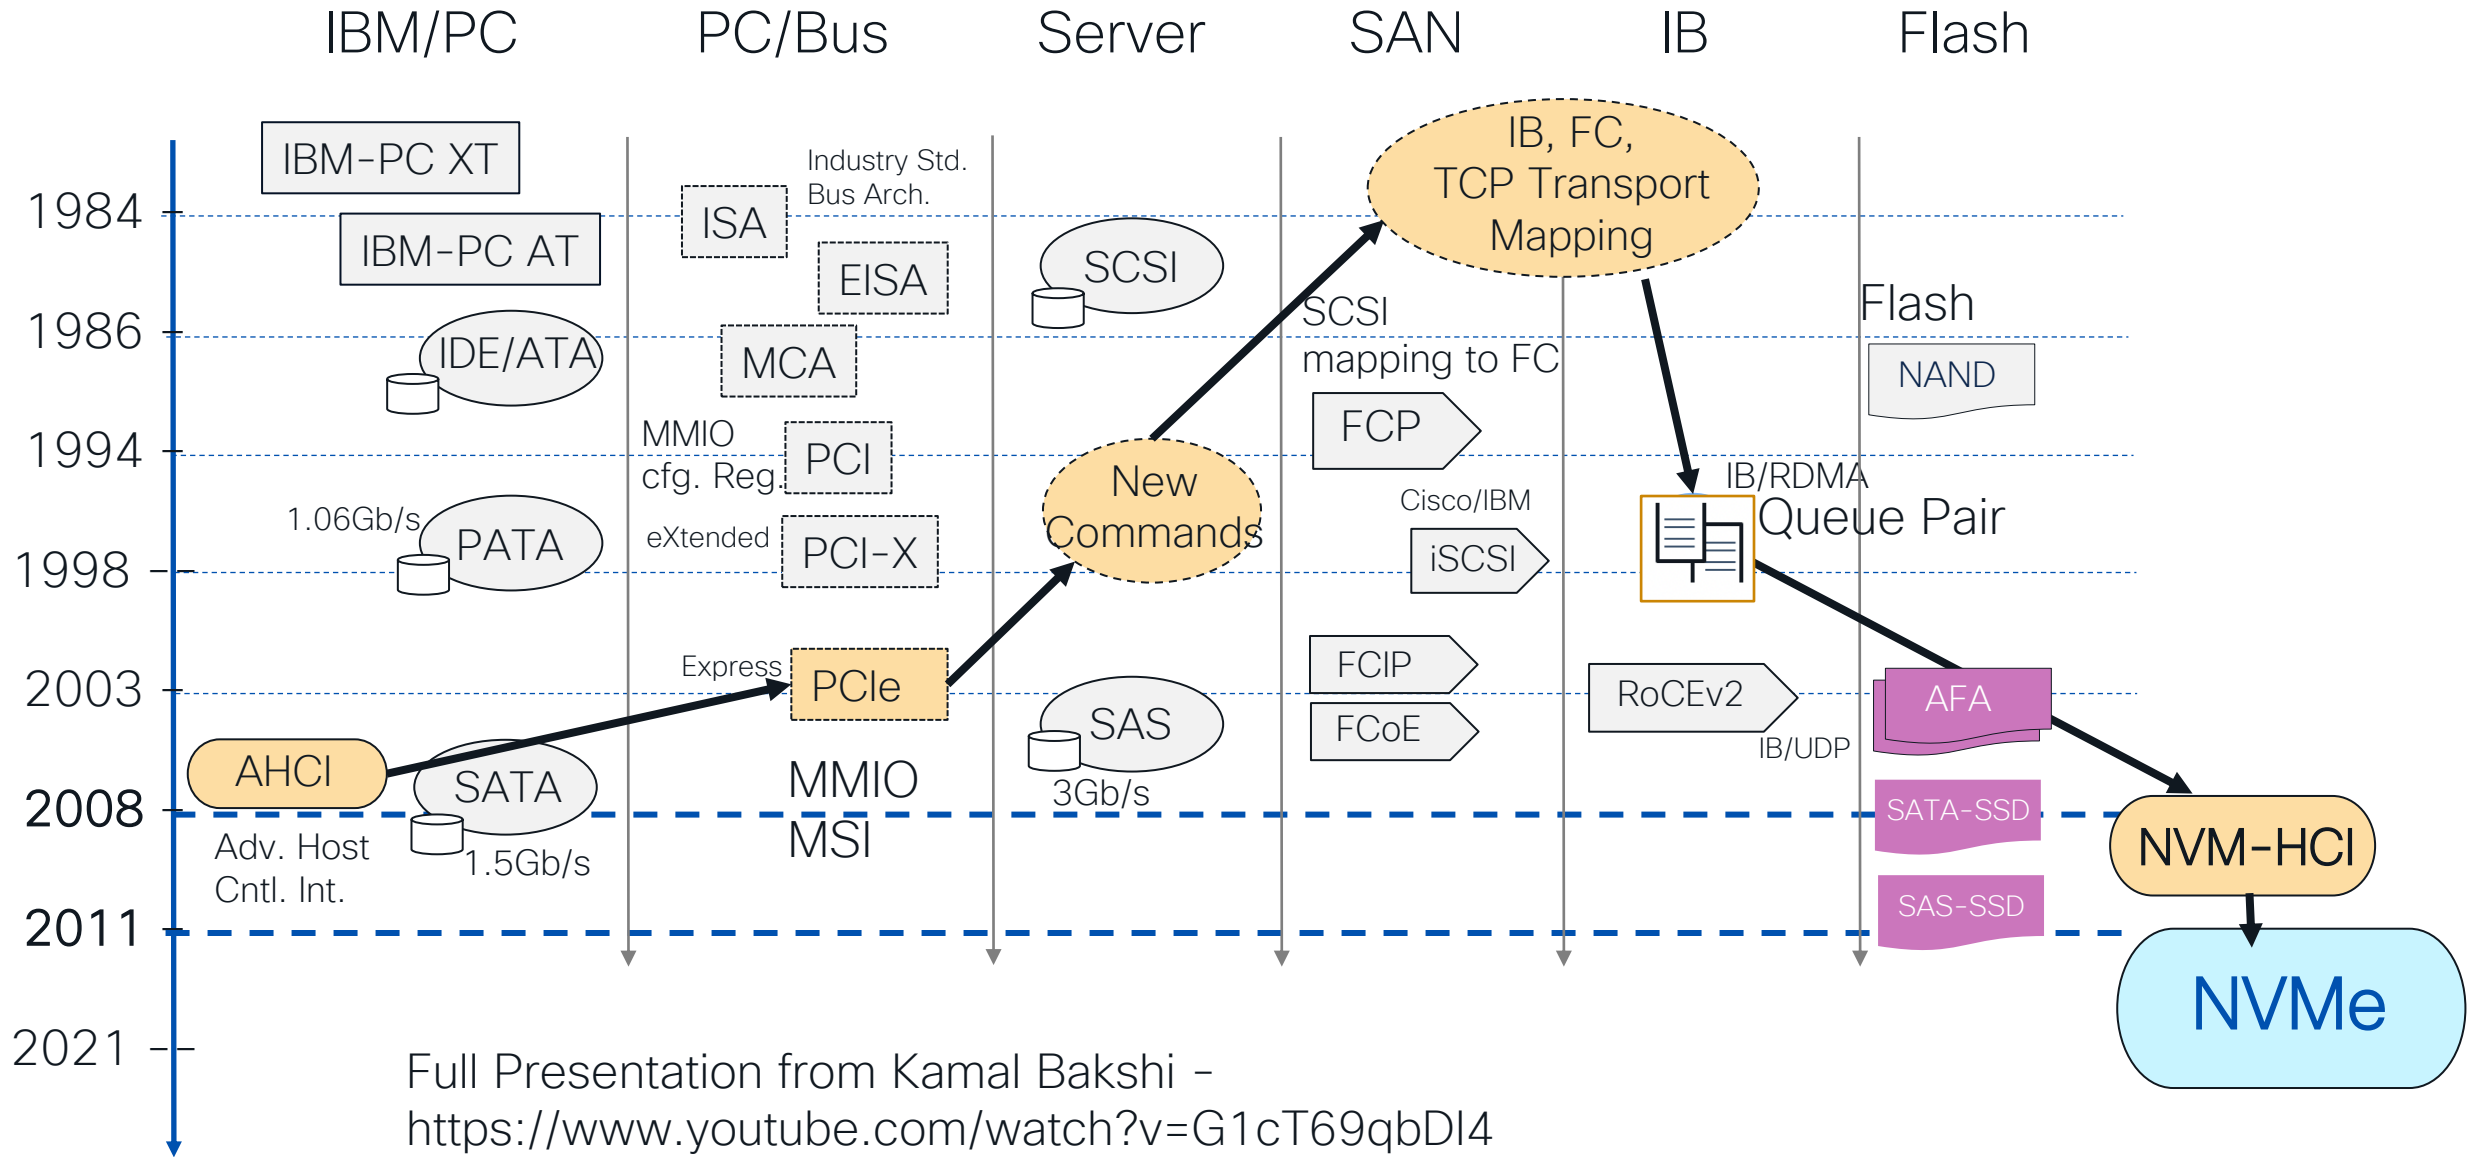
My customer, IBM, needs a second vendor with Cisco MDS.

I'm invited to an MDS training class with my customer.

I found out.

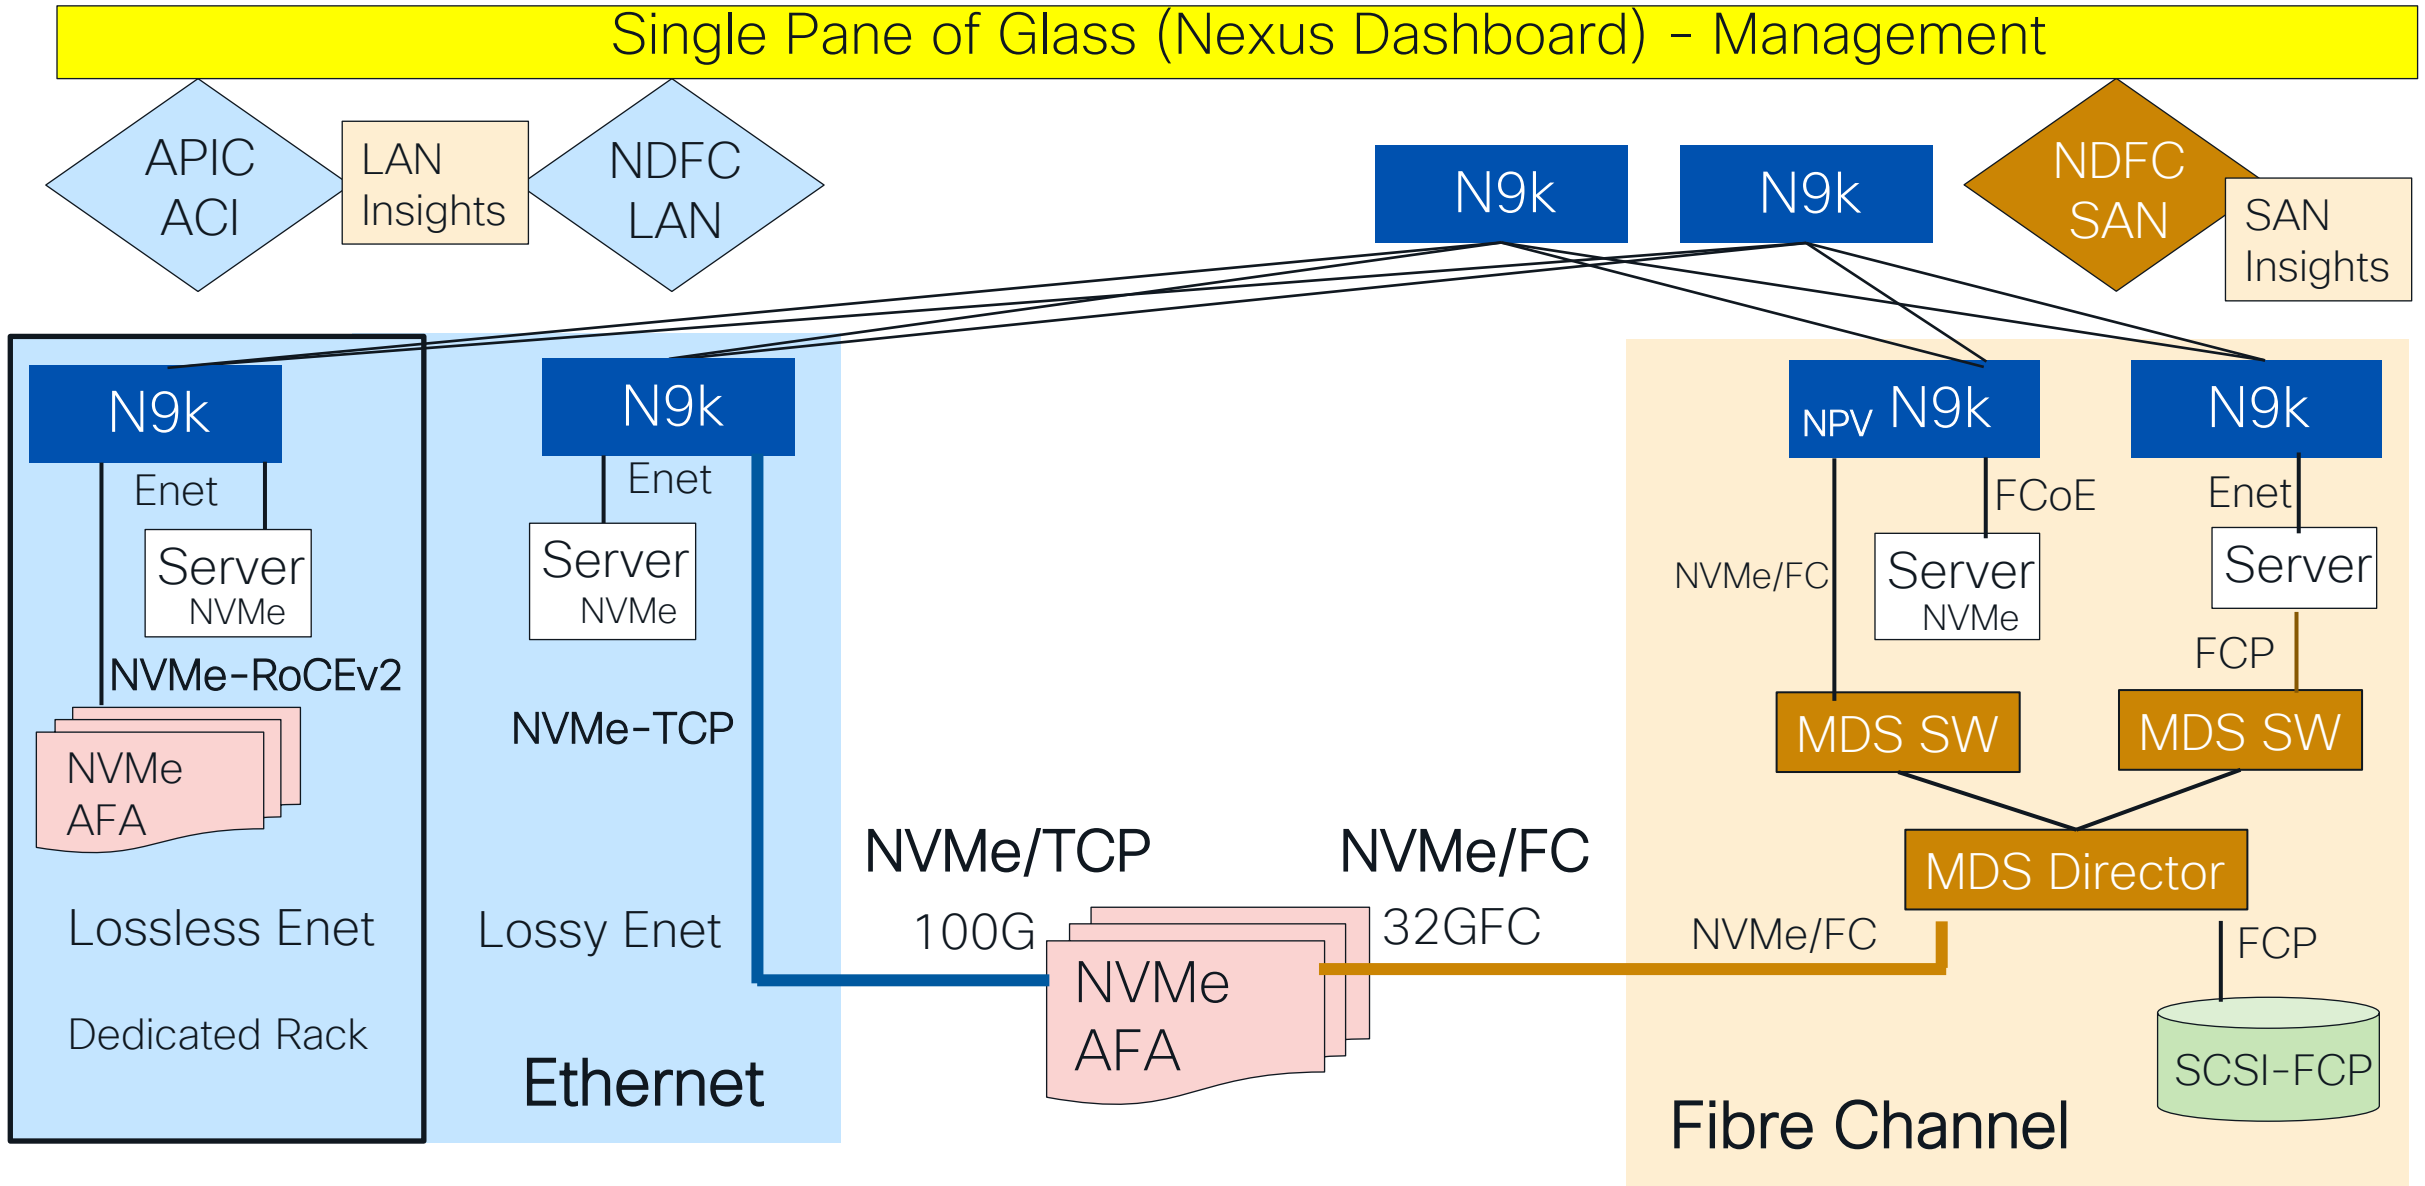


NVMe - Best of all worlds....



Full Presentation from Kamal Bakshi - <https://www.youtube.com/watch?v=G1cT69qbDI4>

Cisco NVMe-Anywhere (Fibre Channel / Ethernet)



VIC 15000 Series

- 5th Gen VIC card for X- , B- *, C- Series
- Supports 10G/25G/40G/50G/100G/200G*
- CNA, Single Wire Mgmt
- Dynamic FC and Ethernet virtual interfaces
- x16 PCIe Gen 4
- NVMeoF: FC-NVMe, RoCEv2
- Overlays: NVGRE, VXLAN, Geneve
- RSS, NetQueue, VMQ, VMMQ, RSSv2*
- SR-IOV*, SIOV*, usNIC, DPDK
- PTPv2, L3ECN*, 16K Rx Ring Size
- VIC 15231 (x210c MLOM) and VIC 15428 (rack-server MLOM) at FCS
- VIC 15231 (2x100G) and VIC 15428 (4x10/25/50G)



15231



15428

* Various VIC PIDs and features will be available post FCS

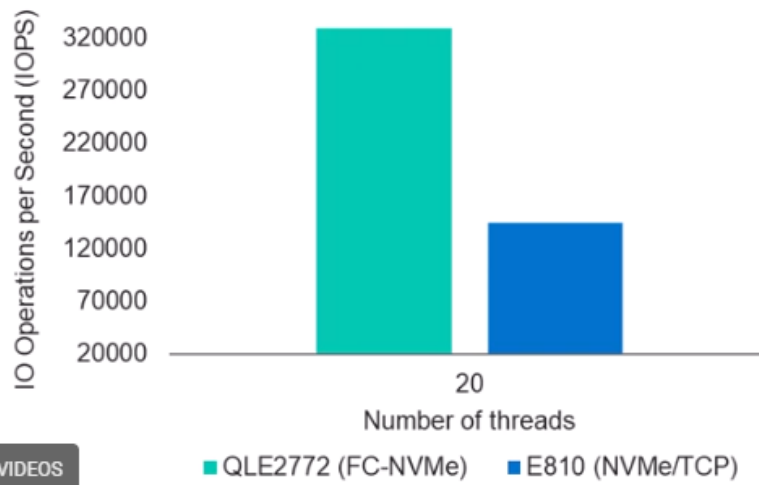
Marvell/Qlogic testing (FC vendor Bias?)

<https://www.youtube.com/watch?v=4SYVgNtx0Fs>

NVMe-oF™ transactional performance under medium load

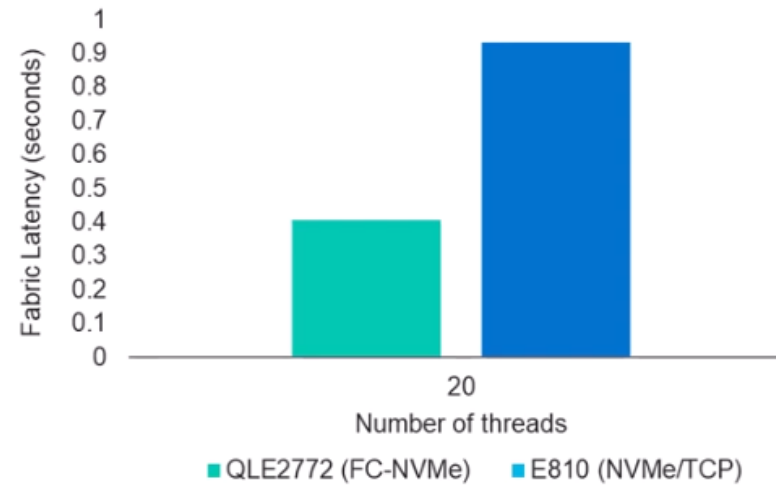
FOLLOW
THE WIRE

NVMe over Fabrics IOPS comparison
ESXi 7.0 U3 Win VM 4KB Read (IOPS)
[with 2xPVNVMe]
32GFC-NVMe vs. 25GbE NVMe/TCP



MORE VIDEOS

NVMe over Fabrics latency comparison
ESXi 7.0 U3 Win VM 4KB Read (Latency)
[with 2xPVNVMe]
32GFC-NVMe vs. 25GbE NVMe/TCP



SUSE Performance on NVMe-TCP

<https://www.youtube.com/watch?v=D-TzcH99AF8>



SUSE Labs Conference 2025

NVMe multiqueue design


The block layer assigns a set of CPUs to each queue.

So eg CPU1-4 are assigned to Queue A, and CPU 5-8 to Queue B.

Okay.

But NVMe-TCP is scheduling a workqueue function to handle each queue, and this function is running on a specific CPU, so Queue A will be handled from CPU 1, and Queue B handled from CPU 5.

Not optimal, but okay.



6



Dell (EMC) NVMe transport performance Comparison

<https://www.delltechnologies.com/asset/en-gb/products/storage/industry-market/h18892-nvme-transport-performance-comparison.pdf>

Dell Technologies

Technical White Paper

NVMe Transport Performance Comparison

Dell PowerStore: Fibre Channel (FC), NVMe/FC, NVMe/TCP and iSCSI

Abstract

This technical white paper provides the results of NVMe transport performance testing using Dell PowerStore with the same workload configurations to determine performance tradeoffs.

VSAN with NVMe drives

<https://blogs.vmware.com/cloud-foundation/2022/10/18/designing-vsan-networks-2022-edition-vsan-esa/>



Technical/How-To | VCF Storage (vSAN)

Designing vSAN Networks – 2022 Edition vSAN ESA



John Nicholson

October 18, 2022

Share on:



Back in 2019, I wrote an update on common vSAN network design choices. I thought it would be a good idea to by discussing how these questions have aged, and what new questions have come up. VMware vSAN Express Storage Architecture™ (ESA) is helping customers drive new levels of extreme performance and with faster storage processing comes increased demands on the networking stack.

What is new for vSAN Express Storage Architecture (ESA) Networking?

Flashstack with Nutanix



- Hypervisor choice and licensing flexibility with Nutanix AHV. Unified operations using Intersight (compute), Prism (virtualization), and Pure1 (storage).
- A validated, prescriptive architecture that reduces deployment risk.
- Independent scaling of compute and storage for efficient lifecycle management.
- Consistent, predictable performance through FlashArray SAN and **NVMe/TCP transport**.

This is a next-generation FlashStack design built for customers needing performance, simplicity, and modernization without disruption



DCB in RoCEv2

RoCEv2 utilizes the same [Data Center Bridging \(DCB\)](#) standards originally developed to support **FCoE** (Fibre Channel over Ethernet).

Both protocols require a "lossless" Ethernet environment because they do not have the same robust retransmission mechanisms as TCP/IP. They rely on these specific DCB components:

- **Priority Flow Control (PFC - 802.1Qbb)**: The most critical piece. It allows the network to [pause specific traffic classes](#) (like storage or RDMA) instead of dropping packets during congestion, while letting normal web traffic continue unaffected.
- **Enhanced Transmission Selection (ETS - 802.1Qaz)**: Used to [allocate guaranteed bandwidth](#) to different traffic types so that RoCEv2 or FCoE doesn't get "starved" by a burst of background data.
- **DCB Exchange (DCBX)**: A discovery protocol that ensures both the [server and the switch](#) agree on the lossless configuration.



DCB in RoCEv2

While they share the same underlying DCB tools at the hardware level, their "reach" differs:

- **FCoE** is strictly a **Layer 2** protocol; it cannot be routed beyond a single local network.
- **RoCEv2** encapsulates its data in **UDP/IP**, allowing it to be [routed across different subnets](#) while still using DCB for the "local" hops.
- In modern setups, RoCEv2 often adds **Explicit Congestion Notification (ECN)** to these DCB features to proactively manage congestion before it even triggers a "pause".

Cisco MDS Switch Hardware

SAN directors

16/32/64G

MDS 9718
768 ports

MDS 9710
384 ports

MDS 9706
192 ports



Fabric switches

32G

MDS 9132T

MDS 9148T

MDS 9396T

64G

MDS 9124V

MDS 9148V

MDS 9396V

FCIP

MDS 9220i

MDS 9250i

Modules

64G

48 x 64G FC

32G

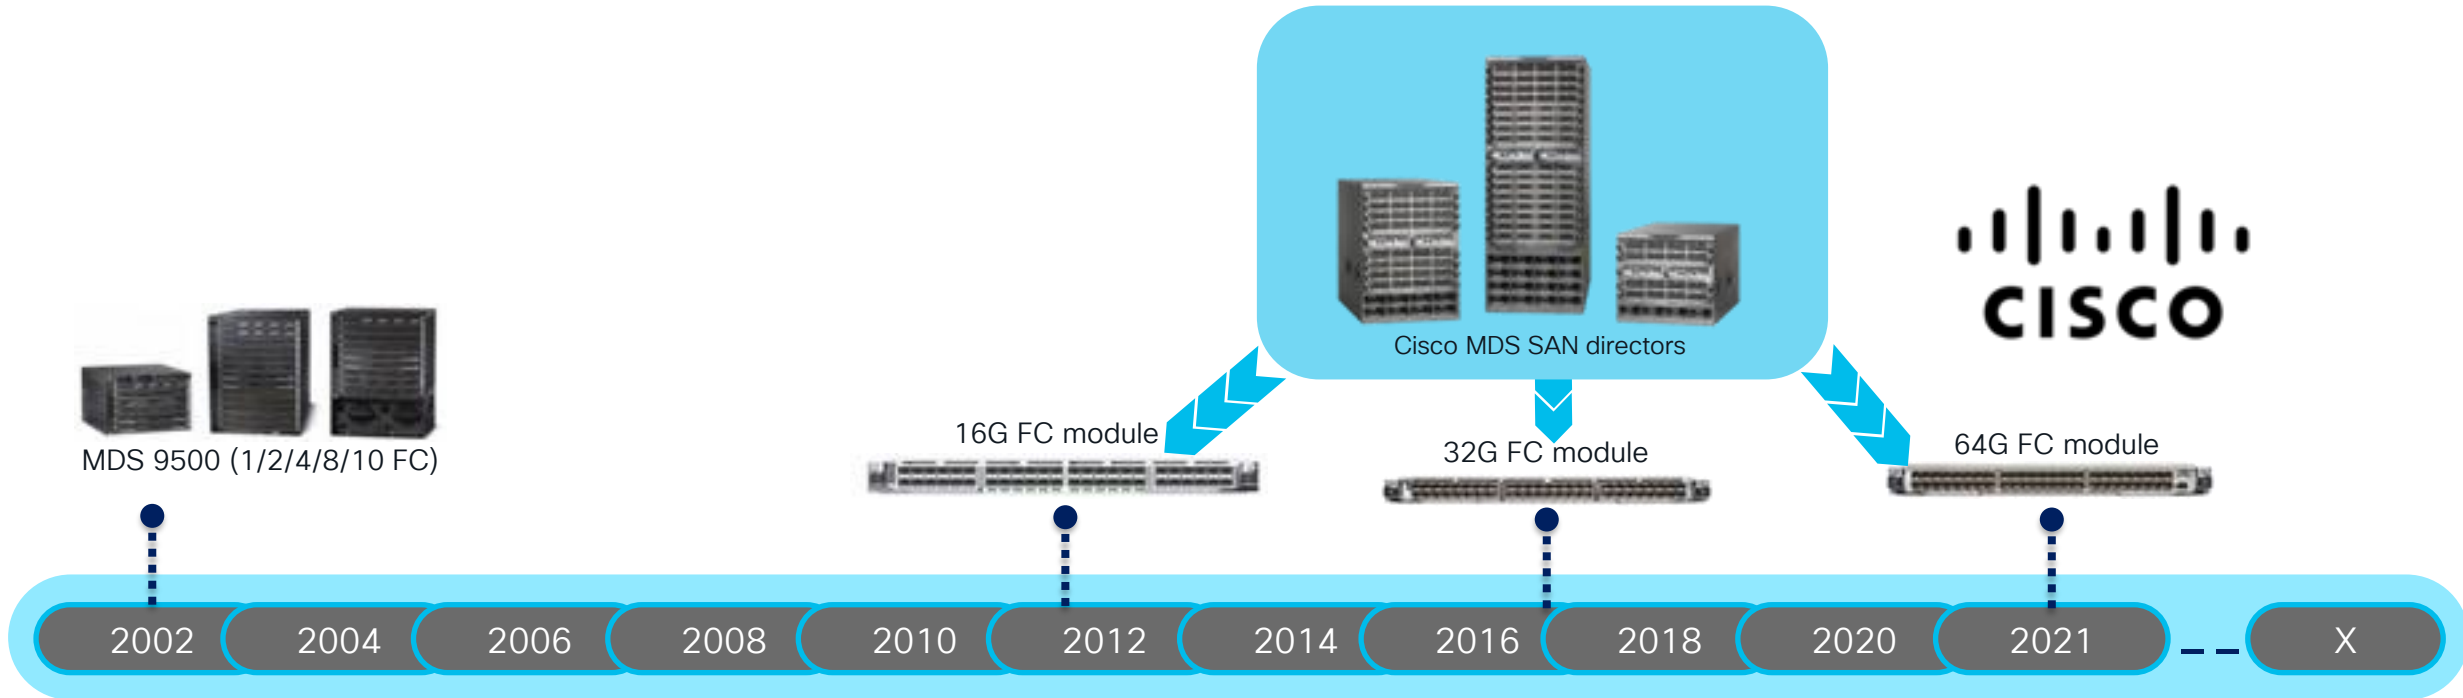
48 x 32G FC

FCIP

24 x 16G FC, 8 x 1/10 GE & 2 x 40 GE

- 8/16/32/64G FC
- 4/8/16/32G FC
- 2/4/8/10/16G FC
- SAN Analytics
- FCIP
- FICON

Cisco MDS Switch Hardware



Non-disruptive multi-generation upgrades

Multi-generation optics support

8/16/32/64G FC optics

Multi-generation module support

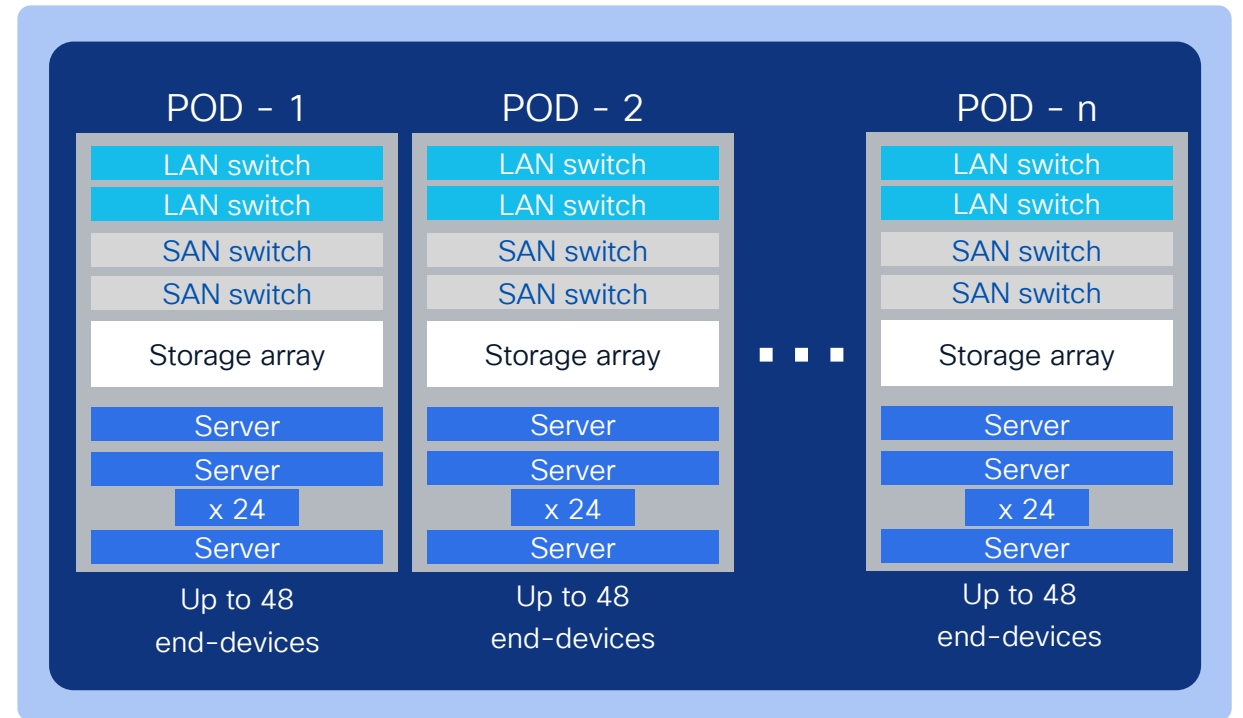
16/32/64G FC module

Multi-generation FC speed support

2/4/8/10/16/32/64G FC

Cisco MDS SAN directors: Single chassis multi-generation support

POD Fabric



Benefits

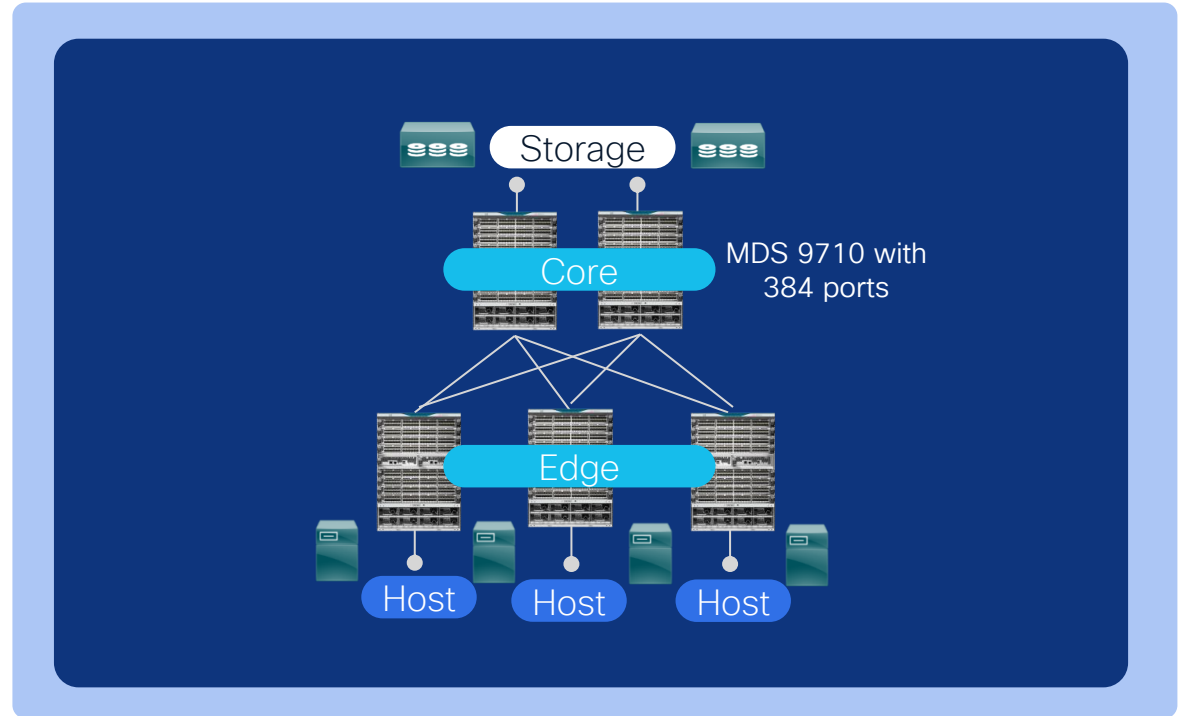
Best at containing SAN congestion / slow drain within the rack

No effect on other racks while one rack is under maintenance

Challenges

Requires a different data center operational model

Edge-Core Design



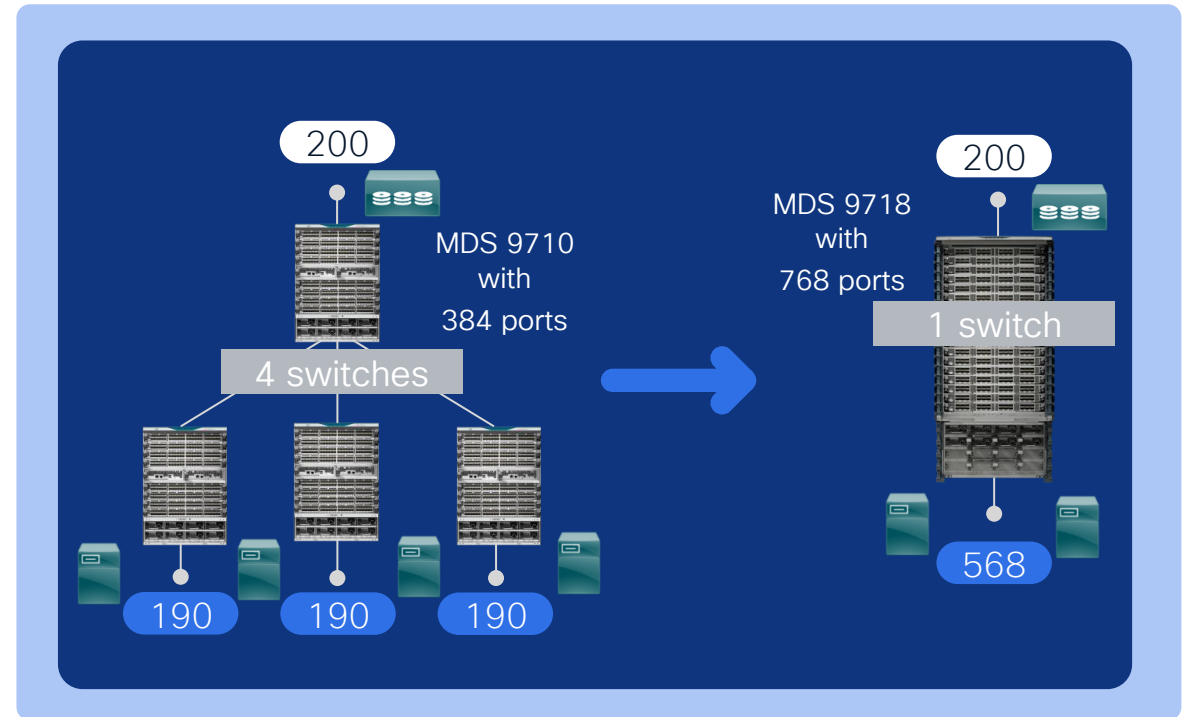
Benefits

Most Deployed. Tried and Tested

Challenges

Prone to congestion and slow drain

Consolidated Switch Design



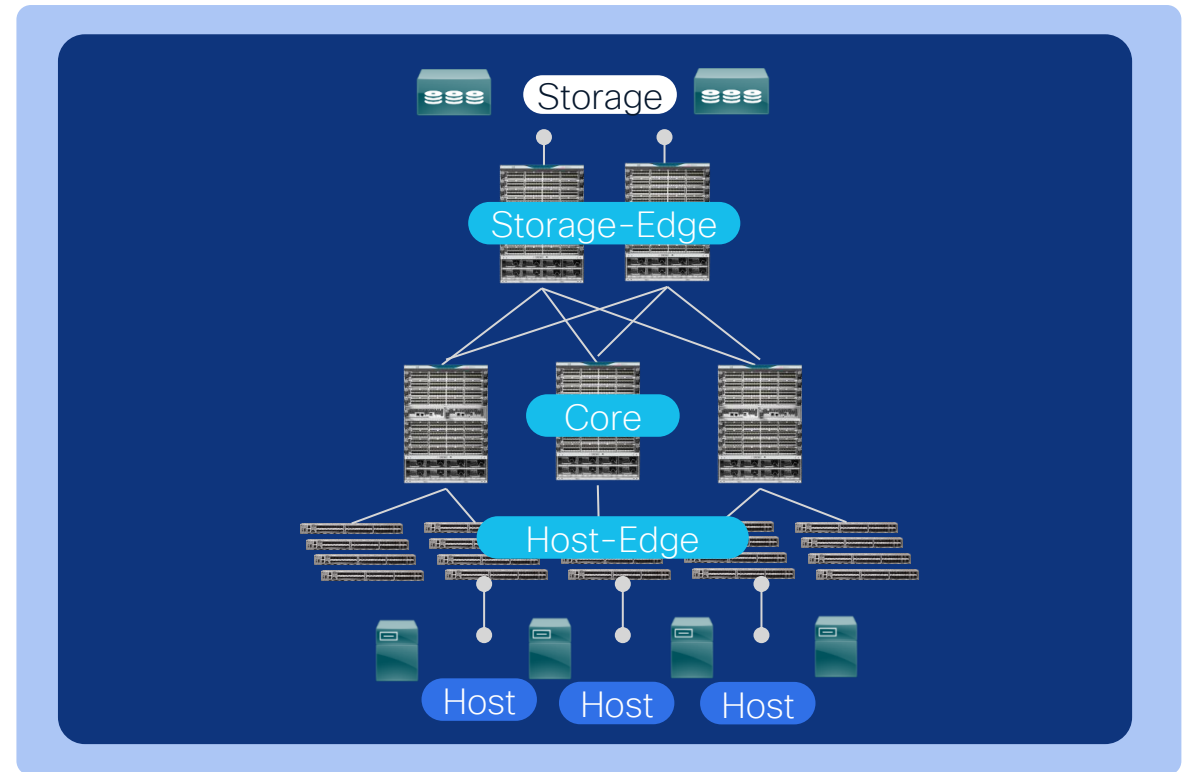
Benefits

- Less operational overhead
- Reduced Congestion/Slow drain

Challenges

- Cable length and scaling beyond 768 devices

Large Scale Edge-Core Design



Benefits

Easy to Scale and Investment protection

Reduced Congestion/Slow drain


Challenges

Prone to Congestion/Slow drain

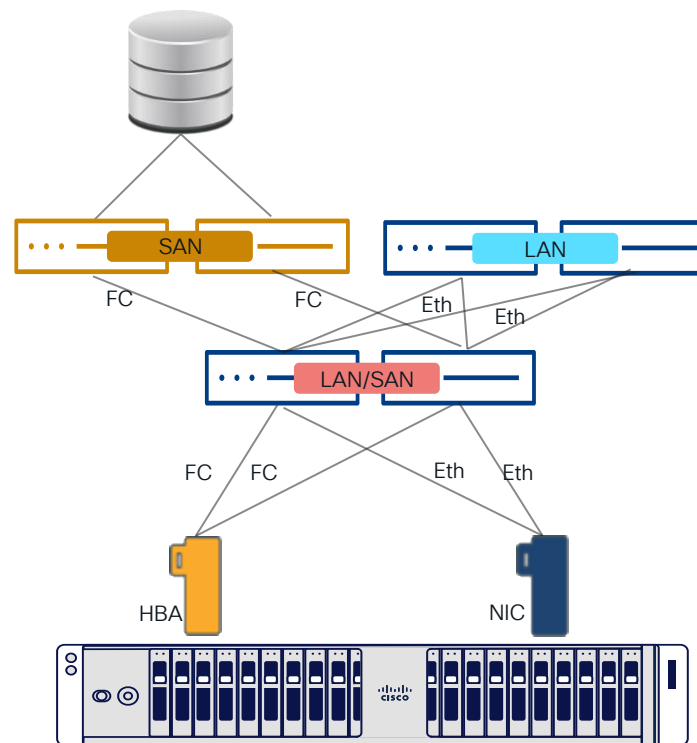
Strict operational control

Typical LAN and SAN Designs


LAN and SAN separation



Convergence at ToR switch



Converged at ToR switches and host



Understanding SAN Congestion

Understanding SAN Congestion

Fibre Channel Buffer-to-Buffer Flow Control - The Basics

Fibre Channel is a 'lossless' network protocol

Sender does not send a frame unless the receiver has a buffer

'Fibre Channel utilizes Buffer-to-Buffer(B2B) Credit based flow control

Each side of link informs adjacent side of the number of buffers/credits

Each frame sent requires a B2B credit to be returned

B2B credits are also called 'R_RDYs'

Frame receivers can slow rate of ingress traffic by 'withholding' credits

If a sender runs out of credits it must stop sending until it receives one



Understanding SAN Congestion

There are 4 reasons for congestion in a Fibre Channel SAN

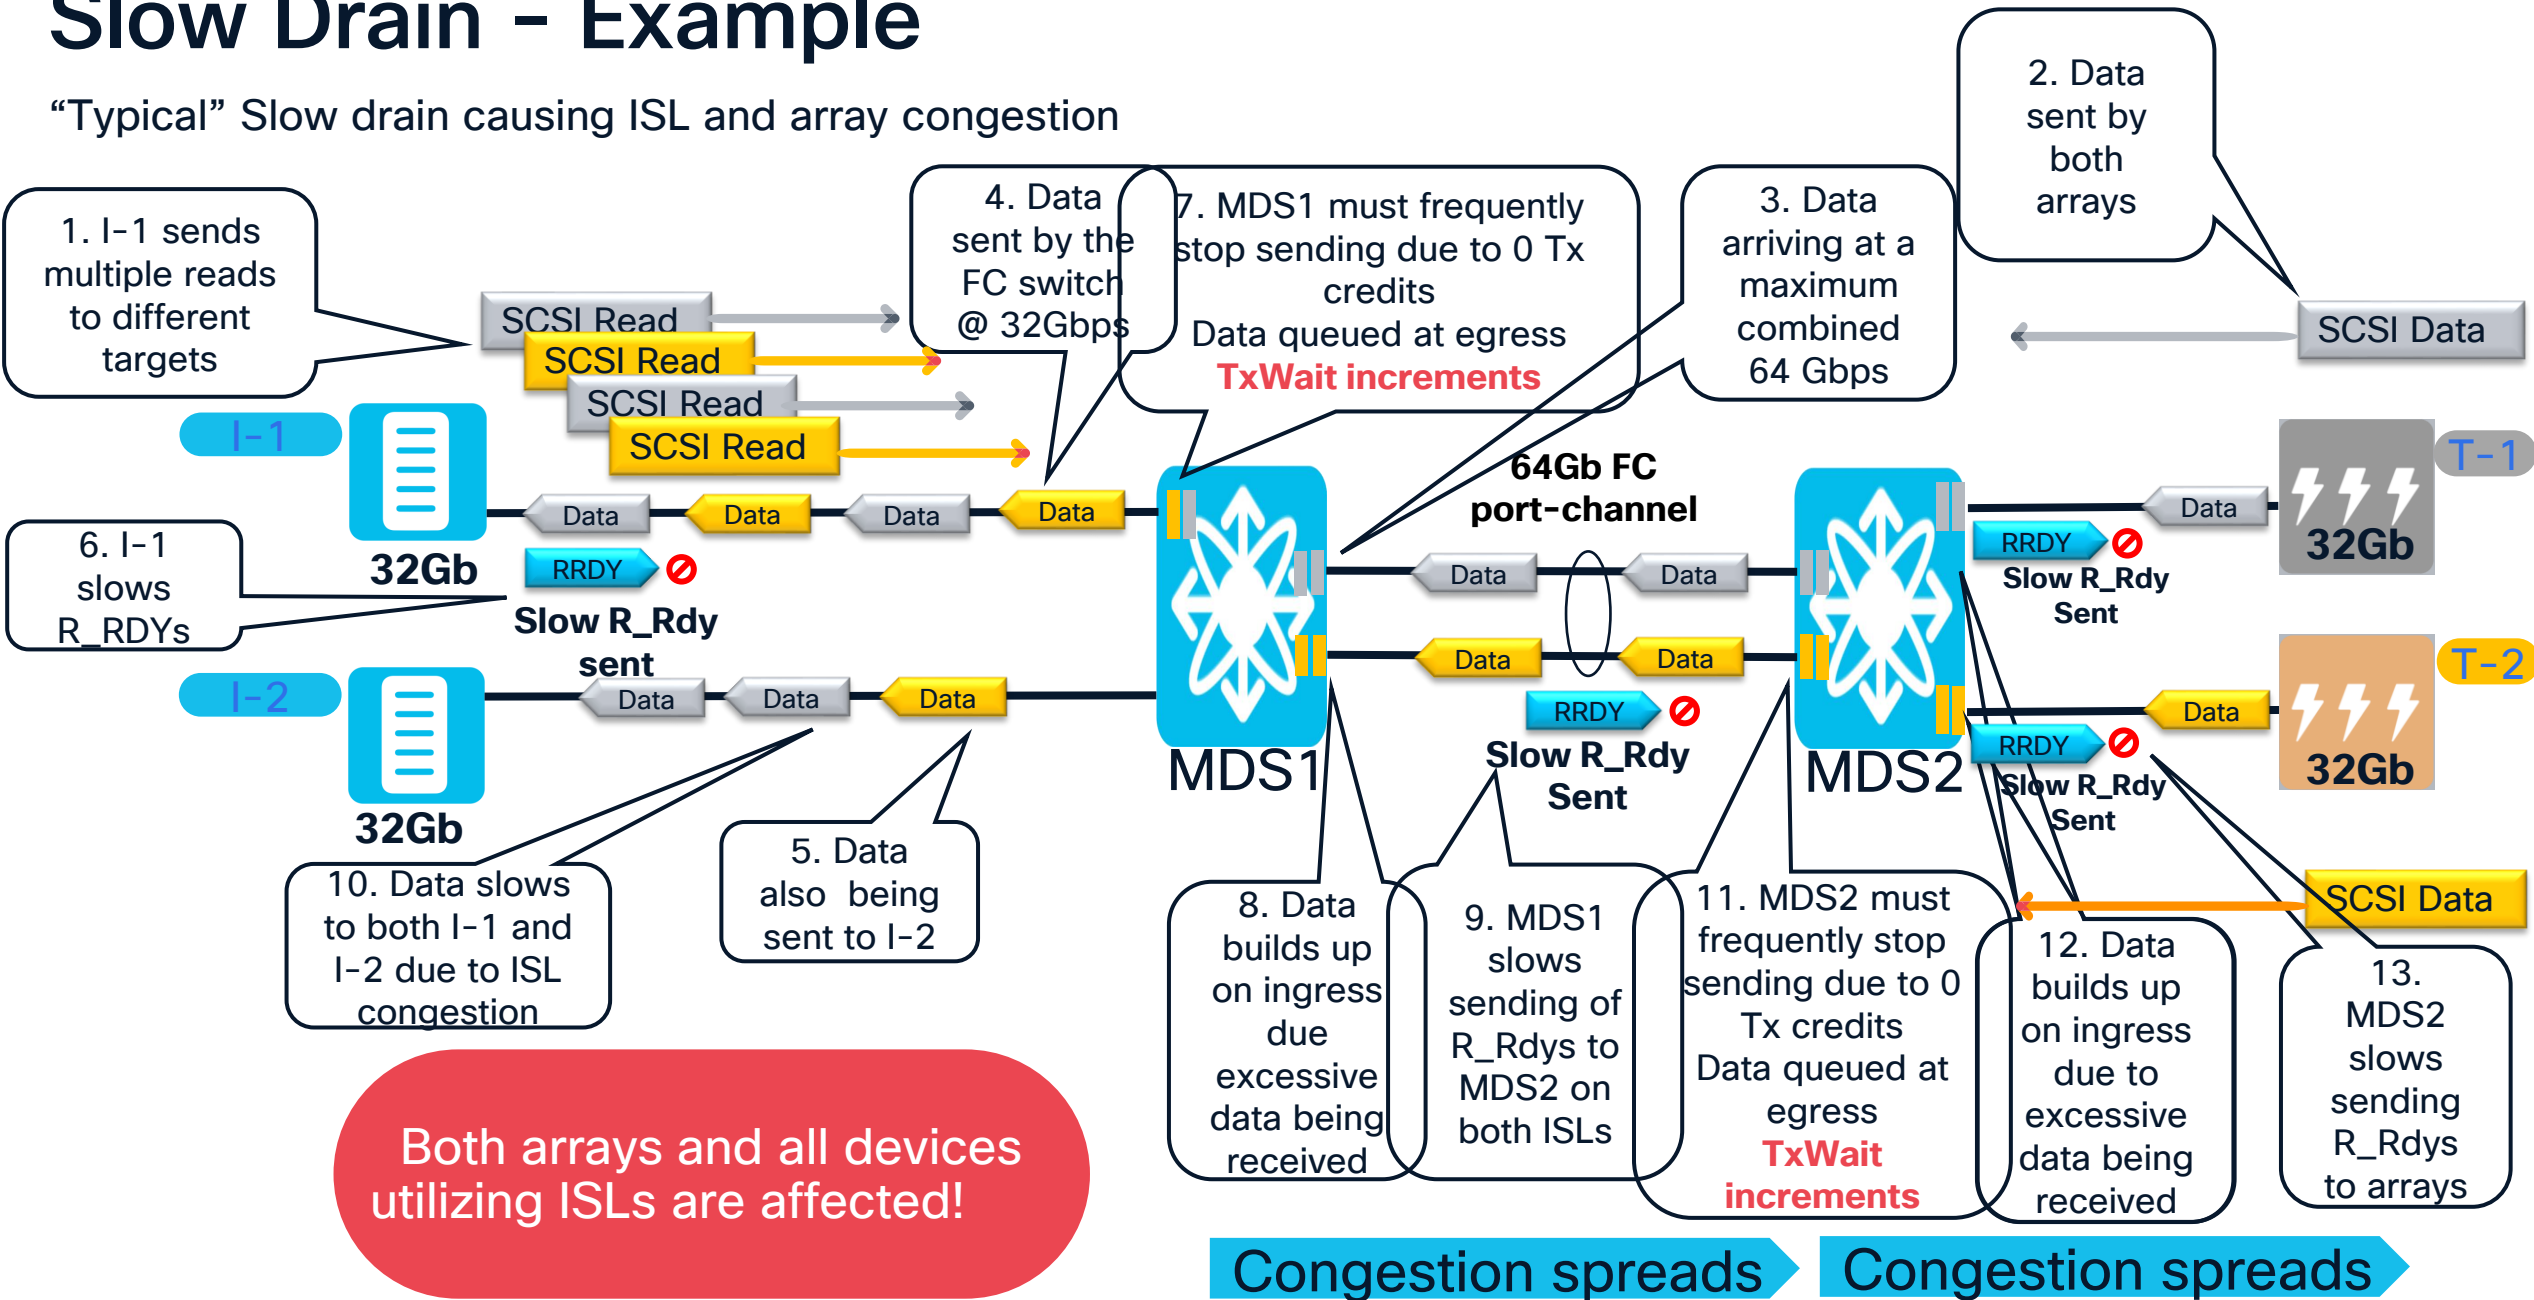
1. Slow Drain – Receiver purposely slowing down traffic by withholding R_RDYs
2. Over-Utilization – Receiver requesting more data than can be transmitted
3. Insufficient B2B credits for the link's distance(latency), speed and frame size
4. B2B credits lost due to bit errors or Invalid Transmission Words(ITW)

#1, #2 are the focus of this presentation

#4, See BRKDCN-3641 from CL2024 On Demand Library

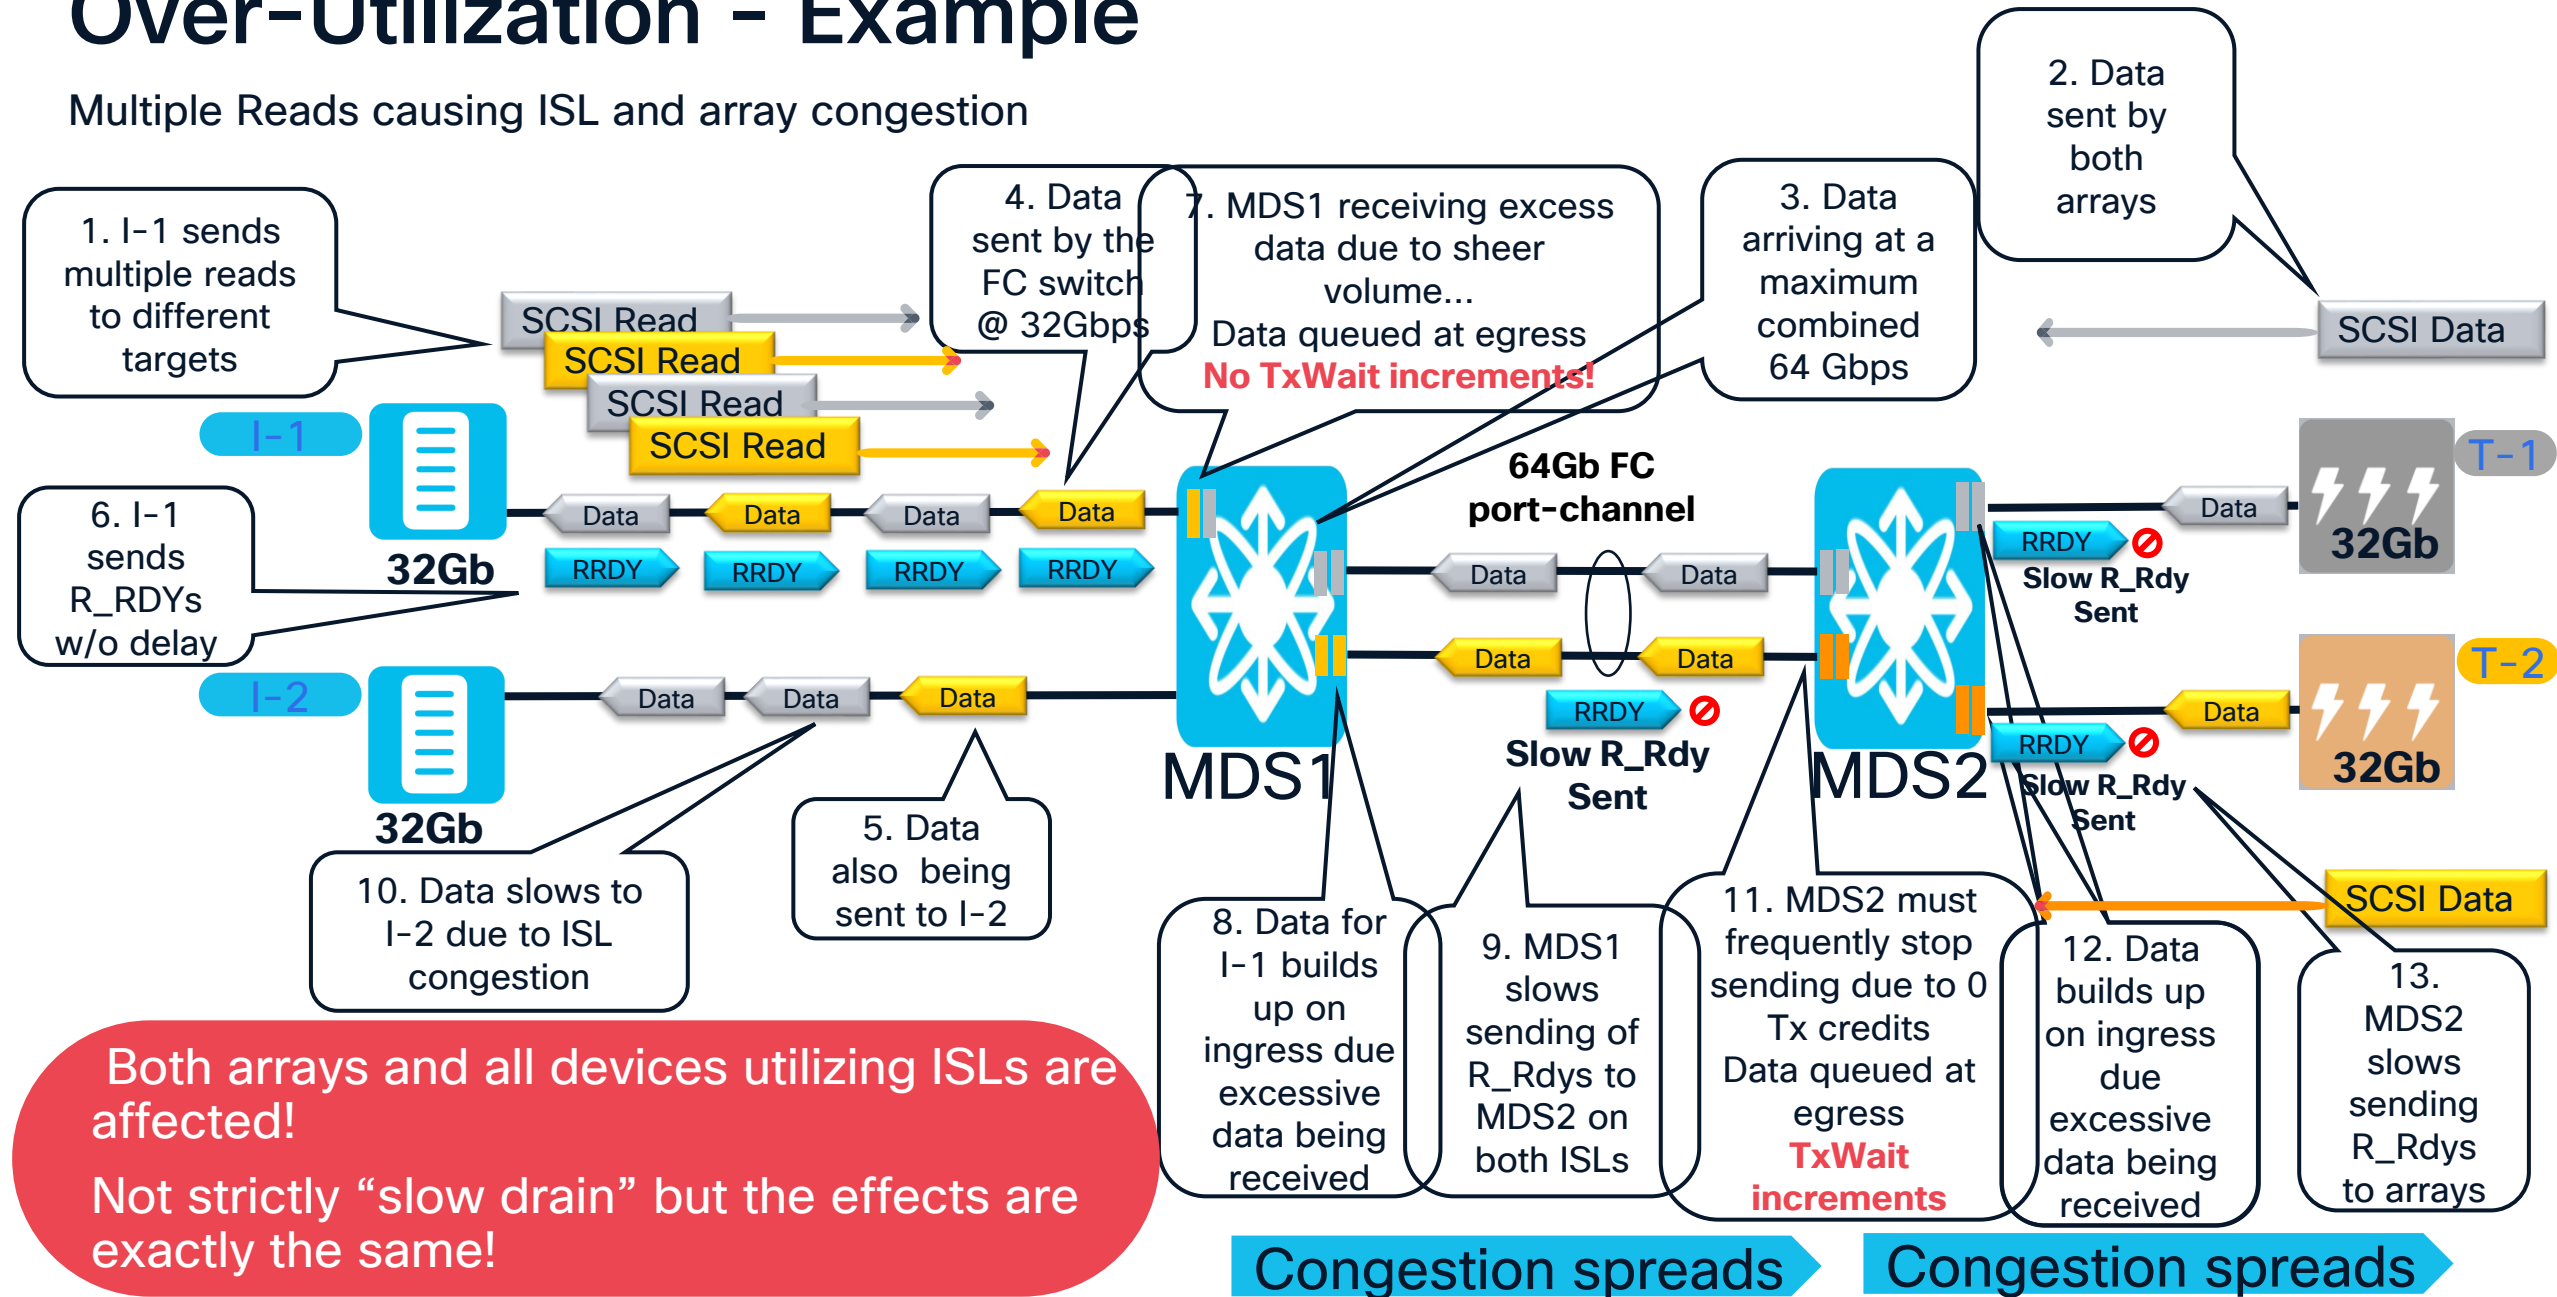
Slow Drain - Example

“Typical” Slow drain causing ISL and array congestion

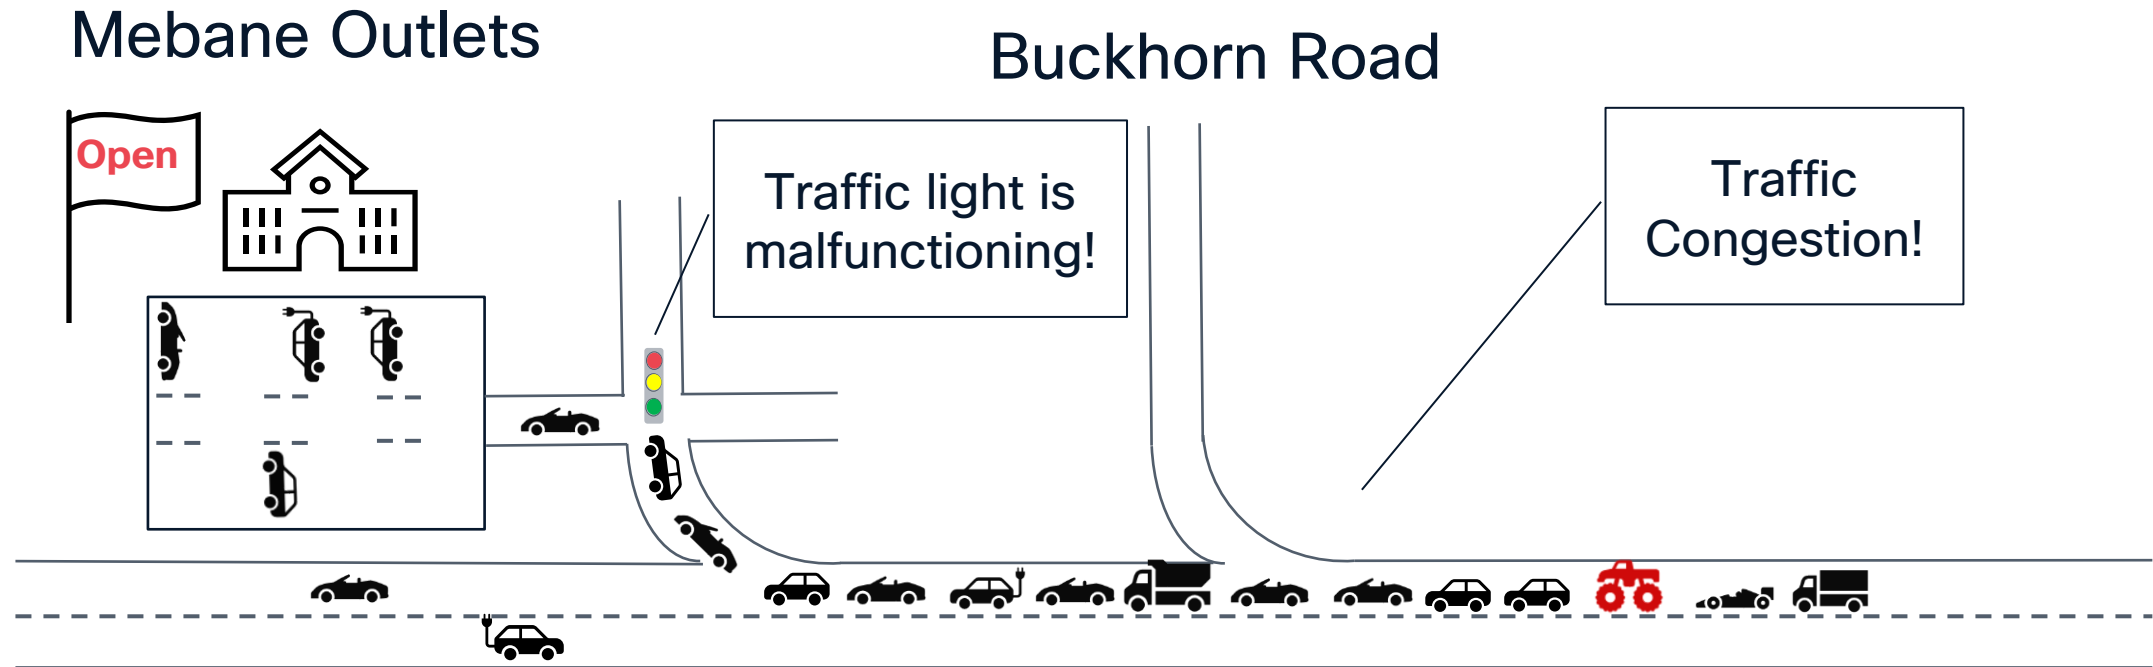


Over-Utilization - Example

Multiple Reads causing ISL and array congestion




Congestion Road Analogy – Slow Drain



Slow Drain: A number of vehicles, within the design limits of the exit, trying to go to the Mebane Outlets exit. Stoplight malfunctions causing congestion. **Red truck** cannot get to the Buckhorn Road exit.

What's the 'fix' here? Fix traffic light so it functions normally.

Congestion Road Analogy - Over-Utilization



Over-utilization: More vehicles trying to go to the Mebane Outlets than the exit was designed to handle. This causes congestion. **Red truck** cannot get to the Buckhorn Road exit.

What's the 'fix' here? More exit lanes, change stoplight timing, increase exit capacity, etc.

Comparison of Slow Drain vs. Over-Utilization

Slow Drain


Tx B2B credit starvation

I-1 is busy

I-1 slows down its ingress traffic rate by slowing down sending of R_RDY to MDS1

Tx Utilization % (tx-datarate) 

Tx Congestion % (txwait) 



Over-utilization


Receive data rate on ISL port is faster than the host port speed

I-1 is receiving at full capacity

MDS1 slows down traffic rate to I-1 by slowing down sending of R_RDY to MDS2

Tx Utilization % (tx-datarate) 

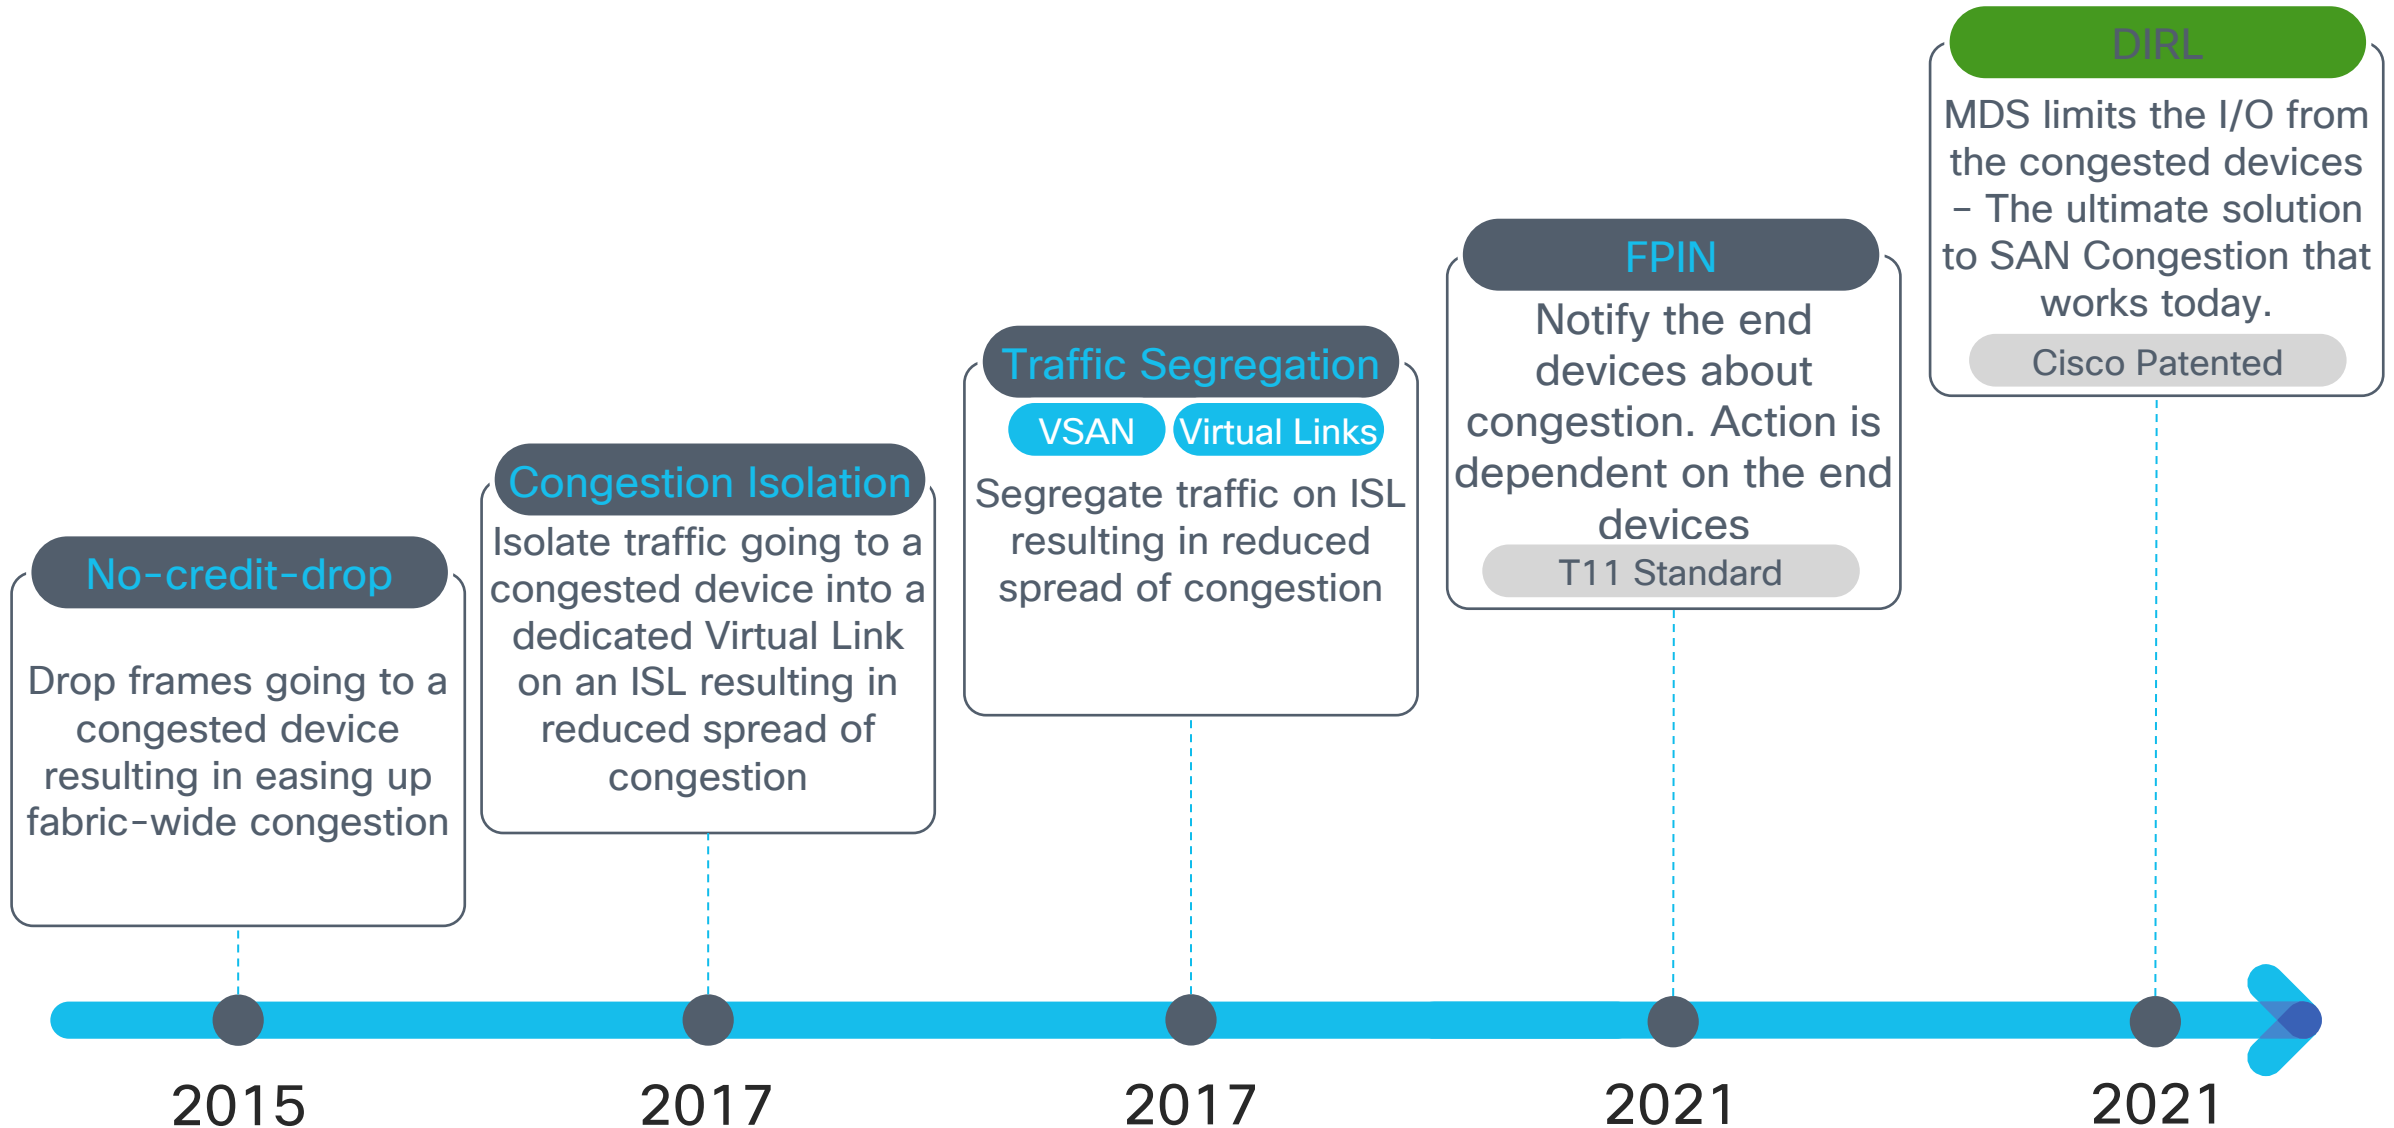
Tx Congestion % (txwait) 



Frames are not dropped in FC fabric. Rather, they consume switch buffers causing a fabric-wide congestion spreading

SAN Congestion Prevention

SAN Congestion Innovation on Cisco MDS



The Solution

8.5(1)

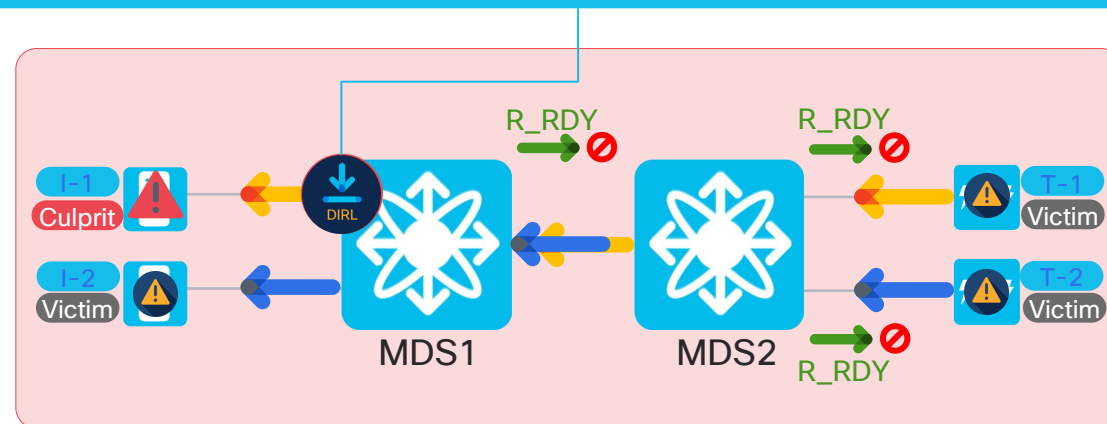
Cisco Dynamic Ingress Rate Limiting



I-1 is asking for more than it can ingest

DIRL limits I-1's asking rate to reduce its receiving rate

DIRL dynamically changes I-1's asking rate to adapt to its traffic profile



The Solution


Cisco Dynamic Ingress Rate Limiting



I-1 is asking for more than it can ingest

DIRL limits I-1's asking rate to reduce its receiving rate

DIRL dynamically changes I-1's asking rate to adapt to its traffic profile



No more victims!

DIRL prevents SAN Congestion due to slow-drain and over-utilization.

Cisco Dynamic Ingress Rate Limiting



End-device independent

Upgrading of end-devices is not needed



Adaptive

DIRL dynamically adjusts as per the traffic profile of the host



No side effects

Rate limits congested hosts only. Other non-congested hosts and storage ports are not impacted



Easy adoption

DIRL is available on MDS switches after a software-only upgrade.



Gradual Rollout

DIRL can be implemented one switch at a time



Affordable


No additional license needed




Topology independent

DIRL works in edge-core, edge-core-edge, or collapsed core (single switch fabric) topologies

Without Cisco DIRL




With Cisco DIRL



No more victims!

Back pressure alleviated!

Notifications and Congestion Signals in FC



Available on Cisco MDS in NX-OS 8.5(1) onwards

- Exchange Diagnostic Capabilities (EDC) (for Congestion Signals)
- Register Diagnostic Functions (RDF) (for FPINs)
- Congestion Signals (Primitives)
- Fabric Performance Impact Notifications (FPIN)

DIRL vs FPIN

DIRL helps today. FPIN readiness will take a few years.

DIRL is available on existing MDS switch after a software-only upgrade, without any dependency on end devices

Although FPIN is supported on MDS switches, action is dependent on the end devices

DIRL is affordable

DIRL and FPIN work on existing MDS switches and don't need an additional license

Must upgrade end-devices to benefit from FPIN

In the future, when you are ready for FPIN, DIRL will can be a complementary technology

What if a few devices don't react to FPIN and still cause congestion? DIRL within MDS switches will be the protection



SAN Congestion Management - Recommendations

Reactive

- Gather 'show tech-support slowdrain' from **all** switches
- Use OBFland and other commands to identify culprit and victims
- **TAC can help!**

Proactive

- Schedule NDFC/DCNM Congestion Analysis to run daily for 24 hours.
- Important for troubleshooting
- Configure MDS port-monitor (PMon) for automated alerts and actions.
- Important for congestion prevention using DURL.

Predictive

- Enable SAN Analytics and SAN Insights for getting visibility into application I/O traffic patterns.
- Important for finding the underlying root cause and predicting congestion

*Slow-drain Analysis is renamed to Congestion Analysis in NDFC 12.1(1e)


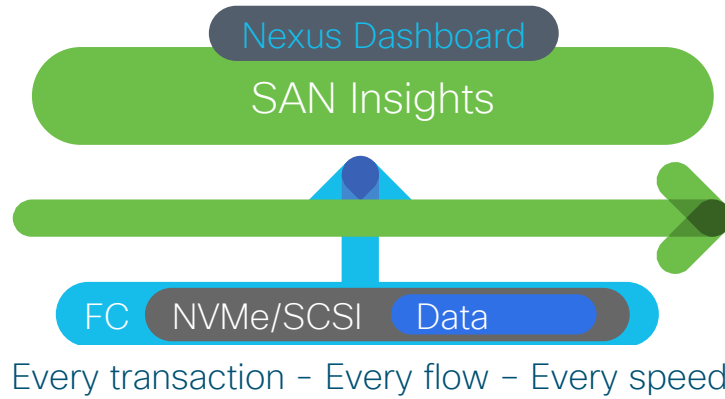


Part 2

SAN Analytics and NDFC SAN Insights

NVMe and SCSI I/O Visibility Using SAN Analytics

Is the app slowing down due to storage access issues?




Use-cases

<p>Anomaly Detection</p>	<p>Application Performance Insights</p>	<p>Storage provisioning recommendation</p>
<p>Auditability</p>	<p>Infrastructure Optimization</p>	<p>Storage Performance Insights</p>
<p>Change Management</p>	<p>Health Report</p>	<p>Faster Troubleshooting</p>

Solution Components

‘Cisco SAN Analytics’ is the umbrella name for the overall solution



‘SAN Analytics’ is also the name of the feature to enable flow metric collection on MDS switches (NX-OS command: `feature analytics`)

‘SAN Telemetry Streaming’ is an efficient mechanism to export metrics from MDS switches (NX-OS command: `feature telemetry`)

‘SAN Insights’ is an analytics and visualization engine within NDFC/DCNM

Cisco 32G SAN Analytics – Architecture

Traffic Inspection

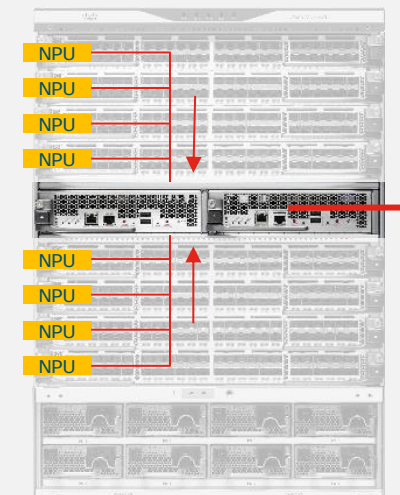
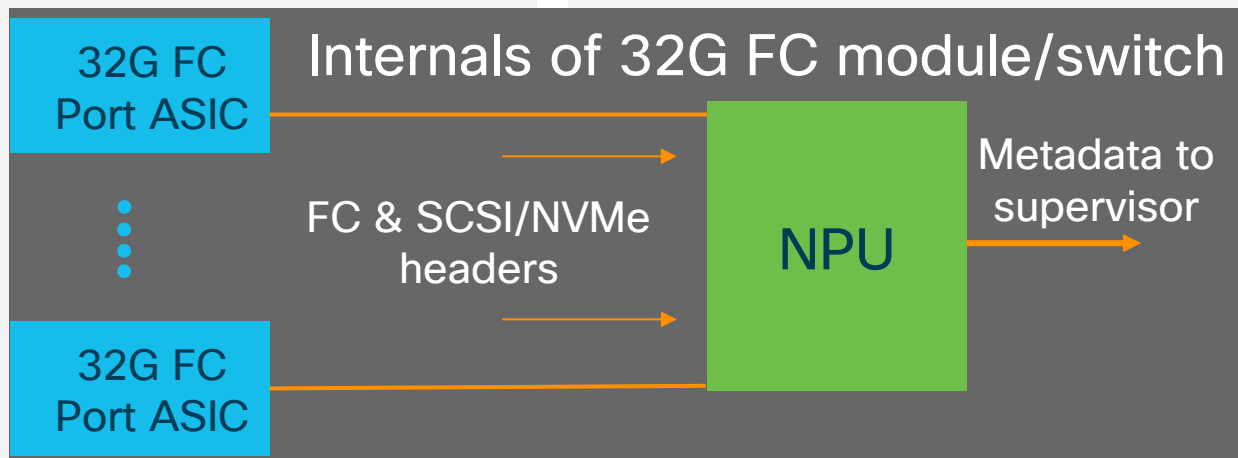
- Inbuilt tap in 32G FC port-ASIC
- Traffic inspection capability on all ports
- Zero impact to switching functionality
- Inspects only FC & SCSI/NVMe headers, not data

Metric Calculation

- Network Processing Unit (NPU) on 32G FC products
- Receives headers of specific frames from port-ASIC
- Extracts metrics from headers
- Stores metrics in multiple views

Metric Export

- SAN Telemetry Streaming (STS) exports flow metrics to external receivers
- Extremely efficient mechanism
- Works using existing mgmt. port
- On-switch CLI and remote RESTful access available



External Receivers:
DCNM SAN Insights
Or Virtual Instruments or
any other 3rd party app

9.2(2)

2022

Cisco SAN Analytics

using
Cisco MDS 64G Switches



2017

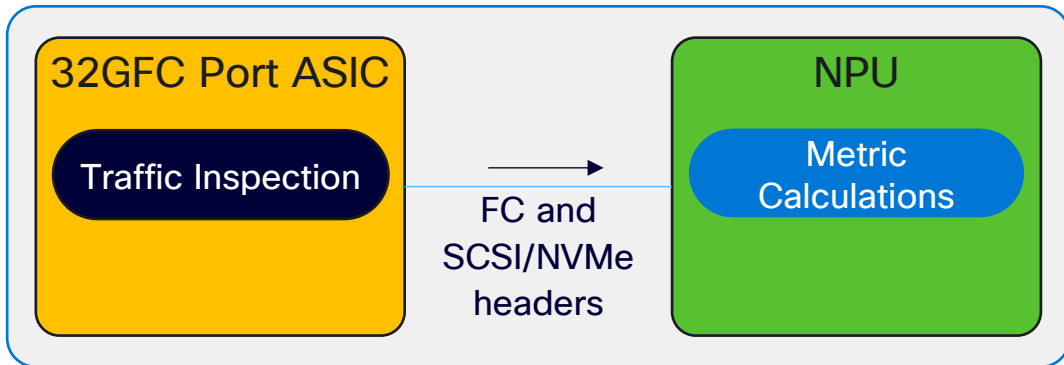
Cisco SAN Analytics

using
Cisco MDS 32G Switches


Cisco SAN Analytics Architecture

Cisco SAN Analytics Architecture

Using
Cisco MDS 32G switches



Using
Cisco MDS 64G switches



Cisco SAN Analytics

Using
Cisco MDS 64G switches



Analytics for Billions of IOPS

Traffic inspection and metric calculation in ASIC



Software Programmability

On-board Network Processing Unit



Additional flow metrics

Host Response Latency, First Burst, Optimized Read




Investment Protection

1GbE streaming port on the 64GFC module

SAN Analytics Deployment

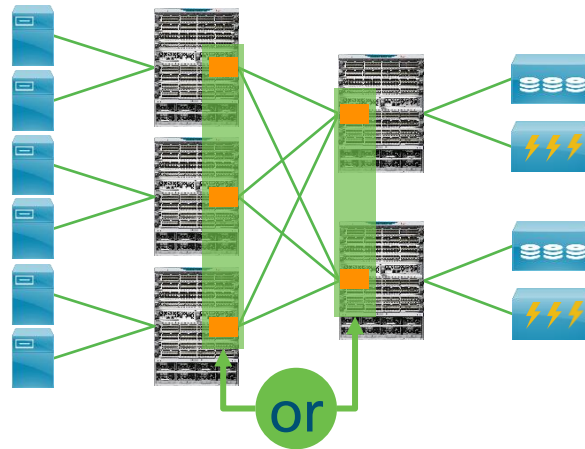
Deployment Models

Storage Ports




Closest to storage

ISL Ports



High capacity 64G ISLs

Host Ports



Closest to apps


- Inspection of traffic at least once in the end-to-end data path is enough
- Avoid double inspection of traffic
- Design for uniform utilization of the NPU

Most customers are enabling analytics on **storage ports**

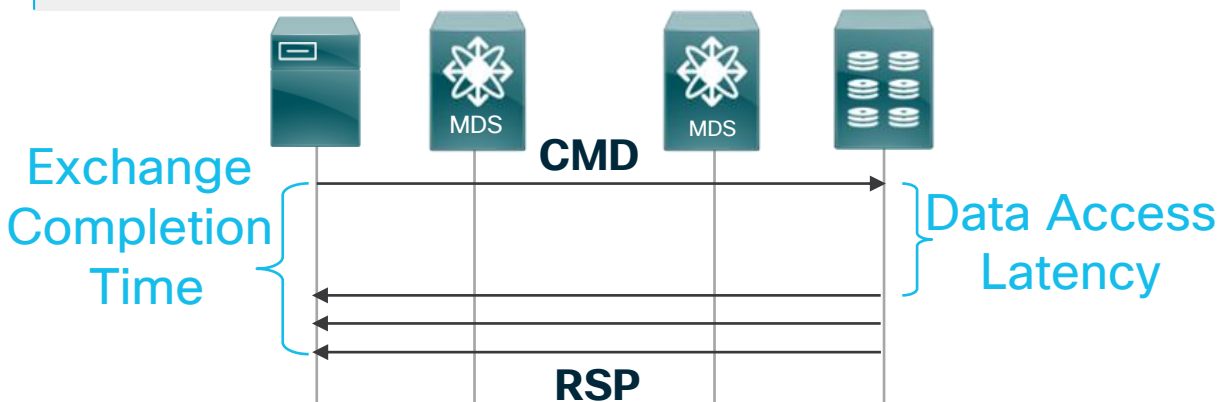
IO Flow Metrics

80+ Metrics per SCSI or NVMe Flow

Write CMD



Read CMD



Exchange Completion time

Data Access Latency

Outstanding IO

IOPS

I/O Size

Failed exchanges, IO retransmissions

Error conditions (Aborts, Rejects, etc.)

Check conditions, queue-full, etc.

Measured at Initiator-Target-LUN/Namespace (ITL or ITN) level

Metric Export

On-switch CLI

Flexible SQL-like CLIs

- Pulls raw data
- Output in JSON format (key-value, similar to dictionaries)
- Basic trending & correlation

Programmable ShowAnalytics

- Output in nice-to-read tabular format, similar to any other NX-OS show command output

NX-API

- Capability to query switch remotely
- Response in JSON format
- Response depends on the SQL-like CLI input
- Works best to track specific metrics
- Extremely flexible

Streaming Telemetry

- MDS streams out the metrics regularly
- Data transport and encoding in industry leading open formats
- Optimized performance for continuous data export at high frequency

SAN Insights

3rd party apps

Metric Export

On-switch CLI

Flexible SQL-like CLIs

- Use for low-level troubleshooting in real-time
- Required detailed understanding of metrics, their format and units

Programmable ShowAnalytics

- Use for low-level troubleshooting in real-time
- Common use-cases – already available
- Special use-cases – Write your own

NX-API

- Use for occasional access of selective metrics only. Example: a custom script which is invoked occasionally.
- Requires a simple http(s) call to the switch (NX-API)
- No need to setup a gRPC receiver
- Performance may be an issue at scale. Use streaming telemetry with large number of flows

SAN Insights

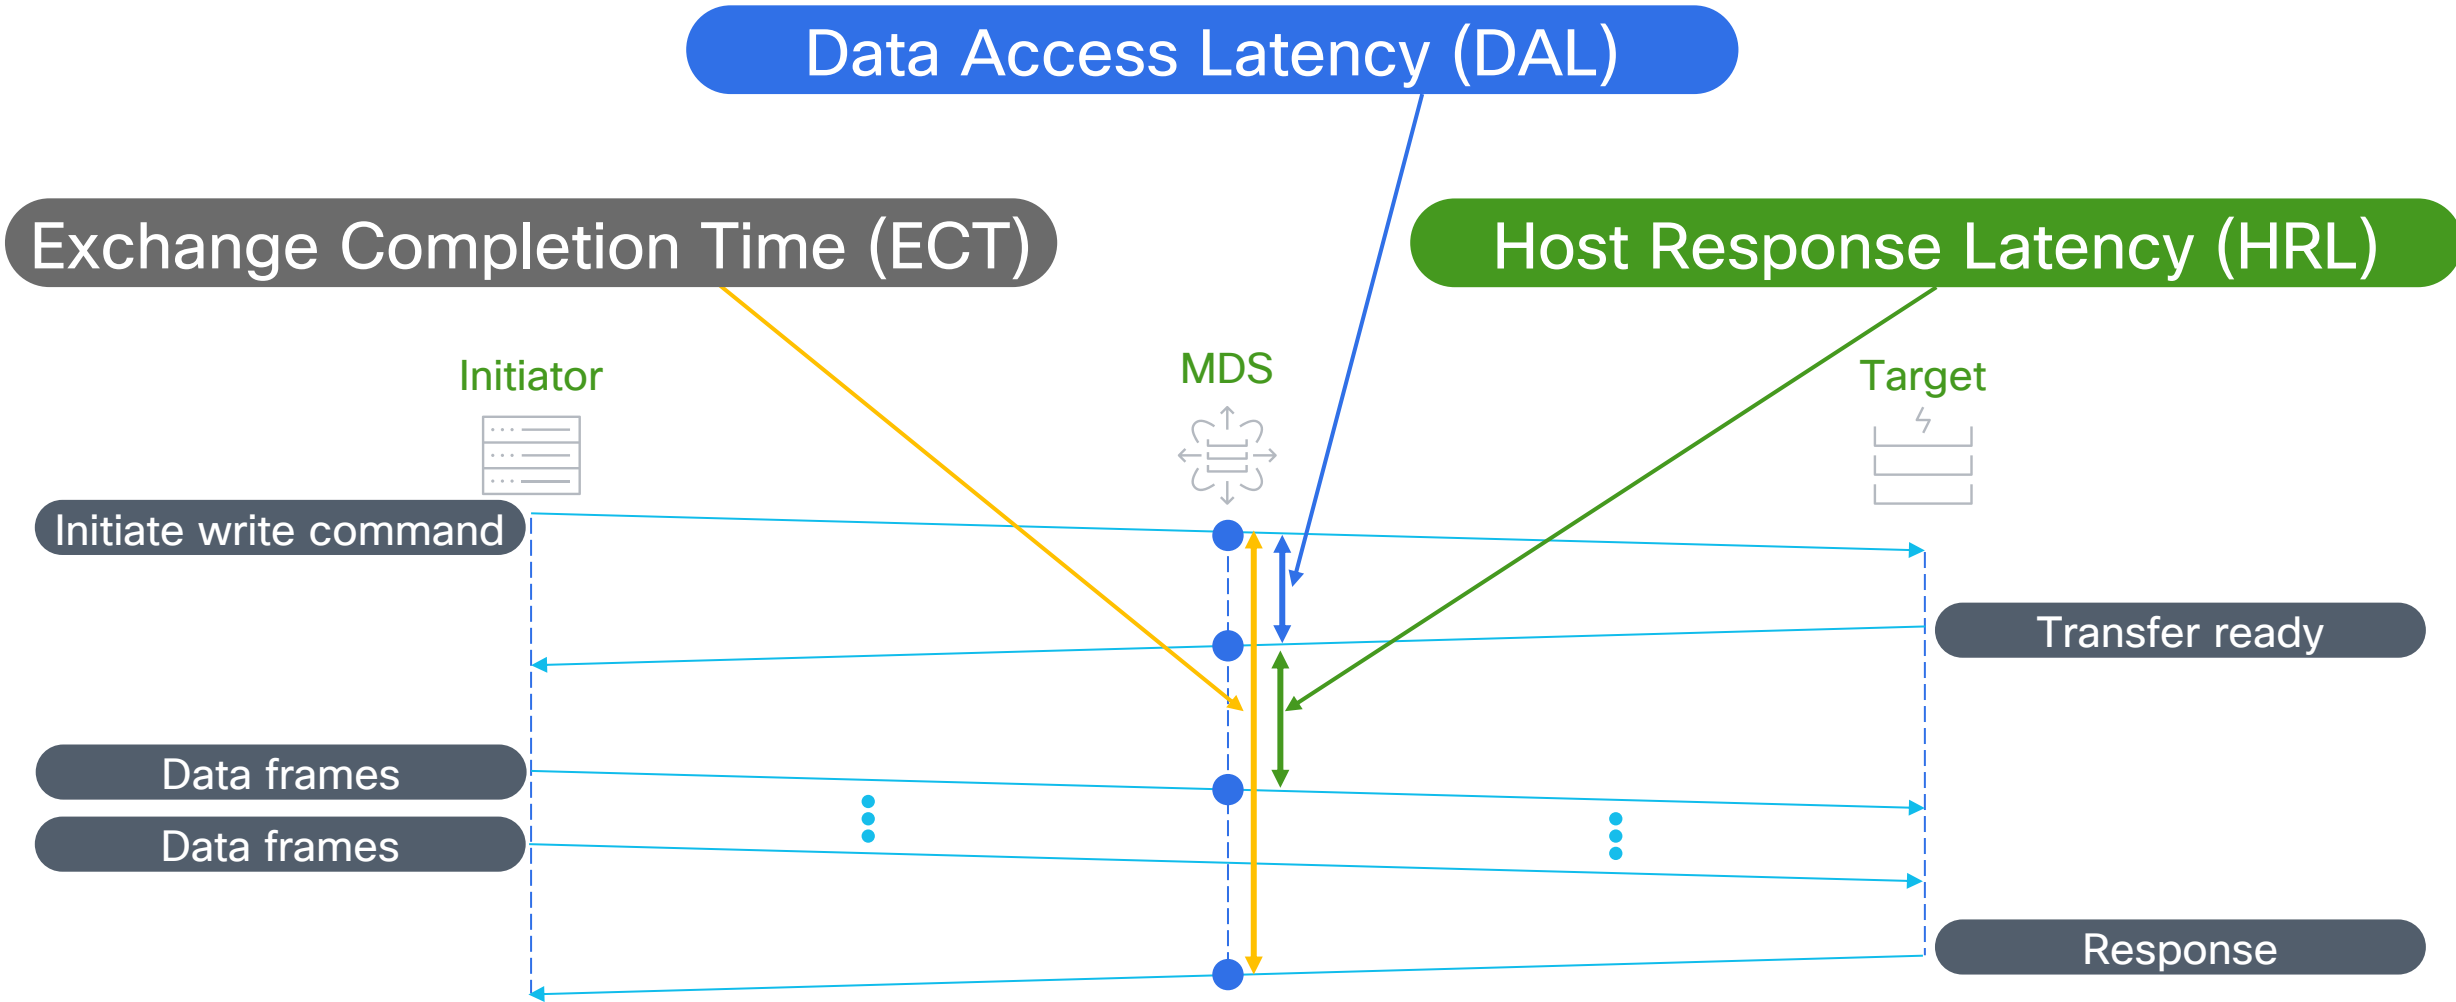
3rd party apps

Streaming Telemetry

- Use for continuous export of metrics. Example: A continuous monitoring app
- Requires a special receiver to understand gRPC.

Cisco SAN Analytics Flow Metrics

Write I/O Operation



HRL is available only on 64G MDS switches. ECT and DAL are available on 32G MDS as well.

Cisco SAN Analytics Flow Metrics

Write I/O Operation

New with Cisco MDS 64G switches - NX-OS 9.2(2)

Host Response Latency (HRL)

Exchange Completion Time (ECT)

Data Access Latency (DAL)


Higher ECT indicates degraded storage performance

Higher HRL indicates a slow host


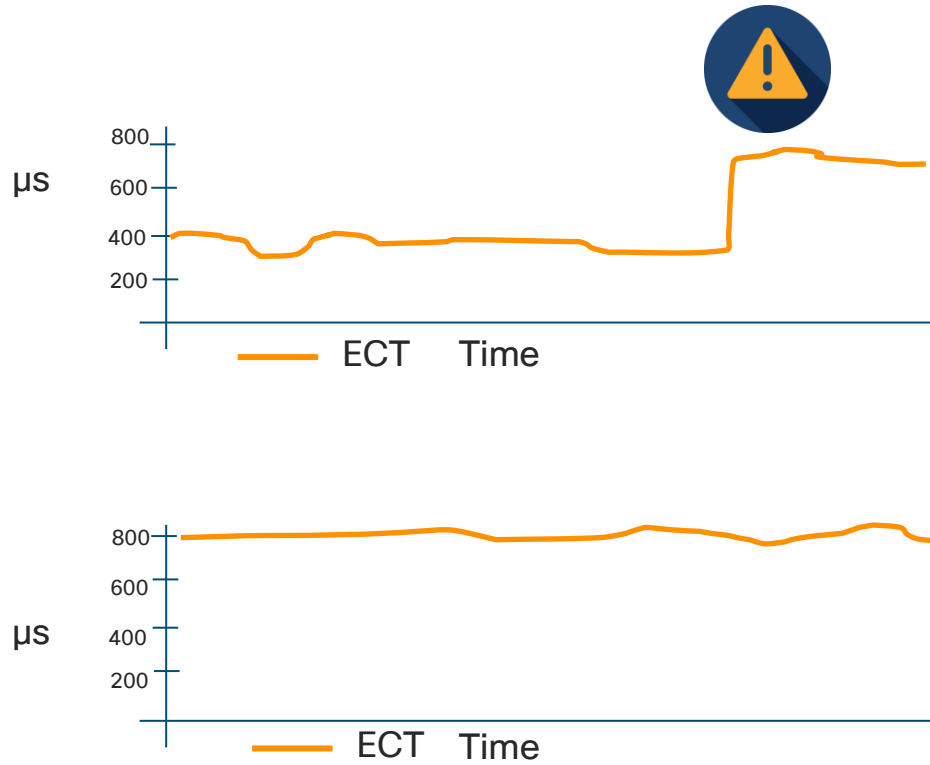
Higher DAL indicates a slow storage array

Diagnosing Storage I/O Performance Issues

Which server is performing better ?



Diagnosing Storage I/O Performance Issues



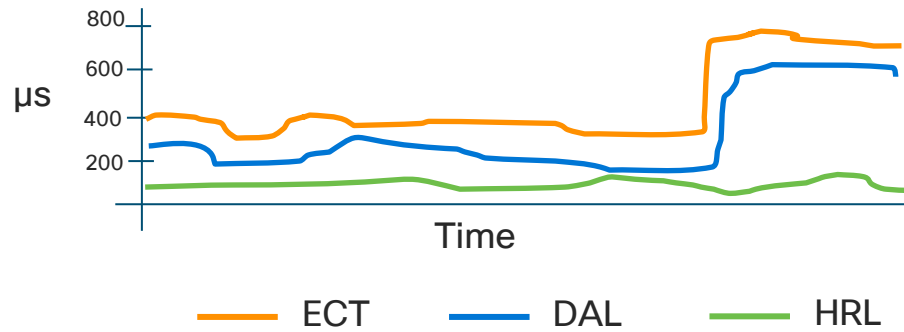
Which server is performing better ?

Diagnosing Storage I/O Performance Issues


Where is the bottleneck



Storage performance is degraded due to delay caused by **storage array**



Storage performance is degraded due to delay caused by **host**

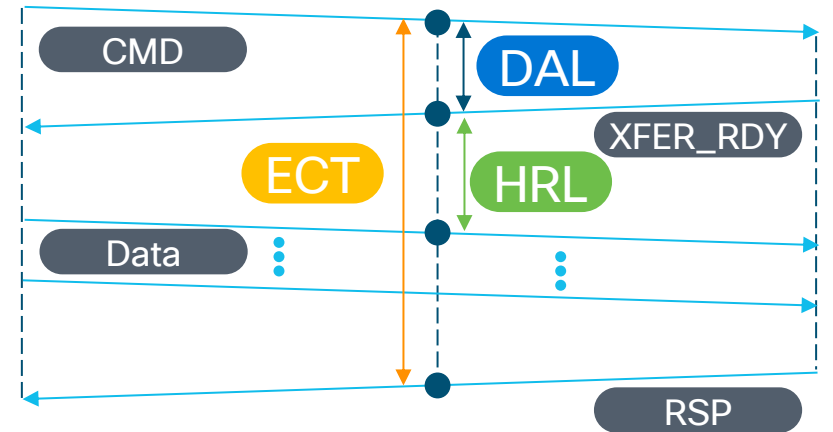



Write transaction

Initiator

MDS

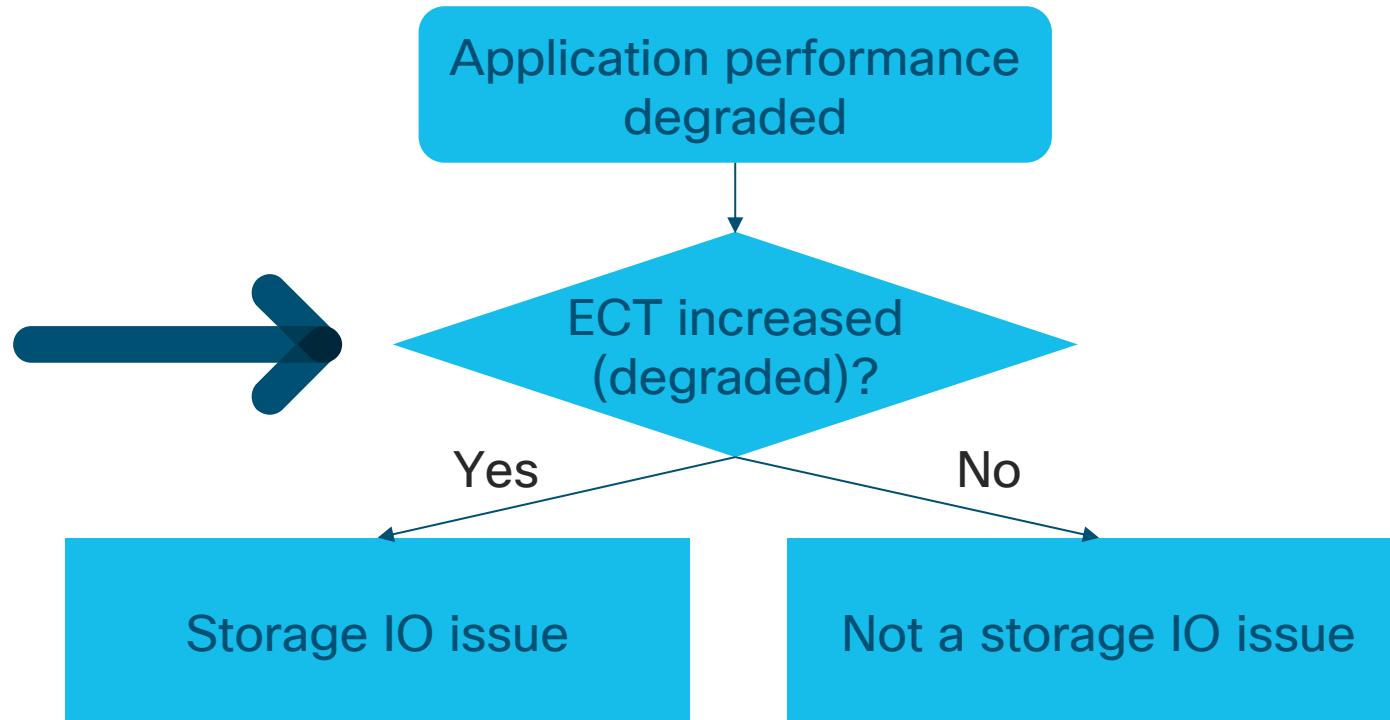
Target



Using ECT, DAL, and HRL for diagnosing delays

1st level diagnosis

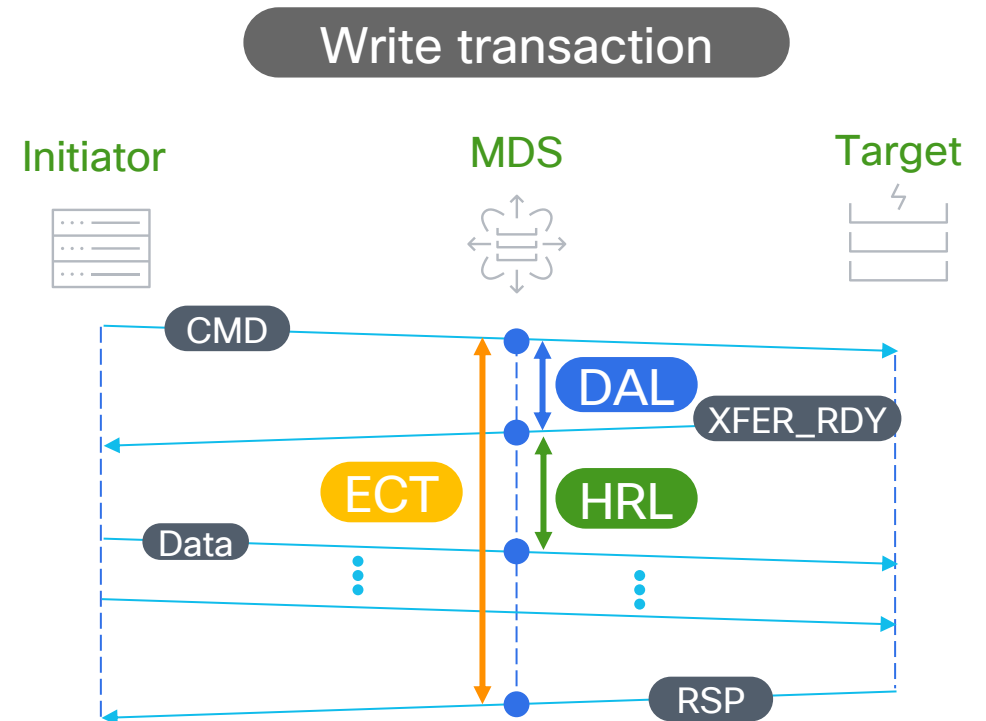
- Increase in ECT may directly lead to application slowdown and is the first level of pin-pointing towards storage IO issue



Using ECT, DAL, and HRL for diagnosing delays

2nd level diagnosis

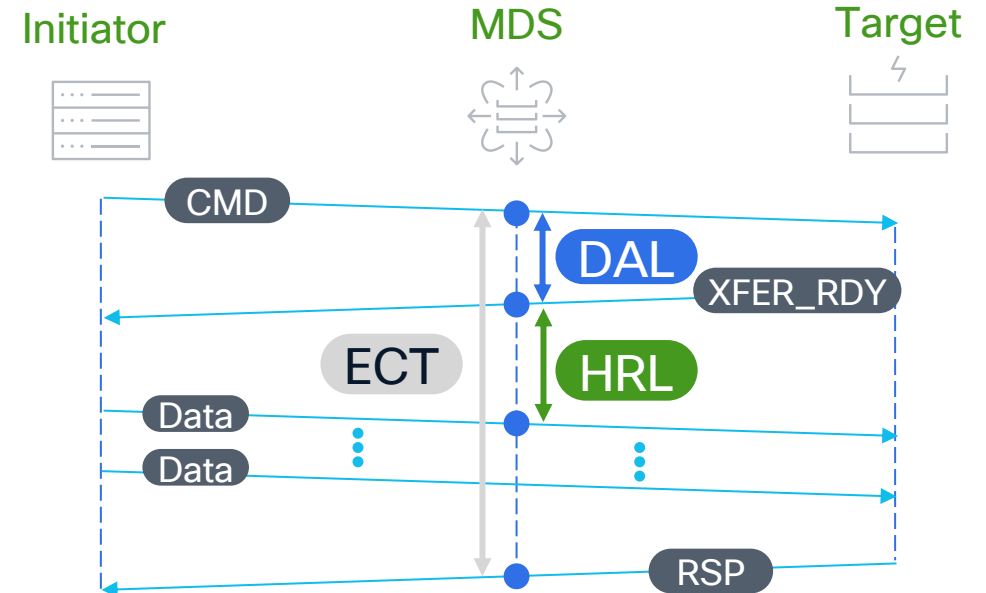
- ECT may increase (degrade) due to
 - Internal delay within storage array
 - Delay in the fabric (SAN Congestion)
 - Internal delay within host




Using ECT, DAL, and HRL for diagnosing delays

2nd level diagnosis

Write transaction





Yes




Yes

No



- Delay in the fabric (SAN Congestion)
- Internal delay within host (HRL)


Case study 1 – A Large European Bank



1st level diagnosis


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

Case study 1 – A 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 1 - A Large European Bank




3rd level diagnosis

- 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
- No fabric congestion detected
- Delay within host → Resulted in detection of an unpatched Oracle app on host

Case Study 2 - Congestion due to Over-utilization

Find the culprit VM




Problem statement

- Host-1 is virtualized. It uses LUN/Namespaces/volumes from many storage arrays that are connected via SAN.
- Host-1 is the cause of congestion due to over-utilization
- Goal – Which VM and volumes are the top contributors to high link utilization?

Case Study 2 - Congestion due to Over-utilization


Case Study - Many customers use this approach today



Background

- Network link utilization depends on I/O throughput
- Find I/O throughput using SAN Analytics at flow granularity
 - VM-I-T-L: If VE ID is supported (VE = Virtual Entity (Container or VM))
 - ITL: If VE ID is not supported (Most deployments)
- Traffic towards Hosts (Initiators) is mostly read I/O throughput, whereas traffic towards storage arrays (targets) is mostly write I/O throughput

Top 10 Host by Avg Read Throughput for Enclosure



Graph Table

Filter by attributes Download

Initiator Enc	Initiator	Target Enc	Target	LUN	Switch IP Address	Port	Timestamp	Read Throughput (MB/s)	Write Throughput (MB/s)
STG-TME-ESXI-B230-4	B230-4-A1F	A300	A300-SVM_001_0F	0018-0000-0000-0000	172.22.163.20	fc1/35	2023-06-6 10:00:00	3.2164	3.2198
STG-TME-ESXI-B230-4	B230-4-A1F	A300	A300-SVM_001_01	001c-0000-0000-0000	172.22.163.20	fc1/33	2023-06-6 10:00:00	3.215	3.2337
STG-TME-ESXI-B230-4	B230-4-A0F	A300	A300-SVM_001_0F	001d-0000-0000-0000	172.22.163.20	fc1/35	2023-06-6 09:55:00	0.0767	0.0005
STG-TME-ESXI-B230-4	B230-4-A1F	A300	A300-SVM_001_0F	001d-0000-0000-0000	172.22.163.20	fc1/35	2023-06-6 10:00:00	0.0766	0.0005
STG-TME-ESXI-B230-4	B230-4-A1F	A300	A300-SVM_001_06	0006-0000-0000-0000	172.22.163.20	fc1/34	2023-06-6 09:55:00	0.0752	0.0005
STG-TME-ESXI-B230-4	B230-4-A1F	A300	A300-SVM_001_06	0004-0000-0000-0000	172.22.163.20	fc1/34	2023-06-6 09:55:00	0.0745	0.0005

100 Rows Page 1 of 1 << < 1-100 of 100 >>

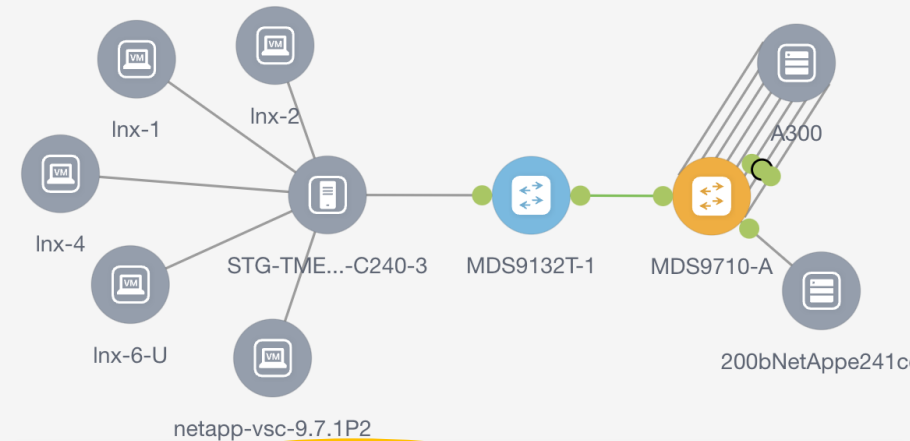
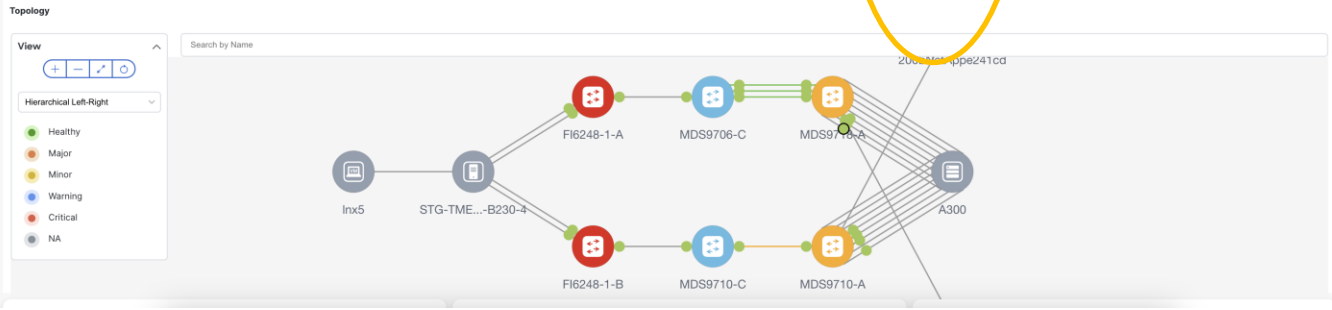
Host Enclosure - STG-TME-ESXI-B230-4

Viewing SCSI Metrics Showing Data from 6/6/2023, 9:59:15 AM

Initiator Target Pairs < 1% 1 - 15% 15 - 30% > 30% Total Read IOPS 21301 Total Write IOPS 20902 Total Read Throughput 103.8480 Total Write Throughput 82.6525

Filter by attributes

Source Alias	SID	Destination Alias	DID	Fabric	ECT (% dev)		IOPS		Throughput (MBps)		ECT (ms/IO)		DAL (ms/IO)	
					Read Avg.	Write Avg.	Read Total	Write Total	Read Total	Write Total	Read Avg.	Write Avg.	Read Avg	Write Avg
B230-4-A0F	32006c	A300-SVM_001_01	50021	Fab-A			2656	2610	12.8351	10.3291	0.2536	0.5453	0.2424	0.0355
B230-4-A0F	32006c	A300-SVM_001_0F	50101	Fab-A			2665	2614	13.0261	10.3341	0.2061	0.5443	0.1951	0.0354
B230-4-A0F	32006c	A300-SVM_001_06	50041	Fab-A			2670	2617	13.1422	10.3361	0.2312	0.5613	0.2205	0.0369
B230-4-A0F	32006c	A300-SVM_001_10	50121	Fab-A			2664	2611	12.9929	10.3252	0.2230	0.5533	0.2120	0.0355
B230-4-A1F	32006d	A300-SVM_001_06	50041	Fab-A			2666	2614	13.0900	10.3320	0.2359	0.5507	0.2250	0.0358
B230-4-A1F	32006d	A300-SVM_001_10	50121	Fab-A			2661	2613	12.9000	10.3314	0.2207	0.5456	0.2100	0.0364
B230-4-A1F	32006d	A300-SVM_001_01	50021	Fab-A			2658	2611	12.9057	10.3317	0.2156	0.5531	0.2045	0.0358



VM Name	VM IP
Inx-2	172.22.163.205
Inx-1	172.22.163.217

Predicting SAN Congestion

Case Study – A trading company selectively upgraded SAN in phases using predictive capabilities of SAN Insights

Problem statement

- Large SAN environment with thousands of ports per fabric. Many fabrics.
- Upgraded to all-flash storage, without upgrading the host speeds at the same time
- Observed increased occurrences of congestion due to over-utilization of host links
- Aware that the ultimate solution was to upgrade end-to-end connectivity
- But resources weren't enough for an overnight upgrade
- Goal – Among thousands, which server to upgrade first?

Predicting SAN Congestion

Case Study – A trading company selectively upgraded SAN in phases using predictive capabilities of SAN Insights

Solution

- The trading company enabled SAN Analytics on storage ports
- Collected the peak and average read and write I/O sizes for all hosts
 - Peaks are important.
- Made a sorted list and started upgrading the hosts first that **have larger read I/O size**
- The data collected by SAN Analytics gave them predictive insights for an informed upgrade plan.
 - Without SAN Analytics they would have run into many more congestion issues

Additional Relevant Sessions

Storage Area Networking



- **Cisco Live 2025 San-Diego**

BRKDCN-2941 - Advanced Storage Area Network Design - Tuesday June 10 4:00 PM

BRKCOM-2657 - Understanding, Identifying and Resolving Congestion in a UCS
Converged Infrastructure - Wednesday June 11th 1:30 PM

- **On Demand Library**

BRKDCN-2648 - What's New in NX-OS 8.x and 9.x for MDS and Nexus 9000 for SAN

BRKDCN-3002 - Dynamic Ingress Rate Limiting as a Real Solution to SAN Congestion

BRKDCN-3645 - SAN Insights-Real-time and Always-On NVMe Visibility at Scale

BRKDCN-3641 - Manage, Operate, and Optimize High-Performance Storage Area
Networks



Cisco Tech Day
Denver

Questions?

Thank you to our sponsors!



World Wide Technology, Inc.

PRESIDIO®



