

STORAGE DEVELOPER CONFERENCE



Fremont, CA
September 12-15, 2022

BY Developers FOR Developers

A **SNIA** Event

The Quest for an Autonomous Storage Fabric

Fibre Channel Fabric Notifications Technology

Shyam Sundar and Nishant Lodha, Marvell

Agenda



Storage fabrics
Getting more intelligent?



Self driving SANs
Under the hood

Future of storage fabrics

Leading the transition from one **heterogeneous** world to another



Fibre Channel Connectivity
Purpose built for Business-critical apps



iSCSI Acceleration
Standard Ethernet Based infrastructure



FCoE Acceleration
Blades and Converged



Hyper Converged Systems
vSAN, AzureStack on standard networking



Concurrent FC-NVMe and FCP

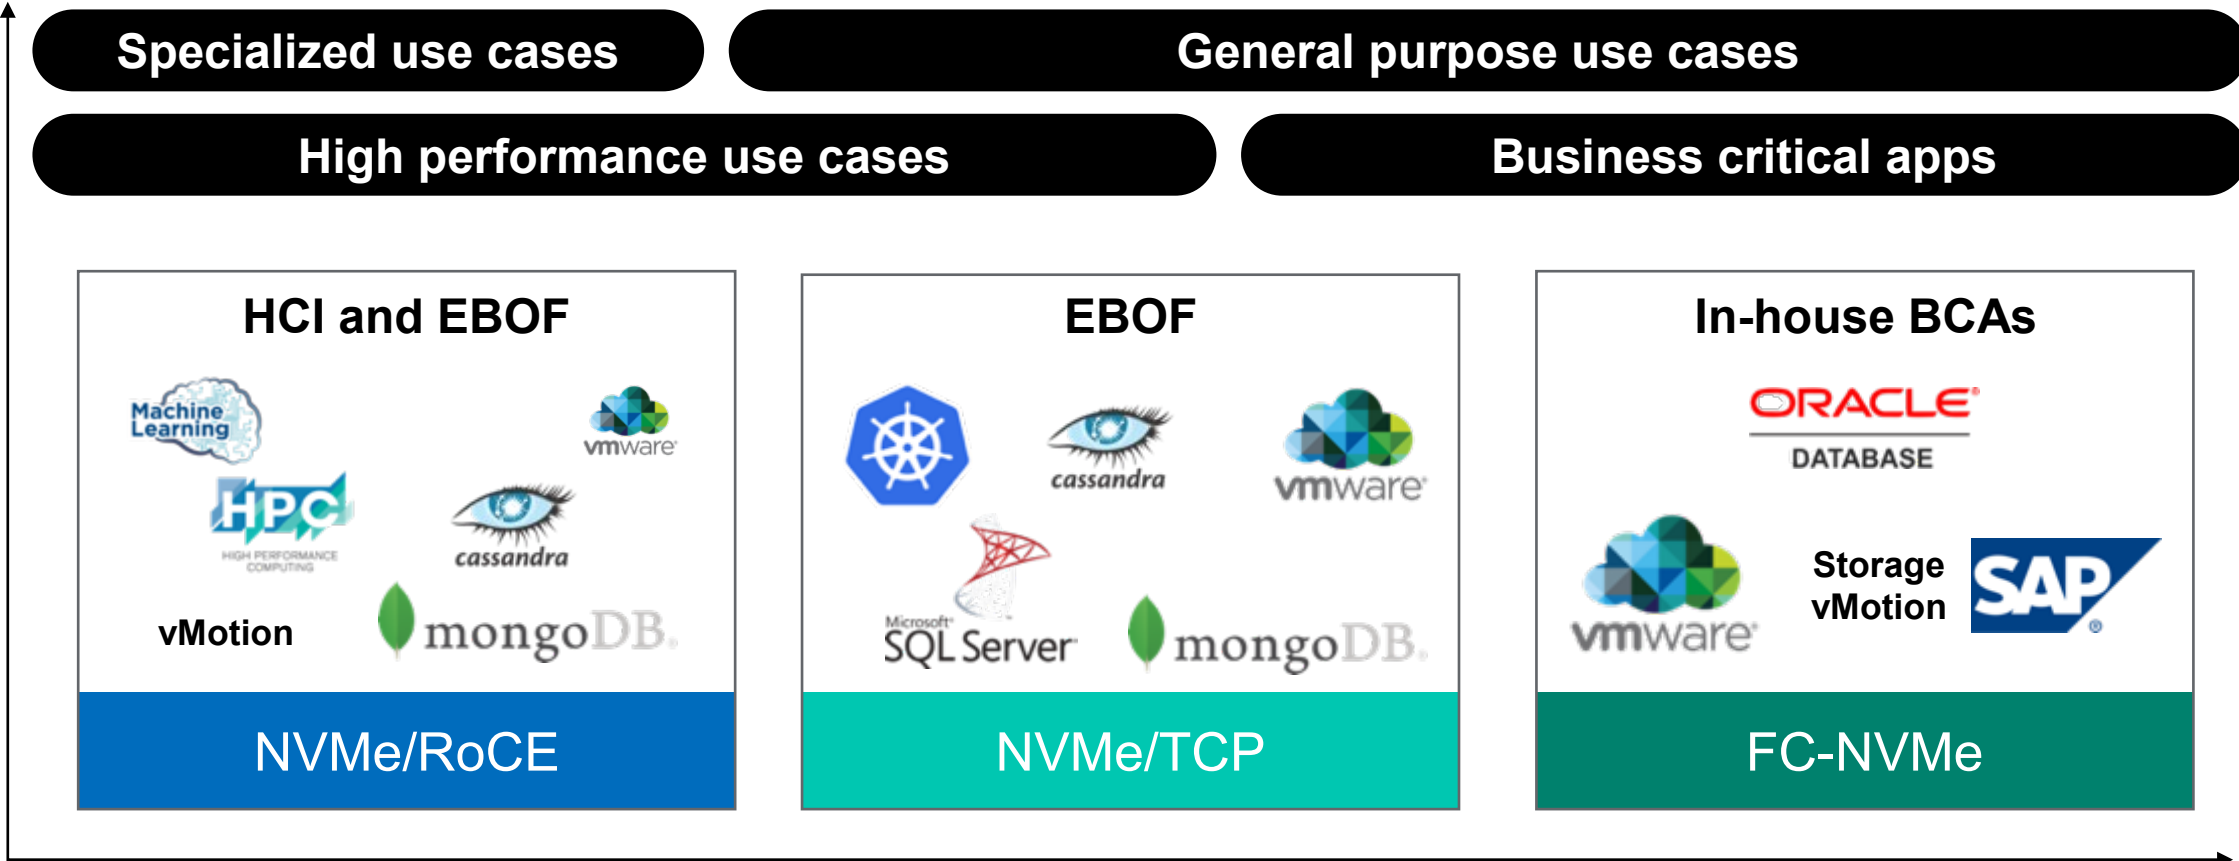


NVMe/TCP and NVMe/RoCEv2



RDMA Acceleration for HCI

Storage fabrics use cases



Logos are indicative of workload characteristics only.

HCI = Hyper Converged Infrastructure
EBOF = Ethernet BO Flash
BCA = Business Critical Apps



Autonomous fabrics? Need to go beyond day 0!

iSCSI

Automated discovery

- iSNS
- Proprietary methods

NVMe/TCP

Automated discovery

- Centralized discovery controller
- Direct discovery controller



FC-NVMe

Discovery, metering, congestion, auto routing

- Discovery
- Virtual machines IDs
- Fabric notifications (FPINs)

Sensing the road ahead – foundation for the autonomous SAN

“Awareness” of FC SAN conditions

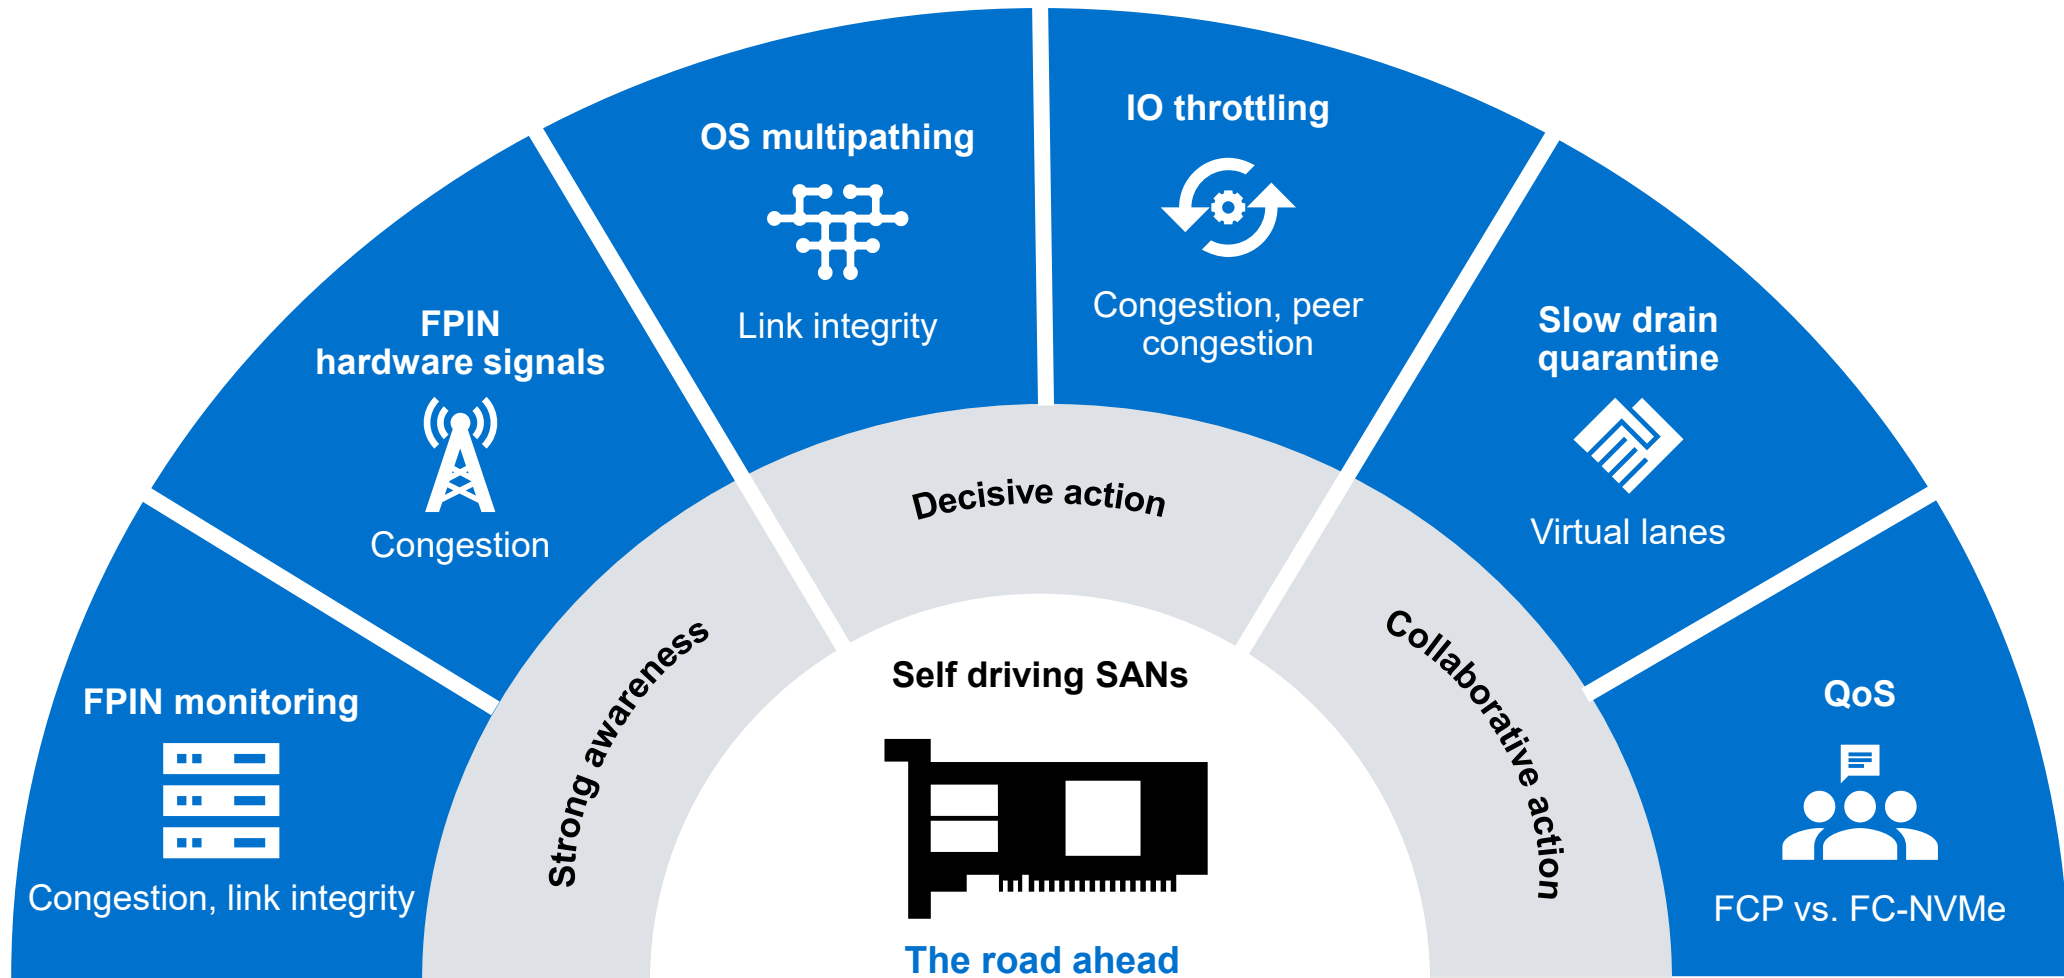
- **Link integrity (Potholes)**
 - Awareness: Physical layer errors exceed a threshold
 - Decisive Action: Use an alternative path or adjust for the error rate
- **Congestion (NB Traffic)**
 - Awareness: Fabric has identified the congestion source (it's you)
 - Decisive Action: Reduce the rate of IO requests or quarantine flows
- **Peer congestion (SB Traffic)**
 - Awareness: Fabric has identified the congestion source (not you)
 - Decisive Action: Use another path or reduce the transfer rate
- **Delivery (can't make it)**
 - Awareness: Fabric transmit timeout, cannot deliver
 - Decisive Action: Report IO error immediately (short-circuit protocol timeout)



**INDUSTRY
STANDARD**

“Decisive Action” – fight or flight!

The path to the autonomous SAN



Road to Autonomous SANs

Thinking beyond Day 0

What it takes to build an autonomous SAN?

Understanding road conditions and driver behavior



Adoption and co-existence of heterogeneous speeds

Longer shelf life of a single FC speed
8/16/32FC in datacenters concurrently



Increased pressure on CapEx and OpEx

IT asked to improve utilization
Workload migration



Flash storage

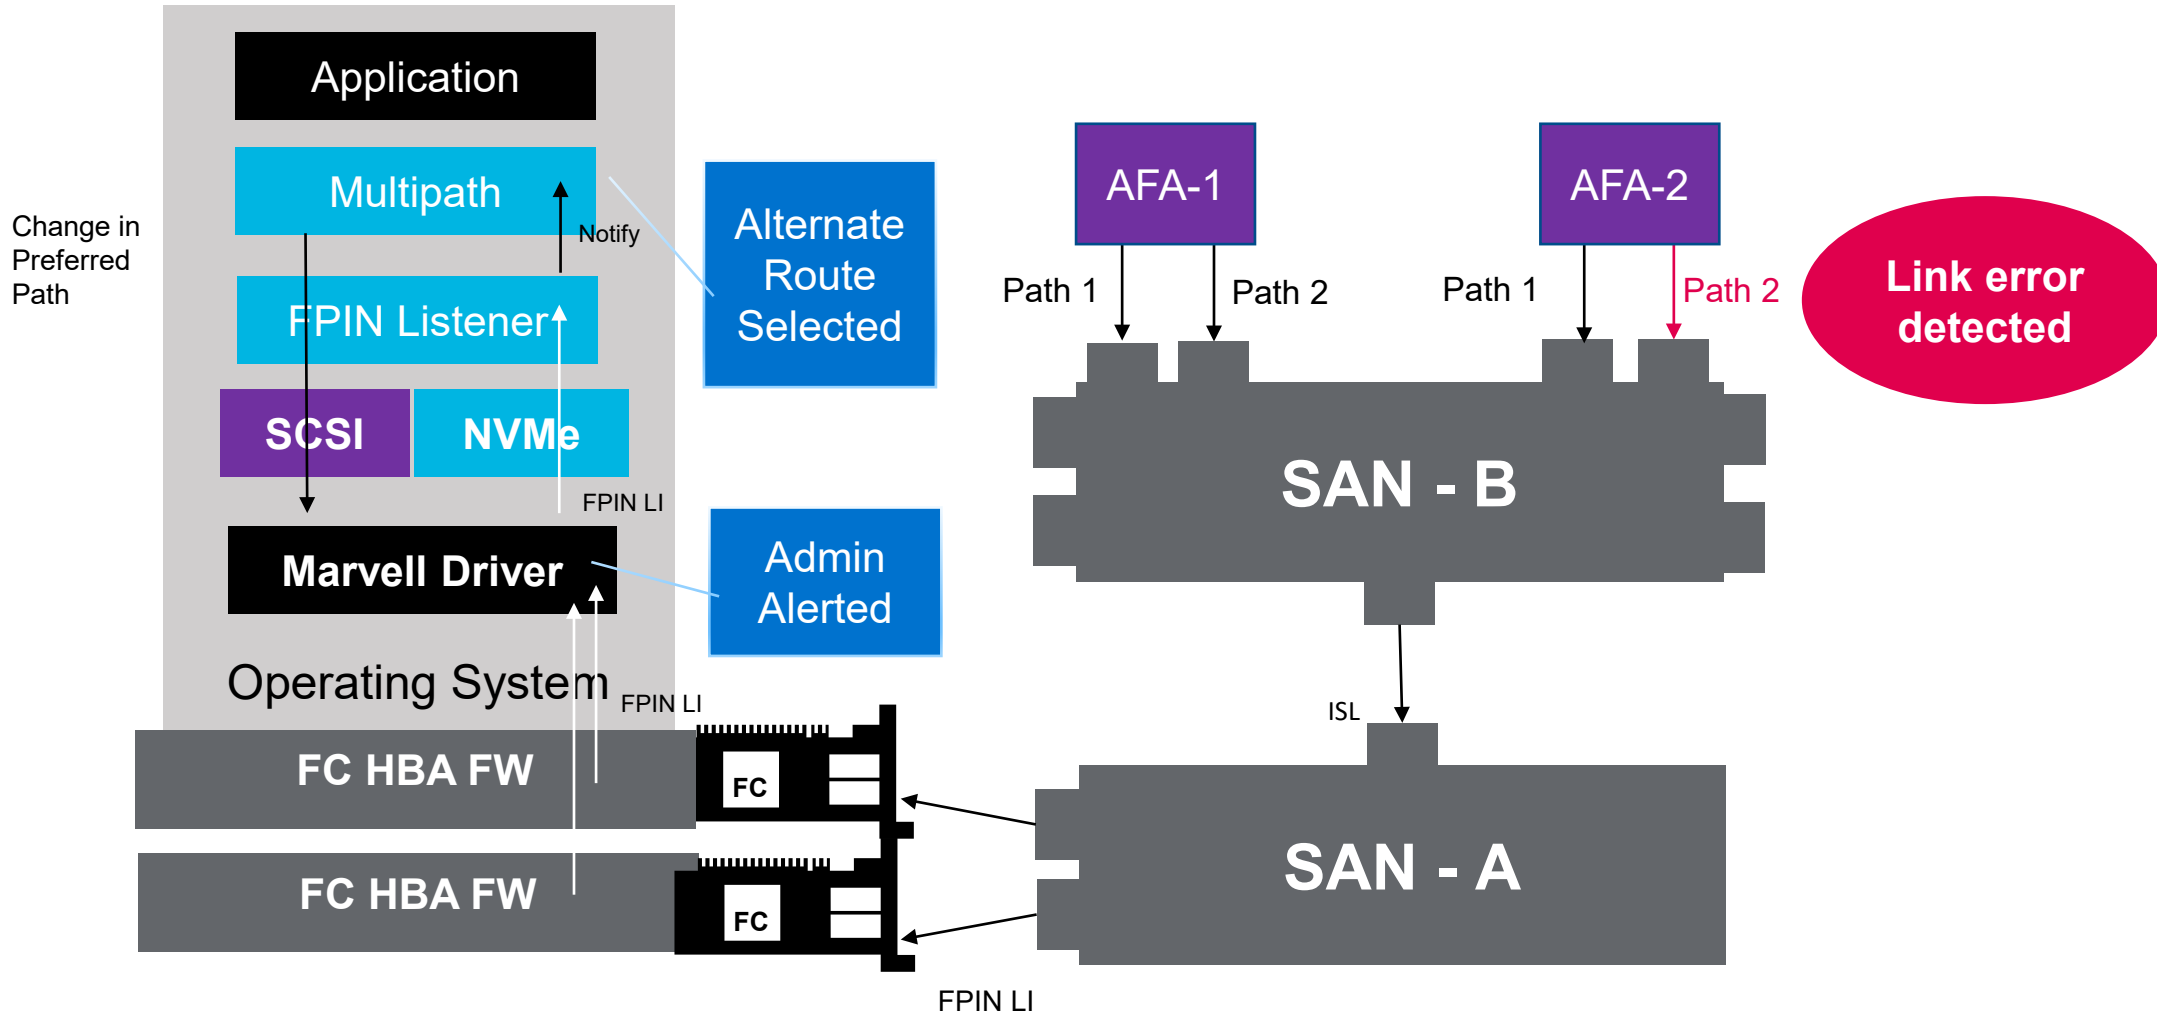
Storage no longer throttling the SAN
Pushing SAN infrastructure to its limit



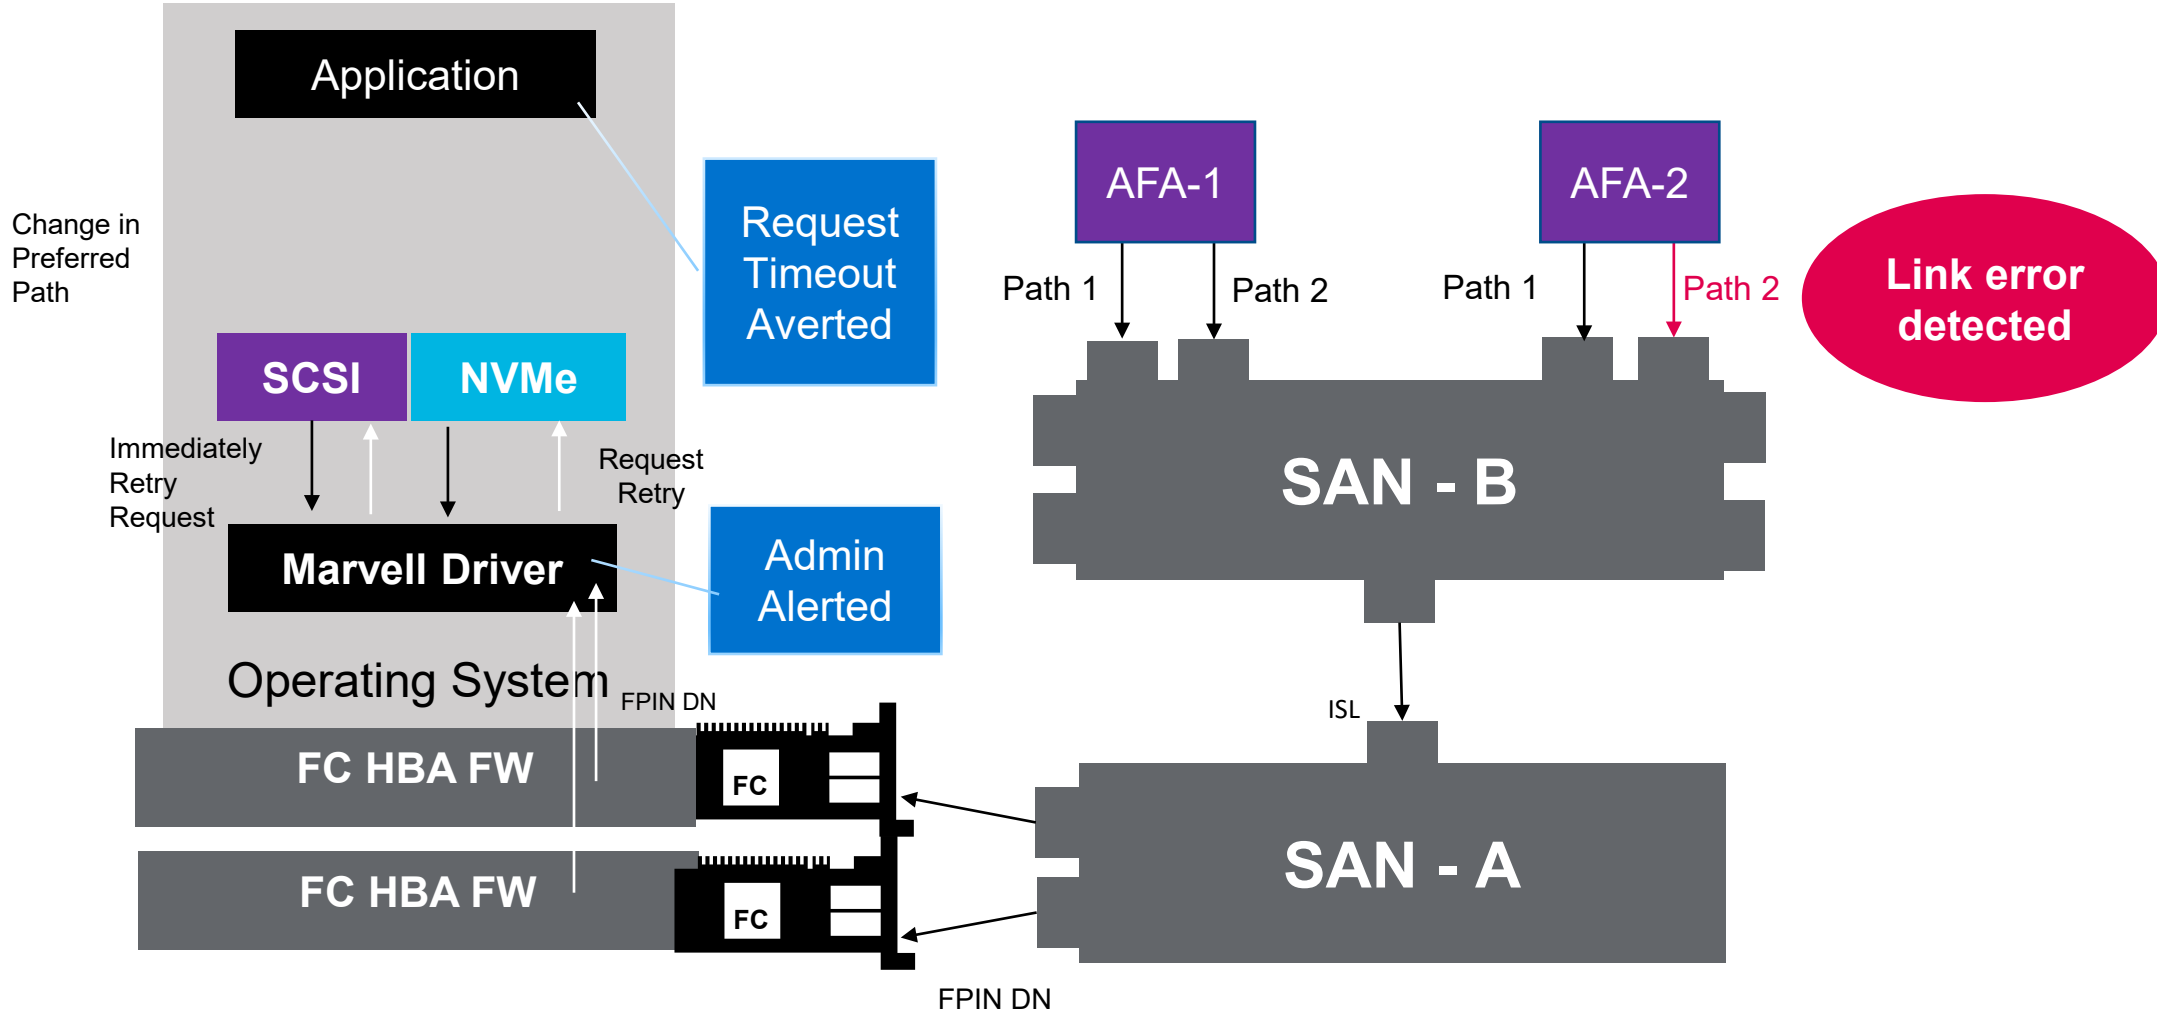
Degraded Physical Links

Consistent performance required smarter reporting that goes beyond a hard link down

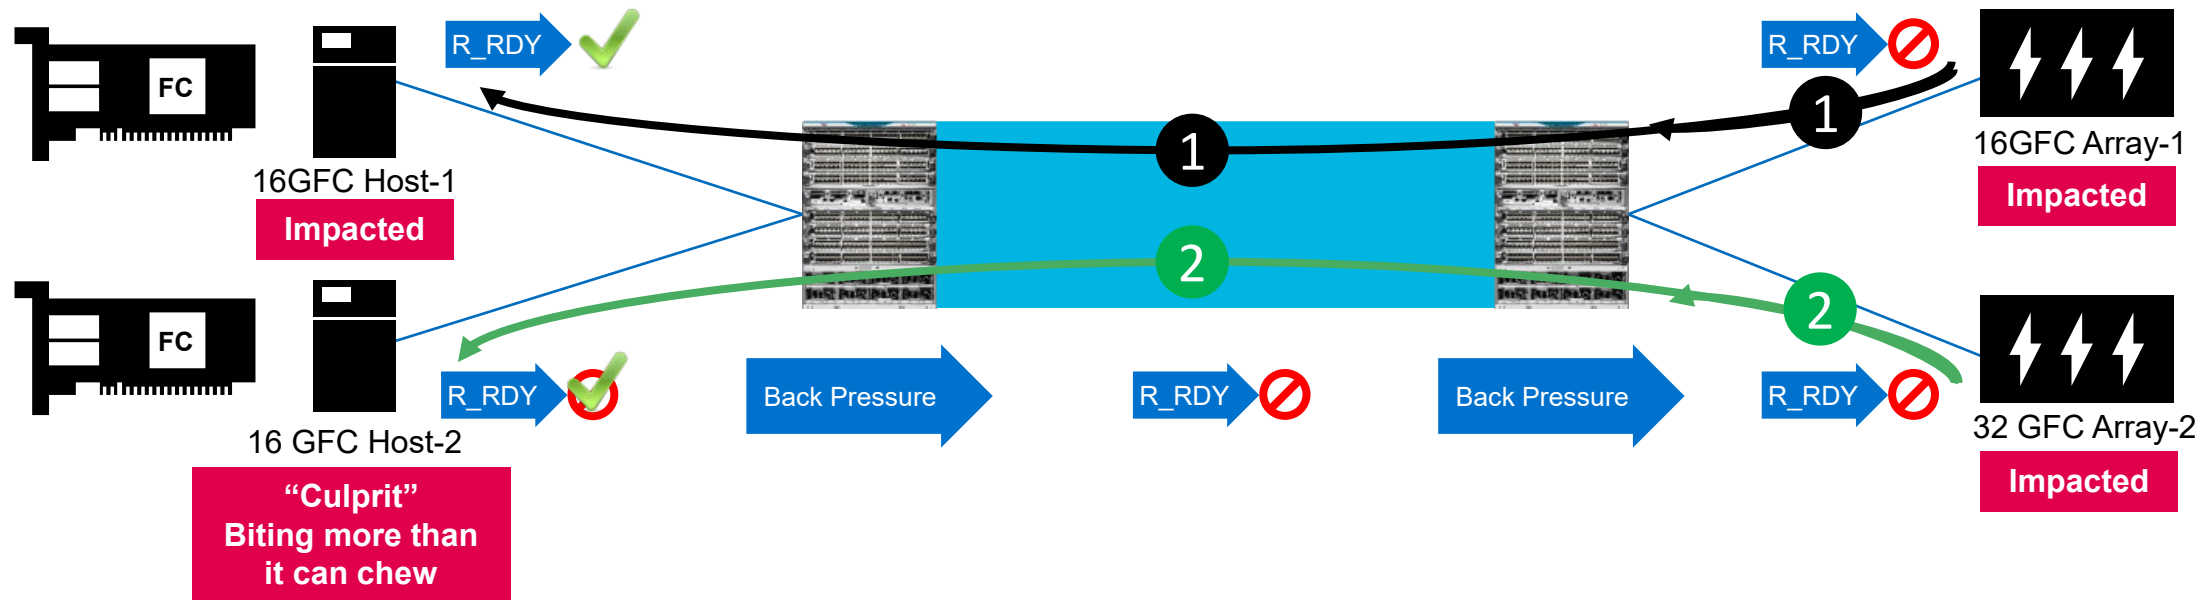
Mitigating temporal link / layer 1 issues



Mitigating temporal link / layer 1 issues



Flash arrays and the oversubscription problem



- A single target upgrade (Array 2), with the ability to send faster responses can cause slow drain / oversubscription
- Switchport connected to a slow drain device (Host 2) is starved for B2B credits
- Impact is seen on multiple end-devices sharing the same pair of switches and ISLs

Congestion mitigation solutions



Switch only

Independent of FC HBA Capability

- Cisco DIRL
- Brocade SDDQ, Port Fencing, Traffic Optimizer

HBA driven

Switch driven awareness, HBA driven Action

- Brocade and Cisco FPIN
- QLogic USCM IO Throttling (initiator driven)

Collaborative

Switch driven awareness, HBA + SW action (isolation)

- Brocade and Cisco FPIN
- Brocade and Cisco Virtual Channel

Level of mitigation



Switch only

- Triaging requires an analysis of fabric statistics.
- Issue noticeable only after a severe impact.
- Mitigation options are limited.



Switch + host

- Simplified triaging with alerts and warnings on impacted hosts.
- Issue noticeable before severe impact.
- Mitigation options are greater (but not perfect).



Switch + host + target

- Simplified triaging with alerts and warnings on impacted hosts and targets.
- Issue noticeable before severe impact.
- Potential of eliminating Congestion altogether.

What comes next?

You!

Calling the cavalry!

- 1 Spread the awareness – literally and programmatically.
- 2 Determine who is best positioned to “take action” – device or stack or app
- 3 Plug the gap in standards – wants vs. need?
- 4 Visibility for workload balancers / schedulers
- 5 Test drive, this is in our GA code (“actions” are disabled by default)

Thank You!



Please take a moment to rate this session.

Your feedback is important to us.