Introducing Fabric Notifications, From Awareness to Action

Live Webcast
April 28, 2021
10:00 AM PT/1:00 PM ET

Today's Panel

Moderator
Kiran Ranabhor
Cisco

Panelist
Howard Johnson
Broadcom

Panelist
Rupin Mohan
HPE

Panelist
Mark Jones
Broadcom

Panelist
Nishant Lodha
Marvell

About the Fibre Channel Industry Association (FCIA)

25+ Years
Promoting Fibre Channel Technology

Industry Leading
Member Companies

142M+ FC Ports
Shipped Since 2001

Agenda

- SAN Automation
- Fabric Notifications
- Questions and Answers Panel Discussion

SAN Automation – Enabling Deployment & Managing SANs Simple

Cloud Orchestration Platforms

ANSIBLE

Industry Tools

HPE OneView

HPE Smart SAN
HPE Network Orchestrator

HITACHI
Inspire the Next

IBM Spectrum Control

NetApp OnCommand Insight

openstack

Fibre Channel SAN Automation

Python Libraries

REST API

Fabric Notifications, Peer Zoning, Target Driven Peer Zoning

The Problem

Flakey paths

Persistent, intermittent errors

- Significant role in customer escalations
- Difficult for traditional solutions to resolve
- Required manual intervention increases mitigation costs
- MPIO solutions struggle with resolution, which impacts the dual fabric paradigm

Causes

- Marginal cables, SFPs, connections, etc
- Congestion due to lost credit, credit stall, or oversubscription

Why now?

- Fibre Channel solutions are mature and diversified
- Identification and mitigation tool have evolved
- Customers are demanding more automation

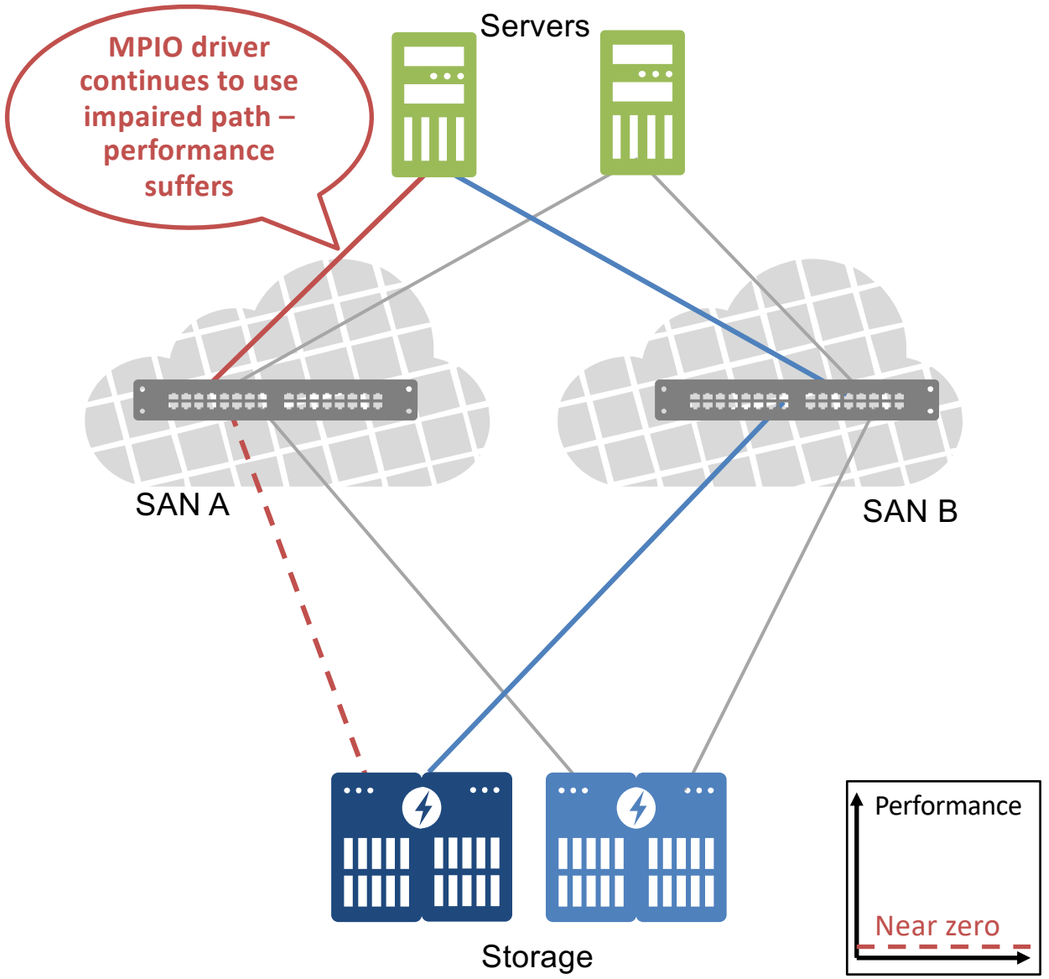

The Solution

Fabric Notifications

Fabric Notifications

- Notifications and signals
  - Generated by the fabric
  - Inform devices of impairments

Notifications

- Reporting: Events sent to registered devices
- Diagnostics: Helps efficiently evaluate errors
- Operation: Extended Link Services (ELS)

Signals

- Signaling: Report resource depletion to registered devices
- Diagnostics: Transmitter indicates resource usage
- Operation: Link level Primitive Signal

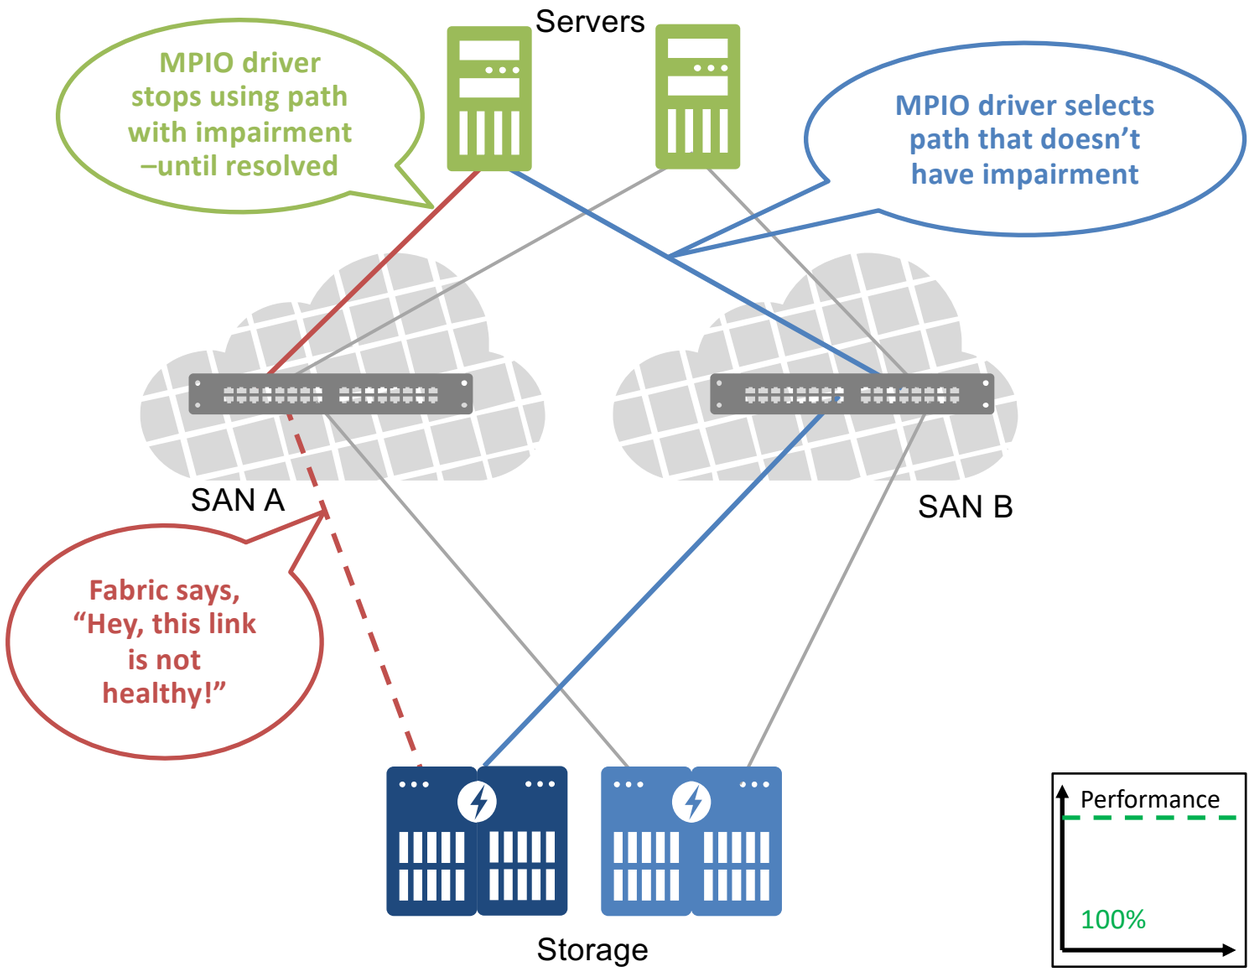

Fabric Notification

History

November 2014

– Fibre Channel ecosystem investigations

2015-2017

– Research and experimentation

2018

– Fibre Channel ecosystem collaboration

– Standardization starts

2019-2021

– Accepted into the T11 Standards

  – FC-FS-6: Congestion Signals (r0.3)

  – FC-LS-5: Notifications (r5.01)

  – FC-SW-8: Fabric detection and generation (r1.01)

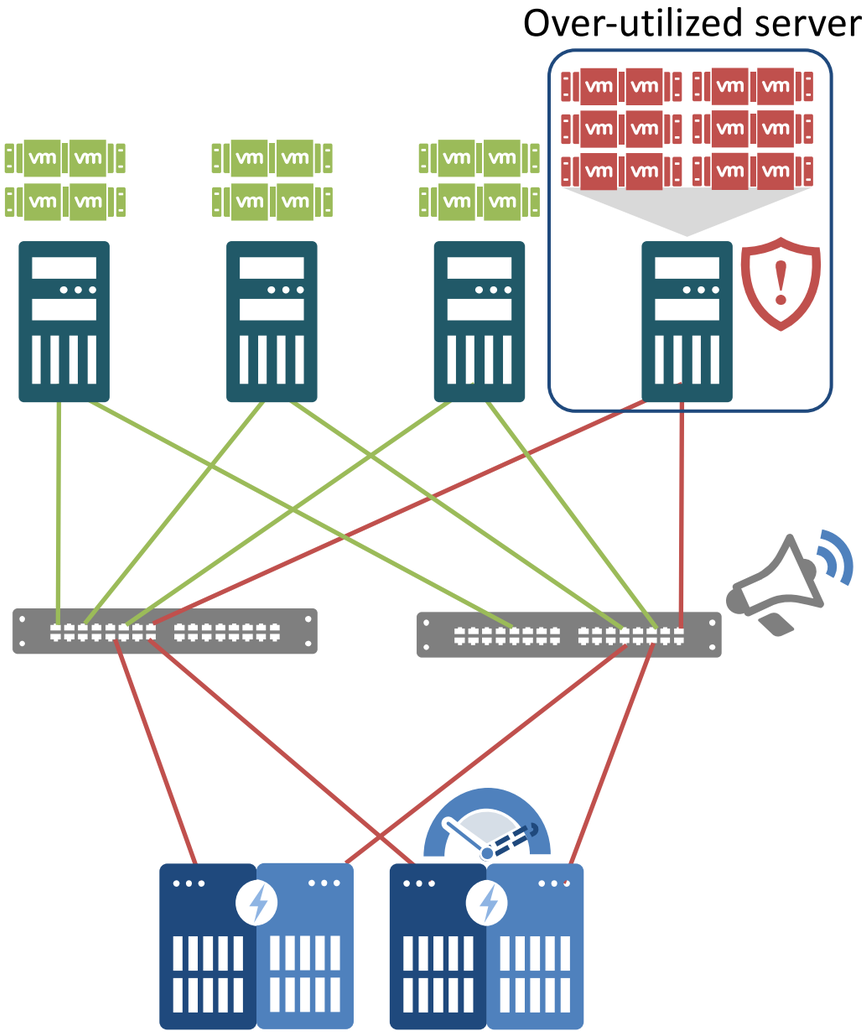

FC-SW-8 (r1.01)

Fabric Notifications Overview and Scope

Fabric Notifications overview

- Describes error detection, signaling and notification, and registration
  - See Clause 19
- Specifies scope

Fabric Notifications examples

- Provides use case examples
  - See Annex E (Informative) Fabric Notification information and examples
  - In-progress (r1.02)

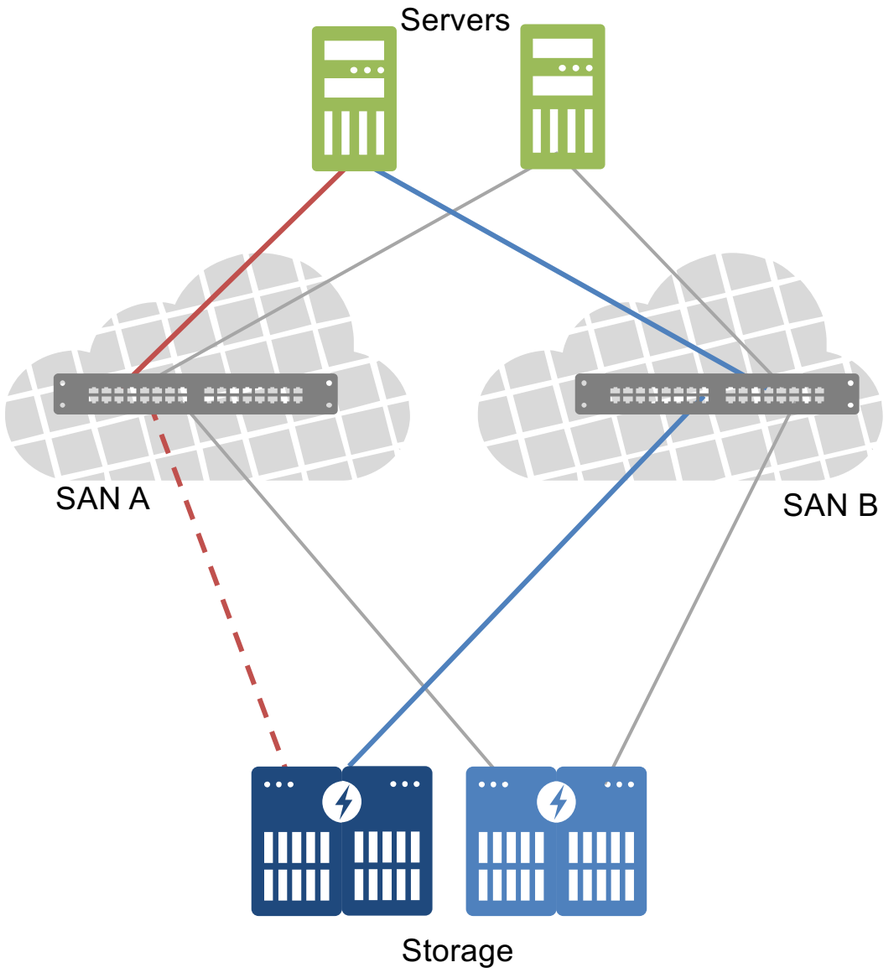

FC-FS-6 (r0.3)

Congestion Signals and F_D_TOV

Congestion Signal definitions

- Defines Warning/Alarm signals
  - See Tables 8 and 14
- Defines congestion signal use
  - See Clause 25 Congestion Signal

Congestion Signal examples

- Describes resource consumption
- Provides example of signal generation
  - See Annex L (Informative) Congestion Signal Examples

Frame Discard Timeout definition

- Defines F_D_TOV value and use
  - See Clause 22.3.6 F_D_TOV (r0.4)

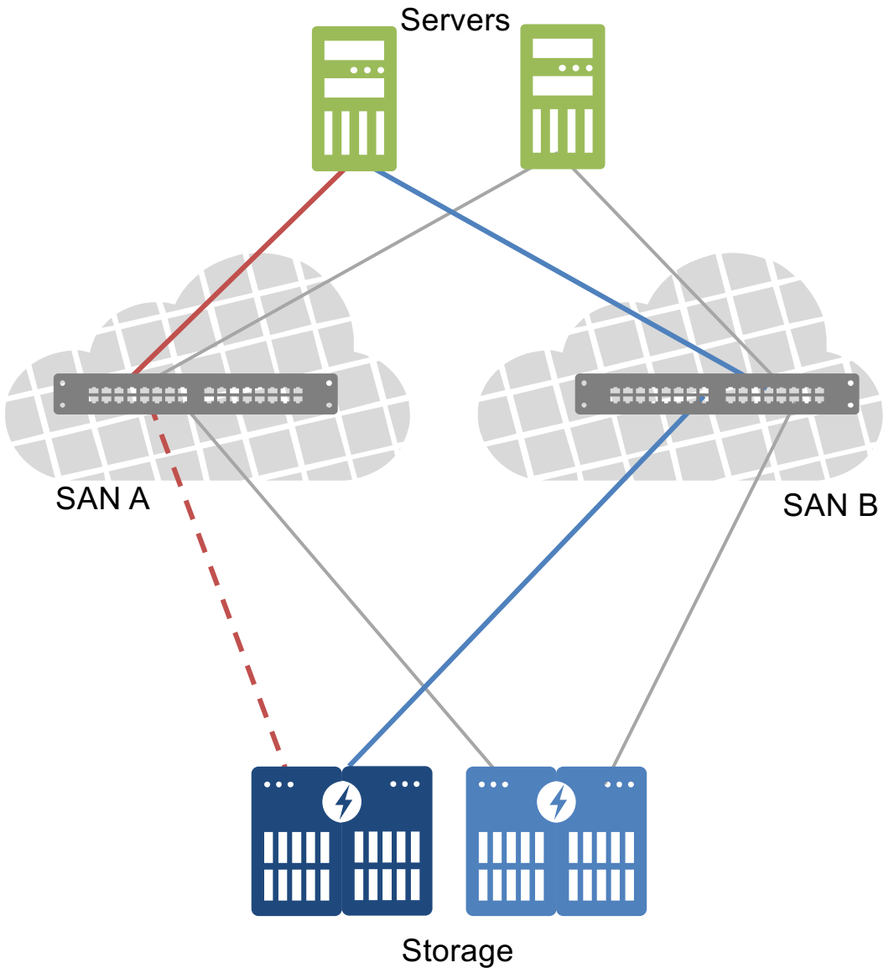

FC-LS-5 (r5.01)

Fabric Notification ELSs and Informative Annex

Fabric Notification ELS definitions

- Congestion Signal capability exchange
  - See clause 4.3.52 Exchange Diagnostic Capabilities (EDC)
- FPIN registration
  - See clause 4.3.53 Register Diagnostic Function (RDF)
- FPIN event descriptions
  - See clause 4.3.54 Fabric Performance Impact Notification (FPIN)
- Event type definitions (descriptor types)
  - See Tables 6 and 9

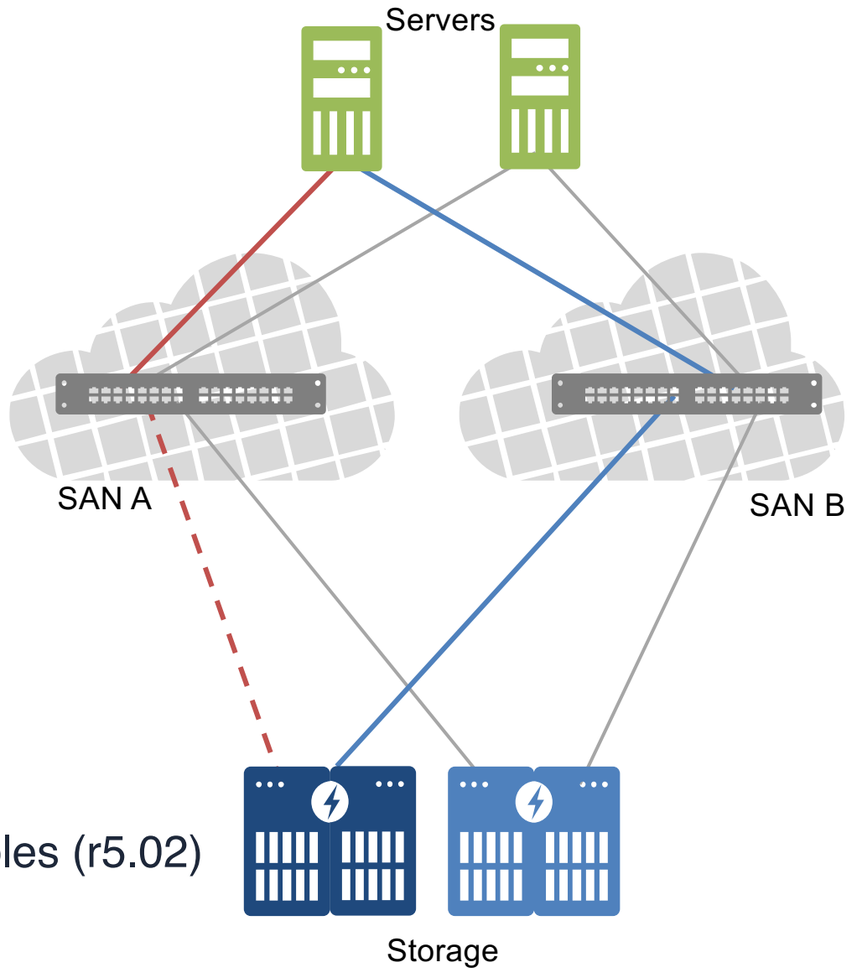

Fabric Notifications examples

- Provides use case examples and definitions
  - See Annex A (Informative) Fabric Notification information and examples (r5.02)

Fabric Notifications

Component Summary

Congestion Signal

- Primitive sent from transmitter to receiver
- Signifies resource depletion at the transmitter
  - I.e., frames are backing up

Notification ELSs

- Exchange Diagnostic Capabilities (EDC)
- Register Diagnostic Function (RDF)
- Fabric Performance Impact Notification (FPIN)
  - Link Integrity Notification (FPIN-LI)
  - Congestion Notification (FPIN-CN)
  - Peer Congestion Notification (FPIN-PN)
  - Delivery Notification (FPIN-DN)

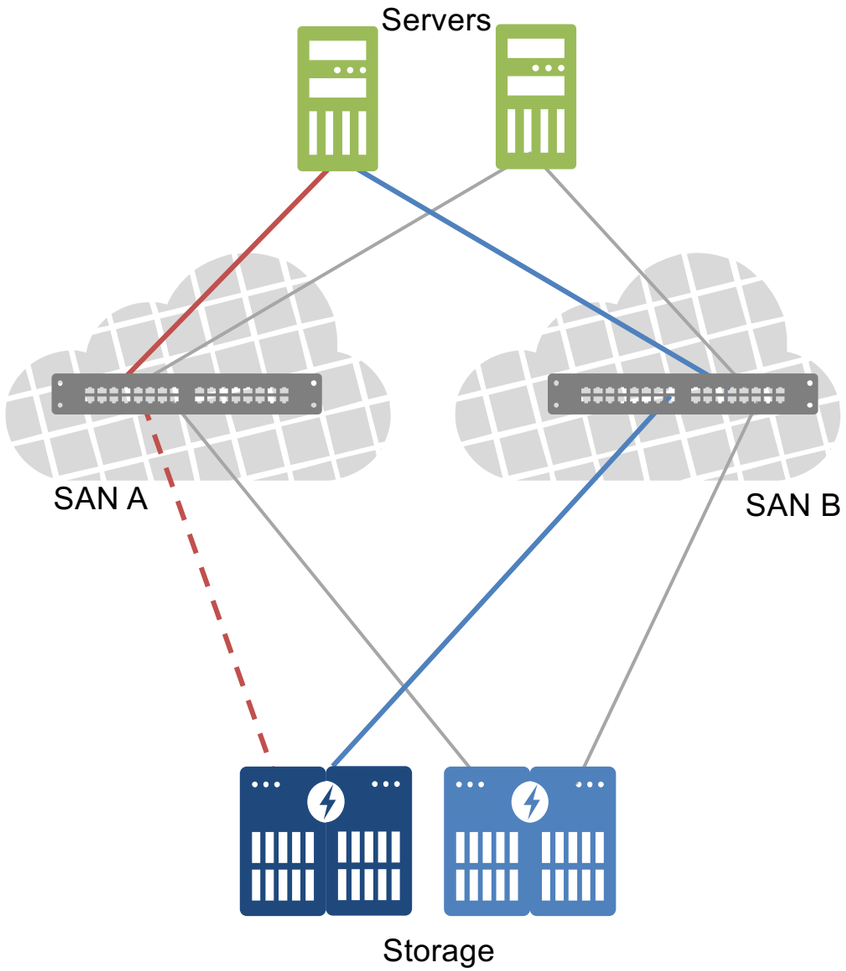

Problem Isolation and Determination

Register for Fabric Notifications and Log Events

Problem

- Link issues are difficult to isolate and resolve
- Fabrics and devices have different views of link issues

Solution

- Devices register for events and log notifications

Benefit

- Logged events provide detailed problem determination and isolation information
- Administrators gain insight into issues and are able to isolate and mitigate issues faster

Examples

- Server or storage array logs marginal link events
- Storage array logs events identifying an oversubscribed server
- Server logs events identifying an oversubscribed storage array

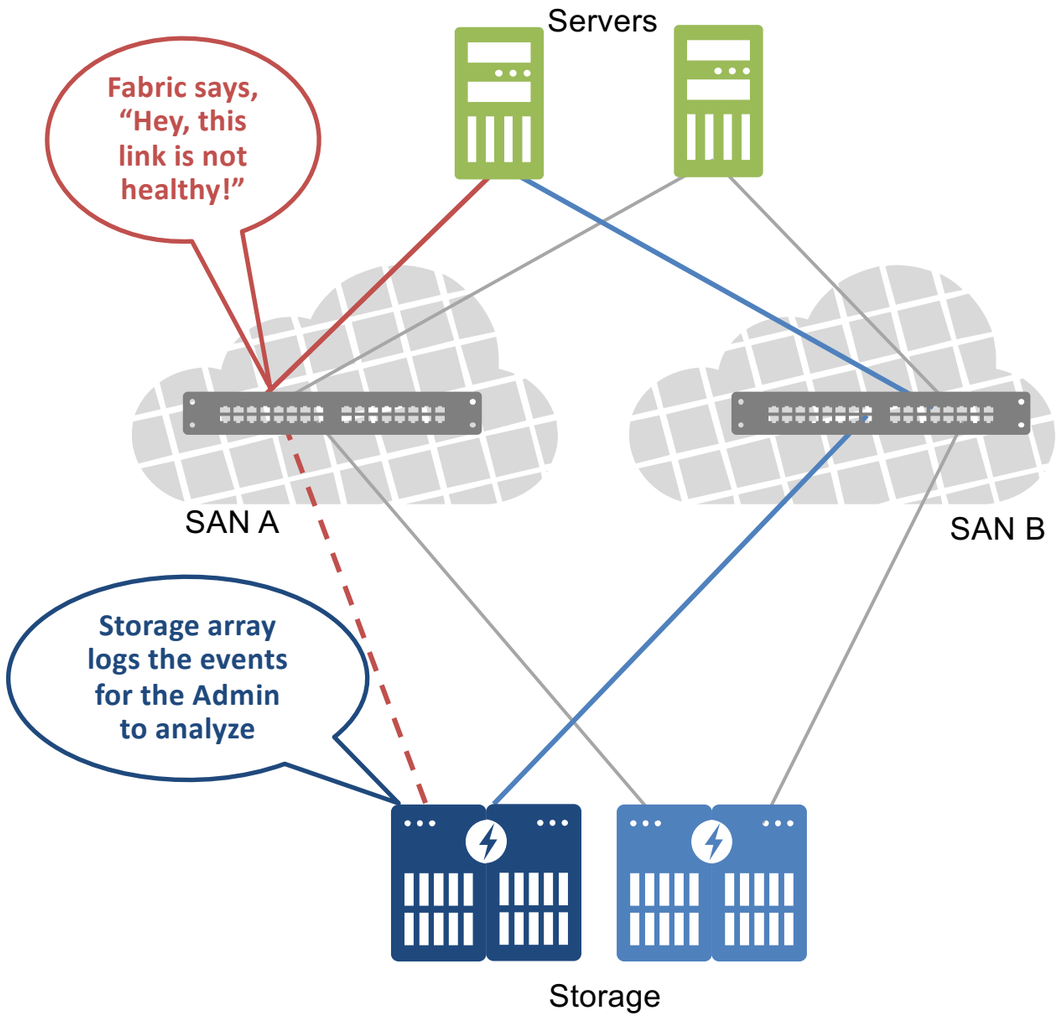

Link Integrity Isolation

Process and Report Link Integrity Events

Problem

- Link integrity issues disrupt Fabric operations
- Persistent, intermittent problems are difficult to isolate and resolve
- Fabric and devices have different views of link integrity issues

Solution

- Devices register for events and report detected link integrity events

Benefit

- Switches and devices monitor the link for marginal operation issues
- Significantly improves resiliency and reliability
- Servers and storage arrays automatically notify MPIO solutions

Examples

- Fabric detects physical errors and sends notifications to devices
- Device detects physical errors and sends notifications to the Fabric
- Initiators surface Link Integrity notifications to MPIO layer

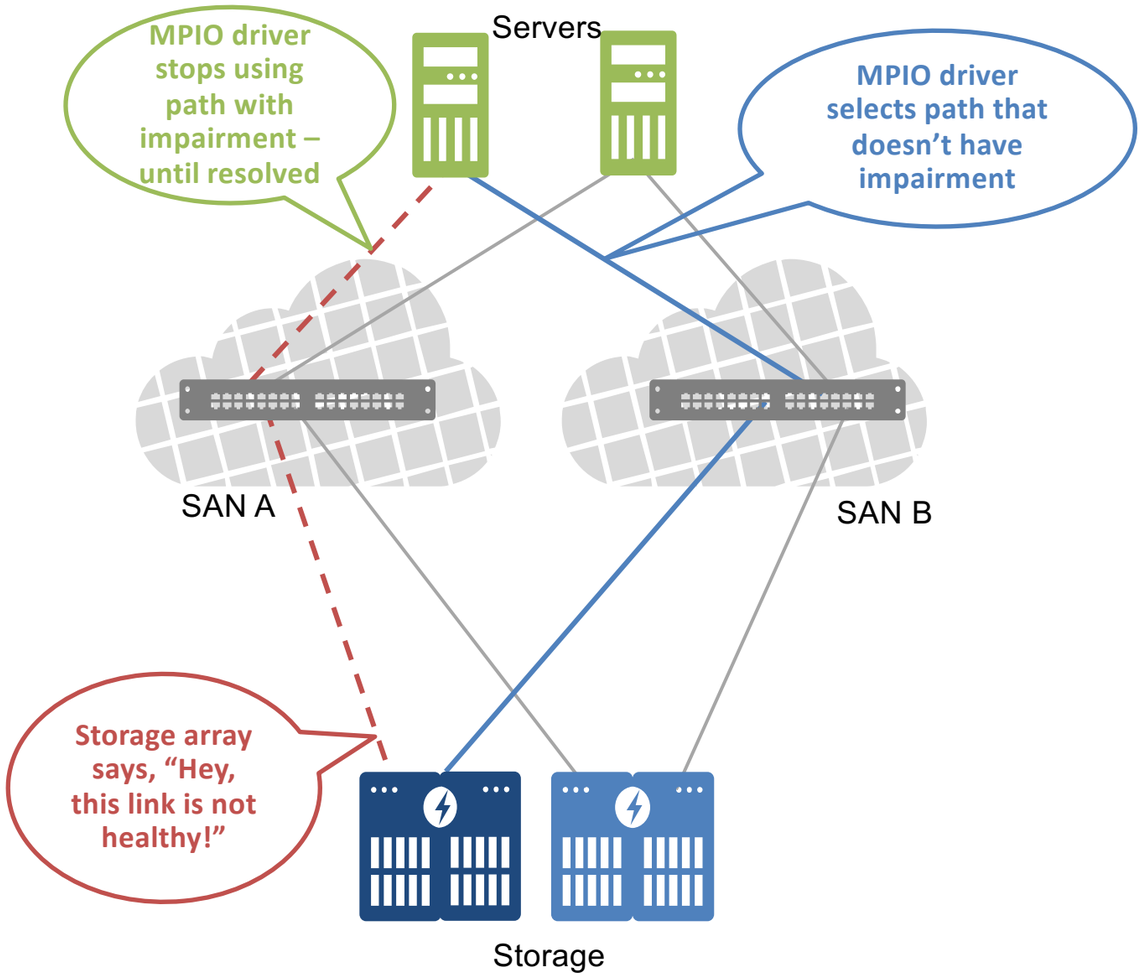

Target Credit Stall

Identify Internal Resource Constraints and Notify Initiators

Problem

- Target credit stall occurs when unsolicited commands fill the queue
  - "Unsolicited command queue" is fixed length, which causes backup into HBA buffers leading to Target credit stall conditions

Solution

- Targets register for events and sends throttling notifications to Initiators
- Targets use FDTOV to determine when to discard unprocessed requests

Benefit

- Devices automatically respond to internal constraints that lead to the Target Credit Stall condition

Examples

- Storage array sends notification to stop unsolicited requests
- Storage array discards unsolicited requests based on FDTOV
- HBA surfaces notification to MPIO layer to use an alternate path

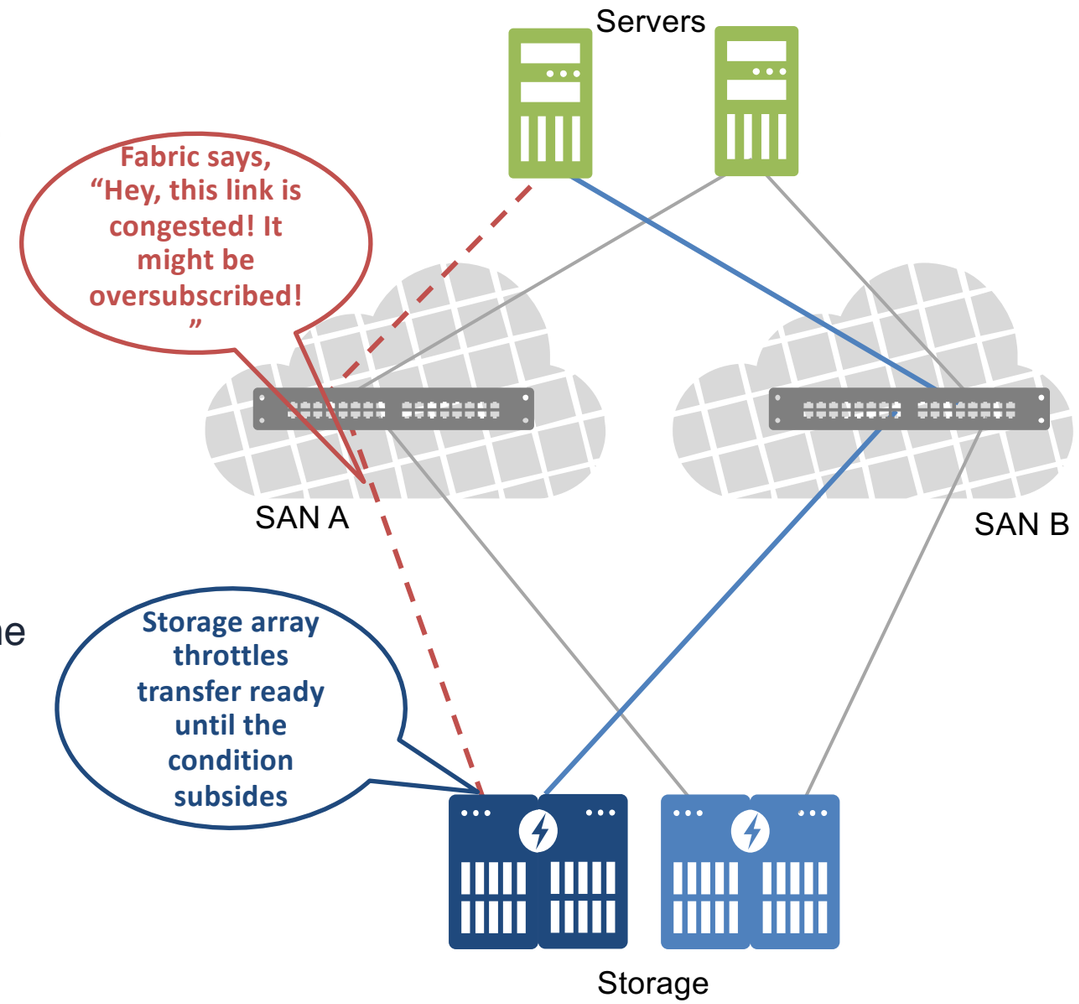

Read Oversubscription

Detect Oversubscription and Throttle Data Requests

Problem

- Read oversubscription occurs when Initiators are overrun by Target(s)
  - Initiators requesting more data than they can consume, Speed mismatches, multiple Targets zoned with a single Initiator, etc

Solution

- Initiators register for events and throttle incoming I/O
- Targets register for events and perform speed matching

Benefit

- Devices automatically responds to read oversubscription

Examples

- HBA throttles read requests to match the capacity of the local port
- Storage array reduces the rate of transmission to match the speed of the requesting Initiator(s)

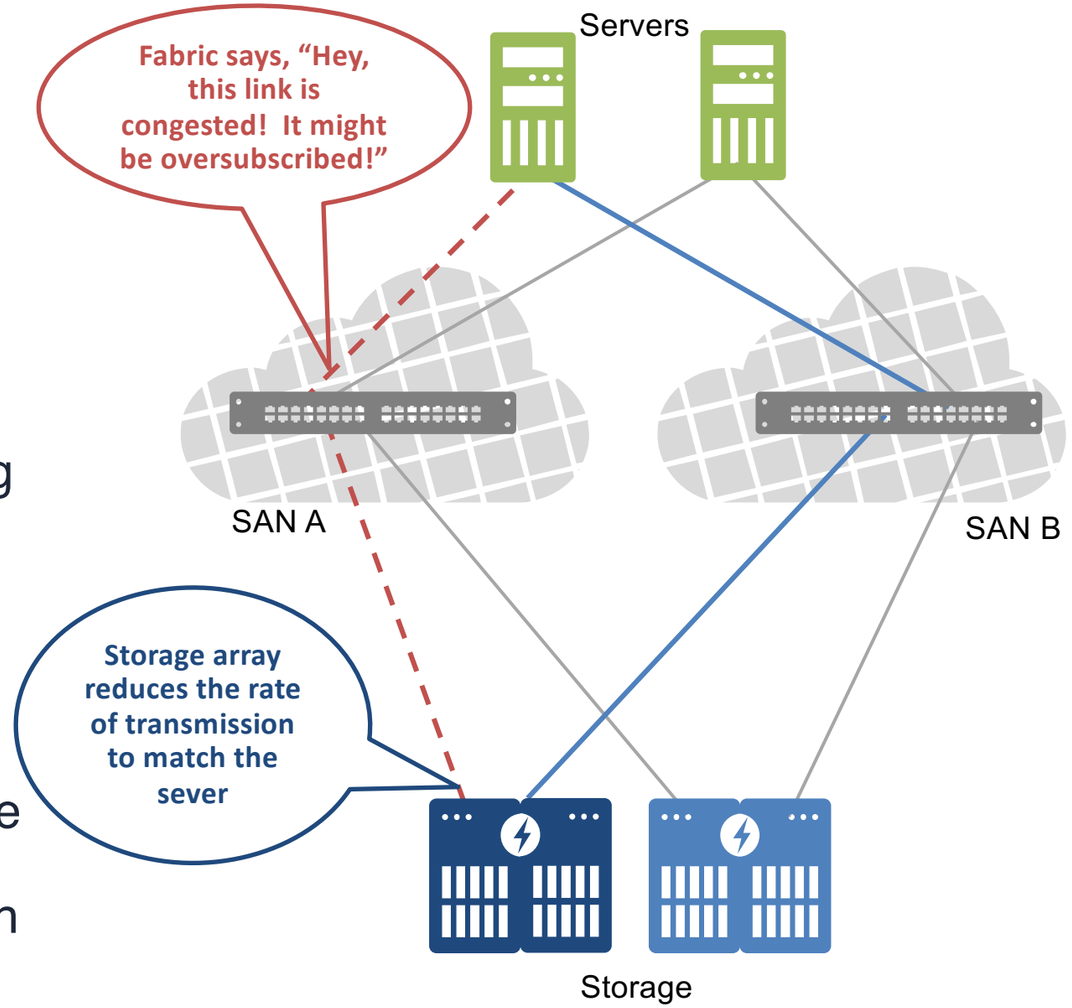

Write Oversubscription

Detect Oversubscription and Throttle Data Requests

Problem

- Write oversubscription occurs when Targets are overrun by Initiators
  - Speed mismatches, multiple Initiators zoned with the same Target, etc

Solution

- Target registers for events and throttles data transfers
  - May discard unprocessed requests
- Initiators register for events and favor uncongested paths

Benefit

- Devices automatically respond to write oversubscription

Examples

- Storage array throttles transfer ready until congestion notifications cease
- Storage array discards unsolicited requests after FDTOV
- Storage array sends notification to limit unsolicited requests from Initiators
- HBA surfaces notification to MPIO layer to use an alternate path

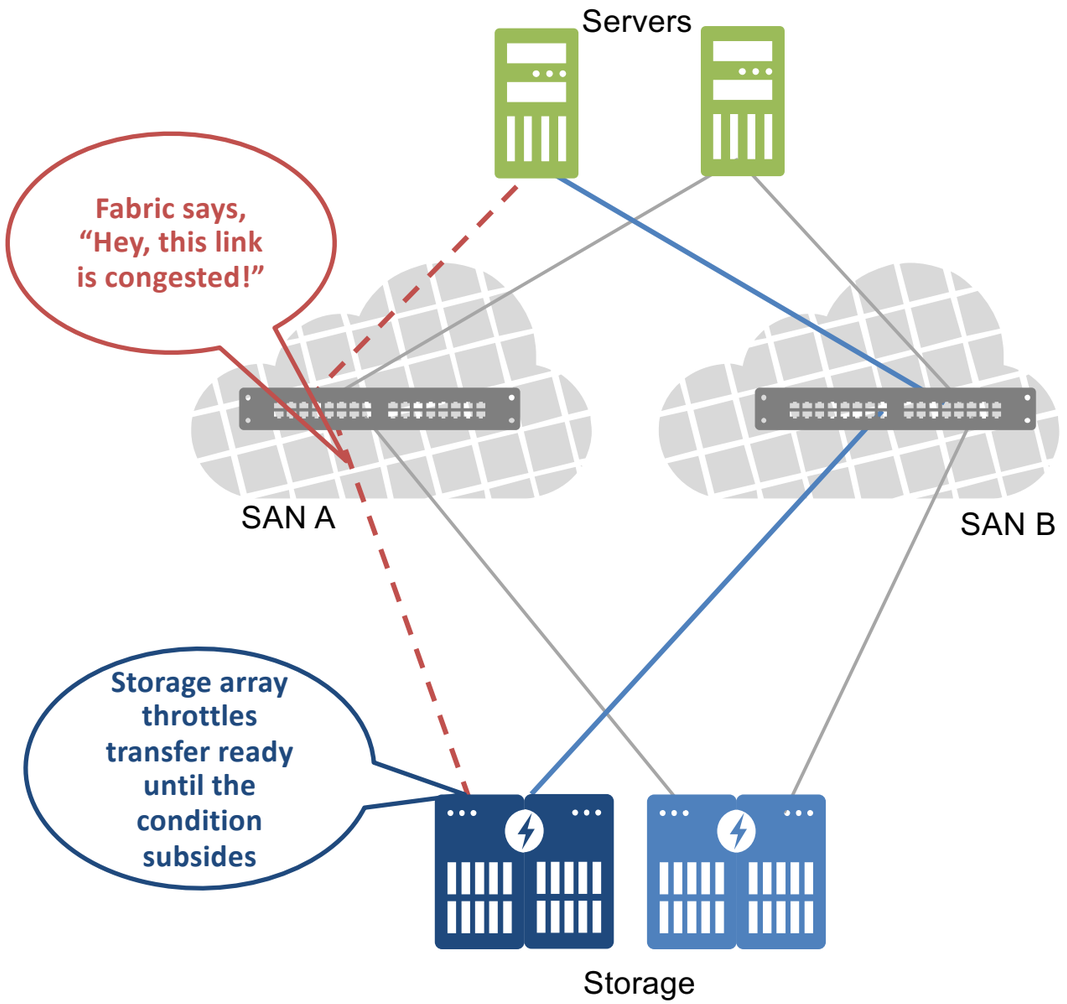

Array to Array Replication

Detect and React to Link Integrity and Congestion Events

Problem

- Array to array replication performance is impacted by link issues
- I/O based detection and recovery is incomplete

Solution

- Targets register for events and array adjusts replication behavior automatically

Benefit

- Storage array automatically responds to link integrity and congestion events
- Replication applications are more resilient to Fabric issues

Examples

- Storage array shifts the replication traffic to more reliable links
- Storage array favors alternative paths to the remote array to balance the replication traffic
- Remote array reduces the request rate to favor less used alternative paths

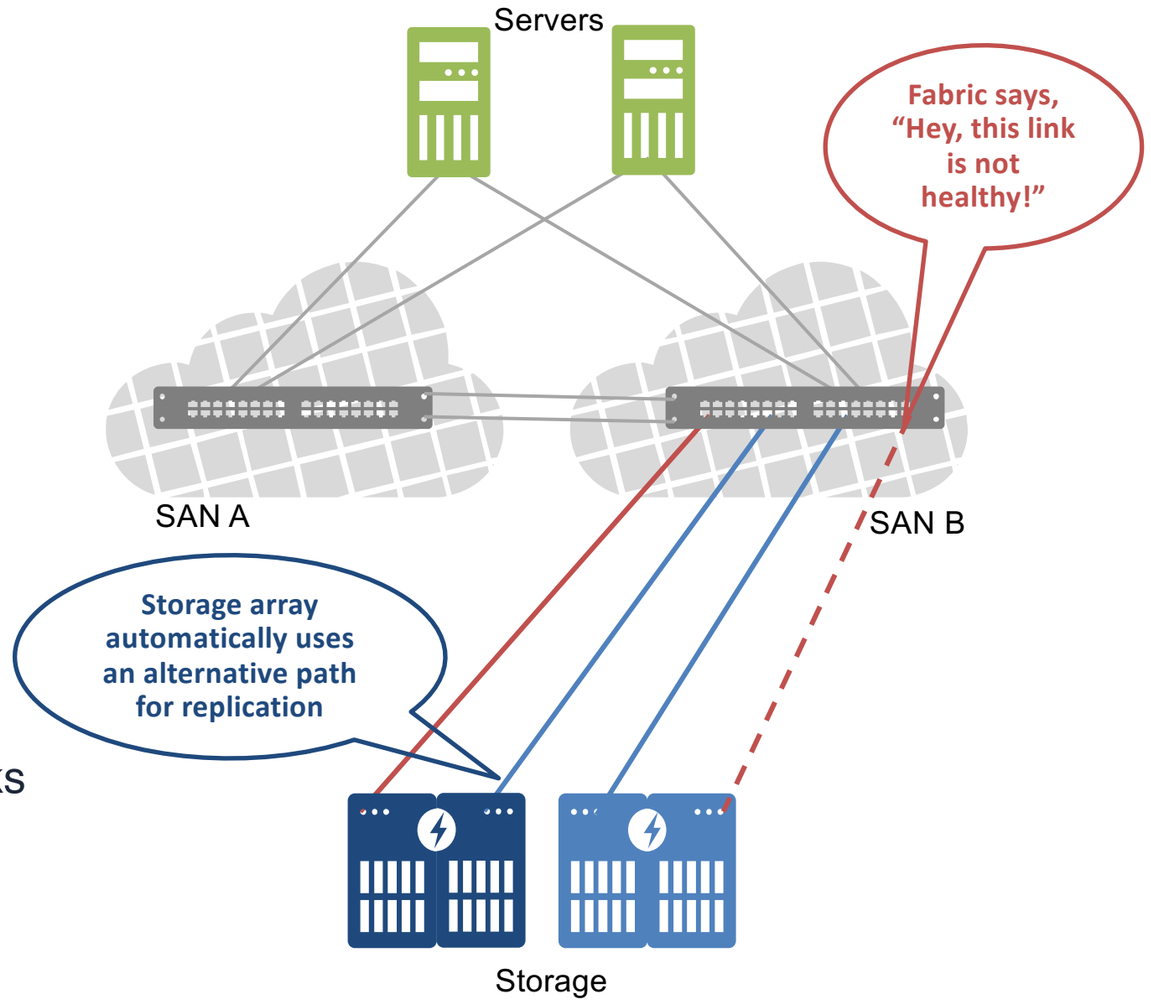

Panel Discussion

Moderator
Kiran Ranabhor
Cisco

Panelist
Howard Johnson
Broadcom

Panelist
Rupin Mohan
HPE

Panelist
Mark Jones
Broadcom

Panelist
Nishant Lodha
Marvell

More FCIA Resources

Visit our library of FCIA on-demand webcasts at http://fibrechannel.org/webcasts/ to learn about:

- Fibre Channel Fundamentals
- FC-NVMe
- Long Distance Fibre Channel
- Fibre Channel Speedmap
- FCIP (Extension): Data Protection and Business Continuity
- Fibre Channel Performance
- FICON
- Fibre Channel Cabling
- 64GFC
- FC Zoning Basics

Download at:
https://fibrechannel.org/fibre-channel-solution-guide-2020/

After this Webcast

- Please rate this event – we value your feedback
- We will post a Q&A blog at http://fibrechannel.org/ with answers to the questions we received today
- Follow us on **Twitter @FCIAnews** for updates on future FCIA webcasts
- Subscribe to our YouTube Channel: https://www.youtube.com/c/FibreChannelIndustryAssociationFCIA

Introducing Fabric Notifications
From Awareness to Action

Thank You!