



# Trucks Keep Right!

Maximizing Efficiency with Fibre Channel Virtual Lanes

Nishant Lodha, Marvell

# Fibre Channel

Reliable, secure, **block storage connectivity** for business-critical systems




Supply chain



CRM



Banking



Manufacturing



Telco



Cloud



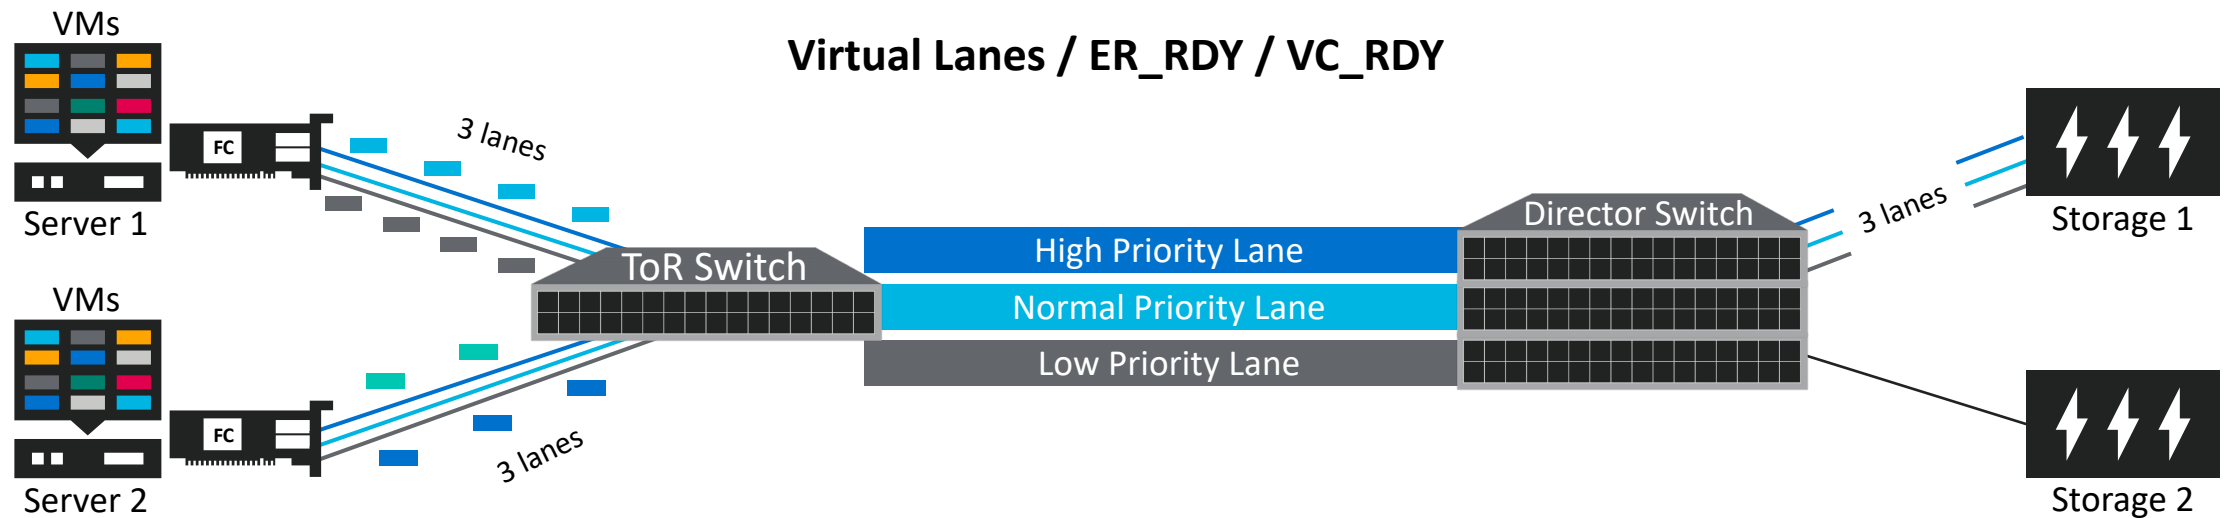

Accounting



Healthcare

Work on a “Credit”-based mechanism

# Credit based flow control in Fibre Channel



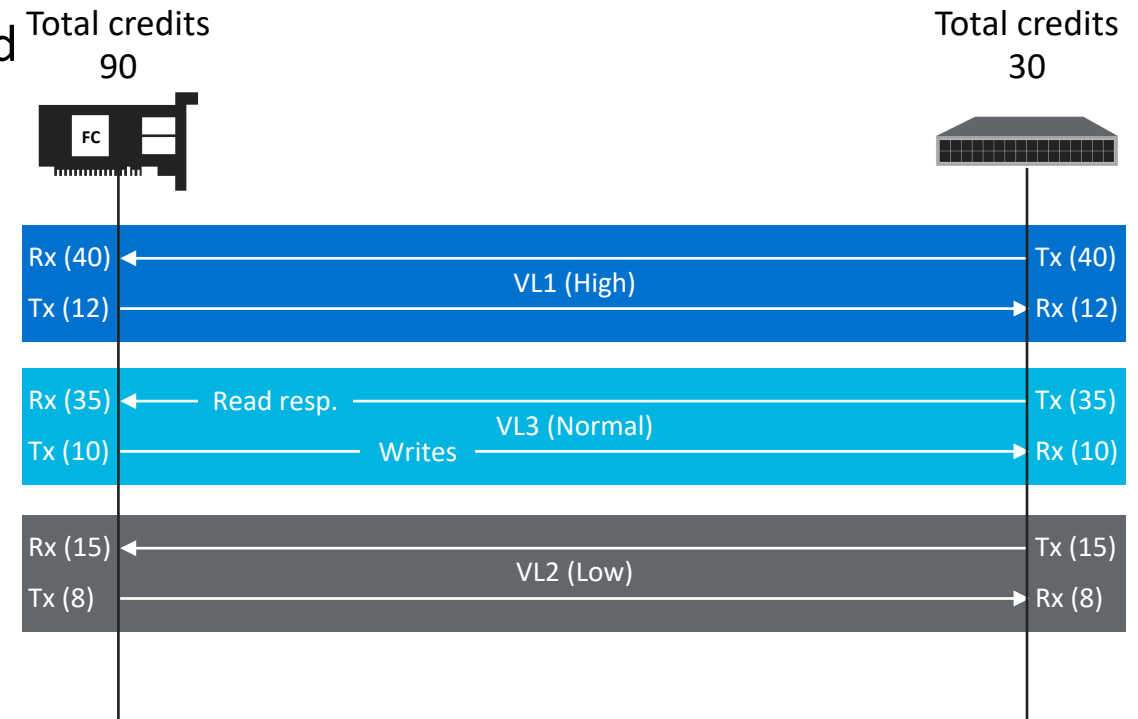
# FC Virtual Lanes

- HBA to Switch, typically a one lane highway
- Slow traffic impedes everyone else
- Very difficult to root cause, impacts SLAs
- Marvell Virtual Lane Technology (standards based)
  - Three lane highway
  - Slow traffic in slow lane, fewer credits
  - No head of line blocking
  - Dedicated lane from initiator to target
- Many use cases around segregation, QoS, traffic priorities



# How do Virtual Lanes work?

- **Configure:** Physical FC Link (and credits) divided into multiple independent lanes
  - Cisco MDS Virtual Links (VLs)
  - Brocade FC Virtual Channels (VCs)
- **Negotiate:** HBA ↔ Switch ↔ Storage @Login
  - Do you VL?
  - How many?
  - Rx BB\_CRs per VL?
  - Tx BB\_CRs per VL?
  - Priority Range values per VL?
- **Operate:** Frames are tagged (CS\_CTL for Cisco, Primitives for Brocade) and transmitted on specific VLs
- **Availability:** Marvell 32G and 64GFC HBAs



# Use case: Target congestion

Read and write traffic in progress, storage device is congested



## Storage congested

Peer Congestion Notification (FPIN PCN)  
(Target in zone is congested)

**Situation:** Writes to slow target. Target not sinking to media fast enough.

**Trigger:** Switch sends PCN to FC HBA via Fabric Notifications


**Action:** Move ALL IOs of this slow flow to a slow VL (isolate it)

**Clear:** Move flow back to normal/fast VL

**Use case in production using Marvell® QLogic® FC HBAs and Cisco MDS switches**

# Use case: Prioritizing Flash traffic

Application I/O spans disk and NVMe storage



## Hybrid storage access

Ability to prioritize flash traffic

**Situation:** DB server transacting with flash for queries and disk for logs


**Trigger:** HBA driver detects destination media / protocol

**Action:** Move all NVMe flows to fast VL and HDD flows to slow VL

**Clear:** No head of line blocking of NVMe IOs

# Use case: Mapping VM priority

Multiple VMs share the same fabric, map VM priority profile to fabric priorities



## Virtualized

Ability to prioritize individual applications

**Situation:** Virtual Machines sharing the physical storage infrastructure

**Trigger:** HBA driver receives VM priority along with VM-ID tags

**Action:** I/Os to high priority VMs mapped to fast and normal VLs

**Clear:** Storage fabric that is aware of administrator set VM priority



# Virtual Lanes avoid delays

- **Congestion** issues and lack of a workload aware storage fabric introduces delays
- **Virtual lanes** enable physical layer priority enforcement for Fibre Channel traffic
- **Futures:** New use cases, dynamic credit allocation, hypervisor and array integration



Q&A

# Thank you