



Education

iSCSI : A loss-less Ethernet fabric with DCB

Jason Blosil, NetApp
Gary Gumanow, Dell

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

- A look at different storage technologies and a brief look at each; benefits/tradeoffs, making the case for a SAN. Then looking at making your SAN a lossless, high-performance, predictable, resource for your business. We'll look at different storage protocols and how they compare to the OSI model, and the new DCB protocols. And lastly, I'll present some findings of using these technologies.

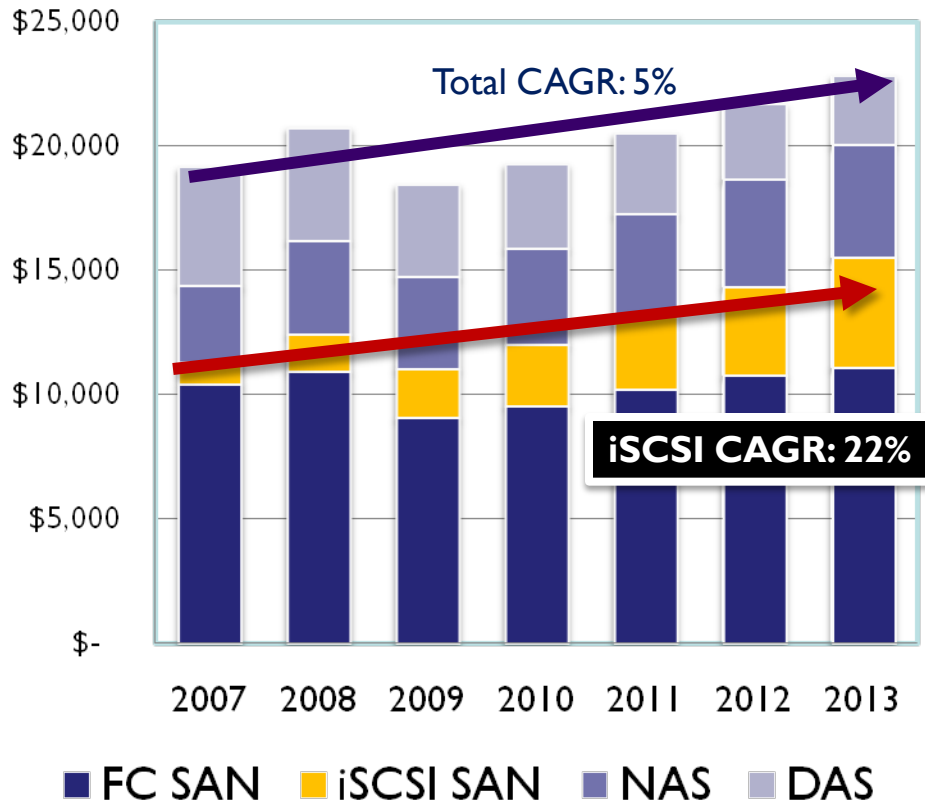
- **Learning Objectives**
 - ◆ A detailed look and comparison between three storage models; iSCSI, FC, and FCoE. Comparing all three to the OSI model. A discussion of the requirements for each layer of the networking stack.
 - ◆ A look at the benefits that DCB can provide iSCSI, as a lossless Ethernet fabric. A review of performance and contention issues.

What we'll talk about today...


- How is iSCSI doing today
- What is Data Center Bridging
- What's the roadmap for Ethernet
- A protocol comparison
- iSCSI over DCB performance
- Does DCB really matter for iSCSI

IDC Forecast (Dec 2009)

Revenue (\$M)



Revenue Share



➤ Rich set of tools

- ◆ Disaster Recovery / Backup
- ◆ Data Virtualization
- ◆ Thin provisioning
- ◆ Snapshots
- ◆ Replication
- ◆ VM integration
- ◆ Disk de-dupe

➤ SW Initiators available from all major OSes

- ◆ MPIO support
- ◆ Low CPU utilization without HW offload

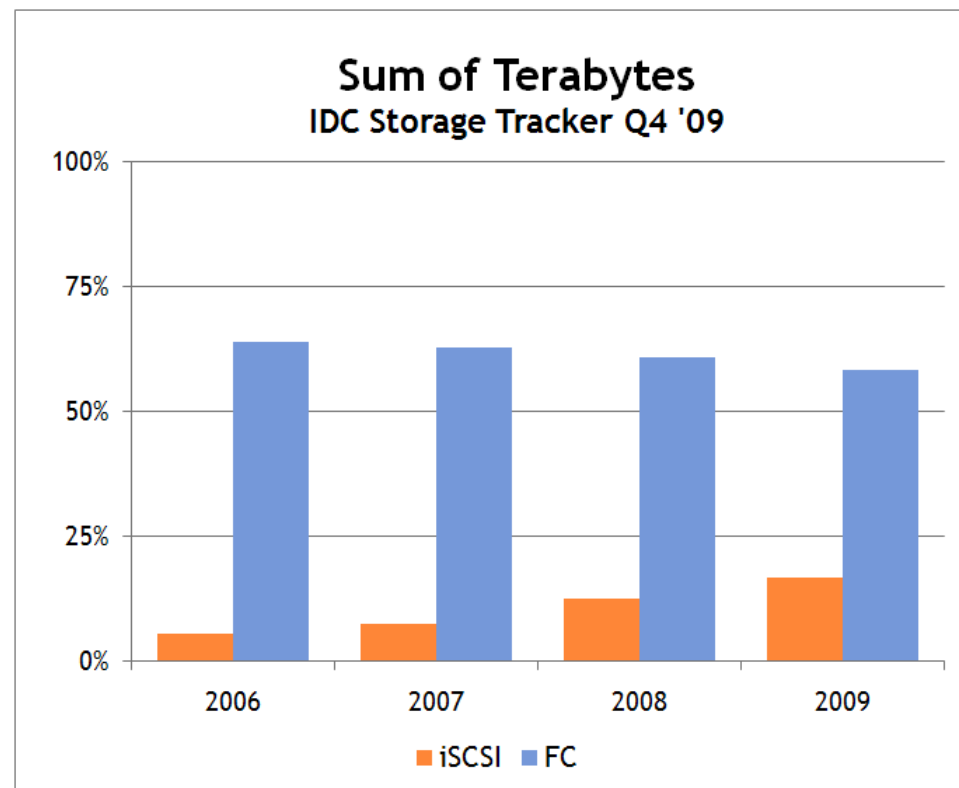
➤ Architecture choice

- ◆ Scale-out & Scale-up

➤ 10GbE available today

Sum of Terabytes	CAGR 2006-2009
iSCSI	43.6%
Storage Growth	29.0%

Source: IDC Storage Tracker Q4 2009



➤ Networking

- ◆ Data Center Bridging (DCB)
- ◆ 10 Gigabit Ethernet
- ◆ Top of Rack switching

➤ Storage


- ◆ Increased focus on Ethernet-based storage

➤ Server

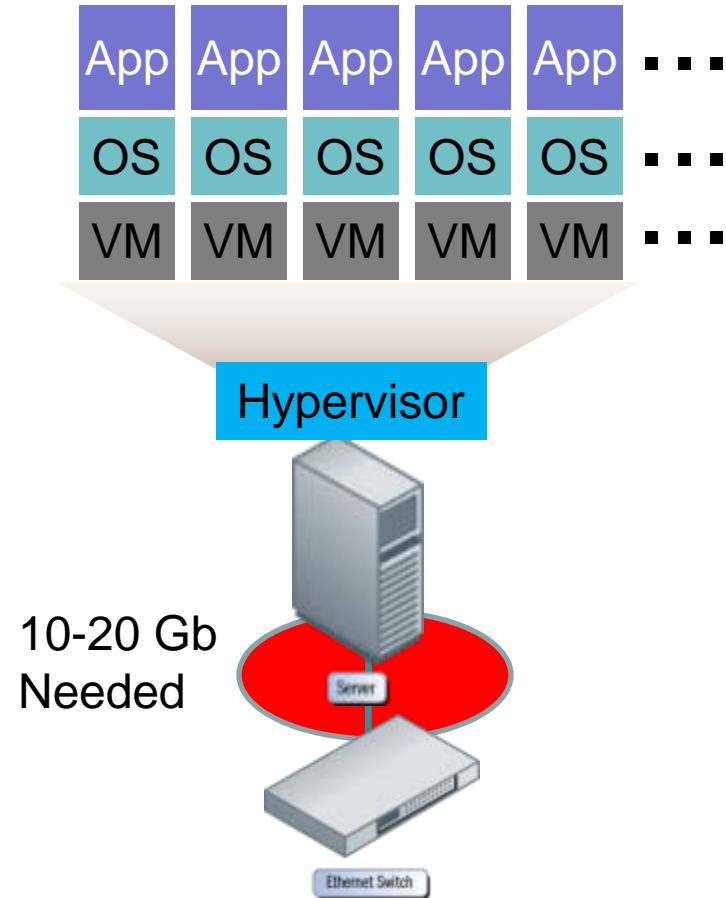
- ◆ PCIe allows server architectures to support > 50Gbps
- ◆ Increased density with virtualization

Virtualization Drives Bandwidth

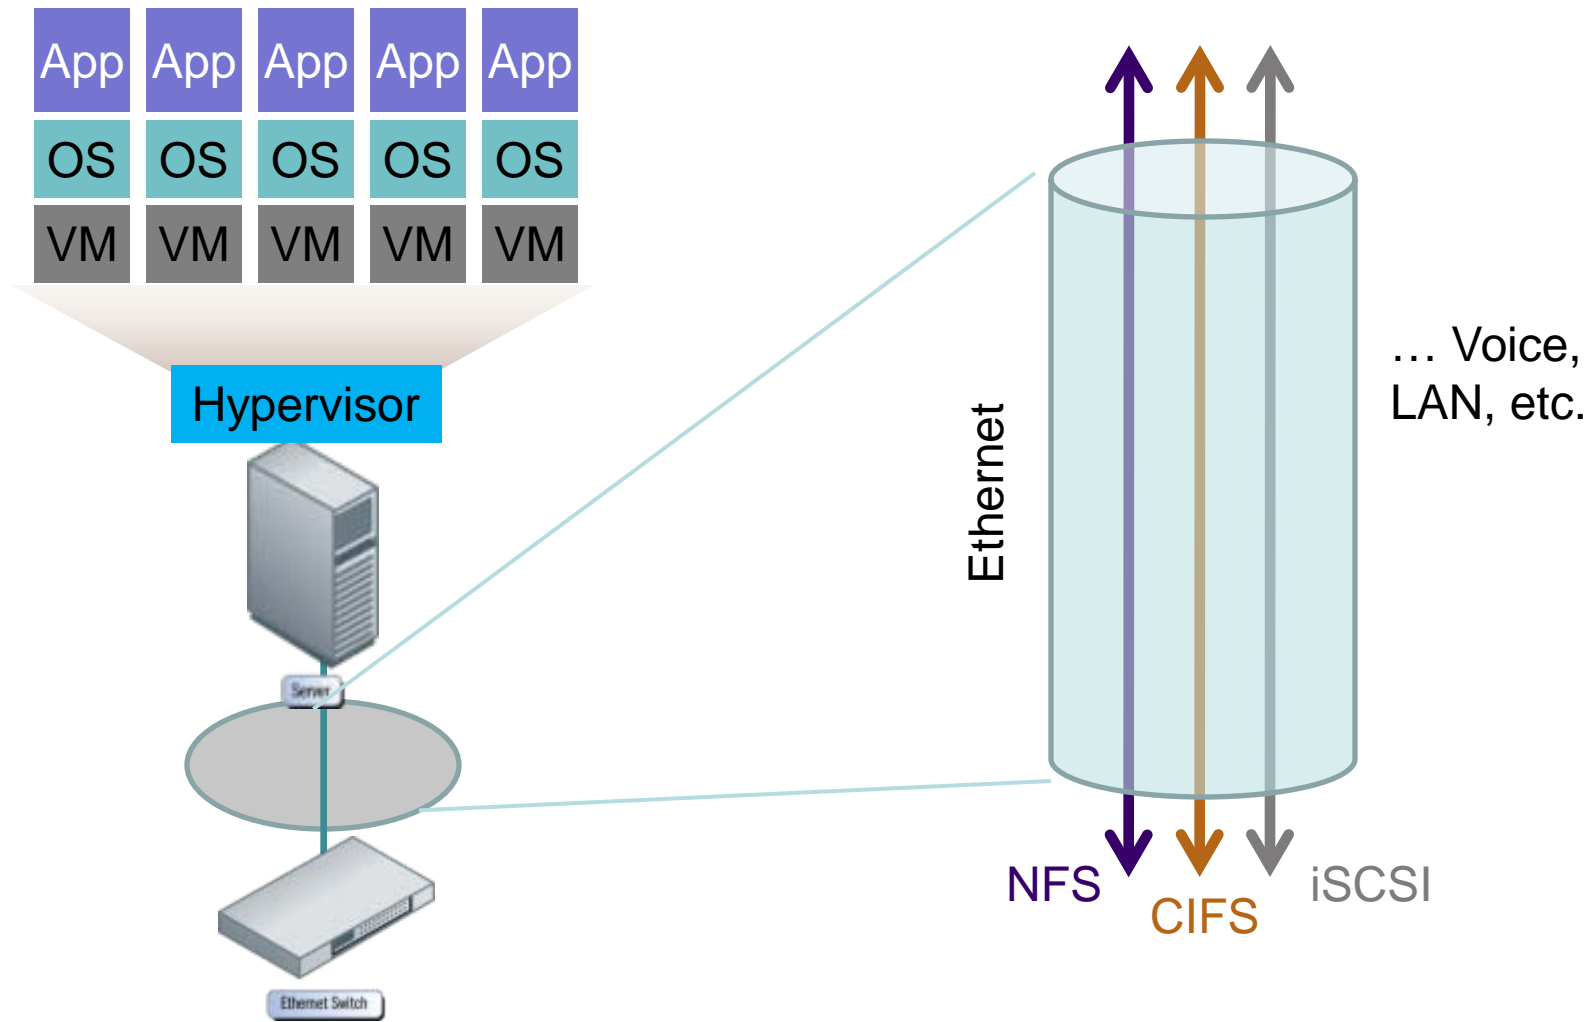
From



To



Ethernet Flexibility – Shared Network



Virtualization

- VMs requires more network capacity
- 80% CPU increases demand on bandwidth
- 10GbE & iSCSI offers the bandwidth needed today
- 10GbE w/ DCB enables enhanced fabric virtualization by offering enhanced manageability and QoS


Consolidation

- Enabling server virtualization allows for increasing levels of server consolidation
- Unifying LAN & SAN consolidates network fabric
- Enables reductions in HW, power and OpEX

Application Performance

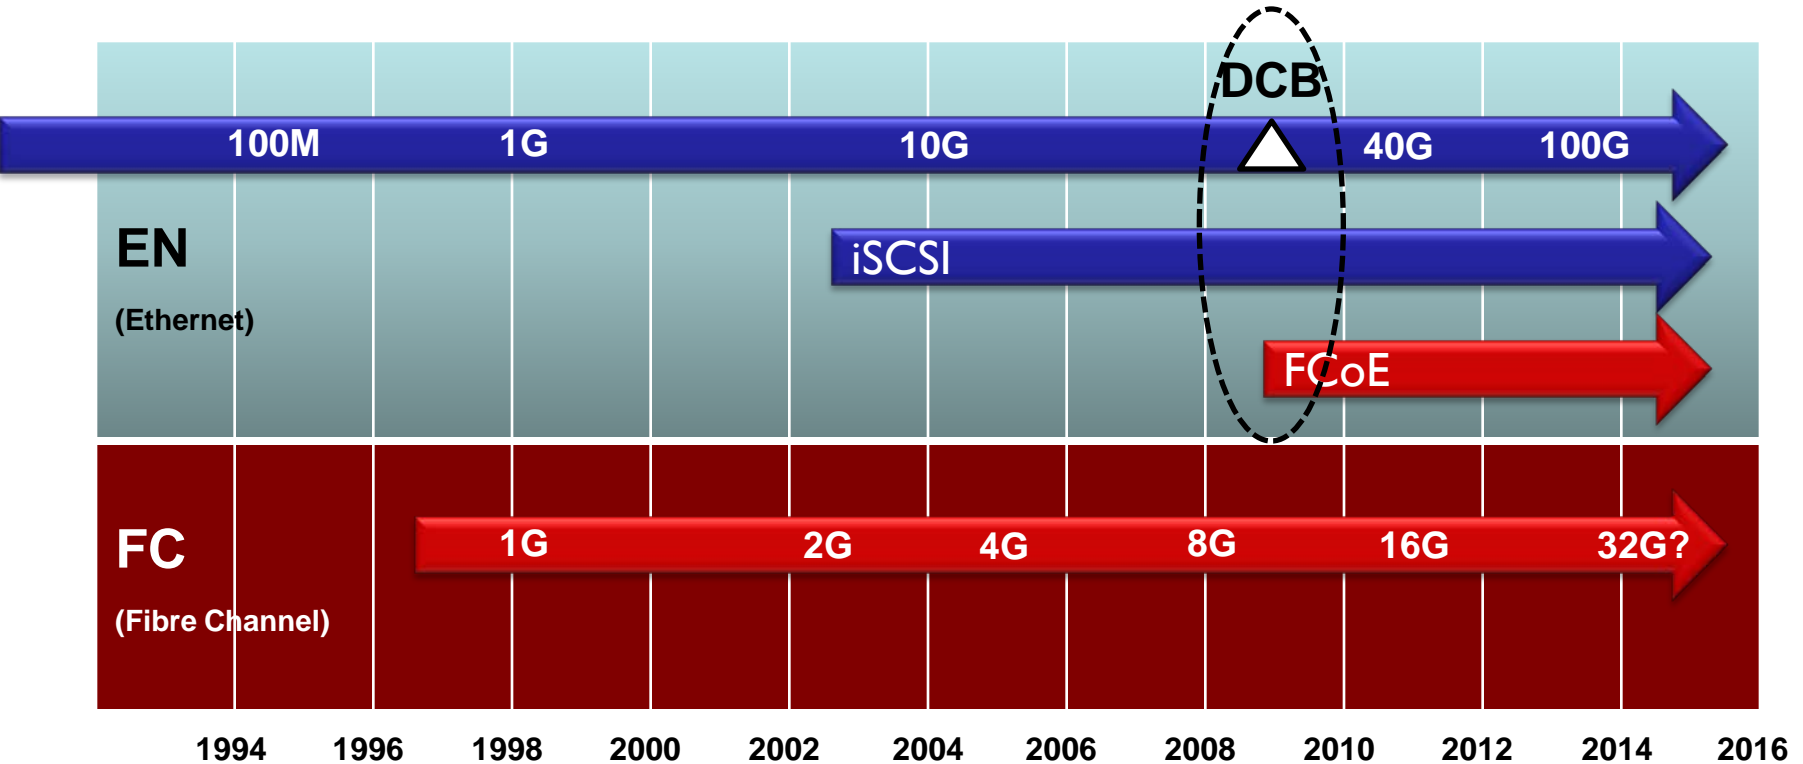
- Instant access to data anytime anywhere
- Larger database, streaming media and other high I/O applications demand data center performance
- Enables thin client and desktop virtualization

Fabric Roadmaps



- Infrastructure migrates to Ethernet irrespective of storage protocols
- Convergence and virtualization are driving enhancements to Ethernet

Fabric Roadmaps



- Infrastructure migrates to Ethernet irrespective of storage protocols
- Convergence and virtualization are driving enhancements to Ethernet

Data Center Bridging

➤ Enhancements to Ethernet


- ◆ Provides enhanced QoS support to Ethernet
- ◆ Not just storage, but all applications

➤ What constitutes DCB Standards?

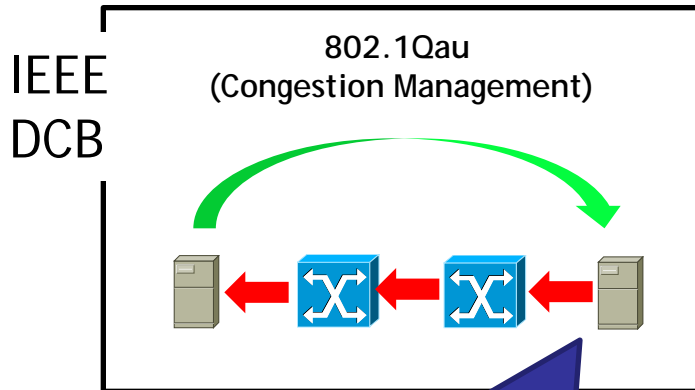
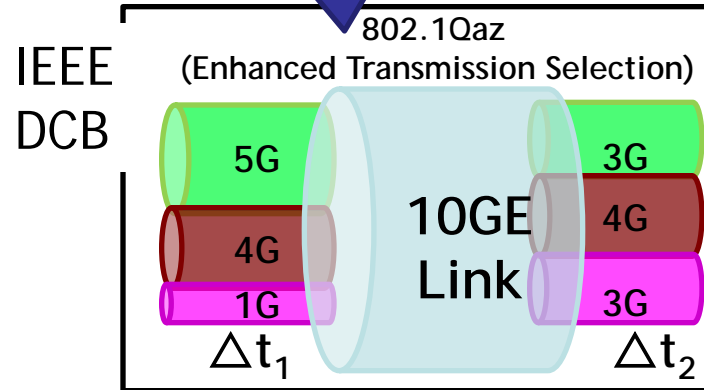
- ◆ PFC aka Priority based Flow Control (802.1Qbb)
- ◆ ETS aka Enhanced Transmission Selection (802.1Qaz)
- ◆ CN aka Congestion Notification (802.1Qau)
- ◆ DCBX aka Data Center Bridging capability eXchange
 - › LLDP vs. DCBX
 - LLDP: Primarily a link level information exchange protocol
 - DCBX: Neighbors can configure parameters based on info exchange and state machine

DCB Components

HALT an individual lane, but NOT all of them!

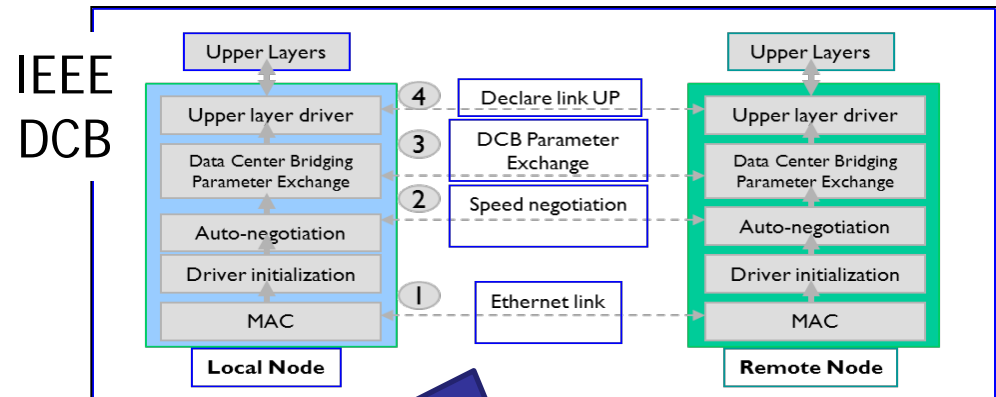


Allocate bandwidth based upon predetermined classes of traffic



End-to-End Communication between end-points. Tells the end-point to BACK OFF!

Data Center Bridging Exchange




End-nodes exchange DCB capabilities

Priority-based Flow Control

IEEE 802.1Qbb – PFC


- **During periods of heavy congestion**
 - Ensures delivery of critical data without loss
 - Latency sensitive traffic continues normal operation



Enhanced Transmission Selection

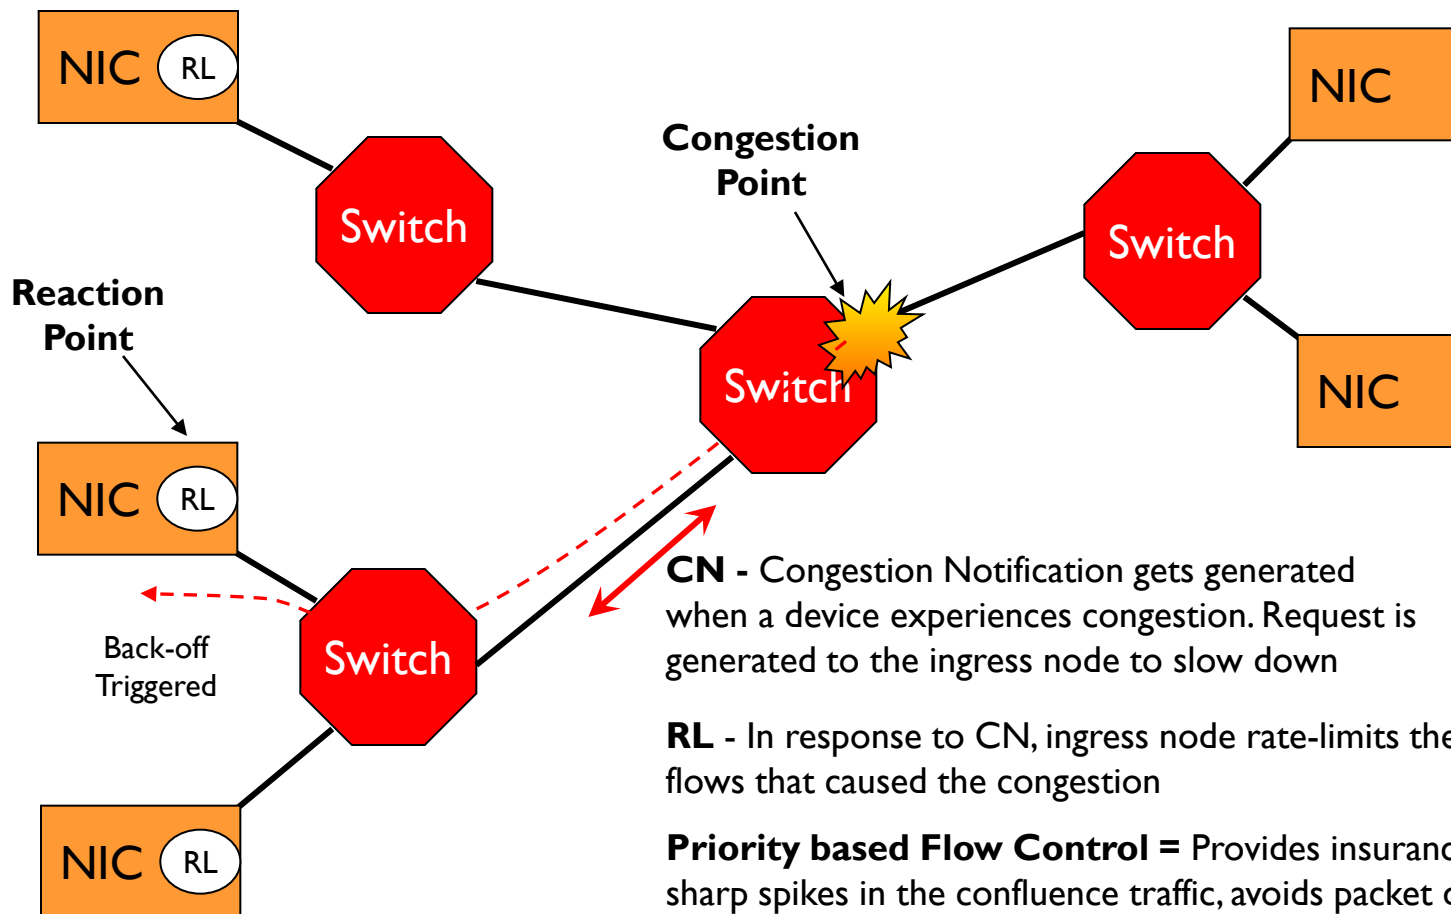
IEEE 802.1Qaz – ETS

- Capability to apply differentiated treatment to different traffic within the same traffic class enabled by ETS



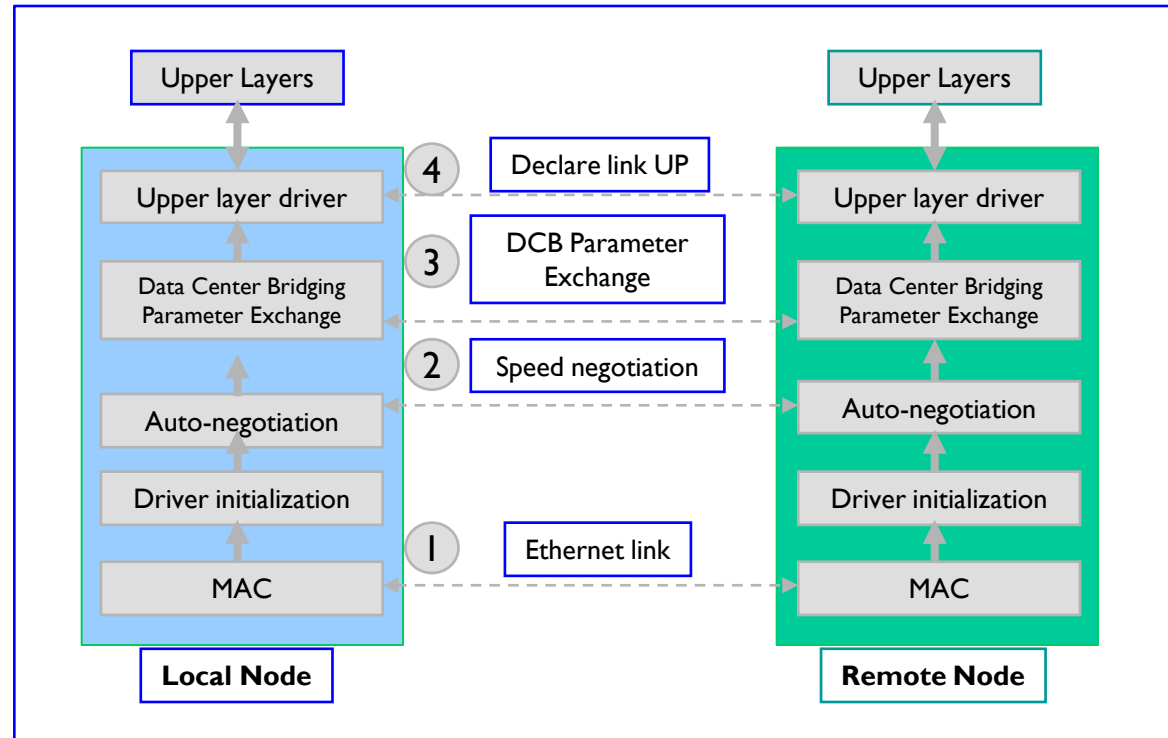
Congestion Notification

(IEEE 802.1Qau)



Discovery & Negotiating Capabilities


- DCBX allows two sides to exchange PFC and ETS attributes
- It also allows exchange of Application capabilities like iSCSI
- It uses LLDP as transport to exchange the above capabilities
- DCBX only works for point-to-point interfaces




Transparent configuration of end-nodes

iSCSI performance at 10GbE

Read/Write IOPs and Throughput Test



Read/Write IOPs and CPU Test



1,030,000 IOPs

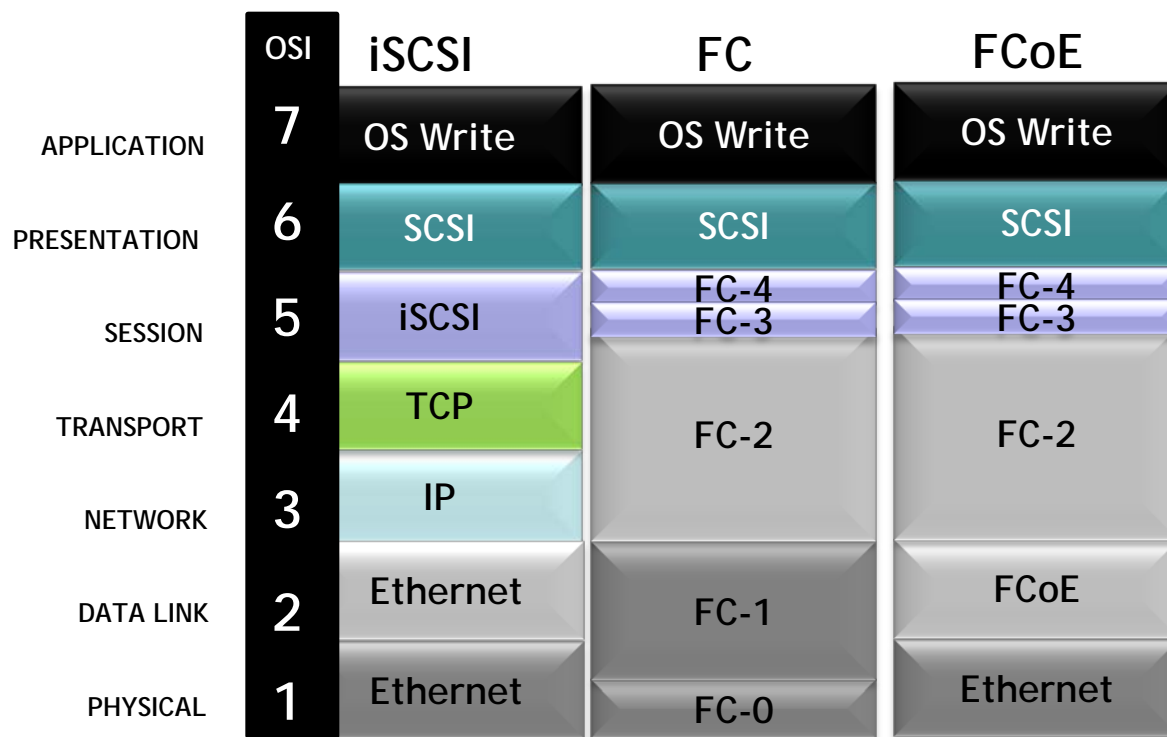
- Single Port
- 10GbE line rate
- 10k IOPs per CPU point
- Performance for real world apps
- Future ready: Performance Scales

552k IOPs at 4k represents

- 3,100 Hard Disk Drives
- 400x a demanding database workload
- 1.7m Exchange mailboxes
- 9x transactions of large eTailers
- Jumbo frames: >30% CPU decrease is common for larger IO size (jumbo frames not used here)

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Microsoft and Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing.

iSCSI: A Lossless Ethernet Fabric with DCB



Fibre Channel is a layered protocol. It consists of 5 layers, namely:

- FC0 The physical layer, which includes cables, fiber optics, [connectors](#), [pinouts](#) etc.
- FC1 The data link layer, which implements the [8b/10b encoding](#) and decoding of signals.
- FC2 The network layer, defined by the FC-PI-2 standard, consists of the core of Fibre Channel, and defines the main [protocols](#).
- FC3 The common services layer, a thin layer that could eventually implement functions like encryption or RAID.
- FC4 The Protocol Mapping layer. Layer in which other protocols, such as SCSI, are encapsulated into an information unit for delivery to FC2.

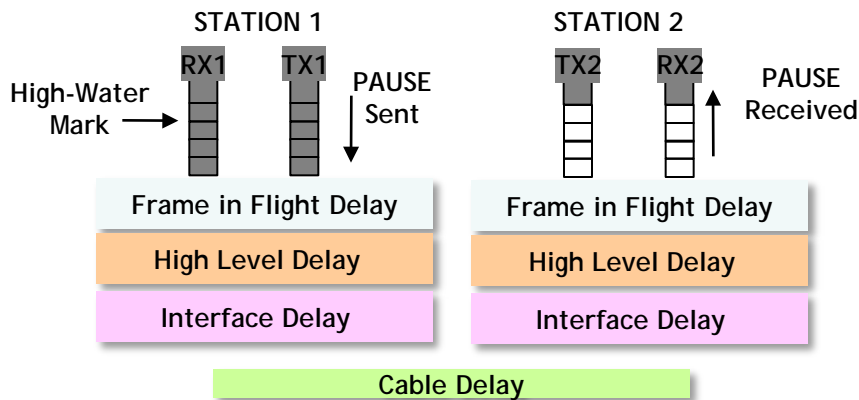
Comparing iSCSI To FCoE

Feature	iSCSI	FCoE
Routable (across LANs and WANs)	Layer 3 & 4	Layer 2 ONLY
MTU	1500 But 9K (Jumbo) Preferred	2500 (Baby Jumbo Frame)
Works over Standard Ethernet	YES	NO
Requires CEE/DCB	No (but equally benefits)	YES in a converged network
Leverages FC Management Tools	NA	Yes
Requires FC Gateway for Legacy FC	NA	Yes
Flow Control Type	Rate based	Rate based

FLOW Control Types

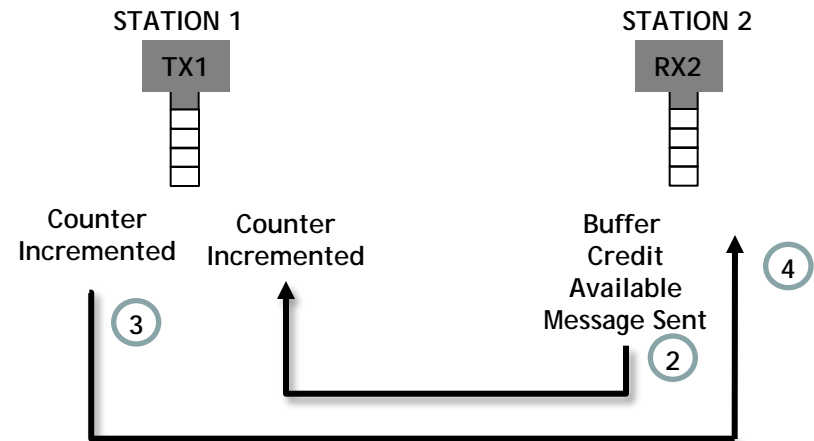
Reactive-Time/Link Dependent

PAUSE



Proactive-Time/Link Independent

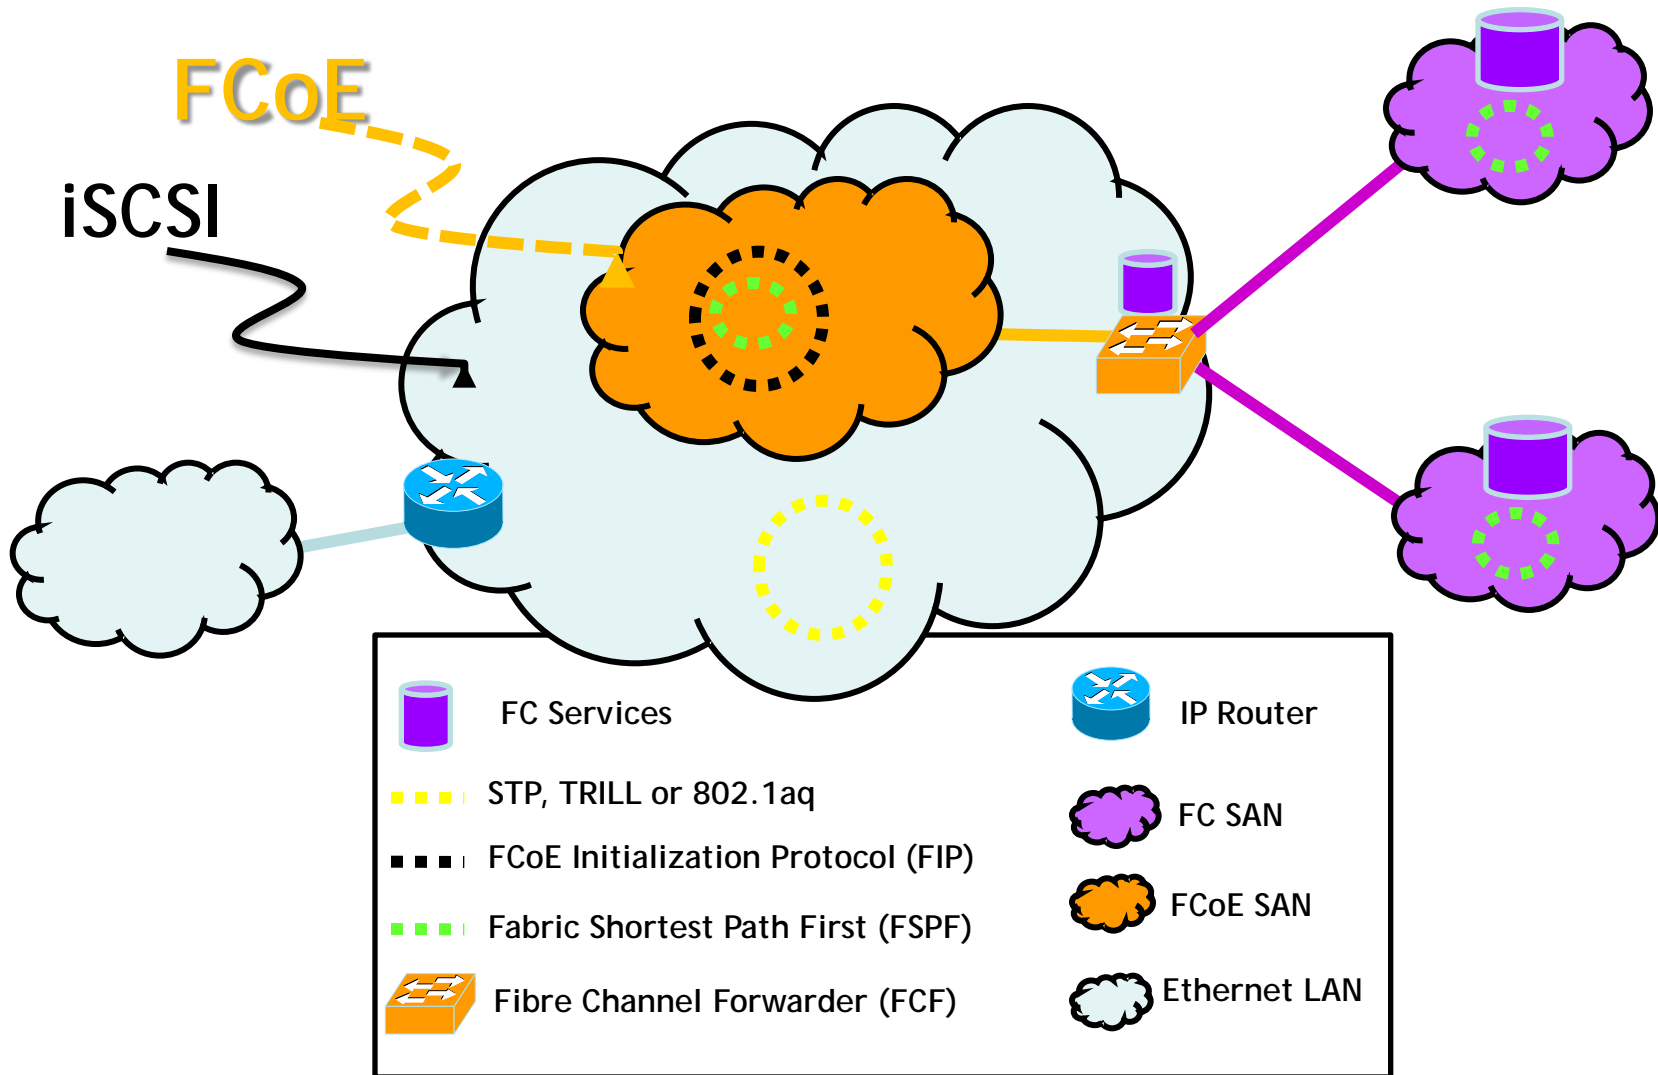
CREDITS



① Ex: Buffer_Credit_CNT= 4

- PAUSE: Pause command sent to sender to halt I/O traffic to avoid overflow of receiving buffers
- Buffer-Credit: Source and destination negotiate credit count to regulate I/O traffic flow

Storage Network Components



Is DCB relevant to iSCSI?

Improves application responsiveness

Provides QoS improvements in mixed traffic environments

Increases robust of Ethernet as a storage fabric

Improves performance of some IP Protocols including iSCSI

Questions?



- Please send any questions or comments on this presentation to SNIA: tracknetworking@snia.org

**Many thanks to the following individuals
for their contributions to this tutorial.**

- SNIA Education Committee

**Jason Blosil
Gary Gumanow
Jordan Plawner**