



Education

ETHERNET ENHANCEMENTS FOR STORAGE

Sunil Ahluwalia, Intel Corporation

- 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.

➤ Ethernet Enhancements for Storage

This session discusses the Ethernet enhancements required for storage traffic. It reviews an end-to-end view to evaluate FCoE benefits from a host and switch perspective.

- Ethernet Everywhere!

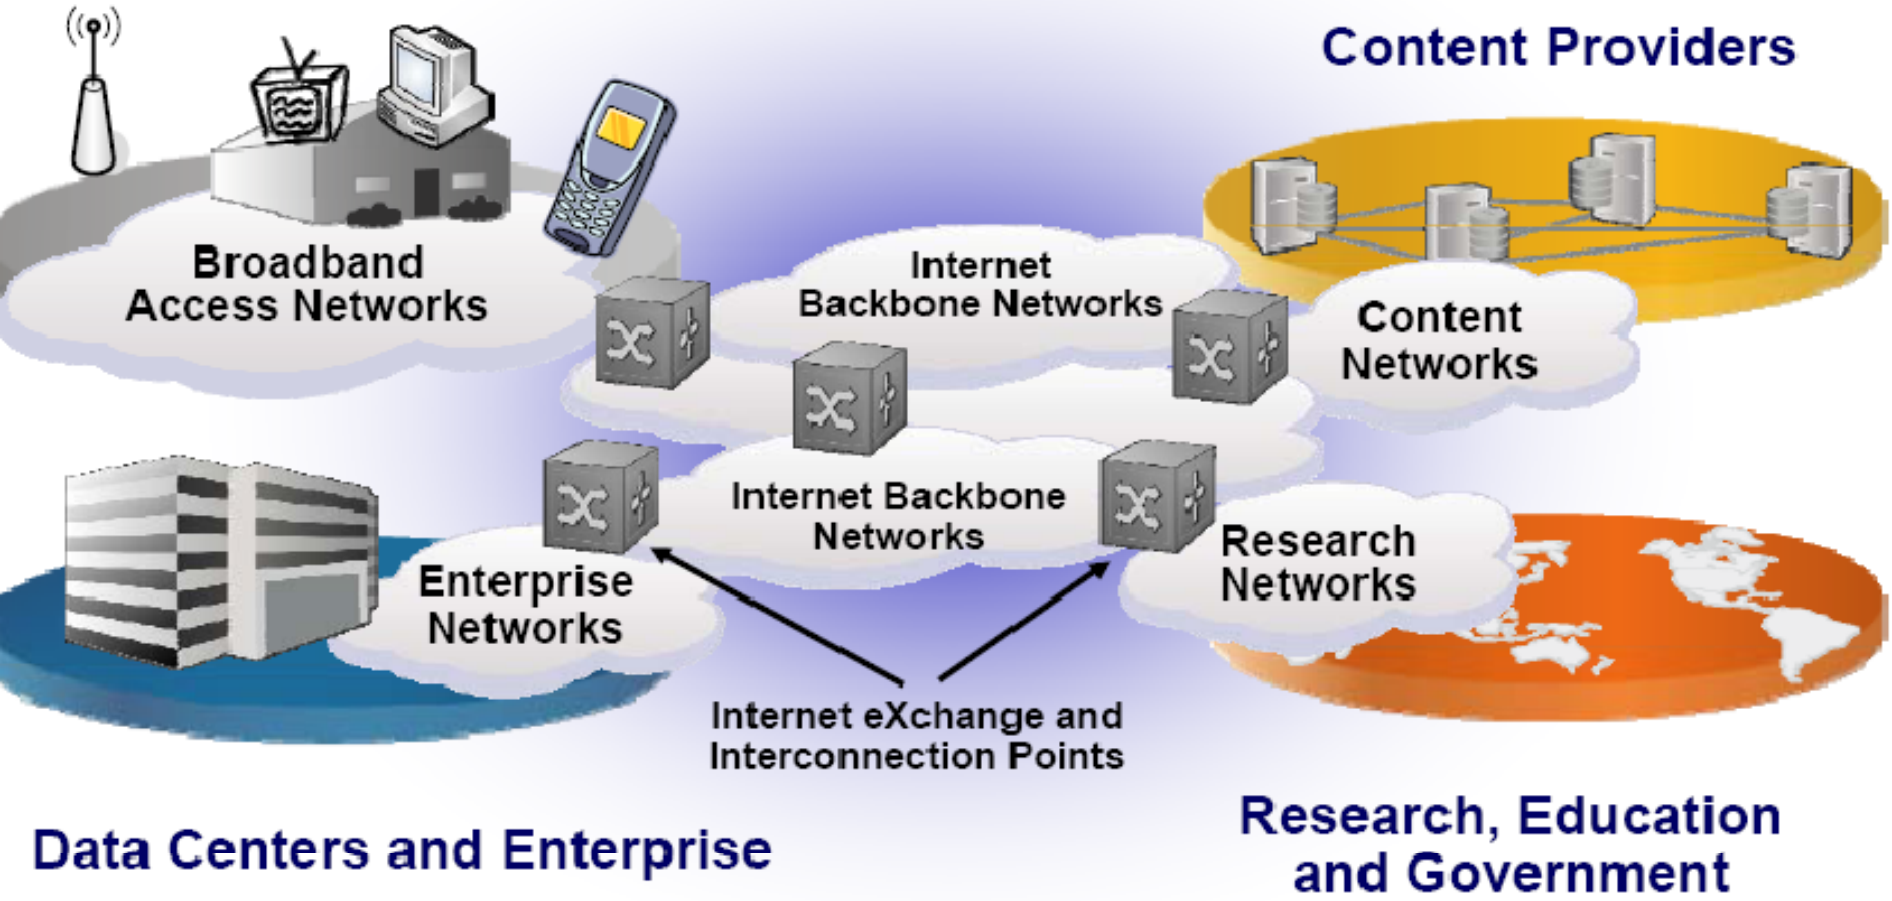
- Data Center Requirements

- Ethernet Enhancements
 - ◆ Data Center Bridging

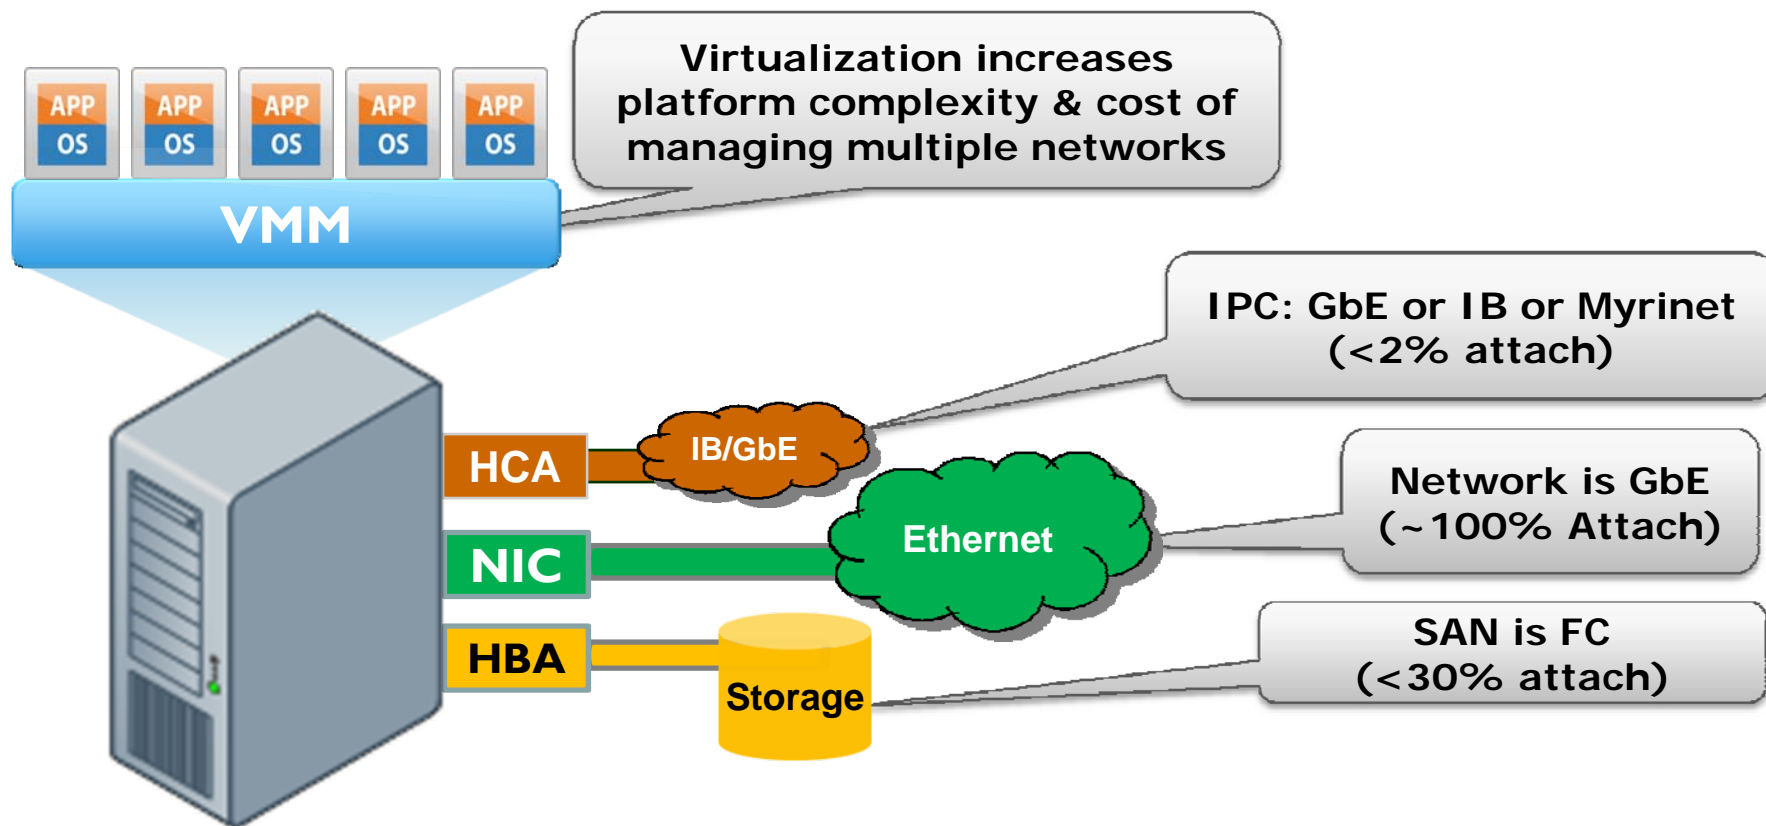
- FCoE Deployment

Ethernet Everywhere!

Broadband Access



Nearly all of the traffic on the Internet either originates or terminates with an Ethernet connection



➤ Multiple networks, one per traffic class

- ◆ IPC over an InfiniBand network
- ◆ IP and other LAN protocols over an Ethernet network
- ◆ SAN over a Fibre Channel network

LAN / IP

- Must be Ethernet!
 - Too much investment
 - Too many applications that assume Ethernet
 - Pervasive LAN technology

Storage

- FC SAN implementations
 - lossless requirement
- IP SAN assumes IP and Ethernet with IP recovery mechanisms

IPC

(Inter-Process Communication)

- Transparent to underlying network, provided that
 - It is cheap
 - It is low latency
 - It supports APIs like OFED, MPI, sockets

Ethernet Enhancements [Data Center Bridging]

- **Data Center Bridging (DCB)** is an architectural collection of Ethernet extensions designed to improve Ethernet networking and management in the Data Center.
- Sometimes also called
 - ◆ CEE = Converged Enhanced Ethernet
 - ◆ DCE = Data Center Ethernet (Cisco Trademark)
 - ◆ EEDC = Enhanced Ethernet for Data Centers

- Effort underway to provide DC enhancements in IEEE
 - ◆ 25+ companies actively working on standards in IEEE
 - ◆ Work is called Data Center Bridging (DCB)
 - ◆ <http://www.ieee802.org/1/pages/dcbridges.html>
- IEEE projects necessary for I/O Consolidation in Data Center
 - ◆ Congestion Notification: Approved project IEEE 802.1Qau
 - ◆ Shortest Path Bridging: Approved project IEEE 802.1aq
 - ◆ Enhanced Transmission Selection: Approved project IEEE 802.1Qaz
 - ◆ Priority based Flow Control: Approved project in IEEE 802.1Qbb
 - ◆ DCB Capability Exchange Protocol: Part of various projects above
- DCB Standards trending for ratification in 2010

Ethernet Enhancements for Data Centers

➤ Traffic Differentiation

- ◆ Provides end-to-end traffic differentiation for LAN, SAN and IPC traffic

➤ “Lossless” Fabric

- ◆ Transient congestion - Priority Based Flow Control
- ◆ Persistent congestion - Congestion Notification

➤ Optimal Bridging

- ◆ Allow shortest path bridging within Data Center
- ◆ Eliminates the need to shut off links to prevent loops

➤ Configuration management

- ◆ Exchange parameters and work with legacy systems

Ethernet Enhancements [Data Center Bridging]

- Traffic Differentiation**
- “Lossless” Fabric**
- Optimal Bridging**
- Configuration Management**

➤ Link Sharing (Transmit)

- ◆ Different traffic types may share same queues/links
- ◆ Large burst from one traffic should not affect other traffic types


➤ Resource Sharing

- ◆ Different traffic types may share resources (buffers)
- ◆ Large queued traffic for one traffic type should not starve other traffic types out of resources

➤ Receive Handling

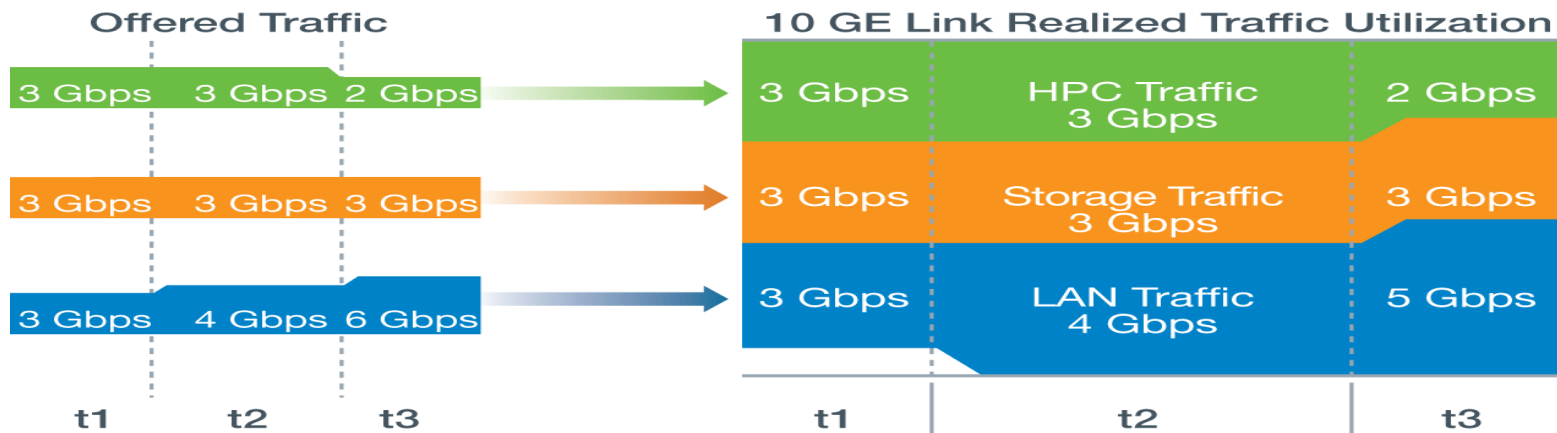
- ◆ Different traffic types may need different receive handling (eg. interrupt moderation)
- ◆ Optimisation for CPU utilisation for one traffic type should not create large latency for small messages for other traffic types

Traffic Differentiation



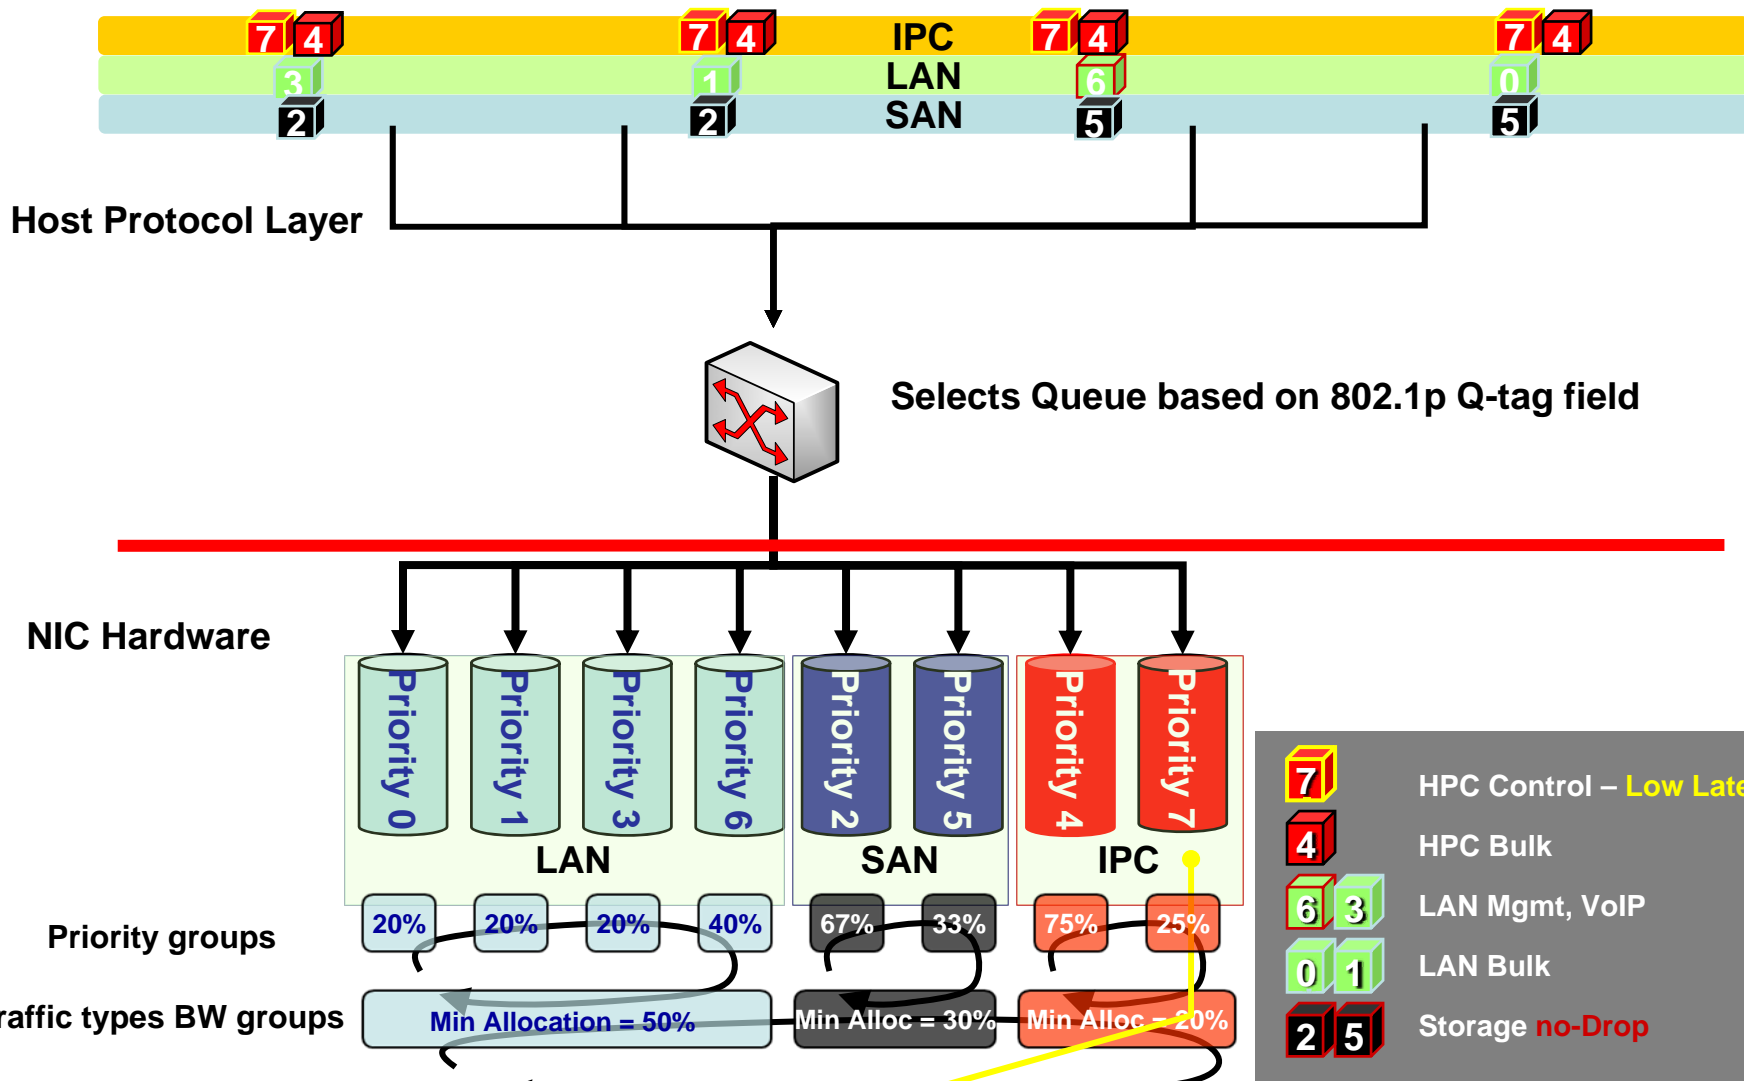
- Multiple Link Partitions, One per traffic class
- Resource allocation and association
- Provisioning “aggregate flow bundles”

Priority based Bandwidth Management



Enables Intelligent sharing of bandwidth between traffic classes control of bandwidth

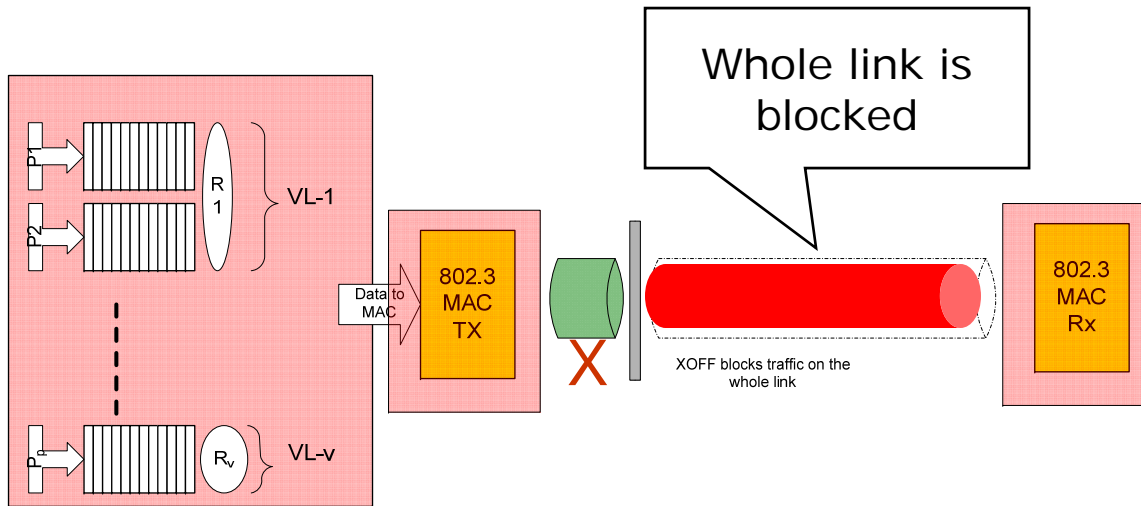
Packet Flow



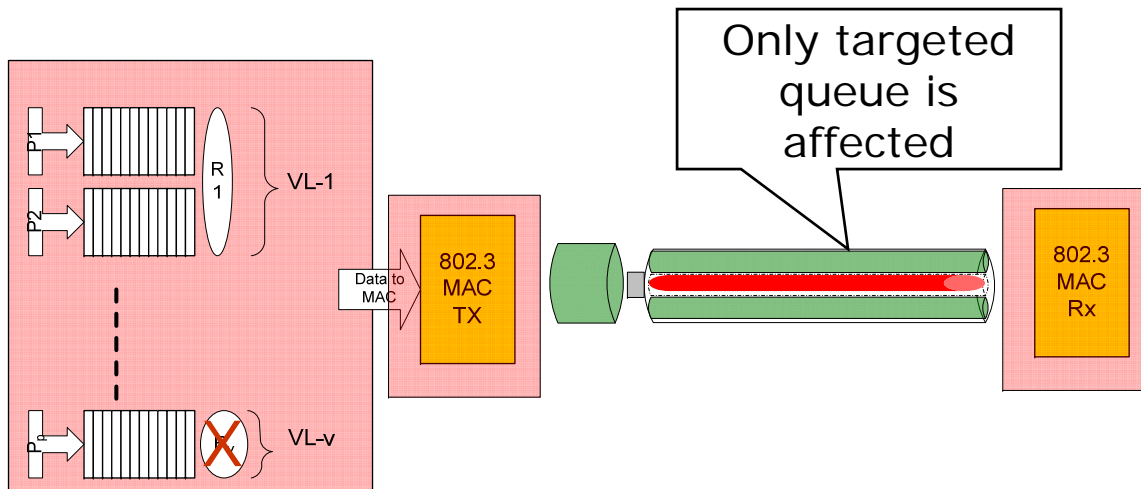
Ethernet Enhancements [Data Center Bridging]

- Traffic Differentiation**
- “Lossless” Fabric**
- Optimal Bridging**
- Configuration Management**

Priority-based Flow Control (IEEE 802.1Qbb)



Link
Pause




Granular
Pause

PFC and BB_Credits

- IEEE 802.3x Pause provides no drop flow control
 - ◆ similar to BB credits for FC
- Priority Flow Control is a finer grained mechanism of flow control over standard pause or link level BB credits
- Priority Flow Control uses .Ip CoS value mapping to a system class to send appropriate pause to previous hop
- The Pause frame is handled by the MAC layer
 - ◆ Similar to the R_RDY handling by the FC-I level
- The BB_Credit mechanism allows to not lose frames over any link
 - ◆ Under-utilizing a link if the credits are not enough
 - ◆ Requiring to handle the buffer in maximum frame size units

Congestion Notification (IEEE 802.1Qau)



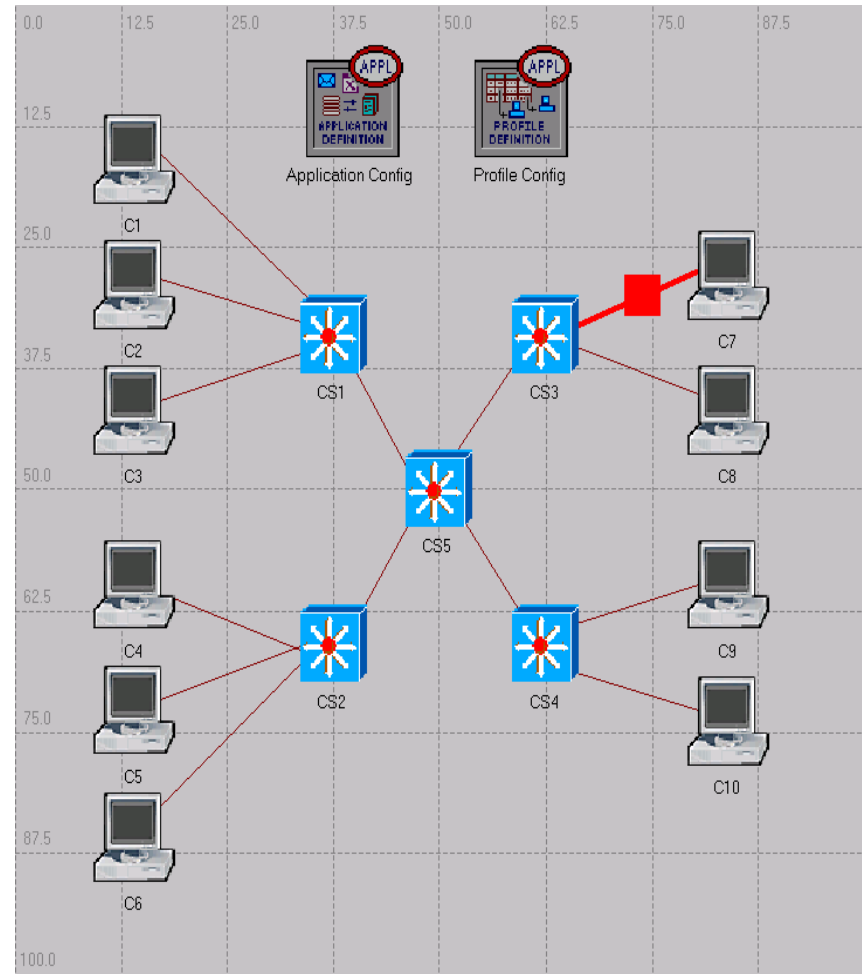
CN - Congestion Notification gets generated when a device experiences congestion. Request is generated to the ingress node to slow down

RL - In response to CN, ingress node rate-limits the flows that caused the congestion

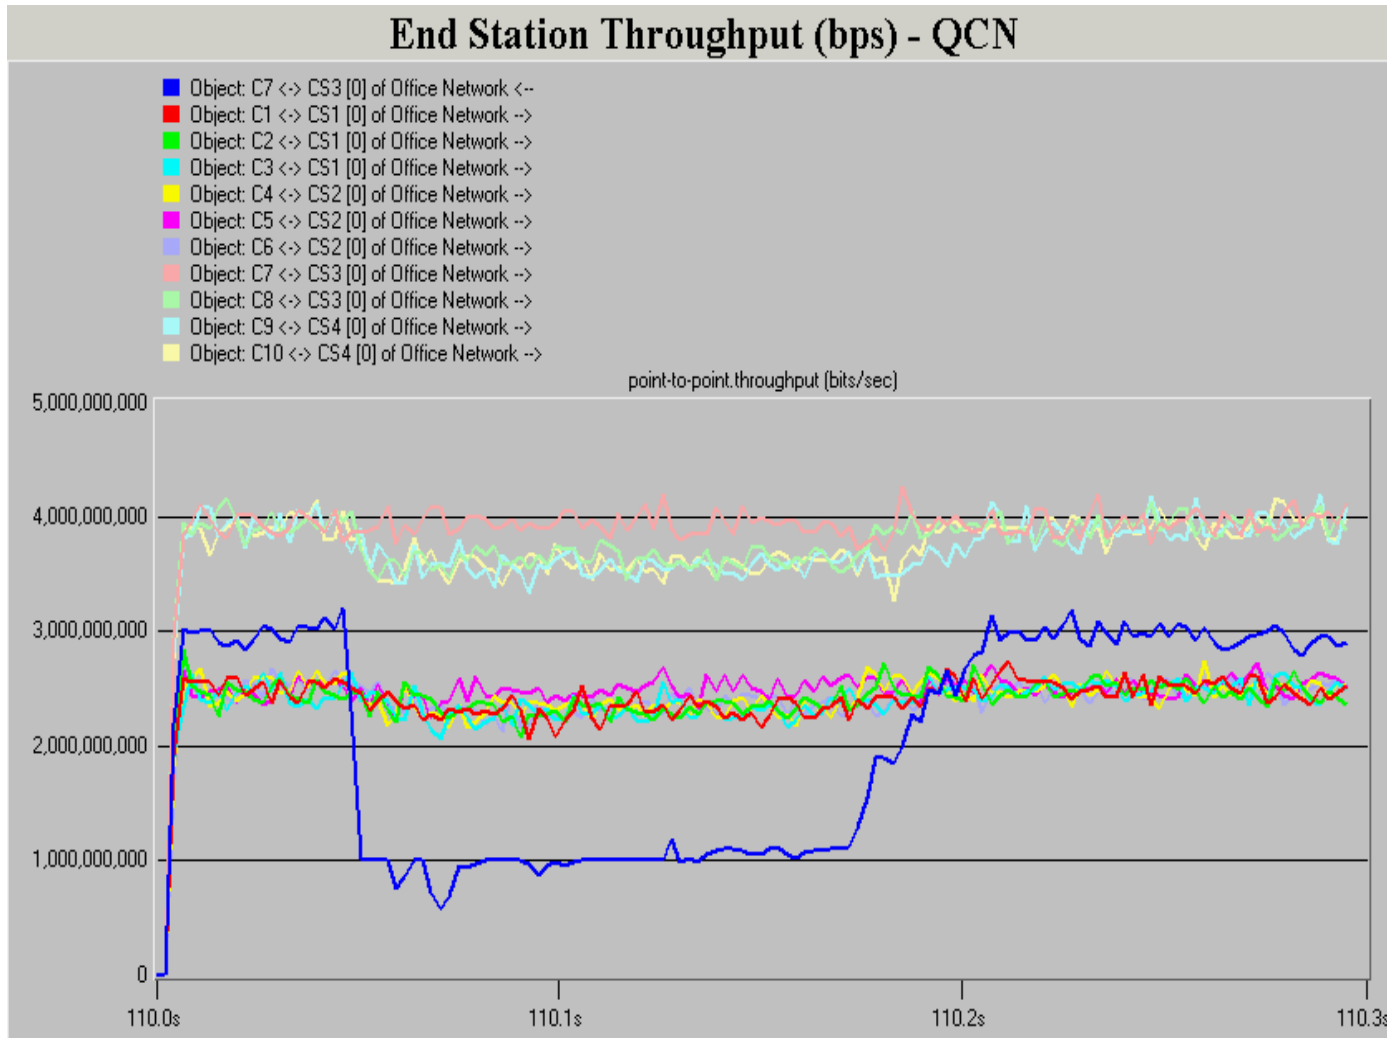
Priority based Flow Control = Provides insurance against sharp spikes in the confluence traffic, avoids packet drops

Simulation: Topology & Workload

- 1 Gbps OG hotspot for 80 ms @ CS3→C7
- **Traffic pattern:**
 - ◆ 10 Gbps links, 20 μ s loop latency
 - ◆ Nodes 1-6 @ 25% (2.5 Gbps)
 - ◆ Nodes 7-10 @ 40% (4 Gbps)
 - ◆ Spatially uniform (except self)
 - ◆ Bernoulli arrival time distribution



Simulation: Throughput results (QCN)



➤ PFC is good for transient congestion

- ◆ Reacts to avoid packet loss but doesn't diminish congestion.
- ◆ Congestion spreads upstream which can affect innocent flows.
- ◆ Unfair to low-priority latencies when higher priorities are culprit flows.

➤ QCN works on persistent congestion


- ◆ Rates reduced at source to eliminate congestion.
- ◆ Increased aggregate throughput.
- ◆ Improves Fairness
- ◆ Reduced egress buffer usage limits congestion spreading.

Ethernet Enhancements [Data Center Bridging]

- Traffic Differentiation**
- “Lossless” Fabric**
- Optimal Bridging**
- Configuration Management**

Why?

- Spanning Tree Protocol (STP) and its variants have a bad reputation with customers
 - ◆ Non-optimal forwarding
 - ◆ Parallel paths cannot be leveraged
- These problems can be solved at L3
- But L3 cannot be deployed in many scenarios such as clusters, metro Ethernet, virtualized servers (VM's) etc,



Shortest Path Bridging - 802.1aq

Another Approach to Shortest Path Bridging

➤ What is it

- ◆ Enhancement to 802.1Q to provide Shortest Path Bridging (Optimal Bridging) in L2 Ethernet topologies
- ◆ Provides for each bridge to be the root of its own topology and hence uses the “best” path to any destination

➤ Benefits

- ◆ Resolves issues related to root disappearance
- ◆ Fast convergence – no count to infinity

➤ Does not require a link state protocol


➤ Resources

<http://www.ieee802.org/1/files/public/docs2005/aq-nfinn-shortest-path-0905.pdf>

<http://www.ieee802.org/1/files/public/docs2006/aq-nfinn-shortest-path-2-0106.pdf>

802.1aq - Spanning Tree per bridge

How does it work



Each bridge is the root of a separate spanning tree instance.

Bridge G is the root of the green tree

Bridge E is the root of the blue tree

Both trees are active at all times

“TRILL” WG in IETF

➤ What is it

- ◆ “Transparent Interconnection of Lots of Links” Internet Drafts (also called “Routing Bridges” or “RBridges”)
- ◆ IETF effort to solve L2 STP forwarding limitations
- ◆ TRILL is a solution intended for data centers (and campuses) to provide connectivity among end stations with ease of current bridges but without using spanning tree protocol
- ◆ Replaces STP with a link-state routing protocol to discover the topology

➤ Benefits

- ◆ Shortest-Path Frame routing in multi-hop 802.1-compliant networks
- ◆ Permits Load Splitting among multiple paths
- ◆ Forwarding based on destination bridge-id – smaller tables than conventional bridge systems

➤ Must be backward compatible with 802.1d – inter-working at the edge


➤ Resources

- ◆ <http://www.ietf.org/html.charters/trill-charter.html>
- ◆ <http://www.ietf.org/internet-drafts/draft-ietf-trill-rbridge-protocol-03.txt>

Ethernet Enhancements [Data Center Bridging]


- Traffic Differentiation**
- “Lossless” Fabric**
- Optimal Bridging**
- Configuration Management**

Configuration Management



- **Link level capability and configuration exchange**
 - ◆ Similar to FLOGI and PLOGI in Fibre Channel
 - ◆ Allows either full configuration or configuration checking
- **Based on LLDP (Link Level Discovery Protocol)**
- **DCBX TLVs**
 - ◆ Priority Groups, Priority-based Flow Control, Congestion Notification
 - ◆ Application (frame priority usage)

Example DCBX Deployment Model




- Detects configuration mismatches between link peers and notifies Management
- Discovers DCB related peer capability
- Detect boundaries of congestion management

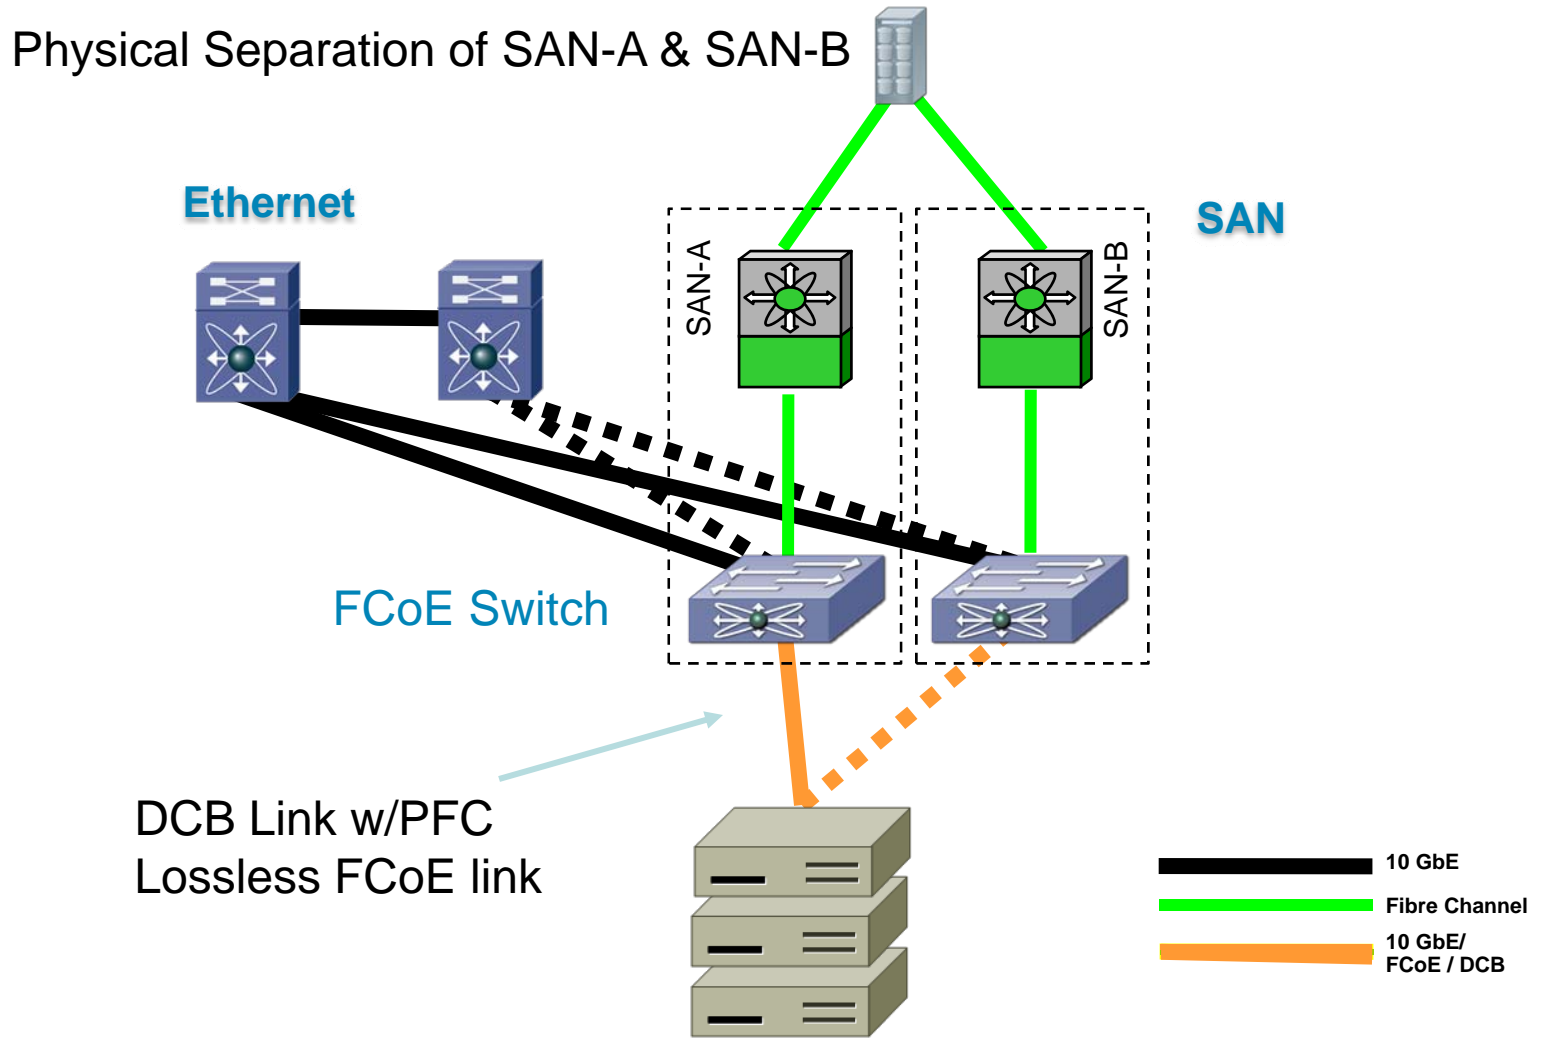
FCoE Deployment

Data Center View: Today

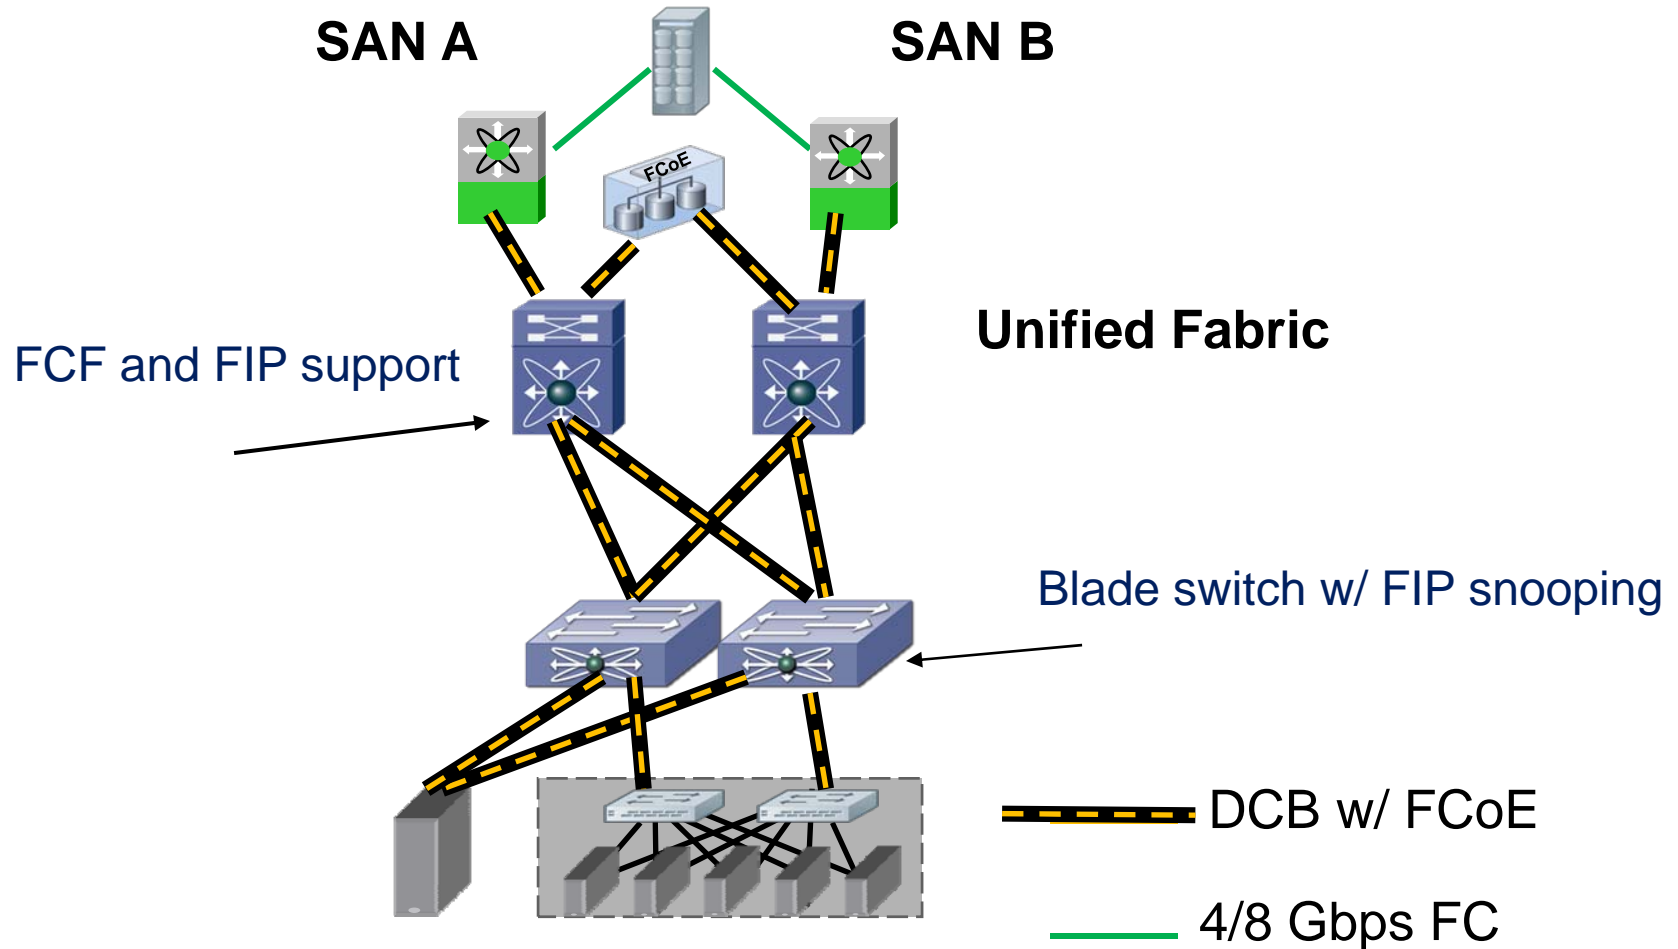
Separate
LAN and **SAN**
Environments



Data Center View: Access Layer Consolidation



Data Center View: Multi Tiered



- Datacenter requirements driving the need for one network fabric
- Ethernet Enhancement for carrying multiple traffic types
 - ◆ Traffic Differentiation
 - ◆ “Lossless” Fabric
 - ◆ Optimal Bridging
 - ◆ Configuration Management
- FCoE is the first big use case of these Ethernet Enhancements

- 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

**Rob Peglar
Walter Dey
Steve Wilson
Joe White**