



Education

# Design and Implementations of FCoE for the DataCenter

Mike Frase, Cisco Systems

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- 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.**

## Design and Implementations of FCoE for the DataCenter

This session will discuss the technologies and solutions for implementing FCoE in an end to end design for the Unified Data Center. The terms and technical requirements of the 10 GigE Network along with Multi-hop FCoE and Unified I/O Compute will be explained, allowing anyone new or old to the technology the ability to engage in the discussion. We will focus on what the Data Center Engineer can deploy in today's designs, intergrading FCoE with current Storage Area Networks (Native Fibre Channel) or complete new FCoE deployment. Attendees shall learn what new skills are required for the implementation and management of an FCoE virtualized network.


### Learning Objectives:

- ◆ - FCoE technologies
- ◆ - Skills requirements for Unified I/O Networks
- ◆ - 10 GigE designs for FCoE

- **Network, Storage, and Server Admin -  
Positioning your skills**
- **FCoE Technology**
- **FCoE Deployment Designs – Today**
- **FCoE Deployments Designs - Tomorrow**


## ➤ Network Data Engineer

- ◆ 1, 10 GigE
- ◆ Spanning-Tree Protocols
- ◆ Highly Available Networks
- ◆ Low Latency Switching
- ◆ Layer 2 & 3
- ◆ Network adaptors and drivers
- ◆ IP addressing IPV4 & IPV6
- ◆ IP Routing Protocols
- ◆ Standard Network services ARP, DHCP, PXE, LLDP
- ◆ QOS
- ◆ Security



## ➤ SAN Engineer


- ◆ 2, 4, 8 Gig FC
- ◆ Fibre Channel
- ◆ FSPF
- ◆ Highly Available Fabrics
- ◆ Low Latency Switching
- ◆ Fabric Login Services
- ◆ FCID
- ◆ Domain Services
- ◆ Zone Service
- ◆ Storage
- ◆ HBA's and drivers



## ➤ Data Center Engineer

- ◆ LAN Engineer
- ◆ SAN Engineer
- ◆ DCB (Lossless) GigE
- ◆ CNA adaptors & drivers

Technologies for LAN & SAN  
iSCSI  
FCoE

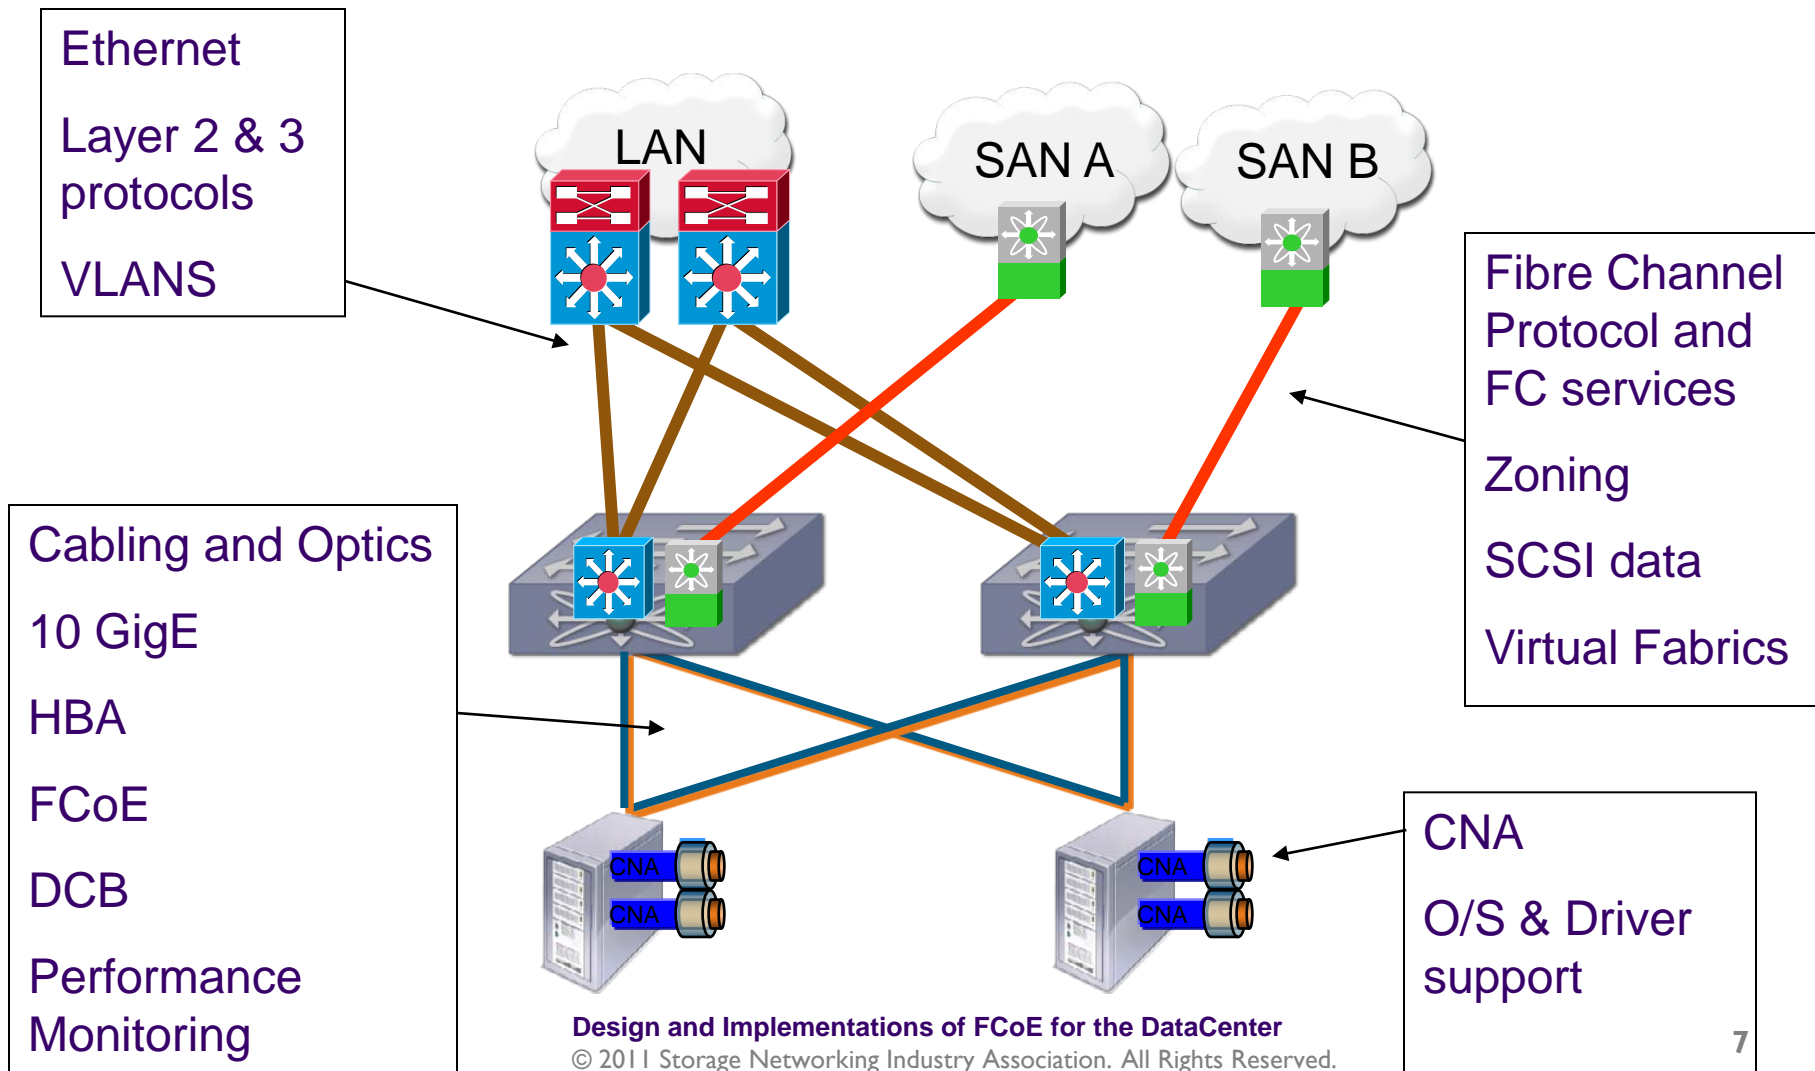


## ➤ Add these extras for **Advanced DC Engineer**

- ◆ Virtualization
- ◆ Server Hardware (Memory, CPU, PCIe)
- ◆ Data Center Interconnect Technologies

# Learning Curve

## Unified Network Technologies



# Unified I/O: Network Operations View

## Server Ops

- Work Processes: Same as today  
Converged Network Adapter (CNA)  
config equivalent to NIC and HBA.
- Mgmt: No impact
- Network Mgmt System:  
Existing tools

## SAN Ops

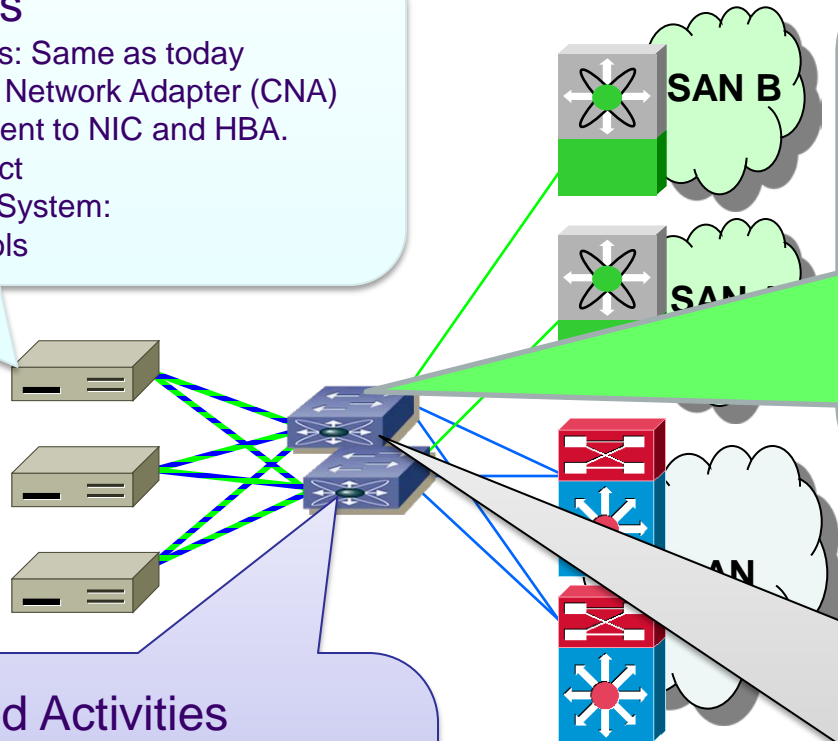
- Work Process: Same as today  
Configure FC Interface mode, speed,  
Virtual Fabrics, Port Channels, Trunks,  
NPV etc.  
Monitor SAN specific MIBs, Syslog,  
RMON.
- Mgmt :  
CLI (Role based)  
XML API
- Network Mgmt System:  
Existing tool – SAN Fabric Managers

## Converged Activities

- Work Processes:  
Switch reboot  
Image upgrade  
User roles  
Initial creation of virtual interfaces
- Mgmt: Creation of Mgmt Interface
- Day-2 System level support: Memory,  
CPU error, Call Home features, Span ports.

## Network Ops

- Work Process: Same as today  
Configure Ethernet port, VLAN, STP,  
Port Channels, Trunks etc.  
Monitor LAN specific MIBs, Syslog,  
RMON.
- Mgmt :  
CLI (Role based)  
XML API
- Network Mgmt System:  
Network/Device Management



— 1 GbE  
— 10 GbE  
— FC  
— 10GbE FCOE



# FCOE TECHNOLOGY

## ➤ Data Center Bridging eXchange (DCBX):

- ◆ **DCBX**, is a management protocol that extends the Link Layer Discovery Protocol (LLDP) defined in IEEE802.1AB. DCBX allows the Fibre Channel Forwarder (FCF) to provide Link Layer configuration information to the Converged Network Adaptor (CNA) and allows both the CNA and FCF to exchange status.

## ➤ FCoE Protocols:

- ◆ **FCoE Ethernet Frame Type 8906 Hex**  
FCoE data protocol requires lossless Ethernet and is typically implemented in hardware, and is used to carry most of the FC frames and all the SCSI traffic.
- ◆ **FIP (FCoE Initialization Protocol) Ethernet Frame Type 8914 Hex**  
A Discovery Protocol whose purpose is to present a multi-access Ethernet network as a set of point-to-point Virtual Links for FCoE to use.

# DataCenter Bridging (DCB) for FCoE


## ➤ Part of the IEEE 802.1 Work Group

- ◆ Lossless Ethernet 802.1Qbb
- ◆ Priority Grouping 802.1Qaz
- ◆ Configuration Verification 802.1Qaz

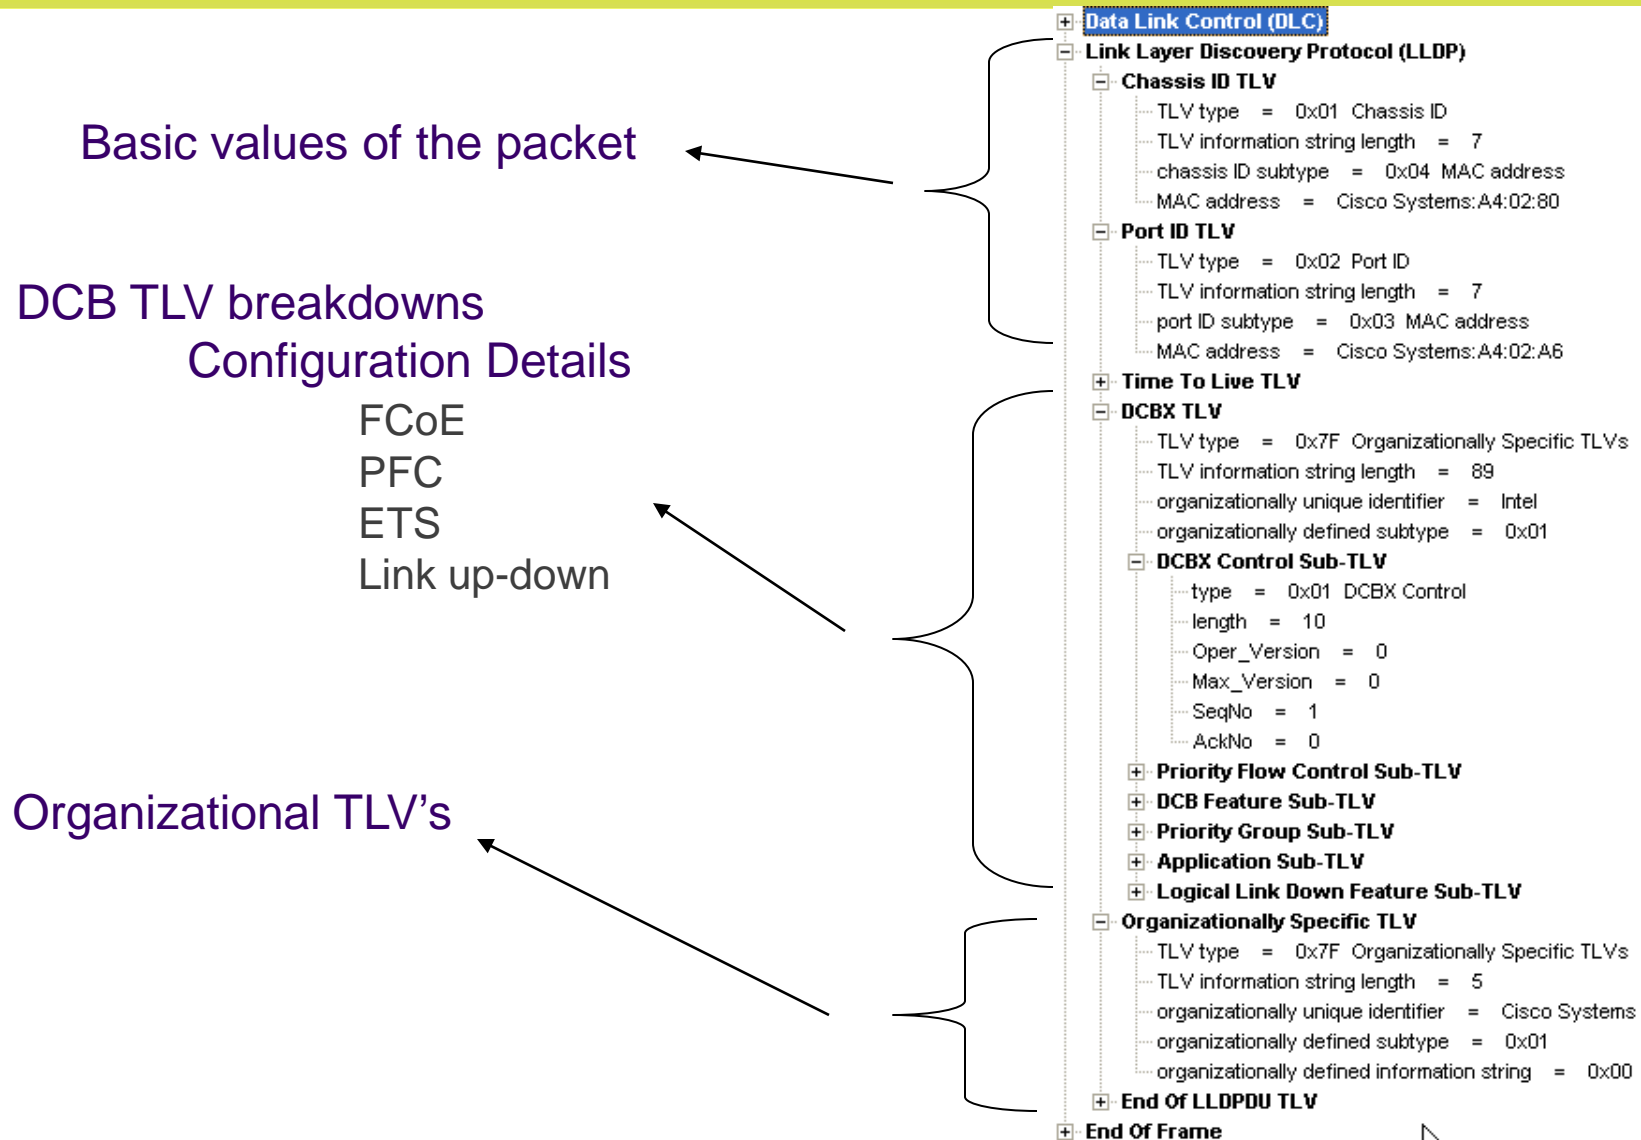
## ➤ Auto-negotiation of capability and configuration

## ➤ Allows network devices to advertise their capabilities over a link

- ◆ Enables hosts to get proper configuration from network
- ◆ Enables switches to verify proper configuration



# LLDP & DCB inspection



- FIP VLAN Discovery - used by the ENode to discover FCOE VLAN
- FIP Discovery Solicitations - used by the ENode to discover a FCF
- FIP Discovery Advertisement - used by FCF to advertise itself to ENodes
- FIP FLOGI, FIP LS\_ACC and FIP LS\_RJT - used by FCF to login and get itself an FCID
- FIP LOGO - sent by ENode to logout
- FIP NPIV FDISC - sent by VN ports to login
- FIP KeepAlives - liveliness mechanism between ENodes and FCF
- FIP Clear Virtual Link - used by the FCF to signal to the ENode to bring down the virtual link

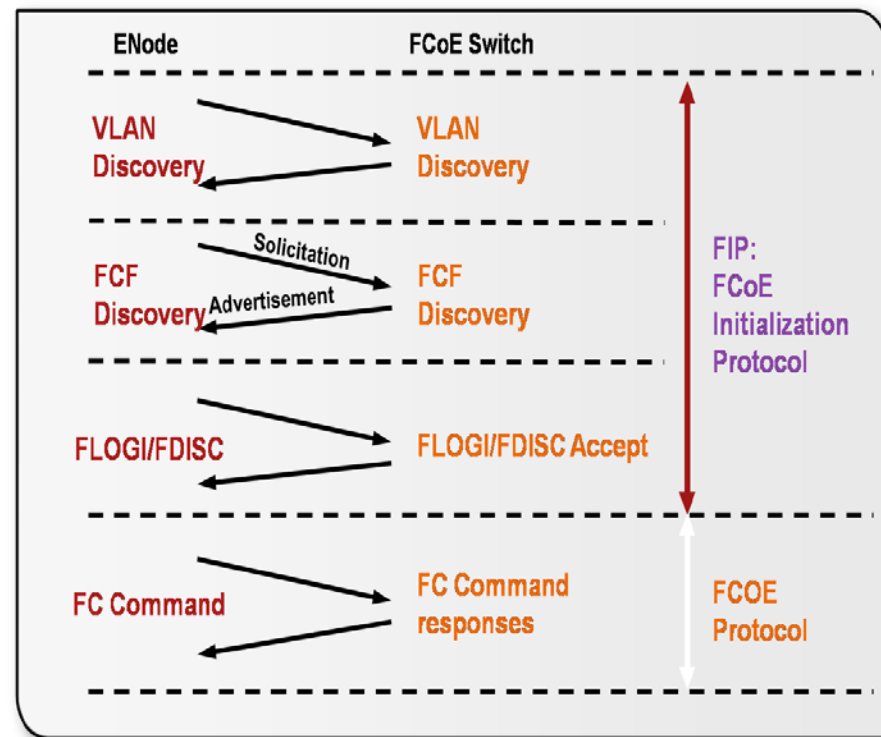
# FIP: FCoE Initialization Protocol

## ➤ FCoE VLAN discovery

- Automatic discovery of FCoE VLANs

## ➤ Device discovery

- ENodes discover VF\_Port capable FCF-MACs for VN\_Port to VF\_Port Virtual Links
- VE\_Port capable FCF-MACs discover other VE\_Port capable FCF-MACs for VE\_Port to VE\_Port Virtual Links
- Verifies the Ethernet network supports the required Max FCoE Size along with many other settings like FC map value using Descriptor TLV's



## Virtual Link instantiation

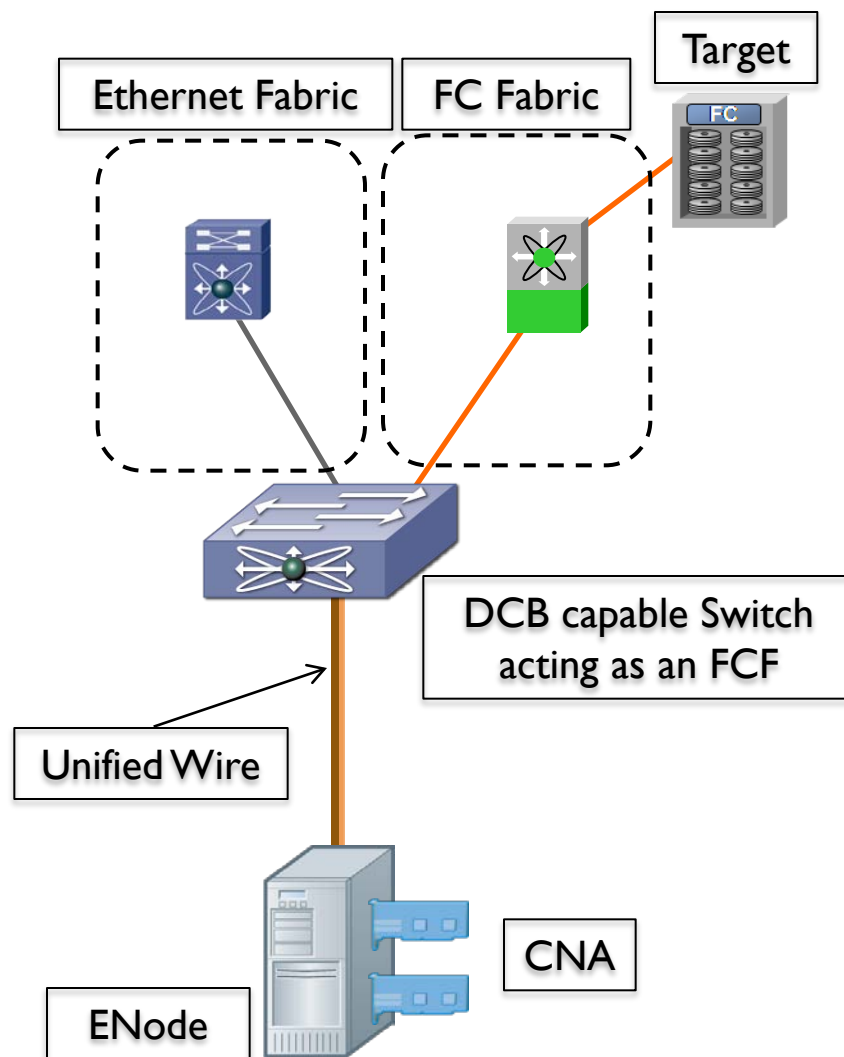
- Builds on the existing Fibre Channel Login process
  - Fabric Provided MAC Address (FPMA),

## Virtual Links maintenance

- Timer based

# FCoE, Same Model as FC

- **A-B SAN separation** – picture showing one side
- **Same host to target communication**
  - ◆ Host has 2 CNA's (one per fabric)
  - ◆ Target has multiple ports to connect to fabric
- **Connect to a DCB capable switch**
  - ◆ Port Type Negotiation (FC port type will be handled by FIP)
  - ◆ Speed Negotiation
  - ◆ DCB Negotiation
- **Access switch is a Fibre Channel Forwarder (FCF)**
- **Dual fabrics are still deployed for redundancy**




# FCOE NETWORK DEPLOYMENTS




# Basic Network Study

Enhanced Ethernet and FCoE Ethernet FC

## 10 GigE Ethernet

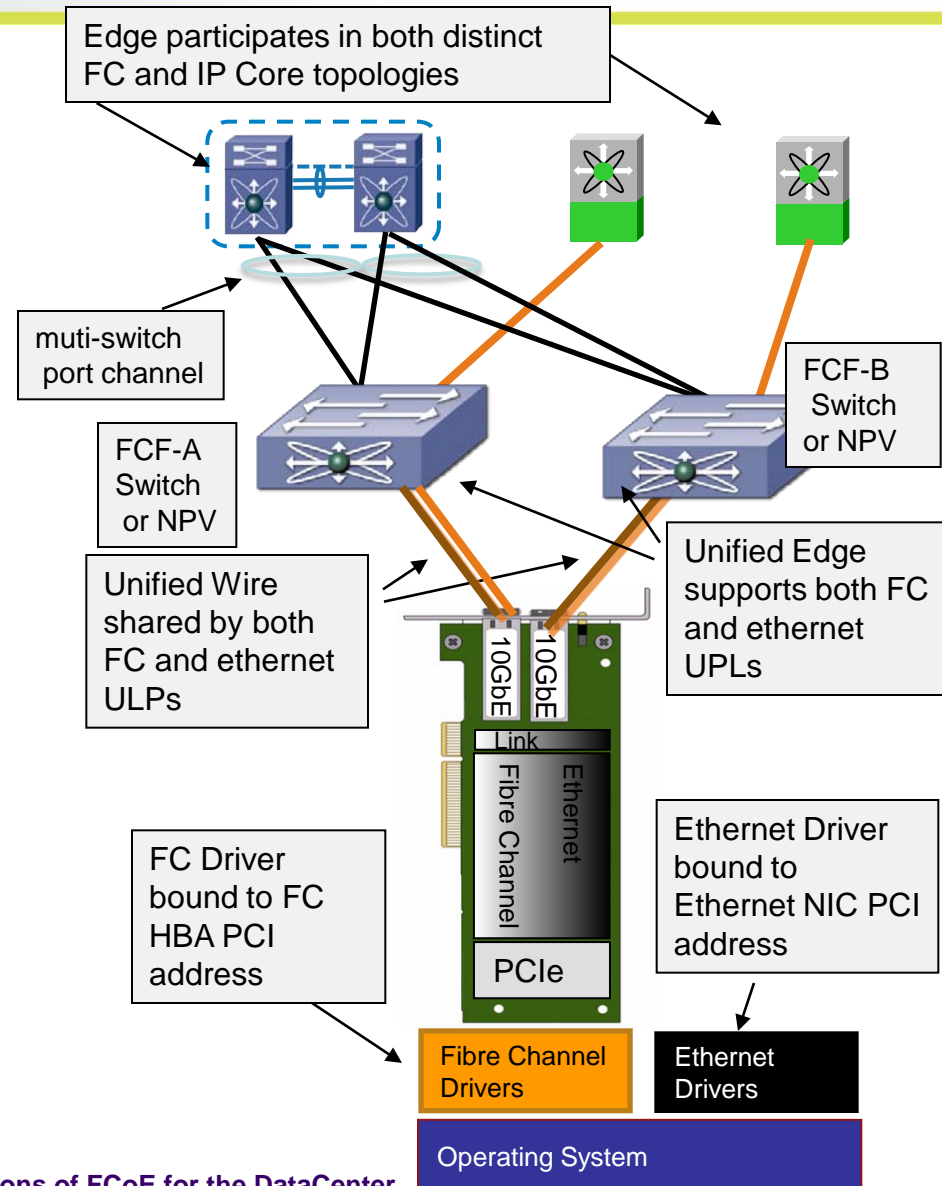


## I/O Consolidation FCoE



# Deployments with FCoE

- Converged Network Adapter (CNA) presents two PCI address to the Operating System (OS)
- OS loads two unique sets of drivers and manages two unique application topologies
- Server participates in both topologies separately
- Two stacks and thus two different views of the same **'unified wire'**
  - ◆ SAN Multi-Pathing provides failover between two fabrics (SAN 'A' and SAN 'B')
  - ◆ NIC Teaming provides failover within the same fabric (VLAN)




# What is being deployed

Single Hop Design

## ➤ Generation 2 CNA

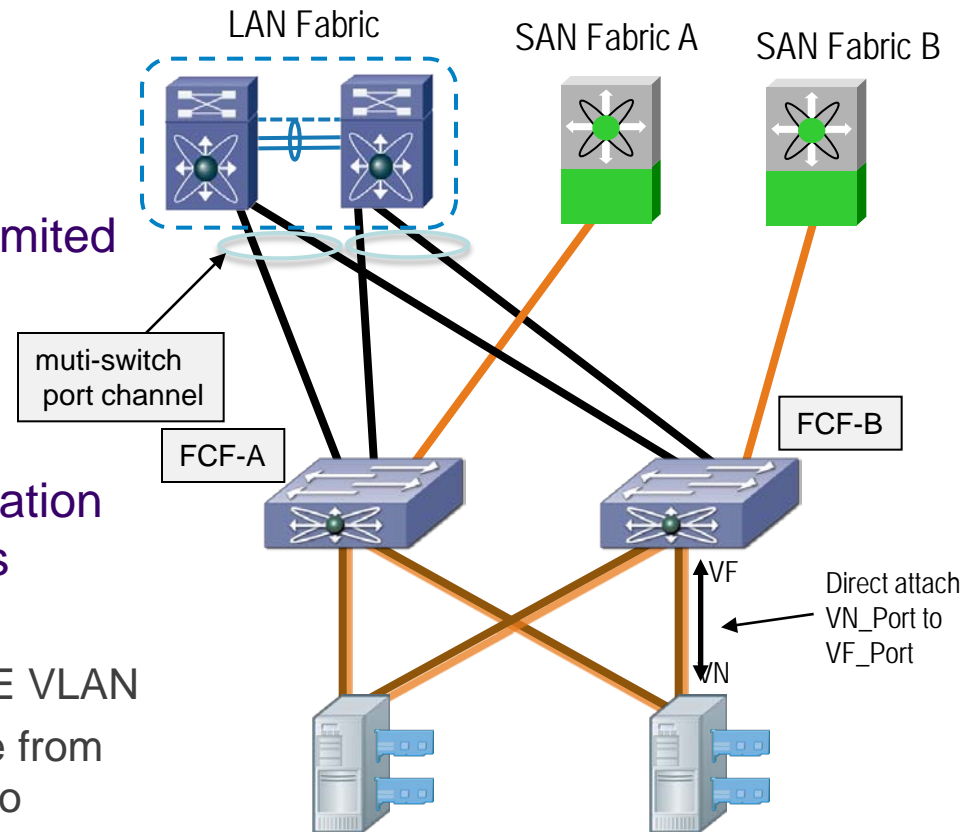
- Utilizes Converged Enhanced Ethernet & Data Center Bridging Exchange protocol (DCB)
- ◆ Utilizes FCoE Initialization Protocol (FIP) as defined by the T.11 FC-BB-5 specification
- ◆ Supports both direct and multi-hop attachment (*i.e. through a FIP Snooping Bridge or over VE-Ports*)



# What is being deployed

Initiator connection


- Physical link is brought online (today requires 10GE)
- DCBX negotiation – discovers DCB capable devices, including but not limited to lossless Ethernet.
- FIP Process – discovery and negotiation of FCoE devices and characteristics
  - ◆ FCoE VLAN Discovery
  - ◆ FCF Discovery on the specific FCoE VLAN
  - ◆ Fabric Login - builds the logical wire from the end node to the FCF (VN\_port to VF\_port)




Direct Attach Topologies

# Single Hop Design FCoE VLAN

- A VLAN can be dedicated for every Virtual Fabric in the SAN
- FIP discovers the FCoE VLAN and signals it to the hosts
- Trunking is not required on the host driver – all FCoE frames are tagged by the CNA
- FCoE VLANs can be pruned from Ethernet links that are not designate for FCoE
- Maintains isolated edge switches for SAN 'A' and 'B' and separate LAN switches for NIC 1 and NIC 2 (standard NIC teaming or Link Aggregation)




- In order to maintain the integrity of FC forwarding over FCoE, FCoE VLANs should be treated differently than LAN VLANs
  - ◆ No flooding, MAC learning, broadcasts, etc.
- The FCoE VLAN should be configured as a native VLAN
  - ◆ FIP uses native VLAN
- Separate FCoE VLANs are used for FCoE in SAN-A and SAN-B
- Unified Wires can be configured as trunk ports and STP edge ports



# Single Hop Design

Unified Wires and multi-switch port channels

- Optimal layer 2 LAN design often leverages multi-switch port channel
- A multi-switch port channel can provide network based load sharing and redundancy and some implementations can do this without introducing layer 2 loops in the topology
- Multi-switch port channel results in diverging LAN and SAN high availability topologies
  - ◆ FC maintains separate SAN 'A' and SAN 'B' topologies
  - ◆ LAN utilizes a single logical topology



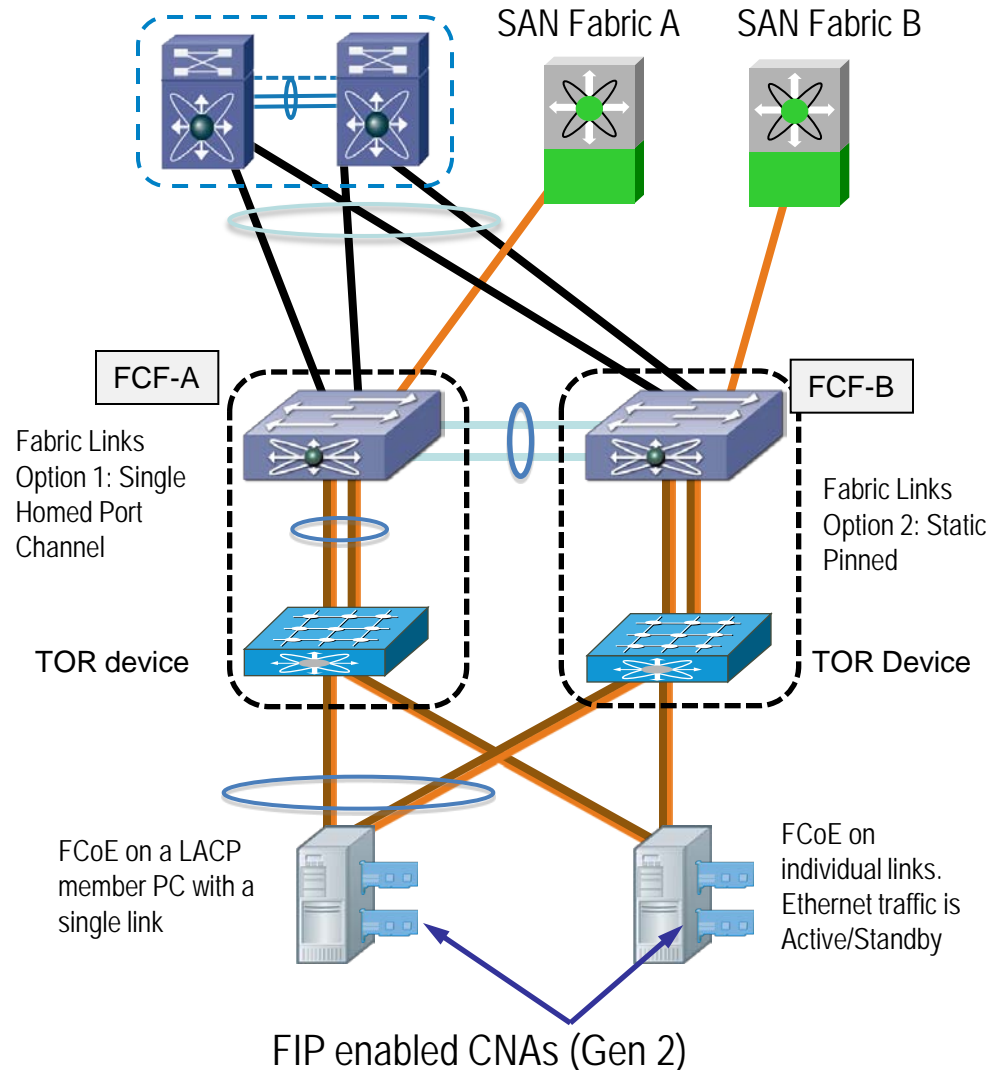
Direct Host Dual Attached Topology



# Single Hop Design

## Extending the FCoE Edge to Top of Rack

- Server Ethernet driver connected to TOR device with NIC Teaming
- FCoE can connect over multi member port with a single link from server to TOR
- TOR device *single* homed to upstream core FCF
  - ◆ Top of Rack fabric links can be connected to FCF with individual links or a multi-switch port channel
  - ◆ Can be oversubscribed
- Consistent with separate LAN Access and SAN Edge Topologies





# Extending FCoE with FIP Snooping

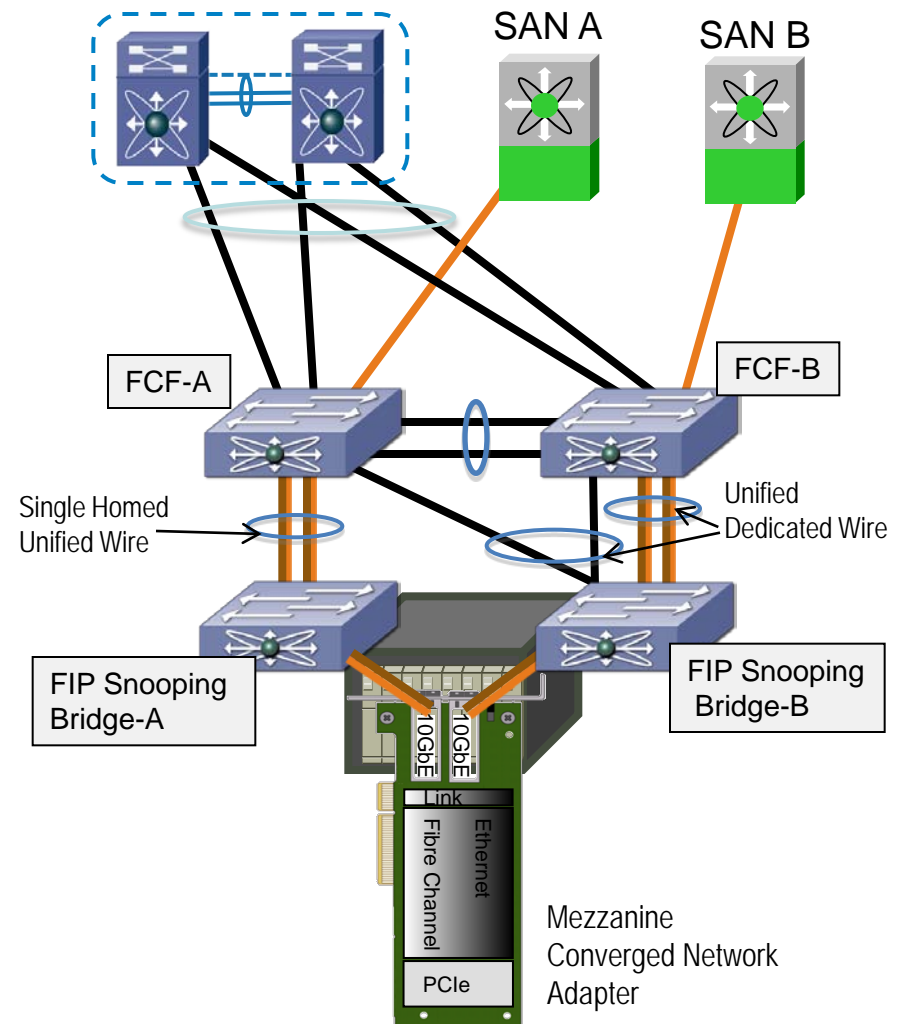
## ➤ FIP snooping Bridge

- ◆ DCB enabled

## ➤ Dual Topology Capable

## ➤ Still need to maintain the A-B SAN separation


- ◆ Unified Dedicated Wires from FIP bridge to FCF
- ◆ Single Unified Wire Port Channel from FIP Bridge to FCF



# Multi - Hop Designs (VE)

## Considerations for FCoE Multi-hop


- **Virtual E\_Ports (VE)**
- **Design considerations when extending FCoE beyond the Unified Edge?**
  - ◆ High Availability for both LAN and SAN
  - ◆ Oversubscription for SAN and LAN
  - ◆ Ethernet Layer 2 and STP design
- **Where does Unified Wire make sense over Unified Dedicated Wire?**
- **Unified Wire provides for sharing of a single link for both FC and Ethernet traffic**



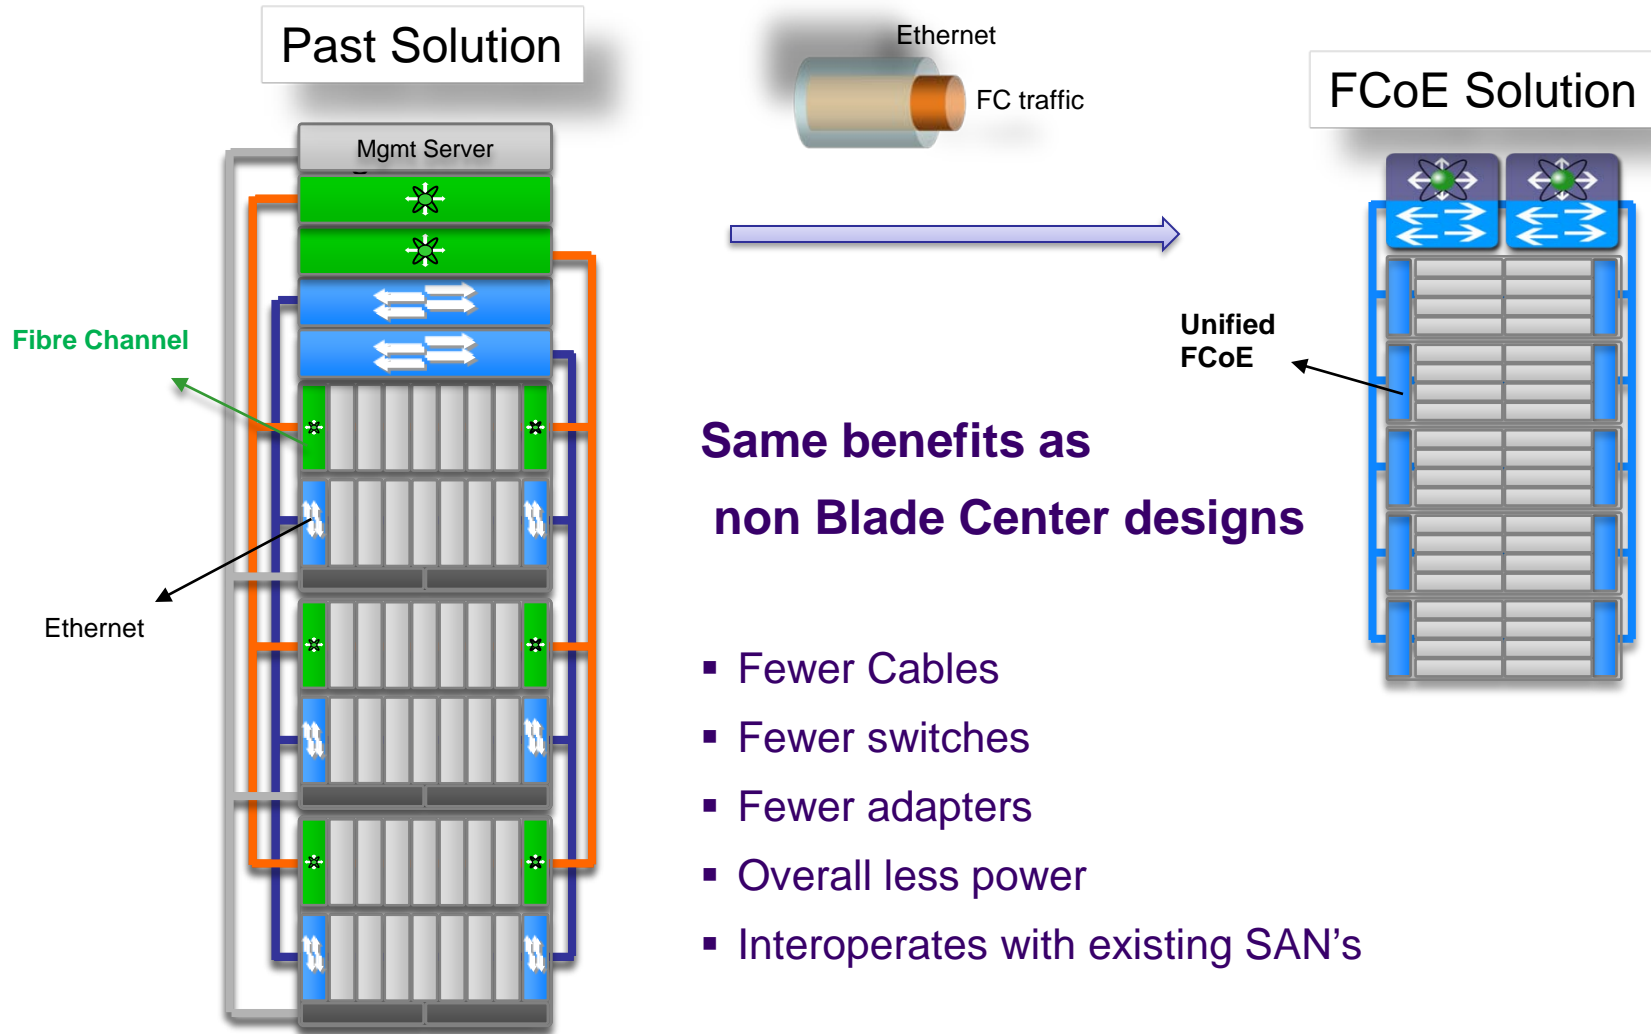
# UNIFIED I/O IN COMPUTE BLADE CENTER

# FCoE in the Compute Blade Center

## Fabric Consolidation



# FCoE in the Compute Blade Center



# Check Out the Hands-On Lab



**Check out the Hands-On LAB at SNW:**

**Unified Storage Infrastructure  
Unified Storage IP Solutions**

---

**Hands-On**  
**LAB** COMPUTERWORLD SNIA  
**SNW**

---

- Please send any questions or comments on this presentation to SNIA: [tracknetworking@snia.org](mailto:tracknetworking@snia.org)

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

**- SNIA Education Committee**

Simon Gordon

Joseph White