Cisco Expo
2012

Vše o FCoE co jste chtěli vědět, ale báli jste se zeptat

René Raeber

Datacenter Architect

IEEE 802.1DCB Architect

Prosíme, ptejte se nás

- Twitter www.twitter.com/CiscoCZ
- Talk2cisco www.talk2cisco.cz/dotazy
- SMS 721 994 600

Detailed Agenda

FC over IP Building Blocks

- Addressing Methods
- Converged adaptors
- Cisco Products & FCoE Features

FCoE Requirements Ethernet

- DCB Protocol
- FCoE Standards
- QOS & configurations

FCoE Protocol Fundamentals

- Data & Control Plane
- FIP Exchange

FCoE Network Requirements

- FIP Snooping
- N-Port ID Virtualization
- Zoning

Single Hop Design Details

- Unifying the Wire
- VLAN to VSAN mapping
- Designs with vPC
- Designs with N2K FEX
- Designs B22 FEX for HP C7000/C3000

VE_Ports

- Dedicated ISL links
- Protocol Handshake

Multi-Hop Design Details

- N7K Storage VDC
- Share Converged Ports
- MDS SAN Designs
- FCoE Core Edge Designs with N2k, N5K, N7K & MDS

FCoE Connectivity

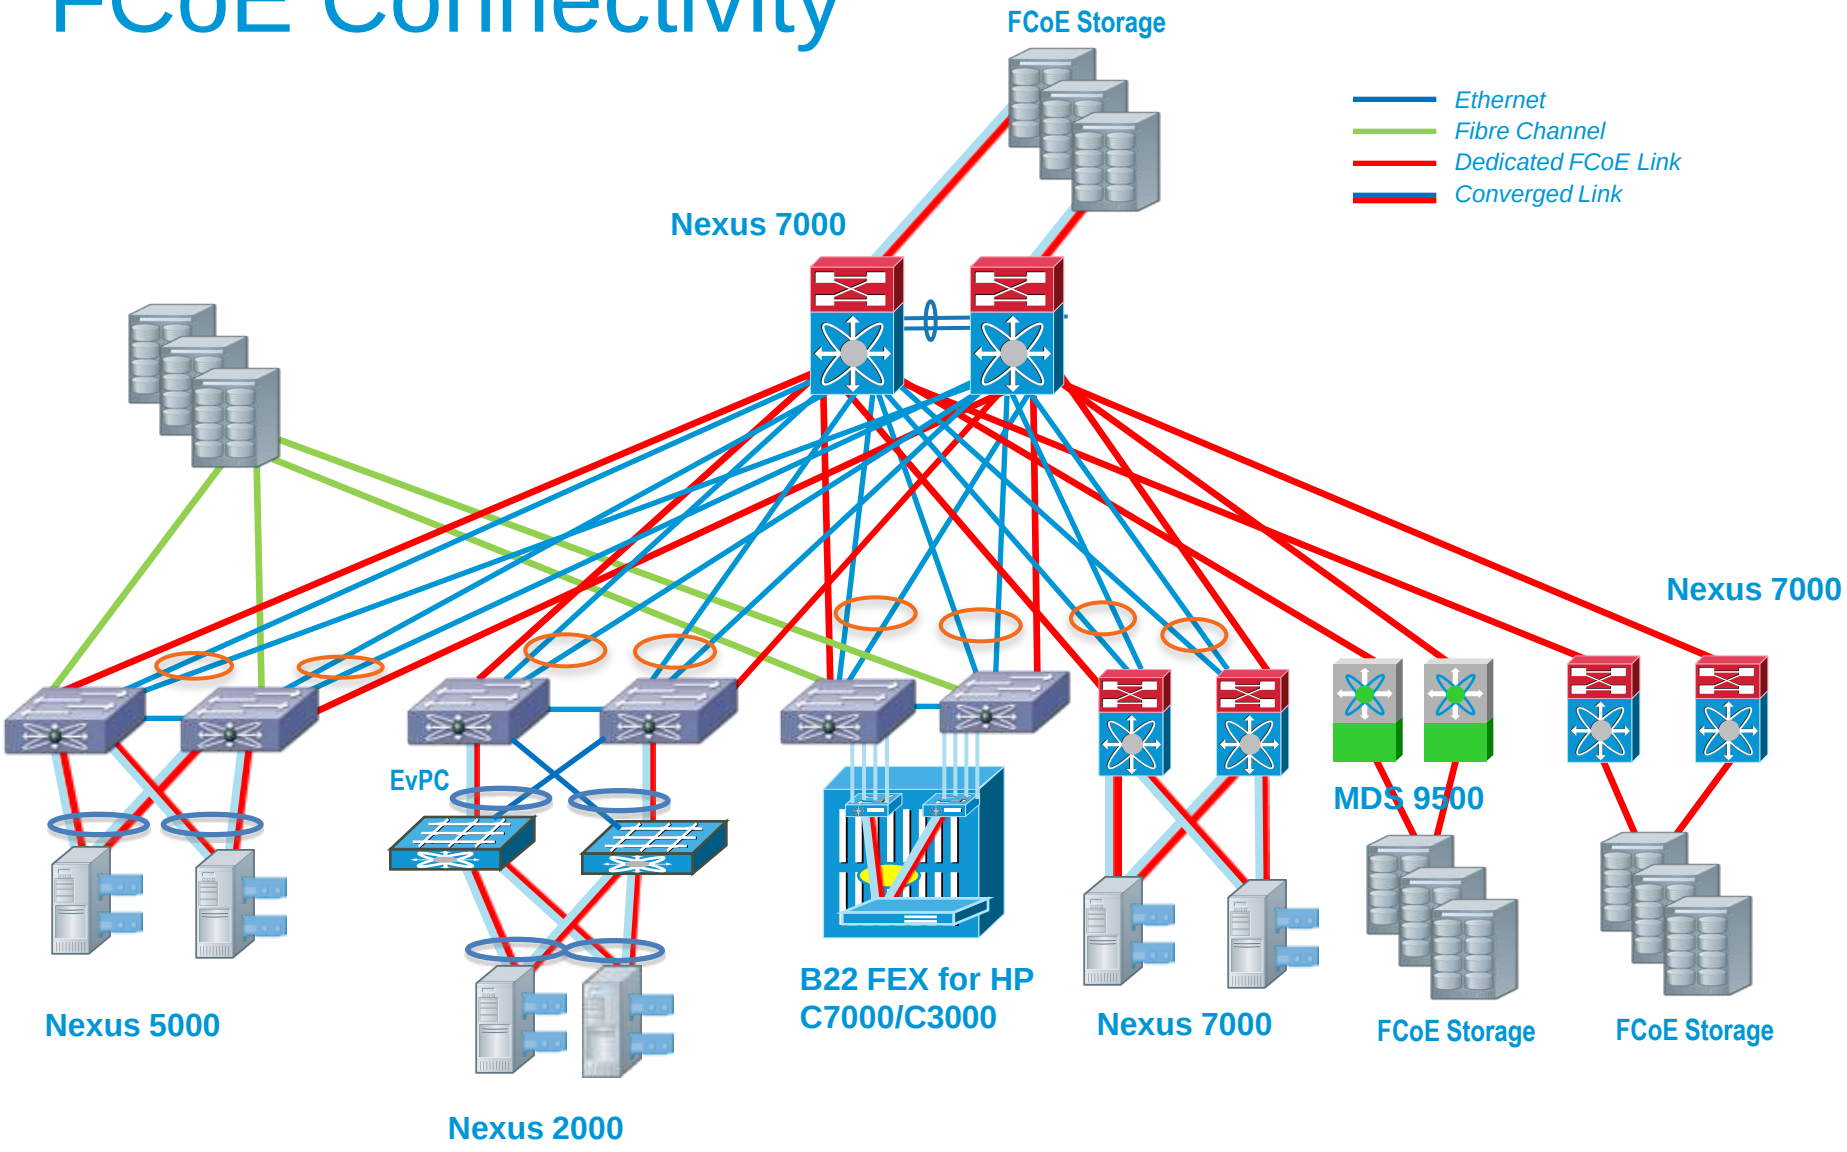

Terms Defined

- **CE -** Classical Ethernet (non lossless)
- **DCB & DCBx** - Data Center Bridging, Data Center Bridging Exchange
- **FCF** - Fibre Channel Forwarder (Nexus 5000, Nexus 7000, MDS 9000)
- **FIP** – FCoE Initialization Protocol
- **Enode**: a Fiber Channel end node that is able to transmit FCoE frames using one or more Enode MACs.
- **FIP snooping Bridge**
- **FCoE-NPV** - Fibre Channel over IP N_Port Virtualization
- **Single hop FCoE** : running FCoE between the host and the first hop access level switch
- **Multi-hop FCoE** : the extension of FCoE beyond a single hop into the Aggregation and Core layers of the Data Centre Network
- **Zoning -** Security method used in Storage Area Networks
- **FPMA –** Fabric Provided Management Address

FCoE Building Blocks
Fibre Channel over Ethernet Addressing Scheme

- Enode FCoE MAC assigned for each FCID
- Enode FCoE MAC composed of a FC-MAP and FCID

  FC-MAP is the upper 24 bits of the Enode's FCoE MAC

  FCID is the lower 24 bits of the Enode's MAC
- FCoE forwarding decisions still made based on FSPF and the FCID within the Enode MAC

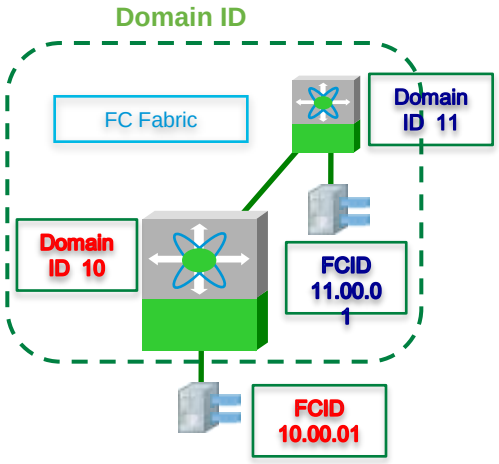

Fibre Channel FCID Addressing

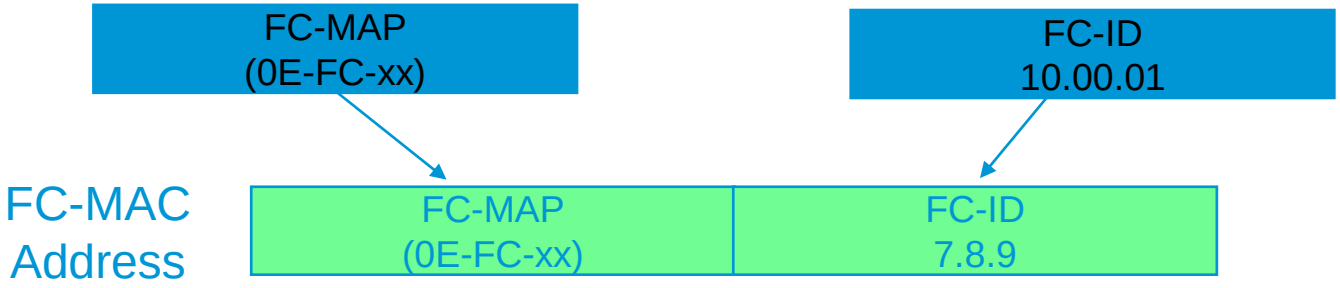

FCoE Building Blocks

Converged Network Adapter (CNA)

- Replaces multiple adapters per server, consolidating both Ethernet and FC on a single interface
- Appears to the operation system as individual interfaces (NICs and HBAs)
- CNAs support PFC, DCB as well as FIP
  - Single chip implementation
  - Half Height/Length
  - Low power consumption
- Installer vendor choice sometimes based off Fibre Channel Host Bus Adaptor use in current enterprise SAN

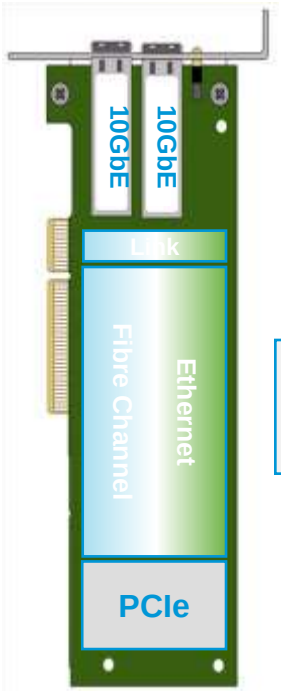

FC Driver bound to FC HBA PCI address

Ethernet Driver bound to Ethernet NIC PCI address

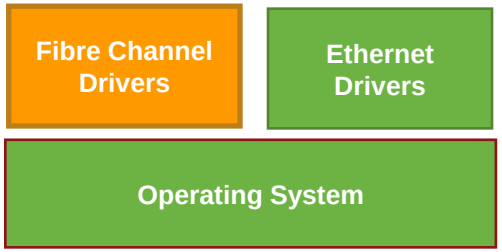

FCoE Building Blocks
Fibre Channel Forwarder

- FCF (Fibre Channel Forwarder) is the Fibre Channel switching element inside an FCoE switch
  - Fibre Channel logins (FLOGIs) happens at the FCF
  - Consumes a Domain ID
- FCoE encap/decap happens within the FCF
  - Forwarding based on FC information

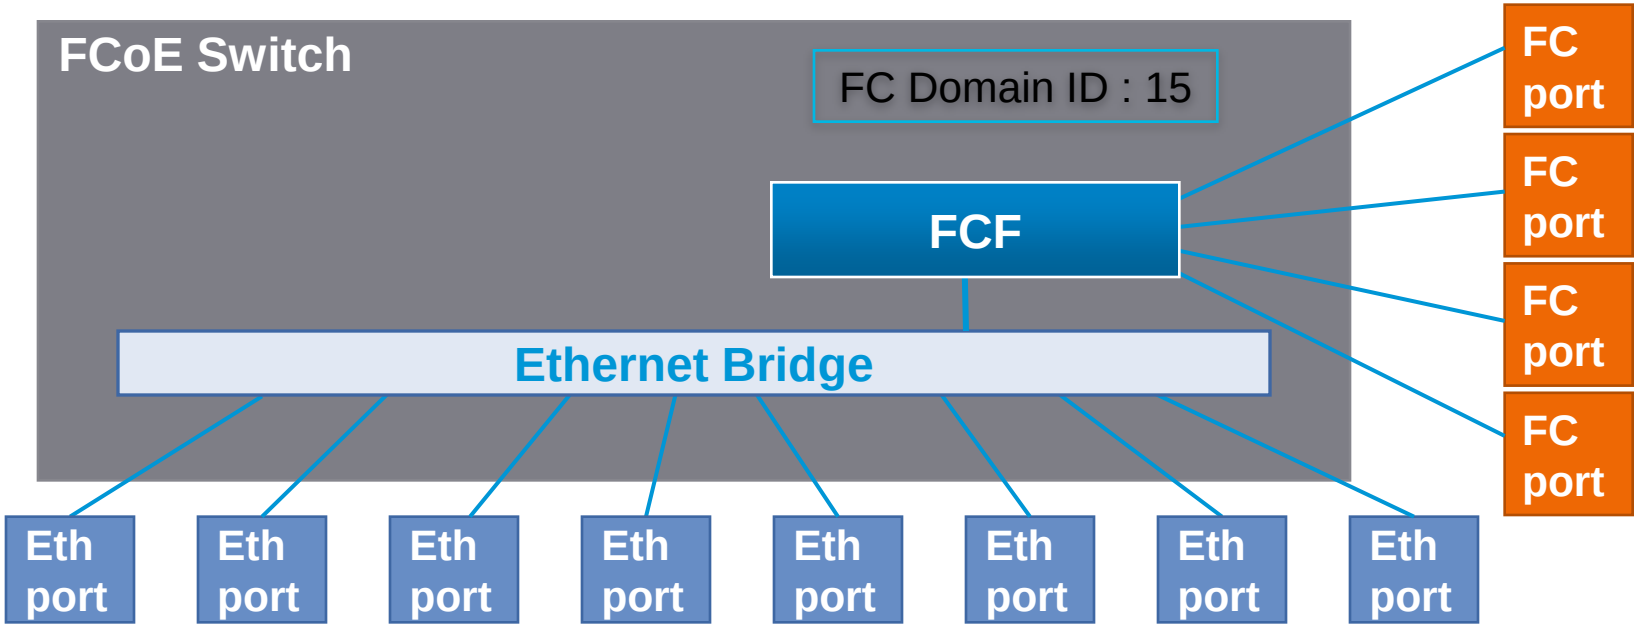

Fibre Channel Forwarder

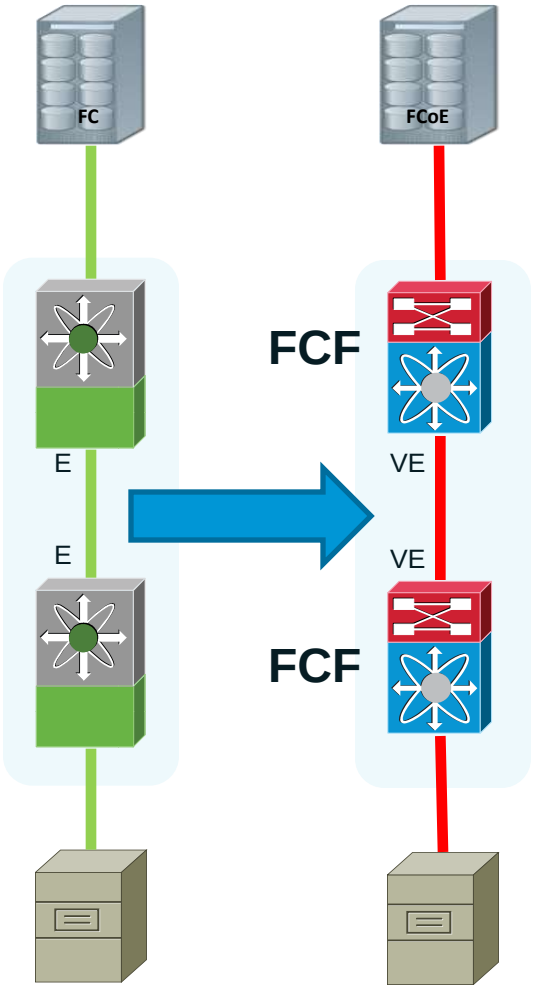

Allows switching of FCoE frames across multiple hops

FCF supports all **FC** Functionalities:

- Support up to 7 hops
- Support up to 10,000 logins per fabric
- Supports up to 8,000 Zones per switch
- Supports up to 500 Zonesets per switch

E-Ports w/ FC

VE-Ports w/ FCoE

FC-BB-5 (FCOE) Terminology

- Unchanged from previous FC standard
  - VN_Port: Virtual N_Port
  - VF_Port: Virtual F_Port
  - VE_Port: Virtual E_Port

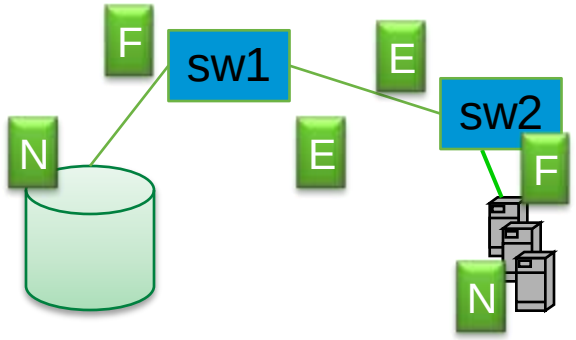

- Added to support FCoE
  - FCoE_LEP (FCoE link endpoint): The data forwarding component that handles FC frame encapsulation/decapsulation, and transmission/reception of FCoE frames
  - FCoE Controller: the entity that implement the FIP protocol

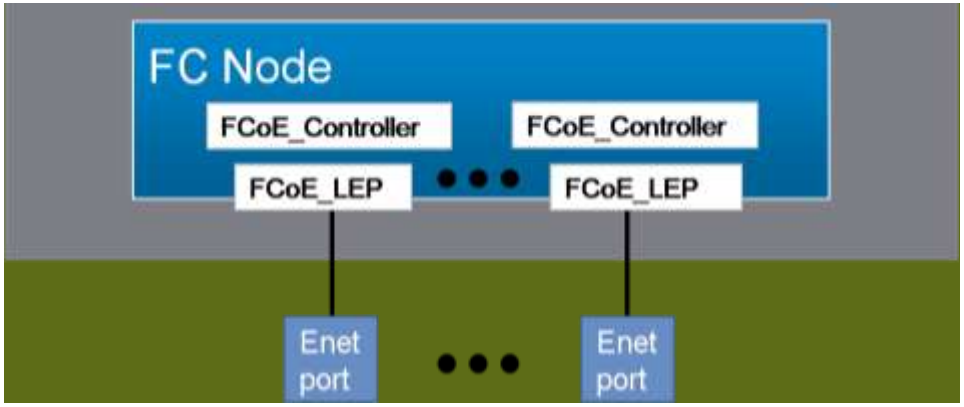

FCoE Building Blocks

FCoE Port Types

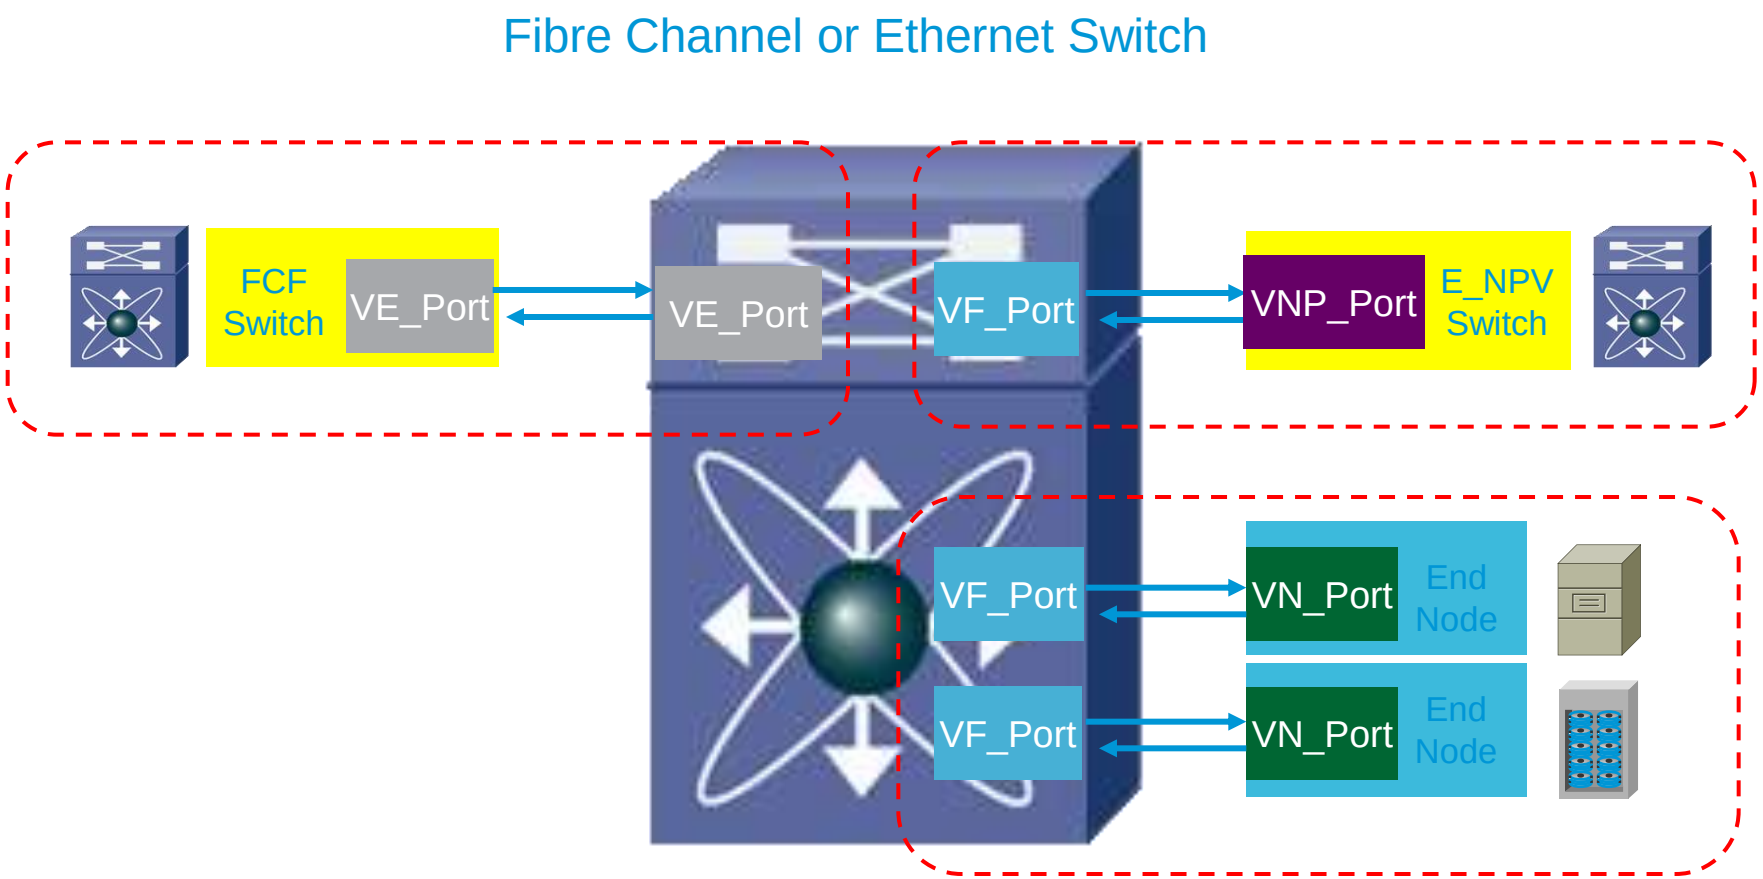

FCoE Switch : FCF

Converged Networks

- **Unifying the I/O** – using Ethernet as the transport medium in all network environments -- no longer needing separate cabling options for LAN and SAN networks

Converged Link – a single DCB Ethernet link actively carrying both LAN and Storage (FC/FCoE/NAS/iSCSI) traffic simultaneously

Dedicate Link -- a single DCB Ethernet link capable of carrying all traffic types but actively dedicated to a single traffic type for traffic engineering purpose

FCoE Building Blocks

- **Converged Link** to the access switch
  - Cost savings in the reduction of required equipment
  - “cable once” for all servers to have access to both LAN and SAN networks
- **Dedicated Link** from access to aggregation
  - Separate links for SAN and LAN traffic - both links are same I/O (10GE)
  - Advanced Ethernet features can be applied to the LAN links
  - Maintains fabric isolation

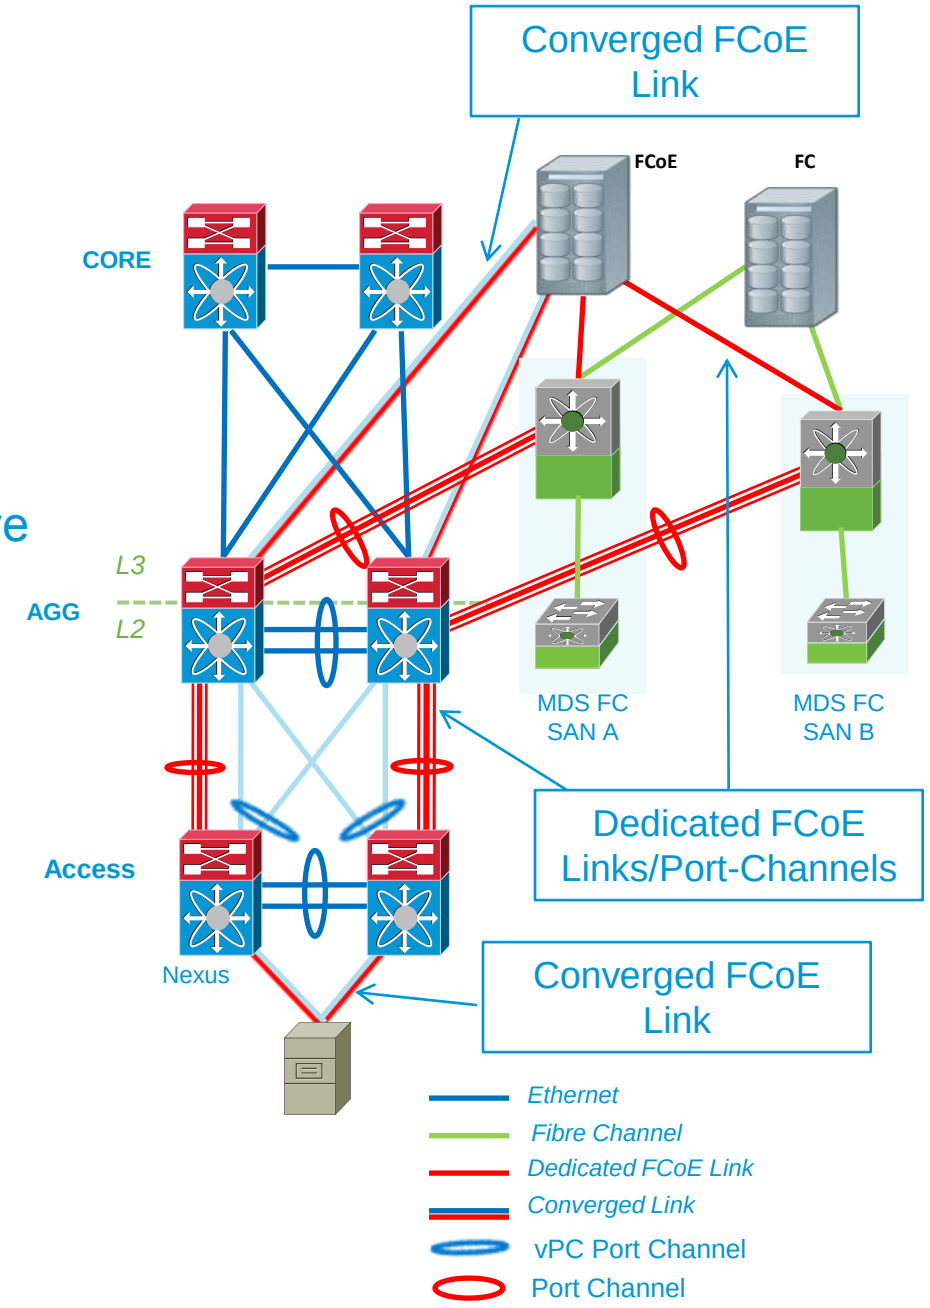

Converged Network Design

A Single Fabric

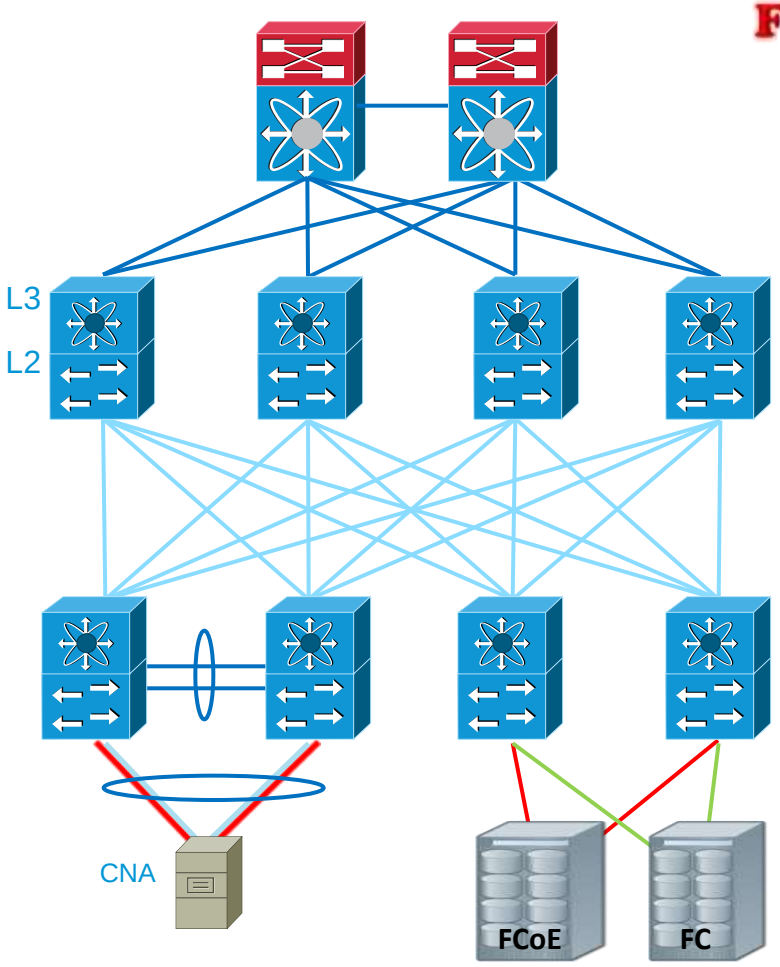

FORWARD LOOKING

- Multi-Purpose and Resilient Ethernet Fabric
- Link Level (ECMP) and Switch level Redundancy
- Converge on to this Fabric, all your LAN and SAN traffic
- Logical SAN A&B separation
- Currently Not supported:
  1. Majority of customers still require physical A&B separation
  2. Both Fabric Path and MultiHop FCoE are at their infancy. Perceived Risk is High
  3. As customer appetite for FCoE matures, in due course this model will be supported

Unified Port Overview

- **Ports on Nexus 5548UP & Nexus 5596UP and Unified port GEM can be configured to be in Ethernet or FC mode.**

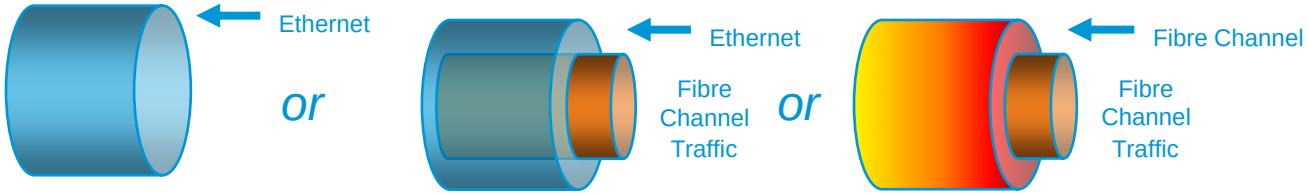

Configuration example:

```
n5k(config)# slot 1
n5k(config-slot)# port 41-48 type fc
n5k(config-slot)# port 1-40 type ethernet
```

- **Supported HW**

  Nexus 5548UP & Nexus 5596UP baseboard ports
  Unified GEM Ports

Any Nexus 5500UP port can be configured as 1/10GE, DCB (lossless Ethernet), FCoE on 10GE (dedicated or unified wire) or 8/4/2/1G native Fibre Channel

Implement with native Fibre Channel today, enables smooth migration to FCoE in the future

Cisco Nexus Unified Ports

Unified Platform

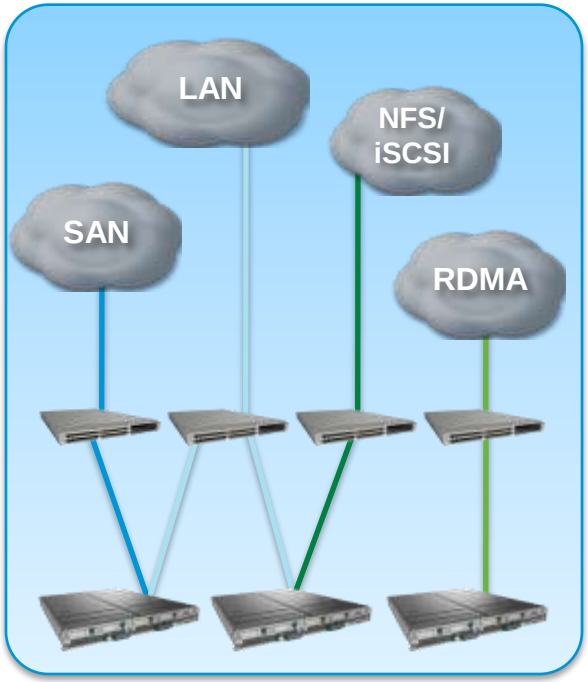

Unified Device

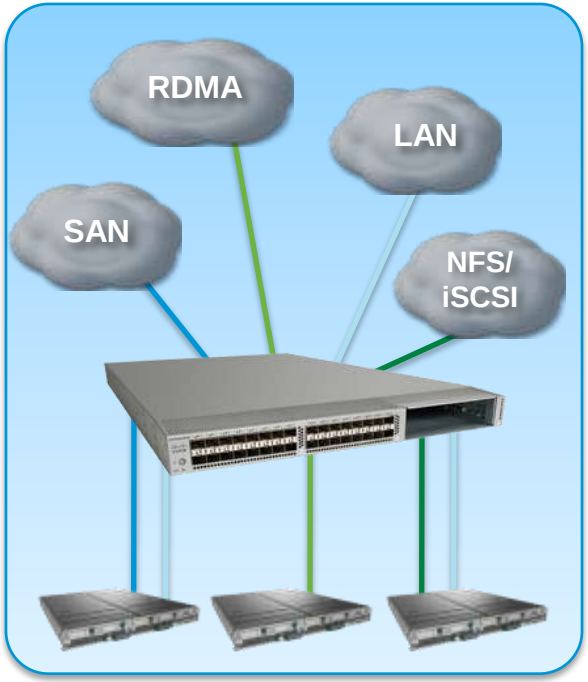

Unified Wire

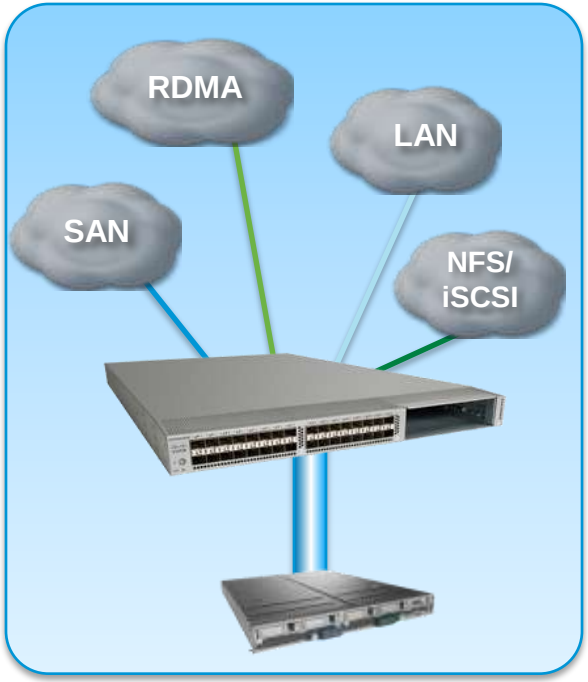

Unified Platform – Discrete services over separate cables with a consistent choice of platform throughout an organization

Unified Device – Multiple Services to drive Device Consolidation over separate cables

Unified Wire – Multiple Services to drive Device & Cable Consolidation

Dynamic and Efficient Port

Availability

- 16-port Expansion Module on the Nexus 5548 and 5596
- All Ports on the Nexus 5596

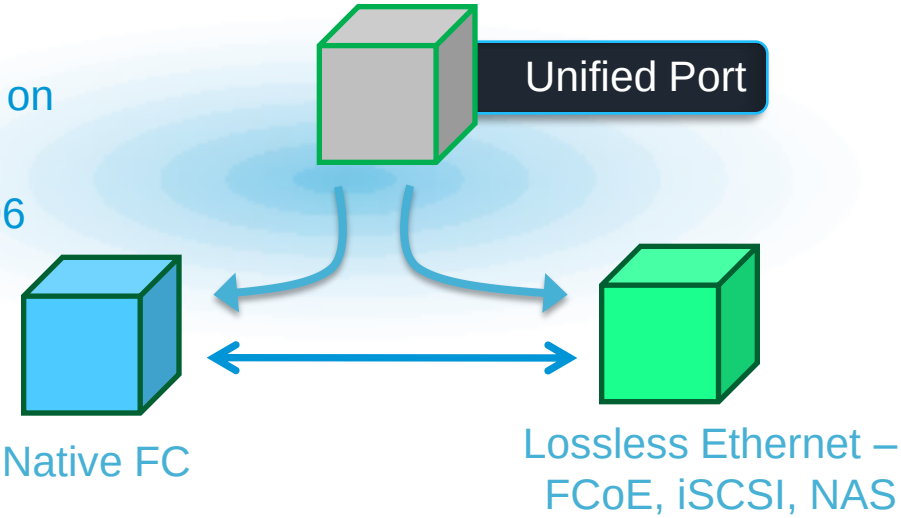

Benefits

- Simplify switch purchase - remove ports ratio guess work
- Increase design flexibility
- Remove specific protocol bandwidth bottlenecks

Use-cases

- Flexible LAN & storage convergence based on business needs
- Service can be adjusted based on the demand for specific traffic

Nexus 55XX Port Range Restrictions

- In each module, continuous set of ports can be in configured as Ethernet or FC.
- Eth ports have to be the first set and they have to be one contiguous range. FC ports have to be second set and they have to be contiguous as well.
- Number of Ethernet and FC ports on a module can be changed.

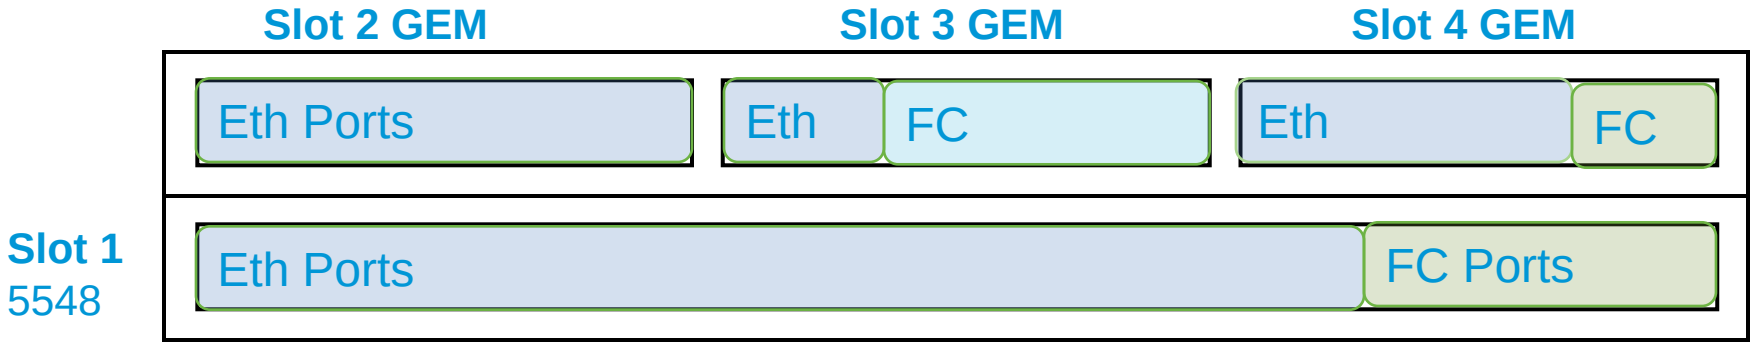

5548

5596

10Gig/FCoE Fabric Extender

Nexus 2232PP

Remote Line Card (FEX)

FEX-2232

- 32 server facing 10Gig/FCoE ports
  - T11 standard based FIP/FCoE support on all ports
- 8 10Gig/FCoE uplink ports for connections to the Nexus 5K
- Management and configuration handled by the Nexus 5K
- Support for DCBx
- Part of the Cisco Nexus 2000 Fabric Extender family
- N7K will support FCoE on 2232 (Future)

FCoE Support on N7K & MDS

- N7K FCoE capable line card: Module (N7K-F132XP-15)

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9402/data_sheet_c78-605622.html

- MDS FCoE capable line card: Module (DS-X9708-K9)

http://www.cisco.com/en/US/partner/products/hw/ps4159/ps4358/products_data_sheets_list.html

8 ports at 10 GbE (line-rate)
Loss-Less Ethernet Capable (PFC, ETS, DCBX)
Requires: Supervisor 2A Module and Fabric2
FCoE only ports

F1 Series
32 ports at 1/10 GbE
Loss-Less Ethernet Capable (PFC, ETS, DCBX)
512 ports per system

F2 Series
48 ports at 1/10 GbE
FCOE support with FEX 2232 and F2 LC is planned

- N7K Requirements
  - Storage VDC
  - FCoE License per Line card
- MDS Requirements
  - FCoE License

Agenda

FCoE Requirements for Lossless Ethernet

Data Center Bridging eXchange

Control Protocol – the “handshake”

- Defined in IEEE 802.1Qaz
- Point to point per link discovery
- Negotiates capability’s – PFC, ETS, Application’s (FCoE)
- Simplifies management of DCB nodes
  - Allows for configuration and distribution of parameters from one node to another
- Responsible for Logical Link Up/Down signaling of Ethernet and Fibre Channel
- Uses Link Layer Discovery Protocol (LLDP) defined by 802.1AB to exchange and discover DCB capabilities
- DCBX negotiation failures result in:
  - per-priority-pause not enabled on CoS values
  - vfc not coming up – when DCBX is being used in FCoE environment

https://www.cisco.com/en/US/netsol/ns783/index.html

DCBX Protocol History

- **CIN:** (Cisco Intel Nuova)
  PFC,ETS, Applications (e.g., FCoE), BCN, Logical Link Down, NIV (Network Interface Virtualization)
- **CEE:** (Converged Enhanced Ethernet)
  PFC, ETS, Applications (e.g., FCoE)
- **IEEE:**
  PFC, ETS, Applications (e.g., FCoE)

IEEE-DCB Complete – awaiting publication

Current implementations (hardware & software) support CIN/CEE DCBX, when IEEE-DCBX is implemented, DCB exchange will begin by sending an IEEE-DCBX message

- If an IEEE DCBX response is received, continue to talk IEEE-DCBX
- If a CIN/CEE DCBX response is received come up as CEE

Standards for I/O Consolidation

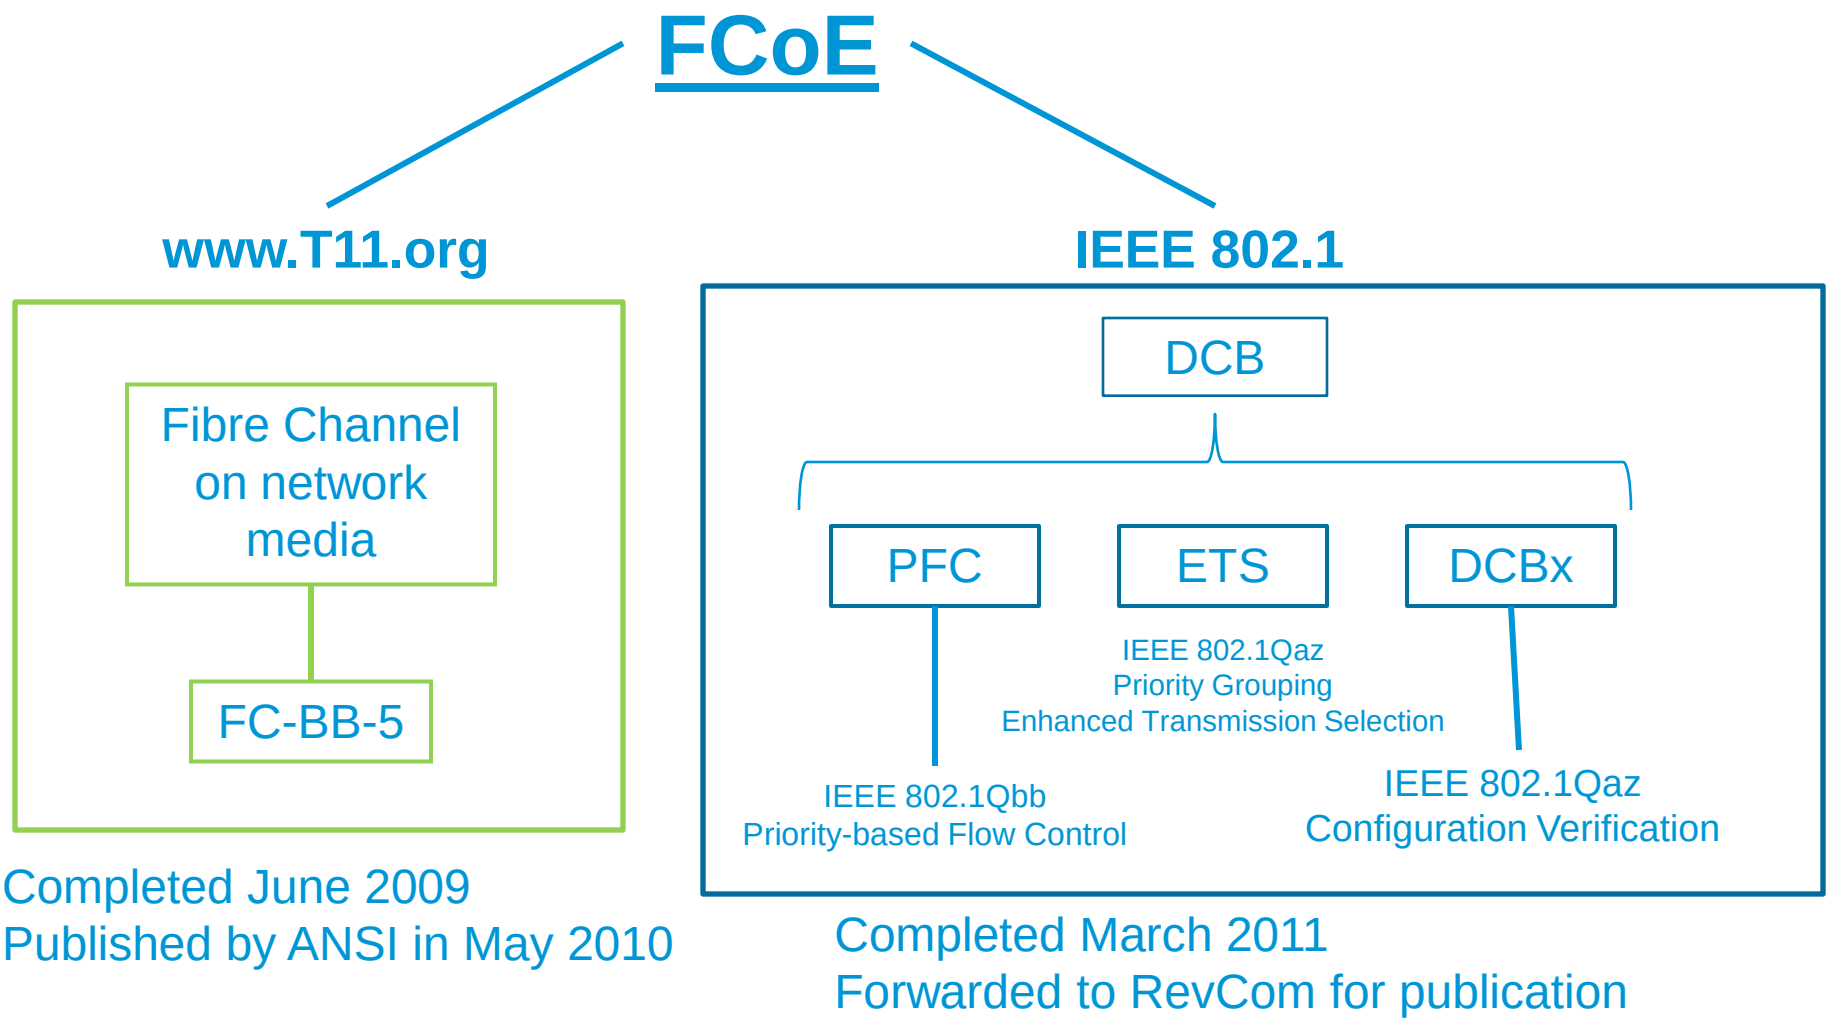

Priority Flow Control

FCoE Flow Control Mechanism – 802.1Qbb

- Enables lossless Ethernet using PAUSE based on a COS as defined in 802.1p
- When link is congested, CoS assigned to “no-drop” will be PAUSED
- Other traffic assigned to other CoS values will continue to transmit and rely on upper layer protocols for retransmission
- Not only for FCoE traffic

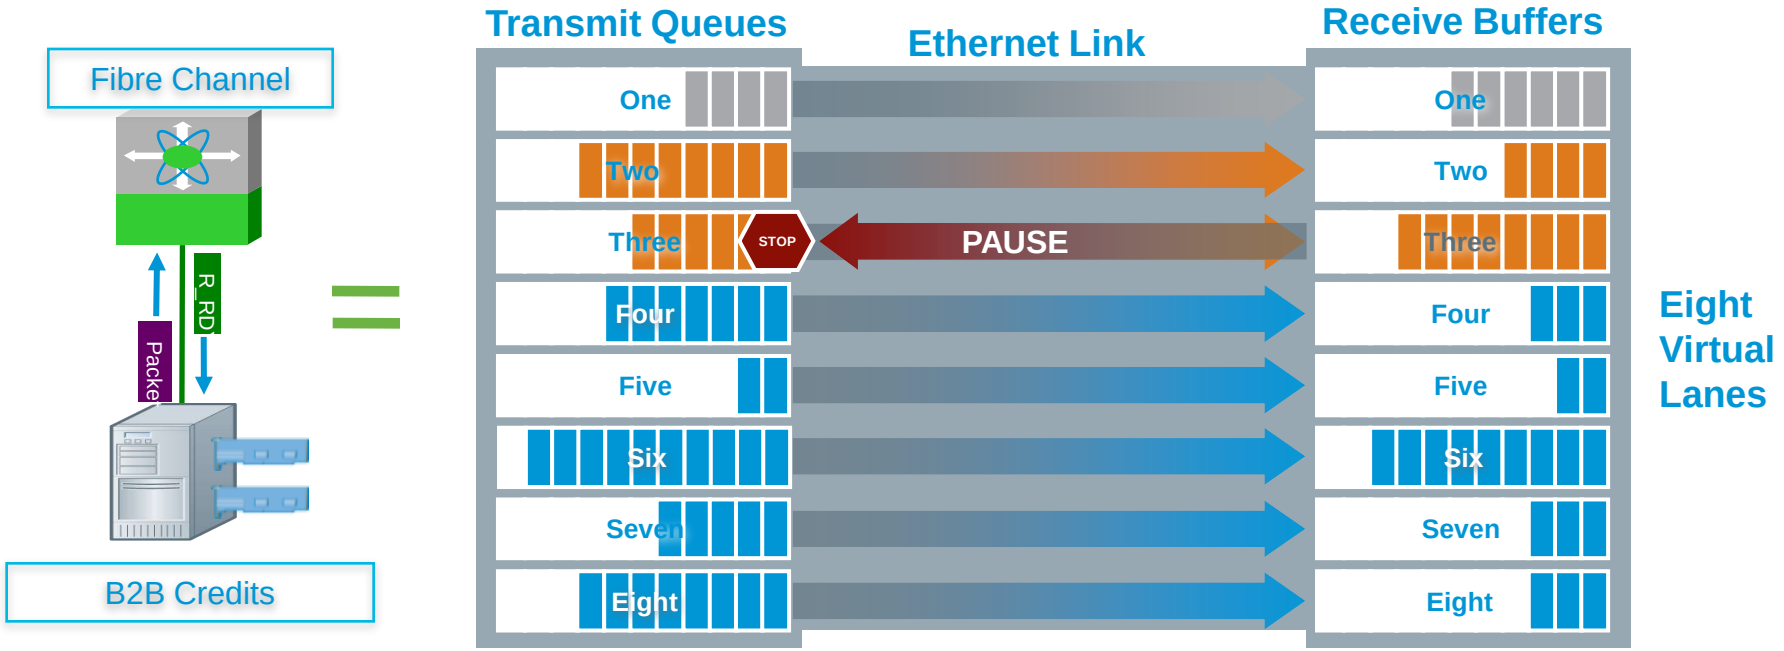

Priority Flow Control – Nexus 5K - Release 5.1(3)

Operations Configuration – Switch Level

- Once **feature fcoe** is configured, 2 classes are made **by default**

```
policy-map type qos default-in-policy
  class type qos class-fcoe
    set qos-group 1
  class type qos class-default
    set qos-group 0
```

- **class-fcoe** is configured to be **no-drop** with an MTU of 2158

```
policy-map type network-qos default-nq-policy
  class type network-qos class-fcoe
    pause no-drop
    mtu 2158
```

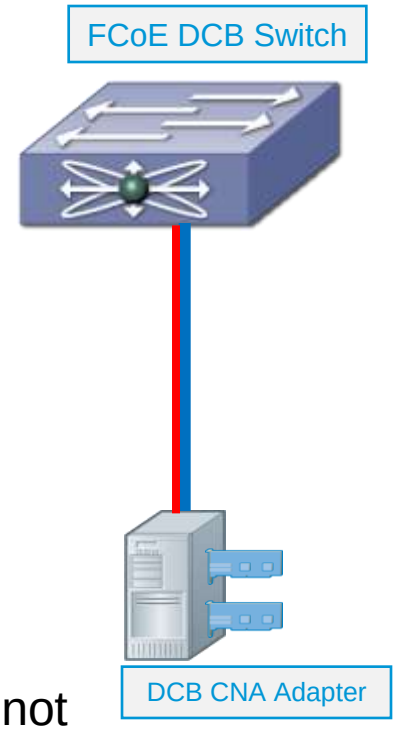

- If a custom FCoE QOS policy is configured and FCoE feature is not enabled, it will turn feature on and use the custom FCoE config

- **Best Practice** - use the default COS value of **3** for **FCoE/no-drop** traffic
- Can be changed through QOS class-map configuration

```
Type qos class-maps
===================================
  class-map type qos match-any class-fcoe
    match cos 3
  class-map type qos match-any class-default
    match any
```

Policy-maps for N5K FCoE

Applied to the System

```
system qos
  service-policy type qos input fcoe-default-in-policy
  service-policy type queuing input fcoe-default-in-policy
  service-policy type queuing output fcoe-default-out-policy
  service-policy type network-qos fcoe-default-nq-policy
```

Type **network-qos** policy-maps

```
policy-map type network-qos fcoe-default-nq-policy
  class type network-qos class-fcoe
    match qos-group 1
    pause no-drop
    mtu 2158
```

QOS

```
Service-policy (qos) input:   fcoe-default-in-policy
  policy statistics status:   disabled

  Class-map (qos):   class-fcoe (match-any)
    Match: cos 3
    set qos-group 1
```

Bandwidth IN

```
Service-policy (queuing) input:   fcoe-default-in-policy
    policy statistics status:   disabled

    Class-map (queuing):   class-fcoe (match-any)
      Match: qos-group 1
      bandwidth percent 50
```

Bandwidth OUT

```
Service-policy (queuing) output: fcoe-default-out-policy
  policy statistics status:   disabled

  Class-map (queuing):   class-fcoe (match-any)
    Match: qos-group 1
    bandwidth percent 50
```

Priority Flow Control Nexus N5K

PFC Settings to create other no-drop class-maps

- Create classification rules first by defining and applying policy-map type qos

```
pod3-5010-2(config)# class-map type qos class-lossless
pod3-5010-2(config-cmap-qos)# match cos 4
pod3-5010-2(config-cmap-qos)# policy-map type qos policy-lossless
pod3-5010-2(config-pmap-qos)# class type qos class-lossless
pod3-5010-2(config-pmap-c-qos)# set qos-group 4
pod3-5010-2(config-pmap-uf)# system qos
pod3-5010-2(config-sys-qos)# service-policy type qos input policy-lossless
```

DCBX protocol to negotiate PFC for priority 4

- Define and apply policy-map type network-qos

```
pod3-5010-2(config-pmap-qos)# class type network-qos policy-lossless
pod3-5010-2(config-cmap-uf)# match qos-group 4
pod3-5010-2(config-cmap-uf)# policy-map type network-qos policy-lossless
pod3-5010-2(config-pmap-uf)# class type network-qos class-lossless
pod3-5010-2(config-pmap-uf-c)# pause no-drop
pod3-5010-2(config-pmap-uf)# system qos
pod3-5010-2(config-sys-qos)# service-policy type network-qos policy-lossless
```

Priority Flow Control – Nexus 7K & MDS

Operations Configuration – Switch Level

N7K-50(config)# *system qos*
N7K-50(config-sys-qos)# *service-policy type network-qos default-nq-7e-policy*

- **No-Drop PFC w/ MTU 2K set for Fibre Channel**

```
show policy-map system
  Type network-qos policy-maps
  ====================================
  policy-map type network-qos default-nq-7e-policy
    class type network-qos c-nq-7e-drop
      match cos 0-2,4-7
      congestion-control tail-drop
      mtu 1500
    class type network-qos c-nq-7e-ndrop-fcoe
      match cos 3
      match protocol fcoe
      pause
      mtu 2112
```

```
show class-map type network-qos c-nq-7e-ndrop-fcoe

  Type network-qos class-maps
  ===========================================
  class-map type network-qos match-any c-nq-7e-ndrop-fcoe
      Description: 7E No-Drop FCoE CoS map
    match cos 3
    match protocol fcoe
```

Policy Template choices

Template	Drop CoS	(Priority)	NoDrop CoS	(Priority)
default-nq-8e-policy	0,1,2,3,4,5,6,7	5,6,7	-	-
default-nq-7e-policy	0,1,2,4,5,6,7	5,6,7	3	-
default-nq-6e-policy	0,1,2,5,6,7	5,6,7	3,4	4
default-nq-4e-policy	0,5,6,7	5,6,7	1,2,3,4	4

Enhanced Transmission Selection (ETS)
Bandwidth Management – 802.1Qaz

- Prevents a single traffic class of “hogging” all the bandwidth and starving other classes
- When a given load doesn’t fully utilize its allocated bandwidth, it is available to other classes
- Helps accommodate for classes of a “bursty” nature

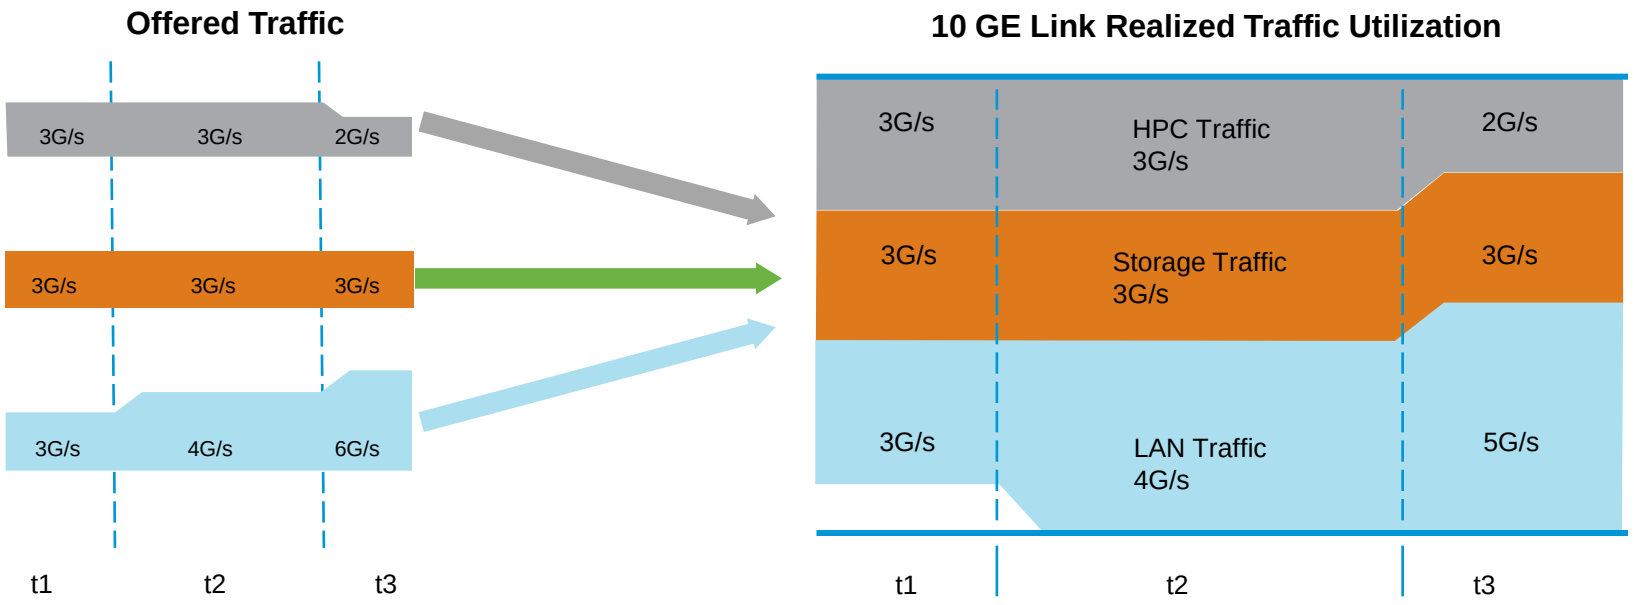

Enhanced Transmission Selection - N5K
Bandwidth Management

- Once **feature fcoe** is configured, 2 classes are made **by default**
- By default, each class is given **50%** of the available bandwidth

```
N5k-1# show queuing interface ethernet 1/18
Ethernet1/18 queuing information:
  TX Queuing
    qos-group  sched-type  oper-bandwidth
        0       WRR            50
        1       WRR            50
```

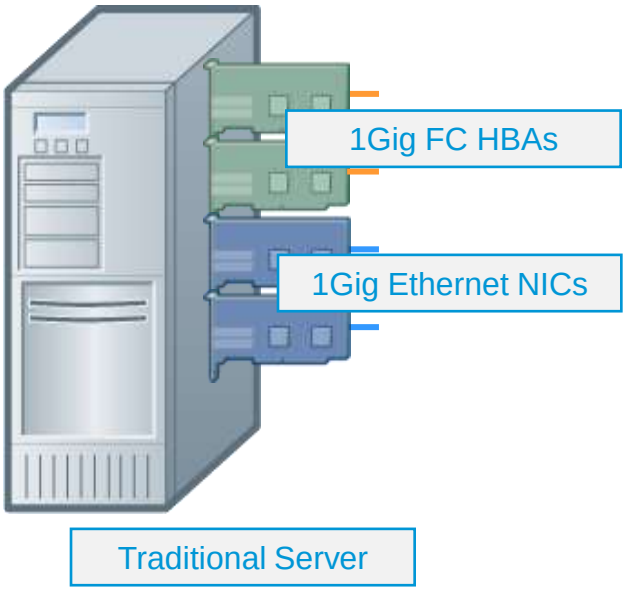

A typical server has equal BW per traffic type

- Best Practice on N5K : FCoE and Ethernet each receive 50%
- Can be changed through QoS settings when higher demands for certain traffic exist (i.e. HPC traffic, more Ethernet NICs)

Enhanced Transmission Selection – N7K
Bandwidth Management

n7k-50-fcoe-2# show queuing interface ethernet 4/17

Egress Queuing for Ethernet4/17 [System]

Template: 4Q7E

Group	Bandwidth%	PrioLevel	Shape%
0	80	-	-
1	**20**	**-**	**-**

Que#	Group	Bandwidth%	PrioLevel	Shape%	CoSMap
0	0	-	High	-	5-7
1	**1**	**100**	**-**	**-**	**3**
2	0	50	-	-	2,4
3	0	50	-	-	0-1

Ingress Queuing for Ethernet4/17 [System]

Trust: Trusted

Group	Qlimit%
0	70
1	**30**

Que#	Group	Qlimit%	IVL	CoSMap
0	0	45	0	0-1
1	0	10	5	5-7
2	**1**	**100**	**3**	**3**
3	0	45	2	2,4

MDS switch ETS

- MDS FCoE linecard does not compete with the Fibre Channel bandwidth and received 100% of the ethernet bandwidth

```
9513-71# sh policy-map interface ethernet 13/1

Global statistics status :   enabled

Ethernet 13/1

  Service-policy (queuing) input:   default-4q-7e-in-policy
    policy statistics status:   enabled (current status: enabled)

    Class-map (queuing):   1q4t-7e-in-q-default (match-any)
      queue-limit percent 100
      bandwidth percent 100
      queue dropped pkts : 0

  Service-policy (queuing) output:   default-4q-7e-out-policy
    policy statistics status:   enabled (current status: enabled)

    Class-map (queuing):   1q1t-7e-out-q-default (match-any)
      bandwidth remaining percent 100
      queue dropped pkts : 0
```

```
9513-71# sh policy-map system type queuing


  Service-policy (queuing) input:   default-4q-7e-in-policy
    policy statistics status:   disabled (current status: disabled)

    Class-map (queuing):   1q4t-7e-in-q-default (match-any)
      queue-limit percent 100
      bandwidth percent 100

  Service-policy (queuing) output:   default-4q-7e-out-policy
    policy statistics status:   disabled (current status: disabled)

    Class-map (queuing):   1q1t-7e-out-q-default (match-any)
      bandwidth remaining percent 100
```

Agenda

FCoE Protocol Fundamentals

Fiber Channel over Ethernet

FC-BB-5 Protocol

incits InterNational Committee for Information Technology Standards
Where IT all begins

FCoE

Mapping of FC Frames over Ethernet

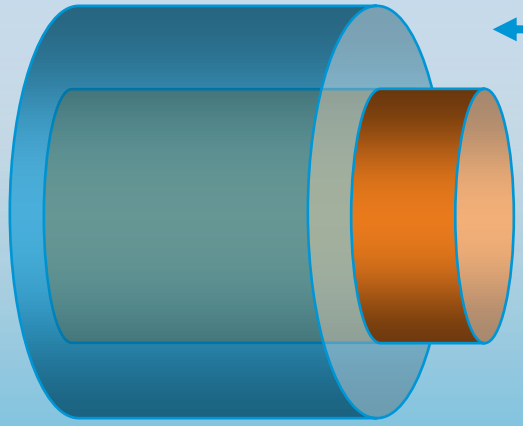

- Completely based on the FC model
- Same host-to-switch and switch-to-switch behavior as FC
- WWNs, FC-IDs, hard/soft zoning, DNS, RSCN

FCoE is Fibre Channel

Cisco	Dell	EMC²	EMULEX
HP	IBM	Microsoft	NetApp
Intel	QLOGIC	Redhat	VMWARE

Fiber Channel over Ethernet

FC-BB-5 Protocol

InterNational Committee for Information Technology Standards
Where IT all begins

FCoE

Mapping of FC Frames over Ethernet

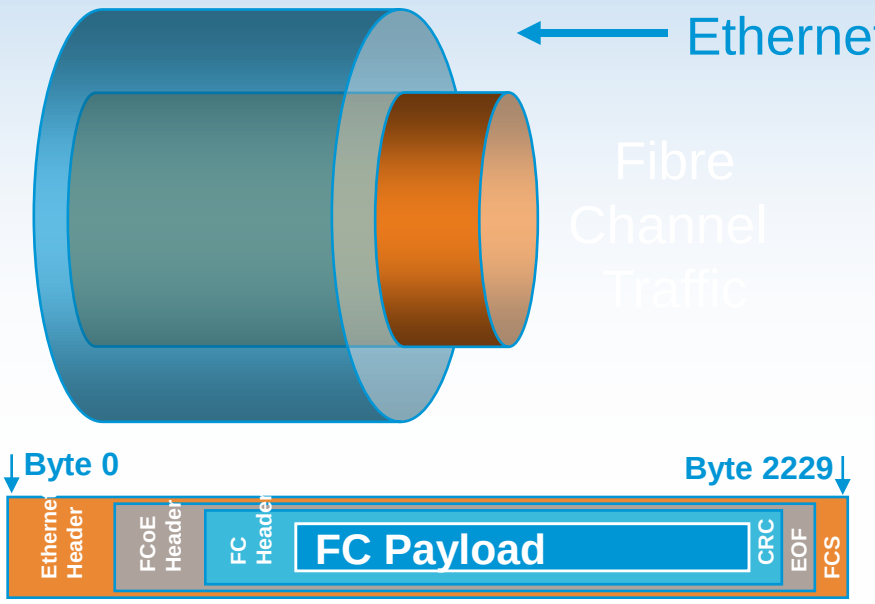

- Roadmap of Ethernet
- Economy of Scale
- Massive industry investment

FCoE is Ethernet

FCoE Frame

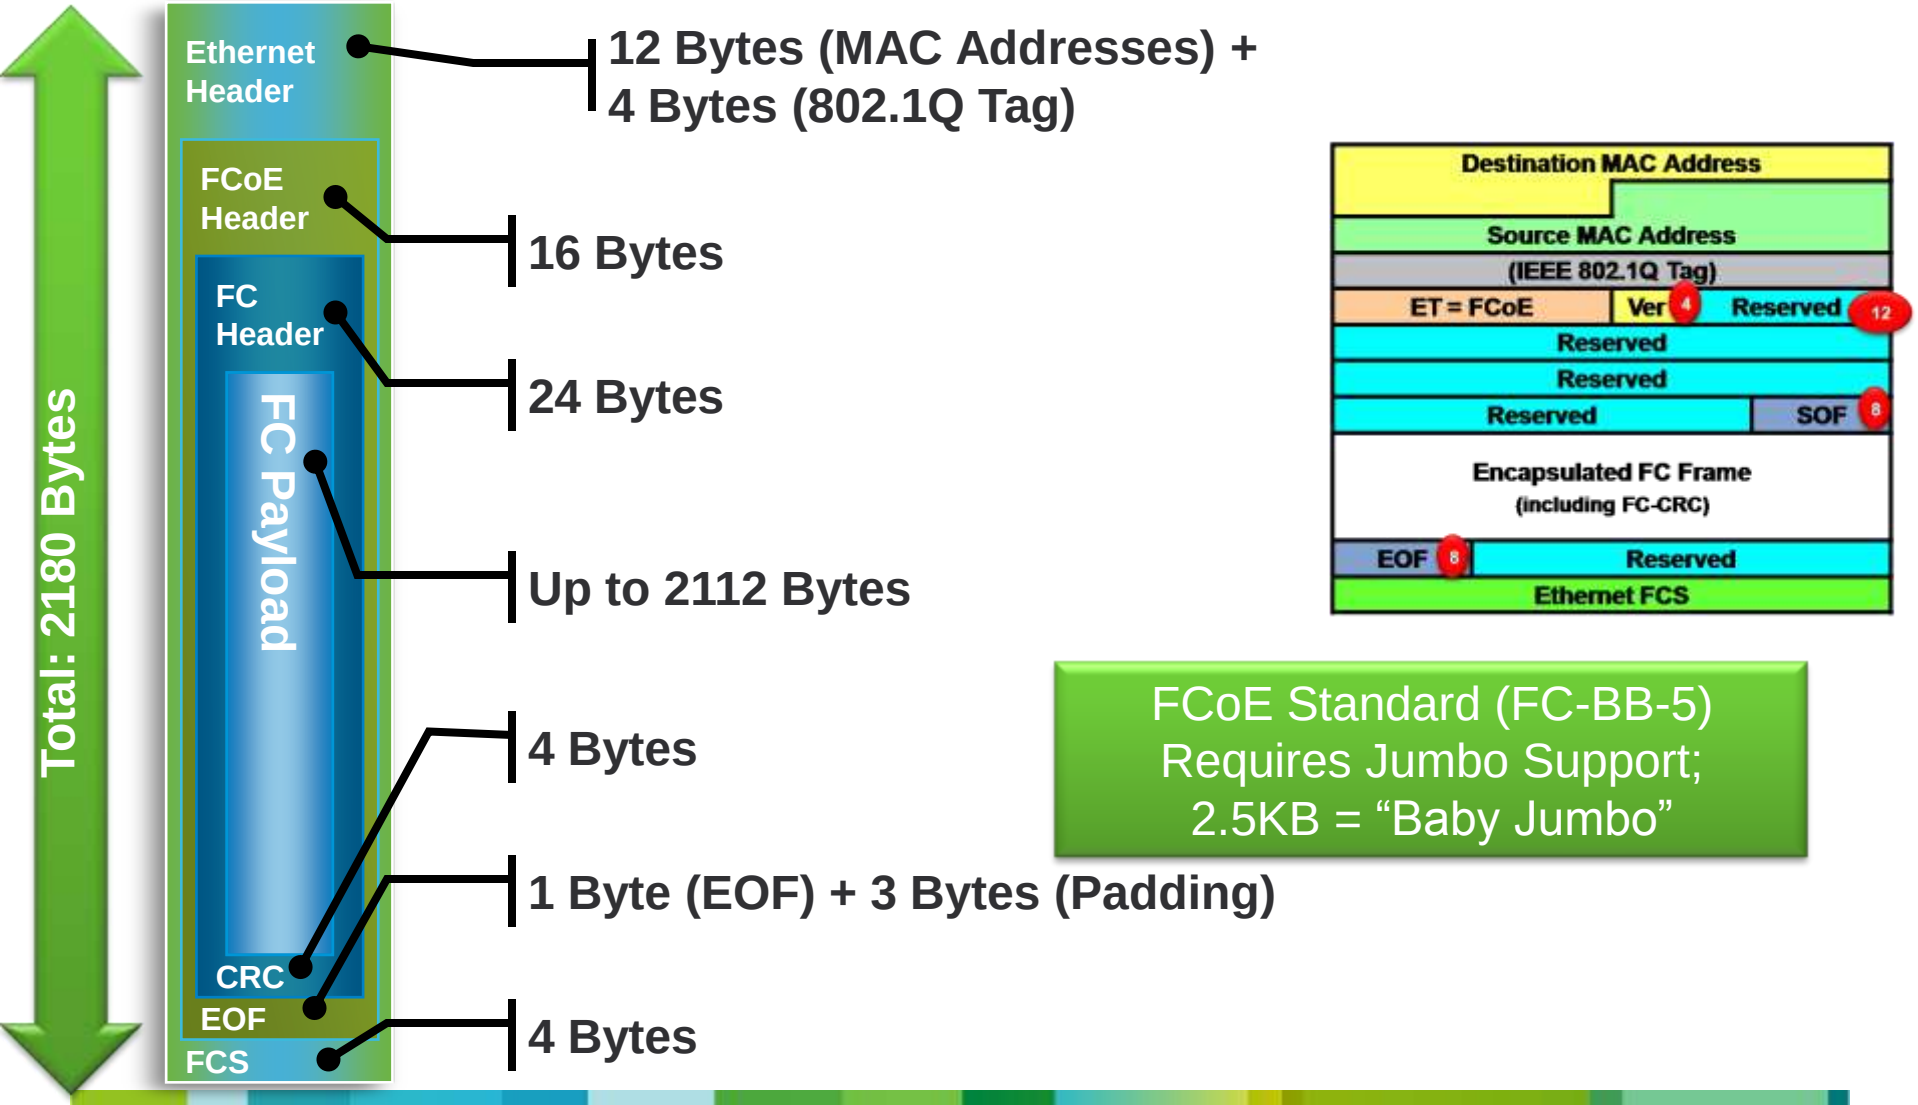

FCoE Standard (FC-BB-5)
Requires Jumbo Support;
2.5KB = “Baby Jumbo”

Storage Protocol Technologies

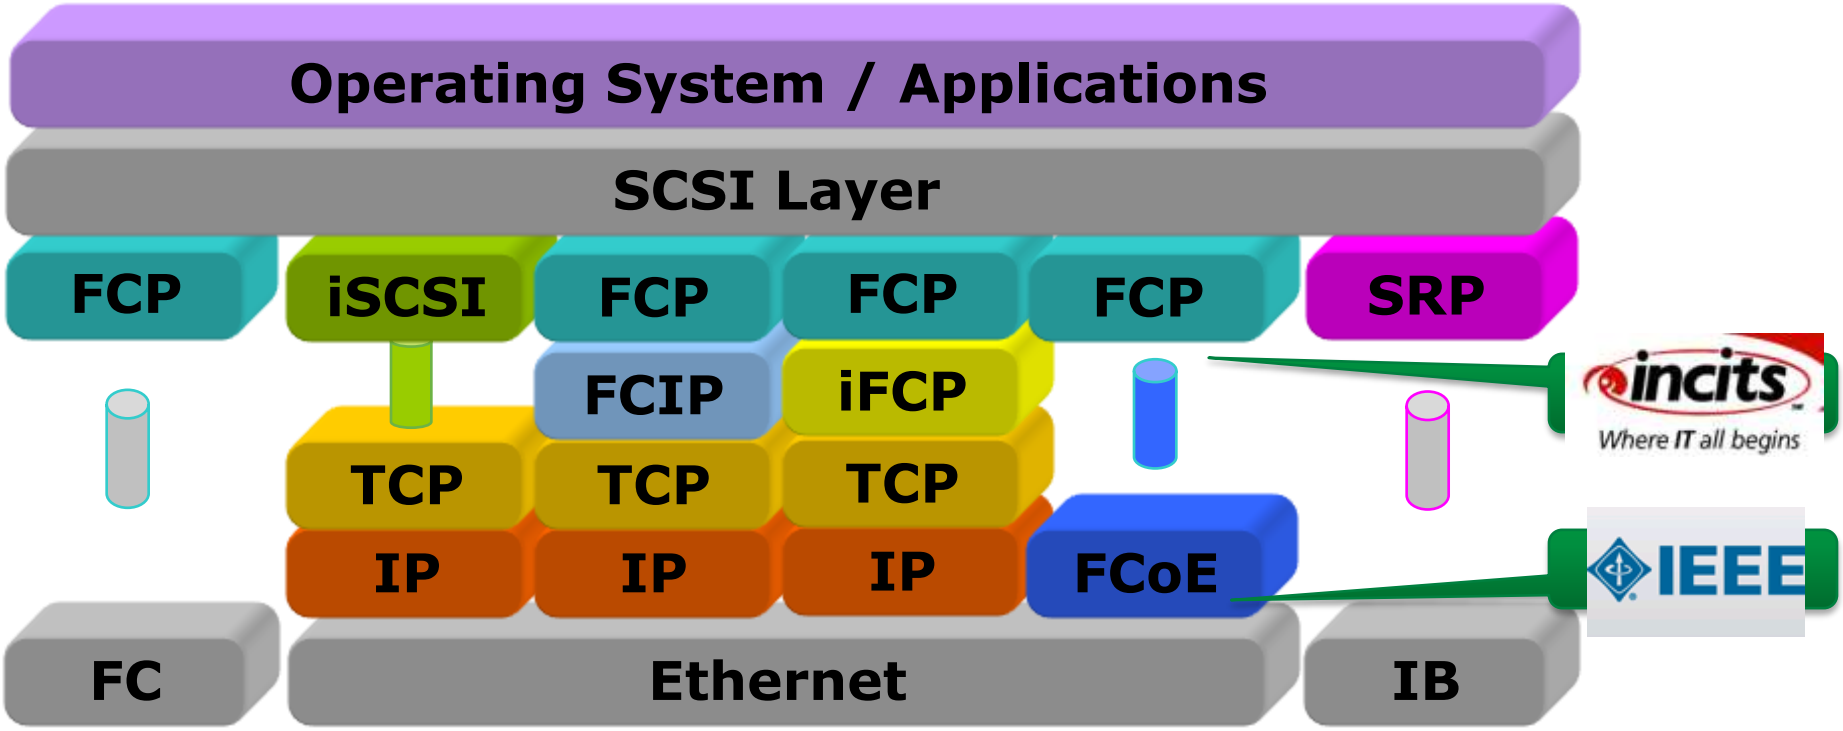

FCoE Is Non Routable, Localized DC Transport Solution with Lower Protocol Overhead than FCIP or iSCSI

FCoE Frame vs. iSCSI and Ethernet

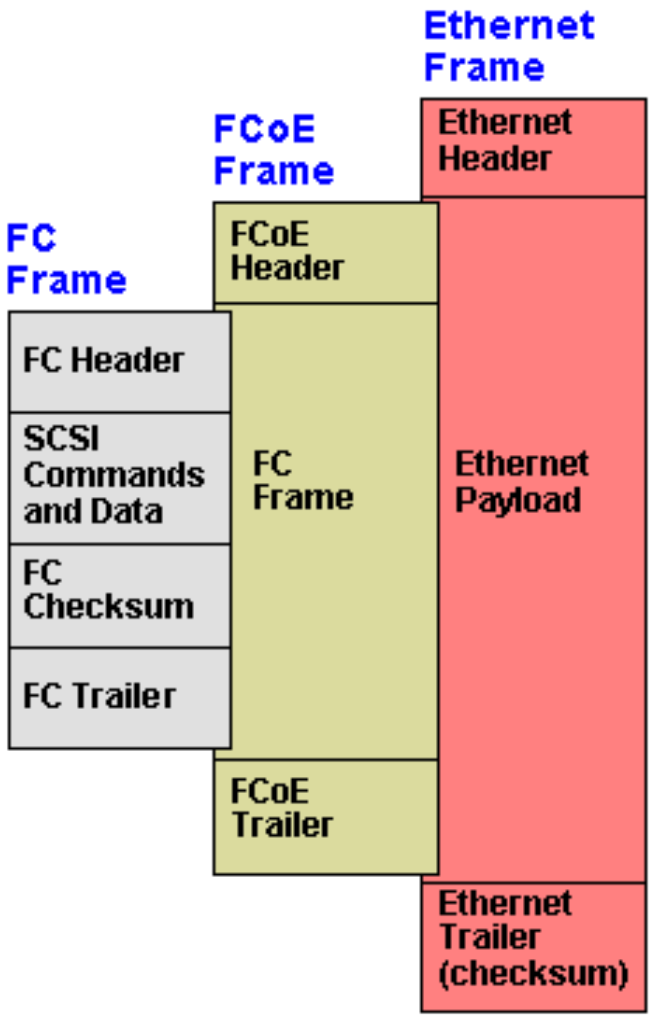

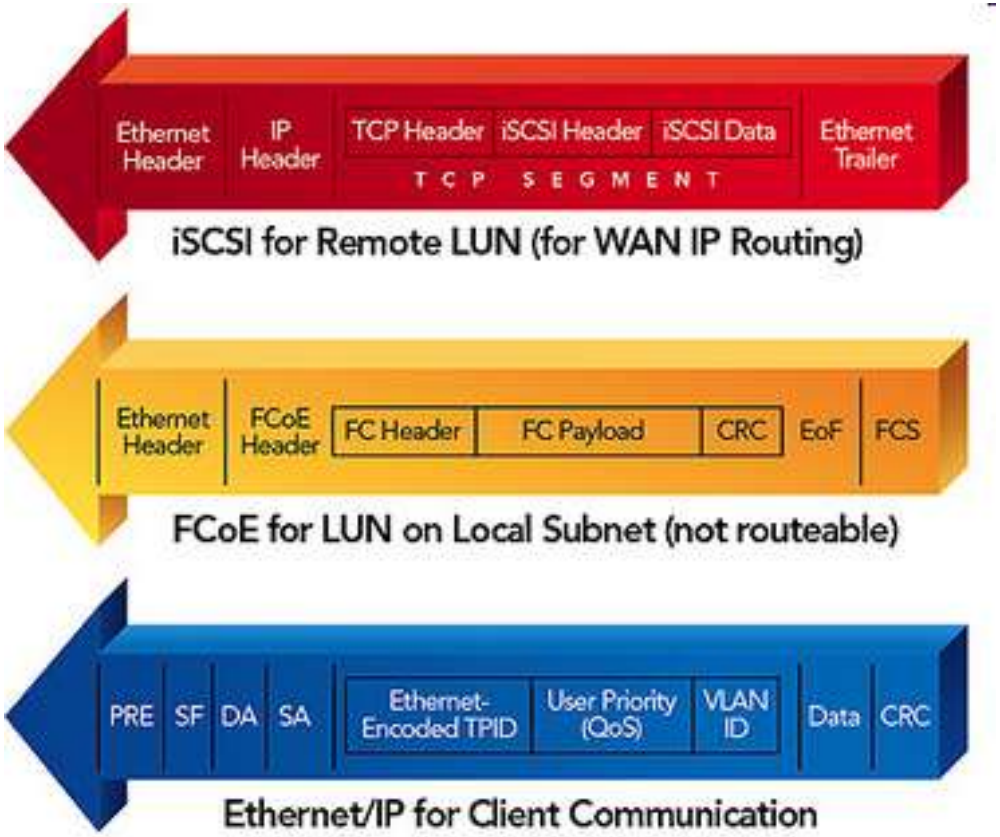

Fiber Channel over Ethernet
Protocol Mapping

- **From a Fibre Channel standpoint it is:**

  FC connectivity over a new type of cable called… Ethernet

- **From an Ethernet standpoints it is:**

  Yet another ULP (Upper Layer Protocol) to be transported

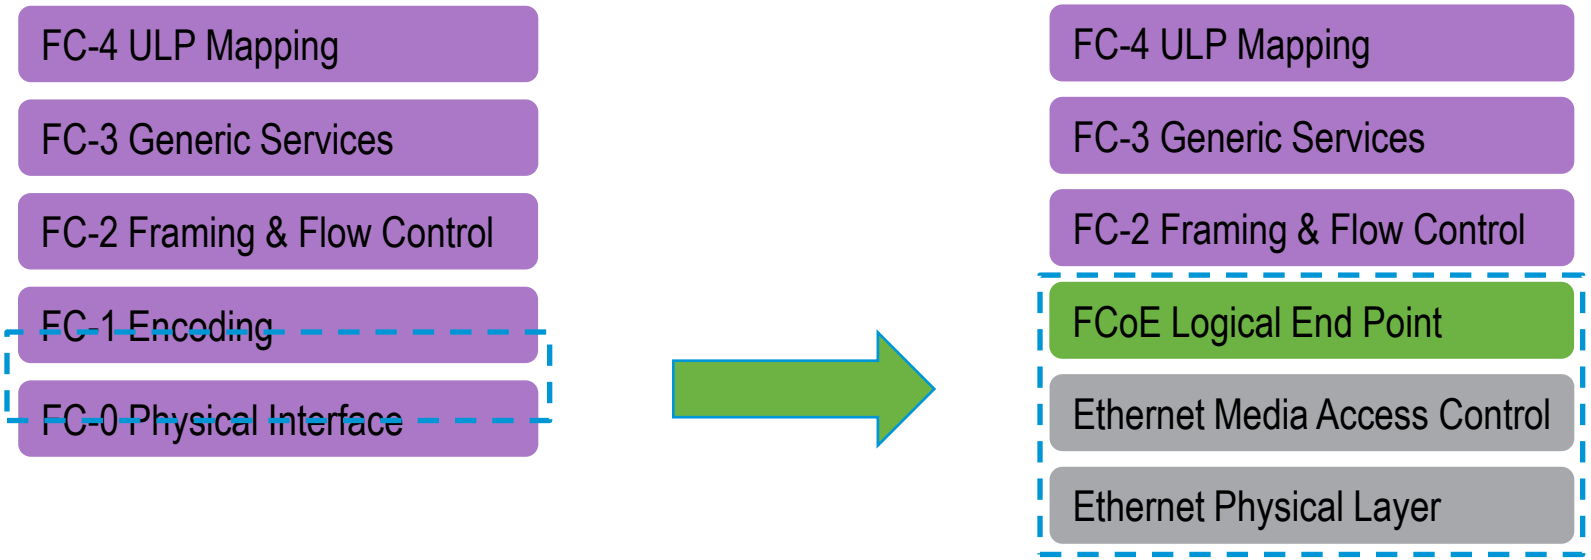

Fiber Channel over Ethernet

Data and Control plane

FCoE is really two different protocols:

FCoE

- Is the data plane protocol
- It is used to carry most of the FC frames and all the SCSI traffic
- Uses Fabric Assigned MAC address (dynamic)

FIP (FCoE Initialization Protocol)

- It is the control plane protocol
- It is used to discover the FC entities connected to an Ethernet cloud
- It is also used to login to and logout from the FC fabric

Both Protocols Have…

Two different Ethertypes, 0x8906 = FCoE, 0x8914 = FIP

Two different frame formats

Both are defined in FC-BB-5

http://www.cisco.biz/en/US/prod/collateral/switches/ps9441/ps9670/white_paper_c11-560403.html

Fibre Channel over Ethernet Protocol

FIP: FCoE Initialization Protocol

- FIP discovers other FCoE capable devices within the Ethernet Cloud
  - Enables FCoE adapters (CNAs) to discover FCoE switches (FCFs) on the FCoE VLAN
  - Establishes avirtual link with between the adapter and FCF or between two FCFs
- FIP frames use a different Ethertype from FCoE frames making FIP-Snooping by DCB capable Ethernet bridges
- Building foundation for future multi-hop FCoE topologies
  - **Multi-hop refers to FCoE extending beyond a single “hop” or “access” switch**

Fiber Channel over Ethernet Protocol

FCoE Initialization Protocol (FIP)

- **Step 1:** FCoE VLAN Discovery

  FIP sends out a multicast to ALL_FCF_MAC address looking for the FCoE VLAN

  FIP frames use the **native VLAN**

- **Step 2:** FCF Discovery

  FIP sends out a multicast to the ALL_FCF_MAC address on the FCoE VLAN to find the FCFs answering for that FCoE VLAN

  FCF'sresponds back with their MAC address

- **Step 3:** Fabric Login

  FIP sends a FLOGI request to the FCF_MAC found in Step 2

  Establishes a virtual link between host and FCF

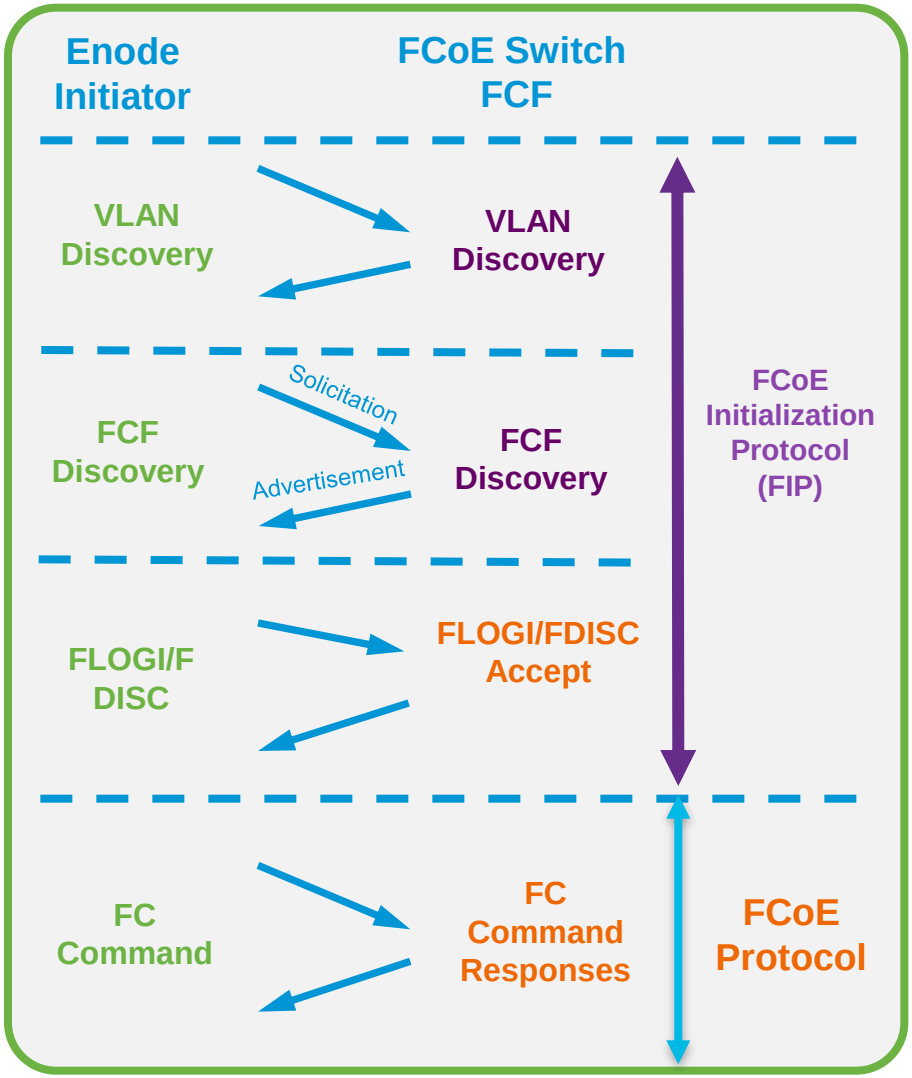

Fiber Channel over Ethernet Protocol

FCoE Initialization Protocol (FIP)

- The FCoE VLAN is manually configured on the Nexus 5K

```
tme-n5k-2# conf t
Enter configuration commands, one per line.  End with CNTL/Z.
tme-n5k-2(config)# vlan 2
tme-n5k-2(config-vlan)# fcoe vsan 2
tme-n5k-2(config-vlan)# show vlan fcoe
VLAN      VSAN      Status
--------  --------  --------
2         2         Operational
```

- The FCF-MAC address is configured on the Nexus 5K by default once `feature fcoe` has been configured

  This is the MAC address returned in step 2 of the FIP exchange

  This MAC is used by the host to login to the FCoE fabric

```
tme-n5k-2# show fcoe
Global FCF details
        FCF-MAC is 00:0d:ec:df:5f:80
        FC-MAP is 0e:fc:00
        FCF Priority is 128
        FKA Advertisement period for FCF is 8 seconds
tme-n5k-2#
```

** FIP does not carry any Fibre Channel frames

Fiber Channel over Ethernet Protocol

FCoE Initialization Protocol (FIP)

- Login process: show flogi database and show fcoe database show the logins and associated FCIDs, xWWNs and FCoE MAC addresses

```
tme-n5k-2# show flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
vfc1             2     0xb00000  21:00:00:c0:dd:11:29:1d 20:00:00:c0:dd:11:29:1d
vfc2             2     0xb00001  21:00:00:c0:dd:11:2c:61 20:00:00:c0:dd:11:2c:61
vfc14            2     0xb00004  21:00:00:c0:dd:12:13:8f 20:00:00:c0:dd:12:13:8f
vfc15            2     0xb00005  21:00:00:c0:dd:12:13:b3 20:00:00:c0:dd:12:13:b3
vfc16            2     0xb00006  21:00:00:c0:dd:12:14:23 20:00:00:c0:dd:12:14:23
vfc25            2     0xb00008  50:0a:09:83:87:d9:6e:b7 50:0a:09:80:87:d9:6e:b7
vfc26            2     0xb00009  50:0a:09:87:87:d9:6e:b7 50:0a:09:80:87:d9:6e:b7
                         [netapp_fcoe1]
vfc30            2     0xb00007  50:0a:09:85:87:d9:6e:b7 50:0a:09:80:87:d9:6e:b7

Total number of flogi = 8.

tme-n5k-2# show fcoe database

------------------------------------------------------------------------
INTERFACE        FCID           PORT NAME               MAC ADDRESS
------------------------------------------------------------------------
vfc1             0xb00000       21:00:00:c0:dd:11:29:1d 00:c0:dd:11:29:1d
vfc2             0xb00001       21:00:00:c0:dd:11:2c:61 00:c0:dd:11:2c:61
vfc14            0xb00004       21:00:00:c0:dd:12:13:8f 00:c0:dd:12:13:8f
vfc15            0xb00005       21:00:00:c0:dd:12:13:b3 00:c0:dd:12:13:b3
vfc16            0xb00006       21:00:00:c0:dd:12:14:23 00:c0:dd:12:14:23
vfc25            0xb00008       50:0a:09:83:87:d9:6e:b7 00:c0:dd:0a:b7:82
vfc26            0xb00009       50:0a:09:87:87:d9:6e:b7 00:c0:dd:11:41:21
vfc30            0xb00007       50:0a:09:85:87:d9:6e:b7 00:c0:dd:11:0d:89
tme-n5k-2#
```

Agenda

FCoE Designs Requirements

Design Requirements
Ethernet vs Fibre Channel

- **Ethernet is non-deterministic.**
  - Flow control is destination-based
  - Relies on TCP drop-retransmission / sliding window

Networks
- Connectionless
- Logical circuits
- Unreliable transfers
- High connectivity
- Higher latency
- Longer distance
- Software intense

- **Fibre-Channel is deterministic.**
  - Flow control is source-based (B2B credits)
  - Services are fabric integrated (no loop concept)

Channels
- Connection service
- Physical circuits
- Reliable transfers
- High speed
- Low latency
- Short distance
- Hardware intense

Design Requirements

Classical Ethernet

- Ethernet/IP

  Goal is to provide any-to-any connectivity

  Unaware of packet loss – relies on ULPs for retransmission and windowing

  Provides the transport without worrying about the services -*services provided by upper layers*

  East-west vs north-south traffic ratios are undefined

- Network design has been optimized for:

  High Availability from a transport perspective by connecting nodes in mesh architectures

  Service HA is implemented separately

  Takes in to account control protocol interaction (STP, OSPF, EIGRP, L2/L3 boundary, etc…)

Fabric topology and traffic flows are highly flexible

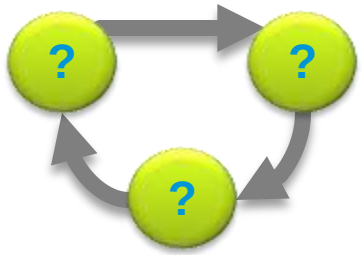

Client/Server Relationships are not pre-defined

Design Requirements

LAN Design – Access/Aggregation/Core

- Servers typically dual homed to two or more access switches
- LAN switches have redundant connections to the next layer
- Distribution and Core can be collapsed into a single box
- L2/L3 boundary typically deployed in the aggregation layer
  - Spanning tree or advanced L2 technologies (vPC) used to prevent loops within the L2 boundary
  - L3 routes are summarized to the core
- Services deployed in the L2/L3 boundary of the network (load-balancing, firewall, NAM, etc)

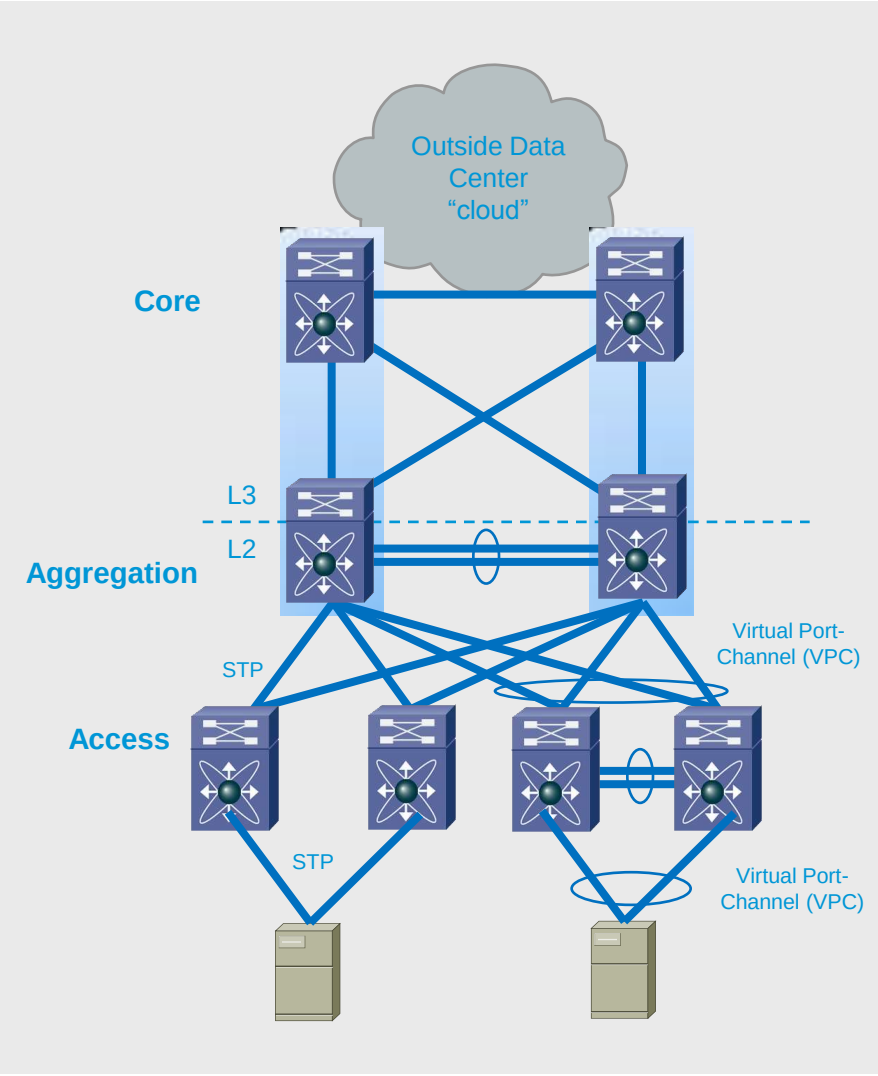

Design Requirements

Classical Fibre Channel

- **Fibre Channel SAN**
  - Transport and Services are on the same layer in the same devices
  - Well defined end device relationships (initiators and targets)
  - Does not tolerate packet drop – requires lossless transport
  - Only north-south traffic, east-west traffic mostly irrelevant
- **Network designs optimized for Scale and Availability**
  - High availability of network services provided through dual fabric architecture
    - SAN ‘A’ and SAN ‘B’ : physically separate and redundant fabrics
  - Strict change isolation - end to end driver certification

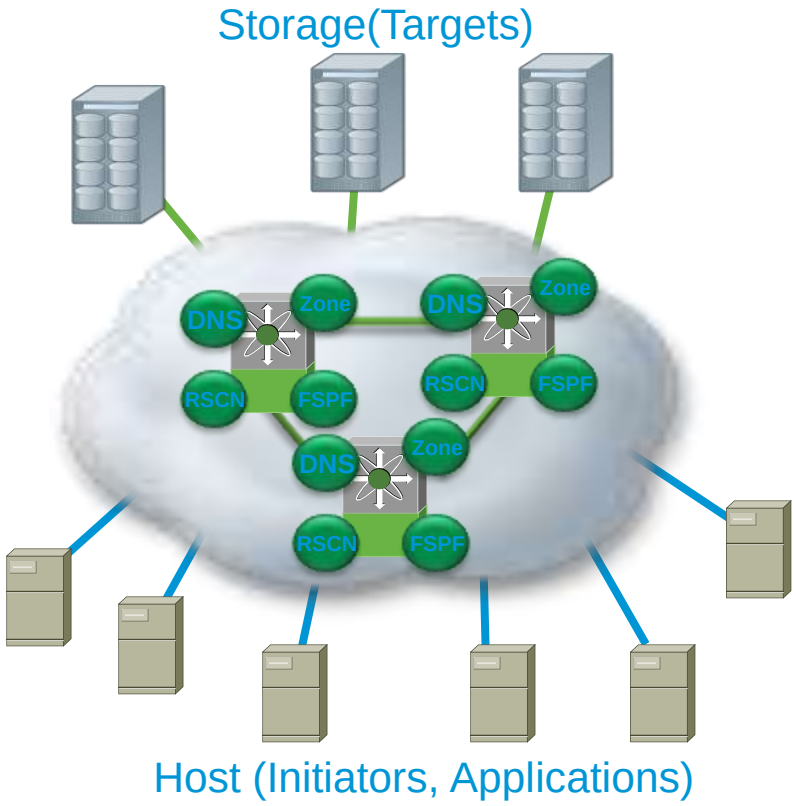

Fabric topology, services and traffic flows are structured

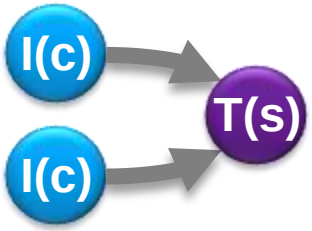

Client/Server Relationships are pre-defined

Design Requirements
SAN Design – Two Tier Topology

- "Edge-Core" Topology
- Servers connect to the edge switches
- Storage devices connect to one or more core switches
- Core switches provide storage services to one or more edge switches, thus servicing more servers in the fabric
- ISLs have to be designed so that overall fan-in ratio of servers to storage and overall end-to-end oversubscription are maintained
- HA achieved in two physically separate, but identical, redundant SAN fabrics

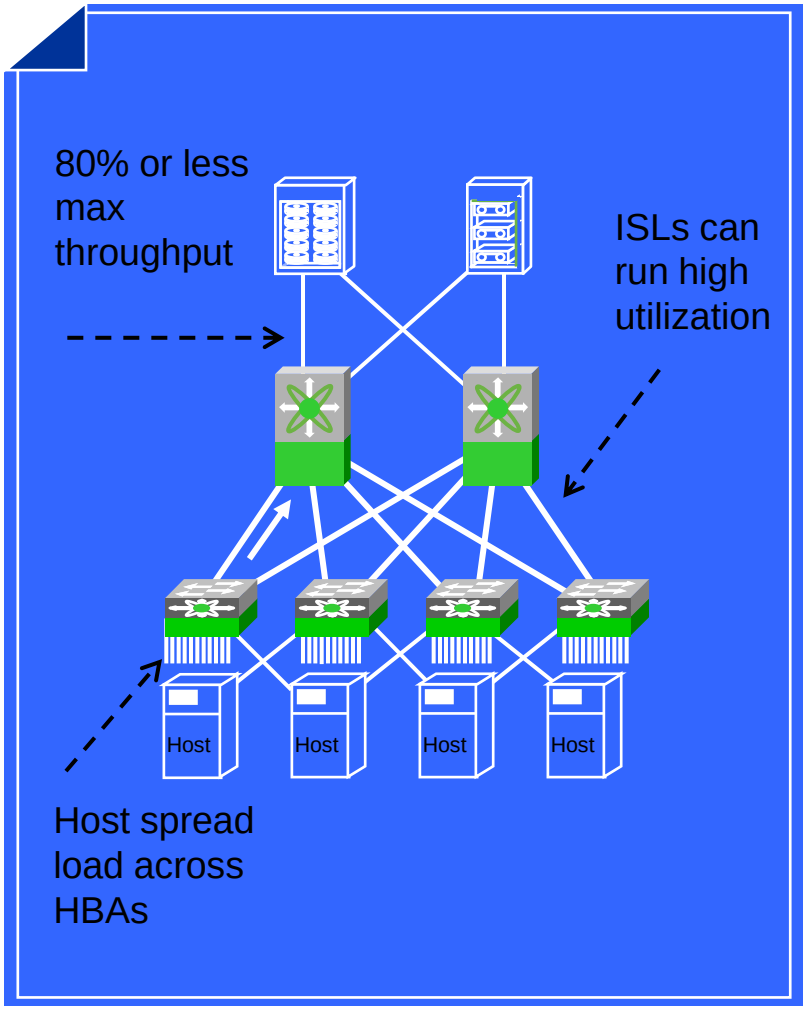

Design Requirements
SAN Design – Three Tier Topology

- "Edge-Core-Edge" Topology
- For environments where future growth of the network has the number of storage devices exceeding the number of ports available at the core switch
- A set of edge switches dedicated to server connectivity and another set of dedicated for storage devices
- Extra edge can also be "services edge" for advanced network services
- Core is for transport only, rarely accommodates end nodes
- HA achieved with dual fabrics

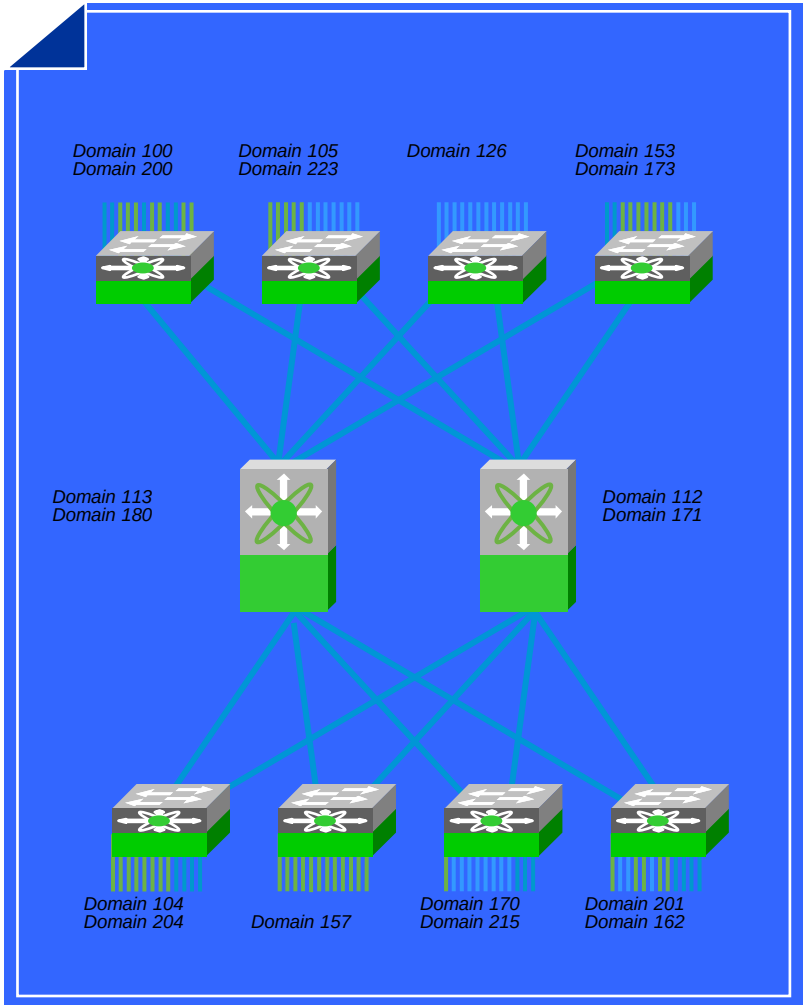

Design Requirements

Fibre Channel Protocol & Ethernet

- Unified Fabric design has to incorporate the super-set of requirements

Network -- Lossless **‘and’** Lossfull Topologies

Transport – undefined (any-to-any) **‘and’** defined (one-to-one)

High Availability – redundant network topology (mesh/full mesh) **‘and’** physically separate redundant fabrics

Bandwidth – FC fan-in and oversubscription ratios **‘and’** Ethernet oversubscription

Security – FC controls (zoning, port security,…) **‘and’** IP controls (CISF, ACL, …)

Manageability and Visibility – Hop by hop visibility for FC **‘and’** the cloud for Ethernet

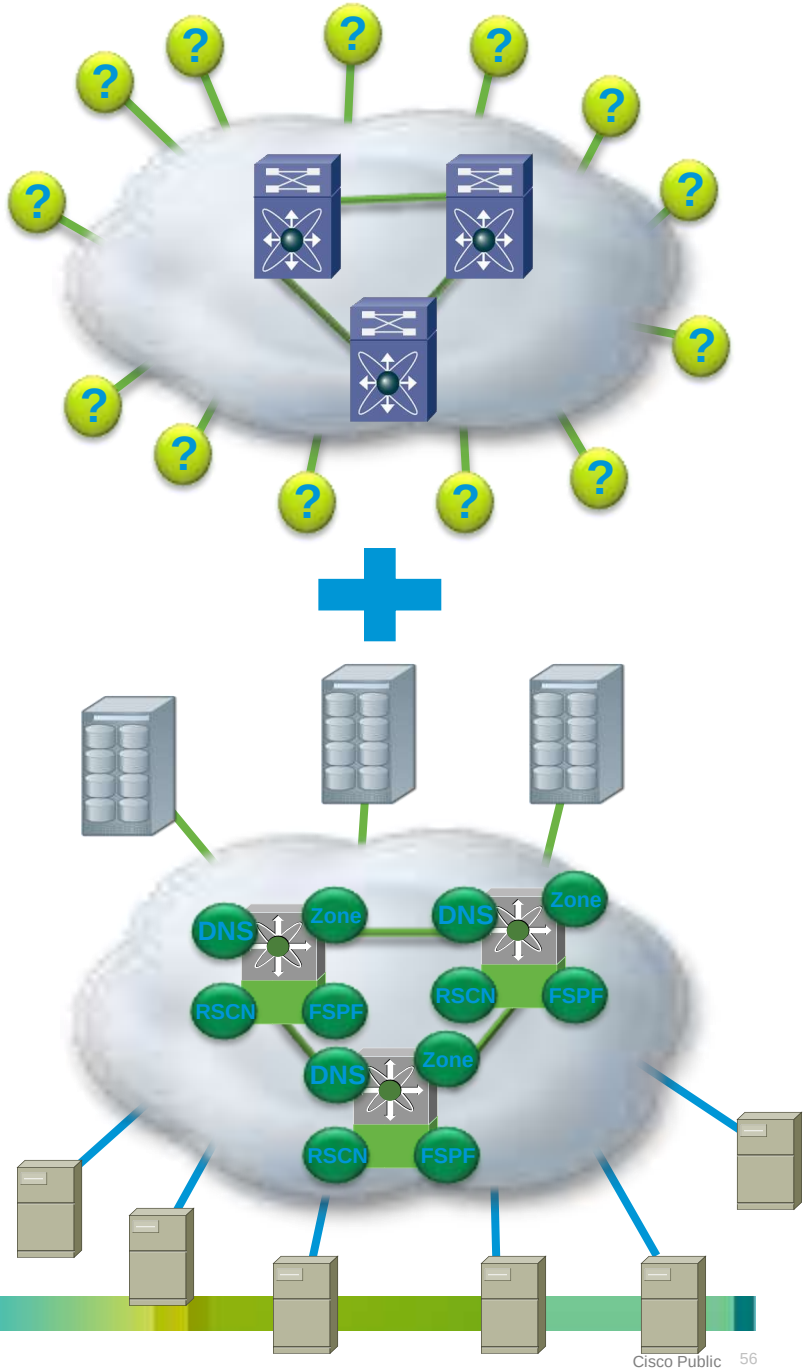

What is Fibre Channel NPIV

- N-Port ID Virtualization (NPIV) provides a means to assign multiple FCIDs to a single N_Port

  Limitation exists in FC where only a single FCID can be handed out per F-port. Therefore and F-Port can only accept a single FLOGI

- Allows multiple applications to share the same Fiber Channel adapter port
- Usage applies to applications such as VMWare, MS Virtual Server and Citrix

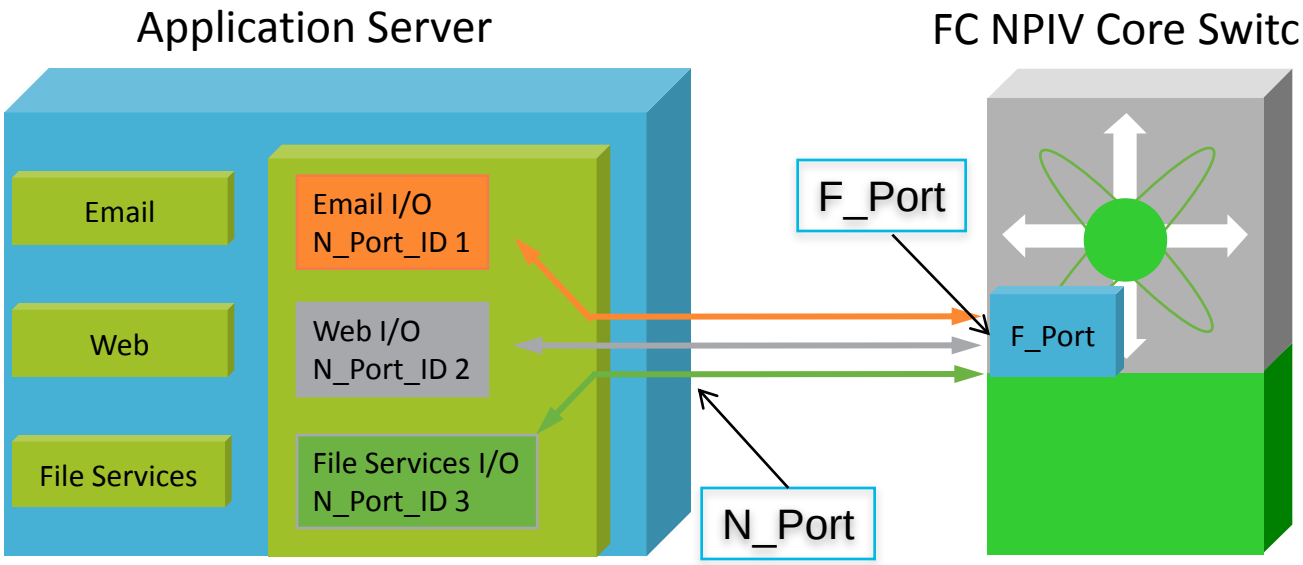

What is Fibre Channel NPV

- N-Port Virtualizer (NPV) utilizes NPIV functionality to allow a "switch" to act like a server performing multiple logins through a single physical link
- Physical servers connected to the NPV switch login to the upstream NPIV core switch
- No local switching is done on an FC switch in NPV mode
- FC edge switch in NPV mode does not take up a domain ID

  Helps to alleviate domain ID exhaustion in large fabrics

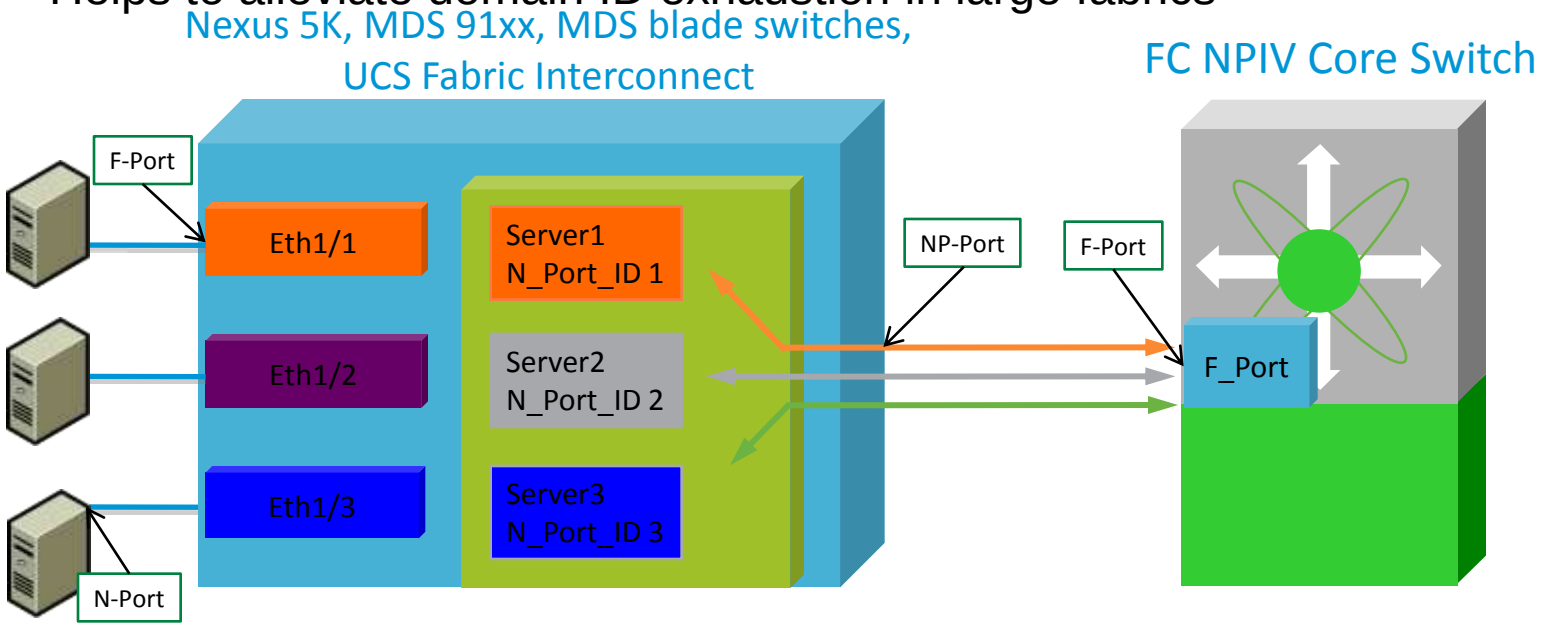

What is FCoE-NPV

- **FCoE Pass – through device**
  - All FCoE Switching is performed at the upstream FCF
  - Addressing is pass out by the upstream FCF
- **More FCoE connectivity to hosts without:**
  - Running into the domain ID issue
  - Less-expensive
  - Consistent management
- **Proxy's FIP functions between a CNA and an FCF**
  - FCoE VLAN configuration and assignment
  - FCF Assignment
- **FCoE-NPV load balance logins from the CNAs evenly across the available FCF uplink ports**

  FCoE-NPV will take VSAN into account when mapping or 'pinning' logins from a CNA to an FCF uplink
- **Operations and management process are in line with today's SAN-Admin practices**
- **Similar to NPV in a native Fibre Channel network**

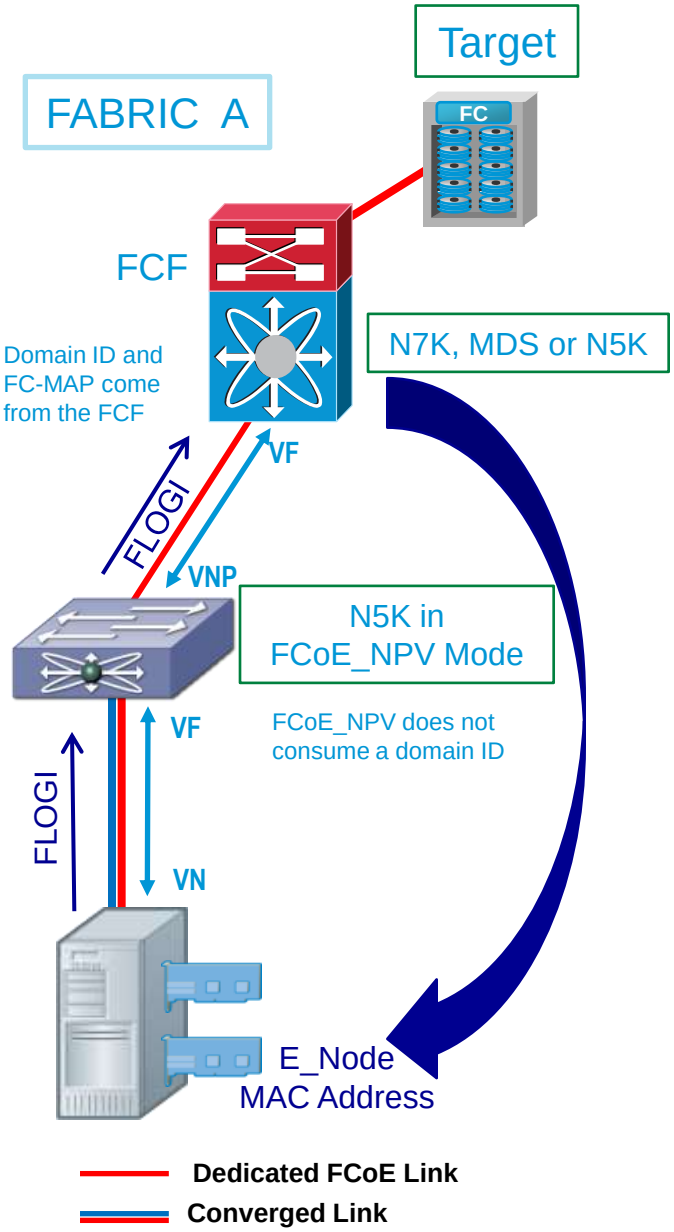

FCoE-NPV Feature

- On the control plane (FIP ethertype), an Fibre Channel over Ethernet-NPV improves over a "FIP snooping" by intelligently proxying FIP functions between a CNA and an FCF
  - takes control of how a live network will build FCoE connectivity
  - makes the connectivity very predictable, without the need for an FCF at the next hop from the CNA
- On the data plane (FCoE ethertype), an “FCoE-NPV” offers more ways to engineer traffic between CNA-facing ports and FCF-facing ports
- An “FCoE-NPV” knows nothing about Fibre Channel, and can’t parse packets with FCoE ethertype

FCoE-NPV configuration Details

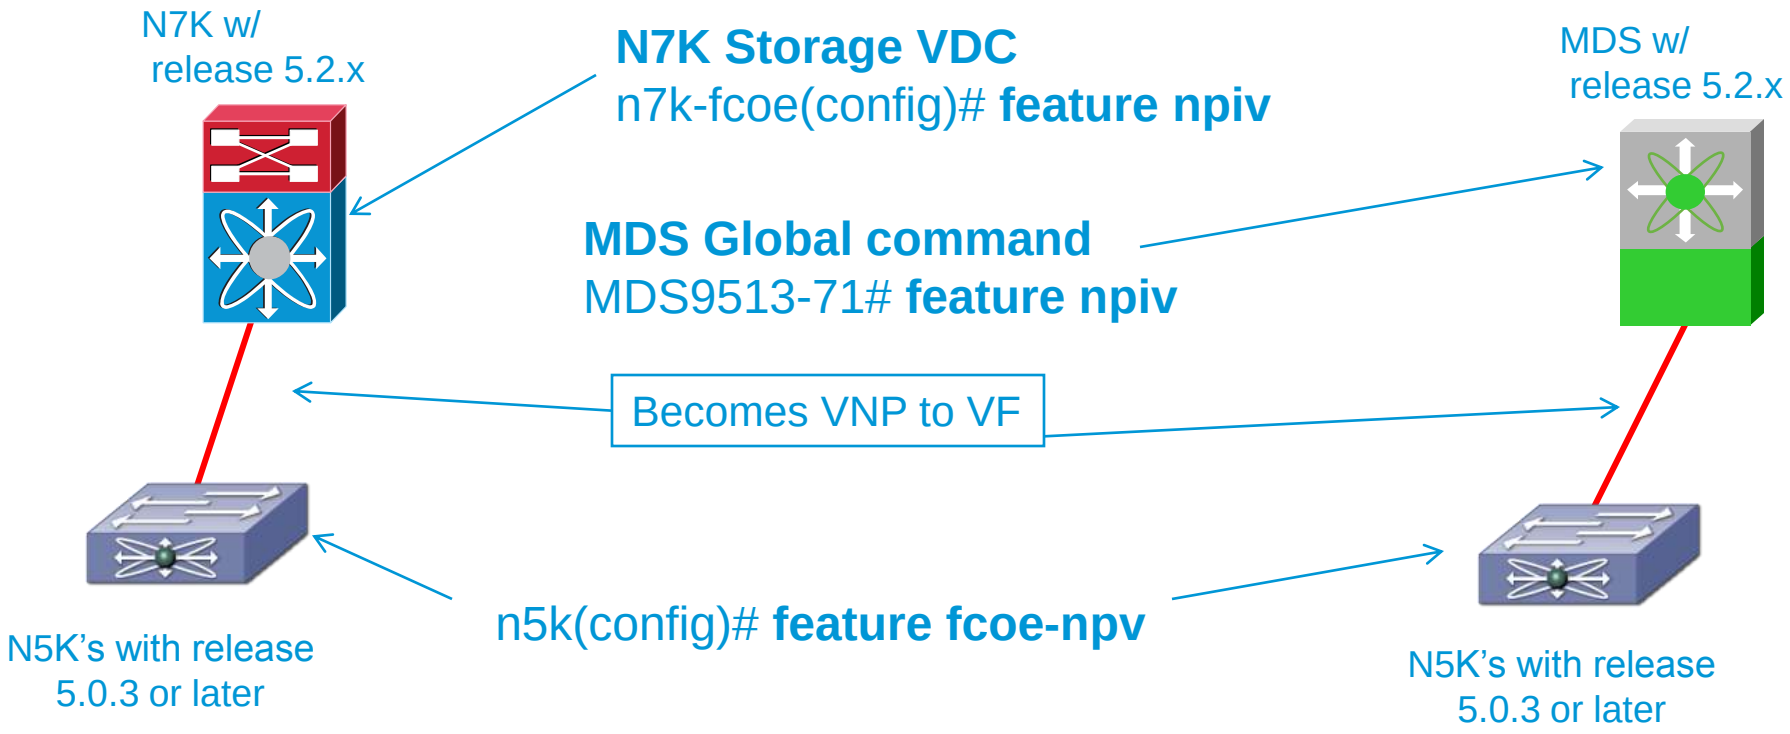

Proper no drop QOS needs to be applied to all NPIV VDC's and NPV switches as shown in earlier slides

LACP Port-channels an be configured between switches for High availability.

Zoning

- Zoning feature is required for all Fibre Channel Networks
- Default (No zoning) can be applied, but is not best practice in most all storage networks
- All Cisco FCoE switches support full T11 standard and enhanced zoning
- Best practice of Zoning methods use in existing SAN networks can be applied to Cisco FCoE switches. Most applied zoning method is use of pWWN.
- Storage networking Zoning Best Practice is Initiator to target zoning (Single host to Storage port) Best security for LUN protection.
- Zoning can be done via CLI, scripts, or Cisco Zone Management tools found via Cisco Fabric Manager (DCNM)

What is Zoning?

Zoning is a way to limit which devices can see and communicate with each other over a SAN comprised of 1 or more fabric switches.

Only devices in the same zone can communicate with each other. Commonly used to protect disk access from an undesired host.

Limits change notification to only devices that might be affected.

Disk port is a member of both zones

Hosts are members of a single zone

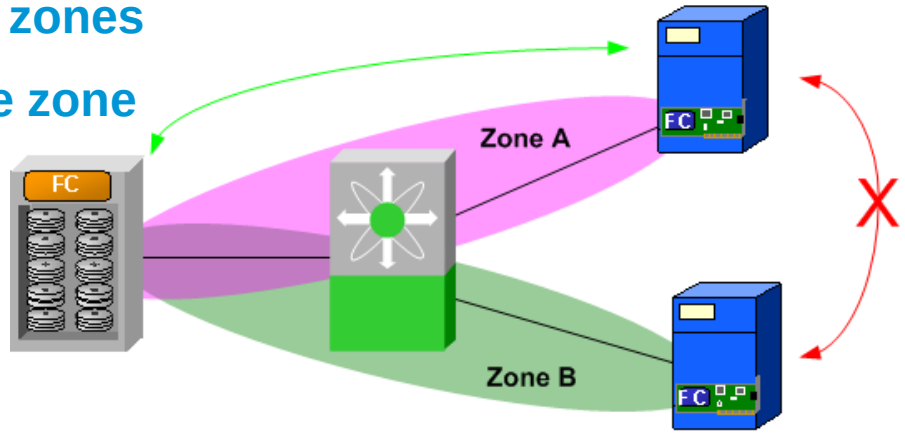

Agenda

Single Hop Designs

Single Hop Design
Converged Access, Bridging to FC SAN

- Host connected over **Converged Link** to first hop access switch
  - Access switch (Nexus 50XX) is the FCF
  - Fibre Channel ports on the access switch can be in NPV or Switch mode for native FC traffic
- DCBX is used to negotiate the enhanced Ethernet capabilities
- FIP is use to negotiate the FCoE capabilities as well as the host login process
- FCoE runs from host to access switch FCF – native Ethernet and native FC break off at the access layer

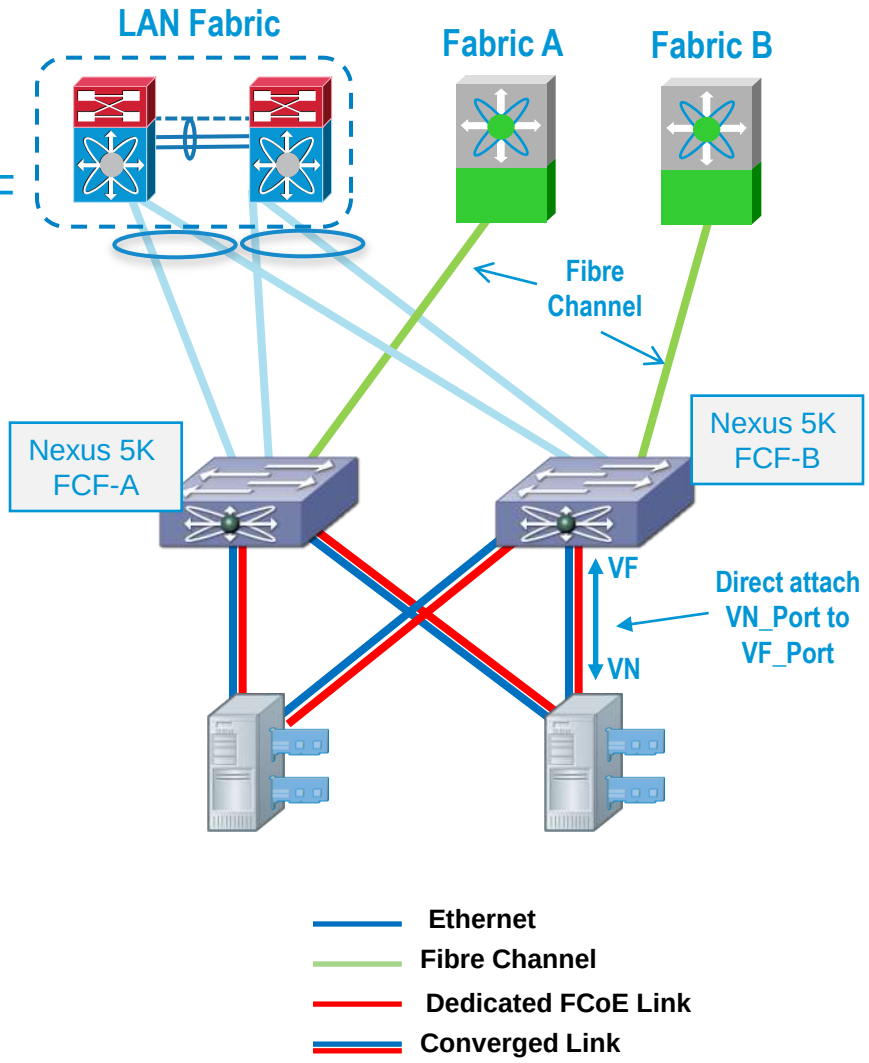

Single Hop Design
Unified Wire at the Access

- The first phase of the Unified Fabric evolution design focused on the fabric edge
- Unified the LAN Access and the SAN Edge by using FCoE
- Consolidated Adapters, Cabling and Switching at the first hop in the fabrics
- The Unified Edge supports multiple LAN and SAN topology options
  - Virtualized Data Center LAN designs
  - Fibre Channel edge with direct attached initiators and targets
  - Fibre Channel edge-core and edge-core-edge designs
  - Fibre Channel NPV edge designs

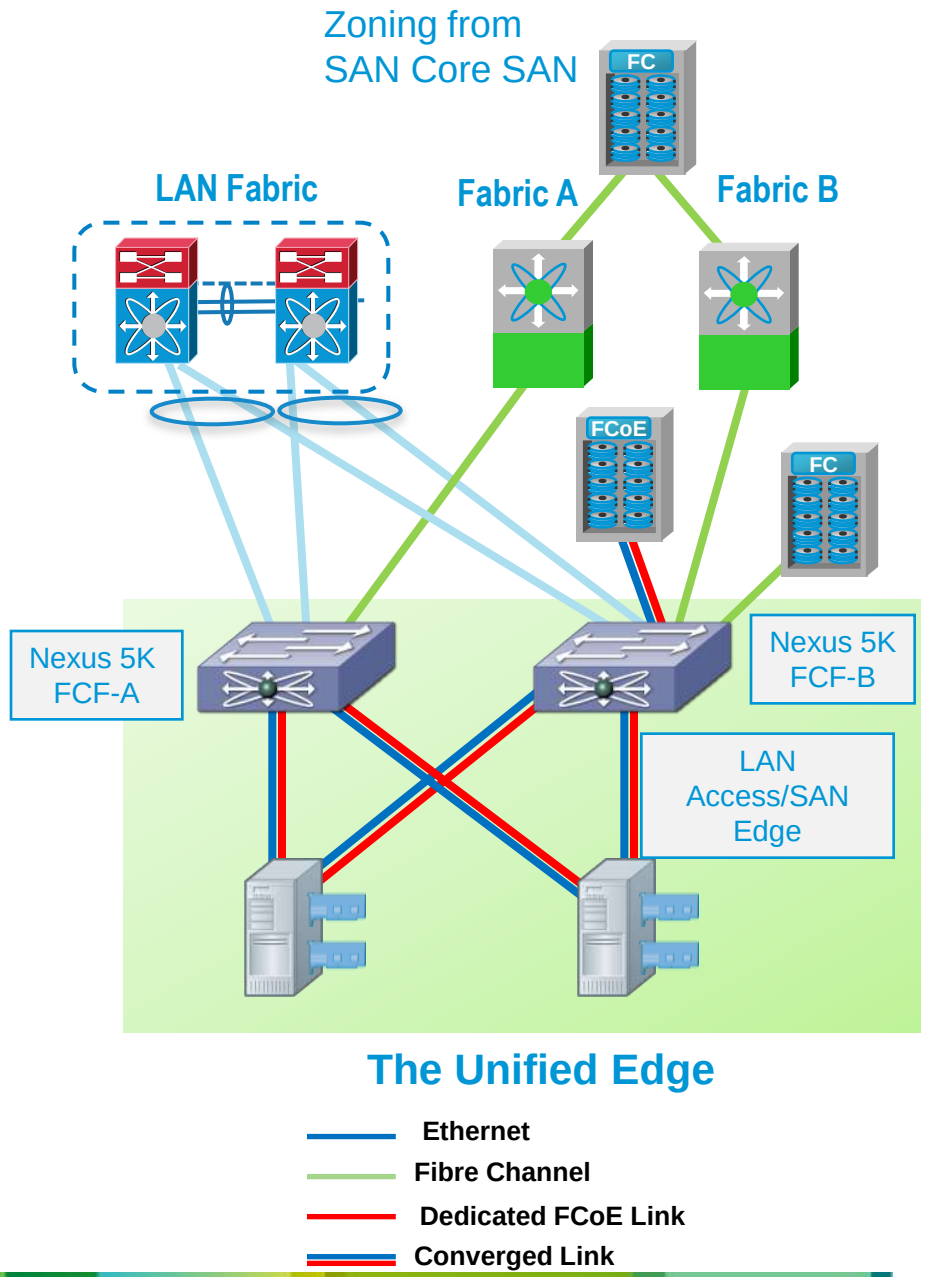

Single Hop Design
The CNA Point of View

- 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 views of the same **'Converged Link'**

  SAN Multi-Pathing provides failover between two fabrics (SAN 'A' and SAN 'B')

  NIC Teaming provides failover within the same fabric (VLAN)

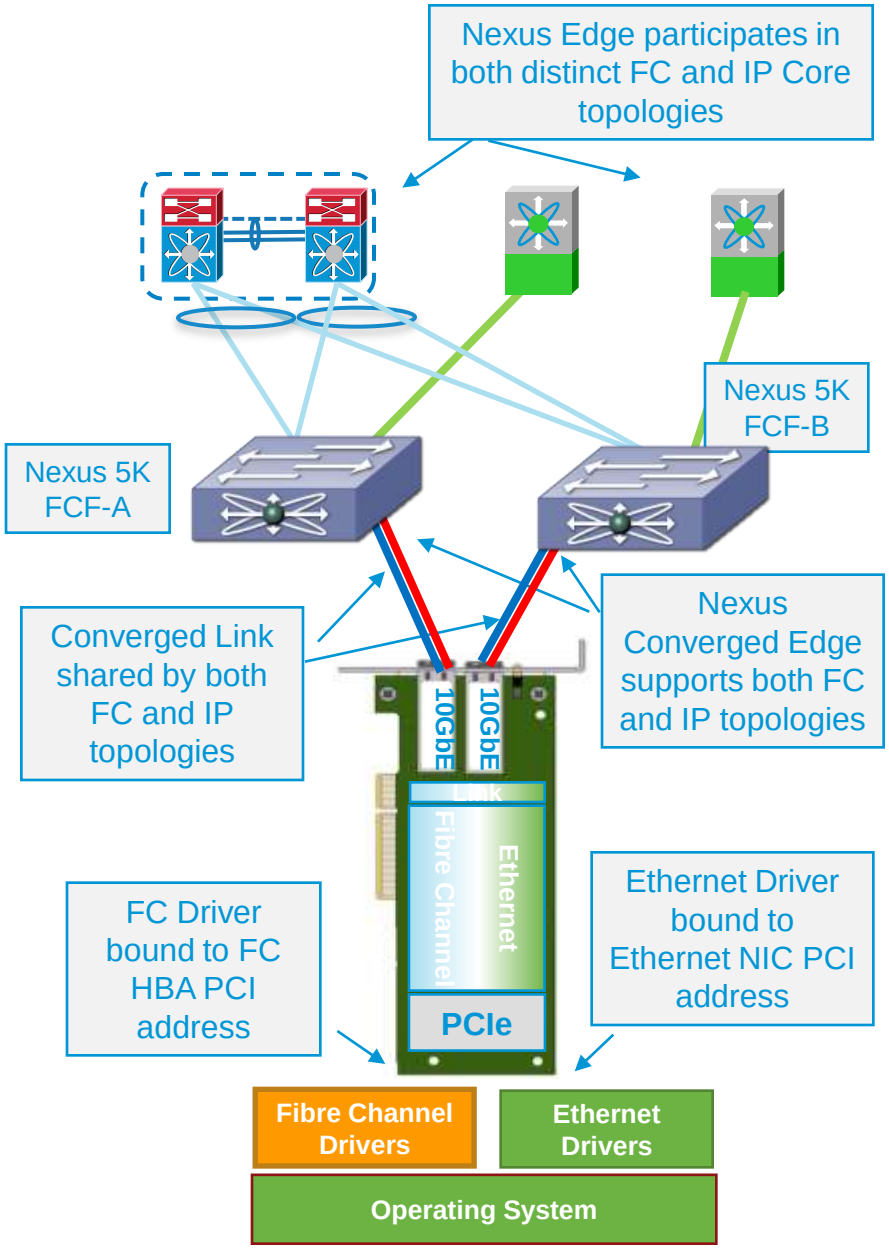

Single Hop Design
The FCoE VLAN

- A VLAN is dedicated for every VSAN in the fabric
- 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 must not be configured on 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)
- In order to maintain the integrity of FC forwarding over FCoE, FCoE VLANs are treated differently than LAN VLANs

  No flooding, MAC learning, broadcasts, etc.

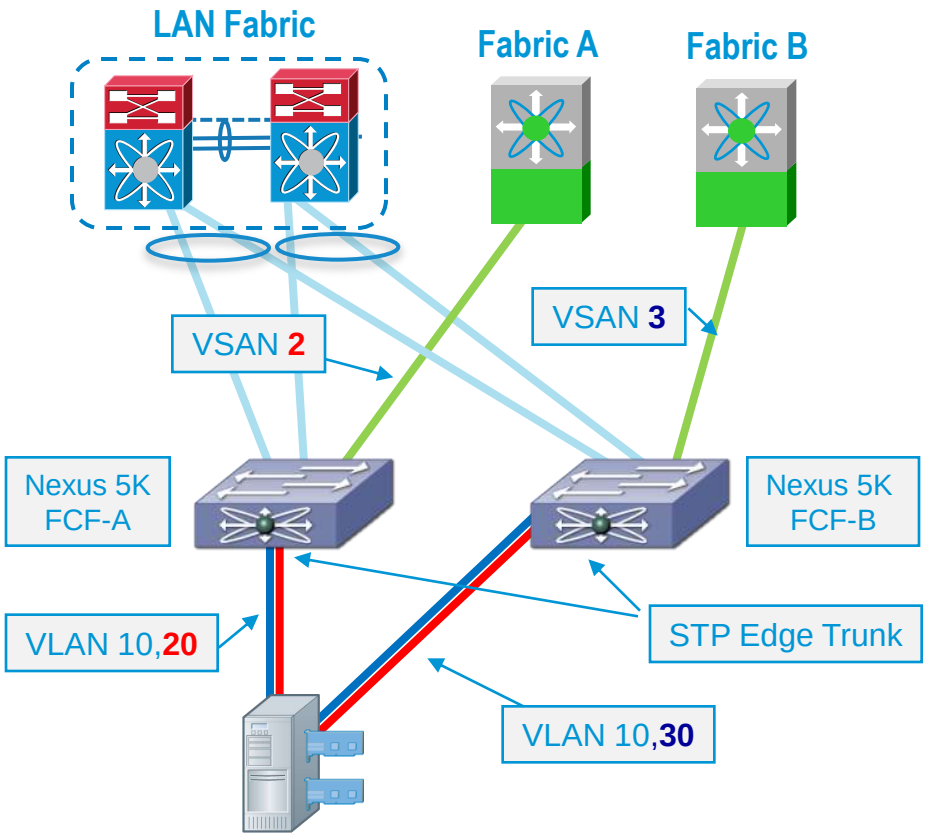


```
! VLAN 20 is dedicated for VSAN 2 FCoE traffic
(config)#  vlan 20
(config-vlan)# fcoe vsan2
```

Single Hop Design
The FCoE VLAN and STP

- **FCoE Fabric 'A' will have a different VLAN topology than FCoE Fabric 'B' which are different from the LAN Fabric**
- PVST+ allows unique topology per VLAN
- MST requires that all switches in the same Region have the same mapping of VLANs to instances
- MST does **not** require that all VLANs be defined in all switches

A separate instance must be used for FCoE VLANs

Recommended: three separate instances – native Ethernet VLANs, SAN 'A' VLANs and SAN 'B' VLANs

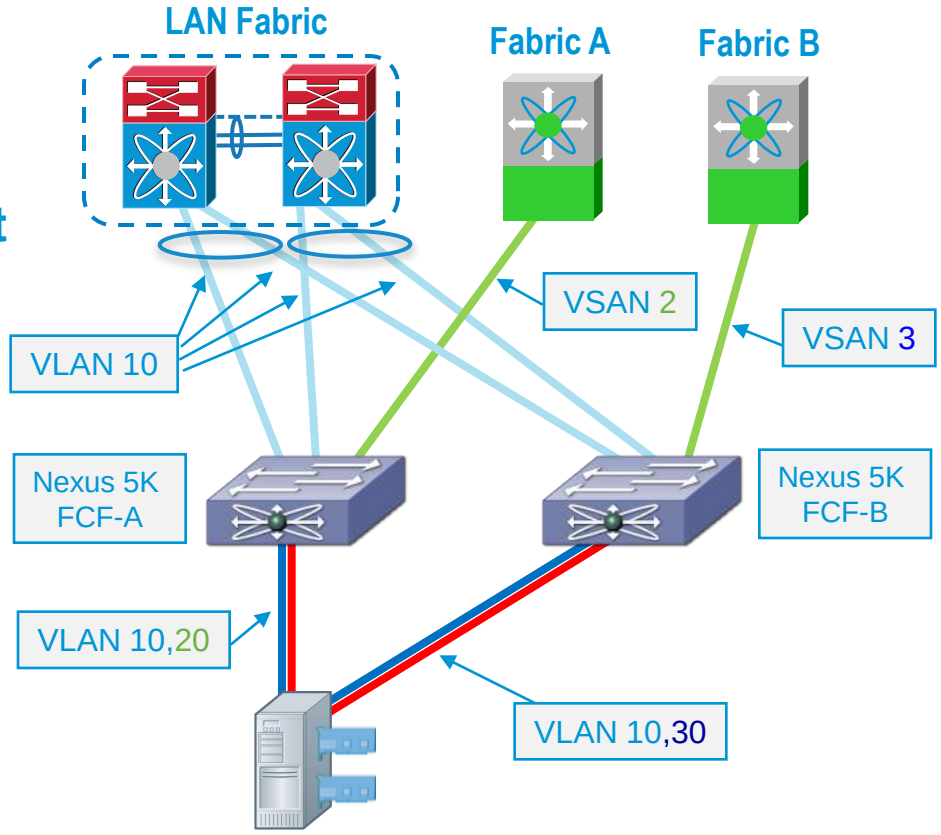


```
spanning-tree mst configuration
  name FCoE-Fabric
  revision 5
  instance 5 vlan 1-19,40-3967,4048-4093
  instance 10 vlan 20-29
  instance 15 vlan 30-39
```

Single Hop Design
Converged Links and vPC

- Optimal layer 2 LAN design often leverages Multi-Chassis Etherchannel (MCEC)
- Nexus utilizes Virtual Port Channel (vPC) to enable MCEC either between switches or to 802.3ad attached servers
- vPC provides network based load sharing and redundancy without introducing layer 2 loops in the topology
- vPC results in diverging LAN and SAN high availability topologies
  - FC maintains separate SAN 'A' and SAN 'B' topologies
  - LAN utilizes a single logical topology

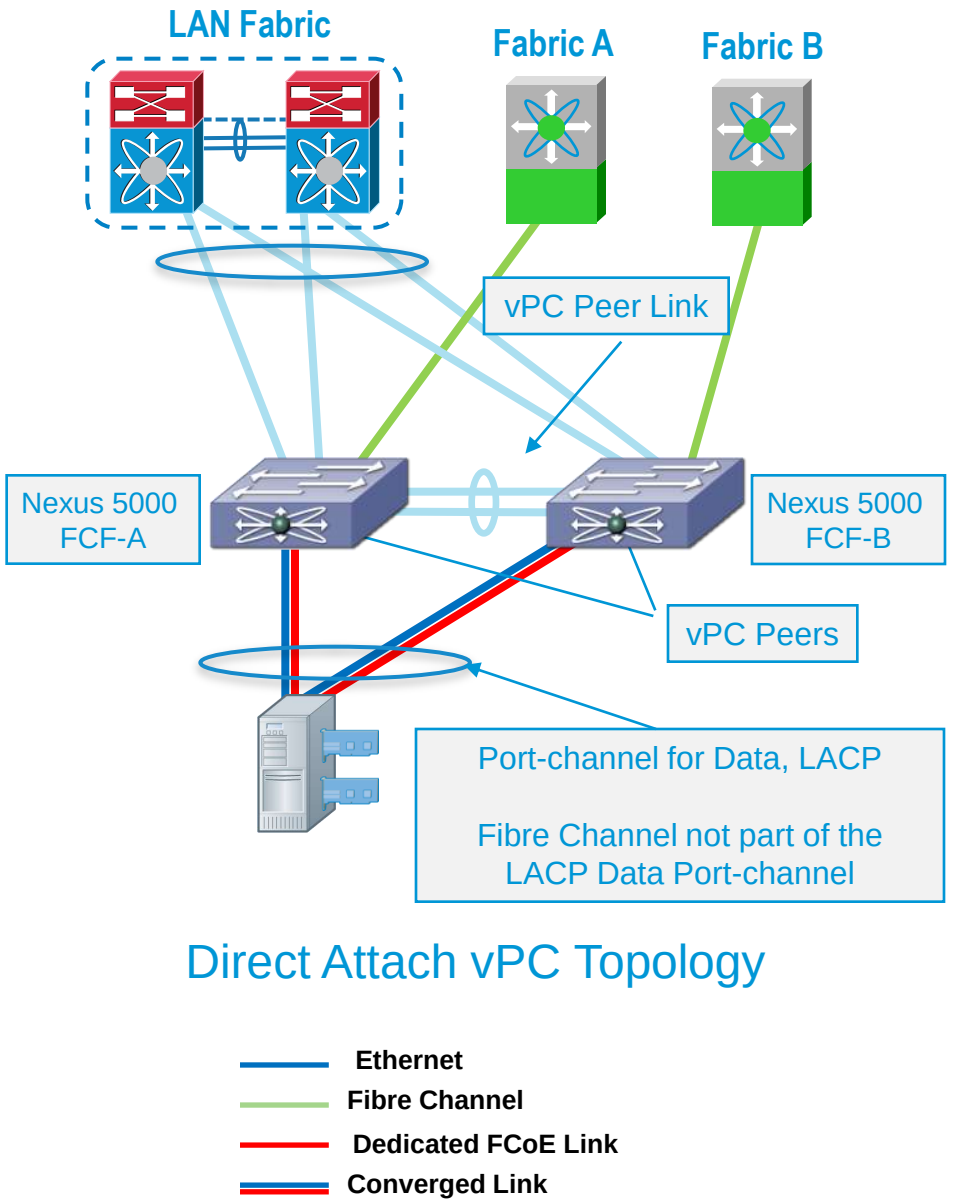

Direct Attach vPC Topology

Single Hop Design
Converged Access and vPC

- vPC enabled topologies with FCoE must follow specific design and forwarding rules
- With the NX-OS 4.1(3) releases a 'vfc' interface can only be associated with a vPC which has a single [***one (1)***] CNA port attached to each edge switch
- While the port-channel is the same on N5K-1 and N5K-2, **the FCoE VLANs are different**
- vPC configuration works with Gen-2 FIP enabled CNAs ONLY
- FCoE VLANs are *'not'* carried on the vPC peer-link
- FCoE and FIP ethertypes are *'not'* forwarded over the vPC peer link

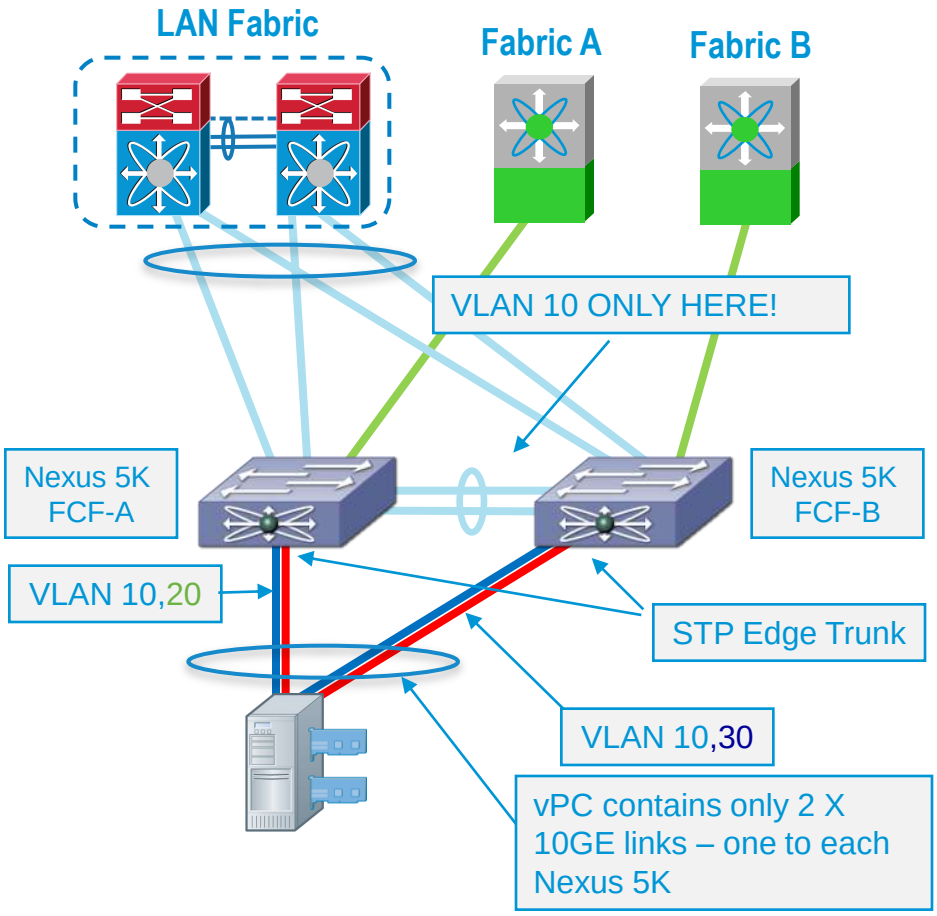

Direct Attach vPC Topology

Ethernet
Fibre Channel
Dedicated FCoE Link
Converged Link

Single Hop Design

Extending the FCoE Edge – Nexus 2232

- Server Ethernet driver connected to the FEX in NIC Teaming (AFT, TLB) or with vPC (802.3ad)
- FCoE runs over vPC member port with a single link from server to FEX
- FEX **single** homed to upstream Nexus 5000

  FEX fabric links can be connected to Nexus 5000 with individual links (static pinning) or a port channel

  Oversubscribed 4:1
- Consistent with separate LAN Access and SAN Edge Topologies

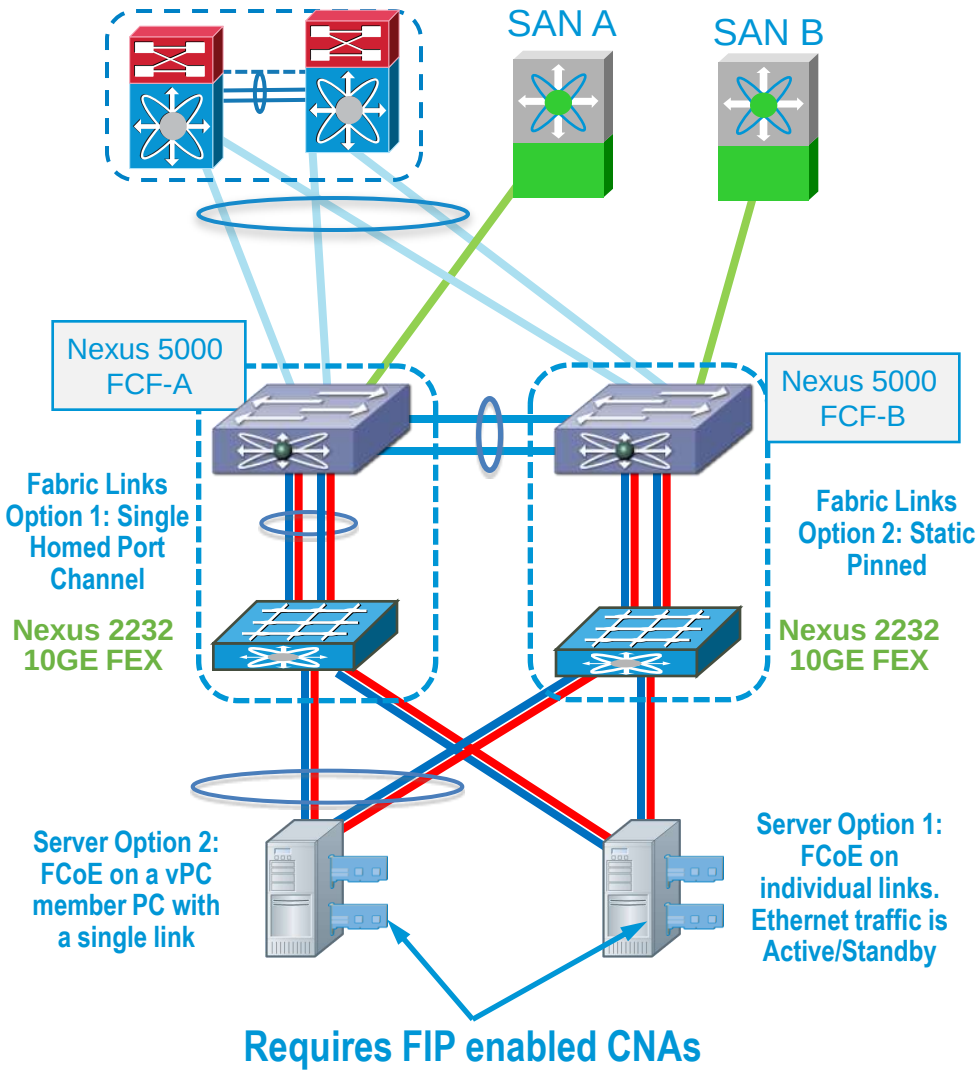

Cisco UCS Adaptor-FEX

UCS P81E Virtual Interface Card for C-Series Rackmount

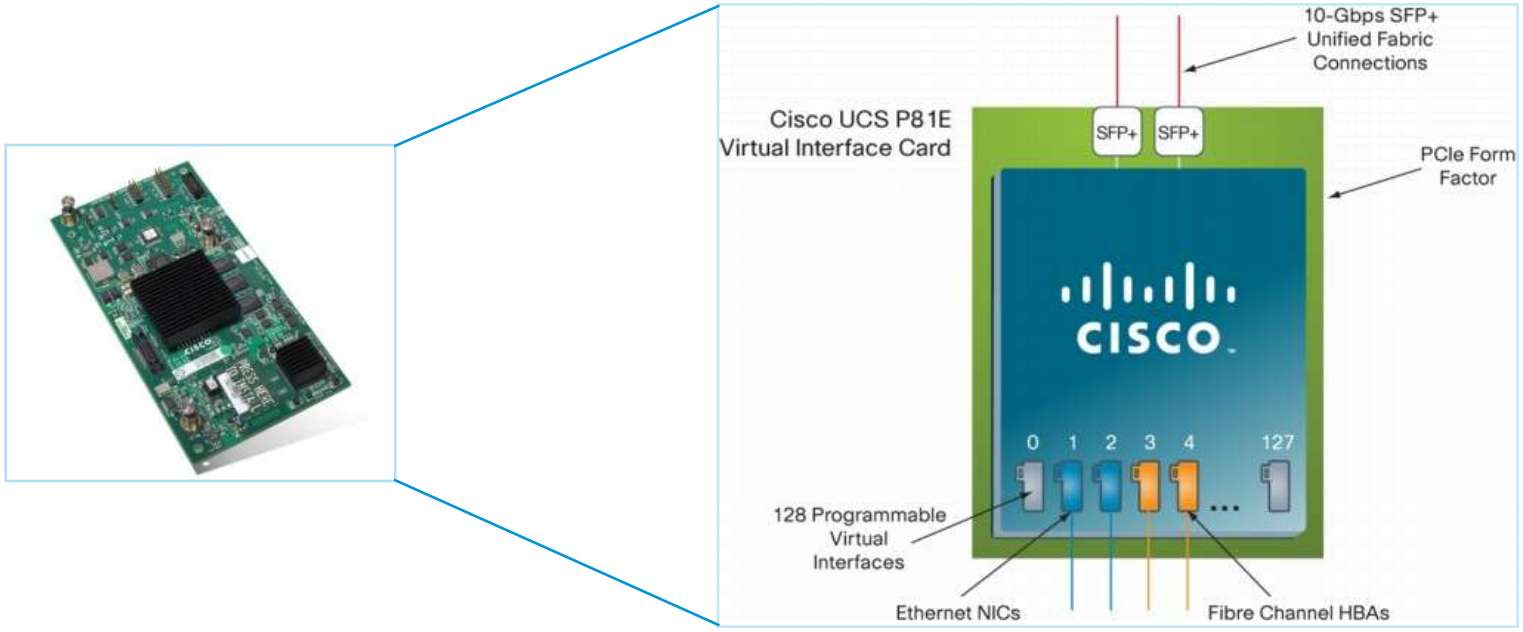

- It supports NIC partitioning to the OS and 802.1BR to the switch
  - In A-FEX mode: support for up to 16 Eth vNIC and 2 FC vHBA
  - In VM-FEX mode: support for up to 96 vNics
- Adapter Failover feature: in case of failure on the primary path, the vNIC is mapped to the standby port transparently to the OS
- Security and scalability improvements: no need of trunking all VLANs to the server interface

Broadcom Adaptor-FEX

Cisco 10-Gigabit Ethernet BCM57712 Converged Network Adapter

- Supports for NIC partitioning to the OS, 802.1BR to the switch
  - Up to 8 Virtual Interfaces total in A-FEX mode (8 vEth, 2 vHba)
  - VM-FEX not supported at FCS
- Adapter failover is not available
- Supported on 3rd Party servers

Standards Status

FEX based on IEEE 802.1BR

Working Group Ballot of Bridge Port Extension (P802.1BR), the IEEE standard for VNLink, **has reached 100% approval by the voting members of the IEEE 802.1** committee.

November 10, the IEEE 802.1 committee passed a motion to advance the draft standard to Sponsor Ballot. This is the final stage for ratification of the standard.The first Sponsor Ballot is expected to take place in late November

Ratification of the standard is currently predicted for **March 2012**

The same is true for P802.1Qbg, which is the standard the includes some of the protocols that support Bridge Port Extension as well as the VEPA device being promoted by HP

Both standards are expected to be ratified in March.

Enhanced vPC

FCoE Deployments using Adaptor-FEX NIV

NX-OS 5.1(3)N1(1))

- Enhanced vPC is supported on Nexus 5548P, Nexus 5548UP and Nexus 5596UP chassis.
- The previous generation Nexus 5010 and Nexus 5020 hardware don't have the capability to support Enhanced vPC.
- The Enhanced vPC is implemented on Nexus 5500 switch and has no special requirement from FEX point of view.
- As a result all different type of FEX can be deployed in Enhanced vPC topology.

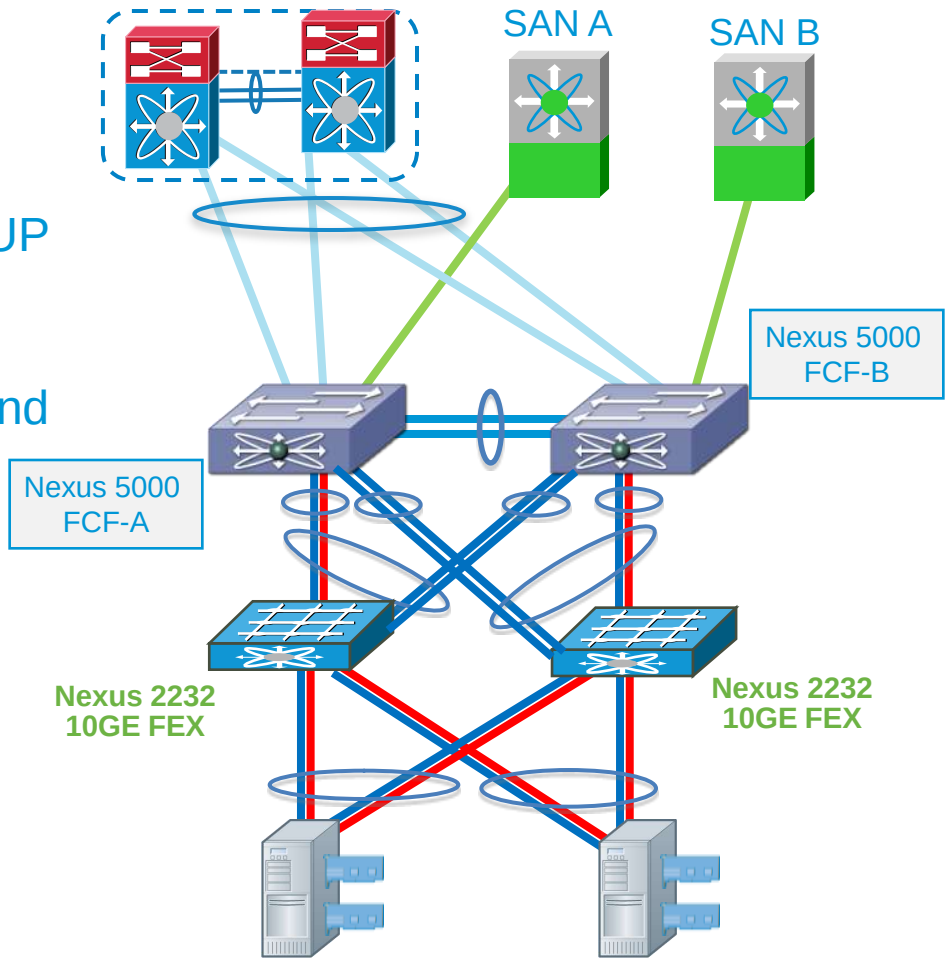

Enhanced vPC

FCoE Deployments using Adaptor-FEX NIV

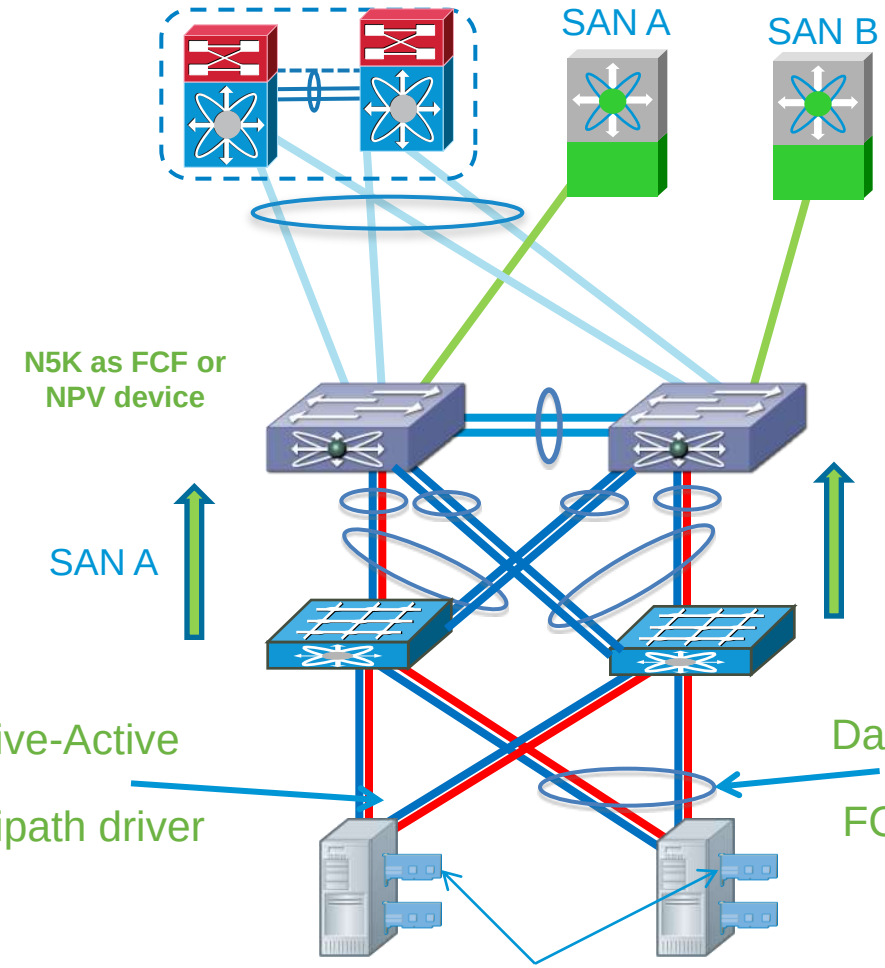

Cisco Nexus B22HP Fabric Extender

Overview

- 16 x 1/10 GE server interfaces
- 8 x 10 GE network interfaces
- 8 QoS queues (6 configurable)
- Host vPC (virtual Port-Channel)
- DCB and FCoE in 10G mode
- Upstream Nexus 5000 supports FEX mix & match

Requires upstream Nexus 5000/55xx

Single Hop Design

Extending the FCoE Edge with B22 HP-FEX (N2K-B22HP-P)

Active-Standby attached Host (HP C7000)

Static Pinning (Fabric Links)

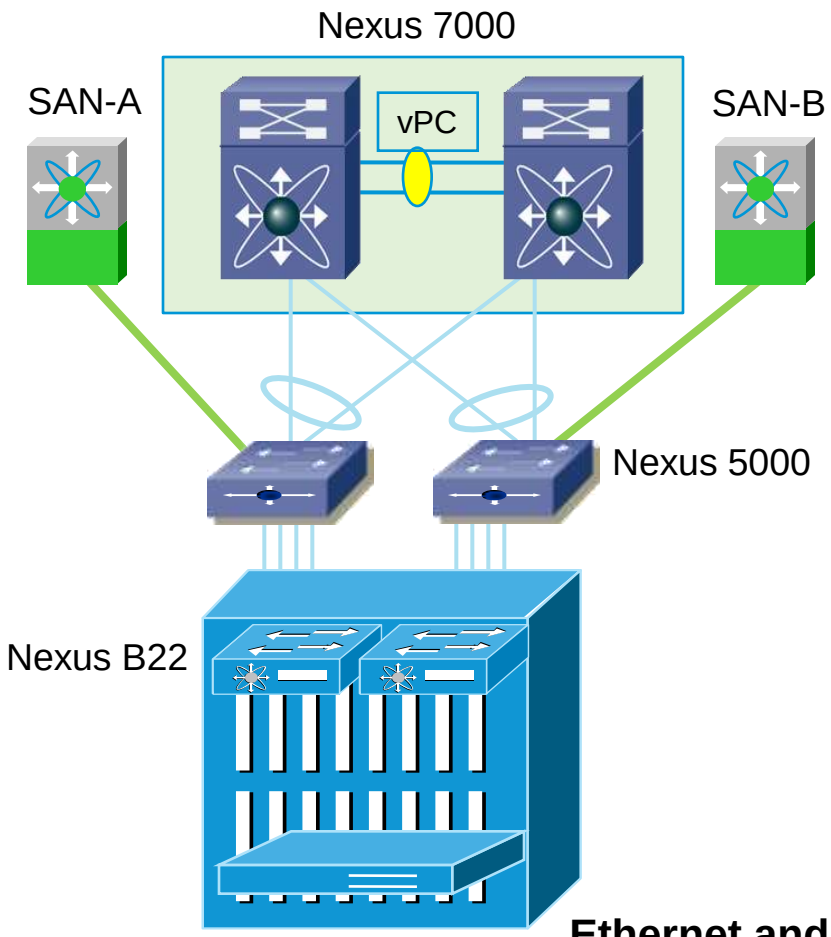

Port Channel (Fabric Links)

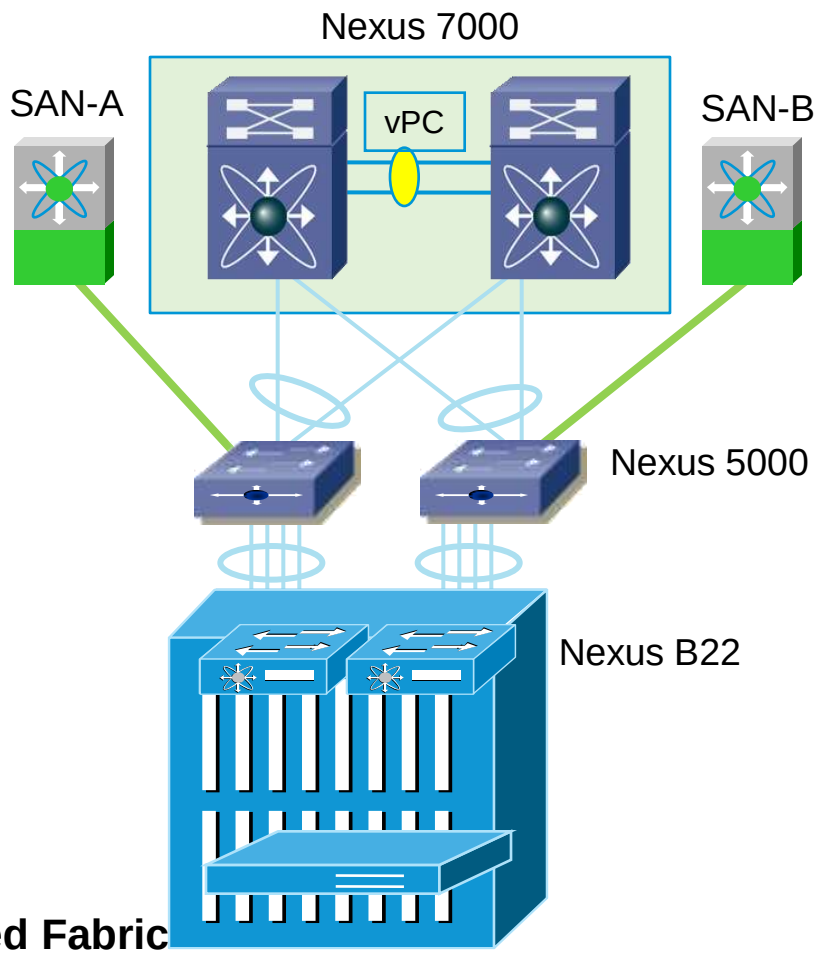

Ethernet and Unified Fabric

Single Hop Design

Extending the FCoE Edge with B22 HP-FEX (N2K-B22HP-P)

Active-Active attached Host (C7000)

Static Pinning (Fabric Links)

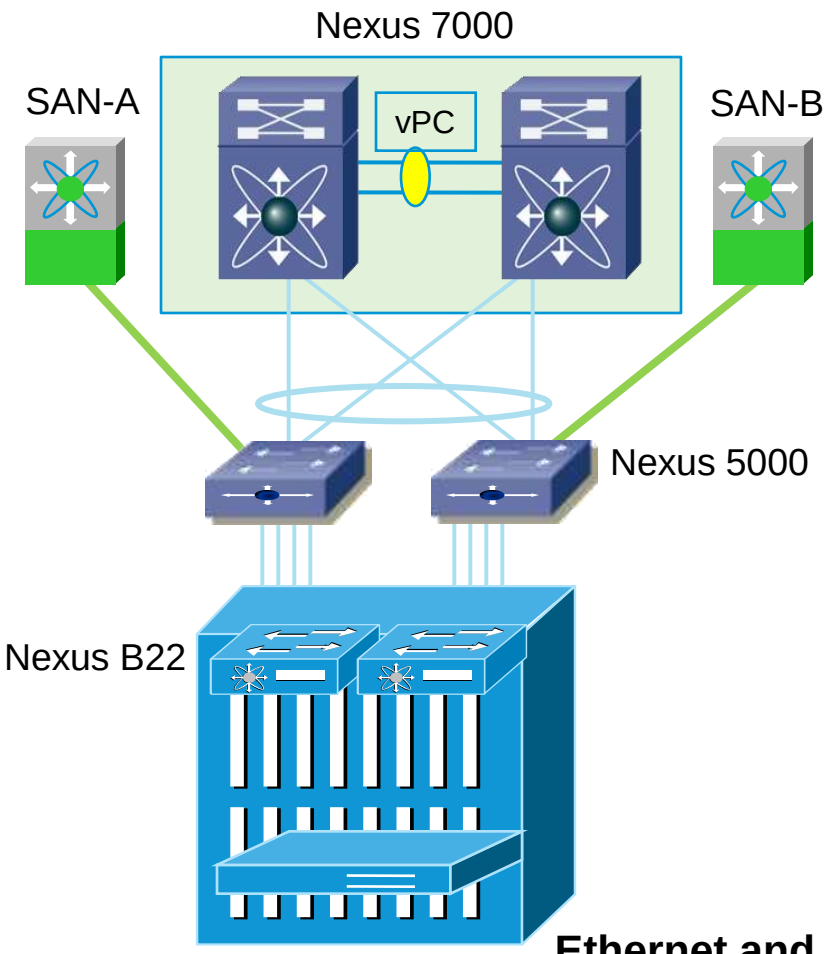

Port Channel (Fabric Links)

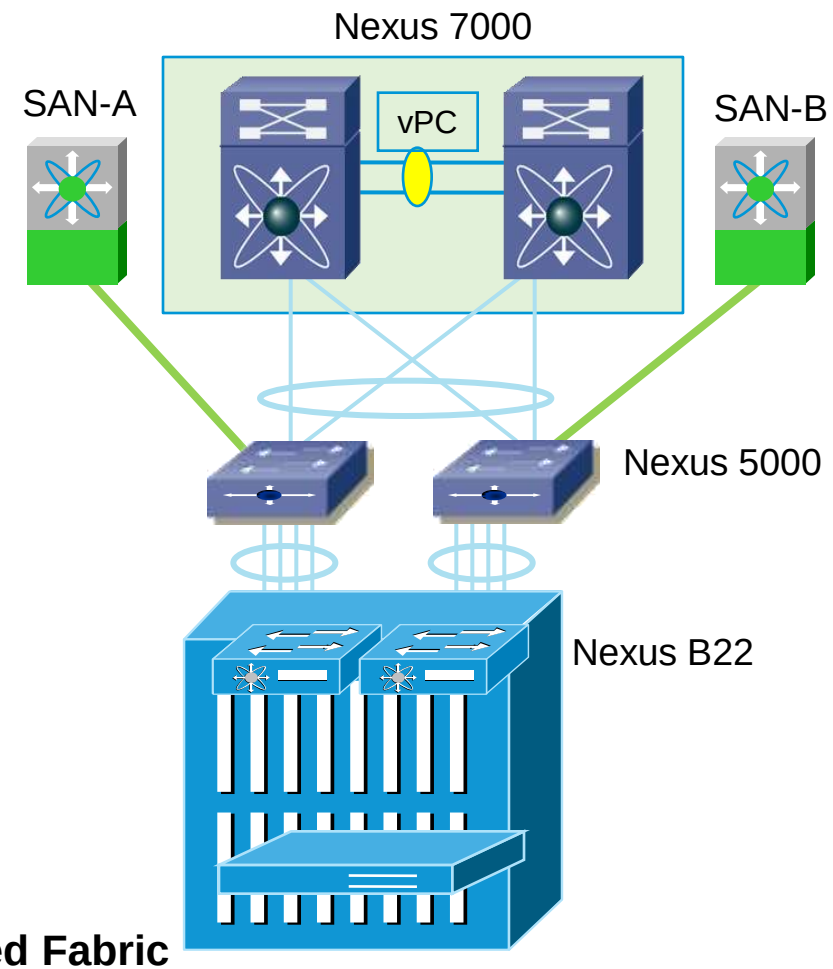

Ethernet and Unified Fabric

Agenda

Virtual E_Ports

What happens when FCF's are connected via VE_Ports

- 10Gig Ethernet ISL or LACP port-channel must first be established between FCF Switches expanding the L2 ethernet network
- LLDP frames with DCBx TLV's, sourcing the MAC addresses of each switch are exchanged across the ethernet Link to determine abilities.
- FIP Control exchange is done between switches
- FSPF routing established
- Fibre Channel Protocol is exchanged between the FCF's and a Fibre Channel merge of Zones is accomplished building out the FC SAN.

You now have established a VE_Port between two DCB switches

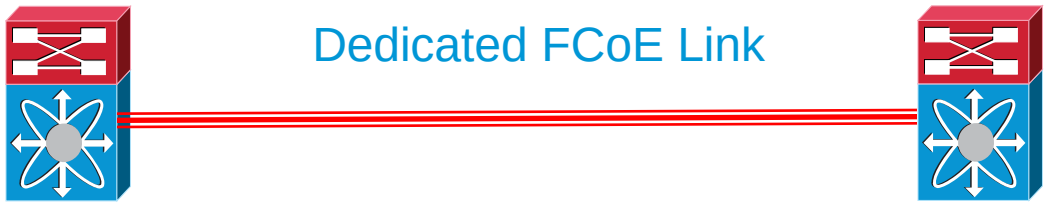

VE_Port FIP exchange

A FIP ELP (Exchange Link Parameter)is sent on each VLAN by both the N5K, N7K or MDS. A FIP ACC is sent by the switch for each VLAN.

Protocol	Summary	Source [MAC - FC]	Destination [MAC - FC]	VLAN/VSAN
FIP	Virtual Link Instantiation Request; FC4UCtl; FC FCSS; ELP;	Cisco Systems:20:A9:40 - Fabric Controller	00:26:98:0A:DF:02 - Fabric Controller	200
FIP	Virtual Link Instantiation Request; FC4UCtl; FC FCSS; ELP;	Cisco Systems:20:A9:40 - Fabric Controller	00:26:98:0A:DF:02 - Fabric Controller	100
FIP	Virtual Link Instantiation Request; FC4UCtl; FC FCSS; ELP;	Cisco Systems:20:A9:40 - Fabric Controller	00:26:98:0A:DF:02 - Fabric Controller	10
FIP	Virtual Link Instantiation Request; FC4UCtl; FC FCSS; ELP;	00:26:98:0A:DF:02 - Fabric Controller	Cisco Systems:20:A9:40 - Fabric Controller	200
FIP	Virtual Link Instantiation Request; FC4UCtl; FC FCSS; ELP;	00:26:98:0A:DF:02 - Fabric Controller	Cisco Systems:20:A9:40 - Fabric Controller	100
FIP	Virtual Link Instantiation Request; FC4UCtl; FC FCSS; ELP;	00:26:98:0A:DF:02 - Fabric Controller	Cisco Systems:20:A9:40 - Fabric Controller	10
FIP	Virtual Link Instantiation Reply; FC4SCtl; FC FCSS; Accept ELP;	Cisco Systems:20:A9:40 - Fabric Controller	00:26:98:0A:DF:02 - Fabric Controller	200
FIP	Virtual Link Instantiation Reply; FC4SCtl; FC FCSS; Accept ELP;	Cisco Systems:20:A9:40 - Fabric Controller	00:26:98:0A:DF:02 - Fabric Controller	100
FIP	Virtual Link Instantiation Reply; FC4SCtl; FC FCSS; Accept ELP;	Cisco Systems:20:A9:40 - Fabric Controller	00:26:98:0A:DF:02 - Fabric Controller	10

Discovery Solicitations & Advertisements from the FCF are sent both ways across the VE_Port, one for each FCoE mapped VLAN that is trunked on the interface.

Port	Count - Type	Count - Type	Protoc	Summary
GE Port(1,1,2)		1 - Ether Frar	LACP	LACP ver=1, A-Key=414, A-Port=271, P-Key=414, P-P
GE Port(1,1,2)		1 - Ether Frar	FIP	Discovery Solicitation;
GE Port(1,1,2)		1 - Ether Frar	FIP	Discovery Solicitation;
GE Port(1,1,2)		1 - Ether Frar	FIP	Discovery Solicitation;
GE Port(1,1,1)	1 - Ether Frar		FIP	Discovery Solicitation;
GE Port(1,1,1)	1 - Ether Frar		FIP	Discovery Solicitation;
GE Port(1,1,1)	1 - Ether Frar		FIP	Discovery Solicitation;
GE Port(1,1,2)		1 - Ether Frar	FIP	Discovery Advertisement;
GE Port(1,1,2)		1 - Ether Frar	FIP	Discovery Advertisement;
GE Port(1,1,2)		1 - Ether Frar	FIP	Discovery Advertisement;
GE Port(1,1,1)	1 - Ether Frar		SNAP	PVSTP+; Cisco;
GE Port(1,1,1)	1 - Ether Frar		SNAP	PVSTP+; Cisco;
GE Port(1,1,1)	1 - Ether Frar		SNAP	PVSTP+; Cisco;
GE Port(1,1,1)	1 - Ether Frar		FIP	Discovery Advertisement;
GE Port(1,1,1)	1 - Ether Frar		FIP	Discovery Advertisement;
GE Port(1,1,1)	1 - Ether Frar		FIP	Discovery Advertisement;
GE Port(1,1,2)		1 - Ether Frar	FIP	Discovery Advertisement;
GE Port(1,1,2)		1 - Ether Frar	FIP	Discovery Advertisement;
GE Port(1,1,2)		1 - Ether Frar	FIP	Discovery Advertisement;

FCoE VE - Fibre Channel E_Port handshake

Protocol	Summary
FC	FC4UCtl; FC FCSS; EFP;
FC	FC4UCtl; FC FCSS; BF;
FC	FC4SCtl; FC FCSS; Accept;
FC	FC4UCtl; FC FCSS; MRRA;
FC	FC4UCtl; FC FCSS; EFP;
FC	FC4UCtl; FC FCSS; EFP;
FC	FC4SCtl; FC FCSS; Accept;
FC	FC4SCtl; FC FCSS; Accept;
FC	ABTS; Basic Link Service; Abort Exchange;
FC	FC4UCtl; FC FCSS; MRRA;
FC	FC4UCtl; FC FCSS; DIA;
FC	FC4SCtl; FC FCSS; Accept;
FC	FC4UCtl; FC FCSS; RDI;
FC	FC4SCtl; FC FCSS; Accept;
FC	FC4UCtl; FC FCSS; EFP;
FC	FC4SCtl; FC FCSS; Accept;
FC	FC4UCtl; FC FCSS; MRRA;
FC	FC4SCtl; FC FCSS; Accept;
FC	FC4UCtl; FC FCSS; MR;
FC	FC4SCtl; FC FCSS; Accept;
FC	FC4SCtl; FC FCSS; Accept;
FC	FC4UCtl; FC FCSS; MR;
FC	FC4SCtl; FC FCSS; Accept;
FC	FC4UCtl; FC FCSS; HLO;
FC	FC4UCtl; FC FCSS; HLO;
FC	FC4UCtl; FC FCSS; HLO;
FC	FC4UCtl; FC FCSS; HLO;
FC	FC4UCtl; FC FCSS; LSU;
FC	FC4UCtl; FC FCSS; LSU;
FC	FC4UCtl; FC FCSS; LSA;
FC	FC4UCtl; FC FCSS; LSU;

Exchange Fabric Parameters

Build Fabric

Enhanced Zoning **M**erge **R**equest **R**esource **A**llocation

Domain ID Assign by Existing Principal Switch

Request Domain ID from New Switch

Zone Merge Request

FSPF exchanges

Differences in Trunking VSANs with FCoE VE_Ports

- In FC on the MDS, trunking is used to carry multiple VSANs over the same physical FC link. With FCoE, a physical link is replaced by a virtual link, a pair of MAC addresses.

- FCoE uses assigned MAC addresses that are unique only in the context of a single FC fabric. Carrying multiple fabrics over a single VLAN would then mean having a strong possibility for duplicated MAC addresses.

- In FCoE there cannot be more than one VSAN mapped over a VLAN. The net result is that trunking is done at the Ethernet level, not at the FC level.

- FC trunking is not needed and the Fibre Channel Exchange Switch Capabilities(ESC) & Exchange Port Parameters (EPP) processing is not required to be performed as on the MDS

Agenda

Multi-Hop Designs

FCoE Extension Options

Short Distance Options

- **Requirement:** Maintain loss-less behavior across the point-to-point link
- Supported distance is governed by the ingress buffer size available on the switch

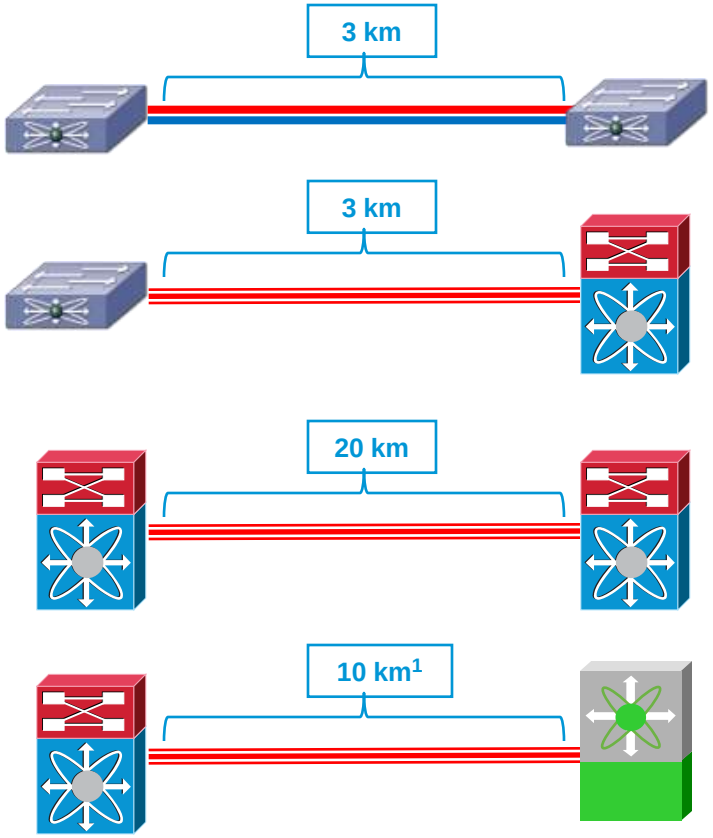

FCoE Convereged

FCoE Dedicated

FC

Longer Distance Options

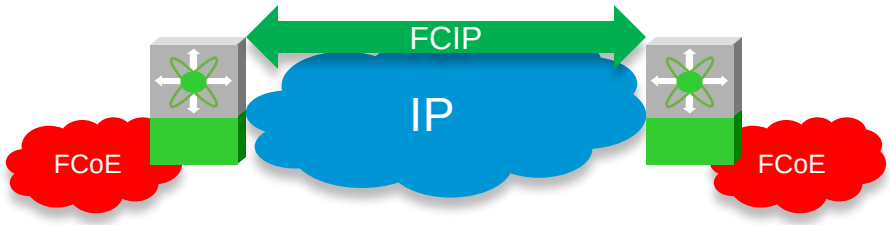

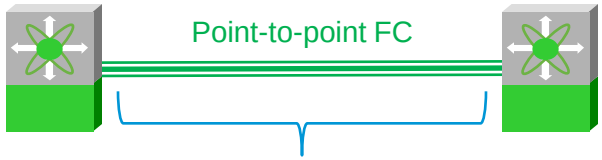

Speed (Gbps)	Max Distance (KM)
1	8000
2	4000
4	2000
8	1000
10	680

1. Limited by supported Optics

Multi - Hop Design

FCoE Pass-through options

- Multi-hop FCoE networks allow for FCoE traffic to extend past the access layer (first hop)
- In Multi-hop FCoE the role of a transit Ethernet bridge needs to be evaluated
  - Avoid Domain ID exhaustion
  - Ease management
- **FIP Snooping** is a minimum requirement *suggested* in FC-BB-5
- **Fibre Channel over Ethernet NPV** FCoE-NPV is a new capability in N5K release 5.2.1 intended to solve a number of design and management challenges as NPV does in FC SAN's

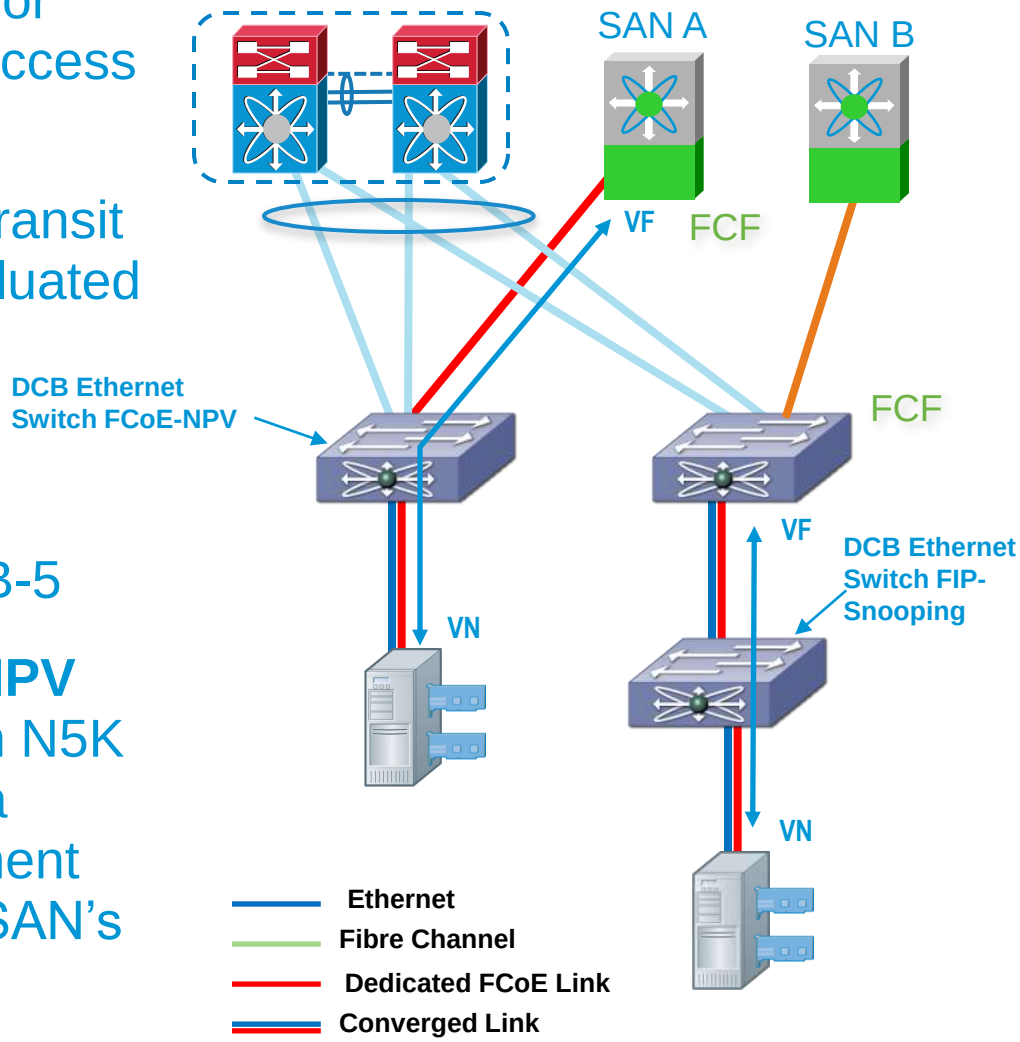

Multi - Hop Design

Extending FCoE with FIP Snooping

- Nexus 4000 is a Unified Fabric capable Blade Switch
  - DCB enabled
  - FIP Snooping Bridge
- Dual Topology requirements for FCoE multi-hop
- Servers IP connection to the Nexus 4000 is Active/Standby

  MCEC is **not** currently supported from blade server to Nexus 4000
- Options 1: **Dedicated Links** from Nexus 4000 to Nexus 5000
- Options 2: **Converged Link** Port Channel from Nexus 4000 to Nexus 5000
- Not support connected to N7K or MDS

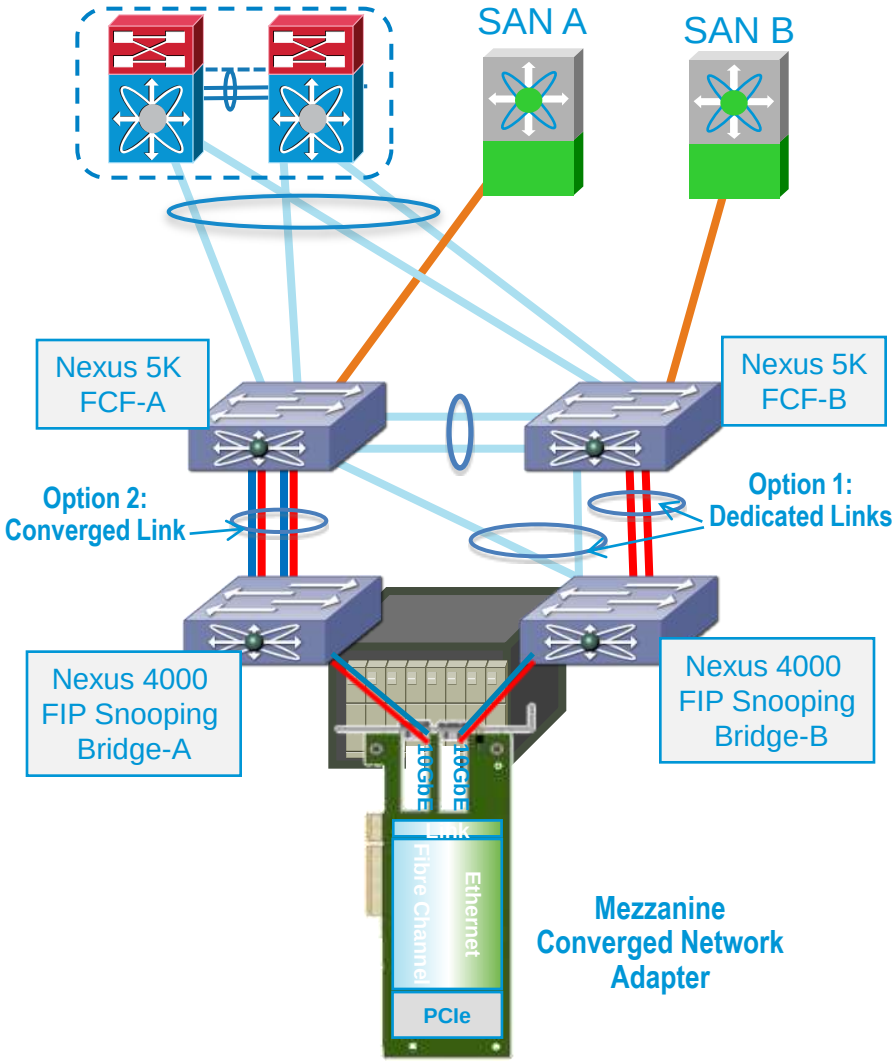

Multi - Hop Design

VE_Ports

- Standards-based multihop
- Supported between Nexus 5000, Nexus 7000, and MDS 9500
- Fully interoperable with FC SANs
- Maintains A/B fabric separation
- High performance Ethernet switches for storage area networks
- Leverage a single network infrastructure for LAN and SAN
- Can be designed with Port-channels

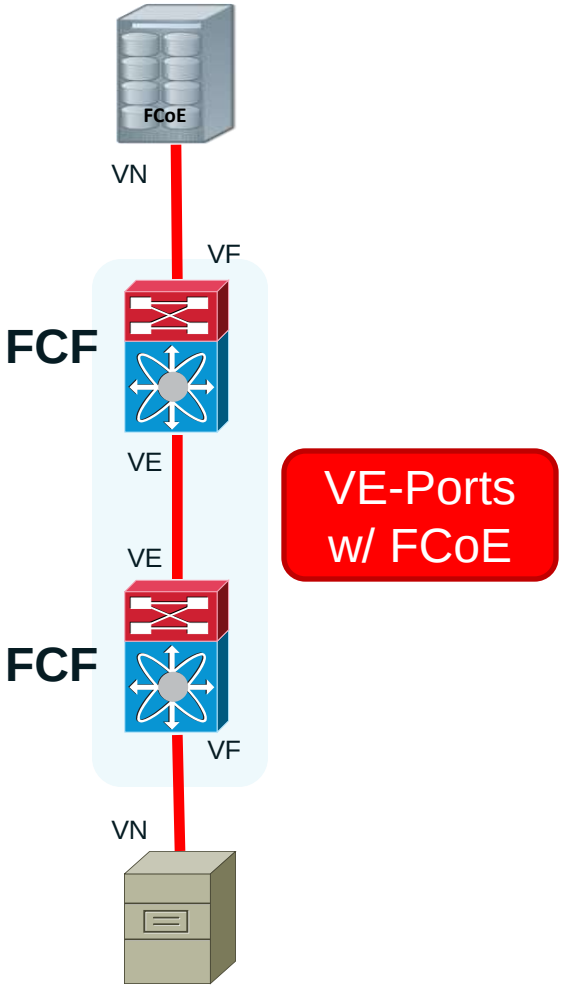

Storage VDC on the Nexus 7000
Supported VDC models

- **Separate VDC running ONLY storage related protocols**
- **Storage VDC: a *virtual* MDS FC switch**
- **Running only FC related processes**
- **Only one such VDC can be created**
- **Provides control plane separation**

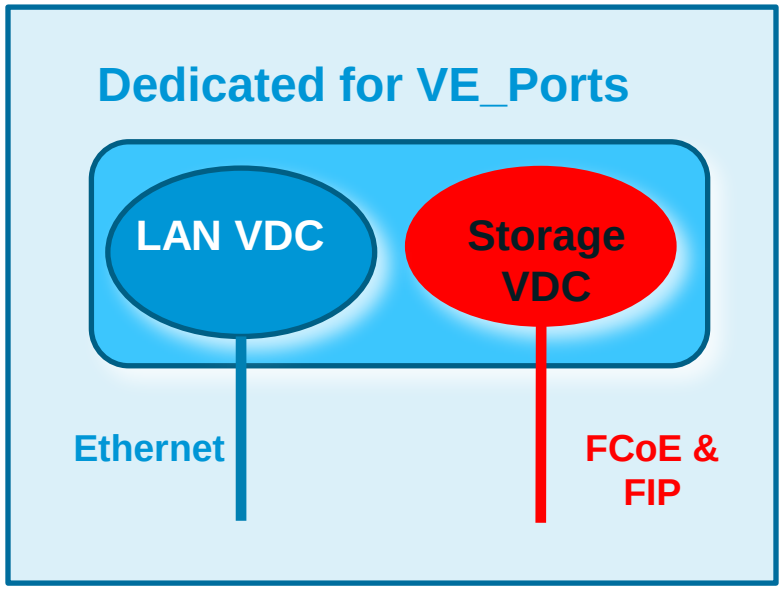

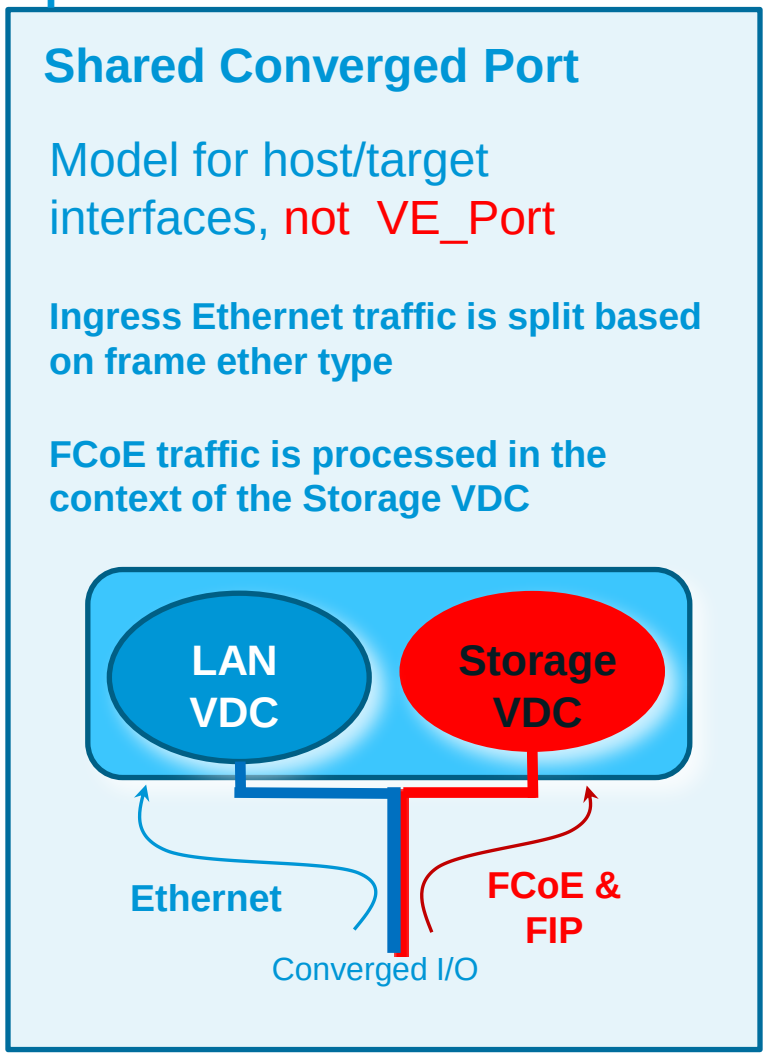

Creating the Storage VDC Port and VLAN allocations

- **Create VDC of type storage and allocate non-shared interfaces:**

N7K-50(config)# vdc fcoe id 2 type storage
N7K-50(config-vdc)# allocate interface Ethernet4/1-16, Ethernet4/19-22

- **Allocate FCoE vlan range from the Owner VDC to the Storage VDC. This is a necessary step for sharing interfaces to avoid vlan overlap between the Owner VDC and the Storage VDC**

N7K-50(config) vdc fcoe id 2
N7K-50(config-vdc)# allocate fcoe-vlan-range 10-100 from vdcs n7k-50

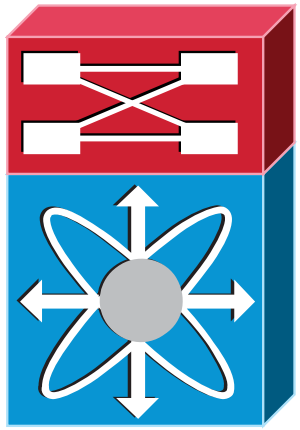

- **Allocated the shared interfaces:**

N7K-50(config-vdc)# allocate shared interface Ethernet4/17-18

- **Install the license for the FCoE Module.**

n7k-50(config)# license fcoe module 4

Shared Converged Port

- **Preparing for converged shared interfaces.**

  Interfaces must be from the same Orion pair. This means two consecutive 10G ports must be configured at a time.
  For example, ports 1-2, or ports 3-4, and so on.

  Interfaces have to be in trunk mode.

  Interfaces have to be no shut in the Parent VDC.

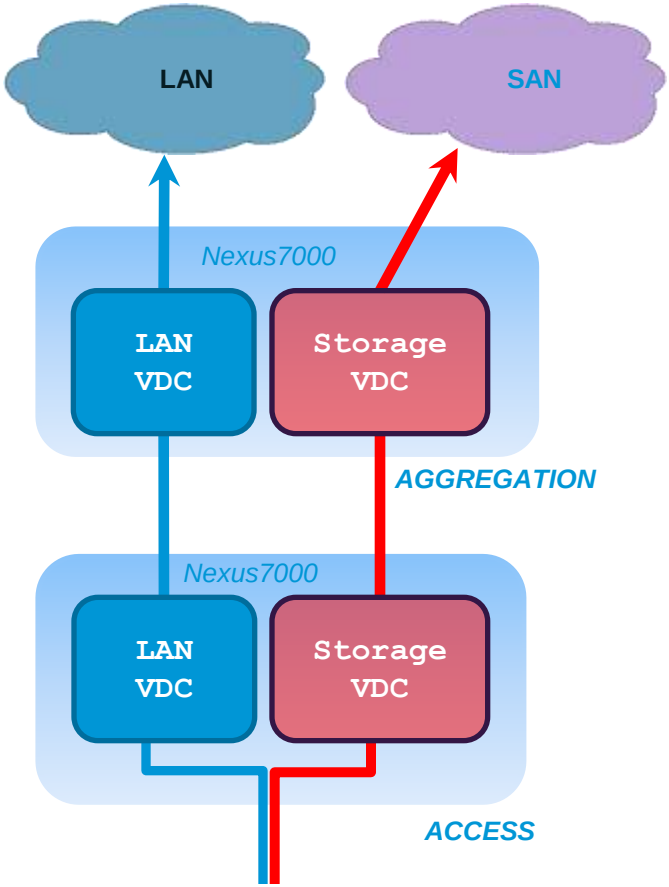


```
N7k-50(config)# interface ethernet 4/17-18
N7k-50(config-if)# switchport mode trunk
N7k-50(config-if)# switchport trunk allowed vlan 1
N7k-50(config-if)# no shut
```

MDS Configuration Details

There is no need to enable FCoE explicitly on the MDS switch. The following features will be enabled once an FCoE capable linecard is detected.

Install feature-set fcoe
feature-set fcoe

feature lldp
feature vlan-vsan-mapping

Create VSAN and VLAN, Map VLAN to VSAN for FCoE

```
pod3-9513-71(config)# vsan database
pod3-9513-71(config-vsan-db)# vsan 50
pod3-9513-71(config-vsan-db)# vlan 50
pod3-9513-71(config-vlan)# fcoe vsan 50
```

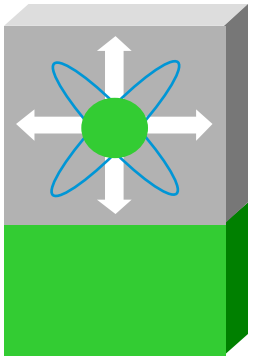

Build the LACP Port Channel on the MDS

Create VE port and assign to the LACP Port-channel

```
pod3-9513-71(config-if-range)# interface vfc-port-channel 501
pod3-9513-71(config-if)# switchport mode e
pod3-9513-71(config-if)# switchport trunk allowed vsan 50
pod3-9513-71(config-if)# no shut
```

Storage on MDS

Extending FCoE to MDS SAN

- Migration to 10G FCoE in place of 4/8G FC links (Ethernet price per bit economics)
- FSPF forwarding for FCoE traffic is end-to-end
- Hosts will log into the FCF which they are attached
- Storage devices will log into the MDS at the core/storage edge
- Maintains HA requirements from both LAN and SAN perspective
- Same SAN Management tools and best practice

Multi - Hop Design

Extending FCoE to MDS SAN from Aggregation

- **Converged Network into the existing SAN Core**

  Leverage FCoE wires between Fibre Channel SAN to Ethernet DCB switches in Aggregation layer using Dedicated ports

  Maintain the A – B SAN Topology with Storage VDC and Dedicated wires

  Using N7K director Class Switches at Access layer

  Dedicated FCoE Ports between access and Aggregation, vPC's for Data

  Zoning controlled by Core A-B SAN

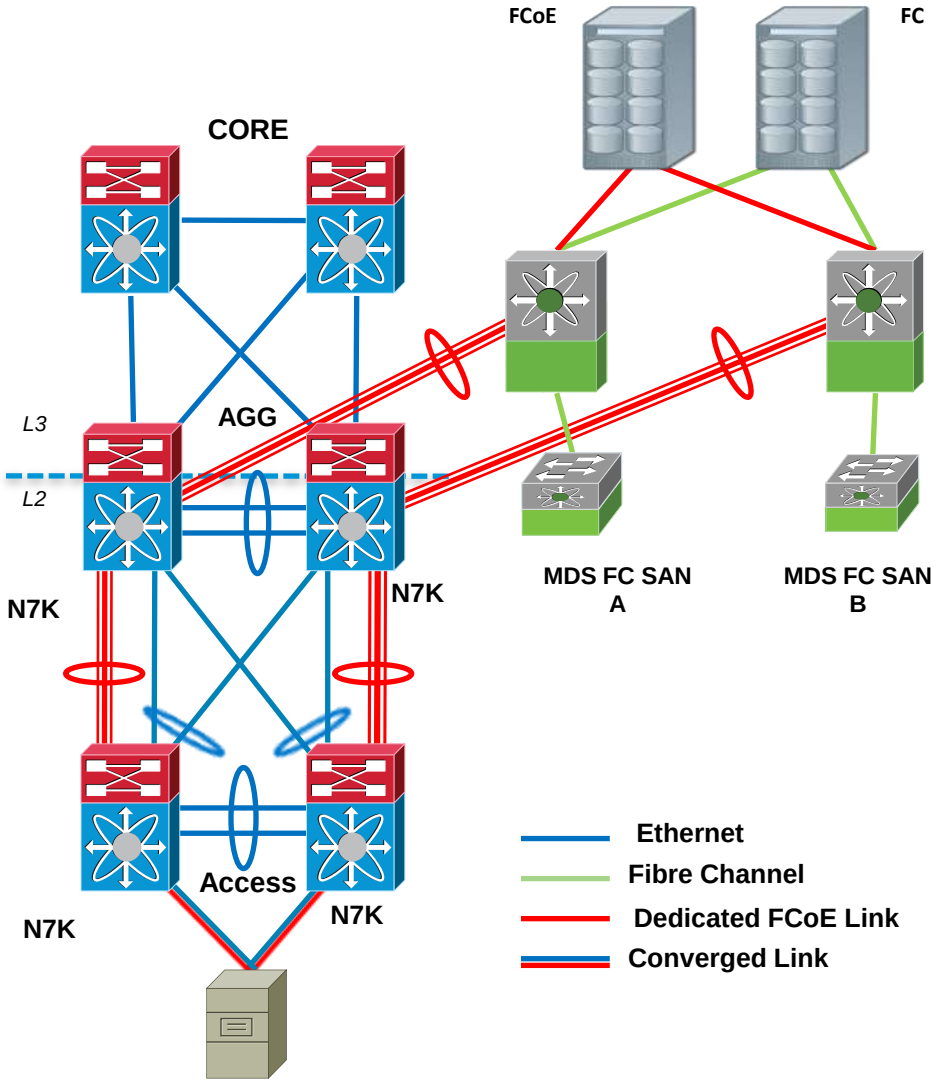

Storage on MDS

Extending FCoE to MDS SAN from Access

- **Converged Network capabilities to the existing SAN Core**

  Leverage FCoE wires between Fibre Channel SAN to Ethernet DCB switches (VE_Ports)

  N5K access switches can be in Fibre Channel switch node and assigned Domain ID , or

  N5K access switches can run in FCoE-NPV mode, no FC services running local.

  Zoning controlled by Core A-B SAN

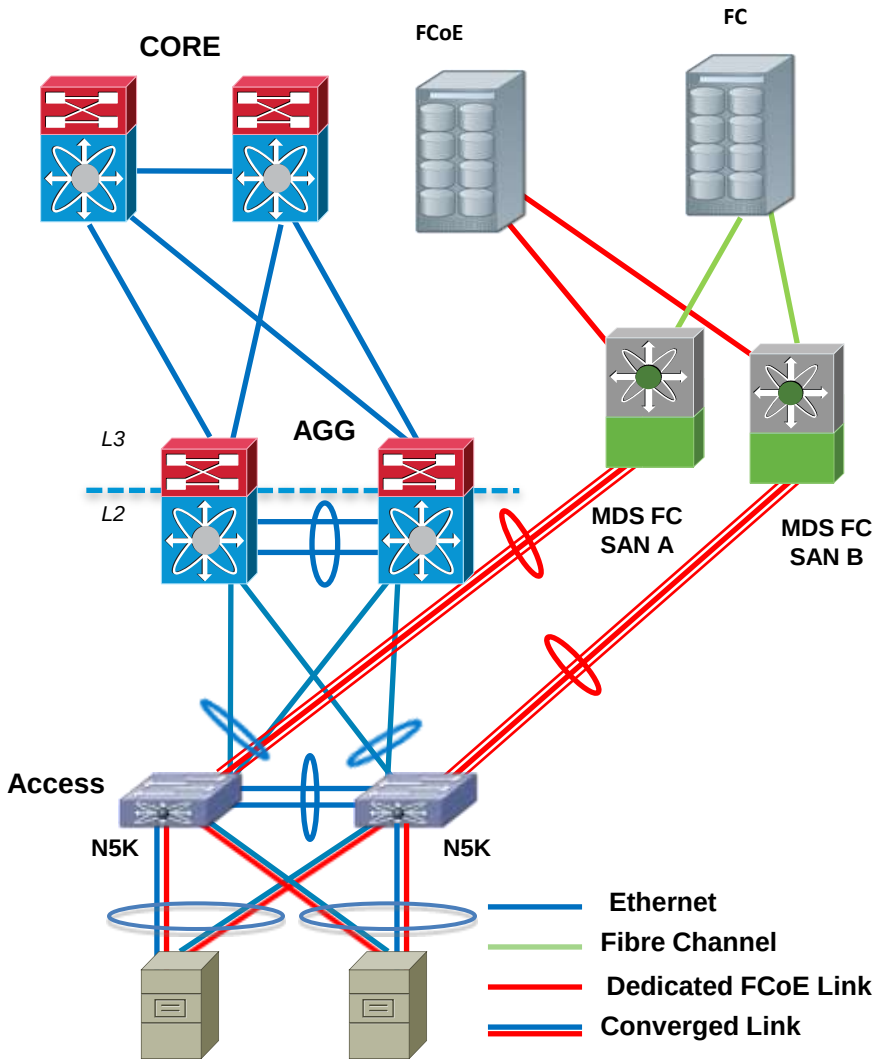

Migration of Storage to Aggregation

- **Different requirements for LAN and SAN network designs**
- **Factors that will influence this use case**
  - Port density
  - Operational roles and change management
  - Storage device types
- **Potentially viable for smaller environments**
- **Larger environments will need dedicated FCoE 'SAN' devices providing target ports**
  - Use connections to a SAN
  - Use a "storage" edge of other FCoE/DCB capable devices

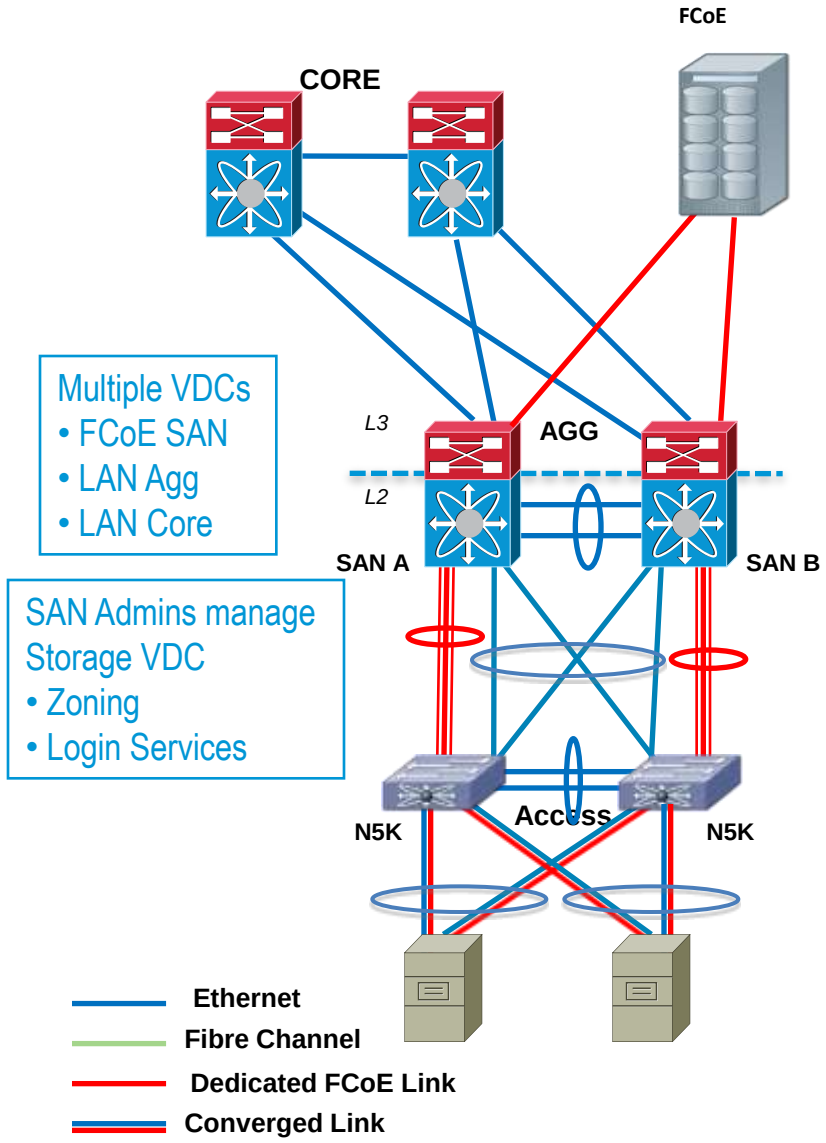

Data Center Design with E-SAN

Ethernet LAN and Ethernet SAN

- Same topologies as existing networks, but using Nexus Unified Fabric Ethernet switches for SANs
- Physical and Logical separation of LAN and SAN traffic
- Additional Physical and Logical separation of SAN fabrics
- Ethernet SAN Fabric carries FC/FCoE & IP based storage (iSCSI, NAS, …)
- Common components: Ethernet Capacity and Cost
- Standardize on OS, I/O and Platform
- Storage administrators in **Large Data Centers** almost always prefer this model (distinct storage management plane)

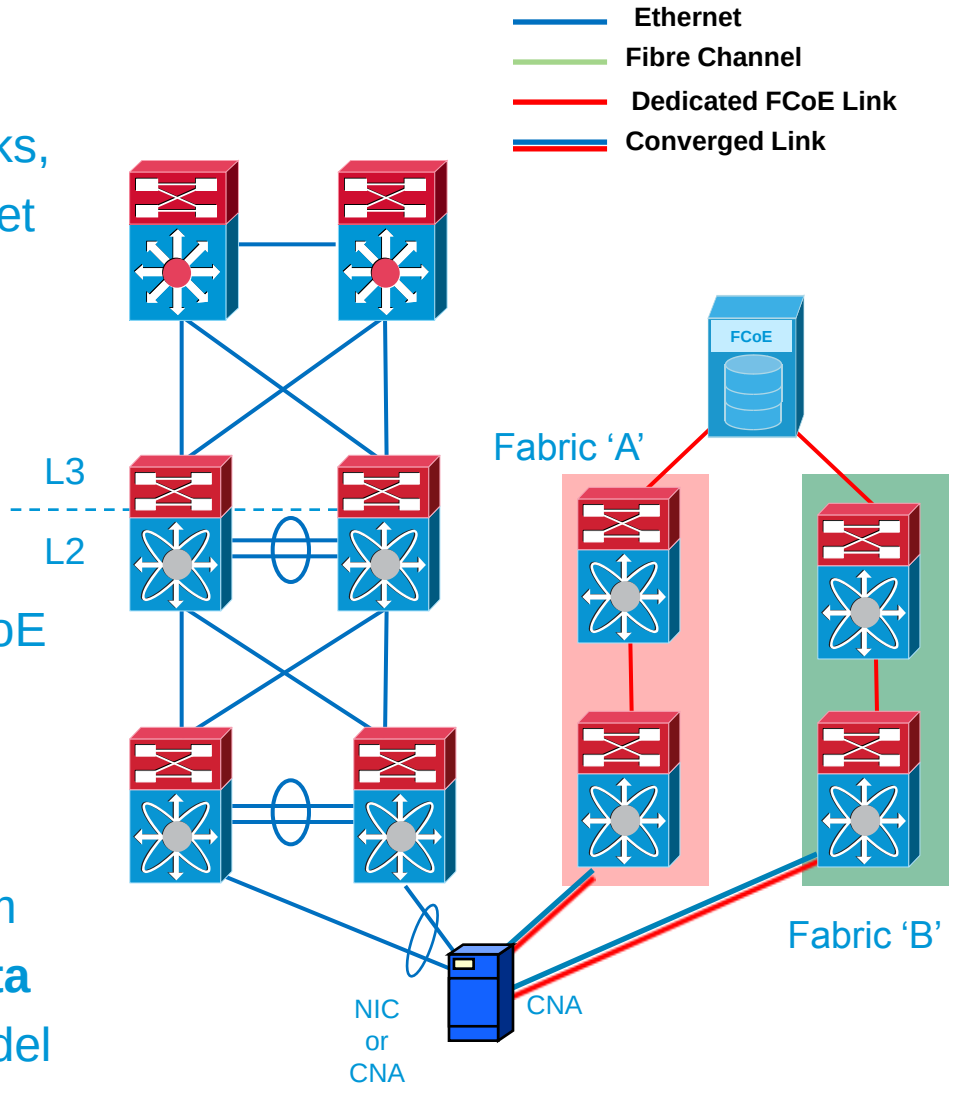

FCoE Deployment Considerations
Shared Aggregation/Core Devices

- Does passing FCoE traffic through a larger aggregation point make sense?
- Multiple links required to support the HA models
- 1:1 ratio between access to aggregation and aggregation to SAN core is required
- Need to plan for appropriate capacity in any core VE_Port link
- When is a direct Edge to Core links for FCoE are more cost effective than adding another hop?
- Smaller Edge device more likely to be able to use under-provisioned uplinks

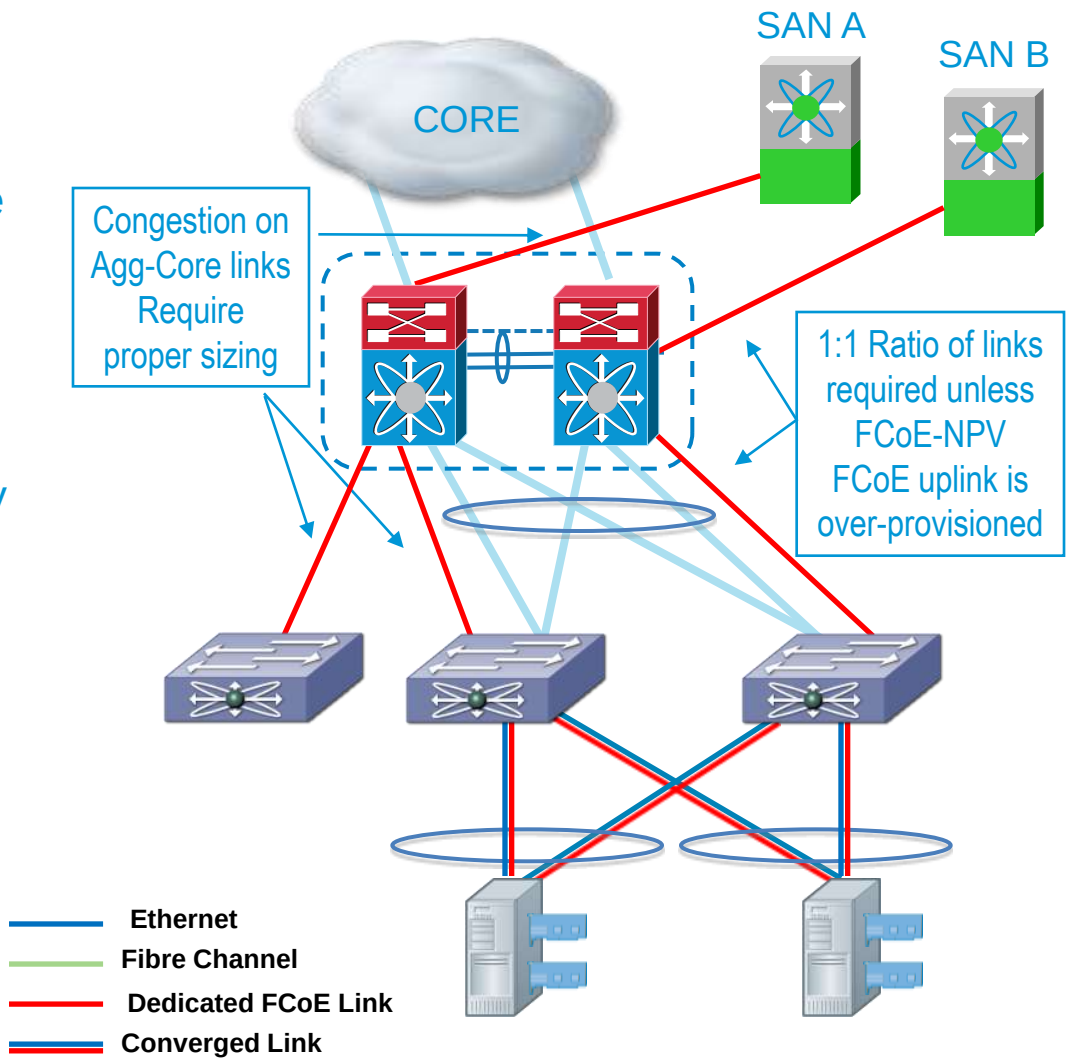

FCoE Deployment Considerations

- **Multi-hop edge/core/edge topology**
- **Core SAN switches supporting FCoE N7K with DCB/FCoE line cards MDS with FCoE line cards**
- **Edge FC switches supporting either:**

  N5K - FCoE-NPV with FCoE uplinks to the FCoE enabled core (VNP to VF)

  N5K or N7K - FC Switch with FCoE ISL uplinks (VE to VE)

- **Scaling of the fabric (FLOGI, …) will most likely drive the selection of which mode to deploy**

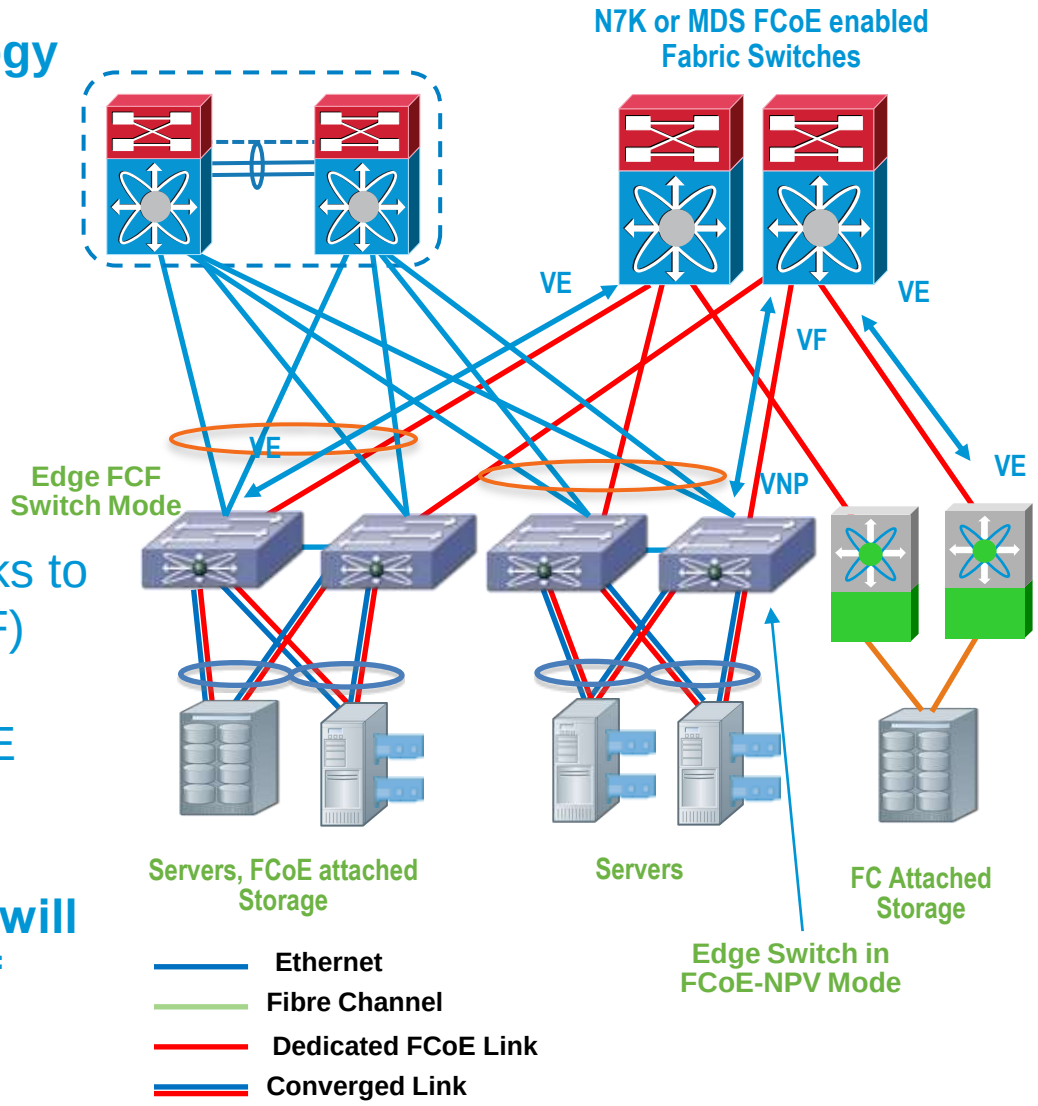

Data Center Network Manager

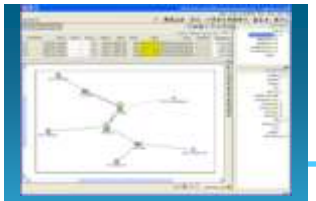

Fabric Manager FMS

DCNM

DCNM (Converged)

DCNM-SAN DESKTOP CLIENT

DCNM-LAN DESKTOP CLIENT

One converged product

- Single pane of glass (web 2.0 dashboard)
- Can be licensed for SAN and/or LAN.
- Common operations (discovery, topology)
- Single installer, Role based access control
- Consistent licensing model (licenses on server)
- Integration with UCS Manager and other OSS tools

LAN/SAN Roles

Data Center Network Manager

- Collaborative management
- Defined roles & functions
- FCoE Wizards

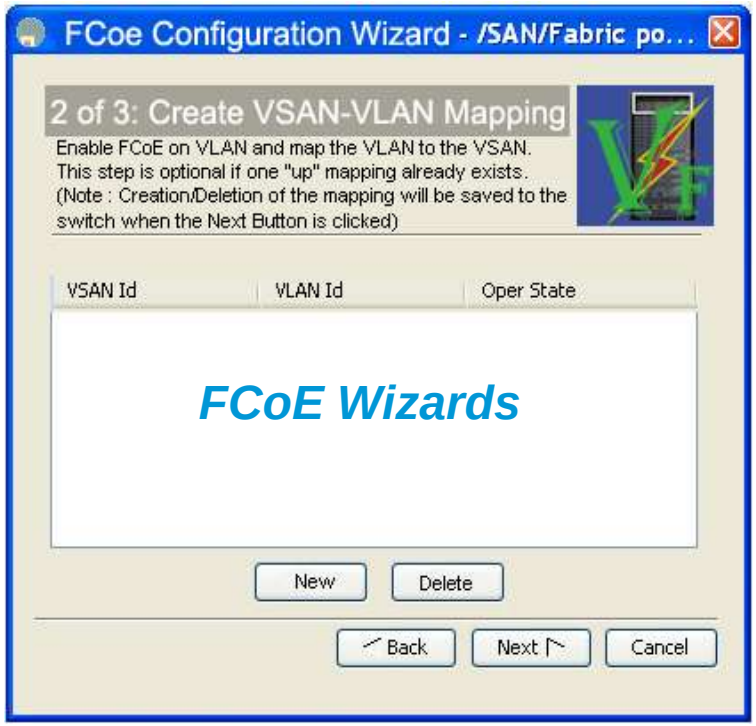

Nexus 7000	Tasks	Tools
LAN Admin	Storage VDC provisioning VLAN management Ethernet config (L2, network security, VPC, QoS, etc. DCB Configuration (VL, PFC, ETS Templates)	DCNM-LAN
SAN Admin	Discovery of Storage VDCs VLAN-VSAN mapping (use reserved pool) *Wizard* vFC provisioning *Wizard* Zoning	DCNM-SAN

Cisco Recommended Reading

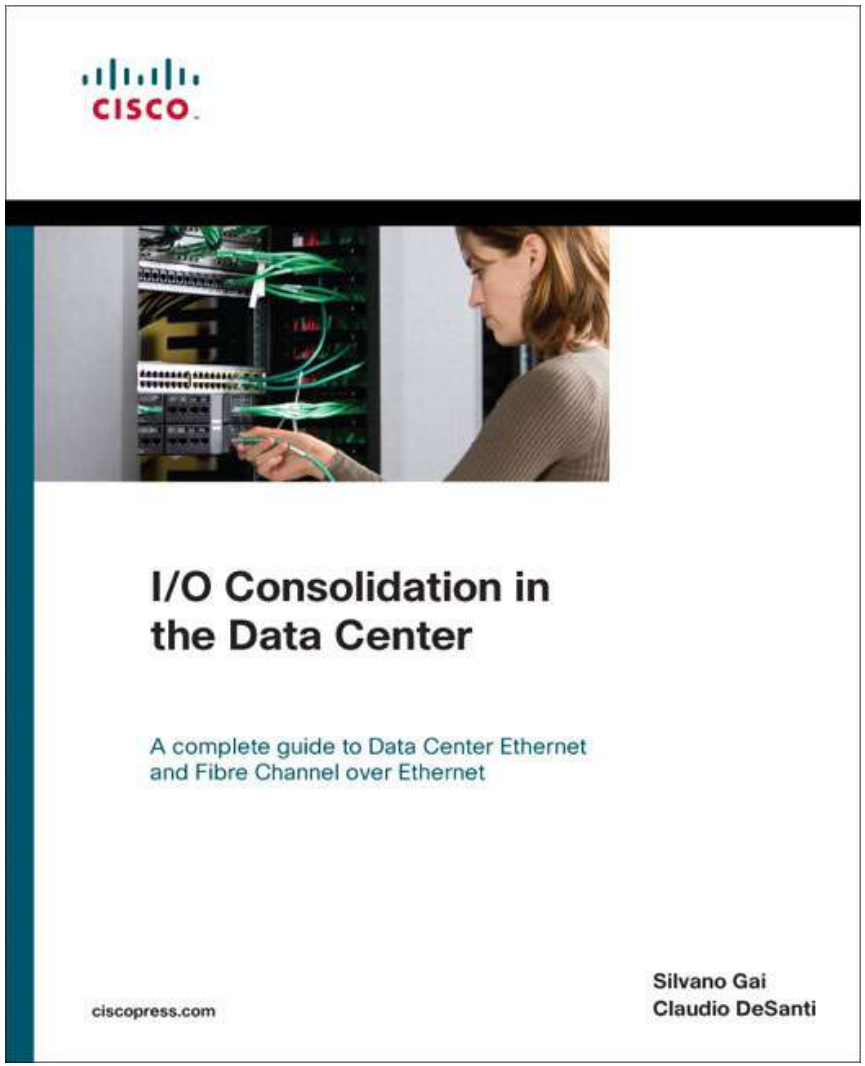

Otázky a odpovědi

- Twitter www.twitter.com/CiscoCZ
- Talk2Cisco www.talk2cisco.cz/dotazy
- SMS 721 994 600

- Zveme Vás na **Ptali jste se…** v sále **LEO**
  1.den 17:45 – 18:30
  2.den 16:30 – 17:00

Kód přednášky

Prosíme, ohodnoťte tuto přednášku.