10 Gbps Cabling

DC Facilities Top of Mind

Complexity, Cost, Power, Cooling
Standards Compliance
Reliability, Availability
Management, Security
Future Proofing

Increased Efficiency,
Simpler Operations
Scalability, Flexibility
Technology
Modularity, Mobility

10GE Server Evolution in the DC access

- The 10GbE LOM market will start in 2011 (Intel server arch. release cycle)
- 10G LOMs will show up first on premium servers.
- From 2011 it might take about 2 years to see the majority of servers 10GBASE-T LOM capable.
- Fiber & CX1 today (Multiple optics choices)
- Fiber, CX1, 10GBASE-T tomorrow

10 Gigabit Transmissions

Different Standards

- 10GBase-T (IEEE 802.3an)
- 10GBase-CX4 (IEEE 802.3ak)
- 10GBase-*X (IEEE 802.3ae)
- SFF 8431 (SFP+ Fiber & cu)

Applications

- Server Interconnects
- Aggregation of Network Links
- Switch to Switch Links
- Storage Area Networks (SAN)

10 Gbps Transciever Options

10G Sever Connectivity Options–UTP/F-UTP, MMF, SMF, TwinAx, CX4

Connector (Media)	Cable	Distance	Power (each side)	Transceiver Latency (link)	Standard
SFP+ CU* copper	Twinax	<10m	~ 1.5W	~ .1 µs	SFF 8431**
X2 CX4 copper	Twinax	15m	4W	~ 0.1 µs	IEEE 802.3ak
SFP+ USR MMF, ultra short reach	MM OM2 MM OM3	10m 100m	1W	~ 0	none
SFP+ SR MMF,short reach	MM OM2 MM OM3	82m 300m	1W	~ 0	IEEE 802.3ae
RJ45 10GBASE-T copper	Cat6 Cat6a/7 Cat6a/7	55m 100m 30m	~ 6W*** ~ 6W*** ~ 4W***	2.5 µs 2.5 µs 1.5 µs	IEEE 802.3an

* Terminated cable ** Draft 3.0, not final *** As of 2008; expected to decrease over time

10GE SFP+ Optical

- Smallest 10GE form factor
- Low Power
- Low Latency
- Hot swappable
- High density
- Optical SFP+ interoperates with other 10GE modules

  XFP

  XENPAK

  X2

SFP+ Optical Module

Optics Positioning for Data Center Applications

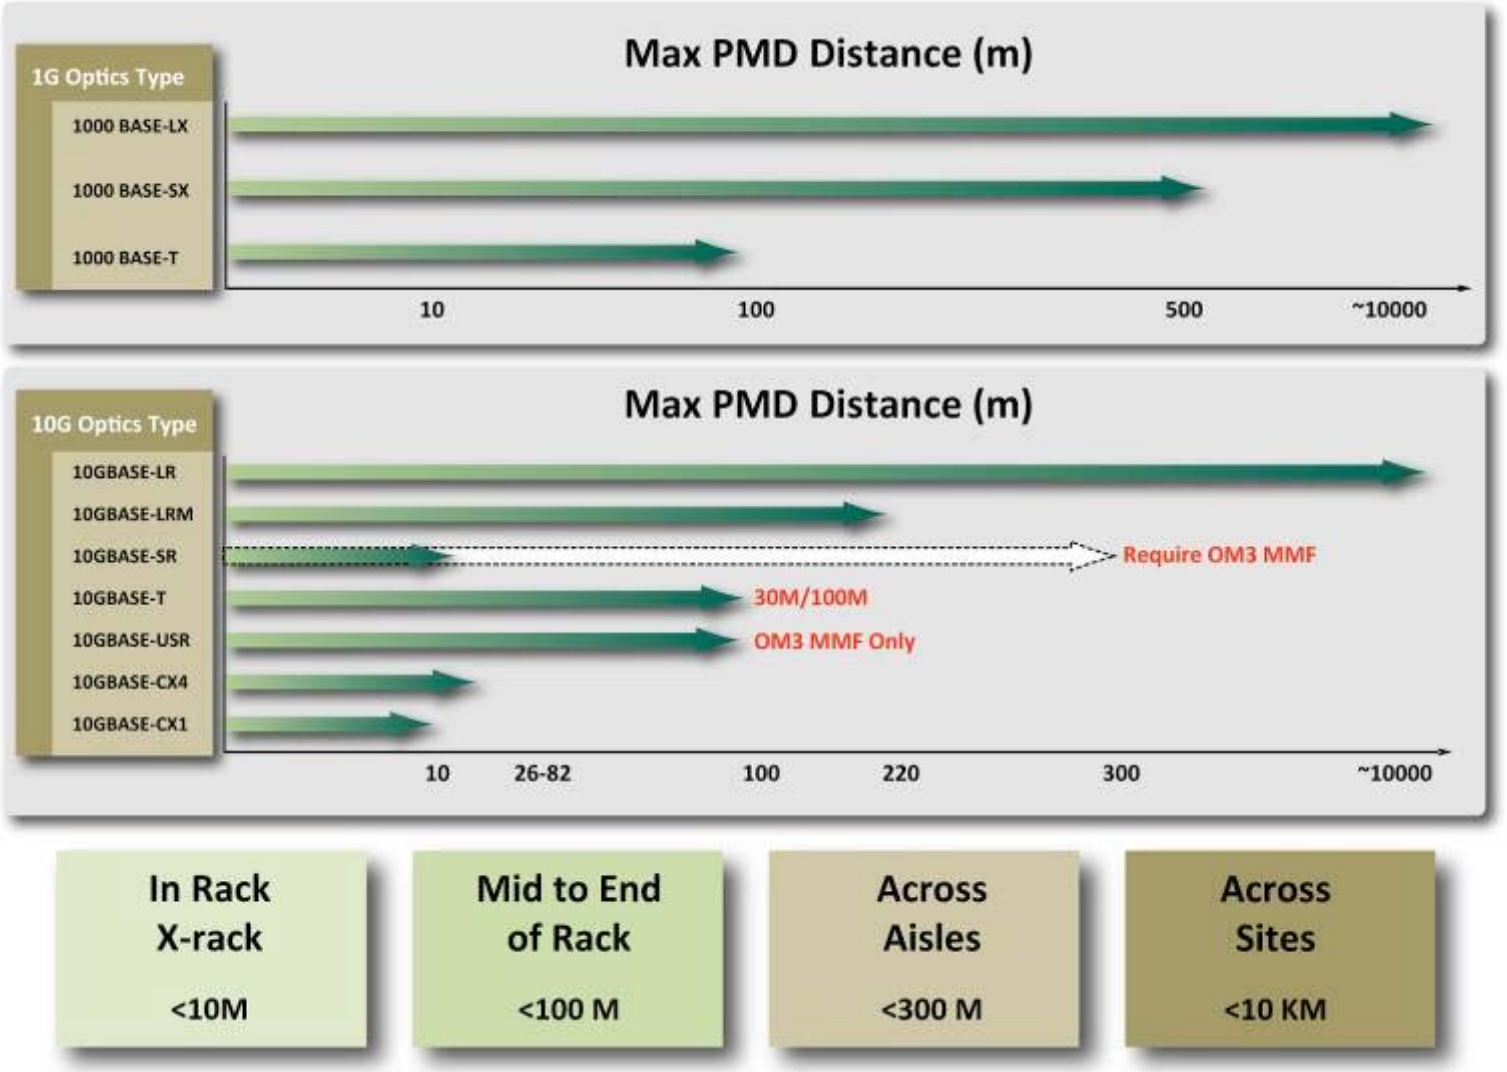

A look at the 100M > 1G transition

(Desktop and Server driven)

- **1G/10G server transition expected to be within ~2-years once LOM is available (slower than at 1G – source Intel)**
- **Switch edge port adoption of high speed technology historically took longer than server transition (hence the need for 100/1000/10000 technology on LOM)**

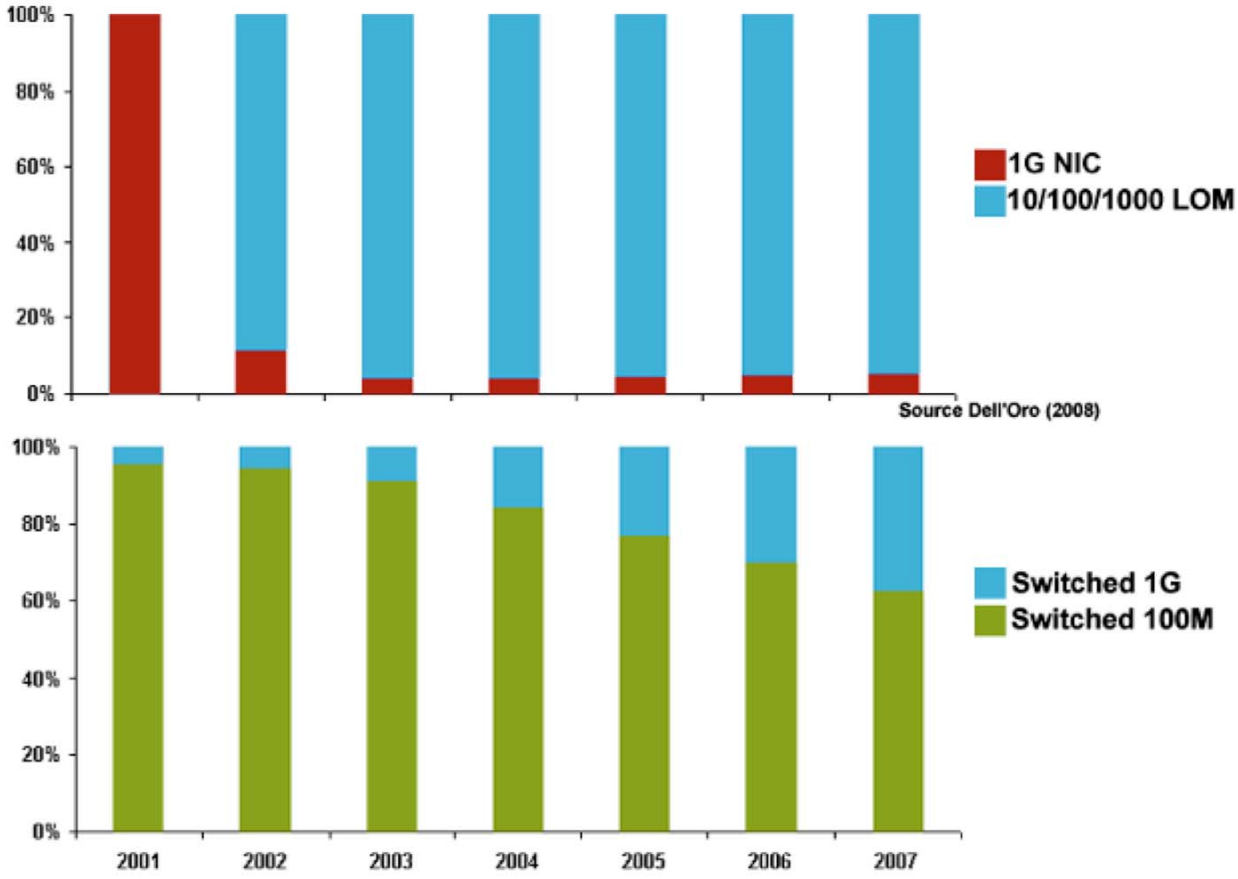

10GBASE-T PHY Evolution

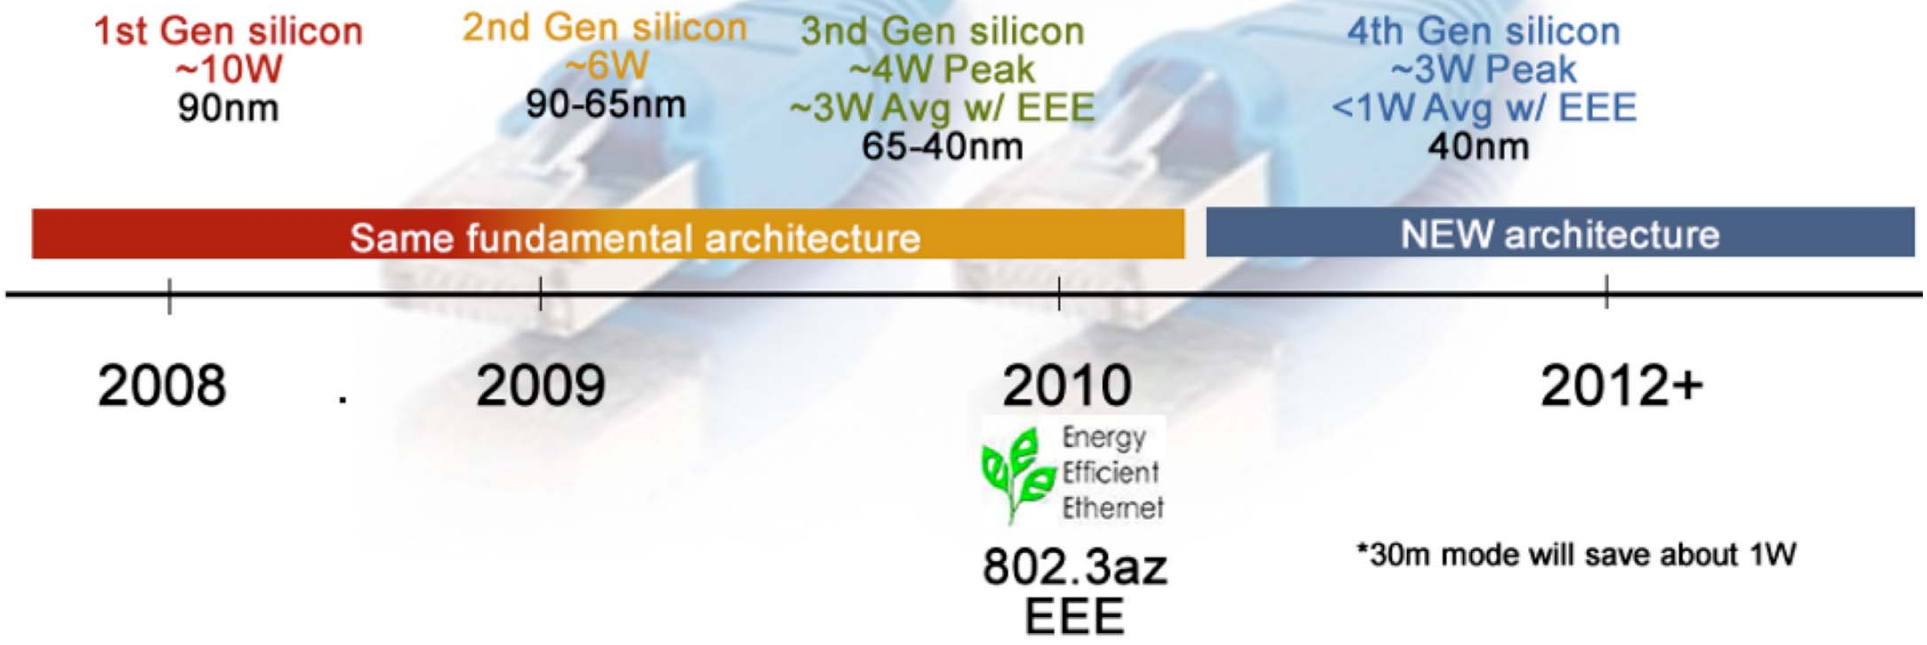

10GBASE-T PHY Evolution – Switch View

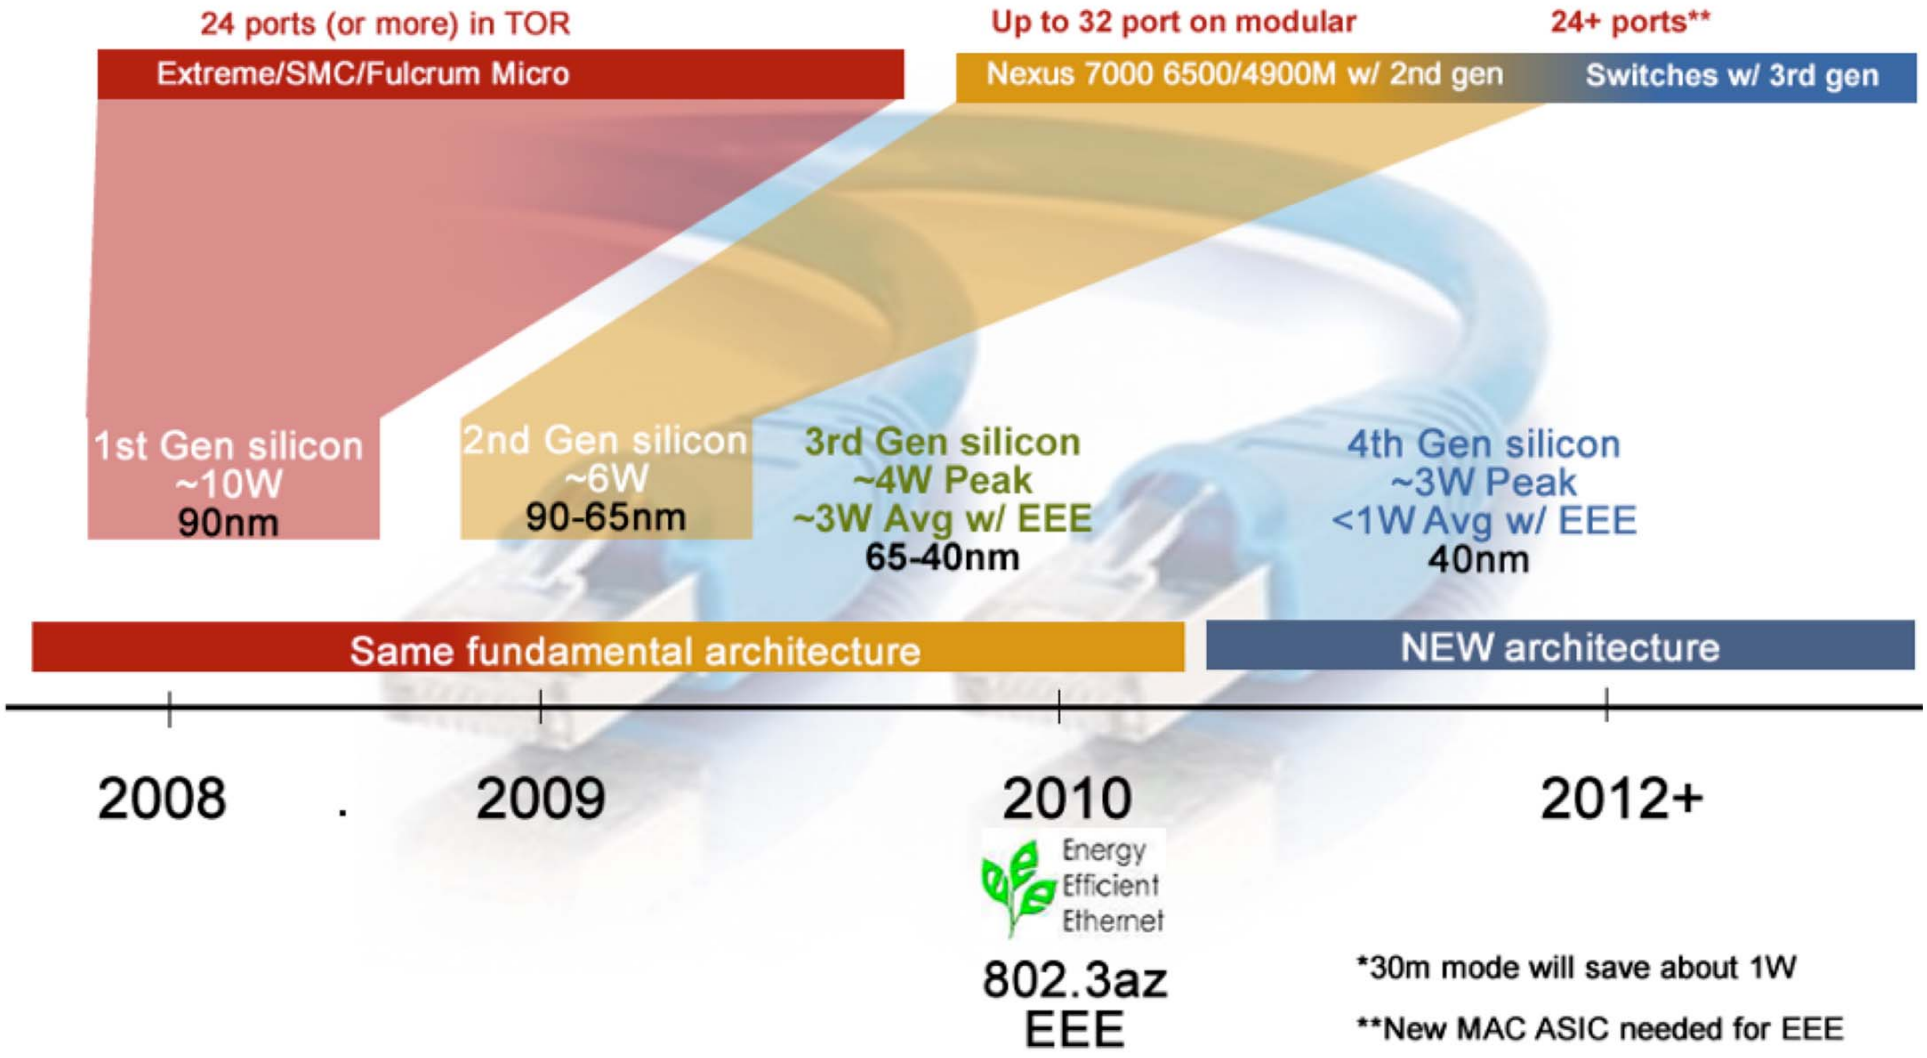

10GBASE-T PHY Evolution – Server View

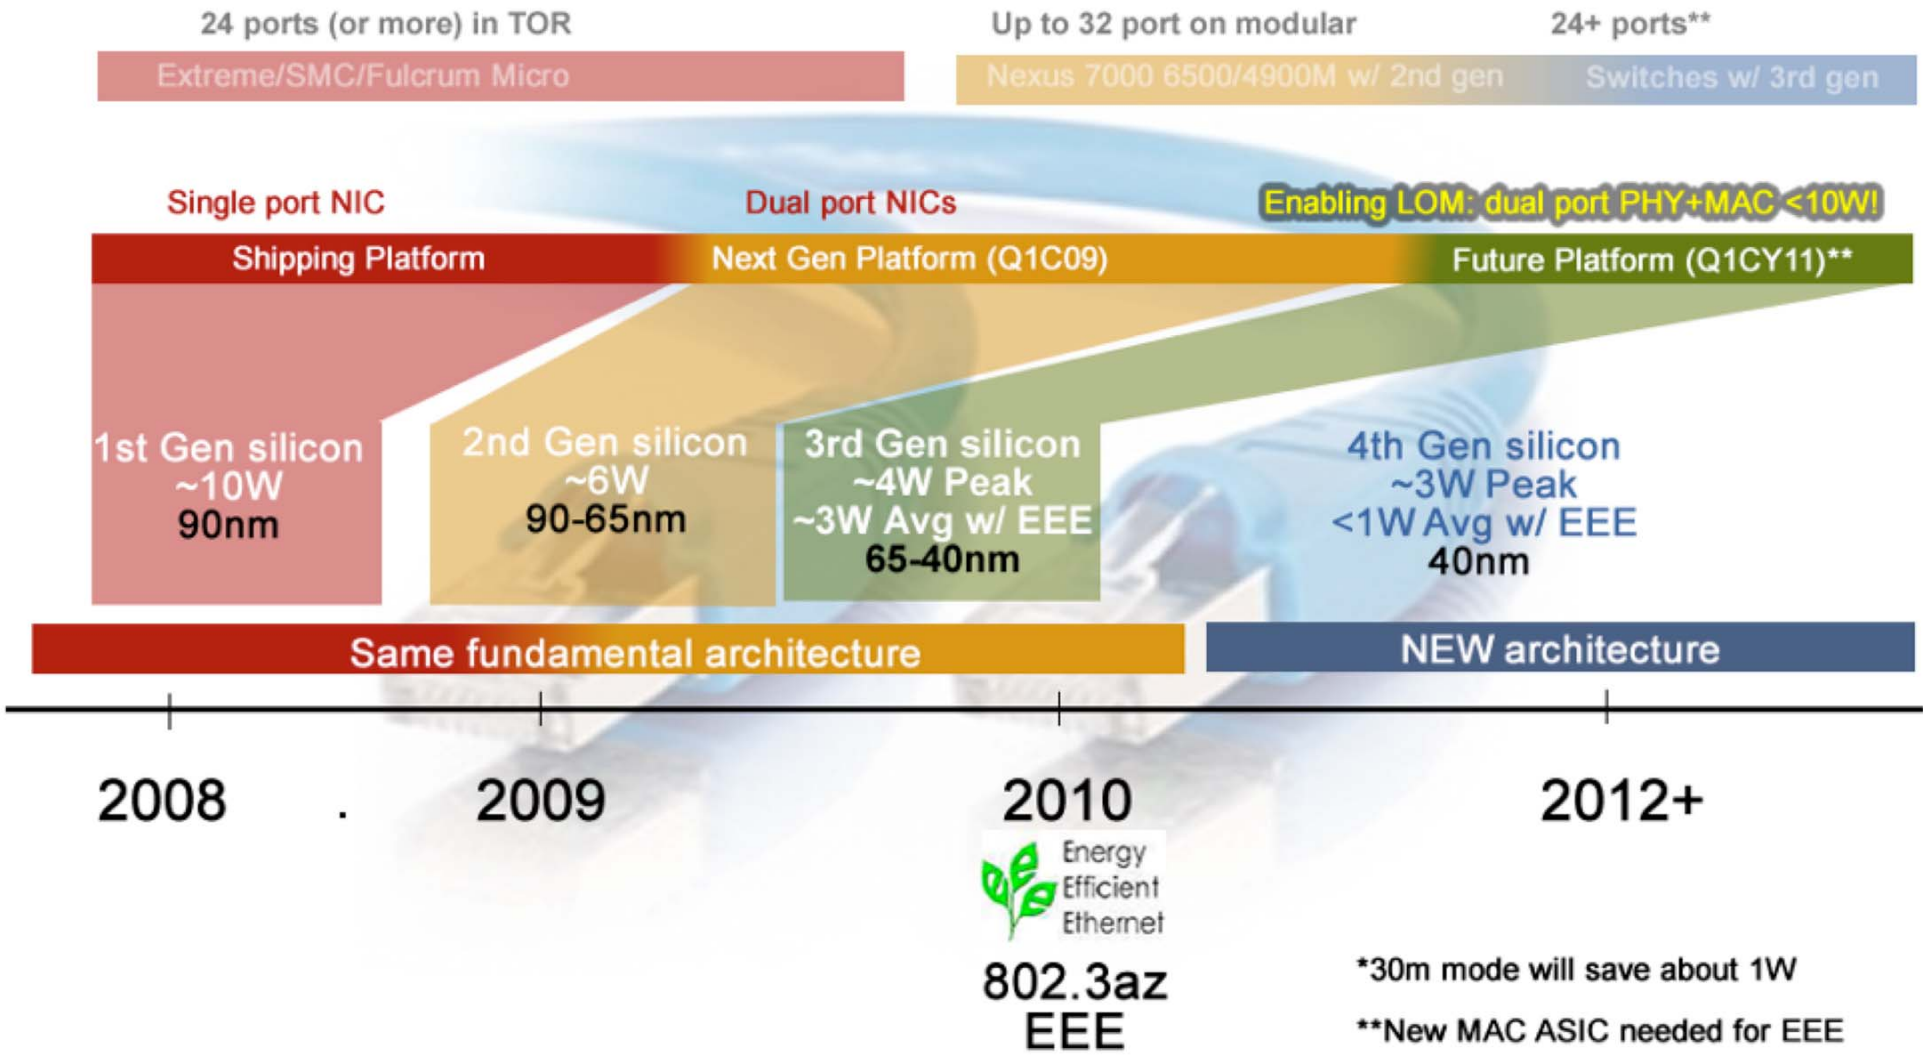

Twisted Pair Cabling For 10GBASE-T
(IEEE 802.3an)

U/UTP (Old designation UTP)
Outer Unshielded/Inner Pairs Unschielded

F/UTP (Old designation FTP)
Outer Foil Shielded/Inner Pairs Unshielded

S/FTP (Old designation S/STP)
Outer Foil Shielded/Inner Pairs Foil shielded

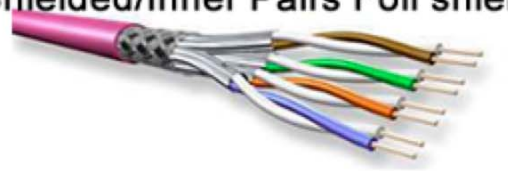

Cat 6a:
*100m 10GBASE-T
**largest diameter up to 0.354 in

Cat 6:
*55m 10GBASE-T
**larger diameter than Cat5 (~0.3 in)

Cat 6/6a:
*100m 10GBASE-T
**More flexible/easier to manage than Cat6a U/UTP
***Equivalent diameter to Cat6

Cat 7:
*100m 10GBASE-T
**Most expensive
***Smaller diamter than Cat6a
****Not popular in North America

Shielded (F/UTP) vs. Unshielded (6a U/UTP)

F/UTP

- Alien XT reduction w/ shielding
- Diameter 0.3in/7.6mm
- Up to* 40% cable conduit fill-ratio higher than U/UTP
- Cabling choice in Cisco's new Richardson Data Center

U/UTP

- Alien XT reduction w/ spacing
- Diameter 0.354in/9mm (worst case)
- Stiffer/less flexible cable than F/UTP

F/UTP **U/UTP**

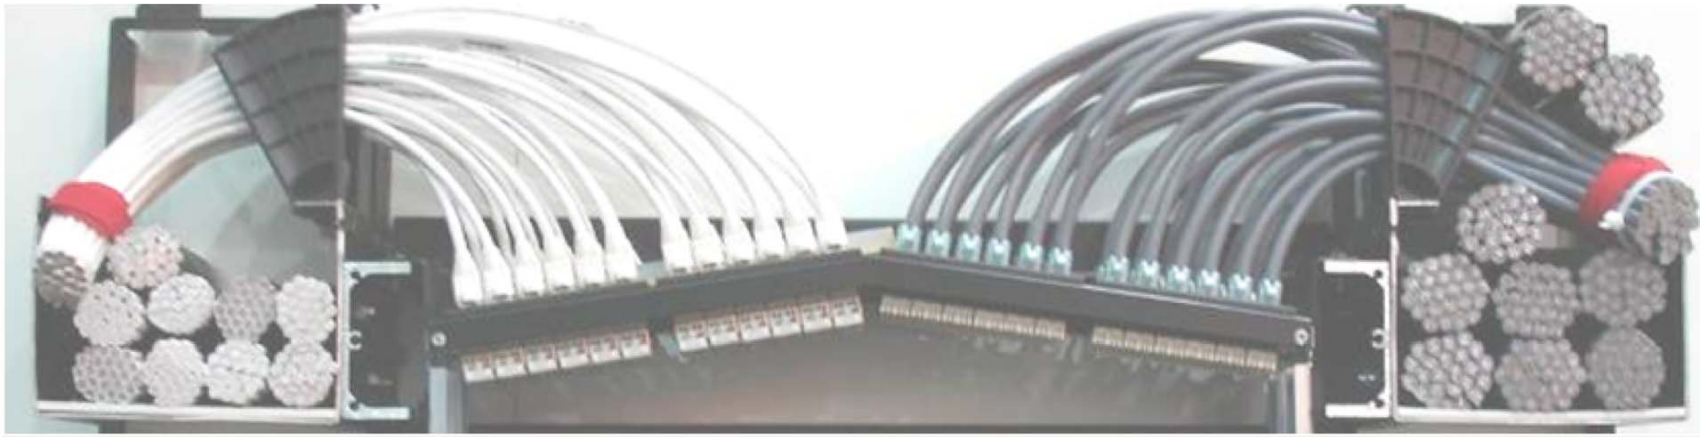

*Assumption of 40% conduit fill ratio and various conduit sizes (Source Tyco)

Copper Cabling for Data Centers

Frequency Bandwidth	TIA (Components)	TIA (Cabling)	ISO (Components)	ISO (Cabling)
1 - 100 MHz	Category 5e	Category 5e	Category 5e	Class D
1 - 250 MHz	Category 6	Category 6	Category 6	Class E
1 - 500 MHz	**Category 6A**	**Category 6A**	**Category 6A**	**Class E_A**
1 - 600 MHz	n/s	n/s	**Category 7**	**Class F**
1 - 1,000 MHz	n/s	n/s	**Category 7_A**	**Class F_A**

Application	Category 5e Class D	Category 6 Class E	Category 6A Class E_A	Class F	Class F_A
4/16 MBPS Token Ring	x	x	x	x	x
10BASE-T	x	x	x	x	x
100BASE-T4	x	x	x	x	x
155 MBPS ATM	x	x	x	x	x
1000BASE-T	x	x	x	x	x
TIA/EIA-854		x	x	x	x
10GBASE-T			**x**	**x**	**x**
Broadband CATV				x	x

Charts extracted from www.siemon.com

10G SPF+ Cu

- SFF 8431
- Supports 10GE passive direct attached up to 10 meters
- Active cable options to be available
- Twinax with direct attached SFP+
- Primarily for in rack and rack-to-rack links
- Low Latency, low cost, low power

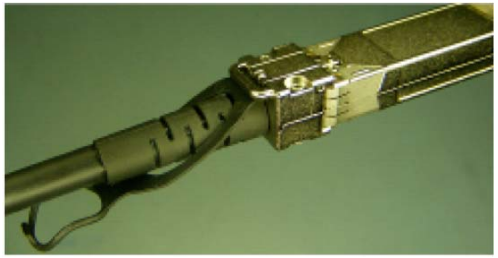

10G Copper Infiniband - 10GBase-CX4

- IEEE 802.3ak
- Supports 10G up to 15 meters
- Quad 100 ohm twinax, Infiniband cable and connector
- Primarily for rack-to-rack links
- Low Latency

HSSG: Higher Speed Study Group

HSSG Objectives

- Support full-duplex operation only
- Preserve the 802.3 / Ethernet frame format utilizing the 802.3 MAC
- Preserve minimum and maximum FrameSize of current 802.3 standard
- Support a BER better than or equal to 10^{-12} at the MAC/PLS service interface
- Provide appropriate support for OTN

- Support a MAC data rate of 40 Gb/s
- Provide Physical Layer specifications which support 40 Gb/s operation over:
  - at least 100m on OM3 MMF
  - at least 10m over a copper cable assembly
  - at least 1m over a backplane

- Support a MAC data rate of 100 Gb/s
- Provide Physical Layer specifications which support 100 Gb/s operation over:
  - at least 40km on SMF
  - at least 10km on SMF
  - at least 100m on OM3 MMF
  - at least 10m over a copper cable assembly

3

Access Layer Network Model

End of Row, Top of Rack & Blade Switches

What it used to be…

GE Access / 10GE

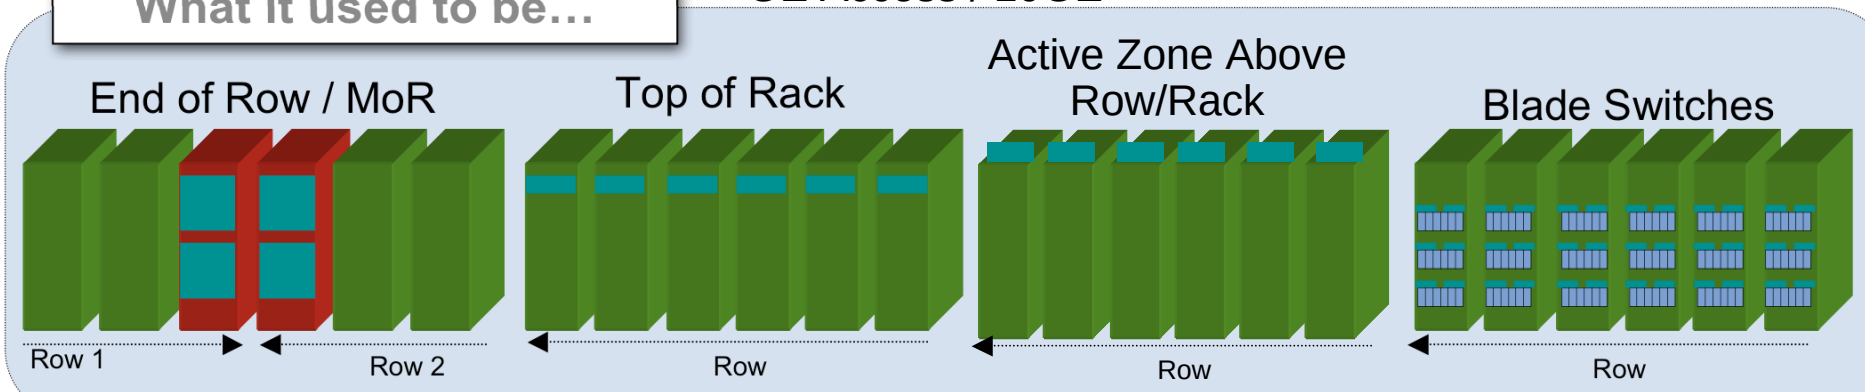

What else is coming…

What influences the physical layout…

- Primarily:
  - Power
  - Cooling
  - Cabling
- Secondarily
  - Access Model
  - Port Density

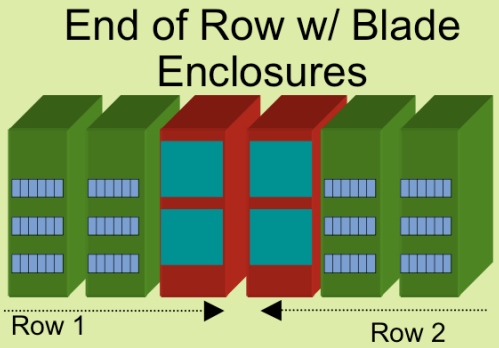

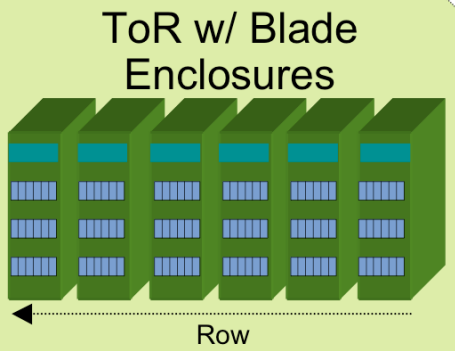

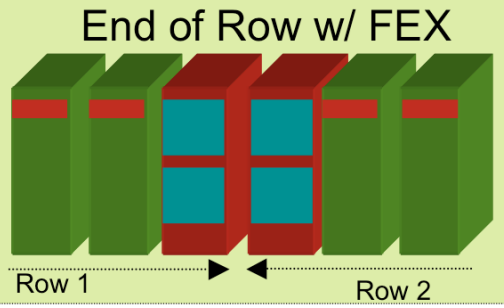

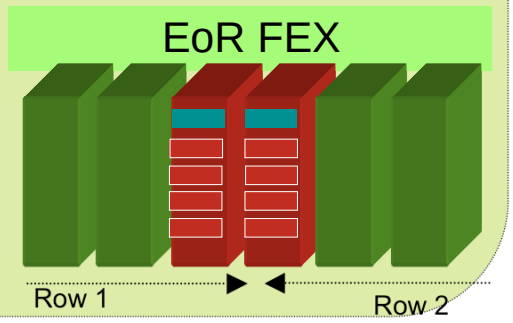

Pod Concept - Sizing

Pod Concept

Network Zones and Pods

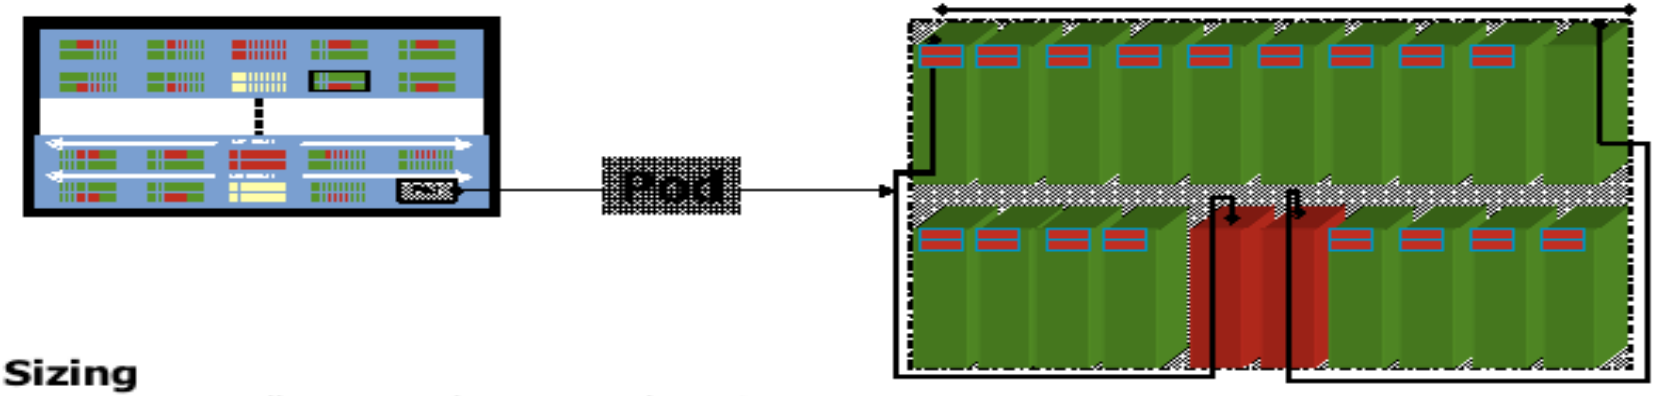

Sizing

_Zone: Typically mapped area in the DC
_Pod: Typically mapped to agg pair with ToR switches in EDA/ access cabinets
  _Size: determined by distance and density
  _Cabling distance from server racks to network racks
    _Fiber
    _200-500m Fiber
  _Cabling density: # of servesr per rack and I/Os per server
_Rack
  _Server: 6 -30 Servers per rack – limited by power
  _Network: Depends on access model: Modular, ToR or Blade
  _Storage: special Cabinets

Physical Infrastructure and Network Topology

Modular Cabling Architecture Methodology

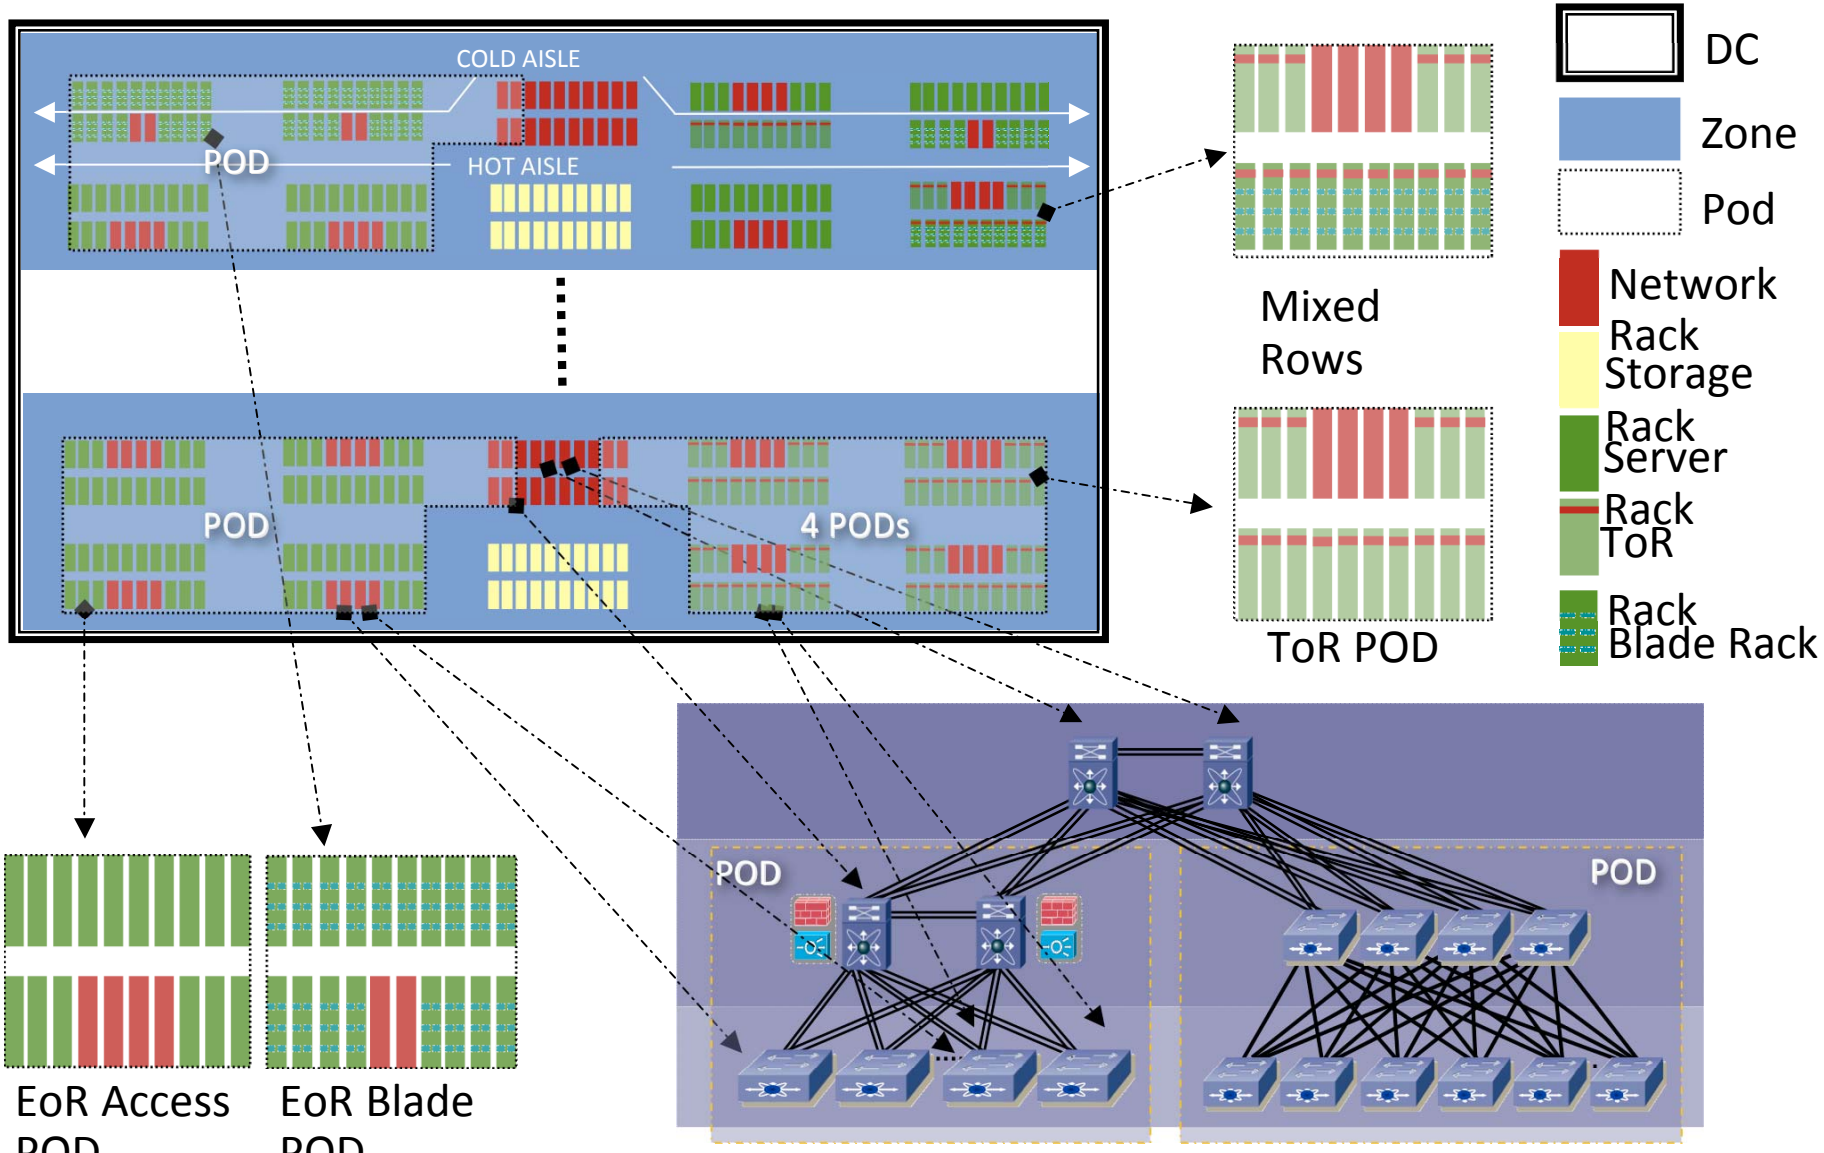

TIA-942 Logical Layout

EIA/TIA 568 Copper & Fiber Cabling

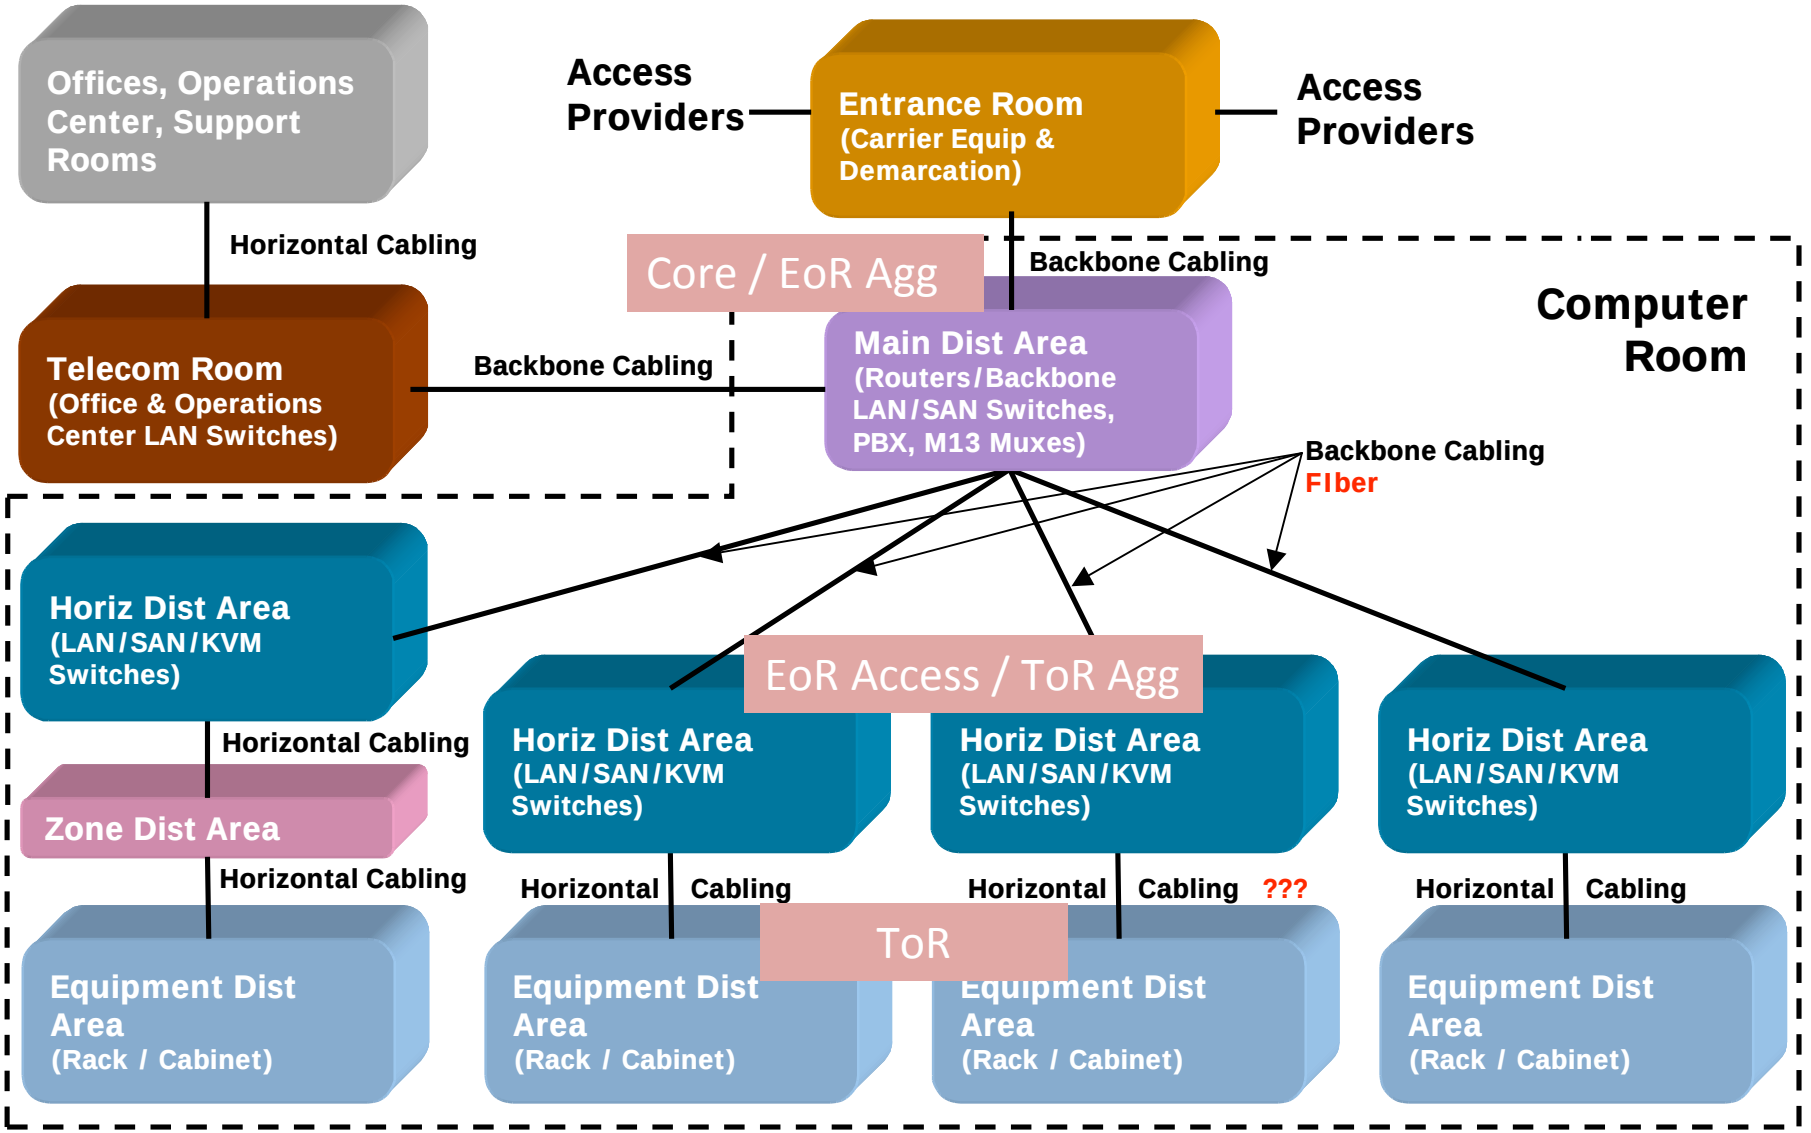

Dual-Homed 1-G Attached Servers

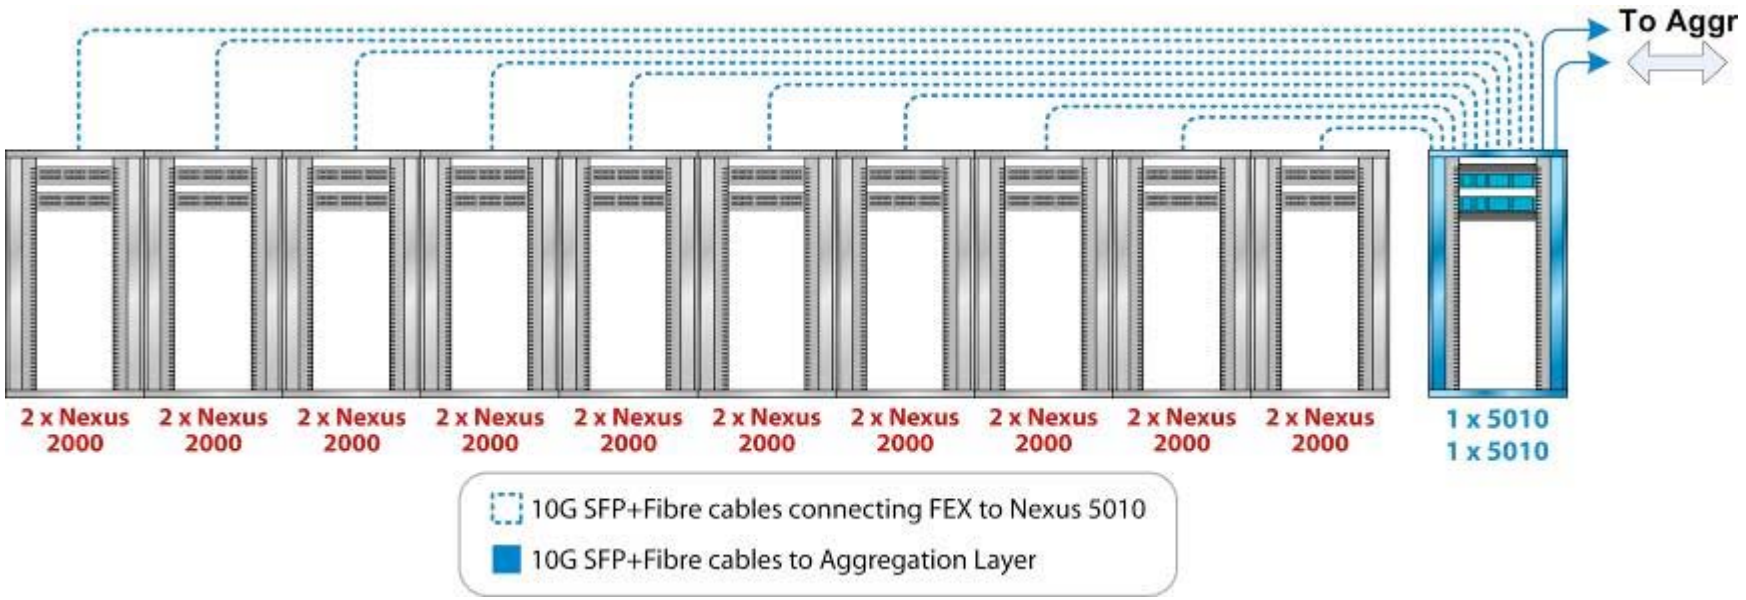

ToR Configuration

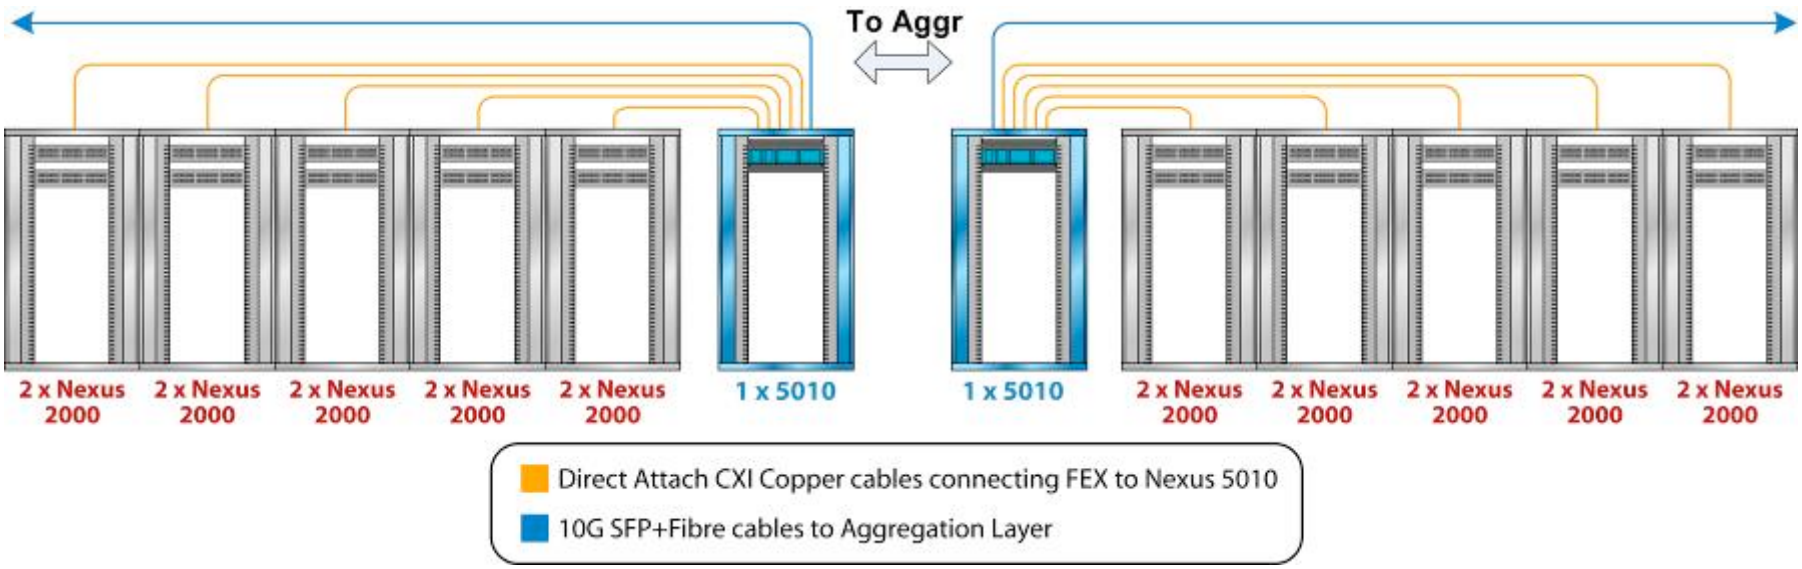

Mixed 1-G and 10-G Attached Servers

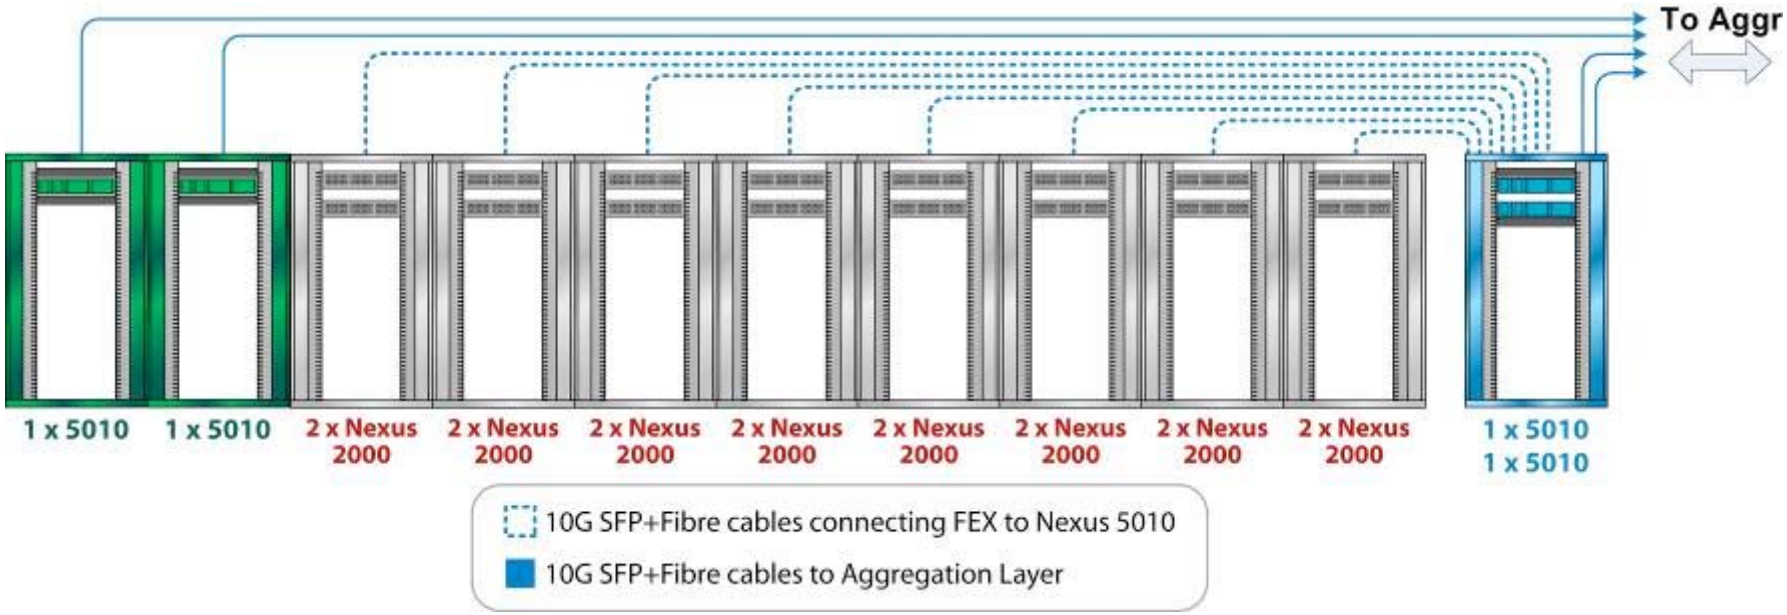

Review of the Nexus 2000 Fabric Extender

Nexus 5000 Family Overview

Nexus 5010

20 Fixed Ports 10G/FCoE/Data Center Ethernet
Line Rate, Non-Blocking 10G
1 Expansion Module slot

Nexus 2000 Fabric Extender (Q1CY09)

48 Fixed Ports 1G Ethernet
4 Fixed Ports 10G Uplink
Distributed Virtual Line Card
Hardware VN-Link Technology

Nexus 5020

40 Fixed Ports 10G/FCoE/Data Center Ethernet
Line Rate, Non-Blocking 10G
2 Expansion Module slots

Ethernet

6-Port 10G/FCoE/Data Center Ethernet

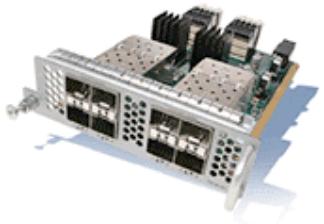

Ethernet + Fibre Channel

4-Port 10G/FCoE/Data Center Ethernet
4-Port 1/2/4G Fibre Channel

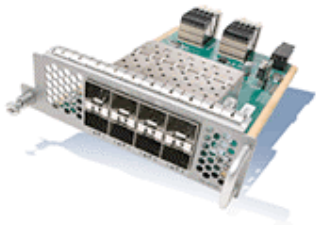

Fibre Channel

8-Port 1/2/4G Fibre Channel

Fibre Channel (Q3CY09)

2/4/8G Fibre Channel

NX-OS, DCNM, CiscoWorks and Fabric Manager

Nexus 5020 Overview

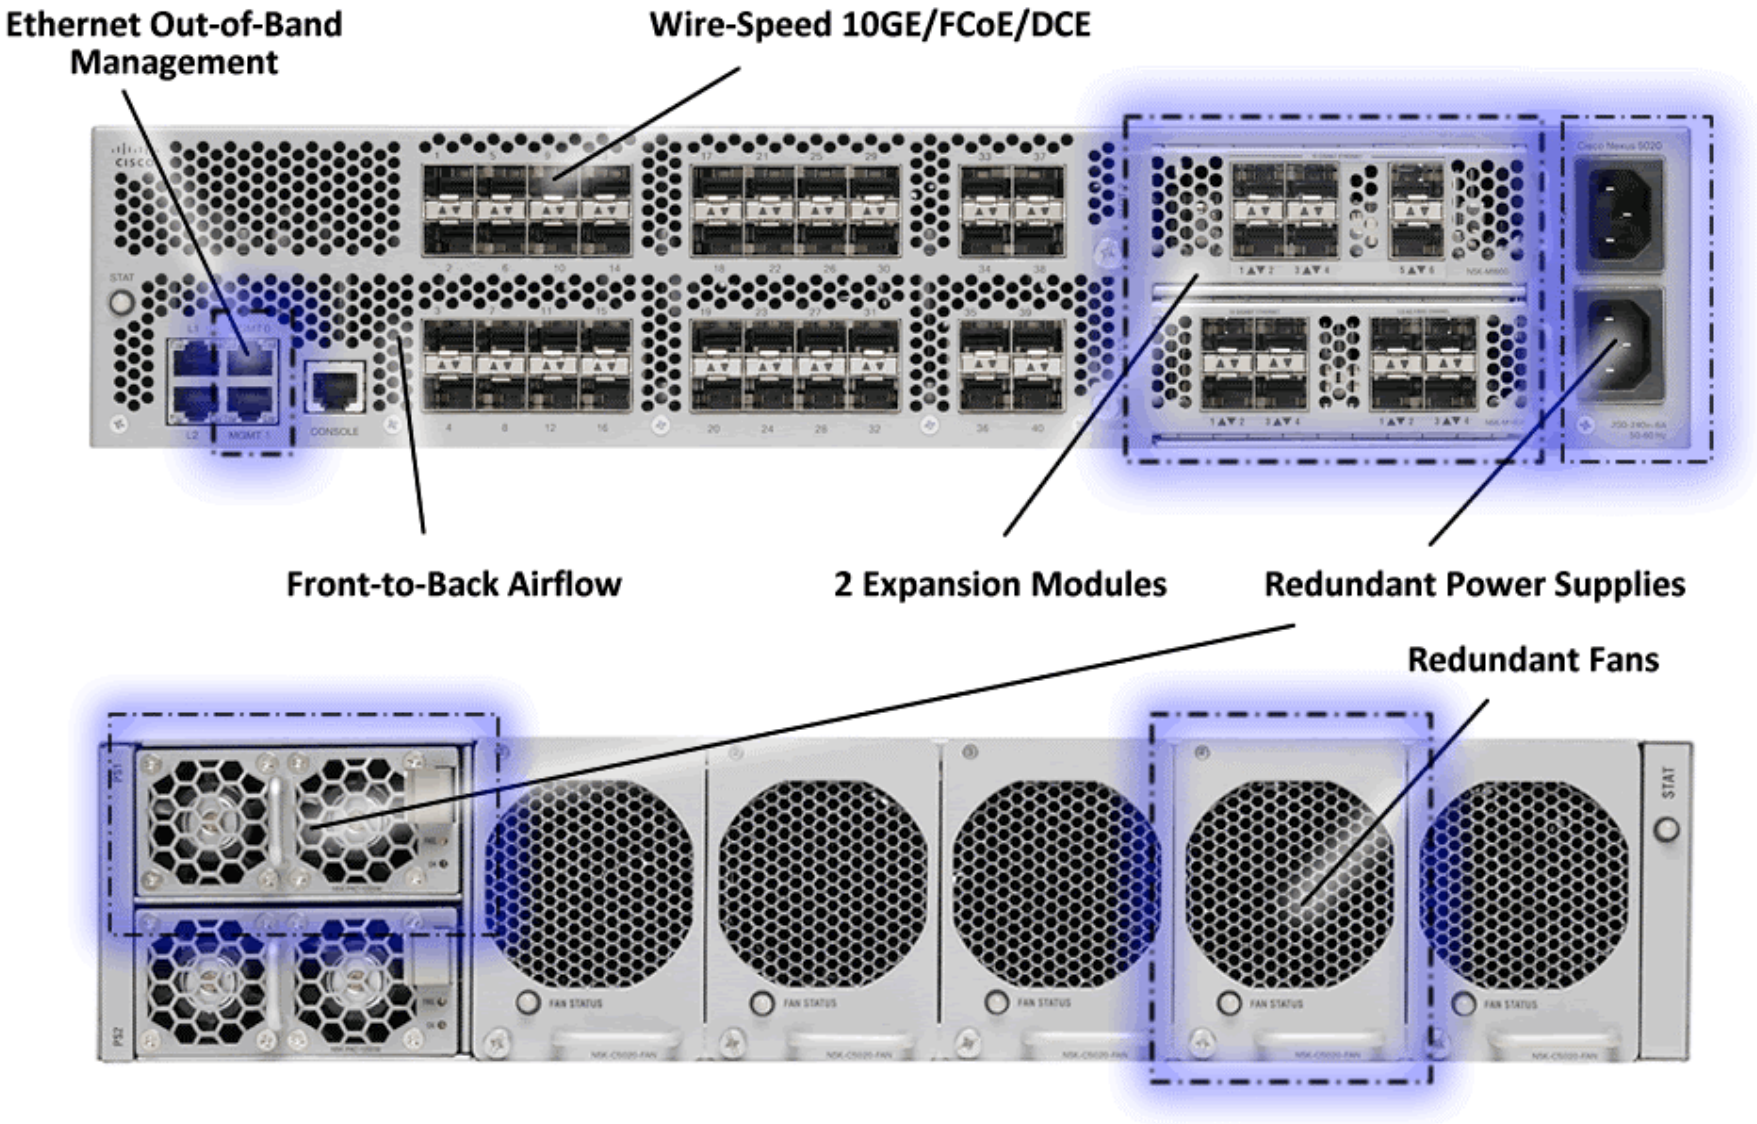

Nexus 5010 Overview

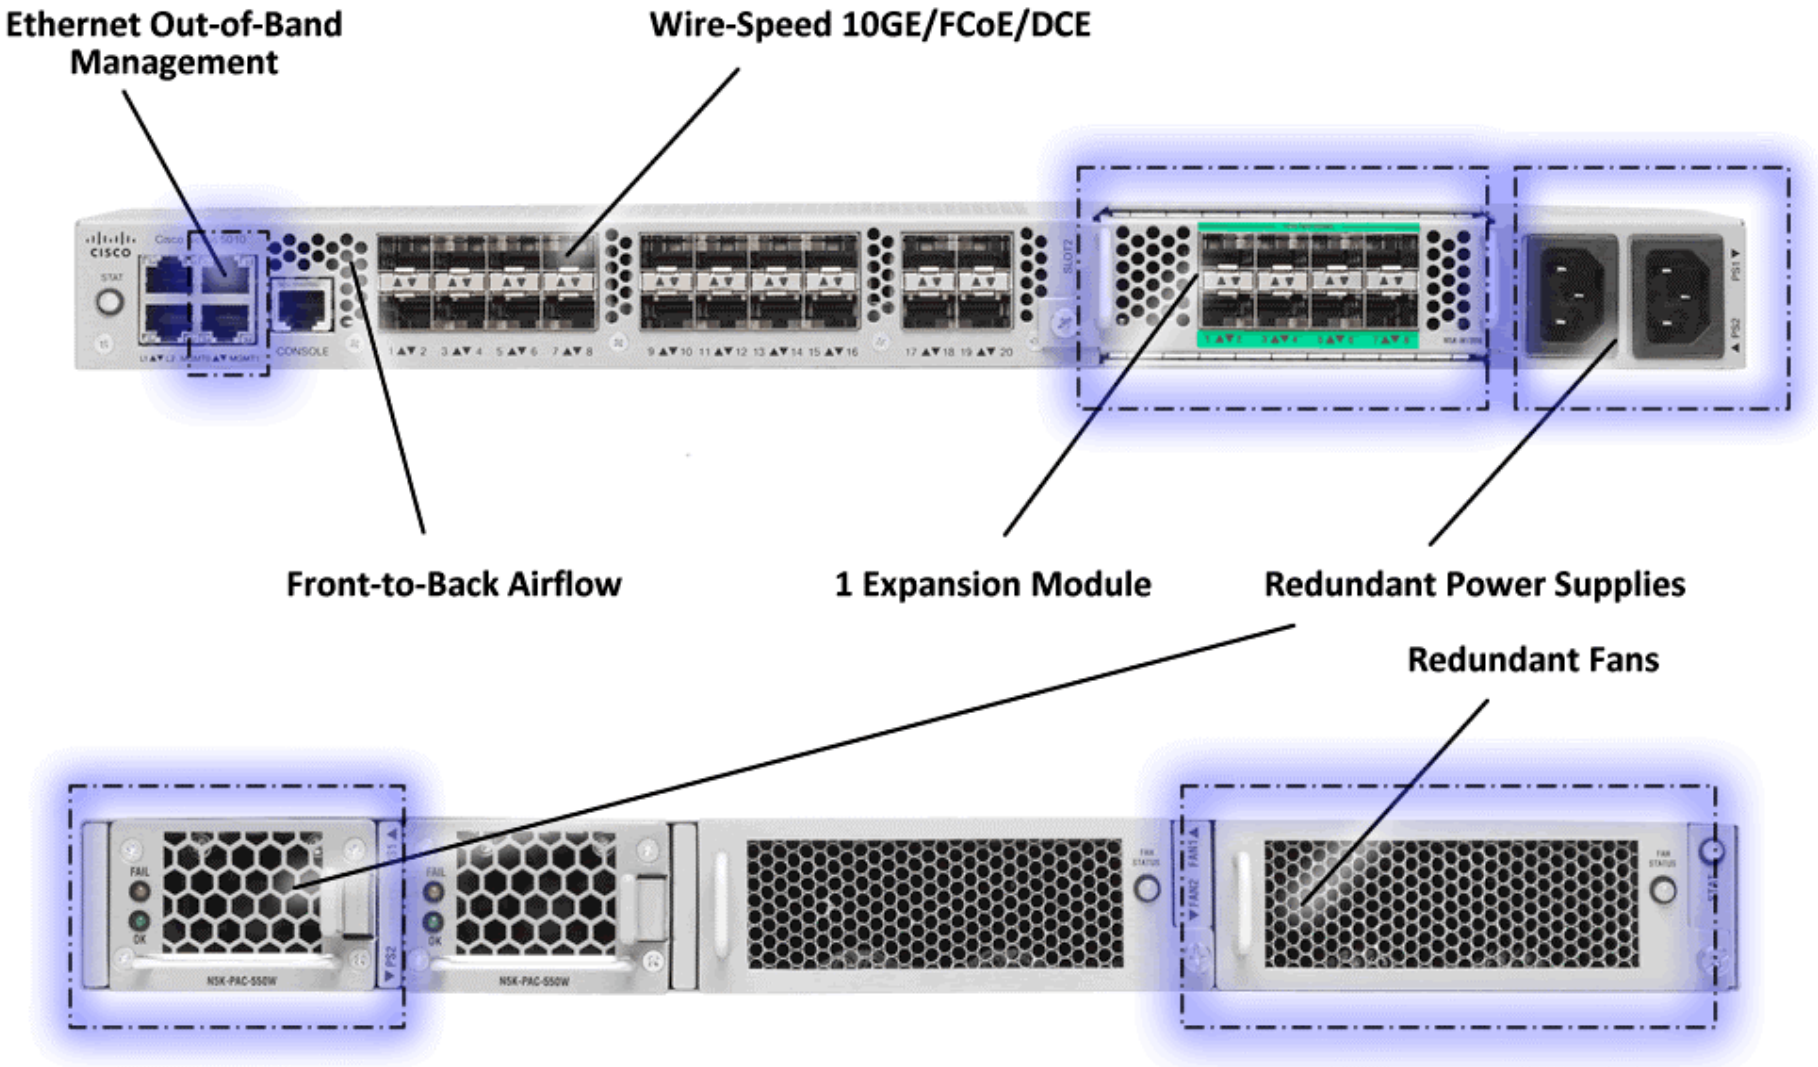

Nexus 2000 Fabric Extender
Virtual Chassis

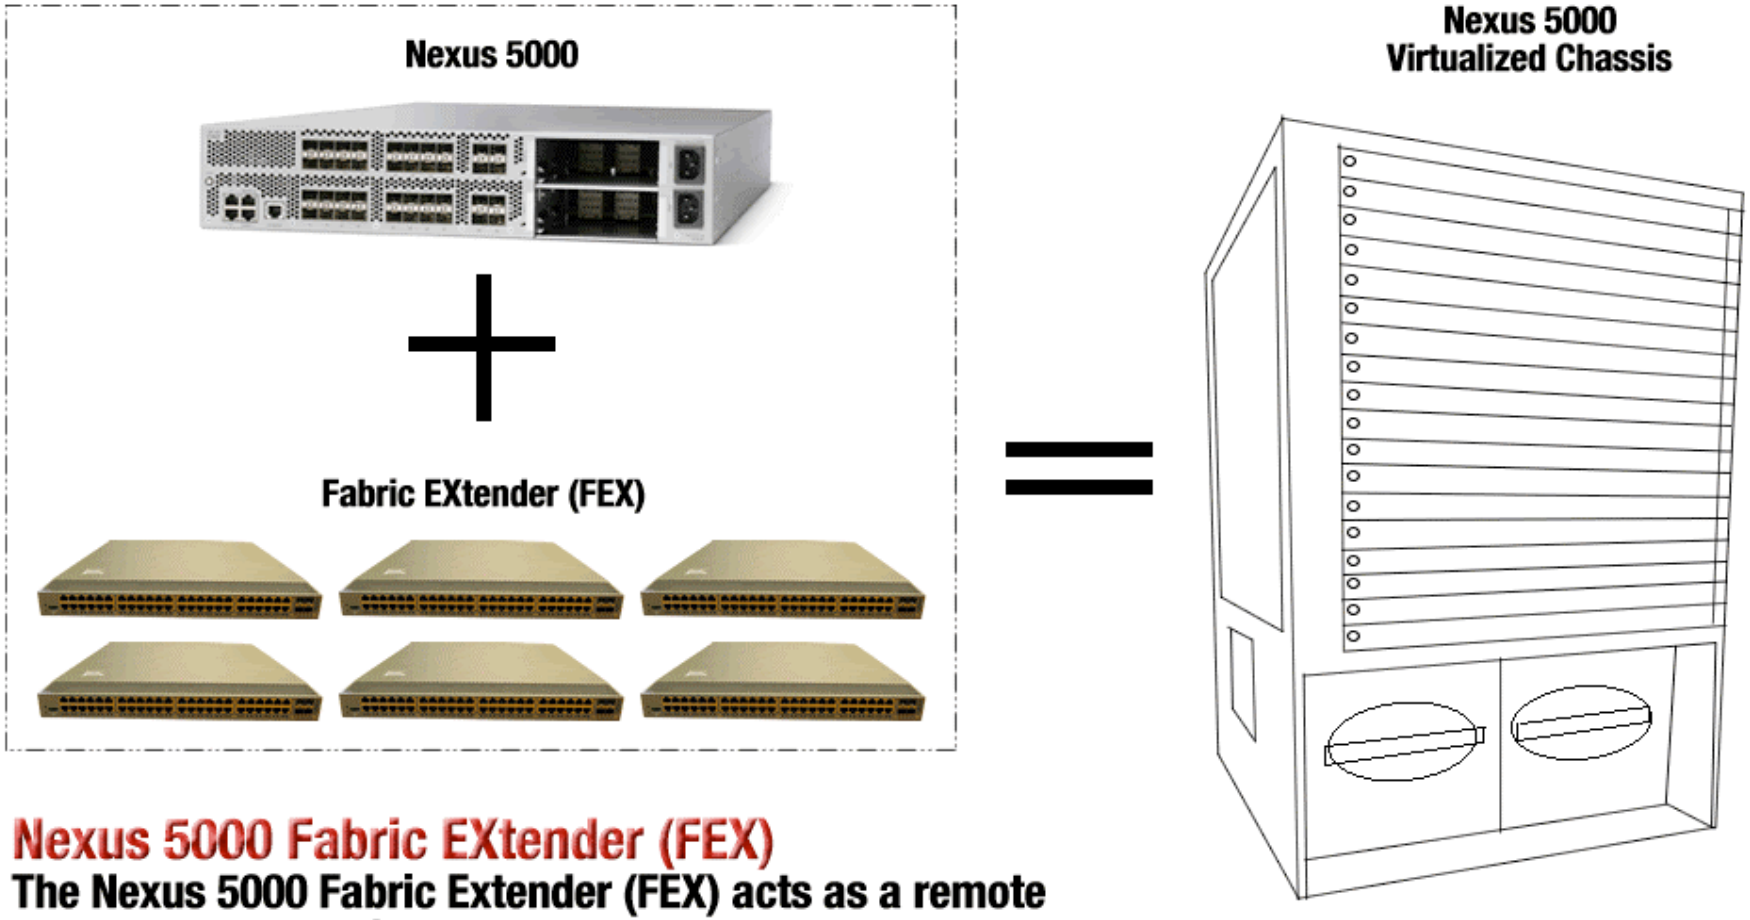

Nexus 5000 Fabric EXtender (FEX)

The Nexus 5000 Fabric Extender (FEX) acts as a remote line card (module) for the Nexus 5000, retaining all centralized management and configuration on the Nexus 5000, transforming it to a Virtualized Chassis

Nexus 2000 Fabric Extender
1GE Connectivity

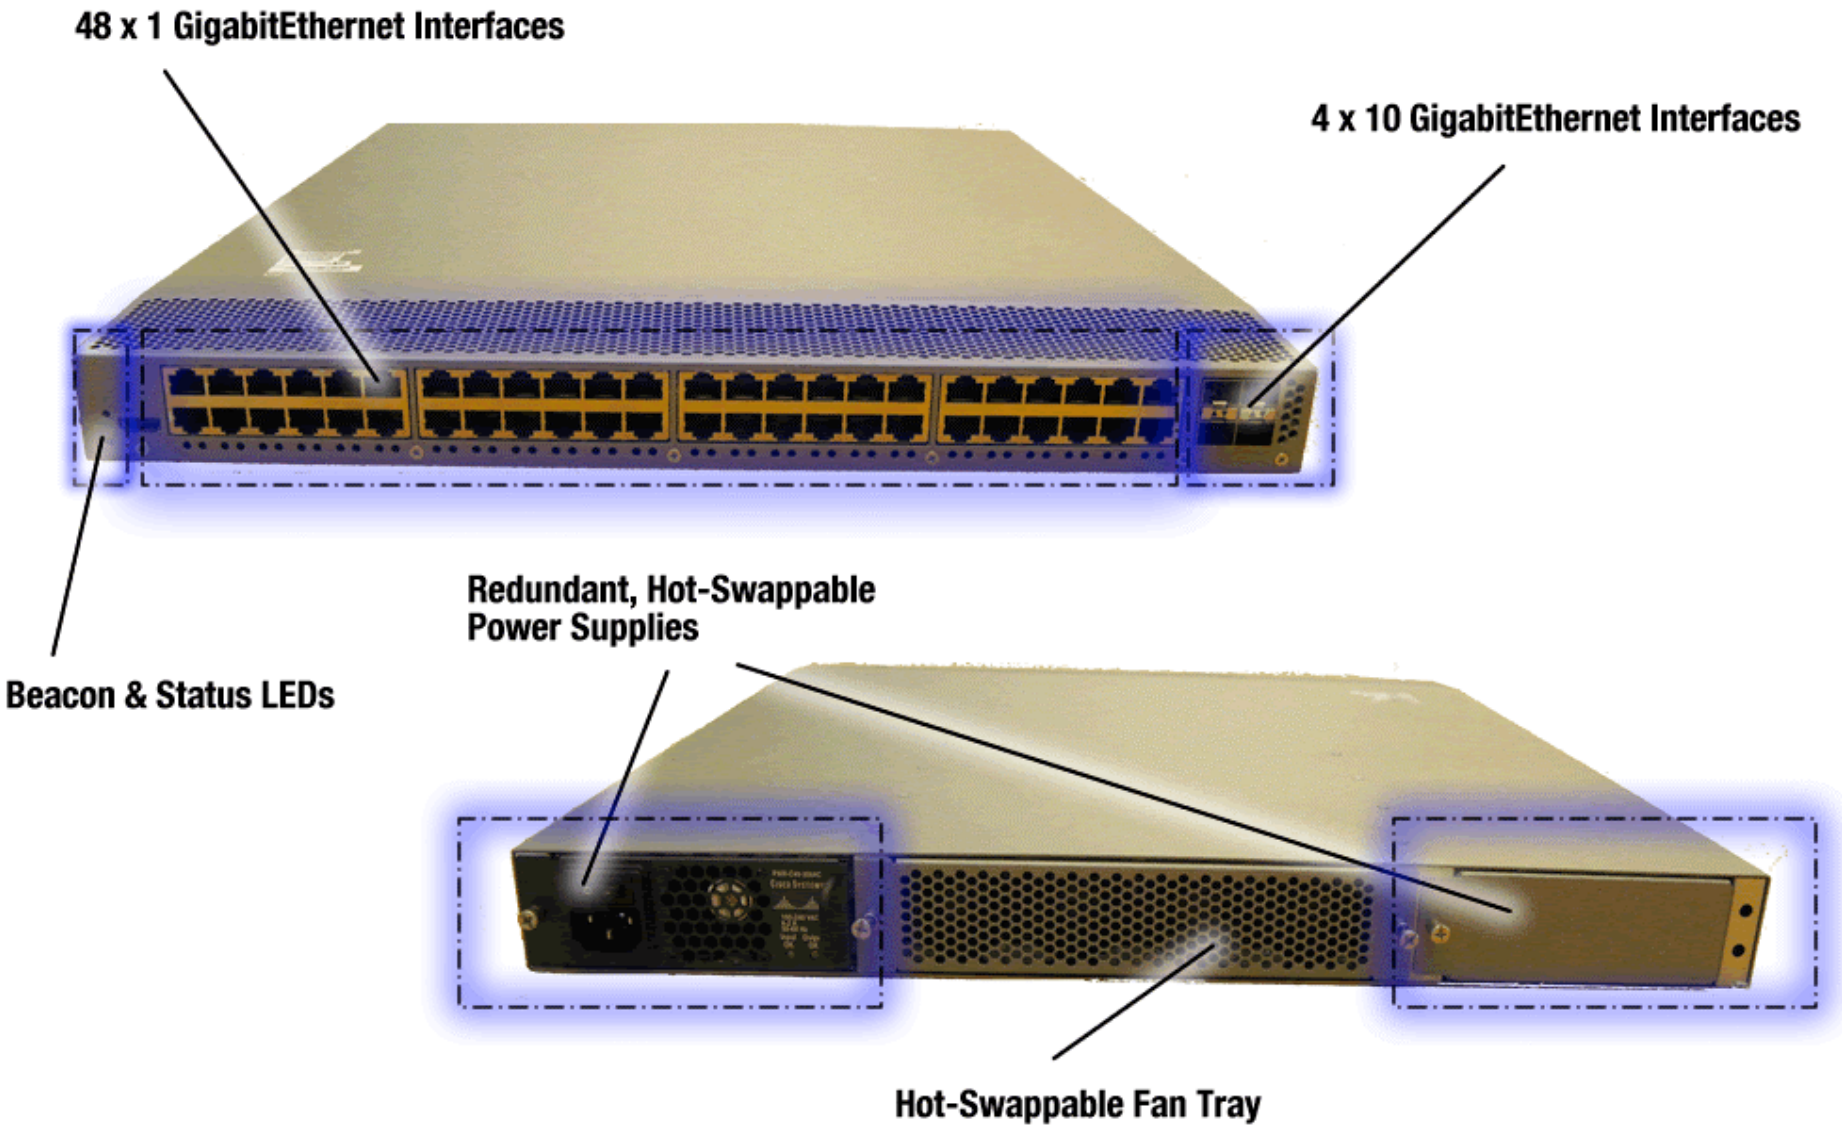

Nexus 2000 Fabric Extender
Virtual Chassis Mgmt

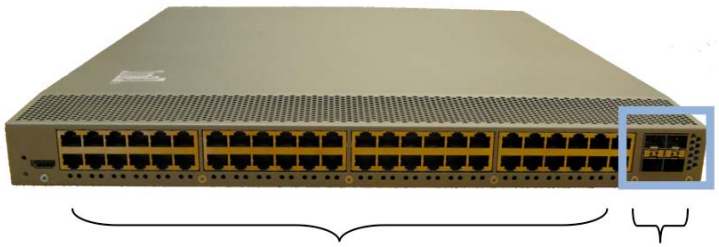

- From a management perspective, fabric extender is modeled as a NX5K linecard
- Nexus 5000 and Fabric extender communicate via inband connectivity
- There is no configuration stored on the Fabric Extender
- Fabric extender is ***not*** an independent manageable entity

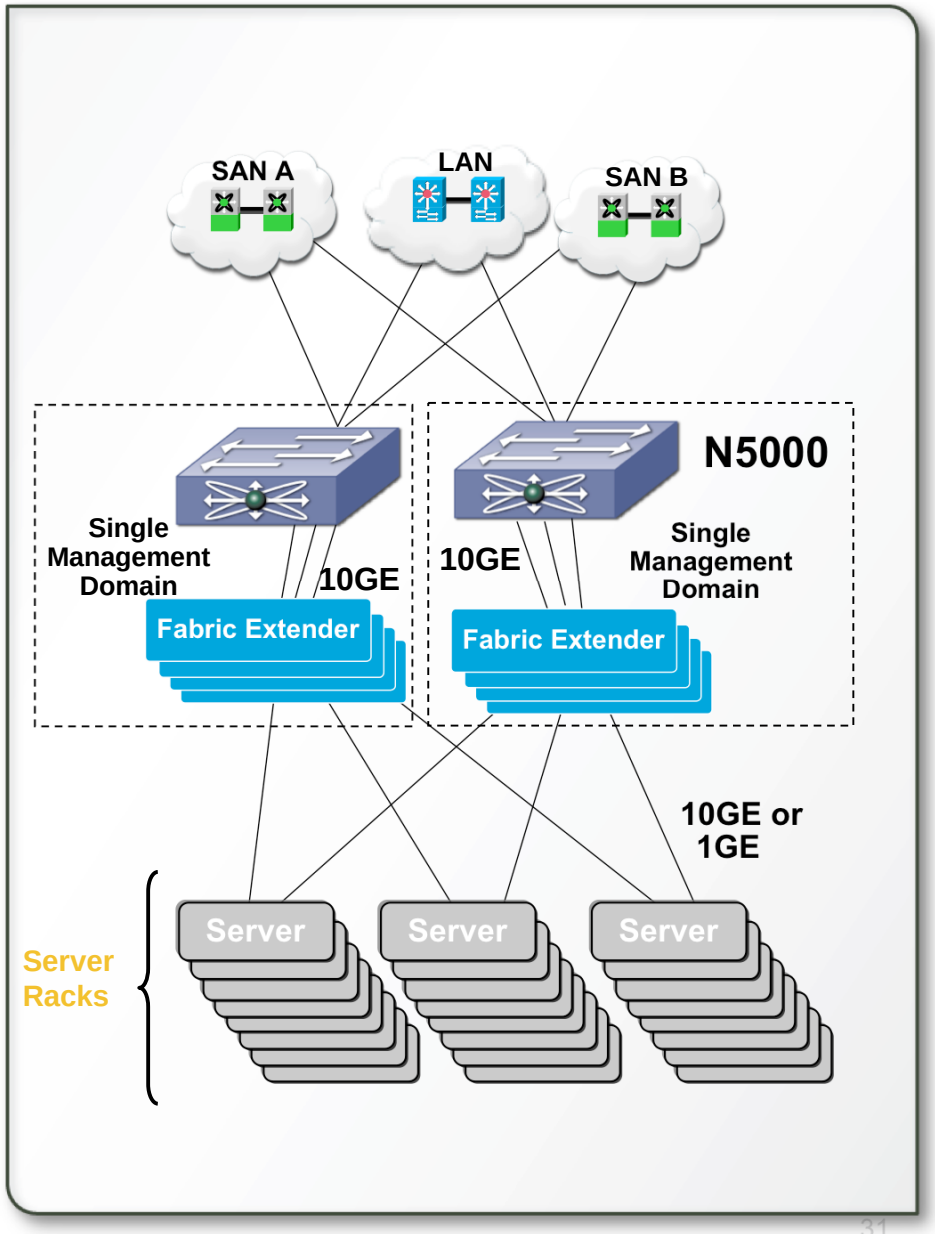

Nexus 2000 Fabric Extender

Network Topology – Physical vs. Logical

Physical Topology

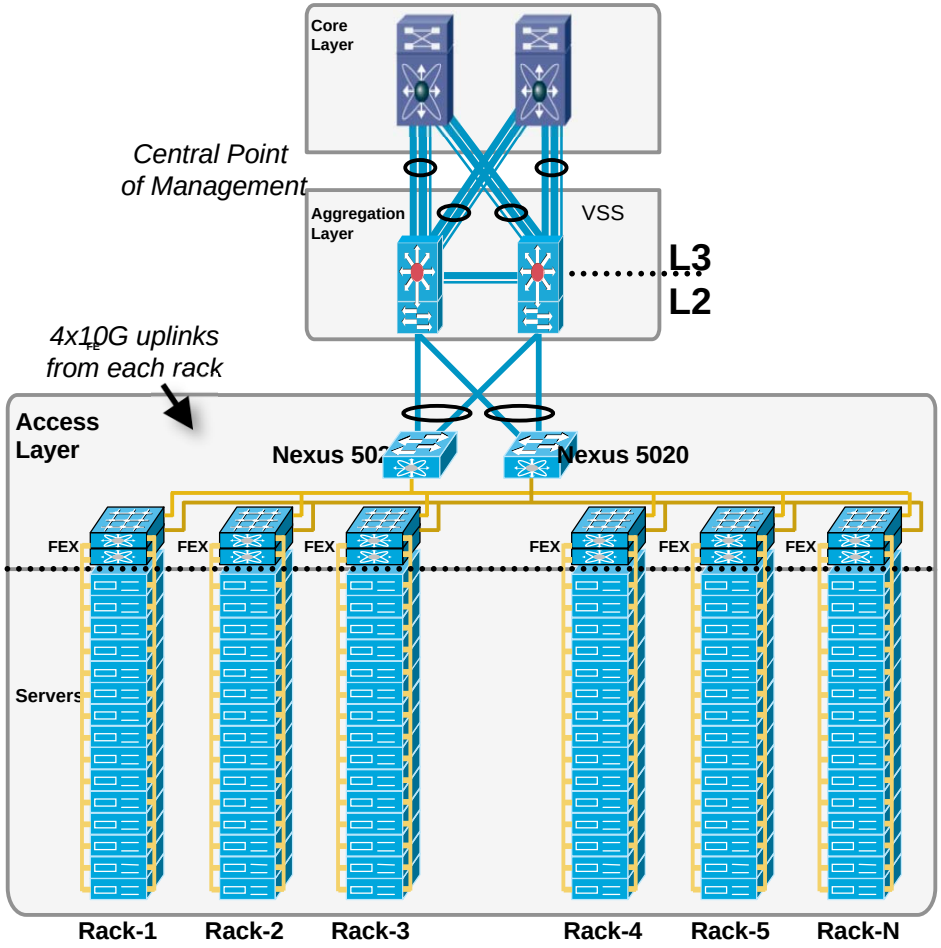

Logical Topology

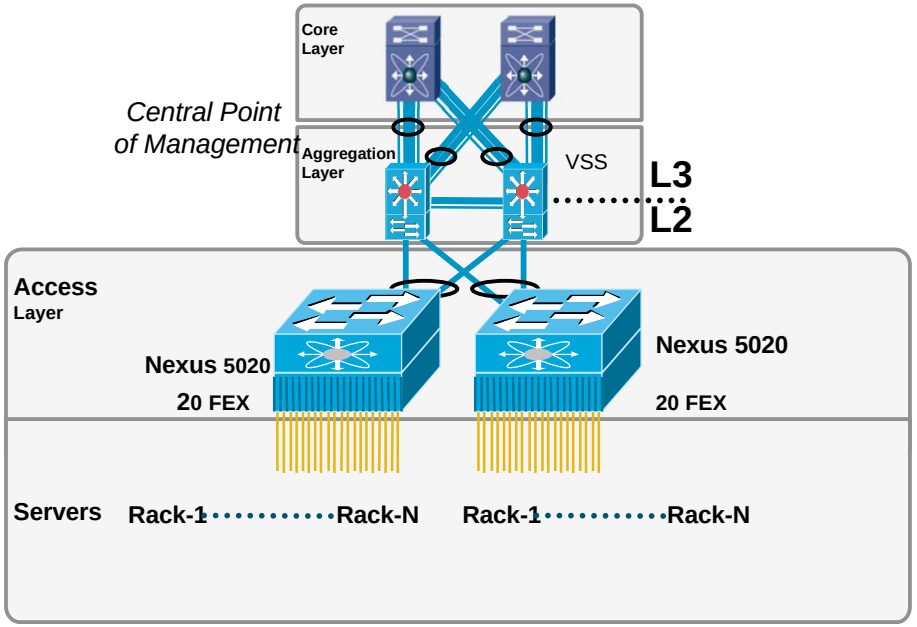

Nexus 5000 and Fabric Extender are one logical access switch in the layer 2 topology ToR flexibility ***without*** a larger STP topology

Fabric Extender
Design Option 1 – Dual NIC

- Dual NIC (redundant NIC) designs can leverage server side redundancy
- A common configuration for VMWare servers
- Two vSwitches configured with redundant NIC's
  - vSwitch 1 - Virtual Machine port groups
  - vSwitch 2 – VMkernel and service console port groups
- Common practice is to assign specific NIC's to specific traffic types

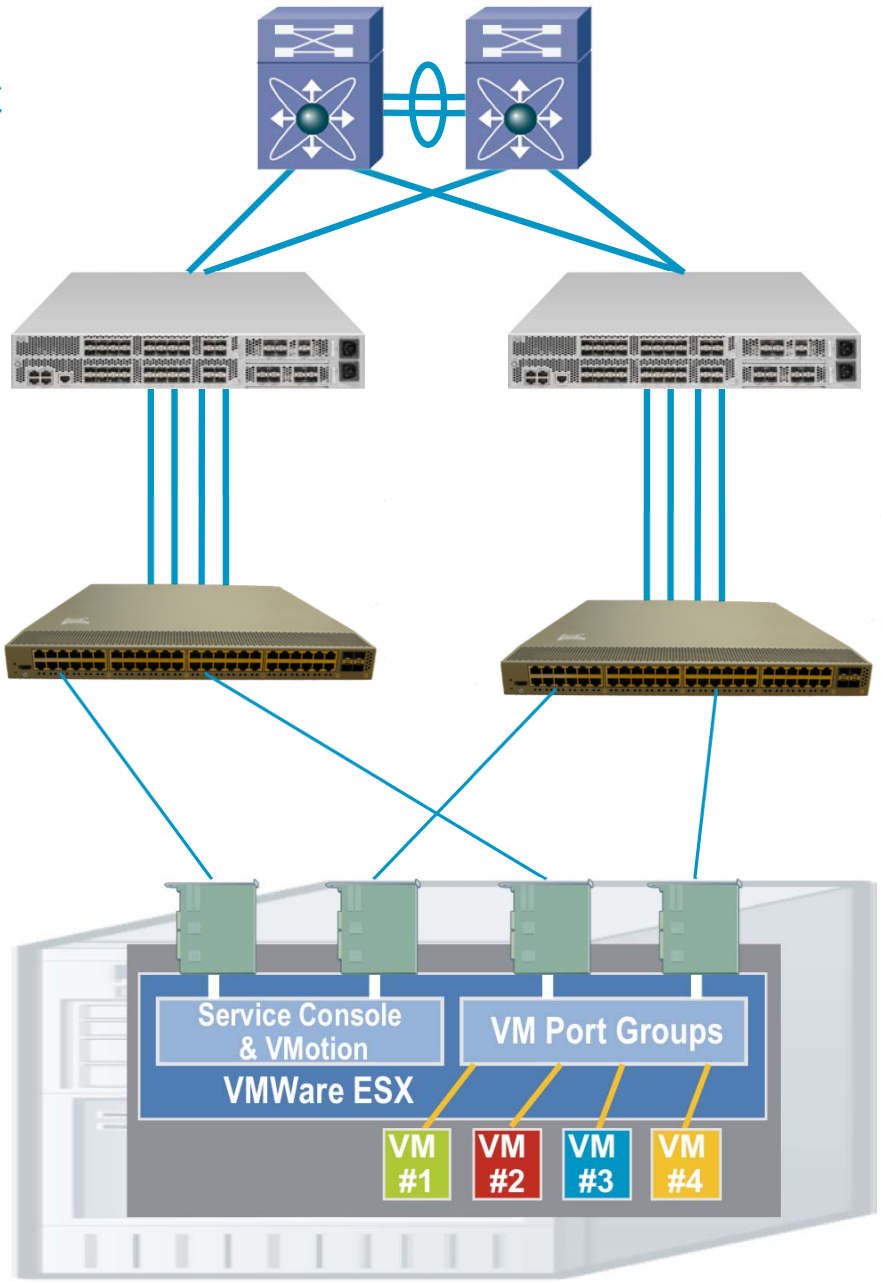

Fabric Extender
Design Option 2 – Single NIC

- Single NIC (non-redundant NIC) designs usually are used:
  - For non critical systems
  - As compute nodes in a larger application architecture leveraging ACE and application level redundancy
- All upstream traffic is multiplexed over a single NIC
- Note: Dual NIC servers often leverage ACE and application redundancy in addition to the physical NIC redundancy

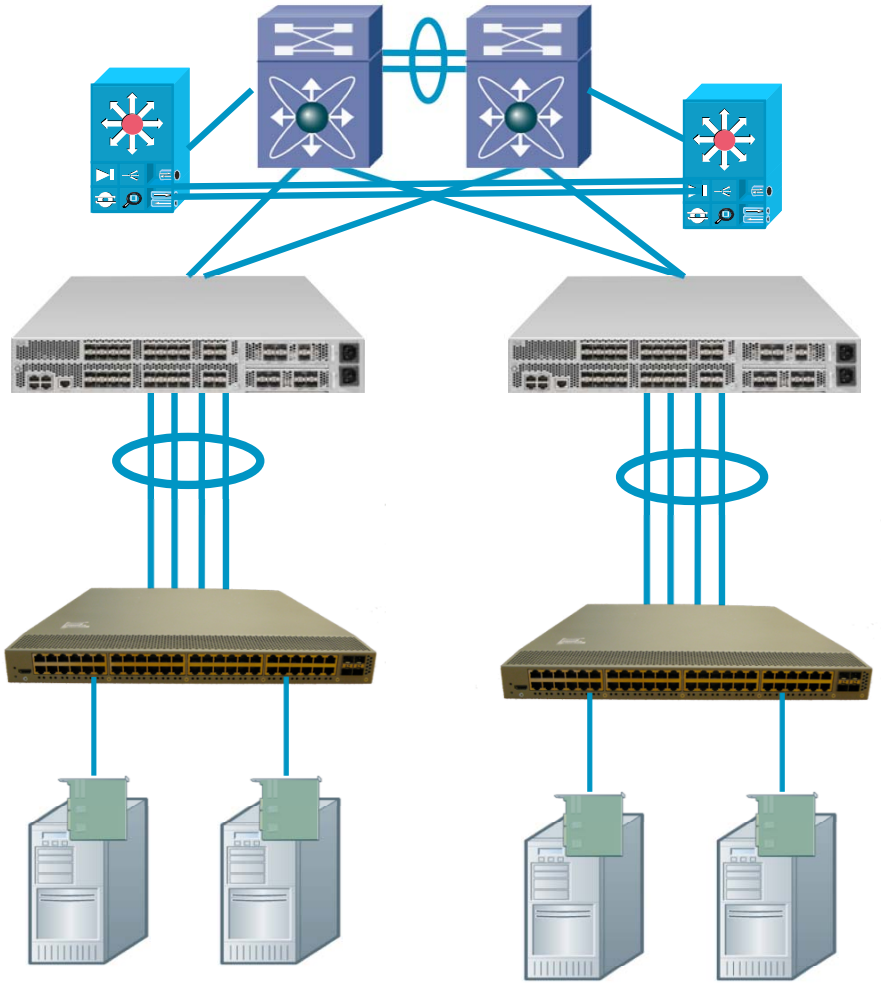

Fabric Extender
MCEC and FEX uplink

- Design Option 2 - Nexus 5000 will support a vPC uplink configuration from the Fabric Extender
- Multiple uplinks from FEX to Nexus 5000 provide data plane redundancy
- MCEC provides redundancy for the control plane
- Two elements of the MCEC configuration
  - vPC which provides the multi-chassis capabilities
  - Physical Port Channel on each of the individual Nexus 5000

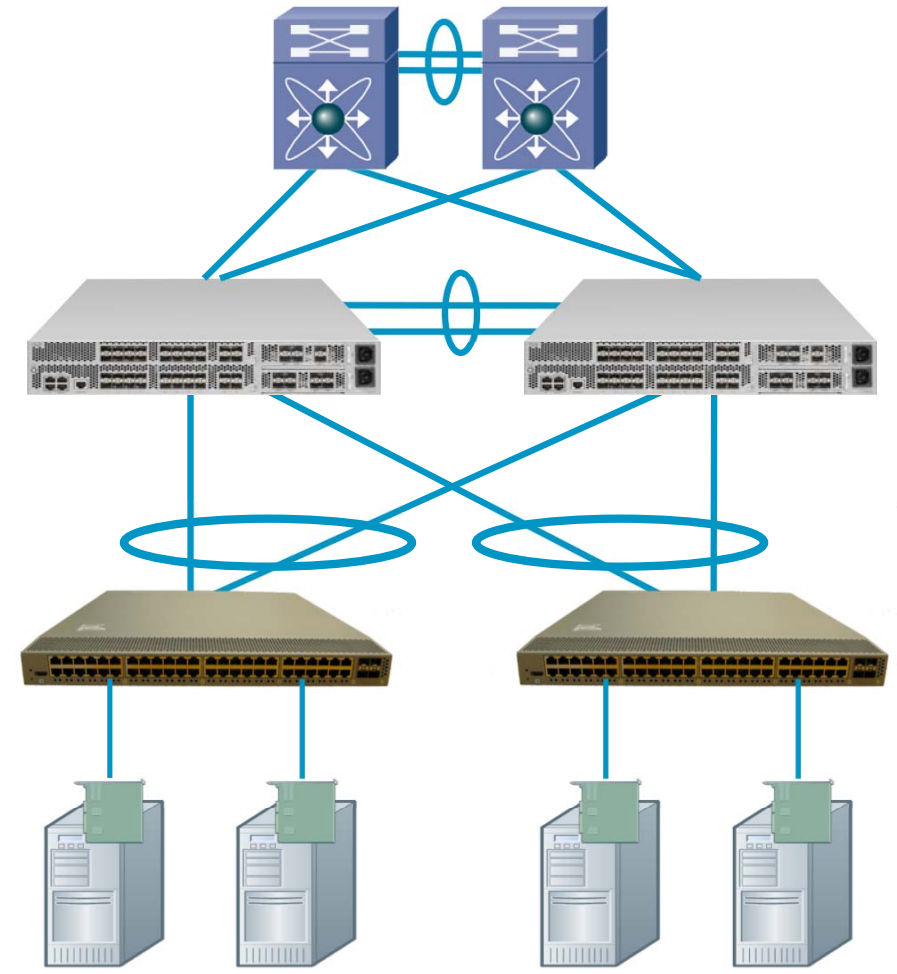

In this example FEX is configured with two 10GE vPC uplinks

Fabric Extender
MCEC and FEX uplink

- The Fabric Extender MCEC vPC configuration can either be built from
  - Two individual 10GE links
  - Two physical port channels
- NX-OS can theoretically support up to 768 virtual port channels (vPC's) per pair (what will be supported at FCS still to be determined)
- Nexus 5000 will support up to 12 Ethernet physical port channels per switch

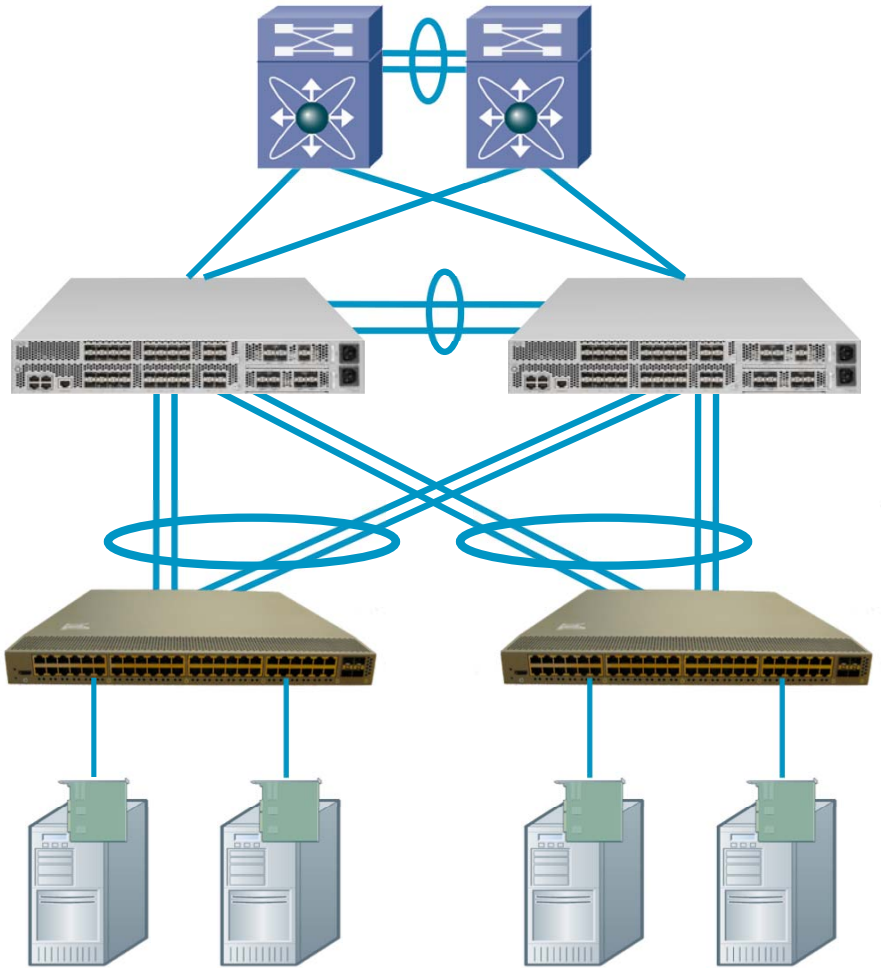

In this example FEX is configured with two 20GE 'Port Channel' vPC uplinks