Questnet
July 2010

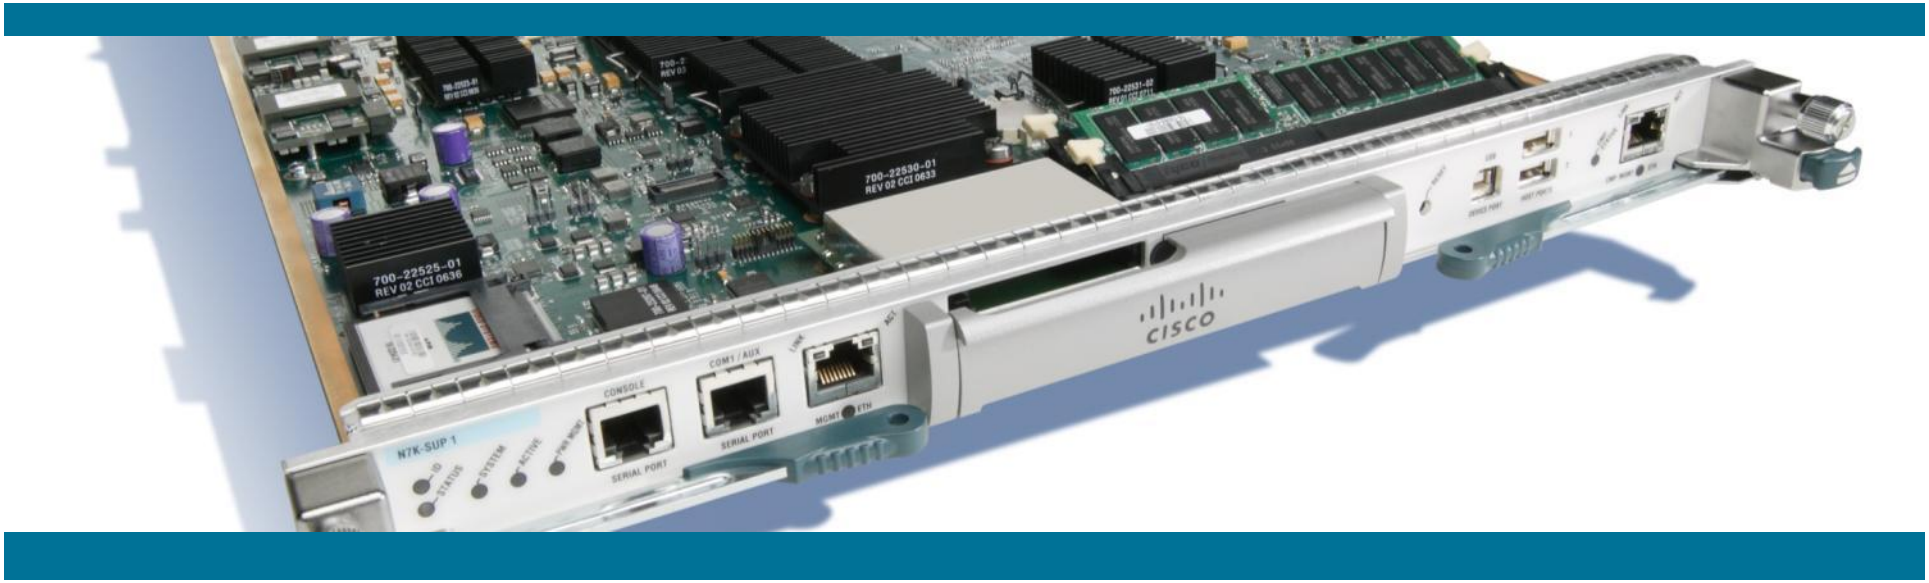

Feeds, Speeds and Futures in Ethernet Switching

The evolution of the network in a data centre and beyond

Lincoln T Dale <ltd@cisco.com>
Distinguished Engineer, Data Center Switching Technology Group

Agenda

- Overview of traditional network designs
- FEX Link
- Virtual Port Channel
- Overlay Transport Virtualization (OTV)
- Cisco FabricPath / Rbridges / TRILL
- Transceivers
- 40GbE and 100GbE
- Management

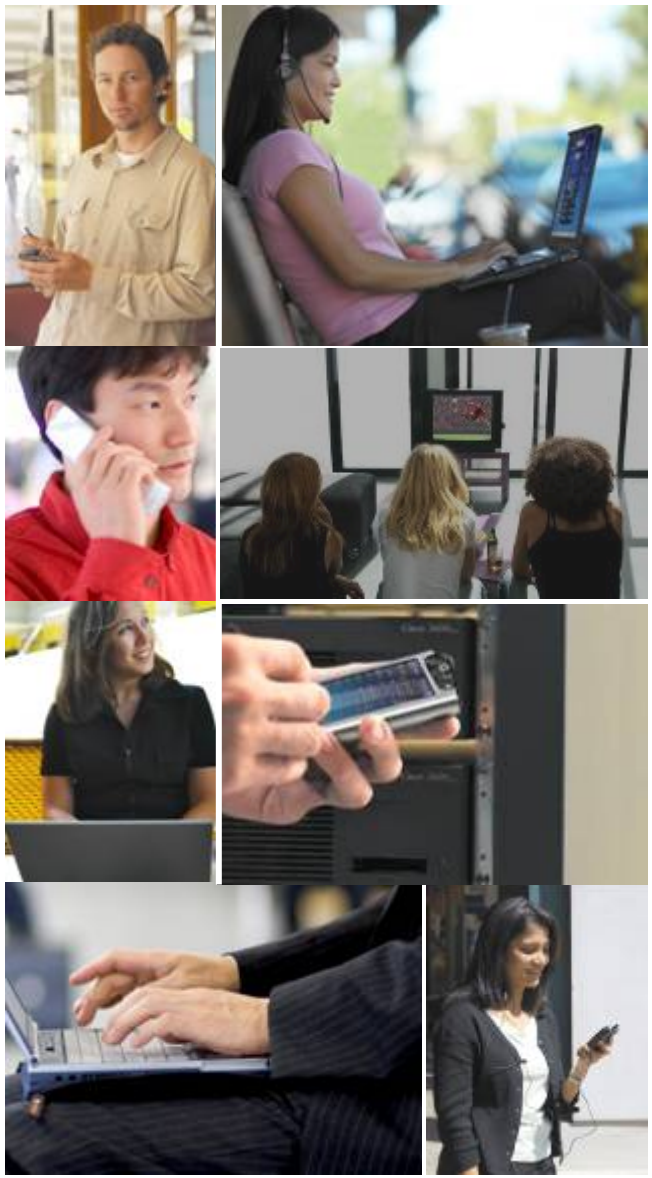

Agenda

- **Overview of a traditional network designs**
- FEX Link
- Virtual Port Channel
- Overlay Transport Virtualization (OTV)
- Cisco FabricPath / Rbridges / TRILL
- Transceivers
- 40GbE and 100GbE
- Management

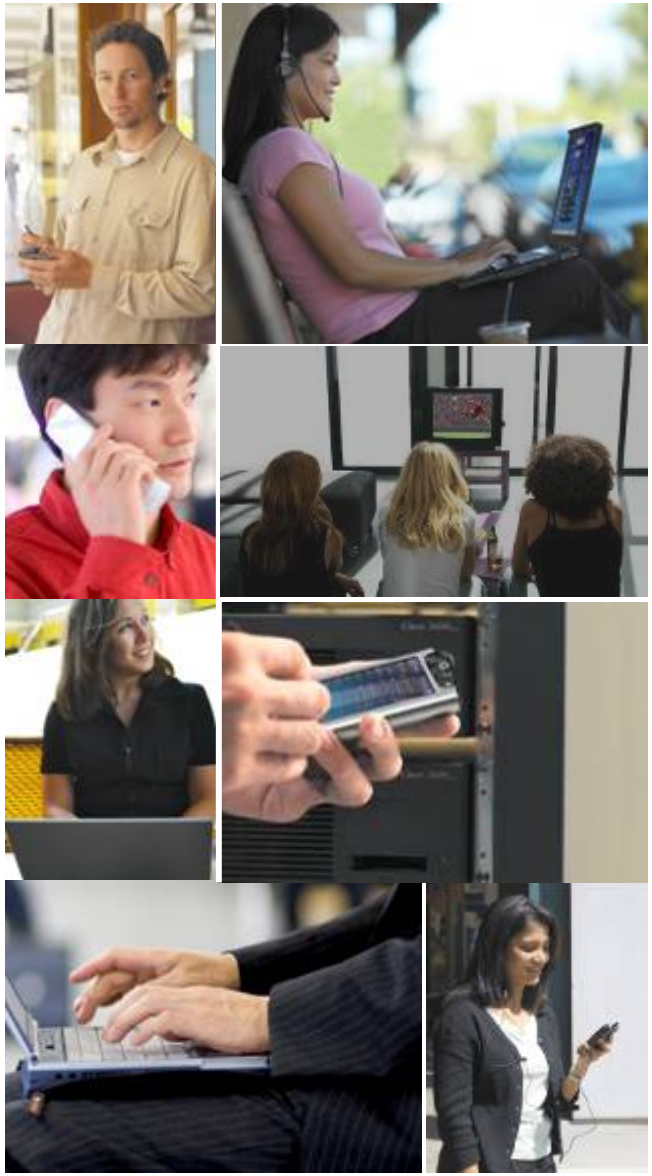

Data Center Designs
Typical 3-Tier Network Design

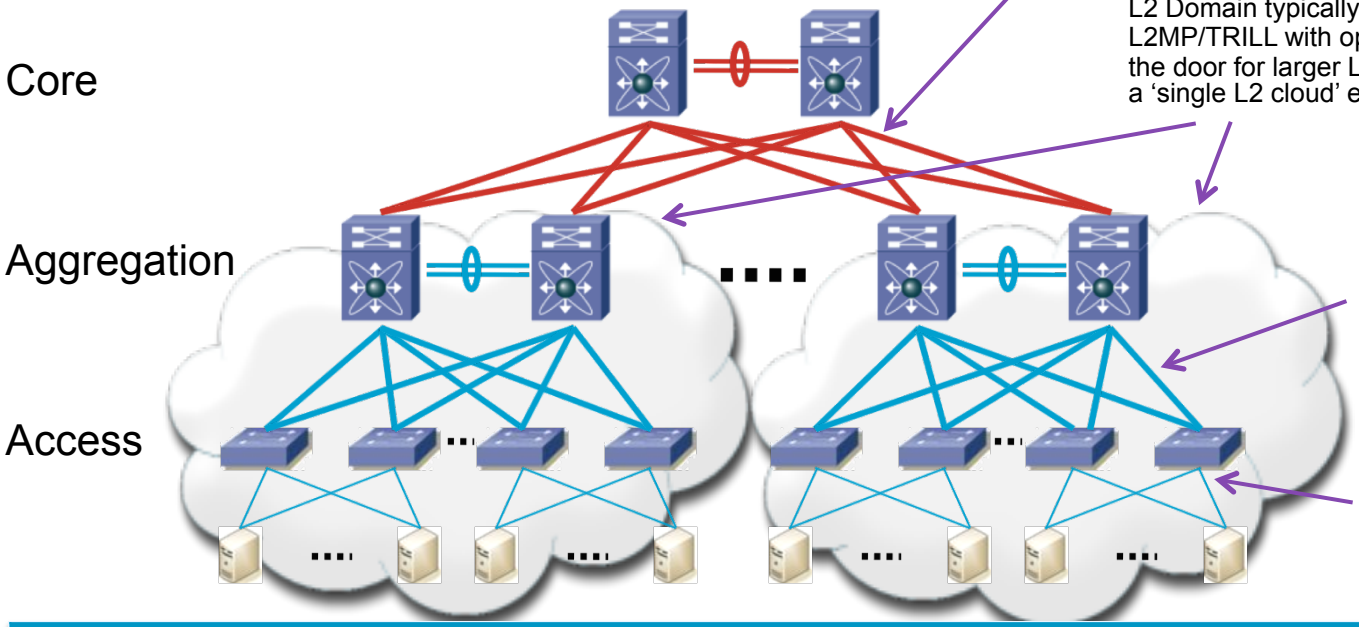

Core always L3
Core in a 3-tier design always L3. 2-tier design (collapsed core/agg) effectively same as 3-tier except with smaller maximum scalability

Scope of L2 Domain
L2 Domain typically remains within a 'pod'. L2MP/TRILL with optimized MAC learning opens the door for larger L2 pods but not necessarily a 'single L2 cloud' except in smaller designs

Classical Ethernet today, L2MP/TRILL tomorrow
Best practice today is vPC for full bisectional bandwidth. Expectation this will migrate towards L2MP/TRILL over time (e.g. some switches may not support L2MP/TRILL, 'services' may remain CE etc.)

Classical Ethernet
Best practice active/active vPC to A/A 'NIC teaming' on server

Layer	Switch Type	Port Speed	Configuration	Oversubscription	Other
Core	Modular	10GE	Layer 3 only	Low to medium	Campus hand off
Aggregation (Distribution)	Modular	10GE	L2/L3 boundary	Medium to high	Services (optional)
Access	Fixed or Modular	GE/10GE	Layer 2 only	Medium to high	ToR, MoR, Blade Switch

Table values are considered "typical" for a green field deployment

Data Center Designs

Larger port count switches enables collapsed Core +Agg for many networks with smaller requirements

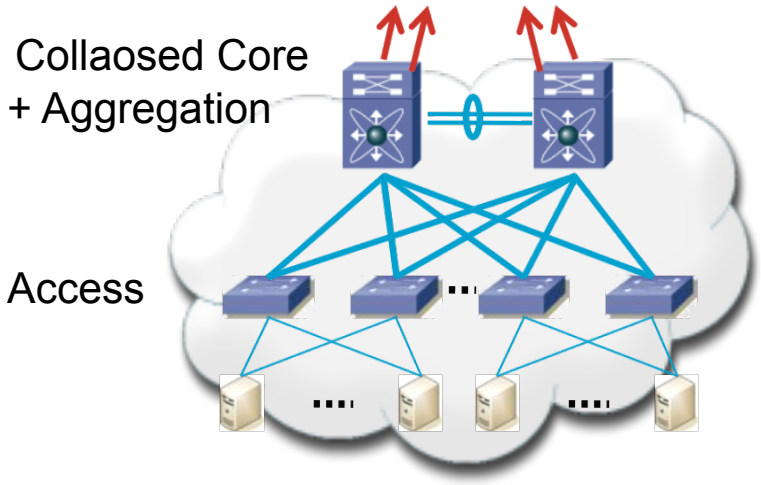

Layer	Switch Type	Port Speed	Configuration	Oversubscription	Other
Collapsed Core + Aggregation	Modular	10GE	L3 Northbound L2 Southbound	Low to medium	Campus hand off, Services (optional)
Access	Fixed or Modular	GE/10GE	Layer 2 only	Medium to high	ToR, MoR, Blade Switch

Spanning Tree Standards and Features
IEEE 802.1D, IEEE 802.1s, IEEE 802.1w

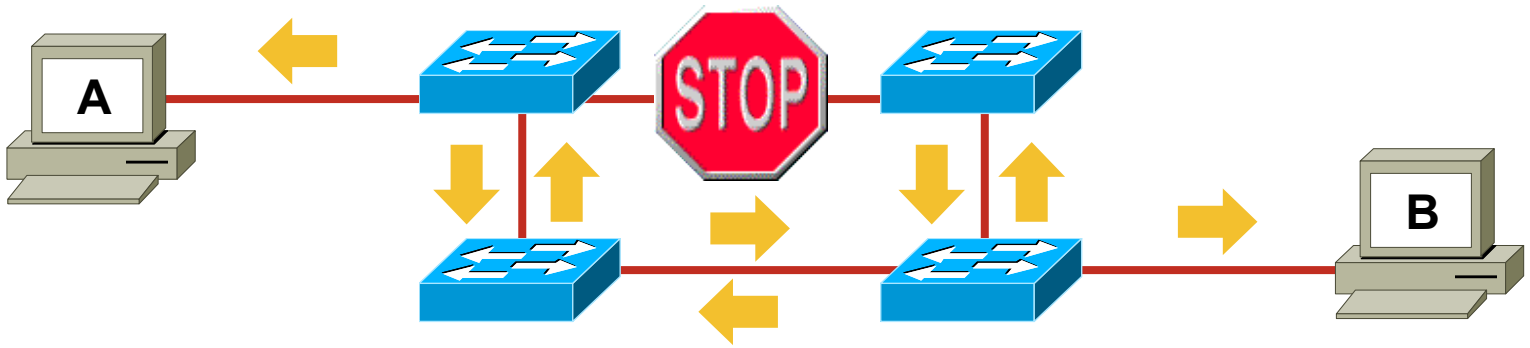

- **802.1D/1998**: legacy standard for bridging and Spanning Tree (STP)
- **802.1D/2004**: updated bridging and STP standard; includes 802.1s, 802.1t, and 802.1w
- **802.1s**: Multiple Spanning Tree Protocol (MSTP)—maps multiple VLANs into the same Spanning Tree instance
- **802.1t**: MAC address reduction/extended system ID—moves some BPDU bits to high-numbered VLANs from the priority field, which constrains the possible values for bridge priority; unique “MAC” per chassis not port
- **802.1w**: Rapid Spanning Tree Protocol (RSTP)—improved convergence over 1998 STP by adding roles to ports and enhancing BPDU exchanges
- **Cisco Features**: Per VLAN Spanning Tree (PVST), PVST+, UpLinkFast, BackboneFast, BPDU Guard, RootGuard, LoopGuard, Bridge Assurance, UDLD

Agenda

- Overview of traditional network designs
- **FEX Link**
- Virtual Port Channel
- Overlay Transport Virtualization (OTV)
- Cisco FabricPath / Rbridges / TRILL
- Transceivers
- 40GbE and 100GbE
- Management

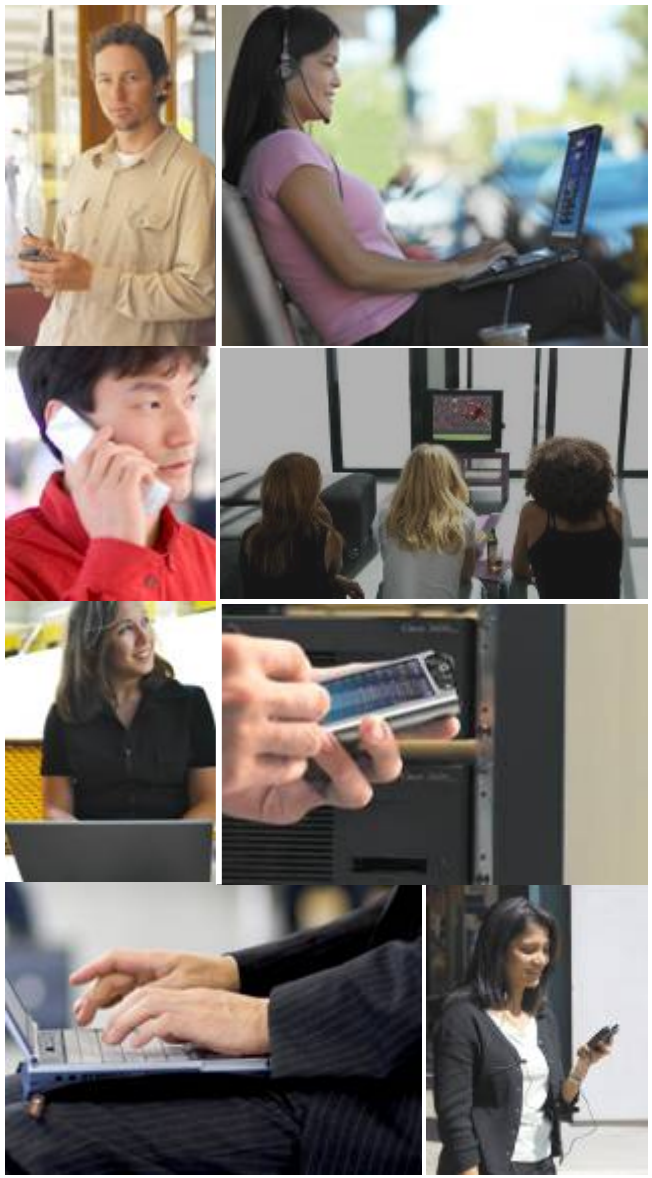

Access Layer Options

Animated Slide!

Top of Rack (ToR)

- Typically 1-RU servers
- 1-2 GE LOMs
- Mostly 1 sometimes 2 ToR switches
- Copper cabling stays within rack
- Low copper density in ToR
- Higher chance of East-West traffic hitting aggregation layer
- Drives higher STP logical port count for aggregation layer
- Denser server count

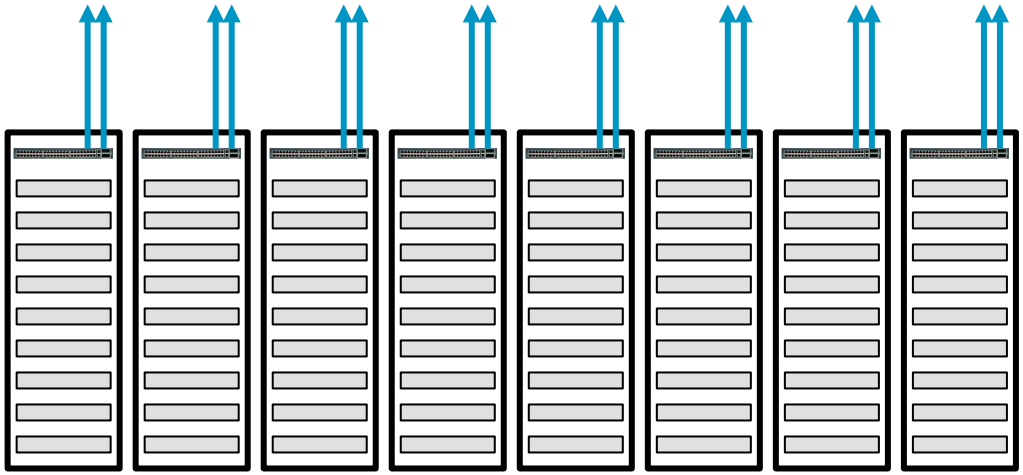

Middle of Row (MoR aka EoR)

- 1-RU or multi-RU servers
- Multiple GE or 10GE NICs
- Horizontal copper cabling for servers
- High copper cable density in MoR
- Larger portion of East-West traffic (Server-to-Server) stays in access
- Larger subnets → less address waste
- Keeps agg. STP logical port count low (more EtherChannels, less trunk ports)
- Lower # of network devices to manage

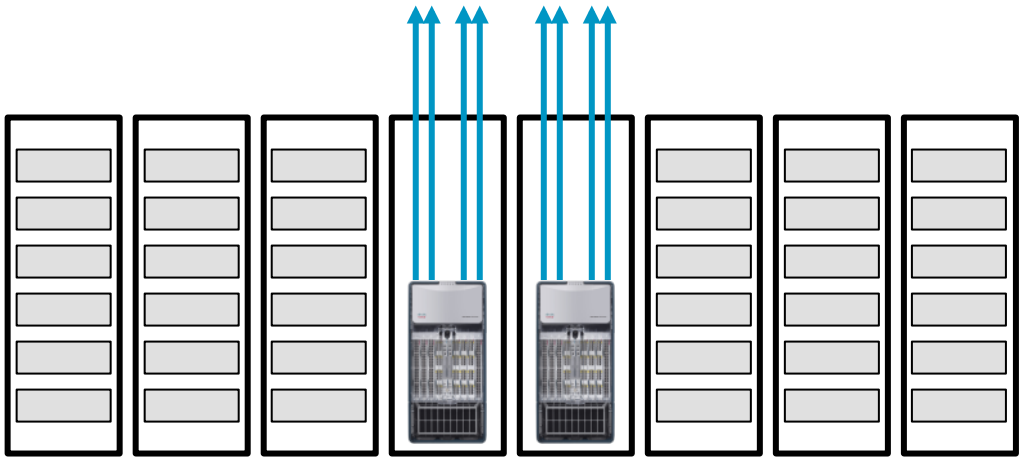

Animated Slide!

Access Layer Options

Hybrid Top of Rack (ToR) AND Middle of Row with FEX Combination!

- Copper cabling stays within rack
- Low copper density in ToR
- Lowest # of network devices to manage
- Denser server count

Combines benefits of Top of Rack AND End of Row (EoR) network architectures

- ***Physically*** **resides on the top of each server rack**
- ***Logically*** **acts like an end of access row device**
- ***Scales*** **in a manner that enables collapsing of Agg + Access layers in many networks**

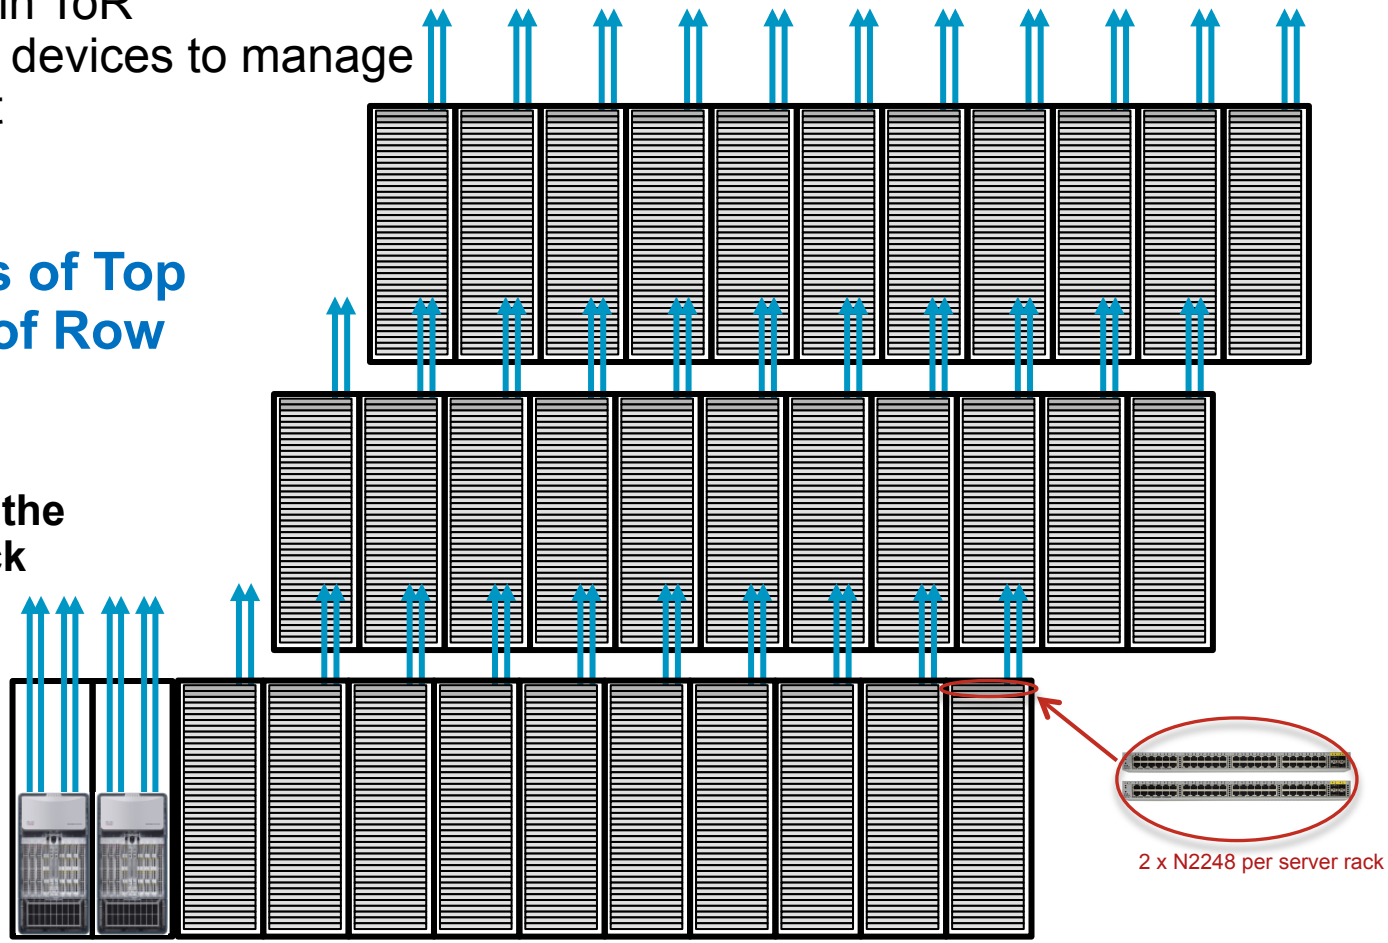

Nexus 2000 FEX Benefits

Combines benefits of Top of Rack (ToR) and End of Row (EoR) network architectures

- **Physically resides on the top of each server rack, Logically acts like an end of access row device**

- **Reduces cable runs**
  Majority of physical cabling is within the rack, ≤5 foot cable
- **Reduce management points in the network**
  1,500 1G port HA network with 48-port ToR access switches: 34 management points
  1,500 1G port HA network with N7K + FEX: 2 management points
- ***Easier* to ensure feature consistency across hundreds or thousands of server ports**
- **Investment protection**
  - **FEX can be used across multiple Cisco Nexus switching platforms**
- **Achieves all the benefits of higher port/module density (e.g. MRJ21) without the associated non-standard cabling & investment protection issues**

Data Center Designs
Typical 3-Tier Network Design

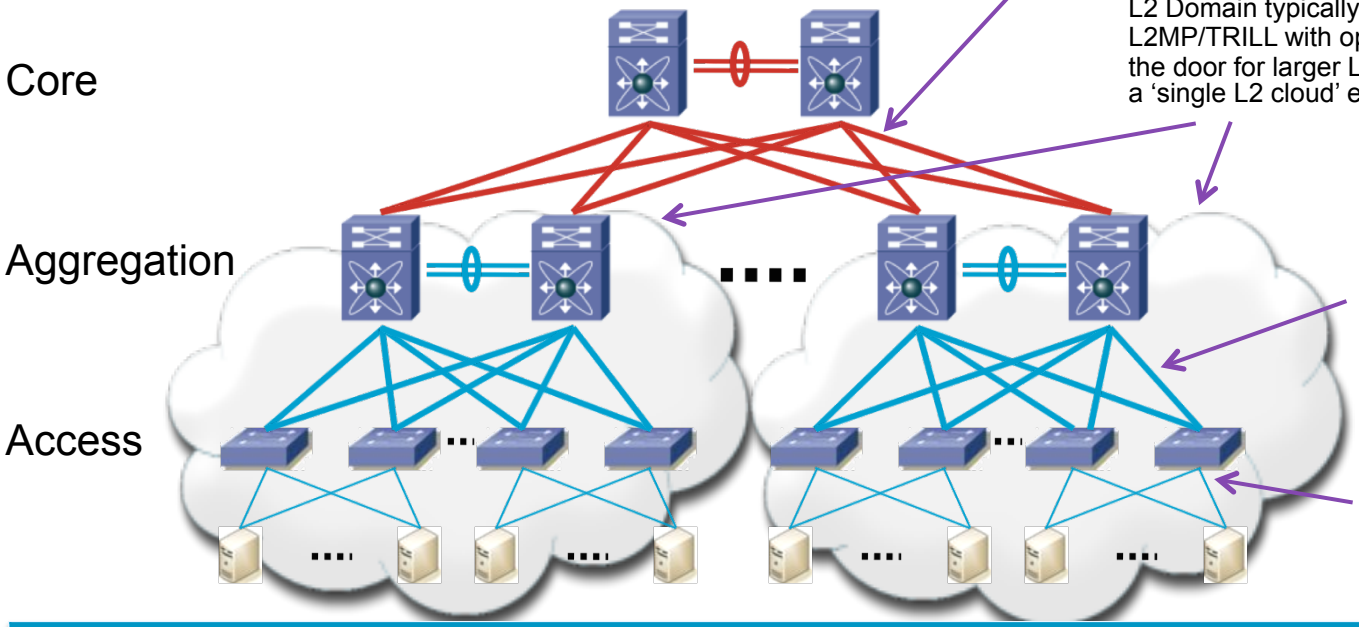

Core always L3
Core in a 3-tier design always L3. 2-tier design (collapsed core/agg) effectively same as 3-tier except with smaller maximum scalability

Scope of L2 Domain
L2 Domain typically remains within a 'pod'. L2MP/TRILL with optimized MAC learning opens the door for larger L2 pods but not necessarily a 'single L2 cloud' except in smaller designs

Classical Ethernet today, L2MP/TRILL tomorrow
Best practice today is vPC for full bisectional bandwidth. Expectation this will migrate towards L2MP/TRILL over time (e.g. some switches may not support L2MP/TRILL, 'services' may remain CE etc.)

Classical Ethernet
Best practice active/active vPC to A/A 'NIC teaming' on server

Layer	Switch Type	Port Speed	Configuration	Oversubscription	Other
Core	Modular	10GE	Layer 3 only	Low to medium	Campus hand off
Aggregation (Distribution)	Modular	10GE	L2/L3 boundary	Medium to high	Services (optional)
Access	Fixed or Modular	GE/10GE	Layer 2 only	Medium to high	ToR, MoR, Blade Switch

Table values are considered "typical" for a green field deployment

Data Center Designs enabled with N7K+FEX

FEX on N7K enables collapsed Agg+Access

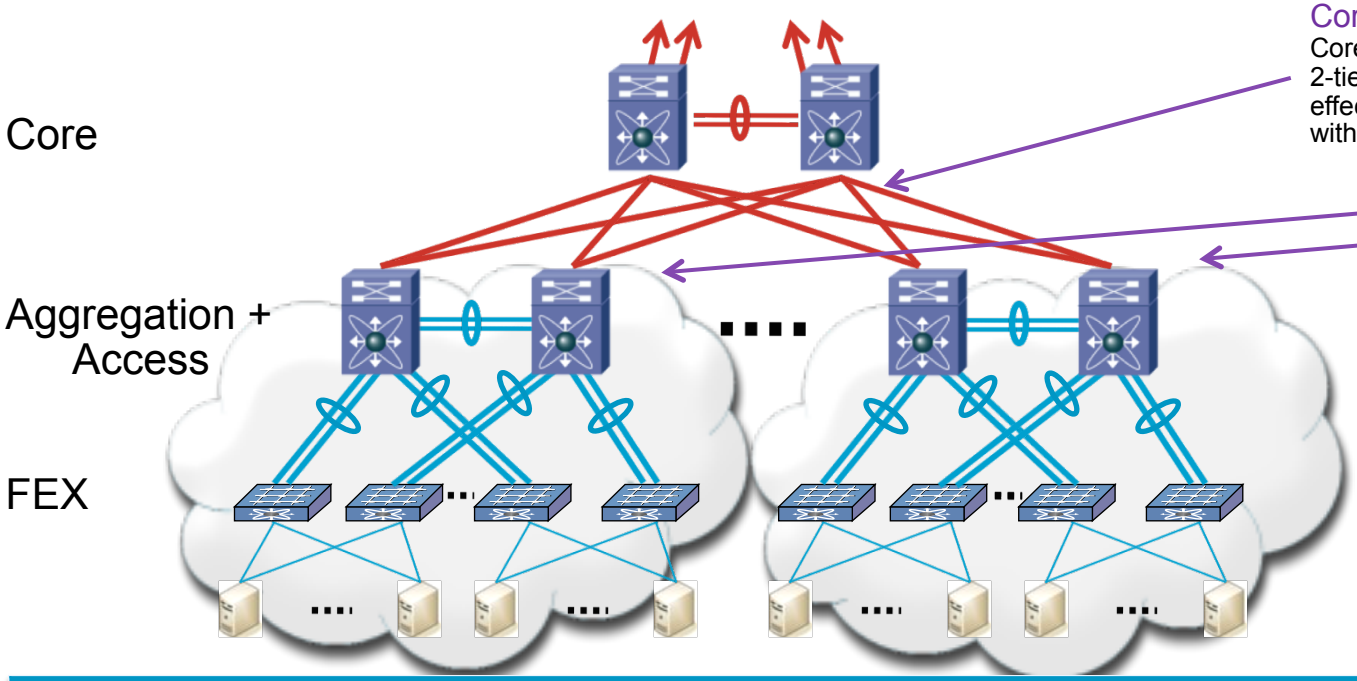

Core always L3

Core in a 3-tier design always L3. 2-tier design (collapsed core/agg) effectively same as 3-tier except with smaller maximum scalability

Scope of L2 Domain

L2 Domain typically remains within a 'pod. Rather than using switches for access the access-layer switching is collapsed into Agg layer here and a 'pod' potentially becomes 2 switches total.
Larger designs in future may make use of L2MP/TRILL to scale this significantly.

Layer	Switch Type	Port Speed	Configuration	Oversubscription	Other
Core	Modular	10GE	Layer 3 only	Low to medium	Campus hand off
Aggregation + Access	Modular	10GE	L2/L3 boundary	Medium to high	Services (optional)

Data Center Designs

Large port count switches enables collapsed Core +Agg for many networks…

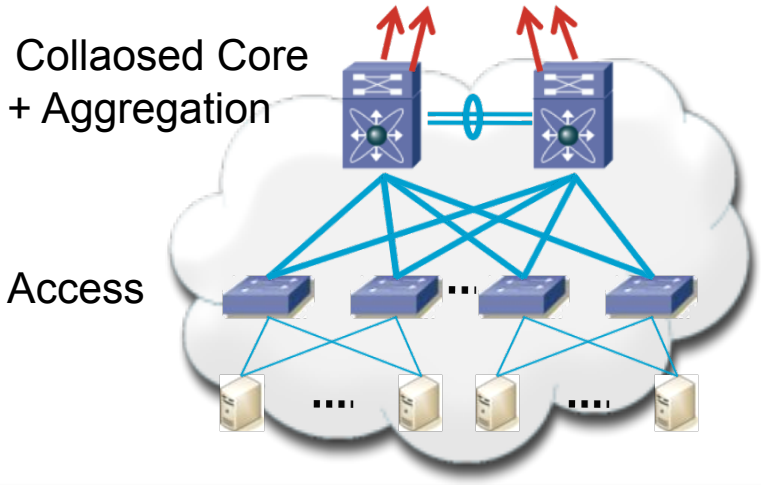

Layer	Switch Type	Port Speed	Configuration	Oversubscription	Other
Collapsed Core + Aggregation	Modular	10GE	L3 Northbound L2 Southbound	Low to medium	Campus hand off, Services (optional)
Access	Fixed or Modular	GE/10GE	Layer 2 only	Medium to high	ToR, MoR, Blade Switch

Data Center Designs enabled with N7K+FEX
(or even Collapsed Core+Agg+Access)

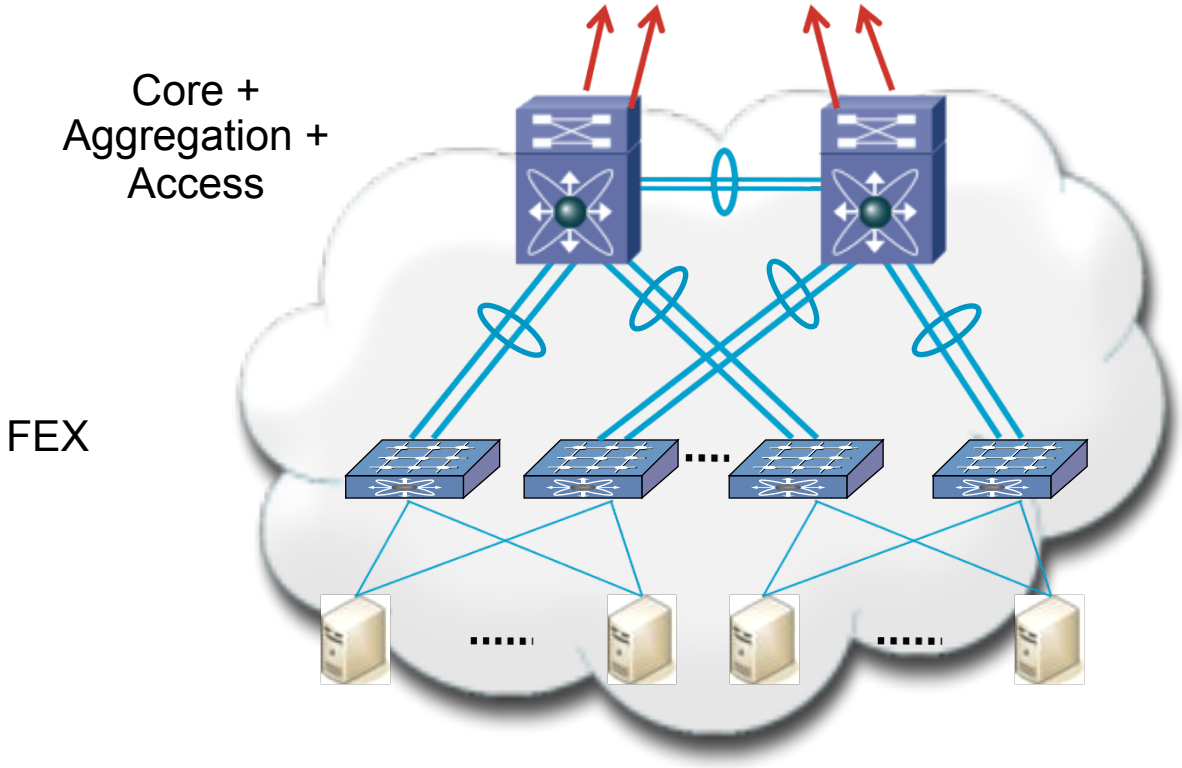

Agenda

- Overview of traditional network designs
- FEX Link
- **Virtual Port Channel**
- Overlay Transport Virtualization (OTV)
- Cisco FabricPath / Rbridges / TRILL
- Transceivers
- 40GbE and 100GbE
- Management

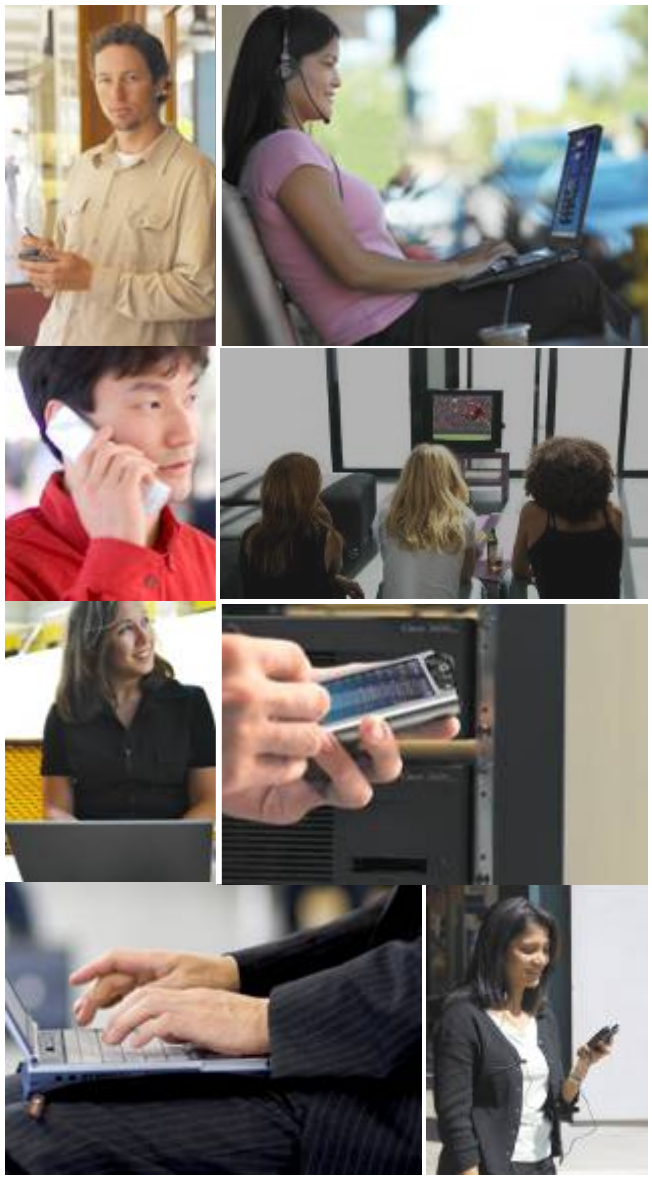

Virtual Port Channel (vPC)

- Allow a single device to use a port channel across two upstream switches
- Eliminate STP blocked ports
- Uses all available uplink bandwidth
- Dual-homed server operate in active-active mode
- Provide fast convergence upon link/device failure
- Reduce CAPEX and OPEX

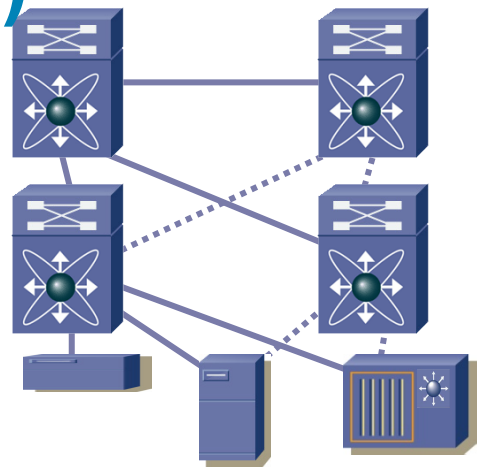

Logical Topology without vPC

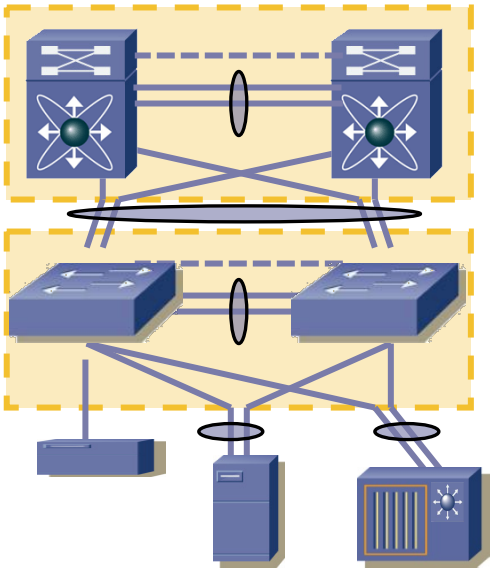

Logical Topology with vPC

Virtual Port Channel (vPC)

Following steps are needed to build a vPC
(Order **does Matter**!)

- Define domains*
- Establish Peer Keepalive connectivity
- Create a Peer link
- Reuse port-channels and Create vPCs
- *Make Sure Configurations are Consistent*

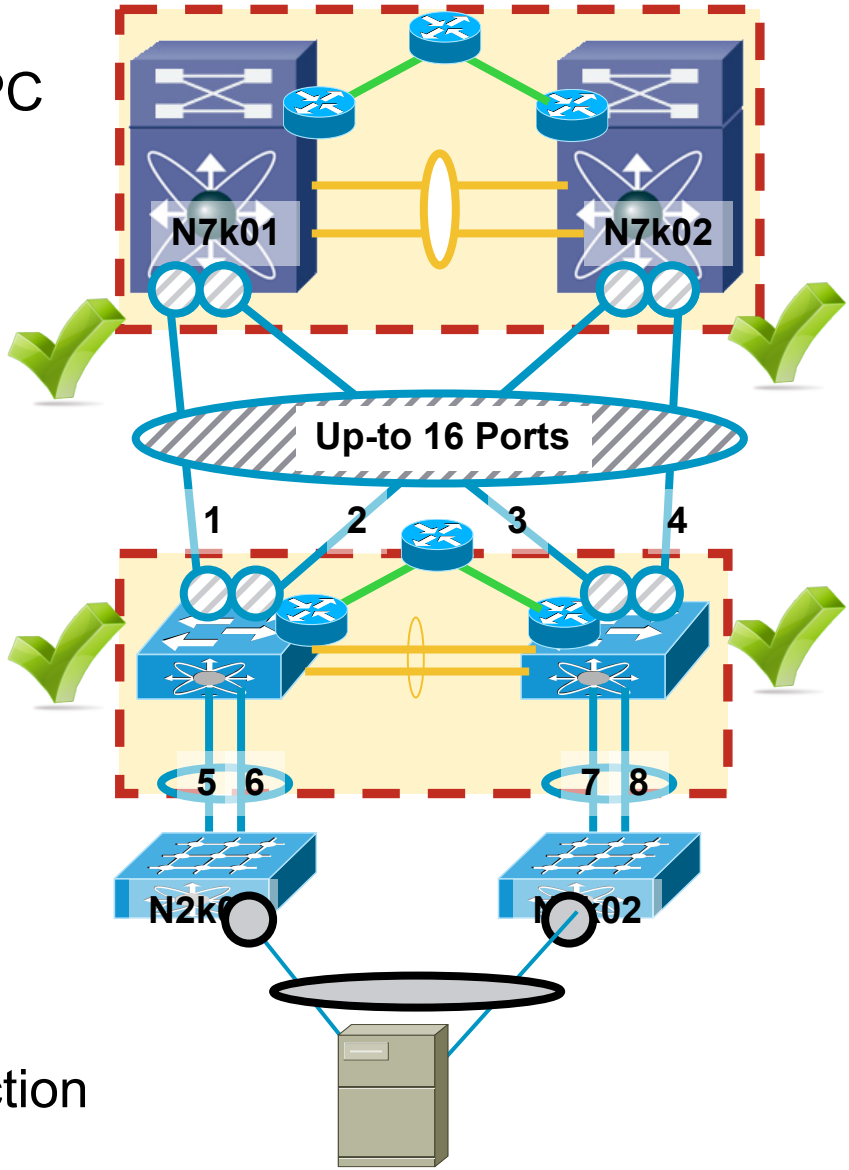

*See Configuration details in the note section

Feature Overview

vPC and VSS Comparison

Capability	vPC (Virtual Port Channels)	VSS (Virtual Switching System)
Multi-Chassis Port Channel	Yes	Yes
Loop-free Topology (no blocking ports)	Yes	Yes
STP as a “fail-safe” protocol only	Yes	Yes
Control Plane	Two Independent Nodes, both active	Single Logical Node
Support for Layer 3 port-channels	No	Yes
Switch Redundancy (sup failover)	Intra-chassis	Inter-chassis
Control Plane Protocols	Instances per Node	Single instance
Switch Configuration	Common Configs (w/ consistency checker)	Combined Configs
Maximum Physical Nodes	2	2
Non Disruptive ISSU Support	Yes	No
Inter-switch Link Hardware	All 10GE Shipping modules: 32 port 10GE SFP+, 8 port 10GE Module X2	PFC3C mode, Sup 70 10G, 6708, 6716

Agenda

- Overview of traditional network designs
- FEX Link
- Virtual Port Channel
- **Overlay Transport Virtualization (OTV)**
- Cisco FabricPath / Rbridges / TRILL
- Transceivers
- 40GbE and 100GbE
- Management

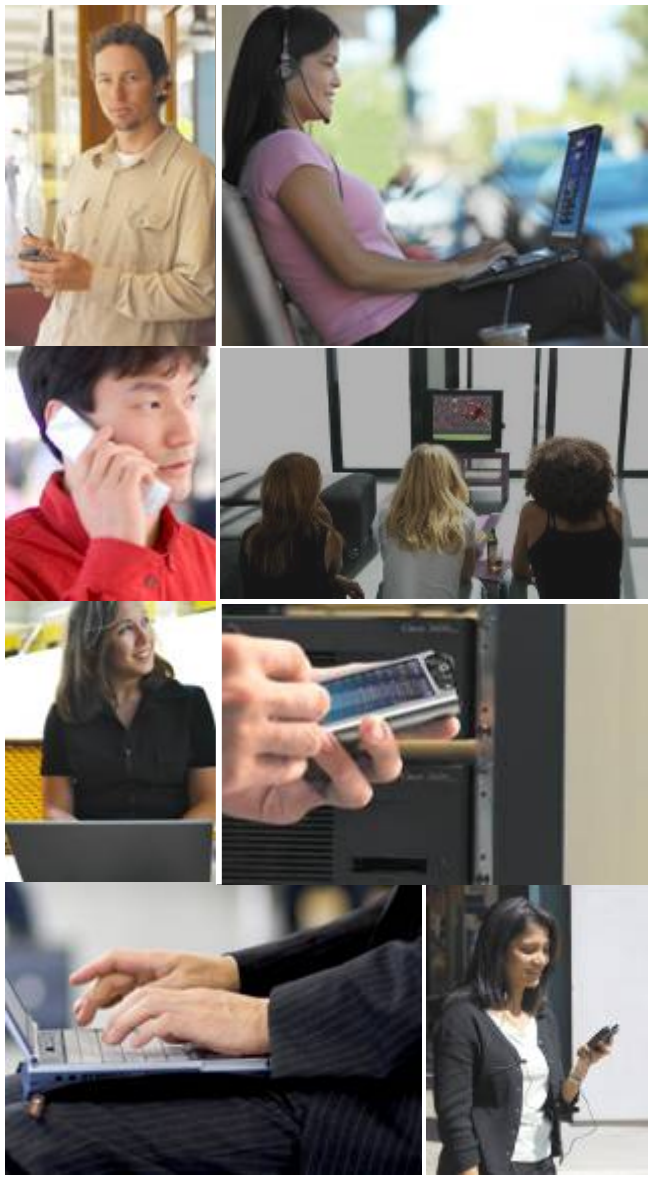

Concept - LAN extensions: Inter/intra Data Center

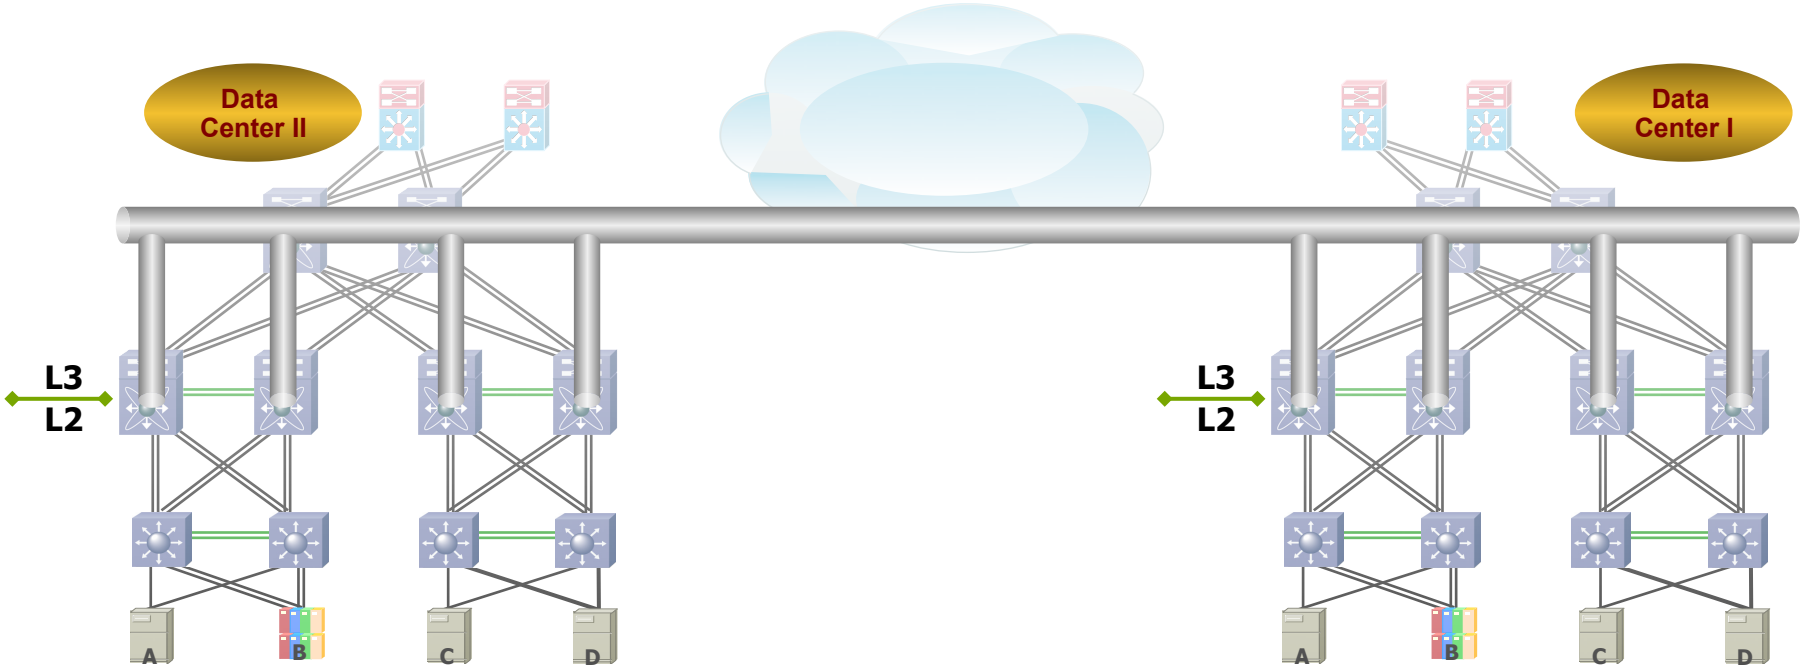

- Certain Applications require L2 connectivity among peers
  - Clusters (Veritas, MSFT)
  - vMotion
  - Home-brewed apps, non-IP sync
- Server migrations
- Disaster recovery and resiliency
- High rate encryption may require an L2 transport between sites
- Distributed Active-active DCs

Intelligent L2 Domains
POD Innovation and Interconnection

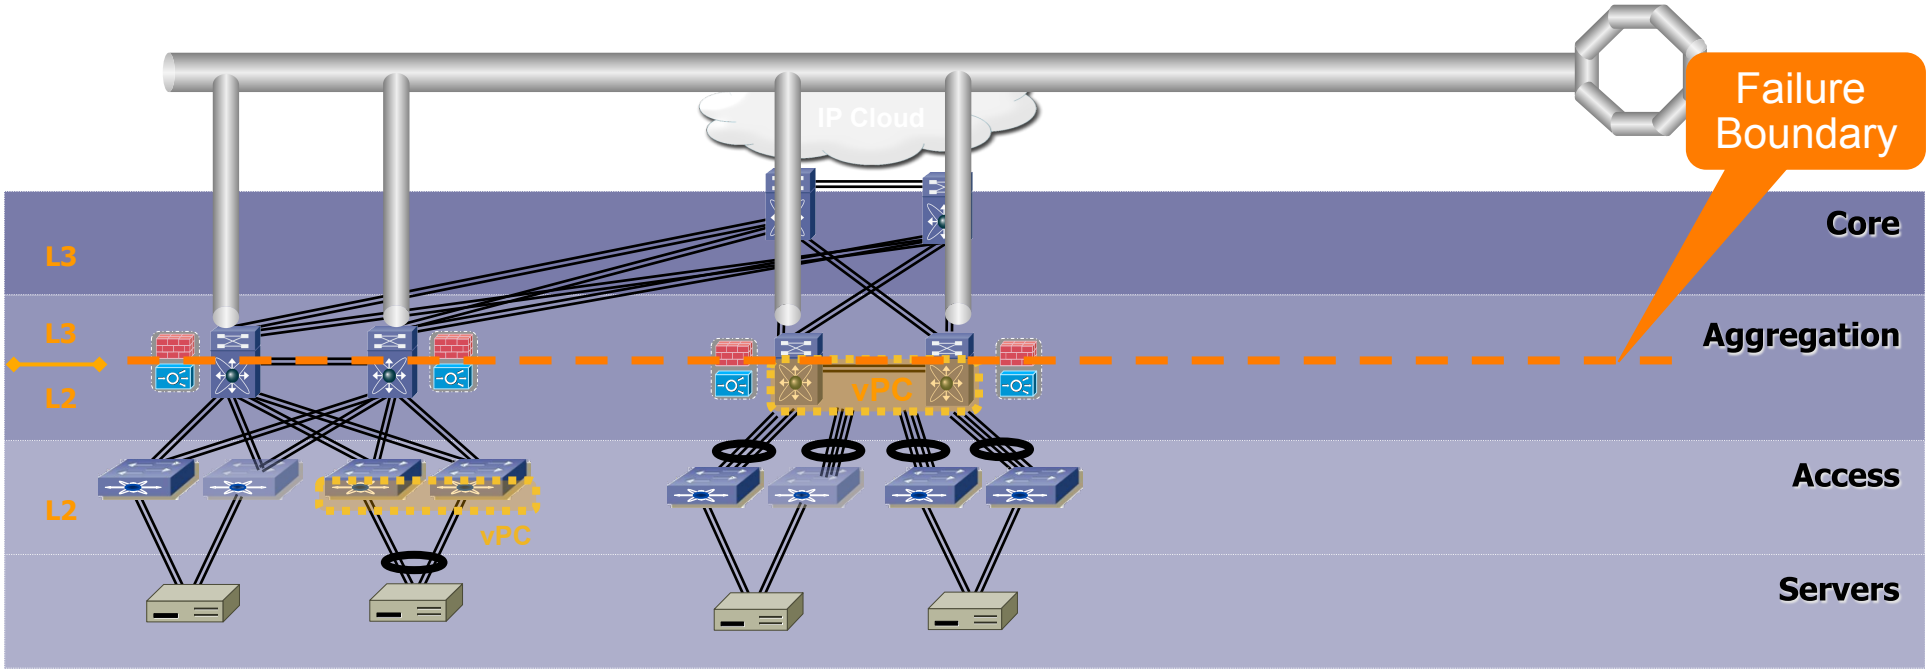

STP+

STP Enhancements

Bridge Assurance

vPC

NIC Teaming

Simplified loop-free trees

2x Multi-pathing

OTV

Inter-POD Connectivity across L3

Failure Boundary Preservation

Concept – Overlay Transport Virtualization Enterprise Class Ethernet VPNs

- IP based Ethernet (L2) VPN solution
  - MAC routing
  - IP encapsulated forwarding
- 'Smart' Ethernet Pseudo-node
- Core and Site Transparency

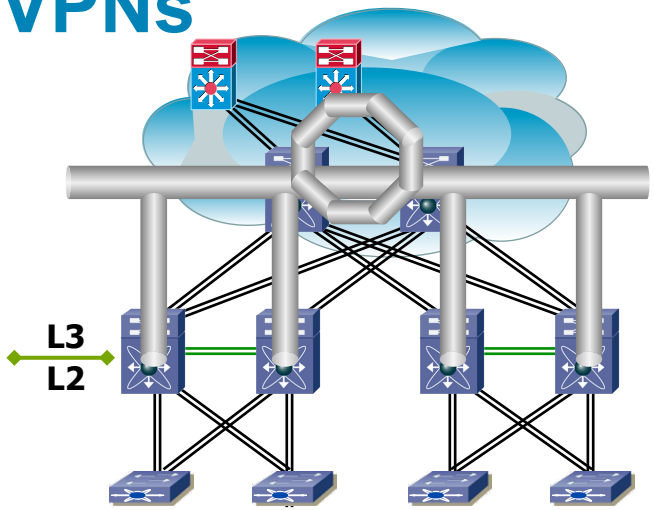

Protocol Learning

Built-in loop prevention (no STP)

Failure domain is bound

Floods/b-casts can be suppressed

Seamless adds/removes & multihoming

Full cross-sectional BW

- Equal Cost Multi-Pathing
- All-active Multi-homing

Packet Switching

Multi-point connectivity

Minimal State

- No PW state preserved

Optimal m-cast replication

Point-to-cloud model for improved manageability

Overlay Transport Virtualization – OTV Ethernet in IP Dynamic Encapsulation

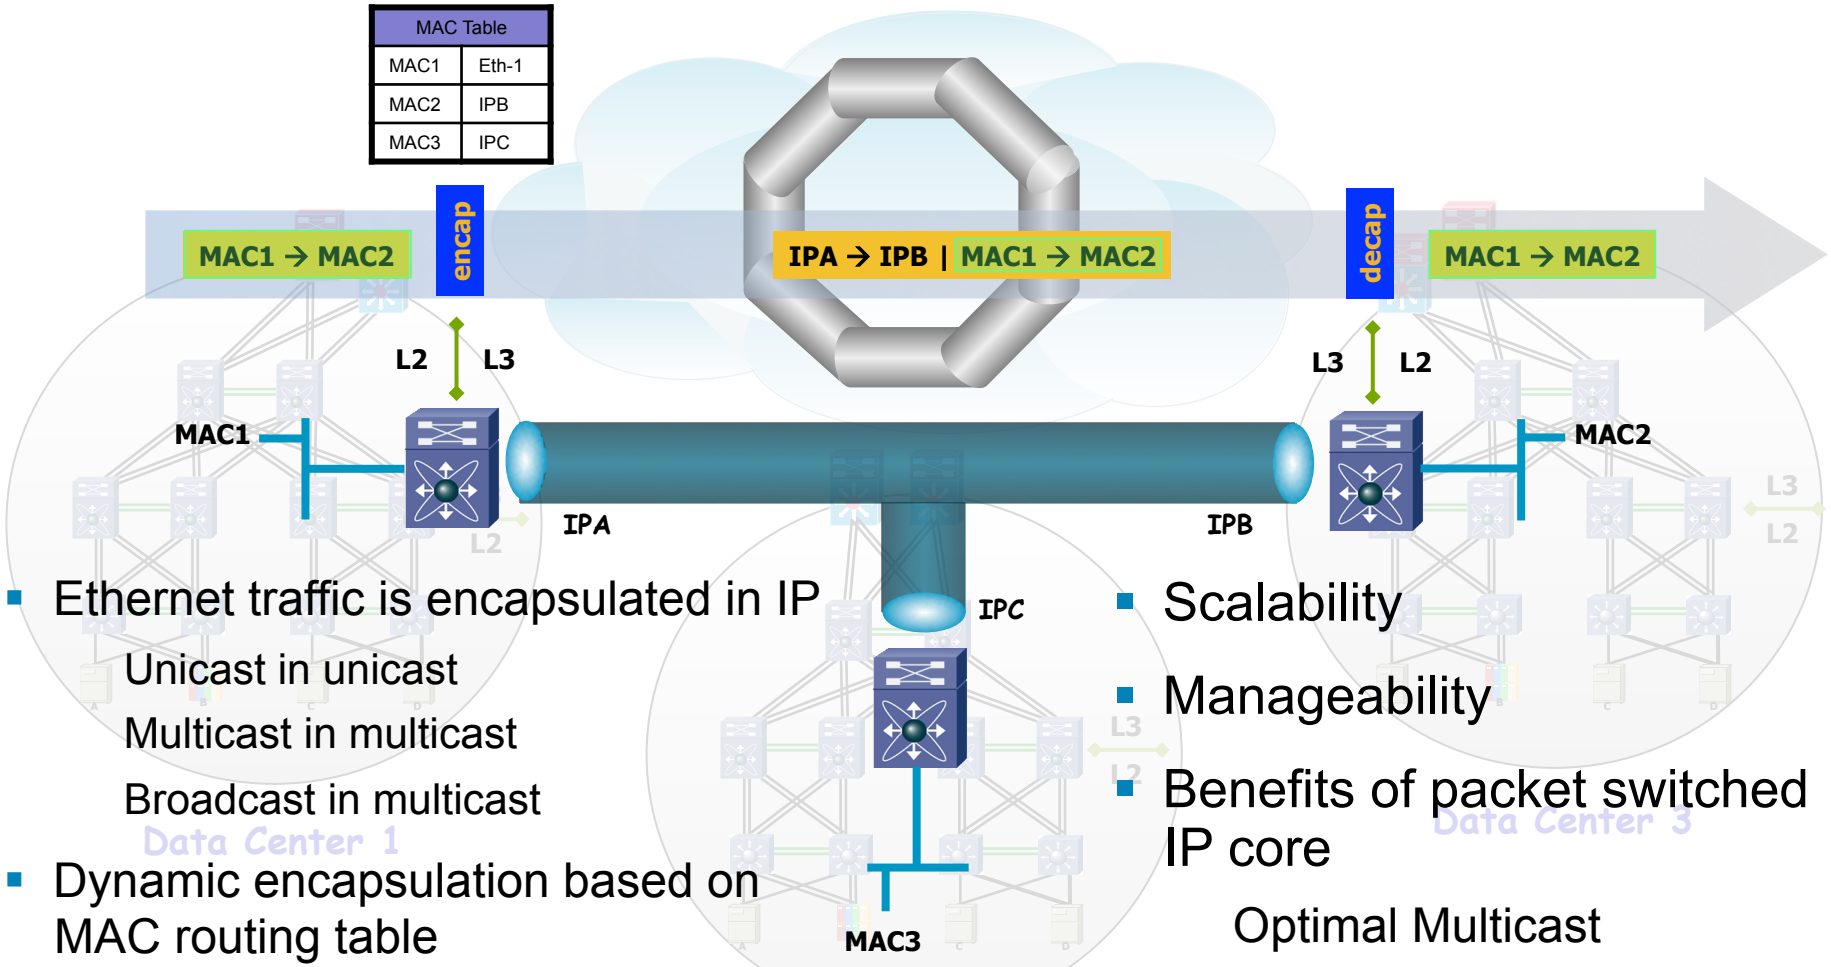

- Ethernet traffic is encapsulated in IP
  - Unicast in unicast
  - Multicast in multicast
  - Broadcast in multicast
- Dynamic encapsulation based on MAC routing table
  - No PW state maintained
  - These are NOT tunnels

- Scalability
- Manageability
- Benefits of packet switched IP core
  - Optimal Multicast
  - Fast convergence
  - ECMP

Overlay Transport Virtualization – OTV MAC Routing

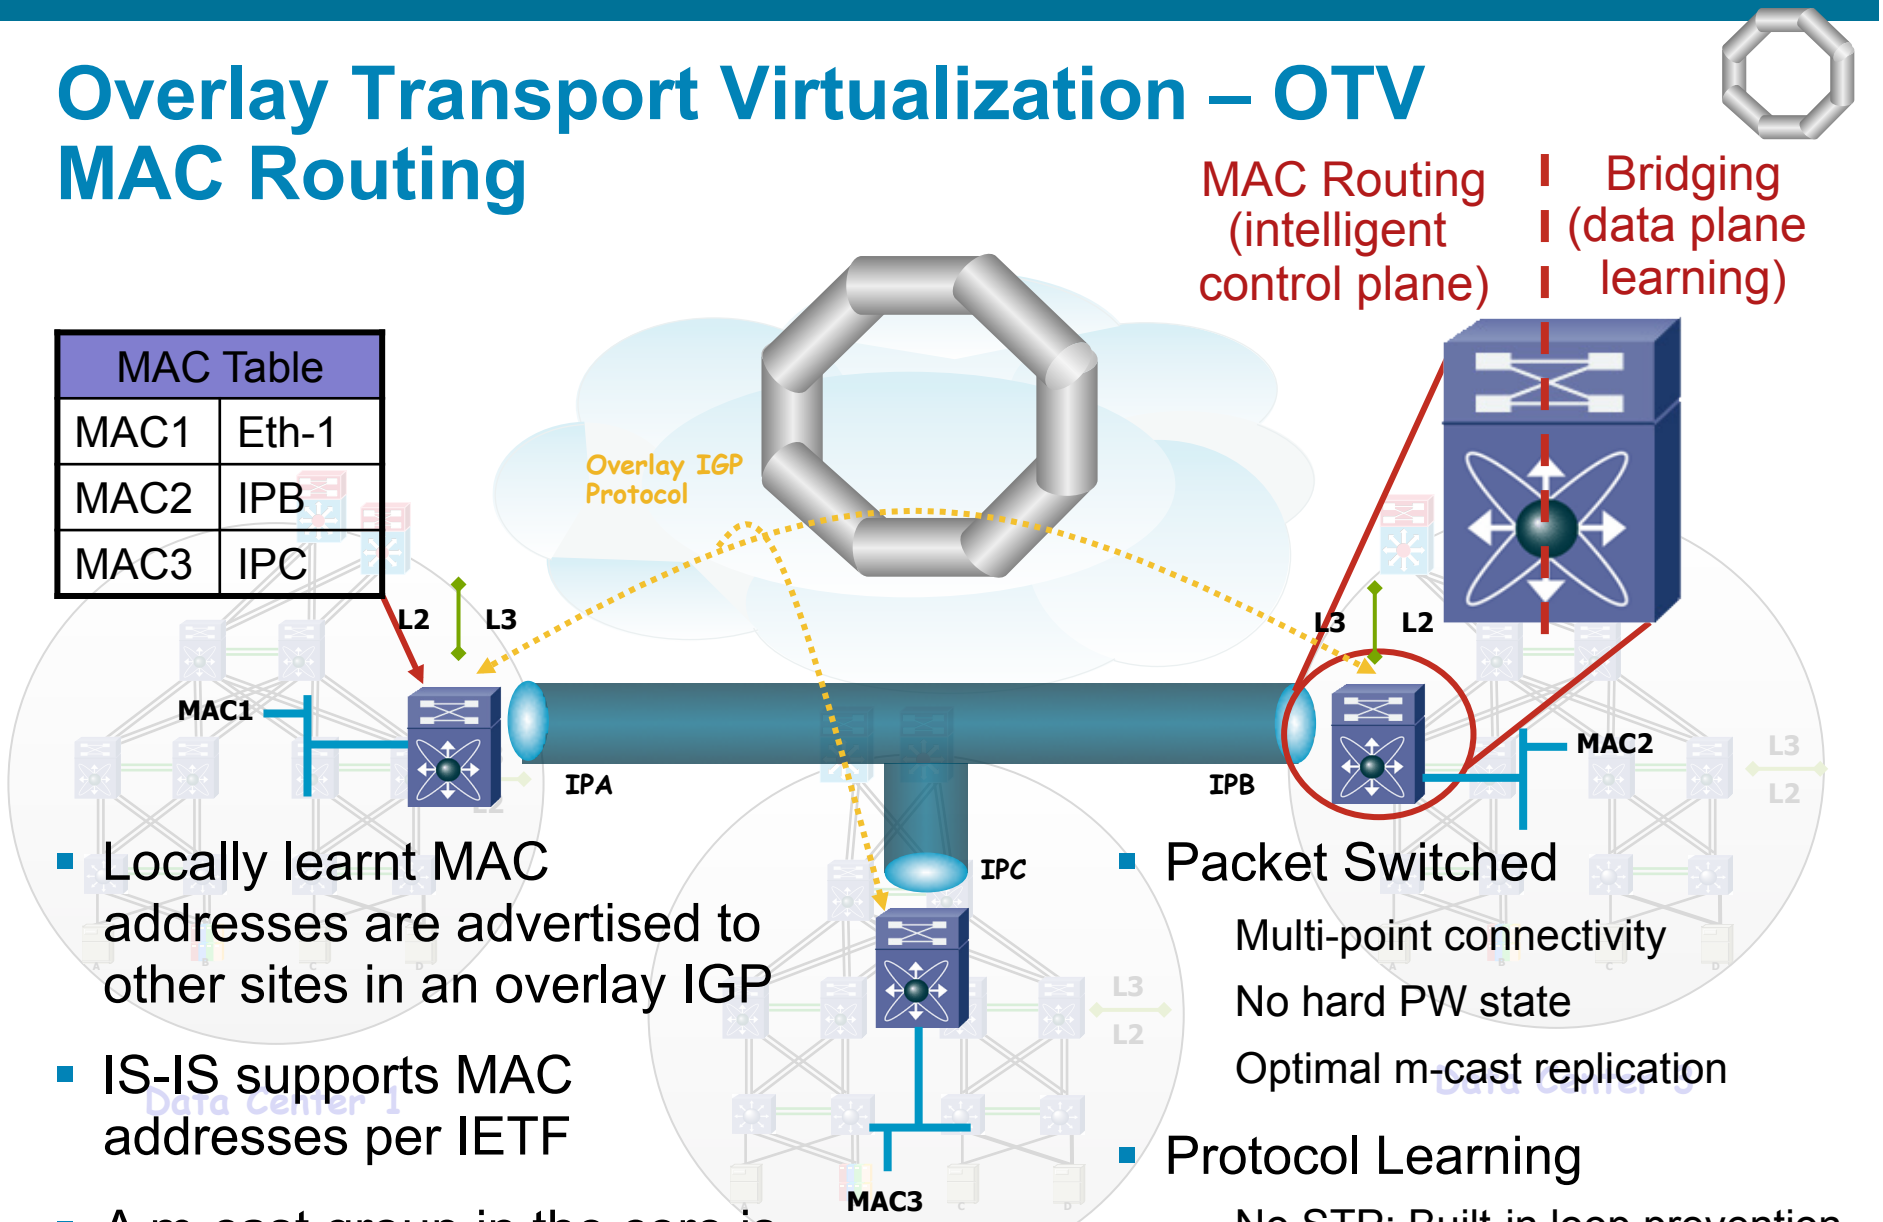

- Locally learnt MAC addresses are advertised to other sites in an overlay IGP
- IS-IS supports MAC addresses per IETF
- A m-cast group in the core is used to establish IGP peering

- Packet Switched
  - Multi-point connectivity
  - No hard PW state
  - Optimal m-cast replication
- Protocol Learning
  - No STP: Built-in loop prevention
  - Failure domain is bound
  - Flood/b-cast suppression

Core Transparency & Site Transparency

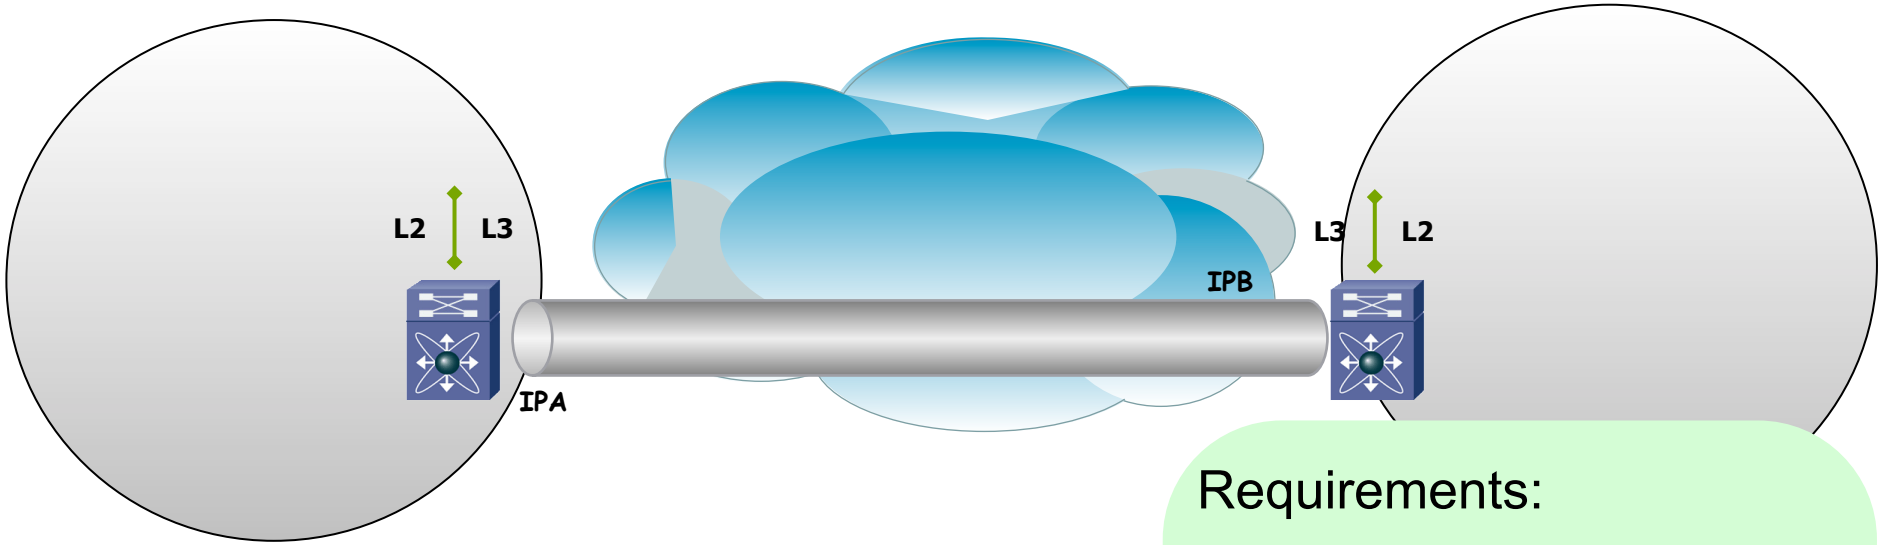

- No upgrades to the core
- No state in the core
- Core can leverage any transport: IP, MPLS, etc.
- Solution could leverage the benefits of an IP core
  - Native IP multicast

Requirements:

- Core Transparency

Optimal Multi-point Connectivity
No Circuits → Packet Switched

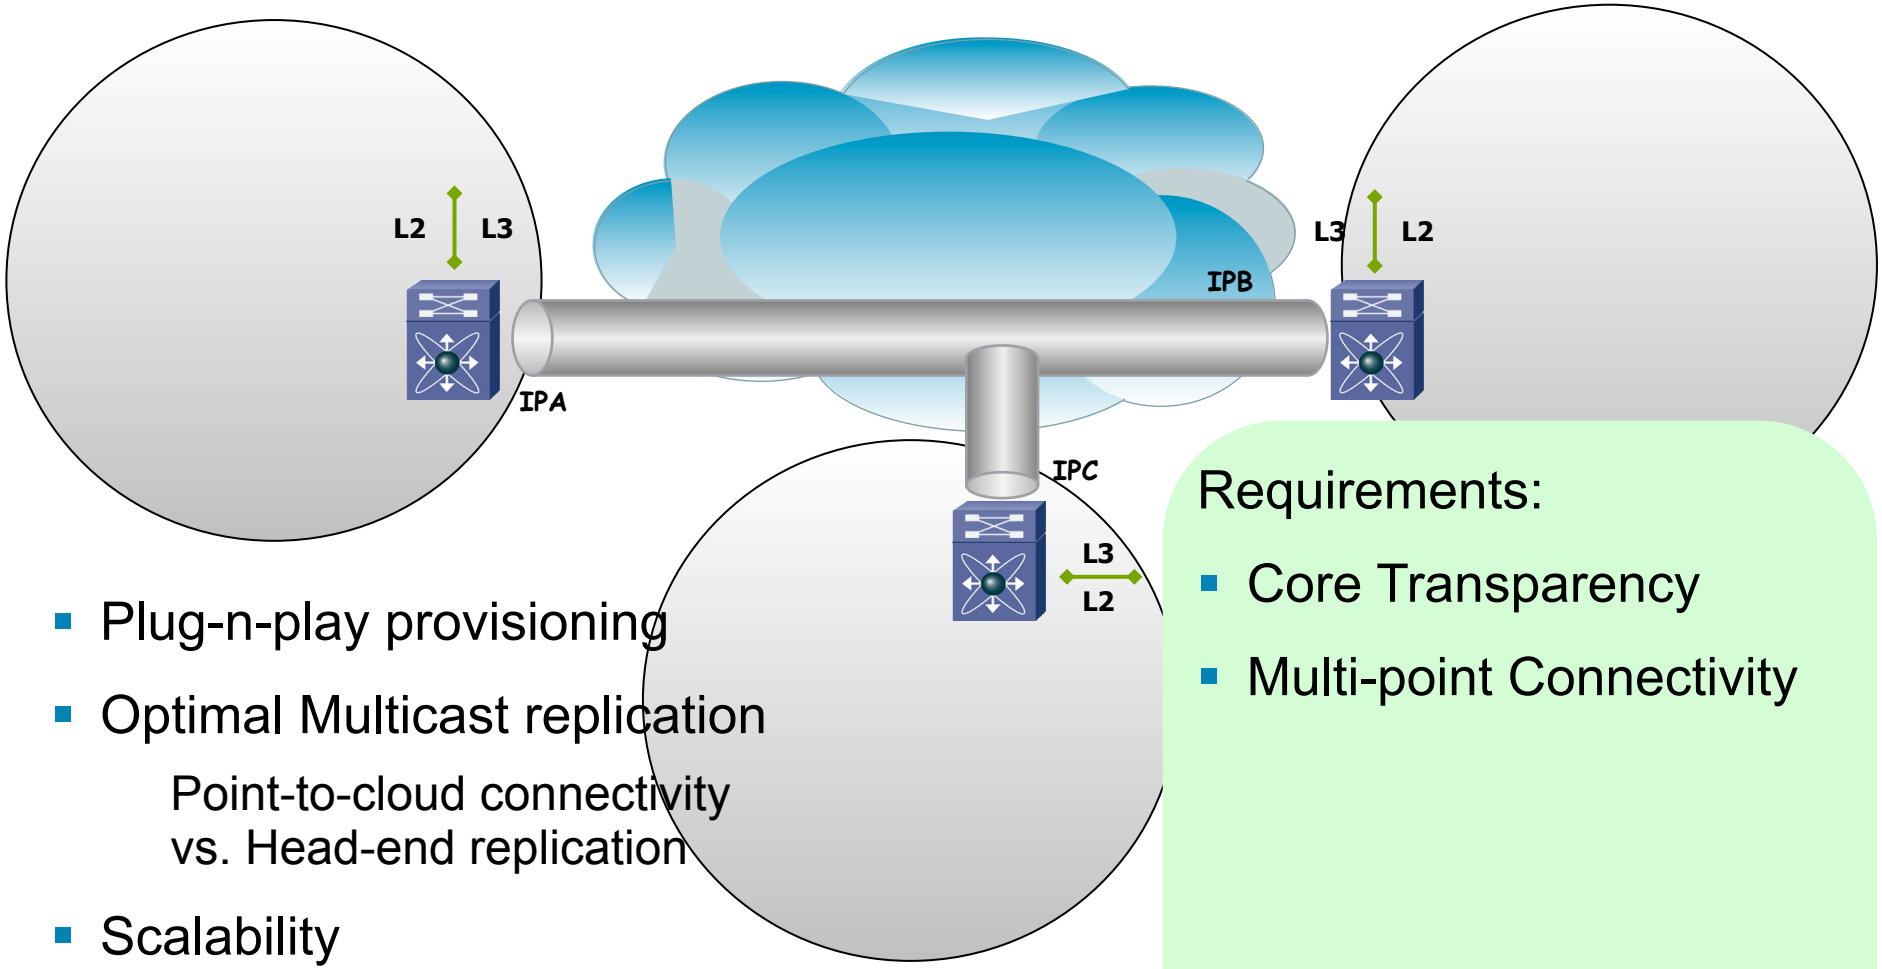

- Plug-n-play provisioning
- Optimal Multicast replication

  Point-to-cloud connectivity
  vs. Head-end replication
- Scalability

  Packet Switching vs.

  Circuit State

Requirements:

- Core Transparency
- Multi-point Connectivity

Automatic Multi-homing

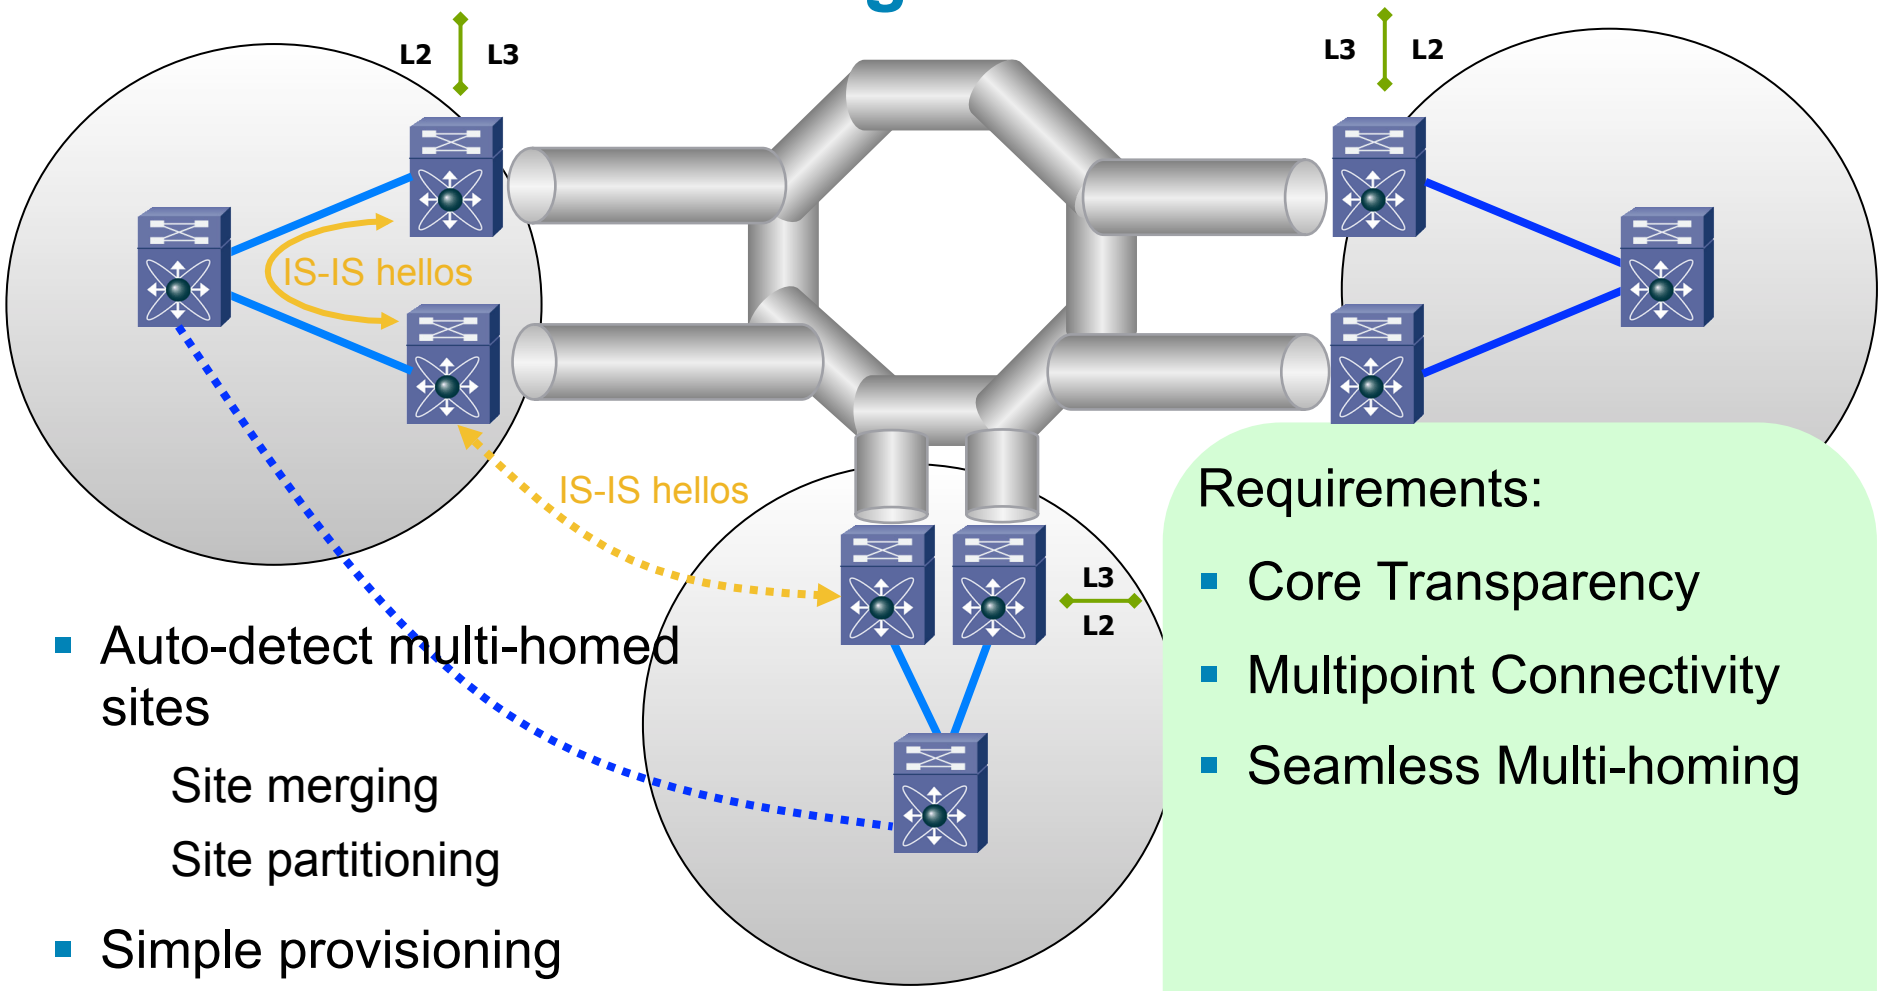

- Auto-detect multi-homed sites
  - Site merging
  - Site partitioning
- Simple provisioning
- Enabled by Protocol Learning

Requirements:

- Core Transparency
- Multipoint Connectivity
- Seamless Multi-homing

Optimal BW utilization
All paths are active

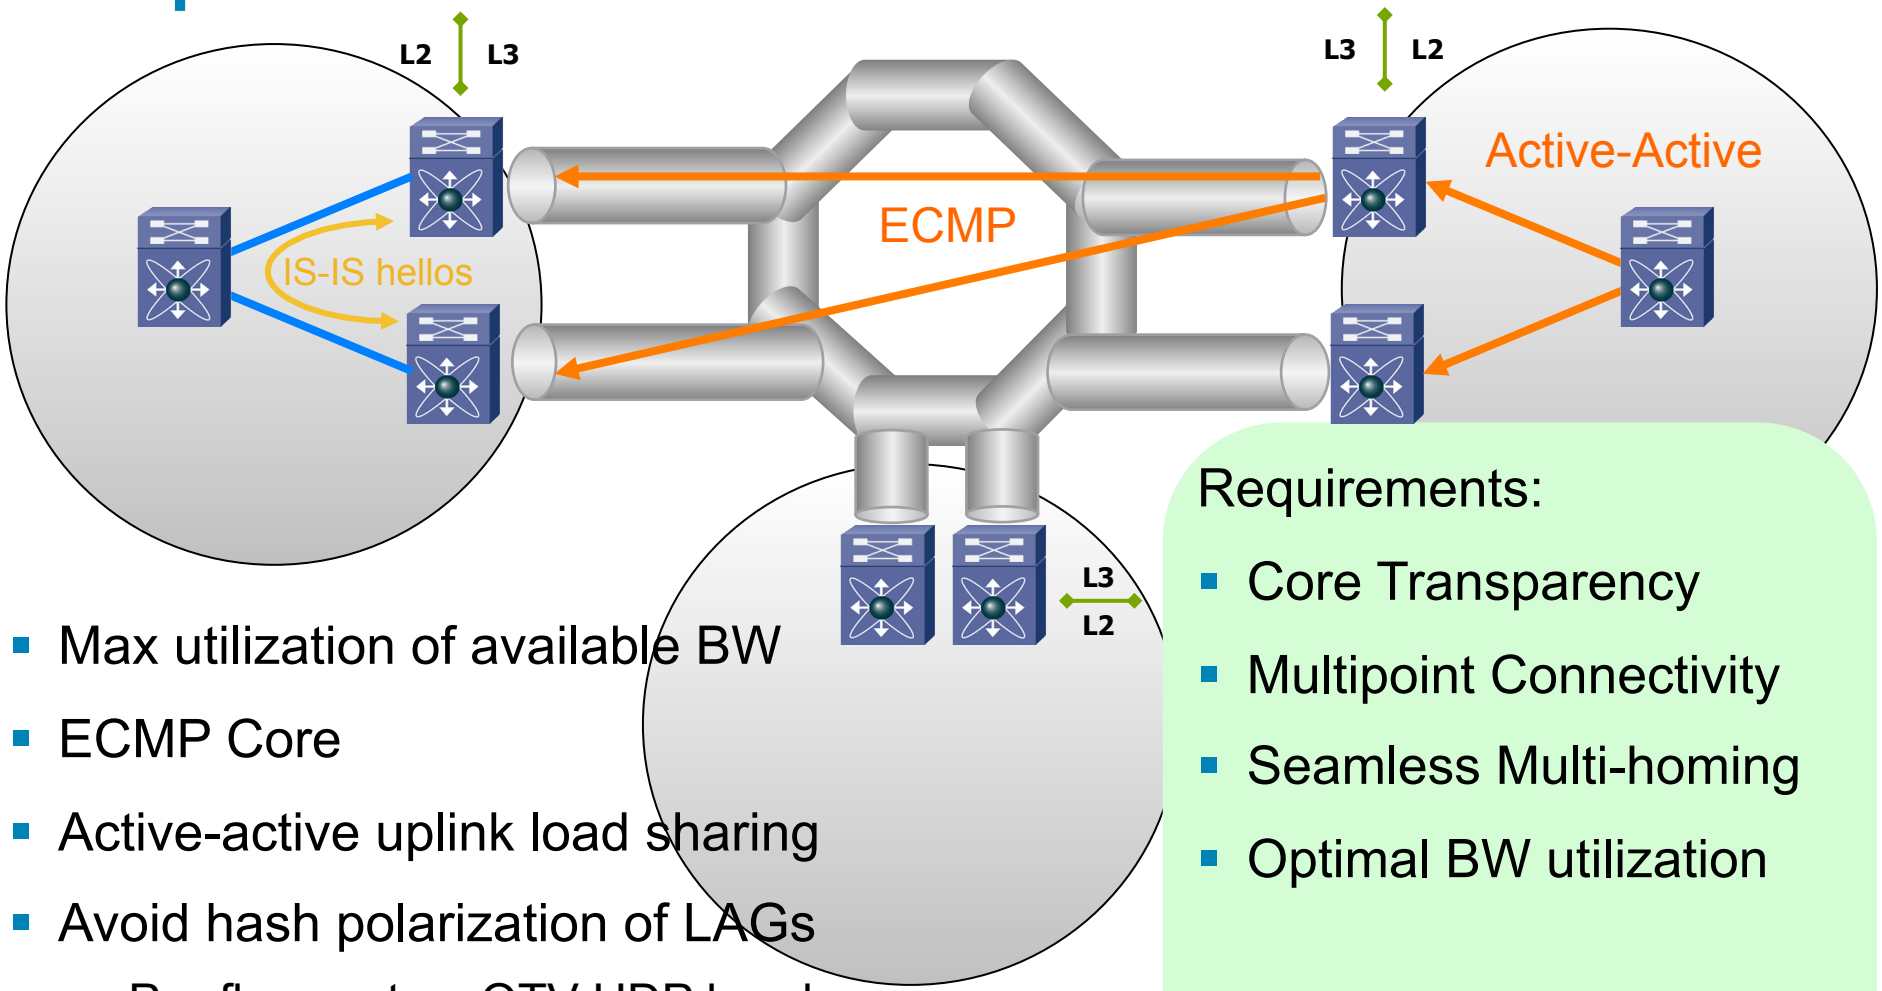

- Max utilization of available BW
- ECMP Core
- Active-active uplink load sharing
- Avoid hash polarization of LAGs
  - Per-flow port on OTV UDP header
  - Hash on five-tuple
- Enabled by Protocol Learning

Requirements:

- Core Transparency
- Multipoint Connectivity
- Seamless Multi-homing
- Optimal BW utilization

End-to-end Loop Prevention
Core loops and Induced loops

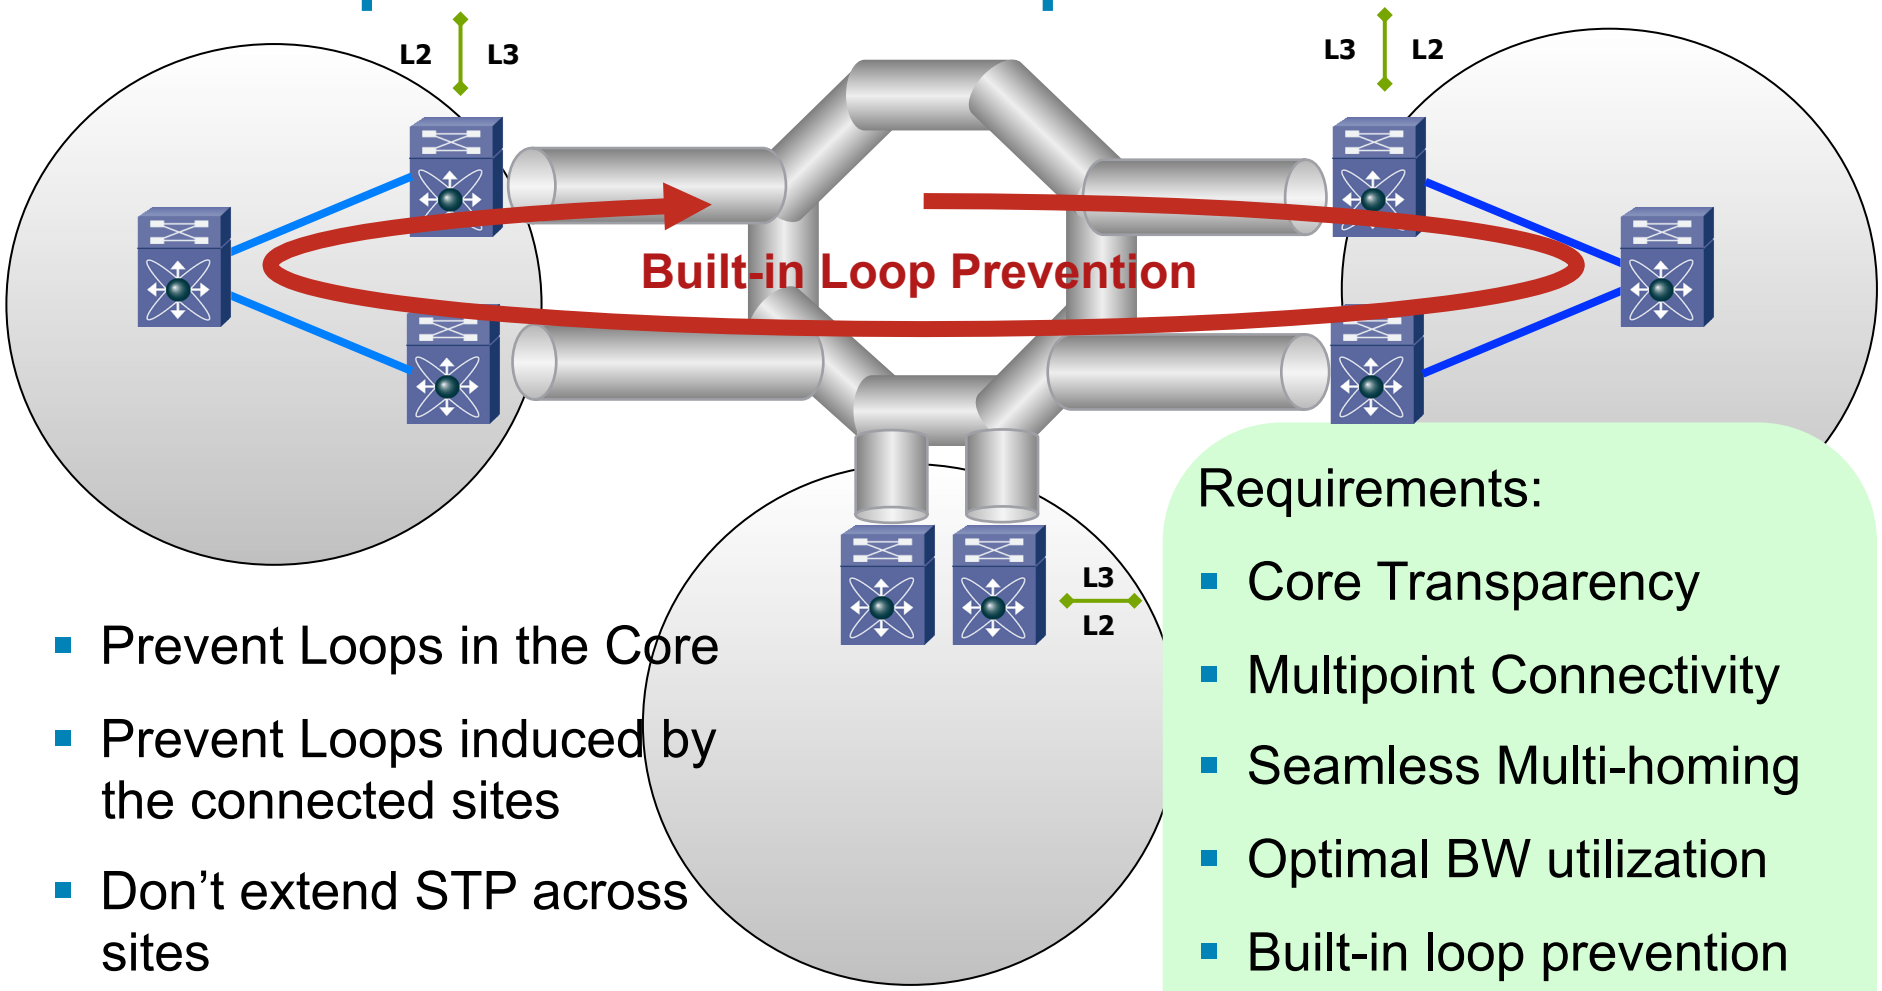

- Prevent Loops in the Core
- Prevent Loops induced by the connected sites
- Don’t extend STP across sites

  Independent STP domains
- Enabled by protocol learning

Requirements:

- Core Transparency
- Multipoint Connectivity
- Seamless Multi-homing
- Optimal BW utilization
- Built-in loop prevention

Flood Control

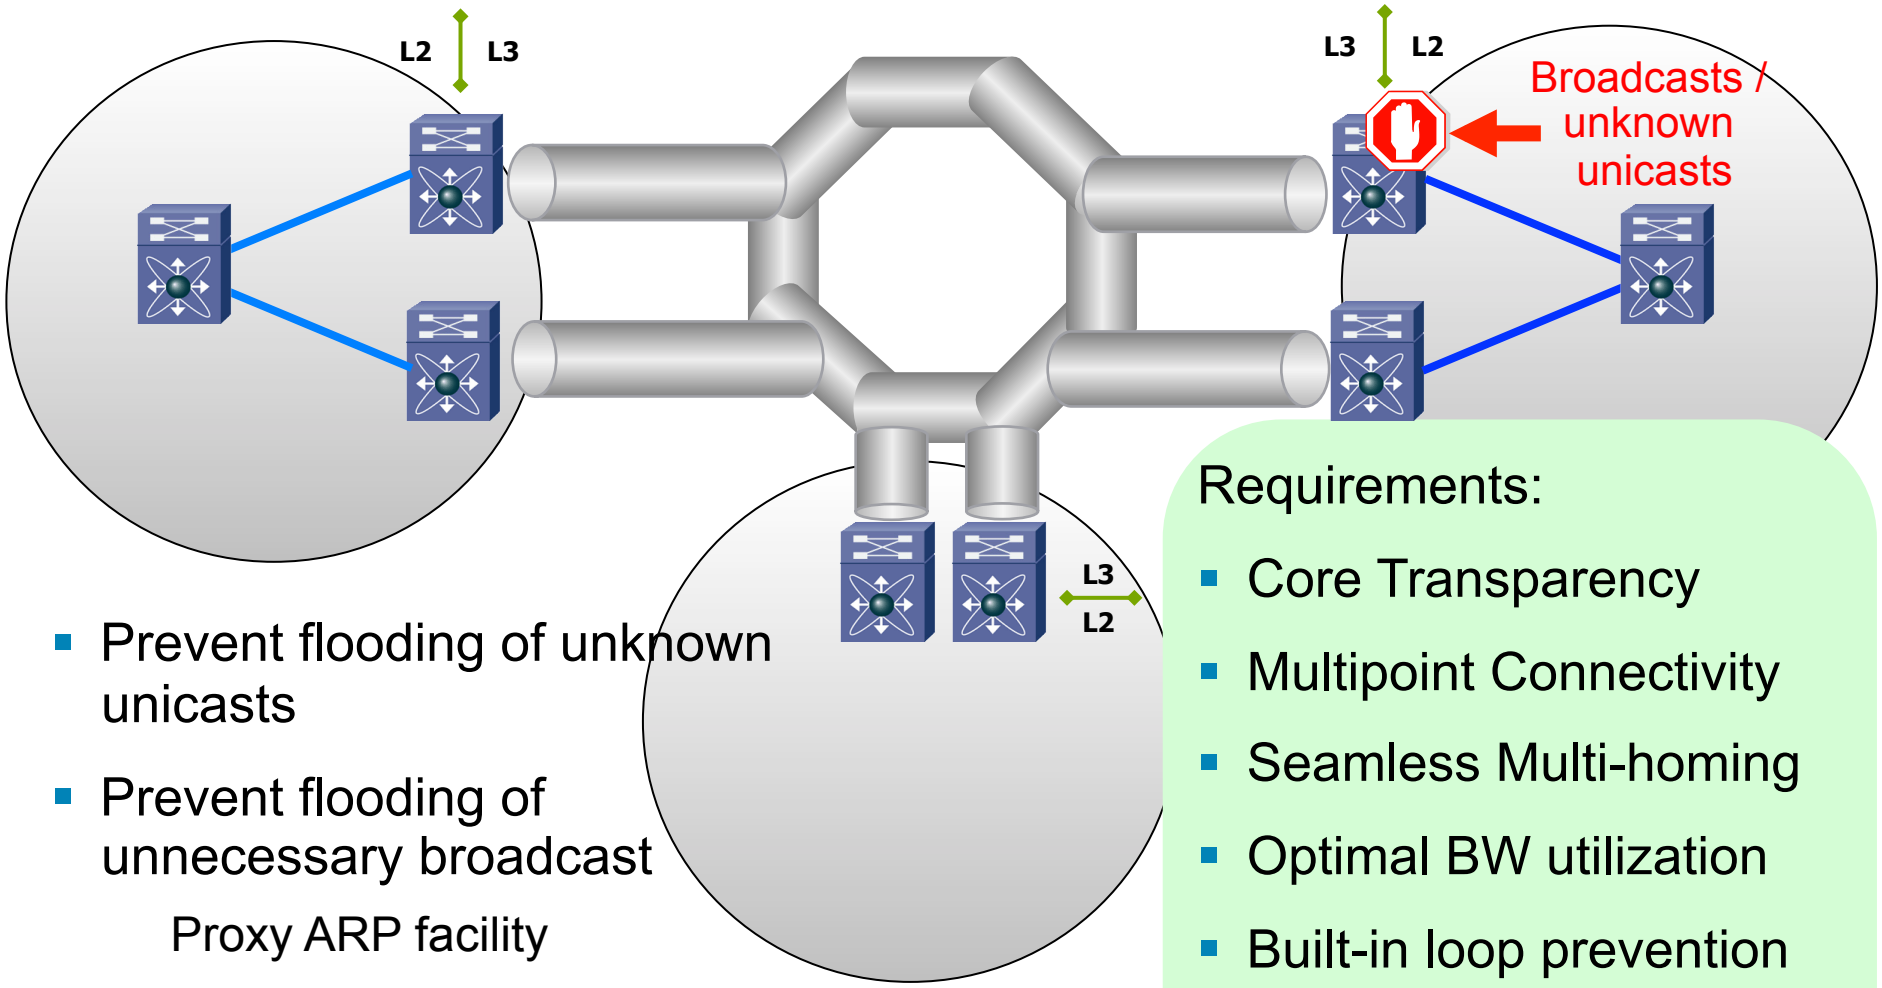

- Prevent flooding of unknown unicasts
- Prevent flooding of unnecessary broadcast
  - Proxy ARP facility
- Enabled by protocol learning

Requirements:

- Core Transparency
- Multipoint Connectivity
- Seamless Multi-homing
- Optimal BW utilization
- Built-in loop prevention
- Flood Control

Optimal Routing and LAN Extensions

- Routing challenges of extended LANs:
  - Sub-optimal routing
  - Asymmetric routing
- Sub-divide the extended subnet:
  - Home subnet per site
  - Advertise more specific subnets per site
  - Inject host-routes on server moves
- OTV enables:
  - ARP interception
  - HSRP filtering/partitioning
  - Move detection
  - SPF signal to gateways

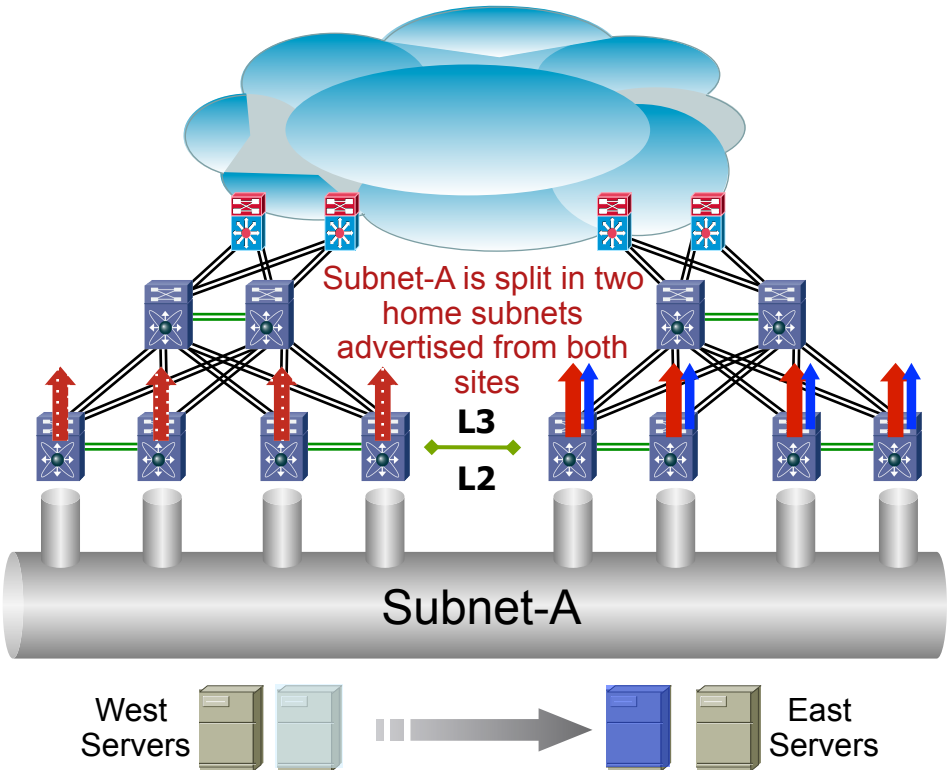

OTV's arp filtering, HSRP filtering and proxy ARP facility guarantee outbound traffic uses the optimal egress gateway

OTV-edge signals the routers to advertise more specific routes when a server move is detected (SPF Forwarding Address)

Seamless Mobility

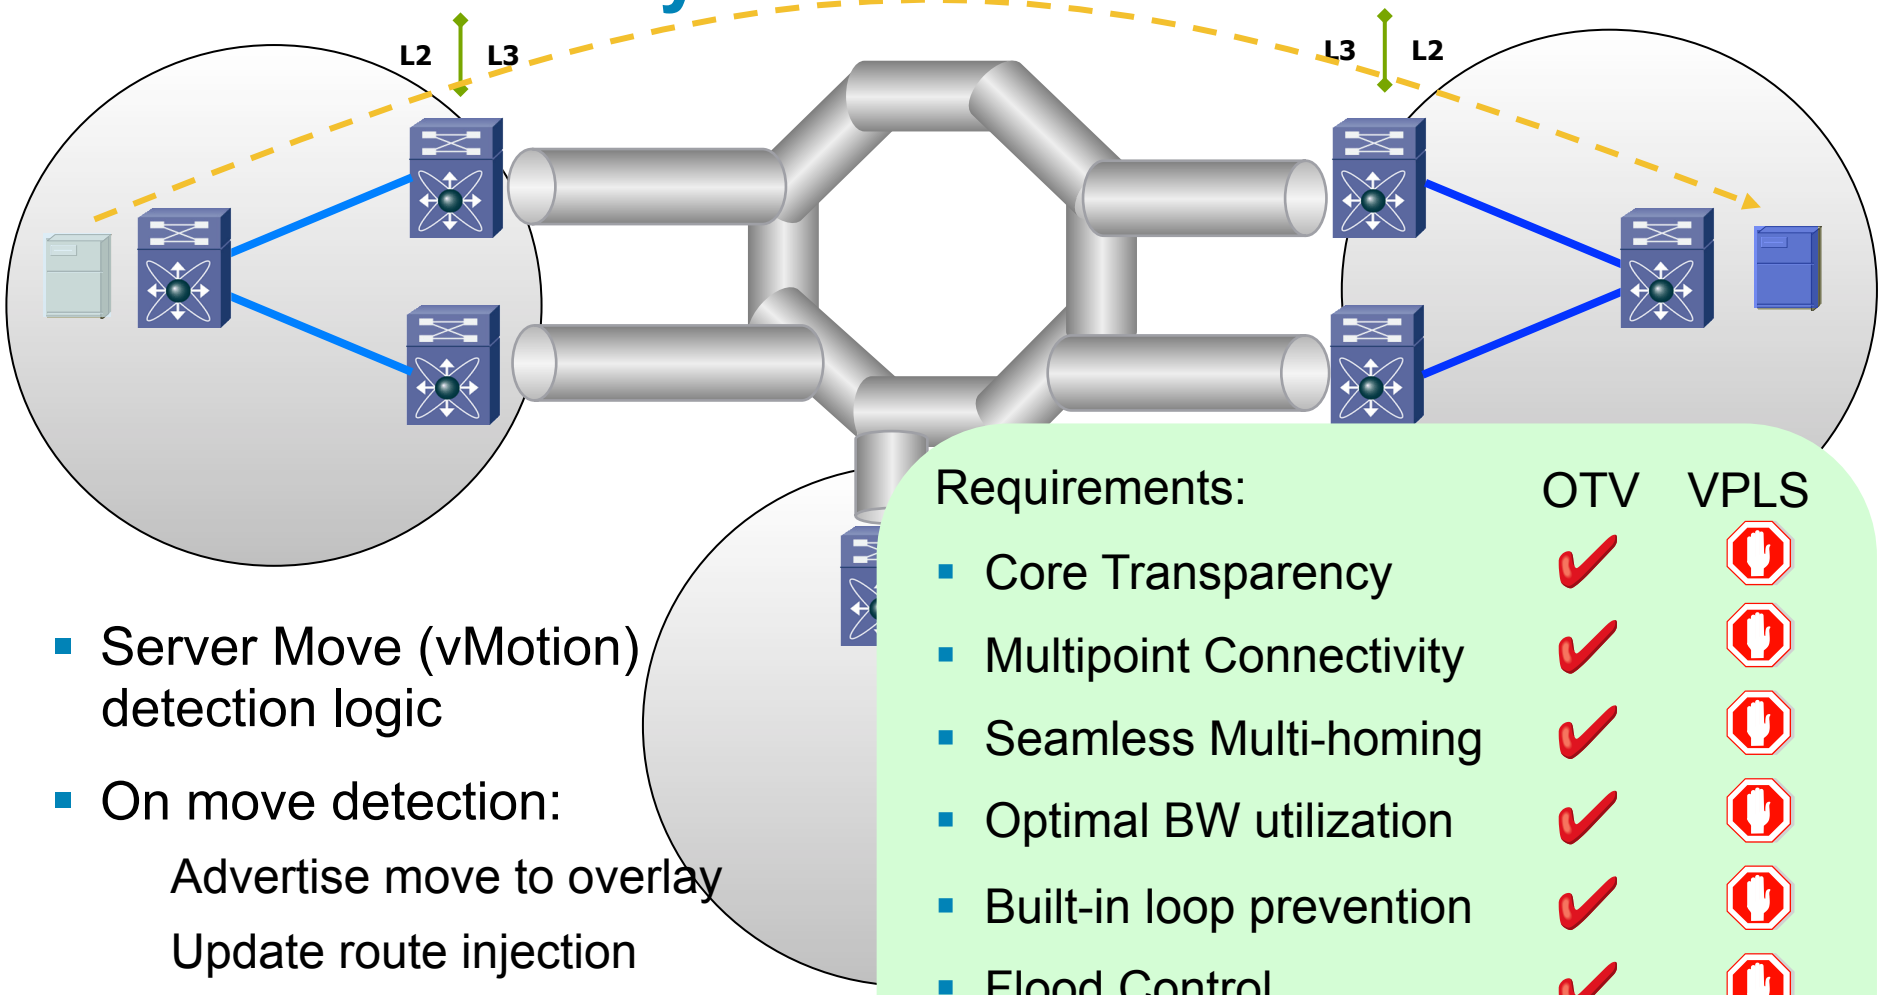

- Server Move (vMotion) detection logic
- On move detection:
  - Advertise move to overlay
  - Update route injection

Requirements:	OTV	VPLS
Core Transparency	✔	
Multipoint Connectivity	✔	
Seamless Multi-homing	✔	
Optimal BW utilization	✔	
Built-in loop prevention	✔	
Flood Control	✔	
Optimal routing	✔	
Seamless Mobility	✔	

Overlay Transport Virtualization Technology Comparison

Requirement	PLSB	VPLS	EoMPLS	IP PW (L2TPv3)	OTV
Core Transparency	802.1ah	MPLS	MPLS	IP unicast	IP unicast & multicast
Multi-point Connectivity	Meshed	Full mesh	No	No	Stateless
Multi-homing	No	Static with BGP	No	No	Built-in
Optimal BW utilization	Single path	Standby circuits	Standby circuits	Standby circuits	All paths active
Loop prevention	STP	Large STP	Large STP	Large STP	Built-in
Optimal traffic handling	Distribution Trees Flooding	Head-end replication Flooding	Head-end replication Flooding	Head-end replication Flooding	Native multicast No flood
Scalability	Poor	Poor	Poor	Poor	Good
Manageability	Many trees Config intensive add/delets	Many protocols Config intensive add/deletes	Many protocols Very poor add/deletes	Many protocols Config intensive add/ deletes	Single protocol Single command

Agenda

- Overview of traditional network designs
- FEX Link
- Virtual Port Channel
- Overlay Transport Virtualization (OTV)
- **Cisco FabricPath / Rbridges / TRILL**
- Transceivers
- 40GbE and 100GbE
- Management

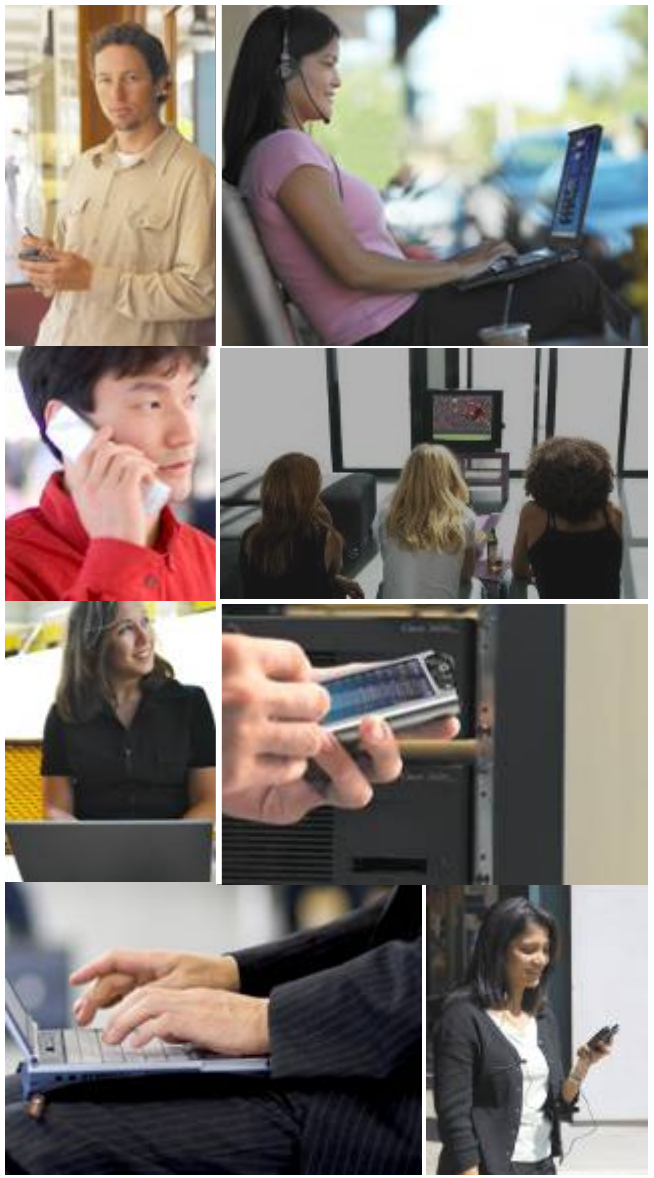

Eternal Debates on Network Design

Layer 2 or Layer 3?

Both Layer 2 and Layer 3 are required for any network design

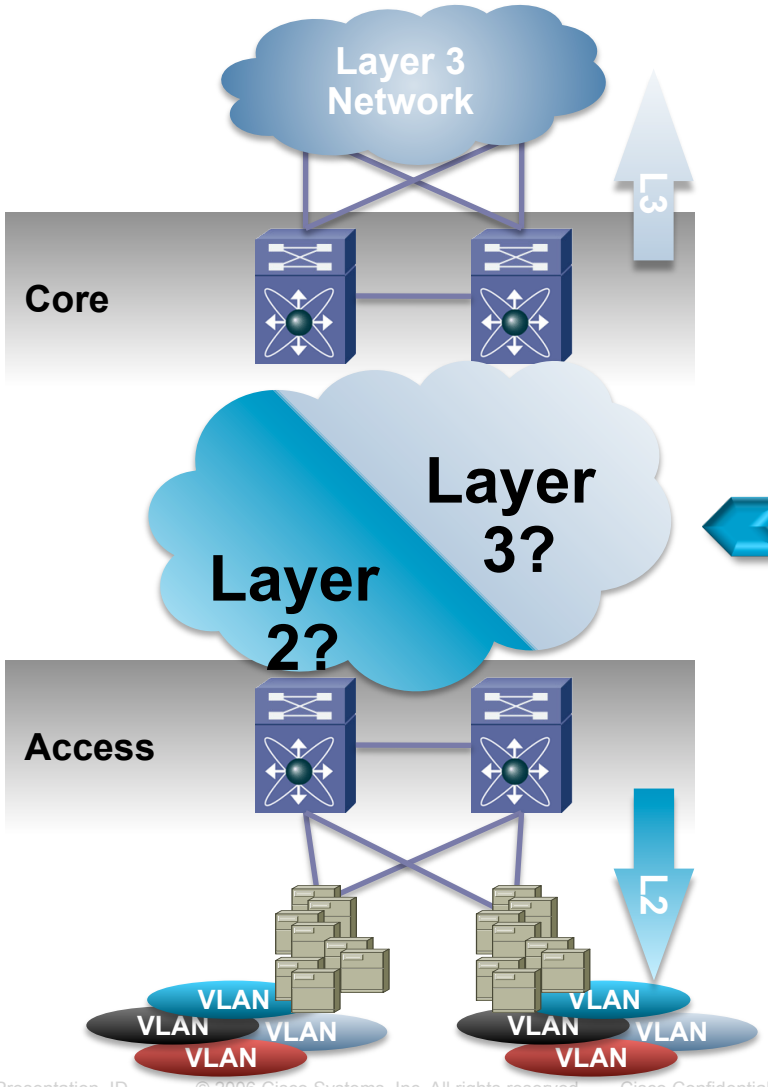

- Subnet provide fault isolation
- Scalable control planes with inherent provision of multi-pathing and multi-topology
- HA with fast convergence
- Additional loop-mitigation mechanism in the data plane (e.g. TTL, RPF check, etc.)

Cisco has solutions for both Layer 2 and Layer 3 to satisfy Customers' requirements

- Simplicity (no planning/configuration required for either addressing or control plane)
- Single control plane protocol for unicast, broadcast, and multicast
- Easy application development

L2 Network Requirements inside DC

- Maximize Bi-Sectional Bandwidth
- Scalable Layer 2 domain
- High Availability
  - Resilient control-plane
  - Fast convergence upon failure
  - Fault-domain isolation
- Facilitate Application Deployment
  - Workload mobility, Clustering, etc.
- Multi-Pathing/Multi-Topology

Constraints for Scaling Layer 2 Network

- Port Density on Switches
- Over-subscription Ratio
- Complex STP Configuration
- MAC Table Size
- X-Chassis Port-Channel

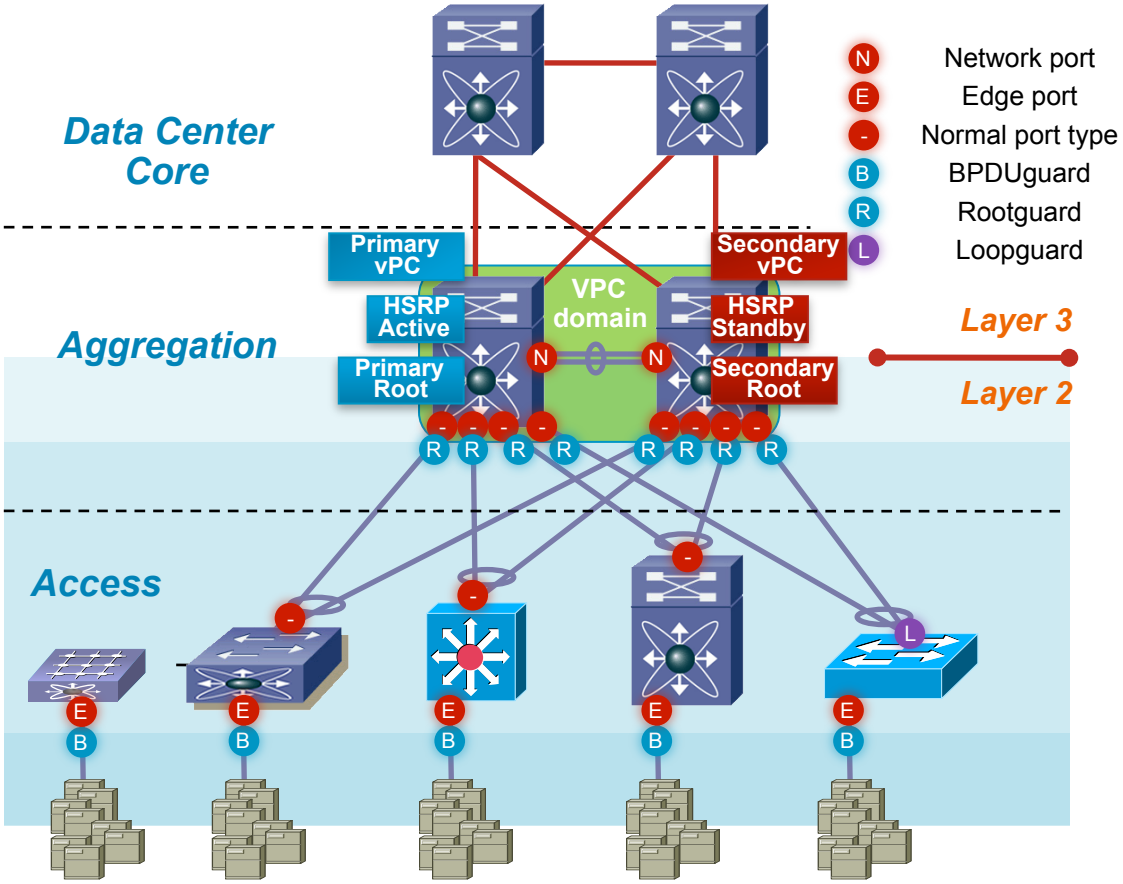

Existing Options for Layer 2 Expansion

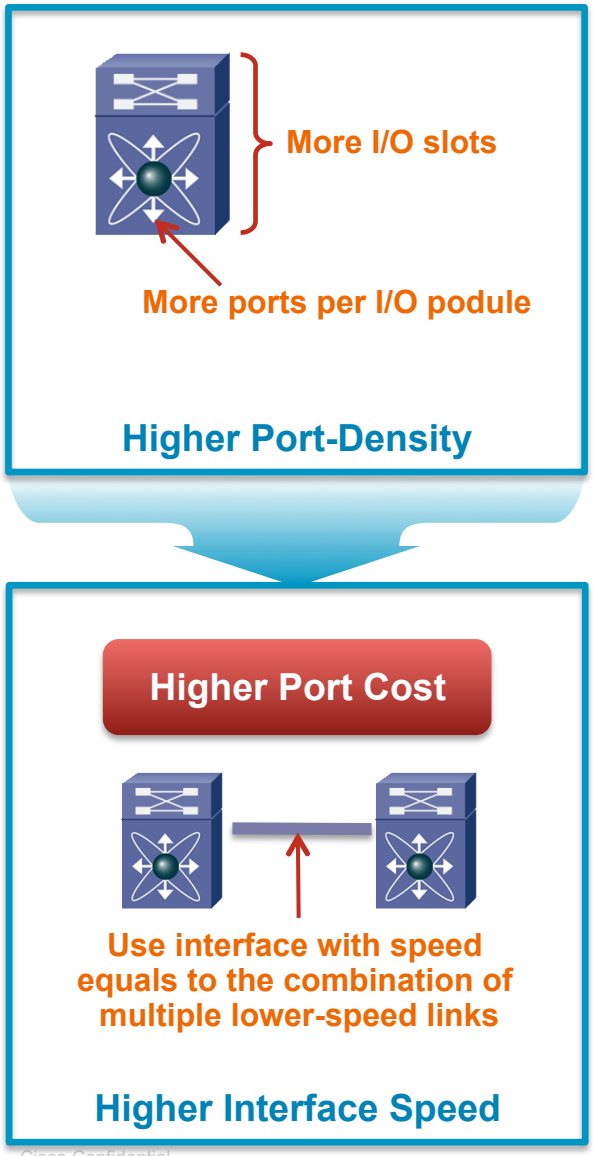

Spanning Tree and Over-subscription

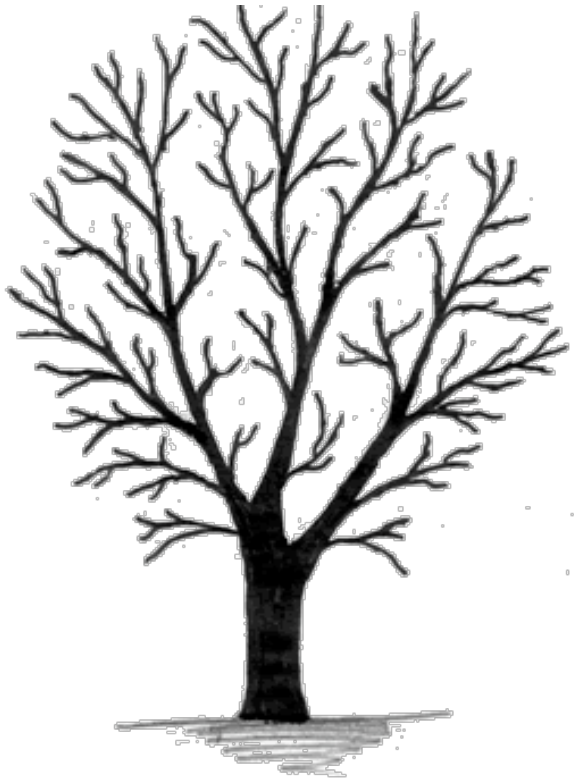

- Branches of trees never interconnect (no loop!!!)

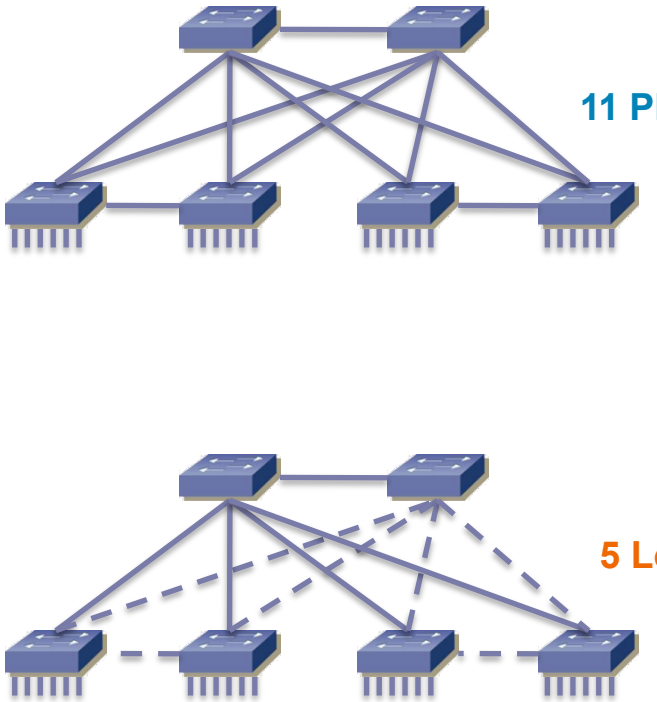

- Spanning Tree Protocol (STP) uses the same approach to build loop-free L2 logical topology
- **Over-subscription ratio exacerbated by STP algorithm**

Limitations of Spanning Tree Protocol

- Sub-optimal path selection
  - Single path between any 2 bridges in the same L2 network
  - Shortest path only from Root Bridge's perspective
- Under-utilized bandwidth
  - Ensure loop-free L2 logical topologies by blocking redundant links
  - Increased waste of available bandwidth as link-speed getting faster and faster
- No control plane security
  - Root election purely based on switch-ID, which is prone to problems caused by operator errors
- Slow and unreliable reconvergence upon link failure
  - Up to seconds of service disruption even with RSTP

Is Over-subscription Acceptable?

Campus Network

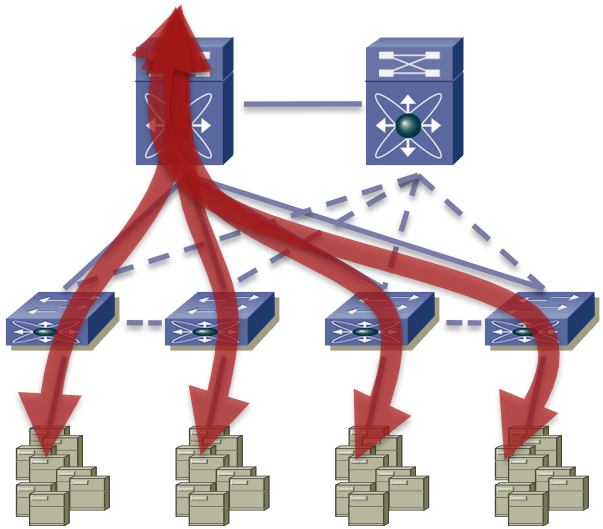

- Mostly North-South traffic flows
- Over-subscription acceptable for client-server type of applications

Data Center

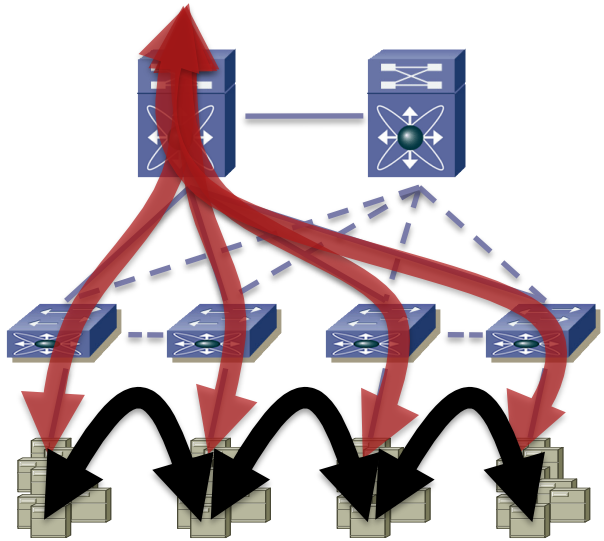

- Mix of North-South and East-West traffic flows
- Often demands special design consideration to minimize bandwidth limitation imposed by over-subscription

Nature of Layer 2 Bridging

- **Transparent** – act like “shared media” to end devices
- **Plug-N-Play** – No user configuration is required to build forwarding database
- **Data plane learning** – Forwarding database built based on frame contents
- **Flooding** – Default forwarding behavior for frames with unknown unicast destination is to flood the whole broadcast domain
- **Every MAC, Everywhere!!!** – All unicast MACs need be learn by all bridges in the same bridge domain to minimize flooding

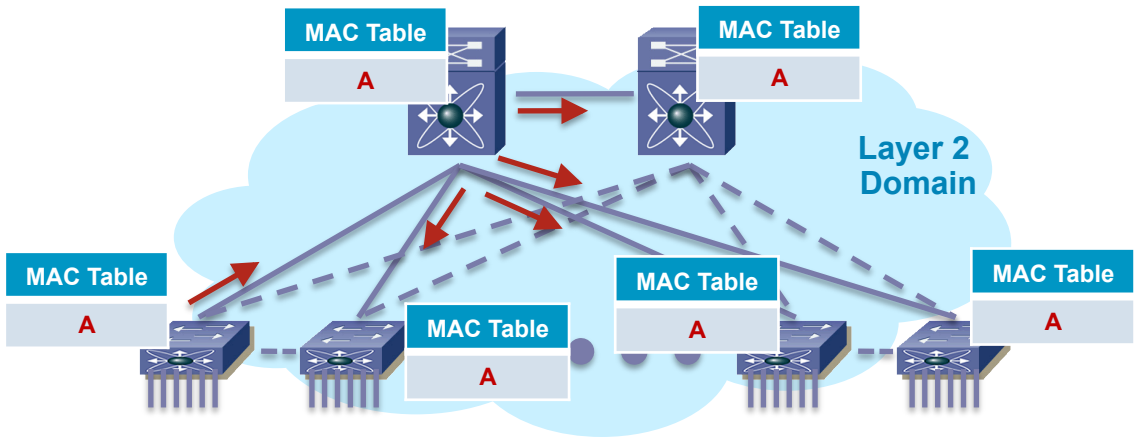

Network Addressing Scheme

MAC v.s. IP

0011.1111.1111

Non-hierarchical Address

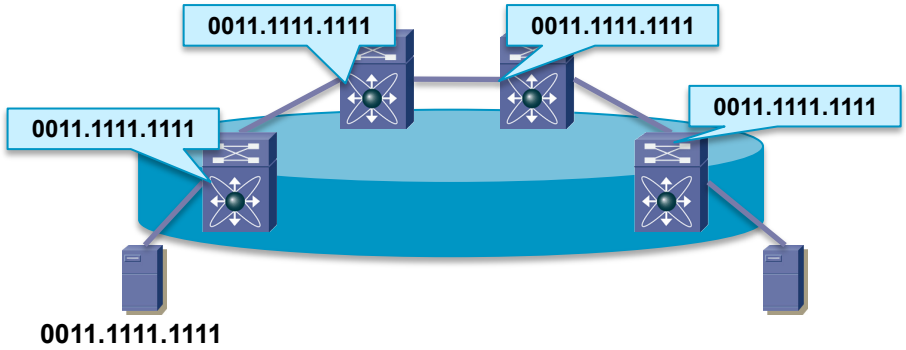

Network Address
10.0.0.0/24

10.0.0.10 /24

Host Address
10.0.0.10

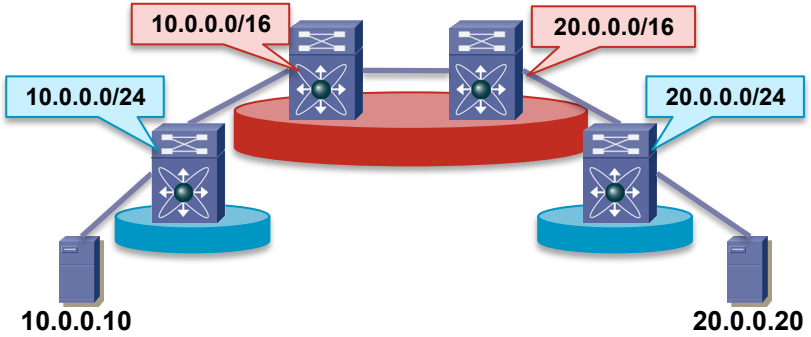

L2 Forwarding (Bridging)

- Data-plane learning
- Flat address space and forwarding table (MAC everywhere!!!)
- Flooding required for unknown unicast destination
- Destination MACs need to be known for all switches in the same network to avoid flooding

L3 Forwarding (Routing)

- Control-plane learning
- Hierarchical address space and forwarding
- Only forwarding to destination addresses with matching routes in the table
- Flooding is isolated within subnets
- No dependence on data-plane for maintaining forwarding table

The Next Era of Layer 2 Network

What Can Be Improved?

- **Network Address Scheme:** Flat ➔ ***Hierarchical***

  Additional header is required to allow L2 “Routing” instead of “Bridging”

  Provide additional loop-prevention mechanism like TTL

- **Address Learning:** Data Plane ➔ ***Control Plane***

  Eliminate the needs to program all MACs on every switches to avoid flooding

- **Control Plane:** Distance-Vector ➔ ***Link-State***

  Improve scalability, minimize convergence time, and allow multipathing inherently

The ultimate solution needs to take both control and data plane into consideration this time!!!

Cisco FabricPath
Revolutionary Solution for Scalable, Highly-Available Layer 2

Introducing Cisco FabricPath

NX-OS software innovation for multi-pathing Ethernet

"FabricPath brings Layer 3 routing benefits to flexible Layer 2 bridged Ethernet networks"

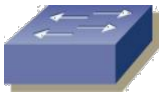

Switching

- Easy Configuration
  - Plug & Play
- Provisioning Flexibility

Routing

- Multi-pathing (ECMP)
- Fast Convergence
- Highly Scalable

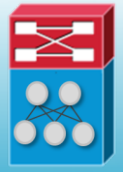

FabricPath

Auto-Discovery Simple Operation

Flexible Workload Mobility

Scalable Bandwidth

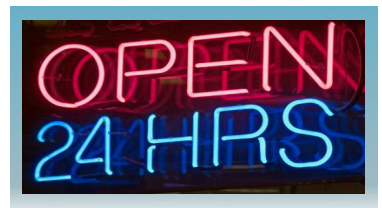

High-availability

Cisco FabricPath Overview

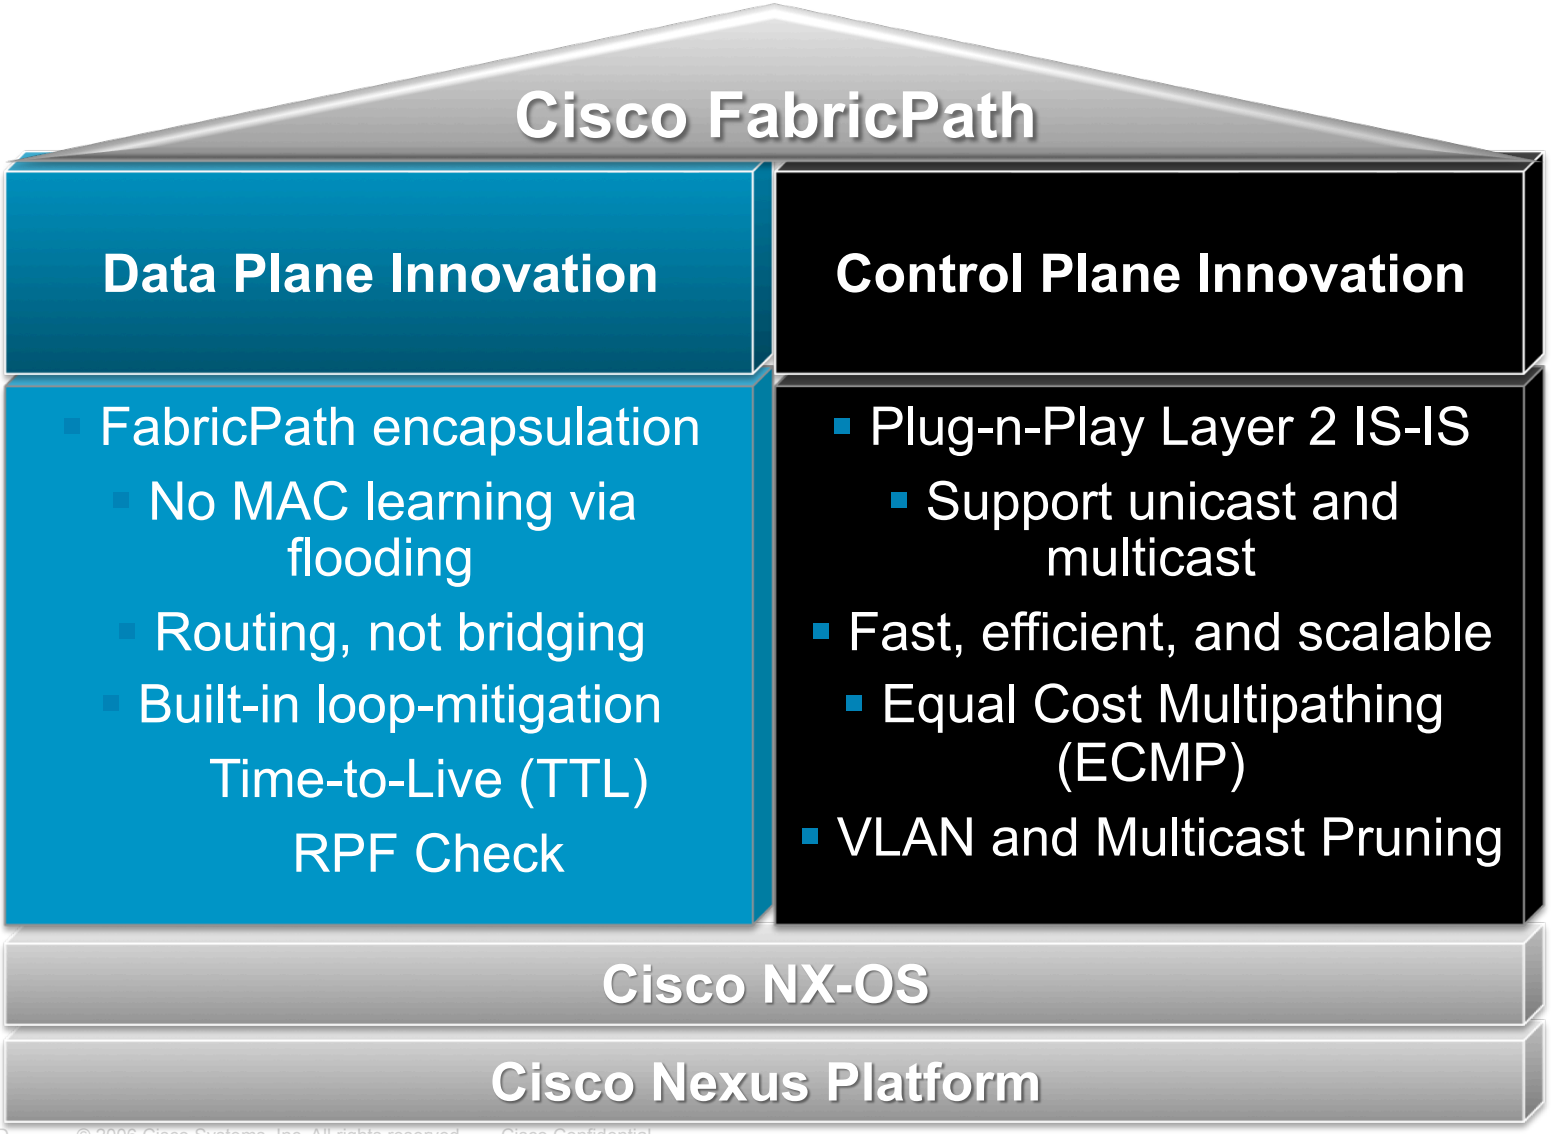

Data Plane Operation

Encapsulation to creates hierarchical address scheme

- FabricPath header is imposed by ingress switch
- Ingress and egress switch addresses are used to make "Routing" decision
- No MAC learning required inside the L2 Fabric

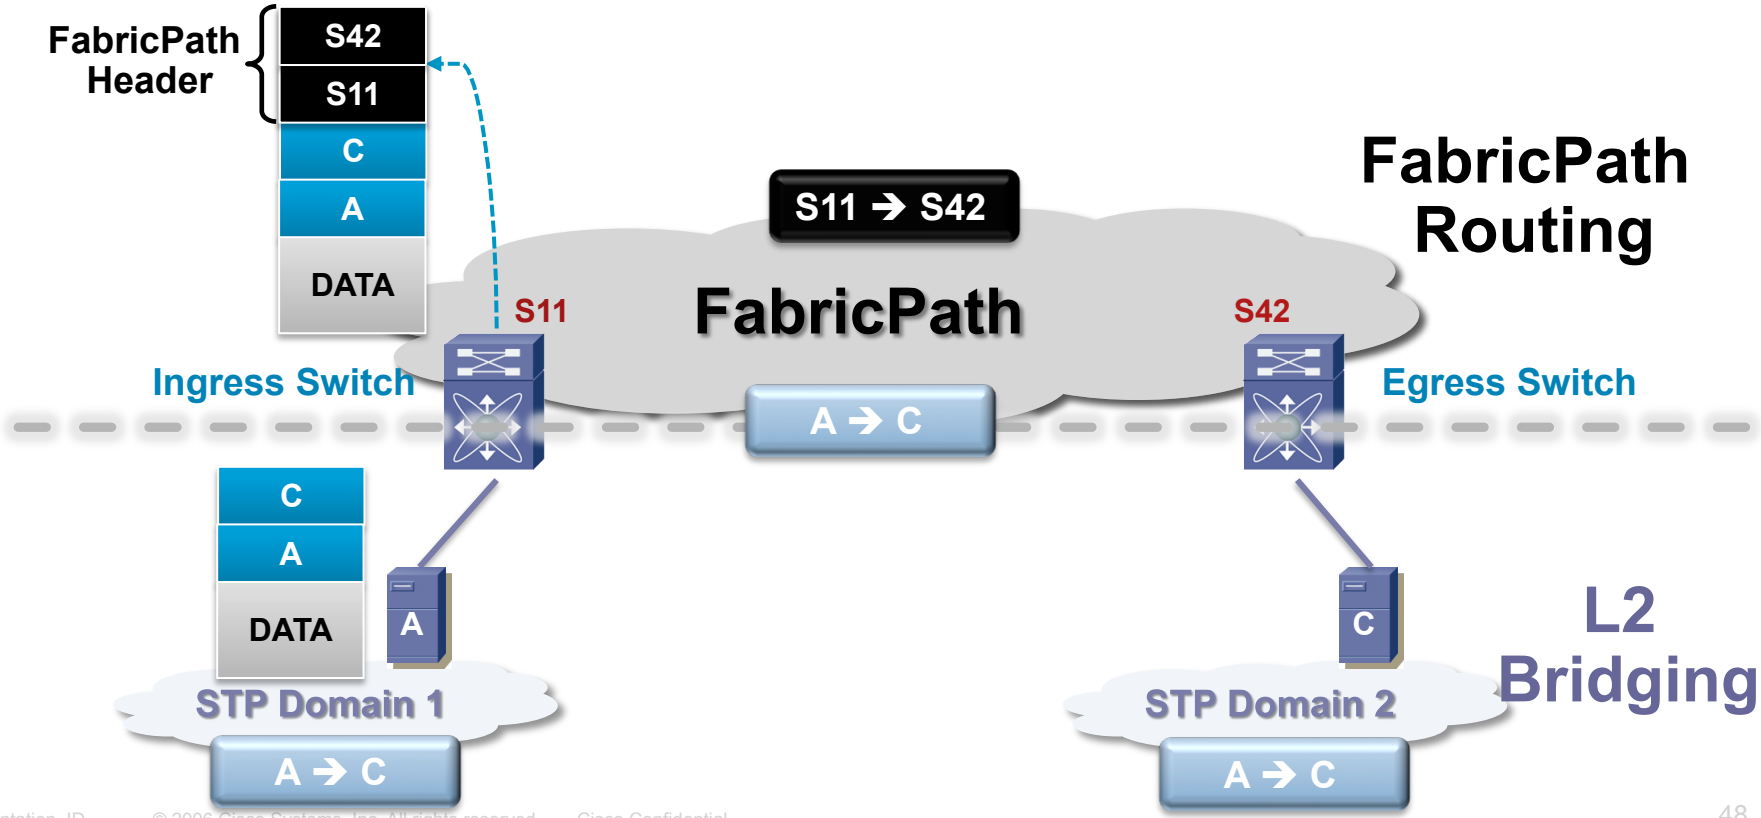

Control Plane Operation

Plug-N-Play L2 IS-IS is used to manage forwarding topology

- Assigned switch addresses to all FabricPath enabled switches automatically (no user configuration required)
- Compute shortest, pair-wise paths
- Support equal-cost paths between any FabricPath switch pairs

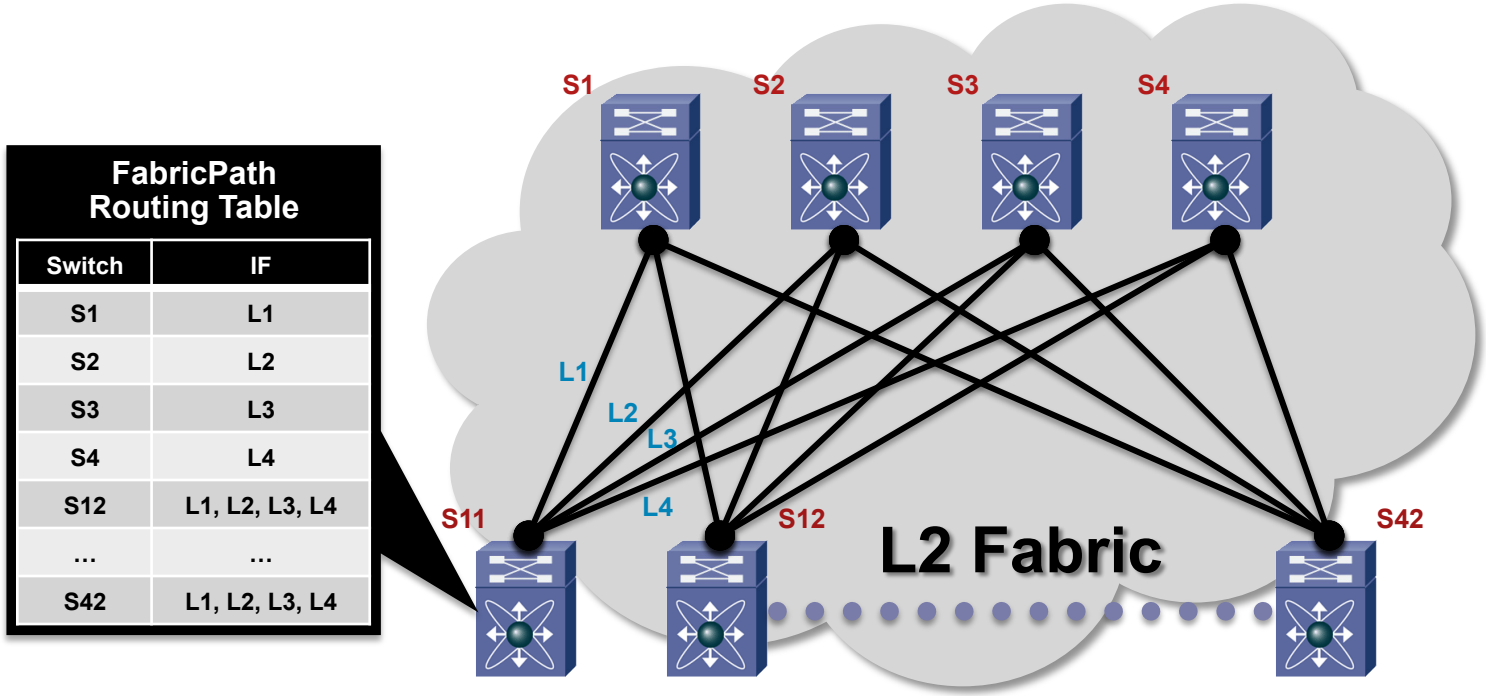

Unicast with FabricPath

Forwarding decision based on ‘FabricPath Routing Table’

- Support more than 2 active paths (up to 16) across the Fabric
- Increase bi-sectional bandwidth beyond port-channel
- High availability with N+1 path redundancy

Switch	IF
…	…
S42	L1, L2, L3, L4

MAC	IF
A	1/1
…	…
C	S42

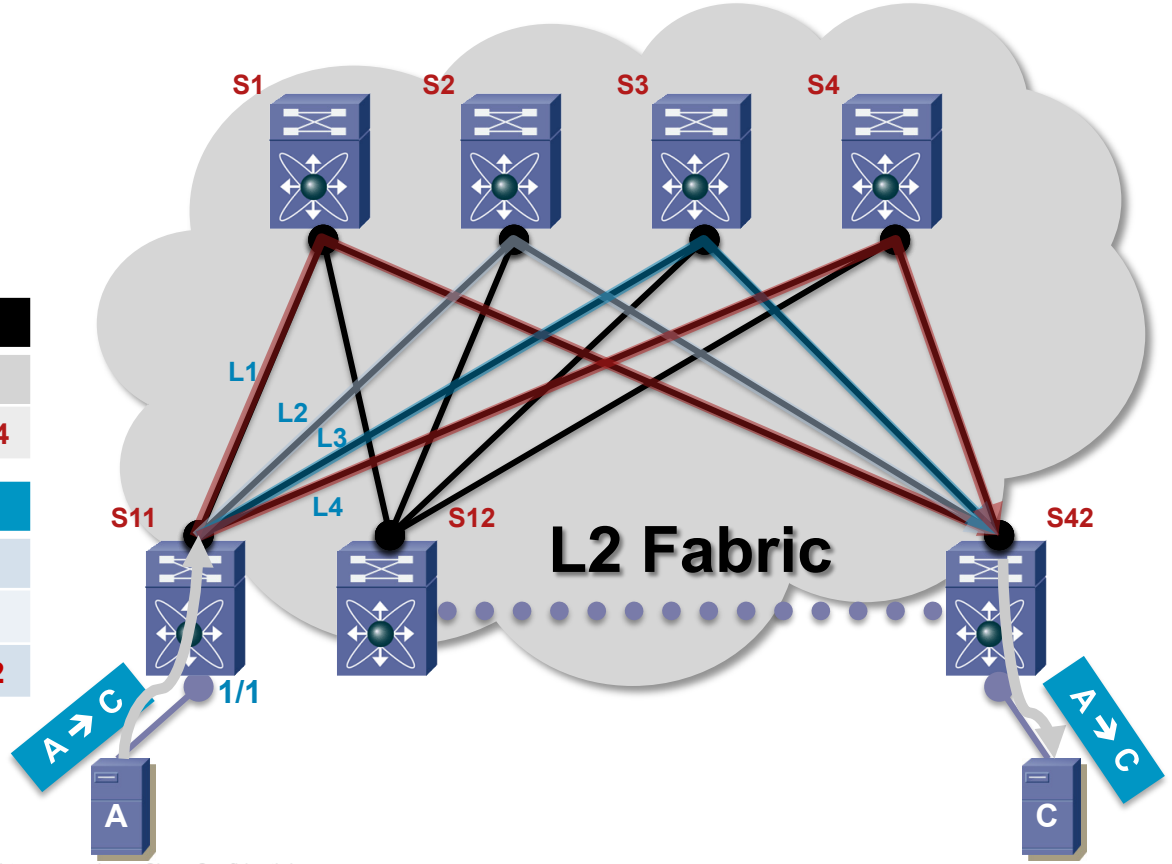

Loop Mitigation with FabricPath

Minimize impact of transient loop with TTL and RPF Check

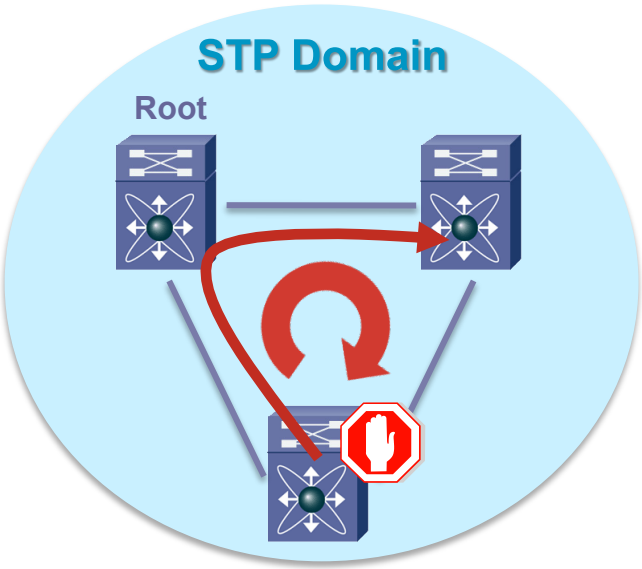

- Block redundant paths to ensure loop-free topology
- Frames loop indefinitely if STP failed
- Could results in complete network melt-down as the result of flooding

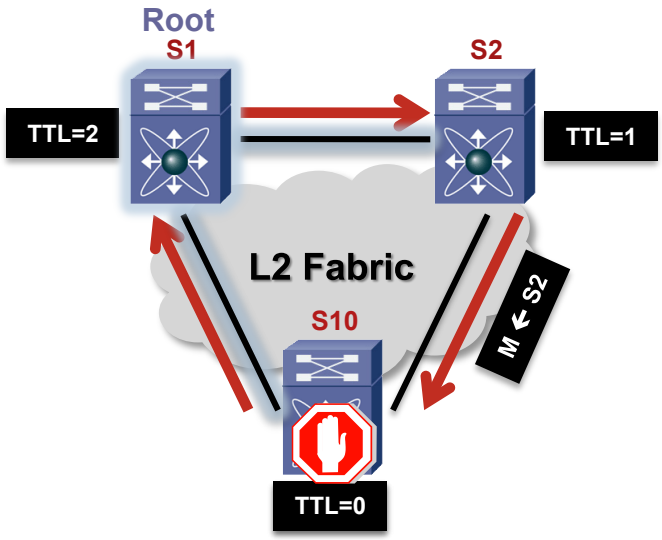

- TTL is part of FabricPath header
- Decrement by 1 at each hop
- Frames are discarded when TTL=0
- RPF check for multicast based on “tree” info

VLAN Pruning in L2 Fabric

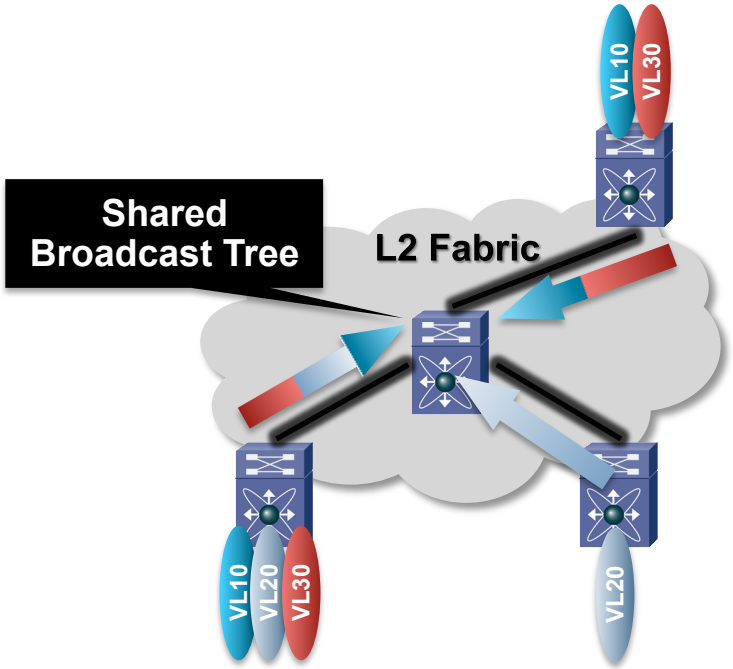

- Switches indicate *'locally interested VLANs'* to the rest of the L2 Fabric
- Broadcast traffic for any VLAN only sent to switches that have requested for it

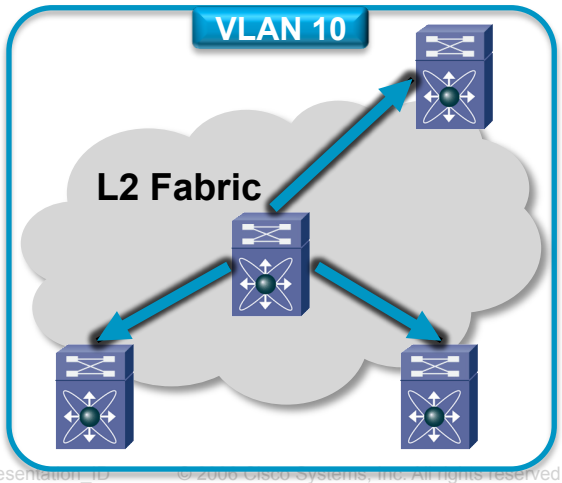

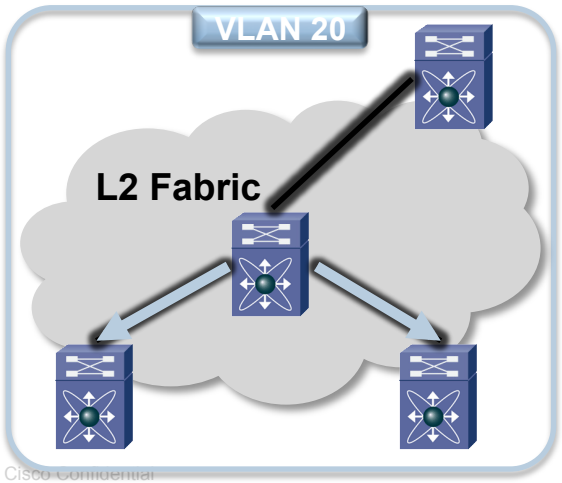

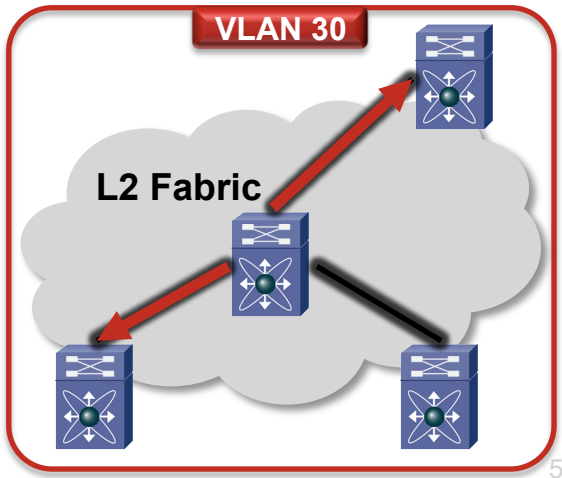

FabricPath Summary

- **FabricPath is simple**, keeps the attractive aspects of Layer 2
  - Transparent to L3 protocols
  - No addressing, simple configuration and deployment
- **FabricPath is scalable**
  - Can extend a bridged domain without extending the risks generally associated to Layer 2 (frame routing, TTL, RPFC)
- **FabricPath is efficient**
  - High bi-sectional bandwidth (ECMP)
  - Optimal path between any two nodes

Agenda

- Overview of traditional network designs
- FEX Link
- Virtual Port Channel
- Overlay Transport Virtualization (OTV)
- Cisco FabricPath / Rbridges / TRILL
- **Transceivers**
- **40GbE and 100GbE**
- Management

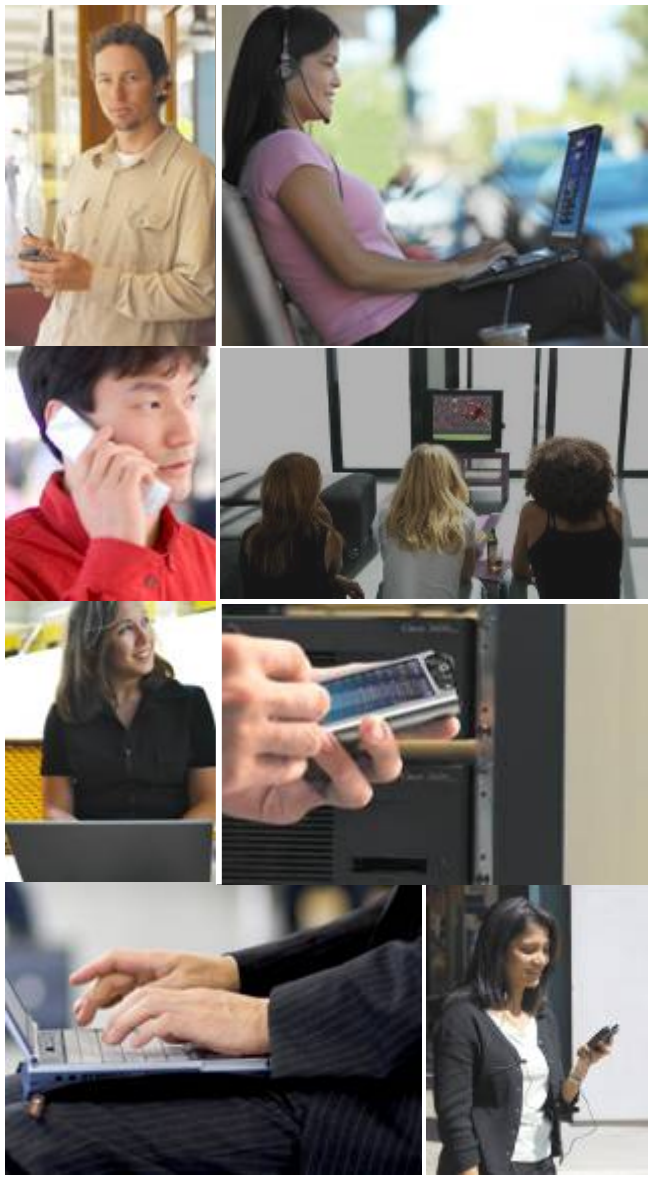

Historical Adoption of High Speed Ethernet

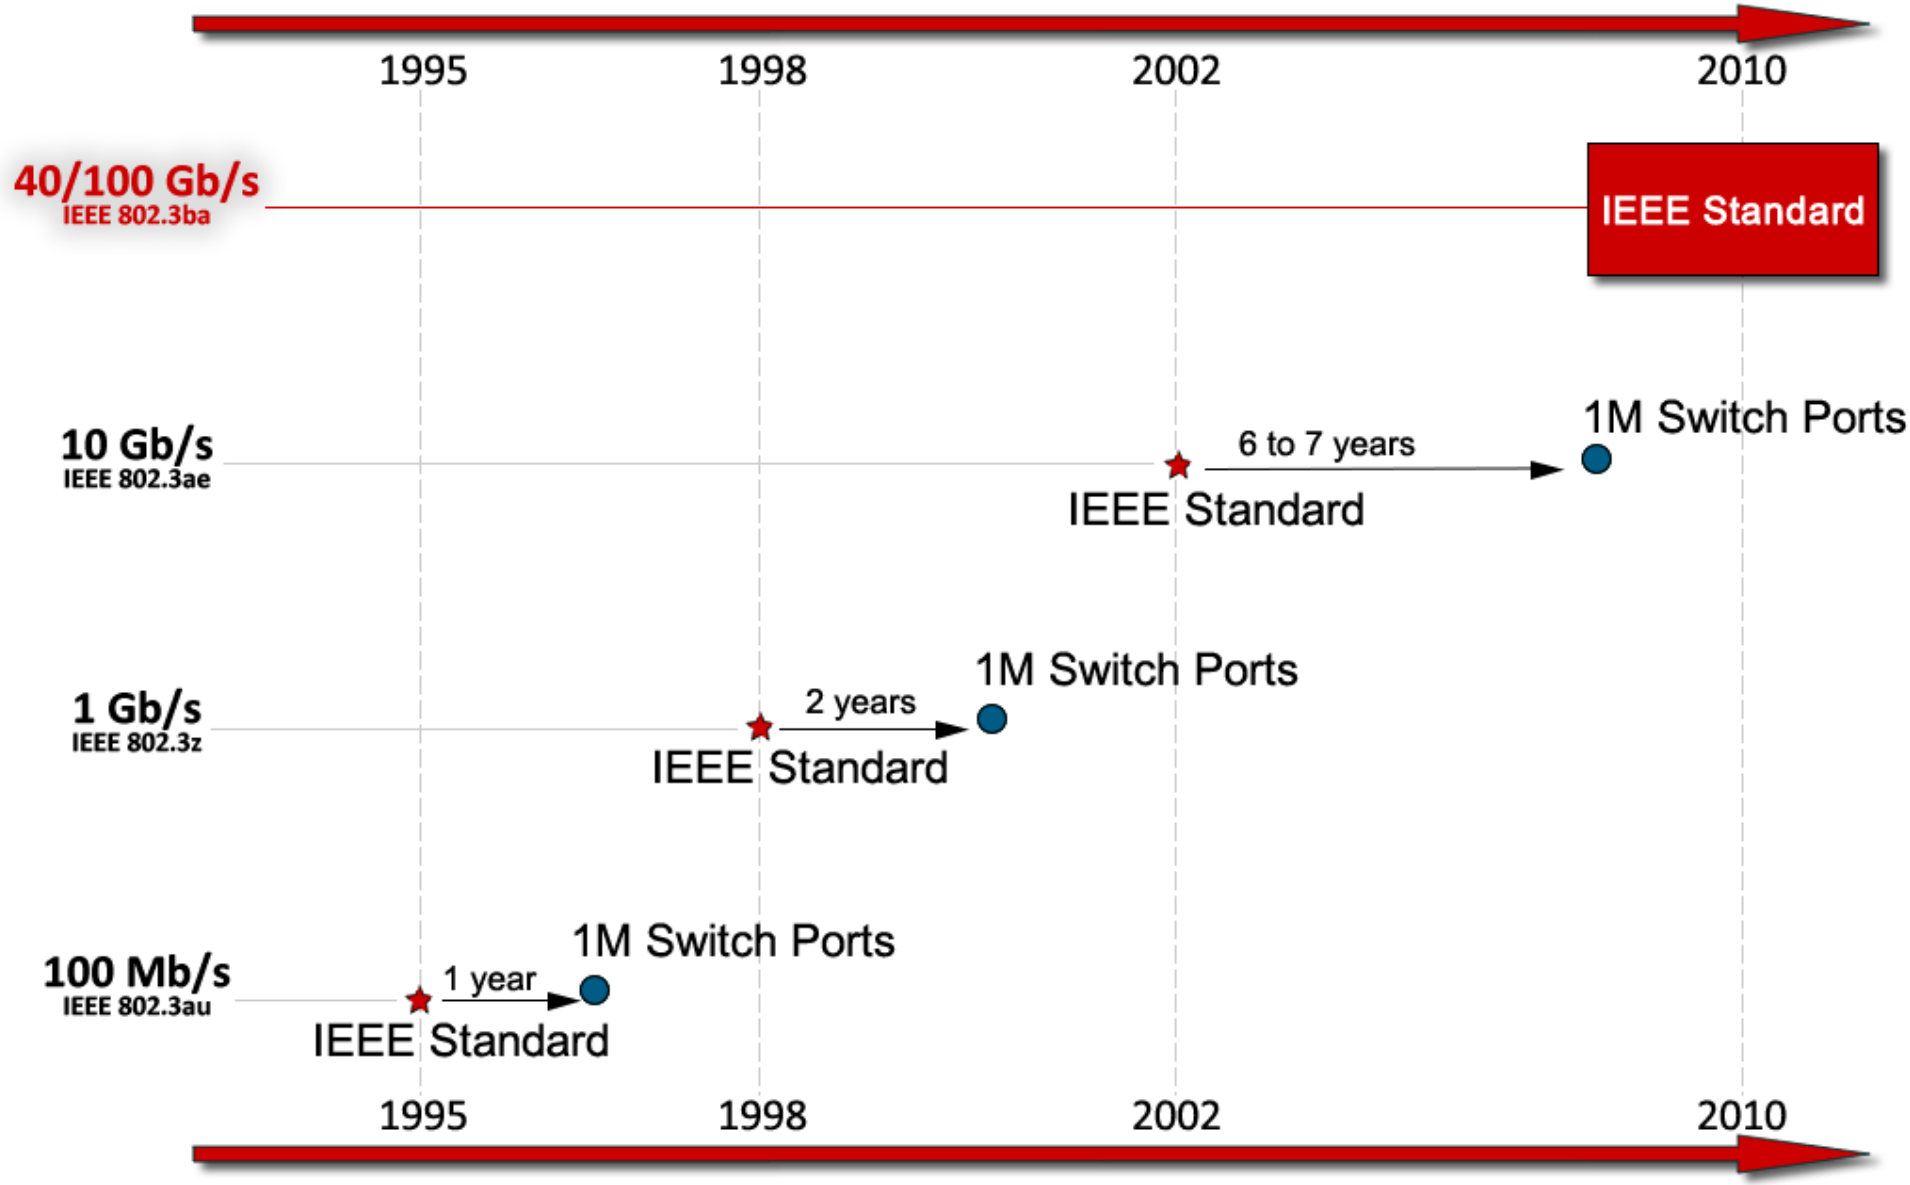

Standard Update

- **IEEE 802.3ba D3.0 released (It's a DONE DEAL)**

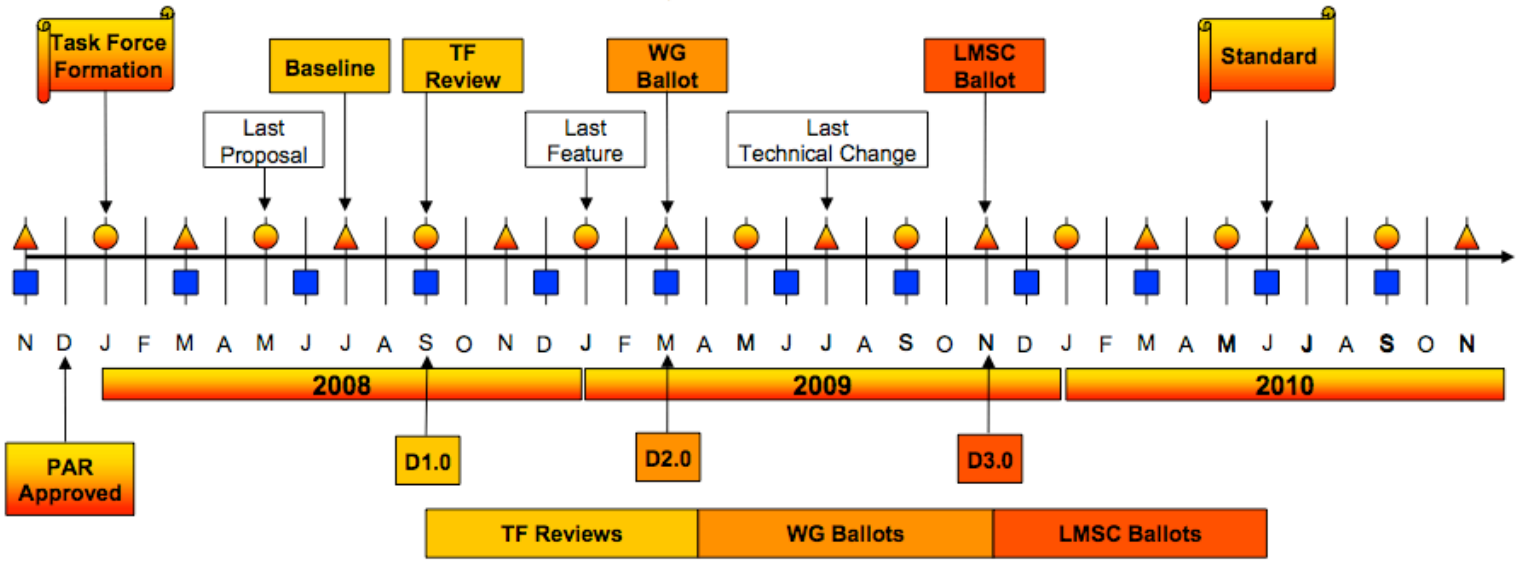

High Speed Ethernet Adoption on Servers

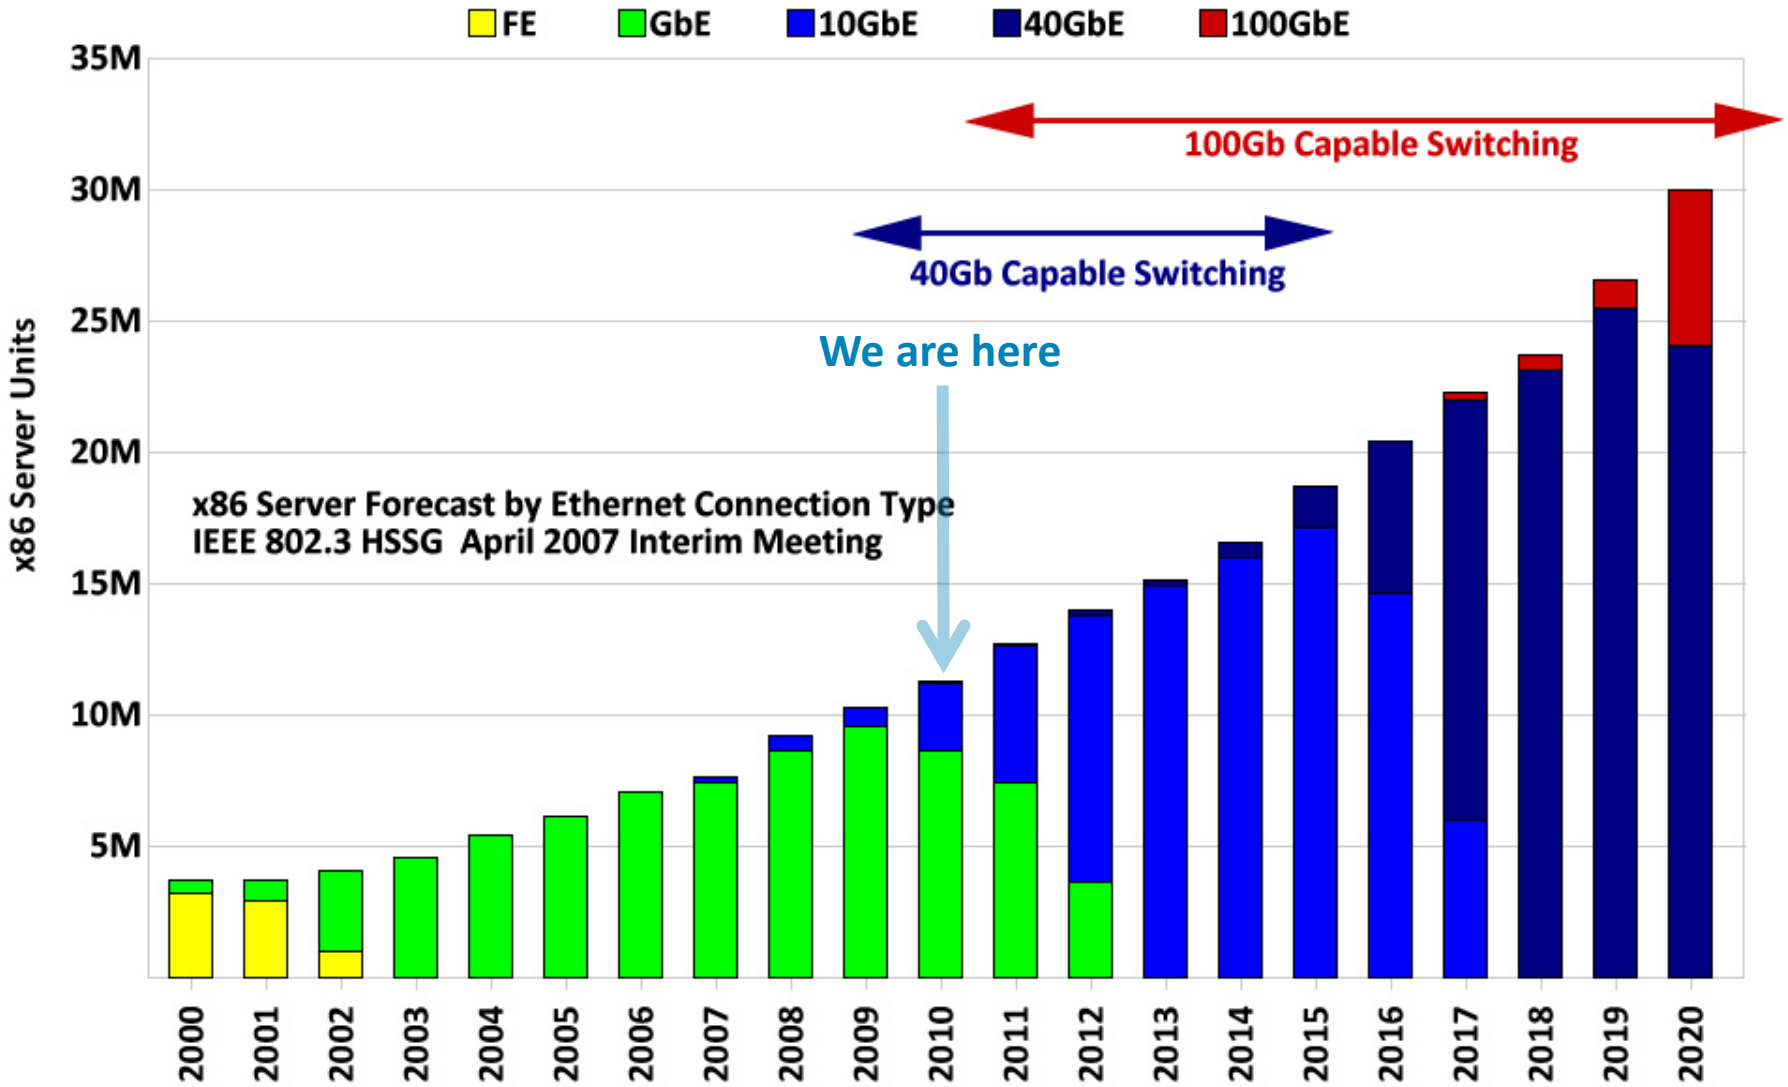

Data Center Deployments for 40G/100G

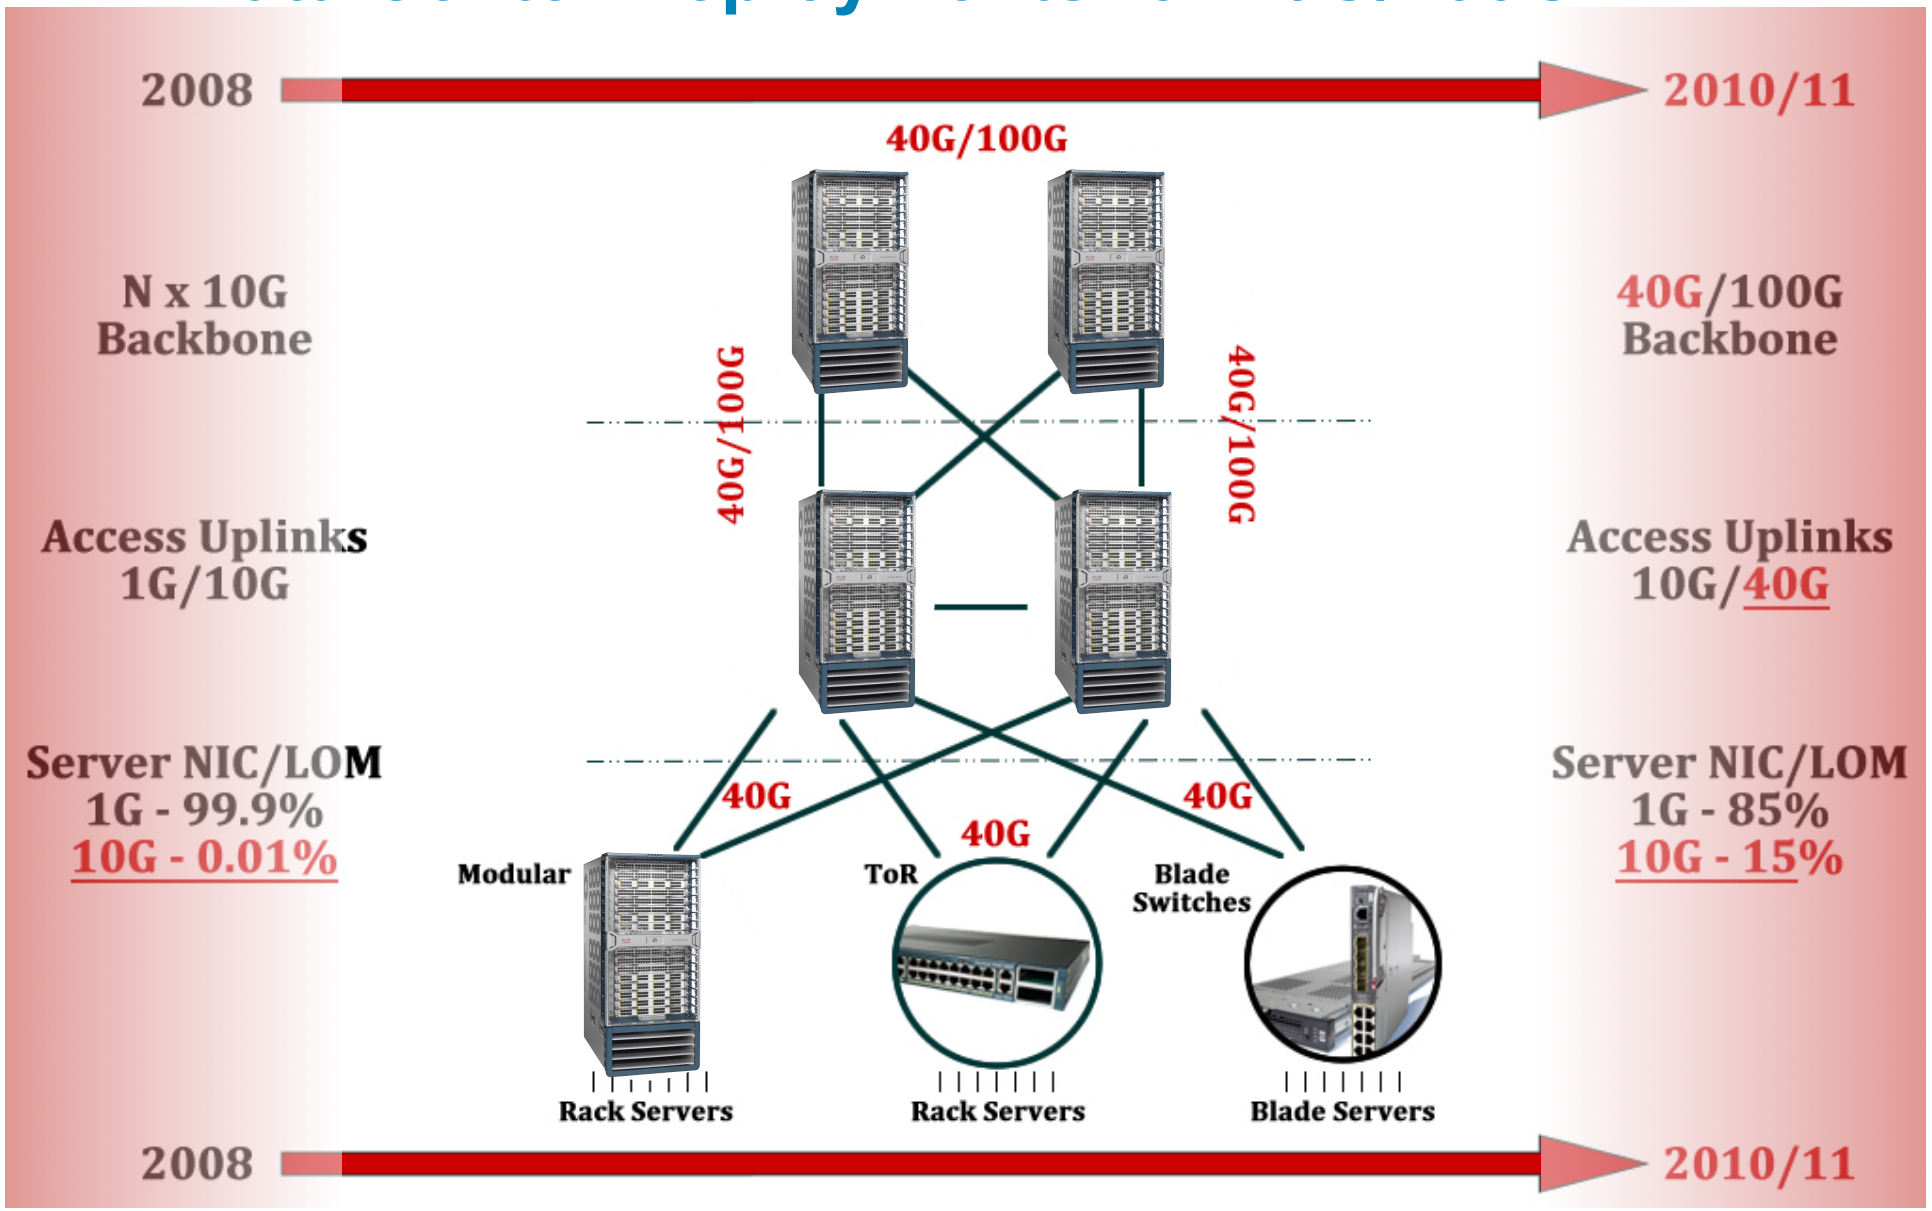

High Speed Ethernet Standard Interfaces

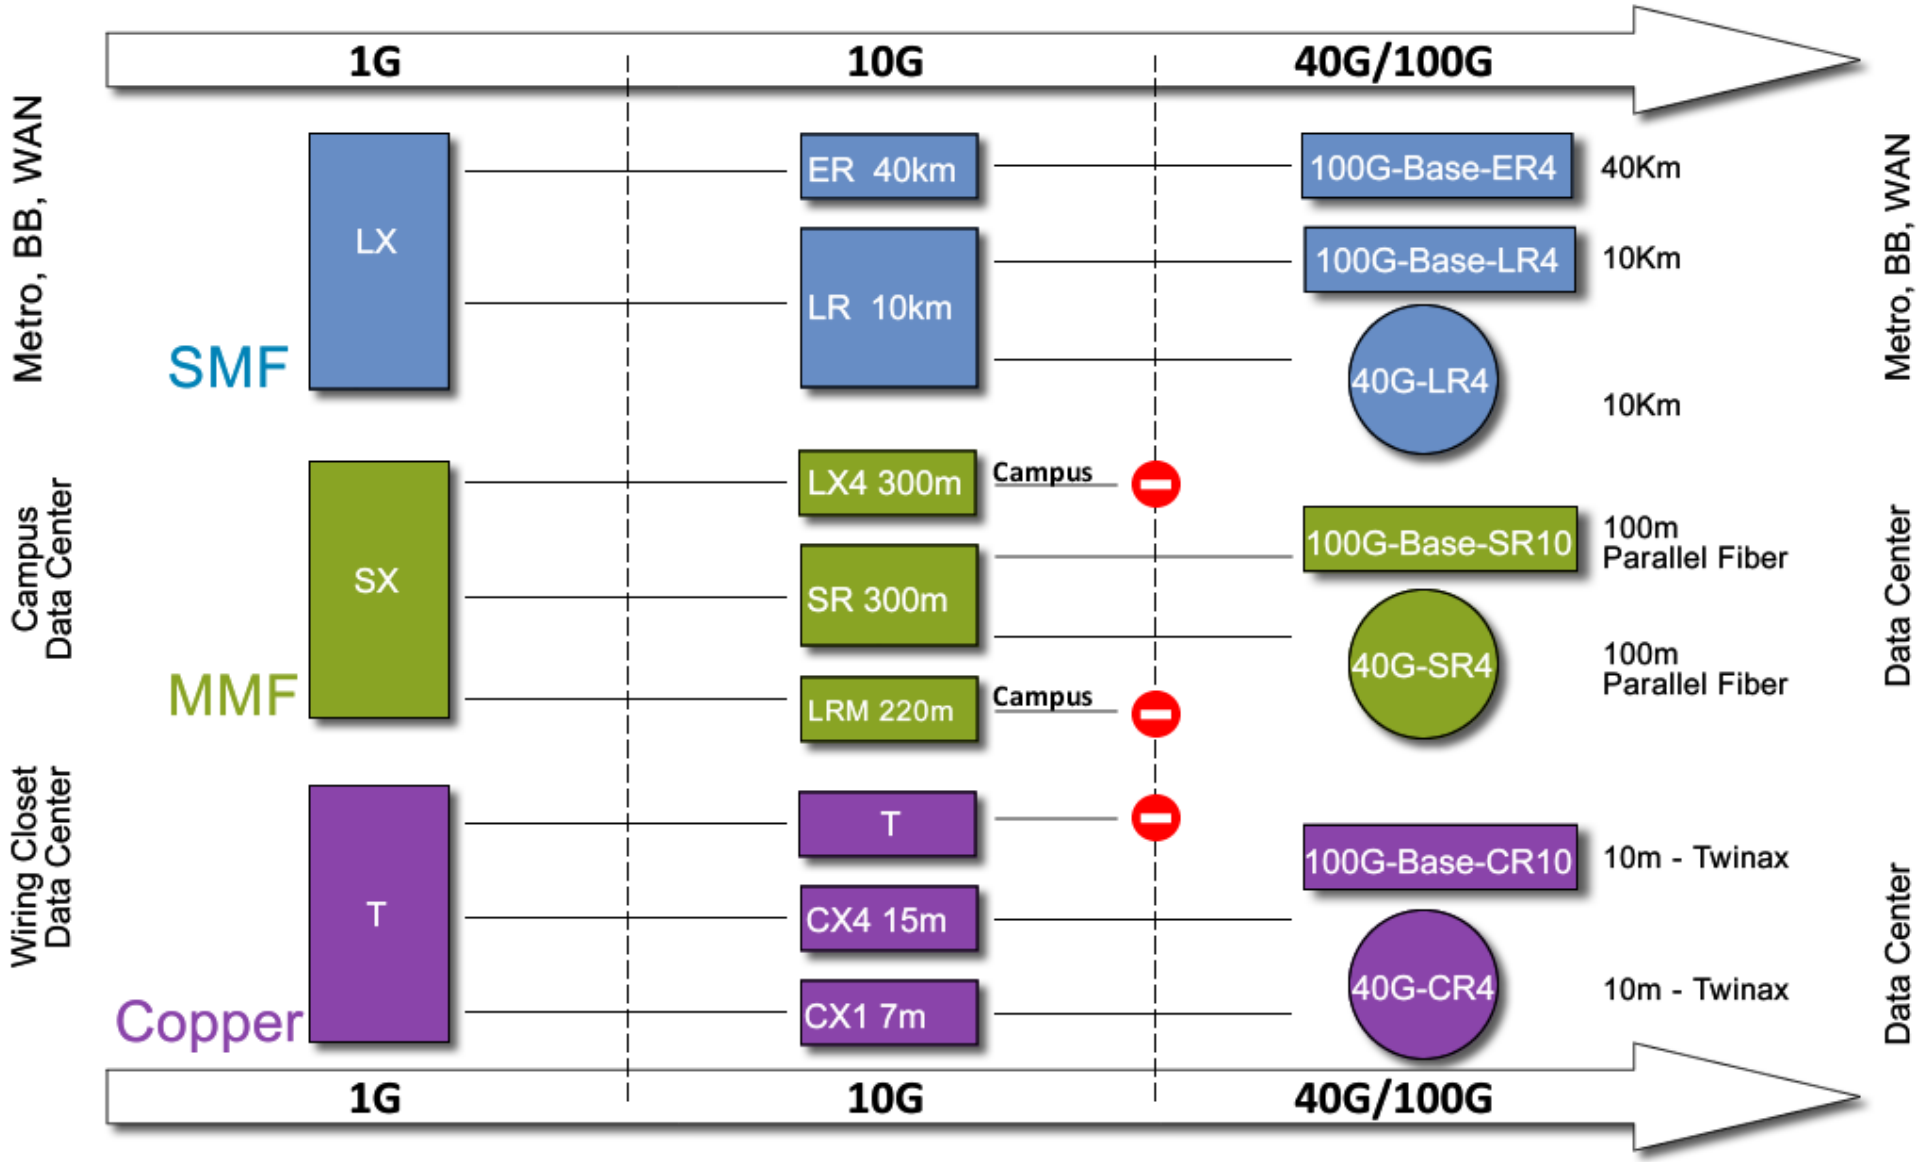

High Speed Ethernet – Not To Exceed Pricing

Estimated Not to Exceed List Prices (Industry-wide)

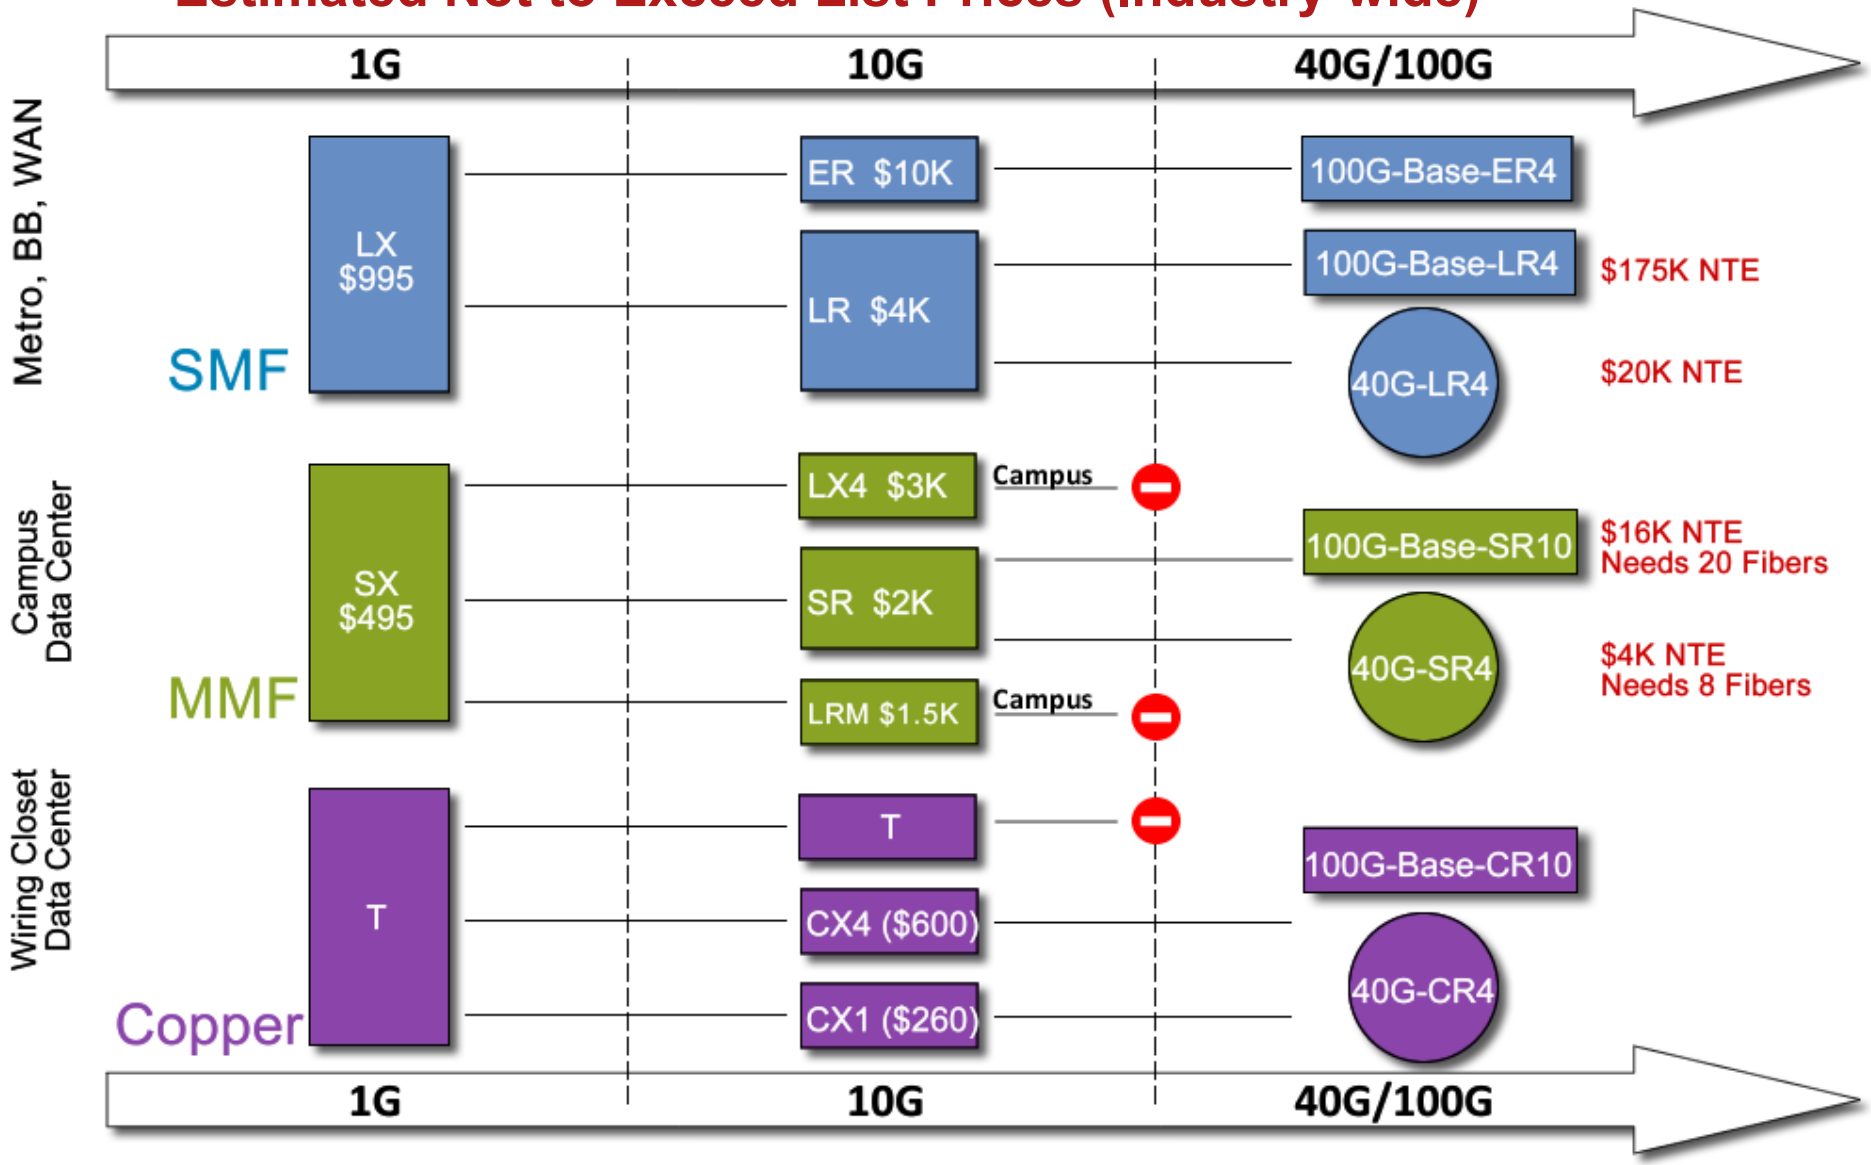

High-Speed Transceivers Form Factors

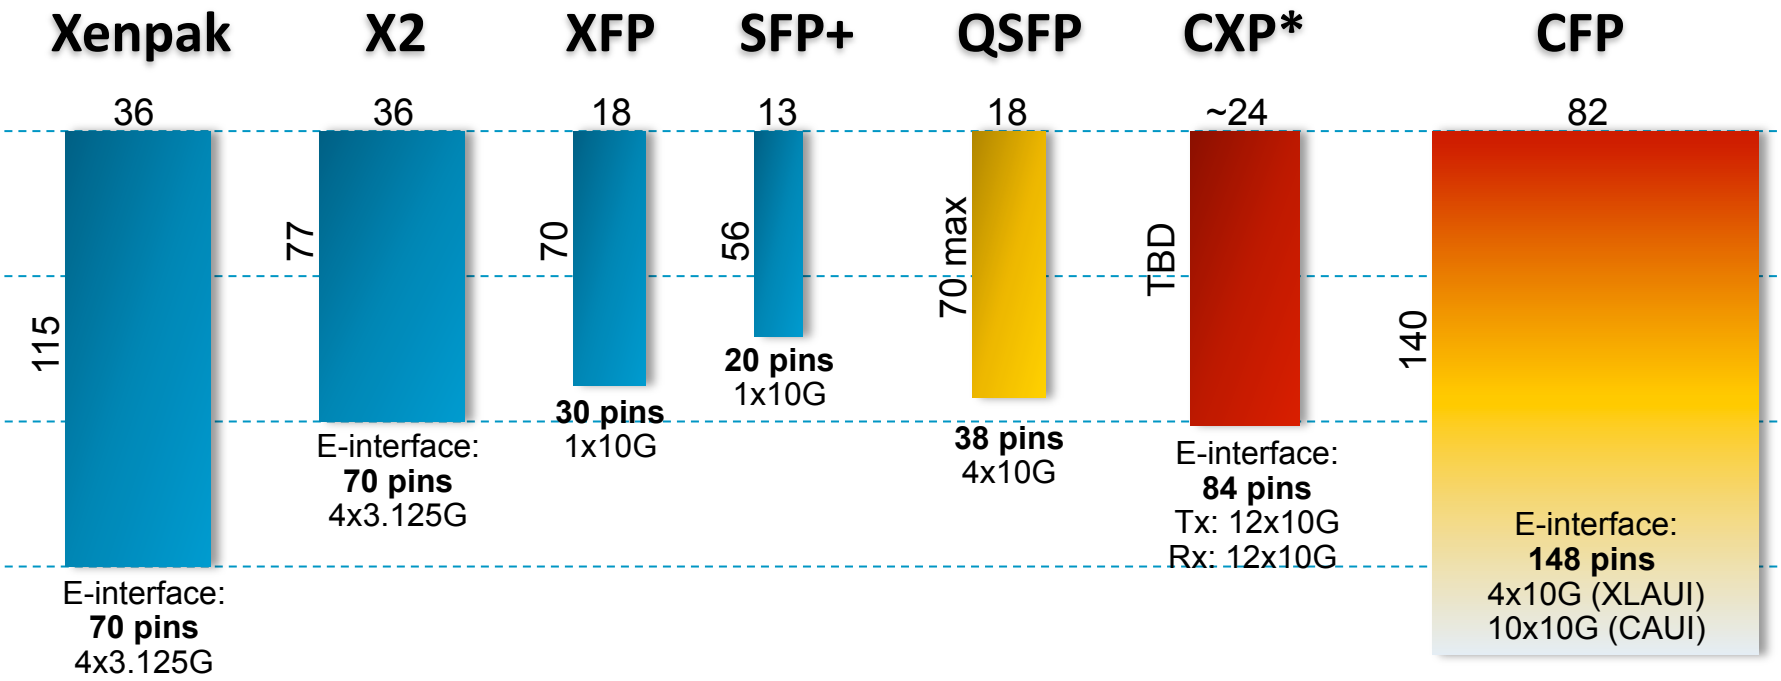

All units are in millimeters and round numbers

High-Speed Ethernet Transceiver Landscape

40G/100G CFP

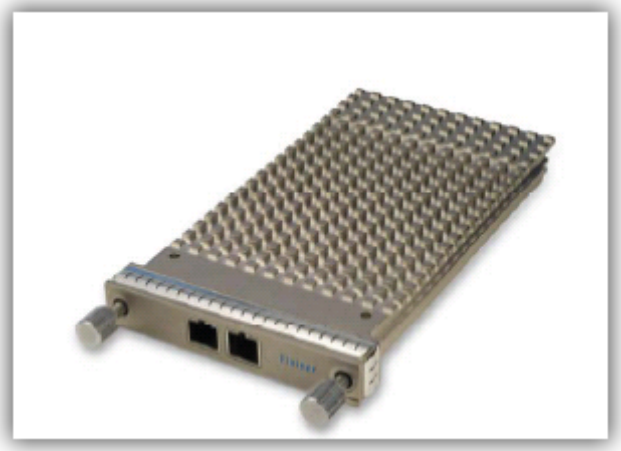

Applications:

Single Mode Fiber 10-40+Km
Multimode Parallel Fiber
Twinax Copper
Convertible in to 4x10GbE (SFP+)

Power Consumption:

Up to 8W @ 40GbE
Up to 25W @ 100GbE

40G QSFP

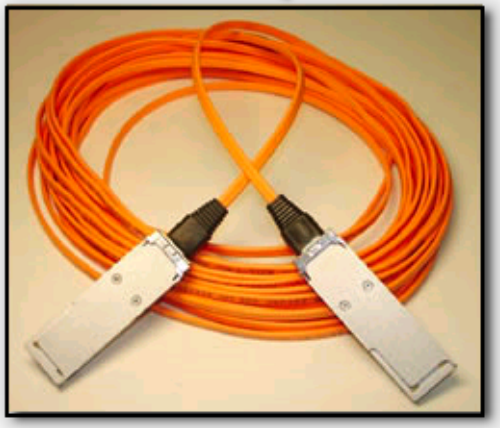

Applications:

Multimode Parallel Fiber
Twinax Copper
Future (?) 10 Km Single Mode

Power Consumption:

Up to 3.5W

High-Speed Ethernet Transceiver Landscape

100GbE CFP requires "Riding HeatSink" SMF optimized

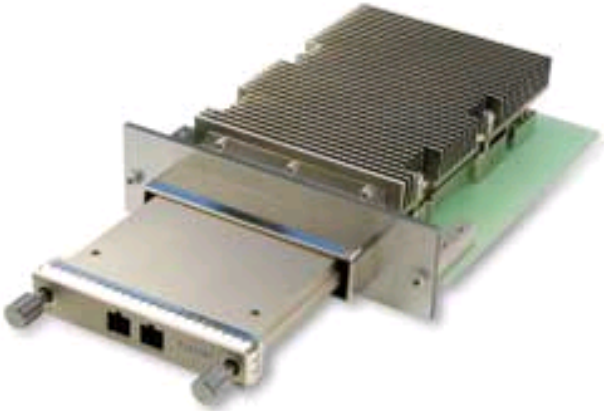

CFP features a new concept known as the riding heat sink, in which the heat sink is attached to rails on the host card and "rides" on top of the CFP, which is flat topped.

100GbE QXP MMF/Twinax optimized

CXP was created to satisfy the high-density requirements of the data center, targeting parallel interconnections for 12x QDR InfiniBand (120 Gbps), 100 GbE, and proprietary links between systems collocated in the same facility. The InfiniBand Trade Association is currently standardizing the CXP.

40G/100G Multimode OM3 Fiber Array Cables

- Custom-length cabling delivered with factory-installed connectors on both ends
- Cable is plugged into the back of patch panels. At 40GbE/ 100GbE it will plug directly into QSFP/CFP

MPO Plugs
(12-fiber array connectors)

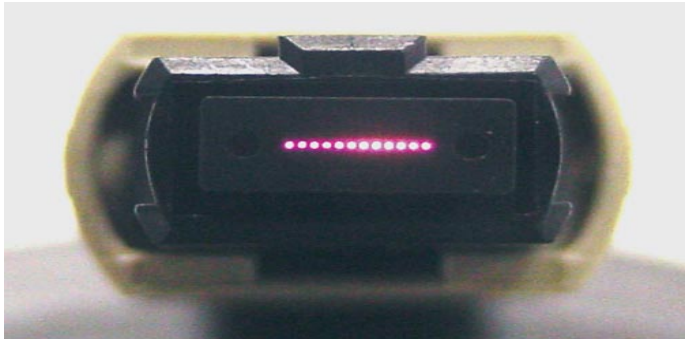

Agenda

- Overview of traditional network designs
- FEX Link
- Virtual Port Channel
- Overlay Transport Virtualization (OTV)
- Cisco FabricPath / Rbridges / TRILL
- Transceivers
- 40GbE and 100GbE
- **Management**