Cisco IOS Advantage Webinars

Efficient Data Center Design with FabricPath

Babi Seal and Patrick Warichet

We'll get started a few minutes past the top of the hour.

Note: you may not hear any audio until we get started.

Speakers & Panelists Introduction

Speakers

Babi Seal
Product Manager
sseal@cisco.com

Patrick Warichet
Technical Marketing Engineer
pwariche@cisco.com

Panelists

Francois Tallet
Product Manager Engineering
ftallet@cisco.com

Vimala Veerappan
Technical Marketing Engineering
vveerapan@cisco.com

Housekeeping

- Submit questions in Q&A panel and send to "All Panelists"
  Avoid CHAT window for better access to panelists

- For Webex audio, select COMMUNICATE > Join Audio Broadcast

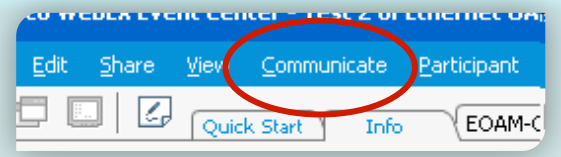

- For Webex call back, click ALLOW Phone button at the bottom of Participants side panel

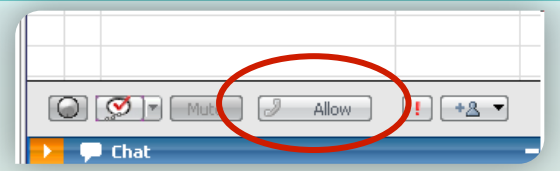

- Where can I get the presentation?

Or send email to: ask_iosadvantage@cisco.com

- Please complete the post-event Survey

- Join us on September 5 for our next IOS Advantage Webinar:
  A Closer Look: Comparing Benefits of EIGRP and OSPF
  www.cisco.com/go/iosadvantage

Agenda

- Why Layer 2 in the Data Center?
- FabricPath solution overview
- Inside the Fabric
- FabricPath Unicast Details
- FabricPath Multicast Details
- FabricPath designs

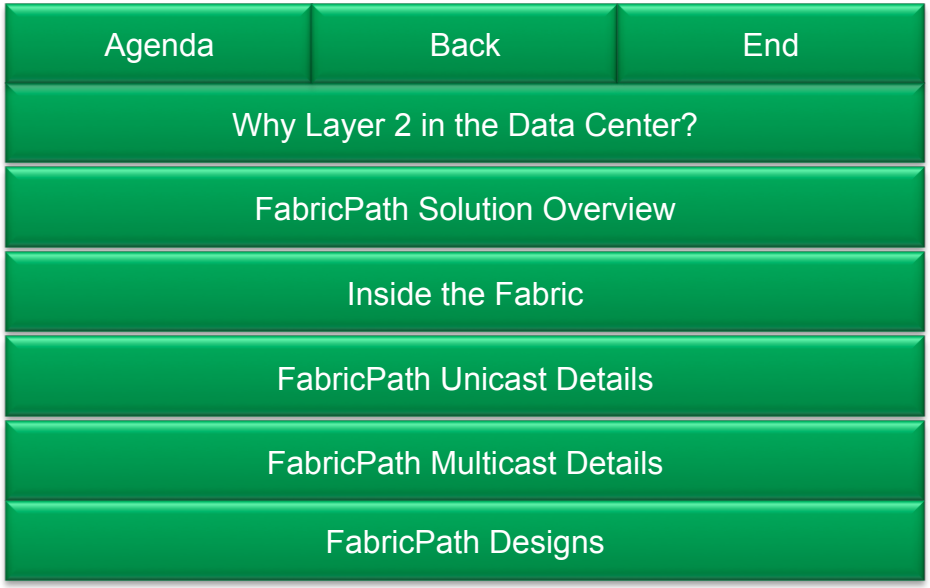

Why Layer 2 in the Data Center?

Why Layer 2 in the Data Center?

Key Decision Factors

- Some protocols rely on the functionality
- Simple, almost plug and play
- No addressing
- Required for implementing subnets
- Allows easy server provisioning
- Allows virtual machine mobility

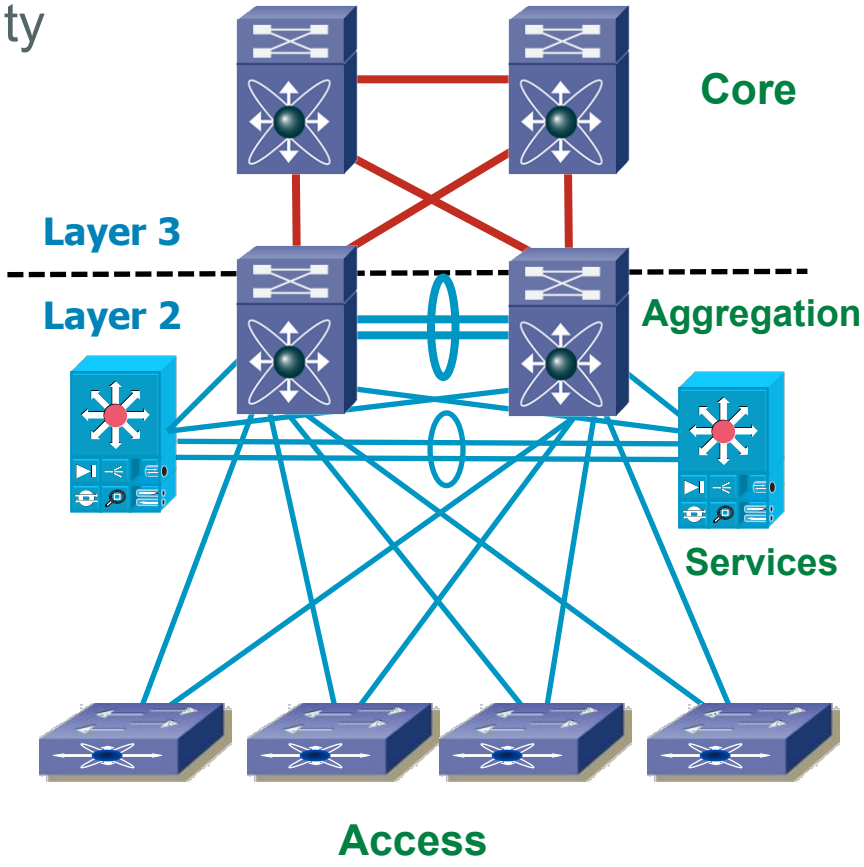

So what changed ?

Change # 1 – Workload Virtualization
Flexibility & Provisioning

- Partitioning
- Physical devices partitioned into virtual devices

- Clustering
- Applications distributed across multiple servers

Virtual Machines

Physical Servers

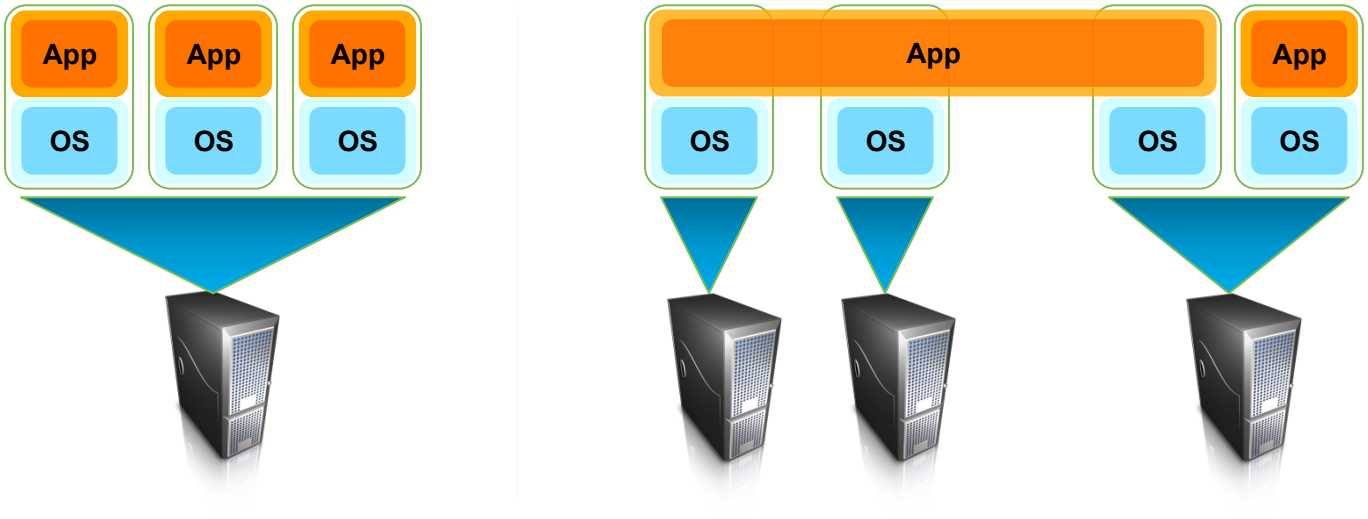

Change # 2 – Web 2.0 & Big Data
Big Data – Hadoop, NoSQL & HDFS

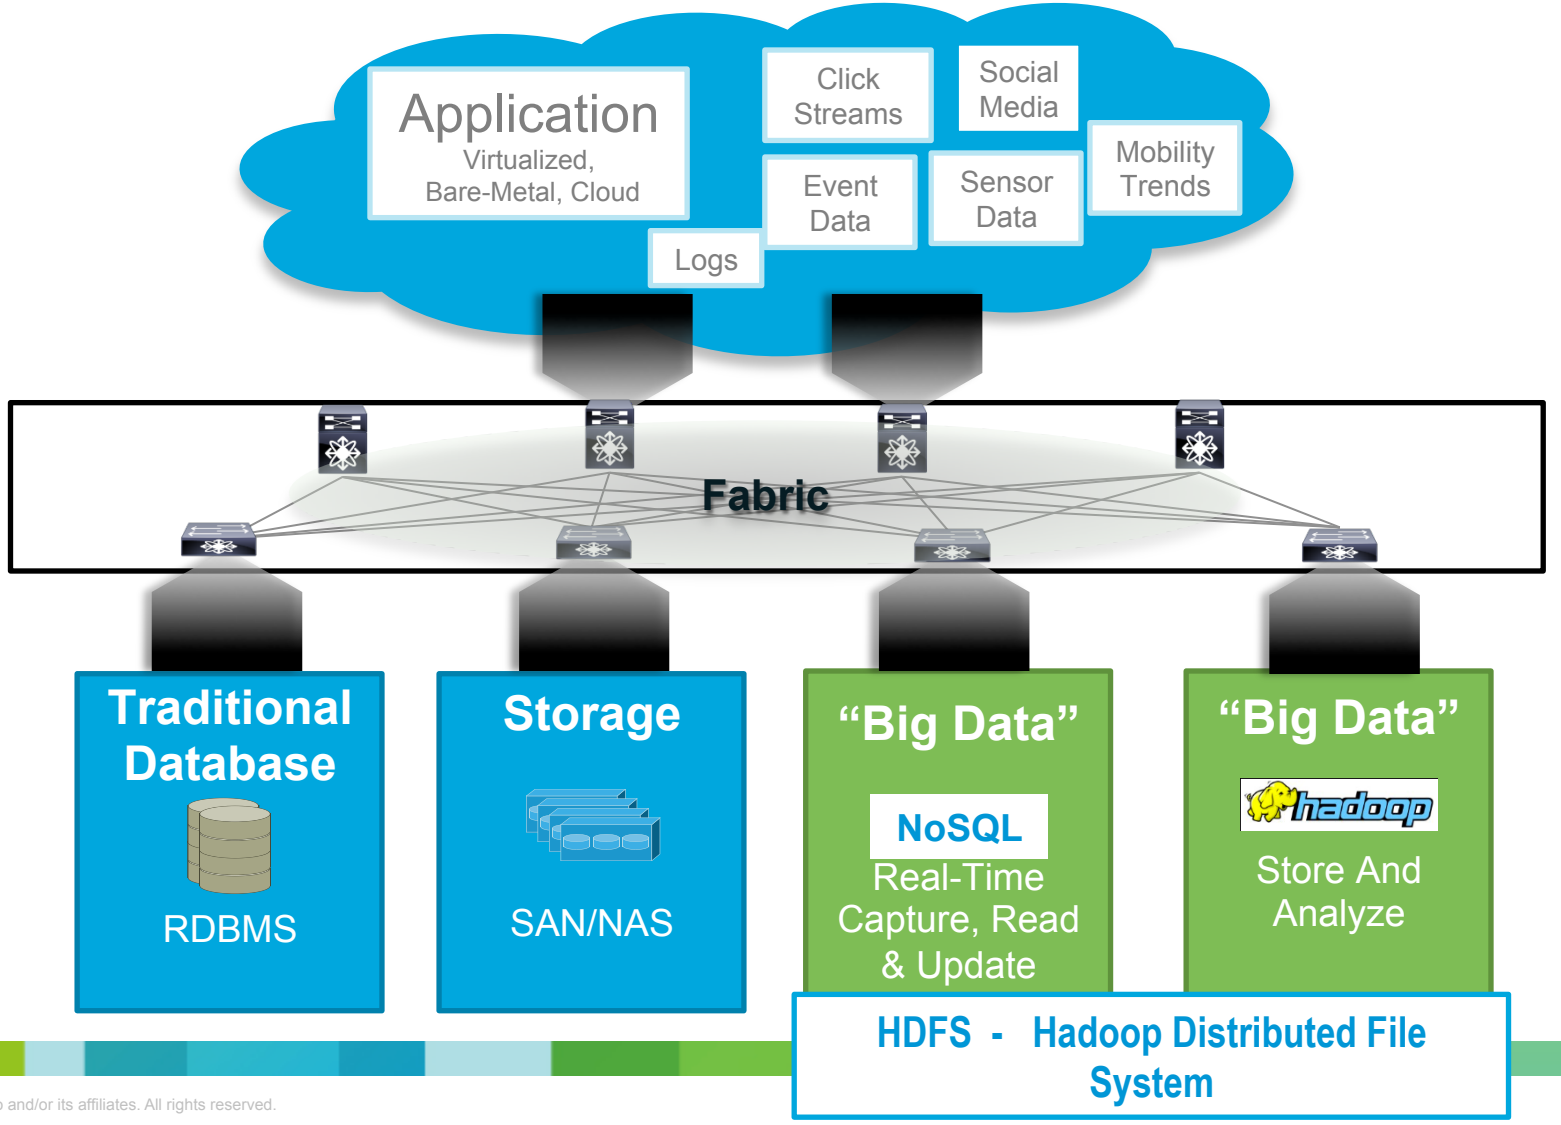

Possible Solution for End-to-End L2?

Just extend STP to the whole network

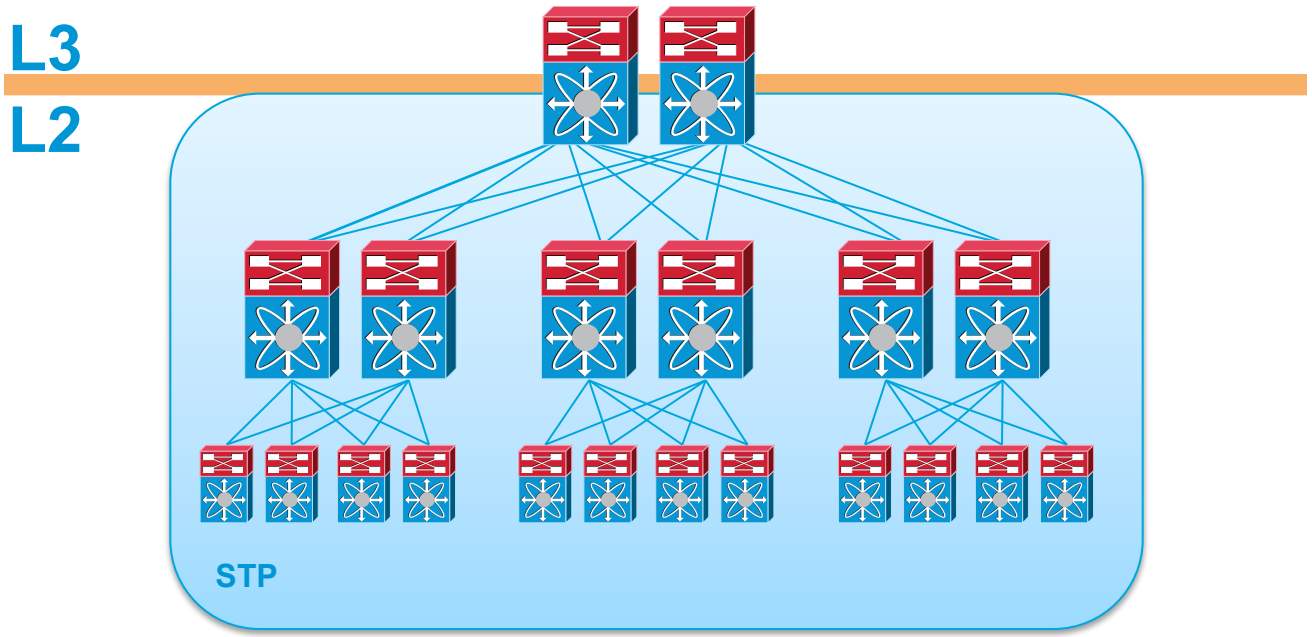

Some Layer 2 Limitations

- Local problems have network-wide impact, troubleshooting is difficult
- Tree topology provides limited bandwidth
- Tree topology introduces sub-optimal paths
- STP convergence is disruptive
- MAC address tables don’t scale
- Host flooding impacts the whole network

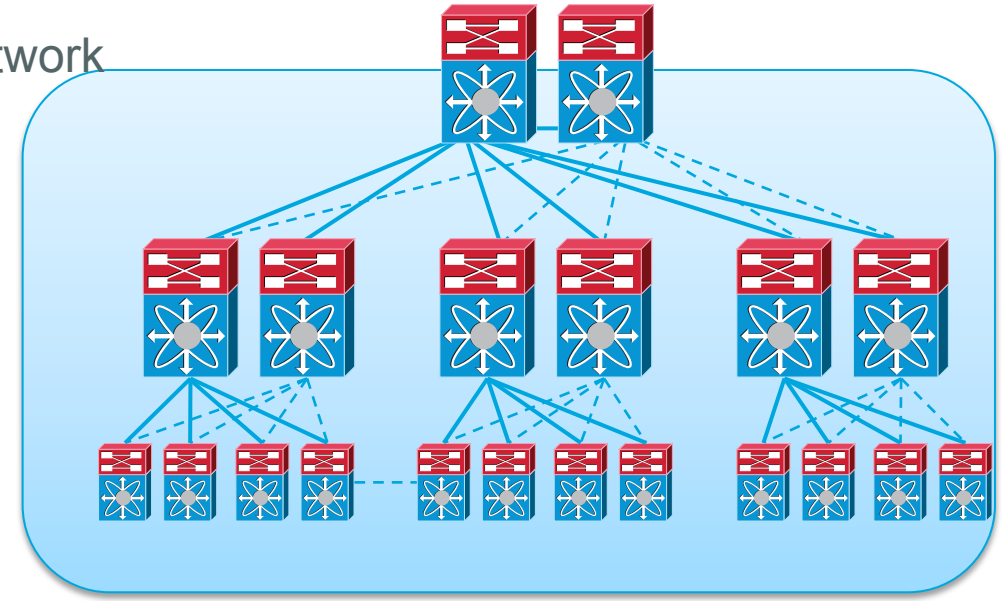

Introducing Cisco FabricPath

An NX-OS Innovation Enhancing L2 with L3

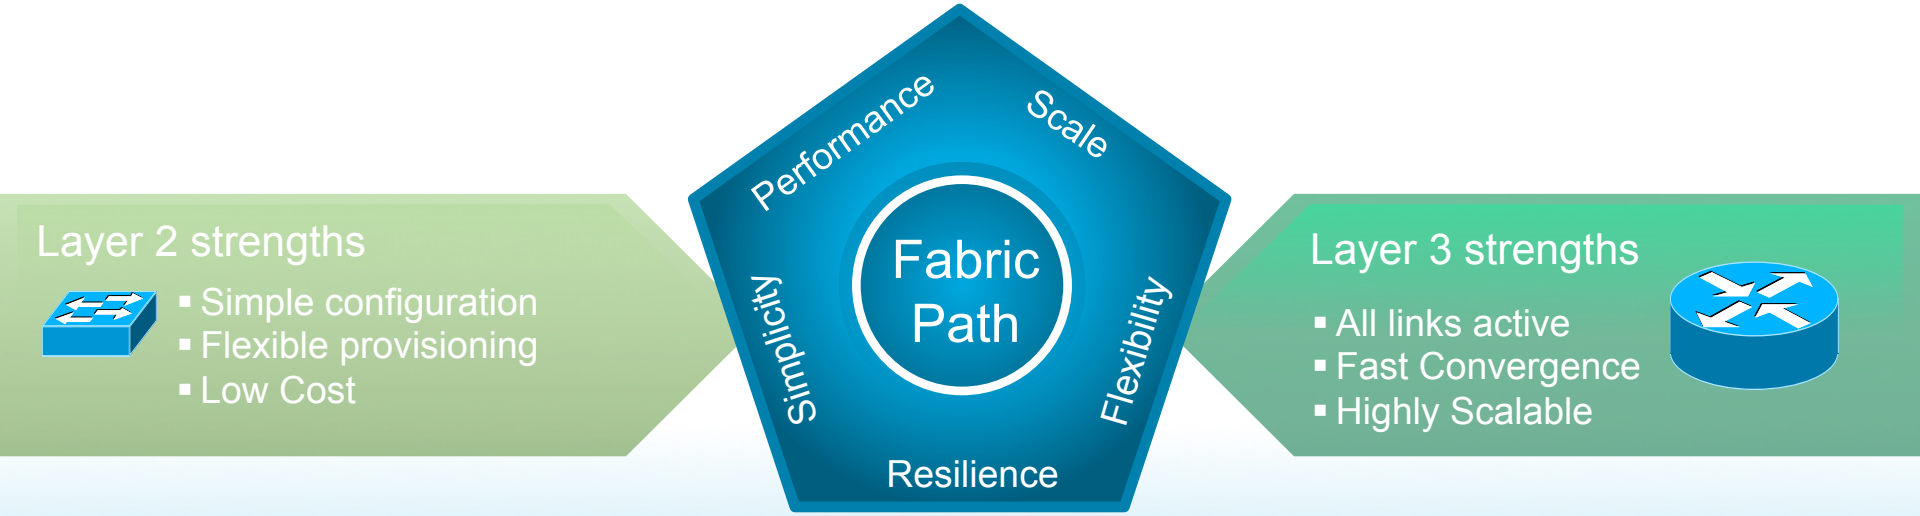

"The FabricPath capability within Cisco's NX-OS offers dramatic increases in network scalability and resiliency for our service delivery data center. FabricPath extends the benefits of the Nexus 7000 in our network, allowing us to leverage a common platform, simplify operations, and reduce operational costs."

Mr. Klaus Schmid, Head of DC Network & Operating,
T-Systems International GmbH

T · · Systems · · ·
Business flexibility

FabricPath
NX-OS Solution Roadmap

	CY2011	1H CY2012	1H CY2013

Nexus 7000	• Scale: 128 Switch IDs
	• FabricPath on F2 I/O modules
	  • Introducing native Fabric Extender support
	  • No conversational learning
		• vPC+ on Fabric Extender host interfaces
		• Conversational learning on F2
		• PVLAN on F2 (FabricPath and CE)
		• FabricPath Traceroute (PONG)
			• F1/M1 L2 proxy learning
			• **Multiple topologies**
			• Overload bit (for least disruptive convergence)
			• route-map and mesh groups for scale
			• Vlan pruning enhancement (based on CE forwarding ports)
			• Support on F2 10GbaseT I/O Modules
			• Affinity to the closest rooted tree for multidestination traffic
			• Anycast HSRP
			• Scale 4K vlans, 512 switch IDs

Nexus 5500	• Cisco FabricPath • 16-Way ECMP • 128 Switch IDs • STP Boundary & Termination • vPC+ (FabricPath Boundary) • L2 proxy learning • **PIM SSM over vPC+**	**SHIPPING NX-OS 5.2** • **Multi-Topology Support (2)**	• Anycast HSRP • Overload bit (for least disruptive convergence) • FabricPath Traceroute (PONG)

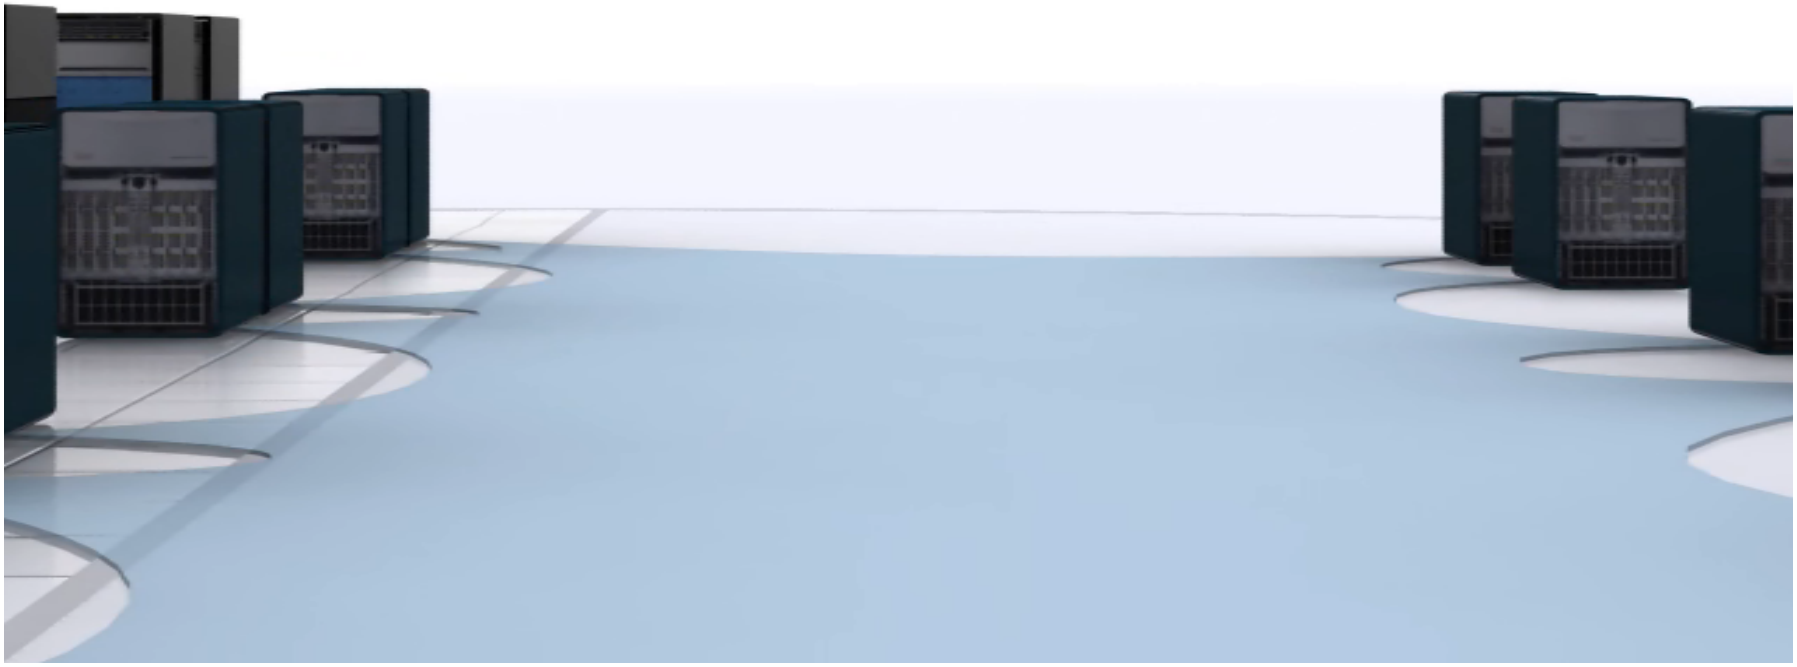

FabricPath Solution Overview

Agenda | Back | End

- Why Layer 2 in the Data Center?
- FabricPath Solution Overview
- Inside the Fabric
- FabricPath Unicast Details
- FabricPath Multicast Details
- FabricPath Designs

FabricPath, an Ethernet Fabric

Shipping Since 2010: turn your network into a Fabric

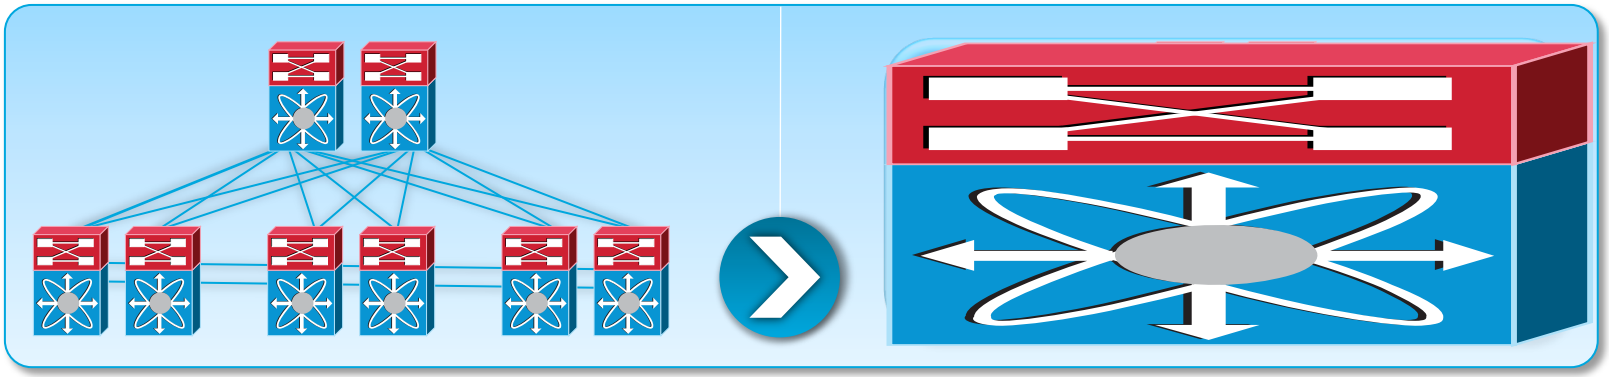

- Connect a group of switches using an **arbitrary** topology
- With a simple CLI, aggregate them into a Fabric:

```
N7K(config)# interface ethernet 1/1
N7K(config-if)# switchport mode fabricpath
```

- An open protocol based on Layer 3 technology provides Fabric-wide intelligence and ties the elements together

Optimal, Low Latency Switching

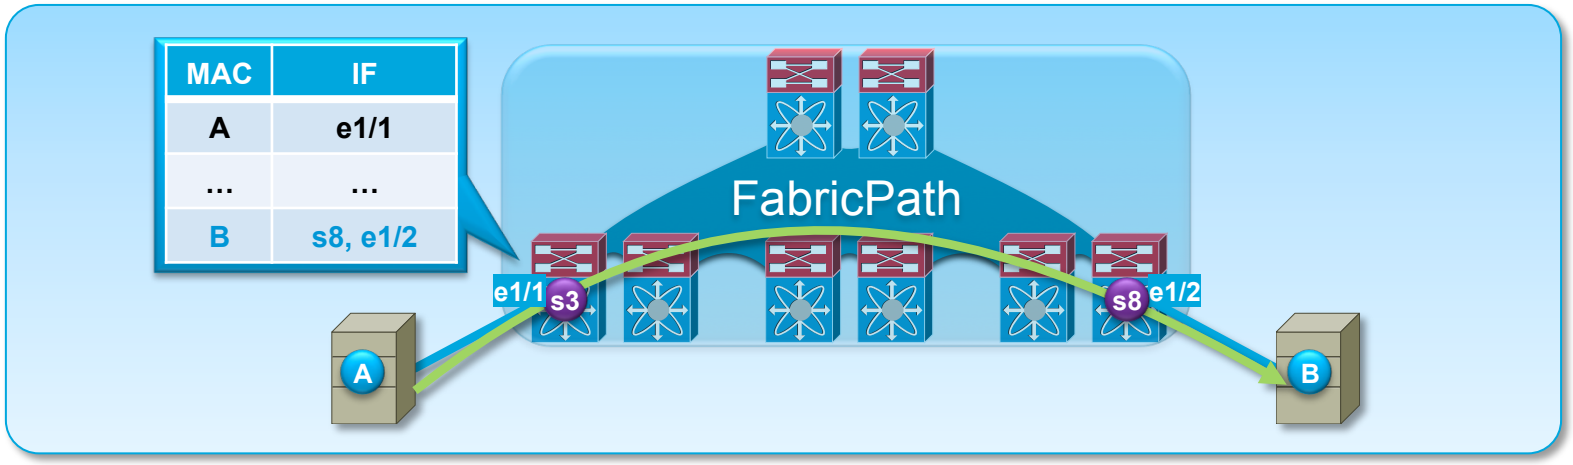

- Single address lookup at the ingress edge identifies the exit port across the fabric
- Traffic is then switched using the shortest path available
- Reliable L2 connectivity any to any
  (as if it was the same switch, **no STP inside**)

High Bandwidth, High Resiliency

Equal Cost MultiPathing (ECMP)

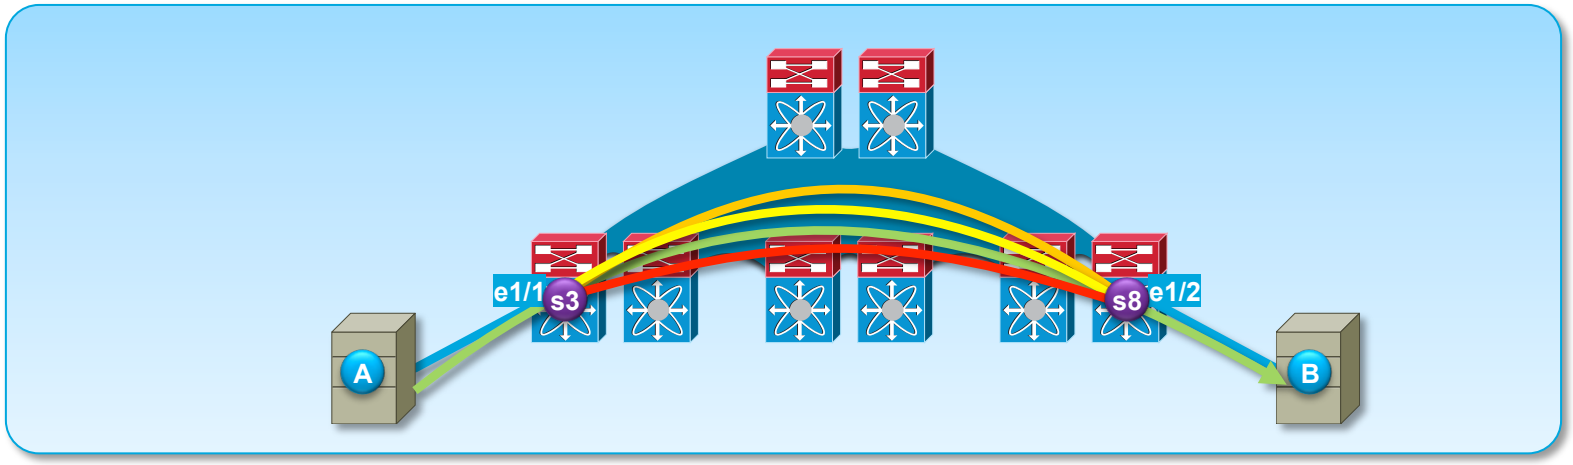

- Mutipathing (up to 256 links active between any 2 devices)
- Traffic is redistributed across remaining links in case of failure, providing fast convergence

Mac Address Table Scale

Conversational Learning

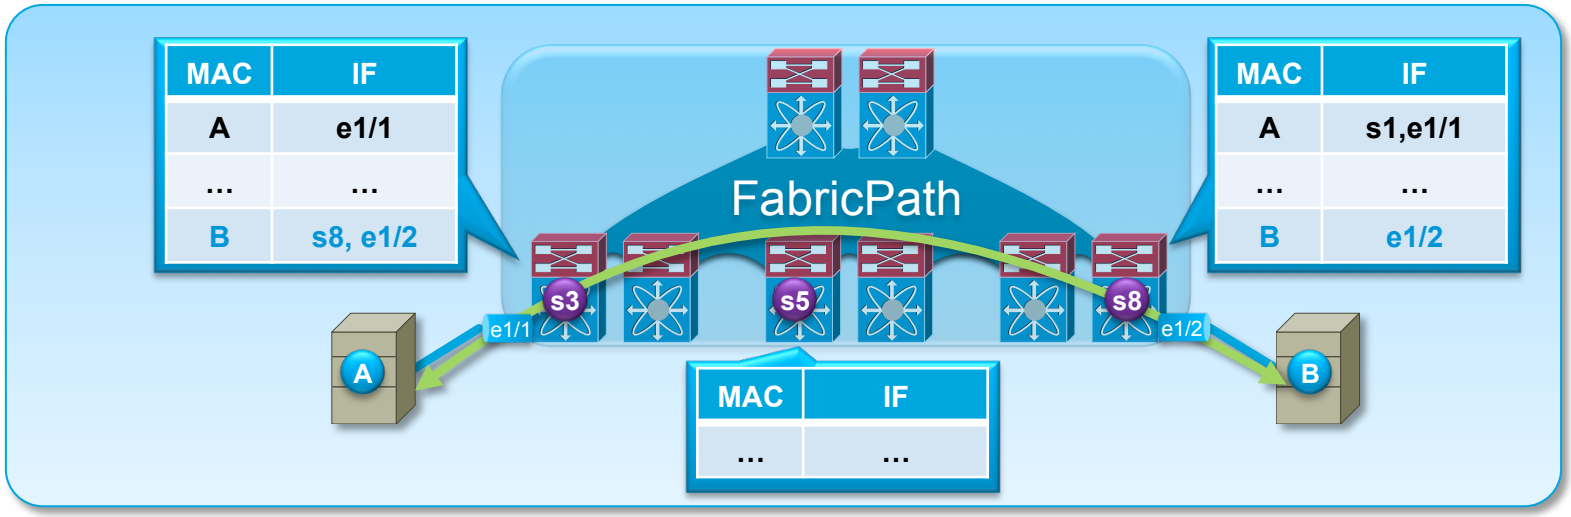

- Per-port mac address table only needs to learn the peers that are reached across the fabric

A virtually unlimited number of hosts can be attached to the fabric

Layer 2 integration

vPC+

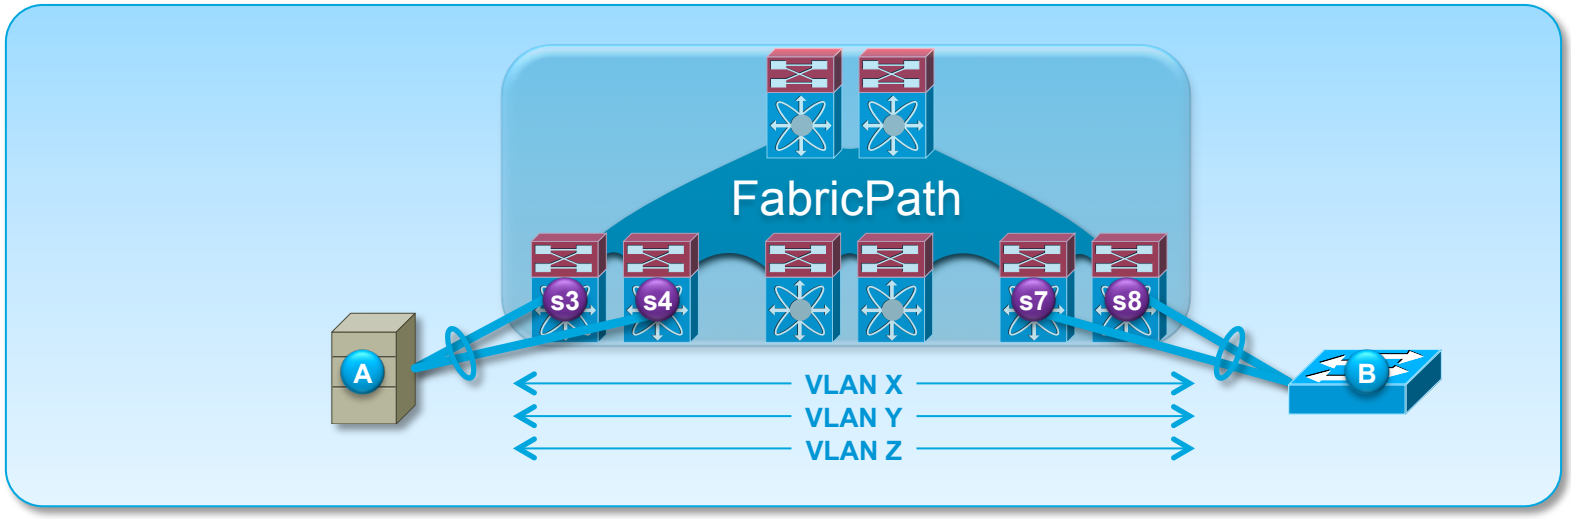

- Allows extending vlans with no limitation (no risks of loop)
- Devices can be attached active/active to the fabric using IEEE standard port channels and without resorting to STP

Edge Devices Integration

Hosts Can Leverage Multiple Active L3 Default Gateways

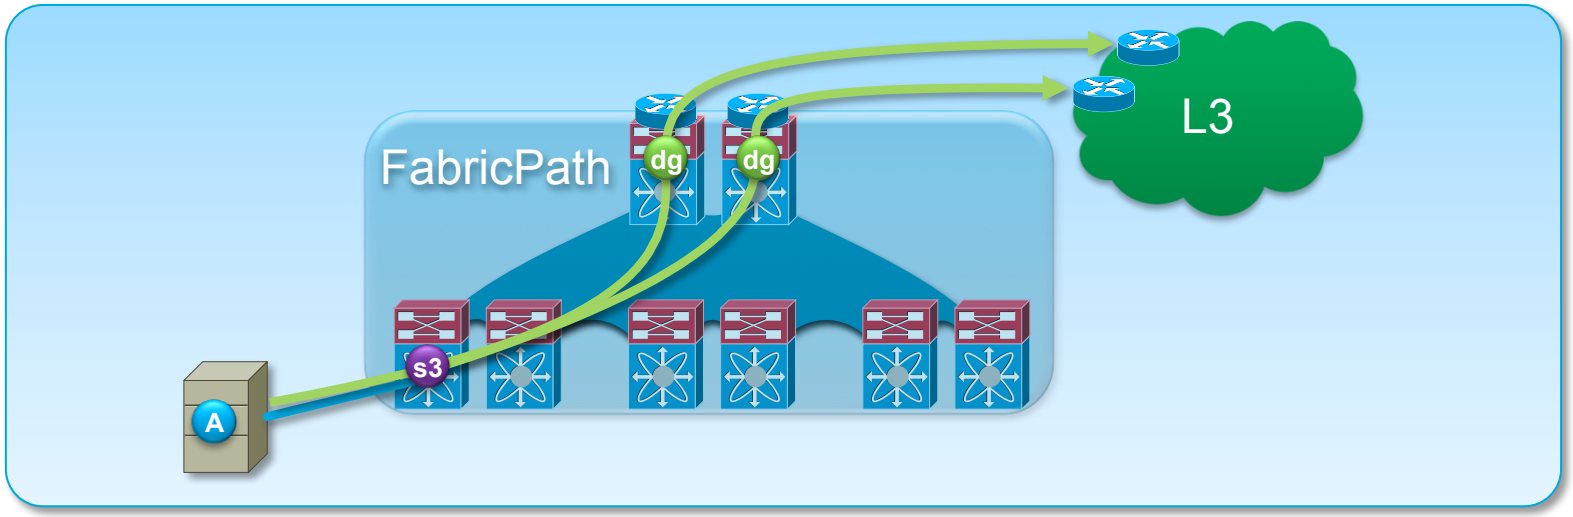

- Hosts see a single default gateway
- The fabric provide them transparently with multiple simultaneously active default gateways
- Allows extending the multipathing from the inside to the fabric to the L3 domain outside the fabric

Layer 3 Integration

Large tables (XL Hardware), SVIs Anywhere

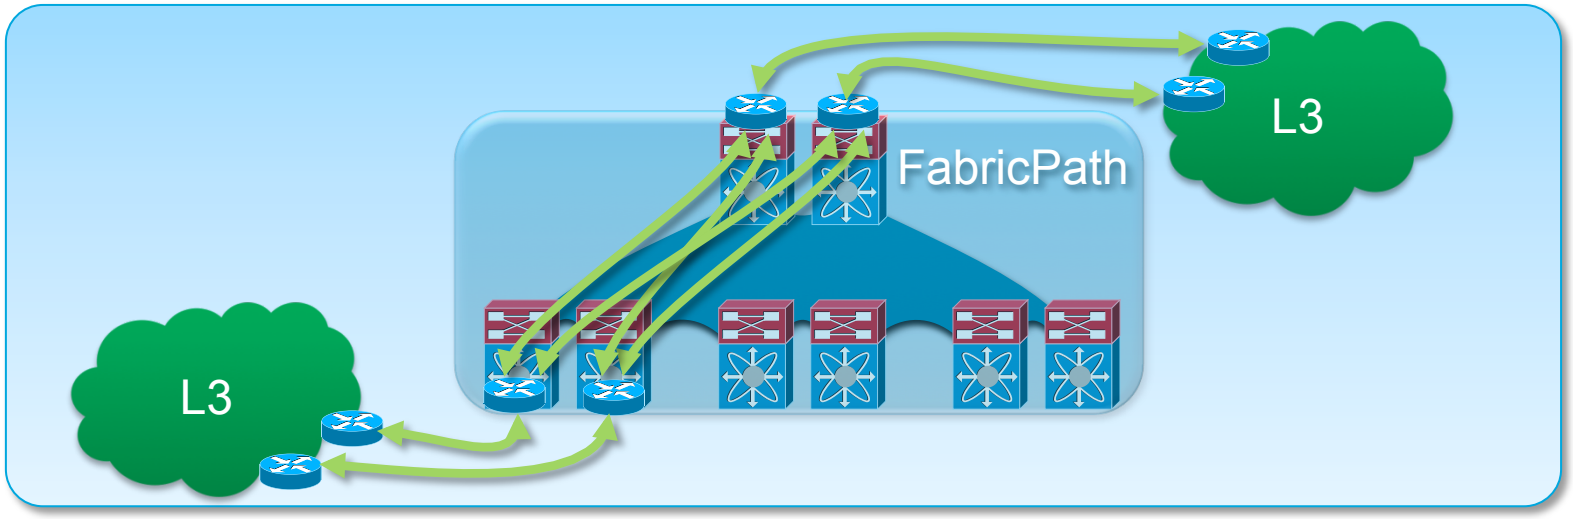

- The fabric provides seamless L3 integration
- An arbitrary number of routed interfaces can be created at the edge or within the fabric
- Attached L3 devices can peer with those interfaces
- The hardware is capable of handling million of routes

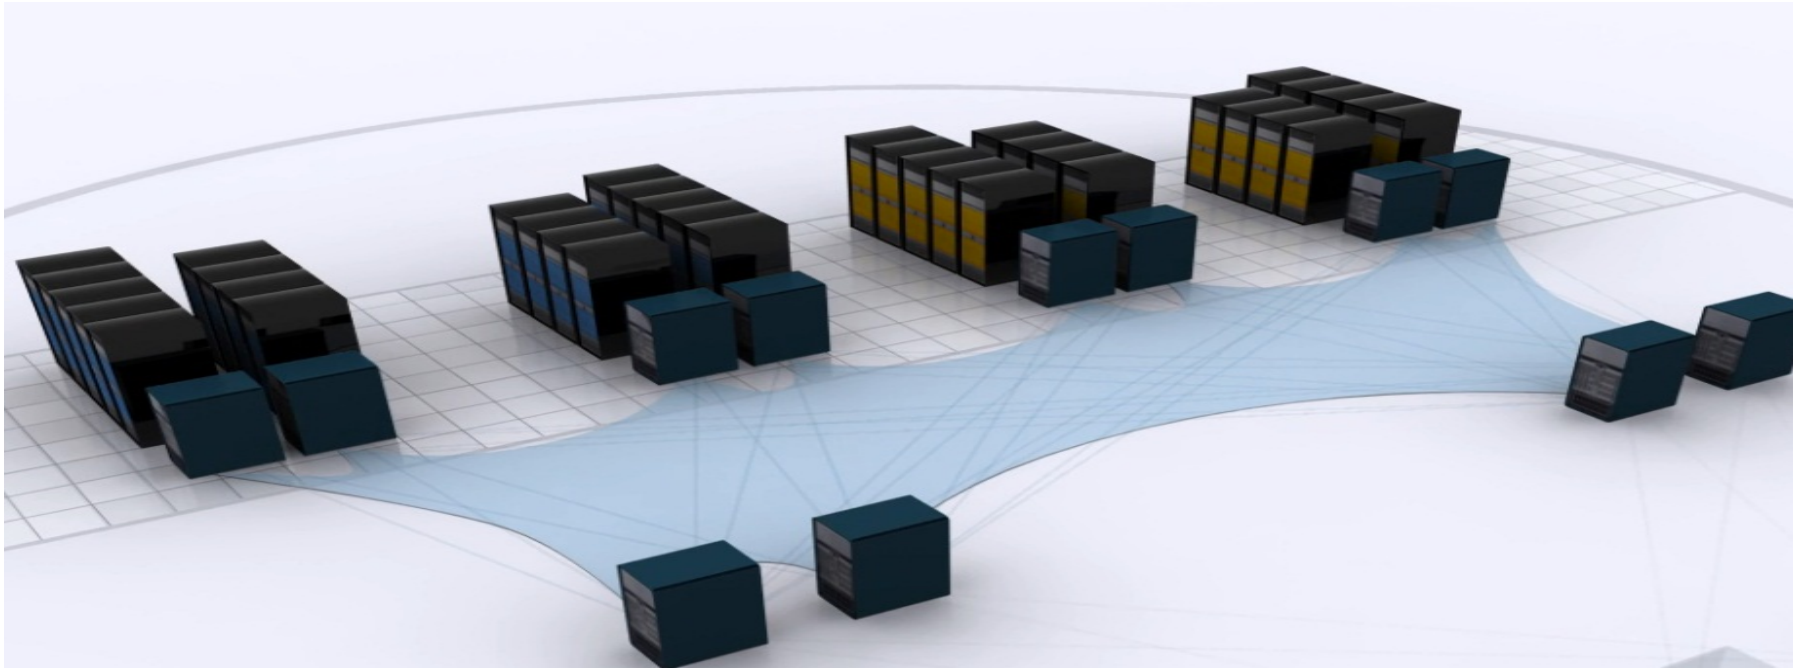

Inside the Fabric

FabricPath & Clos Fabric Terminology

- Send/receive FabricPath frame
- No STP, no MAC learning, no MAC address table
- Using a routing table computed by IS-IS

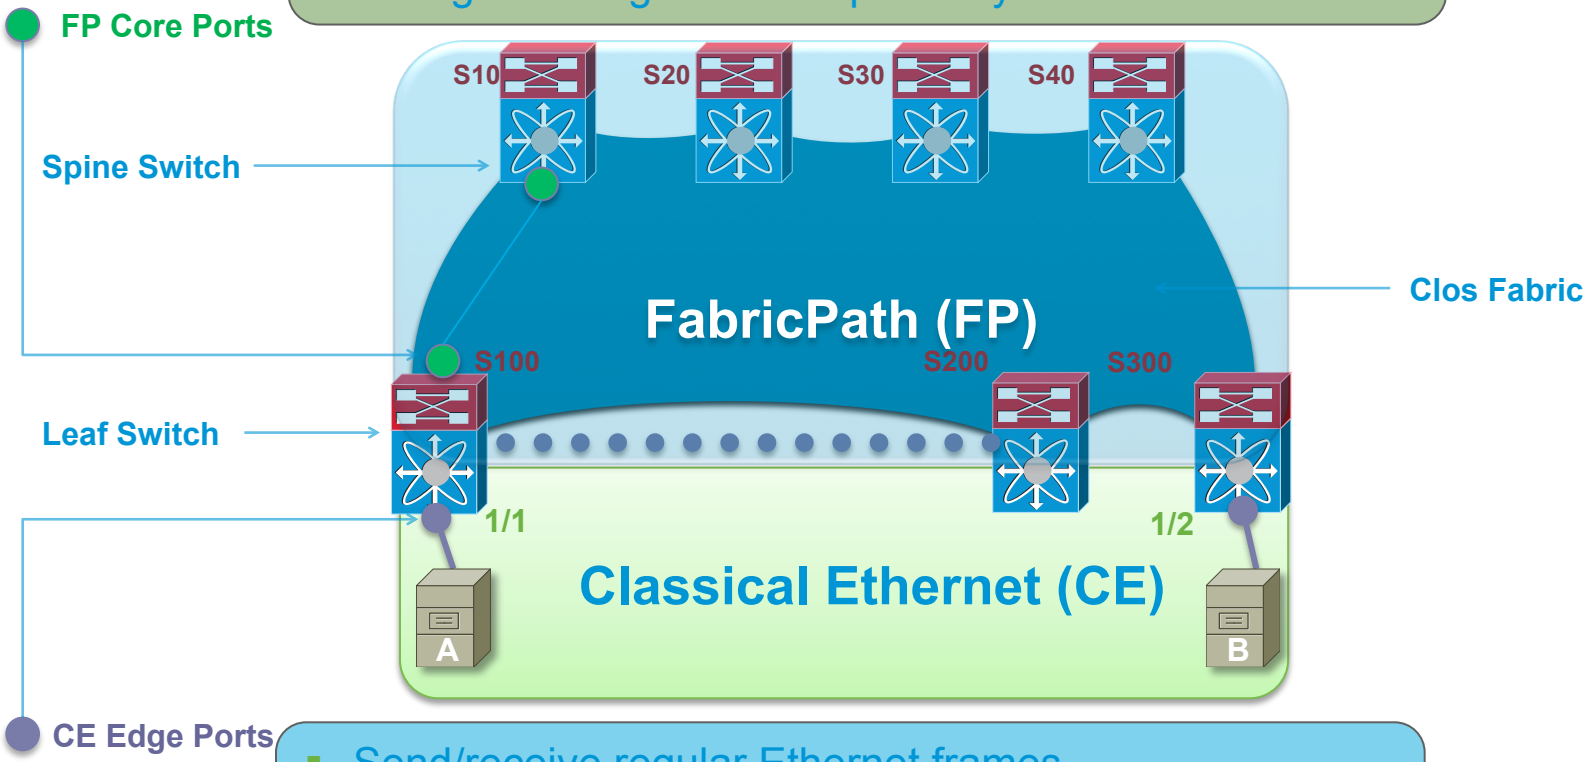

- Send/receive regular Ethernet frames
- Run STP, do MAC address learning using a MAC address table

New Control Plane

Plug-n-Play L2 IS-IS Manages Forwarding Topology

- IS-IS assigns addresses to all FabricPath switches automatically
- Compute shortest, pair-wise paths
- Support equal-cost paths between any FabricPath switch pairs

FabricPath Routing Table

Switch	IF
S10	L1
S20	L2
S30	L3
S40	L4
S200	L1, L2, L3, L4
…	…
S400	L1, L2, L3, L4

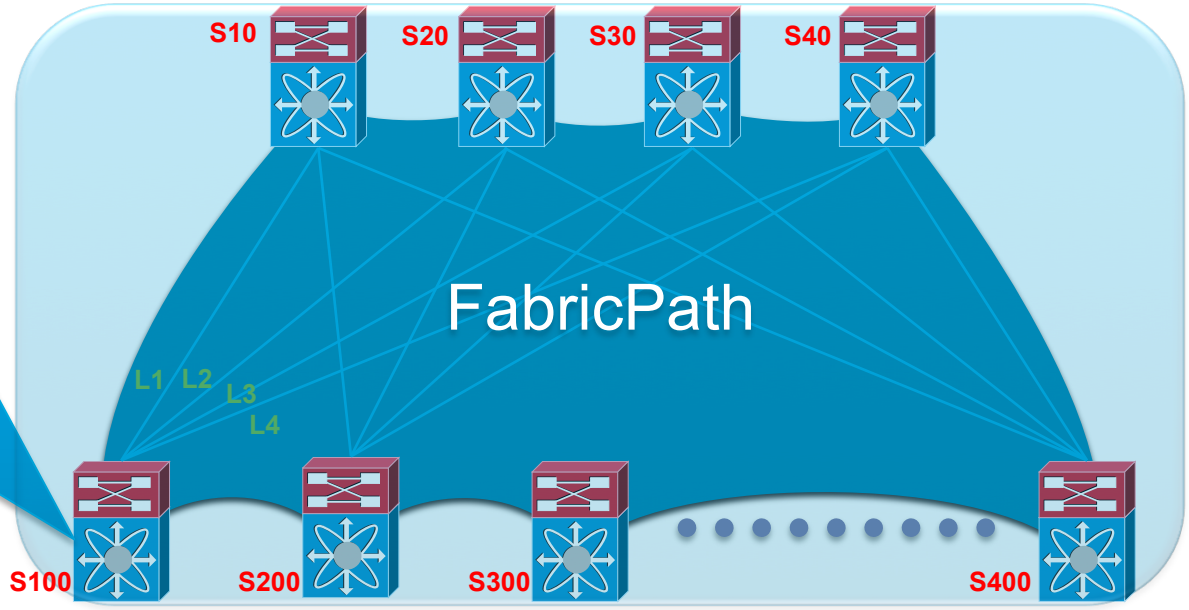

New Data Plane

- The association MAC address/Switch ID is maintained at the edge

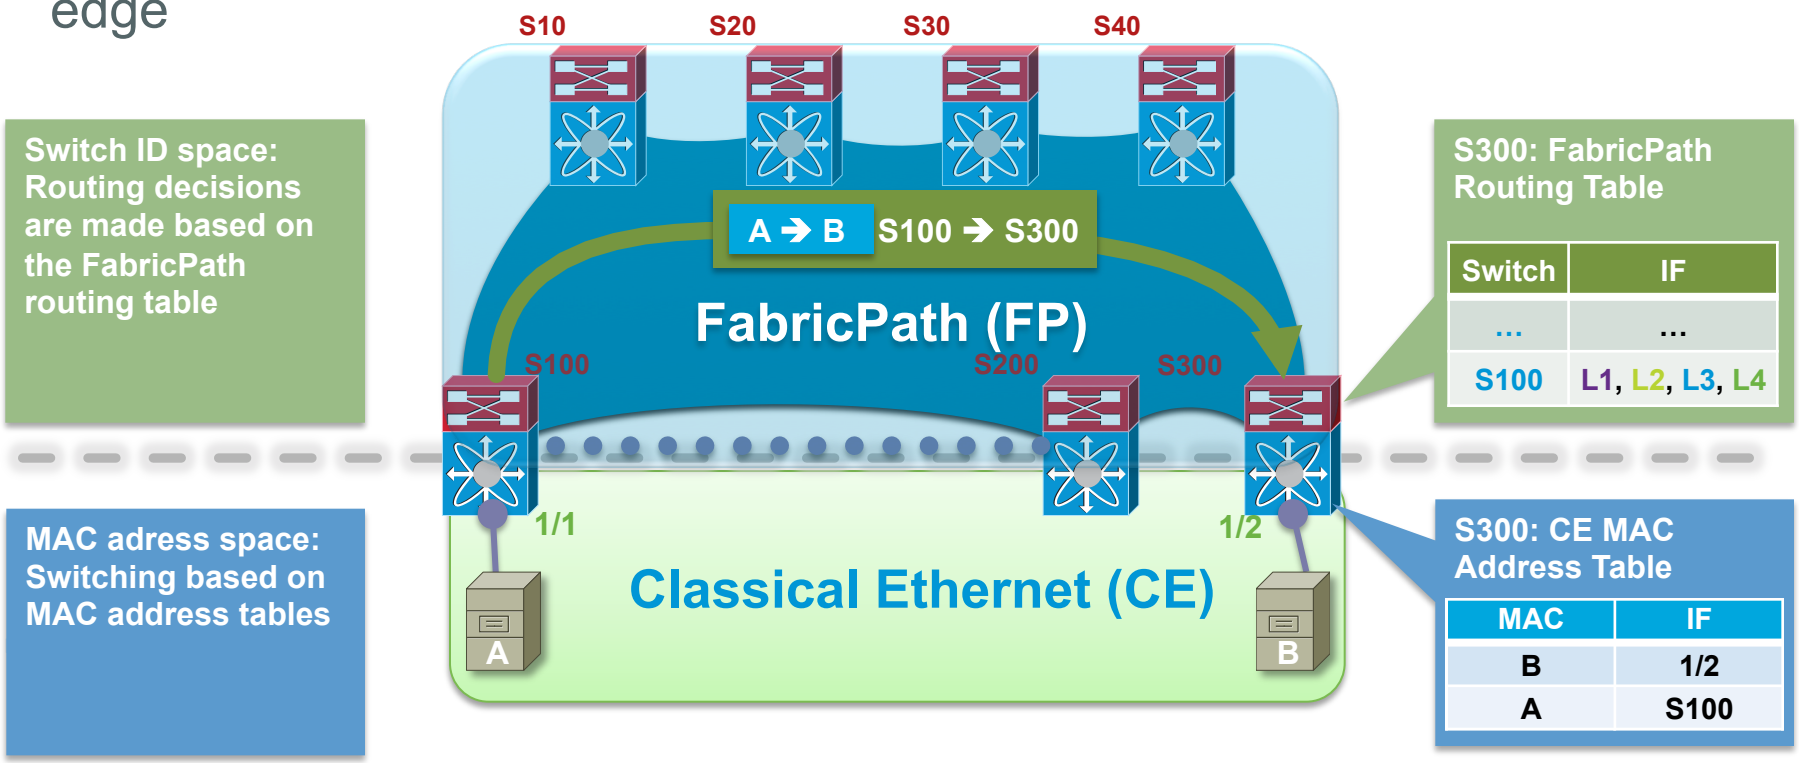

- Traffic is encapsulated across the Fabric

FabricPath Encapsulation

16-Byte MAC-in-MAC Header

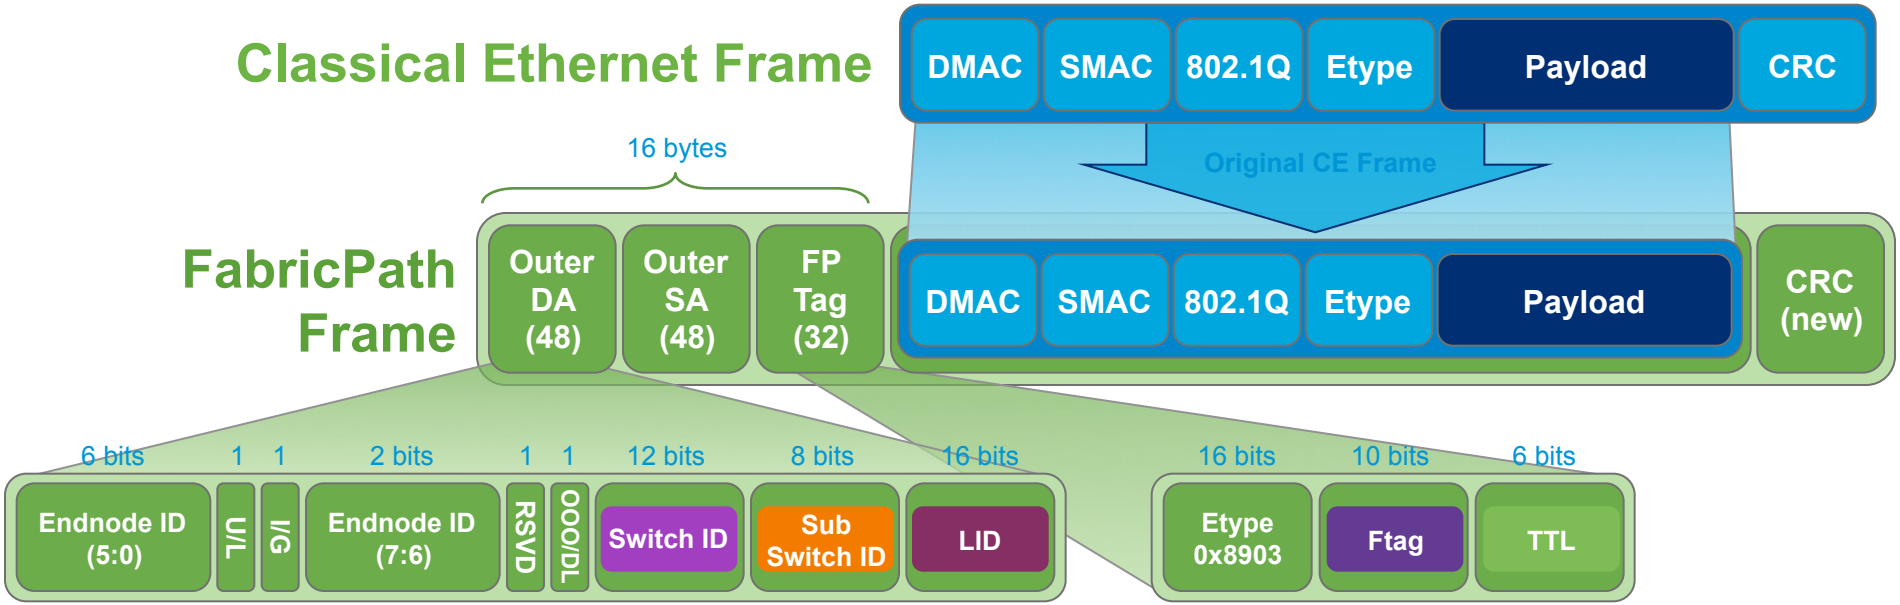

- **Switch ID** – Unique number identifying each FabricPath switch
- **Sub-Switch ID** – Identifies devices/hosts connected via VPC+
- **LID** – Local ID, identifies the destination or source interface
- **Ftag** (Forwarding tag) – Unique number identifying topology and/or distribution tree
- **TTL** – Decremented at each switch hop to prevent frames looping infinitely

Unknown Unicast

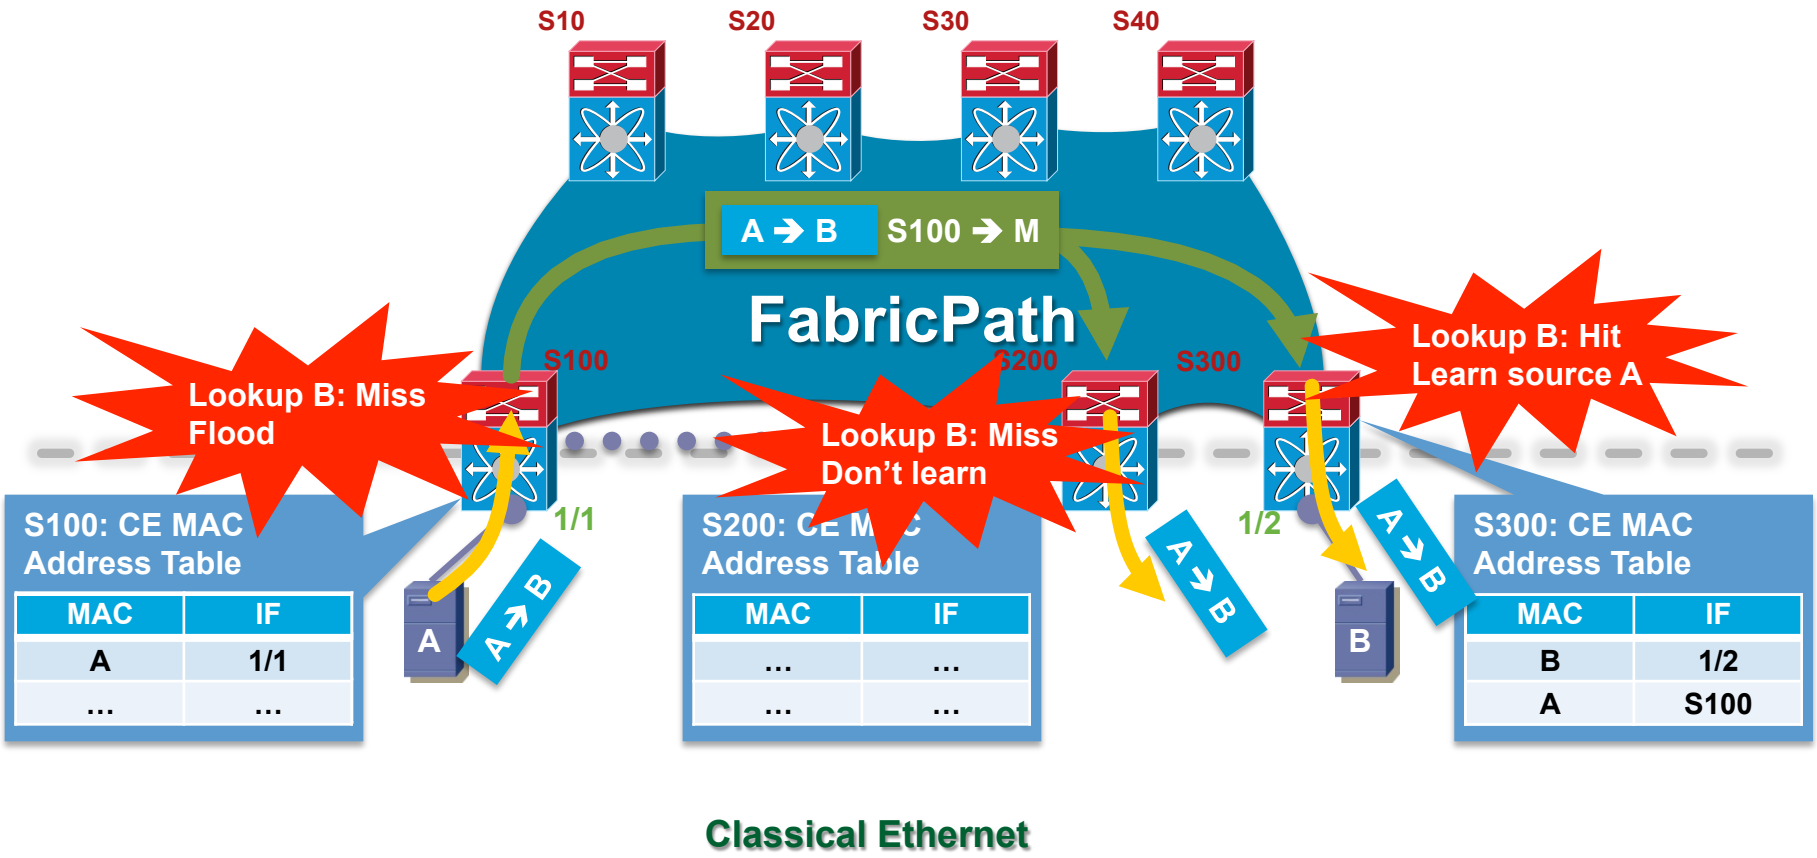

Classical Ethernet

Known Unicast, Conversational Learning

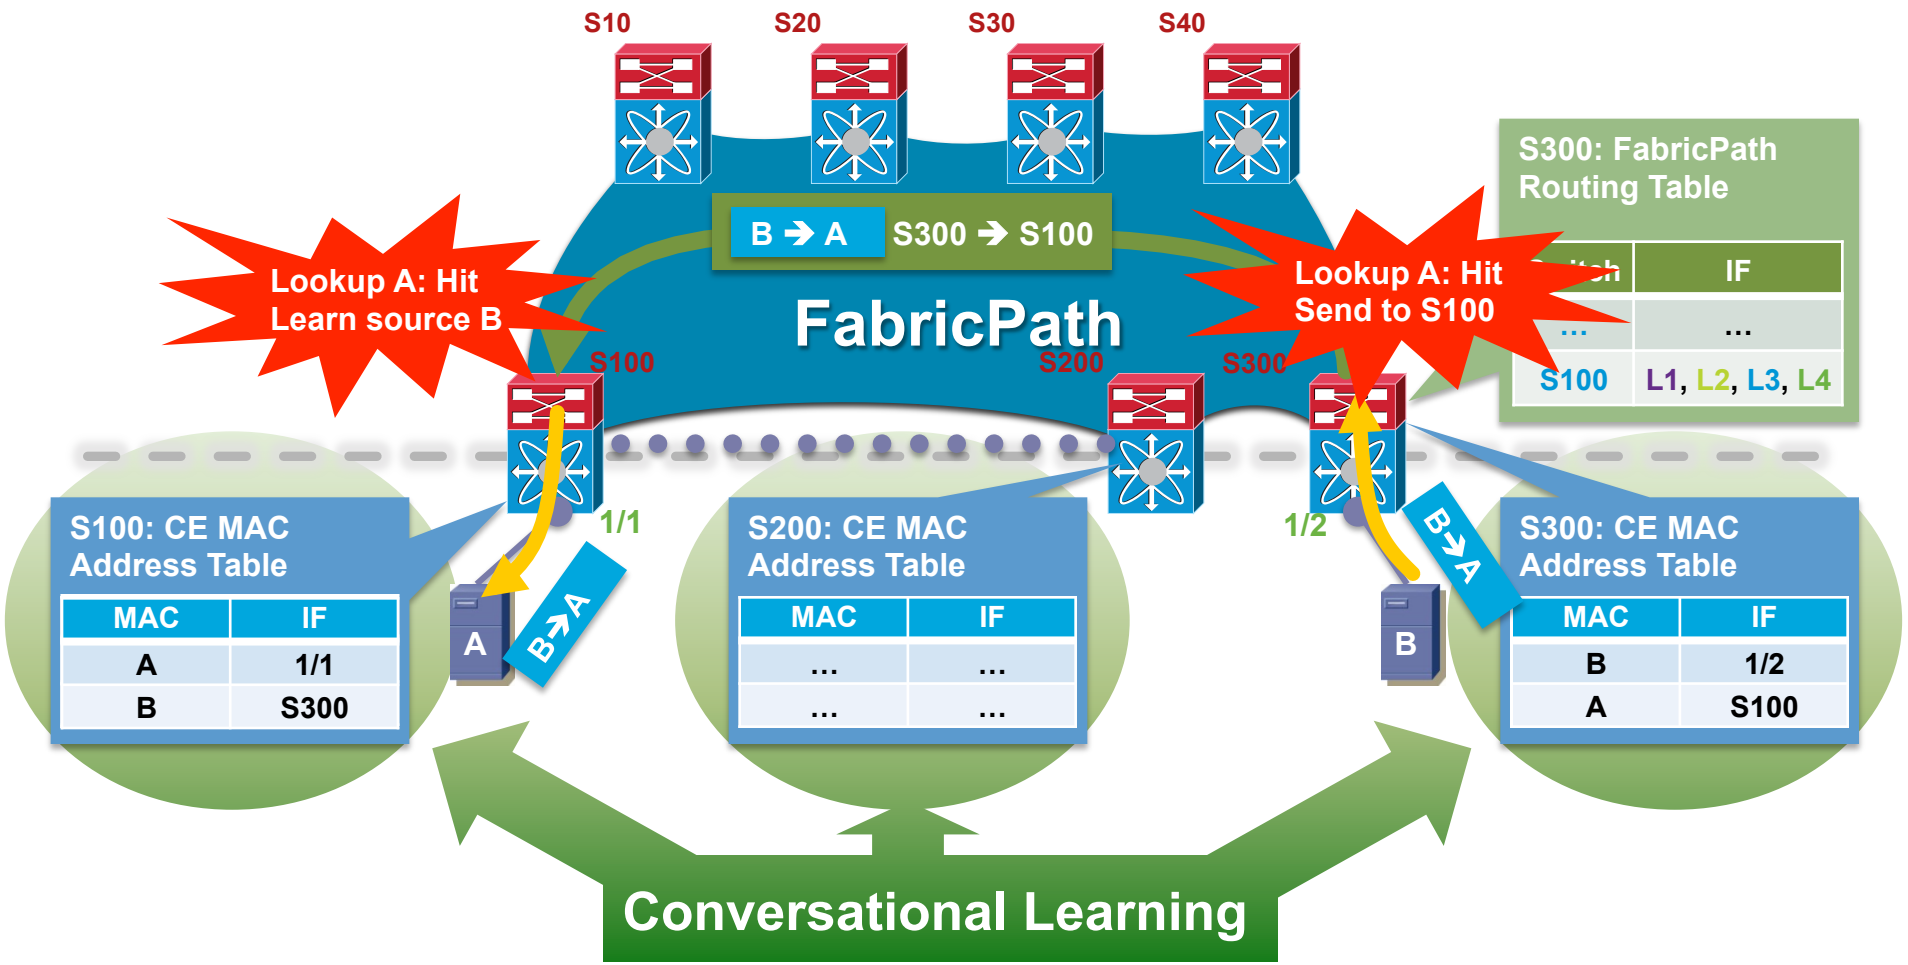

FabricPath VLANs

- The Nexus 7000 features M and F I/O Modules
- **FP Core and CE Edge ports must be on an F module**

- New FabricPath/CE locally significant VLAN mode:

```
S100(config)# vlan 10
S100(config-vlan)# mode ?
  ce           Classical Ethernet VLAN mode
  fabricpath   Fabricpath VLAN mode

S100(config-vlan)# mode fabricpath
S100(config-vlan)#
```

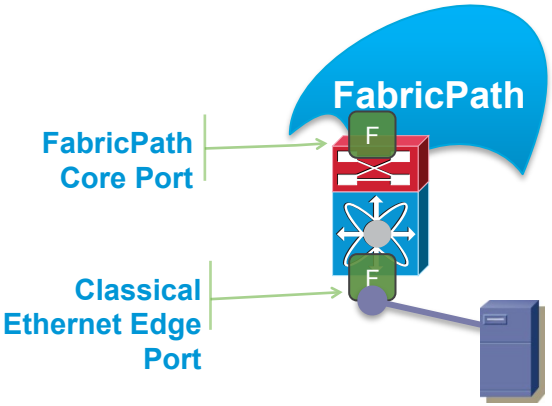

- FabricPath VLANs can only be enabled on F modules
- FabricPath VLANs are also relevant to the Nexus 5500

Equal Cost Multipathing

Traffic Forwarding Based on a Routing Table

- IS-IS assigns addresses to all FabricPath switches automatically
- Compute shortest, pair-wise paths
- Support equal-cost paths between any FabricPath switch pairs

FabricPath Routing Table

Switch	IF
S10	L1
S20	L2
S30	L3
S40	L4
S200	L1, L2, L3, L4
…	…
S400	L1, L2, L3, L4

S100: CE MAC Address Table

MAC	IF
A	1/1
B	400

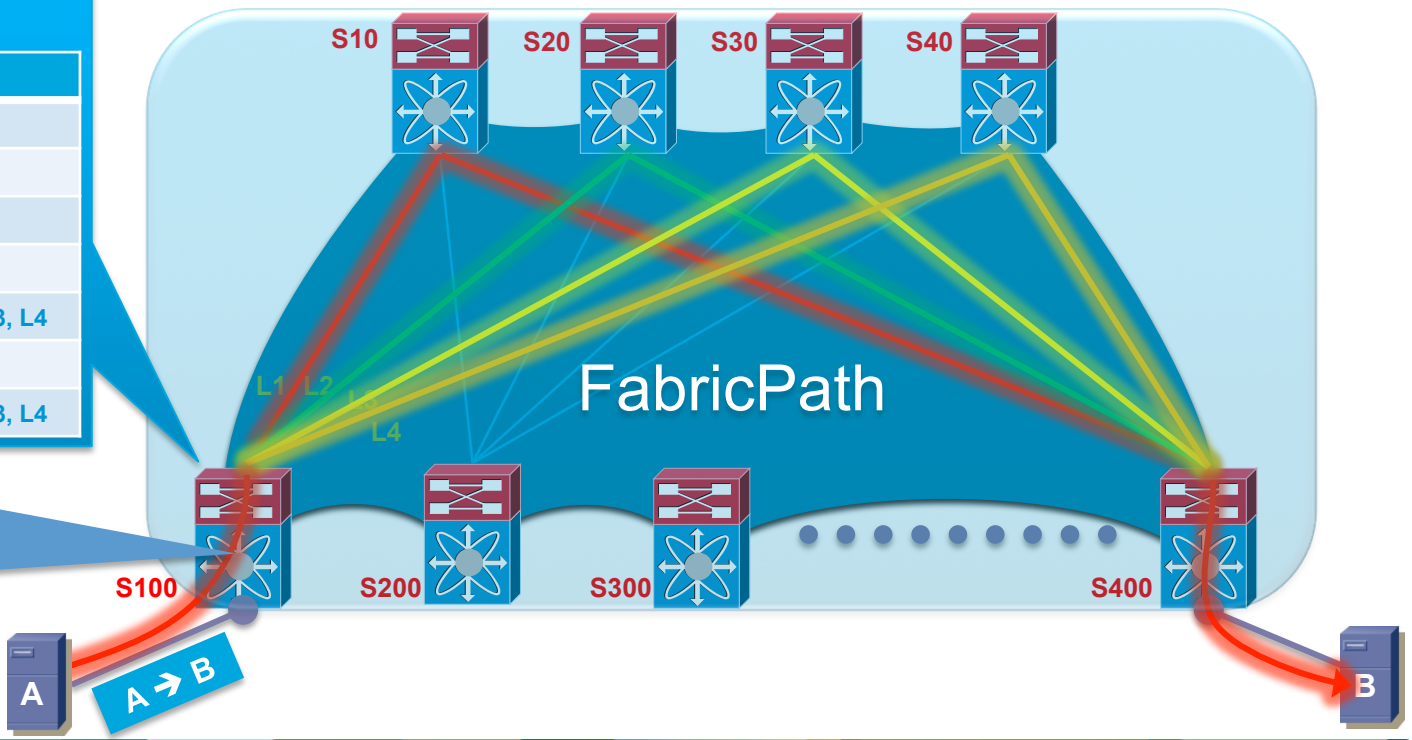

Multicast Traffic

Load Balancing on a Per-Tree Basis

- IS-IS computes several trees automatically
- Location of the root switches can be configured
- Multicast traffic is pinned to a tree at the edge

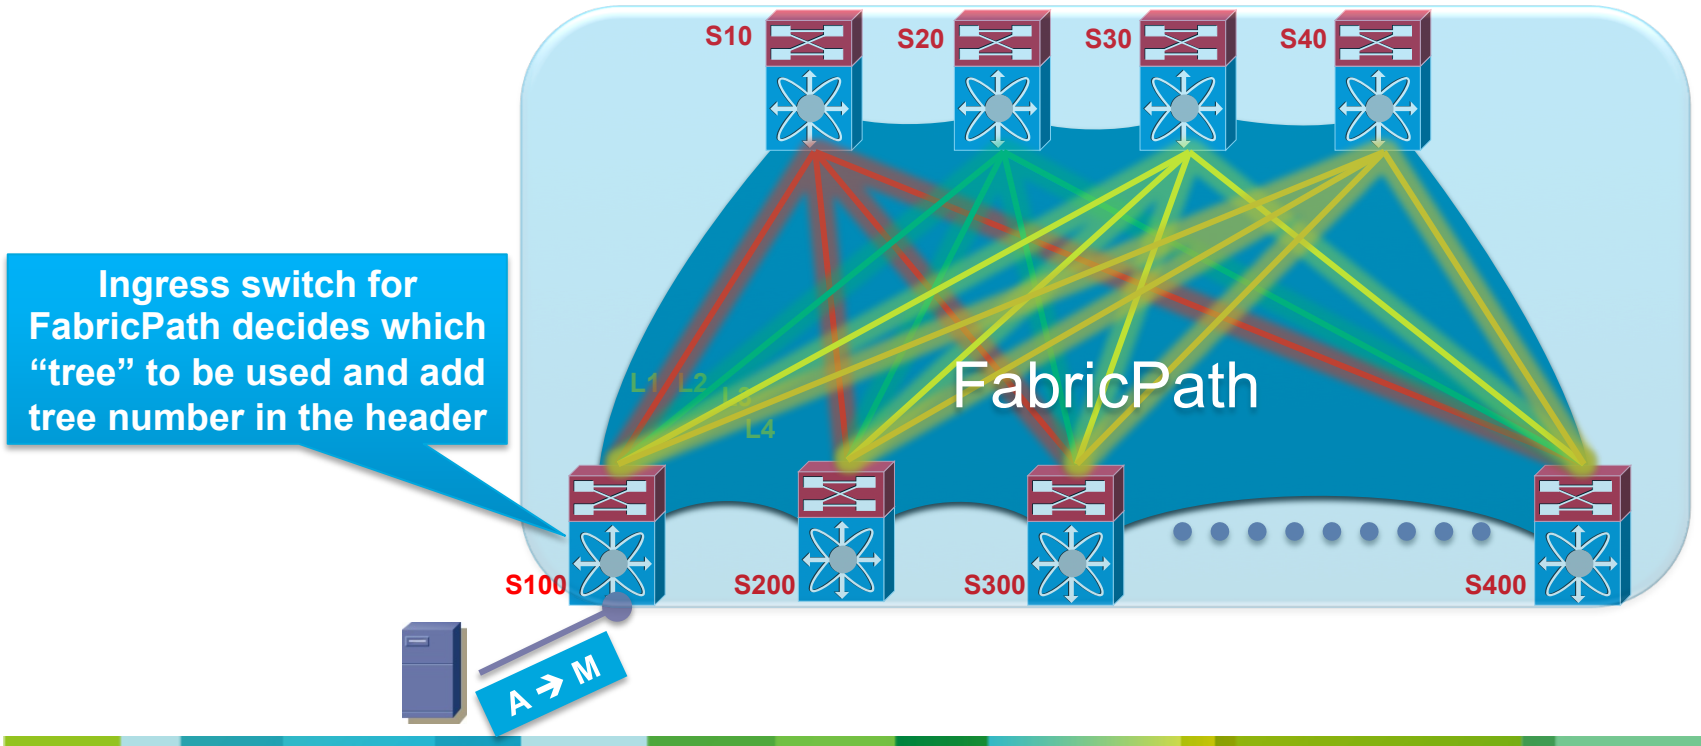

VLAN Pruning

Automatically Handled by IS-IS

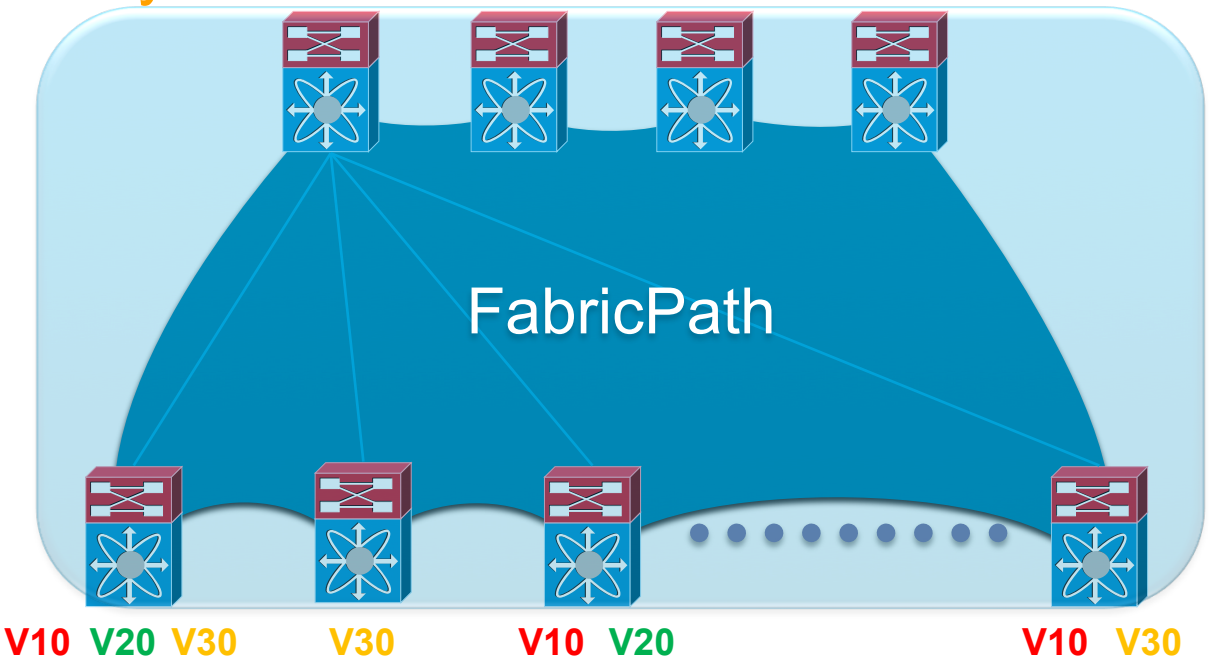

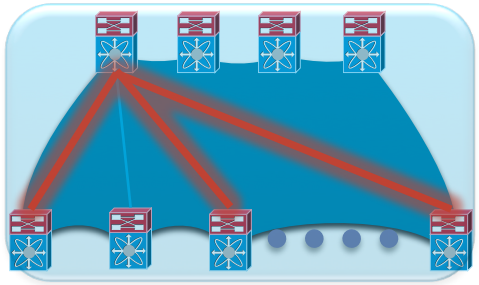

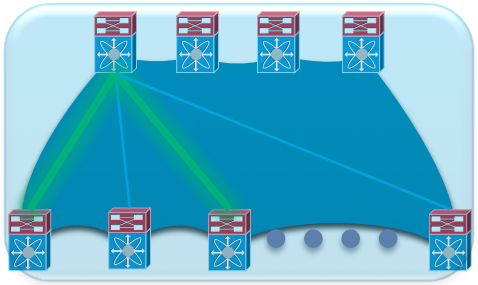

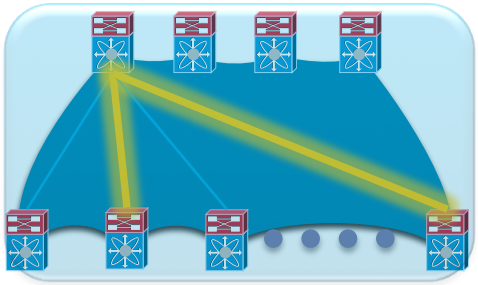

VPC+

Virtual Port Channel in FabricPath Environment

- Allows non FabricPath capable devices to connect redundantly to the fabric using port channels
- Provides active/active HSRP
- Configuration virtually identical to standard VPC

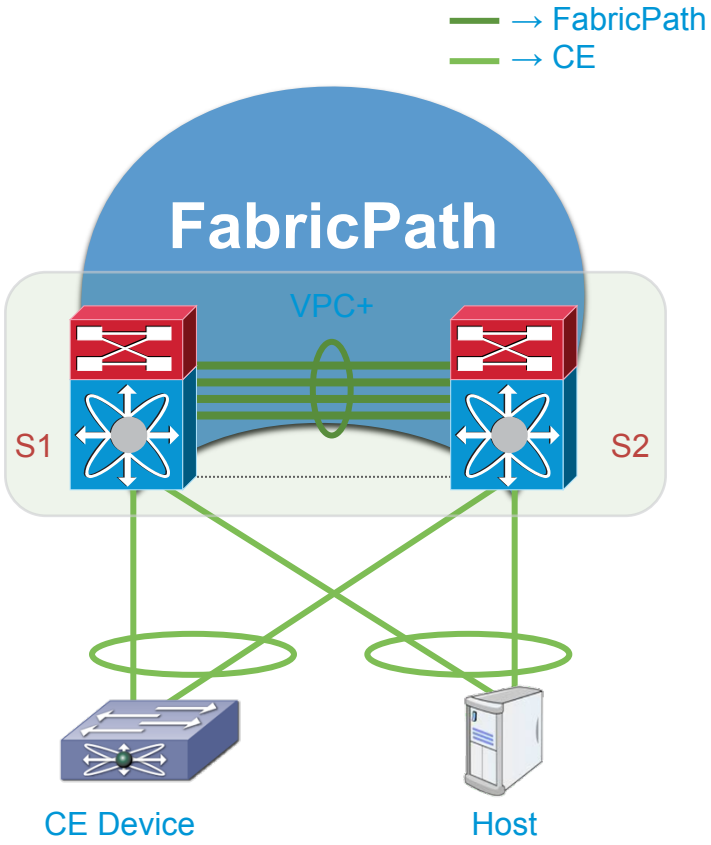

VPC+ Technical Challenges

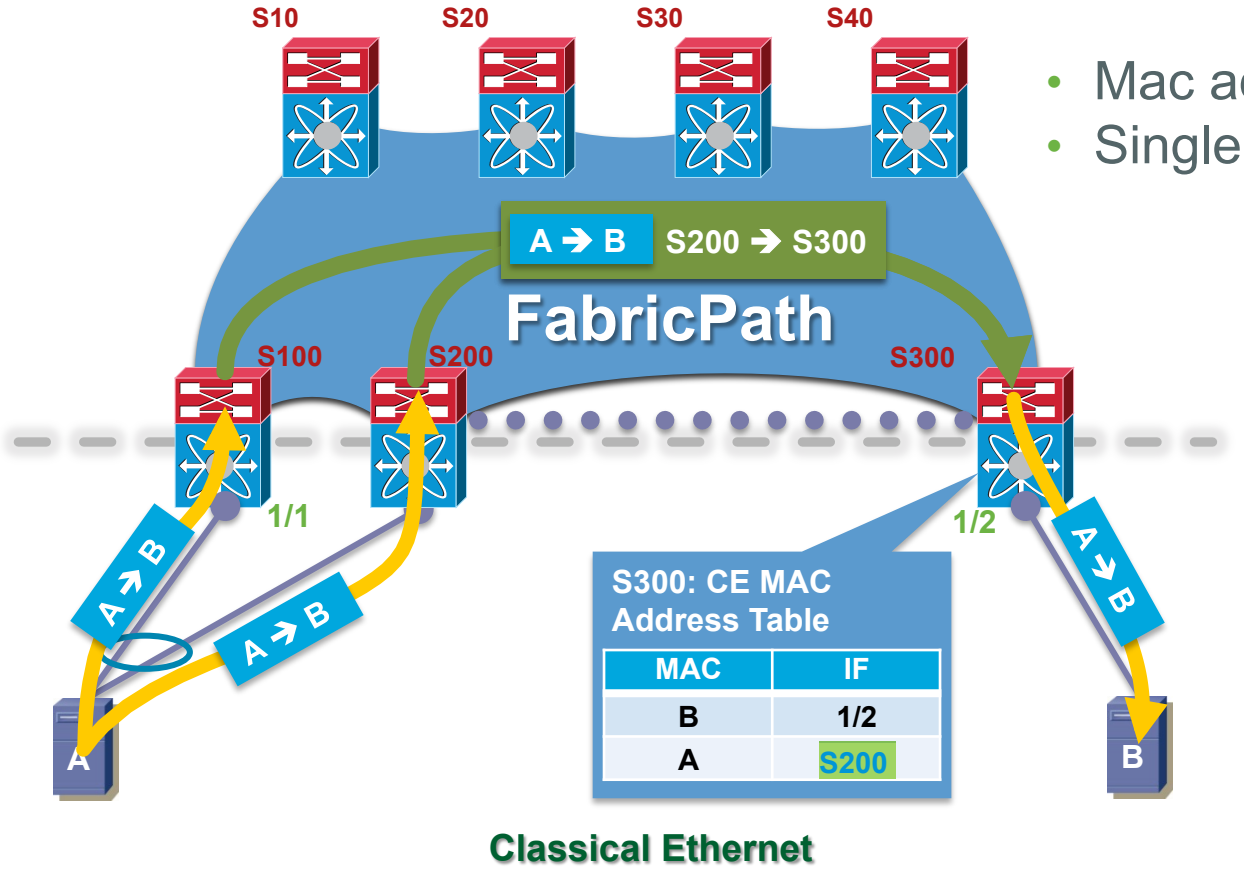

- Mac address flapping on S300
- Single path to A

VPC+ Virtual Switch

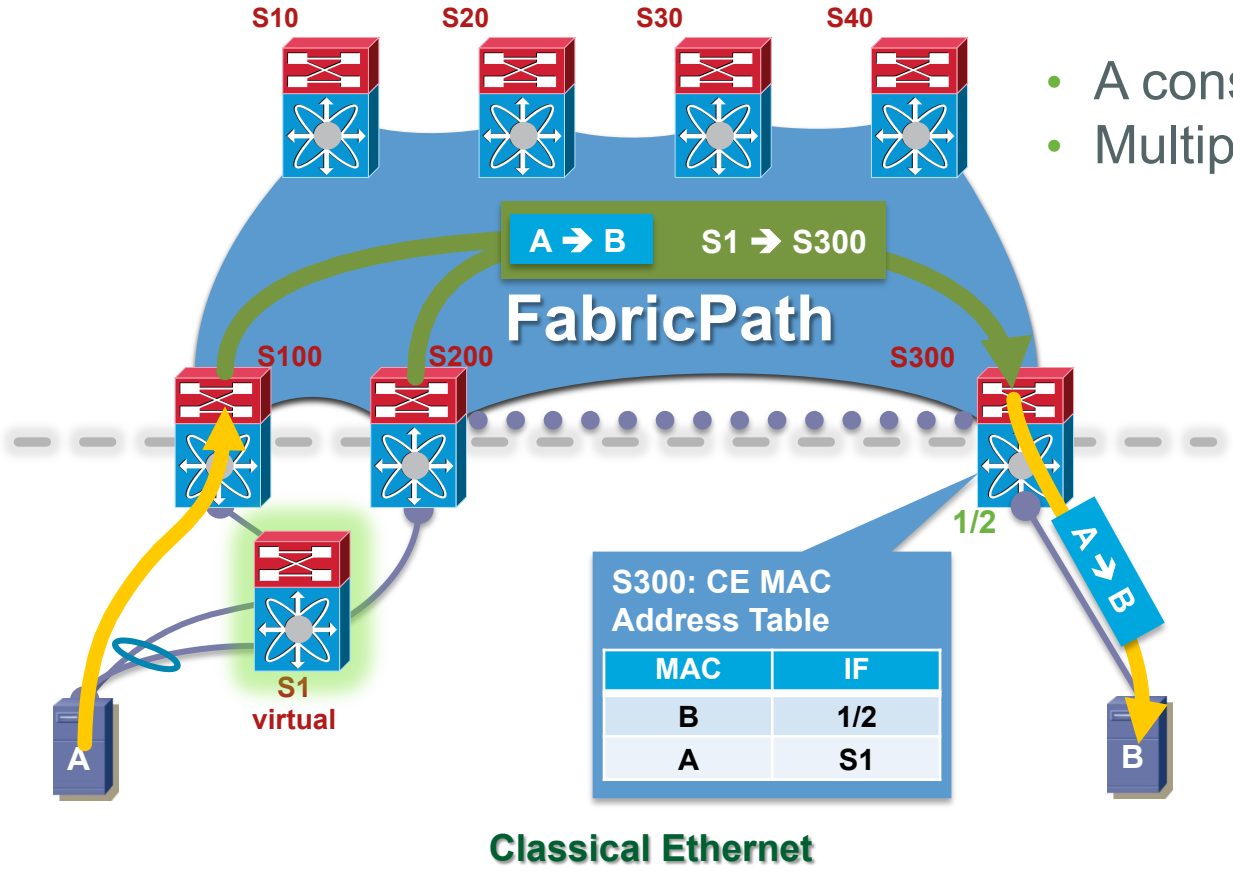

- A consistently associated to S1
- Multipathing to A

Sub-Switch ID

Identifies a VPC Off a Virtual Switch

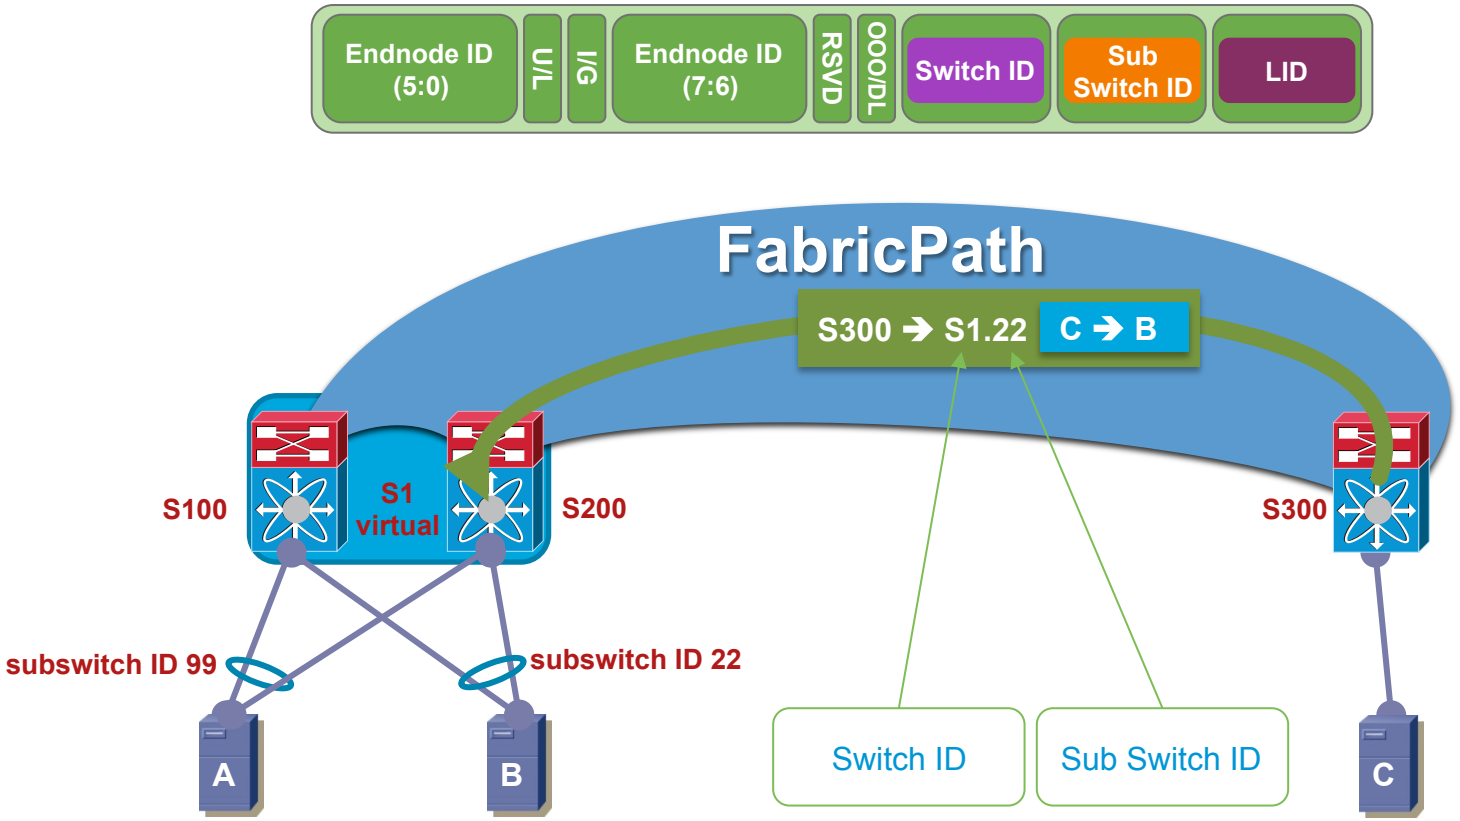

Loop Mitigation with FabricPath

Time To Live (TTL) and Reverse Path Forwarding (RPF) Check

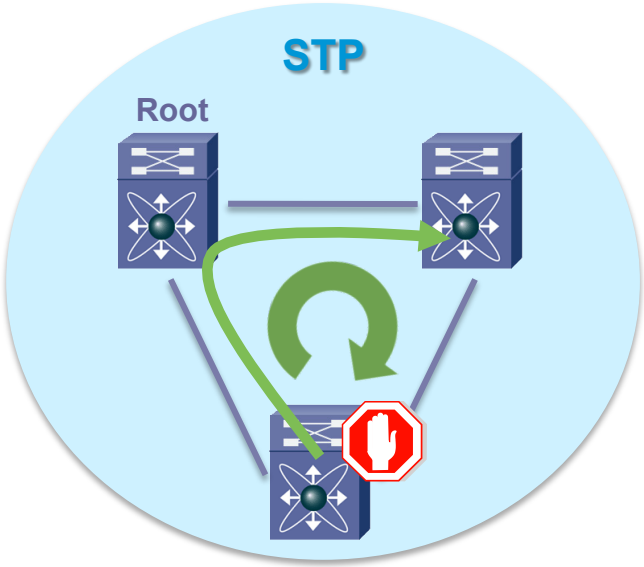

- The control protocol is the only mechanism preventing loops
- If STP fails →infinite loop
  - No backup mechanism in the data plane
  - Flooding impacts the whole network

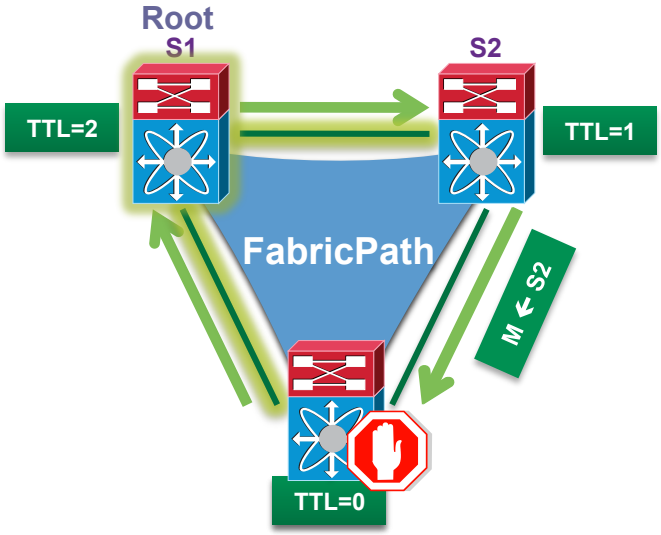

- TTL in FabricPath header
- RPF Check for multi-destination traffic
- The data plane is protecting against loops

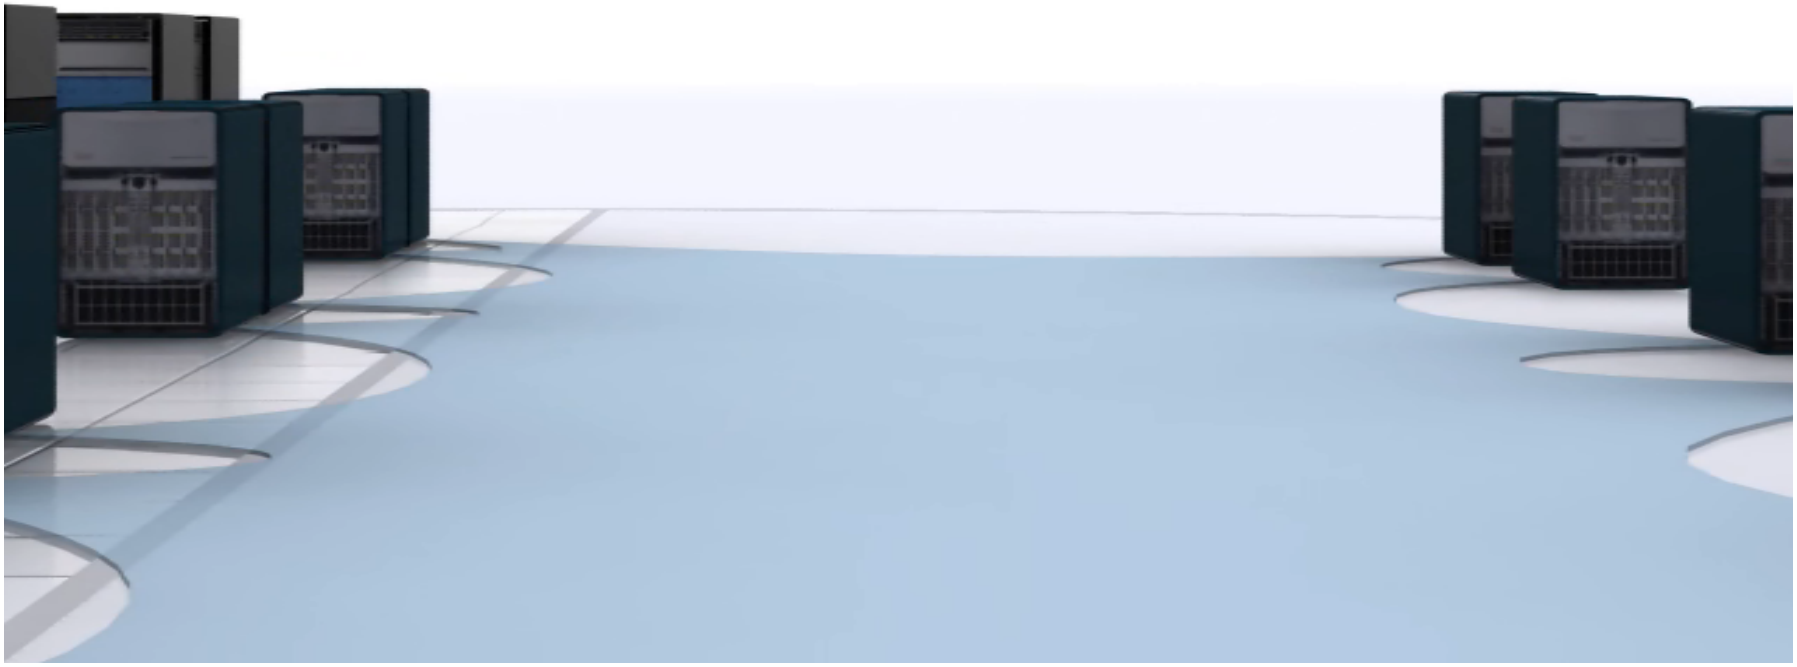

Unicast Forwarding Details

FabricPath Unicast System Architecture

- Complete separation of control plane and data plane
- Fully modular software implementation of control plane and infrastructure components
- Fully distributed data plane forwarding with hardware-based MAC learning / forwarding and hardware SID / ECMP lookups

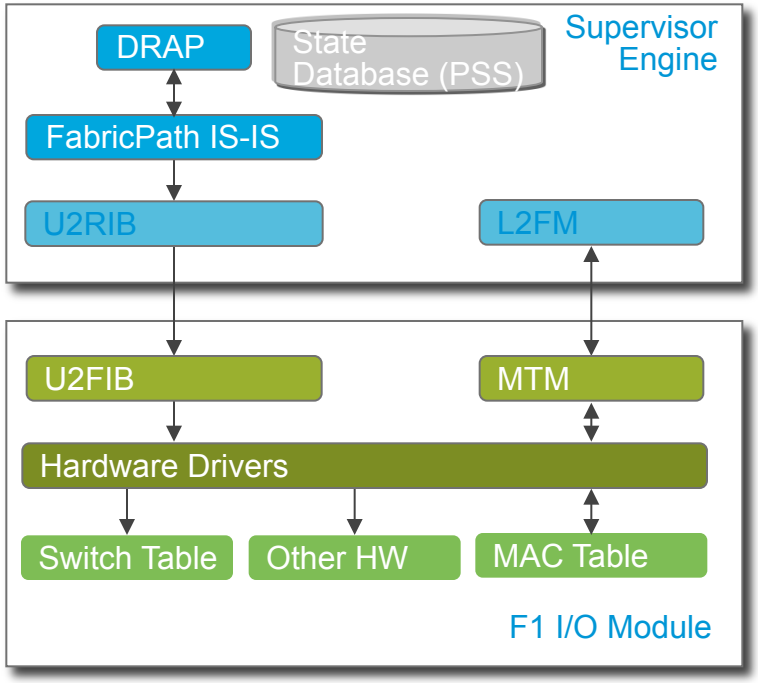

FabricPath Routing Table

- Describes shortest (best) paths to each Switch ID based on link metrics
- Equal-cost paths supported between FabricPath switches

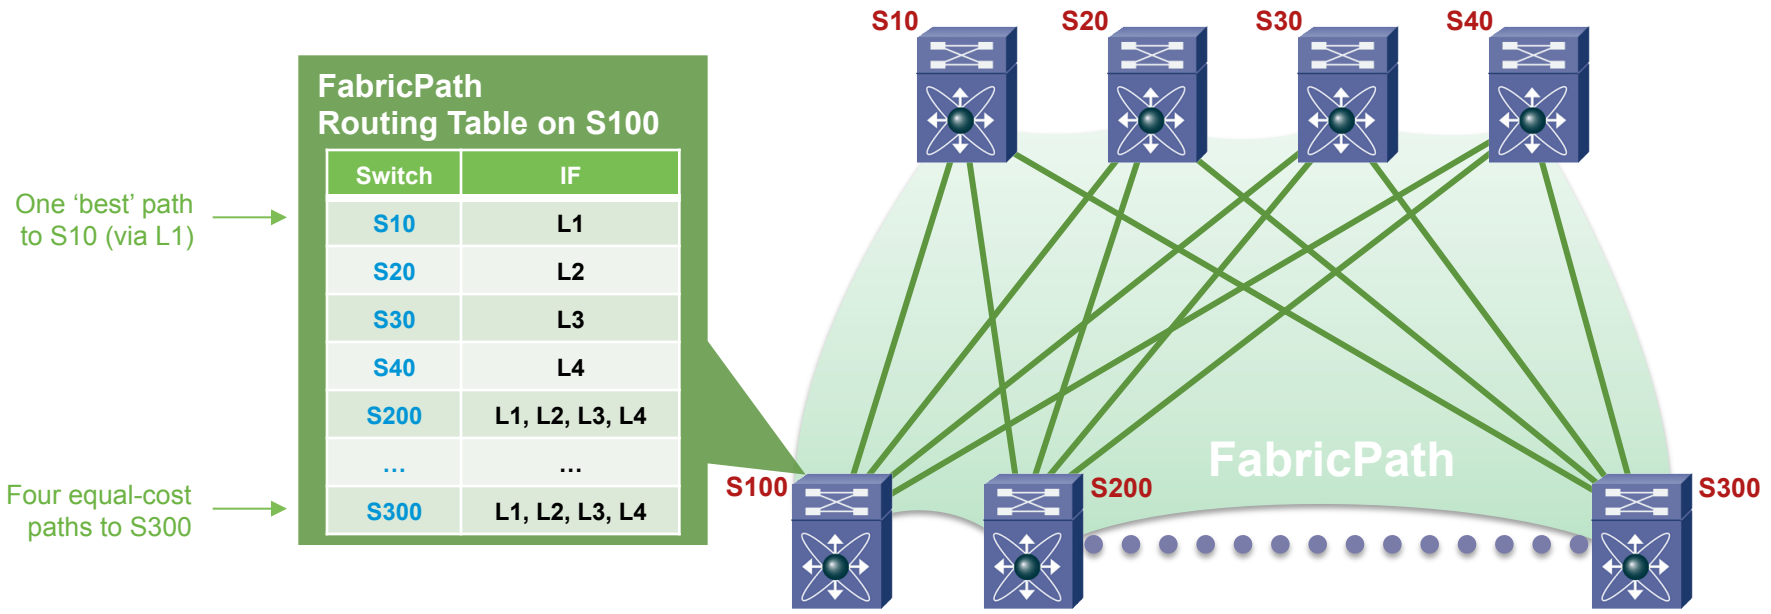

Display IS-IS View of Routing Topology

show fabricpath isis route

Destination Switch ID

Routing metric

Next-hop interface(s)

```
S100# sh fabricpath isis route
Fabricpath IS-IS domain: default MT-0
Topology 0, Tree 0, Swid routing table
10, L1
 via port-channel10, metric 20
20, L1
 via port-channel20, metric 20
30, L1
 via port-channel30, metric 20
40, L1
 via port-channel40, metric 20
200, L1
 via port-channel30, metric 40
 via port-channel40, metric 40
 via port-channel20, metric 40
 via port-channel10, metric 40
300, L1
 via port-channel30, metric 40
 via port-channel40, metric 40
 via port-channel20, metric 40
 via port-channel10, metric 40

S100#
```

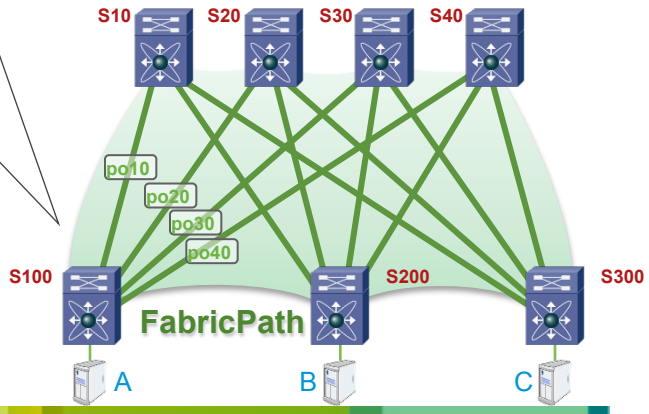

Display U2RIB View of Routing Topology

show fabricpath route

```
S100# sh fabricpath route
FabricPath Unicast Route Table
'a/b/c' denotes ftag/switch-id/subswitch-id
'[x/y]' denotes [admin distance/metric]
ftag 0 is local ftag
subswitch-id 0 is default subswitch-id


FabricPath Unicast Route Table for Topology-Default

0/100/0, number of next-hops: 0
        via ----- , [60/0], 0 day/s 04:43:51, local
1/10/0, number of next-hops: 1
        via Po10, [115/20], 0 day/s 02:24:02, isis_fabricpath-default
1/20/0, number of next-hops: 1
        via Po20, [115/20], 0 day/s 04:43:25, isis_fabricpath-default
1/30/0, number of next-hops: 1
        via Po30, [115/20], 0 day/s 04:43:25, isis_fabricpath-default
1/40/0, number of next-hops: 1
        via Po40, [115/20], 0 day/s 04:43:25, isis_fabricpath-default
1/200/0, number of next-hops: 4
        via Po10, [115/40], 0 day/s 02:24:02, isis_fabricpath-default
        via Po20, [115/40], 0 day/s 04:43:06, isis_fabricpath-default
        via Po30, [115/40], 0 day/s 04:43:06, isis_fabricpath-default
        via Po40, [115/40], 0 day/s 04:43:06, isis_fabricpath-default
1/300/0, number of next-hops: 4
        via Po10, [115/40], 0 day/s 02:24:02, isis_fabricpath-default
        via Po20, [115/40], 0 day/s 04:43:25, isis_fabricpath-default
        via Po30, [115/40], 0 day/s 04:43:25, isis_fabricpath-default
        via Po40, [115/40], 0 day/s 04:43:25, isis_fabricpath-default
S100#
```

Topology (ftag), Switch ID, Sub-Switch ID

Administrative distance, routing metric

Route age

Client protocol

Next-hop interface(s)

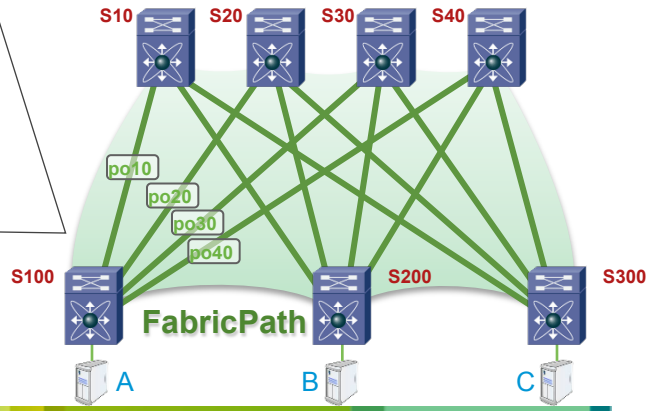

Putting It All Together – Host A to Host B

(1) Broadcast ARP Request

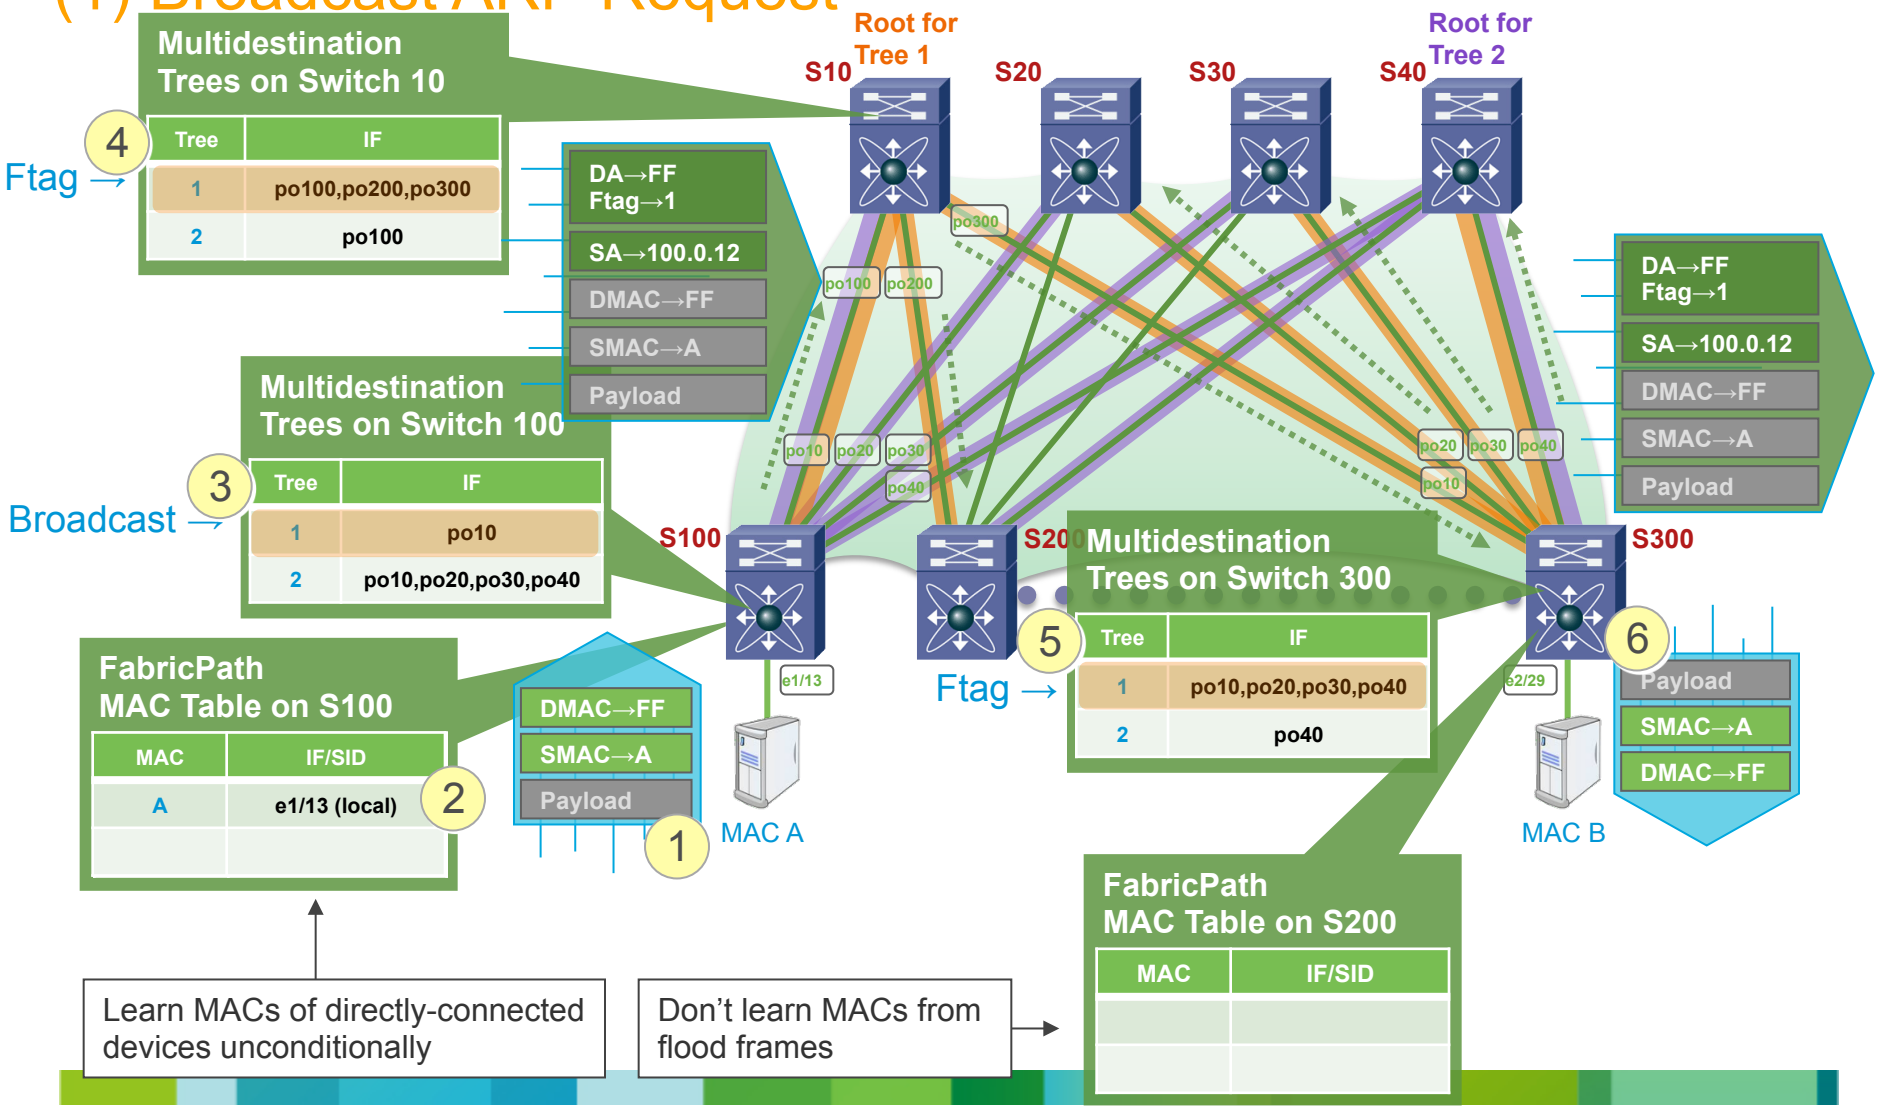

MAC Address Tables After Broadcast ARP

- S100:

```
S100# sh mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0000.0000.000a    dynamic   0          F     F  Eth1/13

S100#
```

MAC A learned as local entry on e1/13

- S10 (and S20, S30, S40, S200):

```
S10# sh mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------

S10#
```

MAC A not learned on other switches

- S300:

```
S300# sh mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
```

Putting It All Together – Host A to Host B

(2) Unicast ARP Reply

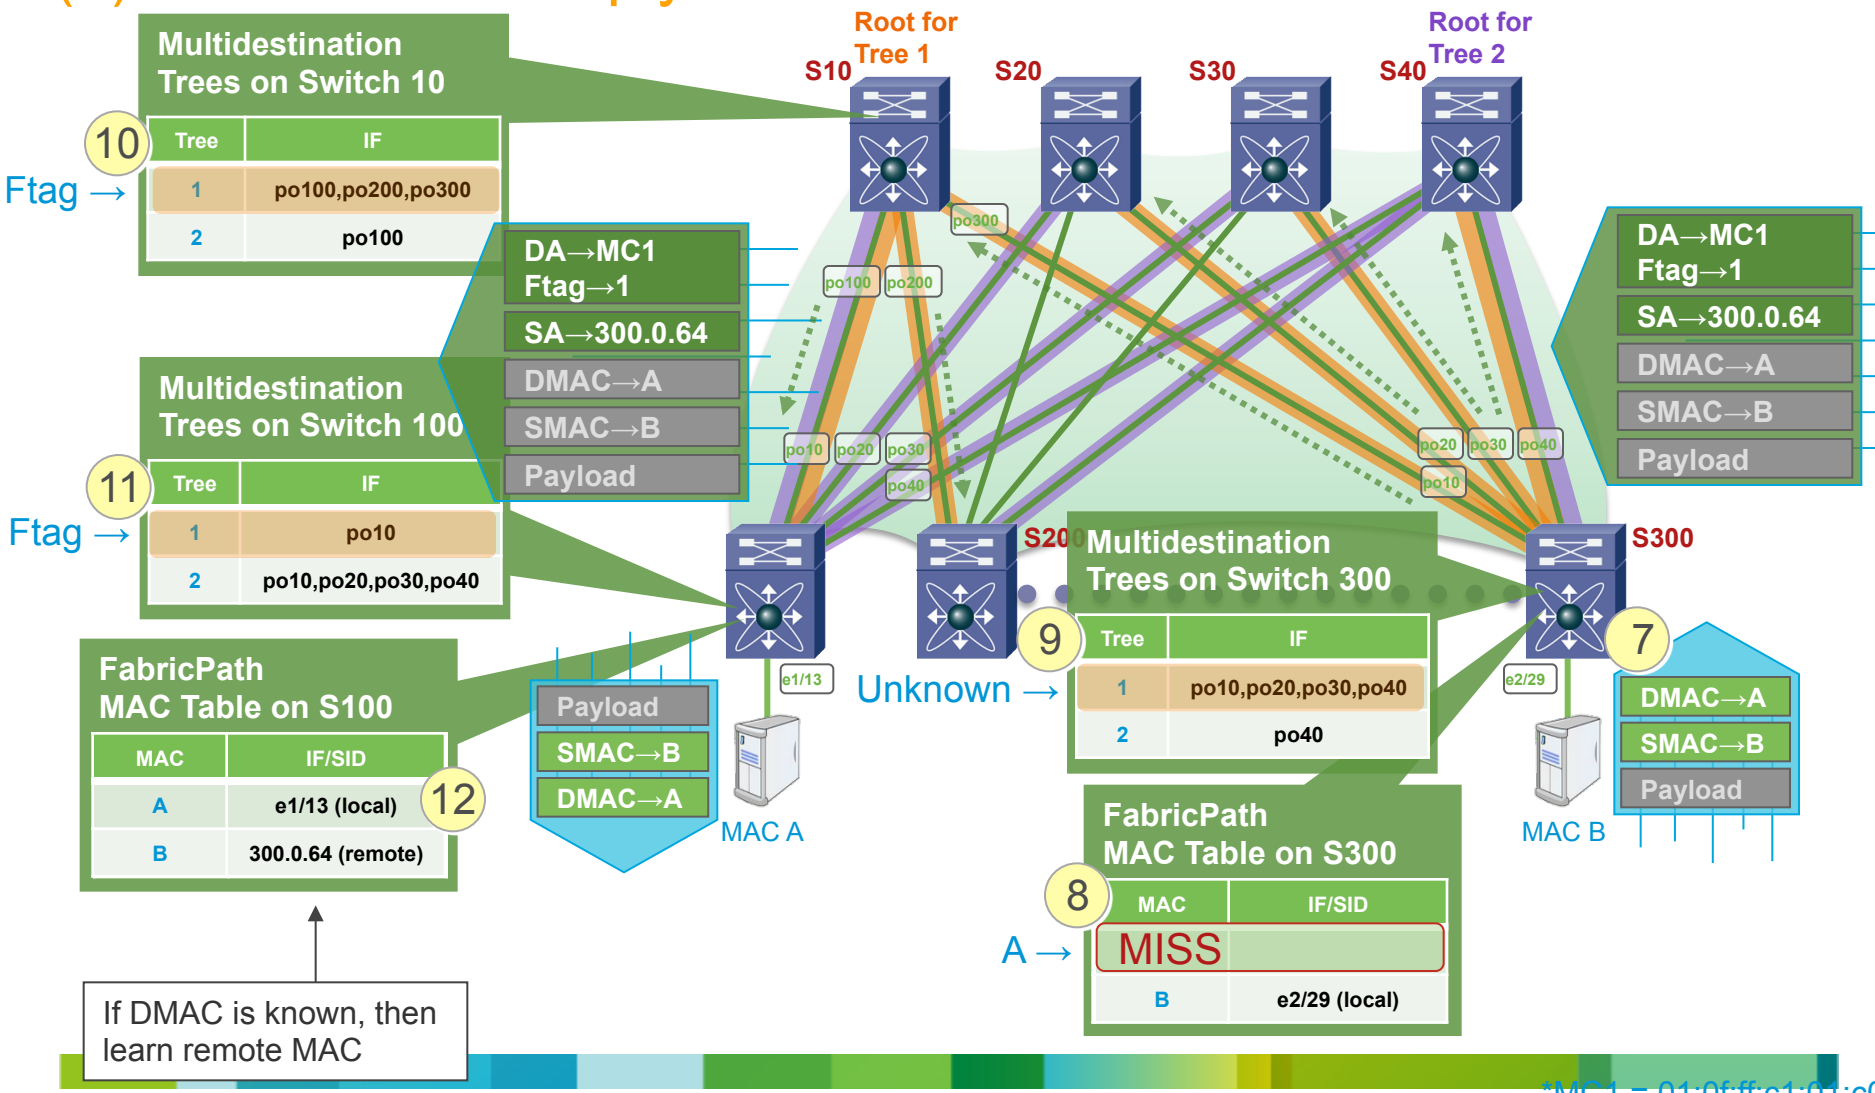

*MC1 = 01:0f:ff:c1:01:c0

MAC Address Tables After Unicast ARP Reply

- S100:

```
S100# sh mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0000.0000.000a    dynamic   90         F    F  Eth1/13
  10       0000.0000.000b    dynamic   60         F    F  300.0.64

S100#
```

S100 learns MAC B as remote entry reached through S300

- S300:

```
S300# sh mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0000.0000.000b    dynamic   0          F    F  Eth2/29

S300#
```

MAC B learned as local entry on e2/29

Putting It All Together – Host A to Host B

(3) Unicast Data

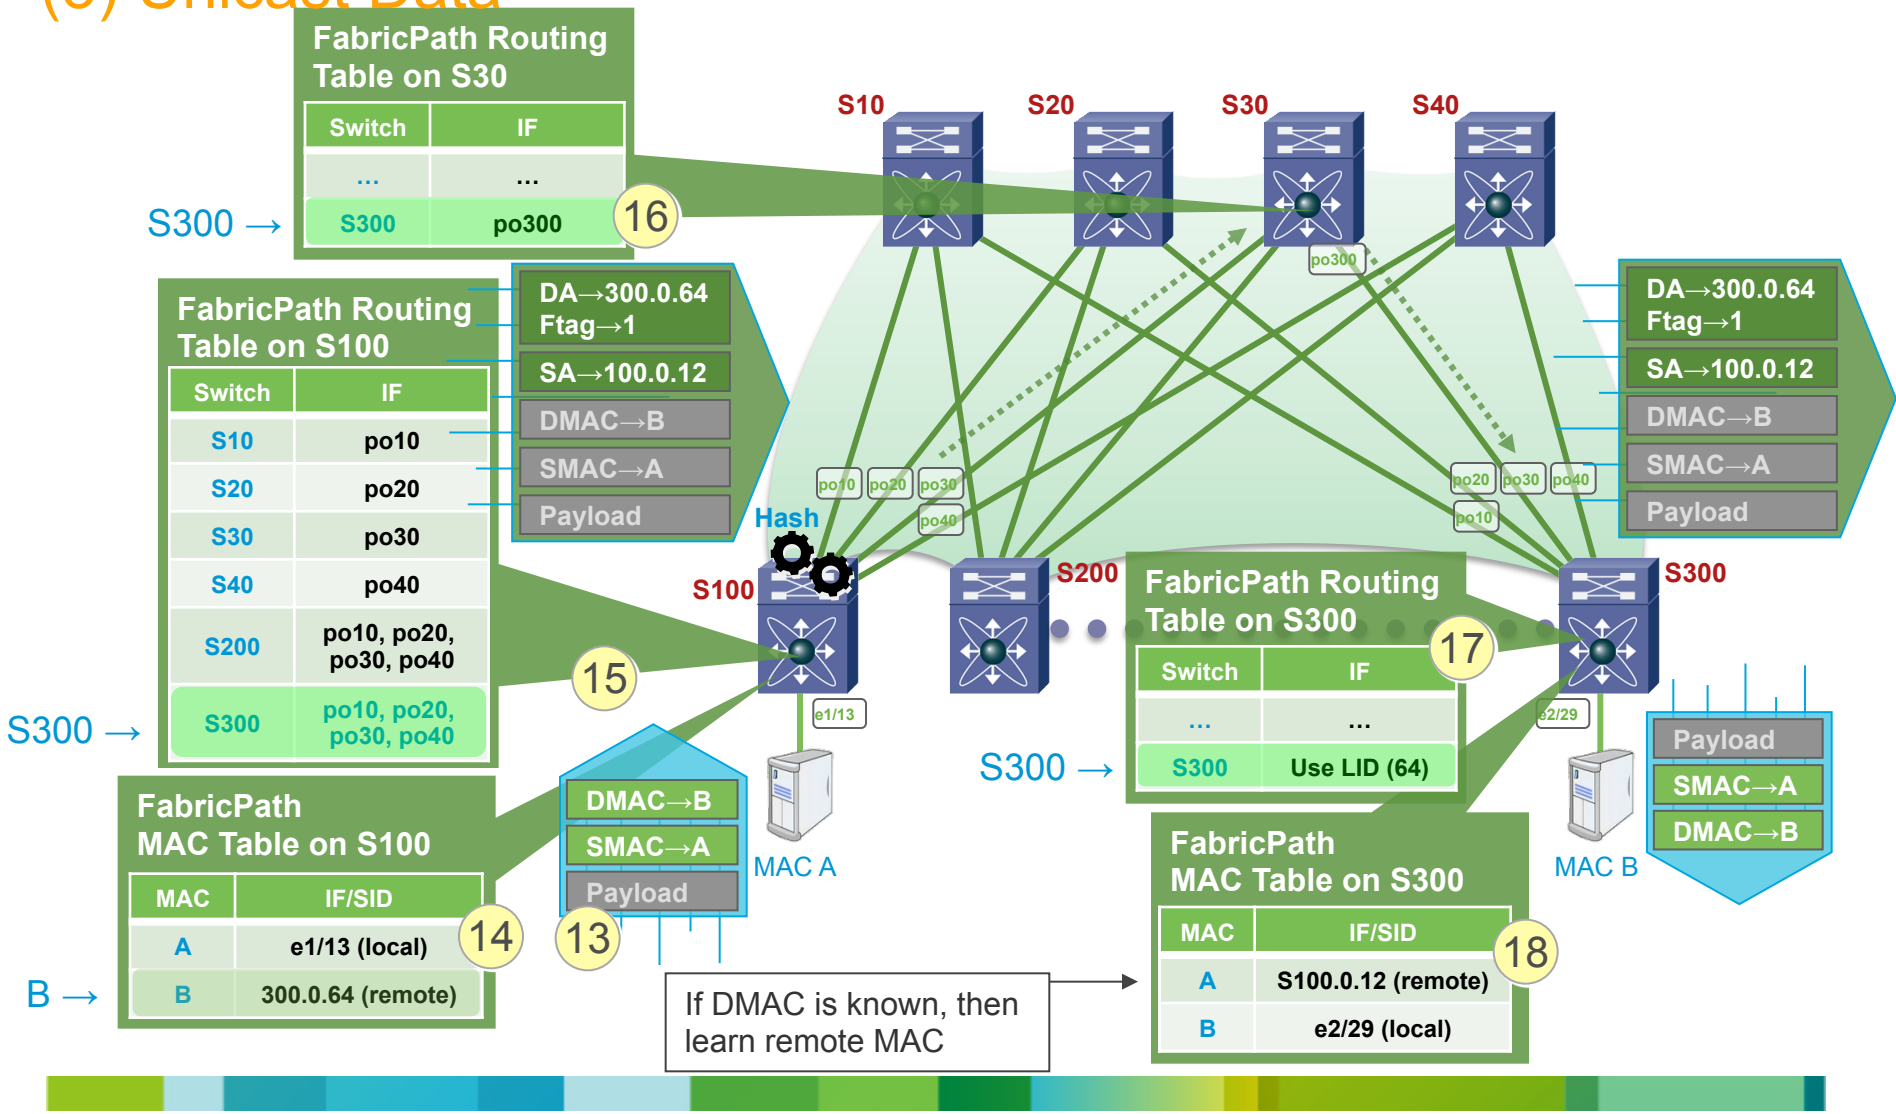

MAC Address Tables After Unicast Data

- S100:

```
S100# sh mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
* 10       0000.0000.000a    dynamic   90         F    F  Eth1/13
  10       0000.0000.000b    dynamic   60         F    F  300.0.64

S100#
```

- S300:

```
S300# sh mac address-table dynamic
Legend:
        * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
        age - seconds since last seen,+ - primary entry using vPC Peer-Link
   VLAN     MAC Address      Type      age     Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
  10       0000.0000.000a    dynamic   30         F    F  100.0.12
* 10       0000.0000.000b    dynamic   90         F    F  Eth2/29

S300#
```

S100 learns MAC A as remote entry reached through S100

Multicast Forwarding Details

FabricPath IP Multicast

- **Control plane:**
  - Build several multidestination trees
  - Run IGMP snooping on FabricPath edge switches
  - Advertise receivers location with dedicated LSPs
- **Data plane (hardware):**
  - Selects which multidestination tree or each flow based on hash function
  - Forward traffic along selected tree

FabricPath IP Multicast System Architecture

- Complete separation of control plane and data plane
- Fully modular software implementation for control plane and infrastructure components
- Fully distributed data plane forwarding with hardware-based MAC learning

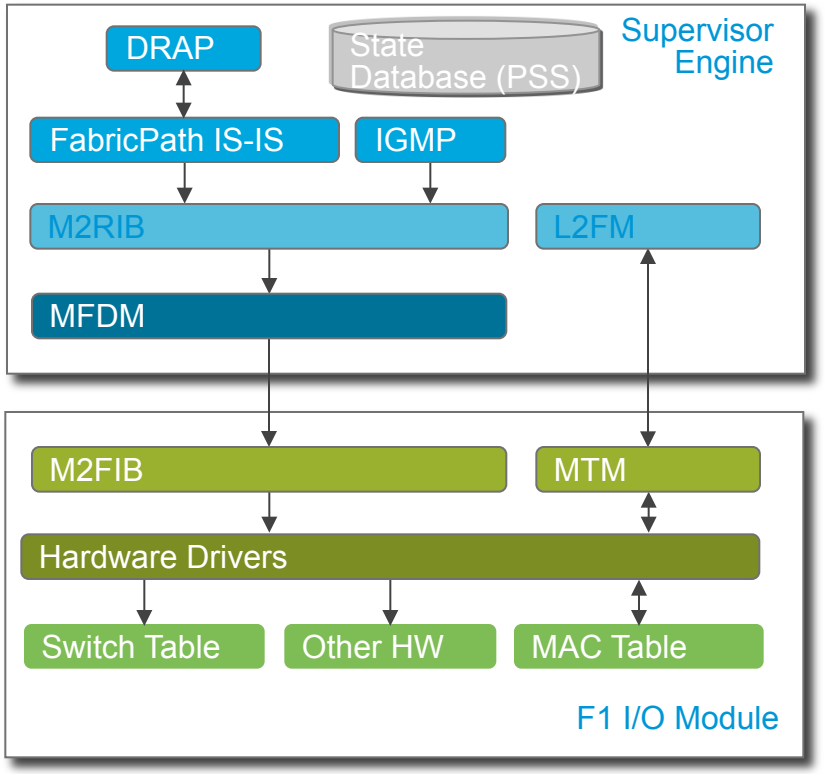

Multicast Trees Determination

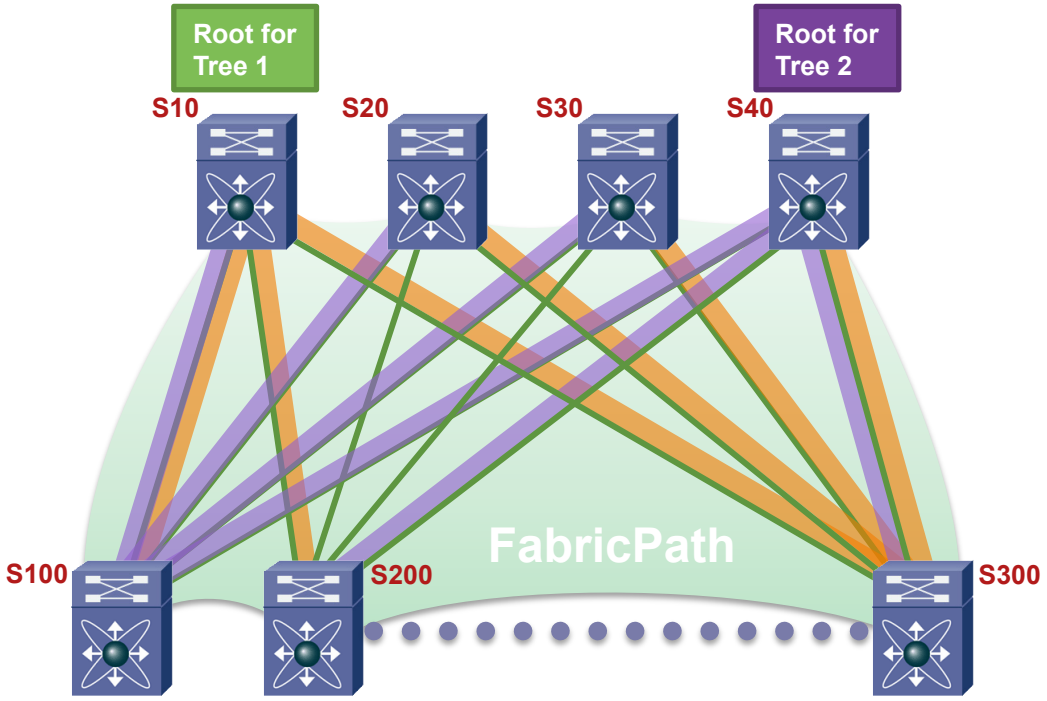

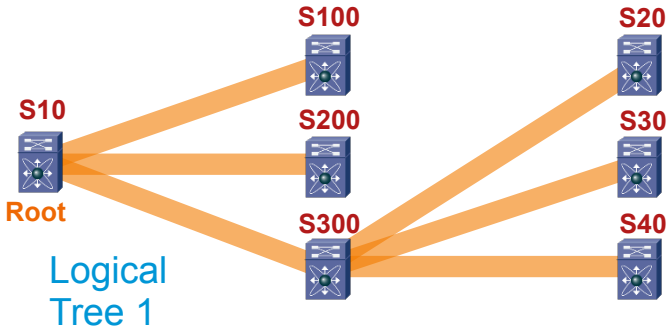

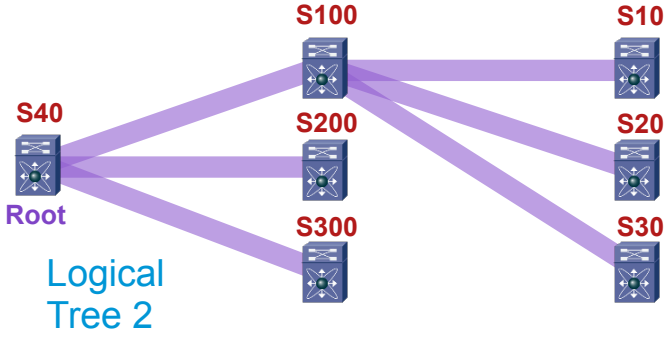

- Switch with highest priority value becomes root for primary tree
  Highest system ID, then highest Switch ID value, in case of a tie
- Primary root designates different secondary root(s) ensuring path variety.

Multicast Tree Pruning

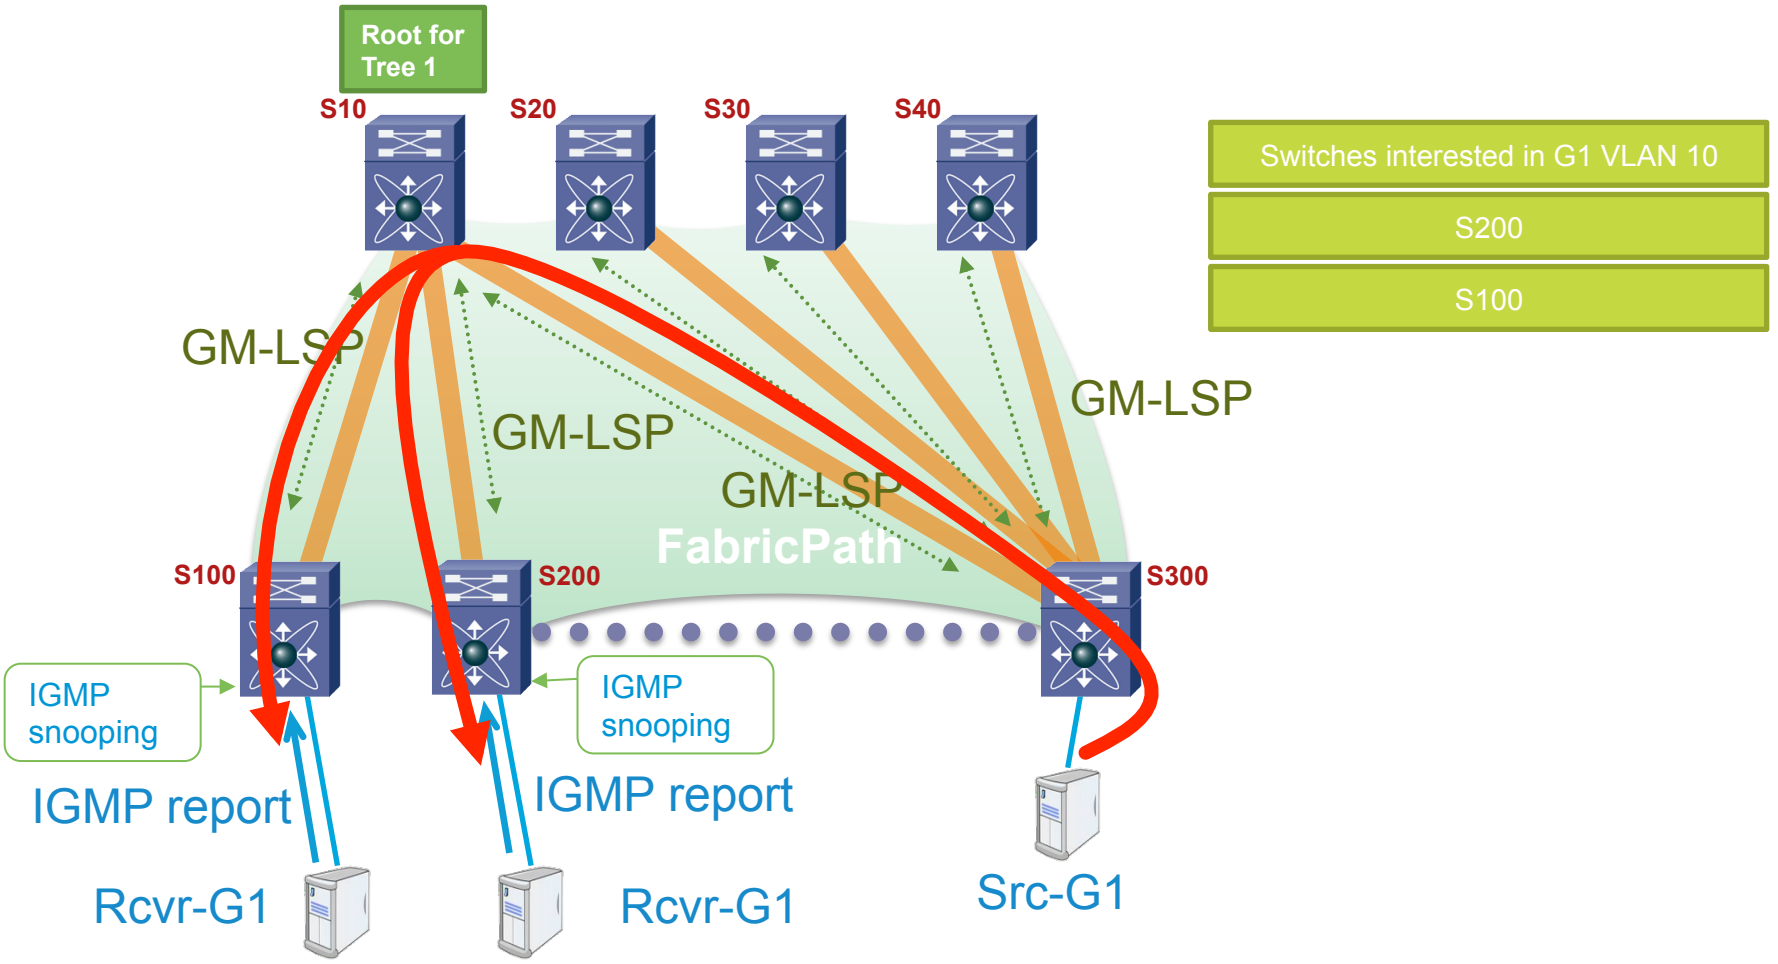

- IS-IS Group Membership LSPs contain multicast forwarding information

Multicast Load Balancing

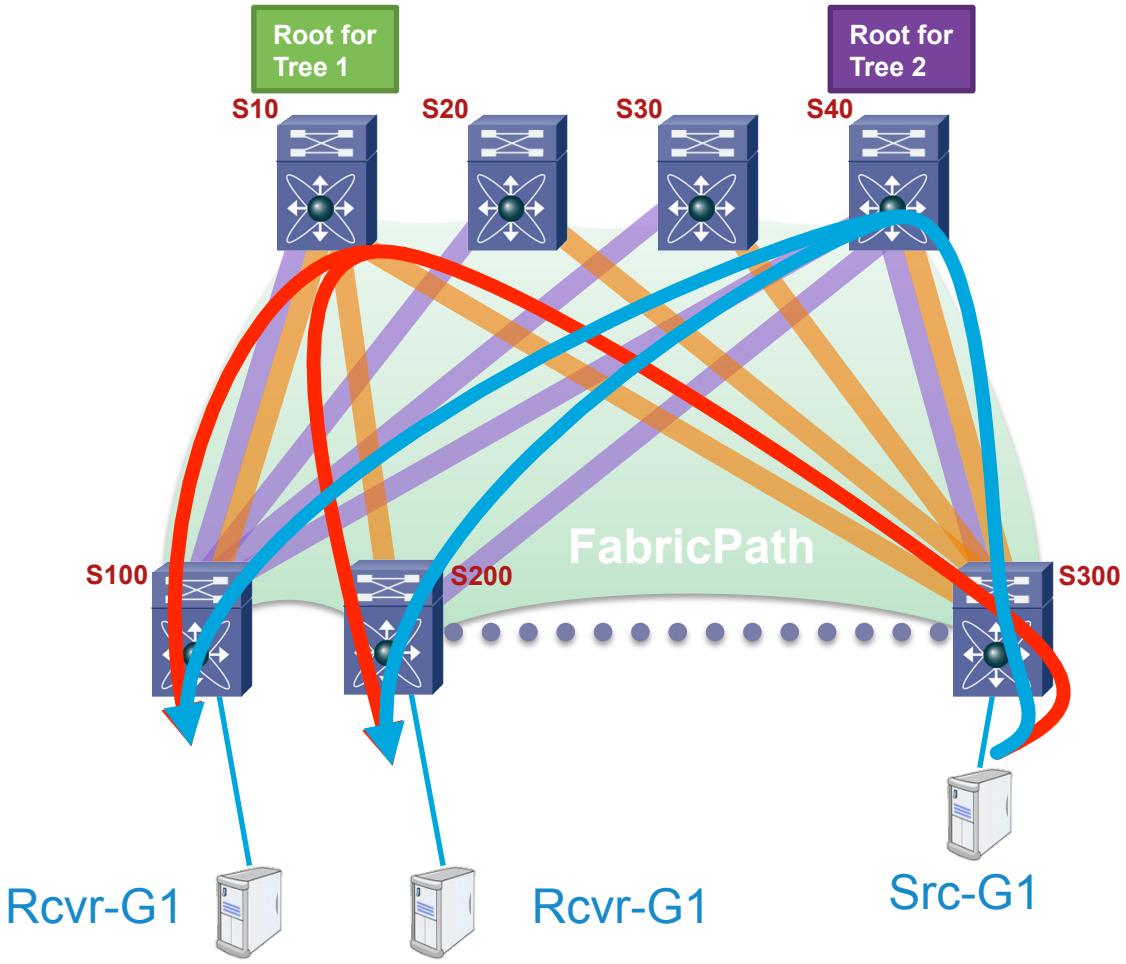

Multicast Data Plane Step by Step

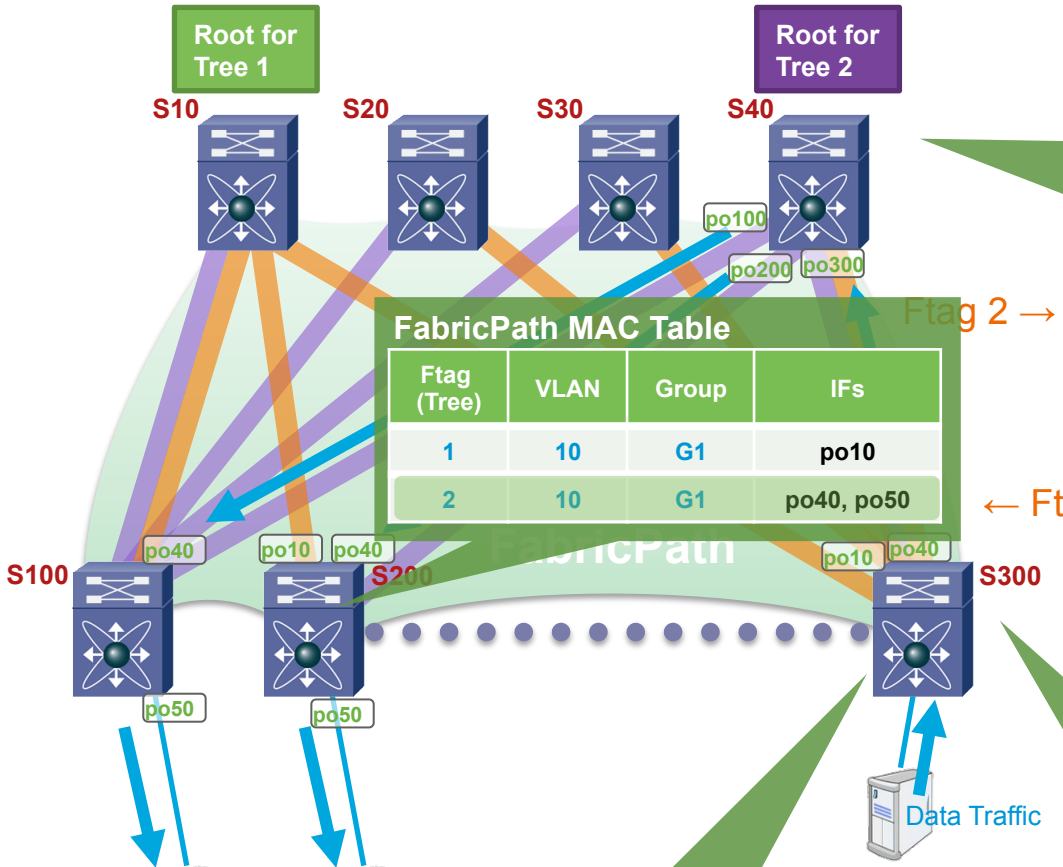

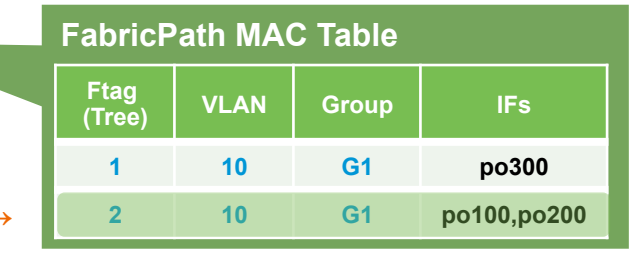

- **Ftag: Forwarding tag** – Unique 10-bit number identifying topology and/or distribution tree

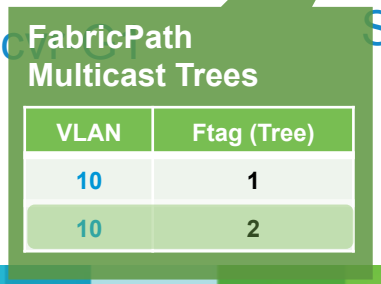

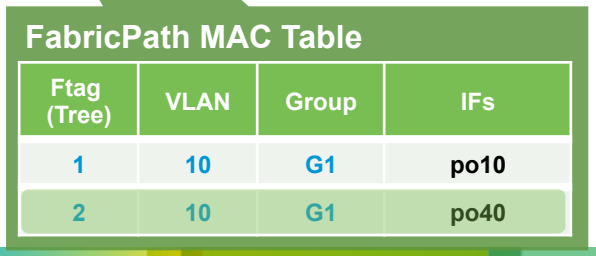

Packet data → → → Ftag 2 →

FabricPath Designs

Classical POD with FabricPath

FabricPath vs. vPC/STP

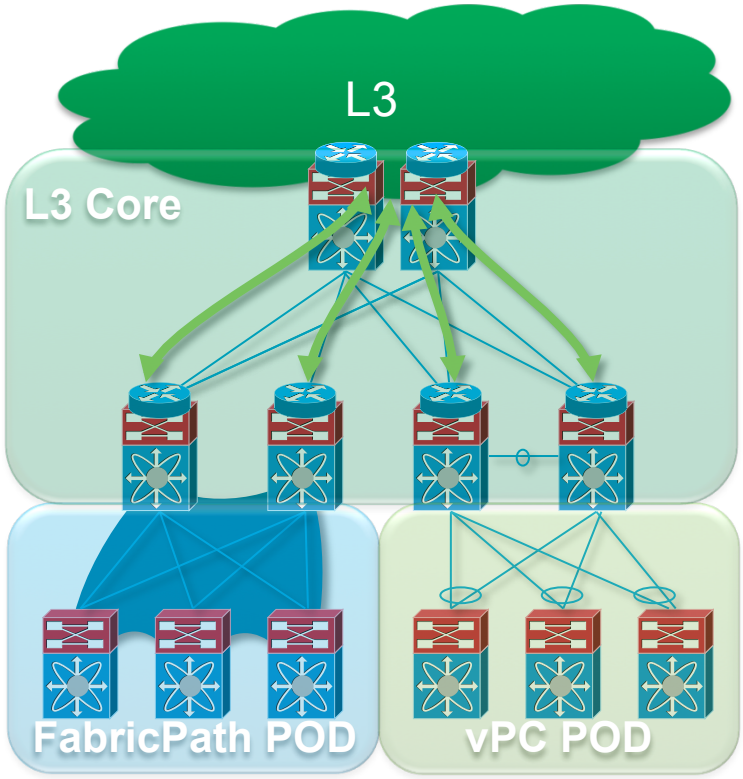

- Simple configuration
- No constraint in the design
- Seamless L3 integration
- No STP, no traditional bridging
- Mac address table scaling
- Virtually unlimited bandwidth
- Can extend easily and without operational impact

FabricPath Core

Efficient POD Interconnect

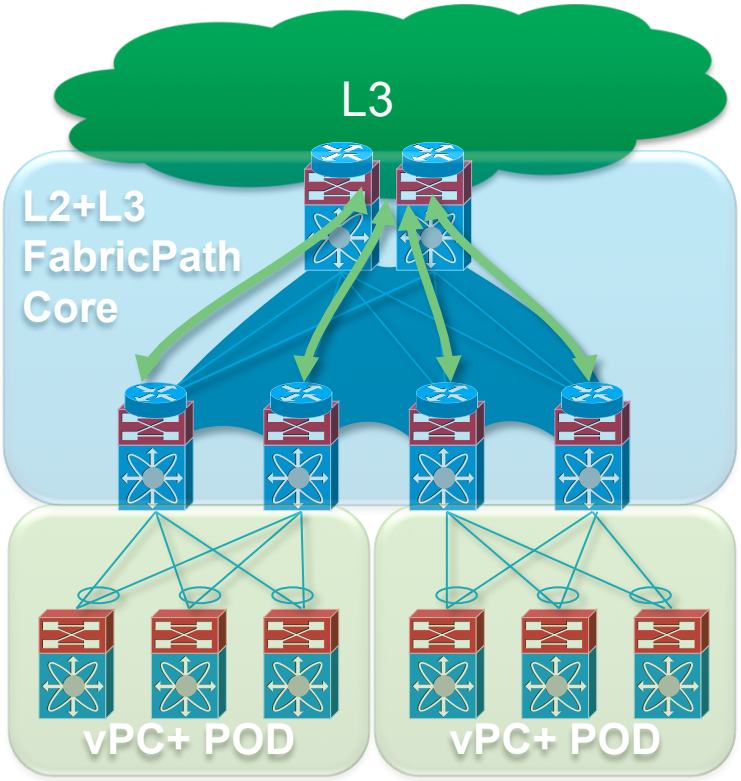

- FabricPath in the Core
- VLANs can terminate at the distribution **or** extend between PODs.
- STP is not extended between PODs, remote PODs or even remote data centers can be aggregated.
- Bandwidth or scale can be introduced in a non-disruptive way

FabricPath as Site Interconnect

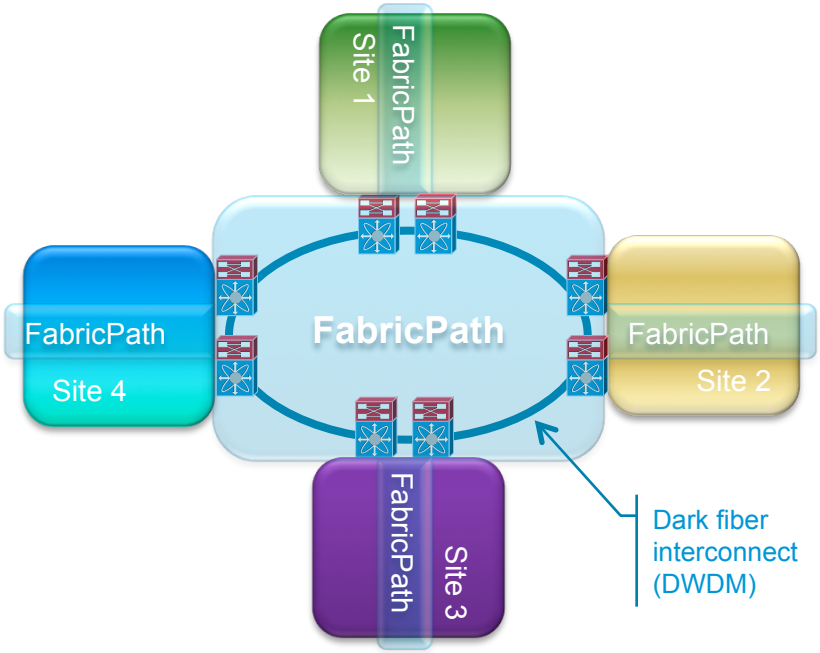

- Requires dark fiber
- Arbitrary interconnect topology (not dependent of port channels)
- Any number of sites
- High bandwidth, fast convergence
- Spanning tree isolation
- Mac address table scaling
- VLANs can be selectively extended/terminated

Internet Exchange Point (IXP)

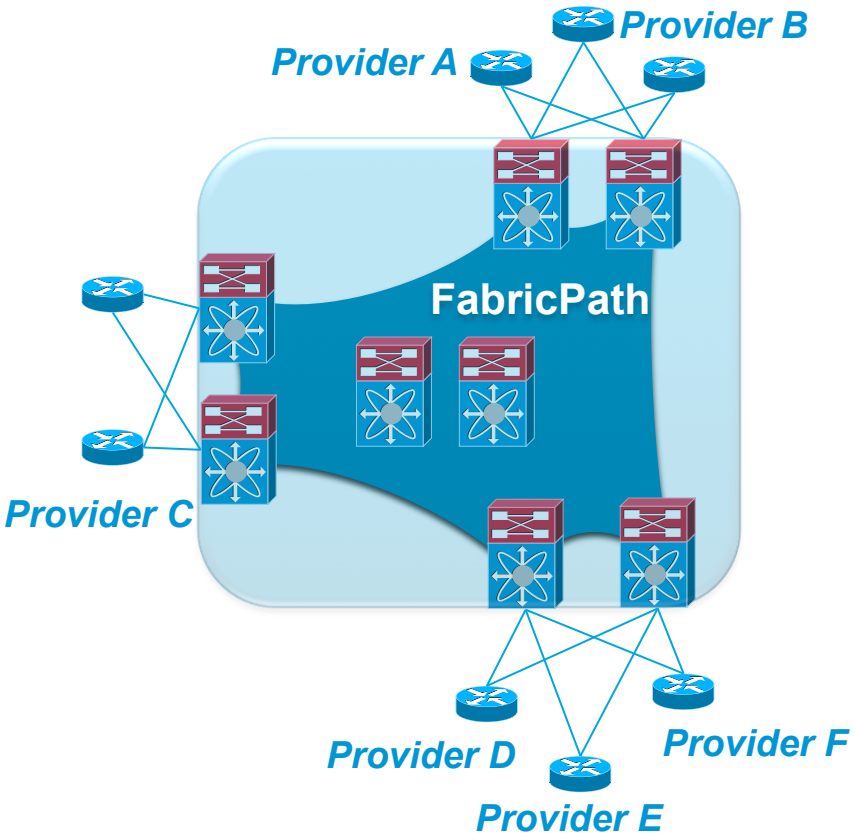

IXP Requirements

- Layer 2 Peering
- 10GE non-blocking Fabric
- Scale to thousands of ports

FabricPath Benefits for IXP

- Layer 2 Fabric
- Non-blocking up to thousands 10GE ports
- Simple to manage
- No design constraint, easy to grow

FabricPath Evolution

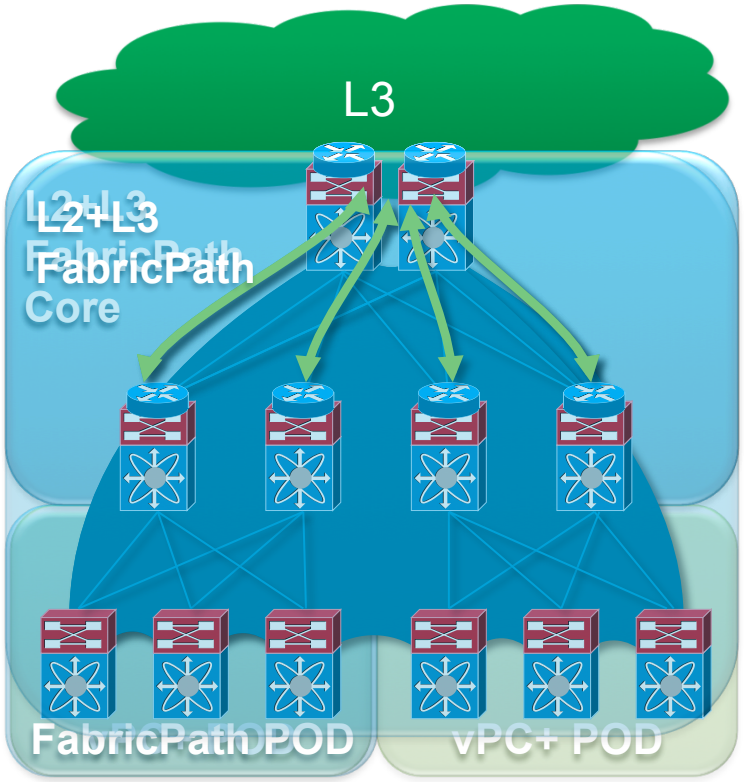

- FabricPath in the Core
- FabricPath extended down to the leaves

Tier Consolidation with FabricPath

"Flattening"

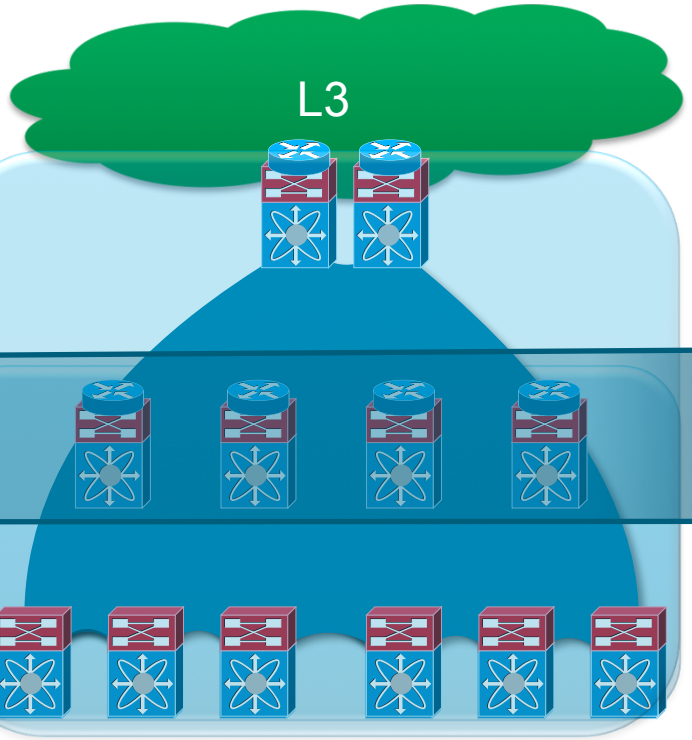

- FabricPath in the Core
- FabricPath extended down to the leaves

- There is enough bandwidth and port density on the core Nexus 7000s for aggregating the whole network.
- There is no need for a distribution layer for POD isolation

FabricPath Flexibility

The Network Can Evolve With No Disruption

- Need more edge ports? → Add more leaf switches
- Need more bandwidth? → Add more links and spines

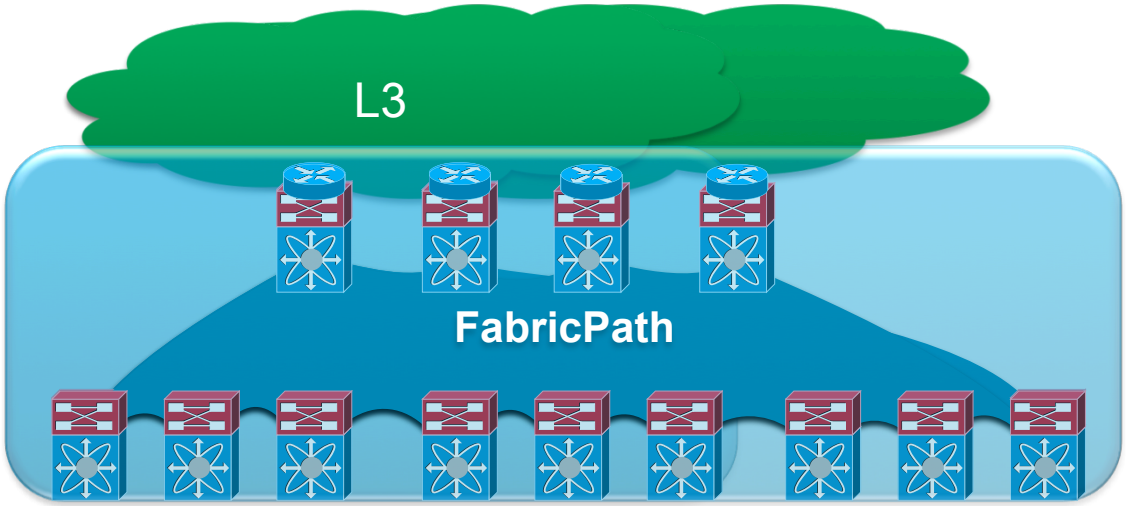

Scaling with FabricPath

Example: 2,048 x 10GE Server Design

- 16X improvement in bandwidth performance
- 6 to 1 consolidation (from 74 managed devices to 12 devices)
- 2X+ increase in network availability
- Simplified IT operations (fewer devices, vlans anywhere)

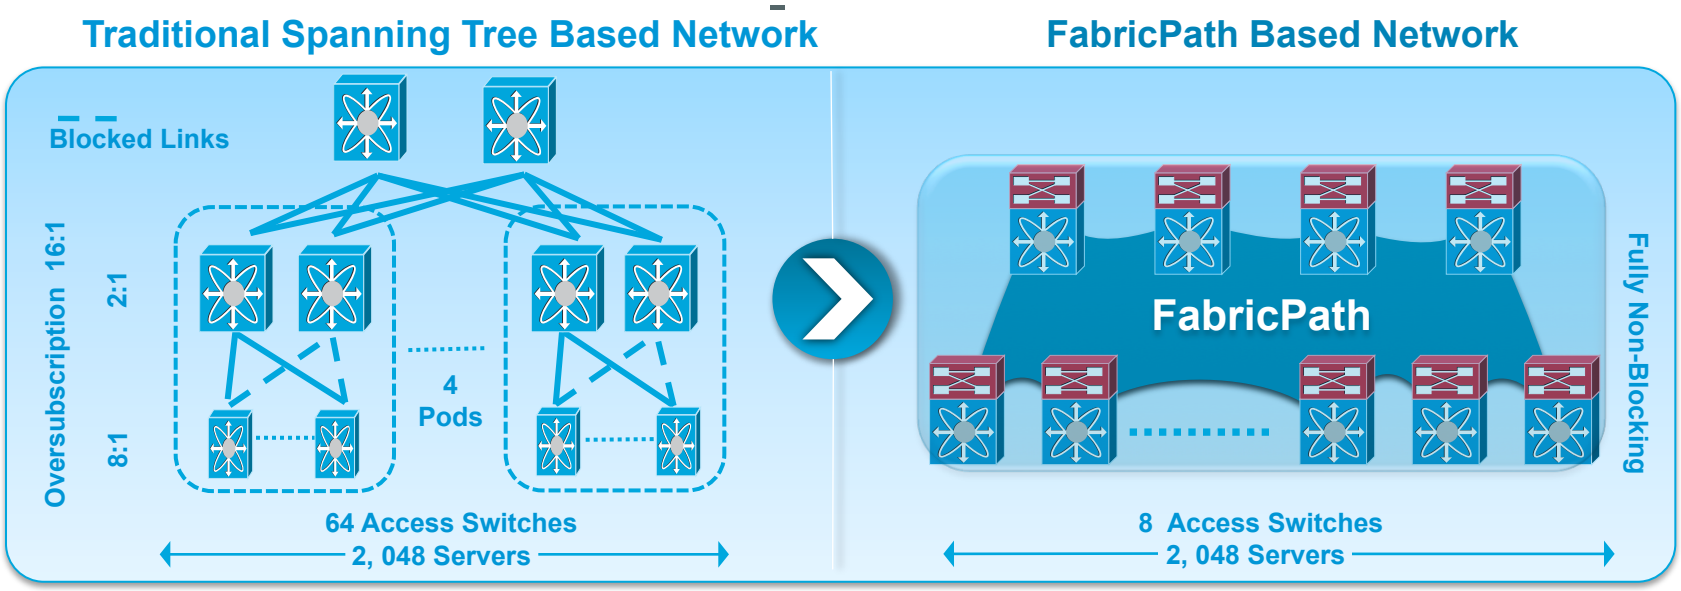

Layer 3/Service Integration, 2 Spine Switches

Business as Usual…

- Existing designs and best practices apply easily

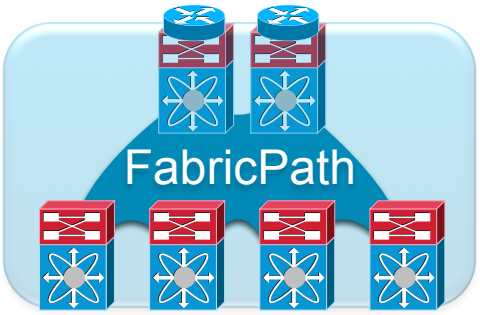

SVIs, active/active

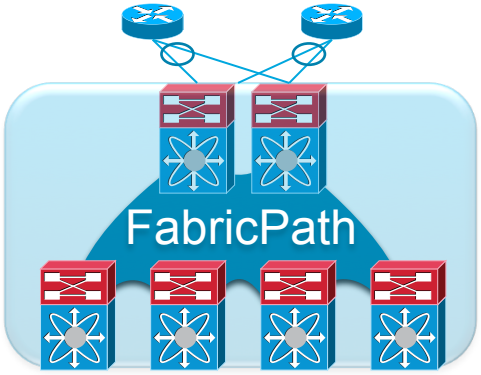

vPC+ external

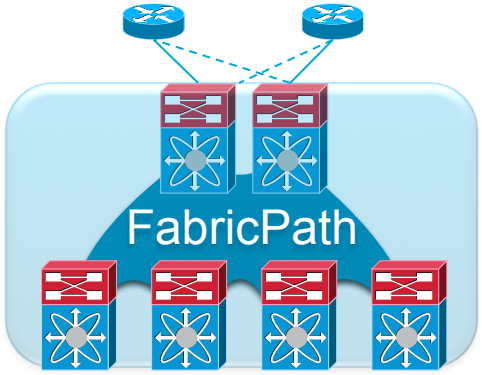

Active/Standby

Layer 3/Service Integration, >2 Spine Switches

Trying to adapt the existing model

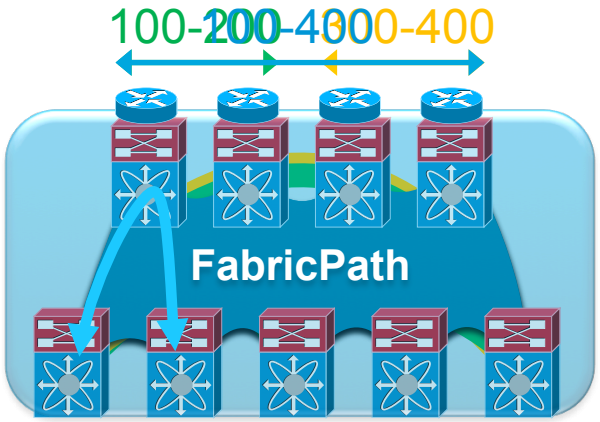

Split SVIs

- Host is pinned to a single VLAN
- Inter-VLAN traffic can be suboptimal

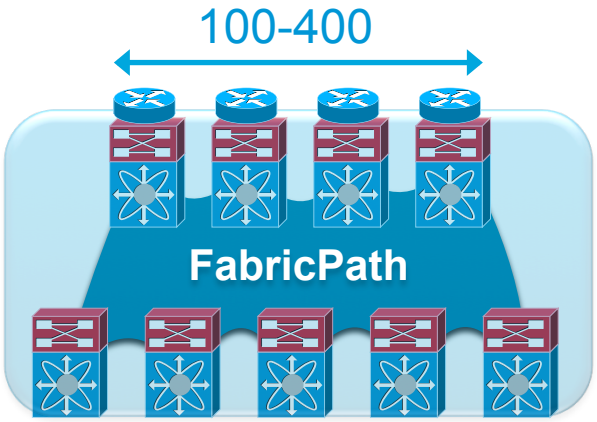

GLBP

- Host is pinned to a single VLAN
- Inter-VLAN traffic can be suboptimal
- Load balancing

Anycast FHRP

- All active
- Available in the future for routing
- Might not be available soon for service modules/appliances

Other L3/Service Integration Options

N-Way VPC+

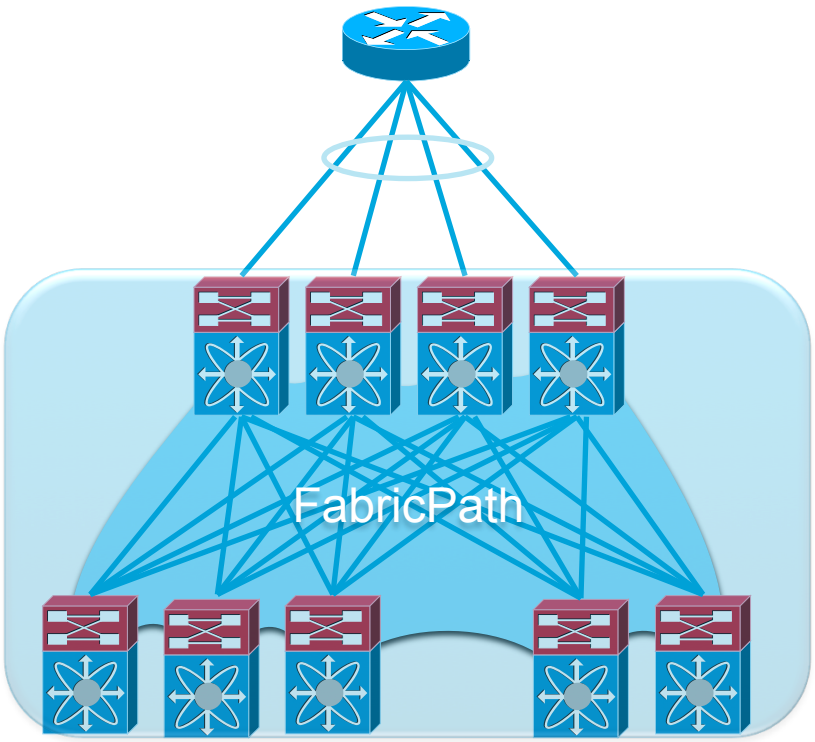

- Would allow to connect any device supporting port channeling
- Not in the roadmap

Other L3/Service Integration Options

FabricPath Capable Router/Appliance

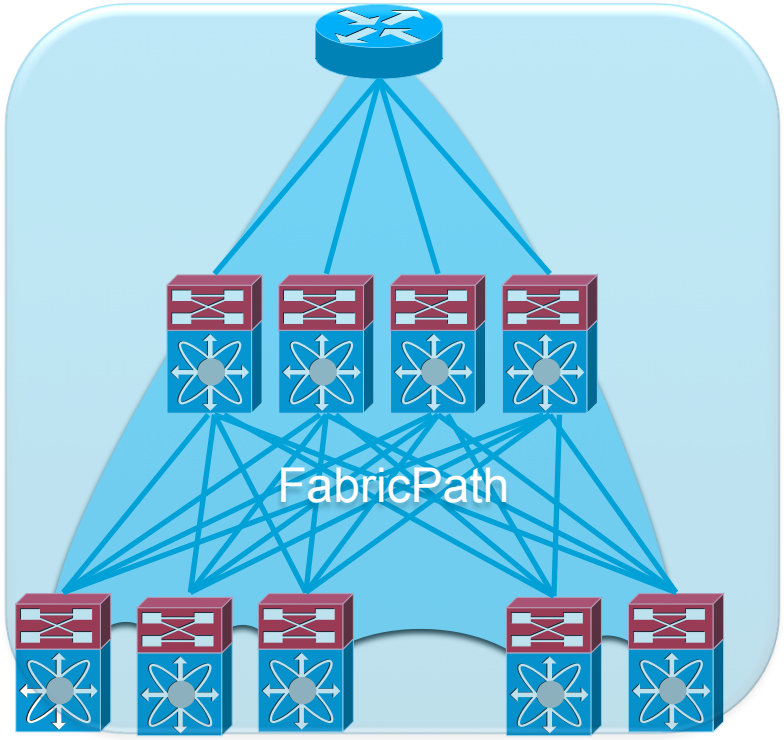

- Possible with Nexus platforms as routers
- Other appliances not FabricPath capable

Other L3/Service Integration Options

Insert a FabricPath Switch

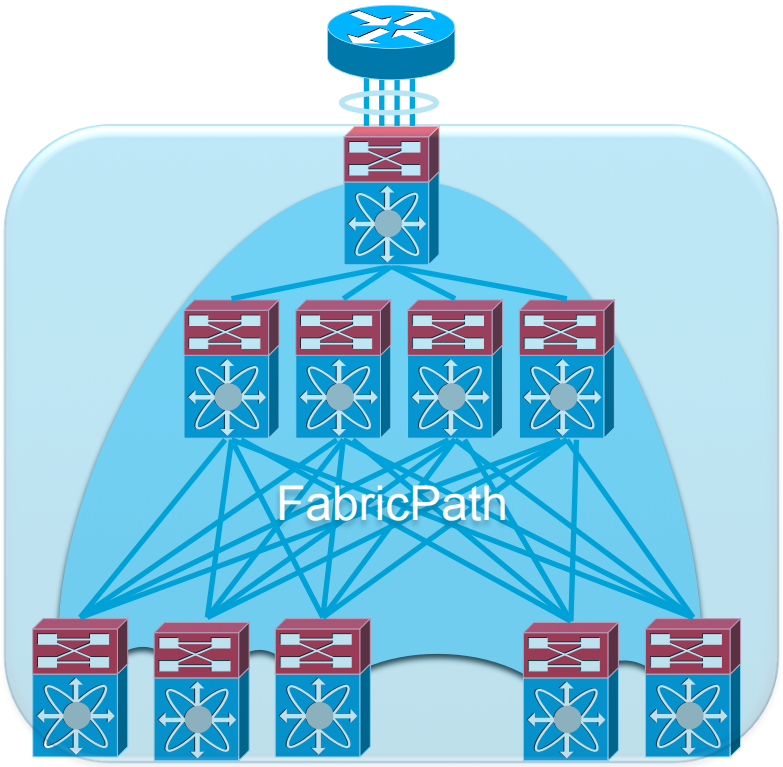

- Allows non-FabricPath capable devices to be attached to the fabric in an optimal way
- 2 hops to reach the router/ appliance
  → Nexus 7000 F series I/O modules and Nexus 5500 have lower latency anyway
- Thanks to FabricPath, enough bandwidth can be provisioned to the router/ appliance

Other L3/Service Integration Options

Attach to a Leaf Switch

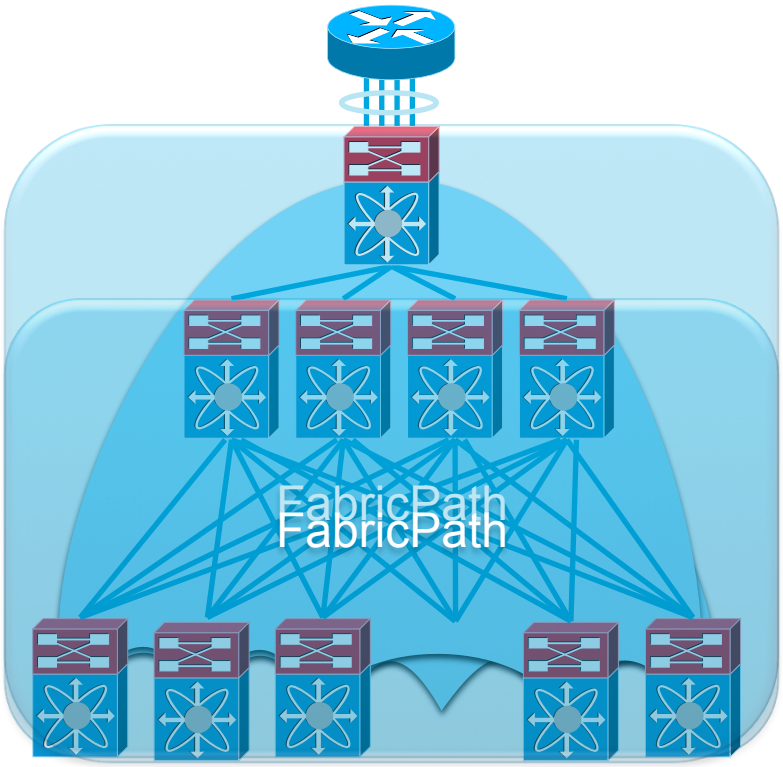

- Allows non-FabricPath capable devices to be attached to the fabric in an optimal way
- 2 hops to reach the router/ appliance
  → Nexus 7000 F series I/O modules and Nexus 5500 have lower latency anyway
- Thanks to FabricPath, enough bandwidth can be provisioned to the router/ appliance

FabricPath Edge Connectivity Options

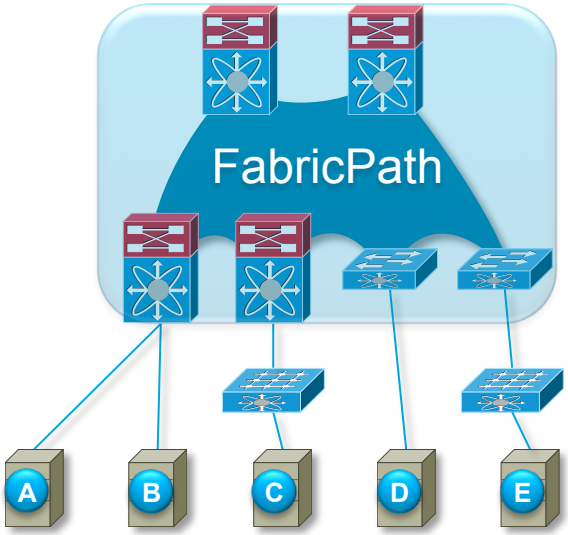

- **A** 1/10G connectivity to Nexus 7000 F1 I/O Module CX-1 cabling provide cost effective solution
- **B** 1/10G connectivity to Nexus 7000 F2 I/O Module
- **C** 1/10G connectivity to Fabric Extender attached to Nexus 7000 F2 I/O Module
- **D** 1/10G connectivity to Nexus 5500
- **E** 1/10G connectivity to Fabric Extender attached to Nexus 5500

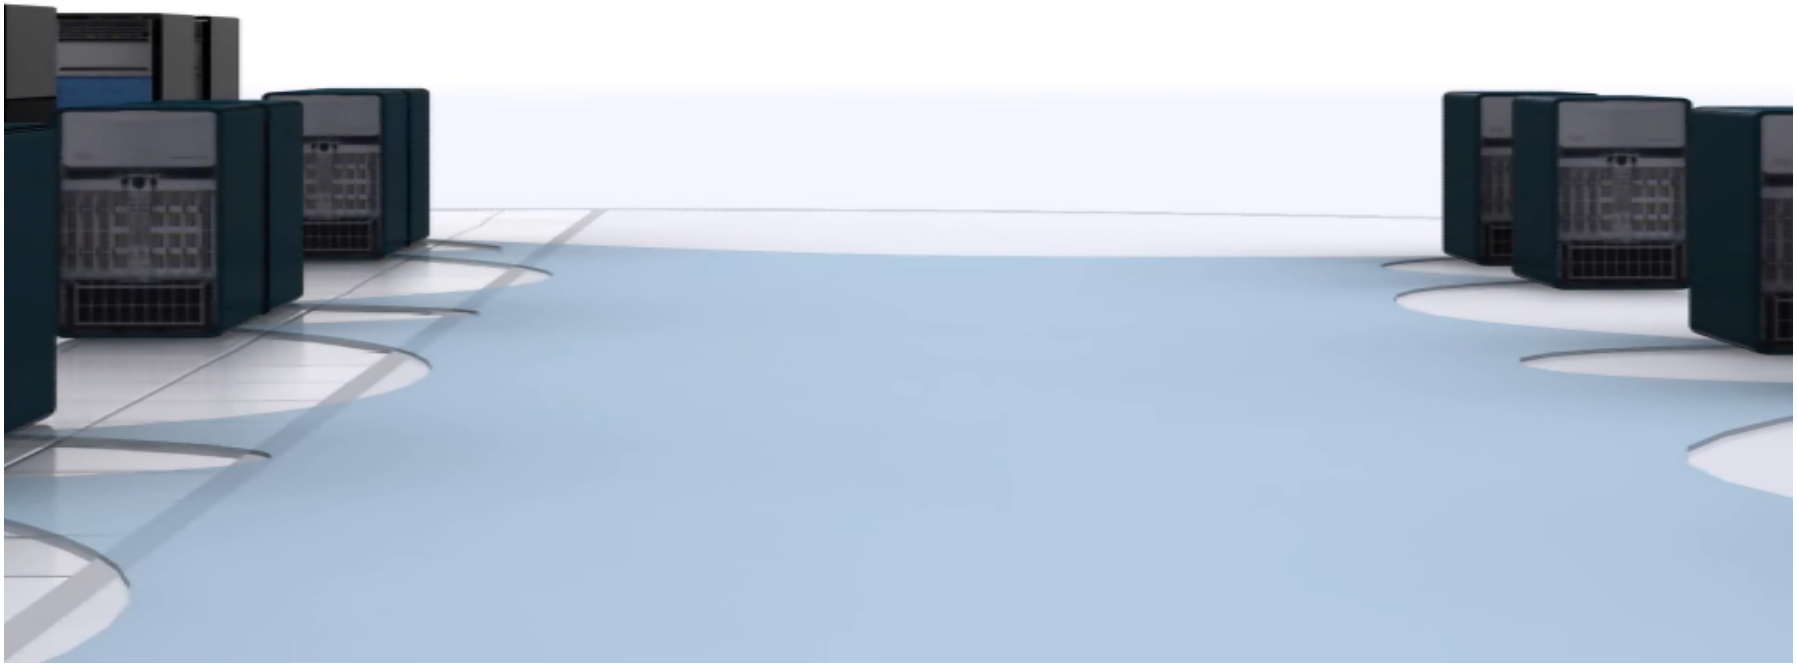

Conclusion

Agenda | Back | End
Why Layer 2 in the Data Center?
FabricPath Solution Overview
Inside the Fabric
FabricPath Unicast Details
FabricPath Multicast Details
FabricPath Designs

Key Takeaways

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

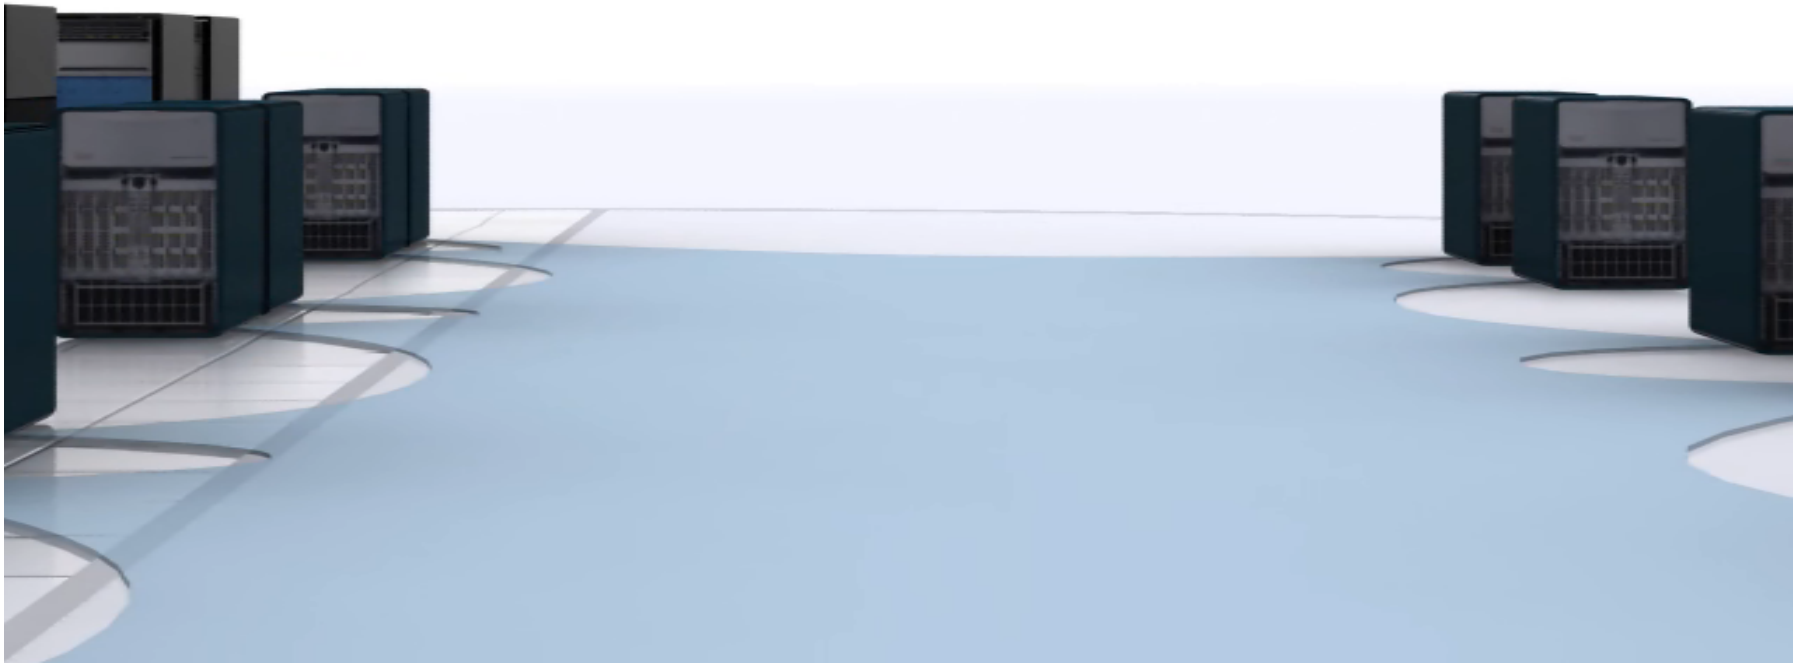

TRILL

TRILL

What is TRILL?

http://datatracker.ietf.org/wg/trill/

- TRILL, Transparent Interconnection of Lots of Links
- IETF standard for Layer 2 multipathing (since summer 2011)
- Pushed by Cisco but by other competitors too.

Why aren't we doing TRILL already?

- We delivered FabricPath before the standard, and some critical features are still missing to TRILL
- Our FabricPath capable hardware is also TRILL capable
- We will provide TRILL as a FabricPath mode

FabricPath vs. TRILL Overview

	FabricPath	TRILL
Frame routing (ECMP, TTL, RPFC etc…)	Yes	Yes
vPC+	Yes	No
FHRP active/active	Yes	No
Multiple topologies	Yes	No
Conversational learning	Yes	No
Inter-switch links	Point-to-point only	Point-to-point OR shared

- Cisco will push FabricPath specific enhancements to TRILL

- Thank you!
- Please complete the post-event survey.
- Join us September 5 for our next webinar:

A Closer Look: Comparing Benefits of EIGRP and OSPF

To register, go to www.cisco.com/go/iosadvantage

Thank you.