Massively Scalable Data Center

Tommi Saxelin
Data Center Specialist

14 September 2012

Why “Massively Scalable”?

The Good Ol' Internet

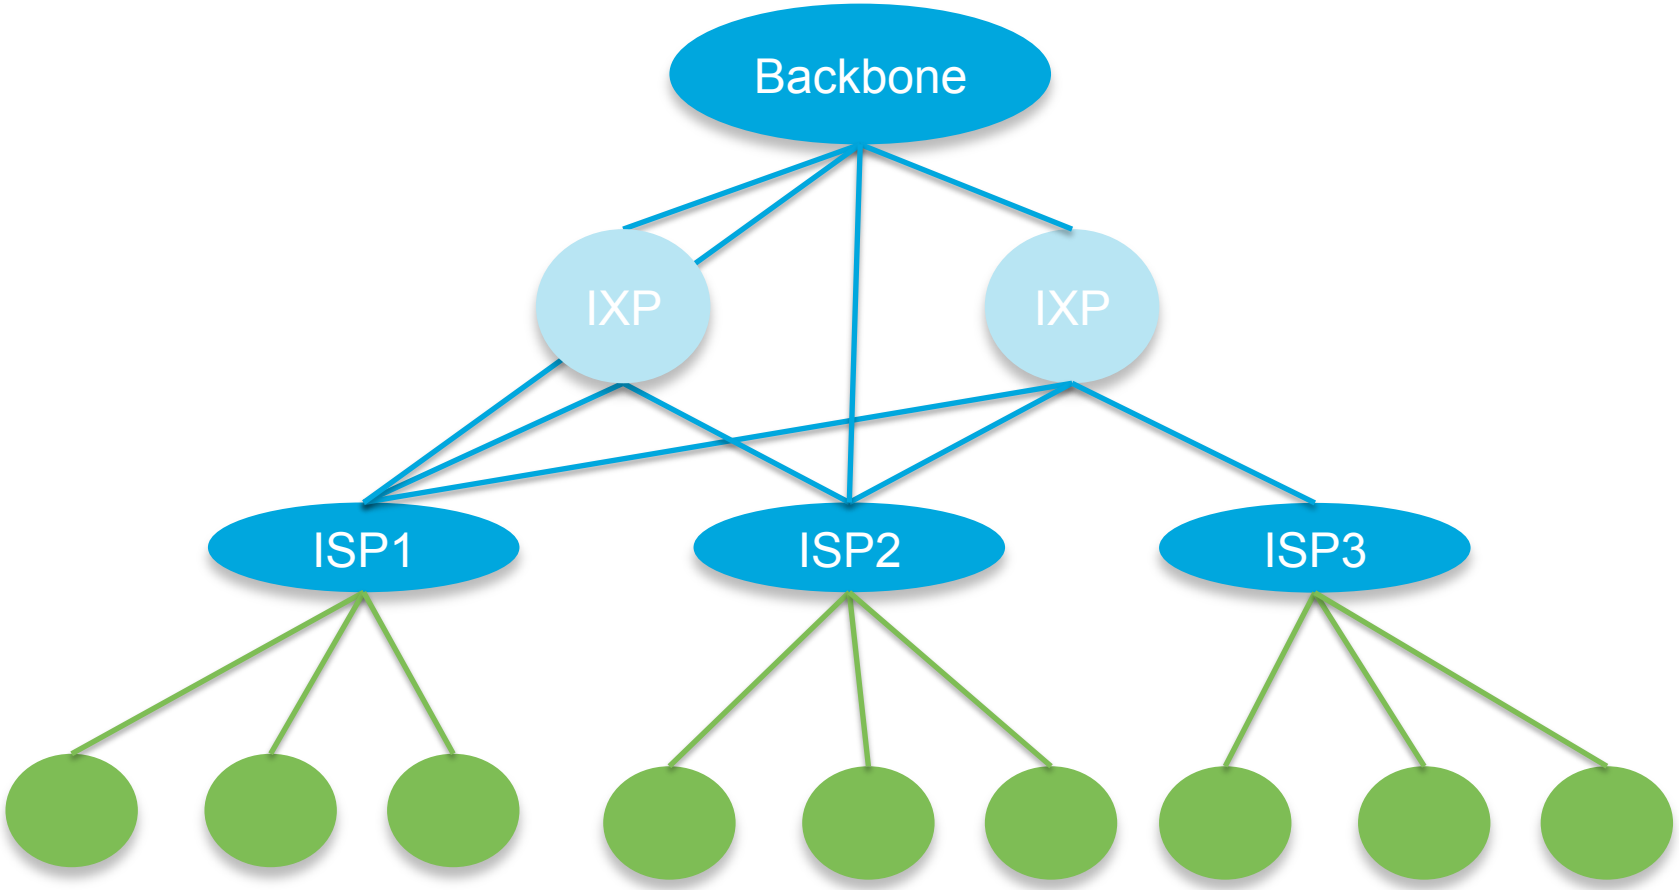

Warehouse-Scale Data Center Players

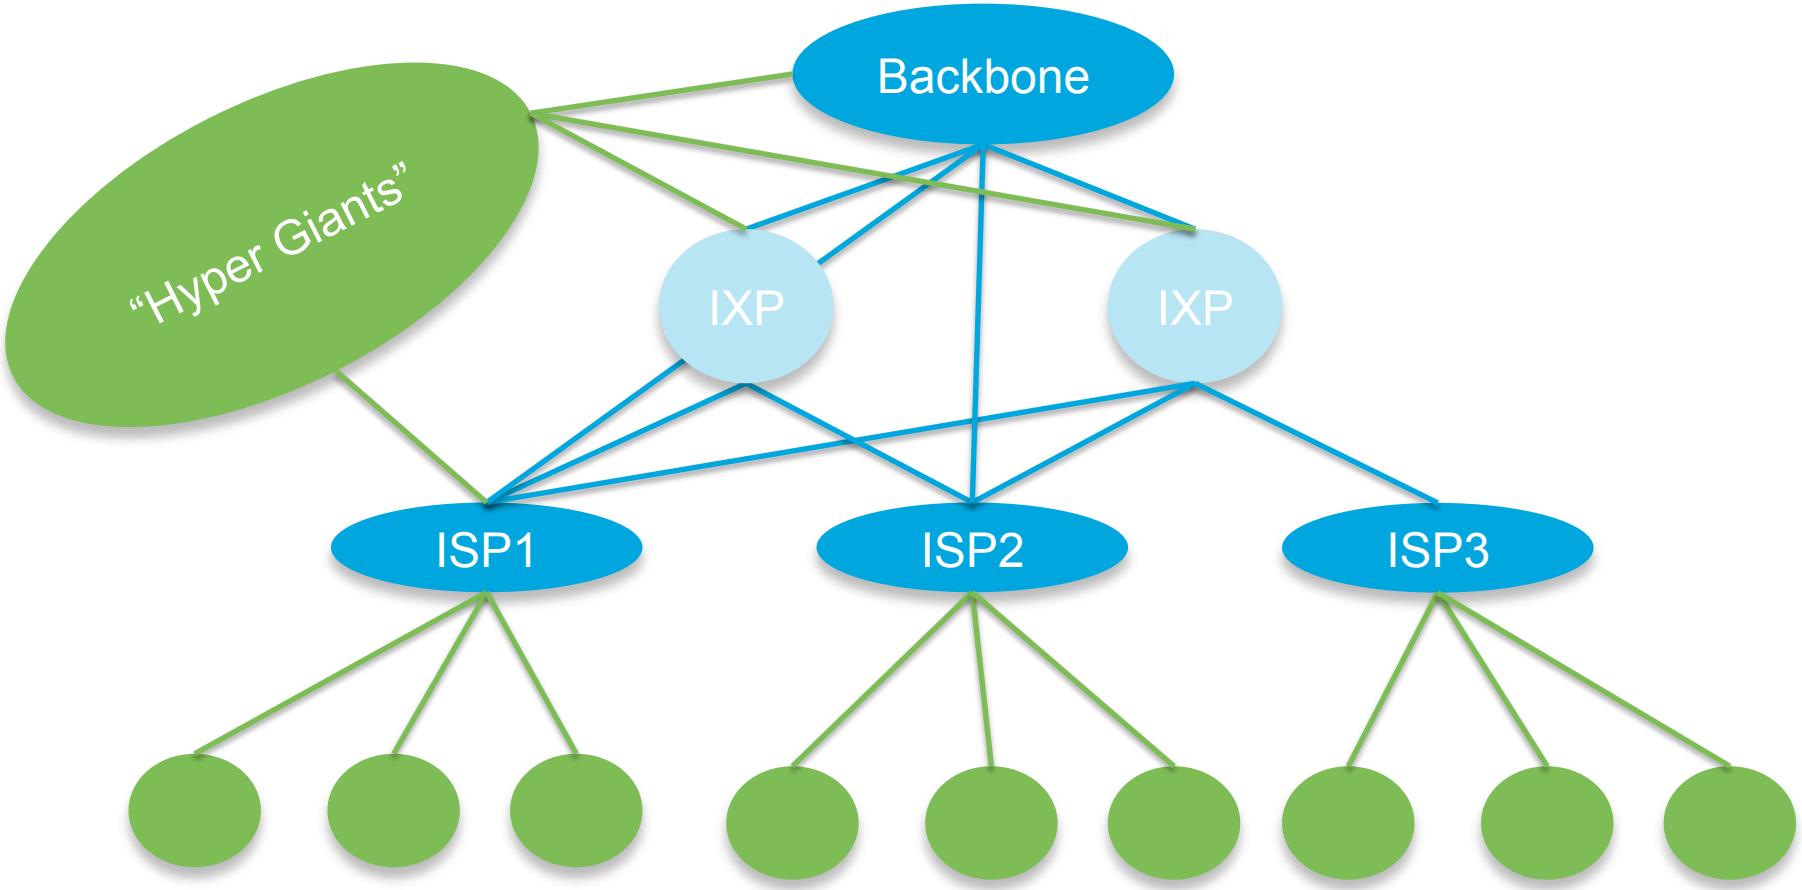

Massive Scale?

30,000+

physical servers / ports

How to Scale?

New Topologies Needed?

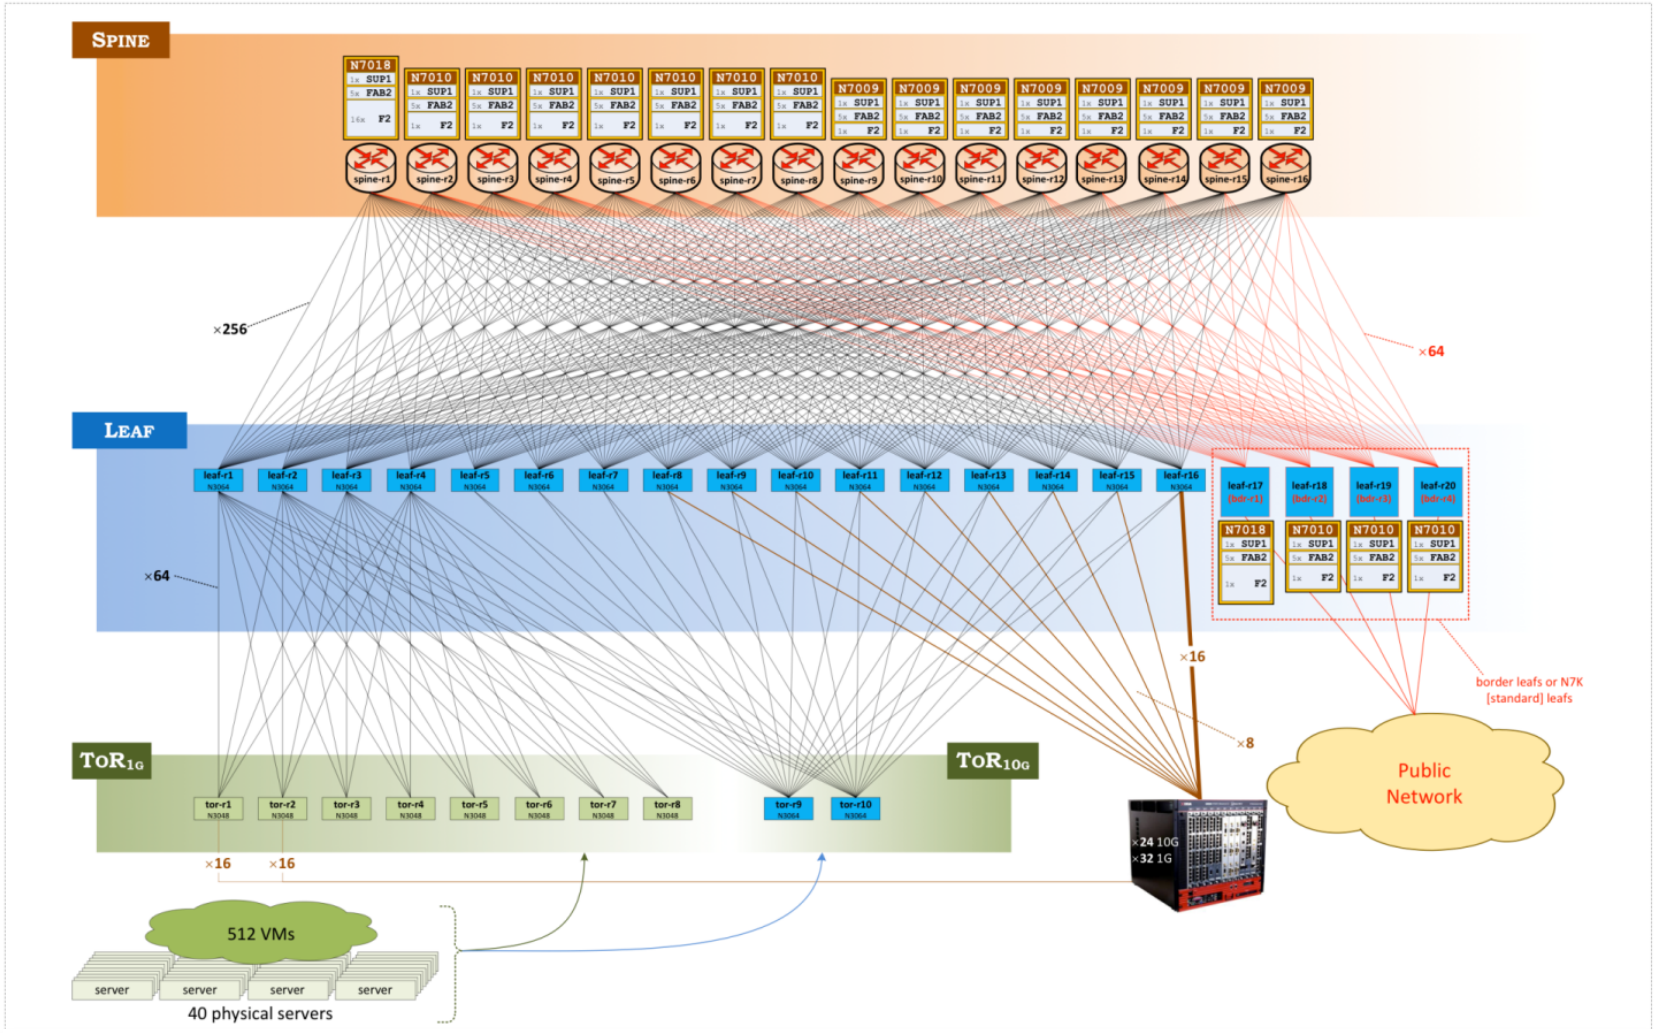

How About The Non-Giants?

Layer 2 Needed in The Data Center

- 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

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

FabricPath

Cisco FabricPath Goal

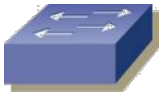

Switching

- Easy Configuration
- Plug & Play
- Provisioning Flexibility

Routing

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

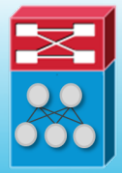

FabricPath

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

FabricPath, an Ethernet Fabric

Shipping Since 2010: turn your network into a Fabric

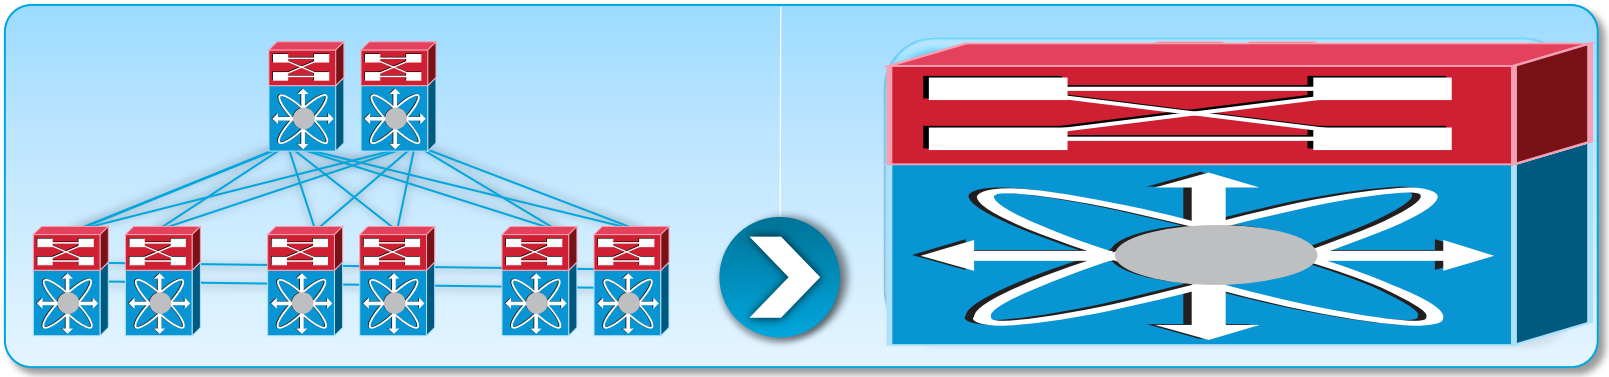

- Connect a group of switches using an **arbitrary** topology
- With a simple CLI, aggregate them into a Fabric:

```
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

Mac Address Table Scale

Conversational Learning

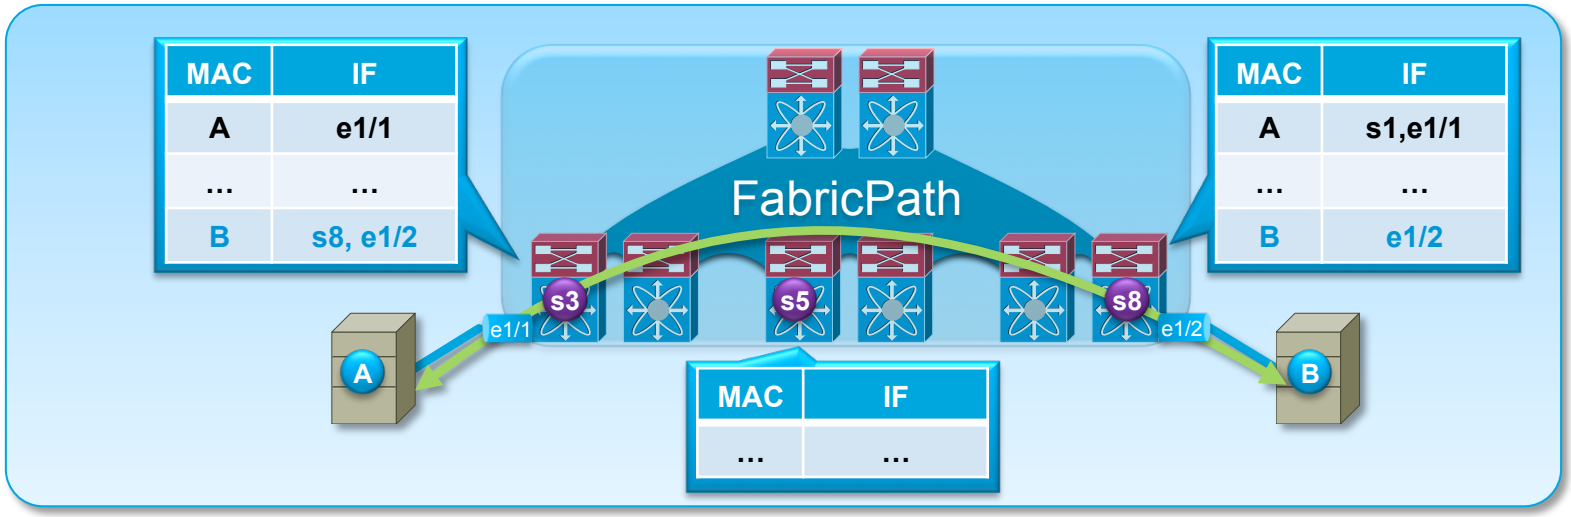

- Per-port mac address table only needs to learn the peers that are reached across the fabric

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

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

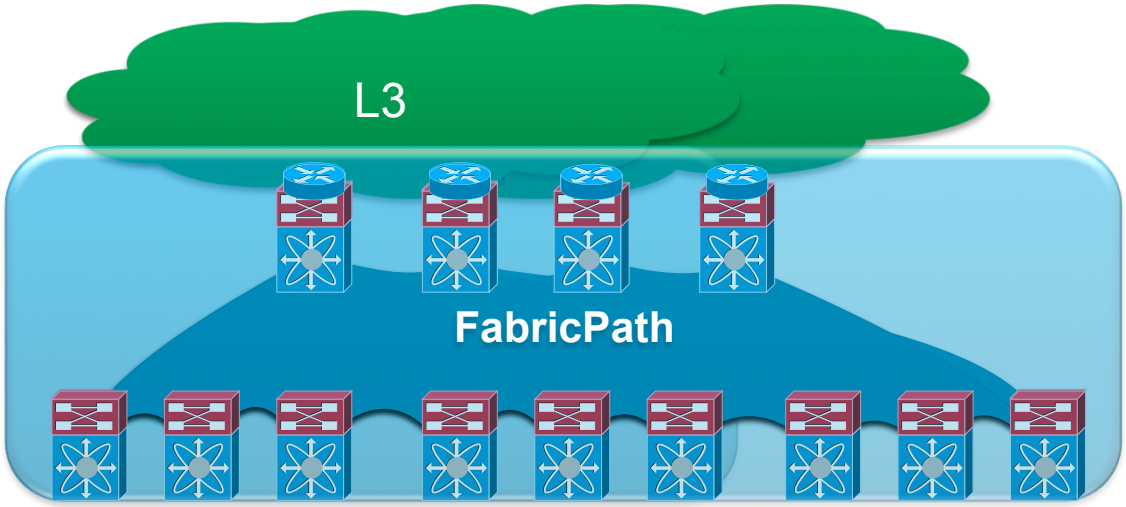

Scaling with FabricPath

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)

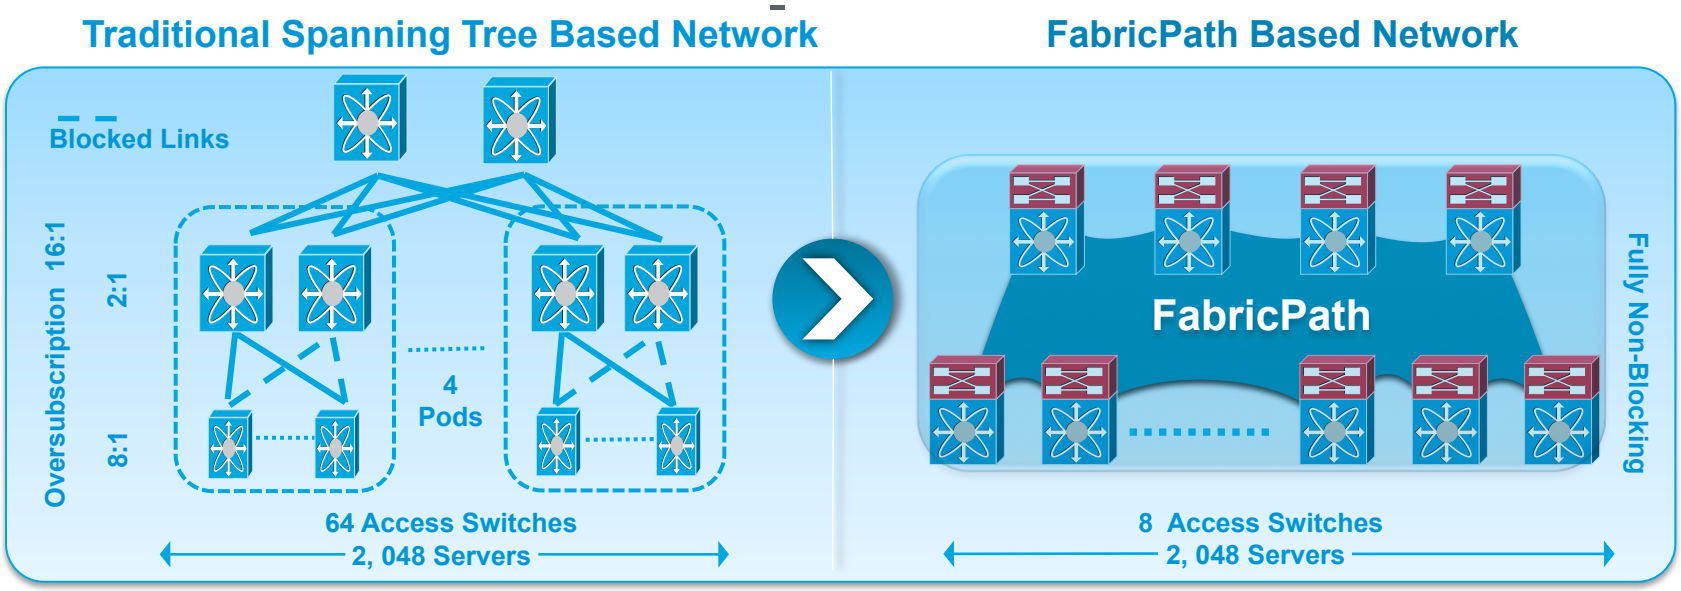

FabricPath vs. TRILL Overview

	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

- FabricPath will provide a TRILL mode with a software upgrade (hardware is already TRILL capable)
- Cisco will push FabricPath specific enhancements to TRILL

Example Use Case: Internet Exchange Point (IXP)

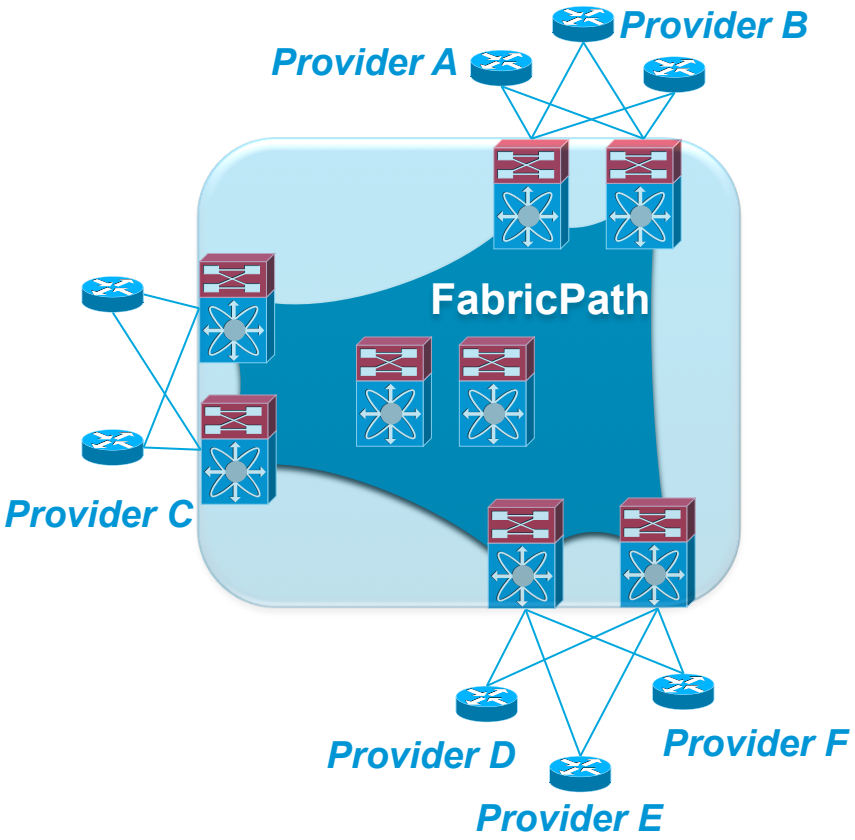

IXP Requirements

- 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

Software-Defined Networking

Keyword: Programmability

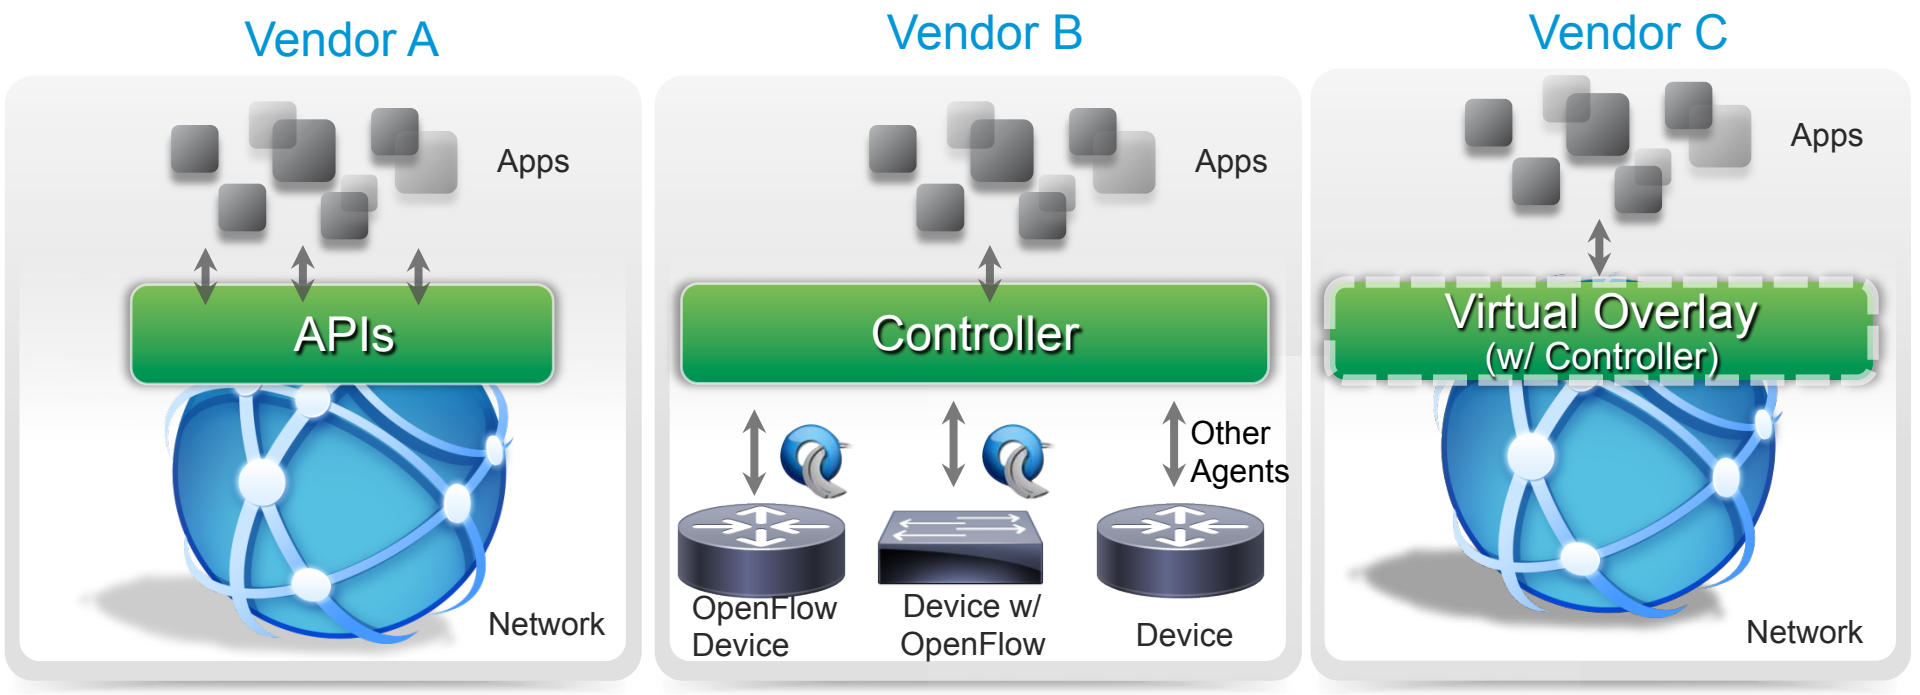

Cisco Approach : Flexibility to choose - The Power of "AND"

Cisco Open Network Environment

Platform APIs	Controller/Agents	Virtual Overlay Networks
onePK	**SDN Controller OpenFlow Agents**	**Nexus 1000V**
Comprehensive Developer Kit across IOS, IOS-XR and NX-OS	Proof-of-Concept Controller software for SDN research OpenFlow Agents on certain switching platforms	OpenStack and REST API Multi-Hypervisors VXLAN Gateway Security, Services Chaining

Thank you.