

Cisco Nexus Yeni Nesil Switching Çözümleri

Celal Milletinisever
Senior Network Consultant, CCIE
Biltam



Ajanda

- DC Networking and Design
- Data Center Fabric Journey
- Nexus Family Architecture
- NX-OS Features
- Cisco ACI

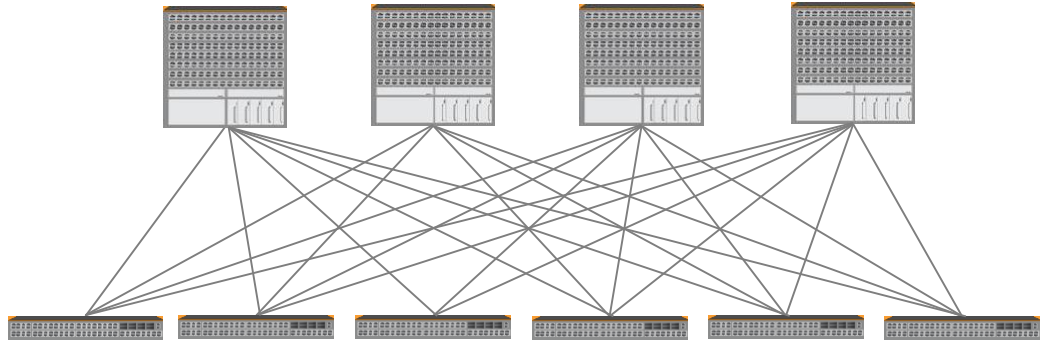


DC Networking and Design



Design Requirements

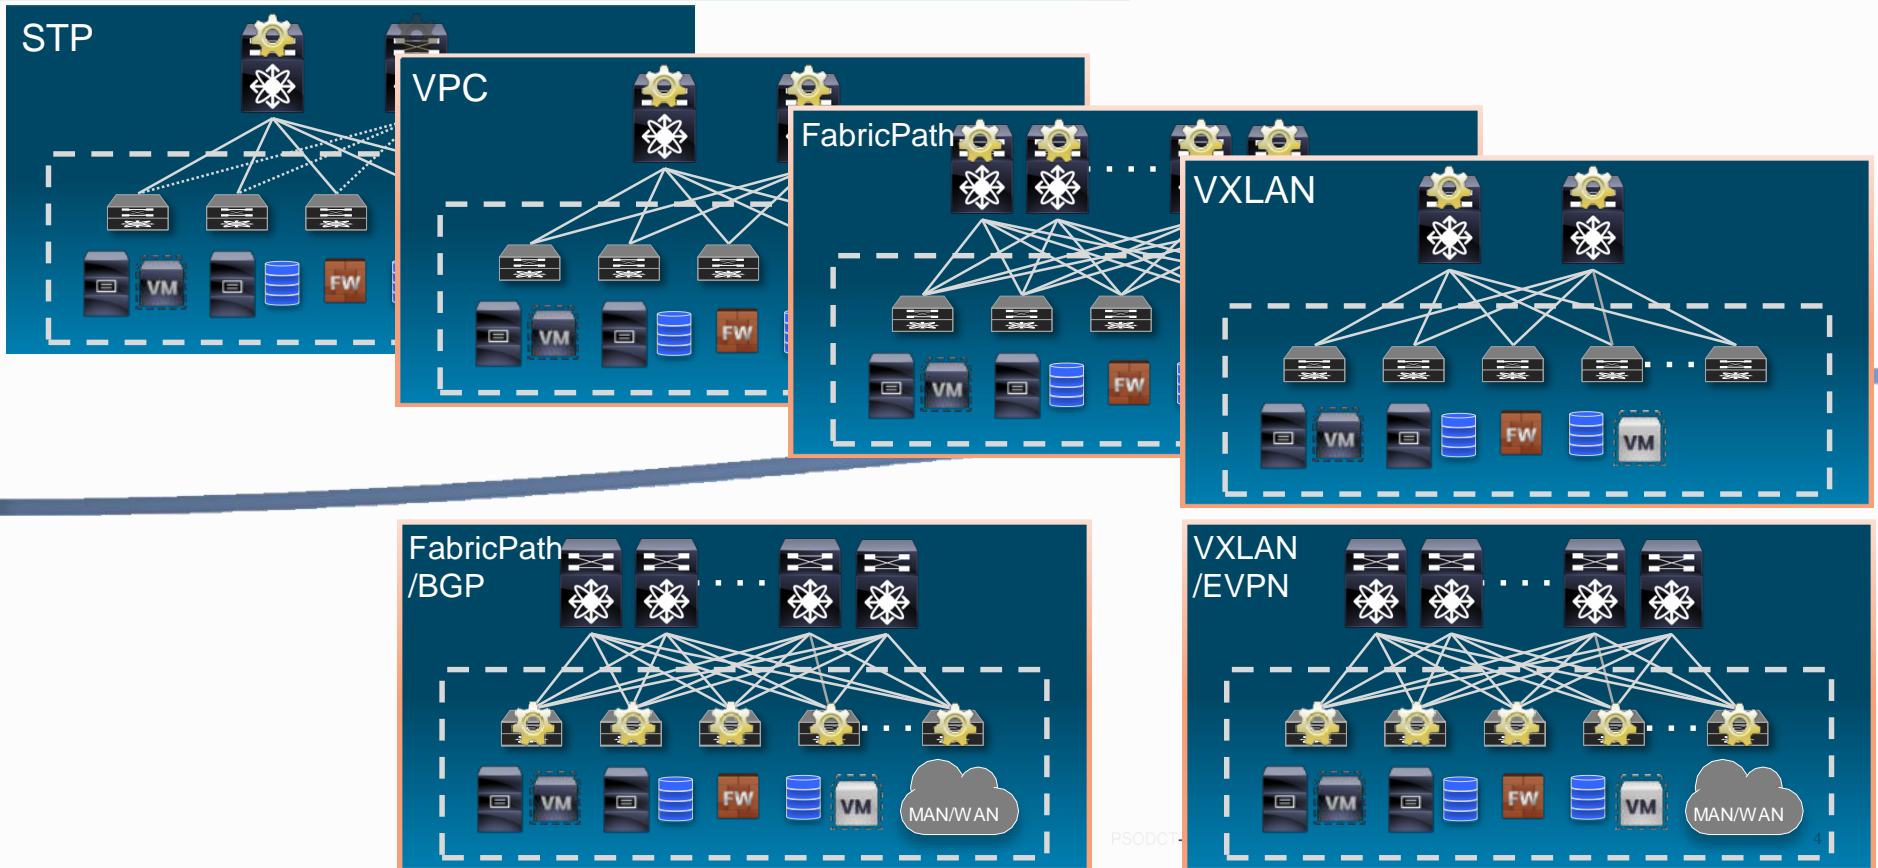
- Redundancy
- Performance
- Flexibility
- Scalability
- Security
- Predictably grow the data center



Data Center Fabric Journey



Common DC Topologies



Spanning-Tree vs vPC

How does vPC help with STP?

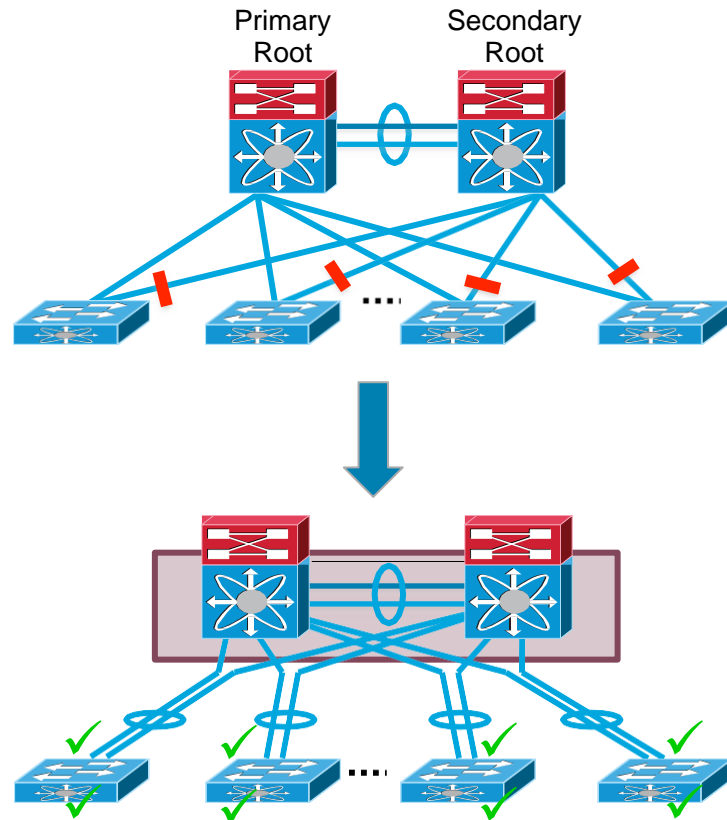
- **Before vPC**

- STP blocks redundant uplinks
- VLAN based load balancing
- Loop Resolution relies on STP
- Protocol Failure



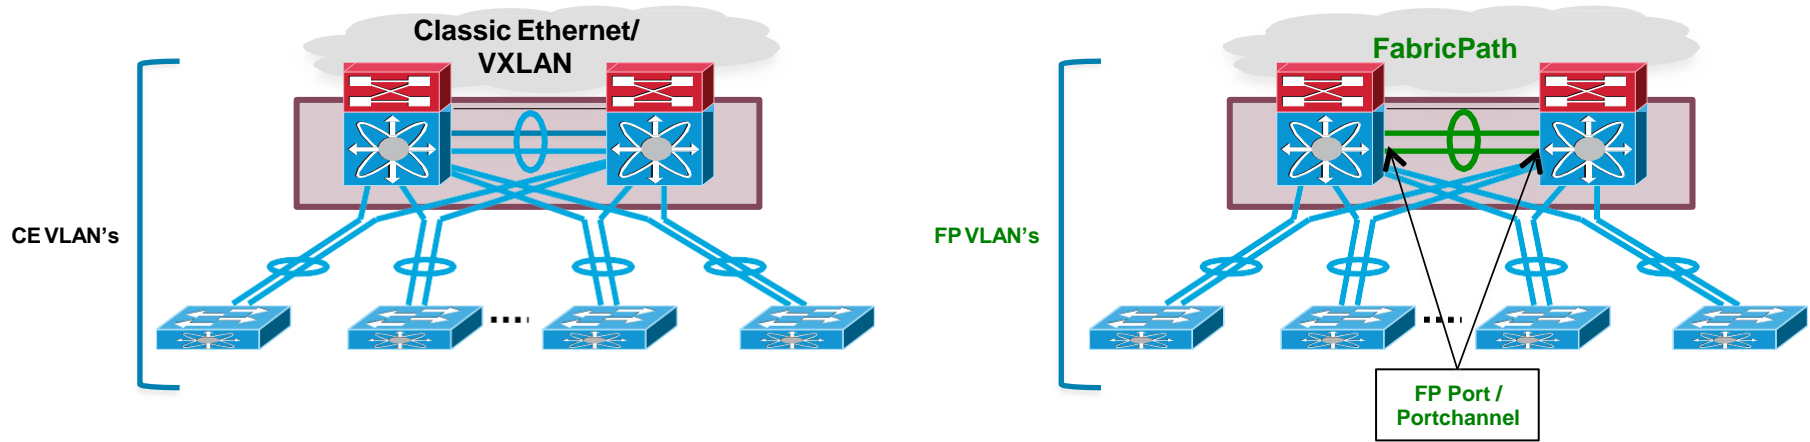
- **With vPC**

- No blocked uplinks
- EtherChannel load balancing (hash)
- Loop Free Topology
- Lower oversubscription



vPC vs vPC+

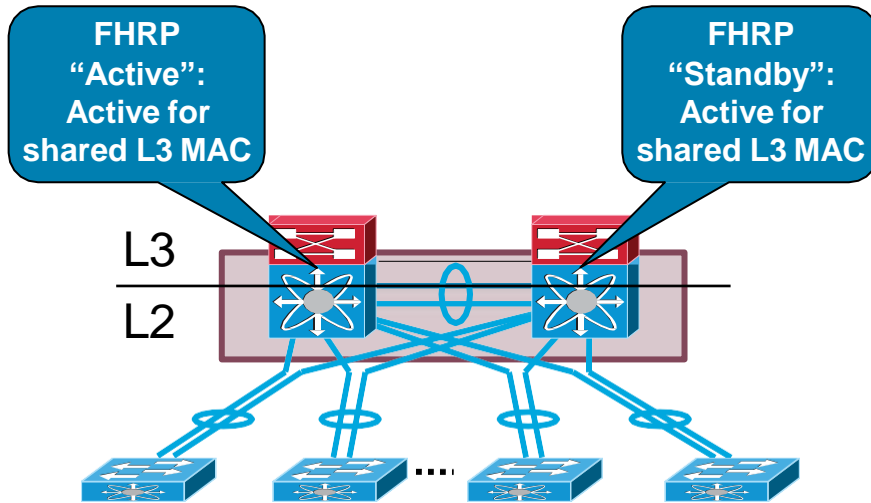
Architecture of vPC and FabricPath with vPC+



- Physical architecture of vPC and vPC+ is the same from the access edge
- Functionality/Concepts of vPC and vPC+ are the same
- VXLAN uses VPC while FabricPath uses VPC+

FHRP with VPC/VPC+

HSRP / VRRP/ GLBP Active/Active



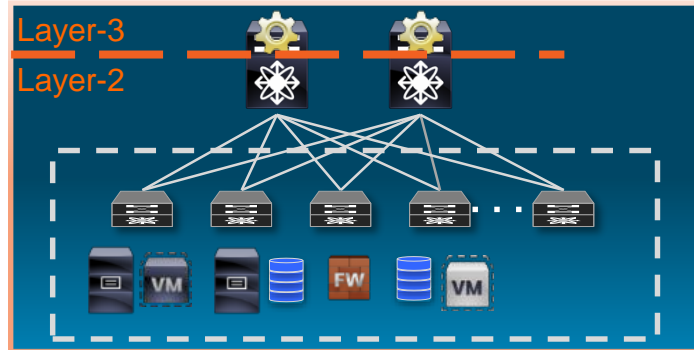
- Support for all FHRP protocols in Active/Active mode with vPC/VPC+
- No additional configuration required
- Standby device communicates with vPC manager to determine
- if vPC peer is “Active” FHRP peer
- ‘peer-gateway’ command allows a vPC peer to respond both the FHRP virtual and the real MAC address of both itself and it’s peer

VxLAN

What is VXLAN?

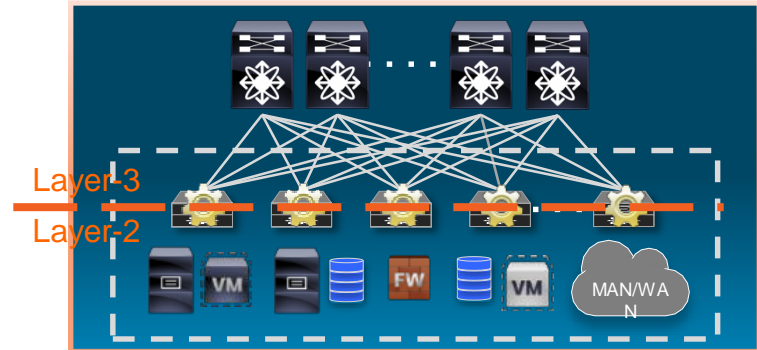
- VXLAN is a Network Virtualization and Overlay Technology
- VXLAN builds Layer-2 & Layer-3 overlay Networks on top of an IP routed Network
- VXLAN uses MAC in UDP encapsulation (UDP destination port 4789)

VXLAN Flood & Learn



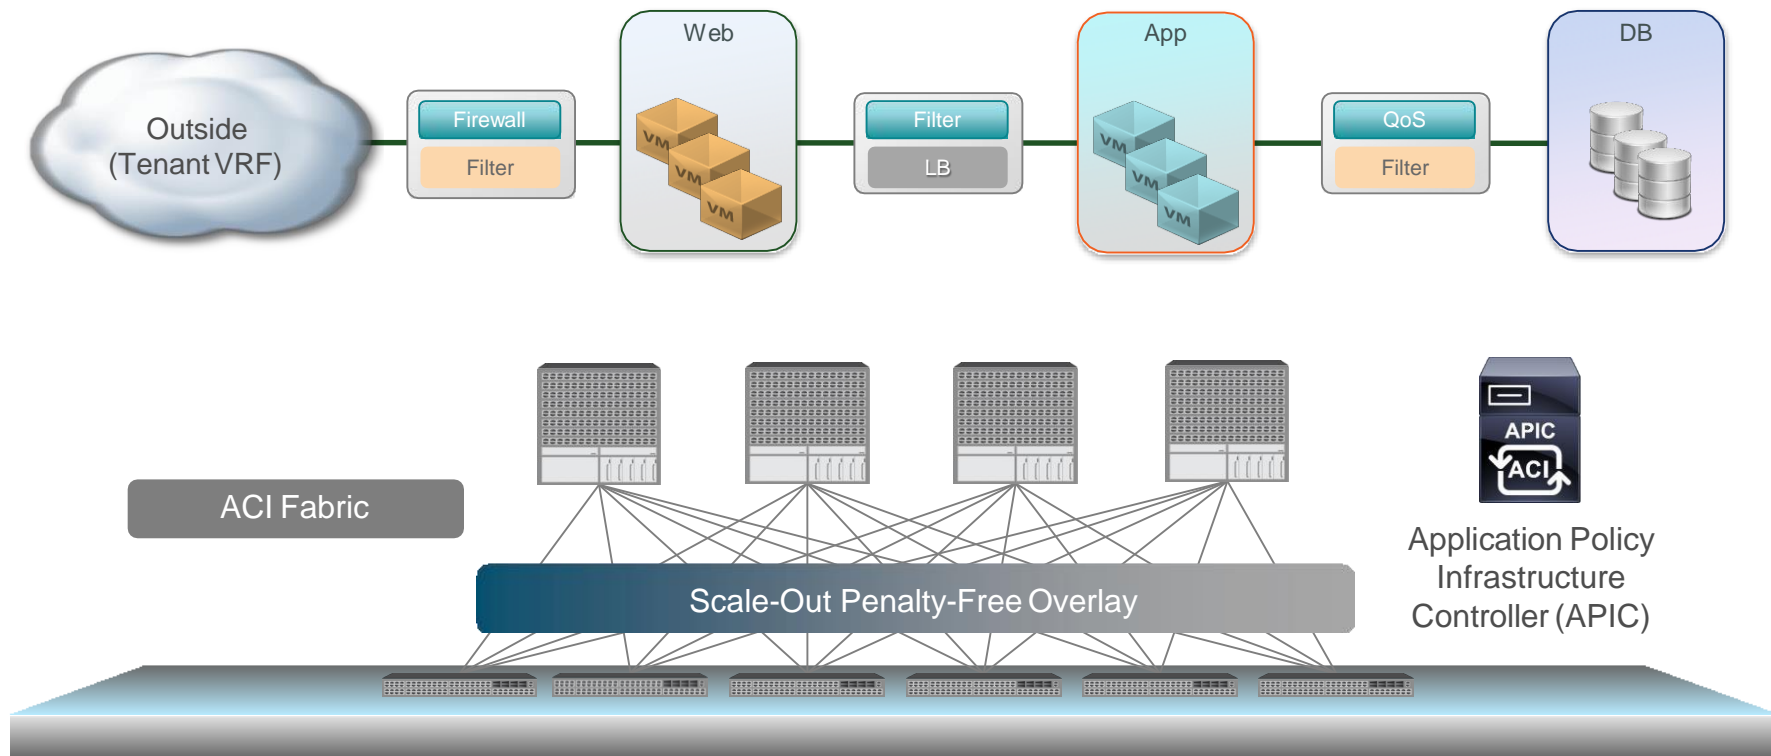
SVIs, active/active

VXLAN with BGP EVPN



SVIs with Distributed Anycast Gateway

ACI – Application Centric Infrastructure



Nexus Family
Architecture



Cisco Nexus Data Center Switching Portfolio

Virtual Switching

Fixed Switching Platforms



Nexus 3100



Nexus 3000



Nexus 5600



Nexus 5000/5500



Nexus 6000

Modular Switching Platforms



Nexus 7000



Nexus 7700



Application Centric Infrastructure (ACI)



Nexus 1000V

Nexus 2200 / 2300



Nexus 9000 & APIC

Nexus 7000 Series



Cisco Nexus® 7000 Series

Cisco Nexus 7000 Series
Switches

Cisco Nexus 7700
Platform Switches

Common Cisco® NX-OS

Modular design, grow as the network grows
High-density 10, 40, and 100 Gigabit Ethernet

Open source programmable tools for automation and orchestration

Industry's Most Proven Data Center Switching Platform

F3 Module Portfolio

Industry's Most Feature Rich Data Center

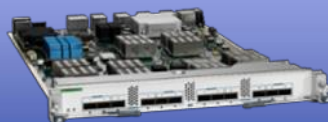
Nexus 7000 10G

F3 48-port 10G Module
N7K-F348XP-25



Nexus 7000 40G

F3 12-port 40G Module
N7K-F312FQ-25



Nexus 7000 100G

F3 6-port 100G Module
N7K-F306CK-25



Nexus 7700 10G

F3 48-port 10G Module
N77-F348XP-23



Nexus 7700 40G

F3 24-port 40G Module
N77-F324FQ-25



Nexus 7700 100G

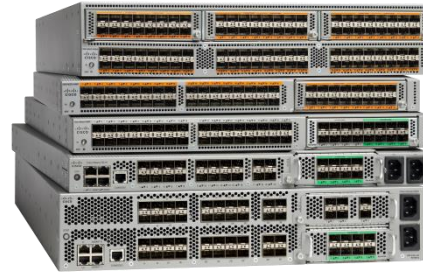
F3 12-port 100G Module
N77-F312CK-26



Nexus 5000 Series

Expanding the Family with the Nexus

Cisco Nexus® 5600 Series







Common Cisco® NX-OS

Modular and fixed port designs
High-density 10, and 40 Gigabit Ethernet

Open source programmable tools for automation and orchestration

Nexus 5600 is the 3rd generation of 5000 series switches

Nexus 5K/6K Fixed Switching Platforms

				
Model	Nexus 5672UP	Nexus 56128P	Nexus 5648Q	Nexus 5696Q
Switch Fabric Throughput	1.44Tbps	2.56Tbps	3.84Tbps	7.68Tbps
10Gbe Ports	72x 10G (48x 10G + 6x 40G with Breakout)	128x 10G (96x 10G + 8x 40G with Breakout)	192x 10G (48x 40G with Breakout)	384x 10G (384x 10G + 96x 40G with Breakout)
40Gbe Ports	6x 40G (true 40G or 4x 10G-Mode)	up to 8x 40G (true 40G or 4x 10G-Mode)	48x 40G (true 40G or 192x 10G-Mode)	up to 96x 40G (true 40G or 4x 10G-Mode)
Layer-3 Capability	Native line rate L3	Native line rate L3	Native line rate L3	Native line rate L3
Unified Ports	✓	✓	✗	✓
Switch Footprint	1RU	2RU	2RU	3RU

Nexus 3000 Series

Cisco Nexus® 3500 Series



Common Cisco® NX-OS

Exceptional switching latencies of less than 200 nanoseconds
Over 1000 times lower than Nexus 9000

Open source programmable tools for automation and orchestration

Cost-effective, power-efficient line-rate switches

Nexus 3000 Series

Expanding the Family with the Nexus 3500

Feature	Nexus 3172PQ	Nexus 31128PQ	Nexus 31108TC-V	Nexus 3132Q-V	Nexus 3232C	Nexus 3264Q	Nexus 3524-X	Nexus 3548-X
VXLAN capability	Yes	Yes	Yes	Yes	Bridging	Bridging	No	No
Rack units	1	2	1	1	1	2	1	1
Switching capacity	1.44 Tbps	2.4Tbps	2.8Tbps	2.8Tbps	6.4 Tbps	5.12 Tbps	480 Gbps	960 Gbps
Interface type	48 SFP+ and 6 QSFP+	96 SFP+ and 8 QSFP+	48 RJ-45 and 6 QSFP28	32 QSFP	32 QSFP28	64 QSFP+	24 SFP+	48 SFP+
Max number of 10 ports	72	96	72	128	128	128	24	48
Max number of 40G ports	6	8	6	32	32	64	6	12
Switch latency	-	-	650 ns	650 ns	450 ns	515 to 700ns	≤ 250ns	≤ 250ns

Nexus 2000 Series Fabric Extenders

Cisco Nexus® 2000 Series



Low OPEX Low CAPEX

Lower Devices to Manage Central Point of Management

Enhanced Buffers Low over subscription High Port Density

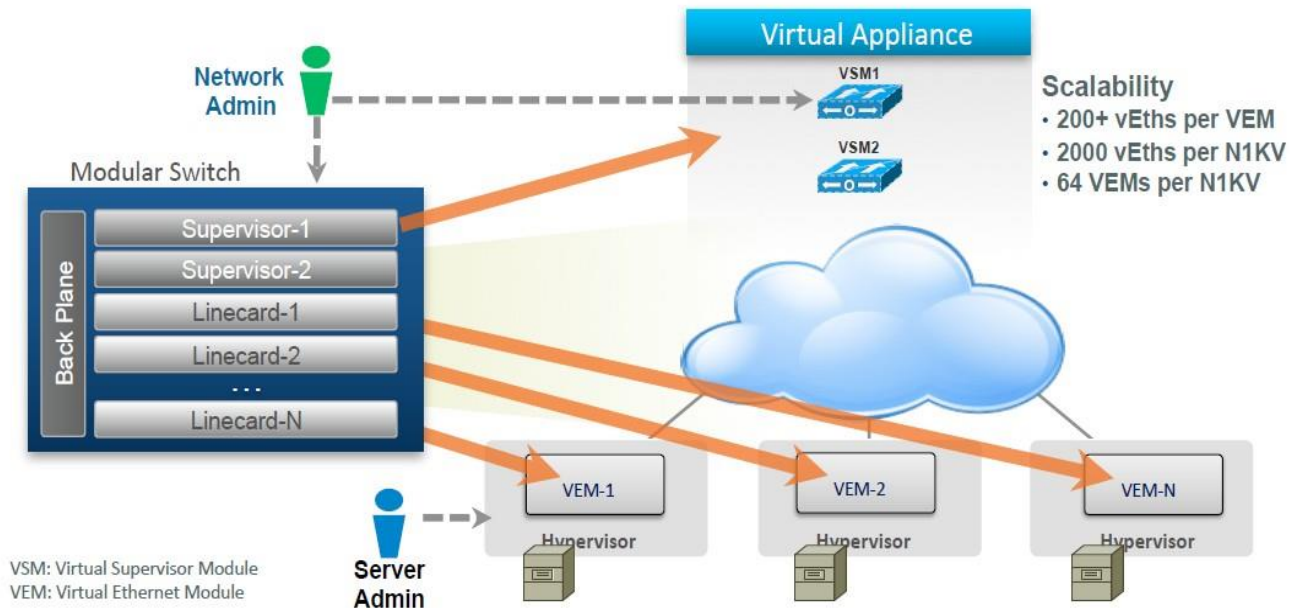
Increased Efficiency, Simpler Operations Scalability, Flexibility, Technology Modularity, Mobility

Nexus 2000 Series

	Nexus 2232TM-E	Nexus 2332TQ	Nexus 2348TQ	Nexus 2232TM-E	Nexus 2232PP	Nexus 2348PQ
Form Factor (RU)	1	1	1	1	1	1
Network Interfaces	8 (SFP+)	4 (QSFP+)	6 (QSFP+)	8 (SFP+)	8 (SFP+)	6 (QSFP+)
Host Interfaces	32	32	48	32	32	48
Performance	560Gbps	1.2Tbps	1.2Tbps	560Gbps	560Gbps	1.2Tbps
FCoE Support	No	No	Yes (FCoE)	No	Yes	Yes
Max. Power (W)	240 W	350 W	350 W	240 W	210 W	350 W
Oversubscription	4:1	2:1	2:1	4:1	4:1	2:1
Parent Switches	Nexus 5000, 6000, 7000,9000	Nexus 5000, 6000, 7000, 9000	Nexus 5000, 6000, 7000, 9000	Nexus 5000, 6000, 7000,9000	Nexus 5000, 6000, 7000,9000,UCSFI	Nexus 5000, 6000, 7000, 9000

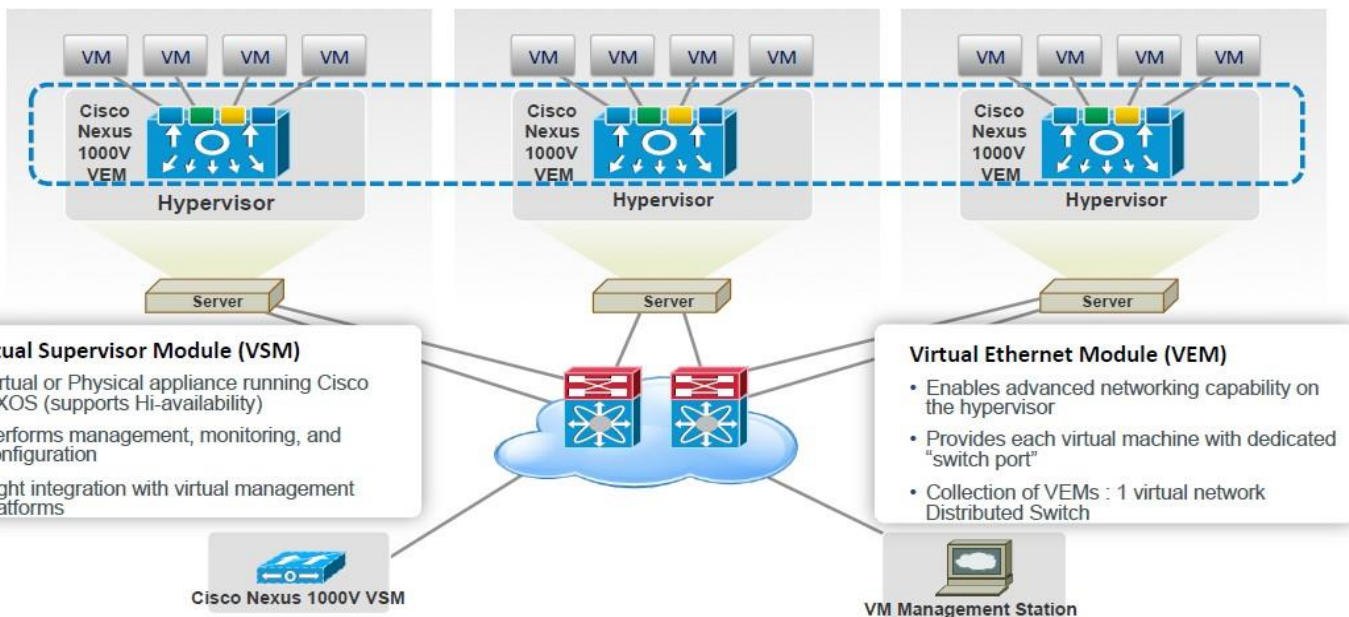
Cisco Nexus 1000V Overview

Architecture consistent with other modular switches



Cisco Nexus 1000V Overview

A Feature-rich Distributed Virtual Software Switch

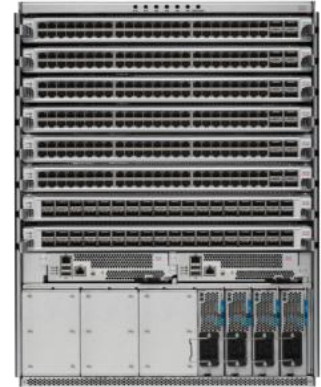


Advanced Switching Capability on the hypervisor

Provides Each VM with Dedicated "Switch Ports"

Nexus 9000 Series

Cisco Nexus® 9000 Series



Designed for Reliability

Designed for Future Scale

Designed for Power and Cooling Efficiency

Proven by theory and experiments in real operational networks

Nexus 9000 Switches

2nd Generation

Nexus 9300-EX



48p 10/25G SFP + 6p 40/100G QSFP

Nexus 93180YC-EX



48p 1/10GT + 6p 40/100G QSFP

Nexus 93108TC-EX

Dual personality – **ACI and NX-OS mode**

Industry's first native 25G VXLAN capable switch

Flexible port configurations – 1/10/25/40/50/100G

Up to 40 MB shared buffer

Native Netflow

Nexus 9200



36p 40/100G QSFP

Nexus 9236C



56p 40G + 8p 40/100G QSFP

Nexus 92304QC



72p 40G QSFP

Nexus 9272Q



48p 10/25G SFP + 4p 100G/6p 40G QSFP

Nexus 92160YC-X

NX-OS switches

Industry's first 36p 100G 1RU switch

Industry's first native 25G VXLAN capable switch

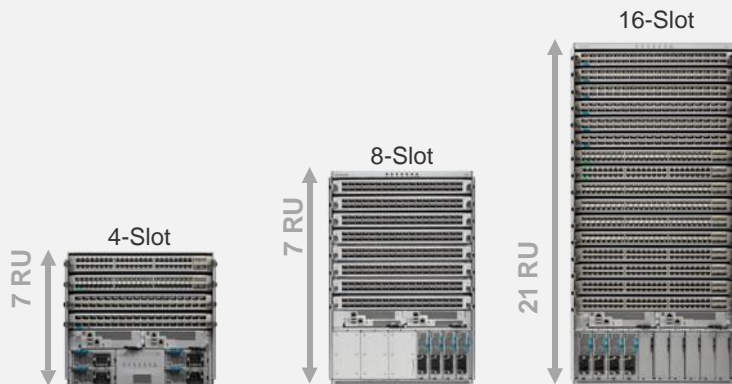
Up to 30 MB shared buffer

High density compact 40/100G aggregation

Cisco Nexus 9500 Platform Switches

Density in DC Optimised Footprint

Cisco Nexus® 9500



	Nexus 9504	Nexus 9508	Nexus 9516
Payload Slots	4	8	16
Cloud Scale	Shipping	Shipping	Mid CY17
BRCM TH	Shipping	Shipping	No Plans
BRCM T2	Shipping	Shipping	Shipping
BRCM Jericho	Q2CY17	Shipping	Future

Common Components

Chassis, Supervisor, System
Controller, Power Supply, Fan Tray

Deployment Options

Choice of ACI and NX-OS,
Choice of BRCM and Cisco ASIC

Multi-Generation

Investment Protection:

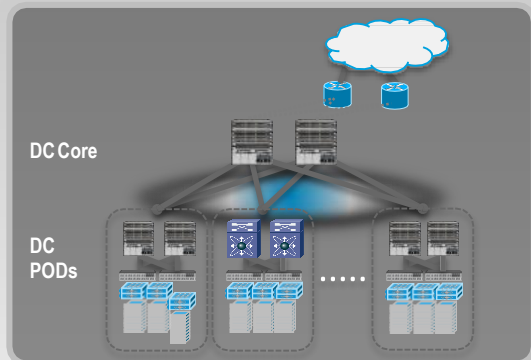
No Mid-plane, Power Supply Headroom
for 100/400G and Line rate encryption

Hardware/Software DC Topology Capabilities

	STP	VPC	FEX	FabricPath	VXLAN	ACI
Nexus 9x00	✓	✓	✓	✗	✓	✓
Nexus 7x00	✓	✓	✓	✓	✓	✗
Nexus 5x00	✓	✓	✓	✓	✓	✗
Nexus 3x00	✓	✓	✗	✗	✓	✗
Nexus 2x00	✓	✓	✓	✓	✓	✓
Nexus 1000v	✓	✓	✓	✓	✓	✓

Catch the Transition & Innovate

EXISTING 2/3-TIER DESIGNS



Existing 2-Tier & 3-Tier Designs

VPC

FEX

Nexus 3000, 5x00, 6000

Nexus 3x00, 5600, 7x00

Nexus 9000

PROGRAMMABLE SDN OVERLAY MODEL

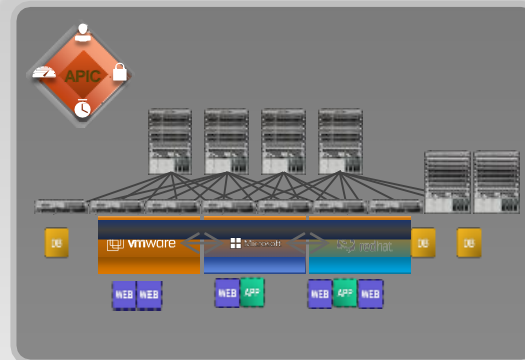


VXLAN Bridging & Routing

Integrated Network Virtualization

SDN Controllers

APPLICATION PROFILES & POLICIES



Application Centric Infrastructure

Policy Model

Automation

NX-OS Features



Ajanda

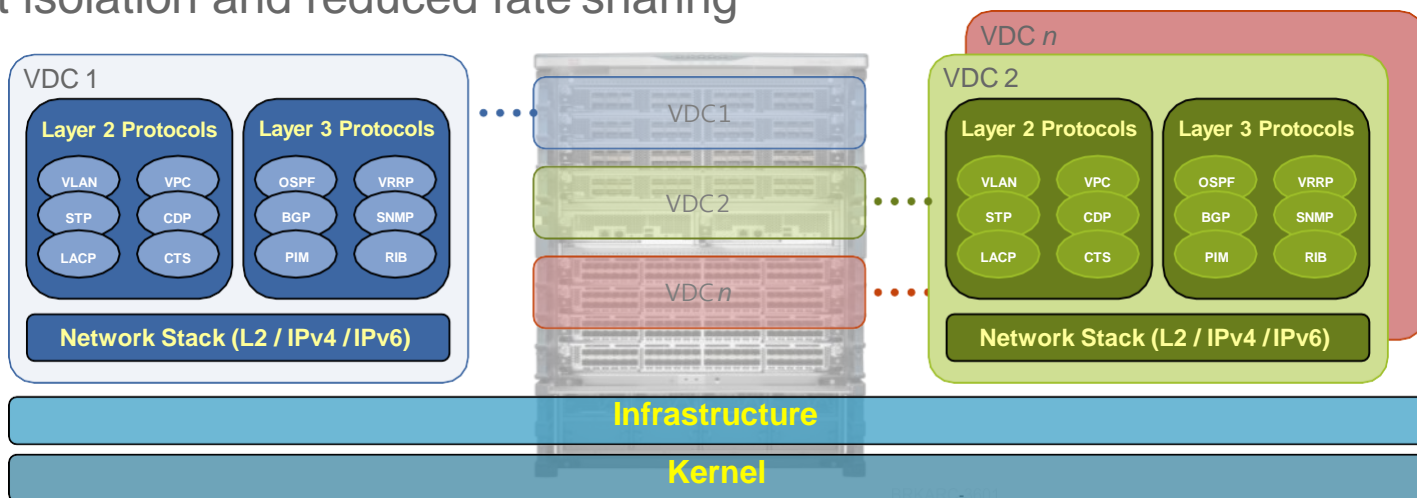
- Virtual Device Context (VDC)
- Virtual Port Channel (vPC)
- FabricPath (vPC+)
- Overlay Transport Virtualization
- ITD

Virtual Device Contexts (VDC)

Virtual Device Contexts

- Create multiple logical devices out of one physical device
- Provide data-plane, control-plane, and management-plane separation
- Fault isolation and reduced fate sharing

Note: VDCs do *not* provide a hypervisor capability, or ability to run different OS versions in each VDC



Virtual Port Channel (VPC)

How does vPC help with STP?

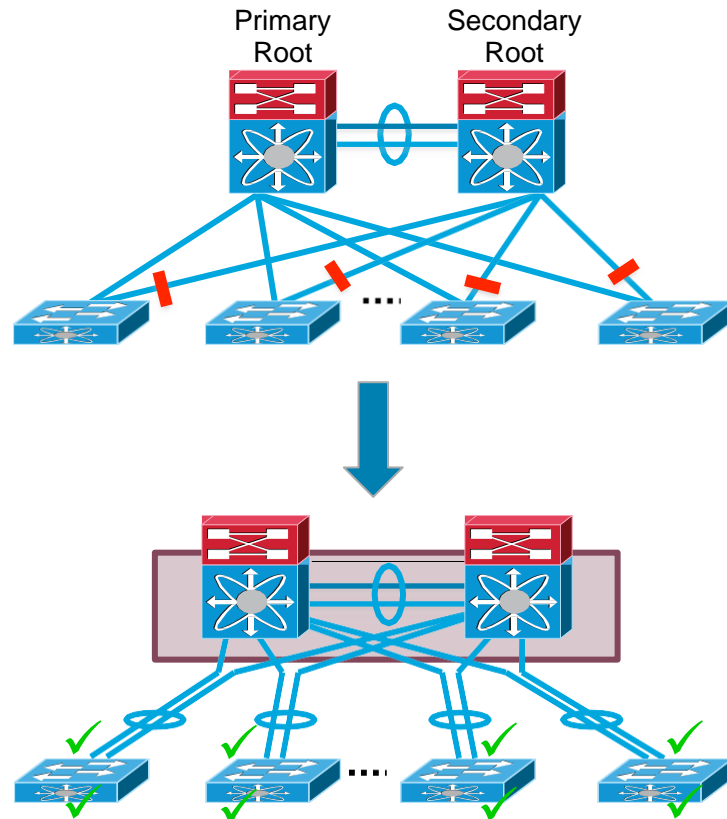
- **Before vPC**

- STP blocks redundant uplinks
- VLAN based load balancing
- Loop Resolution relies on STP
- Protocol Failure

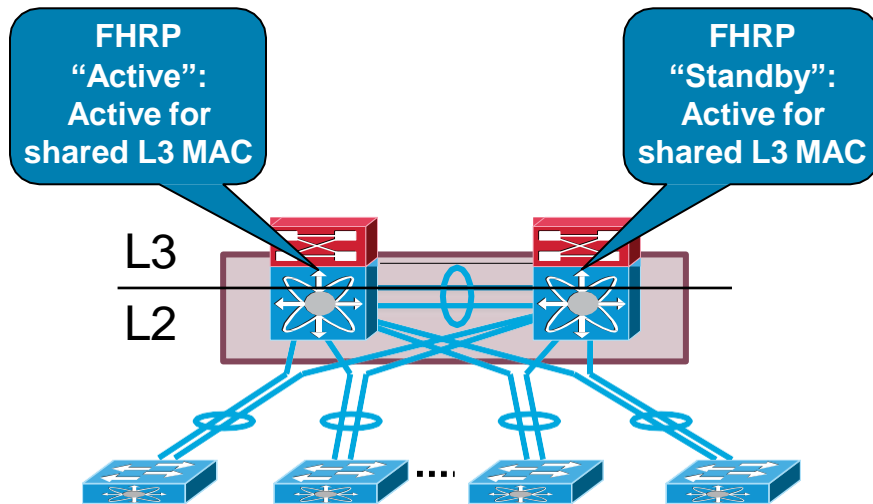


- **With vPC**

- No blocked uplinks
- EtherChannel load balancing (hash)
- Loop Free Topology
- Lower oversubscription



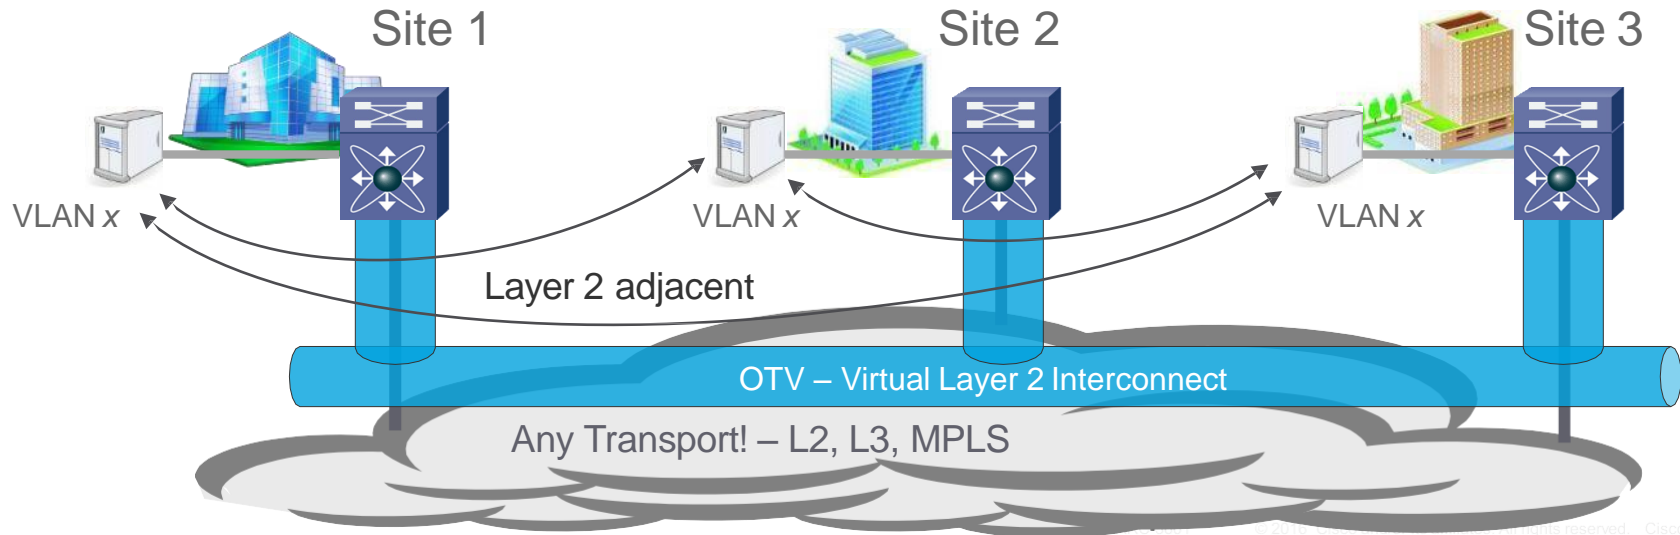
FabricPath



- With L2
 - 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
- With L3
 - Hosts see a single default gateway
 - An arbitrary number of routed interfaces can be created at the edge or within the fabric
 - The hardware is capable of handling million of routes

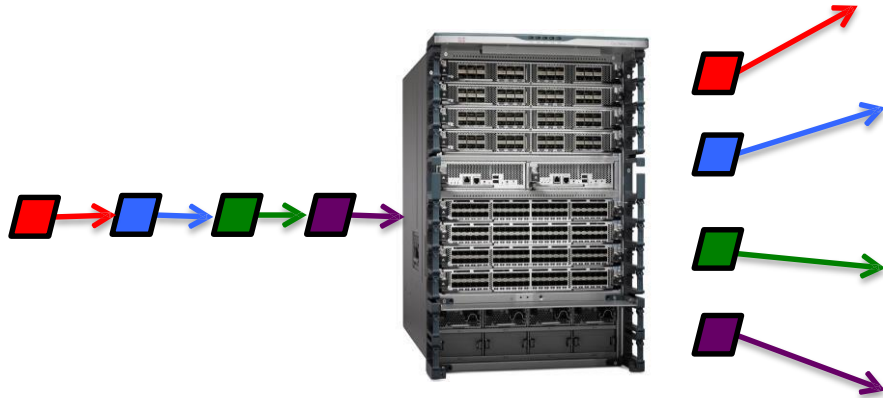
Overlay Transport Virtualization (OTV)

- Overlay Transport Virtualisation (OTV) provides multi-site Layer 2 Data Centre Interconnect (DCI)
- Dynamic “MAC in IP” encapsulation with forwarding based on MAC “routing” table
- No pseudo-wire or tunnel state maintained



Intelligent Traffic Director (ITD)

Traffic distribution through packet redirection

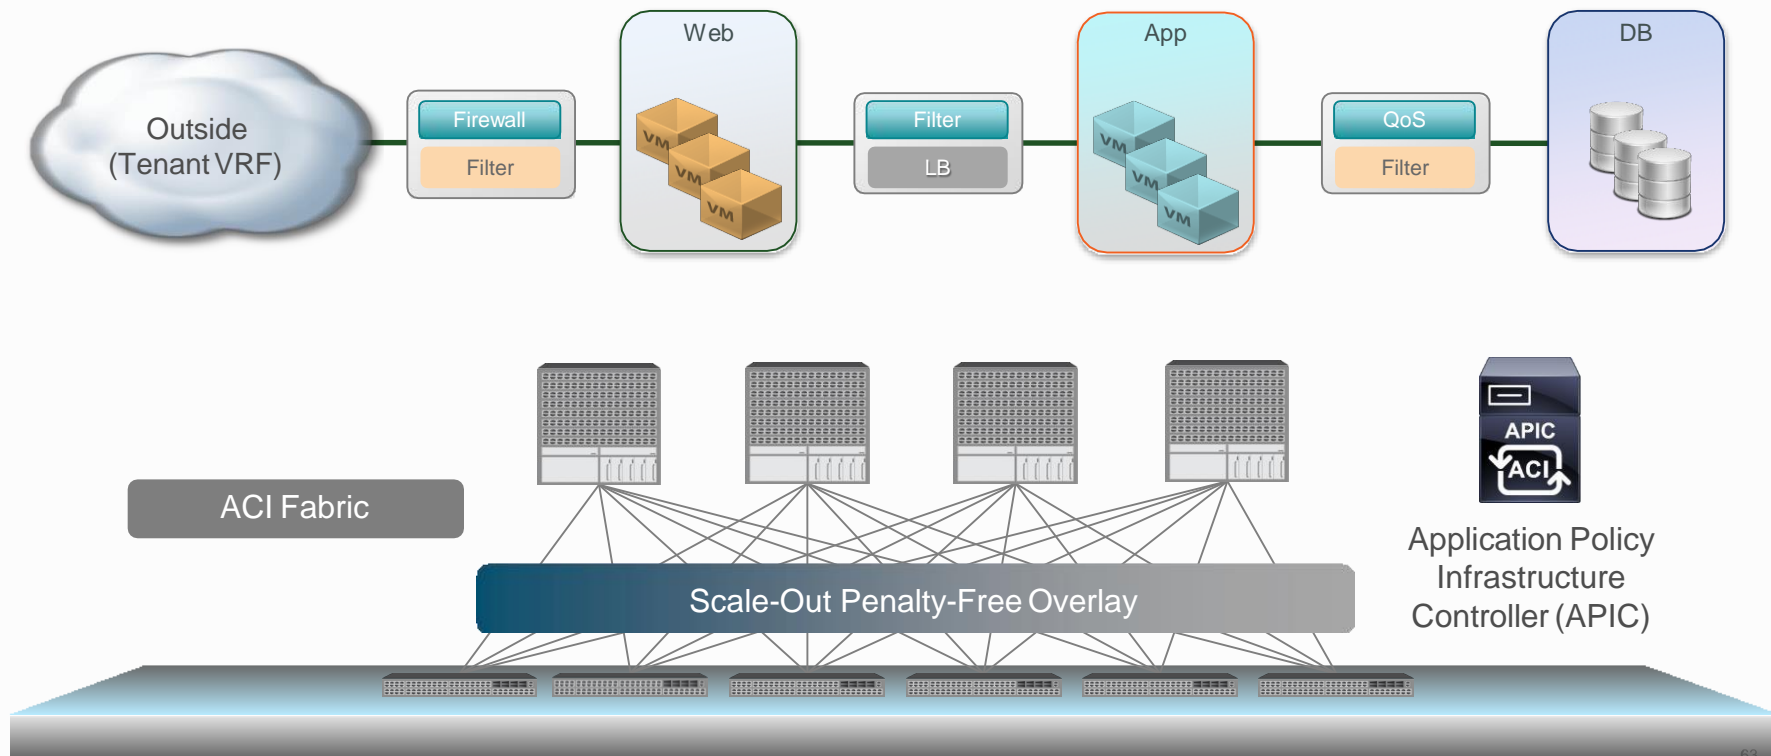


- Traffic distribution and redirection
- ASIC based solution(HW-switched)
- Caters to multi-terabit traffic
- Works on Nexus switches – 9/7/6/5k

Application Centric Infrastructure

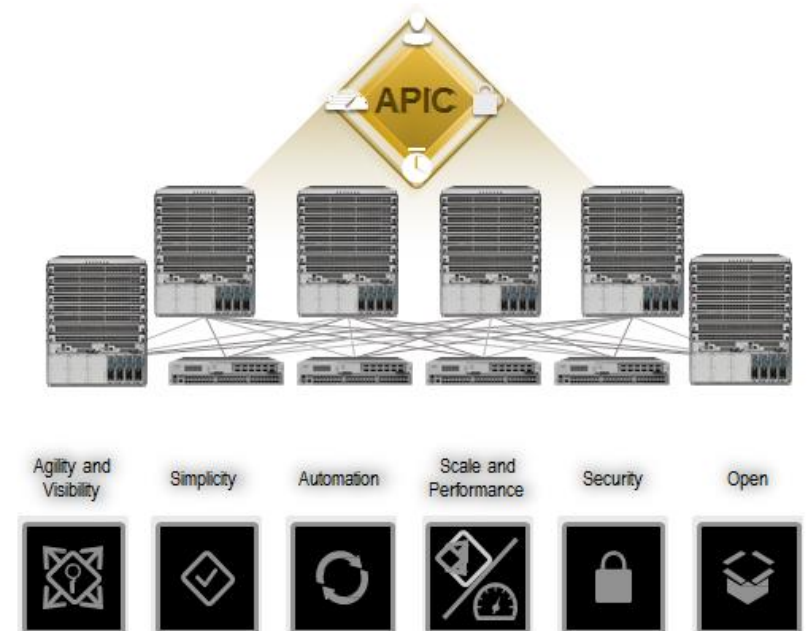


Application Policy Infrastructure



ACI FABRIC

- Agile, open, and highly secure architecture
- Application-based policy model
- Reducing errors and accelerating application deployment
- Provides transparent support of heterogeneous physical and virtual endpoints
- Increased visibility of the entire infrastructure
- Supports open APIs, open source, and open standards to optimize customer choice and flexibility



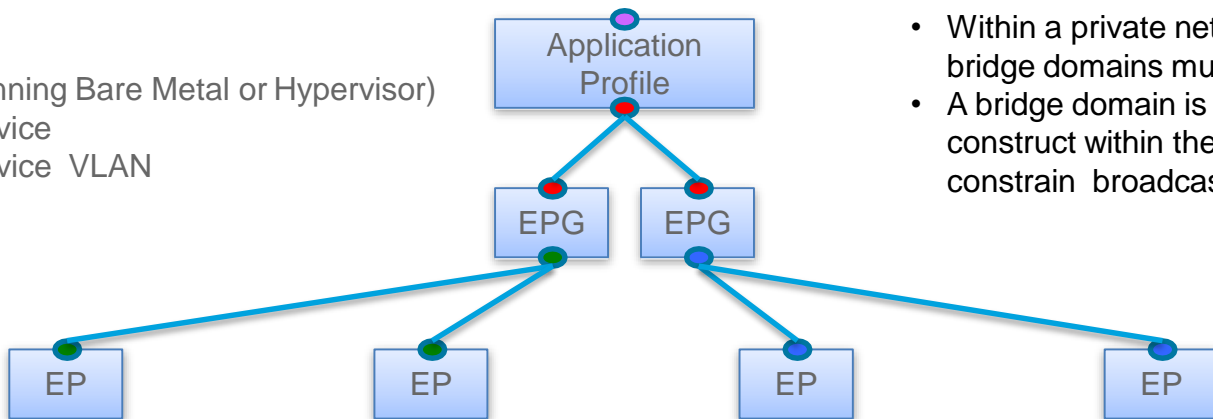
ACI Components

- A **Tenant** is a container for all network, security, troubleshooting and L4 – 7 service policies.
- Tenant resources are isolated from each other, allowing management by different administrators.

Private networks (also called VRFs or contexts) are defined within a tenant to allow isolated and potentially overlapping IP address space.

EPG

- Virtual Machine
- Physical Server (running Bare Metal or Hypervisor)
- External Layer 2 device
- External Layer 3 device VLAN
- Subnet Firewall
- Load balancer



Bridge Domains

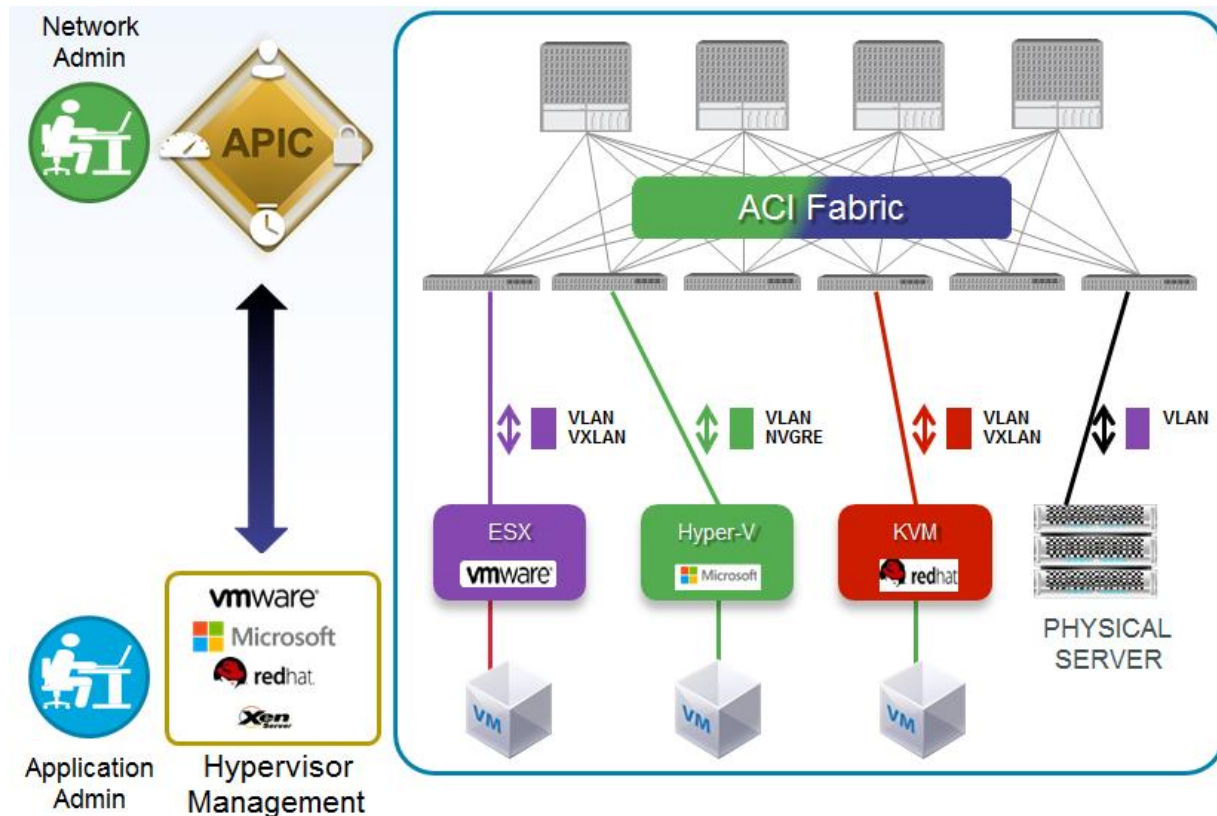
- Within a private network, one or more bridge domains must be defined.
- A bridge domain is a L2 forwarding construct within the fabric, used to constrain broadcast and multicast traffic.

Virtual Port, Physical Ports, External L2 VLAN, External L3 subnet

Multi-Hypervisor-Ready Fabric

Virtual Integration

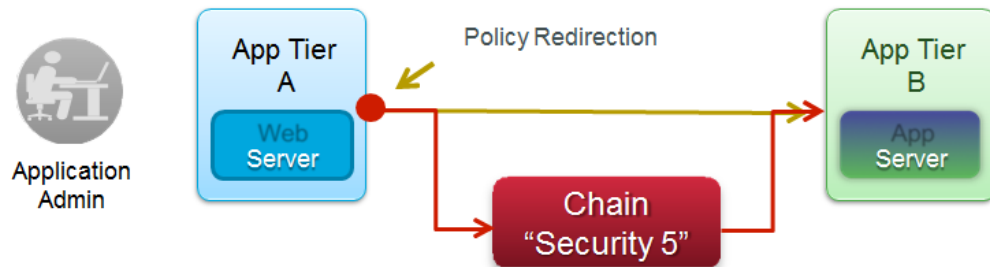
- Integrated gateway for VLAN, VxLAN, NVGRE networks from virtual to physical
- Normalization for NVGRE, VXLAN, and VLAN networks
- Customer not restricted by a choice of hypervisor
- Fabric is ready for multi-hypervisor



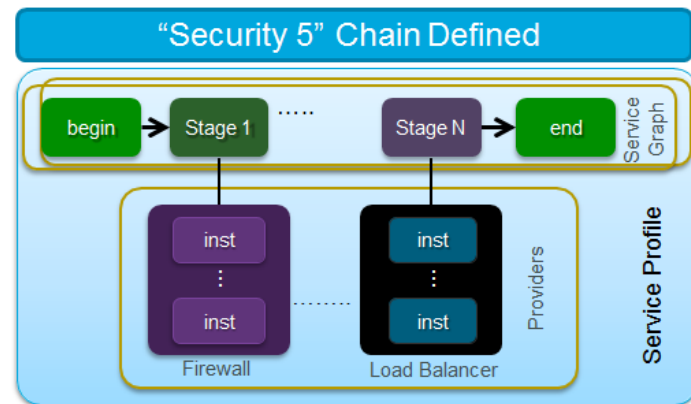
ACI Layer 4 – 7 Service Integration

Centralized, Automated, and Supports Existing Model

- Elastic service insertion architecture for physical and virtual services
- Helps enable administrative separation between application tier policy and service definition
- APIC as central point of network control with policy coordination
- Automation of service bring-up / tear-down through programmable interface
- Supports existing operational model when integrated with existing services
- Service enforcement guaranteed, regardless of endpoint location



Service Admin



Teşekkür Ederiz.

