

# What You Make Possible



# End-to-End Data Centre Virtualisation

BRKVIR-2931

# Setting the Stage

What's the Meaning of Virtual?

- If you can see it and it is there
  - It's real
- If you can't see it but it is there
  - It's transparent
- If you can see it and it is not there
  - It's virtual
- If you can not see it and it is not there
  - It's gone



# Today's Data Centre Virtualisation "Journey" ...



# Today's Data Centre Virtualisation "Journey" ...



# Our Journey in a Nutshell




Applications team



Servers team



Networking team




Virtualisation Team



Storage team



Facilities team



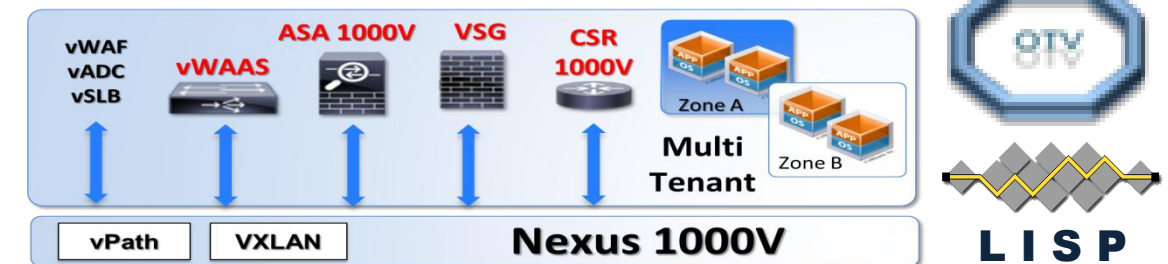
## Front-End Virtualisation & Spine Layer



## Leaf and Access Layer – Physical and Virtual



## Virtual Networking & Cloud Network Services



## Virtual Compute & IO Virtualisation



## Virtualised Storage and SAN



## Software Defined Network (SDN) & Orchestration





# Ultimate Goal: Remove the Virtualisation ...



... so, Please ...




Potential  
Collateral  
Effect





# Agenda

- **Data Centre Virtualisation Overview**
- **Front-End Data Centre Virtualisation**
  - DC Edge Layer
  - Aggregation / Spine Layer
  - Access / Leaf Layer
- **Virtual Networking & Cloud Network Services**
  - Virtual Access Layer: Nexus 1000v
  - Cloud Network Services
  - Inter DC Services: OTV & LISP
- **Virtual Compute & IO Virtualisation**
  - Unified Computing System
  - Hypervisors & Virtual IO @ Cisco UCS
- **Virtualised Storage & SAN**
  - Unified Fabric
  - Virtualised SAN and Storage
- **Software Defined Network & Orchestration**
  - SDN & Cisco Open Network Environment (ONE)
  - Cloud Orchestration
- **Implementation Examples**
- **Q&A**



## Front-End Virtualisation & Spine Layer




## Leaf and Access Layer – Physical and Virtual



## Virtual Networking & Cloud Network Services



## Virtual Compute & IO Virtualisation




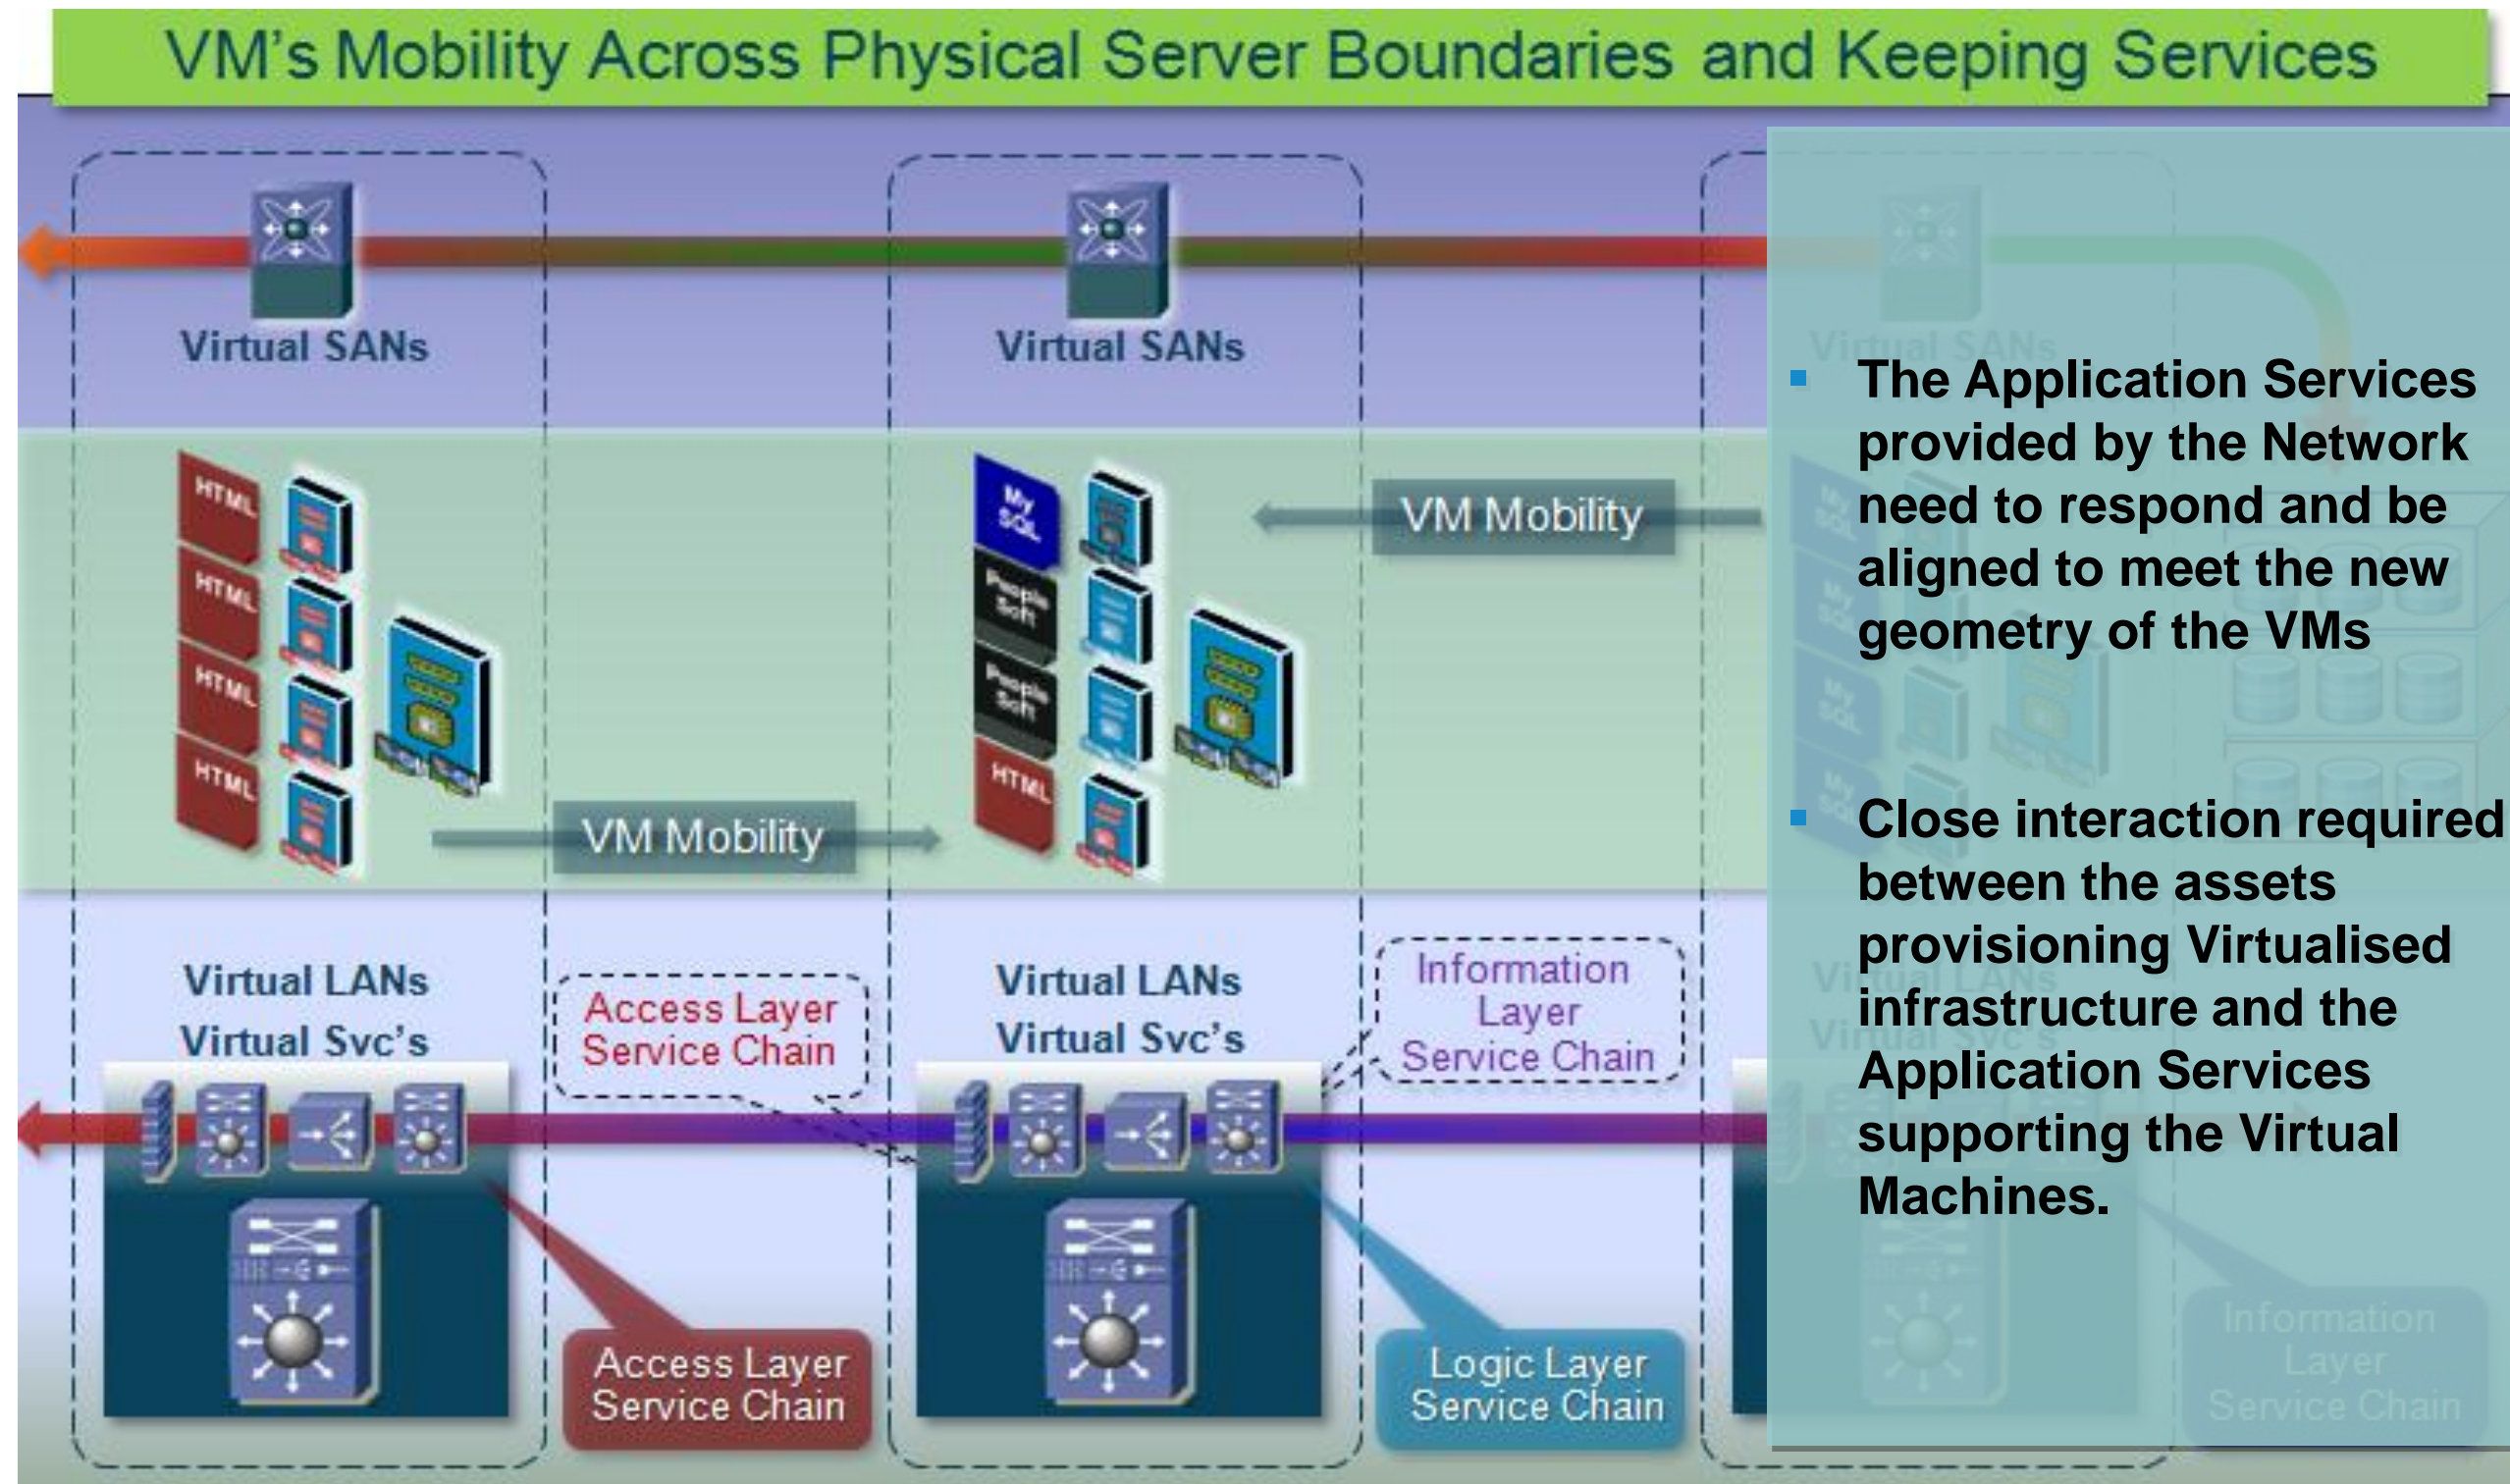
## Virtualised Storage and SAN




## Software Defined Network (SDN) & Orchestration



# The “Virtual Data Centre” Approach

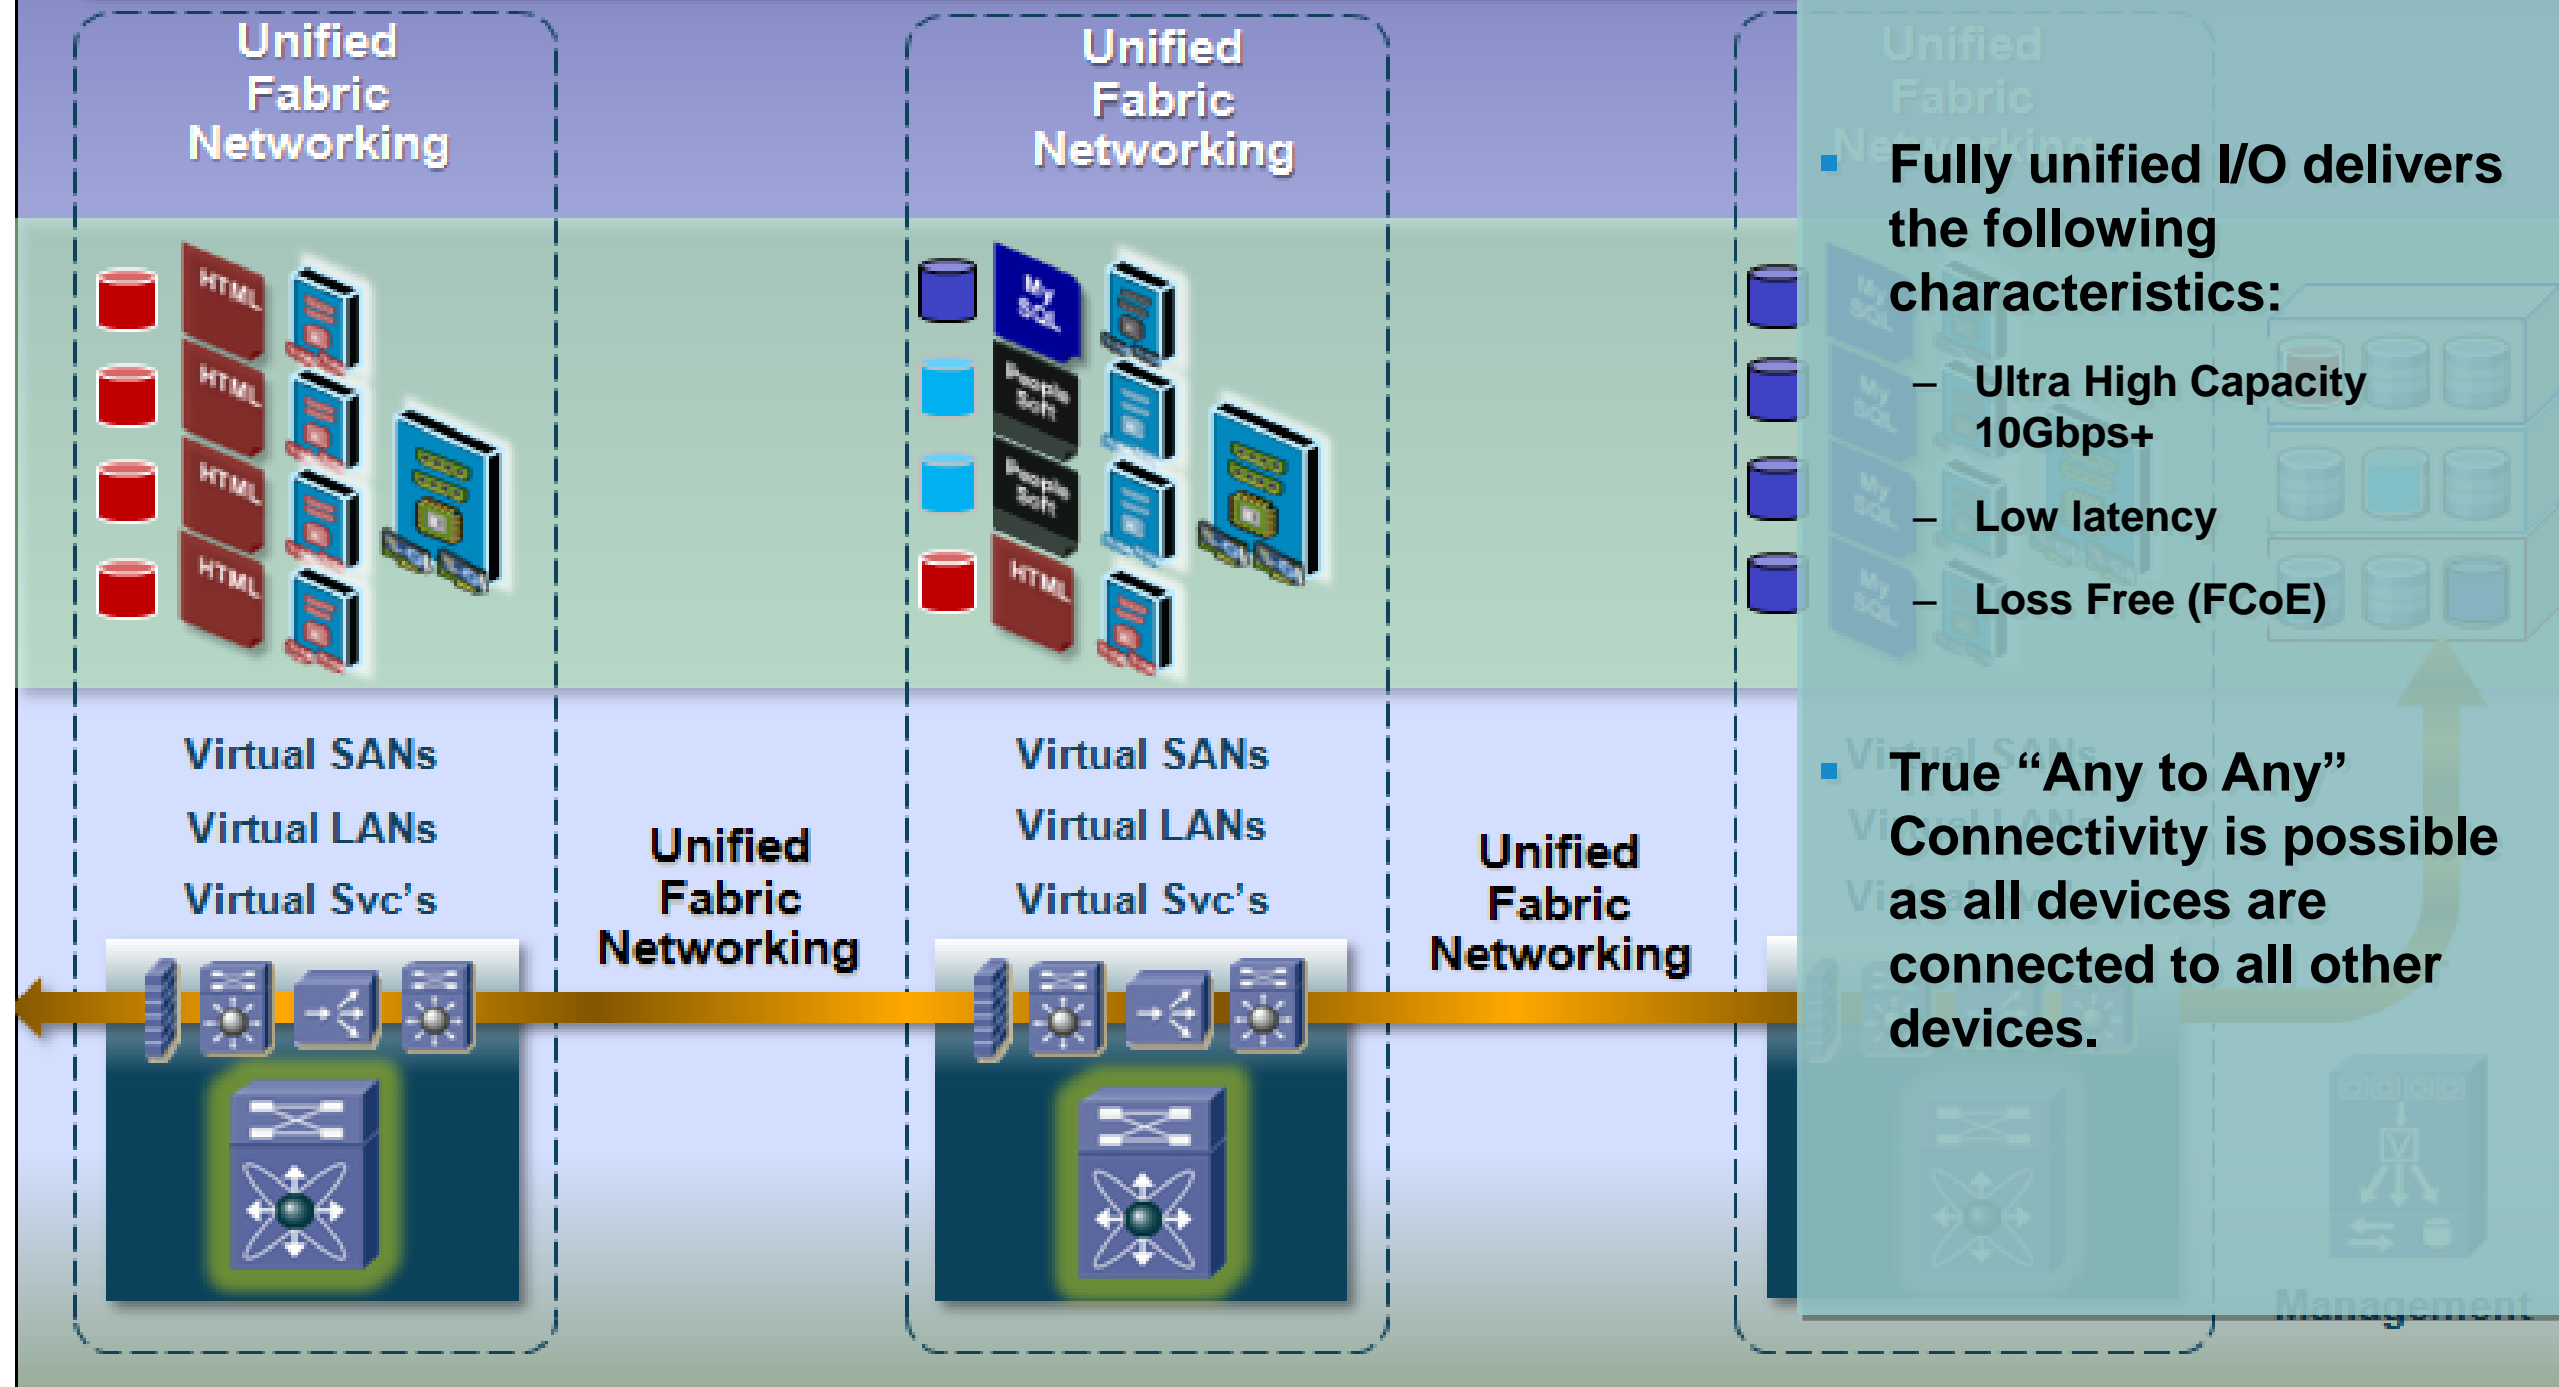


H  
I  
S  
T  
O  
R  
Y



# Moving to a Unified Fabric

Moving to a fully virtualised Data Centre, with Any to Any Connectivity



# From Multiple Parallel Paths to Virtualised Unified Fabric

1GE  
1GE

1GE  
1GE



10/40 Gigabit/s

10/40 Gigabit/s

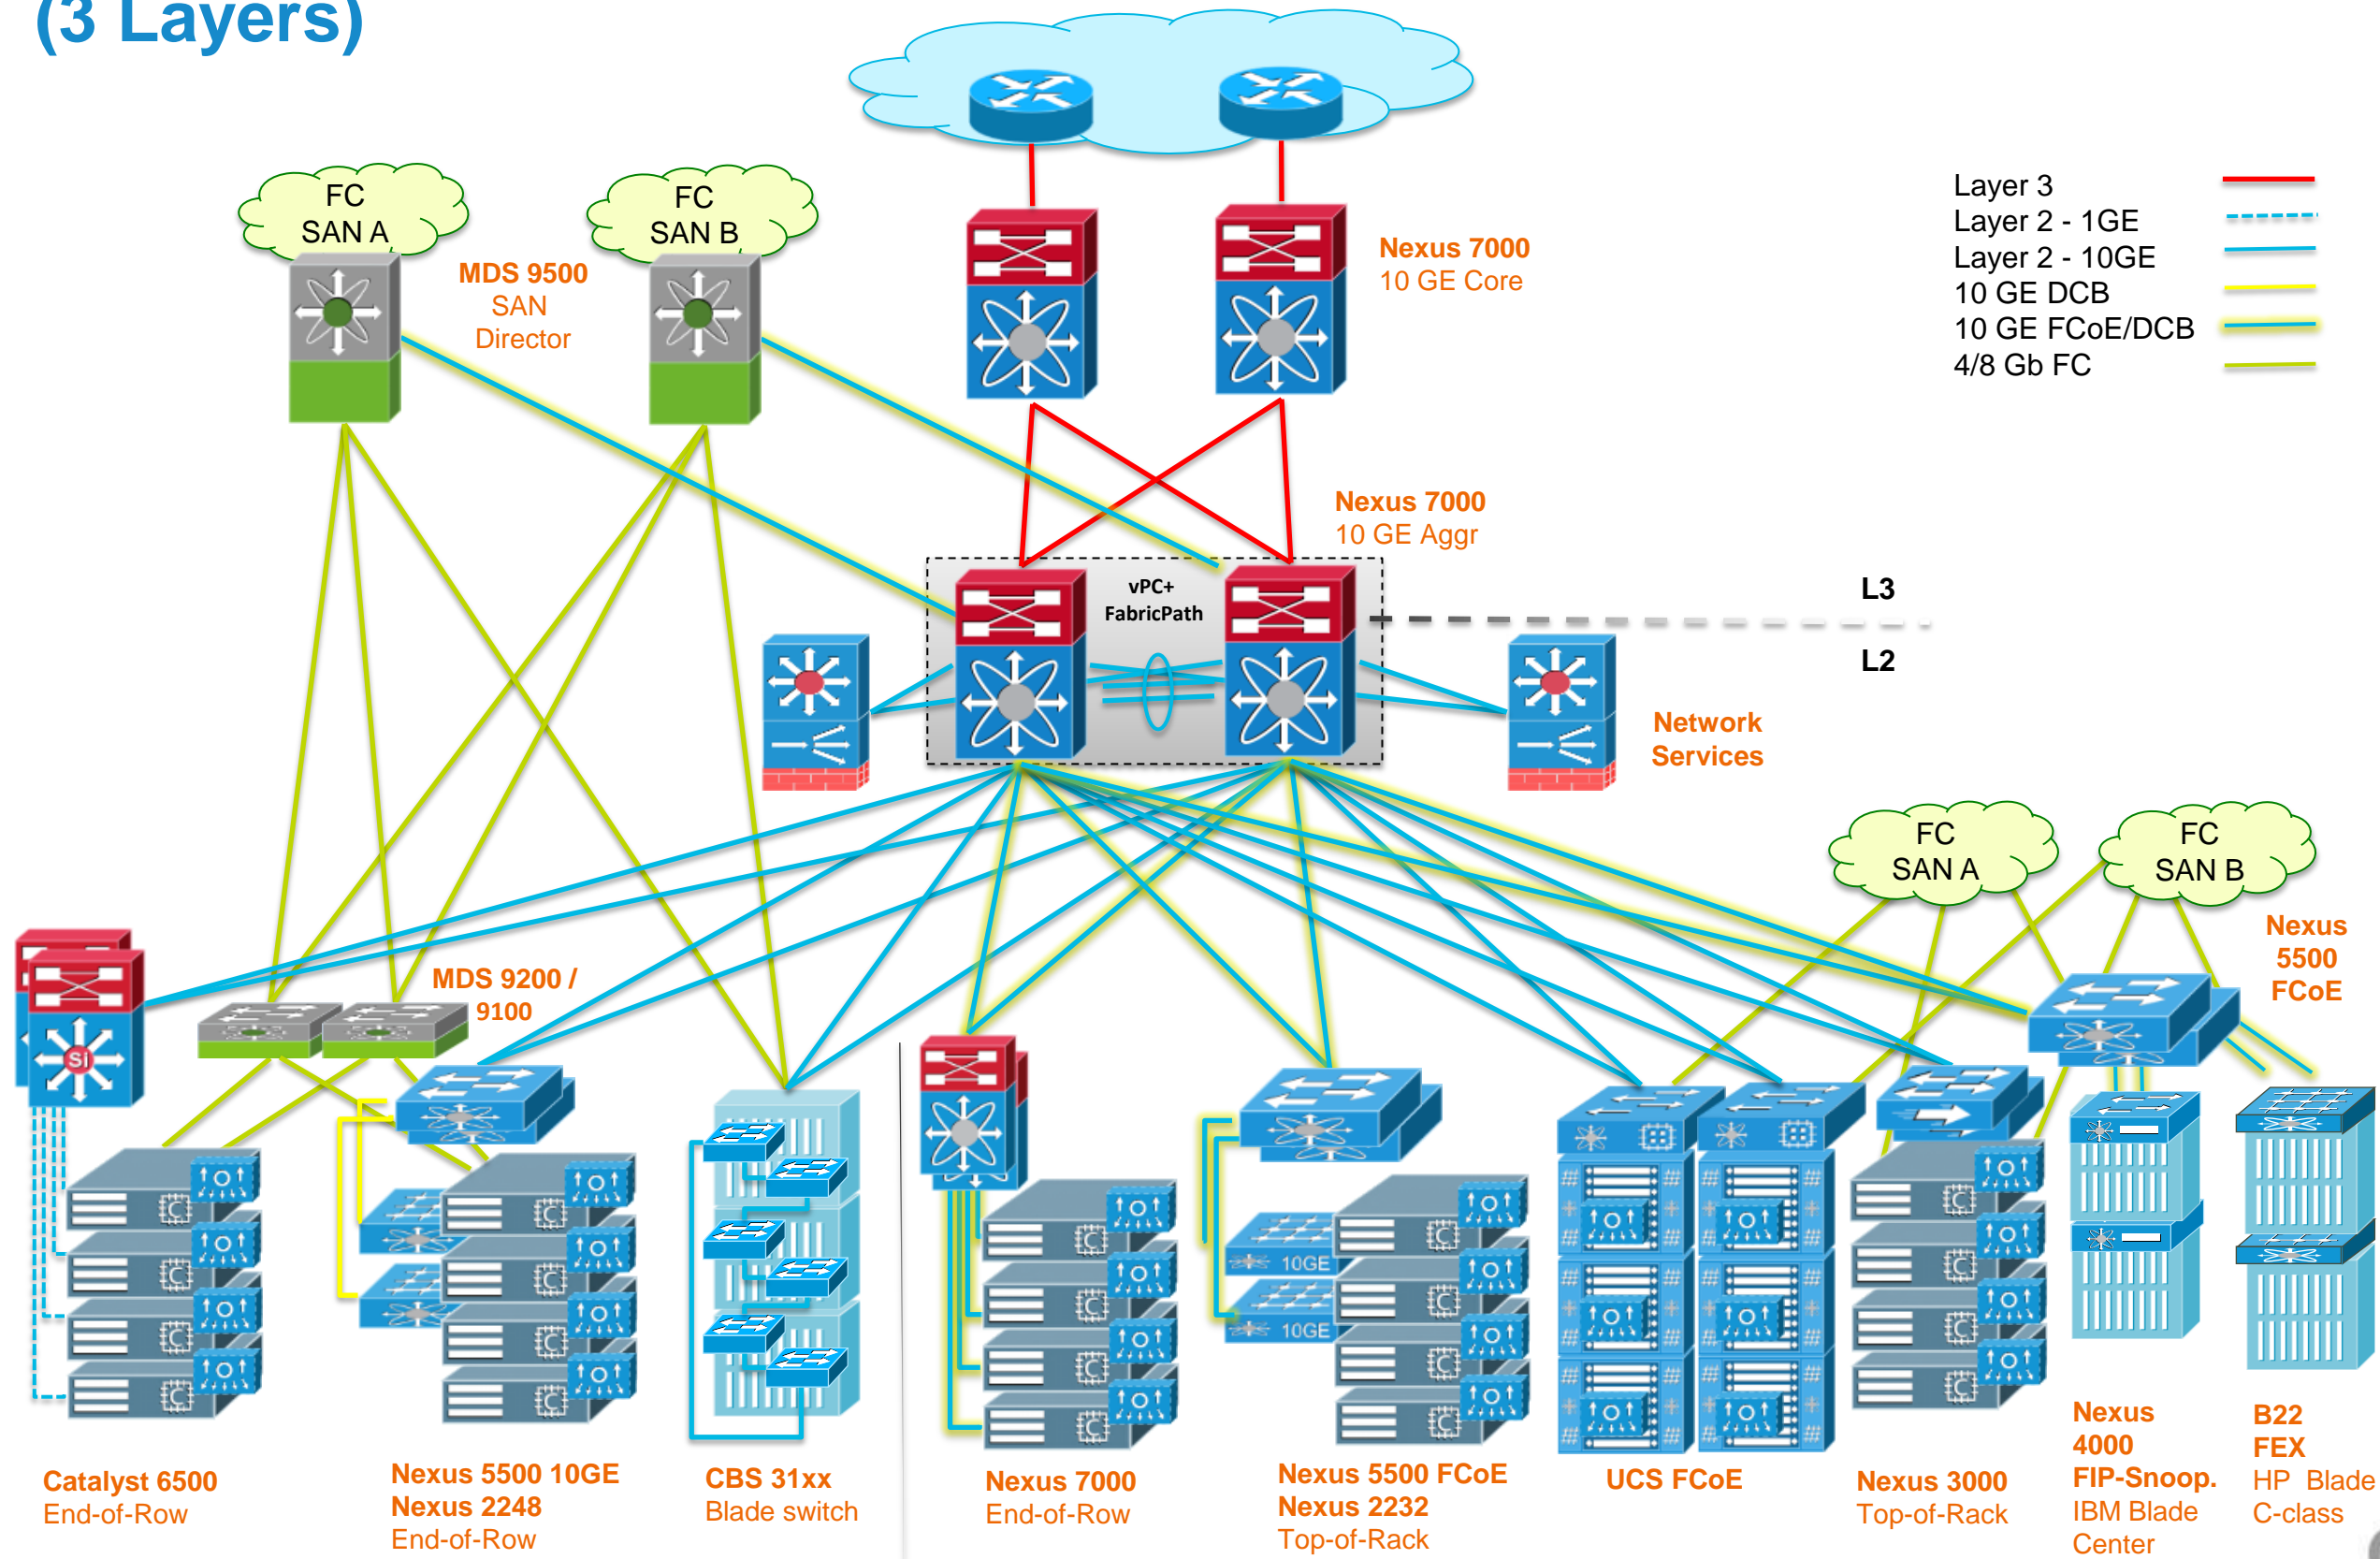
# Typical Virtualised Data Centre Infrastructure (3 Layers)

WAN Edge Layer

DC Edge Layer (LAN & SAN)

Aggregation & Services Layer

Access Layer SAN Edge



- Layer 3
- Layer 2 - 1GE
- Layer 2 - 10GE
- 10 GE DCB
- 10 GE FCoE/DCB
- 4/8 Gb FC

1 GbE Server Access & 4/8Gb FC via dual HBA (SAN A // SAN B)  
BRKVIR-2931

© 2013 Cisco and/or its affiliates. All rights reserved.

10Gb DCB / FCoE Server Access or 10 GbE Server Access & 4/8Gb FC via dual HBA (SAN A // SAN B)  
Cisco Public



# The Evolving Data Centre Architecture

## Challenges for the Classical Network Design

- Hypervisor based server virtualisation and the associated capabilities (vMotion, Live Migration, etc.) are changing multiple aspects of the Data Centre design
- Where is the server now?
  - Where is the access port?
  - Where does the VLAN exist?
  - Any VLAN Anywhere?
  - How large do we need to scale Layer 2?
- What are the capacity planning requirements for flexible workloads?
- Where are the policy boundaries with flexible workload (Security, QoS, WAN acceleration, ...)?
- What about programmability and automation ?


Data Centre Row 1



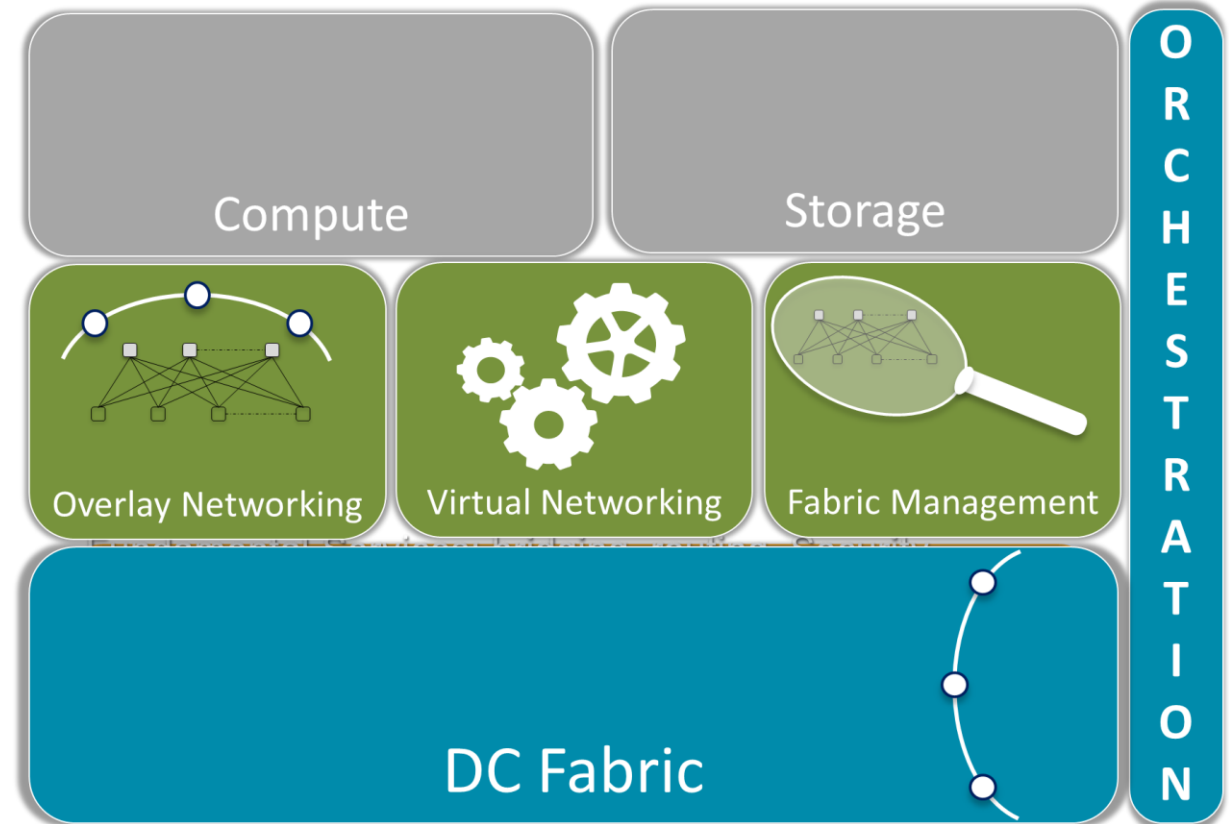
Data Centre Row 2



# Data Centre Infrastructure Strategy & Trends



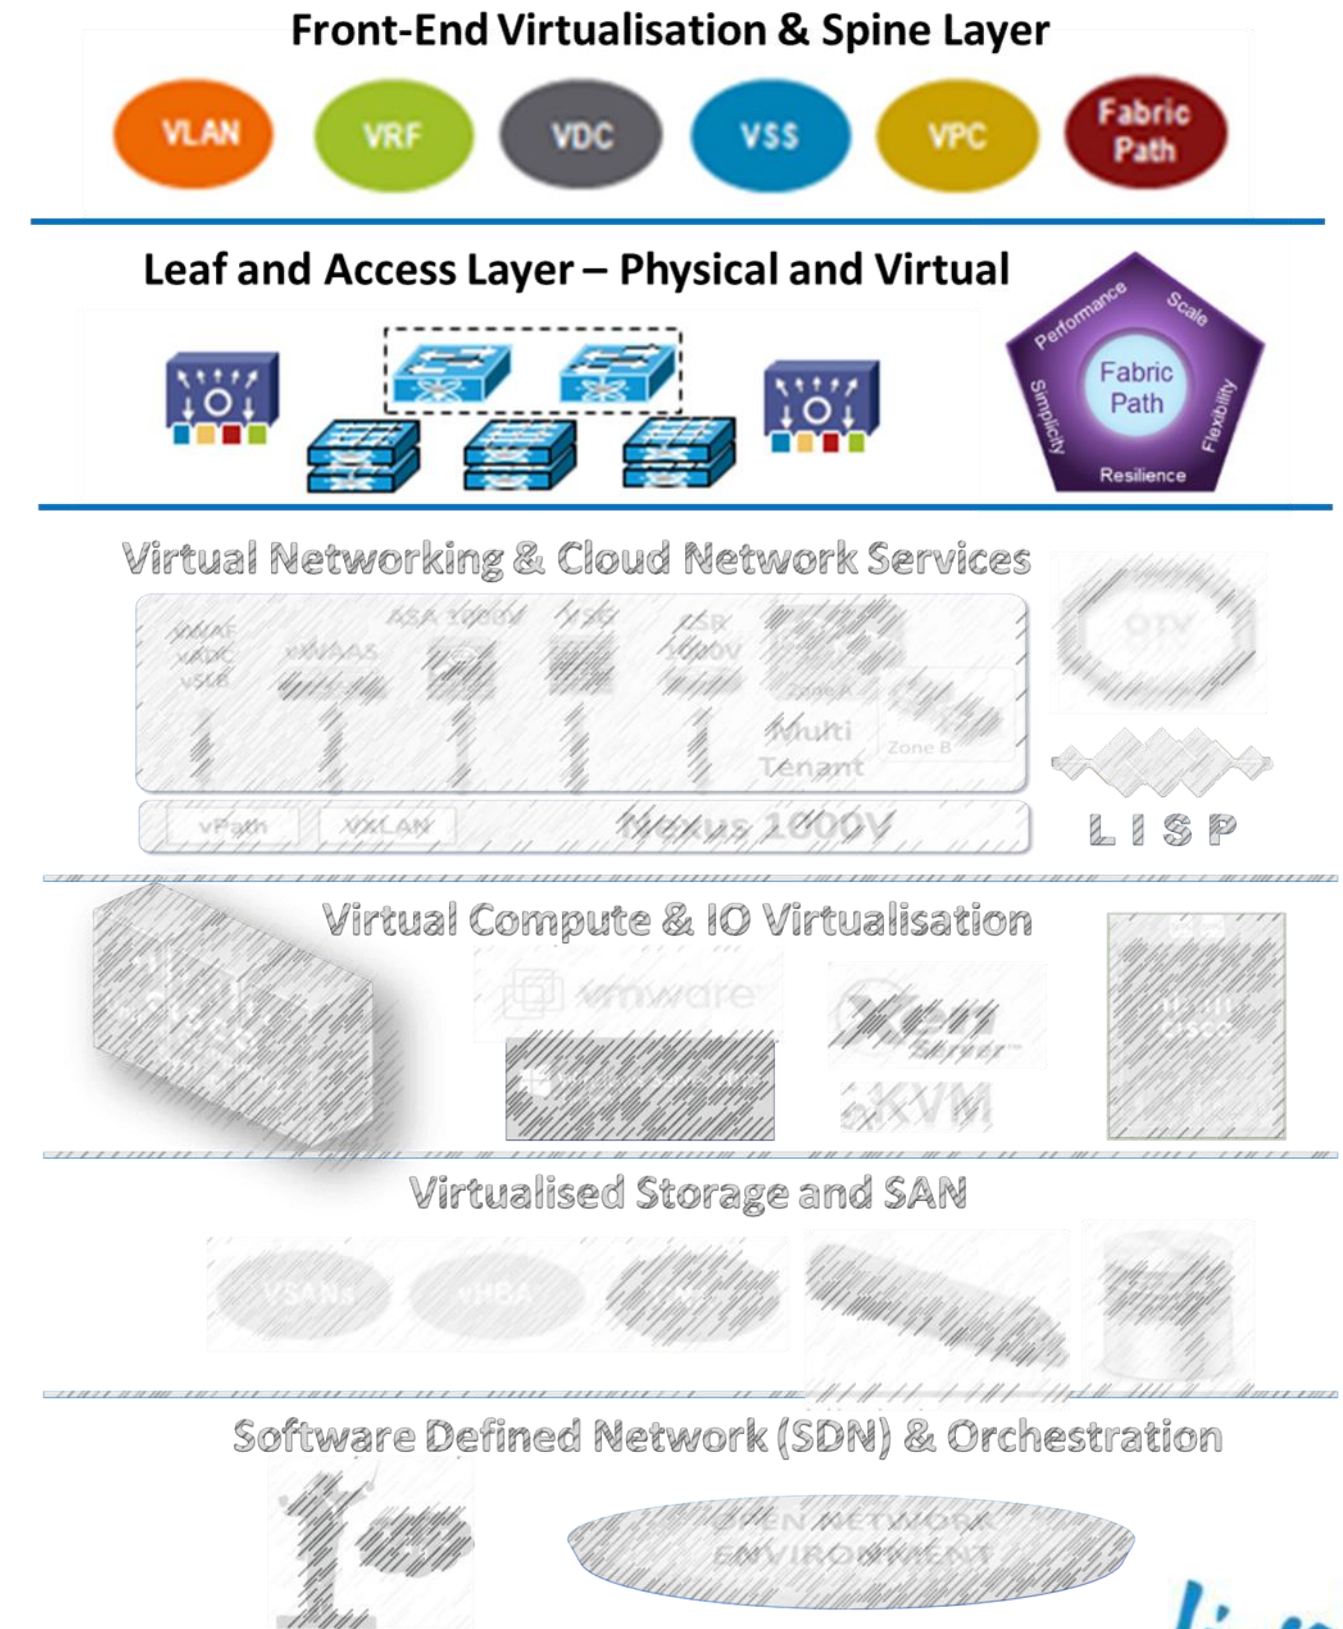
**TRENDS  
2013**



- Virtualisation & Cloud
- Software Defined Network (SDN)
- Virtual Networking / Overlay
- Data Centre Fabric
- Architecture vs Technology

# Agenda

- **Data Centre Virtualisation Overview**
- **Front-End Data Centre Virtualisation**
  - DC Edge Layer
  - Aggregation / Spine Layer
  - Access / Leaf Layer
- **Virtual Networking & Cloud Network Services**
  - Virtual Access Layer: Nexus 1000v
  - Cloud Network Services
  - Inter DC Services: OTV & LISP
- **Virtual Compute & IO Virtualisation**
  - Unified Computing System
  - Hypervisors & Virtual IO @ Cisco UCS
- **Virtualised Storage & SAN**
  - Unified Fabric
  - Virtualised SAN and Storage
- **Software Defined Network & Orchestration**
  - SDN & Cisco Open Network Environment (ONE)
  - Cloud Orchestration
- **Implementation Examples**
- **Q&A**



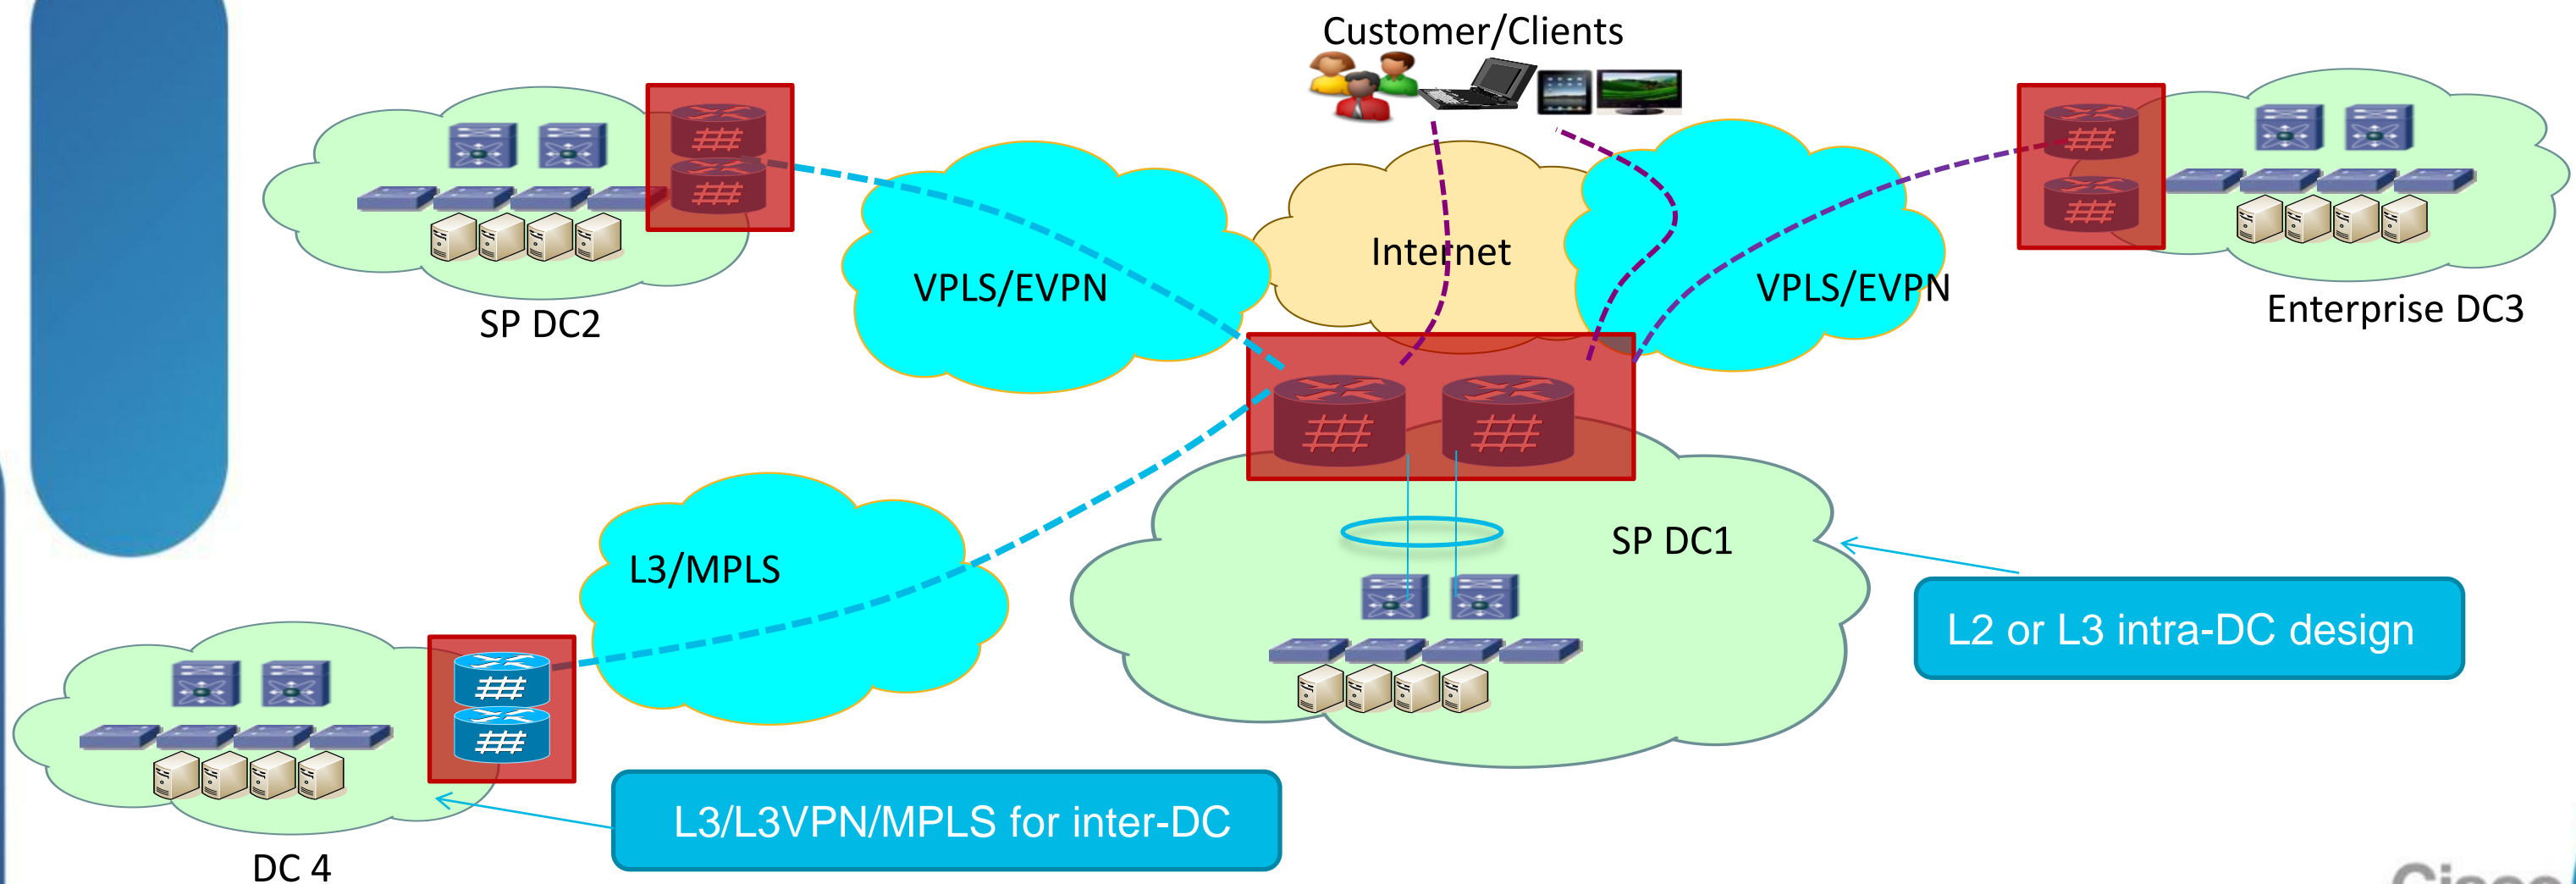


# Next-generation DC Edge – L2 & L3 Foundation

ASR 9000 - Flexible and Scalable:

Layer 2: Any service any port, Any VLAN to any VLAN

Layer 3: Proven BGP/MPLS Feature Set, with High Scale



# ASR 9000 Universal DC Gateway Investment

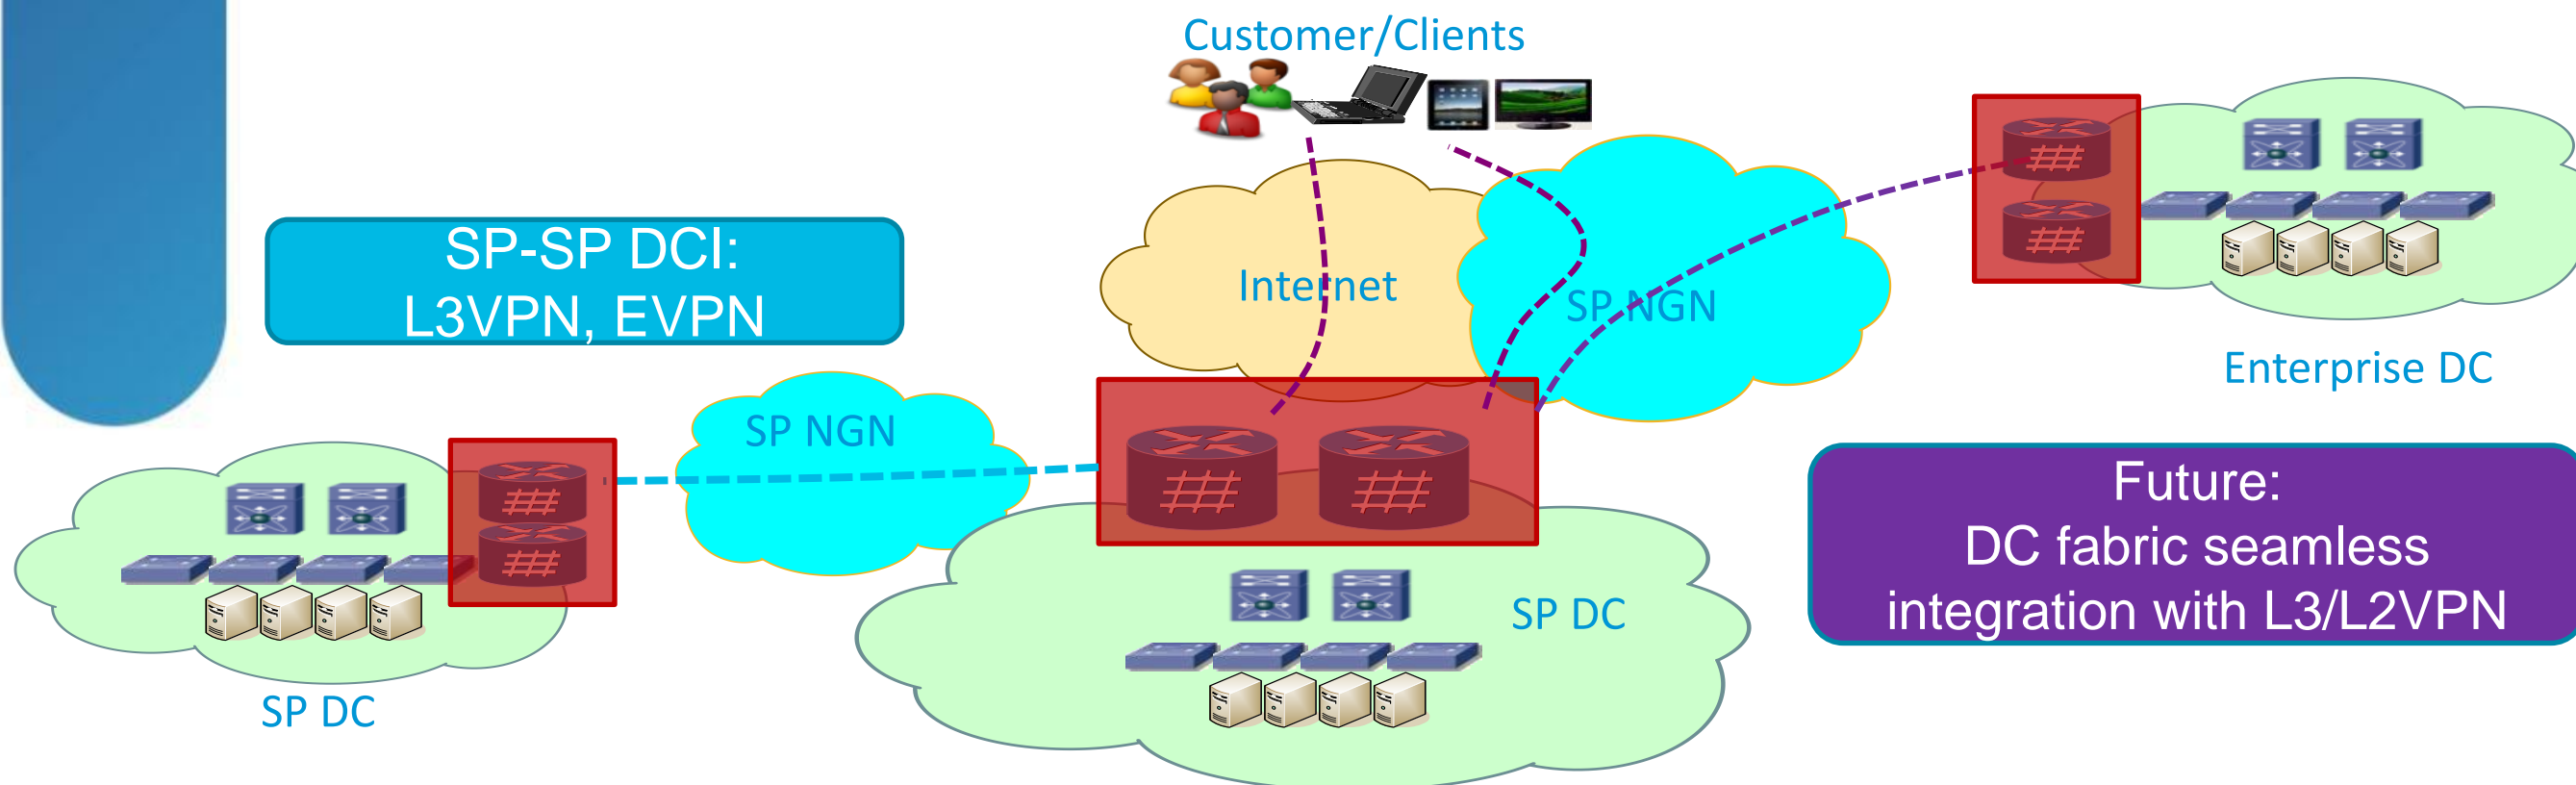
Flexible, Scalable, Integrated and Intelligent

- NPS: Optimised DC resource navigation
- LISP: optimised client-server path, and DCI
- SDN: programmable network, service aware networking

- Internet peering
- Client access internet or VPN

SP-Enterprise DCI:  
L3VPN, EVPN/VPLS


SP-SP DCI:  
L3VPN, EVPN



Future:  
DC fabric seamless  
integration with L3/L2VPN

Cisco *live!*

# Network Planes of Operation





<b>Policy Plane</b>	The business glue of the network. Rules execution, decision making, Service Manager and all the other components to make a productise service.
<b>Services Plane</b>	Overlay “Layer 7” application flow built on the foundation of the other layers. Dependent on the other layers.
<b>Management Plane</b>	The management plane is the logical path of all traffic related to the system management of the platform.
<b>Control Plane</b>	It’s the brain of any networking platform and the technical glue of the network. The control plane is where all routing, switching, other protocols and control information are exchanged
<b>Data Plane</b>	The data plane receives, processes, and transmits network data between network elements, and represents the bulk of network traffic that passes to and through the gear.



# Virtual Port Channel (VPC)

Active — Active Layer 2 Links



## Virtual Port Channel

### Bi-sectional bandwidth with vPC

- ❑ vPC is a Port-channelling concept extending link aggregation to two separate physical switches
- ❑ Allows the creation of resilient L2 topologies based on Link Aggregation.
- ❑ Eliminates the need for STP in the access-distribution Layer
- ❑ Enable seamless VM Mobility, Server HA Clusters
- ❑ Scale Available Layer 2 Bandwidth
- ❑ Dual-homed server operate in active-active mode (via Multi-chassis Etherchannel – MCEC)
- ❑ Simplify Network Design
- ❑ Available on Nexus 7000, Nexus 5000 / 5500 and Nexus 3000

# Cisco FEXlink: Virtualised Access Switch

Nexus 2200 Fabric Extender (FEX)

Cisco Nexus® 5500



+



Cisco Nexus® 2200 FEX

Cisco Nexus® 6000



+



Cisco Nexus® 2200 FEX



Cisco Nexus® 7000



+




Cisco Nexus® 2200 FEX

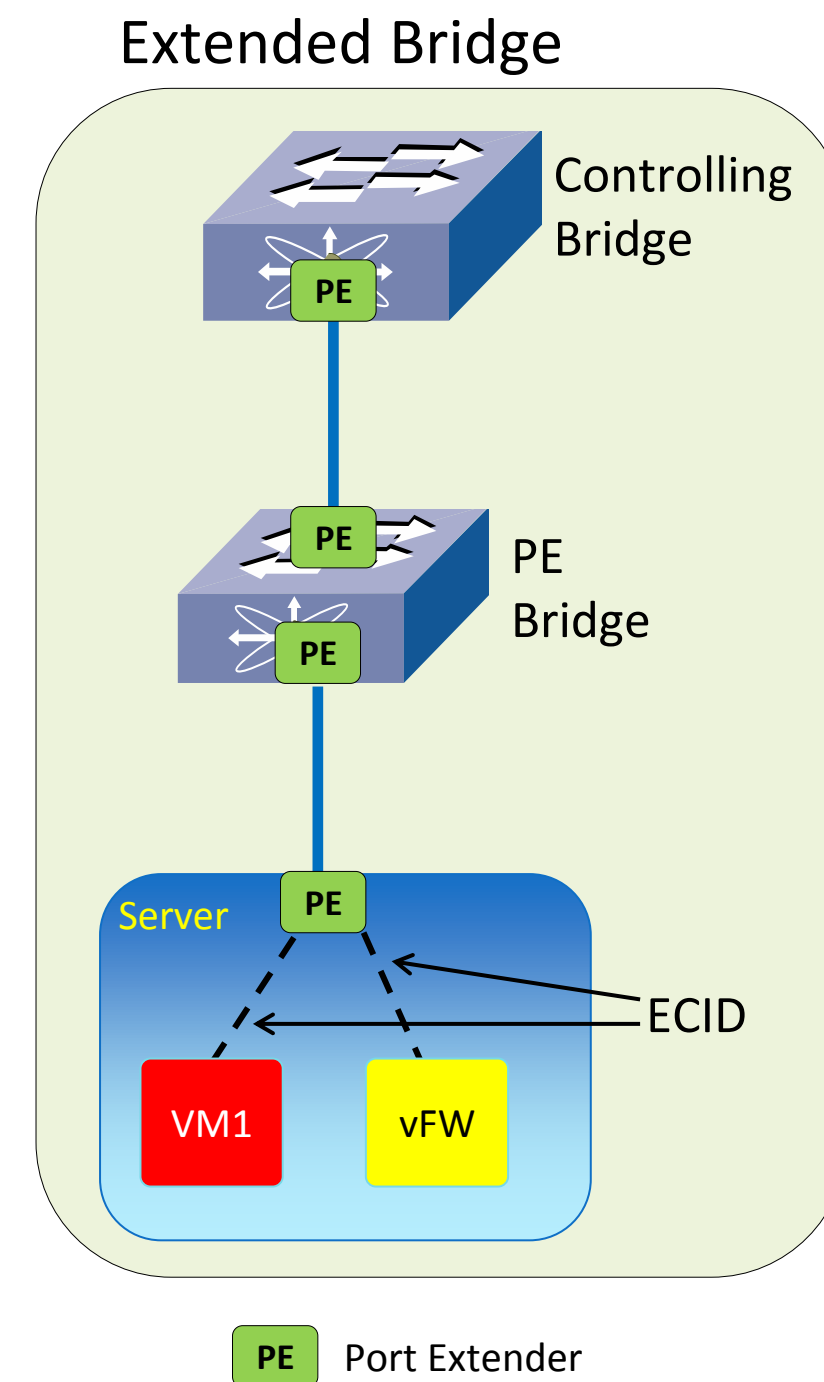


Distributed High  
Density Edge  
Switching System

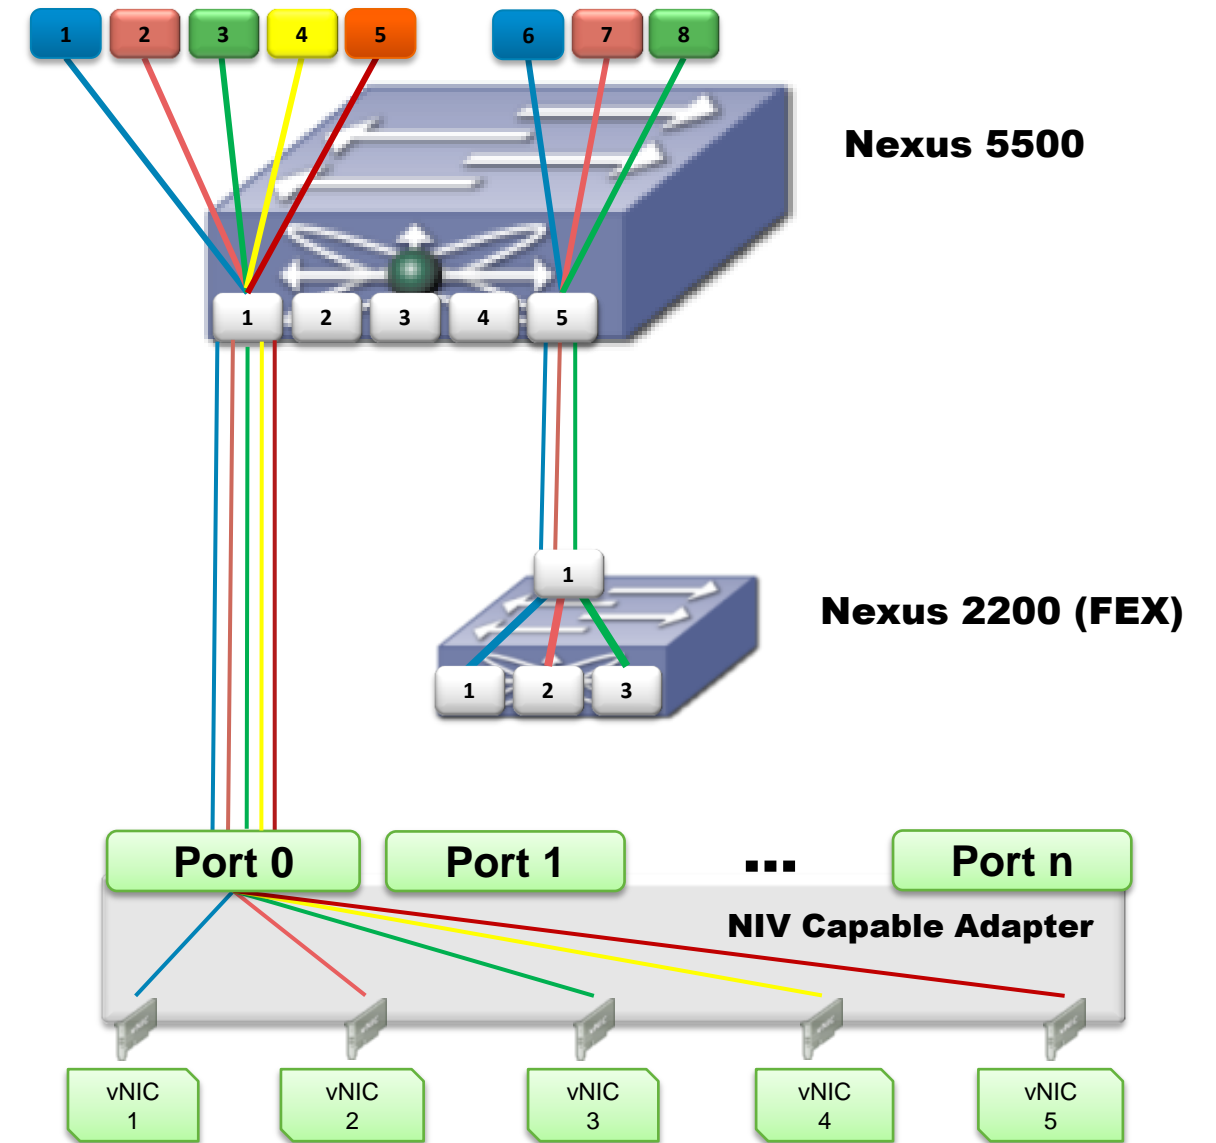
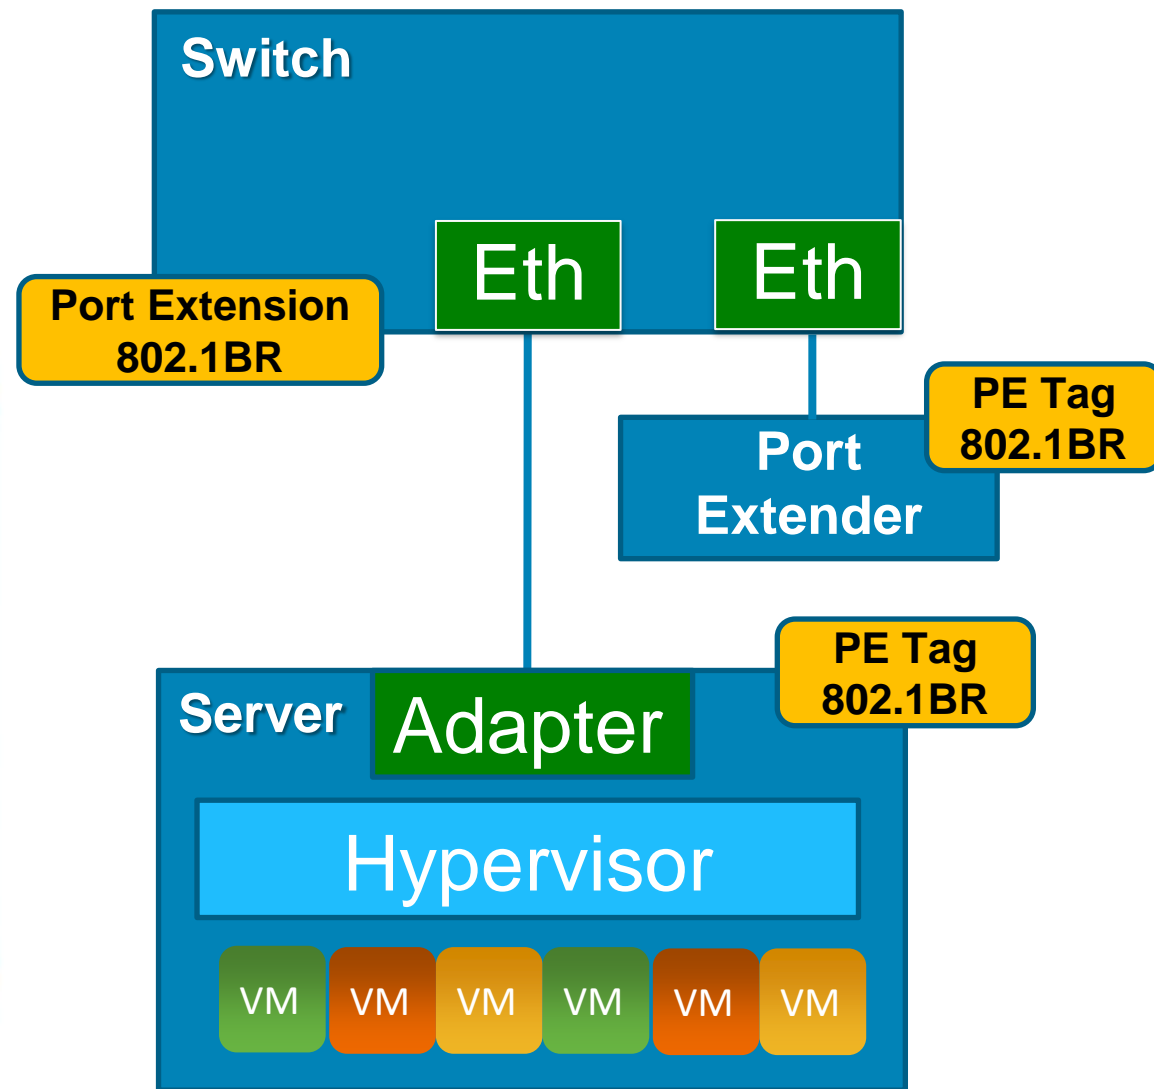
# IEEE 802.1BR: Bridge Port Extension



- Fully specifies a Port Extender (FEX Equivalent)
  - Extends ports of a switch to lower entities in a network
- Port Extenders are not individually managed
  - Their ports become ports of the controlling switch
- Cascading Port Extenders
  - Allows one to choose the appropriate controlling switch
  - Frame replication supported for efficient multicast / flooding
- Traffic from each “Extended Port” is reliably segregated to an E-channel and identified by a tag containing an E-channel identifier (ECID)
  - Does not require prior knowledge of MAC addresses; switch performs standard learning functions
  - Works with all devices including VEBs, VEPAs, individual VMs, physical services, and devices providing transparent services
- Controlling Bridge + PE = Extended Bridge
  - Single Point of Management

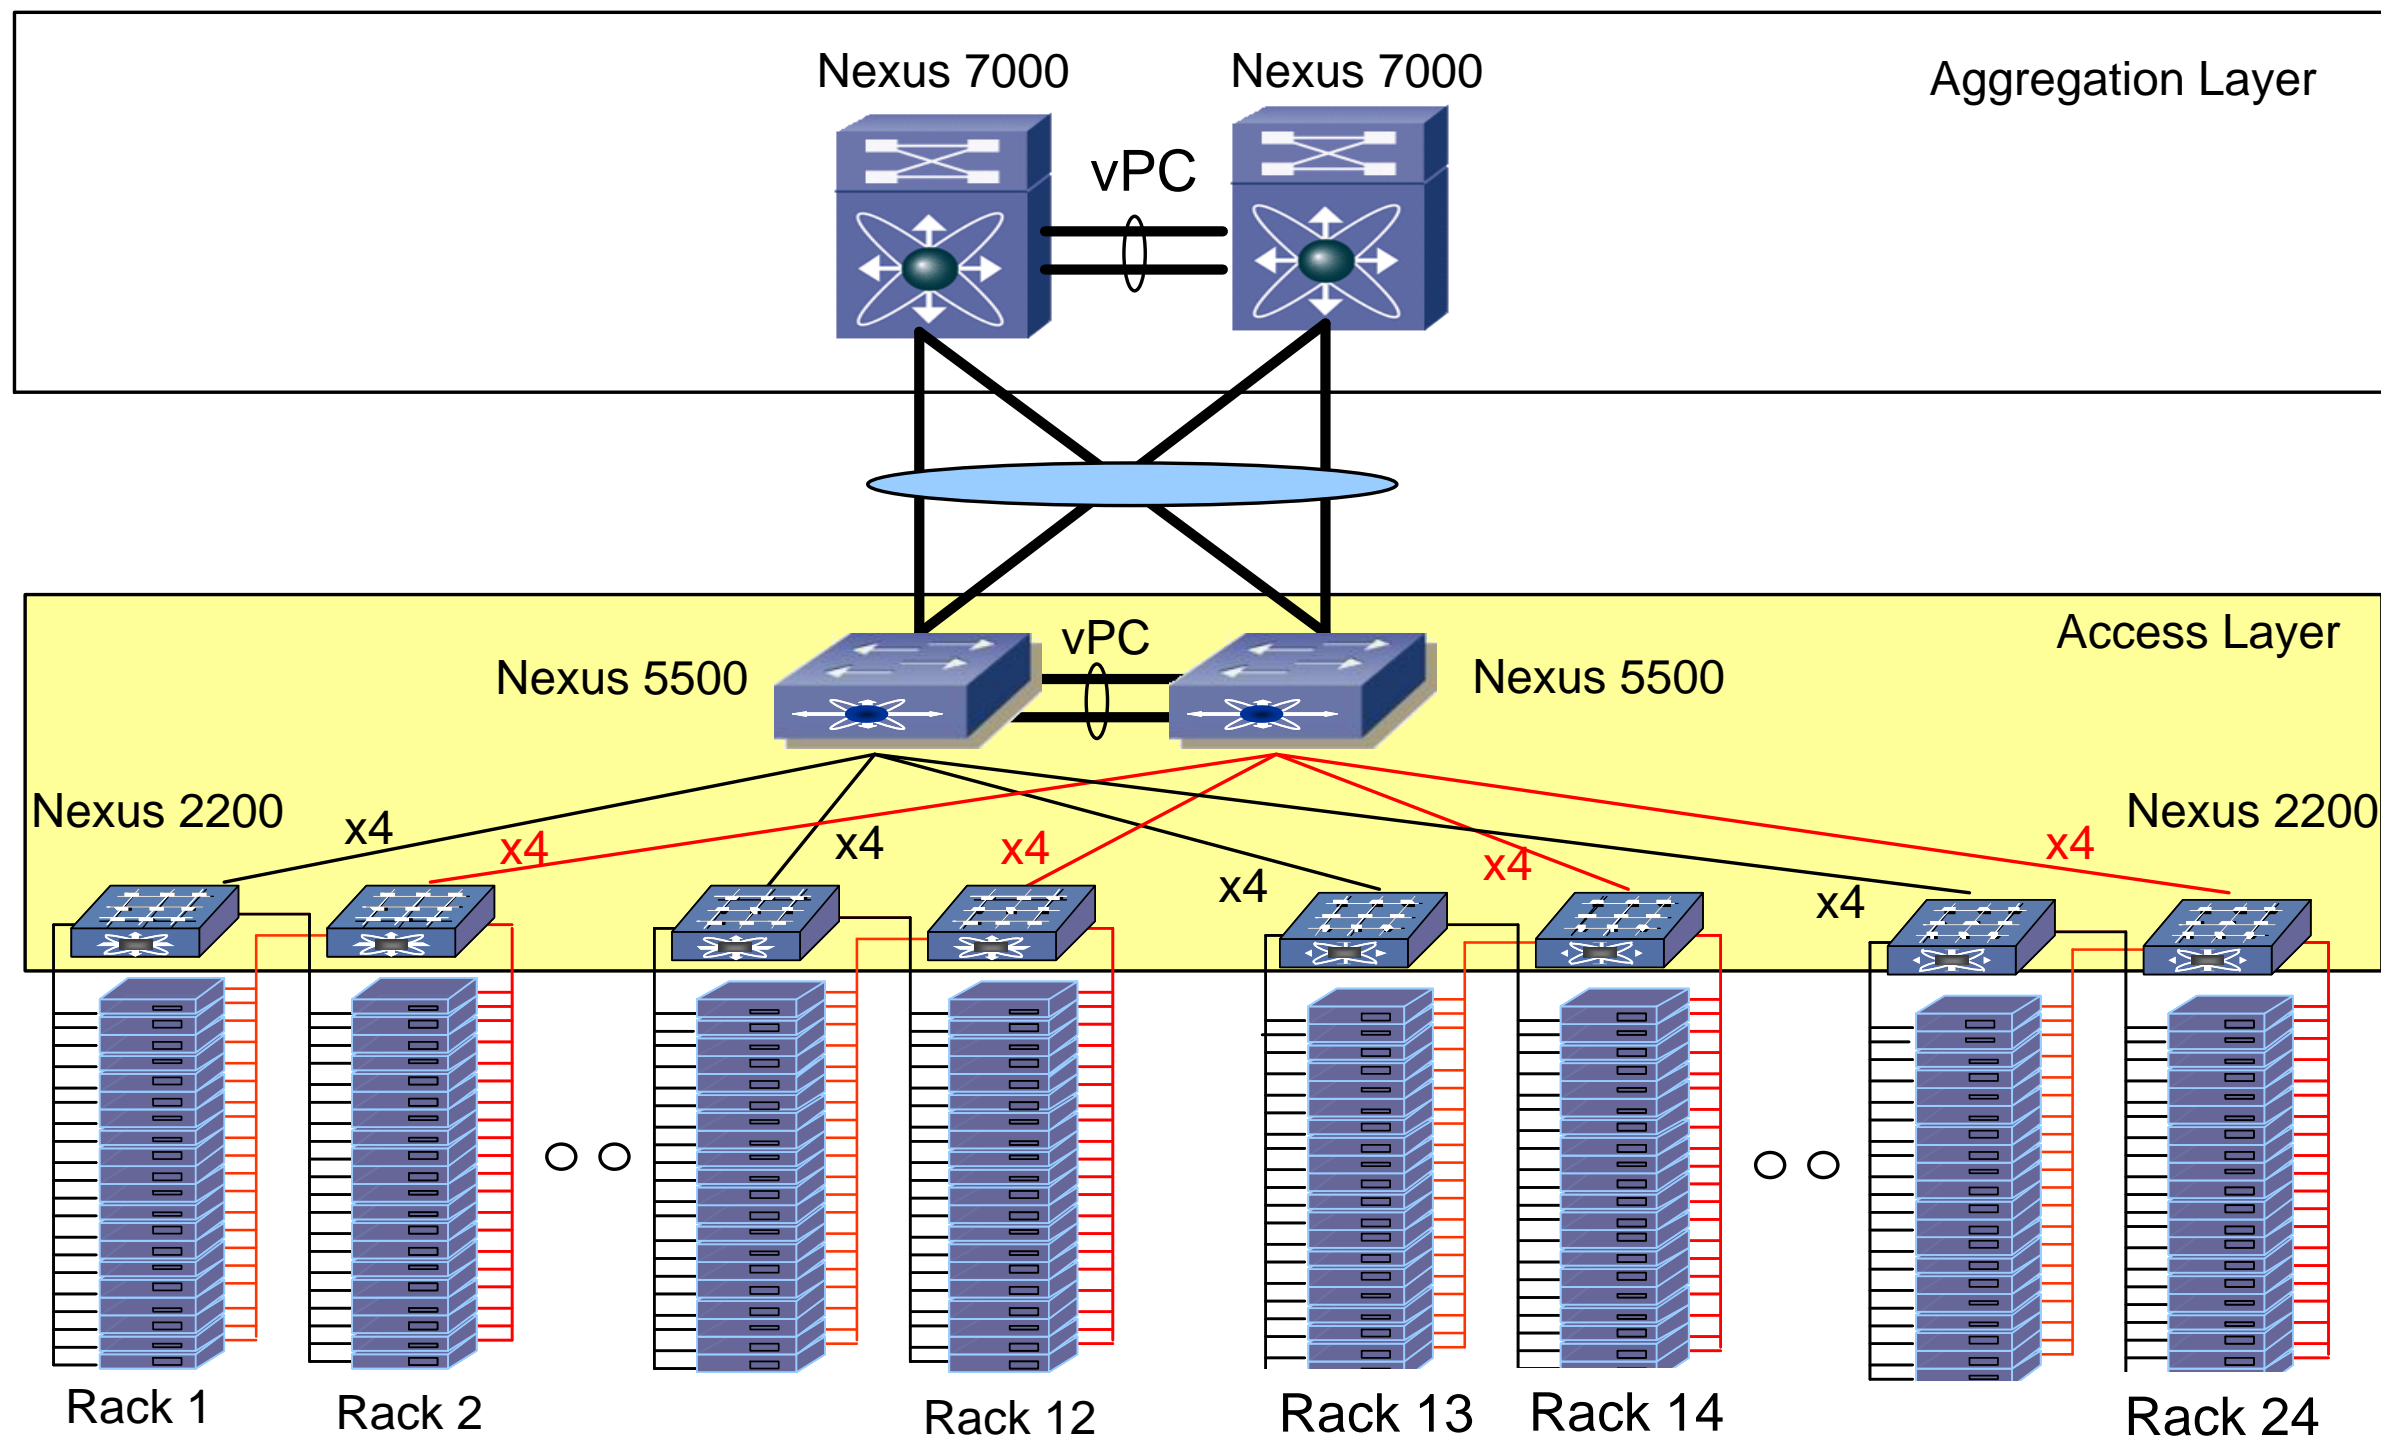


# Representation of Port Extender & Fabric Extender



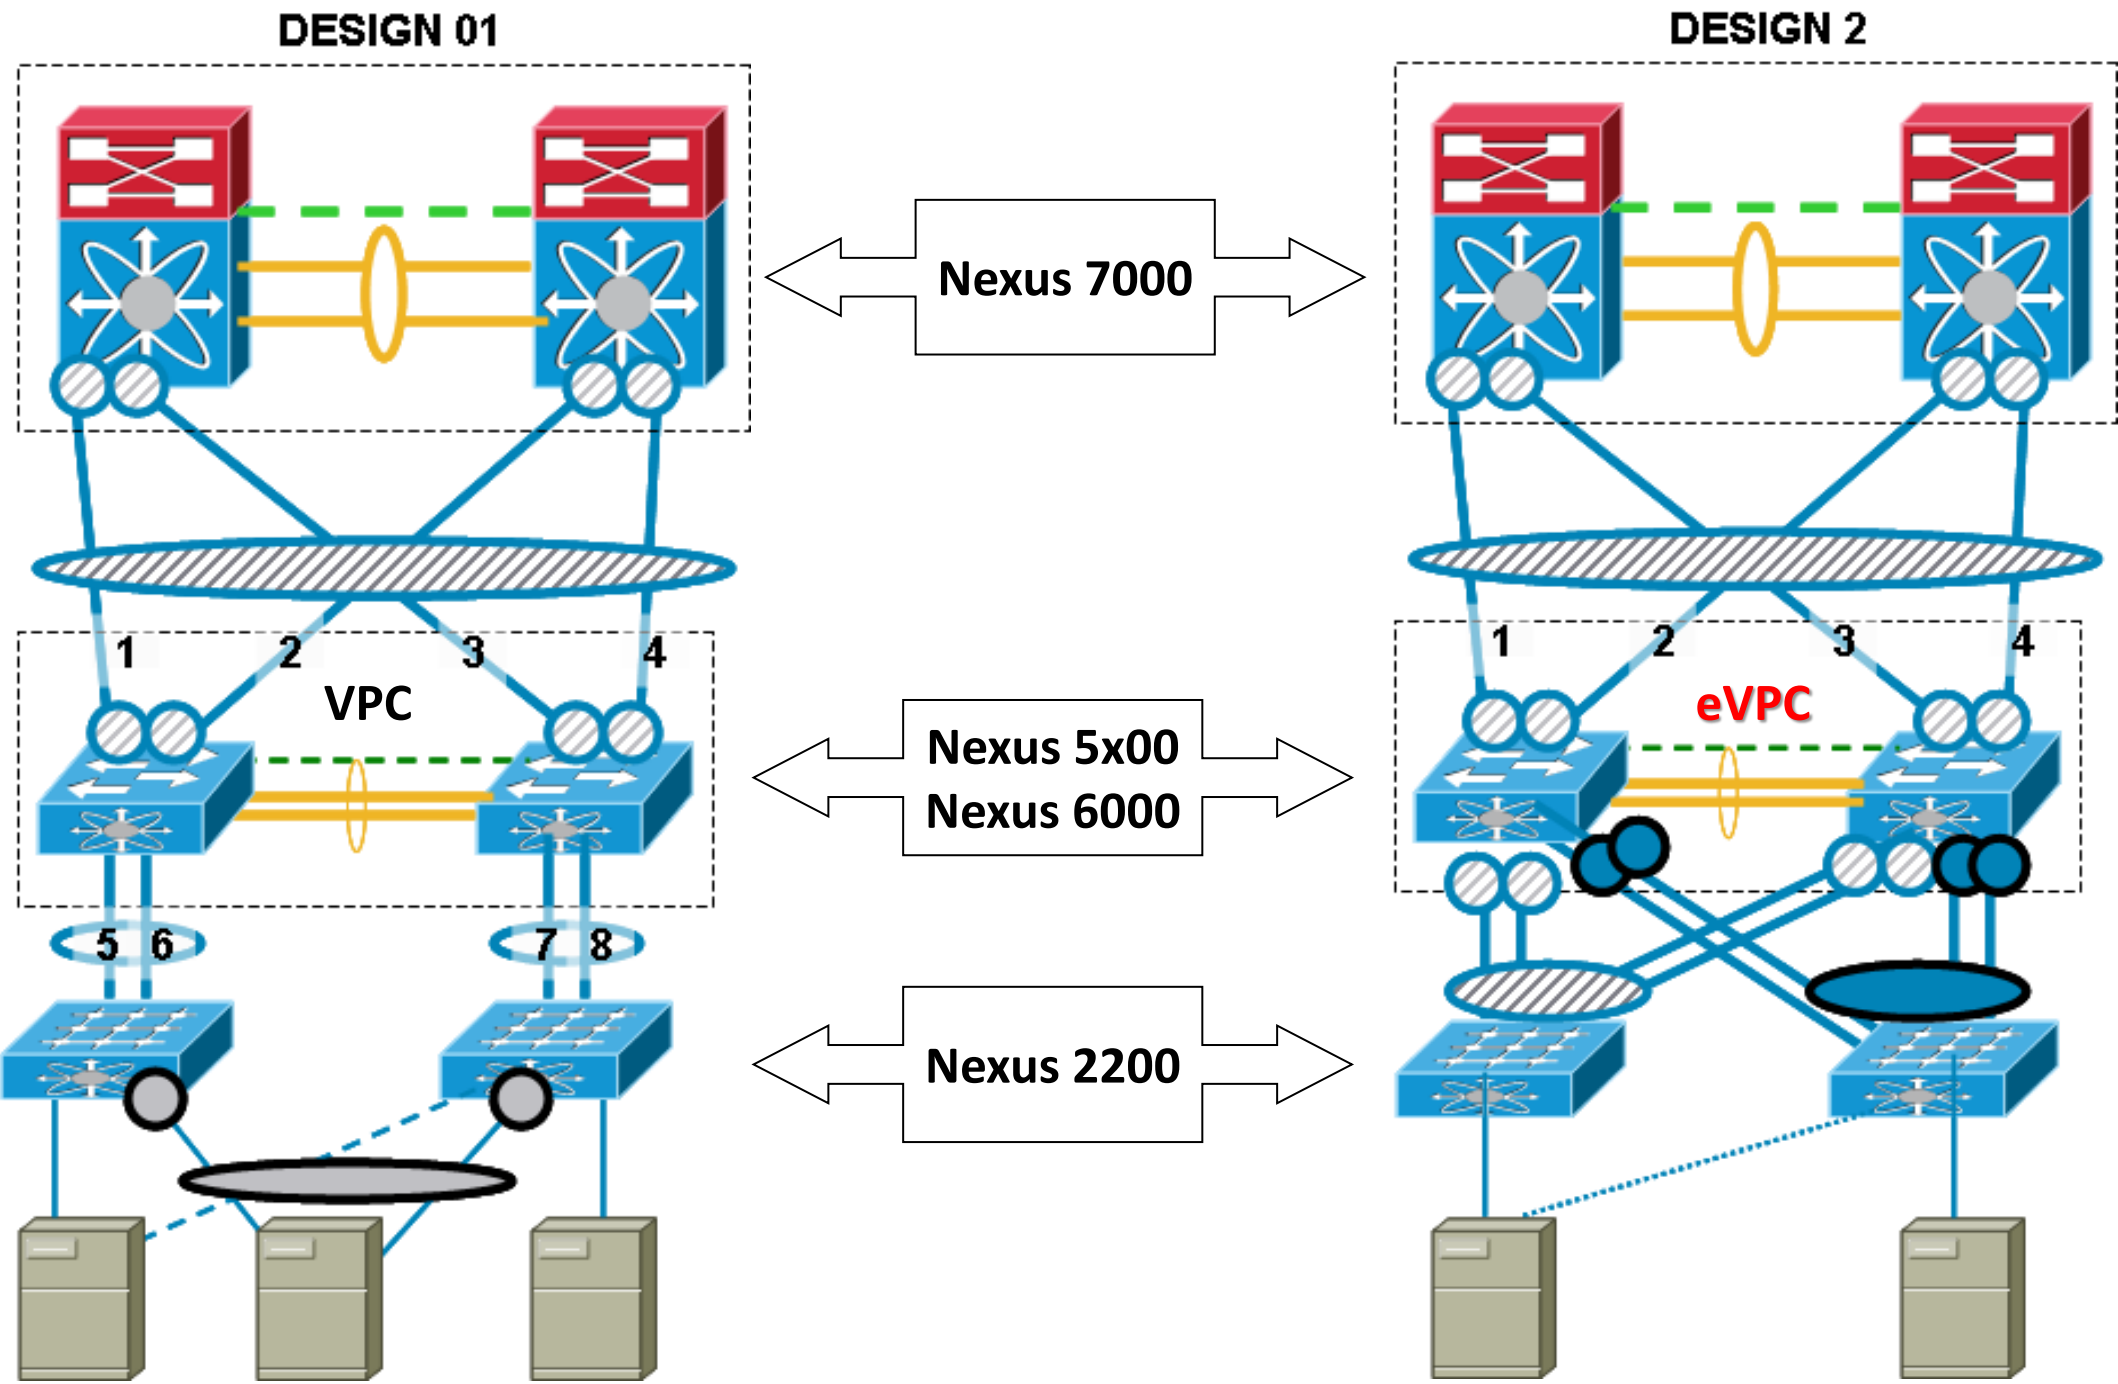
IEEE Bridge Port Extender = Cisco FEX (Fabric Extender)

# To2R: Nexus 2200 Deployment Example



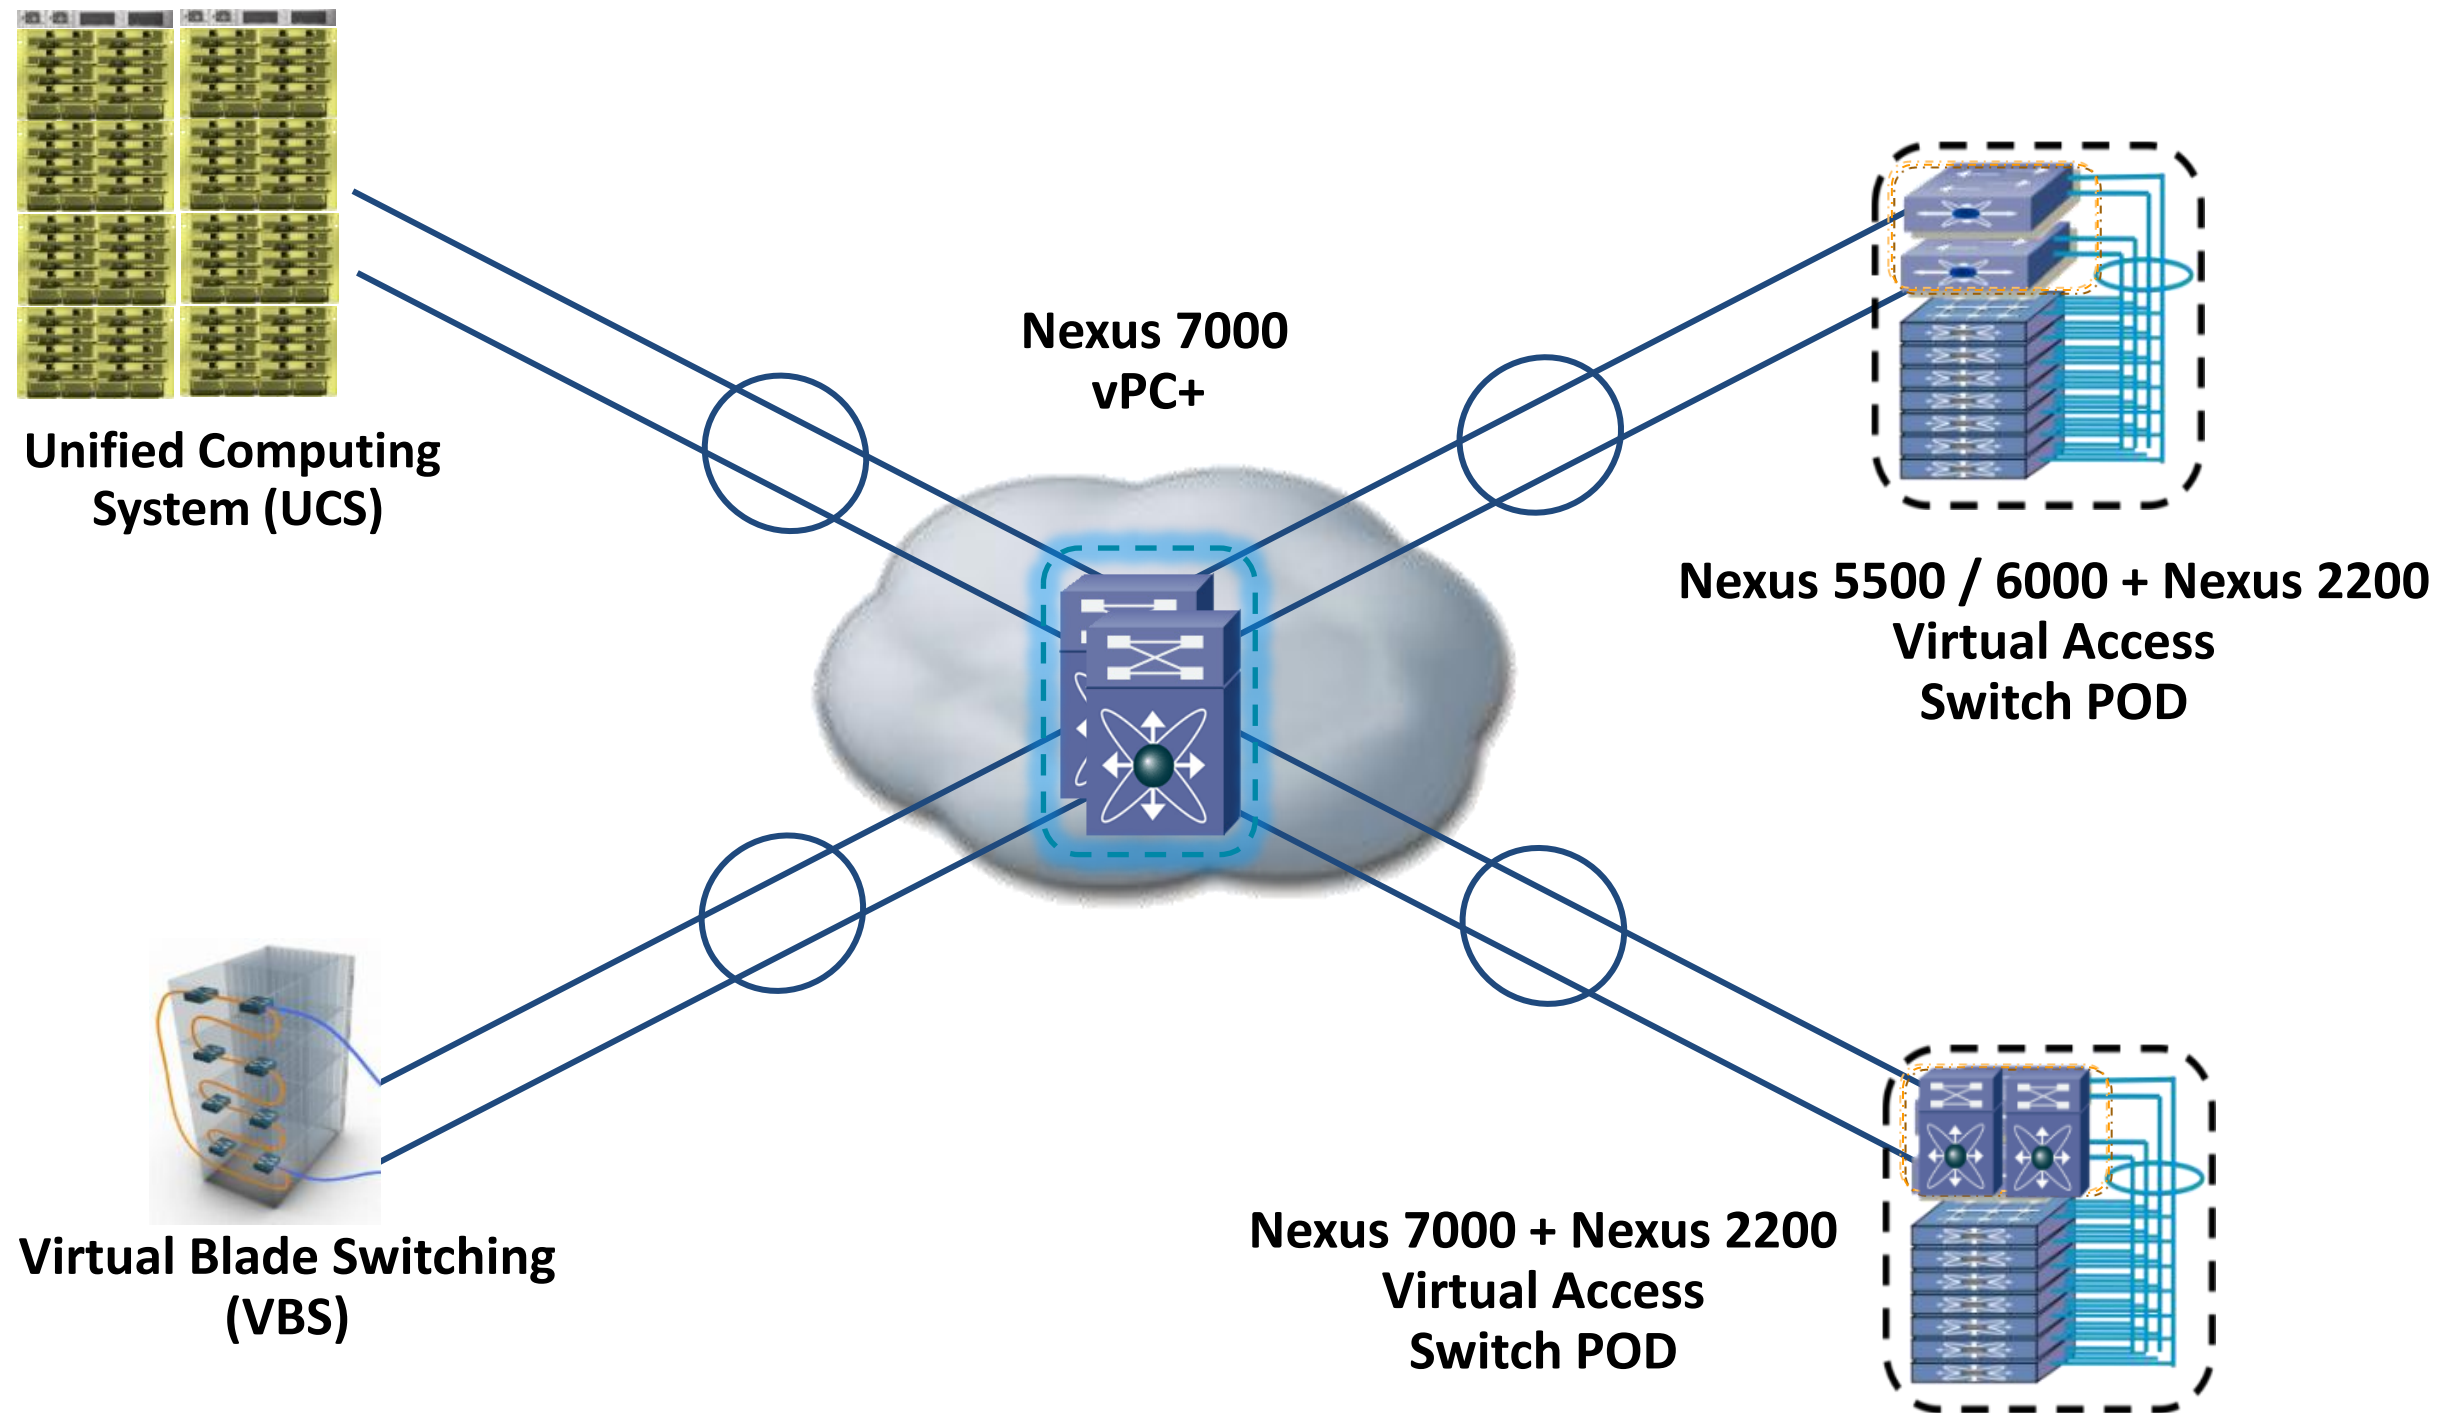


# vPC Between Nexus 7000 and Nexus 5500/6000 and Nexus 2200 — Active/Active




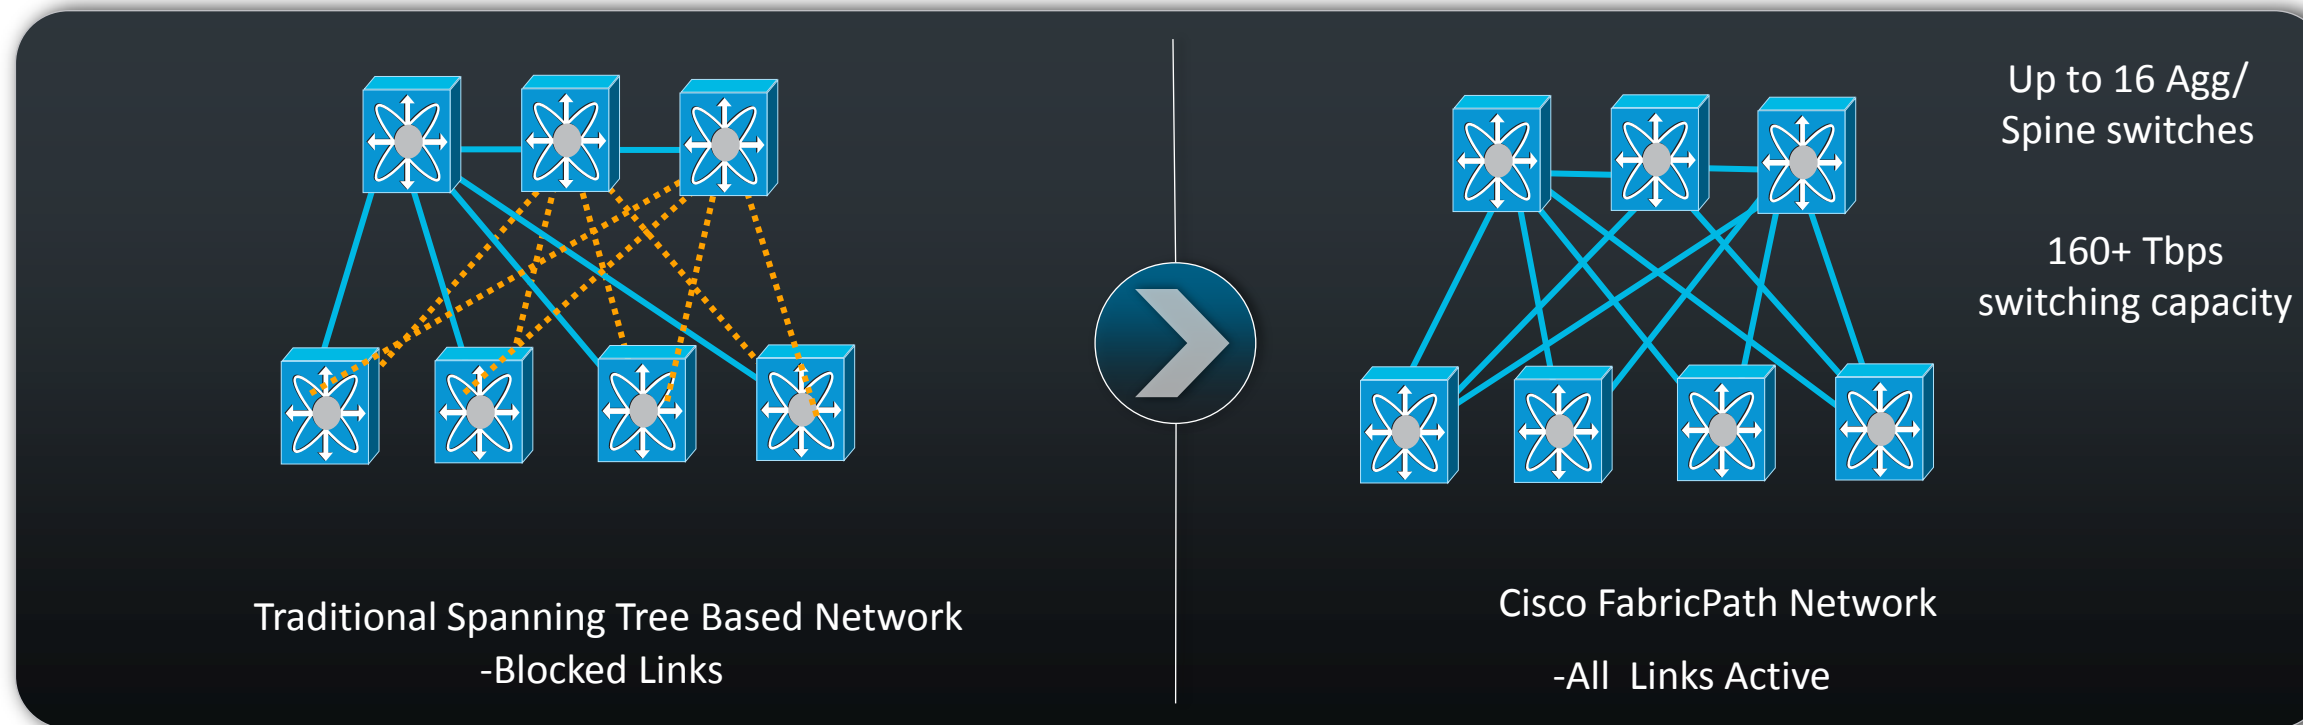
Enhanced VPC (eVPC) supported with Nexus 5500 only

# Logical View: Star Topology Without L2 Loops



# Cisco FabricPath

## Scaling and Simplifying Layer 2 Ethernet Networks



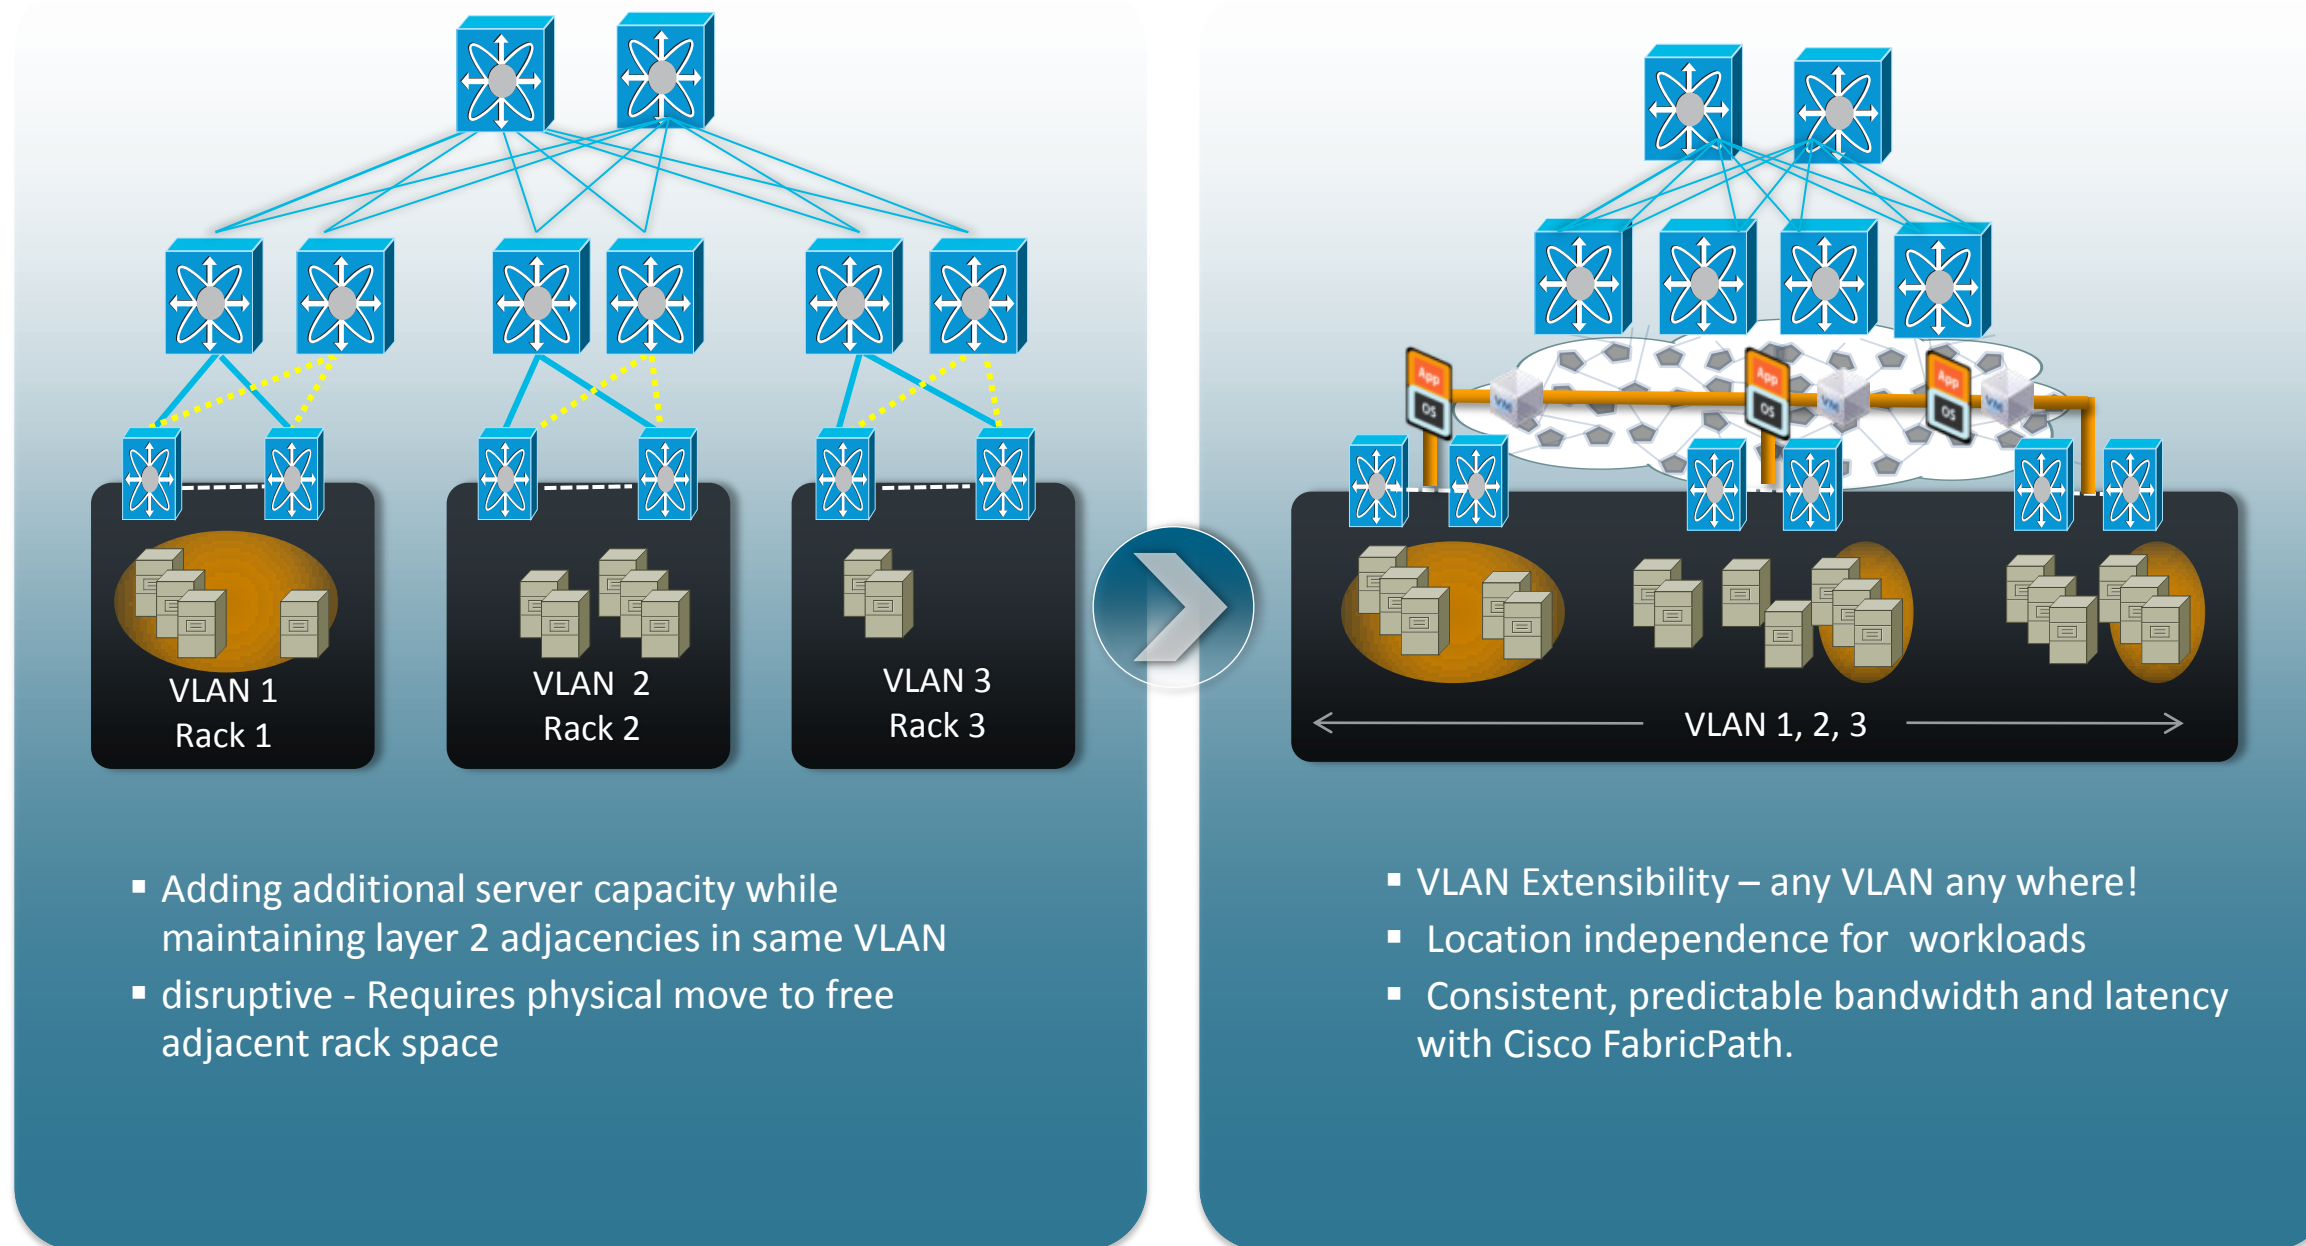
- Eliminate Spanning tree limitations
- Multi-pathing across all links, high cross-sectional bandwidth
- High resiliency, faster network re-convergence
- Any VLAN, any where in the fabric eliminate VLAN Scoping





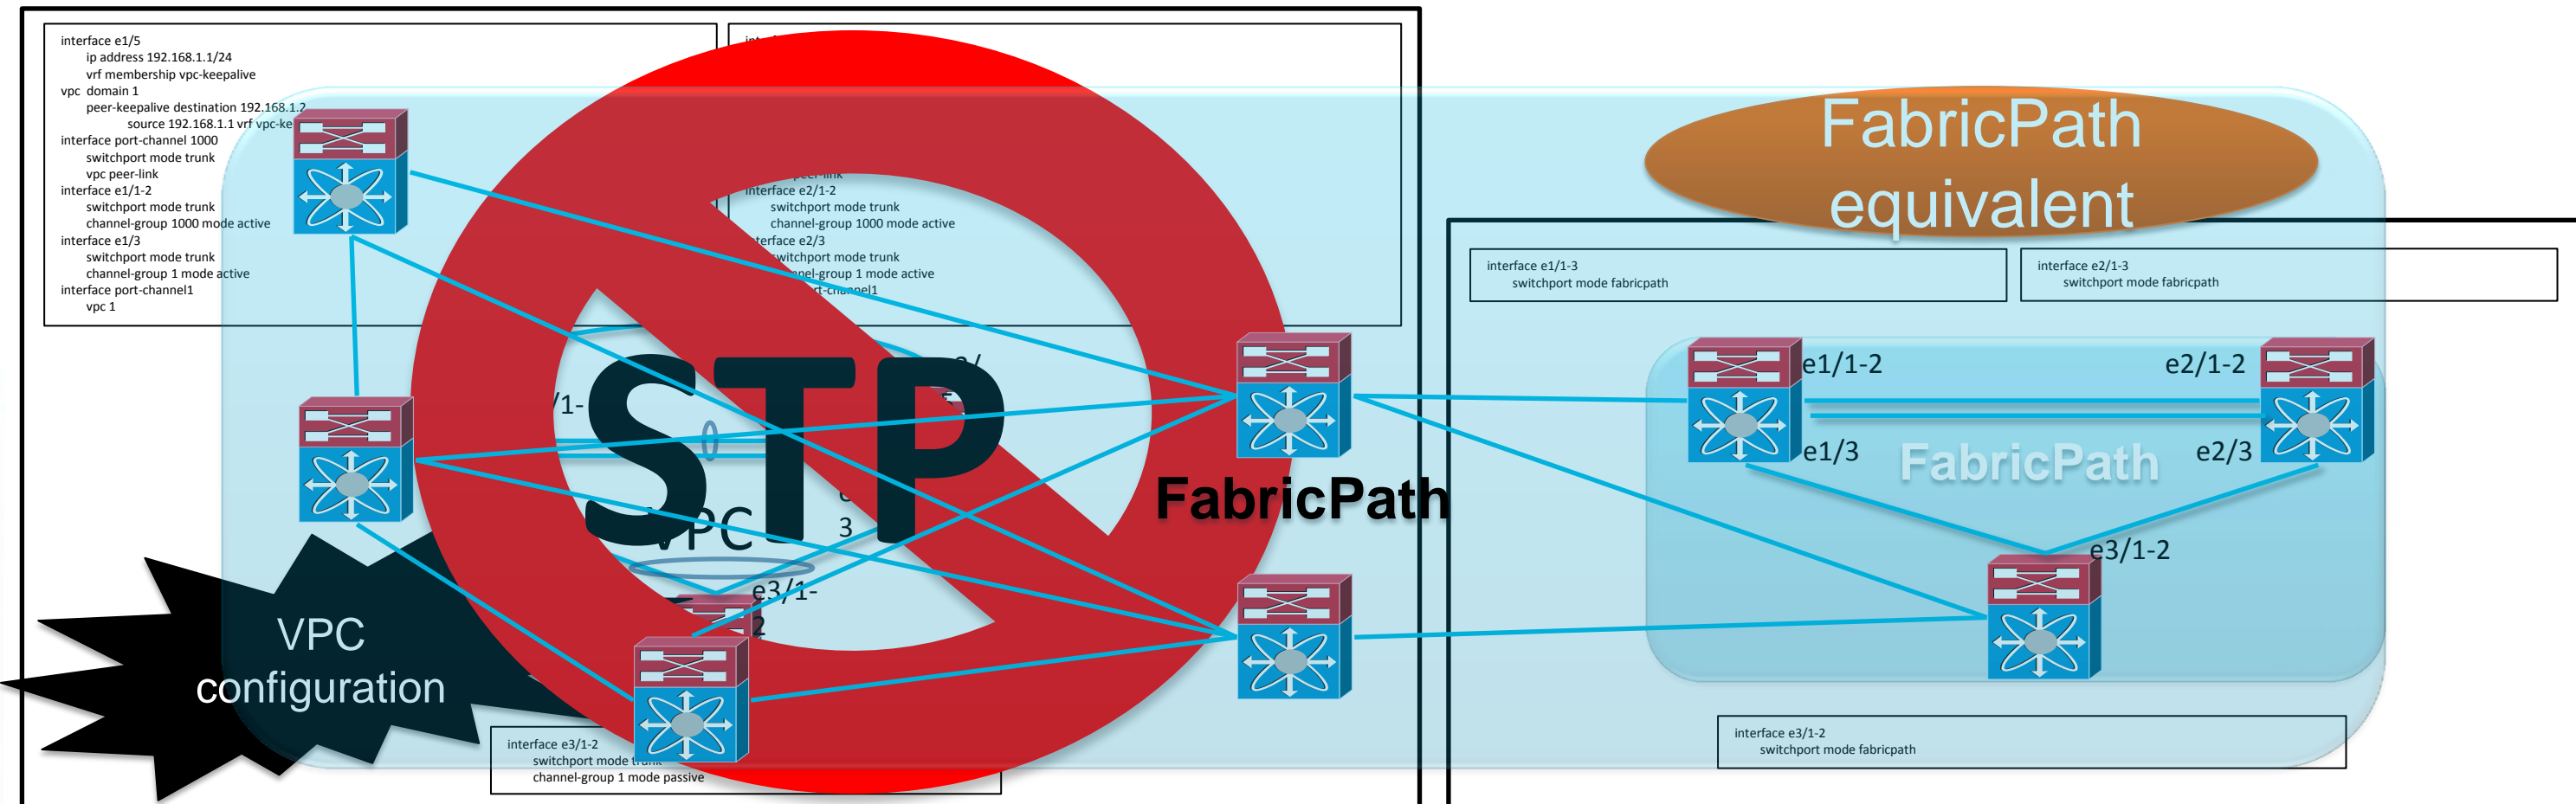
# Example: Handling Application Growth

**Scenario:** Application grows beyond currently compute capacity and allocated rack space causing network disruptions and physical changes





# Why Migrate from vPC to FabricPath?

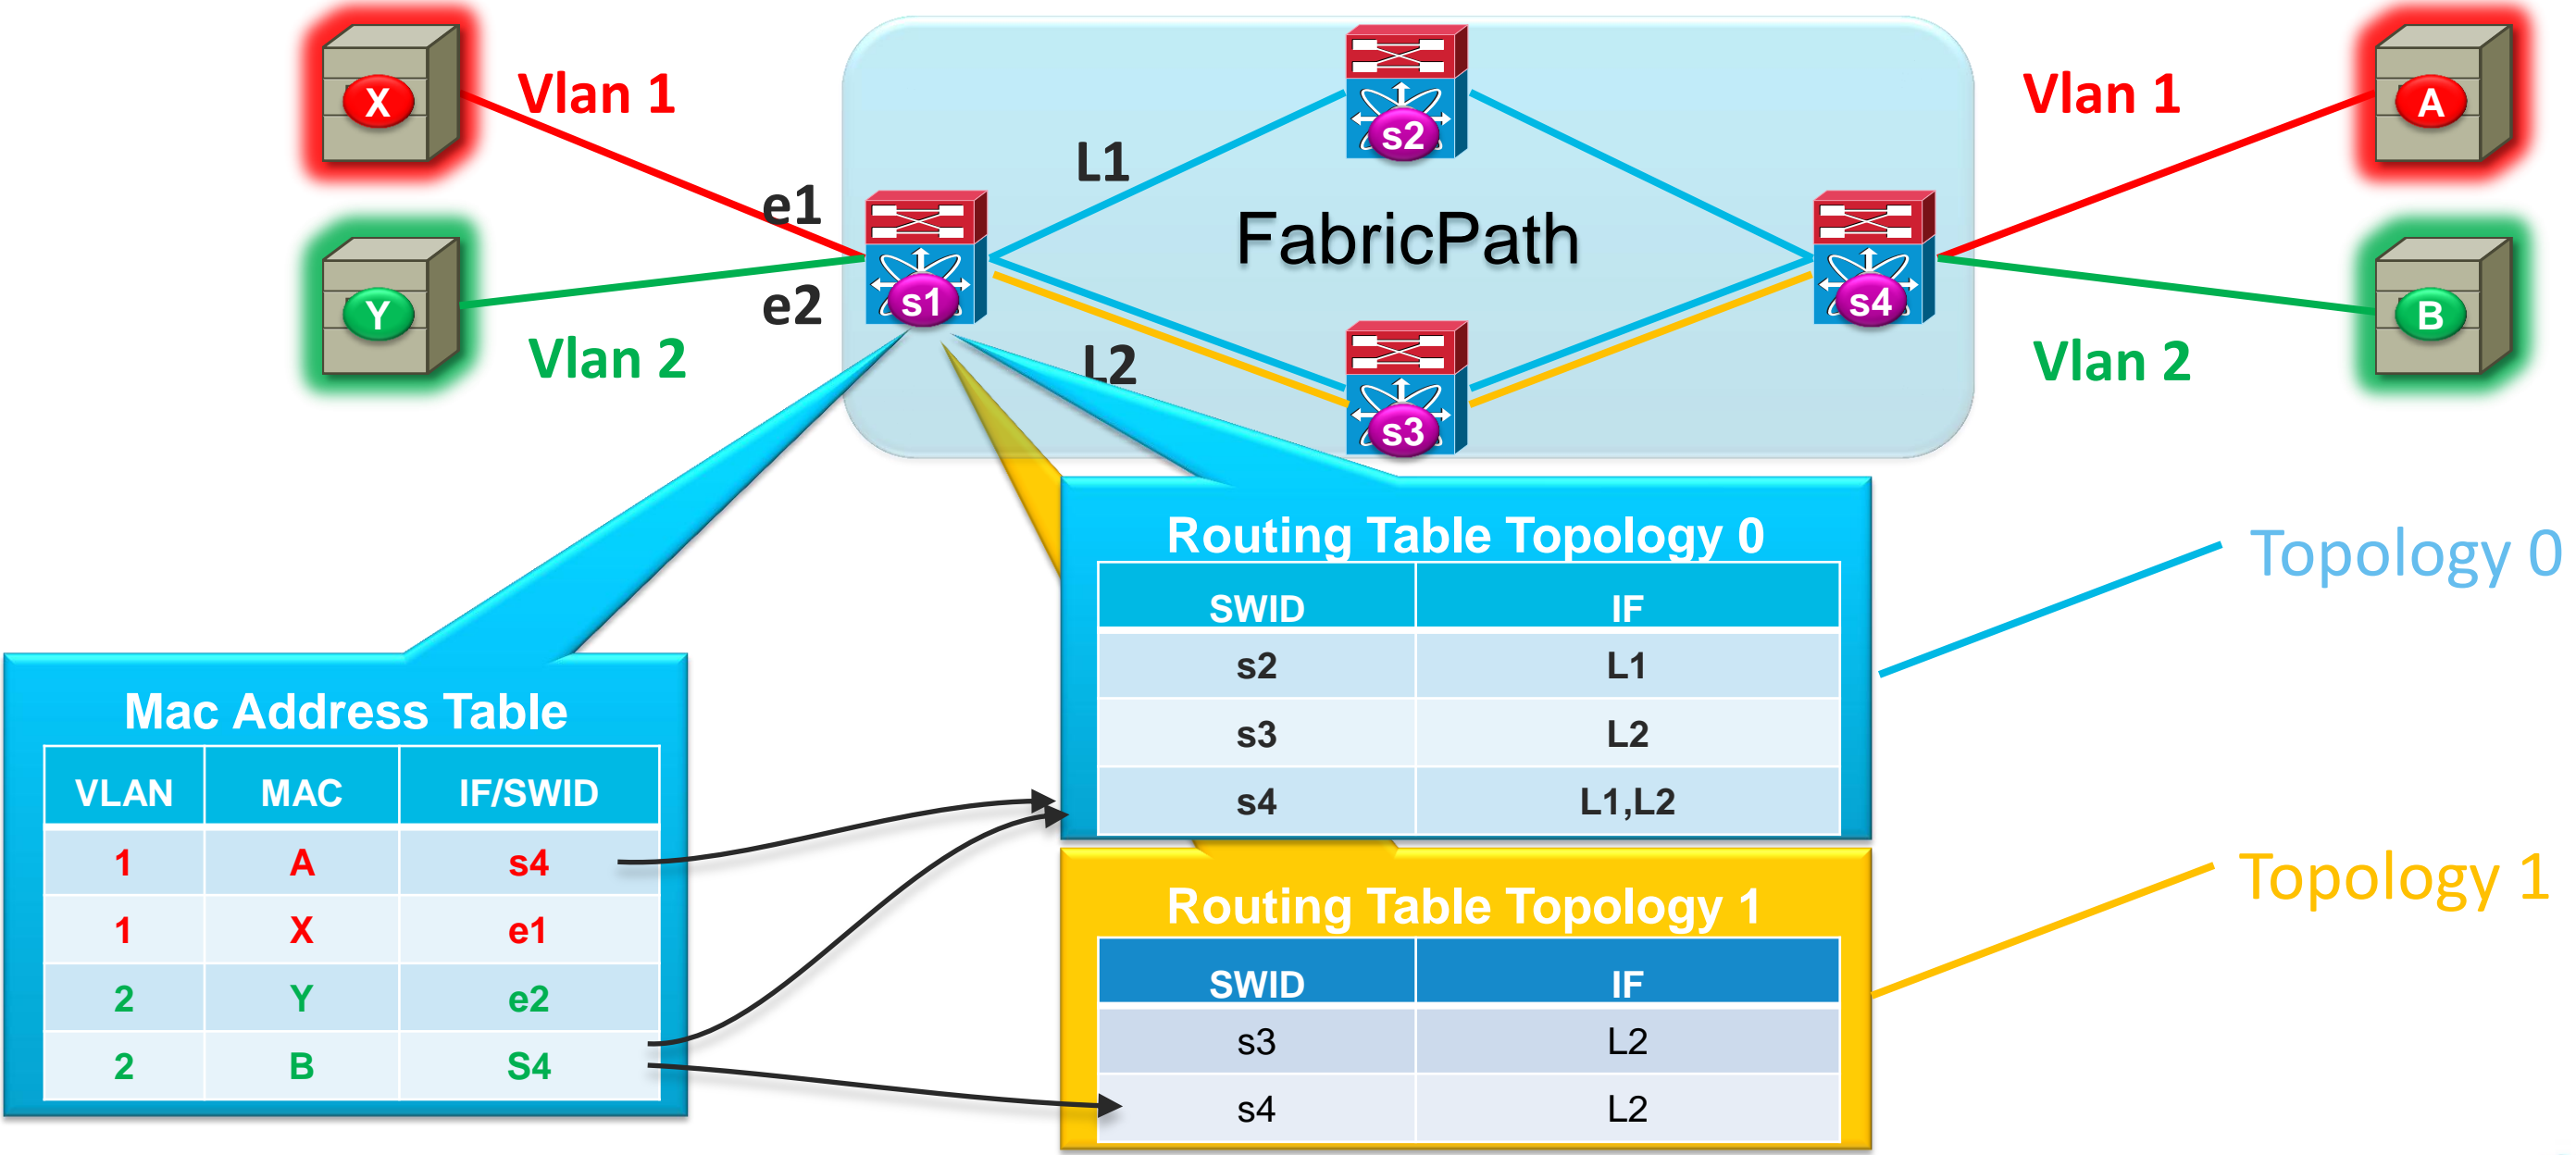


- Much Simple and straight configuration
- Absolutely no Spanning Tree, forwarding based on routing
- Any mesh possible, non-disruptive cabling changes

# FabricPath Forwarding

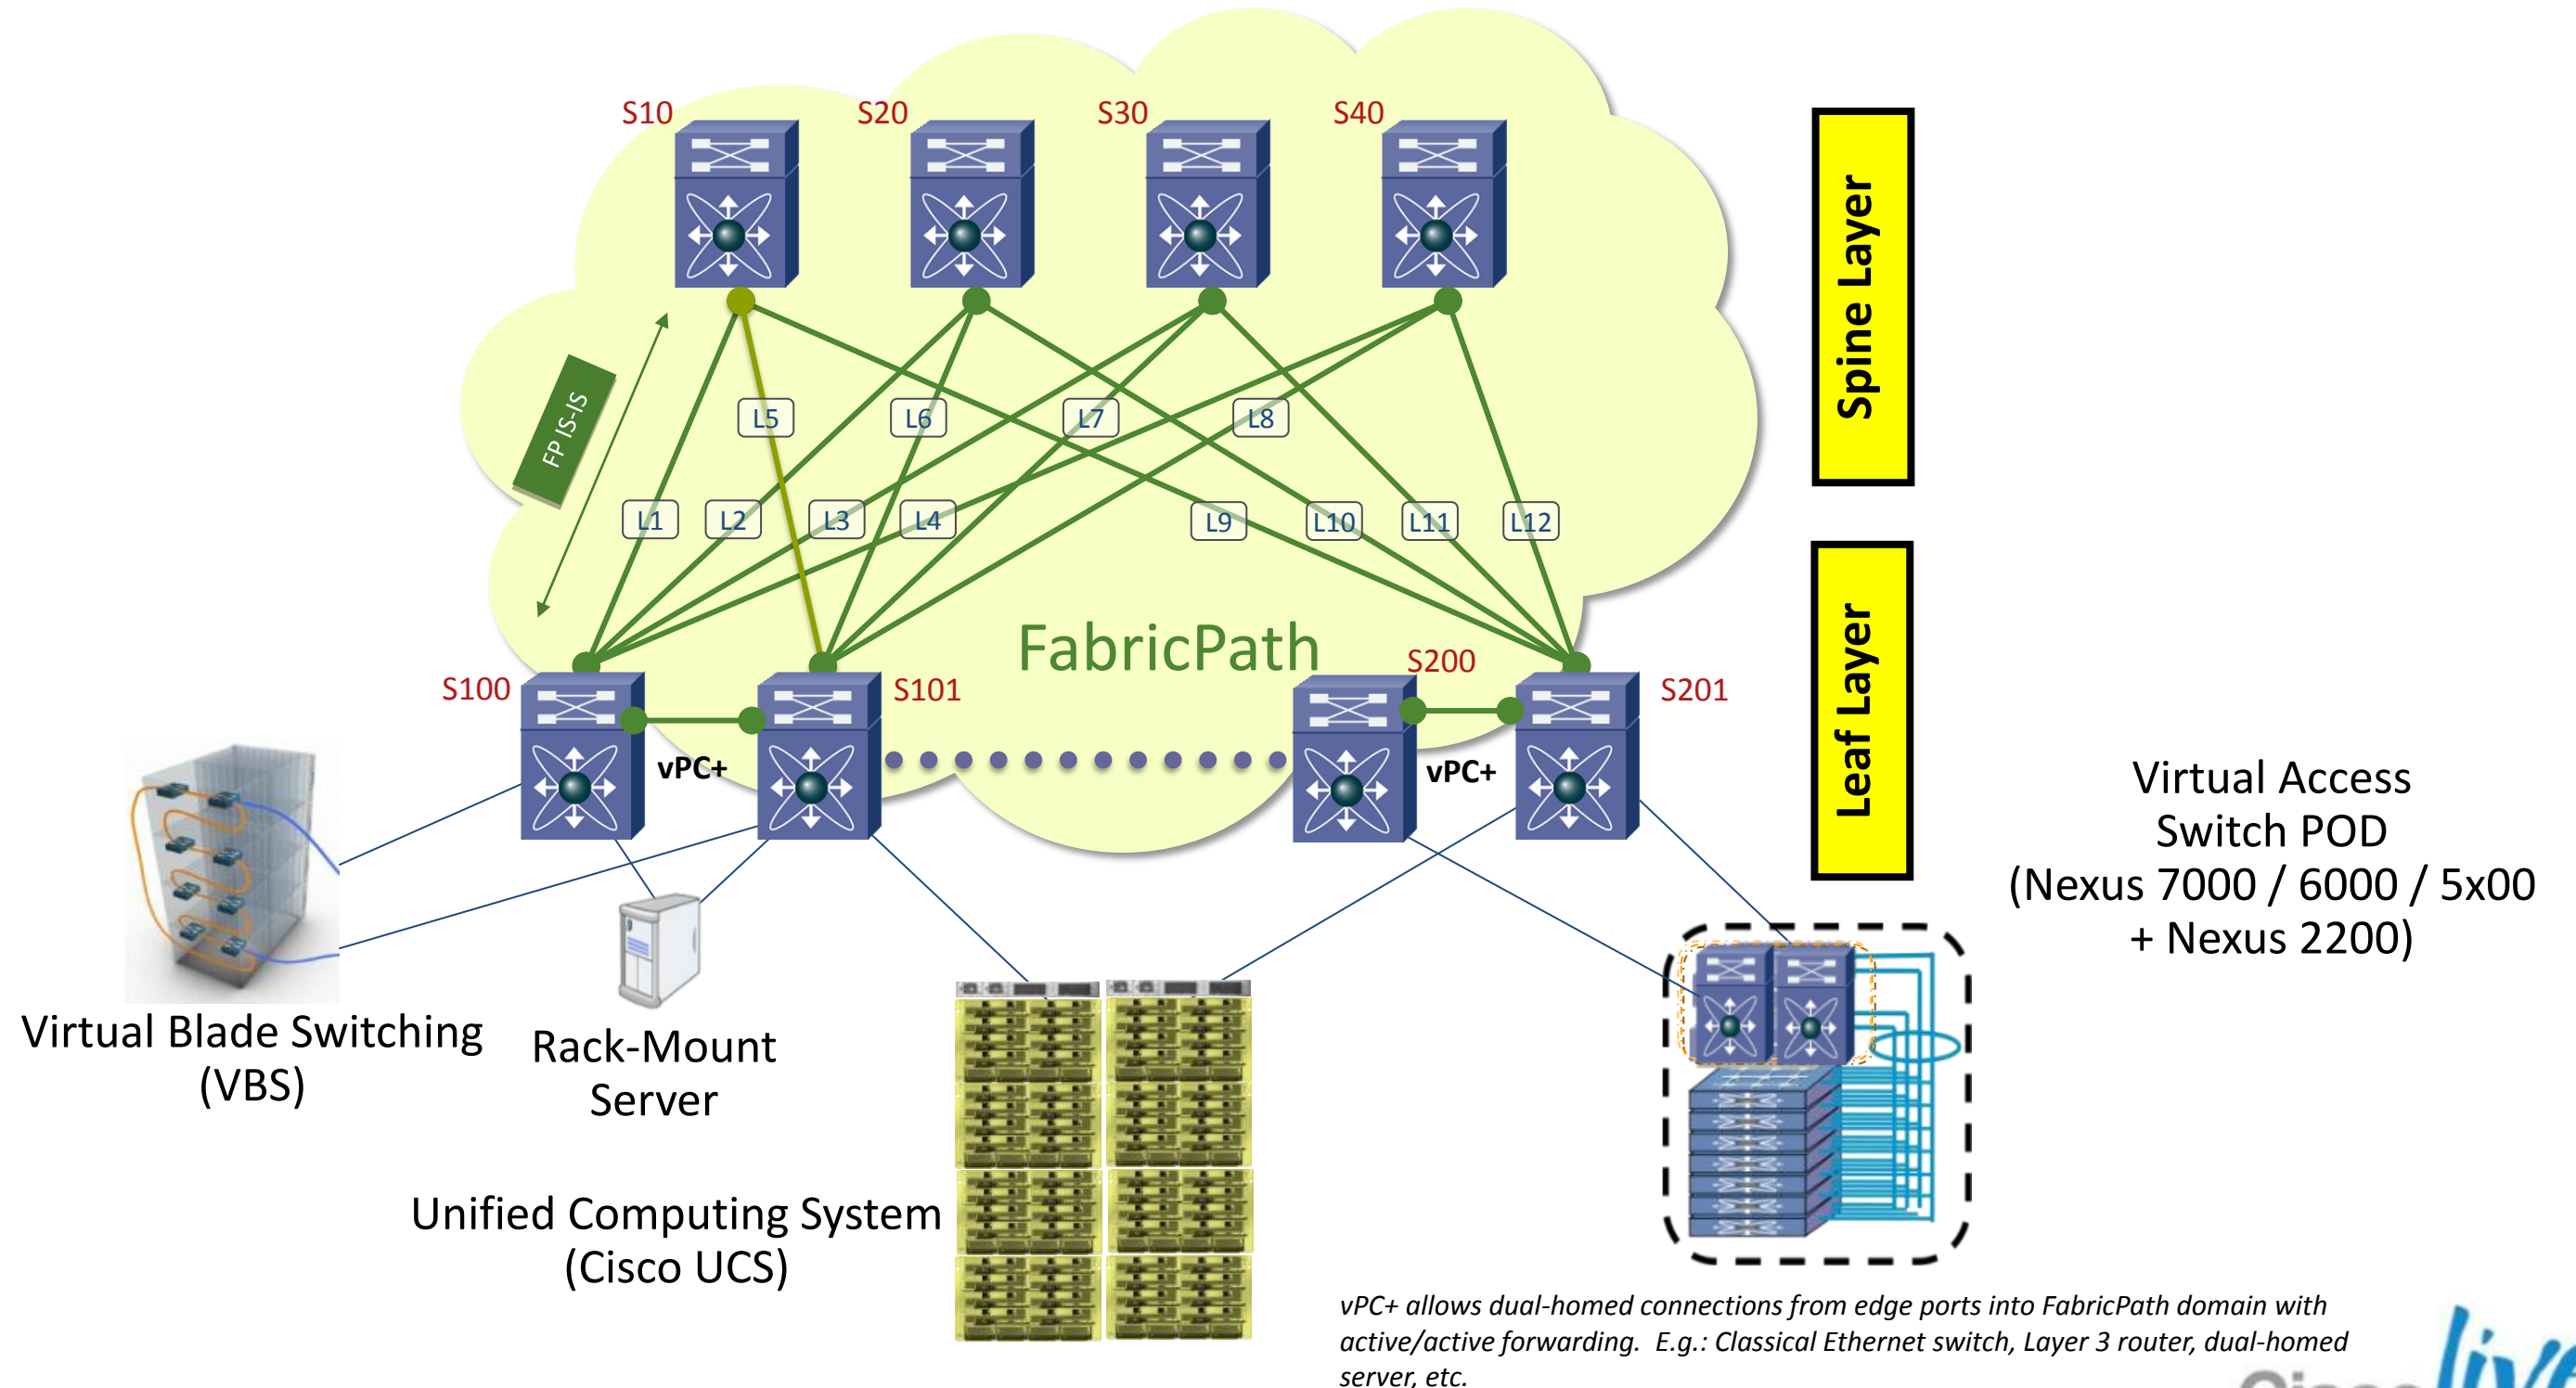


Based on routing tables. VLAN pruning automatic




# Logical View with FabricPath:

## Distributed Topology without L2 loops



# Agenda

- **Data Centre Virtualisation Overview**
- **Front-End Data Centre Virtualisation**
  - DC Edge Layer
  - Aggregation / Spine Layer
  - Access / Leaf Layer
- **Virtual Networking & Cloud Network Services**
  - Virtual Access Layer: Nexus 1000v
  - Cloud Network Services
  - Inter DC Services: OTV & LISP
- **Virtual Compute & IO Virtualisation**
  - Unified Computing System
  - Hypervisors & Virtual IO @ Cisco UCS
- **Virtualised Storage & SAN**
  - Unified Fabric
  - Virtualised SAN and Storage
- **Software Defined Network & Orchestration**
  - SDN & Cisco Open Network Environment (ONE)
  - Cloud Orchestration
- **Implementation Examples**
- **Q&A**






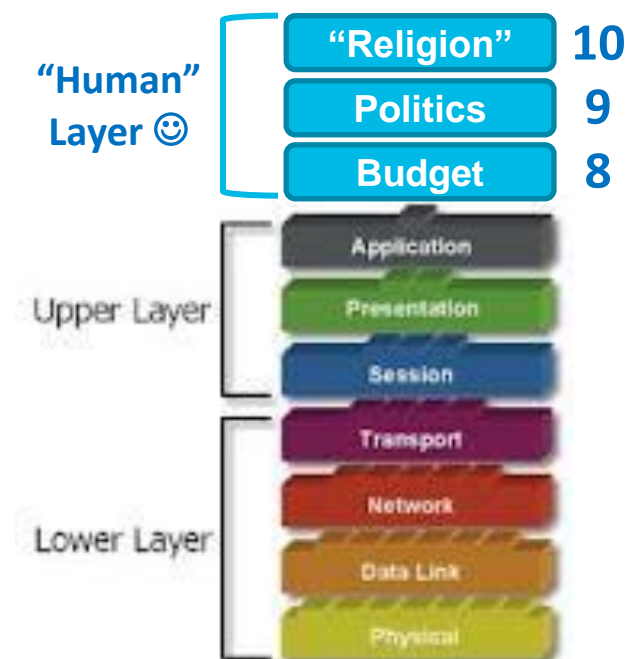
# What Happens When We Mix Network and Server Virtualisation?




# Current View of the Access Layer with VMs



- Typically provisioned as trunk to the server running ESX
- No visibility to individual traffic from each VM
- Unable to troubleshoot, apply policy, address performance issues




# Networking for Server Virtualisation



## Problems:

- Dynamic Migration of VMs may move them across physical server ports—policy must follow
- Impossible to view or apply policy to locally switched traffic
- Need collaboration between network and Virtualisation admin




# Cisco VN-Link


What is that and which problems does it solve ?



- VN-Link (or Virtual Network Link) is a term that refers to a VM specific link that is created between the VM and Cisco switch.
- Logical equivalent & combination of a NIC, a Cisco switch interface and the RJ-45 patch cable that hooks them together.



# VN-Link View of the Access Layer



- Nexus 1000V and VN-Link provide visibility to the individual VMs
- Policy can be configured per-VM
- Policy is mobile within the ESX cluster

Nexus 1000V  
Distributed Virtual Switch

# Cisco Virtual Networking Vision

Nexus 1000V

Multi-Hypervisor

VMWare vSphere  
Microsoft Hyper-V @ Windows Server 2012  
KVM  
Xen Server

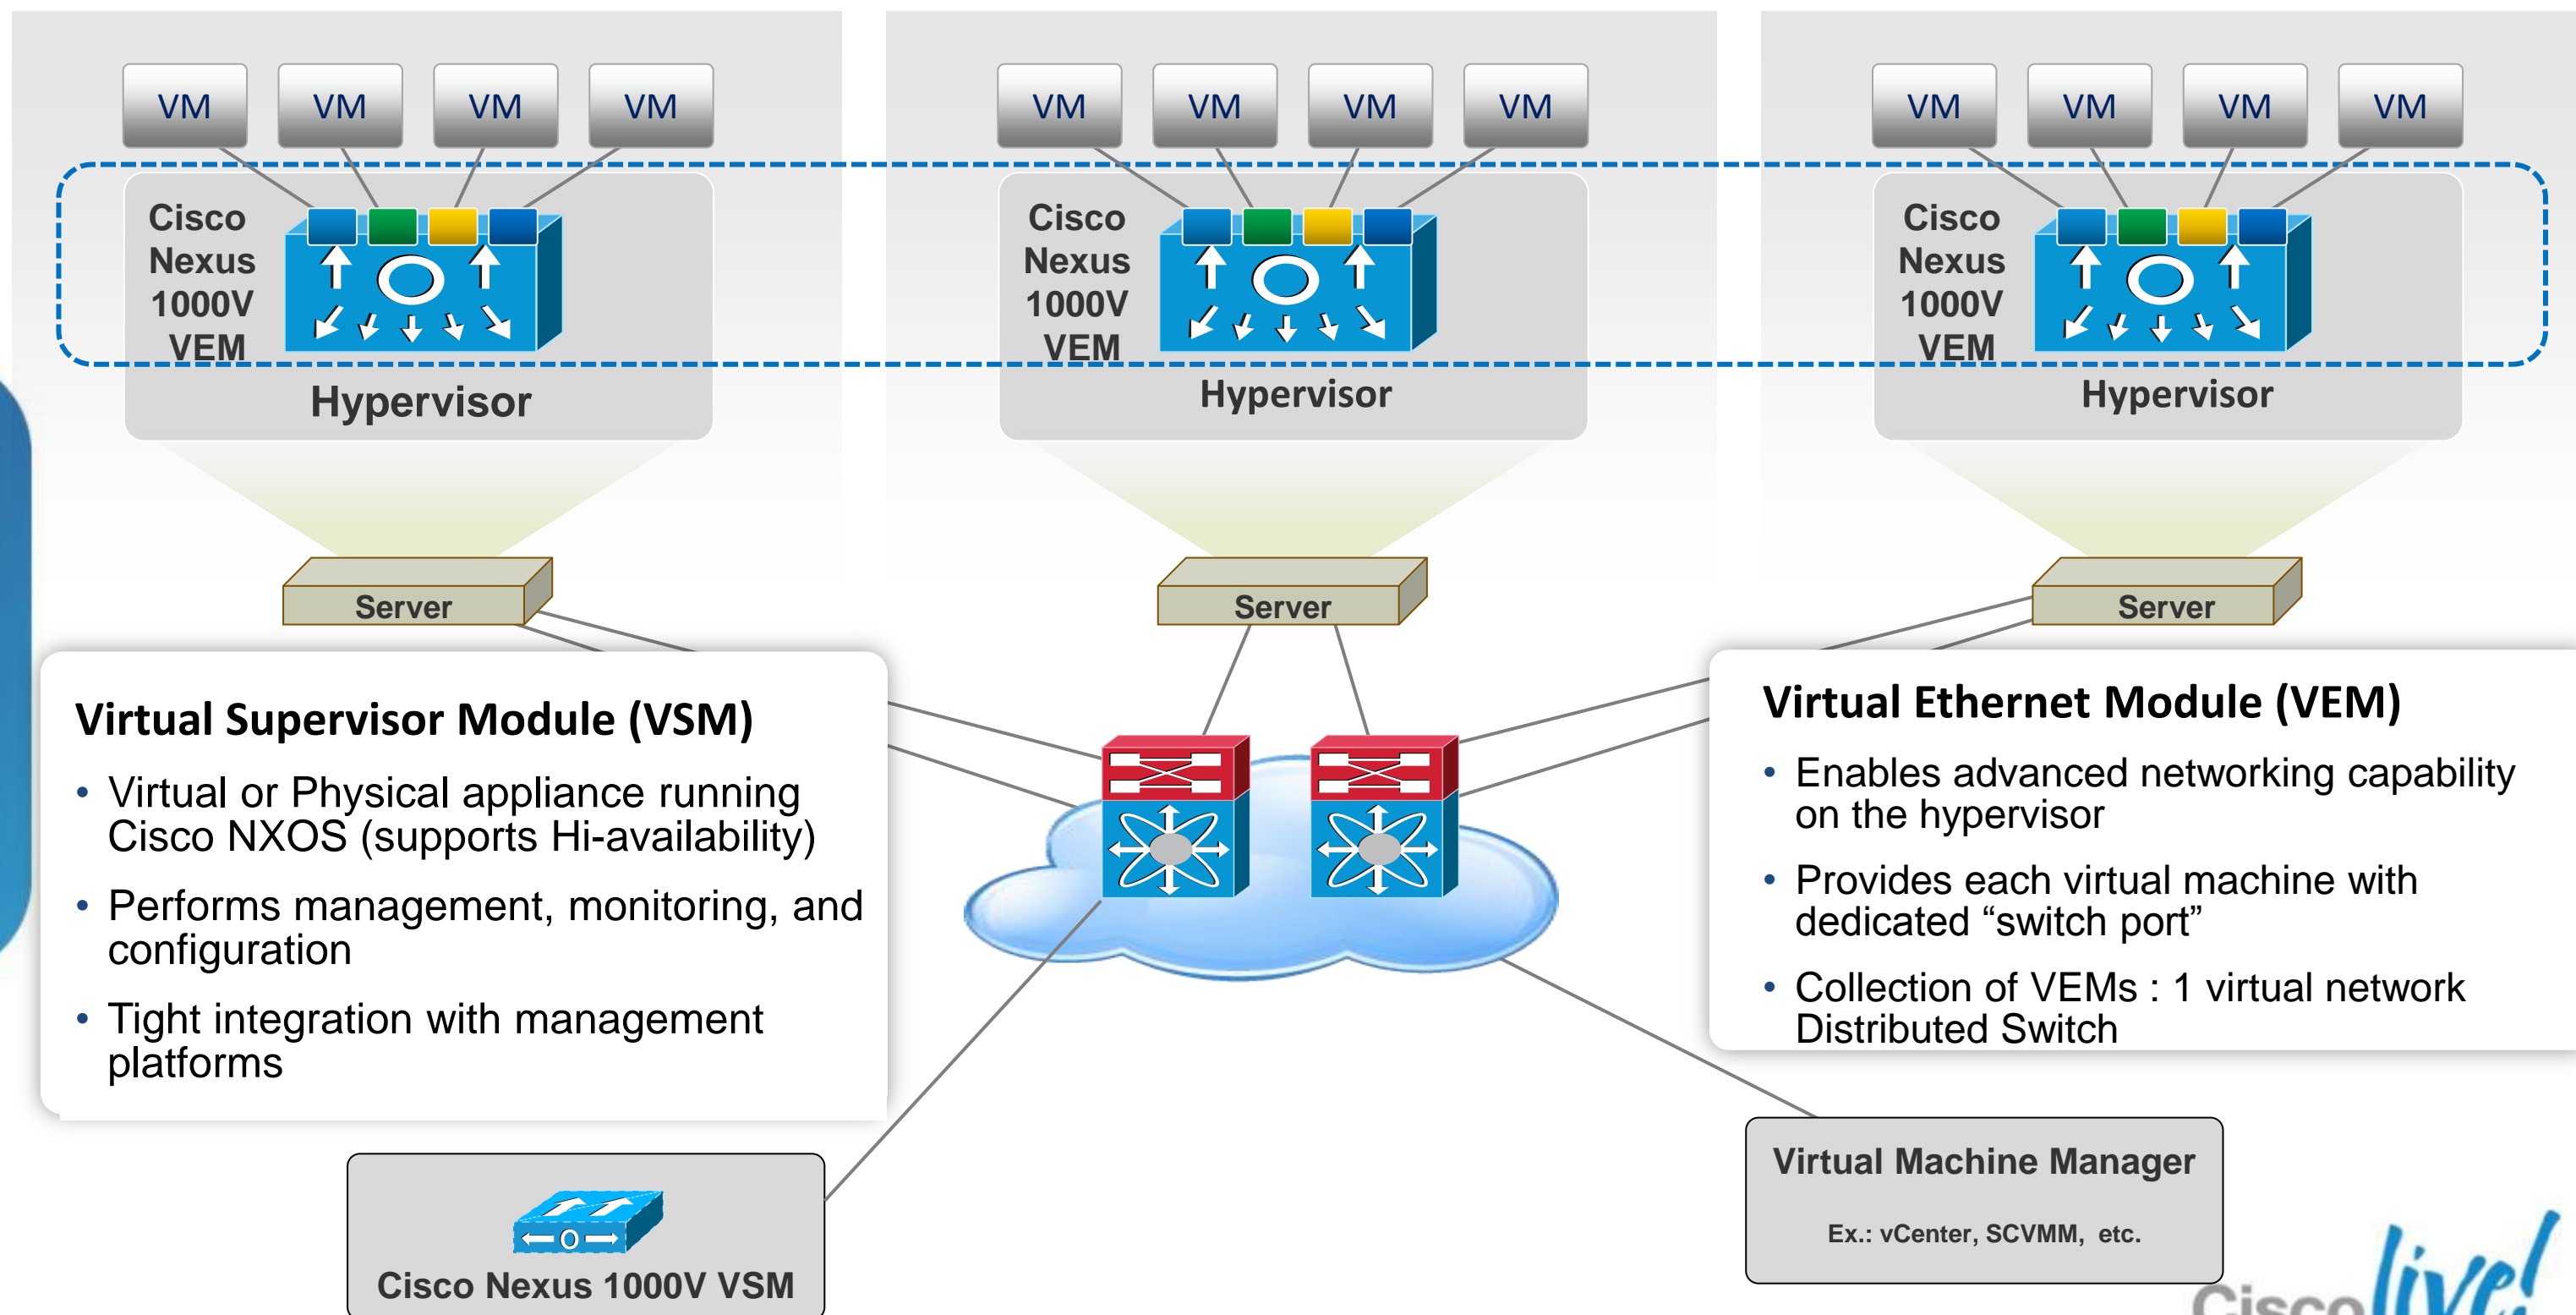
Multi-Services

Firewall, WAN/App optimisation, ADC,  
Cloud Router, WAF, VM Segmentation

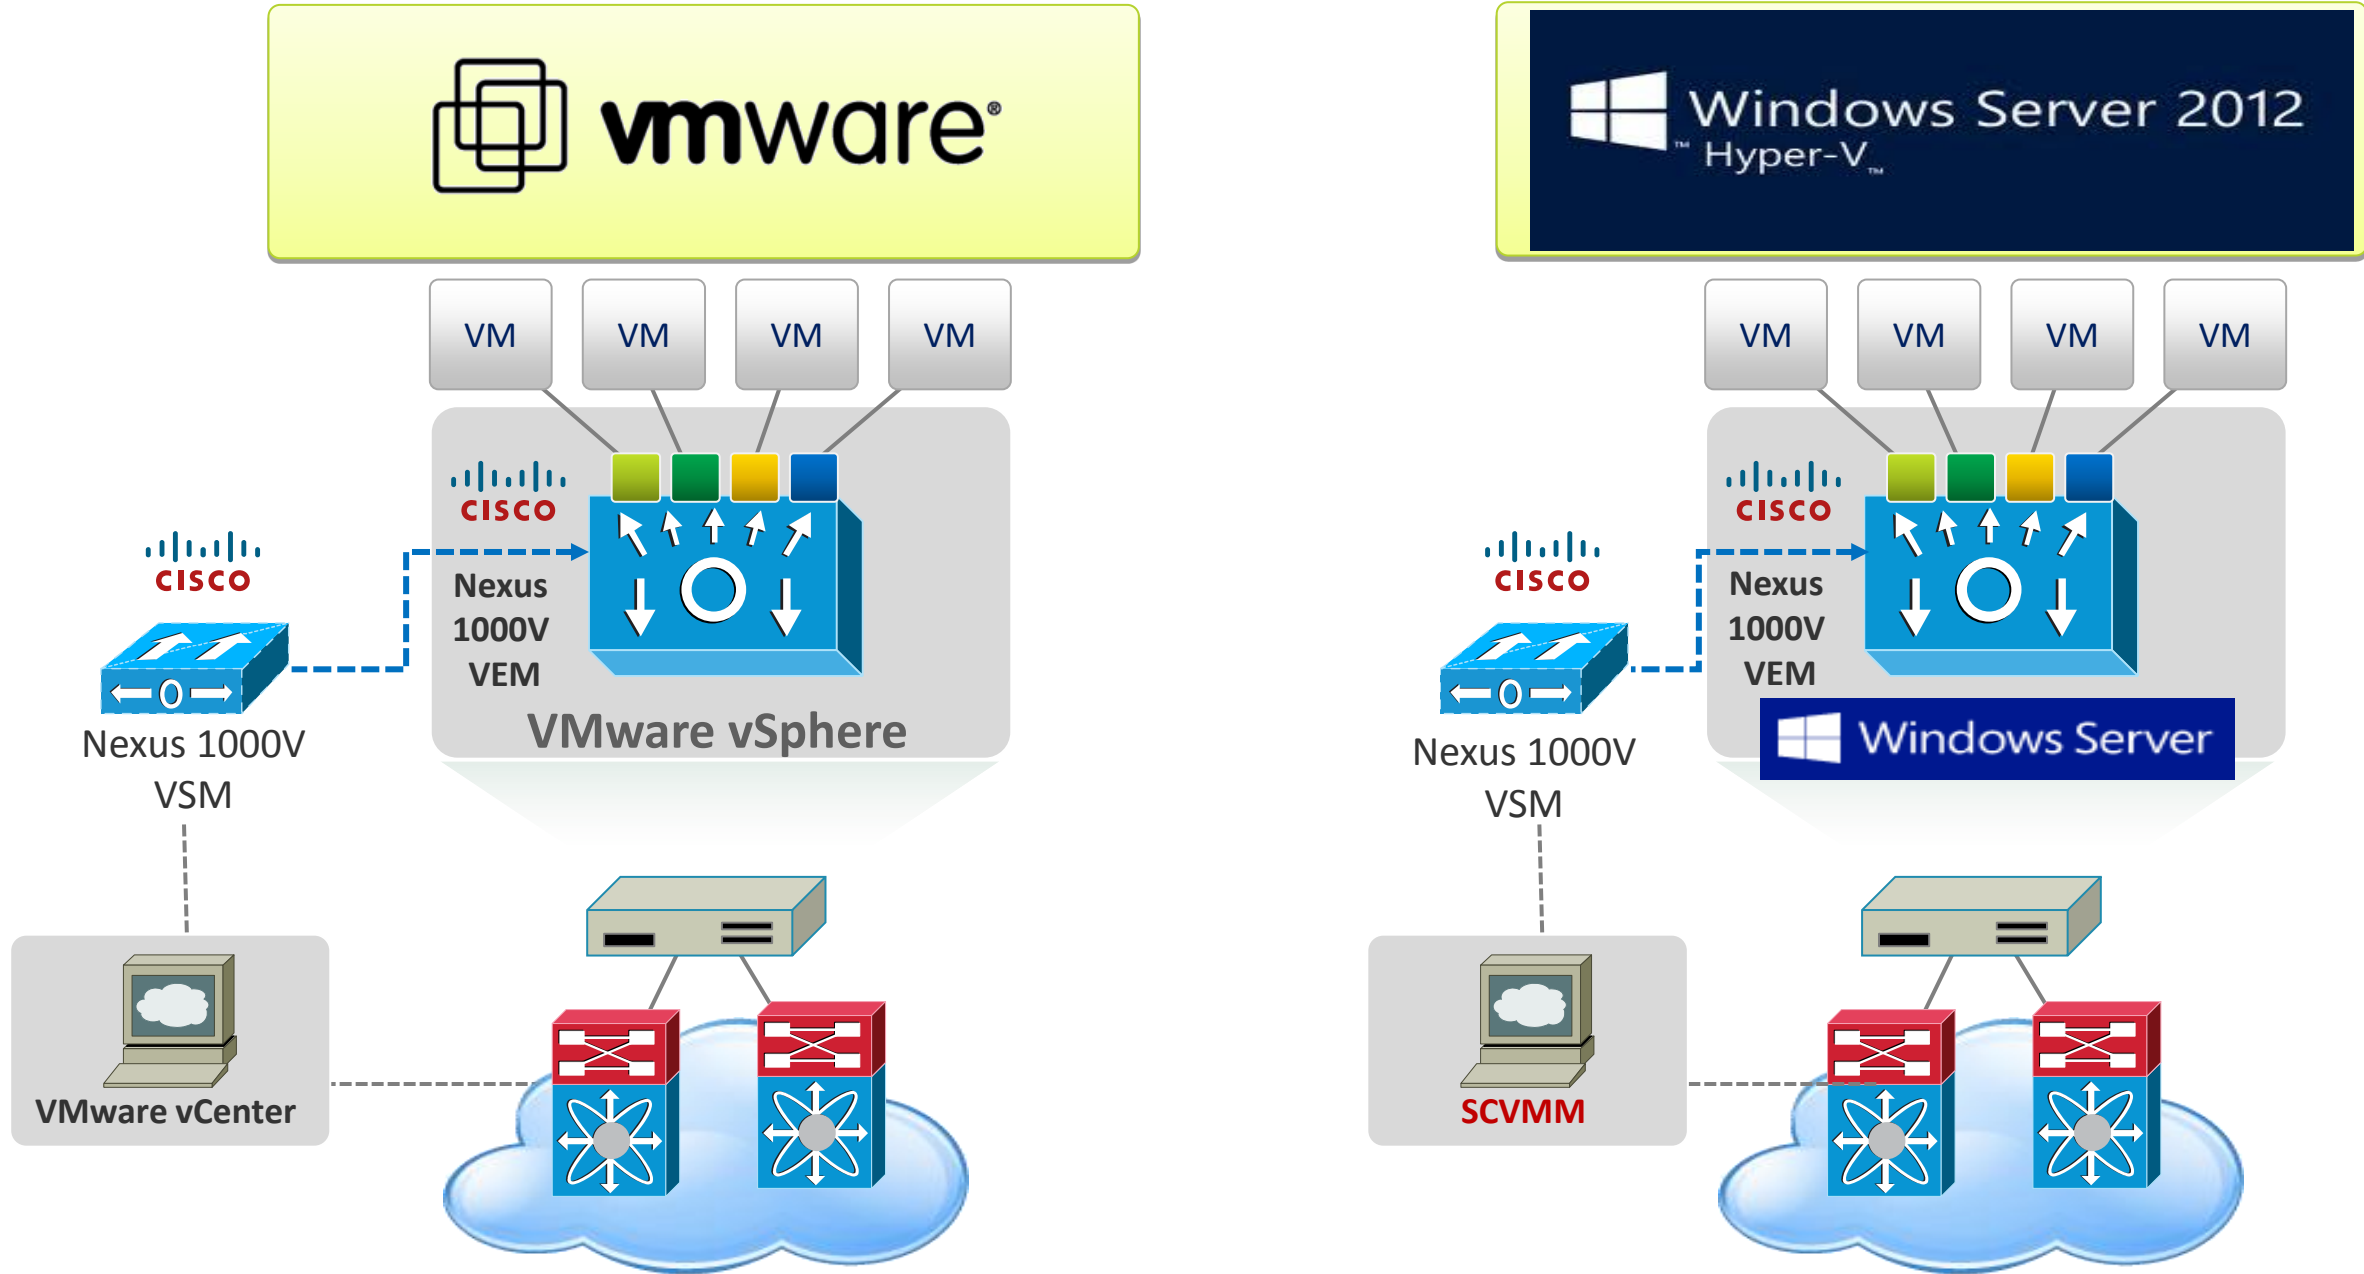
Multi-Cloud

Private, Public, Hybrid, Community

# Cisco Nexus 1000V Architecture



# Cisco Nexus 1000v Multi-hypervisor Support

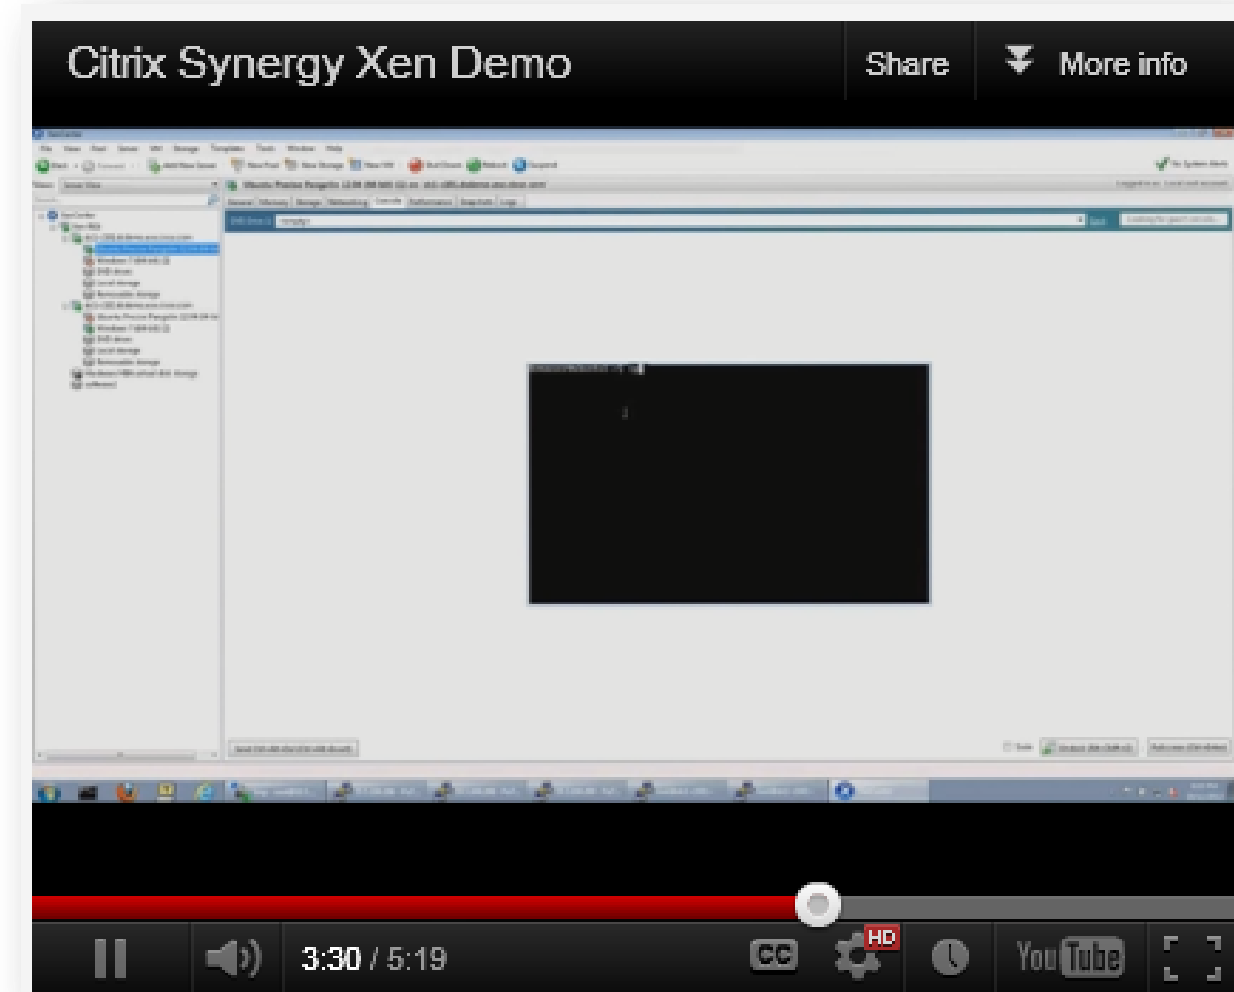
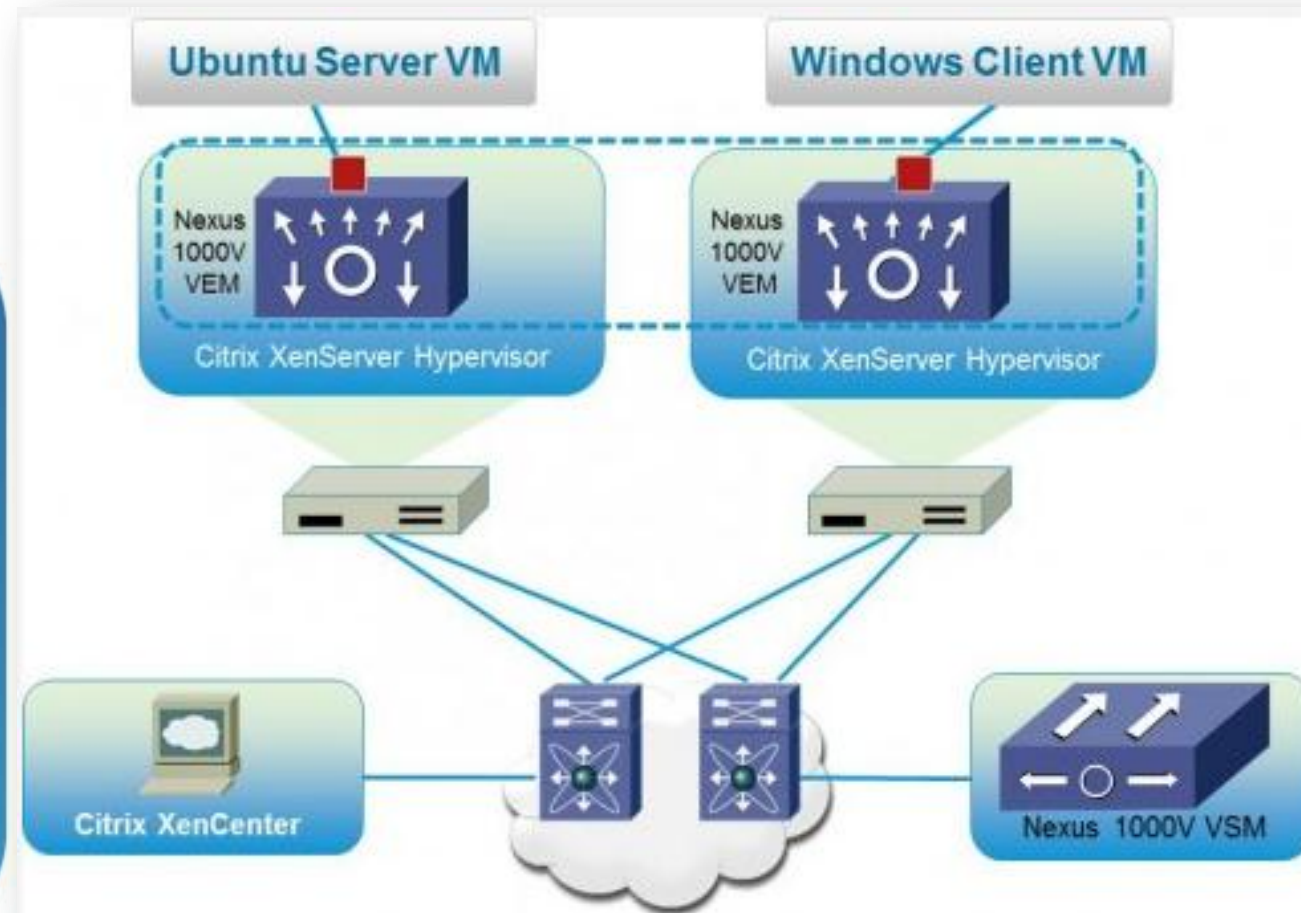


Consistent architecture, feature-set & network services ensures operational transparency across multiple hypervisors.



# Cisco Nexus 1000v on Xen Server

Demo at Citrix Synergy Europe 2012



Cisco Blog > Data Center and Cloud

Citrix Synergy Round-up: Cisco Previews Nexus 1000V on XenServer

October 18, 2012 at 11:00 am PST



Gary Kinghorn



<http://youtu.be/tS8bnbh38ls>

# Cisco Nexus 1000v on KVM

Demo at Cisco Live 2012 in San Diego (for Openstack integration)

Cisco Blog > Open at Cisco

## Nexus 1000V On Open Source Hypervisors



Kyle Mestery | June 20, 2012 at 2:26 pm PST

3

Tweet 22

Like 3

Last week at Cisco Live, Cisco unveiled the Cisco ONE strategy. I won't go into detail on Cisco ONE in this blog post, there has been plenty of blog and analyst coverage of this elsewhere. One piece of the announcement I would like to talk about is the Nexus 1000V and it's move to running on Open Source hypervisors, along with OpenStack Quantum integration.

### Nexus 1000V on KVM With OpenStack: The Cisco Live Demo

At Cisco Live, we demonstrated the Nexus 1000V on KVM with integration into OpenStack. The demo included both the Nexus 1000V Virtual Supervisor Module (VSM), as well as the Virtual Ethernet Module (VEM). The VSM is a virtual machine running Cisco NX-OS software. For the demo, the VSM was running on a Nexus 1010 physical appliance. The VEM was running on the Linux host itself, which was running Fedora Linux, version 16. The OpenStack version we demoed was OpenStack Essex. We were running Nova, Glance, Keystone, Horizon and Quantum. We also wrote a Nexus 1000V Quantum plugin which handles interaction between Quantum and the Nexus 1000V VSM. This is done via a REST API on the Nexus 1000V VSM.

What we demonstrated was the ability for providers to create networks using the standard "nova-manage" CLI in OpenStack. These networks were then mapped to port-profiles on the Nexus 1000V VSM. When a tenant then powered up a VM, the VM was placed on the provider network, and ultimately had it's VIF attached to the port-profile associated with the provider network. The network administrator, through the VSM, is now able to see the virtual interfaces attached to veth ports, and can apply policies on them. We demoed ACLs on the virtual ports, to demonstrate a Nexus 1000V feature in use with OpenStack. What the demo ultimately showed was the Nexus 1000V operational model separating network and server administrator in an OpenStack deployment.



## Openstack at Cisco

### Cisco Edition of OpenStack

Cisco announces availability of a simplified OpenStack installation.



### Related Links

OpenStack is one part of a Cisco's broader strategy to create open, programable networks.

[Learn More](#)

See what's possible with new APIs, data on demand, and smarter, faster networks.

[Meet onePK](#)

Join the Cisco OpenStack team and help shape the industry.

[Contact Us](#)

The Cisco Edition of OpenStack is a packaged and tested version of the OpenStack Folsom release with open source components. It delivers automated deployment and additional functionality such as high availability, monitoring, and Cisco enabled networking through Quantum.

### Installation

Get instructions for all-in-one, multi-node, and high-availability models for Essex and Folsom OpenStack releases.

[Get Instructions](#)

### Documentation

Download administrative, user management, and reference documentation for OpenStack Cisco Edition.

[Learn More](#)

### Open@Cisco

Hear from and engage with the engineers behind the Cisco OpenStack efforts on the open@cisco blog.

[Read Now](#)

# New Nexus 1000V Freemium Go-To-Market Model

No-Cost Version

## Nexus 1000V Essential Edition

The world's most advanced virtual switch

- Full Layer-2 Feature Set
- Security, QoS Policies
- VXLAN virtual overlays
- Full monitoring and management capabilities
- vPath enabled Virtual Services

\$695 per CPU MSRP

## Nexus 1000V Advanced Edition

Adds Cisco value-add features for DC and Cloud

- All Feature of Essential Edition
- VSG firewall bundled (previously sold separately)
- Support for Cisco TrustSec SGA policies
- Platform for other Cisco DC Extensions in the Future

**Freemium Pricing Model Offers Flexibility for Customers to Deploy Cisco Virtual Data Centre**

# Essential and Advanced Edition - Features Comparison

**Fully Available !!!**

## Essential Edition (\$0)


- VLAN, ACL, QoS
- VXLAN, vPath
- LACP
- Multicast
- Netflow, ERSPAN
- Management
- vTracker
- vCenter Plug-in

## Advanced Edition (\$695 LIST)

- Cisco TrustSec support
- CISF: DHCP snooping, IP Source Guard, ARP Inspection
- **VSG (previously \$495 List)**

## Essential Edition

- VLAN, ACL, QoS
- VXLAN, vPath
- LACP
- Multicast
- Netflow, ERSPAN
- Management
- vTracker
- vCenter Plug-in



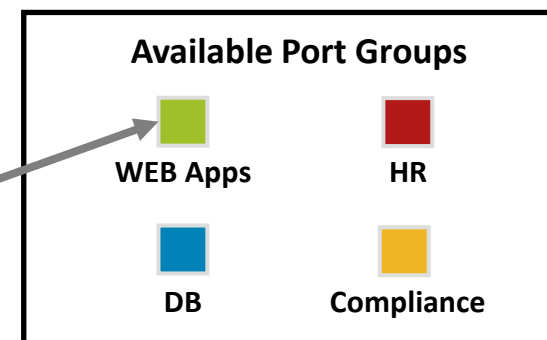
# Port Profiles “How to”

1. Nexus 1000V automatically enables port groups in vCenter via API
2. Server Admin uses vCenter to assign vnic policy from available port groups
3. Nexus 1000V automatically enables VM connectivity at VM power-on

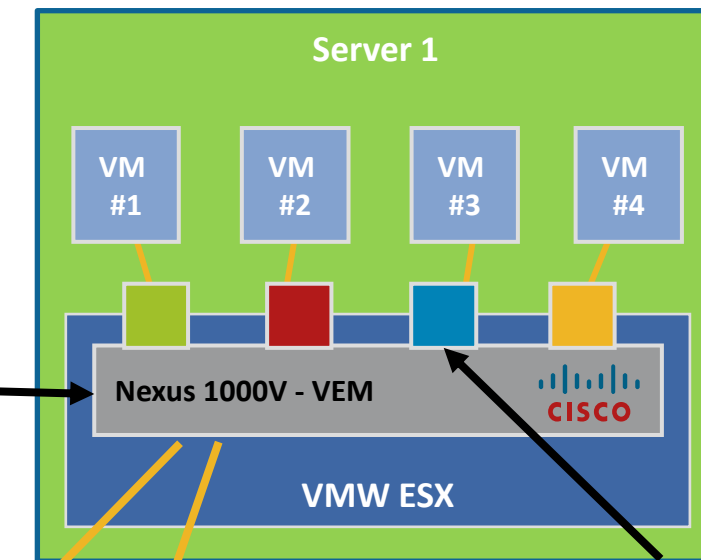
**“WEB Apps” Port Profile:**

- PVLAN 108, Isolated
- Security Policy = Port 80 and 443
- Rate Limit = 100 Mbps
- QoS Priority = Medium
- Remote Port Mirror = Yes


1. Virtual Machine Manager (VMM)




2.



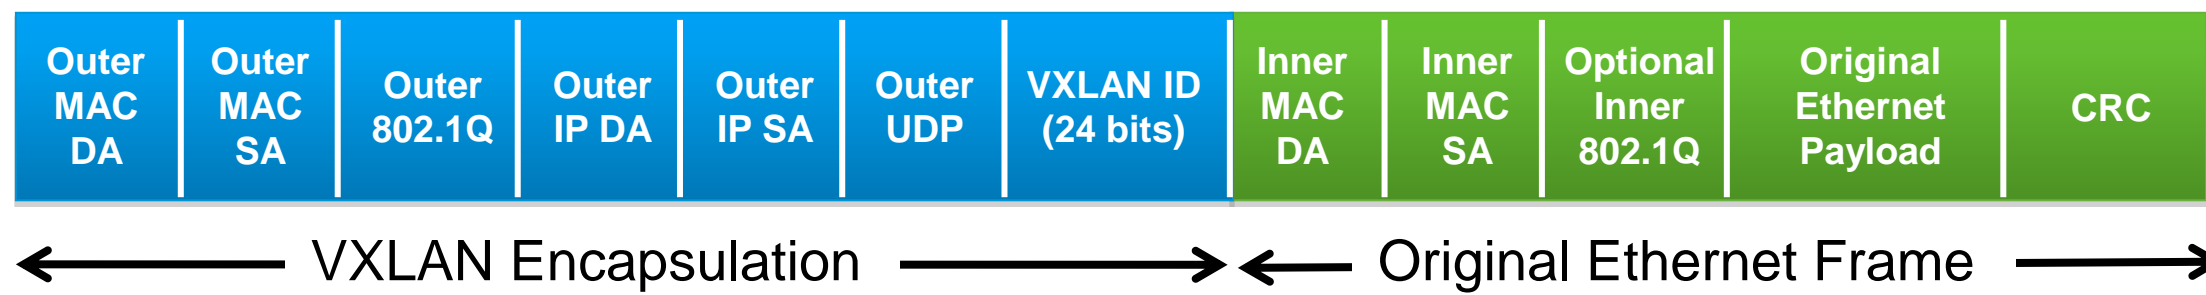
3.






# Virtual Extensible Local Area Network (VXLAN)

- Ethernet in IP overlay network
  - Entire L2 frame encapsulated in UDP
  - 50 bytes of overhead
- Include 24 bit VXLAN Identifier
  - 16 M logical networks
- VXLAN can cross Layer 3
- Tunnel between VEMs
  - VMs do NOT see VXLAN ID
- IP multicast used for L2 broadcast/multicast, unknown unicast
- Technology submitted to IETF for standardisation
  - With VMware, Citrix, Red Hat and Others

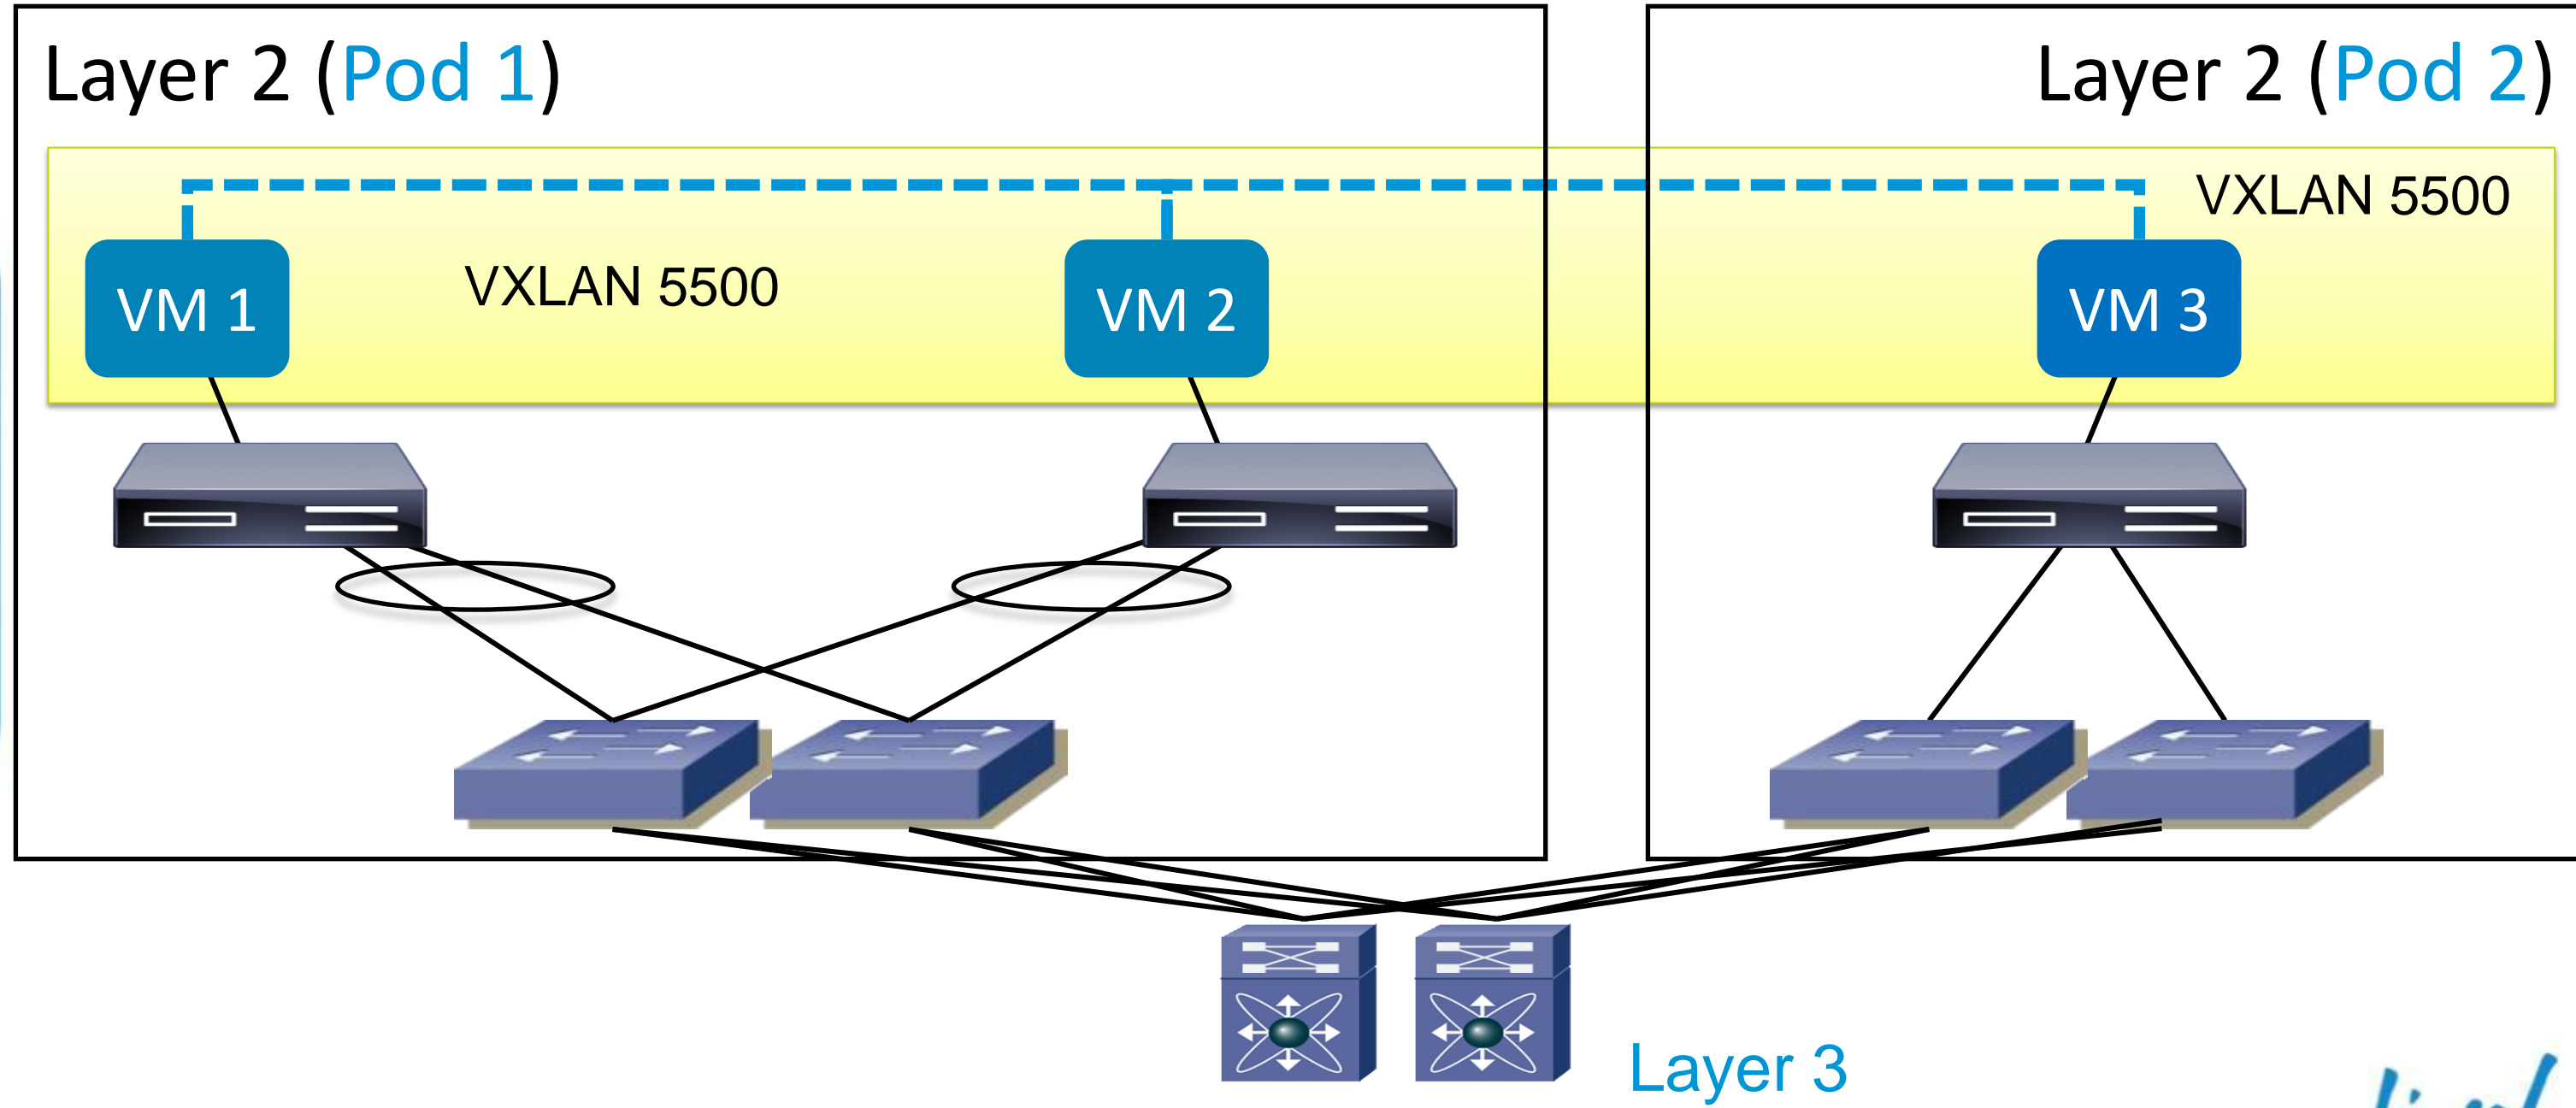


# Existing Solution: Reachability of VMs Within VLAN

Limited Scalability with 4k VLANs



# VXLAN: Crossing L3 Subnets within the DC

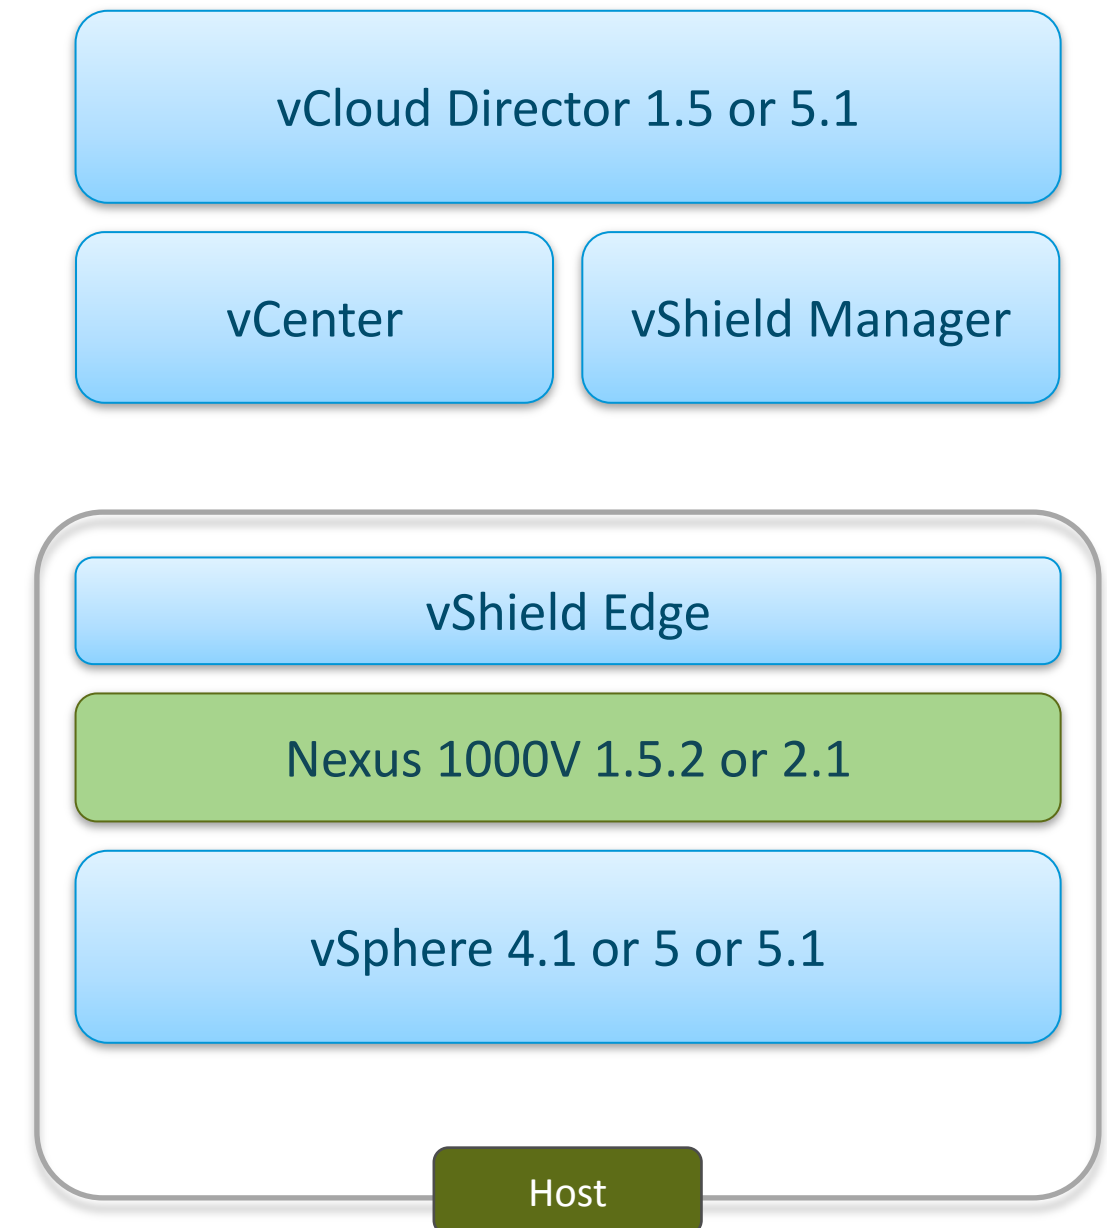




# Nexus 1000V @ ESX: vCloud Director Integration

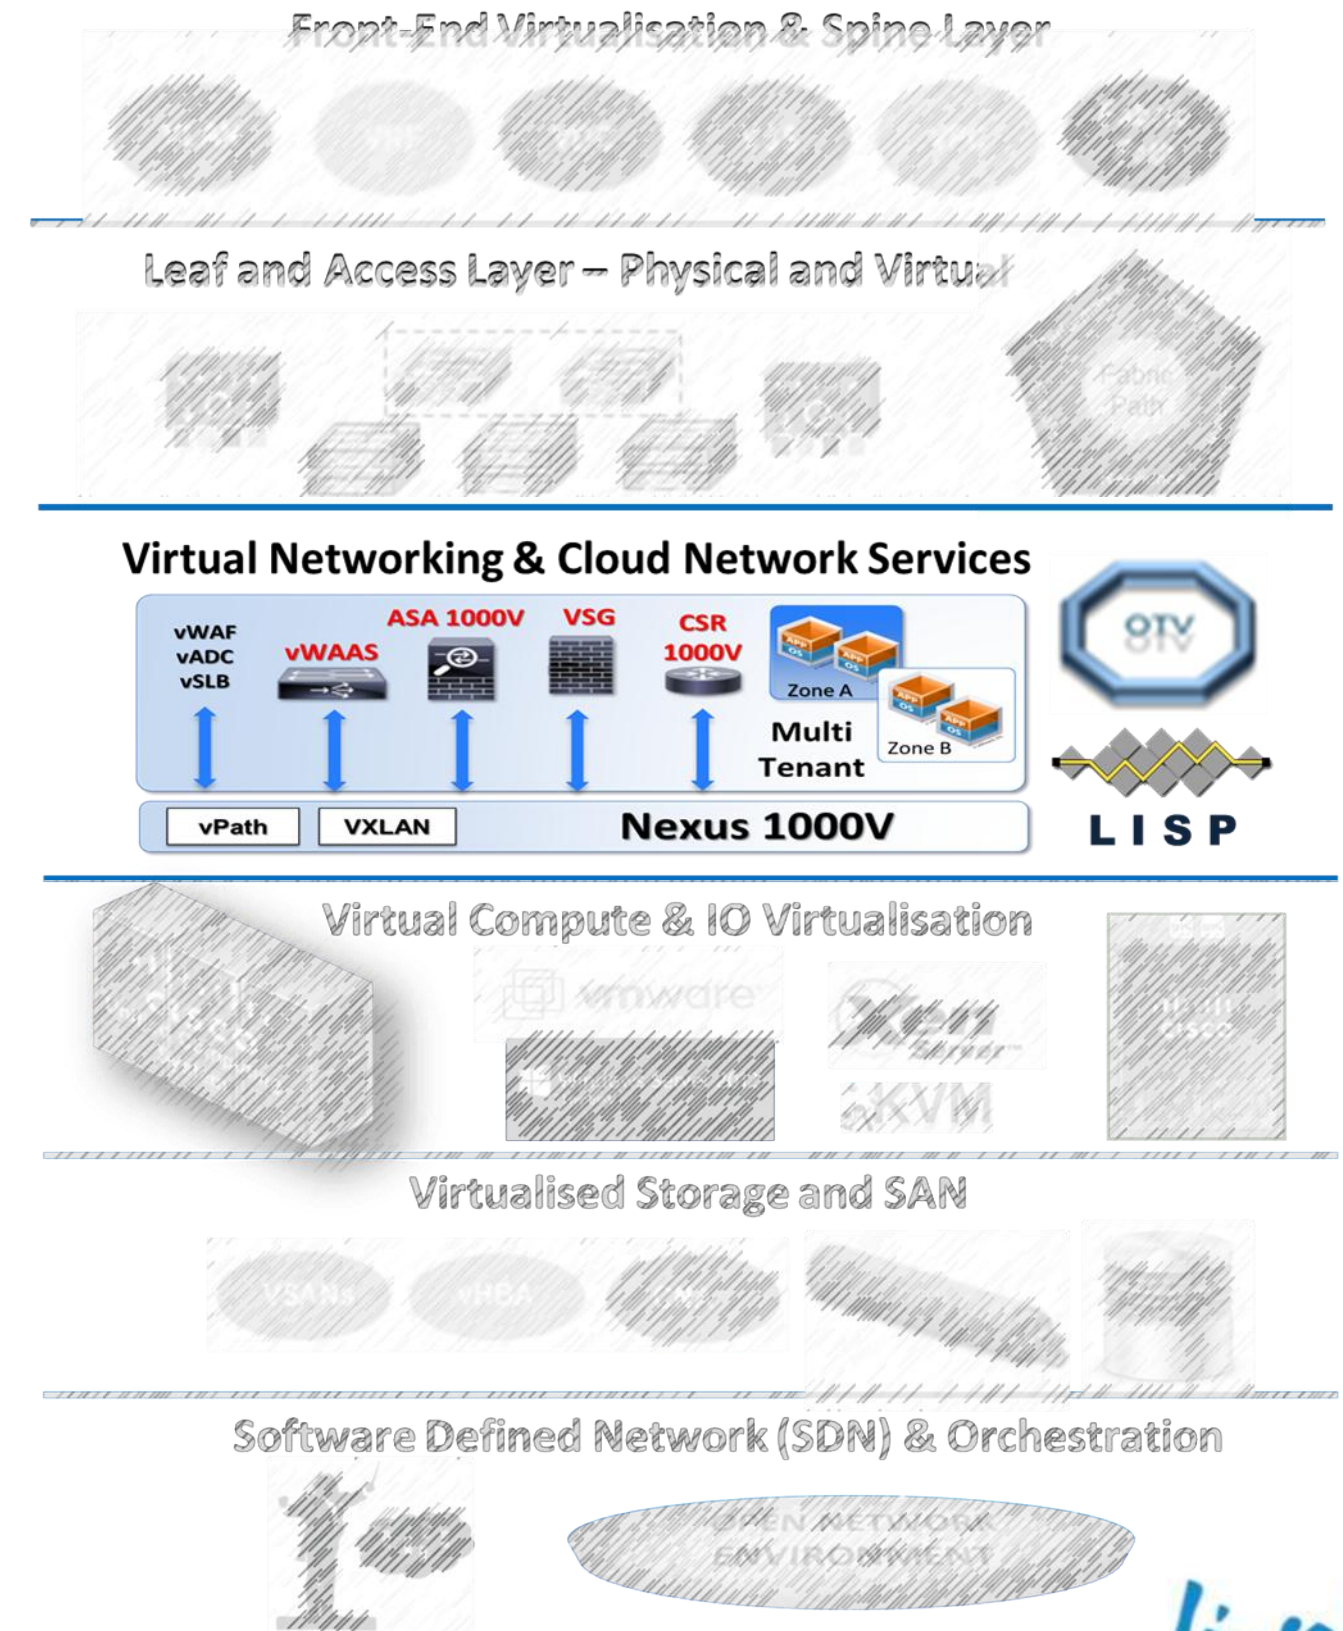
- N1KV to fully integrate since vCD 1.5
- Support dynamic network provisioning
  - Port-group backed pools
  - VLAN-backed pools
  - Network isolation backed pools (via VXLAN)
- Choice of vSphere 4.1 or 5 or 5.1

**Nexus 1000V & vCD**

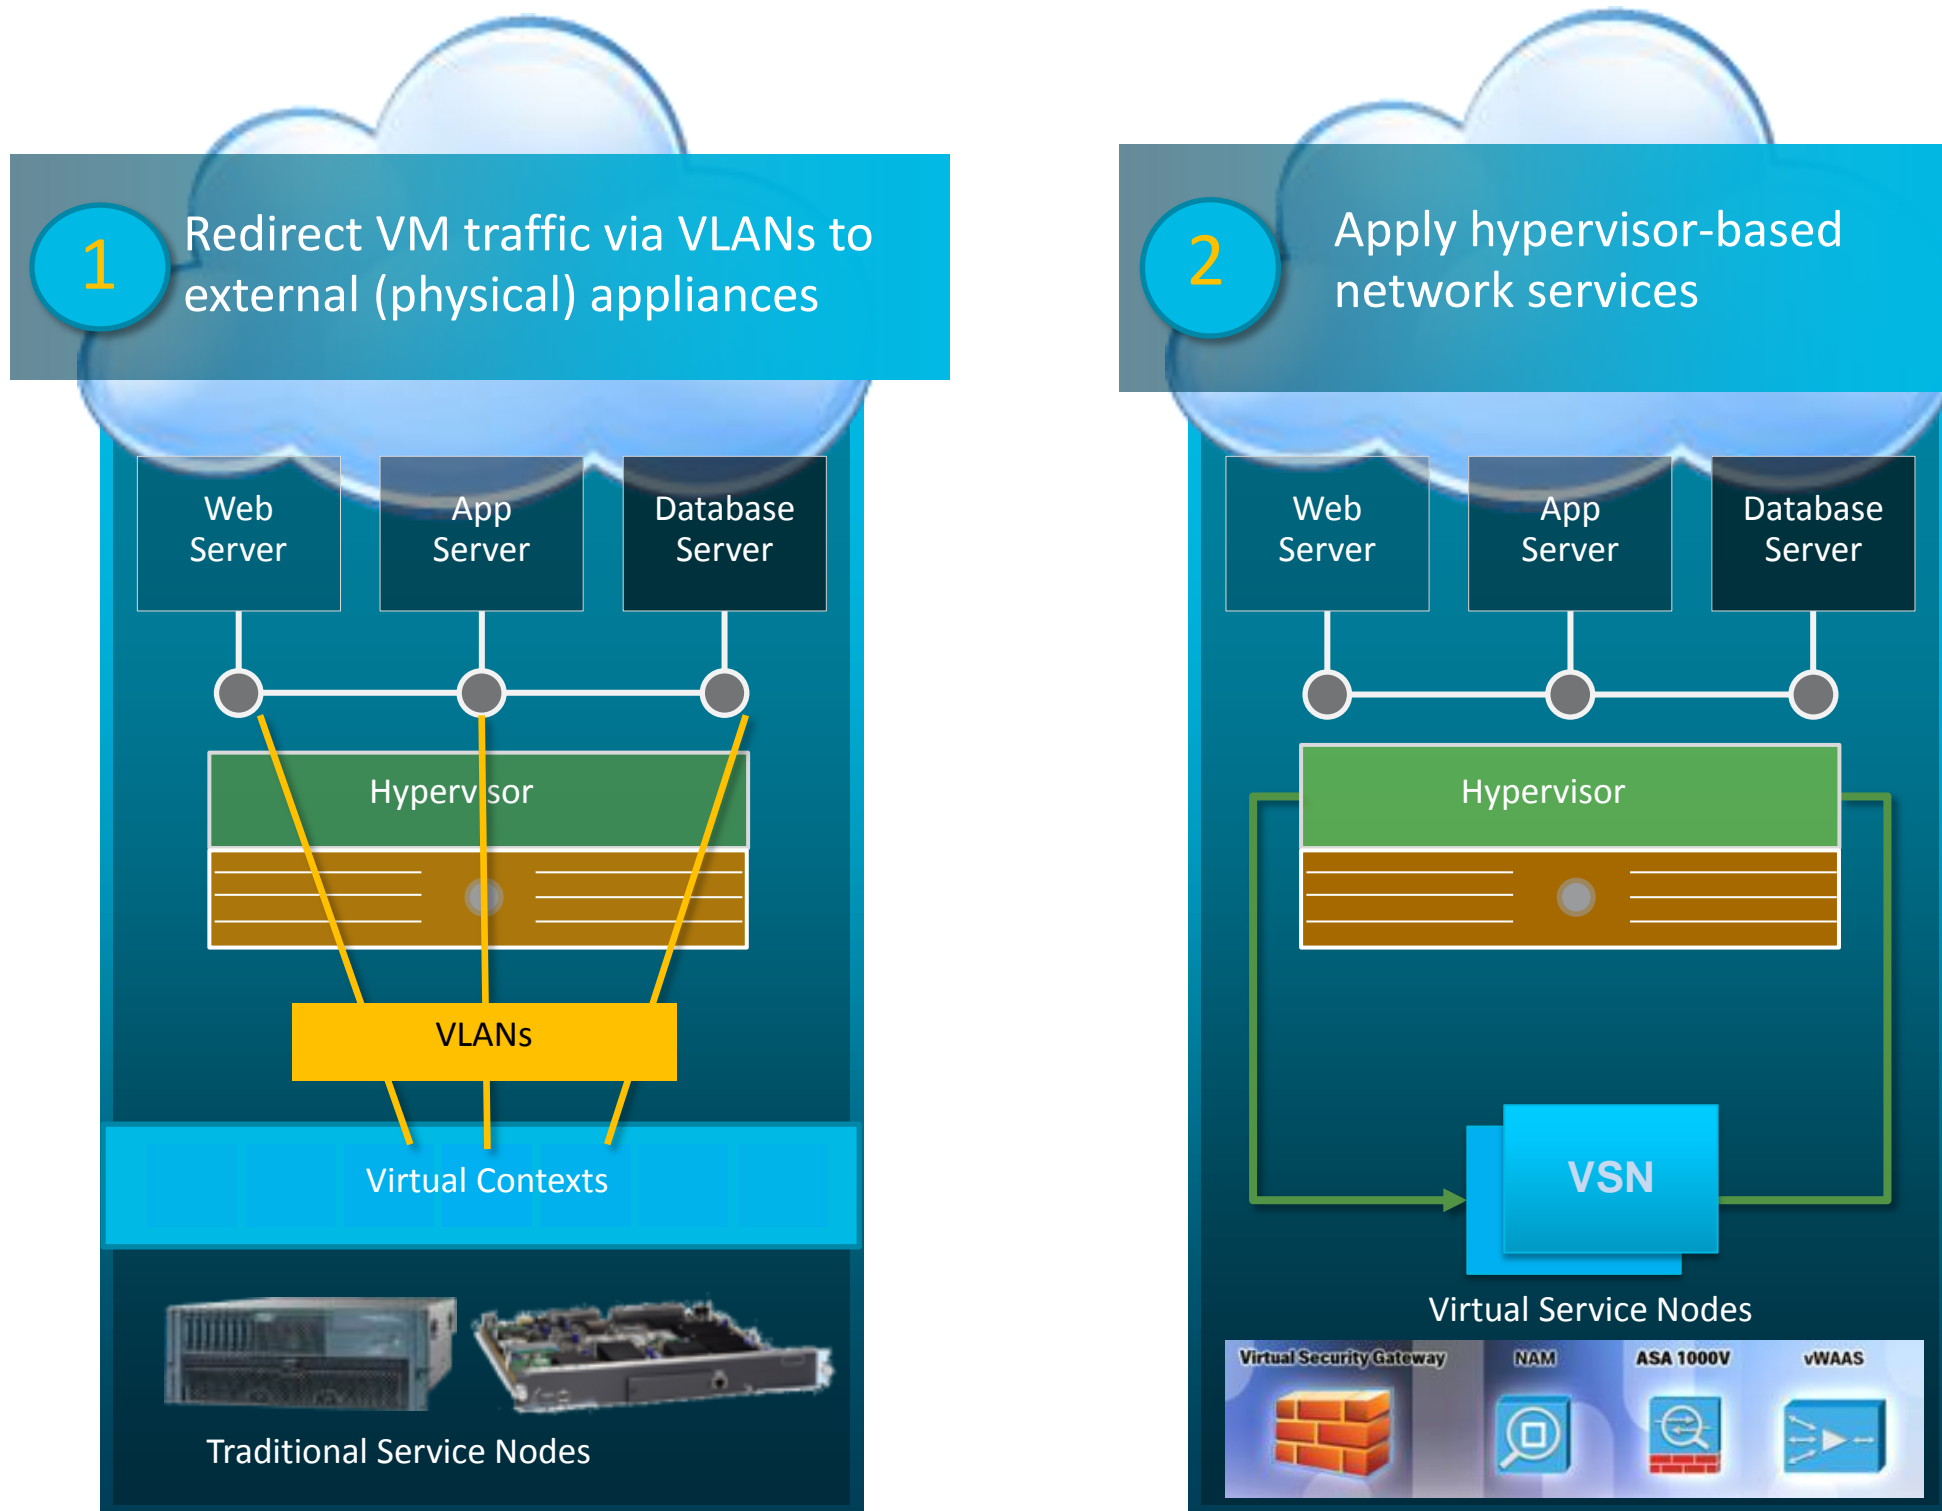


# Agenda

- **Data Centre Virtualisation Overview**
- **Front-End Data Centre Virtualisation**
  - DC Edge Layer
  - Aggregation / Spine Layer
  - Access / Leaf Layer
- **Virtual Networking & Cloud Network Services**
  - Virtual Access Layer: Nexus 1000v
  - **Cloud Network Services**
  - **Inter DC Services: OTV & LISP**
- **Virtual Compute & IO Virtualisation**
  - Unified Computing System
  - Hypervisors & Virtual IO @ Cisco UCS
- **Virtualised Storage & SAN**
  - Unified Fabric
  - Virtualised SAN and Storage
- **Software Defined Network & Orchestration**
  - SDN & Cisco Open Network Environment (ONE)
  - Cloud Orchestration
- **Implementation Examples**
- **Q&A**

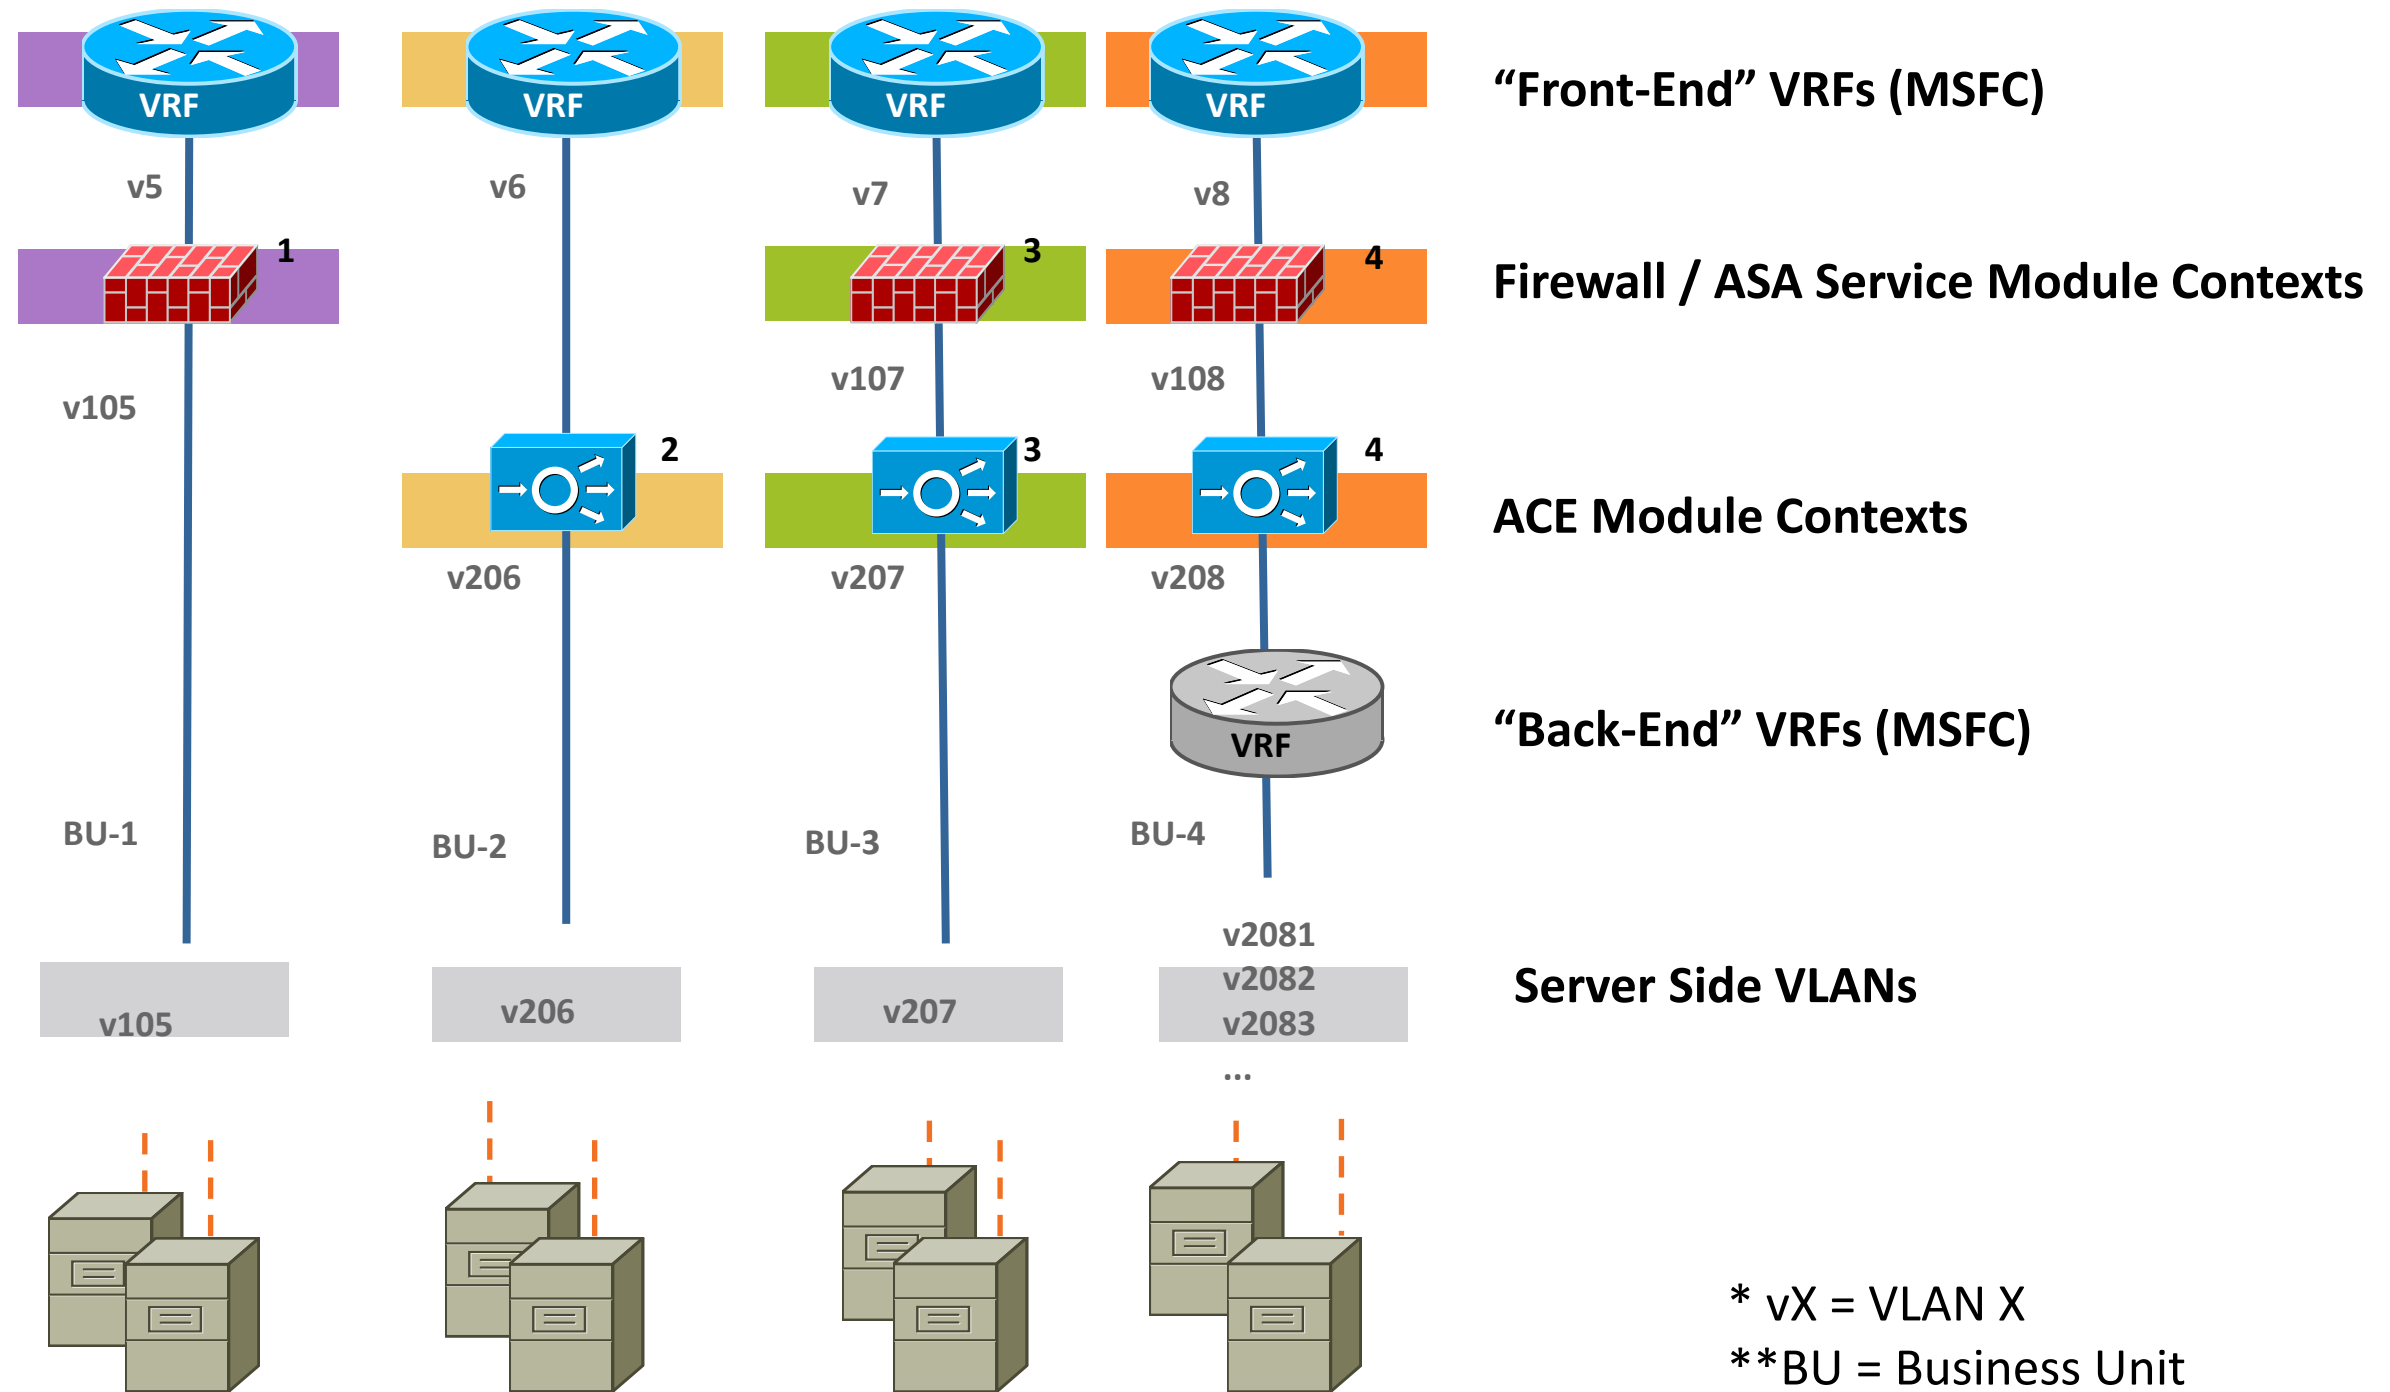


# Deployment Options for Virtual Network Services

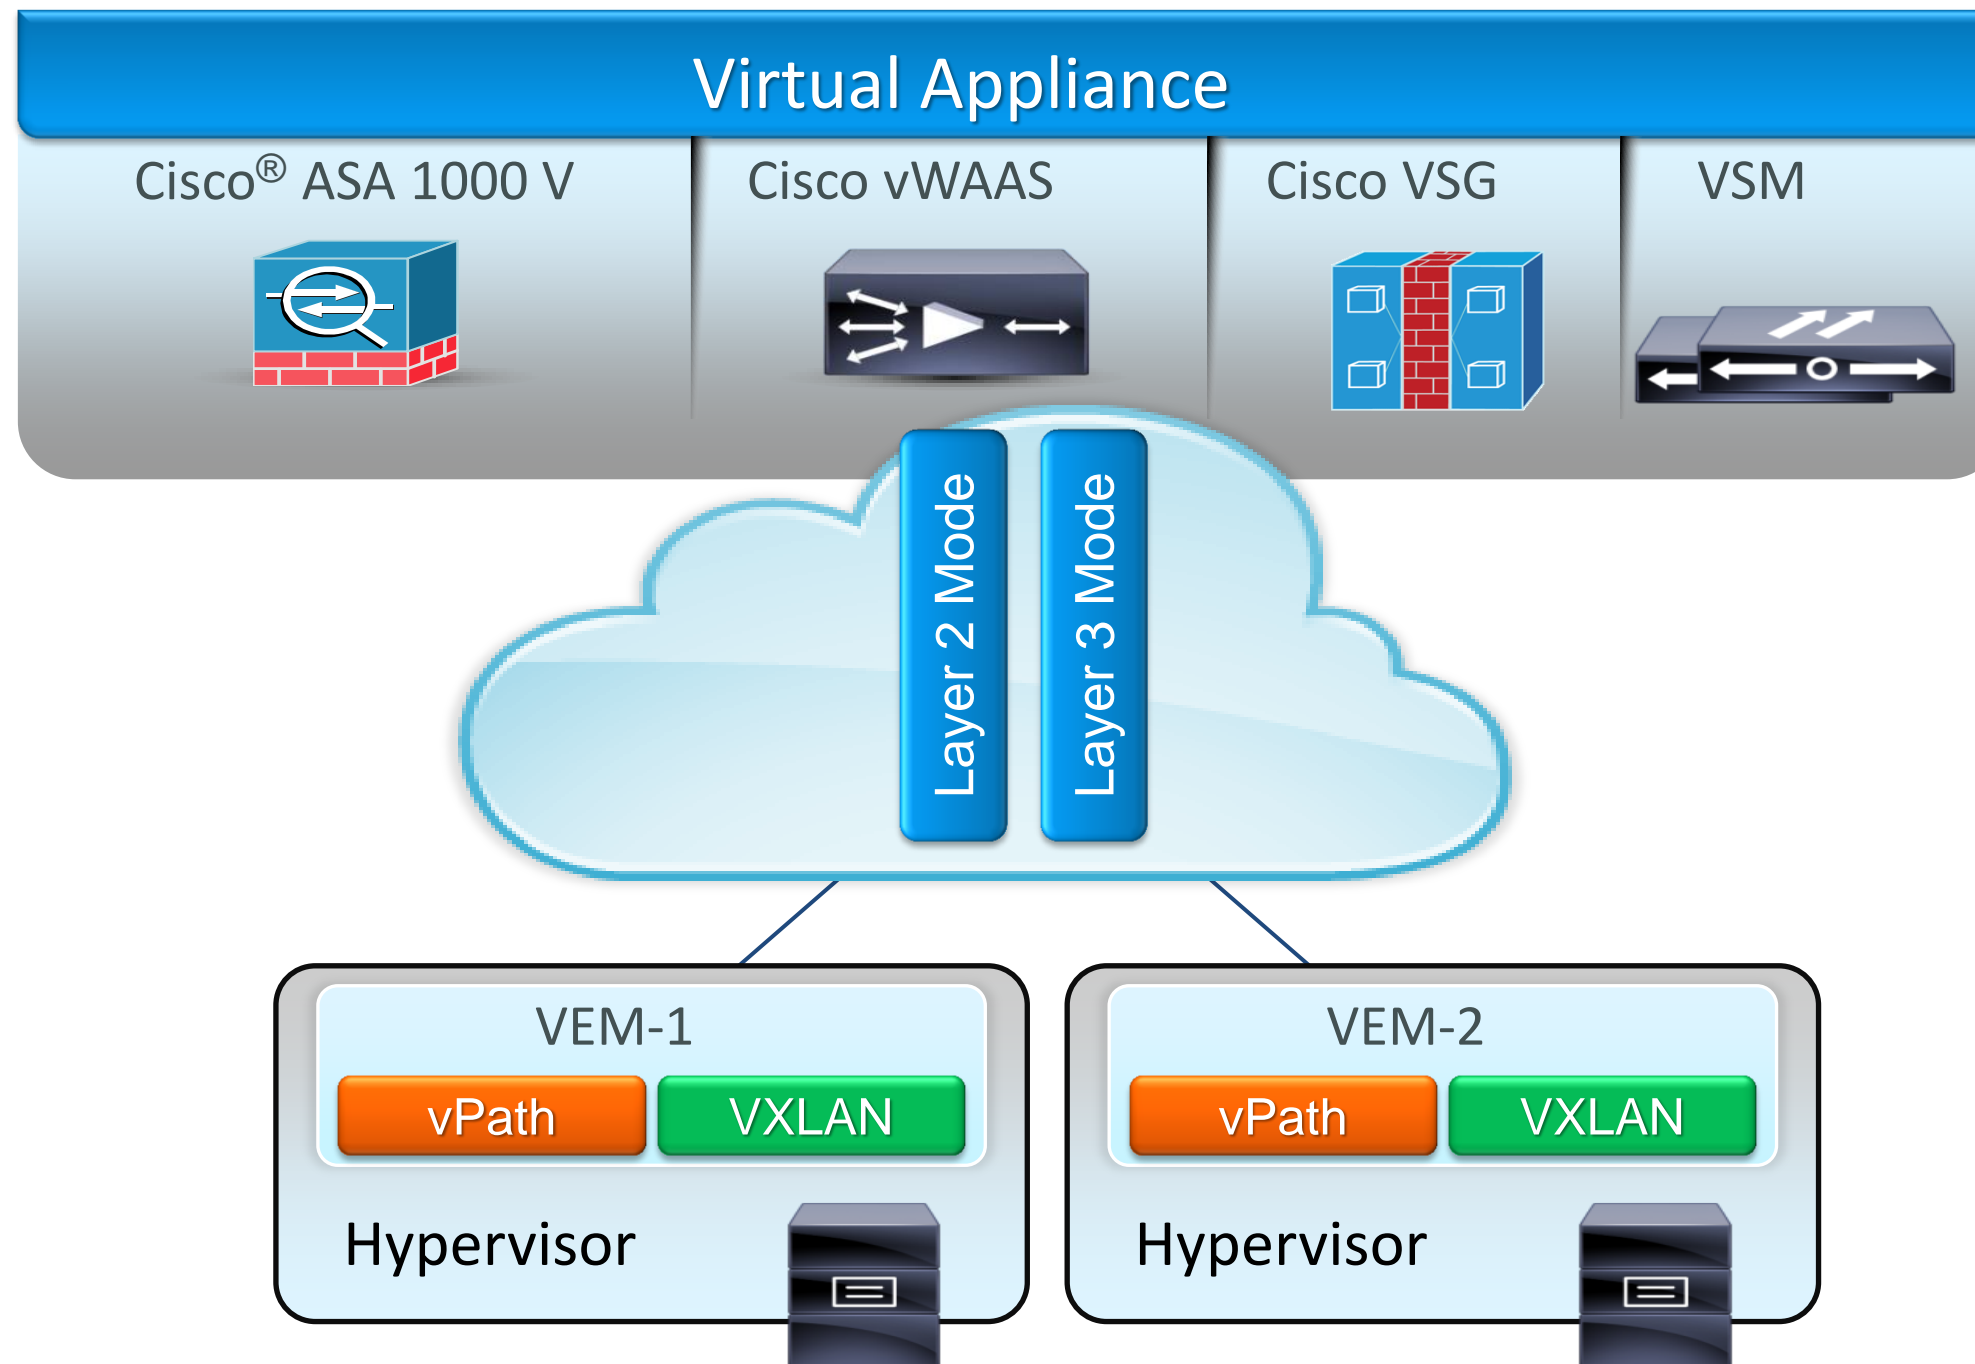


# Data Centre Virtualised Services via Contexts

## Physical Appliances/Modules Context Combination Example



# Cisco Nexus 1000V Architecture: vPath




### Virtual Service Data Path (vPath)


Embedding intelligence for virtual services

- Service chaining (traffic steering)
- Fast-path offload
- VXLAN aware


# vPATH Interception in Nexus 1000v




- vPATH Interception is configured on Server VM's Port Profile in both directions to redirect to a VSN
- Server traffic is intercepted by vPATH interception in VEM and redirected to a Virtual Service Node - VSN (or multiple via service chaining on vPATH 2.0)
- VSN egress traffic forwarded without further vPATH interception.



# vPATH – Application Based Interception



Network Admin view



Nexus 1000v VSM

```

port-profile type vethernet Opt-Exchange-Server
  vmware port-group
  switchport mode access
  switchport access vlan 3185
  vn-service ip-address 2.8.2.90 vlan 3002 mgmt-ip-address 2.8.2.90 fail open
  no shutdown
  state enabled
    
```

vPATH interception

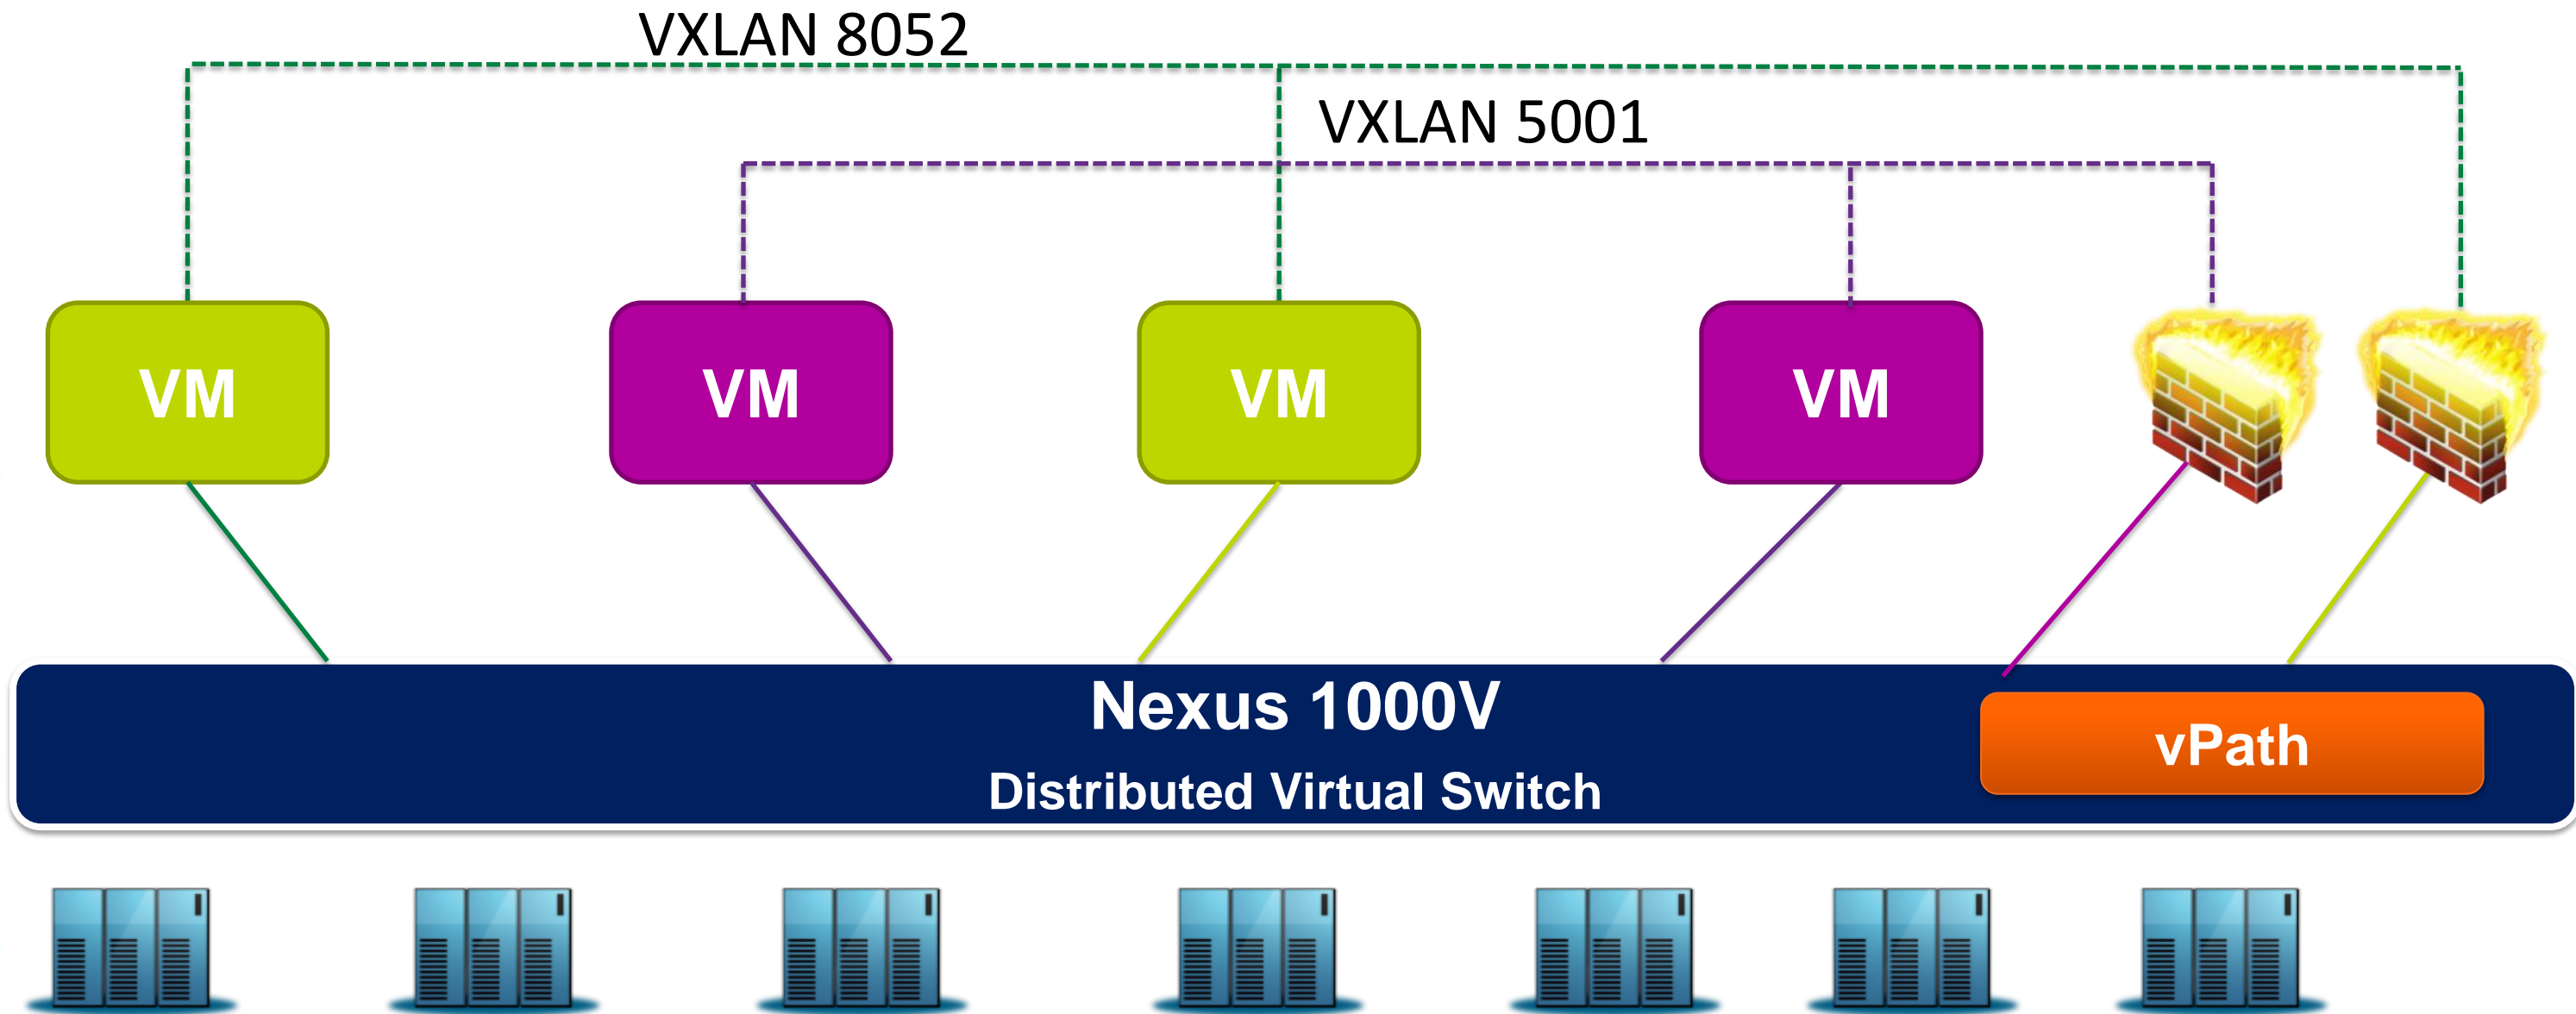
Server Admin view



vSphere client

Attach Opt-port-profile to server VMs

# Nexus 1000V vPath2.0: VSNs on VXLANs

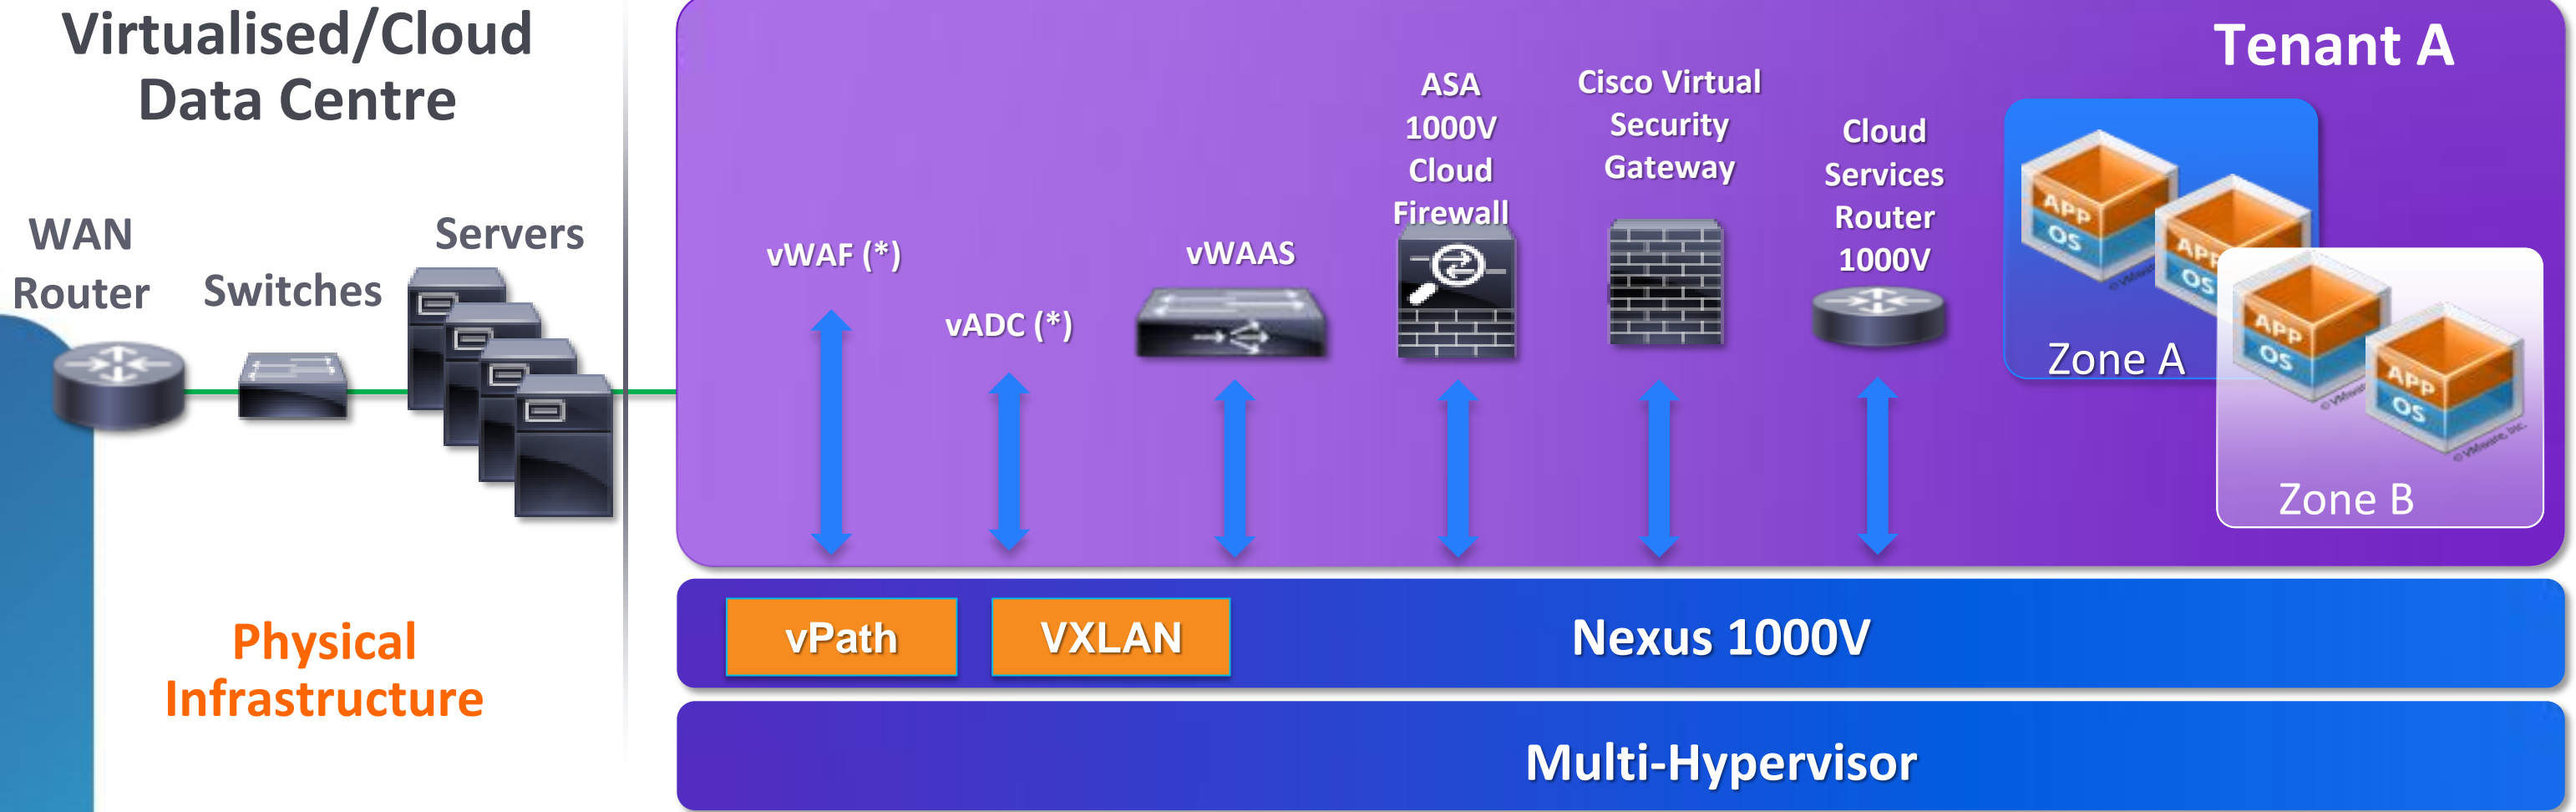


- Deployment- VMs and Virtual Service Nodes, ASA 1000V, VSG, vWAAS, etc. on VXLANs
- Same VSG can protect VMs on multiple VXLANs with overlapping IP addresses



# Cisco Virtual Networking and Cloud Network Services

## Cloud Network Services



### Nexus 1000V

- Distributed switch
- NX-OS consistency

**6000+ Customers**

### VSG

- VM-level controls
- Zone-based FW

**Shipping**

### ASA 1000V

- Edge firewall, VPN
- Protocol Inspection

**Shipping**

### vWAAS

- WAN optimisation
- Application traffic

**Shipping**

### CSR 1000V (Cloud Router)

- WAN L3 gateway
- Routing and VPN

**Beta**

### Ecosystem Services (\*)

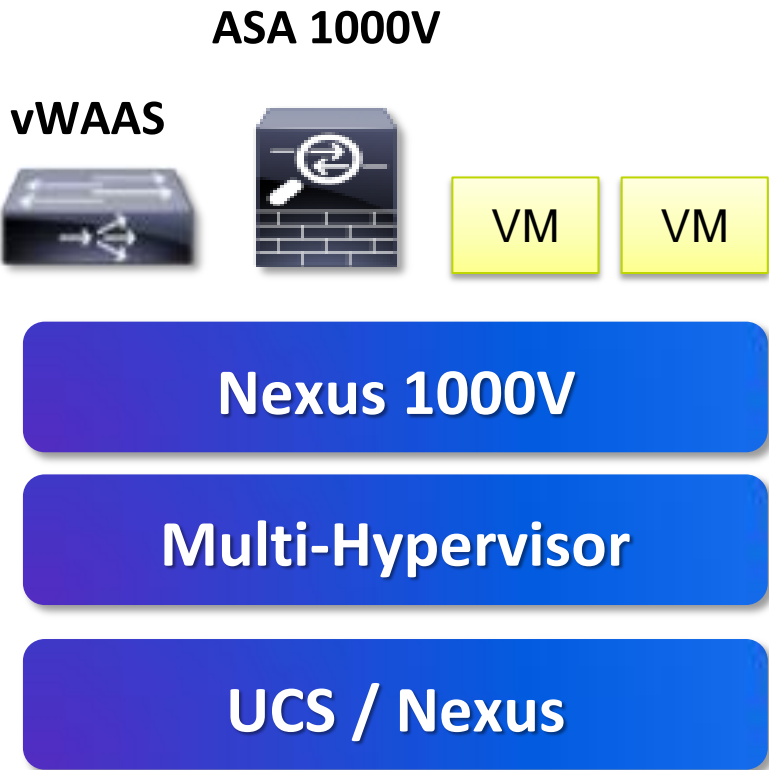
- Virtual ADC
- Web App. Firewall

**FUTURE**


**Cisco live!**

# Cisco Nexus 1110 Virtual Services Appliance

Enabling Physical-Virtual Consistency across Cisco and Ecosystem Partners Products



Virtual Services as VM's on Multiple Hypervisors



Virtual Services Deployed on Cisco Virtual Service Appliance

# vWAAS with vPATH


## Elastic vWAAS Deployment

### Feature

Automatic application of vWAAS service when a new 'Web Server' VM gets provisioned  
vWAAS services associated with 'Web server' VMs using Nexus 1000V policies.

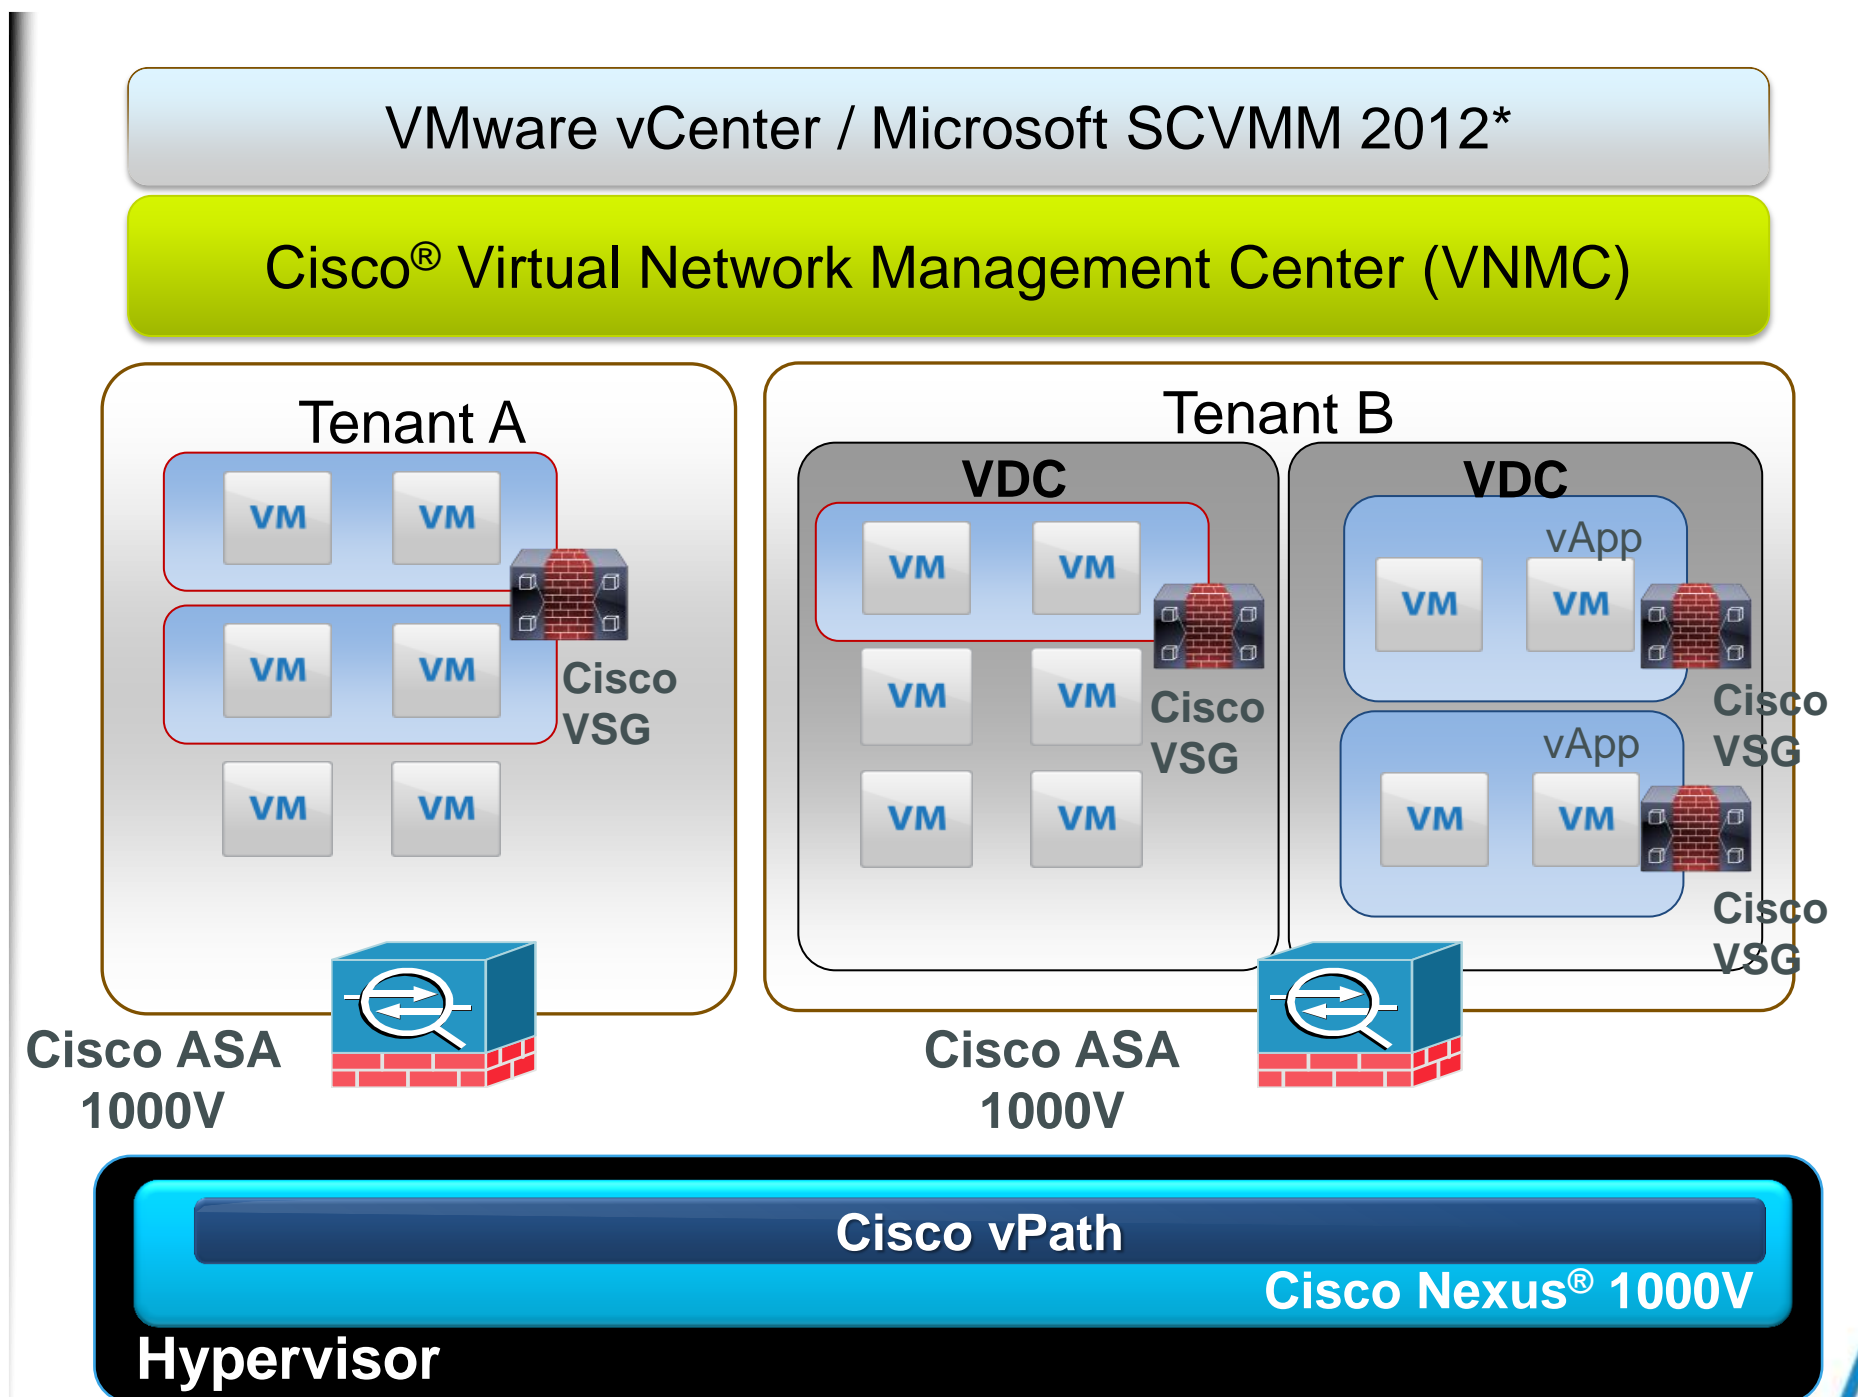
### Benefit

Elastic vWAAS deployment  
Scale-out Virtual Web Server farm by provisioning additional VMs while applying WAN optimisation



# Cisco Virtual Security Solution

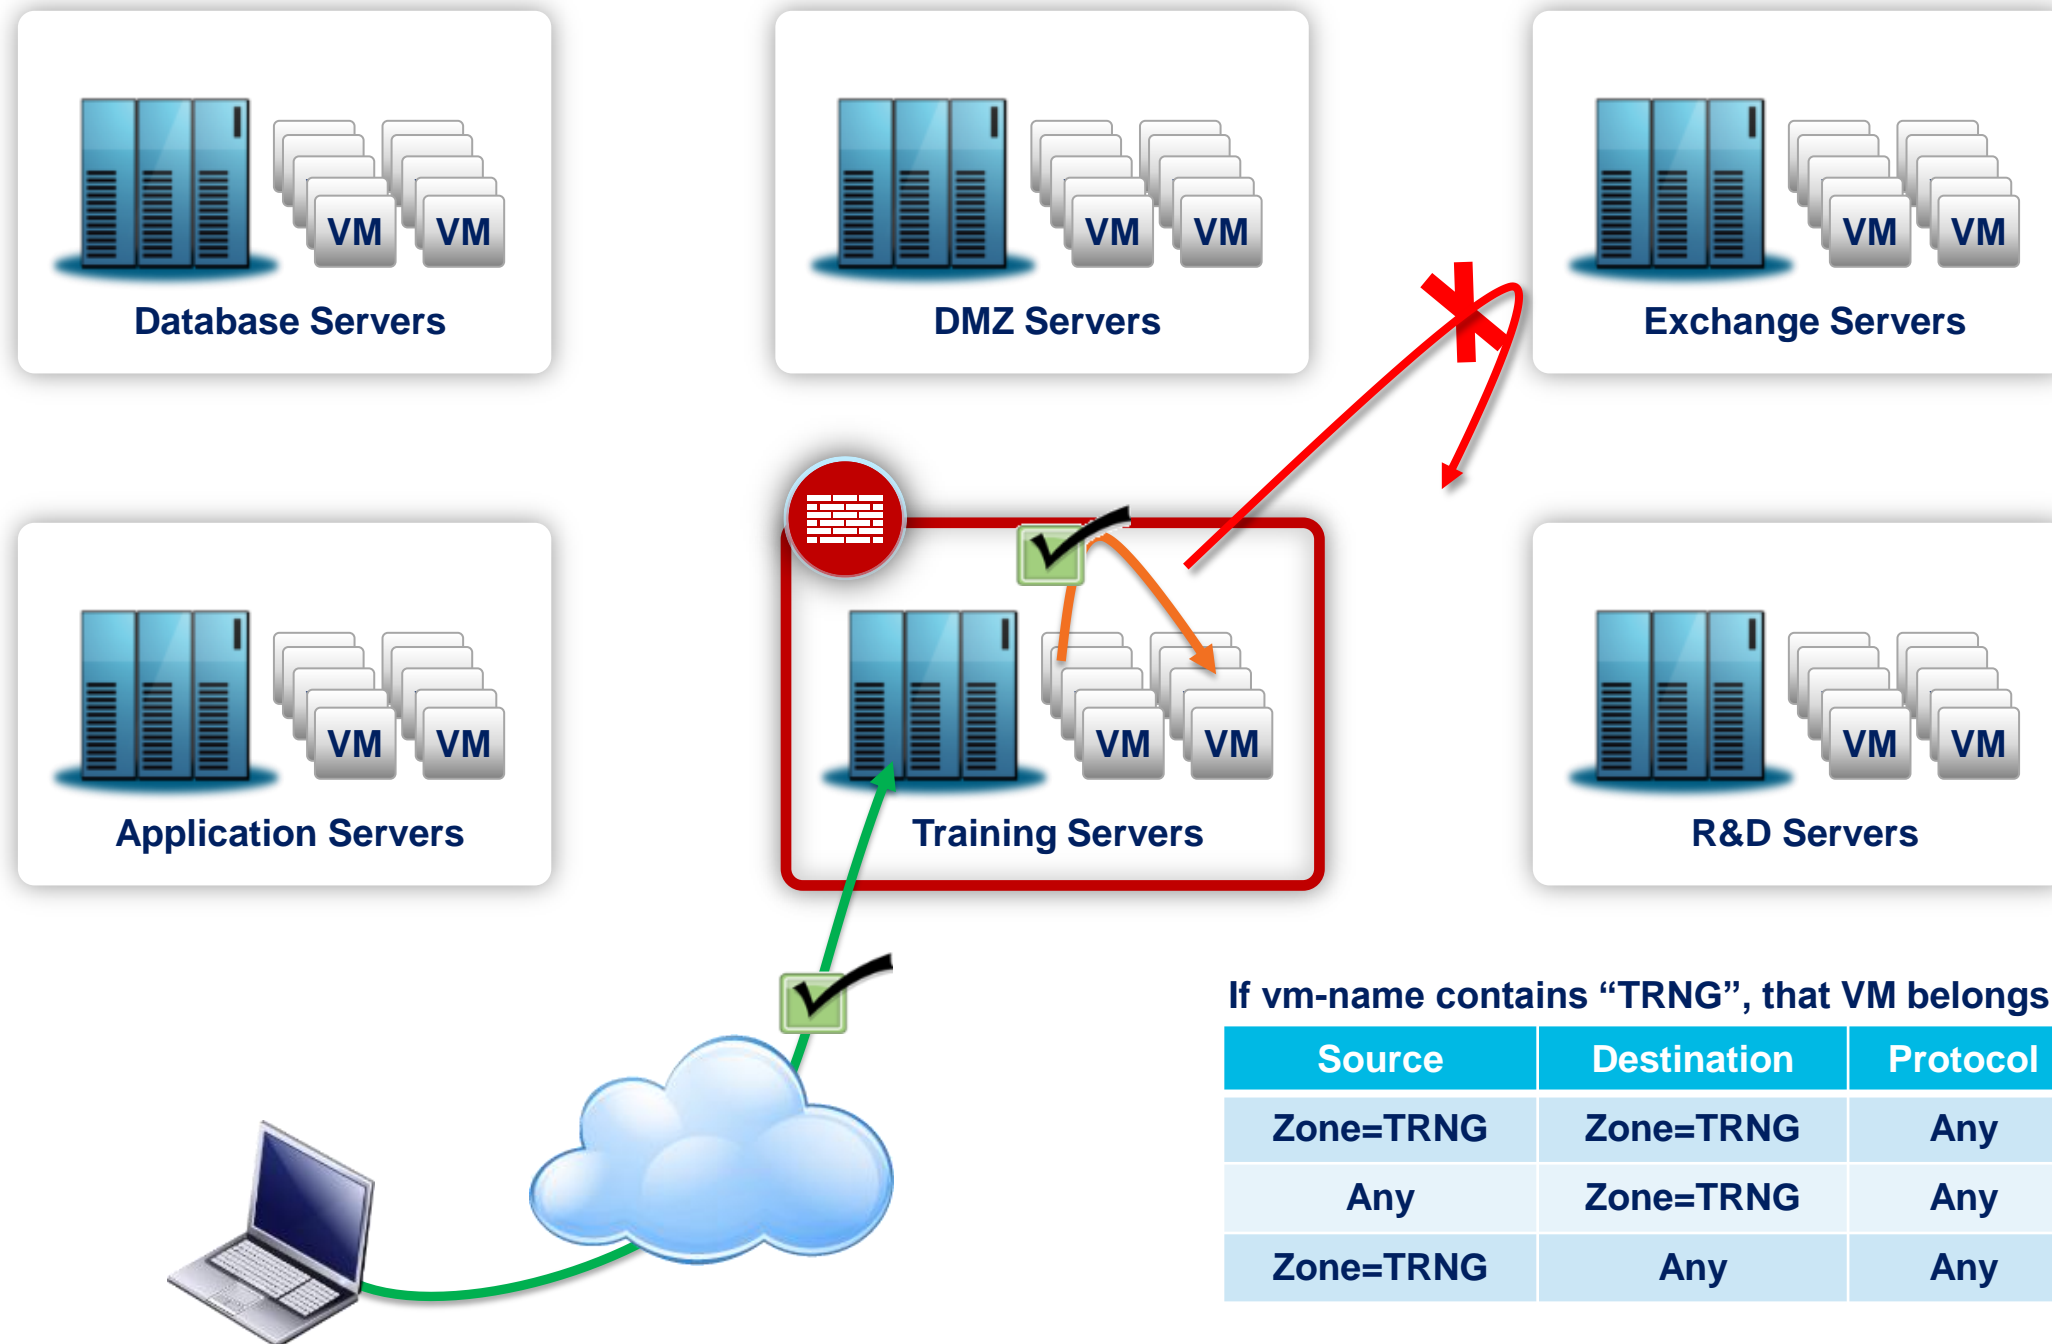
- Proven Cisco® security: virtualised physical and virtual consistency
- Collaborative security model
  - Cisco Virtual Secure Gateway (VSG) for intra-tenant secure zones
  - Cisco ASA 1000V for tenant edge controls
- Transparent integration
  - With Cisco Nexus® 1000V Switch and Cisco vPath
- Scale flexibility to meet cloud demand
  - Multi-instance deployment for scale-out deployment across the data centre



\* Planned for future

CiscoLive!

# VSG Deployment with Logical Zoning, vMotion Support & Scalable Solution



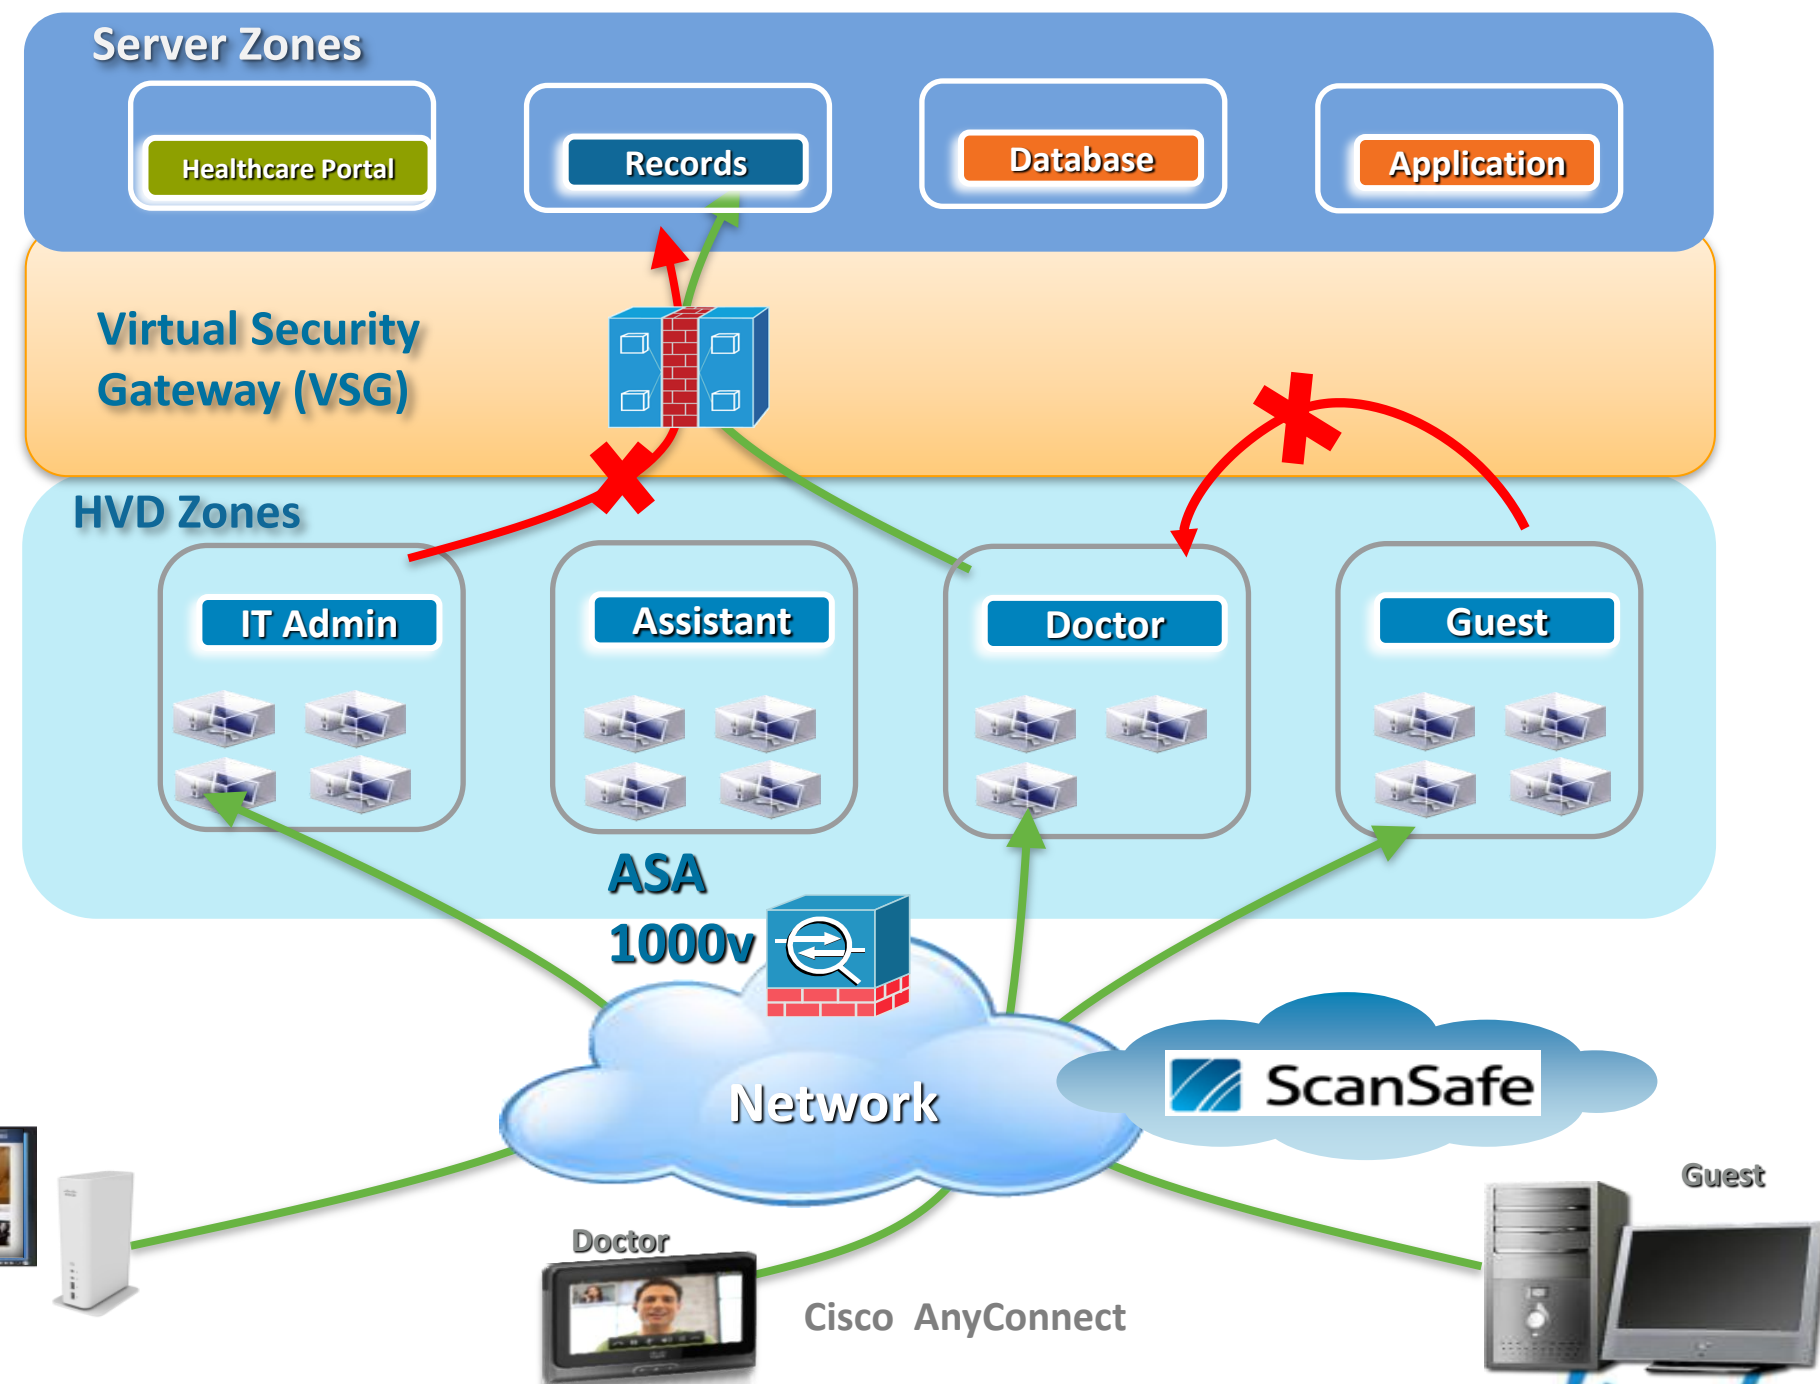
If vm-name contains "TRNG", that VM belongs to TRNG zone

Source	Destination	Protocol	Action
Zone=TRNG	Zone=TRNG	Any	Permit
Any	Zone=TRNG	Any	Permit
Zone=TRNG	Any	Any	Drop

# Securing VDI with Cisco VSG and ASA1000v

- Persistent virtual workspace for the doctor
- Flexible workspace for Doctor's assistant
- Maintain compliance while supporting IT "consumerisation"

Leverage VM context (eg VM-name) to create VSG security policies

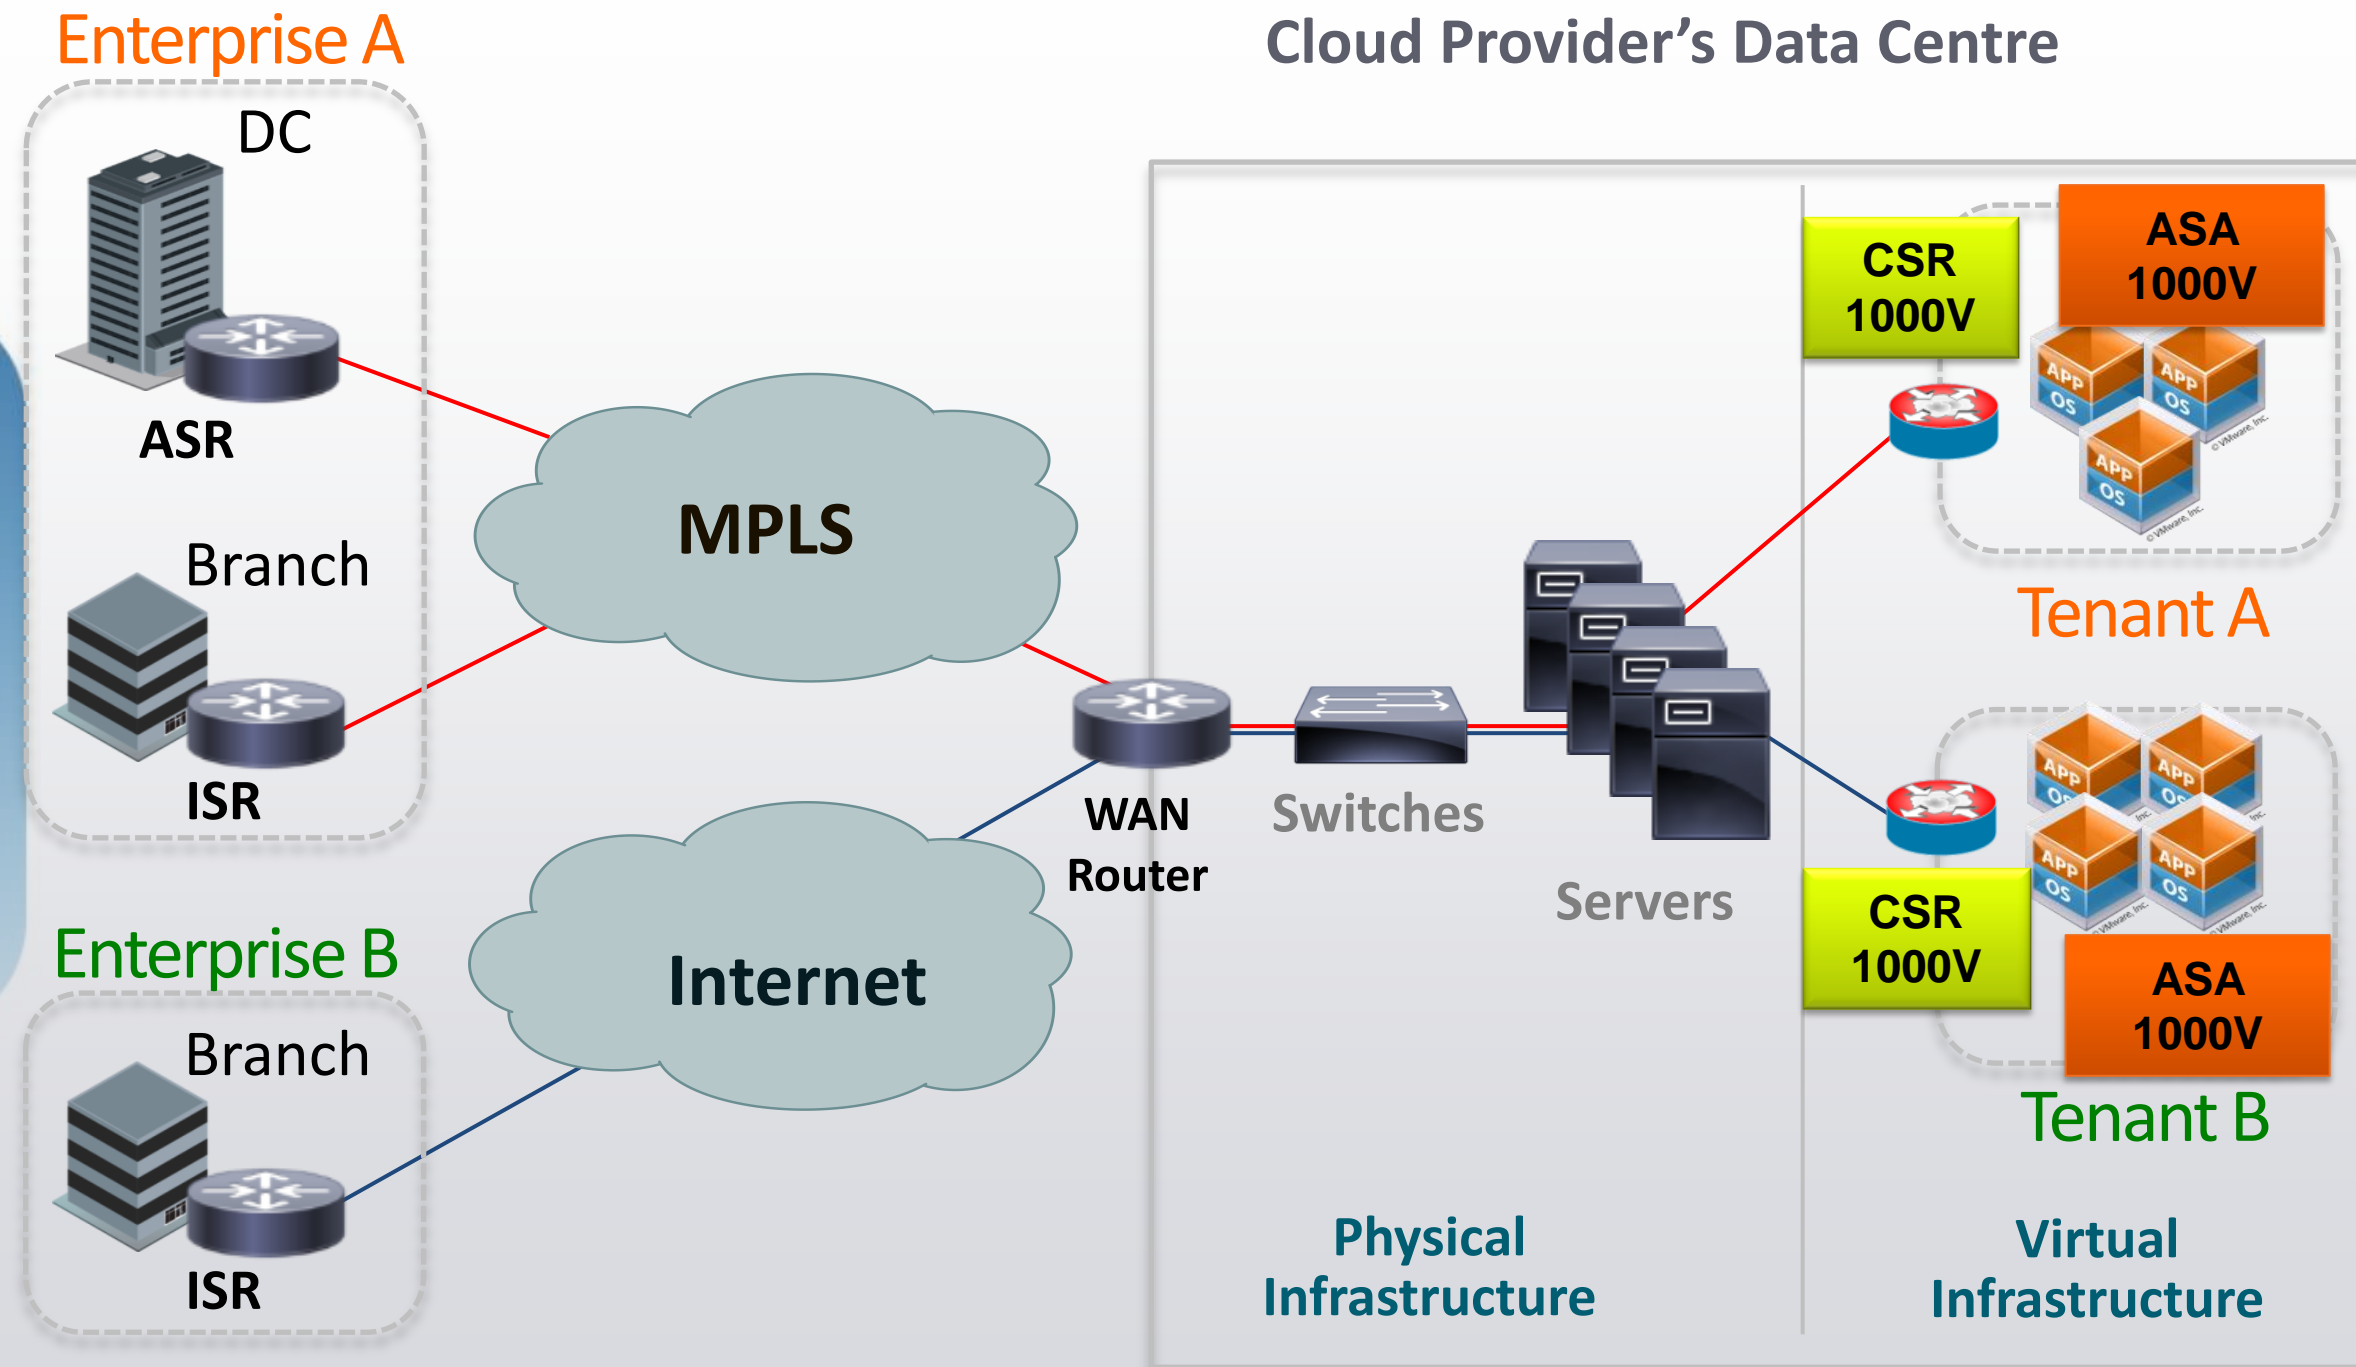


Reference Architecture:

- [1000V and VSG in VXI Reference Architecture](#)

# Single-Tenant WAN Gateway in Shared Multi-tenant Clouds

Can be deployed by Enterprises or Cloud Providers



## Enterprise Use Cases

- Secure VPN Gateway
- L3 Extension
- Tenant Firewall

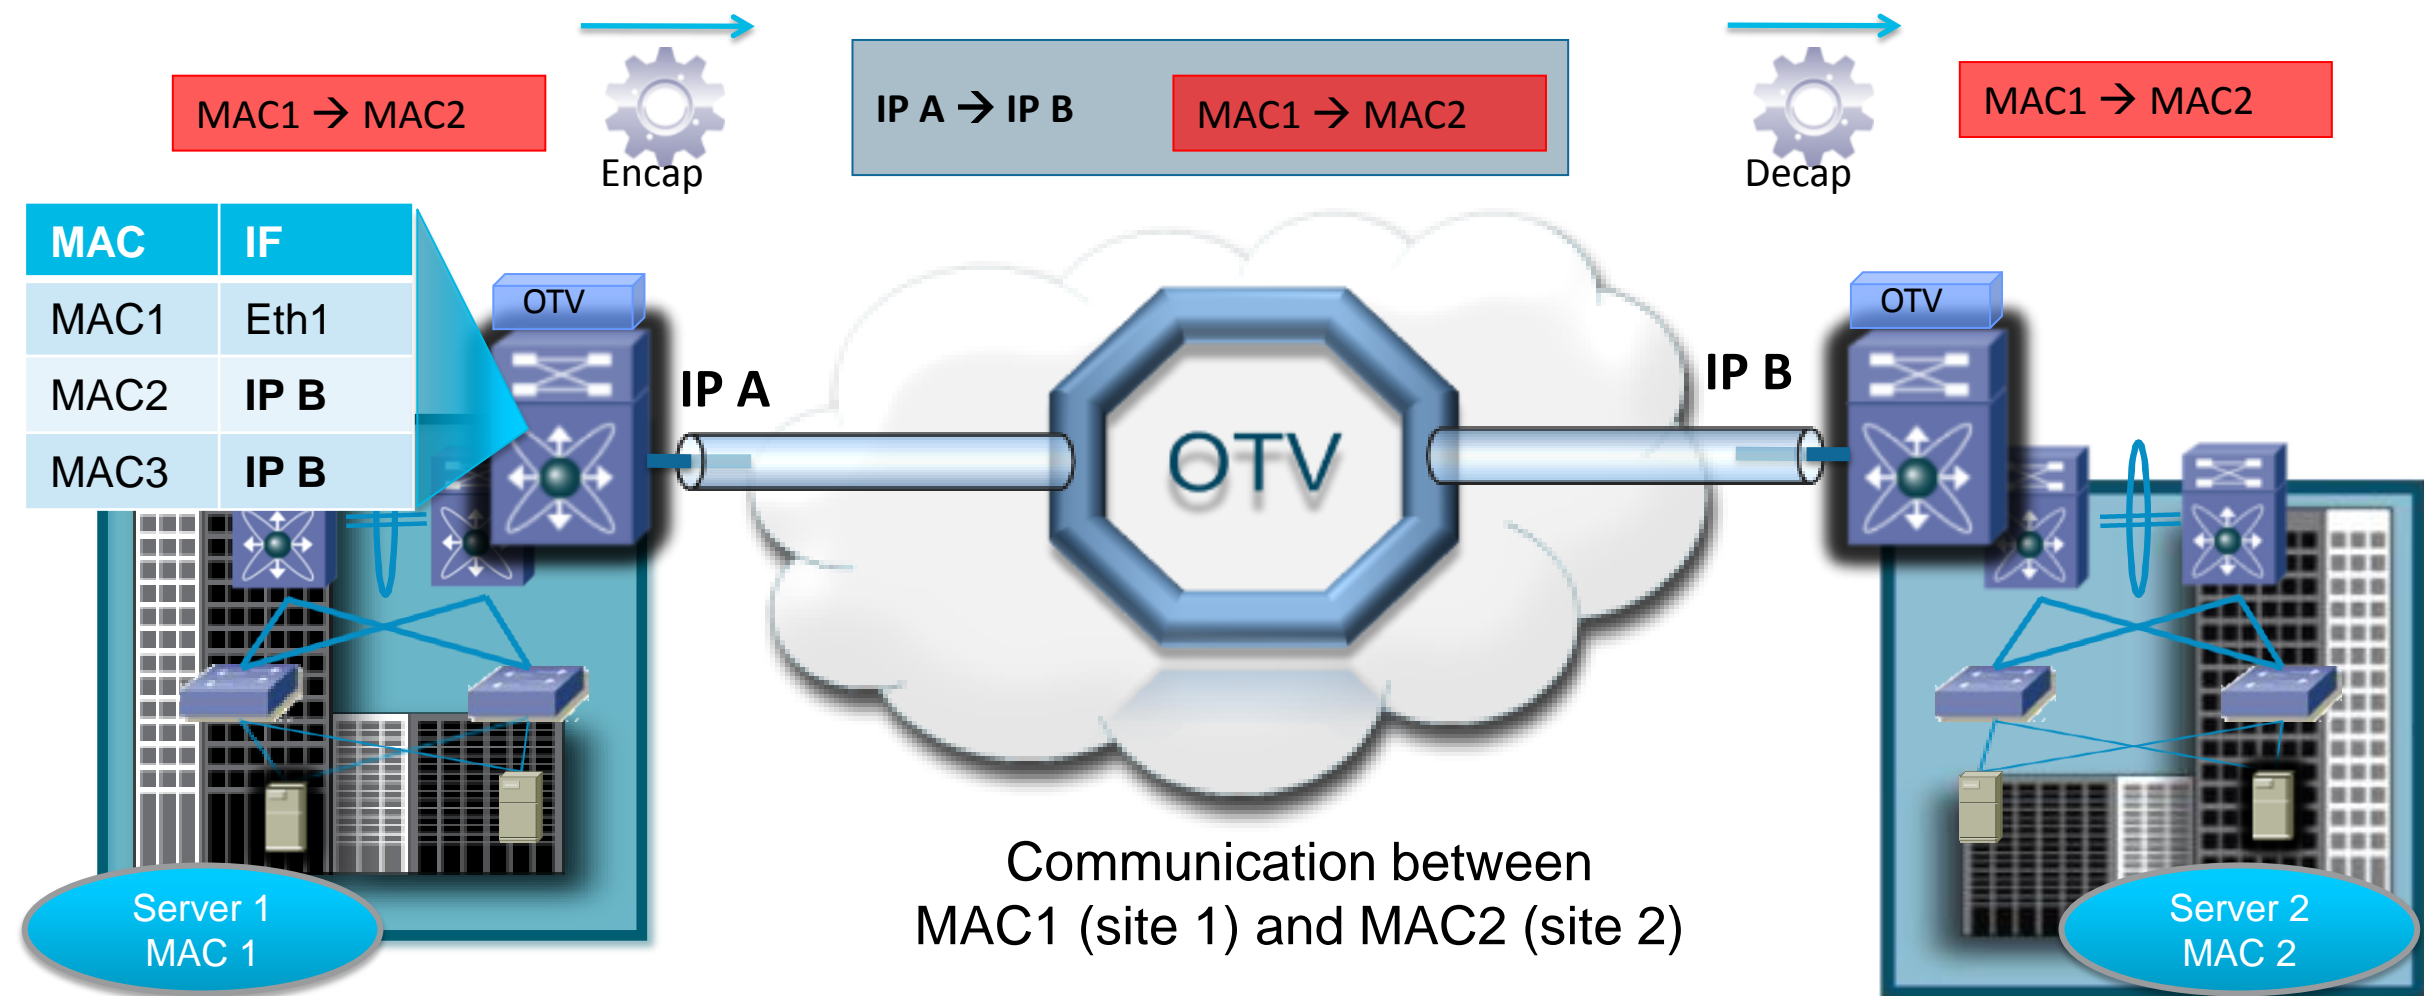
## Cloud Provider Use Cases

- Secure VPN Gateway
- MPLS Extension
- Tenant Firewall

# Overlay Transport Virtualisation (OTV)

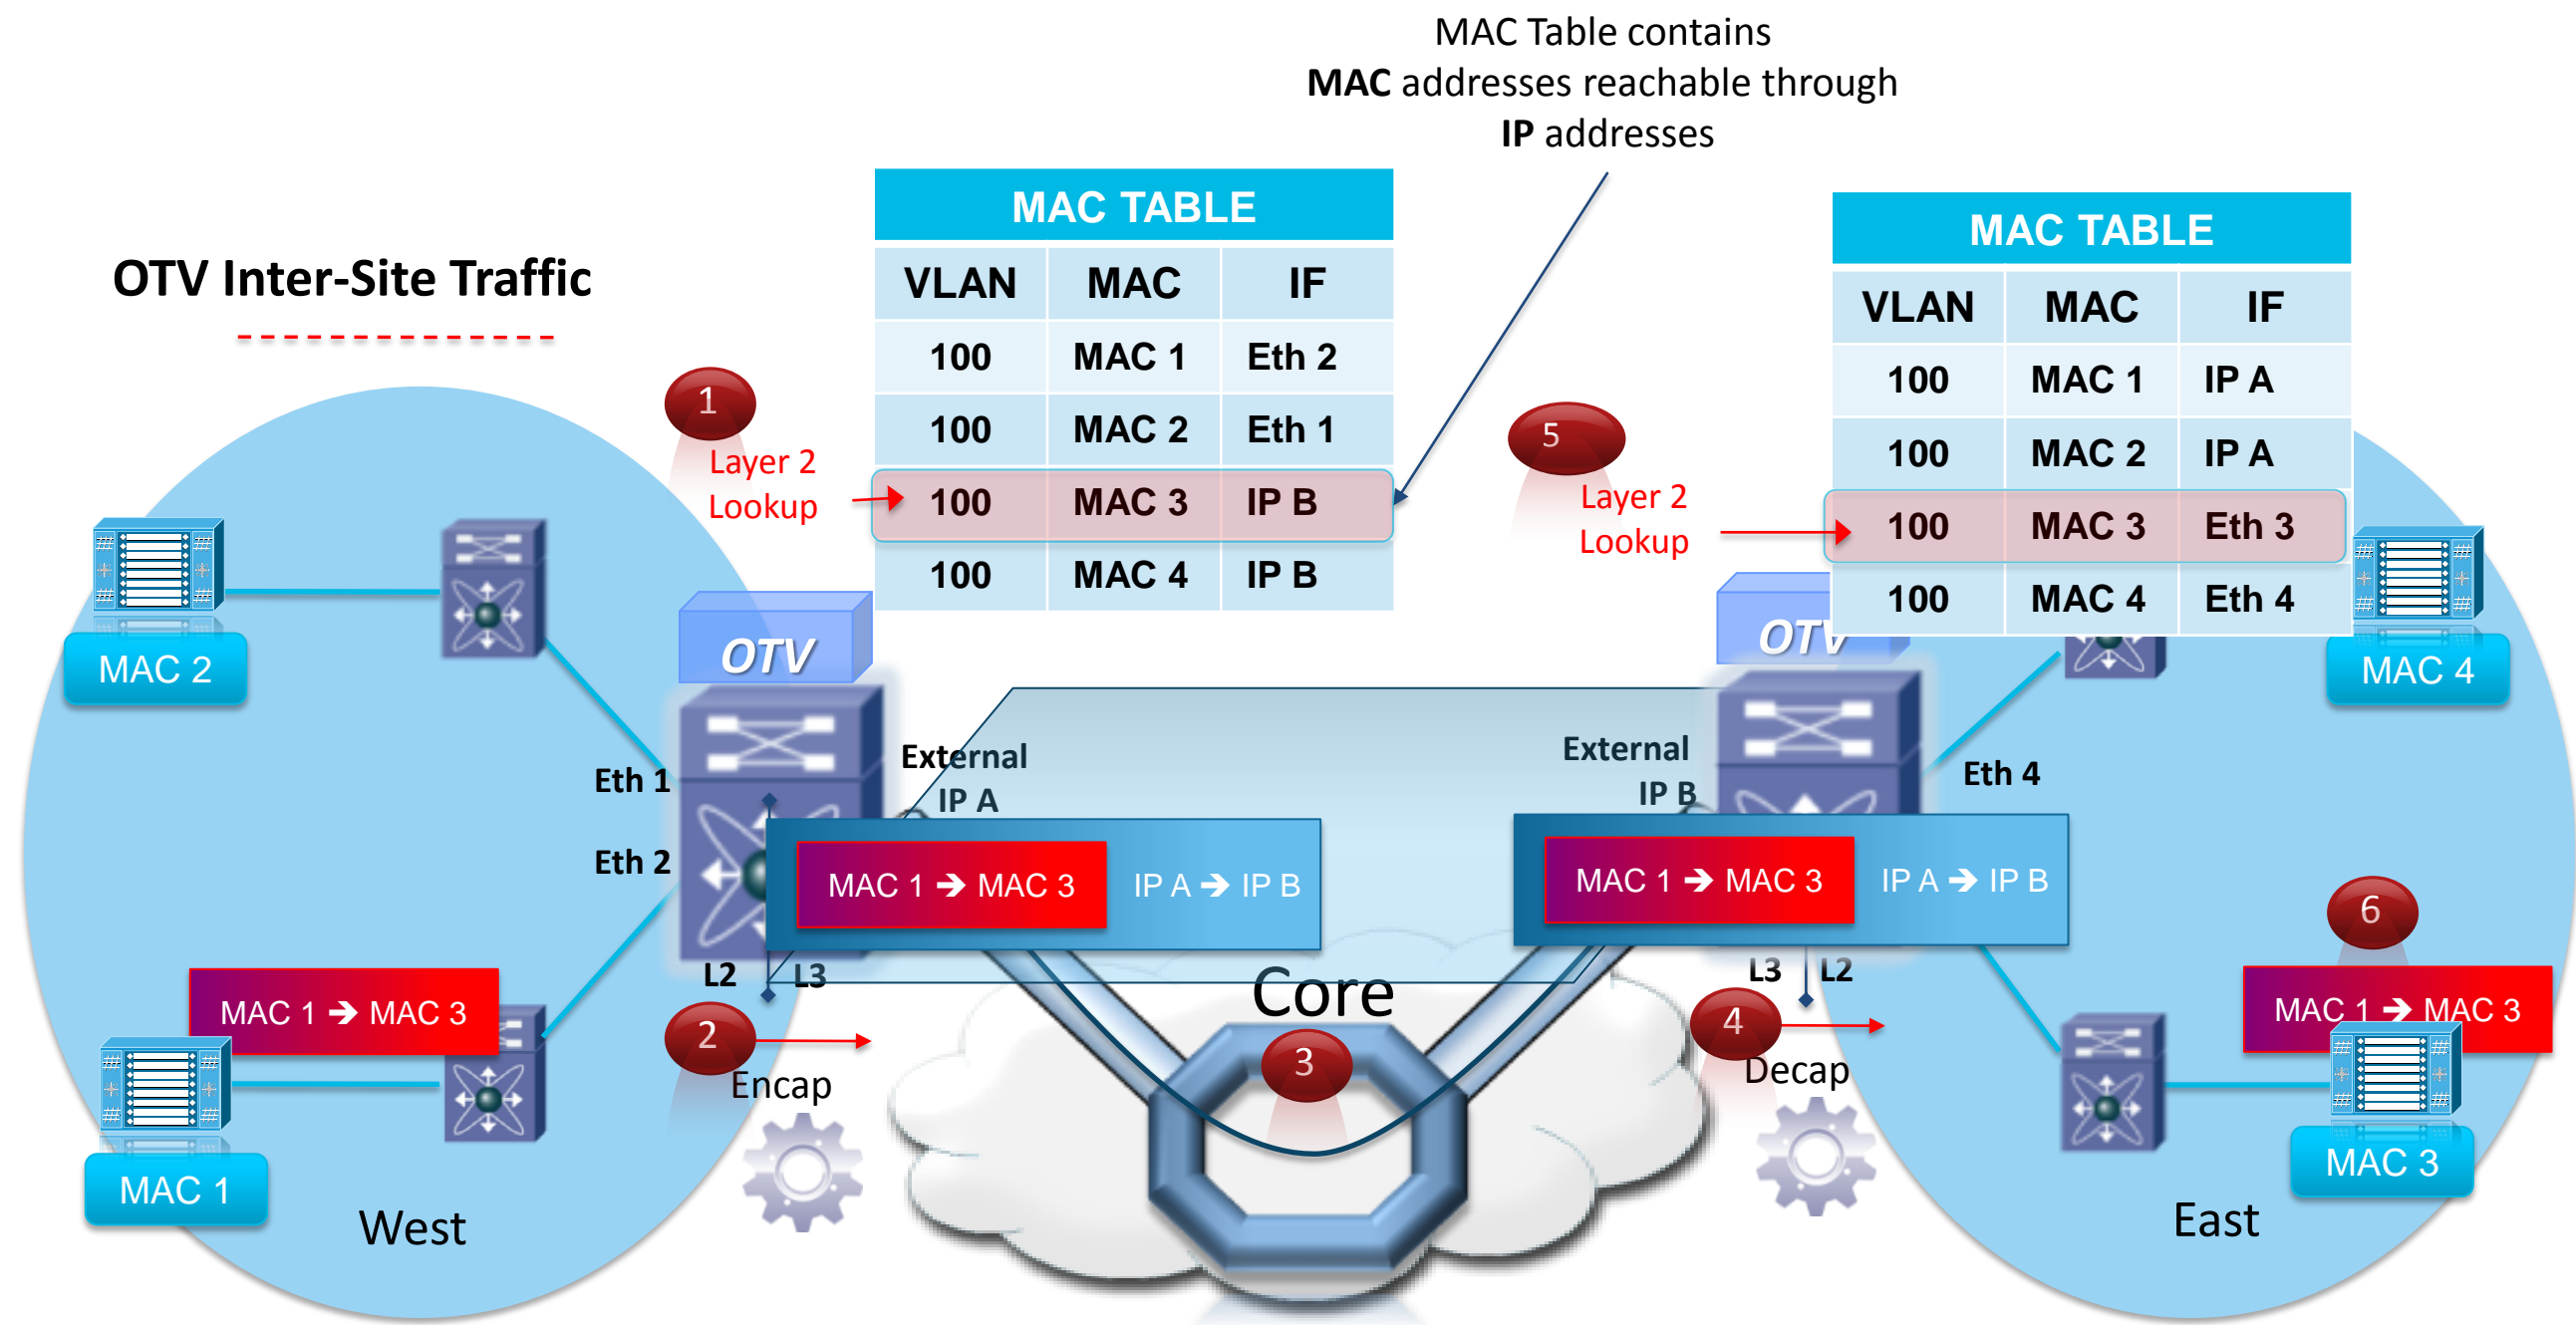
## OTV at a Glance

- Ethernet traffic between sites is encapsulated in IP: “MAC in IP”
- Dynamic encapsulation based on MAC routing table
- No Pseudo-Wire or Tunnel state maintained







# OTV Data Plane: Unicast



- No Pseudo-Wire state is maintained.
- The encapsulation is done based on a Layer 2 destination lookup.

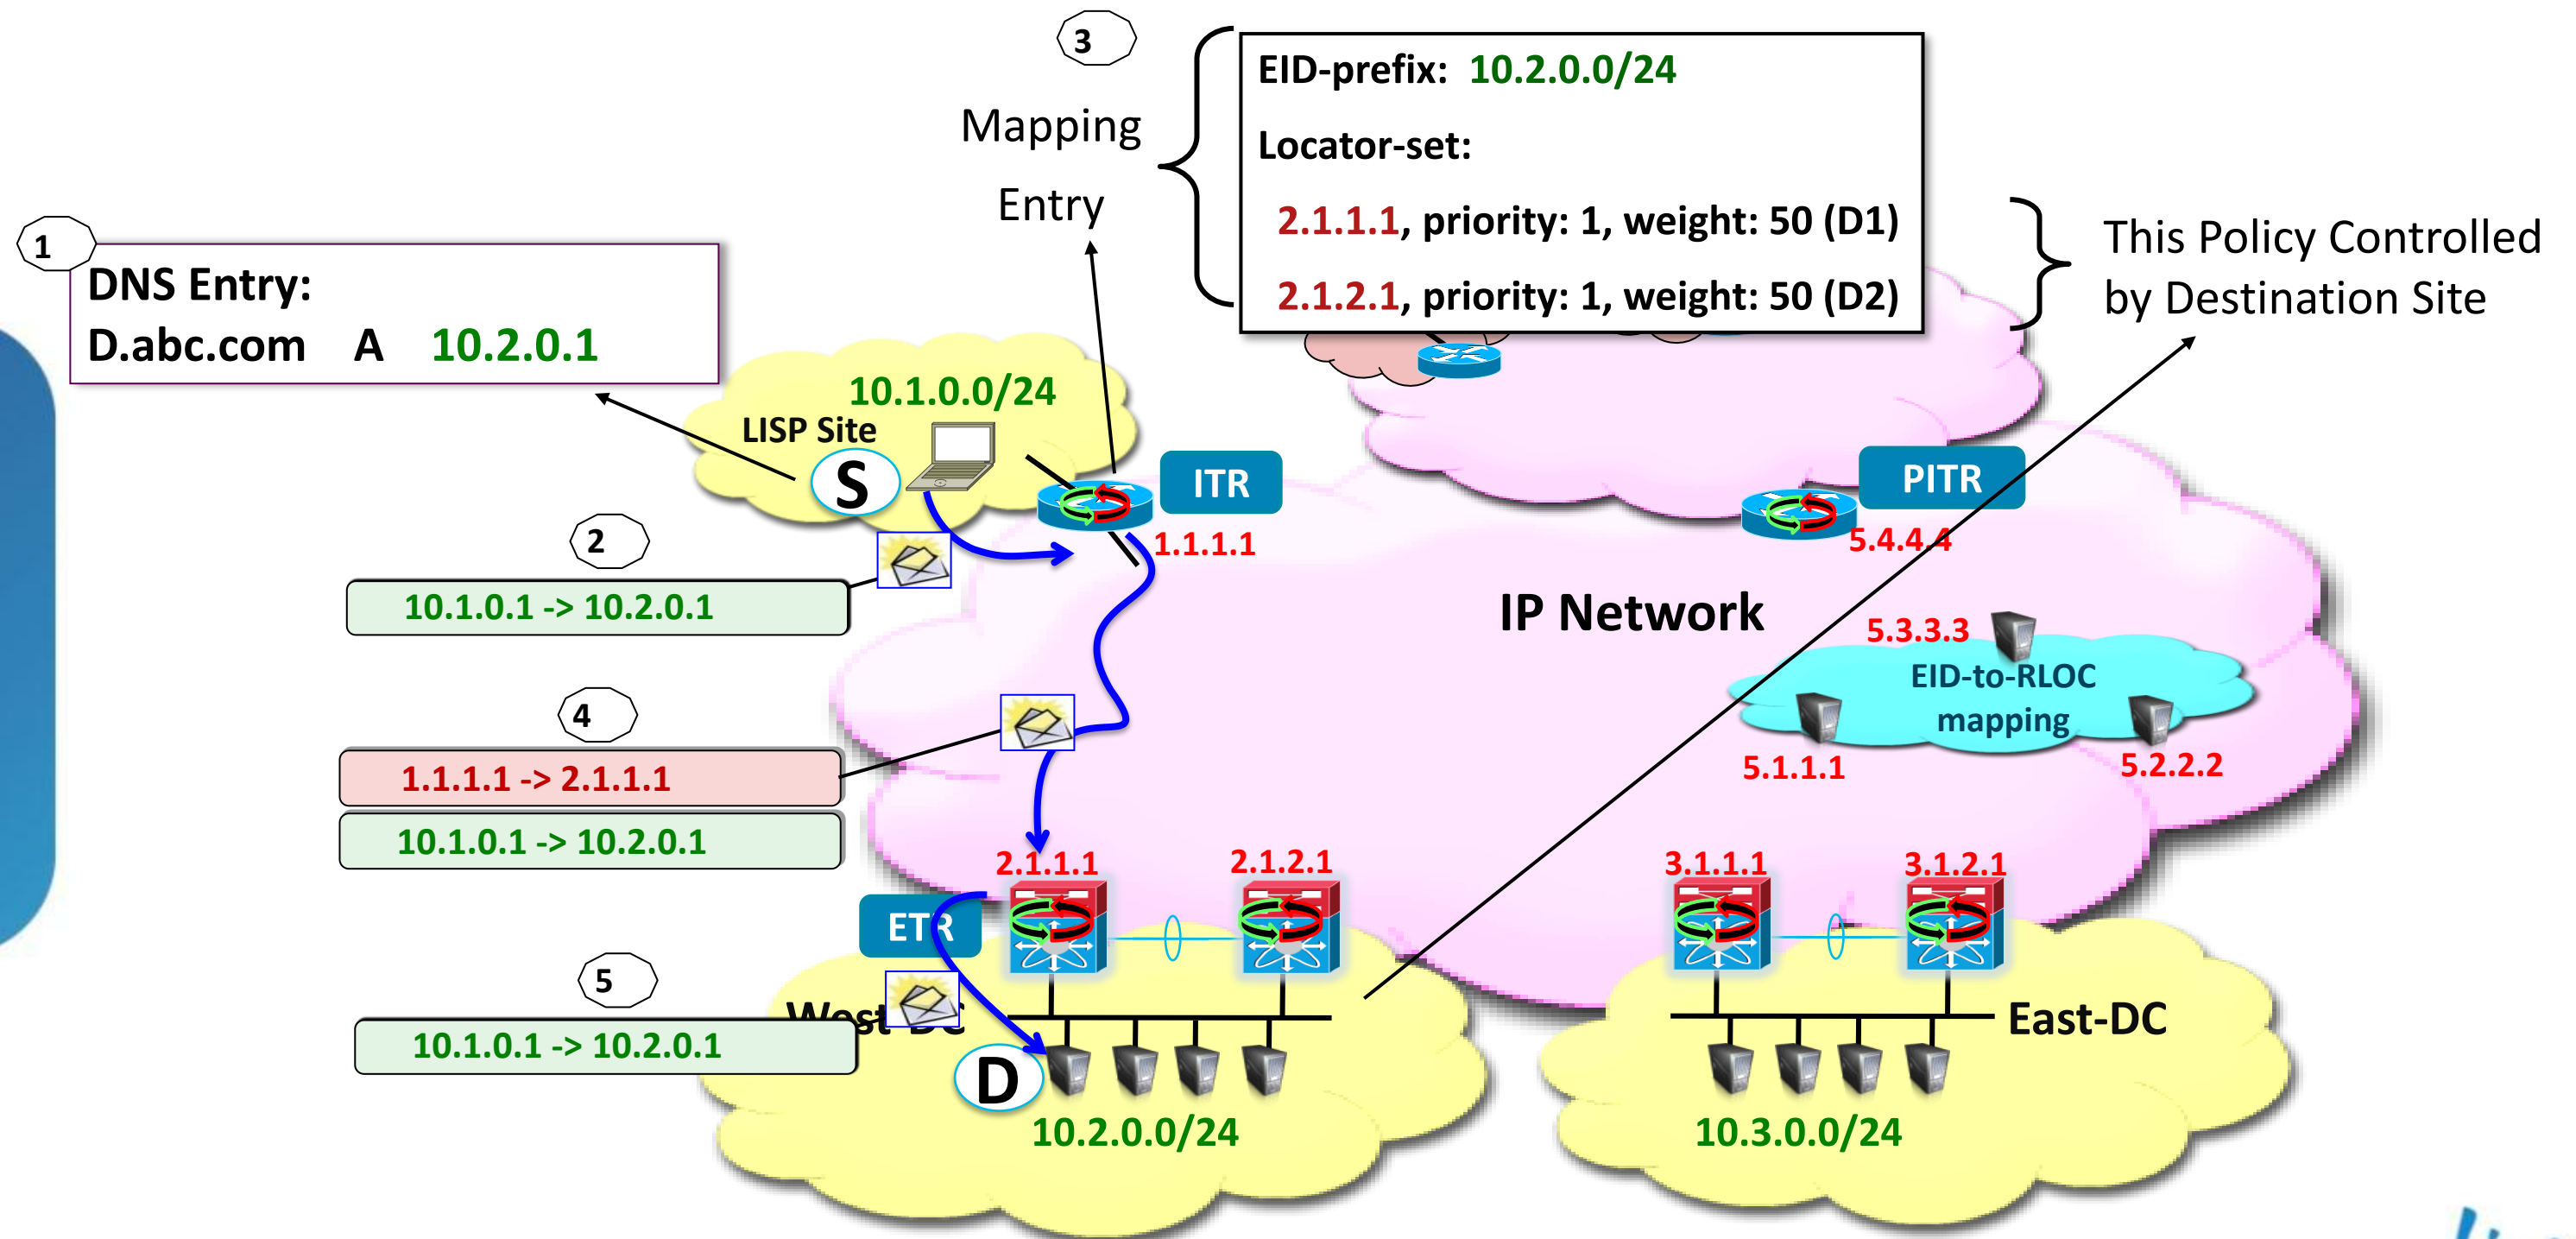
# Location Identity Separation Protocol (LISP)

What Do We Mean by “Location” and “Identity”?



# A LISP Packet Walk

## How Does LISP Operate?

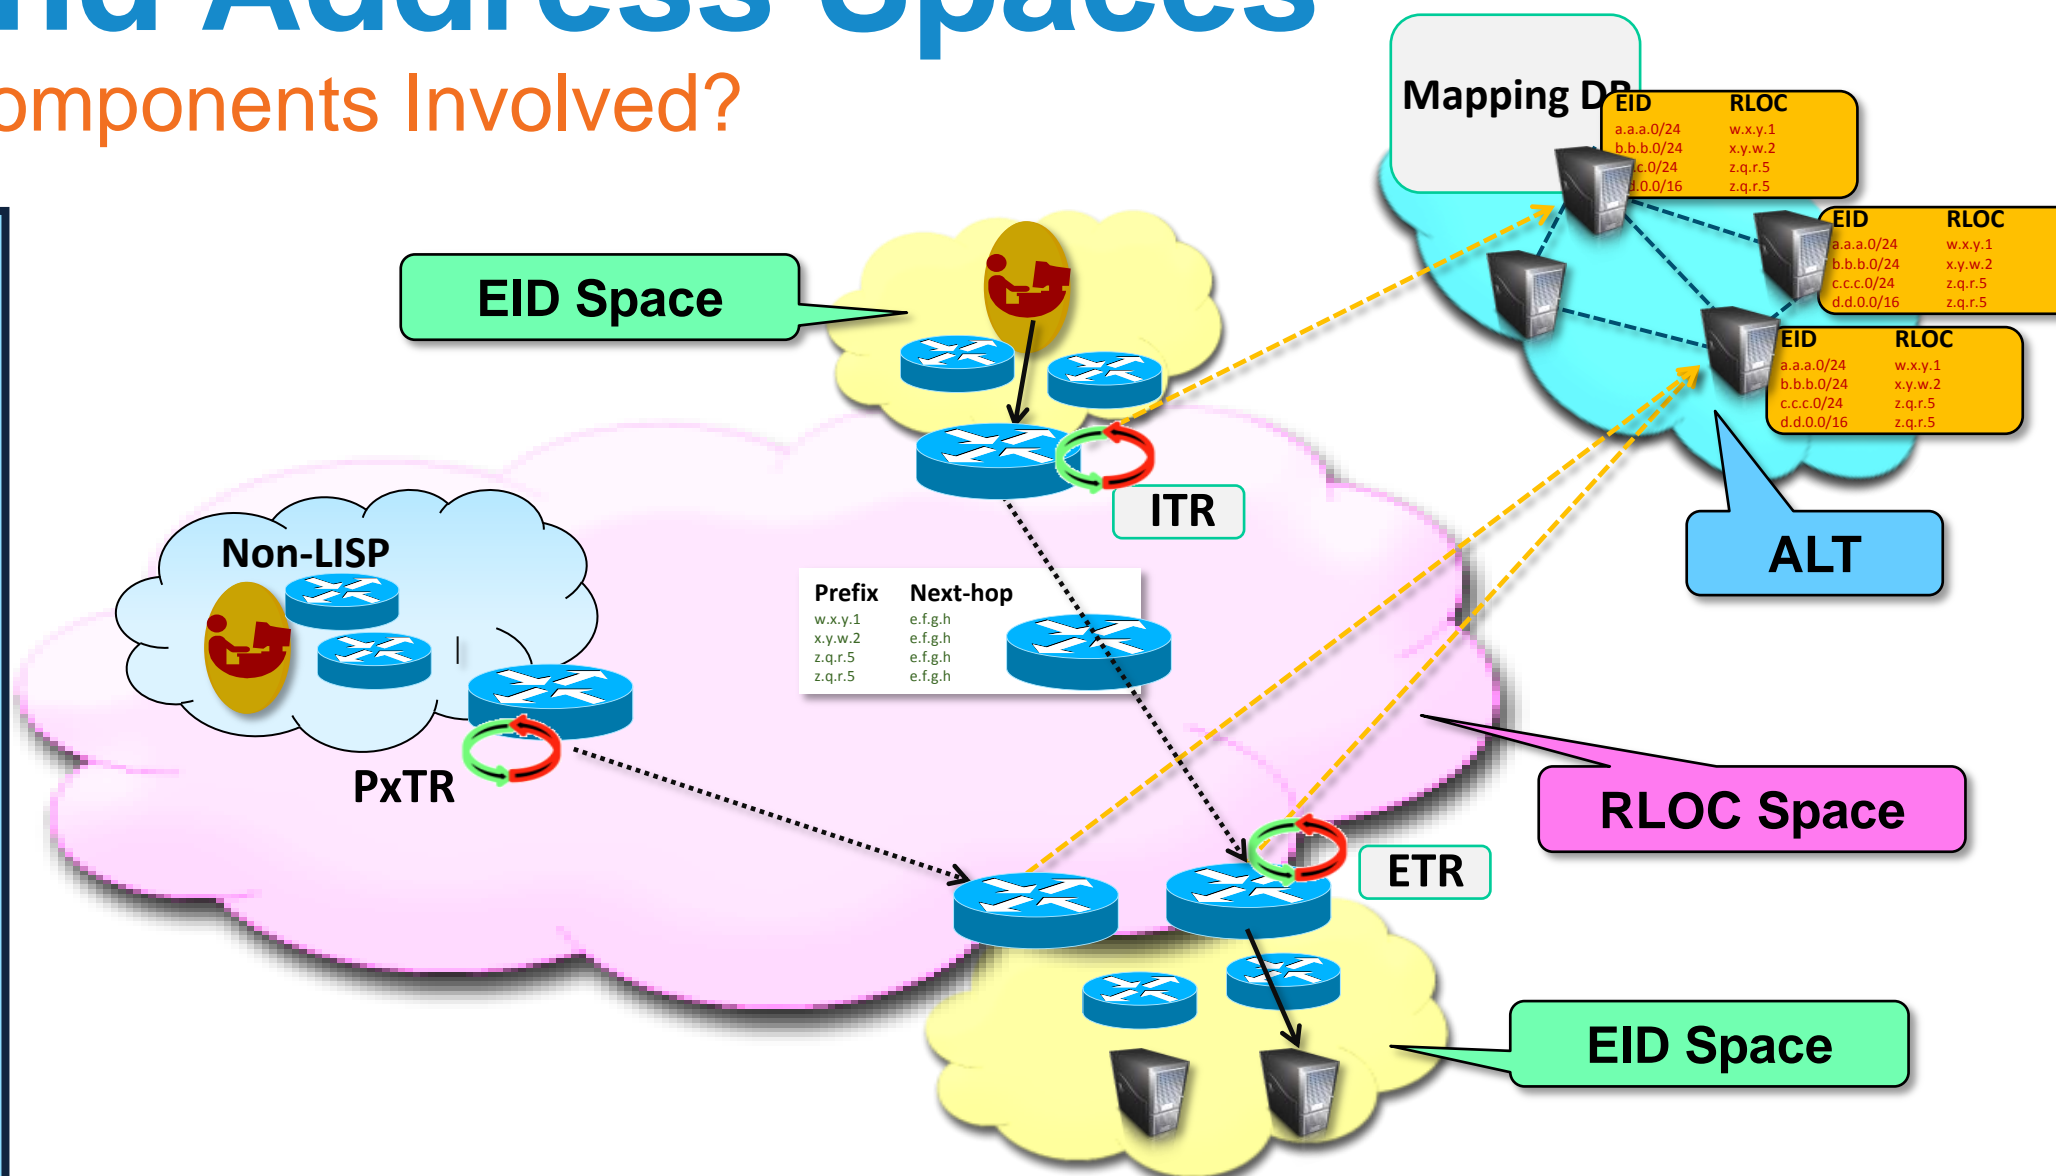


# LISP Roles and Address Spaces

What Are the Different Components Involved?

## LISP Roles

- **Tunnel Routers - xTRs**
  - Edge devices in charge of encap/decap
  - Ingress/Egress Tunnel Routers (ITR/ETR)
- **EID to RLOC Mapping DB**
  - Contains RLOC to EID mappings
  - Distributed across multiple Map Servers (MS)
  - MS may connect over an ALT network
- **Proxy Tunnel Routers - PxTR**
  - Coexistence between LISP and non-LISP sites
  - Ingress/Egress: PITR, PETR



## Address Spaces

- **EID = End-point Identifier**
  - Host IP or prefix
- **RLOC = Routing Locator**
  - IP address of routers in the backbone

# LISP Host-Mobility

## Needs:


- Global IP-Mobility **across subnets**
- Optimised routing **across extended subnet sites**

## LISP Solution:

- Automated **move detection** on xTRs
- Dynamically update EID-to-RLOC mappings
- **Traffic Redirection** on ITRs or PITRs


## Benefits:

- Direct Path (no triangulation)
- Connections maintained across move
- No routing re-convergence
- No DNS updates required
- Transparent to the hosts
- Global Scalability (cloud bursting)
- IPv4/IPv6 Support



# Host-Mobility Scenarios

## Moves Without LAN Extension




### IP Mobility Across Subnets

Disaster Recovery  
Cloud Bursting

Application Members in One Location

## Moves With LAN Extension



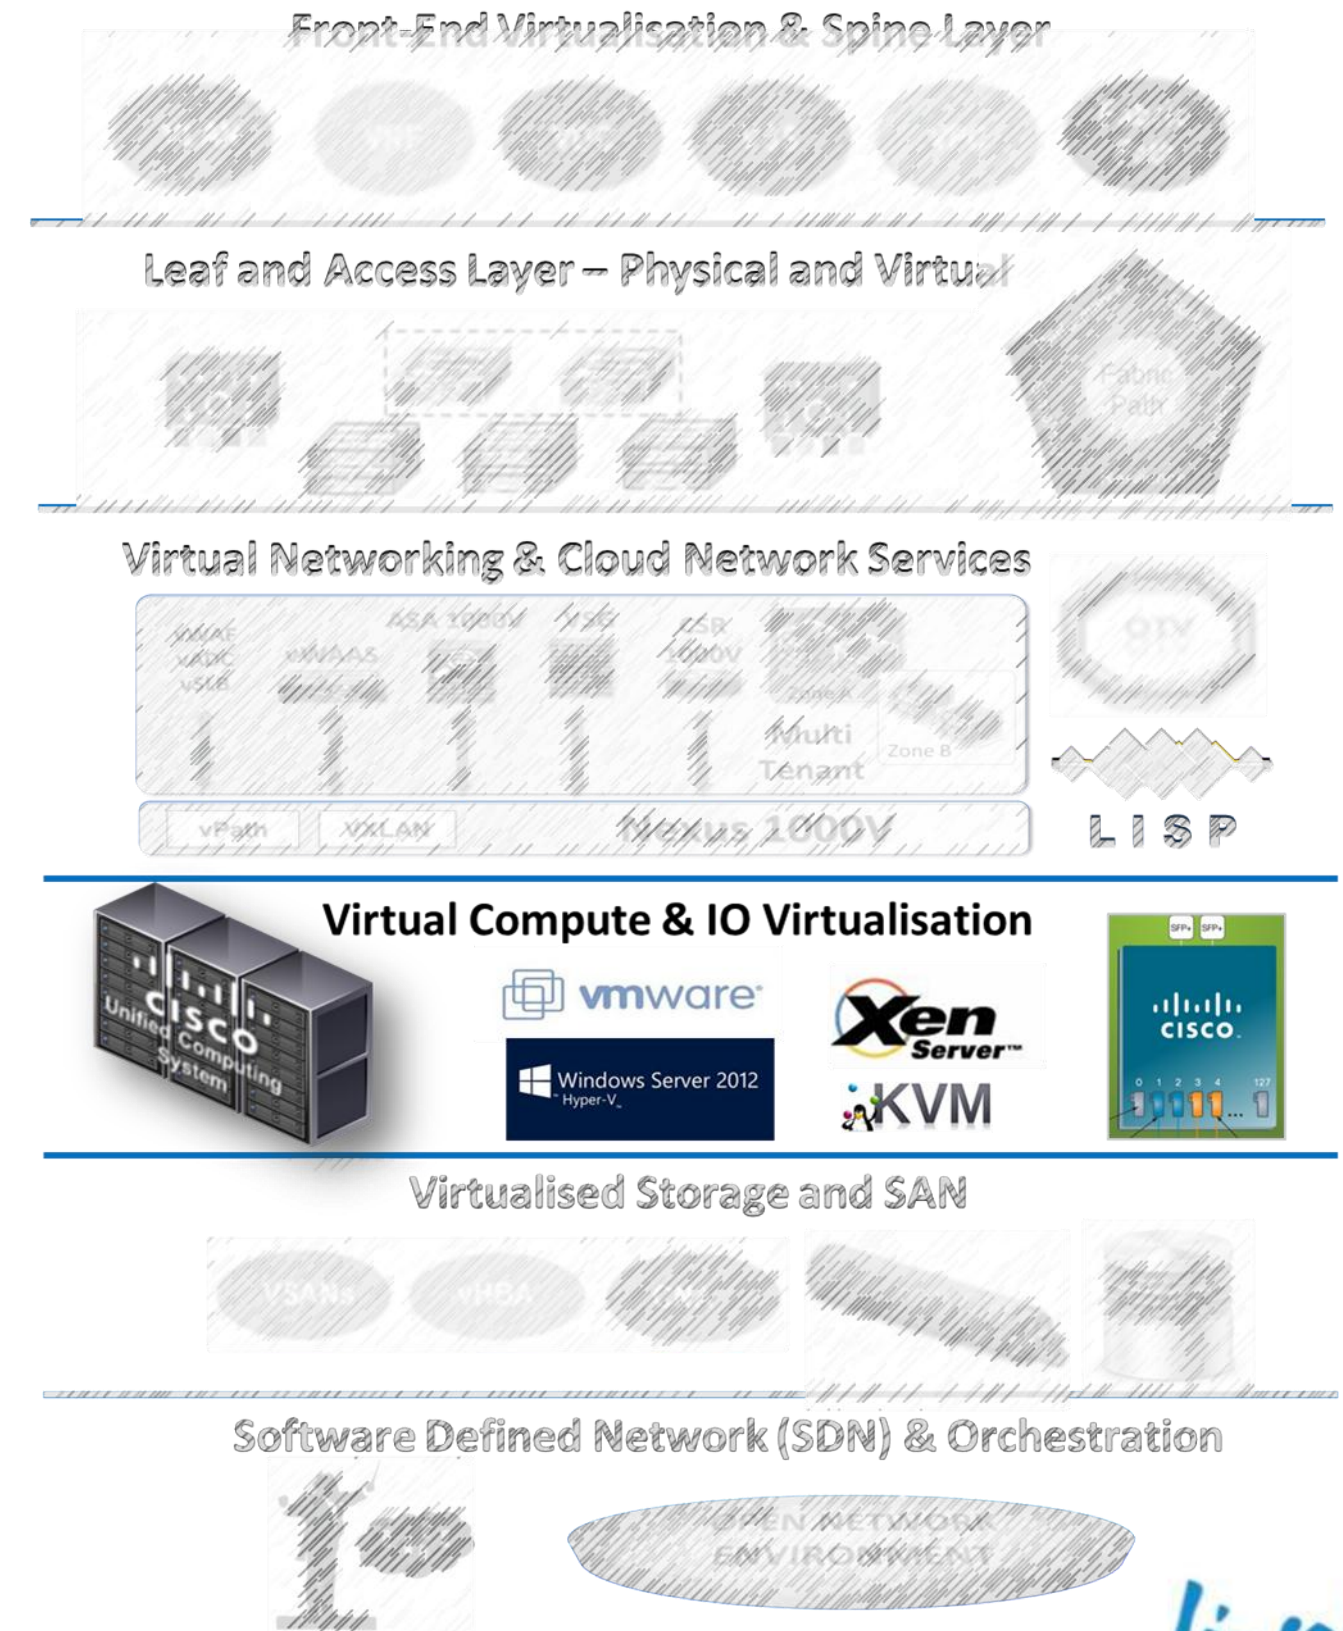
### Routing for Extended Subnets

Active-Active Data Centres  
Distributed Clusters

Application Members Distributed  
(Broadcasts across sites)

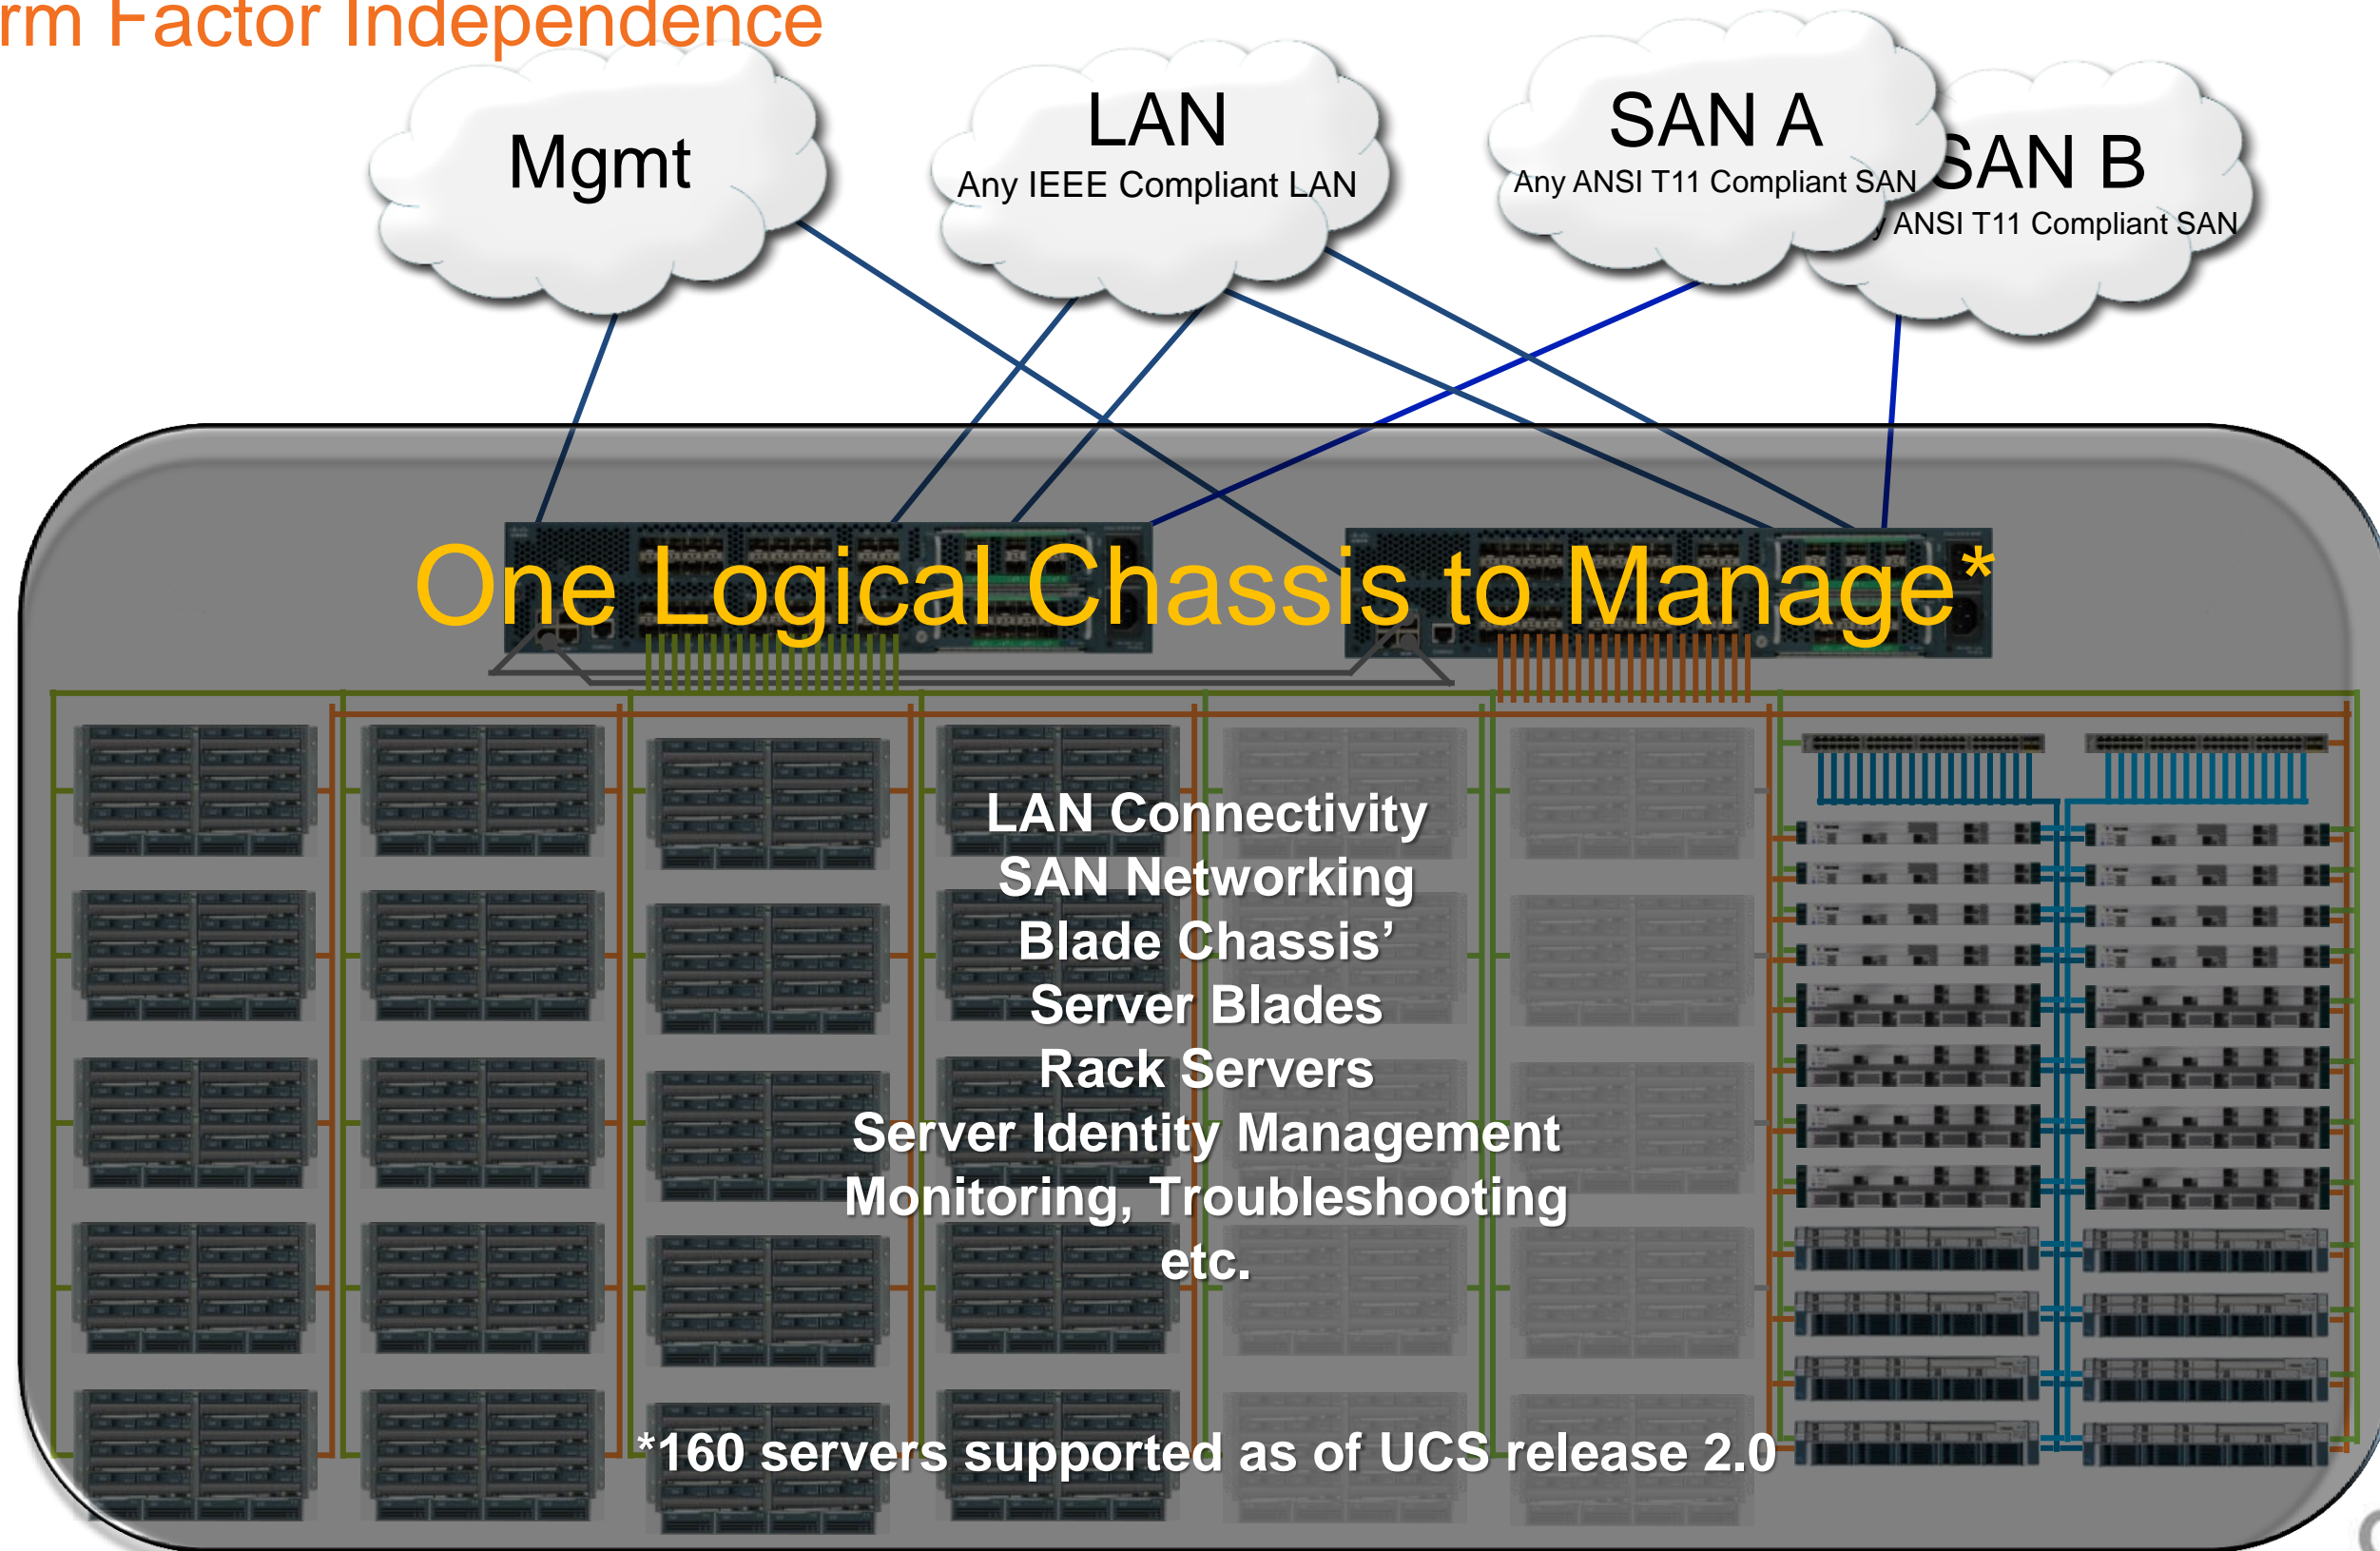
# Agenda

- **Data Centre Virtualisation Overview**
- **Front-End Data Centre Virtualisation**
  - DC Edge Layer
  - Aggregation / Spine Layer
  - Access / Leaf Layer
- **Virtual Networking & Cloud Network Services**
  - Virtual Access Layer: Nexus 1000v
  - **Cloud Network Services**
  - Inter DC Services: OTV & LISP
- **Virtual Compute & IO Virtualisation**
  - **Unified Computing System**
  - **Hypervisors & Virtual IO @ Cisco UCS**
- **Virtualised Storage & SAN**
  - Unified Fabric
  - Virtualised SAN and Storage
- **Software Defined Network & Orchestration**
  - SDN & Cisco Open Network Environment (ONE)
  - Cloud Orchestration
- **Implementation Examples**
- **Q&A**




# Cisco Unified Computing System (UCS)

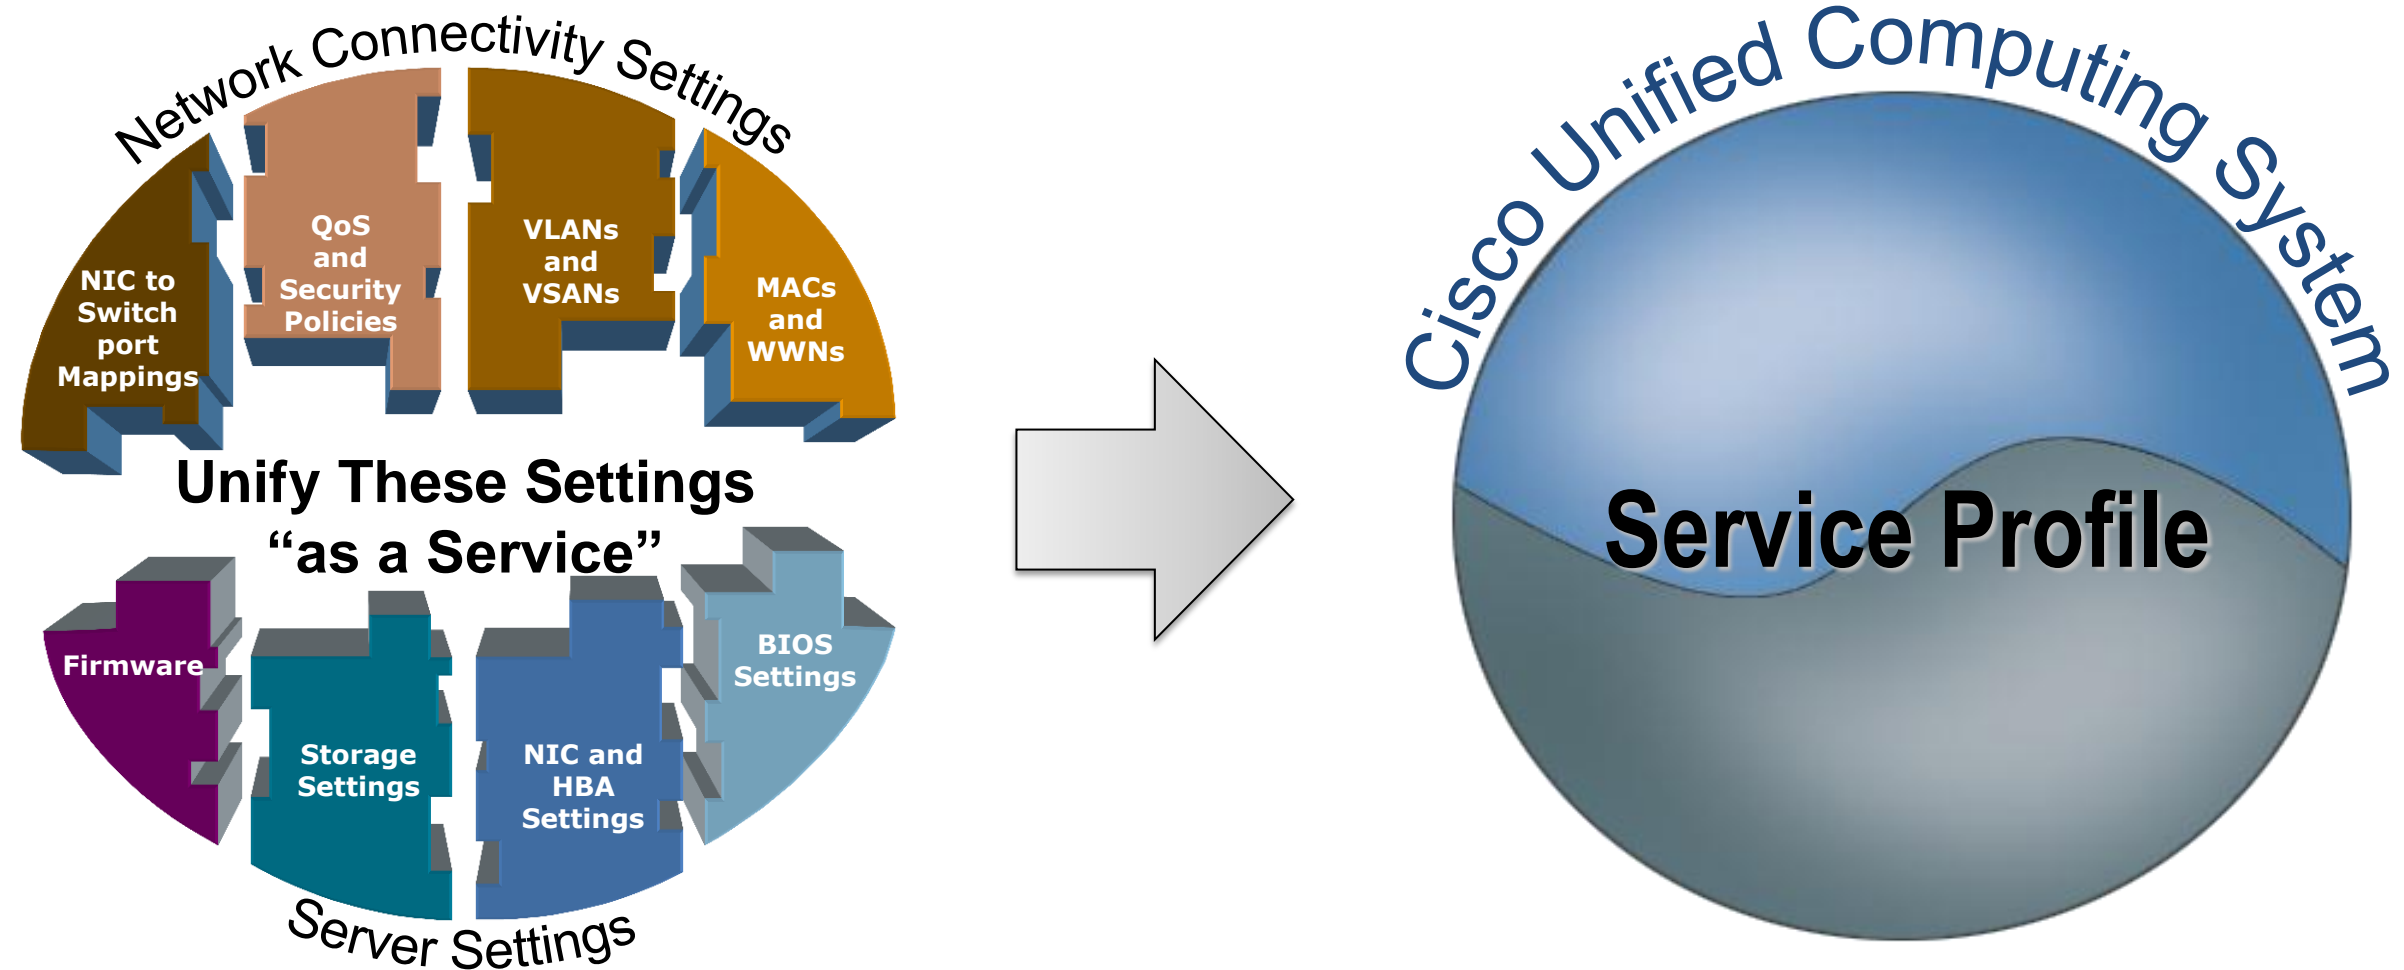
Form Factor Independence







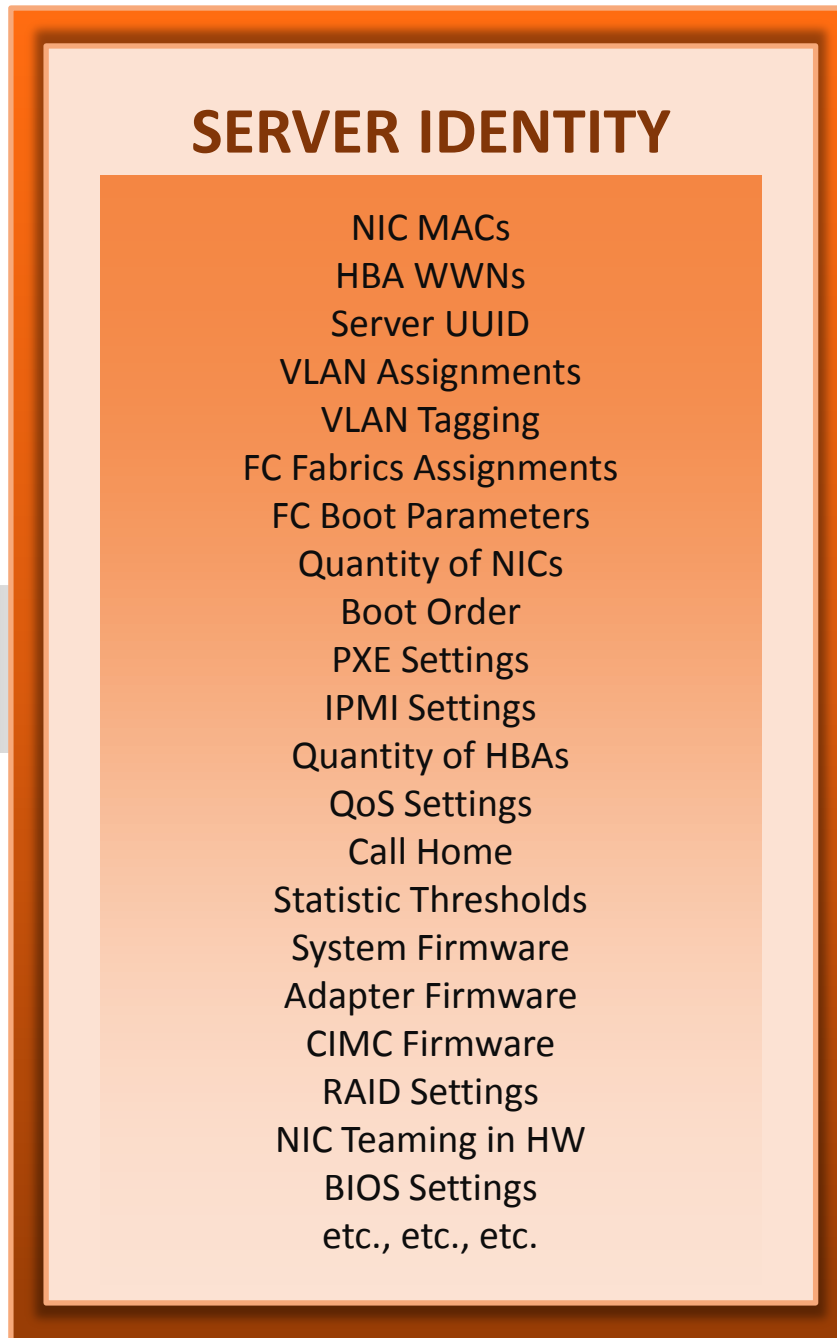



# Unifying Configuration Settings to Deliver a “Server as a Service”





# Stateless Computing: Service Profiles



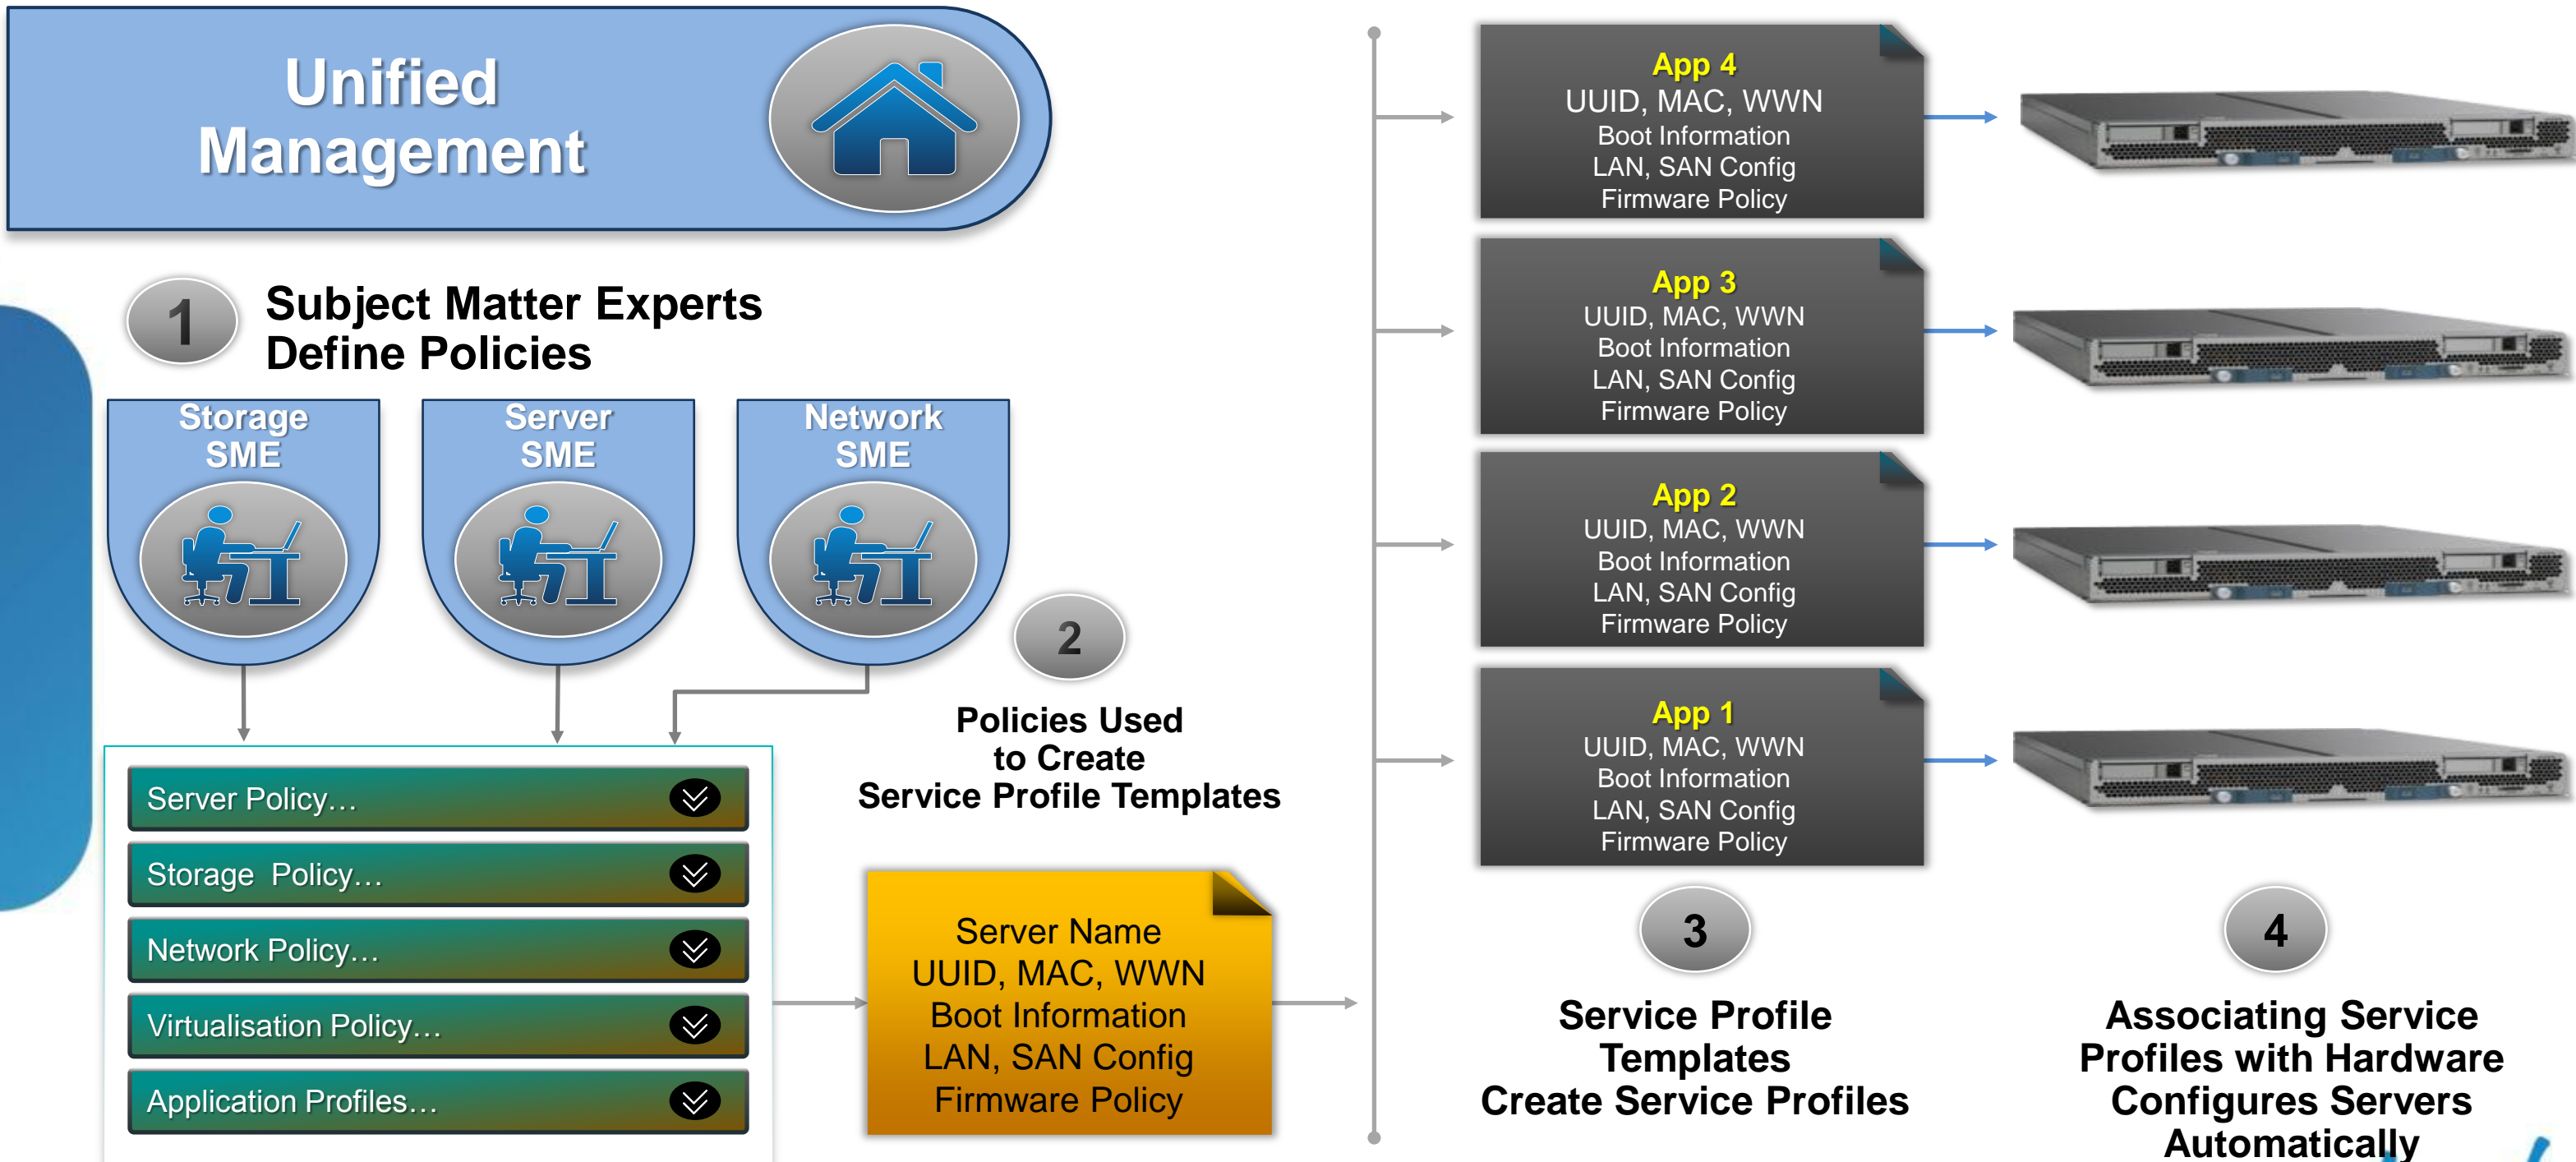
**Adds:**

- **Portability**
- **More flexibility**
- **Improved uptime**



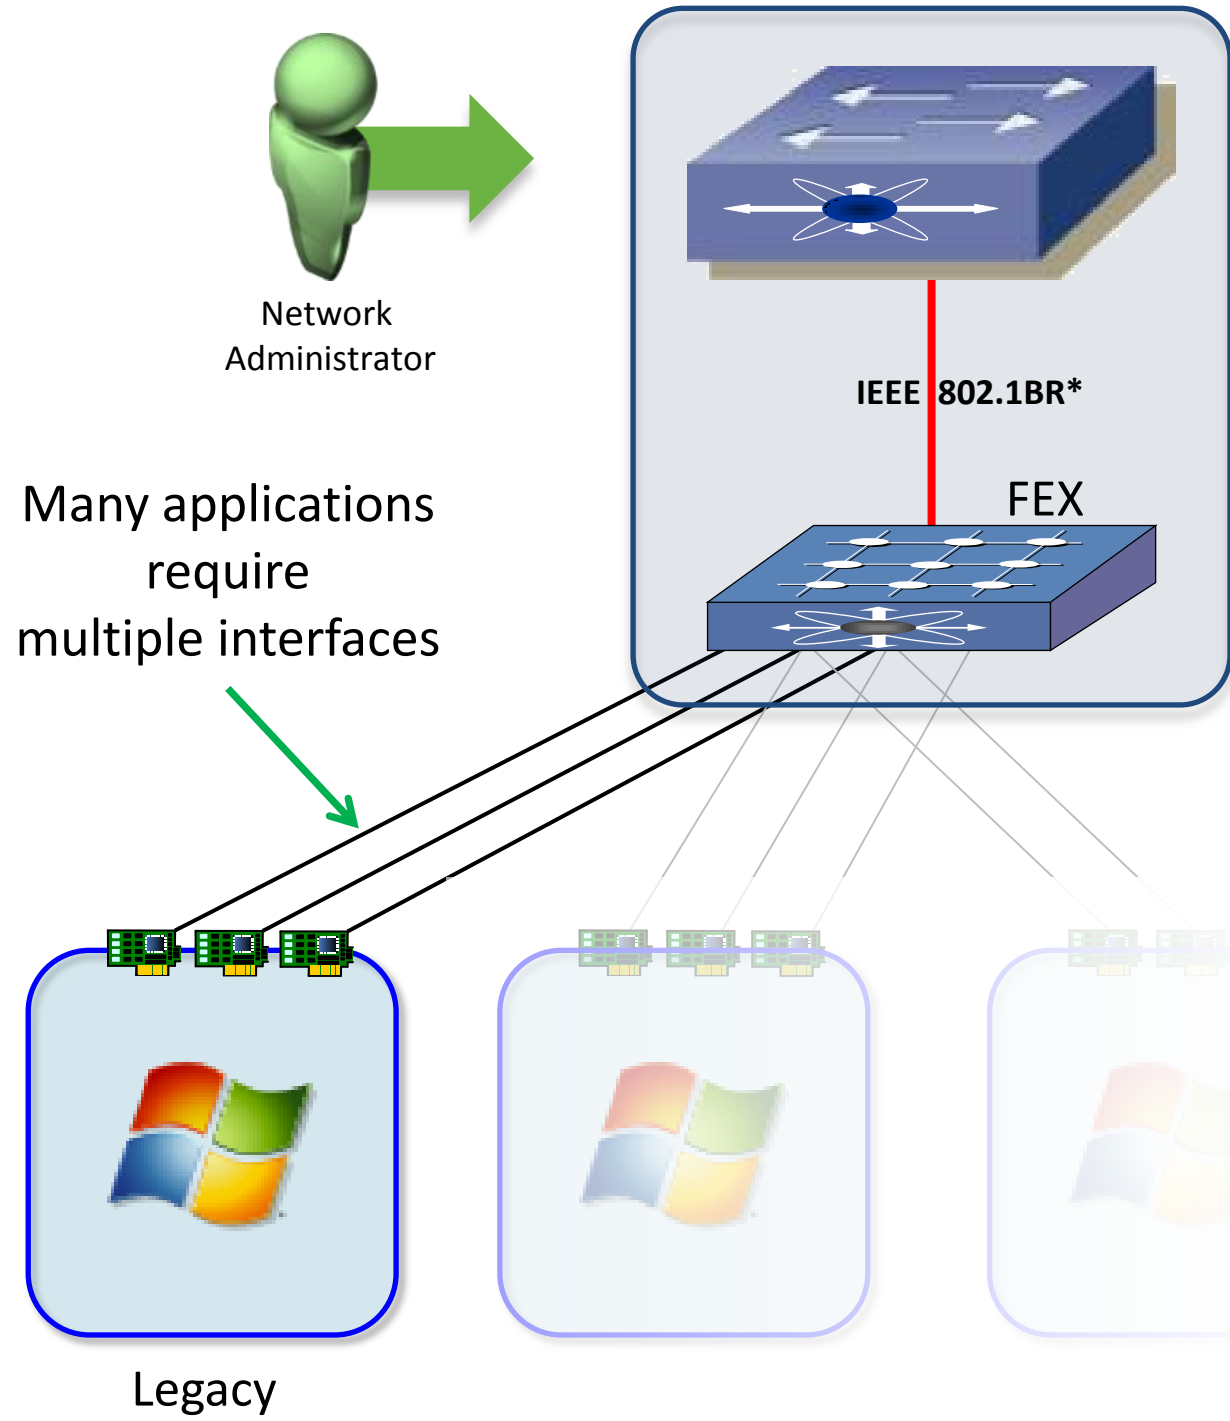

# Unified & Stateless via UCS Service Profiles

Aligns People, Policy, and Configuration With Workload



# Fabric Extender Evolution

Distributed Modular System to the ToR, Server, and Virtual Machine



Legacy

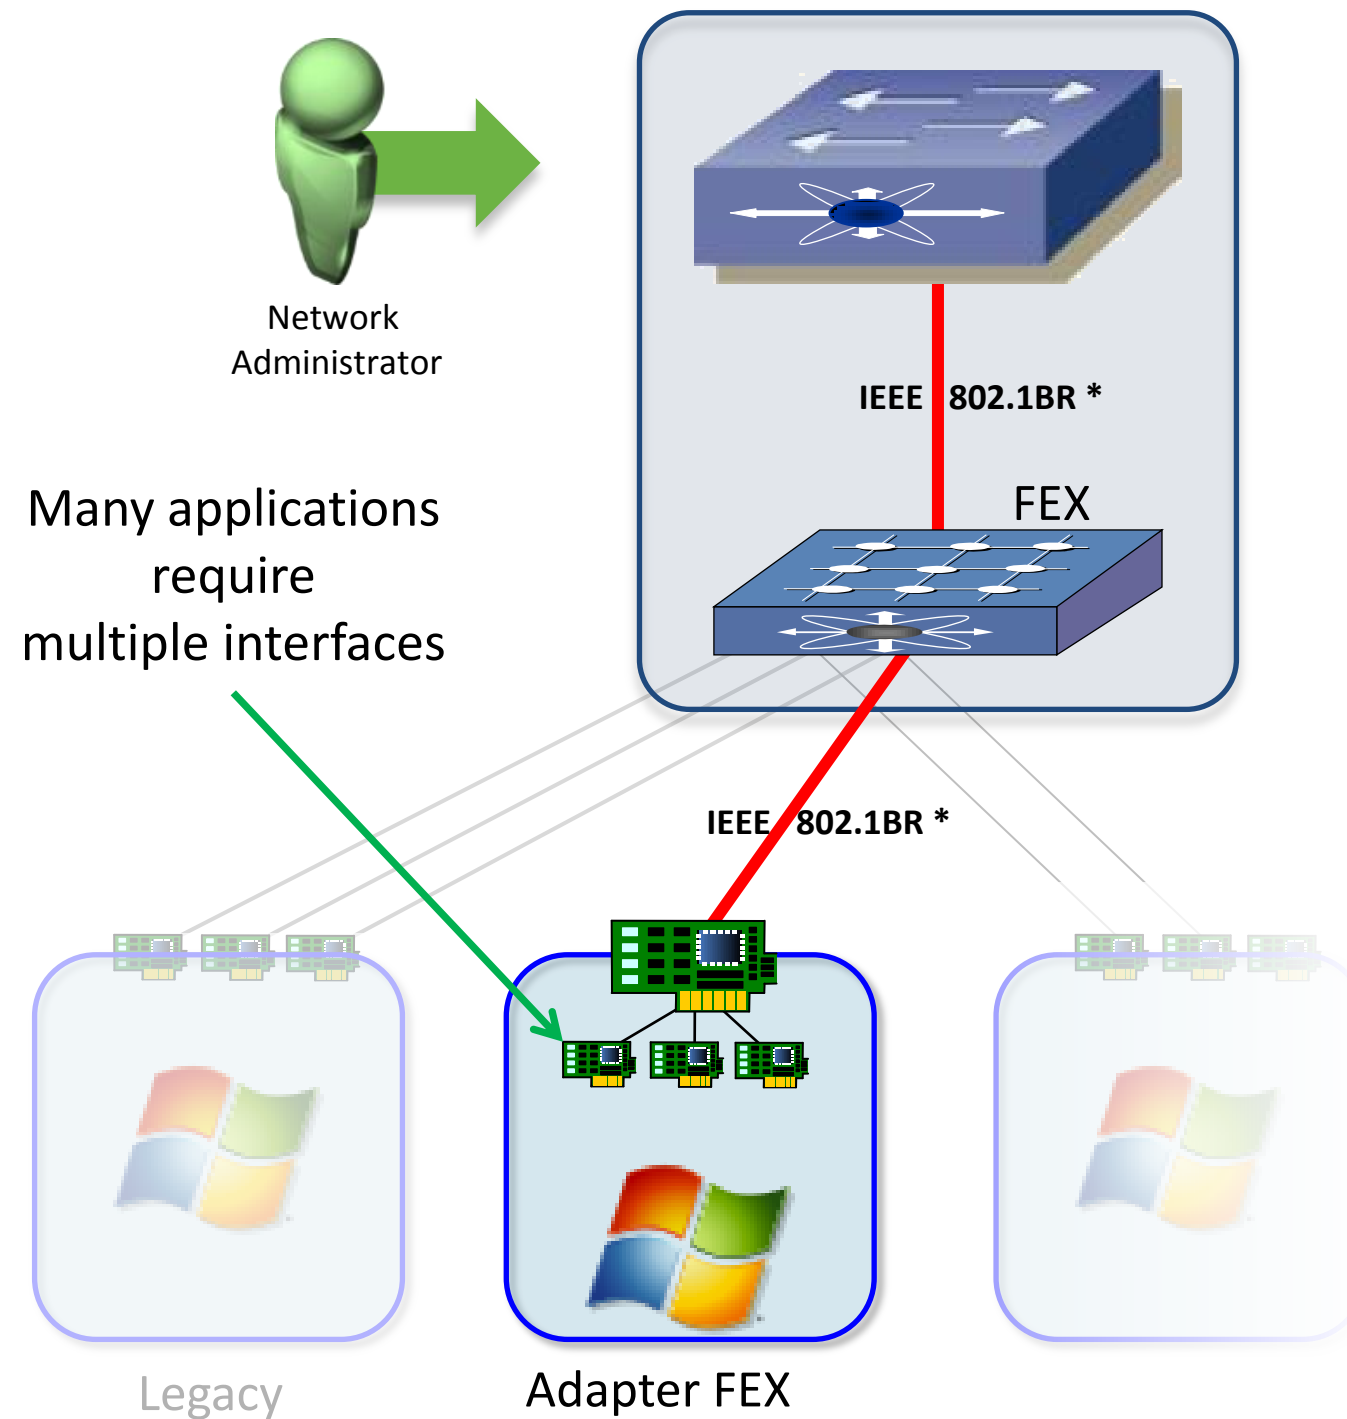

\*IEEE 802.1BR Pre-Standard

## One Network Parent Switch to Top of Rack FEX Architecture

- **Consolidates** network management
- FEX managed as line card of parent switch
- Uses Pre-standard IEEE 802.1BR

# Fabric Extender Evolution

Distributed Modular System to the ToR, Server, and Virtual Machine



Many applications require multiple interfaces

Legacy

Adapter FEX

\*IEEE 802.1BR Pre-Standard

## One Network

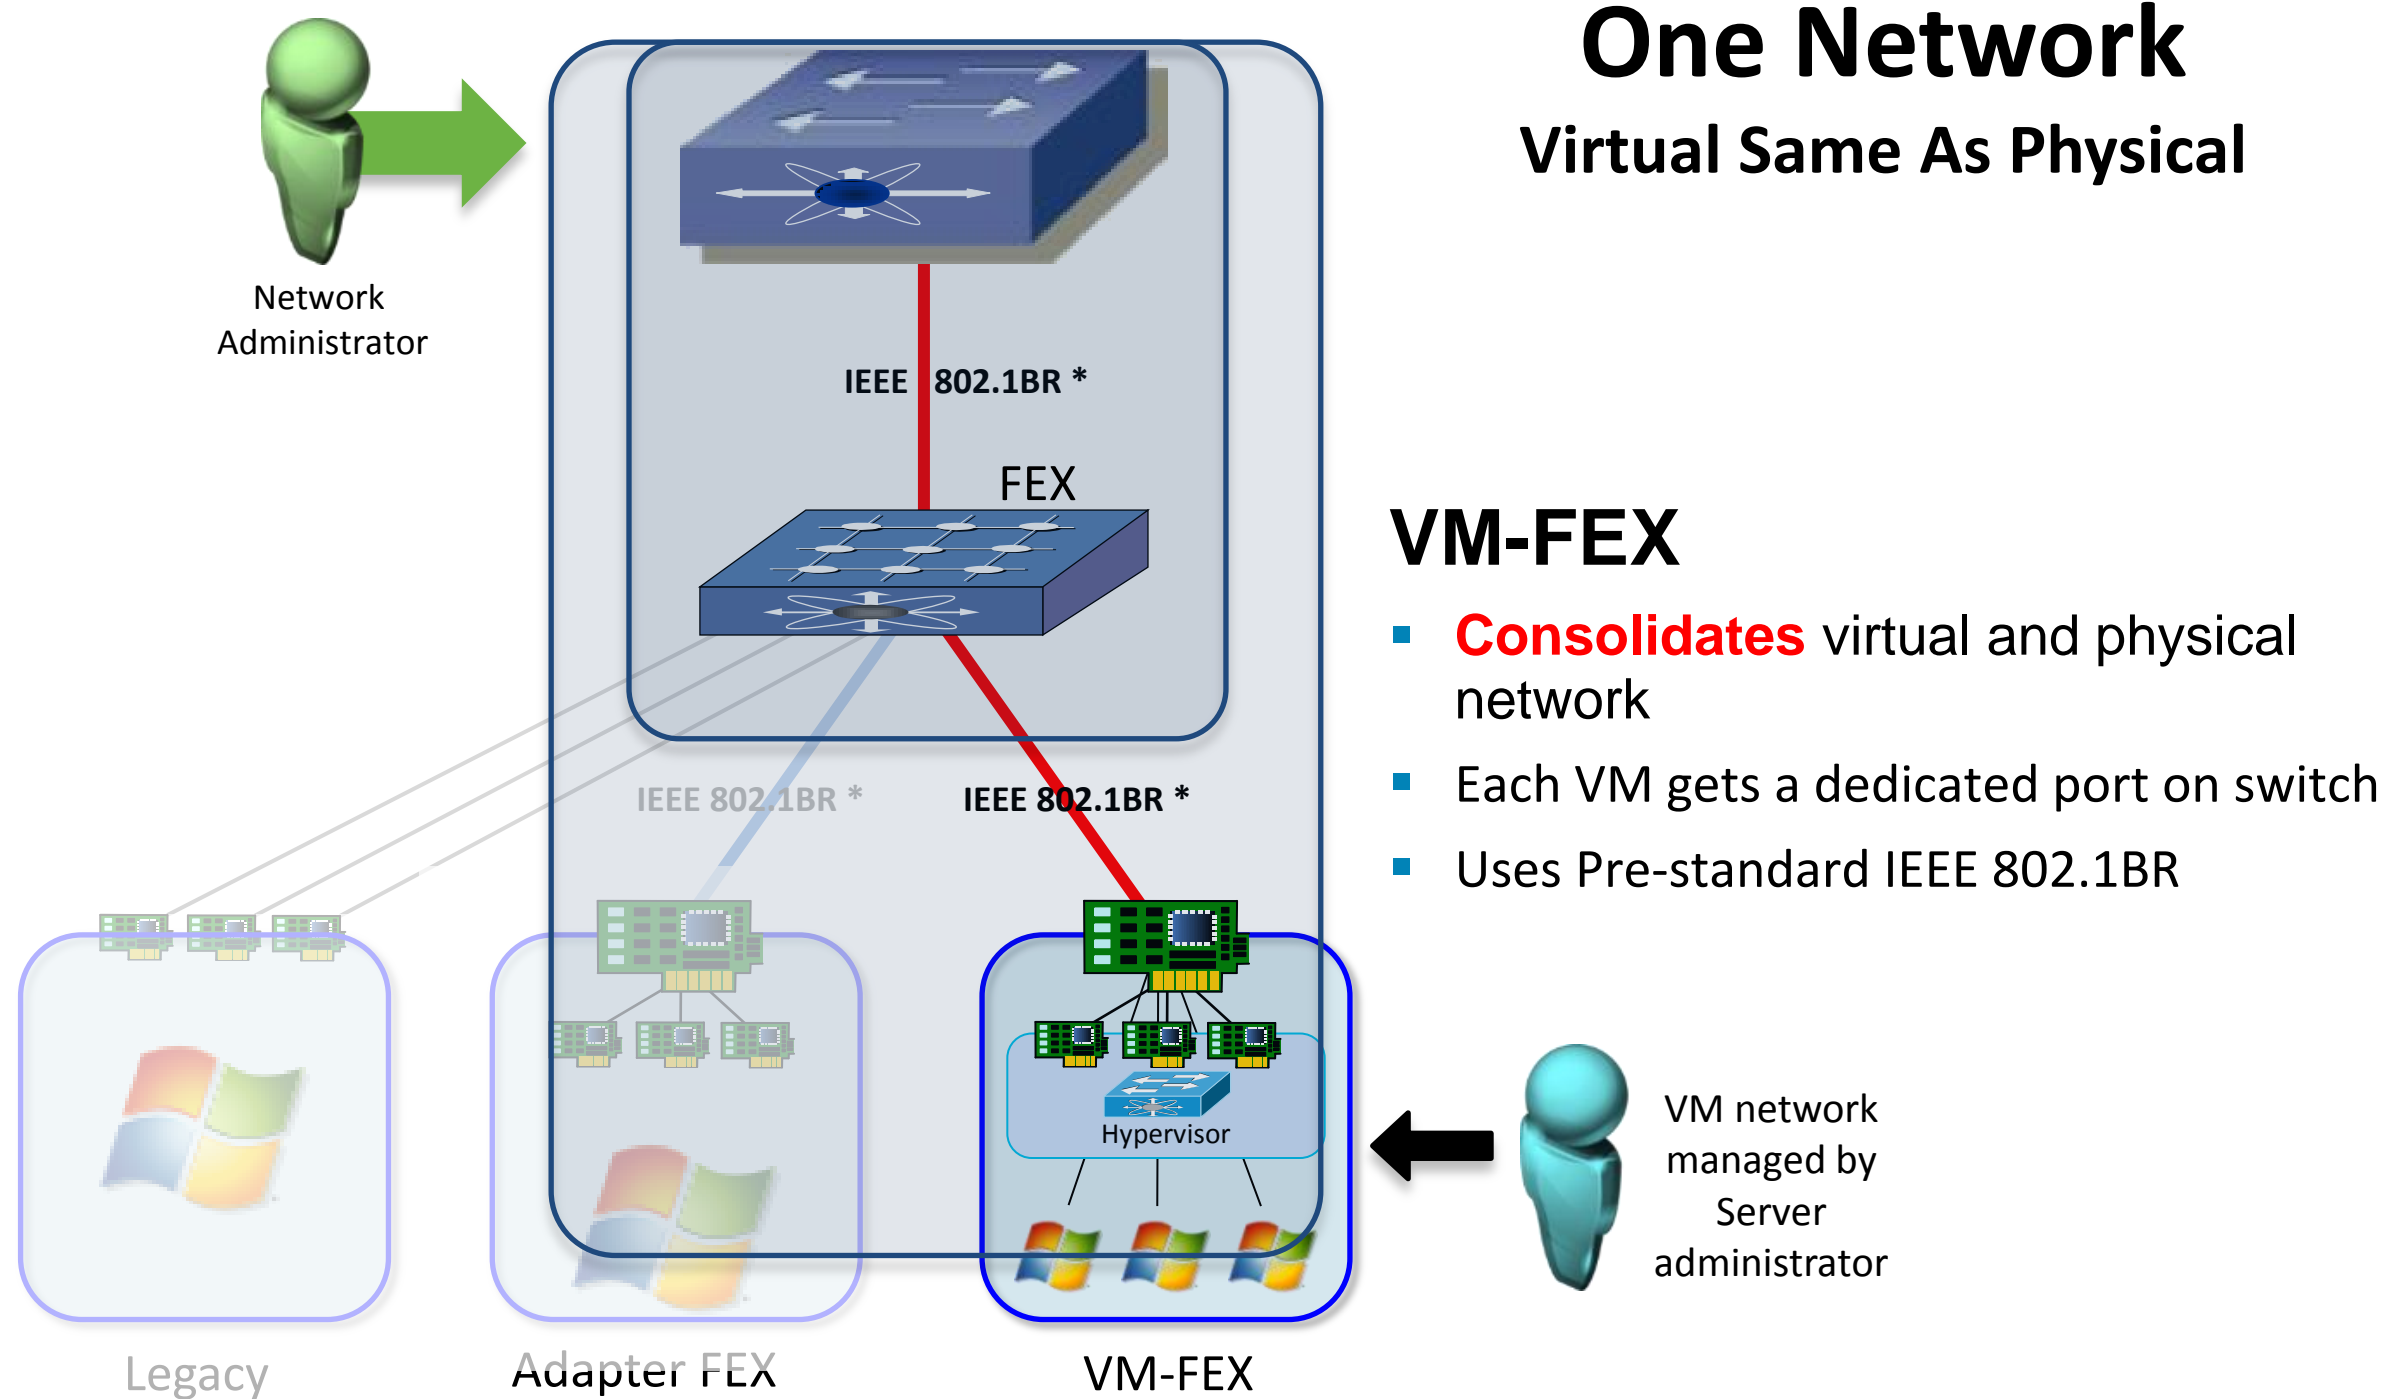

Parent Switch to Adapter

### Adapter FEX

- **Consolidates** multiple 1Gb interface into a single 10Gb interface
- Extends network into server
- Uses Pre-standard IEEE 802.1BR

# Fabric Extender Evolution

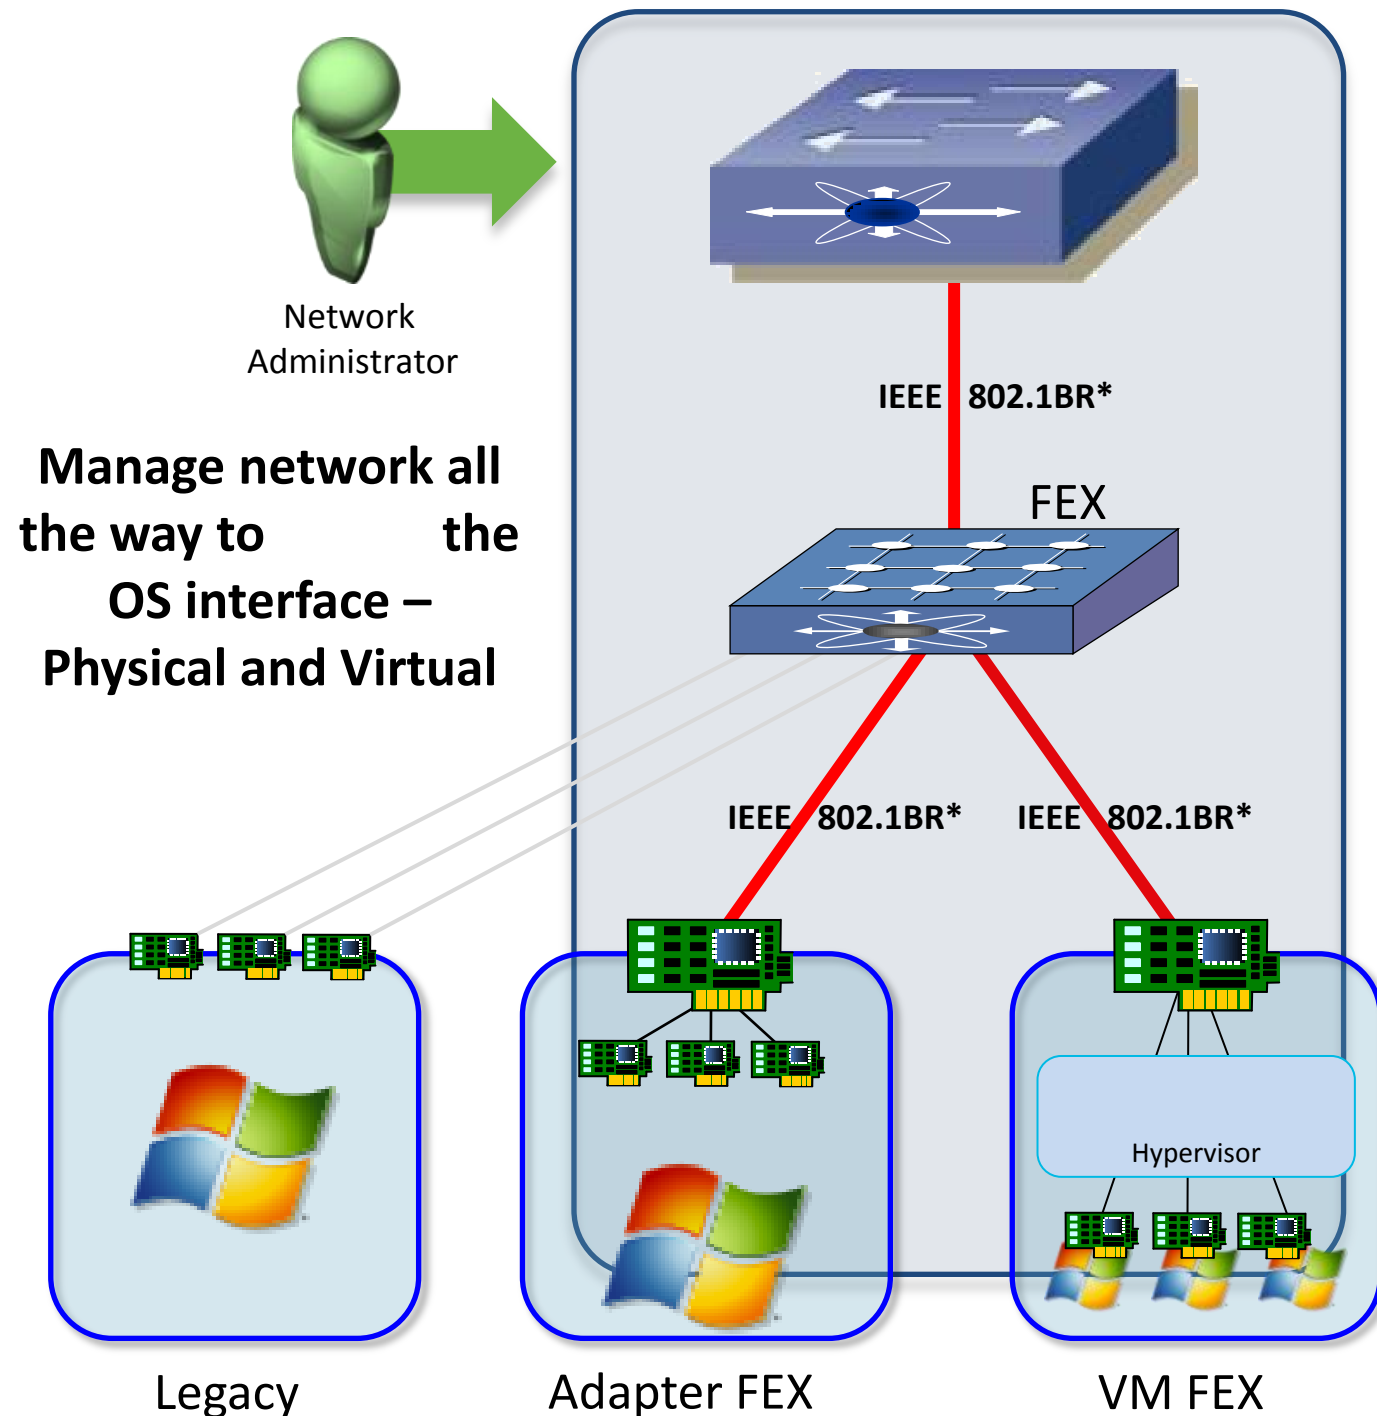

Distributed Modular System to the ToR, Server, and Virtual Machine



\*IEEE 802.1BR Pre-Standard

# Fabric Extender Evolution

Distributed Modular System to the ToR, Server, and Virtual Machine



\* IEEE 802.1BR Pre-Standard

## One Network

Parent Switch to Application  
Single Point of Management

### FEX Architecture

- **Consolidates** network management
- FEX managed as line card of parent switch

### Adapter FEX

- **Consolidates** multiple 1Gb interface into a single 10Gb interface
- Extends network into server

### VM-FEX

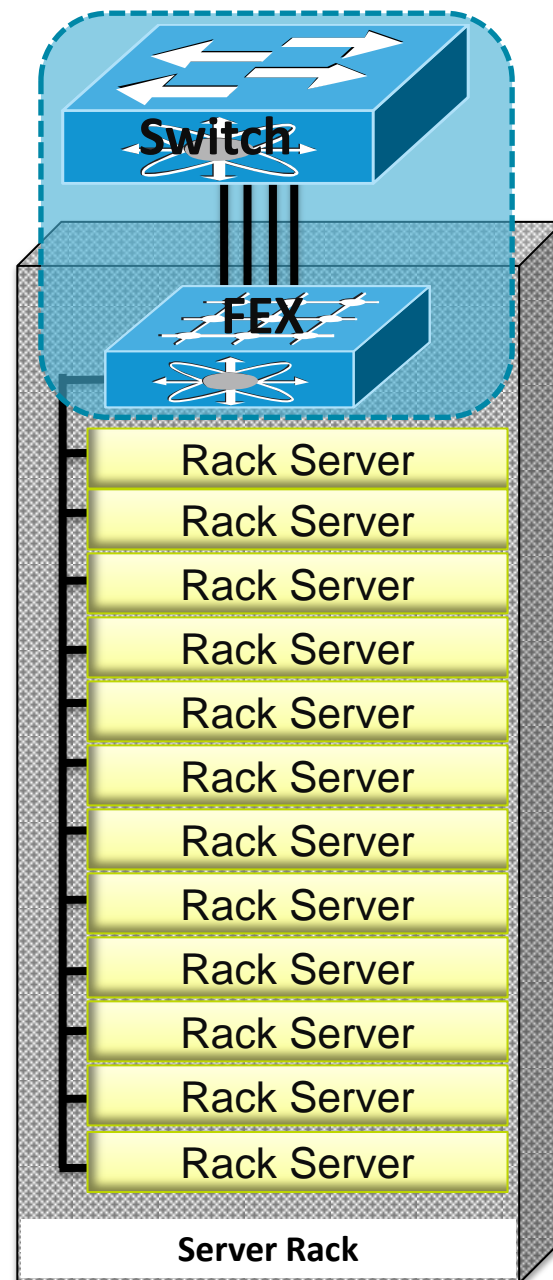
- **Consolidates** virtual and physical network
- Each VM gets a dedicated port on switch

# 4 Deployments of Cisco's FEX Technology

1

## Rack FEX

Nexus 5000/5500/7000  
+ Nexus 2200

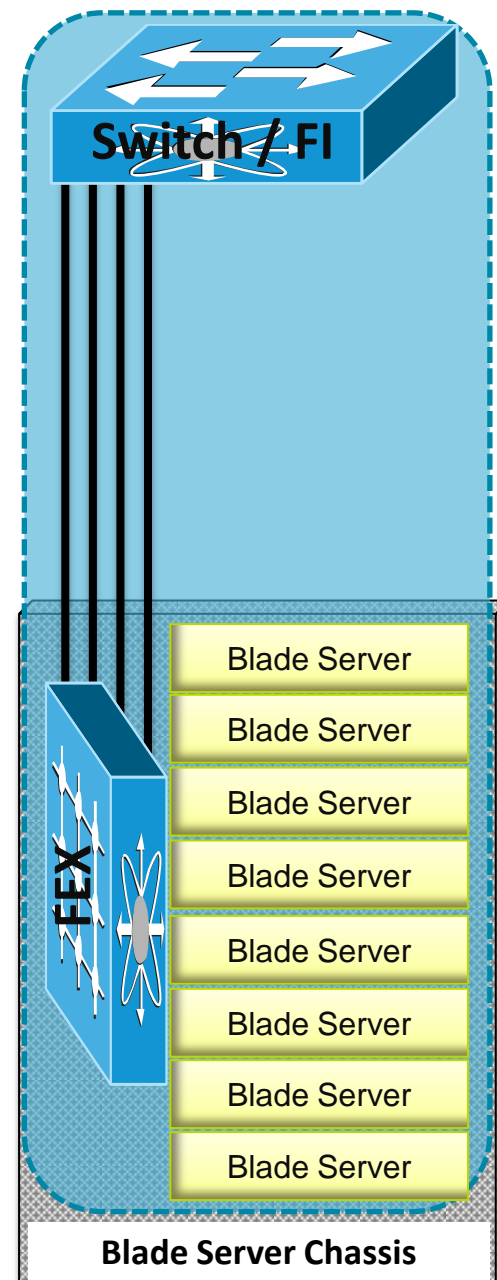


BRKVIR-2931

2

## Chassis FEX

UCS 6100/6200 + IOM 2k  
B22H with Nexus 5500 (HP)




© 2013 Cisco and/or its affiliates. All rights reserved.

3

## Adapter FEX

UCS 6100/6200 + VIC 1 or 2  
Nexus 5500 + VIC P81E

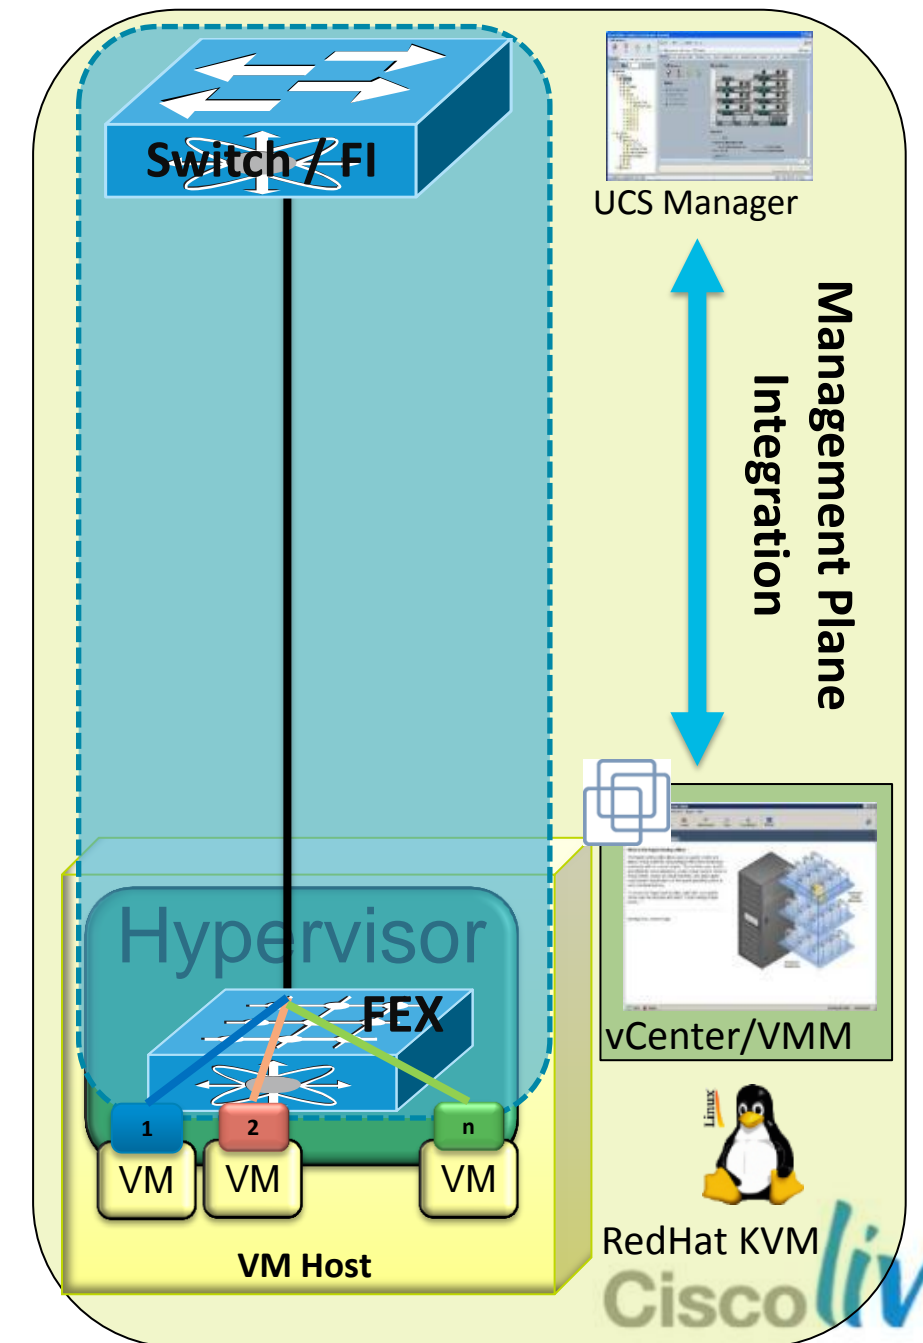


Cisco Public

4

## VM-FEX

UCS 6100/6200 + VIC 1 or 2 + VM Mgmt Link  
Nexus 5500 + VIC P81E

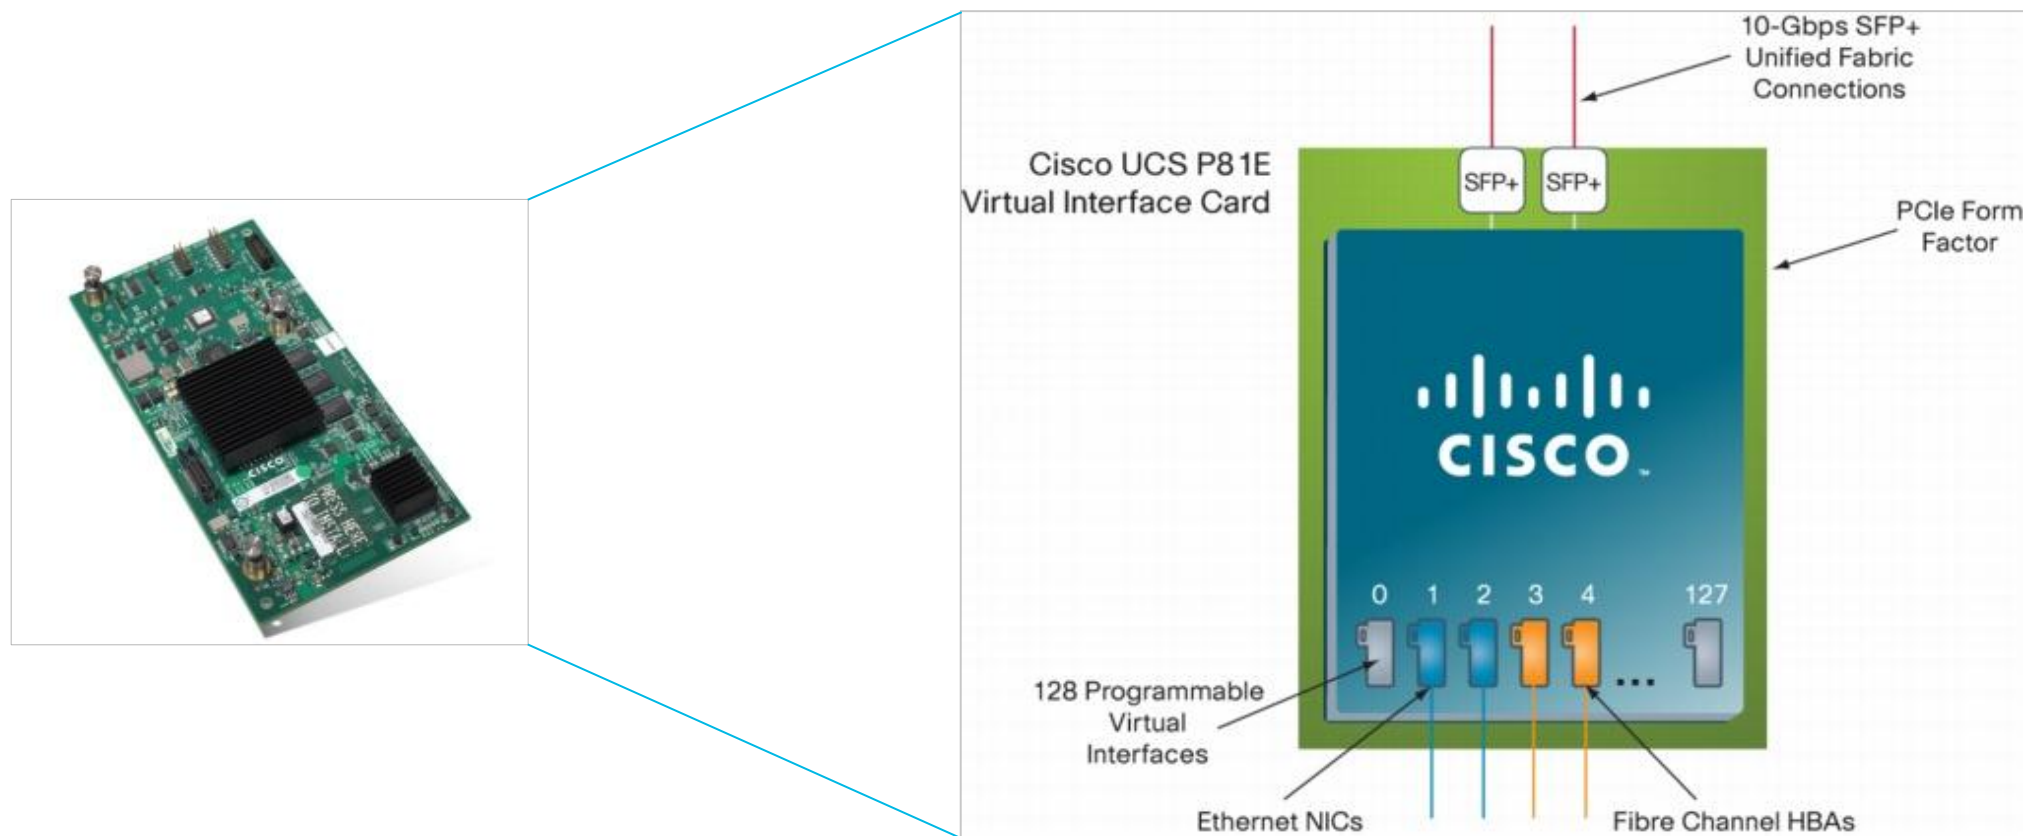


CiscoLive!



# Cisco UCS C-Series Adapter-FEX and VM-FEX

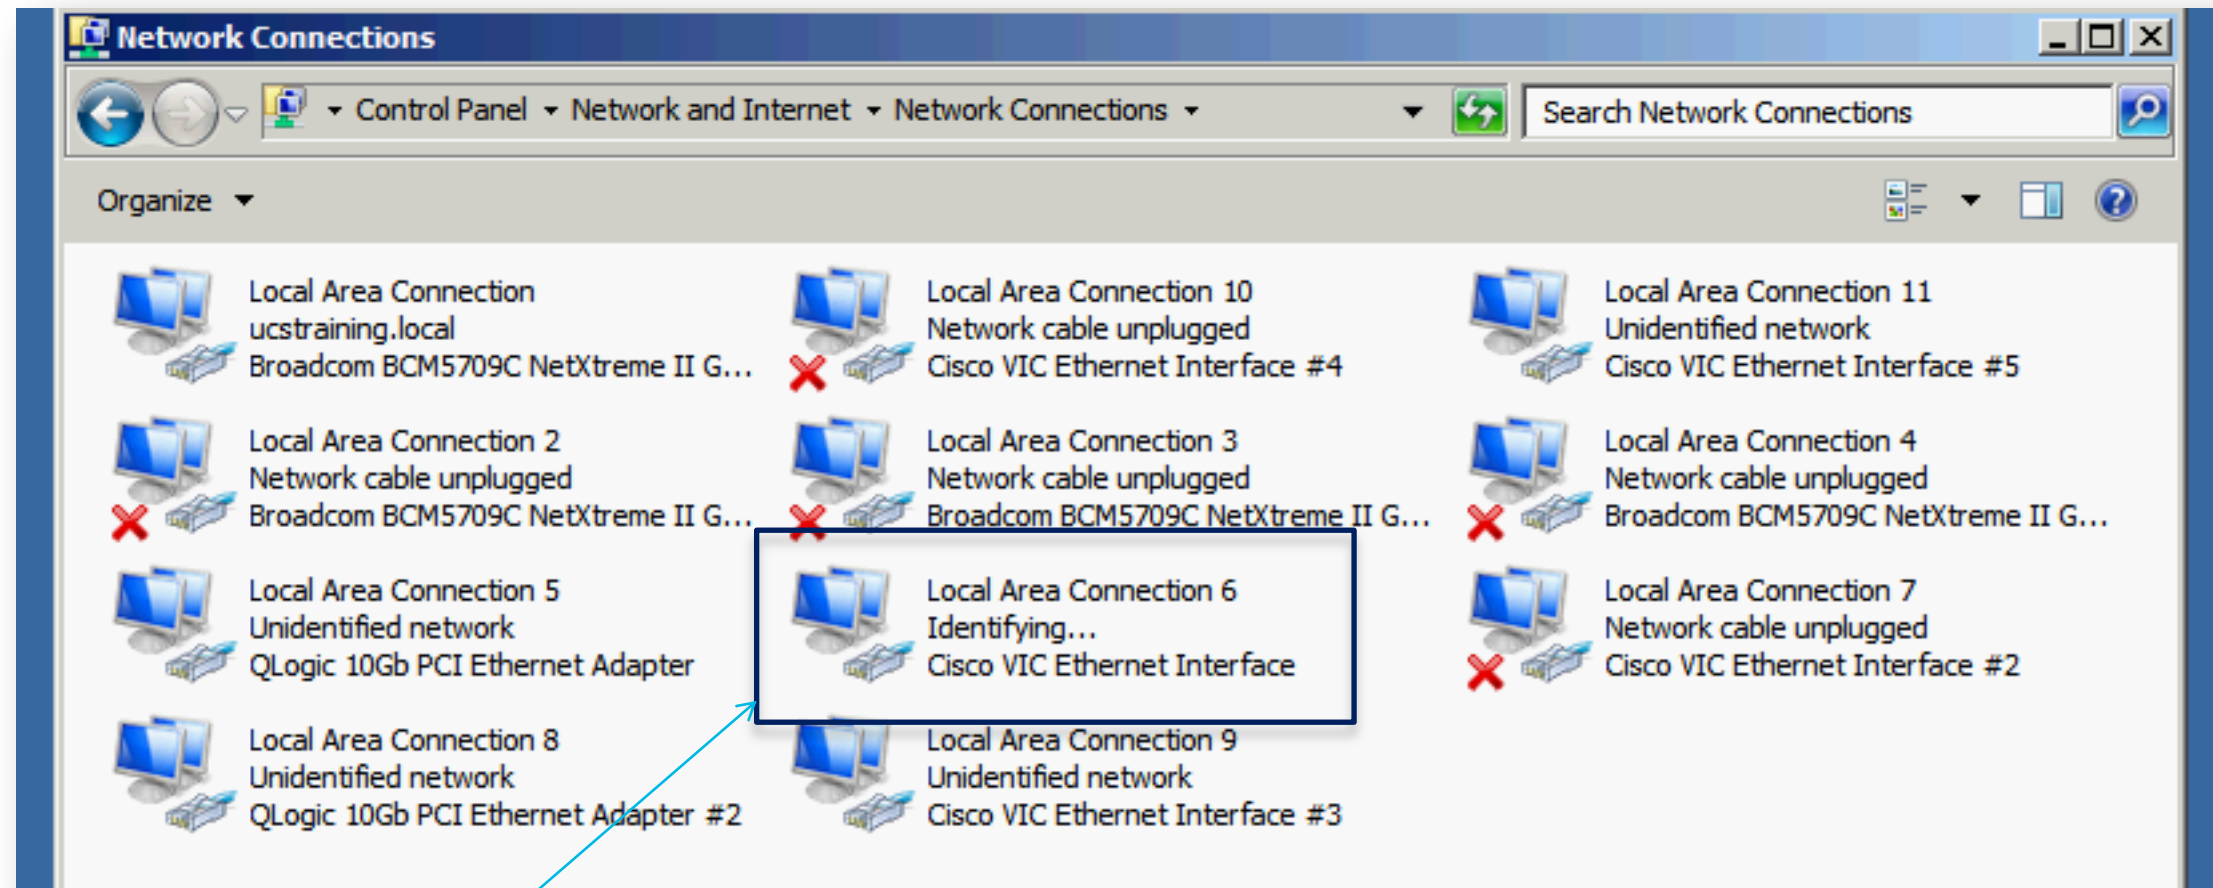
## UCS P81E Virtual Interface Card



- It supports NIC partitioning to the OS and 802.1BR to the switch
  - In Adapter-FEX mode: support for up to 16 Eth vNIC and 2 FC vHBA
  - In VM-FEX mode: support for up to 96 vNics
- Adapter Failover feature: in case of failure on the primary path, the vNIC is mapped to the standby port transparently to the OS
- Security and scalability improvements: no need of trunking all VLANs to the server interface

# Adapter-FEX at UCS C-Series Servers

Network admin can control the veth configuration and, as a result, the server network adapter



Support matrix at Nexus 5500 (NX-OS 5.1(3)N1(1)) and UCS C-Series Servers

```
Nexus-5548(config)# int veth6
Nexus-5548(config-if)# shut
Nexus-5548(config-if)# no shut
```

A-FEX	Adapter	Server	OS	Features	Scalability
Cisco	Cisco UCS P81E VIC	UCS C series	Windows, Linux RH, ESX	vNic, vHBA	16 vNIC + 2 vHBAs per adapter, 1K per N5500 pair
Broadcom	BCM57712 CNA	UCS C series Third party*	Windows, Linux RH, ESX	vNic, vHBA	Max of 8 Virtual Interfaces Up to 8 vNICs, up to 2 vHBA (vNICS + vHBA <=8), 1K per N5500 pair




# UCS B-Series: Virtual Interface Card (VIC) 1280

## Customer benefits

- Dual 4x 10 GE (80 Gb per host)
- VM-FEX scale, up to 112 VM interfaces /w ESX 5.0

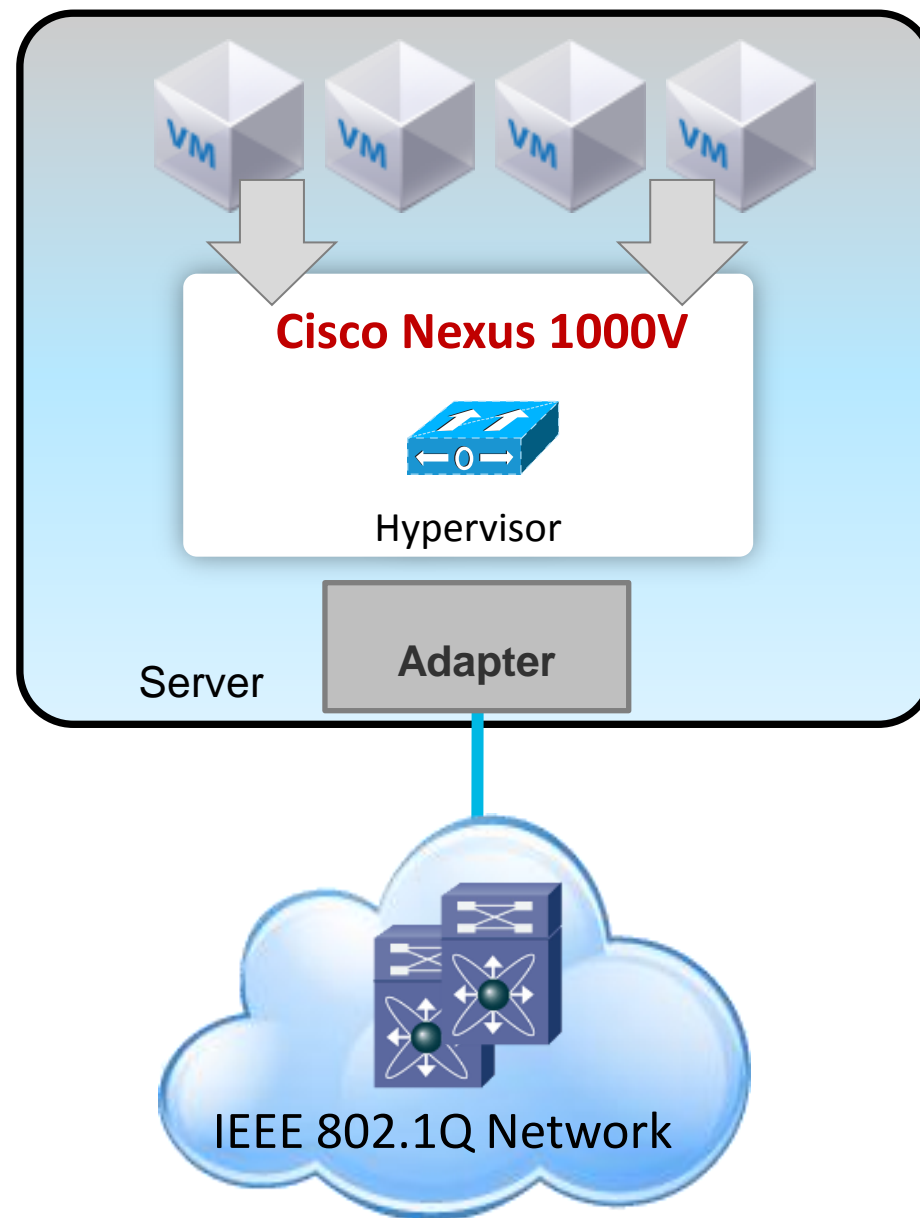
## Feature details

- Dual 4x 10 GE port-channels to a single server slot
- Host connectivity PCIe Gen2 x16
- PCIe Gen 2 x16 bandwidth limit is 32 Gbps
- HW Capable of 256 PCIe devices
  - OS restriction apply
- PCIe virtualisation OS independent (same as M81KR)
- Single OS driver image for both M81KR and 1280 VIC
- FabricFailover supported
- Eth hash inputs : Source MAC Address, Destination MAC Address, Source Port, Destination Port, Source IP address, Destination IP address, and VLAN
- FC Hash inputs: Source MAC Address, Destination MAC Address, FC SID and FC DID

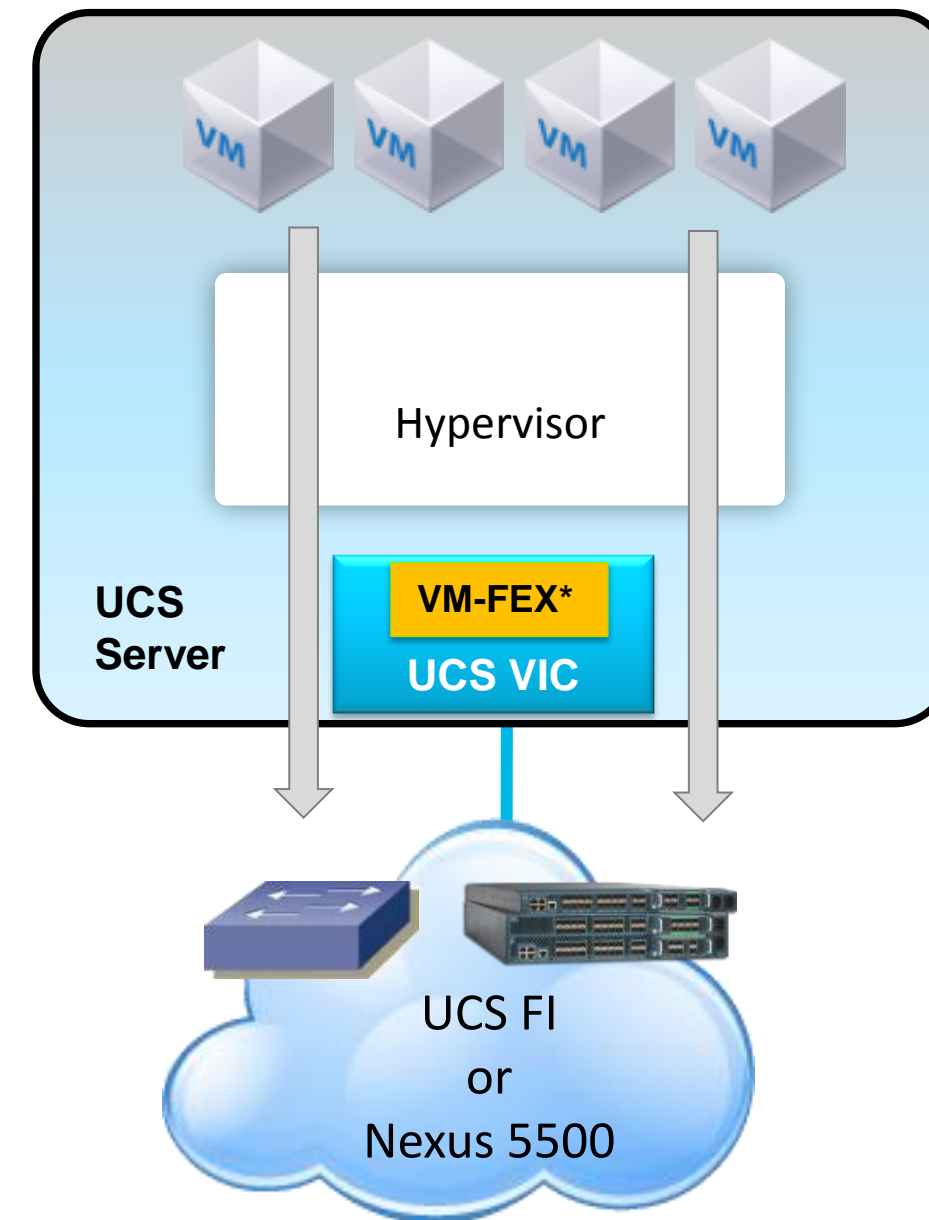


# Cisco Nexus 1000V and UCS VM-FEX

Bring network to the hypervisor  
(Cisco Nexus 1000V Switch)



Bring VM awareness to physical network  
(Cisco UCS VM-FEX)



# Cisco Virtual Machine Networking Options

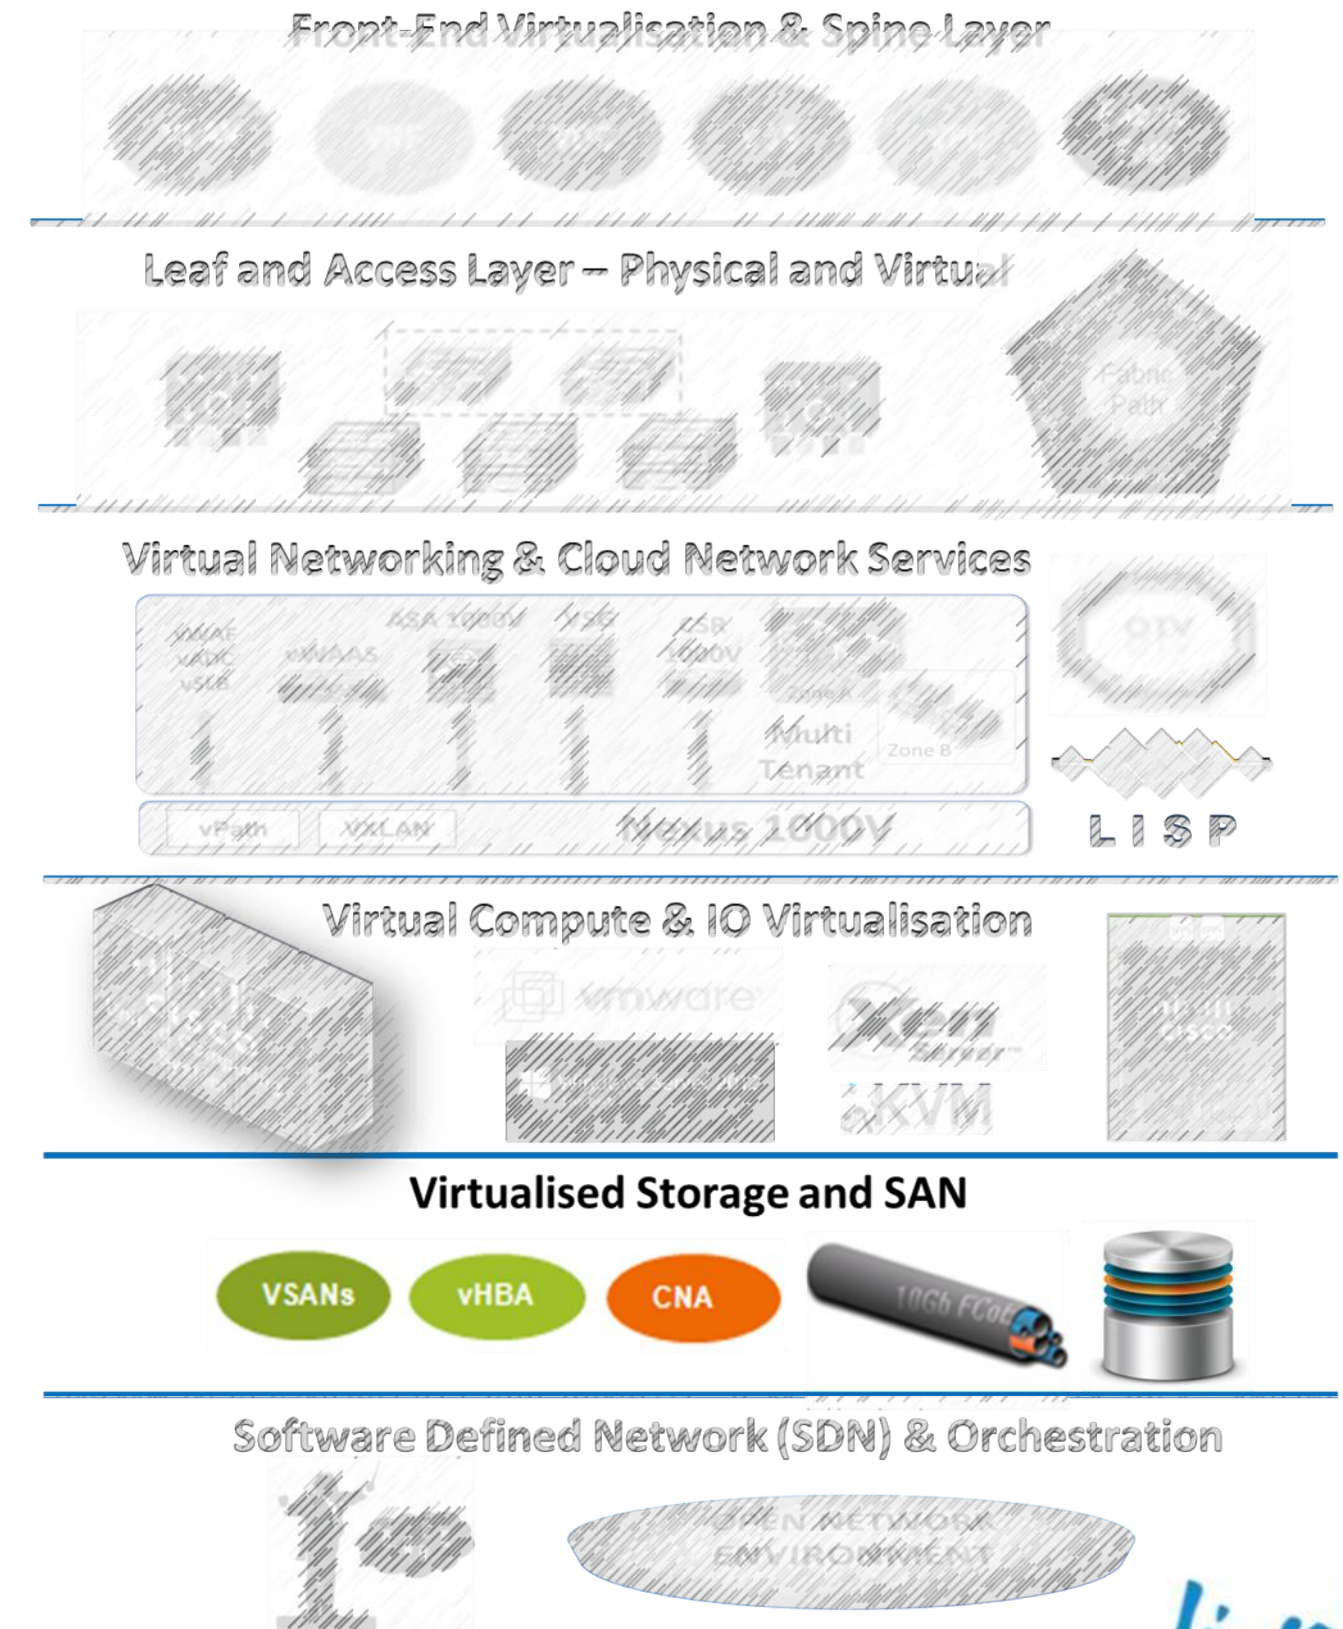
Criteria	Nexus 1000V	Nexus 5500 VM-FEX	UCS VM-FEX
Server + Network management			Y
Heterogeneous Servers	Y	Y (future)	
Heterogeneous 1 <sup>st</sup> Hop Network	Y		
Physical/Virtual Network Consolidation		Y	Y
Feature Richness/Velocity	Y		
Hardware Performance		Y	Y
VM Density per Server	216 VMs	VIC P81E: 96 VMs	VIC M81KR: 56 VMs VIC 1280: 116* VMs

Model	Total Interfaces	Interface Type	Ethernet Interfaces	Fibre Channel Interfaces	VM-FEX	Failover Handling
<a href="#">Cisco UCS Virtual Interface Card 1280</a>	256	Dynamic	0-256	0-256	Hardware/Software	Hardware, no driver needed
<a href="#">Cisco UCS P81E Virtual Interface Card</a>	128	Dynamic	0-128	0-128	Hardware/Software	Hardware, no driver needed
<a href="#">Cisco UCS M81KR Virtual Interface Card</a>	128	Dynamic	0-128	0-128	Hardware/Software	Hardware, no driver needed

\* Limitation of OS (VIC 1280 hardware is capable of up to 256)

# Agenda

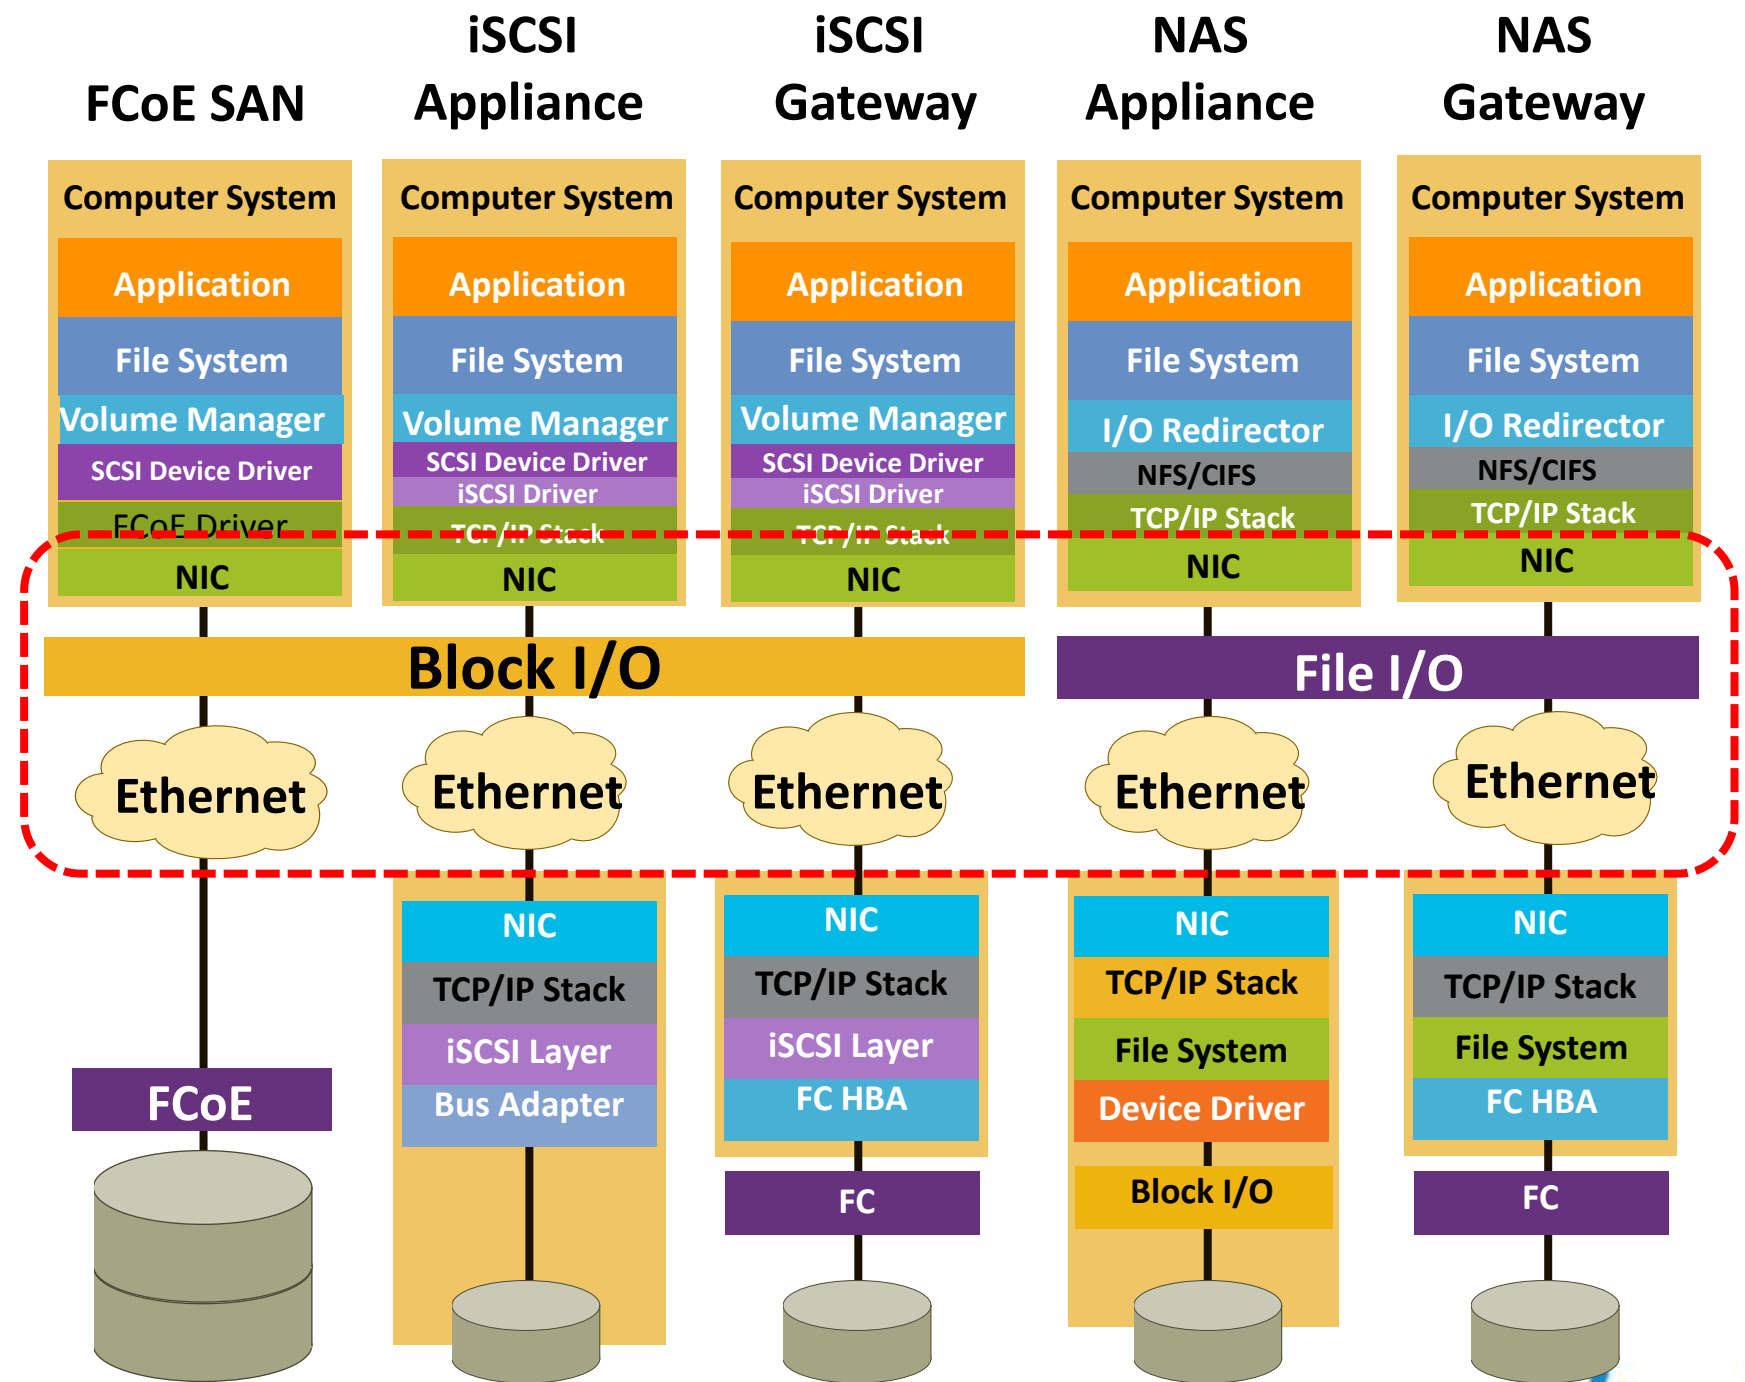
- **Data Centre Virtualisation Overview**
- **Front-End Data Centre Virtualisation**
  - DC Edge Layer
  - Aggregation / Spine Layer
  - Access / Leaf Layer
- **Virtual Networking & Cloud Network Services**
  - Virtual Access Layer: Nexus 1000v
  - **Cloud Network Services**
  - **Inter DC Services: OTV & LISP**
- **Virtual Compute & IO Virtualisation**
  - Unified Computing System
  - Hypervisors & Virtual IO @ Cisco UCS
- **Virtualised Storage & SAN**
  - **Unified Fabric**
  - **Virtualised SAN and Storage**
- **Software Defined Network & Orchestration**
  - SDN & Cisco Open Network Environment (ONE)
  - Cloud Orchestration
- **Implementation Examples**
- **Q&A**



# Unified Fabric

## Storage I/O—Flexibility and Serialised Re-Use

- Ability to re-provision any compute unit to leverage any access method to the data stored on the ‘spindle’
- Serialised Re-Use – Boot from SAN and Run from NAS
- Virtualisation requires that the Storage Fabric needs to exist everywhere the IP fabric does



# Network Behaviour and Characteristics

- **Ethernet is non-deterministic.**
  - Flow control is destination-based
  - Relies on TCP drop-retransmission / sliding window
- **Fibre-Channel is deterministic.**
  - Flow control is source-based (B2B credits)
  - Services are fabric integrated (no loop concept)

“Lossless Ethernet” became a requirement =>  
Data Centre Bridging (DCB)

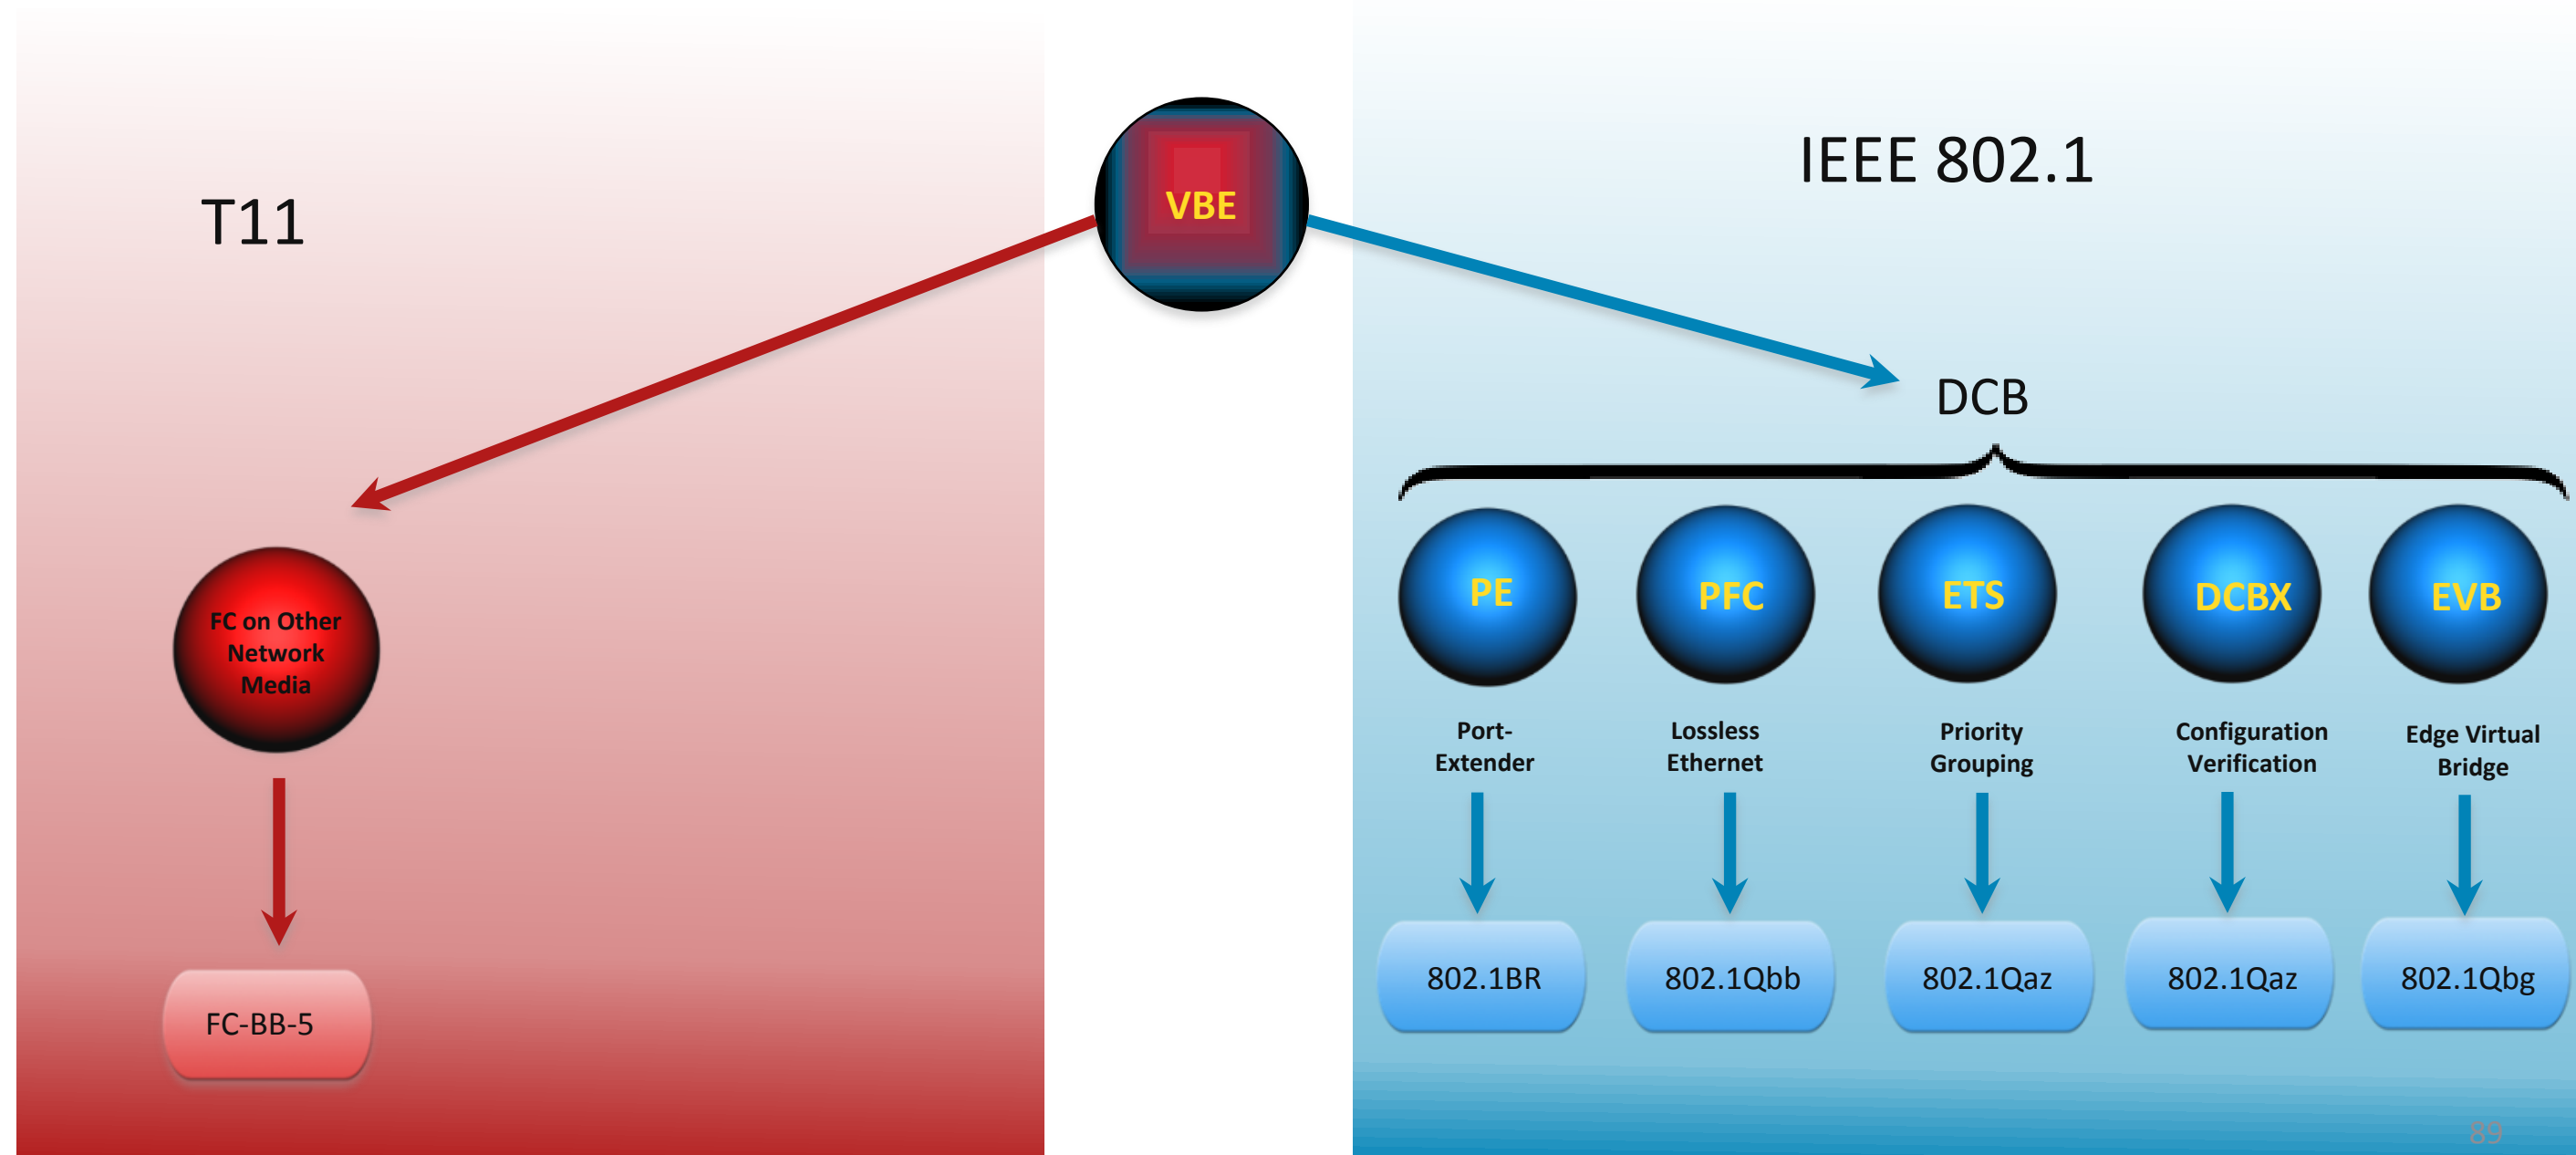






# Standards for Unified I/O with FCoE and VBE

## I/O Consolidation with FCoE

- FCoE is *fully* defined in FC-BB-5 standard
- FCoE works alongside additional technologies to make I/O Consolidation a reality



# Unified I/O Architecture Consolidation Initial Goal



# Example: Embedded FCoE at Cisco UCS

Cisco UCS



From ad hoc and inconsistent...




...to structured, but siloed, complicated and costly...



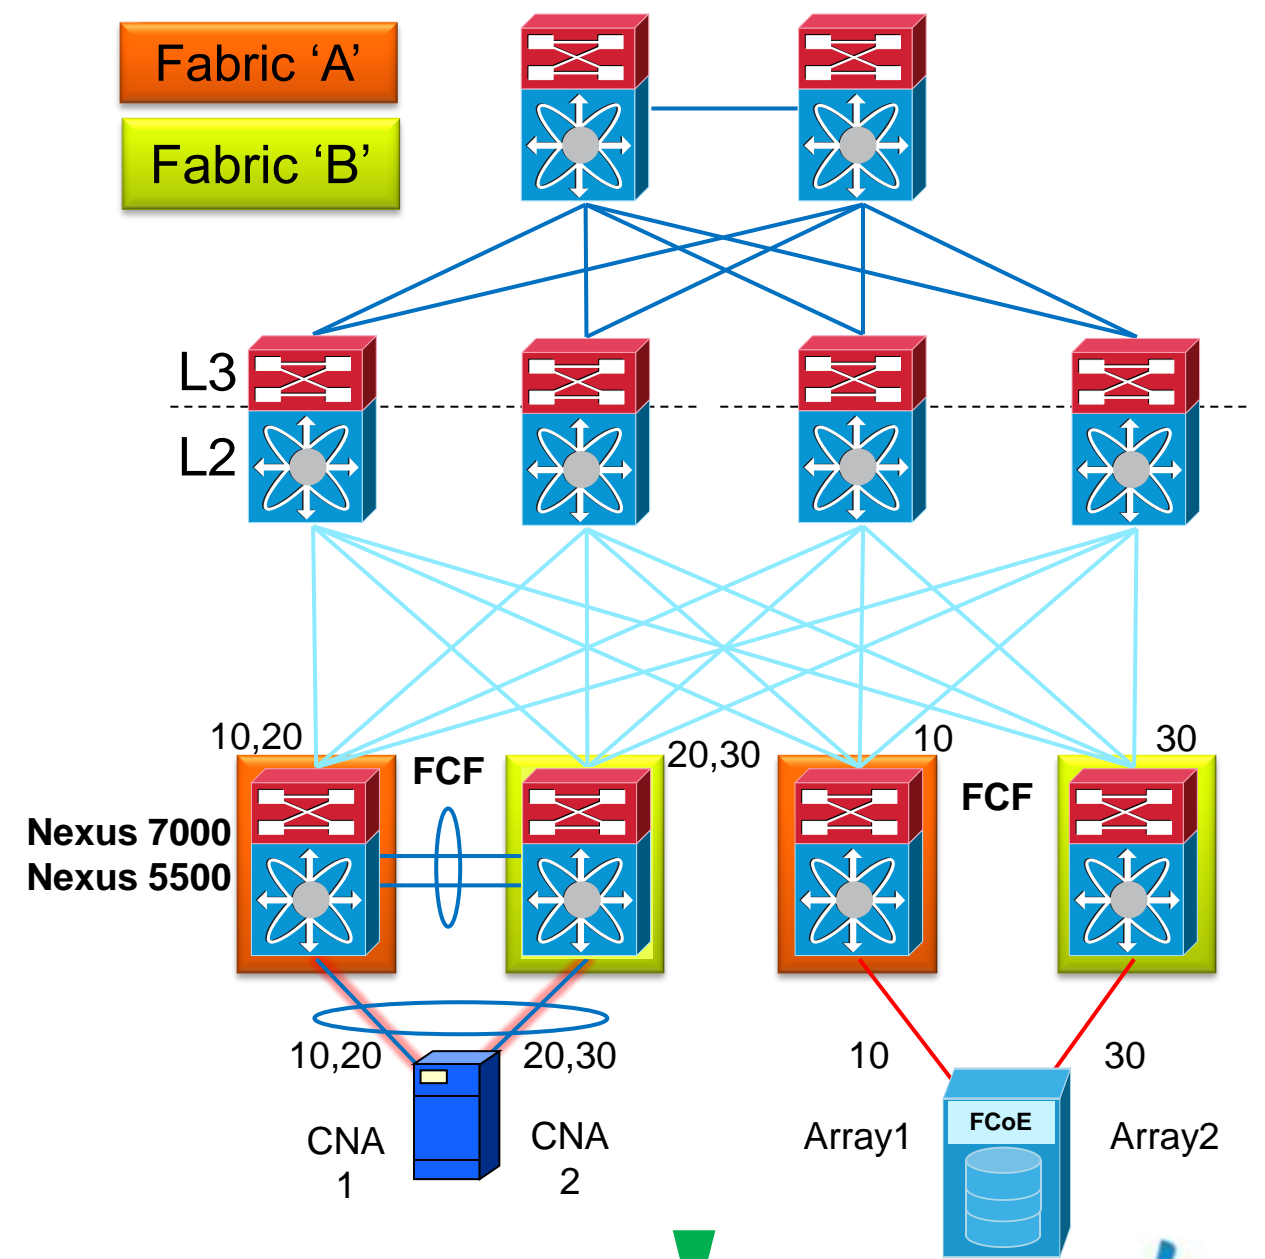
...to simple, optimised and automated

# Looking forward: Full Unified Fabric

## SAN Separation at the Access Switch (and FC-BB-6, when available)



- ✓ LAN and SAN traffic share physical switches and links
- ✓ FabricPath enabled
- ✓ All Access switches are FCoE FCF switches
- ✓ VE\_Ports to each neighbour Access switch
- ✓ Single process and database (FabricPath) for forwarding
- ✓ Improved (N + 1) redundancy for LAN & SAN
- ✓ Sharing links increases fabric flexibility and scalability
- ✓ Distinct SAN 'A' & 'B' for zoning isolation and multipathing redundancy
- ✓ With FC-BB-6 on the future:
  - Access switches for server connectivity to behave as FDF
  - FDF to FCF transparent failover
  - VA\_Ports to each neighbour FCF switch



Isolation



Convergence



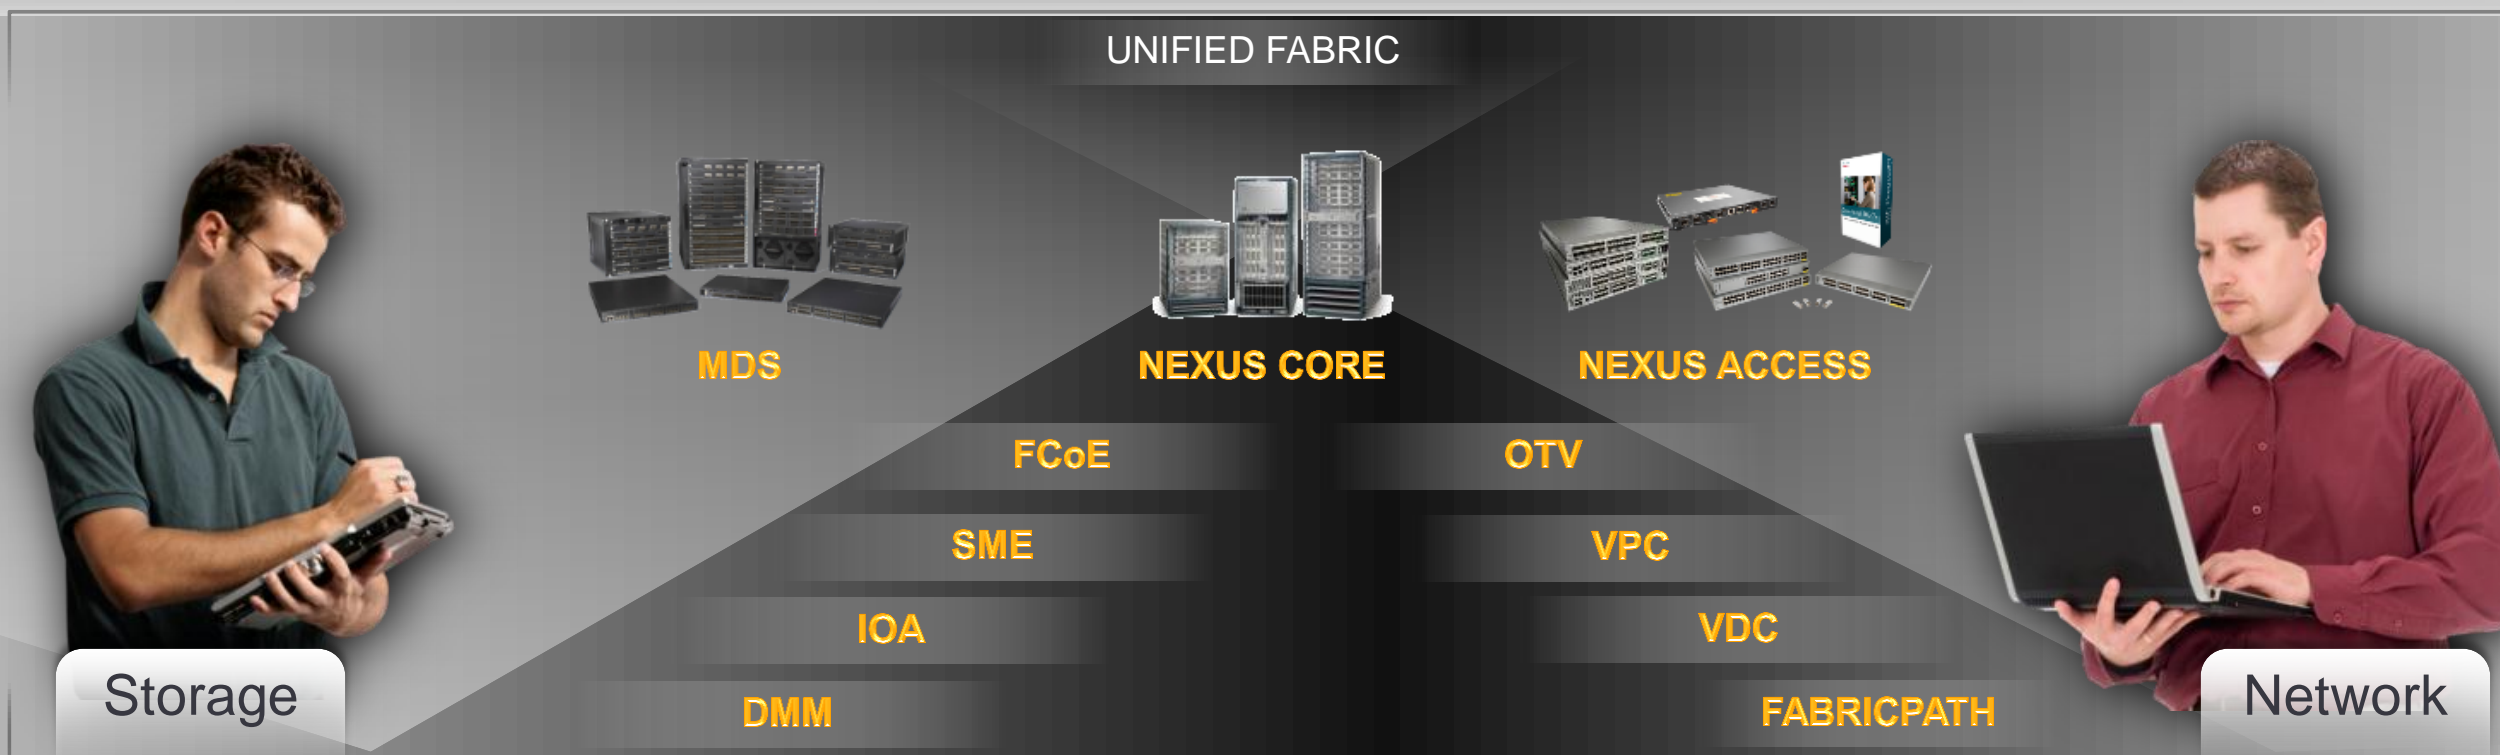
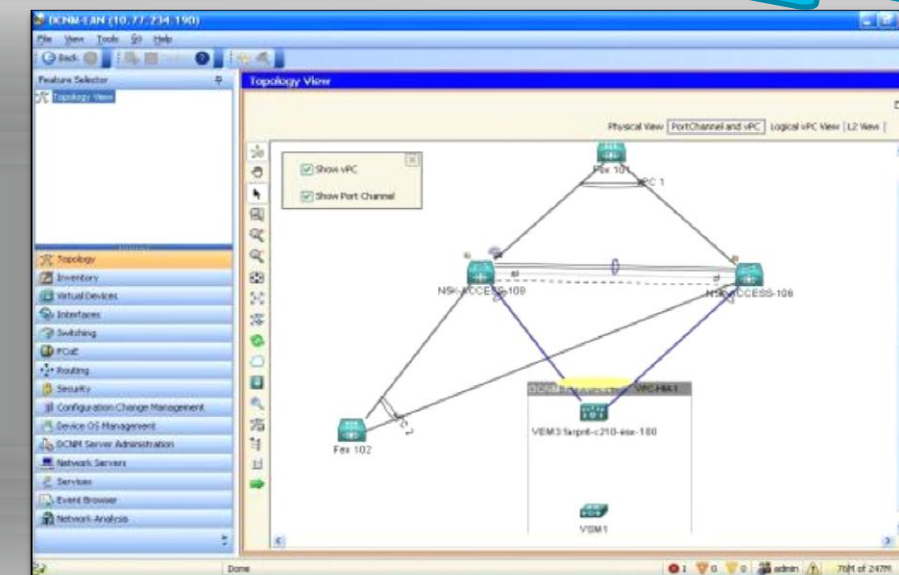
# Converged Cisco Data Centre Network Manager

Available Today !!!

## Cisco Prime Data Centre



- Fault Management
- Resource & Change Mgmt
- Virtual Machine level integration
- Provisioning/Configuration
- API Integration
- Security RBAC




# Server-Based Storage for Cloud Scale-Out

- By 2015, annual shipments of enterprise storage are expected to exceed 175,000 PB<sup>1</sup>, of which approximately 60% can be addressed by Scale-Out architectures (server and non-server based)<sup>2</sup>.
- Enterprise Strategy Group estimates scale-out systems will make up 80% of new external storage shipments by 2015<sup>3</sup>.

ESG estimates annual revenue for scale-out architectures (including non-server based) will exceed \$18B by 2015.

Nexenta claims it has generated ~\$300M in revenue for its partners in CY2011<sup>4</sup>.

\$ / PB forecast based on IDC forecast of DAS storage \$ / GB pricing.



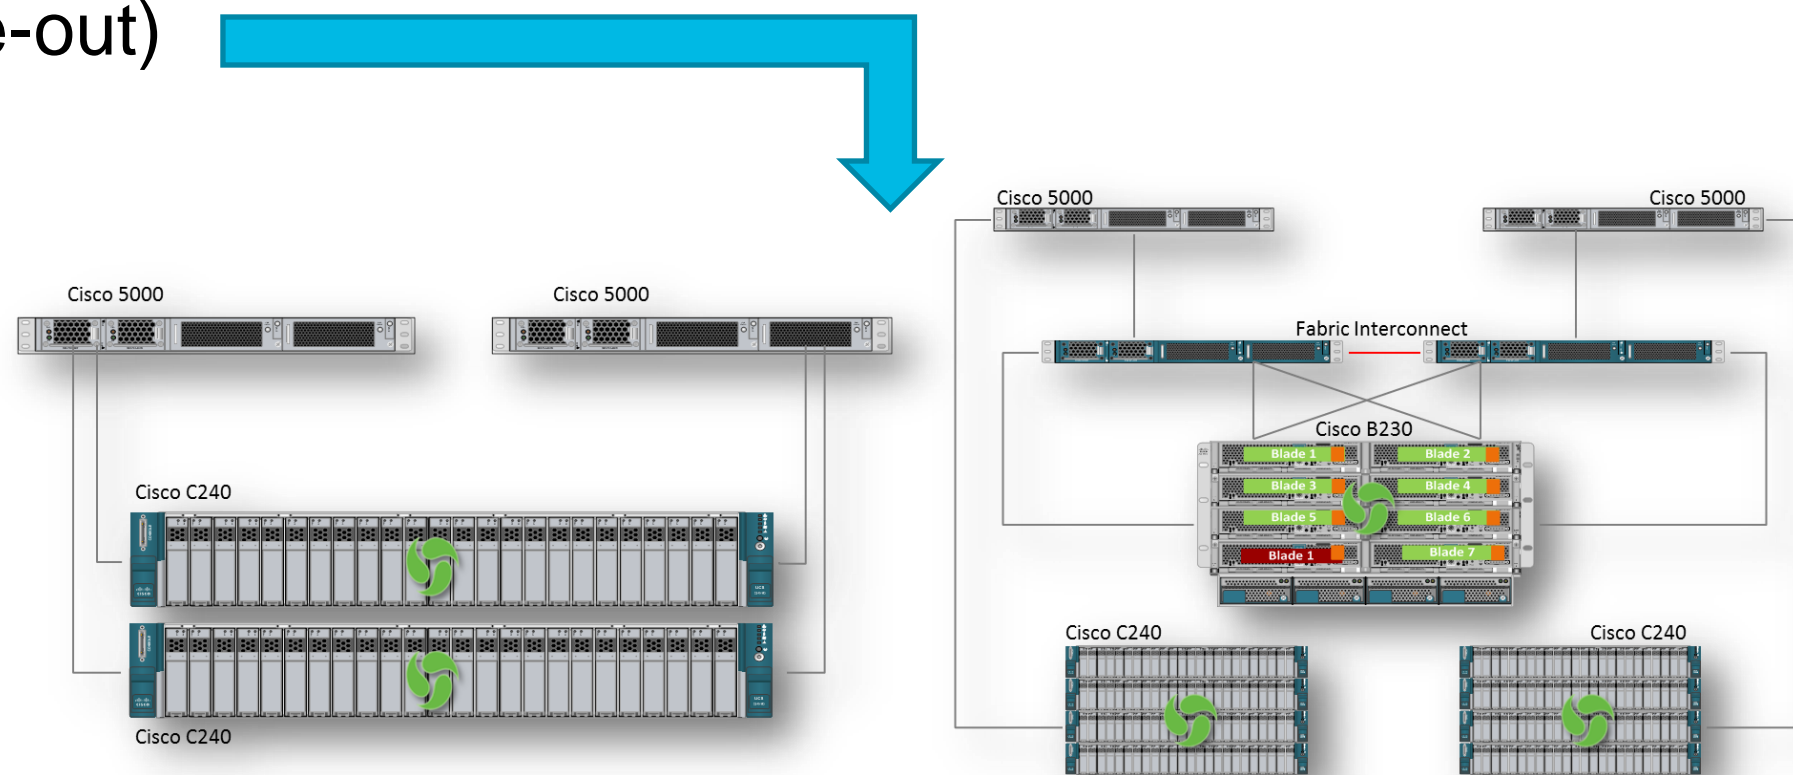
Assuming < 5% market penetration, Server-Based Scale-Out storage will be a **\$2.4B opportunity by 2015.**

Source: IDC, Enterprise Strategy Group, 451 Group, and Cisco CBD.

1. IDC, "WW Storage Forecast 2011-2015"
2. Assumed Public Cloud PB (enterprise & direct) to be addressable.
3. Current forecast assumes 5% technology penetration by 2015.
4. 451 Group, 4/27/2012

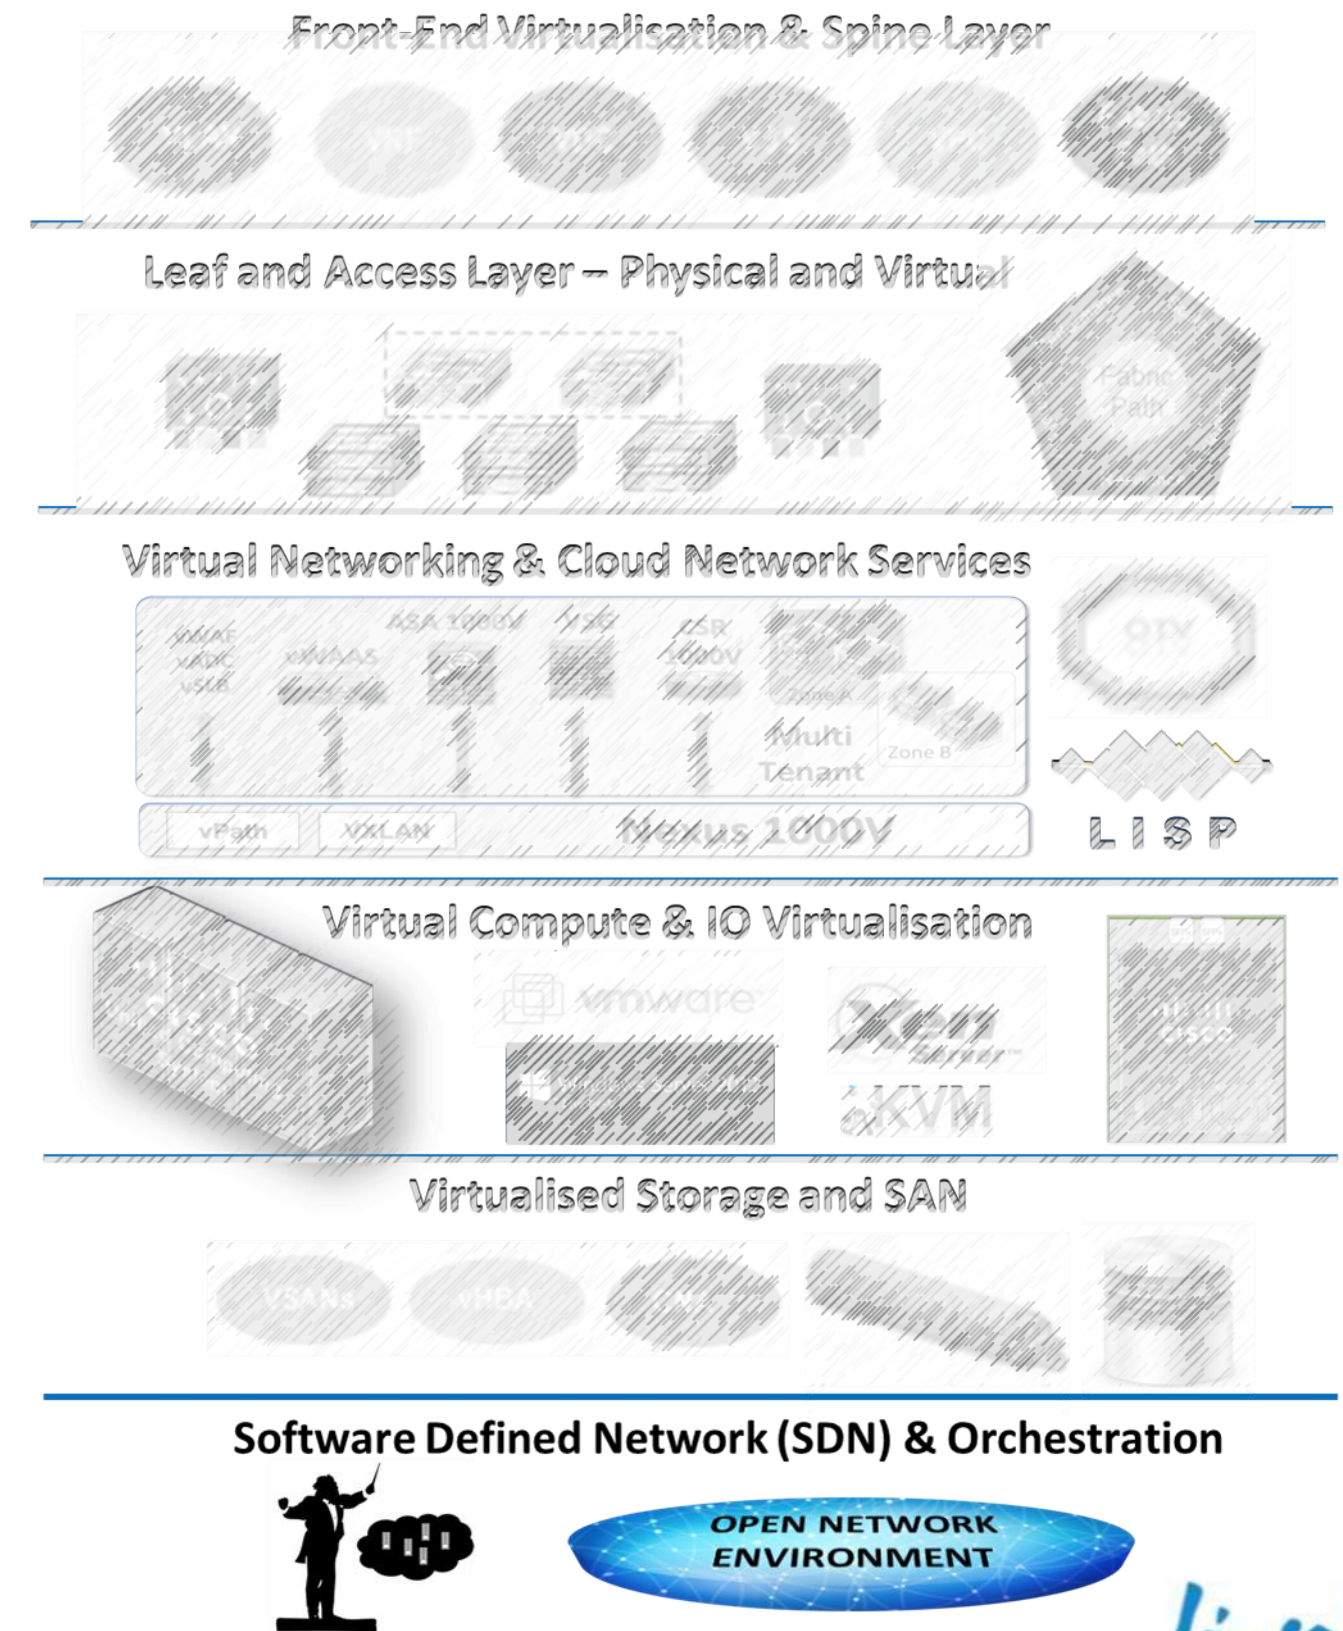
# Some Server-based Storage Solutions Validated with Cisco UCS

- StorMagic SvSAN (Storage Virtual Appliance)
- VMWare Virtual Storage Appliance (VSA)
- Nexenta (VM-aware and scale-out)
- Nimble Storage CS-Series
- NetApp OnTap Edge
- Red Hat Storage Server



# Agenda

- **Data Centre Virtualisation Overview**
- **Front-End Data Centre Virtualisation**
  - DC Edge Layer
  - Aggregation / Spine Layer
  - Access / Leaf Layer
- **Virtual Networking & Cloud Network Services**
  - Virtual Access Layer: Nexus 1000v
  - **Cloud Network Services**
  - **Inter DC Services: OTV & LISP**
- **Virtual Compute & IO Virtualisation**
  - Unified Computing System
  - Hypervisors & Virtual IO @ Cisco UCS
- **Virtualised Storage & SAN**
  - Unified Fabric
  - Virtualised SAN and Storage
- **Software Defined Network & Orchestration**
  - **SDN & Cisco Open Network Environment (ONE)**
  - **Cloud Orchestration**
- **Implementation Examples**
- **Q&A**






# Before We Get Started: “SDN-related” Definitions

## What Is Software Defined Network (SDN)?

“...Software Defined Networking (SDN) is a network design concept in which the network control plane is centrally accessible through an API to applications wanting to better use the network...”

Source: wikipedia



## What Is OpenFlow?

“...open standard that enables researchers to run **experimental protocols** in campus networks. Provides standard hook for researchers to run experiments, without exposing internal working of vendor devices...”

Source: [www.opennetworking.org](http://www.opennetworking.org)



## What is OpenStack?

**Open source software** for building public and private Clouds; includes Compute (Nova), Networking (Quantum) and Storage (Swift) services.

Source: [www.openstack.org](http://www.openstack.org)

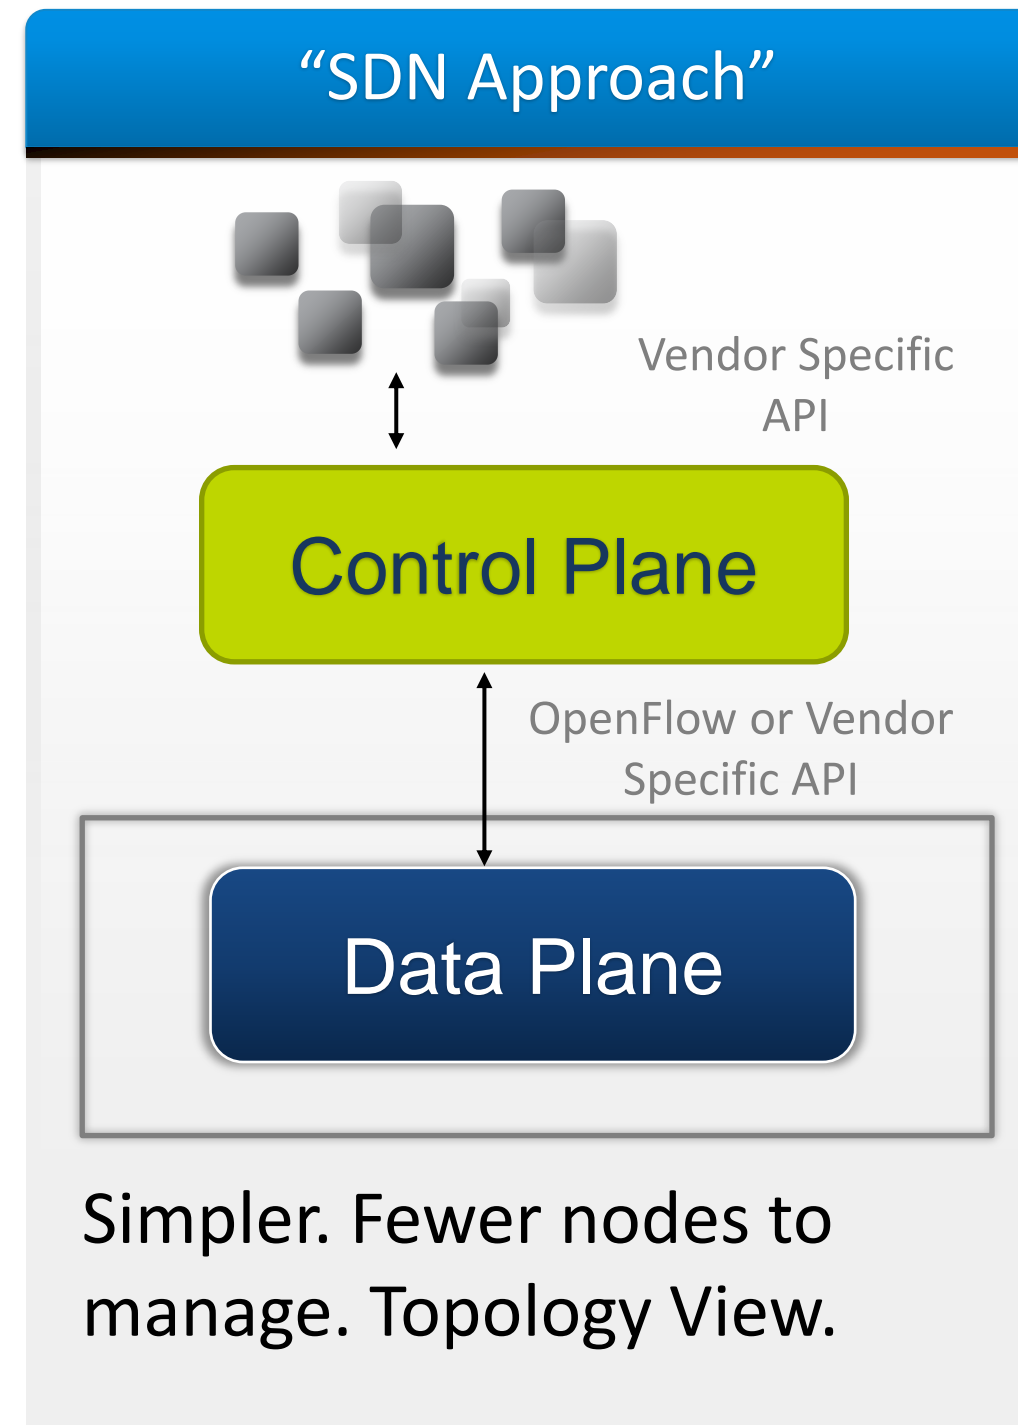
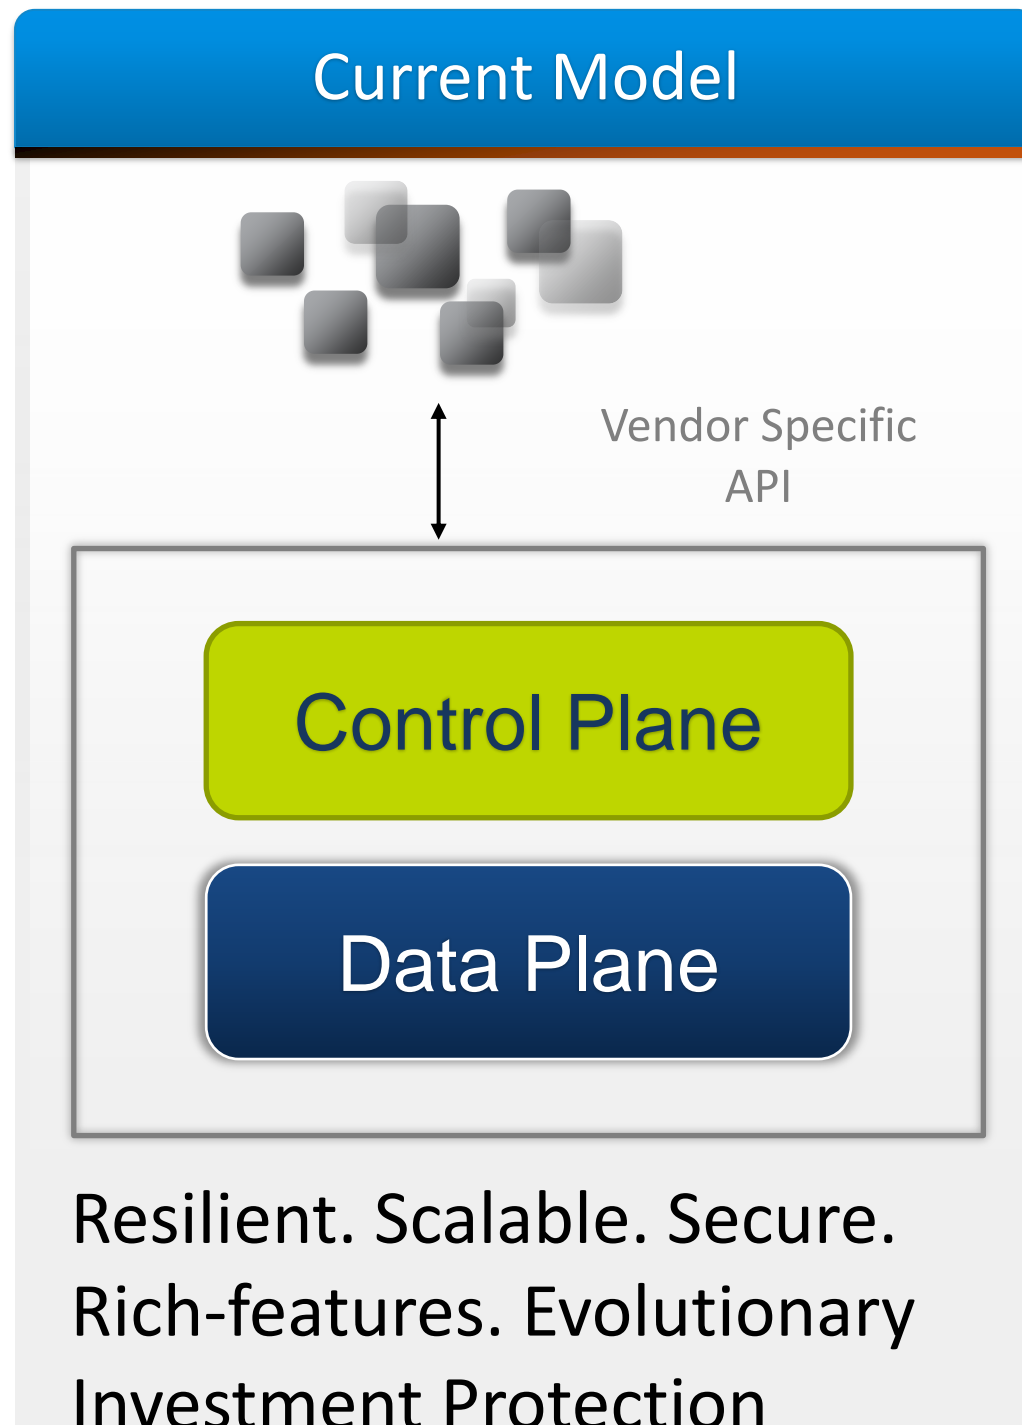


## What is Overlay Network?

Overlay network is created on existing network infrastructure (physical and/or virtual) using a network protocol. Examples of overlay network protocol are: MPLS, GRE, IPSEC, LISP, OTV and VXLAN

# Network Programmability Models

Physical or Virtual



# Cisco Software Defined Network (SDN) Strategy

Evolution of the Intelligent Network

## Preserve what's working

- Resiliency
- Scale
- Rich feature-set



## Evolve for emerging requirements

- Operational Simplicity
- Programmability
- Application Centricity

Evolve the Network for the next wave of application requirements

Cisco *live!*

# Cisco Open Network Environment (ONE)

Industry's Most Comprehensive Networking Portfolio

Hardware + Software

Physical + Virtual

Network + Compute

Applications

**OPEN NETWORK ENVIRONMENT**

Platform APIs

onePK

Virtual Overlays

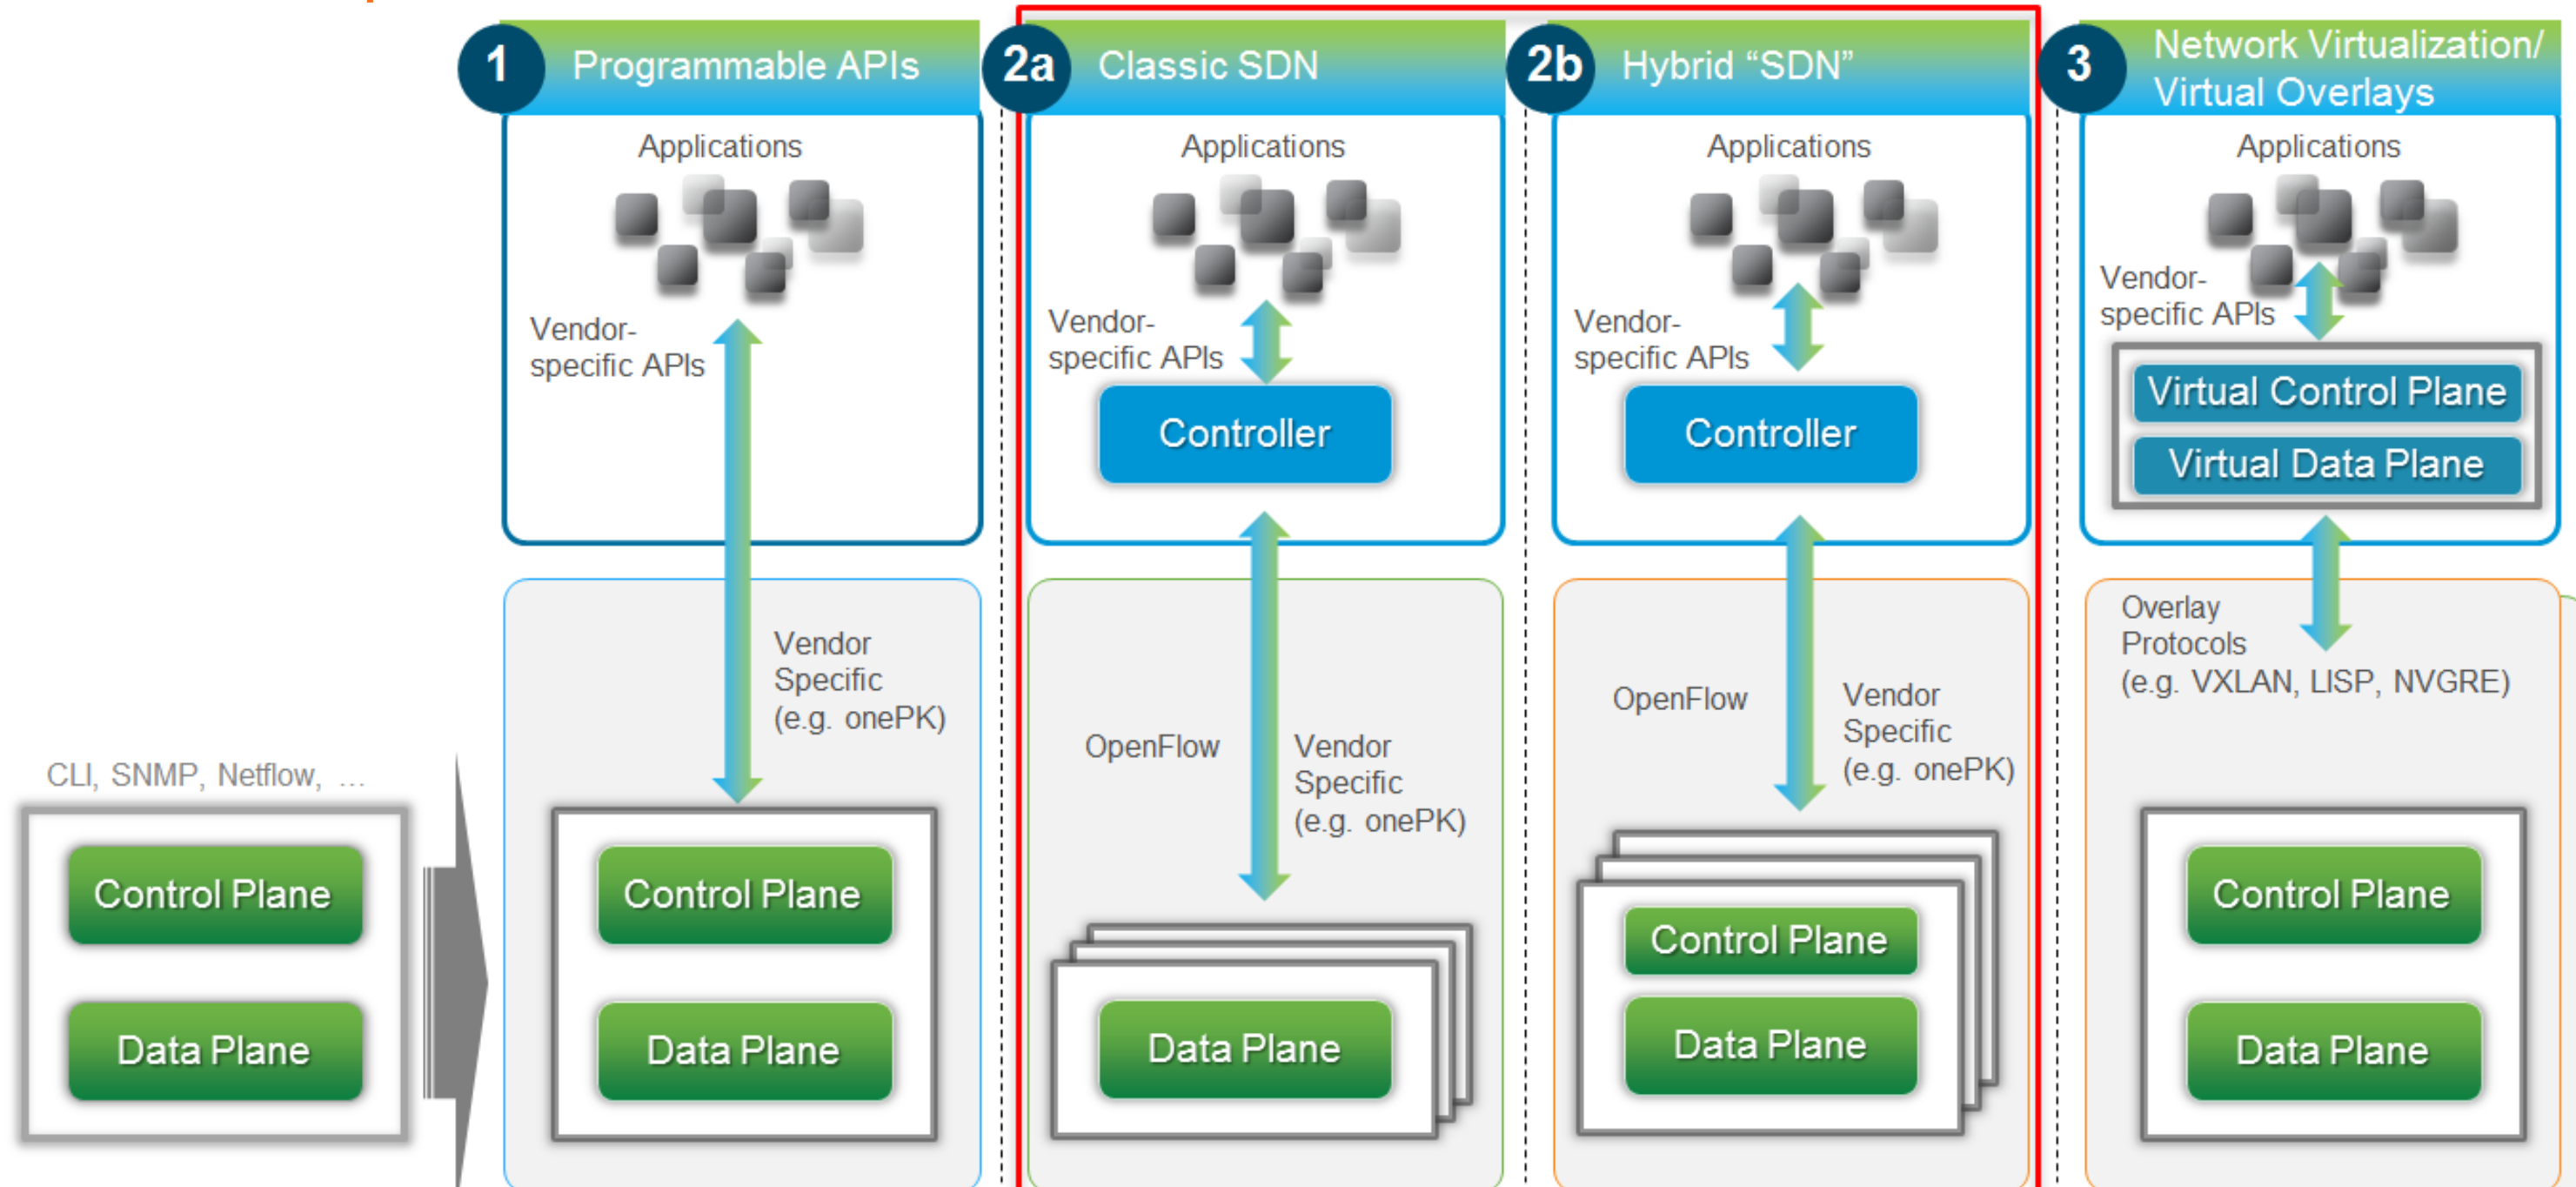
Nexus 1000v Enhancements

Controllers and Agents

SDN:  
SW Controller  
OpenFlow agents

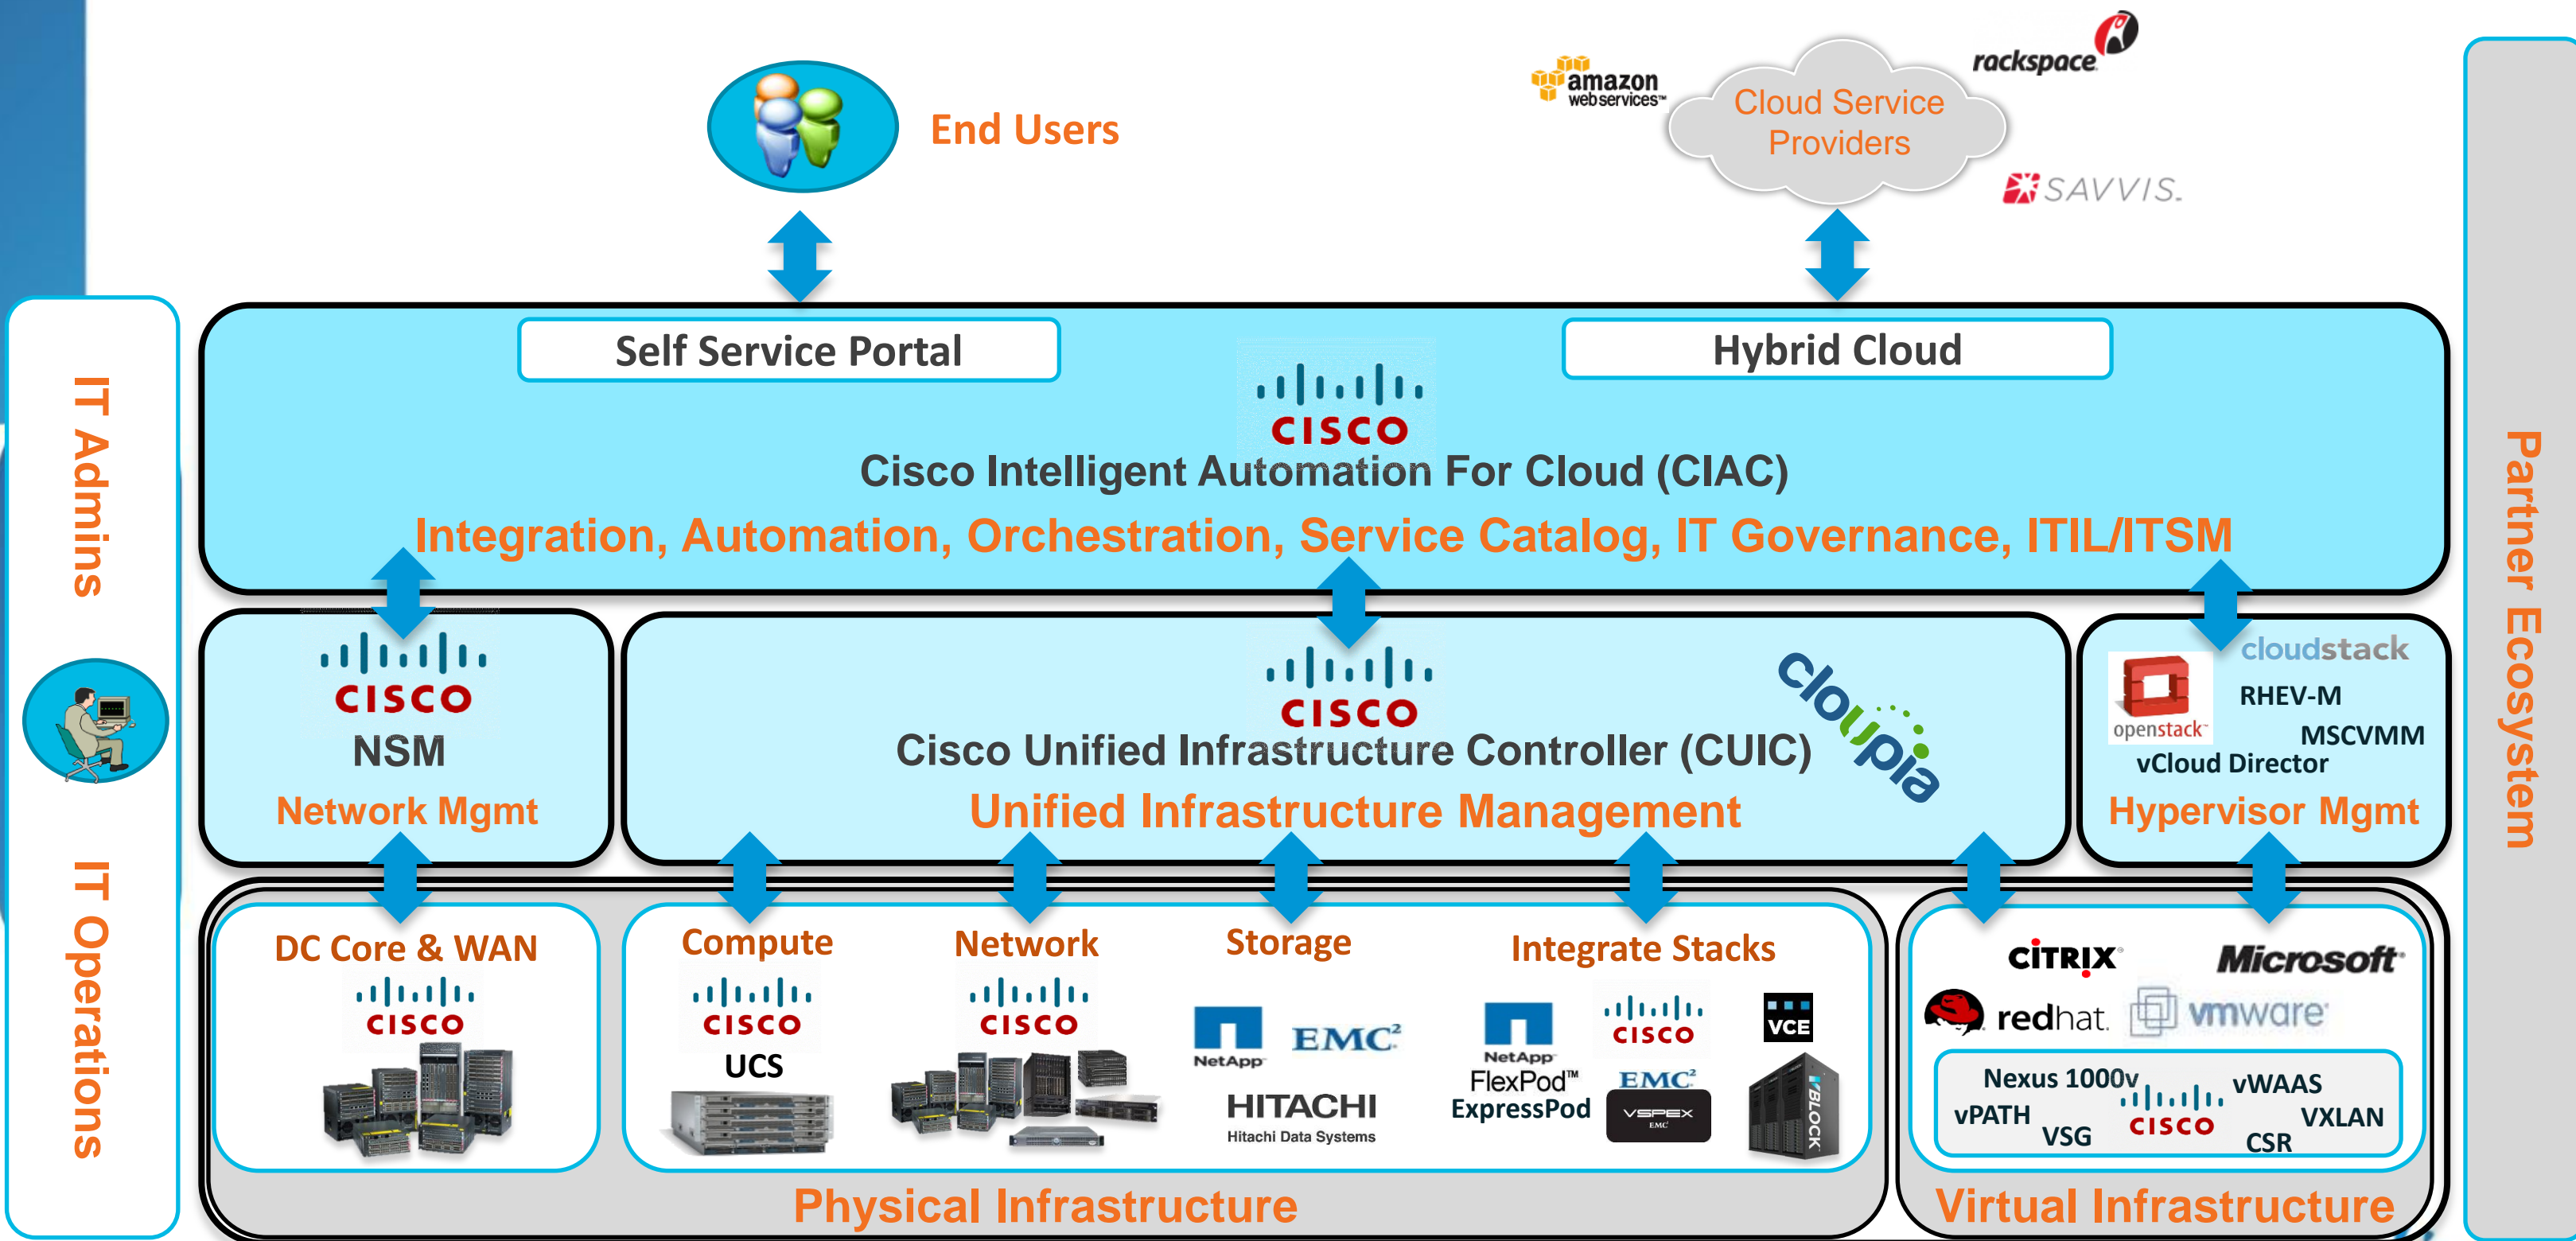
# Network Programmability Models

## Cisco's Deeper View



Openstack and Network Overlays Apply to All Models (Physical/Virtual)  
Custom Features Can Be Built

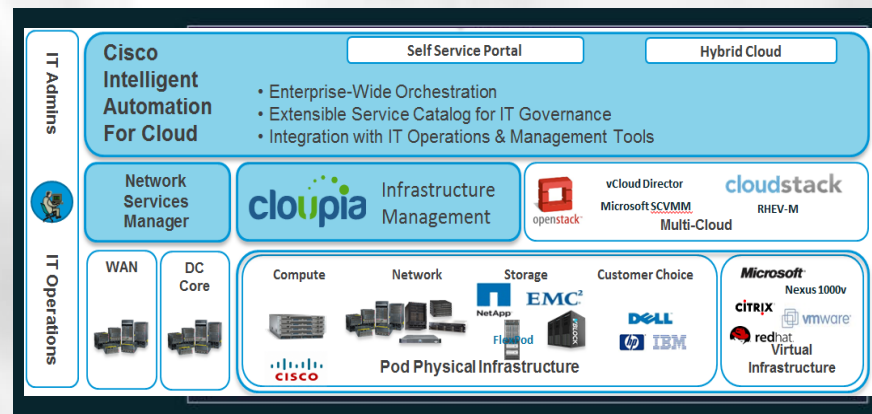
# Cisco Cloud Stack and Components Overview





# The Datacenter

Selecting the  
**BEST SERVICE PROVIDER  
& PUBLIC CLOUD**

## PRIVATE CLOUD

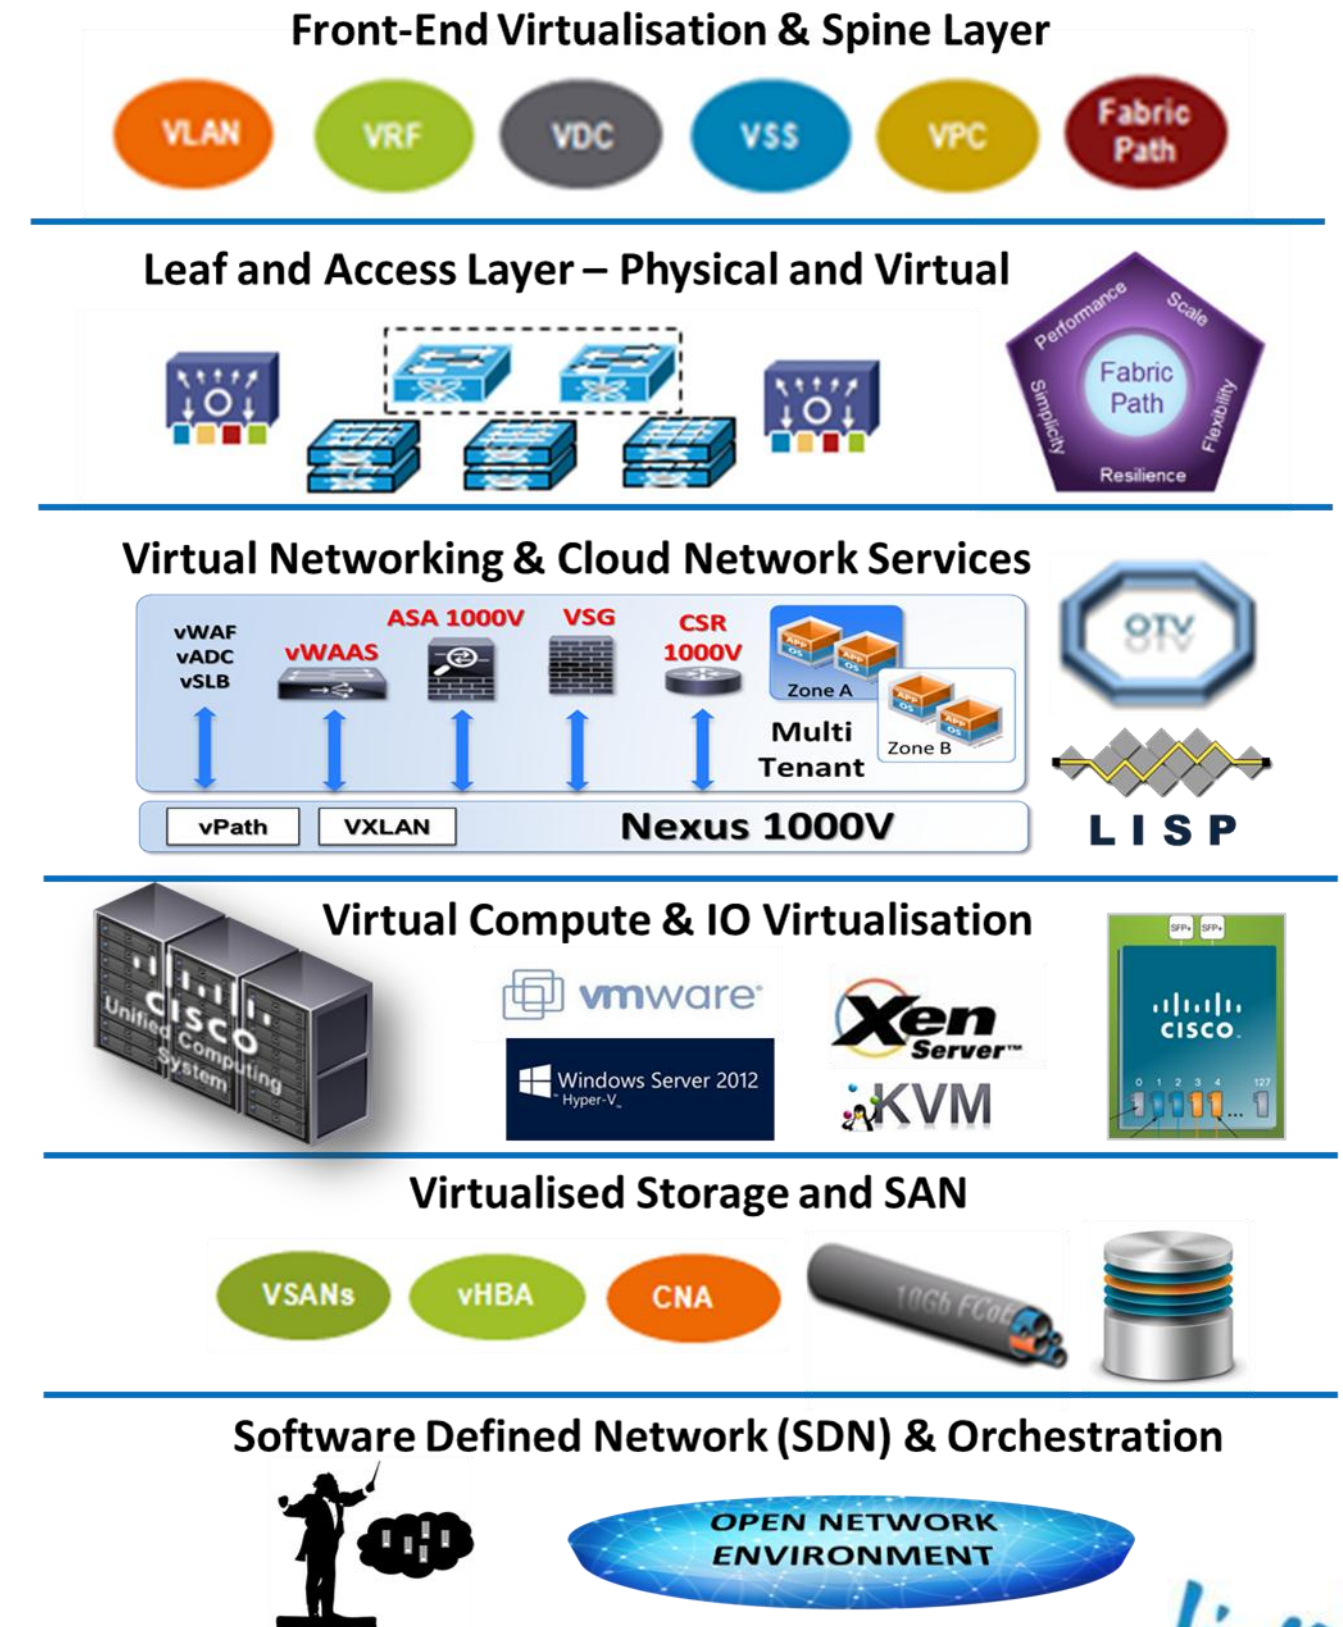


**L2 & L3 HYBRID CLOUD via  
Cisco Nexus 1000v  
InterCloud**



# Agenda

- **Data Centre Virtualisation Overview**
- **Front-End Data Centre Virtualisation**
  - DC Edge Layer
  - Aggregation / Spine Layer
  - Access / Leaf Layer
- **Virtual Networking & Cloud Network Services**
  - Virtual Access Layer: Nexus 1000v
  - **Cloud Network Services**
  - Inter DC Services: OTV & LISP
- **Virtual Compute & IO Virtualisation**
  - Unified Computing System
  - Hypervisors & Virtual IO @ Cisco UCS
- **Virtualised Storage & SAN**
  - Unified Fabric
  - Virtualised SAN and Storage
- **Software Defined Network & Orchestration**
  - SDN & Cisco Open Network Environment (ONE)
  - Cloud Orchestration
- **Implementation Examples**
- **Q&A**





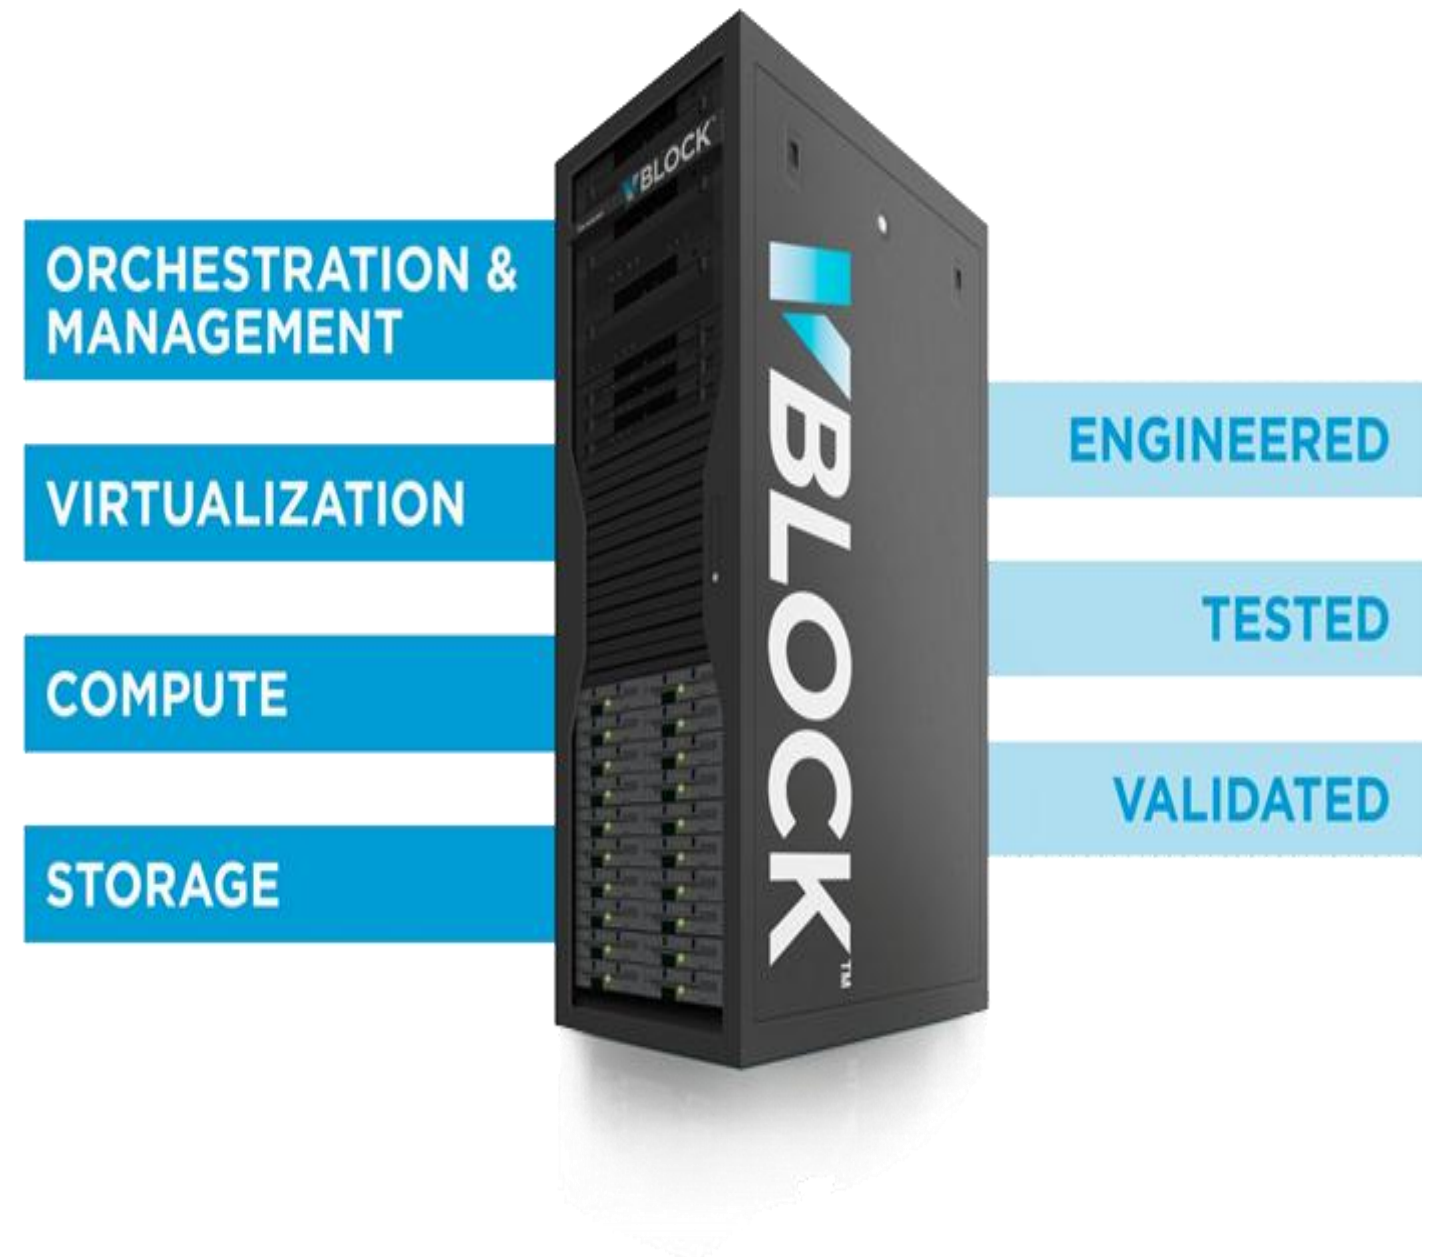
# VCE VBLOCK

## Vblock™ Infrastructure Platforms

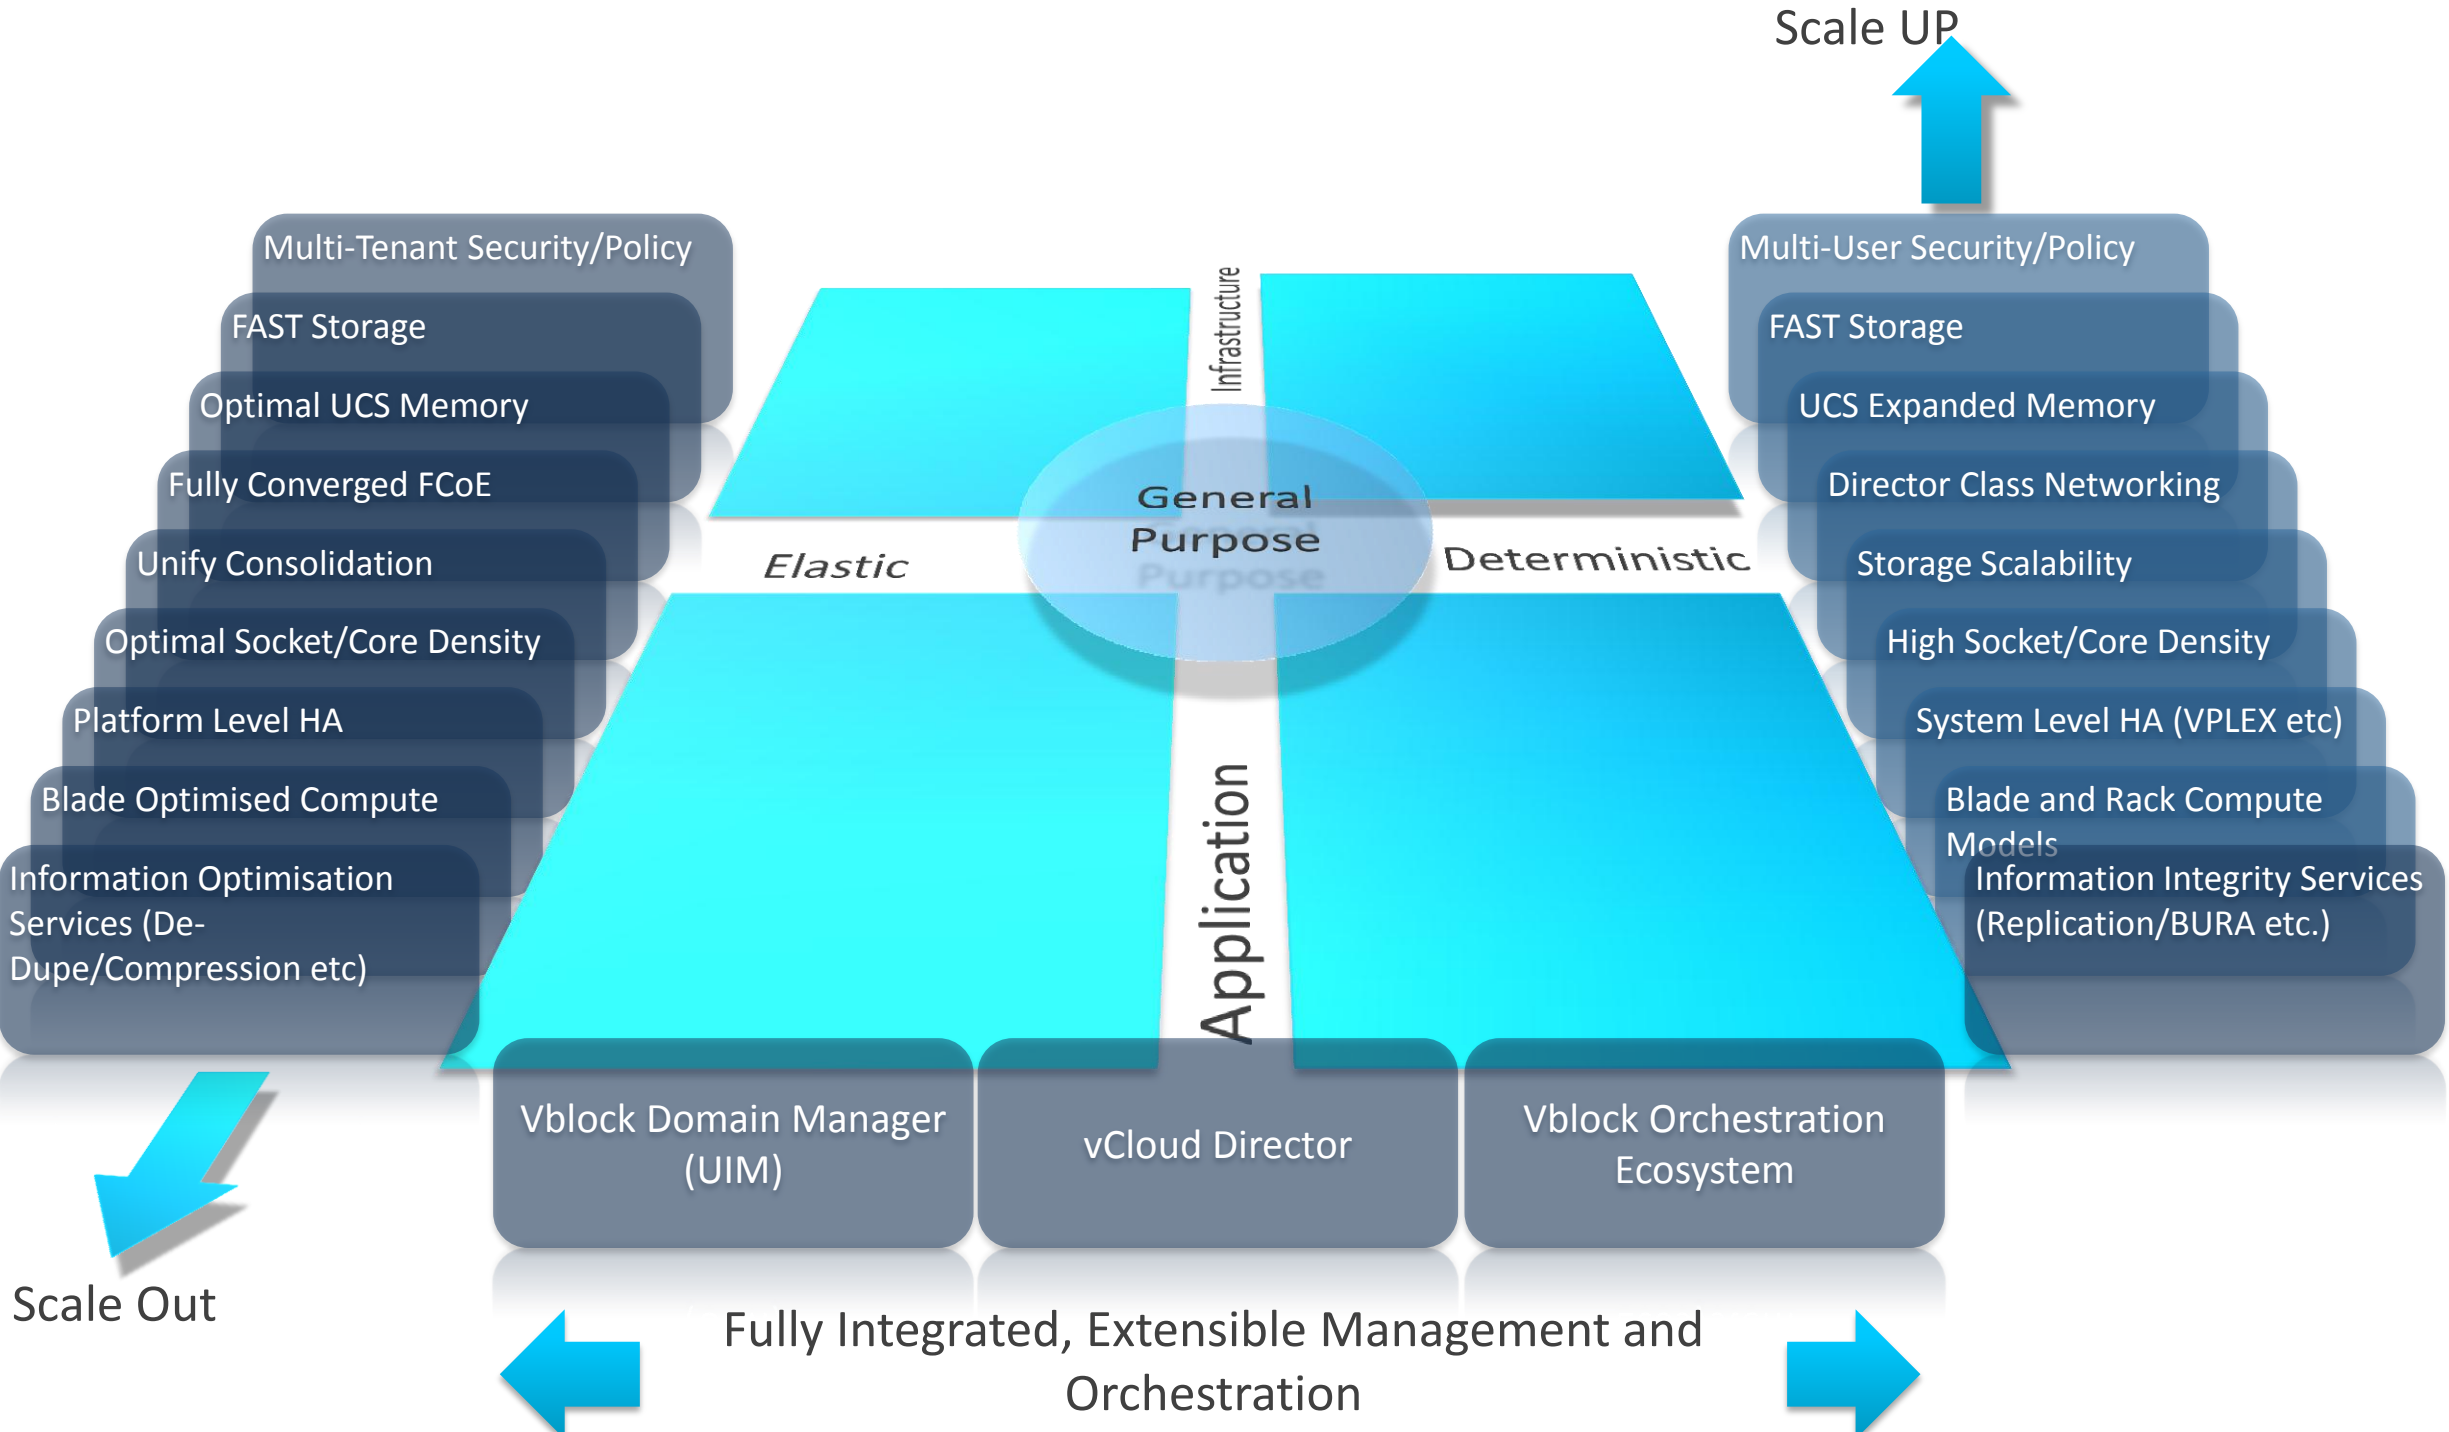
- Management and Orchestration: Unified Infrastructure Manager (UIM) framework
- Virtualisation: VMware
- Compute: Cisco UCS
- Network: Cisco Nexus and MDS switching
- Storage: EMC Symmetrix VMAX or Unified

## Solutions and Services

- Accelerate time to value of business applications
- Seamless Support



# Vblock Capabilities




# FlexPod

The Most Efficient Journey to the Next-Generation Data Centre




Cisco® UCS B-Series  
Blade Servers and  
UCS Manager



Cisco Nexus®  
Family Switches




NetApp® FAS  
10GE and FCoE



- Standard, prevalidated, best-in-class infrastructure building blocks
- Flexible: One platform scales to fit many environments and mixed workloads
  - Add applications and workload
  - Scale up and out
- Simplified management and repeatable deployments
- Design and sizing guides
- Services: Facilitate deployment of different environments

# Scale Out with FlexPod

- **Scale out with standard and proven configurations**
  - Predictable and highly efficient
    - Capacity and performance
    - Floor space, power, cooling
  - Or scale up within a FlexPod
- **Benefits**
  - Reduce effort for architecture, design, deployment, and test
  - Reduce infrastructure deployment cycle time by up to 50%
  - Manage resource pools not individual systems
    - Storage, compute, network



# Q & A



# Complete Your Online Session Evaluation

## Give us your feedback and receive a Cisco Live 2013 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site [www.ciscoliveaustralia.com/mobile](http://www.ciscoliveaustralia.com/mobile)
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm



Cisco *live!* 365

Don't forget to activate your Cisco Live 365 account for access to all session material,

communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.

[www.ciscoliveaustralia.com/portal/login.wv](http://www.ciscoliveaustralia.com/portal/login.wv)

Cisco *live!*

