



Virtualized Workload Mobility Latest Design Guidance

Jake Howering, Product Manager - Data Center Virtualization & Cloud
Systems Development Unit (SDU)

Syed Ghayur, Technical Marketing Engineer (TME)
Server Access Virtualization Technology Group (SAVTG)

Gunnar Anderson, Nexus 1010 Product Manager [HOST]
Server Access Virtualization Technology Group (SAVTG)

Nexus 1000V Fall Webinar Series

(www.cisco.com/go/1000vcommunity)

Date	Technical Track Topic	Webinar	Preso	Q&A
10/05	Nexus 1000V, VXLAN, and vCloud Director	Play	PDF	PDF
10/12	Virtualized Multi-Tenant Data Center (VMDC)	Play	PDF	
10/19	Nexus 1010 v1.3 - What's New?	Play	PDF	
10/26	Virtualized Workload Mobility - Latest Design Guidance	Register		
11/02	UCS and Nexus 1000V - Best Practices	Register		
11/09	Virtual Security Gateway (VSG) v1.2 - What's New?	Register		

Reference Solutions

Solution	Nexus 1000V	Nexus 1010	Virtual Security Gateway
vBlock	✓		
FlexPOD	✓	✓	
Virtual Desktop	✓	Implicit Support	✓
Virtualized Multi-tenant DC (VMDC)	✓	Implicit support	Target Nov '11
Long-distance vMotion (Virtualized Workload Mobility)	✓	Implicit support	✓
PCI 2.0	✓	Implicit support	✓

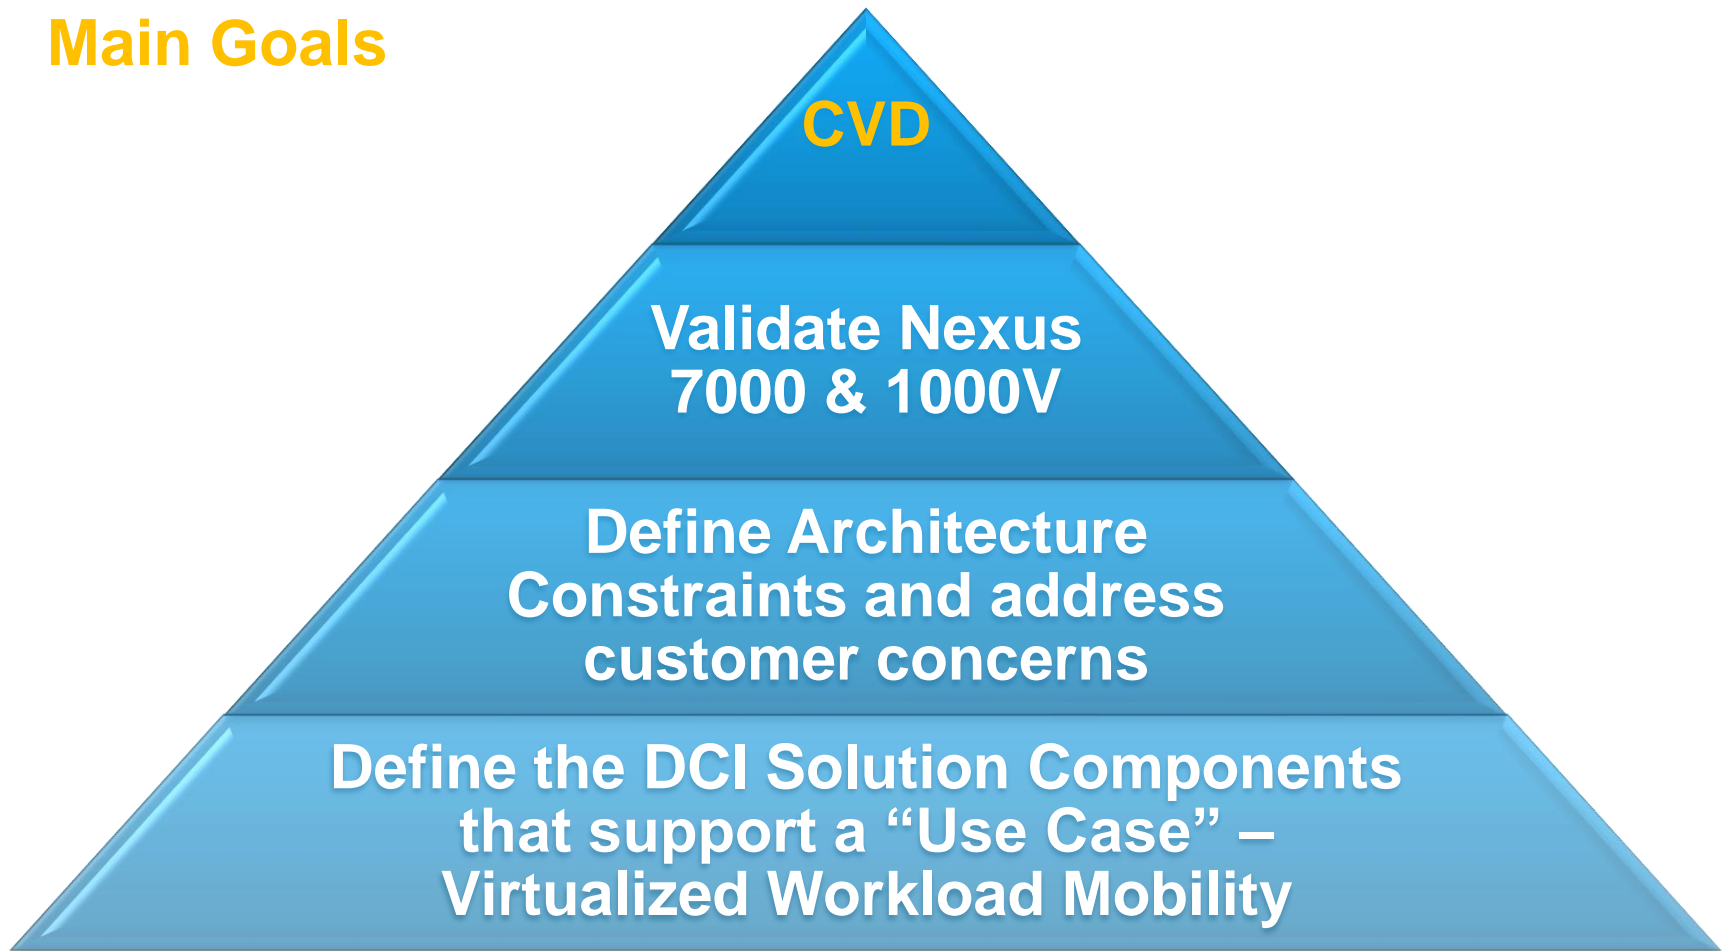
Agenda

- Business Drivers
 - Systems Development Unit
 - DCI Use Case: Virtualized Workload Mobility
- Virtualized Workload Mobility Components
 - LAN Extensions
 - Path Optimization
 - Storage Elasticity
 - Workload Mobility – Nexus 1000V/1010 and VSG
- Nexus 1000V/1010 & VSG and Virtualized Workload Mobility
- VXLAN
- Resources
- Q&A as we go...



SDU Developing a Solution: Virtualized Workload Mobility

Main Goals



CVD = Cisco Validated Design

Agenda


- Business Drivers
 - Systems Development Unit (SDU)
 - DCI Use Case: Virtualized Workload Mobility
- Virtualized Workload Mobility Components
 - LAN Extensions
 - Path Optimization
 - Storage Elasticity
 - Workload Mobility – Nexus 1000v and VSG



Private Cloud – Building Blocks

Virtualized Multi-Tenant Data Center

Integrated Compute Stacks



Integrated Compute Stacks

- Reference Architectures
- Unified Compute Systems

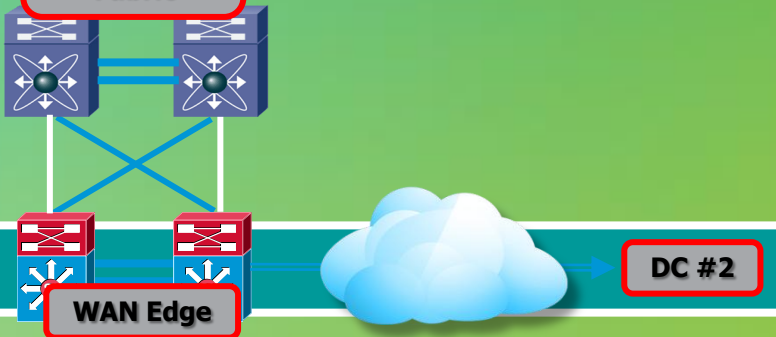
Unified Data Center Networking



Unified Data Center Networking

- Converged LAN & SAN
- Network Services

Data Center Interconnect



Data Center Interconnect

- Application Mobility
- Clustered Solutions


Workload Mobility Solution Requirements

LAN Extension

Enables Application Mobility with OTV, VPC, VPLS

Path Optimization

Efficient traffic flow from client-to-server with LISP, GSS



Virtual Networking & Services

Stateful Service stitching between physical and virtual networks With Nexus 1000V & VSG

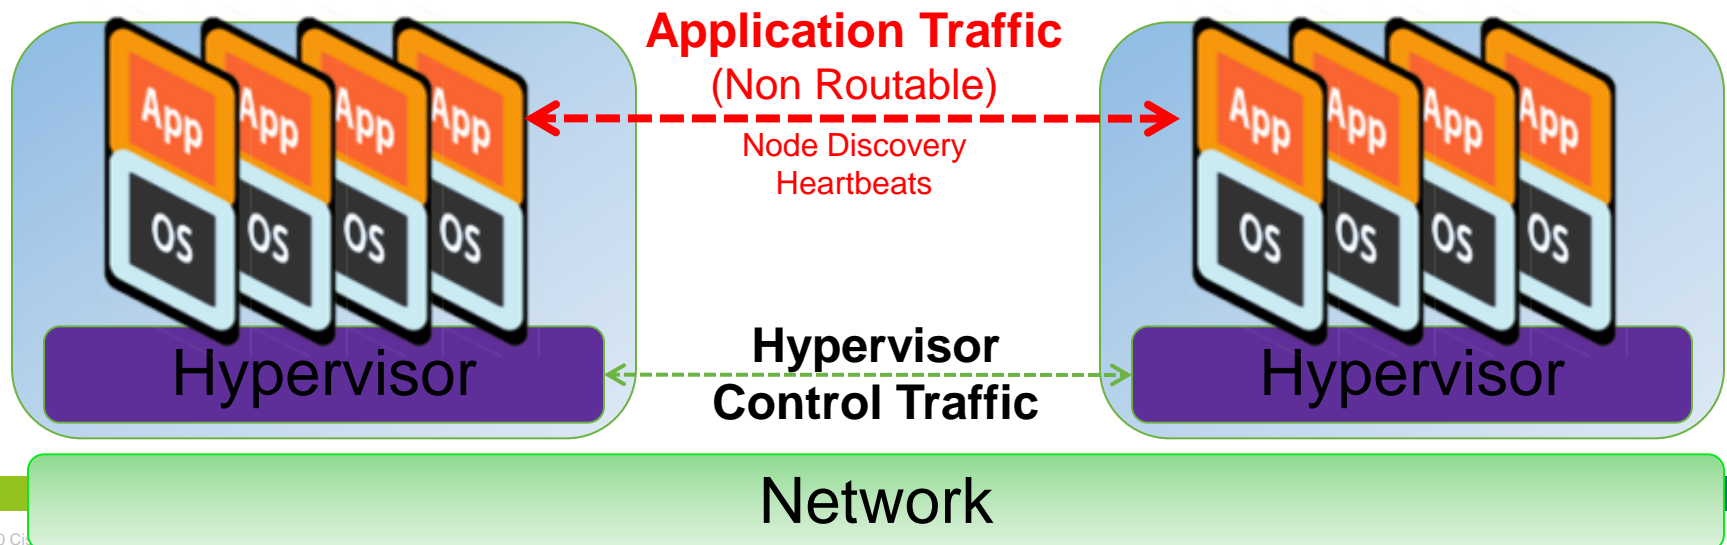
Storage Elasticity

Continuous Data Availability throughout the virtual data center With NetApp or EMC

Key Point: LAN Extensions – Foundation Component

Why the Need for LAN Extensions ?

- Applications running in the VMs use non-routable traffic
 - e.g. Node Discovery & Heartbeats in clustered Applications
- With Virtualization, application members may be distributed across PODs/Data-centers
- Moving and distributing application members across locations should not break the application



Workload Mobility: Products and Technologies

Virtualization

- Distributed Virtual Switching - Nexus 1000V
- Distributed Virtual Security - Cisco VSG
- Hypervisor - VMware ESXi 4.1 or 5.0

LAN Extension


- Overlay Transport Virtualization (OTV)
- Virtual Port Channels or
- VPLS

Storage Extension

- Synchronous Replication distances
- Shared Storage Model
- Active / Cache - Netapp FlexCache
- Active / Active - EMC VPLEX


Path Optimization

- Egress the Virtual Data Center – FHRP Filtering
- Ingress the Virtual Data Center – LISP or ACE/GSS integration with vCenter



Workload Mobility

Component View



Cisco-VMware With
EMC & NetApp
Validated Design for
Virtualized Workload
Mobility

DC 1

DC 2



LAN Extension

ESX-A source

ESX-B target



Nexus 1000V with
distributed switching



Storage Elasticity

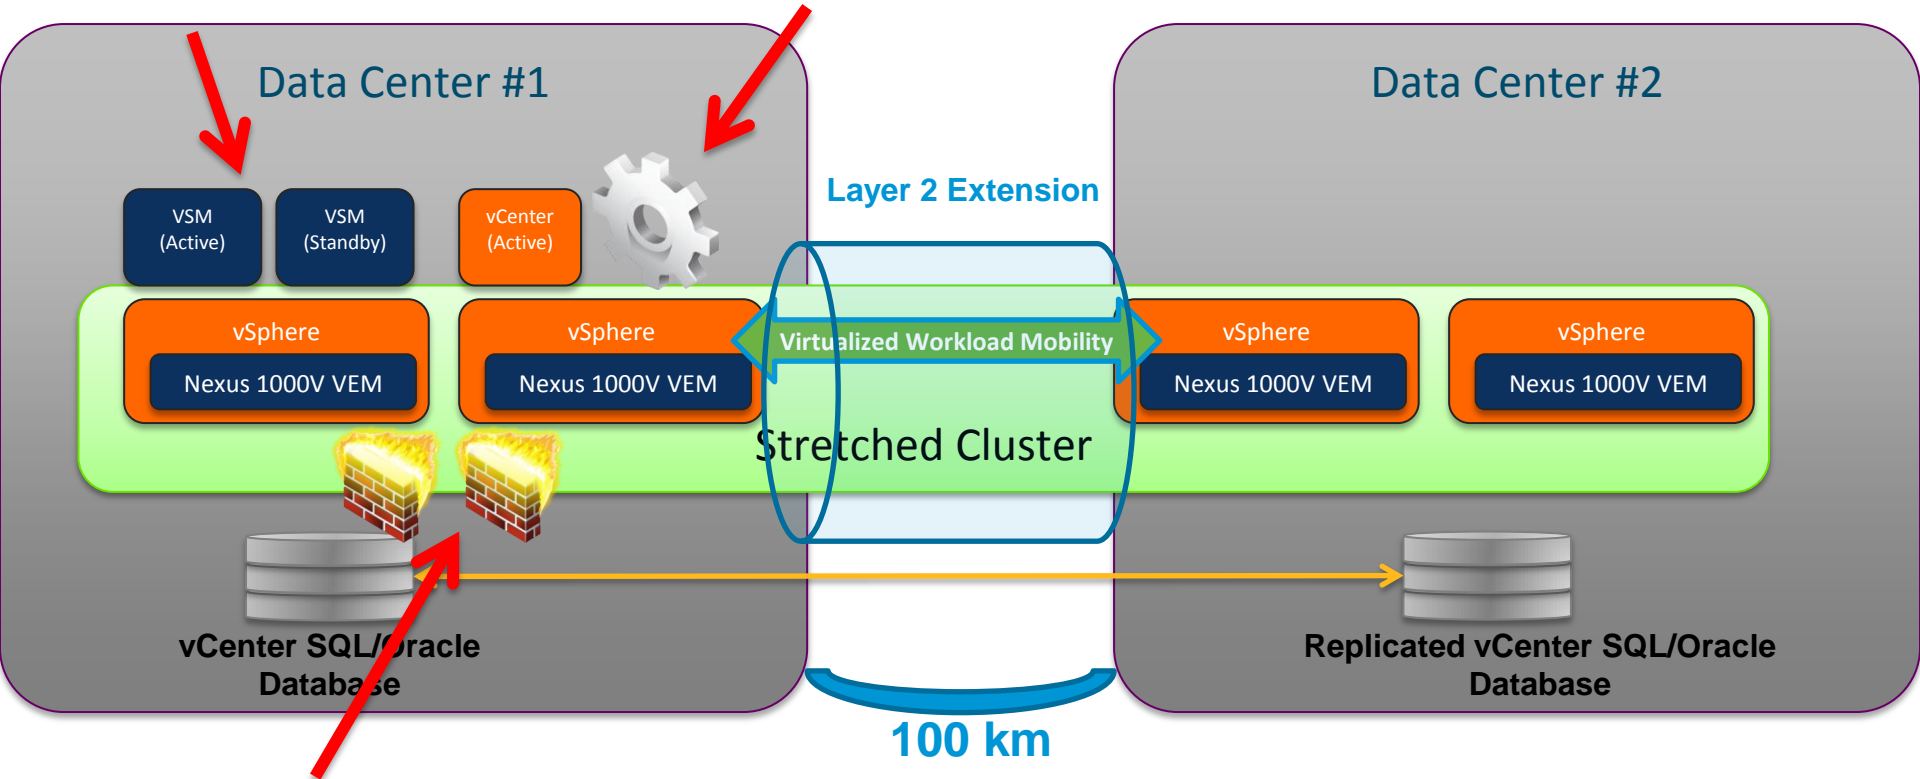


Key Point: Nexus 1000V Deployment Model

Stretching the ESX Cluster to 100km apart

Nexus 1000V VSM Pair

VNMC



VSG Pair

Agenda

- Business Drivers
 - Systems Development Unit (SDU)
 - DCI Use Case: Virtualized Workload Mobility
- Virtualized Workload Mobility Components
 - LAN Extensions
 - Path Optimization
 - Storage Elasticity
 - Workload Mobility – Nexus 1000V and VSG

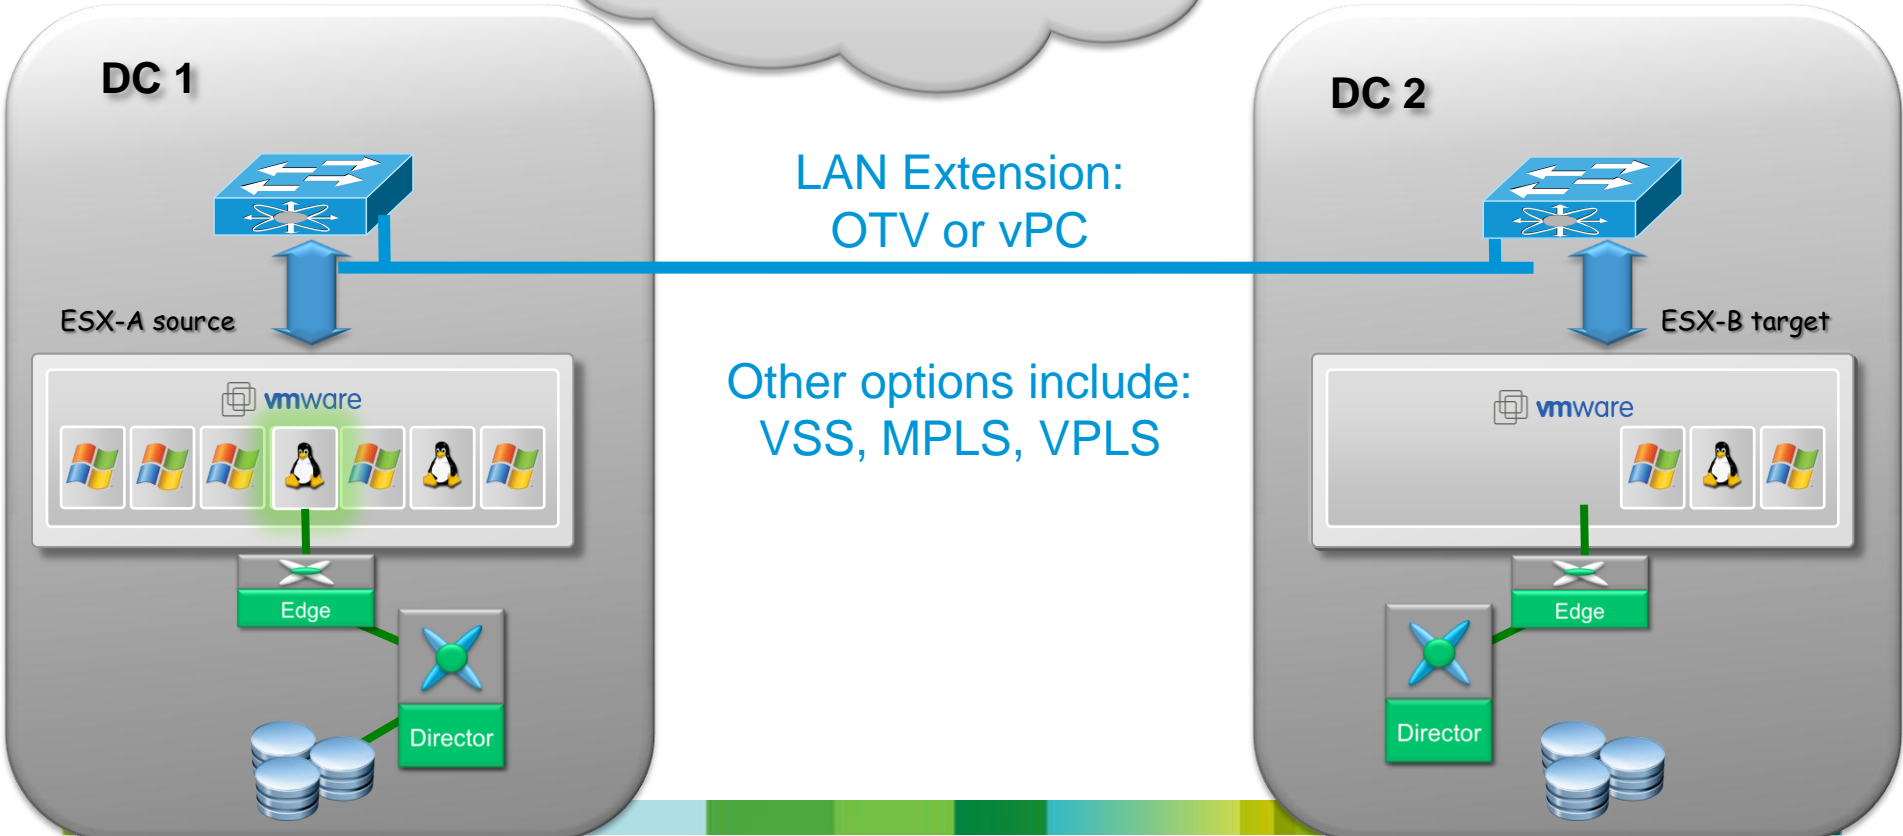



Data Center Interconnect

LAN Extension




- STP Isolation is the key element
- End to End Loop Avoidance



LAN Extension

vPC



STP Topology


vPC Topology

STP - Half of the Ports are blocked

vPC – Loop Free topology and use all B/W

LAN Extension

vPC Extended over Dark Fiber

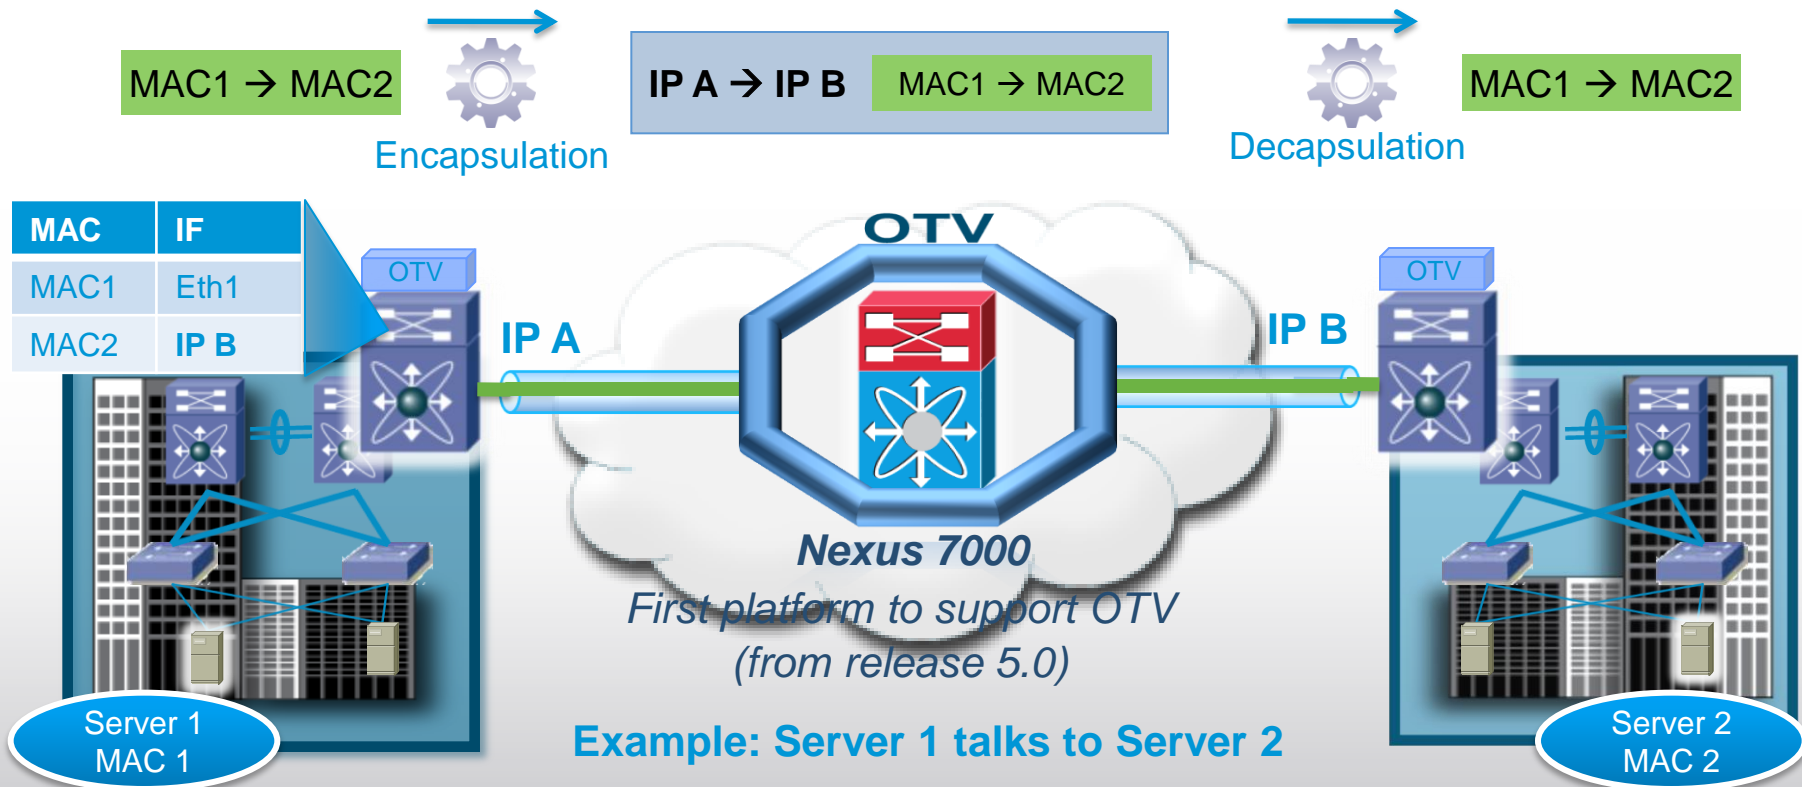


Extend Loop Free Topology over distance
Separate L3 link for dynamic IGP

Overlay Transport Virtualization

Dynamic Encapsulation


- 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



http://www.cisco.com/en/US/prod/switches/ps9441/nexus7000_preso.html

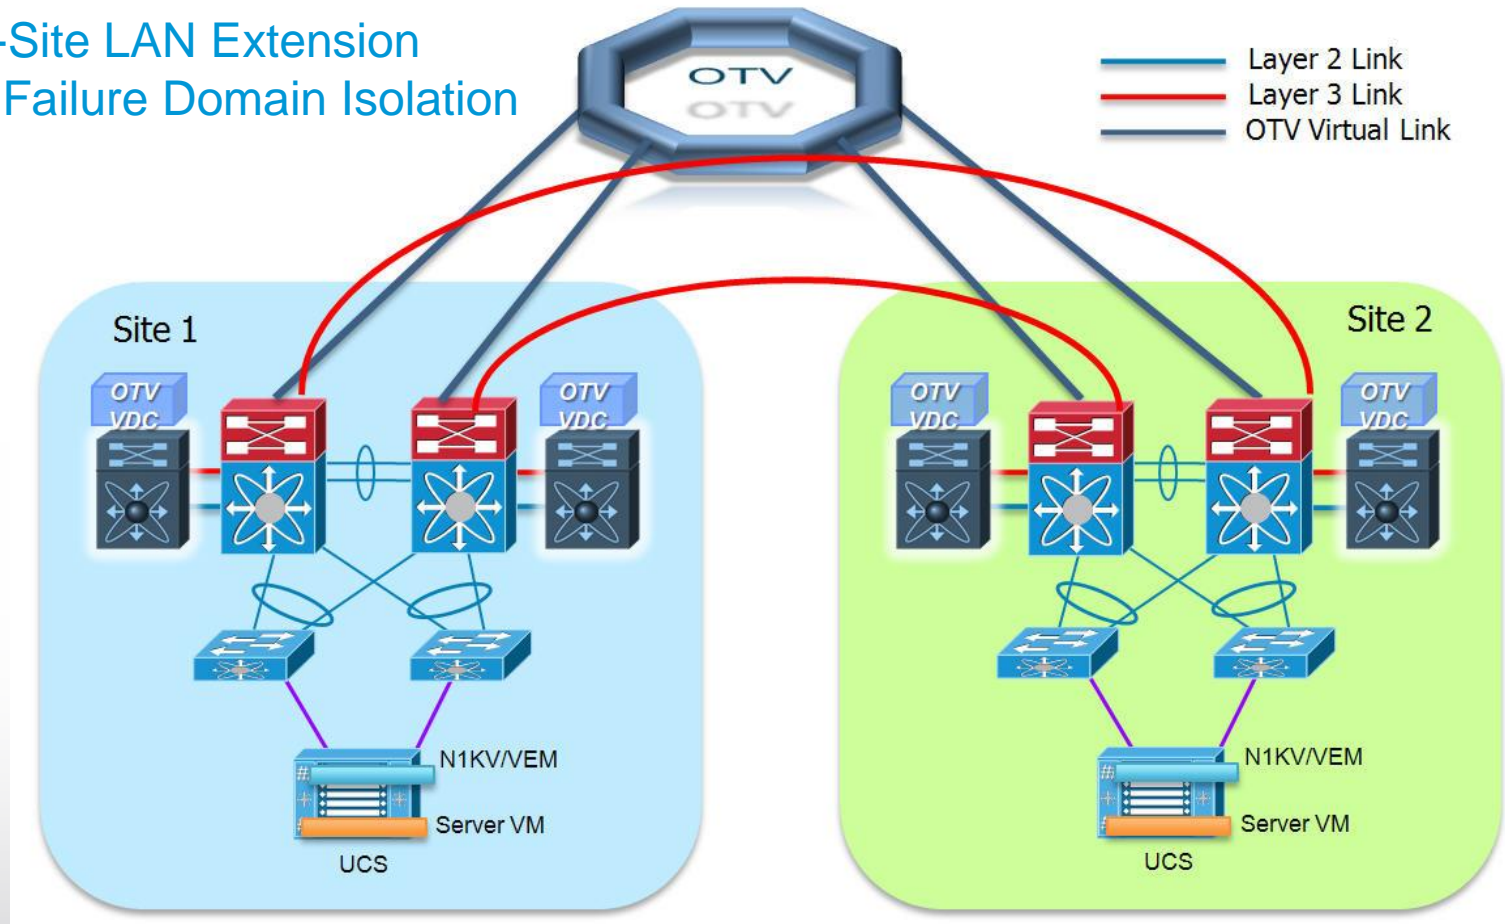
Real Problems Solved with OTV

- Extensions over any transport
- No Re-design required
- Failure boundary preservation
- Site independence / isolation
- Resiliency/multihoming
- Built-in end-to-end loop prevention
- Operations simplicity (5 commands)



Overlay Transport Virtualization

Site-to-Site LAN Extension
Native Failure Domain Isolation




Agenda

- Business Drivers
 - Systems Development Unit (SDU)
 - DCI Use Case: Virtualized Workload Mobility
- Virtualized Workload Mobility Components
 - LAN Extensions
 - Path Optimization
 - Storage Elasticity
 - Workload Mobility – Nexus 1000V and VSG



Data Center Interconnect

Path Optimization



Options

- Egress
 1. Addressed by FHRP Filtering
- Ingress:
 1. LISP
 2. DNS Redirection with GSS

DC 1



ESX-A source



Edge



Director



DC 2



ESX-B target



Edge

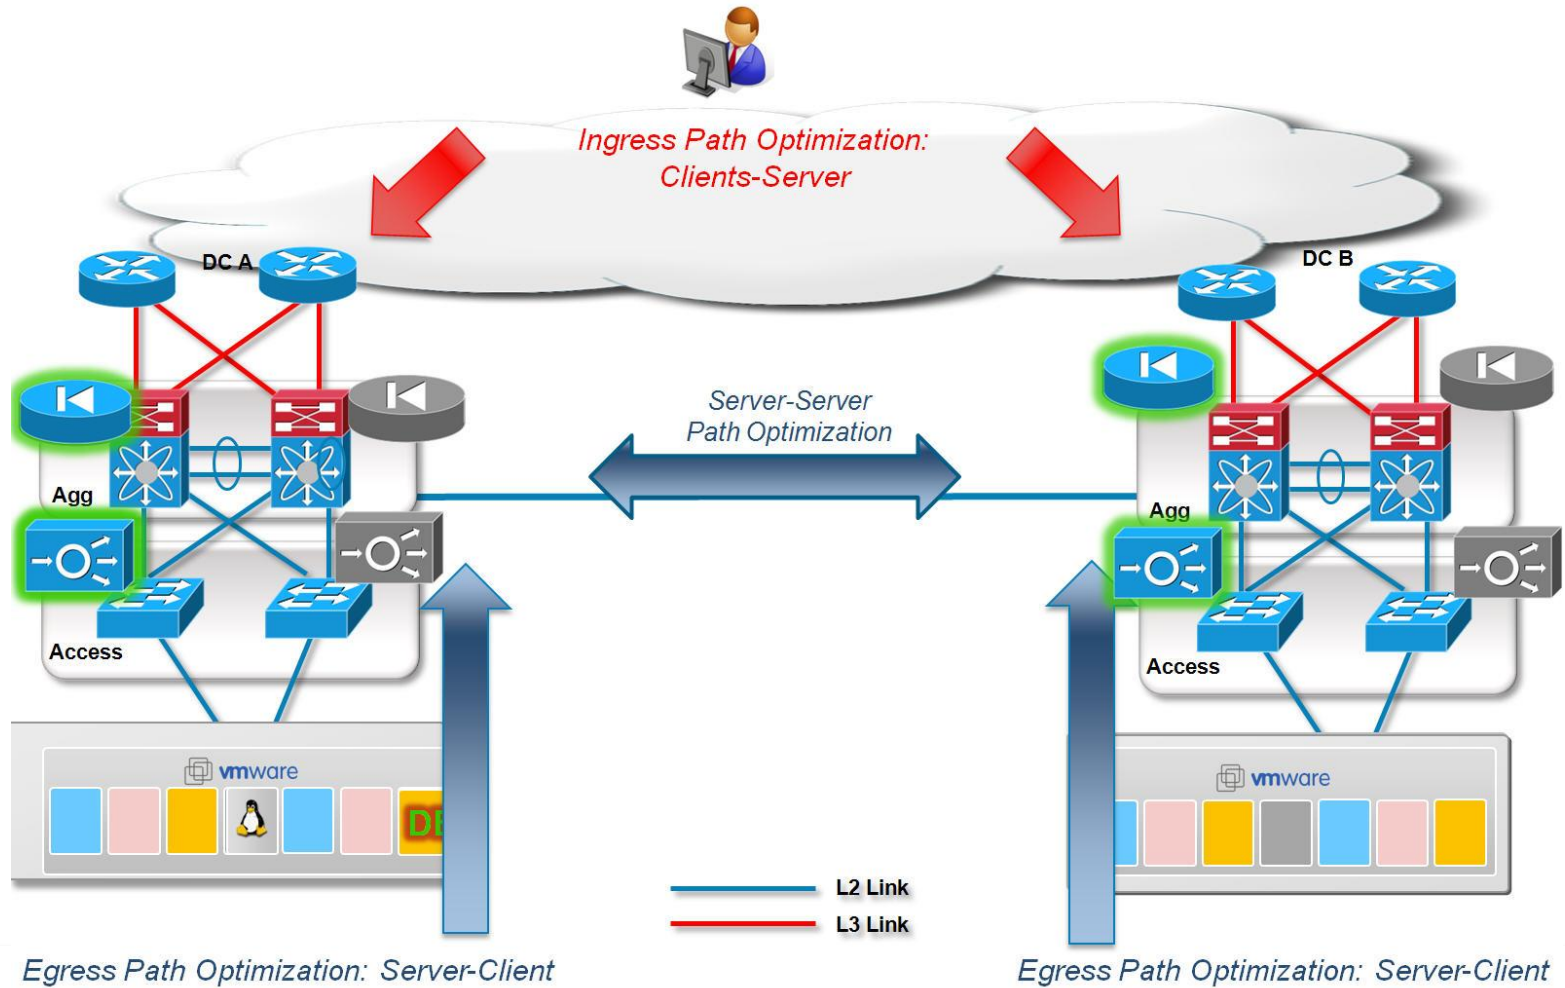


Director



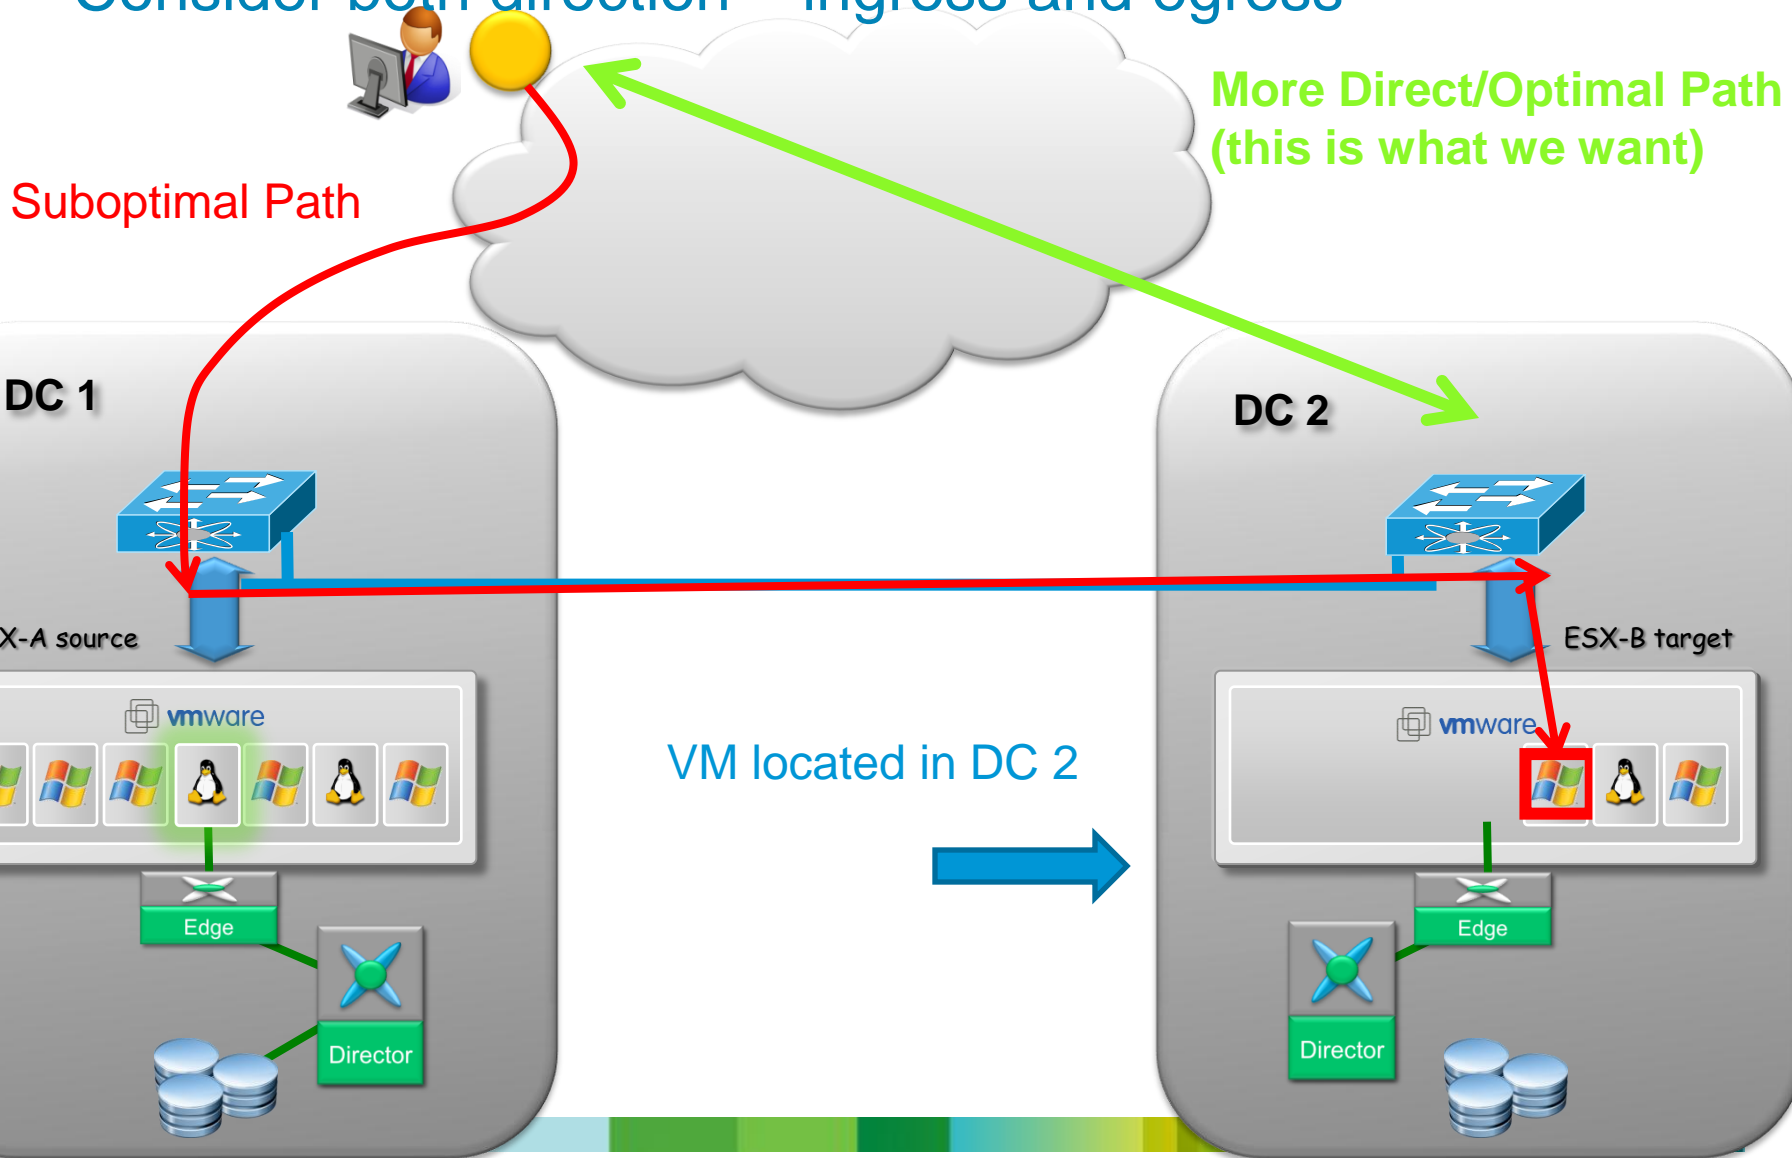
Path Optimization

3 Different Paths to Optimize



Path Optimization

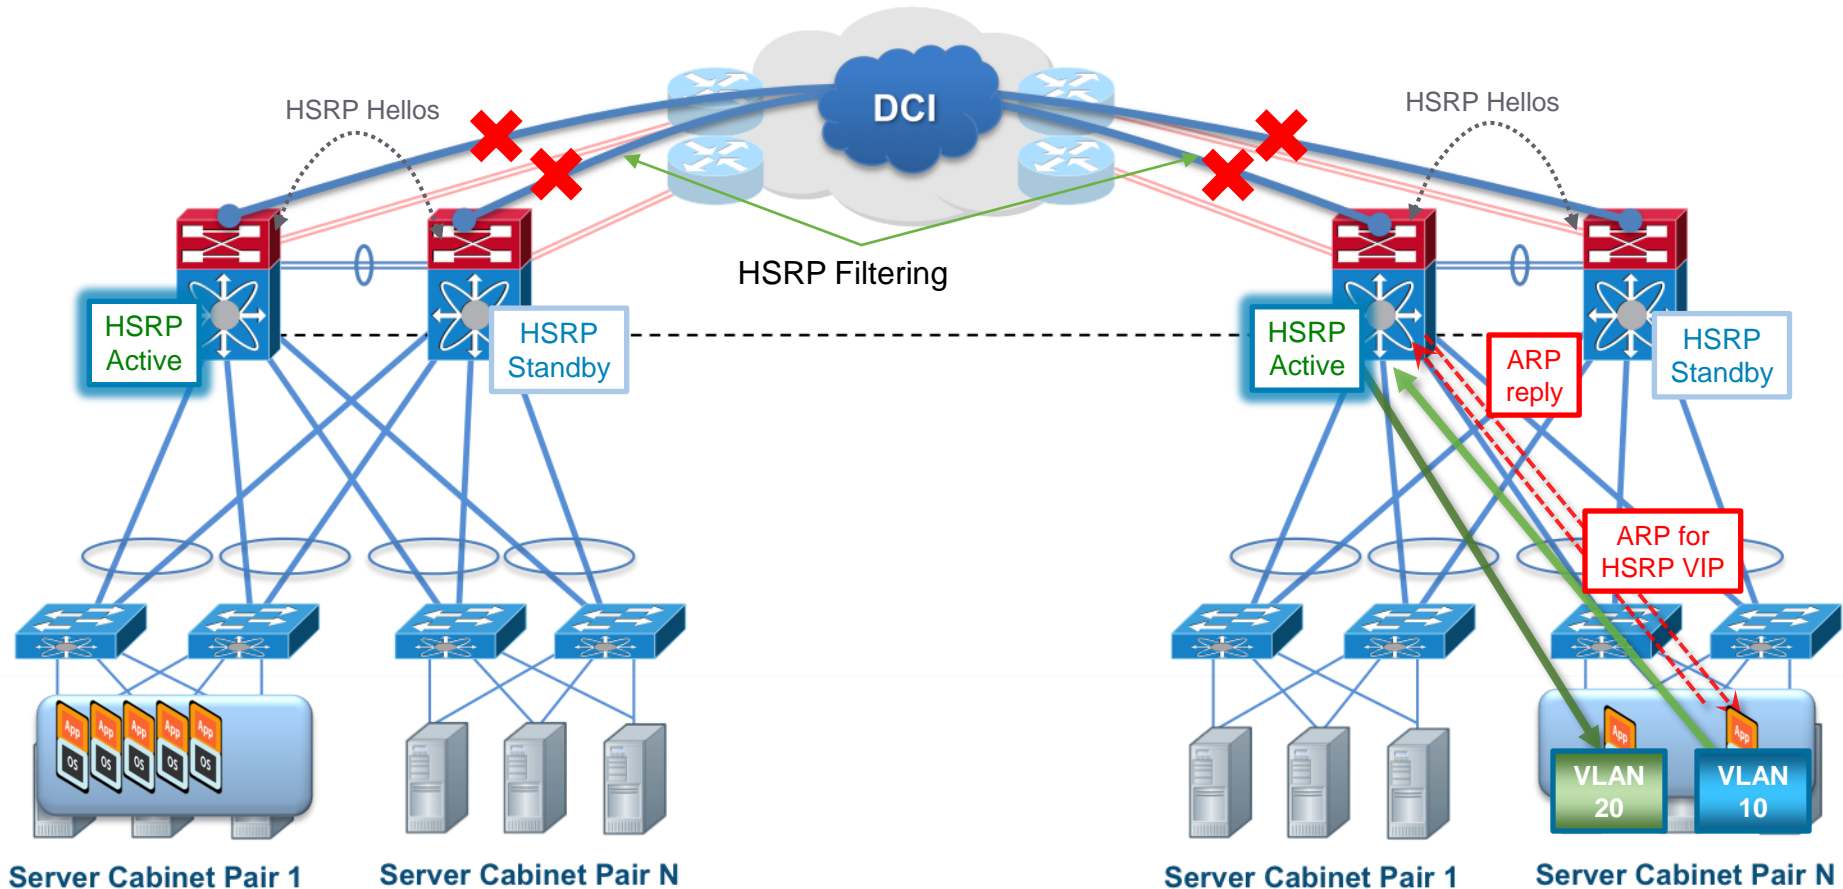
Consider both direction – ingress and egress



Egress Path Optimization

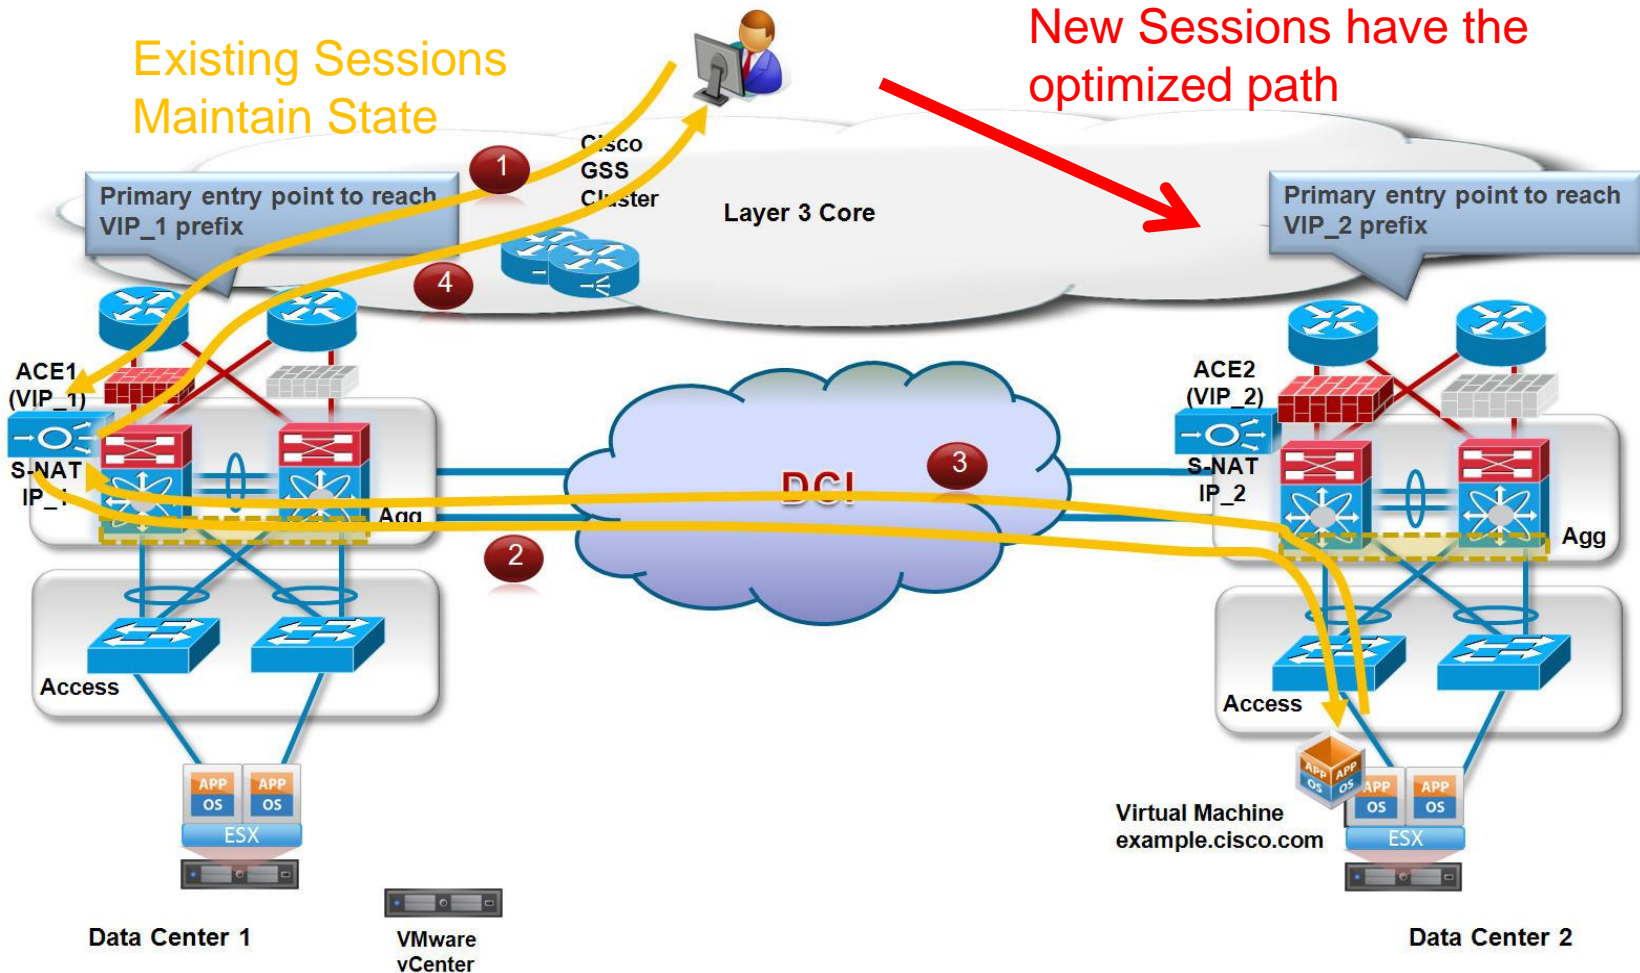
FHRP Filtering Solution

- Filter FHRP with combination of VACL and MAC route filter
- **Result: Still have one HSRP group with one VIP**, but now have active router at each site for optimal first-hop routing



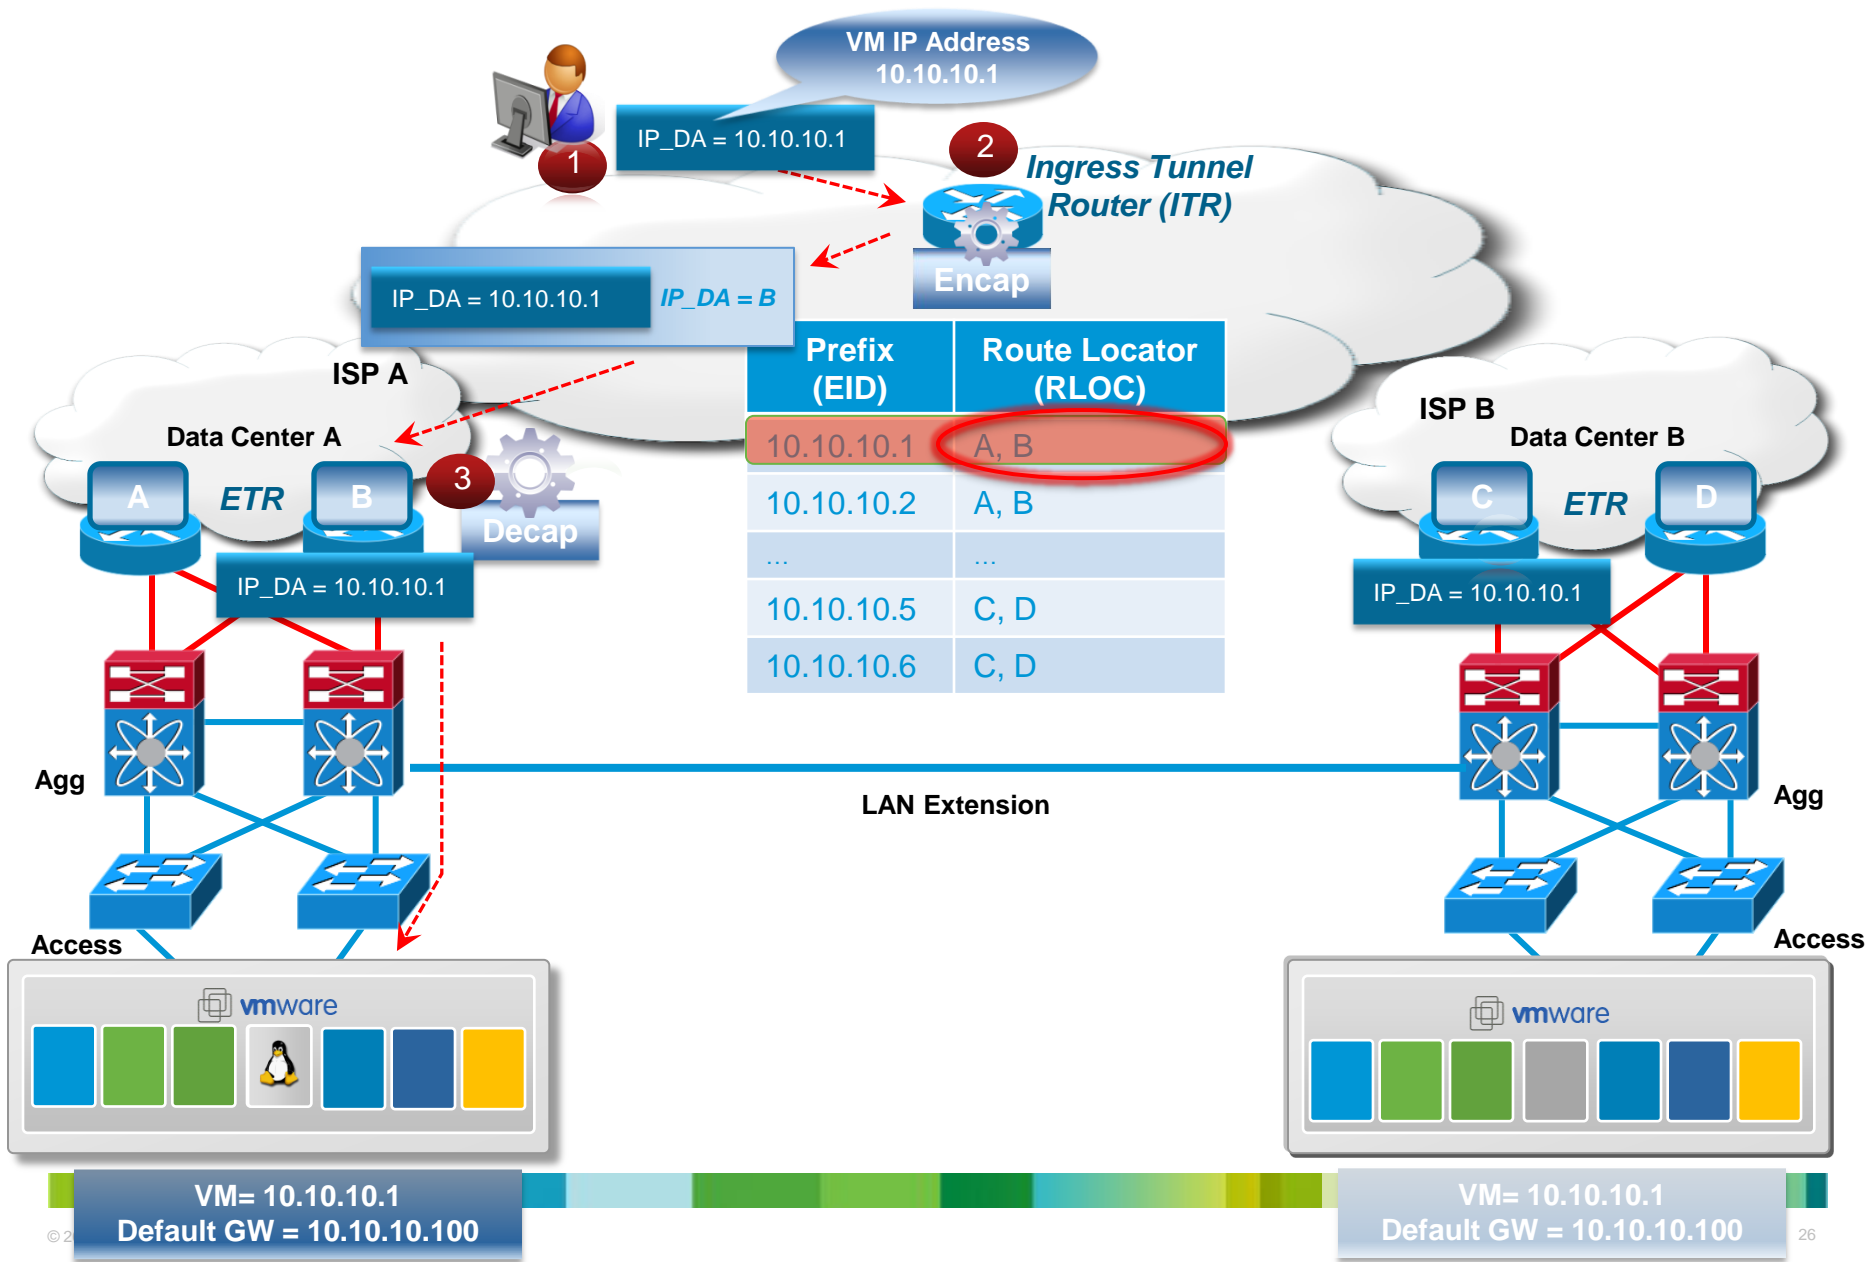
Ingress Path Optimization

DNS Based



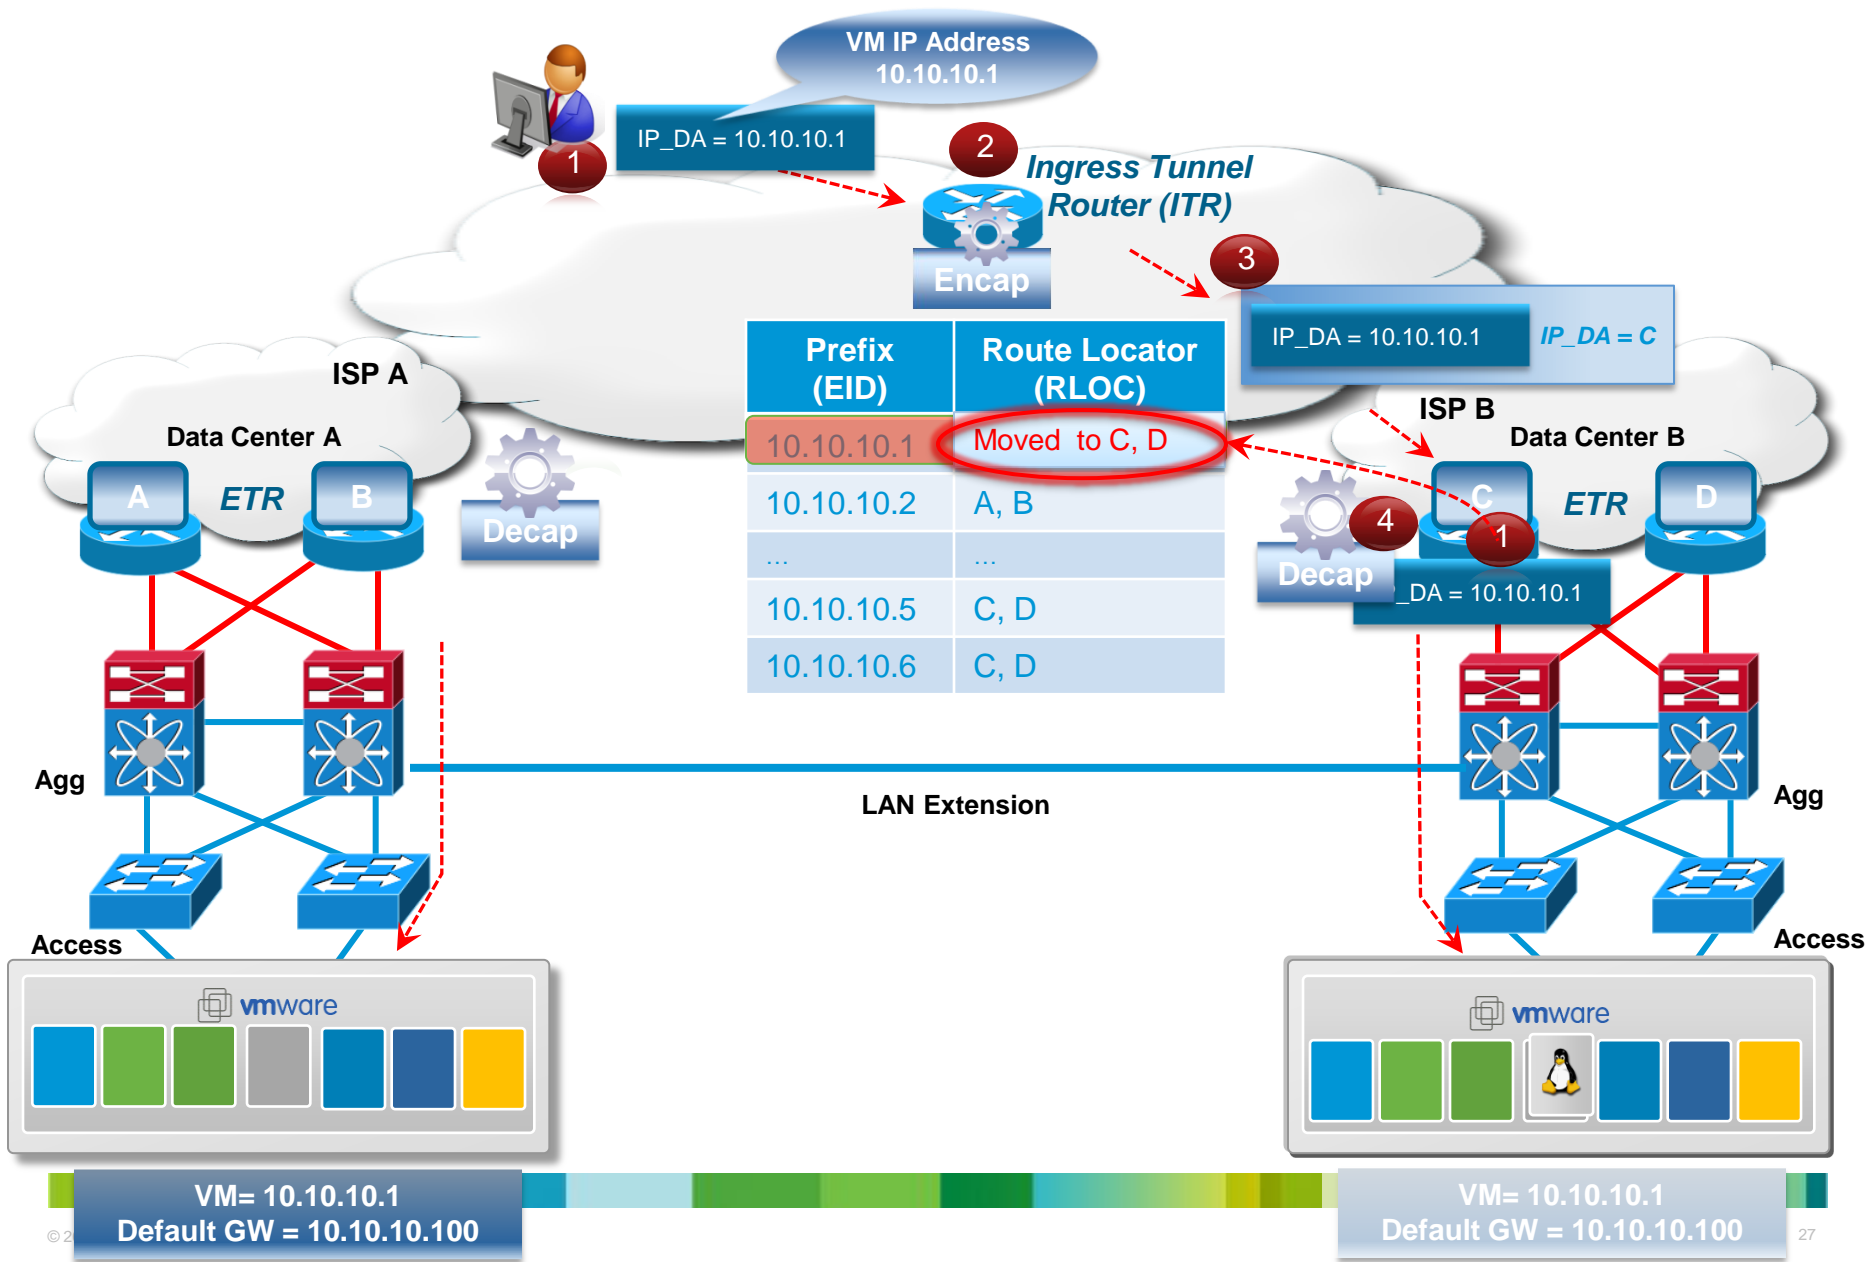
Ingress Path Optimization

LISP – VM Mobility



Ingress Path Optimization

LISP – VM Mobility



Agenda

- Business Drivers
 - Systems Development Unit (SDU)
 - DCI Use Case: Virtualized Workload Mobility
- Virtualized Workload Mobility Components
 - LAN Extensions
 - Path Optimization
 - Storage Elasticity
 - Workload Mobility – Nexus 1000V and VSG



Data Center Interconnect

SAN Extension



- Synchronous I/O implies strict distance limitation
- Localization of Active Storage is key

DC 1



ESX-A source



Edge



Director



DC 2



ESX-B target



Edge



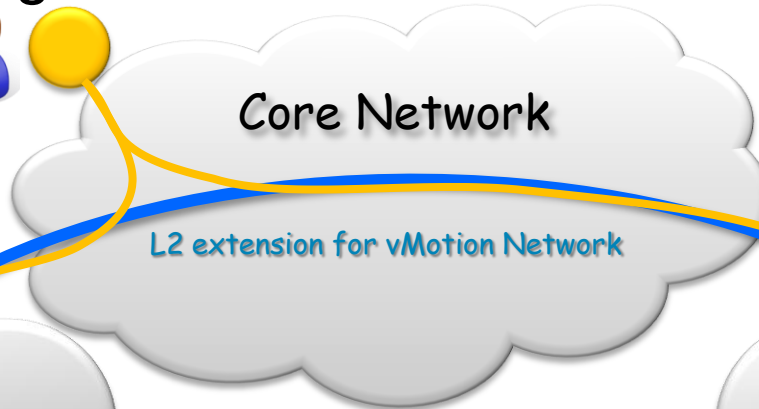
Director



Synchronous Data Replication
Asynchronous Data replication

Storage Elasticity

Shared Storage



L2 extension for vMotion Network

DC 1



ESX-A source



vmware



Edge



Director



Virtual Center



VM and Storage in Different locations

DC 2



ESX-B target




vmware



Storage Elasticity

NetApp FlexCache (Active/Cache)



Core Network

L2 extension for vMotion Network

DC 1



ESX-A source



Origin Volume


vmware



DC 2




ESX-B target



Cache Volume

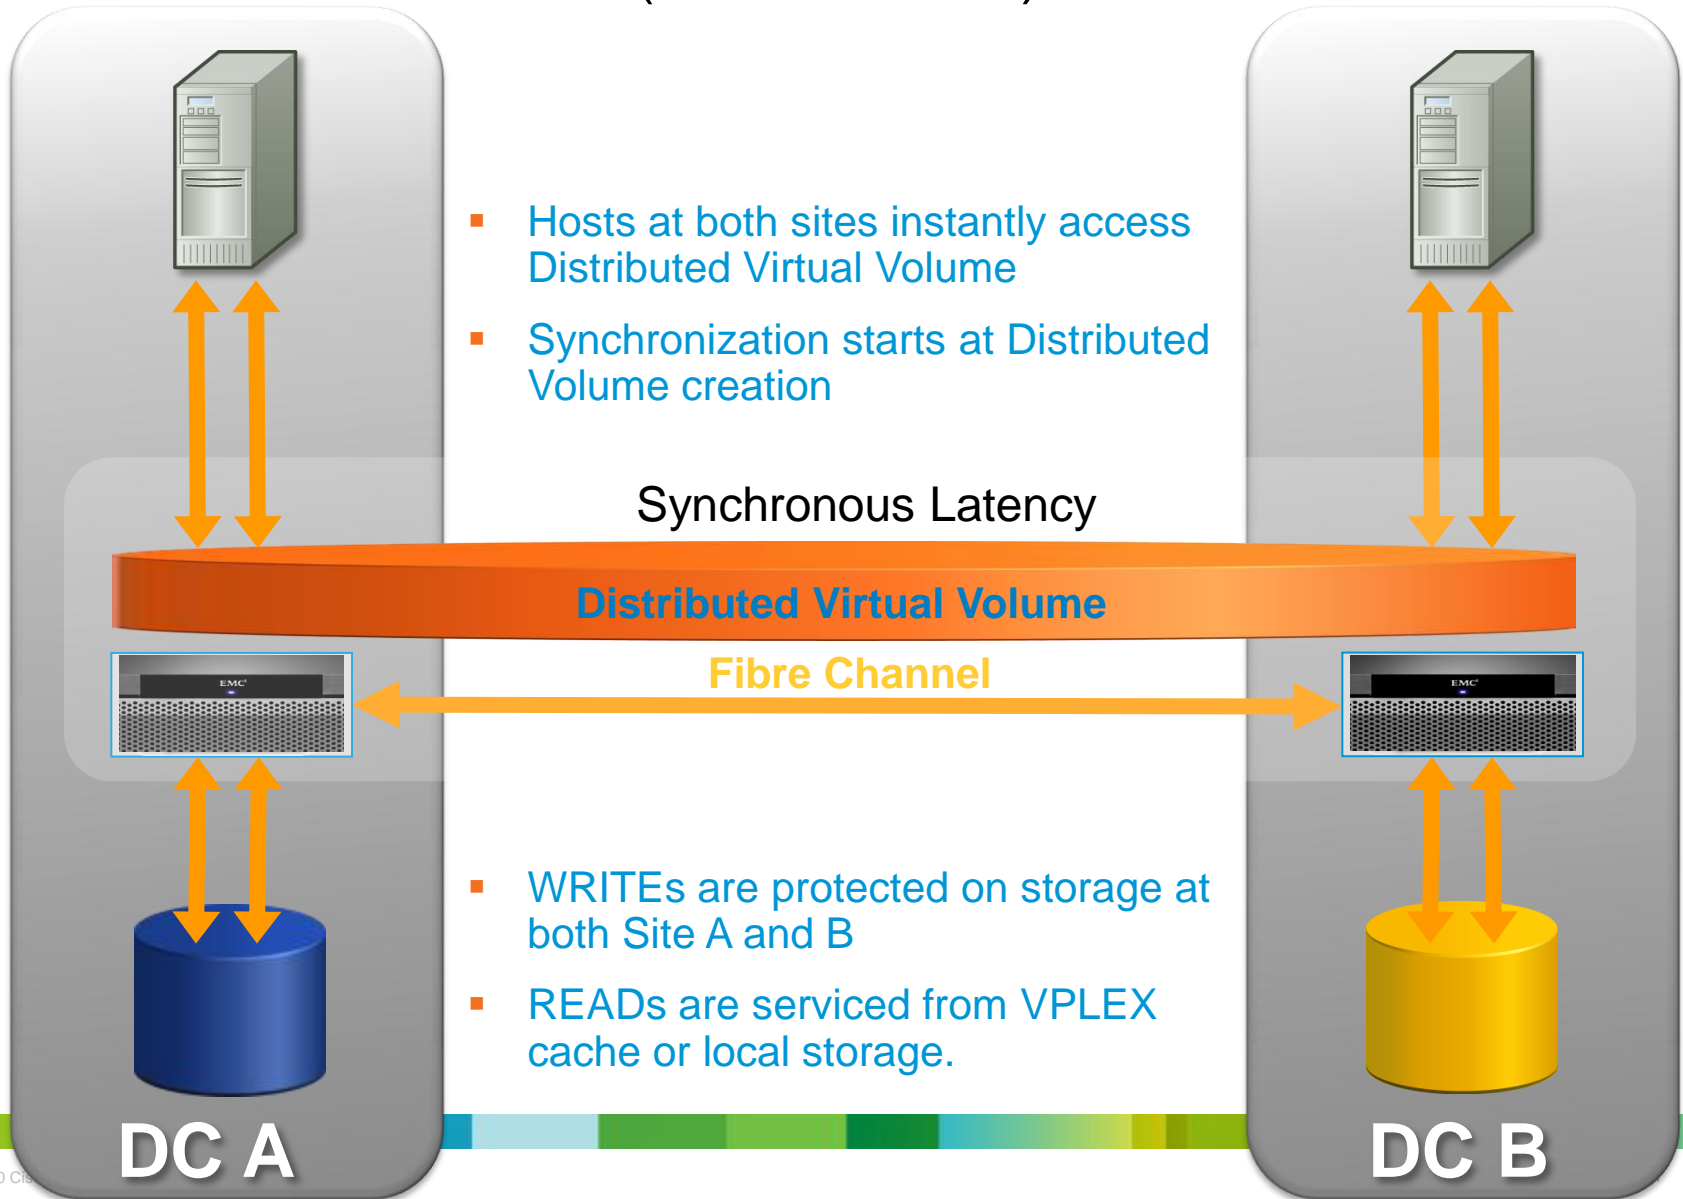
vmware



Virtual Center

Storage Deployment in DCI

EMC VPLEX Metro (Active/Active)

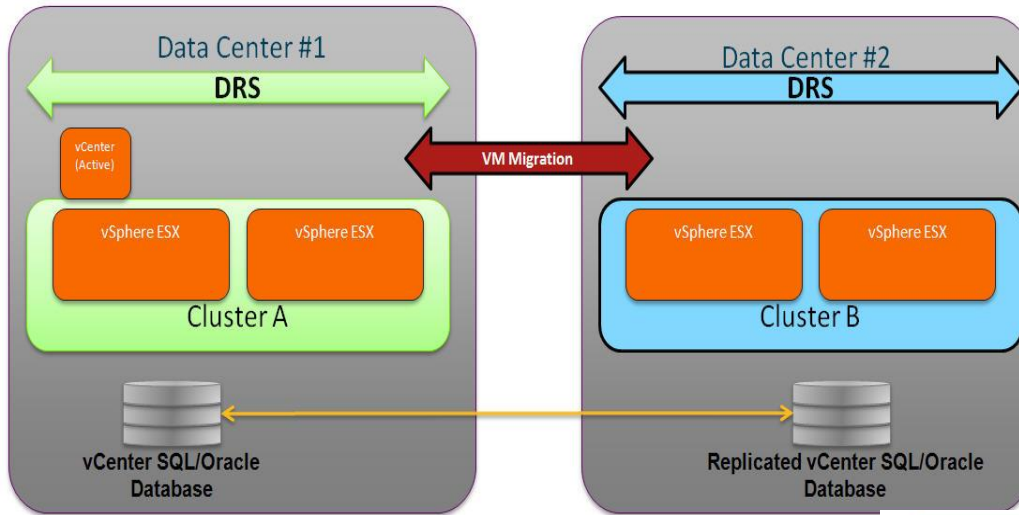


Agenda

- Business Drivers
 - Systems Development Unit (SDU)
 - DCI Use Case: Virtualized Workload Mobility
- Virtualized Workload Mobility Components
 - LAN Extensions
 - Path Optimization
 - Storage Elasticity
 - Workload Mobility – Nexus 1000V and VSG

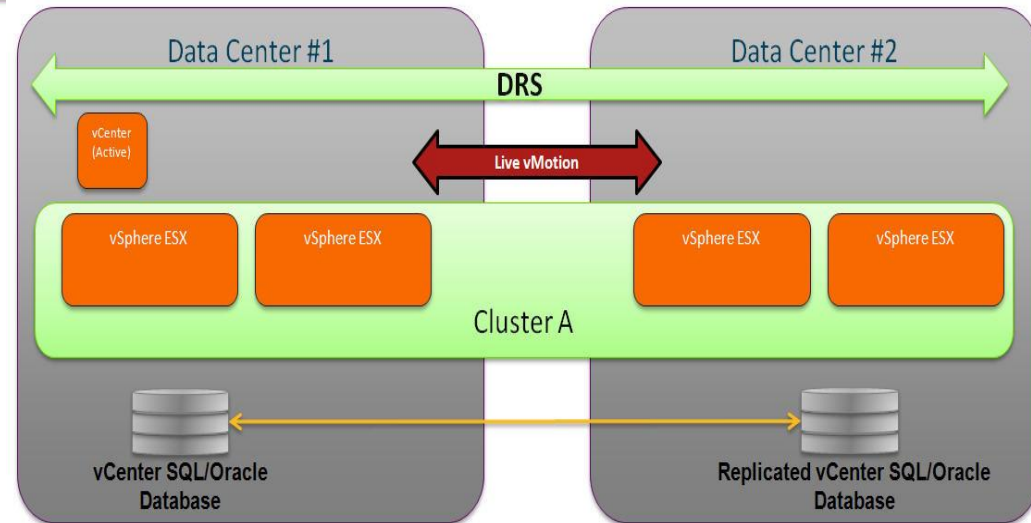


Workload Mobility: ESX Cluster requirements ?



Solution works with Separate clusters

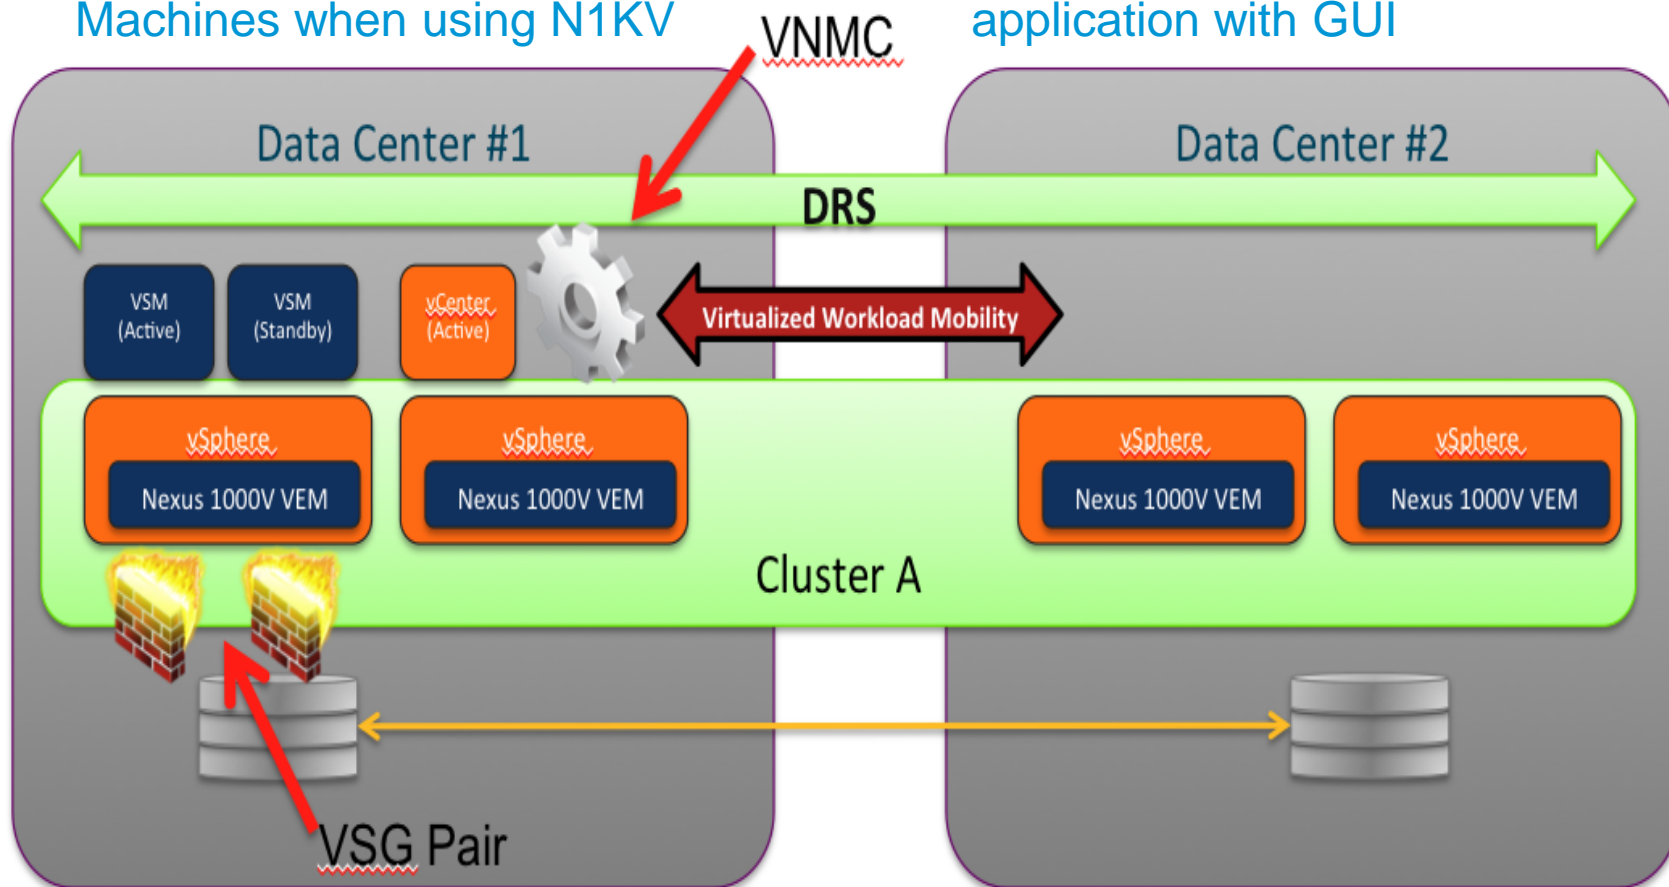
Or Stretched Clusters



Workload Mobility: Stretched VSM & VSG Model

VSM and VSG are Virtual Machines when using N1KV


VNMC is a Network Management application with GUI



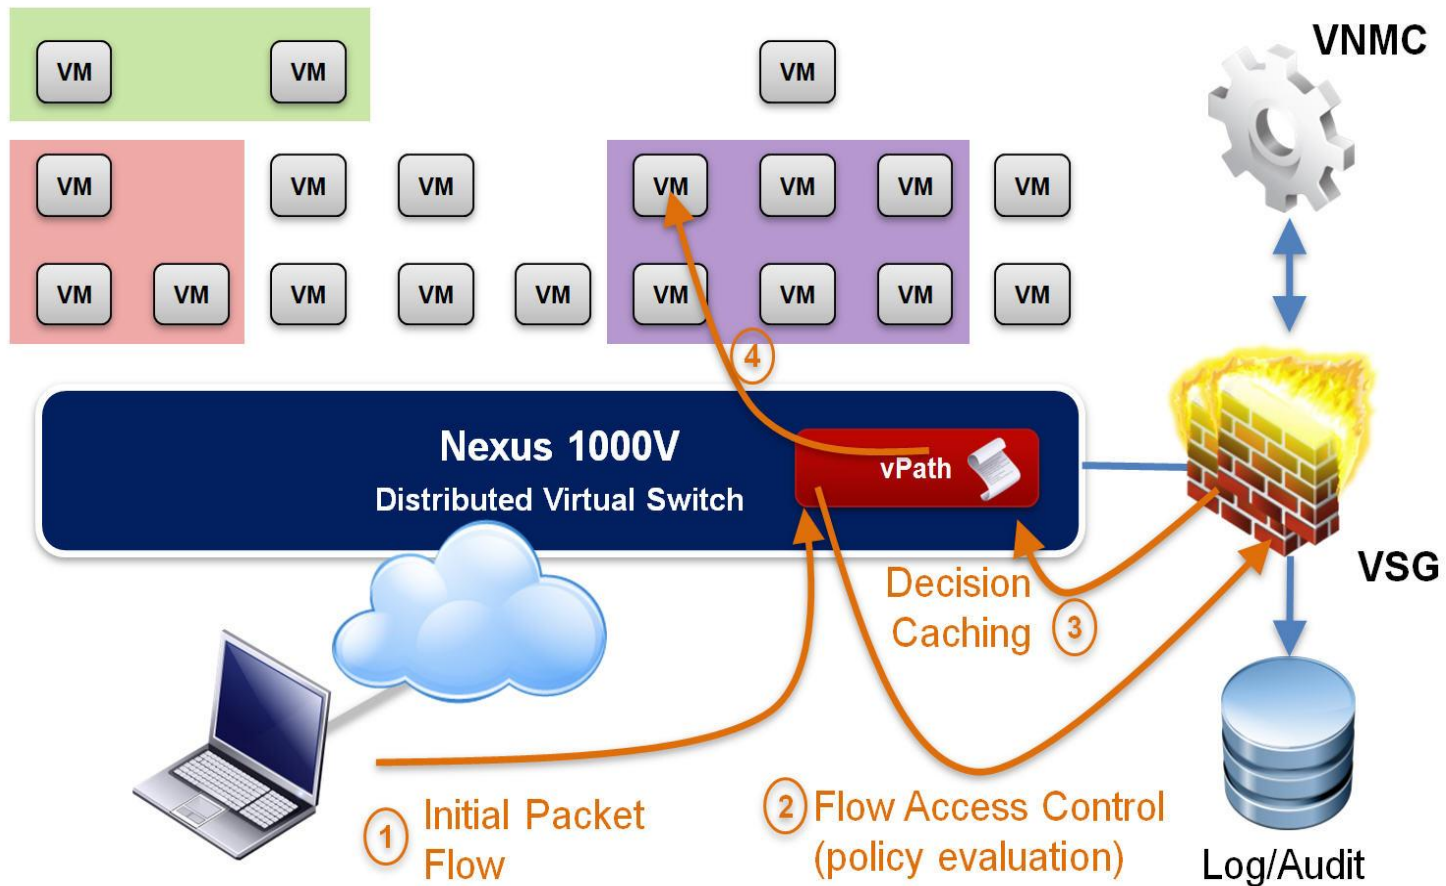
Workload Mobility: VSM to VEM Control Plane

Nexus 1000V
VSM Pair

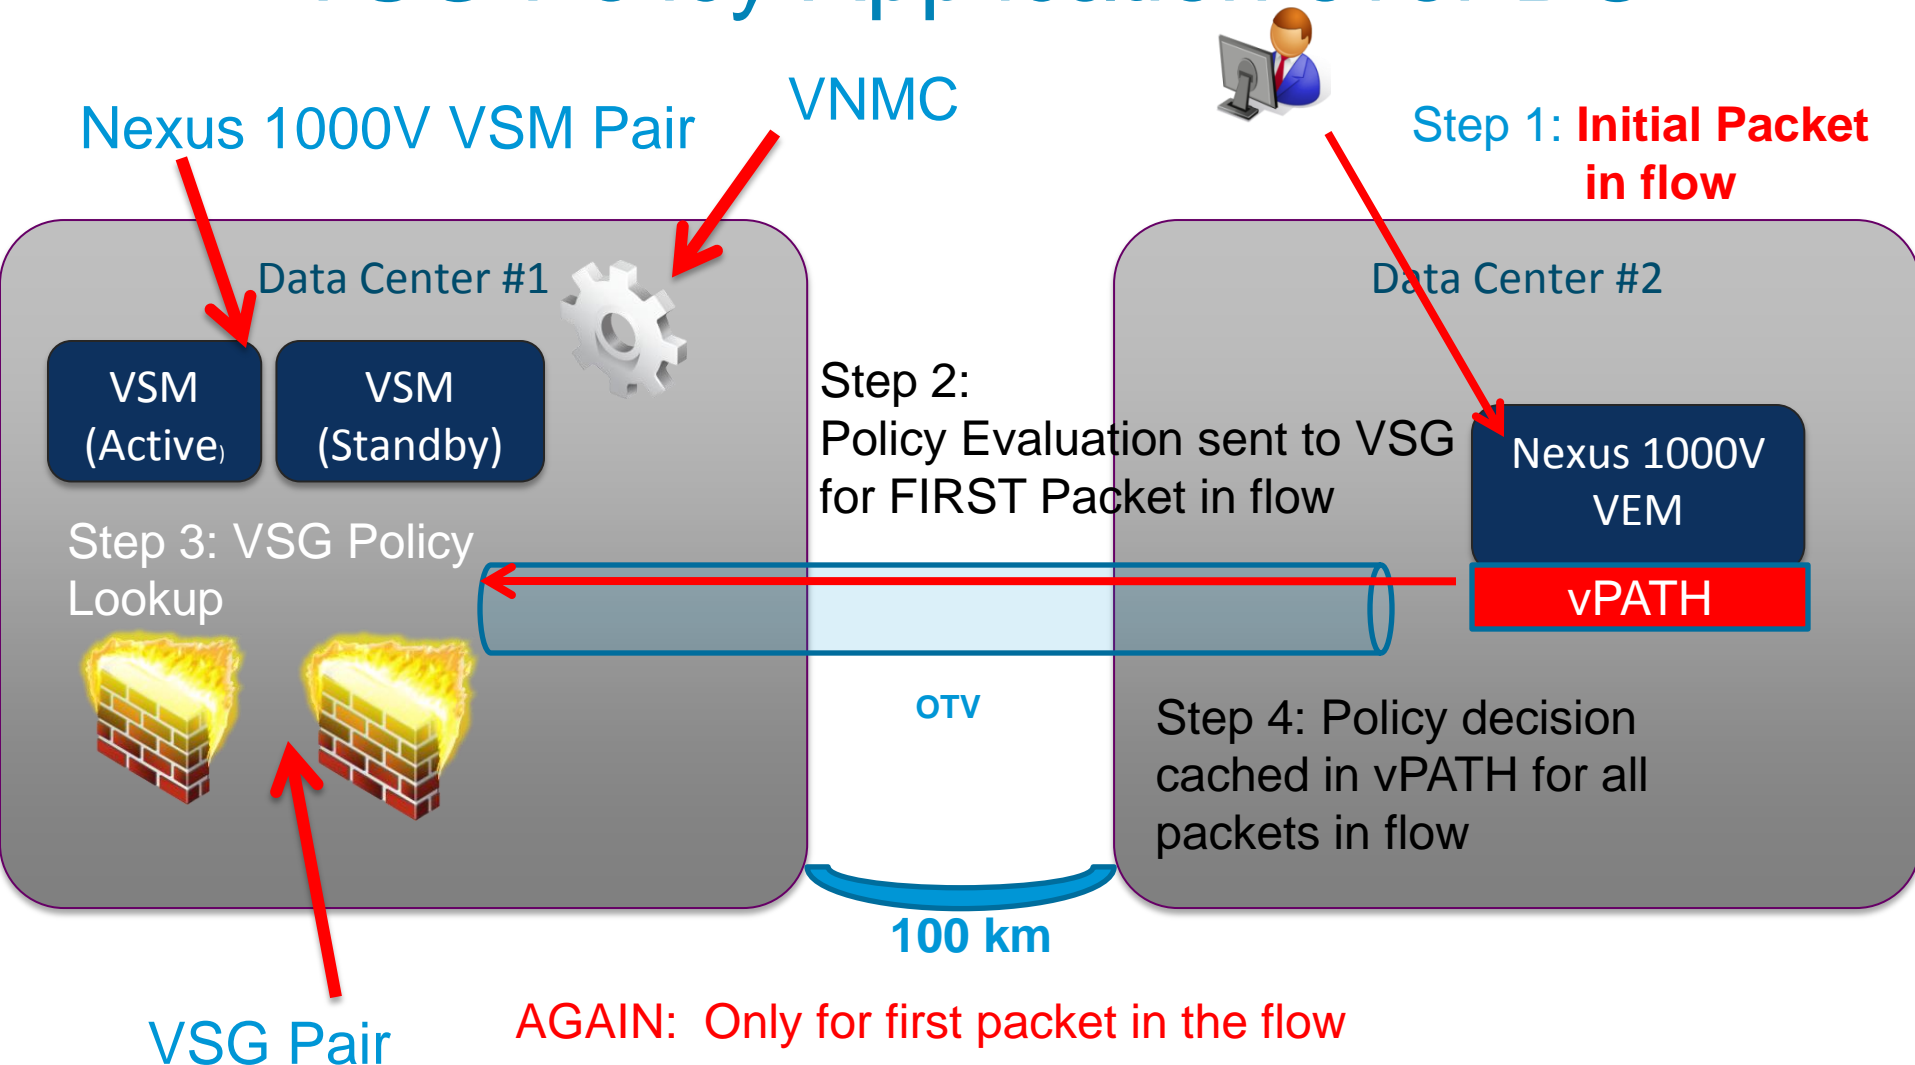
VSM to VEM Control Plane
Connectivity – Can be L2 or L3
Recommended is L3



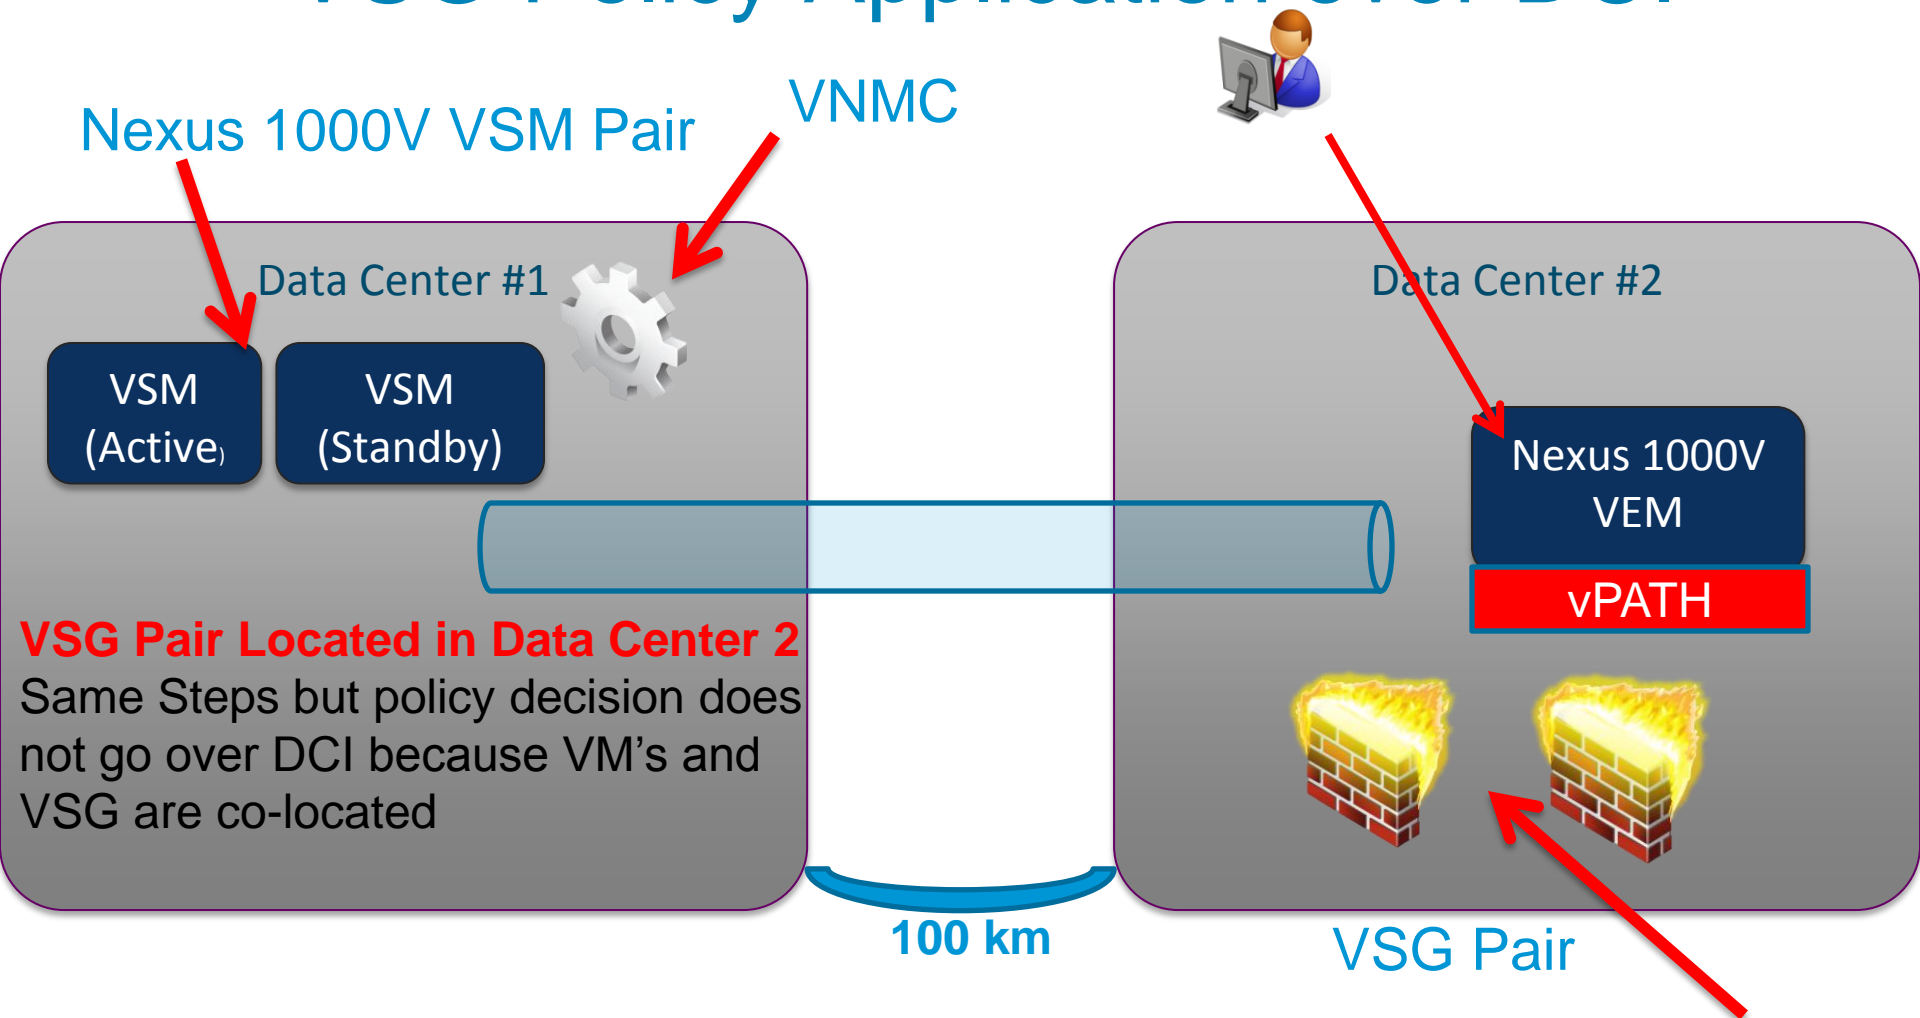
Workload Mobility: VSG Policy Application



Workload Mobility: VSG Policy Application over DCI



Workload Mobility: VSG Policy Application over DCI



AGAIN: Only for first packet in the flow goes to VSG

Workload Mobility - Services

Data Center Services – In Progress

Services Include:
Data Center Perimeter FW
Server Load Balancing
Quality of Service

Live workload mobility

- Maintain state
- Maintain client-server session
- Maintain VM / App policies
- Active / Active Data Centers

Non live workload mobility

- Maintain VM / App policies
- Active / Backup Data Centers (Disaster Recovery sites)

Nexus 1000V/1010 & VSG and Virtualized Workload Mobility

Syed Ghayur



N1K Deployment Requisites

- The following components must be in the **same physical data center**:

Component	Version
Nexus 1000V VSM HA Pair, including VSM VM storage	4.2(1)SV1(4) or later
vCenter Server and vCenter Server Heartbeat pair	VMware vSphere 4.1 or later
VMware Update Manager	VMware vSphere 4.1 or later
Nexus 1010 HA Pair	4.2(1)SP1(3)
Virtual Security Gateway	4.2(1)VSG1(1)
Virtual Network Management Center	VNMC release 1.0.1
IP Storage, if vmk is on N1KV	Depends on underlying storage technology

N1K Deployment Requisites

- The following components can be in the **different physical data centers**:

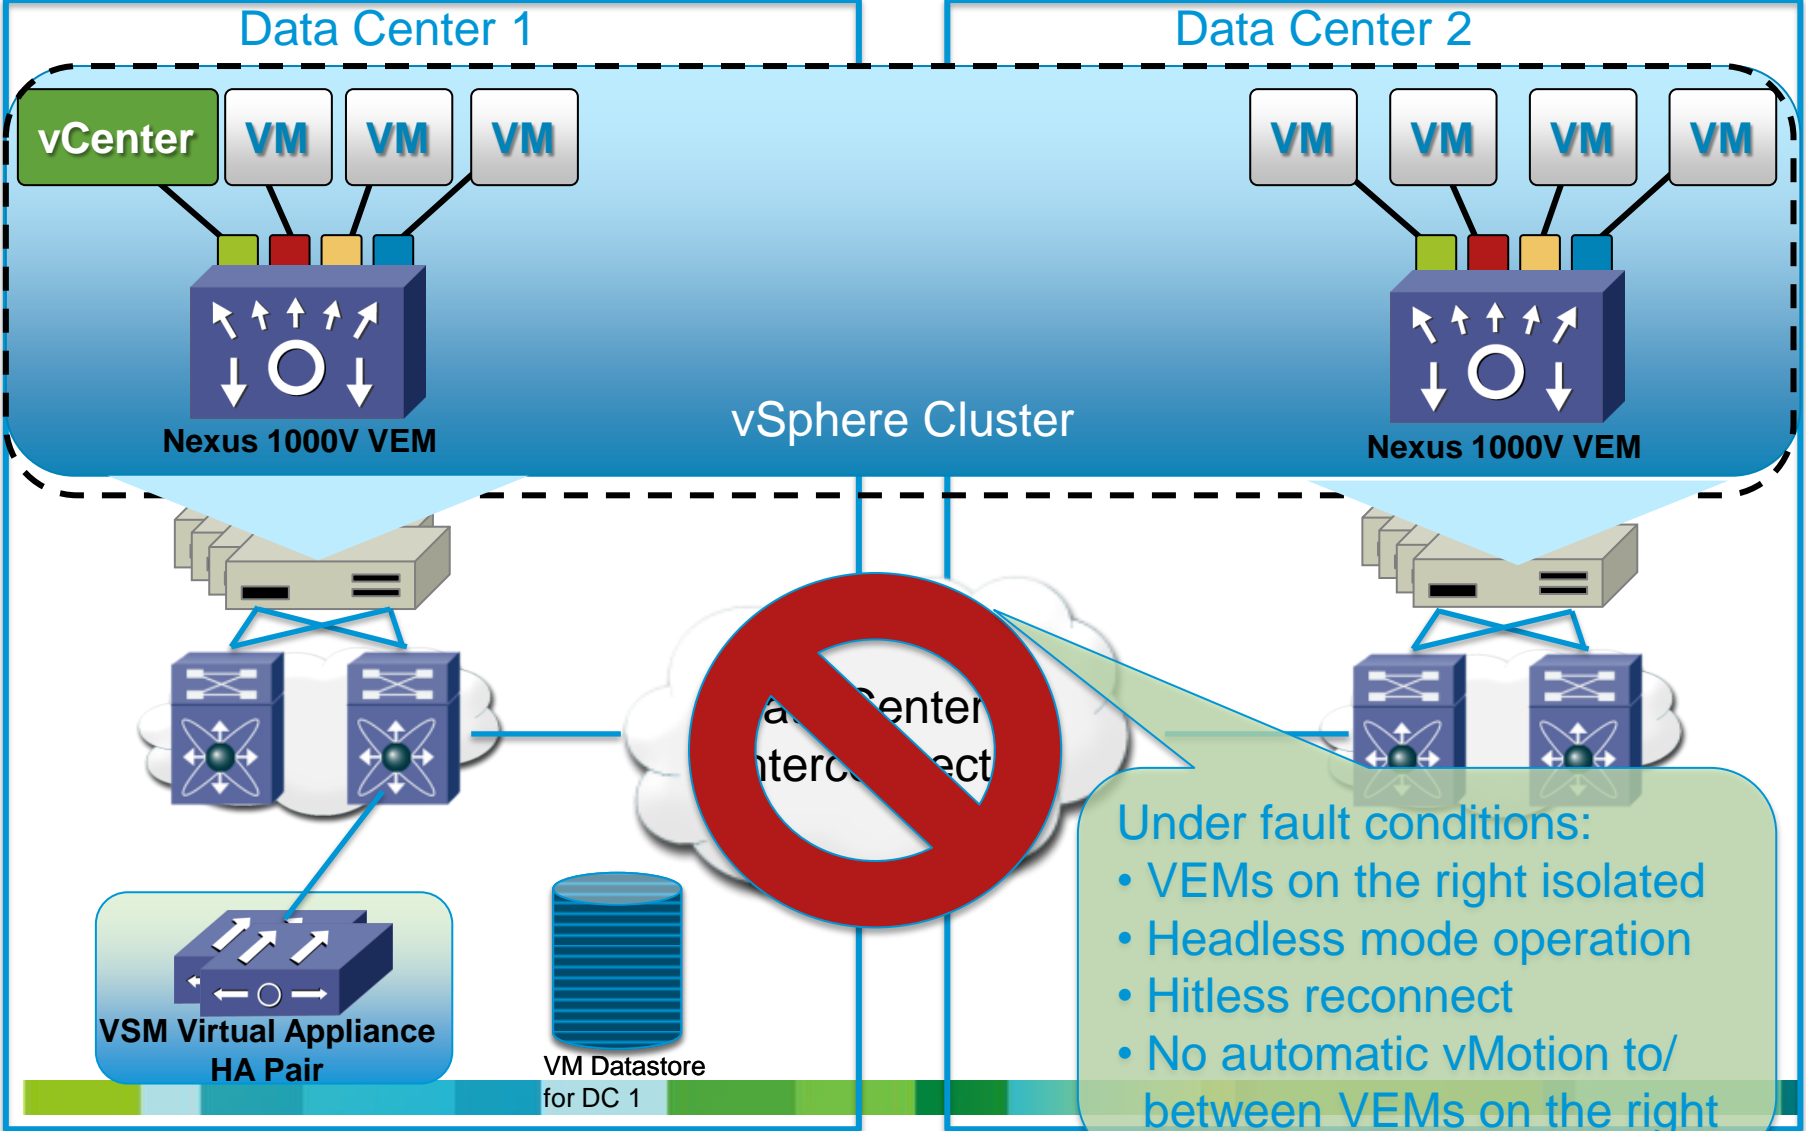
Component	Minimum Version
ESX/ESXi hosts	VMware vSphere 4.1 or later
VEM	4.2(1)SV1(4) or later
VM physical datastores, but must be local to associated VMs	Depends on underlying storage technology

N1K Networking Requisites

- L2 must be extended across data centers
- A variety of interconnect technologies are available
 - OTV and vPC validated by SDO
- DCI link must provide 5 ms or better latency between 2 sites
- Routing in data centers should be robust. This is important for L3 control from VSM to VEM.
 - ie. Routing protocol convergence

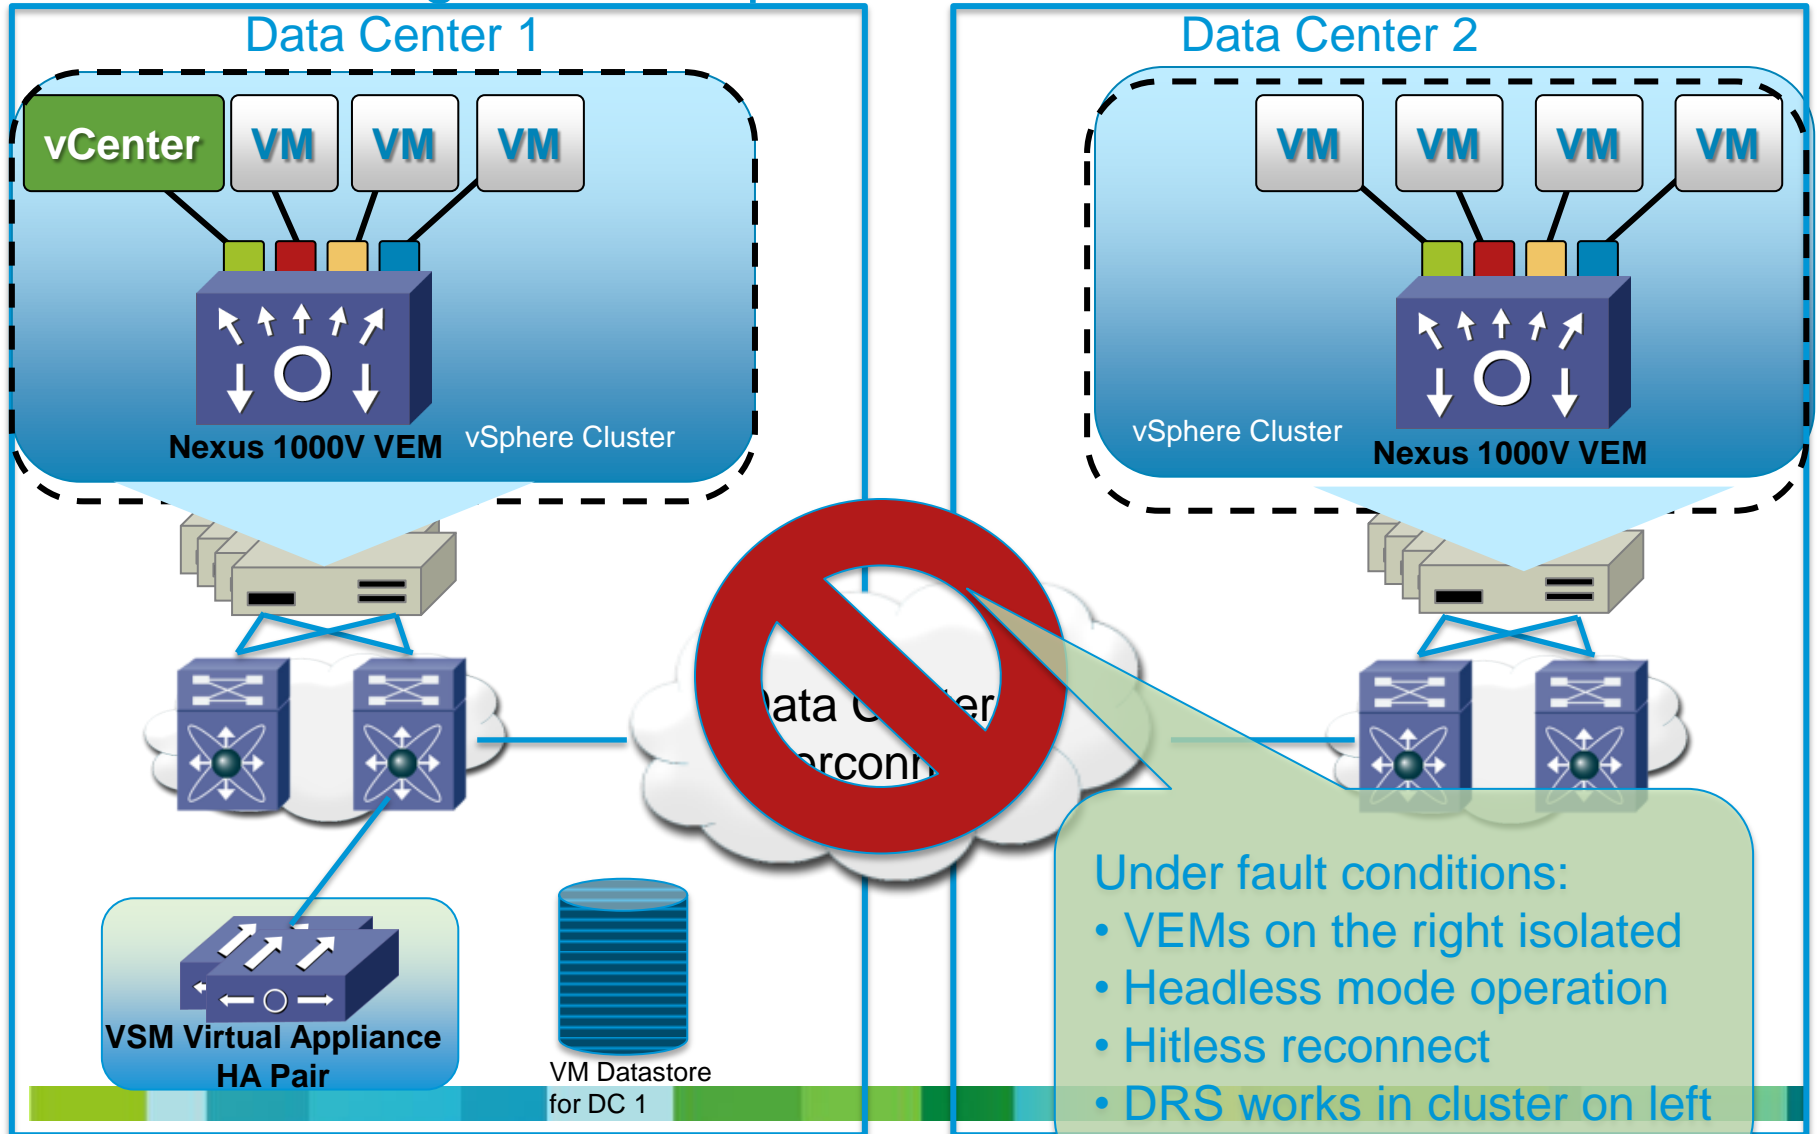
Failure Scenario – Single, Stretched Cluster

ESX/ESXi Stretched Cluster



Failure Scenario – Multi Cluster

ESX/ESXi Single Cluster per Site



N1K Scalability requisites

- Stretched Cluster
 - 32 hosts per cluster (16 hosts/site), if using manual HA primary nodes election
Total of 32 hosts per instance
 - 8 hosts per cluster (4 host/site), if using default HA settings
Maximum 4 stretched clusters/ N1K instance, total of 32
- Single Cluster/Site - 32 total hosts (max 16 hosts/site)
- 1024 ports (across each N1KV DVS)
- Feature scalability at 50% feature numbers
 - ie. 1024 VLANs and port-profiles instead of 2048, 256 Mcast groups instead of 512, etc.

Deployment Best Practices

- Use L3 control for VSM-VEM (more flexibility)
- Prioritize N1KV control/management/packet traffic with QoS across DCI
- Use LACP offload (if topology requires LACP)
- Review 4.2(1)SV1(4a) **release notes** for headless mode caveats)

Verifying Latency

For L3 Control From VMS to VEM, can get estimated min/avg/max latency, should be $\leq 5\text{ms}$

- Ping from VSM to L3 control vmk intf
- Ping from VEM with 'vmkping' to VSM L3 control intf (mgmt0 or control0)

```
~ # vmkping 10.29.170.100
PING 10.29.170.100 (10.29.170.100): 56 data bytes
64 bytes from 10.29.170.100: icmp_seq=0 ttl=254 time=0.495 ms
64 bytes from 10.29.170.100: icmp_seq=1 ttl=254 time=0.555 ms
64 bytes from 10.29.170.100: icmp_seq=2 ttl=254 time=0.386 ms

--- 10.29.170.100 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.386/0.479/0.555 ms
```

Verifying Latency

For L2 Control, you can use “mping” to estimate latency, which should be $\leq 5\text{ms}$

```
BL1-VSM# mping broadcast
64 bytes from node 0x0101 (msg id = 0x816e052 1) (time=0 sec, 275 usec)
64 bytes from node 0x0201 (msg id = 0x0cfcdb 1) (time=0 sec, 904 usec)
64 bytes from node 0x0302 (msg id = 0x957607 1) (time=0 sec, 1179 usec)
10.29.176.159
64 bytes from node 0x0402 (msg id = 0x92a1e5 1) (time=0 sec, 1740 usec)
10.29.176.160
64 bytes from node 0x0101 (msg id = 0x816e086 2) (time=0 sec, 609 usec)
64 bytes from node 0x0201 (msg id = 0x0cfcbe 2) (time=0 sec, 624 usec)
64 bytes from node 0x0302 (msg id = 0x957609 2) (time=0 sec, 635 usec)
10.29.176.159
64 bytes from node 0x0402 (msg id = 0x92a1e7 2) (time=0 sec, 648 usec)
10.29.176.160
64 bytes from node 0x0101 (msg id = 0x816e097 3) (time=0 sec, 250 usec)
64 bytes from node 0x0201 (msg id = 0x0cfcbf 3) (time=0 sec, 630 usec)
64 bytes from node 0x0302 (msg id = 0x95760b 3) (time=0 sec, 643 usec)
10.29.176.159
64 bytes from node 0x0402 (msg id = 0x92a1e9 3) (time=0 sec, 653 usec)
10.29.176.160
64 bytes from node 0x0101 (msg id = 0x816e0aa 4) (time=0 sec, 249 usec)
64 bytes from node 0x0201 (msg id = 0x0cfcc0 4) (time=0 sec, 672 usec)
64 bytes from node 0x0302 (msg id = 0x95760d 4) (time=0 sec, 684 usec)
10.29.176.159
```

Common Verification Commands

This is the same as local DC deployment

```
show system redundancy status
```

```
show inventory
```

```
show module
```

```
show module uptime
```

```
show module vem mapping
```

```
show module vem missing
```


```
show svs connection
```

```
show svs neighbors
```

VXLAN

Virtual Extensible Local Area Network (VXLAN)

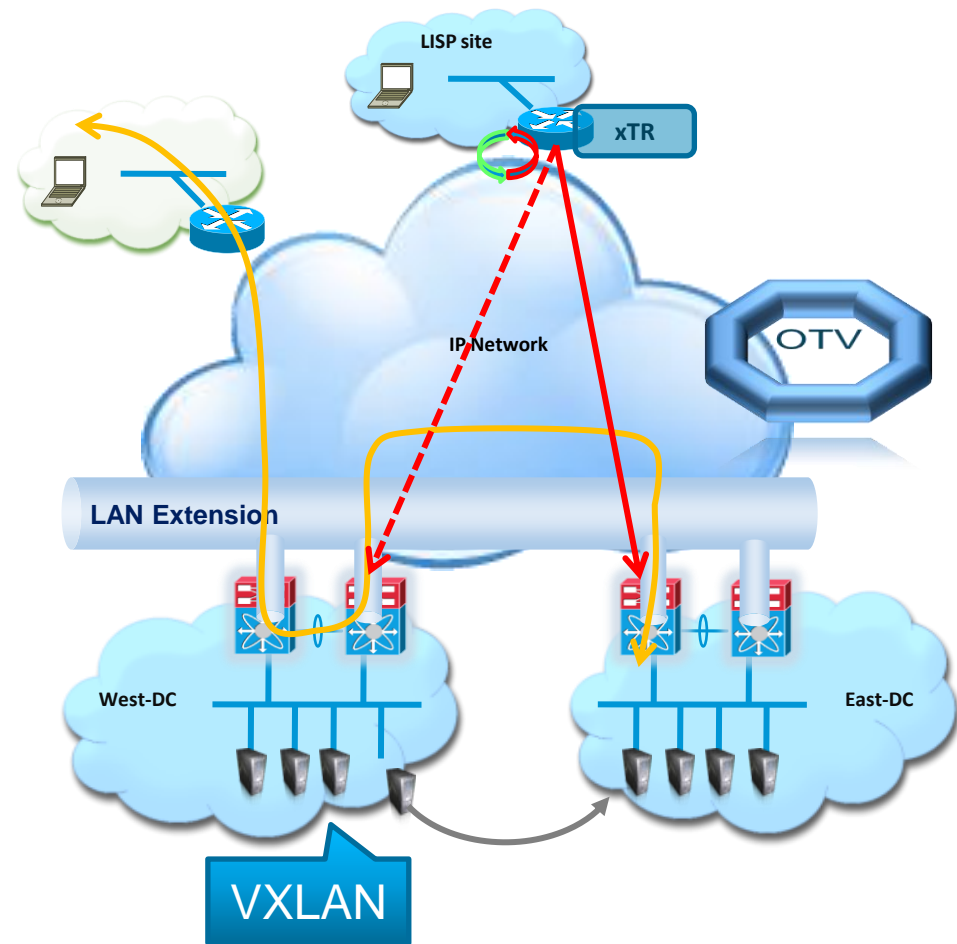
- Ethernet in IP overlay network
 - Entire L2 frame encapsulated in UDP
 - 50 bytes of overhead
- Include 24 bit VXLAN Identifier
 - 16 M logical networks
 - Mapped into local bridge domains
- VXLAN can cross Layer 3
- Tunnel between VEMs
 - VMs do NOT see VXLAN ID
- IP multicast used for L2 broadcast/multicast,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



VXLAN, DCI, and LISP

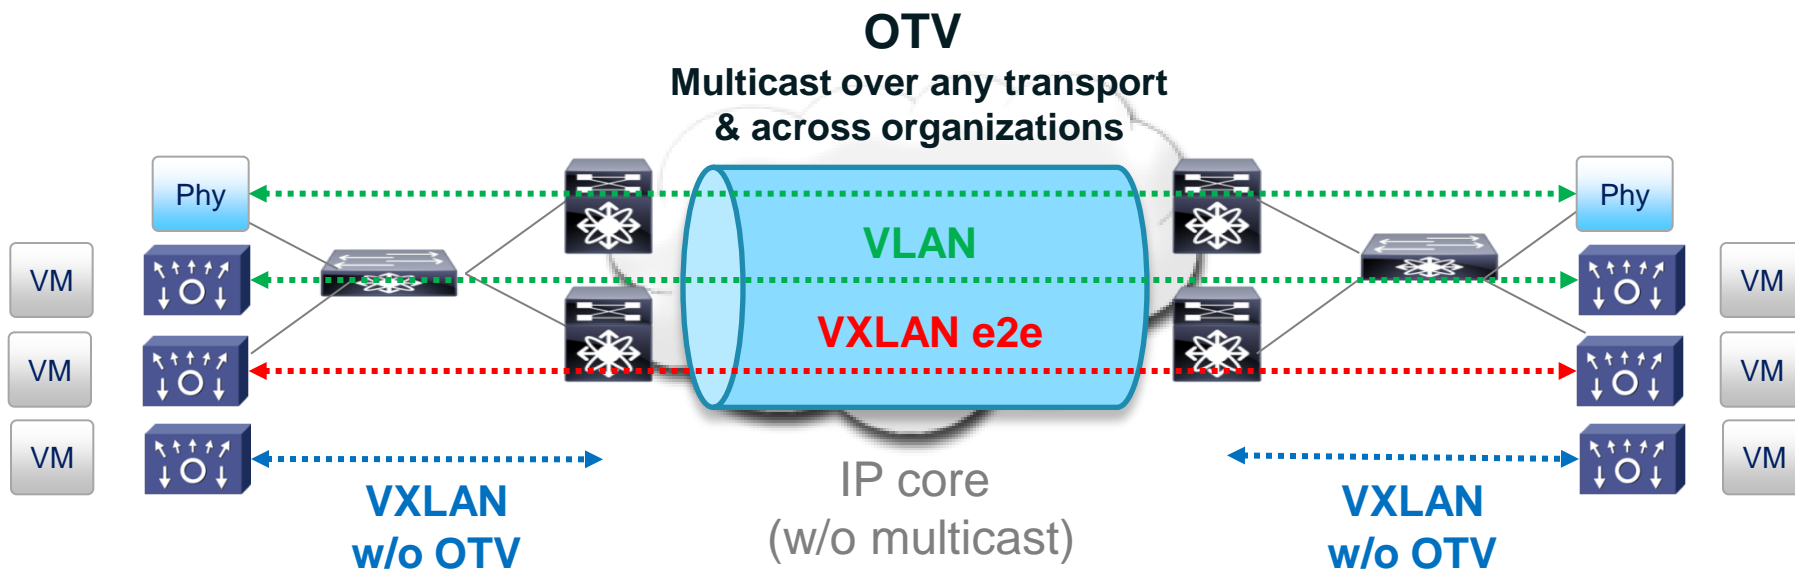
- VXLAN: For scaling logical networks and crosses L3 boundaries within a data center
- Data Center Interconnect, such as OTV, extend Layer 2 domains between data centers
- Locator ID Separation Protocol (LISP): Dynamically reroute IP address based on location of VM

Complementary Advanced Networking Capabilities



LAN Extensions: VXLAN over OTV

- OTV is part of the IP core and provides
 - An e2e multicast transport over any network
 - Connectivity across different Autonomous Systems (organizations)
- VXLAN is transported over OTV to enable e2e connectivity
- Non-VXLAN hosts use OTV LAN extensions natively



Resources



Data Center Interconnect

Where to Go for More Information

Data Center Interconnect General:

<http://www.cisco.com/go/dci>

http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/white_paper_c11_493718.html

Nexus 7000 OTV Specific:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/DCI/whitepaper/DCI3_OTV_Intro_WP.pdf

TechWiseTV and Youtube Videos

http://www.cisco.com/en/US/solutions/ns340/ns339/ns638/ns914/html_TWTV/twtv_episode_69.html

<http://www.youtube.com/watch?v=H5f5Q6UgvnQ&feature=related>

Cisco Validated Designs

<http://www.cisco.com/go/datacenter>

http://www.cisco.com/en/US/partner/solutions/ns340/ns414/ns742/ns743/ns749/landing_site_selection.html

Nexus 1000V Resources

<http://www.tinyurl.com/N1k-Resources>

<https://communities.cisco.com/docs/DOC-24984>

Reference Guides

With Nexus 1000V, Nexus 1010 and VSG

- [vBlock with Nexus 1000V](#)
- [FlexPOD with Nexus 1000V and Nexus 1010](#)
- [Virtual Multi-tenant Data Center with Nexus 1000V](#)
- Virtual Desktop
 - [1000V and VMware View](#)
 - [1000V and Citrix XenDesktop](#)
 - [1000V and VSG in VXI Reference Architecture](#)
- Virtual Workload Mobility (aka Long-distance vMotion)
 - [Cisco, VMware and EMC \(with 1000V and VSG\)](#)
 - [Cisco, VMware and NetApp \(with 1000V and VSG\)](#)
- [PCI 2.0 with Nexus 1000V and VSG](#)

On Cisco.com: 1000V / 1010 / VSG

- CCO Links
 - 1000V: www.cisco.com/go/1000v
 - 1010: www.cisco.com/go/1010
 - VSG: www.cisco.com/go/vsg
 - VNMC: www.cisco.com/go/vnmc
 - vWAAS: www.cisco.com/go/waas
 - NAM on 1010: <http://www.cisco.com/en/US/products/ps10846/index.html> (or www.cisco.com/go/nam)
- My Cisco Community: www.cisco.com/go/1000vcommunity
- Deployment Guides
 - [Nexus 1000V Deployment Guide](#)
 - [Nexus 1000V on UCS – Best Practices](#)
 - [Nexus 1010 Deployment Guide](#)
 - [VSG Deployment Guide](#)
- White papers:
 - [Nexus 1000V and vCloud Director](#)
 - [N1K on UCS Best Practices](#)
 - [Nexus 1000V QoS White paper \(draft\)](#)
 - [VSG and vCloud Director \(draft\)](#)
 - [vWAAS Technical Overview, vWAAS for Cloud-ready WAN Optimization](#)

N1K Public Webcasts

Date	Business Track Topics	Webinar	Preso	Q&A
3/22	Nexus 1000V/1010 Overview and Update	Play	PDF	PDF
4/05	Virtual Network Services: Virtual Service Datapath (vPath), Network Analysis Module (NAM), Virtual Application Acceleration (vWAAS)	Play	PDF	PDF
4/19	Virtual Security Gateway (VSG) Overview (Installation Videos: Link)	Play	PDF	PDF
5/03	Journey to the Cloud w/ N1KV: vCloud Director & Long Distance vMotion	Play	PDF	PDF
5/17	Secure Virtual Desktop with Nexus 1000V & VSG	Play	PDF	PDF

Date	Technical Track Topics	Webinar	Preso	Q&A
3/29	Nexus 1000V v1.4 Features & Install Overview (Installation Screencasts Link)	Play	PDF	PDF
4/12	Nexus 1010 Overview & Best Practices	Play	PDF	PDF
4/26	Virtual Security Gateway (VSG) Technical Overview	Play	PDF	PDF
5/10	Nexus 1000V Key Features Overview	Play	PDF	PDF
5/24	Nexus 1000V Troubleshooting	Play	PDF	PDF
7/27	Long Distance vMotion with Nexus 1000V and VSG	Play	PDF	
8/10	PCI Reference Architecture with Nexus 1000V and Virtual Security Gateway	Play	PDF	

Webinar Link: www.cisco.com/go/1000vcommunity

Additional Links

- N1K Download and 60-day Eval: www.cisco.com/go/1000vdownload
- N1K Product Page: www.cisco.com/go/1000v
- N1K Community: www.cisco.com/go/1000vcommunity
- N1K Twitter www.twitter.com/official_1000V
- N1K Webinars: www.tinyurl.com/1000v-webinar
- N1K Case Studies: www.tinyurl.com/n1k-casestudy
- N1K Whitepapers www.tinyurl.com/n1k-whitepaper
- N1K Deployment Guide: www.tinyurl.com/N1k-Deploy-Guide
- VXI Reference Implementation: www.tinyurl.com/vxiconfigguide
- N1K on UCS Best Practices: www.tinyurl.com/N1k-On-UCS-Deploy-Guide

Cisco Cloud Lab

Hands On Training & Demos

- Hands on labs available for Nexus 1000V and VSG in Cloud Lab
- <https://cloudlab.cisco.com>
- Open to all Cisco employees
 - Customers/Partners require sponsorship from account team for access via CCO LoginID
 - Extended duration lab licenses for 1000V and VSG are available upon request



Welcome to Cisco CloudLab

Please select one of the available labs, by clicking on its name. Hover over the lab name content.

Available labs:

- Cisco Nexus 1000V - Basic Introduction (N1K-000111)
- Cisco Nexus 1000V - Installation (N1K-000211)
- Cisco Nexus 1000V - Upgrade to 1.4 (N1K-000310)
- Cisco Virtual Security Gateway (VSG) - Introduction (VSG-000110)
- Cisco Nexus 7000 - Introduction to NX-OS (N7K-000110)
- Cisco Overlay Transport Virtualization (OTV) (N7K-000210)
- Demo: Cisco Nexus 1000V (Pre-Configured) (N1K-100111)
- Demo: Cisco Virtual Security Gateway (VSG)(Pre-Configured) (VSG-100110)

Thank you.

