INX

Cloud Computing: IPMA
May 18, 2010

INX

Agenda

Introduction

Where We Have **BEEN**: A Look **BACK**

Where We **ARE**: Taking A Look **AROUND**

Where We **ARE GOING**: Taking A Look **AHEAD**

What: Live Demo

How: Data Center Considerations

Q&A

INX

Introduction

Joe Rinehart, CCNP/DP/VP, CCIE#14256

- **Senior Consulting Engineer for INX**
- **12+ Years in Networking & Telecommunications**
- **AT&T for 5.5 Years**
- **Lived in Northwest for 14 Years**
- **Married with 7 Children**
- **President of the Seattle Cisco Users Group**
- **Outdoor Enthusiast**
- **Science Fiction & Nonfiction Writer**

2

INX

The Evolution of Computing

In the Beginning…There Was the MAINFRAME!

2

INX

The Evolution of Computing

The Mainframe Age

- Centralized Computing
- Resource Intensive
- Limited Memory
- Extensive Disk Operations
- Dumb Terminals for Communication
- Networking Added Later
- Proprietary
- Expensive

IBM

3

INX

The Evolution of Computing

The Personal Computer Age

- Distributed Computing (servers)
- Less Resource Intensive
- Less Memory Demand
- Moderate Disk Operations
- Smart Terminal Capability
- Networking Assumed (Ethernet, Token Ring)
- Open Standards

IBM hp COMPAQ DELL

4

The Evolution of Computing

The Virtual Computer Age

- Centralized Distributed Computing (servers)
- Less Resource Intensive per server
- Less Memory Demand per server
- Extensive Disk Operations
- Smart or Dumb Terminal Access
- Networking Mandated (Ethernet)
- Open Standards

vmware Microsoft

5

The Evolution of Networking

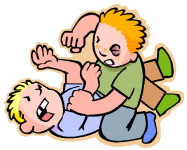

Competing Standards

- Early Networking was Proprietary
- Several Different LAN Media Emerged:
  - ARCNET
  - Ethernet
  - Token Ring
- Several Protocols Emerged Also:
  - Novell NetWare (IPX)
  - TCP/IP
  - Banyan Vines
  - DECNet
  - Appletalk

Cisco Systems hp Novell Apple Computer digital 3Com

6

The Evolution of Networking

TCP/IP and Ethernet Dominance

- Ethernet Eventually Won Out Because:
  - Generally Cheaper to Implement
  - Inexpensive Cabling
  - Much Higher Bandwidth
- TCP/IP Also Had Greater Adoption
- Increased Proliferation of LANs
- Development of Wide Area Networks
- Widespread Adoption of Personal Computers
- Growth of the Internet

hp intel CISCO

7

The Evolution of Networking

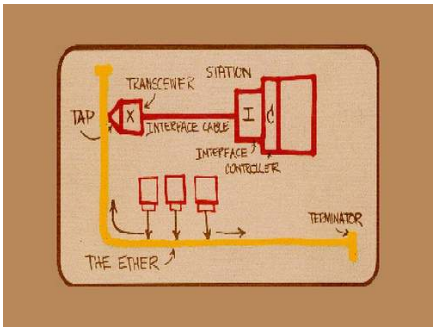

7

The Evolution of Networking

Virtualized Networking

- Logical Components Replacing Physical
- Networks Using Various Connection Methods:
  - Wired (Copper, Fiber LANs)
  - Wireless (802.11a/b/g/n)
  - Cellular (CDMA/WiMAX/UMTS)
  - Broadband (Cable/DSL)
- Networks Using Various Devices (PC, phone)
- Increased Use of the Internet & Remote Access (Telecommuters, Travel, etc.)

hp intel CISCO

8

The Evolution of Networking

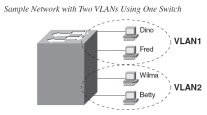

VLANs: Logical/Virtual Separation

- Virtual LANs create ***logically*** separate networks that mimic ***physically*** separate networks
- Ports on the specific access layer switch are associated with specific VLANs
- Broadcasts, IP Subnets, etc., are contained within the specific VLAN
- A Layer 3 Device (L3 switch or router) is required for communication between VLAN segments
- VLANs can represent:
  - Geography: PC's in the same area
  - Function: Workstations/Servers
  - Department: Accounting vs. Marketing
  - Etc.

10

The Evolution of Networking

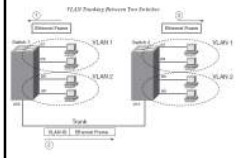

VLAN Trunking: Joining "Islands"

- VLANs can exist across multiple devices
- Communication across devices requires VLAN information to be shared through TRUNKS:
  - ISL Trunk: Cisco Proprietary, New Frame
  - 802.1Q Trunk: IEEE Standard, Tagging
- Trunks carry traffic from multiple VLANs across a single physical wire
- Only a Layer 3 Device can route between VLAN segments
- This is just a simple example to drive the point that ***virtualized networking is not new!***

11

The Evolution of Networking

The "Grid" Compute Model - *solving large compute problems with parallel processing*	Utility Computing – *Offering compute resources as a metered services*	Software and Platform as a Service – *Network based subscription to software & application*	Cloud Computing *Anywhere, anytime access to Virtual Resources delivered as a metered services Foundation for Internal and External Clouds*
Mainframe	ISPs	Single App or Service	Today

Abstraction –> Simplification -> Capability Aggregation

What **IS** Cloud Computing?

RELEVANCE: Cloud Computing is a Popular Topic

164,000,000 results for cloud computing (Yahoo)

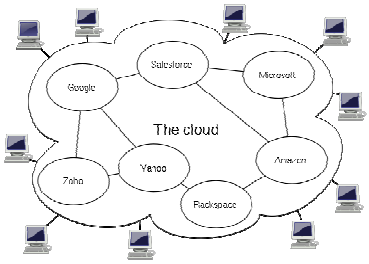

http://en.wikipedia.org/wiki/Cloud_computing

What **IS** Cloud Computing?

"Abstraction, Elasticity, Always On & Pay as You Grow"

- ***Abstraction:***
  - applications from server hardware
  - users from a specific workstation,
  - Networks from specific wires &
  - Data from pre-allocated disk
- ***Elasticity:***
  - *The ability to rapidly to grow or shrink, move and the IT Infrastructure*
- ***Always On:***
  - *Enhanced local and remote recoverability built like a submarine that can sustain outages with no downtime*
- ***Pay as You Grow:***
  - *Add Capacity and Capability without re-wiring or retraining and provision only what is needed when it's needed*

What **IS** Cloud Computing?

- **A Cloud Infrastructure can be in or out of the firewall and comes in 3 *"Flavors"***

1. **Infrastructure as a Service (IaaS)–** All Internal, INX's VHS & VRS, TerraMark, Hosting.com, Rackspace, AT&T and Savvis
2. **Software as a Service –** SalesForce.com, WebEx, Expedia, Google, Microsoft (CRM, Exchange, SQL), Amazon
3. **Platform/Development as a Services –** Azure, EnginYard, Amazon EC2, GigaSpace, Blue Gecko, CohesiveFT

What **IS** Cloud Computing?

The Efficiency of Cloud Computing…

Business

- Fast provisioning response times – improved productivity
- Contractual and auditable SLAs
- Usage base "pay-as-you-go" financial model

50% reduction in OpEx Cost

60% reduction in CapEx Cost

IT

- Economies of shared infrastructure
- High performance, highly available
- Policy-driven automation

An *"Elastic"* IT Infrastructure Appl, Information and People Centric

- Lower Cost – Saves Energy
- Quickly add capacity & capability w/o adding hardware
- Improved recoverability for local and site outages
- Compatible with any existing or future application
- Choice of underlying physical infrastructure
- Future proof – no lock in to specific architectures

What IS Cloud Computing?

View in slideshow

Efficiency	Control	Choice
...to cut capital & operational costs by up to 50% for all applications	...while automating & guaranteeing quality of service	...for flexibility in SW stack, HW, & service providers

Primary Issues – Federation & Compatibility

- *Lack of Compatibility between Clouds*
- *Inability to move application to/from Cloud*
- *Inadequate Security & Governance*
- *Service Level Agreement Remediation*

It's a delicate balance between the security of a dedicated facility and the economies of a shared facility

What IS Cloud Computing?

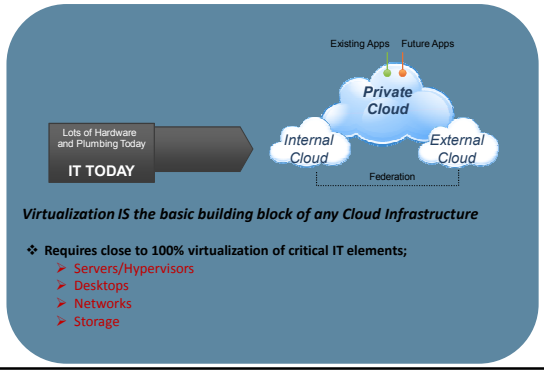

Virtualization IS the basic building block of any Cloud Infrastructure

- **Requires close to 100% virtualization of critical IT elements;**
  - Servers/Hypervisors
  - Desktops
  - Networks
  - Storage

What IS Cloud Computing?

Internal Cloud (IssA)
- Large Server Farm
- Fully Redundant Enterprise SAN
- Replication and De-duplication
- Backups to Secondary Disk Arrays

External IaaS Cloud (Virtual Recovery Service)
- Large Multi-tenant/Shared Server & Storage Infrastructure
- Multiple Carrier Fibre WAN Connections
- Uninterruptable Power and Cooling
- 24x7 System Monitoring and Management

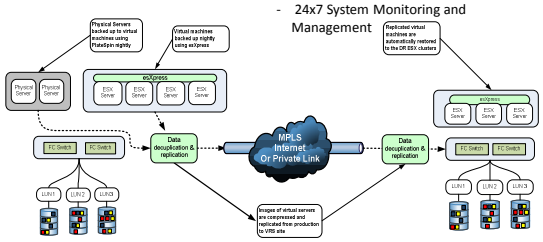

Closer Look: Hybrid Cloud

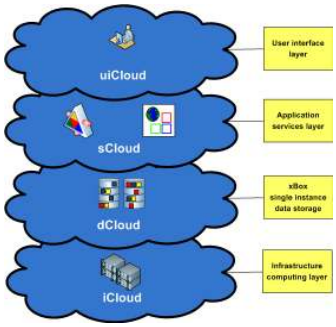

Cloud Risks

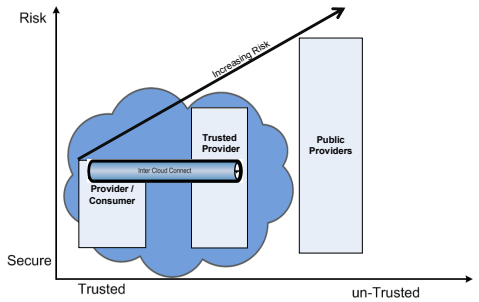

Connecting Clouds

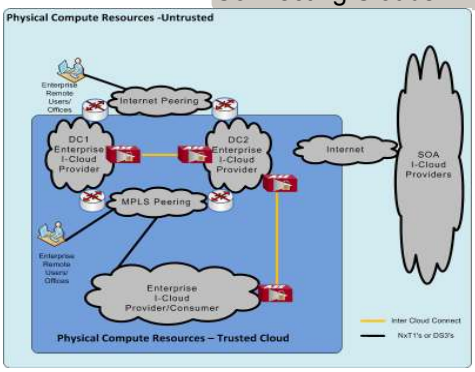

Connecting Clouds

Virtual Fiber - Simplify Wide Area Connectivity

- Leverage lower cost IP, MPLS & Ethernet WANs.
- Save costs compared.
- Multi-Protocol transport for Ethernet and Fibre Channel.
- Encrypt all critical data In-Flight.
- Mitigate packet loss and effects from WAN errors.

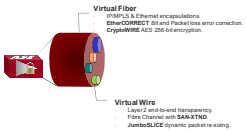

Virtual Wire - Transparent, Optimized Bandwidth

- End-to-end virtual machine migration.
- Reduce protocol latency and round trip delay effects.
- Drive more throughput with existing bandwidth.
- Enable High Availability, Fault Tolerant architectures.
- Optimize path selection on latency and packet loss.

Connecting Clouds

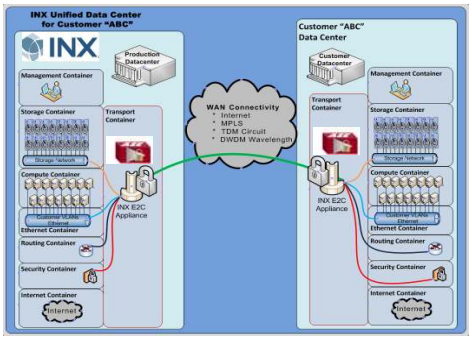

Connecting Clouds

Virtual Machine Live Migration requirements:

- Servers must reside on the same IP subnet or LAN.
- Layer 2 procedures are used to manage VM migrations.
- ARP used to update the network of IP address moves.
- Both source and destination must access common.

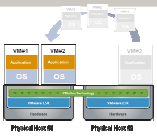

Wide Area Networking Challenges

- IP WANs block Layer 2 protocols.
- LAN & SAN jumbo packets not supported by WAN.
- Inefficient tunneling techniques (i.e. GRE, L2TP & MPLS).
- Data center protocols are intolerant to packet loss or re-ordering.
- SAN protocols must be enhanced to extend over distance.
- WAN traffic may need to be encrypted for compliance.

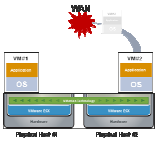

Live Demo

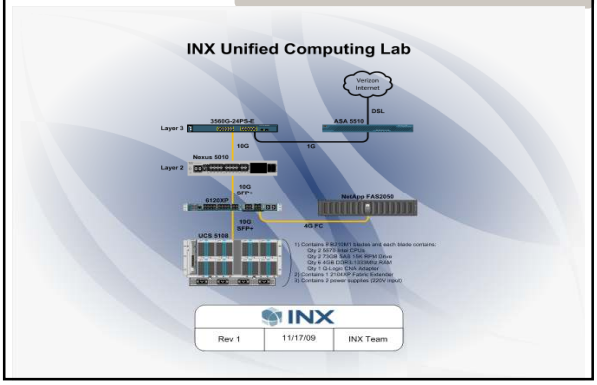

Live Demo: UCS Manager

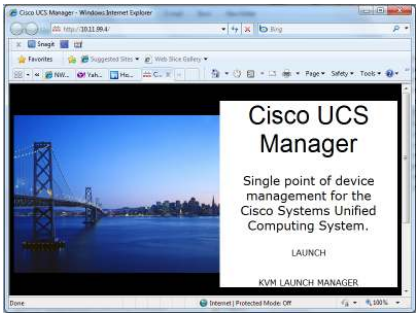

http://10.11.99.4/

Live Demo: Cisco Device Mgr

Live Demo: VCenter/VMware

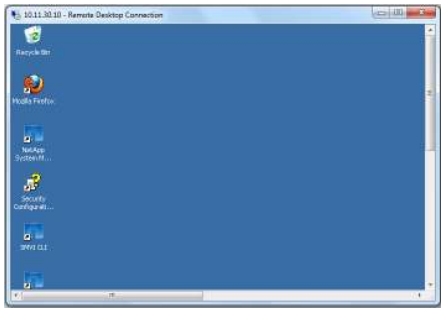

RDP://10.11.30.10

Data Center Considerations

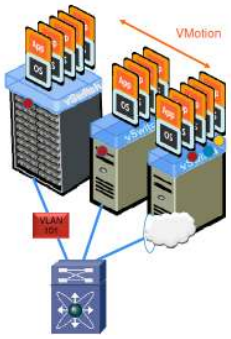

Evolving Storage/Compute Needs

- Server, Storage, Application & Facilities are driving Layer 2 Scalability requirements
- Server Virtualization & Clustering driving the need for every VLAN everywhere based designs
- Facilities requirements defining the network topology "No watt left behind"
- Only a Layer 3 Device can route between VLAN segments
- Increasing drive for SAN connectivity – VM requirements along with Data Storage growth mandating a need for more efficient and pervasive network based storage

12

Data Center Considerations

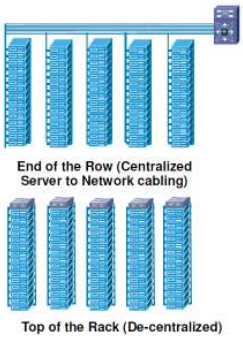

Evolving Facilities Requirements

- Data Center environments responding to facilities demands & technology changes
- Technology changes will impact cabling plant design(s)
  - Migration to 10GE as default technology
  - Migration to 40GE & 100GE within 5 years
- Requirement for a power efficient & flexible access architecture
- Not a choice ***between*** End of Rack and Top of Rack placements

13

Data Center Considerations

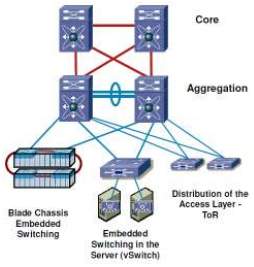

Evolving Access Layer Requirements

- The Access Layer in the Data Center is evolving ***rapidly***
  - Blade Switches
  - Blade Servers
  - Distribution of the Access Layer
- These changes bring challenges:
  - Increase in the size of the access topology (the dreaded STP growth)
  - Understanding High Availability design requirements
  - Increased number of management points
  - Policy Boundaries
  - Growth in Complexity AND Scope

14

Data Center Considerations

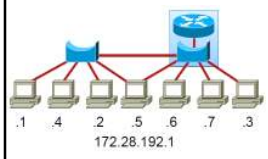

Evolving Access Layer Requirements (why Layer 2 instead of Layer 3?)

- Some protocols require it
- Extend a subnet across Devices
  - For port density (not enough port on device)
  - For provisioning flexibility (add devices without changing L3 network configuration)
  - For redundancy (NIC teaming)
  - Virtual machine mobility

15

Data Center Considerations

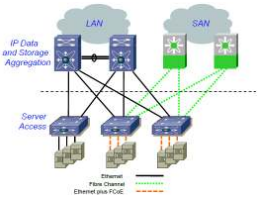

VMOTION L 2 Requirements

At the network level, there are two major requirements for VMotion. ***First, a dedicated, gigabit Ethernet VLAN is needed for VMotion traffic.*** This dedicated VLAN and bandwidth ensures VMs can be moved without impact to users. ***Second, and most importantly, the group of servers that VMs can be balanced between must be in the same Layer-2 domain.*** When a VM moves it cannot change any attributes, like its IP addresses. Thus, all target servers must have connections to the same VLANs as the source VMWare server.

This requirement will definitely lead to an expansion of Layer-2 domains, and thus spanning-tree. VMotion is just too cool and useful to not use. Server teams are going to demand large "VMware Farms" to balance and move VMs around. All of these servers require several physical connections so switch ports will be at a premium. This means more switches and, since all VMware servers must be in the same VLANs, larger spanning-tree fault domains. Spanning-tree and Layer-2 domains were being minimized, or at least significantly controlled, in the last few years in network design best practices. That is likely to begin changing as VMware becomes more mainstream. New technologies like Cisco's VSS will become very important to limit spanning-tree's impact (when it goes wrong...and it will go wrong...or at least one of your engineers will go wrong with it).

http://www.networkworld.com/community/node/24653

16

Data Center Considerations

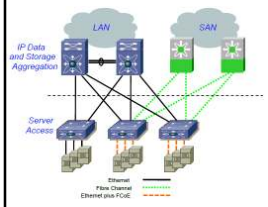

VMOTION L 2 Requirements

When setting up virtual machine networks, consider whether you want to migrate the virtual machines in the network between ESX Server 3i hosts. If so, be sure that both hosts are in the same broadcast domain—that is, the same Layer 2 subnet. ***ESX Server 3i doesn't support virtual machine migration between hosts in different broadcast domains*** because the migrated virtual machine might require systems and resources that it would no longer have access to by virtue of being moved to a separate network. Even if your network configuration is set up as a high availability environment or includes intelligent switches capable of resolving the virtual machine's needs across different networks, you may experience lag times as the ARP table updates and resumes network traffic for the virtual machines.

http://www.vmware.com/pdf/vi3_35/esx_3i_e/r35/vi3_35_25_3i_server_config.pdf

17

Data Center Considerations

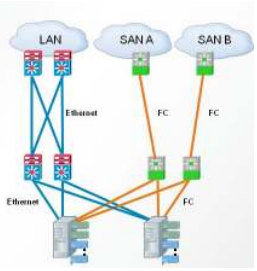

Storage Area Network Requirements

- Storage of data is a major requirements in today's networks
- Many storage products (SAN) utilize block data writing technologies using Fiber Channel
- Fiber Channel requires entirely separate devices and networks
- Fiber Channel data does not tolerate loss or failed transmissions the way that TCP/IP does
- If block data is lost during transmission it corrupts the entire data structure of the storage array
- Separate FC networks are expensive

18

Data Center Considerations

SAN Requirements (Topologies)

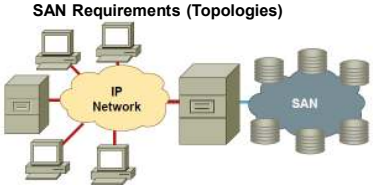

IP Front-End Network

- Host to host
- Application to file system
- Client to server
- NFS, SMB, CIFS, NCP
- NAS, WAFS

SAN Back-End Network

- Host to storage
- Block I/O to storage
- Storage to storage
- SCSI, IDE, NTFS, FAT
- SAN

19

Data Center Considerations

Storage Area Network Requirements

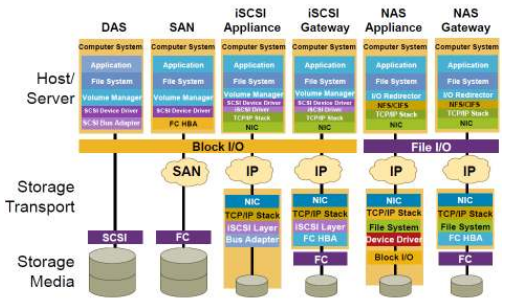

20

Data Center Considerations

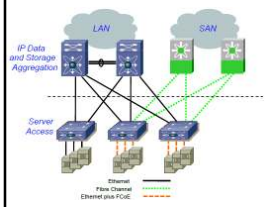

Unified Fabric

- Nexus 5000 Series switches support integration of both IP data and Fiber Channel over Ethernet at the network edge
- FCoE traffic may be broken out on native Fiber Channel interfaces from the Nexus 5000 to connect to the Storage Area Network (SAN)
- Servers require Converged Network Adapters (CNAs) to consolidate this communication over one interface, saving on cabling and power

21

Data Center Considerations

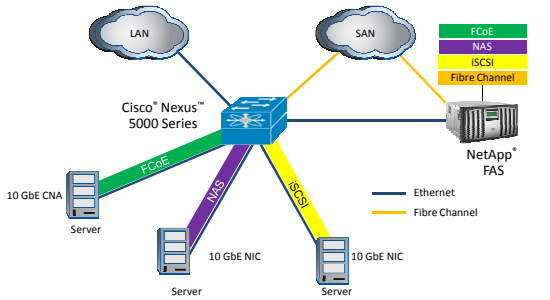

22

Data Center Considerations

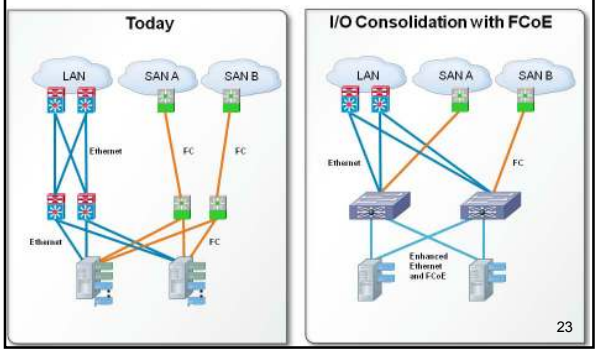

23

Building Blocks of DCVN

Overview: Where Do Resources "Live?"

- **One Physical Network Supports Many VIRTUAL Networks**
- **End Users Are Unaware of the Underlying Infrastructure**
- **Must have a rock-solid campus design in place before adding virtualization to the network**

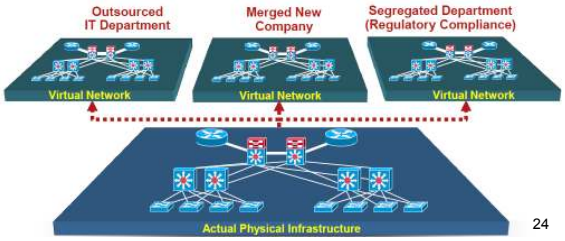

24

Building Blocks of DCVN

Overview: Where Do Resources "Live?"

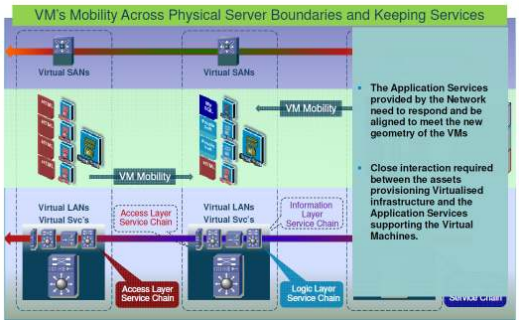

25

Building Blocks of DCVN

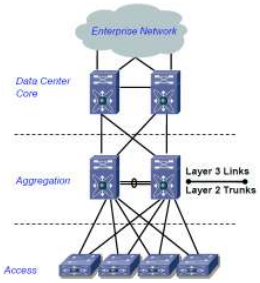

Design Reference Architecture

- **Data Center Core Layer**
  - Routed layer distinct from enterprise network core
  - Provides scalability to build multiple aggregation blocks
- **Aggregation Layer**
  - Provides boundary between L3 routing & L2 switching
  - Point of connectivity for service devices (firewall, SLB, etc.)
- **Access Layer**
  - Provides point of connectivity for servers & shared resources
  - Typically layer-2 switching

26

Building Blocks of DCVN

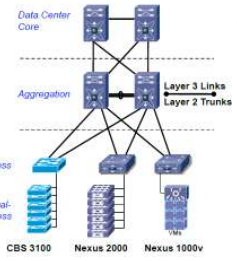

Switching Infrastructure (Access)

- **Virtual Access Layer**
  - Still a single logical tier of L2 switching
  - Common control plane with virtual hardware & software based I/O modules
- **Nexus 2148 Fabric Extender**
  - Switching fabric extender module
  - Acts as a virtual I/O module supervised by Nexus 5000
- **Nexus 1000V**
  - Software-based Virtual Distributed Switch for server virtualization environments
  - Replaces vSwitch in ESX environment

27

Building Blocks of DCVN

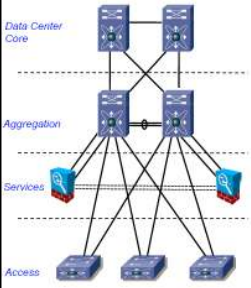

Services Infrastructure

- **Standalone Appliances**
  - Adaptive Security Appliances
  - WAAS
  - WLAN Controllers
  - WAN Access (ISR Series)
  - Etc.
- **Design Implications**
  - Ability to support EtherChannel (for example, ASA does not)
  - May require more cabling for fault tolerance
  - Facility Issues (power, cooling, rack space) possible
  - Additional costs possible

32

Building Blocks of DCVN

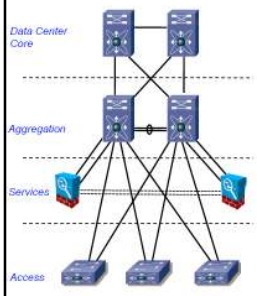

Services Infrastructure

- **Chassis-Based Modules (65XX)**
  - Firewall Services Module
  - ACE Module
  - Wireless Controller
  - Etc.
- **Design Implications**
  - Dual 6500 Chassis design supports multiple service modules
  - Full support for VLANs, EtherChannel, etc.
  - Depending on design/density can alleviate or add facility requirements
  - Dedicated integration points(s) for multiple data center services

33

Building Blocks of DCVN

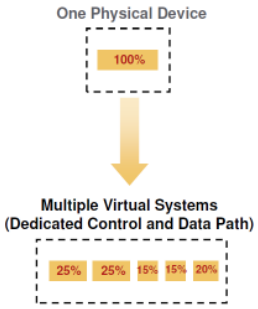

Services Infrastructure

- **Device Partitioning**
  - Distinct context configuration files
  - Separate routing tables
  - RBAC with contexts, roles, domains
  - Independent application rule sets
- **Examples**
  - Server Load Balancing (ACE)
  - Appliance
  - ACE Module

34

Building Blocks of DCVN

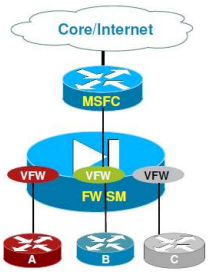

Services Infrastructure

- **Example: Virtual Firewalls**
  - VLANs can be shared if needed
  - Each context has its own policies (NAT, access-lists, inspection engines, etc.)
  - FWSM concurrently supports routed or transparent / stealth virtual firewalls

35

Building Blocks of DCVN

Services Infrastructure: Combined Example

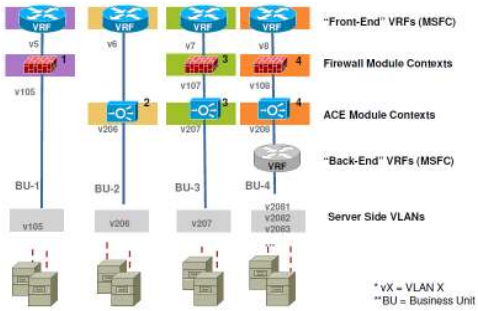

36

Wrap-Up

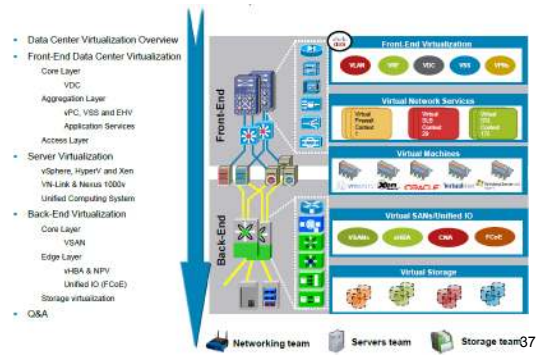

37

Q&A

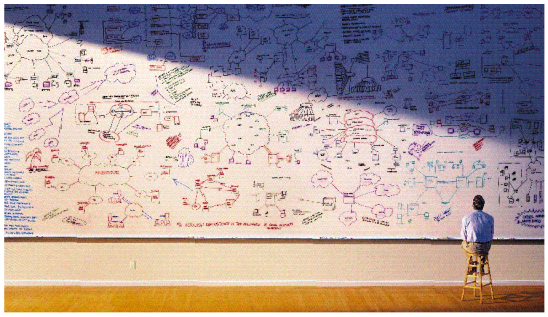

38


Forecast: CLOUDY