

Data Center Segmentation and Virtualization

Dr. Peter J. Welcher,
Chesapeake Netcraftsmen

Cisco University, MD and VA
November, 2007

Chesapeake NETCRAFTSMEN 1 Copyright 2007

About the Speaker

- **Dr. Pete Welcher**
 - Cisco CCIE #1773, CCSI #94014, CCIP
 - Specialties: Network Design, QoS, MPLS, Wireless, Large-Scale Routing & Switching, High Availability, Management of Networks
 - Customers include large enterprises, federal agencies, hospitals, universities, cell phone provider
 - Taught many of the Cisco router/switch courses
 - Reviewer for many Cisco Press books, book proposals
 - Designed and reviewed revisions to the Cisco DESGN and ARCH courses
 - Presented lab session on MPLS VPN Configuration at Networkers 2005, 2006, 2007
- Over 140 articles at <http://www.netcraftsmen.net/welcher/>

Chesapeake NETCRAFTSMEN 2 Copyright 2007

Agenda

- Introduction and Motivation
- Case Study: Cisco
- Case Study: High Availability Enterprise
- Techniques for Segmentation
- Enterprise Case Study – Design
- Data Center and Layer 2
- Controlling User Access to the Data Center
- More Segmentation and Virtualization for the Data Center
- Summary

Chesapeake NETCRAFTSMEN 3 Copyright 2007

Introduction and Motivation

Why Segment, What Objectives?

Chesapeake NETCRAFTSMEN 4 Copyright 2007

Why Segment the Data Center?

- Hardware proliferation – consolidation
- Governance
- Security
- Flexibility and speed of provisioning
- Data Center as Co-lo for business units or customers
- Government data center or shared DR site
- Virtualization

Chesapeake NETCRAFTSMEN 5 Copyright 2007

Design Concerns in Segmenting the Data Center

- Performance
- Scalability
- Complexity
 - MPLS: concern about technical skills and complexity
 - L2 spaghetti with QinQ?
- Ease of management
- L2 reliability, MAN reliability
- Re-addressing servers – NOT!
- High Availability
- Maintenance windows
 - Critical network boxes with many stakeholders get to where they cannot ever be touched for maintenance

Chesapeake NETCRAFTSMEN 6 Copyright 2007

What Are Your Objectives?

- Driving factor: often Governance and Security
 - Key question: exactly how much do you need?
- Goal: Controlling server to server accesses
 - Are your needs more in the ACL / Firewall space?
 - PVLAN, VACL, transparent mode FWSM can segment servers for governance and security purposes, to some extent
 - ACL or firewall is needed anyway to control server access in larger segmentation and virtualization projects: *is more needed?*
- Goal: Grouping servers and other resources (firewalls, server load balancers) by function
 - E.g. separating Production from Dev and Test environments
 - Reduce potential impact of mistakes, changes, Dev and Test work
 - Especially on a Business Unit or critical application basis
- Goal: Controlling user to server traffic
 - Keeping unauthorized users from sending packets to certain servers
 - Governance, Sarbanes Oxley (SoX), HIPAA, PCI

User Segmentation

- Not a data center topic
- We will touch on the topic of user segmentation in passing
- NAC and 802.1x can do this
 - Dynamic role-based VLANs
- Voice / IPT is like another user segment



What's Going On Here?

- “The network is the problem”
- “Server guys don't plan and are disorganized”
- New applications and requirements keep coming...
- Reality: complexity, time challenging everyone
- **NEED NEW APPROACHES, flexibility**



Audience Survey

- How many looking at Data Center Segmentation?
- How many already segmented?
- Reasons?



Case Study: Cisco

Cisco Service-Oriented Data Center

Cisco: The Winner! - Richardson


Richardson, Texas

Runner-ups

- Phoenix
- Boulder


Tipping Points


- Leverage \$21M RDC9 Capital investment
- Accelerate Data Center Business value by 12 months
- Cisco Community and Campus
- Multiple land options at optimal distances at right size
- Fiber Infrastructure
- Skilled IT resources



Cisco: Design Phases


- Consolidate
 - Optimize Data Center Resources
 - Increase Resource Utilization
- Virtualize
 - Virtual Resource Pools
 - Increase Availability and Agility
- Automate
 - Adaptive Orchestration
 - Rapid Delivery of Services



Chesapeake NETCRAFTSMEN  13 Contents of this slide copyright Cisco, used with permission Copyright 2007


Cisco: Data Center Evolution


	Legacy Data Center	Consolidated Data Center	Virtual Data Center	Service Oriented Data Center
Compute	4 Tier Silos Heterogeneous OS	Standardization Virtual Machines	Server Repurposing VM Mobility	Infrastructure Aligned to Application Services
Storage	Storage Silos Low Utilization	SANs, VSANs Tiered Storage	Storage Virtualization	Policy Based Management
Network	IP Connectivity	Consolidated Network Services	Virtualized Network Services	Intelligent Data Management
Security	Perimeter Security	Secure Each Application Tier	Virtual Firewalls	Tiered Recovery Usage and SLA-based Funding Model
Application	Application Silos Distributed	Consolidate, Centralize	Optimization	
	2004	2005	2006 - 2007	2008 - 2010
	Consolidation Phase		Virtualization Phase Automation Phase	

Chesapeake NETCRAFTSMEN  14 Contents of this slide copyright Cisco, used with permission Copyright 2007

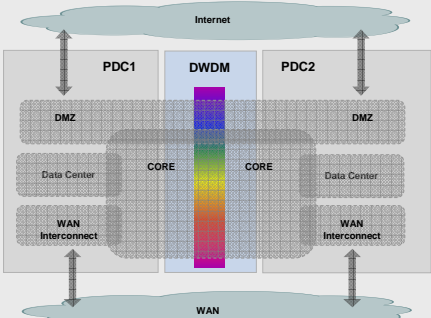
Cisco: SODC Server Virtualization


- Data Center Server Consolidation
 - Improve Operational Agility
 - Lower Data Center Operating Expense
- Increase Utilization of Physical Servers
 - Optimize TCO
 - Improve Data Center Capacity Management
- Reduce Service Provisioning Times
 - Rapid deployment of Operational Environments
- Increase Operational Efficiencies
 - Ease Support of Environments
 - Reduce Planned and Unplanned downtime




Chesapeake NETCRAFTSMEN  15 Contents of this slide copyright Cisco, used with permission Copyright 2007


Cisco: Network Foundation Architecture



Chesapeake NETCRAFTSMEN  16 Contents of this slide copyright Cisco, used with permission Copyright 2007

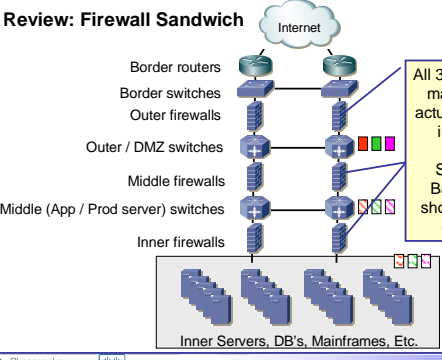
Case Study: High Availability Enterprise



Chesapeake NETCRAFTSMEN  17 Copyright 2007


Traditional: Non-Virtual, Non-Segmented Design

Review: Firewall Sandwich

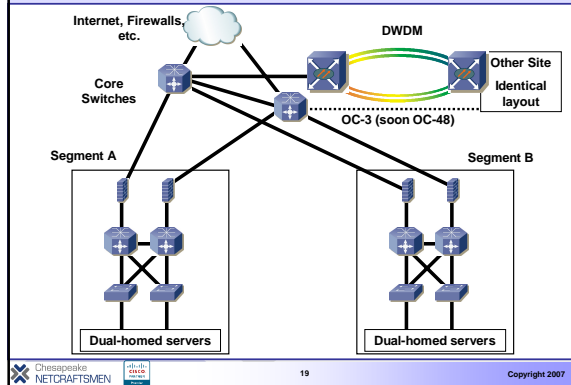


All 3 firewall layers may be virtual, actually one blade in a chassis

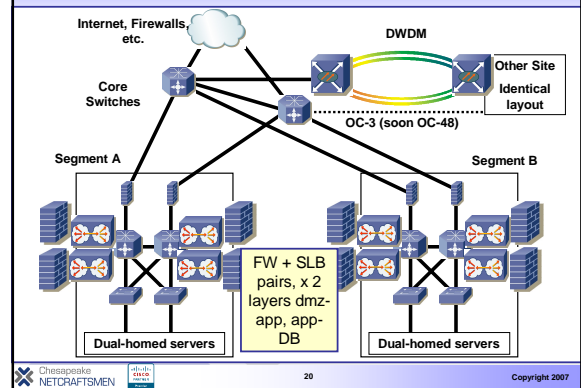
Server Load Balancers not shown to reduce complexity

Chesapeake NETCRAFTSMEN  18 Copyright 2007

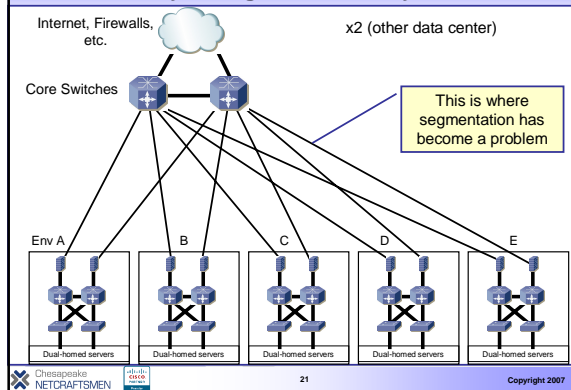
Case Study – High-Level (Simplified) Design



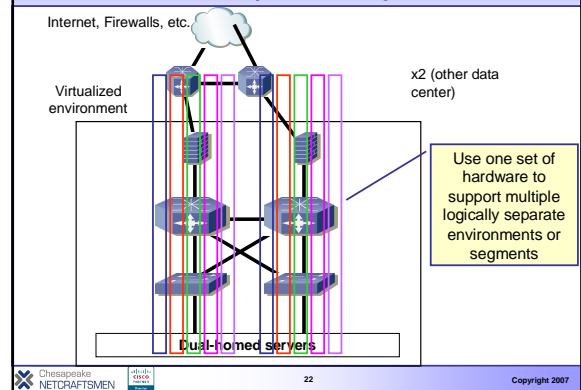
Case Study – High-Level (Un-Simplified) Design



Case Study – Segmentation by Hardware



Case Study – The Objective



Techniques for Segmentation



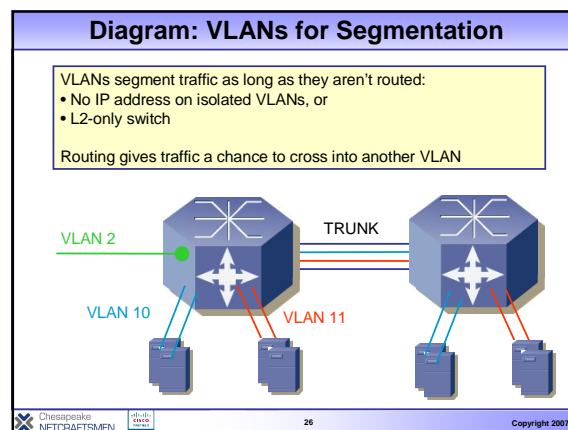
Techniques for Segmentation

Technique	Scale	Complexity
VLAN and PVLAN	Low	Easy
L2: Ethernet over Something ("EoX")	Medium	Medium
VACL's or RACL's in switch	Medium	Medium
ACL's in firewall(s) – routed or transparent	Medium	Medium
Client to Server IPsec solutions (Apani, Microsoft)	Large, but...	Easy
VRF Lite (Multi-VRF)	Medium	Easy
VRF Lite and IPsec or GRE Tunnels	Medium	Medium
MPLS VPN	Large	Harder

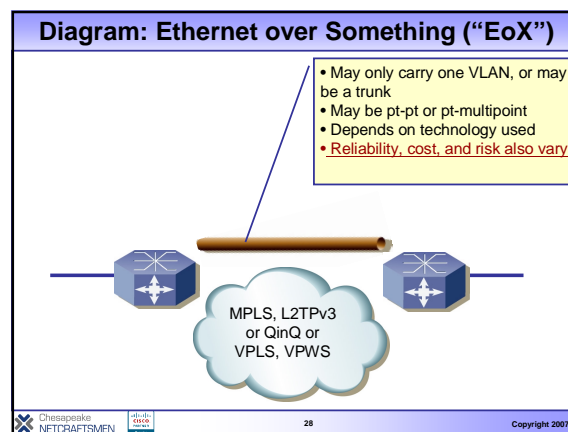
Chesapeake NETCRAFTSMEN

Copyright 2007

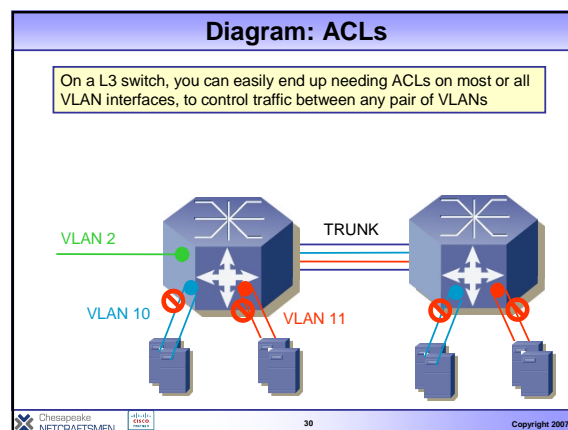
Technique: VLAN	
<ul style="list-style-type: none"> Techniques in this category: <ul style="list-style-type: none"> VLAN PVLAN QinQ 	
Pro	Con
<ul style="list-style-type: none"> Simple Accommodates VMotion 	<ul style="list-style-type: none"> Small scale only STP diameter grows too large STP risks (failures tend to affect entire VLAN or STP domain) Extended VLANs or QinQ require hop-by-hop verification



Technique: Ethernet Over Something ("EoX")	
<ul style="list-style-type: none"> Ethernet over: <ul style="list-style-type: none"> MPLS L2TPv3 VPWS or VPLS 	
Pro	Con
<ul style="list-style-type: none"> Ethernet over MPLS or L2TPv3 somewhat localizes VLAN impact on core / distribution 	<ul style="list-style-type: none"> Risks of over-subscription, statistical muxing STP loop = high traffic still, impact? More complex to troubleshoot VP*S: protection against other customers' problems? Danger of creating "virtual cabling tangle"



Technique: ACLs (Switch or Firewall)	
<ul style="list-style-type: none"> VACL at L2 or RACL at L3 in switch ACLs in firewall(s) / FWSM(s) 	
Pro	Con
<ul style="list-style-type: none"> Need ACLs to control what enters an environment / segment anyway 	<ul style="list-style-type: none"> ACLs at any L3 switch in data center hard to manage But without ACLs, any L3 hop is a chance to be routed out any interface Works ok for core+dist limited L3, not if larger amount of L3



Technique: Transparent Mode FWSM

- FWSM or FW in bridging mode
- Can be used to “split” a VLAN, isolate one group of servers from another, without re-addressing


Pro

- Helps when re-addressing servers is a non-starter*
- *Can often be detected by loud sounds when suggested*
- Handy to have in bag of design tricks

Con

- Doesn't scale
- Messy
- Doesn't virtualize anything

Diagram: Transparent Mode Firewall



Technique: Client-Server IPsec

- Apani and Microsoft selling the idea of “zones” with client-server IPsec encryption
- Reduces all access to server authentication / login and role assignments

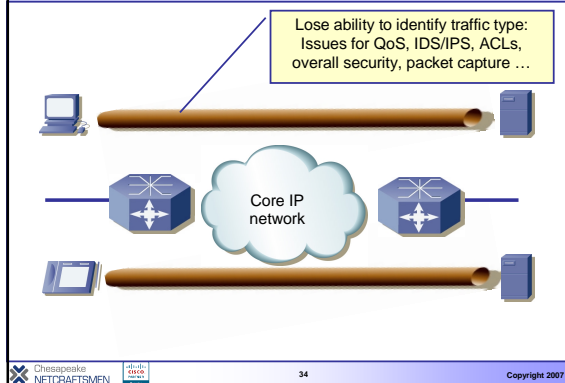
Pro

- Simple to deploy from server-side perspective

Con

- Defeats use of NAM, IPS/IDS, Sniffer
- Makes troubleshooting a lot harder
- Possible MTU issues or performance impact?
- Sales doesn't mention: encryption burden on server (add how many more?)

Diagram: Client-Server IPsec or GRE



Technique: VRF Lite

- Cisco devices allow “VRF Lite” or Multi-VRF
- Use of VRF virtual routing table without MPLS, MBGP
- Create virtual routing tables that interconnect VLANs
 - Think of VLANs as pipes, VRF as the plumbing connector that ties them together at L3
- Acts like “Layer 3 VLANs”

Pro

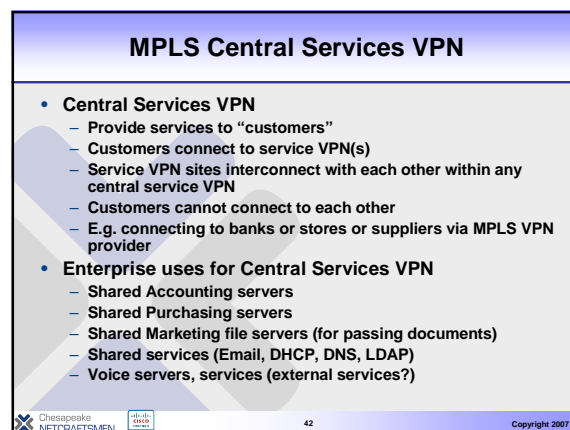
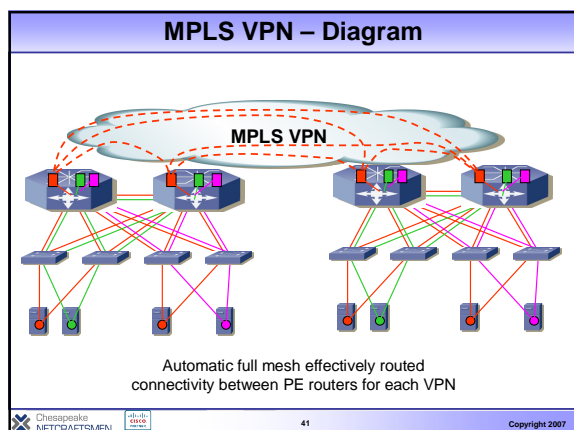
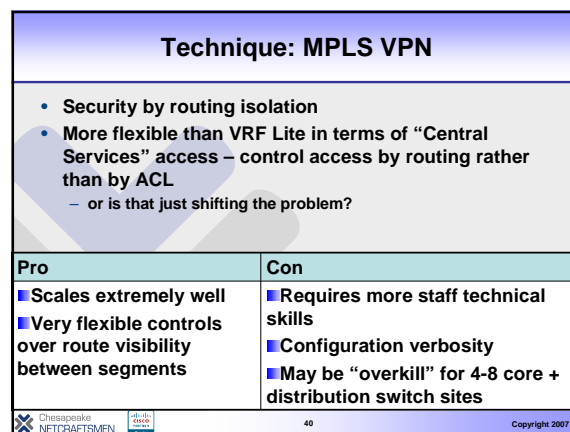

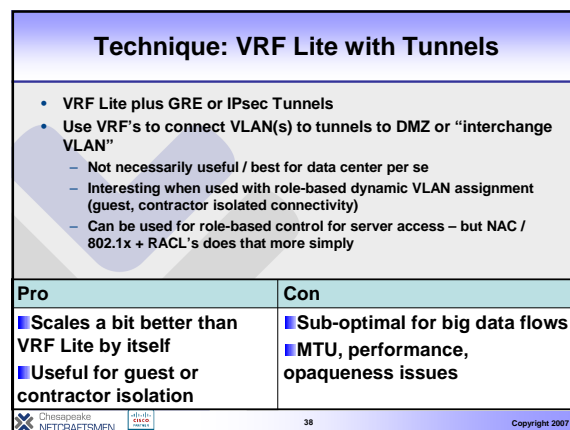
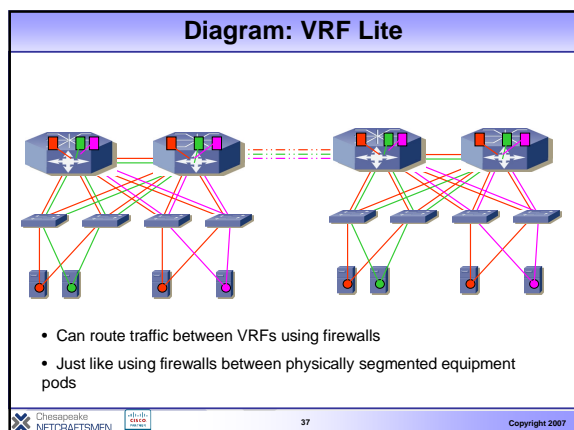
- Breaks up STP domains
- Provides benefits of L3

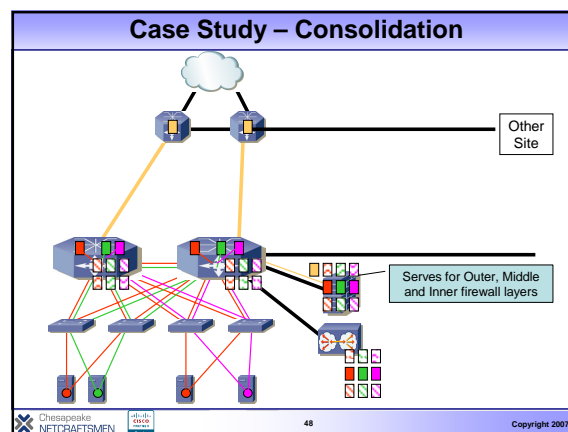
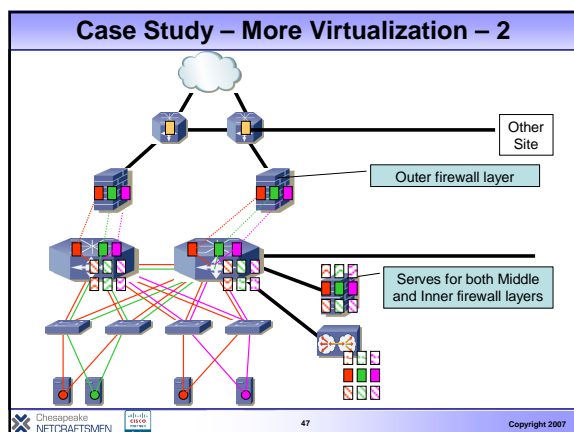
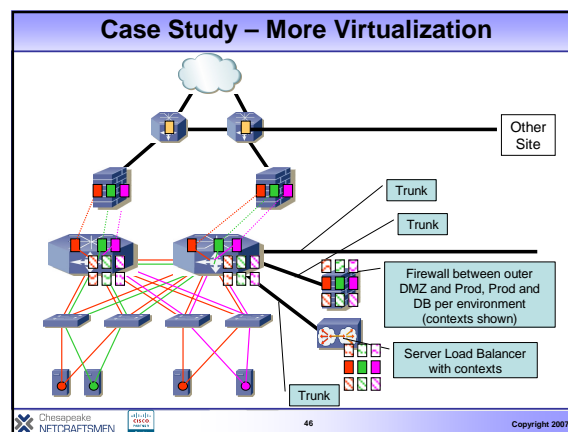
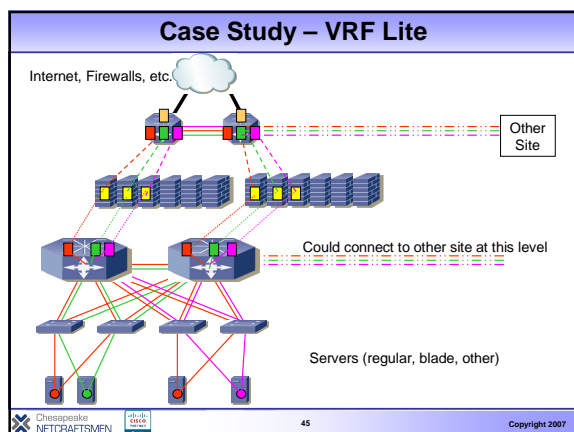
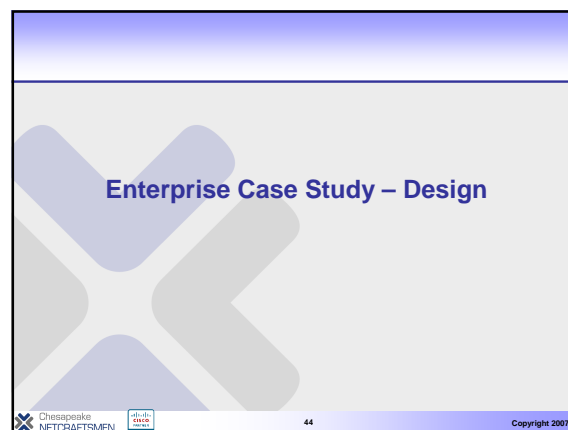
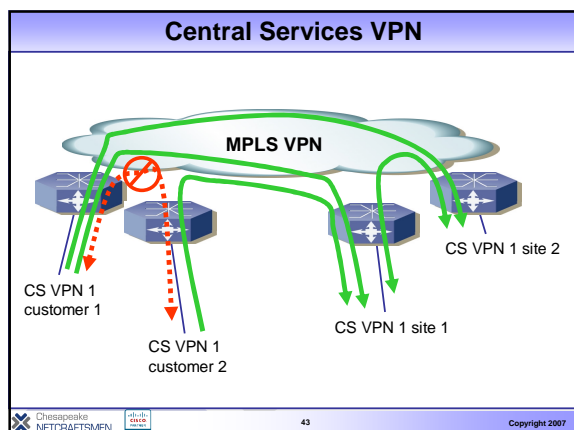
Con

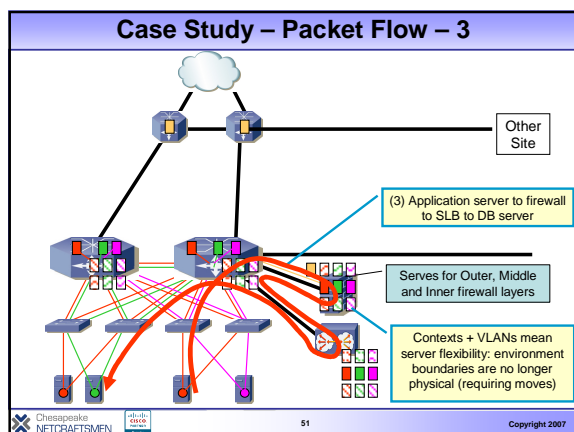
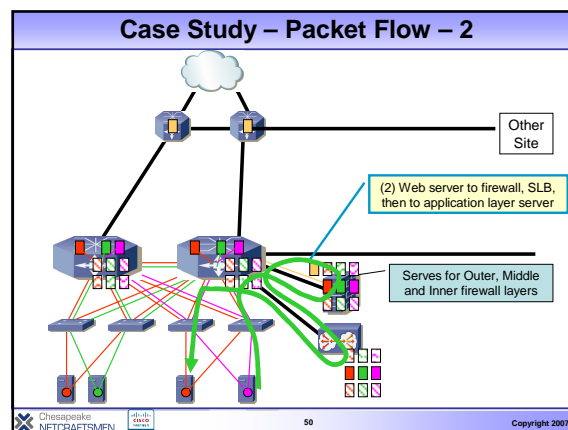
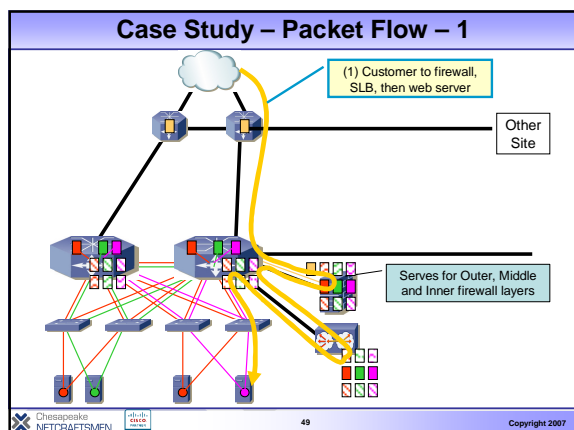
- “Plumbing” – has to be configured and verified hop by hop
- Logical topology usually inefficiently compared to physical topology
- E.g. have to route to core firewall to get between two servers on same switch in different segments

VRF's Are Virtual Routing Tables

- VRFs are virtual routing tables, almost like virtual routers
- VRFs can be connected by interfaces
 - Physical interfaces
 - Logical interfaces
- Must use different interfaces to keep the VRFs separate
 - None or one VRF is assigned to each interface
 - None = typical “global” routing



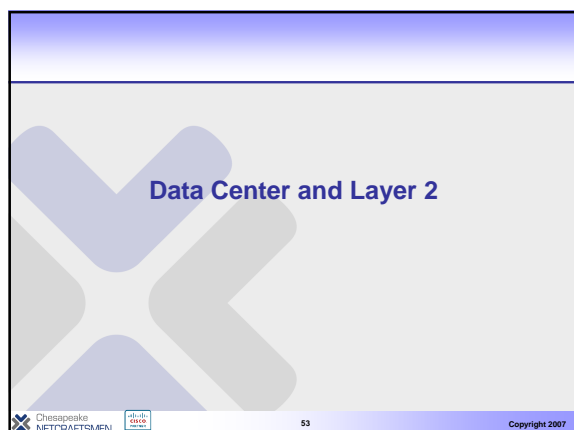




Case Study – Followup

- Think about bringing up a new virtual server for an existing virtual environment...
 - Plumb VLAN out to access switch(es) or blade server
 - Not necessary if already trunking
 - (Trust boundary discussion we'll skip here)
 - Tie VLAN to VRF if not already done
 - Attach new virtual server
 - Assign or re-address per new VLAN**
- The point: VLANs and VRF become “virtual patch panel”
 - Servers don't move
 - Logical server moves do require re-addressing
 - One big flat VLAN everywhere so don't ever have to re-address?
 - **NOT a good idea!**

Chesapeake NETCRAFTSMEN 52 Copyright 2007



Data Center Layer 2: Background

- Many of us have seen the consequences of major L2 Spanning Tree Protocol (STP) issues
 - Can cause large-scale outages (STP domain size)
 - Tough to troubleshoot
- We and certain major sites try to limit use of L2 / STP, L3 to access layer
- There is however a tension in design: Layer 2 has advantages from the server team perspective...
 - You need to find the right balance for your site
 - Layer 9 issue (“political layer”)

Chesapeake NETCRAFTSMEN 54 Copyright 2007

Data Center Layer 2 Challenges: VMotion

- VMotion requires VMotion ports on same VLAN
 - Currently: work is in progress on L3 VMotion
 - Can create pressure for a VLAN that spans rows or the entire Data Center
 - Trades convenience (any v-server anywhere) for risk
 - Don't take on (hidden) risk!

Data Center Layer 2 Challenges: Clusters

- Server Clusters create the temptation to geographically split the cluster
 - GeoCluster: until recently this meant a L3 WAN cluster, specialized DB synch applications, etc.
 - Data Center to Data Center L2 clustering was done using SONET/DWDM (i.e. robust, dedicated links, no statistical muxing of traffic)
 - Recent enabler: MAN Ethernet allows inexpensive high-speed Layer 2
 - You (sometimes) get what you pay for: Best Effort?
 - Question: has the server team done their homework, understands risks, understands cluster behavior with packet loss or intermittent conditions, does vendor provide WAN (routed) alternatives, why not use them, etc.?
 - Result: now have VLANs extending between two sites, wider STP risk, odd traffic patterns, may have to manually intervene for optimal failover
- Conclusion: may be OK to do it, by no means a “no-brainer”

L2 and Segmentation

- Another approach: Ethernet over Something (“EoX”)
 - Leaves underlying infrastructure Layer 3
 - Still carries VLANs between rows of servers or sites
 - Does it mitigate risk?
- Easy to create EoX “spaghetti”, hard to maintain
 - It is about the same as pulling fiber between rows of servers to directly connect dedicated for servers in a VMotion pool – ad hoc, unstructured
- My answer to date:
 - There's a reason we call it “bleeding edge”
 - Do you want to be the first to learn about new technology with your most critical apps and servers? (Especially clusters)

Controlling User Access to the Data Center

Controlling User Access to the Data Center

- Governance in some cases means that only selected admins and users can send packets to crucial servers
 - Most exploits are internal
 - Login controls may no longer be “enough”
 - Goal: prevent the average user from probing financial or credit card servers, or confidential HIPAA data servers, for vulnerabilities
- Such roles typically administered via LDAP or Active Directory groups of users

Controlling User Traffic to Data Center

- Two approaches (at least):
 - In-band Cisco NAC Appliance between users and Data Center or selected servers
 - Allows per-role ACLs
 - Caution: Watch your scaling!
 - Out of Band Cisco NAC Appliance, or 802.1x / IBNS
 - Authentication identifies group role
 - Role-based dynamic VLAN assignment
 - Switch ACL controls accesses on a per-source VLAN subnet basis
 - Scales much better
- Could also do VACLs on the VLANs sensitive servers are on
 - Maintaining VACLs all over the data center could get ugly
- Conclusion: something between users and Data Center has to block dis-allowed packets
 - Your choice of what device you want the ACL's on
 - Your choice of which approach fits your needs better

More Segmentation and Virtualization for the Data Center

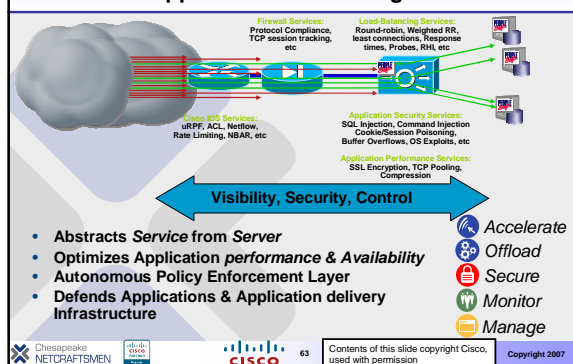
Chesapeake NETCRAFTSMEN 61 Copyright 2007

More Segmentation and Virtualization for the Data Center

- What about virtual firewalls?
 - Cisco PIX, ASA, FWSM can now do contexts
 - Some limitations, but no barrier to current typical uses
 - FWSM nominally 5 Gbps throughput
- What about Server Load Balancers?
 - Cisco ACE does contexts (and ACLs, and HTML verification...)
 - ACE nominally 15 Gbps throughput
- What about SAN?
 - VSAN, VSAN routing, zones, etc.
 - No discussion here due to time/space constraints
- Evolving fast!
- **Reducing hardware per environment has green consequences!**

Chesapeake NETCRAFTSMEN 62 Copyright 2007

Business Continuity Application Control Engine



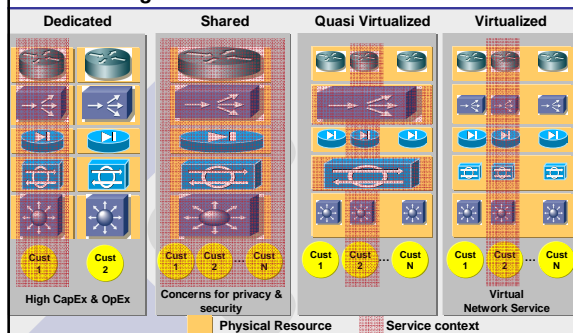
Virtualized Services

Services: Firewall, Load Balancing, SSL Encryption/Decryption

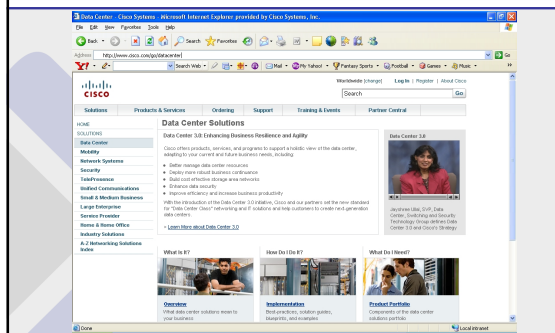
- L4-L7 services integrated in Cisco Catalyst® 6500
- Server load balancing, firewall and SSL services may be deployed in:
 - Active-standby pairs (CSM, FWSM 2.X)
 - Active-active pairs (ACE, FWSM 3.1)
- Integrated blades optimize rack space, cabling, mgmt, providing flexibility and economies of scale
- Influences many aspects of overall design

Chesapeake NETCRAFTSMEN 64 Copyright 2007

Service Integration and Virtualization Evolving towards Virtual Network Services



www.cisco.com/go/datacenter



Cisco's Vision for Virtualization (and Segmentation)

- Cisco VFrame
- Automated control over deployment and configuration of data center infrastructure
 - Network connectivity
 - Firewalls
 - Server Load Balancers
 - SAN: virtual fabric
 - Virtual machines
- <http://www.cisco.com/go/vframe>


Chesapeake NETCRAFTSMEN 67 Copyright 2007

Summary


- Virtualization in the form of VRF's, with or without MPLS, can help consolidate equipment in the data center
 - Main server farm
 - DMZ and e-commerce complexes
- Virtualization can also segment to provide better logical and security separation between environments (Prod, Dev, Test, etc.)
 - Un-tangling environments reduces complexity, chance of mistakes
 - Potentially mistakes only knock out one environment
- Virtualization of Firewalls and Server Load Balancers enhances the benefits
- Thanks for coming!

Chesapeake NETCRAFTSMEN 68 Copyright 2007

Any Questions?



- For a copy of the presentation, email me at pjw@netcraftsmen.net
- References: see web article I will post at <http://www.netcraftsmen.net/welcher/papers/index.htm>
- About Chesapeake Netcraftsmen:
 - Cisco Premier Partner
 - Cisco Customer Satisfaction Excellence rating ★
 - Highly certified technical experts
 - We wrote the original version of the Express Foundations courses required for VAR Premier Partner status (and took and passed the tests)
 - Cisco Advanced Specializations:
 - Advanced Unified Communications (and IP Telephony)
 - Advanced Wireless
 - Advanced Security
 - We have deep expertise in Routing and Switching (several R&S CCIE's)
 - We do network / security / unified communications Design and Assessment
 - Expertise and experience in many other areas as well



Chesapeake NETCRAFTSMEN 69 Copyright 2007