



DATA CENTER NETWORKING

OVERVIEW, ARCHITECTURE AND DESIGN GUIDELINES

Mikkel Brodersen
Systems Engineer
Cisco Systems, Denmark

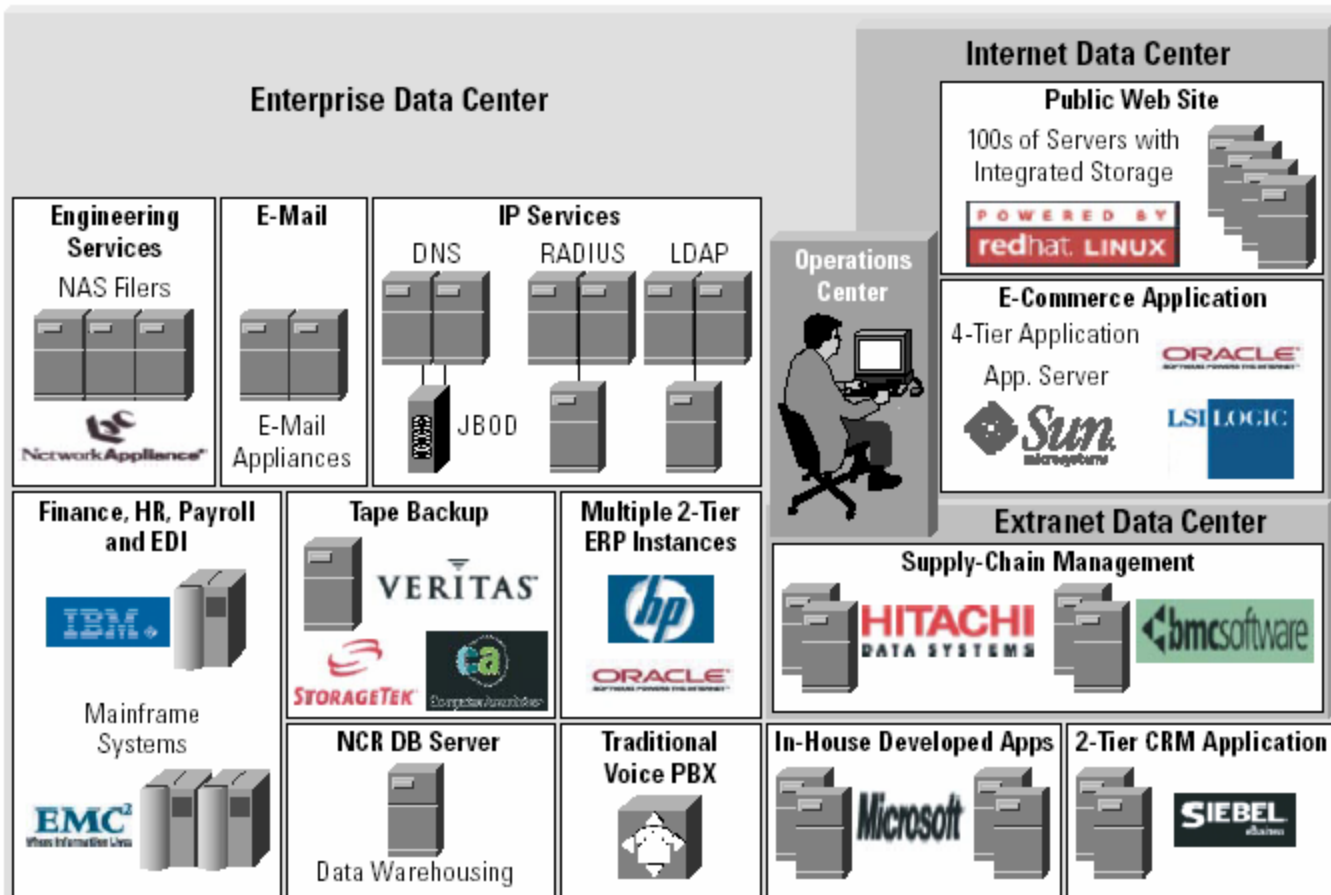
Agenda

Data Center Infrastructure

- **Data Center Introduction**
- **Server Farm Architecture Overview**
- **Design Requirements in the Server Farm**
- **Access Layer Design Models**
- **Density and Scalability Implications**
- **Scaling B/W with Gigabit EtherChannel[®] and 10GE**
- **Spanning Tree Design and Scalability**
- **Summary**

Enterprise Data Centers Today: The Reality—Disparate Silos, Major Inefficiencies

Cisco.com

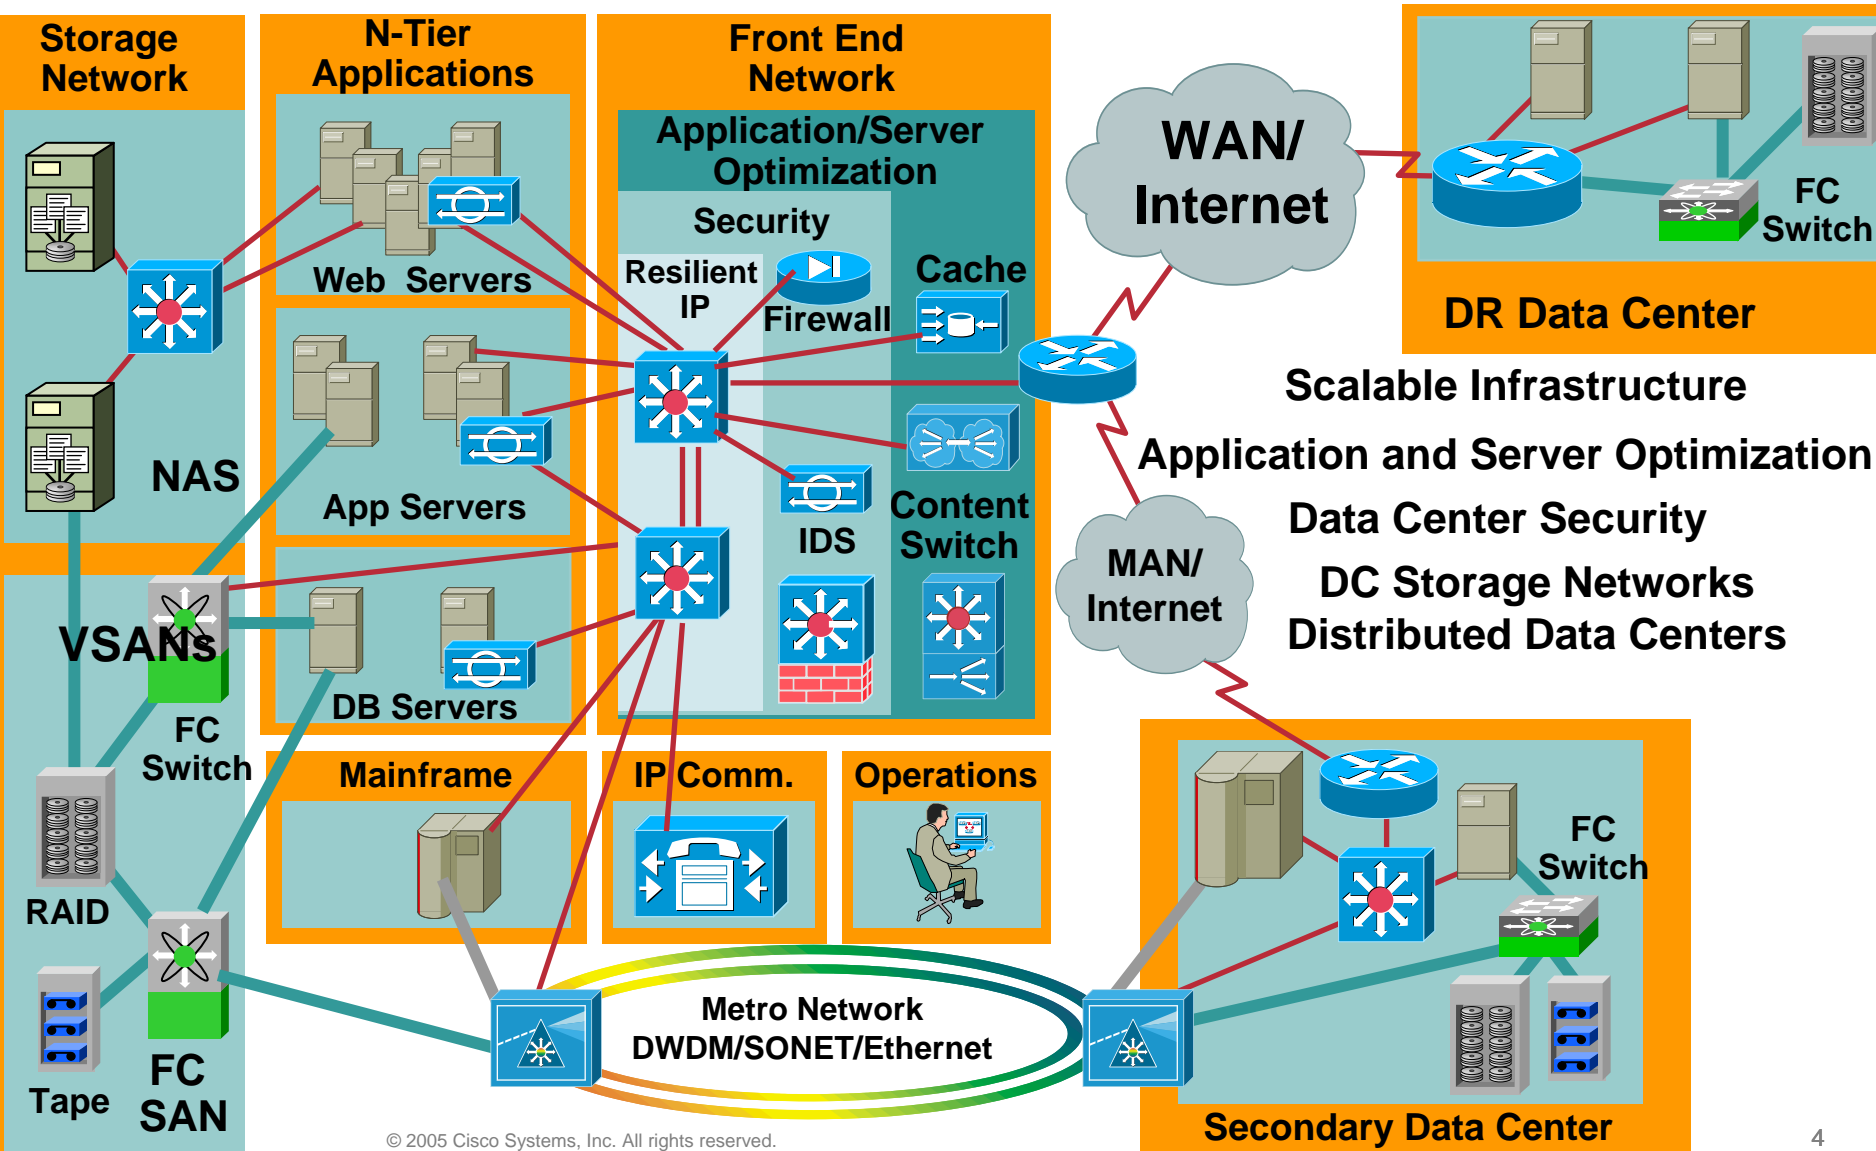


- Enterprise Applications and Services
- Communications Applications
- Computing Infrastructure
- Storage Infrastructure
- Networking Infrastructure
- Facilities (Power, Cooling, Cabling, Physical Security)

Today's Data Center

Integration of Many Systems and Services

Cisco.com




Scalable Infrastructure
 Application and Server Optimization
 Data Center Security
 DC Storage Networks
 Distributed Data Centers

Systems and Solutions

Tactical Execution

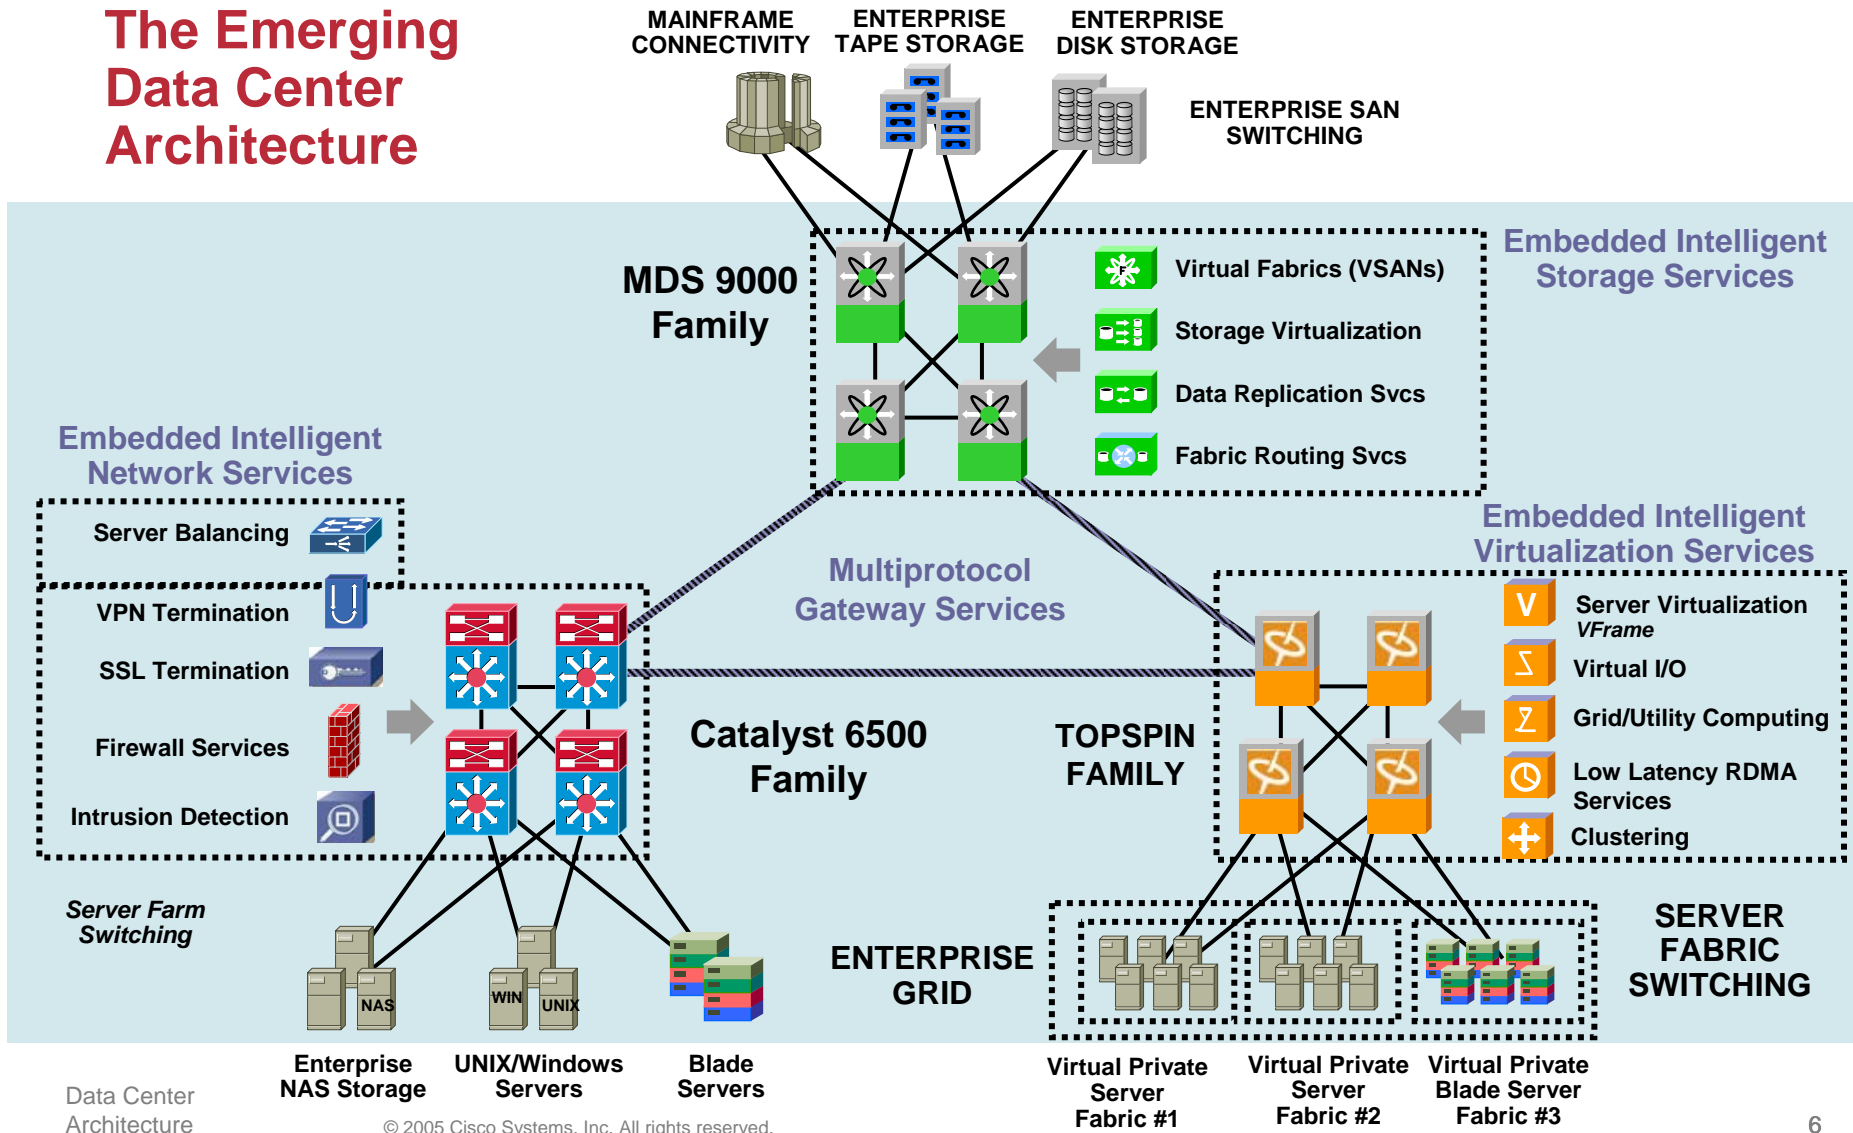
Cisco.com



IP Network Infrastructure	Application Optimization	DC Security	Storage Network Infrastructure	Business Continuance Networking
Routing	Content Switching	Traffic Filtering	FC Switching	DC Interconnectivity
Switches Routers IGP and BGP Protocols	Server Load Balancing	Firewalls ACL—RACLs, VACLs	Director Class Switches Stackable Switches	DWDM, SONET/SDH, CWDM GE and 10GE MetroE and IP WAN Services L2 and L3 VPN Service
Switching	SSL Offload	Intrusion Detection/Protection	iSCSI/FCIP	Site Selection
Modular Stackable RPVST+	SSL Acceleration Mgmt Simplification Monitoring Encrypted Traffic	Network IDS Host IDS	FC to IP Ethernet Gateways IP Services in FC switches	Failover and Load Balancing DNS Base Site Selection Route Health Injection IGP and BGP Site Selection
Server Farm Topologies	Caching	L3 and L2 Features	SAN Topologies	Data Replication and SAN Extension
NIC Teaming Clustering iSCSI, FC, NAS,	Reverse Proxy Caching WCCP and SLB Redirection Content Prepositioning	PVLANS, Static ARP, Port Security, MD5 Authentication AAA, SSH, Root and BPDU Guard, ARP Spoofing, DHCP Spoofing, VLAN Hopping	Best Practices Intra-DC Inter-DC	Synchronous and Asynchronous FC Over Campus and MAN FCIP Over WAN

The Big Picture—The Cisco Data Center

The Emerging Data Center Architecture




Data Center Architecture

Intranet Core
Internet Edge

DC Aggregation

Server Access



IP Infrastructure

Highly Available, Adaptable, Predictable, Deterministic and Service Ready

Cisco.com

- **Integration with the routed network**

 - Intranet and Internet peering

 - DC Isolation from external events

- **Server farm topologies**

 - Layer 2 Adjacency requirements

 - Layer 3 Boundary

 - Service Location

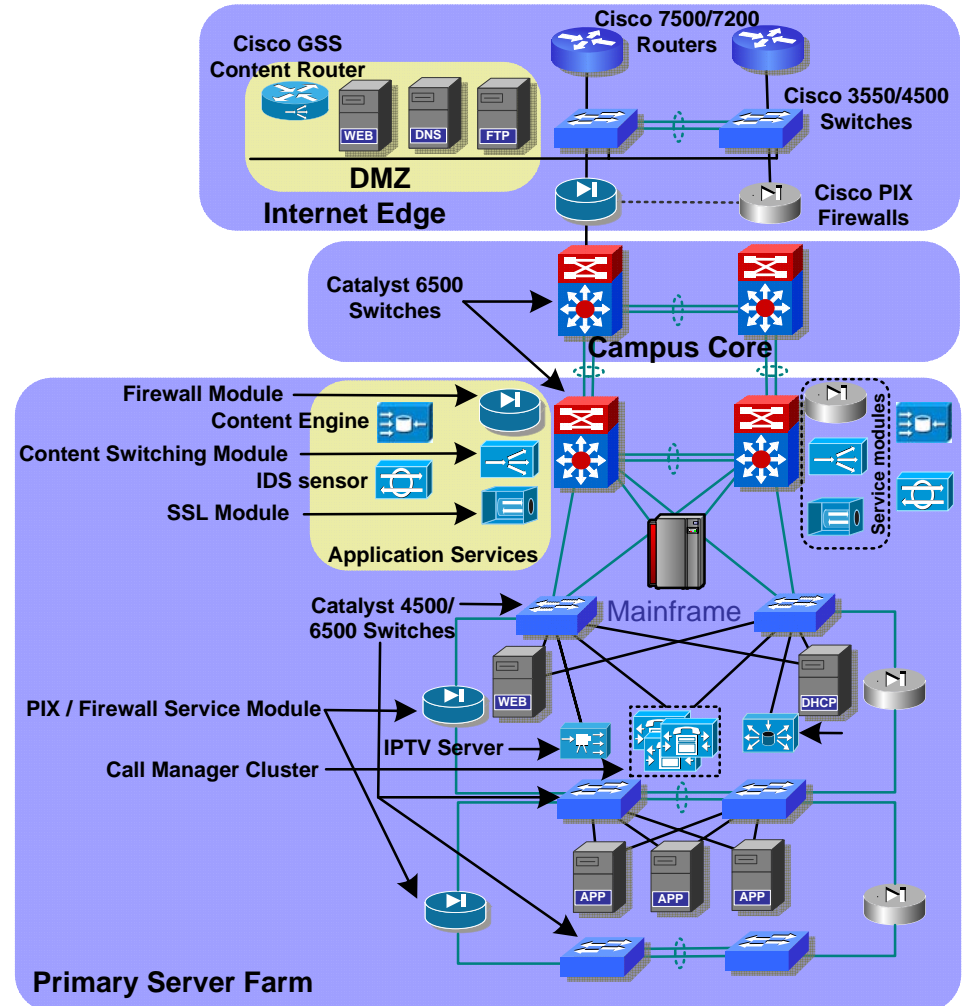
 - Multi-tier Topologies

 - Scalability

- **1RU and Blade Servers Integration**


 - Design Alternatives

 - Service Integration




Multi-tier Topologies

Physical vs. logical Topology



Logical Topology

- Used in Hosting Services
- Dedicated service devices
- Highly scalable



Physical Topology

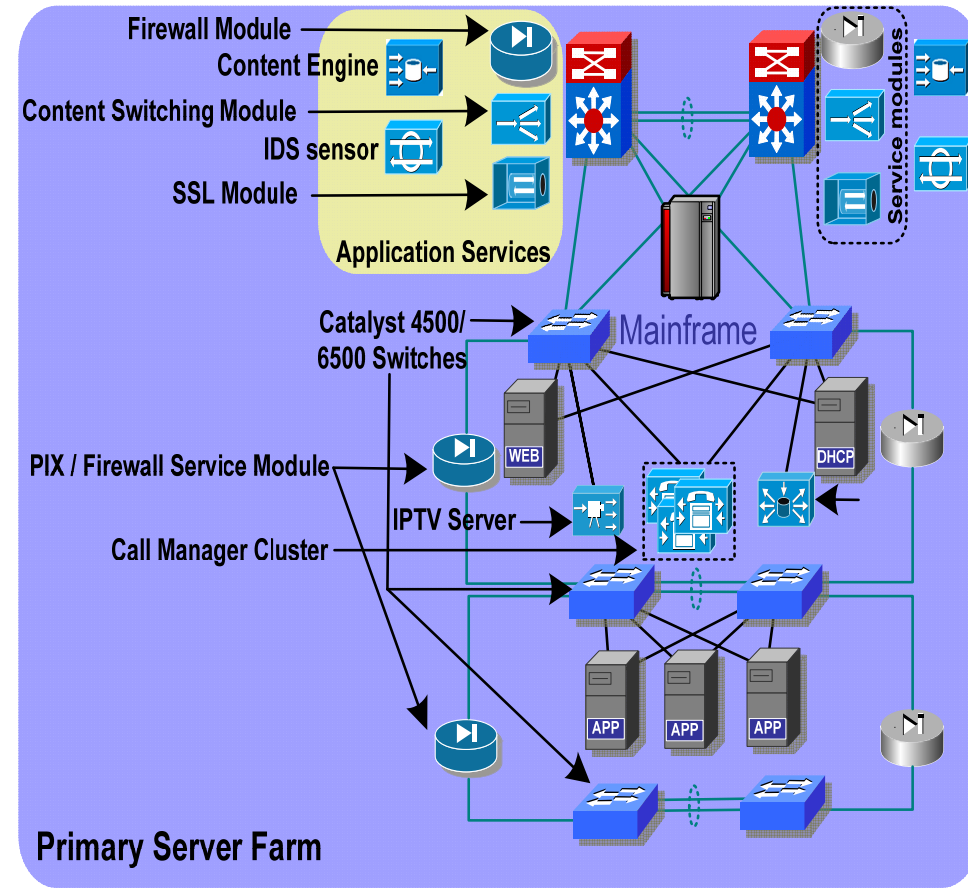
- Service devices are shared (transparent and virtual)
- Greater service efficiency
- Easier manageability
- More cost effective

Application Optimization

High Availability, Scalability and Health of Application Services

Cisco.com

- **Load Balancing**
 - Application Availability
 - Scalability
- **SSL Offloading**
 - SSL Scalability
 - Centralized PKI Management
 - Security/Traffic Inspection
- **Caching**
 - Server I/O and CPU reductions

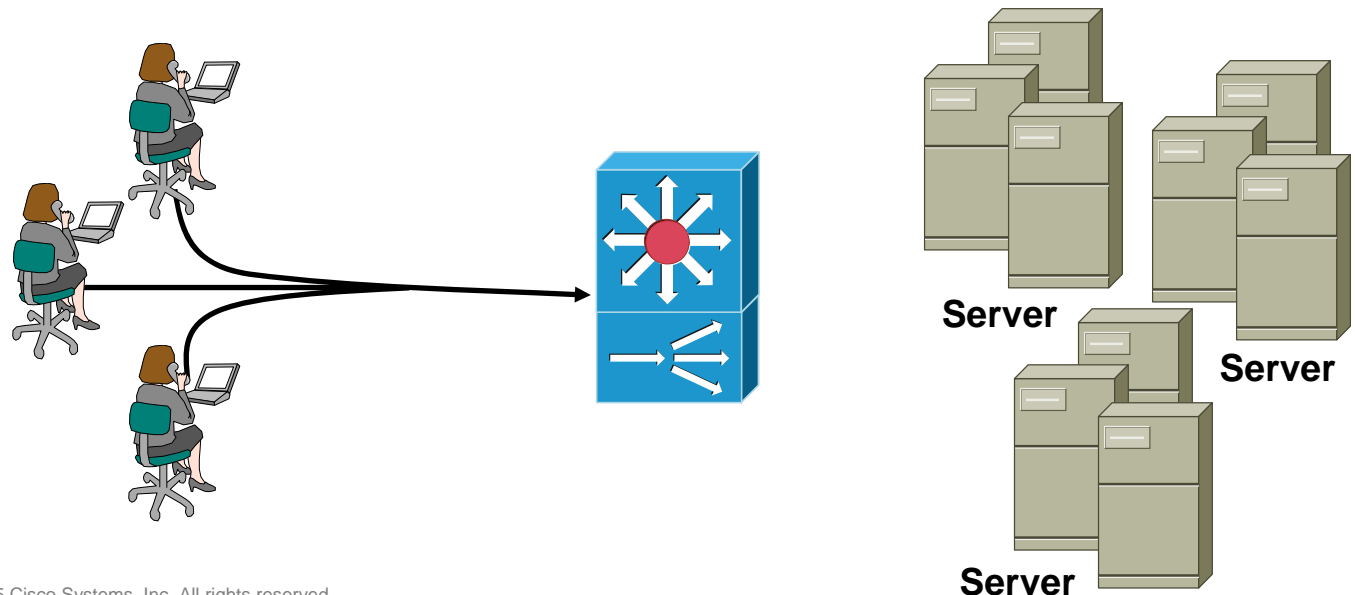


Server and Application Scalability

Improving and Guaranteeing Service Levels

Load balancing and Content Switching Technology

- **Distribute Traffic Load**
- **HW alternative to clustering technologies**
- **Avoiding misbehaving apps/server: app health checking**
- **Allows seamless scalability**
- **Enables any-window maintenance change control**



SSL Offloading

Scaling Application Layer Security

Hardware Alternative to Software or per Server SSL Transactions

With SSL Offloader

- Improves server scalability
- Enables L5+ load balancing and user session persistence via L4-7 switch integration
- Simplifies certificate management
- Provides traffic inspection visibility hidden in SSL sessions

Without SSL Offloader

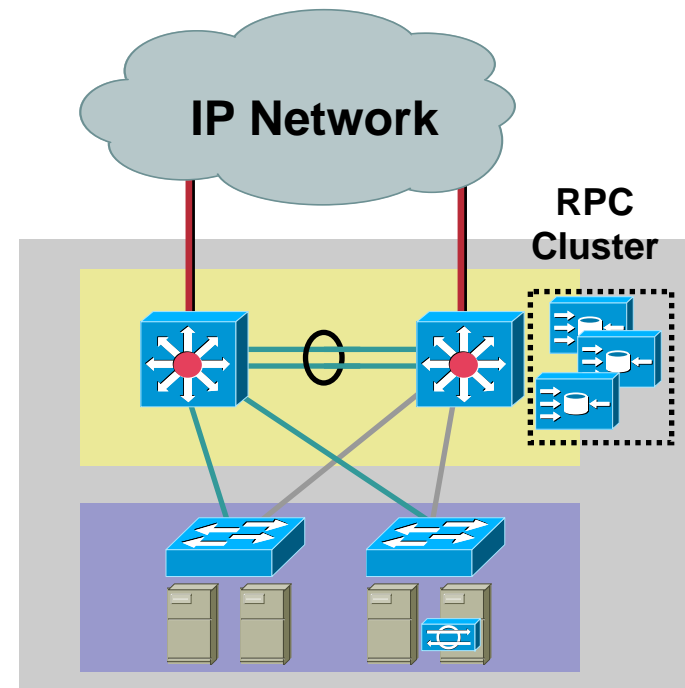
- SSL processing within servers requires high-end servers to scale SSL capacity
- Hides L5+ info for intelligent load balancing, user persistence, and security
- Managing SSL certificates on individual servers is a heavy operational burden

Caching in the Data Center

Offloading Static Content Serving

Caching Technology

- **Offload processing of requests for static content**
- **Reduces server farm CPU cycles**
 - Increases Scalability
- **Done close to the server farm**
 - Reverse Proxy Caching mode
 - Aggregation layer
- **Offloading Redirection Alternatives**
 - Web Cache Control Protocol (WCCP)
 - Content Switch



DATA CENTER SERVER FARM SWITCHING ARCHITECTURE OVERVIEW




Defining the DC Access Layer

Layer 2, Layer 3 Server and Mainframe Connectivity

Cisco.com

- L2 and L3 requirements
- Dual and single attached
- High performance, low latency L2 switching
- Mix of over-subscription requirements
- Many uplink options
- STP processing for configured VLANs only
- Utilizes services in the aggregation layer

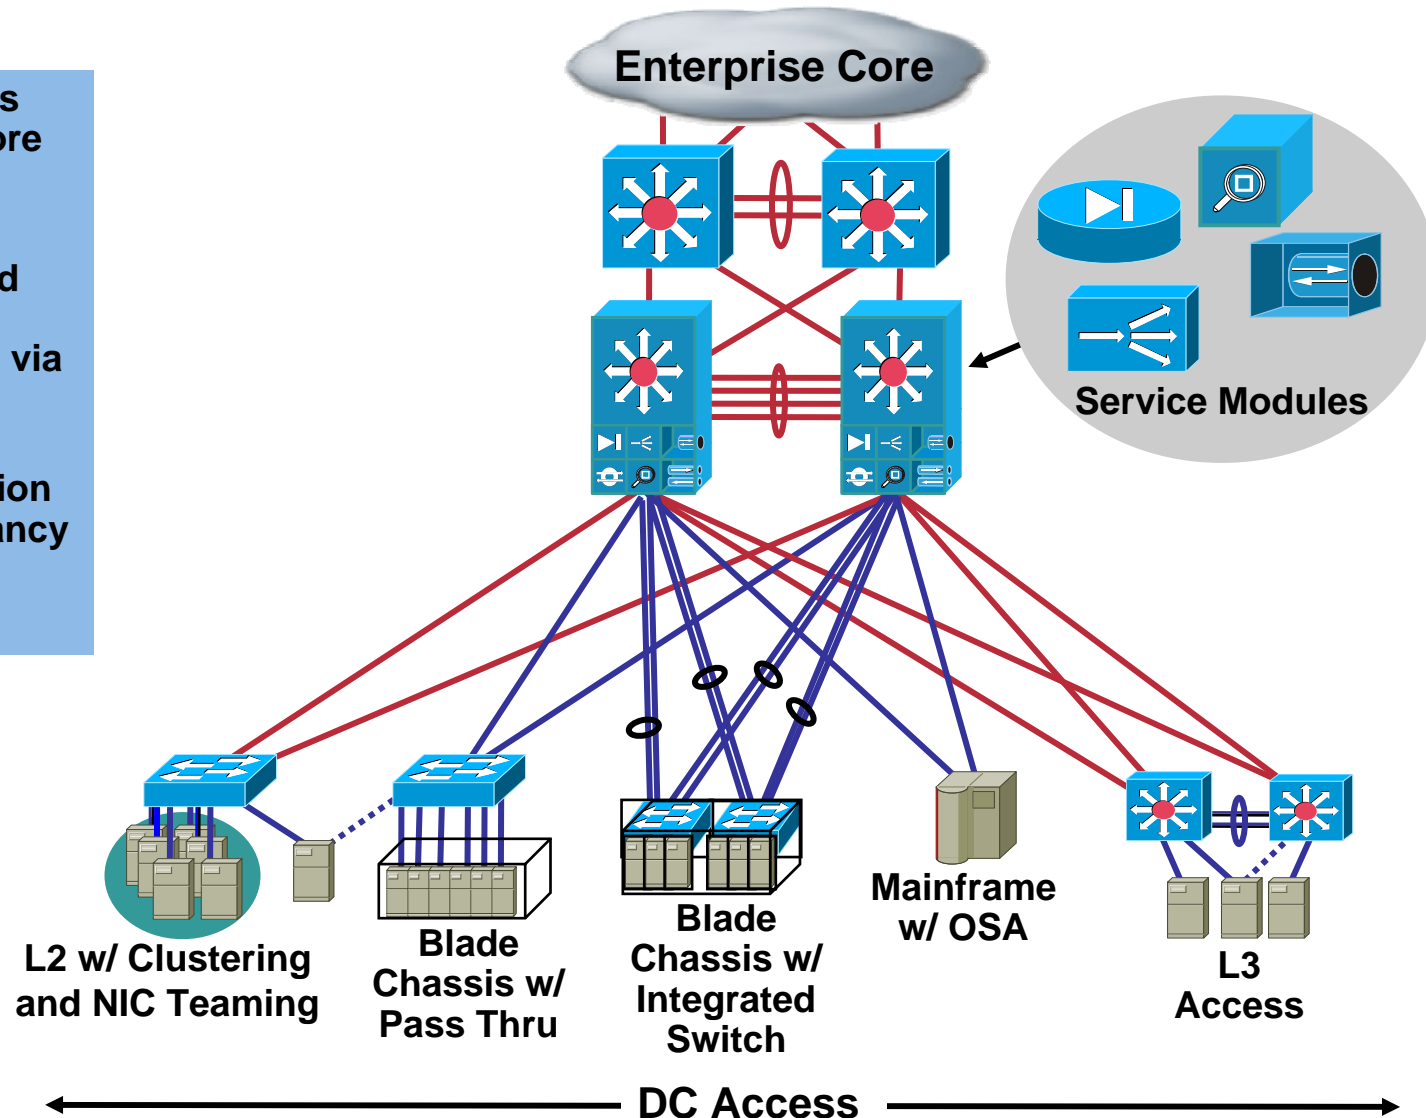


Defining the DC Aggregation Layer

Providing Shared Application/Security Services


Cisco.com

- Aggregates access uplinks into DC core
- Large STP processing load
- Provides advanced application and security functions via service modules
- Maintains session state and connection tables for redundancy
- What are these **services**?



Defining the DC Aggregation Layer Integrated Services

Layer 4–7 Services: FW, SLB, SSL, IDS




- **Application and security services can be deployed as:**
 - Appliances
 - Service blades
- **Service blades such as firewall blades and load balancing blades...provide hardware-based stateful functions**
- **Integrated blades optimize rack space, cabling and configuration mgmt**
- **Provide highest flexibility and economies of scale**

Defining the DC Aggregation Layer Server to Server Communications Path

What Types of Server to Server Traffic Will Exist?

Multi-Tier Interaction, Backup, Replication, Cluster Messaging, Storage over IP




- The aggregation module may provide the primary communication path for server to server traffic
- Non traditional traffic emerging
- Driving lower oversubscription and 10GE uplinks
- Servers now ship with PCI-X NIC's and GE
- **Plan bandwidth for future server true capacity**

Defining the DC Core Layer

High Speed Switching Fabric for Aggregation Modules

Cisco.com



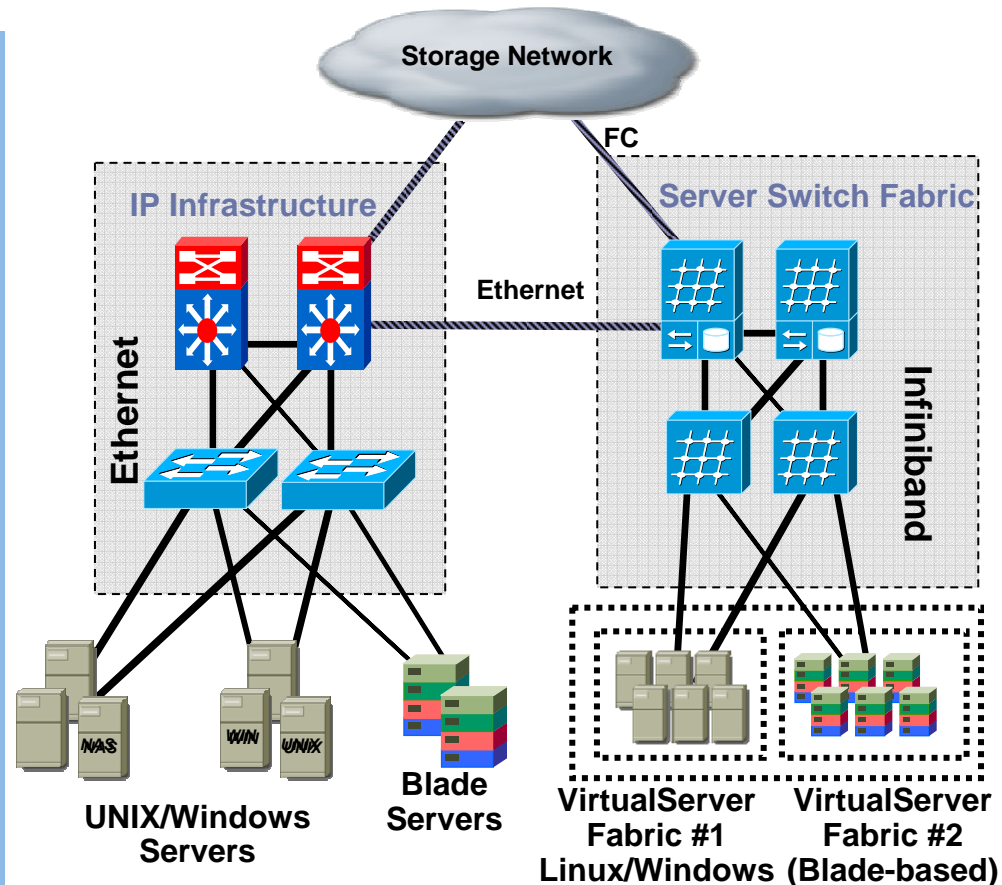
- Interconnects AGG modules
- Isolates failure domains
- Scales large STP diameters
- Improves 10GE scaling
- **Plan and build DC core up front**

Defining the DC Server Switch Fabric

Server Switch Fabric for Inter-Server Traffic

Cisco.com

- Purpose built server switching fabric enabling:
 - Low latency RDMA
 - Server virtualization
 - GRID/Utility computing
- Clustering environments
 - Database clustering
 - HA clustering
 - HPC clustering
- Gateway to IP switching and storage layers
- New, leading edge, still maturing



Enterprise GRID

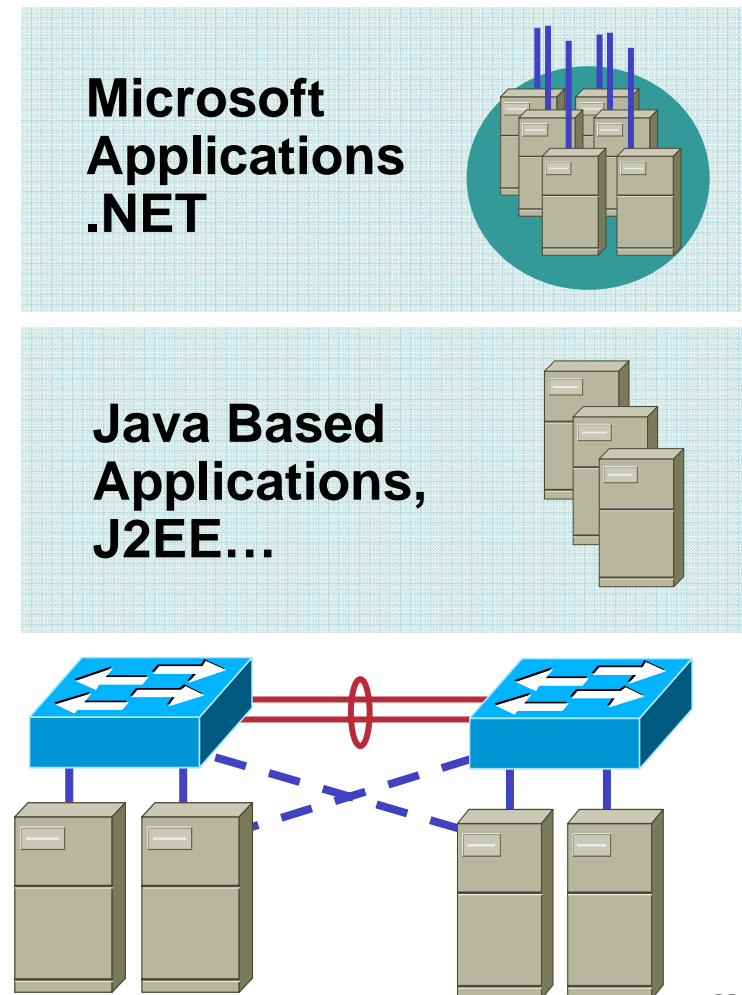
DESIGN REQUIREMENTS IN THE SERVER FARM



When Is Layer 2 Adjacency Required? Meeting Server Farm Application Requirements

Cisco.com

- **Clustering**: applications often execute on multiple servers clustered to appear as a single device; common for HA and load balancing requirements; (Windows MSCS and NLB)
- **NIC teaming** software requires layer 2 adjacency between teamed NICs



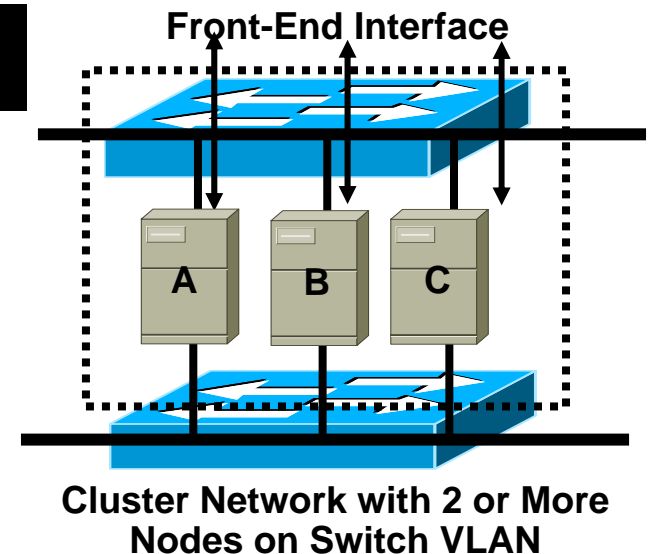

Defining Layer 2 Adjacency

“Layer 2 adjacency between servers means that the servers are in the same broadcast domain. When servers are Layer 2 adjacent, each server receives all broadcasts and multicast packets from another server.”

**Packet Magazine: Second Quarter 2005
Designing the Data Center Access Layer**

Defining Clustered Servers Clusters for High Availability

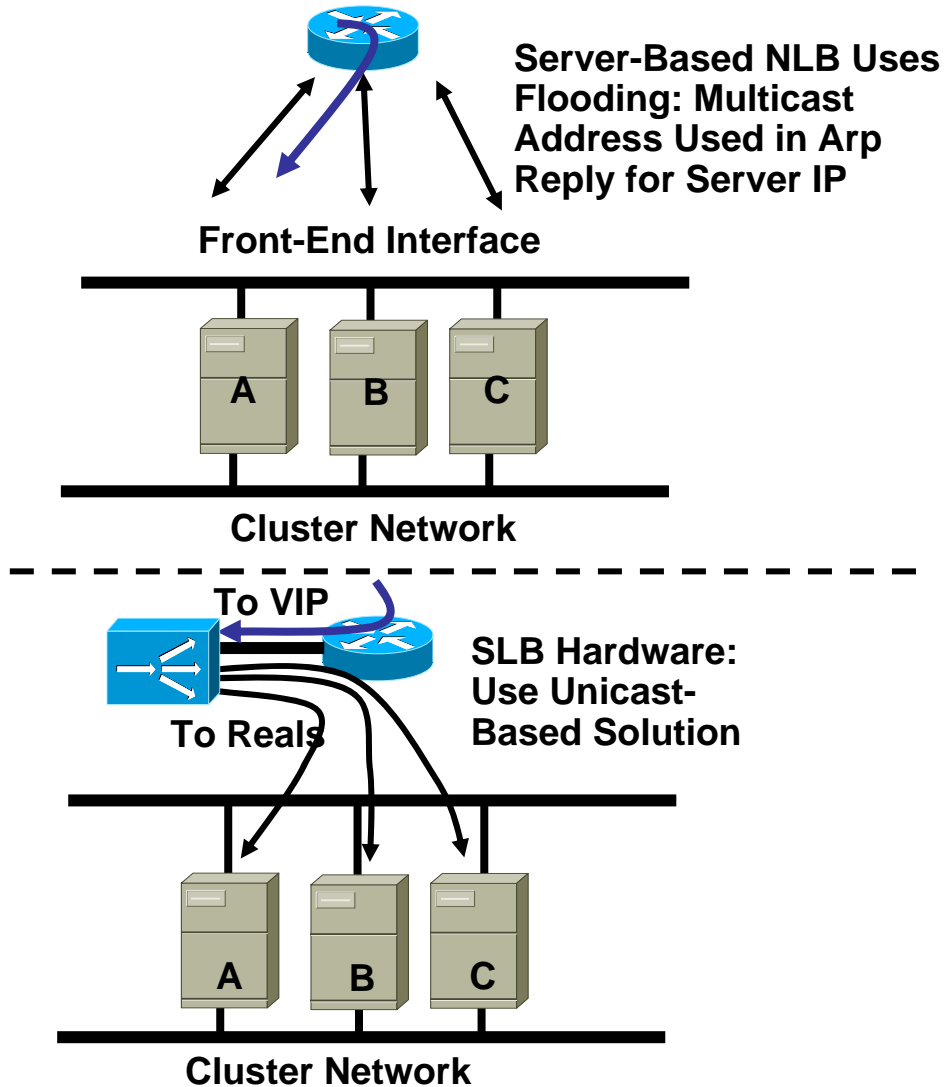
- Common goal: combine multiple servers to appear as a unified system through special s/w and network interconnects
- A 2 Node HA cluster can use a dedicated crossover cable for exchange of data, session state, monitoring...
- Two or more servers use a switch to provide the interconnect on an isolated layer 2 segment/VLAN
- Examples: MS-Windows 2003 Advanced Server 2003 Cluster Service (MSCS), for Exchange and SQL Servers (up to eight nodes)
- Veritas Clustering for HA
- L2 Adjacency is required



Defining Clustered Servers

Clusters for Load Distribution

- Transactional/HTTP-based applications are “clustered” for scalability purposes (MSCS NLB example, up to 32 nodes)
- Layer 2 segment used to “multicast” all incoming packets to all hosts in cluster (L2 Adjacency is required)
- Single IP address associated with **a multicast MAC address** in the cluster arp reply (Windows)
- **Purpose built load balancers provide a standard hardware based unicast solution supporting hundreds of nodes**



Defining Clustered Servers

Database Clustering

- Objective: improve DB lock times and enable more efficient parallel scans

Database Cluster Examples

Oracle 10g RAC

IBM DB2 Parallel

SQL Server

MySQL Cluster

Storage Approaches

Shared Everything

FS locking

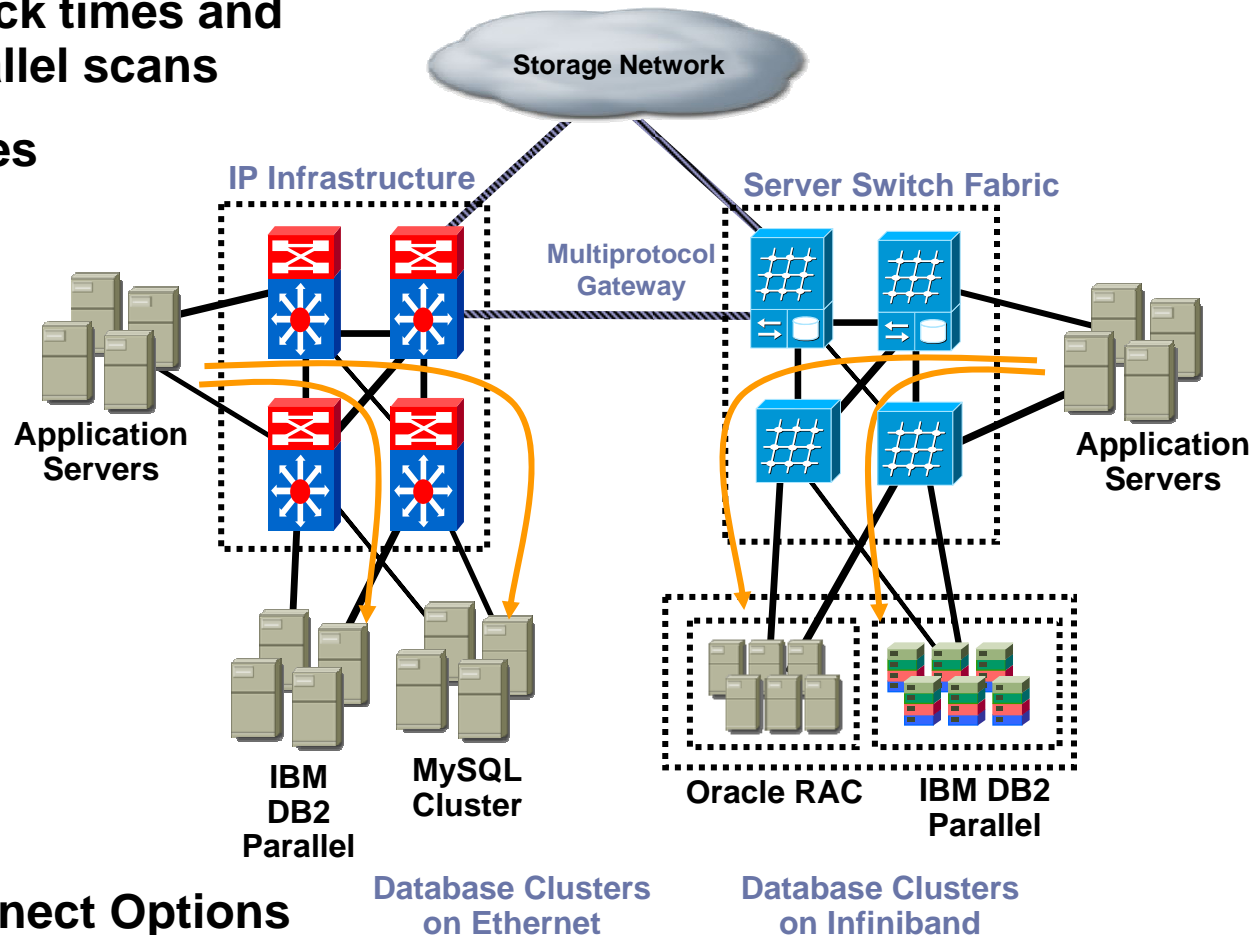
Shared Nothing

Slices up database

Standards based Interconnect Options

Ethernet, Infiniband

Low Latency + High B/W



Defining Clustered Servers


High Performance Computing Clusters

Specific Applications

- Animation rendering
- Seismology
- Oil exploration
- Biochemistry
- Financial analysis
- Common practice to lease HPC environment



Implementations are very customized and rarely alike

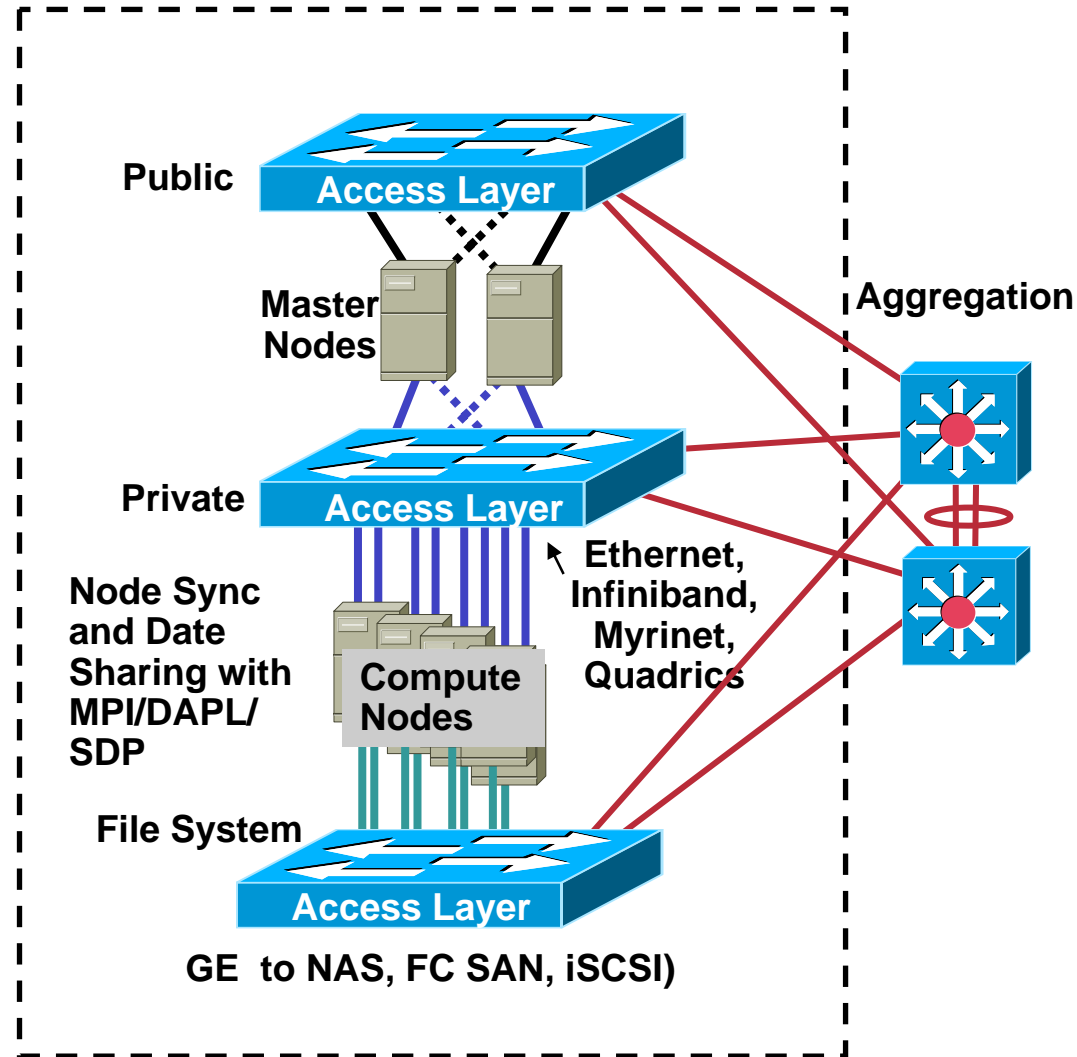


Defining Clustered Servers

Requirements/Implications of HPC Clusters

Cisco.com

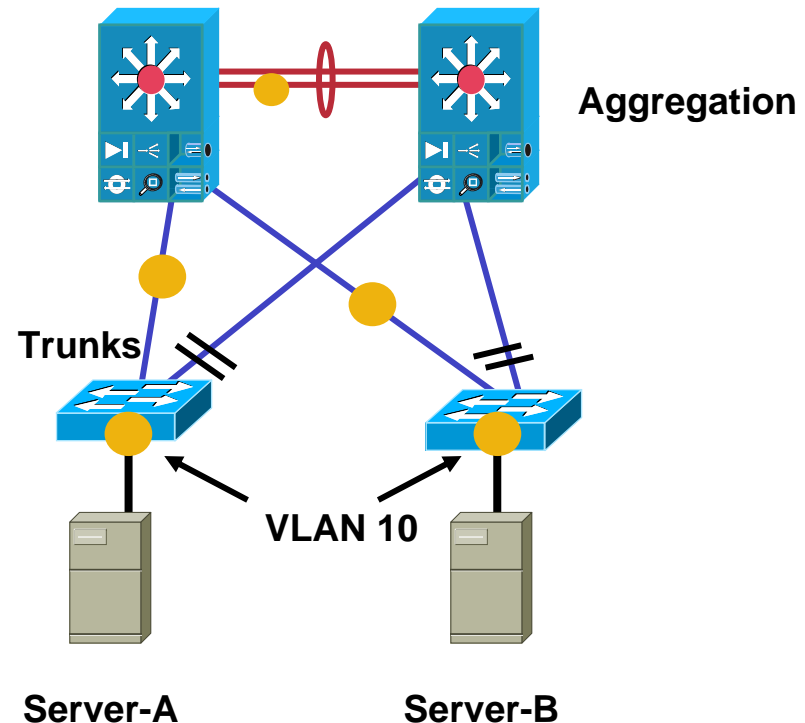
- L2 adjacency between compute nodes may be required (application dependent)
- Common file system used by compute nodes
- Latency is critical to performance
- Network and system staff usually don't communicate clustering needs well
- Who determines which servers to include in a cluster? Same rack? row?
- Will there be an impact on access layer uplinks?



Clustered Servers

Network Design Implications

- **Server-A and server-B communicate at layer 2 to exchange state, session and other information**
- **Servers (2 or more) in cluster may be across different access switches—extending VLANs and Spanning Tree diameter**
- **Server to server cluster fabric may require higher b/w uplinks (GEC, 10GE)**

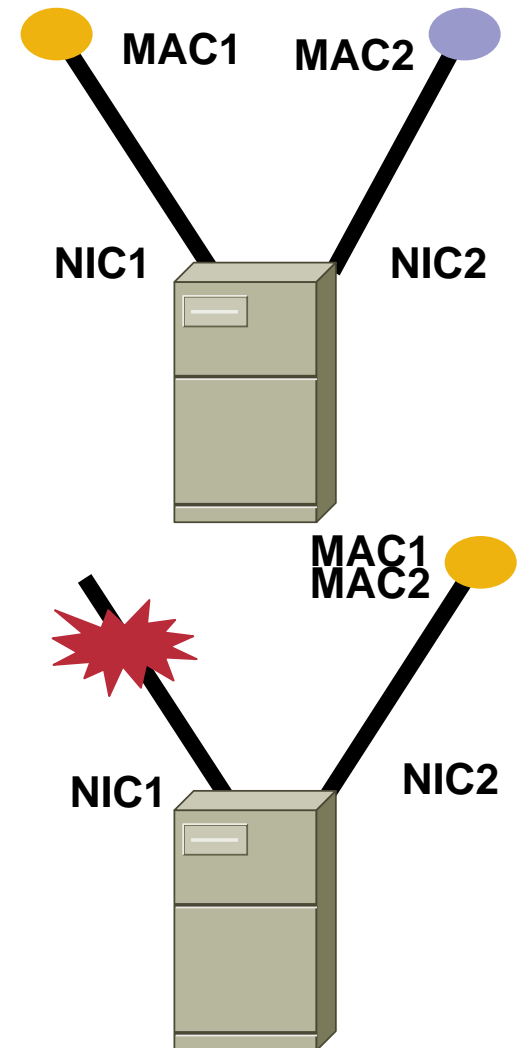


NIC Teaming Requirements

Teaming Requires NICs to Be on the Same Segment

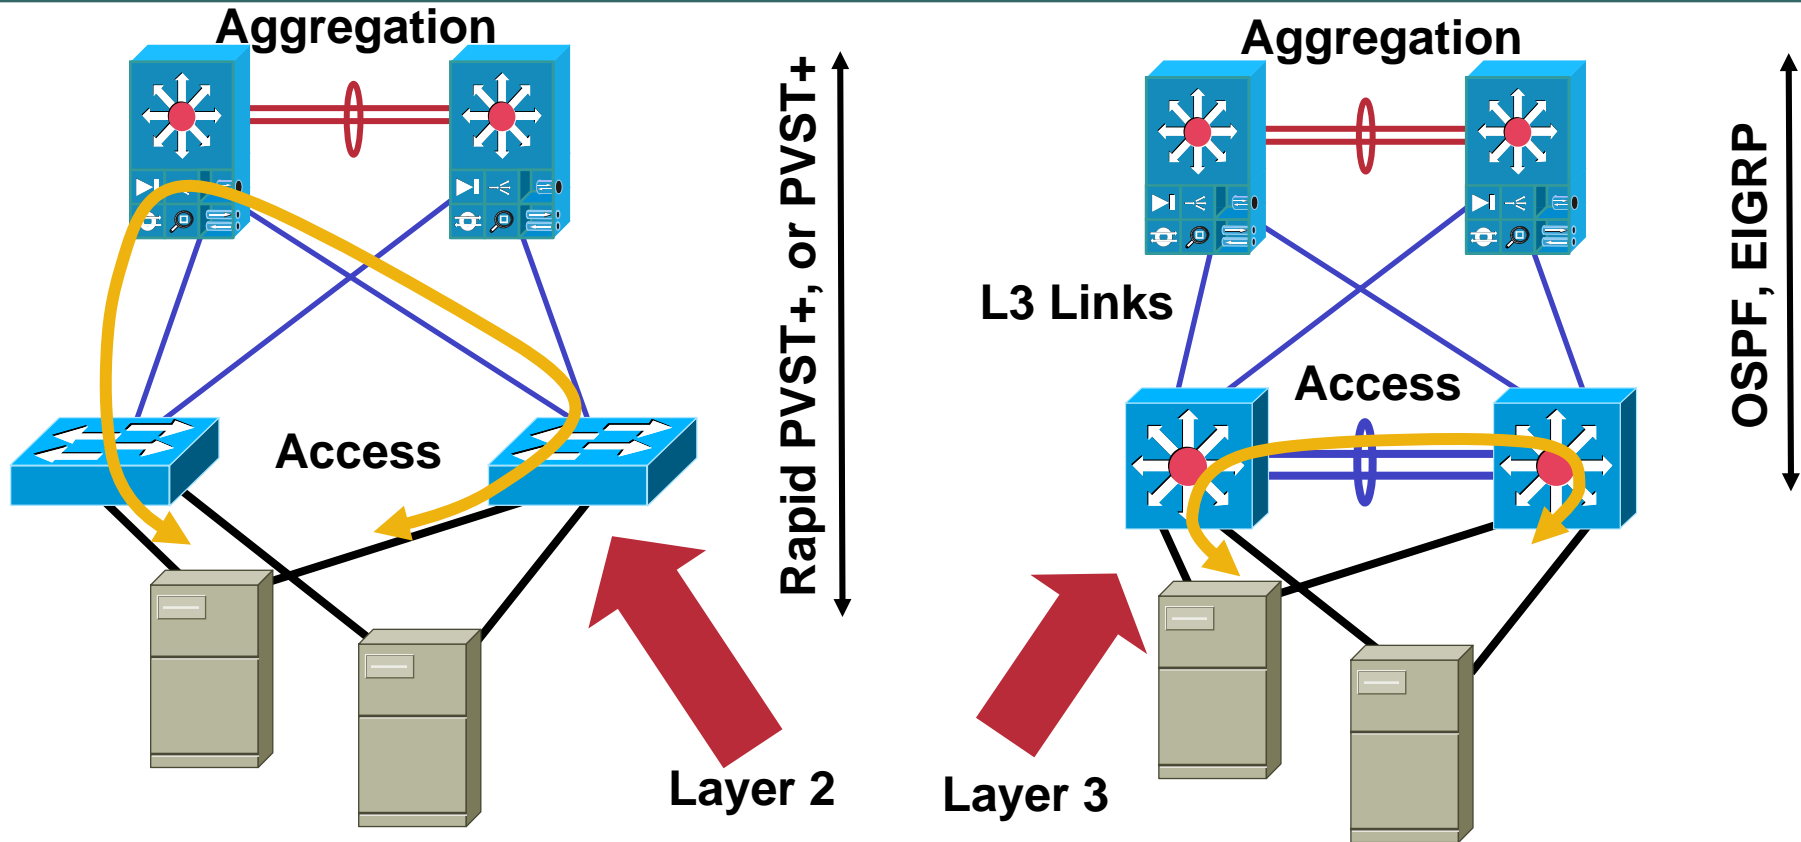
Cisco.com

- Servers are often dual connected for high availability purposes
- The NIC driver bundles multiple NIC cards as if they were a single interface
- If one NIC loses connectivity the redundant NIC becomes active and inherits the same MAC address as the primary one
- The server is always reachable at the same IP address
- This means that both NIC's need to belong to the same BROADCAST domain—same subnet
- Optional probes/heatbeats for monitoring are multicast-based



NIC Teaming Requirements


Integration of Teaming in a Data Center




- The access switches need to provide layer 2 adjacency between the NIC cards of servers with NIC teaming configured; a layer 2 path must exist between such servers

Blade Server Requirements Connectivity Options

Using Pass-Through Modules



Using Integrated Ethernet Switches



Connecting to the Aggregation Layer

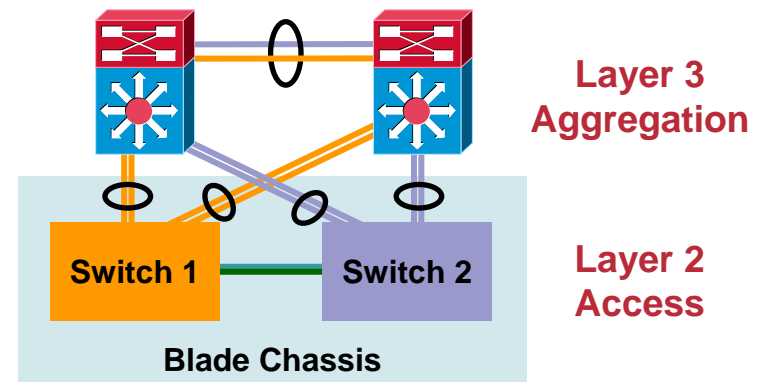
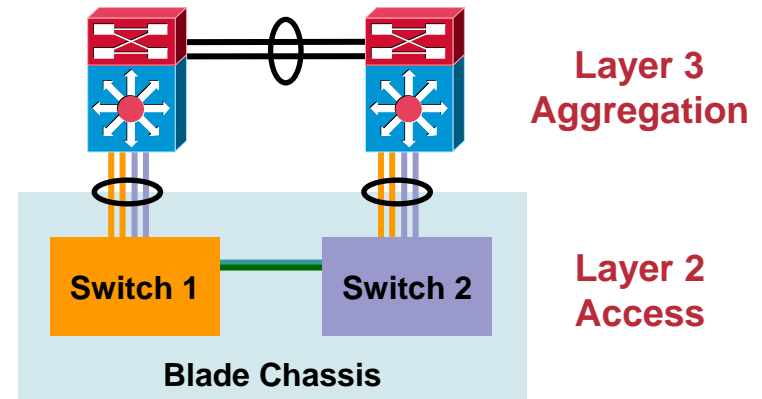
HA Considerations

Uplinks to the Same Aggregation Switch

- **Single STP Loop**
- **Traffic path is not optimal**
 - No direct path to STP root
- **Topology is deterministic**
- **Topology may be loopless**

Uplinks to Different Aggregation Switches

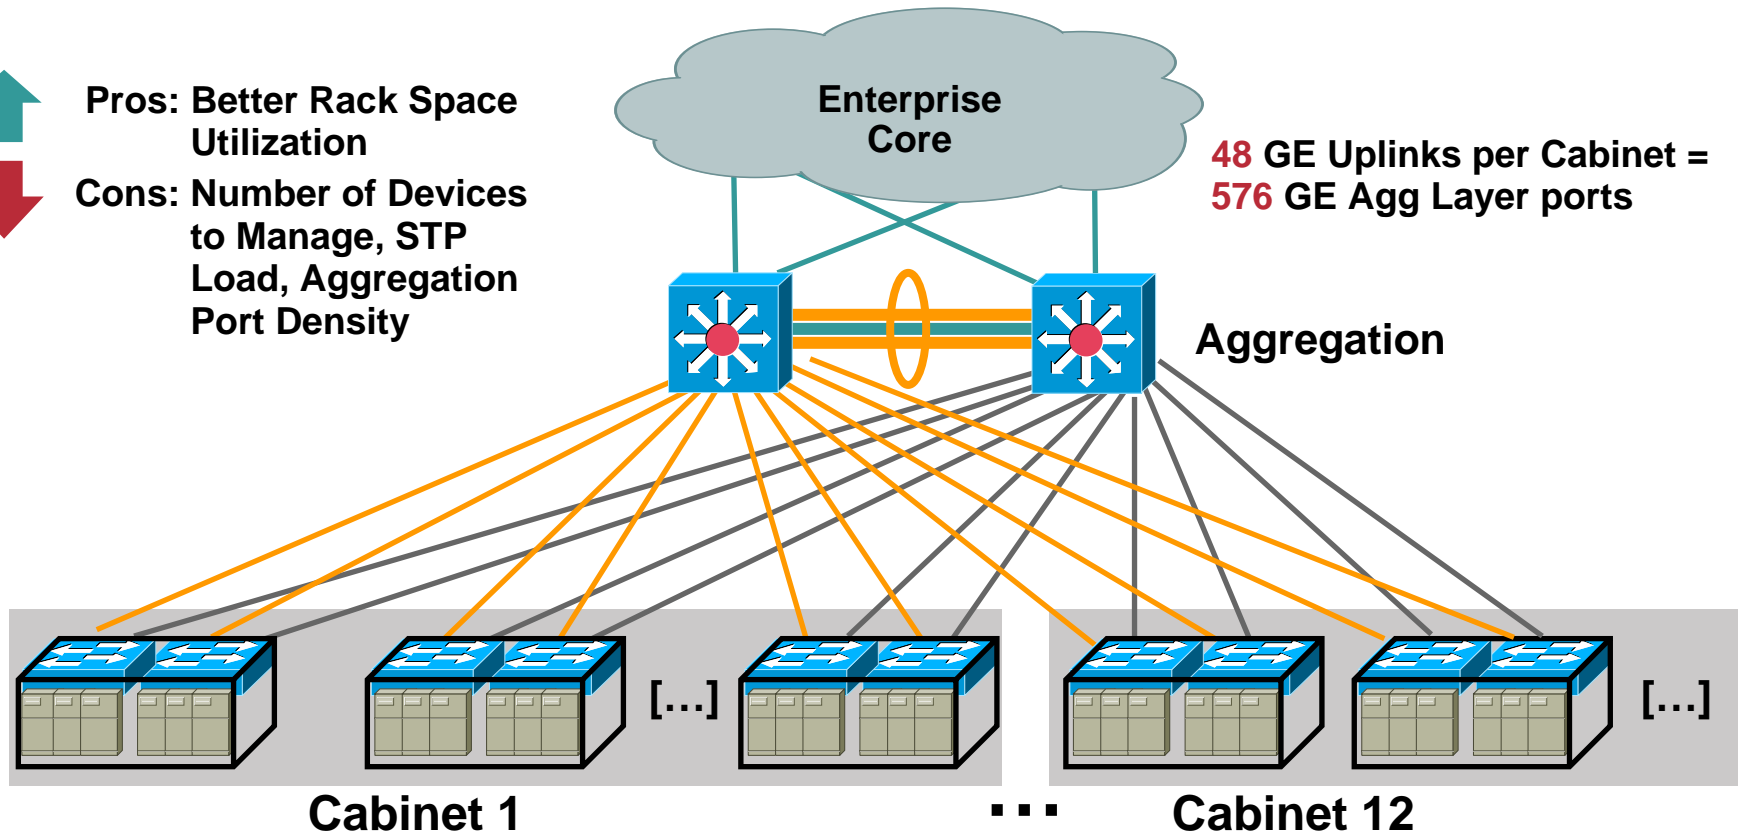
- **A few STP Loops**
- **Optimal traffic path**
 - Direct path to STP root
- **Topology is deterministic**



Network Topology with Built-In Switch



- Pros: Better Rack Space Utilization
- Cons: Number of Devices to Manage, STP Load, Aggregation Port Density



With ~1,000 Servers in 12 Cabinets = ~140 Built-in Switches to Manage

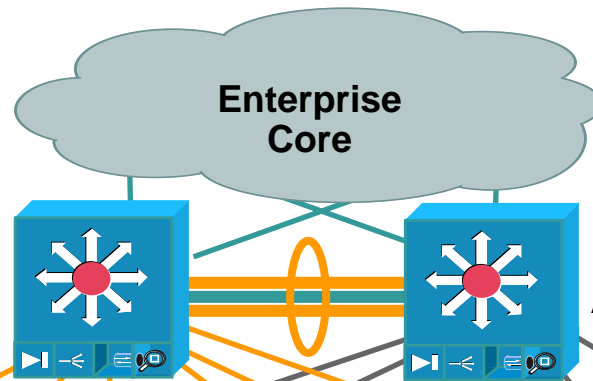
140+ Switches in a Spanning Tree Domain

Who Manages the Internal Switches; Network or Server Admins?

Network Topology with Modular Switches




Pros: Fewer Devices/mgmt
Cons: Inter-rack Cabling,
Separate Switch
Cabinets



4 10GE Uplinks per Cabinet =
16 10GE Agg Layer Ports

Aggregation




With ~1,000 Servers in 12 Cabinets = 8 Switches in 4 Cabinets

~8 Access Switches to Manage

Network Topology with 1RU Switches



Pros: Efficient Inter-rack Cabling
Cons: Number of Devices/Mgmt, STP Load




Enterprise Core

16 GE uplinks per cabinet = 224 GE Agg Layer ports



Aggregation

Access



Cabinet 1

Cabinet 2

...

Cabinet 14

With ~1,000 servers in 14 cabinets = 56 switches to manage


ACCESS LAYER DESIGN MODELS



Design Model with Layer 2 Access

Defining Layer 2 Access

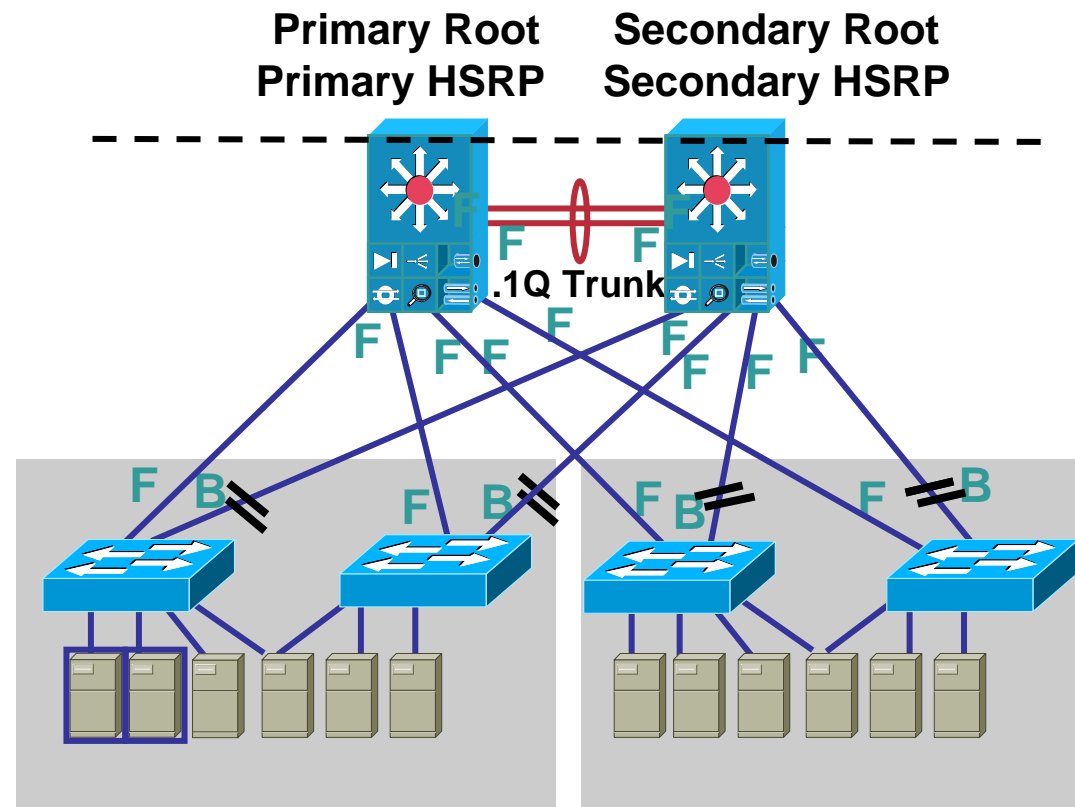
- Layer 2 access provides layer 2 adjacency between servers in the access switches
- It DOESN'T mean carrying all VLANs unnecessarily across all access switches
- L3 processing is first performed in the aggregation layer
- L2 topologies consist of **looped, loop free, and hub and spoke**



Design Model with Layer 2 Access


Looped Design

- VLANs are extended between aggregation switches, creating the looped topology
- Spanning Tree is used to prevent actual loops (Rapid PVST+, MST)
- Redundant path exists through a second uplink that is blocking
- The backup link goes forwarding when the primary link is lost



Design Model with Layer 2 Access Looped Design

- Services like firewall and load balancing can be deployed at the aggregation layer and shared across multiple access layer switches
- VLANs are primarily contained between **pairs** of access switches
- A VLAN may be provisioned on a different access switch if administrative reasons require this
- **NIC teaming and clustering** can be supported across access layer modules



Design Model with Layer 2 Access

Drawbacks of Layer 2 Looped Design


- Main drawback: if a loop occurs the network may become unmanageable due to the infinite replication of frames
- New features plus best practices improve stability and prevent loop conditions

UDLD

Loopguard


Rootguard

BPDUGuard



Design Model with Layer 2 Access Loop Free Design


- Each pair of access switches are assigned a set of VLANs specific to that pair
- No VLANs are trunked between aggregation switches
- Spanning Tree is enabled but no port is blocking
- **All links are forwarding**
- NIC teaming and clustering can be supported **within access layer modules**



Design Model with Layer 2 Access

Loop Free Design and Autostate

Cisco.com




- If the uplink connecting access and aggregation goes down, the VLAN interface on the MSFC goes down as well (autostate, i.e. when there is no port forwarding on a given VLAN, the VLAN interface on the RP goes down)
- There can be service module implications as state is not conveyed
- See new tracking features for CSM and FWSM

Design Model with Layer 2 Access

Drawbacks of Loop Free Design

Cisco.com


- If the trunk between access switch pairs is broken, the return IP path may be broken
- VLANs must be restricted to access switch pairs
- If VLAN's are extended between access layer modules then STP blocking will occur



Design Model with Layer 3 Access

Defining Layer 3 Access

- L3 access switches connect to Aggregation with a dedicated subnet
- L3 routing is first performed in the access switch itself
- L2 links between pairs of L3 access switches support **L2 adjacency** requirements (limited to access switch pairs)
- All uplinks are active, no spanning tree blocking
- Convergence time is usually better than Spanning Tree (Rapid PVST+ is close)
- Provides isolation/shelter for hosts affected by broadcasts

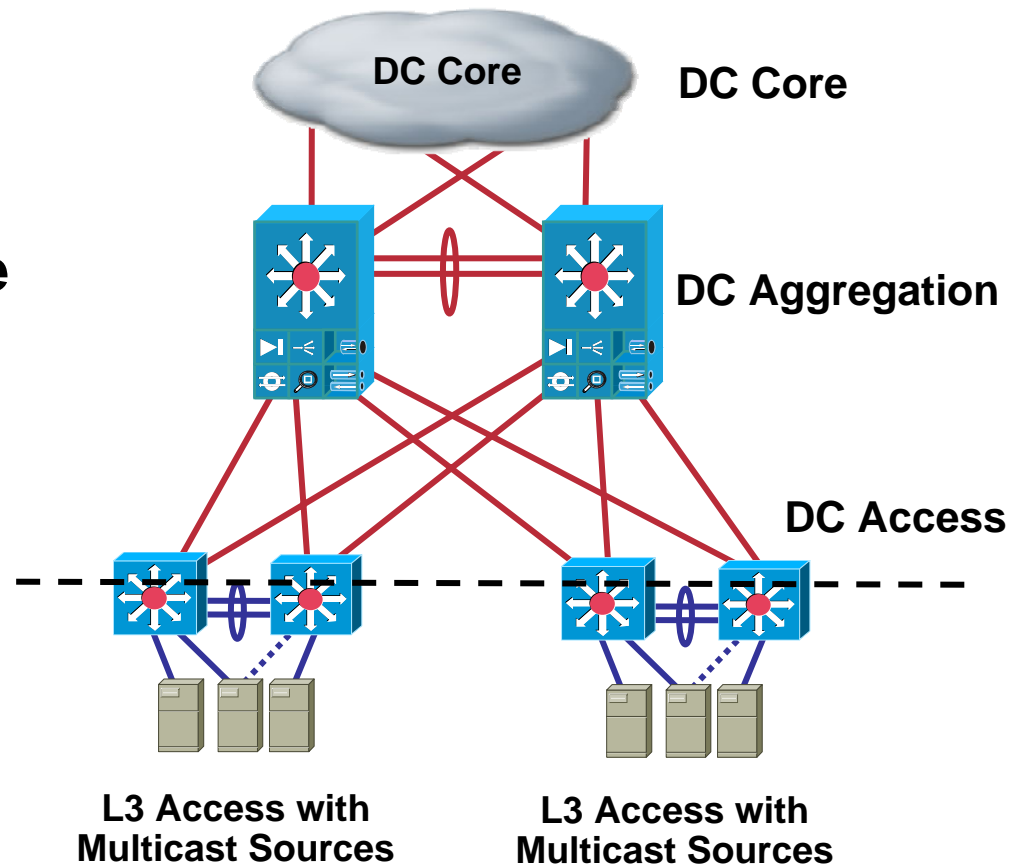


Design Model with Layer 3 Access Supporting Multicast Sources

- Multicast sources on L2 access works well when **IGMP snooping** is available
- Use L3 when IGMP snooping is not available or when particular L3 administrative functions are required

IGMP snooping at access switch automatically limits multicast flow to interfaces with registered clients in VLAN


PIM provides Auto-RP and further administrative abilities to the access layer



Design Model with Layer 3 Access

Other Layer 3 Benefits


- Minimizes broadcast domains attaining high level of stability
- Meet server stability requirements or isolate particular application environments
- All uplinks are available paths, no blocking (up to ECMP maximum)
- Load balance uplink path selection with GLBP or manual HSRP configuration
- Very good convergence time can be attained



Design Model with Layer 3 Access


Drawbacks of Layer 3 Design

- **Clustering and NIC teaming limited to access pairs**
- **IP address space management**
- **Service Module implications (next slide)**



When to Use Layer 2 or Layer 3 Access Service Module Implications


- L2 adjacency requirements between servers and service modules may be required
- Service module active/standby vs. active/active operation considerations
- Service modules require L2 adjacency for state and session synchronization
- Utilize service module interface tracking and monitor features



Layer 2 and Layer 3 Access Compared

What Are My Requirements?

Cisco.com



The Choice of One Design Versus the Other One Has to Do With:

- Difficulties in managing loops
- Staff skillset—time to resolution
- Convergence properties
- NIC teaming—adjacency
- HA Clustering—adjacency
- Specific application requirements
- Broadcast domain sizing
- Oversubscription requirements
- Link utilization on uplinks
- Ability to extend VLANs

DENSITY AND SCALABILITY IMPLICATIONS IN THE DATA CENTER



Density and Scalability Implications

Modular and 1RU Access Layer Switching Models



Cisco.com

- **What are the issues?**
 - **Server density**
 - **Management**
 - **Oversubscription**
 - **Equipment sparing**
 - **Redundancy**
 - **Cabling**
 - **STP scalability**
 - **Environmentals**
- **The right solution is completely based on business requirements**
- **Hybrid implementations can and do exist**



Density and Scalability Implications

Example Server Farm Cabinet Layout




How Many Interfaces per Server Should Be Planned For?

How to Provide Network Connectivity?

Density and Scalability Implications Cabinet Design with 1RU Switching

Servers Connect Directly to a 1RU Switch in the Rack

- Minimizes the number of cables to run from each cabinet/rack
- All servers are dual homed: 2 -1RU switches per rack are required
- Will 2 1RU switches provide enough port density?
- Cooling requirements may not permit a full rack of servers




Density and Scalability Implications Network Topology with 1RU Switching Model

↑ Pro: Efficient Cabling
↓ Con: Number of Devices/Mgmt
 Con: Spanning Tree Load

4 Uplinks per Cabinet = 100
Uplinks Total to Agg Layer

GEC or 10GE Uplinks?

Aggregation



Cabinet 1

Cabinet 2

...

Cabinet 25


With ~1,000 Servers/25 Cabinets = 50 Switches to Manage

Density and Scalability Implications

Cabinet Design with Modular Access Switches

Cisco.com

Servers Connect Directly to a Modular Switch Outside the Rack




- Cable bulk at cabinet floor entry can be difficult to manage
- Cable bulk can block air flow
- Typically placed at ends of cabinet row
- May need to space switches out within the row and at ends

Density and Scalability Implications

Network Topology with Modular Switches in the Access

Cisco.com

↑ Pro: Fewer Devices/Mgmt
Con: Cabling Challenges
Con: Cooling Challenges



With ~1,000 Servers/9 Slot Access Switches= 8 Switches

~8 Access Switches to Manage

Density and Scalability Implications

Density: How Many NICs?

- **Three to four NICs per server are common**

Front end or public interface

Storage interface (GE, FC)

Backup interface

Back end or private interface


integrated Lights Out (iLO)
for OOB mgmt

- **May require up to 4 1RU switches per rack to meet port density requirements**

30 servers with 4 active ports = 120 ports
required in a single cabinet (3x48 port
1RU switches)

May need hard limits on
cabling capacity to

Avoid cross cabinet cabling



Density and Scalability Implications

Oversubscription and Uplinks

- **What is the oversubscription ratio per uplink?**

Develop an oversubscription model

Identify by application, tier, or other means

- **Consider future true server capacity (PCI-X, PCI- Express)**

Server platform upgrade cycle will increase levels of outbound traffic

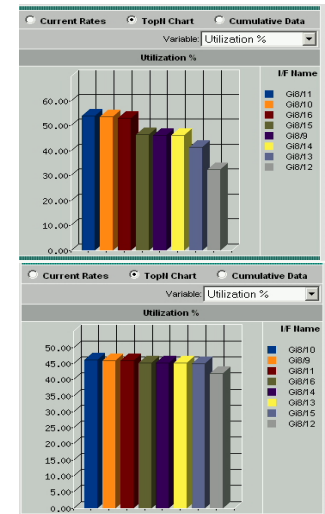
- **Consider uplink choices that will scale with your business**

Gigabit EtherChannel 10GE

10Gig EtherChannel

- **Consider flexibility in adjusting oversubscription ratio**

Can I upgrade to 10GE easily? 10G EtherChannel?



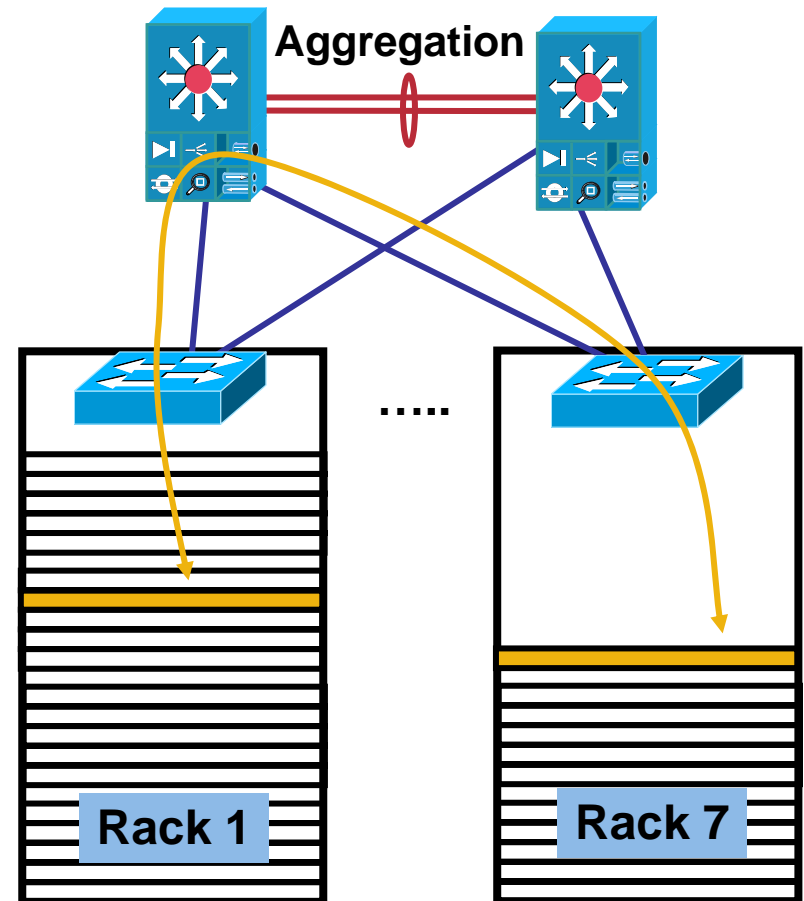
CSI1 Mark Noe 4/29/2005
fine graphic that shows gec
to 10GE

is this a request for CIR? Dana x65463
Cisco Systems, Inc.; 20-05-2005

Density and Scalability Implications

L2 Adjacency Requirements

- What if NIC teaming or clustering requirements grow later?
- How far will the VLAN need to be extended?
- How many additional STP logical ports will be added at the agg layer?
- **How to do this if a layer 3 access is used?**



Density and Scalability Implications Spanning Tree

Cisco.com

- **1RU switching increase chances of larger spanning tree diameter**
- **A higher number of trunks will increase STP logical port counts in agg layer**
- **Determine spanning tree logical and virtual interfaces before extending VLANs or adding trunks**
- **Use aggregation modules to scale STP**
(Covered in More Detail Later in Presentation)



SCALING B/W WITH GIGABIT ETHERCHANNEL AND 10GE



Scaling B/W with GEC and 10GE

Choosing GE(C) or 10GE Uplinks

What Needs to Be Considered?

- **Server NIC Improvements**
 - PCI-x and PCI-Express, RDMA, TOE
- **EtherChannel Hashing Algorithm**
- **GLBP Designs**
- **ECMP Designs**
- **10GE NICs**
 - Server consolidation efforts
- **10GE Density**
 - How to scale the aggregation layer



Scaling B/W with GEC and 10GE Server Improvements

- **PCI Improvements**
 PCI, PCI-x and PCI-Express
- **PCI-X now pervasive, PCI-Express now shipping**
- **PCI to PCI-X: 8:1, PCI-X to PCI-Express: 4:1 increase**

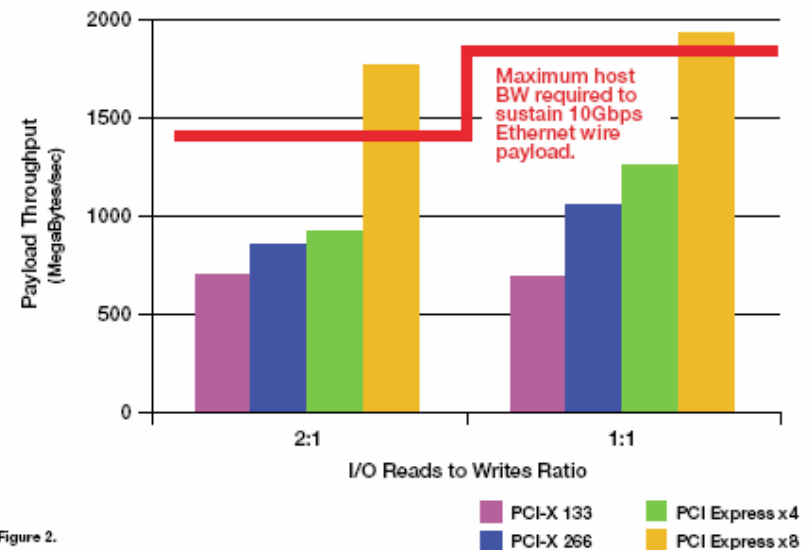
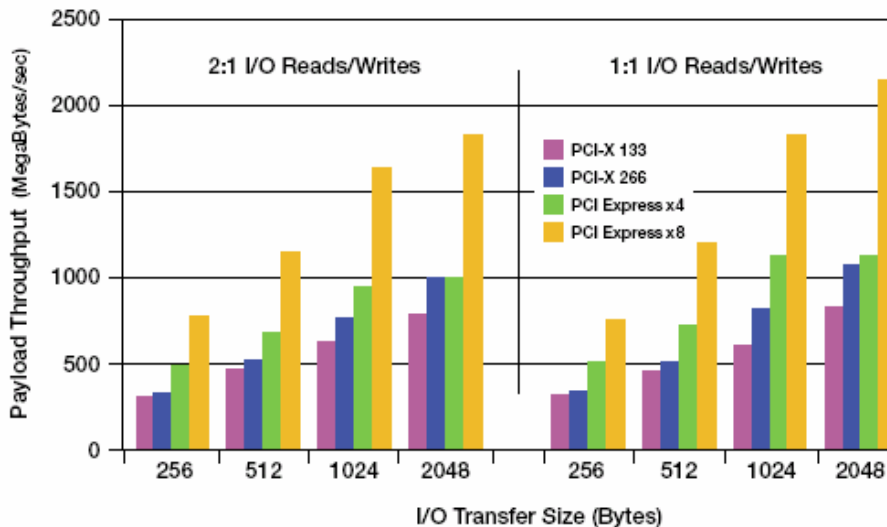


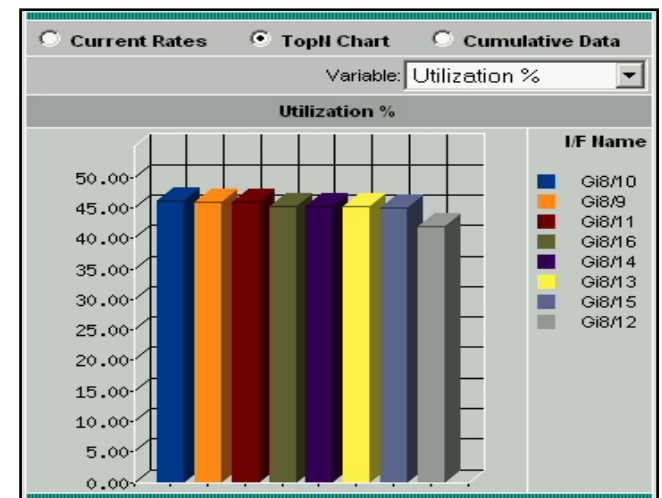

Figure 2.

Scaling B/W with GEC and 10GE

EtherChannel Hash Algorithms

Cisco.com

- EtherChannels provide link redundancy and increased bandwidth
- Can be on **physically different** line cards decreasing failure impact
- Uses two main aggregation protocols: PagP and LACP
- Optional hash algorithms available: default is L3 IP source/dest address
- The goal is to have link utilization equally distributed on all the links as in the picture on the right



Scaling B/W with GEC and 10GE EtherChannel vs. 10 Gigabit Ethernet

- Analyze the traffic flows in and out of the server farm:

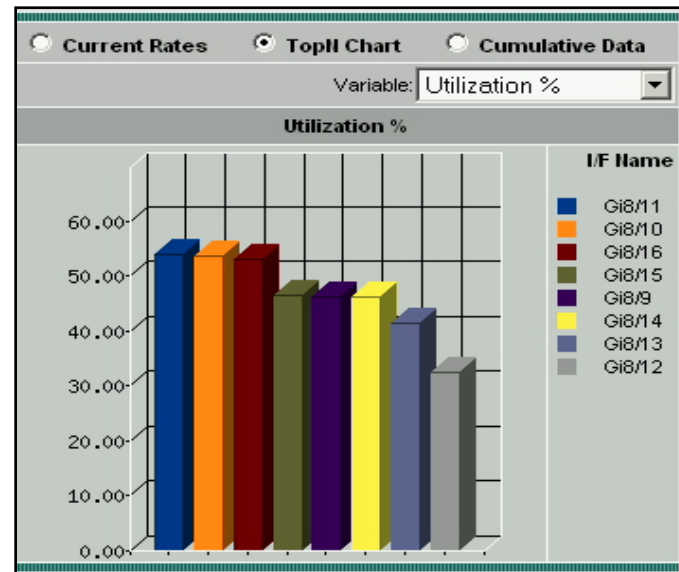

IP addresses (how many)

L4 port numbers (randomized?)

- Default L3 hash may not be optimal: look at L4 hash

```
agg(config)# port-channel load balance  
src-dst-port
```


- Ideal is graph on top right
- Bottom left graph more typical
- 10 GigE gives you effectively the full bandwidth



Scaling B/W with GEC and 10GE

Optimizing EtherChannel Paths (1)

- Under normal conditions:
 - etherchannel to the root is active
 - etherchannel to the sec-root is blocking
- What if only one link in channel is broke?
- Will STP converge to higher b/w path?
- Is it possible to put the remaining link in blocking mode and the alternate port in forwarding (the 2G channel)?



Scaling B/W with GEC and 10GE


Optimizing EtherChannel Paths (2)

- LACP-configured channels change path cost when ports join or leave the channel
- The default STP cost of EtherChannels is as follows:

- 1 Gb = 4
- 2 Gb = 3
- 3 Gb = 2
- 4 Gb = 2

- Hence if one link fails the traffic takes the single Gigabit link
- If using *spanning-tree pathcost method long*

- 1 Gig link = 20000
- 2 Gig link = 10000
- 3 Gig link = 6660
- 4 Gig link = 5000



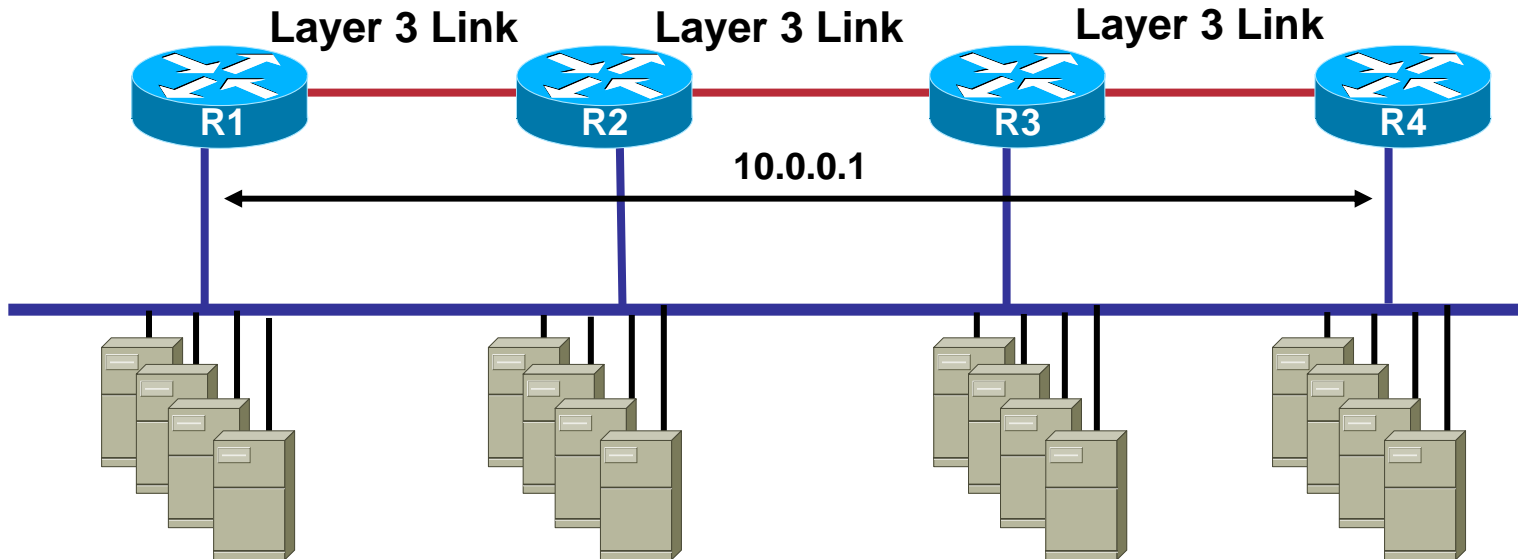
Scaling B/W with GEC and 10GE GLBP Refresher

AVG/AVF
 IP: 10.0.0.251
 MAC: MAC-R1
 vIP: 10.0.0.1
 vMAC: 0007.B400.0101

AVF
 IP: 10.0.0.252
 MAC: MAC-R2
 vIP: 10.0.0.1
 vMAC: 0007.B400.0102

AVF
 IP: 10.0.0.253
 MAC: MAC-R3
 vIP: 10.0.0.1
 vMAC: 0007.B400.0103

AVF
 IP: 10.0.0.254
 MAC: MAC-R4
 vIP: 10.0.0.1
 vMAC: 0007.B400.0103



IP: IP1, IP2, IP3, IP4
 GW: 10.0.0.1
 ARP: 0007.B400.0101

IP: IP5, IP6, IP7, IP8
 GW: 10.0.0.1
 ARP: 0007.B400.0102


IP: [...]
 GW: 10.0.0.1
 ARP: 0007.B400.0103

IP: [...]
 GW: 10.0.0.1
 ARP: 0007.B400.0104

Scaling B/W with GEC and 10GE

Using GLBP to Scale Bandwidth


- Looped topology may not be desirable: only one uplink active
- Use GLBP to distribute default gateway: maximum four gateways



Scaling B/W with GEC and 10GE

Using Layer 3 ECMP to Scale B/W

- All links are layer 3 links in this picture
- Relies on Equal Cost Multiple Path (ECMP) with DCEF load balancing
- Scales to equal cost routes that the CEF hardware supports (currently eight)
- Permits very low oversubscribed designs
- Popular for large clusters and HPC designs



Example Above:
9,216 Servers Attached with GE
3.6:1 Oversubscription
278Mbps per Server
All 6700 dCEF Enabled Modules

Scaling B/W with GEC and 10GE

Connecting Servers with 10GE

- **What is driving 10GE at the server level?**

 - Storage and clustering requirements**

 - Server consolidation efforts**

 - Virtual machine solutions on large SMP's with 802.1Q trunks**

- **Influences access layer design**

 - Oversubscription ratios are even more important**


- **Influences aggregation layer scaling**

 - Port Density**

Scaling B/W with GEC and 10GE

Migrating Access Layer Uplinks to 10GE


- How do I scale as I migrate from GEC to 10GE uplinks?
- How do I increase the 10GE port density at the agg layer?
- Is there a way to regain slots used by service modules?



Scaling B/W with GEC and 10GE


Aggregation Modules: Scaling Horizontally

A Data Center Core Provides Interconnectivity Between Multiple Aggregation Modules Which Adds Ports Required to Support 10GE to the Access Layer



Scaling B/W with GEC and 10GE Service Layer Switch Introduction

- **Where to deploy DC service modules?**
CSM, SSL, IDSM, VPNSM, etc...
- **Are all service modules required to be physically placed in the aggregation layer switches?**
- **Can I move certain service modules to other locations?**




Scaling B/W with GEC and 10GE

Service Layer Switch Introduction

Service Layer Switch

- Move certain services out of aggregation layer
- Particularly works well for CSM modules
- Extend only necessary L2 VLANs to service switches via .1Q trunks
- Opens slots in agg layer for 10GE ports
- Service Layer can utilize GEC uplinks
- Use QOS or separate GE links for FT paths



SPANNING TREE DESIGN AND SCALABILITY




Spanning Tree Scalability

Common Questions

Which STP Protocol Should Be Used?

- How many VLANs can I support in a single aggregation module?
- How many servers can I support per complex?
- How many access switches can I support in each aggregation module?
- What is the recommended oversubscription rate?
- What are the maximum number of logical ports?
- Are there any STP hardware restrictions?



Spanning Tree Scalability Common Questions (Cont.)

- **Why are these questions important?**

 - Scalability

 - Convergence

 - Throughput/performance

 - Manageability


- **Datacenter L2 designs are getting larger**

 - NIC teaming/dual homing

 - Clustering

 - Applications requiring L2 adjacency

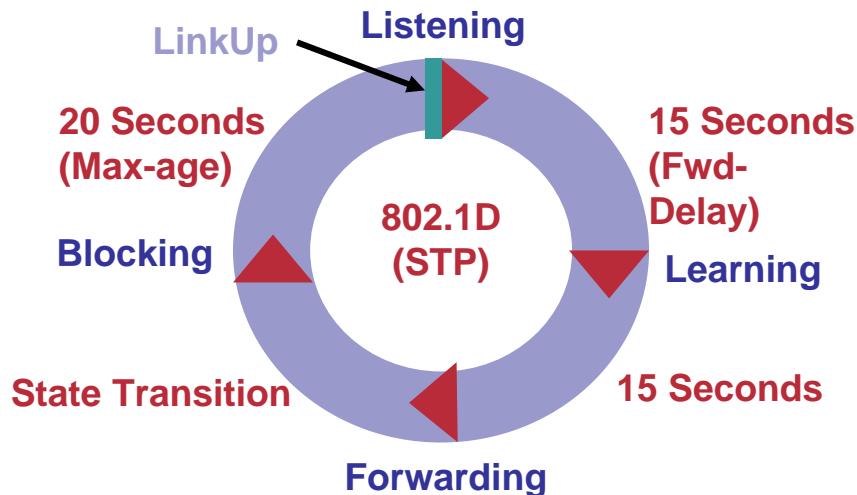
 - Server growth: adoption of blade and 1RU server technology



Rapid PVST+ vs PVST+ Protocol Highlights

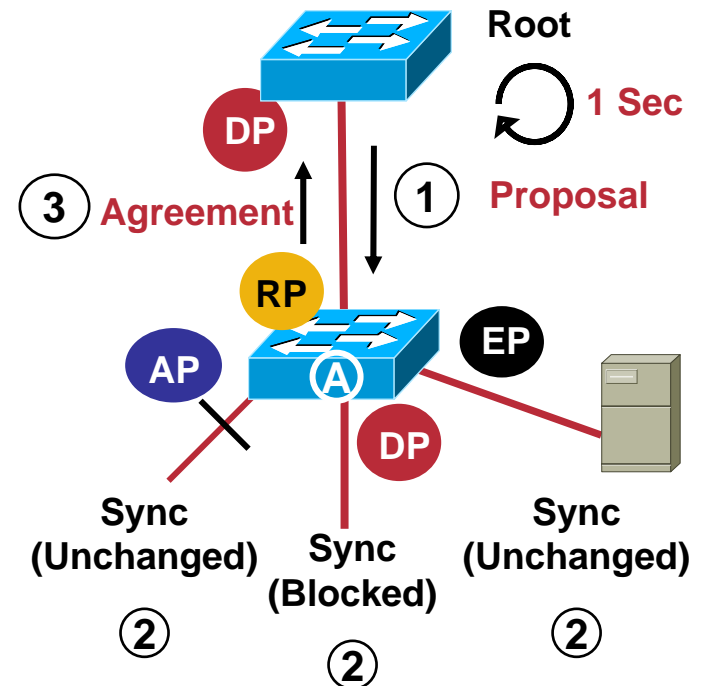
PVST+:

- **Timer-based aging protocol:**
Default convergence time ~30s for direct failures and ~50s for indirect failures
- **Uplink failure:**
Uplinkfast = few seconds
- **Indirect link failure:**
Backbonefast cuts down Maxage



Rapid PVST+:

- **Not timer based: explicit handshake between switches**
- **Uplink failure: < 1s**
- **Indirect link failure: < 1s**
- **New port roles and port states**



Spanning Tree Scalability

Spanning Tree Protocols Used in the DC

- **Rapid PVST+ (802.1w)**

 - Most common in data center today

 - Scales to large size (~10,000 logical ports)

 - Coupled with UDLD, Loopguard, RootGuard and BPDU Guard, provides a strong-stable L2 design solution

 - Easy to implement, proven, scales

- **MST (802.1s)**

 - Permits very large scale STP implementations (~30,000 logical ports)

 - Not as flexible as Rapid PVST+

 - More common in service providers and ASPs


This Focuses on the Use of Rapid PVST+

Spanning Tree Scalability

Comparing PVST+ and MST

Root for Odd VLANs

Root for Even VLANs




Load Distribution on Uplinks



If You Have 100 Vlans, the Cpu Handles 100 Topologies

Root for Instance 1

Root for Instance 2



If You Have 100 Vlans, the CPU Needs to Maintain Only Two Topologies




Less Design Flexibility

Spanning Tree Scalability Design Guidelines

General Guideline for Rapid PVST+ Implementation Using Modular Access:

- Maximum eight access switches per agg module
- Maximum 100 VLANs per modular access switch
- Add aggregation modules to scale beyond this
- **Does not consider HSRP/GLBP scalability**




Eight Access Switches per Agg Module

Spanning Tree Scalability Design Guidelines

General Guideline for Rapid PVST+ Implementation Using 1RU Access:

- Maximum 30 VLANs per access switch
- Maximum 20 access switches
- Leaves more buffer space as VLAN extension is more likely
- **Does not consider HSRP/GLBP scalability**




Hardening Spanning-Tree against Loops

Loopguard


- Rapid PVST+ include 802.1w, which is better than 802.1d, but still doesn't solve the problem of layer 2 loops.
- Layer 2 loops are still possible, but **UDLD and Loopguard** combined make their occurrence unlikely.
- Except if you have a bridging device that doesn't bridge BPDUs
- Loopguard detects Root and Alternate Port and put them into inconsistent state when they stop receiving BPDUs

Enable Loopguard on all the switches



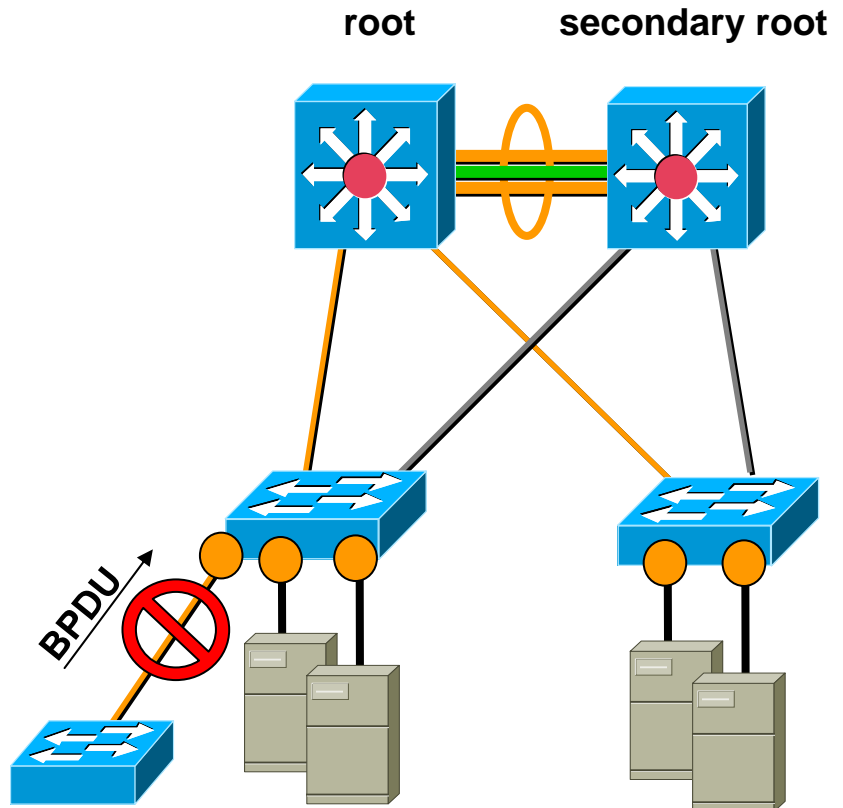
Hardening Spanning-Tree *Rootguard*

- Avoid the chance for a change in the forwarding topology as the result of a misconfiguration
- Rootguard prevents an access switch from taking over the root role
- If somebody connects a switch with a lower priority
- Or if the priority of the access switch is changed, the interface is brought down



Hardening Spanning-Tree *BPDU-Guard*

- Servers should not be sending BPDUs!
- Use BPDUguard on the access ports
- It helps against accidental daisy chaining of switches
- It helps detecting when somebody accidentally connects a switch in place of a server



Spanning Tree Scalability

Spanning Tree Best Practices Summary


Rapid PVST+
 Loopguard Default
 UDLD Enable
 Spanning Tree Pathcost Method Long
 Limit Access Nodes per Agg Module
 Scale Aggregation Horizontally

STP Root
HSRP Primary
HSRP Preempt and Delay

STP Secondary Root
HSRP Secondary
HSRP Preempt and Delay

LACP
 L4 Hashing

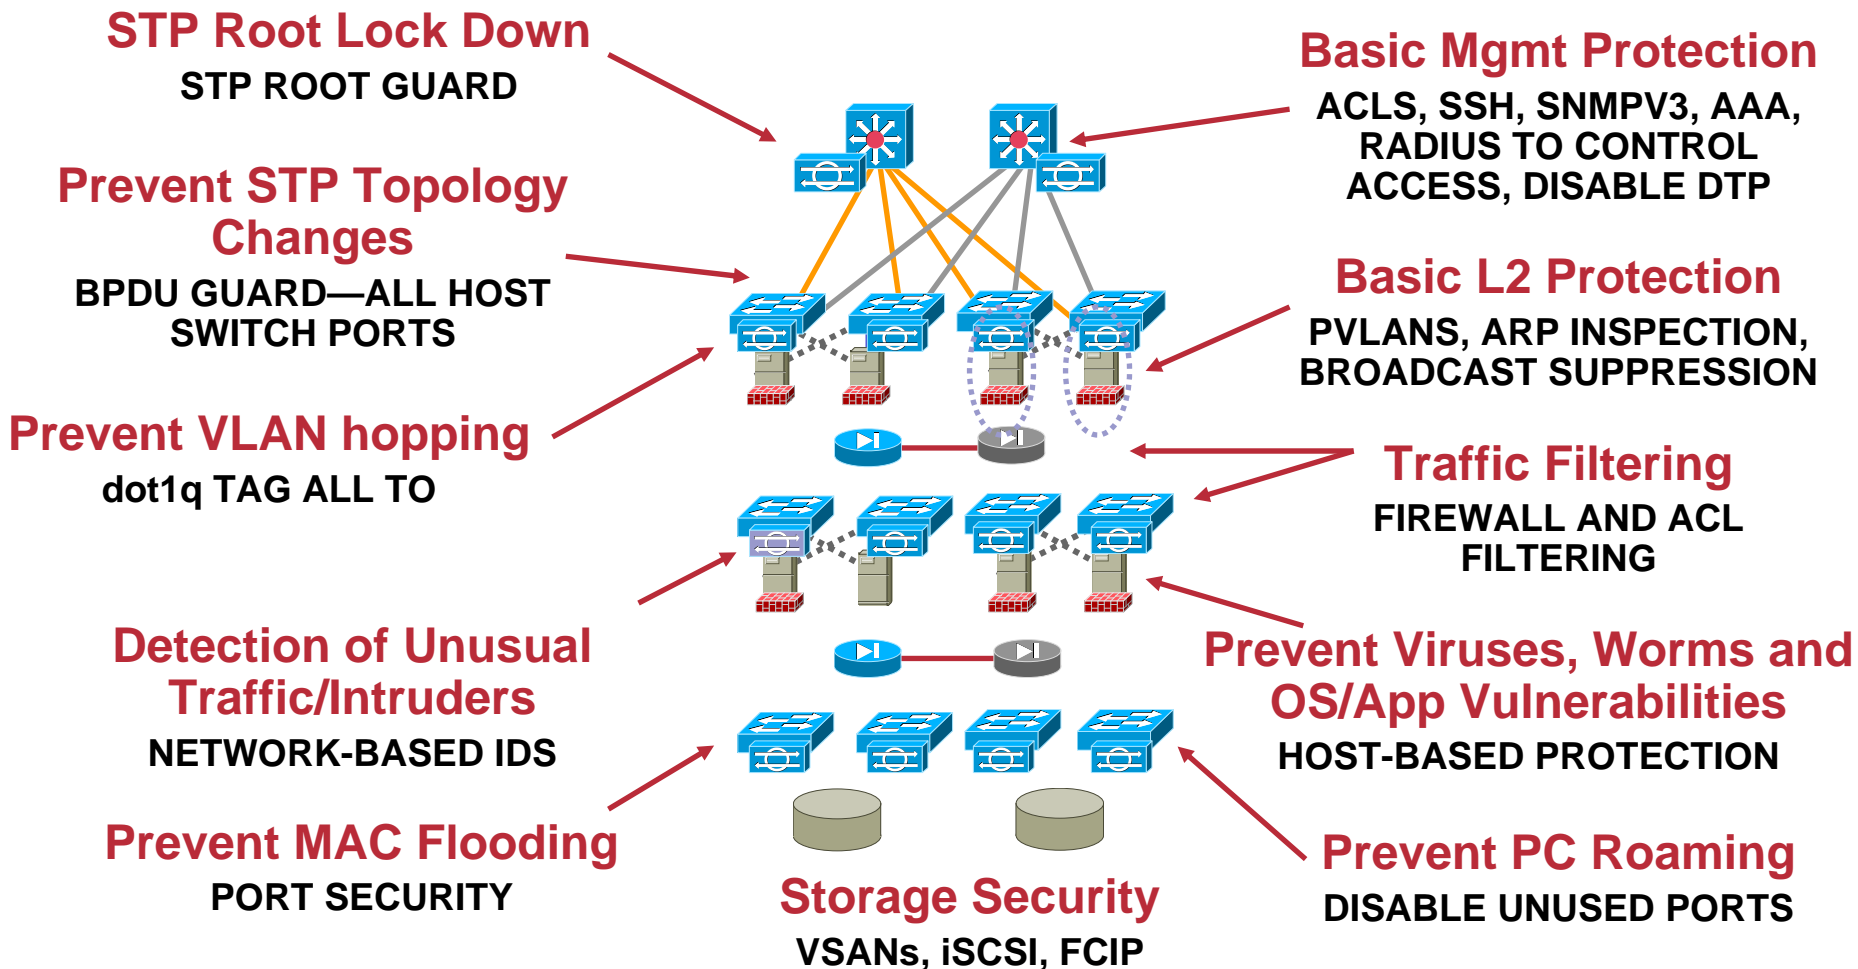
- Rootguard
- Portfast+
BPDUguard



Maximum Number of STP Active
 Logical Ports
 Virtual Ports Per Linecard Is 1800

Network Infrastructure Security Best Practices Highlights

Using the Security Features Throughout the Data Center




HIGH AVAILABILITY IN THE DC



High Availability in the Data Center

Server High Availability

Common Points of Failure



Without Data Center HA Recommendations

With Data Center HA Recommendations

1. Server network adapter
2. Port on a multi-port server adapter
3. Network media (server access)
4. Network media (uplink)
5. Access switch port
6. Access switch module
7. Access switch

These Network Failure Issues Can Be Addressed by Deployment of Dual Attached Servers Using Network Adapter Teaming Software

High Availability in the Data Center Fault Tolerance Modes

- **ACTIVE/STANDBY**



 - Adapter Fault Tolerance (AFT)

 - Switch Fault Tolerance (SFT)

 - Network Fault Tolerance (NFT)

- **Single IP address and MAC address move from active to standby adapter in the event of network failure**

 - Failover occurs in less than one second with no TCP session loss



High Availability in the Data Center Fault Tolerance Modes

- **ACTIVE/ACTIVE**


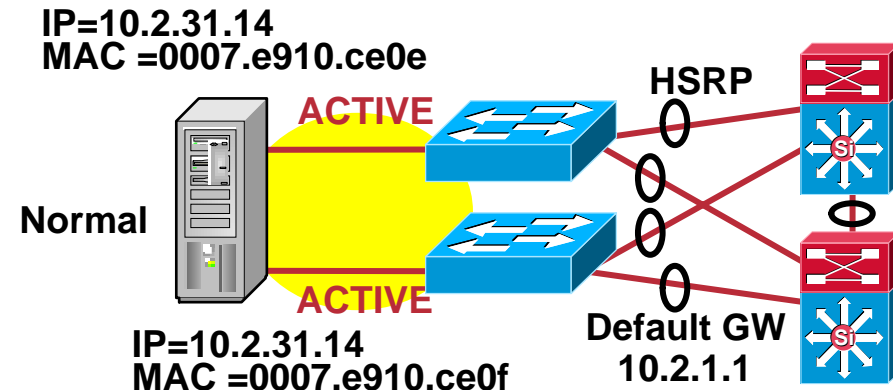
Adaptive load balancing

(IP and IPX traffic only)

- **One port receives, all ports transmit based on source dest IP; incorporates fault tolerance**

- **Multiple MAC addresses exist for the same server IP address**

- **Failover occurs in less than one second with no TCP session loss**



High Availability in the Data Center

Link Aggregation Modes

- **ACTIVE/ACTIVE**



 - EtherChannel

 - Fast EtherChannel

 - Gigabit EtherChannel

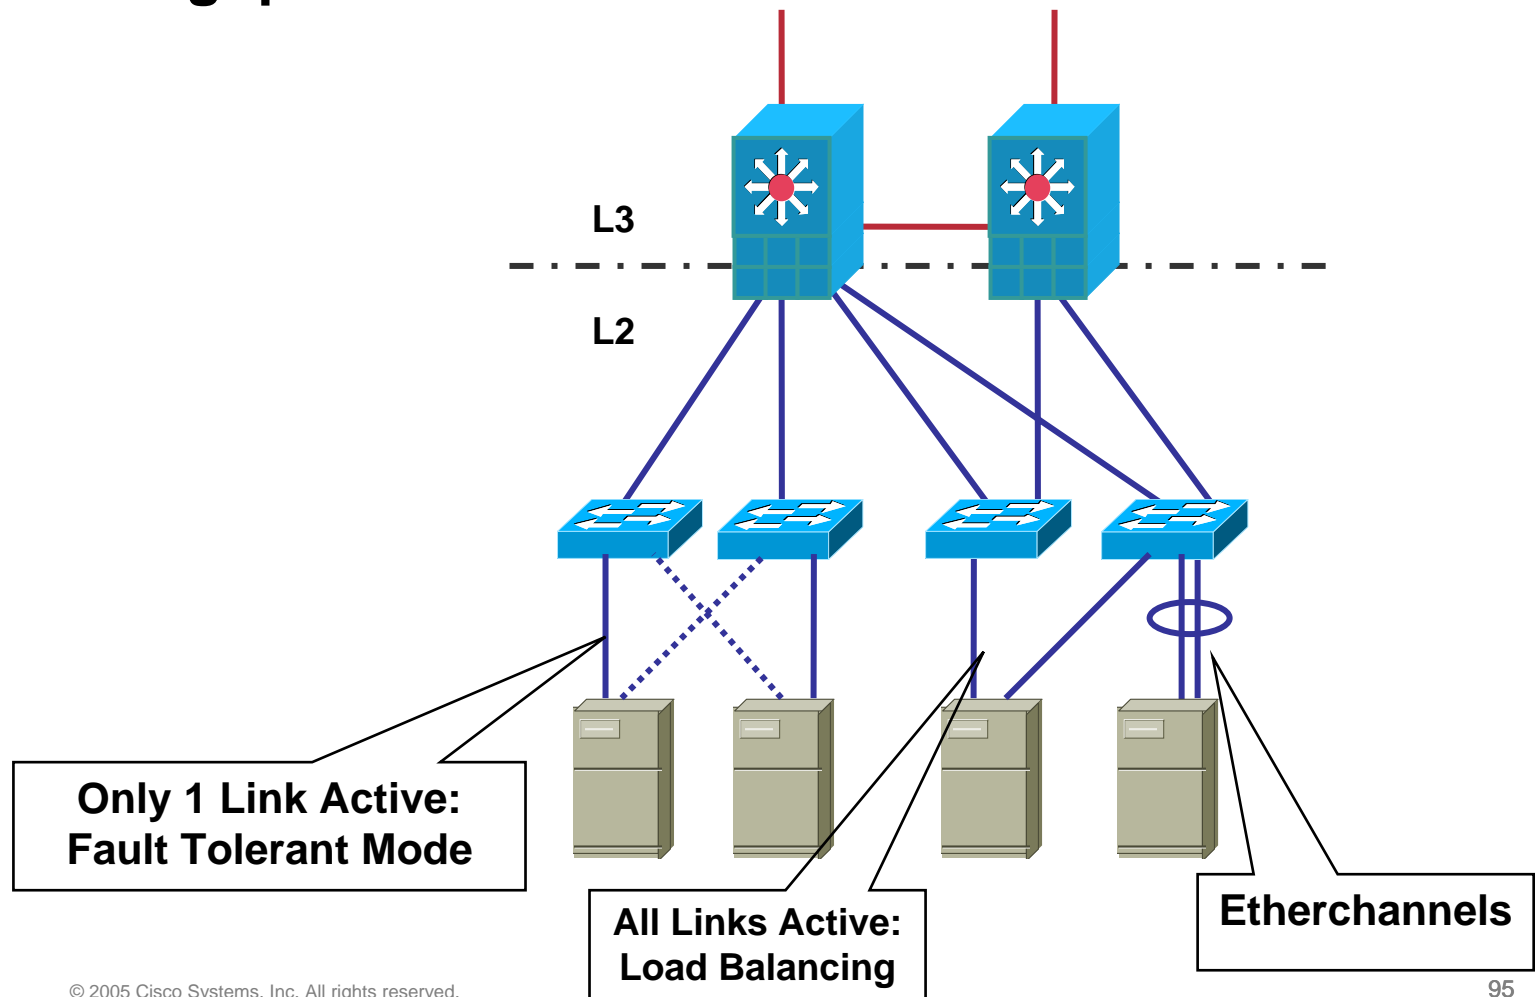
 - IEEE 802.3ad (LACP)

- Multiple physical links operate as one logical link; incorporates fault tolerance and load balancing based on source Dest IP
- Failover occurs in less than one second with no TCP session loss




High Availability in the Data Center Server Attachment: Multiple NICs

You Can Bundle Multiple Links to Allow Generating Higher Throughputs Between Servers and Clients



High Availability in the Data Center EtherChannel


- **EtherChannel load distribution works for many-to-many communications**
- **EtherChannels are ineffective for 1-to-1 communication**
- **Backup software creates multiple streams**
- **Using port aggregation of Ethernet links provides benefits for the aggregate traffic not for individual stream**



High Availability in the Data Center

Failover Times

- The overall failover time is the combination of convergence at L2, L3, L4
- Stateful devices replicate connection information and typically failover within 3s
- EtherChannels << 1s
- STP converges in <1
- HSRP can be tuned to <1s but do you need to?
- Fallback converges in ~4–5s




High Availability in the Data Center

Supervisor Failover in the Data Center

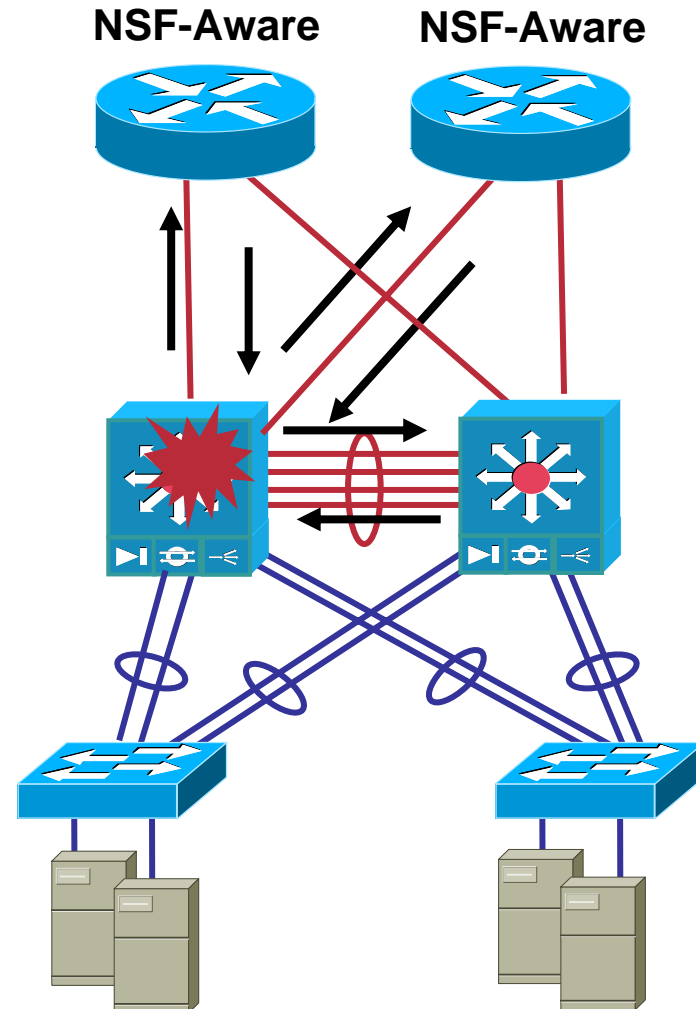
If the Supervisor in One Aggregation Switch Fails, the Failover Time Accounts For:

- The failure of the spanning tree root
- The HSRP primary
- The recovery time of load balancers and firewalls



High Availability in the Data Center NSF/SSO

- NSF/SSO is a supervisor redundancy mechanism for intrachassis failover
- SSO synchronizes layer 2 protocol state, hardware L2/L3 tables (MAC, FIB, adjacency table), ACL and QoS tables
- SSO synchronizes state for: trunks, interfaces, EtherChannels, port security, SPAN/RSPAN, STP, UDLD, VTP
- NSF with EIGRP, OSPF, IS-IS, BGP makes it possible to have no route flapping during the recovery



SUMMARY



Summary

- **Understand the applications, clustering and server HA requirements to effectively design the DC access layer**
- **Communicate well with sys-admin staff on areas of STP, VLAN extension, cabling and 10GE requirements**
- **Understand the implications of 1RU and modular access layer switches**
- **NIC teaming can greatly enhance the HA architecture**

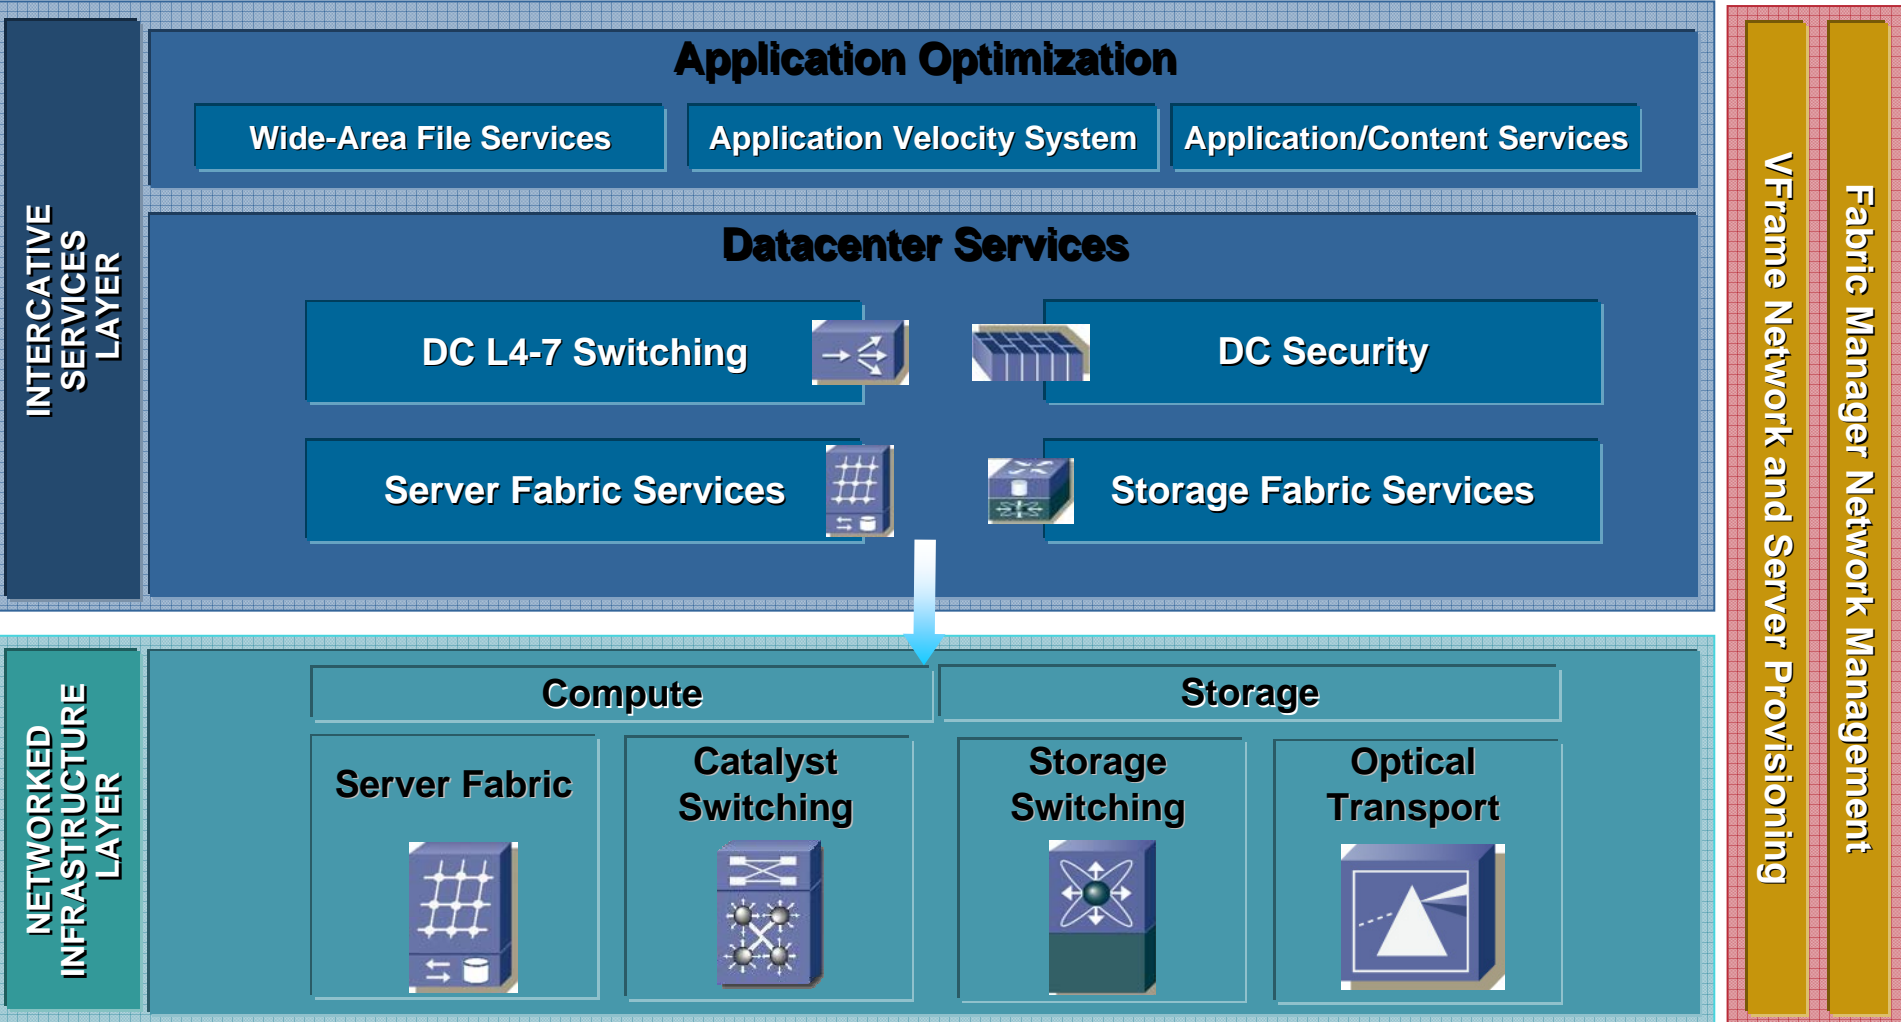


DATA CENTER PRODUCTS

**INFINIBAND SERVER SWITCHING
FIBRE CHANNEL SWITCHING
APPLICATION OPTIMIZATION**

**Mikkel Brodersen
Systems Engineer
Cisco Systems, Denmark**

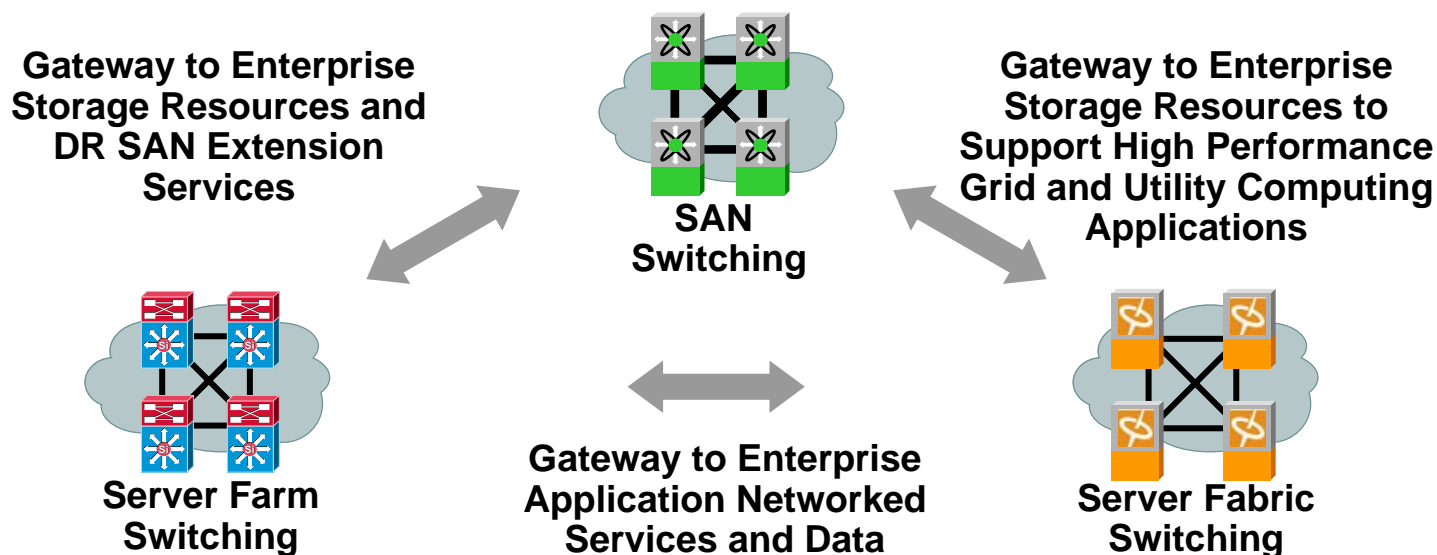
Datacenter DNA Products



Data Center Network Infrastructures

Purposeful Infrastructures—Purposeful Interconnects

Cisco.com

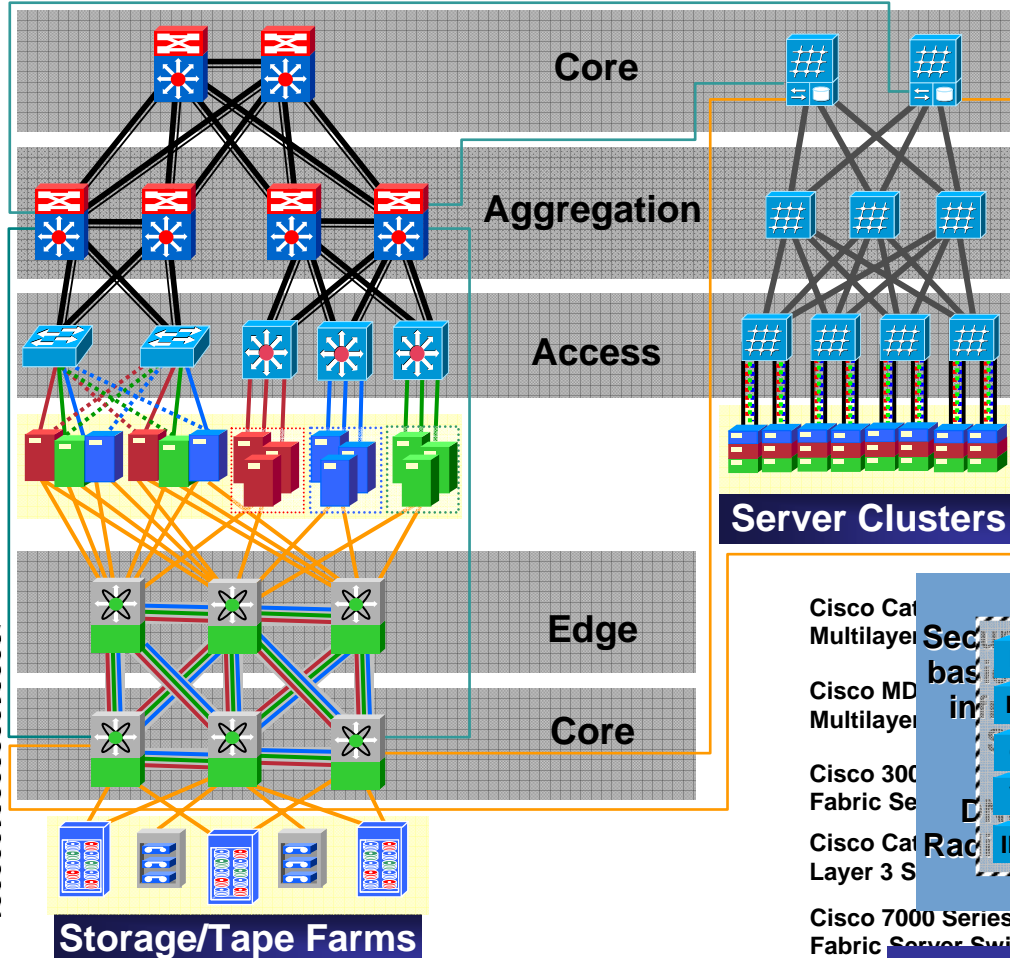


Market Maturity	Mature	Medium	New
Degree of IO Consolidation	High	Medium	Medium to High
Primary Protocols	Ethernet/IP	Fibre Channel AND ETHERNET/IP	InfiniBand ETHERNET/IP
Primary Purpose	Providing a Mature Low Cost, High Performance Server Infrastructure for Enterprise Applications	Providing a Scalable and Secure Networked Storage Infrastructure Supporting Enterprise Applications	Providing a High Performance Fabric for Server Virtualization, Clustered Enterprise Applications, and Grid/Utility Computing

Data Center Architecture

Functional Layers and Services

Layers



- ▶ Firewall Services
- ▶ Intrusion Detection
- ▶ DOS Protection
- ▶ VPN Termination
- ▶ Anomaly Detection
- ▶ Network Analysis
- ▶ App Acceleration
- ▶ AON Analysis
- ▶ Server Balancing
- ▶ SSL Offloading
- ▶ File Caching
- ▶ Content Caching

Server Farms

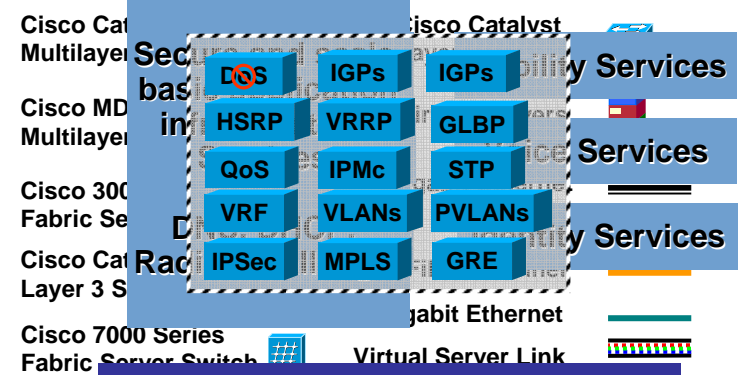
- ▶ Virtual Fabrics (VSANs)
- ▶ Storage Virtualization
- ▶ Data Replication Svcs
- ▶ Fabric Routing Services
- ▶ Fabric Gateway Services

Storage/Tape Farms

- ▶ Server Virtualization
- ▶ Virtual I/O
- ▶ Compute Fabric Services
- ▶ Remote DMA Services
- ▶ Clustering Services
- ▶ Fabric Gateway Services

Application Networking Services

Server Clusters

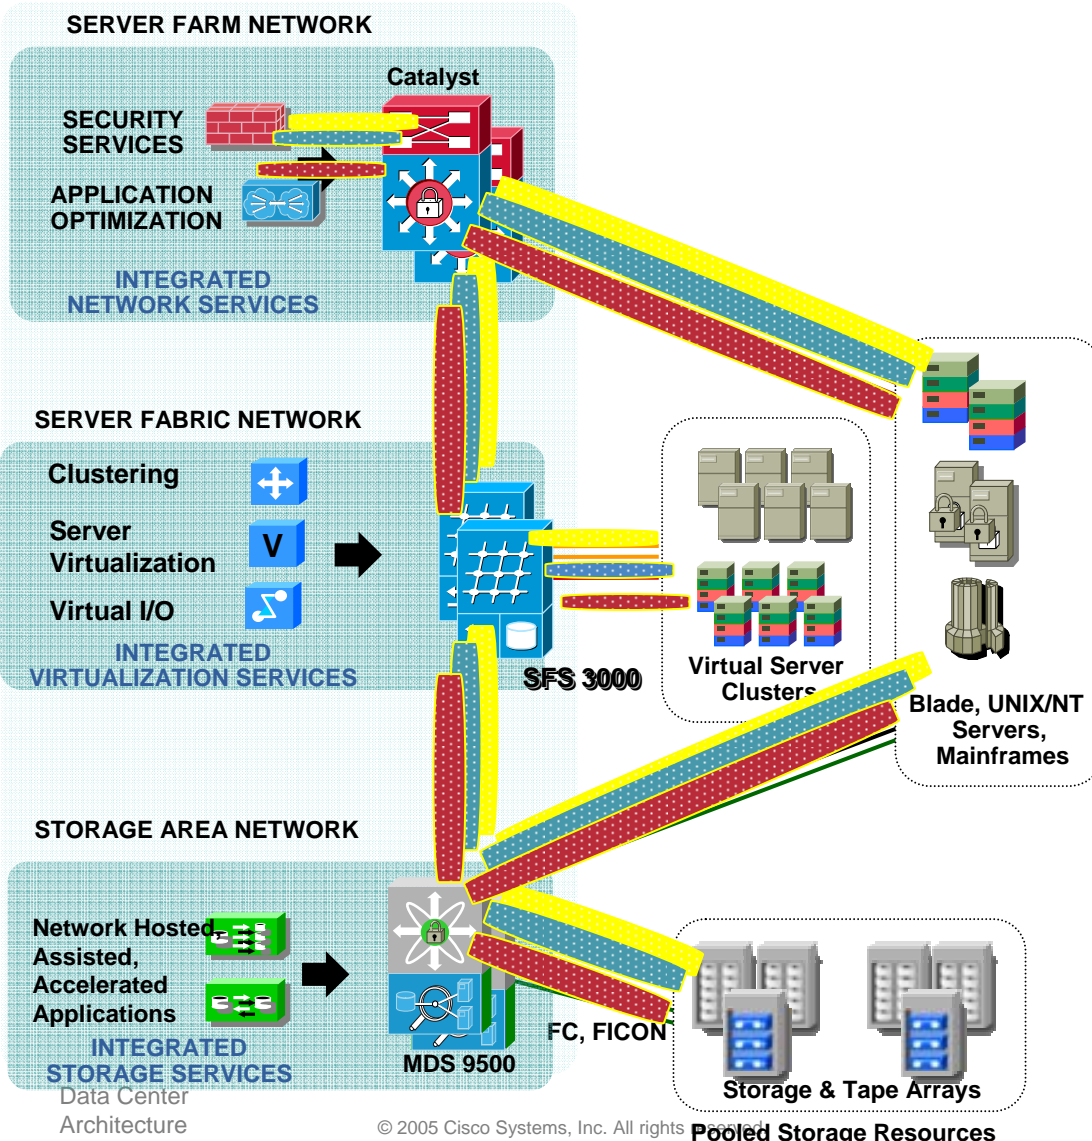


Foundation Infrastructure Functions

Enabling the Virtualized Data Center

Virtualizing Server, Storage and Network Infrastructure

Cisco.com



Compute Virtualization

- Infiniband Switching
- Server I/O Virtualization
- Policy-based Server Virtualization – VFrame

Storage Virtualization

- SAN Fabric Virtualization – VSAN
- Network-based virtualization – e.g EMC Invista, IBM SVC with MDS Storage Service Module

Virtualized Network Services

- Virtual LANs (VLANs)
- Virtual Firewalls
- Service Chaining and Virtualization

Virtualized User Connectivity

- GLB, SSL / IPsec VPN
- MPLS, VRF Mapping

Virtualized DC Interconnect


- DWDM, Sonet,
- InterVSAN Routing

EMERGING TECHNOLOGY INFINIBAND



What is InfiniBand?

- **InfiniBand is a high-speed, low-latency switched communications technology used to interconnect servers, storage and networks**
- **Routable protocol based on IPv6**
- **Standards-based**
InfiniBand Trade Association
<http://www.infinibandta.org>
- **The Cisco Server Fabric Switch (SFS) uses InfiniBand to create a high-speed, low-latency server fabric**



InfiniBand Fabric

Components of an InfiniBand Solution

Cisco.com

On the Host:

- HCA with two InfiniBand ports
- HCA drivers and ULPs

InfiniBand Switches

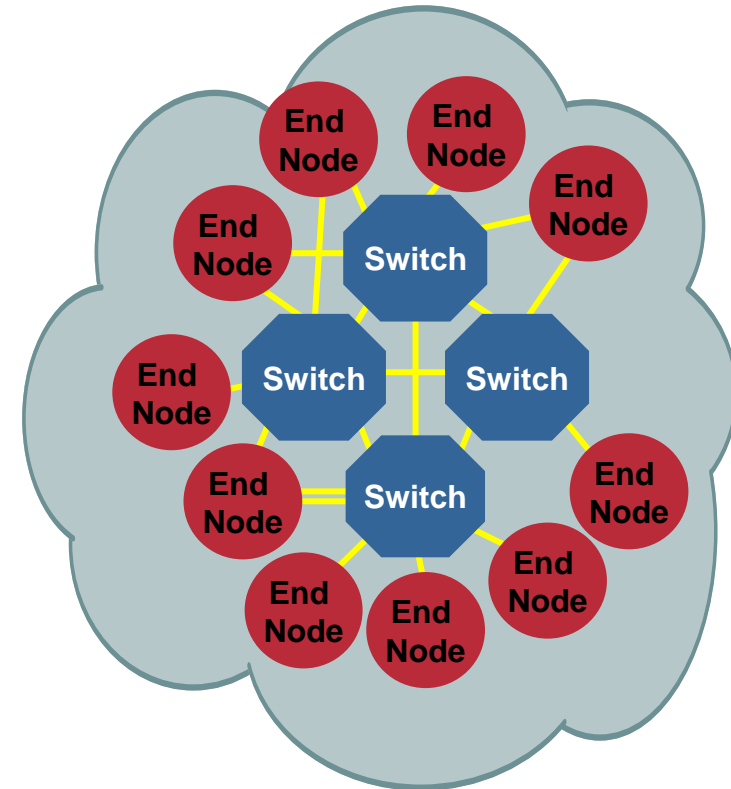
Gateways

- Ethernet Gateway
- FiberChannel Gateway



InfiniBand Architecture (IBA)

- **IBA defines a switched communications fabric allowing many devices to concurrently communicate with high bandwidth and low latency in a protected, remotely managed environment**
- **At the end-nodes, HCA/TCA off-loads from the CPU much of the I/O communications operation:**
 - Provides protected zero processor copy data transfers, with no kernel involvement, and hardware based reliable and fault-tolerant communications**

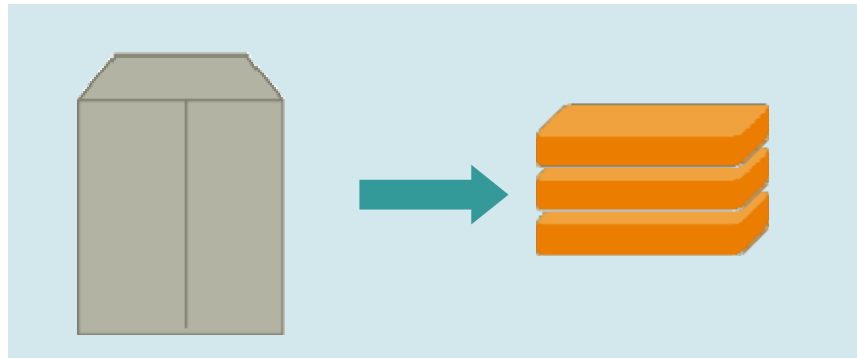


Server Fabric Switch

Infiniband—Purpose Built Server Switching Fabric

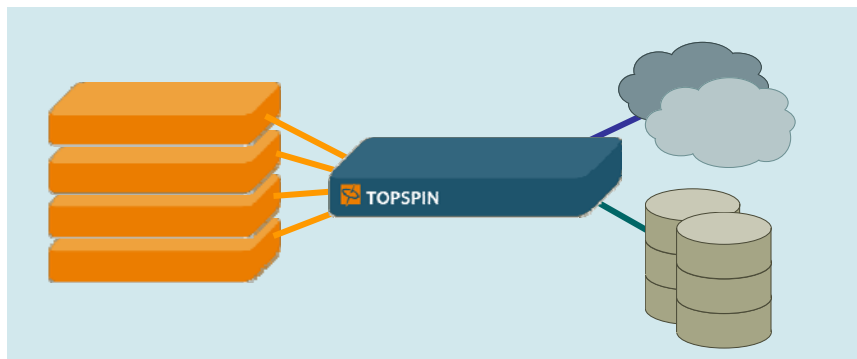
Cisco.com

Server Clustering



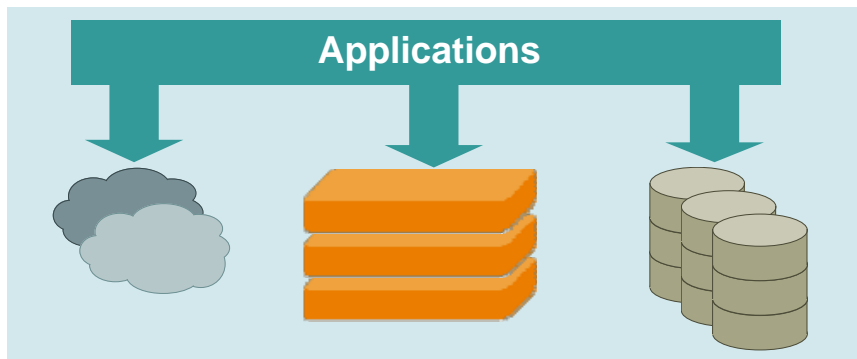
- High Performance Computing (HPC)
- “Enterprise-Class” HPC
- Database Scalability
- Low Latency RDMA

I/O Virtualization



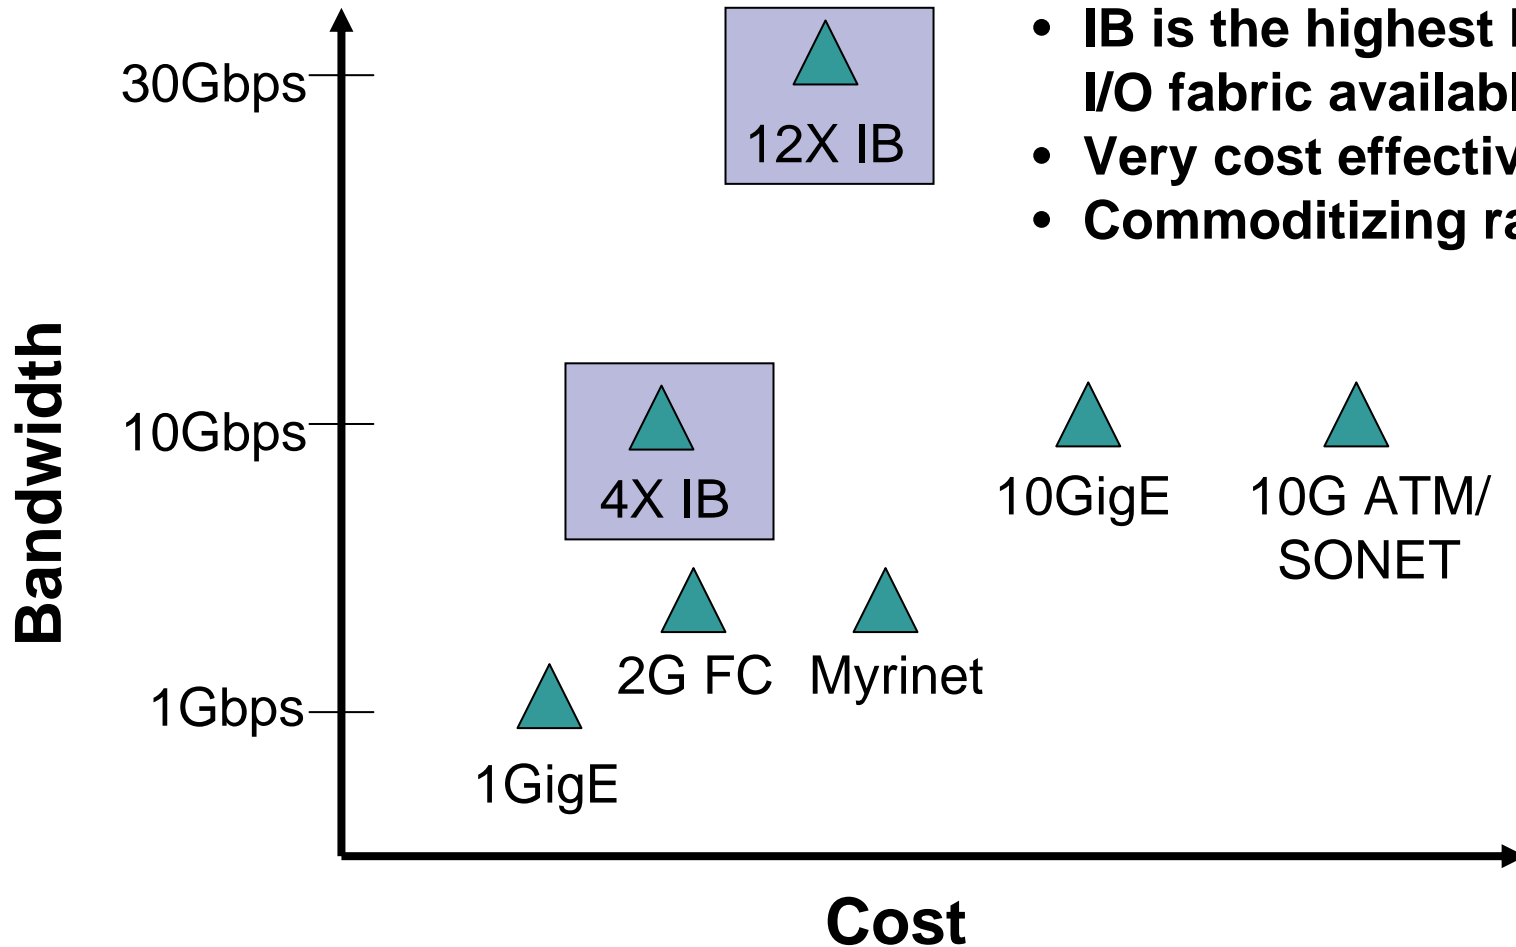
- I/O Consolidation
- I/O Aggregation
- Server Consolidation

Utility/Grid Computing



- Application Provisioning
- Server Re-purposing
- Server Migration

How Does InfiniBand Compare?



SFS Building Blocks

SFS 7000 Series




SFS 3000 Series




Gateway Modules


- InfiniBand to Ethernet
- InfiniBand to Fibre Channel



Host Channel Adapter (HCA)



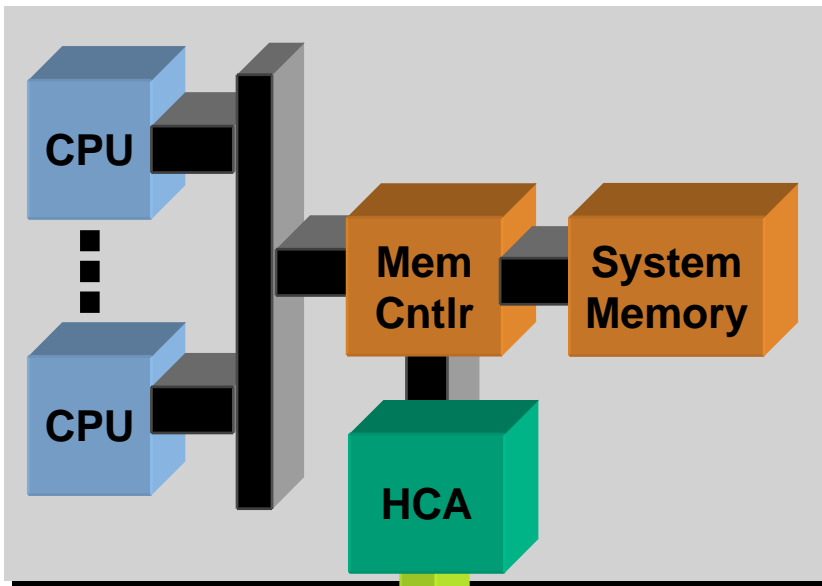
Future: IB for Blade Servers



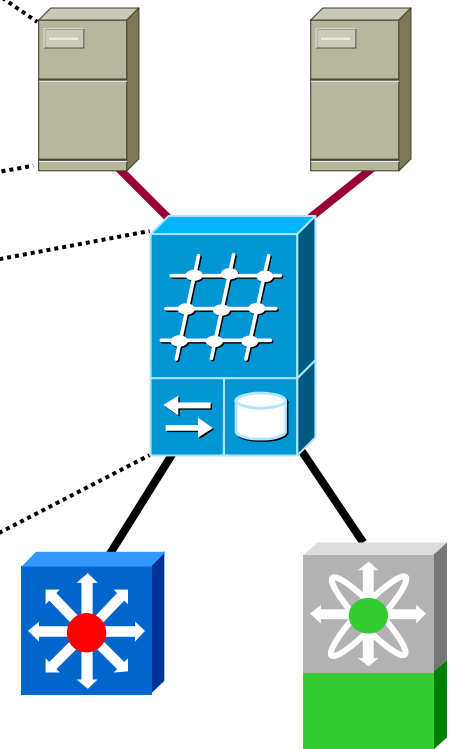

Linux and Windows driver support VFrame Host virtualization



InfiniBand Nomenclature



HCA = Host Channel Adapter
TCA = Target Channel Adapter
SM = Subnet Manager



InfiniBand Physical Layer

Physical Layer

- **Serial links with credit based link level flow control**
- **1.25 GHz (2.5 Gbps) signaling rate, with an effective data rate of 2 Gbps (due to 8b/10b encoding)**
- **1X, 4X and 12X link widths providing 2.5, 10 and 30 Gbps bandwidth**
- **Support for serial data polarity inversion and lane reversal**
- **Copper and fiber support**

Cable

- **Copper option: 4X cable typically used; Turning radius is 4 inches**
- **Fiber option: GBIC like transceiver used to drive an optical ribbon cable of length up to 100 meters**


InfiniBand Links

- **Consists of two unidirectional physical links:**

Each physical lanes supports 2.5 Gbps with 8b/10b encoding

Each physical link is 1X, 4X, or 12X physical lanes wide

	Raw	Effective
1X	2.5 Gbps	2 Gbps
4X	10 Gbps	8 Gbps
12X	30 Gbps	24 Gbps



InfiniBand Physical Interfaces

- **Copper**

 - Up to 15m* for 4x connections

 - Up to 10m for 12x connections

- **Optical**

 - Initial availability via external transceiver

 - Up to 300m with current implementation

 - Long-haul possible in future releases

- **PCB**

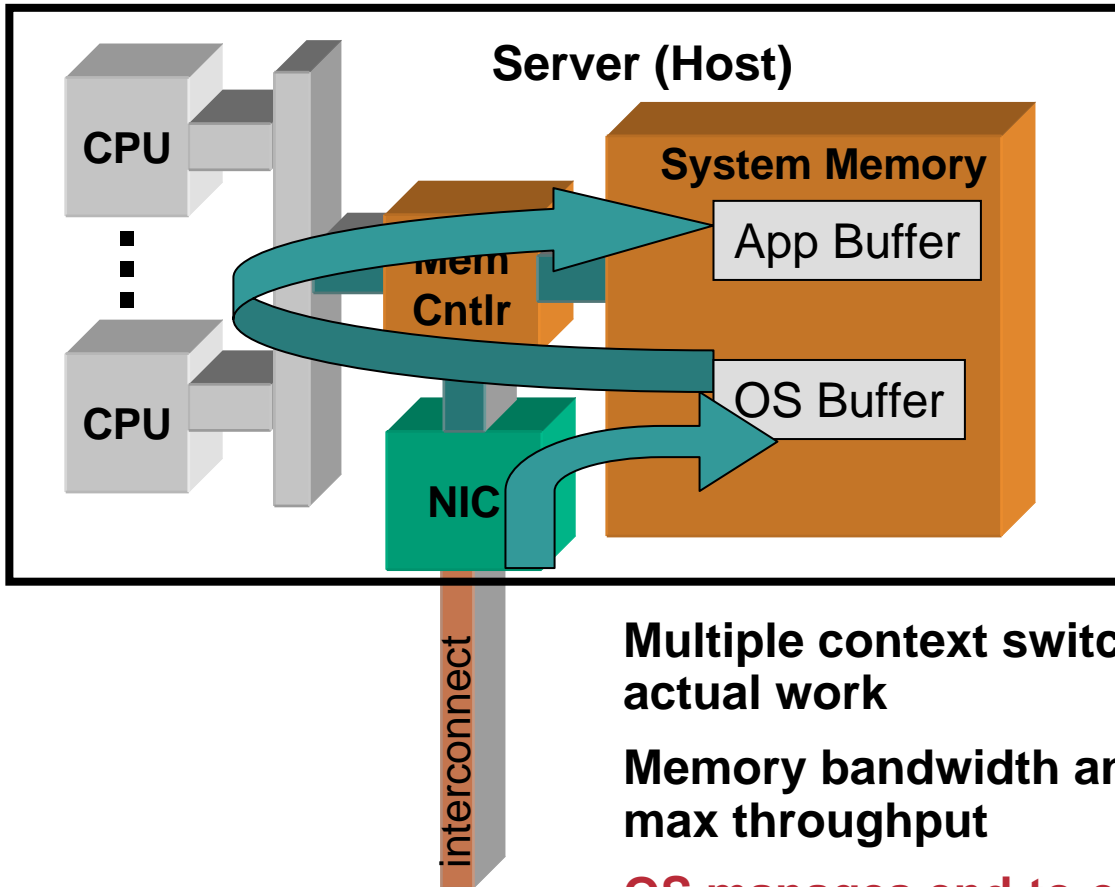
 - Specified for in-box implementations



* 20m in certain circumstances

InfiniBand's Secret Sauce: RDMA

Traditional Copy on Receive



Data traverses bus 3 times

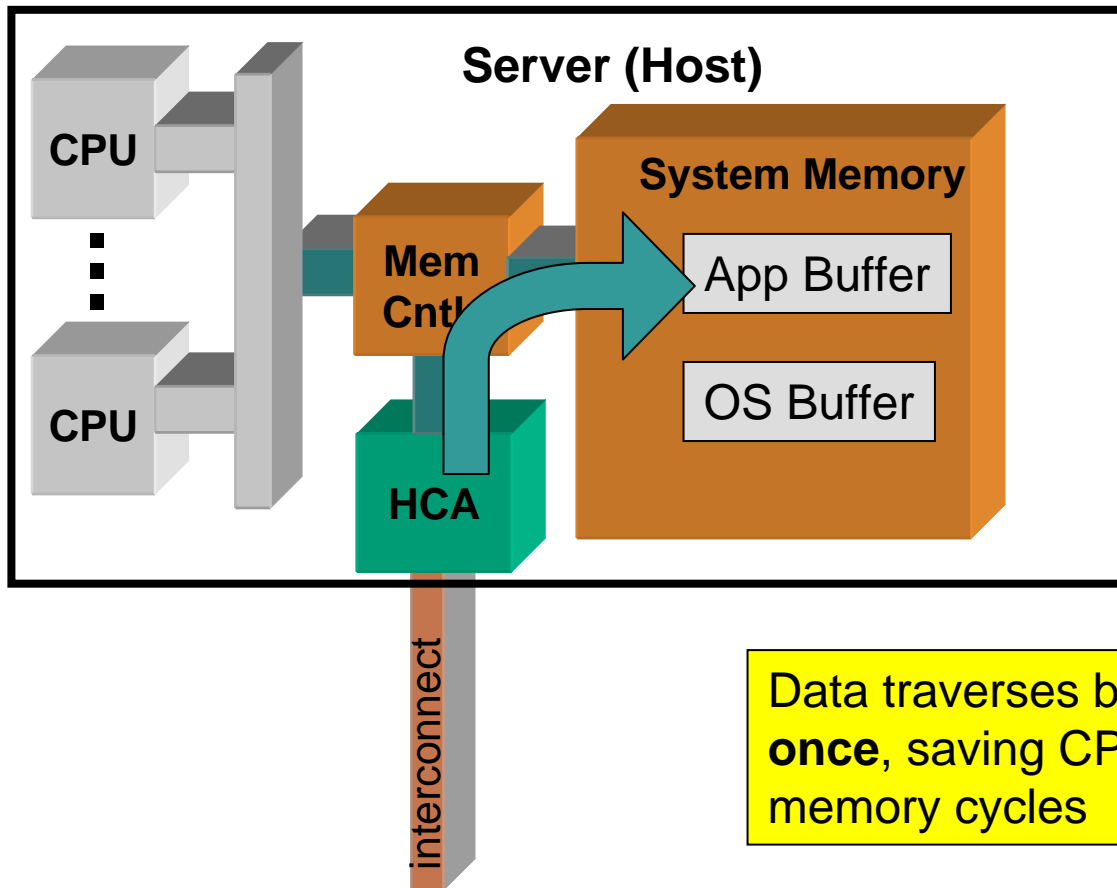
Multiple context switches robs CPU cycles from actual work

Memory bandwidth and per packet interrupts limit max throughput

OS manages end-to-end communications paths

InfiniBand's Secret Sauce: RDMA

InfiniBand Technology: RDMA & OS Bypass



Data traverses bus **once**, saving CPU and memory cycles

The Evolution of I/O Virtualization


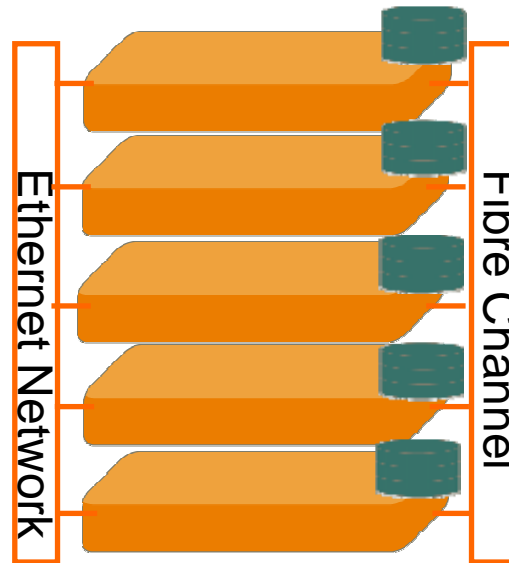
SMP



Pro: Single managed entity, fast backplane

Con: Expensive, proprietary server + backplane


Dis-aggregation



Pro: Standard servers, inexpensive


Con: Lots of managed components, low-performing interconnect

Server virtualization




Pro: Reduced # of managed components, virtual I/O, fast standards backplane

Three Types of Server Virtualization



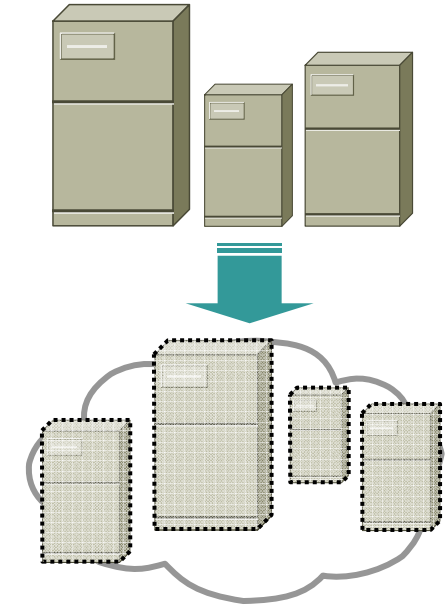
Virtual Machine

- Server partitioning
- Partitions a single server into multiple virtual servers
- Scaling up



Virtual SMP

- Tightly integrated server clustering
- Aggregates multiple homogeneous servers into a single managed entity
- Scaling out



I/O virtualization:

- Makes servers stateless by moving server identity into the network, including storage and I/O subsystem
- Supports infrastructure change in heterogeneous environments

Three Categories of Server Virtualization

- **Virtual Machine:** Splits a servers into independent virtual servers.
VMWare, XEN, MSFT
Main value is higher server utilization.
- **Virtual SMP:** Combines servers together into a single managed powered entity.
Virtual Iron, Qlusters
Main value is scaling mission critical apps on commodity HW.
- **Physical Server Virtualization:** Makes servers stateless by moving server identity into the network, including storage and I/O subsystem.
Cisco VFrame™, Egenera
Main value is making infrastructure change easier in heterogeneous environment.

VFrame Architecture

- **A Virtual Server**

Everything but the physical hardware including:

Ethernet MAC addresses

Fibre Channel WWNs

O/S image

Scripts

- **Virtual Server Group combines:**

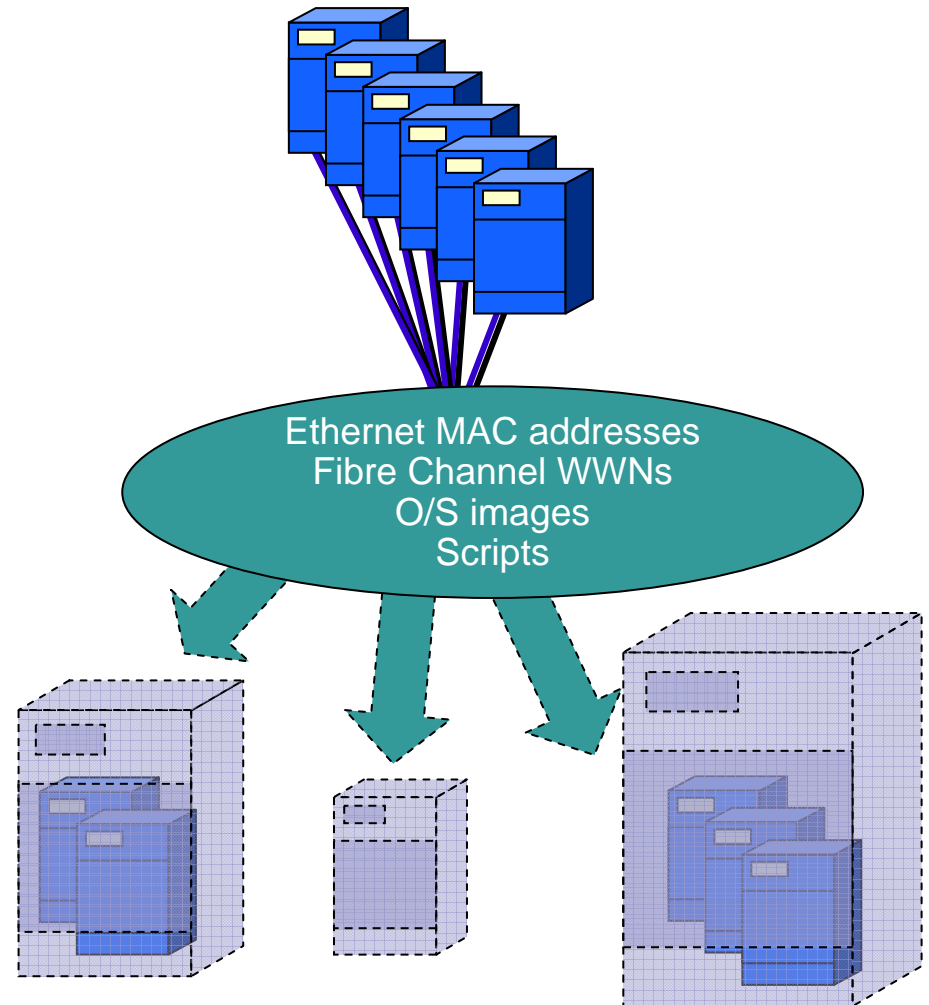
Servers

Storage

Networking

- **Policies**

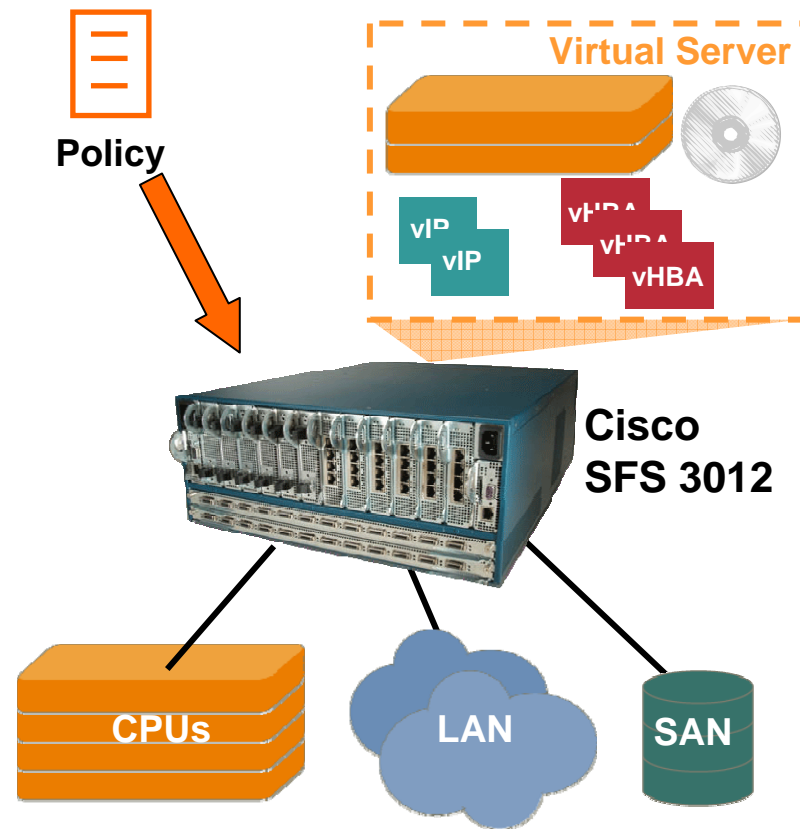
Triggers/Actions



Programmability

VFrame™

- 1) Server Switch receives policy from VFrame™ Director or 3rd party software.
- 2) Based on policy, Server Switch assembles the virtual server
 - Selects server(s) that meet minimum criteria (e.g. CPU, memory)
 - Boot server(s) over the network with appropriate app/os image
 - Creates virtual IPs in servers and maps to VLANs for client access.
 - Creates virtual HBAs in servers and maps to Zones, LUNs, and WWNNs for storage access




I/O Gateways for Network and Storage

Eliminating Technology Islands

Single InfiniBand link for:


- Storage
- Network

Server Cluster



InfiniBand switches for cluster interconnect

- Twelve 10Gbps InfiniBand ports per switch card
- Up to 72 ports total ports with optional modules
- Single fat pipe to each server for all network traffic



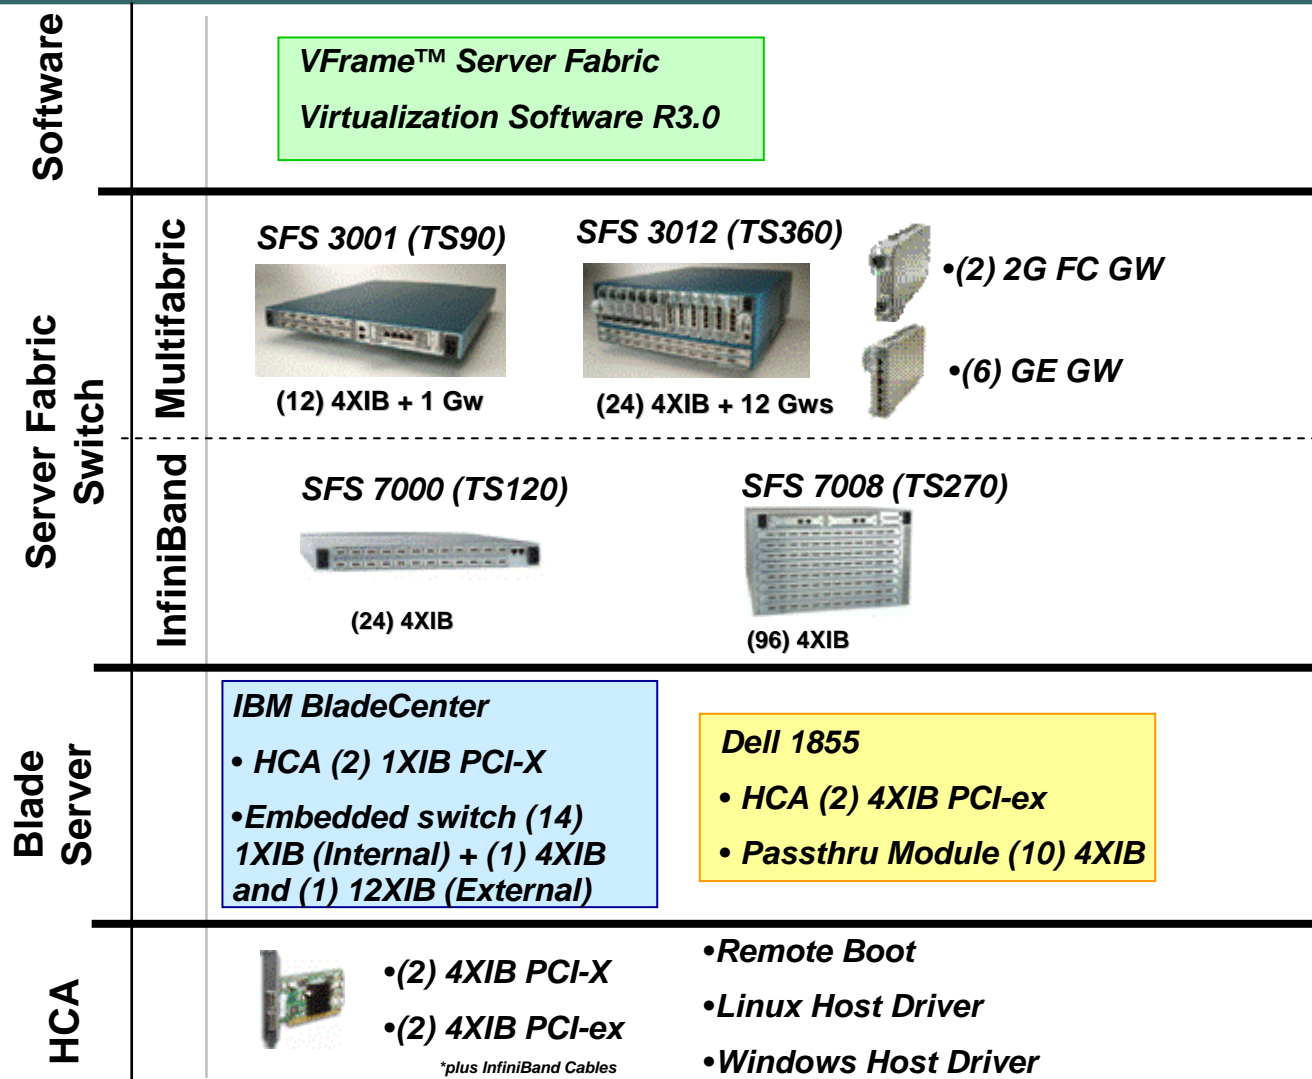
FC to IB gateway for storage access

- Two 2-Gbps Fibre Channel ports per gateway
- 10-Gbps virtual storage pipe to each server

Ethernet to IB gateway for LAN access

- Six Gigabit Ethernet ports per gateway
- Create a virtual GigE pipe to each server


The Cisco SFS Product Line



IBM BladeCenter - InfiniBand

- Cisco (Topspin) InfiniBand Host Channel Adapter Card (HCA)

Two InfiniBand ports to the CPU blade.




- Cisco (Topspin) InfiniBand Switch Module

InfiniBand switching capability to hosts in the BladeCenter chassis

Hosts communicate over InfiniBand within the chassis.

Hosts can communicate with an external InfiniBand fabric.



Dell PowerEdge 1855 and Cisco/Topspin

- One HCA per blade server
 - PCI Express
 - Two 4x IB ports
- HCA(s) connect to Pass-Thru module
 - 10 Ports internally and externally
 - Up to two Pass-Thru modules
- Pass-Thru module connects to external IB switch
- Supports Cisco/Topspin Server Virtualization



*Topspin InfiniBand
HCA Daughtercard*



*Topspin InfiniBand
Pass-Through Module*

Case Study: Leading Research Facility

High Performance Computing Cluster

Cisco.com

- **Application:**

High Performance Computing Cluster

Compute time outsourced to Commercial Enterprises (major oil & gas)

- **Environment:**

520 Dell Servers

3:1 Blocking ratio

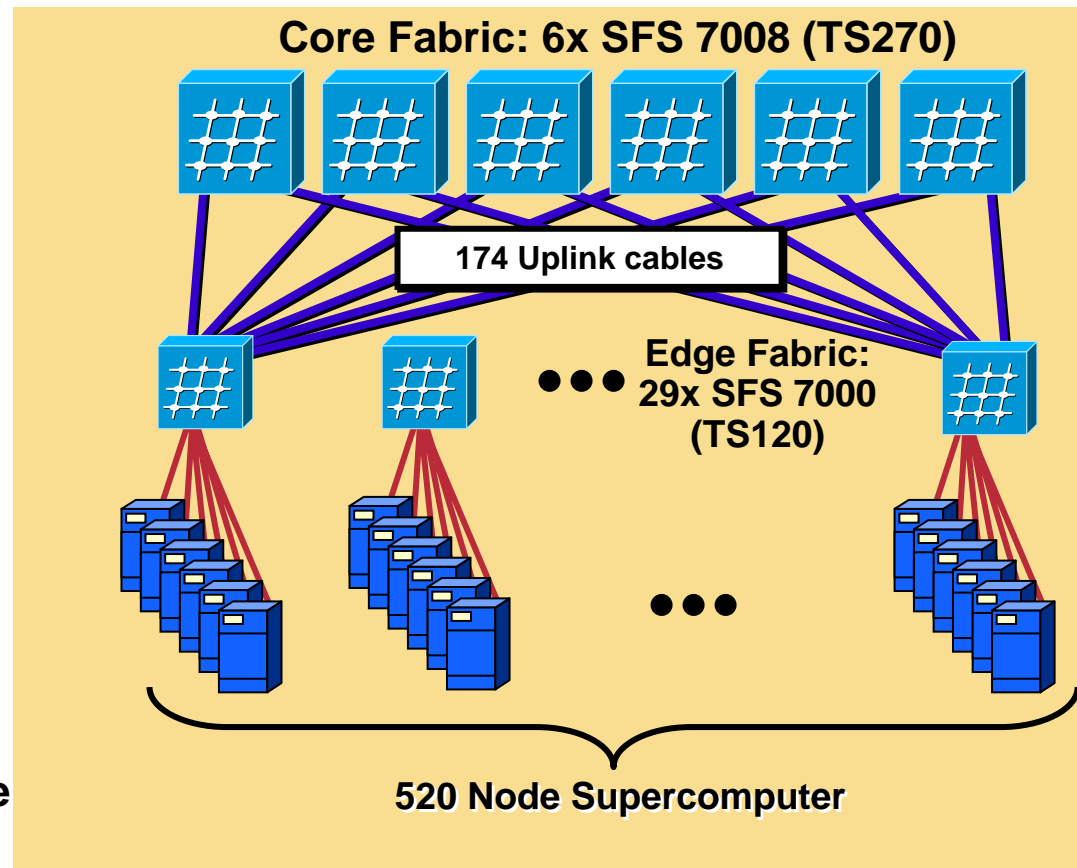
6x SFS 7008 (TS270)

29x SFS 7000 (TS120)

- **Benefits**

Compelling Price & Performance

Measured MPI latency 5.2 μ s



Case Study: Large Wall Street Bank

Enterprise Grid Computing

- **Application:**


- Replace proprietary platforms with standards-based components
 - Build scalable “on-demand” compute grid for financial applications

- **Environment:**

- 500+ Intel Servers per slice
 - Topspin Server Switch with Ethernet and Fibre Channel Gateways
 - Hitachi RAID Storage
 - SAN Switches
 - Ethernet Switches

- **Benefits:**

- 20X Price/Performance Improvement over four years
 - 30-50% Application Performance Improvement
 - Standards-based solution for on-demand computing
 - Environment that scales using 500-node building blocks



Case Study: Major System Vendor Utility Computing Service

Cisco.com

- **Application:**

Build scalable “on-demand” compute service for enterprise customers (license \$/CPU)

Key initiatives around Financial Services and Energy verticals

- **Environment**

1024x Sun V20z Nodes

34 TS270 Server Fabric Switches

Non-Blocking CLOS Network

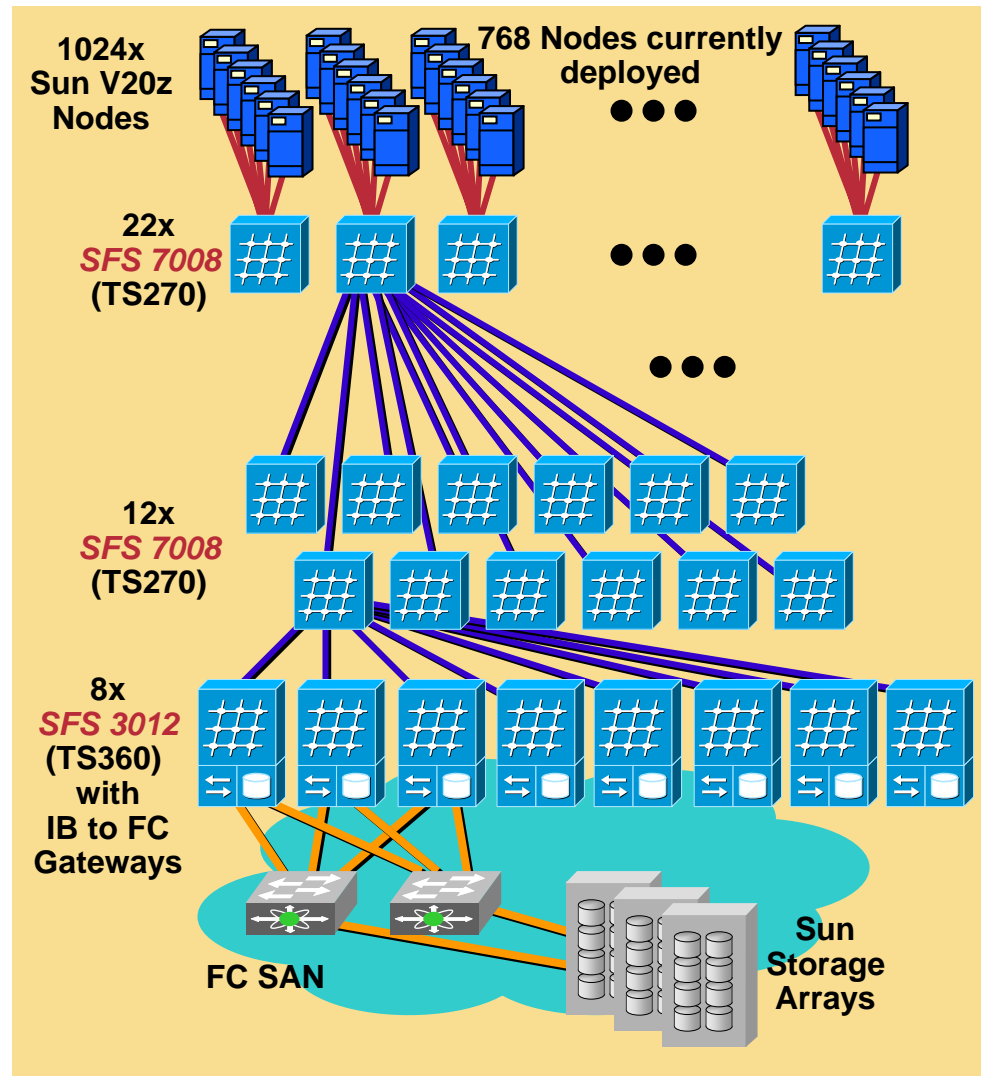
8 TS360's with Gateways

Sun Storage

Enterprise-Class Reliability

- **Benefits:**

Ability to outsource computing services to many customers with common infrastructure



Case Study- Large Government Lab

Worlds 2nd Largest Super Computer

- **Application:**

High Performance
SuperComputing Cluster

- **Environment:**

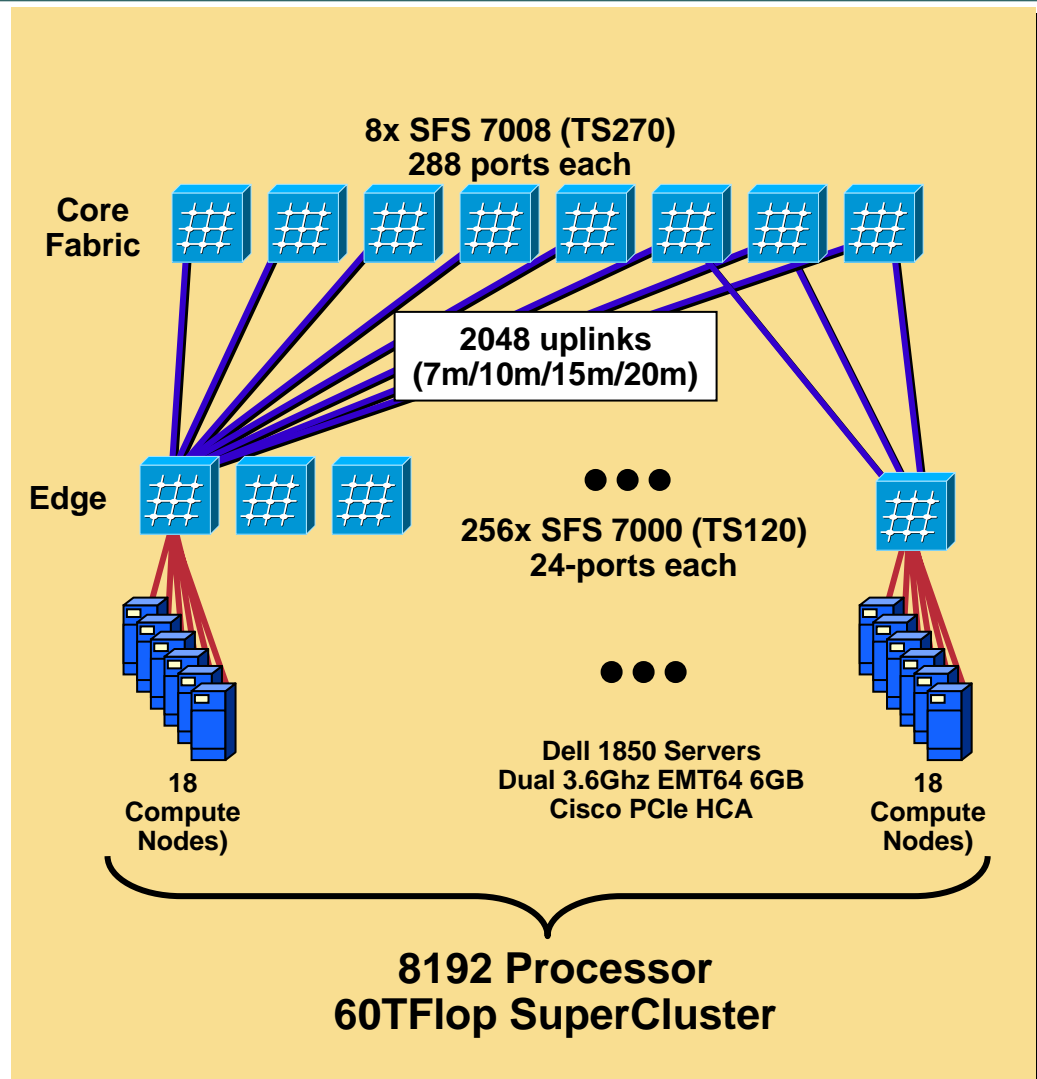
4096 Dell Servers
50% Blocking Ratio
8 TS 740s
256 TS120s

- **Benefits:**

Compelling
Price/Performance

Largest Cluster Ever Built
(by approx. 2X)

Expected to be 2nd
Largest Supercomputer in
the world



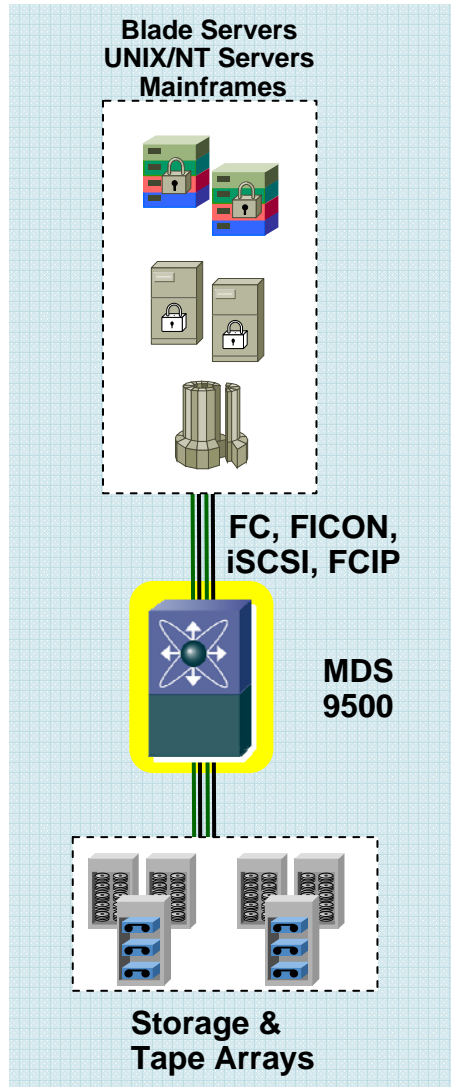
STORAGE NETWORKS FIBRE CHANNEL SWITCHING



Storage Consolidation

Lower TCO, Improved Flexibility, Resilience

Cisco.com



Mission-Critical Availability

- Fully-redundant platform (MDS 9500)
- Non-disruptive software upgrades
- Stateful switchovers
- Fault-Isolation with VSANs

Scalability and Density

- 1.44 Tbps switching bandwidth
- 224 FC ports in 14RU (MDS 9509)
- 10G Ready
- VSANs and Inter-VSAN Routing (IVR)

Troubleshooting

- Embedded FC Analyzer
- FC Ping and FC Traceroute
- Remote SPAN (RSPAN)
- Hot-spot analysis

Traffic Management

- QoS
- Fibre Channel Congestion Control (FCC)
- Hardware-based load balancing across ISLs
- VSAN Trunking

Heterogeneous Fabrics

- Interoperability with legacy switches
- Inter-VSAN Routing (IVR)
- Native-mode interop. with legacy switches

Multi-protocol Connectivity

- Integrated FC, FCIP, FICON, iSCSI
- Common set of features
- Unified provisioning and management

Cisco MDS 9000 Family



MDS 9216i



MDS 9140/MDS 9120



MDS 9216



MDS 9506



MDS 9509

MDS 9000 Modules



Supervisor



16 and 32-port
Fibre Channel



14 FC/2 IP



4-and 8-port IP
iSCSI + FCIP



SSM

Mgmt

CISCO FABRIC MANAGER

OS

Cisco SAN-OS

Common Architecture – *Ease-of-Migration and Investment Protection*

Cisco.com



Generation-1

- Architectural support for up to 256 ports
- Max planned system density of 240 ports
- 1/2Gbps FC interfaces

*Some feature limitations in mixed Current-Gen/New-Gen configurations

Generation-2

- Architectural support for up to 1,024 ports
- Max planned system density of 528 ports
- 1/2/4G, 10G FC interfaces

Multilayer Storage Networking

Reduced TCO via Intelligent Network Capabilities

Cisco.com

Modular Platforms with Industry-leading Performance

Maximize Investment Protection and Scalability

High Port Density with Connectivity Options

Lower Cost-per-Port

Multiprotocol Integration

Supports Tiered Storage Architectures and Enables Data Protection

Virtual SAN Technology

Maximize Storage and Server Consolidation

Advanced Traffic Engineering

Optimize Use of Network Resources

Intelligent Fabric Applications

Centralize Management and Improve Resource Utilization

Integrated Diagnostic Capabilities

Higher Availability through Quicker Problem Resolution

Cisco MDS 9000: The Open Platform for Intelligent Fabric Applications

Cisco.com

32-port Fibre Channel Module

MDS 9200 Fabric Switches
MDS 9500 Directors

Scalability and High Performance

Separate Control Plane Processor
Multiple Data Path Processors
ASIC-Based Innovation

Open Platform

Industry Standards-Based (FAIS API)
SANTap Protocol Support

Hosts Intelligent Applications

Volume Management
Point in Time Copy (Snapshot)
Replication / Data Migration
Server-less Backup

**Storage
Services
Module**



Intelligent Fabric Applications:

Network-Hosted
Network-Assisted
Network-Accelerated






Building Virtual Fabrics



Costly Practice: Building Out SAN Islands

Cisco.com

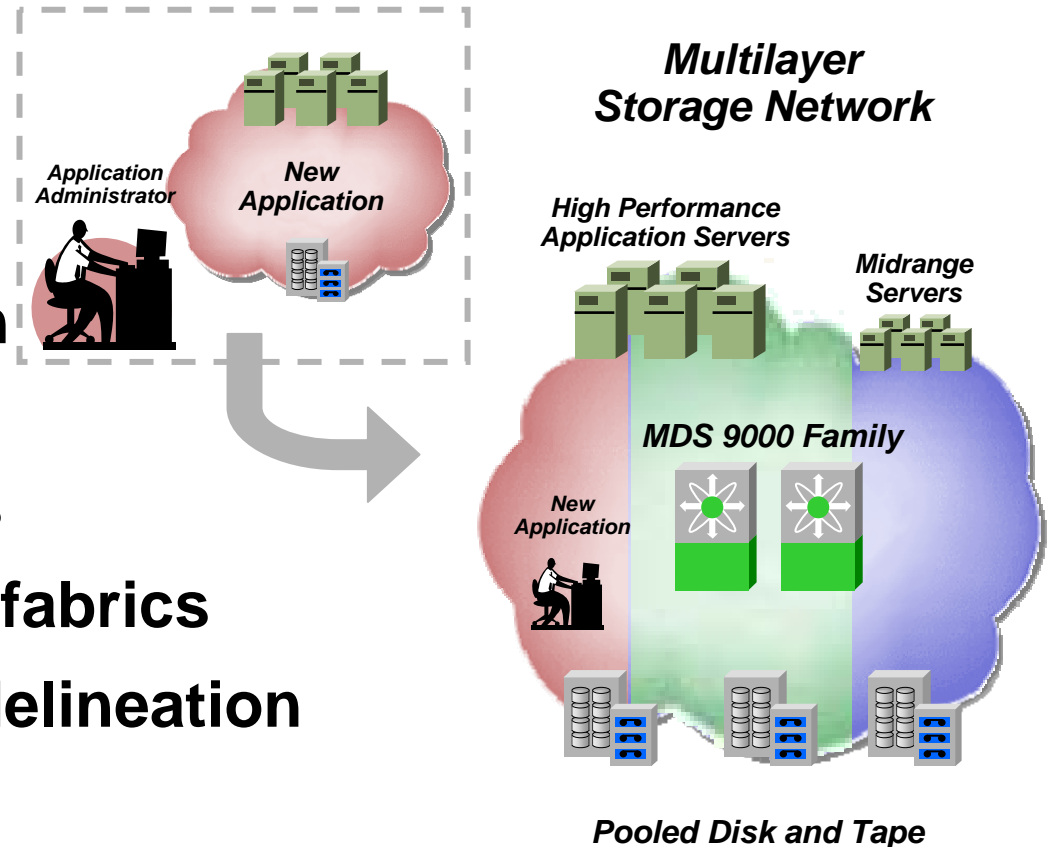
- **SAN islands are built to address several technical and non-technical issues**
 - Maintain isolation from fabric events or configuration errors**
 - Provide isolated management of island infrastructures**
 - Driven by bad experiences of large multi-switch fabrics**
- **However . . . Islands add cost and complexity**
 - Extra ports (\$\$) due to dedicated ISLs**
 - Extra FTEs (\$\$) to manage extra switches**
 - Extra time (\$\$) to provision new islands**
 - Extra time (\$\$) to expand existing islands**
 - Cannot move ports from one island to another**



Fabric Consolidation - Cisco Solution Foundation

Cisco.com



- Provide cost effective solution retaining benefits of SAN islands
- Flexible provisioning with port-based granularity
- Multi-protocol extensions
- Central administration of fabrics
- Secure hardware-based delineation between virtual fabrics
- Fabrics are added and scaled dynamically
- Enabled through Cisco's VSAN technology



Virtualizing the Fabric – The Full Solution

Cisco.com


- To build a cost saving fabric virtualization solution, 7 key services are required:
 - Virtual Fabric Attachment – assign virtual fabric membership at the port level
 - Multiprotocol Extensions – extend virtual fabric service to iSCSI, FCIP, FICON, etc.
 - Virtual Fabric Services – create fabric services per virtual fabric (routing, zones, RSCNs, QoS, etc.)
 - Virtual Fabric Diagnostics – troubleshoot per virtual fabric problems
 - Virtual Fabric Security – define separate security policies per virtual fabric
 - Virtual Fabric Management – map and manage virtual fabrics independently
 - Inter-Fabric Routing – provide connectivity across virtual fabrics – without merging the fabrics



Full Service End-to-End Virtual Fabric Implementation

Virtual SANs Raise Fabric Availability


- Virtual SANs sectionalize fabric to increase availability
- High availability - all fabric services are replicated and maintained per VSAN (eg. name service, etc)
- Fabric events are isolated per VSAN
 - Misbehaving HBA or controller
 - Fabric rebuild event
 - Zone set change
- Fabric recovery from a disruptive event is 'per VSAN' resulting in faster reconvergence – smaller scope



Faults are constrained to the extents of the VSAN and only affect devices within the VSAN

VSANs Allow Sharing of DR Facilities


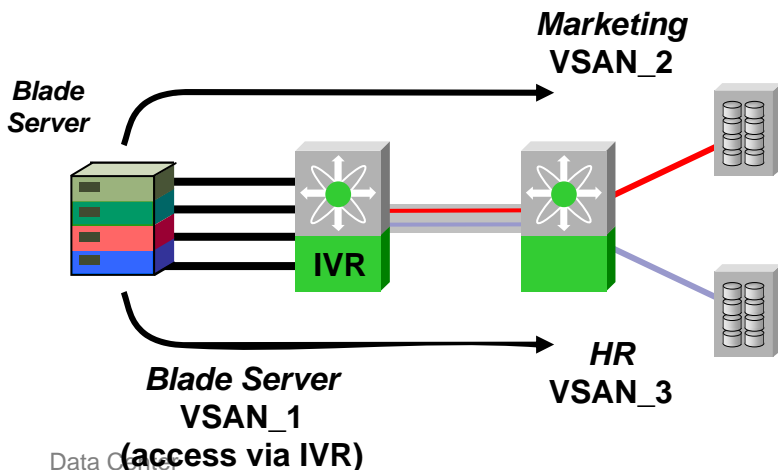
- Virtual SANs can be carried between data centers over various transport facilities
- Cost savings through consolidation of DR facilities while still maintaining isolation
- Various wide and metro area facilities can be used securely: FCIP (PoS, ATM, Metro Ethernet), SONET, DWDM, or CDWM, etc.
- Cisco MDS 9000 family provides traffic statistics per VSAN for chargeback possibility
- Full fabric discovery per-VSAN through Cisco Fabric Manager



Inter-VSAN Routing (IVR)

Sharing Resources Across VSANs

- Allows sharing of centralized resources such as tape and disks across VSANs – without merging separate fabrics (VSANs)
 - Provides high fabric resiliency and VSAN-based manageability
 - Distributed, scaleable, and highly resilient architecture
- Transparent to third-party switches
- Enables blade-per-VSAN architecture for blade servers

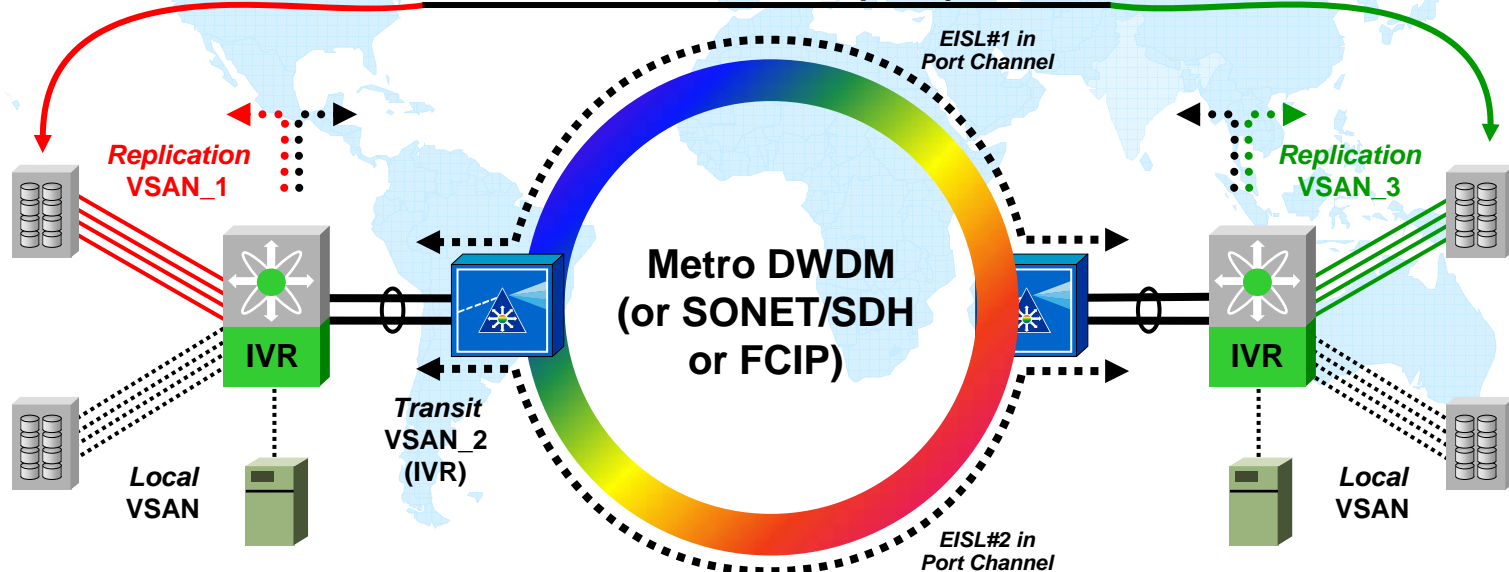


Inter-VSAN Routing (IVR)

Resilient SAN Extension Solutions

- Minimize the impact of change in fabric services across geographically distributed sites
- Limit fabric control traffic such as RSCNs and Build/Reconfigure Fabric (BF/RCF) to local VSANs
- Flexible connectivity with the highest availability
- Works with any transport service (FC, SONET, DWDM/CWDM, FCIP)

Inter-VSAN Connection with Completely Isolated Fabrics




VSAN Based Roles - Just Like SAN Islands

- Enables deployment of VSANs that fit existing operational models
- Network admin configures all platform-specific capabilities and VSAN allocations
- VSAN-admin(s) configure and manage their own VSANs
- The existing “role” definition is enhanced to include VSAN(s)



Network Administrator
VSAN Provisioning



Cisco Fabric Manager

Simplifies Management of Multiple Switches and Fabrics

- Switch-embedded Java Application

- Installed and updated automatically by Java Web Start

- Runs on Windows, Solaris, Linux workstations or laptops

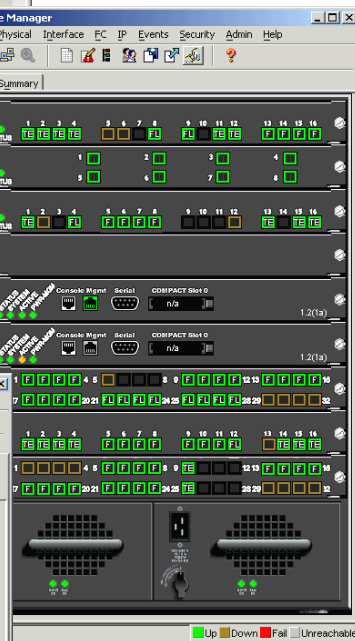
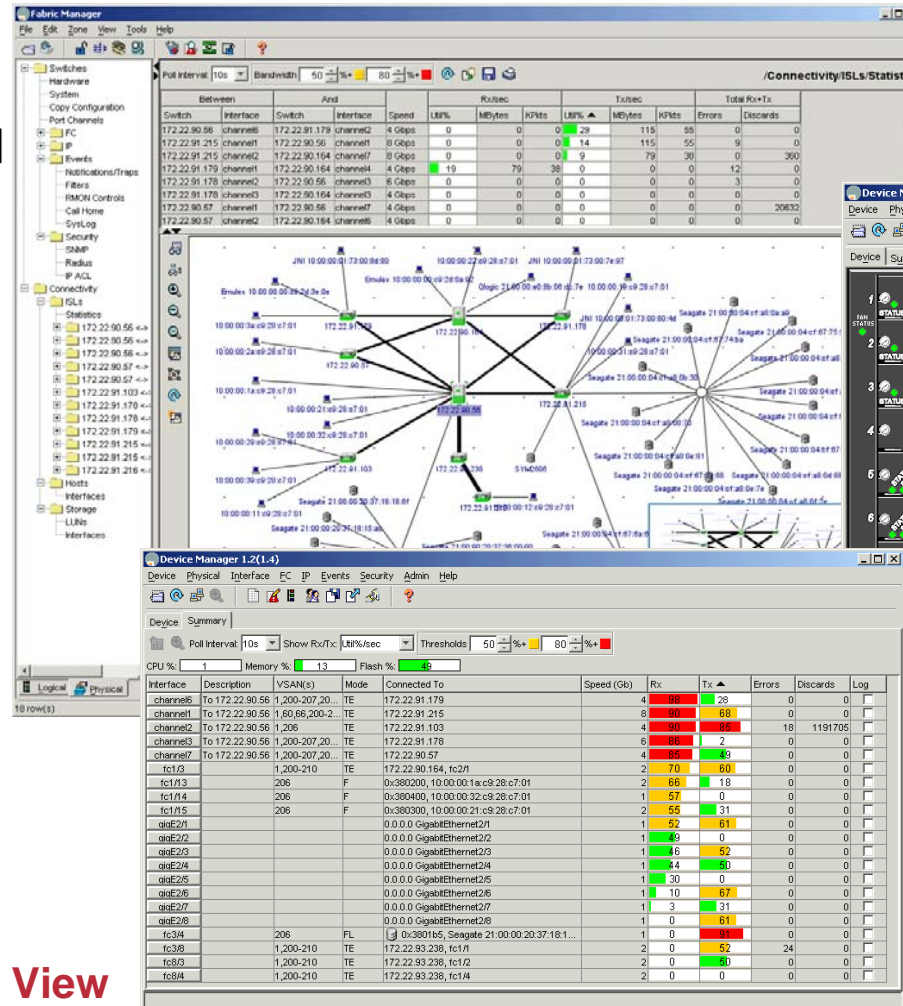
- SNMPv3 for secure management

- Discovers FC fabric and visualizes network topology, VSANs, Zones

- Enables rapid multi-switch configuration and analysis

Summary View

Fabric View



Device View

Cisco Performance Manager Reports

Fabric 172.22.37.55, Vsan 1 - Historical Performance - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Fabric 172.22.37.55, Vsan 1 - Historical Performance

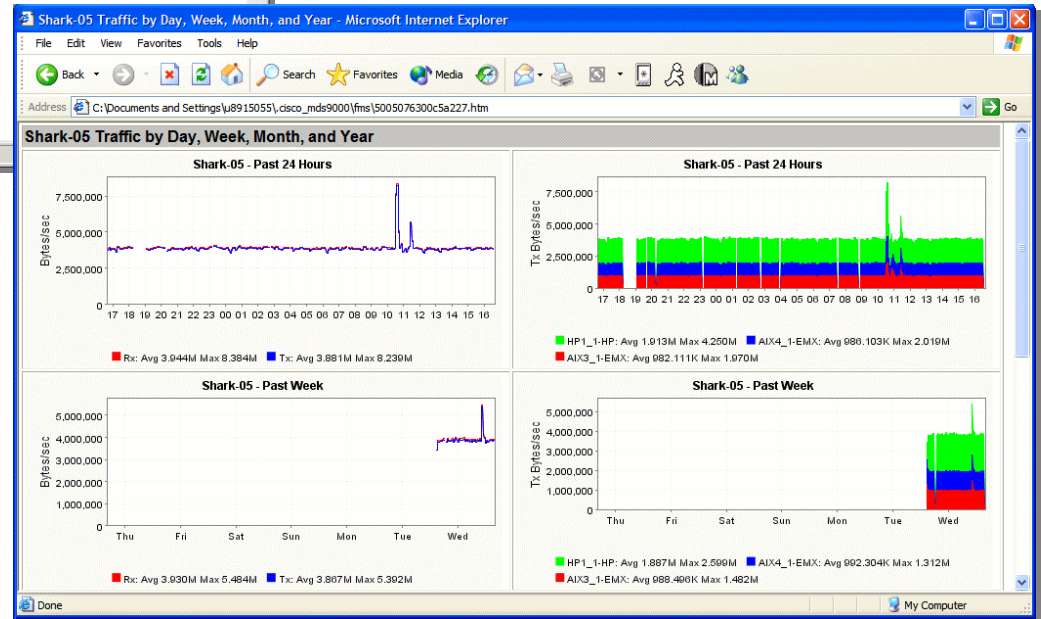
Hosts		Average/sec (24hr)		Max/sec (24hr)		ISLs		Average/sec (24hr)		Max/sec (24hr)	
Top 10	Rx:Bytes	Tx:Bytes	Rx:Bytes	Tx:Bytes	Top 10	Rx:Bytes	Tx:Bytes	Rx:Bytes	Tx:Bytes	Rx:Bytes	Tx:Bytes
iSeries-Disk 2-EMX	29.592K	30.890M	30.293K	31.481M	172.22.37.55 fc1/0 <-> 172.22.37.59 fc1/5	22.653M	36.542M	25.221M	39.505M		
W2k5 2-QLA	11.563M	18.651M	12.119M	19.489M	172.22.37.55 fc1/10 <-> 172.22.37.59 fc1/8	22.697M	36.475M	25.179M	39.471M		
W2k5 1-QLA	11.053M	19.155M	12.942M	21.732M	172.22.37.58 fc1/1 <-> 172.22.37.55 fc1/5	16.448M	12.366M	19.704M	15.642M		
iSeries-Disk 1-EMX	25.459K	26.557M	27.118K	28.420M	172.22.37.58 fc1/2 <-> 172.22.37.55 fc1/8	16.282M	12.133M	19.565M	15.327M		
W2k4 1-QLA	6.972M	11.155M	11.206M	18.003M	172.22.37.58 fc1/1 <-> 172.22.37.55 fc1/7	15.897M	11.961M	19.138M	15.161M		
W2k4 2-QLA	6.890M	11.034M	7.032M	11.373M	172.22.37.60 fc1/1 <-> 172.22.37.59 fc1/7	15.974M	4.202M	16.673M	4.389M		
W2k1 1-EMX	4.407M	9.462M	6.297M	13.345M	172.22.37.60 fc1/4 <-> 172.22.37.59 fc1/4	15.917M	4.210M	16.530M	4.413M		
W2k1 2-EMX	4.375M	9.349M	6.214M	13.095M	172.22.37.60 fc1/2 <-> 172.22.37.59 fc1/8	15.984M	3.966M	16.555M	4.174M		
SUN5 1-JNI	6.358M	6.408M	6.575M	6.626M	172.22.37.60 fc1/3 <-> 172.22.37.59 fc1/3	15.941M	3.965M	16.590M	4.159M		
AIX2 3-EMX	4.020M	4.226M	4.375M	4.568M	172.22.37.55 fc1/10 <-> 172.22.37.59 fc1/2	1.515M	17.023M	1.682M	17.355M		

Storage		Average/sec (24hr)		Max/sec (24hr)		Storage->Host Flows		Average/sec (24hr)		Max/sec (24hr)	
Top 10	Rx:Bytes	Tx:Bytes	Rx:Bytes	Tx:Bytes	Top 10	Rx:Bytes	Tx:Bytes	Rx:Bytes	Tx:Bytes	Rx:Bytes	Tx:Bytes
FAS1700-B2	11.448M	11.425M	11.882M	11.813M	FAS1700-A2->W2k5 2-QLA		7.024M		7.341M		
FAS1700-B1	11.760M	8.353M	13.123M	9.410M	FAS1700-A1->W2k5 1-QLA		6.705M		7.688M		
Shark08	9.477M	9.862M	13.409M	18.206M	Shark08->W2k3 1-QLA		5.445M		11.834M		
Shark11	9.367M	9.797M	13.137M	15.400M	Shark11->W2k3 2-QLA		5.428M		10.600M		
FAS1700-A1	11.617M	6.846M	14.161M	12.235M	FAS1700-B2->W2k5 2-QLA		4.527M		4.922M		
FAS1700-A2	11.175M	7.026M	11.780M	7.457M	Shark08->W2k1 1-EMX						
Shark02	6.147M	5.542M	6.862M	6.200M	Shark11->W2k1 2-EMX						
Shark07	4.056M	4.041M	5.728M	5.685M	FAS1700-B1->W2k5 1-QLA						
Shark10	4.035M	4.022M	5.740M	5.736M	FAS1700-B1->AIX2 1-EMX						
Shark05	3.944M	3.881M	8.384M	8.239M	FAS1700-B2->HP2 2-HP						

Top 10 SAN "Talkers"


- Top devices (hosts, storage, ISLs, flows) based on network utilization

Long term trend view of individual host, storage, ISL or flow utilization over one Day, Week, Month, Year (DWMY)



Cisco Traffic Analyzer - Flow Statistics

- Unprecedented analysis capability from a fabric
- Faster resolution of problems as a result of better 'eyes'
- LUN-level flow statistics between any source-dest pair of devices
- Includes full SCSI performance data - IOPS, response times, etc.
- SCSI session status and mgmt statistics (check conditions, etc)



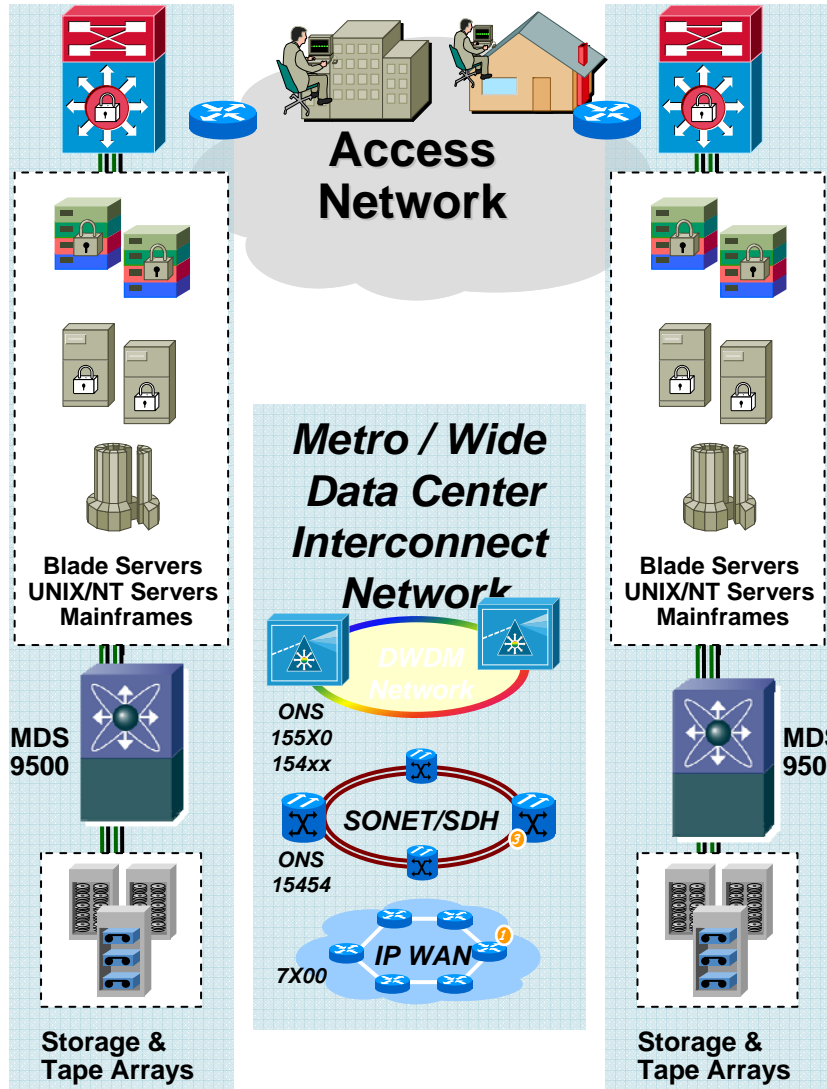
SCSI Sessions (Bytes)																	
VSAN	Initiator	Target	LUN	Total Bytes		Data Bytes			Read Size		Write Size		Xfer Rdy Size		IOPS		
				Sent	Rcvd	Read	Write	Other	Min	Max	Min	Max	Min	Max	Min	Max	
				2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	9	23.7 KB	134.3 KB	111.0 KB	0	3.0 KB	1	2	0	0	0
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	1	23.7 KB	132.1 KB	109.5 KB	0	2.9 KB	1	2	0	0	0	0	5.0	108.0	
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	2	23.7 KB	133.1 KB	110.0 KB	0	3.0 KB	1	2	0	0	0	0	5.0	146.0	
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	3	23.7 KB	133.2 KB	110.0 KB	0	3.0 KB	1	2	0	0	0	0	5.0	108.0	
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	4	23.7 KB	133.6 KB	110.5 KB	0	3.0 KB	1	2	0	0	0	0	5.0	146.0	
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	5	23.7 KB	133.2 KB	110.0 KB	0	3.0 KB	1	2	0	0	0	0	5.0	121.0	
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	6	23.7 KB	133.7 KB	110.5 KB	0	3.0 KB	1	2	0	0	0	0	6.0	164.0	
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	7	23.7 KB	133.2 KB	110.0 KB	0	3.0 KB	1	2	0	0	0	0	8.0	132.0	
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	8	23.7 KB	133.7 KB	110.5 KB	0	3.0 KB	1	2	0	0	0	0	8.0	76.0	
																185.0	

SCSI Sessions (Times)													
VSAN	Initiator	Target	LUN	Cmd-Status RTT		Cmd-XFR_RDY RTT		Cmd-First_Data RTT(Read)		Cmd-First_Data RTT(Write)		Duration(secs)	
				Min	Max	Min	Max	Min	Max	Min	Max	Min	Max
				2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	9	0.2 ms	0.2 ms			0.2 ms	0.2 ms
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	1	0.2 ms	0.2 ms			0.2 ms	0.2 ms				1674
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	2	0.2 ms	0.2 ms			0.2 ms	0.2 ms				1674
2	10:00:00:00:c9:2d:43:e2	50:00:60:e8:02:eb:50:10	3	0.2 ms	0.2 ms			0.2 ms	0.2 ms				1674

Comprehensive Business Continuance

Addressing Risk and Compliance

Cisco.com



Broad Application Support

- Asynch / Synch Replication, Backup, Point in Time Copy
- Multi-vendor Support: EMC, HP, HDS, IBM, Appliances
- 3rd Party Appliance Support: SANTap
- Network Assisted Serverless Backup

Optimized SAN Extension

- Multi-protocol Transport: FCIP, DWDM, Sonet/SDH
- Distance / Application Optimization: Write and Tape Acceleration
- Security: FCIP Encryption and FC-SP Auth
- WAN Utilization: Compression and Large B2B Credits
- Availability: VSANs and Inter VSAN Routing

Continuous User Access

- Global Site Selector
- VPN – IPsec, SSL, MPLS
- Optimized Exit Routing

APPLICATION OPTIMIZATION

WEB SERVICE ACCELERATION

FILE SERVICE ACCELERATION



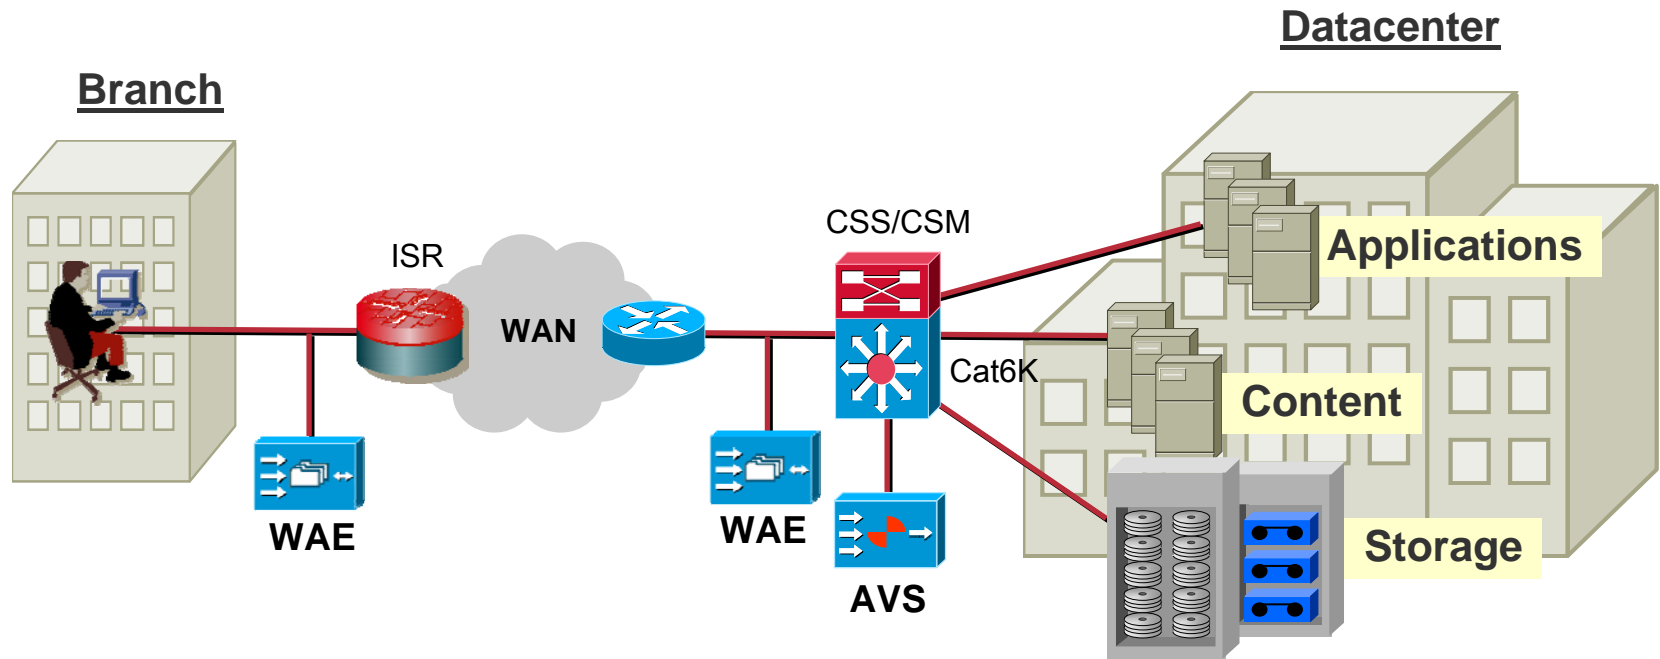
ADBU (CNBU-CSBU) Technology Strategy

Cisco.com

- **Create a foundation of compression technologies and TCP packet flow optimizations**
 - Optimize user access to all applications – **not just those with specific application adapters**
- **Strong Emphasis on Application Adapters**
 - Combine** application adapters from **ACNS and WAFS** into common, optimized foundation
 - Add further application adapters as needed for “problem” applications
- **Network Centric Architecture for Investment Protection**
 - Compliance with all standard network functions (QoS, NBAR, ACLs, IPsec, load-balancing, network management & monitoring)**
 - Integrate into existing network hardware platforms**

Cisco Application Delivery Solutions

Cisco.com



Branch/WAN

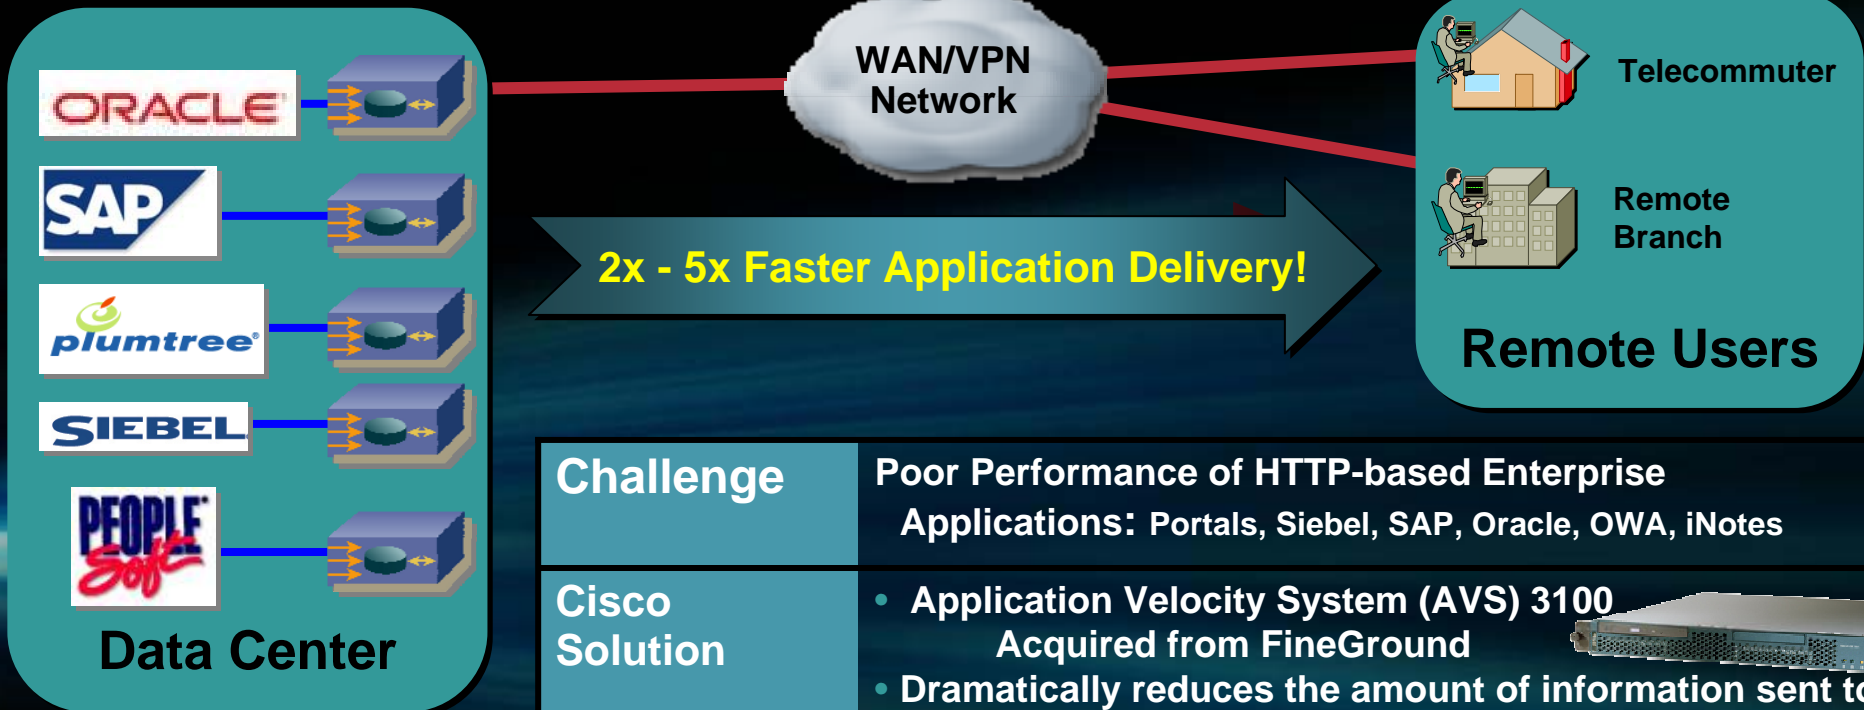
- Wide Area Application Engine (WAE)
- Integrated Services Router (ISR)
- IOS with NetFlow, NBAR, QoS, IP-SLAs...

Data Center

- Application Velocity System (AVS)
- Content Switches (CSS/CSM)
- Catalyst 6500 switches

Optimization of Web-Based Applications

Cisco.com



<p>Challenge</p>	<p>Poor Performance of HTTP-based Enterprise Applications: Portals, Siebel, SAP, Oracle, OWA, iNotes</p>
<p>Cisco Solution</p>	<ul style="list-style-type: none"> • Application Velocity System (AVS) 3100 Acquired from FineGround • Dramatically reduces the amount of information sent to the client from the central site. • Dynamically Caches, Transforms Content, Compresses Content, and Secures the Web With Full-Proxy Functionality. • Reduces latency and improves responsiveness by handling all redirections.
<p>Reference Customers</p>	<ul style="list-style-type: none"> •AIG, AXA, BMW, US Air Force, Royal Mail •American Airlines, Walgreen's, Whirlpool



Application Velocity System (AVS)

Cisco.com



Hardware: 1 U Appliance Server

OS: Red Hat Linux Advanced Server 3

Software Version: Performance Suite 8.0

Hardware Version: 4.0

AVS Delivers Dramatic Performance and Business Impact

Application	Software	FineGrout and Impact	Cost of Likely Alternative	Business Impact
Call center (High tech)	PeopleSoft	↓ 63% (↑270%)	\$4MM (2 new overseas data centers)	▶ Meet support goals with no additional staffing or costs
Purchasing (Manufacturing)	SAP	↓ 71% (↑350%)	\$5MM (Multiple overseas data centers)	▶ Increase procurement automation
Mortgage origination (Financial)	Custom J2EE: WebSphere	↓ 68% (↑300%)	\$2MM (reengineer apps and infrastructure)	▶ 30% more transactions across same infrastructure
Claims Management (Insurance)	Custom J2EE: WebSphere	↓ 55% (↑220%)	\$3MM (No reengineering)	▶ Support “zero-footprint” branch
B2B operations (Retail)	Plumtree	↓ 71% (↑350%)	\$500K annually (upgrade 650 sites)	▶ Move all costly paper-based processes online
CRM (Financial)	Siebel	↓ 66% (↑290%)	\$2.4MM annually (upgrade 200 U.S. locations)	▶ Immediate jump in CRM usage with Improved account retention

- Provides performance

Optimization,
Monitoring, and
Security

- for **web-based** applications.

Conventional Approaches

Network Team


- More bandwidth
- Point products
- Replicate data centers

Operations Team

- More servers
- More management tools
- Re-architect infrastructure

Applications Team




- More testing
- Rewrite applications
- Security patching




Expensive, complex, ineffective

Cisco AVS-3110

Delivery Functions




 Accelerate Best response time on existing infrastructure	 Optimize Minimize required network infrastructure	 Offload Maximize capacity of application infrastructure
---	--	--



**WAN Latency:
100 ms**

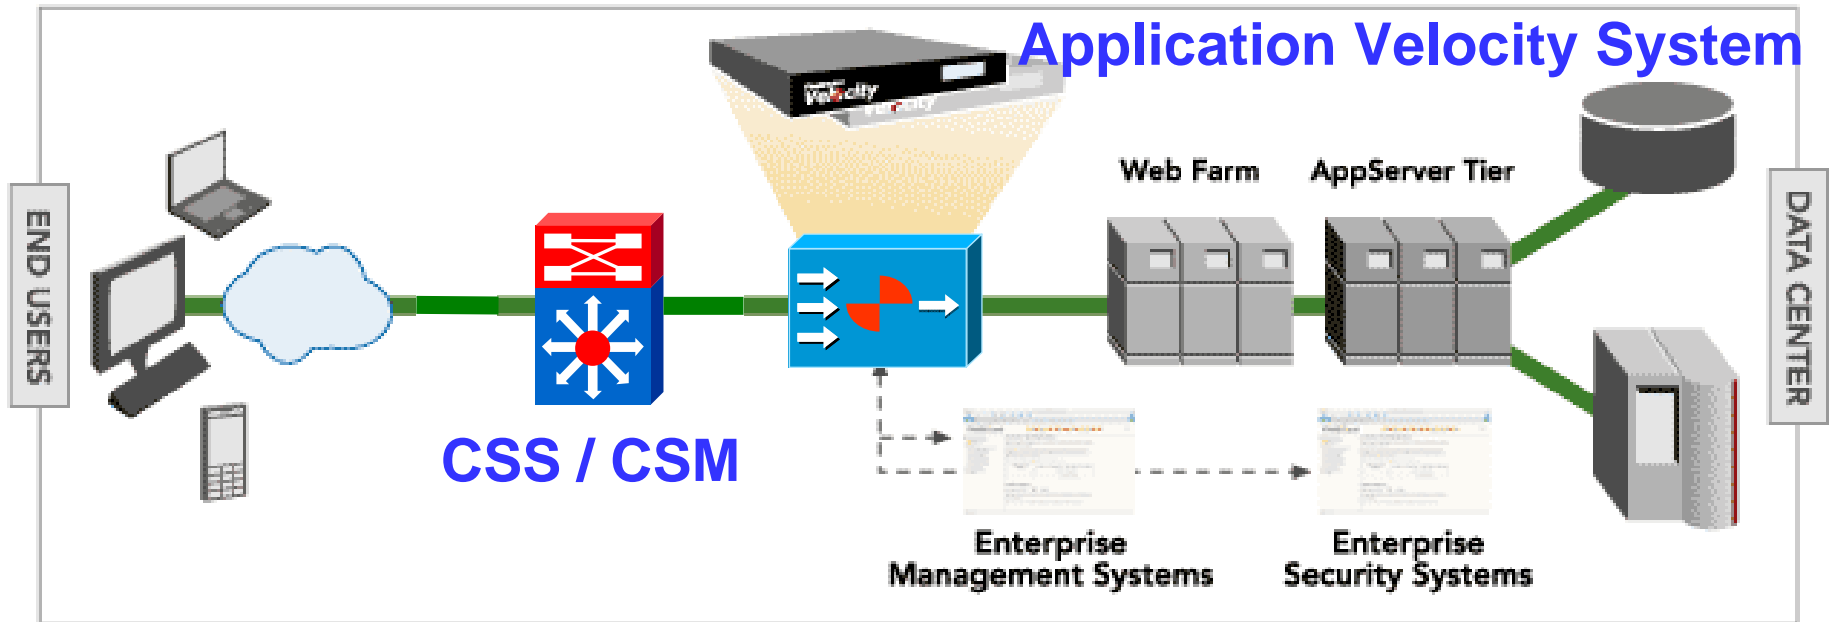


Service Functions

 Monitor Provide end-user quality of service metrics	 Secure Policy-based protection of app infrastructure	 Manage Management and exception handling
--	---	---

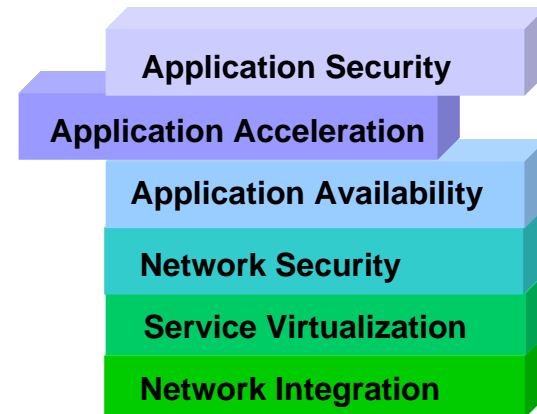
Cisco AVS-3110 Deployment Scenarios

Cisco.com









- AVS 3110 devices are deployed in two configurations:

- “Inline” using internal clustering for scalability and failover
- “Out of band” using Layer 4-7 SLB to manage infrastructure
 - Proven configurations available with Cisco CSS
 - Velocity appears as another web server to the SLB





Technology Advantage

Functional Areas	Basic Capabilities	AVS Capabilities (*= Patented)
 Accelerate Network Latency Management		<ul style="list-style-type: none"> ▪ Request aggregation / browser cache management* ▪ Browser TCP multiplexing* ▪ PDF download optimization ▪ Response redirection control*
 Optimize Bandwidth Reduction	<ul style="list-style-type: none"> ▪ Gzip/DEFLATE compression 	<ul style="list-style-type: none"> ▪ Delta encoding* ▪ Dynamic browser caching* ▪ Dynamic image optimization (JPG, GIF, PNG) ▪ Flexible processing rules
 Offload Server Efficiency	<ul style="list-style-type: none"> ▪ TCP connection multiplexing ▪ SSL offload and acceleration ▪ Static caching 	<ul style="list-style-type: none"> ▪ Configurable dynamic caching* ▪ Load-based caching* ▪ Lazy request evaluation* ▪ Single sign-on optimizations ▪ XML merging/transformation
 Monitor Application QoS	<ul style="list-style-type: none"> ▪ Logging ▪ System health checking 	<ul style="list-style-type: none"> ▪ End-to-end response time monitoring ▪ Business transactions capability ▪ First-line service triage
 Secure Protect Applications and Infrastructure	<ul style="list-style-type: none"> ▪ Rules-based protection 	<ul style="list-style-type: none"> ▪ Out-of-the-box Layer-7 protections ▪ Stateful Content inspection policies ▪ Comprehensive exception handling and monitoring
 Management/ Integration	<ul style="list-style-type: none"> ▪ SNMP access and control 	<ul style="list-style-type: none"> ▪ Application delivery dashboard ▪ Service-level integration with BMC, HP, etc.

Consolidation of Branch File Servers

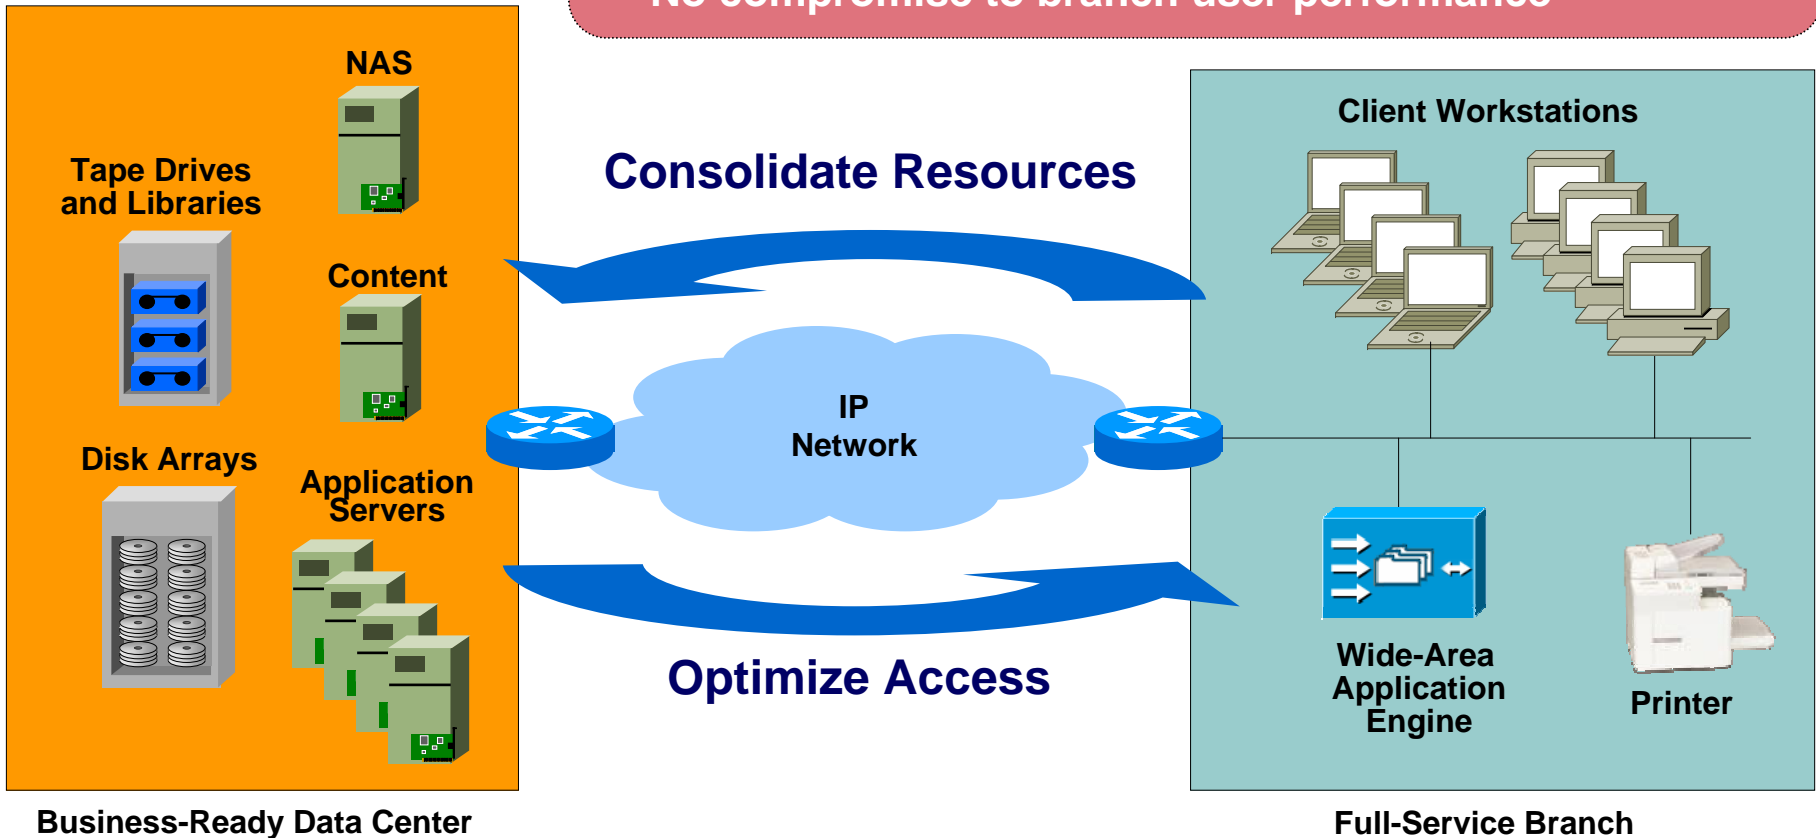
Cisco.com



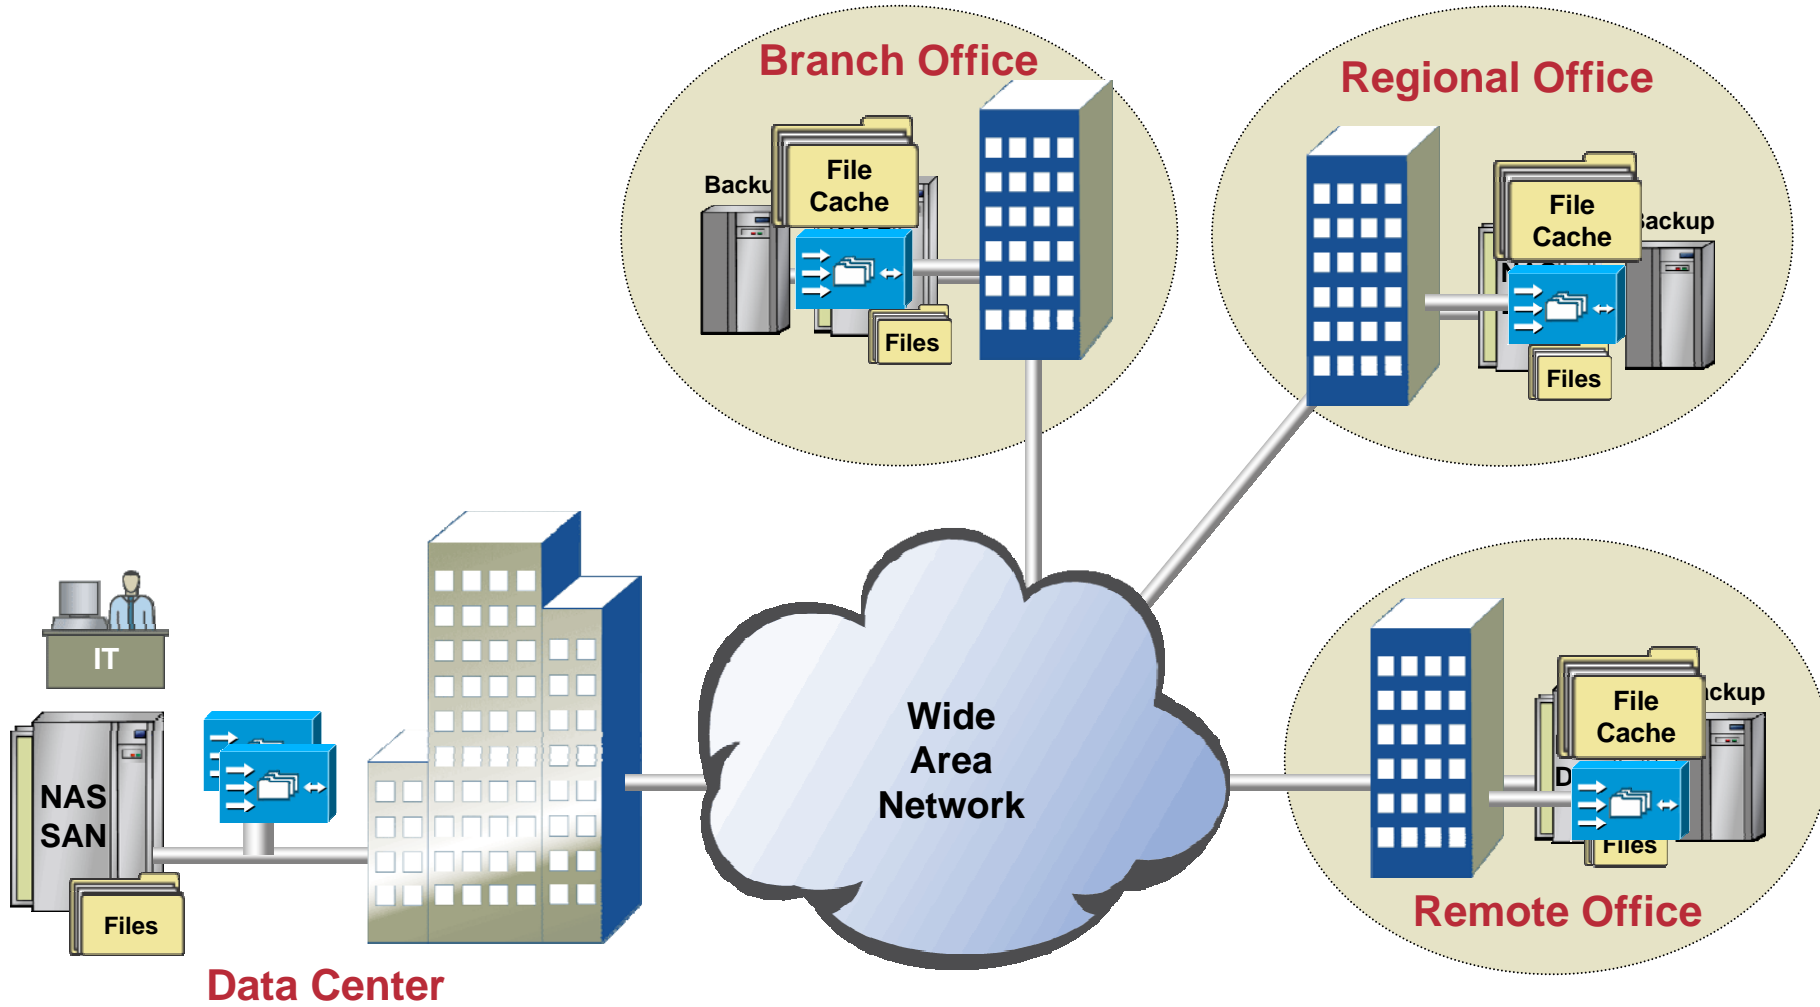
<p>Challenge</p>	<p>High Operational Cost of Managing Branch File and Print Servers</p>
<p>Products Used</p>	<ul style="list-style-type: none"> • Cisco Wide Area File Services • Core File Engine • Edge File Engine • Removes the need for Local File and Print Servers in the Branch • Brings branch office data back to the Data Center • Provides LAN-like file access over the WAN 
<p>Reference Customers</p>	<p>CSC, 3M, NCR, Deloitte, Morgan Stanley, DTS, GMAC, Del Monte Foods, Dimension Data</p>

The Branch Consolidation Vision

- Consolidation of application and service infrastructure
- Centralized management, deployment, and protection
- Single, protected copies of data
- No compromise to branch user performance



WAFS = Enabling Server Consolidation



Data Center

Branch Office

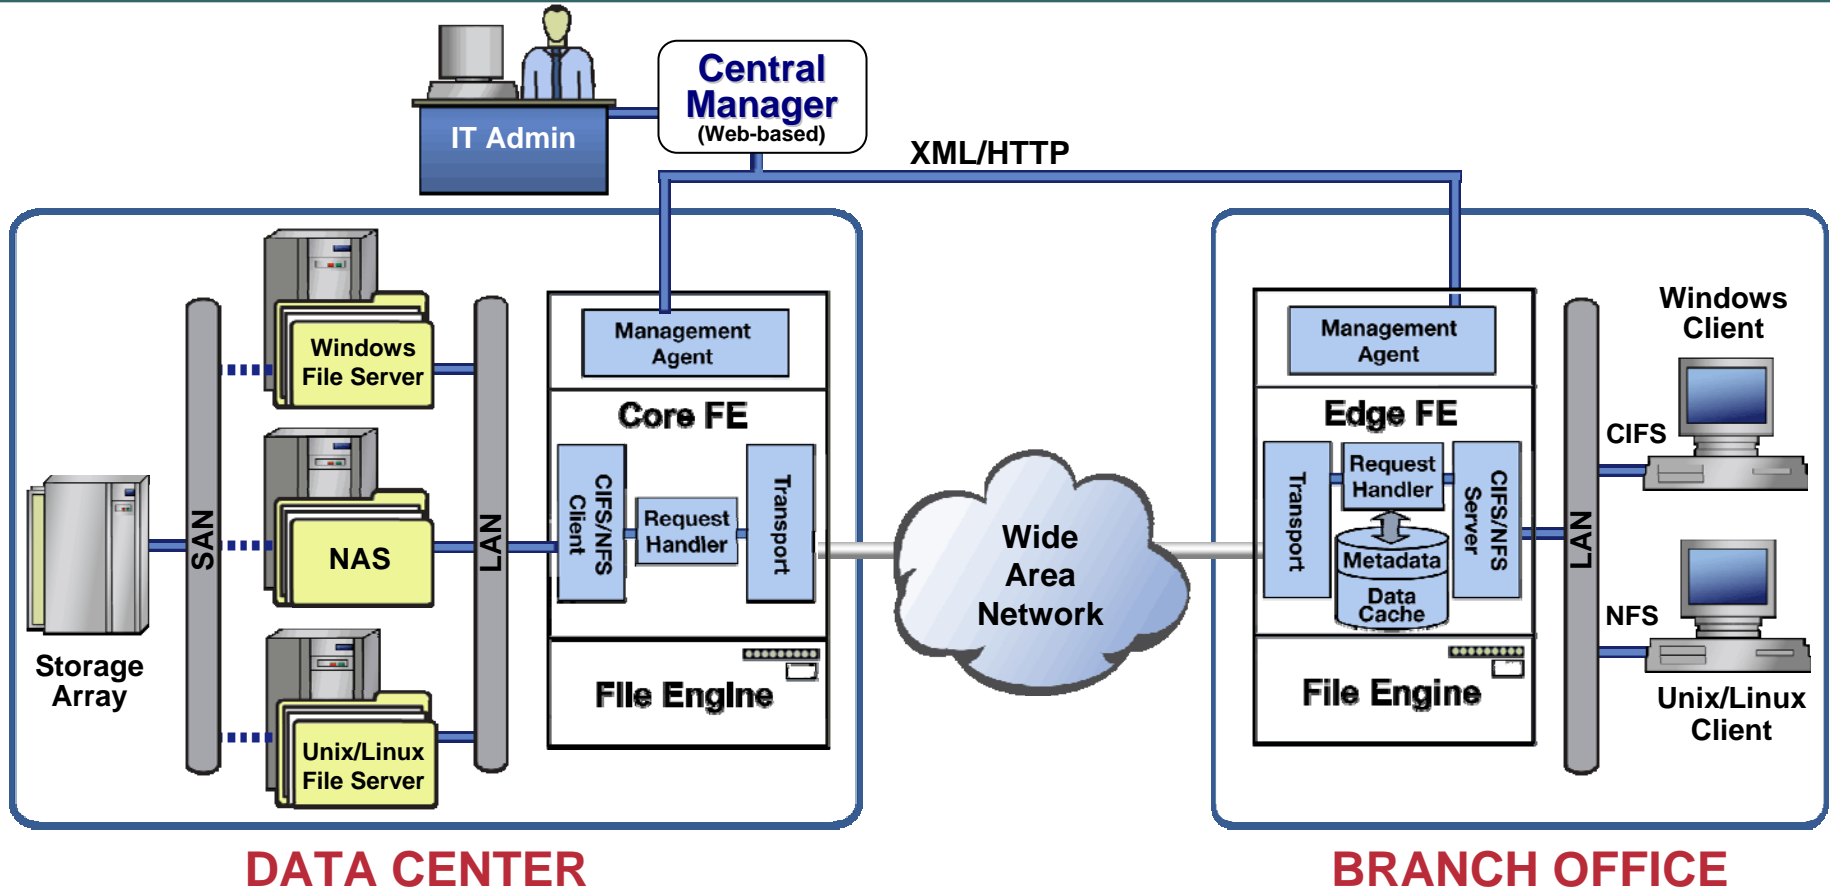
Regional Office

Wide Area Network


Remote Office

Centralized and Sprung

WAFS Architecture Overview




WAFS Performance Improvement




- Overall Microsoft Office performance benchmark
- Cisco WAFS technology provides near-LAN response time
- Orders of magnitude faster than native WAN

Microsoft Office Performance


Word - Time to Open




Word - Time to Save



Excel - Time to Open



Excel - Time to Save



Network Link – E1, 80ms latency

- **Cisco WAFS provides 4x to 9x faster performance vs. the WAN, and near LAN performance for typical operations on MS Office applications**

Today's ACNS & WAFS Solutions



ACNS

Web App.
Acceleration

- Internet/Intranet Portals
- Enterprise App. Acceleration
- Large File Transfer

Business
Video

- Corp Communications
- eLearning
- Live Broadcast + VoD

HTTP
HTTPS
FTP
TFTP

Web Content
Security

- Web Access Policy Enforcement
- Content filtering
- Virus scanning

MMS
RTSP

WAFS

Wide Area
File Services

- File & Print Server Consolidation
- Remote Branch Data Protection
- File Access & Collaboration
- Software Distribution

CIFS/NFS

Current Hardware Platform Wide Area Engines (WAE)

Platform	Hardware	Positioning
NM-CE	<ul style="list-style-type: none"> 500MHz Pentium III 512MB of memory 40GB or 80GB ATA 	<ul style="list-style-type: none"> Integrated branch services for small deployment
WAE-511	<ul style="list-style-type: none"> 2.8GHz Celeron up to 1GB memory 80GB to 500GB SATA 	<ul style="list-style-type: none"> Dedicated appliance for small-medium edge or core deployments
WAE-611	<ul style="list-style-type: none"> 3.0GHz Pentium IV 2GB of memory 288GB SCSI 	<ul style="list-style-type: none"> Dedicated appliance for medium-large edge or core deployments
WAE-7326	<ul style="list-style-type: none"> Dual 3.2GHz Xeon 4GB of memory 73GB to 1.8TB SCSI 	<ul style="list-style-type: none"> Dedicated appliance for enterprise edge or core deployments



Network Module



WAE-511



WAE-611



WAE-7326

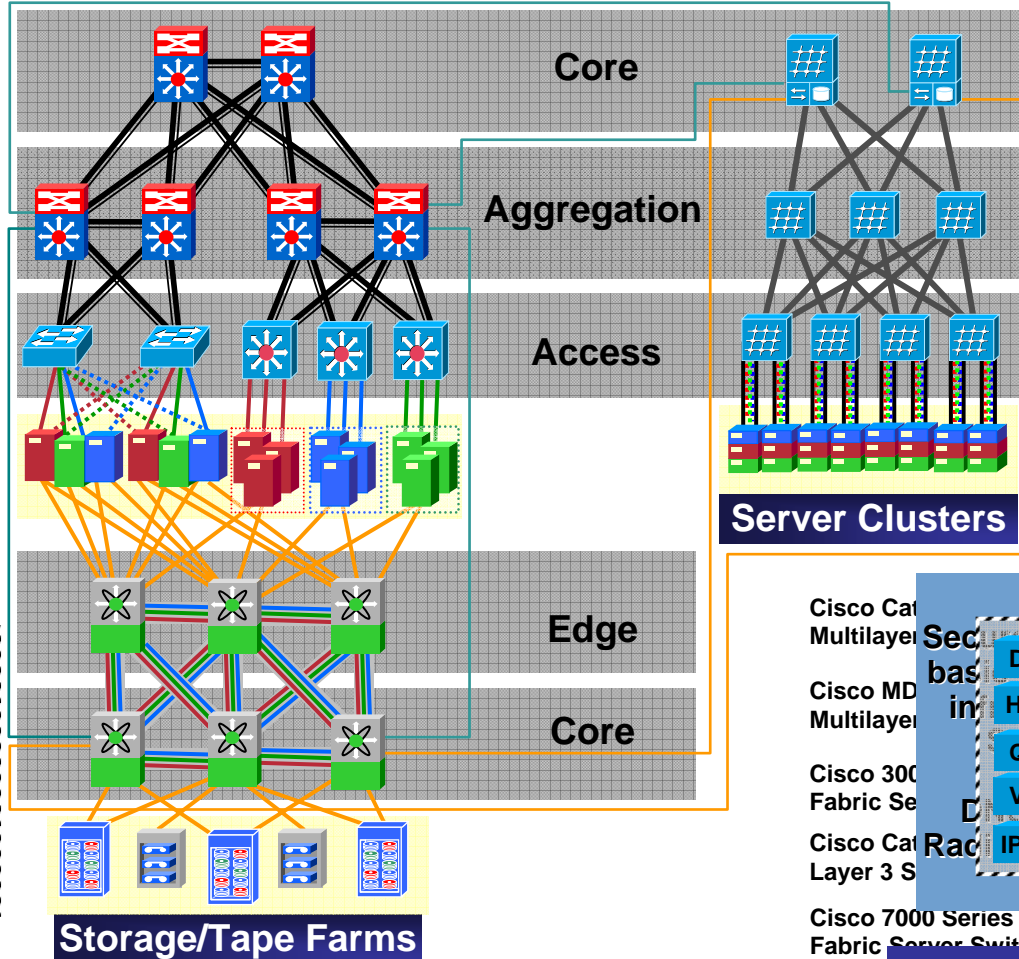
Cisco WAAS Summary

- **Server Consolidation**
 - Consolidate server, storage & backup infrastructure to the data center where it can be centrally managed
 - Provide local services
- **Application Delivery**
 - Deliver centralized applications (ERP, CRM, etc.) to the edge without performance compromise
 - Access any TCP-based data (file, web, video, messaging) across the WAN at LAN-type speeds
 - Seamlessly deliver content (software distribution, patches, etc.) to the edge
- **Network Transparency**
 - Solution integrated transparently within the current packet network

Data Center Architecture

Functional Layers and Services

Layers



- ▶ Firewall Services
- ▶ Intrusion Detection
- ▶ DOS Protection
- ▶ VPN Termination
- ▶ Anomaly Detection
- ▶ Network Analysis
- ▶ App Acceleration
- ▶ AON Analysis
- ▶ Server Balancing
- ▶ SSL Offloading
- ▶ File Caching
- ▶ Content Caching

Server Farms

- ▶ Virtual Fabrics (VSANs)
- ▶ Storage Virtualization
- ▶ Data Replication Svcs
- ▶ Fabric Routing Services
- ▶ Fabric Gateway Services

Storage/Tape Farms

- ▶ Server Virtualization
- ▶ Virtual I/O
- ▶ Compute Fabric Services
- ▶ Remote DMA Services
- ▶ Clustering Services
- ▶ Fabric Gateway Services

Application Networking Services

Server Clusters

Cisco Cat Multilayer Sec bas in
 Cisco MD Multilayer
 Cisco 300 Fabric Se
 Cisco Cat Rac Layer 3 S
 Cisco 7000 Series Fabric Server Switch
 Cisco Catalyst 10 Gigabit Ethernet
 Virtual Server Link

DOS	IGPs	IGPs
HSRP	VRRP	GLBP
QoS	IPMc	STP
VRF	VLANs	PVLANS
IPSec	MPLS	GRE

Security Services
 Services
 Services

Foundation Infrastructure Functions

CISCO SYSTEMS

