Tech update Datacenter

Hans Donnerborg

hans@cisco.com

Agenda

- **ACE**
- **VFrame Datacenter**
- **Fibre Channel update**
  - **NPV, Blades**
- **Ethernet switching update**
  - **Blades, Top of rack, End of rack**
  - **VSS, Fabric failover etc**
- **WAAS**

ACE

ACE Release Timelines – Near Term

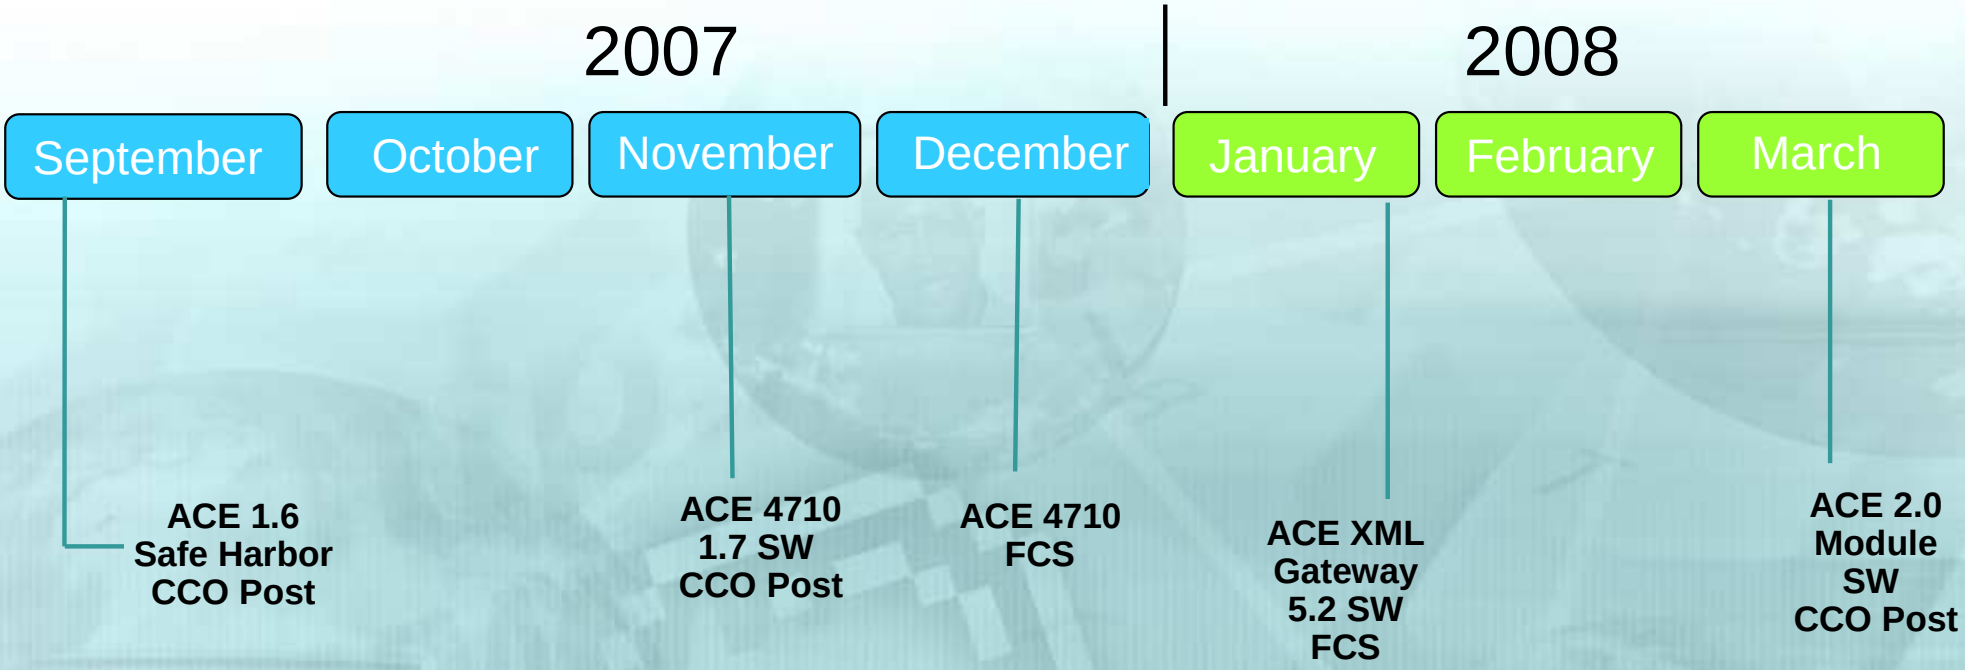

CSS Update

- CSS 11500 **is a critical part of Cisco Data Center Strategy and is marketed as a battle-tested, robust server load balancer and content switch**
- **The** CSS 11500 **product family is the most widely deployed Content Switch in the world with over 54,000 in service**
- **In particular,** CSS 11501 **is a terrific value at $10K~20K list price**
- **Cisco recommends** ACE **product family for new advanced services like virtualization, application acceleration and security**

Cisco ACE
Application Networking Manager (ANM) 1.2

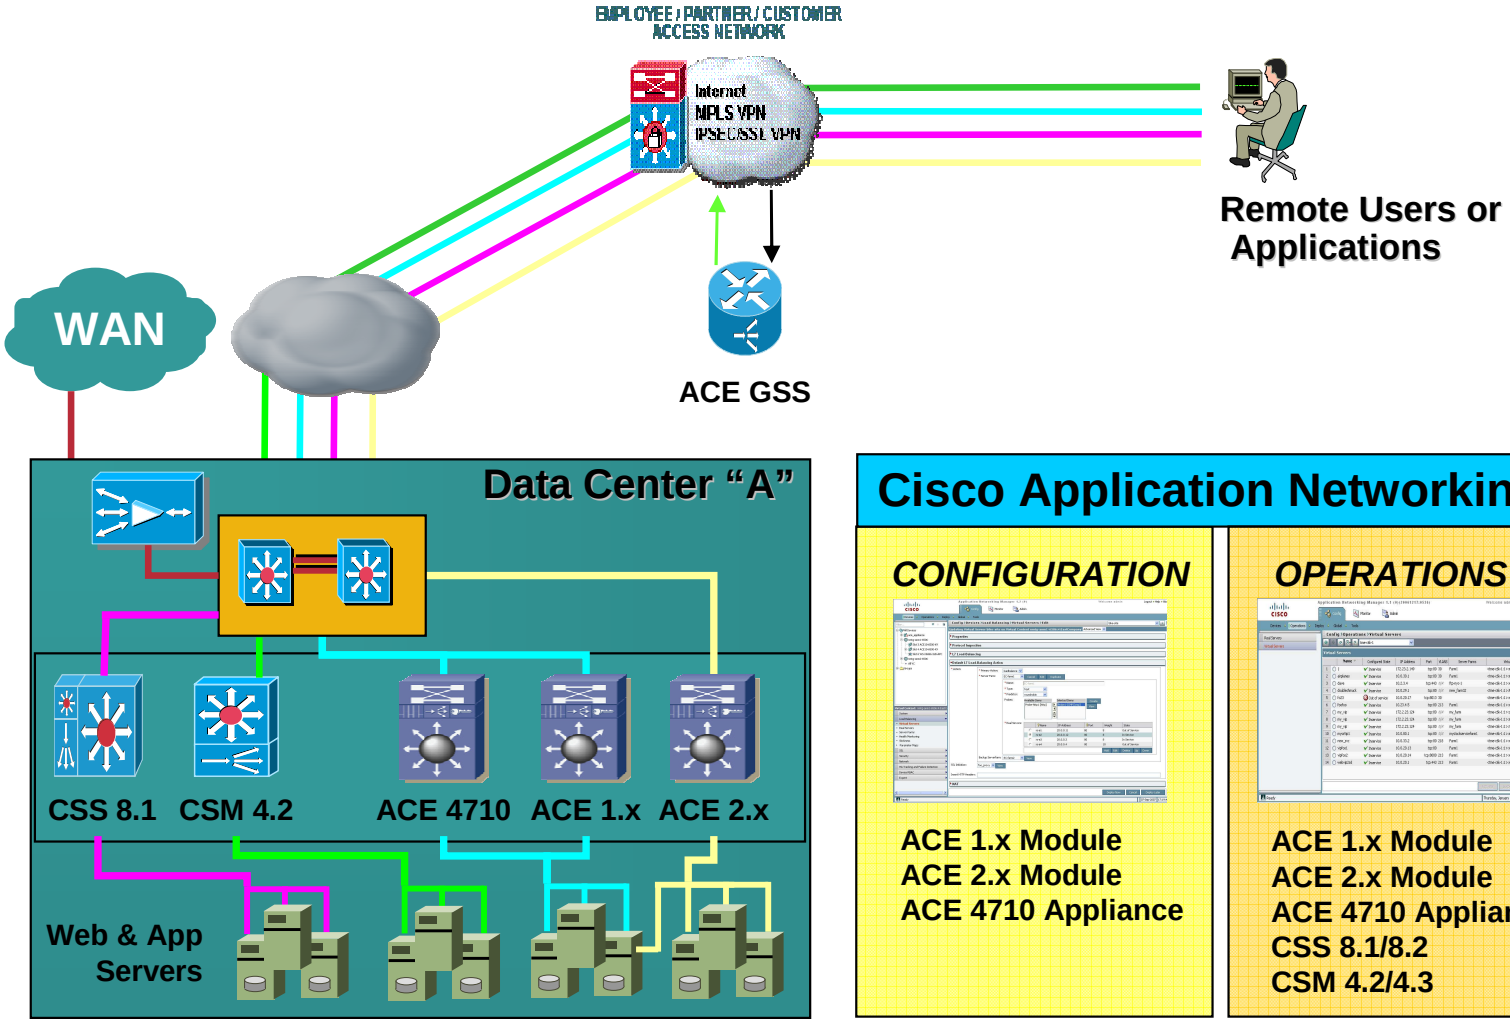

Cisco Application Networking Manager (ANM)

CONFIGURATION

ACE 1.x Module
ACE 2.x Module
ACE 4710 Appliance

OPERATIONS

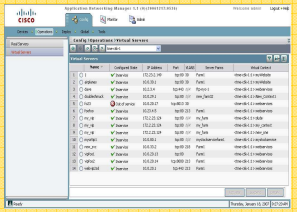

ACE 1.x Module
ACE 2.x Module
ACE 4710 Appliance
CSS 8.1/8.2
CSM 4.2/4.3

MONITORING

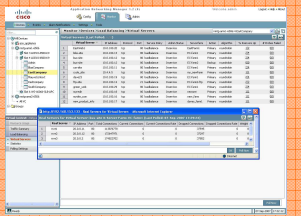

ACE 1.x Module
ACE 2.x Module
ACE 4710 Appliance
CSS 8.1/8.2
CSM 4.2/4.3

VFrame DC-ANM: Integration Methodology and Roadmap

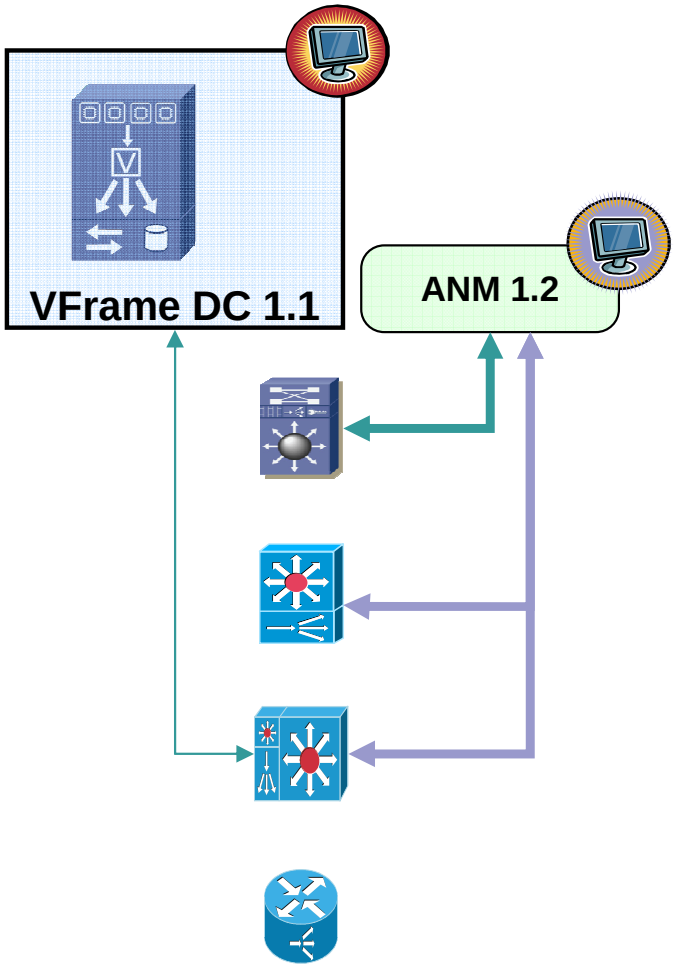

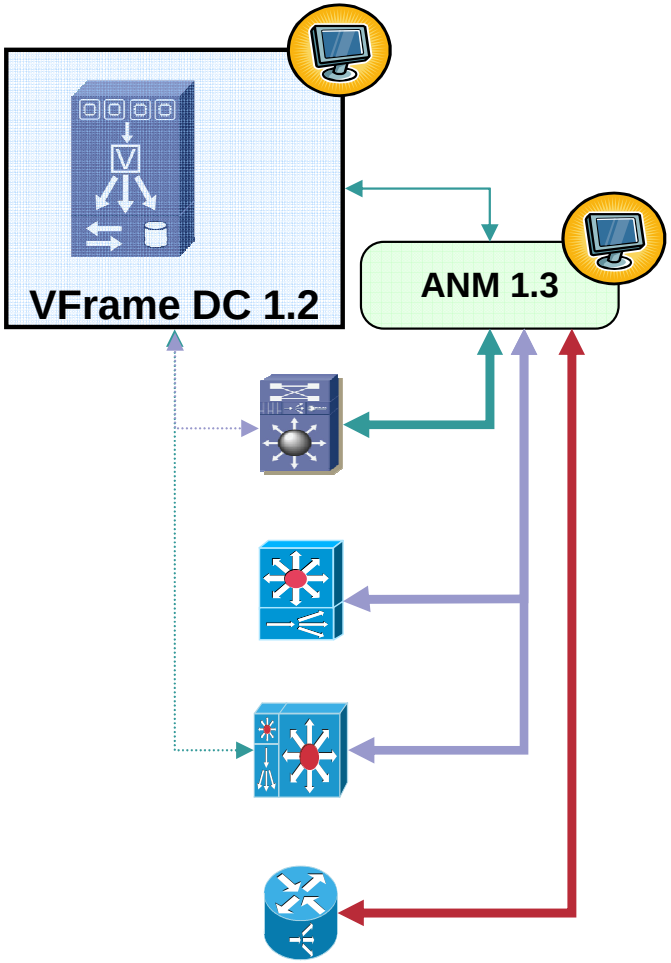

Roadmap

- End-to-end service automation
- Resource pool management
- Capacity & utilization optimization to user policy
- Additional device support
- Common look & feel per DCWG designs
- Automatic synchronization
- Unified RBAC
- Bundling
- Co-residency

- Direct interaction with devices
- VFrame relies on macro scripting

- Automation through APIs for coordinated services deployment of Layer 4-7
- Converging Look & Feel / Taxonomy

VFrame Data Center 1.1

Agenda

- **VFrame DC Status**
- **Feature Recap**
- **New Features & Release Schedule**
- **VMware integration**

VFrame DC Current Status

- **Launched at Networkers July 2007**
- **Shipping since beginning of August 20**
- **Currently at Release 1.1**
- **Current Support Matrix Summary:**

  Catalyst 6500

  MDS

  FWSM, CSM

  x86 Servers from HP, IBM , Dell (Modular addition for servers)

  NAS storage – NetApp

  FC Storage – EMC, IBM (Modular addition for storage arrays)

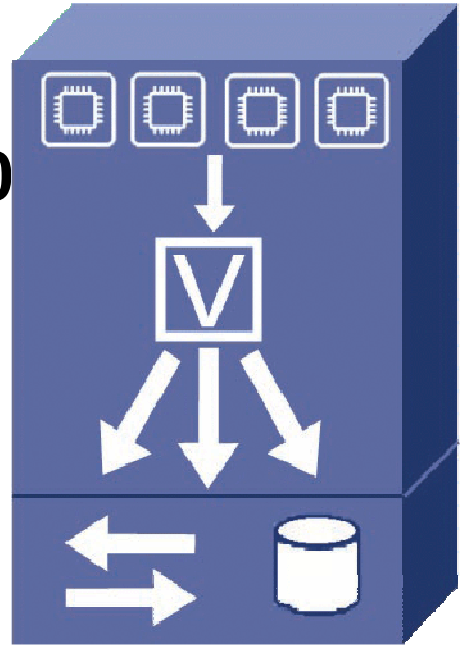

VFrame Data Center 1.1 Overview

- **VFrame Components**
  - –Appliances
  - –Java based GUI
  - –Web Services API
  - –Host Agent
  - –Extensibility Macros
- **VFrame Workflow**
  - Design
  - Discover
  - Deploy
  - Operate

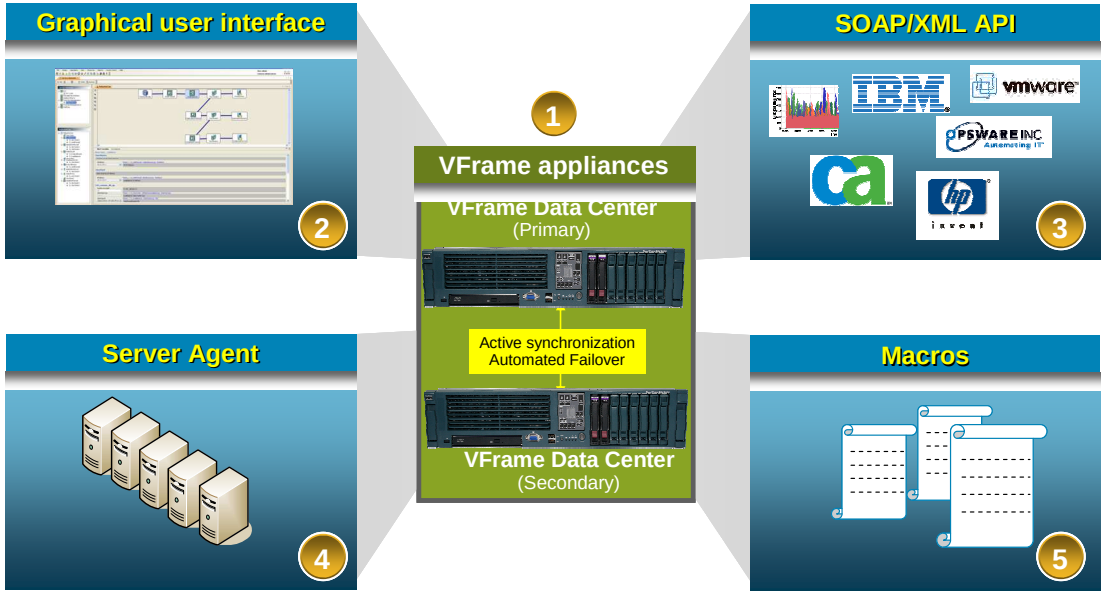

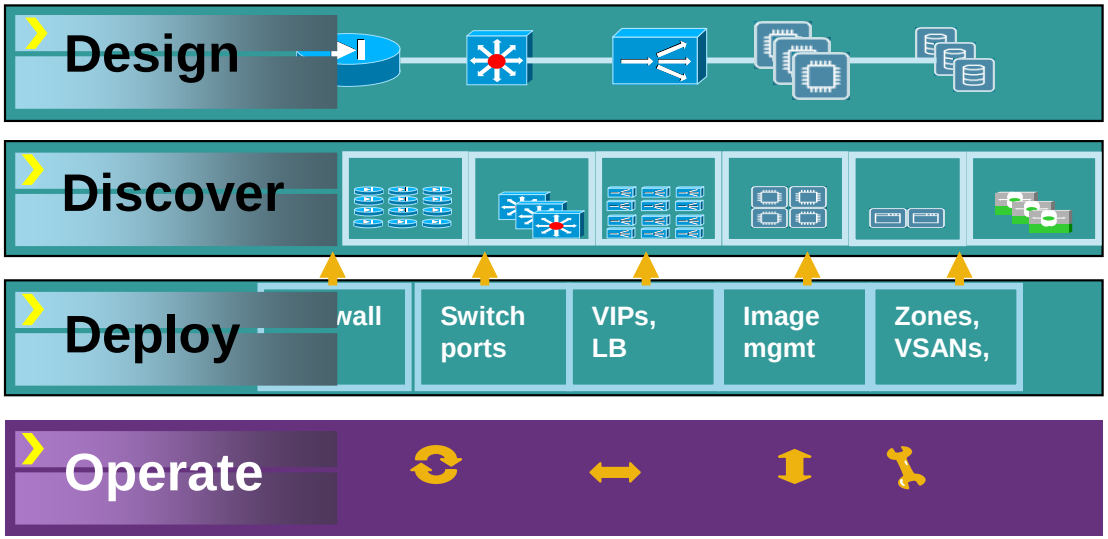

VFrame DC: Service Template Design GUI

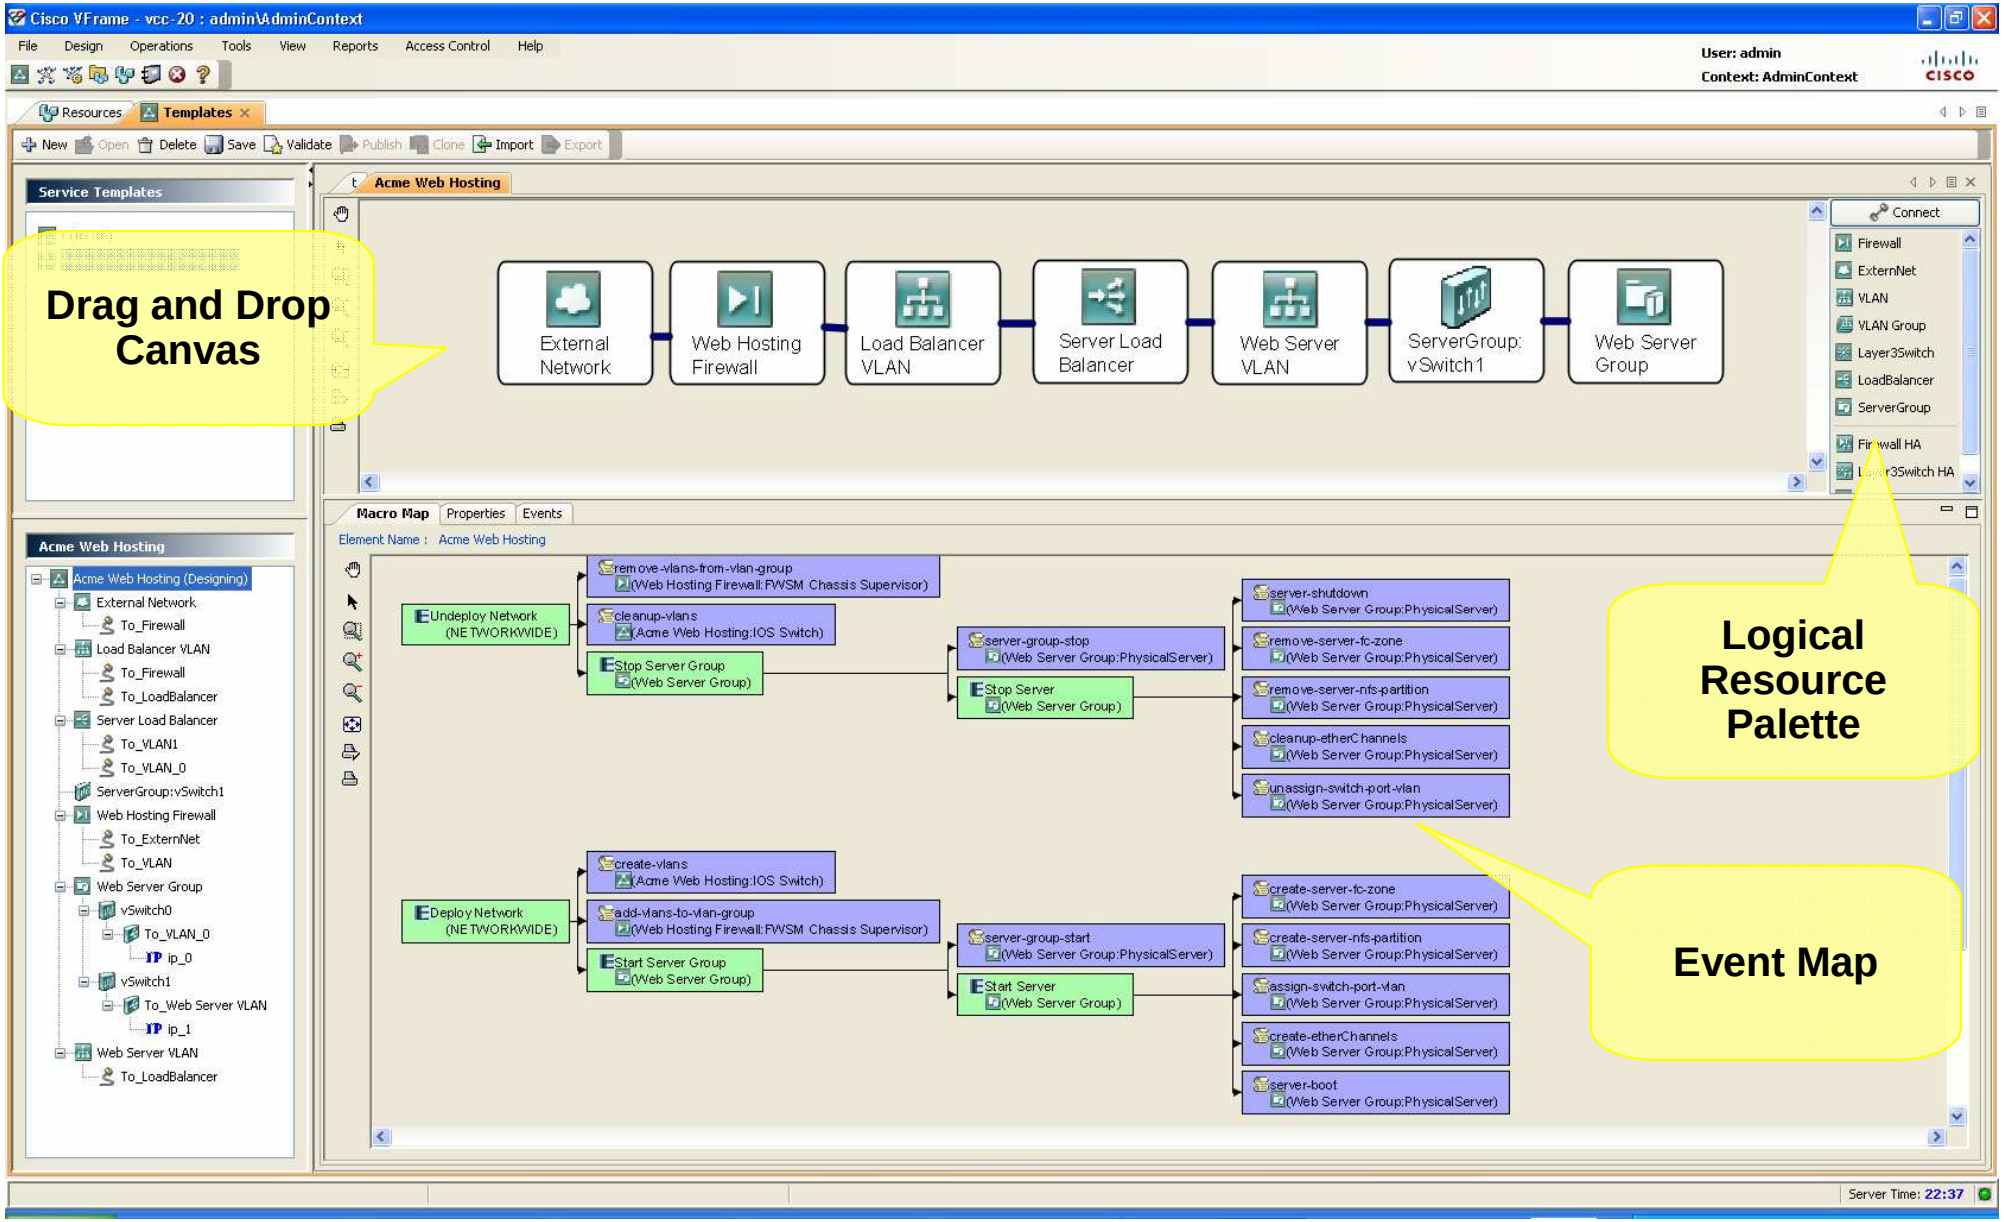

VFrame DC: Service Template: Events & Actions

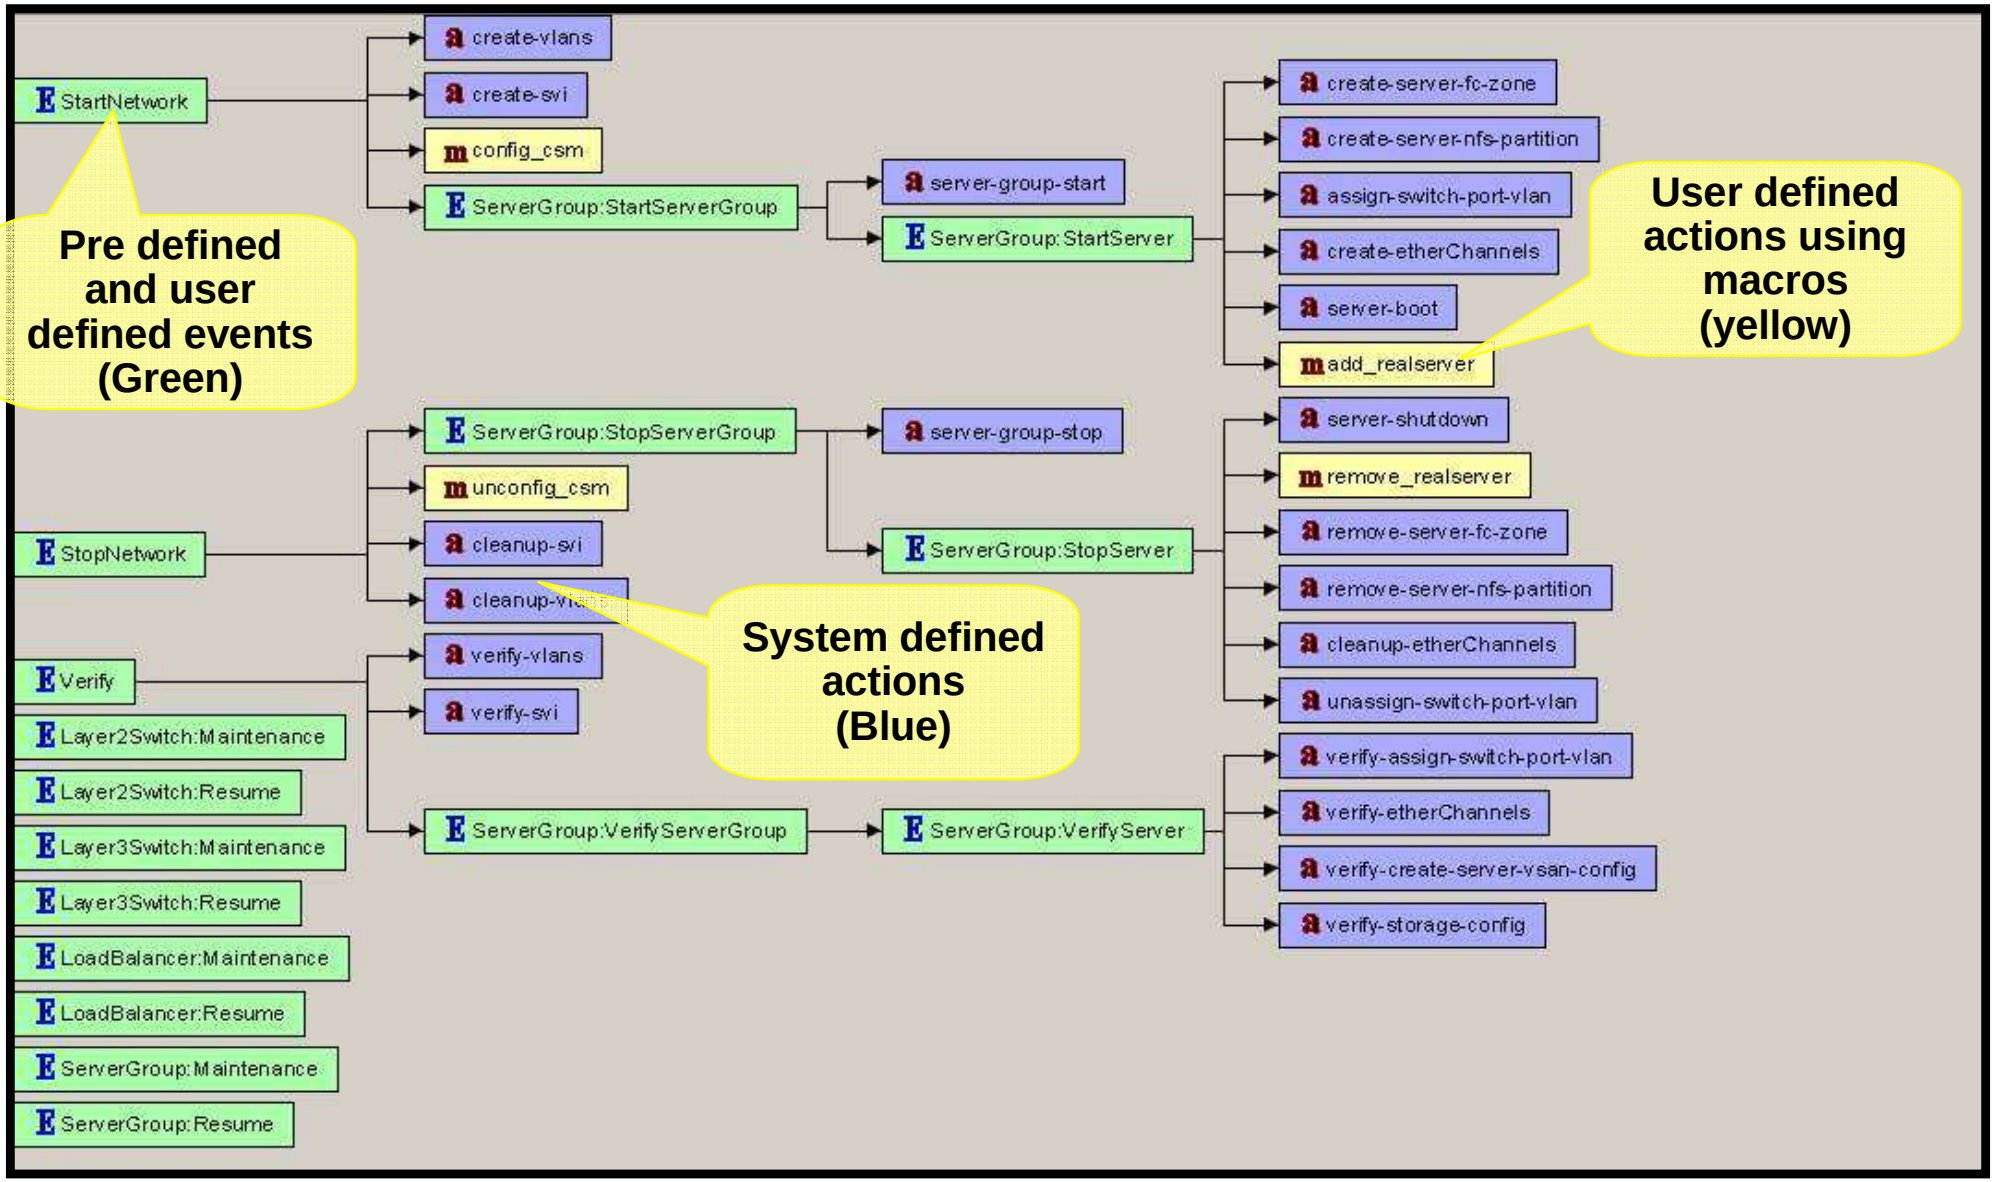

VFrame DC: Discovered Resources

- Server
  - Server discovery and inventory of CPU, Memory, Local hard disk
  - Server to switch port binding for Ethernet and Fibre Channel
  - LOM and server/LOM binding
- Storage
  - Storage LUN Discovery
  - NAS Volume Discovery
- SAN
  - MDS Switches, VSANs, Zones and IVR
- Network
  - Cat6k chassis, L2/L3 connectivity and L2 topology
  - VLANs, SVIs
- Network Services
  - Service modules (FWSM, CSM), pre-created contexts and HA pairs

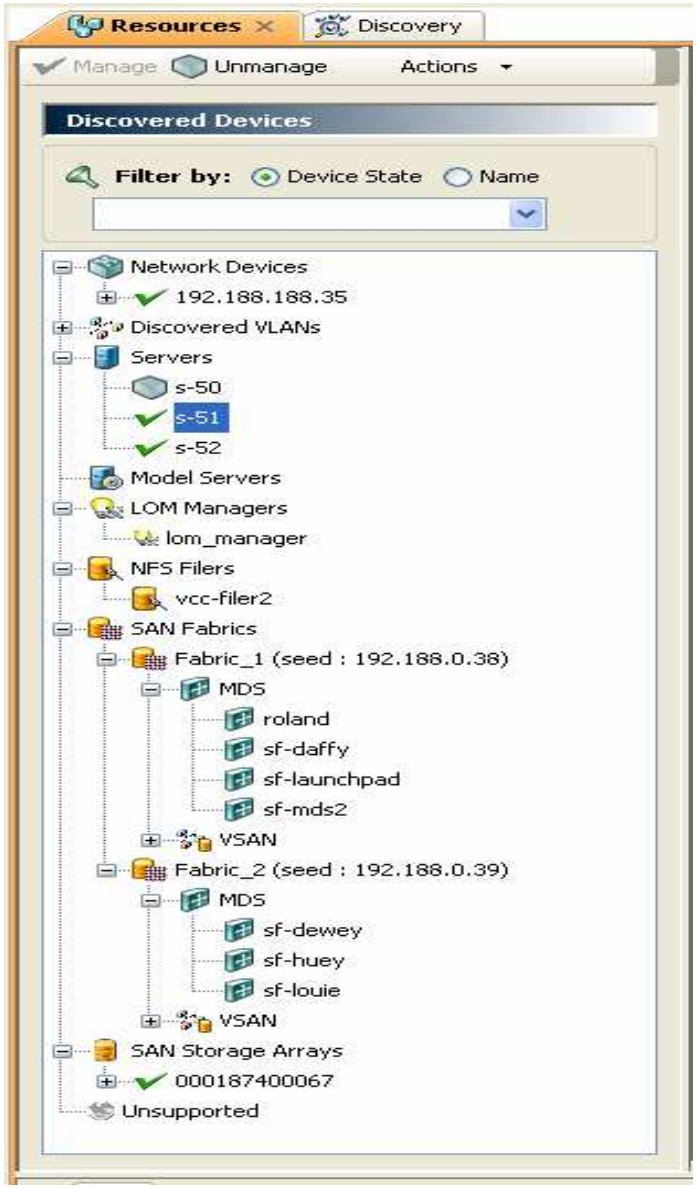

VFrame DC: Roles and Virtual Contexts

- **Delegate resources to a particular customer or department**
- **Securely share global templates**
- **Control access to specific users based on role**
  sysadmins, storage admins, network admins
- **Specify the role for the user**
  Designer, Operator

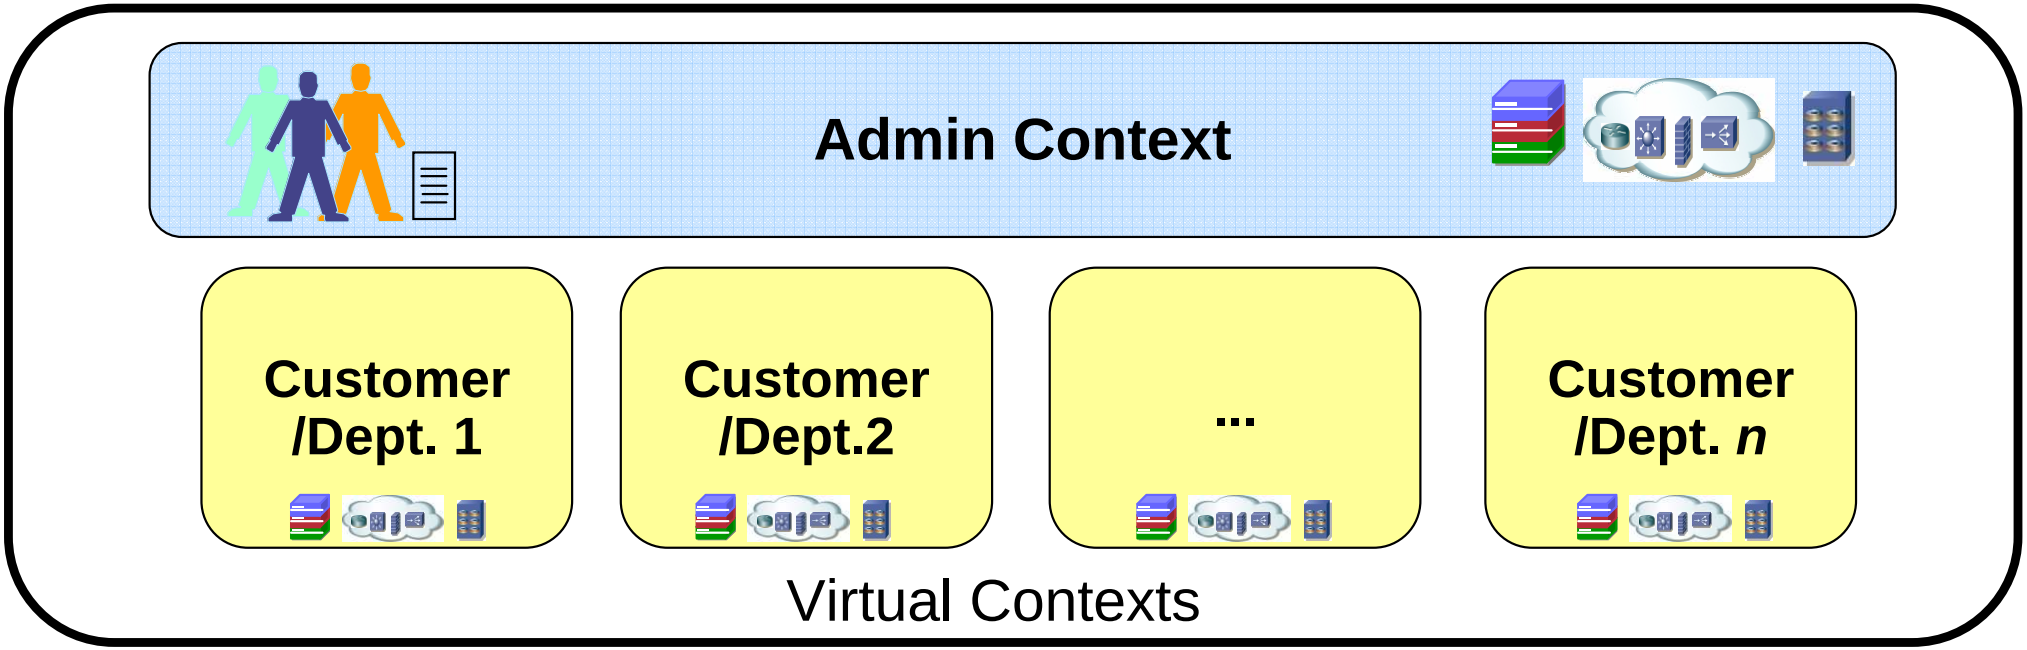

Upcoming release updates

- **Release 1.1.4 – October 3rd week 2007**
  Key Features:
  - **Support for Top of Rack 4948, 3750 Switches**
  - **Support for 3110 and 3120 Blade Switches**
  - **Support for MDS Blade Switches**
  - **Support for IBM Blade Center BC-H with Cisco Switching**
  - **Support for HP c-class Blade Server with Cisco Switching**
  - **New server Lights Out Mgmt. addition mechanism**
  - **SNMPv3 support**
- **Release 1.2 VMware ESX integration beta – December 2007**
- **Release 1.2 – Q2CY08**
  - **VMware ESX Integration**
  - **Flexible SAN provisioning (WWN)**
  - **ACE/ANM integration (eng scoping in progress)**

Unlocking the Benefits of Virtualization

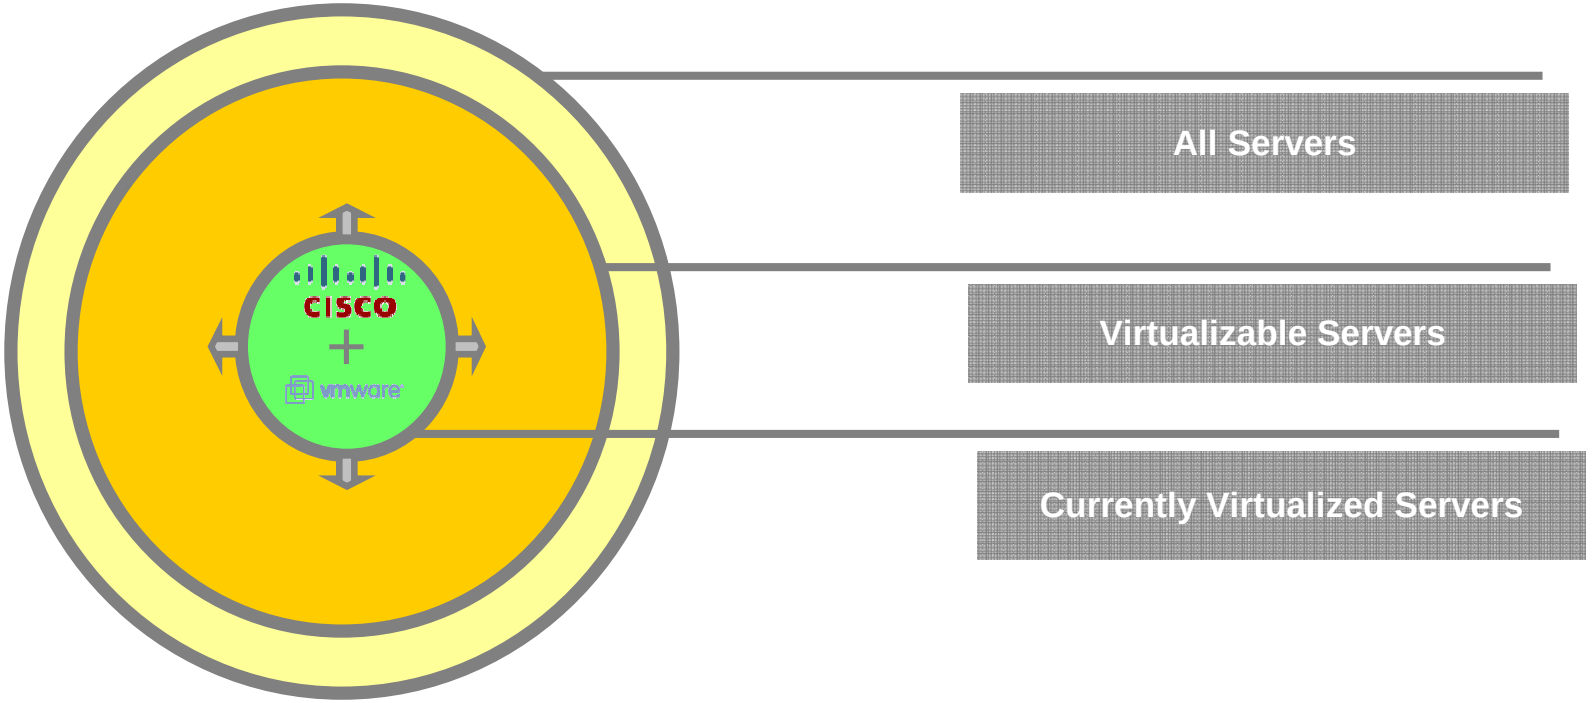

- Kickstarting/Ghost Imaging the ESX software is the easy part
- Configuring the network and storage for the ESX is the hard part
- Repeatability, compliant provisioning of network and storage is key
- VFrame DC provides automated, requirements based provisioning for VMware ESX

VFrame Services with ESX Deployments

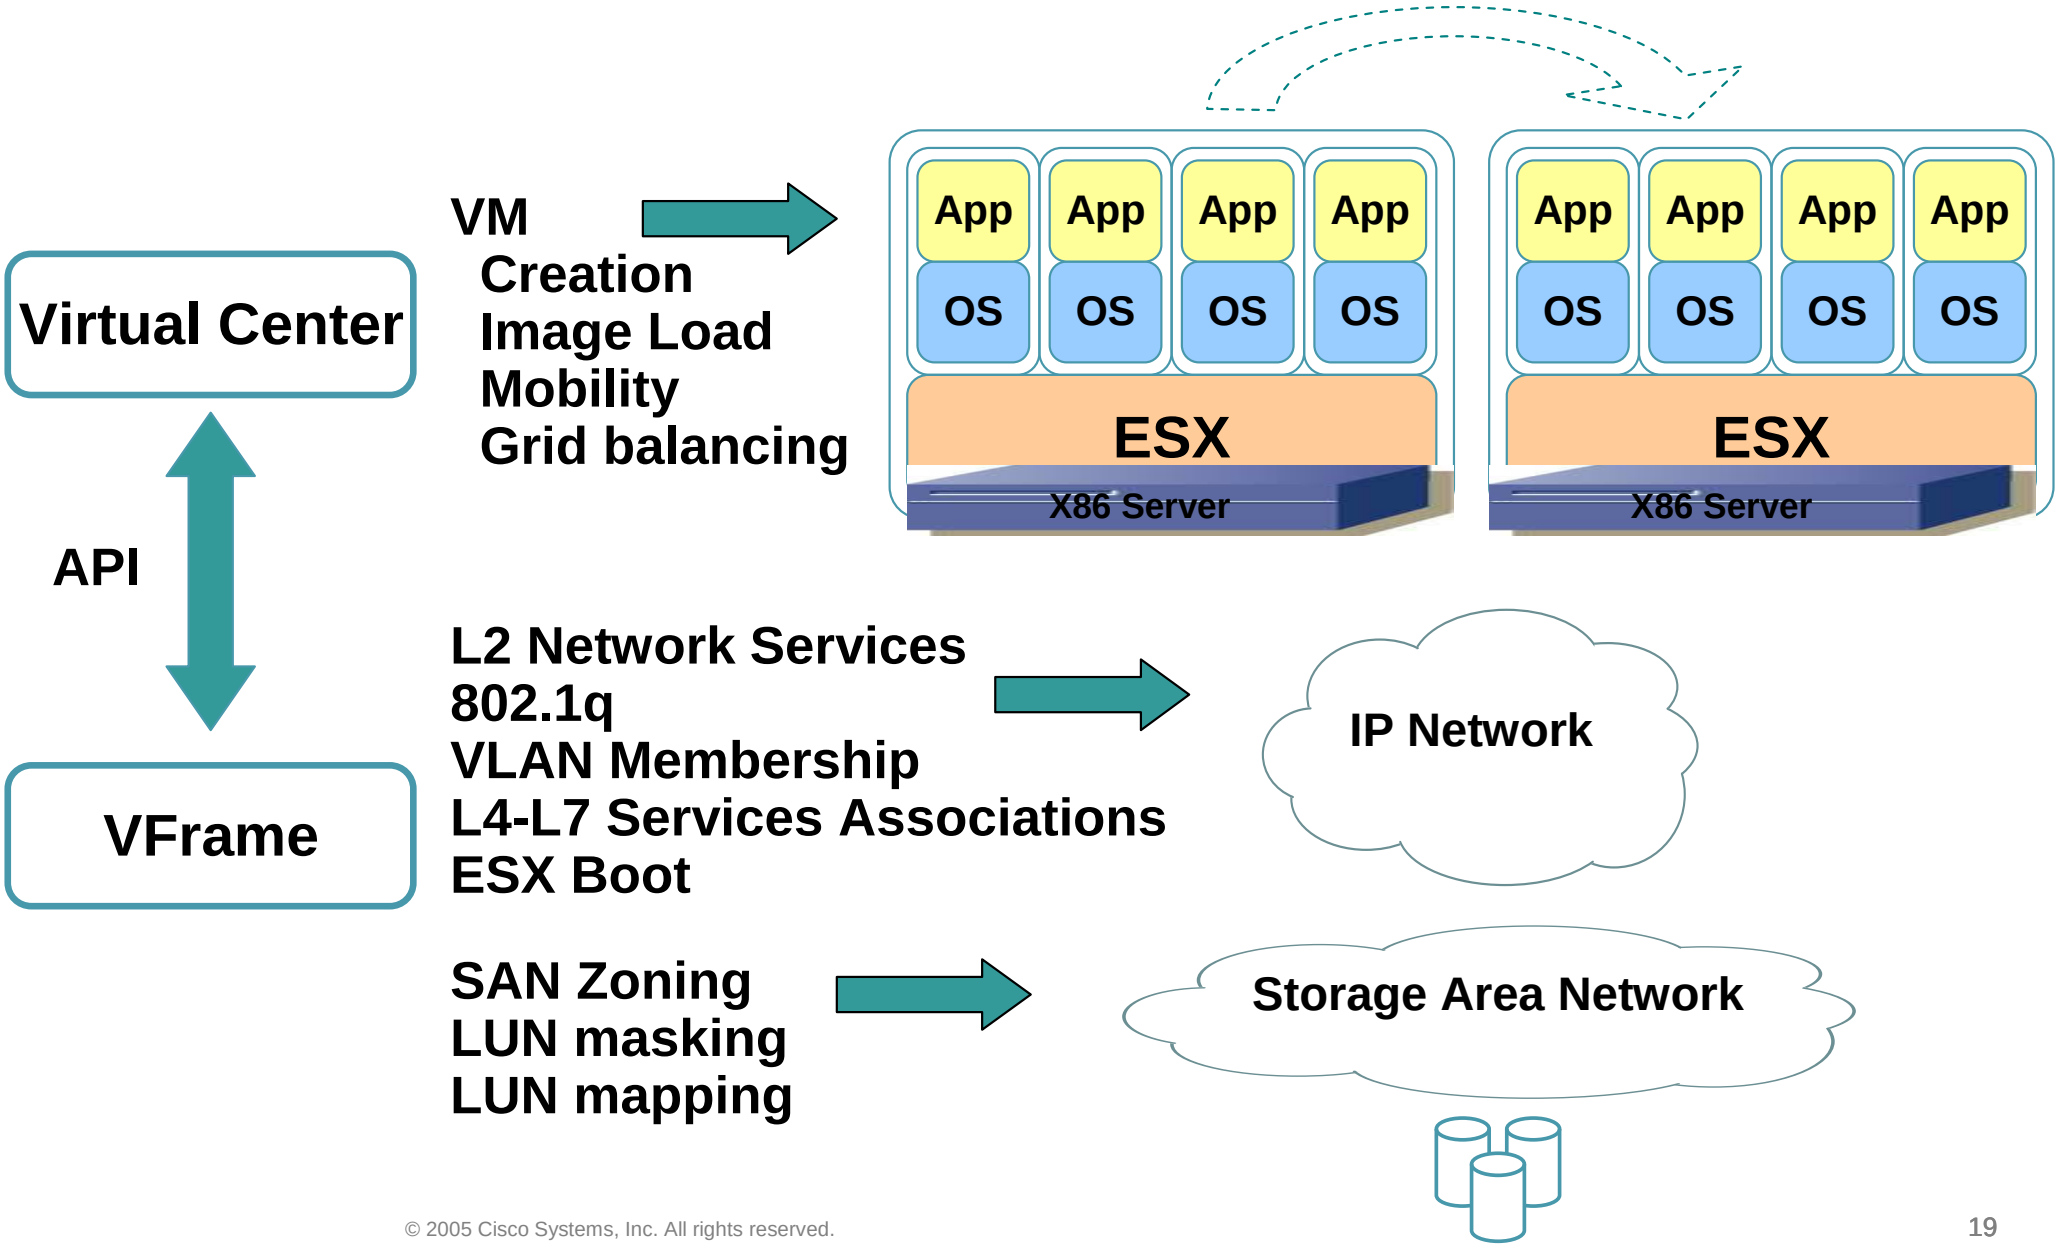

ESX Cluster design in VFrame

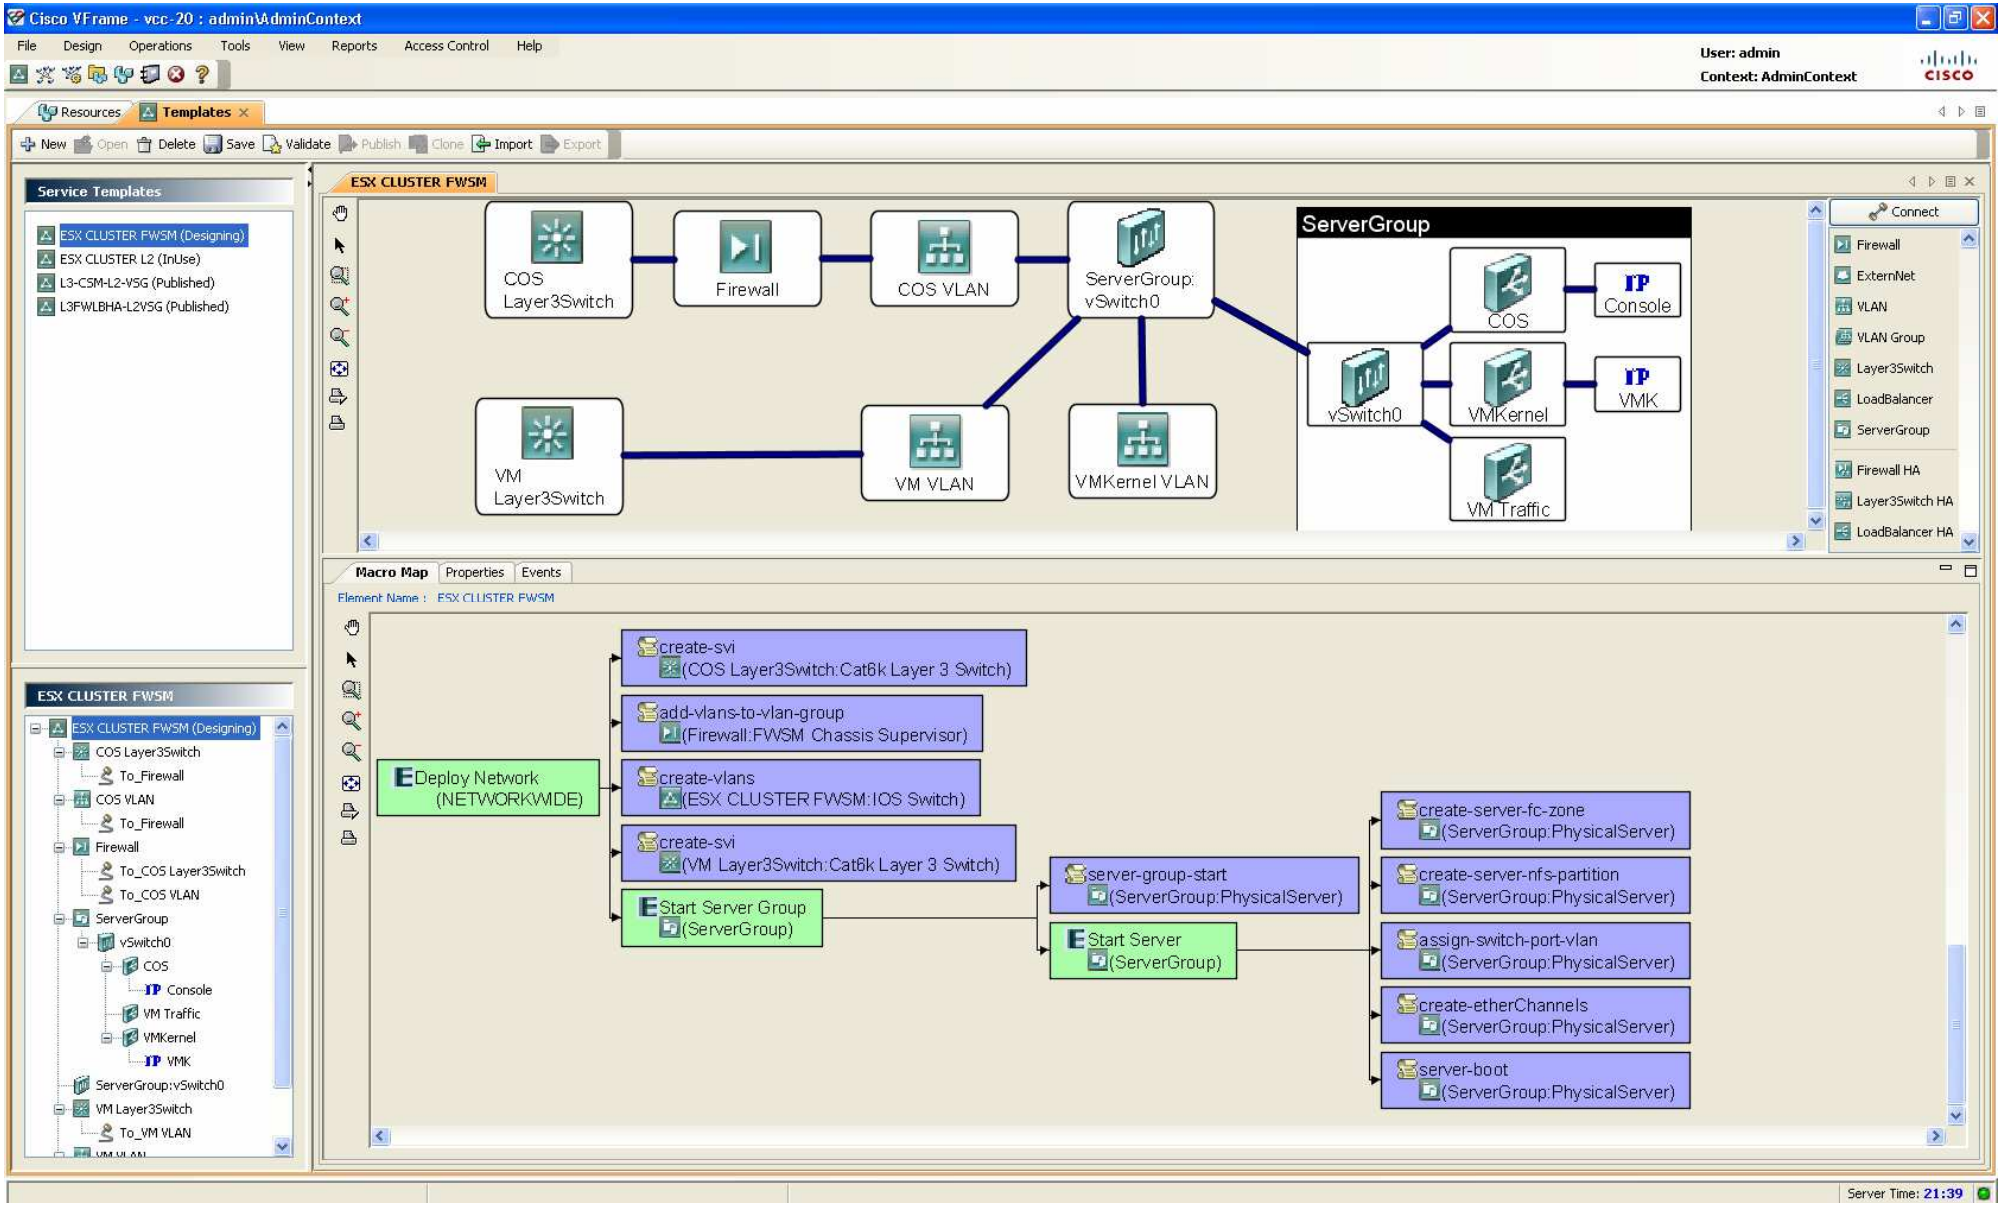

Cisco N-Port Virtualizer (NPV) for Blades and Fabric Switches

Agenda

- **Introduction to NPV**
- **Basics of NPIV**
- **MDS NPV Implementation**
- **NPV Considerations and Limitations**

Introduction to NPV

Blade Switch Explosion Issues

- **Scalability**
  - **– Each Blade Switch uses a single Domain ID**
  - **– Theoretical maximum number of Domain IDs is 239 *per* VSAN**
  - **– Supported number of domains is quite smaller (depends on OSM)**
    - **EMC: 40 domains**
    - **Cisco Tested: 75**
    - **HP: 40 domains**
    - **Other OSM Do Not Post**
- **Manageability**
  - **– More switches to manage**
  - **– Shared management of blade switches between storage and server administrators**

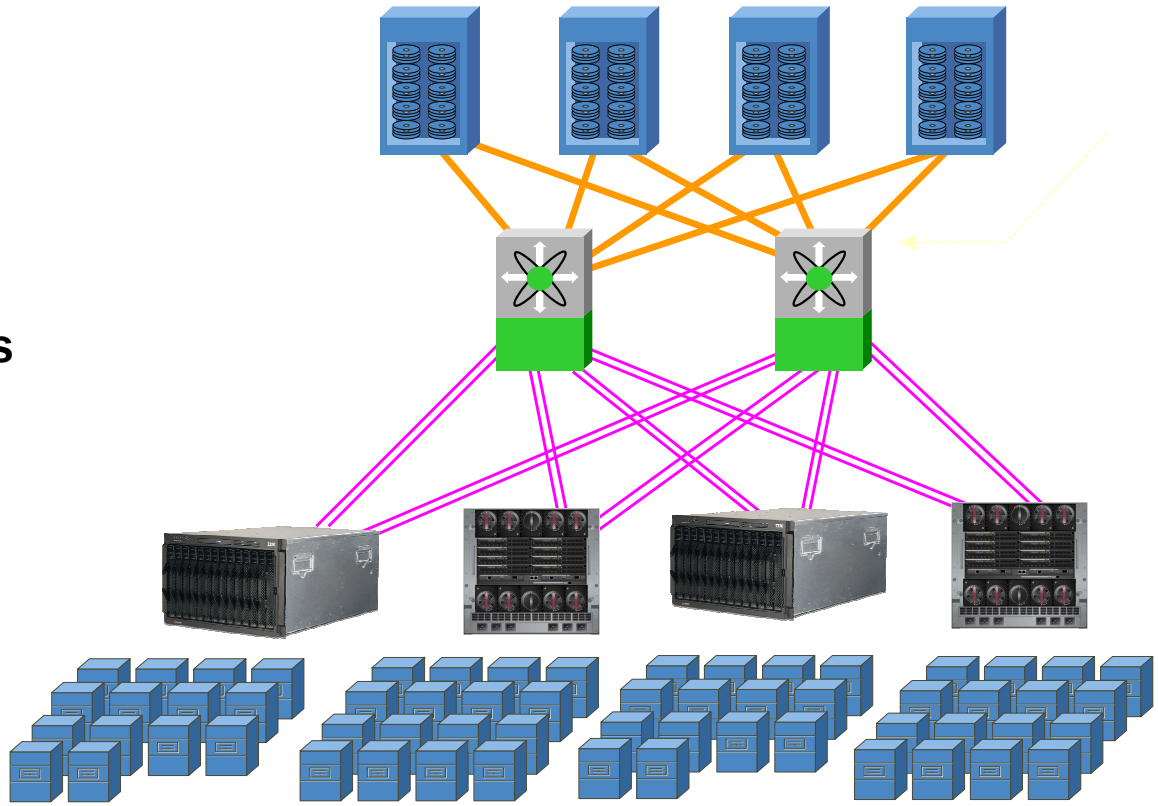

Cisco MDS' NPV Solution

MDS NPV

- NPV enable switch act as a NPIV host
- NPV switch uplink is no longer an ISL
- NPV switch DOES NOT use a Domain ID
  - No longer limited to Domain ID limitation

Manageability

- Less amount of switches to manage
- NPV enable switch is now managed like a NPIV enabled host
- Eliminates the need for server administrators to manage the SAN

NPV Supported Switches

NPV Edge Switches

- MDS 9124 & MDS 9134
- IBM and HP Blade Switches

NPV Core Switches

- MDS 9500 Family of Directors
- MDS 9216A, MDS 9216i & MDS 9222i
- 3rd Party Switches
  - Needs to support NPIV
  - Needs Testing/Qualification

Basics of NPIV

N-Port ID Virtualization (NPIV)

- **Mechanism to assign multiple N_Port_IDs to a single N_Port**
- **Allows all the Access control, Zoning, Port Security (PSM) be implemented on application level**
- **Multiple N_Port_IDs are allocated in the same VSAN**

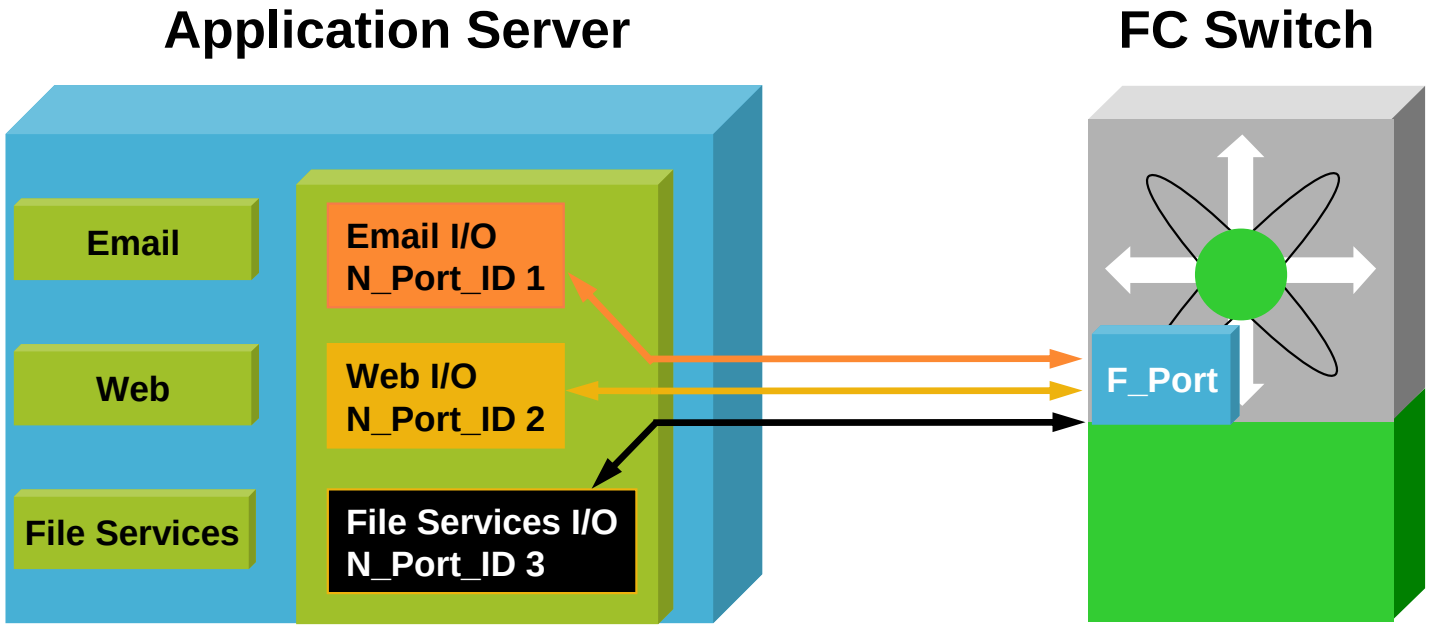

MDS NPV Implementation

NPV Load-Balancing

- **Server FLOGI/FDISC to first available NP link and all data flow for that server goes through the same NP (uplink)**
- **FLOGI/FDISC on all available NP links is load-balance via Round-Robin**

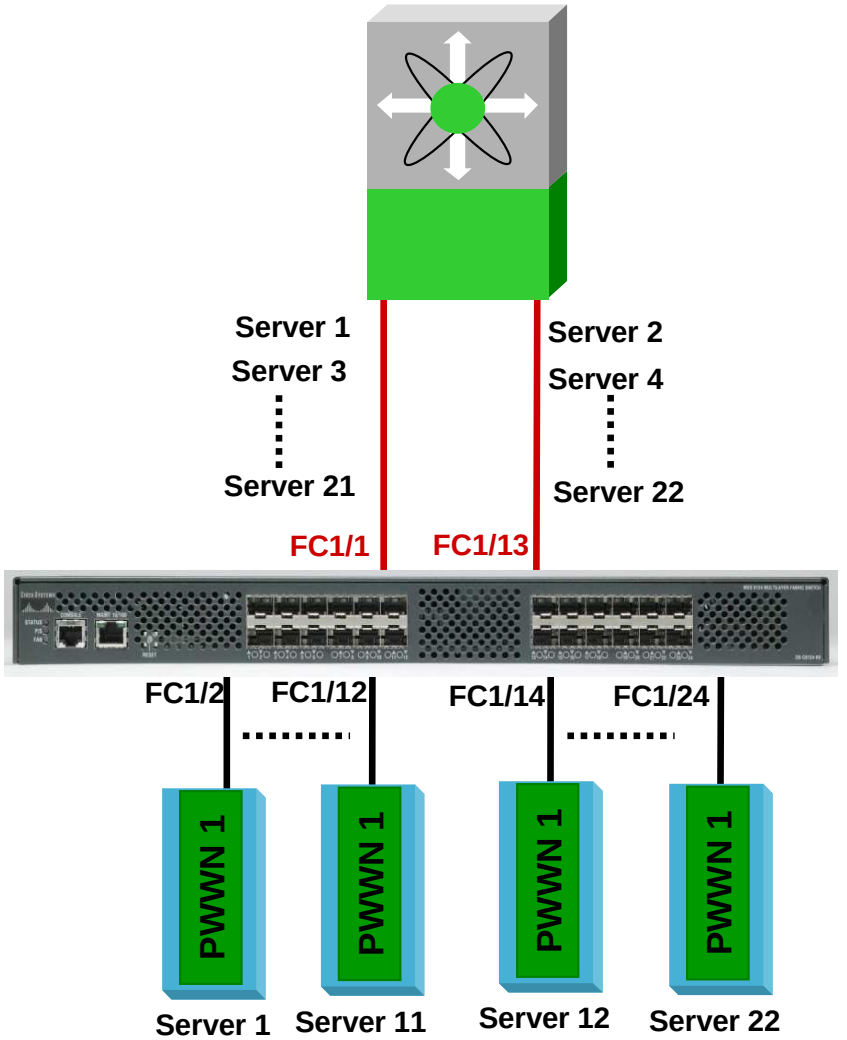

NPV Failover/Failback

Failover

– All servers on failed NP will re-initialize and login to available NP link

Failback

– Initial release of NPV does not provide server preferred NP link to allow failback to original NP link

– Need to re-initialize server port(s) manually from NPV switch or from server

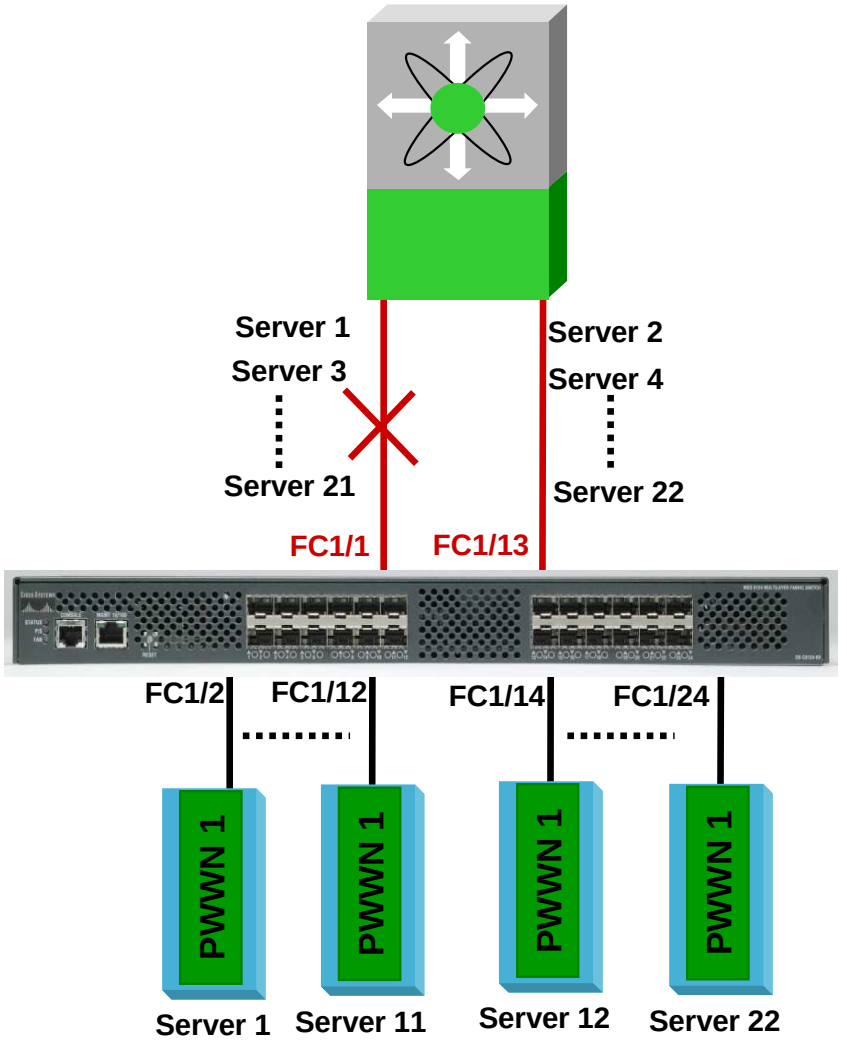

Nested NPIV Load-Balancing

- NPIV Server FLOGI/FDISC to first available NP link and all data flow goes through the same NP (uplink)
- FLOGI/FDISC on all available NP links is load-balance via Round-Robin
- All Nested NPIV logins will flow through the same NP port as the end device is using

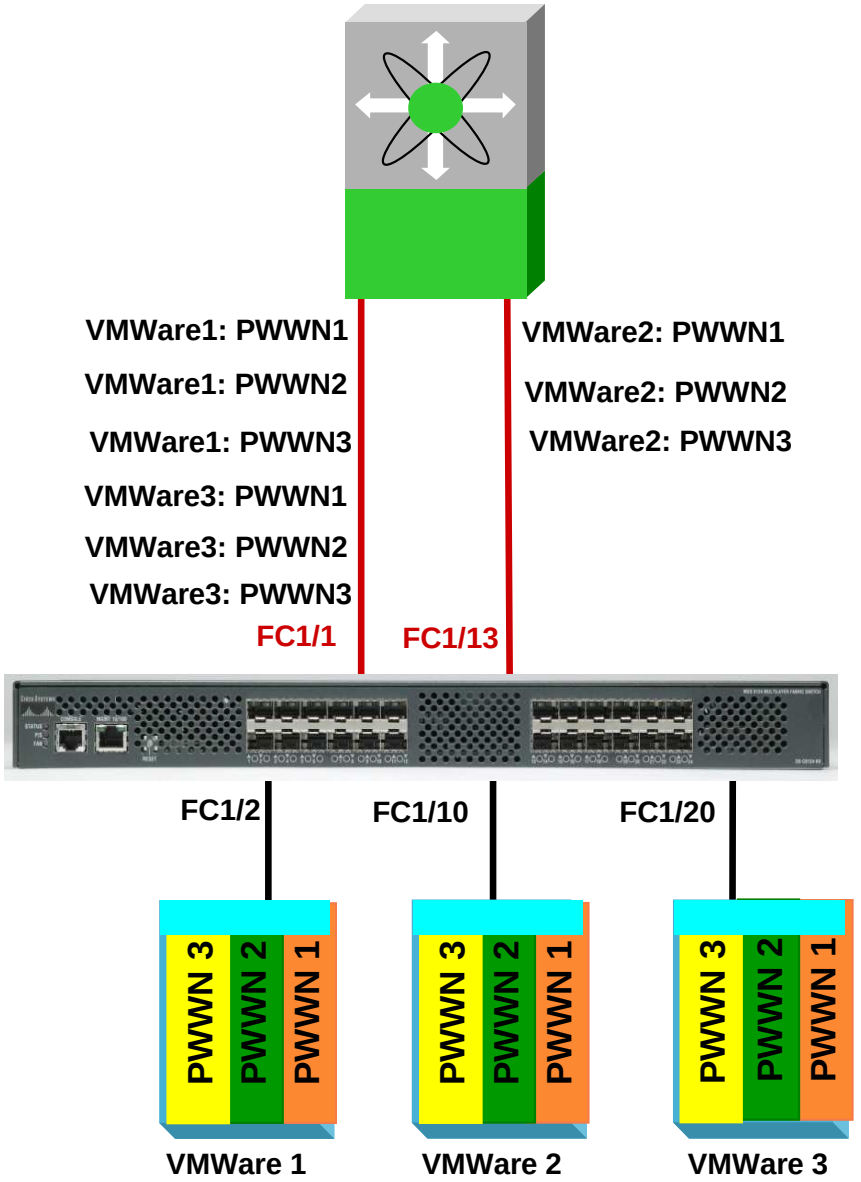

Nested NPIV Failover/Failback

Failover

– All servers on failed NP will re-initialize and login to available NP link

Failback

– Initial release of NPV does not provide server preferred NP link to allow failback to original NP link

– Need to re-initialize NPIV server port(s) manually from NPV switch (affects all PWWN on NPIV server)

– Or from virtual server to specify which PWWN to re-initialize

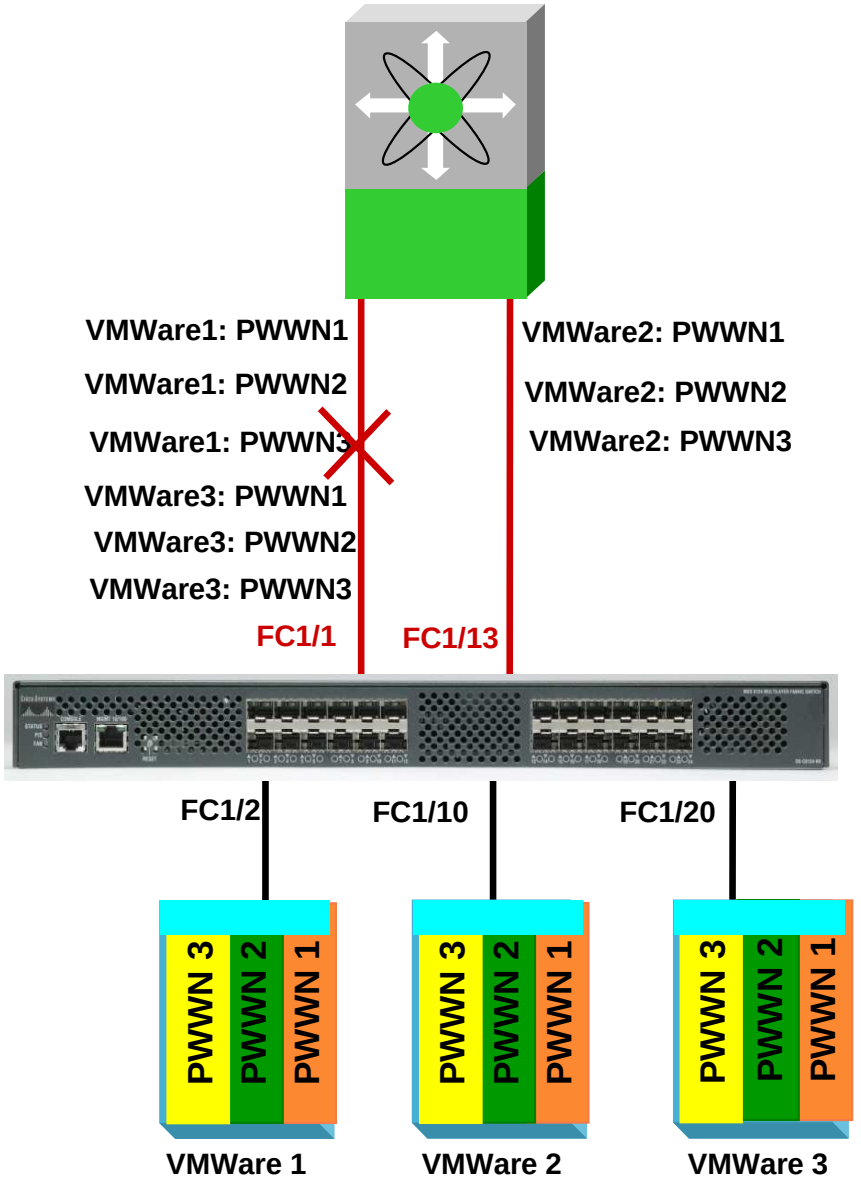

VSAN Support with NPV

- NPV switch supports up to 16 VSANs
- VSAN Port Membership for NPV's NP port and Core switch MUST MATCH

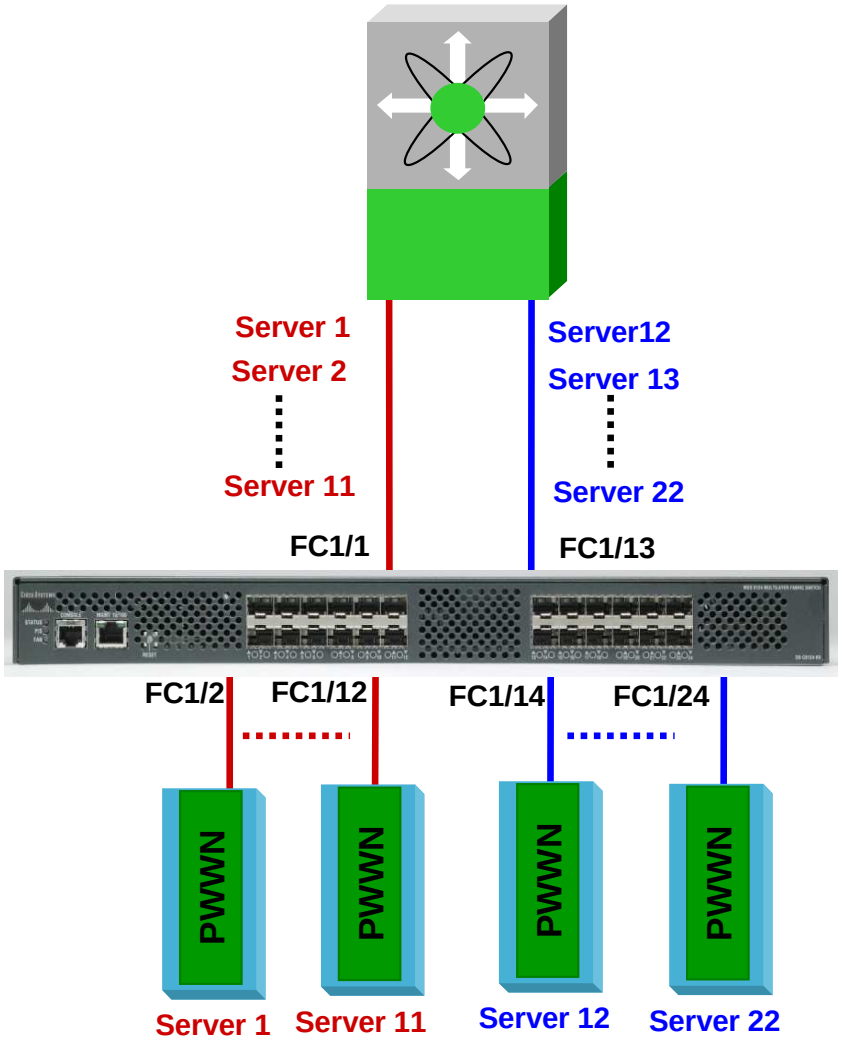

VSAN Load-Balancing on NPV

- **Servers can only utilize NP ports on same VSAN**
- **Load-balancing can not span across other NP ports not in the same VSAN**
- **Round-Robin across same NP ports in the same VSAN**

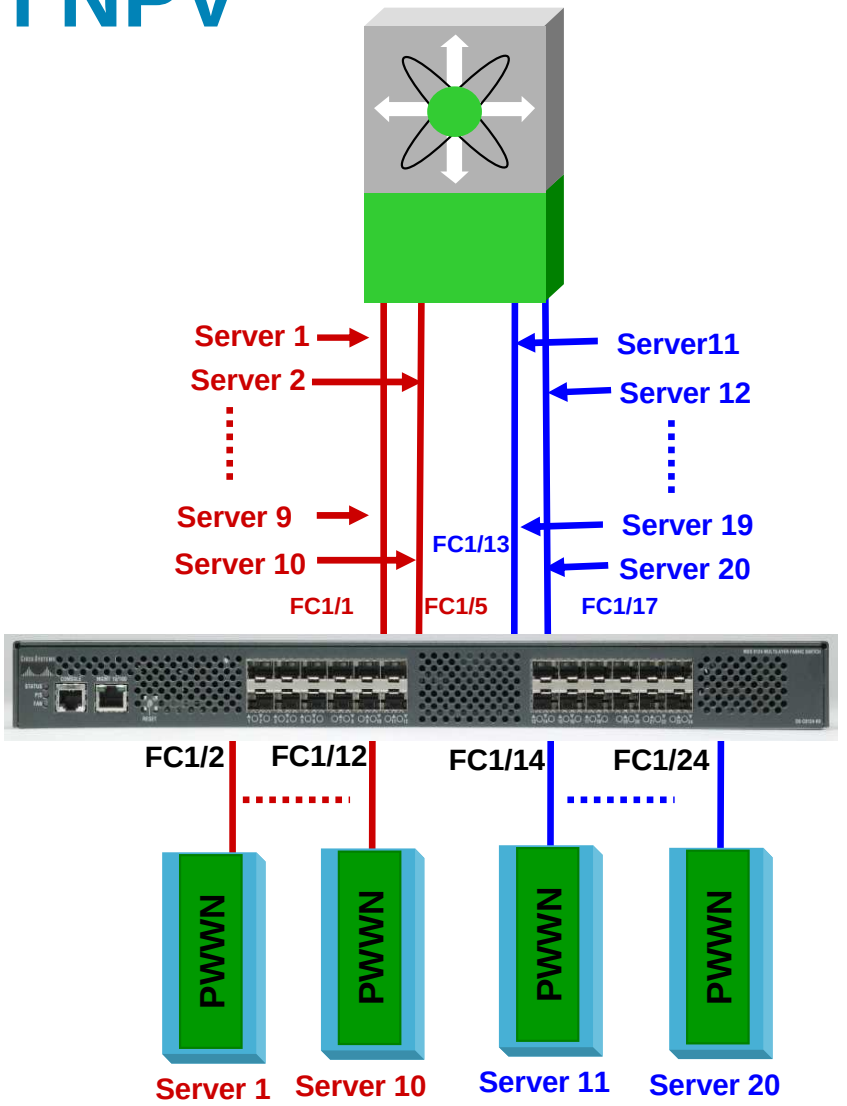

NPV FCID Allocation Examples

Port-Group (Max # of Logins 250)

- F-Port will use 1 entry for regular server (non-nested NPIV)
- F-Port with nested NPIV will take up “X” (amount of logins) + 1 for physical port
- NP-Port will use a FCID entry for each login that uses its uplink

Port-Group1 (ports 1-4)

- FC1/1 (NP-port) → 8 FCIDs
- Total = 8 FCIDs for PG1

PortGroup2 (ports 5-8)

- Total = 0 since no NP uplink for PG2

PortGroup3 (ports 9-12)

- Total = 0 since no NP uplink for PG3

PortGroup4 (ports 13-16)

- Total = 4 FCIDs for PG4

PortGroup5 (ports 17-20)

- Total = 0 since no NP uplink for PG5

PortGroup6 (ports 21-24)

- Total = 0 since no NP uplink for PG6

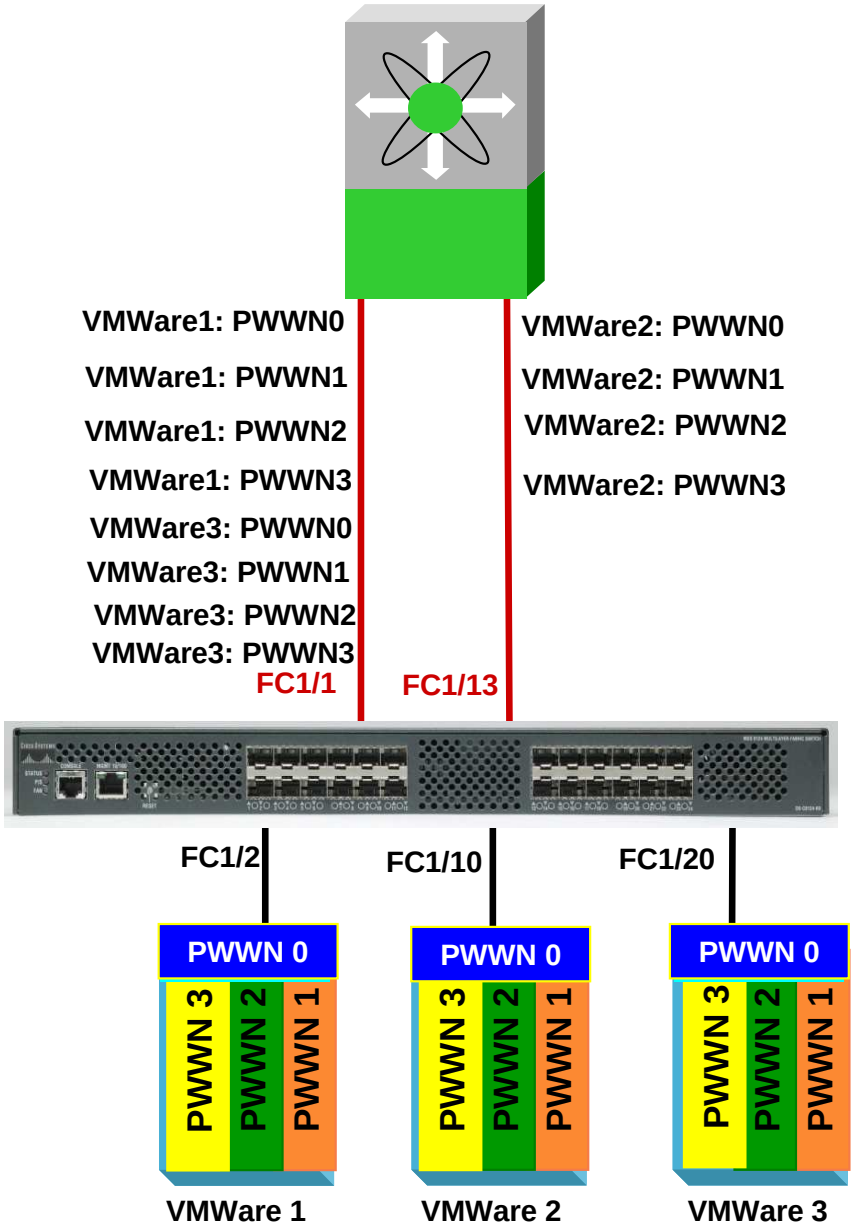

MDS Feature Support

Supported Features in NPV mode

- No needed license for NPV
- Need active port-licenses
- Software hitless upgrade
- Storage Device Virtualization (SDV) - Done at the Core switch
- Dynamic Port-VSAN Membership (DPVM) - Done at the Core switch

Unsupported Features in NPV mode

- No local switching
- No Trunking or Port-Channeling
- No QoS

NPV Considerations & Limitations

NPV Enabled Environment

NPV Edge Switches

- Need to enable switch in “NPV” mode
- Changing To/From NPV mode
  - Disruptive: switch reboots
  - Configuration is not kept
- Supports only F, SD (SPAN Destination) and NP ports
- Supports 16 VSANs
- Local switching is NOT Supported
  - Switching is done at the core

NPV Core Switches

- Must enable NPIV feature
- Supports up to 100 NPV Edge Switches

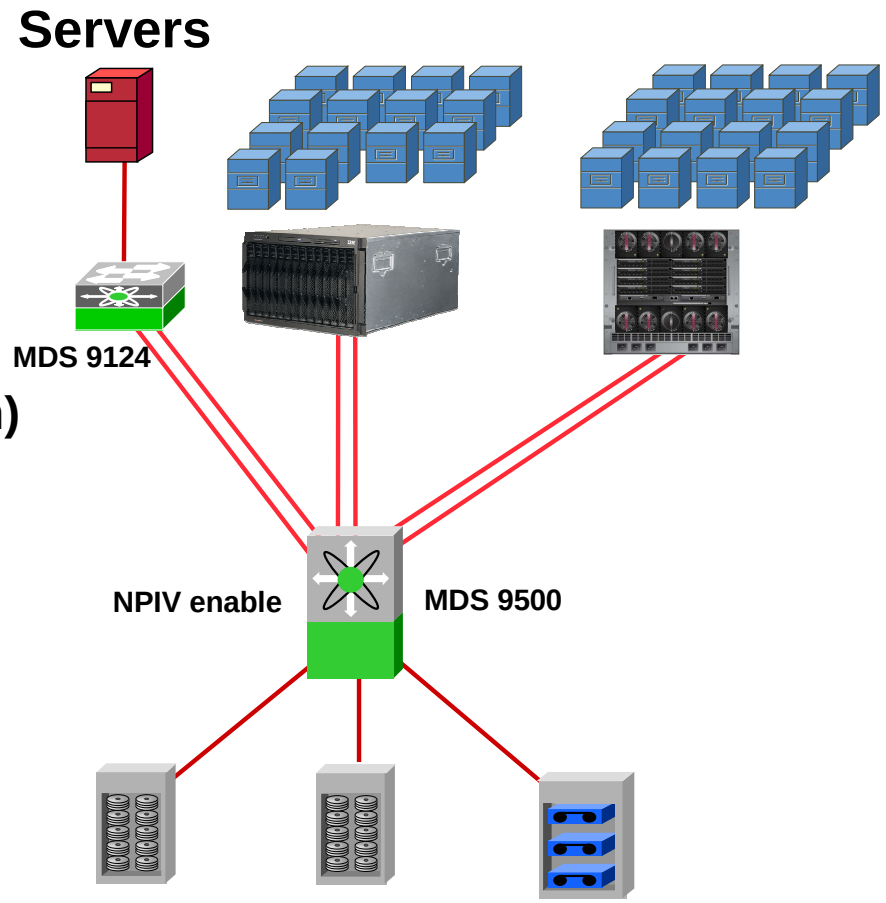

MDS 9124 – NPV Architecture

NPV Architecture

- Total of 6 Port-Groups – every 4 ports
- By default, first port in each Port-Group (ports 1, 5, 9, 13, 17 and 21) is set to “NP” mode for uplink to NPV Core Switch (Can be changed)
- All other ports are set to “F” for device connectivity (DOES NOT SUPPORT FL-Ports)
- A total of 250 FCIDs are allowed in each Port-Group
- Total of 1500 FCID per MDS 9124

Port-Group1: Ports 1 – 4
Port-Group2: Ports 5 – 8
Port-Group3: Ports 9 – 12
Port-Group4: Ports 13 – 16
Port-Group5: Ports 17 – 20
Port-Group6: Ports 21 - 24

MDS 9134 – NPV Architecture

NPV Architecture

- Total of 10 Port-Groups
  - Port-Group consists of 4 ports for 1/2/4Gig ports grouping
  - Each 10Gig port is its own Port-Group
- By default, first port in each Port-Group (ports 1, 5, 9, 13, 17, 21, 25 and 29) is set to “NP” mode for uplink to NPV Core Switch (Can be changed)
- Both 10Gig port is set to “NP” mode
- All other ports are set to “F” for device connectivity (DOES NOT SUPPORT FL-Ports)
- A total of 250 FCIDs are allowed in each Port-Group
- Total of 2500 FCID per MDS 9134

Port-Group1: Ports 1 – 4	Port-Group2: Ports 5 – 8
Port-Group3: Ports 9 – 12	Port-Group4: Ports 13 – 16
Port-Group5: Ports 17 – 20	Port-Group6: Ports 21 – 24
Port-Group7: Ports 25-28	Port-Group8: Ports 29-32
Port-Group9: Port 1 (10G)	Port-Group10: Port 2 (10G)

Port Mapping for IBM Blade Switches

IBM Blade Switch NPV Architecture

- Total of 5 Port-Groups
  - PG 1 -> Uplink 1
  - PG 2 -> Uplink 2
  - PG 3 -> Uplink 3
  - PG 4 -> Uplink 4
  - PG 5 -> Uplink 5 and Uplink 6
- All External Uplink ports are set to “NP” mode
- All Blade Server slot ports are set to “F” mode (DOES NOT SUPPORT FL-Ports)
- A total of 250 FCIDs are allowed in each Port-Group
- Total of 1250 FCID per IBM Blade Switch

Port Mapping for HP Blade Switches

HP Blade Switch NPV Architecture

- Total of 6 Port-Groups
  - PG 1 -> EXT Port 1
  - PG 2 -> EXT Port 2
  - PG 3 -> EXT Port 3 & EXT Port 4
  - PG 4 -> EXT Port 5 & EXT Port 6
  - PG 5 -> EXT Port 7
  - PG 6 -> EXT Port 8
- All External Uplink ports are set to “NP” mode
- All Blade Server slot ports are set to “F” mode (DOES NOT SUPPORT FL-Ports)
- A total of 250 FCIDs are allowed in each Port-Group
- Total of 1500 FCID per HP Blade Switch

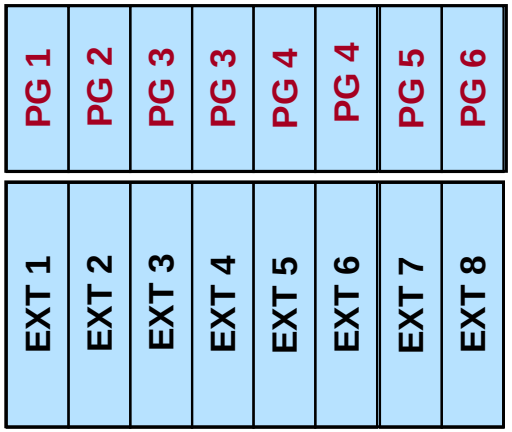

Cisco MDS N_Port Virtualizer Roadmap

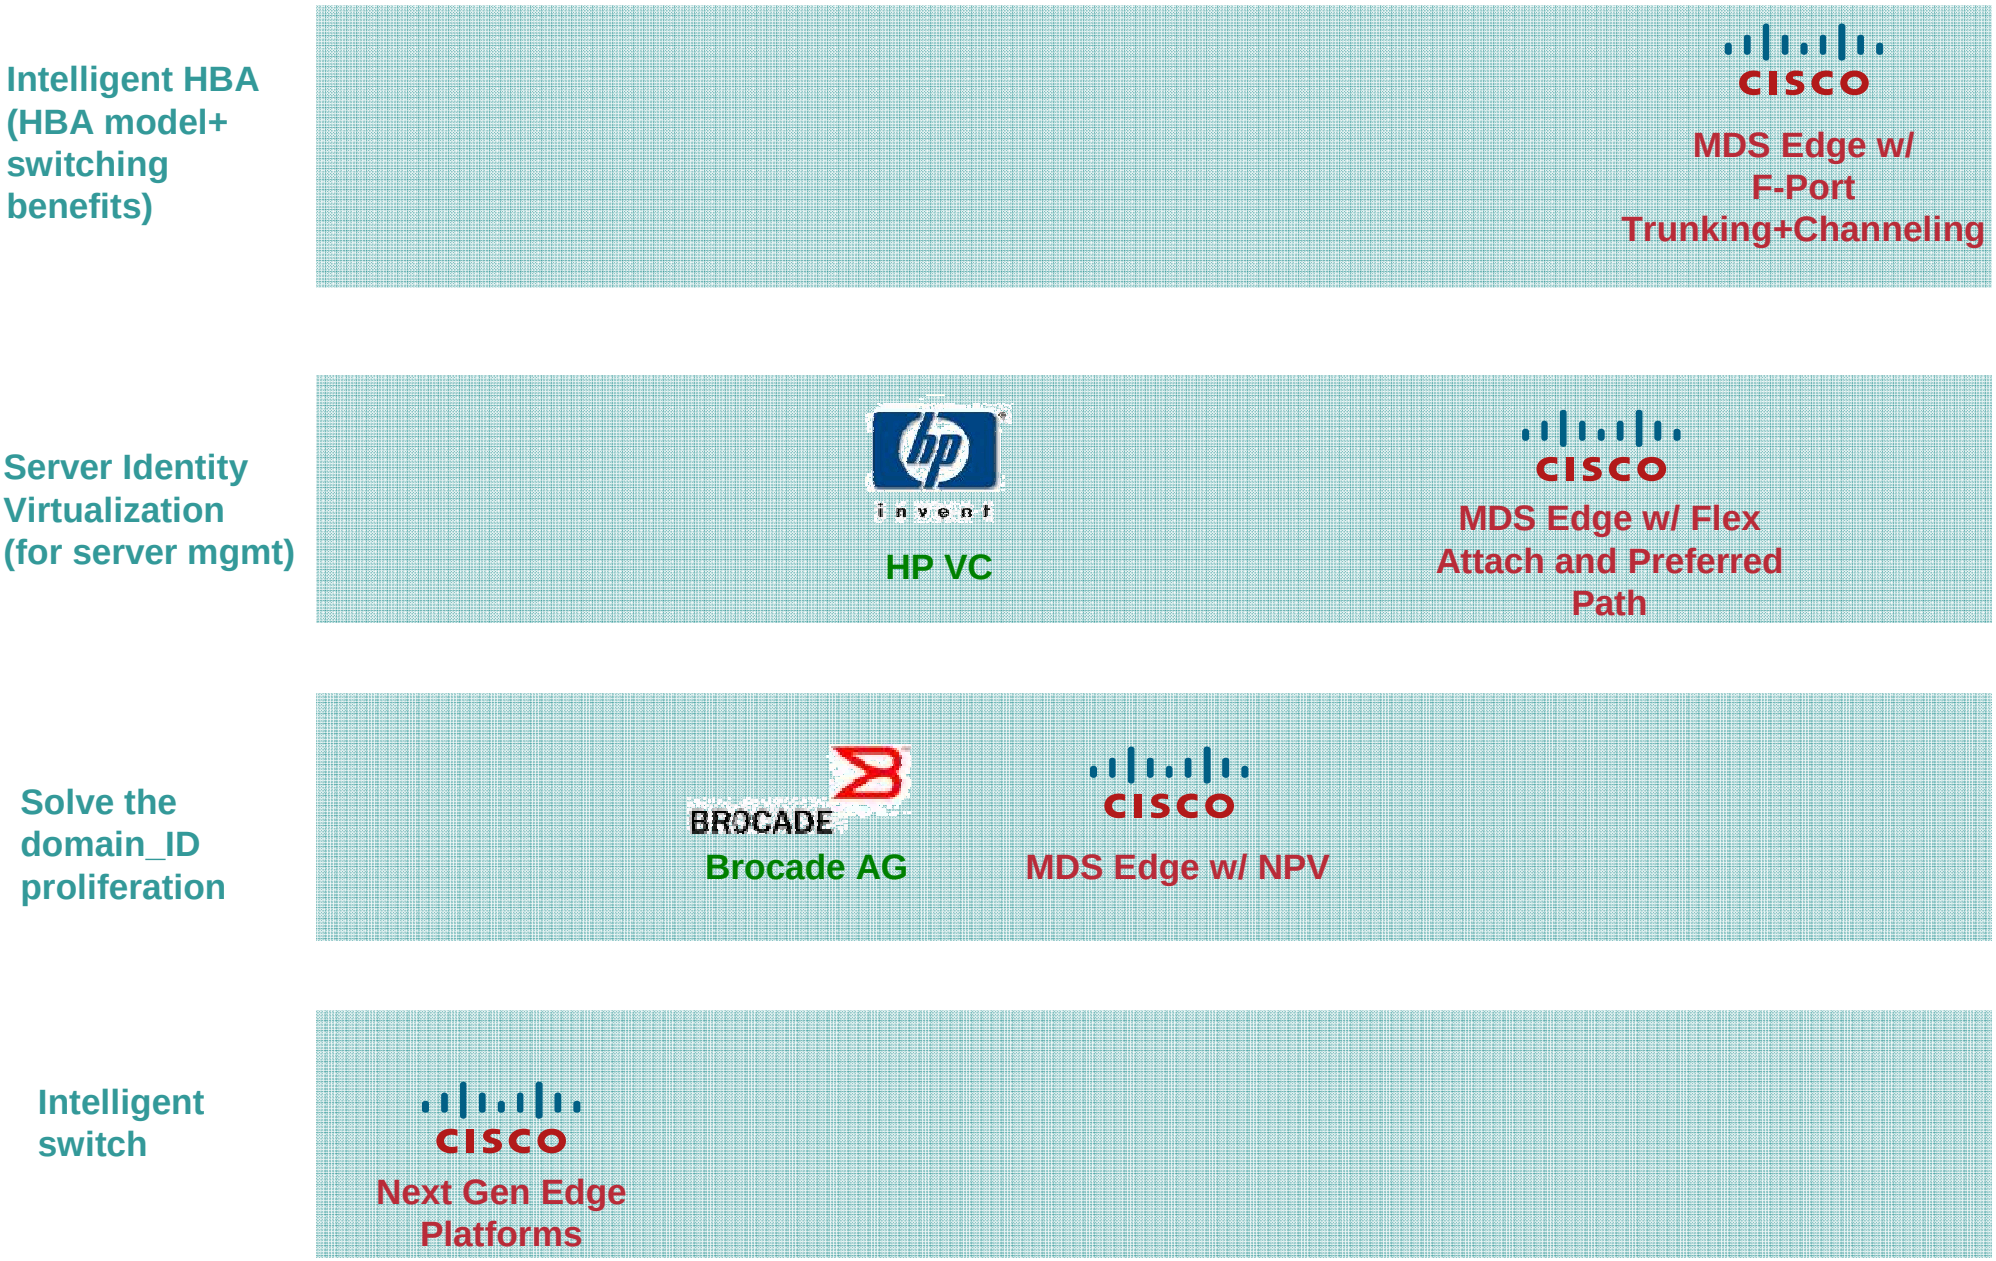

NG Blade Switching Overview

Agenda

- **Blade Server Market Update**
- **Ethernet Blade Portfolio**
- **Fibre Channel & Infiniband Offerings & Roadmap**

Bladeserver Market Growing Very Rapidly

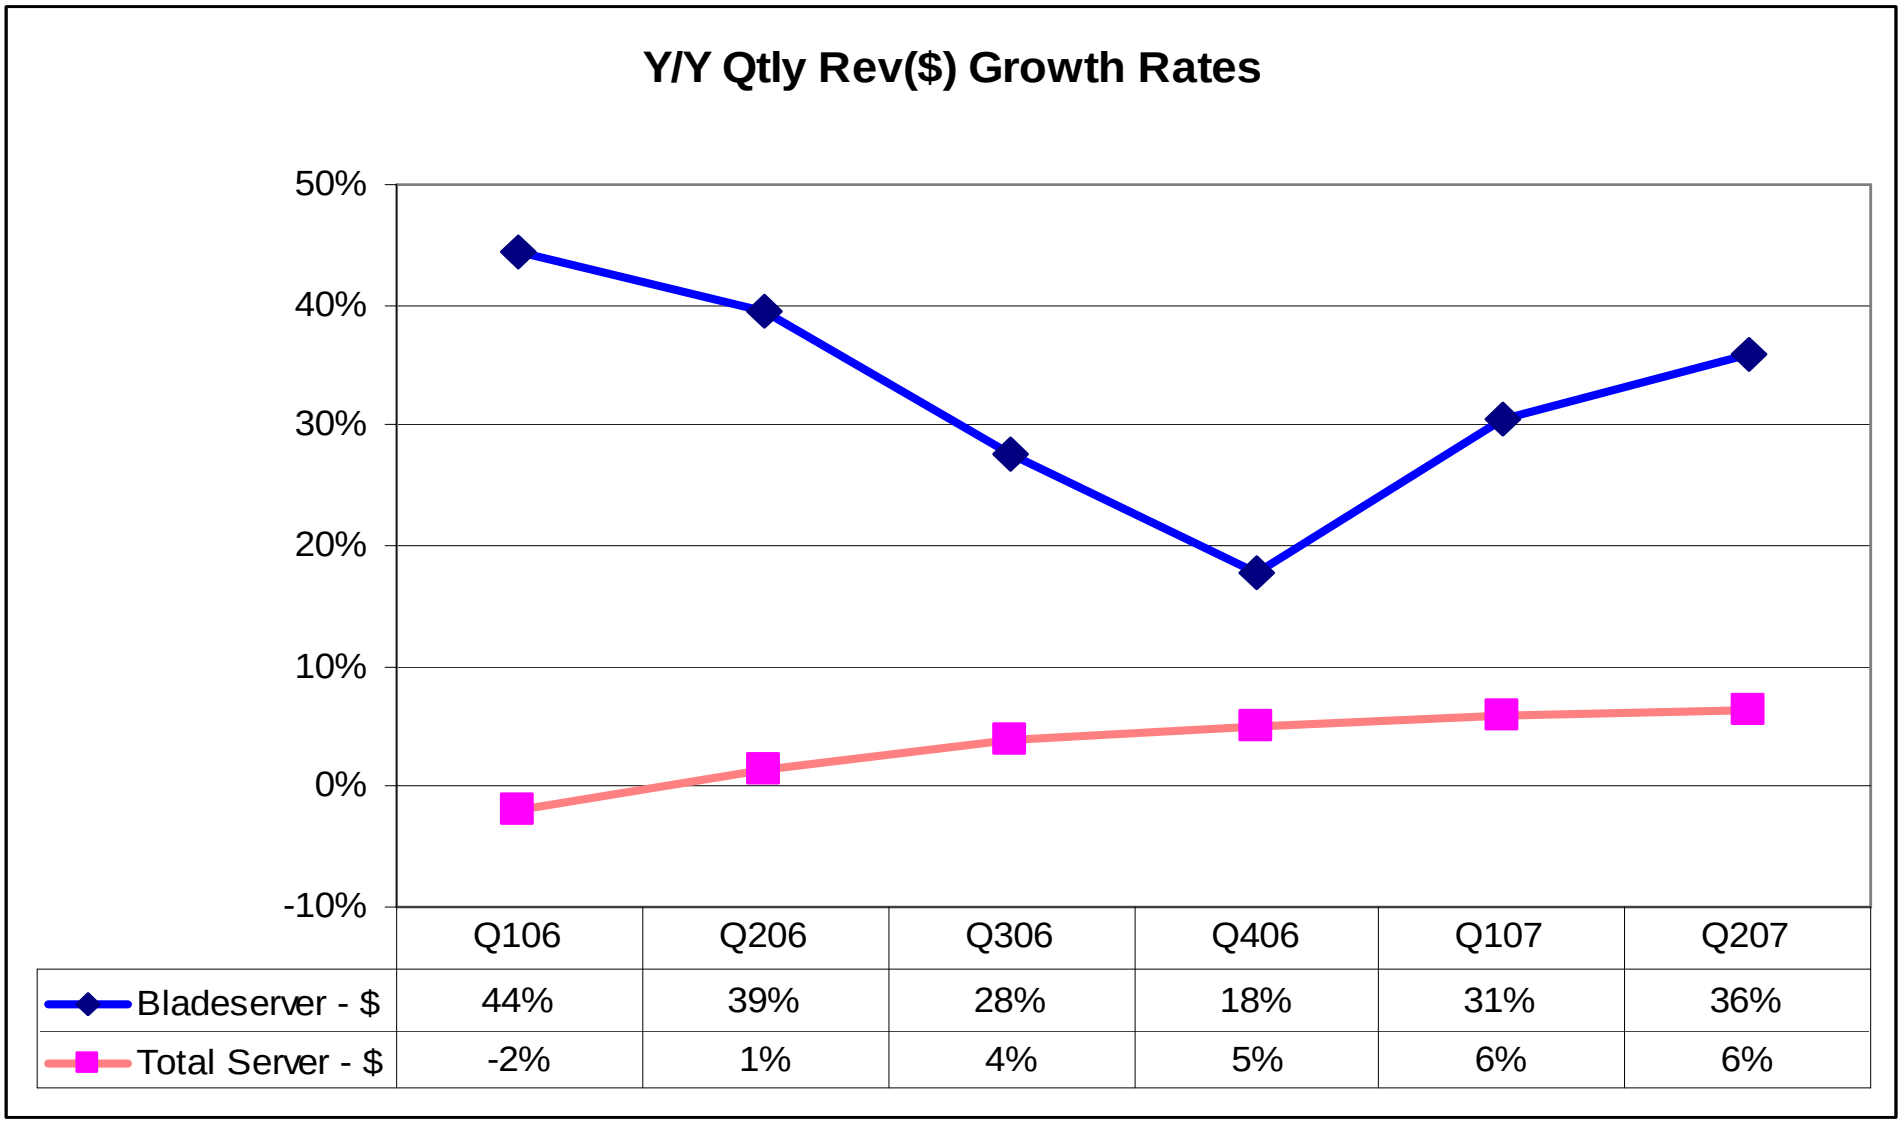

	Q106	Q206	Q306	Q406	Q107	Q207
Bladeserver - $	44%	39%	28%	18%	31%	36%
Total Server - $	-2%	1%	4%	5%	6%	6%

- **Bladeserver market growing much faster than total server market**
- **Significant Market already - Q2'07 run-rate is ~$3.7 billion**

Source: Aug 2007 IDC

Blade Server

Drivers and State of Blade Server Deployments

Biggest DC Concerns	How Bladeservers Help
• Space	• 38% savings*
• Power	• 23% savings*
• Cooling	• 23% savings*
• Dynamic Provisioning / SOA	• Cable once, add/ replace servers
• Increase Utilization	• NA

Massive Penetration in Financial Vertical
(Other Segments/Markets are following them)

Goldman Sachs
Morgan Stanley
LEHMAN BROTHERS
CREDIT SUISSE
BARCLAYS
Merrill Lynch
citi
UBS
A Passion to Perform. Deutsche Bank
JPMorganChase
WELLS FARGO
Bank of America

** 2006 IDC Report - compares HP's c-Class with HP's DL rack servers*

Current Bladeserver Vendor Market Shares

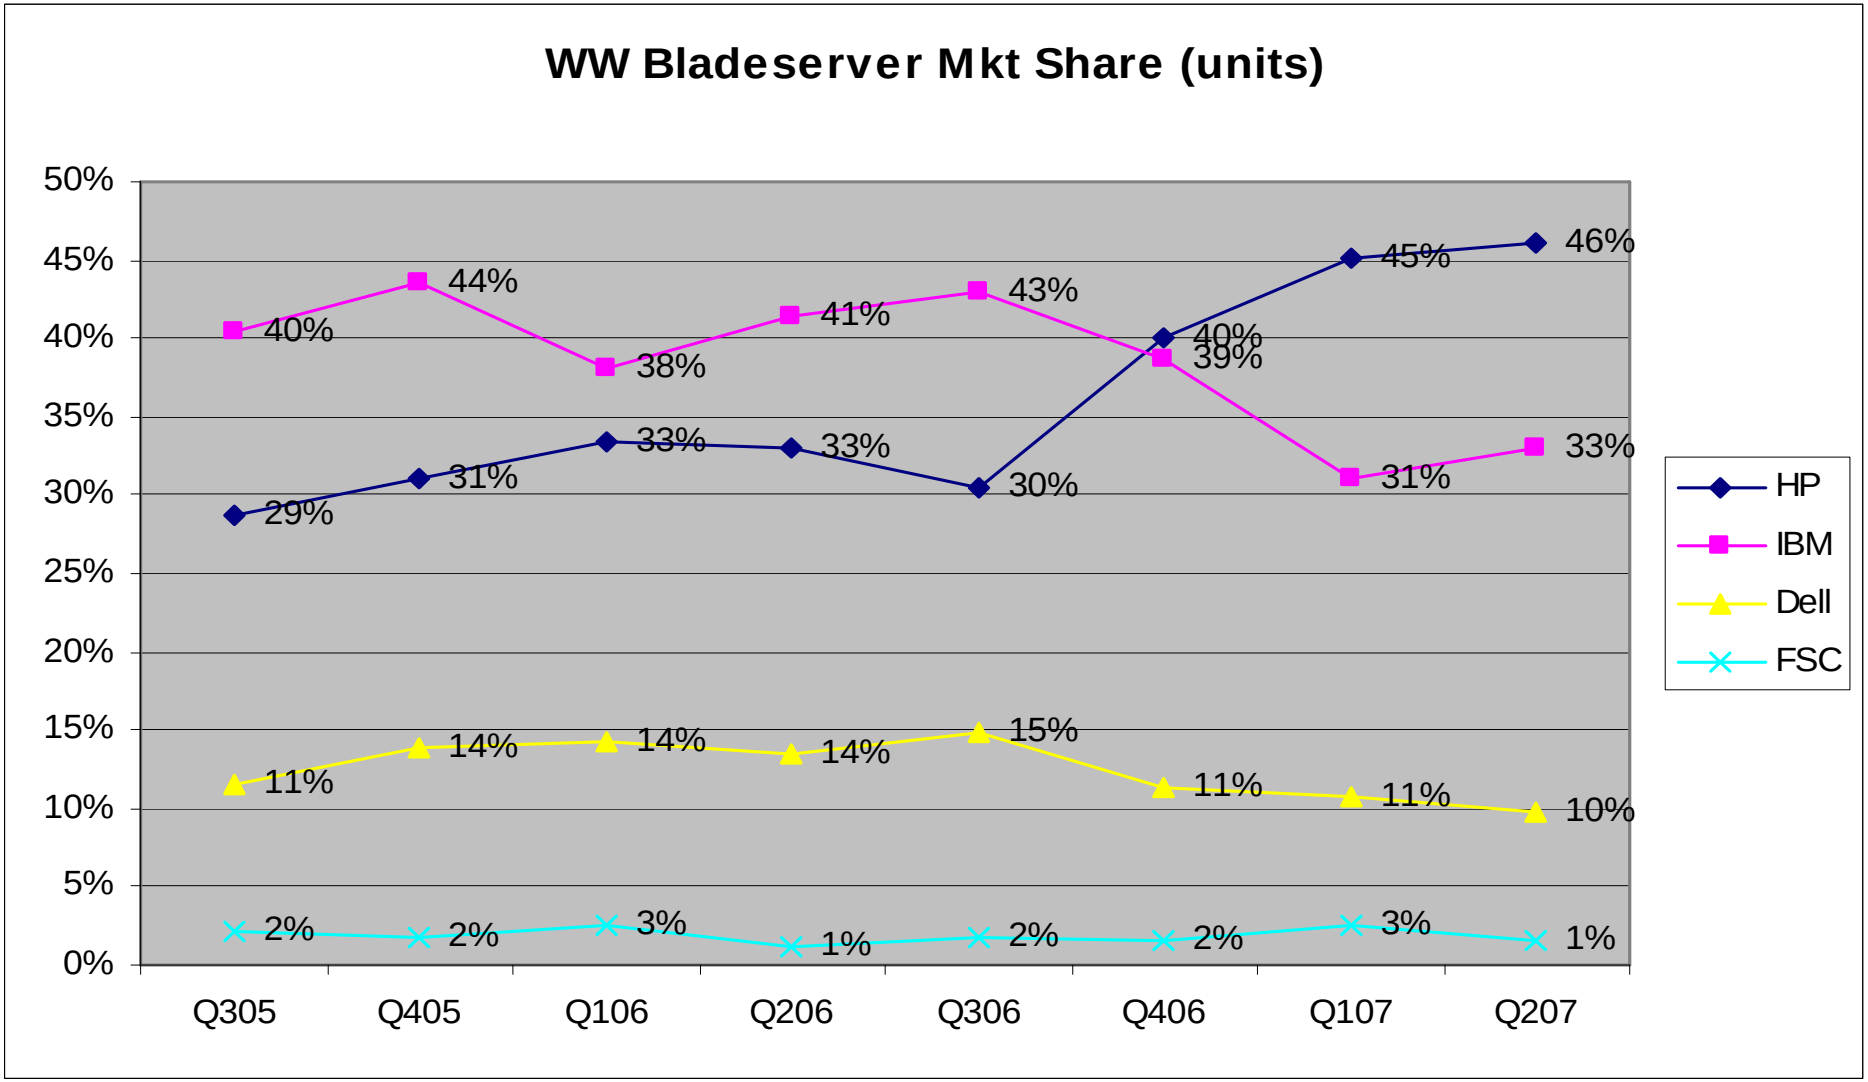

Source: Aug 2007 IDC

Blade Server IO Options

Key Alternatives / Players

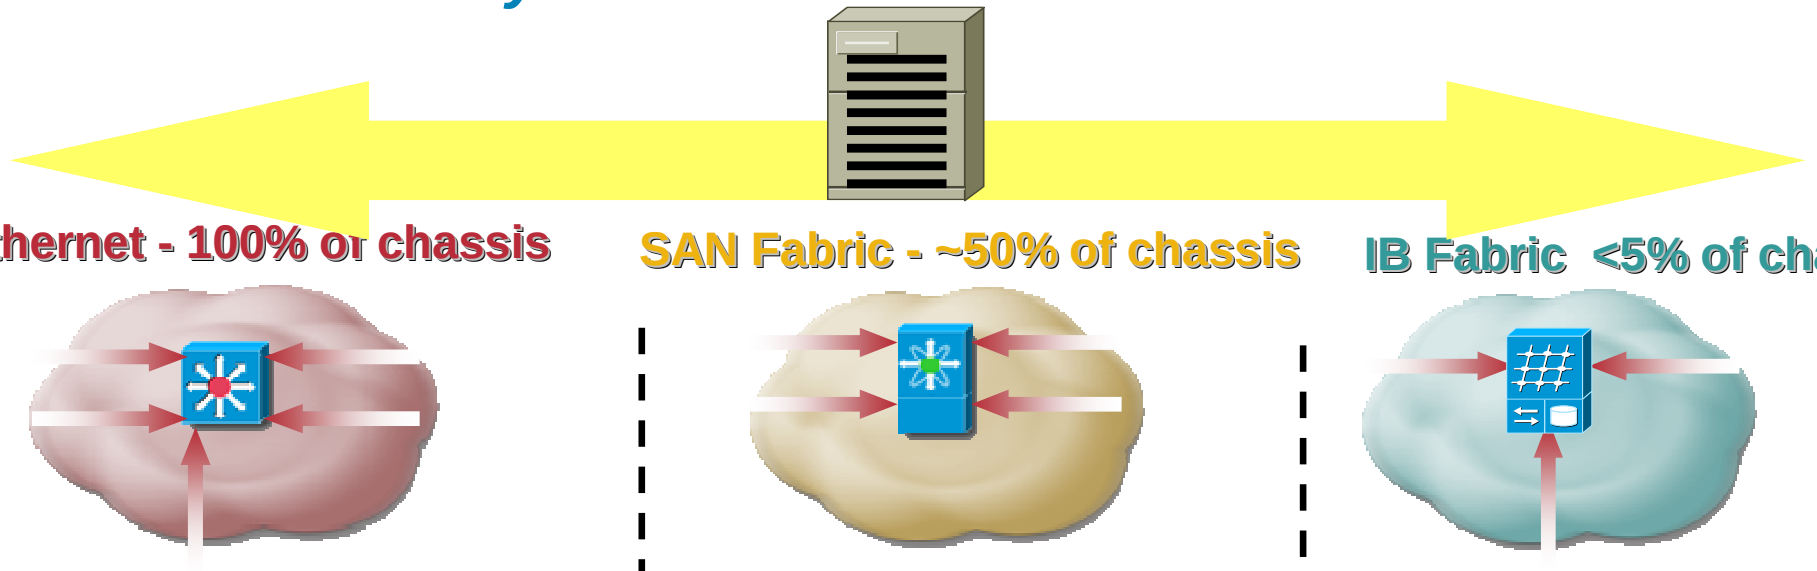

Agenda

- **Blade Server Market Update**
- **Ethernet Blade Portfolio**
- **Fibre Channel & Infiniband Offerings & Roadmap**

Ethernet Blade Switch Roadmap

Comprehensive Offering Today, with Strong Roadmap

Committed

Planning

2007 | **2008/9**

IBM

IGESM

DC Focused

- *NG 10GE VBS**
- *NG GE VBS**
- *Entry Level GE*

10GE to Server

HP invent

CBS3020 for c-class

DC Focused

- *NG 10GE VBS**
- *NG GE VBS**
- *CBS3020*

10GE to Server

DELL

CBS3030

DC Focused

- *NG 10GE VBS**
- *NG GE VBS**
- *Entry Level GE*

10GE to Server

FUJITSU SIEMENS COMPUTERS

CBS3040

**Virtual Blade Switch*

Cisco IOS Release Software Roadmap for CBS30X0

Spring '07 EC
12.2(35)SE

FCS – May CY'07

Key Image Features:

- **LLDP-MED**
- **VLAN Flex Links load balancing**
- **DHCP Snooping stats**

Summer '07 EC
12.2(40)SE

FCS – Q3 CY'07

Key Image Features:

- **PVLAN**
- **Config Rollback**

Just Released!

Fall '07 EC
12.2(XX)SE?

FCS – Q4 CY'07

Key Image Features:

- **CBS3020 moves to IP Base**
- **IPv6 Host Support**

Security:
- *IP Source Guard*
- *Dynamic ARP*
- *RACLS*

Basic Routing
- *Static Routes*
- *RIP*
- *EIGRP Stub*

FCS = First Customer Ship EFT = Early Field Trials EC = Execute Committed CC = Concept Committed NC = Not Committed

PVLAN in Summer 07,12.2(40)SE1

Key Differentiator against Competition

Benefit A Private VLAN is a way to provide layer 2 isolation between target hosts in the same subnet … (segmentation of IP space can waste lot of addresses)

Deployment Scenario: Servers belonging to different departments or customers can be isolated inside the same blade chassis

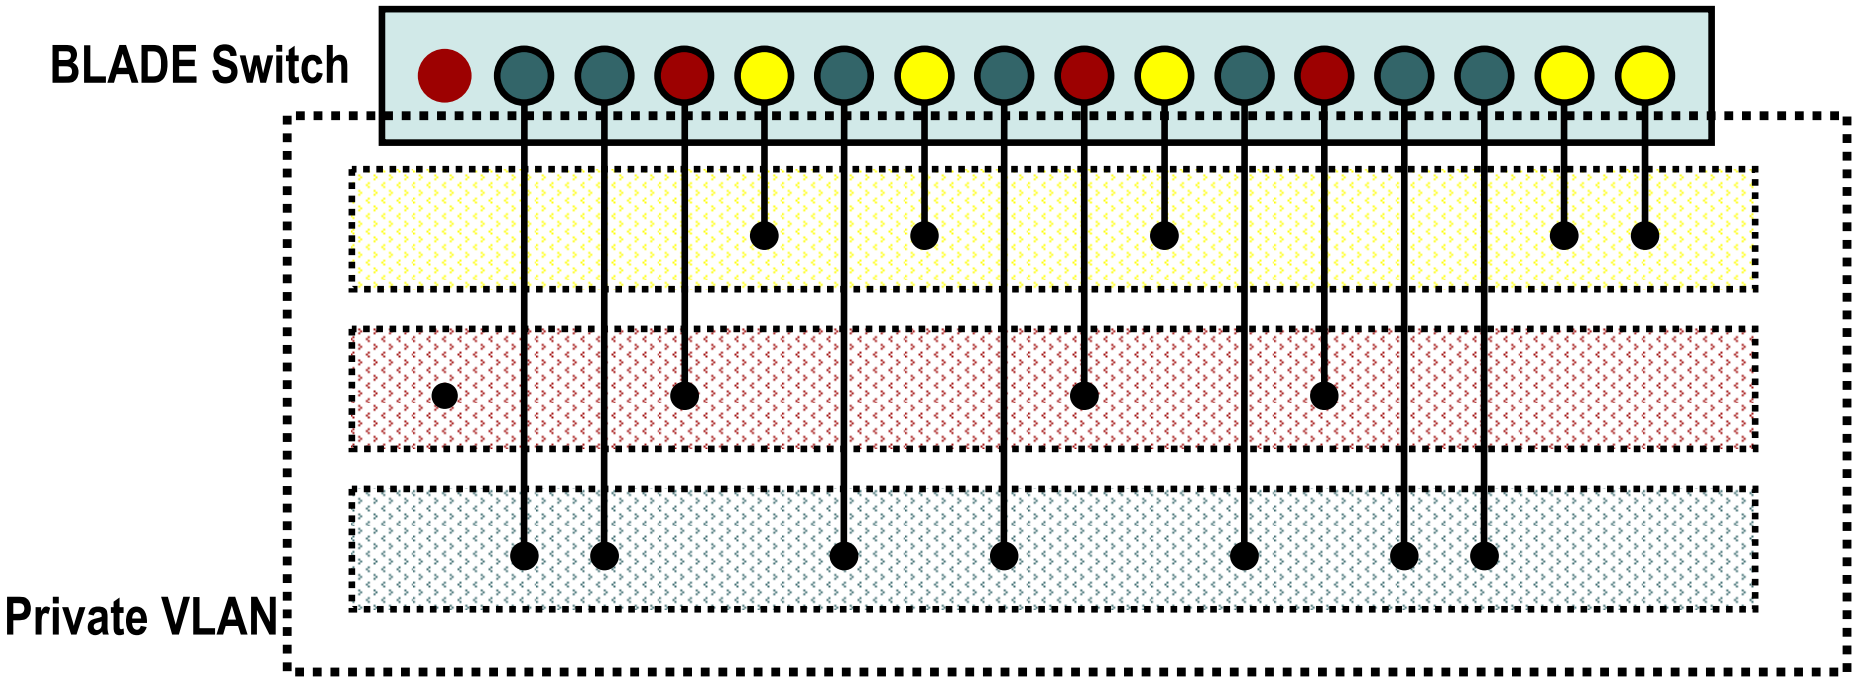

Many Customers Waiting for this Feature

Cisco Network Assistant (CNA 5.0) on Ethernet Blade Switch

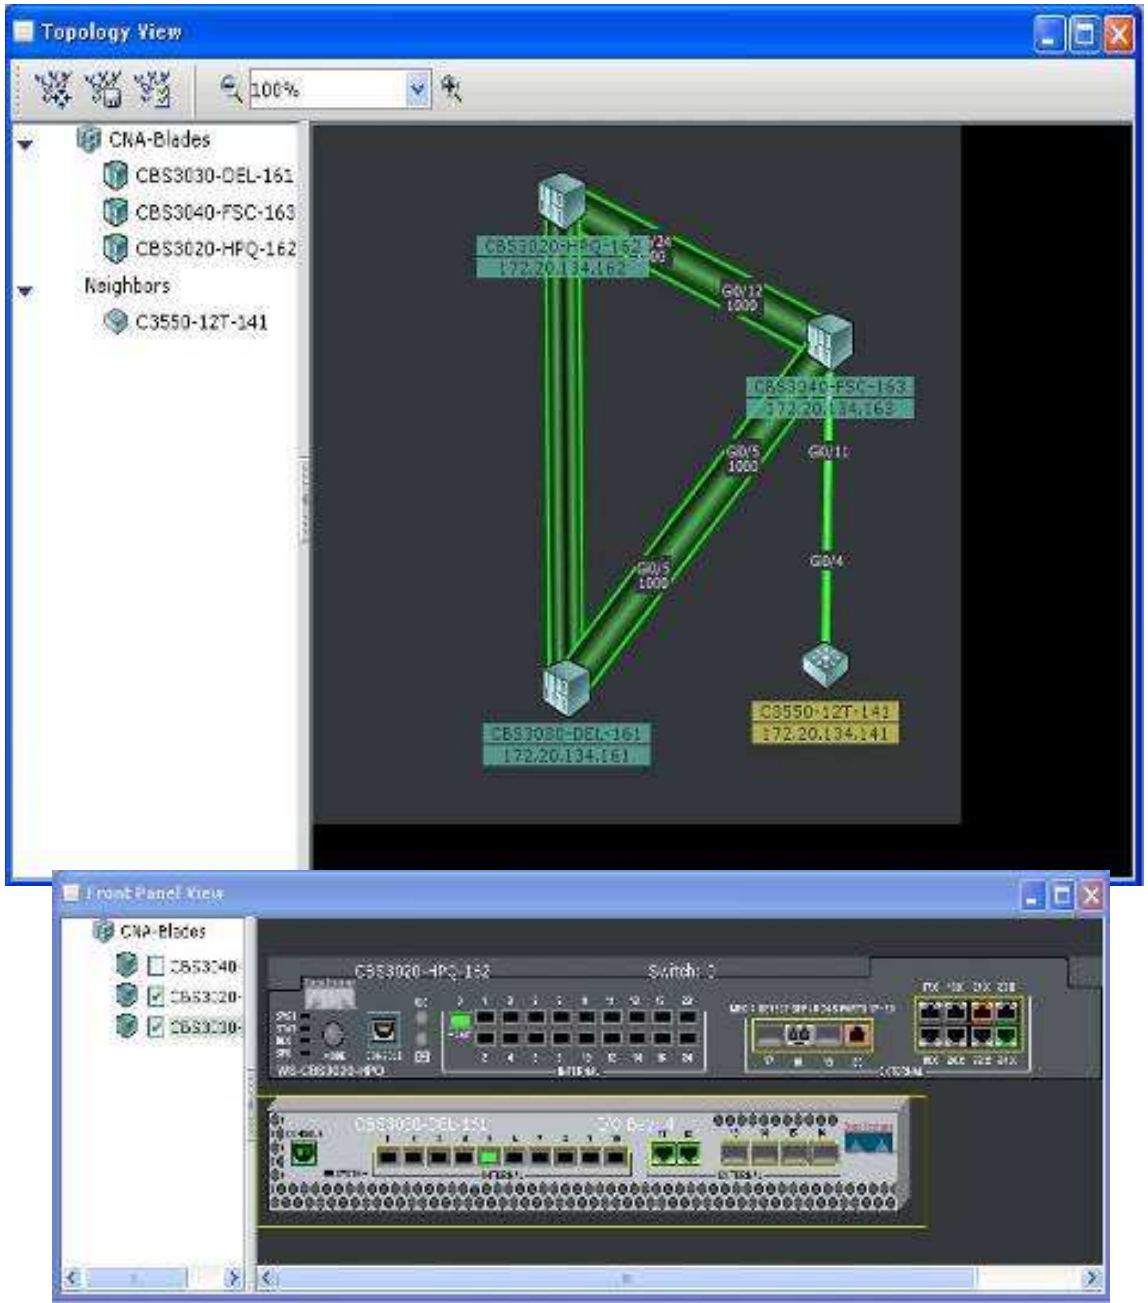

Benefit:

Manage multiple blade switches easily

Cost:

Free Tool

What:

- Mgmt Tool to allow Interactive Configuration, Topology/Front Panel View, Monitoring, Troubleshooting and Network Maintenance
- Supports up to 40 switches
- Examples:
  - Apply multiple port configs on multiple switches
  - Health Monitoring
  - OS upgrades

Feedback from Ethernet Blade Server Customers

- Switch Proliferation

  Increases Network Complexity - Larger L2 and L3 domains where convergence and stability are concerns

  Increases operational Management Headaches

- Growing Bandwidth needs

  Server virtualization and multi-core CPUs are driving more bandwidth from each server and out of the rack.

- Uplink cost

  Each Blade Switch needs at least one 10GE Uplink or multiple GE Uplinks

- Feature Consistency with DC Catalyst Products

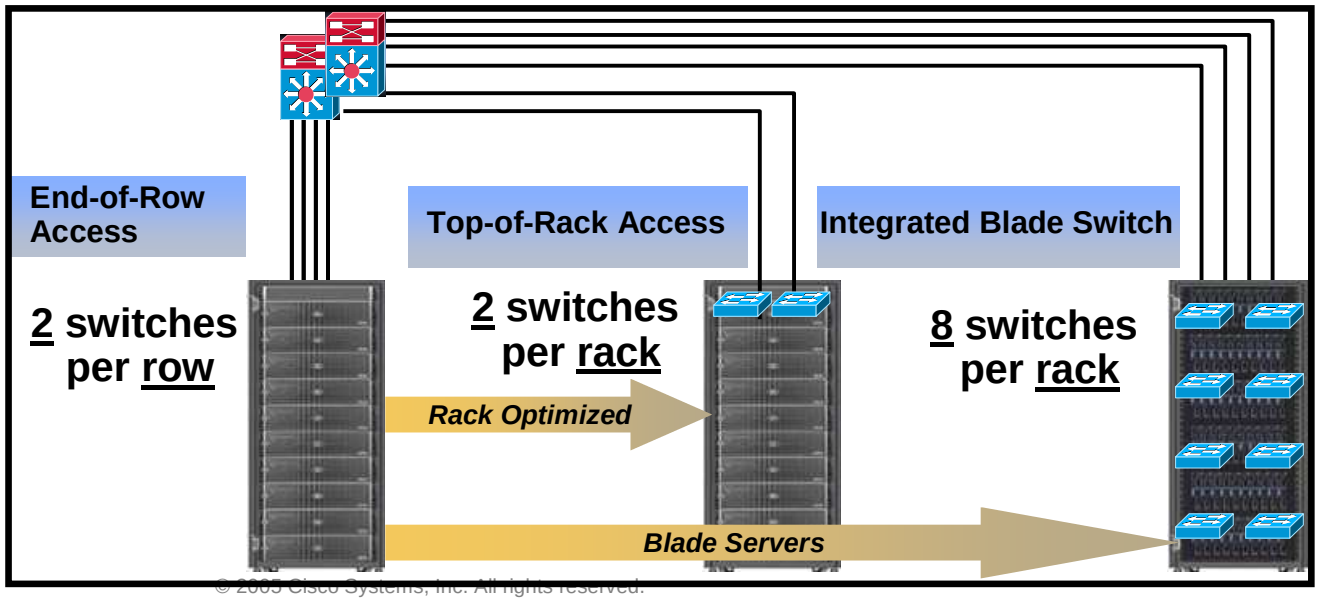

Introducing Cisco Virtual Blade Switch (VBS)

Overview of Concept and Benefits

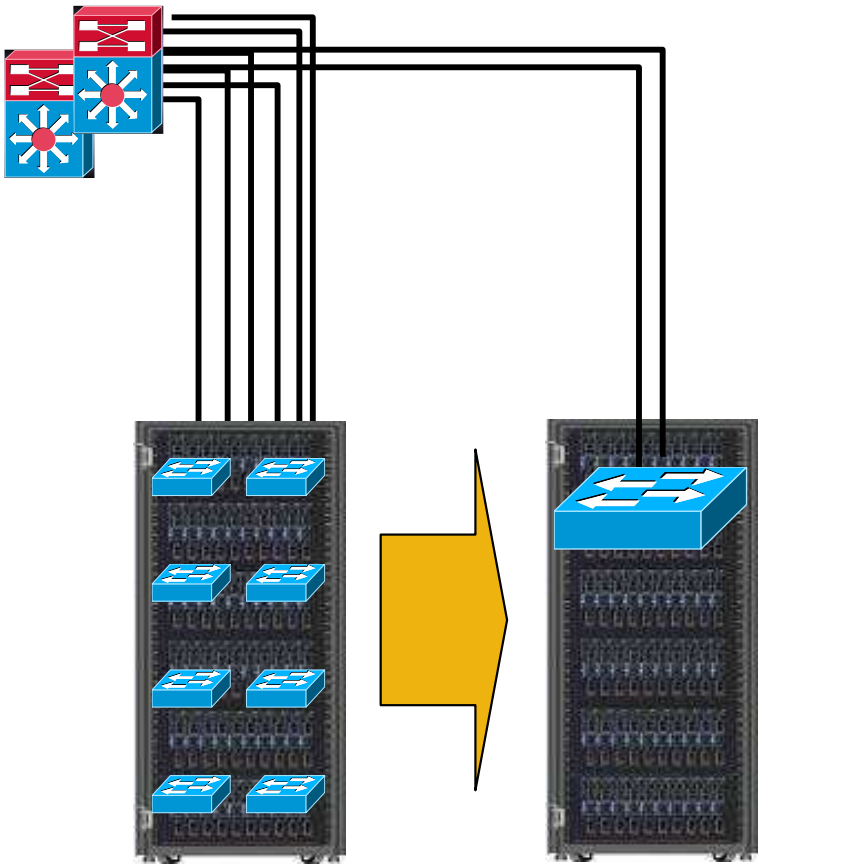

FCS Target:
HP - Dec '07
Dell – Jan '08
IBM – Q1 '08

- *Management Simplification*
  - Operational simplification
    - **Single switch per rack to manage**
    - **True Plug-n-Play of switches**
  - Design Simplification:
    - **Sharing Uplinks helps reduce cables**
    - **Reduction in # of logical nodes in L2/L3 network helps improve network convergence**
  - Operational Consistency
    - **Familiar IOS CLI, MIBs and management tools like CiscoWorks**
    - **Consistent End-to-end features and functionality**
- *Performance & Scalability*
  - **Up to 160G configurable bandwidth out of rack**
  - **VBS allows server to double bandwidth with no additional cost**
- *Reduce CAPEX*
  - **Mix-n-match of GE & 10G uplink switches**
  - **Sharing of uplinks reduces overall network infrastructure cost**

Cisco Catalyst Virtual Blade Switch (VBS)

Switch Overview

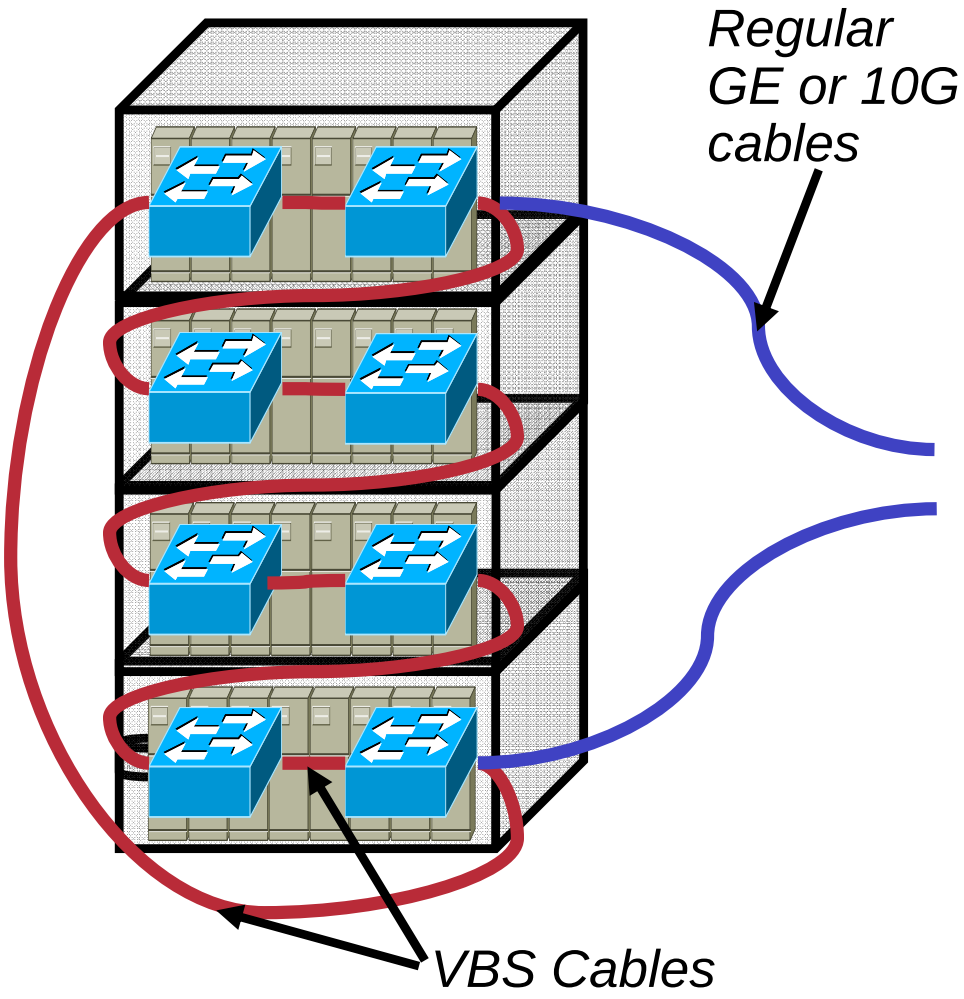

- ***Connect up to 8 switches in a Rack = Equivalent to redundant TOR solution***
- ***Share Uplinks across the rack – GE and/or 10GE***
- ***High Speed VBS Cable***
  - ***– 64G raw bandwidth***
  - ***– Special connector & cable designed for VBS***
  - ***– 3 cable lengths (0.5m, 1m & 3m)***

“Single VBS Switch” Solution – Easy to Manage

Cuts Down Switches to Manage by Factor of 8

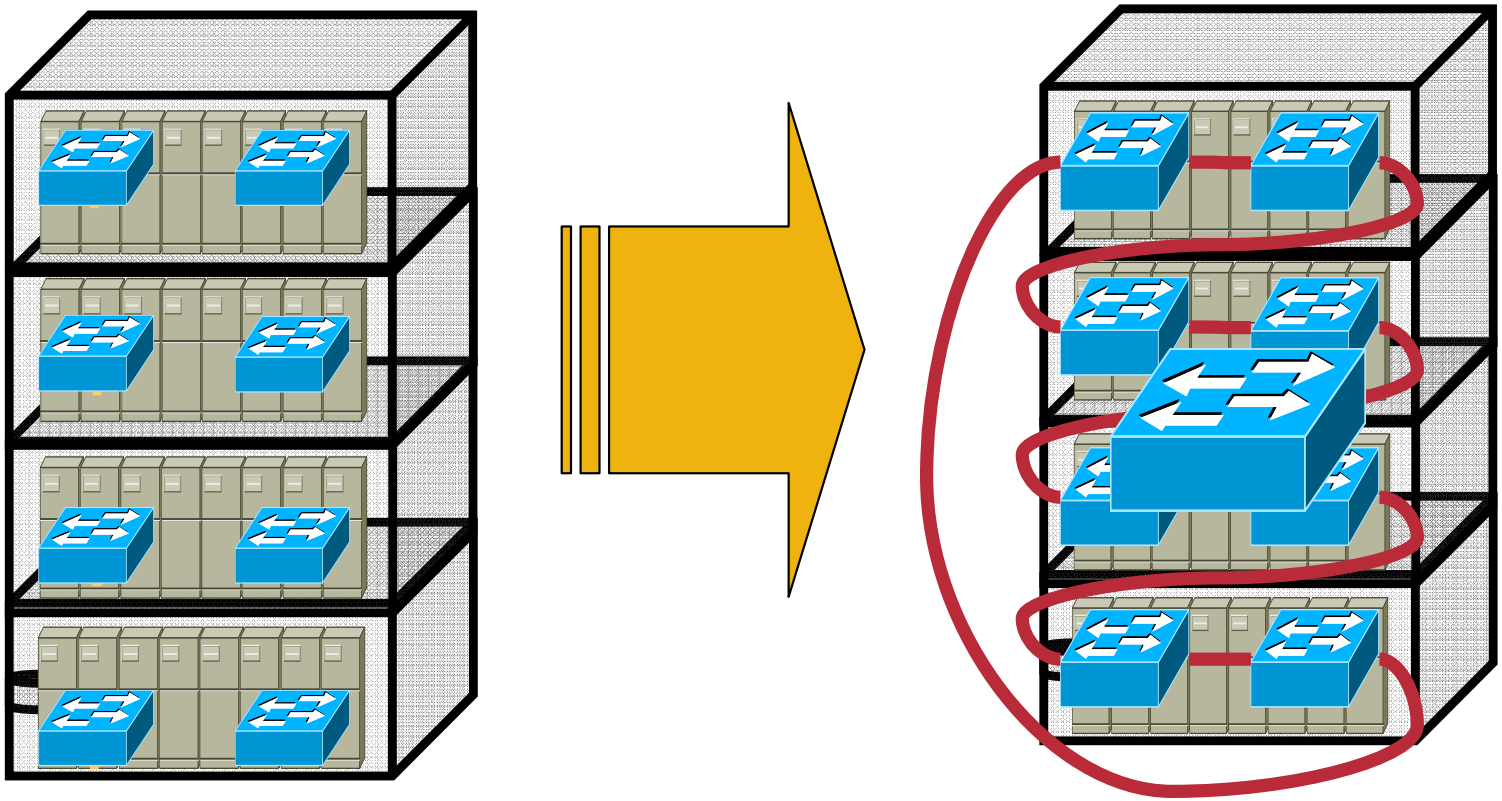

- ***One Logical Switch Now (instead of 8)***
- ***Single IP Address for Entire Rack (for management, MIBs etc)***
- ***1 Node for Spanning Tree (Layer 2)***
- ***1 Node for Routing (Layer 3)***

“Pay-as-you-go Network Infrastructure”

Customers Can Gracefully transition to 10G Uplinks

Mixed Rack of GE and 10G Switches

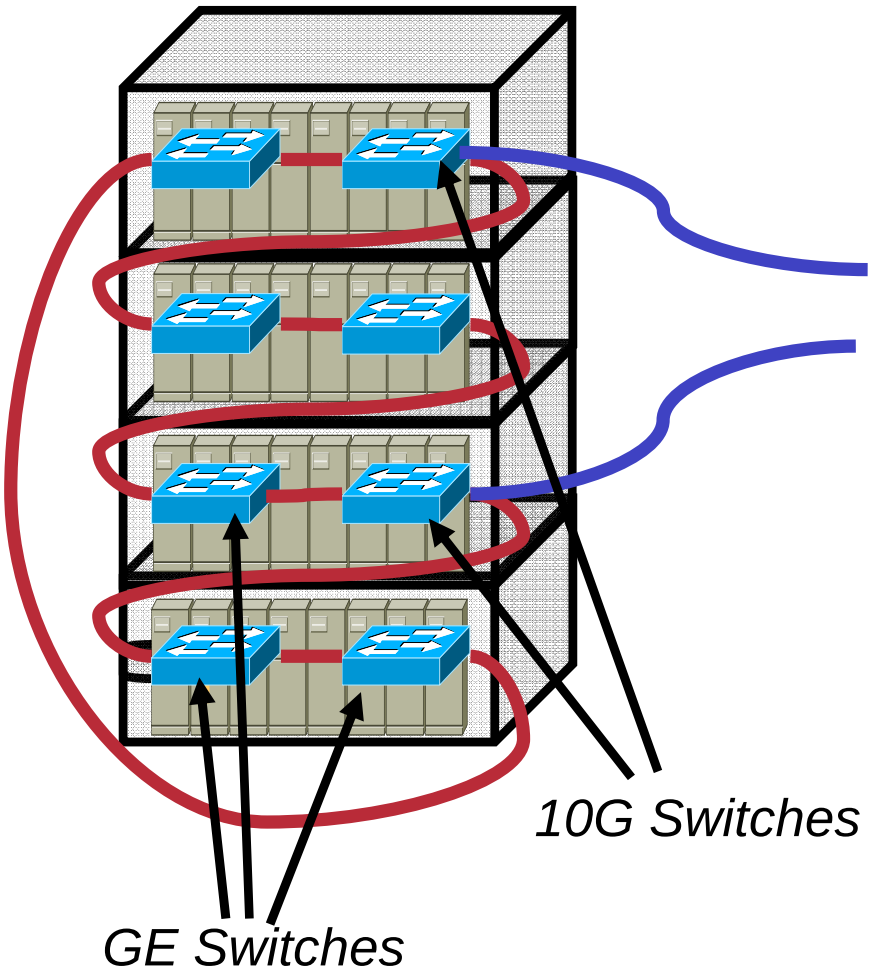

- ***Customers can start with GE switches and add 10G switches when bandwidth need grows***
  - ***Mix GE and 10G switches allowed within a single VBS***
- ***10G X2 ports can be used as GE ports with TwinGig module adapter and SFPs***
  - ***Migrate to 10G X2 modules when bandwidth need grows***

“Multiple Deployment Options for Customers”

Caters to Different Customer Needs

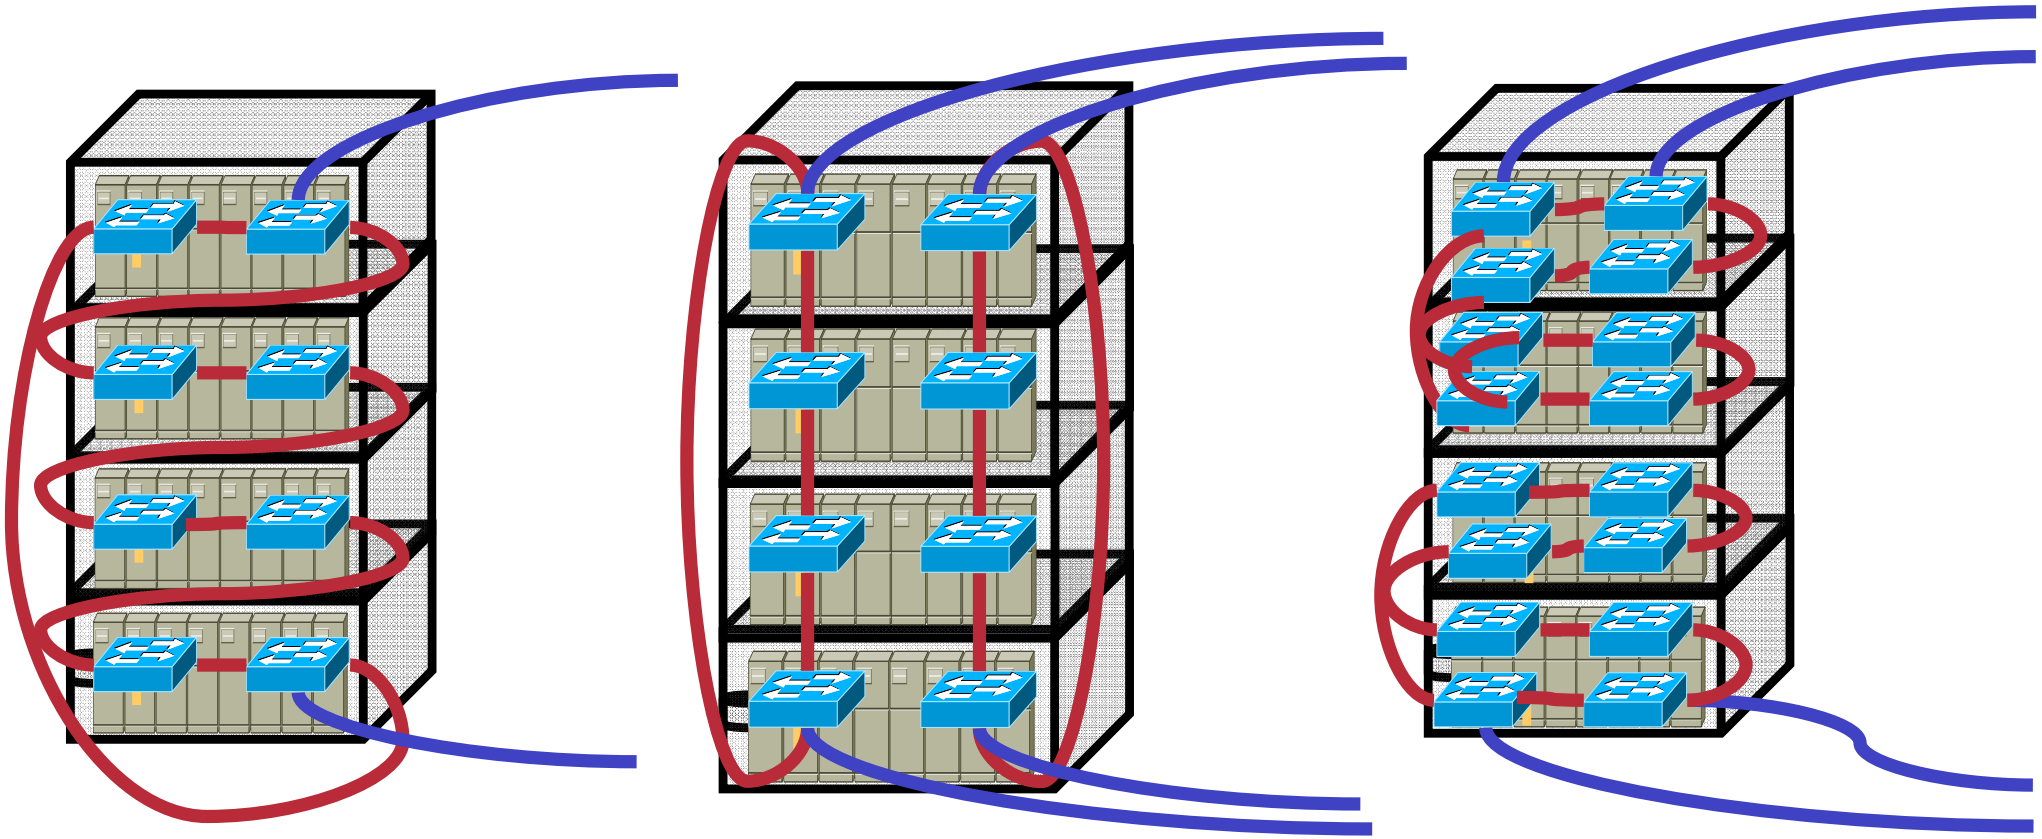

Common Scenario

- ***Single Virtual Blade switch***
- ***Cost Effective***

Redundant ToR Scenario

- ***Separate VBS***
- ***More resilient***

Multi-NIC Scenario

- ***More Server Bandwidth – e.g. VMware***

X2/SFP Optics Support for CBS31xx Family

	Copper	Multi-mode Fiber	Single-mode Fiber
SFP	GLC-T (100m)	GLC-SX-MM – (550m)	GLC-LH-SM – (10km)
X2	X2-10GE-CX-4 (15m)	X2-10GB-SR – (26m) X2-10GB-LRM – (220m) X2-10GB-LX4 – (300m)*	X2-10GB-LR – (10km)*

Note: LX4 and LR modules will not be supported at FCS. We are currently planning on support for these Post FCS. Please let us know the impact of not supporting these.

Cisco Catalyst Virtual Blade Switch

Virtual Blade Switch Architecture

- ***Up to 8 Switches acts as Single VBS Switch***
  - ***Distributed L2/ MAC learning***
  - ***Centralized L3 learning***
- ***Each switch consists of***
  - ***Switch Fabric***
  - ***Port Asics (downlink & uplink ports)***
- ***One Master Switch per VBS***
  - ***1:N Resiliency for Master***
  - ***L2/L3 reconvergence is sub 200 msec***
- ***High Speed VBS Cable (64 Gbps)***

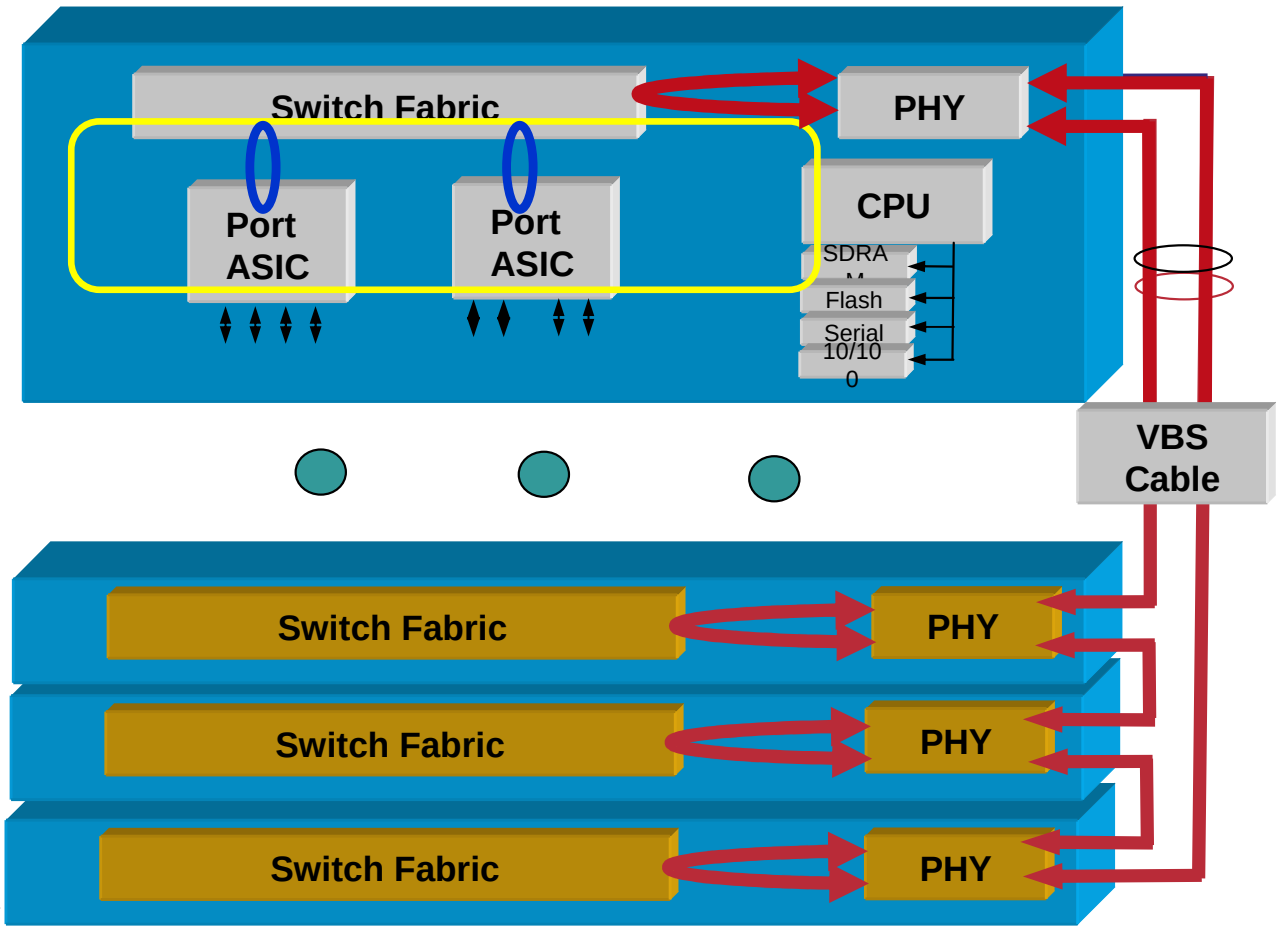

High Performance, Scalable and Manageable Architecture

Day in the Life of A Packet: Unicast Packets

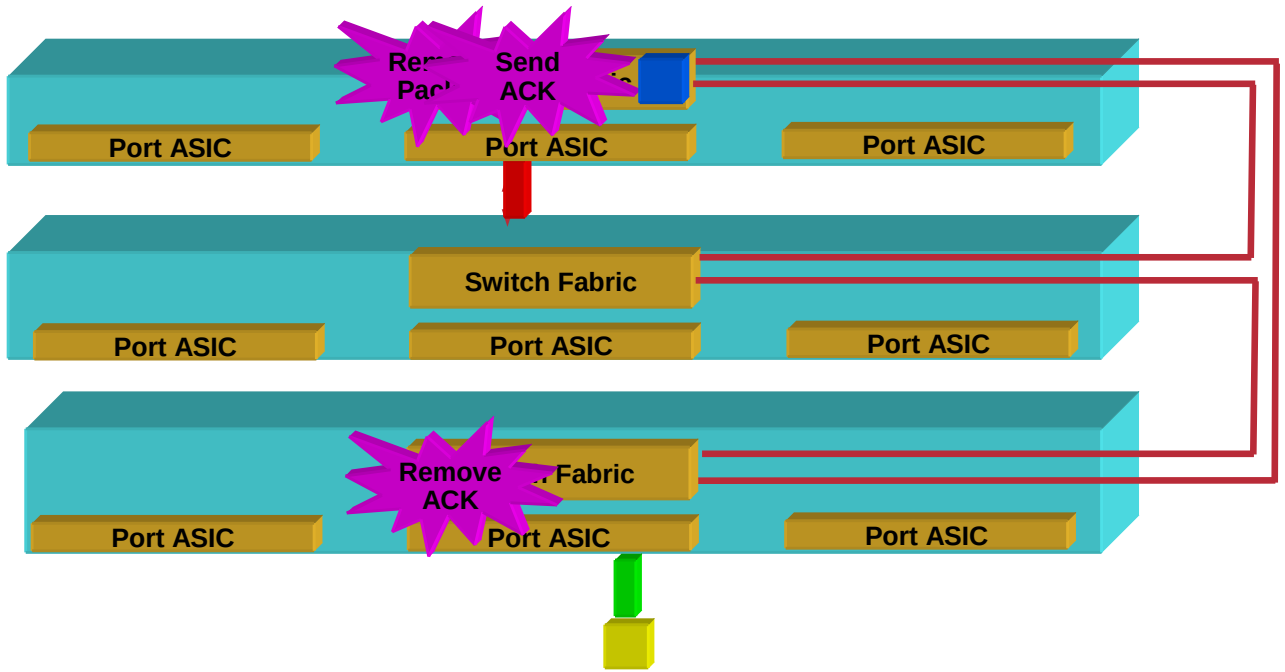

- The Source Port ASIC sends the packet to the Source Switch Fabric and it is switched to the Destination Switch Fabric
- The Switch that contains the destination Virtual Node sends an ACK back to Sourcing Switch. (16 Byte Packet).
- The Originating Switch Fabric receives and removes the ACK

Source
Destination
Data
ACK

Day in the life of a Packet: Multicast Packets

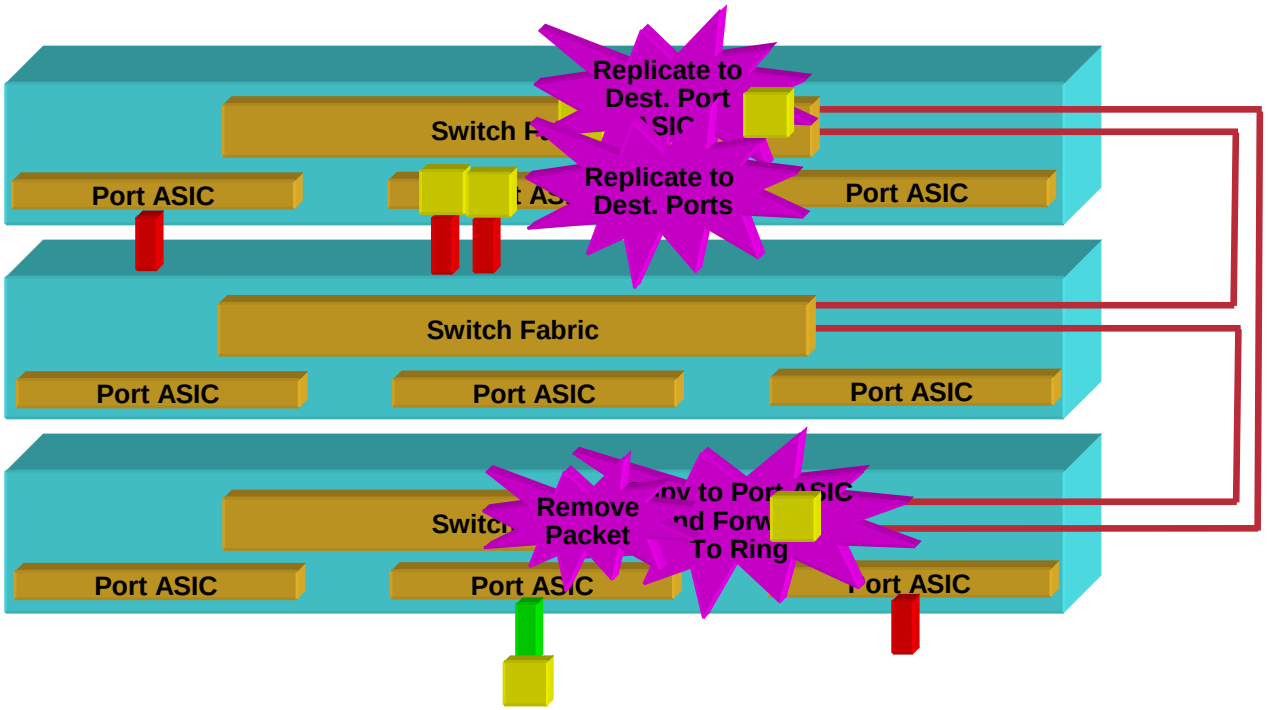

- The packet is passed all the way around the ring
- Switch Fabric with Multicast ports in that group copy the packet
- The originating Switch Fabric removes the packet from the ring
- Note: There is only one packet on the ring per multicast flow, replication only occurs at the local level
- Note: if the sender and all of the receivers are on the same switch no packets are sent to the ring

Ring Capacity

- **Ring raw data rate is 64 Gbps**
- **Packets have 24 Bytes Header and size is rounded up to a multiple of 16 bytes. Therefore a 64 byte packet becomes 96 bytes on the Ring. (64+24 rounded up).**
- **Packets need ACKs (16 Bytes)**
- **Therefore, for 64 Byte packets, the overall ring capacity is reduced to 48 Gbps**

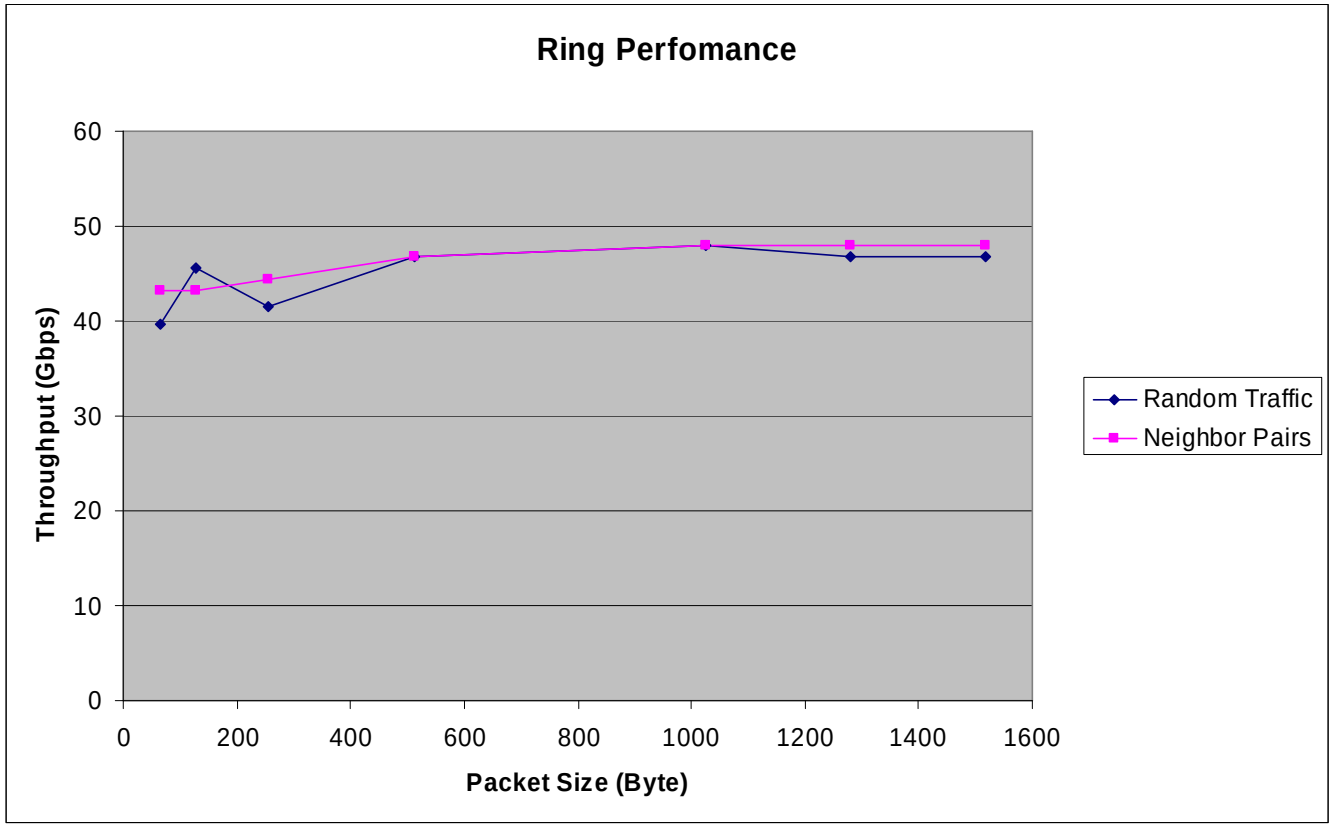

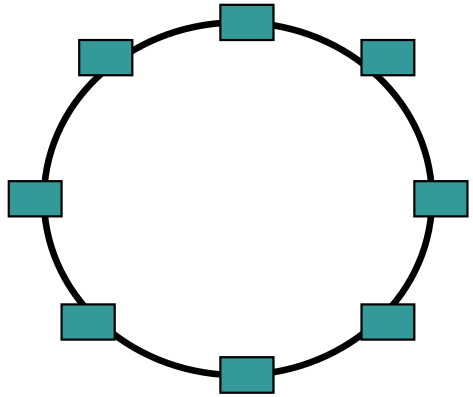

8 Switches
12 ports each

Blade Ring Latency test and Data

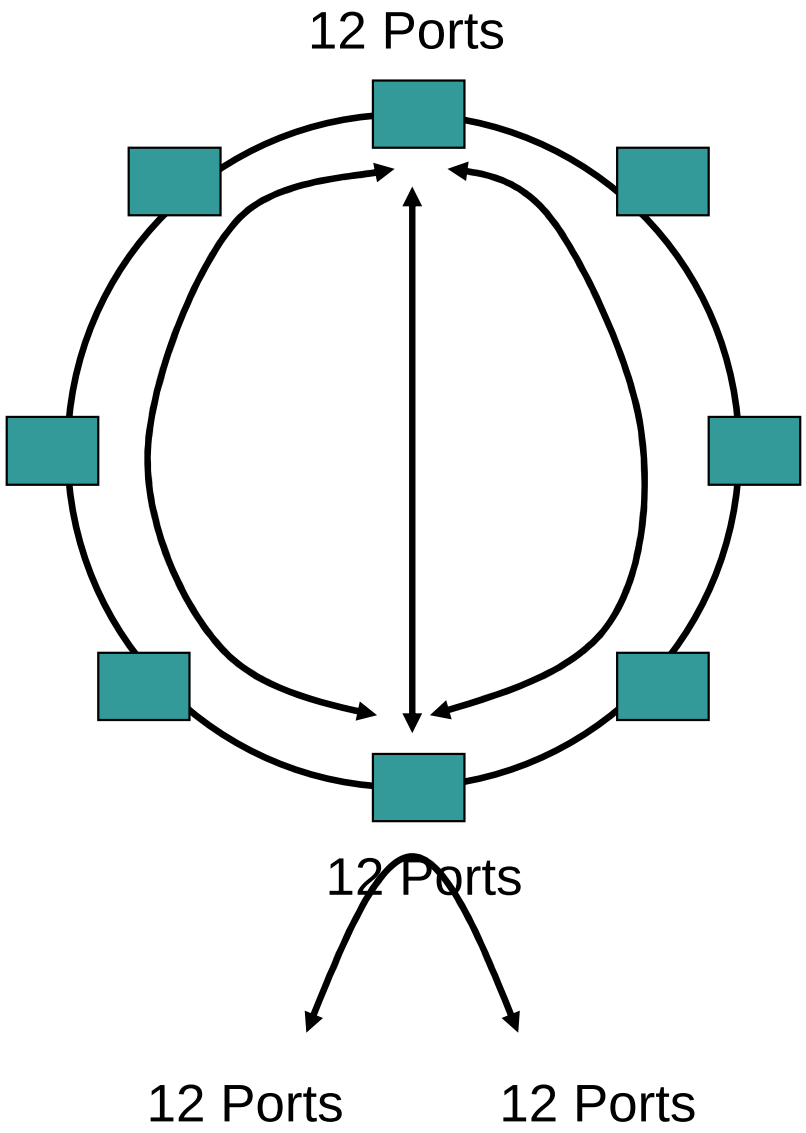

Frame Size	Internal	2 Switch ring	4 Switch ring	8 Switch ring
64	5.6	5.79	6.11	6.7
128	6.44	6.97	7.14	7.69
256	8.93	9.14	9.23	9.67
512	12.35	12.32	12.31	12.8
1024	18.16	19.02	18.27	16.03
1280	22.37	22.16	21.78	22.39
1518	25.43	25.39	24.83	25.41

Blade Ring Latency Graph

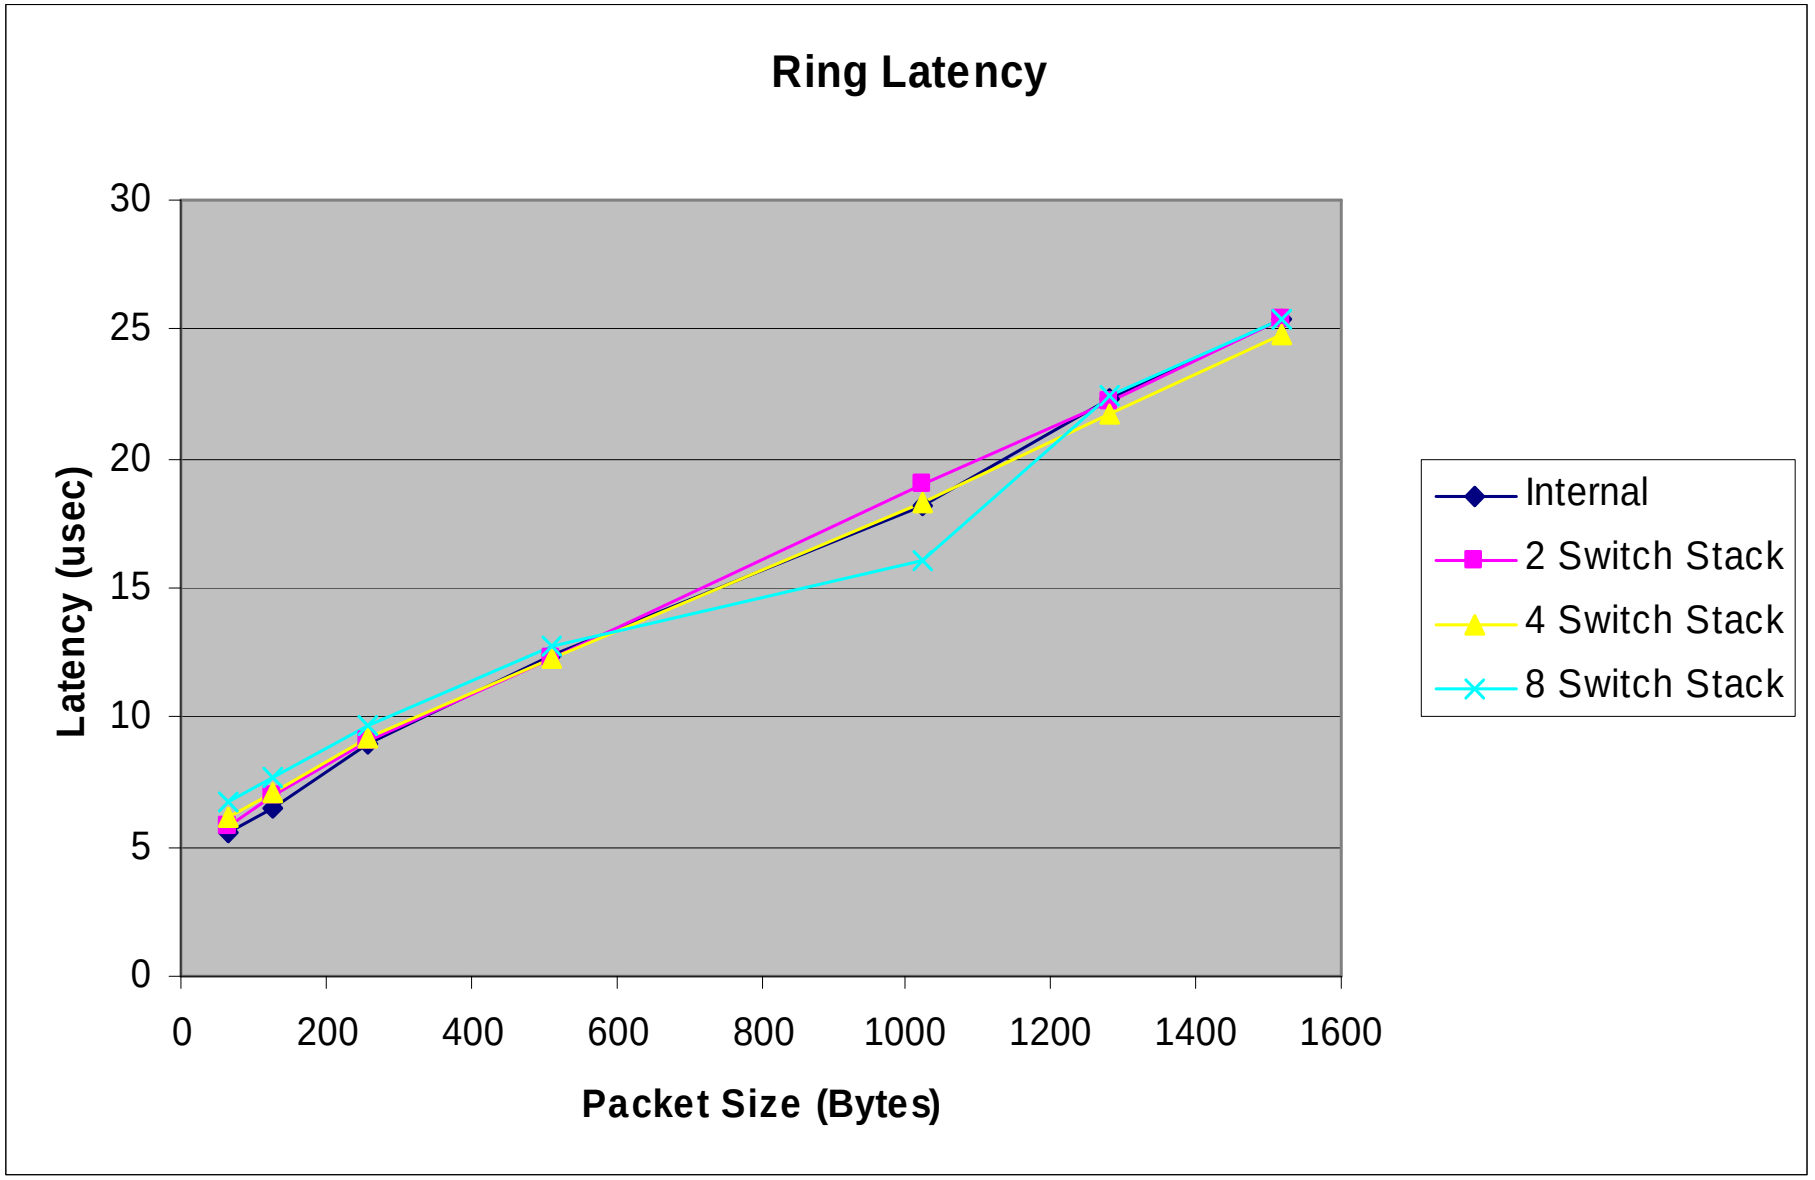

CBS3110 and CBS3120 Product Specifics

CBS3110 Product Overview

- **Works in IBM BC1, BCT, BCH and BCHT Enclosures**
- **Uses Low Speed Slots**
- **Two Product SKUs:**

Both SKUS have:

14 1000Base-X Server Backplane connections

Fa0 Management Interface connected to AMM

2 High Speed Ring Connectors

“USB type connector” Console Port

WS-CBS3110G-E

4 GE Uplink Ports

WS-CBS3110X-E

1 10GE X2 Uplink port

CBS3120 Product Overview

- **Works in HP c-Class Enclosure**
- **Can be Installed in any pair of slots**
- **Two Product SKUs:**

Both SKUS have:

16 1000Base-X Server Backplane connections

Fa0 Management Interface connected to OA

2 High Speed Ring Connectors

RJ45 Console Port

4 Copper GE Uplinks

WS-CBS3120G-E

4 SFP Uplink Ports using TwinGig Modules

WS-CBS3120X-E

2 10GE X2 port

Software Licenses

Feature	IP Base	IP Services	Adv IPS
Standard L2+ Feature Set	**X**	**X**	**X**
IP Source Guard and Dynamic ARP Inspection	**X**	**X**	**X**
Rip/Static Routing, EIGRP Stub	**X**	**X**	**X**
IPv6 Manageability	**X**	**X**	**X**
Multicast Routing, OSPF, BGP		**X**	**X**
WCCP		**X**	**X**
IPv6 Routing			**X**
IPv6 ACLs			**X**

Catalyst Blade Switches (CBS31x0) ship with IP Base Standard

CBS31x0
Ring Features

Autorecovery for Cable (or Switch) Failure

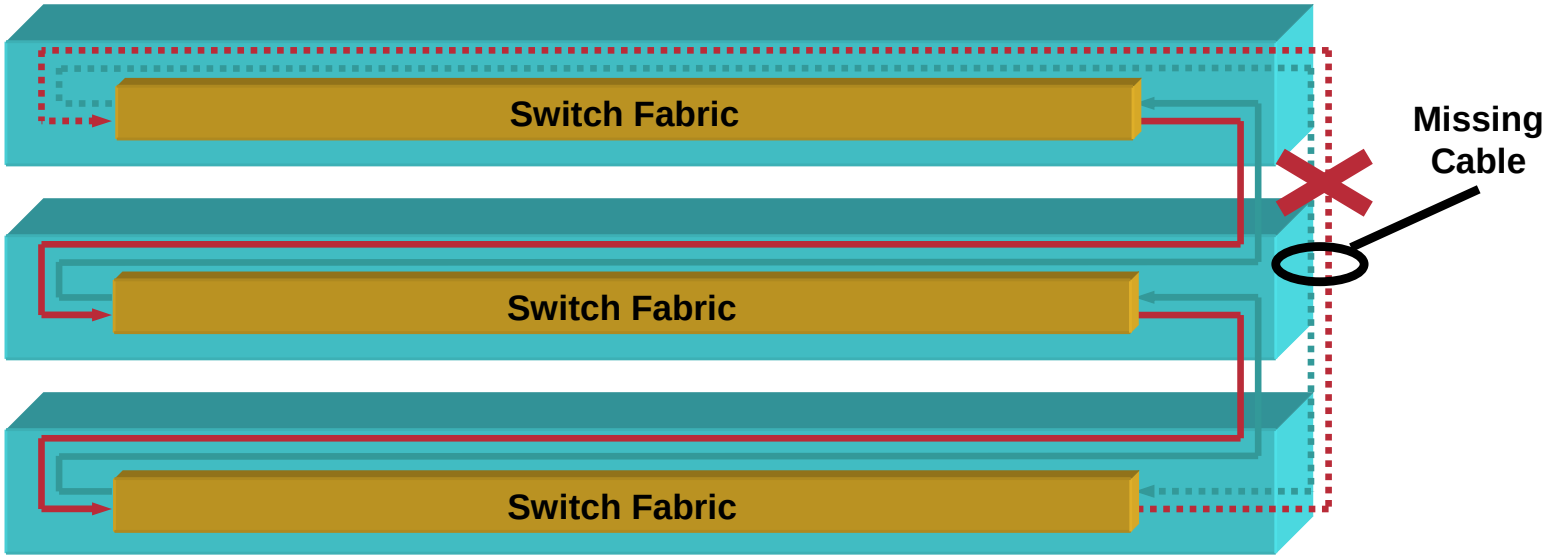

- The Switch Fabric closest to cable detects link down
  - Criteria is coding violations in a period of time
  - Loss of at most one packet that was being transmitted when ring broke
  - Just microseconds for hardware to detect failure
- Each switch signals a bad link to its ring partner
- Both ends of the cable loop back on themselves
  - 32 Gbps backplane when cable is missing
- Failed Switch is treated the same as failed cable

Auto Configuration for New Switch Members

- New members are added automatically as next available member ID
- Auto-upgrade of IOS if necessary
- Auto-configure from Master Switch

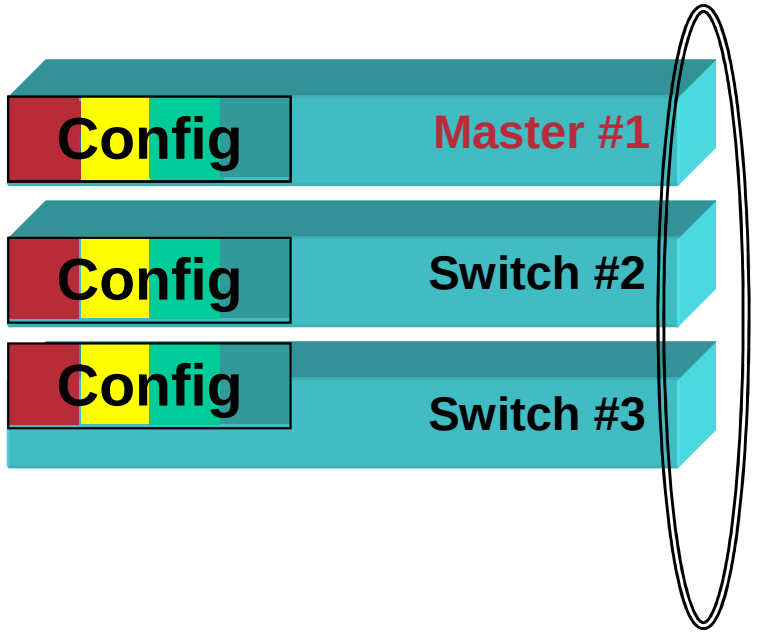

Switch Removal

- **The ring has three members—1, 2, 3**

Switch #3 is removed or powered down
- Neighbor loss is detected by Switches #1 and #2
- Layer 2 and Layer 3 convergence may need to occur
- Now there is a ring of two switches— #1 and #2
- Switch#1 is still the master

Switch #1 is removed or powered down
- Switch #2 takes over as master
- Layer 2 and Layer 3 convergence may need to occur
- Now there is a ring of one switch—#2 which is the master

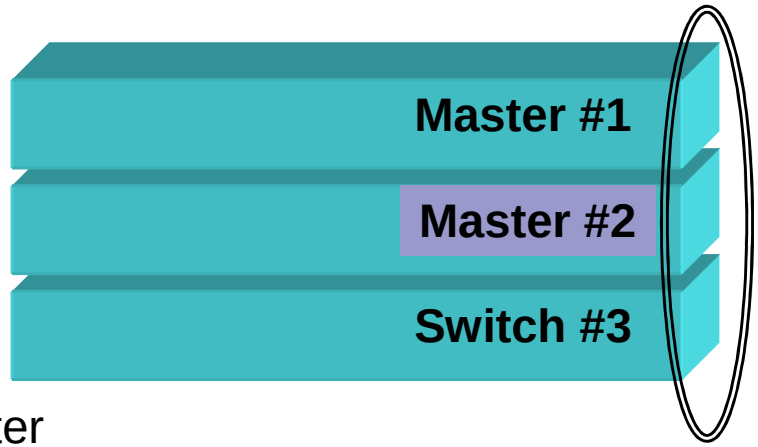

Switch Replacement

- For example, the failed switch with the new switch with the same hardware configuration, then new switch will receive the same port-level configuration of the original unit
- If replaced by a switch with a different hardware configuration, the original configuration is lost and the new switch receives all VBS global configuration

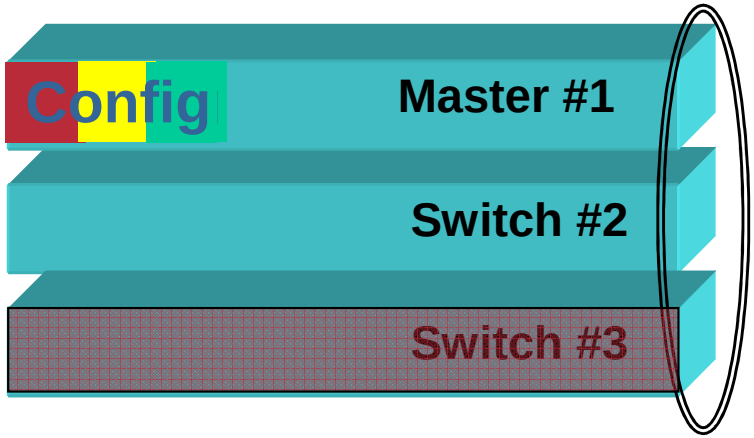

Agenda

- **Blade Server Market Update**
- **Ethernet Blade Portfolio**
- **Fibre Channel & Infiniband Offerings & Roadmap**

Cisco MDS 4Gb Fibre Channel Blade Switches

- 16 internal copper 1/2/4-Gbps Fibre Channel connecting to blade servers through blade chassis backplane
- Up to 8 SFP uplinks
- Offered in 4+8 and 8+16 configurations via port licensing

- 14 internal copper 1/2/4-Gbps Fibre Channel connecting to blade servers through blade chassis backplane
- Up to 6 SFP uplinks
- Offered in 3+7 and 6+14 configurations via port licensing

Blade Switch Roadmap

Comprehensive Offering Today, with Strong Roadmap

Committed

Planning

2007

2008 / 9

Ethernet

HP, IBM, Dell & FSC

DC Focused

10G Virtual Blade Switch

GE Virtual Blade Switch

HP, IBM & Dell

10GE Switch (10G to Server)

Entry Level - IBM & Dell

Fibre Channel

HP & IBM

N-port Virtualizer (NPV)

SanOS 3.2(1)

Flex Attach

SanOS 3.2(2)

F-Port Trunking + Channeling

SanOS 4.X

8G Switch

Infiniband

IBM, Dell & HP

Support for IBM DDR

DDR Switch for NG Dell

Next-gen IB Blade Switches

Catalyst 4900M – The Next Generation TOR Platform

DataCenter Switching Architecture

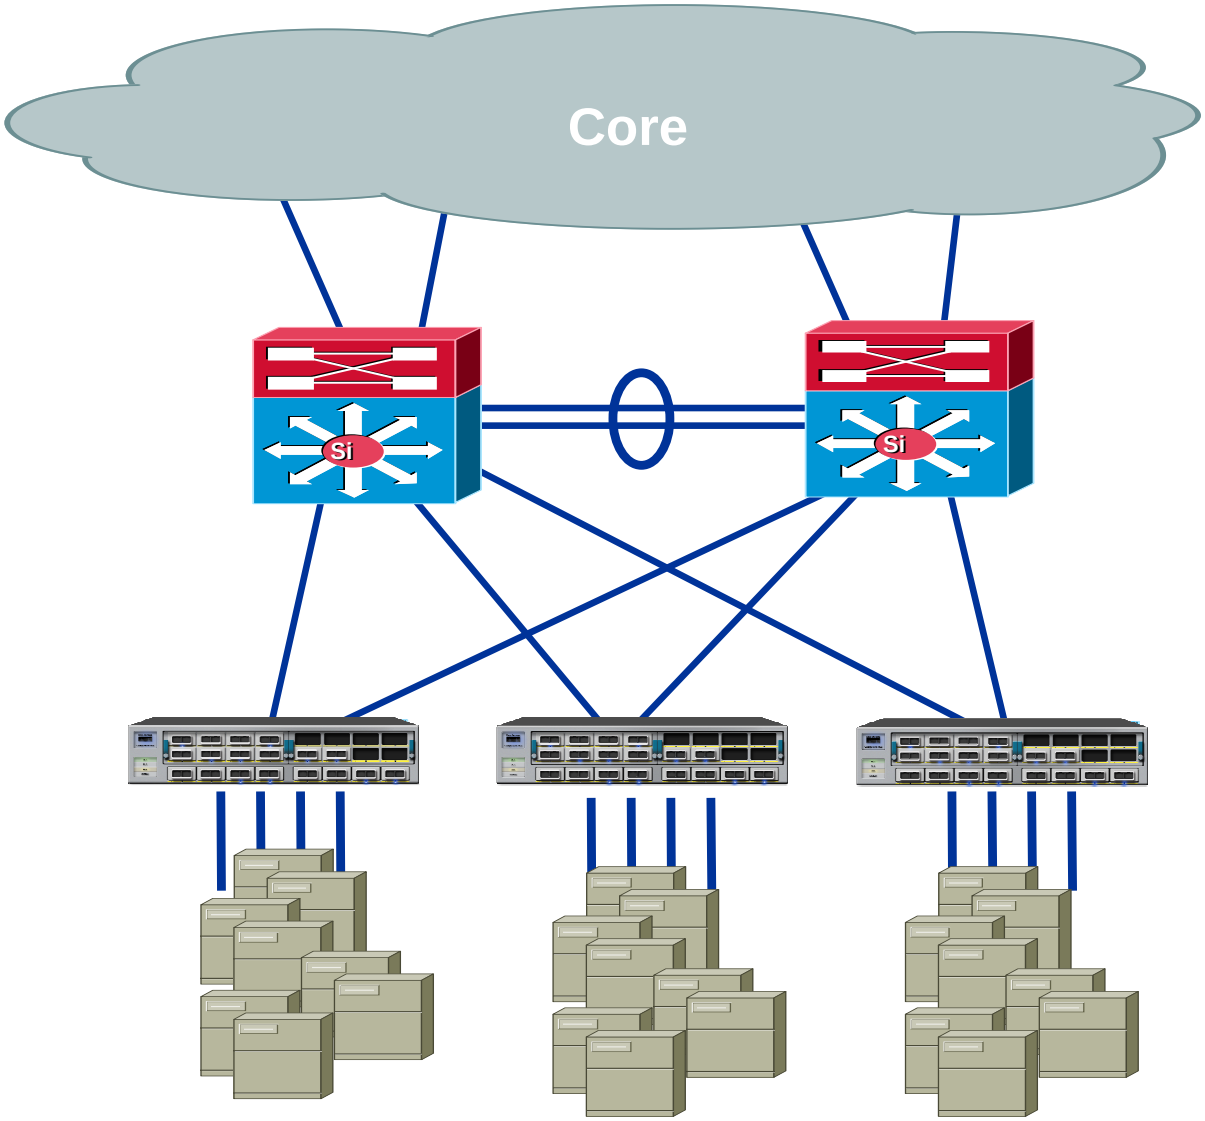

Agenda

- 4900M Overview
- 4900M Architecture
  - CenterFlex
  - Half Cards
- QoS
- Performance

Catalyst 4900M Overview

- 4900M:
  - 320 Gbps
  - 250 (IPv4) / 125 (IPv6) Mpps
  - High 1GE and 10GE Density
- Two RU form factor
- Dual, hot-swappable, internal power supplies (AC or DC options)
- Hot-swappable fan tray
- Jumbo frames on all ports
- Broadcast and multicast suppression in hardware for all ports (L2/3)
- Based on Catalyst 4500-E Series HW and SW Architecture

Catalyst 4900M Configurations

1

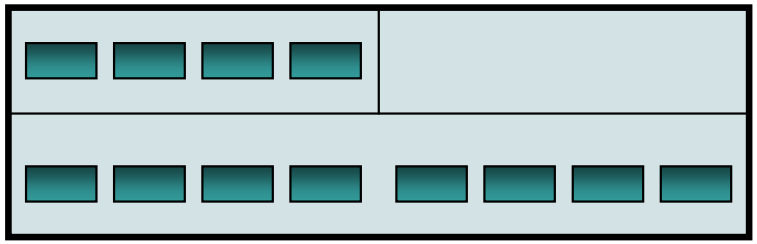

12 total 10 GE Ports (X2)

8 ports wire speed 10GE

+

4 ports wire speed 10 GE

2

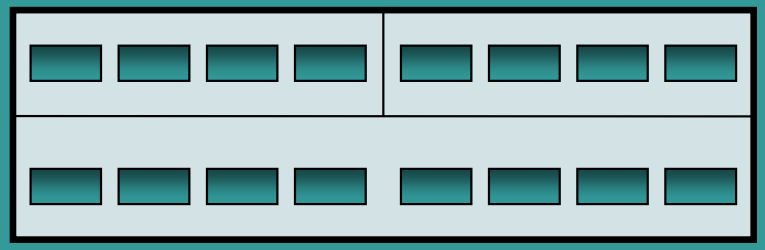

16 total 10 GE Ports (X2)

8 ports wire speed 10 GE

+

8 ports wire speed 10 GE

3

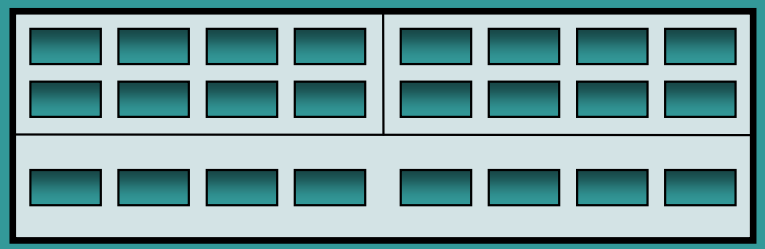

24 total 10 GE Ports (X2)

8 ports wire speed 10 GE

+

16 ports 2:1 oversubscribed 10 GE

– OR –

32 Ports wire speed 1 GE SFP

4

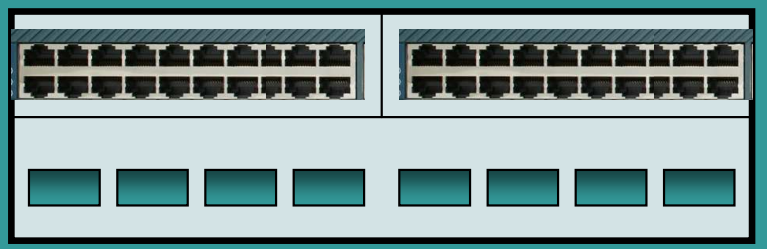

8 10 GE Ports (X2) + 40 10/100/1000

8 ports wire speed 10 GE

+

40 ports wire speed 10/100/1000 (RJ-45)

Current 4948 vs. 4900M Transmit Queuing

- 4948 has a single default queue up to 2336 packets
- 4948 can configure up to 4 queues

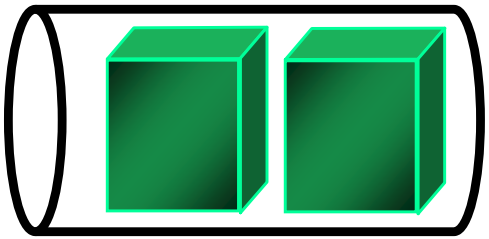

4948

For 1500 Byte Packet
Queue up to
28.4 ms for 1Gig interface
2.84ms for 10Gig interface

- 4900M has a single default queue up to 8184 packets
- 4900M can configure up to 8 queues

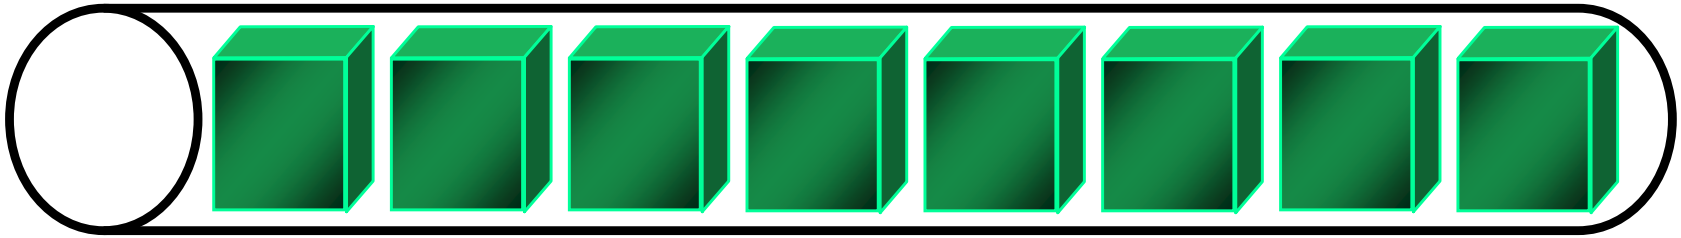

4900M

For 1500 Byte Packet
Queue up to
99.5 ms for 1Gig interface
9.95 ms for 10Gig interface

4900M with wire speed 10 GbE and test card for 10GbaseT

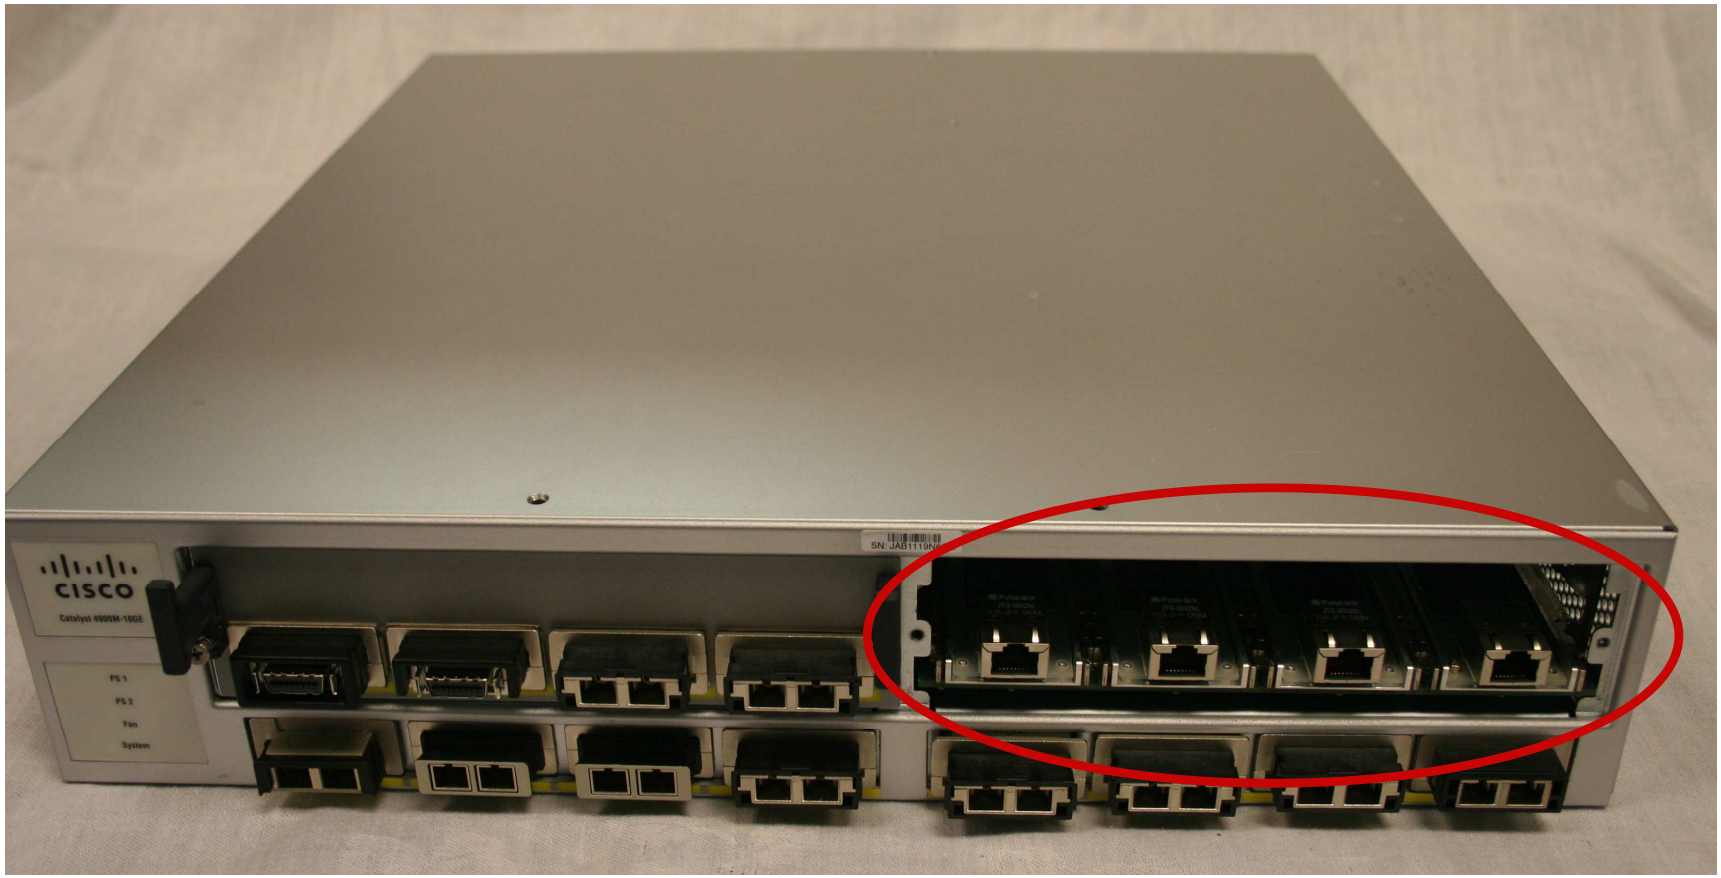

4900M Roadmap

Catalyst 4900M

20 port 10/100/1000
4 port 10 GE
8 port 10GE

8 port 2:1 10GbaseT card 1/10 auto-negotiating

8 port 2:1 SFP+ (Roadmap)

CY08

CY09

Agenda

- 4900M Overview
- 4900M Architecture
  - CenterFlex
  - Half Cards
- QoS
- Performance

CenterFlex

Catalyst 4900M

- **CenterFlex Inside**

High Performance

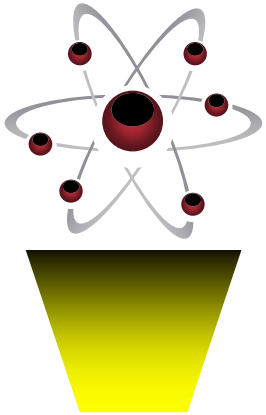

Increased Control and Flexibility

Scalable Services Capacity

IPv4 to IPv6 Migration

TCAM4 – Independent Forwarding and Services

128K Dedicated Services TCAM Entries for:

- Security
- QoS
- Policing

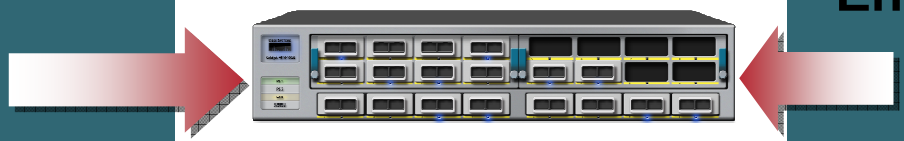

256K Dedicated Forwarding TCAM Entries for:

- IP Route
- Multicast Routes
- IGMP/MLD groups
- PBR (future)

Wire-speed, high-touch services with no performance hit:

- 128K QoS policies & ACL entries dynamically shared
- 16K policers
- 256K unicast/multicast entries

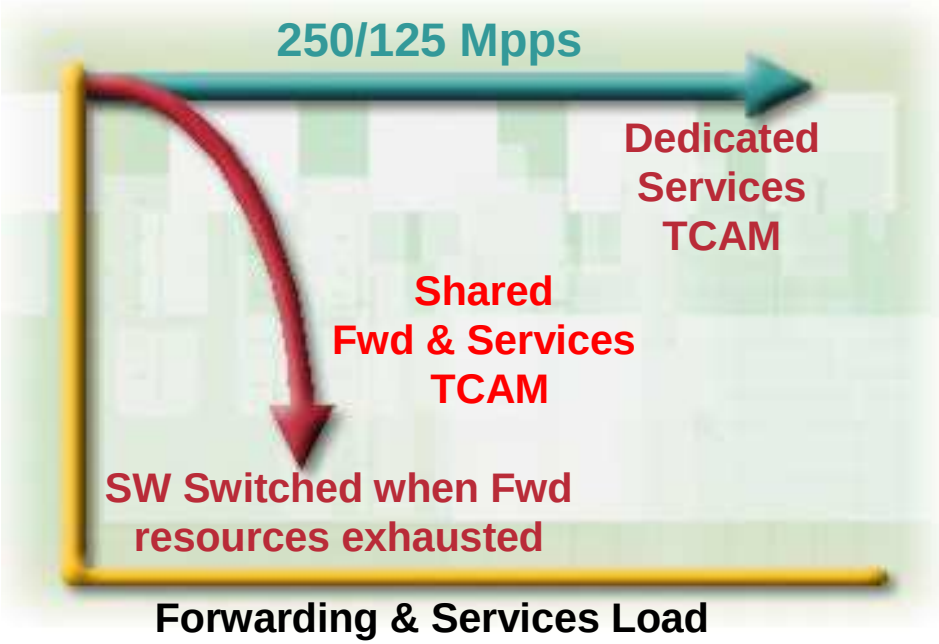

Half Cards

WS-X4930-GB-RJ45

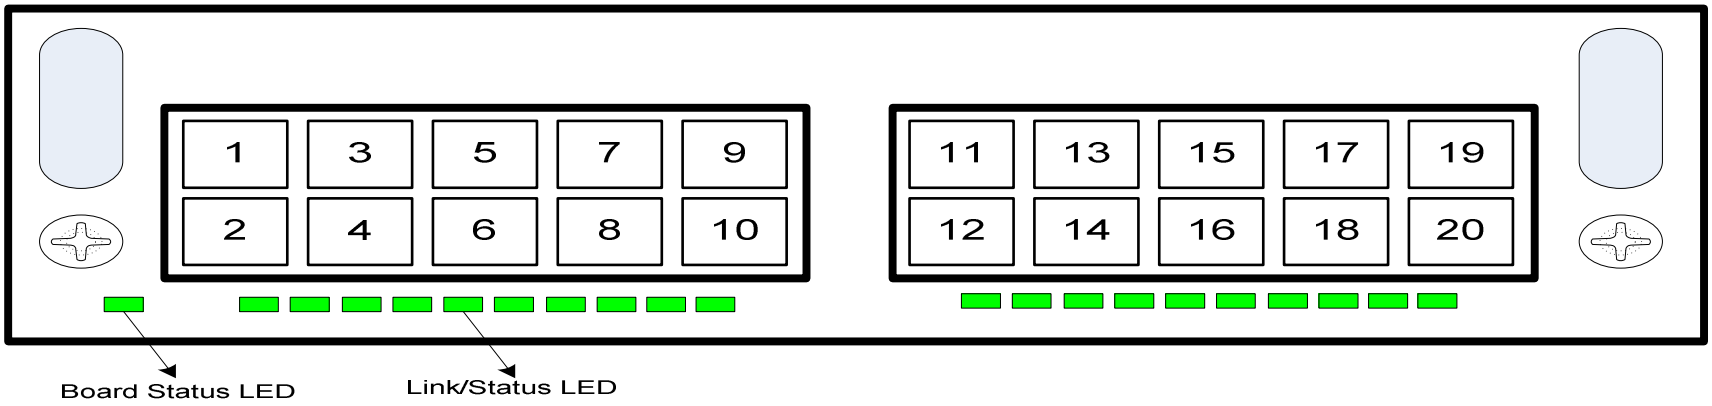

- ✓ **Up to 20 Ports of 10/100/1000 Ethernet**
- ✓ **802.1x Flow Control support**
- ✓ **Up To 8184 Packets per Queue**
- ✓ **8 Queues per Port**

WS-X4904-10GE

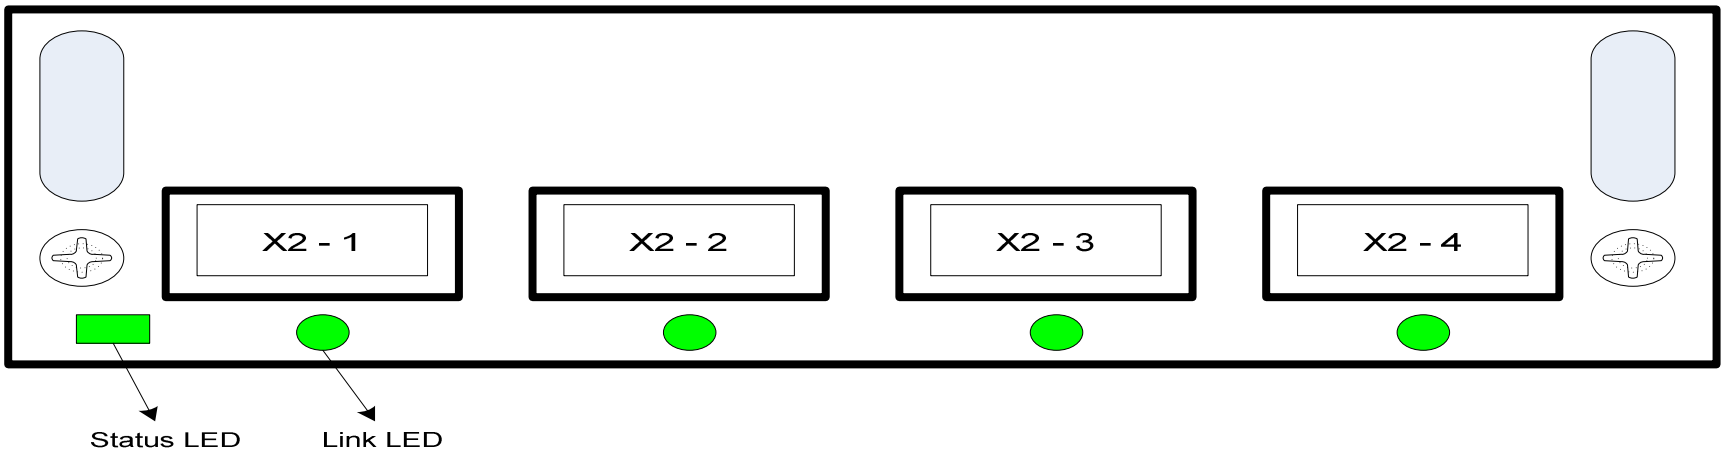

- ✓ **Up to 4 Ports of Wire Rate 10 GbE**
- ✓ **X2 Optics**
- ✓ **Up to 8184 packets per Queue**
- ✓ **8 Queues per Port**

WS-X4908-10GE - Port Numbering

Face Plate in X2 Mode

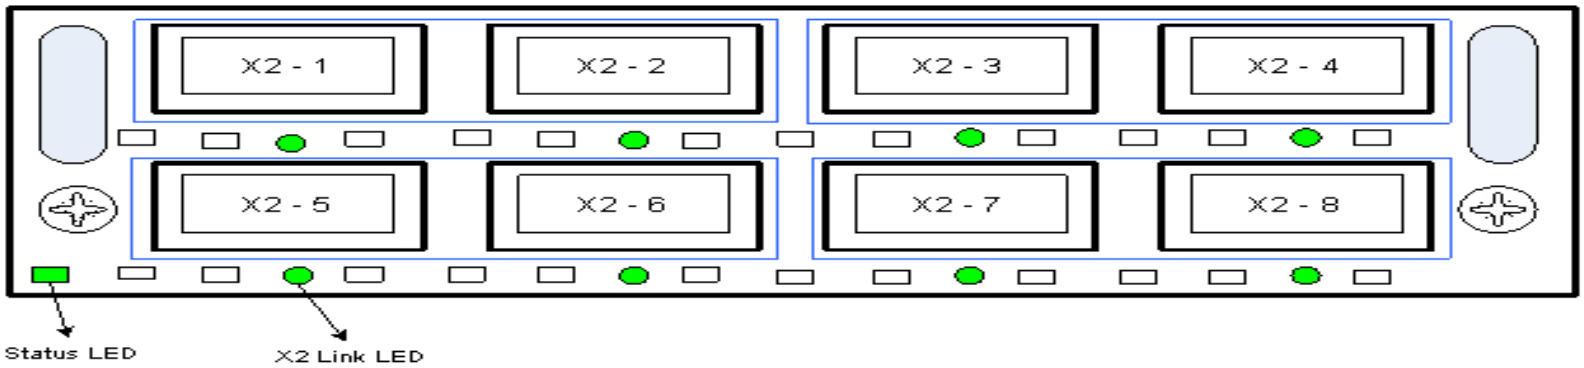

Face Plate in SFP mode

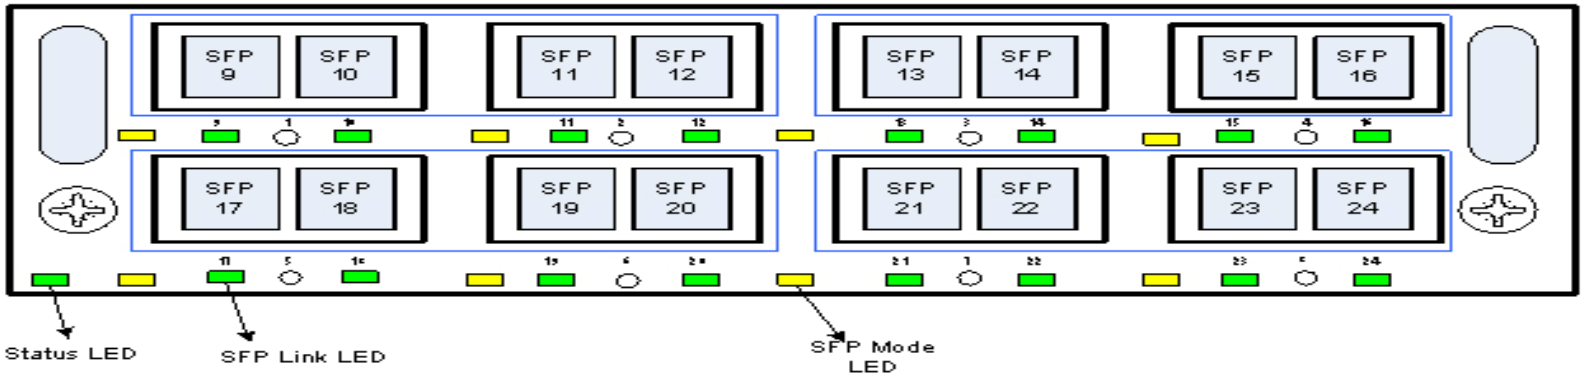

WS-X4908-10GE

- ✓ 8 Ports of 2:1 Oversubscribed 10 GbE with X2 Optics
- ✓ Compatible With TwinGig Converter Module
- ✓ 16 Ports of 1 GE with TwinGig C Module & SFP Optics
- ✓ 802.1x Flow Control support
- ✓ Up to 8184 Packets Per Queue
- ✓ 8 Queues per Port

8 X 10GBaseT (Future)

- ✓ 8 Ports of 2:1 Oversubscribed 10 GBaseT RJ-45
- ✓ IEEE 802.1an
- ✓ AutoNegotiate 1/10 GE
- ✓ Hot Swappable
- ✓ Dynamic Buffer Assignment

TwinGig Converter Module

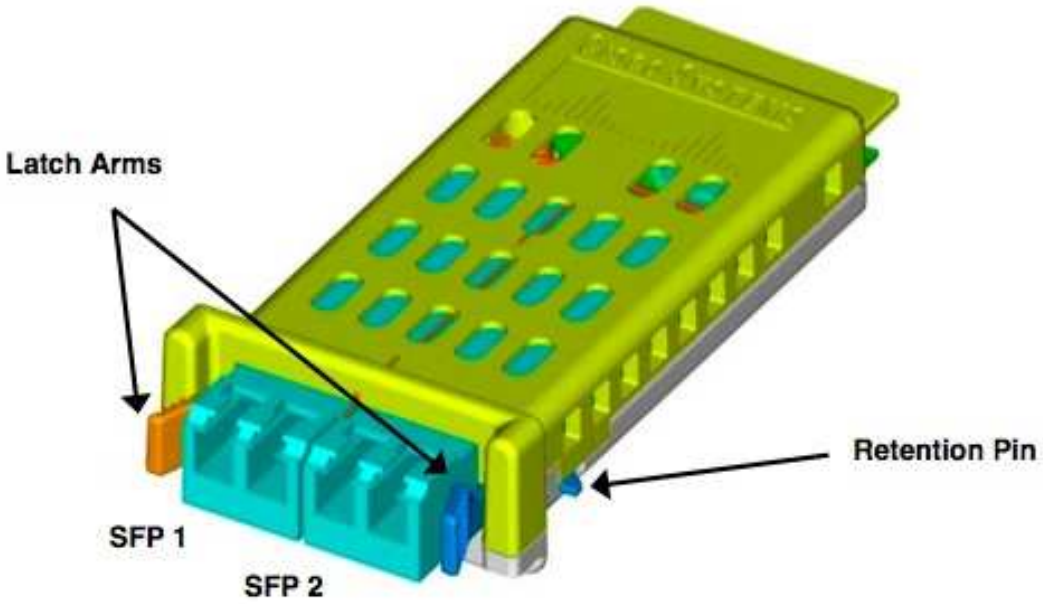

- **Provides Seats for 2 x 1GE SFP slots into a single X2 10GE port**
- **Allows Customers to mix 1GE and 10GE fiber downlinks.**

Agenda

- 4900M Overview
- 4900M Architecture
  - CenterFlex
  - Half Cards
- QoS
- Performance

4900M QoS Feature Highlights

- Configurable to 8 queues per port
- Flexible queue management
- Flexible policers
- Sequential marking and classification
- MQC compliant

4900M QoS Toolkit

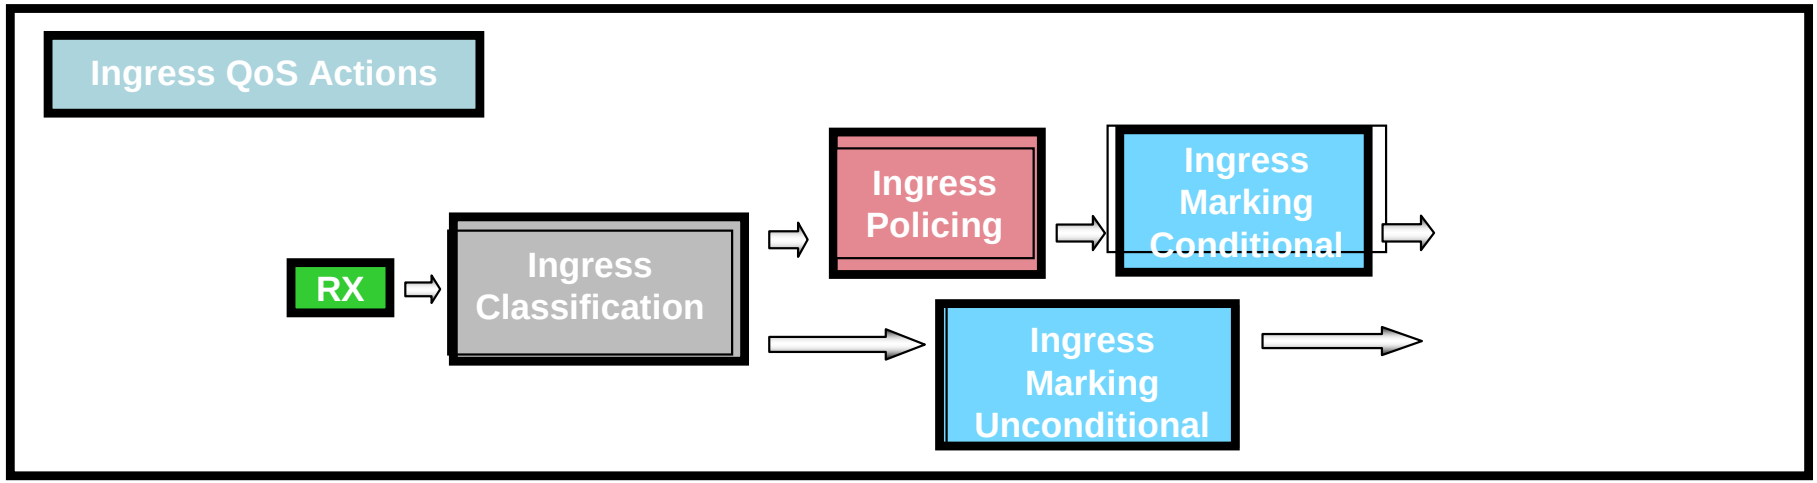

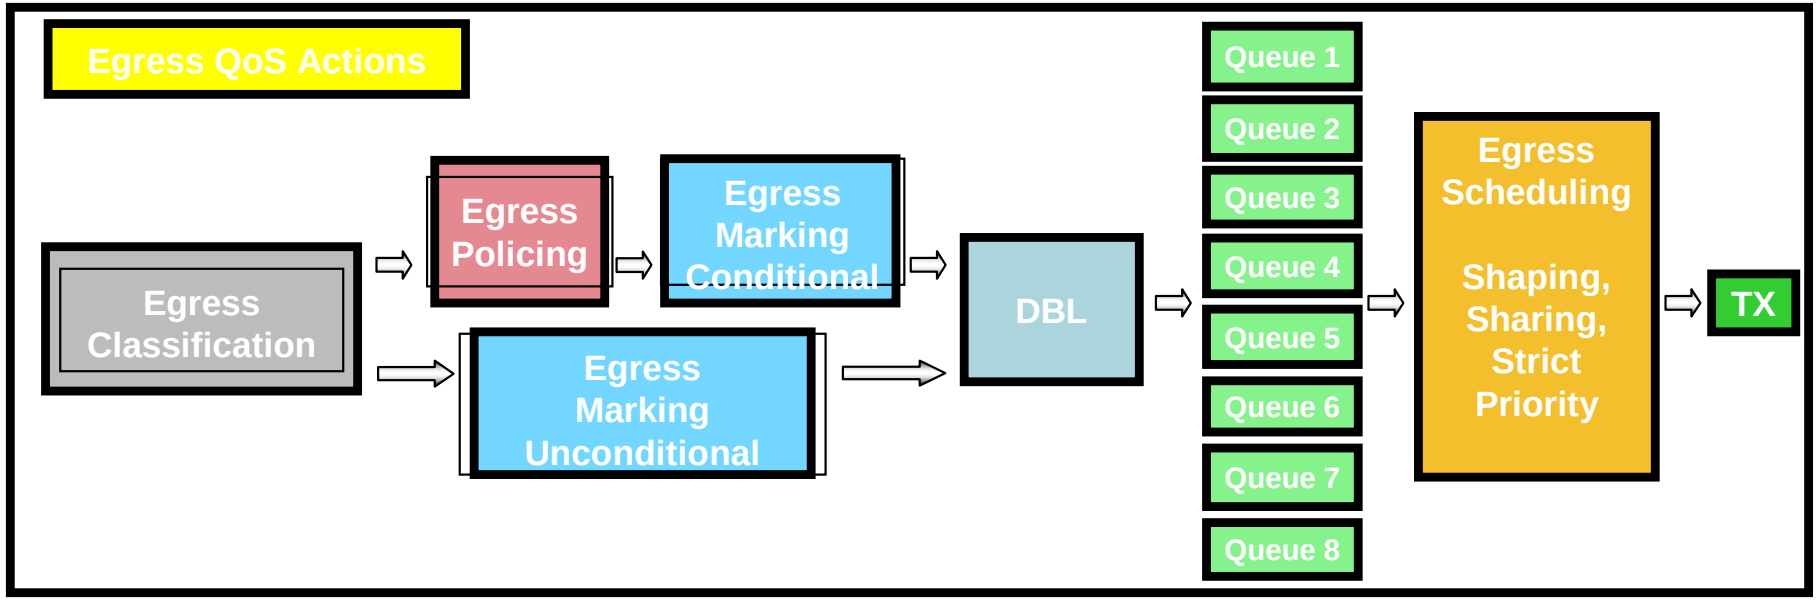

4900M QoS MQC

MQC Compliance = Flexibility

- MQC compliant!
- NO “TRUST” paradigm
- …n for proper queuing
- …on class
- …d of “internal DSCP”

```
table-map cos-to-dscp-mapping
    map from 0 to 0
    map from 1 to 8
    map from 2 to 16
    map from 3 to 24
    map from 4 to 32
    map from 5 to 40
    map from 6 to 48
    map from 7 to 56
    default copy
policy-map trust-cos-policy
    class class-default
       set dscp cos table
   cos-to-dscp-mapping
```

```
class-map priority
   match dscp 40
policy-map egress-queues
    class priority
         priority
```

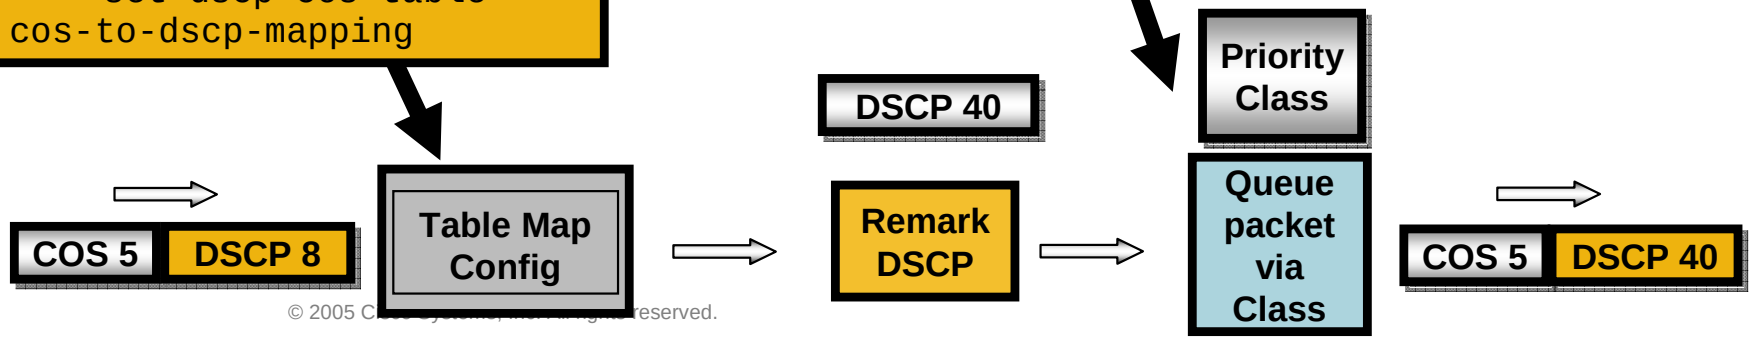

Queue memory allocation – 4900M

1. CPU/Drop queue entries allocated
2. Queue allocation for slots is divided by 3
3. Linecards divide entries equally per port
4. Port entries divided equally per queue

Flexibility and Ease of use

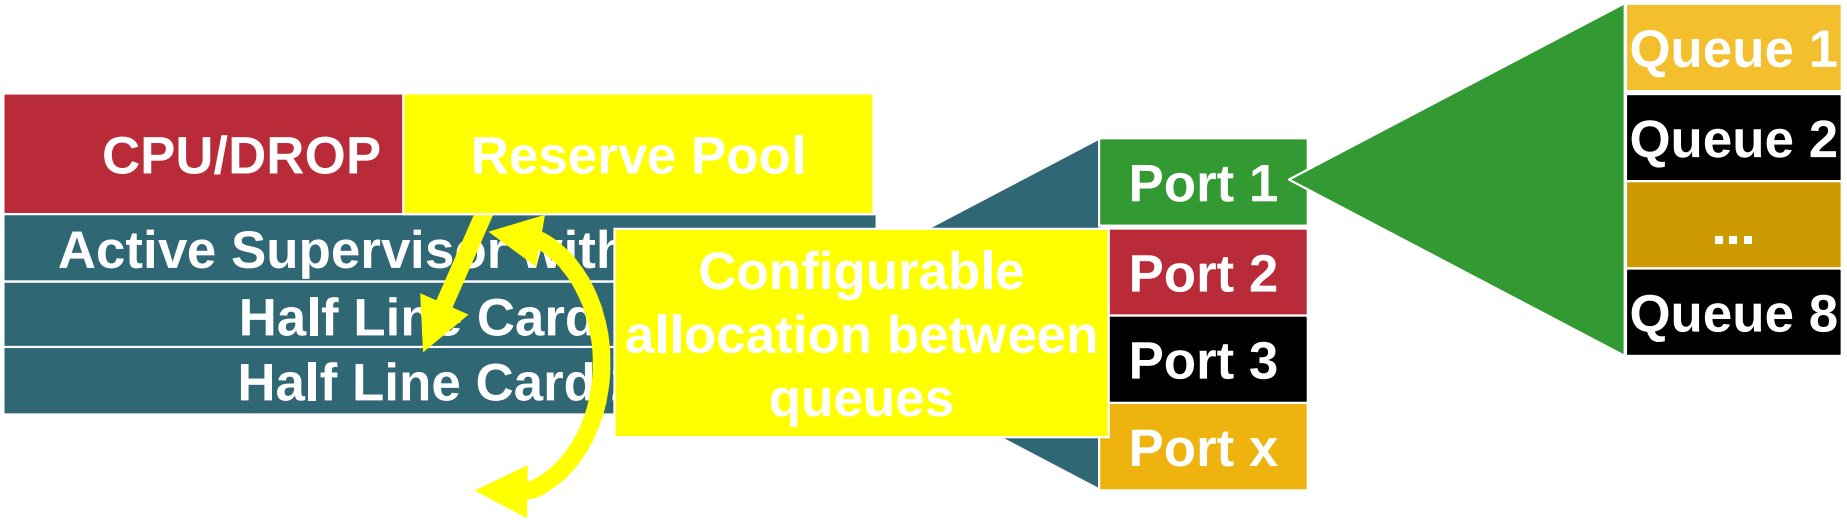

NOTE: **Ports default to 1 queue with NO policy attached. By default the port receives the entire port queuing capacity. Each queue can have a maximum of 8184 packets.**

Agenda

- 4900M Overview
- 4900M Architecture
  - CenterFlex ASICs
  - Port Half Cards
- QoS
- Performance

Catalyst 4900M Layer 2 Latency (Preliminary result on the Uplink ports)

LIFO method: Time from the last bit leaving the Smartbits to the first bit seen back (Recorded at 98% load)

Frame Size	Layer 2	
64 Bytes	2.41 µs	
128 Bytes	2.47 µs	
256 Bytes	2.56 µs	
512 Bytes	2.79 µs	
1024 Bytes	3.18 µs	
1280 Bytes	3.40 µs	
1518 Bytes	3.58 µs	

Cisco Catalyst 4900M —
Rack Optimized Server Switching

Wire speed Performance

Field replaceable fan and power supply

AC or DC Redundant power

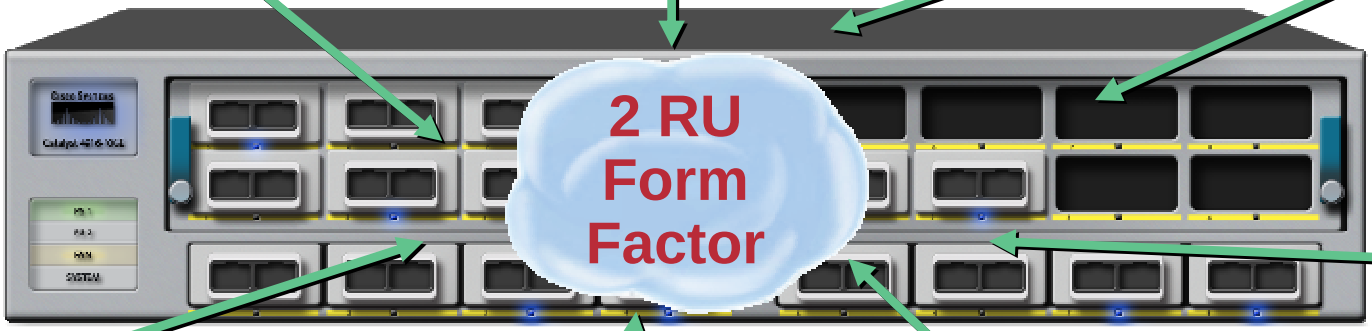

High Density 10-GE

Layer 2/3/4 support

QOS/Traffic Management

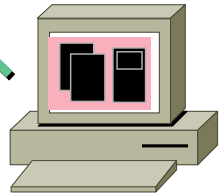

Manageability

4900M & 4948 Series Comparison

Model	WS-C4948	WS-C4948-10GE	WS-C4900M
Switching Capacity	96 Gbps	136 Gbps	320 Gbps
Throughput	72 Mpps	102 Mpps	250/125(IPv6) Mpps
IPv6 Support	Software	Software	Hardware
Multi-Layer Switching	Full L2/3/4 services and routing	Full L2/3/4 services and routing	Full L2/3/4 services
(E)IGRP,OSPF, BGP, ISIS	Yes	Yes	No
CPU	266 MHz	666 MHz	1.3 GHz
IP CEF Entries	32k	32k	48k
SDRAM	256M	256M	512M
Active VLANs	2K	2k	4k
IGMP Snooping	Yes (8K)	Yes (8k)	Yes (16k)
STP Instance	1500	1500	3000
SVI	2k	2k	4k
Broadcast Suppression	Hardware	Hardware	Hardware
Multicast Suppression	Yes	Yes	Yes
QoS Sharing	All Ports	All Ports	All Ports
QinQ	In Hardware	In Hardware	Pass Through
Uplinks	4 SFP	2 10GE (X2)	8 10GE (X2)

DCVT Catalyst 6500
Technical Update

Agenda

- **Catalyst 6500 Portfolio**
- **Catalyst 6500 10 Gigabit Ethernet Portfolio**
  - **Performance analysis**
- **Data Center Features in 12.2SXH**
  - **NSF/SSO Enhancements**
  - **Diagnostics Enhancements**
  - **Operational & Manageability**

Catalyst 6500 Family
E-Chassis Options

Catalyst 6500 "E" Series Chassis Benefits			
Ready for 80G / slot	Power and Cool Dense 10GE	Support up to 14500W of power	Chassis is S/W transparent

Red – roadmap items

Data Center Chassis -6509-V-E Chassis

Shipping In Oct'07

- **Vertical 9 Slot E-Series Chassis**
  - 80 Gbps/slot capable
  - Front-to-Back Airflow
  - Integrated Enhanced Cable Management
  - Redundant and removable fan tray
  - 21 RU (2 Chassis in a 42RU Rack)
  - Supports Sup32 and Sup720 Series
  - Supported in 12.2(18)SXF10 and beyond
- **List Price $11,995 (includes 2 Fan trays)**

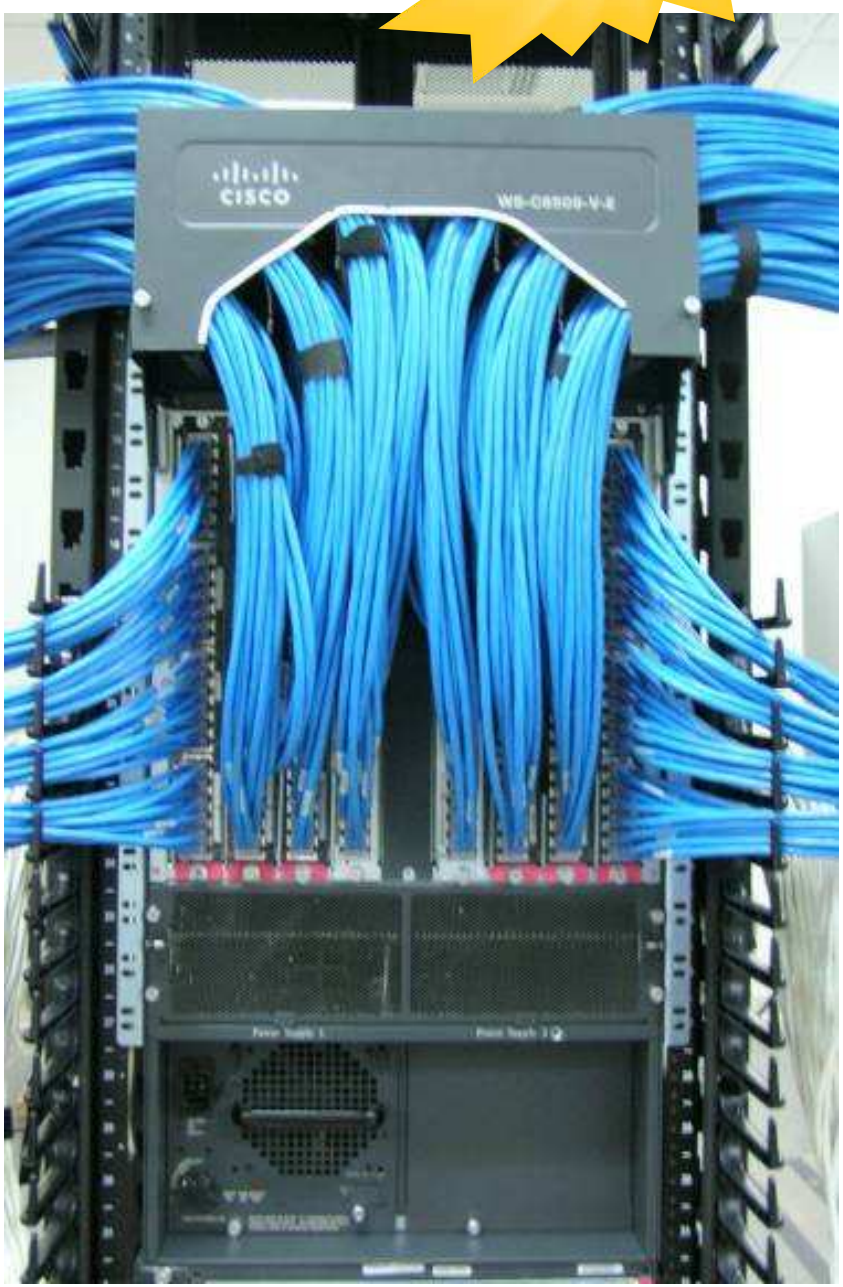

Enabling Hot Aisle/Cold Aisle Designs

Panduit Cabinet

- 45RU (32”W x 40”D x 84”H)
- Up to 20kW/cabinet heat rejection capability
- 3 x Catalyst 6509’s or 3 x MDS 9513’s per cabinet
- Front to back airflow into Hot Aisles
- Integrated Cable Management
- Modular design to support future air handlers or spot cooling
- Part # CN4-1 and CN4-2 for MDS 9513 and # CN4-3 for the Catalyst 6509E

http://www.panduit.com/products/brochures/105309.pdf

Embedded Event Manager
EEM Application Example in Data Center

Available since 12.2(18)SXF4

Upon detecting link to Server down:

- **Display counter error statistics for the link that has gone down**
- **Start a Time Domain Reflectometry (TDR) test**
- **Start a GOLD Loopback test**
- **Send the results using a provided template to a user-configurable address**

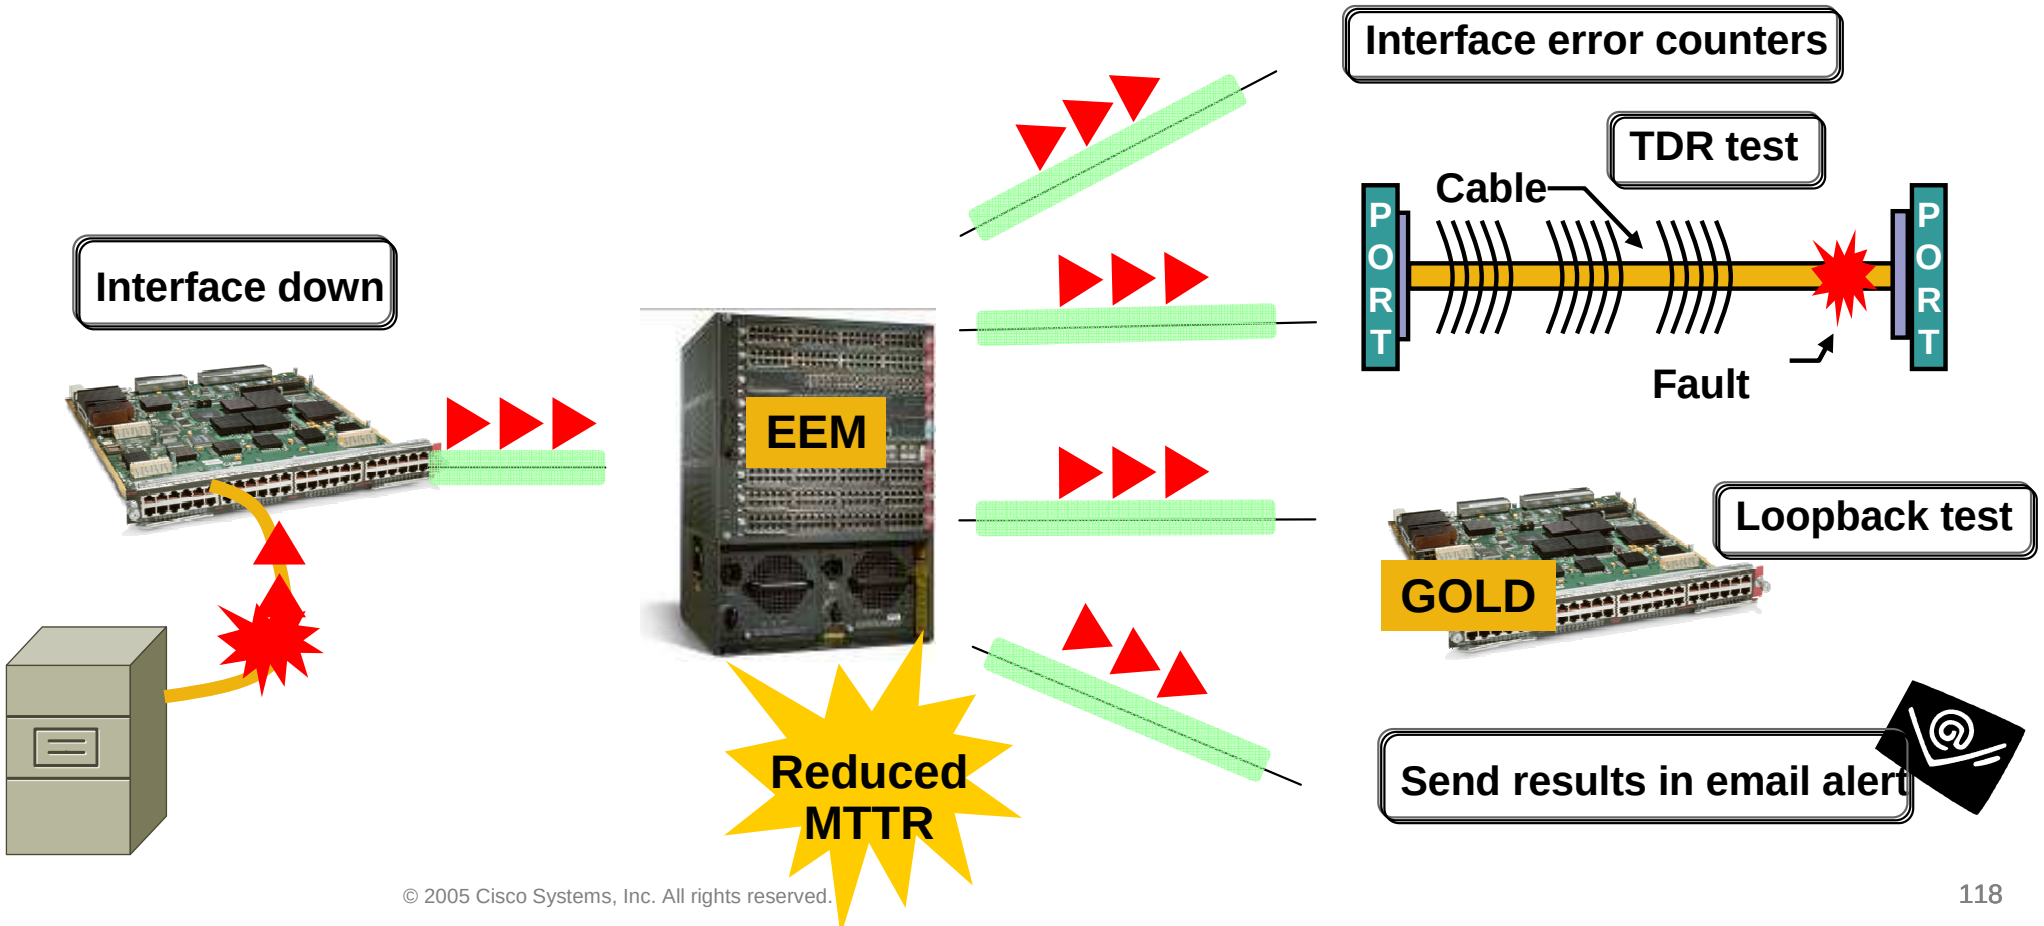

Typical Data Center Configuration Today

Supervisor720: Provides hardware accelerators for mission critical deployments. Throughput of 400Mpps…

WS-X6748-GE-TX :
48ports of GE for data center access

WS-X6704-10GE :
4 ports of 10GE for data center core, distribution, HPCC

WS-X6708-10GE :
8 ports of 10GE for data center core and distribution

Typical Data Center Configuration in the Next 3 – 10 Months

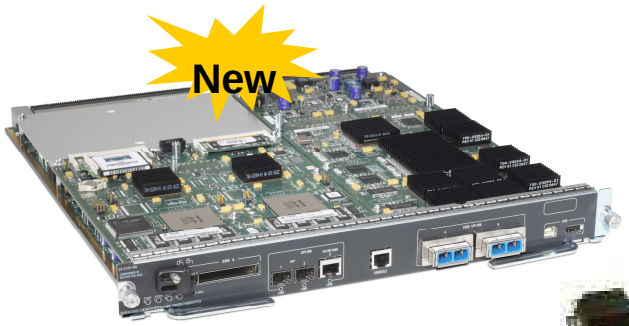

Supervisor720 10GE: Evolution of Sup720 hardware accelerated features with 10GE uplinks. Throughput of 400Mpps…

WS-X6748-GE-TX :
48ports of GE for data center access

WS-X6704-10GE :
4 ports of 10GE for data center core, distribution, HPCC

WS-X6708-10GE :
8 ports of 10GE for data center core and distribution

WS-X6716-10GE :
16 ports of 10GE for data center distribution and access

Virtual Switching Supervisor Engine 720 w/ 10GE Uplinks

New Sup720-10GE Features:

IOS: 12.2(33)SXH

- **Both GE and 10GE Uplink Ports.**
  - **2 x 10GE (X2 Optics) : VSL Capable**
    - **100 MB Per Port Ingress/Egress Buffering**
  - **3 x GE : 2 SFP and 1 Copper (10/100/1000)**
  - **All uplinks active in redundant configuration**
  - **All uplinks share 20G Channel to Backplane**
- **Compatible with all chassis (E- and non-E series)**
- **List : PFC3C - $38,000, PFC3C-XL : $48,000**
- **Default Software: IP Services (VSS included)**

Orderability: Oct'07
FCS: Nov'07

Virtual Switching System 1440

Network System Virtualization

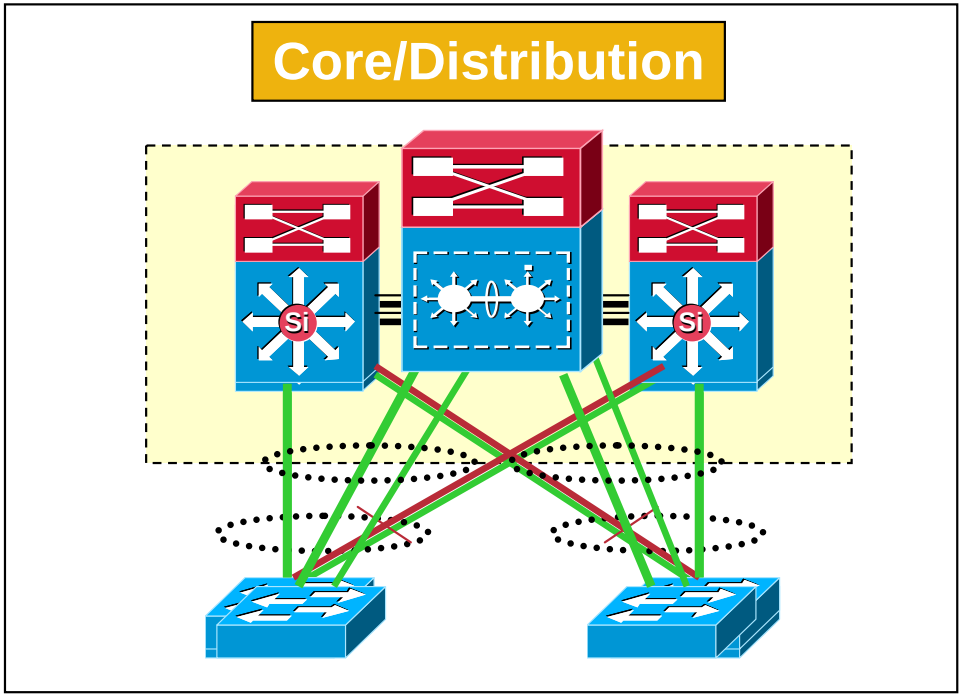

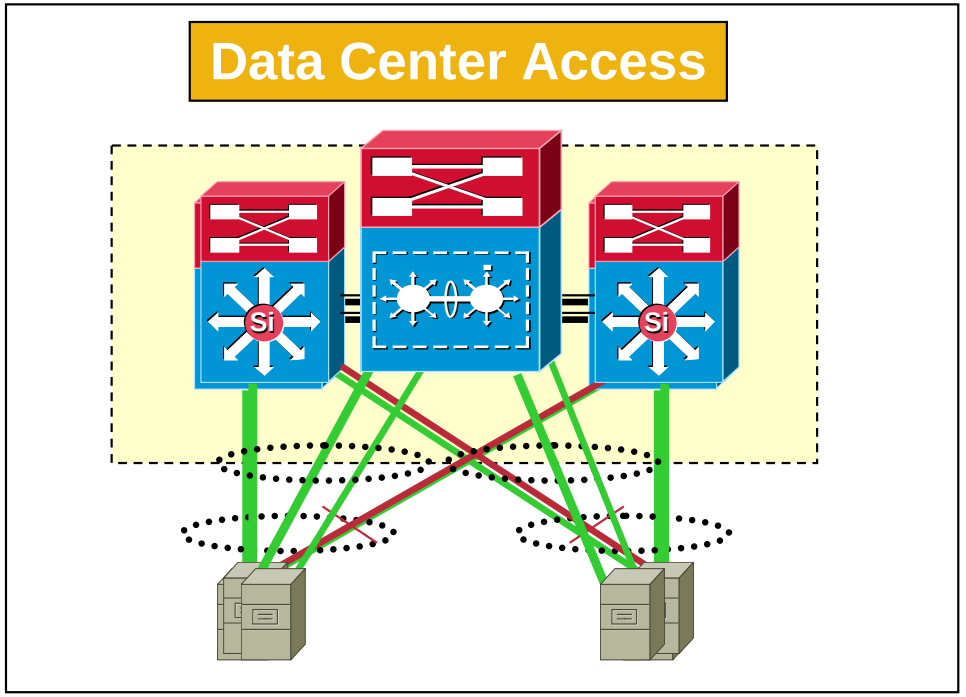

Traditional L2/L3 Design High Availability Challenges

Multiple Nodes and Control Protocols to Manage and Troubleshoot – HSRP / VRRP / STP

Network-level failovers may not be stateful, potentially increasing full recovery times

Underutilized links between Access switches & Core / Distribution switches and Servers & Access switches

Virtual Switching Reduces Latency by 25%

Traditional L2/L3

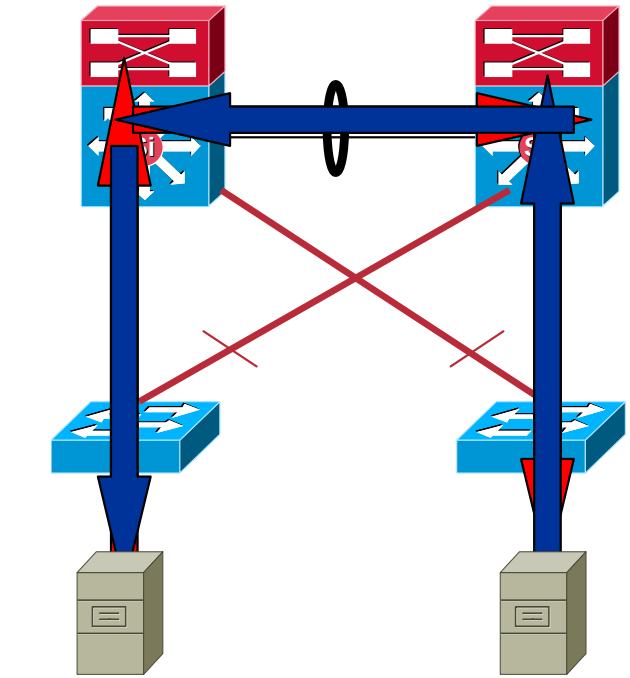

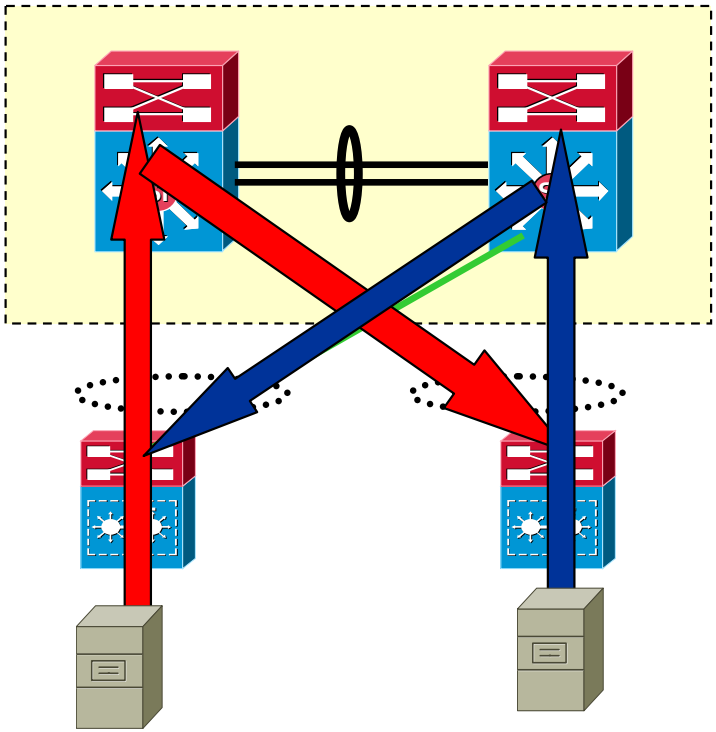

VSS Simplifies Intra-Datacenter Traffic Pattern

- **Minimum Hop to Intra-Data Center Destination**
  - Reduced Latency up to 25% and Simplifies traffic pattern
- **All links forwarding** resulting in simple traffic pattern
- **Etherchannel on virtual Switch member enhanced to prefer local link**

Virtual Switching and Link State Tracking

Link State Tracking

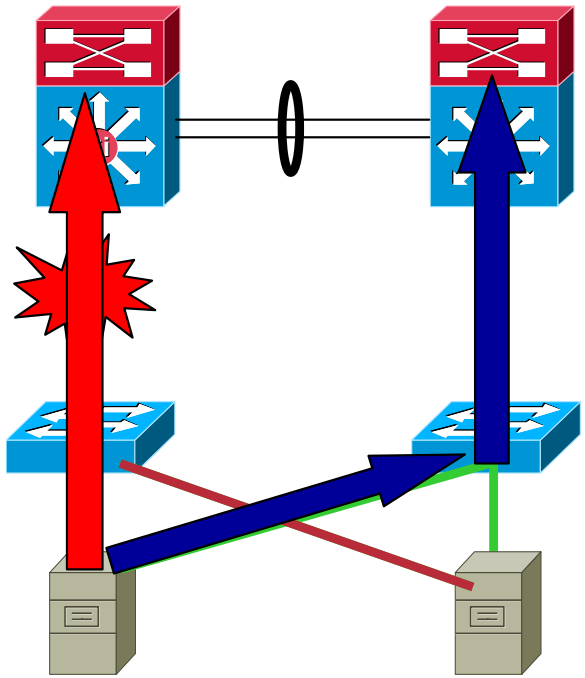

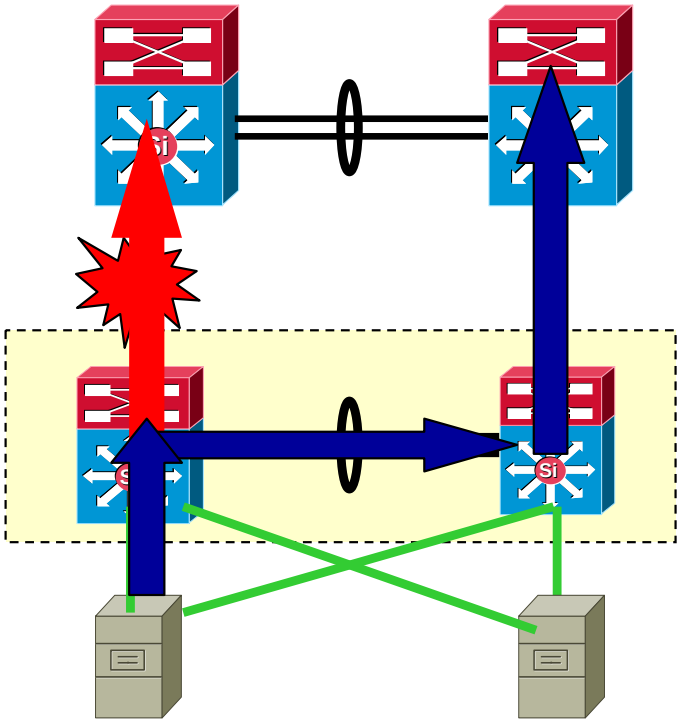

VSS Provides Uplink Resiliency

—All links to servers are maintained – No loss of bandwidth to servers

Recommendation:
To dual home each of the Access switch to the distribution switch

Virtual Switching System 1440 – Where to Deploy

Network Segment	2007	2008
 **Campus**	2 Switch VSS - 67xx Card with CFC/DFC3C - Single Supervisor per Chassis - NAM 1/2	ISSU – True Hitless Upgrade - MPLS and IPv6 - Dual Supervisor per Chassis* - FWSM, WiSM, IDSM
 **Data Center**	2 Switch VSS - 67xx Card with CFC/DFC3C - Single Supervisor per Chassis - 128 MEC - NAM 1/2	- VFRAME DC - ISSU – True Hitless Upgrade - Dual Supervisor per Chassis* 512 MEC ACE 10/20, FWSM, IDSM
 **Wiring closet/WAN**	Under Investigation **(Continue to deploy Supervisor 32/32-PISA and Sup720-10GE in standalone mode)**	Under Investigation **(Continue to deploy Supervisor 32/32-PISA and Sup720-10GE in standalone mode)**

* Roadmap Item, Not committed

6704-10GE Linecard
Overview

Some Facts

- 4 x 10GE ports (XENPAK Optics)
- CEF720 line card
- Optional Integrated DFC
- Upgradeable to DFC3CXL
- Near line rate at all frame sizes
- Egress Shaping (SRR, DWRR)
- 16MB buffer on Ingress and Egress
- Compatible with E and non E chassis
- Support for Fast Switchover (<200msec)
- Ports are not VSL Capable

Deployment Scenarios:

- Core
- Distribution
- High Performance Compute Clusters
- Long Haul Interconnections (XENPAK)

WS-X6704-10GE

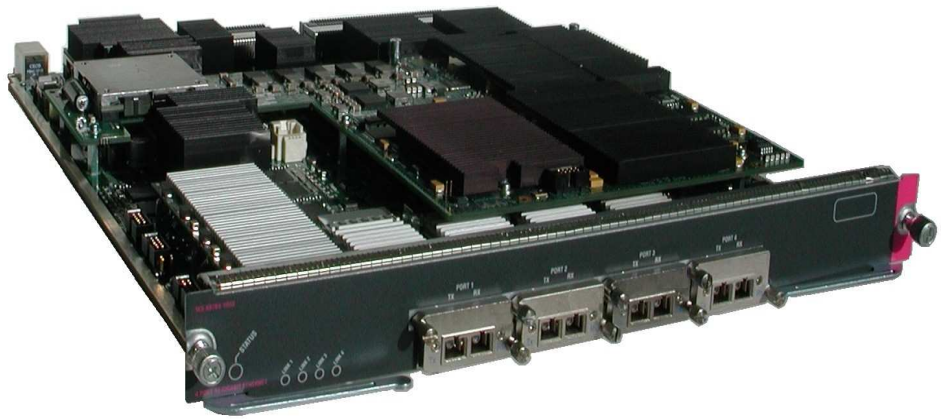

WS-X6704-10GE Cross Fabric Latency

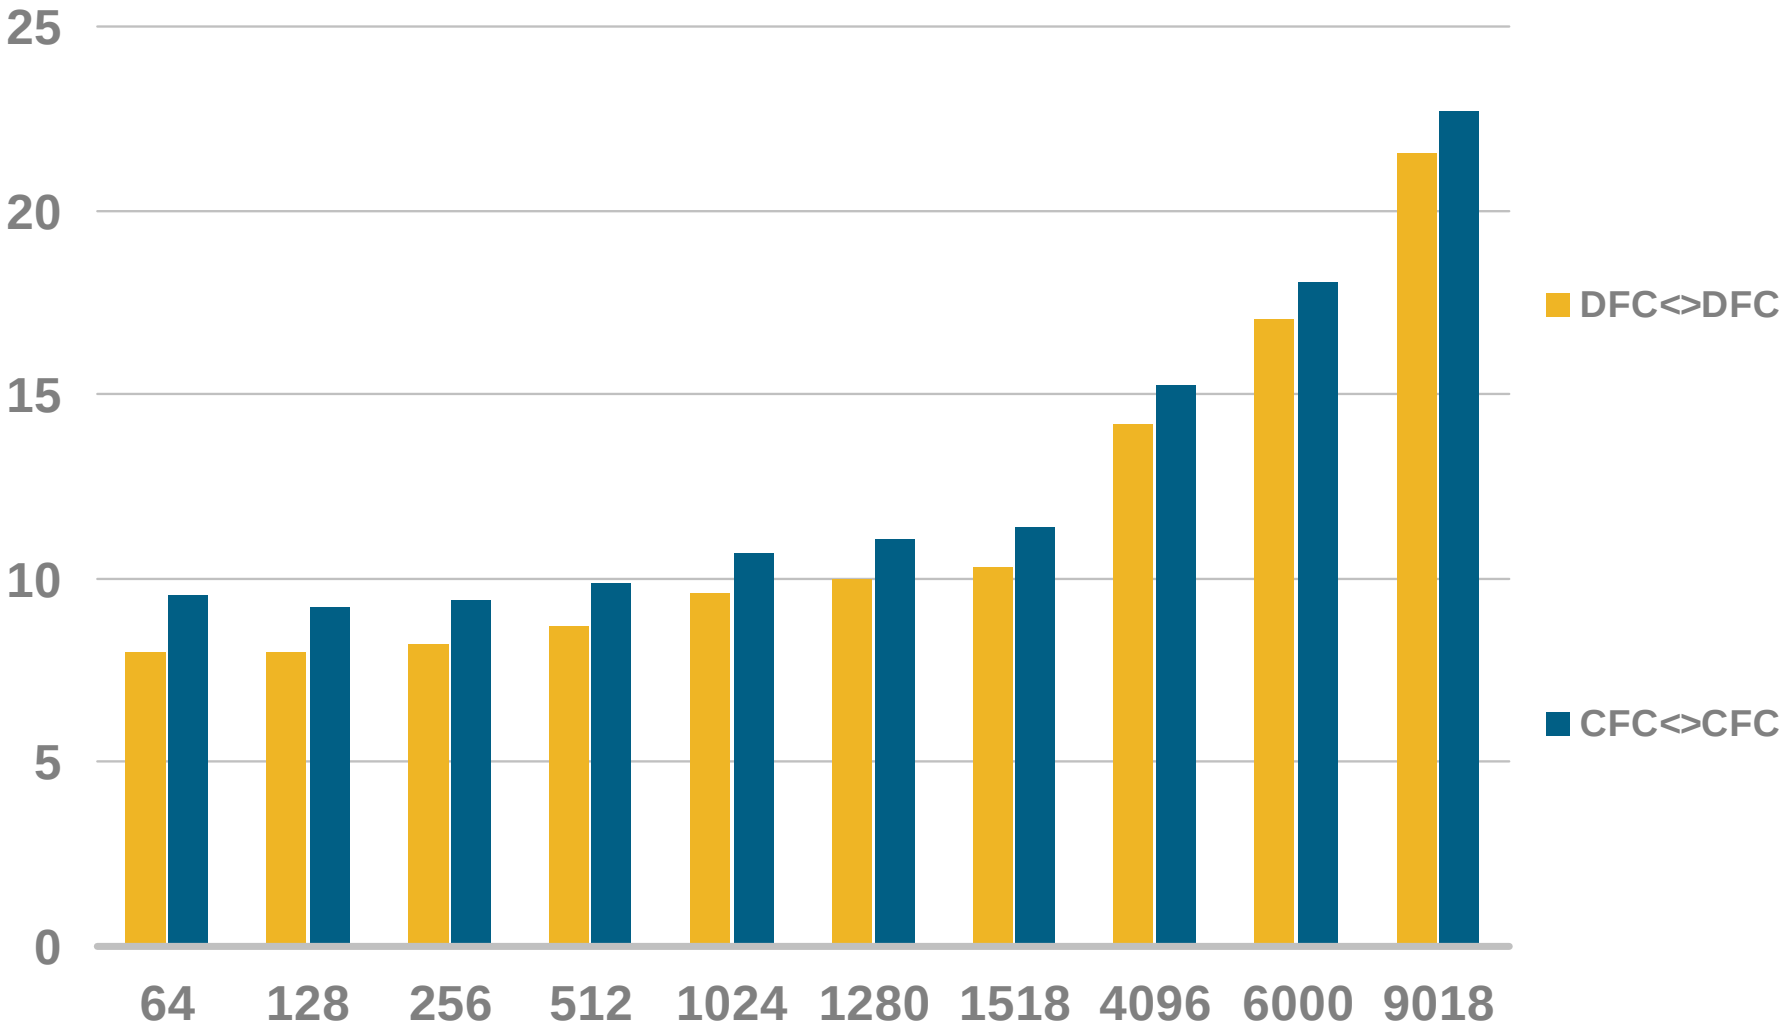

6708-10GE Linecard
Overview

Some Facts

- 8 x 10GE ports (X2 Optics)
- CEF720 linecard
- Integrated DFC3C
- Upgradeable to DFC3CXL
- 2:1 oversubscribed
- Will work with PFC3B and PFC3BXL
- All ports VSL Capable
- Egress Shaping (SRR, DWRR)
- 128MB buffer on Ingress and Egress
- Compatible with E and non E chassis
- Support for Fast Switchover (<50msec)
- Shipping 12.2SXF & 12.2SXH

Deployment Scenarios:
Core
Distribution
High Performance Compute Clusters
VSL Interconnections

WS-X6708-10GE-3C
WS-X6708-10GE-3CXL

WS-X6708-10GE-3C Latency

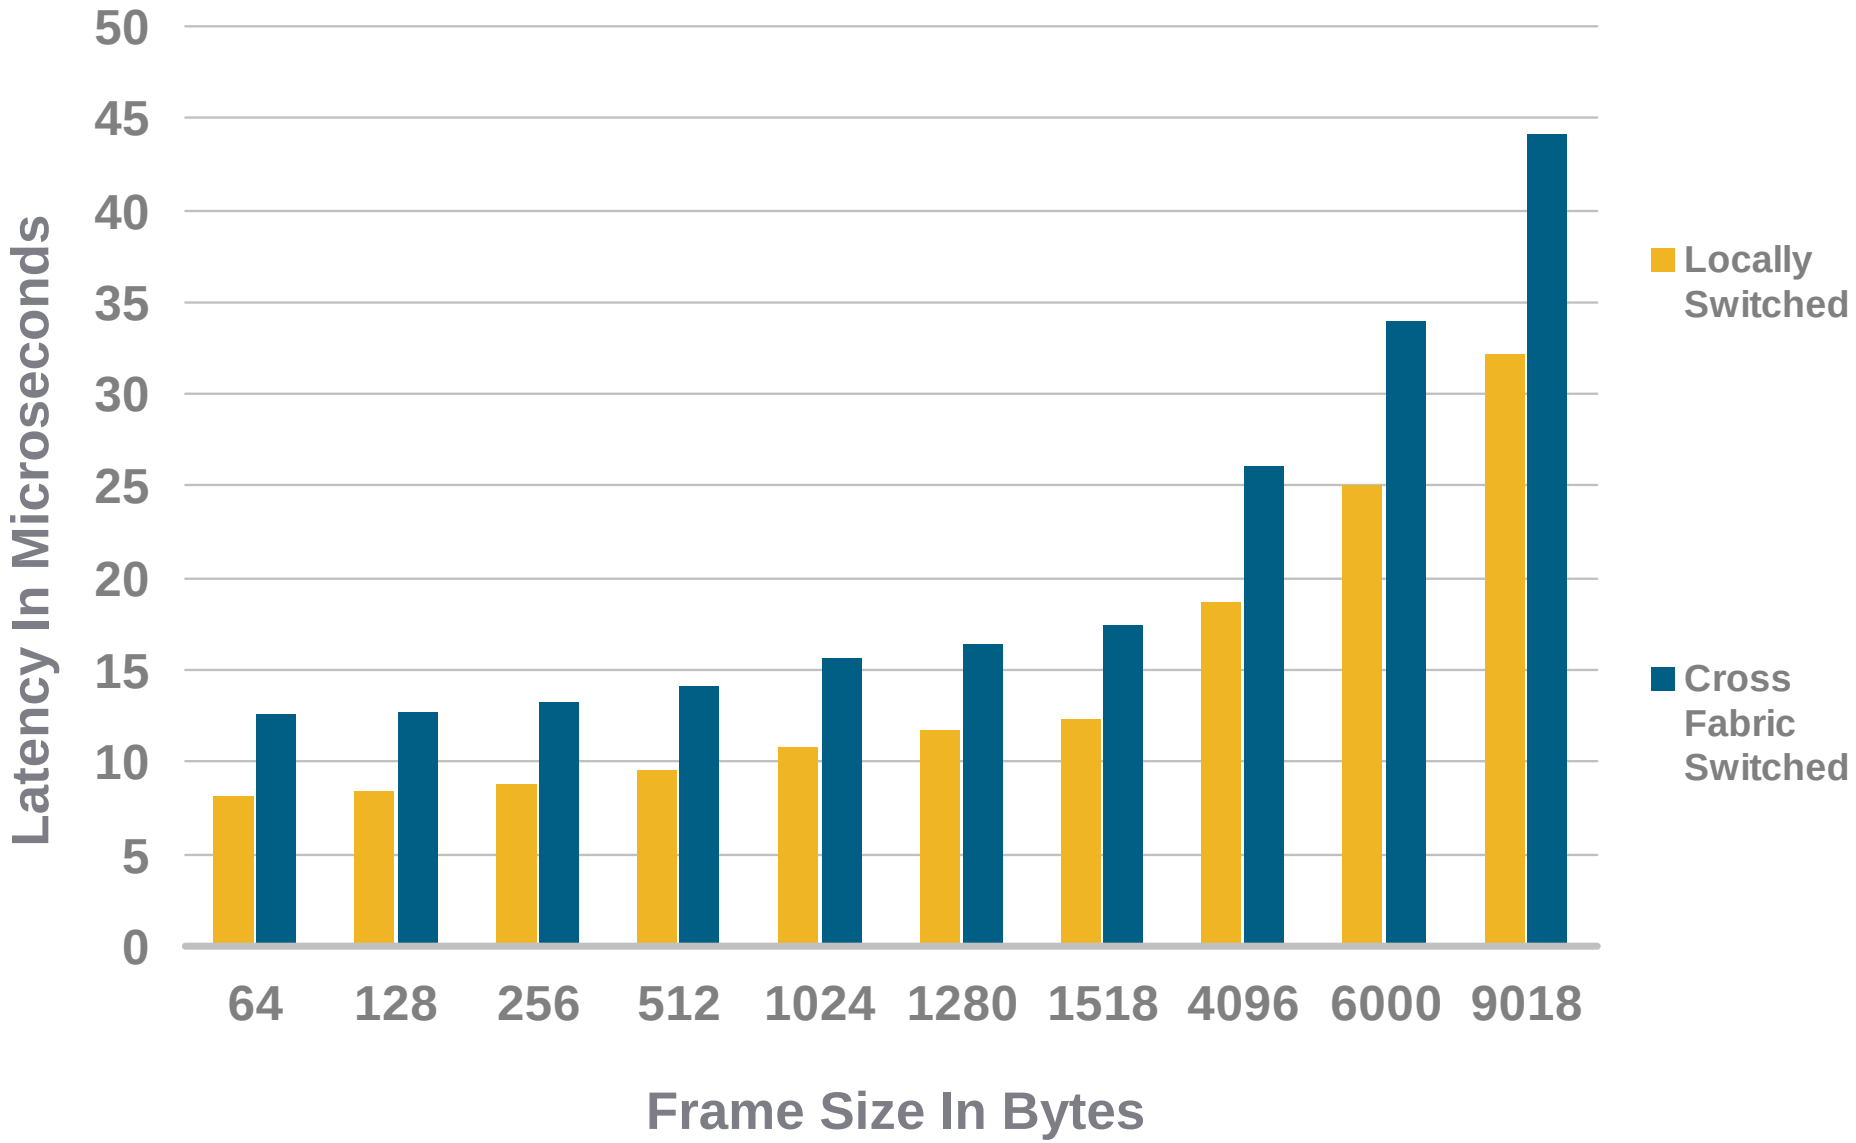

10 Gigabit Ethernet Line Card Comparison

	WS-X6704-10GE	WS-X6708-10GE	WS-X6716-10GE (2Q2008)
Deployment Scenarios	Core, Distribution	Core, Distribution	Enterprise Distribution, DC Access
Performance Ratio	Near non-blocking	2:1 oversubscribed	4:1 oversubscribed
Fabric Connection	40 Gbps	40 Gbps	40 Gbps
Port Density Per Chassis (Catalyst 6509)	32	64	128
VSL Capable	No	Yes	Yes (transparent mode ports)
Transceiver Type	XENPAK	X2	X2
Fabric Switchover (12.2SXH)	Sub-200ms	Sub-50ms	Sub-50ms
Power Requirements	402W DFC3BXL 377W DFC3B 295W CFC	473W DFC3CXL 444W DFC3C	Not finalized
Price	$20,000 CFC $27,500 DFC3C $35,000 DFC3CXL	$37,500 DFC3C $45,000 DFC3CXL	$47,500 DFC3C $55,000 DFC3CXL

10 Gigabit Ethernet Portfolio QoS Comparison

	Sup720-10G (10GE Uplinks in 10G only mode)	Sup720-10G (10G uplinks when GE ports enabled)	WS-X6704-10GE	WS-X6708-10GE	WS-X6716-10GE (Transparent Mode)	WS-X6716-10GE (Mux Mode)
TX Buffer Memory	88MB	88MB	14MB	92MB	90MB	90MB
TX Queue Structure	1p7q4t	1p3q4t	8q8t (DFC) 1p8t (CFC)	1p7q4t	1p7q4t	1p7q4t
TX Queue Scheduling	DWRR or SRR	DWRR or SRR		DWRR or SRR	SRR/WRR	WRR
TX Queue mapping	CoS or DSCP	CoS or DSCP	CoS	CoS or DSCP	CoS or DSCP	CoS or DSCP
RX Buffer Memory	104MB	104MB	2MB	109MB	109MB	950KB
RX Queue Structure	8q4t	2q4t	1p7q8t	8q4t	8q4t	1p7q2t
RX Queue Scheduling	DWRR or SRR	DWRR or SRR	DWRR	DWRR or SRR	WRR	SDWRR
RX Queue mapping	CoS or DSCP	CoS or DSCP	CoS	CoS or DSCP	CoS or DSCP	CoS or DSCP

Catalyst 6500
Data Center
Features in 12.2SXH

Recent Enhancements Now Shipping in 12.2.33SXH

Catalyst 6500 Hot-Sync Standby Fabric

Improves SSO Switchover Times

- **Reduces SSO switchover time to less than 200ms**
- **Standby switch fabric is brought to an online state ready to switch traffic**
- **Data is only switched on the active switch fabric**
- **Supported on 67XX-series line cards**
- **Requires E-series chassis**
- **Available as of 12.2(33)SXH**

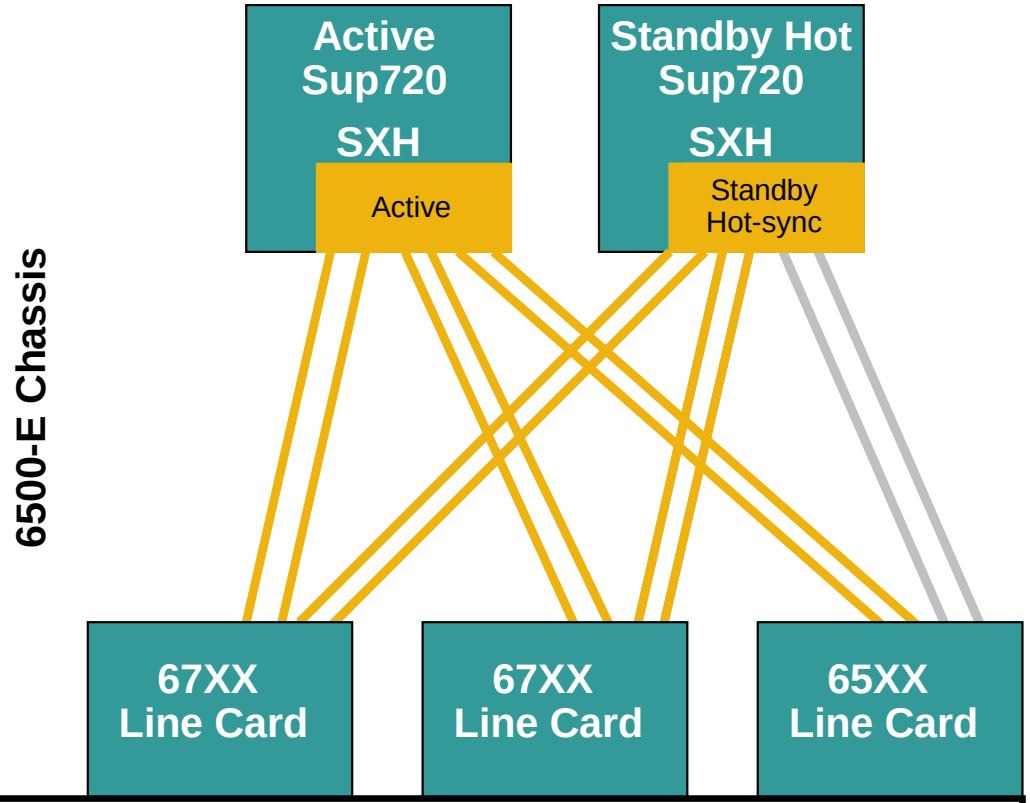


```
Router# show fabric status
 slot  channel speed module   fabric   hotStandby  Standby  Standby
                    status   status      support  module   fabric
    1        0   20G     OK       OK       Y(hot)
    1        1   20G     OK       OK       Y(hot)
    2        0   20G     OK       OK       Y(hot)
    2        1   20G     OK       OK       Y(hot)
    3        0   20G     OK       OK       Y(hot)
    3        1   20G     OK       OK       Y(hot)
    4        0   20G     OK       OK       Y(hot)
    4        1   20G     OK       OK       Y(hot)
    5        0   20G     OK       OK       Y(hot)
    6        0   20G     OK       OK       Y(hot)
```

Switchover Event – Software Notification

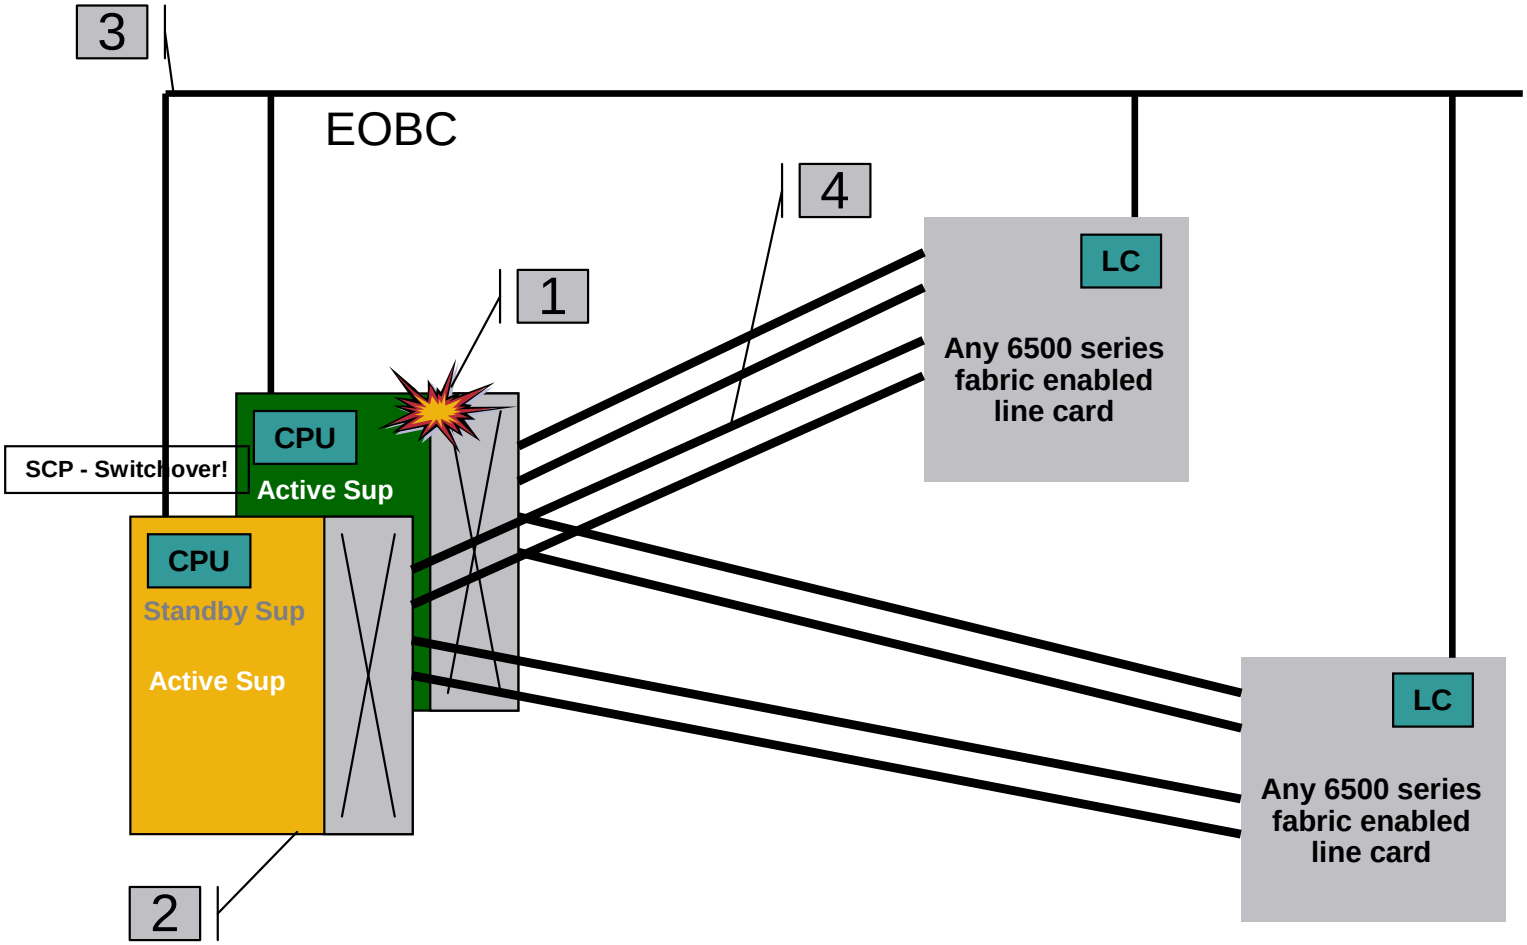

1. Supervisor switchover event triggered
2. Standby Supervisor becomes active
3. Standby notifies line cards of switchover event via an Supervisor Control Protocol message
4. Line cards switch their fabric connections to the new active Supervisor

Switchover Event – Hardware Notification

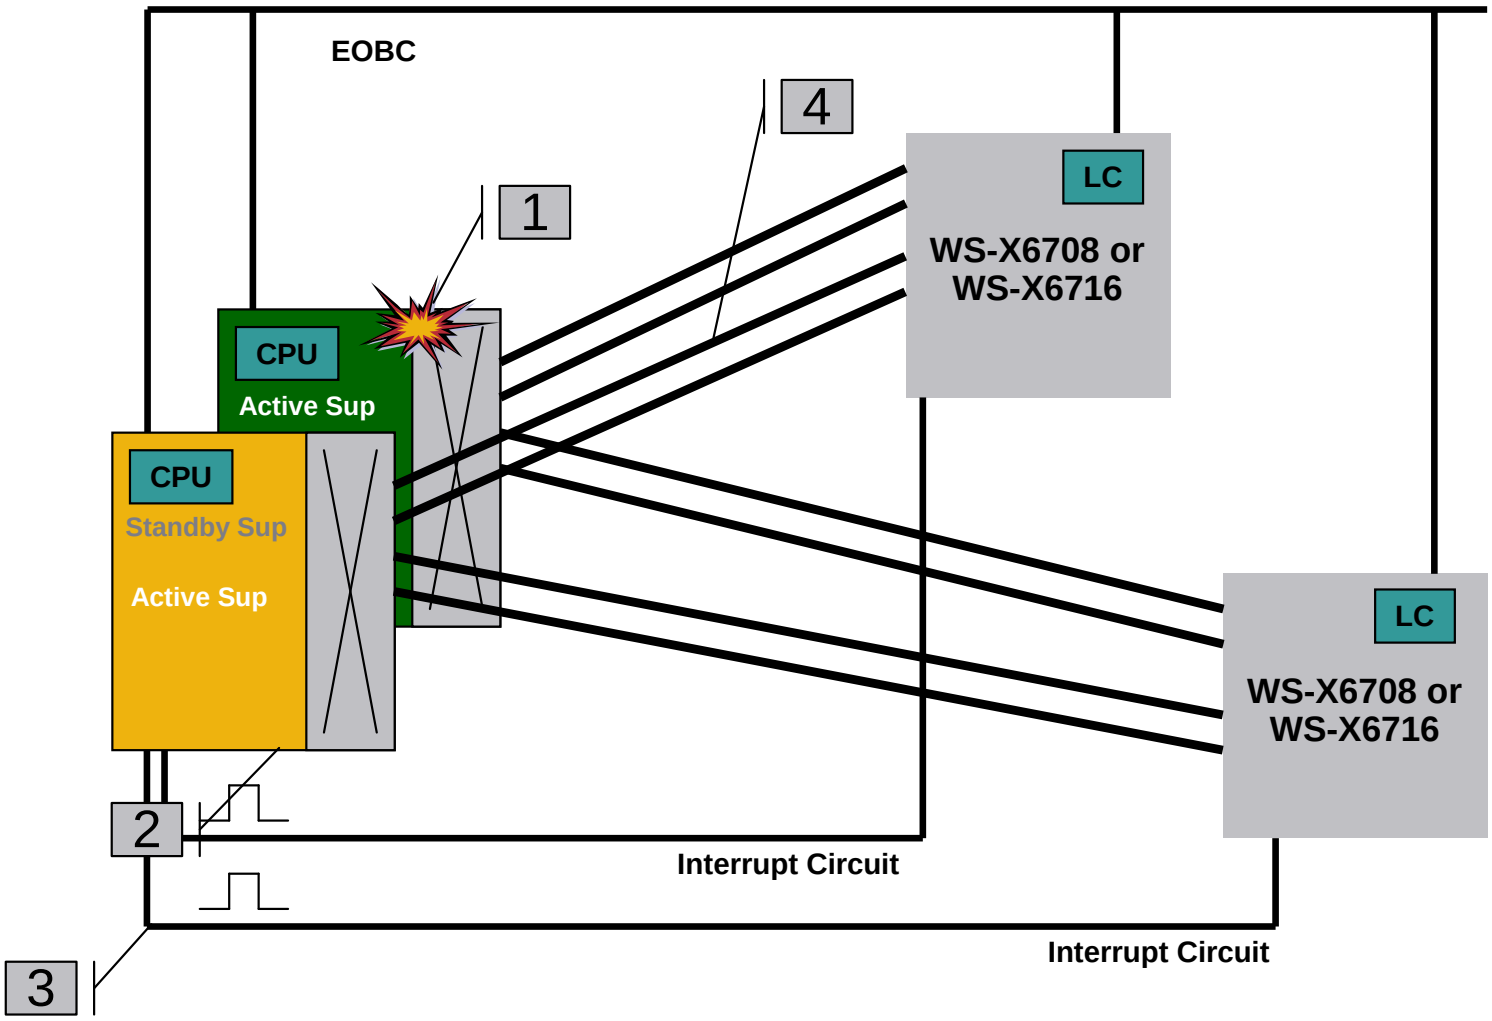

1. Supervisor switchover event triggered
2. Standby Supervisor becomes active
3. Standby notifies line cards of switchover event via hardware interrupt circuit
4. Line cards switch their fabric connections to the new active Supervisor

E-series chassis required
12.SXH required

Catalyst 6500 SSO Switchover Improvements in 12.2(SXH)

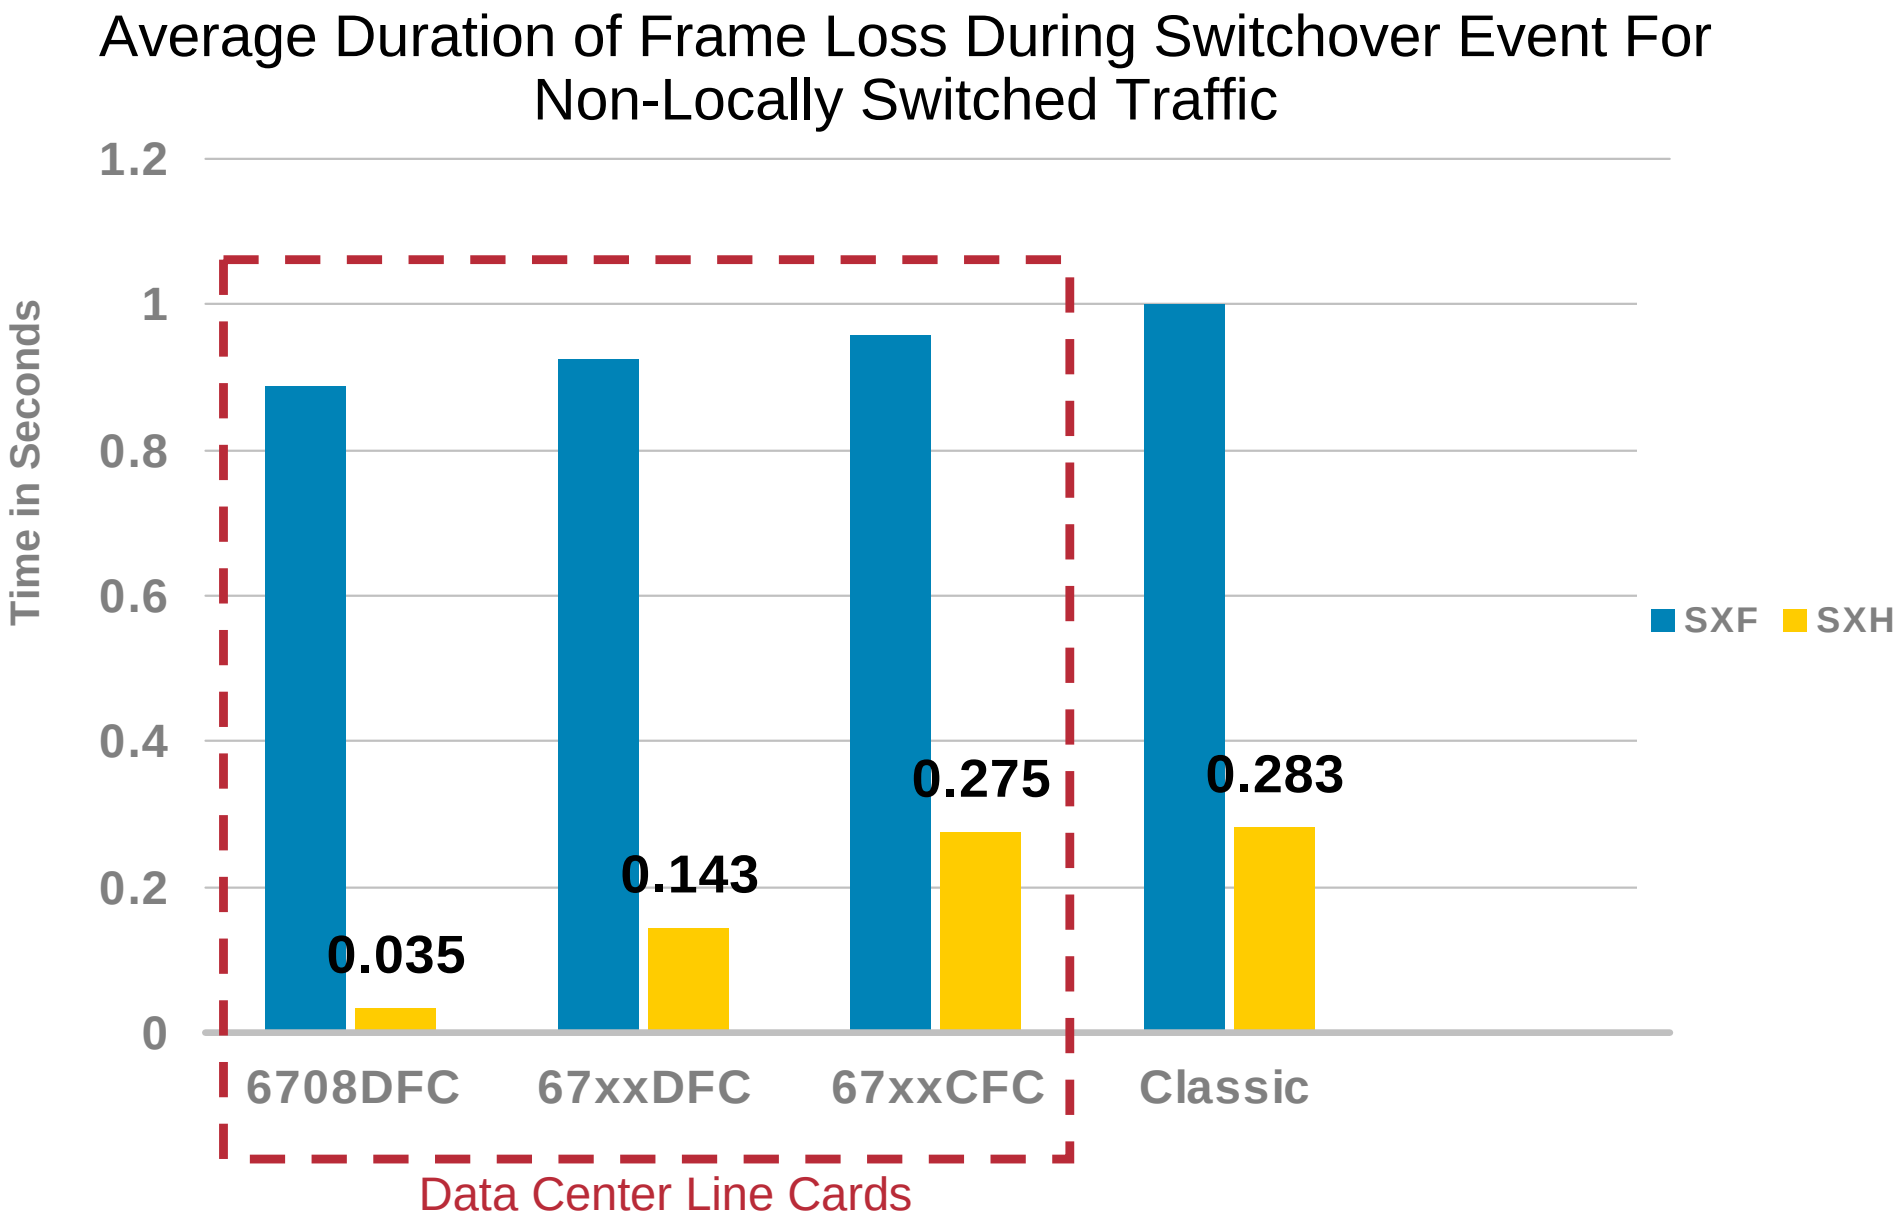

SSO Redundancy Clients

Maintain Protocol States Between Supervisors

RF_INTERNAL_MSG
Cat6k Platform Swove
Redundancy Mode RF
CHKPT RF
Event Manager
Cat6k Fabric Manager
Cat6k OIR
L3 Mobility Manager
Cat6k QoS Manager
Cat6k CWAN HA
CWAN VLAN RF Client
Cat6k Feature Manage
TSPTUN HA
Network RF Client
Cat6k SPA TSM
Cat6k Online Diag HA
XDR RRP RF
CEF RRP RF

Cat6k CWAN Interface
CWAN LTL Mgr HA RF C
Cat6k VPN RF Client
Tableid HA
Config Sync RF clien
Cat6k Local Target L
RF VS Client
LAN-Switch VTP VLAN
802.1x authenticator
LAN-Switch Port Mana
Cat6k Platform
Cat6k Power
Frame Relay
HDLC
LSD HA Proc
MFI STATIC HA Proc

IPROUTING NSF RF cli
GLBP
HSRP
PPP RF
Cat6k Startup Config
C6K_provision_rf_cli
Cat6k IDPROM
MPLS VPN HA Client
SNMP RF Client
ATM
History RF Client
RSVP HA Services
FH COMMON RF CLIENT
SNMP HA RF Client
LDP HA
IPRM
ARP
FH_RF_Event_Detector

Cat6k PAgP/LACP
Spanning-Tree Protoc
Cat6k Multicast Shor
Cat6k Layer3 Manager
Cat6k CAPI
CWAN SRP RF Client
CWAN APS HA RF Clien
Cat6k MLS Multicast
Core <-> satellite u
AC RF Client
AToM manager
SSM
SLB RF Client
Switch SPAN client
Switch Backup Interf
DHCP Snooping
IP Source Guard
MWAM RF Client

LAN Switching IP Host T
IKE RF Client
IPSEC RF Client
IP Admission RF Clie
RF_TS_CLIENT
Network RF 2 Client
Cat6k Clear counter
DATA DESCRIPTOR RF C
DATA DESCRIPTOR RF C
Cat6k HA Vpnsm
RF_LAST_CLIENT

Red font = New in 12.2SXH
Gold font = Key differentiator

```
Router#show redundancy clients
 clientID = 0       clientSeq = 0       RF_INTERNAL_MSG
 clientID = 1319    clientSeq = 1       Cat6k Platform Swove
 clientID = 5030    clientSeq = 2       Redundancy Mode RF
```

HSRP & GLBP SSO-Aware

- **HSRP and GLBP states are synchronized between the active and standby Supervisors**
- **FHRP timers must be set higher than SSO switchover time**
- **Testing for SRND configurations planned**
- **Available in 12.2(33)SXH**
- **VRRP SSO-aware planned for Whitney 2**

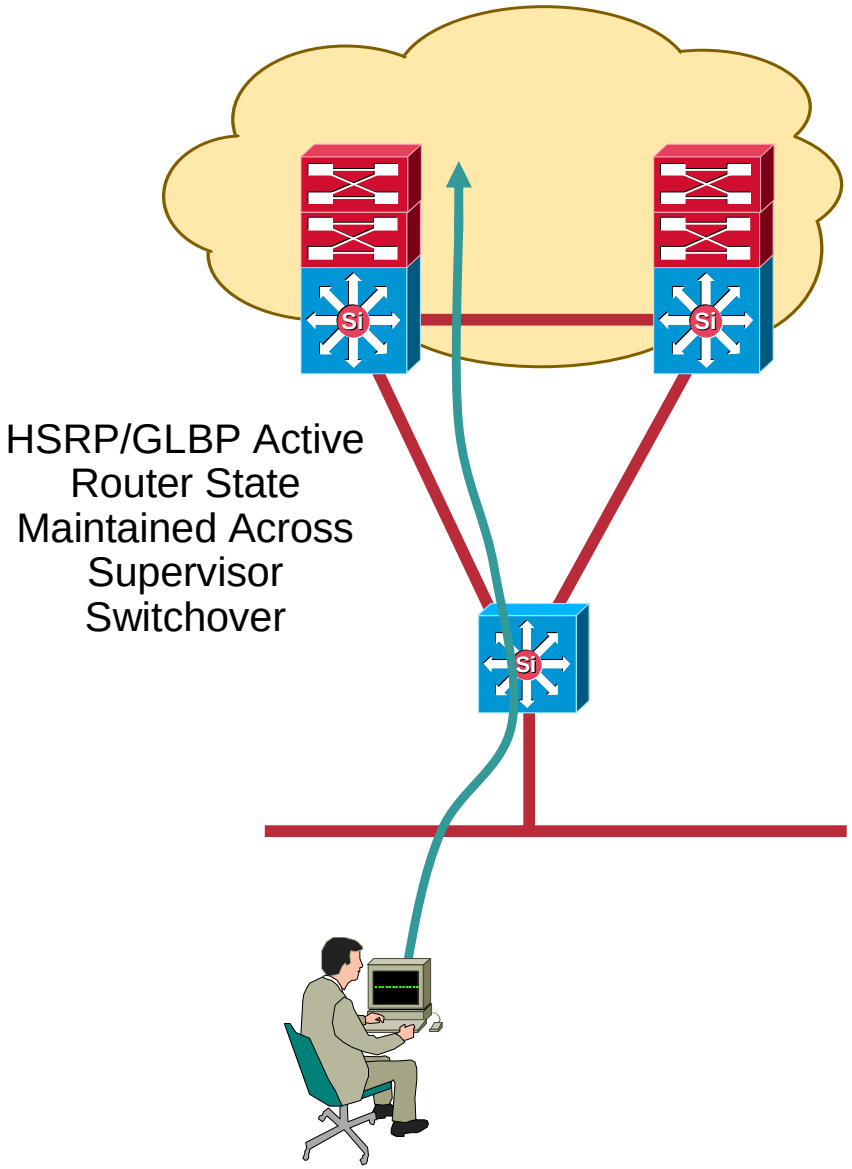

GOLD
TestUnusedPortLoopback

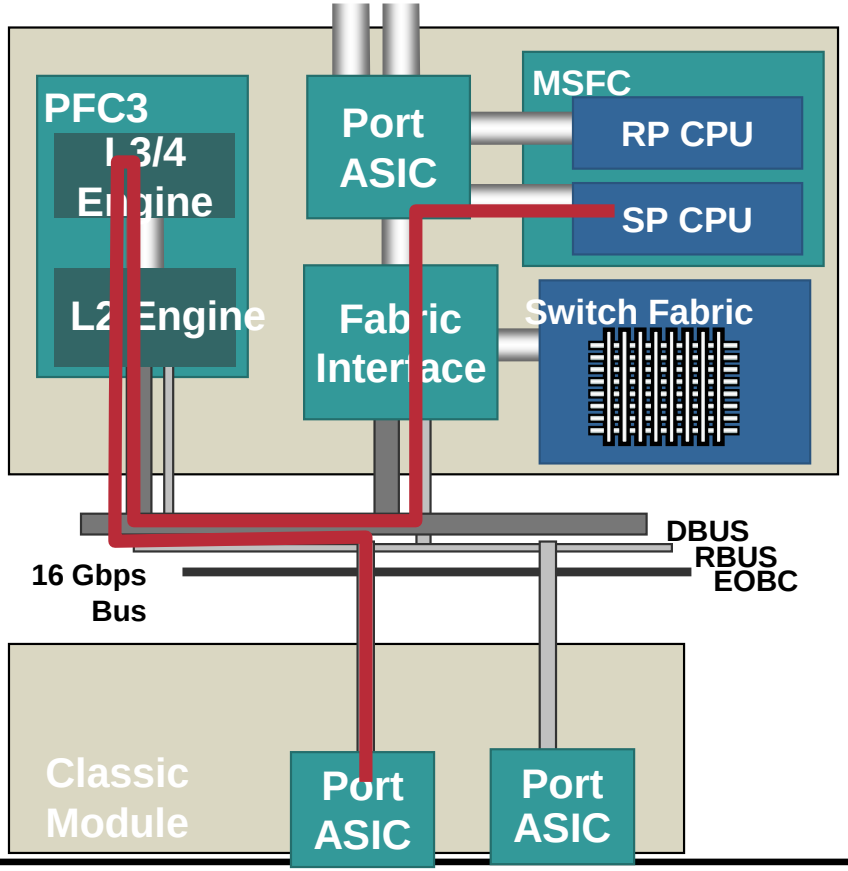

- User selects unused ports for loopback test
- Verifies the data path between the supervisor and the looped back port
- This is helpful for line cards where TestPortLoopback is a disruptive test
- Available in 12.2(33)SXH

```
Router# diagnostic start mod 2 test 32 port 16
Router#
*Apr 13 01:19:24: %DIAG-SP-6-TEST_RUNNING: Module 2: Running TestUnusedPortLoopback{ID=32}
...
*Apr 13 01:19:25: %DIAG-SP-6-TEST_OK: Module 2: TestUnusedPortLoopback{ID=32} has completed
successfully
Router#
```

On Board Failure Logging

```
Router# show logging onboard mod 3
PID: WS-X6724-SFP      , VID:    , SN: SAD08250GDU


----------------------------------------------------------
-
UPTIME SUMMARY INFORMATION
----------------------------------------------------------
-
First customer power on : 01/05/2006 14:34:10
Total uptime            :   1 years  11 weeks   3 days
Total downtime          :   0 years   0 weeks   0 days
Number of resets        : 54
Number of slot changes  : 1
Current reset reason    : 0xA1
Current reset timestamp : 04/06/2007 16:09:22
Current slot            : 3
Current uptime          :   0 years   0 weeks   6 days
----------------------------------------------------------
-
Reset  |        |
Reason | Count  |
----------------------------------------------------------
-
0xA1         54
```

- Like a Black Box on an airplane
- OBFL stores critical diagnostic information on the line card
- Provides key data for troubleshooting and failure analysis
- Information includes
  - Total runtime
  - Boot time
  - Environmental temp/volt
  - Diagnostic failures
  - …more
- Available on 67XX-series line cards
- Phase 1 Q2CY07 (12.2.33(SXH)
- Provides support for Cisco TAC and Engineering commands
- Phase 2 will support user commands

System Event Archive (SEA)

- Allows all modules in the chassis to record events
  - Anomaly conditions such as OIR
  - Major and Critical events
  - Configuration changes
- SEA is a secure file and resides on the active Supervisor SP-bootdisk:
- Line Cards communicate messages to the active SP via the EOBC
- Will consume 10% of flash space
- Running log retained across reboots
- Message throttling is supported

```
Router# show logging system
SEQ: MM/DD/YY HH:MM:SS MOD/SUB: SEV, COMP,    MESSAGE
=================================================

  1: 04/06/07 16:10:07   1/-1 : MAJ, GOLD, online_diag_f
when testing [4/-1/37]
  2: 04/06/07 16:10:03   1/-1 : MAJ, GOLD, online_diag_f
when testing [4/-1/25]
  3: 04/06/07 16:10:00   1/-1 : MAJ, GOLD, online_diag_f
when testing [4/-1/13]
  4: 04/06/07 16:09:56   1/-1 : MAJ, GOLD, online_diag_flush_pak_queue: flushing a packet from [4/1]
when testing [4/-1/1]
  5: 04/06/07 16:09:38   1/-1 : MAJ, GOLD, TestErrorCounterMonitor: ID:42 IN:0 PO:255 RE:200 RM:255
DV:1 EG:2 CF:1 TF:1
  6: 04/06/07 16:09:30   1/-1 : MAJ, GOLD, online_diag_flush_pak_queue: flushing a packet from [3/13]
when testing [3/-1/13]
  7: 04/06/07 16:08:47   1/-1 : MAJ, GOLD, test_rp_fib_sc[1]: missing ping L2 entry [2]
```

Digital Optical Monitoring

X2 XENPAK SFP

- Provides real-time information on key transcievr operating parameters
- TX power
- RX power
- Temperature
- Voltage
- Laser bias current
- Provides default alarm thresholds
- Supported on CWDM and DWDM transceivers
- List of DOM-supported transceivers available in the cli and online

```
Router# show interfaces transceiver
If device is externally calibrated, only c
values are printed.
++ : high alarm, + : high warning, - : low
: low alarm.
NA or N/A: not applicable, Tx: transmit, R
mA: milliamperes, dBm: decibels (milliwatt


Optical Optical
Temperature Voltage Current Tx Power Rx Powe
Port (Celsius) (Volts) (mA) (dBm) (dBm)
------- ----------- ------- -------- -------- --------
Gi1/1 40.6      5.09    0.4  -25.2 N/A
Gi2/1 35.5      5.05    0.1  -29.2 N/A
Gi2/2 49.5      3.30    0.0   7.1 -18.7
Router#
```

Netflow Enhancements in 12.2SXH

Key differentiator versus the competition

- **Netflow Per-interface support**
  - **Reduces data export volumes**
  - `ip flow ingress`
- **Netflow for IPv6 unicast traffic**
- **Netflow MIB and top talkers**
- **Netflow Multicast IP support**

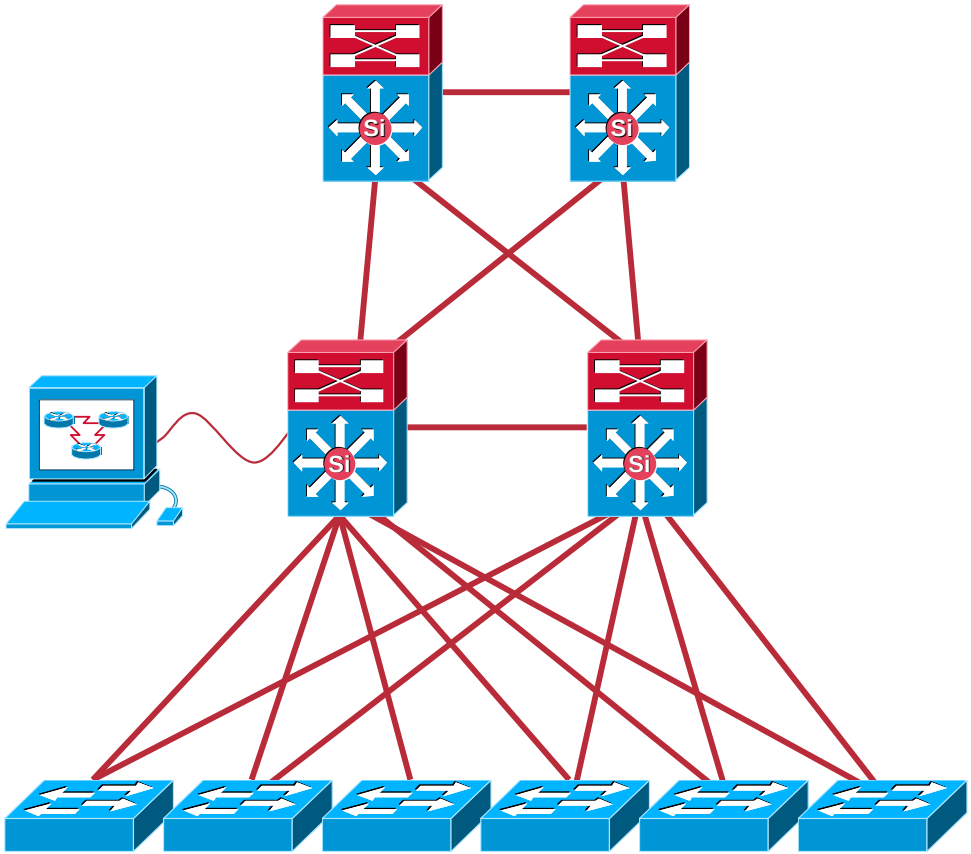

WAAS Update

What differentiates us today

Differentiators	Architecture / Feature	Customer Testimonial
Lowest TCO in Branch & DC; Best Scalability	• Up to 10x fewer servers required in DC • Router Module reduces branch foot print • Lower Support Costs	• ***"WAAS offered us a scalable solution, accelerating all the applications including Exchange, SQL, file and printing services, SAP, and our freight-forwarding system accessed through Citrix,"*** • ***Bill Fraser, VP of Ops*** PANALPINA on 6 continents
Ease of Operation	• Network Transparency • Auto-discovery support for all topologies	• ***"The tunnel-less Cisco WAAS design makes it very easy for me to network-based services so I can remain focused on projects that contribute to the company's success."*** **- Abraham Madha, CIO** ENGlobal
Better w Voice	• Support for Integrated IOS QoS Architecture for Voice & Data	• ***"Riverbed slowed down the VoIP, and voice quality is Number One"*** **- Dave Kizer, Director** nanometrics
Better w Security	• Full Firewall/IPS support w Stateful TCP Inspection • Encryption of Cache Disk	• ***"The latest version of Cisco's WAAS product line includes security features of special interest to us, including disk encryption and support for stateful firewall protection."*** **- Barth Bailey, VP** A BALANCED DESIGN FOR GROWTH FULTON FINANCIAL CORPORATION

Committed to Win the "Technology Race"

- **WAAS 4.0.1** **Aug'06**
  - **L4 optimizations + CIFS adapter**
  - **From WAFS to WAAS**
- **WAAS 4.0.3** **Dec'06**
  - **ISR Network Module**
- **WAAS 4.0.7** **Mar'07**
  - **Inline card + CIFS auto-discovery**
- **WAAS 4.0.13** **Aug'07**
  - **Security, Scale, Performance Monitoring**

4 Release in 12 Months and More to Come…

WAAS 4.0.13: “Trusted WAN Optimization”

- **Secure WAN Acceleration**

  Preserves security while accelerating applications

- **Highest Scalability**

  New WAE-7341/71s offer industry’s highest single device scale. Extend lead in clustered device scale.

- **Monitoring & Performance Management**

  Real-time monitoring; Interoperability with NetQoS for accurate end-to-end application performance reporting

Security

Accuracy

Scalability

NEW

Highest Scalability and Performance

CLEAR CHOICE TEST NETWORKWORLD

Be Aware of Vendor Claim vs. Reality

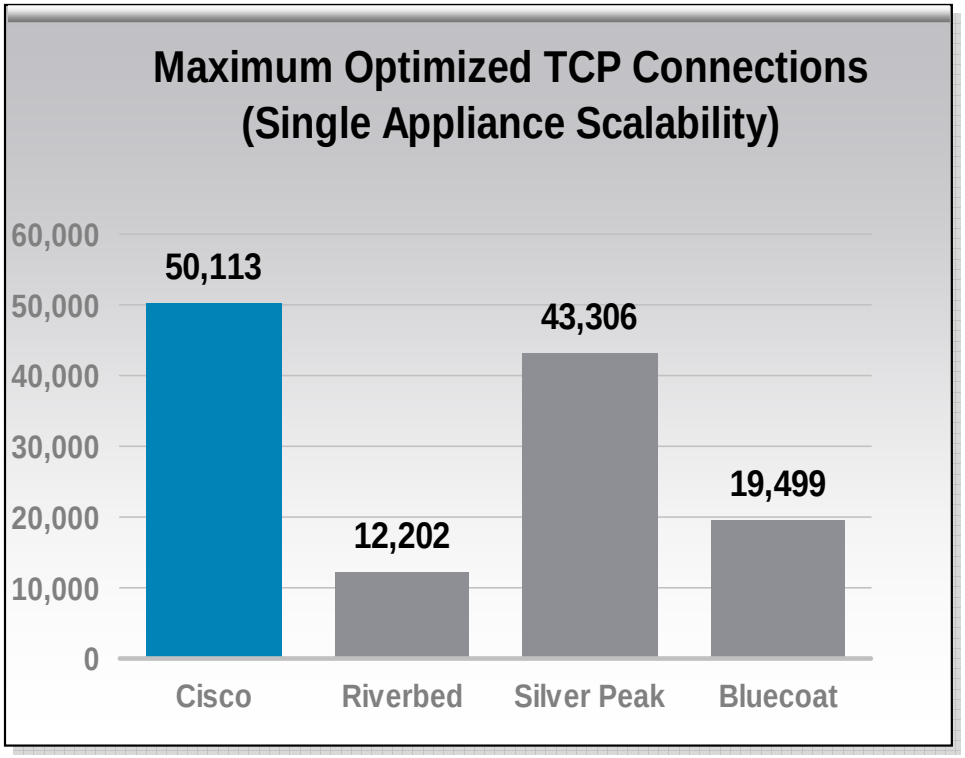

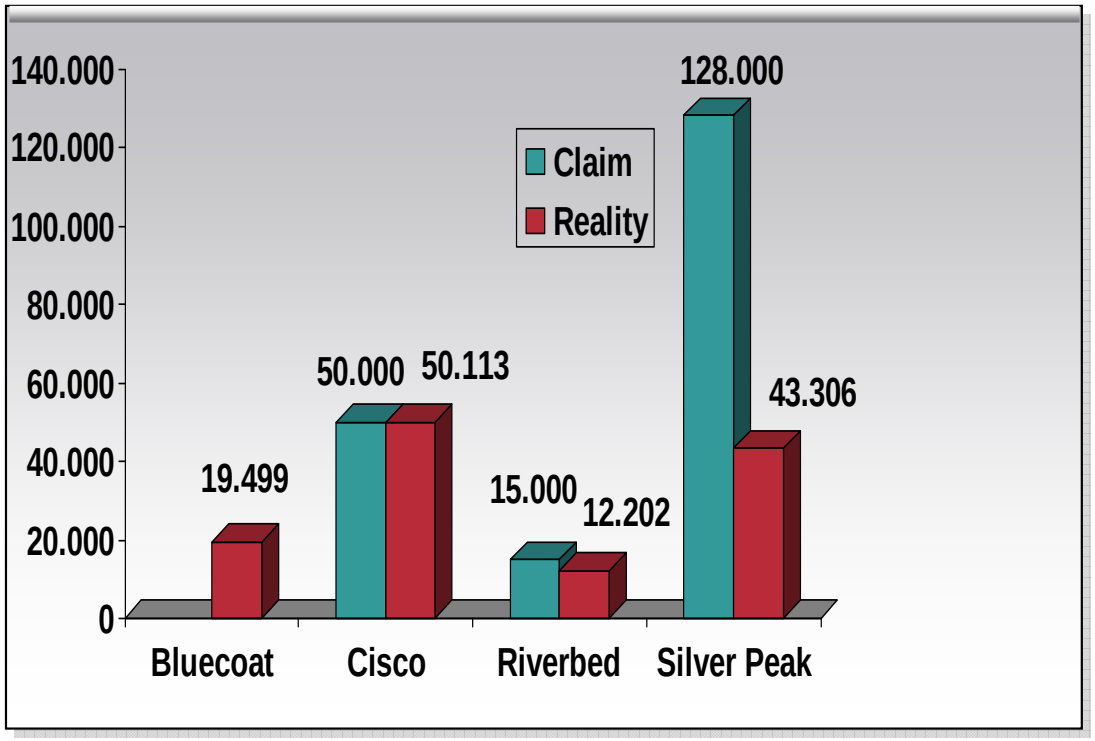

"We urge you to focus on optimization, scalability, security, and mobility criteria to determine the most appropriate solution."

- Robert Whiteley, Forrester Research

Secure WAN Acceleration
At Rest

NEW

The Costs of Data Leakage:

- **The average cost of a data breach per record compromised grew 30%, averaging $4.8 million/breach.**

  Source: Ponemon Institute, "2006 Annual Study: Cost of a Data Breach," October 2006, "768 IT Managers Survey"

- **$355 for a high-profile breach in a highly regulated industry.**

  Source: Forrester Research: "The Cost of a Security Breach"

Cisco WAAS Strong Encryption:

- **256-bit Advanced Encryption Standard evaluated by Common Criteria**
- **No keys left on disks**
- **Centralized key management with backup, restoration and high availability**

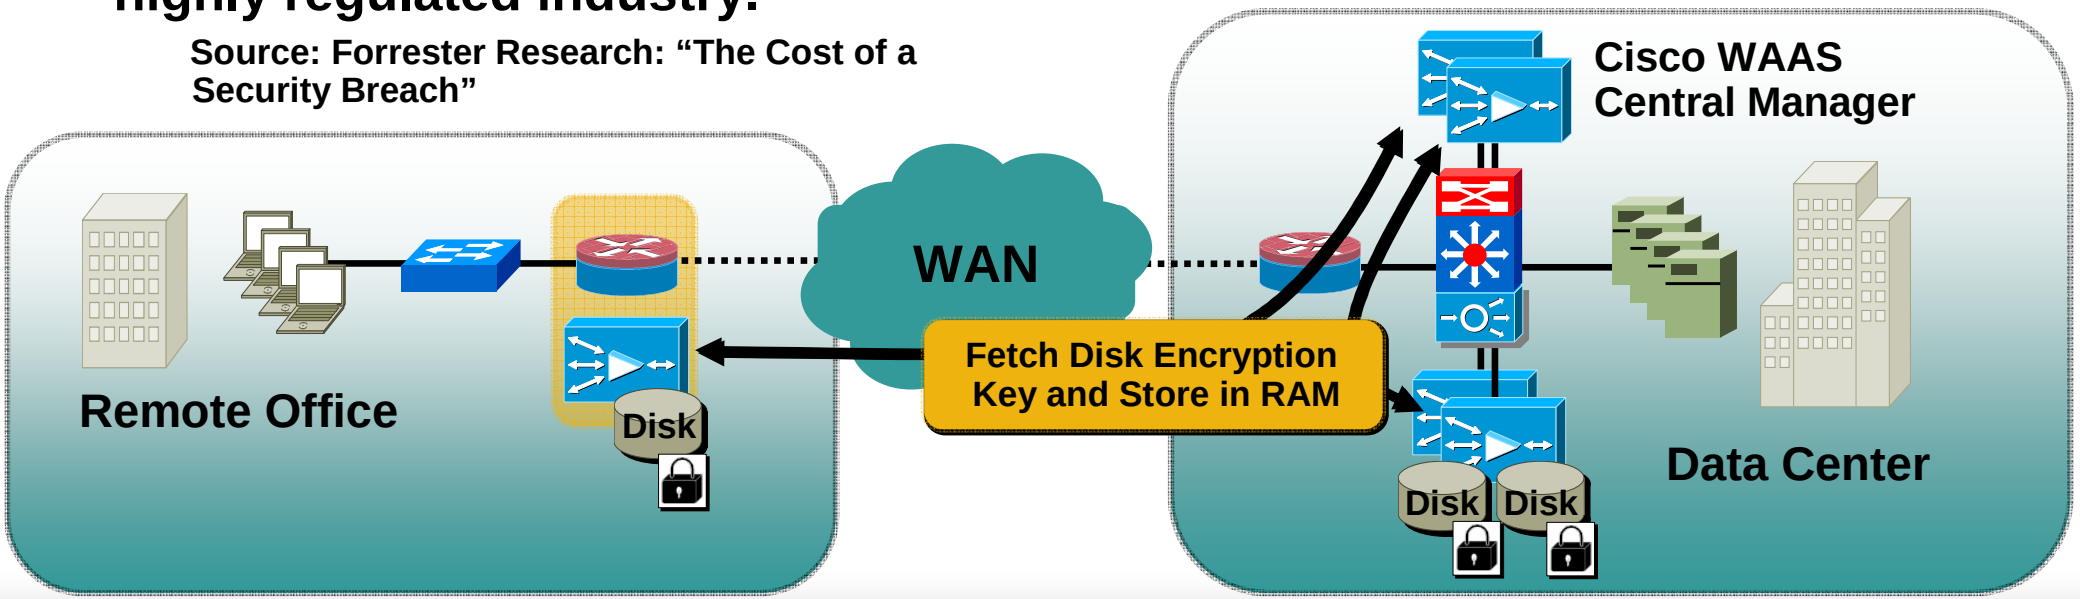

Secure WAN Acceleration
In Flight

NEW

Cisco

Stateful Protection of Each Accelerated Stream

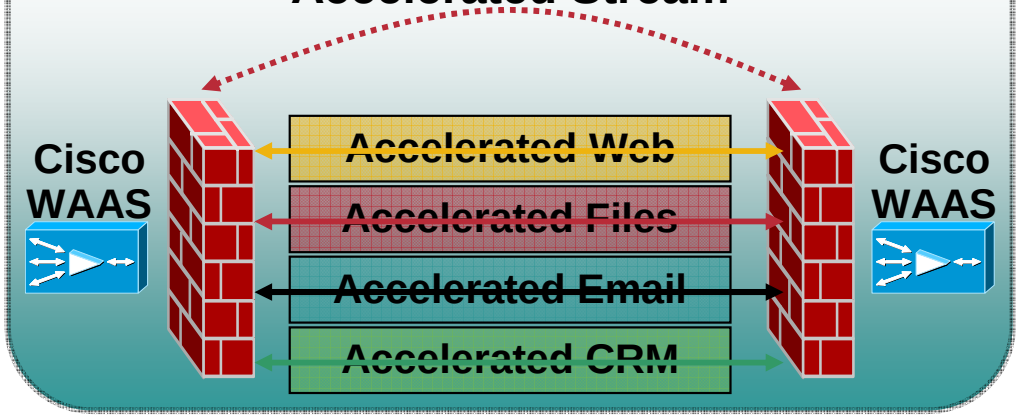

Tunnel Solutions

No Stateful Protection of Each Accelerated Stream

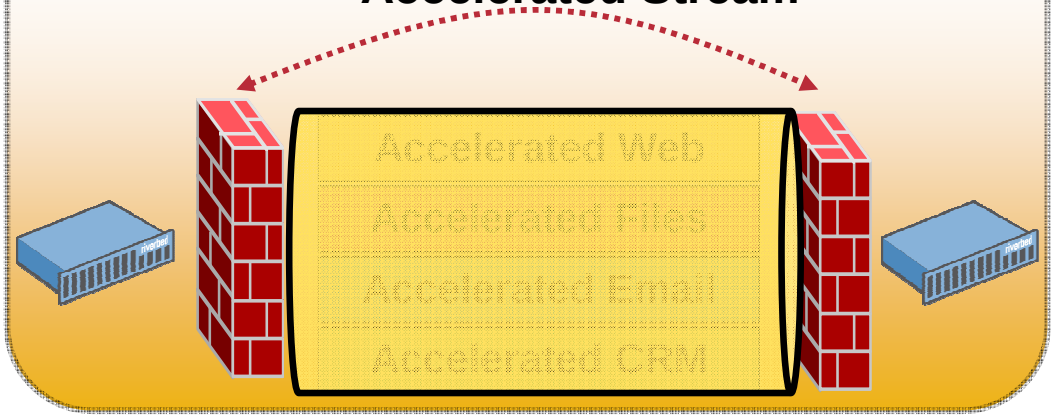

"Cisco WAAS is the most transparent and easiest solution to implement and integrate as it does not require us to overhaul our security infrastructure."

—Mark Drake, Network Engineer, Health Management Associates, Inc

HEALTH MANAGEMENT ASSOCIATES, INC.

NEW

WAAS + NetQoS = E2E App Performance

- **WANopt solutions "segment" TCP flows**

  Performance monitoring tools produce distorted reports

- **WAAS provides "calibration data" to NetQoS**

  WAAS "TCP Flow Export Agent" ensures end-to-end visibility for NetQoS SuperAgent

- **WAAS and NetQoS SuperAgent Provide**

  Accurate end-to-end application response time monitoring

  Reliable SLA management

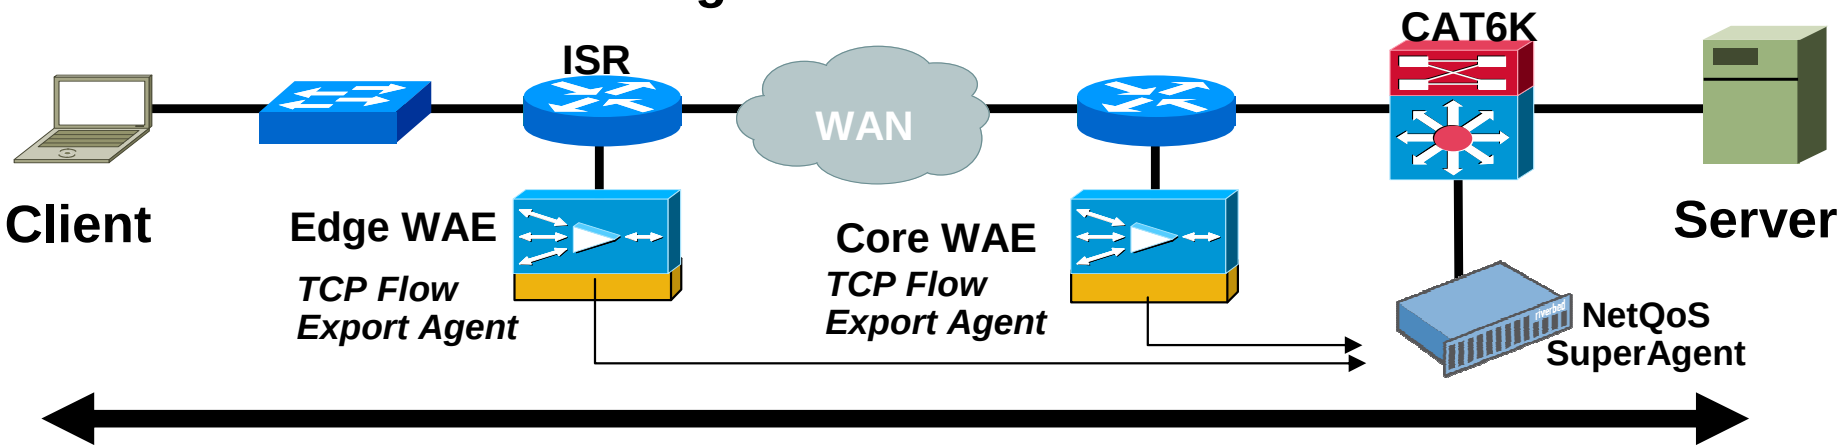

Accurate App Performance Measurement

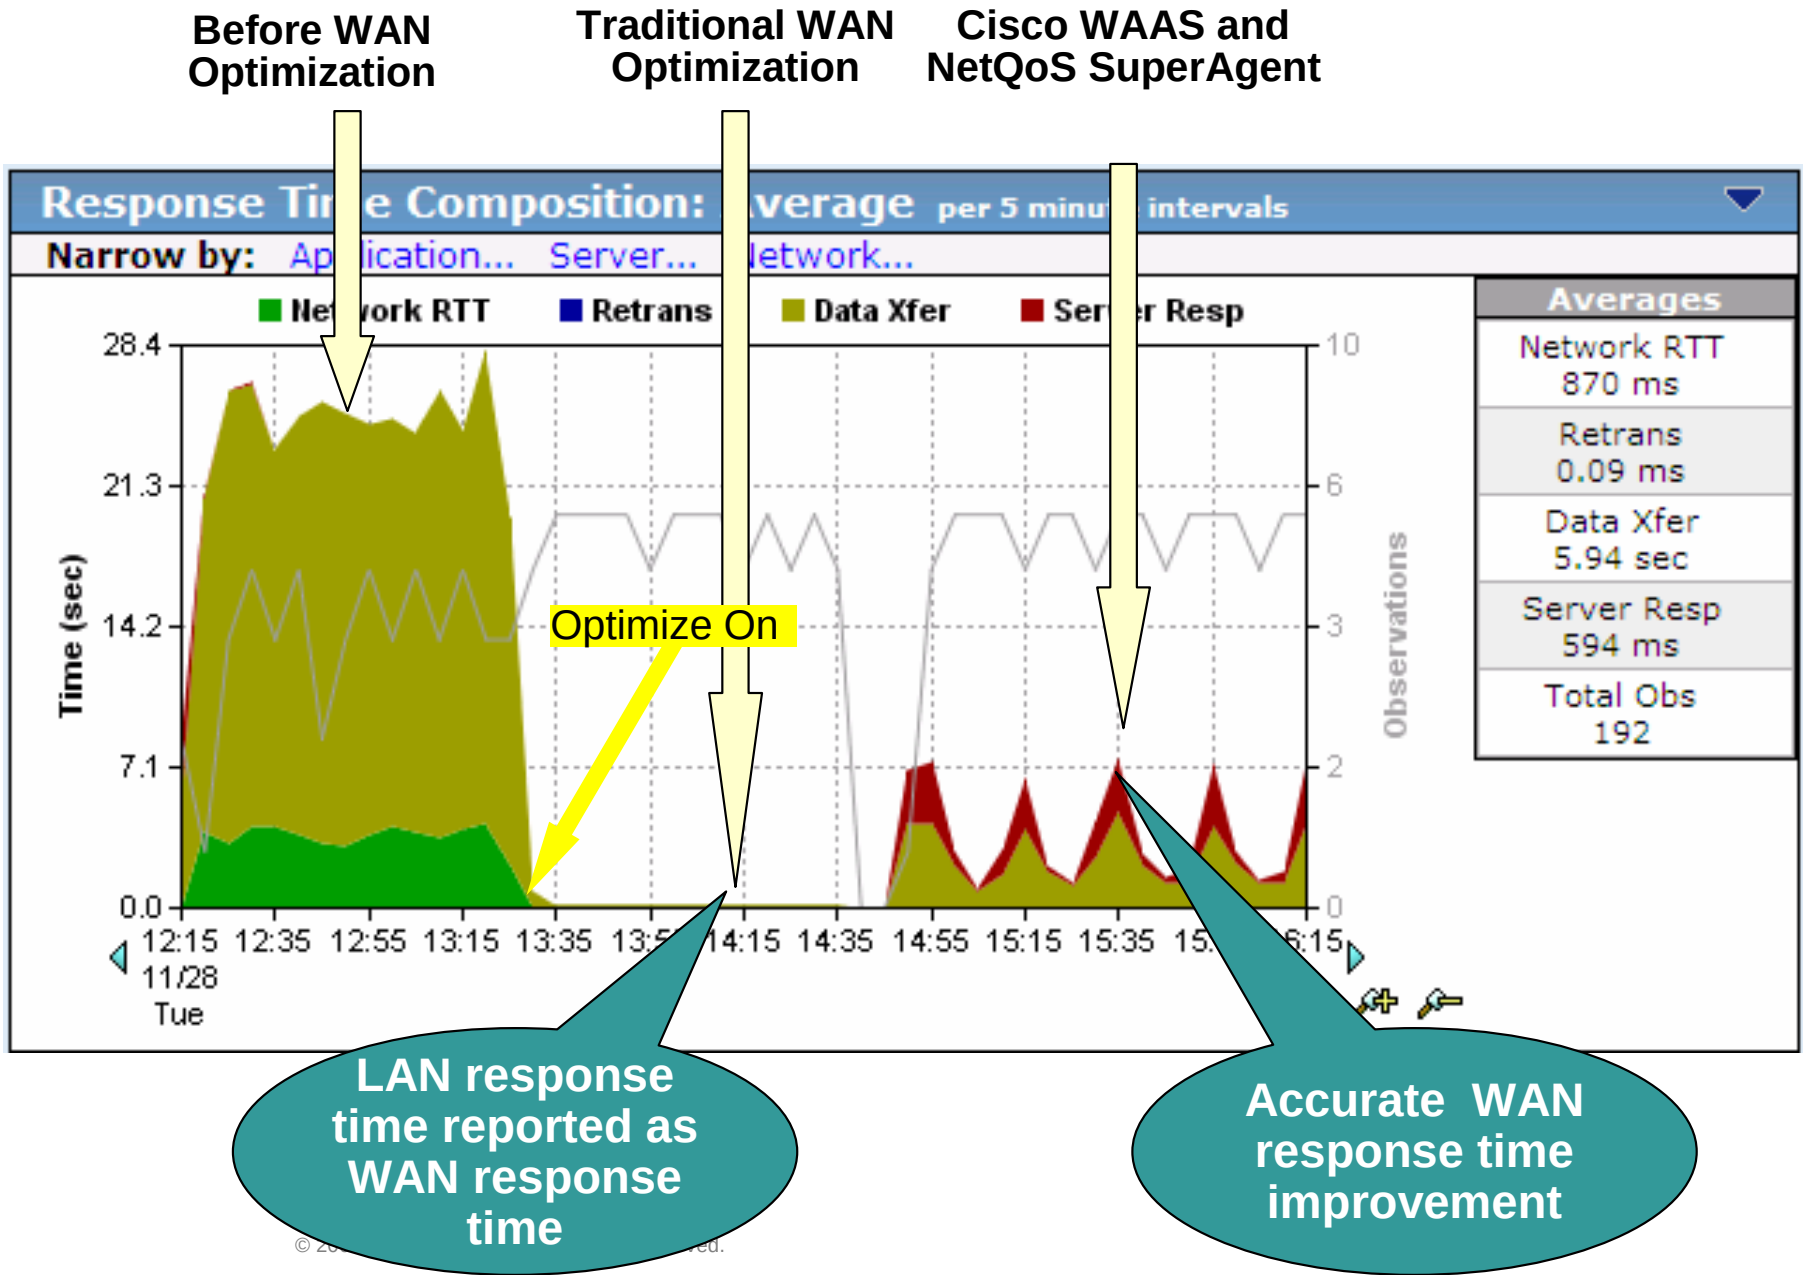

Application Testing and Tuning TCP Optimization and Advanced Compression

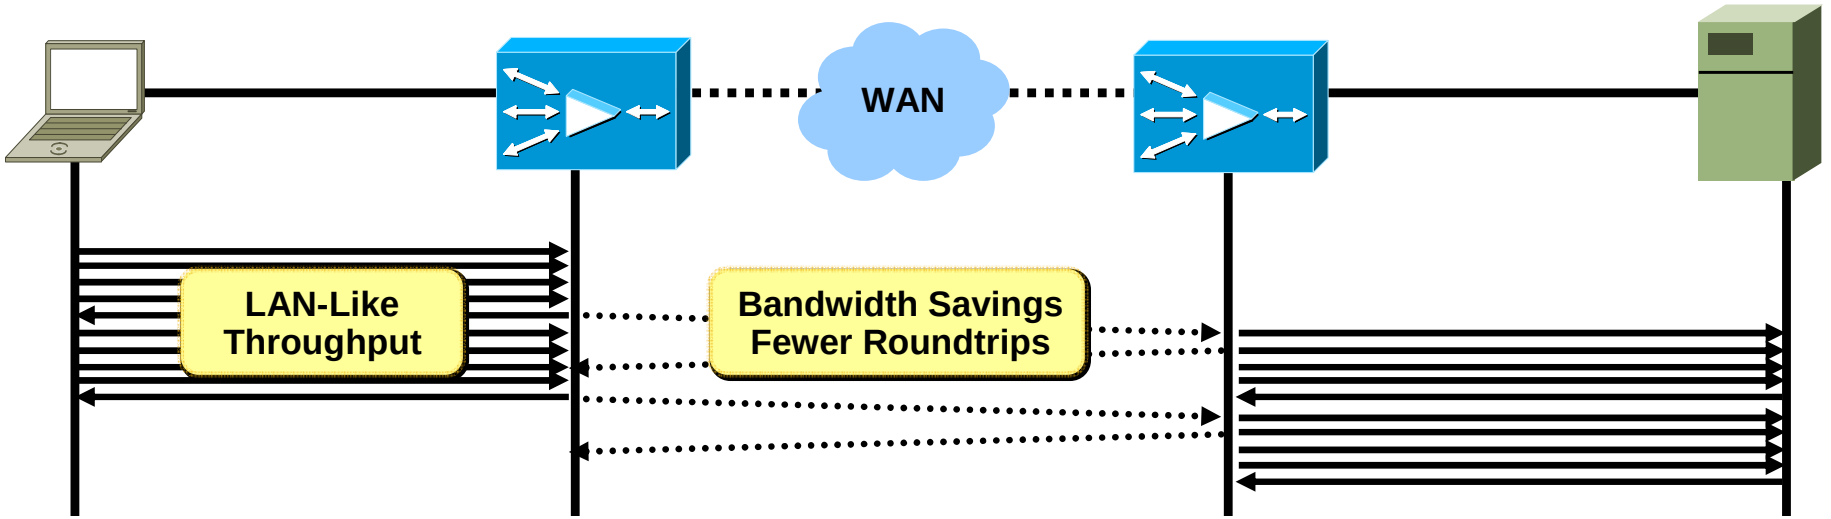

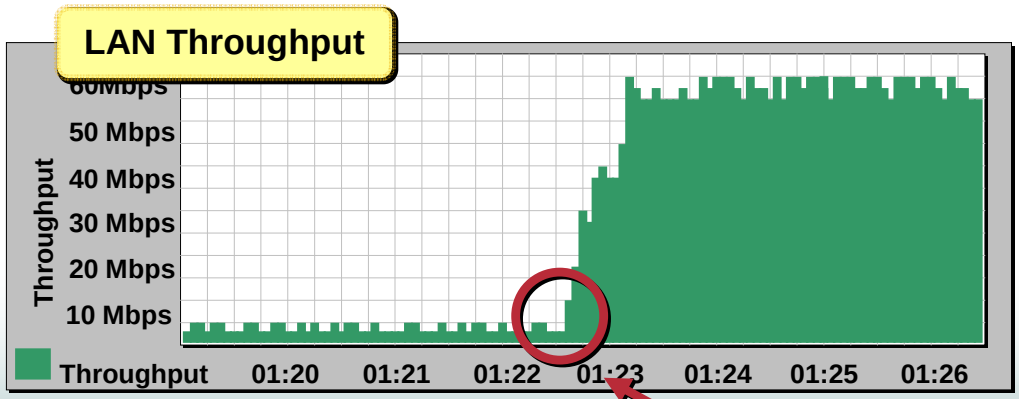

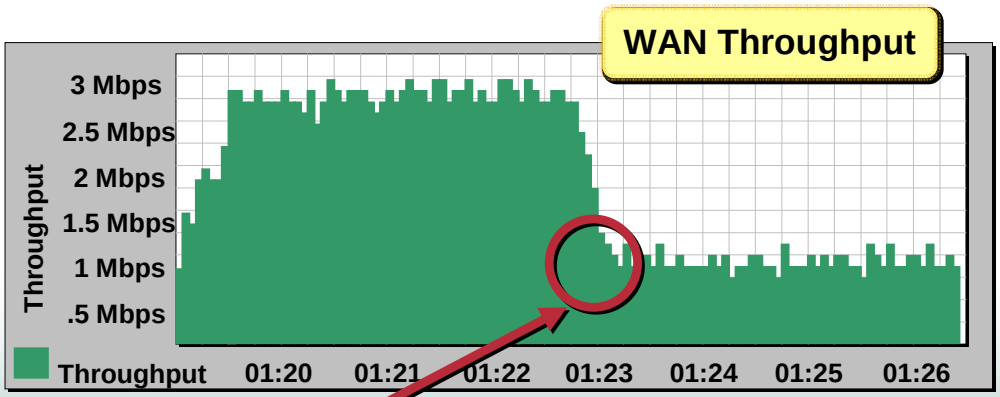

Optimization Enabled

Branch Network Deployments

Non-Redundant Branch

- **Inline**
- **On Host Subnet**
- **Network Module**
- **Off-Router Deployment**

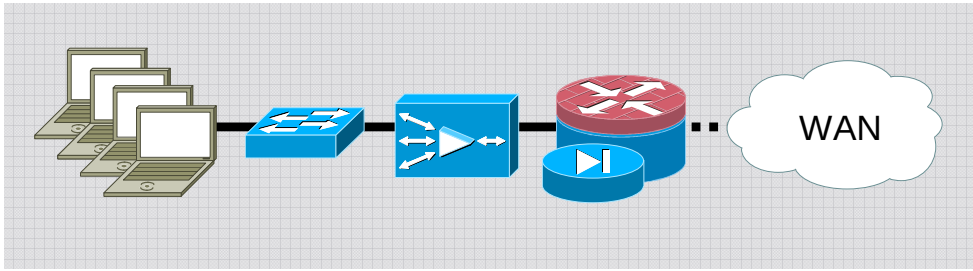

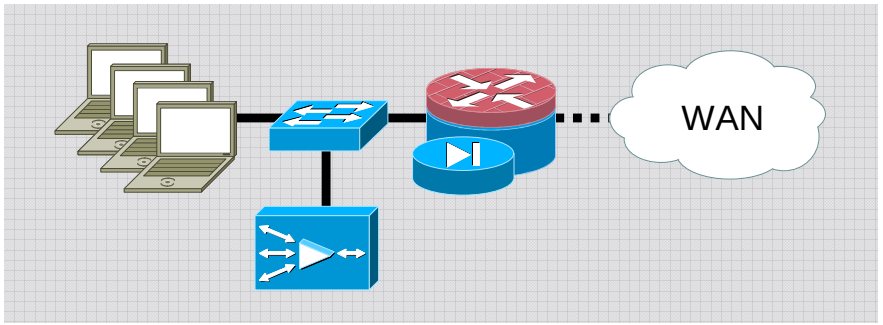

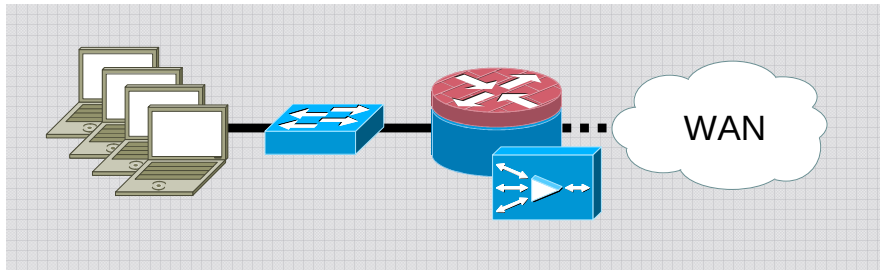

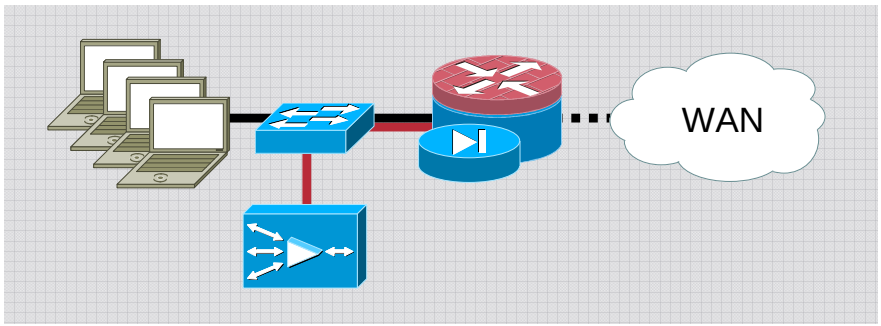

WCCP

Common Configuration Topologies

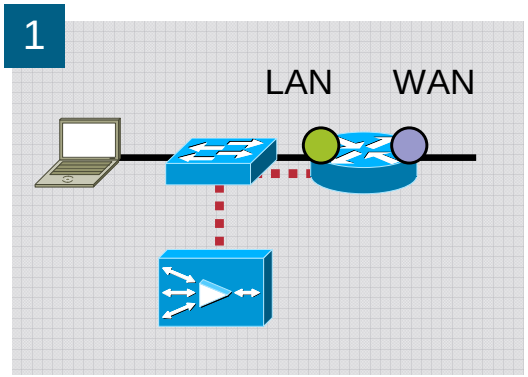

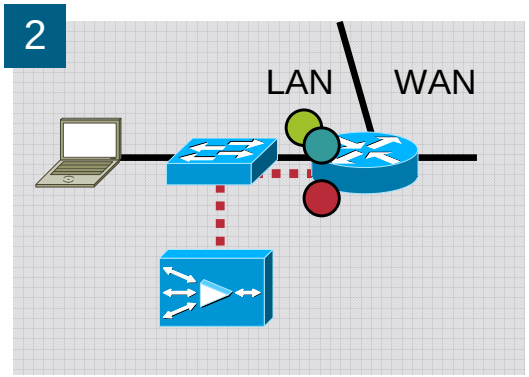

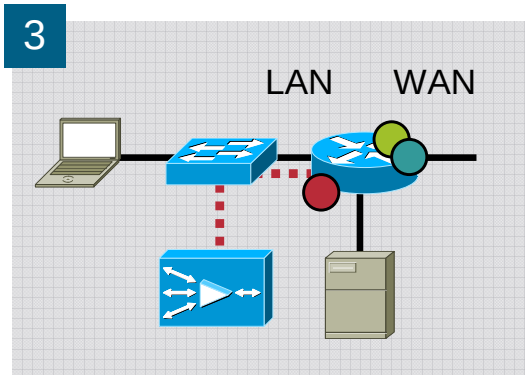

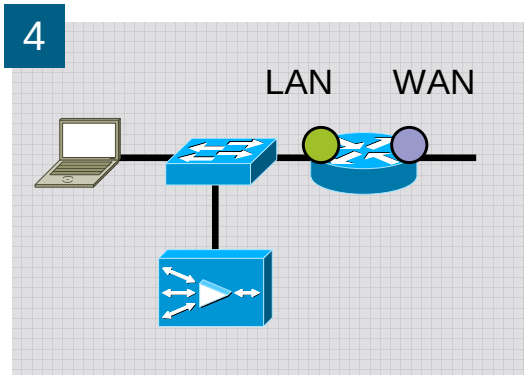

1. **Common one-arm configuration**
2. **WAN router is transit to another site**
3. **Local servers on separate branch subnet**
4. **WAE on host subnet using GRE return in WAAS 4.0.13**

61/out 62/out
61/in 62/in
Redirect Exclude

Branch Network Deployments

Non-Redundant Host Subnet Deployment

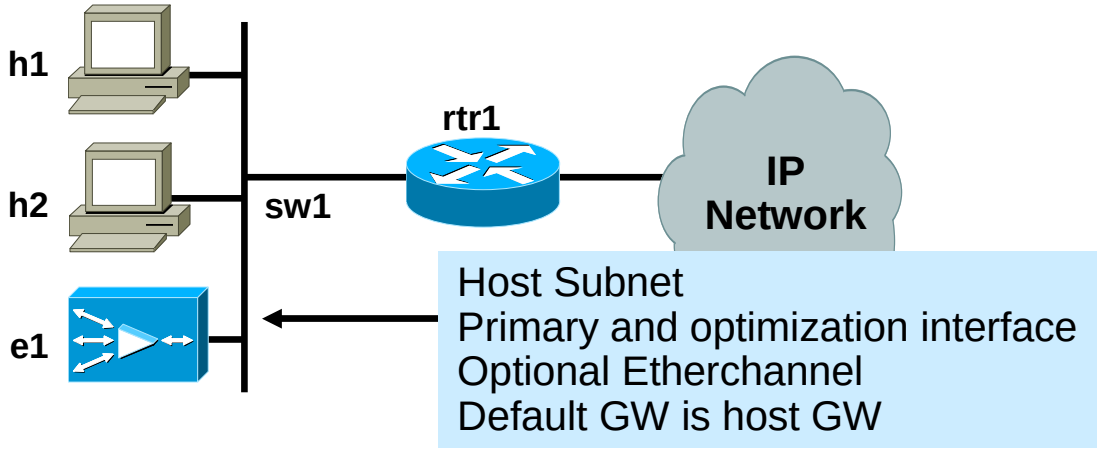

- **Engine**
  - **Requires WCCP with WAAS 4.0.13 or later**
  - **GRE forwarding with stateful negotiated WCCP GRE return**
  - **Primary interface used for management and optimizations**
  - **Default gateway is same as hosts**
  - **PortChannel supported**
  - **IP forward return for router offload**
    - **Local hosts (WAAS 4.0.13)**
    - **Specified hosts (Future)**
- **Router/Switch**
  - **Requires WCCP GRE capable router (no Catalyst 3750 or 4500)**

Branch Network Deployments

Non-Redundant Branch InLine

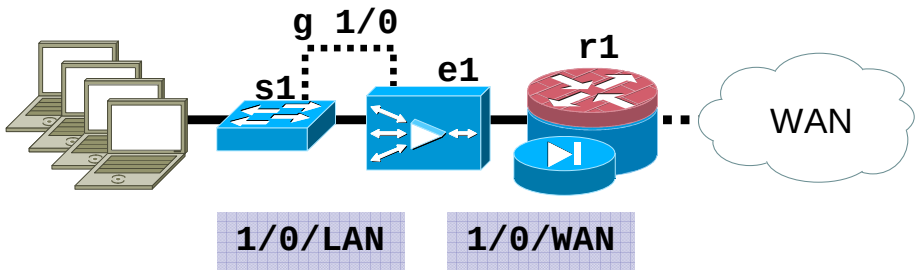

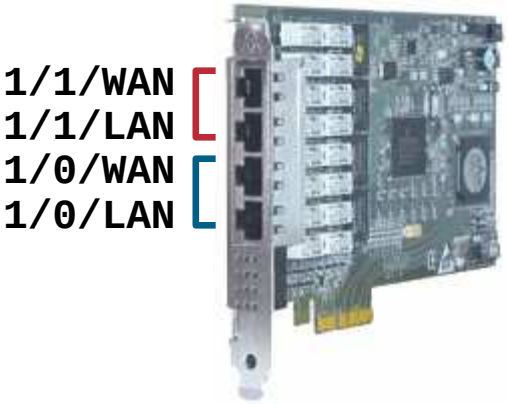

- **Engine**
  - **One InLine NIC per WAE appliance (cannot be used with WCCP)**
  - **Installed in-path between switch and router or firewall**
  - **Use single pair of inline ports (1/0 or 1/1) removing RJ45 port covers**
  - **Ports fail-to-wire upon hardware, software, or power failure**
  - **Support for interception 802.1q trunks**
  - **Ensure that WAN and LAN bandwidth and duplex are identically configured**
  - **Requires GigabitEthernet 1/0 for primary interface**
- **Router**
  - **Crossover cable from router to engine**
- **Switch**
  - **Straight through cable from engine to switch**

Branch Network Deployments

Network Module

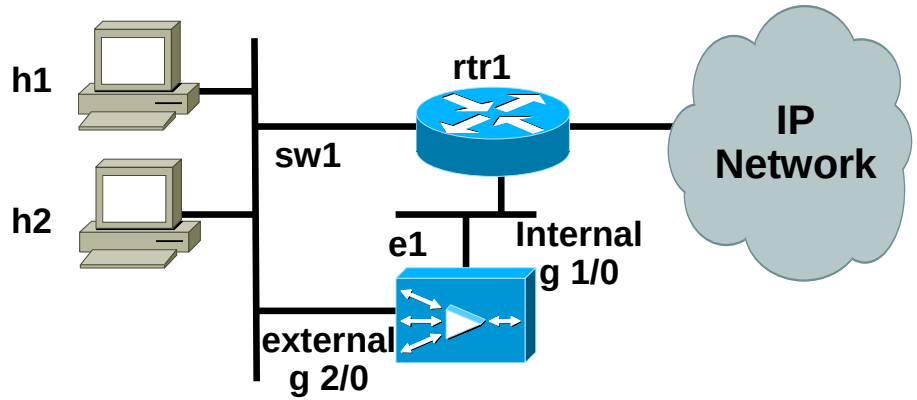

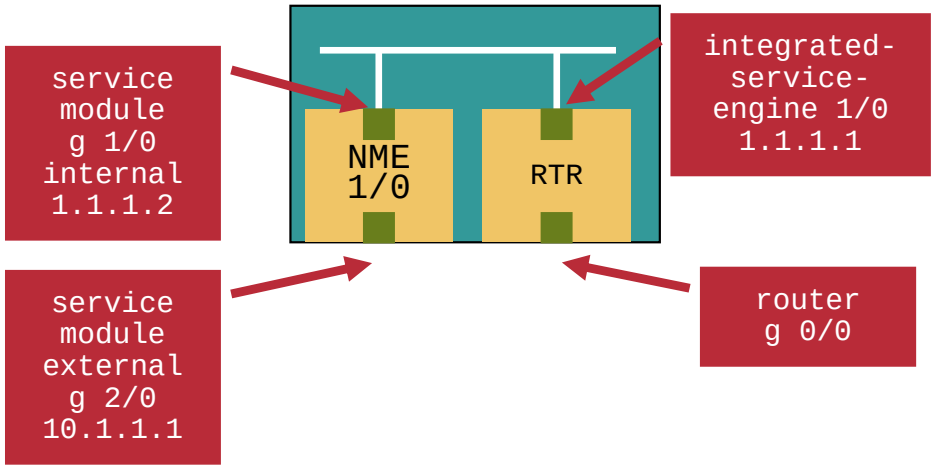

Router Configuration Commands

```
ip wccp 61 redirect-list waas
ip wccp 62 redirect-list waas
interface integrated-service-engine 1/0
 description Origin LAN
 ip address 1.1.1.1 255.255.255.0
 ip wccp 61 redirect in
 service-module IP address 1.1.1.2
   255.255.255.0
 service-module external ip address 10.1.1.1
   255.255.255.0
 service-module ip default-gateway 1.1.1.1
interface Serial0/0
 description Optimized WAN
 ip wccp 62 redirect in
```

Router Service Module Commands

```
service-module integrated-service-engine 1/0
 <status | session | shutdown | reload |
   reset | default-boot | statistics>
show interfaces integrated-services-engine
   1/0
show service-module Integrated-Service-
   Engine1/0
```

NME WAE Configuration

```
WAE(config)#wccp router-list 1 1.1.1.1
WAE(config)#wccp tcp-promiscuous router-list-
   num 1
WAE(config)#wccp version 2
```

Use external interface for Direct host return

Cisco WAAS Router Modules

NME-WAE
Router-Integrated Network Module for the Cisco Integrated Services Router

Cisco Integrated Services Router (ISR) Series

- **Provides the lowest CapEx and OpEx; integrates within the ISR; addresses 80 percent of remote branch offices**
- **Single processor system, can be clustered with WCCPv2, PBR, and is supported in ISR models 2811, 2821, 2851, 3825, and 3845**
- **Model NME-WAE-302**
  - **512MB of RAM, 80GB of disk**
  - **Up to 4Mbps WAN connections and up to 250 optimized TCP connections**

Model NME-WAE-502

- **1GB of RAM, 120GB of disk**
- **Up to 4Mbps WAN connections and up to 500 optimized TCP connections**

- **Model NME-WAE-522**
  - **2GB of RAM, 160GB of disk**
  - **Up to 8Mbps WAN connections and up to 800 optimized TCP connections**

Cisco WAAS Appliance Family

WAE-512 Appliance

WAE-512 Appliance

Single processor system with 250GB of RAID-1 protected SATA2 disk capacity and optional disk encryption

1GB memory configuration supports 8Mbps WAN connections and 750 optimized TCP connections

2GB memory configuration supports 20Mbps WAN connections and 1500 optimized TCP connections

WAE-612 Appliance

- ## WAE-612 Appliance

Dual-core processor system with 300GB of RAID-1 protected and hot-swappable SATA2 disk capacity and optional disk encryption

2GB memory configuration supports 45Mbps WAN connections and 2000 optimized TCP connections

4GB memory configuration supports 90Mbps WAN connections and 6000 optimized TCP connections

Cisco WAAS Data Center Appliances

WAE-7326
Enterprise Data Center Appliance

WAE-7341
Enterprise Data Center Appliance

WAE-7371
Enterprise Data Center Appliance

- **WAE-7326 Appliance:**
  - Dual-core processor w/ 4GB of RAM
  - Up to 155Mbps WAN connections and 7500 optimized TCP connections
  - 900GB RAID-1 protected and hot-swappable SCSI disk capacity with optional disk encryption
- **WAE-7341 Appliance:**
  - Quad-core processor, 8GB of RAM
  - Up to 310Mbps WAN connections and 12000 optimized TCP connections
  - Up to 900GB RAID-5 protected and hot-swappable SAS disk capacity with optional disk encryption
- **WAE-7371 Appliance:**
  - Dual Quad-core processors, 24GB of RAM
  - Up to 1Gbps WAN connections and 50000 optimized TCP connections
  - Up to 1.5TB RAID-5 protected and hot-swappable SAS disk capacity with optional disk encryption