



IBM / Cisco VS HP-VC



Kieran Daly (kidaly@cisco.com)

Technology Solutions Architect – Data Centre

Cisco Systems

Nov 2008 Rev A1

Agenda

- HP Virtual Connect Overview
- Cisco 3012 / 3110
- Cisco FC Switch Module
- Promotions
- Next Generation VMware Networking




HP-VC Overview



Virtual Connect for IBM BladeCenter

What is Virtual Connect?



What

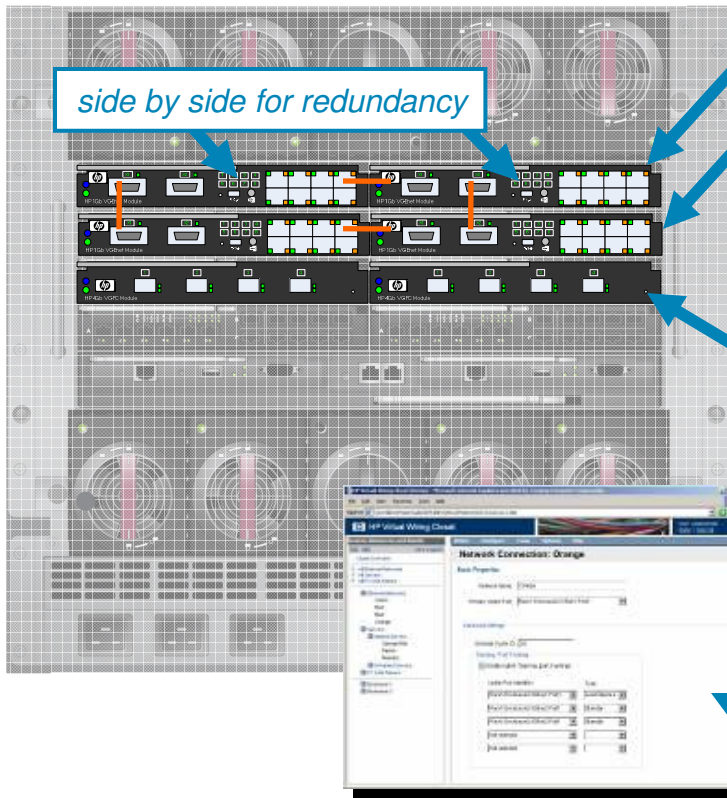
- Virtual Connect is a substitute for embedded Ethernet and FC switches and pass-through.
- Virtualizes the server-edge based on MAC/WWN identity spoofing.
- Effectively acts as a NIC/HBA aggregator.

Our Assessment

- Server folks seem to like the concept.
- However HP VC falls short - we don't believe HP VC meets customer needs

The Fibre Channel VC module requires the Ethernet VC module!

Virtual Connect components



Virtual Connect Ethernet Modules

- 16 downstream GE server ports
- 8 GE uplink ports (RJ-45) & 2 uplink CX4 10GE ports
OR
- 4 RJ45 / 2 SFP GE uplinks & 1 CX4 / 2 XFP 10GE uplinks
- Internal 1x10GE connection to redundant VC
- VC Manager - Built-in Web Manager



Virtual Connect Fibre Channel Modules

- 16 downstream 4G FC server ports.
- 4 uplinks acting as N_Ports with NPIV
- Requires the Ethernet VC
- Requires upstream switch to support NPIV



Virtual Connect Manager (*embedded*)


- Runs on the Ethernet Module
- Administer server MACs, WWNs, and boot parameters.
- Move/upgrade/change servers without coordination with the LAN or SAN admins.

VC Value Proposition

- Simplify Management/Operations/Networking
 - No switches to manage – no STP or Domain Id proliferation
 - Server Admin can provision and manage – similar look and feel as other Proliant mgmt tools
 - More efficient server bring up
 - Cable consolidation
- Server Admin Autonomy
 - Server adds, moves, changes with no impact to the SAN or LAN
- Maintain current operational boundaries
 - Network boundary moved outside the blade chassis
 - Blade chassis and associated I/O devices owned and managed by Server Admin



Cisco 3012 and 3110



Cisco Catalyst 3110 for IBM BladeCenter


Product Overview

CBS 3110X

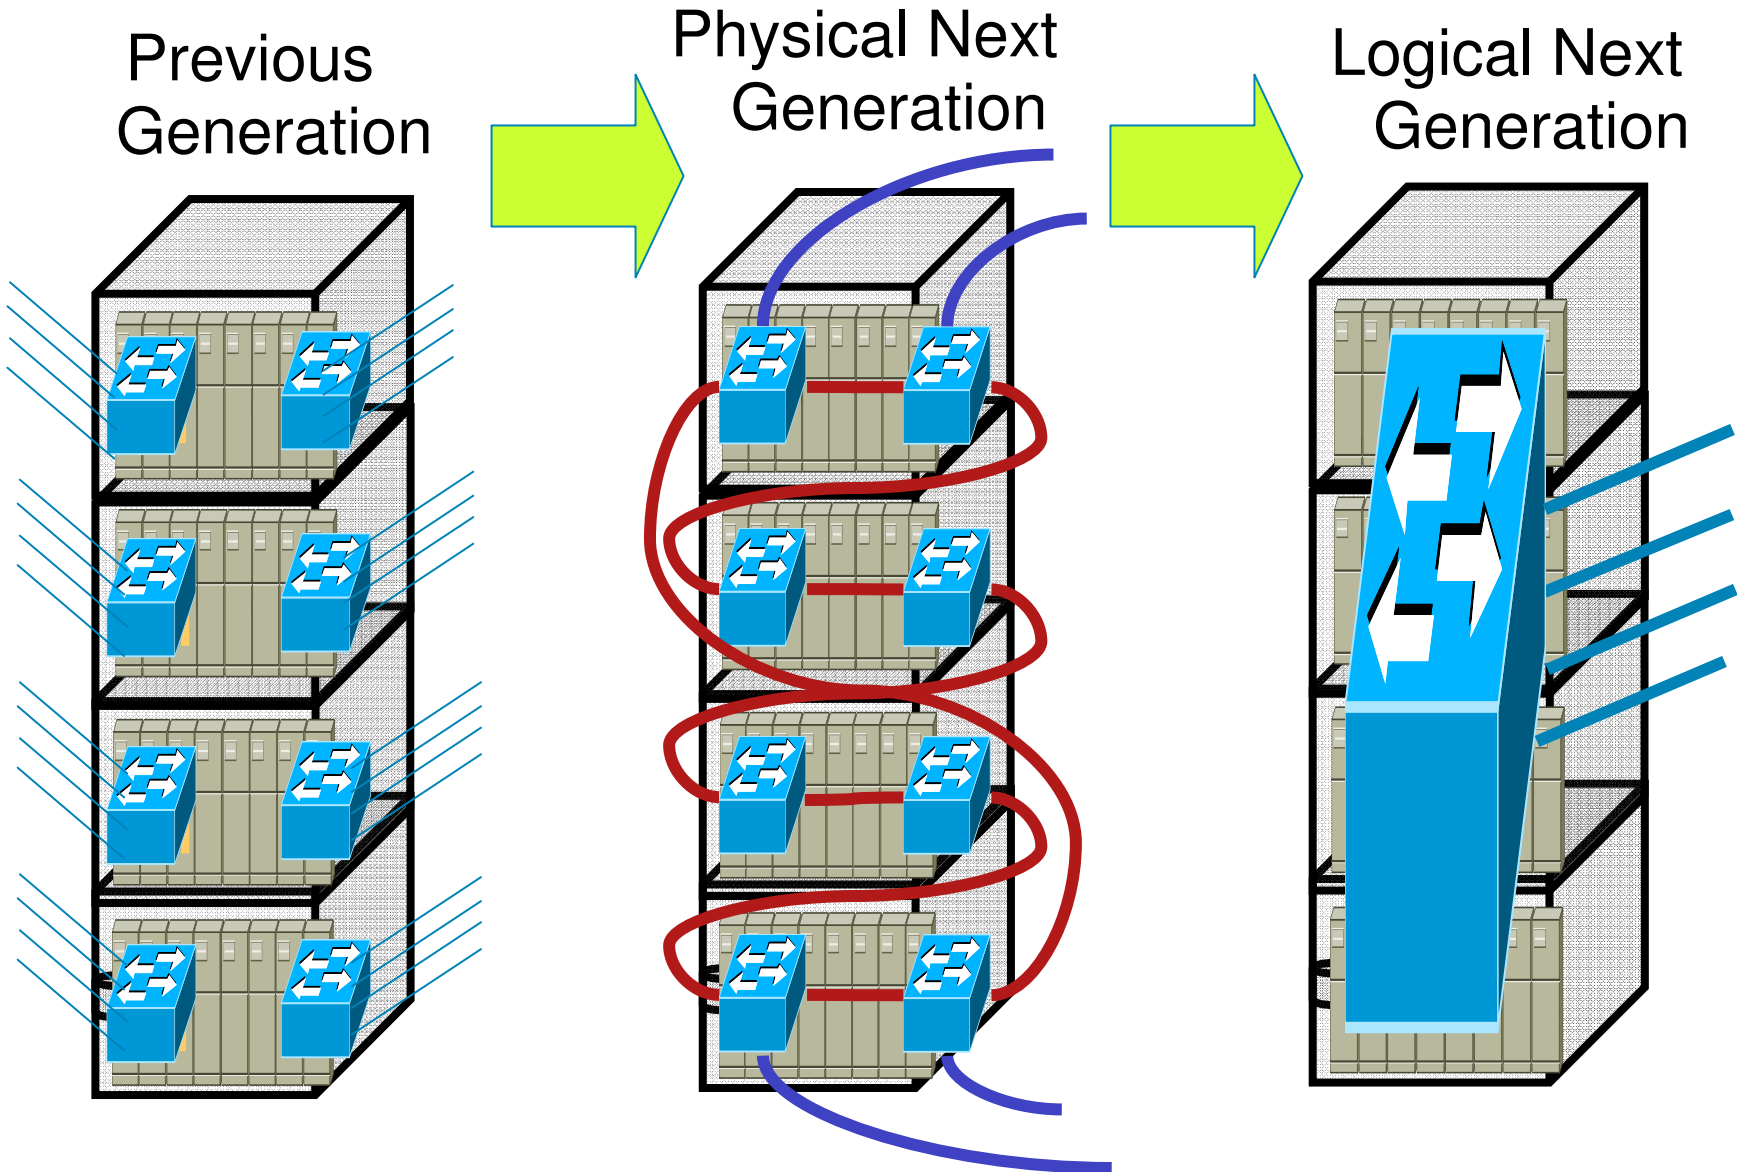


- Based on Cisco 3750E chip set
 - Real field-proven technology!
 - StackWise Plus = 64G stacking links (Cables: .5, 1 and 3 meter lengths)
- Rich feature set
 - Usable L3 support, Full PVLAN, Cross-stack EtherChannel, Flexlink, UDLD, etc.
- Three SKUs to be offered
 - CBS 3110G – 4 x 1G uplinks with stacking (RJ45)
 - CBS 3110X – 1 x 10G uplink with stacking (X2)
 - CBS 3012 – 4 x 1G uplinks without stacking (Entry level targeted at BC S)
- VBS (stacked model) supported in all BladeCenters (except S)
- Requires an Advanced Management Module

CBS 3110G



VBS = Real Evolution


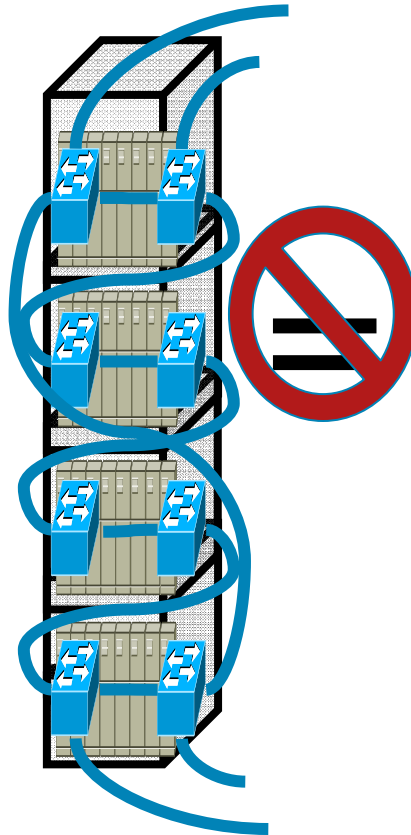


Don't confuse this with the competition

Competition Daisy Chained 10G

Cisco Virtual Blade Switching

- Physically looks similar – But...
- Unproven technology!
- Only 10G cross-connects!
- Many unknowns in the real world



Cisco offers true switch virtualization


High performance clustering of switches

Truly acts as a single switch – for both management and data flow purposes

Based on real world, proven technology!

No one else comes close to the flexibility and power offered by VBS

Multiple Deployment Options



Common Scenario

- Single Virtual Blade switch
- Cost Effective HA

Redundant ToR Scenario

- Separate VBS
- More resilient

Multi-NIC Scenario

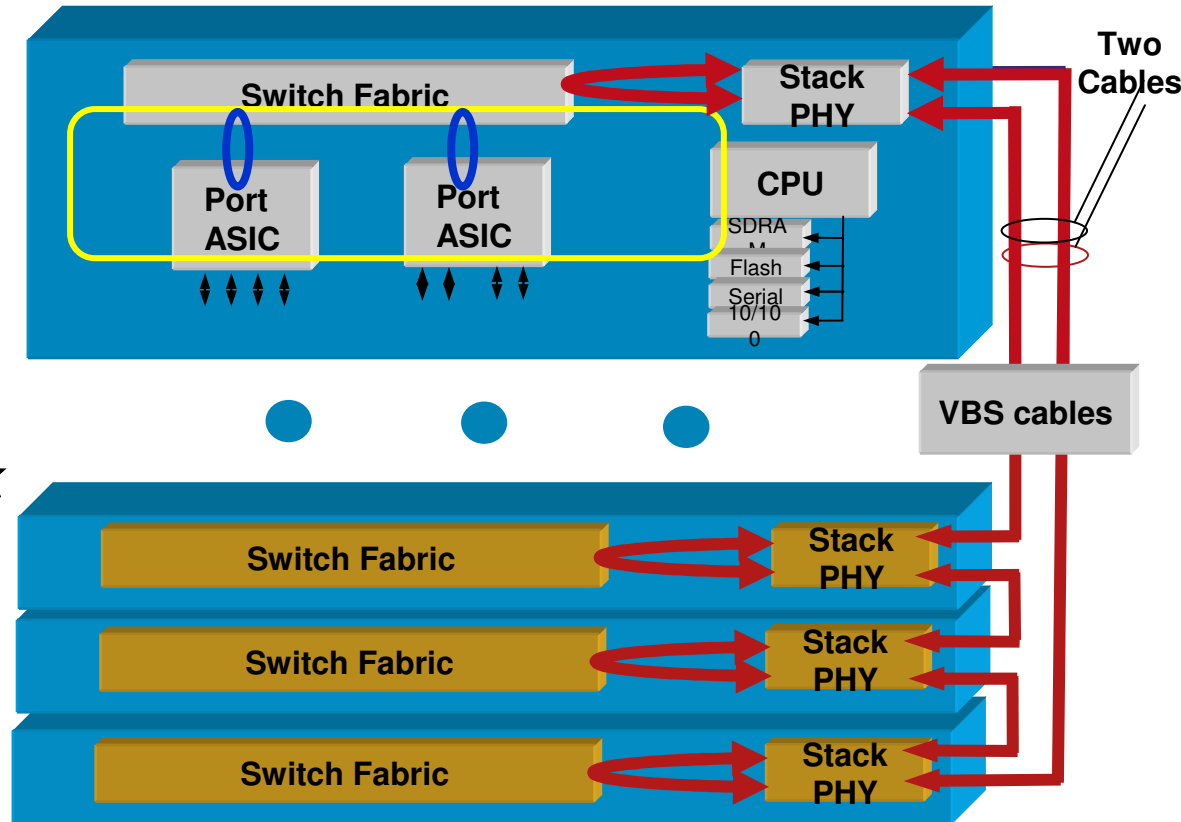
- Separate VBS for paired NICs e.g. HA VMware

Multi-NIC Scenario

- Same VBS for all NICs
- Possible 4 port EtherChannel = Max server bandwidth

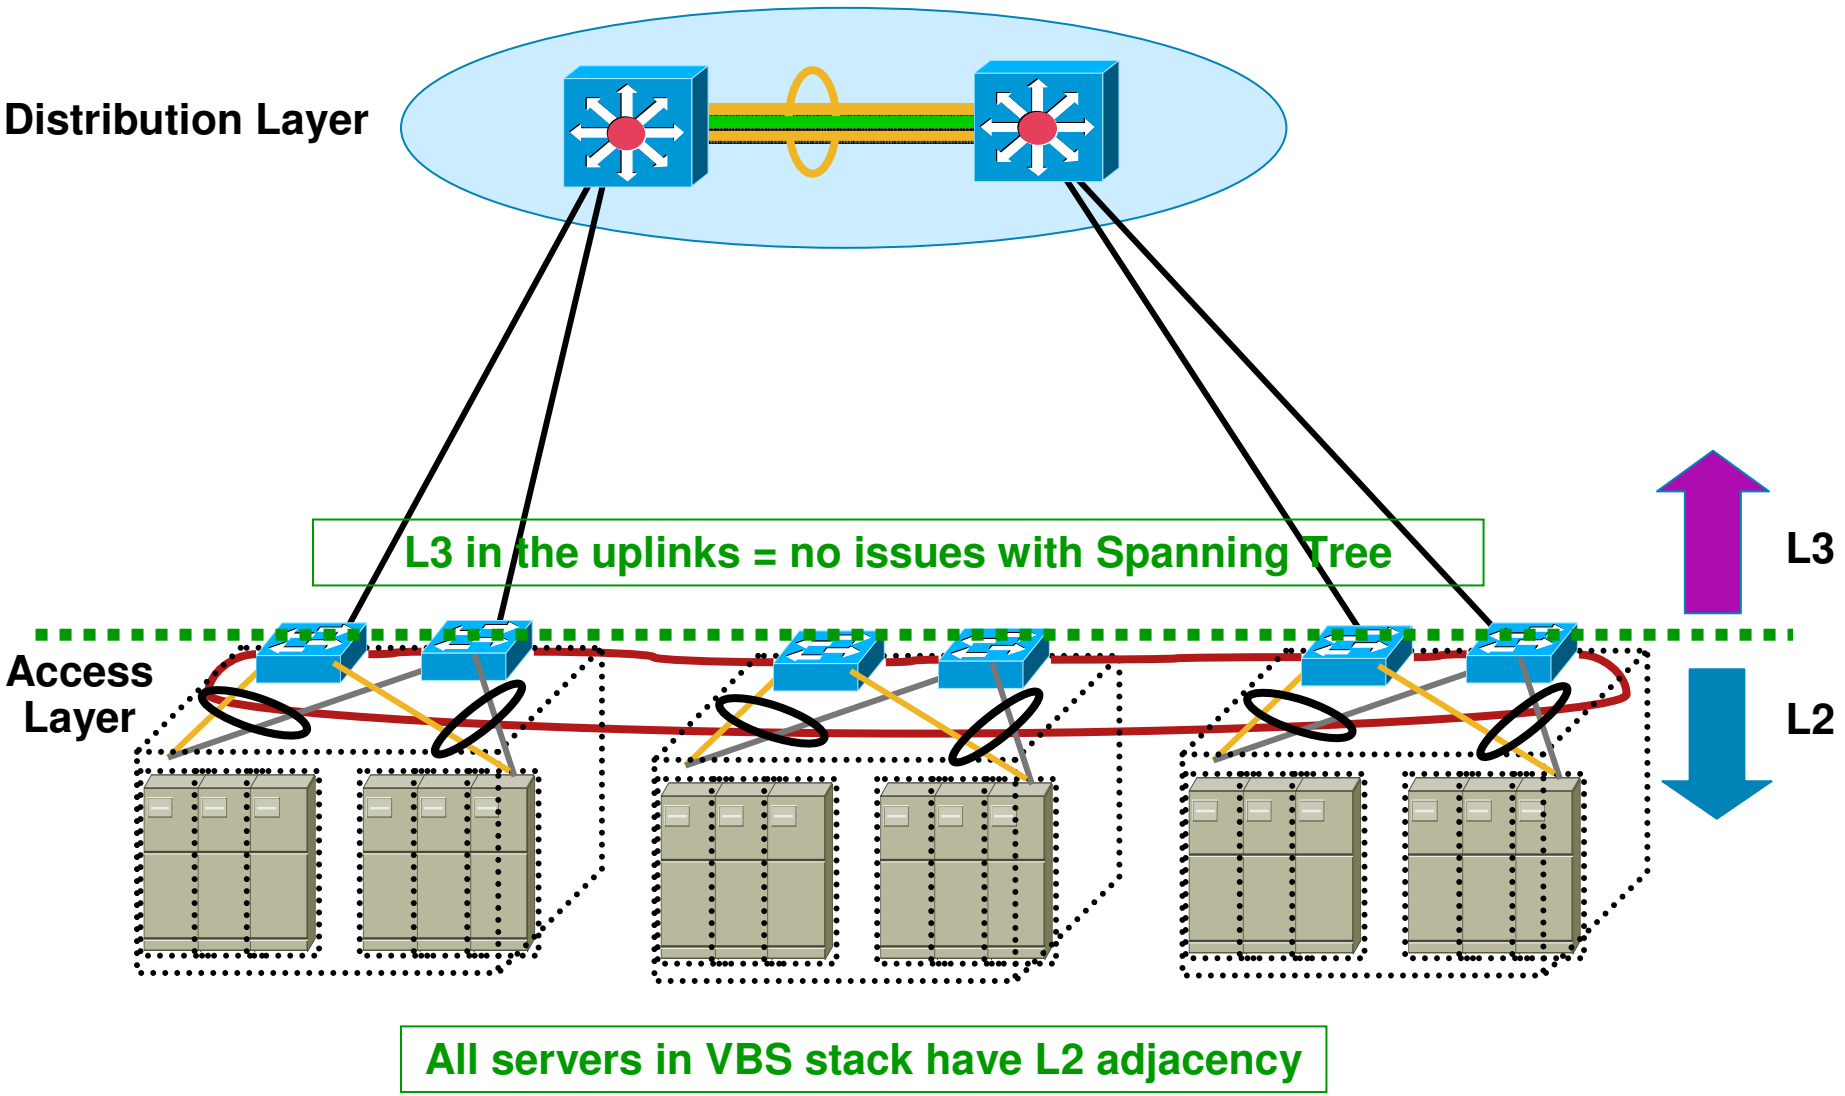
Next-Gen Stacking Switch Architecture

- **Stack of Switch acts as Single Switch**
 - Distributed L2/ MAC learning
 - Centralized L3 learning
- **Each switch consists of**
 - Switch Fabric
 - Port Asics (downlink & uplink ports)
 - Stackwise Plus (for stacking)
- **One Master Switch in stack**
 - 1:N Resiliency for Master
 - L2/L3 reconvergence is sub 200 msec
- **High Speed StackWise Plus (64 Gbps)**
 - Dual counter rotating rings
 - Still functions if not fully wrapped (half speed)



High Performance, Scalable and Manageable Architecture

Example: L2 to servers and L3 to uplinks



Distribution Layer

Access Layer


L3 in the uplinks = no issues with Spanning Tree

All servers in VBS stack have L2 adjacency

L3

L2

Example showing L2 on uplinks of CBS



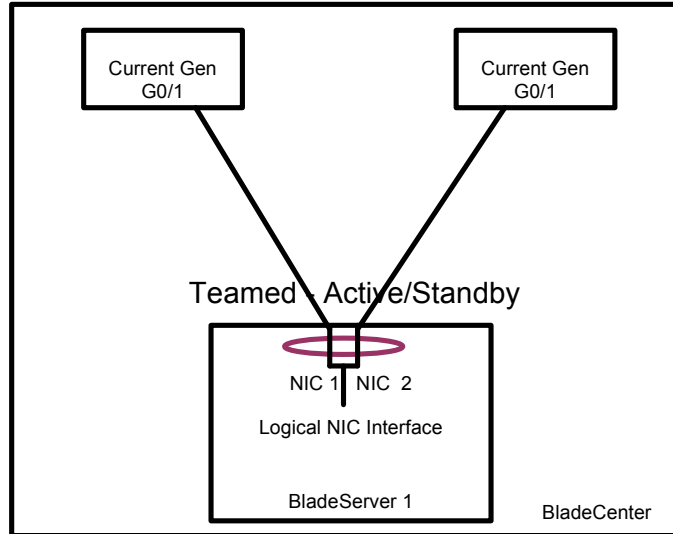
What is Flexlink?

- An alternative to Spanning Tree for redundant L2 connections
- Provides link level redundancy without Spanning Tree Protocols
- STP is automatically disabled on Flexlink ports
- Does not need to be configured on both sides
- Fast convergence time – can be less than 50ms
 - Speed of convergence not impacted by number of VLANs/STP instances running on link
- Configured in pairs (pair of physical ports or pair of EtherChannel ports)
- Configured using ***switchport backup interface xyz*** command on interface
- Use ***show interface switchport backup*** command to view

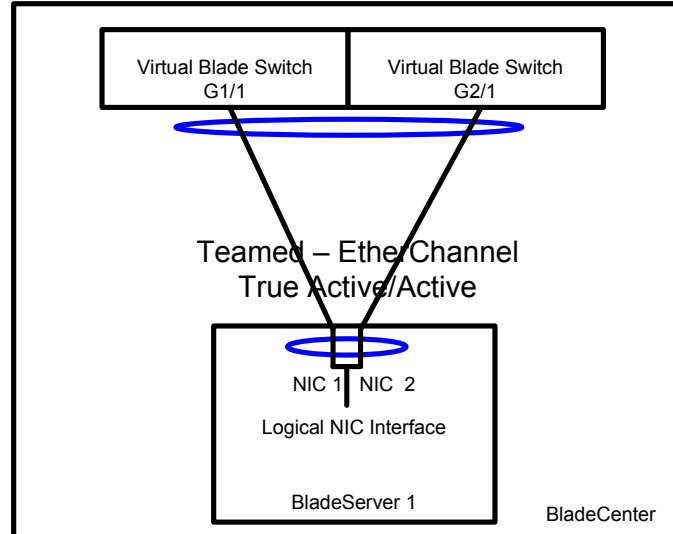
Standards based NIC load balancing

- Most of today's switches only work well with Active/Standby**

It can be done, but does not load balance well and is a nightmare to troubleshoot
- VBS will permit true NIC aggregation using standards based teaming**



Current gen – to different switches
Can not aggregate links



VBS – Stacked = same link switch
Can aggregate links

PVLAN – Cisco-only feature

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

Promiscuous port

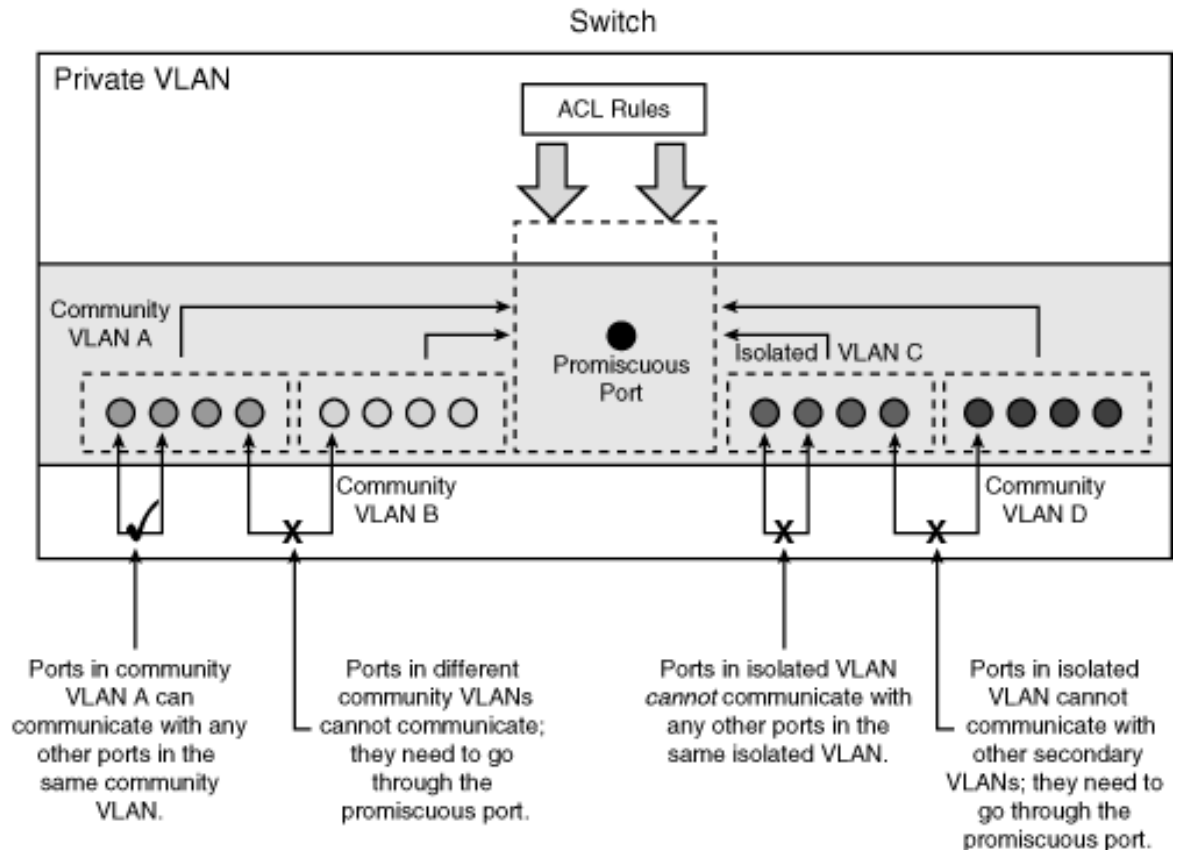
Can communicate with all interfaces, including the isolated and community ports within a PVLAN.

Isolated port

An isolated port has complete Layer 2 separation from the other ports within the same PVLAN, but not from the promiscuous ports.

Community port

Communicate among themselves and with their promiscuous ports.





For more info on PVLAN:

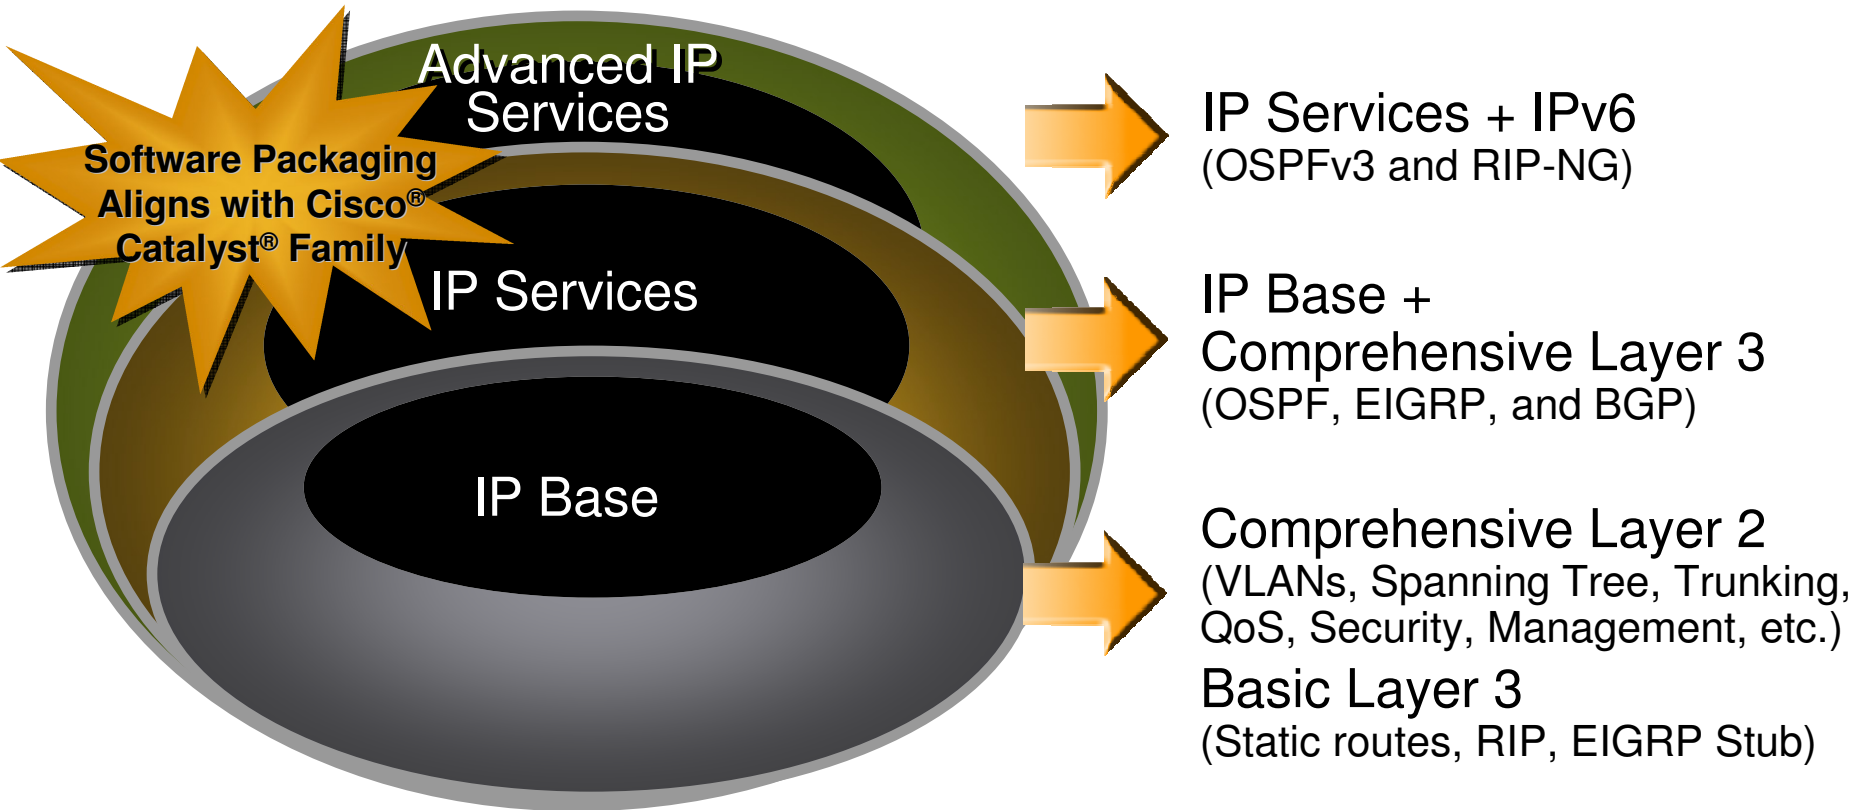
http://www.cisco.com/en/US/tech/tk389/tk814/tk840/tsd_technology_support_sub_protocol_home.html

MSIM – Multi-Switch Interface Module

- Goes in High Speed slots of BladeCenter H and HT
 - Each MSIM takes up 2 HS slots (7/8 or 9/10)
- Allows the use of standard switch modules in High Speed switch slots
 - FC, Ethernet, Pass-Thru
- Connects to servers via CFFh daughter cards on servers
- 3012 and IGESM supported in MSIM at GA. Support for CBS3110 and Cisco FC module in MSIM added mid-July 08 (aMM update required)



Cisco Catalyst Blade Switch 3000 Software Packaging and Offerings



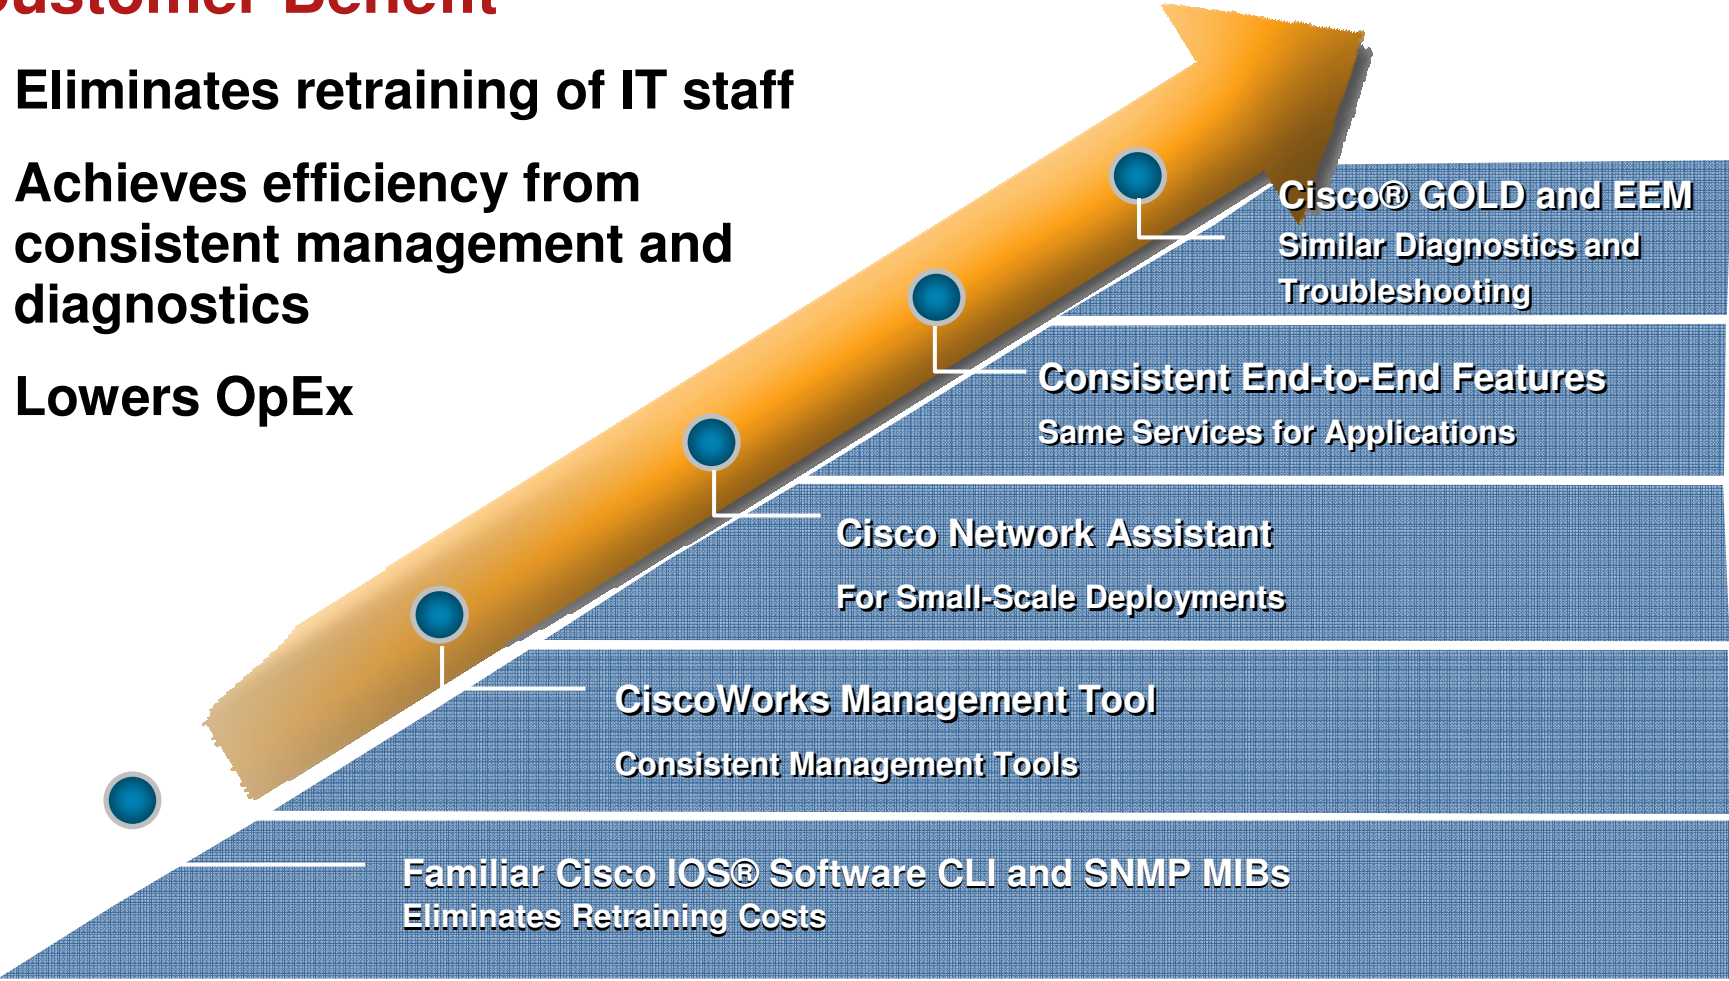
**Software: IP Base image included in each part number
 IP Services and Advanced IP Services images available only for VBS switches
 (extra charge)**

“Operational Consistency” for Customers

Achieving Lower TCO

Customer Benefit

- Eliminates retraining of IT staff
- Achieves efficiency from consistent management and diagnostics
- Lowers OpEx



Device Manager View of a Stacked Switch

- Accessed by pointing your browser at the IP address of the stack
If non-stacked, point browser at IP address of individual switch

The screenshot shows the Catalyst Blade Switch 3110 Device Manager interface. Annotations include:

- Switch selector:** Points to the 'Stack: Master 6' dropdown menu.
- UUID of BladeCenter:** Points to the 'Chassis ID: 67977D0412E7312B8A1C242BCA0B6214' field.
- Slot in BC:** Points to the 'Slot: 1' field.

The interface displays the following information:

Catalyst Blade Switch 3110 Device Manager - Switch

Refresh Print Software Upgrade Legend Help

Uptime: 2 weeks, 3 days, 3 hours, 57 minutes

View: Status Stack: Master 6

Chassis ID: 67977D0412E7312B8A1C242BCA0B6214 Slot: 1

CONSOLE MODE MBR MST

INTERNAL EXTERNAL

STK 1 STK 2

WS-CBS3110G-S

Move the pointer over the ports for more information.

Contents

- Dashboard
- Configure
- Monitor
- Maintenance
- Network Assistant

Dashboard

Switch Information

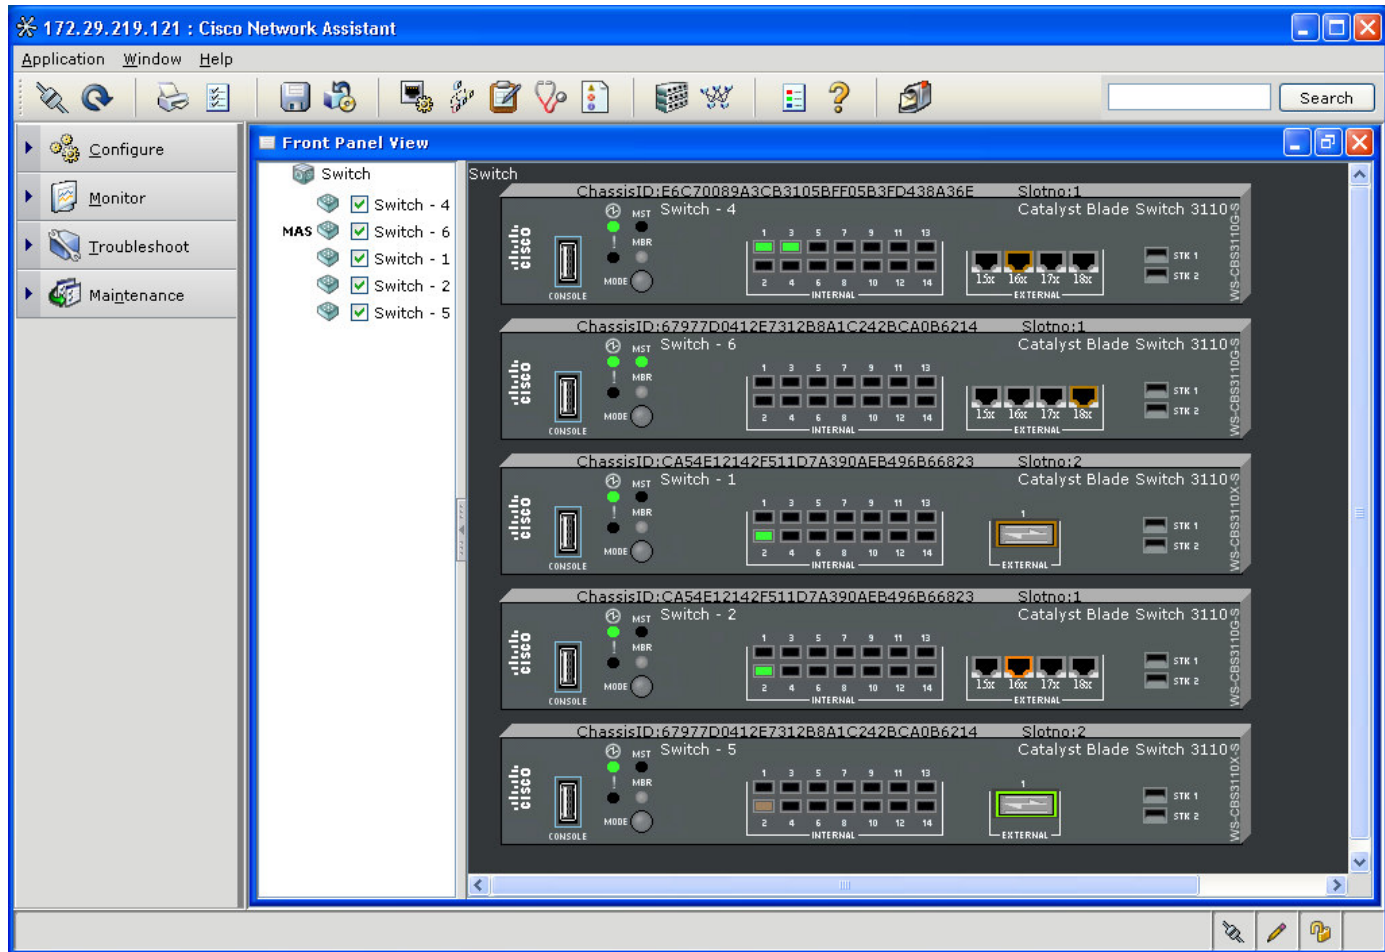
Host Name:	Switch
Product ID:	WS-CBS3110G-S
IP Address:	172.29.219.121
MAC Address:	00:1B:54:DC:58:00
Version ID:	V01
Serial Number:	FOC1122S00E
Software:	12.2(0.0.96)EX
Contact:	
Location:	

Switch Health

Stack Bandwidth Used	Stack Packet I
0%	0%

CNA View of a Stack

- Configure CNA to point at the IP address of the stack
 Good free GUI to configure various basic and advanced options



GOLD (Generic Online Diags)

Online Switch Diagnostics for

- *Hardware Components*
- *Switch Interfaces*

When Do you Run

Health-Monitoring (Run-time)

```
Switch(config) #[no] diagnostic monitor interval { switch <1-9> }
test { test-id | test-id-range | all } hh:mm:ss { ms <0-999> } { days <0-20> }
```

On-Demand

```
Switch#diagnostic start {switch <1:9>} test {test-num | test range
| all | basic | non-disruptive }
```

Scheduled

```
Switch(config) #[no] diagnostic schedule { switch <1-9> } test
{ test-id | test-id-range | all } daily {hh:mm}
```

How to Use:

- *Invoke from Stack Master*
- *Schedules on stack members*
- *Resilient to Stack Master failure*

Offers Basic Diagnostic Tests Two Types of Tests

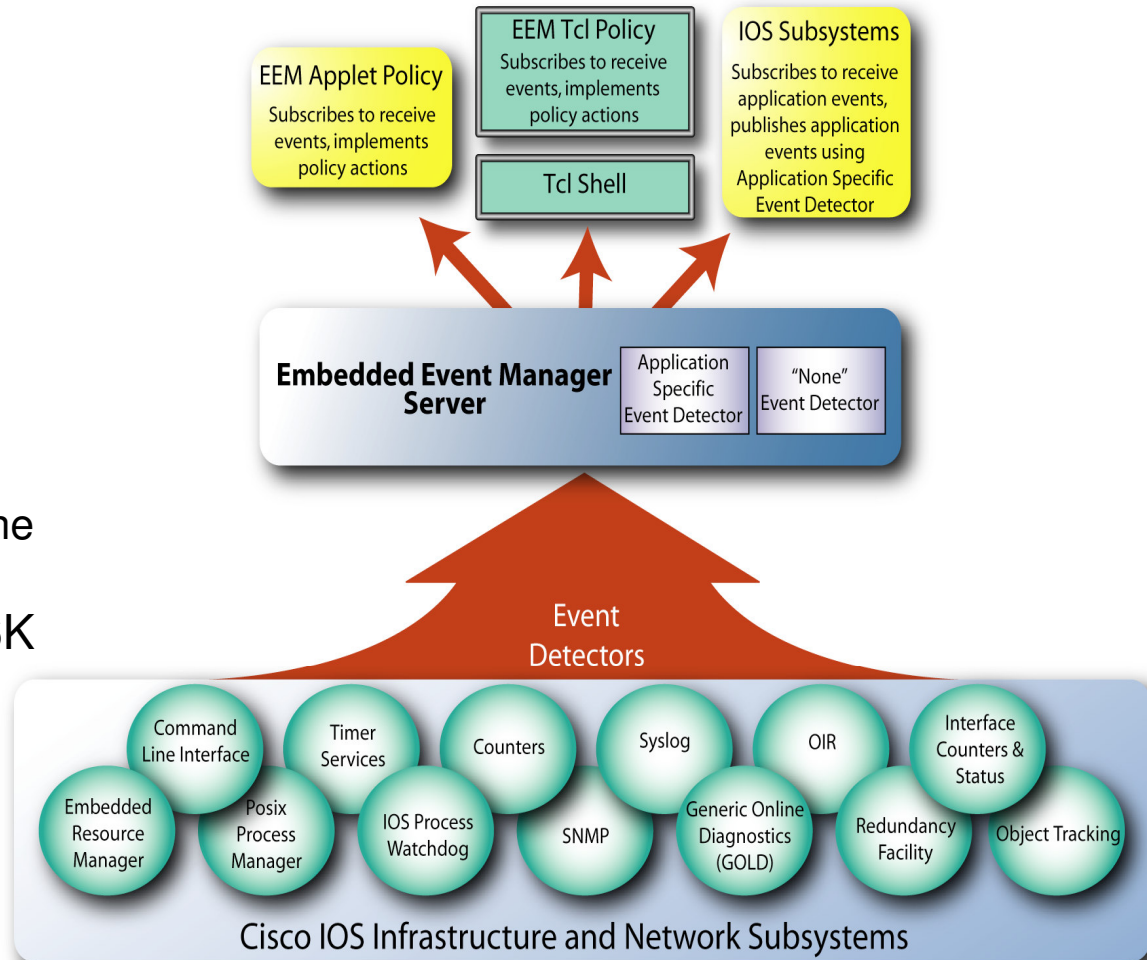
- *Disruptive Tests*
- *Non-disruptive Tests*

*Some differences with Cat 6K
(e.g. No Boot-up test, No
Stop Command)*

Embedded Event Manager

*Very Flexible & Powerful
Capability to Automate
Troubleshooting & Mgmt*

- Program automatic actions based on events
 1. Set policies (scripts) using
 - Applets (IOS CLI)
 - Tcl
 2. Event Detectors watch for events
 3. EEM Server is notified – carries out the actions for the policies
- Feature Consistent with Cat 6K
- Potential application in Blade Environment
 - Notify network admin on insertion of new blade





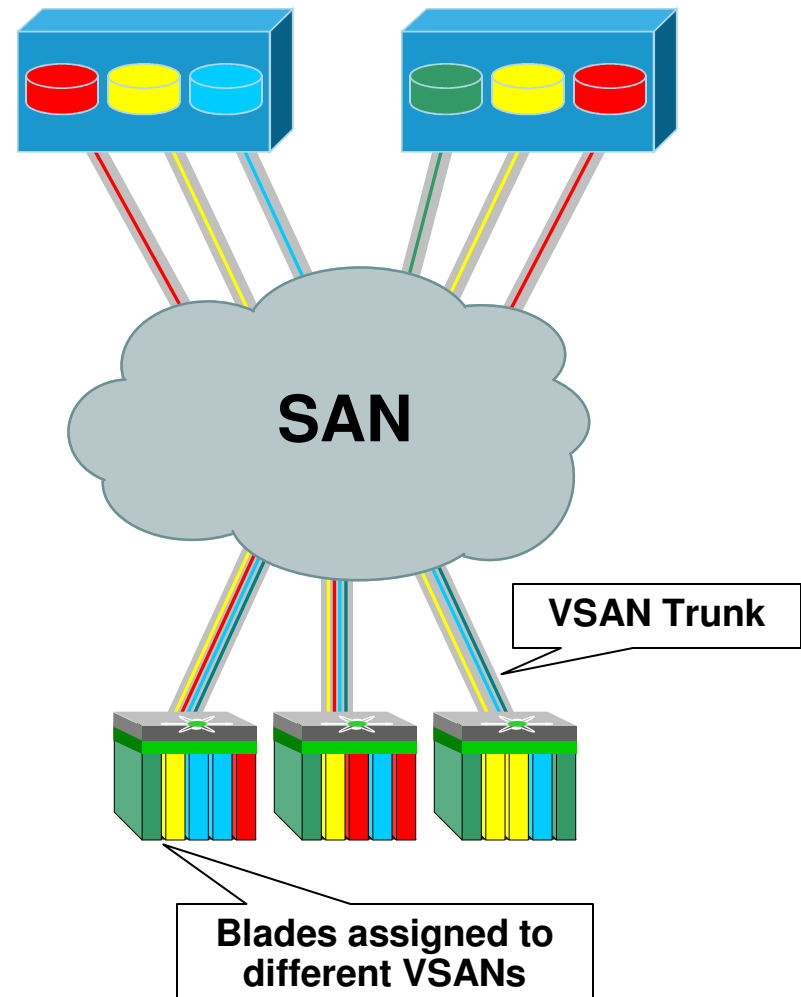
Cisco FC Switch Module



Integrated Virtual SAN Support

- Separates control traffic (Different than zoning)
- Industry leading standards based fabric virtualization technology
- Separate fabric services per VSAN (FSPF, zoning, name services, etc...)
- Up to 16 VSANs per Blade Switch (256 per network)
- Hardware tagging with no loss of performance
- Transparent to the end nodes
- Enables granular network management with RBAC
- VSAN Trunking ensures efficient ISL utilization


VSANs Extended to Blades



Cisco FCSM – Traffic Management

■ Port-Channeling

- Up to 16 physical links
- Up to 24Gbps of bandwidth (4Gb x 6 ports)
- Exchange Based Load-Balancing

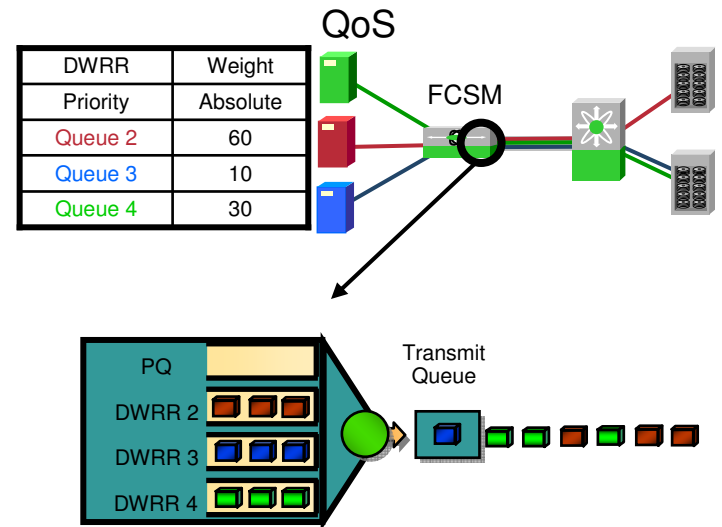


■ Virtual Output Queue

- Non-blocking architecture

■ Quality of Service (QoS)


- 4 different queues
- High priority for critical applications
- Can be on a per zoned based




PortChannel Overview

PortChannels have the following functionality:

- Increases the aggregate bandwidth
- Load balances across multiple links and maintains optimum bandwidth utilization
 - Provides fault tolerance on an ISL
- A PortChannel can include up to 16 interfaces



Single PortChannel Between Two MDS Switches

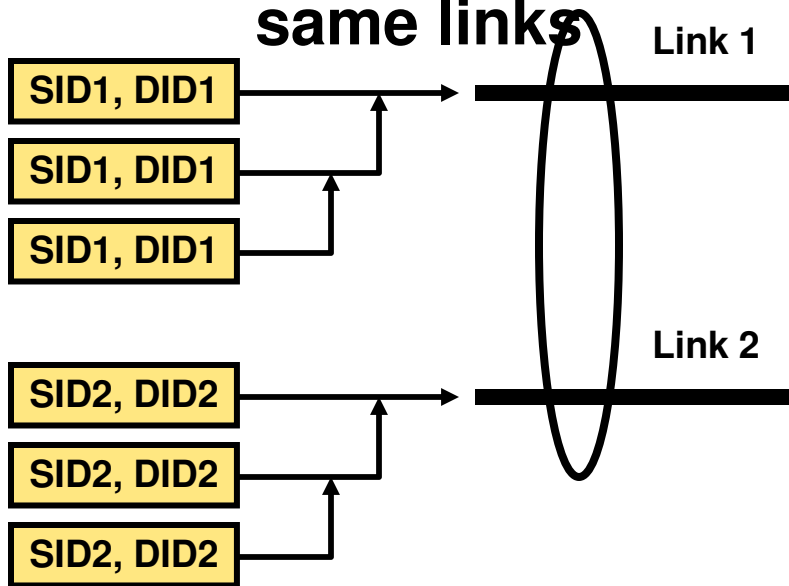


Multiple PortChannels Between Two MDS Switches

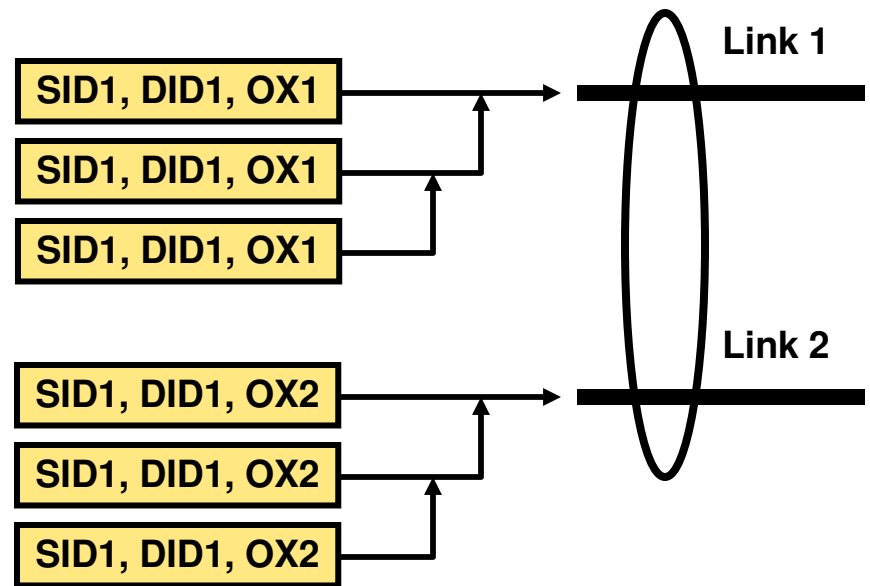
Load-Balancing in PortChannels

Two load balancing mechanisms

Flow based: Frames between source and destination follow the same links




Exchange based: Frames in the same exchange follow the same link



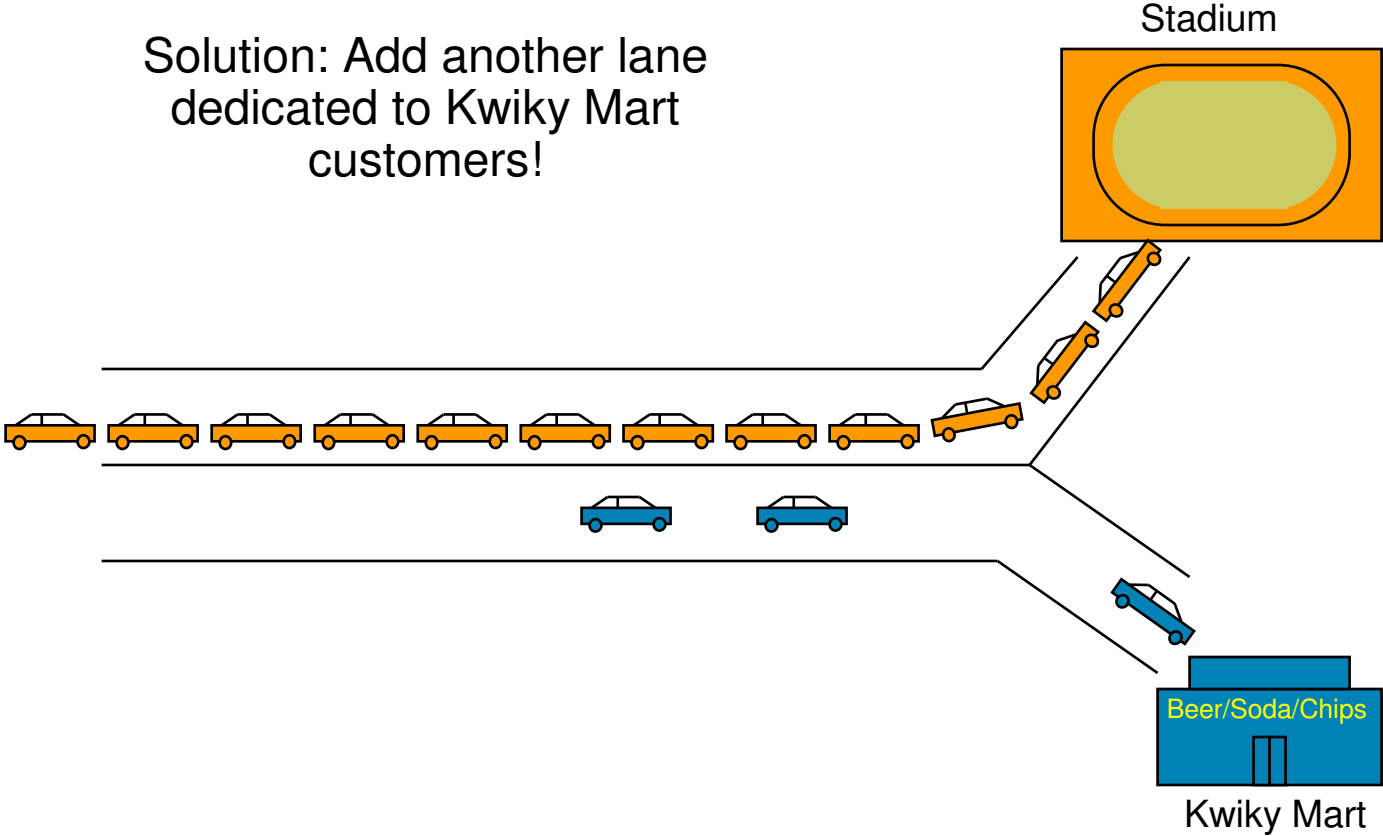
Traffic Engineering - VOQs

HOL blocking



Virtual Output Queues (cont.)

Solution: Add another lane dedicated to Kwiky Mart customers!




Quality of Service to the Blade Server

- End-to-end QoS applied on ingress **and** egress (DWRR)
- **4 queues** (3 user definable, 1 absolute priority)
- **Hardware tagging** with no loss of performance
- **No software required** on the end nodes
- **Mappable to IP** QoS for FCIP or iSCSI
- Simple configuration applied by **Zone** or **VSAN**

User Definable QoS per Blade

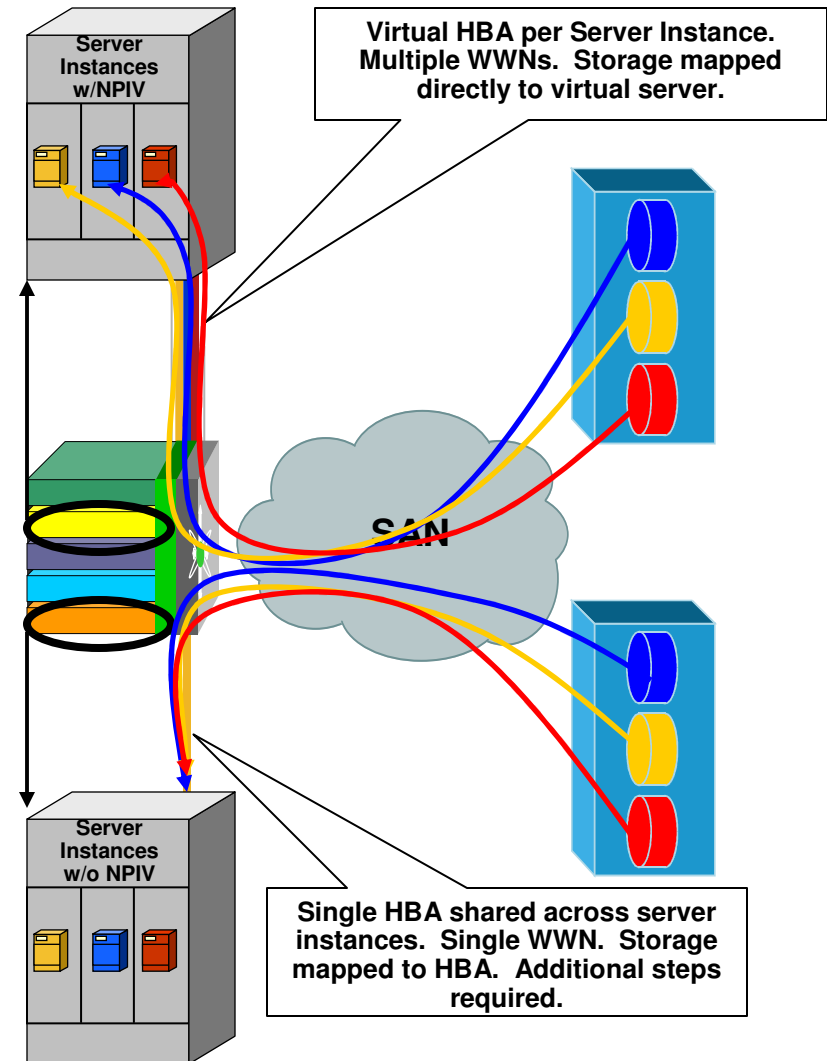
DWRR	Weight
Priority	Absolute
Queue 2	60
Queue 3	10
Queue 4	30



N_Port Interface Virtualization (NPIV)

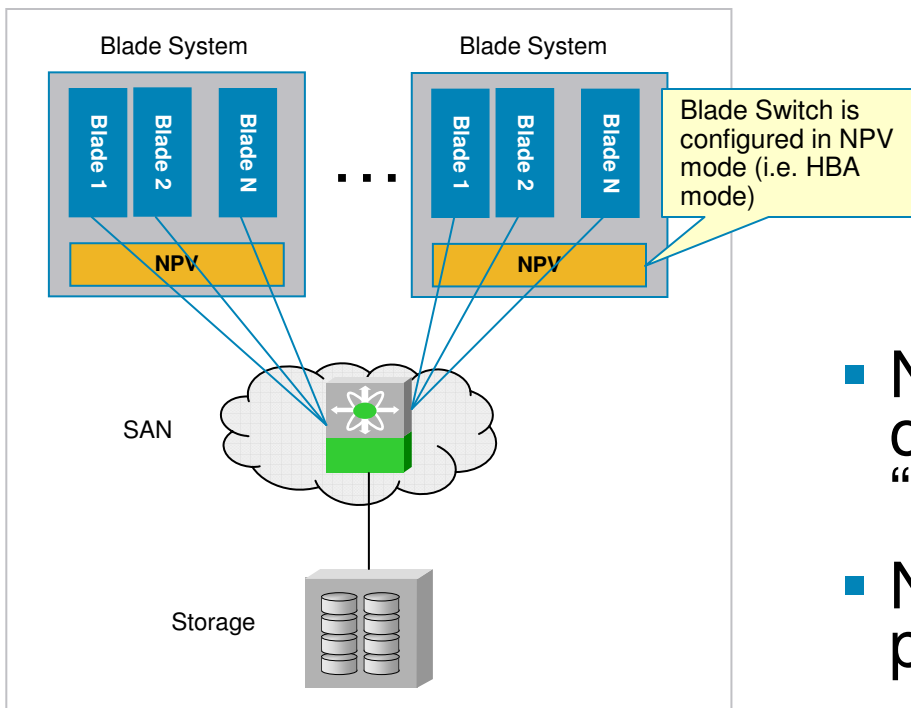
Virtual Servers with and without NPIV

- Creates Virtual N_Port **per Server Instance**
- Based on **T11 Standard**
- **Simplifies** Virtual Server Management
- Allows use of standard zoning, LUN management and security procedures
- **Enables QoS** assignment per virtual server
- **Supported** by Current Generation of HBAs



N-Port Virtualizer (NPV)

Blade Switch Deployment Model – NPV Mode



- NPV simplifies deployment and management of large scale Blade Server environments

Reduces number of Domain IDs

Minimizes interoperability issues with core SAN switches

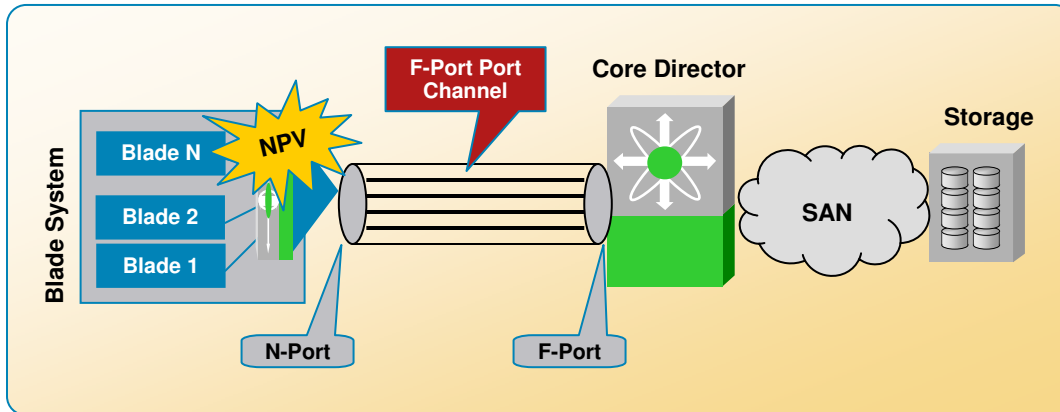
Minimizes coordination between Server and SAN administrators

- NPV converts a Blade Switch operating as “FC Switch” to a “FC HBA”
- NPV is available on the following platforms
 - IBM Blade Switches
 - MDS 9124 Fabric Switches

Enhanced Blade Switch Resiliency

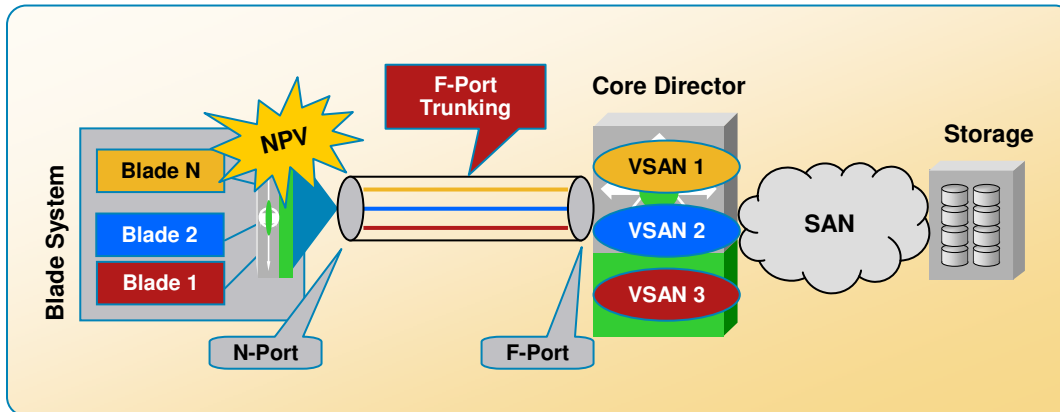


F-Port Port Channel



- F-Port PortChannels
 - Bundle multiple ports in to 1 logical link
 - Any port, any module
- High-Availability (HA)
 - Blade Servers are transparent if a cable, port, or line cards fails
- Traffic Management
 - Higher aggregate bandwidth
 - Hardware-based load balancing

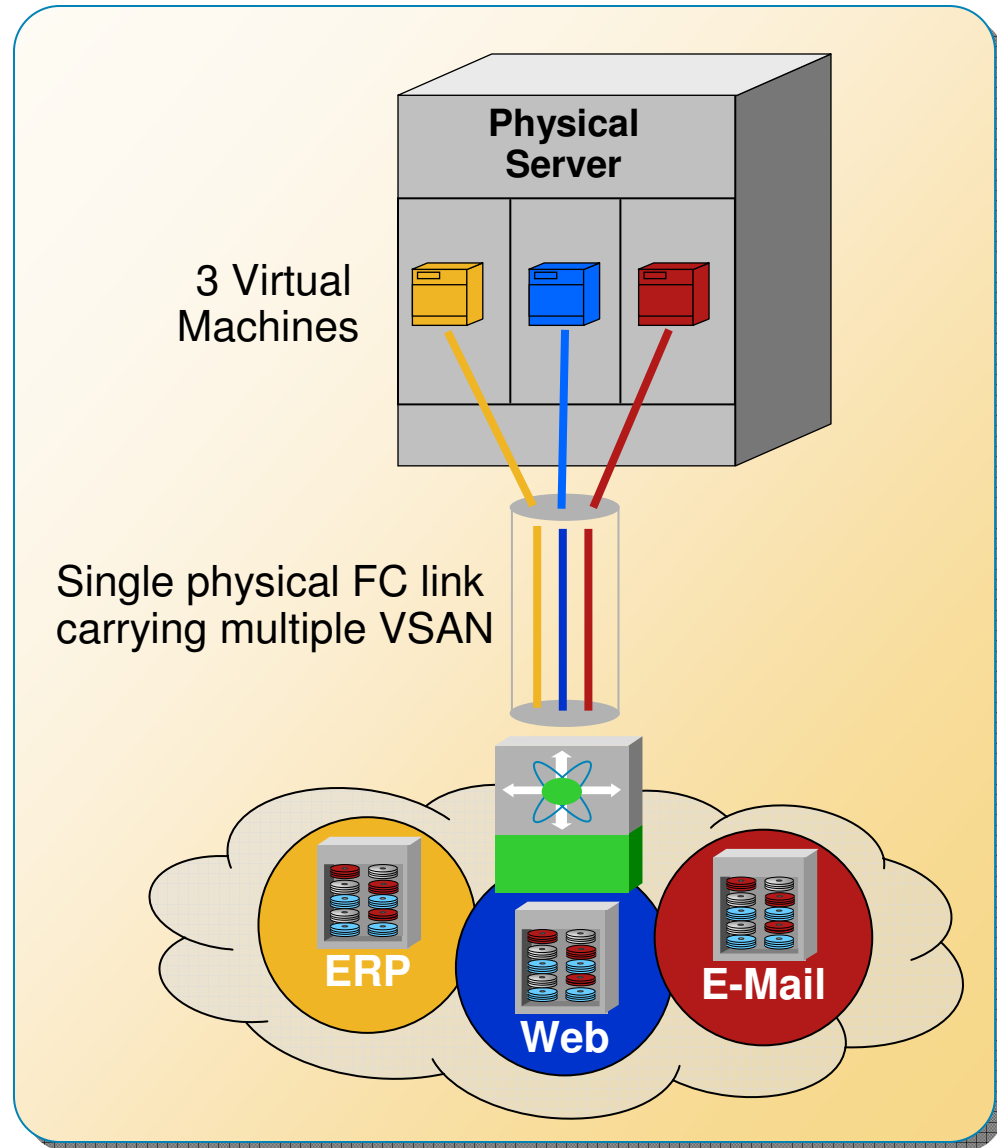
F-Port Trunking



- F-Port Trunking
 - Partition F-Port to carry traffic for multiple VSANs
- Extend VSAN benefits to Blade Servers
 - Separate management domains
 - Separate fault isolation domains
 - Differentiated services: QoS, Security


Fully Extending Fabric Virtualization to VMs

- NPIV allows each virtual machine (VM) to be associated to a unique virtual HBA
 - VMs register independently via unique PWWN and obtain unique FCID
 - Standard-based (ANSI T11)
- Separate fabric login by each VM enables VM level:
 - Zoning
 - Security
 - Traffic mgmt
- Combined with F-Port Trunking, each VM can now belong to a different VSAN




Performance Monitoring of VMs

- Cisco Fabric Manager provides a full set of tools for fabric configuration and performance monitoring.
- The same performance monitoring capabilities available for the physical devices are available for the individual NPIV-enabled virtual machines
- Single monitoring point across the entire end-to-end storage infrastructure






Next Gen VMware Networking



Cisco Nexus 1000V

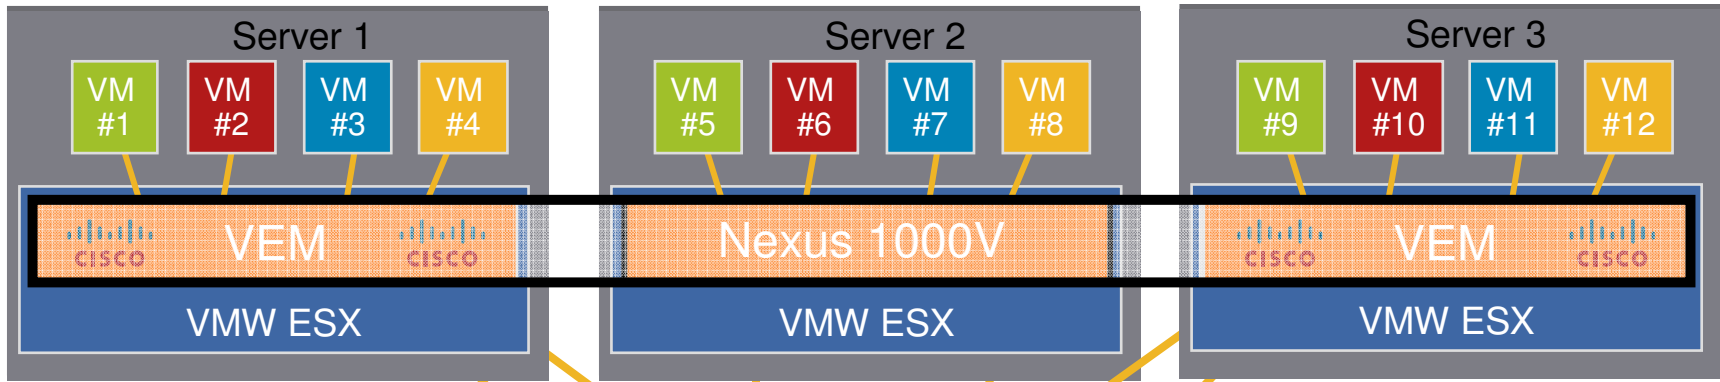
Industry First Third-Party Virtual Distributed Switch

- **Nexus 1000V provides enhanced VM switching for VMW ESX environments**
- **Features VN-Link capabilities:**
 - Policy-based VM connectivity
 - Mobility of network and security properties
 - Non-disruptive operational model
- **Ensures visibility and continued connectivity during VMotion**



Enabling Acceleration of Server Virtualization Benefits


Cisco Nexus 1000V Architecture



Virtual Supervisor Module (VSM)
 Virtual Ethernet Module (VEM)

Cisco Nexus 1000V Enables:


- Policy Based VM Connectivity
- Mobility of Network & Security Properties
- Non-Disruptive Operational Model



Nexus 1000V Deployment Scenarios

Pick Your Flavor

1. Works with all types of servers (rack optimized, blade servers, etc.)
2. Works with any type of upstream switch (Blade, Top or Rack, Modular)
3. Works at any speed (1G or 10G)
4. Nexus 1000V VSM can be deployed as a VM or a physical appliance



Key Features of the Nexus 1000V

Switching

- L2 Switching, 802.1Q Tagging, VLAN Segmentation, Rate Limiting (TX)
 - IGMP Snooping, QoS Marking/Queuing

Security

- Policy Mobility, PVLAN, ACL (L2–4 w/ Redirect), Port Security
- Cisco TrustSec—Authentication, Admission, Access Control

Provisioning

- Automated vSwitch Config, Port Profiles, Virtual Center Integration
 - Optimized NIC Teaming

Visibility

- Historical VMotion Tracking, ERSPAN, NetFlow v.9 w/ NDE, CDP v.2
 - VM-Level Interface Statistics, Wireshark

Management

- Virtual Center VM Provisioning, Cisco Network Provisioning
 - Cisco CLI, XML API, SNMP (v.1, 2, 3)

Virtual Connect – Ethernet Module Features & Trade-off's




Features

- Server profile mobility
- Profile Pre-provisioning
- Single point of management
- Supports MAC address virtualization to enable server identity mobility in the event of a failure
- Supports LACP, .1Q, Shared Uplinks
- Smart link failure identification enables functions such as NIC teaming
- Provides redundancy via multiple modules
- Loop-free via act/stdby on uplinks

Trade-off's

- Redundancy is host based – manual move or blade replacement for repair
- Intra-enclosure stacking only today
- Maximum of 8 Ethernet VC modules
- Could impact VMware Infrastructure 3.0 architecture best practice of supporting 3 networks
- Duplicate MAC's possible
- Smart link disables srvr 2 srvr traffic in an uplink failure
- No IGMP snooping, QOS controls, port security, TACACS
- No ability to tune loop parameters

Virtual Connect – FC Module Features & Trade-off's



Features

- Domain reduction
- Ease of interop
- Server profile mobility
- Profile Pre-provisioning
- Single point of management for all VC modules
- Virtual WWN provides flexibility for server identity mobility

Trade-off's

- Disruptive Software Upgrades
- Diagnostics: SPAN FCPing/traceroute
- No stacking option for FC modules
- FC Switching through uplinks
- Need to consider max number of logins for NPIV per port
- Intra-enclosure switching only
- Can have duplicate WWN's

VBS = True switch virtualization

- **Provides up to 9 to 1 reduction in switches to manage**

 - Can stack up to 9 switches together – looks and acts like a single switch

 - Single IP for stack

 - Single login to manage all switches in the stack

- **More cable flexibility over any of the current gen solutions**

 - Cross-stack EtherChannel

 - Dove-tails into Cisco VSS strategy on 6500

 - Can mix 10G and 4x1G in same stack

 - Can use unused 1G uplinks as optional connections for standalone servers or other devices

- **True plug and play on failure replacement**

 - Replacement automatically gets config from other stack members

 - Replacement automatically gets desired IOS from other stack members

Issues with FC Virtual Connect

HP Fibre Channel VC

Cisco FCSM

- 4:1 HBA-to-Uplink oversubscription
- HBAs do not use all uplinks
- Limited load balancing capability

Bandwidth Utilization

- 2.3:1 HBA-to-Uplink oversubscription
- HBAs use all uplinks
- Optimal load balancing

- If an uplink goes down, HBAs must be manually re-mapped to other uplinks
- HBAs must logout and log back into fabric
- No Port Channeling

High Availability

- Port Channeling protects against uplink failures
- If an uplink goes down, HBAs are dynamically re-mapped to other uplinks
- HBAs do not logout of fabric

- The server admin can effectively bypass zoning by WWN spoofing
- VC puts SAN security in hands of server admin

Security

- Server admin cannot compromise SAN security

- No traffic management for congestion and Head of Line blocking

Traffic Mgmt

- QoS and FCC to manage congestion
- Virtual Output Queuing to mitigate Head of Line blocking

- WWNs can overlap between separate VC domains
- Does not support RADIUS, TACACS+, SSH or SNMPv3

Mgmt

- No WWN overlap
- Supports RADIUS, TACACS+, SSH and SNMPv3

- No VSAN support
- All servers in same VSAN

VSAN

- Full VSAN support

Feature Comparisons Summary

	HP & HPVC	Cisco VBS IBM BOFM	Competitive Advantage
Hardware Requirements	HP VC Ethernet and FC Modules	Cisco VBS3110	HP forces proprietary hardware. IBM BOFM with VBS3110 leverages current environment.
MAC & WWN persistence	Yes, only with VC hardware. Duplicate id possible	Yes	HPVC allows duplicates. Requires VCEM for fabric wide mgmt at an 9k incremental per chassis.
FC Support	NPIV only HP forces a 2 nd or 3 rd element for SAN teams to manage	Switch and NPV modes to match all topologies	NPIV and VMware introduce the potential for many FCID's of a given F-Port. NPIV only implementations limit scale.
Full VM Feature support	No. NPIV and VLAN tagging Fix Claimed	Yes. Including QOS, Security	HP installations do not support VMware best practices. Special config's req'd with exposure
Quality of Service & Security	Huh?	Yes	Ensure that critical applications get more priority in the fabric.


Cisco Catalyst Virtual Blade Switch

Access Layer (Virtual Blade Switch)

Distribution Layer

Mix-n-match
GE & 10GE
switches

Local Traffic
doesn't go to
distribution
switch



Higher
Resiliency
with
Etherchannel

With VSS on Cat
6K, all links
utilized

Single Switch / Node
(for Spanning Tree or
Layer 3 or
Management)

Greater Server BW –
via Active-Active
Server Connectivity

Only Cisco provides
maximum flexibility
and scalability!

