



Cisco Expo
2008

Data Centre Storage Solutions and Virtualization



Bjørn R. Martinussen

European Data Centre Consulting

Data Centers Are Under Increasing Pressure



Collaboration



Empowered User



SLA Metrics



Global Availability



Reg. Compliance

New Business Pressures



Operational Limitations

Power & Cooling



Asset Utilization



Provisioning



Security Threats



Bus. Continuance



Current Problems Demand a New Approach



- Expectation of 24x7 application availability
- 54% of network downtime is caused by human error (Uptime Institute, 2007)



- Continued deconstruction of the server increases demands on the network
- Market transitions between transport technologies and application architectures



- By 2009, 50% of large businesses will spend more on power and cooling than on new servers (Gartner, 2006)
- Quad- & octal-cores will drive significantly more traffic
- Storage is expected to continue to grow at a 40-70% CAGR (Gartner, 2006)

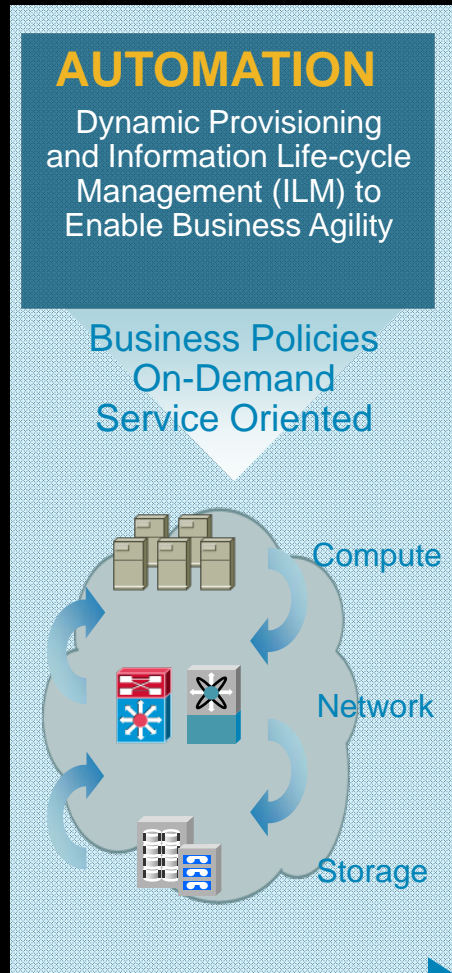
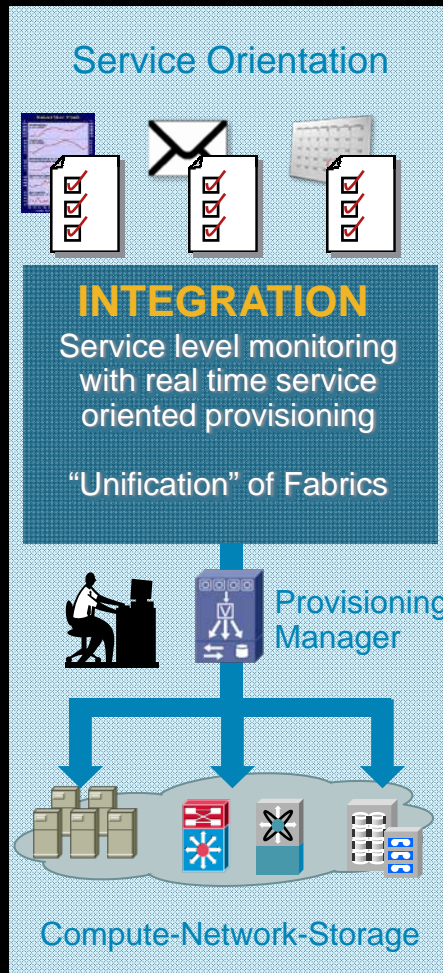
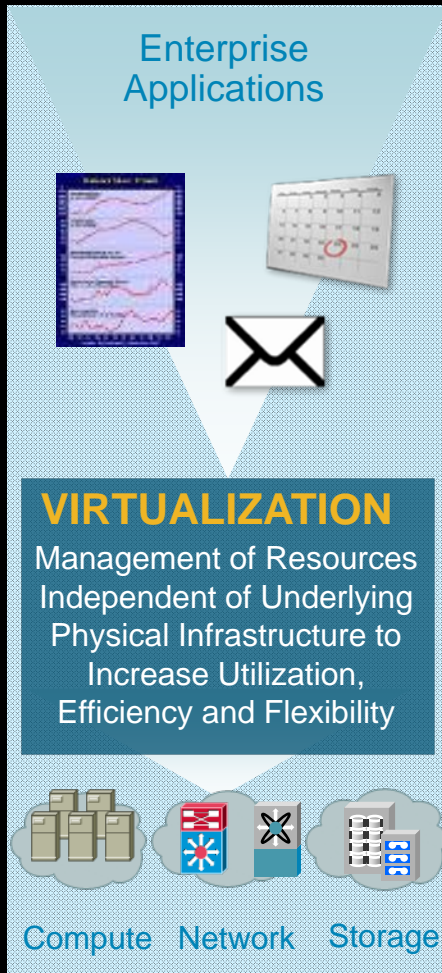
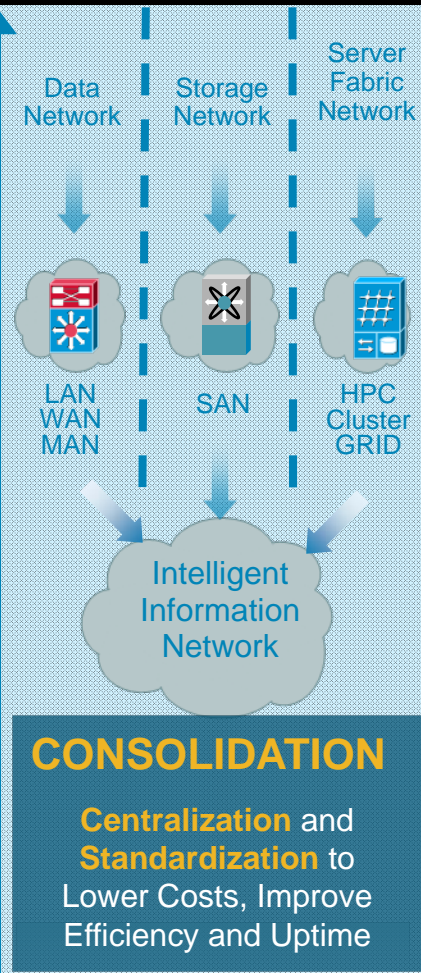
Operational
Continuity

Transport
Flexibility

Infrastructure
Scalability

Evolution of the Data Centre Infrastructure

A Phased Approach...



Data Centre SAN Drivers

Large-scale Physical Consolidation & Virtualization

IT Requirements

- Drive for significant efficiency gains
- Support for tiered storage environments and ILM with data mobility
- Business continuity across a broad range of applications
- Rapidly respond to changing business conditions

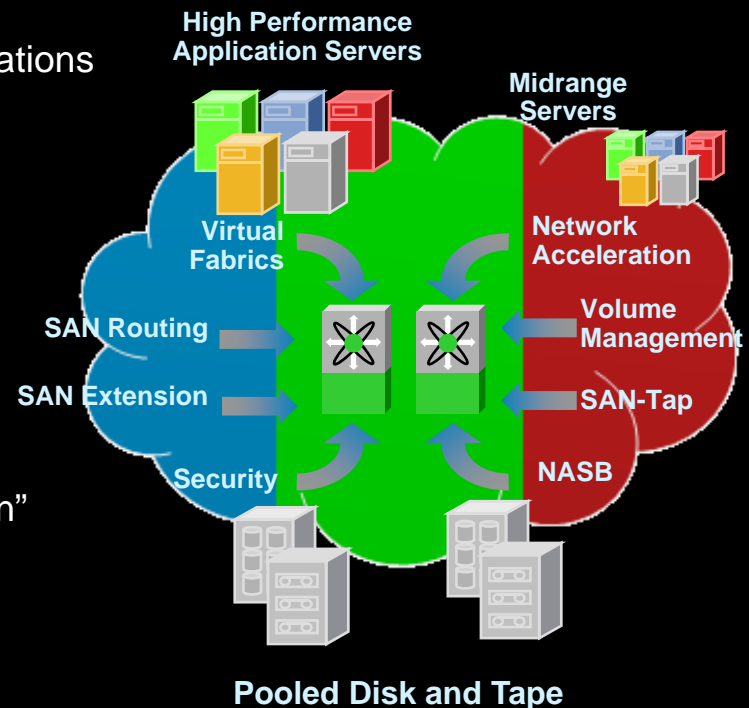
Enabling technologies

- Ultra-scalable directors (500+ ports) with:
 - Virtual fabric (VSAN) AND
 - Integrated routing (IVR)
- Network-hosted volume management, data movement/migration, replication & protection
- Comprehensive set of SAN extension solutions (e.g.. IP & Optical) WITH “I/O acceleration”

Benefits

- Dramatic simplification of storage environment
- Greatly improved service levels
- Significant reduction in downtime, management expense and TCO

Large Scale Physical Consolidation with Resource Virtualization



Network Hosted Intelligence, Optimization & Acceleration

Cisco Enabling Technologies...

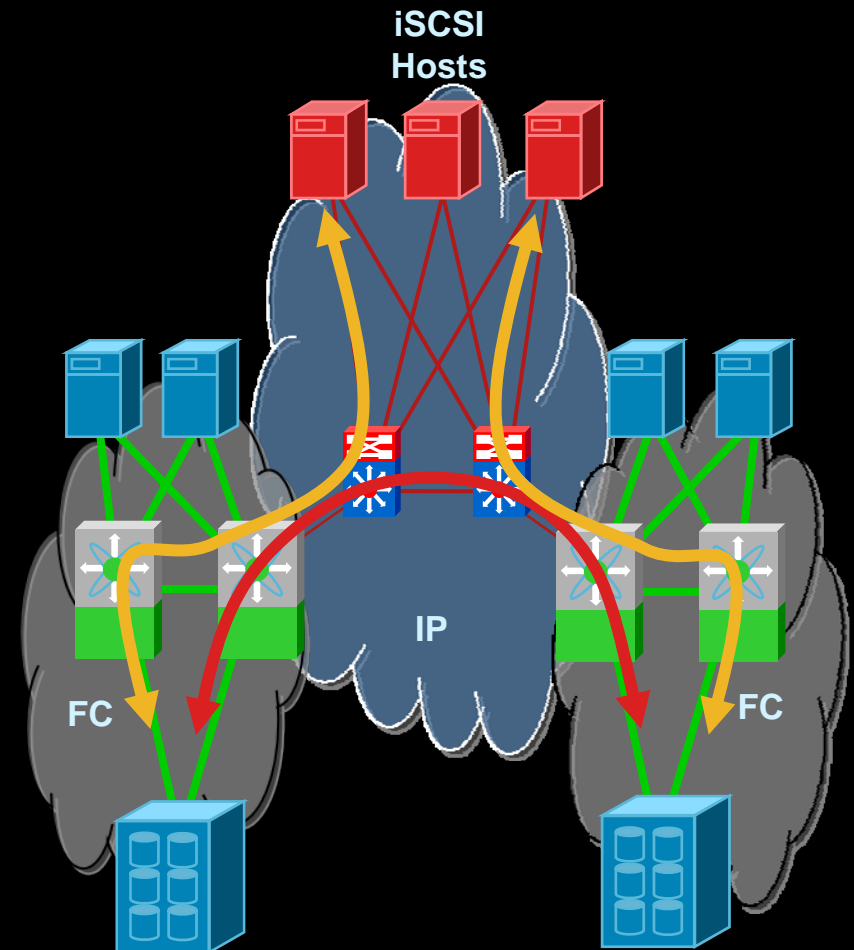


... for Fabric Consolidation

Cisco MDS9000 – Fabric Consolidation

Multiprotocol & Multitransport Support

- Fibre Channel Protocol (FCP)
 - 1,2,4 & 10 Gbps FC-SW and FC-AL support
 - Full range of port modes
- IBM Fiber Connection (FICON)
- Gigabit Ethernet Support, for...
- Small Computer System Interface over IP (iSCSI)
 - Allows Ethernet attached initiators to access FC attached targets
 - Extend SAN benefits to more servers
 - Network Boot over iSCSI
- Fibre Channel over IP (FCIP)
 - Connect FC & FICON fabrics over an IP network
 - Use IP networks for cost-effective business continuance solutions
 - Compression & Encryption support

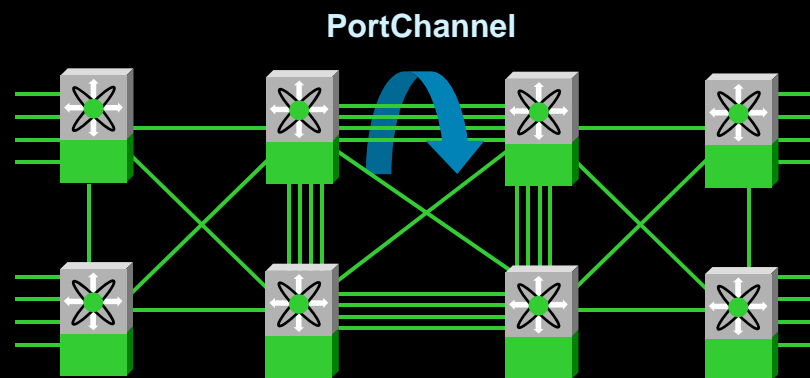


Cisco MDS900 - Fabric Consolidation

Data & Control Plane Scaling

Cisco PortChannel Link Aggregation

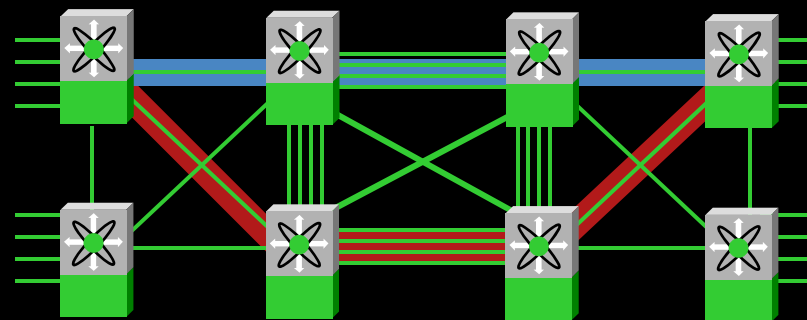
- Adds Performance scalability and resilience
- Group up to 16 links for aggregate of up to 160 Gbps (10G FC interfaces!)
Any port, any line card, no restrictions
- Hardware-based intelligent load distribution, with rapid failover and re-distribution
- PortChannel Protocol (PCP) for simplified auto set-up and configuration validation



Robust & Highly Scalable Control Plane*

- FSPF Routing for up to 16 equal cost paths (1 PortChannel = 1 link)
With traffic engineering based configurable link costs per Virtual SAN
- Support for large scale fabrics of up to 12 hops/fabric
- 239 Switches per Virtual SAN
- 8000 Zones per Switch
- 20000 Zone Members per Physical Fabric
- All zoning in hardware!

FSPF Load Balancing



* Cisco maximums – OSM qualified values maybe smaller

MDS 9500 Family of Ultra Scaleable Directors

Total Flexibility – any module, any chassis, any combination

Chassis Options

MDS 9513



MDS 9509



MDS 9506



Modules Options

4-port 10Gbps FC



12-port 4Gbps FC



24-port 4Gbps FC



48-port 4 Gbps FC



MSM 18/4 & 18/4 FIPS
18 x 4Gbps FC + 4 x GE



Cisco MDS 9100 & 9200 Switches

Scaleable Edge Switching Solutions

Top of Rack & Small Fabric

- Support for 16 Virtual SANs (VSANs)
- Flexible port licensing
- Optional Switch or N_Port Virtualizer (NPV) modes
- NPIV Support



MDS 9124



MDS 9134

Blade Servers

- Support for 16 Virtual SANs (VSANs)
- Flexible port licensing
- Optional Switch or N_Port Virtualizer (NPV) modes
- NPIV Support



**MDS 9124e for
HP Blade Server**



**MDS FCSM for
IBM BladeCenter**

Multiprotocol & Expandable

- 18 4Gbps FC plus 4 x 1GE ports
- Concurrent FCIP & iSCSI Support
- Hardware based Compression & Encryption
- Expansion slot for any MDS module

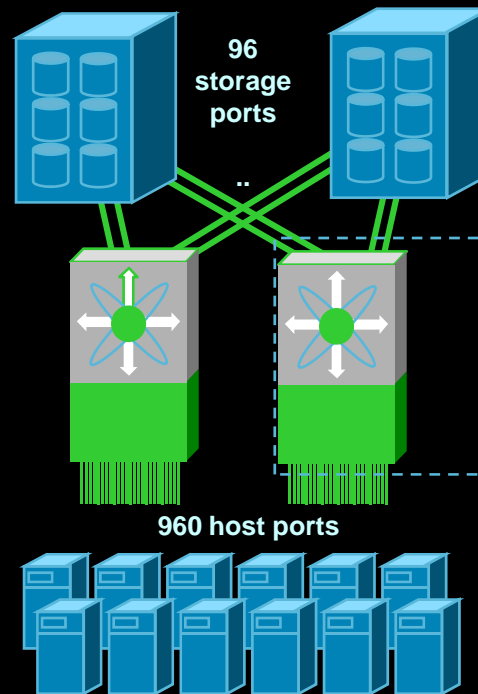


MDS 9222i

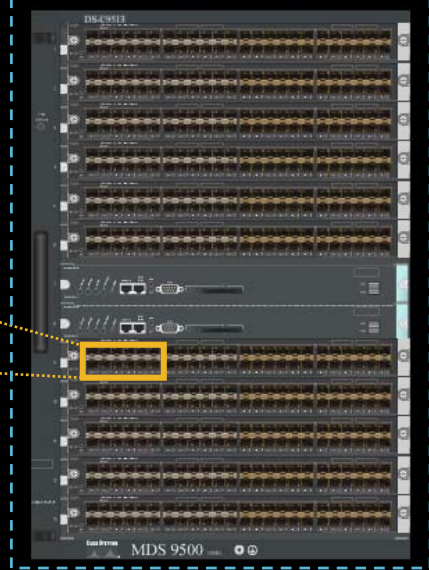
Sample Medium Scale Dual Fabric Collapsed Core Design

Dual Director Switches

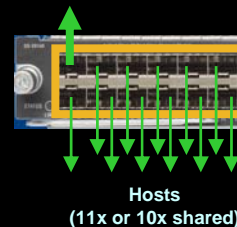
- Up to 528 ports per fabric
- Medium scale design leveraging 48-port modules
- Port bandwidth reservations to provide high density cost-effective solution
- VSAN support enables SAN island consolidation
- Port Bandwidth reservations guarantee performance where needed
- 100% Efficient design – no “wasted” ports



- 11 x 48-port modules
- 528 ports total
- 48 ports for storage
- 480 ports for hosts
- 10:1 oversubscription



Within each port group:
 1 or 2 ports to storage
 11 or 10 ports to hosts
 Port Bandwidth reservations used



Ports Deployed:	528
Usable Ports:	528
Unused (available) ports:	0
Design efficiency:	100%
End-to-end oversubscription:	10 : 1 (480 : 48)

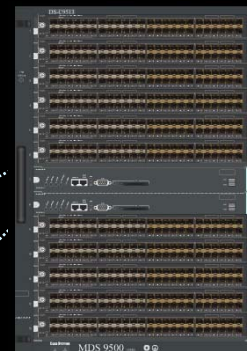
Sample Large-Scale Core-Edge Design

Director Core/Edge Design

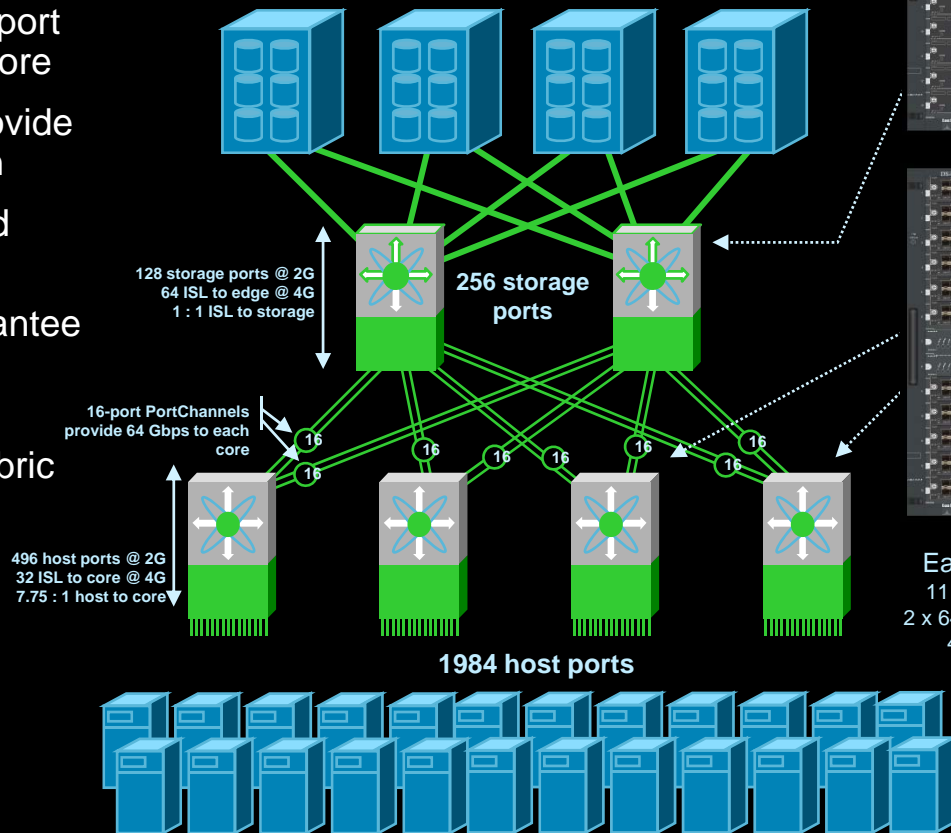
Six Director Switches per fabric

- Up to 2240 useable ports per fabric
- Large scale design leveraging 48-port modules at edge & 12/24-port at core
- Port bandwidth reservations to provide high density cost-effective solution
- VSAN support enables SAN island consolidation
- Port Bandwidth reservations guarantee performance for ISLs
- PortChannels deliver scaleable bandwidth, high availability and fabric stability

Each core switch
6 x 24-port modules (storage)
5 x 12-port modules (ISLs)



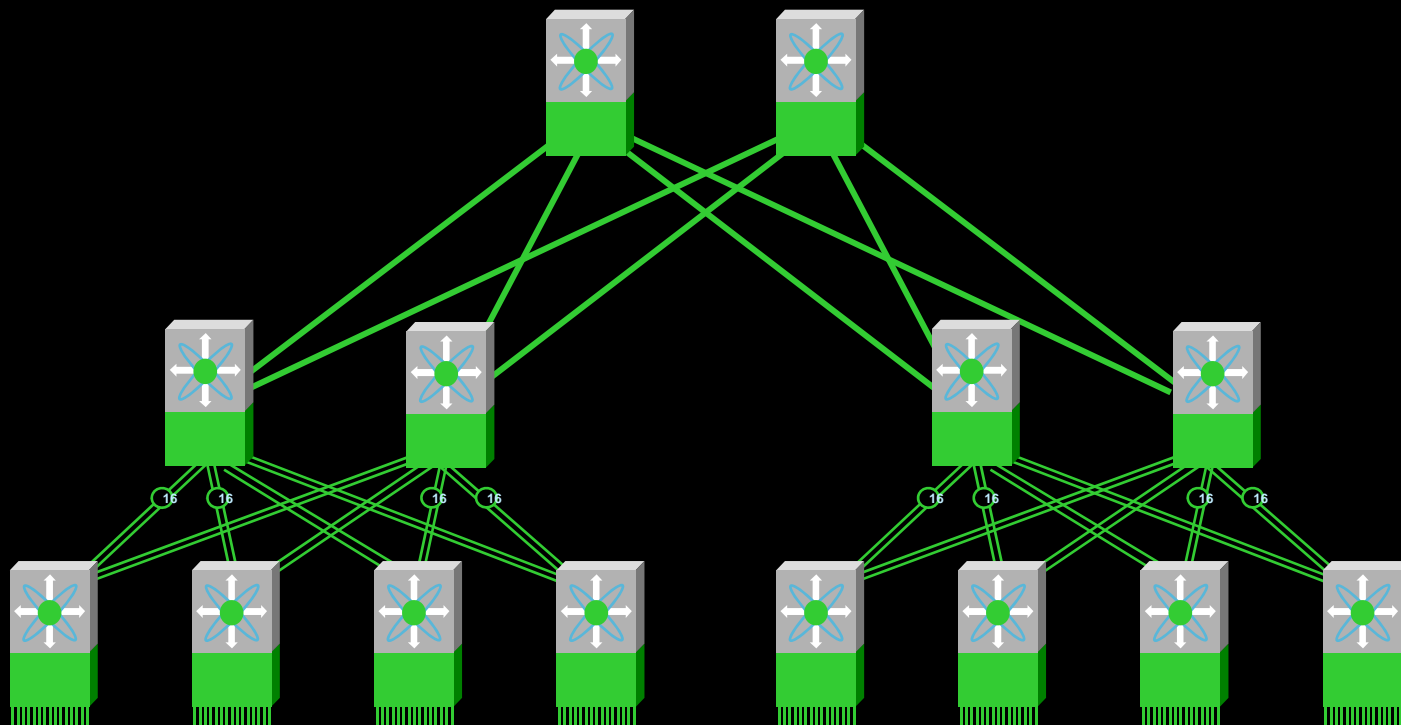
Each edge switch
11 x 48-port modules
2 x 64Gbps PortChannels
496 hosts/switch



Ports Deployed:	2544
Usable Ports:	2240
Unused (available) ports:	24
Design efficiency:	89%
End-to-end oversubscription:	7.75 : 1 (1984 : 256)

Scaleable 3-Tier Core-Edge Design

Scales to over 10,000 ports per fabric



Cisco Enabling Technologies...



... for Fabric Virtualization

Virtualizing the Fabric – The Full Solution

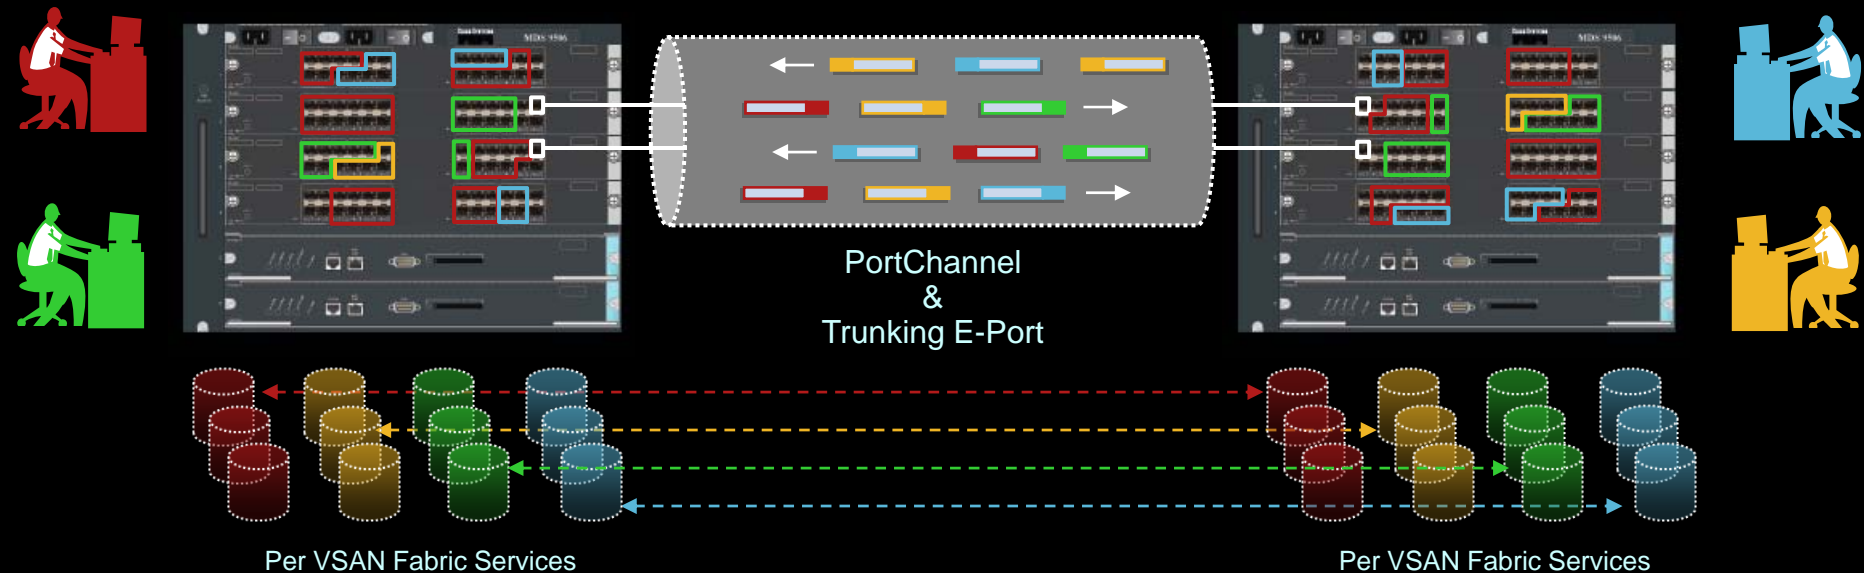


- Provide connectivity across virtual SANs without merging the fabrics
- Map and manage virtual fabrics independently
- Set FC parameters per virtual fabrics (e.g. timer values, FID allocation, DID ranges etc.)
- Define separate security policies per virtual fabric
- Troubleshoot per virtual fabric problems
- Separate fabric services per virtual fabric (e.g. routing, zones, RSCNs, QoS, etc.)
- Extend virtual fabric service to FC ISLs, FCIP, FICON, iSCSI etc.
- Assign virtual fabric membership at the port level

Full Service End-to-End Virtual Fabric Implementation

Cisco VSAN Technology at Work

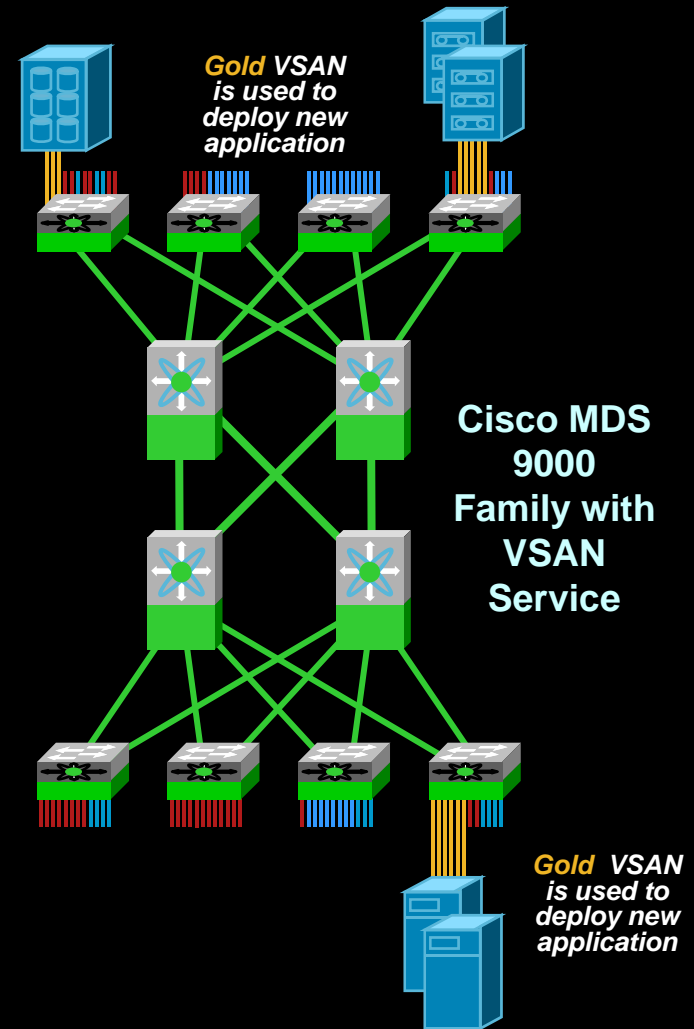
- VSANs create new instances of fabric services - separate policies and control traffic for each VSAN ensure fault isolation
 - VSANs do not exchange any control plane information (e.g. RSCNs, RCF, BF)
 - Each VSAN topology is independent and separate from the next
- Ports are individually assigned to VSAN (manual or automatically with DVPM)
- All frames (data and control) are tagged with VSAN Identifier when passing between Cisco switches providing hardware enforced separation of virtual fabrics
- Role Based Access Control (RBAC) allows for administrators per VSAN



VSANs as Migration Tool to SAN

Virtual SANs can be used to safely migrate applications to the SAN

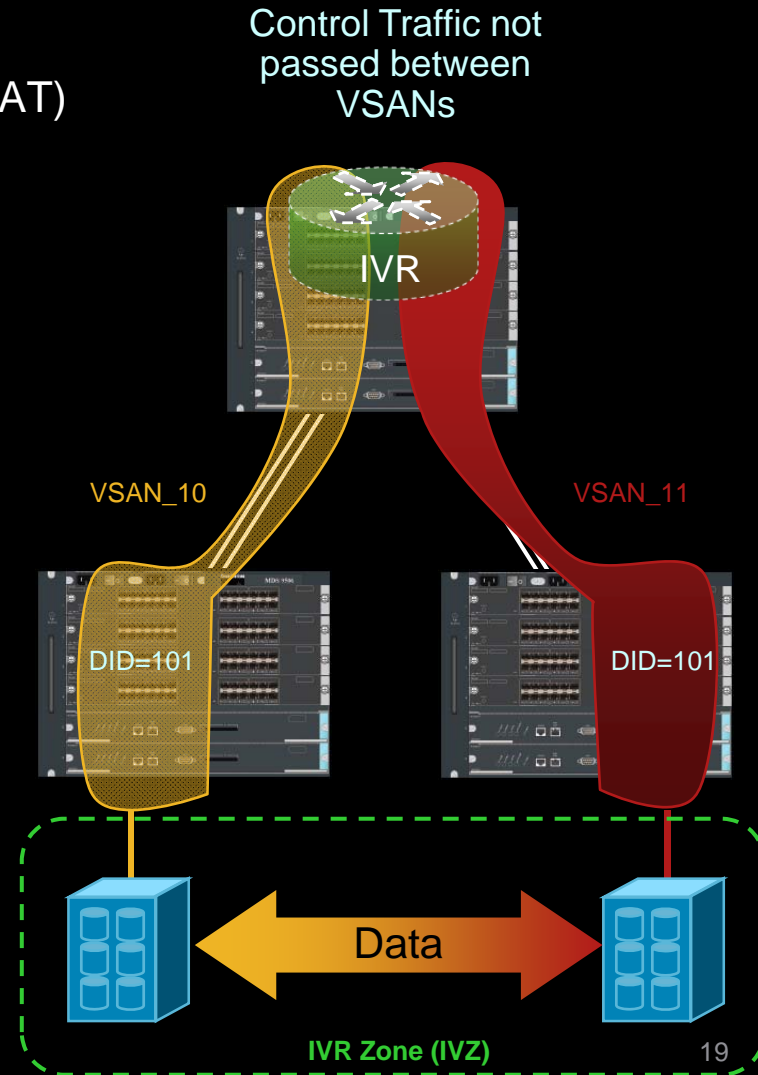
- Use any unused ports in the larger SAN fabric to assign to new application
- Isolate new SAN-attached devices to new VSAN to help restrict fault domain
- Using VSANs, control reach of traffic associated with new VSAN within larger fabric
- Merge VSAN into greater SAN after proper testing and evaluation has been completed
- Can simply use IVR as well to tie in new application to existing storage – no VSAN disruption



Cisco Integrated InterVSAN Routing (IVR)

- Allows resources in different VSANs to communicate
 - But without the need to merge fabrics
- Full bidirectional Network Address Translation (NAT)
 - Connect fabrics with overlapping Domain IDs
- Fully Standards based
 - Transparent to third-party switches
- Simple to set-up and manage
 - Uses well understood zoning principles for define allowed exchanges
- Provides high fabric resiliency and VSAN-based manageability
 - Distributed, scaleable, and highly resilient
- Supported at wire rate on any port on the MDS family!
 - No need for special routing modules or appliances

Wirespeed FC frame rewriting capability on every MDS 9200 & 9500 port is the foundation for delivering scaleable SAN Routing

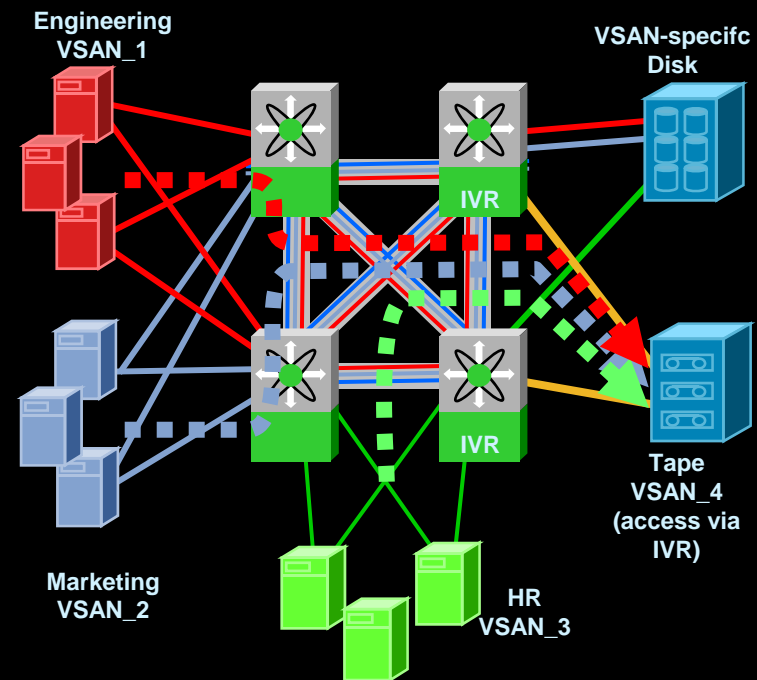


Cisco Inter-VSAN Routing (IVR)

Enables resources sharing across VSANs

IVR Applications in the Data Centre

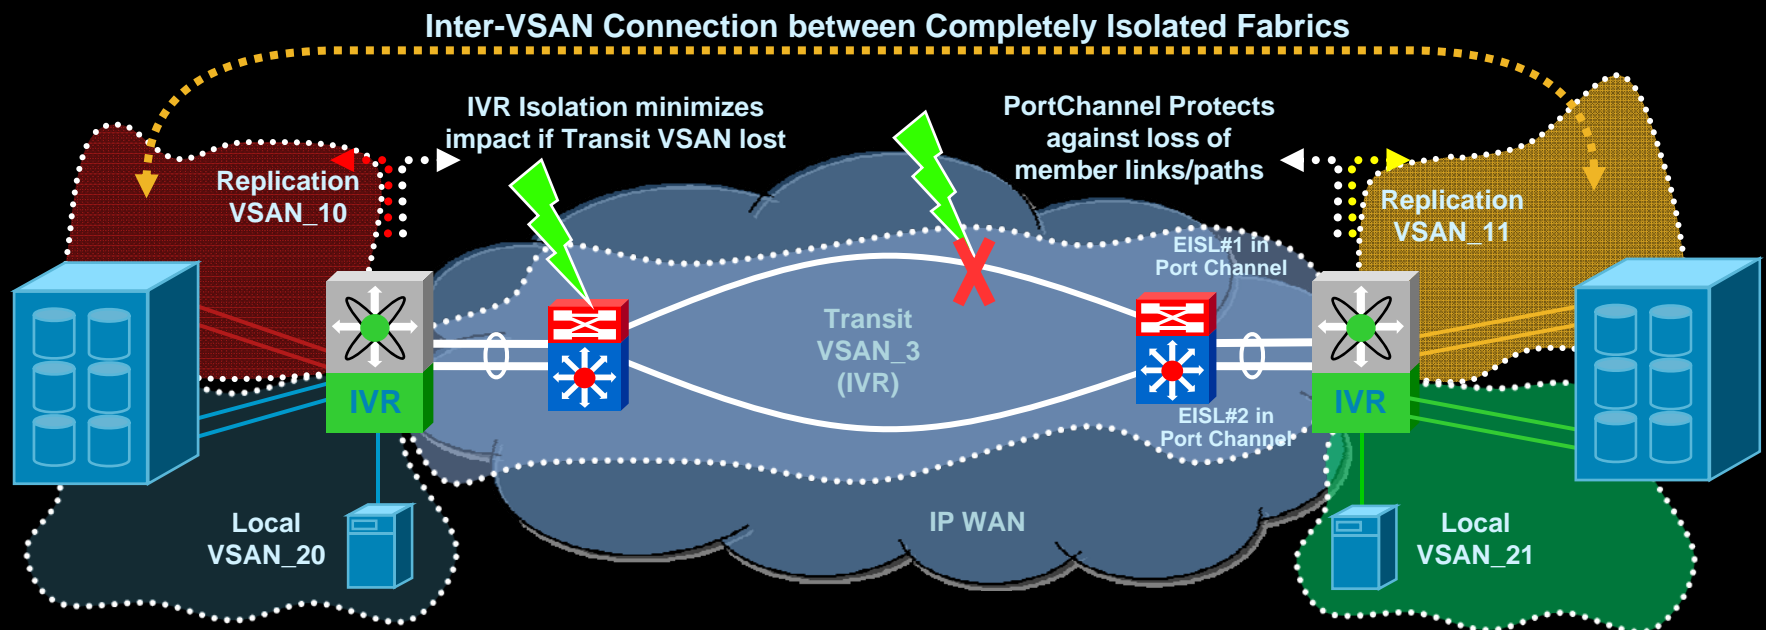
- Enables blade server per VSAN with 3rd party blade switches
- Allows production to back-up fabric connections
 - Allows low end servers (single HBA) to access separate back-up fabric
 - Allows blade servers to access back-up fabric
- Enables Tape Library sharing
- Eases Data migration in mixed vendor environment
 - Cisco per VSAN interop modes allow non disruptive interconnection of Brocade, McDATA, Qlogic & Cisco



Inter-VSAN Routing (IVR)

Enables Resilient SAN Extension Solutions

- Minimize the impact of change in fabric services across geographically distributed sites
- Limit fabric control traffic such as RSCNs and Build/Reconfigure Fabric (BF/RCF) to local VSANs
- Flexible connectivity with the highest availability
- Tune timers for optimal recovery performance per VSAN e.g. increase RA_TOV, ED_TOV for Transit VSAN if needed whilst leaving Local VSAN values as default
- Works with any transport service (FC, SONET/SDH, DWDM/CWDM, FCIP)

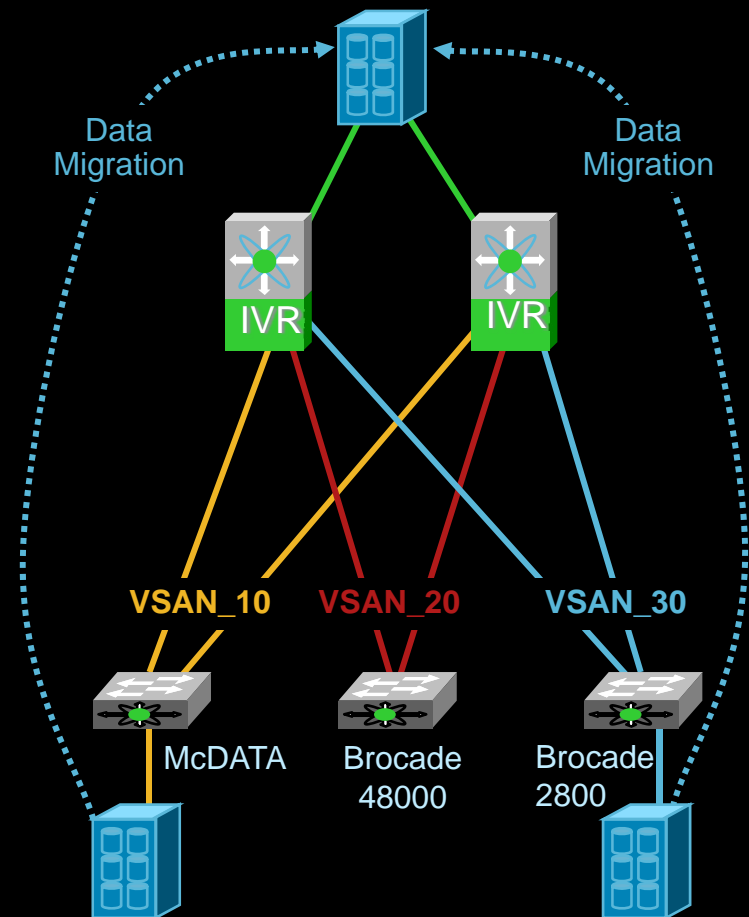


VSANs, IVR & Legacy Interop Modes

Enable fabric and data migration

Cisco Legacy Interop modes

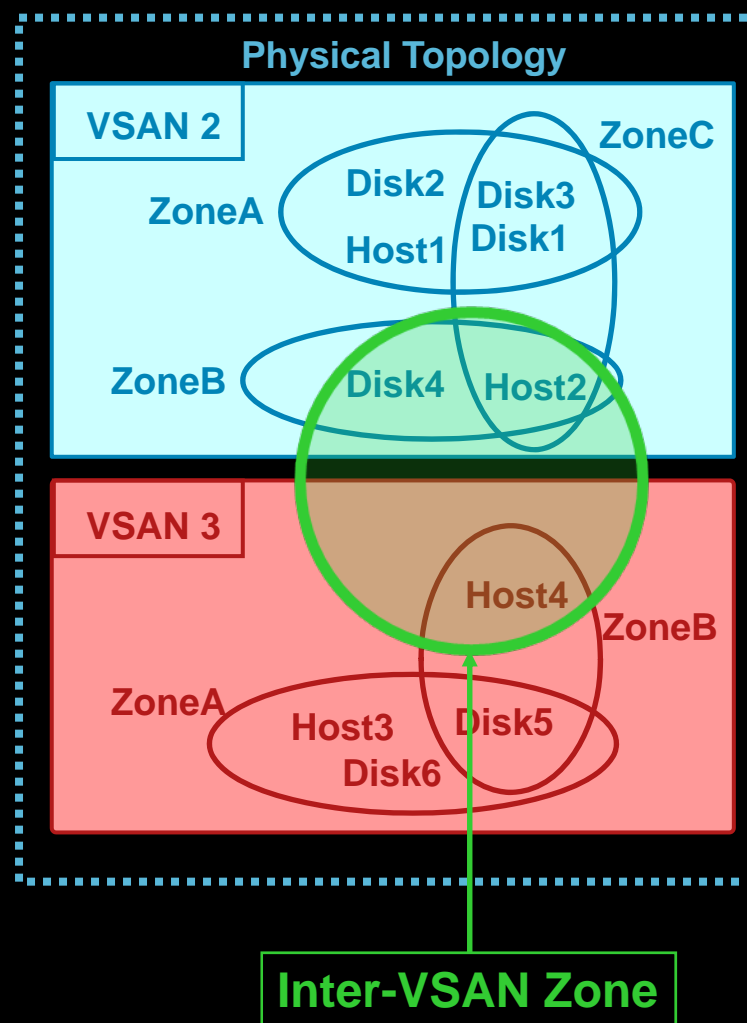
- Enables MDS 9000 family to interoperate with 3rd party switches in their 'Native Mode'
 - Re-use existing legacy fabric switches
 - No impairment to Cisco fabric
 - No change required on legacy switch
- Configurable on a VSAN-by-VSAN basis on MDS 9000
- Enhances standard interop mode
 - Mode 1 - Standard interop mode
 - Mode 2 - Supports Silkorm 2x00, 6400, and 3200/3800 (core_PID=0)
 - Mode 3 - Supports Silkorm 3900, 12000, 48000 (core_PID=1)
 - Mode 4 - Supports all McData platforms



VSANs, Zones, IVR Zones

- Hierarchical relationship:
 - First assign physical ports to VSANs
 - Then configure independent zones per VSAN
- VSANs provide traffic statistics
 - Zones provide added security and allow sharing of device ports
- VSANs only change when ports needed per virtual fabric
 - Zones can change frequently (e.g., backup)
- Ports are added/removed non-disruptively to VSANs
- IVR zone: a container or access control, containing two or more devices in different VSANs
 - Standard zones are still used to provide intraVSAN access
- IVR zoneset: a collection of IVR zones that must be activated to be operational

VSANs and Zoning Are Complimentary



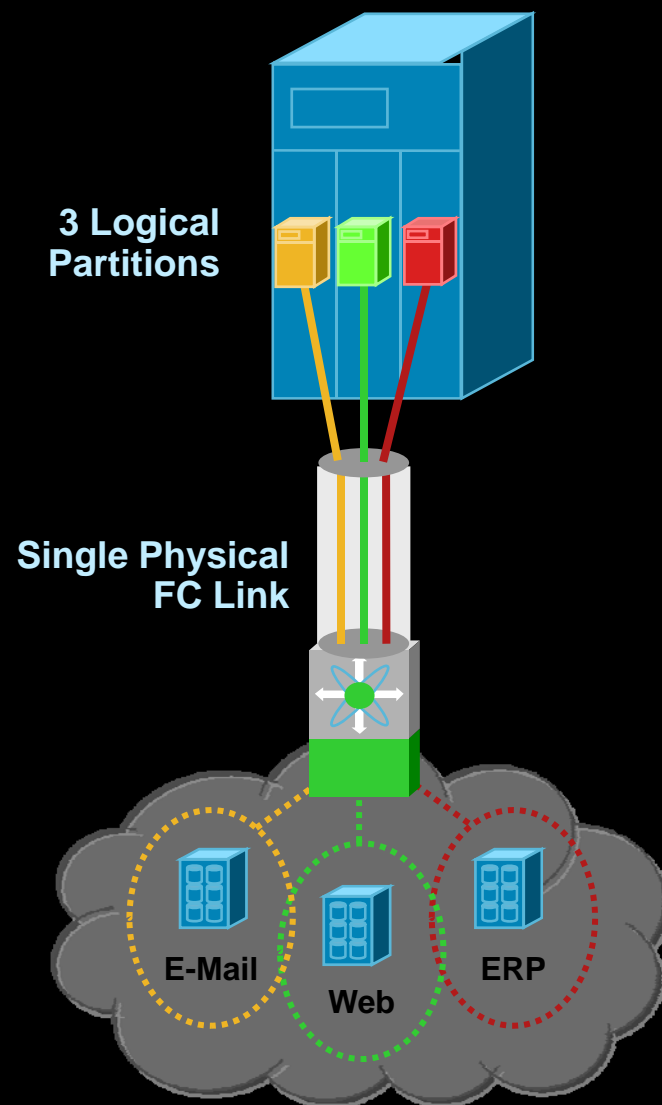
Cisco Enabling Technologies...



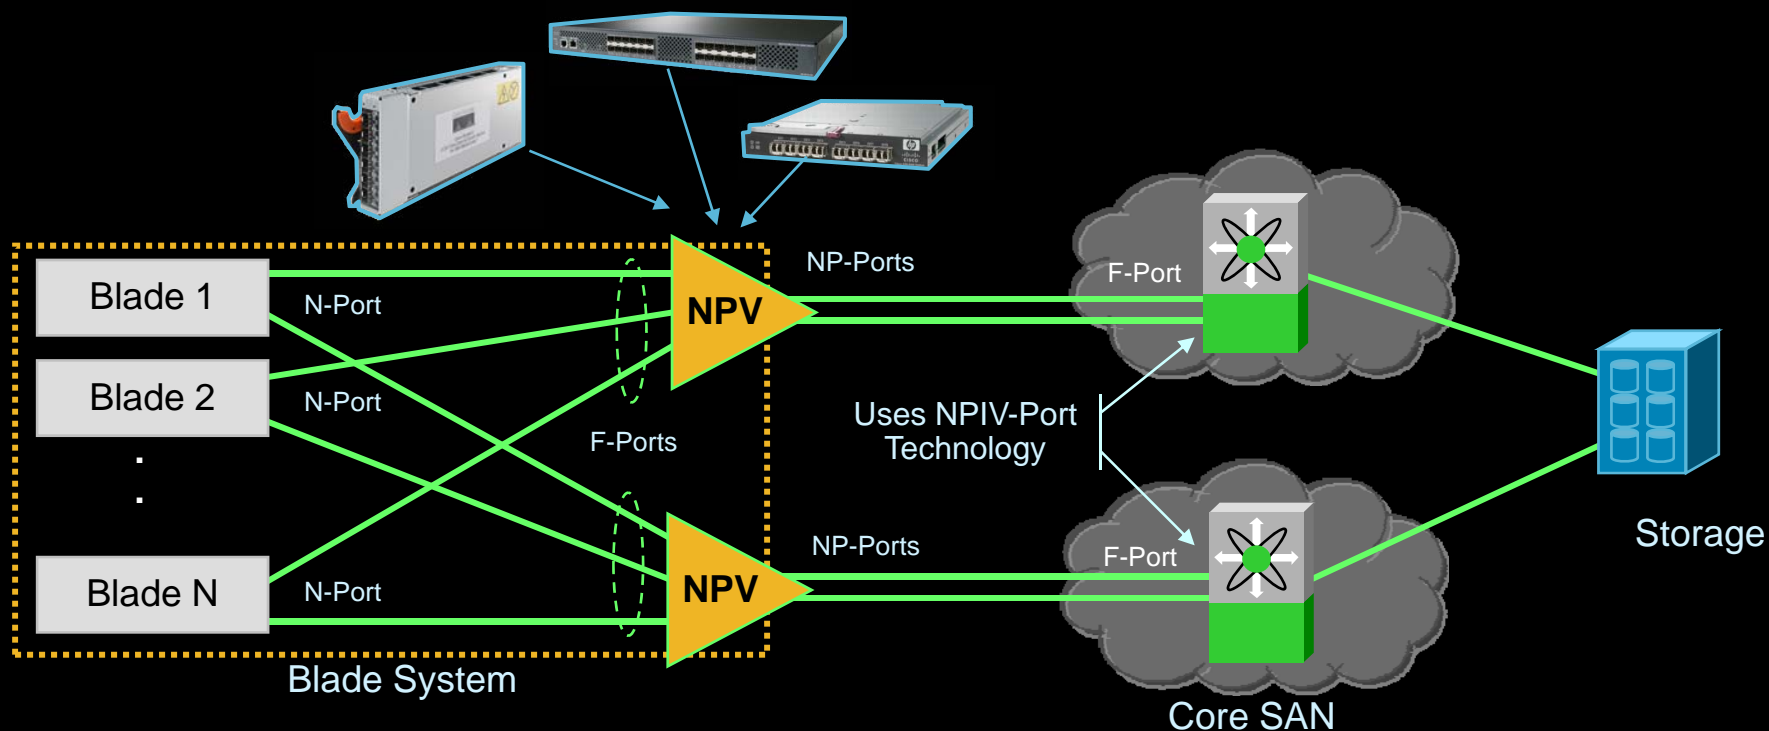
... support for Virtualization of End Devices

N-Port ID Virtualization (NPIV)

- NPIV is a standards-based technology specified by INCITS T11
- Allows HBA port sharing between different virtual machines
- Each virtual device logs into the fabric independently
 - 1st device uses FLOGI (e.g. HBA)
 - Subsequent devices use FDISC
- Each device registers independently with name service via PLOGI
- Enables Independent fabric policies per Virtual Machine e.g.
 - Zoning
 - Security
 - Traffic mgmt (e.g.. QoS)



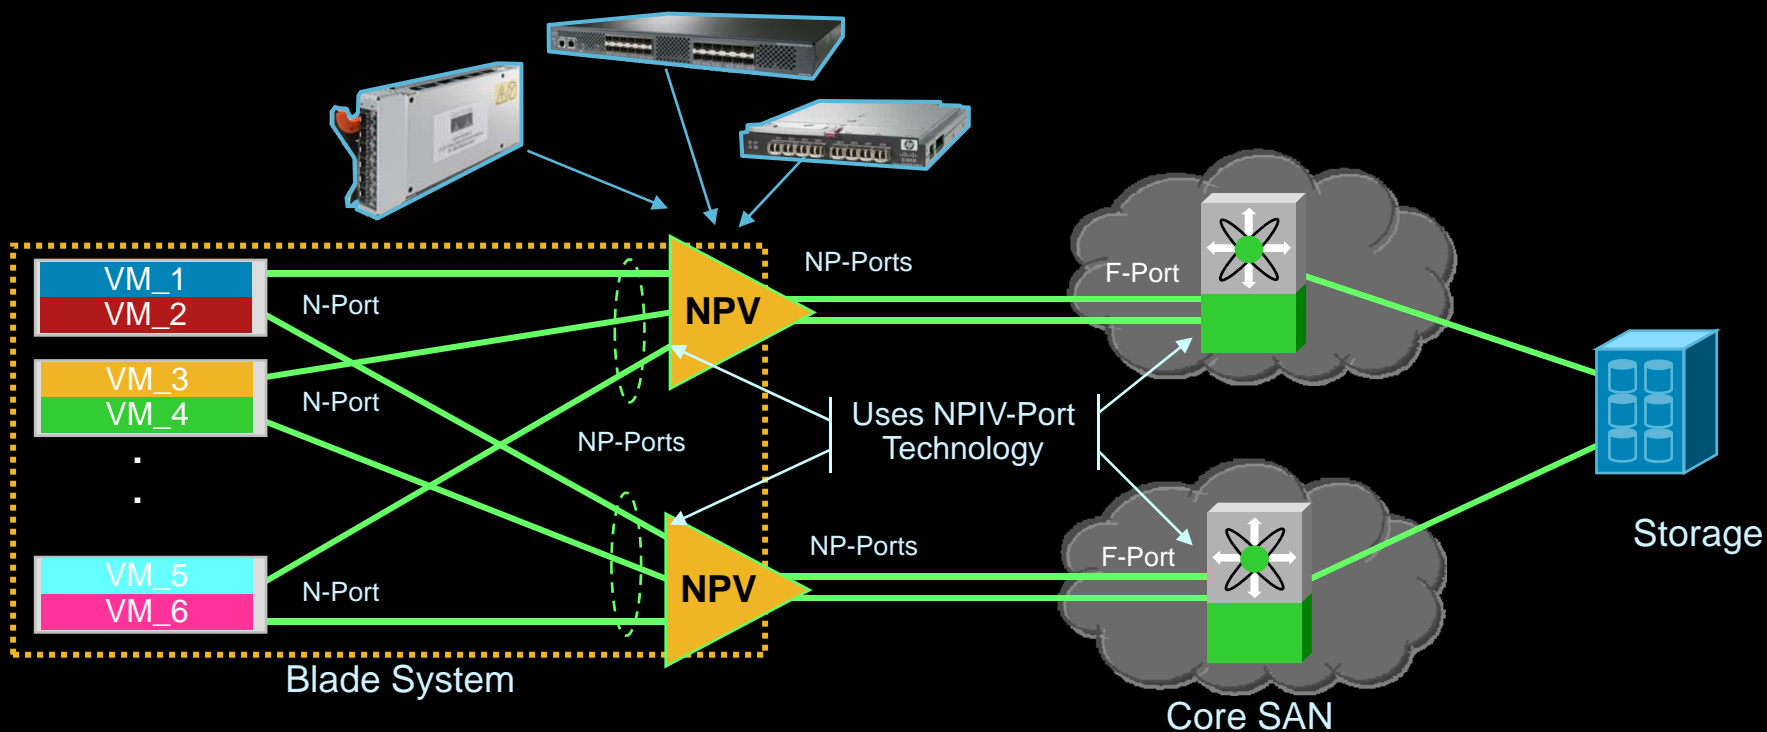
Introducing N-Port Virtualizer (NPV)



Key Benefits of NPV

- Solves the Domain ID issue. With NPV, Blade Switch appears as a HBA to the core
- Addresses the interoperability issues since the Blade Switch presents itself as an HBA
- Simplifies management since the server administrator is not exposed to SAN switch management tasks

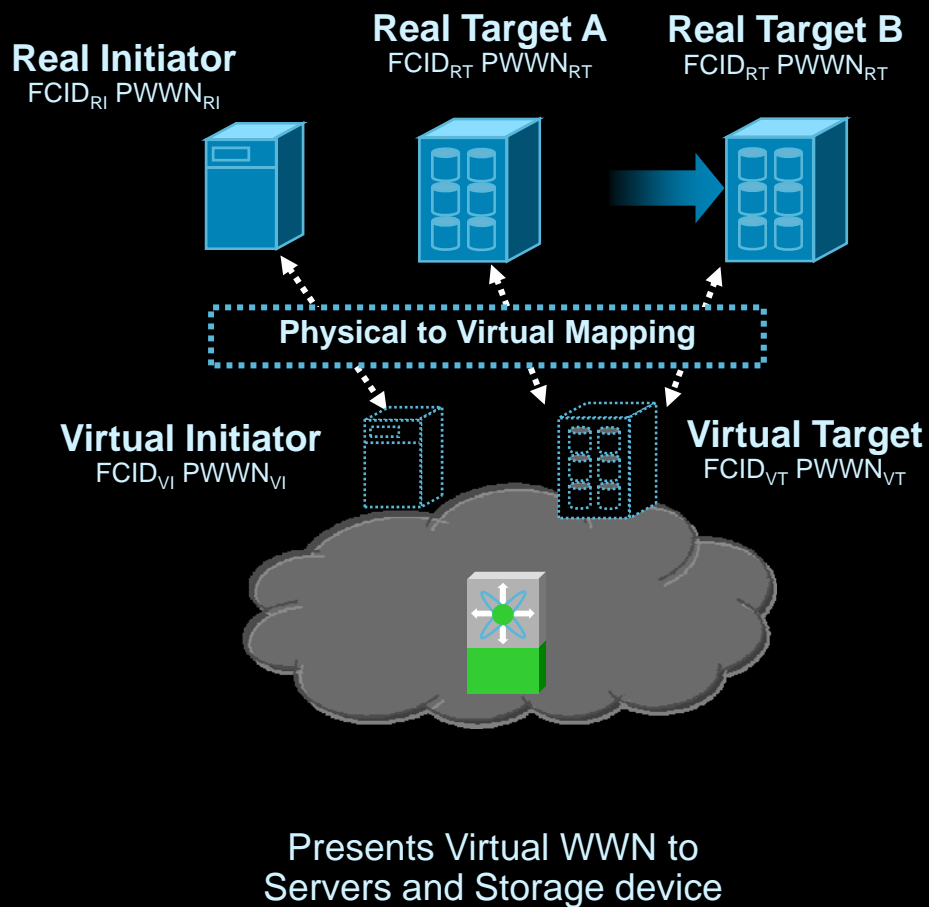
Nested NPIV & NPV – Total Flexibility



Key Benefits of Nest NPIV & NPIV

- Total Flexibility – Administration, Mobility, Control
- Massive Scalability – Build SANs with 1000's of Virtual Machines with extremely simple fabric topologies
- Retain management domains and access control
- Further Enhancements on the way – e.g. WWN virtualization

Cisco SAN Device Virtualization



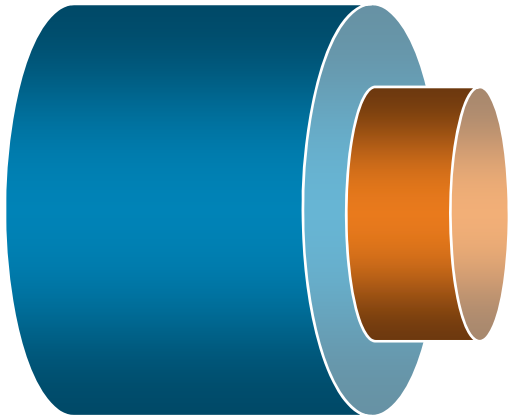
- Allows provisioning with virtualized servers and storage devices
- Significantly reduces time to replace HBAs and storage devices
 - No reconfiguration of zoning, VSANs, etc. required on MDS
 - No need to reconfigure storage array LUN masking after replacing HBAs
 - Eliminates re-building driver files on AIX and HP-UX after replacing storage

On the Horizon



Fibre Channel over Ethernet (FCoE)

-
-



Benefits

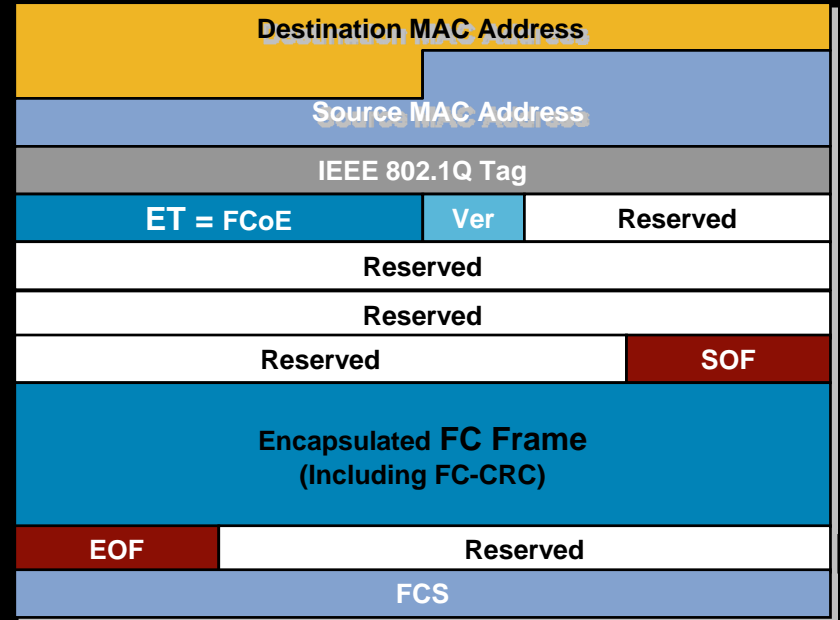
- Fewer Cables
 - Both block I/O & Ethernet traffic co-exist on same cable
- Fewer adapters needed
- Overall less power
- Interoperates with existing SAN's
 - FCoE SAN Management is complete consistent with FC SANs
- No Gateway require
 - Simple encapsulation and de-encapsulation at wire speed

FCoE Frame Format Specification

- An extension of FC over lossless Ethernet
- FCoE Specification in ANSI INCITS FC T11.3
 - Frame Format agreement Aug. 2007
 - Target completion 2H08
- Optional features being defined in IEEE
 - PAUSE enhancements 802.3x mechanisms
 - Priority Flow Control (PFC)
 - Jumbo Ethernet frames – allows 1:1 encapsulation of 2112 byte FC frames

FCoE Frame Format

Bit 0 Bit 31



Final Words



Cisco Consolidated & Virtualized SANs... Yield Tangible Cost and Time Savings for the Enterprise

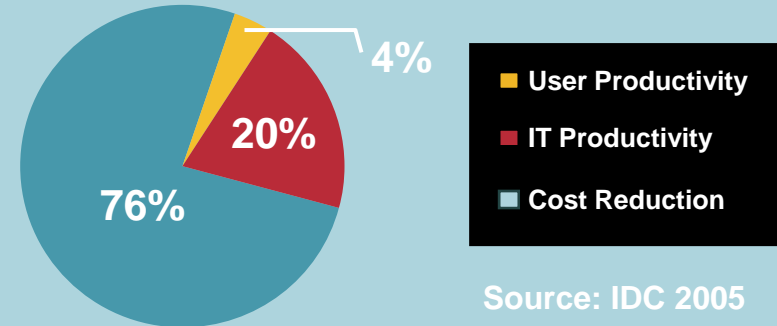
First-generation DAS to SAN savings

- Management simplification
- Storage utilization
- Ease of backup

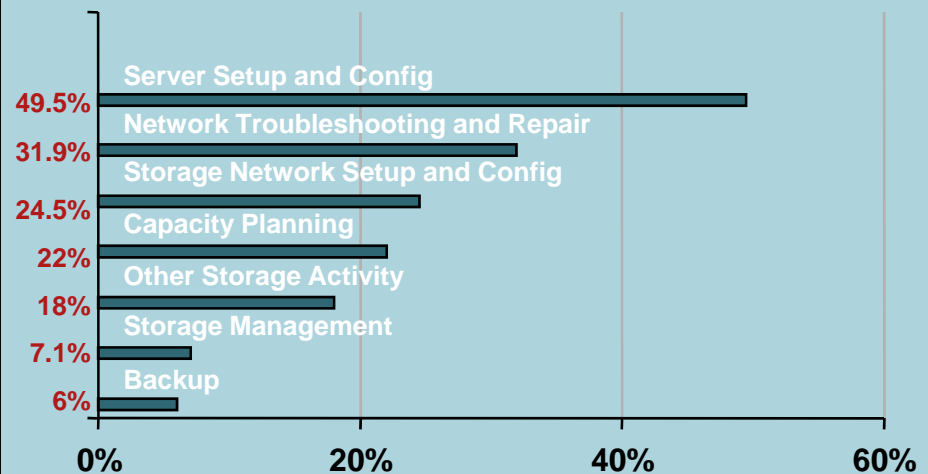
Incremental Savings with Cisco MDS 9000:

- 499% 3-yr ROI
- SAN, storage and server consolidation
- Increased storage utilization
- Increased staff productivity
- Increased reliability
- Automation of Data Center Resource Provisioning
- Integrated SAN Extension
- Network-hosted intelligent services

Incremental Hard Cost Savings with Cisco SAN



Incremental Reductions in IT Staff Time Spent on Storage Support Tasks



Conclusion & Key Points

- The MDS has many unique features designed to enable large scale virtualized fabrics
 - High performance control plane processing – per VSAN
 - Hardware based zoning
 - VSAN frame tagging on ISLs in hardware
 - Hardware based load balancing
 - Robust Quality of Service (QoS) mechanisms
- These features now extend to fabric and blade switches
- Integrated SAN routing on every port enables efficient fabric interconnection without merging SANs
- The MDS offers device identity virtualization designed to ease operational management of the fabric
- Future developments will add additional capabilities in all these areas

