Advanced SAN Fabric & Storage Virtualization

BRKDCT-3008

Tuqiang Cao

Cisco Networkers 2007

HOUSEKEEPING

- We value your feedback, don't forget to complete your online session evaluations after each session and complete the Overall Conference Evaluation which will be available online from Friday.
- Visit the World of Solutions on Level -01!
- Please remember this is a 'No Smoking' venue!
- Please switch off your mobile phones!
- Please remember to wear your badge at all times including the Party!
- Do you have a question? Feel free to ask them during the Q&A section or write your question on the Question form given to you and hand it to the Room Monitor when you see them holding up the Q&A sign.

Data Center Virtualization

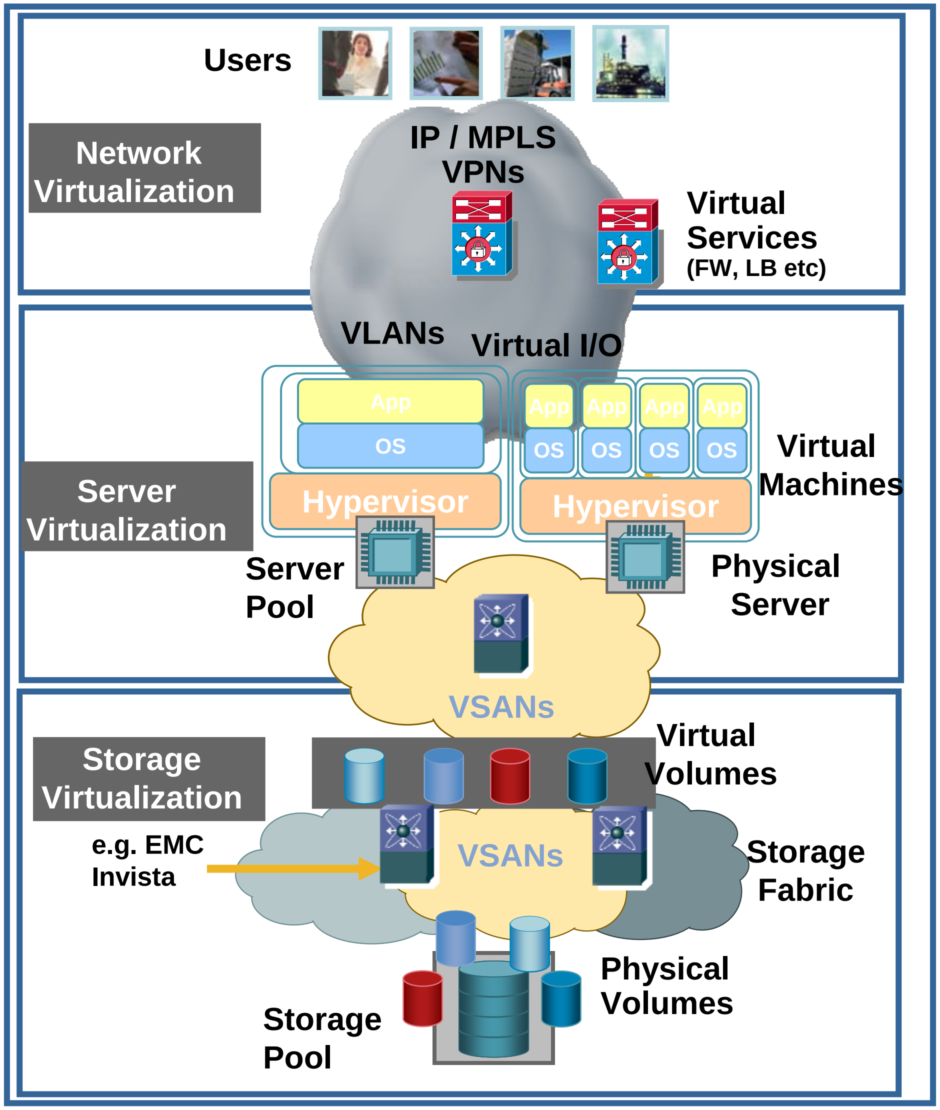

Network Virtualization

- Dynamically creates secure isolated environments for hosting apps on shared infrastructure
- Reduces number of appliances (FW, LB, SSL etc) for lower cost and power consumption
- Unified Fabric and Virtual I/O Reduces Network TCO

Server Virtualization

- Virtual Machine consolidation reduces number of servers to manage, deploy, power and cool
- Non-disruptive virtual machine migration allows upgrades, patches etc. without app downtime
- Flexibility: Bare-iron servers and virtual machines can be rapidly deployed to support existing or new apps

Storage Virtualization

- Network-hosted for scalability, availability and transparency
- Non-disruptive storage provisioning and migration of production data between systems
- Seamlessly upgrade or migrate storage
- Point-in time copy across heterogeneous systems

Outline

SAN Virtualization

- Fabric Virtualization
  - What are Virtual Fabrics?
  - Cisco's Virtual SANs
  - How do VSANs Work?
- Fabric Routing
  - What is Fabric Routing?
  - Cisco's Inter-VSAN Routing

Storage Virtualization

- Storage Services Module (SSM)
- Intelligent Fabric Applications
  - Network Accelerated (NASB and FC-WA)
  - Network Assisted (SANTap)
  - Network Hosted (Storage Virtualization)
- Network Assisted Serverless Backup
  - LAN Backup
  - Serverless Backup
- Fibre Channel Write Acceleration
  - Synchronous Replication
- SANTap
- Storage Virtualization
  - EMC INVISTA

MDS 9000 Fabric Switch Positioning

Cisco Positioned to Extend Reach All Market Segments

Industry-Leading Investment Protection Across a Comprehensive Product Line

Enterprise and Service Provider

Small/Medium Business

MDS 9000 Family Systems

MDS 9120 and 9140

MDS 9216 and 9216i

MDS 9506

MDS 9509

MDS 9513

MDS 9000 Modules

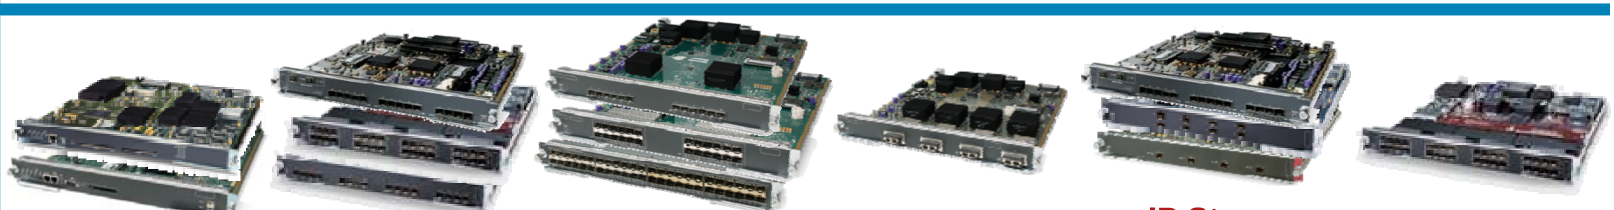

Supervisor 1 and 2

14-Port, 16-Port, 32-Port 1 and 2 Gb FC

12-Port, 24-Port, 48-Port 1, 2 and 4Gb FC

4-Port 10Gb FC

IP Storage Services–iSCSI and FCIP

SSM (Storage Services Module)

Mgmt. — Cisco Fabric Manager

OS — Cisco MDS 9000 Family SAN-OS

Part 1

SAN Virtualization or Virtual Fabrics

Fabric Virtualization and Fabric Routing

Three Key Concepts

- Fabric Virtualization
  Provide independent ('virtual') fabric services on a single physical switch
- Virtual Fabric Trunking
  Ability to transport multiple virtual fabrics over a single ISL or common group of ISLs
- Fabric Routing
  Ability to provide selected connectivity between virtual fabrics without merging them

Fabric Virtualization

SAN Islands Have Purpose: At a Cost

- SAN islands are built to address several technical and non-technical issues:
  - Maintains isolation from fabric events or configuration errors
  - Provides isolated and controlled management of island infrastructure
  - Driven by bad experiences of large multi-switch fabrics
- However…
  - Often over-provisioned port count for future growth—wasteful and costly
  - Very widespread issue today—some analysts still recommending islands

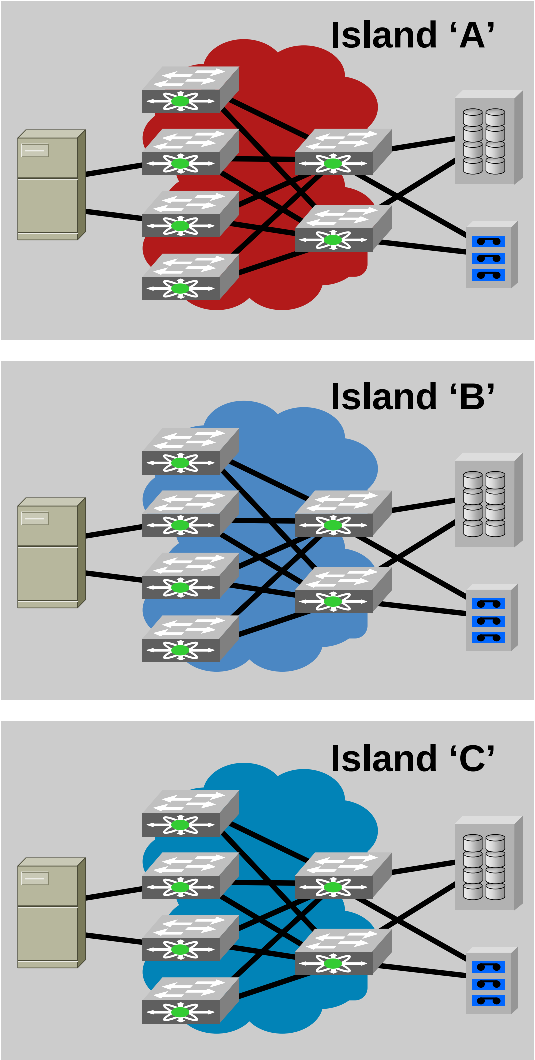

Introducing Fabric Virtualization

Fabric Virtualization Provides:

- A method to divide a common physical fabric into virtual domains
- An infrastructure analogous to VLANs in the Ethernet world
- A method to still isolate virtual fabrics from one another for
  - High availability
  - Security
  - Management
- A method to reduce wasted ports as experiences in the island approach
- A method to charge-back for used resources from the physical fabric

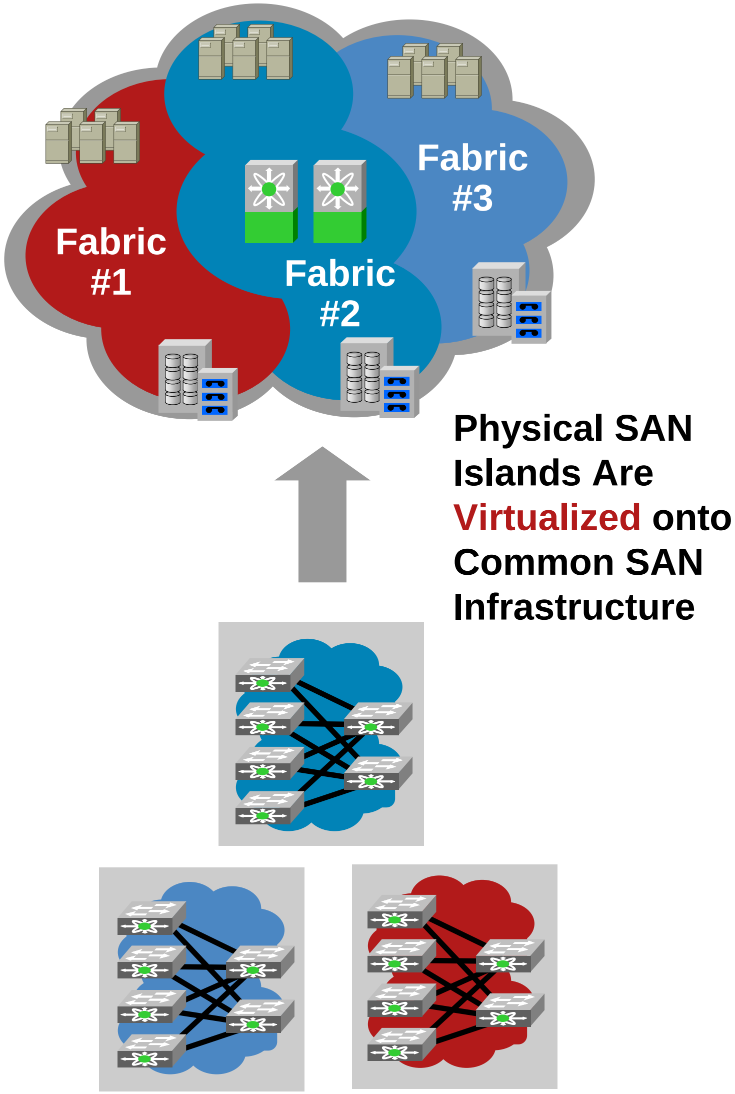

Uses For Fabric Virtualization
Cost-Effective Development, Staging, Backup SANs

- Instead of building separate physical development fabric, build a virtual one
  - Migrate to existing fabric later
  - Use free ports in larger SAN
- Build a virtual tape backup SAN
  - Can be expanded using routing to share tape resources
- Build a staging SAN for new applications or servers
  - Test stability in isolated staging virtual fabric before adding into larger SAN

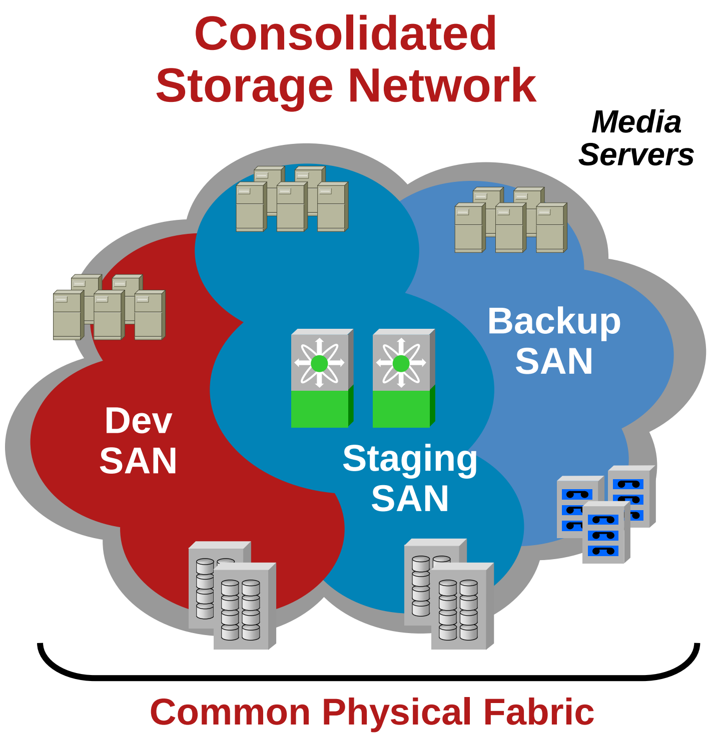

Uses For Fabric Virtualization
Cost-Effective SAN Extension Integration

- Overlay data replication fabric(s) on common physical fabric

  No need for separate pair of switches for each replication connection

  Use one virtual fabric per replication connection

- A *bonus* is to be able to share common SAN extension circuits amongst multiple virtual fabrics
- Fabric routing adds to resiliency of solution

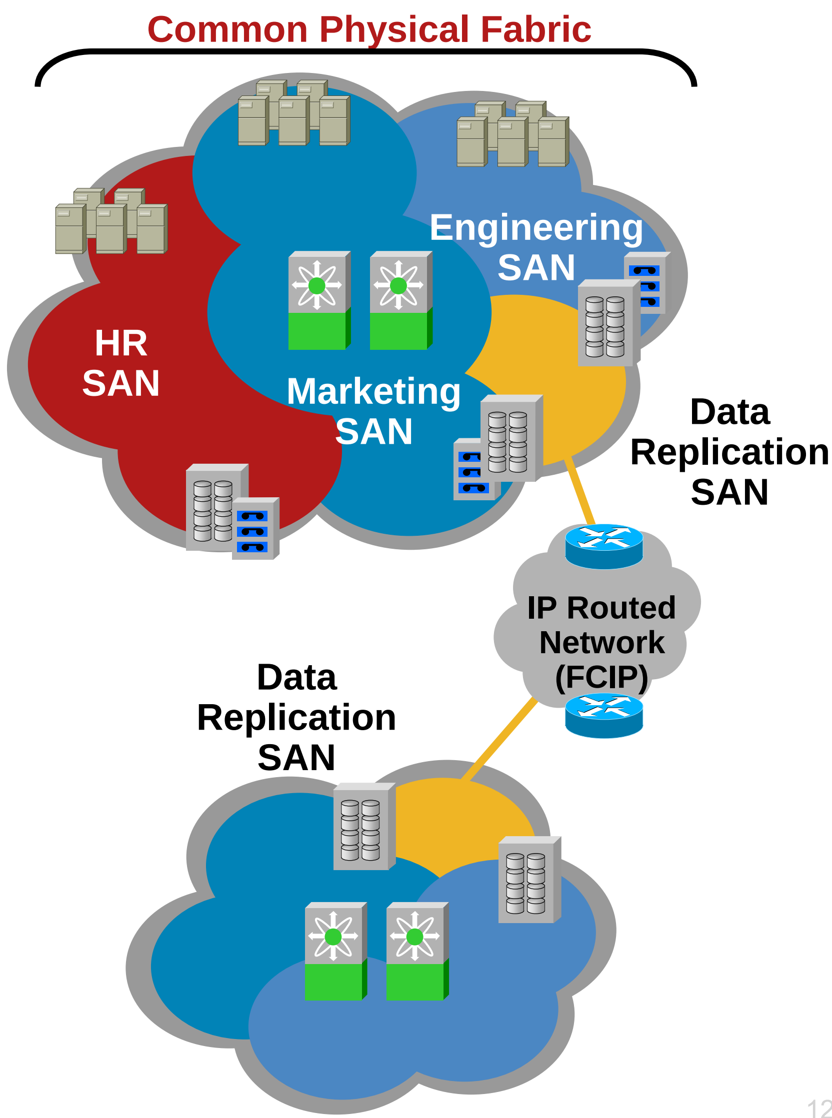

FCIP: Fibre Channel over IP

Three Approaches to Fabric Virtualization

Switch-Based

- Switch line-card partitioning
- Island-level granularity
- No shared ISLs
- Interconnection, but no consolidation

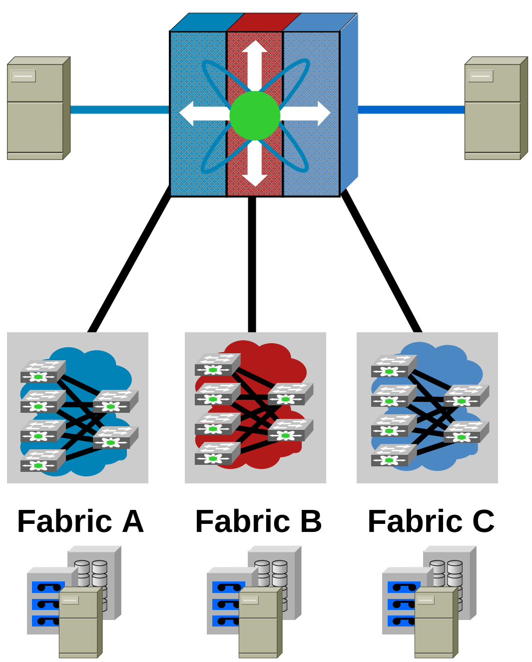

Appliance-Based

- Dedicated appliance provides routing
- Island-level granularity
- No shared ISLs
- Interconnection, but no consolidation

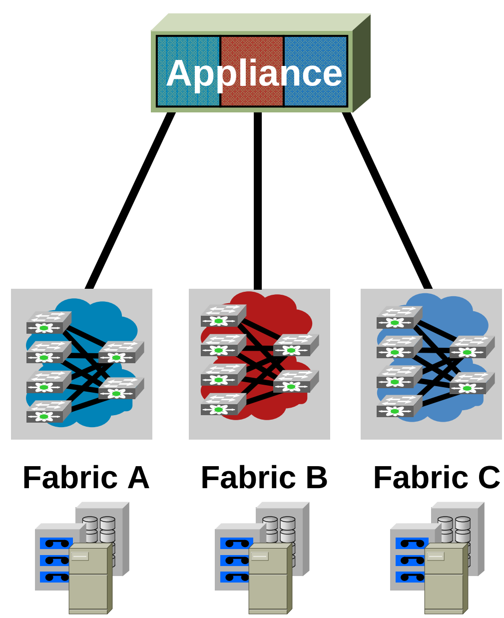

Fabric-Based

- Fabric-wide virtualization via hardware partitioning
- Port-level granularity
- Fully shared ISLs
- Drives consolidation

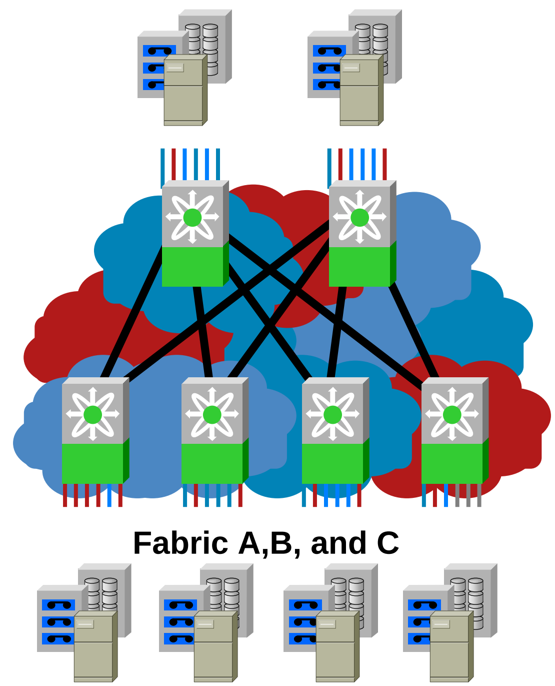

Fabric Virtualization: MDS 9000 Family

- Each port on the MDS 9000 family exists in a VSAN
- Up to 256 VSANs in a single switch (hardware can support up to 4095)
- Logical configuration to move a port from one fabric to another
- WWN-based VSANs can provide automated VSAN membership
- Basis for Virtual Fabric Trunking (VFT) Extended Header (ANSI T11 FC-FS-2 section 10)

Cisco’s VSAN Accepted as Standard

- Cisco’s proposal of a VSAN header specification was accepted by the working group
- Small variation (but compatible) to current VSAN header

Cisco VSAN Header (8 bytes)

R_CTL	Ver	Frame Type	MPLS Present	More Header	User Priority	VSAN Number	CDL Present	TTL	# PAD Bytes	P_VL	Rsvd	OAM	Msg Info
8	2	4	1	1	3	12	1	8	2	4	2	8	8

FC-FS-2 Virtual Fabric Tagging (VFT) Header (8 bytes)

R_CTL (0x50h)	Ver 0x00b	Frame Type	RSVD	RSVD	User Priority	VF_ID	RSVD	HopCt	RSVD
8	2	4	1	1	3	12	1	8	24

VSANs + FCIP for WAN Cost Savings

- Cost savings from multi-application SAN extension consolidation
- Multiple VSANs carried securely over Port Channeled FCIP links
- VSANs can be scaled and provisioned independently of FCIP and WAN link provisioning

2 X FCIP Portchannel with TE (Trunking VE_Port)

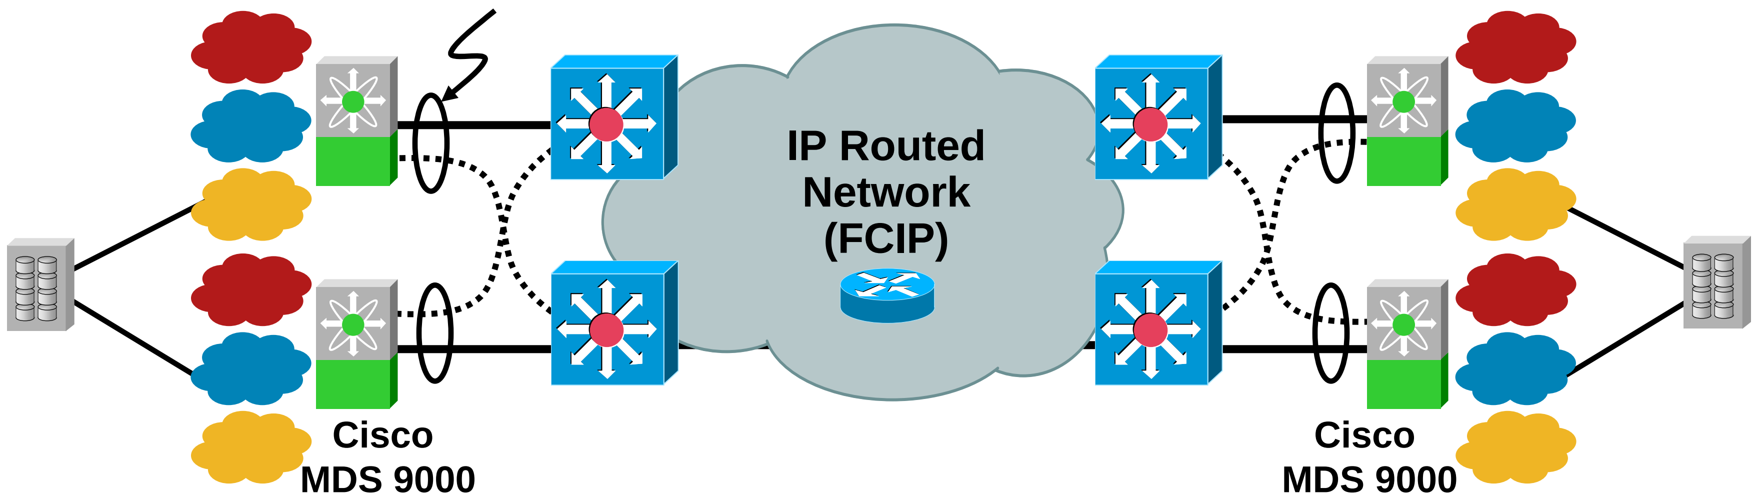

VSANs + FICON for Fabric Consolidation

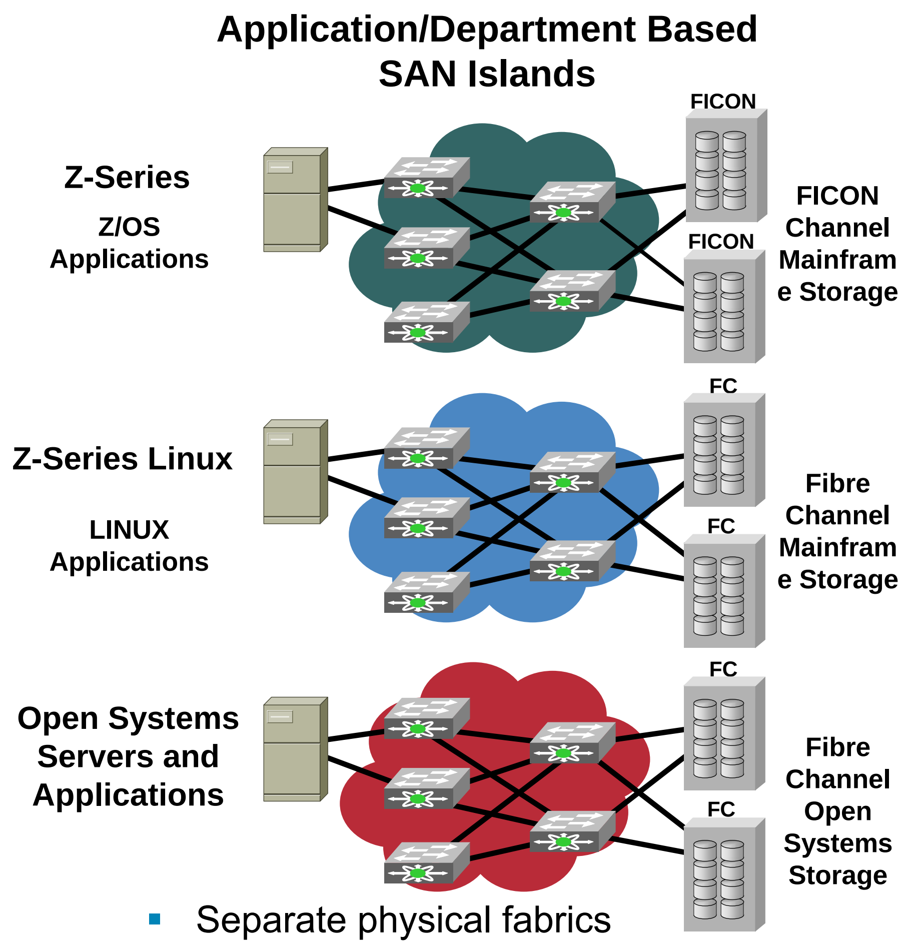

- Separate physical fabrics
- Over-provisioning ports on each island
- High number of switches to manage

FICON: Fiber Connectivity

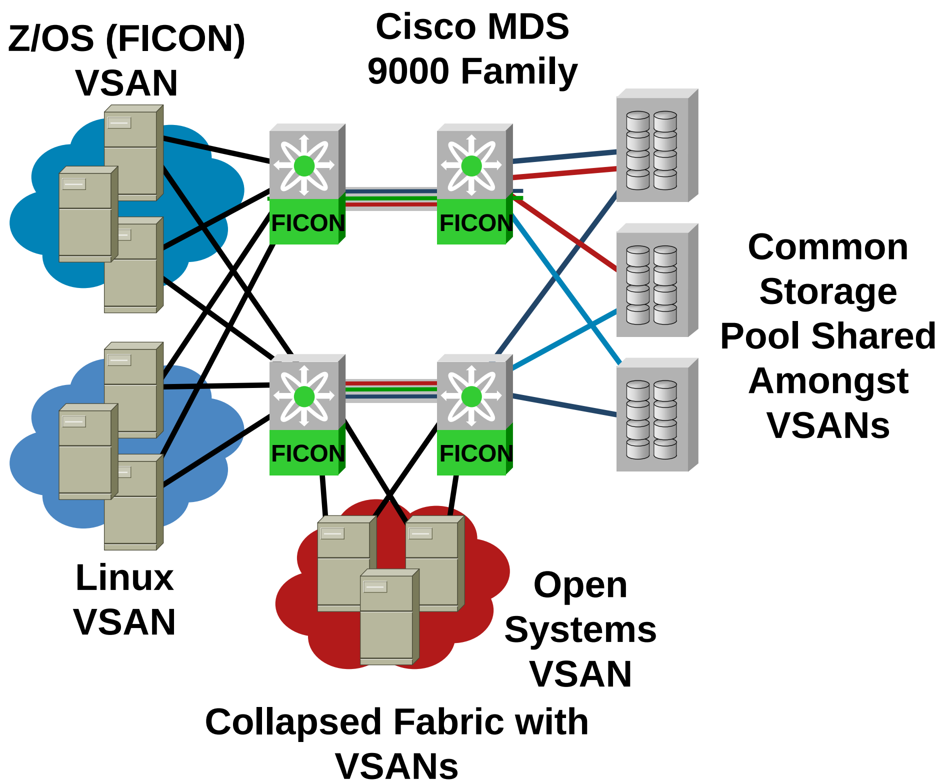

- Clean partitioning of different operating environments (FICON, Z-Series Linux-FCP, Open Systems FCP)
- Significantly more stable and manageable than current zoning+best practices approach

VSANs + Virtualization for Provisioning

- Optimize storage usage while supporting heterogeneous storage
- Virtual Targets with Virtual LUNs are built from discovered physical storage
- Virtual LUNs and targets can be zoned to destined host(s)
- Separate VSAN used to isolate physical storage
- Ability to virtualize across multiple vendors’ storage arrays
- Cisco working with several partners to deliver solutions

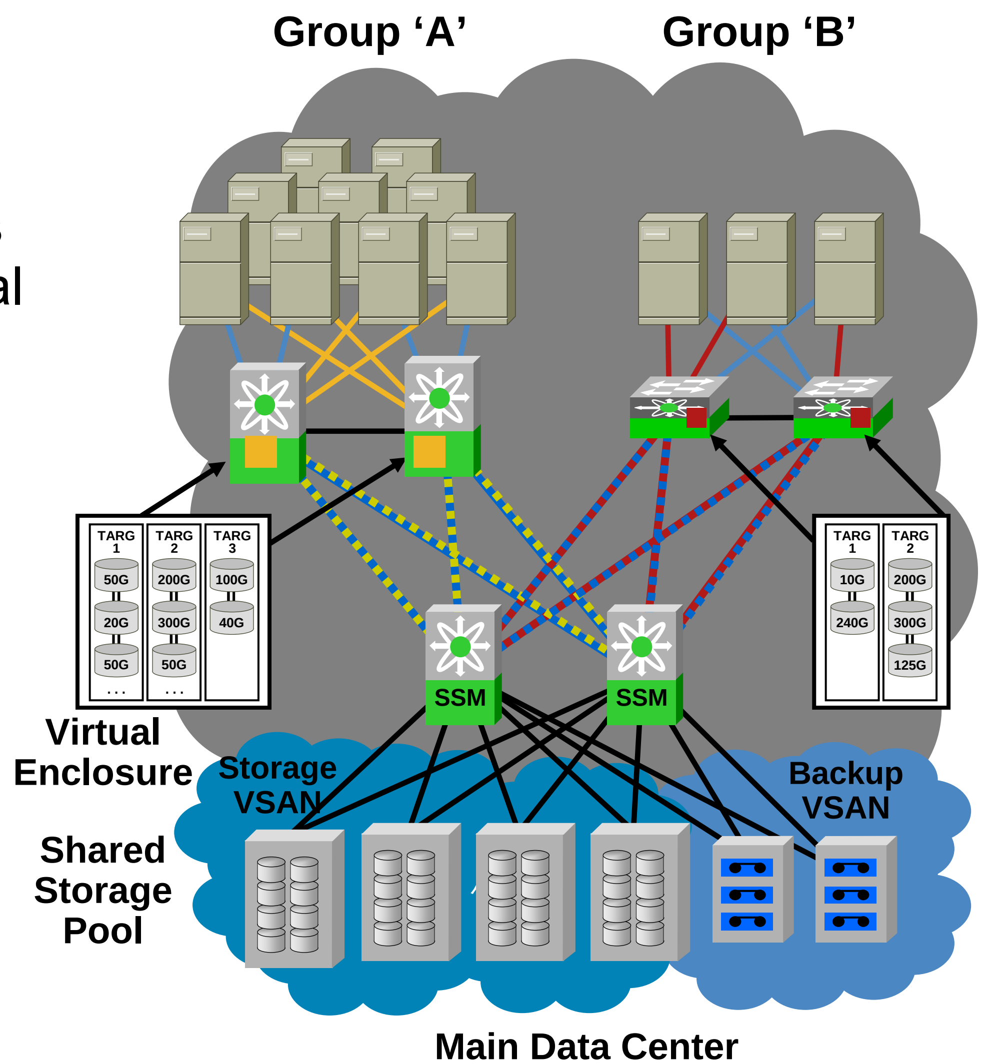

Differentiated Network Service Per VSAN

Virtual SANs enable resource allocation and preference per virtual fabric

- Each VSAN runs a separate instance of FSPF routing – independent forwarding decisions per VSAN
- EISL trunk links can be tuned for preferential routing per VSAN
- Different recovery paths can be configured per VSANs
- VSANs can be carried securely across metro and wide area networks via FCIP, SONET, DWDM, or CDWM
- Quality of Service (QoS) per VSAN to give preferential treatment at points of congestion in network

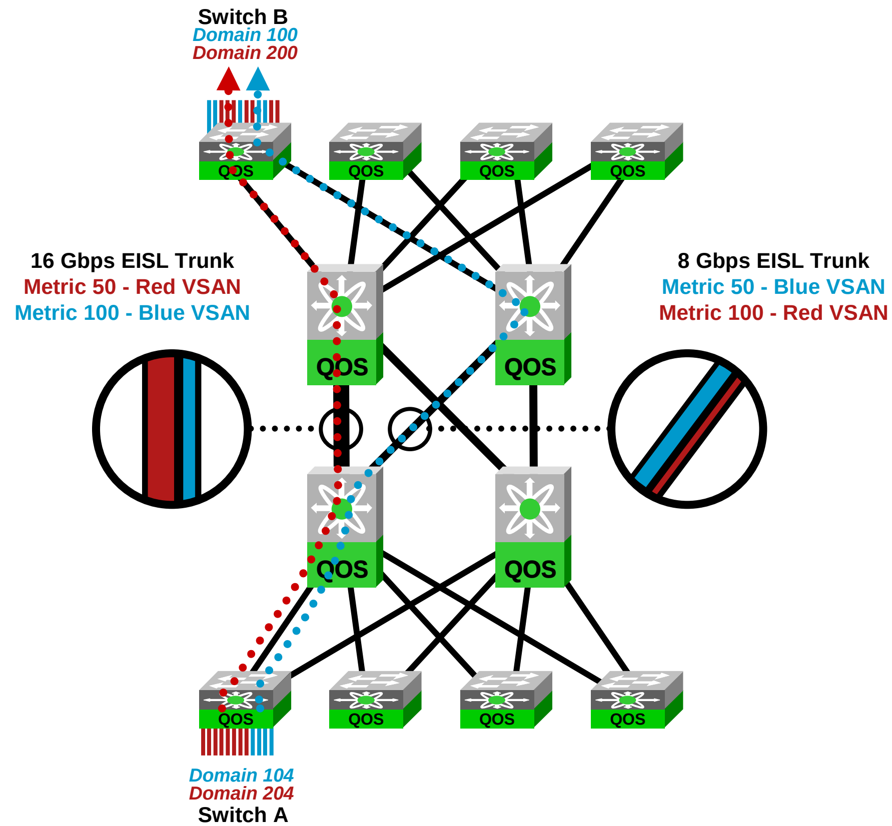

How do I get from switch A to switch B? It depends on which VSAN. Preferential routes can be configured per-VSAN to engineer traffic patterns.

How do VSANs work?

VSANs and Zones: Complimentary

Virtual SANs and fabric zoning are very complimentary

- Hierarchical relationship—
  - First assign physical ports to VSANs
  - Then configure independent zones per VSAN
- VSANs divide the physical infrastructure
- Zones provide added security and allow sharing of device ports
- VSANs provide traffic statistics
- VSANs only changed when ports needed per virtual fabric
- Zones can change frequently (e.g.. backup)
- Ports are added/removed non-disruptively to VSANs

Relationship of VSANs to Zones

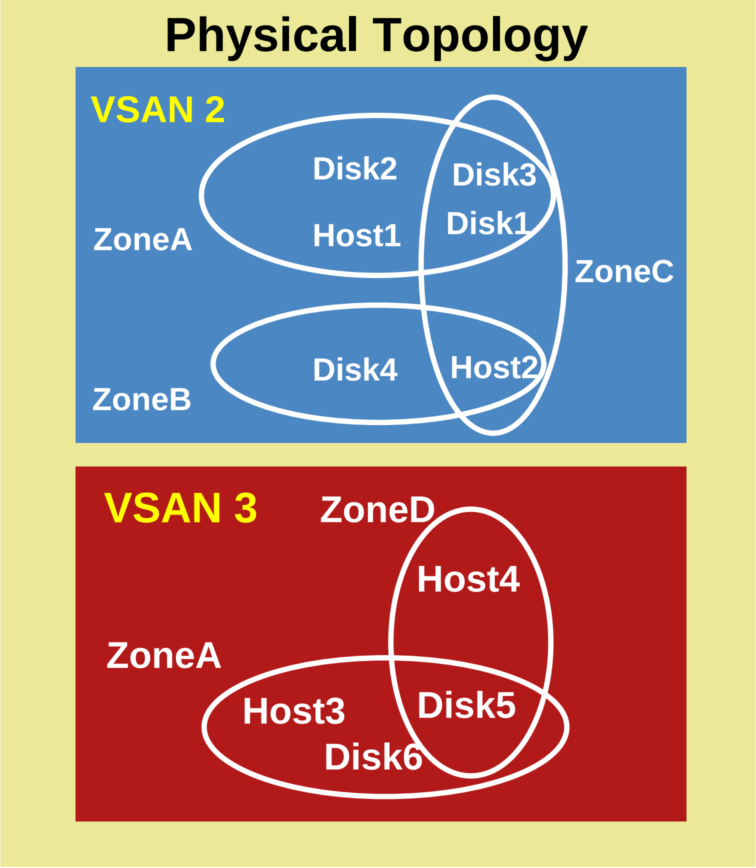

Two Primary Functions of VSANs

The Virtual SANs feature consists of two primary functions:

1. Hardware-based frame tagging of traffic belonging to different VSANs—
   Hardware Isolation

   No special drivers or configuration required for end nodes (hosts, disks, etc)

   Traffic tagged at Fx_Port ingress and carried across EISL (enhanced ISL) links between switches

2. Create independent partition of Fibre Channel services for each newly created VSAN—services include:

   Zone server, name server, management server, principle switch election, etc.

   Each service runs independently and is managed/configured independently

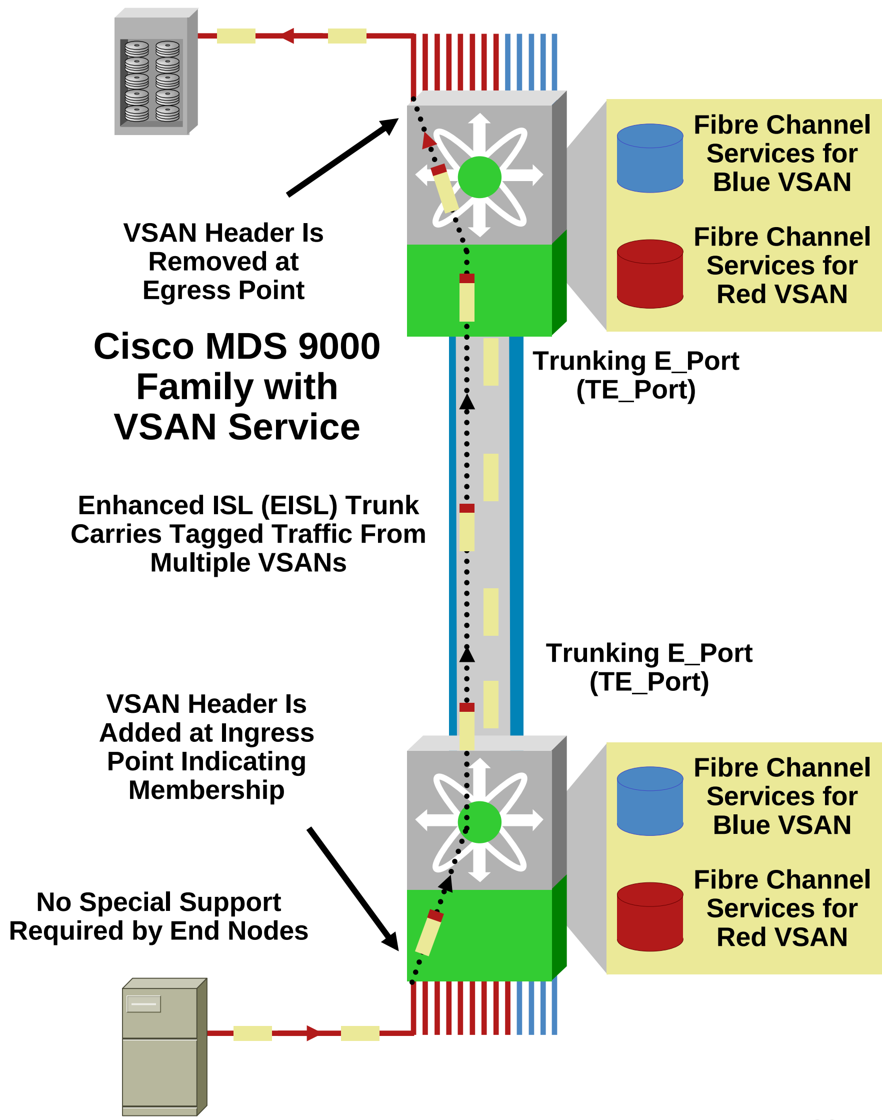

Virtual Fabric Trunking

EISLs and TE_Ports

The Virtual SANs feature introduces two new SAN elements:

1. The Trunking E_Port (TE_Port)
    - Negotiated between MDS 9000 switches—default
    - Carries tagged frames from multiple VSANs
    - Can be optionally disabled to yield E_Port
    - Only understood by Cisco MDS 9000 switches
    - Also has a native VSAN assignment (for E_Port)
    - Trunk all VSANs (1-4093) by default
    - Not to be confused with Brocade ISL aggregation (trunking)
2. The Enhanced ISL (EISL) link
    - The resultant link created by two connected TE_Ports
    - Superset of ISL functionality
    - Carry individual control protocol information per VSAN (e.g.. zoning updates)
    - Can be extended over distance (DWDM, FCIP, etc)

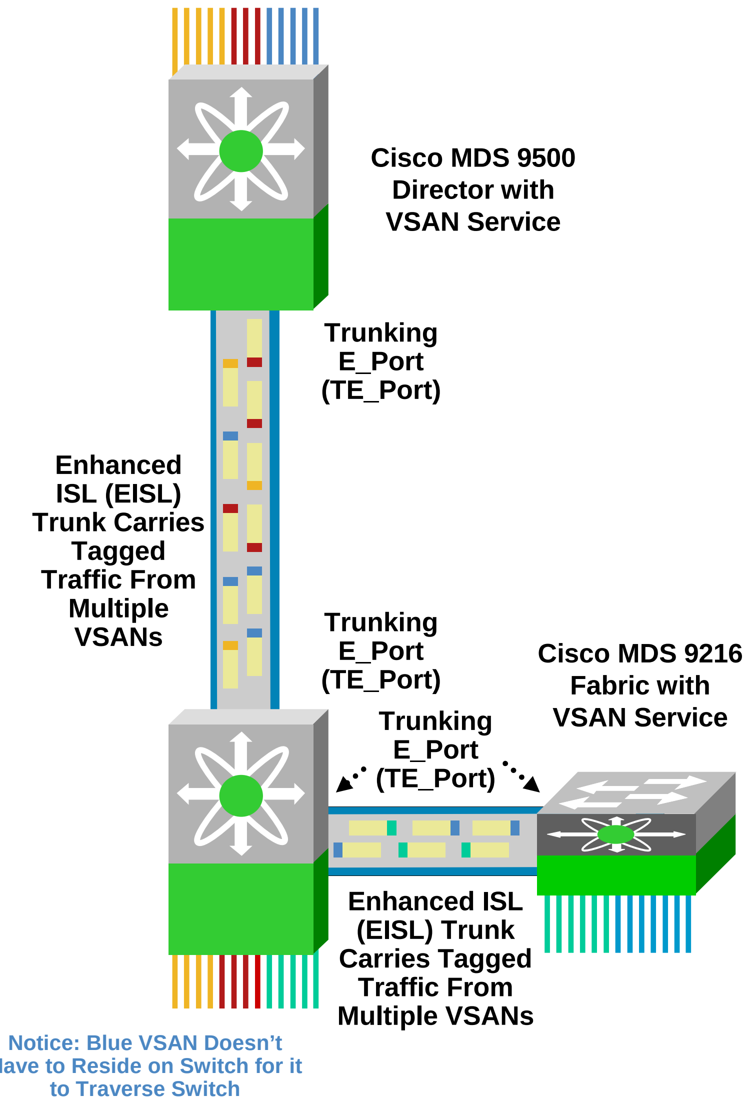

VSAN Number Space

VSAN Numbering Rules:

- VSAN 1 is the default VSAN

    All ports are originally in VSAN1

- VSAN 2 through 4093 can be assigned to 'user' VSANs—VSAN 0, 4094, 4095 are reserved

    A maximum of 256 VSANs can be created from the range of 2-4093

    Your mileage may vary based on other factors (e.g.. #zones, #zone sets, etc)

- VSAN 4094 is a reserved 'special' VSAN

    Called the 'isolated VSAN'

    Used to isolate ports who's port-VSAN has been deleted

    Not propagated across switches

    Always present, can't be deleted

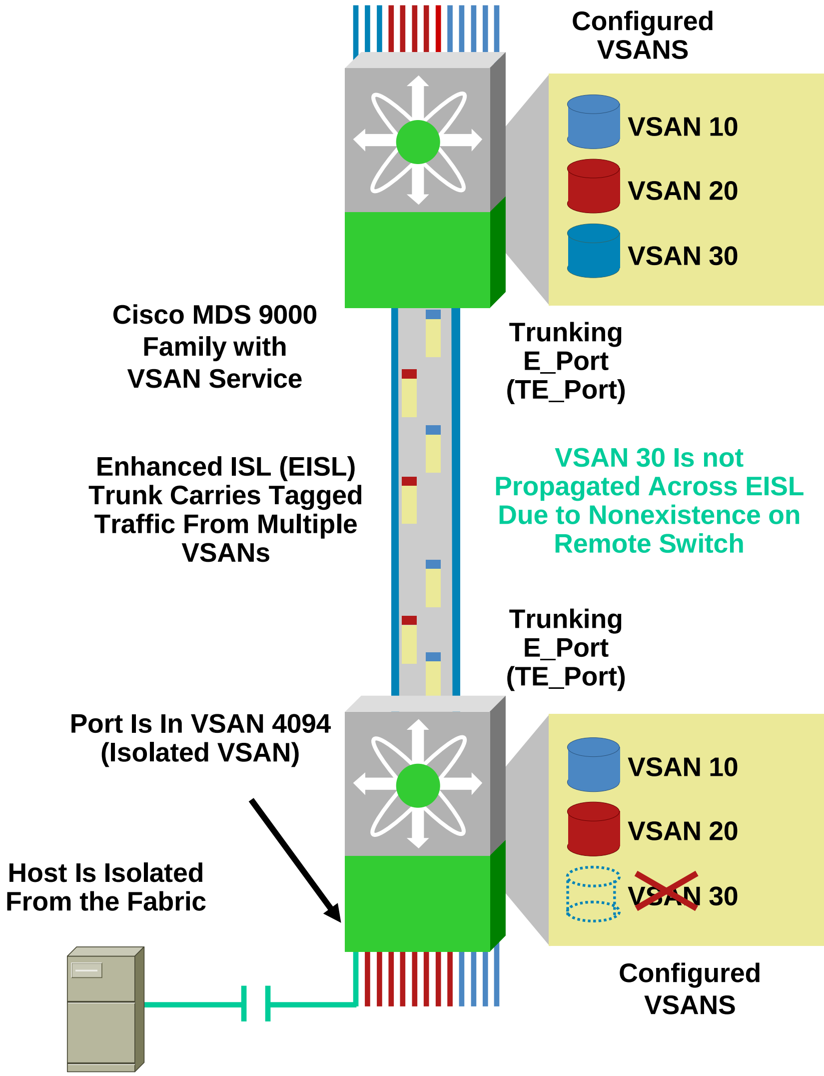

VSANs and Non-Cisco Switches

The Virtual SANs feature involves a tagging mechanism which is not understood by 3rd party fabrics

- Cisco MDS 9000 Family switches do support heterogeneous switch interoperability—non VSAN aware
- Cisco “Interoperability Mode” (required) is configured per-VSAN—no loss of functionality in MDS 9000 switches
- Cisco MDS 9000 switches negotiate an E_Port with non-Cisco switches
- Cisco MDS 9000 E_Ports also have a port VSAN
- Therefore, the entire non-Cisco switch, including all its ports, will reside in the port VSAN of the connecting E_Port

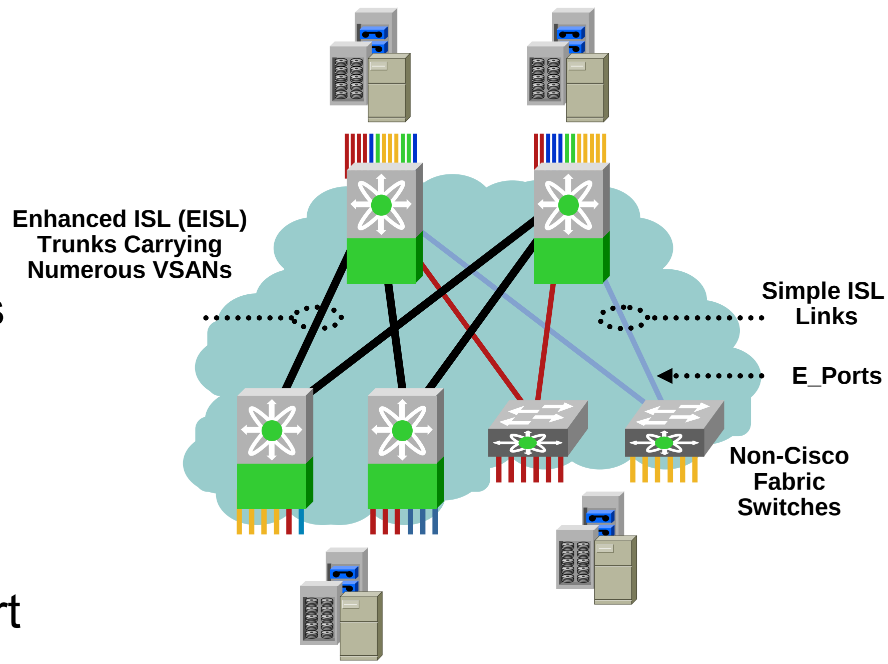

Each Non-Cisco Switch Belongs to Only One VSAN

Fabric Routing

So, What About Fabric Routing?

- We use fabric as an extension of virtual fabrics to enable cross-fabric connectivity
- Done without merging the routed fabrics
  - Without propagation of irrelevant fabric events
  - Without concern for overlapping domain IDs
  - Without concern for fabric interoperability differences
- Follows in footsteps of the Ethernet world
  - Layer-3 Switching ≈ Fabric Routing

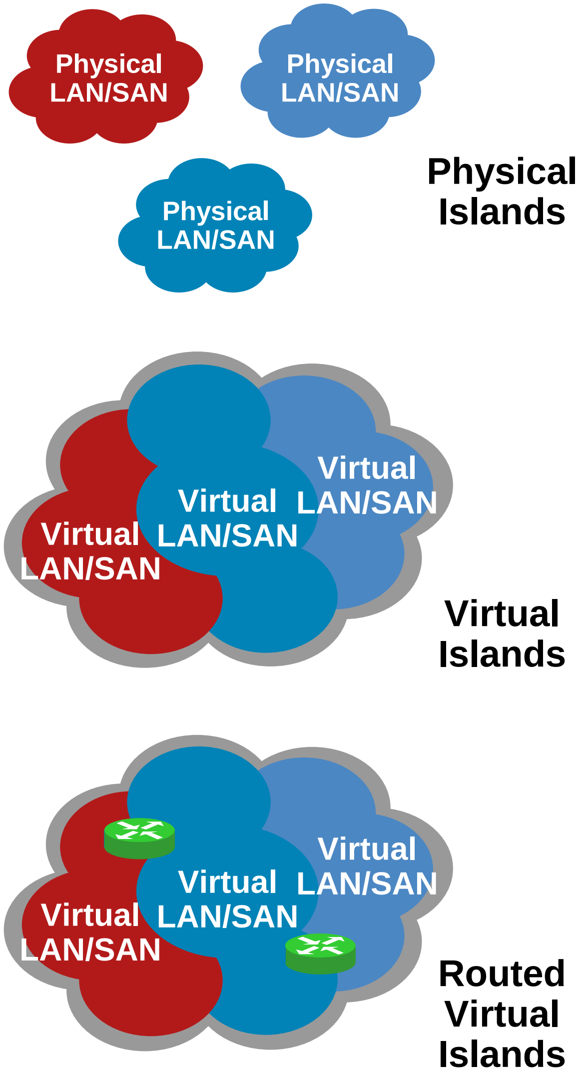

Uses For Fabric Routing
Securely Sharing Common Resource

- Overlay data replication fabric(s) on common physical fabric
  - No need for separate pair of switches for each replication connection
  - Use one virtual fabric per replication connection
- A *bonus* is to be able to share common SAN extension circuits amongst multiple virtual fabrics
- Fabric routing adds to resiliency of solution

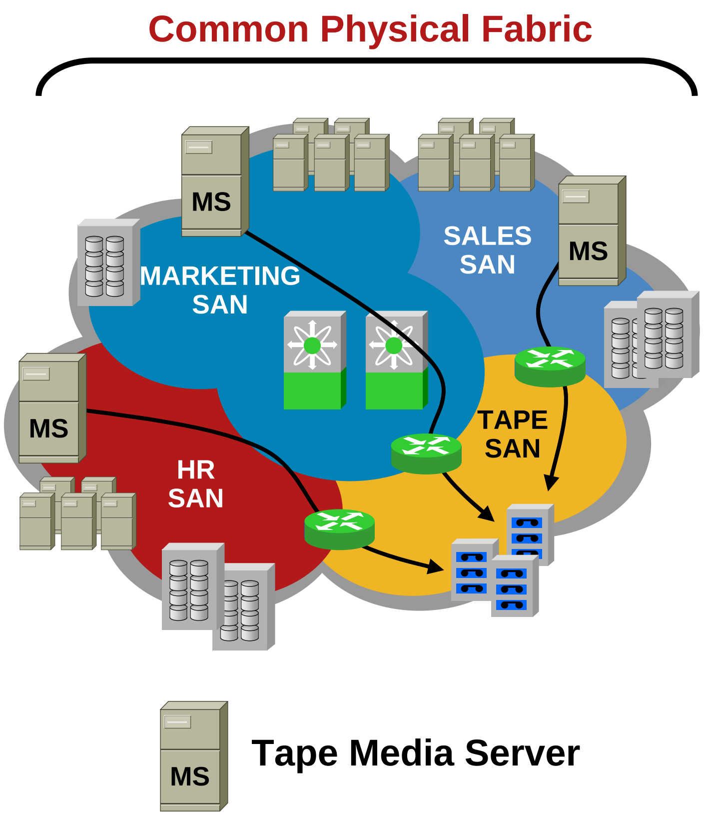

Two Main Approaches to Fabric Routing

External Router

- Dedicated fabric router connected to all fabrics
- Not typically director class—HA concerns
- Performance limited by that of appliance

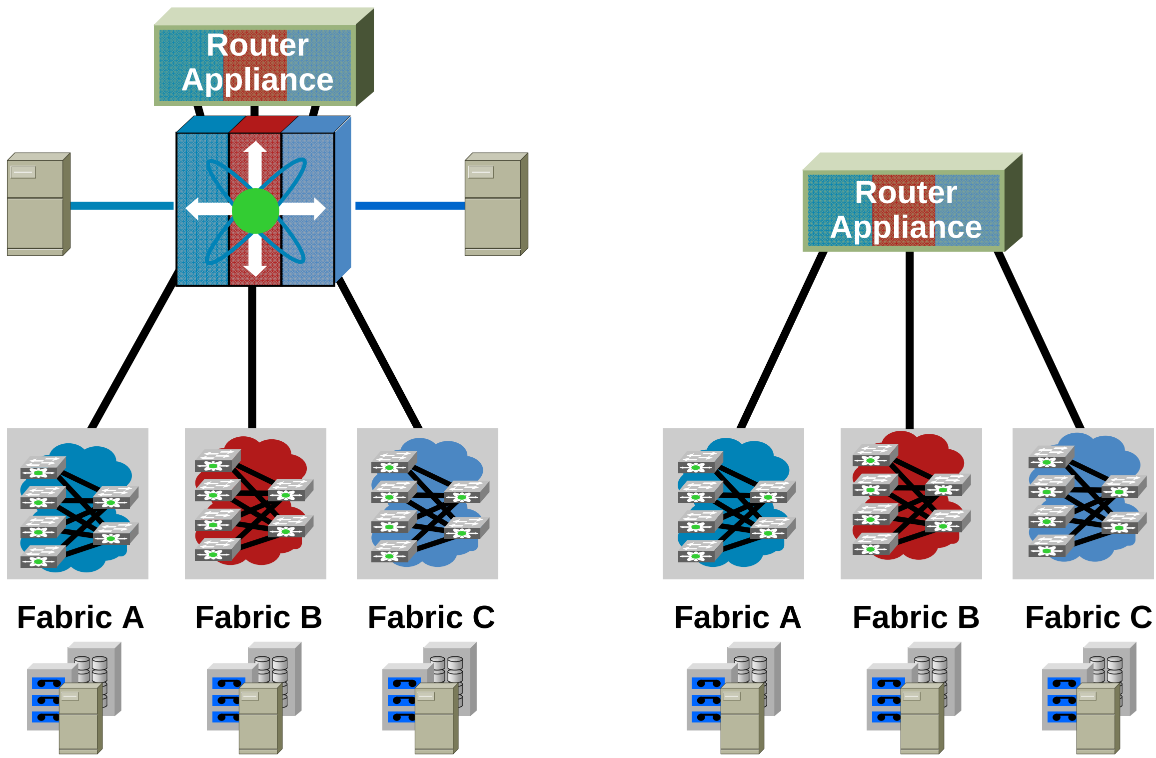

Embedded Routing

- Routing enabled in switch/director hardware
- No performance penalty
- Port-level granularity

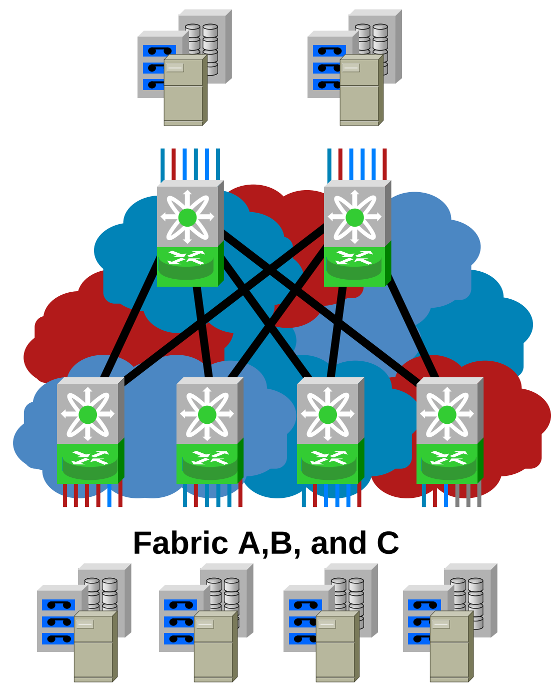

Cisco's Approach to Fabric Routing
Inter-VSAN Routing (IVR)

- Cisco delivers fabric routing through Inter-VSAN routing (IVR)
- Builds on VSAN technology
- Embedded capability in all MDS 9000 Family switch hardware
- No need for external router
- No performance impact due to implementation in hardware
- Leverages any network transport
    - Fibre Channel
    - Optical (DWDM, CWDM, SONET)
    - IP (FCIP)
- IVR includes NAT services

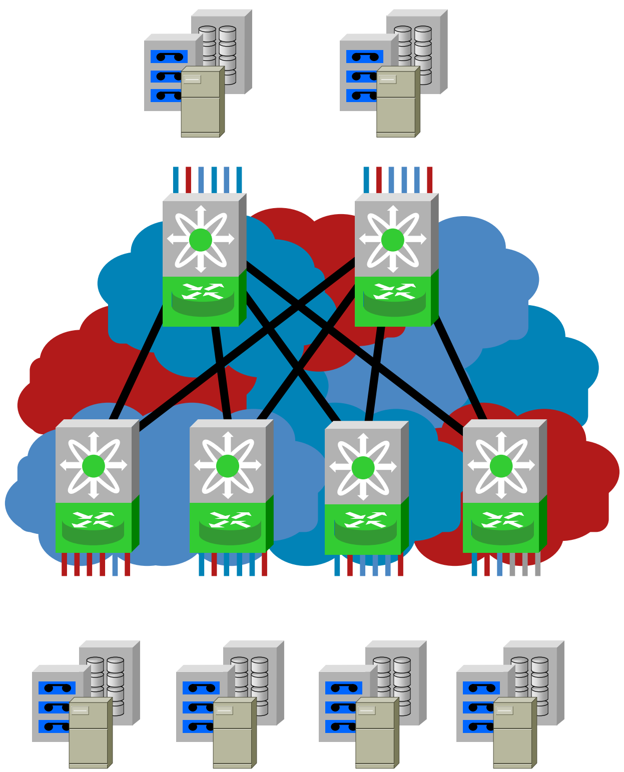

IVR Operation Within a Single Switch

- Effectively turns any MDS 9000 Family switch into giant fabric router
- IVR enabled in any Cisco MDS 9000 Family switch using a license key
- Works with all fabric interoperability modes
- Enabled through zone creation mechanism

ANY CISCO MDS 9000 FAMILY SWITCH

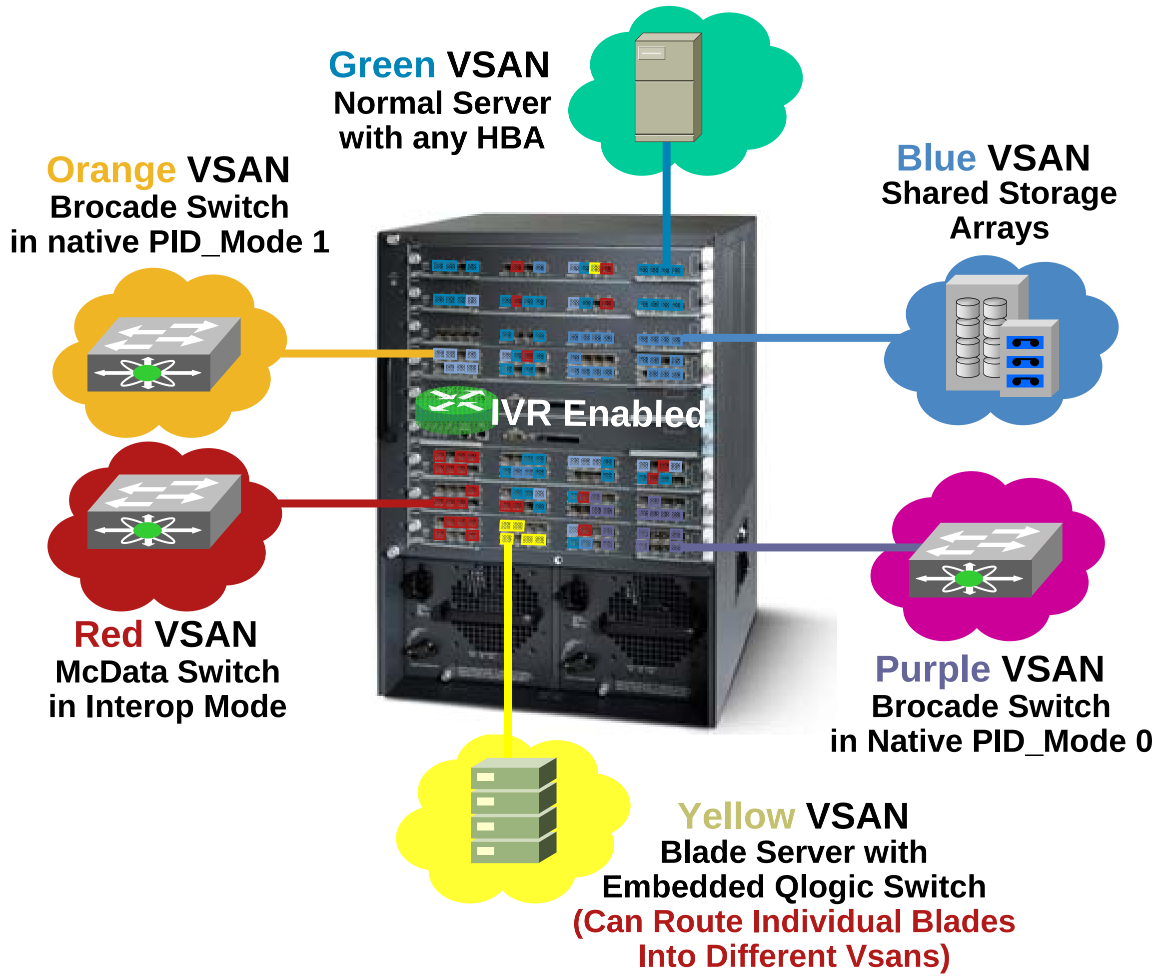

Inter-VSAN Routing (IVR) : Sharing Resources Across VSANs

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

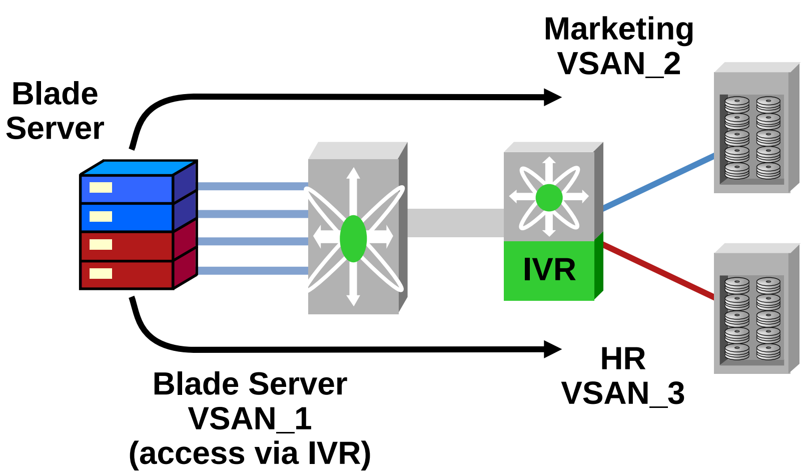

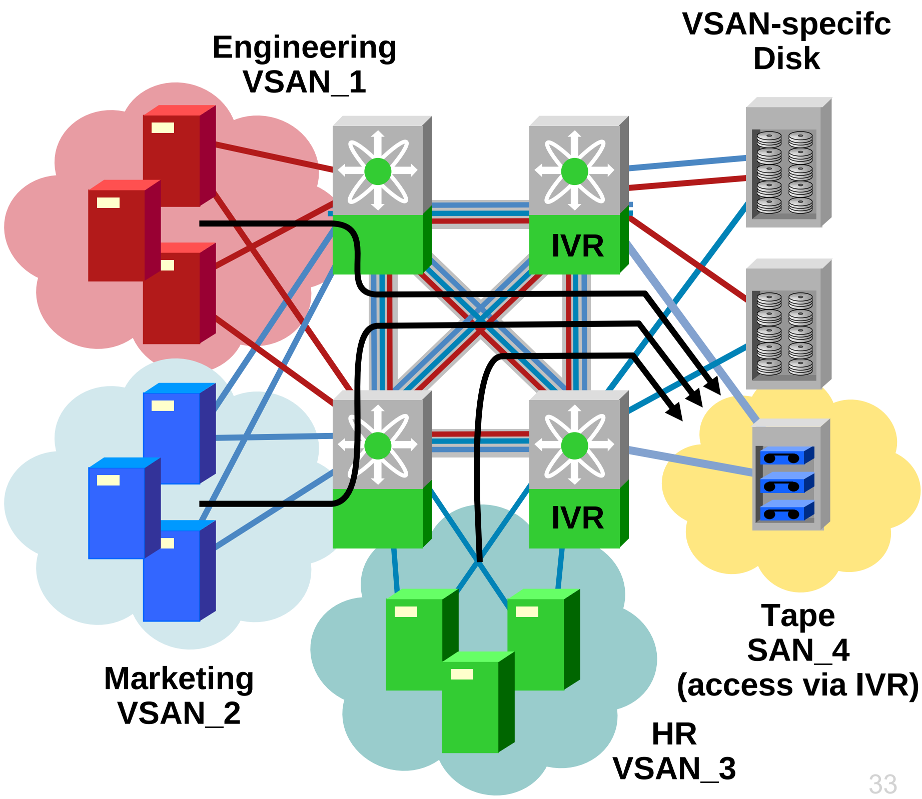

Inter-VSAN Routing (IVR): Resilient SAN Extension Solutions

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

Inter-VSAN Connection with Completely Isolated Fabrics

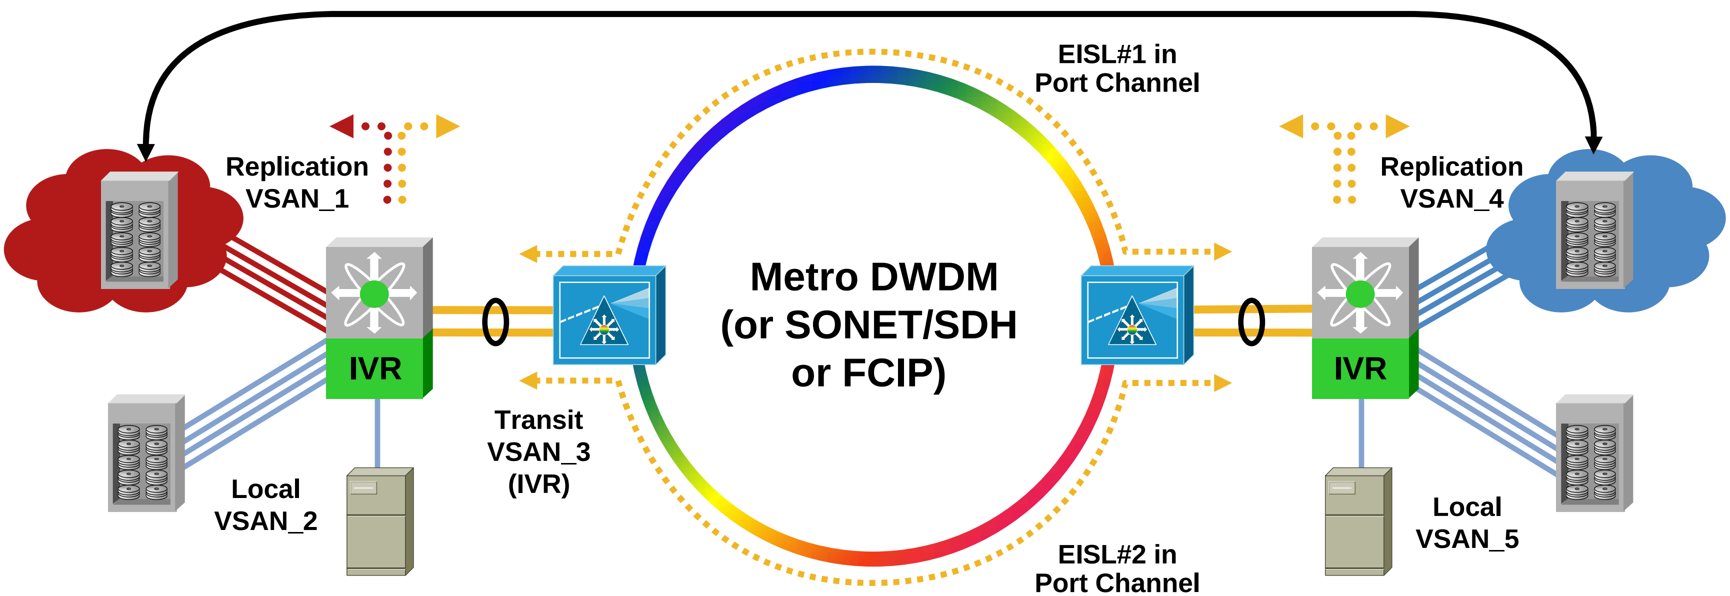

Part 2

Storage Virtualization

Storage Services Module (SSM):
Open Platform for Intelligent Fabric Applications

MDS 9000 Storage Services Module

- ASIC-based innovation
- Open, standards-based platform
- Hosts multiple partner applications
- 9 processors (8 DPP + 1CPP)
- Nine Virtual Initiators
  (one per DPP and one CPP)
- Number of Virtual Targets creation depends on Fabric Application

MDS 9000 Storage Services Module

Network-Hosted	Network-Assisted	Network-Accelerated
FAIS-Based API (T11)	SANTap Protocol	Standard FC Protocols
Volume Mgmt, Data Migration, Copy Services	Async. Replication, CDP	Serverless Backup, FC Write Acceleration
EMC²	Topio, Kashya, FalconStor Software, XIOtech, Alacritus Software, Cloverleaf Communications	Computer Associates, EMC² legato, CommVault® Systems Unified Data Management

FAIS:
Standards-Based Interface for Intelligent Fabric Applications

FAIS: Fabric Application Interface Specification

- ANSI T11 standards-based effort to create a common application programming interface (API) for fabric applications to run on an underlying hardware platform
- The object is to help developers move storage and data management applications off applications, hosts, and storage arrays and onto intelligent storage fabric-based platforms

Network Storage Applications

Potential Network Functions

- Heterogeneous volume management
- Data migration
- Heterogeneous replication/copy services
- Continuous Data Protection (CDP)

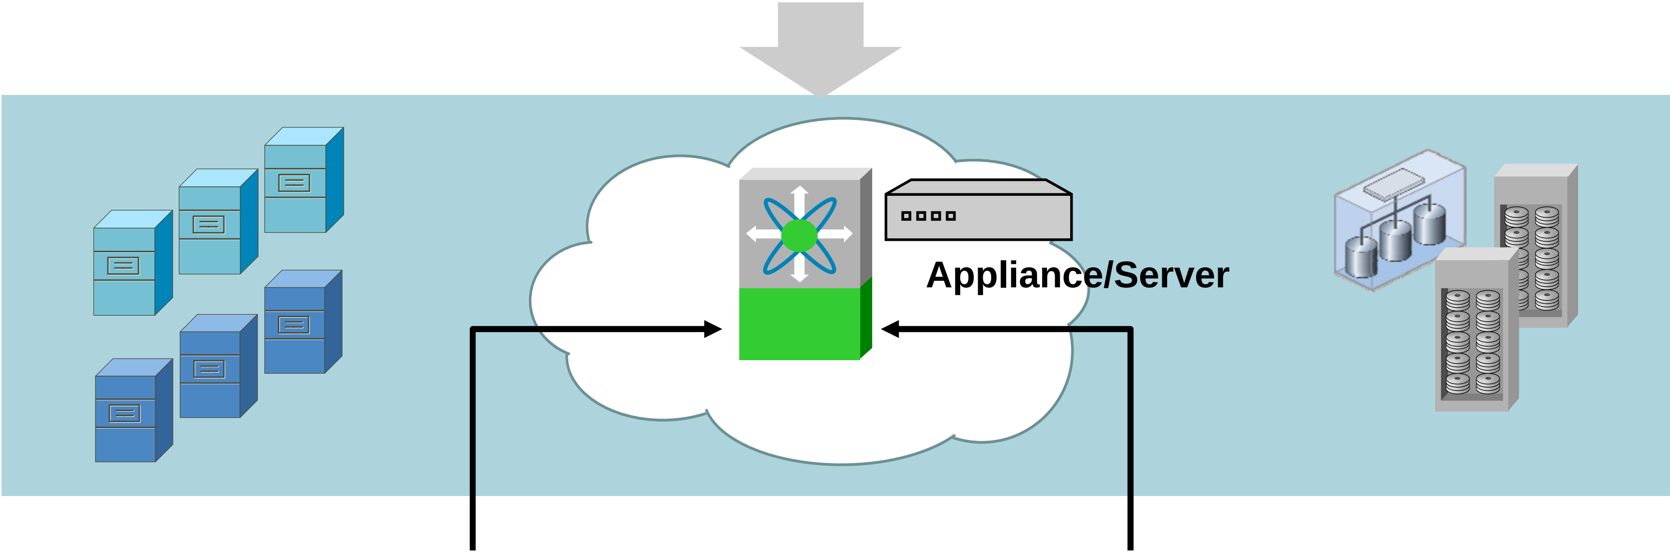

Network-Hosted

- Partner software resides on MDS

Network-Assisted

- Partner software resides on arrays, external server or appliance

FAIS-Based Storage Applications

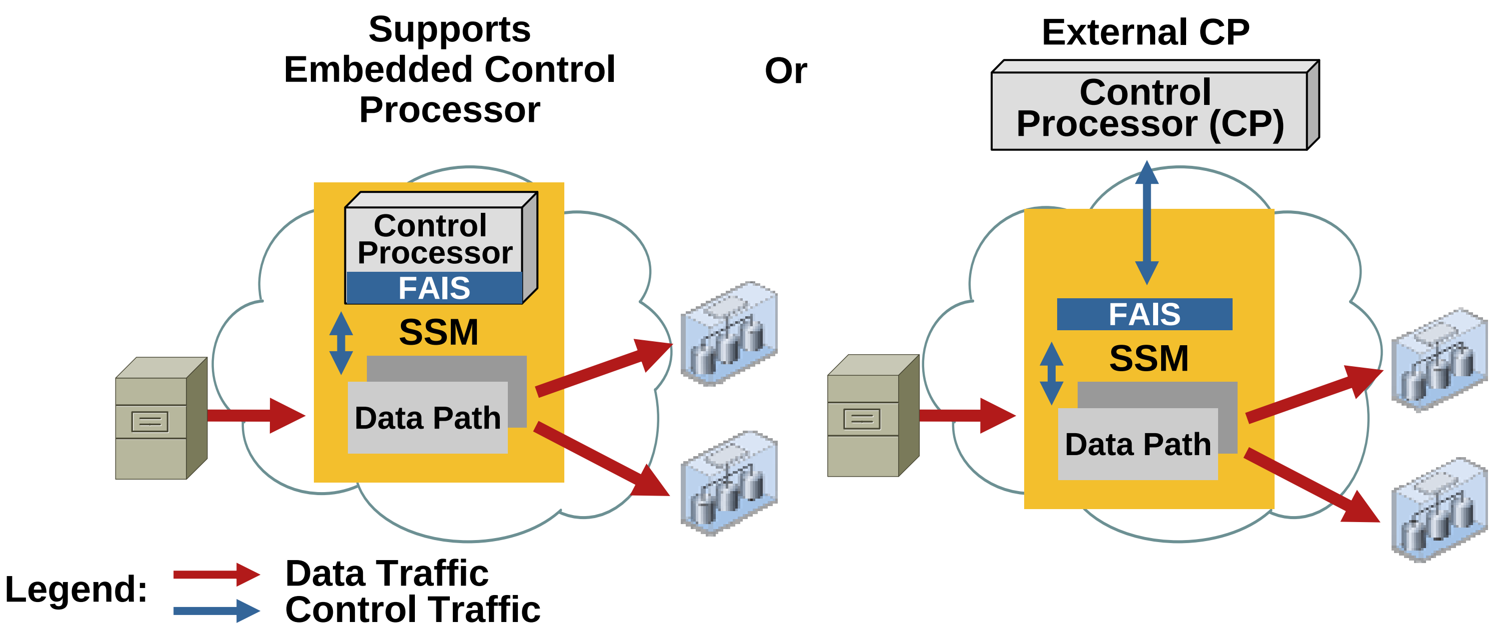

- FAIS standards-based (T11) open-API
- Enables partners to create switch independent storage applications
- Enables Cisco MDS to host partner storage applications
- Independent control and data paths
- Integrated HA architecture
- Fully distributed intelligence
- CP Programs data path and processes exceptions

Network Accelerated Serverless Backup (NASB)

The Evolution of Backup Architectures

LAN-Based	LAN-Free	Server-Free	Fabric-Assisted
- Data moved over LAN	+ Data moved over SAN	+ Data moved over SAN	+ Data moved over SAN
- App server moves data	- App server moves data	+ App server not in data path	+ App server not in data path
		- Dedicated Server moves data	+ Fabric moves data

Serverless Backup

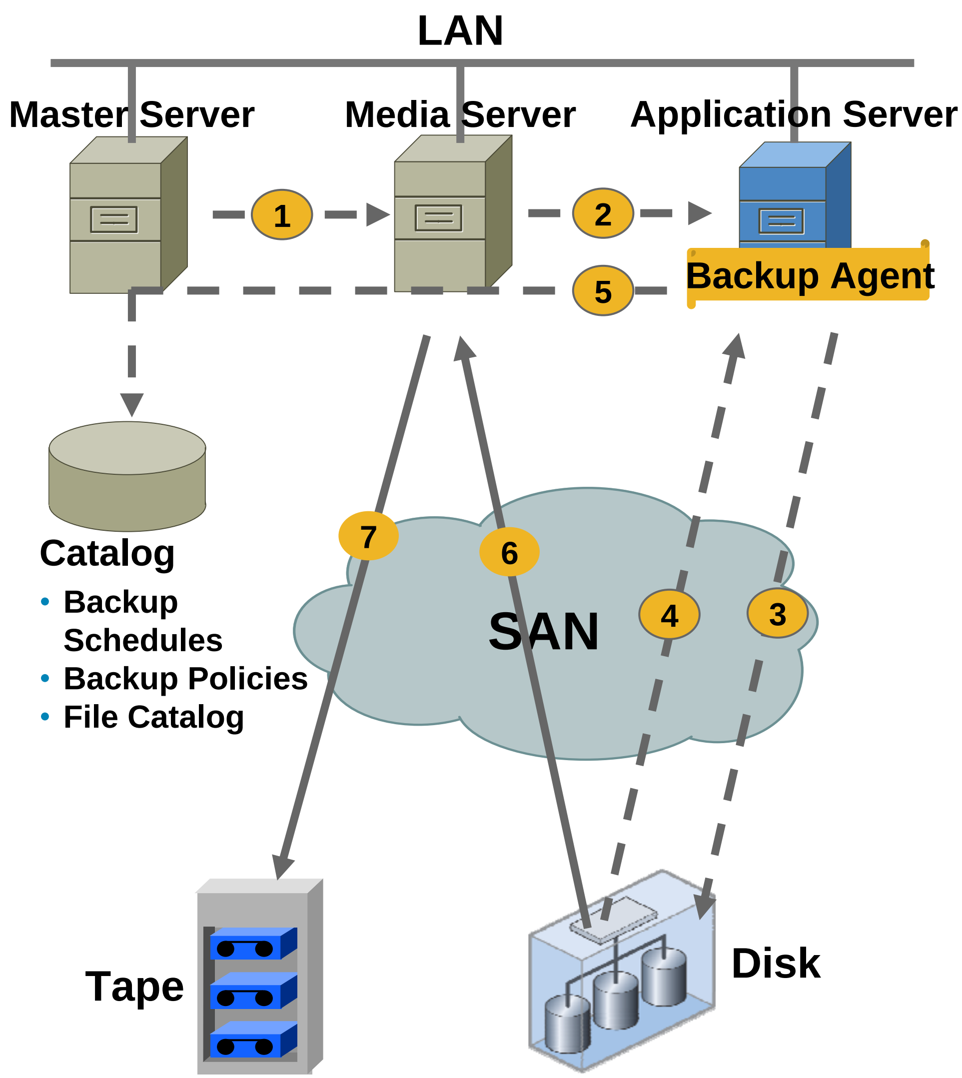

1. Based on policy stored in catalog, the backup process is started
2. Backup node requests the Backup Agent to initiate backup
3. Backup Agent quiesces the app and generates a snapshot
4. Backup Agent creates meta data (extent list) of file/volume mapping to blocks
5. Backup Agent sends meta data to backup node. The backup node sends it to Master Server where it gets written to Catalog.
6. Based on the extent list, backup node reads data blocks from disk
7. Backup node writes these data blocks to tape

Network-Accelerated Backup

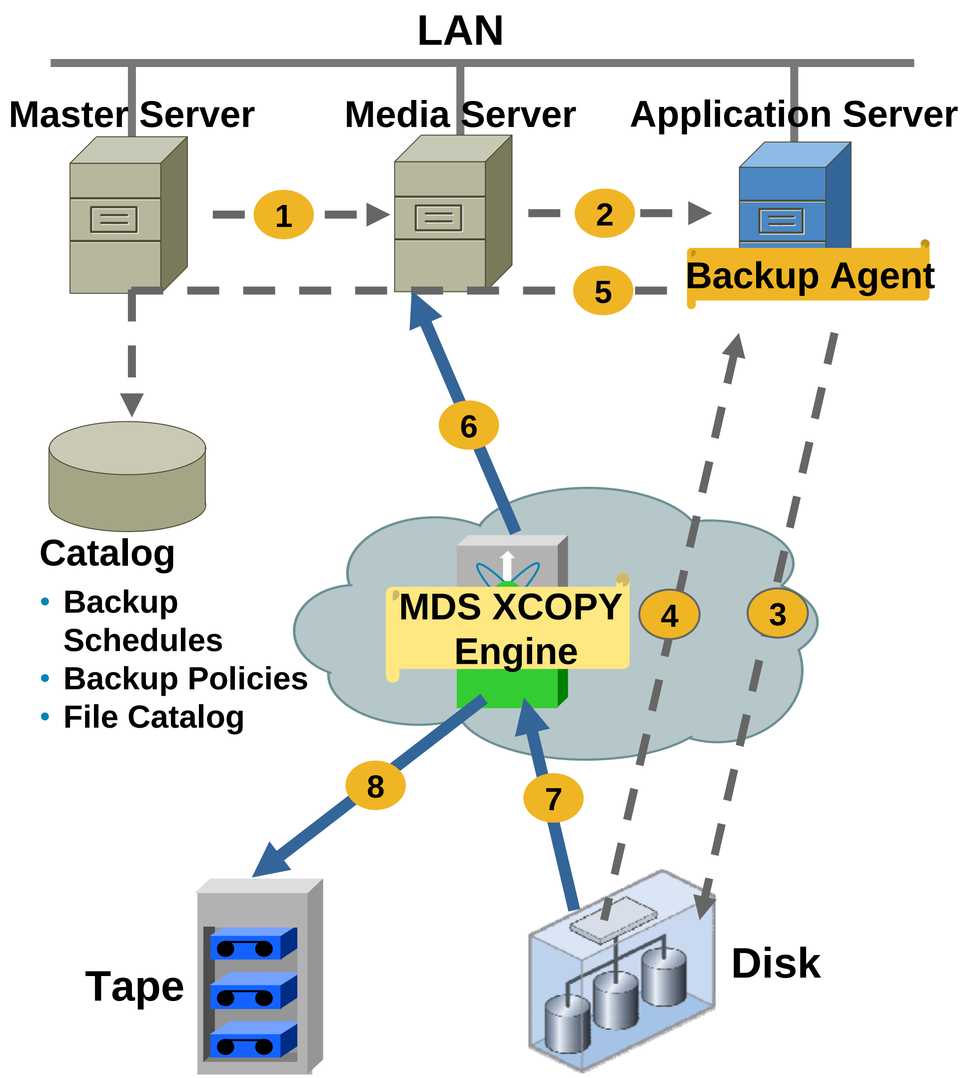

1. Based on policy stored in catalog, the backup process is started
2. Media Server requests the Backup Agent to initiate backup
3. Backup Agent quiesces the app and generates a snapshot
4. Backup Agent creates meta data (extent list) of file/volume mapping to blocks
5. Backup Agent sends meta data to Media Serve. The Media Server sends it to Master Server where it gets written to Catalog
6. Based on the extent list, Media Server sends SCSI XCOPY command to MDS XCOPY Engine
7. MDS XCOPY Engine reads data blocks from disk
8. MDS XCOPY Engine writes these data blocks to tape

Fibre Channel Write-Acceleration (FC-WA)

FC SAN Extension Today

FC-Based Replication: How Far?

- Performance of DR/BC applications inhibited by distance
- Latency degrades with greater distance
  - Databases are very sensitive to latency
- Only write I/Os are affected
  - Increased “Service Time” (response)

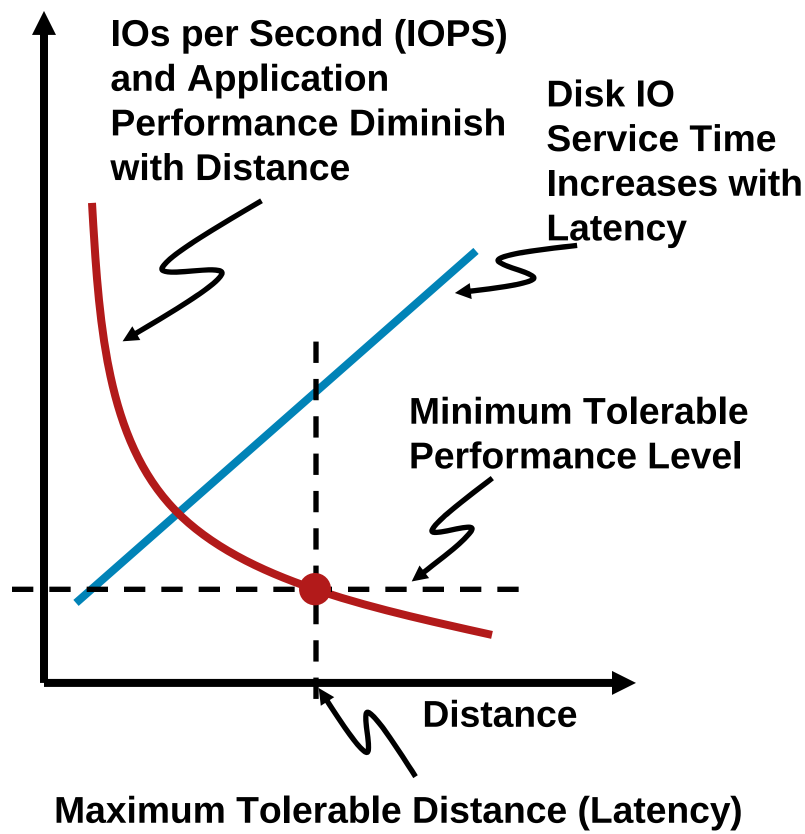

Fibre Channel Write Acceleration (FC-WA)

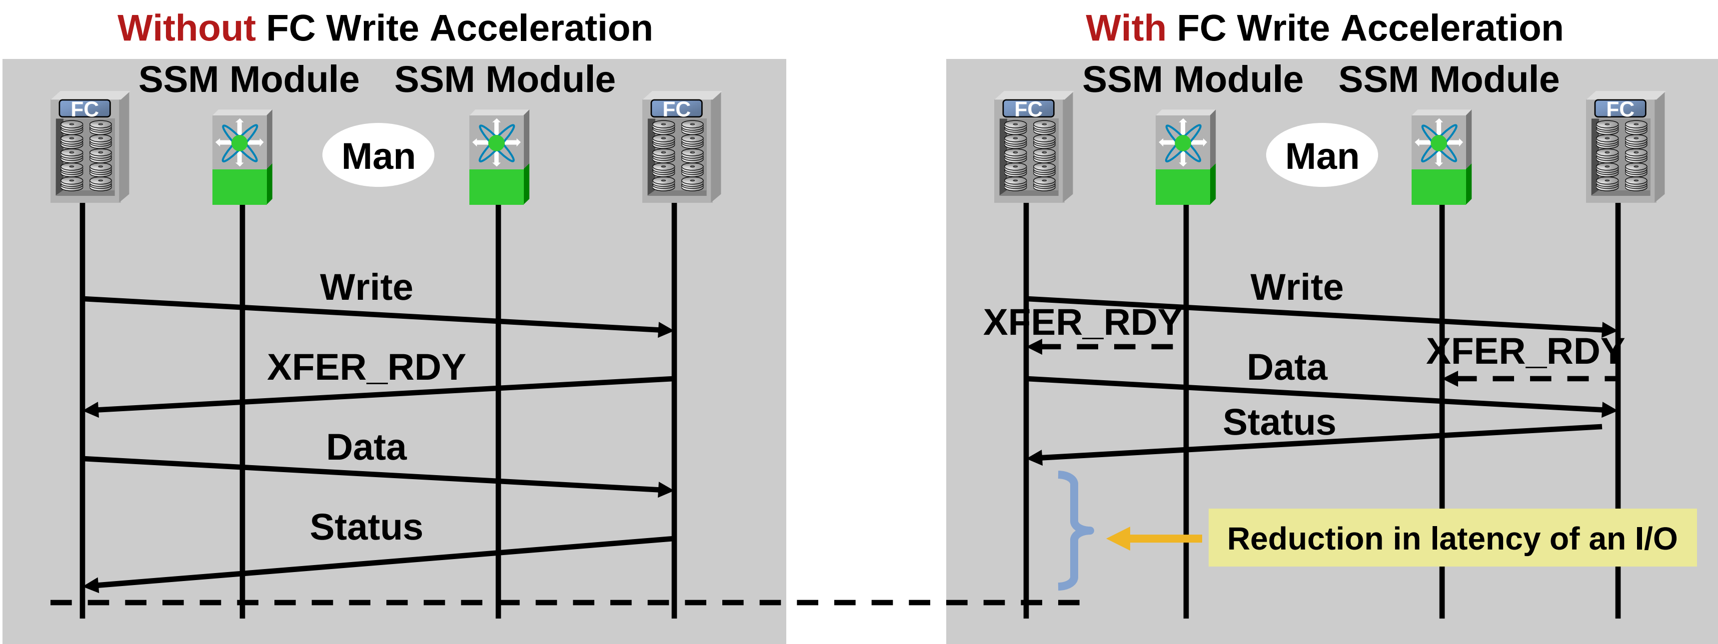

- Benefits of FC Write Acceleration
  - Improves response time for storage applications
  - Extended distance for DR and BC apps
- Requirements for FC Write Acceleration
  - Requires an SSM Module
  - Both initiator and target must be directly attached to the SSM Module

FC-WA Solutions Benefits

- Solution

  Optimize bandwidth for DR

  Increase distance between primary site and remote site

  Minimizes application latency

  Investment protection: transport agnostic (DWDM, CWDM, SONET/SDH,
  dark fiber)

- Primary applications

  Synchronous replication

Extend Distances for DR/BC Applications

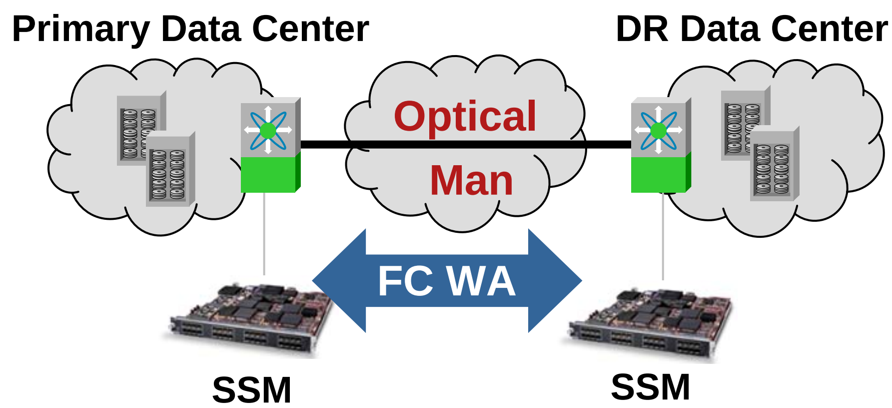

Up to 30%

Performance Improvement Seen by Major Financial Services Company over 125 km Distance

SANTap

Network-Assisted Storage Applications with MDS SANTap

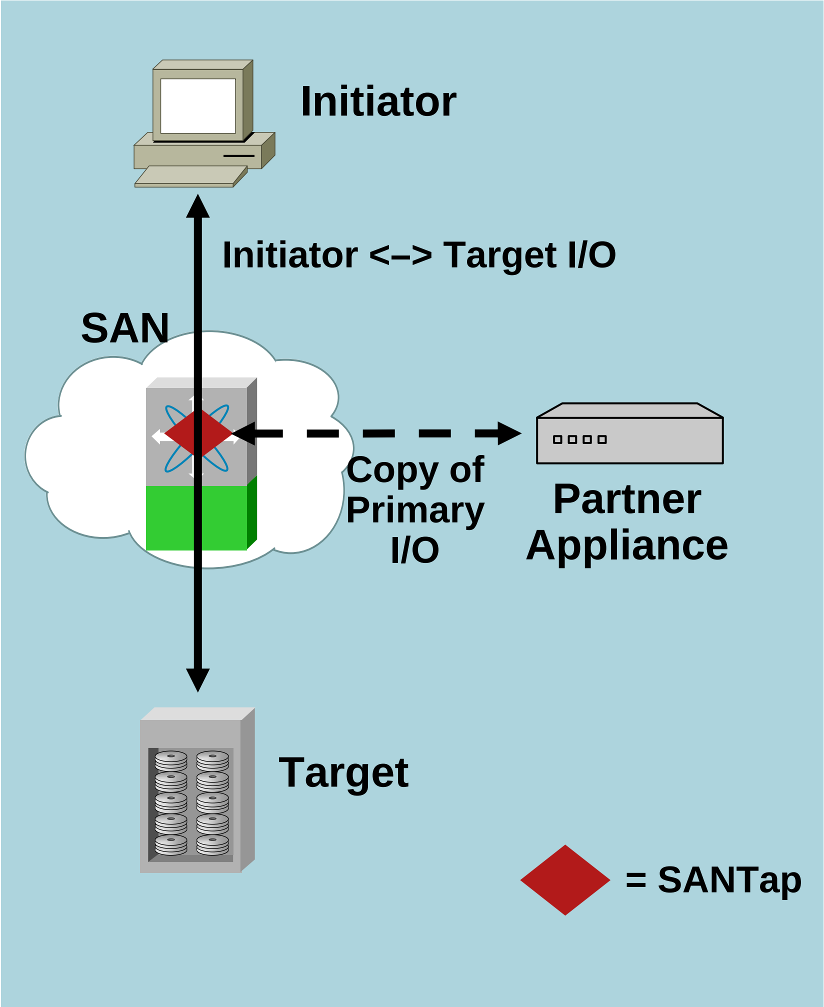

- Enables appliance-based storage applications without compromising SAN integrity
- About SANTap
  - MDS delivers a copy of primary I/O to an appliance
  - Appliance provides the storage application
  - Examples of applications include Continuous Data Protection (CDP), Replication, etc.
- Key customer benefits
  - Preserve integrity and availability of primary I/O
  - No service disruption if appliance fails
  - Investment protection
  - No performance limitations

SANTap-Based Continuous Data Protection (CDP)

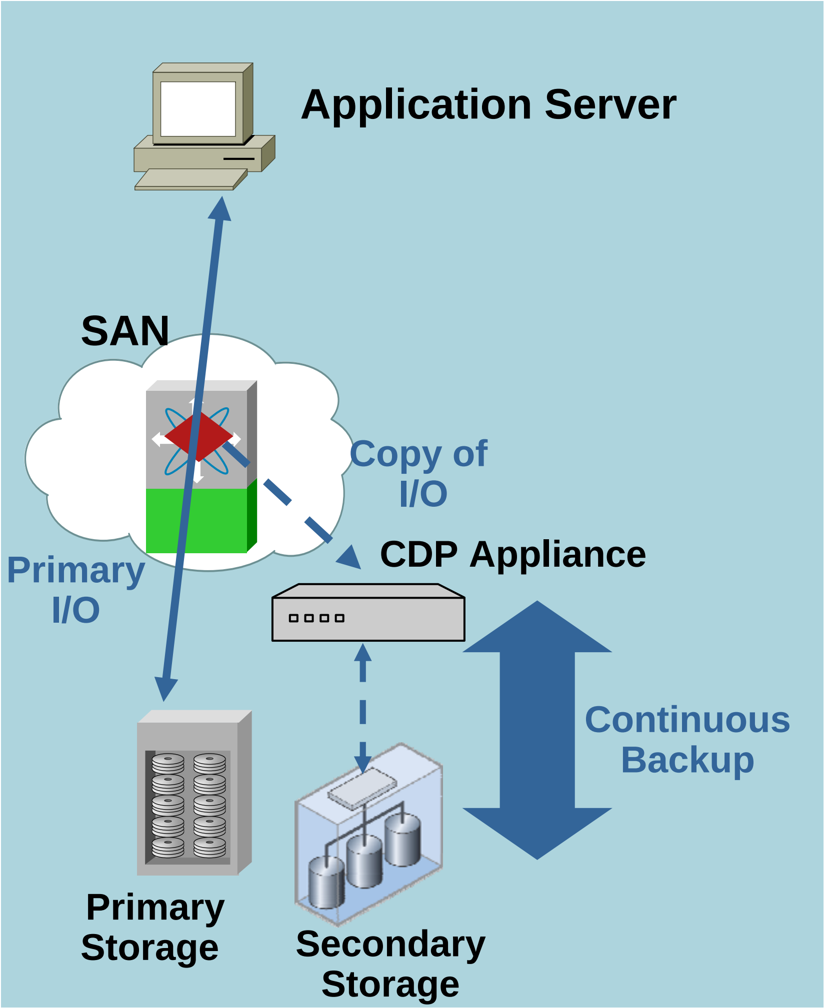

- Preserves the integrity and availability of primary I/O
- Performance of disk, economics of tape
- Zero-backup window
- Instant restore to any point in time
- Tiered storage array for continuous backup
- Policy-based data retention and deletion

Storage Virtualization

Virtualization Implementations

- Host base solution
  - Heterogeneous storage
  - Dynamic volume growth
  - CPU / Operating System bound
  - Multiple points of management
- Appliance solution
  - Heterogeneous storage
  - Limited scalability
  - Performance is appliance bound
- Network/Switch base solution
  - Centralized management
  - Highly scalable
  - Performance scales with network requirements

Alliance Solutions

EMC Invista Features:

- Network volume management
- Heterogeneous data mobility
- Heterogeneous cloning services

Open Platform

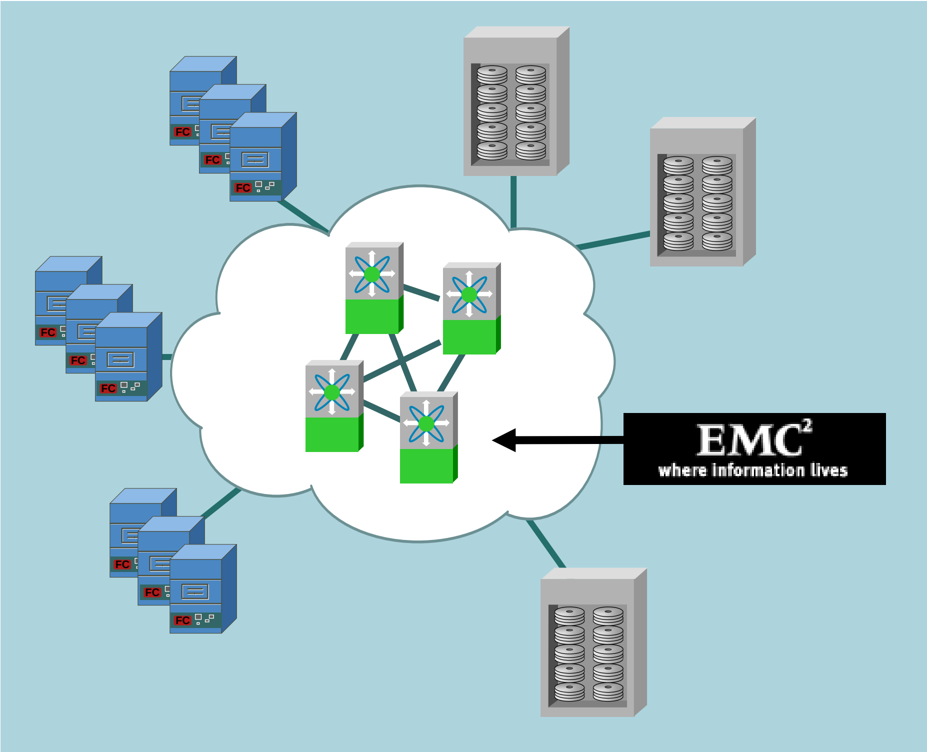

EMC: Invista Logical Communication

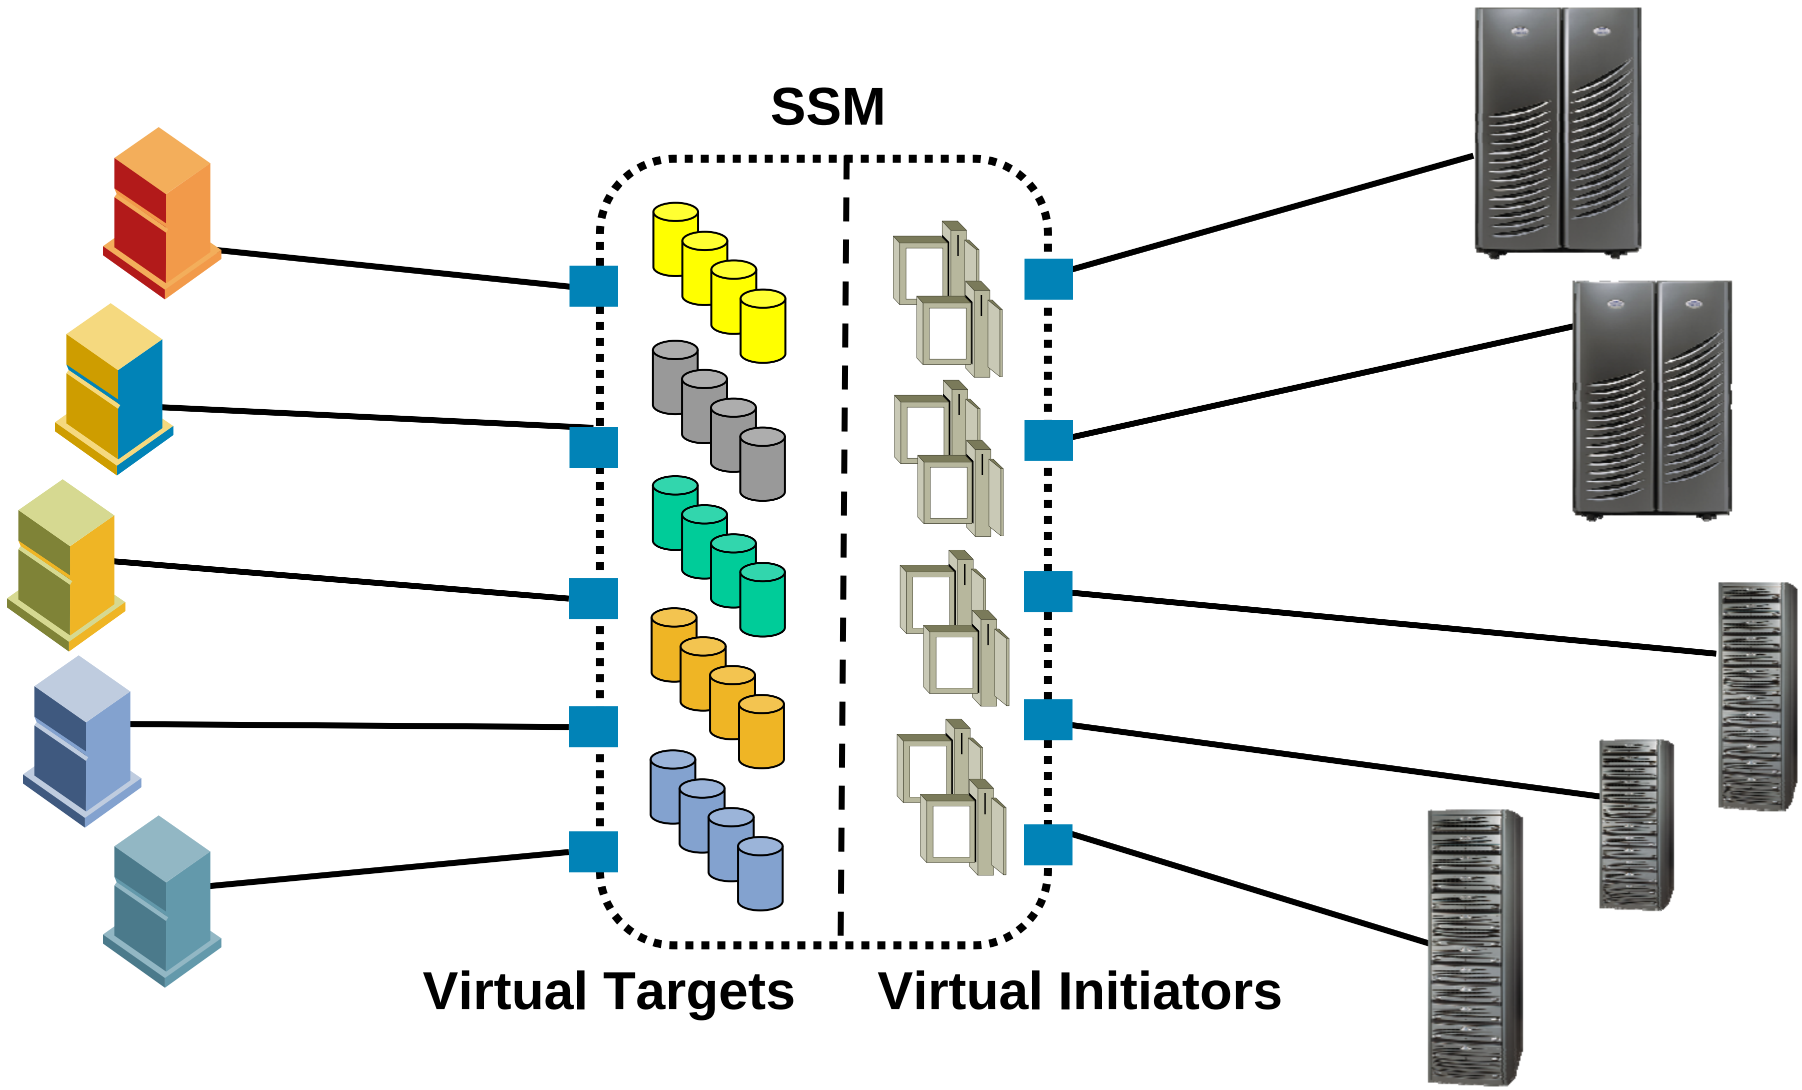

EMC Invista Instance—Hardware Components

Hardware

- Minimum of two SSM in dual Fabric (no ISL)
- Can scale up to two SSM per physical fabric (total of four SSM per Invista Instance)
    - Support up to four MDS 9200 w/SSM in each chassis
    - Support two MDS 9500 w/up to two SSM per chassis
- External Dual Control Processor provided by EMC called “Control Path Cluster” (CPC)
- Two Ethernet routers for communication between CPC and MDS switches (HA)
- Other MDS switches can connect to MDS switches that have SSM/Invista enabled for scalability

A Invista Cabinet

B Metadata Storage
- Three-way mirror
- Contains configuration information and logs

C CPC
- High-availability architecture
- Runs Invista application

D IP Routers
- Internal network to connect CPC to virtualizers

E DPC/Virtualizers
- 1, 2, or 4 virtualizers (MDS SSM) deployed in MDS 9000 chassis

Each configuration includes a standby power supply

EMC: Invista Terminology

- Front-end
- Virtual frames
- Virtual targets
- Virtual volumes
- Storage elements
- Back-end

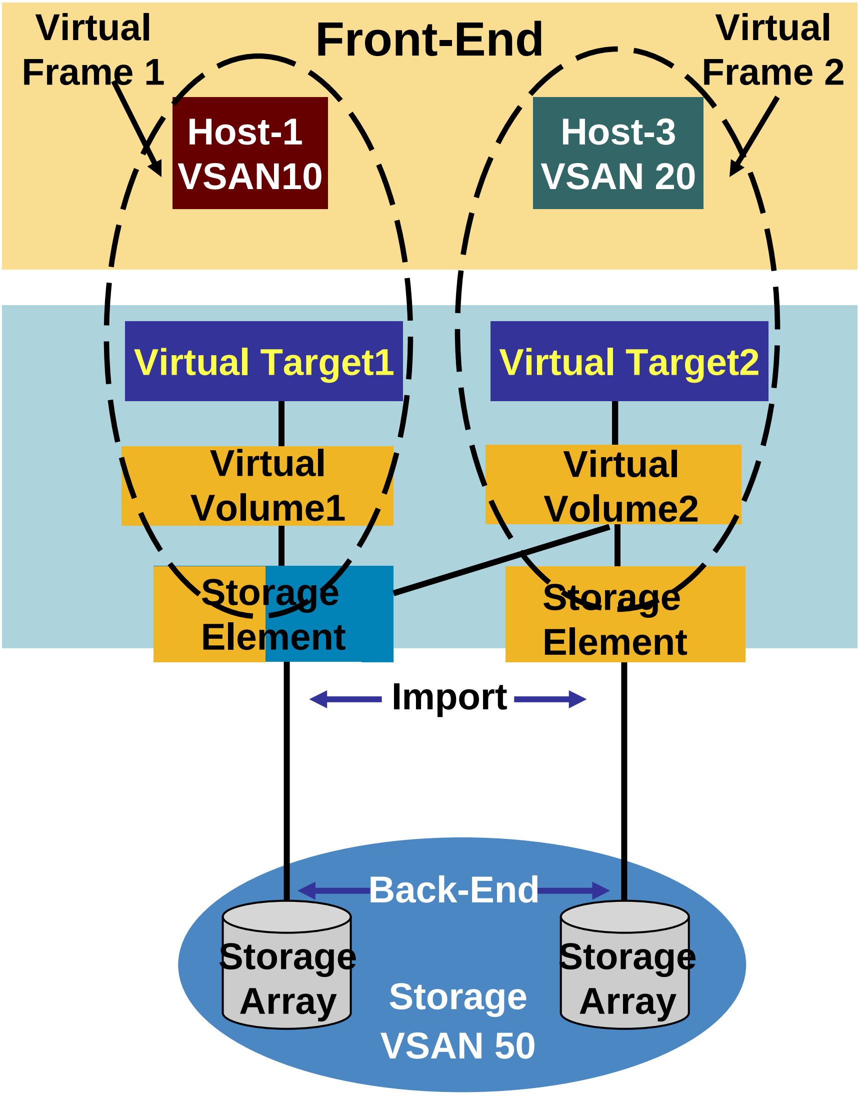

Storage Virtualization Features

Network Volume Management

- Works across heterogeneous arrays
  - Recommend creating volume within same class of storage
- Dynamically grow Virtual Volume
  - Increase size of raw volume but not file system
  - Will need "File System" that allows dynamic growth of file system
  - Does not support dynamic shrinking of volume size
- EMC Invista supports currently RAID 0
  - Stripe and concatenation
  - Future—RAID 1

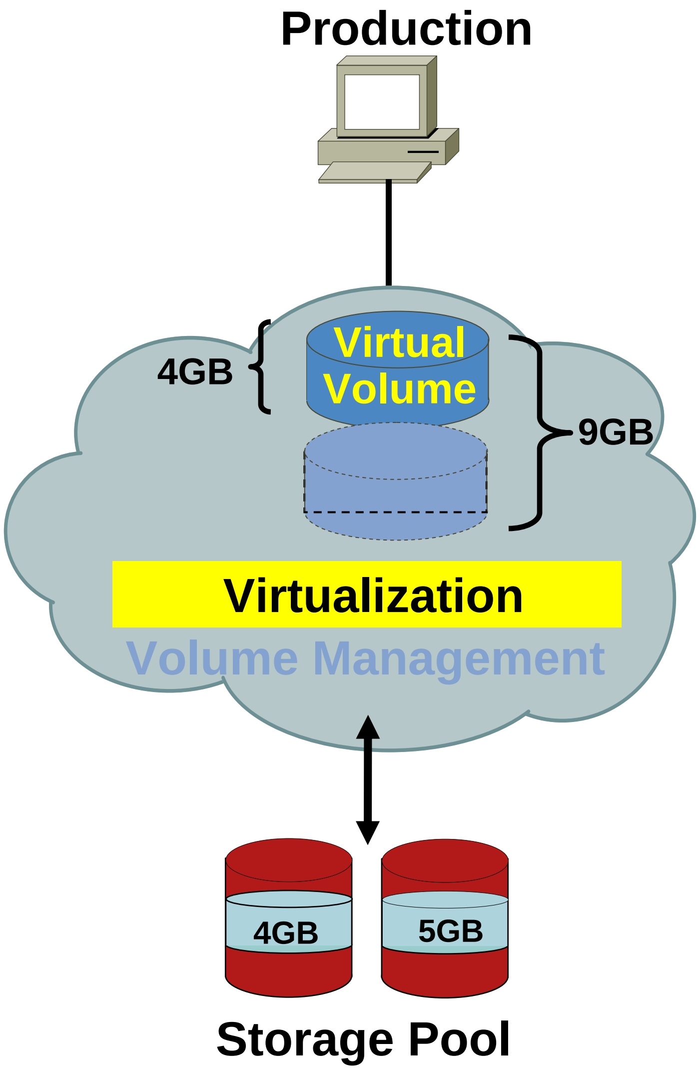

Heterogeneous Copy Services: Cloning

- Works across heterogeneous arrays
- Allocate any class of storage for cloning
- When in sync mode:
  - Host server reads/writes are to both source Virtual Volume and Clone Volume
- When in split (fracture) mode:
  - Host server reads/writes only to source Virtual Volume
  - Clone Volume can be mounted to same or separate host for any purpose

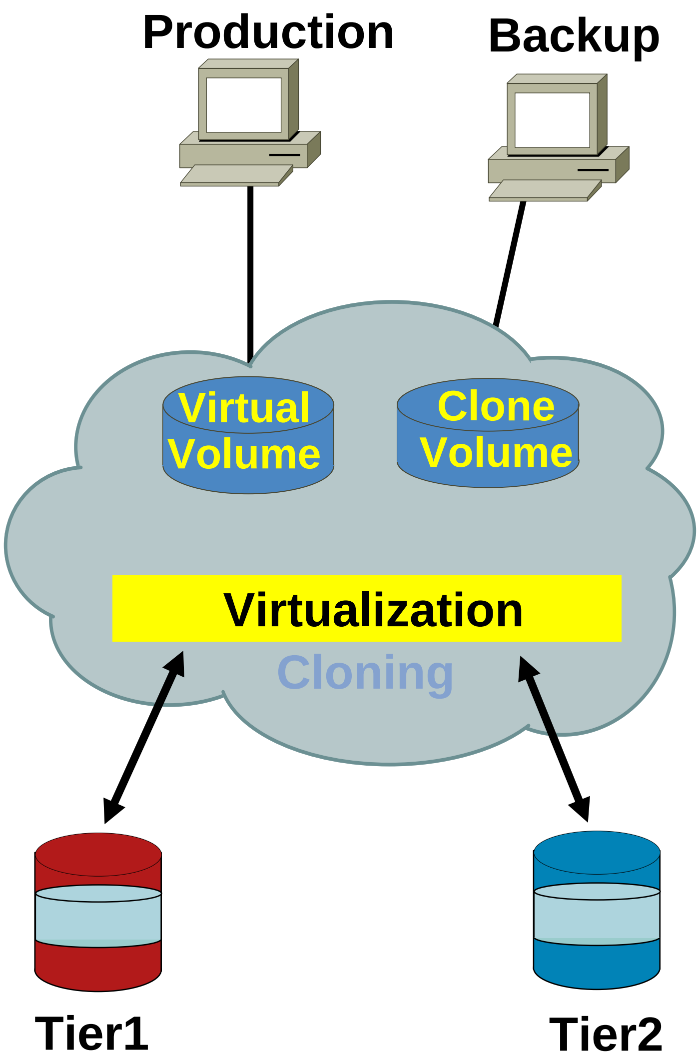

Cloning Data Flow—Initiation Process

1. Admin initiates Cloning process of VLUN 0 to another virtual volume (Clone 1)
2. Invista CPC informs SSM source virtual volume and target "clone" virtual volume
3. Data from VLUN 0 will start the copying process to virtual volume "Clone 1"

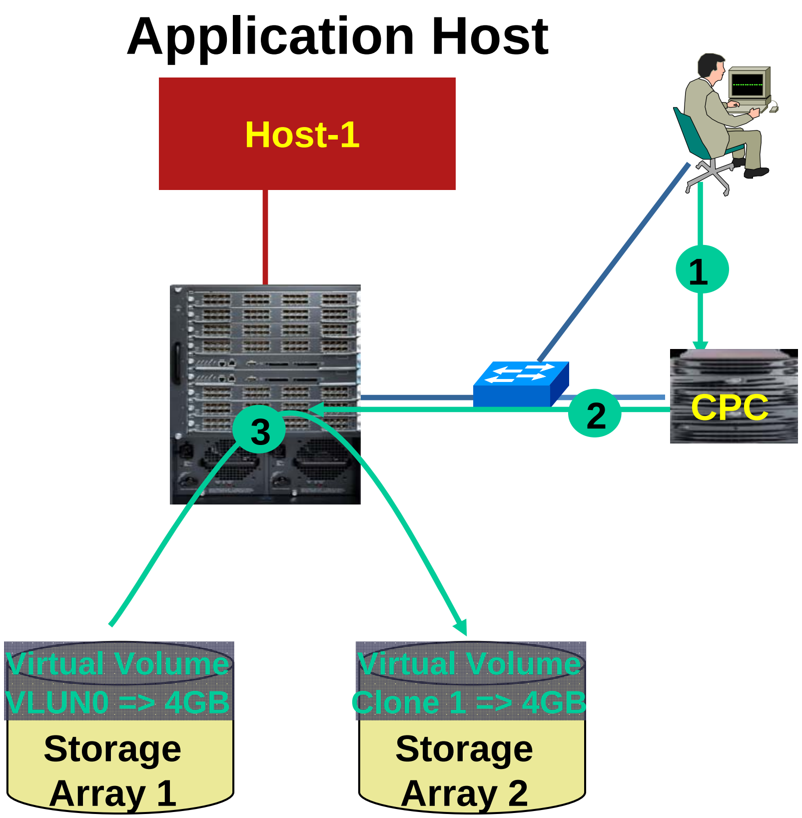

Cloning Data Flow—Application Host Read I/O

1. Host requests read I/O to VLUN 0
2. SSM sends request to appropriate physical storage

   Read I/O during cloning process can only be read from original source volume (i.e., VLUN 0)
3. Physical storage sends data to hosts

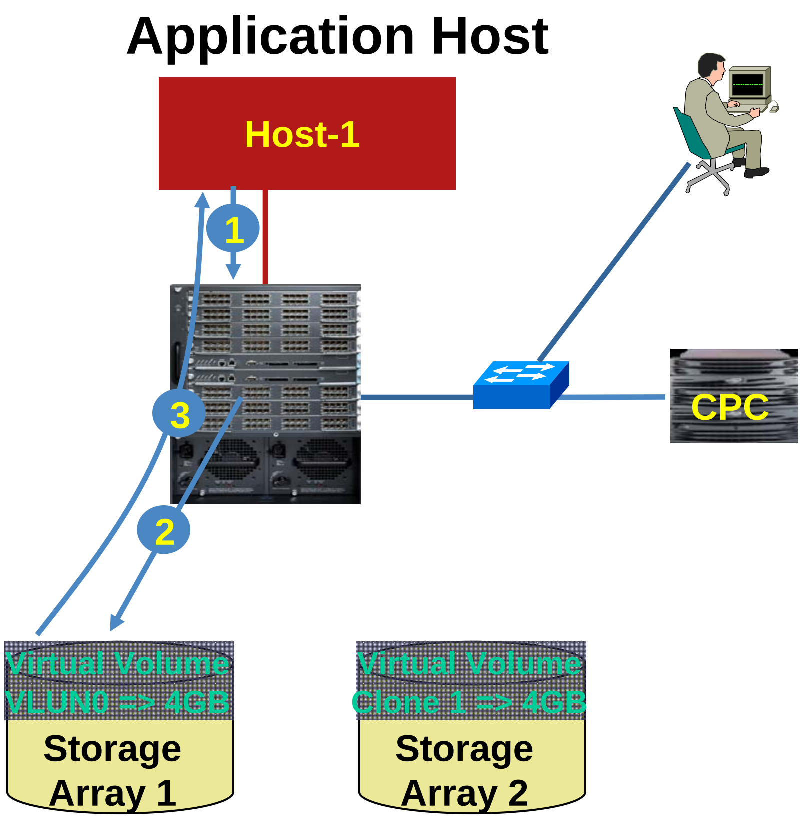

Cloning Data Flow—Application Host Write I/O

1. Host requests write I/O to VLUN 0
2. Host sends data to SSM
3. SSM sends data to both VLUN 0 and Clone 1 physical storage simultaneously
4. Both physical storage from VLUN 0 and Clone 1 sends status=good to SSM
5. SSM receives both acknowledgements and then sends status=good to Host

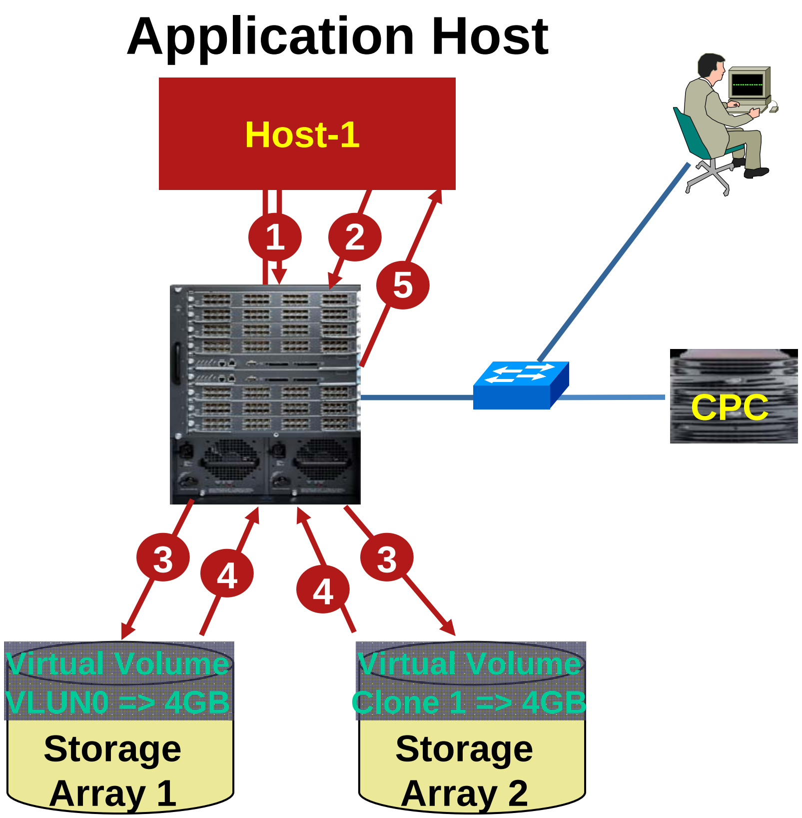

Seamless Data Mobility

- Works across heterogeneous arrays
- Non-disruptive to Application Host
- Can be utilized for “end-of-lease” storage migration
- Movement of data from one tier class to another tier

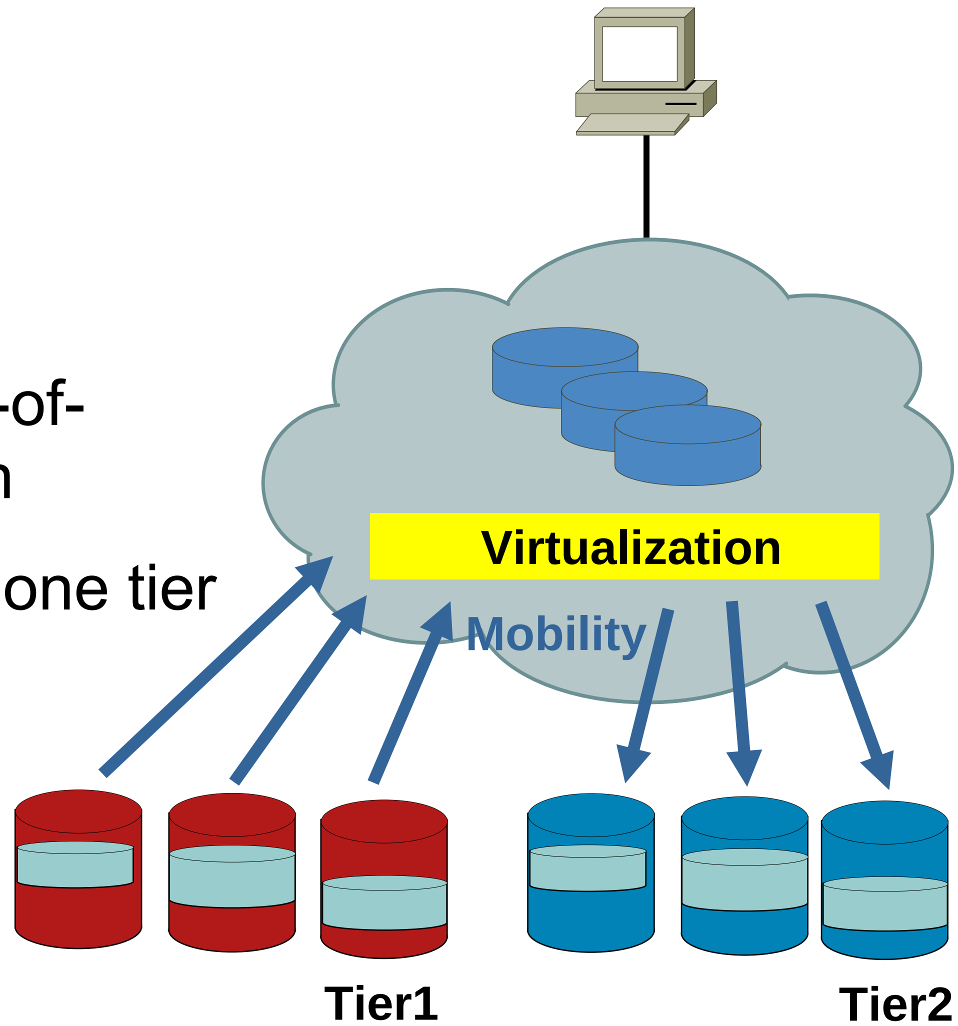

Data Mobility Data Flow—Initiation Process

1. Administrator initiates Data Mobility process
   of VLUN 0 to another virtual volume
2. Invista CPC informs SSM source virtual volume and target virtual volume
3. Data from VLUN 0 will start the copying process to target virtual volume

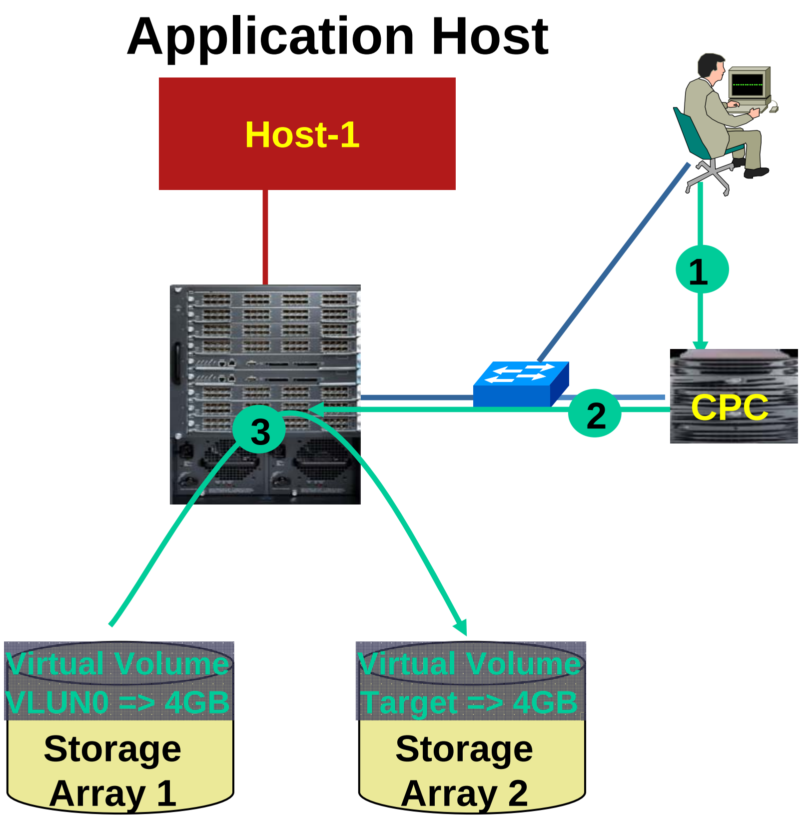

Data Mobility Data Flow—Application Host Read I/O

1. Host requests read I/O to VLUN 0
2. SSM sends request to appropriate physical storage

   Read I/O during migration process can only be read from original source volume (i.e., VLUN 0)

3. Physical storage sends data to hosts

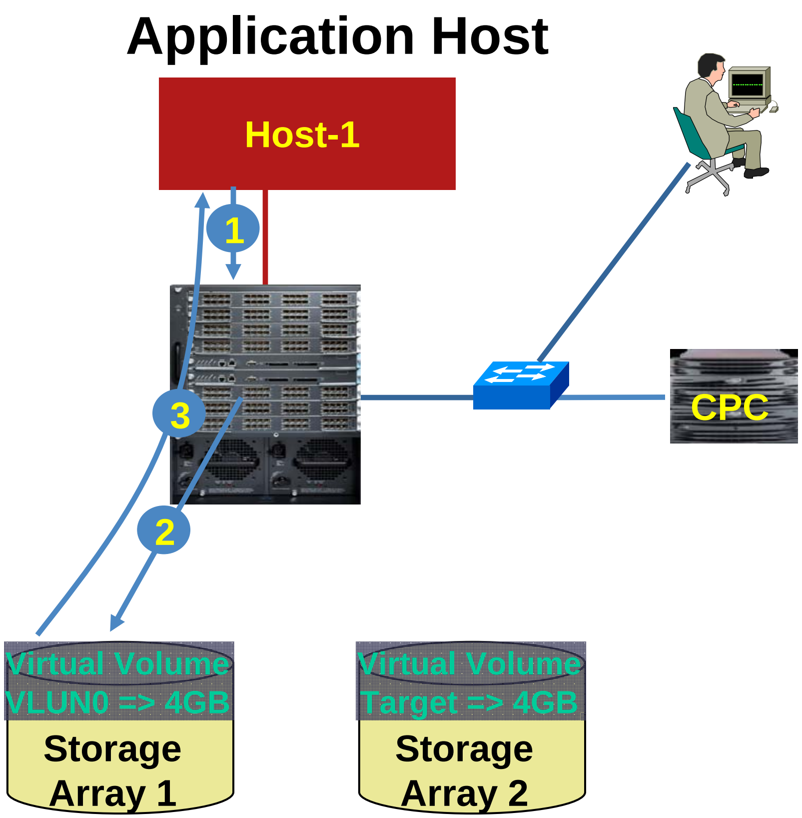

Note: Once migration is completed, read I/O will be read from target virtual volume only

Data Mobility Data Flow—Application Host Write I/O

1. Host requests write I/O to VLUN 0
2. Host sends data to SSM
3. SSM sends data to both VLUN 0 and "Target" physical storage simultaneously
4. Both physical storage from VLUN 0 and "Target" sends status=good to SSM
5. SSM receives both acknowledgements and then sends status=good to Host

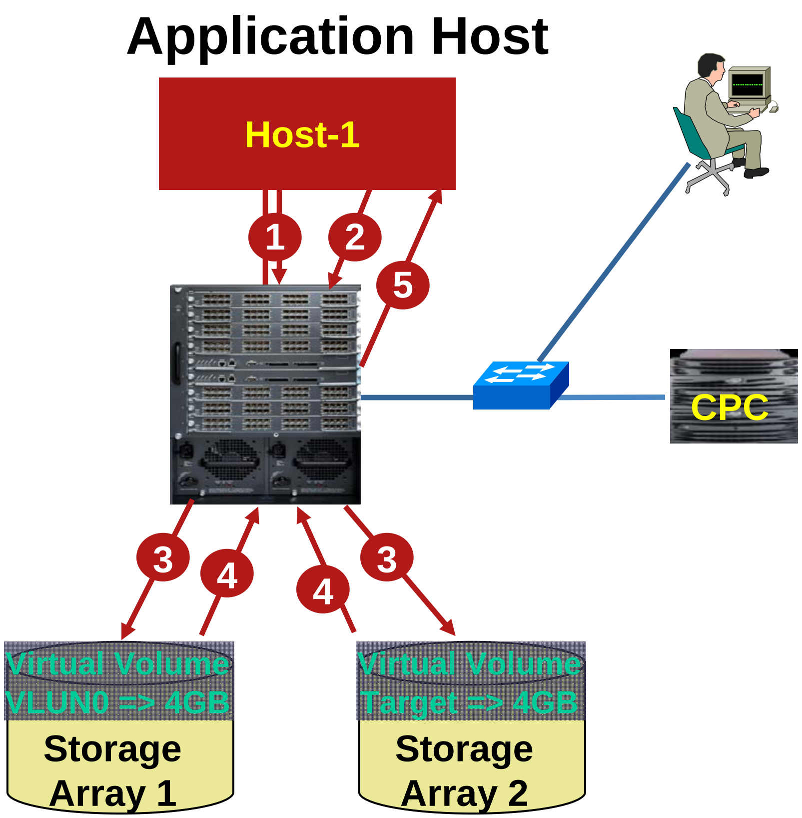

Note: Once migration is completed, switchover for all write I/O will go to only "Target" virtual volume

Cisco MDS Tasks and Best Practice

Cisco MDS Tasks for EMC Invista

Software

- SAN OS 2.1.2b for "Kickstart and System"
- SSI image version 2.1.2j
- EMC binaries uploaded to bootflash
    1. Salagent.bin
    2. Switch_defaults.conf
- Enable SSE license on SSM for EMC Invista

Invista CPC Communication

- Communication done through IP-FC
- IP address for "Control VSAN" and "SSM's CPP"
- Enable IP routing

VSAN and Zoning Consideration

Back-End VSAN

- Support of only one for Invista
- Default status to “Deny”
- Zone all nine VIs to storage ports

Front-End VSAN

- Up to 32 Virtual Targets per SSM
- Default status to “Deny”
- Zone server HBA to one Virtual Target

Control VSAN

- Communication to external CPC
- Default status to “Deny”
- Zone up IP interfaces for VSAN and SSM’s CPP

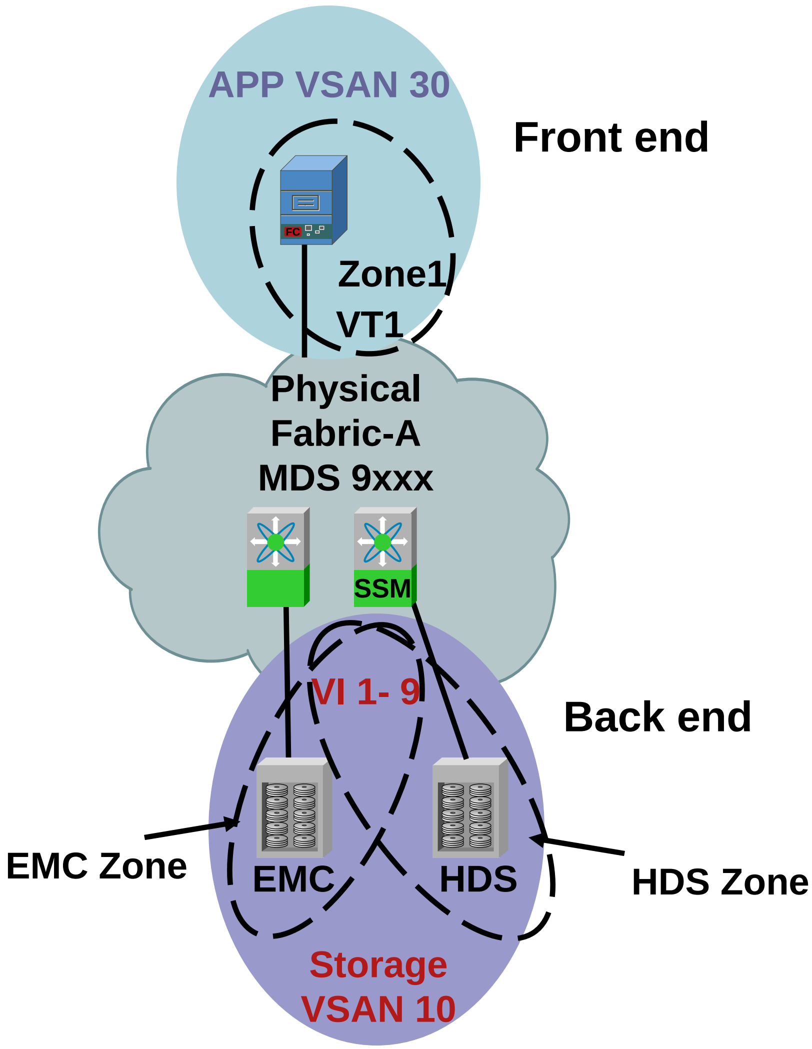

Storage Virtualization Design Consideration

SSM Data Flow—Host Directly Connected to SSM

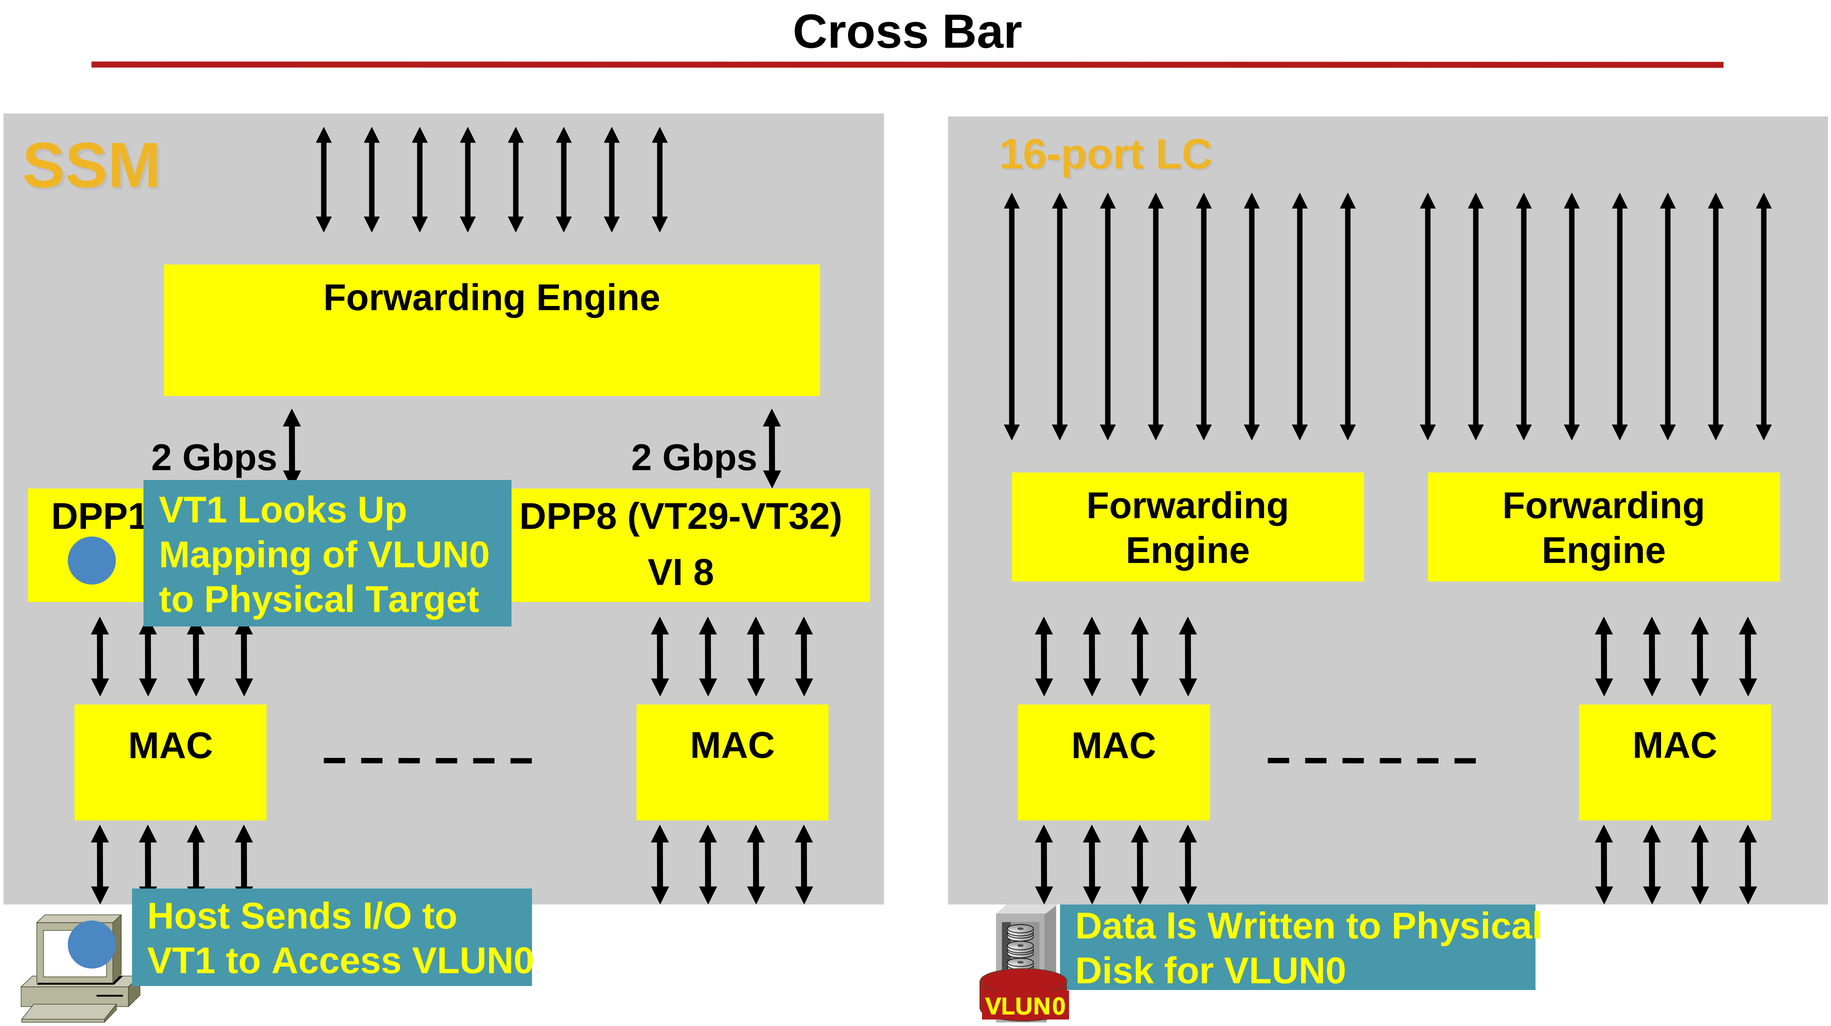

SSM Data Flow—Host Connected on Different Line Card Other than SSM

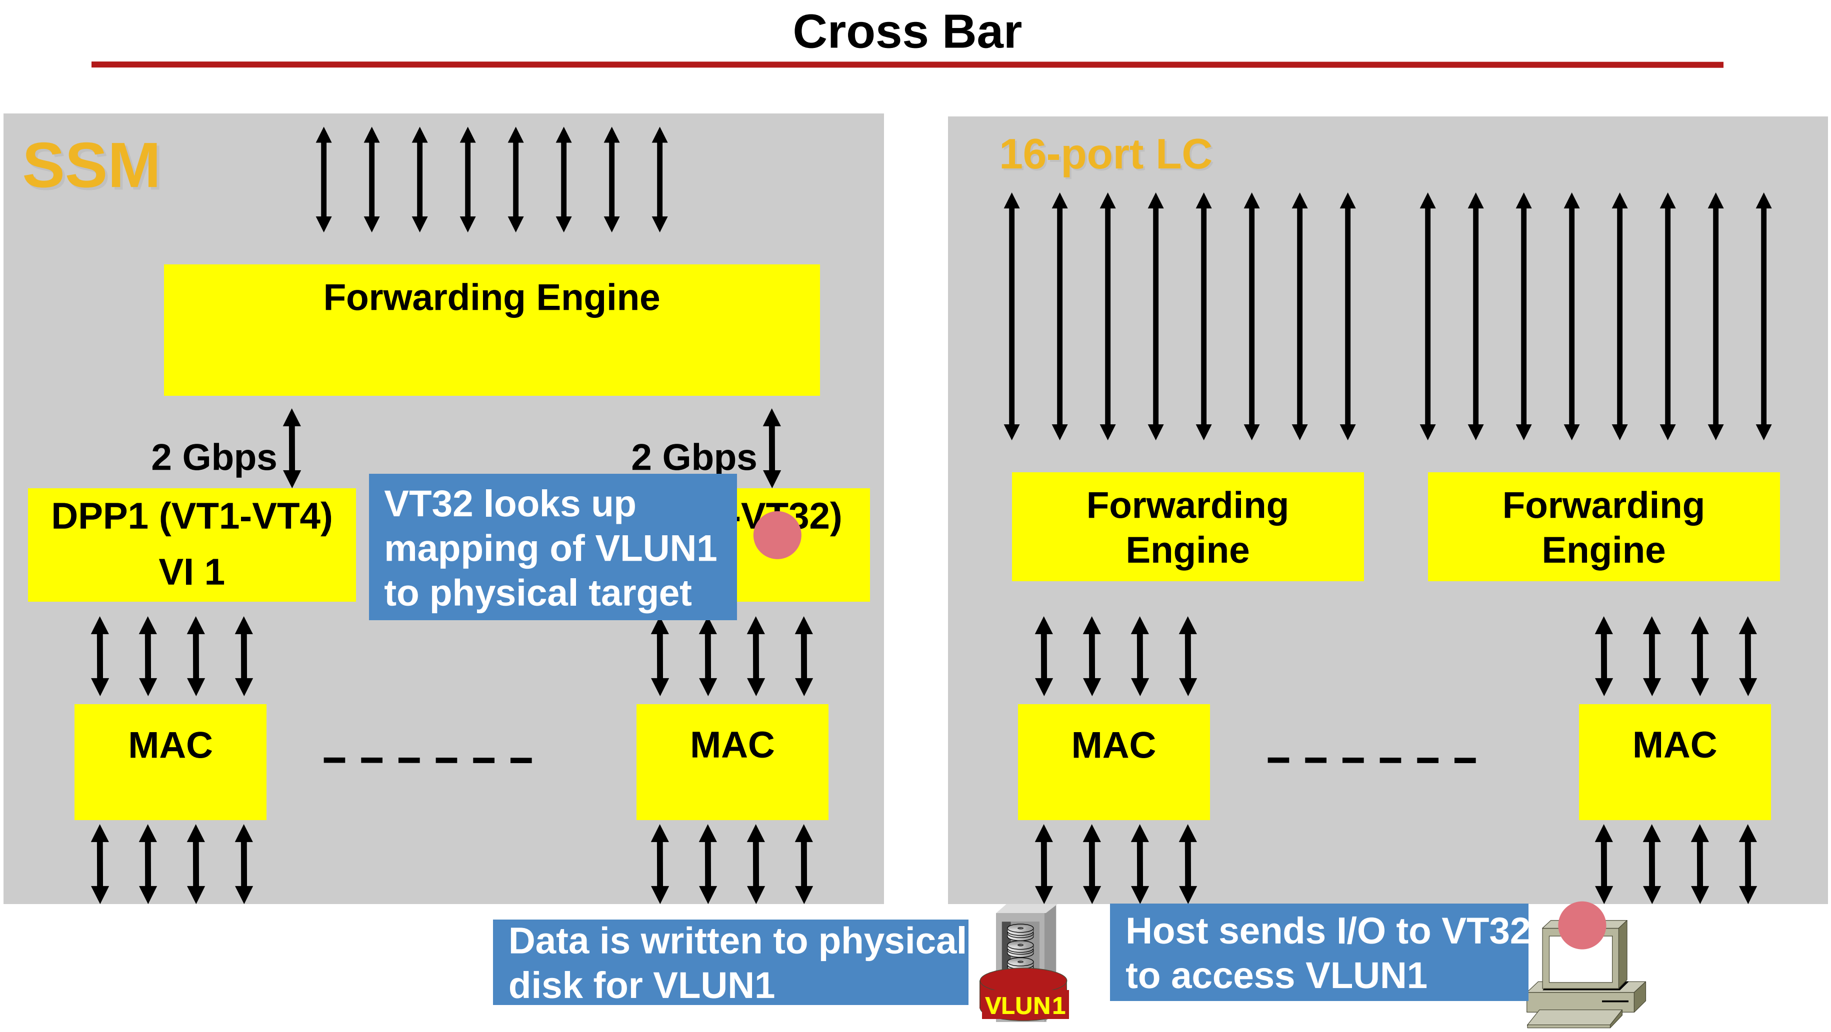

Cisco MDS SAN Topologies

- Collapse Core
- Core-Edge
  - Host and Storage at the Edge—localized traffic
    - Edge switches are either 9200 or 9500
    - Core switch used strictly for ISLs
  - Hosts at the Edge and Storage at the Core
    - Edge switches are 9100, 9200 or 9500
    - Hosts at the Edge
    - Storage and ISLs at the Core
- Edge-Core-Edge
  - Only ISLs at the Core
    - Edge switches are 9100, 9200, or 9500
    - Hosts at one side of the Edge
    - Storage at the other side of the Edge

Collapse Core

- Benefits
  - Can virtualize everything in Collapse Core
    - 32 ports for storage
    - 160 ports for hosts
  - Simple management
  - No hops
- Concerns
  - Not scalable

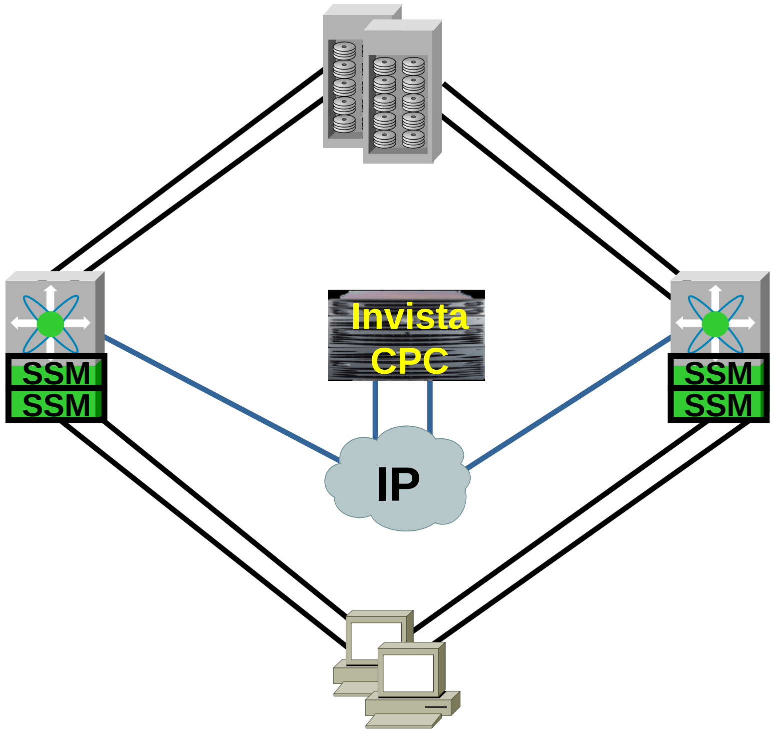

Core-Edge: SSM at the Edge

- Benefits
  - Can virtualize a larger fabric
    - Can virtualize more storage
    - Can virtualize more hosts by X number of edge switches
- Concerns
  - Not centralize for virtualize I/O on hosts
  - Multiple hops to access SSM for Edge switches without local SSM
  - More complex to manage more SSMs on different switches

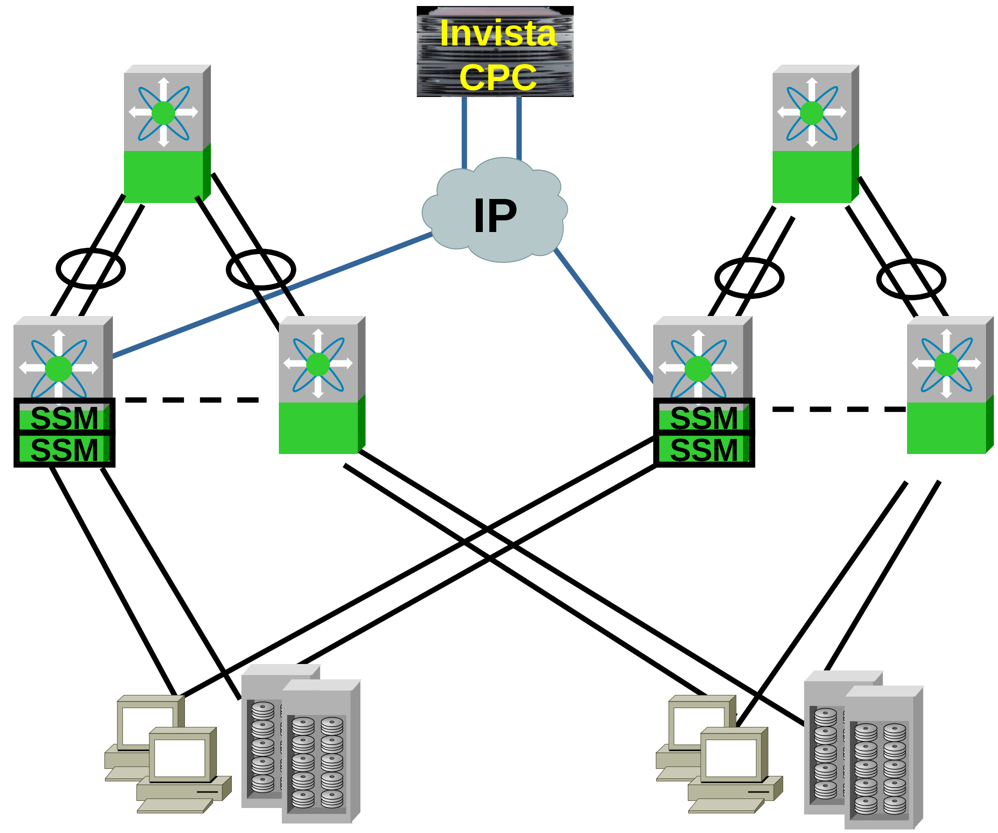

Core-Edge: SSM at the Core

- Benefits
  - Can virtualize a larger fabric
    - More storage can be virtualized
    - More hosts by X # of edge switches
  - Centralize virtualize network
  - Single hop for all hosts to SSMs
- Concerns
  - SSM not recommended for Storage or ISLs connectivity
  - Not fully utilizing 32 ports on SSMs other than virtualization traffic

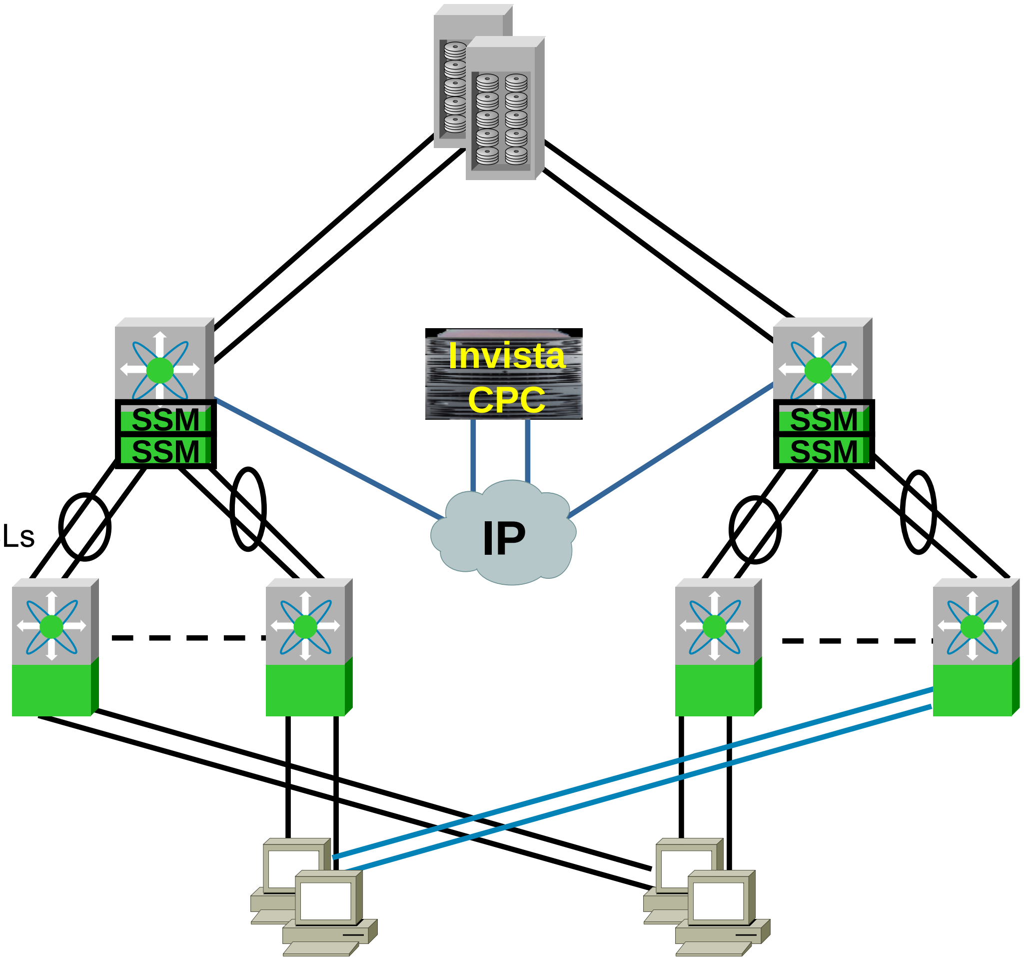

Edge-Core-Edge: SSM at the Core

- Benefits
  - Can virtualize a larger fabric
    - More storage by X # of edge switches
    - More hosts by X # of edge switches
  - Centralize virtualize network
  - Single hop for all hosts/storage to SSMs
- Concerns
  - SSM not recommended for Storage or ISLs connectivity
  - Not fully utilizing 32 ports on SSMs other than virtualization traffic
  - Additional hop for Storage to SSM

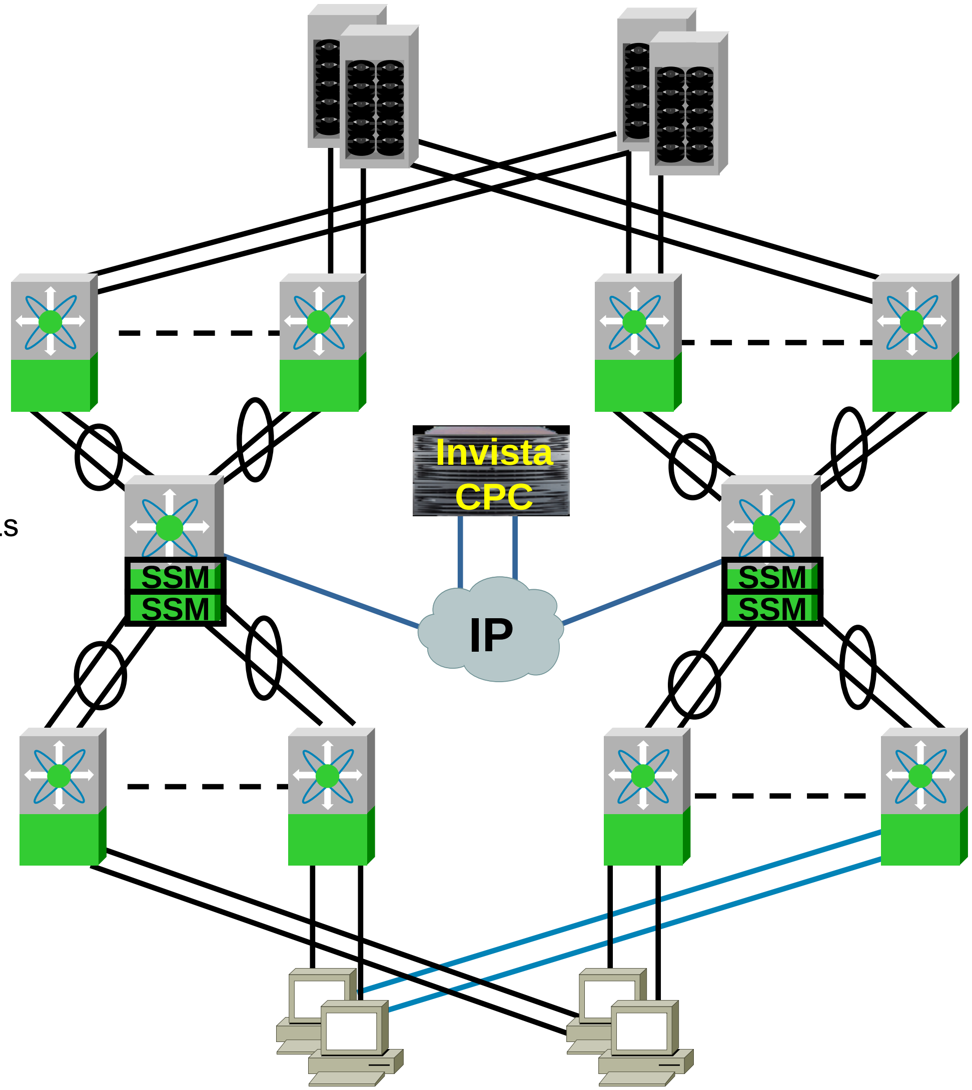

Conclusion

- Virtual SANs and VSAN routing are analogous to Virtual LANs and IP routing
- VSANs now form basis of ANSI T11 standard
- Virtual fabrics and fabric routing help scale fabrics and provides extra security
- All virtualization and routing is done in hardware—no performance impact

Summary

- Storage Services Module (SSM)
    - Focuses data movement processing in the network instead of server or storage
    - Scales as storage area networks scales
- Hosting of Intelligent fabric applications
    - Fibre Channel Write Acceleration (FC-WA)
    - Network Accelerated Serverless Backup (NASB)
    - SANTap
    - Switch virtualization with EMC Invista

Meet the Experts
Data Centre

- Victor Moreno
  Technical Leader

Recommended Reading

BRKDCT -3008

- Data Center Fundamentals
- Storage Networking Fundamentals
- Network Virtualization

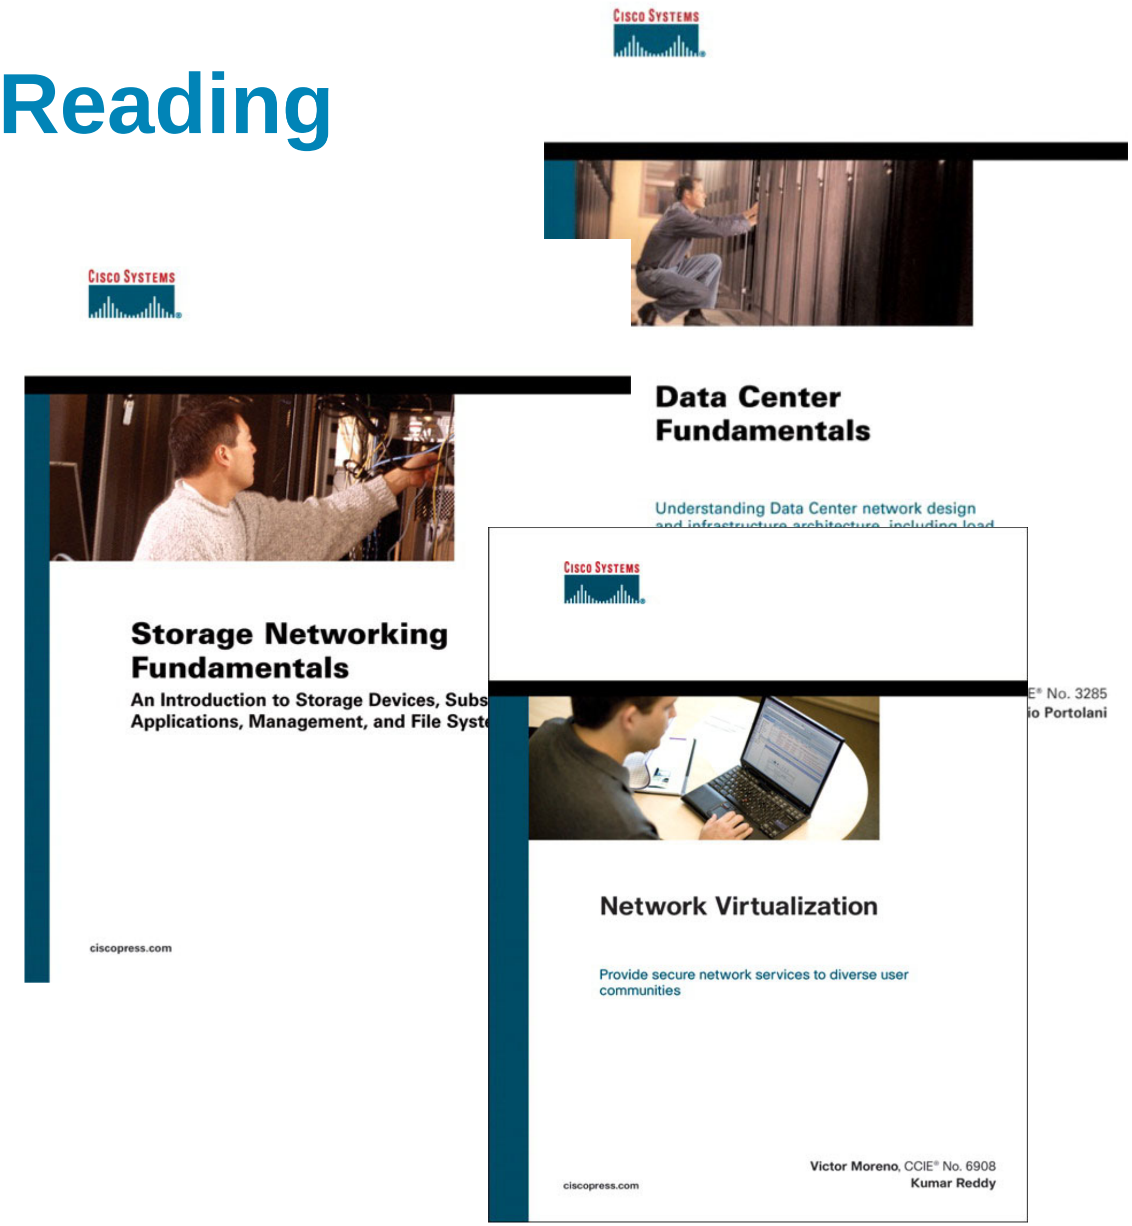

Available in the Cisco Company Store