# System z FICON and FCP Fabrics – Intermixing Best Practices

Mike Blair – mblair@cisco.com
Howard Johnson – hjohnson@brocade.com

10 August 2011 (9:30am – 10:30am)
Session 9864
Room Europe 7

# Abstract

In this jointly presented session, the major players in storage networking will discuss:

1. Decision criteria for deciding whether to merge Open and FICON fabrics
2. DOs and DONTs for FICON/OPEN merged fabrics - best practices
3. How NPIV plays into the fabric definition and how to best to zone for zSeries Linux environment.
4. Management options / best practices for merge Fabrics.

At the end, there will be time for Q&A.

# Agenda

- Merging Open System and FICON Fabrics
  - Intermix Mode
  - Converged / Merged Fabrics
  - Consolidation / Virtualization
- Considerations when consolidating fabrics
  - Asset Utilization
  - Management
  - Human Factors
  - Application Goals
- Managing Merged Fabrics
  - Virtualization (NPIV)
  - Fabric Virtualization (VSAN / Virtual Fabrics)
  - Isolation of Resources (Zoning)
  - CUP

Intermix Mode

Converged / Merged Fabrics

Consolidation / Virtualization

# MERGING OPEN SYSTEM AND FICON FABRICS

# Intermix and Replication

- Production Operations
  - Access ECKD
  - FICON channel
- Replication
  - Array to Array
  - FCP Interface

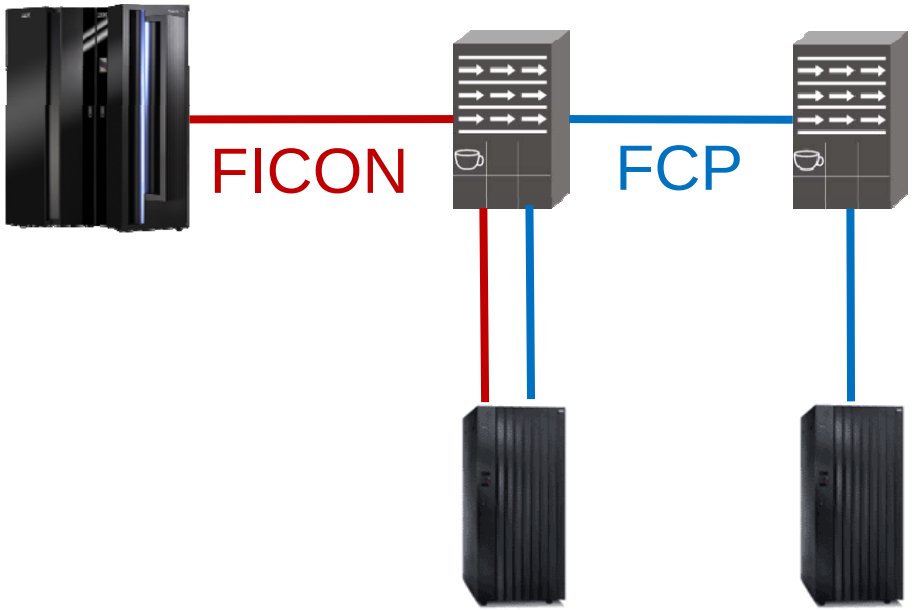

# Intermix and FCP Channel

- FICON
  - zOS
  - ECKD Storage
- FCP
  - Linux for System z
  - Open Systems Storage

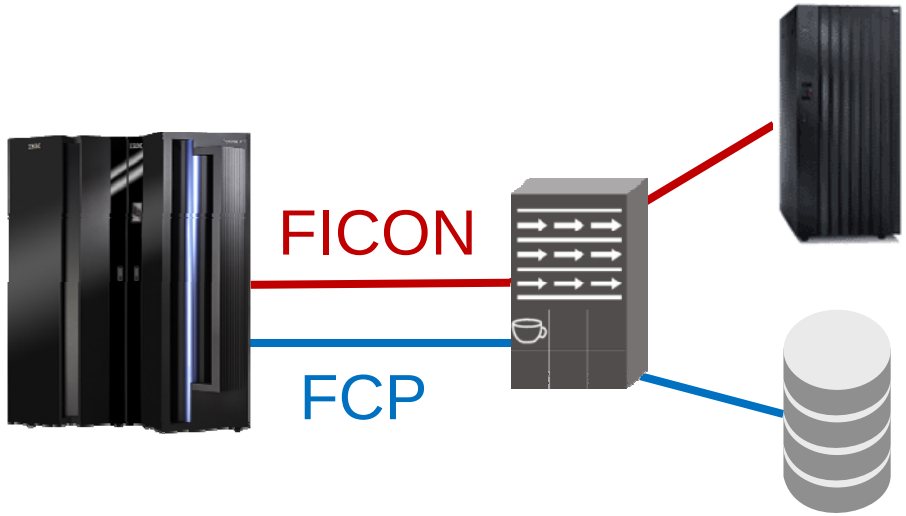

# Intermix for Tape Archive

- Limited Resource
- Shared Across LOBs
- Dual Personality
  - FICON
  - FCP

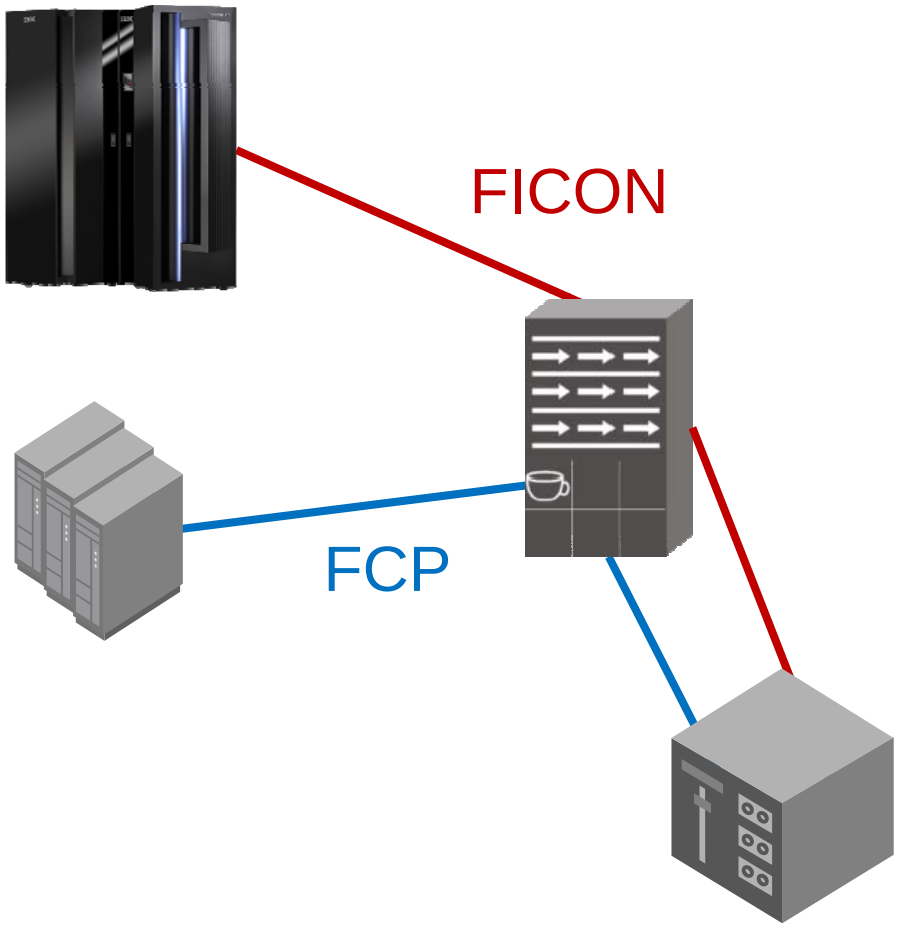

# Fabric Migration Stages

## Independent Processing, Storage, and Fabrics

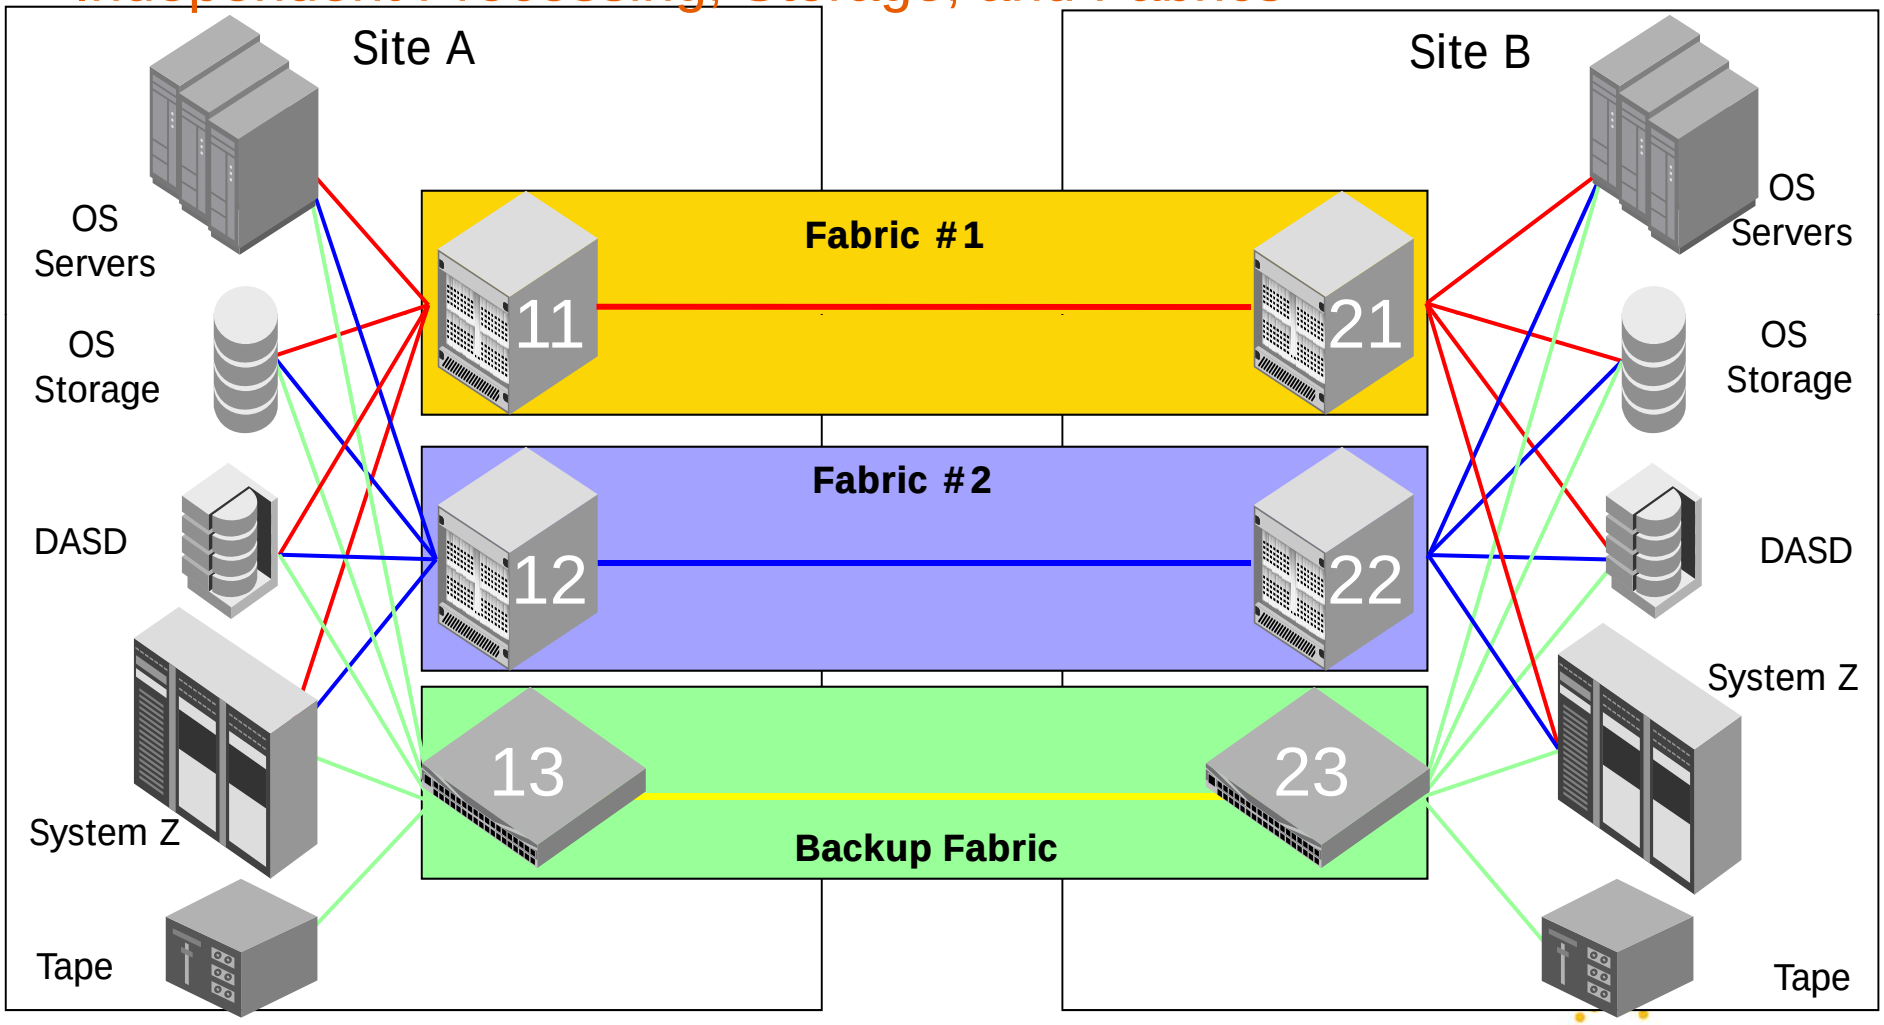

# Fabric Migration Stages

## Processor Optimization for Application Consolidation

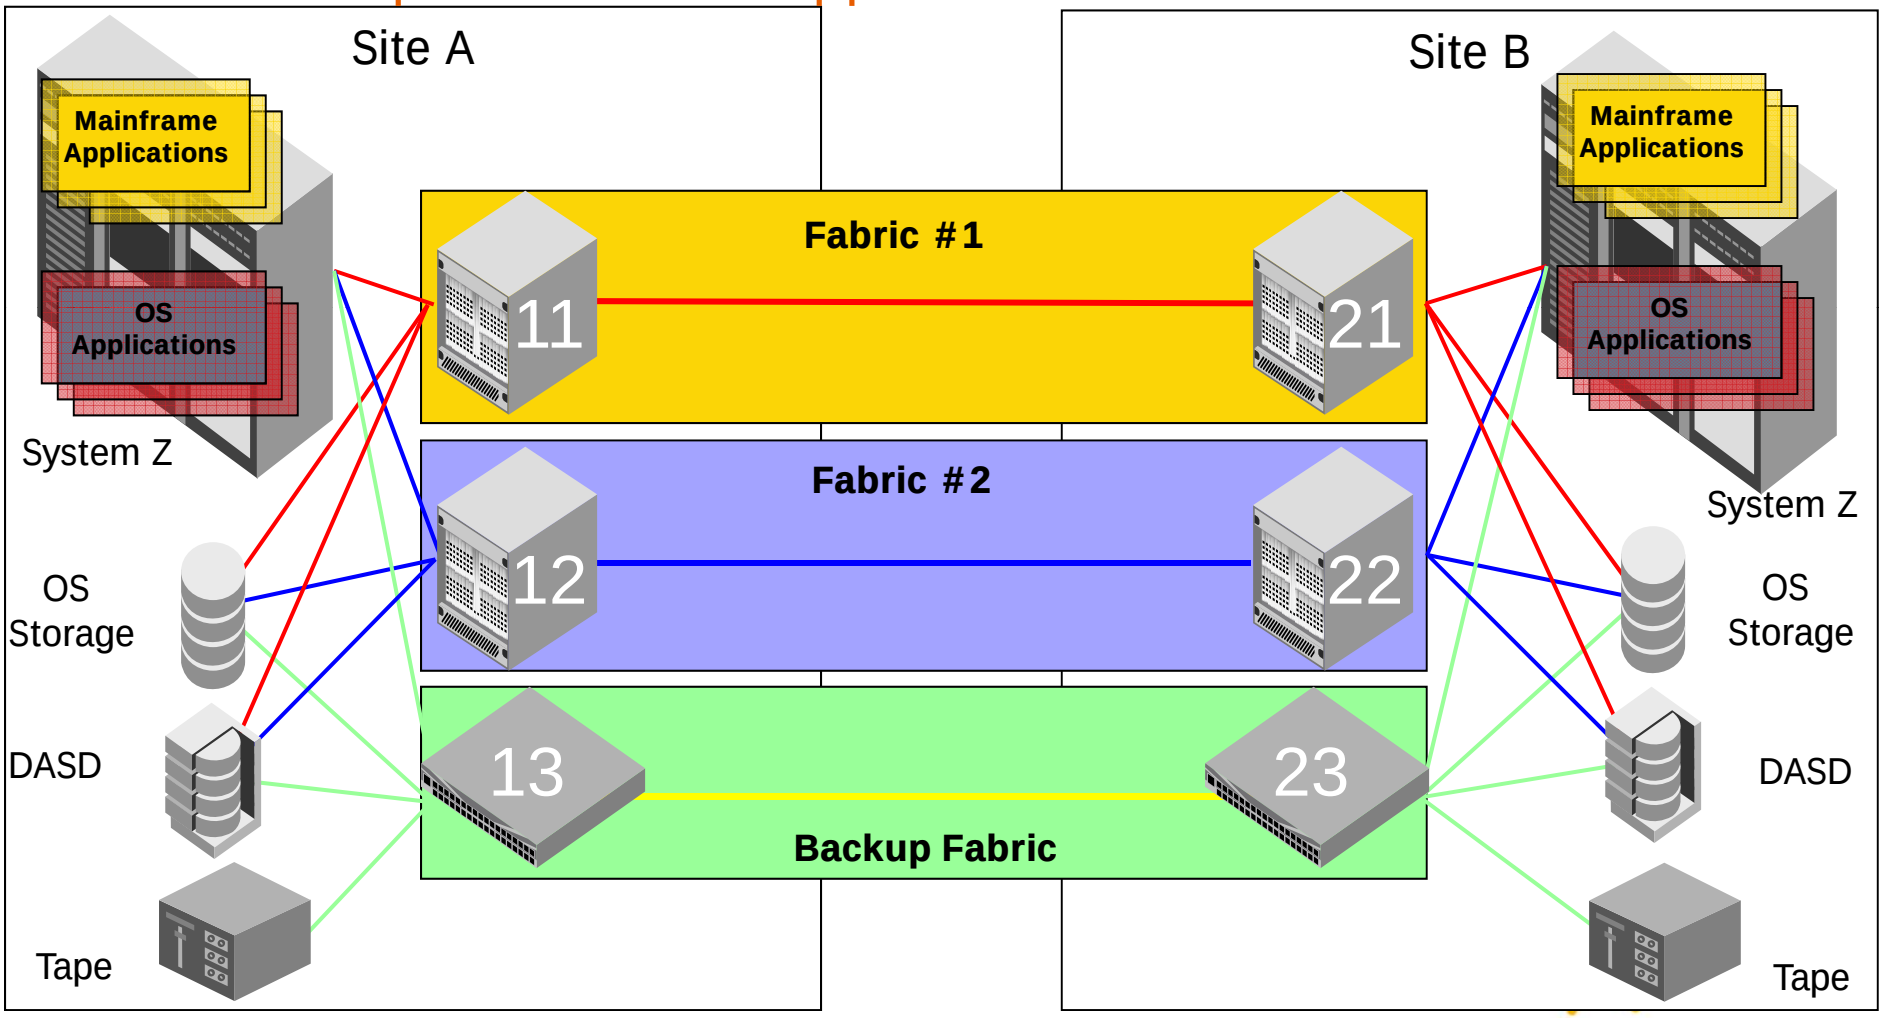

# Fabric Migration Stages

## Network Optimization using Virtual Fabrics

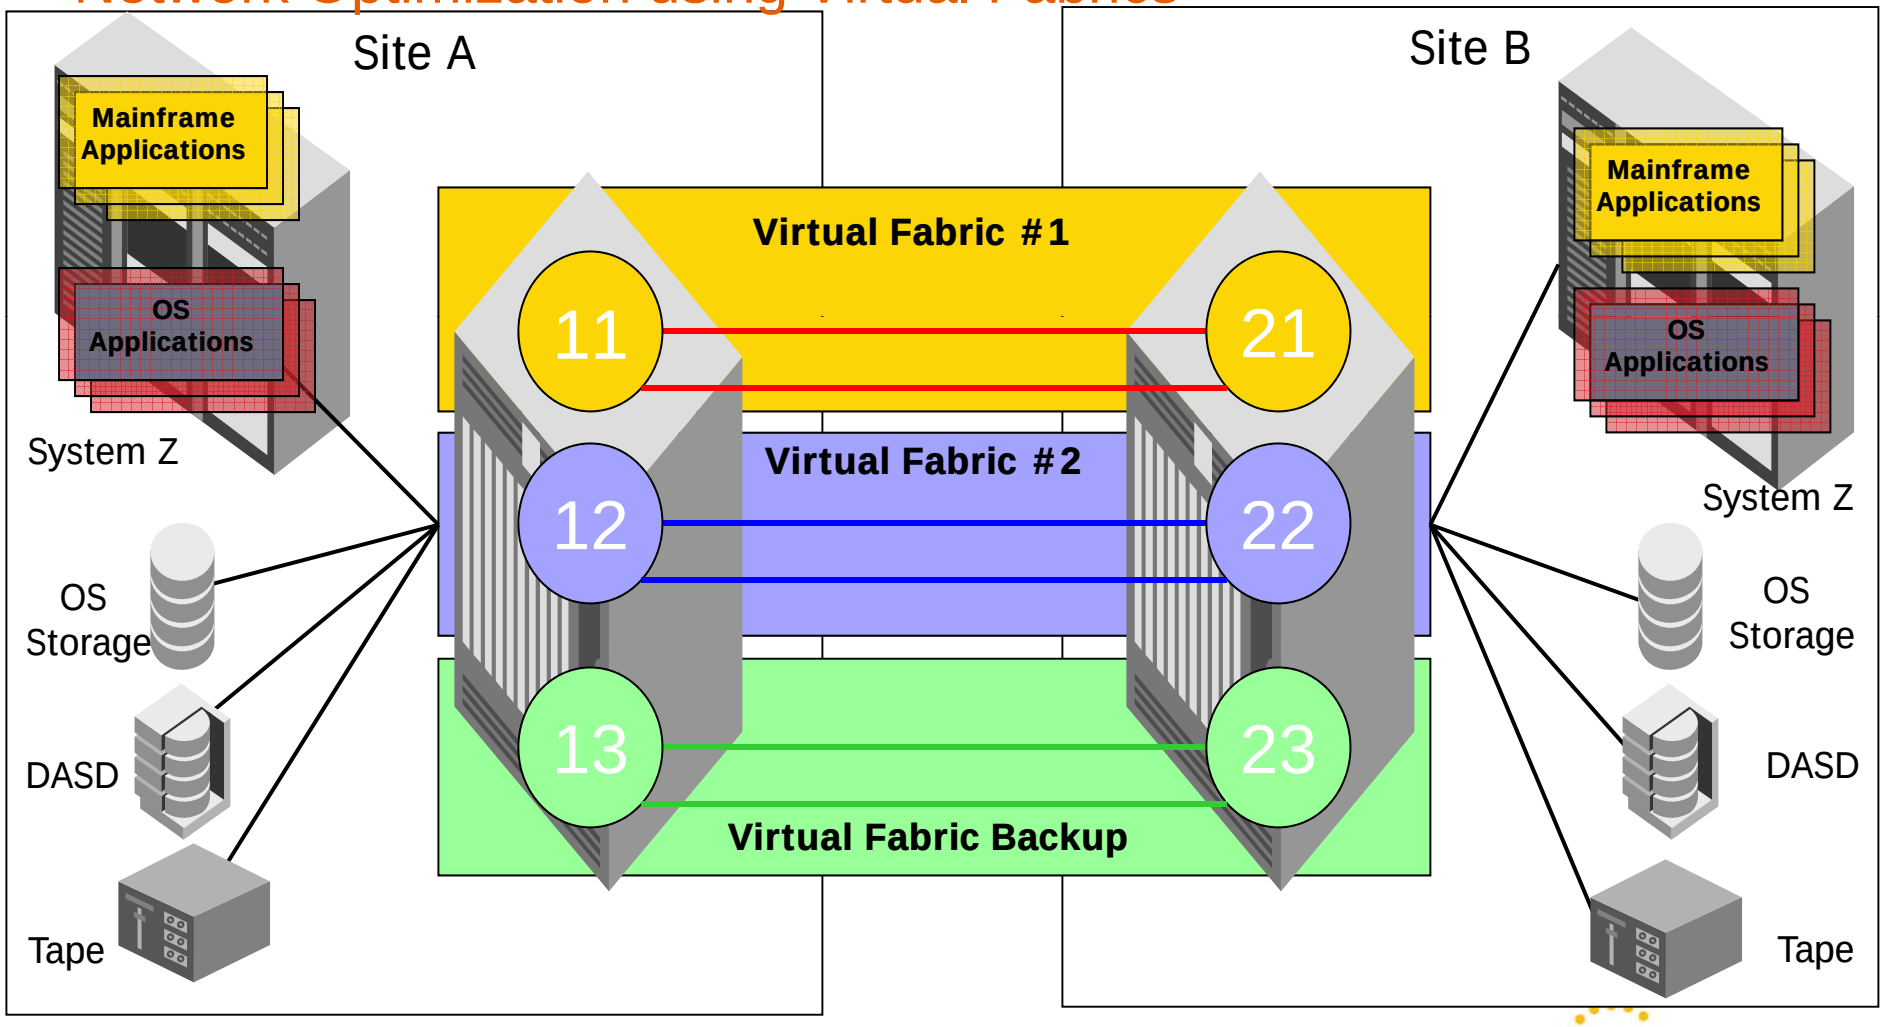

# Fabric Migration Stages

## I/O Optimization

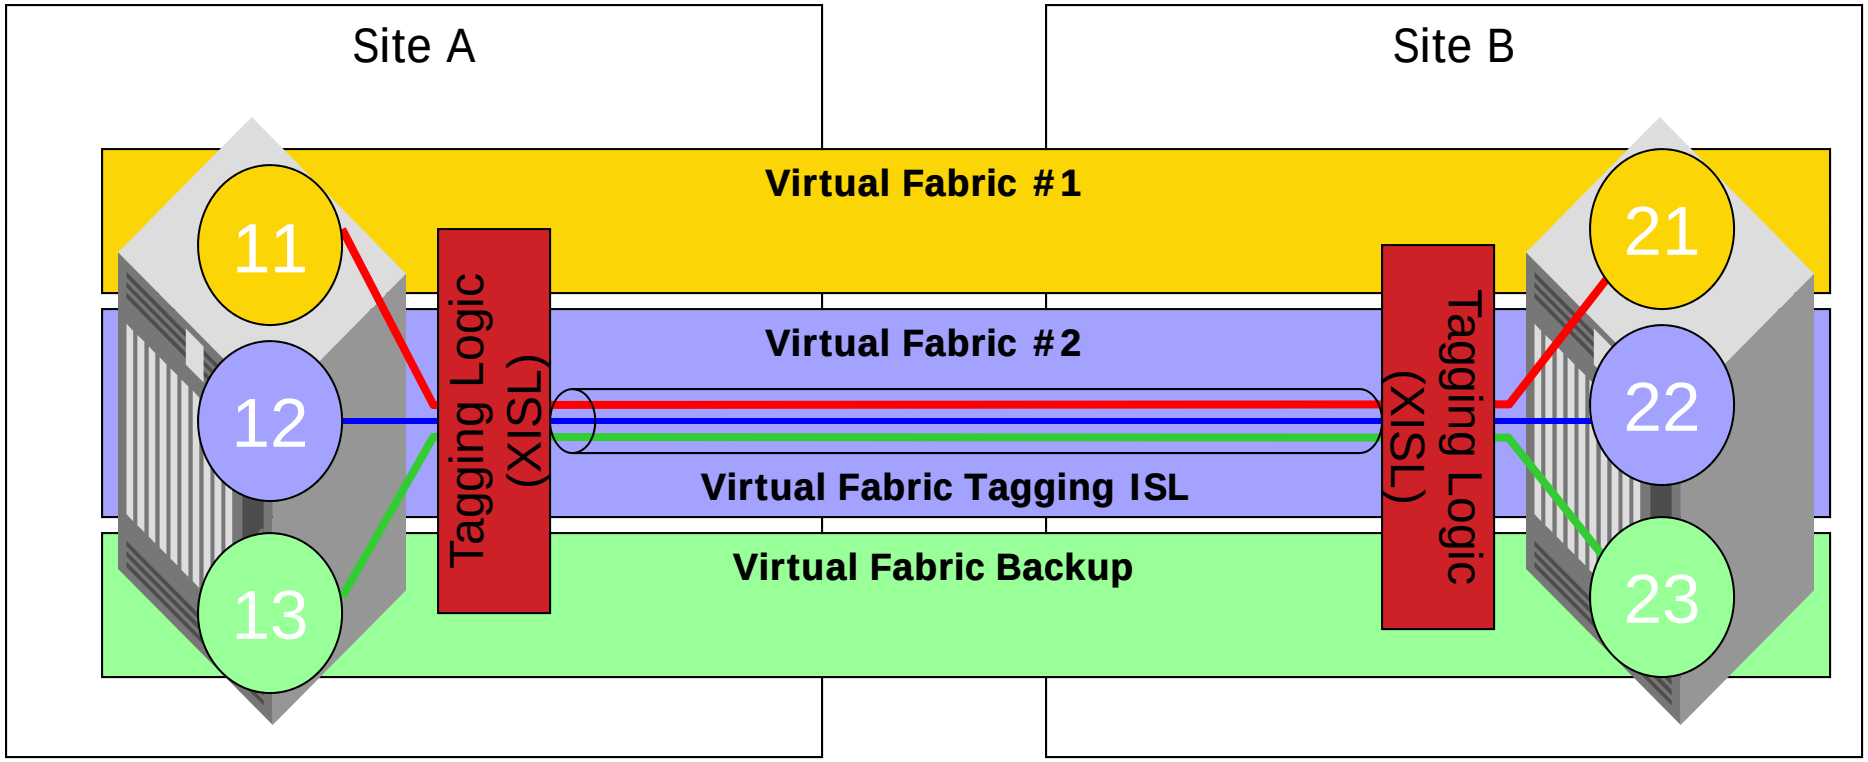

# Integrated or Isolated ISLs per VSAN / VF

- Isolated

  Production, Dev/Test and Linux are all running on a z9

  Isolate the ISL usage between these systems

  4 ISLs per each of the three environments

## Solution:

- Define 4 ISLs only used by single VSAN
- Port Channel for High Availability

Gives Total Isolation – no interference

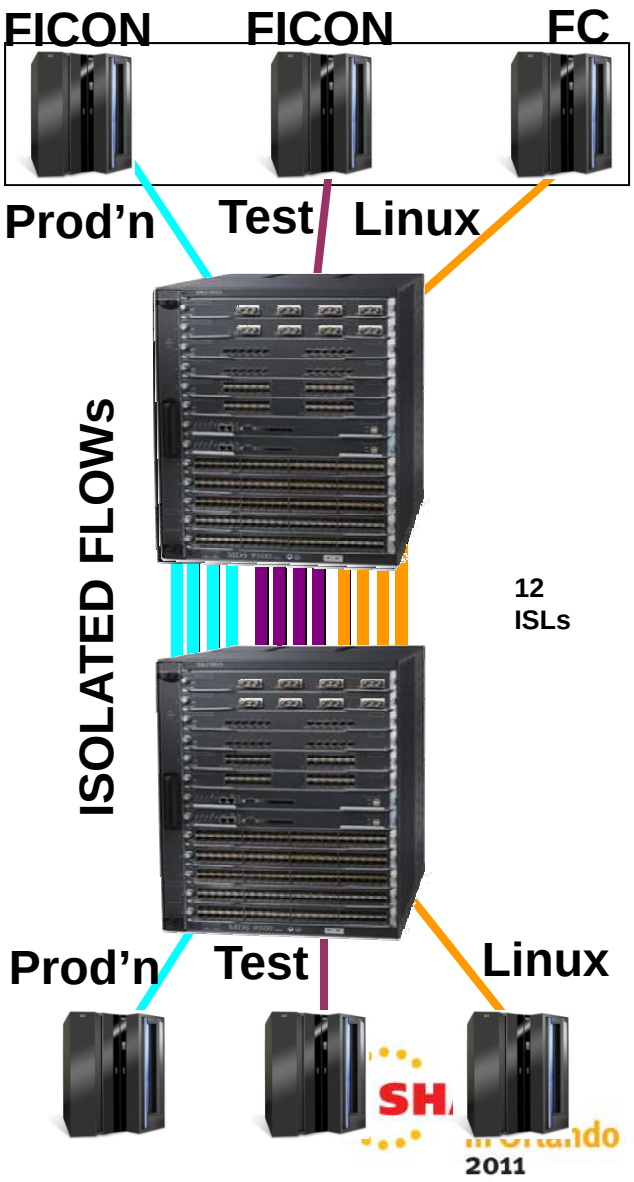

# Integrated or Isolated ISLs per VSAN / VF

- Integrated
  - Production, Dev/Test and Linux are all running on a z9
  - Integrate ISLs for highest availability
  - All 12 ISLs bandwidth are available for peak usage

## Solution:

- Define 12 ISLs carrying all 3 VSANs
- Port Channel for High Availability
- Potentially use QOS to prioritize

Great for if peak usage at different times

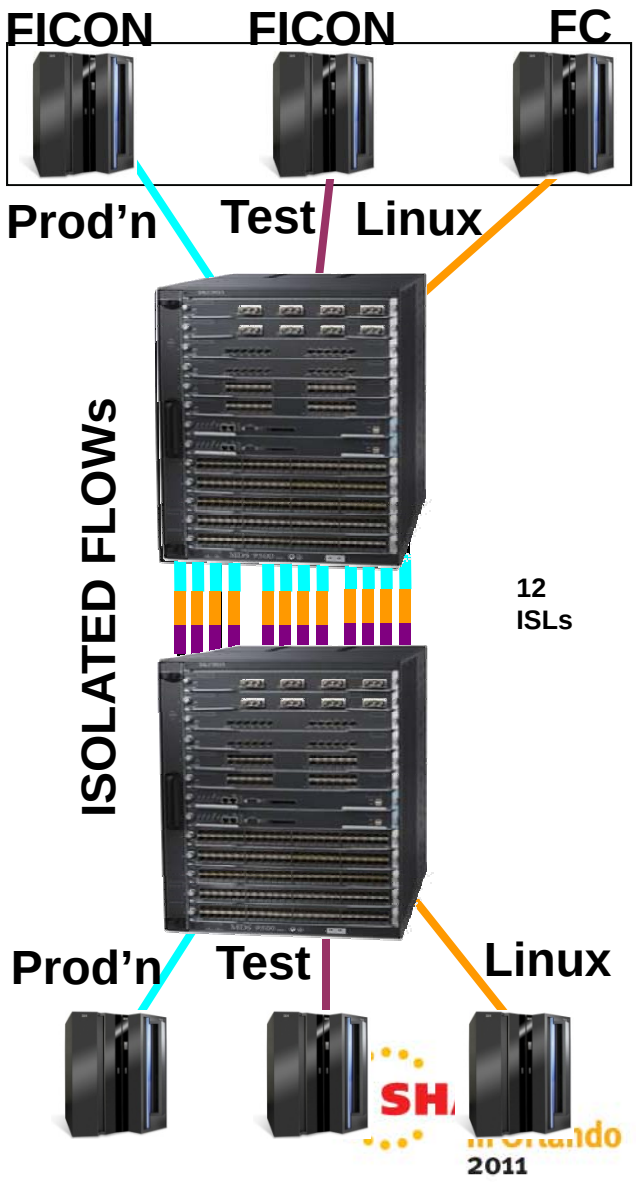

# Virtual Fabric Tagging

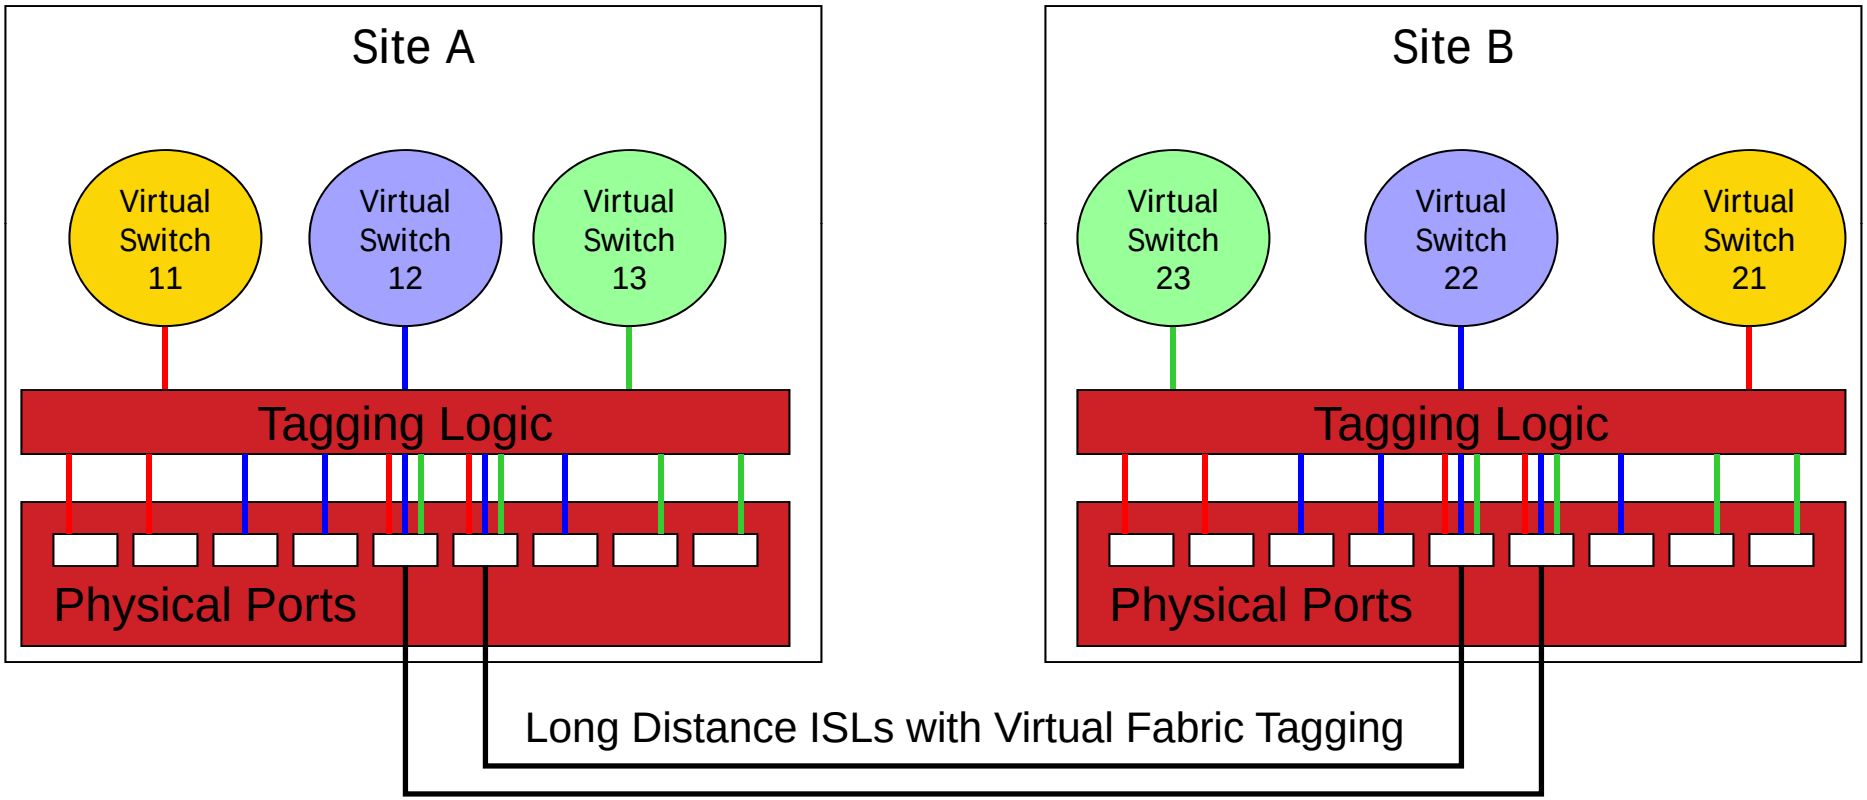

Asset Utilization

Management

Human Factors

Application Goals

# CONSIDERATIONS WHEN CONSOLIDATING FABRICS

# Considering Intermix

## Fabrics overview, enablers, and challenges

- Isolation of heterogeneous FC fabrics has been a standard practice in the data center
- Why–or why not–deploy intermix in your data center now?

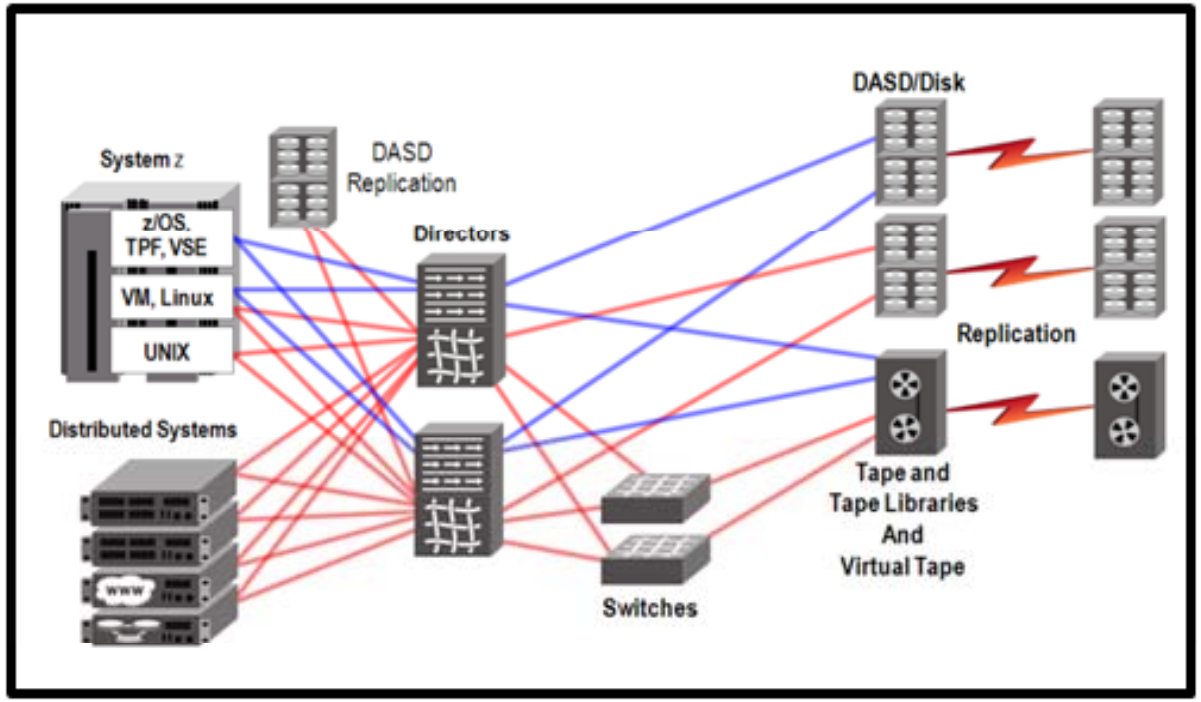

# Customers Today

## Separate Mainframe and Distributed Environments

- Distributed processing
  - It's different ;-)
  - Supports a broad mix of availability levels
  - Mainframe supports only the "gold standard"
- Availability requirements are different
  - Risk adverse environments
    - Call home if errors happen
  - Risk tolerant environments
    - Retry/reboot
  - Both run mission-critical applications
- Merged environments
  - Strictest requirements are deployed for all
  - Application goals and expectations must be understood

# Mainframe Environments

## PREDICTABILITY

- Avoid risk and design for redundancy
  - Directors instead of switches
  - Avoid unscheduled outages
  - Minimize or eliminate scheduled outages
- Workload predictability and stability
  - Moved from one set of resources to another
  - Measure what's currently going on
- I/O Behavior
  - Influenced or directed by Operating System / Hypervisor
  - Predictability in path selection (RMF)
  - Orchestrate "network connectivity" to optimize performance
- Conservative deployment of new features

# Distributed Environments

## EASE of DEPLOYMENT and MANAGEMENT

- Accept risk and design for flexibility
  - Switches instead of Directors
  - Accommodate unplanned outages
  - Cost sensitive
    - Redundancy deploy only when mission critical
  - Regularly scheduled outages
- Workload flexibility
  - Performance can vary – not typically a concern
  - Movement is allowed to be disruptive
- I/O Behavior
  - Layer-2 and layer-3 routing exploited for connectivity
  - Path selection or influence is a low priority
    - It's the SAN's job

# Consolidation Drivers

## Enablers that make it worth your consideration

- Reduce operational cost
  - Footprint
  - Energy
  - UPS demand
  - Increase efficiency
- Optimize asset cost and utilization
  - Virtualization
- Consolidate applications to fewer platforms
  - Linux on System z
- Long-term IT Strategy
  - Cloud computing

# Protocol Intermix Considerations

## A different way to think about provisioning I/O

- Common Architecture
  - Accommodate most highly available applications
  - Avoid single points of failure (Five-9's)
  - Avoid over provisioning (optimal performance)
- Component Utilization
  - Port density
  - Feature options (wide area capabilities)
- Change Management
  - Consumption of resources
  - Isolation of change management
  - Authorization of changes
    - Set and "don't touch"
    - "Plug and play" environment

Virtualization (NPIV)

Fabric Virtualization (VSAN / Virtual Fabrics)

Isolation of Resources (Zoning)

CUP

# MANAGING MERGED FABRICS

# What is NPIV?

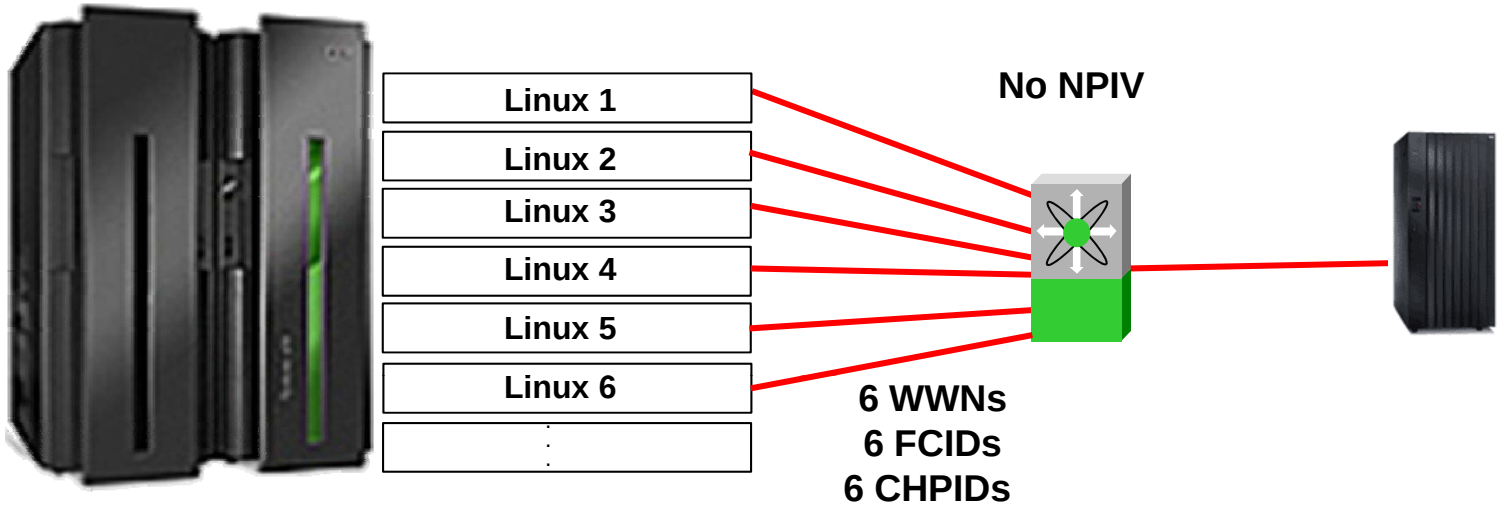


- Each Linux Guest / Machine has their own CHPID
  - Each has own WWN / FCID
  - Each can be zoned independently to protect for data Isolation
  - Wasteful of channel resources (ie. does this guest push 8G ?)
  - Limits the amount of consolidation to the System Z
  - Higher cost per image when considering consolidation

# What is NPIV?

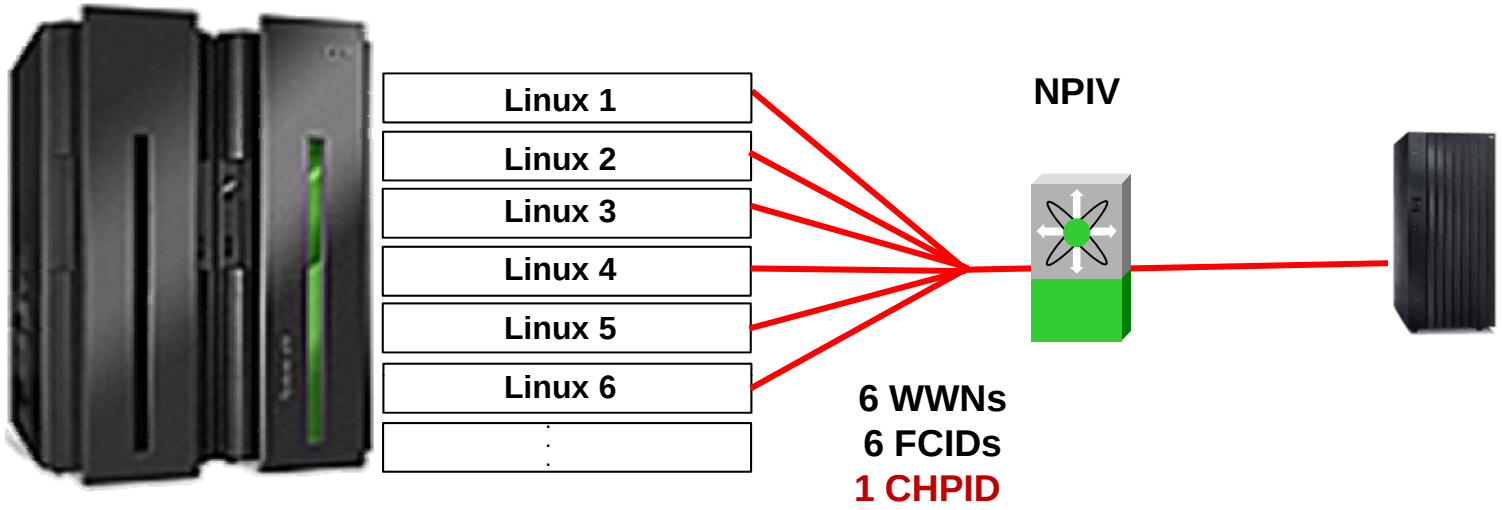


- All 6 Linux Guests / Machines share the same CHPID
  - Each has own WWN / FCID
  - Each can be zoned independently to protect for data Isolation
  - Good utilization of channel resources
  - Number of channels is no longer limiting factor to consolidation
  - Lower cost per image when considering consolidation

# How does NPIV work ?

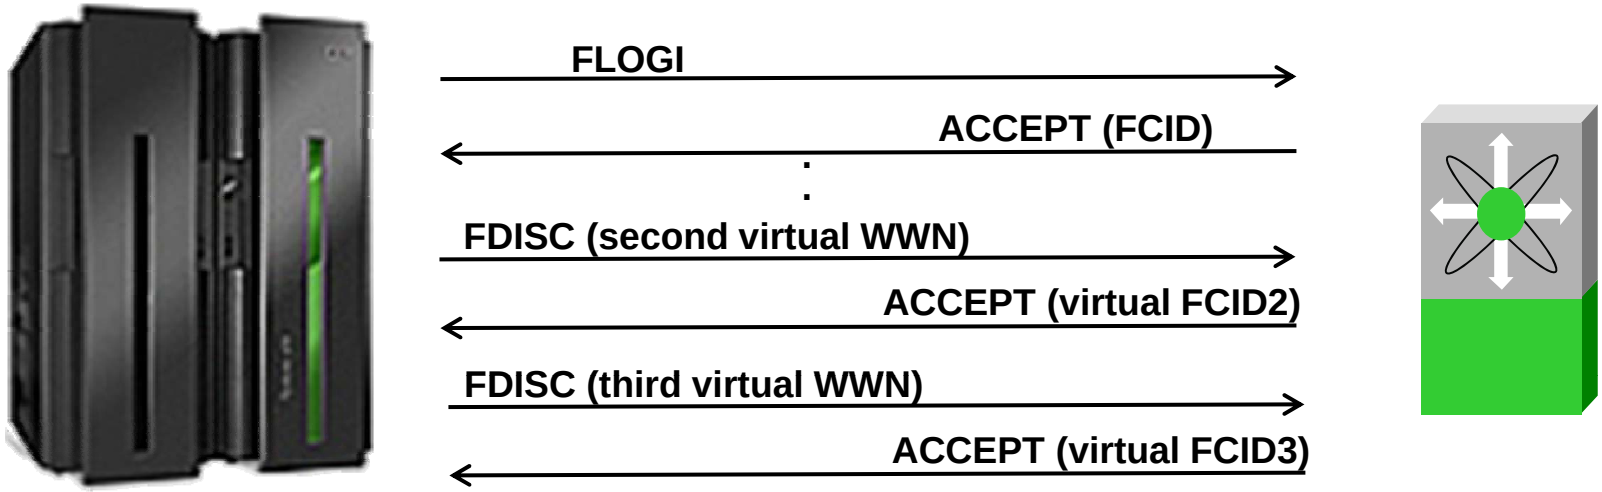


- Both System z and adjacent FC Director must me NPIV enabled
- System Z has a pool of virtual WWNs for each NPIV defined CHPID
- Switch will create unique FCID per FDISC
  - Based FCID will be 0xDDPP00 (DD = Domain PP = Port 00 is constant
  - NPIV FCIDs will be 0xDDPPxx (xx is 01, 02, 03 ….)
- Number of NPIV virtual connections per real is variable

# System z N-port ID Virtualization

**FC-FS 24 bit fabric addressing – Destination ID (D_ID)**

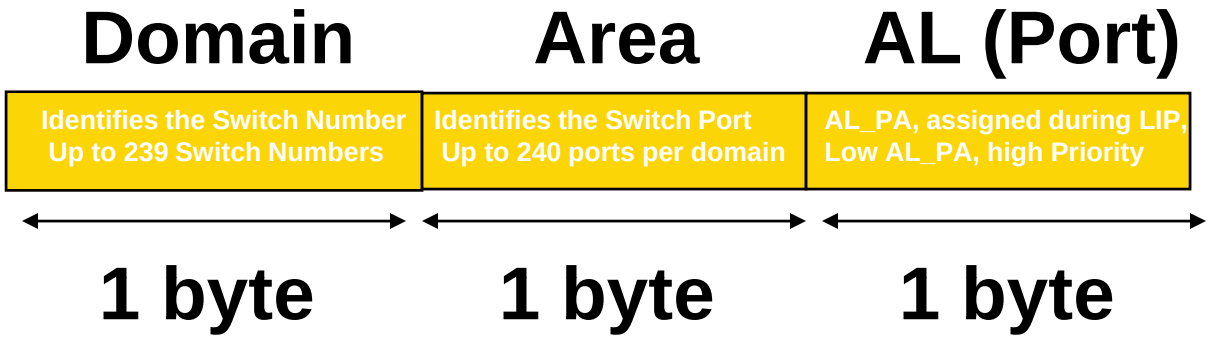


**FICON Express2, Express4 and Express 8 adapters now support NPIV**

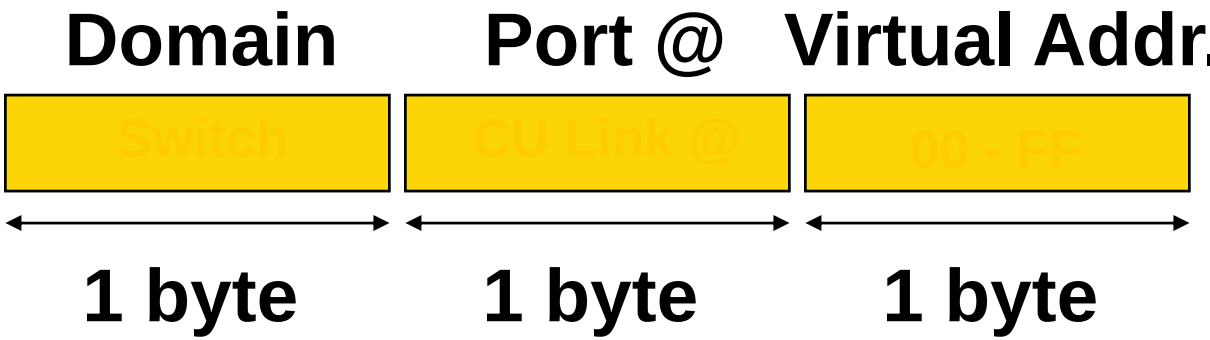

# Virtual Fabric (VSAN)

**A way to Partition a Switch or SAN into a Virtual/Logical environment**

- **Virtual SANs created from larger cost-effective redundant physical fabric**
- **Reduces wasted ports of the older “island” approach**
- **Hardware-based isolation**
- **Statistics can be gathered per VF**
- **Management per VF**
- **Unique Serial Number / CUP per FICON VF**

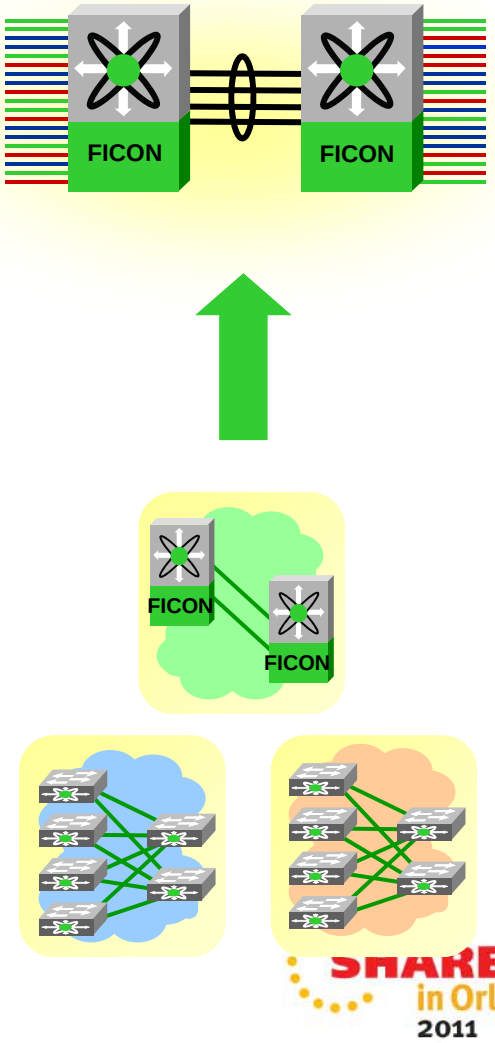

# Mixing FICON AND FCP

## Zoning

- Zoning is a method used with a FICON/FCP switching devices to enable or disable communication between different attached devices
- Zoning can be done by WWN, domain/index (sometimes called port zoning), or a combination of both
  - FCP typically uses WWN zoning
  - FICON typically uses D/I zoning
- A best-practice recommendation is to continue to segregate FICON devices in one zone and FCP devices in one or more other zones
- You would normally continue to use D/I zoning for FICON while using WWN for FCP traffic zoning even on the same switching device and fabric

# Zoning

- A logical grouping of fabric connected devices within a SAN (or virtual fabric)
- Zoning establishes access control
  - Devices within a zone can access each other
- Zoning increases security
  - Limiting access prevents unauthorized access
- Zone membership might be configured by:
  - Port World Wide Name (pWWN)—device
  - Fabric World Wide Name (fWWN)—fabric
  - Fibre Channel Identifier (FCID)
  - Fibre Channel Alias (FC_Alias)
  - IP address
  - Domain ID/port number
  - Interface

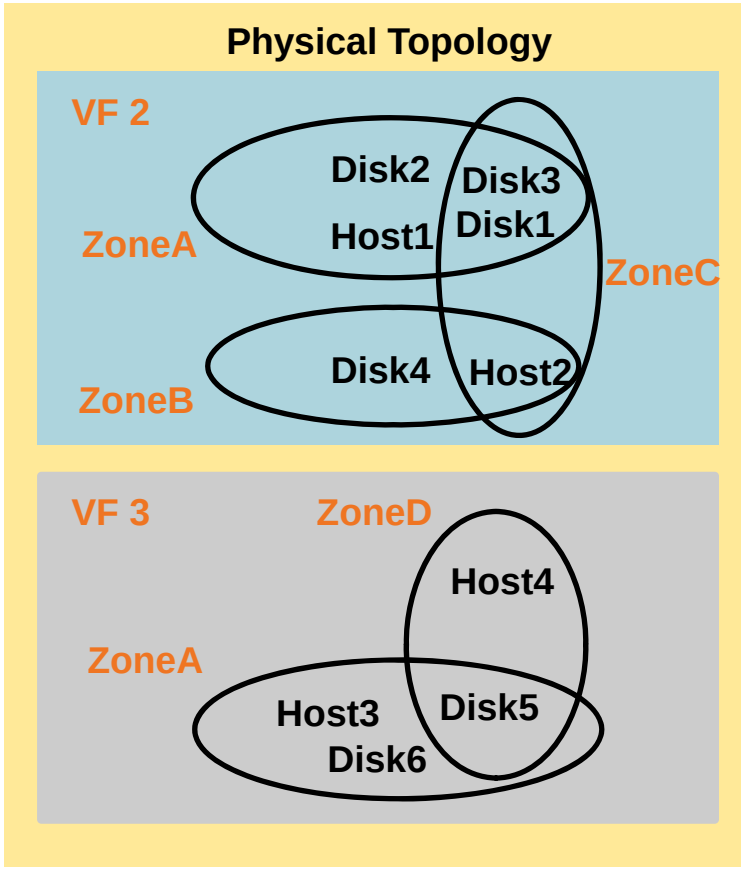

# Mixing FICON AND FCP: Zoning

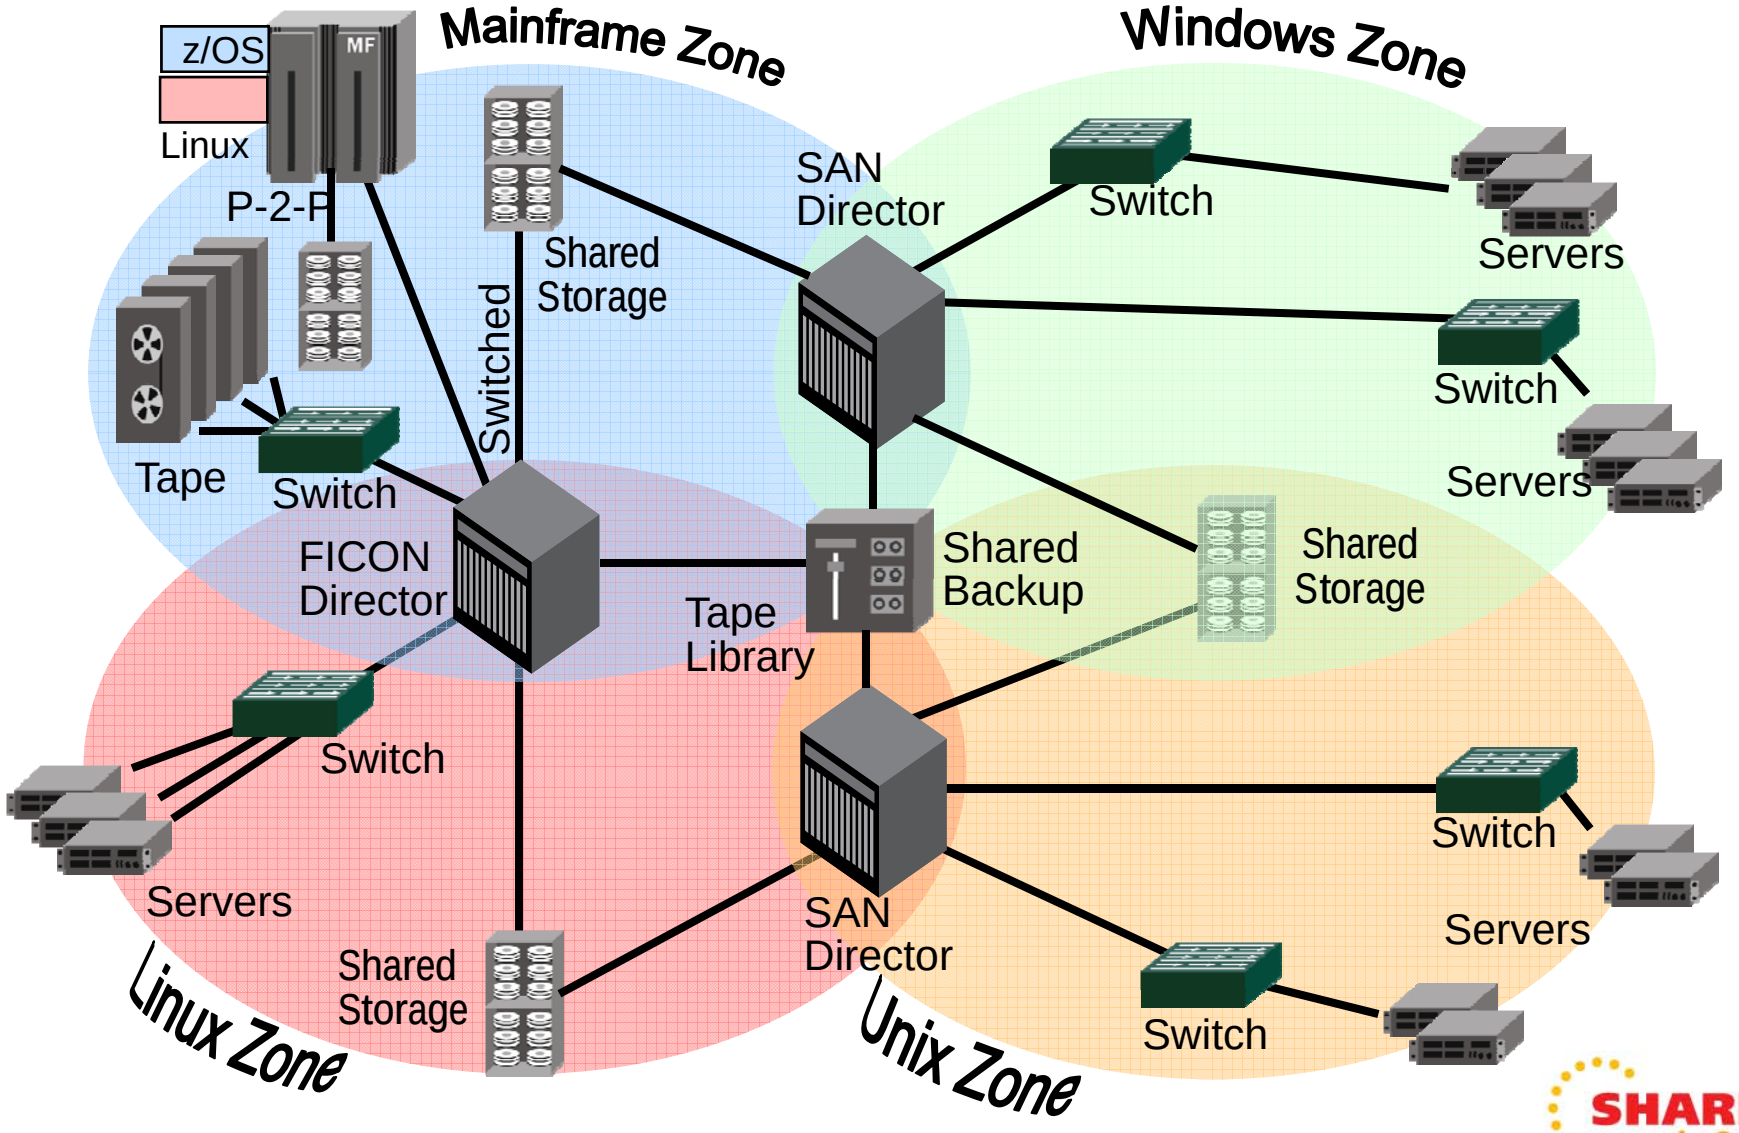

# Fabric Binding for Enhanced Cascading Security

- Two Switches / One Hop
- Based on Switch WWNs
- Only authorized switches can connect to a secure fabric
  - Unauthorized switches result in attachment port being placed in 'Invalid Attachment' state
  - Query Security Attributes and Exchange Security Attributes ensure compliance
- Predictable error recovery
- Requires Insistent (static) Domain IDs

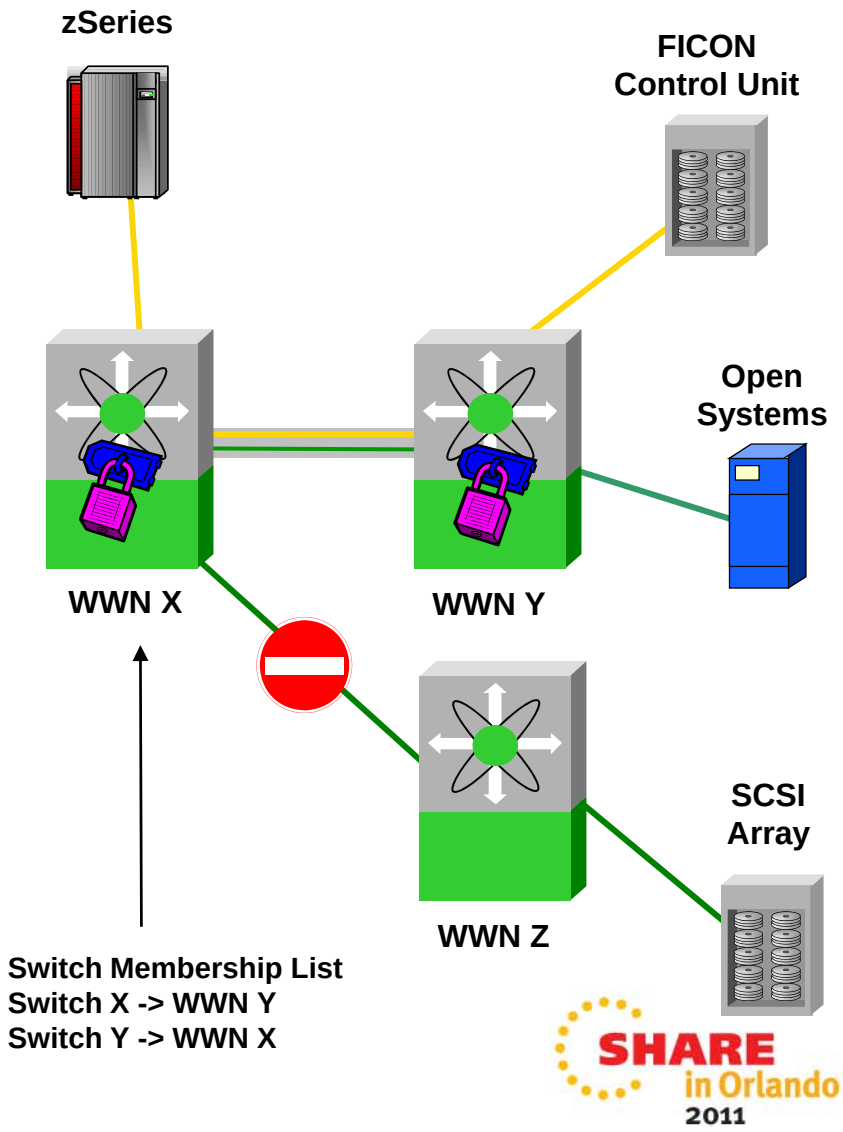

# FICON Cascade Topologies

- Only One “Hop” is allowed for FICON
  - “Multi-hop” is not supported but does work – testing and support are why not supported
  - FCIP links are supported
  - Port Channels are supported for FICON

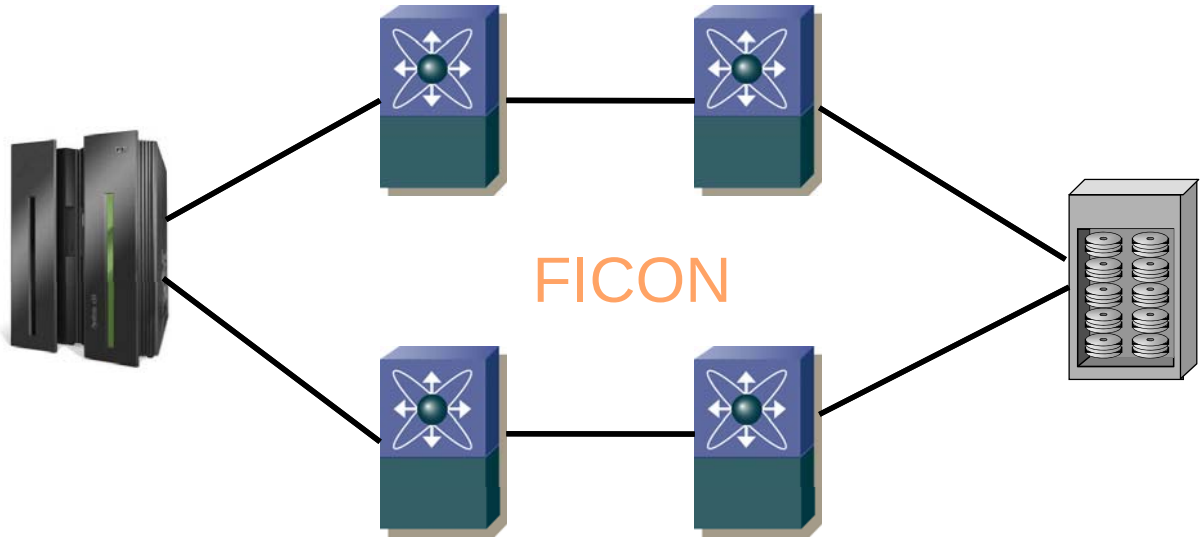

# FCP Topologies

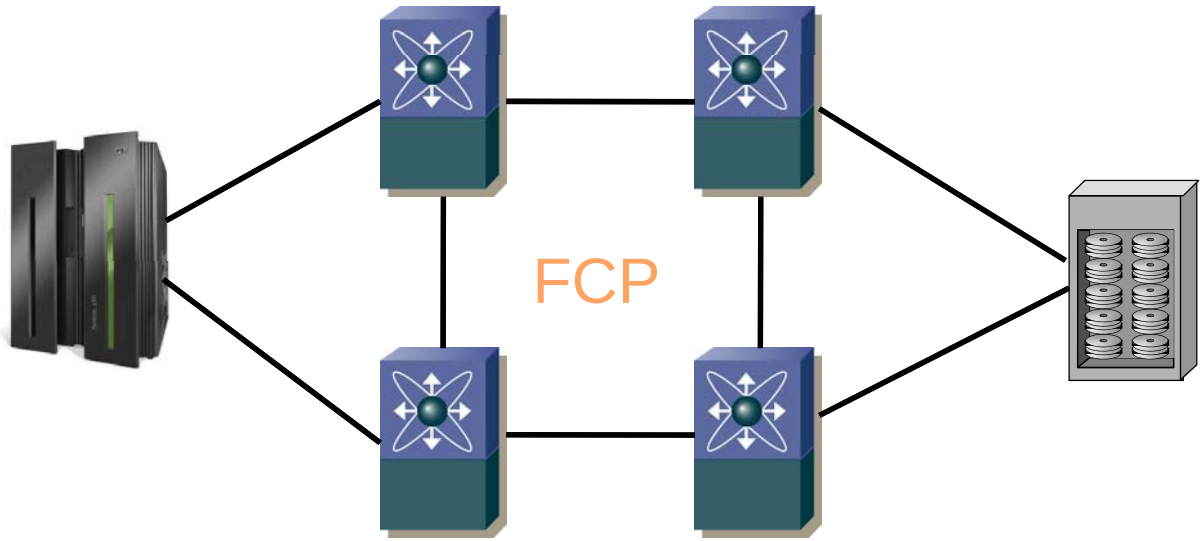

# VSAN / Virtual Fabric Based Roles

**System Administrator**
Configures and manages all platform-specific capabilities

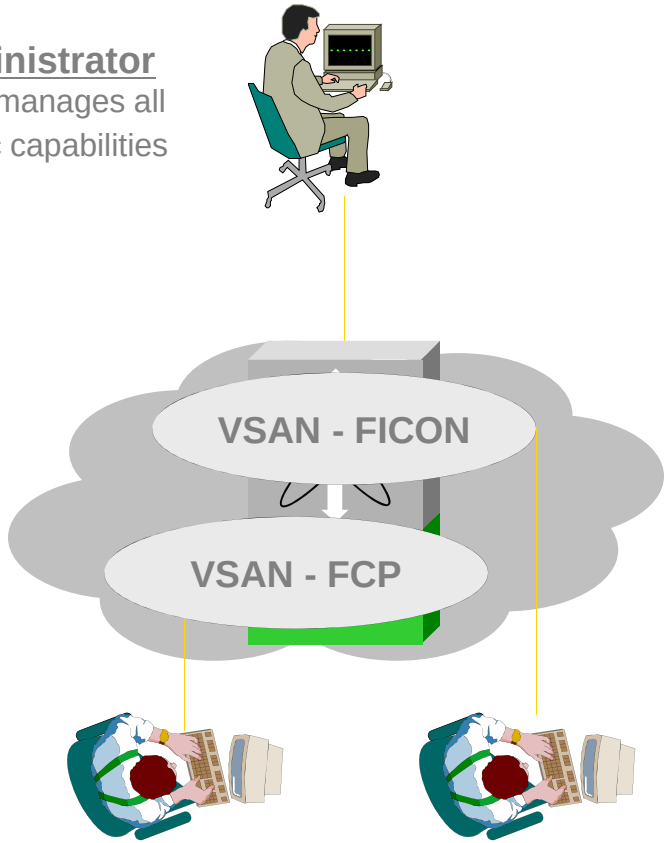


**VSAN Administrators**
Configure and manages only their VSANs

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

# CUP in Consolidated Fabric

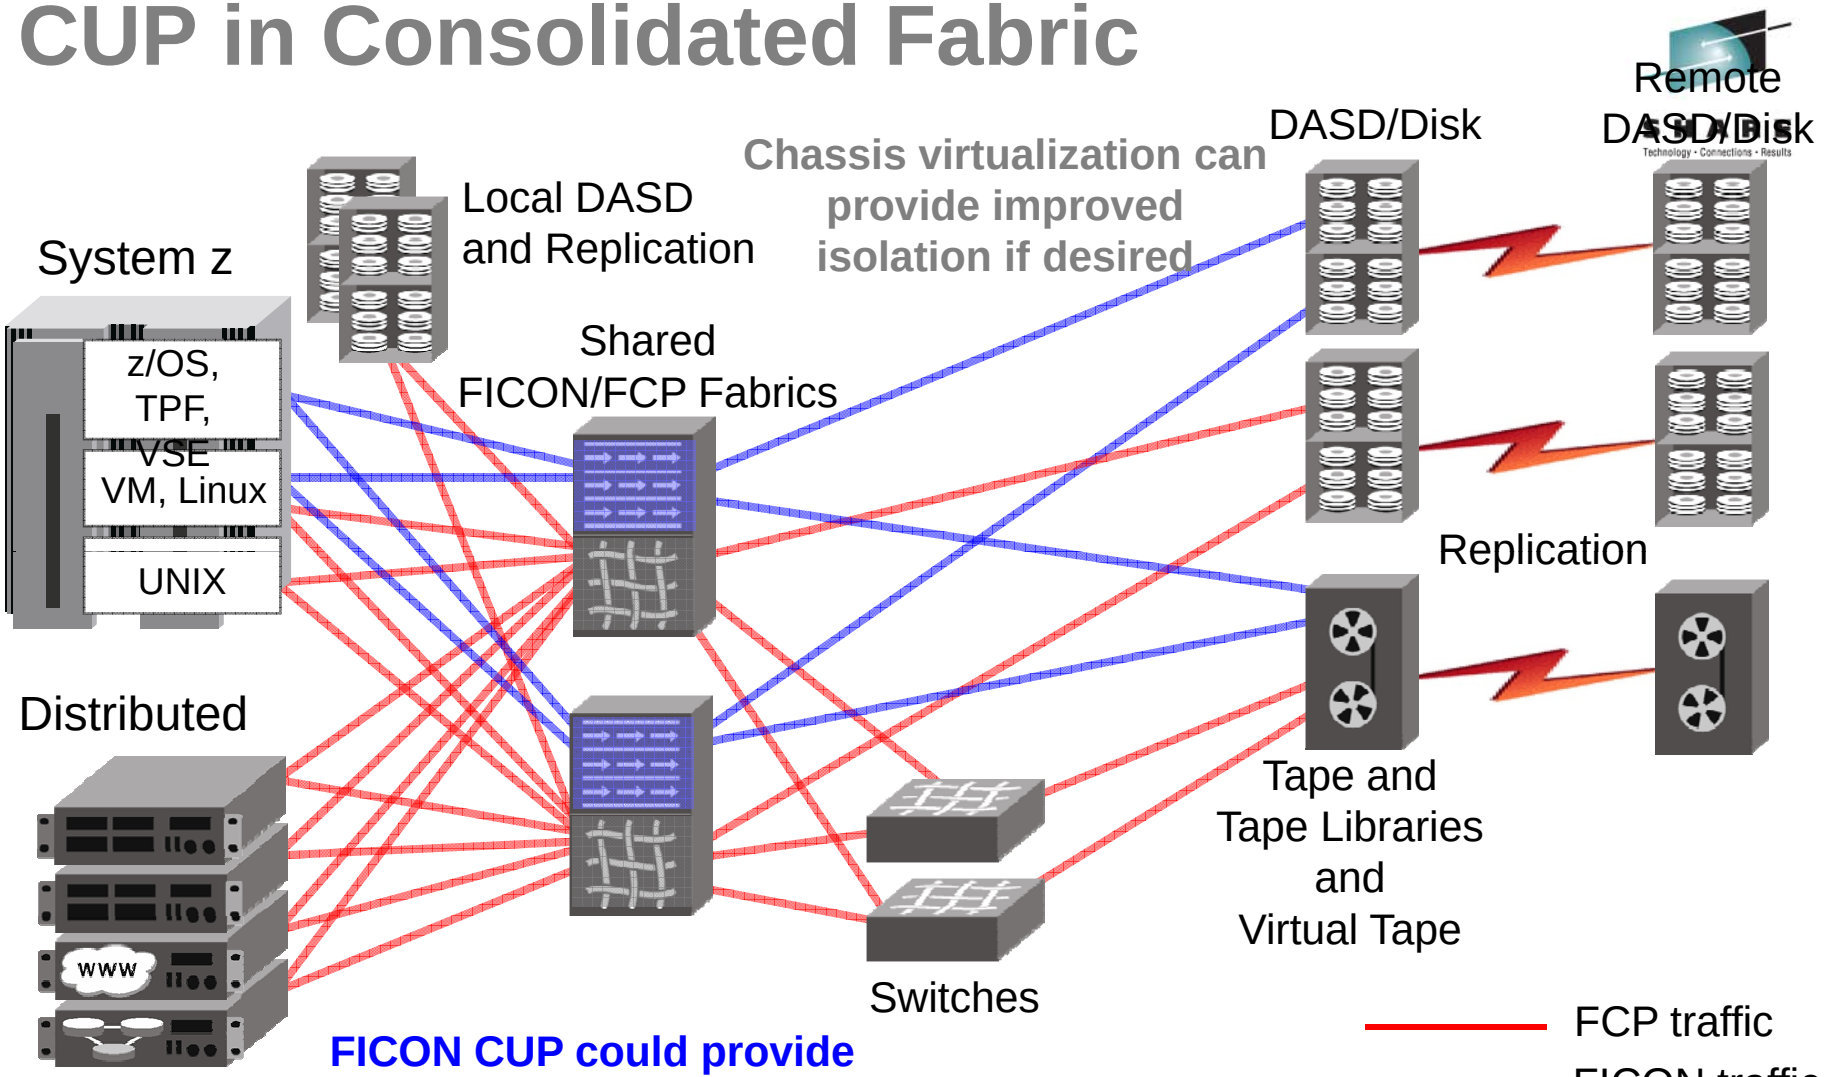


**FICON CUP could provide I/O statistics for both System z and distributed platforms**

SHARE in Orlando 2011

## SHARE, Orlando, August 2011

System z FICON and FCP Fabrics – Intermixing Best Practices

Session 9864

# THANK YOU!

# Standards and NPIV

- FC-FS
  - Describes FDISC use to allocate additional N_Port_IDs
    - Section 12.3.2.41
  - NV_Ports are treated like any other port
    - Exception is they use FDISC instead of FLOGI
- FC-GS-4
  - Describes
    - Permanent Port Name and Get Permanent Port Name command
      - *Based on the N_Port ID (G_PPN_ID)*
    - The PPN may be the F_Port Name
- FC-LS
  - Documents the responses to NV_Port related ELSs
    - FDISC, FLOGI and  FLOGO
    - Reference 03-338v1

# More Standards on NPIV

- FC-DA
  - Profiles the process of acquiring additional N_Port_IDs
    - Clause 4.9
- FC-MI-2
  - Profiles how the fabric handles NPIV requests
    - New Service Parameters are defined in 03-323v1
    - Name Server Objects in 7.3.2.2 and 7.3.2.3

# Virtualizing the Fabric – The Full Solution

To build a cost saving fabric virtualization solution, 7 key services are required:

- Virtual Fabric Attachment – the ability to assign virtual fabric membership at the port level
- Multiprotocol Extensions – the ability to extend virtual fabric service to iSCSI, FCIP, FICON, etc.
- Virtual Fabric Services – the ability to create fabric services per virtual fabric (Login, Name, RSCNs, QoS, etc.)
- Virtual Fabric Diagnostics – the ability to troubleshoot per virtual fabric problems
- Virtual Fabric Security – the ability to define separate security policies per virtual fabric
- Virtual Fabric Management – the ability to map and manage virtual fabrics independently
- Inter-Fabric Routing – the ability to provide connectivity across virtual fabrics – *without merging the fabrics*

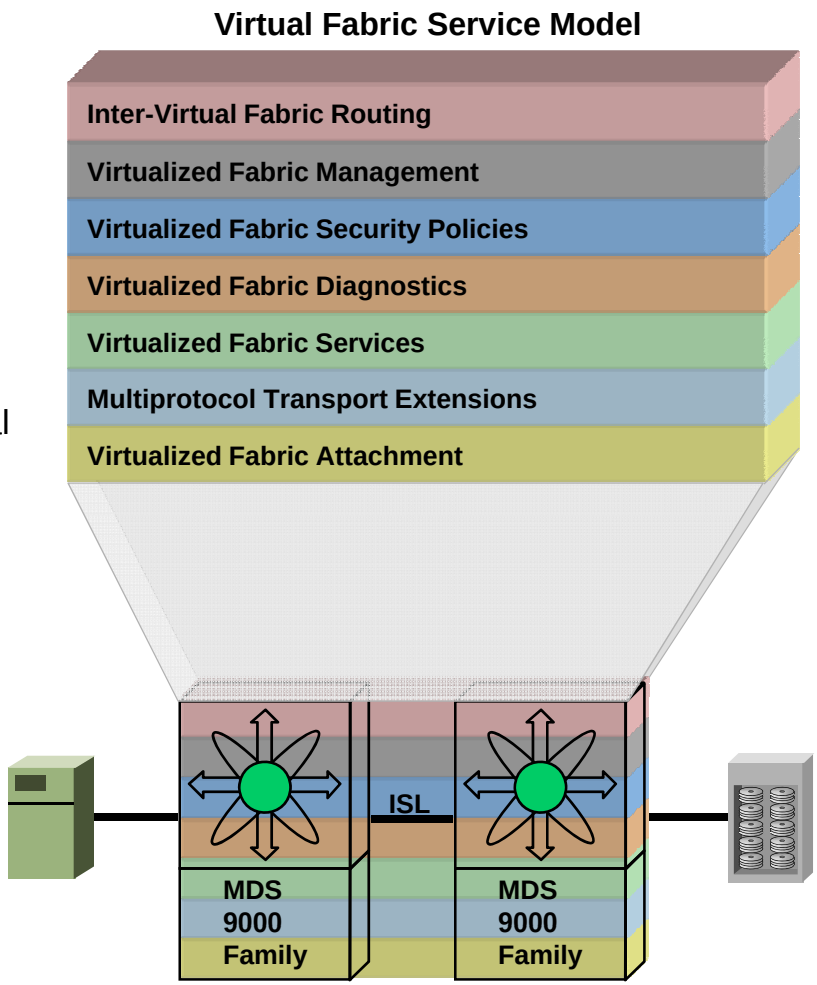


**Full Service End-to-End Virtual Fabric Implementation**


# THIS SLIDE INTENTIONALLY LEFT BLANK