BROCADE

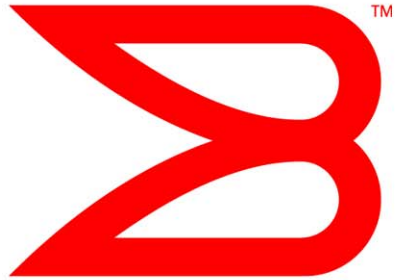

The IBM System z9 and Virtualized Storage:

Consolidation Drives Virtualization in Storage Networks

Stephen R. Guendert, PhD(cand)
Principal Engineer
Brocade

26 November 2007

UK CMG Forum

Agenda

- Fundamentals
  - Forces driving virtualization
- Concepts
  - Consolidation Drives Virtualization in Storage Networks
- Application
  - Fibre Channel Improves Utilization and Scalability
- References
  - Articles and Standards

Introduction

- Explore the fundamentals of virtualization
  - Server, fabric, link, and address/access
    - Consolidation Drives Virtualization in Storage Networks
      - December 2006 / January 2007
    - Fibre Channel Improves Utilization and Scalability
      - February / March 2007
- Examine fundamentals in theory
  - Server
    - N_Port Identifier Virtualization (NPIV)
  - Fabric
    - Virtual Fabric Identifier
    - Inter-fabric Routing
- Apply theory to a practical application

BROCADE

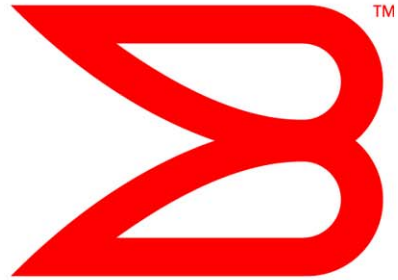

Fundamentals

NPIV

Virtual Fabrics

Frame Tagging

Translation and Routing

What's Driving Server Virtualization?

- zSeries support of zLinux
  - Mainframe expanded to address open system applications
  - Linux promoted as alternative to Unix
  - Mainframe OS virtualization benefits
    - Availability, serviceability, scalability, flexibility
- Initial constraints
  - FCP requests are serialized by the OS
    - FCP header doesn't provide image addresses
    - FICON SB2 header provides additional addressing
  - Channel ports are underutilized
  - Resulting cost / performance benefit is not competitive

Technology Behind Server Virtualization

- N_Port Identifier Virtualization (NPIV)
  - N_Port becomes virtualized
    - Supports multiple images behind a single N_Port
  - N_Port requests more than one FCID
    - FLOGI provides first address
    - FDISC provides additional addresses
  - All FCID's associated with one physical port

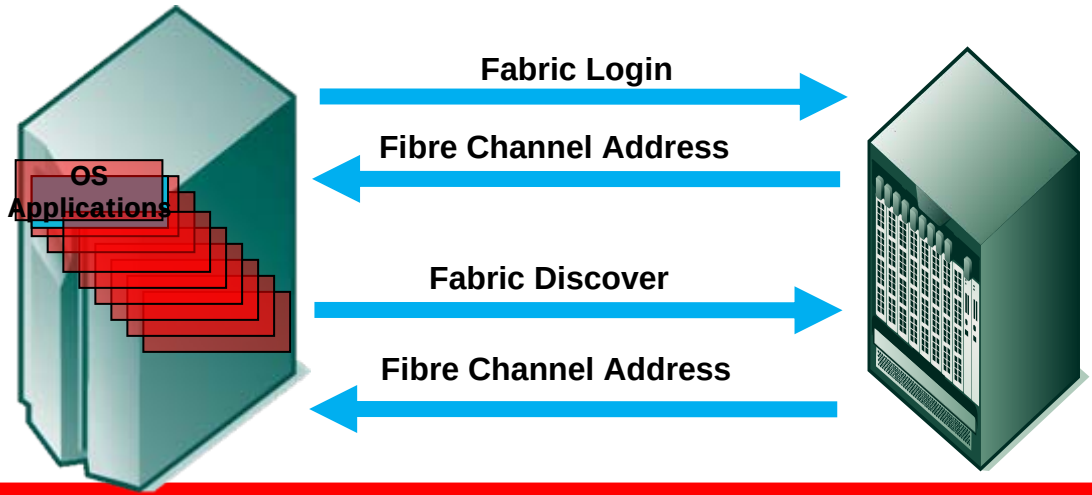

What's Driving Fabric Virtualization?

- SAN Sprawl
  - Organic growth of SANs is creating large physical SAN infrastructures
  - The need to merge data centers produces larger SANs
  - Acquisition of data centers forces SAN expansion
- Controlling the sprawl motivates virtualization
  - Simplified management
  - Local administration
  - Access to centralized services

Technology Behind Fabric Virtualization

- Virtual Fabric Identifier (VFID)
  - Fabric becomes virtualized
    - Supports multiple common domains on the same switch
- Inter-fabric Routing (IFR)
  - Fabric routing becomes virtualized
    - Address Translation supports multiple overlapping device addresses
- Virtual Fabric Configuration
  - Administrative domains
    - Utilizes existing zoning enforcement techniques to create virtual fabrics
  - Virtual Storage Area Networks
    - Utilizes expanded addressing to create virtual fabrics
  - Partitions
    - Utilizes frame tagging to create virtual fabrics and virtual links

Expanded Fibre Channel Addressing

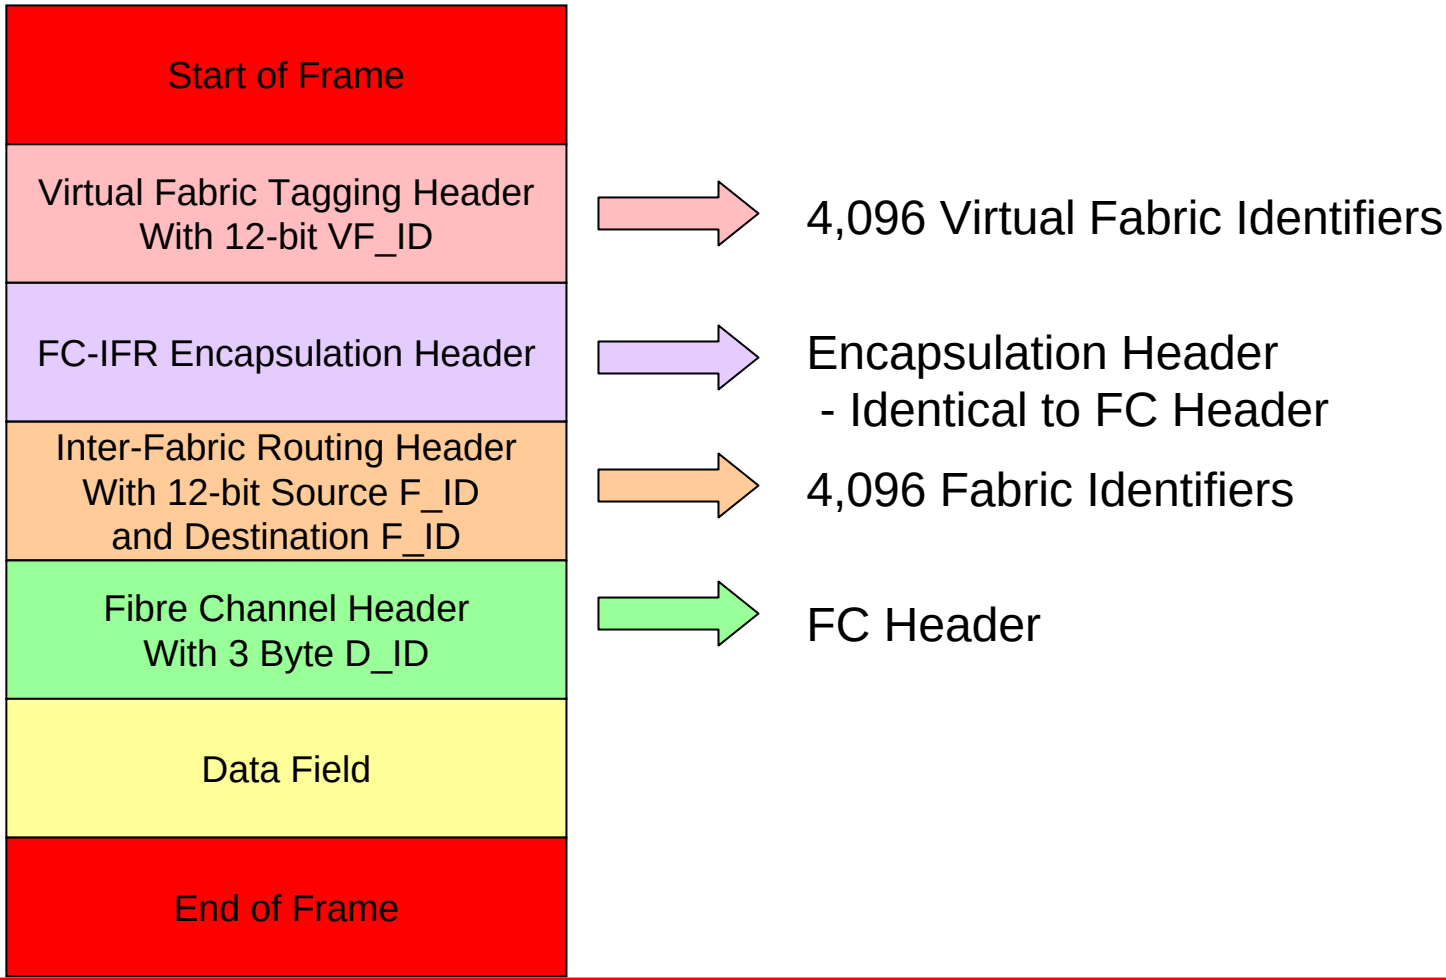

BROCADE

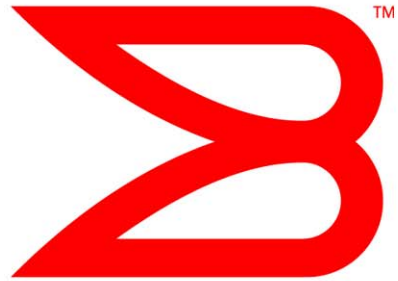

Concepts

Consolidation Drives Virtualization in Storage Networks

Example: Consolidation through Virtualization

- Applications
  - Dual production fabrics
    - Four 64-port Directors
    - Cascaded across two sites
  - Backup fabric
    - Two 24-port Switches
    - Cascaded across two sites
- Devices
  - Open System Servers and Storage
  - Mainframe Server and Storage

Dispersed Storage Network

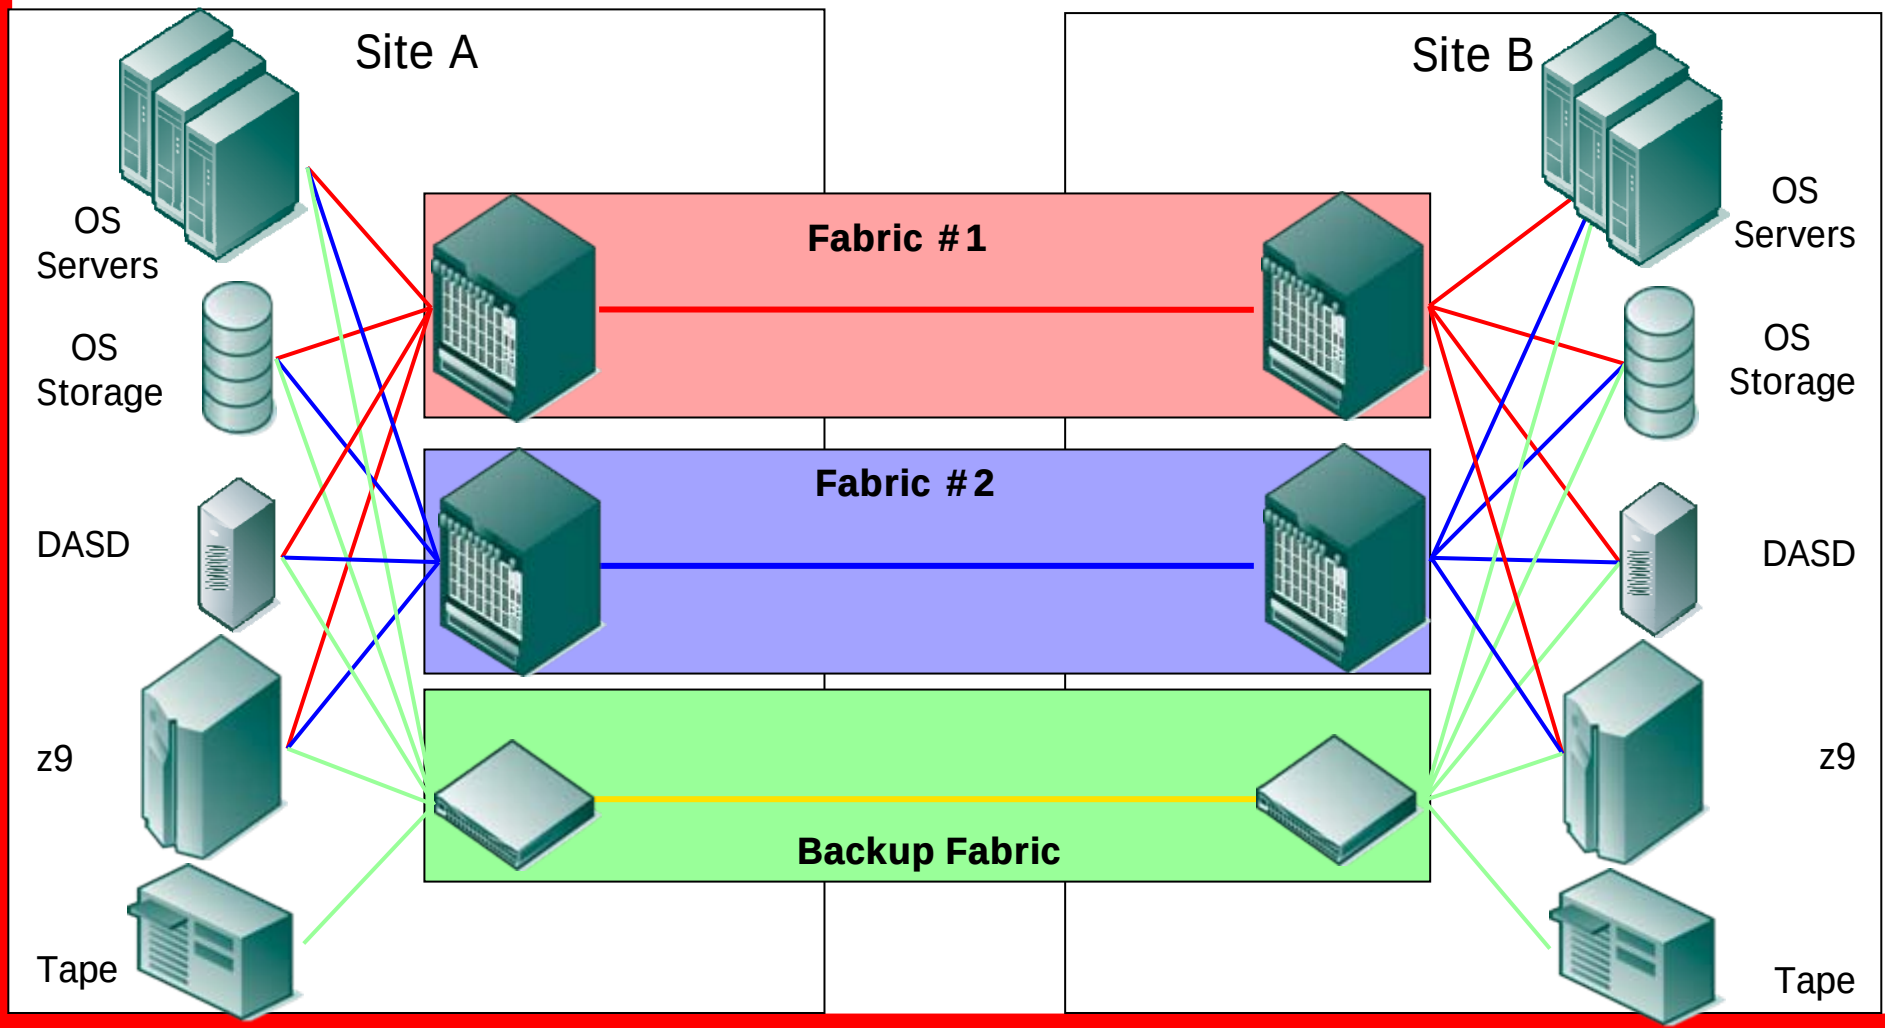

Distributed Storage Network Port Count and Utilization Rate

	FABRIC	PORT COUNT	PORTS USED	LINK SPEED (GFC)	AVERAGE THROUGHPUT (MB/s)	UTILIZATION RATE (%)
Site 1 Switch	1	64	48	1	42	42
	2	64	34	2	36	18
	Backup	24	23	2	28	14
2 Switch	1	64	43	1	25	25
	2	64	28	2	36	18
	Backup	24	23	2	18	9
Total			199		30	20

Server Consolidation through N_Port Identifier Virtualization

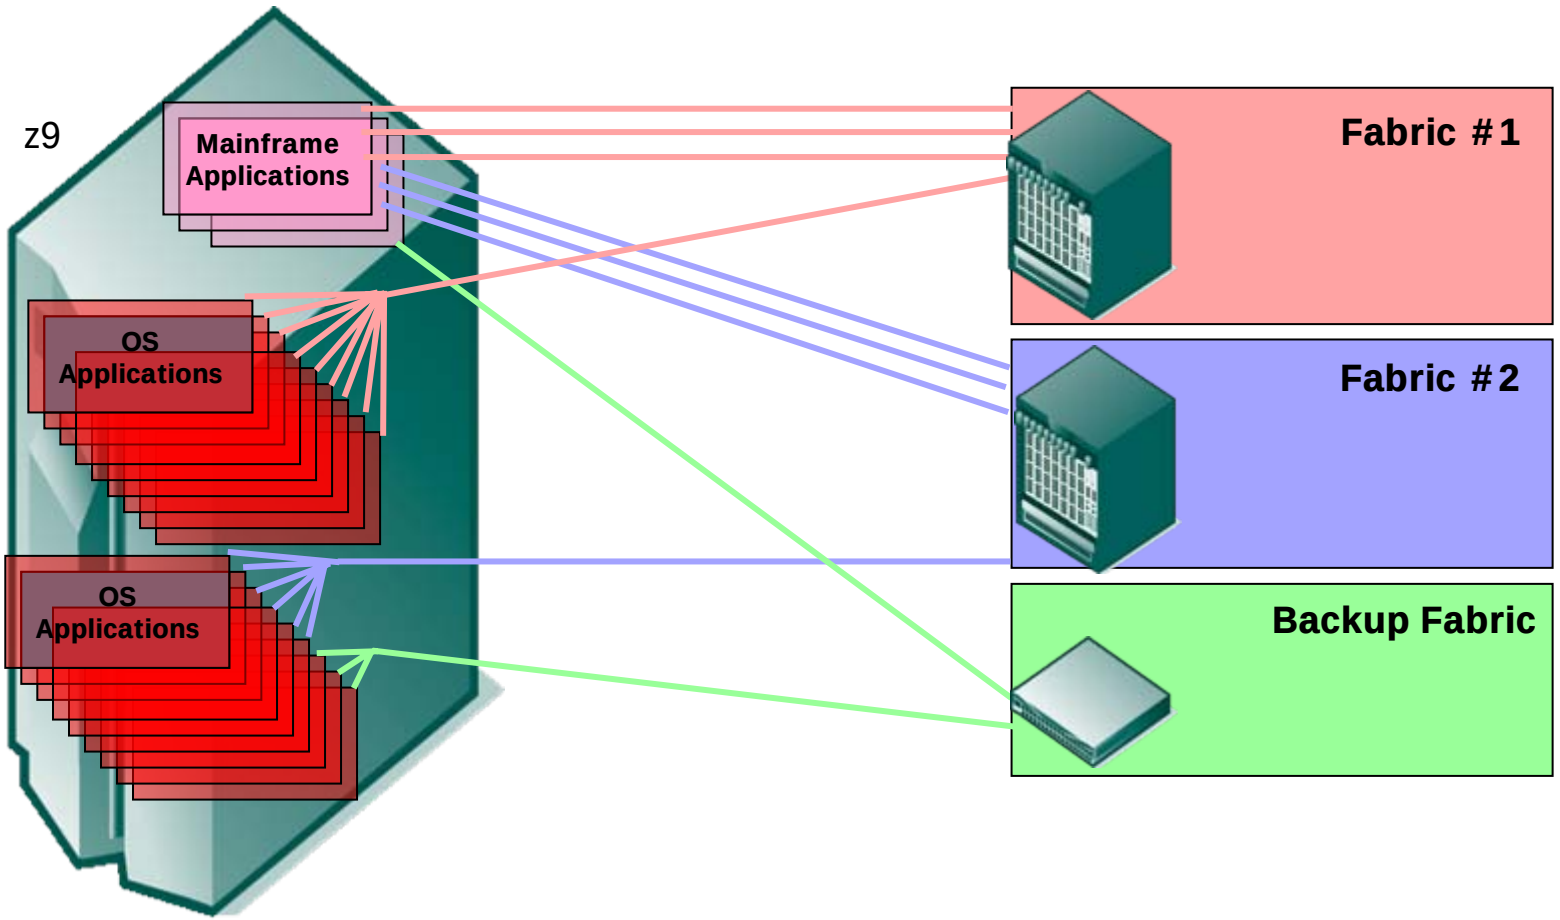

Consolidated Servers

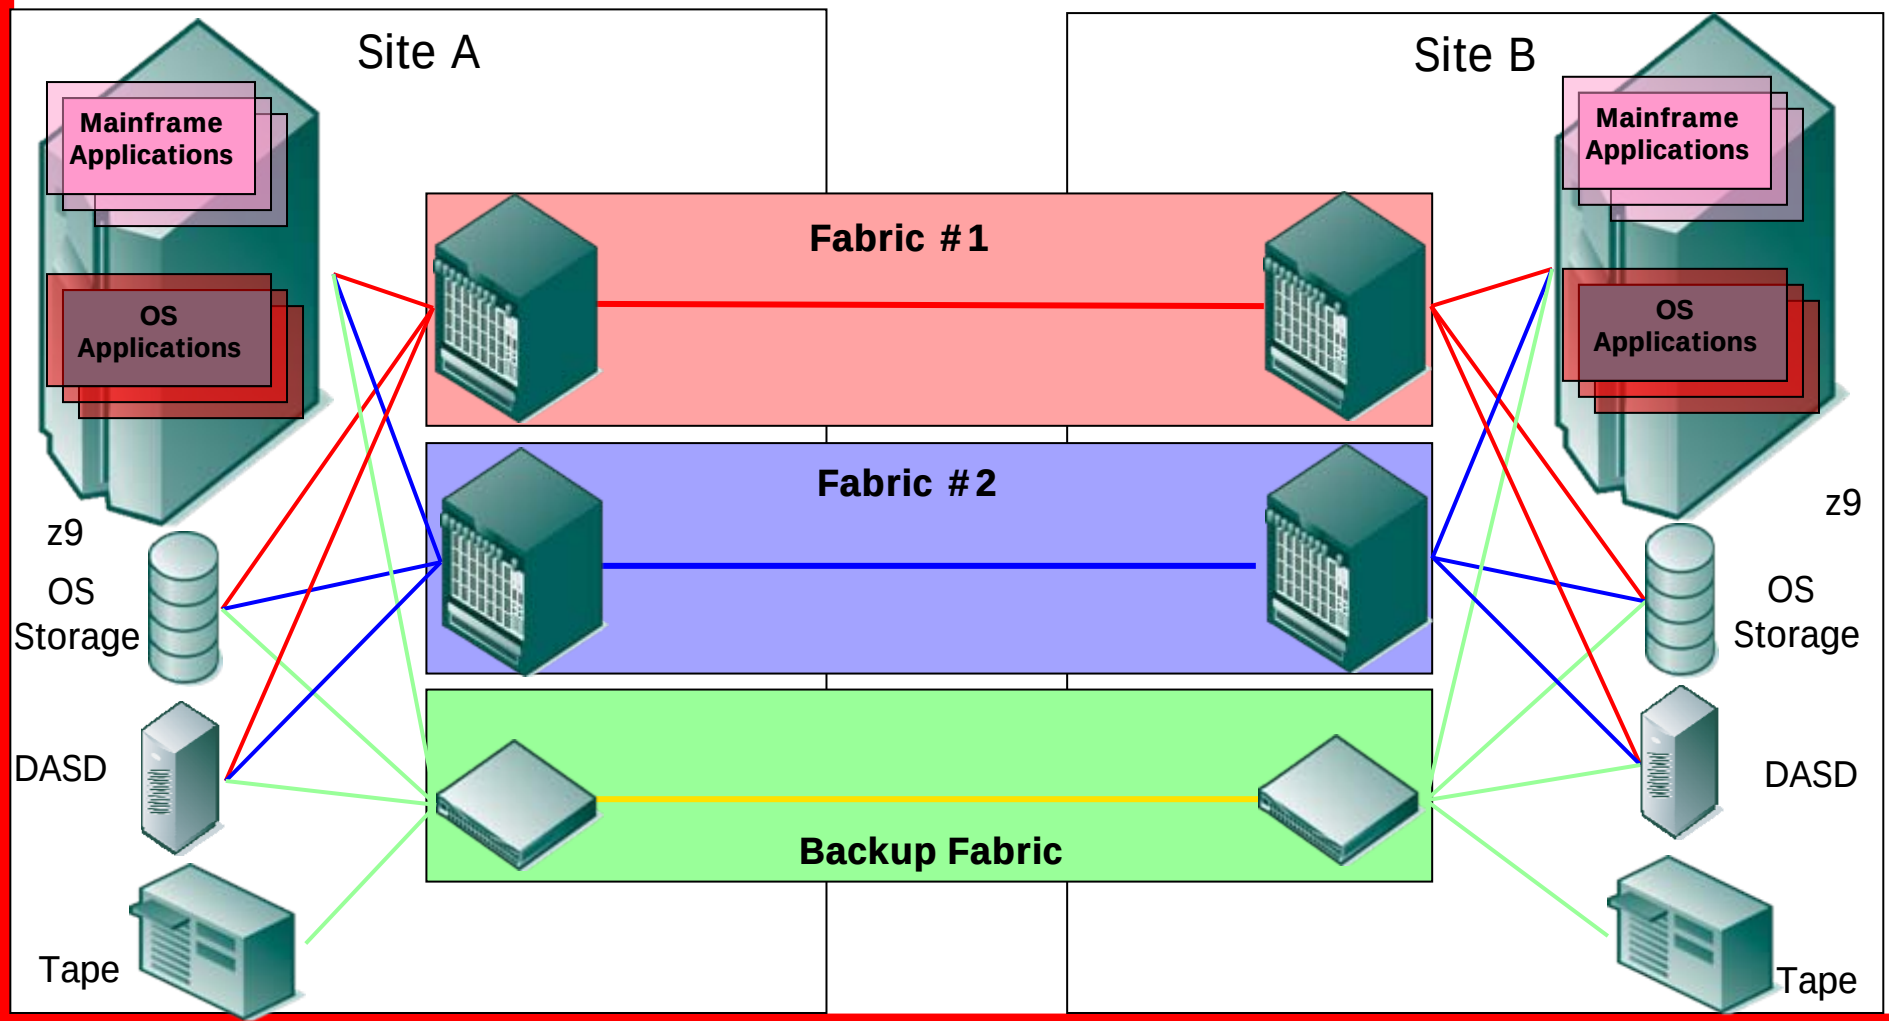

Virtualized Storage Network

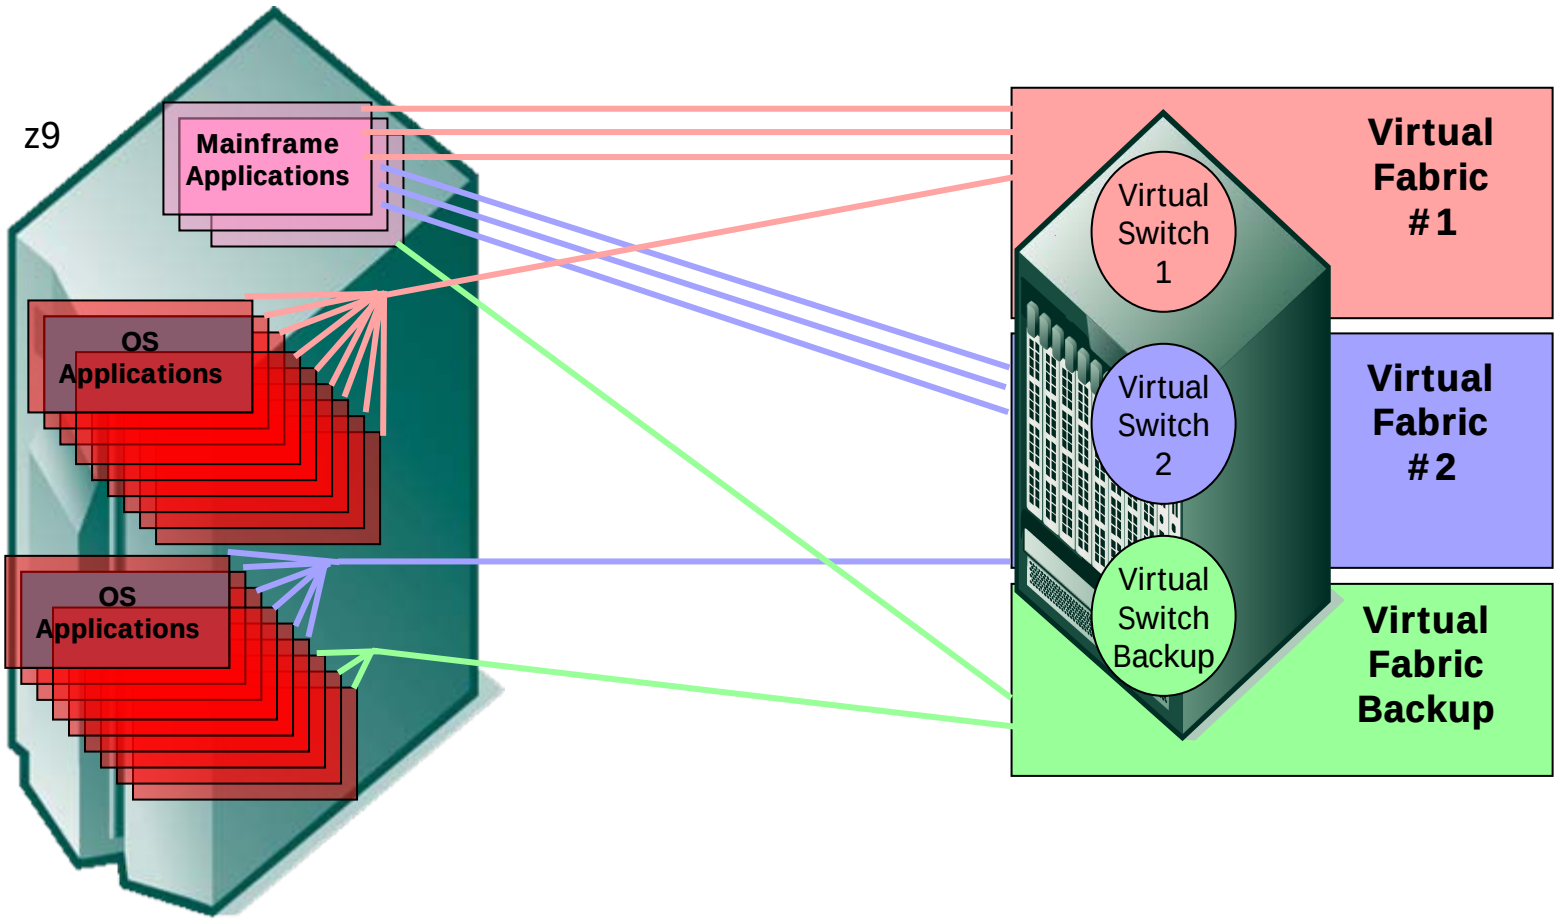

Consolidated Storage Network

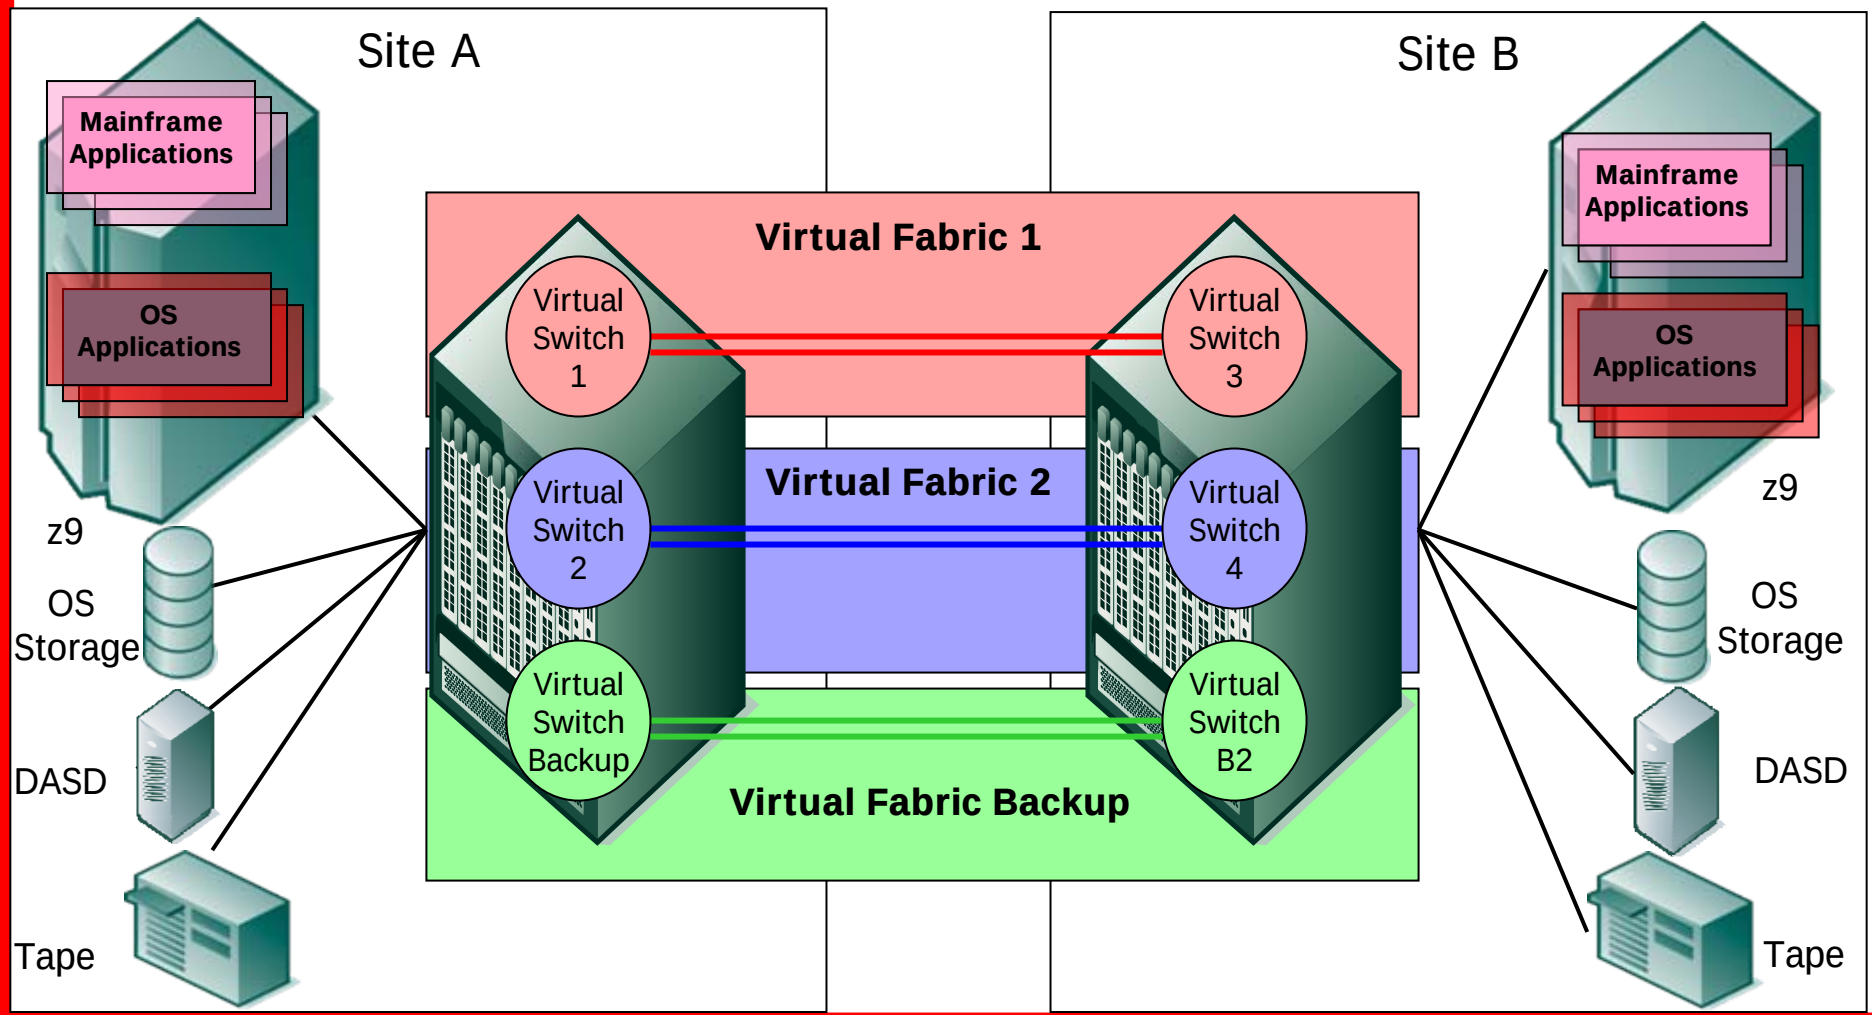

Consolidated Storage Network Port Count and Utilization Rate

	FABRIC	PORT COUNT	PORTS USED	LINK SPEED (GFC)	AVERAGE THROUGHPUT (MB/s)	UTILIZATION RATE (%)
Site 1 Switch (256 Port Director)	1	Up to 256	26	4	288	72
	2	Up to 256	18	4	226	56
	Backup	Up to 256	16	4	166	41
Site 2 Switch (256 Port Director)	1	Up to 256	28	4	313	78
	2	Up to 256	22	4	265	66
	Backup	Up to 256	17	4	156	39
Total			127		247	61

BROCADE

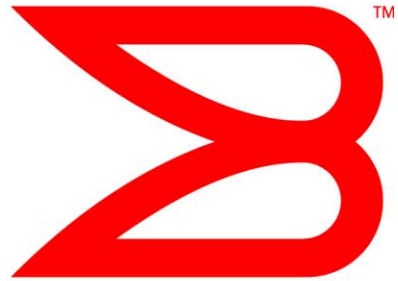

Application

Fibre Channel Improves Utilization and Scalability

Consolidated Storage Network

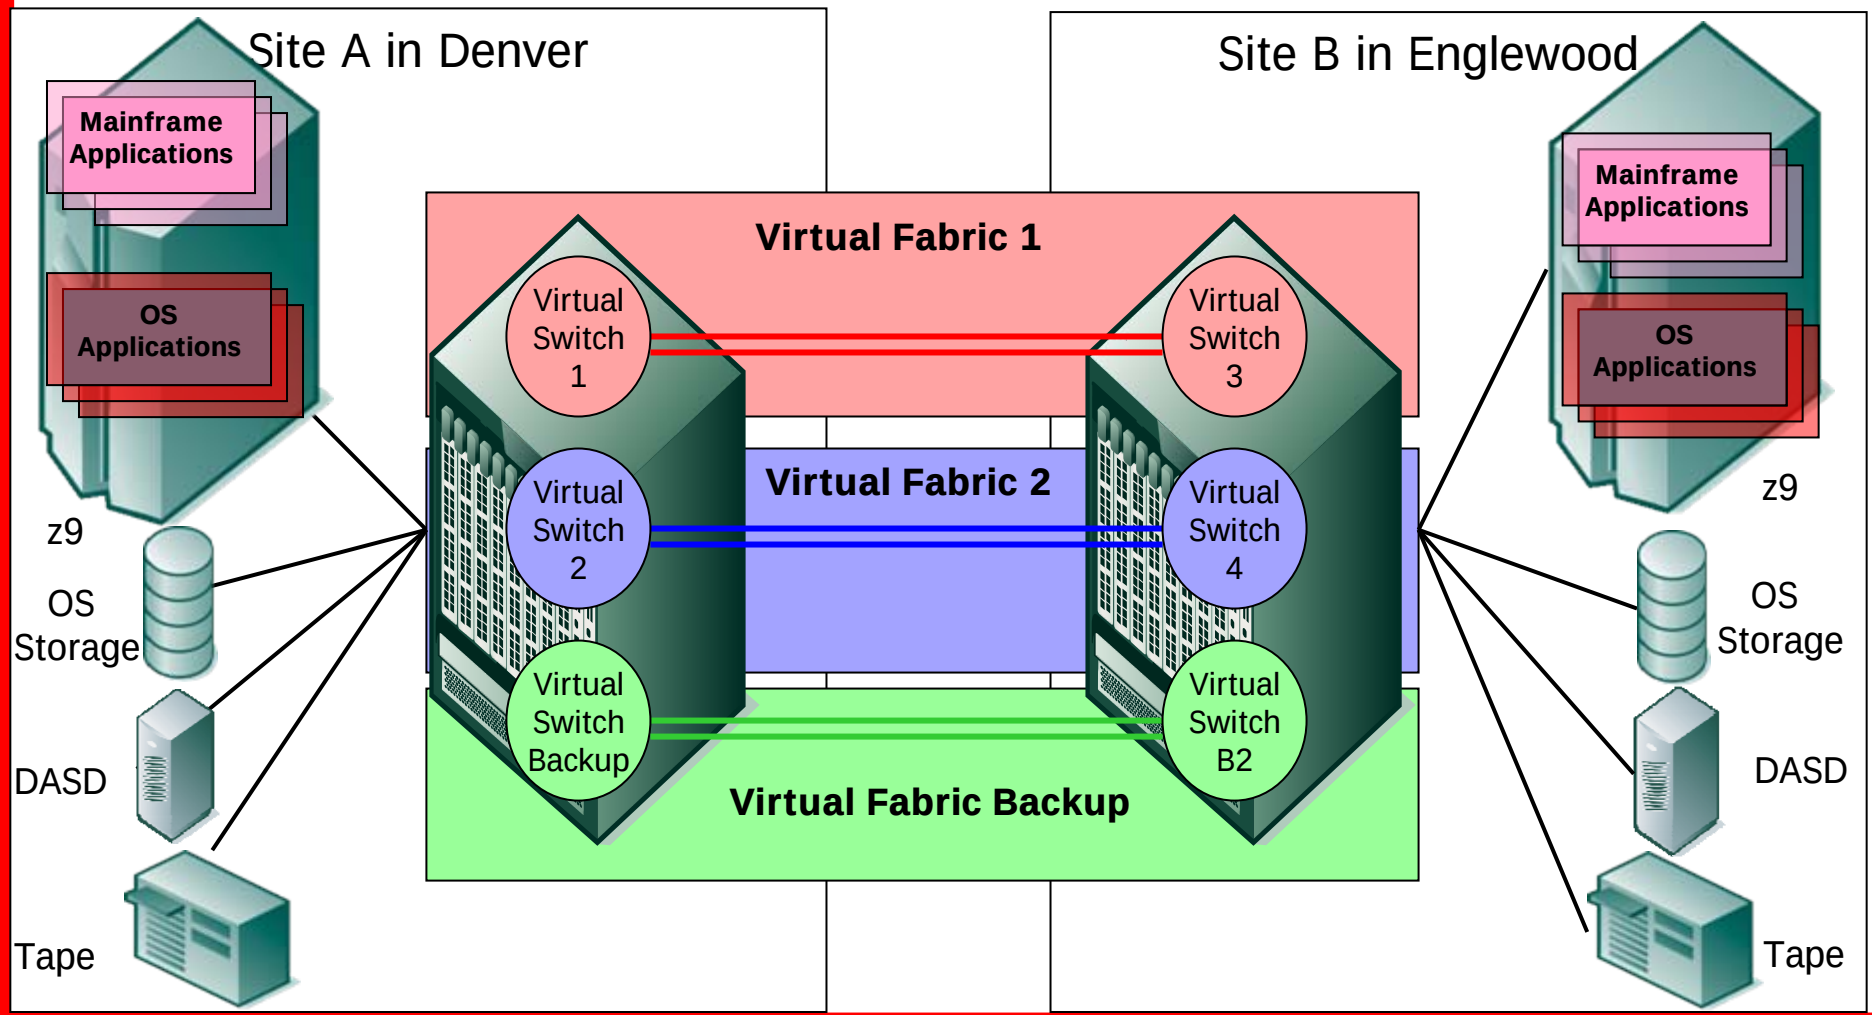

Virtual Fabric Tagging

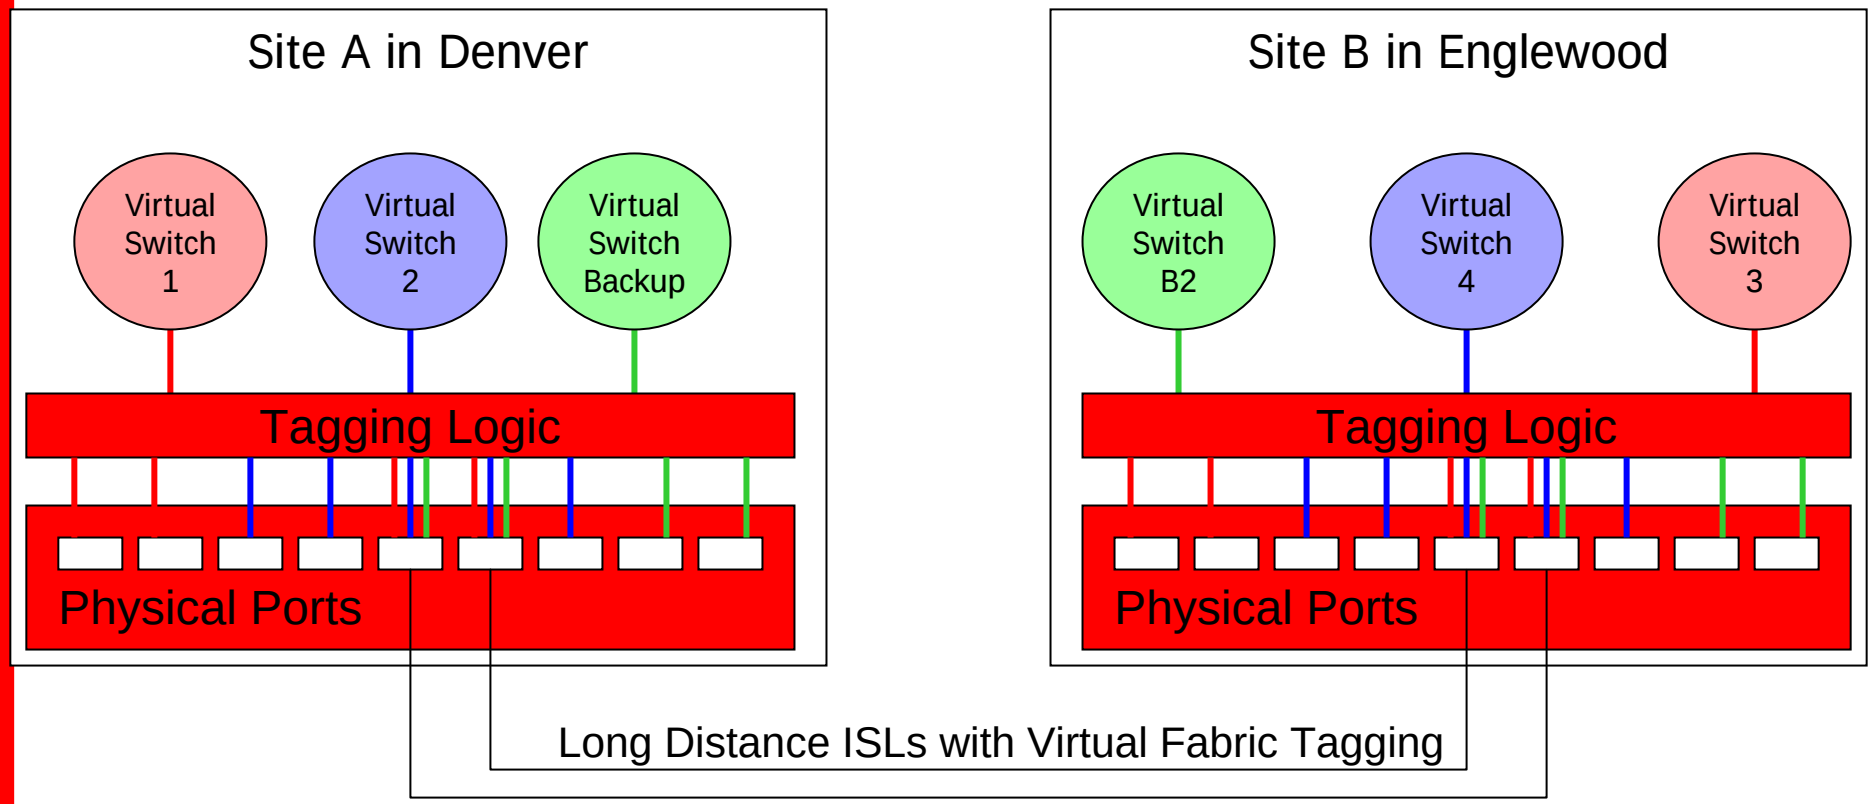

Virtual Fabric Tagging

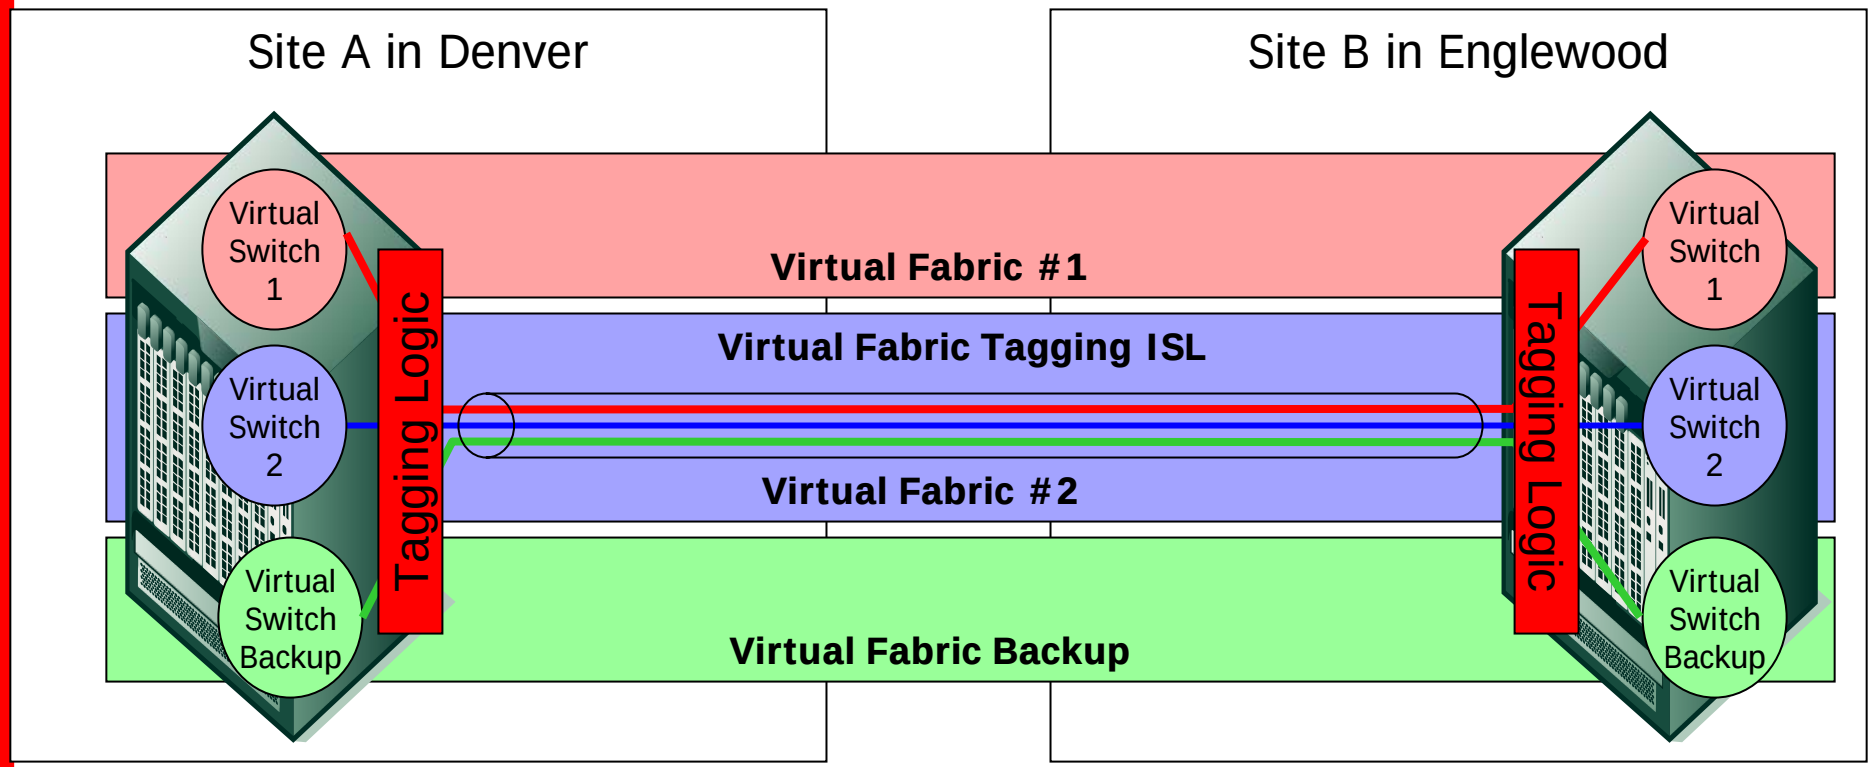

Expanding Data Centers

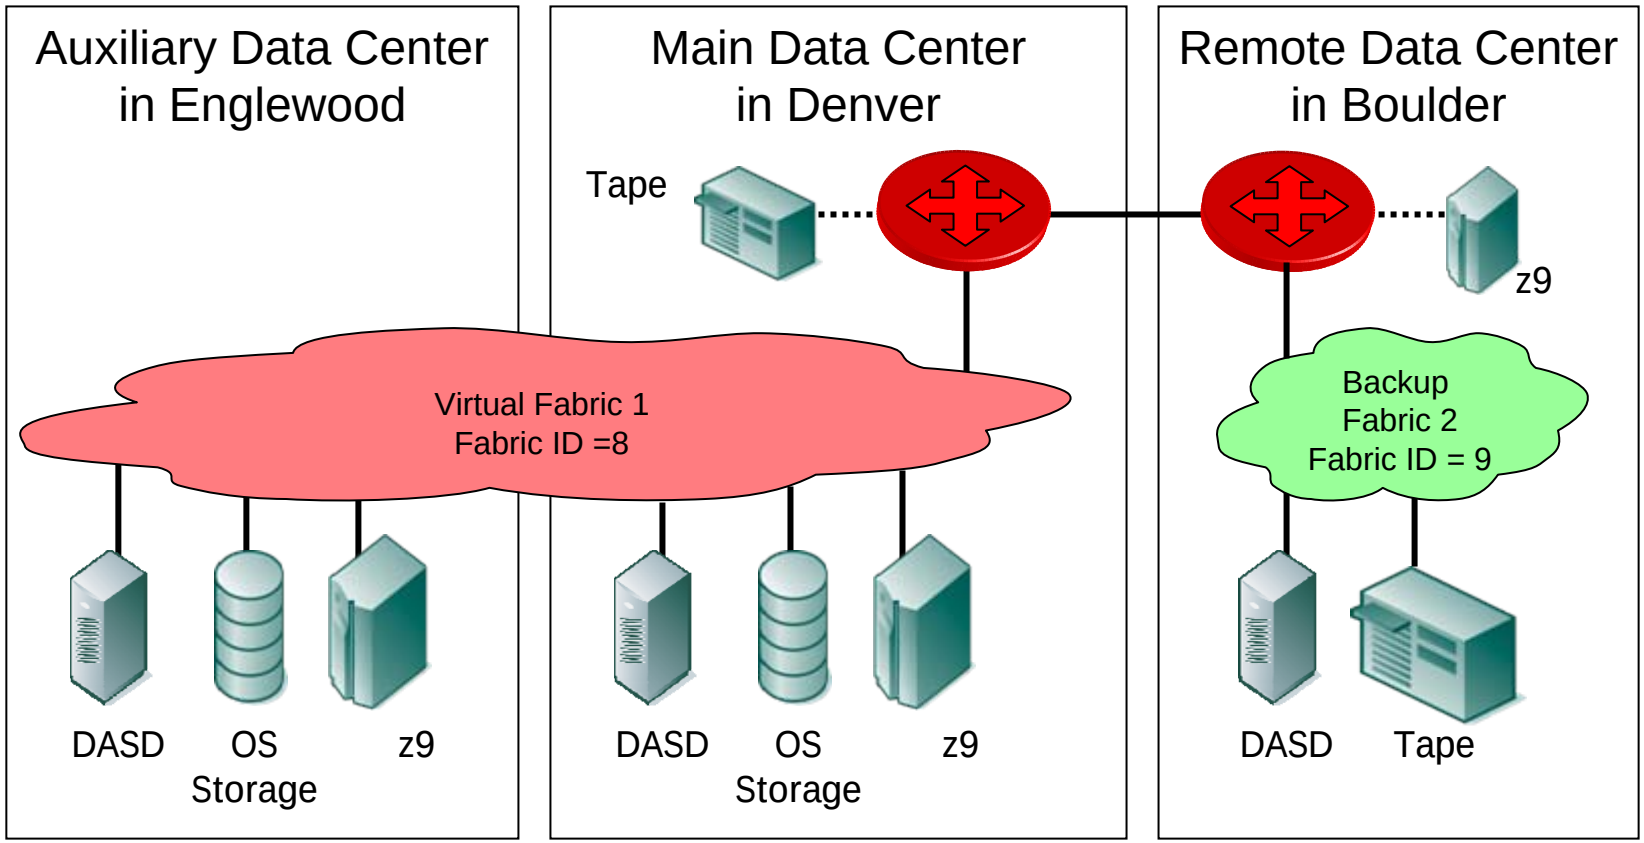

Proxy N_Ports

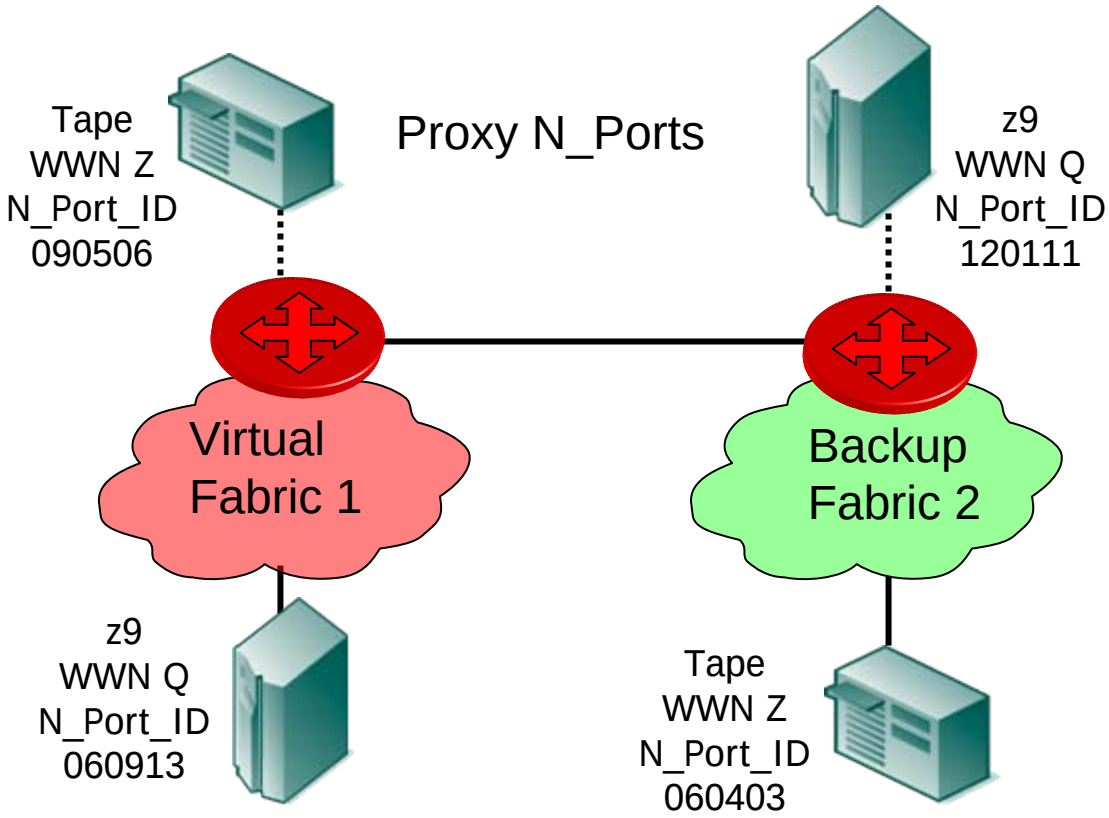

Inside an IFR

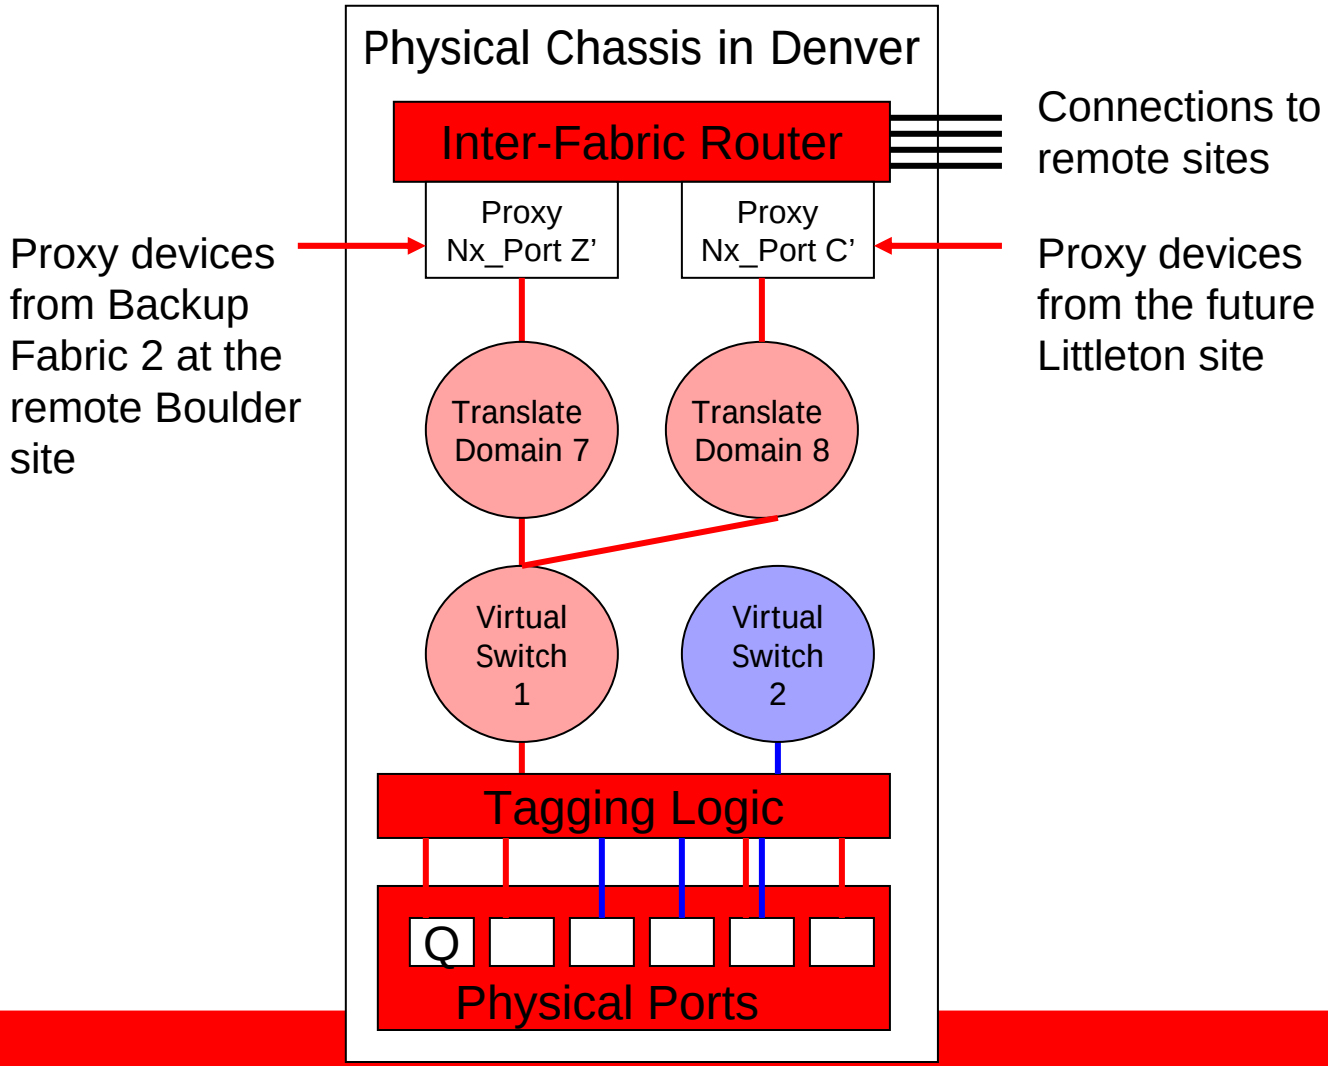

Expanded Fibre Channel Addressing

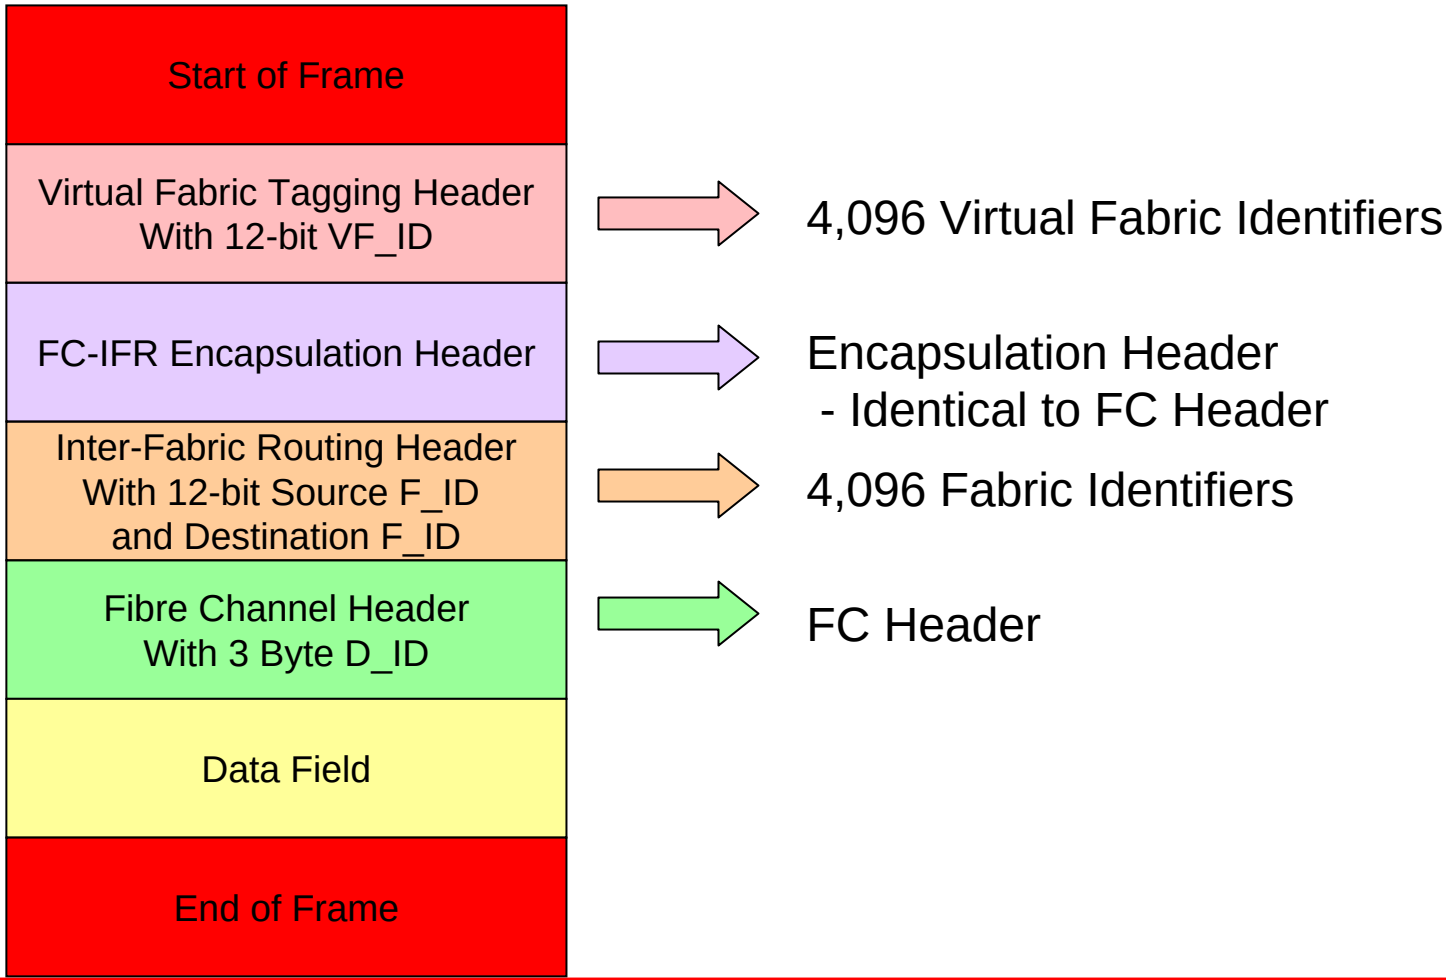

Summary

- Explored the fundamentals of virtualization
  - Server, fabric, link, and address/access
    - Consolidation Drives Virtualization in Storage Networks
      - December 2006 / January 2007
    - Fibre Channel Improves Utilization and Scalability
      - February / March 2007
- Examined fundamentals in theory
  - Server
    - N_Port Identifier Virtualization
  - Fabric
    - Virtual Fabric Identifier
    - Inter-fabric Routing
- Apply theory to a practical application

BROCADE

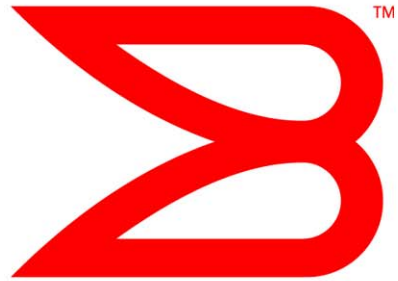

References

zJournal Articles

- Consolidation Drives Virtualization in Storage Networks
  - December 2006 / January 2007
  - http://www.zjournal.com/index.cfm?section=article&aid=739
- Fibre Channel Improves Utilization and Scalability
  - February / March 2007
  - http://www.zjournal.com/index.cfm?section=article&aid=764

Standards and NPIV

- FC-LS
  - Describes FDISC use to allocate additional N_Port_IDs in Section 4.2.32
  - Service Parameters for FDISC are described in Section 6.6
  - NV_Ports are treated like any other port
    - Exception is they use FDISC instead of FLOGI
  - Documents the responses to NV_Port related ELSs in section 6.4.5
    - FDISC, FLOGI and FLOGO
  - http://www.t11.org/t11/docreg.nsf/ufile/06-393v6

- FC-GS-5
  - Describes Name Server queries in 5.2.5
    - Permanent Port Name and Get Permanent Port Name command
      - Based on the N_Port ID (G_PPN_ID)
    - The PPN shall be the F_Port Name in FC-GS-6
  - http://www.t11.org/t11/docreg.nsf/ufile/06-393v6

- FC-DA
  - Profiles the process of acquiring additional N_Port_IDs in Clause 4.9
  - http://www.t11.org/t11/docreg.nsf/ufile/04-202v2

- FC-MI-2
  - Profiles how the fabric handles NPIV requests
    - New Service Parameters are defined in Section 6.3
    - Name Server Objects in 7.3.2.2 and 7.3.2.3
  - http://www.t11.org/t11/docreg.nsf/ufile/04-109v4

BROCADE

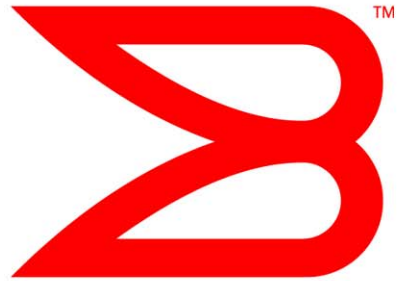

THANK YOU

About the Authors

- Scott Kipp
  - Scott develops standards, including NPIV and virtual fabrics, for Brocade. He has written several books for the Fibre Channel industry Association and speaks with storage network user groups about his latest book, *Fibre Channel Advances*.
  - skipp@brocade.com
- Steve Guendert
  - Steve is Brocade's Mainframe Solutions principal engineer and is an industry expert on ESCON and FICON. He has an MBA and an MS in Management Information Systems and is currently completing his Ph.D.
  - sguender@brocade.com
- Howard Johnson
  - Howard is a Software Architect at Brocade and is responsible for furthering Brocade's lead in FICON technology. His expertise encompasses Brocade's ESCON and FICON products and includes an extensive relationship with IBM's zSeries I/O development team.
  - hjohnson@brocade.com

Speaker Biography: Steve Guendert

- Brocade Principal Engineer, focused on mainframe
- SHARE Board of Directors (Director of IT)
- Nominee for Computer Measurement Group (CMG) BoD
- 2007 CMG Storage Subject Chair
- Academic
  - PhD Coursework completed (CS/MIS) (stats-performance)
  - M.S. in MIS
  - dissertation topic on Enterprise I/O subsystems and designing continuous availability data centers
- Industry experience
  - IBM, McDATA, CNT, Brocade
- Ohio Valley CMG Regional Chairman
- CMG Editorial Review Board (ERB), zJournal ERB
  - Published papers in zJournal, CMG, NaSPA Technical Support, DRJ