#SHAREorg

SHARE
Technology · Connections · Results

Enhanced Availability and IT Resilience: An Integrated TS7700 Grid

Dr. Steve Guendert
Brocade Communications

February 7, 2013
Session 12734

Stephen.guendert@brocade.com
@BRCD_DrSteve

SHARE in San Francisco 2013

Abstract

The IBM Virtualization Engine TS7700 multi-cluster grid is an important component of many enterprise business continuance strategies. The TS7700 grid network is used for the TCP/IP based replication between TS7700 clusters at different sites. The performance and management of this network is crucial to meeting RTO and RPO objectives. This presentation will discuss some best practices and techniques to have the optimal TS7700 grid network and will include a case study. The presentation is based on the IBM Redguide co-authored by the speaker.

- Redguide publication date is 1 November 2012

Enhanced Data Availability and Business Continuity with the IBM Virtual Engine TS7700 and Brocade Integrated Network Grid Solution

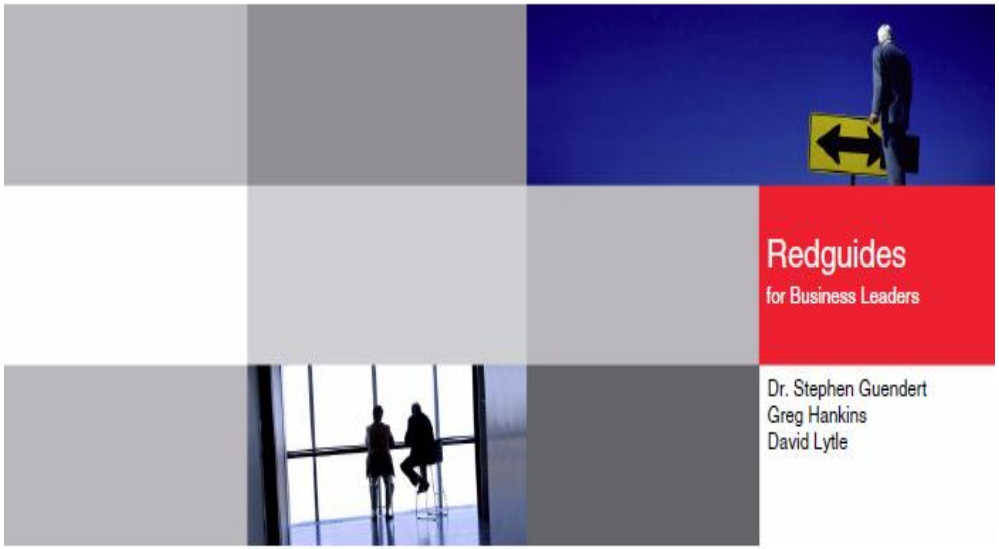

- Learn why a resilient IT architecture is essential to today's businesses
- Understand how to apply TS7700 Grid configurations for a resilient IT architecture
- See how to resolve network bandwidth issues with Brocade solutions

Agenda-Overview

- Introduction-why worry/ why be prepared?
- IT resilience vs. traditional DR/BC planning
- TS7700 and TS7700 Grid Network basics
- Example TS7700 Grid configurations
- Data Center Networks: FCP/FICON and IP synergy
- Case study
- Conclusion and questions.

Why Worry?

Lightning Strikes

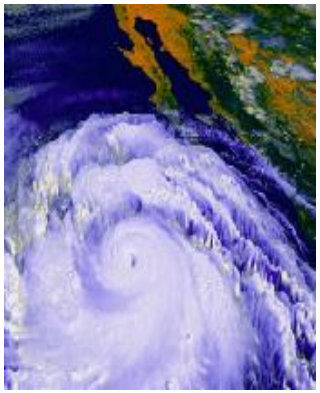

Hurricanes / Cyclones

Earthquake

Tornadoes and other storms

Tsunami

Data theft and security breaches

Overloaded lines and infrastructure

Cut cables and power

Terrorism

Be prepared

- Data availability and business continuity offer a vital competitive edge that is crucial to the success of organizations.
- Organizations must adopt proven business continuity and recovery management strategies, in addition to storage and network technologies, to successfully address operational risk, availability, and security challenges.
- One way to address these challenges is to plan and implement a resilient IT architecture.

IT resilience

- The ability to rapidly adapt and respond to any internal or external disruption, demand, or threat and continue business operations without significant impact.
  - Continuous/near continuous application availability (CA)
  - Planned and unplanned outages
- Broader in scope than disaster recovery (DR)
  - DR concentrates solely on recovering from unplanned events
- **Bottom Line**
  - **Business continuance is no longer simply IT DR**

4 factors creating a challenge

- Planning a resilient architecture has become more challenging over the past several years because of the following primary factors:
  - Increasingly competitive business environments
  - More stringent government regulations
  - More complex applications
  - Increasingly advanced technology

Resilient architecture vs. traditional DR planning

- Although it is related to planning for disaster recovery, planning for a resilient IT architecture is much broader in scope.
- A resilient IT architecture requires organizations to go beyond planning for recovering from an unplanned outage.
- Planning a resilient IT architecture implies planning for avoidance of outages, that is planning for business continuance.
- A key component in a resilient IT architecture for IBM® zEnterprise® customers is the IBM Virtualization Engine TS7700 family of mainframe virtual tape solutions with their multi-cluster grid configuration.

IBM's TS7700 Grid Virtual Tape

- The IBM mainframe offering is the TS7700 Virtualization Engine
- This comprises a virtual tape grid system (think cloud) with synchronous and asynchronous modes of replication.
- A grid can have up to six clusters, but the architecture can support up to eight.
- Management is an integral part of the TS7700 Virtualization Engine (VE) and the intelligence is outboard of the mainframe servers.
  - This means that the location and number of virtual volumes are virtualized, which will assist with content-based access.
- Using virtualization and disk cache, TS7700 is able to operate at disk speeds while maintaining compatibility with existing tape operations.
  - One of the TS7700 configurations completely eliminates physical tape through the use of a deep disk cache repository.
- TS7700 uses IP links for replication. It also has separate channels (FICON) for the front end host communications and (SCSI) for backend tape drive communications.

TS7700 Grid basics

- The TS7700 Virtualization Engine grid configuration is a series of clusters that are connected by a network to form a high availability, resilient virtual tape storage architecture.
- Logical volume attributes and data are replicated by using Internet Protocol (IP)across the clusters that are joined by the grid network.
- A grid configuration looks similar to a single storage subsystem to the hosts. This configuration ensures high availability and ensures that production work continues if one or more clusters become unavailable.
- A TS7700 Grid configuration helps a zEnterprise customer achieve a resilient IT architecture.

TS7700 Grid Solutions

WAN interconnected TS7700s form a Grid configuration

Optimized for Recoverability and Automatic Failback

- Each TS7700 provides
  - –256 virtual devices
  - –4 connections to FICON Channels
  - –Up to 84TB of Cache (3:1 C/R) TS7740
  - –Up to 1.8PB of Cache (3:1 C/R) TS7720
- 4,000,000 logical volumes
- Interconnect is standard TCP/IP using up to four 1Gbit links or two 10 Gbit links
  - –Supporting 1000s of miles separation
- Data written to one is transparently replicated to the other
  - –Policy control for location of copies and how copies are made
- A volume's data can be accessed through any TS7700
- Can be configured for disaster recovery and/or higher availability environments
  - –One to six site configurations

TS7700 Grid Configurations

- Up to six TS7700 clusters together to form a Grid configuration
  - B10/B20 P2P VTCs have been eliminated
  - Hosts attach directly to the TS7700 clusters
- Clusters in a Grid can be any combination of TS7740s and TS7720s
- Any volume accessible through any TS7700 cluster in the Grid configuration
- I/P based replication
  - Two 1Gbps Ethernet links
    - Dual-port copper RJ45 FC1032
    - Dual port shortwave optical fibre FC1033
  - Option to have 4 x 1G (SW optical or Copper) or 2x 10Gb (LW Optical Grid Ports)
  - Dynamic Grid network load balancing
- Policy-based replication management
- Grid Merge[1]
  - Merge two existing multi-cluster grids to form a single larger grid

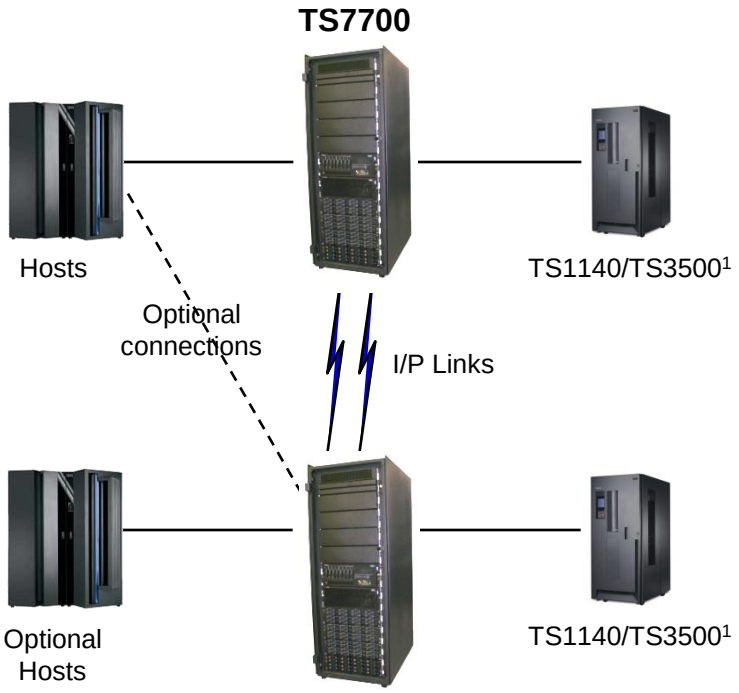

Grid Configuration

Belongs to the R2.1 code

TS7700 Grid Configuration Examples

TS7700 Two Cluster Grid Configuration for Disaster Recovery

- Two TS7700s are located at two sites geographically separated
- Interconnected through a Wide Area Network
- Only the disaster recovery host is connected to the remote TS7700
- If local TS7700 is unavailable, data is only available at remote TS7700
  - Ownership Takeover Manager enabled automatically when one of the TS7700 fails

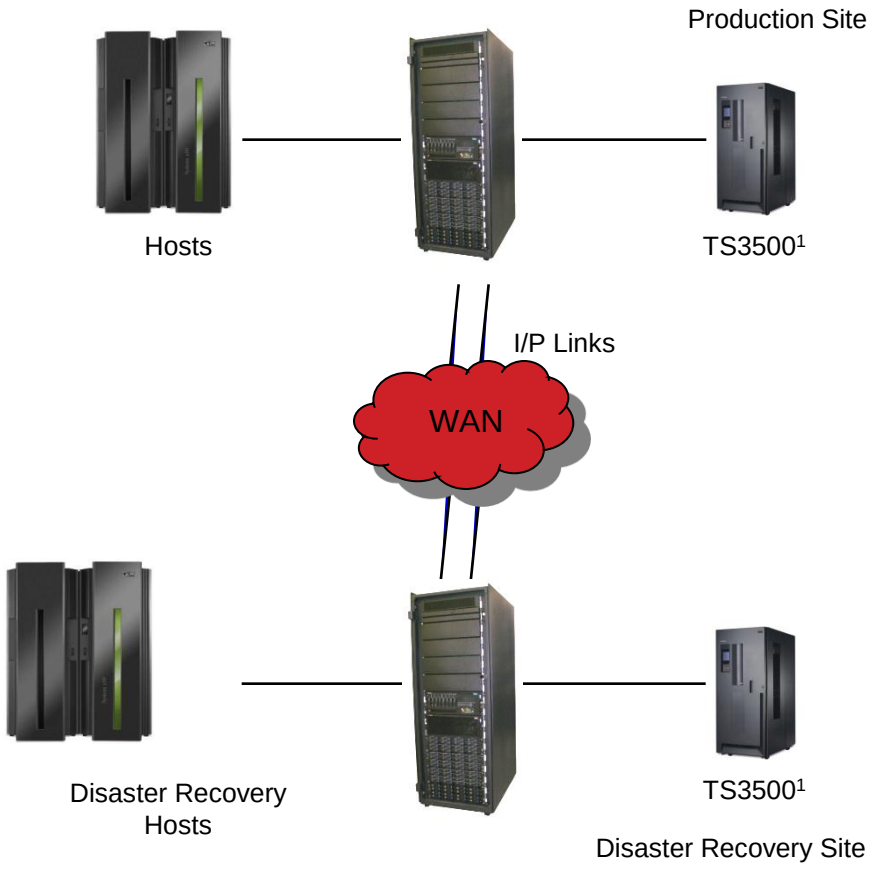

[1]Supported by TS7740 Model only

TS7700 Three Cluster Grid Configuration for Disaster Recovery

- Two TS7700s in independent data centers which may be remote
- Hosts are not connected to each other
- TS7700s do not replicate data between themselves
- Third TS7700 receives primarily deferred copies for DR

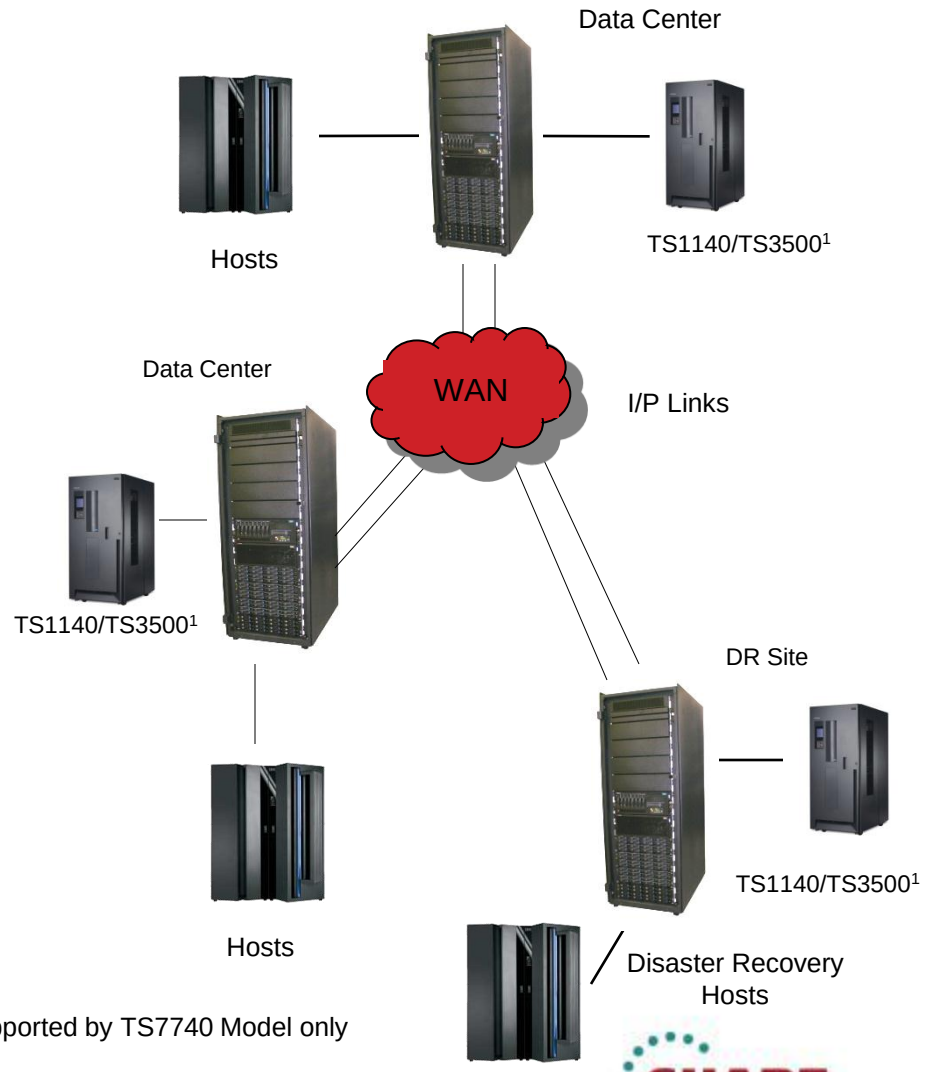

[1]Supported by TS7740 Model only

TS7700 Four-Way Grid Configuration for High Availability and High Performance

- High availability, high performance for all data through remote access
- Higher cache hit percentage with production cache size of 400TB

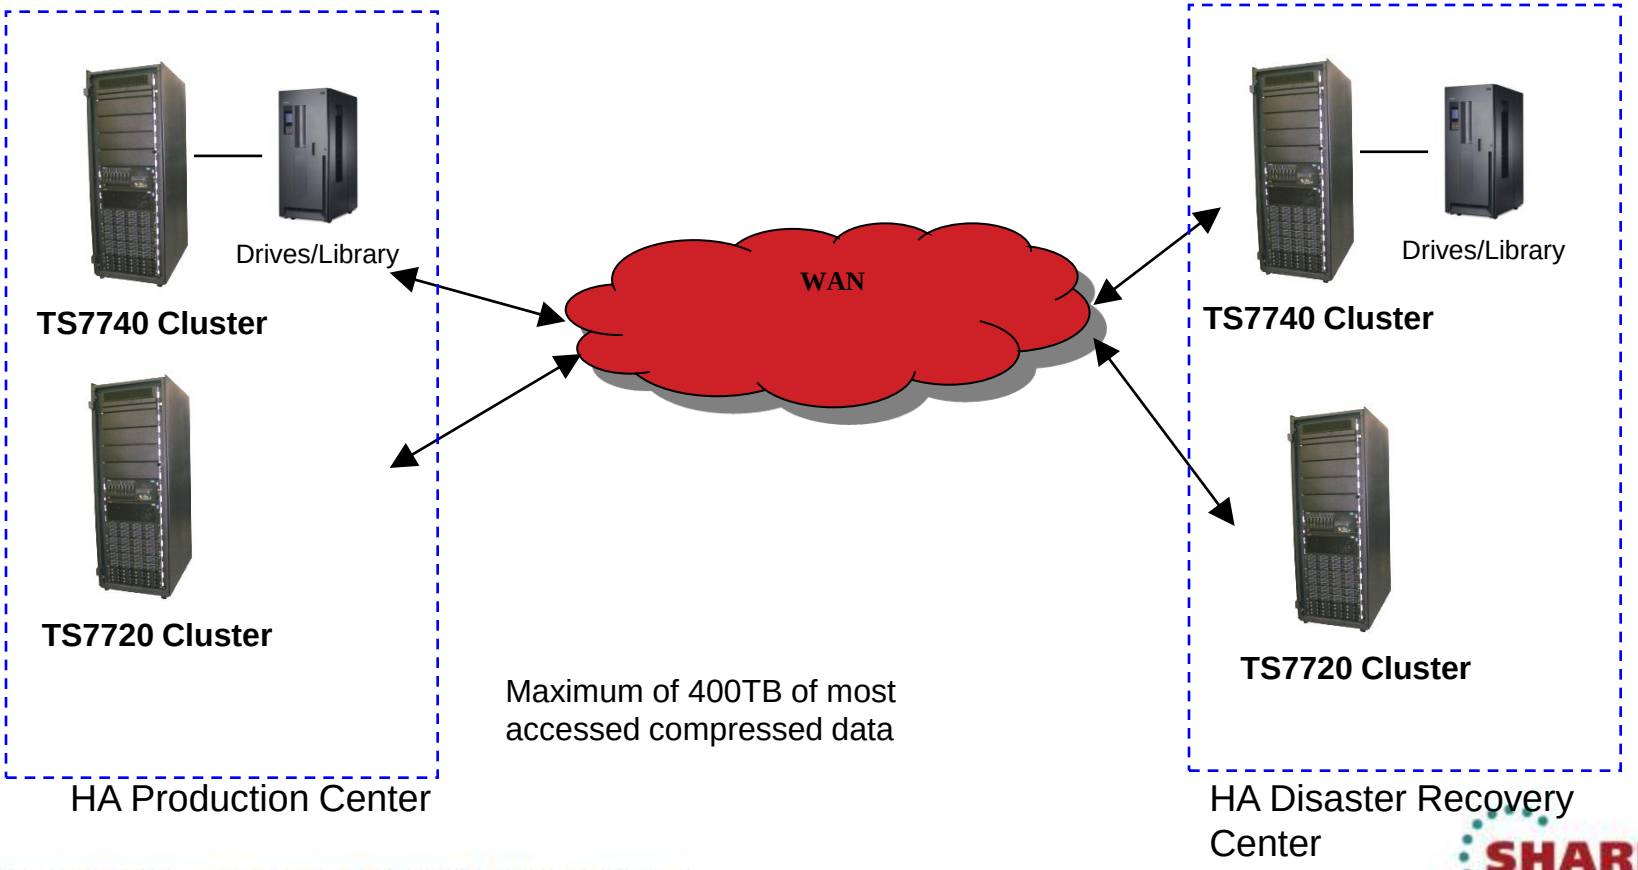

The Value of Grid Interconnect Method

- Compared with older VTS PTP
  - Move from FICON to I/P based interconnections
  - Integrate peering functions into base TS7700 Virtualization Engine
  - Reduces infrastructure requirements
    - Eliminates Virtual Tape Controllers
    - Eliminates channel extension hardware
  - Simplifies management
    - Fewer elements to manage (no VTCs, channel extension hardware)
    - Single management interface for standalone and Grid configurations

Grid Computing Supports Information Infrastructure

- Grid Computing is about virtualizing and sharing resources
- The TS7700 Virtualization Engine
  - Supports infrastructure optimization
    - Facilitates workload management and consolidation
    - Reduces time to information
  - Increases access to data and collaboration
    - Facilitates access to information
    - Supports global distribution
  - Provides resilient highly available infrastructure
    - Business continuity, recovery and failover
  - Provides information security
    - Supports tape encryption
    - Supports WORM

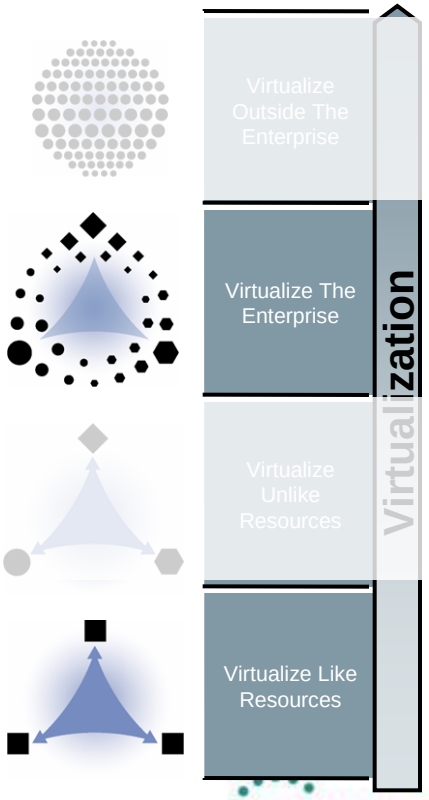

SHARE
Technology · Connections · Results

DATA CENTERS: FC AND ETHERNET SYNERGY

Provisioning for Distance

- The distance between two enterprise data centers could be within a campus, metro, or wide-area span. SAN extension solutions include:
- Direct fiber (aka Dark fibre) as well as:
  - Wavelength services such as Dense Wavelength Division Multiplexing (DWDM) and Course Wavelength Division Multiplexing (CWDM)
  - Synchronous Optical Networking (SONET) and Synchronous Digital Hierarchy (SDH services)
  - Fibre Channel over IP (FCIP)
- Direct fibre (up to 10km), DWDM (up to 300km with amplification), and CWDM (up to 100km) are typically meant for interconnecting SAN segments within metro or regional spans.
- SONET/SDH (thousands of kilometers) is meant for the national level and FCIP (ten of thousands of kilometers) is targeted toward global WAN ranges.

Should Mainframe Professionals Worry About IP?

Absolutely!

- As mainframe professionals we control EVERYTHING that it is within our power to control – stability is our keystone
- Why?
- Because management requires – our heritage demands – that we provide the most **robust, highly available and performance predictable environment** we possibly can
- In the past we have usually only had to worry about those characteristics as they applied within a data center
- But more and more it is necessary to provide that robust, five-9s of availability, predictability and scalability of our environment across and **between** data centers
- LAN / WAN and FCIP technologies have historically not been highly available or predictable so **YOU** must control it and not allow networking people to affect your operations!

You must take charge of all of the components, including IP, that bring stability, availability and performance to your data center enterprise!

Survey Says: Customer Key Business Objectives

Availability and Predictability are KING in the Mainframe World

- Ensure infrastructure uptime
  - Availability is essential to business continuity
  - A combination of reliability, availability and resiliency provides non-stop, predictable performance for mainframe enterprises

Survey Says: Customer Key Business Objectives

Availability and Predictability are KING in the Mainframe World

- Ensure infrastructure uptime
  - Availability is essential to business continuity
  - A combination of reliability, availability and resiliency provides non-stop, predictable performance for mainframe enterprises

- Reduce costs, especially operating costs
  - Gain efficiencies
  - Focus IT on business value needs

Survey Says: Customer Key Business Objectives

Availability and Predictability are KING in the Mainframe World

- Ensure infrastructure uptime
  - Availability is essential to business continuity
  - A combination of reliability, availability and resiliency provides non-stop, predictable performance for mainframe enterprises

- Reduce costs, especially operating costs
  - Gain efficiencies
  - Focus IT on business value needs

- Plan for business growth and scalability
  - Accelerate new technology deployment
  - Scale to meet growth and application needs

Synergy – All Components Working in Harmony

Multi-Sourcing / Multi-Vendors simply cannot provide a synergistic environment

- Synergy : two or more things functioning together to produce a result not independently obtainable.
- Meeting up-time requirements, reducing costs and ensuring efficiencies – while at the same time architecting for growth and scalability – is a very difficult proposition
- Unless all of the data center components are themselves synergistic (designed to work together), it is almost impossible to meet all of your strategic goals while monitoring, managing and troubleshooting a disparate collection of vendors, hardware and services.
- Achieving high availability in that environment is challenging enough but then consider the complexity of change management, incident management, reporting and reviews.

Resilience of Data Center Components

Considerations for Single Vendor versus Multi-Sourcing / Multi-Vendors

- Not everything performs at five-9s
  - Components that are usually capable of five-9s:
    - Mainframes
    - Fibre Channel Director-class switching devices
    - DASD arrays
  - Components that are not usually capable of five-9s:
    - Tape
    - LAN/WAN
- In particular, customer's face a major challenge with downtime on their Local Area Networks (LAN) and Wide Area Networks (WAN)
- This often leads to improper functioning of their applications.

Maximum Data Center Hardware Availability Potential

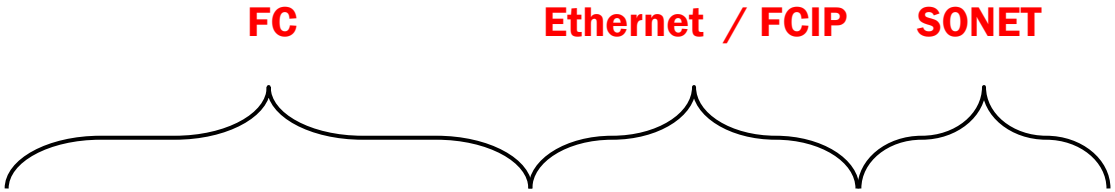

Five-9s of Availability:
- Mainframes
- Director-class switching
- DASD arrays

Five-9s of Availability:
- MLXe

Four-9s of Availability
- Motherboard-based switching

Four-9s of Availability
- FX8-24 blades
- 7800 ext. switch

Four-9s of Availability:
- Telco xWDM

Three-9s of Availability
- Older IP switching

Two-9s of Availability
- Physical tape drives
- Printers, etc.

- **Reliability is a measurement of the dependability of a customer's system, fabric and/or devices**
- Reliability really talks about "how long before something breaks" but in the data center we never want FICON to "break" and that is where "Availability" comes into play
- **Availability is the ability of a system, fabric and/or device to continue to provide services when they are needed without delay even after a component failure**
- In the data center this is typically discussed as a percentage such as two-9s, three-9s, four-9s or five-9s (99.999%) of availability
- It is really an up-time ratio!

Resilience of Data Center Components

Considerations for Single Vendor versus Multi-Sourcing / Multi-Vendors

- Customers often have various tier 2 and tier 3 hardware and service providers and many of them are finding it more and more difficult to manage that environment.
- Engineering varies widely by vendor leading to higher costs in terms of energy, floor space and management.
- Many of the hardware and service vendors are hard pressed to provide proactive support and services
  - This can negatively impact a customer's daily operations by causing delays in provisioning/scaling, troubleshooting and also to non-compliance against Service Level Agreements (SLAs).
- Cooperation and troubleshooting among multi-sourced vendors is **always** a challenge for customer personnel.

Resilience of Data Center Components

Considerations for Single Vendor versus Multi-Sourcing / Multi-Vendors

- Further, end users usually do not have an integrated suite of monitoring and/or management tools for their systems that cross vendor lines. So there often is no single view of the performance and health of all of the links and devices on the FC as well as the Ethernet LAN/WAN.

- In many cases we find that end users do not have the internal strength to monitor, manage and maintain their SAN/LAN/WAN in terms of people, skills, products and tools and are often looking for a provider who can help them to sort out these issues with ease.

TS7700 Grid network infrastructure challenges

- The network components need to individually provide reliability, high availability, and resiliency. The overall solution is only as good as its individual parts.
  - A TS7700 Grid network requires non-stop predictable performance with components that have “five-9s” availability, that is 99.999% uptime.
- A TS7700 Grid network must be designed with highly efficient components that minimize operating costs.
- Rapidly changing and growing amounts of data require a TS7700 Grid network that meets the following specifications:
  - Highly scalable components that support business and data growth and application needs
  - Helps accelerate the deployment of new technologies as they become available

TS7700 Integrated Grid Network Case Study

IT Environment “The Before”

- Previous IT environment
  - Ciena CN 2000 for WAN
  - Cascaded McDATA 6140s for FICON - XRC
  - CNT Edge 3000s, USDX, and Brocade 7500 for Tape Duplexing Environment
  - Cisco 6509 used for IBM 7700 Tape Grid
  - Several Managements Platforms to manage these elements
- Compelling event
  - Lease expirations
- Business / Technology drivers
  - Increased Performance
  - Consolidated solution went from 3 vendors to 1
  - Simplified Management from 5 separate tools to 1

Multi-Sourced: HSM, XRC and Virtual Tape, etc.

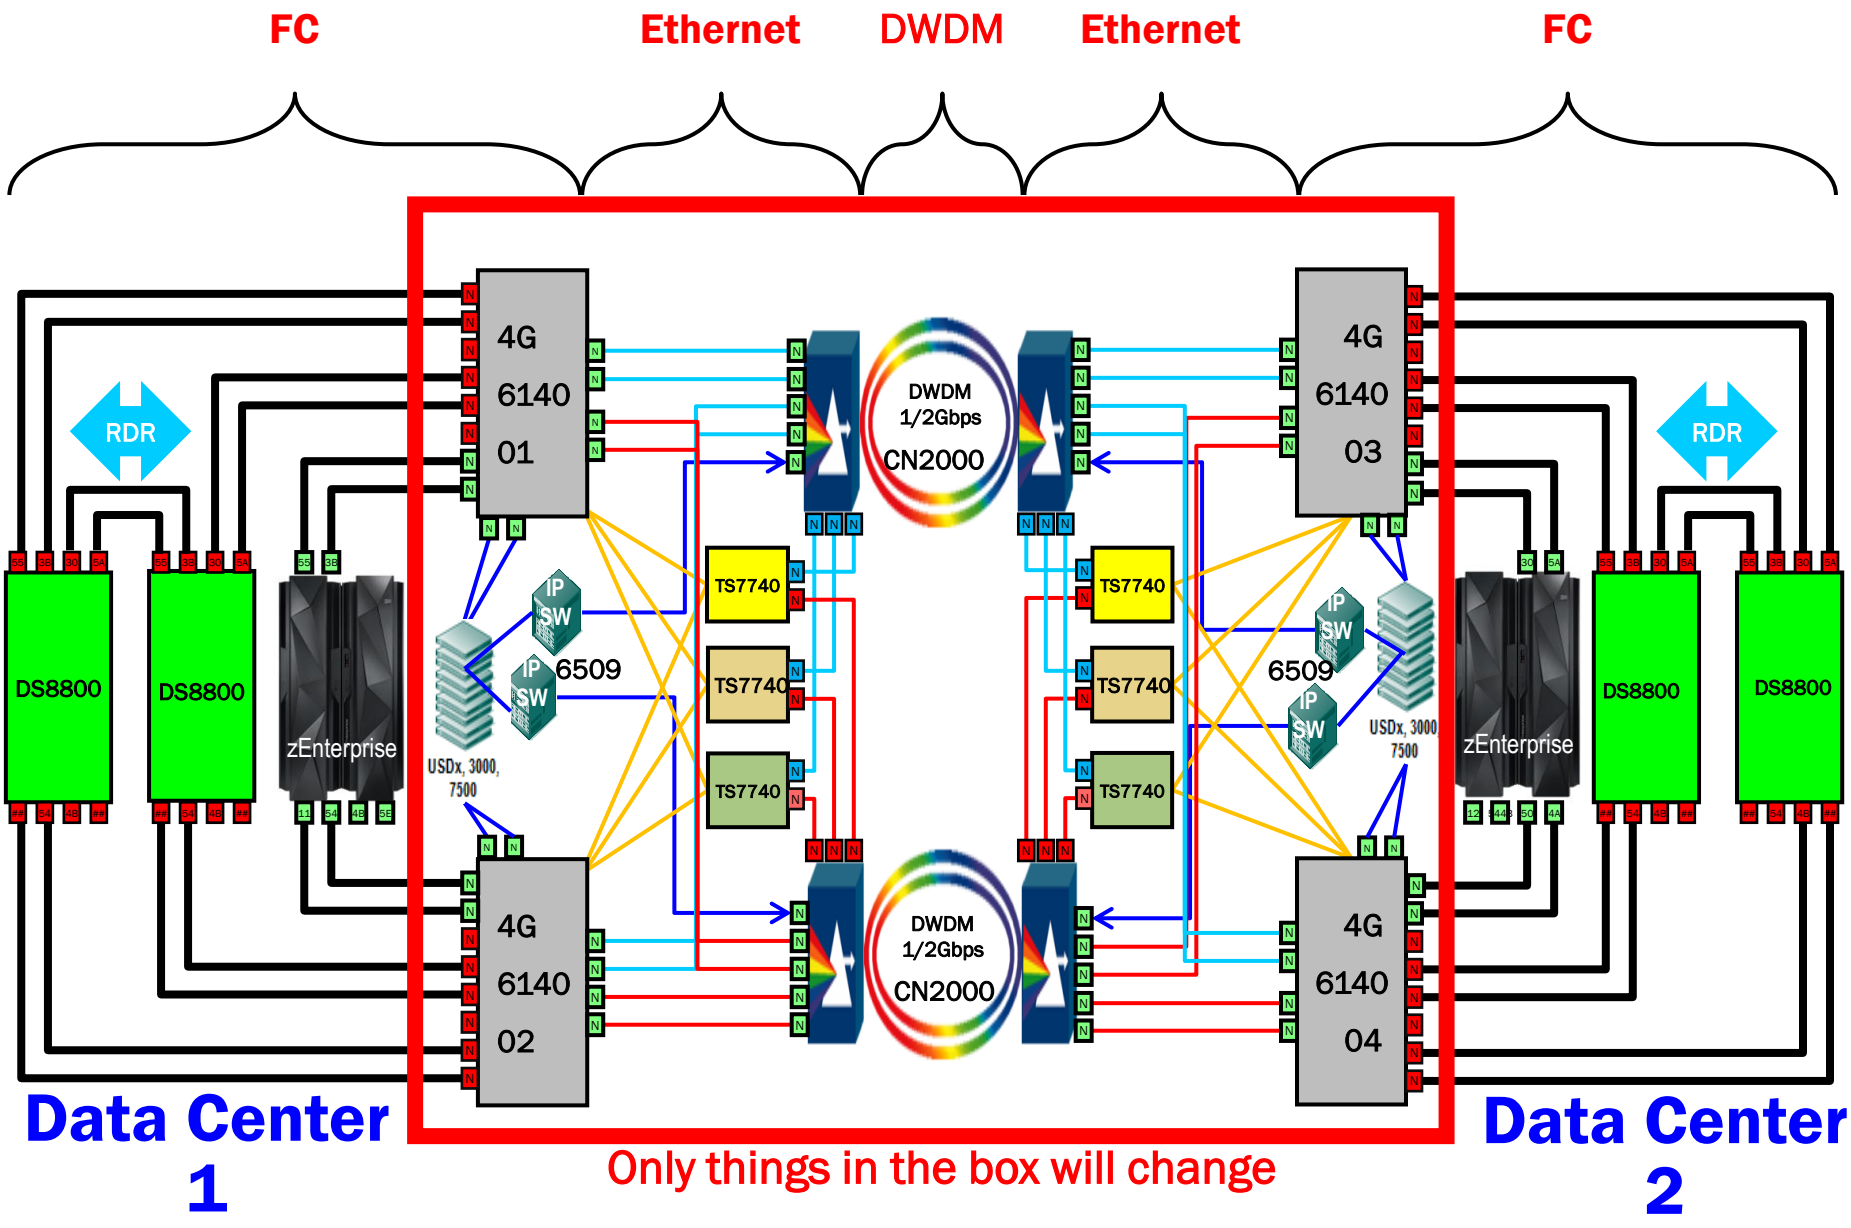

Single-Sourced: HSM, XRC and Virtual Tape, etc.

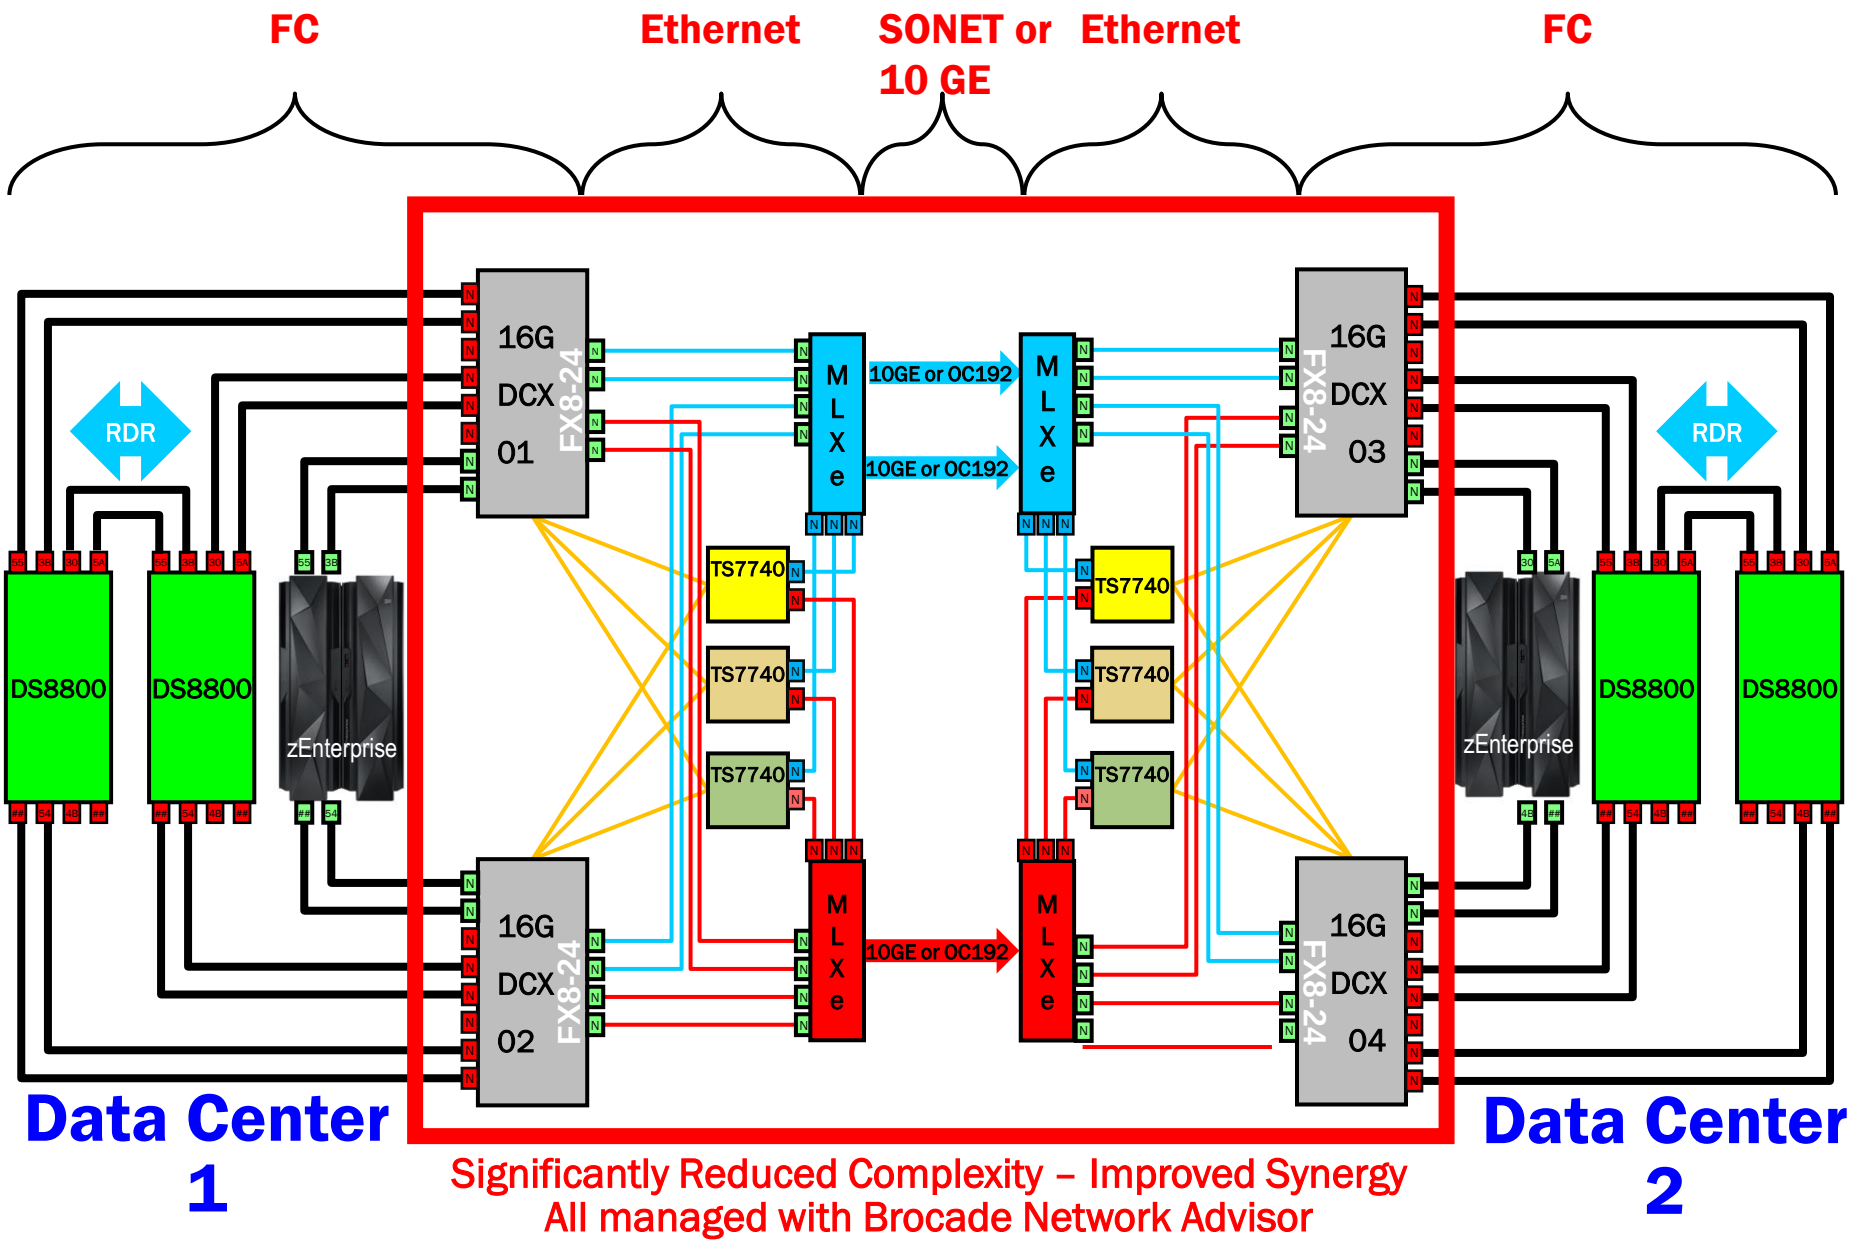

New Environment “The After”

Performance, Availability, and Operation

- Simplified management
  - One management platform and server (BNA)
- Lower costs
  - Lower monthly maintenance costs on hardware
  - Newer, “greener” technology and less floor/rack space
- Better performance-new technology
  - Increased performance, line rate on all ports
  - Up to 10 GE ports in a single trunk
  - Adaptive rate limiting allows all bandwidth to be used
  - Trunking continues emulation even during GE failures
  - Advanced compression improves efficiency

Single-Sourced: HSM, XRC and Virtual Tape, etc.

Benefits of Single Vendor versus Multi-Sourcing

- Streamline the enterprise through consolidation of resources and new technology
  - provides end-to-end high availability while lowering the overall monthly maintenance cost and other significant operating costs (e.g. energy, footprint, management, etc.)
- "One throat to choke"
- Took back control of their data

Brocade's Integrated TS7700 Grid Network

- The integrated TS7700 Grid network solution from Brocade offers zEnterprise and TS7700 customers a high performance, high efficiency, FICON storage area network (SAN).
- This SAN is coupled with high performance, high RAS, high efficiency Brocade IP routers.
- The integrated network is managed by one common network management software platform, which is the Brocade Network Advisor.

Additional Mainframe Virtual Tape Solutions

EMC's DLm Virtual Tape

- The EMC mainframe offering is the DLm8000
- The DLm8000 requires use of VMAX and provides SRDF synchronous replication as well as FICON channel extension:
  - SRDF synchronous replication allows data to be replicated over any distance continuously to a secondary data center or disaster recovery site providing an RTO and RPO of just about zero.
  - Synchronous replication allows for concurrent writes to two separate data centers to mitigate such concerns making disaster recovery with predictable Recovery Time Objectives (RTOs) a reality.
- EMC z/OS Storage Manager (EzSM) offers a common management interface for both VMAX and DLm. The net result is a mainframe-centric view of storage including VMAX monitoring and DLm management functions
- DLm uses IP links for replication and channel extension and FICON for mainframe attach

EMC DLm8000 System Z Virtual Tape

- Disk Speeds
- FICON and IP Paths
- Single Tape Drive attach Support
- Deduplication
- Encryption

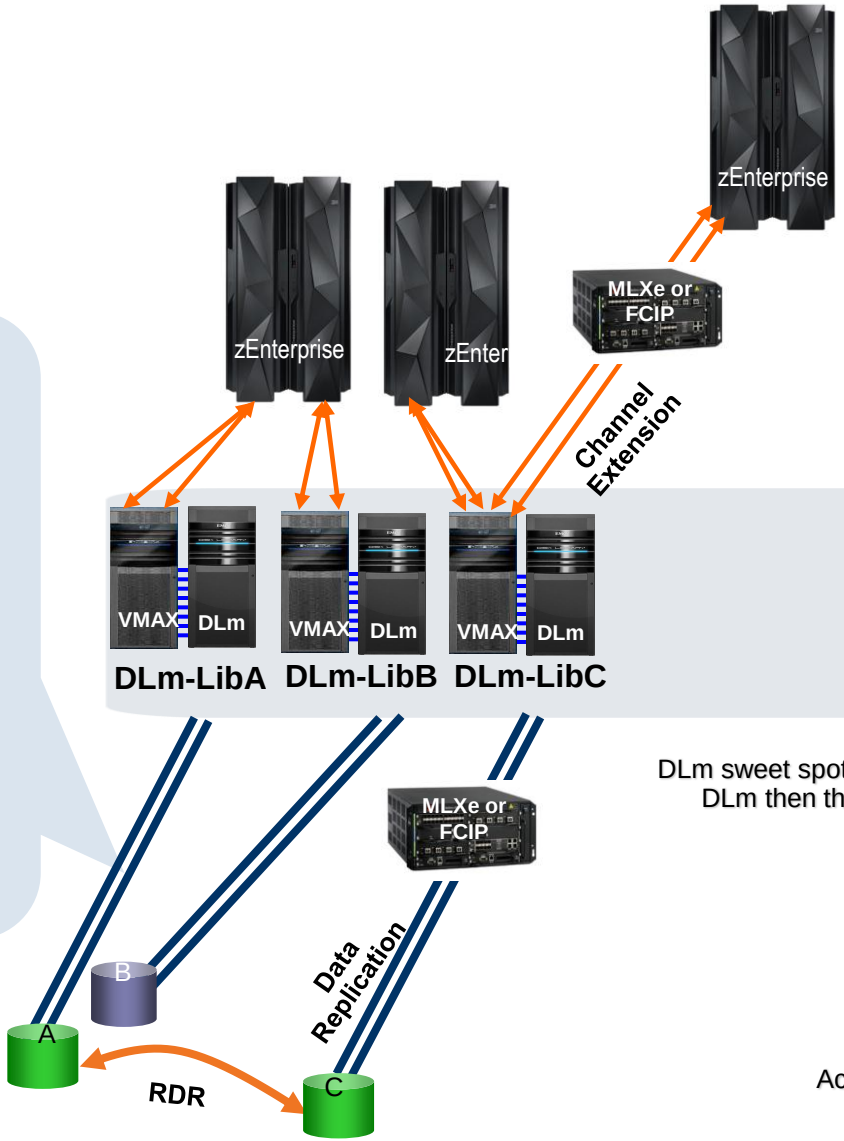

The DLm8000 can replicate data over IP from a source mainframe to one or two remote sites.

168MB/s per FICON channel

Hosts see 3 independent libraries

DLm sweet spot is 4 VMAX engines so if 4 VMAX engines are used for DLm then there would be up to 4x1GbE connections for replication available

Accesses local 'home data' through its local cache

Summary and Conclusions

- Trend towards planning a resilient architecture
- TS7700 grid is a essential component of a resilient architecture
- Take control of your network and decisions
- An integrated network has many advantages
- PPPPPP

Supplemental slides

- The slides which follow are for additional information on the hardware in the TS7700 Integrated Grid network

Components of the Brocade Integrated TS7700 Grid Network

Brocade DCX 8510 16Gbps FICON Directors

16Gbps DCX 8510 is the latest in a long line of award winning FC Directors for SAN and FICON

- **Scalability**
  - 32 and/or 48 port blades
  - 192/384 non-blocking ports at 8Gbps
  - 128/256 non-blocking ports at 16Gbps
  - Inter-Chassis Links double port counts
- **Connectivity**
  - 4Gbps and 8Gbps auto-negotiation
  - 8,10 and 16Gbps Optics in ports
  - 10Gb and 16Gb native ports
  - SW (OM2/OM3/OM4 cables)
  - LW (OS1 cables)
  - FX8-24 Extension Blade
- **Performance**
  - Non-blocking internal/external design
  - Low, deterministic latency (.7 - 2.4 ms)
  - 64b66b data encoding for efficiency
  - 8 buffer credits (BCs)/port – local
  - 8,192 BCs/ASIC for distance
  - Cut-through frame routing

- **New 16Gbps DCX ISL Features**
  - 64b66b data encoding for efficiency
  - D-Port for better ISL provisioning
  - Forward Error Correction (FEC)
  - Improved Bottleneck Detection
  - ISL Port De-Comissioning
  - Automatic Buffer Credit Recovery
- **Management**
  - Brocade Network Advisor; CUP
  - SanHealth Check Free Utility Pgm
  - Network Monitoring Service (NMS)

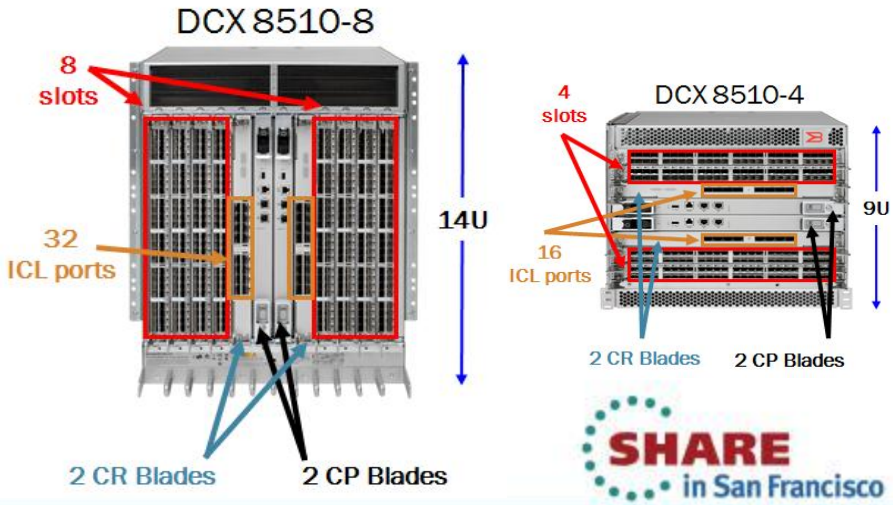

Brocade MLX Series To Create The Mainframe's IP WAN: High-Performance, Multiservice Router

- **Optimum flexibility**
  - 4-, 8-, 16-, 32-slot chassis
  - Field-programmable packet processors
- **Unparalleled operational efficiency**
  - Designed to support environments requiring high availability
  - Best-in-class power efficiency supporting hot aisle-cold aisle
  - Reduced floor space requirements vs competitive offerings
  - Users can provision with xGbE and/or SONET connections
- **Leading performance and scalability**
  - Over 15 Tbps forwarding capacity in a chassis
  - Up to 32 100 GbE, 576 10 GbE, 1,536 1GbE wire-speed ports
- **Ideal For:**
  - Connecting SAN/FICON to unlimited distance connectivity
  - Deploying IP in Large data centers and at the Core
  - Deploying Metropolitan Area Networks (MANs)

Data Center Edge and Core Routing

Brocade FX8-24 Extension Blade

Replicate data across any distance

Advantages
- No DID to manage
- 10Gbps ports

- **Optimum flexibility**
  - Mainframe or open systems
  - Twelve 8 Gbps Fibre Channel ports, 10×1 Gigabit Ethernet (GbE) ports
  - Optional 10 GbE port upgrade with up to 20 Gbps bandwidth
  - Up to four blades per Brocade Director chassis
- **Best-in-class performance**
  - Highest port density, bandwidth, and throughput
  - Industry's first and only 10 Gbps FCIP solution
- **Value-added solutions**
  - Extends replication and backup over any distance
  - Enables faster, more reliable data replication, backup, and recovery
- **Advanced Accelerator for FICON**

Brocade 7800 Extension Switch

- Platform options

Advantages
- No DCX ports consumed
- Scalability

Two models
- **Brocade 7800 4/2**: Four 8G FC ports, two GE interfaces
- **Brocade 7800 16/6**: Sixteen 8G FC ports, six GE interfaces

- Functionality enabled through optional licenses
  - **Brocade 7800 16/6 Upgrade License**: Enables all ports.
  - **Enterprise Bundle**: Advanced Performance Monitoring, Adaptive Networking, Fabric Watch, Brocade Trunking, and SAO
  - **Advanced Extension**: FCIP Trunking
  - **FICON Management Server**: CUP enables host control of switches in mainframe environments
  - **Advanced Accelerator for FICON: Next Page**

Brocade Network Advisor

Simplified Management for SAN, IP and Converged Networks

- Unified Network Management product for SAN, IP, Application Delivery, and Converged Networks
  - One management GUI across FC, IP, FCoE protocols
  - Custom views based on Operator specialization
  - Flexible user management with Role Based Access Control
- Standards-based architecture
- Provides seamless integration with leading partner Orchestration frameworks

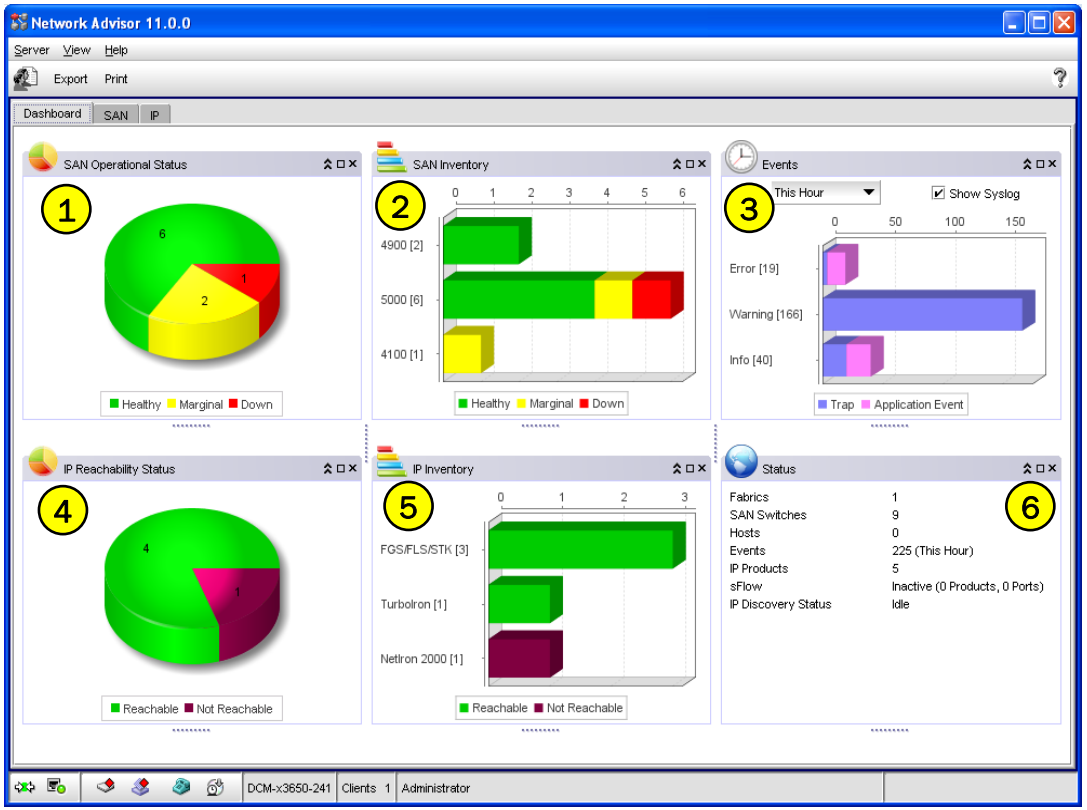

1. SAN Operational Status
2. SAN Inventory
3. Events Summary
4. IP Reachability Status
5. IP Inventory
6. Status Summary

BNA and integration with IBM TPC

- Of substantial benefit to the TS7700 Grid user, Brocade Network Advisor provides support for IBM Tivoli® Storage Productivity Center (TPC).
- Brocade Network Advisor and IBM TPC integration is focused on administrators within the same domain, in this case, storage and SAN management.
- The IBM TPC solution is an industry-leading storage resource management solution that helps customers with the management of storage environments.

Network Monitoring Service (NMS)

Non-Stop Networking Starts Here

Why do customers deploy Brocade NMS?

- Legacy of success in the extension arena
- Key component for business continuity recovery services
- Augments staffing and expertise
- For your Peace of mind and to sleep better at night

Monitored Characteristics:

- Device status
- Authentication failures
- Port, link status (physical or virtual)
  - Network transport status
  - Channel and WAN errors/status
  - Network circuit status
- Threshold alerts for several traffic and performance fault characteristics:
  - Sync loss, link resets, bad CRC
- Fan, power supply status
- Temperature sensor condition
- Power conditions

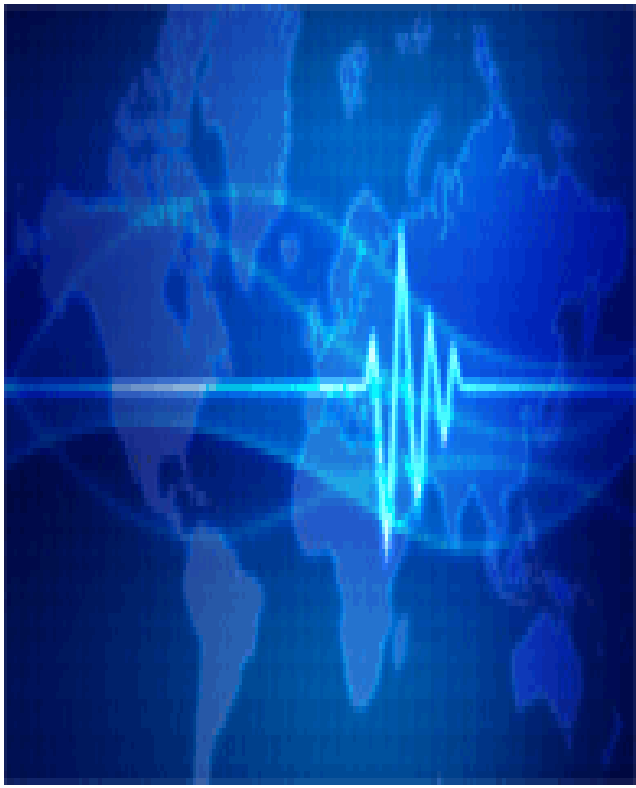

"NMS on the extension products is great. Brocade is telling us about network hits before HP even knows we have any."

—Fortune 100 Financial Institution