4. Data Management & Disaster Recovery

BROCADE COMMUNICATIONS SYSTEMS
Chonghwan Bae
System Engineer

BROCADE
The intelligent platform for networking storage

Causes of Disaster

A/C Failure
Acid Leak
Asbestos
Bomb Threat
Bomb Blast
Brown Out
Burst Pipe
Cable Cut
Chemical Spill
CO Fire
Condensation
Construction
Coolant Leak
Cooling Tower Leak
Corrupted Data
Diesel Generator
Earthquake
Electrical Short
Epidemic
Evacuation
Explosion
Fire
Flood
Fraud
Frozen Pipes
Hacker
Hail Storm
Halon Discharge
Human Error
Humidity
Hurricane
HVAC Failure
H/W Error
Ice Storm
Insects
Lightening
Lost Data
Low Voltage
Microwave Fade
Network Failure
PCB Contamination
Plane Crash
Power Outage
Power Spike
Power Surge
Programmer Error
Raw Sewage
Relocation Delay
Rodents
Roof Cave In
Sabotage
Shotgun Blast
Shredded Data
Sick building
Smoke Damage
Snow Storm
Sprinkler Discharge
Static Electricity
Strike Action
S/W Error
S/W Ransom
Terrorism
Theft
Toilet Overflow
Tornado
Train Derailment
Transformer Fire
Tsunami
UPS Failure
Vandalism
Vehicle Crash
Virus
Water (Various)
Wind Storm
Volcano

The Value of SAN as Disaster Recovery Infrastructure

- Storage Area Network Solutions can tolerate failures that would have been classified as disaster.

- Storage Area Networks provide a platform for “remote” disaster-tolerant solutions.

SANs can Tolerate Non-Expected Failure.

SANs enable failure resilient solutions

- **Component Failure** (switch, HBA, storage, and server)
  - Alternate pathing
  - Redundant Storage Networks
  - Distributed Services
  - Clustering
- **Software Failure** (operating systems, applications)
  - Clustering and Data Sharing between clustered nodes
  - Failure isolation through zoning
- **Human Error**
  - Centralized Control
  - Automatic configuration and reconfiguration without human intervention

SANs can Tolerate Expected Failures

SANs provide alternatives to scheduled down time

- **Maintenance Procedures do not interrupt business functions**
  - Applications rehost without service outage
  - Services can return to original hosts after completion
- **Configuration Changes can occur while systems operate**
  - Servers, storage devices added, upgraded, and replaced
  - Resource purchase and installation deferred until needed

SANs Provide a Platform for "Remote" Disaster-Recovery Solutions

- **SANs provide the infrastructure for business continuance**
  - Data Mirroring
  - Data Replication
  - Electronic Tape Vaulting
  - Clustering and Remote Clustering
  - Remote Disk Access

- **SANs enable these business continuance solutions over long distances**
  - Standard Fibre Channel Networks
  - Longer Distances with
    - Native fibre channel (DWDM and other MAN techniques)
    - Integration with other long distance protocols (IP, ATM)

SANs Provide the Infrastructure for Business Continuance

- **Mirroring**
  physical duplication of data

- **Data Replication**
  logical duplication of data

- **Electronic Tape Vaulting**
  backing up data to off-site tapes

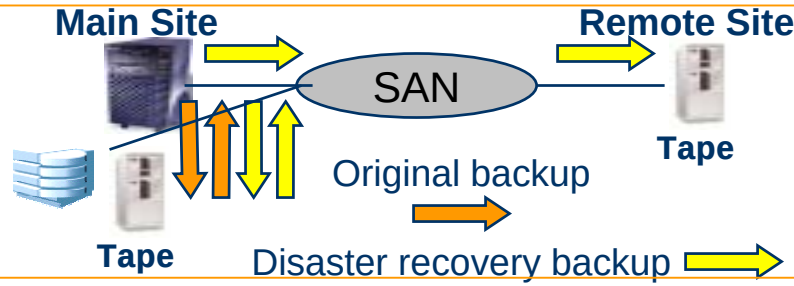

- **Clustering and Remote Clustering**
  multi-hosting services for failover

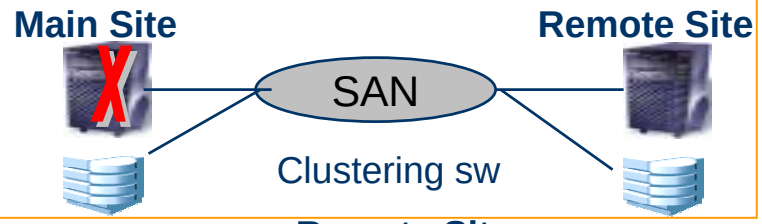

- **Remote disk access**

Business Continuance Considerations

- Time to Restore Business Operations
  - Business Processes
  - Applications
  - Infrastructure
  - Data
- Data Concurrency
- Level of response time or application performance
- Site locations (Distance between main and remote)
- Budget

These factors must be considered in each restoration phase – starting, partial, and full recovery

A Foundation for the Business Continuance Spectrum

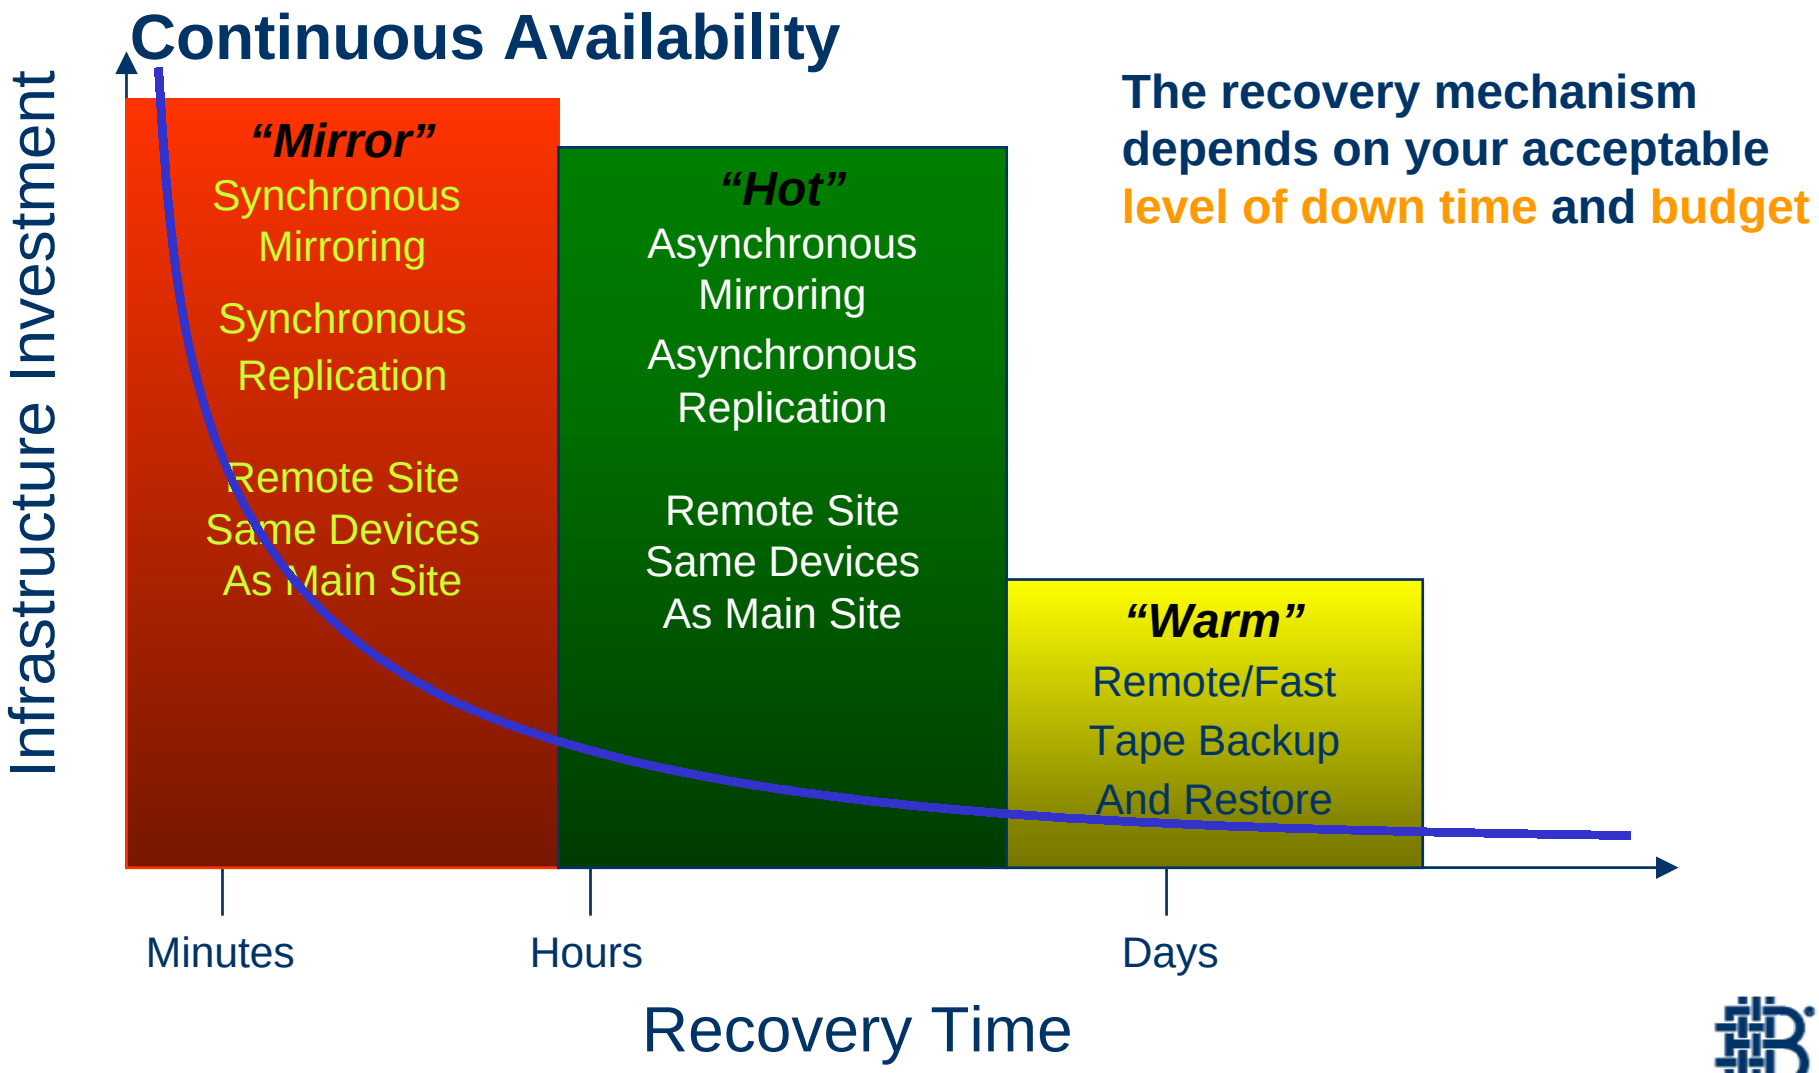

The recovery mechanism depends on your acceptable level of down time and budget

Long Distance SAN Connection: SAN Inter-Networking

- SAN over MAN with Native Fibre Channel (Finisar Extended GBIC, Link Extender, DWDM)
- SAN over WAN with Protocol Translation

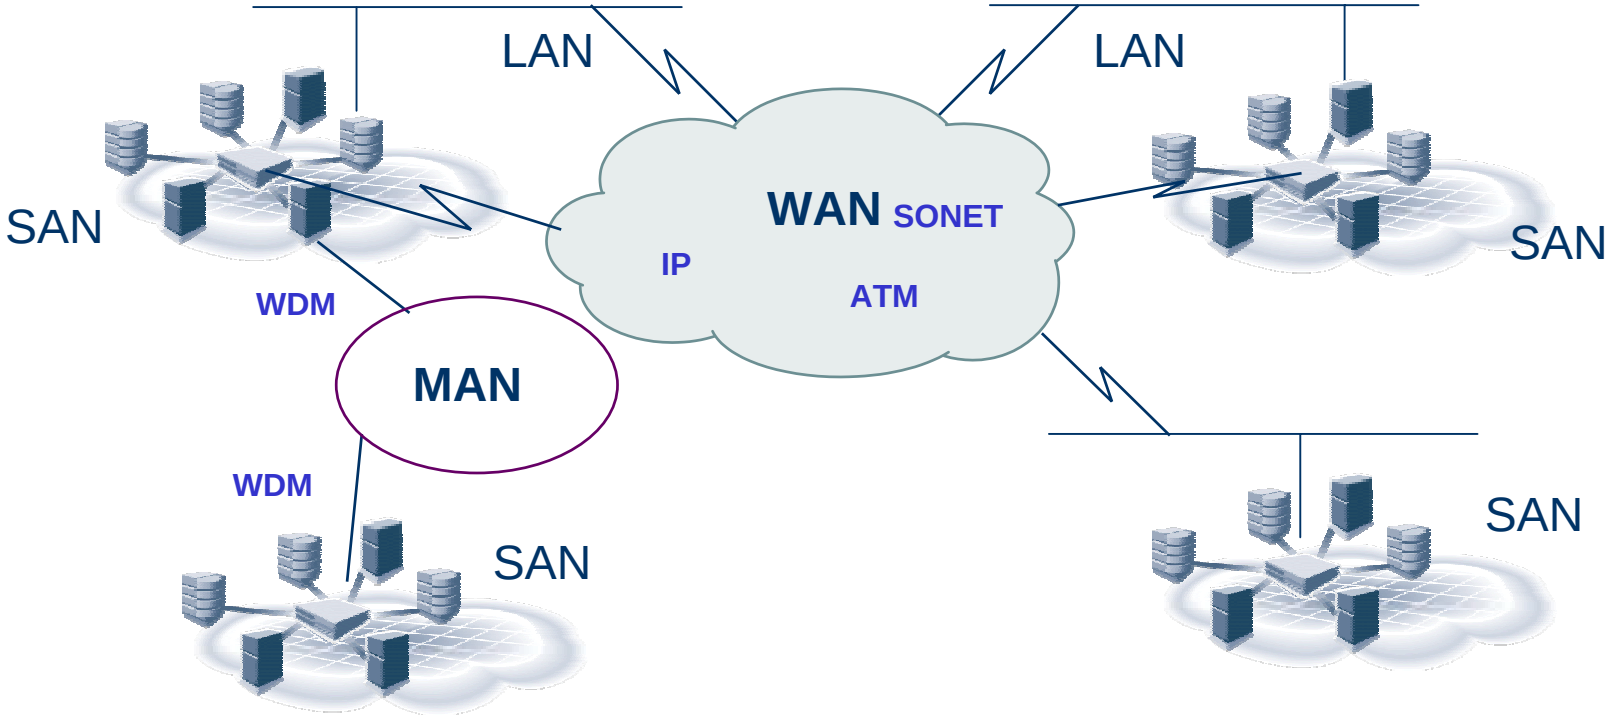

Brocade SAN over MAN Solution: Extended Fabric

- Single Fabric Utilizing Native Fibre Channel – **No Protocol Conversion**
- Distances of up to 120 Km (100 Km Tested)
- Maximum of 239 Switches (44 switches confirmed) in a Fabric
- Line Speed is 100 MB/sec or 200 MB/sec duplex
- Extended Fabric Feature Installed on all Switches in the Fabric
- Requires DWDM or ELWL GBICs or FC link extenders

Brocade Extended Fabric Basic Requirements

- **Optionally licensed product with V2.2 or above for SilkWorm 2xx0 switches**
- **Switch Ports grouped in “Quads”**
  - **0-3, 4-7, 8-11, 12-15**
    - 1 Extended Port and 3 Standard Ports per Quad
- **Configure Switches as;**
  - **Long Distance Extended Fabrics**
- **Each Port must be Configured – i.e. portCfgLongDistance 0- 2**
  - **Level 0 – 10KM (16 buffers)**
  - **Level 1 – 50KM (27 buffers)**
  - **Level 2 – 100KM (60 buffers)**

Extended Fabric Performance

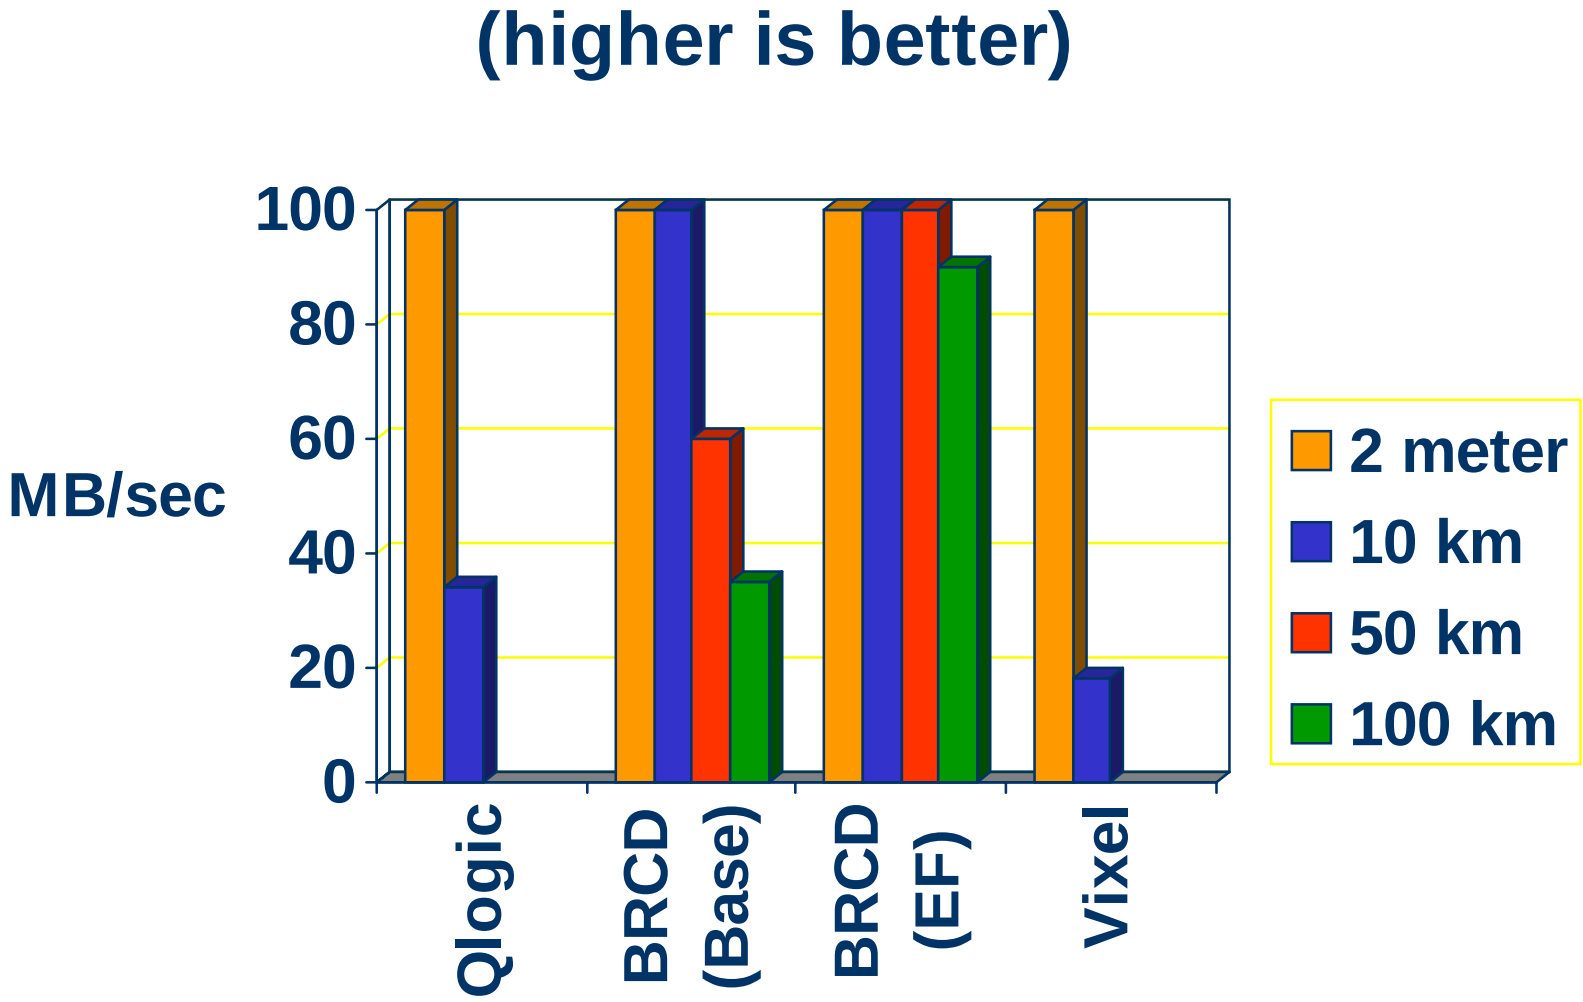

Devices: Fibre Channel Link Extenders and Long Distance GBICs

- Long Distance GBICs Allow Distance up to 80Km
- Link Extenders Allow Distances up to 120 Km (100 Km Tested)
- Line Speed 100 MB/sec or 200 MB/sec Duplex

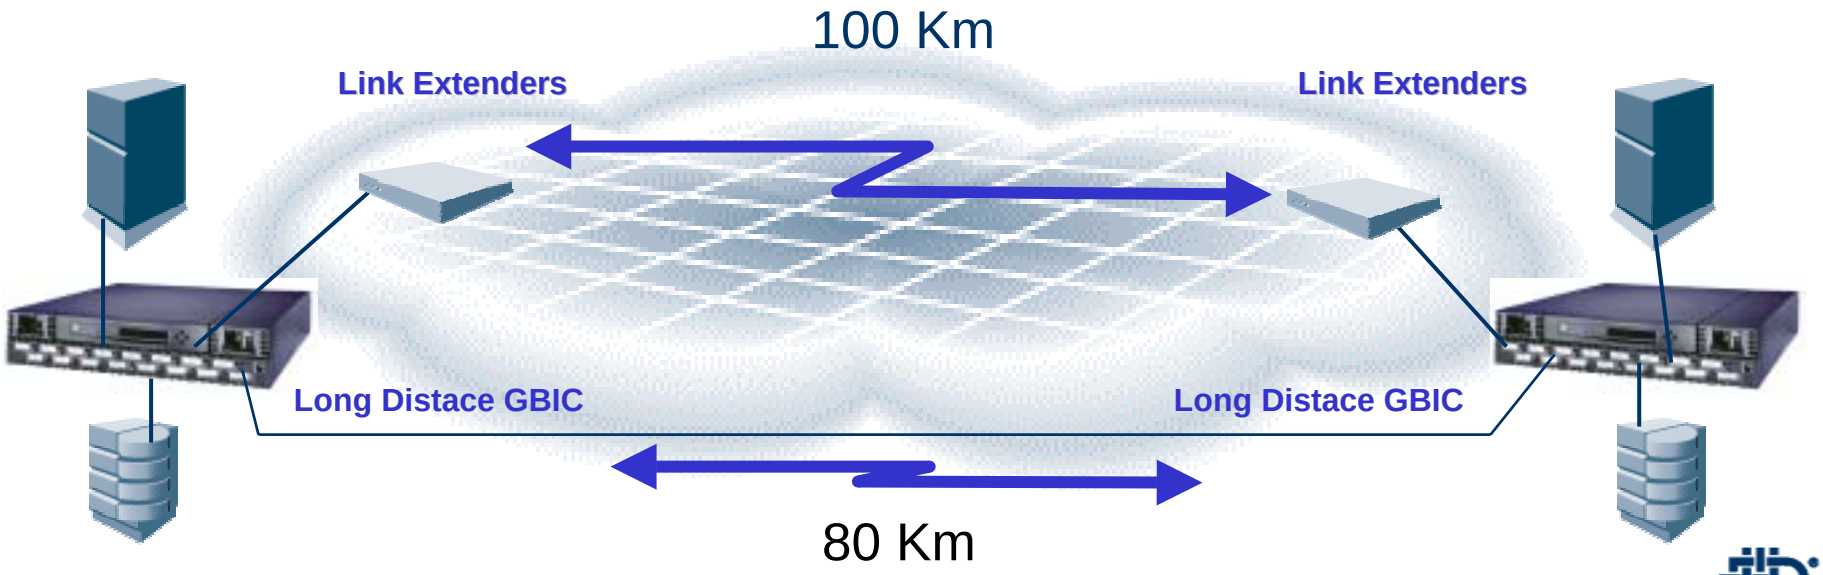

Device: DWDM

- DWDM Dense Wavelength Division Multiplexing– Is a technology that puts data from different sources together on an optical fiber, with each signal carried on its own separate light wavelength. Using DWDM, 80 + separate wavelengths or channels of data can be multiplexed into a light stream transmitted on a single optical fiber.

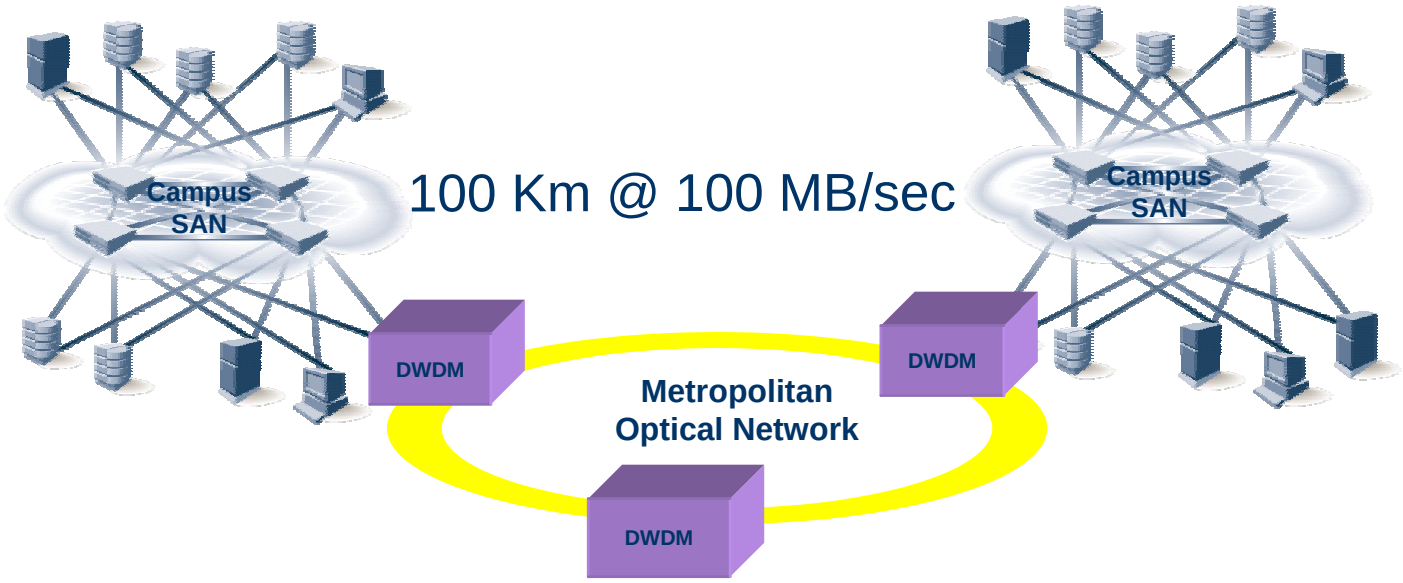

MAN Equipment Partners

- ONI Online 7000 and Online 9000
  - Online 7000 (33 protected mode, Range 100Mbps to 2.5Gbps)
  - Online 9000 (33 protected mode, Range 100Mbps to 10Gbps)

- ADVA FSP-II and FSP-III (Multiplex 32 Channels into one fiber pair)
  - FSP-II tested and certified (50 Km Protected mode, 75 Km)
  - FSP-III not currently tested or certified
- Nortel Optera (Multiplex 32, 80 or 160 channels running at 10Gbps each)
  - Tested and certified (100 Km)
- LuxN WavStation
  - Tested and Certified (50km)
- Finisar
  - FLX2000 tested and certified - Link Extender
  - ELWL GBIC tested and certified
  - Opticity DWDM certification pending

BOCADE SAN over WAN solution: Remote Switch

- Single Fabric Utilizing Fibre Channel Over WAN Gateway (ATM/OC3)
- One Switch Per Site - 2 Sites Total
- Remote Switch Option Installed on Each Switch (Fabric switches only)
- CNT Provides Open Systems Gateway
- Line Speed depend on the WAN's speed

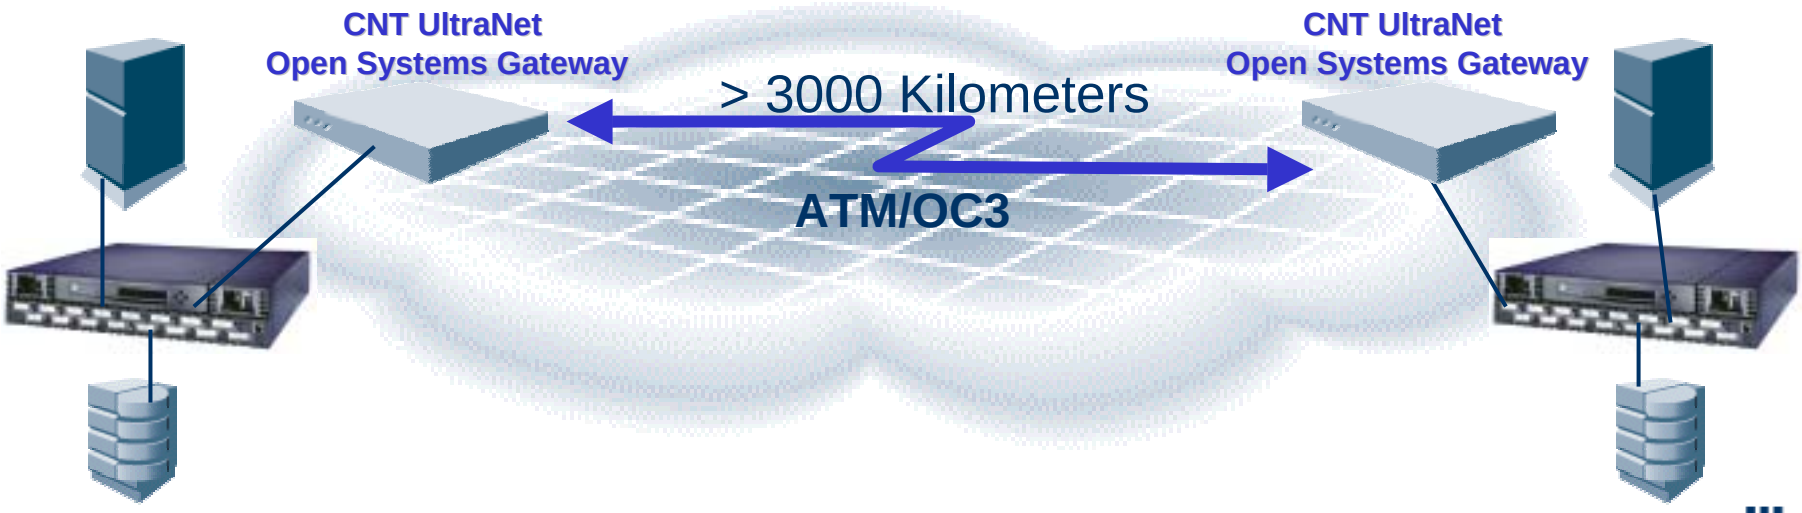

Brocade Remote Switch

- Optionally licensed product with V2.2 or above
  - Requires Silkworm 2000
- Switch to Switch (2 switch configuration)
  - Both Switches must be licensed
- Requires FC to ATM Gateway (i.e. CNT or ADVA OSG)
  - Gateway provides E_port and ATM physical interface
  - FC Frames converted to 53 byte ATM cells within Gateway

Brocade Remote Switch

- **Configure command changes**
  - **R_A_TOV (Resource Allocation TOV)**
  - **E_D_TOV (Error Detect TOV)**
  - **Data Field Size=2048 (max for ATM Gateway)**
  - **Class F Frame Suppression set – CNT can't support F Frame.**
    - Class F Frames converted to priority Class 2 Frames
  - **BB Credits on both switches must be the same**

Remote Switch - Cisco SAN to SAN Over IP

- Single Fabric
- Utilize E-Port Connectivity and Brocade Buffer Management to Provide Performance
- Line Speed Limited to WAN IP Speed

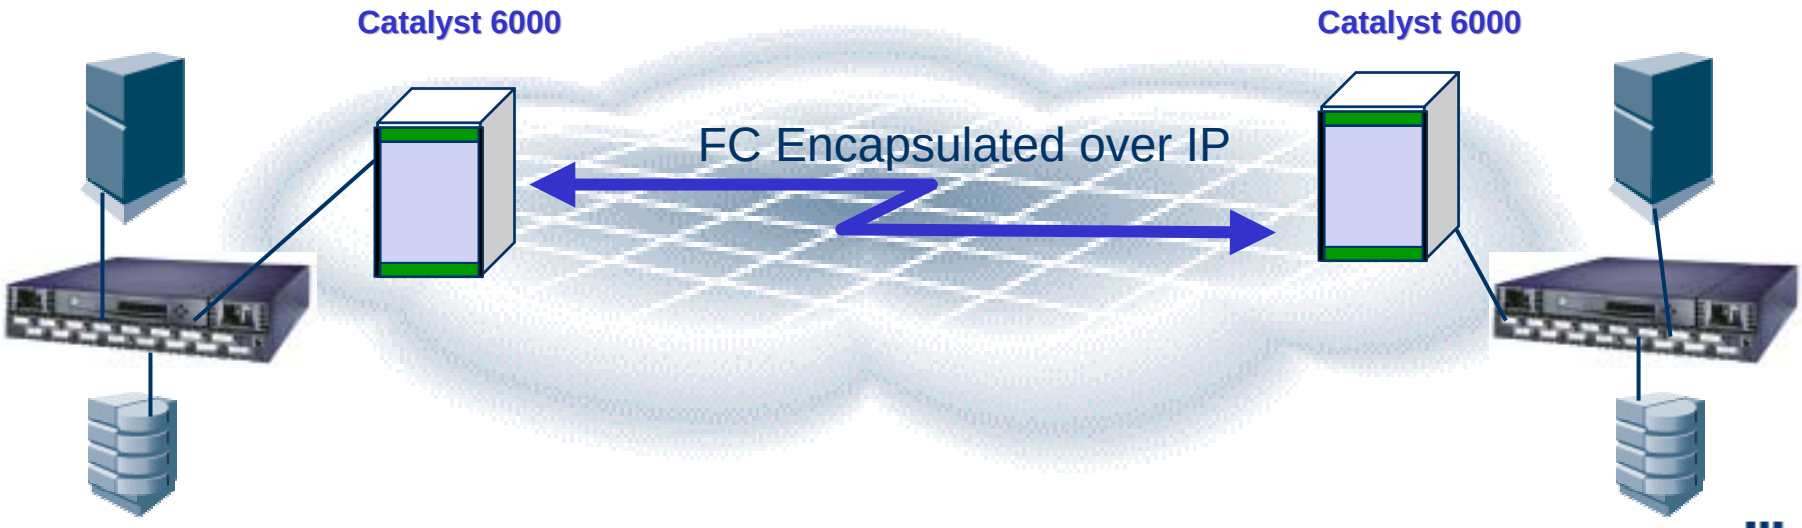

Comparison of Remote Switch and Extended Fabrics

	Remote Switch	Extended Fabrics
Connection Type	FC over WAN Gateway	Native Fibre Channel
Line Speed	155 Mb/s - OC-3 or OC12	1 Gb/s nominal
Maximum Distance	> 3000 km using WAN	< 100 km
Interconnect	ATM or T3 Requires Open Systems Gateway (OSG) at each site.	Extended Distance GBICs; LWL with repeaters; WDM
Topology restrictions	2 switches 1 local / 1 remote	Maximum 239 total switches (local or remote)

EMC SRDF / Brocade Switch / CISCO DWDM
DR 시스템 구성도

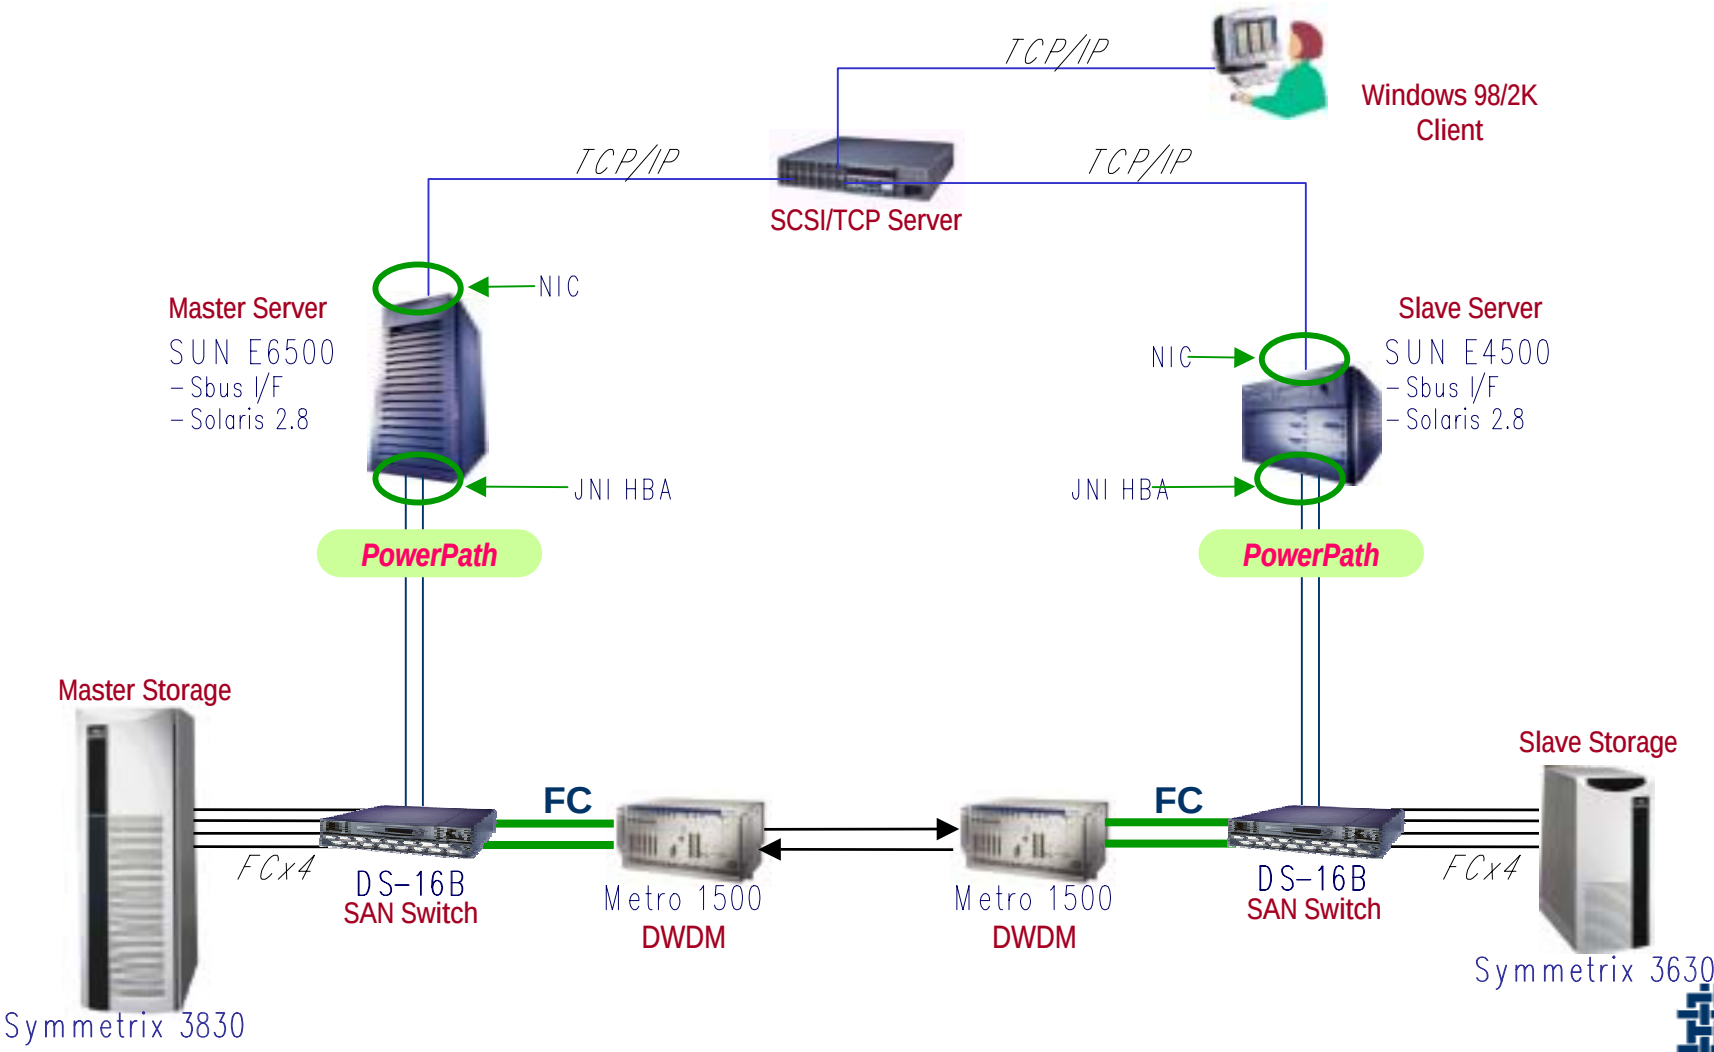

Microsoft Exchange Server High Availability Using EMC SRDF and DWDM

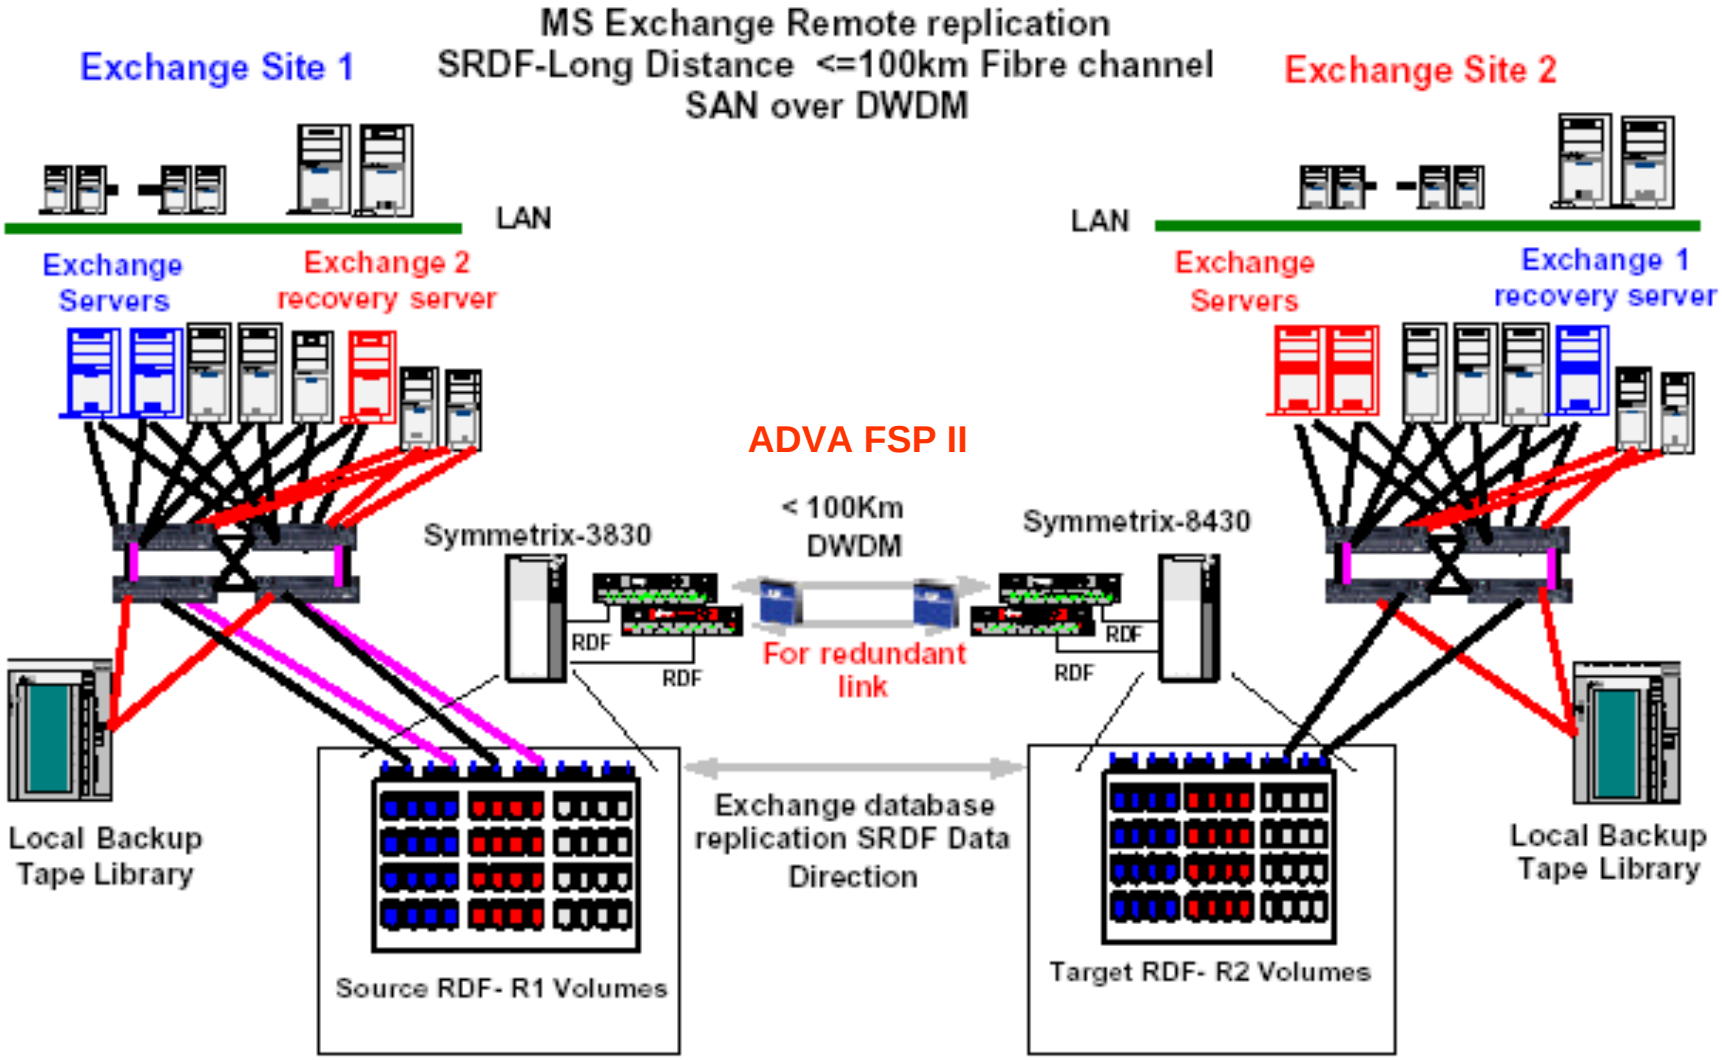

Solutionware: Remote Mirroring (ADVA DWDM)

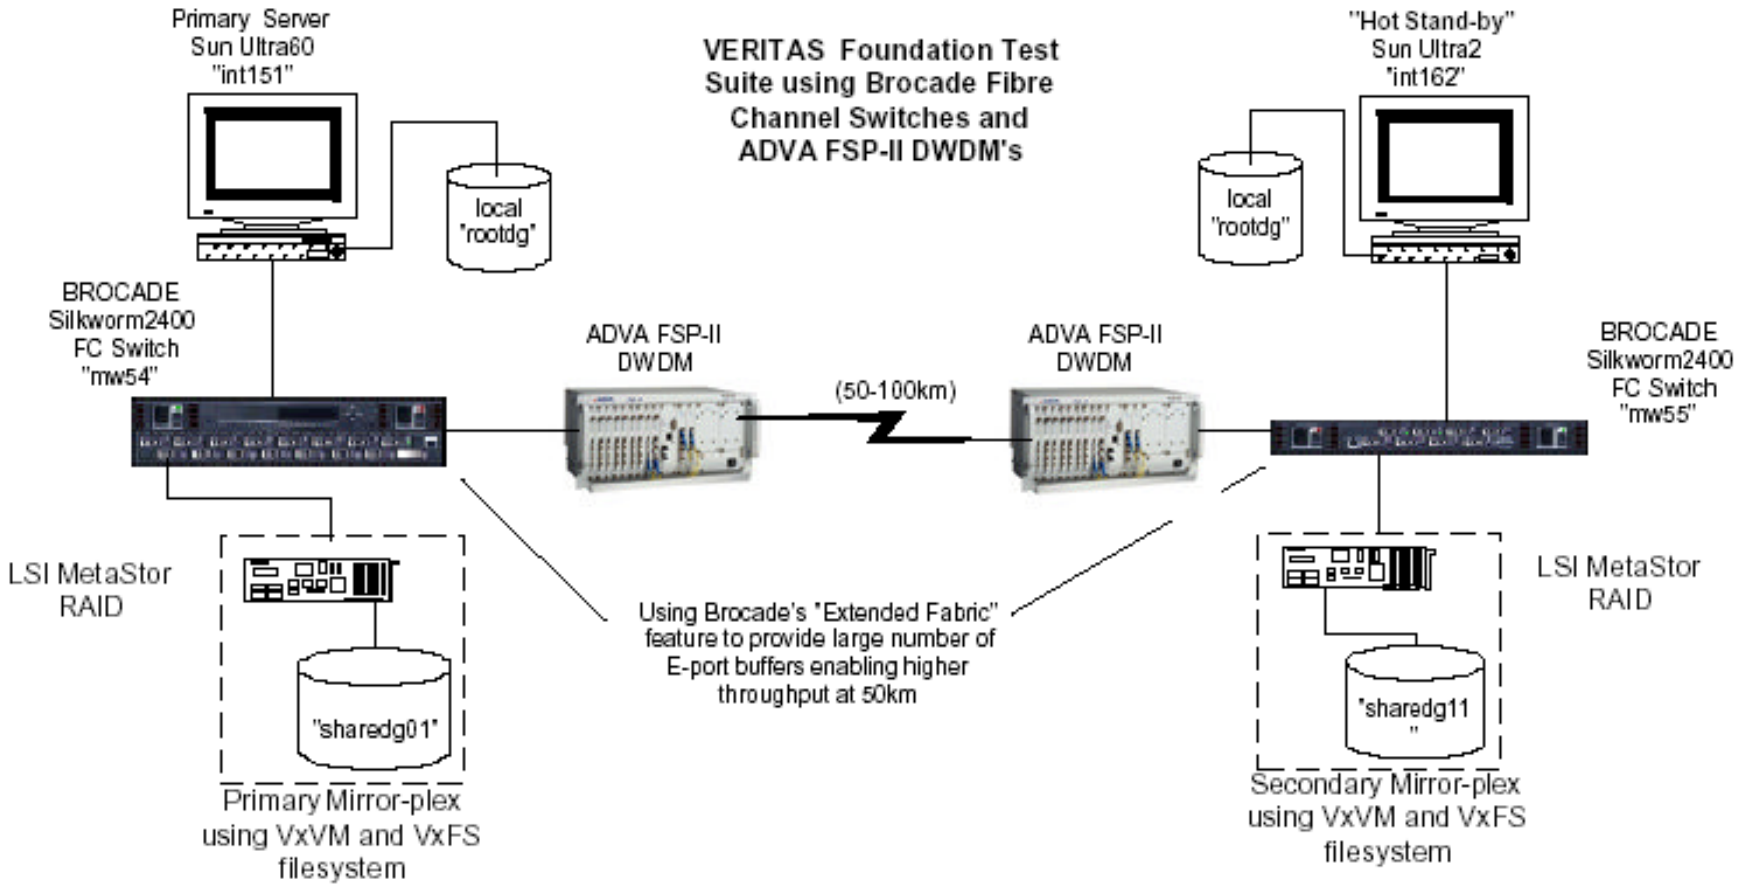

VERITAS Remote Mirroring over MAN using DWDM

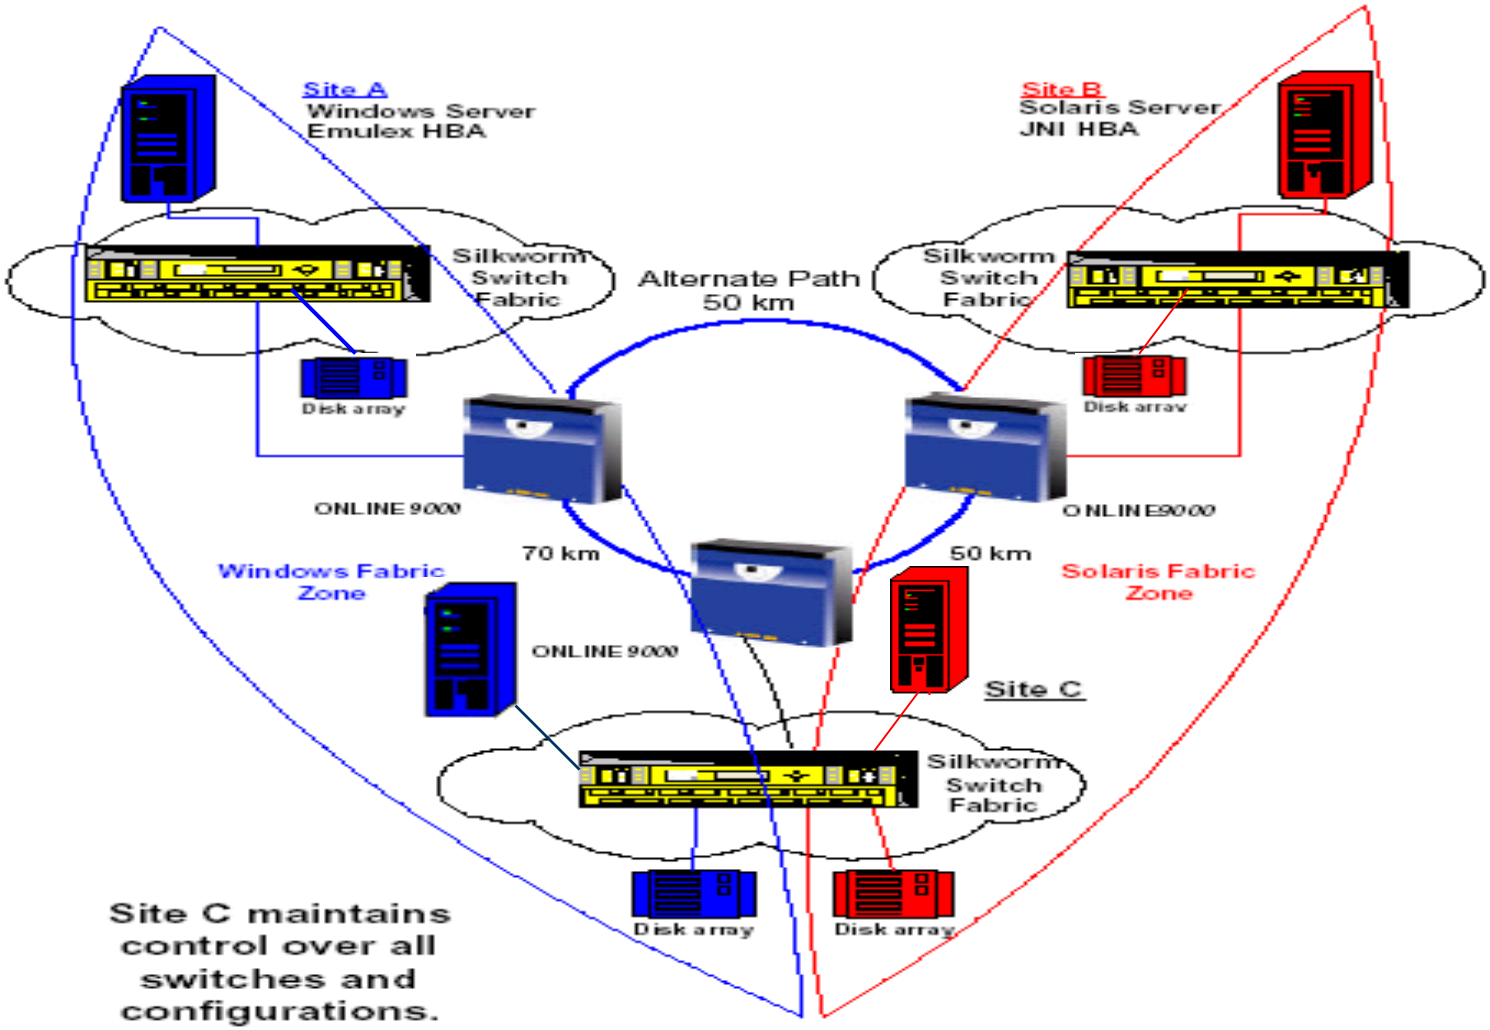

Solutionware: SAN over MAN Data Center

Centralized Remote Backup, Remote Mirroring, Centralized Storage Consolidation

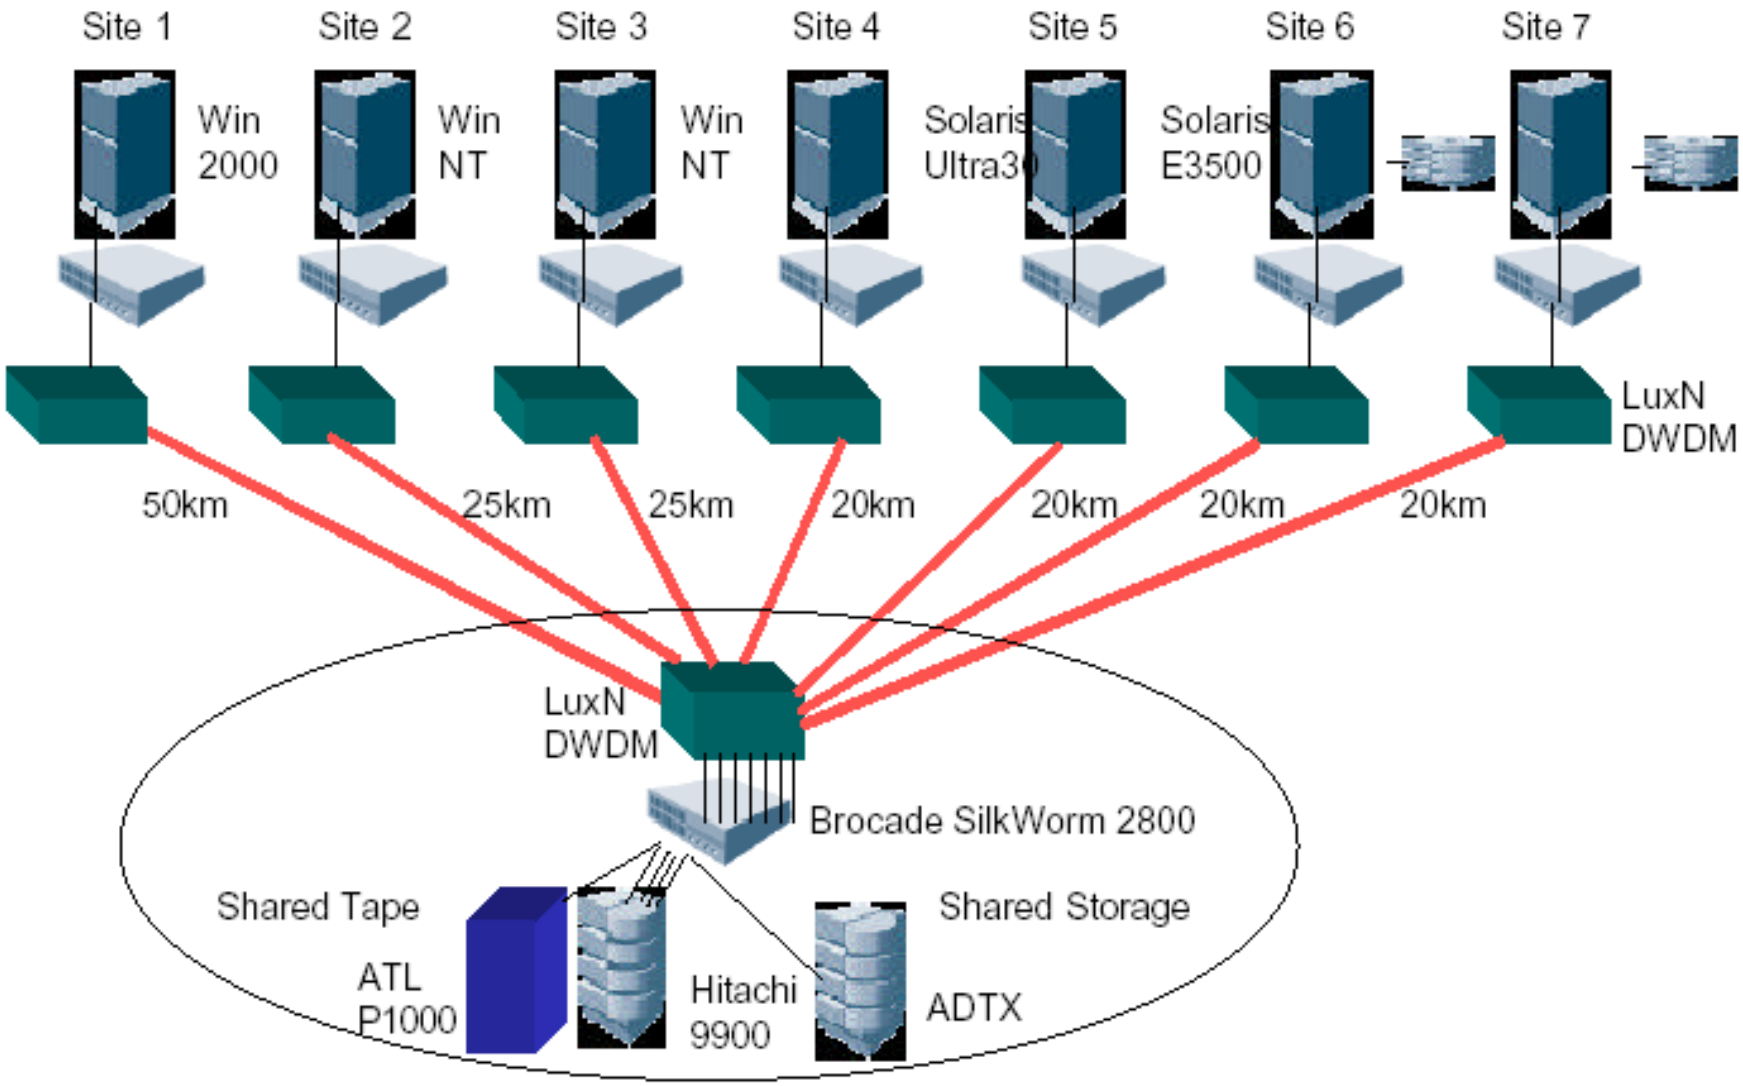

Business Continuance Case Studies

국내 최대 통신사 DR Configuration

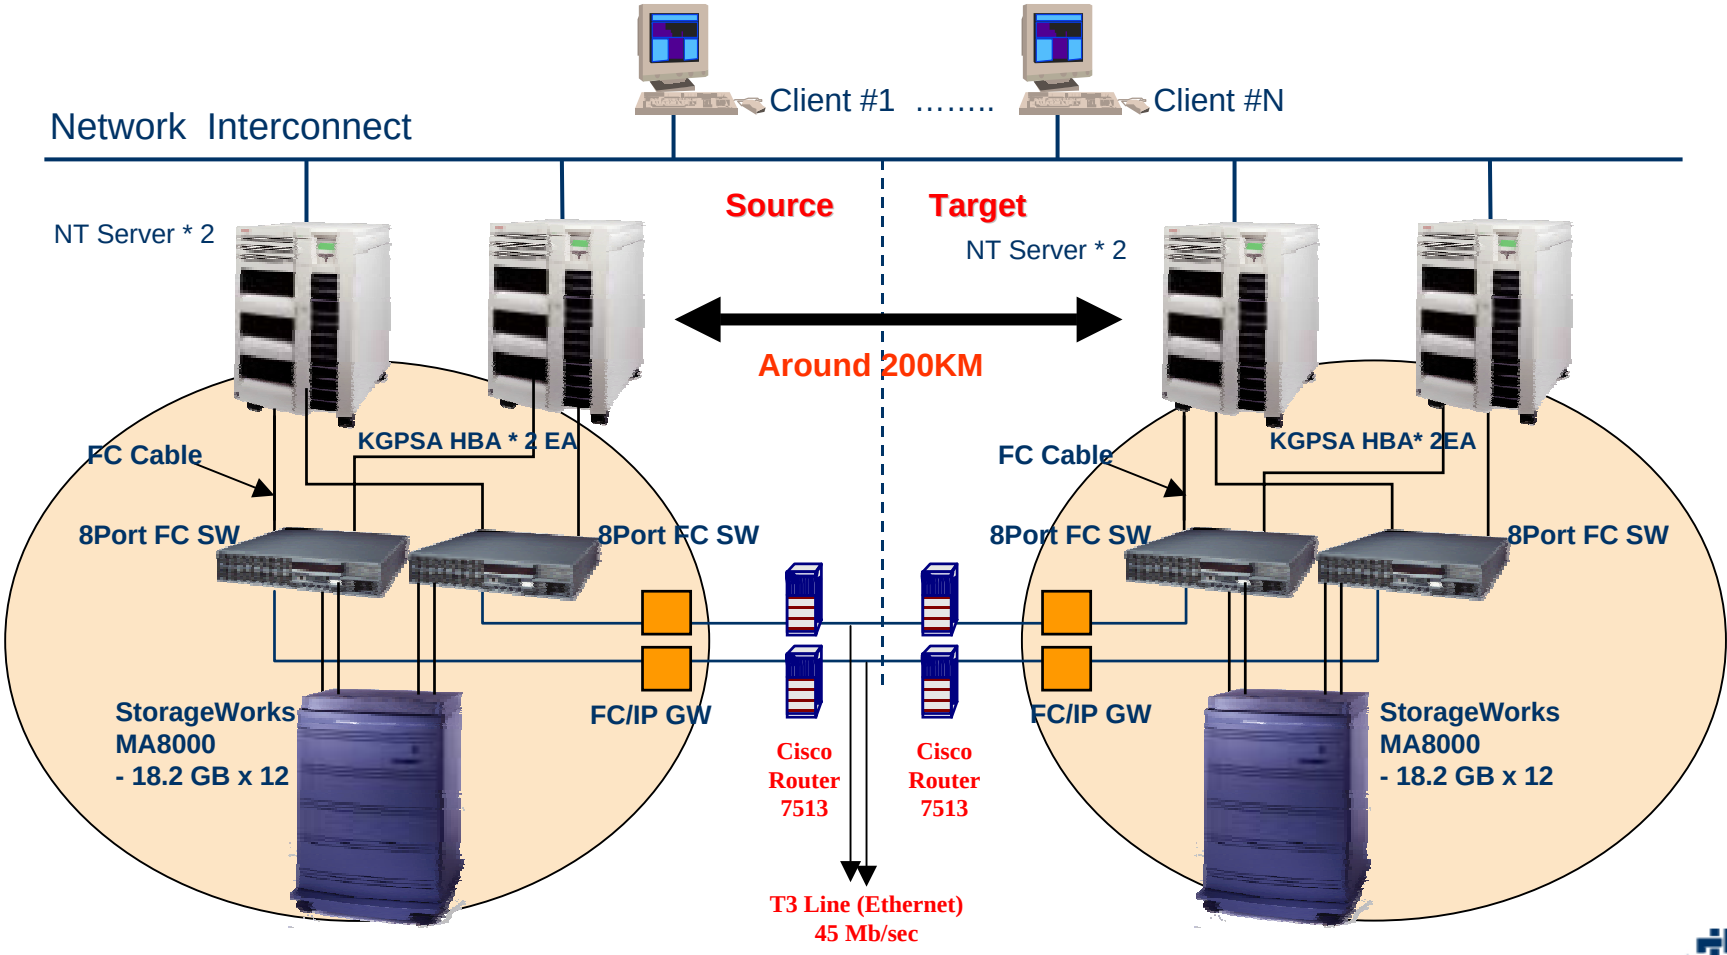

Fiber Channel and Ethernet

Different Networking Approaches for Different Applications

Ethernet	Fiber Channel
Maximum Transmission Unit Size 1,500 bytes	Maximum Transmission Unit Size 64,000 bytes
Transaction Orientated Maximum Sequence 1,500 bytes	Block Orientated Maximum Sequence 128 MBytes
High CPU Utilization IP Transfer 1GB File: ~700 packets Gigabit Ethernet: 80,000 Interrupts /sec	Low CPU Utilization IP Transfer 1GB File: 16 Frames Fiber Channel: 1,600 Interrupts / sec (98%)
Large Processing Overhead	Processing done in hardware
Collision Domain No In-Order delivery	Serial Transport Guaranteed In-Order Delivery

Morgan Stanley

Challenge

- Needed to achieve business continuance in a highly available, disaster tolerant computing environment

Solution

- Dual SAN fabrics, inter-networked over DWDM, server clusters running business critical applications

Result

- Complete fail-over to redundant site possible in under 2 minutes
- Reduced time-to-market: can create storage pools and pre-build capacity for instant availability
- Now supports heterogeneous storage environment; able to select servers and storage of choice

Depository Trust and Clearing Corporation (DTCC)

Background

- World's largest securities depository and clearinghouse for settlement of securities trading
- Processes nearly 100% of broker-to-broker equity and corporate/ municipal bond trades in U.S.

DTC		*NTSC*	
Financial services customers worldwide	11,000	Value of Transactions	$105 Trillion
Value of Securities on Deposit	$23.1 Trillion	Number of Transactions	2.6 Billion
Average Daily Trades	1.3 Million		

Challenge

- 7x 24 x 365 mission critical environment
- Required scalable infrastructure for growth

Solution

- Highly available, fault-tolerant SAN infrastructure
- Dual fabric, inter-networked using DWDM technology

Disk & Tape SAN – Phase 1

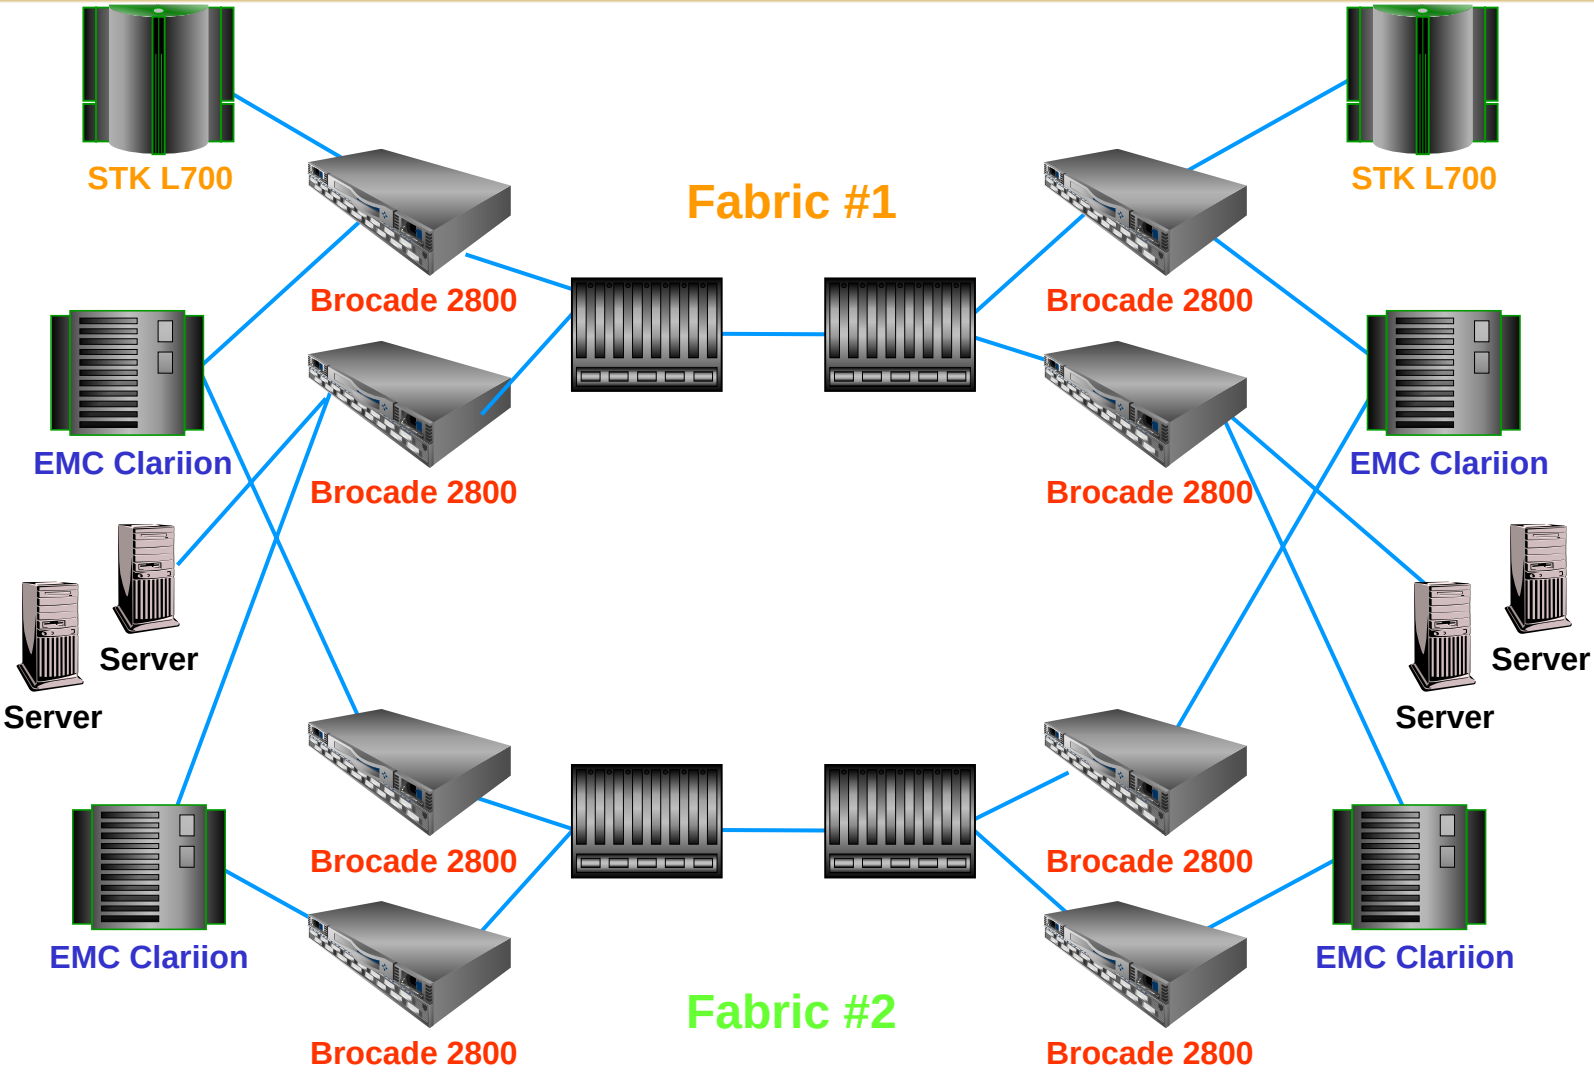

SAN Expansion – Phase 2

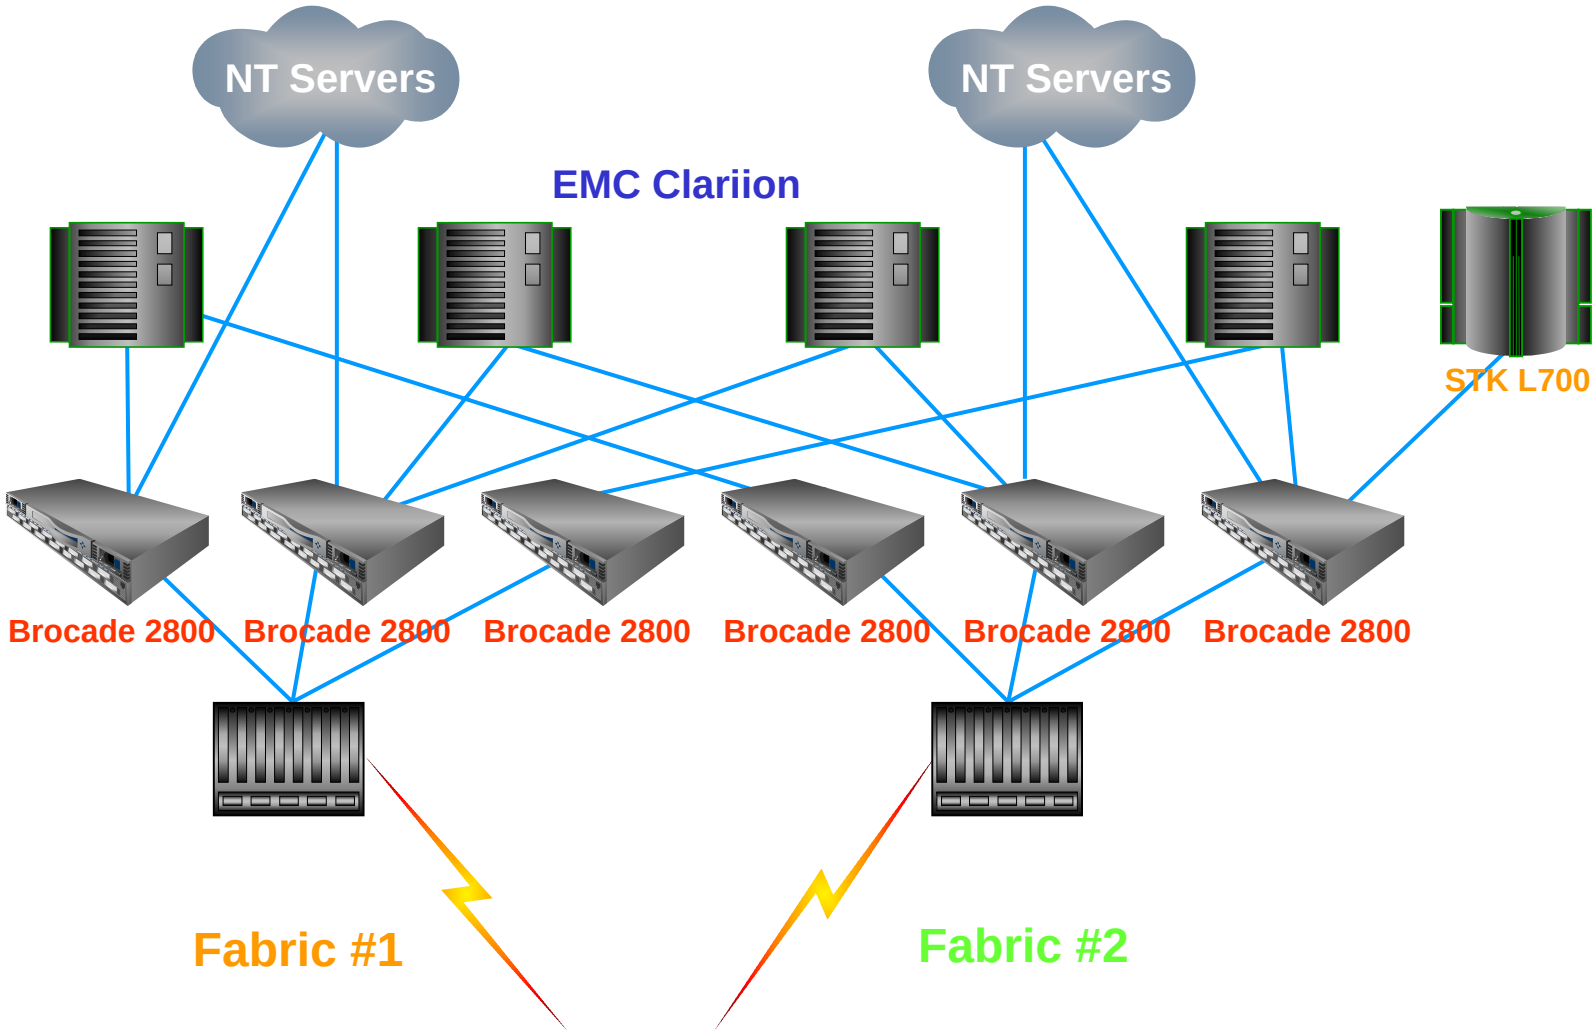

Next Step – Core/Edge Design

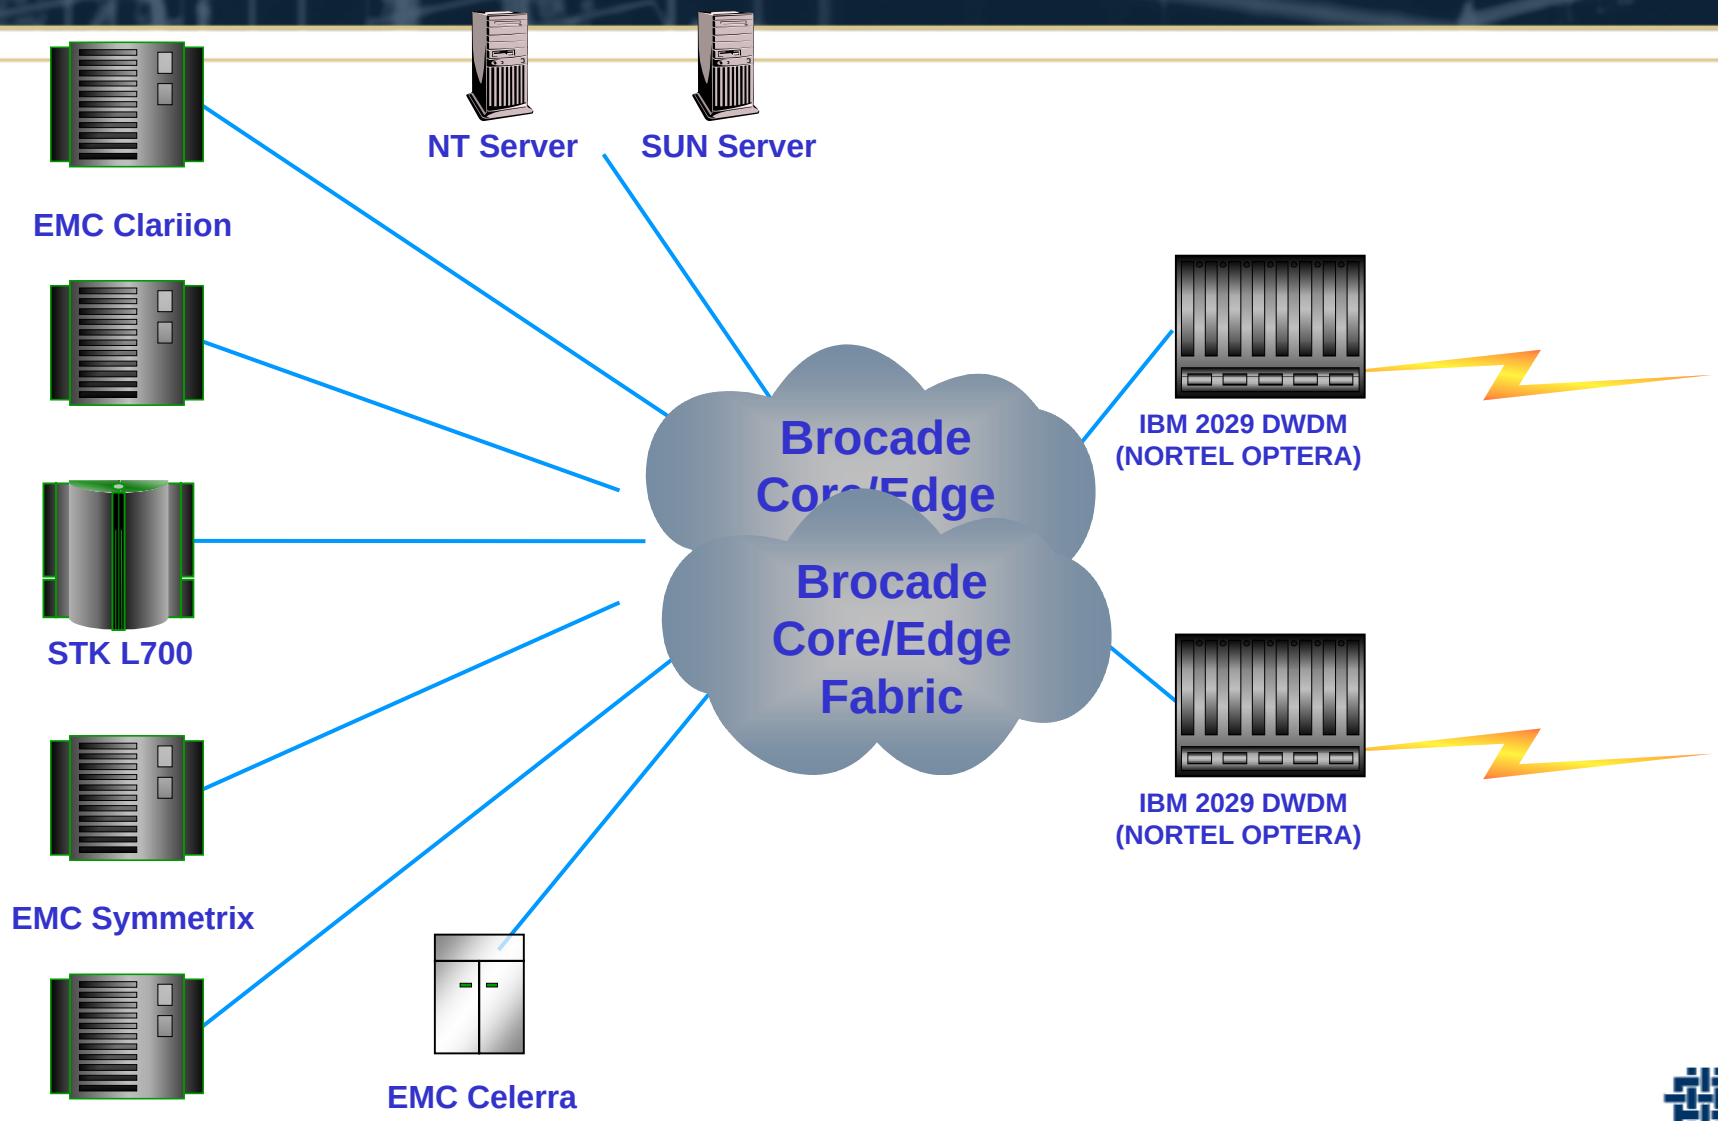

Depository Trust and Clearing Corporation (DTCC)

Result

- **SAN enabled a reliable, fault tolerant storage environment**
  - No single point of failure
  - Enhanced disaster recovery: failover from primary data center to DR site achieved in under 5 minutes
- **Optimized data storage and personnel resources**
  - Scaled easily from 2 to 10 Terabytes of storage in 4 months
  - Can now manage four times as much data with same personnel resources
  - Scaled from 3 servers to 200 without adding operations staff
- **Gained operational benefits**
  - Reduced backup time by 75%

A Leading SSP Case Study

- **Design goals**
  - Support multiple customer storage requirements in a major metropolitan area
  - Long distances (48Km) required for disaster protection
- **SAN backbone connected to ONI DWDM and SONET RING for long distances**
  - SAN islands isolate customer environments
  - SAN backbone (core switches) used for centralized storage
- **Environment**
  - Key applications: Oracle OLTP and Data Warehouse
  - Key platforms: Solaris, NT, AIX, HP-UX
  - Approx. 100TBs storage using EMC, Compaq, and HDS
  - Management: CA Unicenter
  - Tape Backup: Veritas NetBackup

Single Fabric Topology with 44 Brocade 2800 switches

A Leading SSP Case Study (cont'd)

New York City

New Jersey

ONI DWDM
SONET RING

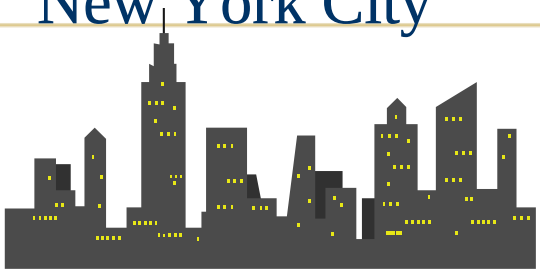

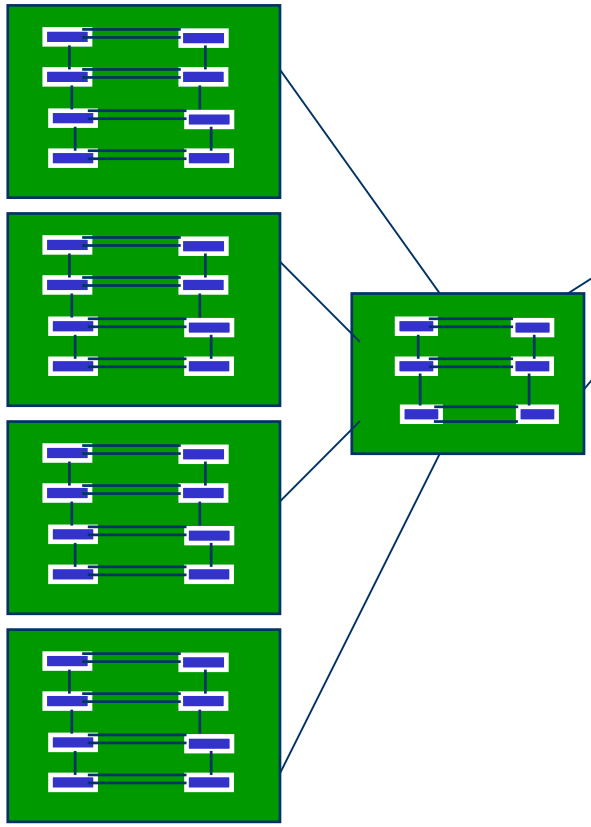

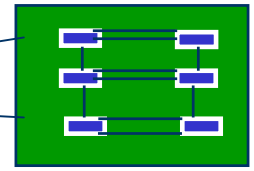

- Four SAN islands NYC
  - 1 island per customer
  - 8 switch meshes
- Six-switch core mesh
- Six-switch SAN island NJ
  - Mirrored EMC and Compaq disk
- Brocade Extended Fabric software

Thank you