WHERE NETWORKS CONVERGE

HP World August 24, 2001

Session #T460

Building, Managing, and Growing High Availability Storage Networks

Greg Schulz – greg.schulz@inrange.com
612-810-9890 (651-777-4559 Fax)

INRANGE Technologies Corporation
100 Mount Holly By-Pass
P.O. Box 440
Lumberton, NJ 08048-0440

www.inrange.com

Building Large Storage Networks

Agenda

- **Why build a large Storage Network?**
- **Storage Networking Architectures**
- **Storage Networking Design**
- **Storage Networking Management**
- **Global or long distance Storage Networks**
- **Storage Networking solutions**
- **Summary and closing comments**

WHERE NETWORKS CONVERGE

Building Large Storage Networks

Why build a large Storage Network?

- Technology exits to support large Storage Networks
- Backup, server and storage consolidation are drivers
- Large Open Systems and OS390 environments
- Consolidation of Storage Networking SAN Islands
- IDC estimates continued SAN port growth per location

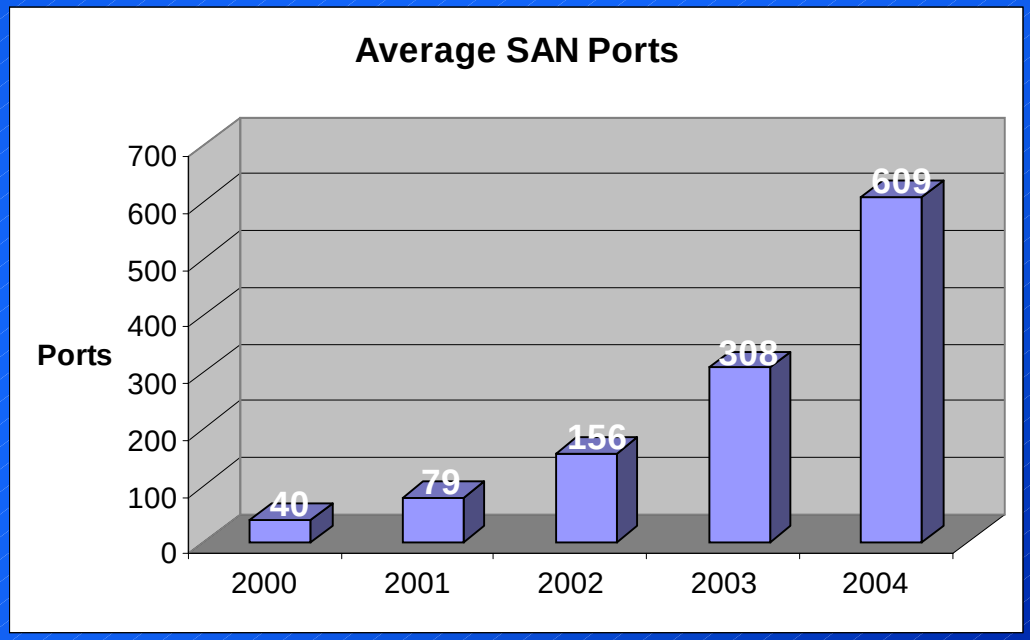

Building Large Storage Networks

Some Storage Networking Terms

- Over-subscription or Congestion
  - More traffic or workload demand than available capacity or bandwidth can handle. Caused by two few ISL's in a fabric, inappropriate use of available technology and topologies.
- Blockage or Blocking
  - Delays experienced by host/servers and storage devices from fabric over-subscription or blockage.
- ISL (Inter Switch Link)
  - Connections between switches or directors to create a mesh or multi-stage fabric to support more ports than a single device is capable of. ISL can interconnect geographically dispersed or separate storage networks.
- ISL Trunking
  - Grouping ISLs to create more aggregate bandwidth to help blocking and inter-switch bandwidth bottlenecks in a fabric. A work-around to using a non-block core director.

Building Large Storage Networks

Some Storage Networking Terms

- Non-Blocking
  - All ports have full bandwidth with no congestion even at 100% workload. Switches and Directors are non-blocking via their internal data paths.
- Switch
  - Smaller Port Count device ranging from 8 to 32 ports in a small self-contained form factor. Multiple switches can be connected using ISLs to create a cascade, mesh, or multi-stage fabric.
- Director
  - Large Port Count device ranging from 32 to over 128 ports performing functions similar to a switch with more scalability, redundancy, non-blocking, and multiple protocols including Fibre Channel, IP, and FICON. A director replaces multiple smaller switches configured in a cascade, mesh or multi-stage fabric. Scalability beyond 100s of ports is possible with in-place upgrades, enhanced security, and simplified management.

Building Large Storage Networks

Some Storage Networking Terms

- Latency
  - How long it takes for data to get from source to destination through a switch, director, or a fabric. Latency can increase with more hops, larger fabrics, and over long distance (WAN & DWDM) links.
  - Latency is measured in microsencds and can become a performance issue for larger storage networks and large fabrics particularly for time sensitive applications like Video and OLTP.
- Bandwidth
  - Amount of traffic measured in Megabytes (MB) that can be moved per second over a given link or “pipe”. This is important for video, backup, and data warehousing or other throughput sensitive applications.
- IOPs (I/O Operations Per Second)
  - Number of Input/Output operations per second on a port or path. This is also the number of messages or I/O commands executed per second per path.

Building Large Storage Networks

Some Storage Networking Terms

- Zoning
  - Isolating ports, traffic, devices, applications, customers, or platforms via hard port, soft name server, or wwn zoning for security and management purposes. Similar to creating networking firewalls or Virtual Private Networks (VPNs).
- Locality or Placement
  - Tuning or configuration activity to reduce ISL traffic, and congestion to improve performance. Places hosts/devices close together to reduce latency for load-balancing.
- Usable Port Count
  - How many usable ports for servers or storage devices on a switch, director, or fabric not consumed by ISLs. The higher the native port count, the simpler management and security are.
- Hop Count
  - How many switches or directors a Fibre Channel I/O or message must travel from source to destination

Building Large Storage Networks

Some Storage Networking Terms

- Protocol Agnostic
  - Today's storage networking director class products support multiple Upper Level Protocols (ULPs) including SCSI_FCP (Open Systems Block data access), IP_FCP (IP over Fibre Channel), VI (Virtual Interface for clustering), and FICON (OS390 Block data access) to name a few. In addition to various ULPs, different physical interconnections can be supported including 1Gbit and 2Gbit Fibre Channel along with emerging interfaces (iSCSI, InfiniBand, etc.) as they become standard and supported.
- ASIC
  - Application Specific Integrated Circuit is a "chip" that has been customer programmed to perform specific functions including implementing. An ASIC can implement and process the Fibre Channel protocol to speed up operations of common functions as opposed to using general purpose processors and other components adding to cost and complexity.

Building Large Storage Networks

Some Storage Networking Terms

- NAS
  - Network Attached Storage is a generic term referring to file or data sharing using common protocols like NFS (Network File System) or Windows CIFS (Common Internet File System) over standard TCP/IP networks. Storage is served from a server or host sometimes called an appliance via the network to clients.
- SAN
  - SAN stands for Storage Area Network and is often associated with Fibre Channel Block type data access. Block level data access enables database, OLTP, video, backup, and other high performance streaming applications to access storage directly in a deterministic manor.
- Core and CLOS architectures
  - CLOS architectures including core and cross-bar type technologies are based on work by Dr. Clos during the late 50's pertaining to high performance, non-blocking technology to create large, low latency, devices.

Building Large Storage Networks

Storage Networking Architecture - Scaling

- What is theoretical
  - 2^{16} possible Fibre Channel Addresses
- What is possible
  - 239 x 128 or 256 Port Directors in a single fabric
  - 30,592 total ports
- What is practical today
  - 128 ports in a single non-blocked director
  - 256 ports in a dual or alternate path director config.
  - 100's to 3,000's of ports per storage network
  - Mixed architectures and platforms (Open and OS390)
  - Mixed protocols and topologies (FICON, SCSI_FCP)
  - Core and edge combinations

Building Large Storage Networks

Storage Networking Architecture – Building Blocks

Some Storage Networking
Building Blocks

8 Port Switch

16 Port Switch

24-64 Port Director

24-128 Port Director

Some Storage Networking Topologies

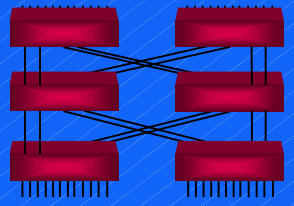

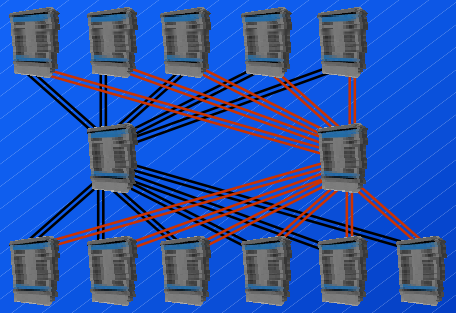

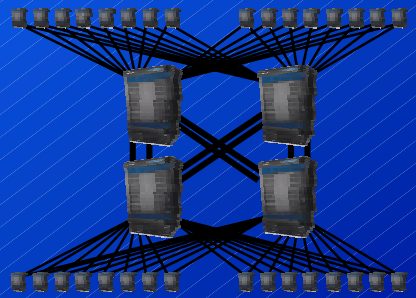

64 Port Mesh

Dual Path Directors 48-256 Usable Ports

1,110+ Port Multi-Stage

3,000+ Port SAN

Building Large Storage Networks

Storage Networking Architecture - Topologies

- Growing to a 128 port storage network
  - Non-blocking for high performance
  - Predictive latency
  - High availability, serviceability, lower TCO
- Some options:
  - 128 Port Non-Blocking Redundant Director
  - 24 x 16 Port layered or multi-stage switches

Simplicity

128 Port Director

Complexity

Meshed Switches

Building Large Storage Networks

Storage Networking Architecture – Building Blocks

- Fibre Channel Switches
  - Similar in function to LAN networking switches for connectivity
  - Utilize ASICs with interconnect busses, or paths for connectivity
  - Typically utilize GBICs (Long and short-wave, copper, optical) for modular link level interoperability and investment protection
  - Can be configured into various topologies including cascade, mesh, multistage, and core & edge fabrics to meet application requirements

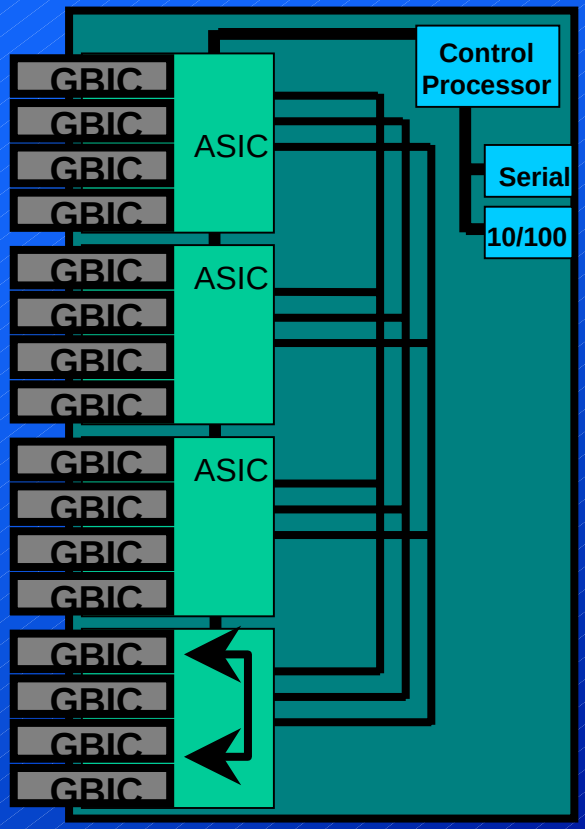

Building Large Storage Networks

Storage Networking Architecture – Building Blocks

- Directors
  - Core or CLOS architecture
    - Protocol agnostic, highly scalable
    - Hot Swappable components
    - Parity & CRC on Data Paths
    - Integrated security (non meshed)
  - User Ports Cards/Blades/Boards
    - Hot swappable GBICS and cards
    - E_Port, FSPF, FC-SW2, FC-MI interoperability
    - Variable port types (Fabric, Loop, Trunk, ISL, FICON)
    - High port count from 24-128 ports with in-place upgrades
    - Multiple Protocols (IP, FICON, SCSI) and interfaces
  - Control or Management Module
    - Non Disruptive Code Load (NDCL)
    - Hot swappable, N+1 redundancy
    - Redundant 100mb Ethernet ports
    - SNMP Fibre Alliance MIB Agent
    - Call home capability

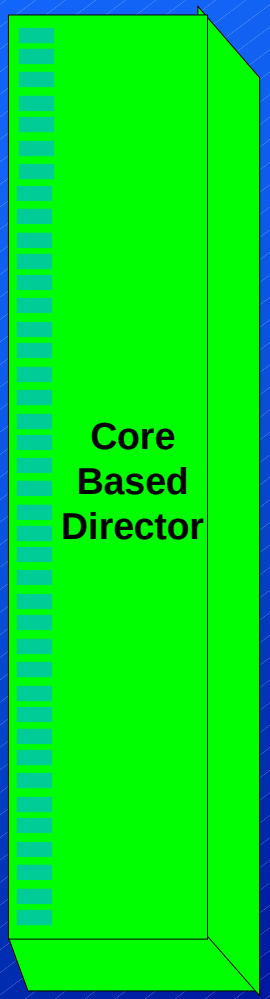

Building Large Storage Networks

Storage Networking Architecture – Multi-Platform

- Open Systems (Unix, NT/W2K, Linux, Novell, Etc.)
  - SCSI_FCP (Fibre Channel) for Block Data Movement
  - IP over Fibre Channel for NAS/Server Area Networking
  - VI emerging and low latency, high speed clustering interface
  - Mixed CPU platforms (HP, Sun, IBM, SGI, Dell, Compaq, Etc.)
  - Mixed Storage platforms (Disk, Tape, EMC, HDS, IBM, Etc.)
- OS390 Mainframe
  - FICON (ESCON over Fibre Channel) for Block Data Movement
- SAN and NAS access methods (Block & File)
- Interoperability and investment protection
  - Servers/HBAs, Switches/Directors, Loop/Fabric, Storage

WHERE NETWORKS CONVERGE

Building Large Storage Networks

Storage Networking Architecture – Multi-Platform

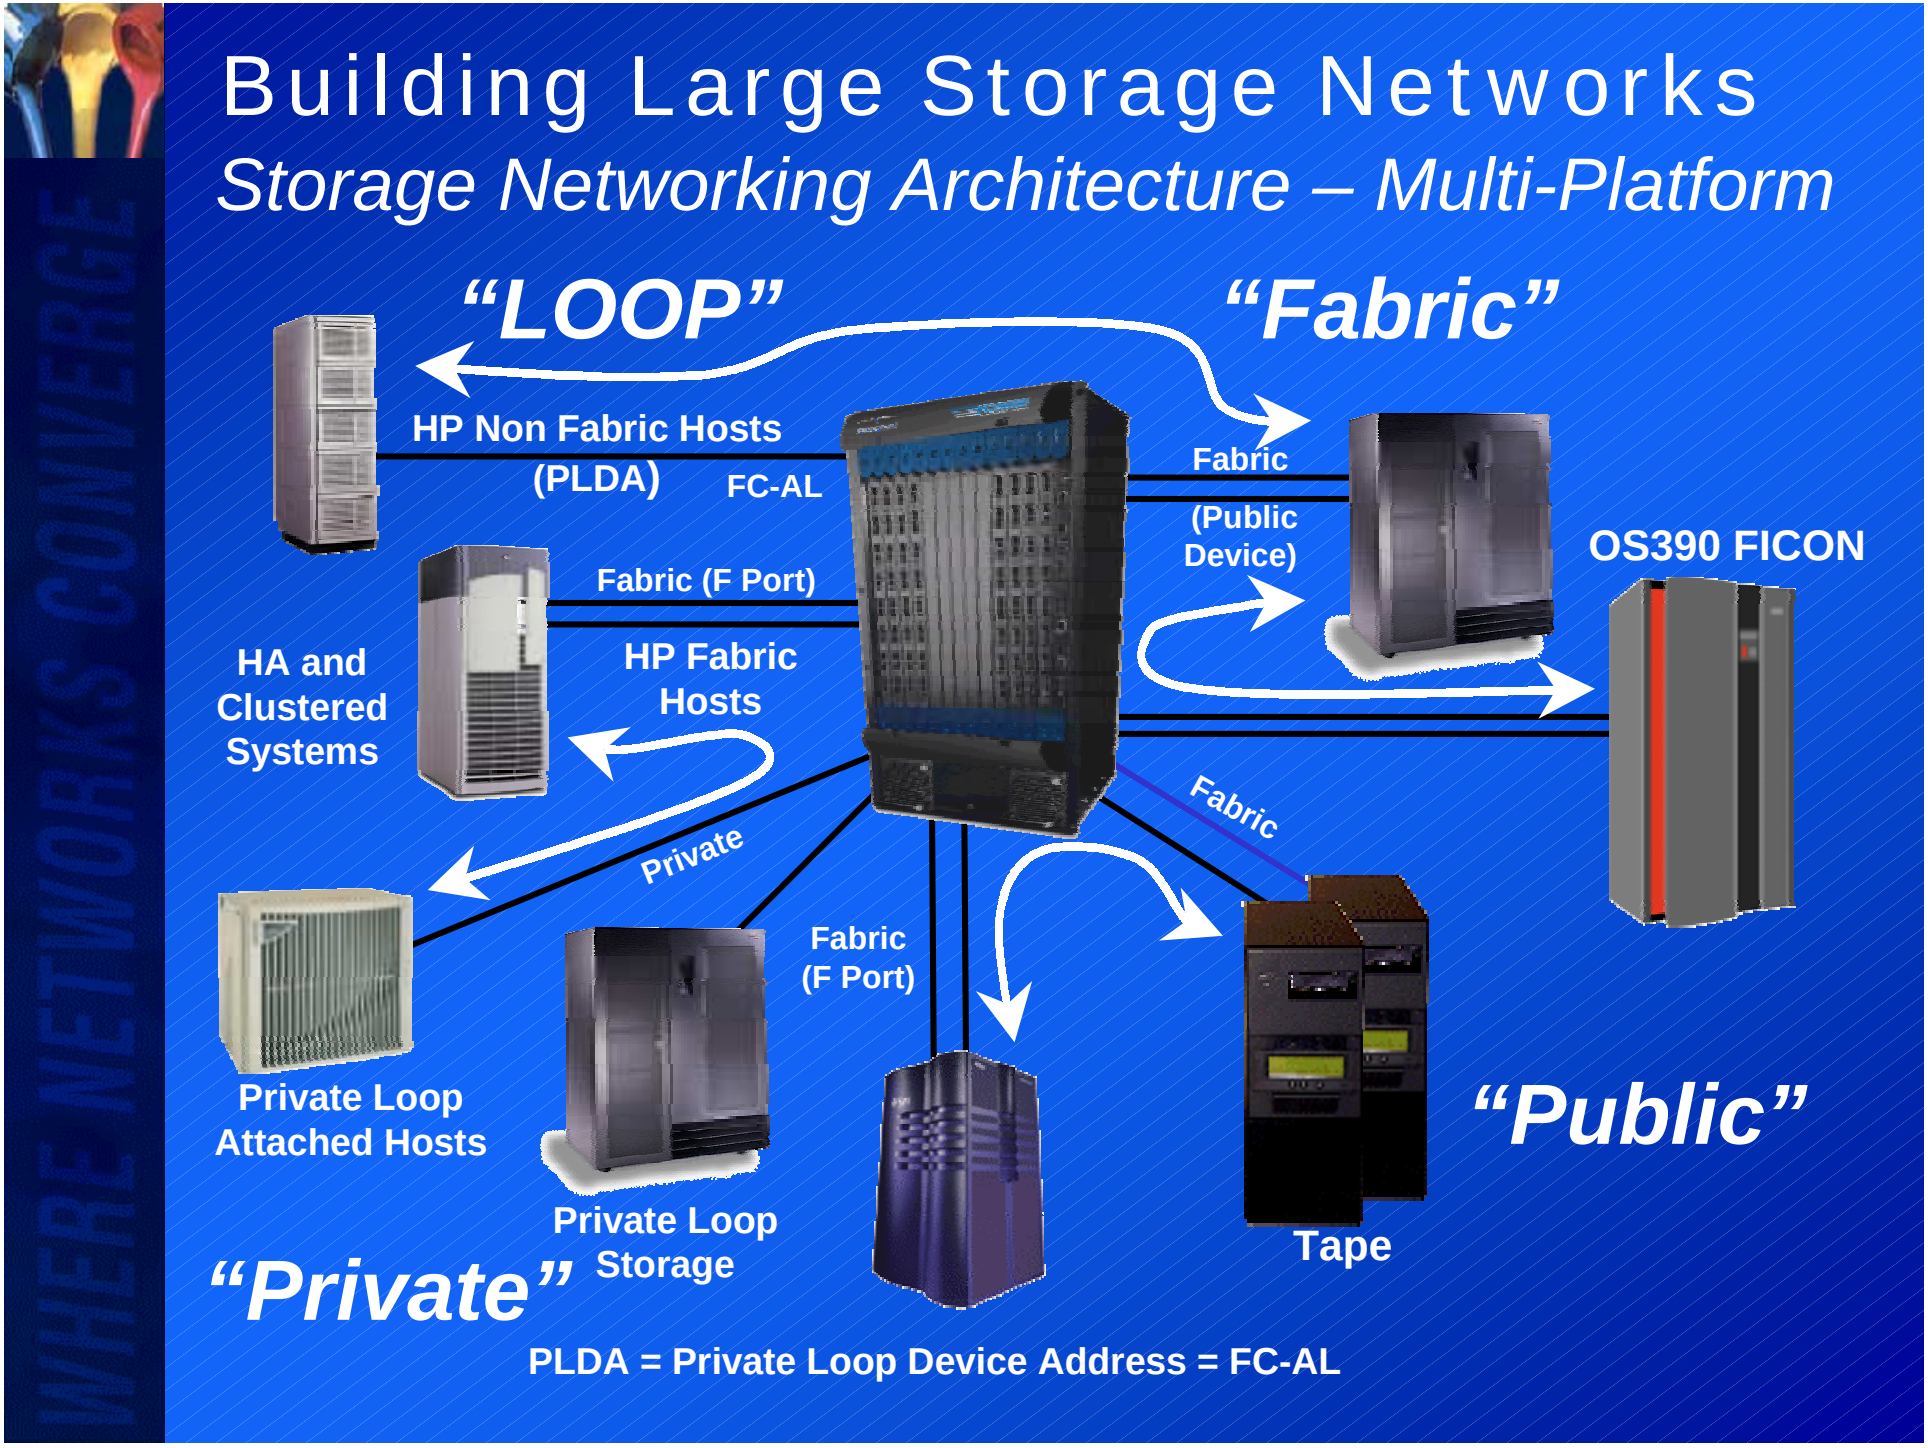

PLDA = Private Loop Device Address = FC-AL

WHERE NETWORKS CONVERGE

Building Large Storage Networks

Storage Networking Architecture - Performance

Over-Subscription, Blockage

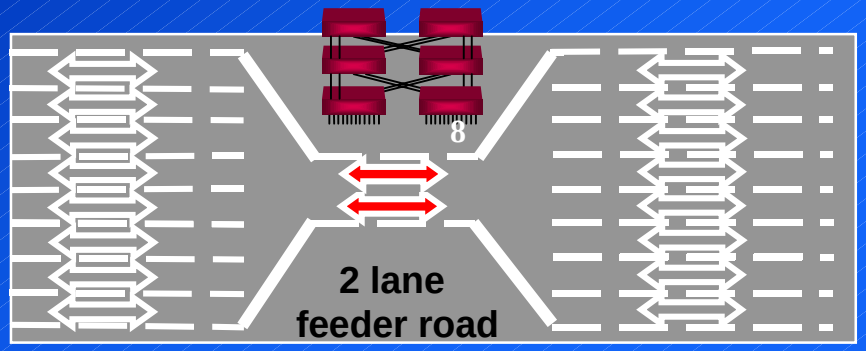

Over-subscription Induces traffic delays, priority juggling, loss of user ports, need for trunking, loss of bandwidth and extra management

Non-Blocking
Performance, Consistent Latency

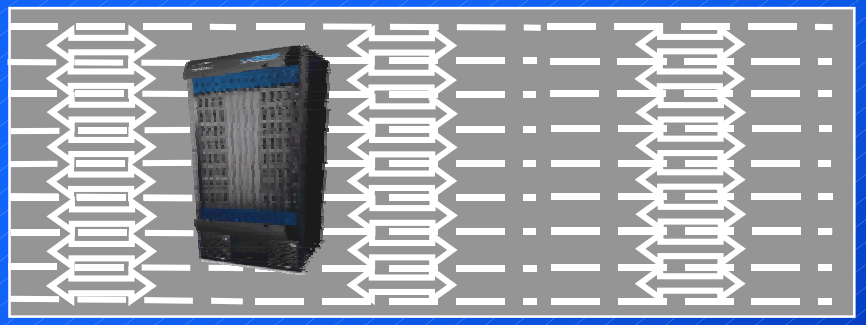

No Blocking, No Over-subscription, No Priority Juggling, No Need for Trunking, No Loss of User Ports

Building Large Storage Networks

Storage Networking Architecture - Performance

- Over subscription results from
  - Use if mesh or multi-stage fabric with ISLs for scalability
  - Similar issues were common and still exist in some LANs
  - Design trade-off of “cheap” ports, complex management
  - To few ISL’s for a given fabric bandwidth for given workload
  - To many hops between storage device and host
  - Backup will place throughput burden on storage network
- Loss of user ports resulting in higher costs per port
  - ISL take away from user ports
  - More ISL’s are needed to support bandwidth
  - Storage requires deterministic performance
  - Servers can talk to Servers (Server Area Networks)
  - New SANs have storage talking to storage

Building Large Storage Networks

Storage Networking Architecture - Performance

- Two approaches for scalability
  - Core director with many any to any high performance ports
    - Non-Blocking core directors remove complexity and enable scaling
    - Single device vs. many smaller devices “tied” together
  - Mesh fabric of multiple smaller switches with Trunked ISLs
    - Enable scaling beyond capacity of single device
    - Leverage the right technology and approach to meet needs
    - Watch for hidden costs and complexity vs. perceived savings
    - Routing algorithms like industry standard FSPF are needed
    - E_Ports are used for ISL and trunking vs. Core architecture
    - Load balancing and priority juggling algorithms and tools needed
    - More components, lower MTBF, more points of failure
- Fibre Channel Interfaces
  - Currently 1Gbit Fibre Channel (100MBs) is the standard
  - 2Gbit Fibre Channel is being deployed (10Gb in design)
    - HBAs, Switches, Devices later in 2001

Building Large Storage Networks

Storage Networking Architecture - Locality

- Locality and reducing blocking, lowering latency
  - Keeping traffic on a single switch or director
  - Utilizing switch/director pairs for performance and redundancy
  - Smaller switches or directors can cause less localization
    - Excessive storage networking hops and latency
    - Complex management and security risks with fabrics
    - Hidden costs of management (Look at TCO)
  - The larger the switch or director, the more localized things are
    - Less congestion resulting in simpler management
    - Better performance or reduced latency
    - Avoid excessive storage networking hops
    - Removes complexity with future expansion
    - Simplify trouble shooting and diagnostics

Building Large Storage Networks

Storage Networking Architecture - Locality

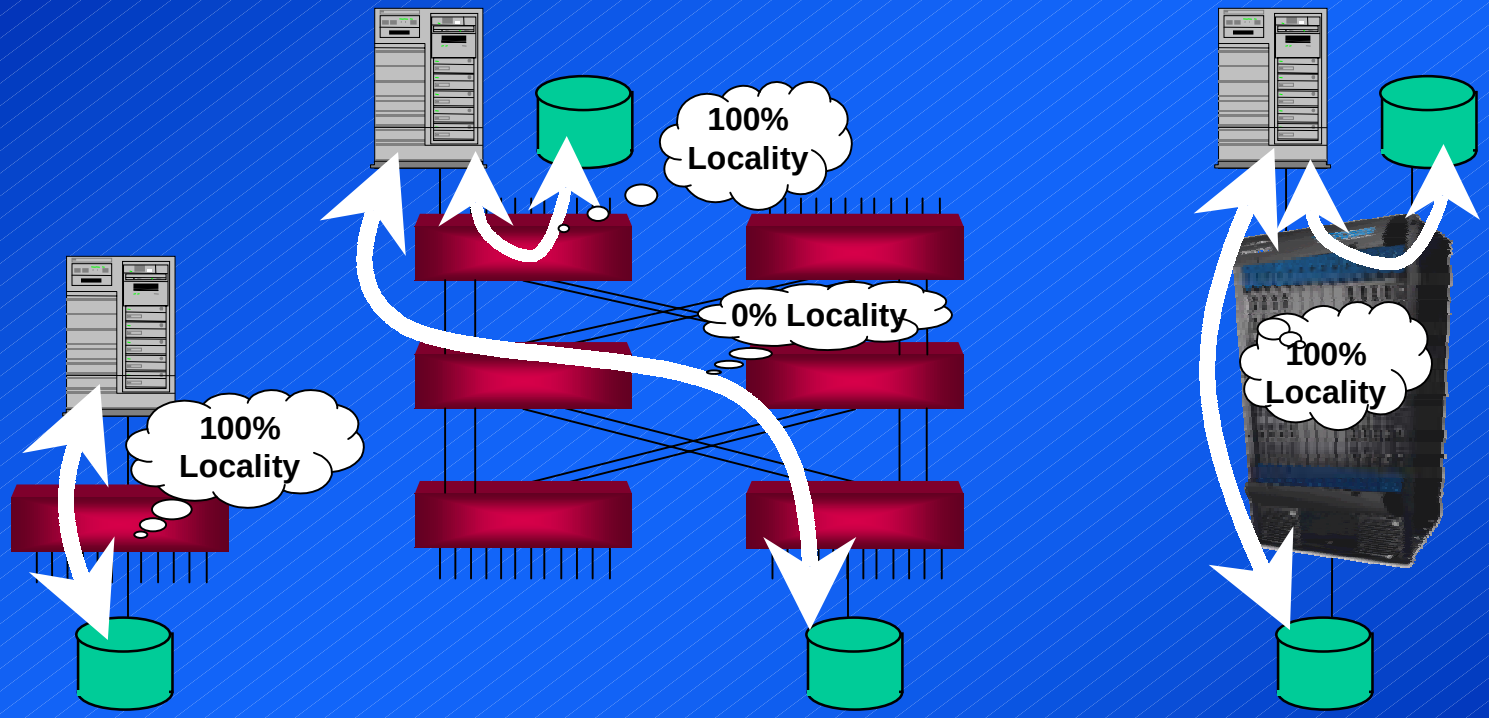

16 Port Switch

Good Locality & Performance 1 Hop

64 Port Six Switch Mesh

Poor Locality & Performance 3 Hops Beyond Local Switch Ports

64 – 128 Port Director

Great Locality & Performance 1 Hop

Building Large Storage Networks

Storage Networking Design – Factors & Variables

- Number and type of hosts/servers (Unix, NT, OS/390)
- Number and type of storage devices and interfaces
- Application types (OLTP, Video, Web Serving, Email)
- Application service criteria
  - Access method (Block, File, Local, Remote, Cluster/Messages)
  - Performance (I/Os, Throughput, Latency)
  - Workload or capacity and performance plan
  - Availability criteria (HA, Cluster, HBA Fail over)
  - Backup, Recovery, Archive and Disaster Recovery
- Expansion and scalability for now and the future
- Investment protection of legacy devices including loop
- Management tools and interfaces

WHERE NETWORKS CONVERGE

Building Large Storage Networks

Storage Networking Design – High Availability

- Dual or Alternate Paths from HBA to Storage Device
- Dual Switches or Directors (separate SANs/Fabrics)
- Path Management, Lun/Volume Mapping, Mirroring
  - Path Management for HBA fail over/alternate pathing
  - Clustering software for server or host fail over
  - Lun/Volume Mapping/Masking for security
  - Remote Mirroring or replication for data protection
- Configure for redundancy to enable
  - Fail over and fault-containment
  - Isolation from name server or software faults
  - Maintenance of various components and software
  - On-line growth and expansion now and in the future

Building Large Storage Networks

Storage Networking Design – Topology Selection

- Single Core Director or large port count Switch
  - Support for Loop, FICON, good Locality, Best Performance
  - Reduces complexity and simplifies management and security
  - Provides scalability in a stable manor to 100's or 1,000's or ports
  - Good fit for test, development, and quality assurance (QA)
  - Provides good entry point to reduce fabric complexity
- Dual Directors or Switches (Alternate Pathing)
  - Best practice for High Availability (HA) environments
  - Provides fault containment and support for maintenance
  - Provides performance and scalability boost beyond fabric
  - Compliments clustering software and related technology
  - Enables easier upgrades and enhancements
  - Enhances security and management of storage network
  - Can be a simple first phase for replacing SAN islands

Building Large Storage Networks

Storage Networking Design – Topology Selection

- Mesh or Multi-Stage fabric of Directors or Switches
  - Small port count switches enables scaling to 100 or so ports
  - For large port count directors, enables scaling to 1,000's of ports
  - Trade off of high port count vs. added management of fabric
  - A mix of small switches and directors improves scalability
  - Match architecture and technology to application requirements
- Edge Devices (Directors or Switches)
  - Leverages Core directors near center of storage network
  - Similar to LAN networking with Cisco 6500 type devices
  - Places smaller switches or even high port directors near edge
  - Storage networks can be tailored to specific needs
    - Performance, availability, low cost per port, backup, etc.
    - Investment protection of older devices and components
    - Matches technology to service or application requirements

Building Large Storage Networks

Storage Networking Design - Considerations

- Strive for 100% locality of host and storage devices
  - Host to storage (Most common using SCSI_FCP or FICON)
  - Storage to Storage (Remote mirroring or replication)
  - Host to Host (NFS/NAS or Clustering using IP or VI)
- Inter Switch Links (ISL) can become bottlenecks
  - Excessive hops increase latency
- Different applications have different profiles
  - Volume managers can buffer or group I/Os together however,
  - Web servers can generate very small I/Os
  - NFS/NAS usually results in 4-8-32Kbyte I/Os
  - OLTP and database can be 2-16Kbyte+ (mixed)
  - Video and backup are throughput or bandwidth sensitive (MB/s)

Building Large Storage Networks

Storage Networking Management – Activities

- Backup and Recovery
- Configuration, Maintence, Allocation
- Performance, Tuning, Monitoring, Notification
- Security including Zoning, Lun Mapping/Masking
- Management Tools
  - Data Management and access tools
  - Configuration/Change/Allocation Management
  - Monitoring and Management
  - Volume and Storage Management
  - Others…

Building Large Storage Networks

Storage Networking Management – Tools

- SNMP Support
  - Fibre Alliance MIB
    - Gets, Sets, Traps
  - Framework Integration
    - Tivoli, Veritas, CA, BMC…
- Centralized Management
  - Core to Edge, Local to Remote
  - Multiple Directors or Edge devices
  - Drill-down capability
  - Configuration, monitoring
- Map of entire Storage Network
- Notification & Diagnostics
  - Call Home, SNMP Alerts
- Real-time & History
  - Performance and Events Logs
  - Export capability

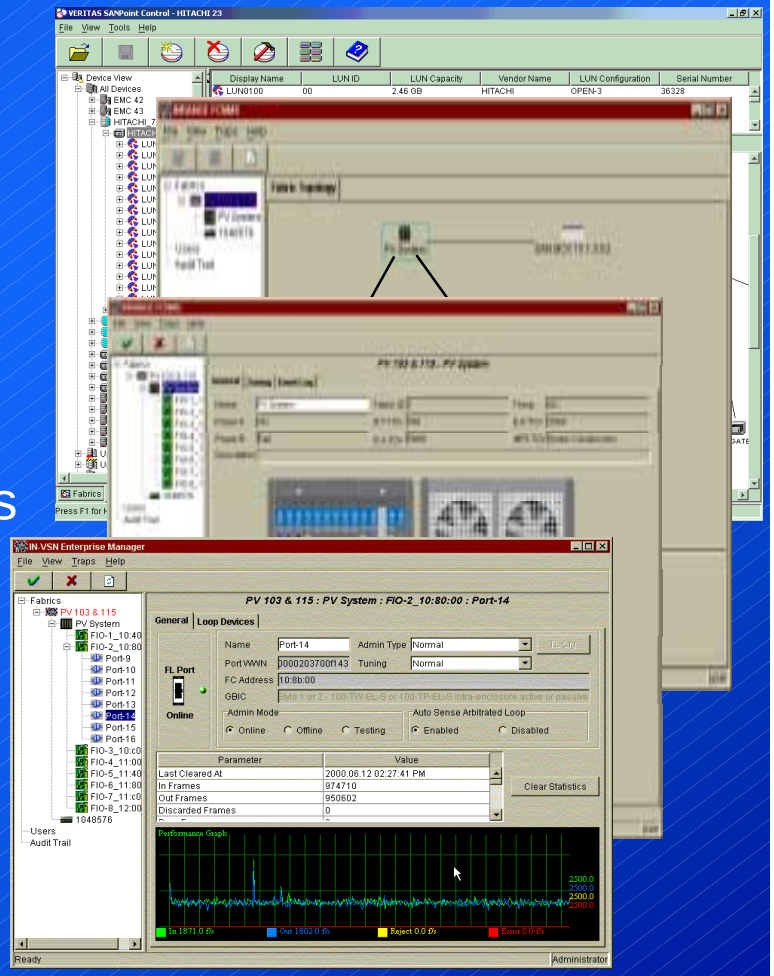

Building Large Storage Networks

Storage Networking Management - Security

- Large Storage Networks like large LANs need security
  - More devices and interconnects introduce risk
  - Risk vs. enablement of scale needs to be managed
- Mesh vs. Large Port Count Core Directors
  - Mesh introduces security complexity (more links and exposure)
    - Additional security measures and software required
    - Forces the need for additional security technology
  - Large port count directors simplify security
    - Address security issues with architecture vs. technology
    - Single device eliminates tampering of ISL's!!!
- Zones are used for security as with regular networking
  - Prevent accidental and intentional access to data
  - Provide security from non-compliant/misbehaving devices
  - Isolate specific customers and data for SLA purposes

WHERE NETWORKS CONVERGE

Building Large Storage Networks

Storage Networking Management - Security

- Multiple levels of Storage Networking security
  - Securing management tools and consoles
    - Preventing un-wanted, or un-authorized access to interfaces
    - Enable management, enable flexibility with safety
  - Securing the ports from host/servers
    - Prevent newly discovered hosts from un-wanted access
    - Isolate various hosts/servers ports or HBAs from storage network
  - Securing the ports from devices
    - Preventing discovered devices from impacting storage network
    - Isolate mis-behaving devices from impacting storage network
  - Securing the ports from other ports/switches
    - Isolate in-compatible devices from each other
    - Isolate hosts/servers or applications from each other
  - Securing hosts and devices
    - To meet SLA or Service Level Agreements
    - Meet other business requirements

Building Large Storage Networks

Storage Networking Management - Security

- Examples and types of zoning include:
  - Storage Networking zoning
    - Host software or persistent binding
      - Maps HBAs to specific storage networking devices/ports
    - Storage or device based zoning
      - Lun or Volume mapping/masking
      - Port or interface mapping/masking
      - Reserve/release techniques
    - Director or switch based zoning
      - Hard Port Zoning
      - Soft or name server zoning
      - WWN zoning
      - IP Broadcast zoning
      - FICON port prohibit
    - Others zoning or security techniques include:
      - Virtualization and volume management tools
      - Tape device and backup library management tools
      - Bridges, gateways, and routers

Building Large Storage Networks

Global or Long Distance Storage Networks

- Dark and dedicated Fiber solutions (80-100km)
  - Long wave GBICs with extended Buffer Credit support
  - WDM/DWDM increases Fiber utilization and distance
  - Well suited for campus and metropolitan applications
  - WDM/DWDM can provide Fiber/Bandwidth on demand
- Long distance or where no dark fiber exists
  - Networking interfaces including ATM/OCx & IP
  - Provides distance out to several hundred Km
  - Enables SAN to SAN connectivity
  - Enables channel extension
  - Latency issues need to be considered

Building Large Storage Networks

Global or Long Distance Storage Networks

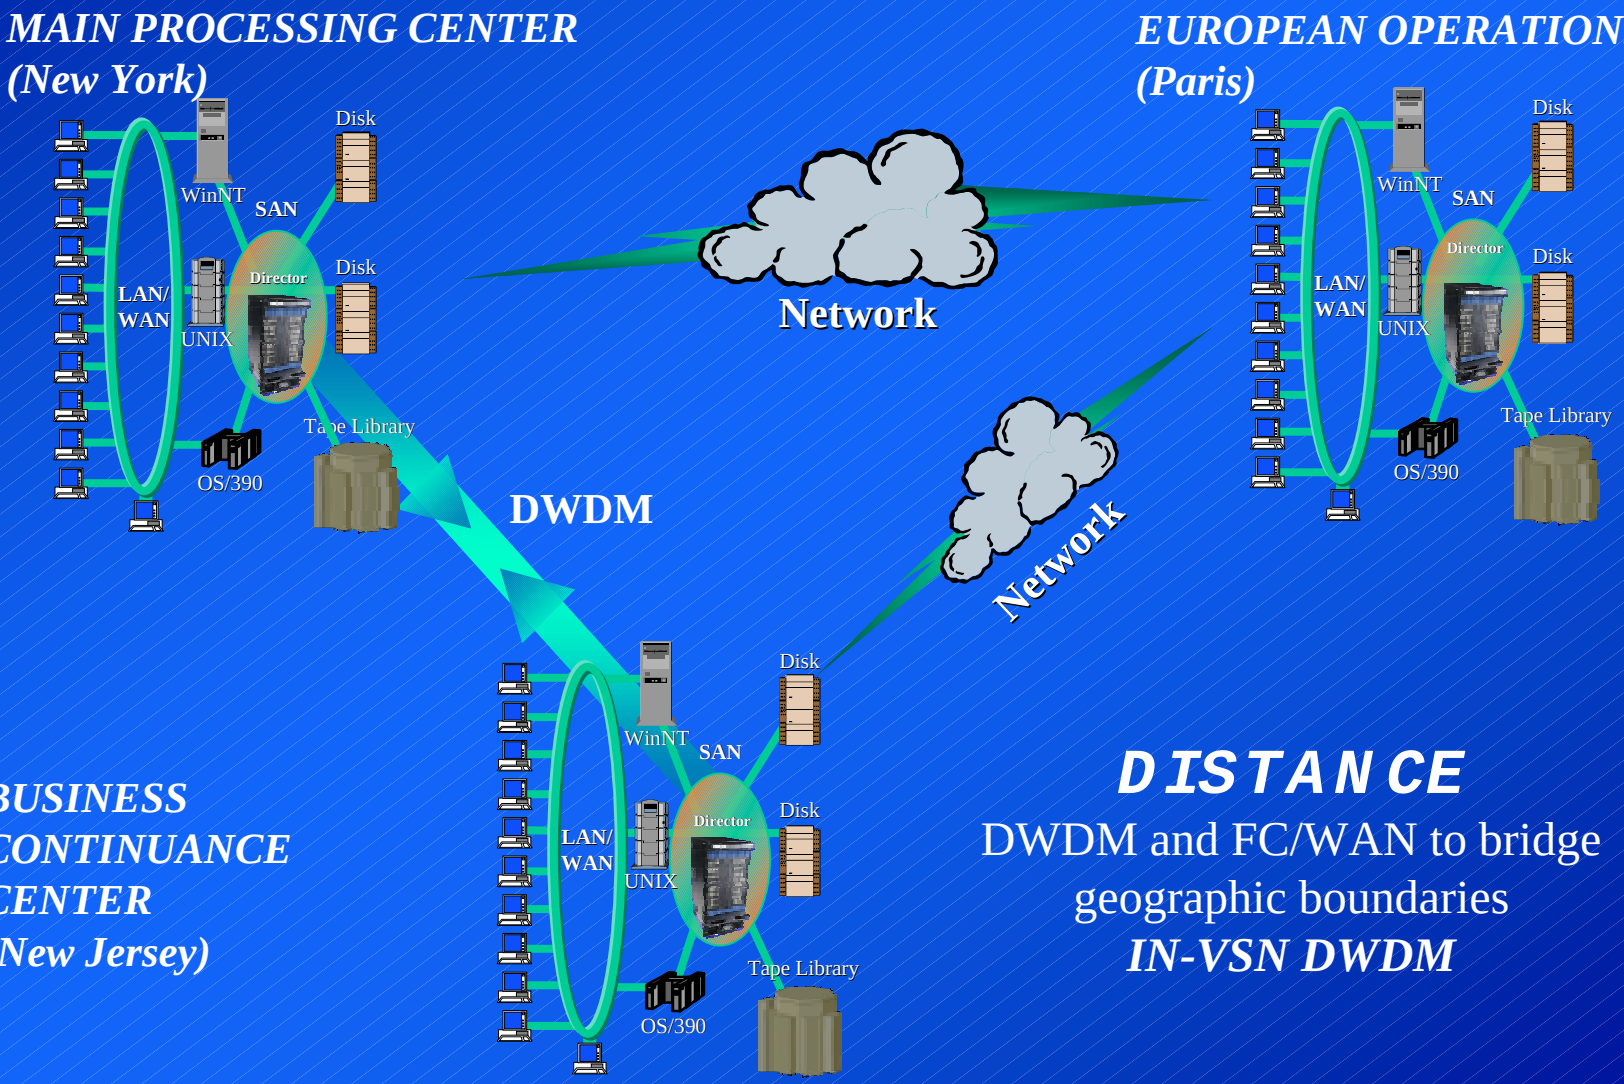

DISTANCE

DWDM and FC/WAN to bridge geographic boundaries

IN-VSN DWDM

WHERE NETWORKS CONVERGE

Storage Networking Solutions

INRANGE™

Building Large Storage Networks

Solution – SAN, NAS, Backup, Server Farms

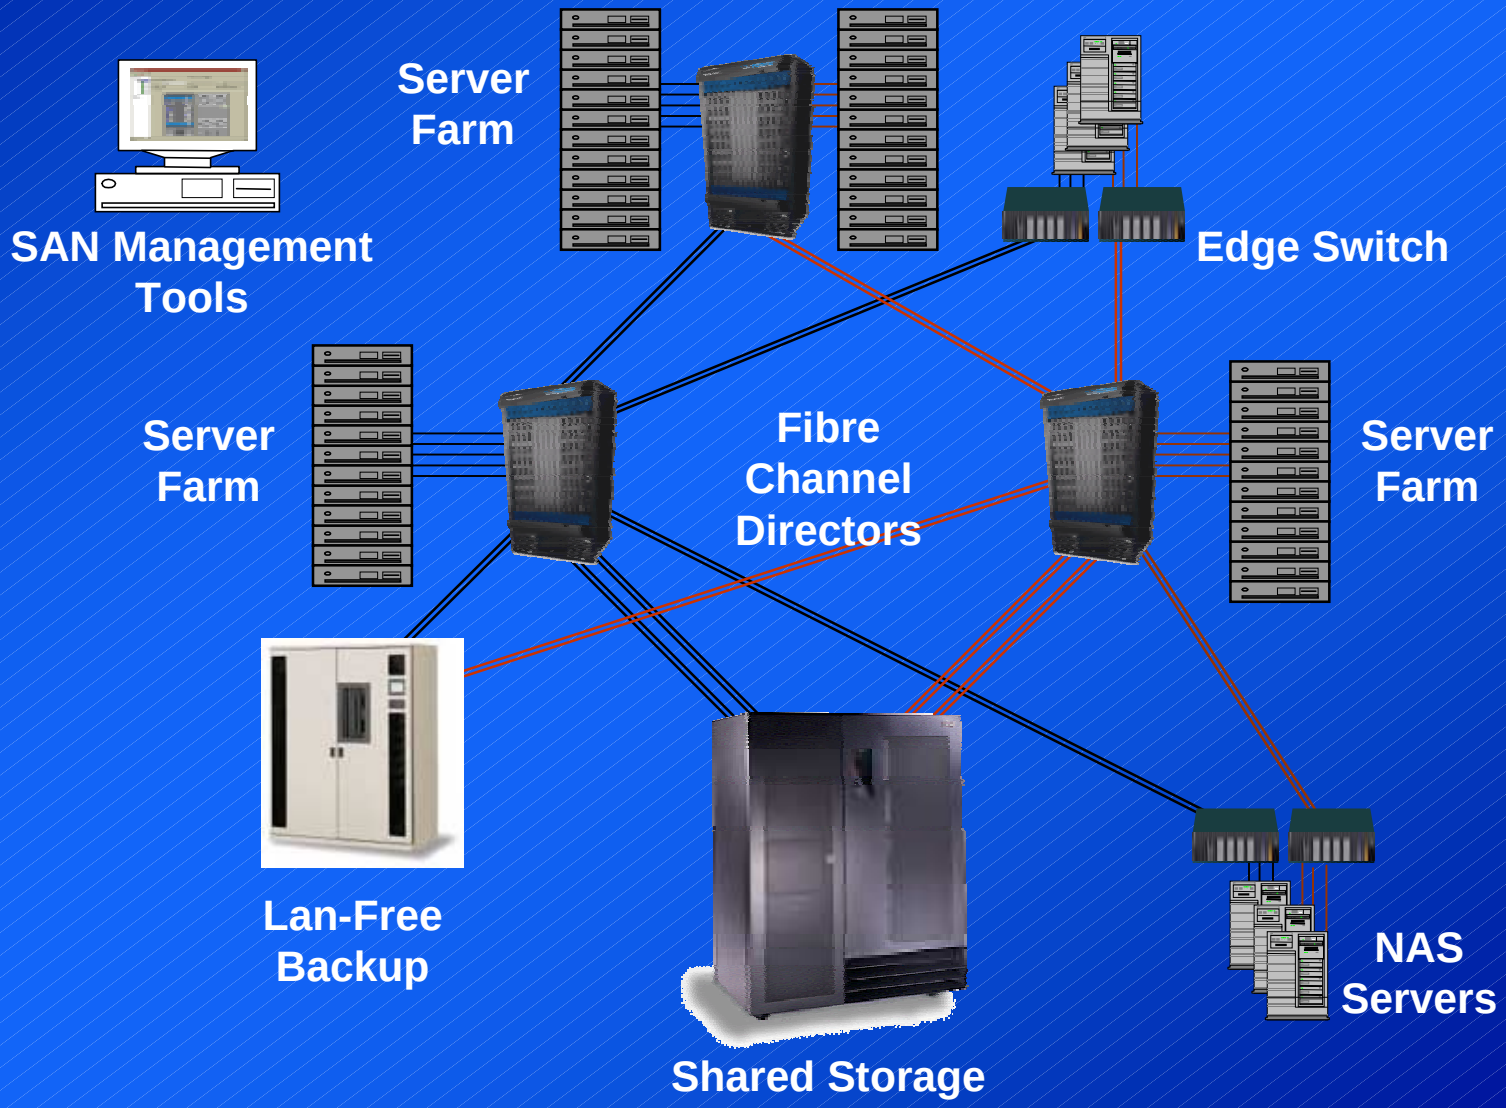

Building Large Storage Networks

Solution – NAS for file share, SAN for Database

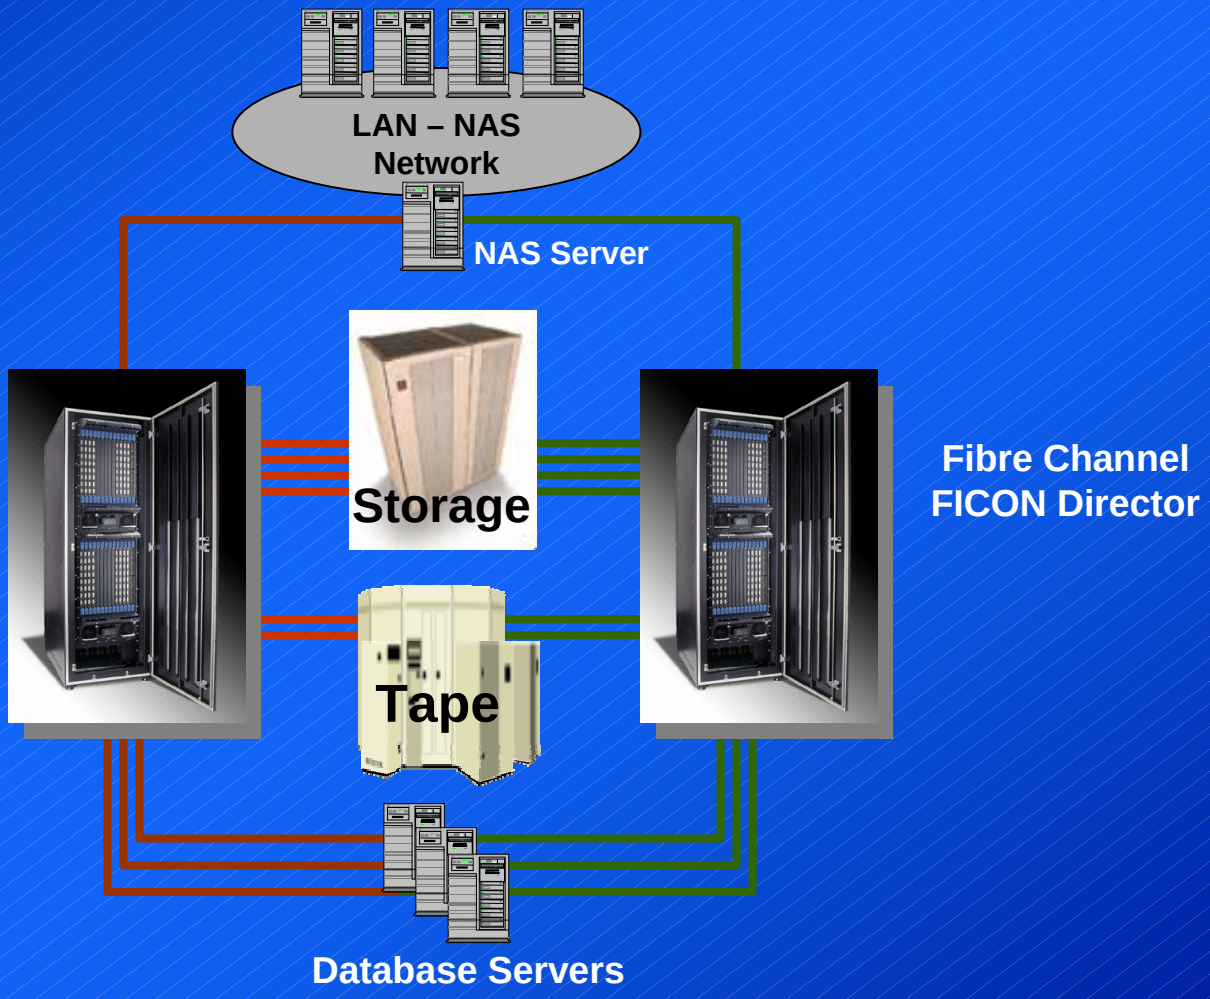

Building Large Storage Networks

Solution – Server Cluster, Backup, Unix & NT

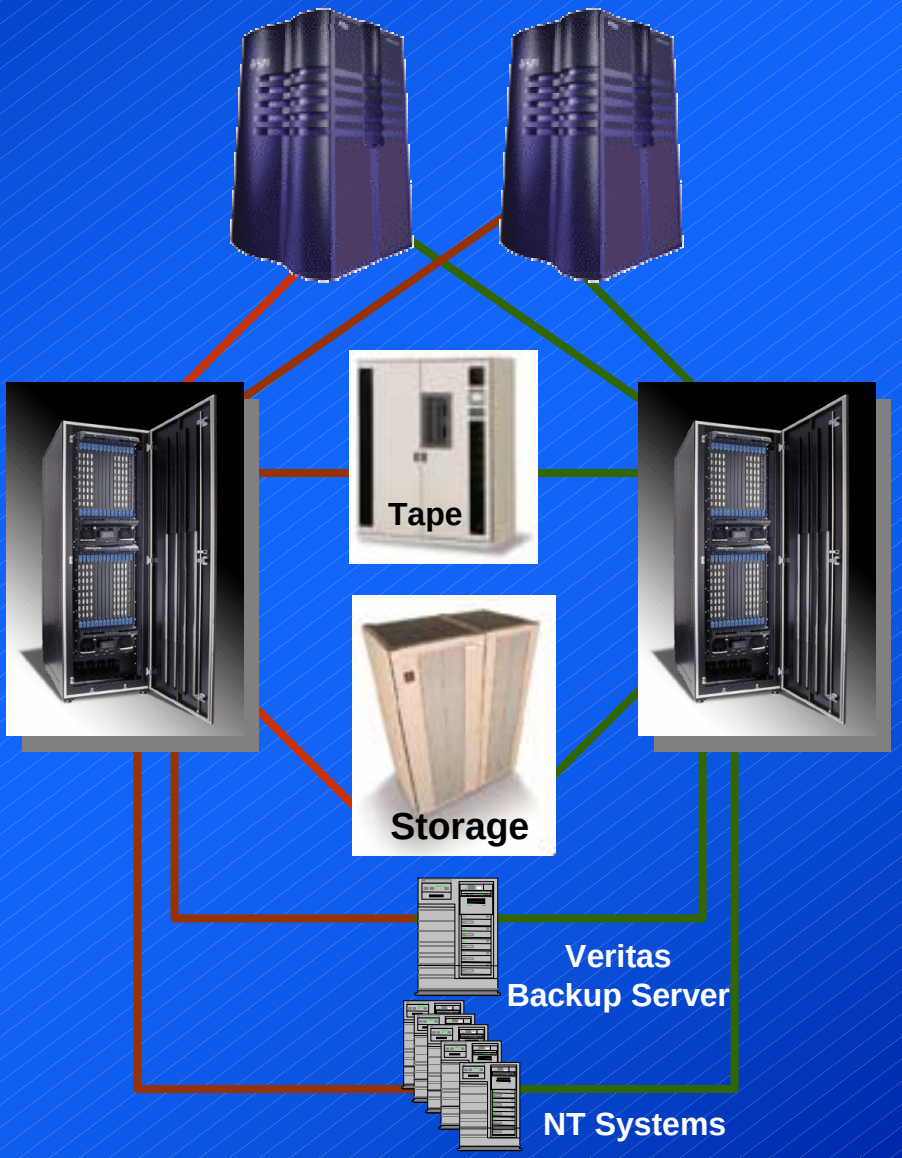

Building Large Storage Networks

Solution – Metro Area Storage Network

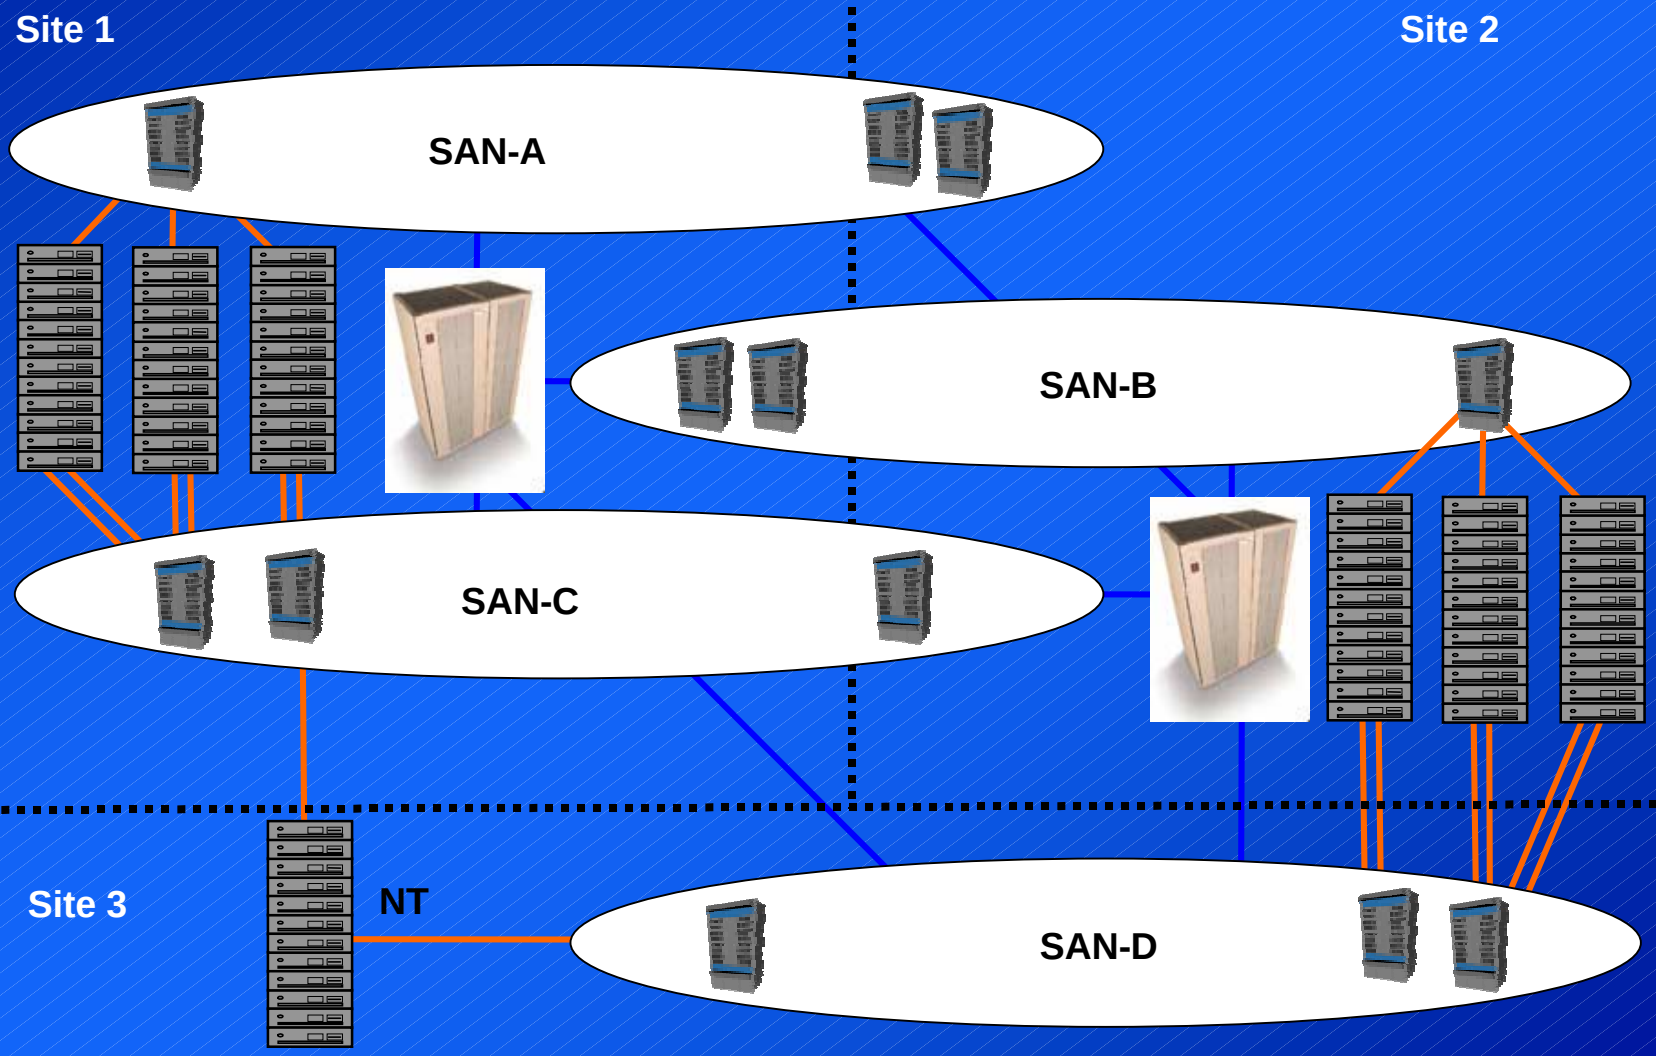

Building Large Storage Networks

Solution – Mixed OS390 and Open Systems

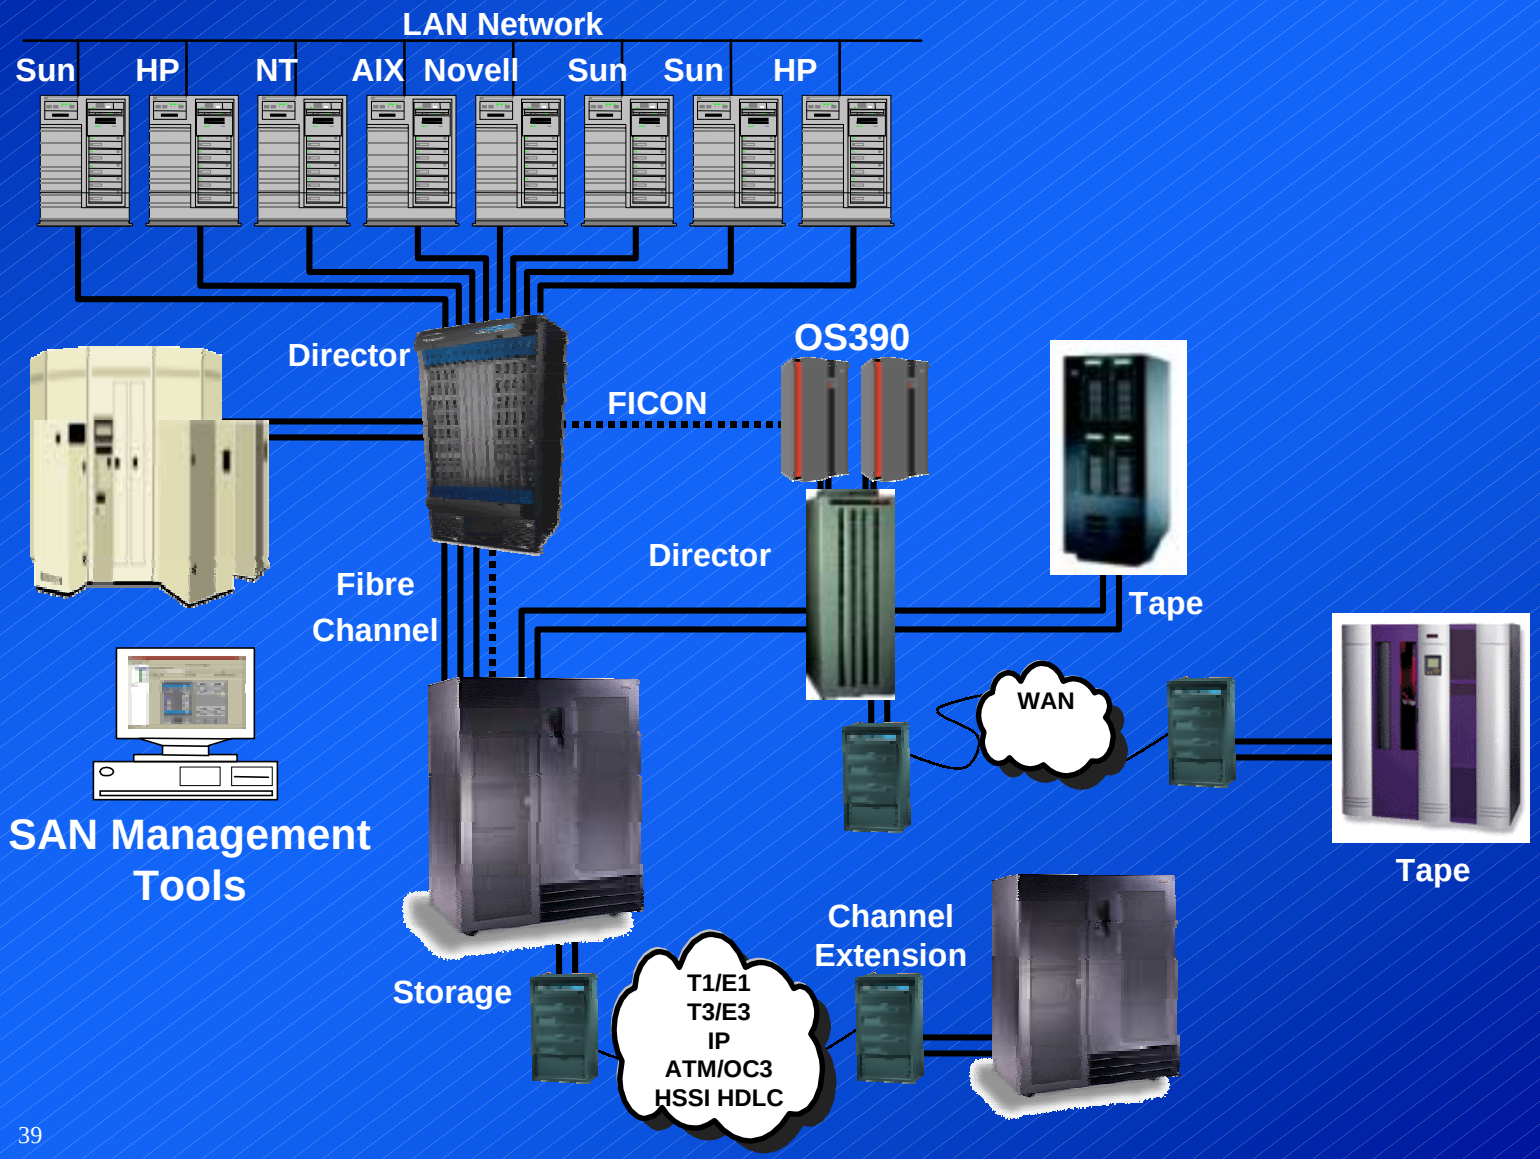

Building Large Storage Networks

Solution – OS390 FICON and Wide Area

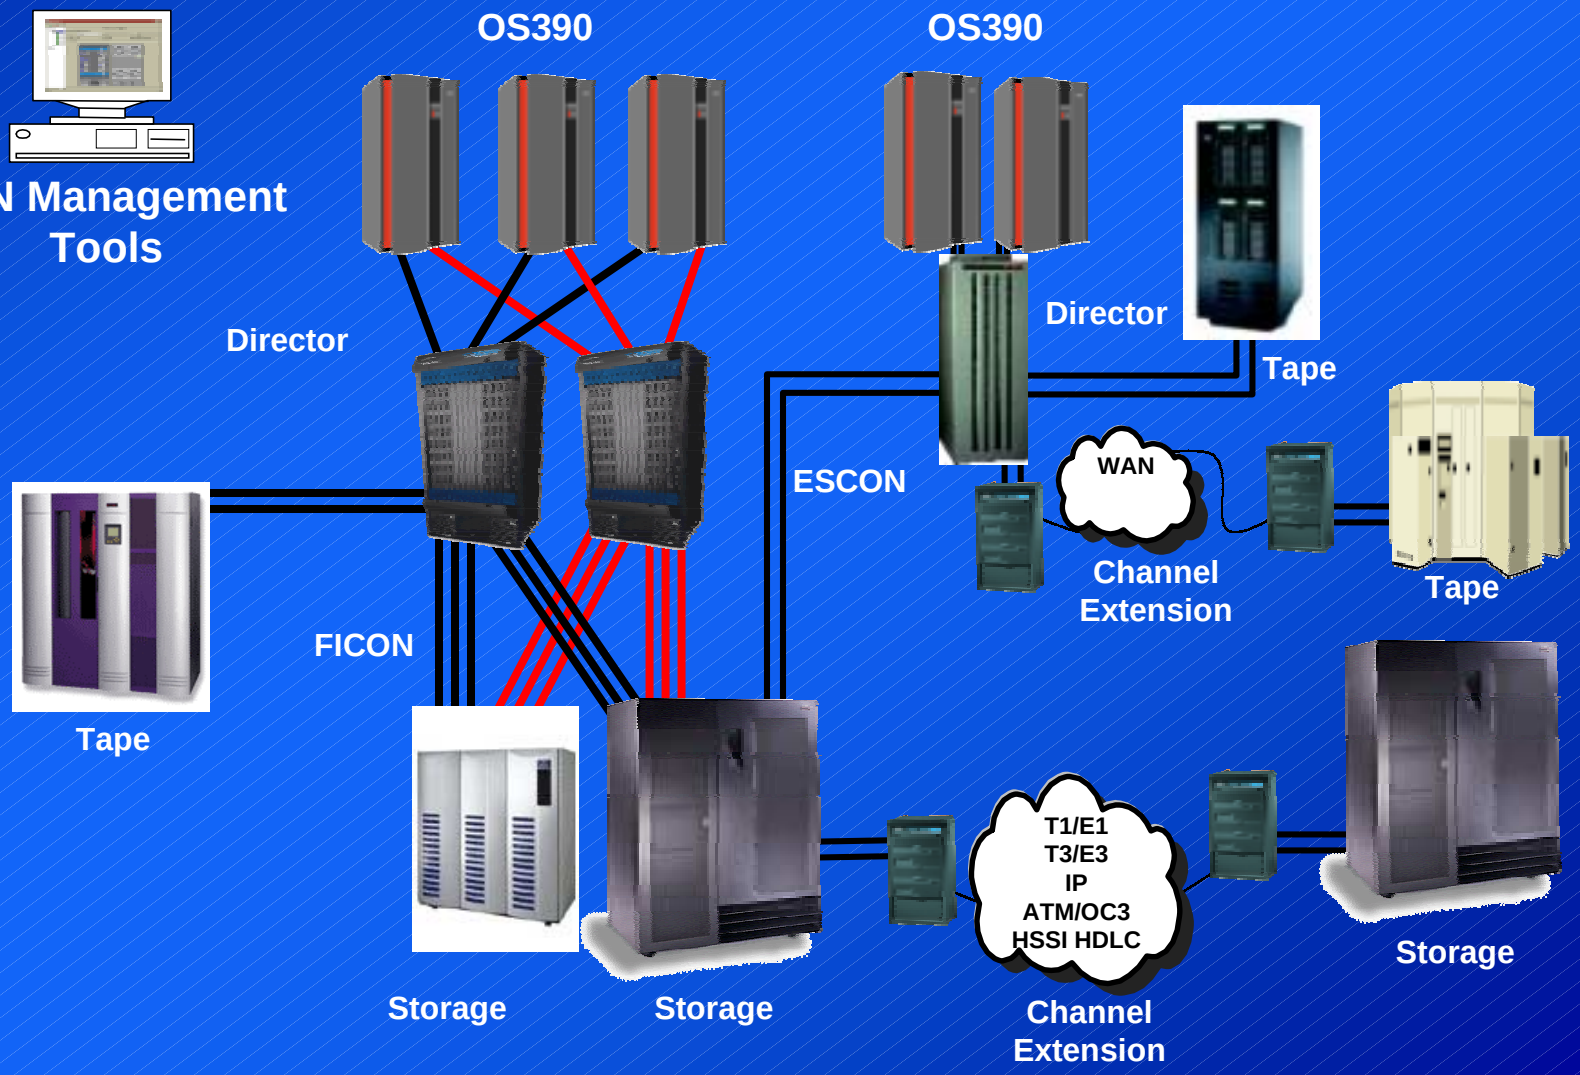

Building Large Storage Networks

Solution – Large Scale Storage Consolidation

- Large Any to Any
  - Non-Blocking
- Multi-Path
  - Load Balance
  - Auto Fail-Over
- This Example
  - 40 Server Ports
  - 32 Storage Ports
  - 8 Tape Ports
  - 48 (112 w/128's) Unallocated Ports

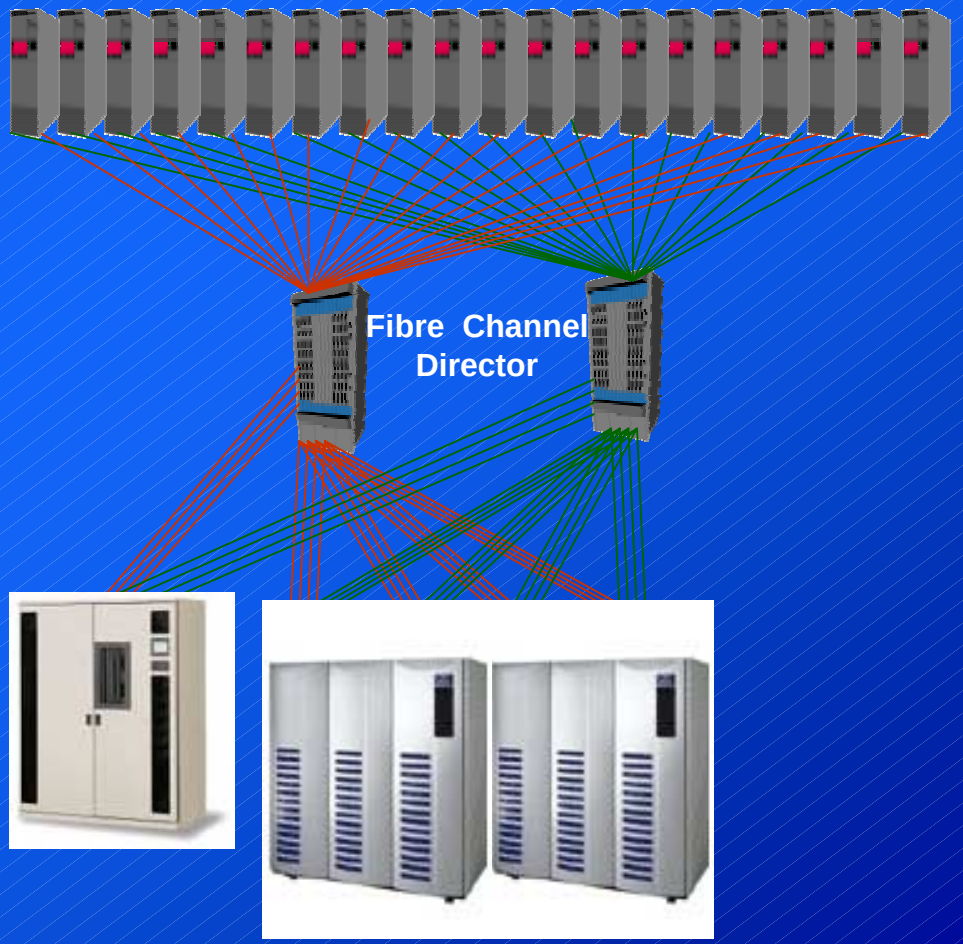

Building Large Storage Networks

Solution – Large Enterprise Storage Network

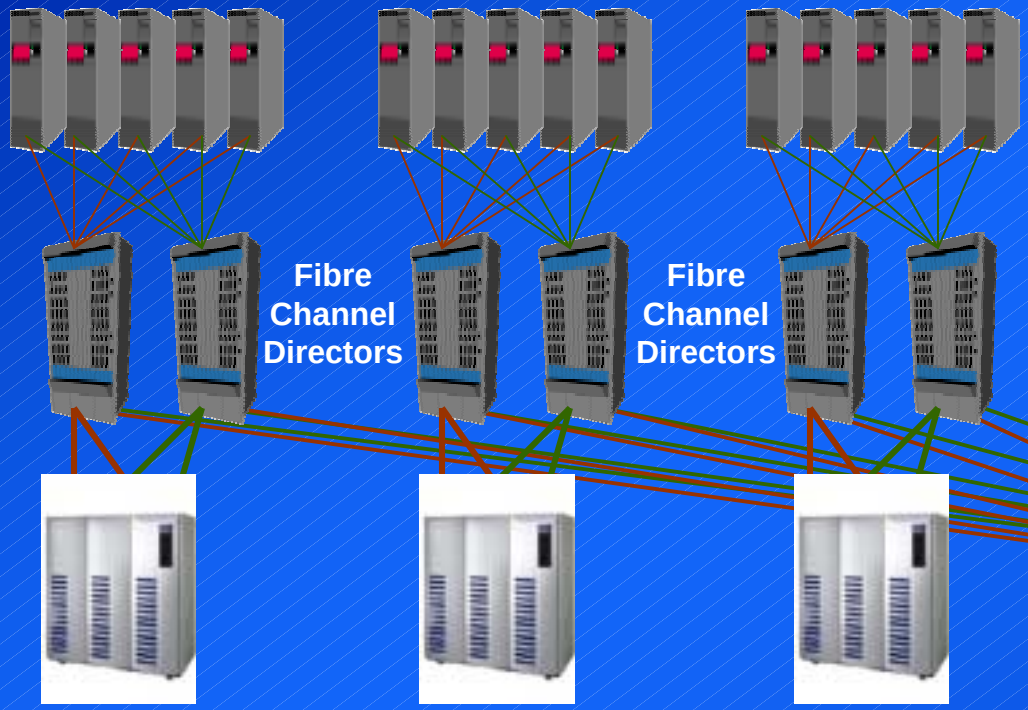

ISL to “Core” Director
- Region to Region Traffic
- Tape Silo on Central Director

Tape Silo on Core Switch

Each Region has Server & Storage
- Organize by Function
- True Any to Any

Building Large Storage Networks

Solution – High Performance 1,100 Usable Ports

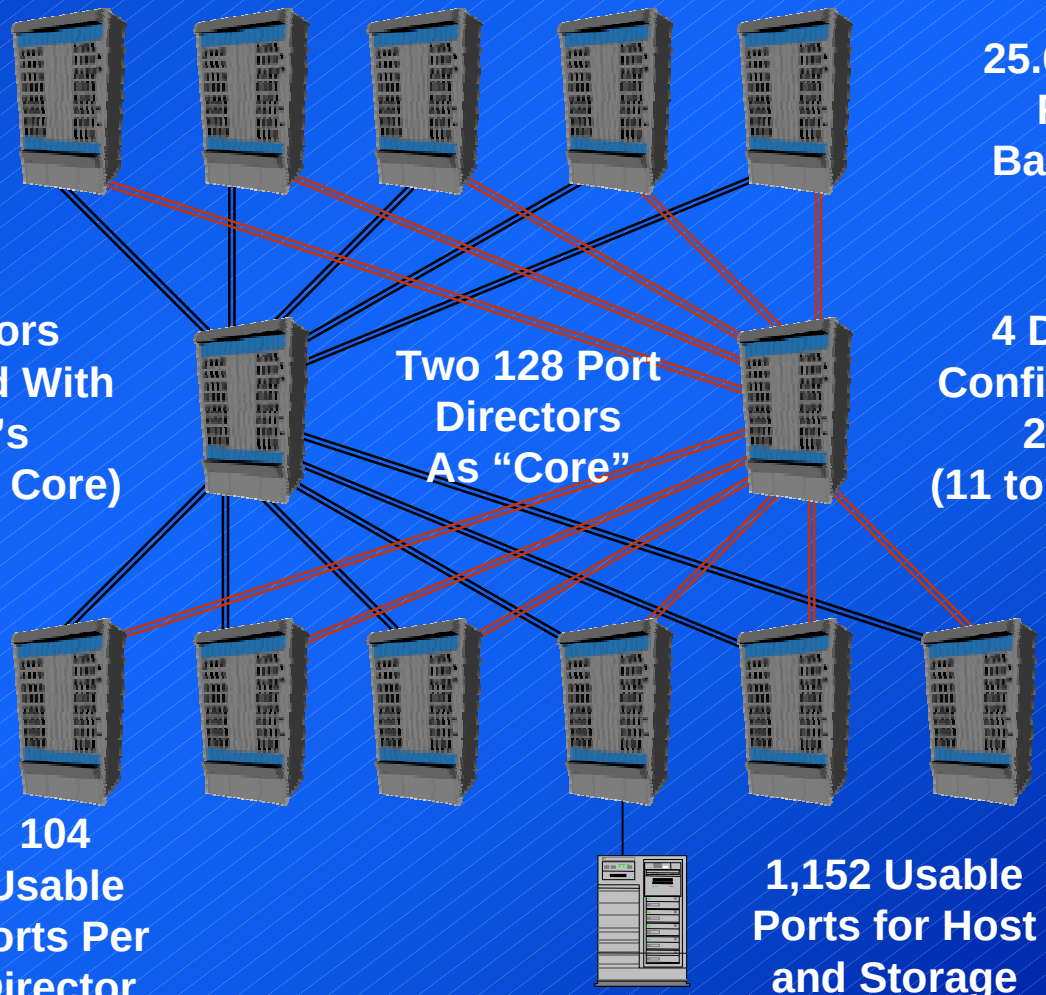

Building Large Storage Networks

Solution – Clustered Local HA Storage Network

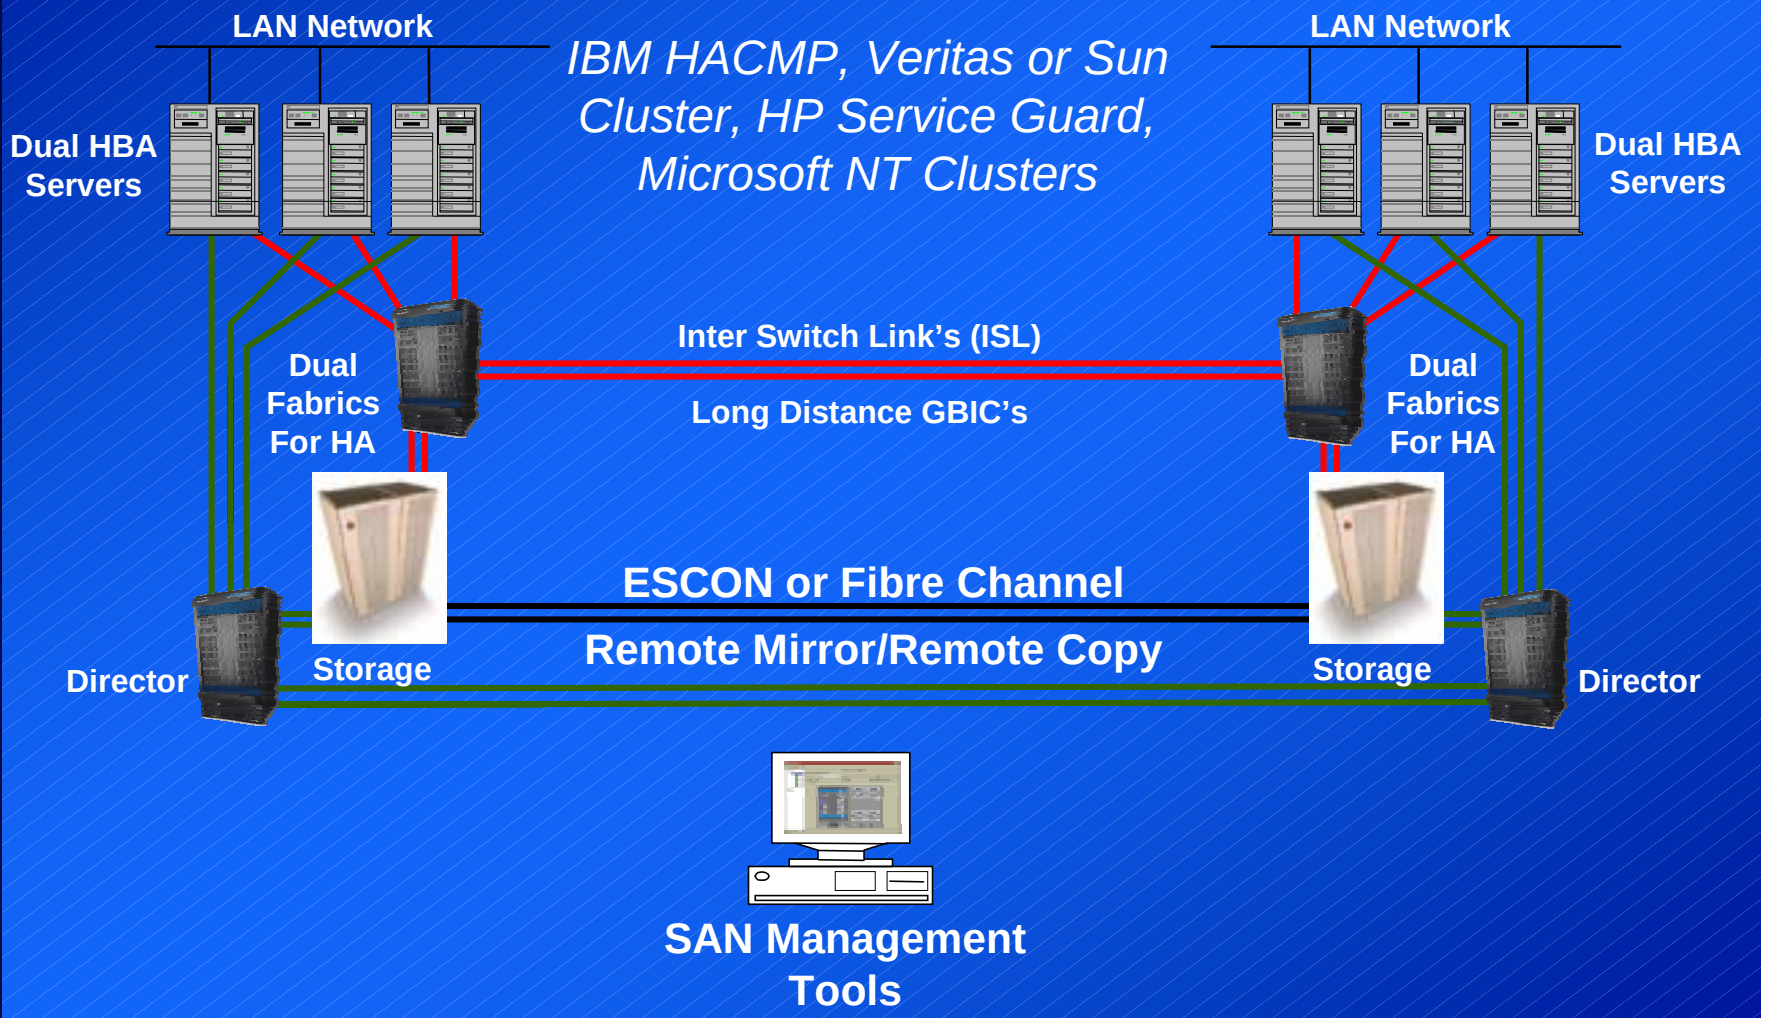

Building Large Storage Networks

Solution – Clustered Metro HA Storage Network

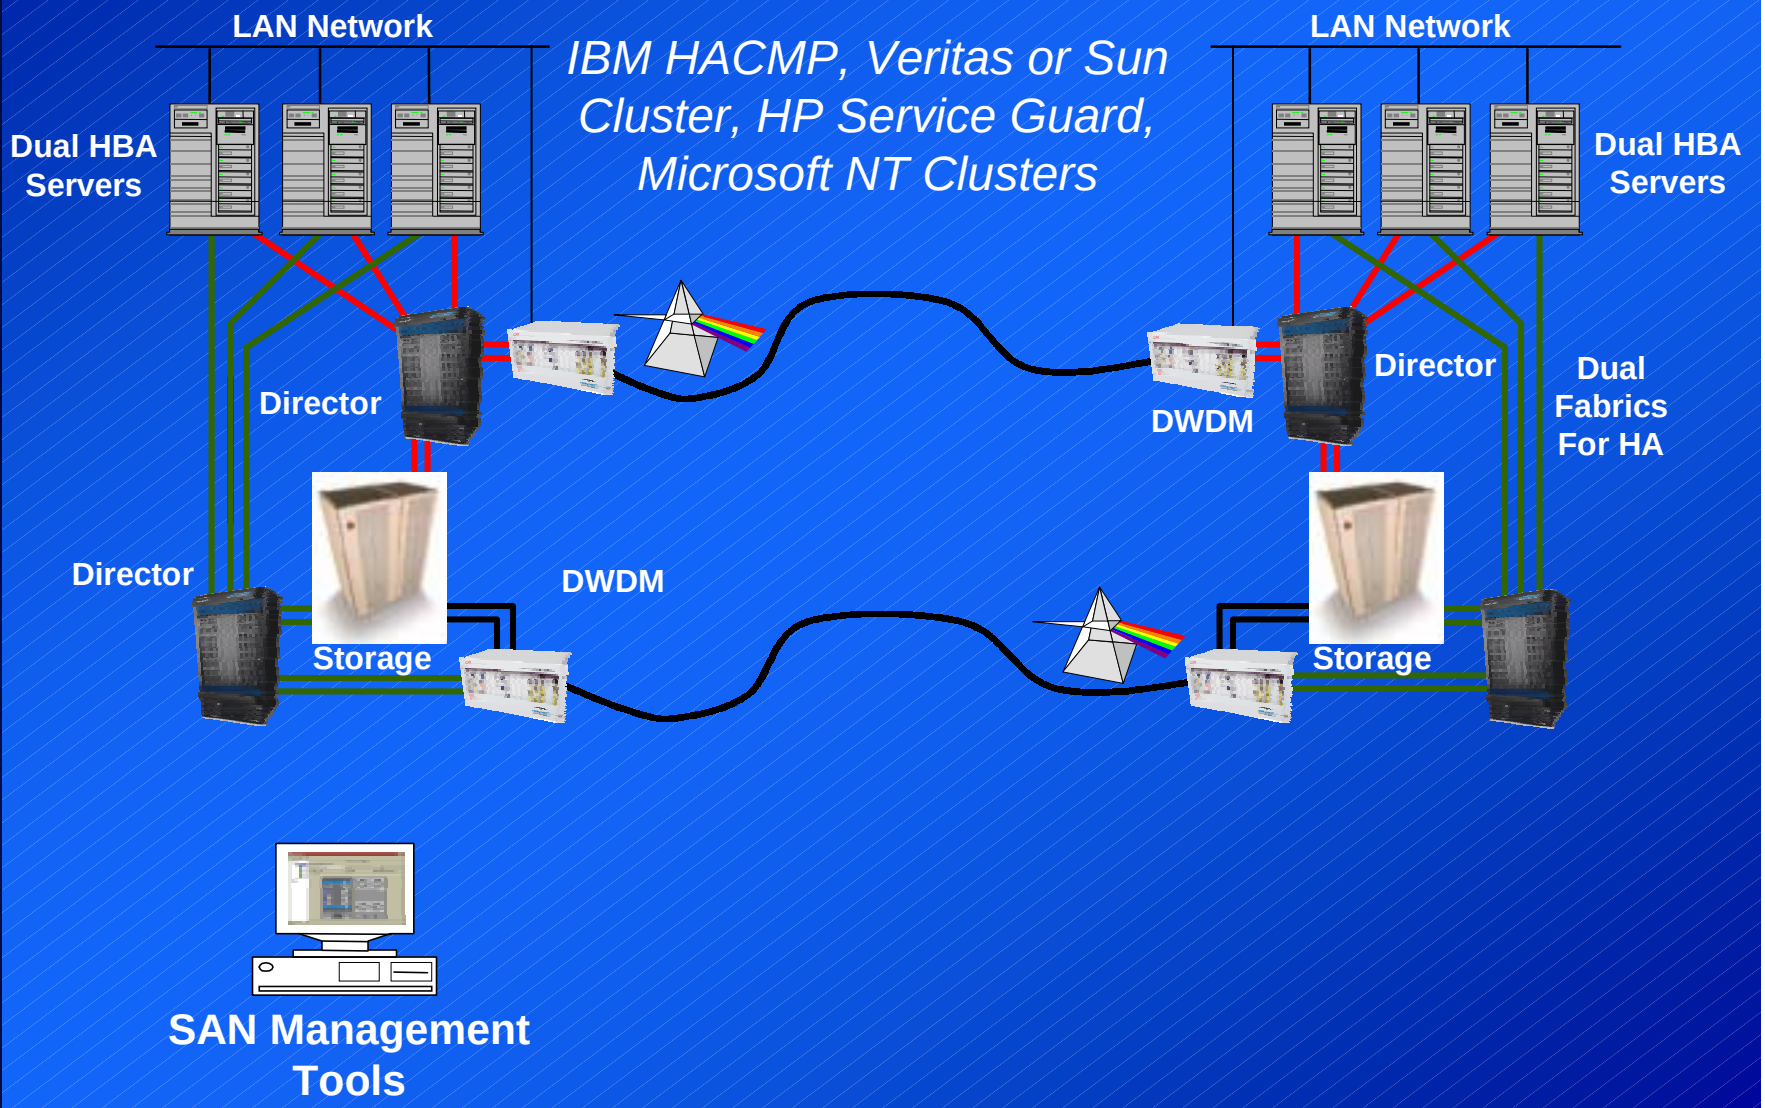

Building Large Storage Networks

Solution – Remote Tape and Storage

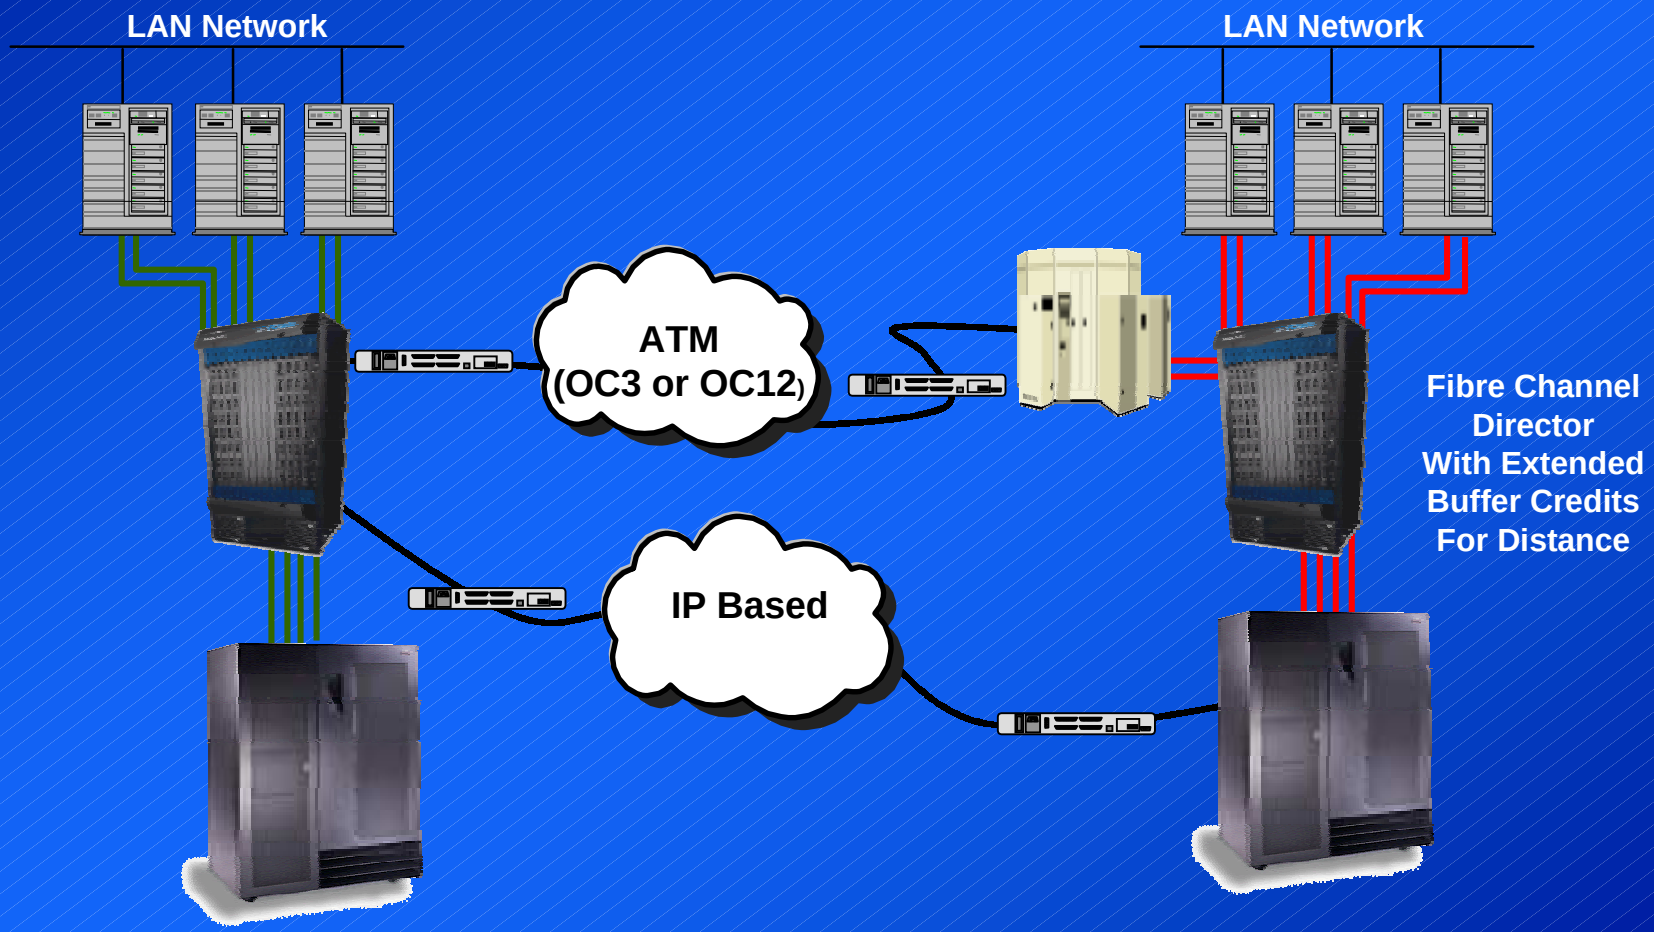

Building Large Storage Networks

Summary (Almost done, one more slide to go…)

- **Large Storage Networks are being built**
- **Storage Networks do not have to be complex**
- **Heterogeneous Storage Networks exists**
- **Plan for scalability and ease of management**
- **Leverage existing storage knowledge**
- **Wide area Storage Networks can be built**
- **Storage Networks will grow larger…**

WHERE NETWORKS CONVERGE

Building Large Storage Networks

Closing Comments

- **Interoperability**
  - Support mixed architectures, investment protection
- **Performance**
  - Non Blocking, Predictive Latency, scales with growth
- **Availability**
  - Redundancy, Concurrent code load, Hot Swap
- **Scalability**
  - Ability to grow on-line in a stable manor including 2Gb support
- **Manageability**
  - Notification, diagnostics, configuration, simplification
- **Flexibility**
  - Local, Campus, Wide Area SANs with fabric and loop support

WHERE NETWORKS CONVERGE

WHERE NETWORKS CONVERGE

IN-VSN Family

for Virtual Storage Networking

Networks for storage
- across the enterprise
- across architectures
- across the globe

INRANGE™