BROCADE
The intelligent platform for networking storage

Decus Symposium Bonn 2003

Brocade SAN Design

Leading the Next Generation of Storage Area Networks

2003

Joachim Meurer Brocade Senior System Engineer

jmeurer@Brocade.com
+49-1713357337

AGENDA

- SAN Building Block
- Backbone with HW Trunking
- Core Edge Design
- Zones, Security, I-Switch

BROCADE
The intelligent platform for networking storage

An Intelligent Platform for Networking Storage

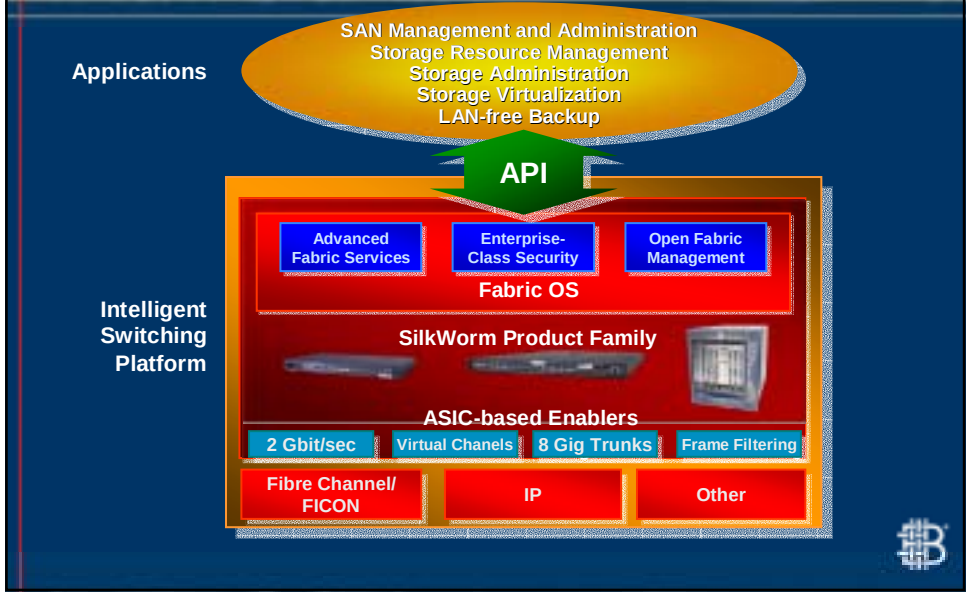

Brocade SilkWorm Fabric Application Platform

Industry's first open, intelligent platform for fabric-based applications

SilkWorm Fabric Application Platform

- Selected by HP, EMC and others
- Fully interoperable and manageable with existing SilkWorm family of products

SilkWorm 12000
Core Fabric Switch
64/128-port configurations

- Enterprise
- Available from all major OEMs

SilkWorm 3900
32-port Fabric Switch

- Mid-range/enterprise
- Available from all major OEMs

SilkWorm 3800
16-port Fabric Switch

- Midrange/enterprise
- Available from all major OEMs

SilkWorm 3200
8-port Fabric Switch

- Entry-level
- Available from all major OEMs

SAN Building Blocks

Know Your Application I/O Profile

Application	Bandwidth Utilization	Read/Write Max	Typical Access	Typical I/O Size
OLTP, e-mail, UFS e-commerce, CIFS	Light	80% read 20% write	Random	8 KB
OLTP (raw)	Light	80% read 20% write	Random	2 KB to 4 KB
Decision support, HPC, seismic, imaging	Medium to Heavy	90% read 10% write (except during "builds")	Sequential	16 KB to 128 K
Video Server	Heavy	98% read 2% write	Sequential	> 64 KB
SAN applications: serverless backup, snapshots, third-party copy	Medium to Heavy	Variable	Sequential	> 64 KB

What does it take to do a SAN Design

- What additional HW/SW plans the customer?
  - New additional Storage Devices, Servers
    - Redundancy Levels
    - Capacity requirements
    - Clustering
  - New Applications
    - LAN free, Server less or remote Backup
    - HA, Mirroring, Remote Copy
    - Storage/SAN Management

FC-AL protocol tape drives ?????????????

Direct-attached to FL-port capable switches

Brocade supports Private & Public Loop !!!!!

What is included in a good SAN Design

- Naming Concept
  - Zones, Switches, Server Nodes, Storage Nodes
- Zoning Concept
  - Hard Zoning or Soft Zoning?
- Scalability Concept
  - capacity, #port, performance, redundancy level…
- Service / Disaster Concept
  - FW-upgrades, Disaster , switch exchange….
- Fabric Design

Cable Management Recommended Device Placement

Distribute High Port Count Devices, Such as Arrays or Tape Libraries Across Multiple Blades

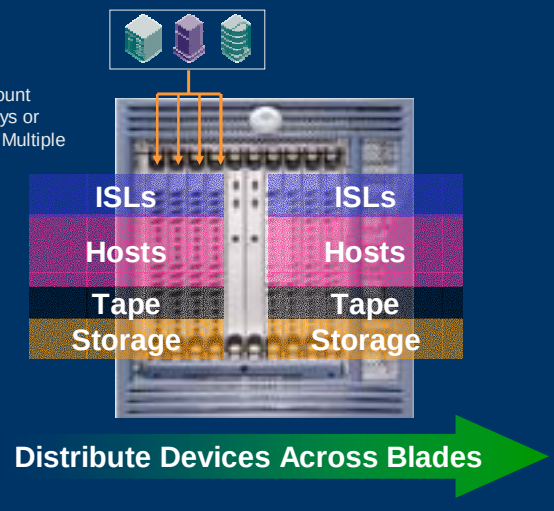

Switch Latency is NOT an issue

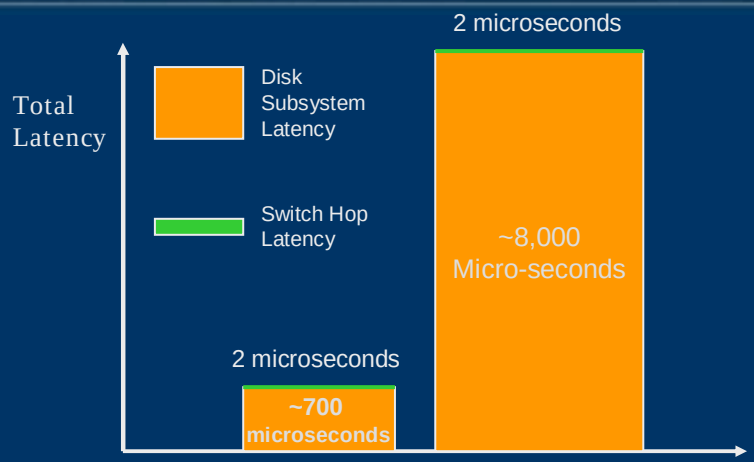

Do you need Any to Any ?
Segment your Problems !

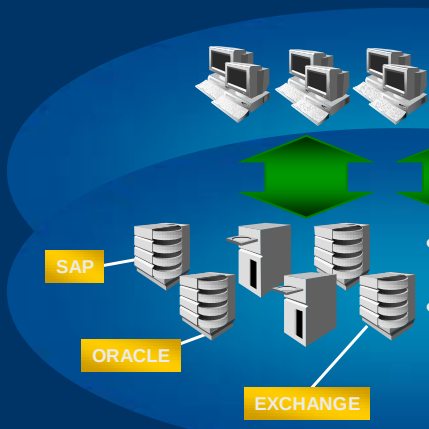

- Data centers increasingly are “facing” the customer
- Business continuity equals continuous information access

Heterogeneous SAN’s - use Zoning ! Hard or Software

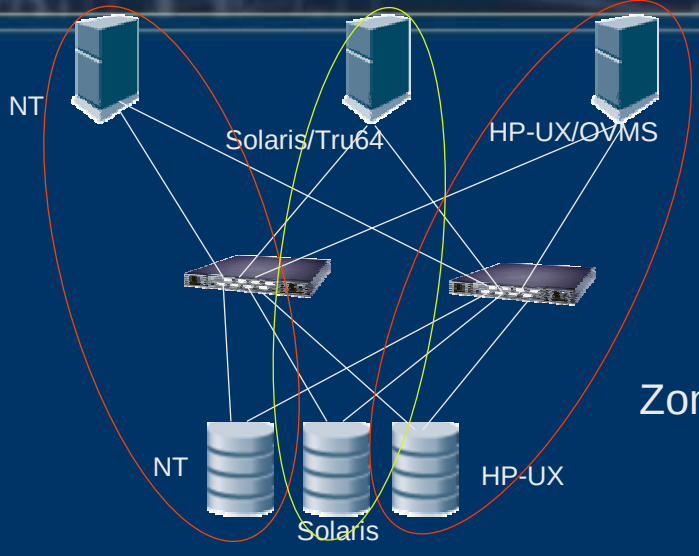

Zoning !

Heterogeneous SAN's – FCAL for Brocade Silkworm Switches No Problem

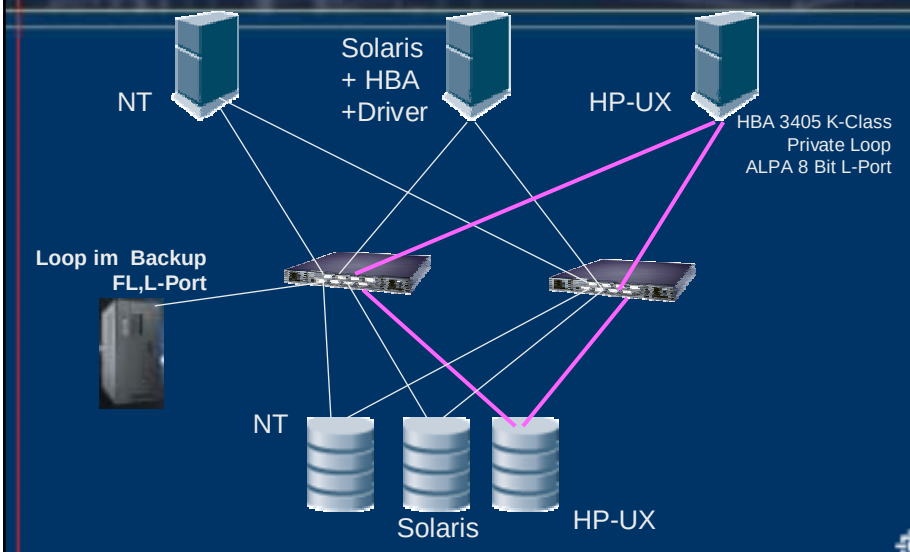

Fabric aware HBA's for Solaris QuickLoop for HP-UX-K Class

Advanced Fabric Services: Security and Access Control

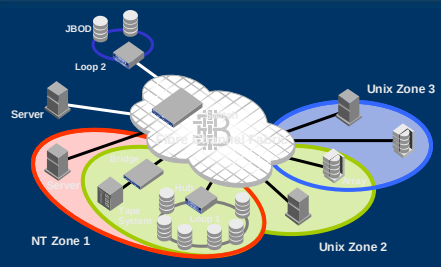

Hardware Enforced *Zoning*

- Isolate server/storage groups by application, OS, organization, etc.
- Protect accidental data corruption in mixed UNIX and Windows Fabrics
- ASIC-enforced (uses Frame Filtering) access control by port or WWN

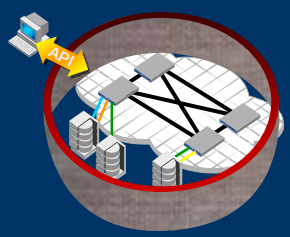

Secure Fabric OS

- Secure control of management, server, and storage access to SAN
- Enables auditable security based on strong PKI authentication
- Only SAN Security product on the market

Secure Fabric OS, Securing The SAN infrastructure

ACL,SSH etc

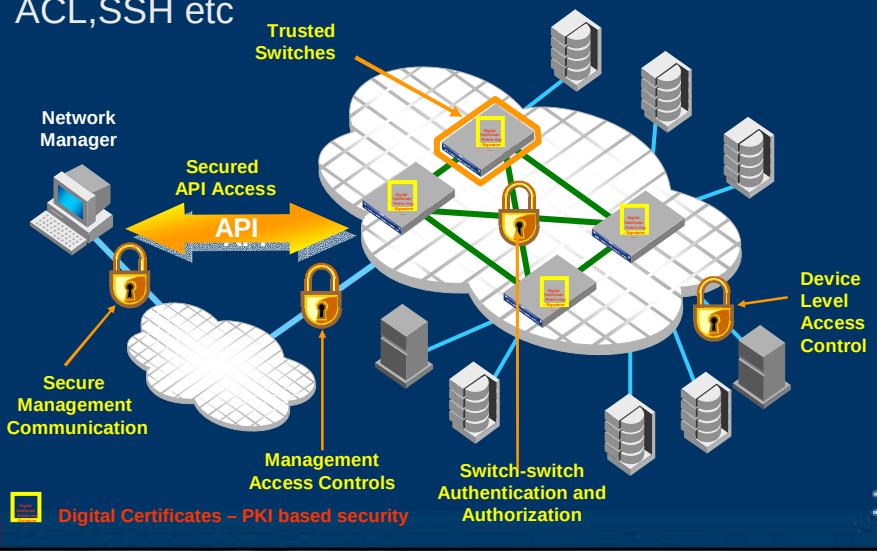

Highest Levels of Availability OVMS,Cluster, Dual Path Raid , DT etc

- True High Availability (HA) means no application downtime, planned or unplanned
- True HA cannot be achieved without ensuring uptime against software failure, human error, and physical locality
- True HA requires a dual-fabric redundant network
  - Centralized management
  - Muli-pathing and failover connection
  - Good operational practices

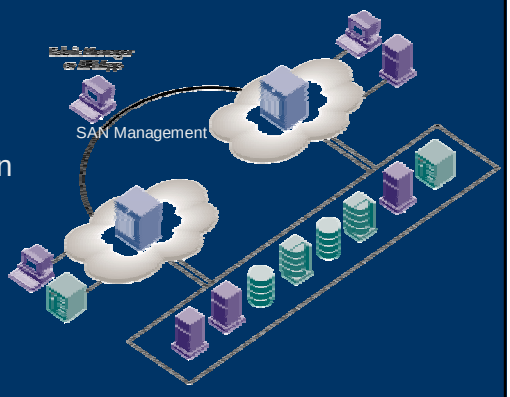

99.999% is 5 minutes of downtime
24hrs/day * 365 days/year * 60mins/hr * .99999 = ~5.25 minutes

Silkworm 12000 Core Fabric Switch

Flexible, Modular Architecture
- Scalable 64/128 port design
- 2Gbit/sec ports; auto-sensing
- 3rd Generation Brocade ASIC
- High port density – 14U enclosure

- 99.999% availability
  - Redundant, hot-swap elements
  - Non-disruptive software updates
  - Two redundant 64 port switches
- Intelligent Fabric Services
  - Existing Fabric OS services
  - 8Gb ISL Trunking
  - Frame Filtering
  - Global Performance Analysis

"Director" **Plus**

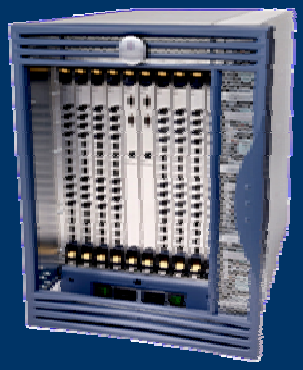

Fully Networkable

High Availibility ✘
Plus, Ultra…This is not a HA SAN Design !!!

A single Director or a single Switch can be a 'single point of failure' :

Hardware (Mid/Backplane), Configuration error, Catastrophes …

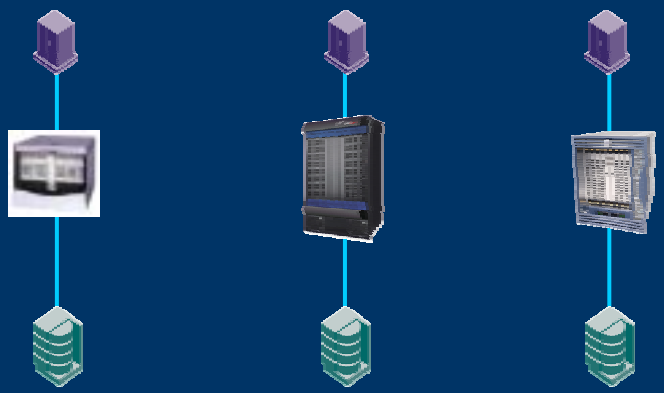

High Availibility ✓

This is a HA SAN Design !!!

High Availability through dual fabrics and switches

Silkworm 12000 is the Perfect Core Switch

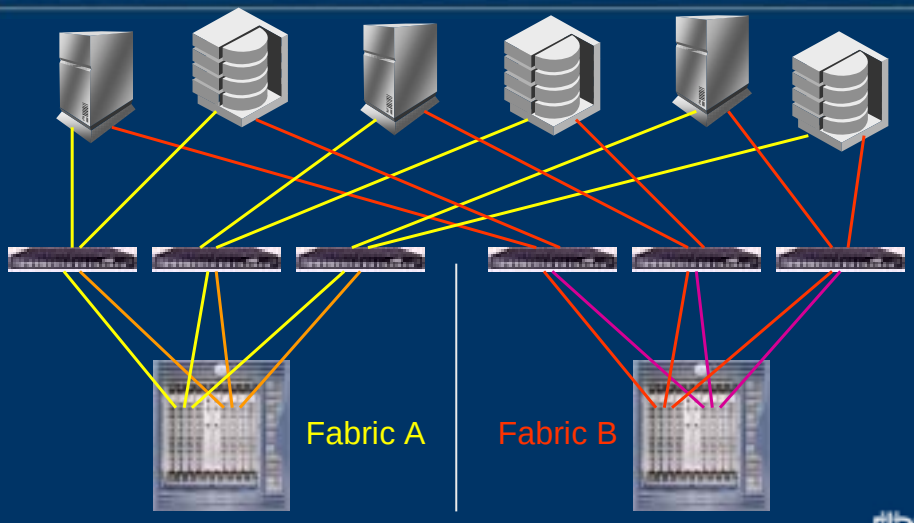

3 Tier Core to Edge Design

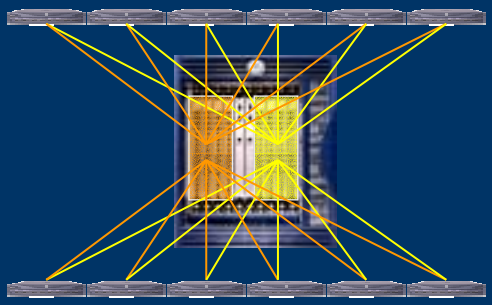

168 Edge Ports + 104 Core Ports = 272 device ports

Core to Edge: Scalability

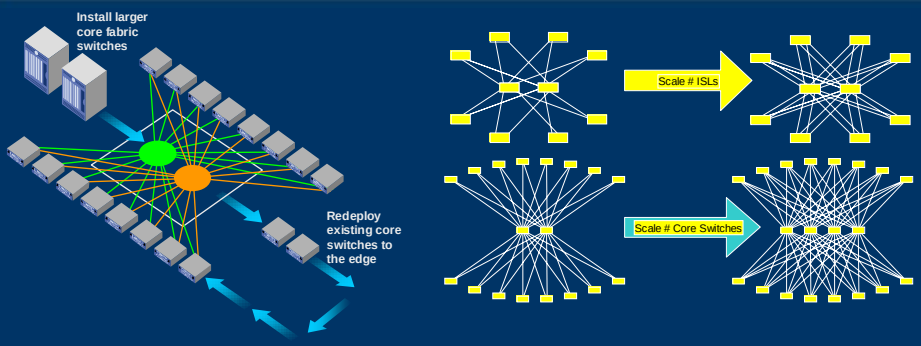

- Scale fabric size by adding switches
- Scale fabric size by replacing existing core with a larger core
- Scale performance by adding ISLs or additional core switches
- Pay as you grow

Advanced Fabric Services
Inter-Switch Link Trunking

Without Trunking (Load Sharing)

- Routes are assigned to ISLs in a round-robin fashion
- Traffic for each route stays on its link
- Some routes can experience congestion while others are underutilized

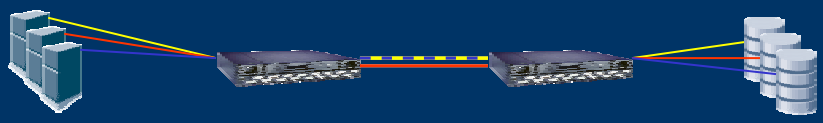

With Trunking (Load Balancing)

- Optimum Bandwidth use
- Double bandwidth with existing lines compared to 1GB Switches
- Managed as a single logical ISL

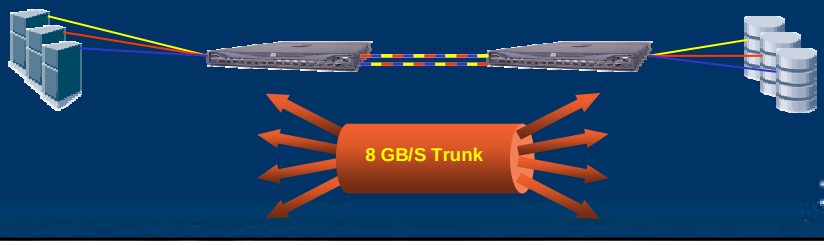

Trunking Flexibility and Reliability

Note that Trunks can be split
Unused ports in one Quad can be used for another Trunk Group

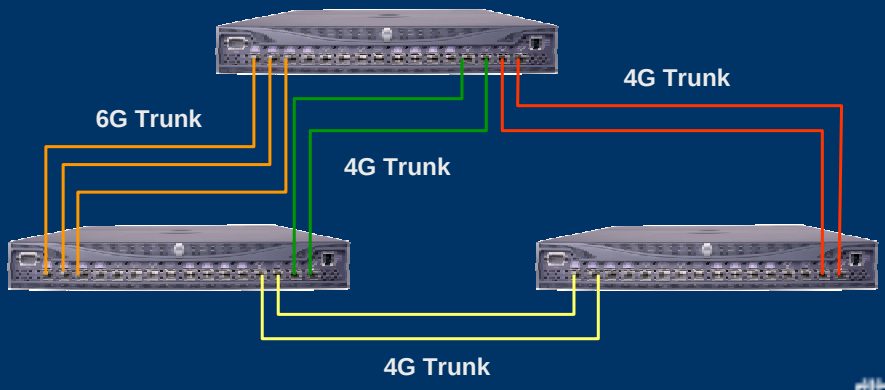

Truth: Trunking Is Even More Important With Larger Switches

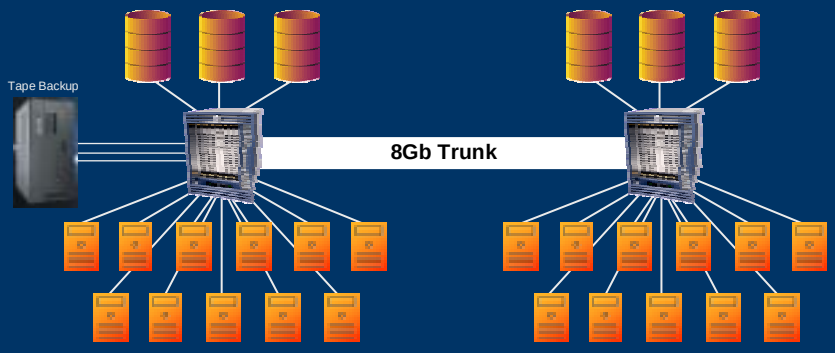

- Building a high performance storage area network
- With HW-Trunking

Business Considerations For Connecting SANs Over Distance

- **Improve Resiliency and Protection of Data**
  - Implement Business Continuity and Disaster Recovery solutions by connecting geographically disparate redundant data/storage centers
- **Improve Efficiency**
  - Centralize data to support storage intensive applications like data warehousing, data mining, email, etc.
- **Improve Profitability**
  - Consolidate sites to achieve operational savings
  - Increase asset utilization

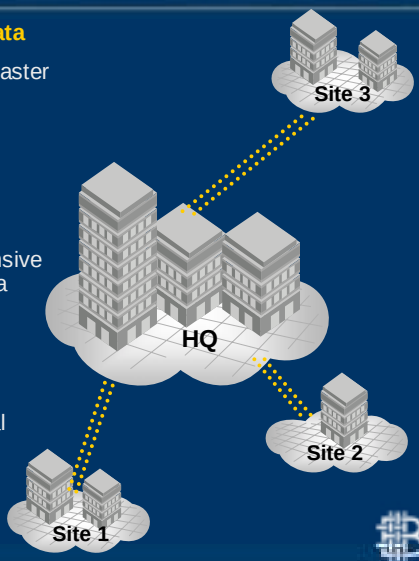

SANs over Campus and Metro Areas

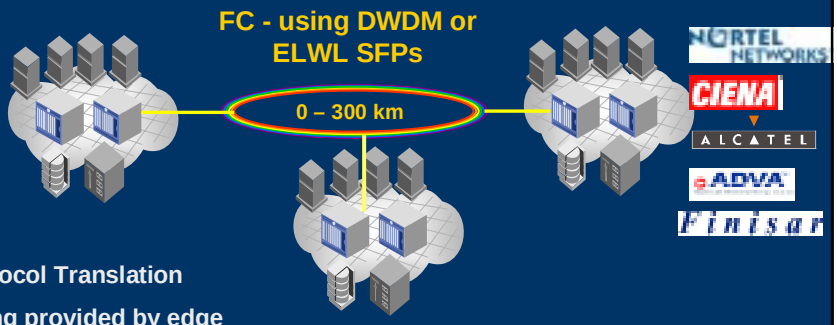

- No Protocol Translation
- Buffering provided by edge switches
- Sync applications (business continuance)
- Proven and Qualified Solutions In Production

Tested Configurations 2GB

Distance [km]	Throughput [MB/sec]
60	200
120	127

Brocade: Extended Fabrics

- Single, distributed fabric
- Spans up to 100 Km with special DWDM xxx KM 6000-BB

  in Production 800KM Frankfurt-London with SW12K
- Reallocates buffers in the ASIC to maintain full 1 or 2 Gbit/sec performance across MAN distances

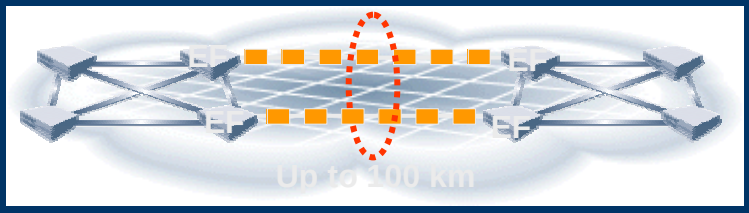

Distances & Transport Technologies

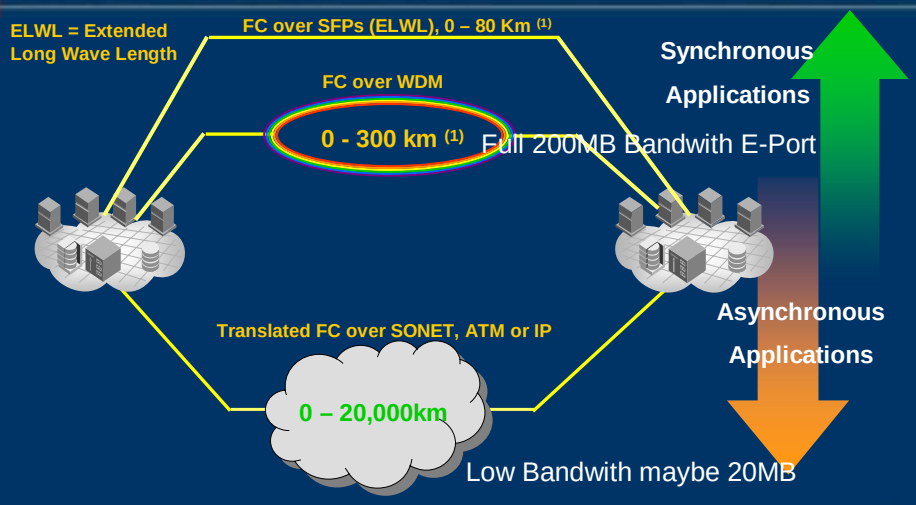

(1) Under optimum conditions, throughput can reach full speed (200MB/sec) at up to 60km and start degrading proportionally from there.

New Fabric Application Platform What means it for SAN Design?

- **NEXT Step Multiprotocol to SAN**
- **Real SAN Subnets with Port Overlay**
- **Applications in the Fabric**

-VTAL,Virtulization,Data Replication
-Data Movement
-Volume Management

BROCADE
The intelligent platform for networking storage

Fabric Application Platform Integration with and SAN Design

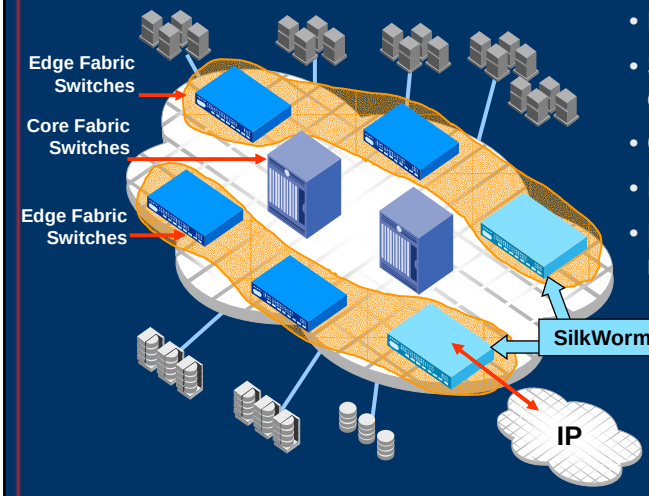

- E-port Interoperability
- Secure Fabric OS compatible
- Common CLI
- Fabric Manager support
- Fabric Access API for management

Brocade Platform for Fabric Applications

- High performance fabric application platform/switch
  - Runs storage applications in the fabric at wire speed
  - XPath™ Technology - future-proof, ASIC-based, multi-protocol fabric
  - Tremendous performance, 50.000, Ios per Port, scalability, economic advantages
  - Xcopy
  - Multi-protocol
  - -VTAL (Virtual Tape Library)
  - More 400 API cycle per sec

- Open software architecture for easy solution migration
  - Robust set of storage services and APIs
  - Affords fast and easy delivery of storage applications in the network

Fabric-based Virtual Tape Library
Example Configuration

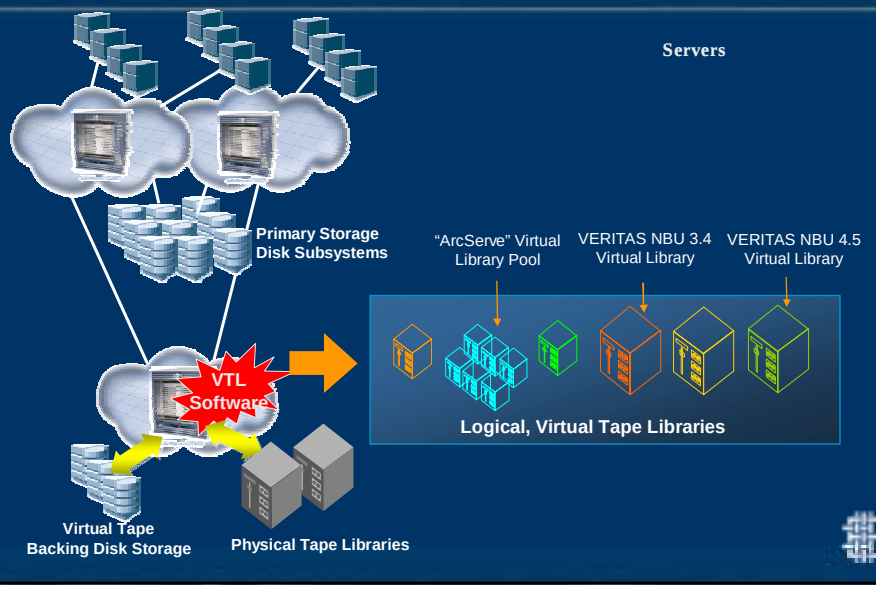

SAN Subnets: SilkWorm Inter-Fabric Switch ("I-Switch") The real SAN Subnet Solution !!

Accelerating Corporate SAN Growth

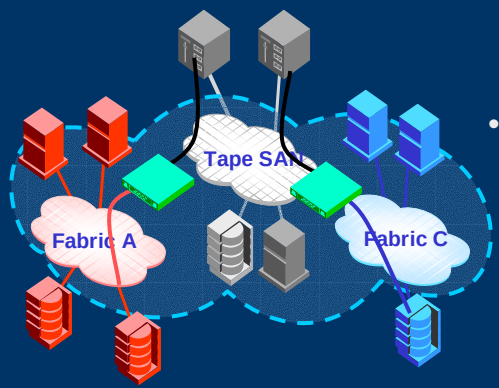

- Create SAN Subnets
  - Connections between devices on separate fabrics
  - Simplifies change management
  - Better fault isolation
- Allows separate organizations to share equipment
  - Common fabric for backup, accessed by separate SAN Islands

Is NO VSAN ; VSAN Only Partion the Box per Port, but no Overlay

Open Fabric Management
Choice of Management Frameworks

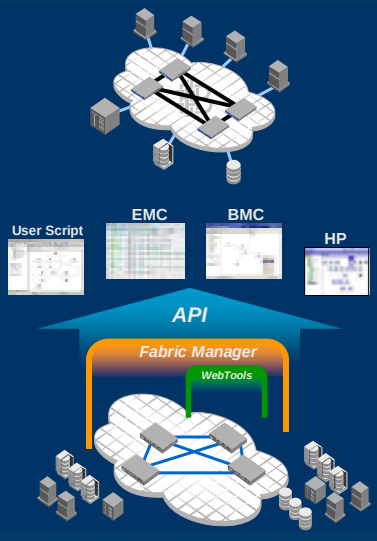

Auto Discovery and Configuration
- SAN discovers all servers, storage, and switches – zero device admin required
- Distributed Name Server enables any-to-any server-to-storage automatically

Powerful, Open Fabric Management
- Intuitive *Webtools* and *Fabric Manager* GUI applications
- Standard SNMP and GS3 mgmt

- **Comprehensive Open API**
  - Mgmt Integration with 40+ leading storage management products
  - Scriptable via end user mgmt tools

Lowest Administration Effort
Greatest Choice for Enterprise Management Integration

Brocade Fabric Manager
Screenshots

Fabric topology focused on BRCD fabric...

- Best-in-class technology partner (ILOG)
- Displays: Switch types, Links, Ports & Devices (clouds)
- Multiple network layouts (Hierarchical, Circular, Core-Edge, ++)
- Real-time fabric: Network, Status, Detail (expand/collapse mode)
- SAN Fabric specific concepts: Maintains non-competitive nature

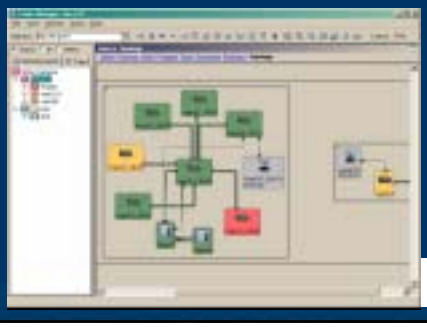

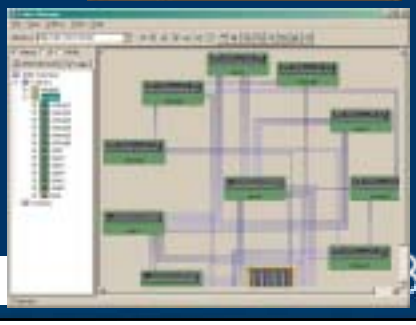

Thank You