



I/O Consistent

Non-Synchronous & Synchronous DB I/O



Enterprise SAN Design End-To-End SWAN Testing




NOTES

1. SilkWorm 12000 comes with dual control units, multiple power-supplies and fans, and has two 16 port switches configured as separate fabrics.
2. CNT UltraNet Edge Storage Router 1101 are used exclusively for the communication between HP SureStore XP 128 devices.
3. Based on the design and the backbone bandwidth limitation, Asynchronous communication between the XP 128 devices was mandated by the original vendor.

Network Considerations For SAN's

Applications over the MAN/WAN

- Moving large amounts of data between sites
- File sharing – supercomputing applications
- Remote mirroring / disaster recovery
- Synchronous mirroring
- Asynchronous mirroring
- ❖ Tape backup
- ❖ Centralizing Storage Management
- ❖ Disk edge caching
- ❖ Virtualization
- ❖ High availability / clustering
- ❖ Video streaming
- ❖ SSP remote primary storage



Moving Data Between Sites

- Fibre channel is the workhorse for moving data inside the data center
- Outside the data center there are two ways to move Terabytes between sites
 - Non real time – i.e. people power
 - Real time – direct connection
- Fundamentally, two things are necessary for a direct connection to efficiently move large amounts of data over distance
 - High Bandwidth interconnection
 - Flow Control - ability to fill and sustain bandwidth flow over time

Why Fibre Channel?

- Storage to storage transfers demands a high QoS
 - Needs “deterministic” communications
 - Generally very sensitive to latency
 - Requires in-order block delivery
- Fibre Channel is deterministic
 - Provides low latency and high bandwidth interconnect
- High performance SAN extension must maintain same QoS features over distance
 - High bandwidth
 - Low latency


How Much Bandwidth is Enough...?

- Moving 10 TB requires:
 - 2.25 hours using OC-192 (10Gb/s)
 - 9 Hours using OC-48 (2.5Gb/s)
 - 14 hours using "2G" FC (1600 Mb/s)
 - 28 hours using "1G" FC (800Mb/s)
 - 35.7 hours using OC-12 (622 Mb/s)
 - 6 days using OC-3 (155 Mb/s)
 - 20 days using T3 (45 Mb/s)
 - 1.6 years using T1 (1.5 Mb/s)



.....If the pipe is fully utilized!

Data Movement Today – A Recent Case Study




Scientists at LANL currently dump 100GB of supercomputing data to tape and FedEx it to SNL because it is faster than trying to use the existing 155Mb/s IP WAN connection

- Actual measured throughput of 16Mb/s! (10% bandwidth utilization)



The Better Way – Directly Between Storage Systems



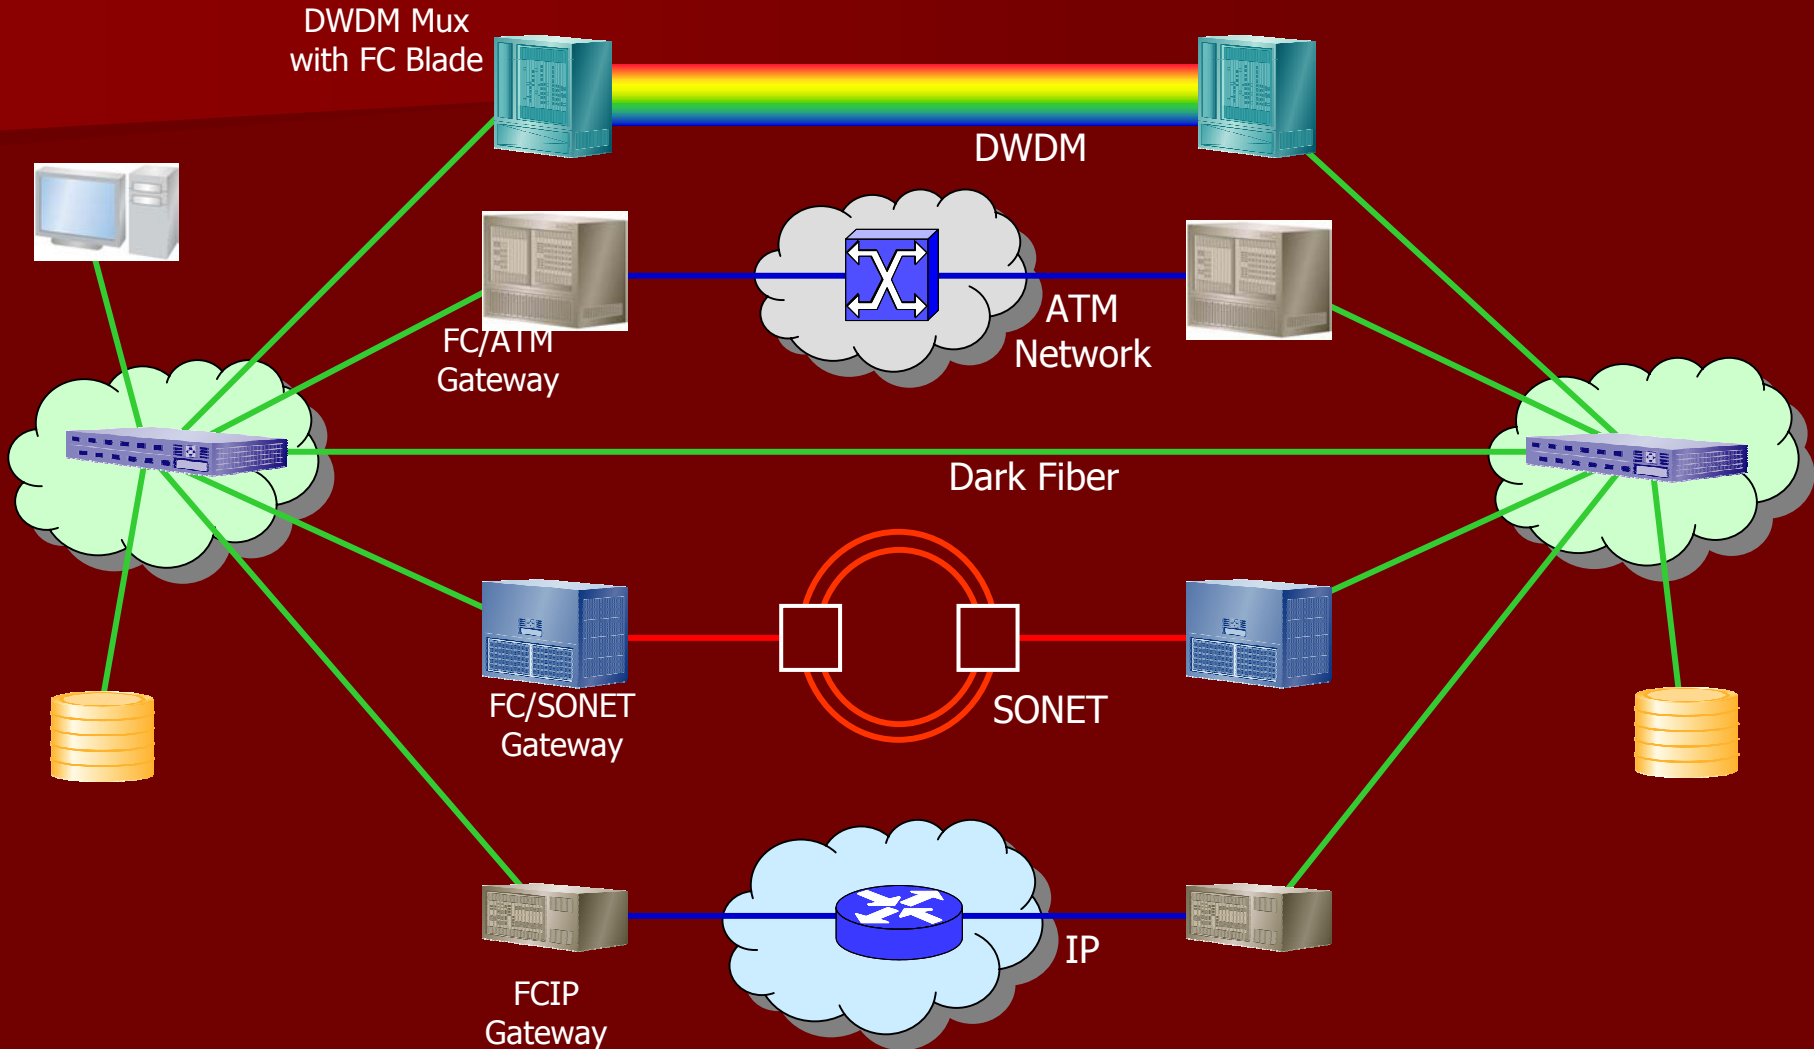
Utilizing the full bandwidth of this interconnection the data could have been moved in just over 80 minutes!

Transport Layer Extension




How do I move the signal?

- FC over Dark Fiber
- FC over DWDM
- FC over ATM
- FC over SONET
- FC over IP

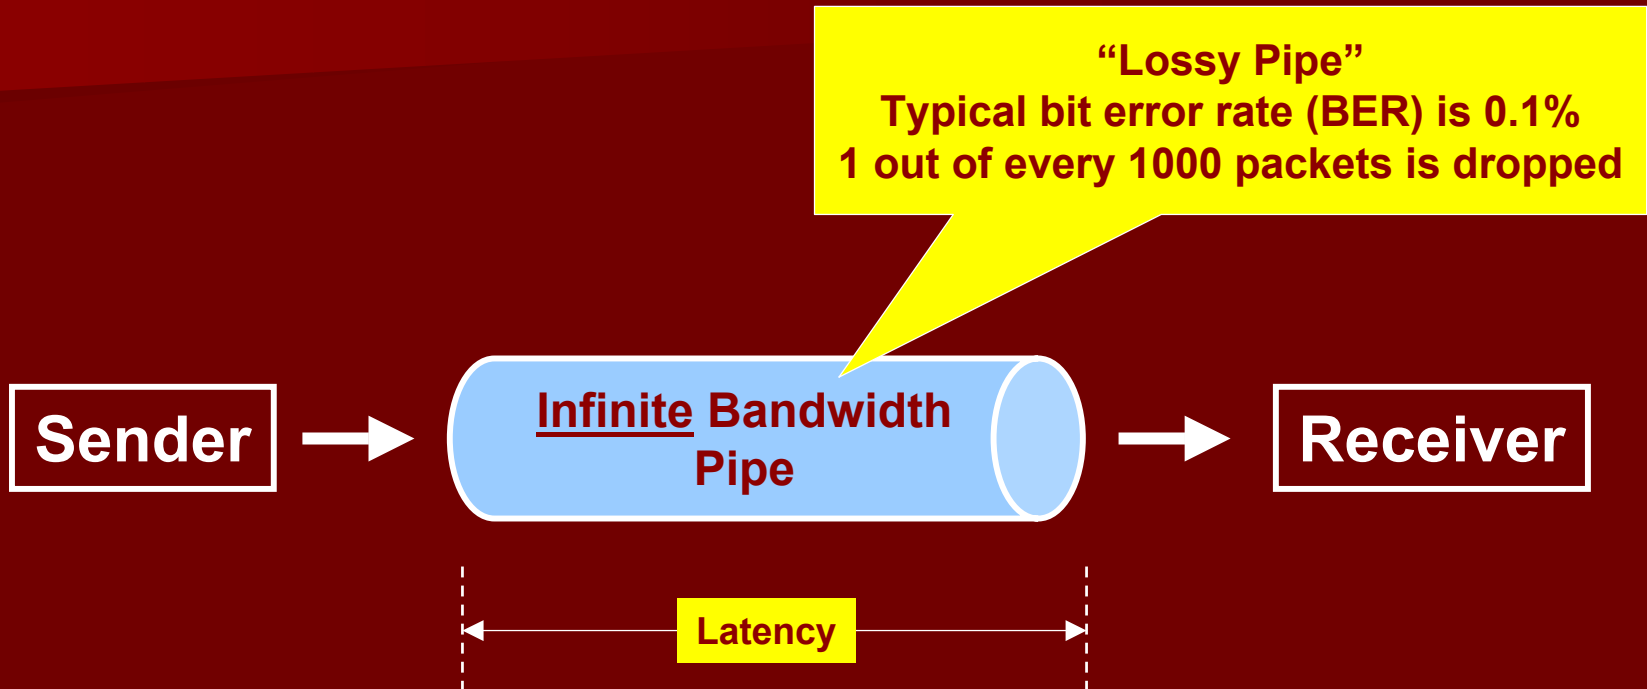
Transport Layer Extension through the MAN and WAN



FC Data Encapsulation Efficiency


	Storage over IP	Storage over SONET	Storage over ATM
Mapping	 <p>FCIP</p>	 <p>FC over SONET</p>	 <p>FC over ATM</p>
Mapping Efficiency	90%	98%	80%

“Fat Pipes” Don’t Guarantee High Throughput with TCP/IP



**With TCP/IP over the WAN,
BER and Latency affect throughput
more than the bandwidth of the pipes**

TCP/IP Performance vs. BER and Latency




What's Wrong with Routed IP in the WAN?

- Routed IP is designed for scalability and connectivity
- Routed IP is NOT designed for performance
- TCP algorithms are designed to make the end user back-off quickly at the first sign of congestion
- Throughput is controlled by latency (geographical distance) and packet loss rate
- IP Carriers will always have packet loss
 - Sell service by peak capacity
 - Allocate equipment based on average capacity


Network Layer Flow Control via Credit Buffering

- Fibre Channel maintains throughput in the data center by using flow control via buffer to buffer credits
 - Nominally FC switches provide credit buffering up to 10km distance
- Any wide bandwidth, long distance movement of FC data must couple flow control over the WAN
 - System requires end-to-end credit buffering
- WAN gateway flow Control options
 - Transparent mode - relies on the FC switches to extend the credits (good up to 130-150km) This method is used in all DWDM gateways and some SONET, FCIP devices
 - Coupled Credit mode – The gateway handles the buffer to buffer credits over the WAN. This is extensible up to great distances (1000's of km). This method is used in most FC over SONET gateways
 - FCIP gateways must provide credit buffering in conjunction with PAUSE and TCP/IP

Multiple Sources for Latency



Fibre Channel Over SONET / WDM




Data re-transmission due to IP packet loss

limits actual IP throughput over

Fibre Channel over SONET

The High Efficiency, Long Distance Alternative



Data re-transmission due to IP packet loss

limits actual IP throughput over

Storage Over IP

Using CNT Solutions




Joint Solution Overview

- Continuous Access XP
 - UltraNet Edge, UltraNet Storage Director
 - HP XP Disk Arrays (XP48/XP256/XP512/XP1024)
 - HP Continuous Access software
 - Cluster Extension XP and Continental Cluster
 - All inter-networking options – IP, ATM, T3, Metro Fiber optics
- StorageWorks DRM
 - Enterprise Modular Array (Enterprise Virtual Array)
 - DRM (Data Replication Manager) Software
 - All inter-networking options – IP, ATM, T3, Metro Fiber optics

Fibre Channel SAN Issues - Extensibility

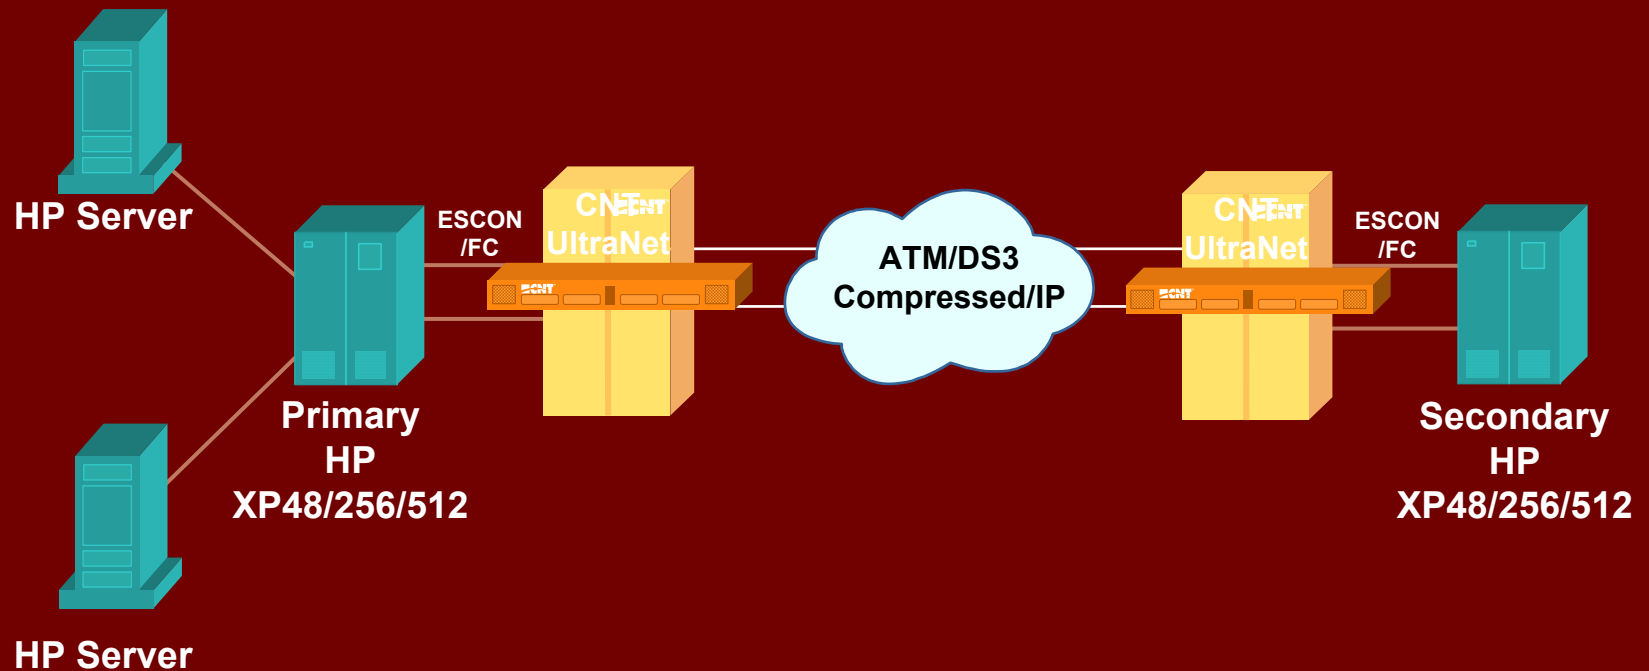
Native FC can not be transported over most service providers infrastructure.

Wide area networks are based on IP or Ethernet and not FC, therefore, costly point-to-point **Dark Fiber** or **DWDM** is required to connect FC SAN.



Fibre Channel SAN island


HP XP CA Remote Disk Mirroring Over WAN/IP



Storage Over IP

Using Nishan IP Storage Solutions

Asynchronous Mirroring – Going the




- **Customer required fast replication with 2 sites 2000 miles apart**
- **FASTWRITE technology excels at long-haul optimized performance**

Extending the Reach of Synchronous Replication

Vtesse Networks - Europe

Synchronous HDS TrueCopy



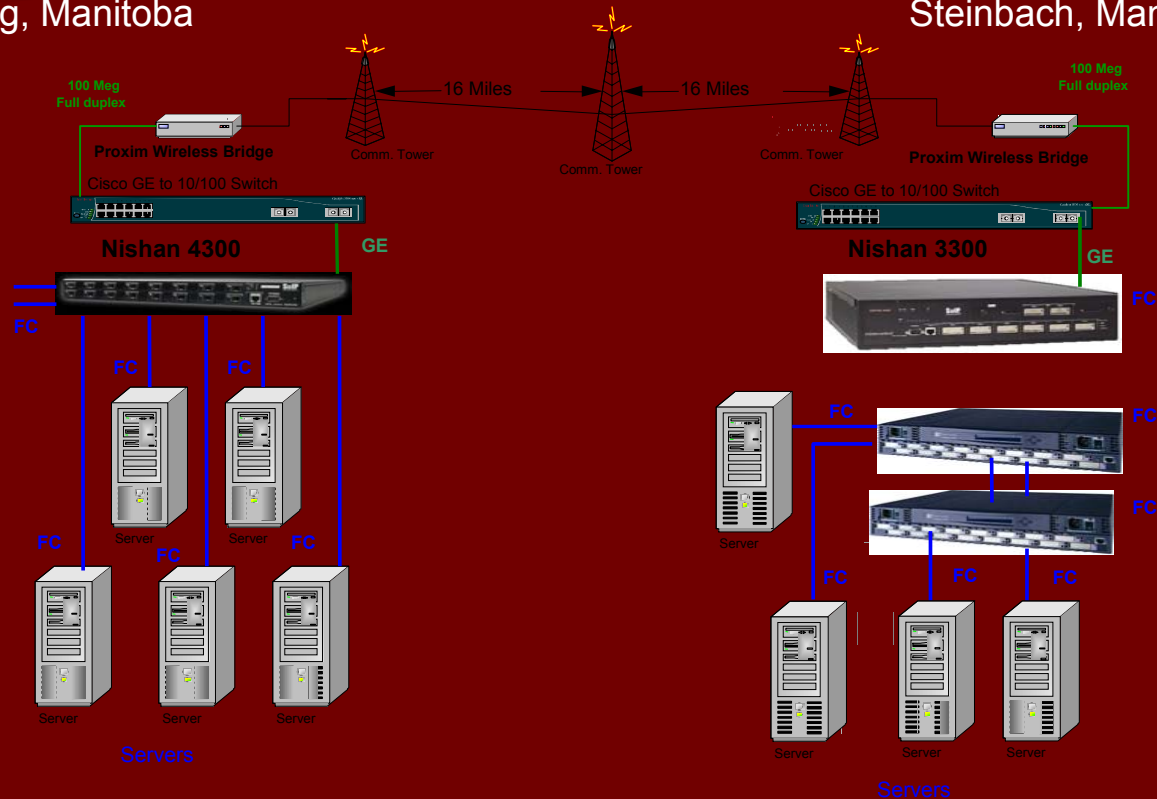
- Nishan switches convert FC mirroring traffic to IP
- Synchronous data replication up to 600KM
- Expands the reach of Sync TrueCopy => more TC addressable market !!

Steinbach Credit Union – Wireless IP SAN


STEINBACH CREDIT UNION

Winnipeg, Manitoba

Steinbach, Manitoba



Enterprise-class Scalability via SAN Internetworking



- **Fault isolation provides router-like scalability – Enterprise-class scalability**
- **Connectivity between Heterogeneous Fibre Channel switches**
- **Extensibility across any distance: campus, metro or WAN**


High Availability And Disaster Tolerant SAN Considerations

Storage WAN (SWAN)


➤ Features

- Manage the WORLD as a single Entity
- Wide area data sharing/migration
- Storage Management/Reporting
- Shadow/Remote Backup
- Wide area DT
- Connectivity within standard infrastructure (Network)

Span: The World




Fabric Extension




*Used to specify the speed of fiber optic networks. The base rate (OC-1) is 51.84 **Mbps**. OC-2 runs at twice the base rate, **OC-3** at three times the base rate (155.52 Mbps), etc. Planned rates are: OC-1, OC-3, OC-12 (622.08 Mbps), OC-24 (1.244 Gbps), and OC-48 (2.488 Gbps)

Data Replication

➤ Replication can be done at many levels



➤ Replication can be done at many levels




•Real Time COPY

- Provide Disaster Recovery
- **NOT to maintain two identical copies**
- Provide I/O consistent copy of data

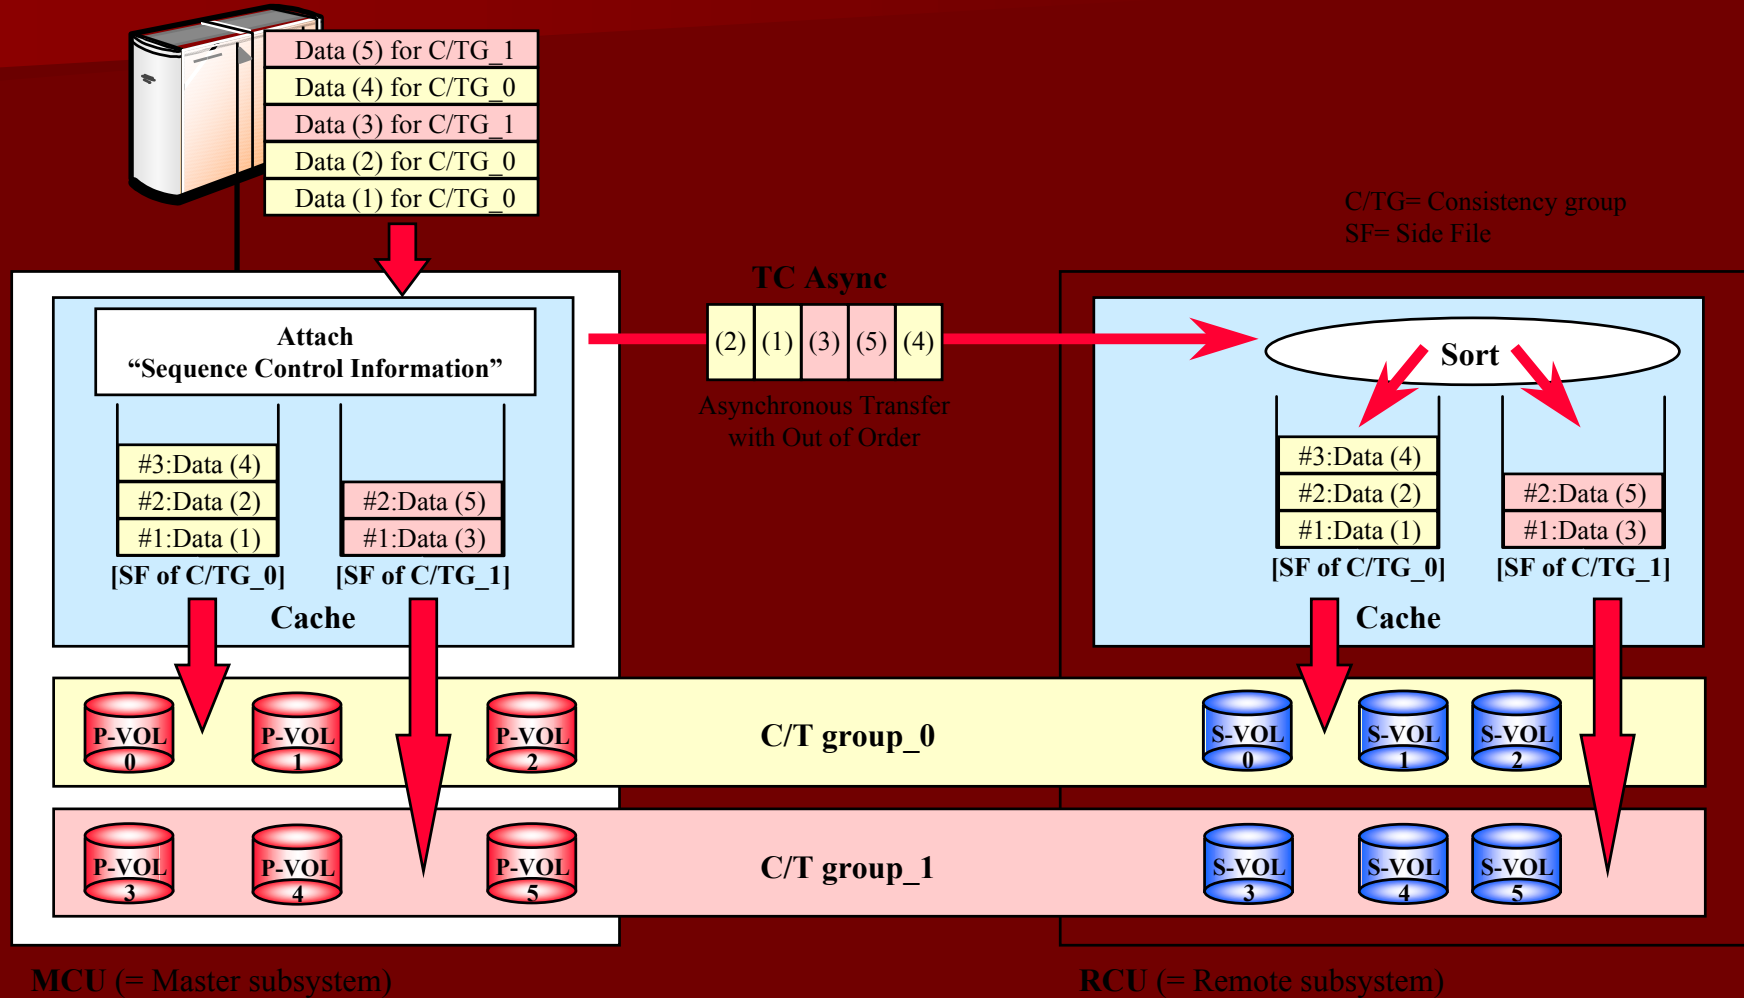
Synchronous Or Asynchronous

- If Within Supported Distance Use Synchronous Because:
 - Data is more secure
 - Best overall performance
- Asynchronous is supported
 - Cases w/low I/O rates and some potentially lost data is acceptable if links are broken

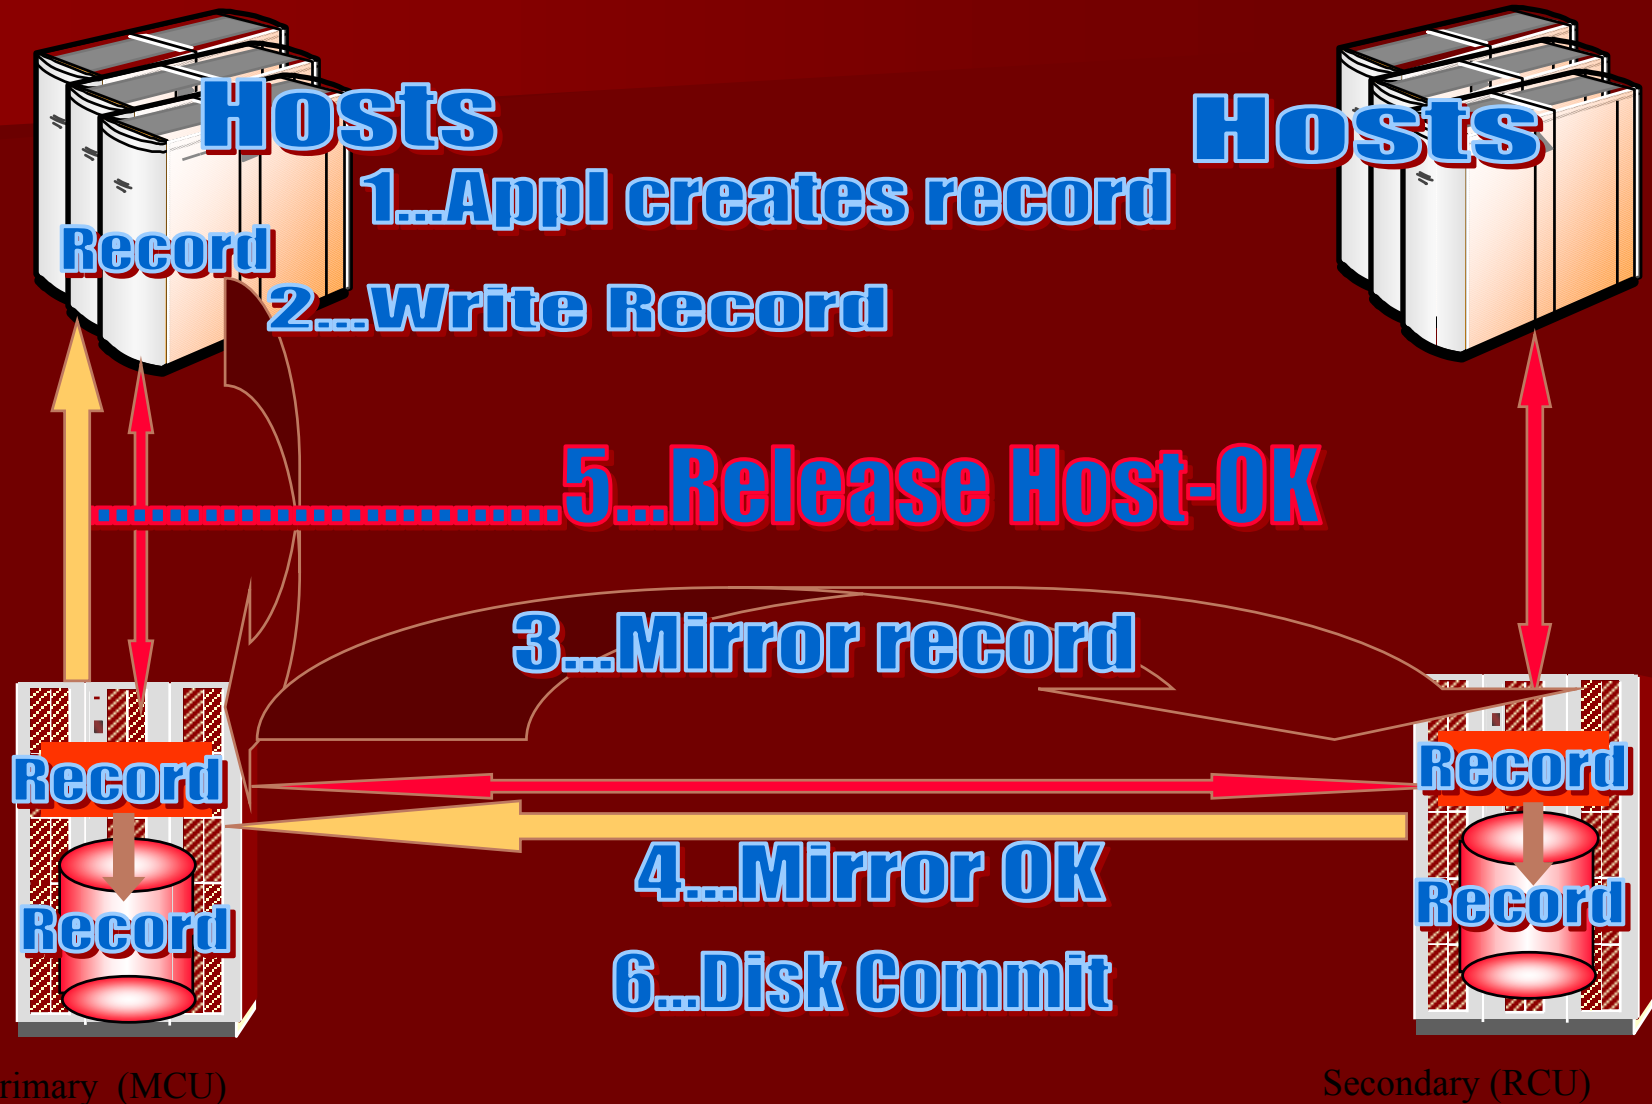
Asynchronous Operations



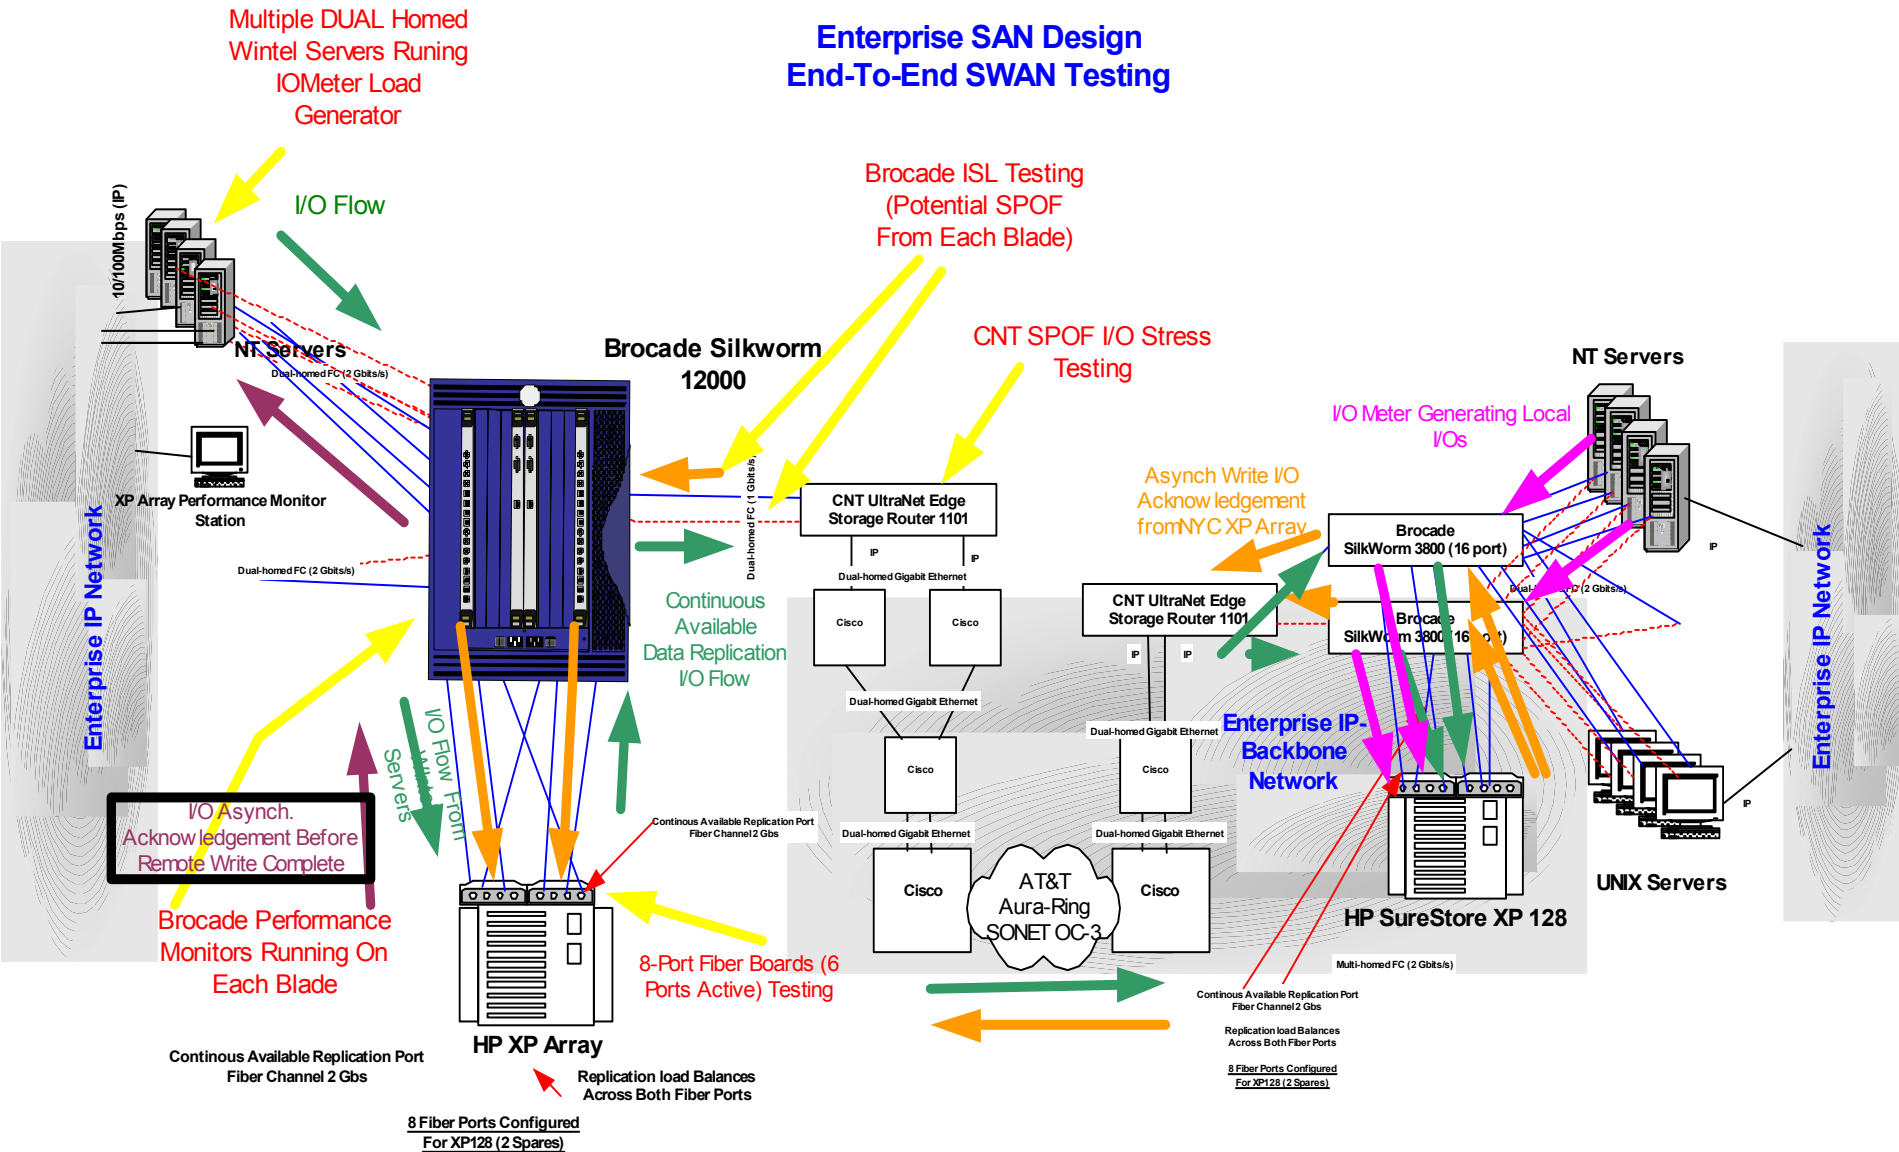
Asynchronous Update Sequence



Synchronous Operations



Enterprise SAN Design End-To-End SWAN Testing



NOTES

1. SilkWorm 12000 comes with dual control units, multiple power-supplies and fans, and has two 16 port switches configured as separate fabrics.
2. CNT UltraNet Edge Storage Router 1101 are used exclusively for the communication between HP XP 128 SAN's.
3. Based on the design and the backbone bandwidth limitation, only Asynchronous communication between the XP 128 SAN's was recommended.
5. The design will work (in Async. mode) with existing Cisco's Enterprise IP-backbone without using QoS, Packet Prioritization, Traffic Shaping, or Layer 3 switching technology.

Questions??

Thank You