



Storage over the MAN/WAN: How to Choose the Best Solution

Paul Schoenau

Title: Senior Product Manager

Email: pschoena@ciena.com

Company: CIENA

AGENDA




Application Overview

- WAN Networking Options & Considerations
- Deployment Examples
- Summary

Data Center Applications over the MAN/WAN

Applications

- Synchronous Disk Mirroring
- Asynchronous Disk Mirroring
- SAN Interconnection
- Tape Backup
- Geo-clustering
- ESCON/FICON Extension
- LAN extension



- To achieve the maximum performance at the lowest cost, every component of the solution from end-to-end must be considered in the design of the network
- This presentation will focus on the cost optimization of the MAN/WAN networking component

Requirements Storage Wide Area Networking

- **Cost**
 - Must meet Budget Constraints
 - Often bandwidth make up >50% of the project cost
- **Security**
 - Guaranteed Isolation of Sensitive Data
 - Guaranteed Data Delivery
- **Performance**
 - Minimal Impact on the Application with a High Throughput, Low Latency, and Rapid Restore Times
- **Capacity**
 - Intelligent utilization of network recourses
- **High Availability**
 - Five 9's, with the Ability to Monitor/Report/Protect to Maximize Performance and Perform Rapid Fault Isolation.
- **Flexibility**
 - Support for **all** Data Types (Storage, Voice, Data, Video) and Applications

AGENDA

- Application Overview

WAN Networking Options & Considerations

- Deployment Examples
- Summary

Choices Connecting Data Centers

Deploy Fiber between buildings

- build you own network (likely using WDM equipment)

Lease a connection from a carrier or service provider

- likely a SONET service, wavelength service, or an IP service
- manage your own network, and the gateways

Lease a Managed Storage Service


- combination of the gateway solution + the connectivity
- customized SLAs and service features focused on extension

MAN / WAN Options

**Private Build or Leased
(Campus / Metro only)**





**Leased Bandwidth
(Campus / Metro / Long-Haul)**



1. WDM / Dark Fiber
2. SONET/SDH
3. ATM
4. IP

Multiple services on a fiber, high capacity, campus or metro
Secure bandwidth, ubiquitous service access
Used by legacy channel extension technology, costly
Storage mapped into IP packets, usually through TCP/IP

Storage over WDM



Example Carrier


Today's Applications

- Fiber relief
- Native protocol carriage (GbE, Fibre Channel, ESCON)
- Virtual Private Network
- Bandwidth Leasing

Characteristics for Storage


- Large capacity (80G in Metro)
- Reach limited to metro distances
- DWDM is cost effective for large # of services and volumes of data (> 10 Gbps)
- CWDM provides a lower cost solution
- Single channel per service regardless of utilization

Channel Extenders / ATM



- Channel extenders represent a legacy approach to storage extension
 - Support ESCON/FC only neglecting overall data center extension requirement likely including GbE
 - ATM mapping adds significant overhead resulting in bandwidth inefficiency
 - Extremely high (starting at \$80K) pricing
 - Adds significant latency to applications not requiring “host emulation” resulting in decreased application performance

IP Storage over IP Networks



- Storage extension over IP networks (e.g. a leased GbE service) is only feasible if very high QoS SLAs can be guaranteed
 - Latency less than 5 ms
 - Packet Delivery Ratio (PDR) of 99.99% or higher
 - TCP can drive a maximum of 100 Mbps over a network with a latency of 5 ms and PDR of 99.99%
- IP networks must be dedicated to SAN extension in order to maximize PDR and minimize latency and even then, they can not support synchronous applications due to their high latency
- Bandwidth must be over-provisioned by up to 50% to account for dropped packets and retransmissions


Storage over SONET / SDH

- What is SONET / SDH?
 - Self-monitored high performance networking technology
 - Ubiquitous network with over 150,000 installed carrier rings
 - Uses Time Division Multiplexing (TDM) to aggregate multiple signals together
 - Standardized rates from 50Mbps – 40 Gbps
- Why SONET / SDH?
 - Guaranteed, high bandwidth
 - Low latency
 - Deterministic
 - Secure, 99.999% availability network
 - Metro and Long Haul networks
 - National and International

SONET/SDH Based Services Perfectly Match the Requirements of Business Continuance Applications

SONET as the Backbone

- 99% of all data traffic goes across the SONET network, including IP
- Native protocols are mapped directly into SONET as soon as they leave the building or campus where they originated



Bit Rate (Mbits/sec)	Electrical SONET	Optical SONET	SDH Equivalent
51.84	STS-1	OC-1	
155.52	STS-3	OC-3	STM-1
622.08	STS-12	OC-12	STM-4
2488.32	STS-48	OC-48	STM-16
9953.28	STS-192	OC-192	STM-64


**Native Services Are Transported
Worldwide Over SONET**

Performance Considerations

- Latency
 - Delay introduced by the intermediate equipment (i.e. switches, transport, speed of light in fiber) which slows down the response time of applications
- Bandwidth
 - Dedicate enough bandwidth to ensure optimum application performance, but only allocate the required to bandwidth to minimize MAN/WAN costs
 - Choose a storage networking technology that efficiently makes use of expensive MAN / WAN resources
- Protocol Flow Control
 - Not have sufficient protocol extension capabilities will leave the application waiting

**Latency, Bandwidth and Protocol Flow Control
can significantly impact application performance**

Latency Sources



1. Storage Infrastructure

- Very small compared with other sources; can be ignored

2. MAN / WAN Gateway

- The time it takes for the gateway to process the frame and send it into the MAN / WAN
- Can be very significant; varies by vendor
- With IP, it is increased significantly whenever MAN / WAN is shared with other applications

3. Network Latency

- Can dominate overall latency
- More detail on next slide


Performance Considerations: Network Latency

Time of Flight

- Time it takes light to traverse the network
- 5 μsec per km
- 1 ms per 125 miles



Element Switching




Technology	Switching Latency	Notes	Impact on Response Time
DWDM	<5 μsec	Lowest latency	None
SONET/SDH	<20 μsec	Low latency	None
ATM	100s of μsec		Significant
IP	1000s of μsec	Varies heavily on traffic load, quality of IP service, bandwidth contention	High

Performance Considerations: Bandwidth


- Bandwidth tradeoff between cost and application performance:
 - Cost:
 - MAN/WAN networking can be up to 50% of the BC/DR application cost
 - Minimize the bandwidth used through:
 - Data Compression
 - Efficient Mapping
 - Application Consolidation over a single WAN
 - Performance
 - Need to consider the actual throughput of data
 - Need to guarantee application isolation and efficient delivery of data
 - If not enough bandwidth is allocated, or too much contention in the network, a bottleneck to the application will occur

Optimum Solution:
Bandwidth Allocated = Application Requirement

Performance Considerations: Packet / Cell Loss




- Lost packet results in cutting throughput in half
- Throughput slowly recovers over time
- Similar effects with other IP and ATM Storage Solutions



**Due to lost packets and retransmissions,
goodput is artificially limited**


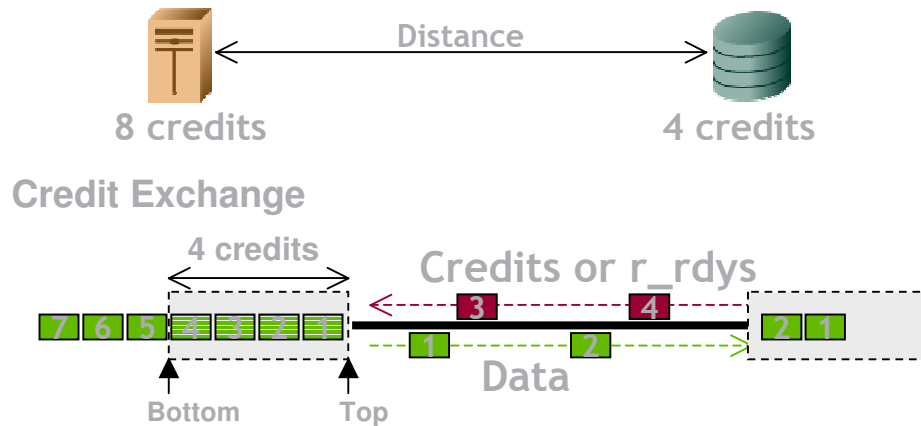
**Performance of Storage Across IP and ATM
Infrastructures is Significantly Affected by Loss**

Performance Considerations: Protocol Efficiency




	Storage over ATM	Storage over IP	Storage over SONET/SDH	DWDM
Mapping	Multiple Layers	Multiple Layers	Direct	Direct
Mapping Efficiency	80%	90%	98%	100%
Protocol Support	FC, ESCON, GbE	no ESCON Support	FC, ESCON, GbE	FC, ESCON, GbE
Retransmission due to loss	Limits performance	Severely impacts performance	0%	0%
Security	Isolated	Shared	Secure Layer 1 Isolation	Secure Layer 0 Isolation

Performance Considerations: Protocol Flow Control



- Each credit(r_rdy) carries one FC/FICON frame
- A FC/FICON frame cannot be sent until an r_rdy is received
- The receiving device allocates the credits to the sending device



- As the distance is increased, a set number of credits can't fill the pipe
- Throughput degrades significantly as distance is increased

Performance Considerations: Cost vs. Performance

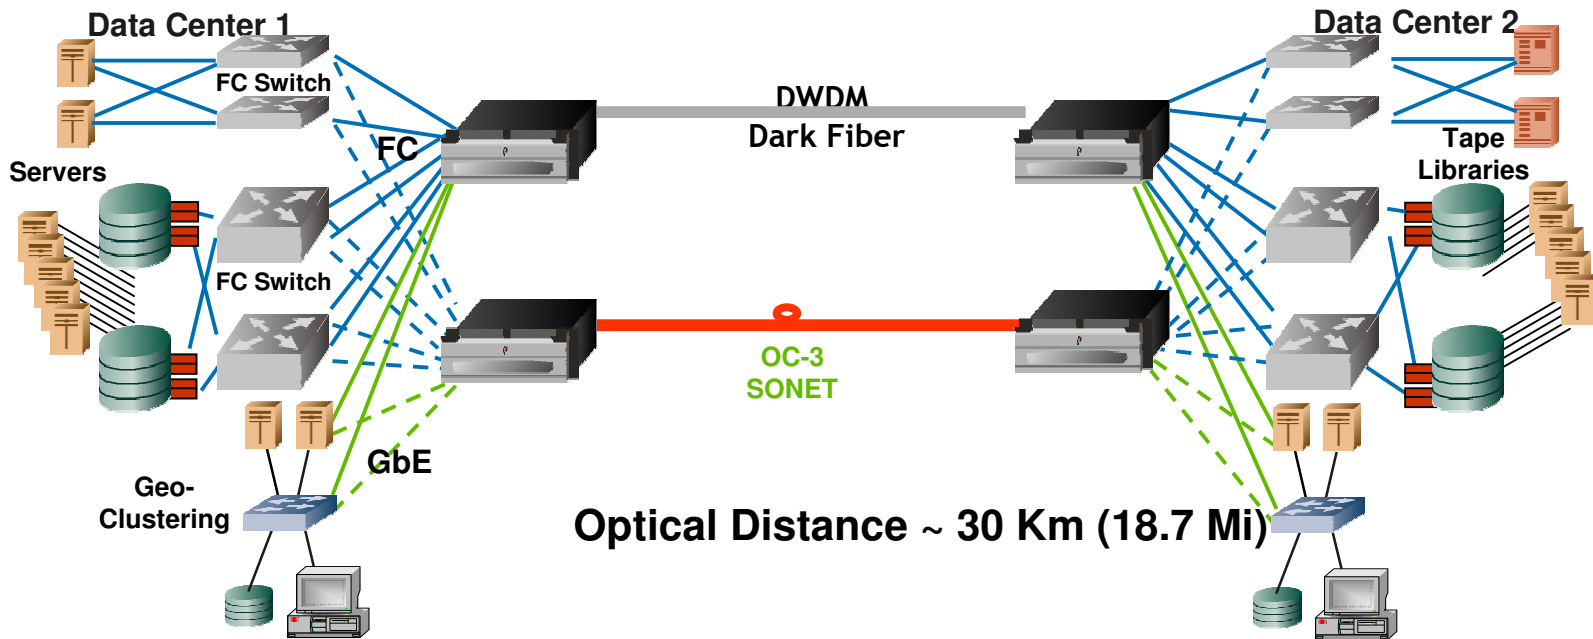
Technology	Cost of Equipment	Cost of Service	Bandwidth	Performance
WDM	\$\$	\$\$	Gbps	High - Guaranteed, never changes Limited to metro distances
SONET/SDH	\$	\$\$	Mbps to Gbps	High - Guaranteed, never changes Suitable for both Metro and Wide Area deployments
ATM	\$\$\$	\$\$	<= 150 Mbps	Med : Varies based on ATM service type (CBR, VBR)
IP – Native IP Service	\$	\$	No guarantees	Low : Varies over time, bandwidth contention of the entire IP network
IP – Router with PoS interface	\$\$	\$\$		Med : Varies over time, bandwidth contention with other services on the same router

Storage over SONET/SDH and WDM are the Best Options for Extending Business Continuance Applications

AGENDA

- Application Overview
- WAN Networking Options & Considerations
- ➔ Deployment Examples
- Summary


Deployment Example: DWDM and Storage over SONET



- Applications:
 - Geo-Clustering, Disk Mirroring and Remote Tape via Fibre Channel and Gigabit Ethernet
 - Enterprise has access to one pair of fiber
- Why the customer choose DWDM & Storage over SONET:
 - DWDM provided extremely low latency and significant capacity
 - 2.5 Gbps today scaling to 80 Gbps
 - Leased a SONET OC-3 service as a backup

Deployment Example: Storage over SONET


Storage over IP over Router with PoS (Packet over SONET)



- Applications:
 - HP Continuous Access via Fibre Channel, LAN via GbE
- Why the customer choose Storage over SONET:
 - > 50 % capital cost savings
 - 2X improvement in WAN utilization
 - Detailed performance monitoring of services and WAN


Deployment Example:

Data Migration with Storage over IP



- Applications:
 - Remote Storage via Fibre Channel
- Why did the customer select Storage over IP:
 - LAN connectivity existed between these two locations
 - No significant time restrictions for the migration
 - Storage over IP allowed the migration to use the WAN when bandwidth is available

Deployment Example: Aggregation over Carrier Managed Service



- Applications:
 - Enterprise is planning on leasing wavelengths from a service provider
 - Utilize Storage over SONET to aggregate multiple services over each DWDM wavelength
- Values to the Enterprise:
 - Cut down reoccurring bandwidth costs by >60%
 - Improve performance monitoring and fault isolation
 - Augment the DWDM with buffer credit extension

AGENDA

- Application Overview
- WAN Networking Options & Considerations
- Deployment Examples



Summary

Summary

- Costs due to WAN extension dominate the business case for BC/DR solutions
- Solutions vary greatly based on cost, performance, security and data availability targets – there is no single answer
- Carriers and Service providers must be leveraged for solutions beyond fiber only
- New products are available that both simplify the deployment of, and help drive down the costs associated with distance solutions for storage extension

HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

