

Moving Mainframe Tape Outside The Data Center

By BRIAN LARSEN

For years, enterprise IT has been faced with a variety of challenges that are almost mutually exclusive: support and enable business growth while reducing costs; increase data protection while improving global data access; meet shorter recovery objectives while managing growing amounts of data, to name a few.

To meet these and other business objectives, IT personnel require new approaches, new technologies and new infrastructure architectures. Many are exploring the use of electronic remote tape backup and remote virtual tape as a cost-effective approach for improved disaster protection and long term data retention. But given the sequential nature of tape operations, can tape be effectively extended over distance? This article focuses on the mainframe enterprise and leveraging FICON attached tape media over long distance to meet data retention and disaster recovery needs. It explores the business requirements, technological challenges, techniques to overcome distance constraints, and considerations for solution design.

Every company in the world has information that keeps its business afloat whether it is billing information, order information, assembly part numbers or financial records. Regardless of company size, it has critical information that keeps its business in operations. What data is deemed critical varies from company to company, but most agree that information on an employee's workstation is less critical than the data stored in a DB2 application. It is that ranking of criticality to application and data that determines what the most cost effective means will be to protect it.

There are many different methods that can be applied to protect information in offsite locations. These range from physically moving tape cartridges to a secure location to synchronized disk replication over distance for immediate failover of operations. Depending on what recovery point objective (RPO) and recovery time objective (RTO) is applied to specific information, different solutions can be implemented to meet those objectives and contain costs like additional equipment and wide area network charges.

RTO and RPO assignments can help in classifying applications or data to a level of protection and solution type to use. Figure 1 shows the guidelines used by a large financial institution to help move much faster toward implementation and deployment. They determined that five classes of data existed and that each classification would use a certain type of recovery or business continuity method. Once an application was classified they did not have to revisit which protection method to use.

Recovery Operations – RTO/RPO Classes

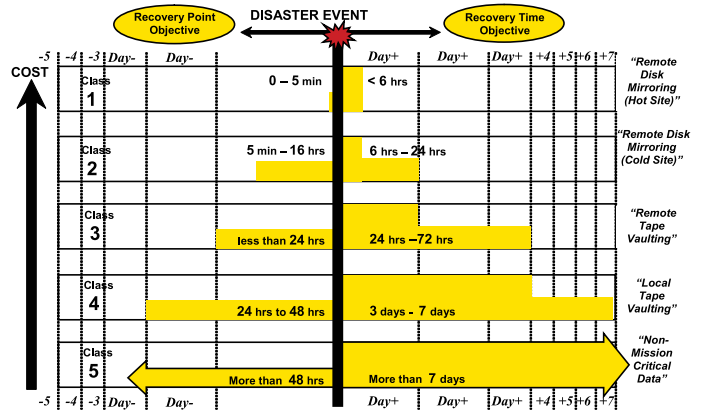


Figure 1 Recovery Point Objective and Recovery Time Objective Assignment to Protection Applications

In addition to disaster recovery needs, some businesses have specific requirements that must be met for retaining information. For instance, a company located in the Central U.S. that does credit card processing and billing for several large financial institutions and was mandated to retain records for no less than seven years. It determined that tape was the most cost effective media to use. It had three processing centers located as far as 1,800 miles from one another and wanted to electronically move tape traffic to its main data center where all the media management could take place. Its application processing was done with a mainframe host in each location and utilized a FICON based virtual tape subsystem that extended the physical tape attachment over a wide area network to its main processing center. Even though this wasn't deployed for disaster recovery, it did provide disaster recovery benefits, such as reduced risk of lost tapes due to handling, allowed centralized tape management and provided immediate access to tape data that was as current as the last backup.

Remote Storage Considerations

Regardless of the business reason for putting data in remote data centers, a few things should be considered when choosing a location. The first consideration should be around the "circle of disaster." The "circle of disaster" has to do with what a company believes is a safe distance between its primary and backup data centers or from one facility to another. This may seem like an easy decision but it adds a whole level of complexity and cost depending on the size of the circle and the resulting level of risk. The horrific events of Sept. 11 caused massive destruction within a few blocks. Some of those affected had enabled backup data centers that were located just across the Hudson River in NJ and their data was protected. However, in August 2003, the entire Northeast Region of the US was impacted by a power outage and a backup data center located across the river was on the same power grid as downtown Manhattan, causing both the production and the backup data center to be affected. What is a safe distance? That is a risk-based business decision that needs to be established by each organization, based on the size and type of business, location and risk tolerance. The implications of increased distance for increased safety have a huge impact on solution costs and what is affordable.

Another consideration is WAN bandwidth requirements. Once a business application is classified, the amount of data to be replicated must be understood and aligned with the wide area network (WAN) circuit that carries the data between locations. Analysts have concluded that nearly 60 percent of the total cost of ownership for a long distance storage application over three years will be the wide area network costs. These costs are determined by distance, location, speed of circuit and quality of service (QoS) requested. Given these large and recurring costs, a business may find its “circle of disaster” much smaller than it would like for optimum protection.

Using Tape in a Remote Storage Solution

If a remote electronic tape solution becomes the method of choice to secure data in another location, users can be confident that it can be done and at the performance levels required. Mainframe tape solutions have been a solid platform to use in local and remote situations for a number of years. In the late 1990s ESCON tape was being extended from the host across metro and long distances over WDM and dedicated leased telco circuits. Long distance solutions required specialized systems that would help maintain performance since the ESCON protocols could not sustain performance beyond 10 miles. As technology evolved from ESCON to FICON, distances supported have increased but still require specialized technology to reach across long distances and outside the “circle of disaster”. Figure 2 provides a comparison of throughput and distances supported as the port speeds have increased.

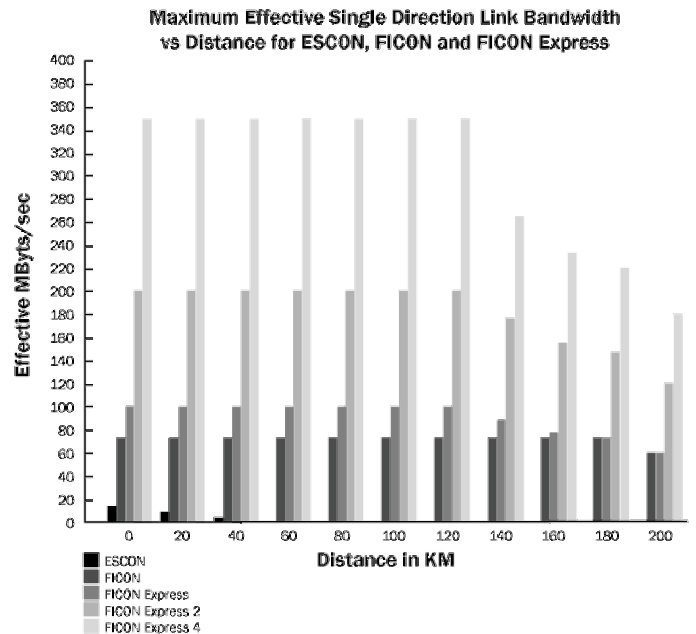


Figure 2
Technology Advancements – Performance vs Distance

FICON Capable Tape

Mainframe tape options vary from vendor to vendor but for the most part there are a variety of standalone drives, automated libraries or virtual tape subsystems that leverage internal disk to house tape data before actually exporting to a real physical media. Selecting the right product for operational needs is a matter of fitting the requirements.

Since all the drives use the FICON standard, techniques needed for long distance deployment have been tested with the major vendors and interoperability issues should be minimal. The FICON standard has evolved along with the Fibre Channel standard since the two use the same lower level protocols at the Layer 1 and 2 of the protocol stack. In addition to the advancements of FICON/FC port speeds there has been a move toward using IP as the transport of choice for long distance connectivity. FICON over IP is based on the FCIP standard but is uniquely modified to support the mainframe channel protocol.

Write operations can be expected to comprise the larger percentage of I/O operations for tape devices (for archival purposes) but given today’s requirement for data retrieval and shorter recovery times, Read operations are becoming more and more critical. Both write and read operations are necessary for stand alone tape extension or remote virtual tape solutions. Virtual tape configurations have a peer-to-peer relationship between the virtual tape controllers and replicate data between them. Peer-to-peer deployments can also be clustered which requires host access to both the primary and secondary virtual controllers. In a clustered virtual tape configuration, write and read operations work together to enhance data access and improve flexibility of system administration. A clustered configuration can take advantage of a primary and secondary subsystem fail-over, allowing for more effective management of maintenance windows. Without both write and read performance enhancements, users could not leverage the fail-over capability and maintain operational performance.

KETCH Consulting

Growing from Disaster Recovery to Business Continuity!

Leading the Way - KETCH Consulting

- *Senior Consultants
- *Certified
- *Experienced
- *Knowledgeable

Contact KETCH Consulting Today (888) 538-2492

- www.ketchconsulting.com -

Handling FICON Tape Protocol

Improved FICON protocol efficiencies reduce the number of end-to-end exchanges required to support tape operations compared with its predecessors ESCON and parallel channel implementations. However, many legacy access methods generate small channel programs consisting of as little as a single read or write CCW, normally preceded in a chain by an operating system supplied mode-set command and in some cases a terminating no-op command. Therefore, small channel programs that support tape operations are still serialized on a device basis by the command-data-status exchanges that typify tape read and write operations.

While these end-to-end exchanges may be considered trivial in native FICON attached tape implementations, they can become a significant impediment to acceptable I/O access times and bandwidth utilization for WAN supported FICON configurations. In addition to the command-data-status exchange required to support tape operations, the effect of IU pacing may also introduce additional undesirable delays in FICON attached tape devices accessed through WAN facilities, particularly for tape write operations where outbound data frames figure significantly in the IU pacing algorithm. Tape pipelining functions (command emulation) reduce the undesirable effects of latency on these exchanges and improve overall performance for WAN extended FICON attached tape devices.

Tape pipelining refers to the concept of maintaining a series of I/O operations across a host-WAN-device environment (not to be

confused with the normal FICON streaming of CCWs and data within a single command chain). Normally tape access methods can be expected to read data sequentially until they reach the end-of-file delimiters (tape marks) or to write data sequentially until either the data set is closed or an end-of-tape condition occurs (multi-volume file). Tape pipelining attempts to optimize performance for sequential reads and writes while accommodating any non-conforming conditions in a lower performance non-emulating frame shuttle.

Tape Pipelining Overview

The concept of pipelining is to reduce the number of host to device command and data sequences over distance. Figure 3 shows a typical configuration where the mainframe host is attached to the FICON tape controller over specialized network systems that convert the FICON tape protocol into an IP network protocol and becomes the long distance transport.

WAN Link	Speed in Mbps	Speed in MBytes/ per second	Effective Speed @ 2:1 Compressed MB/sec
T3	45 Mbps	5.6 MBytes/sec	11.2 MBytes/sec
OC3	155 Mbps	19.4 MBytes/sec	38.8 MBytes/sec
OC12	622 Mbps	77.8 MBytes/sec	155.6 MBytes/sec
Gigabit Ethernet	1000 Mbps	125 MBytes/sec	250 MBytes/sec
OC48	2488 Mbps	311 MBytes/sec	622 MBytes/sec

Figure 3
Typical remote FICON tape configuration

The wide area network could be hundreds or thousands of miles, so that latency is what pipelining will address. In order to reduce the effect of latency, the host must be “virtually” extended to the remote site and the tape control unit (CU) must be “virtually” extended to the host location. If the systems can be virtualized (or emulated) to be in each location, the data and control signals are responded to as if they were local, allowing data to be continuously sent without the impact of latency. Figure 4 provides a simplified view of virtualizing or extending the host to the remote tape location for local-like write operations.

For read operations the tape CU is extended to the host location, providing native read performance. Figure 5 provides a view of an extended CU. When the host and tape CU is virtualized as shown, the impact of latency is overcome. It also provides increased flexibility for storage placement, allowing tape resources to be located where needed, regardless of distance.

Pipelining does remove the latency of the network; however application performance is still a matter of having the bandwidth necessary to move the required data.

Sizing Bandwidth Requirements

Every site will need to analyze the data needing backup and apply that to the wide area network technologies available through carriers in each location. Historical information will need to be understood so that bandwidth requirements are not under- or over-estimated. If the estimate is low, jobs will back up and probably not finish in scheduled production windows. If the estimate is too high, expensive components that are not required add to costs.

1.888.340.8010

IMMEDIATE RESPONSE
FOR DISASTER, EMERGENCY RESOURCES & BUSINESS CONTINUITY

PROVIDING EQUIPMENT, SUPPLIES & SERVICES

GENERATORS
Sales/Rental Generators (Units - 5kw - 22.8 mw)
Portable Power Distribution Panels, Extension Cables, Maintenance, Technical Support & Light Towers.

PORTABLE HEATING/AIR CONDITIONING EQUIPMENT & COOLING TOWER RENTALS

BASE CAMPS/EMERGENCY HOUSING & OFFICES TURN-KEY OPERATIONS
Base Camps, Emergency Housing, We Respond in 100, 500, 1000, 7000, and 100,000 Person-Camps. These Camps are “turn-key” – included Lodging, Food Services, Restrooms, Showers, Laundry, Office/Administration, Recreation & Janitorial Services, 24/7 Security.

BUSINESS CONTINUITY-TURN-KEY
Temporary Office Structures, Heating & AC, Equipped with Office Equipment, Work Stations, Phones, Modular Furniture Setup, Satellite Systems, Food Services & Restrooms.

SPECIAL NEEDS
Blankets, Cots, Sleeping Mats, Clean Up Kits, Hygiene/Comfort Kits, CERT Kits and Supplies.

LOGISTICS, PROCUREMENT, PROJECT MANAGEMENT

FOOD
Catering (complete service), Semi Trailer Kitchen Full Services, Refrigeration Units, MRE (Meals Ready to Eat), & Survival Kits.

WATER PRODUCTS
H/D Purification Equipment and Wastewater Treatment Systems. Ice Blocks and/or Bags, Bottled Water and Water Storage Tanks, Bladders & Containers.

MOBILE ANIMAL SHELTER (BUNKHOUSE DRU)
CRITICAL CARE FACILITIES-HOOKUPS/PUMPS, SEMI TRAILER WATER TANKERS

MOBILE KITCHENS, SHOWER AND LAUNDRY UNITS, HANDWASH STATIONS, MEDICAL CLINICS



MOBILE RESTROOM UNITS (WITH REAL FLUSHING TOILETS), NBC CHEMICAL/ BI-CHEMICAL SUPPLIES

MOBILE COMMAND CENTERS, EMERGENCY GOLF CARTS, FORCE PROTECTION WALLS.

PORTABLE MASS FACILITIES STORAGE FACILITIES AND UNITS

MOBILE RAILROAD SYSTEM HOUSING UNITS

Unlimited Resources, Inc.
Emergency Telephone Number 888.340.8010 or 703.622.6946
Visit our Website at www.unlimitedresources.com

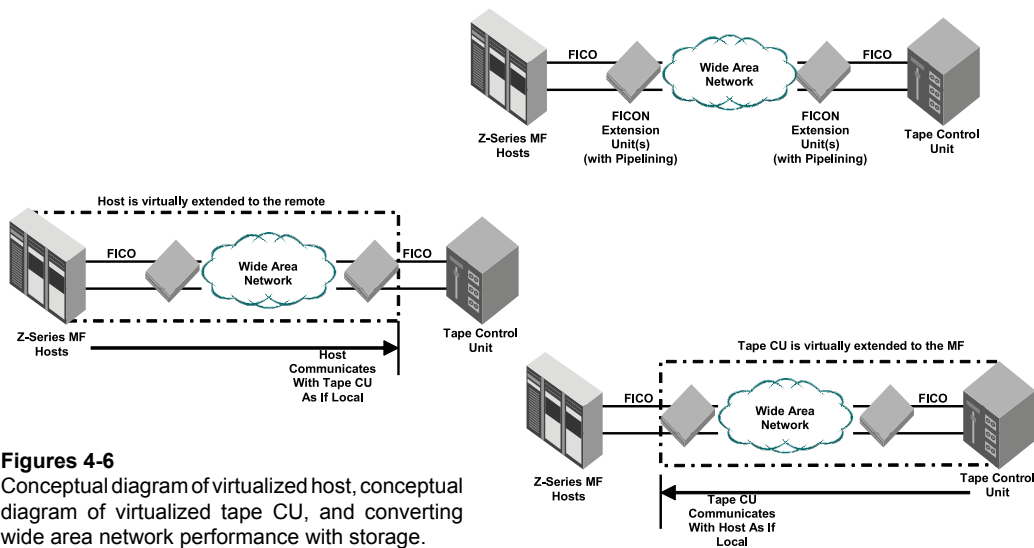



Once performance is understood, as well as the amount of storage to backup, it needs to be converted to telecommunication terms. Once this is done, speed and performance can be matched with similar terminology. Telecommunication companies use “megabits per second” (Mbps) when referring to the performance of their offerings. Figure 6 provides a simple matrix that can help convert storage, typically measured in Gigabytes (or larger) into Mbps. It also provides a performance estimate if the data being transported can be compressed at a 2:1 ratio. Depending on the data, program and drive configuration, data may or may not be able to be compressed over the network.

When going through the design phase of a remote tape solution, proper planning and research needs to be done to design the right solution for today and tomorrow. It is highly recommended

FICON Write Tape Pipelining

FICON Base Tape and Virtual Tape Extension



Figures 4-6
Conceptual diagram of virtualized host, conceptual diagram of virtualized tape CU, and converting wide area network performance with storage.

to work with the storage and network infrastructure vendors to help balance business objectives, performance, scalability, and cost. The following items should be discussed throughout the planning cycle:

- ◆ **Dual or Multi-Site Connectivity – Recovery and Data Access**
- ◆ **Recovery Scenario's – How to bring production back**
- ◆ **Attributes required deploying successfully**
 - **Performance** – can pipelining techniques be used?
 - **Flexibility** – can this solution scale as needed?
 - **Network Optimization** – can compression techniques be used?
 - **Network Resilience** - alternate network routes to remote storage.

Conclusion

When data management techniques are applied, remote mainframe tape is a high-performing option for disaster recovery and remote data access. Pipelining is an important technology to consider when developing business objectives, especially for situations that involve any amount of distance between sites. Read and write operations should have pipelining available to make sure a backup and recovery solution can meet performance needs, and provide investment protection as business requirements change and evolve over time.



Brian D. Larsen joined Brocade in July 1991 and has more than 24 years of professional experience in the high-tech industry. He is responsible for the strategy and marketing activities relating to Brocade's enterprise solutions across all markets with an emphasis on mainframe solutions and specializes in long distance storage applications. Larsen has focused exclusively on storage networking since 1997. He has presented storage networking topics at industry events and authored a number of articles and white papers on storage networking and IP storage. In addition to his 10 years in product management and solutions marketing, Larsen has more than five years of experience in sales and systems consulting, with responsibility for direct customer consulting including network design and project management. He has nine years of experience in technical systems support for worldwide network operations. Prior to joining Brocade, he worked for Unisys Corporation. Larsen holds a management information systems degree from Buena Vista College.

What happens when the phone goes out?

With LIVE 24*7 answering service from Time Communications you'll never have to find out. Lost orders, customer service delays, unhappy customers, are all a thing of the past with Time Communications.

For over 30 years we have provided backup telephone and email support to leading companies.

Call Trisha Stenberg or Michael Eastwood today at 800-486-8581 to learn more.

TIME COMMUNICATIONS

www.TimeCommunications.biz

Because it can happen, it will happen...it's only a question of when!