IBM *e*server **zSeries**

FICON Extended Distance Solution (FEDS)

The Optimal Transport Solution for Backup and Recovery in a Metropolitan Network

Author:
Brian Fallon
bfallon@us.ibm.com

FEDS: The Optimal Transport Solution for Backup and Recovery in a Metropolitan Area Network (MAN)

Traditional means of backup and recovery for customer data centers are typically implemented using channel extension schemes. These connectivity requirements are well understood by networking and data center personnel alike. T1, T3, telco lines are time division multiplexed onto Synchronous Optical Network rings (SONET) between data center sites delivering tape, print, and disk data. SONET rings can typically handle a capacity of 2.5 gigabits per second of data. But as businesses have grown, so have their requirements for bandwidth. This paper focuses on a scalable solution, built to deliver higher capacity at faster speeds for lower cost. FICON Extended Distance Solution (FEDS) is the answer for those customers that require such capacity within 150 KM (90 fiber miles) in the Metropolitan Area Network.

The Problem:

In today's world, terabytes of disk data are used for image and video applications. Traditional channel extension requires transport based on T1, T3, OC/3, over telecommunications lines as shown in Figure 1. These telecommunication line speeds do not have sufficient bandwidth to transport the kind of capacity required for today's Metropolitan Area Networks (MAN).

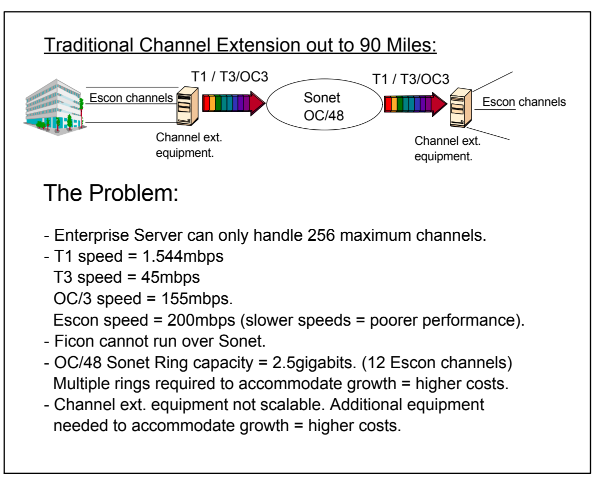

Figure 1: Traditional Channel Extension

As stated previously, legacy channel extension utilizes a SONET ring infrastructure that on average can only handle a capacity between 622 megabits per second (Mbps) to 2.5 gigabits per second (Gbps). This equates to a maximum of about 12 ESCON® channels in a 2.5 Gbps SONET ring and does not take into consideration other types of batch and interactive data traffic. Customers can find themselves trapped in a never ending bandwidth crunch, that requires more and more T1s, T3s, and SONET rings, just to accommodate even the most basic data growth within their organization.

The Answer: FICON Extended Distance Solution (FEDS)

Several finance and securities customers required a more scalable channel extension solution to meet their data needs for backup and recovery within 150 KM (90 fiber miles). To meet this challenge, IBM® constructed a FEDS lab at their Gaithersburg location. The following customer requirements were addressed:

- Reduce I/O channels on IBM S/390® Enterprise Servers and IBM @server zSeries .
- Remove data buffering equipment.
- Provide for transport scalability.
- Provide for seamless migration from an asynchronous disk copy to a synchronous disk copy technology when implementing the Geographically Dispersed Parallel Sysplex™ (GDPS™) solution.

FICON: FEDS utilizes IBM's fiber connection (FICON™) technology to transport significant amounts of data at high speeds between data centers over 150 KM (90 fiber miles) apart. FICON is the next step in channel I/O technology, available on the S/390 Enterprise Server and zSeries. FICON replaces the current ESCON channel implementation. FICON resides at layer 5 of the fibre channel protocol stack (refer to Figure 2). A FICON channel can deliver data at speeds of up to 1.0625 Gbps at distances of over 150 KM (90 fiber miles) as compared to T1 (1.544 Mbps) or T3 (45 Mbps). FICON reduces the number of protocol exchanges, as compared to ESCON, and therefore can transport at longer distances with minimal performance impact. FICON can preserve the customer's investment in S/390 and zSeries I/O capacity by replacing a number of ESCON channels with a fewer number of high capacity FICON channels.

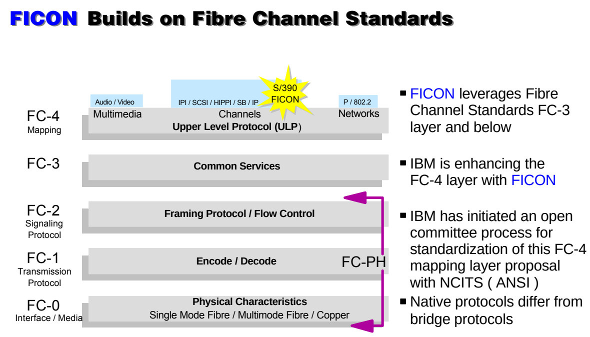

DWDM: Dense Wavelength Division Multiplexing (DWDM) can take multiple data streams (up to 80 gigabytes when using an IBM Fiber Saver 2029 Dense Wavelength Division Multiplexer) and multiplex them over just two pair of optical fibers. This, used in conjunction with Optical amplifiers, can extend the FICON channel out beyond 150 KM (90 fiber miles).

Remote Copy: Remote copy provides for the ability to mirror or copy data from the application or local site to the recovery or secondary site. This can be achieved through two specific implementations:

- eXtended Remote Copy or XRC is a combined hardware and software asynchronous remote copy solution. The application I/O is signaled completed when the data update to the primary storage is completed. Subsequently, a DFSMSdfp™ component, called System Data Mover (SDM), asynchronously offloads data from the primary storage subsystem's cache and updates the secondary disk volumes in the recovery site. Data consistency in a XRC environment is ensured by the Consistency Group (CG) processing performed by the System Data Mover (SDM). The CG contains records that have their order of update preserved across multiple Logical Control Units within a storage subsystem and across multiple storage subsystems. XRC operation results in minimal performance impact to the application systems at the application site. Unlike Peer -to-Peer Remote Copy (PPRC), XRC is not distance dependent.
- Peer-to-Peer Remote Copy or PPRC is a hardware solution that synchronously mirrors data residing on a set of disk volumes, called the primary volumes in the application site, to secondary disk volumes on a second system at another site, the recovery site. Only when the application site storage subsystem receives "write complete" from the recovery site storage subsystem is the application I/O signaled completed. PPRC is distance sensitive.

Test Setup and Test Results:

Test Setup: At the IBM Gaithersburg test setup, two LPARs, LPAR1 and LPAR2 running on an IBM 9672-ZZ7 zSeries simulated the application site and the recovery site as shown in Figure 3. The I/O driver executed in the application site's LPAR 1 copying data from IBM Enterprise Storage Server 1 (ESS1) to ESS2 to simulate a production environment. The eXtended Remote Copy / System Data Mover (XRC/SDM) executed in the recovery site's LPAR2 asynchronously copying data from the primary XRC disk subsystem ESS2 in the application site to the secondary disk subsystem ESS3 in the recovery site. Two FICON channels from the SDM image were Dense Wavelength Division Multiplexed using an IBM 2029 Fiber Saver in the recovery site and another IBM 2029 Fiber Saver in the application site. In the application site, two FICON channels were split out over 16 ESCON channels via an IBM 9032-5 ESCON director with the FICON bridge card feature. The 16 ESCON channels in turn attach to ESS2.

. **Figure 3: FICON / XRC Test Setup**

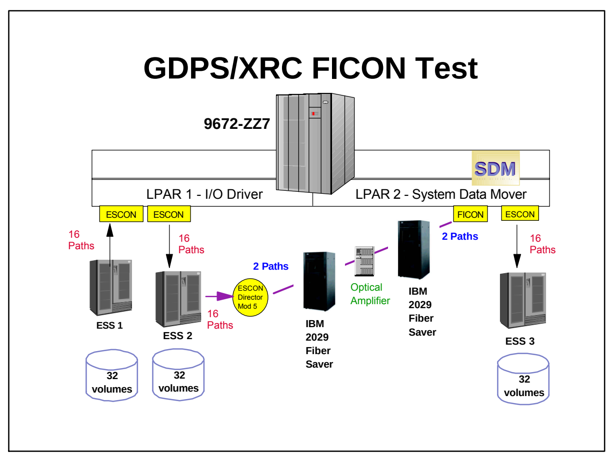

Tests: Various tests, indicative of XRC's and FICON's scalability over distances ranging from 0KM to 150KM, were conducted.

- Elapsed time to initially synchronize 32 volumes from ESS2 to ESS3.
- Sustained throughput achieved during the 32 volume synchronization process from ESS2 to ESS3.
- Elapsed time to copy 100% write data, generated by 32 DSS Copy jobs from ESS1 to ESS2 and copying updated data from 32 volumes of ESS2 to ESS3 using XRC.
- Sustained throughput to copy 100% write data, generated by 32 DSS Copy jobs from ESS1 to ESS2 and copying updated data from 32 volumes of ESS2 to ESS3 using XRC.
- The final test required "breaking" the primary fiber path and dynamically switching to the secondary or backup fiber path without loss of data.

Test Results: Figure 4 shows the test results summarized below:

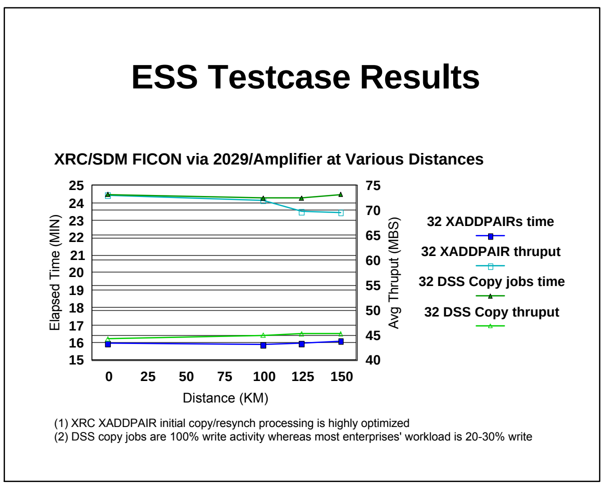

Figure 4: FICON / XRC Test Results

- Elapsed time to initialize 32 volumes (XADDPAIRs) remained consistent from 0 Km to 150 Km. The time ranged from 15:53 min (0 Km) to 16:32 min (150 Km).
- 32 XADDPAIR aggregate throughput achieved a consistent 70 to 74 MBS (mega bytes per second) for the 2 FICON channels from 0 to 150 KM. The average throughput per FICON channel was 35 to 37 MBS.
- 32 DSS copy jobs from ESS1 to ESS2 along with using XRC to copy updates from 32 volumes of ESS2 to ESS3 had an elapsed time of 24.5 min and remained consistent, when the distance was varied from 0 to 150 km.
- 32 DSS copy jobs from ESS1 to ESS2 along with using XRC to copy updates from 32 volumes of ESS2 to ESS3 aggregate throughput achieved a consistent 46 MBS for the 2 FICON channels from 0 to 150 Km. The average throughput per FICON channel was 23 MBS.
- When the primary fiber path failed, the 2029 automatically switched over to the secondary fiber path without any impact to the XRC remote copy application. XRC/SDM continued normal execution after the switchover.

FEDS, The Right Choice For Metropolitan GDPS

Geographically Dispersed Parallel Sysplex (GDPS) is a multisite management facility that is a combination of system code and automation that utilizes the capabilities of Parallel Sysplex® technology, storage subsystem mirroring and databases to manage enterprise servers, storage, and network resources. It is designed to minimize and potentially eliminate the impact of a disaster or planned site outage. It provides the ability to perform a controlled site switch for both planned and unplanned site outages, with no data loss in a synchronous environment (PPRC) and minimal data loss in an asynchronous environment (XRC). GDPS maintains full data integrity across multiple volumes and storage subsystems.

Currently, GDPS in synchronous mode (PPRC), known as GDPS/PPRC, is restricted to 40 km (24 fiber miles) between data centers. In the future, GDPS/PPRC is planned to be able to extend the distance between data centers to 100 km (60 fiber miles). In the interim, customers can implement GDPS for disaster recovery using XRC, known as GDPS/XRC. Both forms of disk copy can use FEDS for optical fiber transport - XRC over FEDS today, and PPRC over FEDS in the future. By using FEDS, the customer can protect their transport infrastructure investment. (Refer to figures 5 and 6).

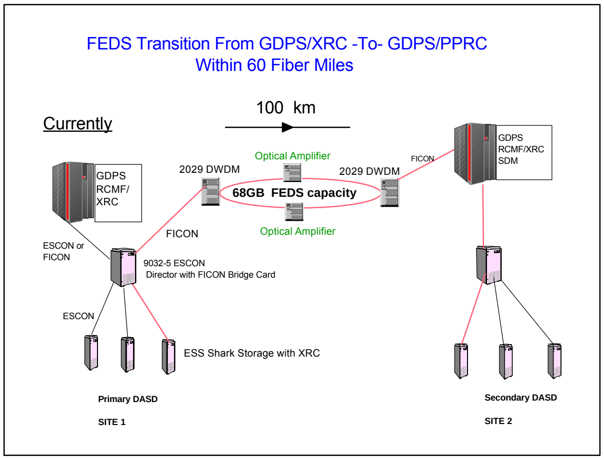

Figure 5: Current GDPS/XRC Implementation

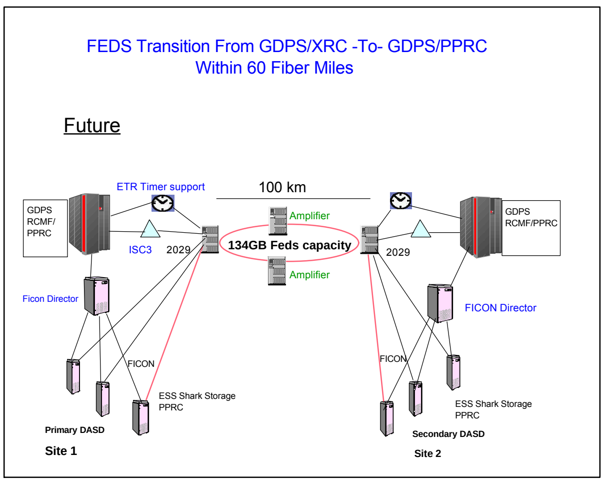

Figure 6: Future GDPS/PPRC Implementation

Summary:

The FICON Extended Distance Solution (FEDS) has solved the traditional channel extension problems of insufficient bandwidth and scalability. The advantages provided by FEDS are:

- Scalable bandwidth upto 80 GB
- Consistent throughput independent of distance from 0 to 150 Km (90 fiber miles)
- Fewer I/O channels on the enterprise server
- Faster data rates
- Lower cost and simpler implementation
- A migration path from GDPS/XRC to GDPS/PPRC.

Additional Information:

FICON:
FICON (FCV Mode) Planning Guide, SG24-5445-00
S/390 FICON Implementation Guide, SG24-5169

IBM Fiber Saver:
Fiber Saver (2029) Implementation Guide SG24-5608
Fiber Saver (2029) Planning and Maintenance
SC28-6801

GDPS:
ibm.com/servers/eserver/zseries/pso/

Acknowledgments:

David B. Petersen
Enterprise Server Group
petersen@us.ibm.com

Noshir Dhondy
Enterprise Server Group
dhondy@us.ibm.com

Ian R. Wright
Enterprise System Connectivity Specialist
iwright@us.ibm.com

IBM Corporation
Integrated Marketing Communications, Server Group
Route 100
Somers, NY 10589

Produced in the United States of America
09-01

GM13-0092-00.
ZSW03336-USEN-0