Proper Sizing and Modeling of ESCON to FICON Migrations
A Comprehensive Methodology for Planning and Cost Justifying for Performance

Stephen R. Guendert, McDATA Corporation

Numerous articles and white papers have been written on FICON technology since its debut in the marketplace in the late 1990s. This paper discusses the basic technology concepts, the relief of bandwidth constraints between processor and storage devices, the performance and distance advantages of FICON over ESCON and how to measure those advantages. In 2005 and beyond, when investing in new mainframe, mainframe disk, tape, or virtual tape storage equipment, companies will most likely be evaluating FICON equipment. Most modern disk, tape, and virtual tape subsystems are so fast that FICON channels are needed to exploit the performance potential they offer or, more importantly, they need FICON channels for to realize the performance capabilities for which the company is paying. A good analogy: purchasing a Porsche 911 Turbo in the U.S – the buyer is paying for thoroughbred performance, but cannot (legally) utilize the full capabilities of the engine. Likewise, purchasing a new mainframe, or new mainframe storage in 2005 and running it on ESCON does not fully tap that DASD array's and/or tape library's performance potential, and hence the full benefits of the investment cannot be realized. Much like buying imported sports cars, moving to FICON can be an expensive undertaking, but from doing a proper analysis, that focuses on the entire environment, including hosts, disk, tape, virtual tape, switching, and physical infrastructure (cabling), it will become apparent that there is significant total cost of ownership savings (TCO) realized by migrating off of ESCON to FICON.

Introduction

FICON has been around for five or more years. All storage vendors are in at least their second iteration of the hardware, software and firmware. With FICON Express2, IBM is in the third iteration of FICON channel technology. From all the performance studies that have been done, it is quite clear that FICON is a major technology improvement over ESCON and that data centers that have migrated from ESCON to FICON have seen improved response times and performance. Additionally, the FICON protocol has many inherent advantages over the ESCON protocol. FICON supports distances well over 10 km with hardly any speed degradation or data droop. FICON supports higher effective data rates per link, fewer links will be required: therefore, a FICON solution becomes a more cost-effective purchase. Even so, only an estimated one third of existing ESCON customers worldwide has migrated to FICON. When the technology advantages are so readily apparent, one must ask why there are so few who have migrated to FICON.

Two major factors are responsible for the seeming lack of rapid adoption of FICON technology. First, many of the concepts associated with FICON, particularly those associated with planning the environment, are diametrically opposed to what the mainframers have learned in the preceding 35 years of parallel and ESCON connectivity. Couple that with the lack of a clear, concise, and statistically scientific planning methodology that emphasizes storage device performance and not just channel utilization statistics, the reason for hesitancy among a very traditionally risk-conscious, conservative user base is obvious. A second major factor in delaying a migration to FICON is the initial cost of entry, including the purchase of new hardware and infrastructure. One thing that is difficult for mainframe end users to recognize and quantify is the cost savings associated with the elimination of some ESCON infrastructure, as well as their ability to leverage other existing ESCON hardware by attaching it to the FICON network.

This paper will touch very briefly on ESCON and FICON as technologies. Should the reader wish to expand on the technical subject matter, the bibliography has a comprehensive reference list. The majority of the paper will focus on how to model and size a new FICON environment and the cost justifications. Disk, tape, and virtual tape will be addressed as individual sections for these explorations. The predominant discussion will focus on disk and the cost savings realized in a DASD environment by migrating from ESCON to FICON.

ESCON configuration planning/technology review

Capacity staff have experience in the planning of ESCON configurations; therefore, it is useful to review some of the history and key points of that technology, as well as ESCON configuration rules and how they influenced the design of current channel and director configurations.

ESCON technology was developed by IBM and announced in 1990. It was the first true storage networking technology and replaced the mainframe parallel channel (bus and tag) which was strictly point to point and slower, with extremely short distance connectivity. Some of the key particulars follow.

<u>Bandwidth</u>

The theoretical ESCON bandwidth limit is 20 MB/sec, but the maximum real world data rate for each ESCON channel is about 17 Mbyte/s during a data transfer peak. While this data rate can be achieved for an individual job like a backup/restore, it is clearly impossible to reach such data rates during a typical peak hour. Most installations plan for a peak channel utilization of less than 60%, resulting in an absolute maximum effective data rate of about 10 Mbyte/s during the peak period. With higher channel utilizations, pend time and other queuing delays slow down the system too much.

<u>I/O limits</u>

There is also a limit to the number of I/O operations that an ESCON channel can handle since the channel is 'connected' during the interpretation of the I/O action. The connect time for a simple I/O like a 4 Kbyte read-hit can range anywhere from 0.5 to 1 ms, mostly related to the disk subsystem year of manufacture. For a typical workload, with larger blocks and cache misses as well. This means that the number of I/O operations per channel will be well below 1000 I/O operations per second (significantly below the theoretical ESCON limit of 1200 I/O operations per second).

<u>ESCON Architecture</u>

ESCON has inherent addressing limitations in its protocol. Up to 1024 different devices per ESCON channel. Thus, for a disk subsystem, this means that a set of channels is required for every 4 logical control units (LCUs) with 256 addresses each. It also means that only 3 to 3.5 terabytes (TB) can be addressed when using 1024 volumes with 3390 Model 3 formats. Modern controllers support 16 or more LCUs, so the number of required ESCON channels will grow proportionally.

<u>Circuit Switched/Unidirectional Protocol (half duplex)</u>

ESCON provides a circuit switched, unidirectional data transfer mechanism. Once a data transfer for an I/O from the channel to the subsystem or from the subsystem to the channel has begun, no other I/O operation is allowed to utilize the channel until the initial transfer has been completed. This leads to an interesting phenomenon: about 75% of the reported connect (CONN) time in an ESCON environment represents time periods in which the channel is occupied but not transferring data.

ESCON channels also retain what is often referred to as the "Mother-May-I" half duplex/unidirectional characteristics that were employed by the parallel channels. Basically, an ESCON channel transmits only a single Channel Command Word (CCW) at a time and waits for the channel-end/device-end notification. This series of "handshakes" continues until the entire channel program has been transmitted one CCW at a time. This protocol dates back to the non-intelligent (memory-less) 360 era storage subsystems and processors and is the reason behind the distance limitations of ESCON and the phenomenon IBM now refers to as ESCON droop.

FICON Considerations and Technology Overview

FICON was developed by IBM and was announced in the 1998/99 timeframe. FICON's specification is contained in the FC-SB2/SB3 command set. FICON is based on a 100 MB/sec (currently FICON is 200 MB/sec with FICON Express2) and is an

upper level protocol of the broader fibre channel standard. (see figure 1)
.

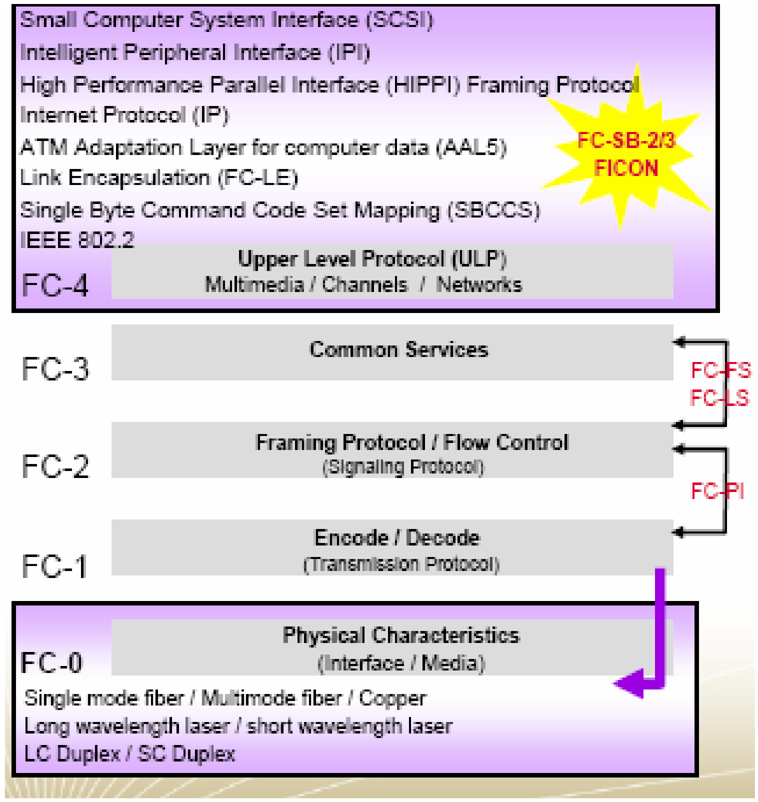

Figure 1. Fibre Channel Standard [2]

<u>Bandwidth and I/O limits</u>

In its latest iteration, FICON Express2 can achieve a theoretical throughput of 270 MB/sec based on IBM performance lab testing. However, published standards are 200 MB/sec with 400 MB/sec likely sometime in 2006. RMF also breaks down FICON measurements into READ and WRITE numbers (in bytes). FICON Express2 is also capable of 13,000 I/Os per second with 4k block sizes.

<u>Packet switched and bi-directional (full duplex)</u>

Unlike ESCON, FICON is a packet switched bi-directional protocol. Multiple bi-directional transfers of data can be simultaneously active over a single FICON channel. This eliminates the “Mother-May-I” issue that is the root cause of ESCON droop over distance. Instead, we now have “assumed completion” where the entire channel program (up to 16 CCWs) is transmitted to the subsystem at the start of the I/O operation. After the storage subsystem completes the entire channel program, it will notify the channel with a channel-end/device-end (CE/DE) transmission. This eliminates the multiple turn-around/protocol handshakes inherent in the ESCON protocol. It allows the FICON architecture to support distances up to 100 KM between the channel and the storage subsystem without droop. Figure 2 shows how this affects performance over distance when comparing ESCON to FICON.

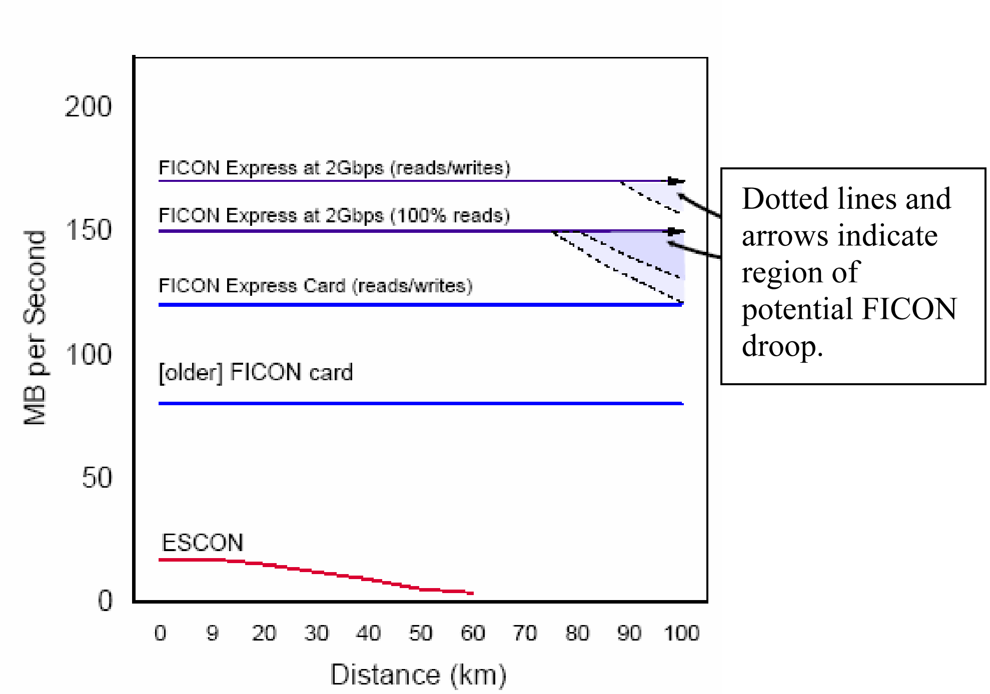

Figure 2 DATA Droop ESCON vs. FICON

The packet switching and interleaving characteristics have an impact on CONN time for a FICON environment. While the handshake overhead and inactivity on the channel is eliminated, another type of overhead, frame pacing delay has been created. Frame pacing delay is the delay due to the interleaving of I/Os on FICON channel paths when those paths become busy. This tends to elongate CONN time when the paths become heavily loaded. The CONN time of any I/O is a function of: 1) the number of packets to be transmitted; 2) the ability of the subsystem to send the packets; 3) the utilization of the path (i.e. the number of available idle packets); and 4) the distance between the subsystem and the channel. In other words, the CONN time can be a function directly correlated with the number of concurrent I/Os being processed by a given subsystem and/or the percentage of bandwidth in use. Thus, as the multi-programming level of a given system increases, we can expect to see its service time elongate proportionally.

FICON and ESCON Performance measurements

Whole papers have been written solely on this subject, and in an effort to avoid trivializing them, it is highly recommended that the readers refer to several of these papers documented in the bibliography section at the end of this paper.

Technical benefits of FICON over ESCON

When performing a comprehensive technical evaluation of each protocol, FICON proves to be a major technological improvement over ESCON. This was true in 1998/1999 at its inception and is even more true in 2005 with 2 Gbps technology (4 Gbps on the way), cascaded FICON, and FICON/FCP intermix. Figure 4 summarizes the comparison of the two protocol technologies. [ELIO92], [TROW02]

Function	FICON (Native and bridge)	ESCON
Switching	Packet Switching	Circuit Switching
Command Execution	Asynchronous	Synchronous Channel/CU handshake
Data Transfer	Simultaneous read and write Full duplex data transfers Connectionless	Read or write Half duplex data transfers Connection oriented
Link Data rate (bandwidth)	200MB/sec (FICON Express)	17 MB/sec
Maximum Throughput (I/O)	3600 I/Os per sec (G5/G6) 130000 I/O per sec (FE2 z990)	1200 I/O per sec
Distance (no repeaters)	10km, 20km(RPQ)	3km
Repeated Distance without significant data degradation (droop)	100km	9km
Frame Transfer Buffer	128KB	1KB
CU Images/CU link	255	16
UAs/channel	16K device addresses	1K device addresses
UAs/Control Unit	4K(FICON Bridge) 16K (FICON Native)	1K

Figure 4 ESCON vs. FICON Technical Comparison Summary

The technical advantages that contribute directly to business and/or cost advantages include FICON's greater addressing limits, improved bandwidth and I/O capabilities, improved distance and enhancements made to the protocol itself allowing for better response times. We will discuss in detail how this translates to cost savings with respect to DASD environments, tape/virtual tape environments, disaster recovery applications, the host, and the physical infrastructure in the appropriate sections. However, at a high level, the costs/benefits are summarized below:

<u>Business Benefits of Migrating to FICON</u>

There are 6 primary business benefits realized when migrating from ESCON to FICON. These will be listed here and discussed in more detail relative to specific device types such as DASD, tape, etc.

1) The technical advantages of FICON over ESCON enhance overall performance in the mainframe environment, meaning more work can be performed in less time. More transactions can be processed in a given amount of time or, looking at it another way, less time goes by between processed transactions. This has been extremely important in the financial sector and is one of the primary reasons that industry moved to FICON *en masse*. In the financial sector, the speed in which transactions can be executed is directly tied to a company's ability to generate revenue. Moving to FICON can translate directly to additional revenue and profit gains.
2) FICON enhances overall enterprise resiliency and disaster recovery planning based on the extended distances that the protocol supports. The bandwidth capability of FICON enables faster recovery over those distances. The faster performance of FICON also allows an enterprise to better meet its Recovery Point and Recovery Time Objectives (RPO and RTO respectively). FICON's improved performance compared with ESCON enables DR site disk volumes to be addressed more rapidly and an

environment to be brought on line from a cold start more quickly. Finally, when recovering from tape, FICON significantly speeds up the recovery process. [SEIT04]

3) FICON enhances a business' accessibility to data with its higher addressing limitations, enabling more disk volumes to be accessed by a given channel path.
4) FICON protocol has room for growth (16,000 addresses supported today with available growth to 65,000 addresses), which allows a business to be better prepared for internal growth, mergers and acquisitions, or consolidation.
5) The cumulative advantages of FICON present businesses with an opportunity to consolidate the IT infrastructure in terms of the overall footprint of the mainframe and mainframe storage environments, as well as, perhaps, the total number of data centers.
6) An intermix of FICON and FCP allows a business to better utilize IT budget dollars targeted for storage networking infrastructure. z/OS and Linux can now be supported on the same mainframe footprint: FICON and open systems SAN networks can leverage the same directors. Using a common infrastructure for all storage connectivity provides significant opportunity for cost savings to be realized.

These business benefits need to be quantified and balanced against the following costs associated with a migration from ESCON to FICON:

1) FICON cabling: A company will need to either put in a new cabling infrastructure (9/125 micron long wave single mode fiber is recommended) or leverage existing ESCON cabling via the addition of mode conditioning patch (MCP) cables. The use of MCP cables will restrict installation to a 1 Gbps FICON network ; therefore, investing in single mode infrastructure is strongly recommended. Either alternative will be a cost to consider, although long wave single mode better positions an enterprise for the future as we see continued improvements in speed/bandwidth.
2) FICON DASD: Most businesses that migrate to FICON do not do so just for the connectivity advantages - generally, there is a more significant driving factor. Often, the key driver is the older ESCON DASD coming off of lease, or a maintenance contract which is nearing expiration. Typically, these businesses will not go to FICON until then. They will likely invest in new DASD array(s) and will go to FICON for attachment -- to purchase new storage and not move to FICON connectivity would be an example of buying a Porsche 911 turbo and fitting it with normal tires to save money.
3) FICON tape drives/libraries/virtual tape: Similar reasons to DASD above.
4) FICON directors/switches: Larger mainframe environments moving to FICON will want to purchase FICON directors to replace their ESCON directors. Direct attachment of FICON only makes sense for the smallest of FICON environments (i.e., 1 host, 1-2 DASD frames). The FICON directors, however can be a significant cost in the equation.
5) FICON processors (possible): Pre 9672 S/390 G5 mainframes are not FICON capable. An enterprise contemplating moving to FICON which has one of these older mainframes on its data center floor will seriously have to think about the consequences of such a decision. Regardless, unless buying all new mainframes, swapping ESCON channel cards for FICON channel cards is a necessity. Be sure to negotiate this with the mainframe vendor prior to purchase. If purchasing a mainframe with ESCON channel cards, it is strongly recommended that the vendor put in writing that those cards can be changed out at a later date at minimal cost.
6) FICON controllers: If there are existing ESCON DASD arrays or tape drives for which the leases have not yet expired, the following options are available: 1) leave them as ESCON; 2) implement FICON bridge cards (FCV) into the model 5 ESCON directors, and run FICON at 1 Gbps; 3) upgrade their controllers to FICON; or 4) implement protocol converter technology that converts native FICON to ESCON for attachment of legacy devices to FICON mainframe channels. This technology would allow a customer to migrate mainframes to FICON, run native FICON channels into the converter (via a FICON director if desired), and run ESCON channels out of the converter to existing ESCON storage.

FICON planning considerations

FICON and ESCON protocols different greatly. The FICON theoretical maximum data rate is much higher: 200 Mbyte/s in each direction for 2Gbs FICON Express channel. The protocol also allows multiple I/O operations to be handled at the same time on a channel because packets can be interleaved. Of course, if multiple operations are running concurrently, they have to share the available bandwidth.

One of the critical questions is “How many operations can be serviced by a FICON channel”?

To fully understand how to determine this, it is important to understand how a FICON channel actually is implemented. For purposes of this discussion assume a z-Series processor. Please refer to figure 5.

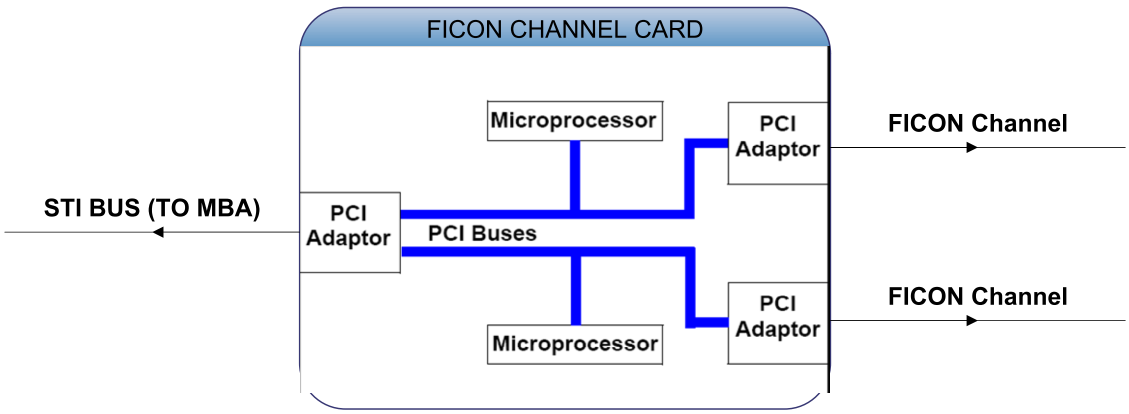

Figure 5 FICON Channel Card

Prior to getting to the FICON channel card, an I/O goes on a journey from the CPU through the “z/box”. The STI (self timed interconnect) provides a path via the MBA (memory bus adapter) from the CPU to the FICON channel card, through a dedicated PCI bus/microprocessor combination and out to one of two FICON (FICON Express) channels (note: on the z990 and z890 this is now four FICON Express 2 channels). Both of the FICON Express interfaces connect to the STI interfaces via a dedicated PCI bus. A dedicated PowerPC microprocessor associated with each interface controls its PCI bus' data transfers and also processes channel command words (CCWs). It is this microprocessor that is actually the FICON channel. What is usually referred to as the FICON channel is technically the FICON channel path.

In February 2002, IBM published a white paper written by Cathy Cronin. [CRON02] This paper introduced the available performance metrics for FICON channels as well as introducing and explaining definitions for FICON channel and bus utilization:

Channel Busy: the measured utilization of the PowerPC microprocessor that manages the fibre channel protocol.

Bus Busy: the measured utilization of the PCI bus over which data is transferred internal to the FICON cards.

I/O activity on a FICON channel card consumes both the PCI bus and the PowerPC microprocessor resources associated with that interface. Dr. H. Pat Artis performed a study on FICON channel and bus busy in 2003 and concluded that there is a strong positive correlation between channel busy and I/O rates, and another strong correlation between bus busy and channel MB/sec. [ARTI03] In other words, small block/short CCW chains consume more microprocessor cycles and large block/long CCW chains consume more PCI bandwidth. RMF reports allow you to easily determine the microprocessor and bus utilization of a FICON channel card. Figure 6 shows an excerpt from an RMF report on FICON channels.

DETAILS FOR ALL CHANNELS

CHANNEL		PATH		UTILIZATION (%)			READ (MB/SEC)		WRITE (MB/SEC)	
ID	TYPE	G	SHR	PART	TOTAL	BUS	PART	TOTAL	PART	TOTAL
7A	FC	1	Y	20.00	30.00	5.00	20.00	30.00	20.00	50.00
7B	FC_SM		Y	15.36	55.86	6.00	15.36	60.00	15.36	60.36
7C	FCV		Y	10.00	30.00	5.00	10.00	50.00	10.00	50.00
7D	FCV M		Y	30.00	45.00	5.00	45.00	50.00	45.00	50.00

Figure 6 Sample RMF report on FICON channels

As FICON channels run multiple operations concurrently, there is a performance penalty due to queuing on the FICON channels. This queuing is called *FICON connect time elongation* and occurs because once an I/O has been started over a particular FICON link; all traffic for that I/O operation will go over this link. Avoiding FICON connect time elongation is important to get the best performance; therefore, FICON channels should be configured such that the link will be no more than 30-50% busy. This may initially seem like a low number, but while the FICON protocol does dynamically route I/O operations based on the observed load, it does not dynamically route individual packets within the I/O. This means that once the path has been selected, queuing is quite likely to occur. It is a like picking a line at the ticket counter; once a choice is made, the customer is stuck.

A different type of elongation can also occur when too many I/Os attempt to share single link. At any point in time each FICON channel can handle only 32 I/O operations (open exchanges) [ARTI03]. This type of elongation happens only when the service time is in the 10 ms range, which is fairly uncommon. In the real world, this is like when the ticket counter lines extend outside the building and around the corner.

Last, but not least, FICON channels can address up to 16K addresses per channel, i.e. typically one (set of) FICON channels would support up to 64 LCUs. This is 16 times as much as for ESCON. Thus the number of channels will no longer be determined by the number of devices, but rather by the throughput and bandwidth requirements.

Focusing merely on the channel utilization percentage and path busy statistics to plan a migration is simply an incomplete methodology. While doing so follows a scientific process, it focuses solely on the mainframe and not the overall system being built. Thus, it is not detailed enough.

Worse yet is utilizing only "rules of thumb" in applying these channel utilization statistics. Such rules of thumb were originally used for 1 Gbps FICON at its inception and were based on no true scientific measurement or performance modeling. Instead, they were based on the 8:1 capabilities of the FICON Bridge (FCV) card that was advertised as being able to take 1 FICON channel from the host and fan out to 8 ESCON channels.
When planning a FICON migration, plan based on what is truly important to stakeholders or internal customers of the IT department. That is, plan SLAs for response times/service times. Having said that, it is unlikely that an SLA for channel utilization percentage will exist. In planning to DASD response time using analytical and simulation modeling tools, what results is a plan for the impact of growth in not only the channel environment but also growth in the DASD environment. In other words, a plan must be developed for the entire system and not just individual components of the system. Such modeling and planning presents a more complete picture, makes a more compelling technology case for moving to FICON to your CTO/CIO, as well as allowing for "what if?" analysis which is key to seeing how DASD storage can be consolidated when moving to FICON.

The sections below will show some detailed real world results of doing such analytical and simulation modeling

Reducing the number of Disk Subsystem images

The following tables show the reduction that can be achieved when migration to new FICON Disk Subsystems. Table 1 below shows the original ESCON configuration:

Subsys ID	I/O Rate (avg)	Usable GB	I/O Intensity	ESCON Channels (17 Mbyte/s)	ESCON Resp (ms)	ESCON Chan Util
ABC7000	2628	8878.2	0.296	24	4.2	26%
ABC8000	2538	8866.1	0.286	32	4.1	17%
ABCE000	2270	4953.1	0.458	16	3.4	27%
ABC6000	1602	4283.8	0.374	16	4.1	24%
ABC1000	1537	1571.9	0.978	16	2.8	14%
ABCF000	1019	3301	0.309	8	3.8	27%
DEFC000	731	2749.4	0.266	16	3.4	12%
ABC9000	424	1618.1	0.262	8	3.4	11%
ABC4000	359	1876.3	0.191	8	3.4	10%
ABCA000	264	826.2	0.320	4	4.6	17%
DEFB000	176	964.7	0.182	16	6	6%
ABCD000	17	289.4	0.059	8	13.4	2%
Overall	1130.4	40178.2	0.028	172	4.7	16%

Table 1 Original ESCON Configuration Response Time Statistics

Overall, there are 172 ESCON channels in this configuration and the largest disk subsystem (ABC8000) has 32 channels. When migrating this configuration to FICON, far fewer channels will be required on each Disk Subsystem. The subsequent sections will discuss how the number of FICON channels can be determined, but for now let's just look at the results of the analysis:

Subsys ID	I/O Rate (avg)	Usable GB	I/O Intensity	ESCON Channels (17 Mbyte/s)	FICON Channels (Express 2Gb)	ESCON Resp (ms)	FICON Resp (ms)	ESCON Chan Util	FICON Chan Util
ABC7000	2628	8878.2	0.296	24	4	4.2	2.3	26%	13%
ABC8000	2538	8866.1	0.286	32	4	4.1	2.2	17%	12%
ABCE000	2270	4953.1	0.458	16	2	3.4	1.8	27%	21%
ABC6000	1602	4283.8	0.374	16	2	4.1	2	24%	16%
ABC1000	1537	1571.9	0.978	16	2	2.8	1.7	14%	12%
ABCF000	1019	3301	0.309	8	2	3.8	2.1	27%	10%
DEFC000	731	2749.4	0.266	16	2	3.4	2.4	12%	9%
ABC9000	424	1618.1	0.262	8	2	3.4	1.7	11%	4%
ABC4000	359	1876.3	0.191	8	2	3.4	1.6	10%	3%
ABCA000	264	826.2	0.320	4	2	4.6	2.5	17%	3%
DEFB000	176	964.7	0.182	16	2	6	3.8	6%	3%
ABCD000	17	289.4	0.059	8	2	13.4	5.5	2%	0%
Overall	1130.4	40178.2	0.028	172	28	4.7	2.5	16%	9%

Table 2 Original Configuration Modeled with FICON

To handle this same workload, only 28 FICON channels would be required, if **each** of the **existing** Disk Subsystems would be **upgraded**. This did not make economic sense however, as many of the Subsystems were reaching the end of their lease periods. So, instead, it was decided to exploit the consolidation potential that FICON offers:

Subsys ID	I/O Rate (avg)	Usable GB	I/O Intensity	FICON Channels (Express 2Gb)	FICON Resp (ms)	FICON Chan Util
subsys 1e€	5350	10800.8	0.495	4	2.2	14%
subsys 89	2962	10484.2	0.283	4	2.2	14%
subsys a7c	2909	9993.7	0.291	4	2.4	15%
subsys bcf	2264	8875.3	0.255	4	2.1	12%
Overall	3371	40154.0	0.331	16	2.2	14%

Table 3 New DASD frames modeled with FICON

The resulting FICON configuration requires only 16 channels to handle all the workload, where we started with 172 ESCON channels. This 11:1 reduction mostly results from the ability to consolidate subsystems because of the improved address capabilities that FICON offers.

Developing planning rules

From the previous sections it can be concluded that, in order to plan the number of FICON channels needed for optimal performance, three factors are important:

1. The number of I/O operations per second,
2. The bandwidth expressed in MByte/s,
3. The number of I/Os that must be concurrently served on each channel.

The first two factors are straightforward and similar to what is known from ESCON configurations, although the FICON link can, of course, carry more traffic. One cannot, however, use a very high percentage of the raw capacity of the FICON link. While the protocol allows up to 200 Mbyte/sec in each direction, what is typically observed in actual implementations are workloads limited to 50 – 100 Mbyte/s peak data rates. While higher numbers can occur in benchmark situations, it is not recommended to plan for them. A good recommendation is to design for no more than 40 Mbyte/s per link during peak production periods. When the data rate becomes higher, significant queuing will occur on the links, the connect times will elongate, and performance degradations will occur in response/service times.

The last factor is, at first, unexpected: when the I/O operations are slower, more channels will be needed to serve them. Let's look at an example to explain how this works:

I/O rate = 1000 I/Os per second, I/O service time (connect + disconnect) = 10 ms = 0.01 second -> 1000 * 0.01 = 10 I/Os concurrently served

This is an illustration of Little's law: When each I/O operation takes 10 ms, and when you serve 1000 I/Os per second, on average, 10 I/Os will be active. In FICON, each active I/O requires an 'open exchange'. Think about this as a session and the number of 'open exchanges' is limited to 32 at any point in time. When this limit of 32 is exceeded, one will start to see elongation on the connect times, which will make the problem worse. Unfortunately there will be even more I/Os!

Dr. H. Pat Artis' recent research has shown that FICON is best managed by measuring with what he refers to as a "Steam Gauge". [ARTI03] The steam gauge measures a computed quantity of open exchanges, which is a logical resource representing an I/O being processed by the channel. For further details please see his outstanding paper on the topic.

So it is very important to make sure that there are enough FICON connections to keep the number of concurrent I/O operations low -- it is recommended there be less than 8 per FICON channel on average. This is equivalent to 1000 I/O operations at 8 ms each or 4000 I/O operations at 2 ms each. This explains why slower I/Os need more channels! Therefore, it is important that the cache hit percentages on FICON controllers are high, even more so than with ESCON. Consider this when using Peer to Peer Remote Copy (PPRC) over long distances, in particular when ESCON links are still being used for PPRC.

To summarize, follow these size rules, based on average values obtained from 15 minute RMF intervals:

- No more than 2500 I/O operations per second per FICON link.
- No more than 25 Mbyte/s continuous for FICON 1 Gbs or 40 Mbyte/s continuous for FICON 2.
- No more than 8 concurrent operations on any FICON channel on average.

It is necessary to determine the peak I/O and MB/s values for each (combination of) subsystems, such that you know what the FICON channels will need to handle. Of course, this will be based on the current workload and anticipated growth.

Table 4 shows how the rules can be applied.

Workload information			Compute..		Number of FICON links required based on ..			Required FICON 2 channels
I/Os per se	Kb/io	MB/s	Conn+Disc	#concurrent	I/Os per se	MB/s	#concurrent	
10000	4000	40	1	10	**4**	1	1.25	4
5000	50,000	250	3	15	2	**6.25**	1.875	7
10000	4000	40	5	50	4	1	**6.25**	7
10000	20,000	200	10	100	4	5	**12.5**	13

Table 4 Planning Rules example

This table shows what drives the numbers of FICON channels: it can be number of operations, the throughput in MB/s or the number of concurrent I/O operations. And often quite different bounds may apply for different shifts in an installation; during the day the number of I/O operations may drive the FICON link requirements and during the night the number of concurrent connections may. Figure 7 shows the I/O rate and MB/s patterns over time for one example installation.

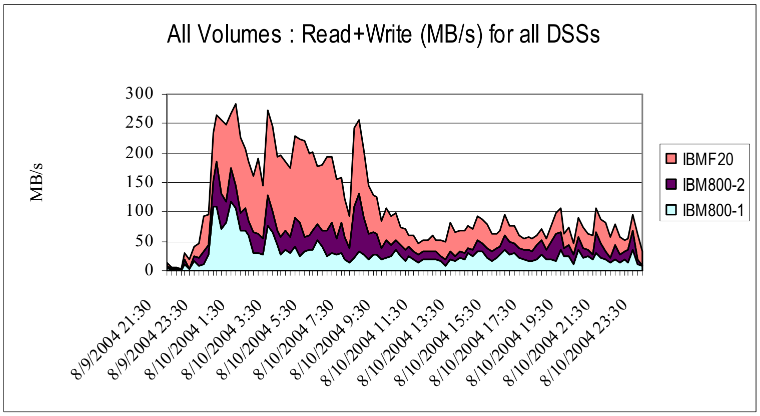

Figure 7.

There is one final step needed before the planning, as described, will really work. Table 4 assumes that the connect and disconnect time and the MB/s values are known before FICON sub subsystems are installed. This is generally not the case, as MB/s is not reported by RMF for ESCON attached subsystems, and the disconnect and connect times for ESCON will change when moving to FICON.

Ideally, one would want to use an analysis tool to determine the MB/s for the subsystems over a longer period. Alternatively, use the SAS reporting based on I/O rate and connect time, assuming that, for example, the average ESCON connect ms yields 12 Mbyte of data transfer.

What to expect

FICON channels will handle the I/O workload faster, in particular during the batch window that tends to use many long transfers where performance was previously ESCON channel bound. This will yield better response time, potentially in more peaked CPU utilization when the workload was previously channel bound. When sizing the processor MIPS to handle the batch peak workload, it is possible to delay processor upgrades because more processor delays are tolerable as a result of faster I/O processing. This, however, does not typically work during the online peak, when the transaction workload is determined by users.

FICON and addressing

Earlier it was discussed that while a FICON link can address very many devices (16 thousand), it can only have 32 operations concurrently active. That is why it was recommended to have no more than 8 operations per link on average. And for most workloads, it turns out the average number of operations is well below 8 -- typically around 2! This means that at any given point in time, only two "disks" (z/OS devices) are performing work.

From a device perspective, it implies that device contention is becoming a rare event; most I/O is handled very quickly (from the cache). This does have interesting sizing consequences for logical volume sizes. As most devices are inactive, it means that consolidating to high capacity logical volumes can easily be done, i.e., migrating from 3390-3 to 3390-27 will not give performance problems. Some PAVs are needed to handle the occasional hot spot, but it does not make sense to assign many more than (maybe) 64 PAVs per FICON link per subsystem, as only 32 addresses (PAV and base) can be active at the same time because of the FICON constraints.

Thus a subsystem with 8 FICON channels might use up to 8 * 64 512 PAVs over 8 LCUs, which is 64 addresses per LCU.

With these ideas in mind, a typical FICON 10 TB building block could look like:

8 channels for 15,000 IOPS
8 LCUs with 64 3390-3 / 32 3390-9 / 32 3390-27 base and 64 alias (1.25 TB each)

More channels may be required when copy services on the configuration.

Cost justifying FICON in a DASD environment

FICON has been available for more than 5 years, yet it has not been widely implemented. Those organizations that have implemented FICON did so primarily because they had a real need to overcome the device address and performance limitations associated with ESCON. In order to fully realize the benefits of FICON, an organization must have CPU, DASD, and directors that all support FICON. In many cases, this can mean a wholesale replacement of the existing environment. Generally, this is a proposition that can be very expensive. Further, most IT departments do not view their existing ESCON performance as a limiting issue, therefore making the case for investment in FICON has been much more difficult – until now.

The industry is coming to a point where older, ESCON capable equipment is being phased out of the support and maintenance matrix, making that equipment very expensive to maintain. One way to reduce ongoing costs is to replace older equipment with new, FICON-capable equipment that comes with a warranty. Once this is done, the move to FICON is inevitable – the hardware is in place and it simply requires some configuration changes.

Today's mainframe environment, in particular the z Series, has strained the limits of ESCON. FICON provides relief from these limitations. FICON allows organizations to use more storage per DASD subsystem without changing to larger volume sizes and to maintain the same or better response times. This translates into requiring fewer subsystems and channels to support a given mainframe DASD environment. In other words, migrating a DASD environment from ESCON to FICON is not just about channel consolidation. It's about a reduction in serial numbers! The ESCON limitations of 1024 addresses in a DASD environment running mod 3 disks has translated on average, into requiring a new DASD array/frame/footprint every 3-4 TB. In the 2000 timeframe, these same ESCON DASD frames had an average advertised capacity of 7-11 TB. It was not possible to fully utilize the frame purchased due to the limitations of ESCON! Today's FICON-capable DASD arrays have an average advertised capacity of 20-80 TB. Much more of that array can be utilized in a mainframe environment due to FICON's relief of the ESCON limitations, as well as the ability to use larger logical volume sizes and PAVs more effectively. Fewer serial numbers translates into real cost savings specific to hardware, floor space, power/cooling, management and software such as advanced copy features. This software is typically licensed per array and the cost varies with the useable storage capacity in the DASD array frame. In other words, DASD vendors tier such software costs based on two factors: the number of arrays and, then, the amount of storage in each array. In general, for "X" TB useable storage total in the environment, it is more expensive in terms of these software licensing costs to have that "X" TB spread over more frames with fewer useable TB/frame than if the "X" TB were consolidated and contained on fewer total frames.

How significant of a consolidation and TCO savings can be achieved? Refer to the analysis/modeling tables from earlier. Obviously every enterprise has a unique environment; however, the author respectfully submits that the results will be similar: TCO savings that justify the expenditure. Being able to consolidate 12 older ESCON attached frames on maintenance onto 4 newer FICON attached frames resulted in significant TCO savings over a 3 year period ($250,000 savings per year in this example).

Planning for Growth

As an illustration of how the analytical and simulation modeling techniques and methods for FICON DASD migrations described above can also help an enterprise plan for growth, please refer to table 5. This is an example taken from a real world case study performed in a small/medium sized mainframe environment. Similar results would be seen for a large environment. The table illustrates how, by using simulation modeling techniques, a channel plan for DASD can be developed from the outset to three years hence.

ESCON Growth	IO Rate	Usable GB	I/O Intensity	Channels	Response	Channel Util
Current	2574	7743.2	0.332	12	2	27%
Mid 2006	3088	9291.8	0.332	12	2.1	33%
Mid 2007	3705	11150.2	0.332	12	2.2	39%
Mid 2008	4446	13380.2	0.332	12	2.3	47%
FICON Growth	**IO Rate**	**Usable GB**	**I/O Intensity**	**Channels**	**Response**	**Channel Util**
Current	2574	7743.2	0.332	4	0.9	9%
Mid 2006	3088	9291.8	0.332	4	0.9	11%
Mid 2007	3705	11150.2	0.332	4	0.9	13%
Mid 2008	4446	13380.2	0.332	4	0.9	15%

Table 5 Growth planning example, 20% YTY DASD growth

Native Tape, VTS, and planning/modeling for FICON performance and cost savings

FICON has allowed for tremendous improvements to be made in performance on the DASD side of mainframe storage. Other components in the mainframe infrastructure have also benefited from FICON technology. While DASD is the focus of performance and the performance improvements/benefits of FICON over ESCON are frequently focused on DASD, a compelling case can be made for tape as well. For space considerations, the tape topic will not be dealt with in as much detail as DASD because the author plans to cover that in a subsequent FICON paper.

Since FICON's inception and adoption in the late 1990s, rapid technological improvements have been made in native tape as well as virtual tape subsystems. The improvements made in DASD during the same time period have led to FICON DASD arrays being able to perform full dumps to disk much more rapidly than in the ESCON era of the early to mid 1990s. The improvements that have been made in native and virtual tape technologies can truly be realized only by migrating to FICON. The remainder of this section discusses these improvements and explains why FICON is necessary to realize the benefits. Summary conclusions from a study done earlier this year are presented: with native tape first and then virtual tape. The implications of I/O rates, MB/sec along the channel paths, load times, mounts, and cartridge capacity will be detailed.

Native Tape and FICON

FICON Express channels allow running 150-170 MB/sec for large data transfers, typically seen with highly sequential tape jobs. Therefore, the metric most crucial to designing the native tape component of FICON infrastructure is that MB/sec number, which can be correlated to the bus busy metric. Aggregating multiple ESCON tape channels onto a single FICON channel can significantly reduce mainframe tape infrastructure by reducing the number of channels, director ports and control unit ports needed for mainframe tape. But, FICON and the new advancements made in tape technology provide another gain: the number of tape drives can be reduced, because each drive handles its work more quickly.

FICON and the new tape technology

The past two-three years have seen significant advances made in enterprise tape technology, both in head to tape transfer rates and cartridge accessibility and capacity. The one most crucial to the FICON discussion is head to tape transfer rates. These new tape drives allow an end user to run at 30+ MB/sec for native head to tape transfer rates, resulting in even higher speeds for uncompressed data. This is several times the capability of the preceding generation of tape drives. Recall from the earlier discussion that the theoretical bandwidth of ESCON is 17-18 MB/sec. That was more than enough to drive an ESCON IBM 3590 or STK 9840 drive at its advertised native speed.

Obviously, FICON is more than capable of running these older 3590 or 9840 drives at their advertised speeds as well. FICON presents several attractive options. FICON significantly reduces the number of channels and ports dedicated to tape by putting several of these tape drives onto one FICON express channel via logical daisy chaining behind a director. Another option that will work well with both IBM's shared control unit type tape drives and STK's direct fibre “1x1” tape drive is to replace several of the older, slower tape drives with the new drive technology that can achieve 30+ MB/sec and still daisy chain multiple tape drives on a FICON channel. An October 2003 STK white paper/study found that it was possible to put 6 of the new 9840B FICON tape drives on one FICON express channel. The same type of study from IBM applies equally to the IBM 3592 tape drives. [STOR04], [HILD04]

In a FICON tape experiment and study conducted, it was found that for a 100% workload requirement requiring 55 3590H ESCON attached tape drives, the same 100% workload requirements could be met using 18 of the new tape technology 3592 FICON drives. For a 97% workload requirement requiring 40 ESCON tape drives, only 14 new FICON tape drives were needed. Implementing the newer tape drives, therefore, required fewer library slots. Also the cartridges for the new drives hold 5x the capacity of the old drives' cartridges, resulting in needing fewer physical cartridges. The installation was able to reduce their costs by consolidating their tape environment, reduce their batch/backup window and improve resource utilization by storing more backup data in a smaller footprint. In subsequent DR testing, they were able to perform restores more rapidly from the new FICON tape drives than they could with their old ESCON tape drives.

Finally, please do not buy these high performance tape drives and run them on ESCON channels. The level of performance paid for cannot be achieved and money is wasted. This is analogous to a US driver in Ohio who decides to purchase a Lamborghini because it's the fastest car out there today; yet when he drives it in the US, he can only drive it at 65 MPH. He leaves a lot of performance that he paid for untapped, as does the end user who purchases 9840B/C or 3592 tape technology and runs it on ESCON channels.

A brief word on virtual tape

While FICON allows the end users of today's tape drives to run them at their rated native speeds, virtual tape systems can also benefit from FICON. How? FICON allows a virtual tape system to accept and deliver more tape workload while using fewer channels. In general (specifics vary by vendor, either IBM or STK), the virtual tape system will have a maximum number of ESCON channels for input. Also, the same virtual tape system will have a maximum number of FICON channels for input; however, the number of FICON channels is less. Using FICON for virtual tape rather than ESCON reduces the number of control units, back end tape drives[1] and channel paths required. But when looking at the bandwidth coming in on those less numerous FICON channels, the MB/sec coming into the FICON virtual tape system is typically 3x what the number was for ESCON. Yet looking solely at this can put the end user in a quandary.

To deal with that bottleneck on the front end by upgrading to a FICON capable virtual tape system and doing nothing with the existing native tape drives on the back end of the virtual tape system may merely shift the bottleneck to the back end. This depends on the hit ratios that are achieved for virtual mounts. When the old ESCON tape drives on the back end of the virtual tape system are the bottleneck, deal with this by adding more of the same tape drives. However, a virtual tape system will have a maximum number of supported native tape drives per system and taking this route will lead to adding more virtual tape systems. Since the reason for the migration to FICON virtual tape was to reduce the number of virtual tape systems, obviously this is not the way to go.
How can this problem be solved? It is elementary. When moving to FICON virtual tape, upgrade to the new FICON native tape technology for the back end of the virtual tape systems. In the study cited earlier for native tape, virtual tape was also analyzed. The conclusion reached by the study showed that for the environment in question, by migrating from ESCON IBM VTS subsystems, they were able to consolidate from 3 VTS subsystems running 12 IBM 3590E ESCON tape drives per VTS, to 2 FICON IBM VTS subsystems running 8 IBM 3592 tape drives.

Tape: Summary comments
Table 6 is a summary of the results of another real world FICON modeling summary that focused on disk and tape. Table 7 shows how moving to new FICON tape drive technology, with the associated cartridges leads to a significant reduction in the number of cartridges required. This study was done for a fairly small size environment to show how migrating both disk and tape to FICON is synergistic to the improvements for the enterprise. It also makes a good point that for an organization without the resources/ability to migrate both disk and tape to FICON in one project, migrating the DASD first is really a pre-requisite to obtain the benefits of FICON tape.

Dump Throughput Comparison	Concurrent dumps	Mins per dump	GB per dump	Carts per dump	MB/sec per Dump	Total MB/sec	Dumps /cart	Carts required	Total dump mins	Total Dump Hours	Hours Saved	Observations
Existing environment	12	8	2.23	1	4.5	54.3	8	66	463.1	7.7	0.0	Limited by Tape subsystem
FICON disk, ESCON tape	12	6	2.23	1	6.8	108.6	8	66	390.8	6.5	1.2	Limited by Tape subsystem
FICON tape, ESCON disk	12	3	2.23	1	12	144	131	4	195.4	3.3	3.3	Limited by Disk subsystem
FICON disk and tape	12	2	2.23	1	20.2	242.5	131	4	130.3	2.2	1.1	Limited by Tape subsystem

Table 6 Dump Throughput Comparisons

Overall Statistics	Current	Projected
Total Specific Mounts	10125	10125
Total Scratch Mounts	3012	71
Total Mounts	13137	10196
Total GB Read	3906	3906
Total GB Written	19917	19917
Total GB	23822	23822
Specific Mounts %	77.1	99.3
Scratch Mounts %	22.9	0.7
GB Read %	16.4	16.4
GB Written %	83.6	83.6
Avg. Vol Size MB	11655.7	27760
Total # of Cartridges	7918	195

Table 7 Overall Statistics

Some final thoughts/summary/conclusions

FICON will continue to be a "hot topic" for some time to come. More and more enterprises are already well familiar with the underlying fibre channel technology through their experience with open systems storage area networks (SAN). FICON and open systems SAN intermix (FICON/FCP intermix), i.e., running one common storage network for both mainframe and open systems, is becoming more and more attractive to enterprises. It makes a lot of financial sense in today's era of cash strapped CIO budgets and, as the "kinks" are worked out by those of us on the "bleeding edge", it will make more and more technical sense as well. It is the author's belief that the world is becoming less and less operating systems centric and more and more data centric. Open source Linux is being run on mainframes in ever increasing frequency. We are heading towards a world that is going to be run on a common systems area network spanning open systems and mainframes. Because FICON is going to be around for a while, don't continue to spend precious budget dollars on ESCON director maintenance, particularly if you're running mainframes, DASD, tape, and virtual tape designed for FICON's performance.

Bibliography/References

1) [ELIO92] J.C. Elliot & M.W. Sachs. *The IBM Enterprise Systems Connection (ESCON) Architecture.* IBM Journal of Research and Development Volume 36, #4 (1992).

2) [TROW02] Ken Trowell, Bill White et. al. *FICON Native Implementation and Reference Guide.* IBM Redbooks (2002). SG24-6266-01. ISBN 0738427365.

3) [HOUT91] Gilbert E. Houtekamer & H. Pat Artis. *MVS I/O Subsystems Configuration Management and Performance Analysis.* McGraw-Hill (1991). ISBN 0-07-002553-3.

4) [CRON02] Catherine Cronin. *FICON and FICON Express Channel Performance, Version 1.0 (February 2002).* www.ibm.com

5) [CRON03] Catherine Cronin and Richard Basener. *FICON and FICON Express Channel Performance, Version 2.0 (November 2003).* www.ibm.com

6) [ARTI03] Dr. H. Pat Artis. *Understanding FICON Channel Path Metrics.* CMG Proceedings 2003, pp 617-625. www.perfassoc.com

7) [SEIT04] Sean Seitz & Steve Guendert. *FICON Migration: Save Money and Improve Performance In Your Mainframe Infrastructure.* Z Journal Oct/Nov 2004, pp 43-47.

8) [ARTI01] Dr. H. Pat Artis & Robert Ross. *FICON I/O Measurement: Everything You Know Is Wrong!* 2001. www.perfassoc.com

9) [GUNT98] Neil J. Gunther. *The Practical Performance Analyst.* McGraw-Hill (1998). ISBN 0-07-912946-3.

10) [JAIN91] Dr. Raj Jain. *The Art of Computer Systems Performance Analysis.* J. Wiley & Sons (1991). ISBN 0-471-50336-3.

11) [HOUT05] Gilbert E. Houtekamer & Steve Guendert. *Sizing Your ESCON to FICON Conversion.* Z Journal. June 2005. pp 22-30.

12) [CRON05] Catherine Cronin. *FICON Express2 Channel Performance Version 1.0 (April 2005).* www.ibm.com

13) [MERR84] Dr. H.W. "Barry" Merrill. *Merrill's Expanded Guide to Computer Performance Evaluation Using the SAS System.* SAS Institute (1984). ISBN 0-917382-54-4. www.mxg.com

14) [SCHU04] Greg Schulz. *Resilient Storage Networks.* Elsevier (2004). ISBN 1-55558-311-3

15) [STOR04] StorageTek Corporation. *9840C FICON Tape Drive Performance.* 2004. www.storagetek.com

16) [HILD04] Justin Hildebrandt. *IBM TotalStorage Enterprise Tape Controller Model J70 Performance.* 2004. IBM Corporation. www.ibm.com

17) [AURA03] Tom Aurand. *How Many Channels? The Effects of FICON on Channel and Device Performance-Parts 1 and 2.* Technical Support. August-September 2003.

18) [NEVI05] Iain Neville, Bill White & Hans-Peter Eckam. *FICON Implementation Guide.* IBM Redbooks. 2005. SG24-6497