Z/JOURNAL

Sizing Your ESCON to FICON Conversion
BY GILBERT HOUTEKAMER, PH.D., & STEVE GUENDERT

Sizing Your ESCON to FICON Conversion

By Gilbert Houtekamer, Ph.D., & Steve Guendert

When you invest in new mainframe disk or tape storage equipment, you're most likely looking at FICON equipment. Most modern disk and tape subsystems are so fast they need FICON channels to exploit the performance potential they offer. FICON has been around for five years and all vendors are in at least their second iteration of the hardware, software, and firmware. This means you can >

expect a smooth, fast ride when investing in FICON.

The FICON protocol also has many inherent advantages over the ESCON protocol previously used. FICON supports distances well over 10 km with hardly any speed degradation or data droop. As FICON supports higher effective data rates per link, you'll need fewer links, so it will be more cost-effective to buy a FICON solution. Of course, with a substantial investment in ESCON technology for your installation, you'll need to make and cost-justify the transition.

This article examines how you can determine the number of FICON channels that will be required, based on the workload in our example installation. We won't discuss the operational challenges of performing this upgrade in your shop.

First, let's recap some ESCON planning observations.

ESCON Configuration Planning

As you're likely used to planning ESCON configurations, it's useful to review some of the ESCON configuration rules and explicitly consider how they influenced the design of current channel and ESCON director configurations. The maximum data rate for each ESCON channel is about 18 Mbyte/second during a data transfer peak. While this data rate can be achieved for an individual job such as a backup or restore, it's clearly impossible to reach such data rates during your typical peak hour. Most installations will plan for peak channel utilization of less than 60 percent, resulting in an absolute maximum effective data rate of about 10 Mbyte/s during the peak period. With higher channel utilizations, pending and other queuing delays would slow down the system too much.

There's also a limit to the number of I/O operations an ESCON channel can handle, as the channel is "connected" during interpretation of the I/O action. The connect time for a simple I/O such as a 4Kb read-hit can range from 0.5 to 1 ms, mostly depending on the disk subsystem year of manufacture. For a typical workload, with larger blocks and cache misses, too, this means the number of I/O operations per channel will be well below 1,000 I/O operations per second.

Finally, ESCON has inherent addressing limitations in the protocol. ESCON can address up to 1,024 different devices per channel. For a disk subsystem, this means a set of channels is required for every four Logical Control Units (LCUs) with 256 addresses each. It also means only 3TB can be addressed when using 1,024 volumes with 3390 Model 3 format. Modern controllers support 16 or more LCUs, so the number of required ESCON channels will grow proportionally.

FICON Considerations

The FICON protocol differs greatly from the ESCON protocol. The theoretical maximum data rate is much higher at 200 Mbyte/s in each direction for the 2Gbs FICON Express channel. The protocol also allows multiple I/O operations to be handled simultaneously on a channel as packets can be interleaved. Of course, when multiple operations are running concurrently, they share the bandwidth.

A critical question to ask is, "How many operations can a FICON channel service?" To determine this, you first need to understand how a FICON channel actually is implemented. (For purposes of this discussion, let's assume a zSeries processor as shown in Figure 1.)

Before getting to the FICON channel card, an I/O journeys from the CPU through "zbox." The Self-Timed Interconnect (STI) provides a path via the Memory Bus Adapter (MBA) from the CPU to the FICON channel card through a dedicated Peripheral Component Interconnect (PCI) bus/microprocessor combination and out to one of two FICON (FICON Express) channels. (On the z990 and z890, this is now four FICON Express 2 channels.) Both 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). This microprocessor is actually the FICON channel; what's 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 that introduced the available performance metrics for FICON channels. This paper also introduced and explained the following definitions for FICON channel and bus utilization:

- Channel busy: the measured utilization of the PowerPC microprocessor that manages the Fibre Channel Protocol (FCP).
- 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 there's a strong positive correlation between channel busy and I/O rates, and another strong correlation between bus busy and channel MB/sec. In other words, small block/short CCW chains consume more microprocessor cycles and large block/long CCW
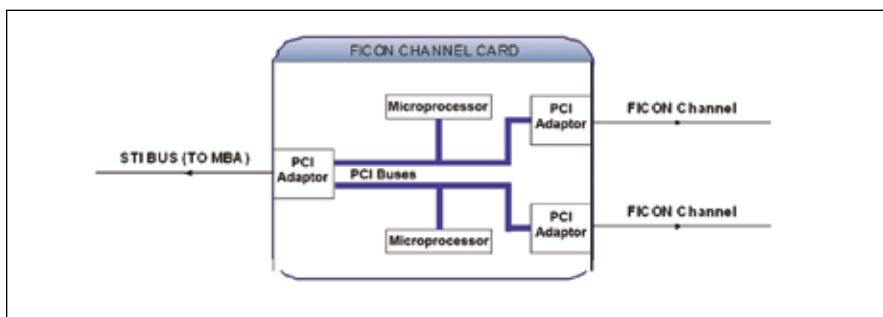

Figure 1: FICON Channel Card

```
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 2: Excerpt From an RMF Report on FICON Channels

chains consume more PCI bandwidth. Resource Management Facility (RMF) reports let you easily determine the microprocessor and bus utilization of a FICON channel card. Figure 2 shows an excerpt from an RMF report on FICON channels.

FICON interface cards have two performance-determining factors: the overhead to start an I/O operation and the bus bandwidth to sustain a transfer. FICON channel sizing therefore needs to consider the activity (I/Os per second) and bandwidth (MB/s).

The high bandwidth of FICON links means single-stream data rates of more than 100 Mbyte/s are possible (during batch processing and backups, for example). Under normal loads, the aggregate throughput on any FICON channel will rarely exceed 50 Mbyte/s.

As FICON channels run multiple operations concurrently, there's a performance penalty because of queuing on the FICON channels. This queuing

is called FICON connect time elongation. It 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; so you should configure FICON channels so the link will be no more than 30 to 50 percent busy. This may initially seem like a low number, but while the FICON protocol does route I/O operations dynamically based on the observed load, it doesn't route individual packets within the I/O dynamically. This means that, once the path has been selected, queuing is quite likely to occur. It's like picking a line at the ticket counter; once you make your choice, you're stuck.

A different type of elongation can also occur when too many I/Os attempt to share a single link; at any time, each FICON channel can handle only 32 I/O operations (open exchanges). This type of elongation occurs only when the ser-

vice time is in the 10 ms range, which is fairly uncommon. In the real world, this is like when the ticket counter lines no longer fit in the building; you have to wait outside.

The following section shows sizing rules based on these three criteria.

Finally, FICON channels can address up to 16KB 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. The number of channels is no longer determined by the number of devices, but solely by throughput and bandwidth requirements.

Reducing the Number of Disk Subsystem Images

Figures 3 and 4 show the reduction possible when migrating to new FICON disk subsystems; Figure 3 shows the original ESCON configuration.

There are 172 ESCON channels in this configuration; the largest disk subsystem has 32 channels. When migrating this configuration to FICON, far fewer channels will be required on each disk subsystem. The subsequent sections discuss how the number of FICON channels can be determined. For now, just look at the results of the analysis in Figure 4.

To handle this same workload, only 28 FICON channels are required if each disk subsystem is upgraded. This doesn't make economic sense, however, as many of the subsystems were reaching the end of their lease periods. The better decision was to exploit the consolidation potential FICON offers (see Figure 5). The resulting FICON configuration requires only 16 channels to handle all the workload; it started with 172 FICON channels. This is a 10:1 reduction that mostly results from the ability to consolidate subsystems because of the improved address capabilities FICON offers.

Developing Planning Rules

To plan the number of FICON channels you need, three important factors to consider are:

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

The first two factors are straightforward and similar to what you know from ESCON configurations, although the FICON link can, of course, carry

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%**

Figure 3: Original ESCON Configuration

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%**

Figure 4: Analysis Results

Subsys ID	I/O Rate (avg)	Usable GB	I/O Intensity	FICON Channels (Express 2Gb)	FICON Resp (ms)	FICON Chan Util
subsys 1e6	5350	10800.8	0.495	4	2.2	14%
subsys 89	2962	10484.2	0.283	4	2.2	14%
subsys a7d	2909	9993.7	0.291	4	2.4	15%
subsys bcf14	2264	8875.3	0.255	4	2.1	12%
Overall	**3371**	**40154.0**	**0.331**	**16**	**2.2**	**14%**

Figure 5: Exploiting FICON's Consolidation Potential

more traffic. You can't, however, use a very high percentage of the raw capacity of the FICON link. While the protocol allows up to 200 Mbyte/s in each direction, there are actual implementations with workloads limited to 50 to 100 Mbyte/s peak data rates. While higher numbers can occur in benchmark situations, you shouldn't plan for them. Plan for no more than 40 Mbyte/s per link during your peak production periods. When the data rate becomes higher, significant queuing will occur on the links, and you'll see connect times elongate and performance degrade.

The last factor is initially unexpected. When the I/O operations are slower, you'll need more channels to serve them. Consider an example to see 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 Little's Law. When each I/O operation takes 10 ms, and when you serve 1,000 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 time. When you exceed this limit, you'll start to see elongation on the connect times, which exacerbates the problem and results in more I/Os.

So, it's vital to ensure you have enough FICON connections to keep the number of concurrent I/O operations low. Less than eight per FICON channel, on average, is recommended. That's equivalent to 1,000 I/O operations at 8 ms each, or 4,000 I/O operations at 2 ms each. This is why slower I/Os need more channels and it's important the cache hit percentages on FICON controllers are high, even more so than with ESCON. You should also 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, you should adopt these size rules, based on average values obtained from 15-minute RMF intervals:

- No more than 2,500 I/O operations per second per FICON link
- No more than 25 Mbyte/s continuous for FICON 1GBs or 40 Mbyte/s continuous for FICON 2
- No more than eight concurrent operations on any FICON channel, on average.

You'll need to determine the peak I/O and MB/s values for each (combination of) subsystem, so you know what the FICON channels will need to handle. Of course, this will be based on your current workload. You'll want to allow for growth as well.
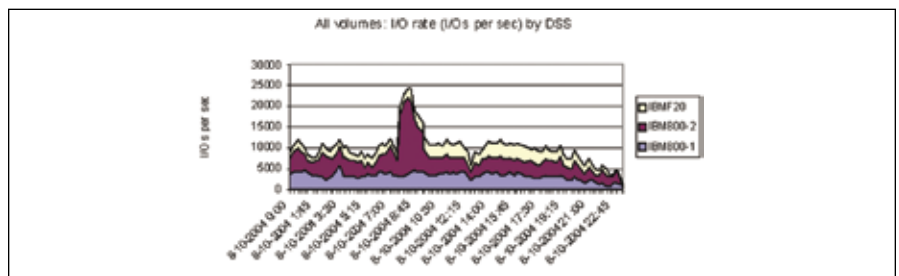
Figure 6 shows how the rules can be

Figure 6: I/O Rate Over Time

applied. This example shows that the number of operations, the throughput in MB/s, or the number of concurrent I/O operations drive the number of FICON channels. Often, different bounds may apply during different shifts in your installation. During the day, the number of I/O operations may drive FICON link requirements; while at night, the number of concurrent connections may drive these requirements. Figure 7 shows the I/O rate and MB/s patterns over time for one installation.

You need to take one final step before the planning, as described, will really work. Figure 7 assumes the connect and disconnect time and the MB/s values are known before FICON subsystems are installed. This is generally not the case, as MB/s isn't reported by RMF for ESCON-attached subsystems and the disconnect and connect times for ESCON will change when moving to FICON.

Ideally, you'd use a commercial product to determine the MB/s for your subsystems over a longer period. Alternatively, you could use SAS-based reporting of I/O rate and connect time, assuming that, for example, the average ESCON connect ms yields 12 Mbyte of data transfer. For ESCON vs. FICON connect time, you can also model with a commercial product or use a rough estimate that connect time halves. Either way, you can estimate the metrics from Figure 7 for your sizing. The next section presents some examples.

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 result in better response time, potentially in more peaked CPU utilization when your workload was previously channel bound. When you size your processor MIPS to handle the batch peak workload, you may even be able to delay processor upgrades because you can tolerate more processor delays due to faster I/O processing. This typically doesn't work during the online peak, as your users determine the transaction workload.

FICON and Addressing

While a FICON link can address up to 16,000 devices, it can have only 32 operations concurrently active. That's why you should have no more than eight operations per link on average. And for most workloads, it turns out the average number of operations is well below eight—typically around two! 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/Os are 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; migrating from 3390-3 to 3390-27 will not give performance problems. You do need some PAVs to handle the occasional hot spot, but it doesn't make sense to assign many more than maybe 64 PAVs per FICON link per subsystem, as only 32
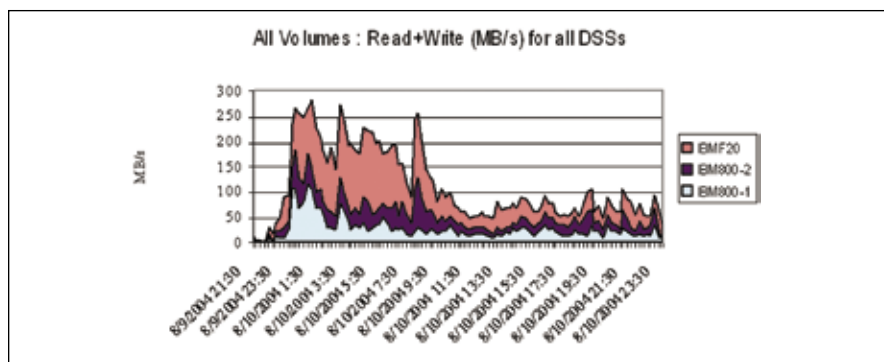

Figure 7: MB/s Pattern Over Time

addresses (PAV and base) can be active at the same time anyhow because of the FICON constraints.

So a subsystem with eight FICON channels might use up to 8 * 64 512 PAVs over eight LCUs, (64 addresses per LCU).

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

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

More channels may be required if your configuration uses copy services.

Configuring Native Tape and VTS

FICON provides valuable performance improvements on the DASD side of mainframe storage. FICON technology has also benefited other mainframe infrastructure components. Since FICON's inception and adoption in the late '90s, rapid technological improvements have occurred in native tape and virtual tape subsystems. The DASD improvements made during the same time led to FICON DASD arrays being able to perform full dumps to disk much more rapidly than in the ESCON era. The improvements made in native and virtual tape can be truly realized only by migrating to FICON. The remainder of this section discusses these improvements and explains why FICON is necessary to realize the benefits. You'll also benefit from findings of a study we conducted and tips on native tape, virtual tape, I/O rates, MB/sec along the channel paths, load times, mounts, and cartridge capacity.

Native Tape and FICON

FICON Express channels let you run 150 to 170MB/sec for large data transfers, typically seen with highly sequential tape jobs. So the metric most crucial to designing the native tape component of your FICON infrastructure is the 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 your mainframe tape infrastructure by reducing the number of channels, director ports, and control unit ports you need for mainframe tape.

But FICON and the new advancements made in tape technology deliver another gain: You can consolidate the number of tape drives, as each drive handles its work faster.

FICON and the New Tape Technology

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

Obviously, FICON can run these older 3590 or 9840 drives at their advertised speeds, too. FICON offers two attractive options:

- You can run FICON and significantly reduce your 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.
- You can replace several older, slower tape drives with new drive technology that can achieve more than 30 MB/sec and still daisy chain multiple tape drives on a FICON channel.

The second option works well with both IBM's shared control unit type tape drives and STK's direct fibre "1x1" tape drive. An October 2003 STK white paper/study found it was possible to put six of the new 9840B FICON tape drives on one FICON express channel. A similar study from IBM applies equally to the IBM 3592 tape drives.

A FICON tape study conducted with a performance analysis tool found that:

- For a 100 percent workload requirement needing 55 3590H ESCON-attached tape drives, the same 100 percent workload requirements could be met using 18 new tape technology 3592 FICON drives.
- For a 97 percent workload requirement needing 40 ESCON tape drives, only 14 new FICON tape drives were needed.

The newer tape drives also required fewer library slots. The cartridges for the new drives hold five times the capacity of the old drives' cartridges, resulting in the need for fewer physical cartridges. The installation was able to:

- Reduce its costs by consolidating its tape environment
- Reduce its batch/backup window
- Improve resource utilization by storing more backup data in a smaller footprint.

In subsequent Disaster Recovery (DR) testing, they also were able to perform restores faster from the new FICON tape drives than they could with their old ESCON tape drives.

Finally, please don't buy these high-performance tape drives and run them on ESCON channels. You won't get what you're paying for. An appropriate analogy is the one with the U.S. driver who decides to purchase a Lamborghini because it's the fastest car available, yet when he drives it in the U.S., he can travel only 65 miles per hour. He leaves a lot of performance he paid for untapped,

as does the user who purchases 9840B/C or 3592 tape technology and runs it on ESCON channels.

Virtual Tape

While FICON lets users of today's tape drives run them at their rated native speeds, virtual tape systems can also benefit from FICON. FICON allows a virtual tape system to accept and deliver more tape workload while using fewer channels. Specifics vary by vendor, but generally 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, but the number of FICON channels is less. When you look at the bandwidth coming in on fewer FICON channels, the MB/sec coming into the FICON virtual tape system is typically three times what the number was for ESCON. However, looking solely at this can put the user in a quandary.

Dealing with that bottleneck on the front-end by upgrading to a FICON-capable virtual tape system while 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 you'll be able to achieve for virtual mounts.

When the old ESCON tape drives on the back-end of your virtual tape system are the bottleneck, you can 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 you to add virtual tape systems. Since the reason you migrated to FICON virtual tape was to reduce the number of virtual tape systems, you can see this isn't the way to go.

How can you solve this problem? It's elementary. When you move to FICON virtual tape, upgrade to the new FICON native tape technology discussed earlier for the back-end of your virtual tape systems. In the tape study cited earlier, the installation in question consolidated from three VTS subsystems running 12 IBM 3590E ESCON tape drives per VTS to two FICON IBM VTS subsystems running eight IBM 3592 tape drives—all by migrating from ESCON IBM VTS subsystems.

Conclusions

FICON technology delivers signifi-

cant performance and configuration benefits, especially for modern disk and tape subsystems that require FICON to exploit their capabilities. By following the sizing rules presented here, you can easily determine the number of FICON links required for your workload. That can help you reduce the number of disk and tape subsystems you need. This reduction in images provides a significant cost savings, as was shown with several actual configuration examples. **z**

References

- "9840C FICON Tape Drive Performance," StorageTek Corp., 2004
- *FICON and FICON Express Channel Performance Version 2.0*, Catherine Cronin, IBM Corp., 2003
- "Understanding FICON Channel Path Metrics," Dr. H. Pat Artis, 2003
- "IBM Total Storage Enterprise Tape Controller 3592 Model J70 Performance," Justin Hildebrandt, IBM Corp., 2004

About the Authors

Gilbert Houtekamer, Ph.D., is president of IntelliMagic BV, a Dutch company specialized in storage performance modeling and sizing software. He has been working in the field of I/O subsystem performance for 20 years, and has published numerous papers on this topic at CMG and other conferences.
e-Mail: gilbert.houtekamer@intellimagic.nl
Website: www.intellimagic.com

Steve Guendert is McDATA Technical Solutions Group's FICON principal consultant and is the director of their FICON practice. He is a PMI certified Project Management Professional (PMP), and he has an M.B.A. and a master's degree in Management Information Systems. He is a well-regarded FICON consultant, and an industry expert on the subject. He has written numerous articles and white papers on FICON, and lectures and consults regularly on FICON technology as well as business case justifications for migrating from ESCON to FICON. He is currently a Ph.D. candidate and is writing his doctoral dissertation on enterprise I/O subsystems.
e-Mail: stephen.guendert@mcdata.com
Website: www.mcdata.com