



Putting Your Data in the Right Place — Using Value to Determine the Right IBM Storage

Analyst: David Reine

Management Summary

Is there one move that your CEO can execute that can guarantee enterprise success? Is there one single plan that his staff can implement to ensure profitability to the business unit? The obvious answer is "no", otherwise we would all run to the bookstore to purchase this guide to corporate success. **Success in business, like in sports, is a game, made up of dozens and perhaps hundreds of individual decisions within each contest.**

In baseball, for example, a manager knows that he cannot win the game in the first inning. However, he can put himself into position to win by using his assets wisely. He must know the strengths, or the value, of each player. He must know who can hit the home run and who can run. A shrewd manager does not hesitate to have one player sacrifice himself in the first inning in order to advance a teammate and, therefore, give his team an immediate advantage. That manager knows that he has to plan the disposition of his pitching assets for at least nine innings. He understands which pitcher he will need for six or seven innings, and who he has to come in for relief. He must know the value of each. When it comes to the implementation of strategy, there is no game older than Chess. With origins dating back 1500 years, Chess is the closest game that we have to enterprise competition. To be successful in chess, you need to understand the value of each piece and each position on the board. Recently, computers have been "taught" to play chess well enough to compete with the great chess champions. These systems have the processing power to think ahead, to consider the options that each move, and each counter-move, provides. **The chess master understands that he must sacrifice some of his assets in order to achieve the prime objective, to capture the king, which is akin to the enterprise strategy.**

In business, the greatest electronic assets of the enterprise lie in the data that they possess. The ability to access and process that data rapidly and with cost efficiency could be the difference between success and failure, between profit and loss. IT departments do not make the decision to migrate this asset from a reliable, legacy platform recklessly, but inevitably, it will be made. Of course, you will do your due diligence. You will identify the vendor that can recognize your needs and deliver products that will provide an immediate return on investment, lower maintenance costs, and provide price/performance at multiple tiers. Because **your enterprise data is continually evolving and has different value depending on the type of data and the phase of its lifecycle**, you need to develop a strategy that contains low-cost storage alternatives along with a high-performance storage solution. Several companies (e.g. EMC, HP, and IBM) can deliver storage solutions to migrate data from older architectures to a consolidated, multi-tiered environment. If you have determined that IBM has the solutions to be your storage vendor, whether on financial or technical grounds, then you must make additional decisions to determine the type of storage platforms that are appropriate for you. Please read on to understand the differences in the IBM storage solutions and how they can contribute to your profitability.

IN THIS ISSUE

> To Migrate or not to Migrate	2
> Migration Strategy.....	2
> Storage Platforms	3
> Summary.....	5

To Migrate or Not to Migrate

It does not require a data processing savant to recognize the signs of a pending data access disaster. The user's complaints become more persistent: availability, response time and throughput are inadequate, there is no space for file growth, and the storage management costs are too high. The IT organization has already recognized that its 5 year-old disk array is obsolete. From a technical standpoint, the bus access is limited, the rotational speed of the drives is inadequate, and the size of the drives (capacity) is too small for the content requirements mandated by government and industry standards. Forget about the management directive for consolidation; you are simply trying to keep your head above the waterline. In addition, maintenance costs are prohibitive, spare parts are hard to find, and the storage management software available when you purchased this albatross does not meet the needs of the 21st century. The reduced maintenance cost associated with a new storage platform alone will produce an immediate return on investment.

However, mission-critical applications are working and the data center is a 7x24 operation connected to the Internet. *How can the IT organization shut down the storage for periods long enough to successfully migrate all of the data and files off the legacy platforms? How long will it take to architect a network of storage platforms designed to support the different requirements inherent in the enterprise storage structures?* The CIO has done a Data Lifecycle Management study. We understand that every database, every file carries with it a value¹. Unfortunately, that value changes from time to time. Data that requires instant access today may not be needed again for 90 days. Even if you had the data movement policies defined, *how long would it take to migrate from a single, simple array to a multi-tiered system?* Furthermore, *how many people would it require?* Perception says that you cannot migrate economically, and we all know that perception is reality. However, is it?

¹ See **The Clipper Group Explorer** dated August 29, 2002, entitled *Tiered Storage Classes Save Money – Getting the Most Out of Your Storage Infrastructure*, at <http://www.clipper.com/research/TCG2002030.pdf>

Exhibit 1. –

Migration Program Objectives

- Evaluate existing storage infrastructure
- Enable scalability beyond today's limits
- Replace old technology with new
- Address need for improved performance
- Provide improved storage infrastructure
- Provide required data migration services
- Enable partner community to deliver migration services
- Improve system resiliency
- Implement flexibility of virtualization
- Lower total cost of ownership

IBM has taken a lot, if not all, of that into consideration, with the announcement of a new market program called *Piper*.² Piper is a data migration service based upon an automated appliance implemented as a tool for the data center. Any authorized IBM service group or IBM Partner trained for data migration and equipped with the cost analysis tools can deliver this service. **Piper enables a transparent migration off a point-to-point or legacy SAN storage platform, maintaining maximum availability of storage resources.** It also removes concern over migration objectives (see Exhibit 1, above) from the CIO's pain list and allows him/her to concentrate on the more important enterprise data lifecycle issues.

Migration Strategy

Once the CIO has made the migration commitment, he needs to adopt a specific strategy. Which hardware and software elements are available to adapt to your environment in order to increase capacity for consolidation, minimize your total cost of ownership, and improve system reliability? Which elements will maximize your return on investment? Will you implement a point to

² See **The Clipper Group Captain's Log** dated December 8, 2003, entitled *IBM Offers New Storage Migration Tool – A Gauntlet for EMC*, at <http://www.clipper.com/research/TCG2003065.pdf>

point architecture or construct a Storage Area Network (SAN) in order to maximize the utilization of the disk resources?

In terms of hardware, the IT department must evaluate the functional advantages that each storage platform provides versus the cost associated to it. *Is the Enterprise Storage Solution (Shark) the correct fit? Should you be considering the FAStT line? Within FAStT, which model is appropriate? Should you look at Fibre Channel or ATA drives? Moreover, should you be looking at a combination of these platforms in order to tailor the costs of the storage to the value of the data?* Fears over network disruption and downtime have already caused you to postpone this migration, **enterprise growth could rapidly exceed the capabilities inherent in the existing architecture. You cannot wait forever.**

Complementary to the selection of hardware, the data center must also select a storage-management application set that meets the strategic objectives of the enterprise. *Is a virtualized storage system in your plan? Are you employing a SAN or a point-to-point architecture? Do you need the local management capabilities of a unified SAN File System? Do you need the capability of the SAN Volume Controller to deploy and administer your storage assets? What is your backup strategy? How can consolidation minimize your software licensing charges?*

Let us look at these various components

IBM Storage Servers	F600	F600T	F900	E.S.S.
F.C. Capacity				
Number of Drives	56	112	224	384
Maximum TB	8.2	16.4	32.8	55.9
SATA Capacity				
Number of Drives	112	112	224	--
Maximum TB	28	28	56	--
I/O Channels				
F.C.	4	4	8	16
SCSI/ESCON	--	--	--	32
Throughput-MB/s	400	768	795	N/A
Cache	512MB	2GB	2GB	64GB
RAID Levels	*	*	*	5,10
Max Partitions	16	64	64	N/A

* = 0,1,3,5,10

Source: IBM

Exhibit 2 – Disk Array Specifications

and try to identify the functionality available from each element along with the enterprise problems that you are attempting to resolve.

Storage Platforms

Whether the data center is migrating or consolidating their storage, the determination of which storage platform(s) to select will depend upon how well they improve performance, increase reliability, enhance productivity, and lower costs. The application servers that are in place and the storage elements that they support will also drive the choice. In order to determine the correct platform, let us look at both the FAStT Family, with the *Model 600 and 900*, and the *Enterprise Storage Server Model 800*. We can see how they compare with regard to capacity, performance, and functionality.

FAStT

The IBM *TotalStorage FAStT600 Storage Server* addresses the needs of a department or small enterprise as an affordable, scalable server. We can look at the metrics in Exhibit 2 (below) and see that there is significant scalability to support migration or consolidation within the FAStT family. Starting with the FAStT600, you can configure the array with up to 8.2TB across 56 Fibre Channel drives. Upgrading to the *Turbo³* capabilities, you can expand that to 112 drives with 16.4TB. Upgrading to the *FAStT900* expands those figures to 224 drives with 32.8TB of capacity.

The drives come in a variety of sizes and speeds, from 36.4GB to 146.8GB at 10K RPM and 18.2GB to 73.4GB at 15K RPM. This provides the flexibility needed to assign different value files to drives with different costs and performance. **More significantly, the FAStT is also available with up to 112 Serial ATA drives, each with a capacity of 250GB.** This not only expands the total capacity of the FAStT600 to 28GB, but it does so with lower cost devices. With a list price of under \$1,500 per disk, ATA has a cost per GB of less than \$6.00. Your CFO can compare that to a cost of almost \$20 per GB

³ See **The Clipper Group Navigator** dated August 31, 2003, entitled *IBM Supercharges FAStT Family*, at <http://www.clipper.com/research/TCG2003041>.

for the 146.8GB F.C. drive and realize the savings. However, with a rotational speed of only 7,200-RPM, ATA disks do not have the same performance or reliability characteristics of Fibre Channel.

This can fit in very well with Data Lifecycle Management requirements. ATA devices are perfectly suited for fixed content and data reference applications that do not have the same access and utilization requirements as Fibre. Migrating data off higher cost devices and on to ATA may not only save your enterprise a significant amount of money, it may also free up higher priced, high-performance devices for mission-critical application needs. **IBM has stated their intent to enable the concurrent attachment of the FAStT EXP100 SATA drive enclosure and the FAStT EXP700 Fibre Channel enclosure to a FAStT storage server. Although not yet available, this capability will be field upgradable and it will further enable the availability of a multi-tiered storage mix.**

In addition to the hardware functionality, the FAStT also has a broad array of software features implemented to improve data management and storage performance. The *FAStT Storage Manager* enables the partitioning of storage into as many as 64 separate partitions in order to optimize its use. Along with support for the IBM *xSeries* and *pSeries*, the FAStT supports a variety of *Windows* operating systems, along with *NetWare*, *Solaris*, *Linux*, *HP-UX*, and *VMWare*, using both 1 & 2Gb connections. Multiple heterogeneous servers thus can access this one consolidated central system, minimizing the number of RAID arrays required. Other FAStT features include *FlashCopy*, *Dynamic Volume Expansion*, and *Volume Copy* to enhance data management and protection.

FlashCopy takes instant point-in-time copies of logical volumes to enable file restoration, backup, and application test while retaining full access to the primary volume. *Dynamic Volume Expansion* allows administrators to resize logical volumes transparently to the user. This is of great value for email and other rapidly expanding applications. The *VolumeCopy* feature provides full volume replication, allowing read-only access to the

source volume.

Enterprise Storage Server

The *Enterprise Storage Server Model 800 (ESS)*, or *Shark*, is available for attachment to the *iSeries* and *zSeries* (with performance accelerators) in addition to the servers attachable to the FAStT. **ESS extends the physical capabilities of the FAStT family in terms of capacity, performance, and functionality, maintaining many of the physical storage options.** As indicated in Exhibit 2, *Shark* can support almost 56TB of data over 384 high-performance disk drives with up to 32 host connections that can be transferring data concurrently, to and from the disk drives. While it does not support ATA like the FAStT, *Shark* does support multiple Fibre Channel devices at varying speeds. This enables the assignment of mission-critical data to higher speed, smaller capacity devices, while assigning less critical data to lower speed, higher capacity devices. In fact, a 15K, 18GB device costs about \$36/GB, while a 10K, 146GB disk costs about \$19/GB, 47% less. If the data center is managing multiple tiers of data with differing values, it may want to consider a multi-array SAN. They can use an ESS for the higher tiers of mission-critical data using a variety of Fibre devices, and a FAStT configured with EXP100 and EXP 700 expansion units for their low-cost storage.

Unlike the FAStT, *Shark* does support FICON and ESCON as well as Fibre Channel connections. Each of the 16 FC connections has the capability for SAN attachment via switch to heterogeneous servers. **This includes support for a NAS gateway to allow the ESS to handle simultaneous block I/O over a SAN and file I/O over a TCP/IP network.**

Moreover, 64MB of cache and 64 unarbitrated/pipelined internal serial (SSA) disk paths contribute to significantly more performance than the FAStT and other similar systems. In fact, with almost 80,000 IOPS in reading 4K blocks, **Shark has 50% more performance than the FAStT900, with 53,000 IOPS at the high end of the FAStT family.** ESS does not employ RAID 0, 1, or 3. RAID 5 and 10 are the only options, with striping done automatically. The ESS also

expands upon the software support provided for the FASTT, with multiple versions of FlashCopy and a Peer-to Peer Remote Copy capability, as follows:

- **FlashCopy V1** – for point-in-time copy capability for logical volumes on ESS;
- **FlashCopy V2** same as V1 with enhancements to FlashCopy for Data Sets, Multiple Relationships, Incrementals, and ease of use;
- **FlashCopy NOCOPY** option allows for flexible reuse of disk capacity with short term requirements, such as backup to tape;
- **Peer-to-Peer Remote Copy V1** – to constantly maintain a current copy of the local application data at a remote site to support a disaster recovery;
- **Peer-to-Peer Remote Copy V2** – same as V1 with additional support for asynchronous cascading to enable high-performance long-distance data replication, disaster recovery, and backup; and Fibre Channel as the communications link between the primary and secondary systems;
- **Peer-to-Peer Remote Copy for Extended Distance** – for remote data migration, off-site backups, and transmission of inactive database logs;
- **Extended Remote Copy** – for the z/OS and OS/390 operating systems.

Moreover, ESS supports a series of features for use specifically with the zSeries in order to support large mainframe data centers:

- **Parallel Access Volume** – to enable multiple I/Os from z/OS systems to access the same volume at the same time;
- **Multiple Allegiance** – enables different operating systems to perform multiple, concurrent I/Os to the same logical volume;
- **Priority I/O Queuing** – allows the assignment of priority access to storage resources.

Summary

As enterprises continue to move toward an e-business on-demand environment, they face an increasing need to have more information on line, 7x24x365, instantaneously. They also are facing the requirement to preserve

increasingly more historical data to satisfy industry and government regulators. This data does not need to have the same access as your mission-critical files, but it does need accessibility. Yesterday's solutions simply will not cut it. As we see in chess, **sometimes you have to sacrifice an important asset, in order to gain your ultimate objectives: improved productivity, increased flexibility, and better service to your customers.** In the case of enterprise warfare, the replaceable asset may be your existing storage solution, too slow and too inflexible for tomorrow's demands.

You do not need a supplier with a disk array solution; you need to find a vendor with a data management solution. In order to make a specific product recommendation, you need to determine the value that you place on each of your data assets, and where they are in their lifecycle. Without that information all we can recommend is to find a vendor who can address them all and tailor a Data Lifecycle Management system for your enterprise.

Combining the FASTT family and Enterprise Storage Servers with transparent migration services, IBM continues to deliver solutions to enable enterprises to consolidate and to improve their productivity in a heterogeneous environment.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine is the Director, Enterprise Systems for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 23. (Please dial “1-23” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group *Captain's Log*, and “clipper.com” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.