THE CLIPPER GROUP
Navigator™

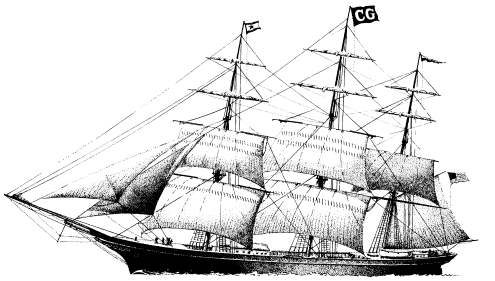

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2003041 August 31, 2003

IBM Supercharges FAStT Family

Analyst: David Reine

Management Summary

From the beginning of time, humans have been learning how to share. In prehistoric times, when they began to live in tribes and develop communal attributes, they learned about the division of labor: building shelters, making clothing, hunting, etc. If each had to do everything independently, then great waste could result. After all, a clan of 20 could have a very nice BBQ from a single mastodon. If everyone hunted and killed a mastodon, then there would be excess, and the mastodon would become extinct. Maybe early humans did not learn that lesson well.

The lessons of sharing continue to this very day. We try to teach our children to share their toys, we try to remind ourselves to share (and spare) our resources, whether personal or communal. From a personal side, we try to be good Samaritans concerning charity; from a communal side, we try to remember to recycle, to preserve our fragile ecology. **When we view our business side, we again must revisit the concept of sharing.**

With the advent of the personal computer in the 1980's, there was a great proliferation of duplicate, often wasted, resources spread across every enterprise. Every PC had its own desktop printer associated with it. The cost of these printers extended well beyond just the acquisition price. It extended to consumables, maintenance agreements and management costs as well. Moreover, the printers were idle the majority of the time. In order to control the costs, reduce that proliferation, and take advantage of a single resource, the print server was invented so that the larger resource could be shared across a workgroup, department, or even an enterprise.

In the 1990's, we saw that same kind of rampant growth affecting the server marketplace. Each server had its own disk devices, not only for the operating system, but also for the storage of data. Moreover, it was not just one or two disk devices. Each server came with anywhere from six to twelve internal disk bays, and what good was an empty disk bay, so we filled them with 9GB, 18GB, 36GB devices and more. The result was an excess of disk space in some servers while others churned, filled almost to capacity. In order to fix this problem, servers connected to external disk arrays, with devices large enough for all of the servers in the office, through a Storage Area Network (SAN). The expected happened next: we filled the arrays and saturated the SAN. Thus, **we continue the cycle of needing arrays and controllers with more throughput, higher capacity and more reliability, availability and serviceability (RAS), than before.** To see how IBM has attempted to get ahead of the curve with the *FAStT* family of disk arrays by providing us with the means to share more data storage, reliably, please read on.

IN THIS ISSUE

The Current FAStT Family

In April 2003, IBM announced the *FAStT600* as an entry level RAID storage array targeted at the mid-range *Windows*, *Linux* and *UNIX* server markets, such as the *IBM eServer xSeries* and *pSeries*. The T600 joined the *T200*, *T700*, and *T900* storage arrays in providing open systems users with an expandable storage platform to satisfy their on-demand expansion requirements. With two active controllers and four 2Gb host interfaces, the T600 was introduced with a capability of 45,500 IOPS (see Table 1, below) and a throughput of 800 MB/s (see Table 2), approx. 4X that of the T200. However, with an entry price of less than $15K, the T600 had supplanted the T200 as the primary offering in the entry position.

Model	IOPS	Drives
T900	148K	140
T700	110K	140
T600	**45.5K**	**42**
T200	11.8K	30

Table 1 - FAStT Family Performance - ***Optimum Performance Level for T600***

Model	MB/s	Drives
T900	800	28
T700	421	14
T600	**800**	**14**
T200	193	10

Table 2 - FAStT Family Throughput - ***Optimum Throughput Level for T600***

IOPS performance is heavily dependent upon the number of drives in the configuration. **This is very relevant for transaction-based applications that involve significant amounts of random I/O on small blocks. The T600 rating for IOPS for the FAStT600 is limited by the ability to install only 42 disk drives for the benchmark. It is also interesting to note that at 800 MB/s, burst rate from cache, the T600 has the same throughput as the FAStT900. This is almost two times the performance of the T700[1].** The throughput rates are heavily dependent on the internal controller bandwidth that is two Gb/s for each I/O interface in the T600 and the T600 Turbo. This factor is important in the transmission of large blocks of sequential data.

Configurable with 14 disk devices in the basic drawer, the T600 supports those 42 devices using two *EXP700* expansion drawers. This scalability makes the T600 an ideal platform for the small and medium business (SMB) markets that desire enterprise-like capabilities to manage up to six terabytes of storage at a reasonable cost. This capacity also enables the T600 for storage consolidation for those SMB customers with the need to sub-divide the array into up to eight partitions.

Introducing the FAStT600 Turbo

With their August announcement, IBM extends the capacity and the throughput of the FAStT600 to limits that were previously unattainable in an entry platform from any vendor. These enhancements include upgrades to the base T600 platform as well as the introduction of a new version, the *FAStT600 Turbo*. This upgrade also addresses a new release to the FAStT array firmware, *Storage Manager v8.4*

FAStT600 Upgrades

Effective with this release, IBM authorizes the attachment of a third EXP700 drawer to the base platform. This attachment extends the scalability of the T600 up to 56 drives, over eight terabytes, the highest scalability for any entry-level platform. In addition, IBM has increased the configurability of the array from eight partitions to 16, a direct result of the requirement to attach the T600 to an IBM *BladeCenter*[2], with the capability to scale to 14 nodes. This is a quick response by IBM to a perceived shortcoming in the original design.

---

[1] With four internal drive loops, the FAStT900 has twice the internal throughput as the FAStT600, which performs at 400MB/s internally.

[2] See *IBM BladeCenter - A Glimpse at the Future of Computing* in **The Clipper Group Navigator** dated October 4, 2002 at http://www.clipper.com/research/TCG2002038.pdf.

IBM has also addressed the firmware side of the platform with the inclusion of *FlashCopy* in the base configuration. This enables the Snapshot capability in the T600.

FAStT600 Turbo Introduction

The FAStT600 Turbo extends the configurability and performance of the base platform from an entry-level array into a mid-range platform for small and medium size businesses or enterprise departments with a need to consolidate both Intel-based operating system servers and UNIX engines into a single SAN. These customers will gain by improving their price/performance profile while at the same time protecting their data, reducing the total cost of running their computer operations. These requirements are common for a wide range of industries from Finance to Transportation, to Health and Retail.

The Turbo option includes:

- One gigabyte of cache per controller up from 256 MB;
- Eight partitions standard with the capability to scale to 64;
- A connection for up to 112 drives, 16 TB of storage.

Because of these upgrades, the T600 Turbo achieves an improvement in the IOPS performance, as shown in a new set of benchmarks to measure array performance (See Table 3).

The attachment of a third EXP700 drawer to increase the device count to 56, as well as an increase in the cache size, improves the array performance by 70%. This is an excellent indicator of what to expect from this family in the future.

Model	IOPS	Drives
T900	148K	140
T700	110K	140
T600 Turbo	**77.5**	**64**
T600	45.5K	42
T200	11.8K	30

Table 3 - FAStT Family Performance - ***Optimum Performance Level for T600 Turbo***

Storage Manager v8.4

The significant new features that are available with *FAStT Storage Manager v8.4* as an option are:

- VolumeCopy – available to make a clone copy of the data for disaster recovery or a test environment;
- Dynamic capacity and volume expansion to grow existing storage;
- An increase to 64 partitions;
- An increase to 256 volumes per partition;
- A new graphic interface tool to facilitate the addition of servers and applications; and
- The extension of an array group to greater than 2TB.

New customers can order the FAStT600 Turbo directly from the factory. **Customers who have already installed the T600 in a base configuration can order a field upgrade kit, thus protecting their investment in this platform technology.**

Product Positioning

Positioned as it is, in the open systems marketplace, the FAStT family is competing with a variety of vendors and array products for attention in a market targeted to the data center. Here, there is a critical need for improved storage utilization and scalability. The arrays need to meet application specific demands for high transaction I/O and data rate performance. For the purposes of this analysis, however, we will confine ourselves to reviewing the capabilities of the FAStT600 to open systems products from the major players in this space.

Competitive Product Capabilities

The major vendors in the open systems arena typically will offer three products to satisfy the needs of their customers. They will offer to the SMB market and, to some extent, the enterprise-level customers that they are courting: an entry-level system, a mid-range capability, and an enterprise-wide solution. Usually these products are the result of a carefully crafted strategy to provide a scalable range of products to their customers. In some

cases, however, they are the result of a never-ending series of mergers and acquisitions that leave a series of independent solutions, which prevents the migration from entry to enterprise.

The first, or entry-level, is typically restricted in the operating systems supported and limited in scalability, typically to 30 drives. In some cases, for example, the RAID arrays will only support Microsoft or Linux operating systems. There may be no UNIX capability. Therefore, these arrays do not have the same capabilities as the FAStT600 Turbo and we will not consider them in this analysis. The high-end, or enterprise-wide solutions, usually have capacity for up to 200-240 drives. That much scalability puts them into the same category as the FAStT700. Therefore, we will ignore these arrays, also, in this analysis.

With expandability to three drawers and 56 disk devices, and with 8TB of storage, the FAStT600 is comparable to the typical array in the mid-range space with scalability to 3 expansion drawers and a total of 55-60 drives. With the Turbo scalability, however, the T600 outstrips these arrays by almost 2 to 1, with 7 drawers and 112 drives. In addition to the physical comparison, the T600 Turbo also has considerably more logical flexibility, with a maximum of 30 disks in a volume group and 1,024 volumes in the array. Some mid-sized arrays are limited to 16 devices in a volume group and 512 volumes in an array.

A similar pattern holds true in a performance comparison. The IOPS rating of 77,500 for the T600 Turbo is almost 30% "FASter" than most of the products, at around 60,000, measuring random cache reads. With a max throughput of 800 MB/s, the T600 Turbo has 18% more performance than those products, doing sequential cache reads. Other products, though, have been measured as high as 100,000 IOPS, although with a throughput limited to < 400MB/s.

There are also architectural advantages for the FAStT family. First, some arrays require a minimum of five drives because they use them for microcode and the configuration database, and to continually save the cache in case of a power failure. The FAStT600 backs up the cache with redundant batteries that can handle a power off condition for 7 days. Second, the FAStT subsystems have an extra RAS feature: each drive contains complete setup information defining which array group it is a part of and which logical volume stripes are contained.

As with the FAStT Storage Manager, other products come with options for snapshot and volume copy features. However, where IBM has positioned the Storage Manager with a fixed price regardless of the amount of storage or the number and type of servers attached, many competitors use a sliding scale to increase the price of the premium features. In complex networks with significant storage, the software can be expensive.

Conclusion

The primary objective for any SAN storage array is to provide a consolidation platform so that a network of heterogeneous servers can efficiently share a common resource. In this context, *efficiently* clearly defines a high performance engine that can provide requested data to the application server in a fast, secure, and reliable manner. At the same time, it must protect the investment that the customer makes in the procured resources with a family of upgradeable arrays.

From the computer room standpoint, IBM has delivered a set of products, the FAStT family, in general, and the FAStT600 Turbo, in particular, which do exactly that. With superior performance, scalable capacity, and outstanding throughput, IBM is also delivering the software functionality required to provide a secure environment in a price-performant package. **The 600 Turbo is positioned, physically, logically, and economically, to meet the growing storage requirements of small, medium, and even larger enterprises.**

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 a Senior Contributing Analyst for the 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 a MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 37. (Please dial "1-37" 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.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.