



IBM Systems Group

T14

FAStT100 Overview

xSeries Technical Conference
August 2004

Agenda

- ATA Positioning
- FAStT100 Overview
- FAStT Positioning



The ATA Appeal Is Growing

- Disk drives can represent the majority of acquisition cost. Customers are looking to ATA drives to reduce the hardware cost per gigabyte of their storage systems
- Understanding the differences between SATA and FC drives will be critical to proper implementation

Drive Comparison

Interface	PATA	SATA	FC
Interface	Single port	Single port	Dual port
Communication	Half duplex	Half duplex	Full duplex
Transmission	Parallel	Serial	Serial
Bandwidth	133 MB/s	150 MB/s	200 MB/s
Command queuing	None	None	Extensive
Connections	Master / slave	Point-to-point	Point-to-point, loop, fabric

Mechanism	PATA	SATA	FC
Rotational velocity	5,400 RPM	7,200 RPM	10K / 15K RPM
Average time to data	~16 ms	~ 13 ms	7.7 / 5.6 ms
Single disk I/O rate	---	88 IOPS	280 / 340IOPS
Max transfer rate – sustained read	46 MB/s	59 MB/s	69 / 76 MB/s
Max transfer rate – sustained write	---	30 MB/s	68 / 71 MB/s

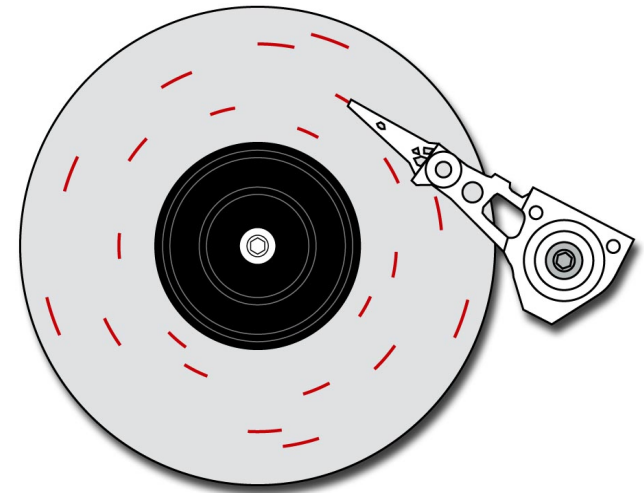
Application Storage Profiles

- FC and SATA will typically be deployed to satisfy very different application requirements

APPLICATION	Read Intensive	Write Intensive	I/O Intensive	Throughput Intensive	Random Access	Sequential Access
OLTP	●	●	●		●	
Data Warehouse	●		●		●	
System (SCP)	●		●		●	
File Serving	●			●	●	
Medical Imaging		●		●	●	
Web/ Internet	●			●	●	
Multimedia Video	●			●		●
Document Imaging		●		●		●
CAD/CAM	●		●		●	
Backup/ Recovery		●		●		●

Application Access Patterns – IOPS

- Random, small-block transfers
 - OLTP, databases, Exchange
 - Vast majority of enterprise applications
- Drive performance enablers:
 - Number of drives
 - More drives equals more IOPS
 - Seek time, latency, rotational velocity, command queuing



SATA


10K FC

15K FC

IOPS -- SATA ~25% of 15K FC

Application Access Patterns – Throughput

- Sequential, large-block transfers
 - Video servers, rich media, seismic processing, HPC
- Drive performance enablers:
 - Data transfer rate, max I/O size, command queuing
 - Max controller bandwidth typically reached with relatively low number of drives, negating small differences in per drive transfer rates



Throughput -- SATA ~ 67% of FC

Application Drive Requirements

- FC and SATA will typically be deployed to satisfy very different application requirements
- FC – performance requirements
 - Maximum number of spindles per box/capacity point, spinning at maximum rotations per minute
- SATA – cost-per-gigabyte requirements
 - Minimum number of spindles per box/capacity point, spinning at acceptably fast rotations per minute

SATA's Target Applications


- SATA is ideally suited for secondary storage applications
 - Disk-to-disk backup and restore / spinning archive
 - Retention managed data / reference data / fixed content
 - Medical imaging, object folders, historical reports
 - “Nearline” – data previously stored on tape
 - Disaster recovery
 - Restoration source of archived information
 - Hot site storage of “like” data from primary site
- SATA as primary storage
 - Bandwidth / streaming applications
 - HPC, video streaming, media / rich content, oil and gas
 - Primary requirement: good throughput at low cost per GB
 - Entry-level SMB
 - Primary requirement: low-cost
 - Entry-level storage is an emerging opportunity for SATA due to its common requirements with secondary storage

Classes of Storage

	Primary	Secondary	Entry-level	Archival
Typical applications	OLTP, databases, ERP, email	Retention managed data, fixed content, D2D backup/restore	1 st external RAID, file / print server	Archive and retrieval
Requirement				
IOPS performance	Highest	Minimal	Minimal	---
MB/s performance	Highest	High	Minimal	High
Time to data	Immediate	~ Immediate	~ Immediate	Prolonged
Media reliability	Highest	High	High	Good
Uptime	24/7	< 24/7 OK	< 24/7 OK	< 24/7 OK
Price	Higher is OK	Low cost-per GB	Lowest	Lowest
Primary media	FC disk	SATA disk	SATA / SCSI disk	Tape

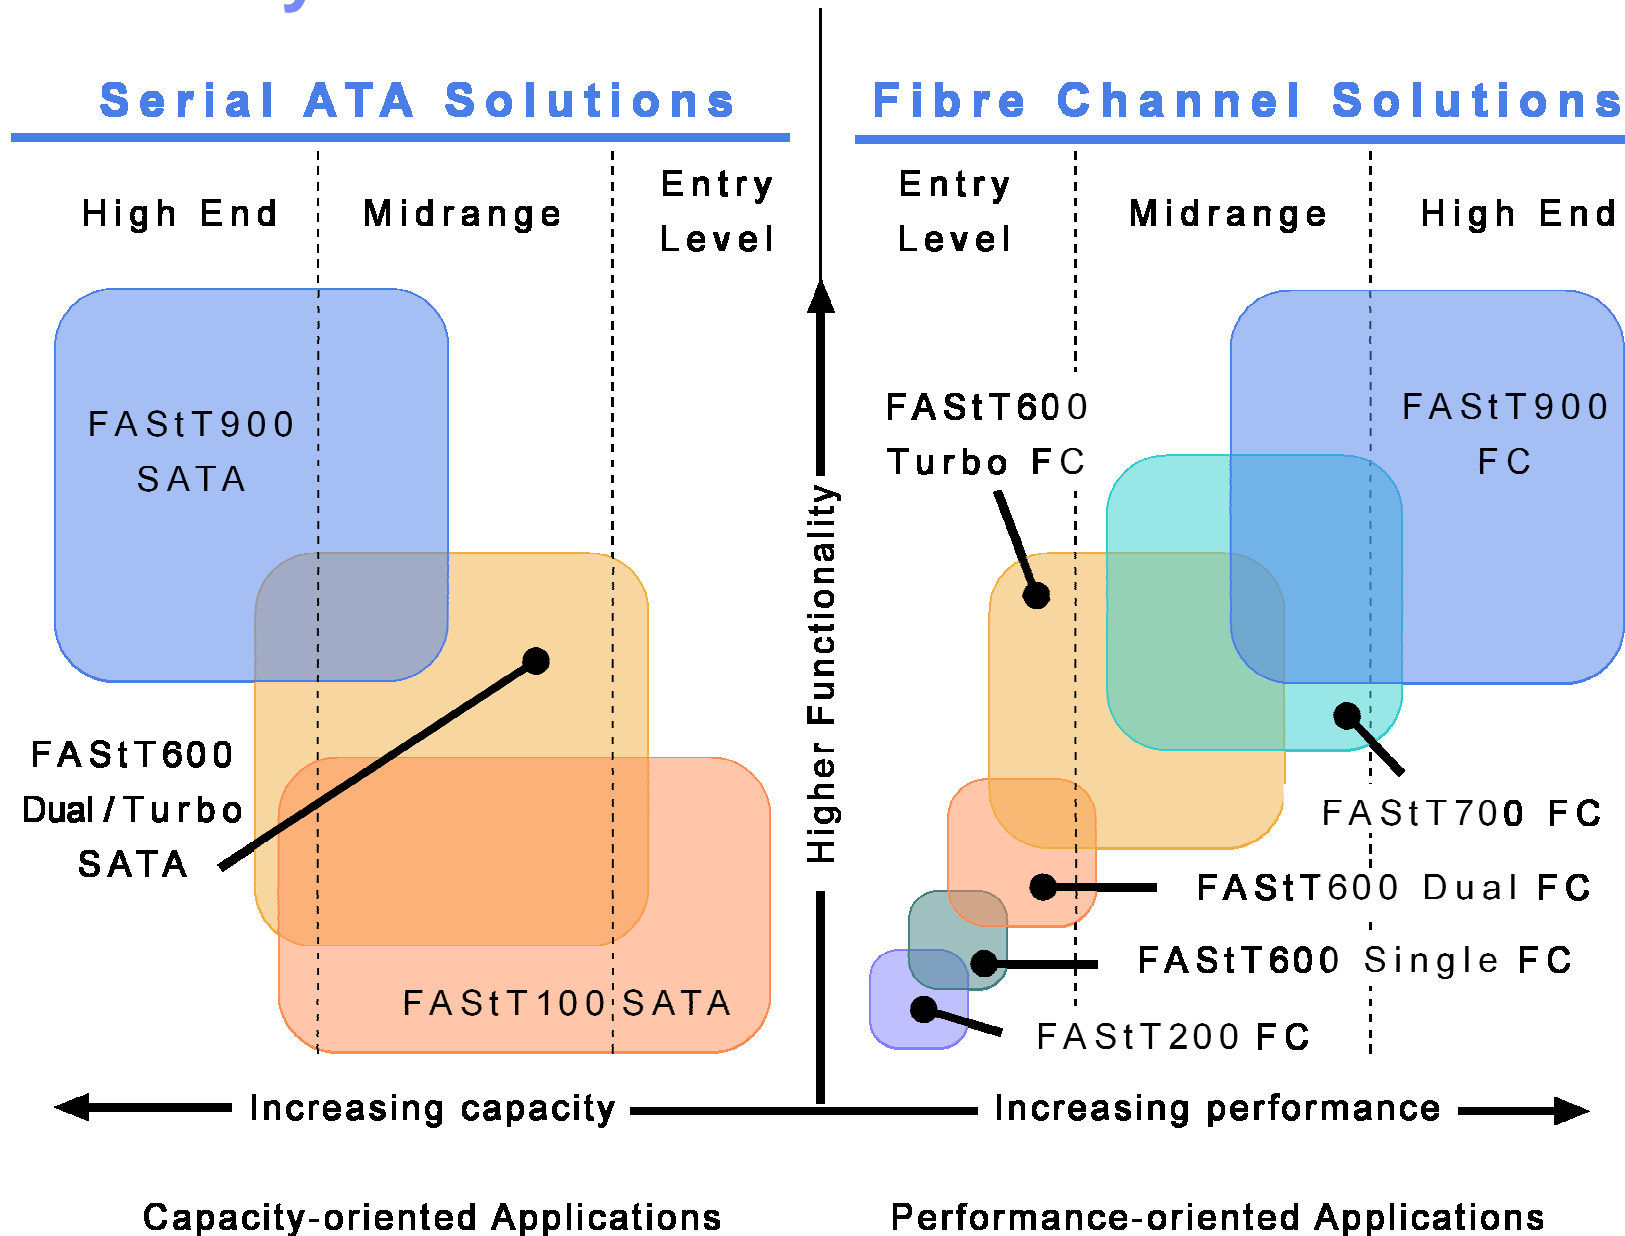
FAStT100 Brings New Opportunity

- Native SATA support enables controllers and up to 14 drives to be integrated in a single enclosure
 - 3U entry point
- The FAStT100 is an entry-level system, with midrange potential and high-end functionality
- FAStT100 continues the FAStT family heritage
 - Supports powerful FAStT features and functionality
 - Centralized administration and a common interface across all FAStT products lowers storage management costs
 - FSM provides robust functionality with an intuitive GUI
 - Online administration and scalability maintain continuous data availability
 - Modular “pay-as-you-grow” scalability lowers acquisition/expansion costs
 - Configuration flexibility creates superior storage utilization



FAStT100 Storage Server


FAStT Family




FAStT100 Product Overview

- **FAStT100 Storage Server***
 - Four 2 Gbps FC host/SAN connections
 - 512 MB cache memory
 - Supports 14 internal SATA drives
 - 250 GB, 7200 RPM
 - Two 2 Gbps FC drive channels for expansion
 - Max of 56 drives (14 TB)
 - Default host group included
 - No SAN partitions
- **EXP100 Disk Enclosure**
 - 14 Serial ATA disk drives
 - 250 GB, 7200 RPM
 - 2 Gbps FC connectivity

*Dual controller specifications



FAStT100 Storage Server



EXP100 Disk Enclosure

Performance Comparison

Random, small-block I/O Transactional applications	FAStT100 SATA	FAStT600 Turbo SATA	FAStT900 SATA
Sustained I/O rate 4k disk reads	5,000 IOPS	10,000 IOPS	19,800 IOPS
Sustained I/O rate 4k disk writes	1,000 IOPS	2,200 IOPS	4,600 IOPS
Drives under test	56 SATA	112 SATA	224 SATA

Sequential, large-block I/O Bandwidth applications	FAStT100 SATA	FAStT600 Turbo SATA	FAStT900 SATA
Sustained throughput 512k disk reads	485 MB/s	395 MB/s	780 MB/s
Sustained throughput 512k disk writes	415 MB/s	310 MB/s	615 MB/s

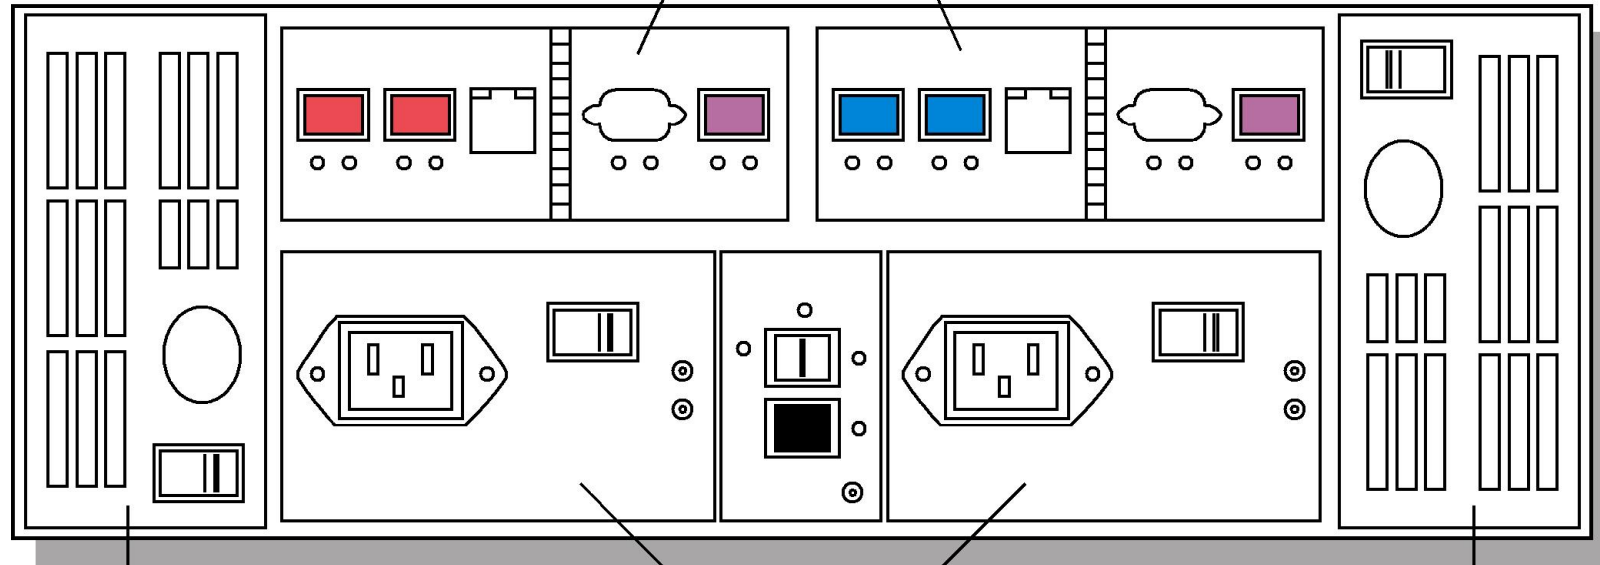
FAStT100 Storage Server

 Controller A

 Controller B

Controller CRUs

 Drive Expansion



Cooling Fan CRU

Power Supply CRUs

Cooling Fan CRU

FAStT100 Controller

- Specialized I/O system core

- Integrated XOR for high-speed RAID parity calculations

- High-speed busses

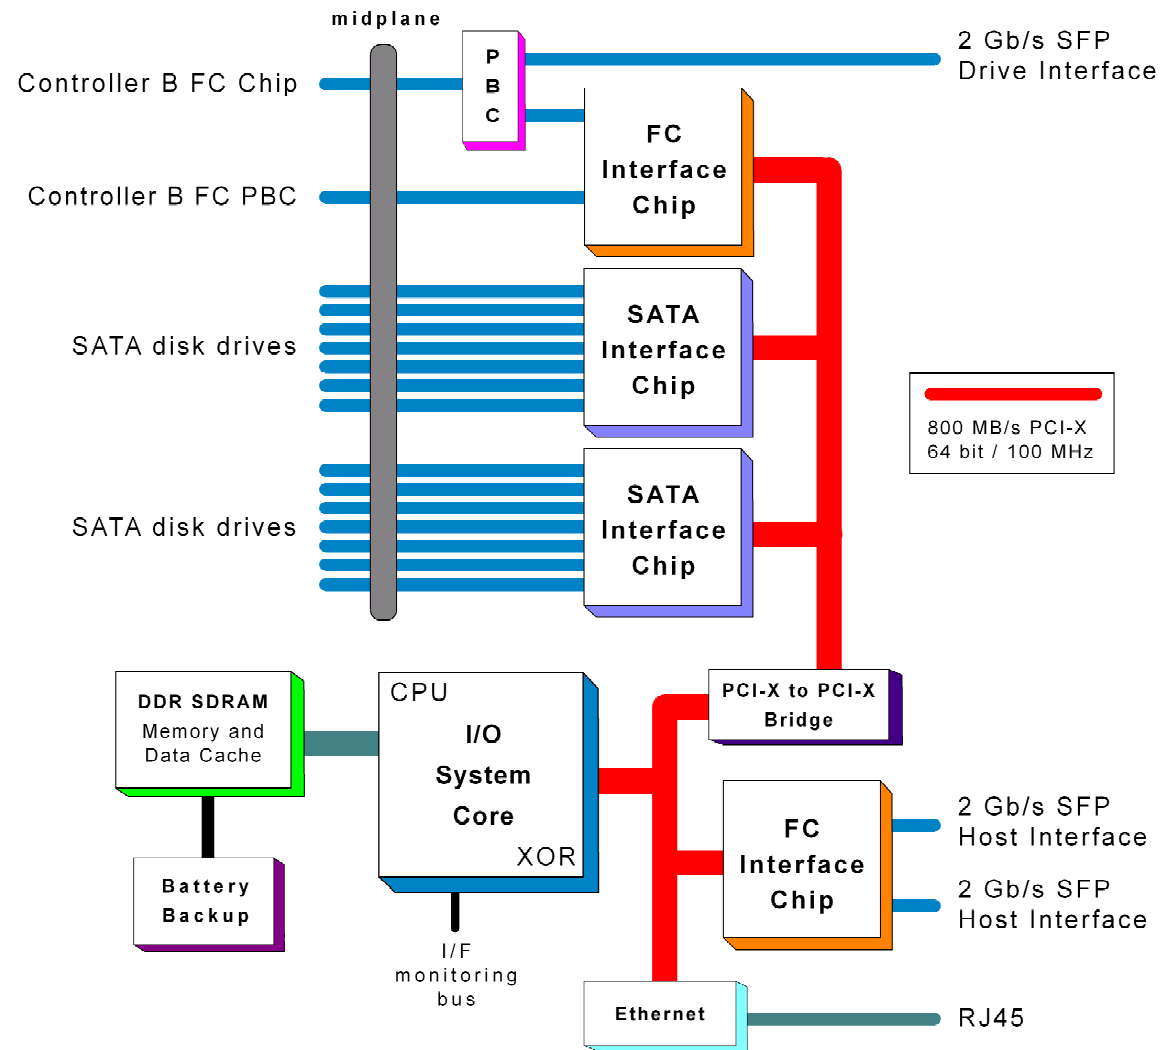
- 2 Gbps host interfaces

- 2 Gbps drive channels with internal PBC Hub

- Links controllers
- EXP100 expansion



- SATA interface chips

- Native support for SATA drives **NEW**



Intelligent SATA Drive Interface

- SATA disk drives are single-ported and lack true multipathing and multi-initiator capabilities
- SATA Interface Card provides the dual-link capability needed to achieve redundant paths
 - Its electronic link switch is in a static position unless failover requirements dictate otherwise.
- Additional benefits:
 - Drive signing
 - Unique World Wide Name
 - Drive activity LED (external to drive)
- The SATA Interface Card is housed inside SATA drive module, attaching directly to SATA drive
- Currently implemented on SATA drives in EXP100





IBM Systems Group

FAStT Positioning When to Sell What Where

xSeries Technical Conference
August 2004

FAStT Storage Servers

- Fully-featured family of solutions designed to address a wide range of storage requirements
 - From entry-level to enterprise
 - From performance-oriented to capacity-oriented applications
 - From direct-attach to SAN-attach
- True family of solutions
 - Common management interface and functionality across all products and platforms
 - Seamless “data intact” upgrades / reconfigurations
 - Common components



Performance Positioning

- Two primary performance metrics – IOPS and MB/s
- IOPS measures random, small-block I/O
 - Important for transaction-based applications such as OLTP, databases, Exchange, web servers, file servers
- MB/s measures sequential, large-block I/O
 - Important for throughput-intensive applications such as video, seismic processing, high performance computing (HPC)
- Performance-oriented applications – either IOPS and MB/s – will be best served by FC drives
 - SATA may be sufficient for throughput applications

IOPS – Transactional Applications

- Achieving high IOPS requires a streamlined controller design to handle large amounts of small I/Os
- Drive factors
 - IOPS is dependent on the # of drives – more drives, more IOPS
 - Drive-limited configurations can result in similar performance between systems with different maximum capabilities
 - FC vs. SATA – key performance enablers are seek time, latency, rotational velocity, command queuing – all FC advantages
- Controller factors
 - All FASTT controllers have hardware XOR engines
 - Custom ASICs and 4 drive loops in FASTT900 / 700 enable higher performance as the # of drives and I/Os increases

IOPS Performance

Results in IOPS	FAStT100	FAStT200	FAStT600	FAStT600 Turbo	FAStT700	FAStT900
Sustained I/O rate 4k disk reads	N/A	3,800	17,000	25,000	40,000	53,200
Sustained I/O rate 4k disk writes	N/A	1,170	4,400	5,200	9,100	10,900
Drives under test	N/A	66 FC	56 FC	64 FC	140 FC	140 FC

Sustained I/O rate 4k disk reads	5,000	N/A	5,000	10,000	N/A	19,800
Sustained I/O rate 4k disk writes	1,000	N/A	1,000	2,200	N/A	4,600
Drives under test	56 SATA	N/A	56 SATA	112 SATA	N/A	224 SATA

Note: Results achieved under ideal circumstances in a benchmark test environment. Actual customer results will vary based on configuration and infrastructure components.

Application Storage Profiles

APPLICATION	Read Intensive	Write Intensive	I/O Intensive	Throughput Intensive	Random Access	Sequential Access
OLTP	●	●	●		●	
Data Warehouse	●		●		●	
System (SCP)	●		●		●	
File Serving	●			●	●	
Medical Imaging		●		●	●	
Web/Internet	●			●	●	
Multimedia Video	●			●		●
Document Imaging		●		●		●
CAD/CAM	●		●		●	
Backup/Recovery		●		●		●

Best served by FC drives

MB/s – Sequential Applications

- Throughput rates are heavily dependent on the internal controller bandwidth and number of host/drive ports
- Drive factors
 - Max rates can typically be reached with a smaller # of drives
 - FC vs. SATA – key performance enablers are data transfer rate, max I/O size, command queuing – slight FC advantages
- Controller factors
 - FASStT900's four host / drive channels give it the advantage
 - FASStT600 and FASStT100 have more internal bandwidth than the FASStT700

Throughput Performance

Results in MB/s	FAStT100	FAStT200	FAStT600	FAStT600 Turbo	FAStT700	FAStT900
Sustained MB/s 512k disk reads	N/A	170	400	400	390	795
Sustained MB/s 512k disk writes	N/A	100	315	315	240	635
Drives utilized	N/A	FC	FC	FC	FC	FC

Sustained MB/s 512k disk reads	485	N/A	395	395	N/A	780
Sustained MB/s 512k disk writes	415	N/A	310	310	N/A	615
Drives utilized	SATA	N/A	SATA	SATA	N/A	SATA

Note: Results achieved under ideal circumstances in a benchmark test environment.
Actual customer results will vary based on configuration and infrastructure components.

Application Storage Profiles

APPLICATION	Read Intensive	Write Intensive	I/O Intensive	Throughput Intensive	Random Access	Sequential Access
OLTP	●	●	●		●	
Data Warehouse	●		●		●	
System (SCP)	●		●		●	
File Serving	●			●	●	
Medical Imaging		●		●	●	
Web/Internet	●			●	●	
Multimedia Video	●			●		●
Document Imaging		●		●		●
CAD/CAM	●		●		●	
Backup/Recovery		●		●		●

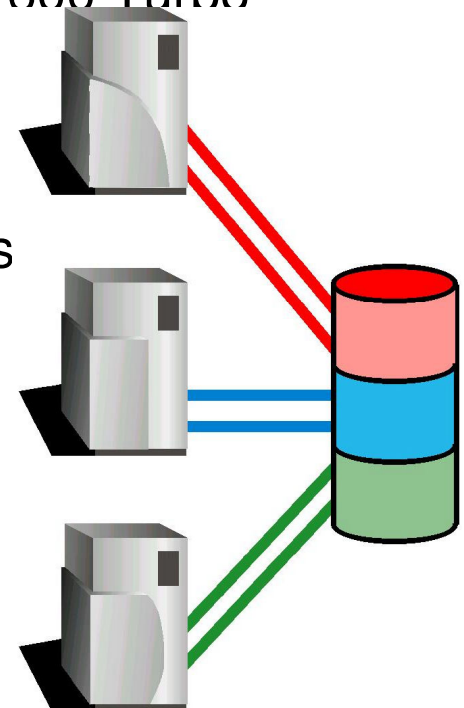
Option for SATA drives

Performance Positioning

- Understand customers requirements
 - Looking for highest performance or best performance value?
 - Is the application best served by FC, or is SATA OK?
- Highest performance always delivered by FAStT900 with FC disk
 - IOPS and MB/s (throughput)
- FAStT600 Turbo (FC) delivers the best IOPS performance value up to about 56 drives
 - Performance value reflects best price/performance ratio, where performance improvement does not offset price difference
- FC for performance-oriented, SATA for capacity-oriented
- Additional customers requirements
 - Looking for highest performance or best performance value?
 - Need future performance “headroom”?
 - Other factors to consider: consolidation, replication, workload

Consolidation Positioning

- FAStT900
 - Enterprise consolidation – largest # of heterogeneous servers
 - High or unpredictable workload
- FAStT700, FAStT600 Turbo
 - Departmental consolidation – “medium” number of servers
 - Heterogeneous – FAStT700, homogeneous – FAStT600 Turbo
 - Consistent workloads
- FAStT600, FAStT200, FAStT100
 - Workgroup consolidation – “small” number of servers
 - Minimal performance requirements



Data Replication

■ FlashCopy

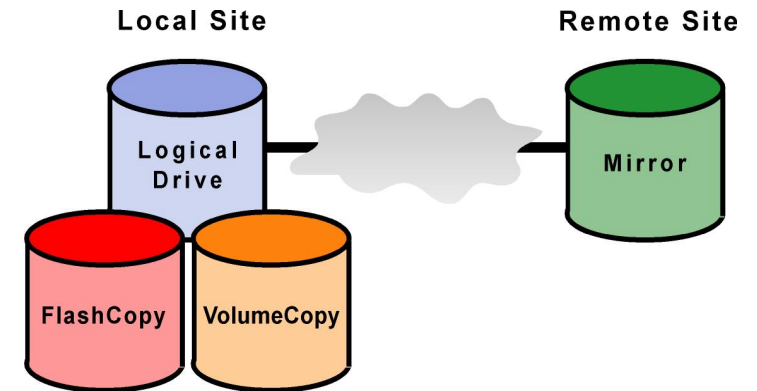
- Creates a point-in-time image
- Ideal uses: backup source, restoration point, checkpoint

■ VolumeCopy

- Creates a complete physical copy (clone)
- Ideal uses: data mining / analysis, PiT archive

■ Remote Mirroring

- Creates a continuously updated copy at a remote location
- Ideal use: disaster recovery



Data Replication Comparison

Attribute	FlashCopy	VolumeCopy	Remote Mirror
Industry term	Snapshot or Image	Clone or Copy	Mirror
Replication type	PiT image of source	Complete PiT copy of source	Continuously updated copy of source
Copy location relative to source	Same storage server	Same storage server	Remote storage server
Time required to create / establish copy	Instantaneous	Dependent on size and prioritization setting	Dependent on size, replication distance and prioritization setting
Performance impact of creating copy	None	Primarily dependent on prioritization setting	Primarily dependent on prioritization setting
Ongoing impact of maintaining copy	Primarily dependent on write percentage	None	Dependent on write percentage and replication distance
Maximum copies	4 per source	8 active copies	32 mirrors
Copy dependence on source	Lost if source is lost	Available if source is lost	Available if source is lost
Data restoration from copy	Designed for easy file and/or source restore	Designed for easy file and/or source restore	Designed for recovery of primary site disaster
Application access to source and target	Full r/w access while FlashCopy remains active	Source is read-only during copy process	Requires mirror be "broken" for r/w access to target

Data Replication Positioning

- FAStT900
 - Best choice when replication is a key piece of storage strategy
 - Least performance impact
- FAStT700, FAStT600 Turbo
 - OK for limited number of concurrent replication process
- FAStT600, FAStT200, FAStT100
 - OK for small number (1-2) of concurrent replication process
 - Largest performance impact

FAStT Positioning Summary

High End	900	<ul style="list-style-type: none"> ◆ Highest IOPS and throughput performance ◆ Best IOPS performance value for large configs (~56+ drives) ◆ Data replication is key piece of storage strategy ◆ Heterogeneous enterprise storage consolidation
Mid Range	700	<ul style="list-style-type: none"> ◆ Price-conscious IOPS performance (~56+ drives) ◆ Mid-range data replication and remote replication site ◆ Heterogeneous departmental storage consolidation
	600 Turbo	<ul style="list-style-type: none"> ◆ Best IOPS performance value for medium-scale configs (~28 to ~56 drives) ◆ Workgroup / departmental storage consolidation ◆ Windows, Linux, NetWare servers
	600 Dual	<ul style="list-style-type: none"> ◆ Best IOPS performance value for small configs (< ~28 drives) ◆ Excellent dual-channel throughput performance ◆ 2 Gb/s entry-level external RAID storage
Entry Level	600 Single	<ul style="list-style-type: none"> ◆ 2 Gb/s entry-level external RAID storage ◆ Upgradeable to dual controller and Turbo features ◆ Robust management software provides feature/functionality
	200	<ul style="list-style-type: none"> ◆ Feature / functionality with limited performance requirements ◆ 1 Gb/s entry-level external RAID storage
	100	<ul style="list-style-type: none"> ◆ Low-cost 2 Gb/s entry-level external RAID storage ◆ Robust management software provides feature/functionality ◆ Excellent MB/s performance value for secondary storage



IBM Systems Group

Questions

