

Filing Information

December 2001
IDC #26241E

Storage Systems

Excerpt

Solid State Disks: The Right Products for Boosting Performance in Tough Times — Product and Market Analysis, 2001

Analysts: Massaki Moriyama and Robert Gray

IDC Opinion

Why is a solid state disk (SSD) a smart choice for transaction-intensive application environments?

Too often, end users faced with performance issues or poor I/O response take the unnecessarily expensive “brute force” approach of simply adding more servers or more disk drives to spread the workload. This ignores the total lifetime cost of ownership (TCO), which is often the dominant expense. We believe that in very high transaction rate applications, solid state disk would be a better solution. Typical instances include the increasing use of cell system short message services and the database-oriented transaction expansion. Sophisticated users will find the approximate 150 I/O per-second transaction limit of a single spinning (rotating) disk drive is the real bottleneck. The bottleneck may be easily overcome by deploying SSDs as the solution to high I/O rate application needs.

Dramatically lower prices and higher capacities make SSDs a match for many choked and throttled applications.

Introduction

End users with transaction-oriented applications should consider the benefits and leverage of SSD in their storage hierarchy and infrastructure. Dramatically lower prices and higher capacities make SSDs a match for many choked and throttled applications.

SSDs make the entire system run more efficiently by ensuring the I/O capability is in balance with the rest of the application system.

SSDs:

- Improve efficiency of application servers by recovering CPU cycles formerly lost in I/O wait loops.
- Improve application response time in such critical applications as message systems and database applications.
- Improve storage system response by off-loading hot files to specialized SSD device and data paths.

This bulletin profiles the market, main suppliers, and products. Forces accelerating and retarding growth are analyzed. As SSDs leverage other IT investments, IDC expects end users' commitment to SSDs to increase the product segment's overall revenue market share.

Worldwide SSD annual spending is approximately \$50 million.

File Cache Versus Block Cache

SSDs are deployed as a disk volume in support of a specific set of files. SSDs are thus most frequently used as a file cache. Typically, multiple frequently accessed (hot) files are stored on the SSD volume. File caches or SSDs provide optimum value when a high percentage of disk I/O requests are directed to a relatively small number of files. Typical applications with such designs include: message queues in collaborative (email) applications, indices, and key tables in databases. Block cache, in comparison, provides best value where there is a high degree of locality of access, but where file-level metadata is not available. Thus, block caches are a standard feature of disk drives and RAID systems. As such, block caches increase disk performance, while SSDs may be thought of as increasing application performance.

Quoting IDC Information and Data: *Internal Documents and Presentations*—Quoting individual sentences and paragraphs for use in your company's internal communications does not require permission from IDC. The use of large portions or the reproduction of any IDC document in its entirety does require prior written approval and may involve some financial consideration. *External Publication*—Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2001 IDC. **Reproduction is forbidden unless authorized.**

For additional copies please contact Cheryl Toffel, 508-935-4389.

Check us out on the World Wide Web!

<http://www.idc.com>
Printed on recycled materials. ♻️

SSD Definitions

- **True SSDs.** This bulletin defines SSD as DRAM, backed by a rotating disk drive with integrated battery backup. Capacities range from 134MB to more than 51GB. SSDs are typically connected to servers or SANs via parallel SCSI or Fibre Channel interfaces and are often used with large enterprise or utility-grade storage systems, or integrated with servers and software for vertical solutions. Typical applications include ebusiness transaction databases, Internet email, wireless messaging, and mobile value-added services. Form factors for SSD systems include both 3.5in. disk size and 19in. rackmount. Prices range from \$5,000 to more than \$300,000, depending on capacity and feature set.
- **Flash storage.** Flash memory is also sold in configurations that are labeled SSD by the supplier. Flash-based storage is generally slower than DRAM-based products. Also, flash memory's destructive-write characteristic limits applications to those with low lifetime duty cycles. Server and enterprise applications typically are not a good fit for flash-based SSD. This research does not include flash-based SSD.
- **Global block cache.** EMC Symmetrix, Hitachi 7700/9900, and a handful of other cache-based storage array products allow users to specify a portion of the shared block cache to function as a file cache (SSD). These permanent caches can only mirror data contained on the disk drives of the array, and this approach reduces the cache efficiency for other storage files. This research does not include storage system SSD functionality utilizing shared block cache.

Market Overview

While SSD has a stable 20-year market history, a number of factors are now revitalizing the market:

- Lower prices
- Applications unfriendly to rotating mechanical disk drives
- The growth of the Internet and digital communications
- The growing gap between disk I/O performance and CPU MIPS performance

Table 1 tracks the relationship between disk drive access time evolution and CPU cycle time. Under the assumptions in Table 1, small, single-CPU servers require the services of 100 disk drives and 3.6TB of capacity for balanced system behavior in 2002. A balanced computer system is one where the I/O workload and the compute workload are in balance and neither is overloaded.

Table 1
Balanced Compute and I/O Microprocessor System Sweet Spot, 1995–2005

	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005
Typical disk drive I/Os	180	190	200	210	220	230	240	250	260	270	290
Typical disk capacity (GB)	2	4	4	9	9	18	18	36	73	73	144
I/Os per GB of disk drive capacity	90.0	47.5	50.0	23.3	24.4	12.8	13.3	6.9	3.6	3.7	2.0
Typical microprocessor (MHz)	100	166	200	275	350	1200	1500	2500	4500	6000	8000
Server I/Os per 100MHz	1,000	1,000	1,000	1,000	1,000	1,000	1,000	1,000	1,000	1,000	1,000
Typical server I/O required (even I/O distribution) per second	1,000	1,660	2,000	2,750	3,500	12,000	15,000	25,000	45,000	60,000	80,000
I/O data rate with 4K blocks (MB per second)	4	6.64	8	11	14	48	60	100	180	240	320
Number of disks required	5.6	8.7	10.0	13.1	15.9	52.2	62.5	100.0	173.1	222.2	275.9
Typical capacity to get I/Os (GB)	11.1	34.9	40.0	117.9	143.2	939.1	1,125.0	3,600.0	12,634.6	16,222.2	39,724.1
Typical capacity to get I/Os (TB)	0.0	0.0	0.0	0.1	0.1	0.9	1.1	3.6	12.6	16.2	39.7

Key Assumptions:

- A price point of \$2,500 is assumed per *Worldwide PC Semiconductor Market Forecast and Analysis, 2000–2005* (IDC #25081, July 2001).
- The I/O per second assumes a typical distribution of seeks within and across disk drives.
- Disk capacity is the assumed capacity sweet spot for the year listed.
- The reduced disk latency assumes a move toward 15K RPM disk drives.
- The workload is spread evenly across all disk drives.
- Accesses result in few cache hits.

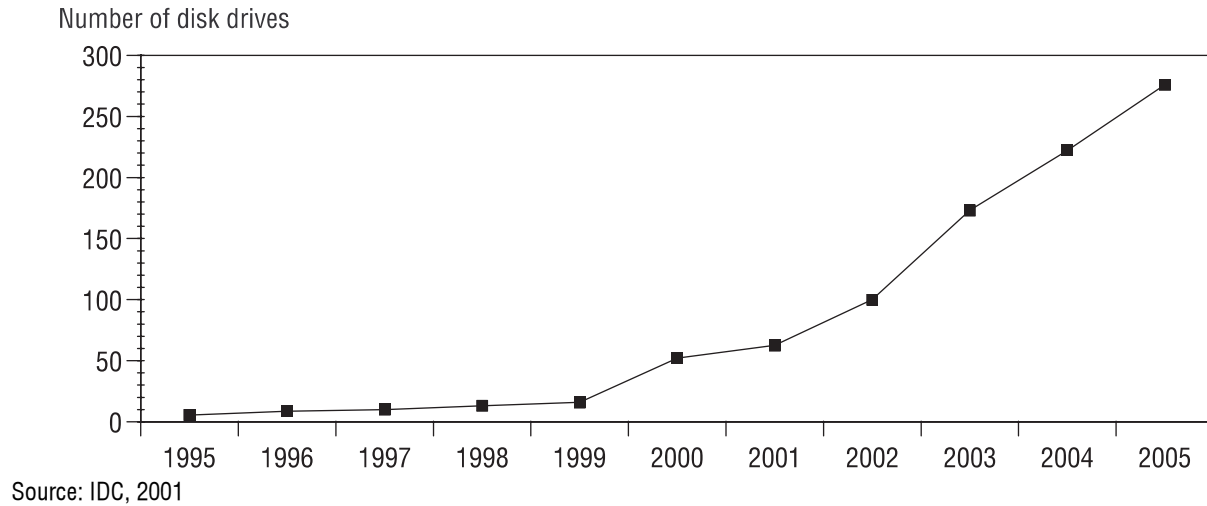
Message in the Data:

- Small servers will increasingly require SSD characteristics for balanced system performance.

Source: IDC, 2001

Figure 1 presents the key relationship in Table 1.

Figure 1
Gap Between I/O Performance and CPU MIPS Performance, 1995–2005



We see future market expansion in the following areas:

- IDC forecasts the evolution of storage area networks (SANs) and the establishment of a virtualization layer to be a significant enabling layer in accelerating SSD revenue growth at the high end.
- The mismatch of disk drive performance with server throughput on the low end will drive low-end expansion. Small multi-MIP servers are best served by a combination of SSD and a few disk drives.

The improved strategic outlook has not gone unnoticed by investors, and nearly all SSD suppliers have attracted additional capital investments within the past year. We expect the opportunity to draw new entrants and more marketing activity.

Solid State Disk Benefits

SSDs deliver value based on the fact that usually just a few files representing 1–5% of capacity represent a very large percentage of all I/O activity.

SSDs deliver value based on the fact that usually just a few files representing 1–5% of capacity represent a very large percentage of all I/O activity. By greatly speeding up access and reducing latency to this transaction-intensive data, the entire application performance improves.

Table 2 and Figure 2 chronicle the number of CPU cycles lost during a single I/O operation if the CPU is unable to do productive work on other tasks. The second set of data captures the same lost cycles, assuming the I/O is handled by an SSD with a 100 microsecond response time. Over 97% of the lost cycles are recaptured for

application use under the assumptions for each year from 1995 to 2005.

In I/O-intensive applications, SSD suppliers claim demonstrated system application performance gains from 200% to 800%. Alternative performance solutions are:

- **Buy more servers.** Application partitioning and server/storage replication are common with 1U and blade servers becoming common. This solution is very capital intensive and very problematic on a TCO cost basis. Adding additional servers is a rational choice when the application is compute bound.
- **Buy big cache-oriented enterprise RAID.** Migration to high drive-count RAID with large shared caches is common with large servers. The storage cost is high and may be inefficient in terms of use of RAID capacity.

While many users take one or both of these paths to scalability, IDC believes that many would be better served and benefit from a thorough performance analysis that includes a review of SSD capabilities.

Table 2
CPU Cycles Consumed in an I/O Wait, 1995–2005

	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005
Average disk access latency (ms)	10.0	9.0	8.0	7.5	7.0	6.0	5.4	5.0	4.6	4.2	4.0
Number of CPU cycles in I/O wait per access (M)	1.0	1.5	1.6	2.1	2.5	7.2	8.1	12.5	20.7	25.2	32.0
Average SSD access latency (ms)	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10
Number of CPU cycles in I/O wait per access (M)	0.01	0.02	0.02	0.03	0.04	0.12	0.15	0.25	0.45	0.60	0.80
Wasted CPU cycles recovered (%)	99.0	98.9	98.8	98.7	98.6	98.3	98.1	98.0	97.8	97.6	97.5

Key Assumption:

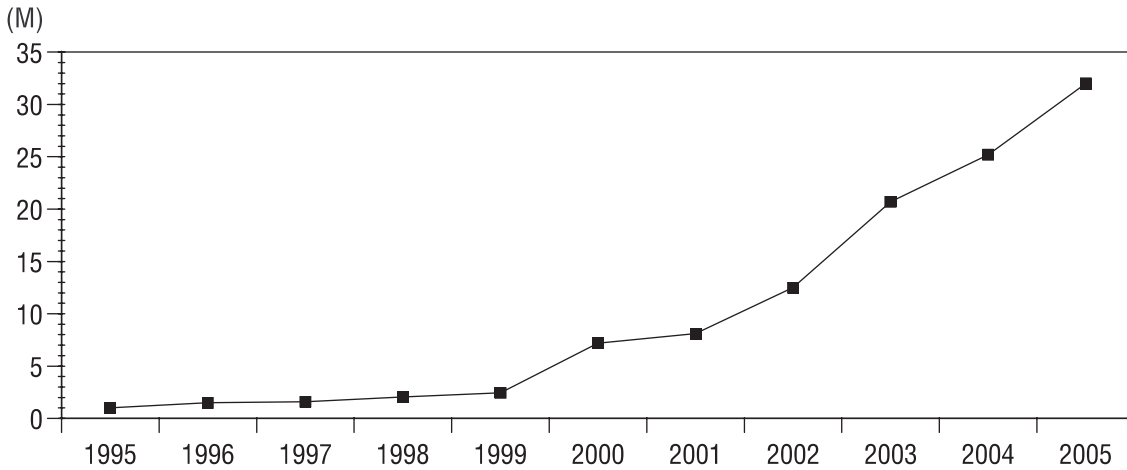
- SSD access is 100 microseconds (0.1 millisecond).

Message in the Data:

- Even with long access times, SSD recovers 97% or more of potential lost CPU cycles.

Source: IDC, 2001

Figure 2
Potential CPU Cycles Lost Waiting for a Disk Transaction, 1995–2005



Source: IDC, 2001

Imperial

Imperial Technology Inc. (El Segundo, California) was founded in 1974. Imperial is a privately held, pure-play SSD supplier. The company supplies what it brands as file cache accelerators and intelligent caching systems. Imperial SSDs support multiple operating systems, interoperate with numerous other hardware suppliers, and are certified with major software applications. The company markets globally through OEMs, VARs, and direct sales channels. Specific value-adds include:

- Scalable Fibre Channel and/or SCSI interconnects
- Flexible single SSD to multiple shared server capability
- Embedded Reed-Solomon ECC
- Broad product range of nonvolatile solutions
- Onsite and extended warranty programs
- 2Gbps Fibre and 160MB/s SCSI capabilities
- Flexible and scalable multiserver and SAN connectivity
- Multiport capability for highly redundant environments
- Remote call-home monitoring and administration
- Minimon performance analysis software tools

Imperial SSDs are branded MegaRam. They are both 3.5in. format and rackmount and offer a mixture of up to eight Fibre Channel ports and/or sixteen SCSI ports using a modular design. High-availability fans, power supplies, and batteries are redundant and

hot-swappable. The backup drive is also hot swappable. Capacities range from 134MB to 52GB. Key application markets include:

- Telecommunications and wireless messaging
- eCommerce
- Financial applications, including trading, banking, and insurance
- Email and Usenet newsgroup services
- Fraud detection, authentication, and security
- Biomedical engineering and imaging

The company lists among its customers: eBay, PacBell, EarthLink, AOL, Bank of America, Raytheon, AT&T, Intel, Cingular Wireless, Boeing, Southwest Airlines, and IBM.

Application Profile: eBay (Supplied by Imperial Technology)

Enormous cost savings that can be realized using SSD not as the performance “patch,” but as a central element essential to role of the SAN. One such company implementing this forward-thinking approach is eBay (NASDAQ: EBAY), the world’s leading online trading community.

Each day, eBay’s site draws over 2.1 million unique visitors, averages over 100 million page hits, and executes over 14.8 million searches. The site — even when coupled with 150 front-end servers, multiple Sun E4500 and E10000 servers utilized for searches, and eleven A3500 storage arrays — began to get taxed by the heavy use. An ever-increasing user base continually pushed the envelope on both hardware performance and possible upgrade paths. Something faster was necessary to continue to support the millions of searches per day, and hopefully at a lower cost than adding more top-end servers.

eBay estimated that 75% of hardware costs were related to supporting the search activity. The performance problem led eBay to Imperial. Imperial’s MegaRam SSDs gave the site the needed performance boost necessary to handle the searches. SSD also permitted eBay to defer server upgrades and bought the company time to organize an architectural redesign of the system.

eBay’s architecture at the end of 1999 included two parts: The first part was a front-end “fetch” system that consisted of 32 Sun E4500 servers, each with a copy of the user database. The second part was the “search” system that included two fully loaded Sun E10000 servers. Attached to these backend servers were 11 Sun A3500 storage units, each connected to mirrored Imperial MegaRam SSDs.

eBay wanted a better architecture to continue its growth. The key was to implement a system that would allow them to scale horizontally rather than vertically. eBay designed a SAN architecture designed around Imperial’s MegaRam systems — using the SSD as a proactive component, rather than as a reactionary performance booster.

The new SAN architecture places Imperial's native Fibre Channel MegaRam-5000 systems at the heart of the SAN, between the front-end and back-end servers. The multiple front-end systems connect directly to Brocade fabric switches, which are attached to the MegaRam SSDs. The backend servers, in turn, are attached directly to the native Fibre Channel MegaRam devices.

This implementation was envisioned to allow eBay to replace the 150 E4500 front-end servers with Zeus Web servers and allow the Sun E10000 servers to be replaced with Sun Blade 1000s. The savings in hardware costs, licensing fees, and colocation floor space further serve as cost-reduction measures.

In addition to cost savings, performance modeling indicated downtime would be halved and the back-end indexing would be done 40% faster.

"The increase in indexing performance alone justified the cost of the MegaRam (SSD) systems," stated the customer.

Forecast and Analysis

This bulletin does not contain a specific forecast of SSD shipments. There are a number of factors that will impact actual customer deployment of SSD (market penetration) over the next five years. The key issues are discussed in the following sections.

DRAM Core Technology Lowering Price Points

The result of lower DRAM prices will be a lowering of the entry price point for SSD technology and a reduction in the cost/GB for all capacity levels.

The core technology of SSDs is basic DRAM chips. The good news is that the cost of DRAMs has been declining rapidly and is forecast to further decline. The result will be a lowering of the entry price point for SSD technology and a reduction in the cost/GB for all capacity levels.

Figure 3 charts the average selling price (ASP) for DRAM chips from 1995 through 2005. From 2001 through 2005, the cost for a gigabyte of DRAM is forecast to decline to one-third of 2001's level.

Lack of Storage Management Staff and Awareness of SSD

We believe the greatest barrier to SSD adoption is the lack of end-user awareness.

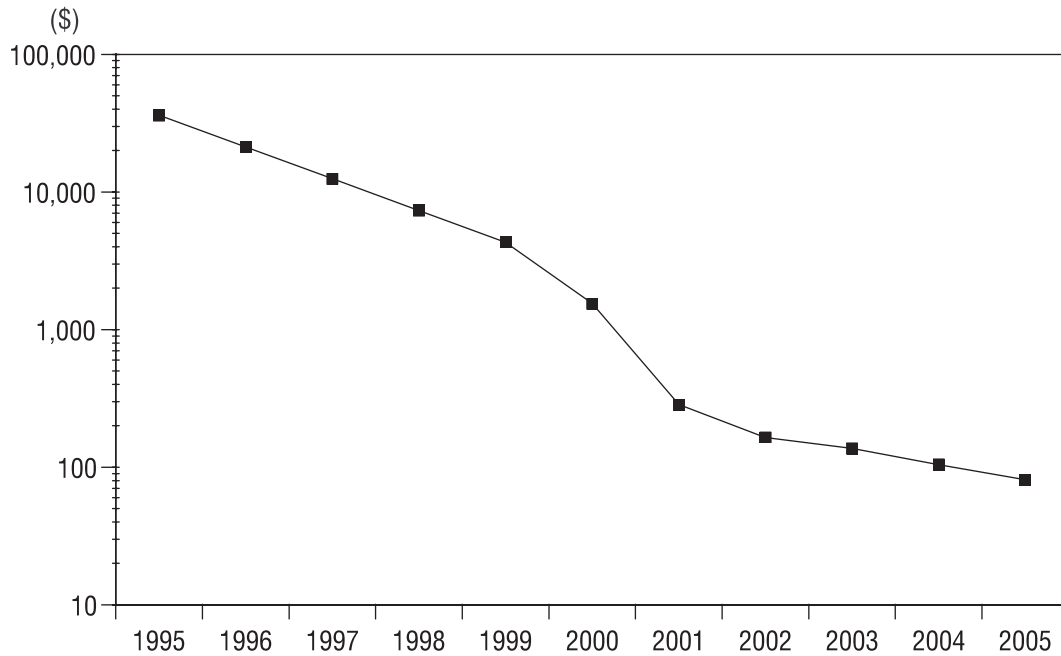
We believe the greatest barrier to SSD adoption is the lack of end-user awareness. Buried under the buzz of storage consolidation (NAS and SAN) and data protection (backup and recovery) issues, end users are, when faced with performance issues, not inclined to seek innovative technology and are more likely to deploy more of what has worked in the past. More servers and more disk array storage is the typical prescription. SSD suppliers need to support market development programs focusing on organizations, consultants, forums, software, and others that are responsible for system performance management and measurement.

Partnering with Vertical Application Suppliers

We believe an essential element of SSD marketing is partnership and joint marketing of SSDs as performance accelerators for specific

applications. Sales and support organizations selling and installing high I/O rate applications need to be aware of SSDs as performance solutions. They also need ready access to recommend specific suppliers and products. Referenceable success stories and collateral marketing materials are also suggested.

Figure 3
DRAM Price per Gigabyte, 1995–2005



Source: IDC, 2001

SAN Virtualization Layer as Inflection Point

The emergence of the SAN virtualization layer and virtualization products will be an important enabler for SSDs.

The emergence of the storage area network (SAN) virtualization layer and virtualization products will be an important enabler for SSDs. SAN-based storage will become 46.7% of the market by the end of 2005. The SAN virtualization management layer will become an ever-present element of SANs in that period. The virtualization layer provides two values: access control and performance management. Both will evolve from manual services to policy-based automation.

The opportunity for SSD suppliers (and SAN virtualization suppliers, too!) is to partner so that products such as so-called SAN appliances come with SSD support. By SSD support, we mean that performance monitoring automatically identifies files for SSD caching or acceleration, and that policy-based data migration to SSDs also be provided for.

The SAN virtualization layer evolution and penetration is a powerful opportunity to have SSDs finally achieve their long-delayed potential.

Conclusion

SSDs will be the smart choice for deployment for increasing numbers of end users. Existing applications with performance issues are immediate targets for SSD evaluation.

New transaction-oriented system deployments would be well served by architecting in performance-measurement tools and SSD technology from the beginning. In many instances, substantial savings may occur from reduced server and storage array requirements.