

➔➔➔ **New Monthly Pass!** [Sign up now](#) and get online access for as low as \$5.95/month!



[Market Watch]

Solid State Storage

Mark Weitz
InstantDoc #25334
July 2002



Is an SSD right for you?

Disk subsystems often are the slowest components in workstations and servers. This performance disparity has become even more apparent as processor and memory speeds have escalated quickly, whereas disk subsystem speeds have increased at a much slower pace. For I/O-intensive applications (e.g., video editing, Web services, messaging systems, database management), relatively slow disk storage severely impacts overall application performance and system scalability. For example, if your application uses only 25 percent of your 8-way server's CPU capacity but I/O requests queue up, you aren't using your expensive server effectively.

Although improvements in magnetic-disk technology have resulted in dramatic capacity increases, I/O performance has improved at a much slower rate. Caching I/O controllers, RAID arrays, and caching algorithms in Windows and database

management systems (DBMSs) improve data-access speeds. Because caches hold much less data than disk subsystems do, the caching algorithms attempt to predict which data blocks the application needs and caches only those blocks. In an effective cache, the correct blocks are present most of the time. But, depending on the locality of the application's read requests, the blocks might reside in the cache for only a few seconds before they're replaced. If the blocks are bumped from the cache, the next time the application requests them, the application must wait to retrieve the blocks from the hard disk or array. And if the requested blocks aren't sequentially located on the disks, the performance hit can be substantial.

Obviously, you can tune application performance to minimize system I/O and improve overall performance. But tuning application performance can be a time-consuming process and might not result in the performance gains you're looking for.

Solid state disk (SSD) products improve overall application performance by permanently storing the most frequently accessed data blocks (typically entire files or groups of files) on the solid state storage device. Systems administrators can identify which files are responsible for the highest percentage of I/O requests, then put those files on the SSD. Access times for the blocks that make up those files then can be hundreds of times faster than similar requests from rotating media. And access times for randomly requested blocks stored on the SSD are just as fast as access times for sequential-block requests.

High SSD prices have limited their use to military and other high-priority applications. However, the recent decline in memory costs has sparked renewed interest in SSD products for mission-critical commercial applications.

On a per-gigabyte basis, SSD prices have fallen dramatically but are still much higher than rotating-disk prices. SSD capacities range from 512MB to 399GB; the least-expensive configurations sell for about \$1000, and high-end systems sell for \$400,000 or more. An SSD large enough to store an entire application and data files could be prohibitively expensive, so systems administrators typically use SSDs to read and write to frequently accessed files (sometimes referred to as *hot files*). Hot files are typically 1 to 5 percent of a database application.

Before you decide to purchase an SSD, you must determine whether your performance problems are I/O-related. Windows Performance Monitor can help with that assessment. In addition, some SSD vendors offer tools that provide a good overview of system performance. If I/O is indeed creating a performance bottleneck, check your CPU utilization. You can use an SSD to boost I/O performance (which will result in overall application-performance improvement), but only if your CPUs can accept the additional workload. Depending on your application, SSDs can improve overall performance 2 to 10 times.

Entry- and Mid-Level SSD Products

Vendors most commonly sell SSD products in three form factors: PCI expansion cards, directly attachable units the size and shape of 2.5" or 3.5" disk drives, or standalone 19" rack-mount products that you can connect directly to one or more servers or use in a Storage Area Network (SAN) environment. For the largest enterprise applications, vendors offer several capacities: PCI expansion cards typically hold up to 8GB of data, the directly attached drive-shaped units I looked at provide up to 77GB of storage, and rack-mounted products support hundreds of gigabytes of storage. Windows sees these products as standard hard disks. I can't cover all the SSD products that more than 15 vendors supply, so I discuss a few products targeted to the Windows audience.

Most SSD products I looked at use SDRAM and therefore require backup power to maintain their contents during commercial power interruptions or, for SSDs installed inside a server or workstation, during system shutdowns. Data-protection features, as well as the number and type of built-in I/O interfaces, are major differentiators that make some products better suited for particular applications and contribute heavily to the price. The interface technology matters even more with SSD products than with rotating disks—because SSD products provide so much bandwidth, in high-volume applications the interface could become a bottleneck. But regardless of these additional features, SSD prices vary widely, so consider several vendor offerings before you make a decision.

PCI card products are the least-expensive SSDs; pricing for Platypus Technology's *QikDRIVE8* starts at \$1470 for 1GB of SDRAM, and CENATEK's *RocketDrive Solid State Disk* starts at \$999 for 512MB of SDRAM. These two products provide external power supplies that protect data when you power down your computer, but they don't provide redundant power sources to protect data during a power outage or a failure of the cards' external power adapters while the computer is shut down. The UPS that keeps your servers alive can also maintain data in the SSD, but unless the UPS software can copy data from the card to a hard disk before the UPS battery dies, you could lose data if an extended power outage occurs after business hours. Losing data might not be a concern when you use the SSD to improve access to a terminal server's swap file, but you probably don't want to use these types of SSDs for files that contain critical data. At press time, CENATEK said it will soon introduce a new product called *RocketDrive XL* that will incorporate a battery, automatic backup software, and a hard disk. CENATEK hasn't set pricing for the product.

Most of the more expensive SSD products provide built-in UPSs and hard disks to protect data. If power fails, batteries maintain the data while the system copies it to the SSD's hard disks. Imperial Technology's *MegaRam-35* provides this technology; its disk, memory, battery, and Ultra Wide SCSI interface components fit into a 3.5" disk bay. This added protection increases the product's price, however; a 1GB internal configuration sells for about \$10,000. The *MegaRam-35*, which supports up to 4GB of SDRAM, is also available in an external configuration. Surprisingly, this high-speed product (which is relatively expensive) is equipped with a 40MBps SCSI Ultra Wide interface rather than an Ultra160 SCSI or Fibre Channel interface. But even with that interface, Imperial claims the *MegaRam-35* can perform close to 10,000 I/O operations per second by using a 4KB frame size—much faster throughput than you'd expect from a large array of rotating disks.

Platypus's *QikDATA* provides similar features in a 1U (1.75") rack-mount chassis that holds 2GB to 16GB of SDRAM, a UPS, and two mirrored hard disks that receive the SDRAM's stored information when power fails. *QikDATA* connects to a workstation or server through the vendor's proprietary 64-bit 66MHz PCI interface card. Platypus says its proprietary interface delivers all the bandwidth your PCI bus can handle; with a 64-bit PCI slot, this is substantially more bandwidth than an Ultra160 SCSI interface can deliver. A *QikDATA* SSD configured with 2GB of SDRAM sells for about \$20,410.

BiTMICRO Networks takes a different approach with its *E-Disk*, which, like *MegaRam-35*, fits into a 3.5" disk bay. The *E-Disk* uses only a small amount of SDRAM as a cache; its primary storage medium is flash memory, which doesn't

require power to maintain stored information. During a power failure, the E-Disk's small battery refreshes the SDRAM while it copies stored data to as much as 77GB of flash memory. Although flash memory's ability to store data without power is a benefit in an SSD, flash memory also has potential limitations: Its typical life span of 100,000 to 300,000 erase/write cycles makes it inappropriate for enterprise-level transaction-processing applications. However, BitMICRO claims to have overcome this limitation (for the company's explanation, see http://www.bitmicro.com/BitMICRO_E-Disk_Features.pdf). You can purchase the E-Disk with SCSI Narrow, SCSI Ultra Wide, or Fiber Channel interfaces. E-Disk with 1GB of flash memory and a SCSI Ultra Wide interface costs about \$1750.

Sharing SSD

The products I discuss in this article connect directly to a workstation or server, but you can share some SSDs—either through direct connections to your servers or through a SAN—to improve the performance of applications that run on several servers. Texas Memory Systems, Imperial, BitMICRO, and other vendors offer products you can use in a SAN environment. These vendors say they've tested their products with several Fibre Channel switch fabrics and host bus adapters (HBAs), but you'll need to contact them to confirm compatibility with your specific configuration.

As with the directly attached MegaRam-35 and QikDATA SSDs, shared-SSD products typically provide built-in batteries and hard disks to protect data, but they also include redundant batteries or power supplies as well as fans to ensure reliability. These 19"-wide rack-mount units also let systems administrators monitor system performance, perform diagnostics, and check battery and fan status and power-supply temperature from a remote location.

In addition to these features, Texas Memory Systems' RamSan-210 can write data to SDRAM and both of its internal disk drives in realtime; it reads data only from SDRAM. The RamSan-210 supports 8GB to 32GB of SDRAM in its 2U (3.5") chassis (an 8GB configuration starts at about \$36,000) and features four Fibre Channel ports you can connect to several Fibre Channel switches for redundancy. For companies that have Network Attached Storage (NAS) systems but haven't implemented SANs, Texas Memory Systems offers a similar model, the NAS-250, that it markets as a storage accelerator for NAS systems. You can purchase the NAS-250 with two Gigabit Ethernet ports, and it supports Server Message Block (SMB) and NFS file systems so that you can access it from Windows and UNIX clients. Prices for NAS-250 start at \$39,000 for 8GB of SDRAM and one Gigabit Ethernet adapter.

You can configure Imperial's MegaRam-2000 and MegaRam-5000 models with a combination of Ultra160 SCSI and/or 1GBps or 2GBps Fibre Channel interfaces for use in clustered or SAN environments. You can connect the MegaRam-2000 to two redundant Fibre Channel switches or up to four servers with Ultra160 SCSI interfaces. Its 2U (3.5") rack enclosure supports up to 38.5GB of SDRAM (a configuration with two 2GBps Fibre Channel interfaces and 1GB of SDRAM sells for about \$20,000). The MegaRAM-5000 supports up to 8 Fibre Channel connections or up to 16 Ultra160 SCSI interfaces, or a combination of the two. Its 5U (8.75") rack-mount chassis holds up to 51GB of SDRAM. MegaRam-5000 configured with 2GB of SDRAM and two Fibre Channel ports sells for about \$48,000.

At press time, BitMICRO was preparing to introduce the E-Disk SAN, which uses up to fifteen 26.6GB E-Disk flash-memory-based Fibre Channel disks and includes four 2GBps Fibre Channel ports. The company hadn't yet set pricing.

Should You Buy an SSD?

The lowest-priced SSD products—those you install in a system's PCI slot—are fairly inexpensive. They're suitable for server applications in which data persistence isn't critical, such as improving swap-file access for Citrix MetaFrame or Windows Terminal Servers, enabling them to support more users. These SSDs are also suitable for desktop systems of engineers, artists, or other creative employees who frequently load and work with particular files. But if these servers or workstations have more memory than you need, you might consider buying RAM disk software, which can improve I/O performance for a low price. (CENATEK sells its RAMDisk NT 4.0/2000 software for just \$35, and SuperSpeed Software sells RamDisk-NT Workstation and RamDisk-2000 Professional for \$89.95. RamDisk XP Professional sells for \$79.95.) Considering SDRAM's relatively low cost, adding system RAM to a desktop PC so that you can use RAM disk software instead of buying an SSD might make sense. The end result won't be quite as fast as an SSD because of the RAM disk's application overhead, but it still should be much faster than a mechanical hard disk. Either way, don't forget to add a small UPS (CENATEK says its RAMDisk product can save its content to a hard disk at regular intervals). For workstations and servers, SuperSpeed also offers the SuperSpeed NT and SuperSpeed 2000 RAM disk products, which it claims can write RAM disk content changes to the system's hard disks as often as several times per second. SuperSpeed NT Workstation and SuperSpeed 2000 Professional sell for \$149. Single-processor versions of SuperSpeed NT and SuperSpeed 2000 sell for \$399 each; two- and four-processor versions sell for \$799 and \$1499, respectively. The company also offers SuperSpeed 2000 Advanced Server (for Windows 2000 Advanced Server) and SuperSpeed NT Server Enterprise Edition (for Windows NT Server, Enterprise Edition—NTS/E); call SuperSpeed to get pricing for your specific configuration.

Although full-featured SSDs are much less expensive than they were a year ago, they still cost about as much as a

new server. The right choice for you will depend on a careful examination of your application's I/O characteristics, your I/O growth expectations for the next few years, the number of servers affected, and whether the servers are due for replacement soon. If you need to improve the performance of just one server and you're sure that your disk subsystem is the limiting factor, test the server with 2GB of additional memory before buying an SSD. The additional RAM could make the application and OS run more efficiently, and additional memory allocated for caching purposes could provide the performance improvement you need. If time is more plentiful than money, performance tuning can also deliver significant performance improvements.

If you're using four mail servers for performance reasons and all of them are I/O-limited, sharing an SSD between two of them could deliver the needed performance with I/O bandwidth to spare, saving you time and the cost of managing superfluous servers. You can avoid buying additional servers by allocating the unused servers to other tasks. In a SAN environment with many servers, SSDs could be even more helpful, but again, careful assessment of server bottlenecks and application I/O characteristics is in order.

CONTACT THE VENDORS

E-DISK, E-DISK SAN

BiTMICRO Networks * 510-743-3475 * <http://www.bitmicro.com>

ROCKET DRIVE SOLID STATE DISK, RAMDISK NT 4.0/2000, ROCKETDRIVE XL

CENATEK * 408-782-1220 * <http://www.cenatek.com>

MEGARAM-35, MEGARAM-2000, MEGARAM-5000

Imperial Technology * 310-536-0018 or 800-451-0666

<http://www.imperialtech.com>

QIKDRIVE8 AND QIKDATA

Platypus Technology * 603-298-7455 or 877-718-8900

<http://www.platypustechnology.com>

SUPERSPEED NT WORKSTATION, SUPERSPEED 2000 PROFESSIONAL, RAMDISK 2000 PROFESSIONAL, SUPERSPEED 2000 ADVANCED SERVER, SUPERSPEED NT SERVER ENTERPRISE EDITION, RAMDISK XP PROFESSIONAL, RAMDISK NT WORKSTATION

SuperSpeed Software * 978-443-5106 or 800-388-8872 <http://www.superspeed.com>

RAMSAN-210, NAS-250

Texas Memory Systems * 713-266-3200

<http://www.texmemsys.com>

Other SSD Vendors

Adtron * 602-735-0300 * <http://www.adtron.com>

Curtis * 763-404-9081 or 800-245-3171

<http://www.mncurtis.com>

DataDirect Networks * 818-700-7600

<http://www.datadirectnetworks.com>

Memtech SSD * 925-294-8483

<http://www.memtech.com>

M-Systems * 510-494-2090

<http://www.m-sys.com>

MTI * 714-970-0300 * <http://www.mti.com>

SanDisk * 408-542-0500 * <http://www.sandisk.com>

Solid Data Systems * 408-845-5700 or

800-287-0373 * <http://www.soliddata.com>

Targa Systems Division * 704-708-4720 <http://www.targaelectronics.com>

Winchester Systems * 781-265-0200 or

800-325-3700 * <http://www.winsys.com>

CENATEK Rocket Drive Solid State Disk

BiTMICRO Networks E-Disk