Managing the information that drives the enterprise

STORAGE

ESSENTIAL GUIDE TO

Solid-State Drives

Solid-state drive technologies that will have the most impact for enterprise data storage administrators are highlighted in this guide.

INSIDE

TechTarget
The Technology Media ROI Experts

Ask 7 smart questions to ensure your next solid state disk is a flash of inspiration and not a regrettable flash in the pan!

The World's Fastest No Regrets SSDs

RamSan-620 SSD — SAN, NAS & Direct Attach

+ 1 to 5 TB SLC Flash
+ Reads & Writes 250,000 IOPS
+ Throughput up to 3GB/s
+ 230 Watts
+ Fibre Channel

RamSan-20 SSD — Workstations, Servers

+ 225 to 450 GB SLC Flash
+ Reads/Writes 120,000/50,000 IOPS
+ Throughput up to 700MB/s
+ 15 Watts
+ PCIe

>> Online Video >> Auto Price Quote

With so many new solid state disks (SSD) on the market, how do you choose one that will be fast and reliable today, and many years from now? You can start by asking your prospective vendors these 7 smart questions:

- ☐ Are the performance specifications based on burst or sustained data rates and how much of the capacity is actually available to me?
- ☐ Who has independently confirmed the performance claims?
- ☐ Is the SSD built with long life enterprise grade SLC Flash or MLC Flash that wears out quickly?
- ☐ Does the SSD have onboard controllers or does it need to tap the resources of my server?
- ☐ Does the SSD have onboard ultra capacitors for near instant recovery after a power outage or could recovery take hours?
- ☐ When did your firm ship its first production SSD and what generation is your current SSD?
- ☐ Can your firm provide a RAM-based SSD, PCIe and SAN-based SSDs when our needs change?

At Texas Memory Systems we'd love to answer those questions for you. Because when you compare the details you'll understand why the world's largest SSD installations, including major financial exchanges and 7 of the world's largest telecoms, rely on our award winning RamSan SSDs – the "World's Fastest Storage."® For over **30 years** RamSan SSDs are the first choice for speed and reliability. They've always been the *World's Fastest No Regrets SSDs.*

TexasMemorySystems

White Papers and Self-Service Quotes | 1-713-266-3200 | TxMemSys.com/flash editorial | rich castagna

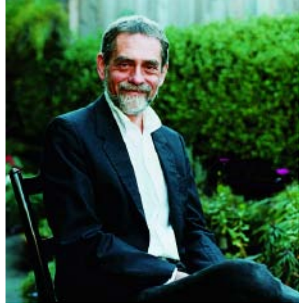

The state of storage

Solid-state storage is poised to revolutionize data center storage with new and inventive implementations of NAND flash already available.

IT'S SMALL, it runs cool, it's power efficient and it's lightning fast. There's an awful lot to like about solid-state storage. Put it head-to-head with traditional spinning media and it wins on nearly all counts, with admittedly some not yet fully resolved issues related to long-term reliability and price. Still, few would question that solid-state is the direction that enterprise storage is headed.

The real question is just how soon the transition from mechanical storage to solid state will take place.

The real question is just how soon the transition from mechanical storage to solid state will take place.

Although new to the data center, NAND flash-based solid-state storage has been around for a long time, used mostly in consumer and mobile products like smart phones, PDAs and MP3 players. The transition to data center quality storage products involves more than just a matter of scale as the technology employed in mobile products isn't durable enough to stand up to the rigors of enterprise computing environments. But even at this early stage in its development, solid-state storage does have something of a track record in enterprise environments—demonstrated perhaps most notably by Texas Memory Systems that has been marketing DRAM-based storage arrays for more than decade. In recent years, the company has also added flash-based systems to its roster.

But NAND flash has been criticized as being too unreliable to be used for enterprise-class storage. It can handle a limited number writes before it simply wears out, and its write performance lags its read performance considerably. Those issues, however, are being actively addressed by scores of solid-state chip makers and integrators with of novel approaches that yield reliability and performance improvements. Already, an impressive array of new NAND flash products is available—products that clearly exceed the performance of hard disks and match their reliability as well.

The other criticism of NAND flash storage—and the one that's expressed most frequently—is that it just costs too much. If you compare NAND with hard disk storage on a dollar per gigabyte basis, the cost critics win the argument hands down. But that kind of comparison really takes both types of storage out of context and also evades the many positive attributes of solid-state storage. Proponents of solid state say that the only valid means of comparison is on a performance per dollar basis. That argument is simple and very convincing. For example, a high-performance application may require dozens of high-speed disks and use only the outer parts of those disks to ensure quick response times. So, a lot of relatively expensive high-speed disks will be used only partially to satisfy the app's needs. A solid-state alternative would require far fewer units and all the capacity on each would be available to the application—and it has an inherent speed advantage from the get go. So which costs more—a lot of pricey spinning disks or a handful of NAND flash drives?

And even if the price was equal, solid state provides enormous advantages over spinning disk in power consumption, heat production and footprint. In some shops, those savings alone are enough to just the move to solid state. The goal for any storage manager is to have storage systems that are reliable, provide sufficient performance, can be easily maintained and are cost effective. ⊙

Rich Castagna (rcastagna@storagemagazine.com) is editorial director of the Storage Media Group.

Experiencing Performance Problems?

Atrato's Application Acceleration delivers real-time management for your high velocity data.

Optimize your storage environment. Atrato's hybrid solid state solution provides the best of both worlds: unmatched performance for I/O hungry applications and cost effective media for less accessed data. Intelligent autonomic tiering places your high velocity data in the correct tier, while eliminating management complexity. Finally, storage that manages itself.

Industry's Highest IOP

Experience the highest available single rack random I/O performance.

Hybrid VLUN

Accelerate each application with SDD and HDD Hybrid VLUNs.

Autonomic Tiered Management

Automatically optimize the storage environment without user intervention.

Interested in storage tiering without complexity.
Contact us today at sales@atrato.com

ATRATO™
DATA MANAGED STORAGE

SSD trends IN ENTERPRISE STORAGE

Here's a look at the latest trends in how solid-state drives are being used in enterprise data storage environments.

By Carol Sliwa

SOLID-STATE DRIVE (SSD) technology has been in use for decades, but it was hardly a factor in enterprise storage circles until approximately 18 months ago, when the price of flash memory started to drop dramatically.

At the present cost of approximately $4 per gigabyte, solid-state drives often make sense only for those applications with the highest performance needs. Although nearly all server and storage vendors have added solid-states drives to their lineups, fewer than 70,000 SSD units shipped last year and fewer than 300,000 will likely ship this year, according to Joseph Unsworth, research director at Stamford, Conn.-based Gartner Inc.

"There's still a lot of challenges out there," Unsworth said. "You need to have the right SSD, properly managed, supported by the right company. Enterprise servers and storage have been architected around hard drives for the last 20, 30, 40 years. All of a sudden you're going to put in a solid-state drive that can access data 100 times faster than even the fastest hard drive? Believe me, these systems aren't architected to take advantage of that and exploit that."

SOLID-STATE DISK VS. FIBRE CHANNEL

Still, the allure becomes apparent even when looking at the test results of one early adopter. Munder Capital Management in Birmingham, Mich., reduced data access time for its Microsoft Corp. Exchange Server by 89% on average during a 25-day test period of solid-state drives vs. 15K rpm Fibre Channel drives in its Compellent Technologies Inc. Storage Center system. The response time for SQL Server plummeted by an average of 86%.

Munder tested 4K, 8K, 16K, 32K and 64K I/O block sizes with 25 different workloads over a 24 hour period, varying between reading and writing, sequential and random. It claimed the performance of one SSD was equivalent to the performance of 10 to 16 Fibre Channel hard disk drives (HDDs), depending on the workload.

Such a performance boost is especially appealing for mission-critical trading applications that need to process large amounts of data and crunch numbers at lightning-quick speed.

Such a performance boost is especially appealing for mission-critical trading applications that need to process large amounts of data and crunch numbers at lightning-quick speed. Munder also plans to use SSDs for Microsoft SharePoint Server and Exchange Server.

"Right now, that's the best bang for the buck," said Michael Dufek, Munder Capital Management's director of information systems. Until the cost of SSD drops, he added, "It's not geared for everything."

With the Compellent system, the cost of a single 15K 300 GB Fibre Channel HDD is $2,100; a single 146 GB SSD is $24,000, although savings in power consumption and software licensing fees help to offset a portion of the cost differential.

The SSD that Munder Capital Management tested in its Compellent array is a type of flash known as single-level cell (SLC)—the technology currently favored by most vendors that produce storage systems aimed at enterprise-level needs.

DIFFERENT TYPES OF SSD TECHNOLOGY FOR ENTERPRISE DATA STORAGE

While users such as Munder can take advantage of SSDs today, it's still early days for the technology in enterprise storage. The technology is still evolving, with different types of solid-state drives available—each with its advantages and drawbacks.

The first type of SSD technology to gain any noteworthy foothold in the enterprise was Dynamic Random Access Memory (DRAM) SSD. Texas Memory Systems Inc. is a leading vendor of DRAM SSD. But the technology generally has been cost-prohibitive for all but those companies whose businesses depend on application performance for success.

The main distinguishing characteristic of DRAM SSD is its volatility. If the power supply is removed, data is lost. DRAM makes use of capacitors, but because the capacitors leak charge over time, they must be refreshed continually to preserve the data. To address those issues, manufacturers typically build-in an alternate power source, such as a back-up internal battery, and the ability to move data to an HDD or a flash-based memory system in the event of a power failure.

Dynamic Random Access Memory SSD's main selling point is performance, and it remains the fastest technology for accessing data and performs reads and writes at approximately the same speed. On the downside, DRAM SSD is expensive and has greater power needs, creating an opening for the cheaper flash-based, energy-efficient SSD technology to gain momentum in the enterprise.

NAND FLASH MEMORY

Unlike DRAM, flash memory is non-volatile, so the storage cells retain data whether the power is on or off. Flash SSDs have floating-gate transistors that can store charge for an extended period of time, even if there's no power supply connected. Oxide insulation surrounding the floating gate traps the electrons there.

Another contrast between flash DRAM SSDs is data access. DRAM is random. Flash is serial and relies on a controller to bring the data out of the chip and correctly present it to the processor.

There are two types of flash memory: NAND and NOR. NAND flash technology strings together floating-gate transistors to achieve greater density. NOR flash has no shared components, is more expensive to produce and is found mainly in consumer and embedded devices, primarily as a mechanism to boot them. NAND flash, by contrast, is used for data storage and is increasingly making inroads into the enterprise.

Unlike DRAM, flash memory is non-volatile, so the storage cells retain data whether the power is on or off.

What may increasingly become an important consideration for enterprise users is whether NAND flash is single-level cell or multi-level cell (MLC). In SLC NAND flash, each memory cell stores an analog representation of the data and two levels of charge. The less-costly MLC flash can store twice—as two or more bits per cell and multiple levels of charge—but it doesn't perform as well and is less reliable than SLC.

As a result, single-level cell is the dominant flash solid-state drive technology in enterprise storage systems today. Gartner's Unsworth estimated that SLC factors into more than 80% of enterprise-grade flash SSDs. But he predicted the breakdown could shift to 60% SLC and 40% MLC by 2011, if vendors sufficiently improve the controller technology and storage management software. MLC is better suited to read-intensive, not write-intensive applications, Unsworth added.

One of the distinguishing characteristics of NAND flash is its process for writing data to a chip. All of the bits in a flash block must be erased before any writing can take place. The term "flash" derives from the violent charge to erase the data. The block size is typically 128 KB to 256 KB, and will likely expand in the future.

NAND FLASH WEAR-OUT

Unfortunately, the erase/program process eventually takes a toll, breaking down the oxide layer that traps the electrons. The gradual deterioration can distort the manufacturer-set threshold value at which a charge is determined to be a zero or a one.

The deterioration is less a problem in SLC flash than it is in MLC flash. In SLC flash, there's only one manufacturer-set threshold value, so the likelihood of a problem is lower. The number of electrons controls the switch-on voltage of the floating gate, and the voltage will either be above the threshold point or below it.

With MLC flash, the manufacturer can set multiple threshold values. As the oxide layer deteriorates, those values can shift across the pre-set threshold points and become difficult to discern, leading to errors.

Both single-level cell and multi-level cell flash rely on error-correction algorithms to ensure the data remains intact, but eventually, NAND flash SSDs wear out. The wear-out figures typically used by the industry are 100,000 program/erase or endurance cycles for single-level cell flash and 10,000 cycles for multi-level cell flash, but those figures vary widely by manufacturer.

> "The flash is going to become a lot more unreliable, but the good news is that error-correction techniques in disk drives will be applied in SSDs, and that will make it more reliable for storing data."
>
> —Michael Cornwell, lead technologist for flash memory, Sun Microsystems Inc.

Michael Cornwell, lead technologist for flash memory at Sun Microsystems Inc., claimed the MLC endurance cycle figure has been worsening to a figure closer to 3,000 program/erase cycles. He said as flash dies get smaller, fewer electrons fit on the floating gate. That trend, coupled with the natural tendency of electrons to escape, will lead to the use of more sophisticated data correction.

"The flash is going to become a lot more unreliable, but the good news is that error-correction techniques in disk drives will be applied in SSDs, and that will make it more reliable for storing data," Cornwell said. He said most flash solid-state drives now use between 8 bits and 16 bits of correction per sector. A hard disk drive, by comparison, uses 500 bits of correction for the same amount of data.

As vendors try to reduce the cost of solid-state drives, they're working to improve MLC to make it more suitable for certain types of enterprise applications. Lower-cost MLC currently accounts for more than 90% of the overall worldwide NAND output, but only a small percentage of the enterprise space.

"There's a race among SSD makers to find a way to make it [MLC] work reliably in an enterprise SSD," said Jim Handy, a Los Gatos, Calif.-based analyst who focuses on memory chips and SSDs for the market research firm Objective Analysis.

"There's a race among SSD makers to find a way to make [MLC] work reliably in an enterprise SSD."

—Jim Handy, analyst, Objective Analysis

Issues such as wear-out and error correction don't factor into the equation with the more expensive DRAM solid-state drives, according to Handy. He said DRAM SSDs can be written to an infinite number of times and feature high-integrity memory.

Manufacturers are also exploring a variety of different mechanisms to improve solid-state drives, such as phase change and resistive memory, in addition to the classic approaches of shrinking the technology and packing more bits into cells.

"There are a dozen different implementations that may extend the life of flash technology as we know it," said Patrick Wilkison, vice president of marketing and business development at STEC Inc., which supplies SSDs to major storage vendors such as Dell Inc., EMC Corp., Hewlett-Packard (HP) Co., Hitachi Data Systems Corp., IBM Corp., NetApp Inc. and Sun Microsystems Inc.

In the meantime, users in need of high-performance storage need to weigh the benefits against the potential pitfalls. Munder Capital's Dufek said he doesn't worry about the solid-state drive wear-out factor because his data is backed up, his Compellent storage system monitors for errors and, if an SSD fails, the system can automatically move data from the SSD tier to a spare HDD. Dufek said that under the terms of his agreement with Compellent, Munder Capital would have to wait no more than four hours for a replacement SSD to arrive.

"With flash, the chances of a chip death are very small," said Robert Ober, an LSI Corp. fellow for business development and strategy, who's driving the company's solid-state drive efforts. Plus, SSD failures don't happen without warning, similar to the way a tire sensor can alert the owner when the tread is wearing out, he added.

THE FUTURE OF SOLID-STATE DRIVES

The term “tier 0” has come into vogue to describe solid-state drive storage, as interest in the technology heats up. Framingham, Mass.-based IDC predicts worldwide revenue for enterprise solid-state drive storage will increase at a compound annual growth rate of 78% from 2009 to 2012. IDC predicts that SSD sales will hit $382 million in 2010, up 132% over 2009.

But several things must happen for SSDs to take off in enterprise storage systems. First, the price must fall. Then storage system vendors must optimize their arrays and management software to work better with solid state.

Jeff Boles, senior analyst and director of validation services at Hopkinton, Mass.-based Taneja Group, said most of the arrays on the market currently support only a limited number of solid-state drive devices because the controller can rapidly become a bottleneck.

“You put a few SSD devices behind an array controller, and you can bottleneck because the array controller can’t keep up with the performance of the SSD devices,” he said. “To work around all the media issues and the array controller performance issues, the array vendors are pretty sensitive on qualifying specific media.”

They also need to make their storage management software more intelligent to exploit the benefits of solid-state drives and enable end users to see true returns on their investments, according to Gartner’s Unsworth.

“This is still a very, very nascent market that’s just starting to get its footing,” he said. “It does have a meaningful, impactful, disruptive role in enterprise storage, but companies need to be careful.” ⊙

Carol Sliwa is features editor for SearchStorage.com.

Wide stripe BEFORE YOU DIVE INTO SSD

Consider whether performance in your environment can be improved through wide striping data across existing disk arrays before investing in new solid-state drive technology.

By George Crump

SOME OF THE ADVANTAGES of short-stroking ensure that you're using only the fastest accessible parts of the drive, thus increasing performance. This also saves on rebuild time, as there'll be less data stored across those drives. This has been used effectively to speed performance in large, slower SATA drives. However, the devil is in the details. The largest single flaw in this practice is that a substantial amount of disk space on your drives will remain hijacked and unusable to your servers.

For example, assume you have a LUN that requires 2,500 IOPS. Using the lowest cost $/IOP disk drive, we'll configure this volume with 146 GB 15,000 rpm disks. To meet the desired performance, 6 RAID-1 sets will be created. Assuming the usable size needed is 200 GB, the actual usable capacity is 876 GB. This has only 23% of each disk being used to achieve the desired performance.

Enter wide striping. Vendors that have internal virtualization of their subsystems along with wide striping are ahead of the game. Because all of the disks are typically treated as one or as a few large storage pools, a user simply carves out the RAID type and LUN sizes from the overall aggregate; the LUN is then laid out across all of the spindles in the storage pool. Each disk should be able to support the following:

Portions of multiple volumes

Different QoS (RAID level, stripe characteristics, availability settings, inner/outer track service times)

Because these storage pools can have a very large number of drives, the performance for even typically slow SATA LUNs will perform well assuming they're spread across a large number of spindles. Now add wide striping to this process and you truly have something. Let's look at the same scenario as we had above. That 200 GB usable volume could be striped across 80 SATA drives and have approximately 8,000 IOPS available to it.

Before you buy SSD, check out wide striping

Now let's go in the other direction. Assume you're considering solid-state drive (SSD) technology for a high-performance application requiring a mix of read and write IOPS totaling 20,000. A much lower cost option is to wide stripe that volume on a pool of 15K SAS or Fibre Channel (FC) disks —120 of these drives can easily generate 27,000 IOPS and can be shared with other LUNs in the array needing lower performance.

Given that the performance for a wide-striped storage subsystem like 3PAR's is very high—approximately 225,000 IOPS according to the Storage Performance Council (SPC)—high IOPS per LUN can be most cost-effectively served by wide striping. Only if the application latency requires a sub-millisecond response does it then make sense to add some SSD into the environment.

If your application necessitates that you buy SSD, don't buy the first thing that comes along. Take the time to evaluate the available products in the SSD space, and spend some time planning exactly what data types need that performance level.

Know your flash memory facts

Most storage system suppliers use flash memory for their SSD strategy. While flash is fine for read-heavy applications, its performance degrades substantially on heavy random-write situations. Flash in a heavy random write isn't much better than standard mechanical drives, and wide-striping systems will perform better for substantially less money.

Keep an SSD future in mind when deciding on a storage system. Features to look for include the ability to allow for the virtualization of SSD into an existing storage system, as well as having automated processes that can move data in and out of SSD storage as needed. Most importantly, look for systems that provide robust performance metrics that can ensure you're utilizing the SSD to its fullest potential, guaranteeing the investment that was made in the technology. You should also have the ability to model the current storage subsystem to determine whether or not you need SSD technology. For example, determining how much or how little of it you can get away with and still provide optimum performance to customers.

The combination of wide striping, internal virtualization and SSD can offer some very attractive benefits. Just be sure to do your homework to ensure the technology you elect to go with is the most cost effective and can provide you with all of the features and functionality suitable to running the storage infrastructure at your organization. ⊙

George Crump is president and founder of Storage Switzerland, an IT analyst firm focused on the storage and virtualization segments.

Gamer's DRAM SSD leaves Fibre Channel in the dust

An online gaming company finds DRAM-based SSDs more effective for boosting performance than Fibre Channel hard disk drives. *By Carol Sliwa*

O**NE COMPANY** with a good vantage point of the merits of Dynamic Random Access Memory (DRAM) solid-state drive (SSD) and NAND flash is CCP hf, a Reykjavik, Iceland-based online game producer that uses both technologies.

CCP began using DRAM SSD technology from Texas Memory Systems Inc. nearly five years ago when it noticed its core product, Eve Online, showing signs of in-game lag due to the increasing stress put on some areas of its database and storage array.

The IT group now reserves its DRAM SSD-based RamSan-400 for the part of the database that's accessed the most and uses its newer "cached flash" RamSan-500—which has 64 GB of DRAM cache and 2 TB of RAID-protected single-level cell (SLC) NAND flash—for the bulk of the database calls. It credits solid-state drives with helping the system to handle 53,850 simultaneous users in mid-May.

"Over time, we have progressively moved away from Fibre Channel hard disk drive (HDD) storage, as the I/O performance just does not keep up with our demands," IT Director Jon-Carlos Mayes wrote in an email.

"To put it simply, for a database like ours, there is just no comparison."

Mayes described the DRAM SSD system's performance as "blindingly fast I/O" for both reads and writes, but at the expense of overall storage capacity. The 2 TB SLC flash-based RamSan-500 handles capacity needs and also provides "very fast read performance," Mayes wrote.

The DRAM solid-state drive-based RamSan-400 system claims to be able to attain 400,000 IOPS, for both reads and writes. Although CCP hasn't verified that number through testing, it appeared fairly accurate based on database usage and percentage usage of the RamSan device, according to Mayes. The RamSan-500, which has both DRAM and NAND flash technology, claims read performance of 100,000 IOPS and write performance of 25,000 IOPS, according to Texas Memory Systems.

The RamSan-500, which has both DRAM and NAND flash technology, claims read performance of 100,000 IOPS and write performance of 25,000 IOPS, according to Texas Memory Systems.

CCP is also testing a flash-only RamSan-20, which has 450 GB of SLC NAND flash storage attached via PCI Express. The RamSan-20 claims to handle 120,000 read IOPS vs. 50,000 write IOPS, illustrating the difference in performance for reads/write that users might expect to see in a dedicated SLC-based flash system.

For systems produced by Texas Memory Systems, the list price of DRAM SSD is $300 per gigabyte, while SLC-based flash SSD is $40 to $70 per gigabyte, according to Woody Hutsell, the company's president.

The 2 TB Flash-based RamSan-500, which has 32 GB of DRAM cache, lists at $150,000, whereas pricing for the all-Flash 5 TB RamSan-620 is $220,000, Hutsell said. The latest DRAM SSD-based 512 GB RamSan-440 lists at $180,000. Today's list price for the RamSan-400 is $61,000, while the RamSan-20 is $18,000.

Carol Sliwa is features editor for SearchStorage.com.

Low-cost MLC NAND flash gains in enterprise solid-state storage

New multi-level cell offerings from SSD vendors are an alternative to more expensive SLC drives for the enterprise.

By Carol Sliwa

NEW SOLID-STATE DRIVES (SSDs) coming onto the market will challenge the notion that only single-level cell (SLC) flash SSDs are fit for true enterprise applications.

Until recently, the conventional wisdom has been that only SLC has the performance and reliability required for the enterprise, while multi-level cell (MLC) NAND flash is more suited for consumer devices.

However, Fusion-io Inc. is planning a new class of enterprise-class NAND flash that it calls single mode level cell (SMLC) that could give customers faster and lower-cost PCI Express-based solid-state storage to plug directly into their servers. And STEC Inc., which supplies SSDs for most of the major storage array vendors, is sampling what it calls enterprise-ready MLC drives with capacities up to 800 GB to its OEM partners.

Now it's a question of how much performance organizations are willing to trade for the cheaper solid-state drive alternatives.

"The best way to take the cost out of enterprise SSD is to use MLC chips that are two to three times cheaper than SLC," wrote Joseph Unsworth, research director at Stamford, Conn.-based Gartner Inc., in an email interview. "The caveat is that the quality must be good enough to last in the enterprise SSD environment."

SLC STILL WRITES FASTER, MORE RELIABLE THAN MLC

The more economical MLC can store two or more bits per cell and multiple levels of charge, giving it at least twice the capacity of SLC. But with slower write speeds and higher bit error rates, MLC lacks the performance and reliability of SLC.

The industry's oft-used wear-out figures are 100,000 program/erase, or endurance, cycles for SLC flash and 10,000 cycles for MLC flash, although industry sources report the MLC figure actually has been dipping well below 10,000 as the flash dies get smaller.

The 160 GB and 320 GB SMLC ioDrive alternatives that Fusion-io plans to deliver this quarter could mitigate some of the MLC shortcomings. The SMLC product uses MLC flash chips, but Fusion-io claims its more sophisticated controllers deliver write speeds equal to SLC Flash and significant endurance improvements over MLC.

David Flynn, chief technical officer (CTO) at Fusion-io, said while Fusion-io's MLC drives could work in the enterprise, the new SMLC drives will give customers a third option alongside MLC and SLC.

"The best way to take the cost out of enterprise SSD is to use MLC chips that are two to three times cheaper than SLC. The caveat is that the quality must be good enough to last in the enterprise SSD environment."

—Joseph Unsworth, research director, Gartner Inc.

"Customers come in and say 'I need this much capacity, performance and endurance,' and they don't want to pay any more for any one of those than they need to," Flynn said. "Right now, they have a tough decision. Do you buy the more expensive SLC just to get the endurance when all you need is small capacity? Or do you buy the large MLC to get the endurance even though you didn't need that much capacity?"

In some ways, Fusion-io treats MLC like SLC. The SMLC product uses single thresholding, rather than multi-thresholding, in a way designed to extend endurance and performance, Flynn said. Lance Smith, Fusion-io's senior vice president of product marketing, noted that the SMLC offering maps a single bit of data to a single bit per cell, whereas MLC maps multiple bits of data to a single cell.

"We get five times the endurance out of MLC," Smith said.

The slower read speeds of Fusion-io's SMLC match those of its MLC technology, and the price tag falls squarely between MLC and SLC.

STEC claims its OEM customers have shown little interest in the hybrid SLC-MLC model, in which a single drive might include both SLC media for write purposes and MLC media to hold data.

"For those guys that are interested in MLC, they're looking for the dollar-per-gigabyte benefit, and SMLC, having half the capacity [of MLC], makes the cost of it two times," said Scott Stetzer, STEC's director of marketing for SSD programs.

STEC's OEM customers include Dell Inc., EMC Corp., Hewlett-Packard (HP) Co., Hitachi Data Systems, IBM Corp., NetApp Inc. and Sun Microsystems Inc. Its MLC offerings support the same SATA, SAS and Fibre Channel interfaces that its SLC technology does. Although STEC has now launched MLC products aimed at the enterprise, company execs say there will always be a large performance gulf between MLC and SLC.

> "MLC in the enterprise is here [and] here to stay. And it will continue to grow in volumes. But, for the foreseeable future, SLC will continue to dominate."
> —Scott Stetzer, director of marketing for SSD programs, STEC

"The interest, especially recently, is picking up," Stetzer said. "MLC in the enterprise is here [and] here to stay. And it will continue to grow in volumes." But, he added, "For the foreseeable future, SLC will continue to dominate."

FINDING A PLACE FOR MLC IN THE ENTERPRISE

One Fusion-io customer is already putting its MLC device to the test for enterprise use. Wine.com Inc. had such success as a beta customer of Fusion-io's 160 GB SLC ioDrives for its enterprise resource planning (ERP) databases that it tried its high-transaction storefront/shopping cart applications on solid-state storage last fall. Wine.com CTO Geoffrey Smalling said the database was too large for Fusion-io's SLC offering, so the online retailer tried the 320 GB MLC option.

Smalling admitted he was nervous after reading blog posts on MLC's 10,000 write/erase, or endurance, cycle limits. But he said Fusion-io's Flynn eased his concerns after explaining the startup's wear leveling algorithm, chip-level redundancy, error correction and system checks.

Smalling said he was especially impressed to hear that each PCIe card has 400 GB of raw capacity, leaving 80 GB free to be substituted when a storage block wears out. Fusion-io also builds in a RAID-like feature called flashback redundancy, which allows a spare flash chip to take over without loss of data in the event of a catastrophic chip failure.

Wine.com can consult Fusion-io's ioManager utility whenever its wants to view the wear-out status of its MLC ioDrives. Smalling said after eight months the chart shows that more than 99% of the MLC drive's capacity remains operational.

Even though Smalling knew MLC's read and write speeds would be slower than SLC, he figured the MLC option's 80,000 IOPS would still provide a "night and day" difference over its prior technology, a lower-end NetApp storage-area network (SAN). An unexpected side benefit was the ability to plug PCIe-based MLC ioDrives directly into its HP ProLiant DL380 G5 servers because neither Wine.com nor its hosting provider had great SAN expertise.

"To me, the SLC-MLC hasn't been a concern. It was more the transactional redundancy," Smalling said. "We took steps to put in RAID on the servers and mirror the servers." He added that the MLC product's five-year projected lifespan exceeds his anticipated three-year server replacement cycle.

MLC SEEN AS ALTERNATIVE RATHER THAN REPLACEMENT TO SLC

Gartner's Unsworth claims while MLC will make inroads in certain segments, he doesn't think it will overtake SLC in the next five years. Mark Peters, an analyst at Milford, Mass.-based Enterprise Strategy Group, added that he expects MLC to remain as an alternative rather than a replacement for enterprise-grade SLC SSDs.

"Generally, the end-user price of SSDs is declining significantly in any case as more vendors enter the market and volumes start to increase," Peters noted. ⊙

Carol Sliwa is features editor for SearchStorage.com.

Channel Perspective:

When to recommend SSD

Here's how to determine which customers can benefit from SSD, and which type you should recommend. *By George Crump*

DESPITE THE BUZZ around solid-state drives (SSDs) over the past few years, companies like Solid Data and Texas Memory have been shipping SSDs for more than a decade. With traditional storage vendors now also in the market, customers interested in SSD have a difficult choice to make: purchase SSD from a traditional storage supplier or buy it from one of the legacy SSD suppliers. To help them make that choice, you'll need to know what performance issue they're trying to address with SSD, as well as the differences between DRAM-based SSD and flash-based SSD.

One key point to consider when guiding customers toward a choice of SSD vendor is that traditional storage manufacturers only offer flash-based SSD, while most of the standalone providers focus on DRAM-based SSD.

There are also important speed differences between the two technologies: Flash SSDs are much faster than spinning disks, but they are not as fast as DRAM-based SSD when handling write I/Os. While the performance metrics are rapidly evolving in this market, today a flash

SSD executes write speeds at a rate of approximately 2 milliseconds and can sustain up to 25,000 random write I/Os per second. RAM SSD, by comparison, performs writes at 15 microseconds (0.015 milliseconds) while sustaining up to 400,000 write I/Os per second. (Mechanical hard disks achieve about 4- to 5-millisecond reads and writes and can sustain about 150 to 300 random I/Os per second.) Numbers are sometimes hard to grasp, so to put it in perspective, DRAM-based SSD could copy an entire DVD movie (assuming you had a DVD device fast enough to feed it the data) in less than a few seconds. A flash-based SSD, on the other hand, would take almost a minute to write the same amount of data.

For most applications, the extra speed of DRAM—and even flash-based SSD—for that matter, may be overkill. However, applications like high-transaction databases that gain a performance enhancement on flash SSD will realize an even greater advantage on DRAM. Needless to say, high-transaction database environments typically are the core revenue generator for many organizations. In fact, some companies can actually correlate every dollar gained for every second they can make the application faster. Most companies that are considering flash-based SSD should also explore RAM-based SSD; they may find it more cost-effective per I/O.

> Most companies that are considering flash-based SSD should also explore RAM-based solid-state disk; they may find it more cost-effective per I/O.

Another important distinction between flash- and DRAM-based SSD is that flash can store data persistently, like a disk drive, but DRAM does not. As a result, data on DRAM SSD could be erased if there's a loss of power or restart of the system. To address that problem, DRAM SSD manufacturers use flash and a battery to back up data on DRAM.

For DRAM-based SSDs, look to Texas Memory Systems, Solid Data Systems and newcomer Violin Memory. (DRAM-based systems are not available from traditional storage manufacturers right now, and at this point none of those companies have announced plans to support a DRAM-based SSD.) For flash-based SSD, look to primary storage manufacturers such as EMC, Sun and IBM or from Texas Memory Systems. Regardless, implementing SSD is going to take some expertise, and integration with an existing storage system won't make that any less difficult. This is an ideal opportunity for storage resellers to add real value to their relationship with a customer.

WHERE DOES SSD MAKE SENSE?

Today, only very specific storage workloads can justify the greater expense of SSD. Almost always, these storage workloads are tied directly to applications that have an impact on the revenue generation capability of an organization. There are three sweet spots to look for among your customers: Companies with databases, network-attached storage (NAS) and environments with a high server count are good candidates to buy SSD.

DATABASES

Databases are by far the most prevalent applications running on SSD. They most often tie directly to revenue creation, and it's relatively easy to dissect and identify bottlenecks in them, often related to storage I/O performance.

But even if a customer can see the value in buying SSD for a particular database, there's confusion about how big the SSD needs to be. Many companies mistakenly think that they have to buy an SSD big enough to store the entire database. This is impractical for most customers, even taking into account the rapid growth in the size of SSDs. Instead, they only need to buy a SSD drive big enough to store the database's frequently accessed files on one SSD. Identifying those files isn't incredibly difficult; most databases have tools to help. Oracle's Statspack, for example, reports the most frequently accessed files and whether the access is read, write or both. After using such a tool, it becomes a manual process to determine exactly which files belong on the SSD: You should look for undo logs, indices and temporary tables. (Most SSD suppliers, especially those that have been in the SSD market for a while, have customized tools that make the process even easier.)

Databases are by far the most prevalent applications running on SSD.

NAS

Surprisingly, the second most common area for buying and deploying SSD is on NAS systems. Using SSD on NAS makes sense when there's a small amount of frequently updated data. A good example of this is in a Web environment with a large farm of servers that need to access similar data on the same file system at nearly the same time—say, a comparison shopping website on which people are searching for iPods. Loading this subset of data onto SSD is ideal, because even though the data is being accessed through the latency-heavy Internet, there could be a thousand servers making the exact same request at almost the exact same time, and those requests are being processed locally at the website's data center. A delay here, and customers will go to another site.

Such a scenario is a perfect example of a read-intensive use of SSD on NAS. There are also write-intensive use cases for SSD on NAS, and these are ideal for DRAM-based systems. A photo sharing site, for example, will get hammered with uploads of files the day after Halloween. Obviously, these are all writes; storing those files to an SSD first and then moving them to hard disk soon after is a good use of SSD in a write-intensive environment. It's important to note that in this scenario, you'd need to use file system tools to identify high-access files or blocks within files because there's no native application to identify that data. There's also the option, for flash SSD, of using data placement software from companies such as IBM, EMC and Compellent; that software automatically moves data to where it needs to be.

Web services isn't the only area where buying SSD makes sense for file-based reads and writes. Most grid-based compute clusters that have hundreds of servers accessing the same data or data areas require fast access to data, and the mechanical nature of hard disk can become a bottleneck. And biometrics is another area where NAS-based SSDs are used. Although a biometrics scan is typically a single operation, speed is critical because an image (thumbprint) needs to be found, read and compared quickly. A delay here, and the person walks into the door rather than having it open for them.

HIGH SERVER AND SPINDLE COUNTS

This sweet spot actually ties back to the first two; in some cases, this one may be the easiest to identify. If a customer is constantly increasing server count on a particular application for better performance, they are quite possibly a candidate for buying SSD. But, before steering them in that direction, it's important to make sure that they do indeed have a storage I/O problem rather than a compute resource problem.

Similarly, high spindle counts are ideal candidates for SSD, especially if your customer is short stroking drives. Short stroking a drive means formatting it so that only the outer sectors of the disk platter are used to store data. Short stroking delivers excellent performance but is a costly way to get there because high-speed (15,000 rpm) and low-capacity drives are typically used, and the low capacity is further reduced by short stroking. In almost every situation, these applications are deployed across a very large number of drives that have been short stroked.

If a customer is constantly increasing server count on a particular application for better performance, they are quite possibly a candidate for buying SSD.

If you identify one of these situations, get ready to be a hero. With today's memory prices, it is highly likely that you can deliver a SSD solution for less than what they've invested in their compute grid or short-stroked disk infrastructure. SSD can deliver more IOPS per GB than hard drives can; as a result, you can offer better performance, at a lower cost, while being more power- and space-efficient.

EVERYTHING ELSE

While the most obvious uses of SSD (and the ones that will be easiest for you to identify and tackle) are the three discussed earlier, applications like data warehousing, file system metadata acceleration, nonlinear video editing and software versioning applications can also be good reasons to buy SSD. In determining where customers can make use of SSD, you should look for a productivity-limiting storage I/O bottleneck that's directly tied directly to restriction of profits.

In today's environment, you shouldn't look to move the entire data set to SSD, nor should you try to move the entire Tier 1 data set to SSD. You should instead look for specific files (or subsets of files, in the case of databases) that are storage I/O-constrained and which could improve the performance of critical applications by being placed on SSD. ⊙

George Crump is president and founder of Storage Switzerland, an IT analyst firm focused on the storage and virtualization segments.

Check out the following resources from our sponsors:

ATRATO™
DATA MANAGED STORAGE

The Greening of IT: Key Capacity and Performance Metrics

High Performance Deduplication for Symantec NetBackup

SSD Applications in Storage and Embedded Systems

STEC

See Why STEC is the Industry's Most Trusted Source for Enterprise Level Solid State Drives.

The Broadest Portfolio of SSD Products on the Market Today. Select the Right Drive for Your Needs.

Case Studies of How Customers have Leveraged STEC's Proven SSD Technology.

TexasMemorySystems
RamSan: The World's Fastest Storage®

The World's Fastest No Regrets SSDs -- RamSan by Texas Memory Systems

A Flash of Inspiration -- RamSan SSDs The "World's Fastest Storage" by Texas Memory Systems

The World's Fastest Storage -- RamSan SSDs by Texas Memory Systems