FLASH: THE NEXT STEP IN STORAGE

Solid-state storage delivers faster, more reliable and more efficient performance to meet today's enterprise storage needs.

Executive Summary

Demand for data storage is growing — and fast. The need to quickly and efficiently process and access data is growing as well. Ubiquitous server virtualization, cloud computing, streaming media and communications, online transaction processing, Big Data, and a host of high-performance computing (HPC) applications – this list barely touches on the numerous scenarios that now rely on fast and responsive storage on an enterprise scale.

Hard disk drives (HDDs) are the backbone of many enterprise data storage and management systems. Yet they remain the last, unbroken bottleneck in the enterprise computing chain. Computing and networking performance have increased greatly over the years, but disk rotation speeds topped out at 15,000 revolutions per minute (rpm) back in 2000, and improved storage interfaces such as SATA and SAS have yielded only incremental performance gains.

Another consideration is that HDDs are increasingly challenged by virtualized environments, which scramble input/output (I/O) operations among multiple virtual machines and pooled storage resources. This randomized I/O is particularly hostile to mechanical spinning-disk drives, which are at their best moving sequential data.

Enter flash-based, solid-state drive storage. SSDs use nonvolatile flash memory to present a disk drive subsystem with no moving parts that is able to deliver big performance gains over traditional spinning media. Storage infrastructures based on SSDs are much more responsive and significantly more reliable and efficient than HDDs. The drawback? Solid-state storage is a great deal more expensive than spinning media.

Table of Contents

Where Flash Fits

Flash-based SSDs began to make their mark in 2012 and have steadily captured market share ever since. A 2013 IHS report projects that shipments of SSDs will increase sharply through 2017, when they are expected to reach 227 million units. A 2014 IDC report projects that sales of enterprise SSD products will grow from $3.3 billion in 2013 to $10.9 billion in 2018.

Behind the growing adoption are the clear performance and efficiency advantages of solid-state storage over spinning media. Flash-based SSDs are faster, smaller, quieter and more energy efficient than HDDs. SSD read performance can outpace HDDs by a factor of 100 in some operations, while drawing a fraction of the power. For large data centers, reduced power and footprint results in lower heat dissipation and cooling load. The result: While SSDs are significantly more expensive, operators can recoup those costs through lower expenses over their useful life.

Flash-based SSDs are quickly finding their way into both existing and new storage infrastructures. Robust solutions from established and upstart vendors enable organizations to tune SSD storage — offering a range of options from tiered or cache-centric flash deployments that accelerate data operations centered on HDDs to flash-only solutions that maximize performance. Here are some common deployment scenarios:

All-flash Arrays: These arrays, such as the modular EMC XtremIO, package fast SSDs in robust arrays that connect to servers via Fibre Channel, Infiniband or 10 Gigabit Ethernet links. Performance specs are impressive, with as many as 1 million input/output operations per second (IOPS) for read operations, and about 500,000 IOPS for writes or mixed reads and writes.

Latencies of about 0.5 millisecond are typical. Capacities range from tens to hundreds of terabytes (though petabyte capacities are available), with technologies such as in-line compression and deduplication maximizing available storage. To ease cost, tiered all-flash arrays intelligently position data between ranks of durable and pricey, single-level cell (SLC) SSDs and more affordable, multilevel cell (MLC) SSDs.

PROS: Top-end performance; reduced power, cooling and space requirements

CONS: Very expensive; limited capacity compared with HDD-based arrays; may lack some storage management features

Hybrid Arrays: Like tiered all-flash arrays, hybrid arrays mix storage media to balance cost and performance — in this case, employing flash-based SSDs and low-cost spinning HDDs. An array with 10 percent flash memory can yield performance gains of 200 percent or more compared with an HDD-only array.

Software manages the disposition of data across spinning and solid-state media to optimize performance. Rack-based arrays, such as the EMC VNX and HP 3PAR StoreServ series, can store petabytes of data while boasting 400,000 or more IOPS. More compact units, such as the Nimble Storage CS-Series, can appeal to small and midsize organizations.

PROS: Reduced cost compared with all-flash arrays; higher capacities into the petabyte range; often supports advanced storage management features

CONS: Lower performance for data committed to HDDs

Flash-based Cache Solutions: Numerous vendors offer flash-based storage accelerator solutions, which place a cache of solid-state storage in front of direct-attached or array-based storage infrastructure. EMC, Fusion-io, NetApp, SanDisk and others can host up to 2TB or more of flash storage on fast PCIe (Peripheral Component Interconnect Express) adapters.

Some products leverage virtualization to pool cache storage across servers, creating a unified cache resource. Software-based solutions pool diverse solid-state memory, including existing PCIe-based flash memory, SSDs and even system DRAM, into a unified storage cache to speed disk IO and throughput.

PROS: Affordable; works with existing storage infrastructure; cache memory sits close to CPU for peak performance

CONS: Limited performance gains compared with other solutions; PCIe implementations require driver software

Direct-attached Storage (DAS): Direct-attached SSD drives connect to servers and PCs via SATA and SAS interfaces, while

Flavors of Flash

The flash memory in SSDs is not all created equal. Many SSD vendors provide both consumer-grade and enterprise-grade SSD solutions, which differ widely in terms of performance, endurance and cost. The base unit of flash memory is the cell, which holds one or more bits of data. An individual memory cell can accept about 100,000 write operations before it begins to degrade, producing risk of data loss and limiting the effective life of the SSD. Here are some common technology architectures in these SSD solutions.

Single-level Cell: As the name implies, SLC flash stores a single bit of information in each cell of the flash memory. This binary state provides for the fastest reads and writes, while minimizing wear on the flash memory. SLC provides the highest performance — and long-term reliability — but is more expensive than other solutions.

Multilevel Cell: MLC technology doubles disk capacity and lowers cost by housing two data segments in each cell. However, the scheme slows reads and writes and results in significantly reduced lifecycles.

Triple-level Cell: TLC is a consumer-grade technology that places three bits of data in each cell. It delivers higher densities and reduced costs, but diminishes useful life and performance and is not generally used for enterprise applications. TLC flash is typically found in MP3 players, tablets and other consumer electronics.

high-performance environments can benefit from PCIe-based drives. In some cases, enterprises can reduce cost by deploying consumer-grade drives that use somewhat less robust MLC flash technology.

PROS: Simple; significantly faster than HDD-based DAS storage; reduced power and cooling requirements

CONS: Expensive; individual drive capacities lower than those of HDDs

Metrics Matter

Is flash-based SSD storage a good fit for a specific data operation? IT shops should carefully assess the targeted computing environment before making any decisions, exploring specific access patterns and use cases and determining the appropriate metrics for addressing them. Any assessment of HDD versus SSD deployment needs to start with a look at the three key performance metrics: IOPS, latency and throughput.

Input/output Operations per Second

IOPS describes the number of I/O transactions that can be performed in a single second. IOPS provides insight into performance in random access scenarios and is particularly relevant in scenarios such as online transaction processing.

If a storage system will be tasked with reading, writing and moving lots of small files or blocks, IOPS is a key metric. Real-time analytics, financial modeling and trading, high-speed messaging, and video editing are other applications that benefit from high IOPS.

Flash-based storage offers large IOPS gains over spinning media. Individual SSD drives can return 10,000 or more IOPS for small file write operations, and upward of 100,000 IOPS for similar read operations. By contrast, even a fast 15,000-rpm HDD will return only hundreds of IOPS for these operations. Flash-based storage arrays can be scaled to more than 1 million IOPS, versus tens of thousands for an HDD array.

Response Time

Response time, also known as latency, is where solid-state storage truly shines. Flash-based SSDs and arrays produce microsecond latencies that are an order of magnitude better than the fastest spinning disks. Even short-stroked HDDs, which sacrifice disk capacity to minimize head movement and thereby speed random access, can't approach the low latencies of flash. An average latency of 500 microseconds (0.5 millisecond) is typical for SSD drives. A fast HDD, by comparison, will produce random access latencies of 5 ms to 10 ms.

As with IOPS, response time is critical in applications and environments where random reads and writes are common. The ubiquity of virtualization in the enterprise, including virtual desktop infrastructure (VDI), can subject server environments to increasingly randomized I/O patterns, putting a premium on low-latency storage. Flash-based storage and SSDs designed for the fast PCIe interface can yield average latencies as low as 50 microseconds (0.05 ms) or lower.

Throughput

Throughput describes how much data a drive or storage system can move in one second. Also referred to as bandwidth, throughput metrics best predict storage system behavior during sequential operations, such as streaming video or moving large blocks of contiguous data.

A SAS- or SATA-based HDD might produce somewhere near 100 megabits per second of throughput, whereas most SSDs are capable of producing two to three times that. Writes generally lag reads in solid-state throughput, but remain substantially better than spinning media.

Beyond throughput, latency and IOPS, IT decision-makers should consider other issues when weighing the value of flash-based storage solutions. These include:

Read-write ratio: SSDs are slower to handle writes than reads, and the slowdown increases as an SSD fills with data. IT groups supporting write-intensive operations can overprovision flash storage capacity to avoid slowed writes, or consider deploying fast, PCIe-based flash storage cache to maximize responsiveness in write-heavy scenarios.

Caching: Flash-based cache solutions can be extremely effective in environments with predictable data usage patterns. Sufficient cache memory should be provisioned to support use cases, as accumulated cache misses (insufficient memory) can degrade overall performance.

Capacity: SSDs are considerably more expensive than HDDs that provide the same capacity, particularly for enterprise-grade SLC SSDs that provide the highest levels of performance and endurance. For large data operations, the best approach is often to deploy tiered storage that leverages spinning media for large amounts of "warm" data and flash-based media for a limited amount of active, or "hot," data.

SSD Use Cases

Few areas benefit from flash storage deployment as much as high-performance computing. Scientific and medical research, oil and gas exploration, weather modeling, and a host of other computing-intensive activities demand urgent access to terabytes of stored data.

Flash storage significantly speeds operations while reducing power consumption and thermal output — key logistical challenges for HPC data centers. Flash storage can extend the useful life of existing data center build-outs and provide better options for new data center deployments.

Flash storage also can deliver performance benefits for Big Data applications. A 2013 survey by The Data Warehouse

Institute (TDWI) shows that more than half of IT departments engaged in Big Data analytics were already using SSDs either as primary (20 percent) or secondary (33 percent) storage.

Flash storage allows frameworks such as Hadoop to create and capture vast amounts of diverse data and support sophisticated data mining and analytics. In many cases, a tiered SSD/HDD array can balance performance and cost for Big Data operations.

Storage Use with Big Data

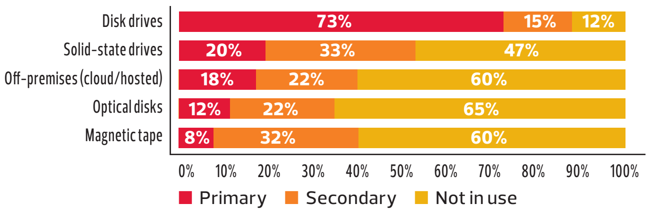

Source: The Data Warehouse Institute (TWDI), 2013, survey of IT/business leaders with Big Data experience

Another common use case is VDI, which serves client PC environments from central servers. VDI offers cost, manageability and security benefits for client PCs in the enterprise, but the random reads and writes its produces can stress spinning media infrastructures. VDI also produces daily "boot storms" when hundreds or even thousands of PCs power up in the morning. Flash storage provides the right mix of random and sequential I/O performance needed for VDI deployments.

CDW: A Storage Partner That Gets IT

Your CDW account manager and certified solution architects are ready to assist with every phase of choosing and leveraging the best storage solutions for an IT environment. CDW has years of experience working with organizations of all sizes. We can help with projects large and small by providing a dedicated account team backed by engineers and specialists.

Our approach includes:

- An initial discovery session to understand your goals, requirements and budget
- An assessment of your existing environment and definition of project requirements
- Detailed vendor evaluations, recommendations, future design and proof of concept
- Procurement, configuration and deployment of the chosen solution
- Ongoing product lifecycle support

SHARE THIS WHITE PAPER