CorData

CorData White Paper

SSD in the Enterprise

August 1, 2014

No spin is 'in'.

It's commonly reported that data is growing at a staggering rate of 30% annually. That means in just ten years, 75 Terabytes grows to over a Petabyte (more than 1,000 Terabytes). But networks haven't usually kept up with the dramatic growth of data storage capacity or user demands on that data.

At 30% growth rate, each gigabit connection five years ago would be four gigabit connections today, and each two gigabit fibre channel connection to a data storage array would be an eight gigabit connection. If your network bandwidth hasn't quadrupled over that time, then performance problems are to be expected. But upgrading your storage network every year is an expensive and disruptive solution. So how do you manage all this growth?

The answer to that question is the subject of this white paper. Midrange storage systems have grown in both capacity and functionality. In this document you'll learn how to deploy modern Solid State Disks (SSD), which represents an effective way to increase performance without continually upgrading your network.

SSD's have come a long way in just the last few years.

The first solid-state disk drives were based on SanDisk MLC flash memory technology. Invented for use in cameras, early flash memory was poorly suited for data storage applications. It didn't store every bit reliably but at the time, this was an acceptable tradeoff for low prices of camera and video memory. Another limitation is that each storage element in the SSD only takes about a thousand write operations before it wears out and stops working. That's why we don't see much USB memory combined with high-write applications. Beat it hard enough, with a write-intensive computer application like a database, and you've got one dead memory stick.

Enter third-generation SSD data storage devices to the rescue! Commonly called 'bricks', they're designed to fit into the same space as a spinning disk drive. Modern eMLC memory bricks feature an 'Enterprise' version of MLC-type solid state disk, with enhanced error correction and wear leveling features. They not only store bits reliably, but they also write blocks of data to different locations to even-out media wear. Rather than fail when a data block error occurs, bricks can self-heal by using extra stand-by blocks to substitute for the failed ones. eMLC bricks have about ten times as much write-wear capacity as the old-style MLC. In cases where a 30% or less write-to-read ratio is required, this can be a reliable and lightning-quick substitute for a spinning disk.

Some data storage usage patterns are better suited for SSD; others less so. For instance, when backup or video data is stored, or continuous streaming data storage is required, SSD can certainly handle the job. But compared to spinning disk, it's simply too expensive to be practical.

On the other hand, when you have highly randomized reads and writes with mostly reads as mentioned earlier, SSD is perfect. Since SSD's have no moving parts, they operate differently than spinning disk drives. There are unavoidable disk drive delays while the disk locates data here and there. The delay, called latency, slows things down. Latency is made up of two unavoidable limits to spinning disk performance; the rotational speed of the drive, and the speed with which the read head can position itself over a particular track. These limits cause delays with every read or write request. SSD's have no such limits. They store data reliably, and they offer very fast read and write performance.

So how do you properly integrate SSD's into your data center to capitalize on their capabilities?

The obvious targets for SSD storage are readily identifiable and separate performance-constrained data sets. For example, data base indexes contain predictable, high-demand data. And with the recent popularity of virtual server environments like VMWare and Hyper-V, storing boot partitions and virtual machines on SSD gives your entire system much better performance. And in VDI, where you might have thousands of virtual displays running on laptops, tablets, and even cell phones, performance is greatly enhanced with SSD. Its performance is relatively unaffected by lots of parallel, randomized read and write requests. In fact, for several years now SSDs have been used with great success in these environments.

Enter Hybrid Storage

There's a new way to use SSD's which is quietly revolutionizing data centers all over the country. It not only gives you better performance, but also makes it possible to easily adapt and tune data storage with greater overall effect and with lower costs than ever before.

A particularly interesting aspect of all these new data center tools such as virtual servers, VDI, and large scale shared storage arrays is their common ability to randomize your data.

In the old days, you had a one-to-one relationship between a server and its storage.

One-to-One Relationship Between
a Server and its Storage

CorData 3

Now, with virtual servers, VDI, and large-scale shared storage arrays, you have a many-to-one relationship between servers, virtual or otherwise, and the storage. When you have lots of parallel read and write transactions, you randomize the storage array's activity.

Virtual Servers and VDI Randomize
the Read/Write Requests to Disk

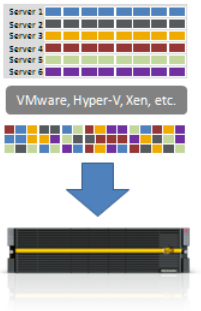

CorData 4

SSD may be used to service these random requests and optimize performance. But SSD isn't very helpful deployed as a dedicated, fixed storage resource in these homogenized environments. It's almost impossible to know what data is 'hot' at any given time. To vastly improve performance, SSD is implemented as a large, adaptive read and write cache for all your data.

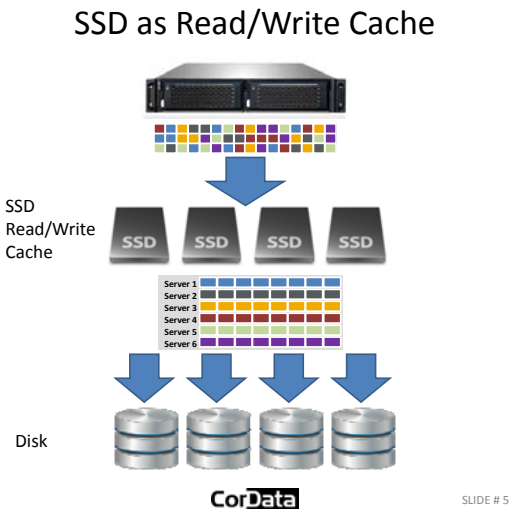

The Nexsan NST is a good example of this type of architecture. The NST utilizes a scalable pool of SSD resources for caching read and write transactions for selected disk resources. When you write data, it's written to SSD first, and then the NST system de-stages that new data from SSD to disk, but without slowing down your servers or the applications that run on them. And the SSD pool is scalable, so you easily expand your SSD resources simply by sliding another SSD 'brick' into a slot when desired. In fact, the NST performs SSD caching on both file and block based storage, serving up data for all the storage needs of a mid-range company or business unit.

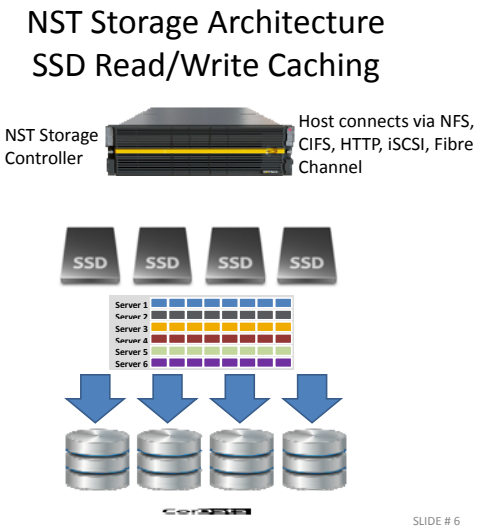

Different types of SSD media for different deployments.

There are several different types of SSD storage, with differing choices of capacity, write wear, and cost:

1. Non-volatile RAM, like system memory, has no wear whatsoever and offers the highest performance but at the highest cost and the lowest capacity. It's typically used for highly active, small write caches.

2. SLC SSD memory is the Cadillac of SSD bricks with the best wear and performance. It comes at a high cost, as you might imagine.

3. eMLC SSD is the latest enhanced SSD for enterprise data storage. This is a hardened version of the economical MLC solid state media, and has become the most popular SSD for the data center with capacities that are continuously increasing while costs continue to drop due to Moore's Law, just like memory chips. We expect eMLC SSD storage to replace most high-performance spinning disk storage applications over the next two-to-three years.

How does Tiered storage fit in?

There's been a lot of discussion recently about Tiered Storage. In the simplest implementations, tiered storage is nothing more than dedicated layers of storage media, each offering different levels of storage performance. To improve overall data performance efficiency, you simply store data at the appropriate level. It's a decision made by the storage administrator.

Simple Tiered Storage

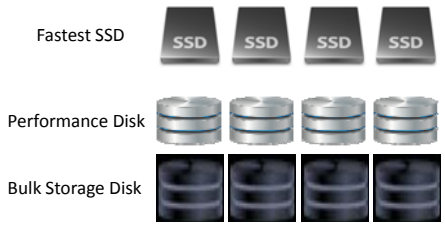

CorData SLIDE # 8

Systems offering “Auto-Tiering” features have recently appeared. Their software attempts to predict which data should be at what level using algorithms which monitor usage, frequency, and even access times. But these systems are either extremely expensive high-end storage solutions, or they’re mostly marketing hype; less-capable systems which are not very effective at predicting which data will be in high demand. And they’re not typically fast enough to make all that work identifying and copying data between tiers worthwhile. These Tiered Storage Systems, particularly the storage media, command high premium prices for all these questionable benefits.

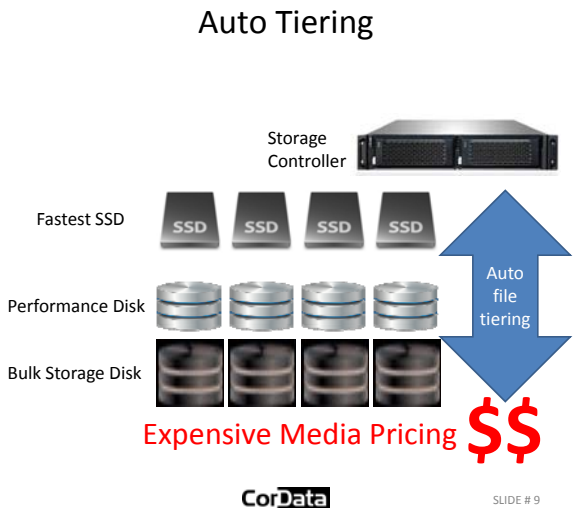

Utilizing SSD as large scale read/write cache is a simple, highly effective, and economical way to increase storage performance in the enterprise. And if the storage media is priced right, you simply buy more high performance storage media to improve performance without having to predict usage patterns at all. Companies like Nexsan are leading the way with systems utilizing these advanced Hybrid Storage architectures; offering an adaptive, expandable cache in front of your spinning disk storage array at an affordable price.

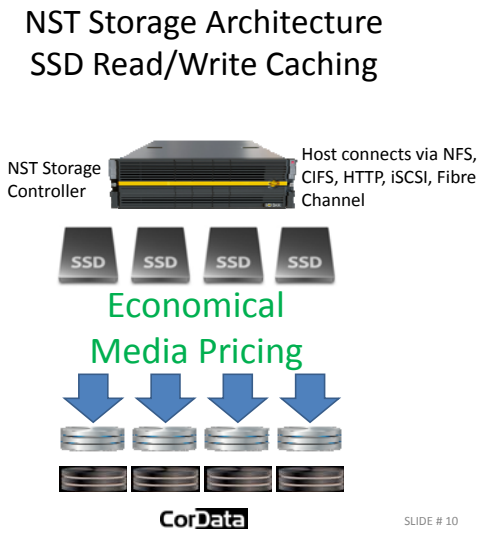

Deploying SSD isn't hard. It just takes predictable data access patterns—or not.

How do you implement SSD in the enterprise as a sharable resource—to economically increase your overall storage performance? If your performance-sensitive data is easily identifiable and isolated, put it onto dedicated SSD storage media.

But if performance is not required for a specific, identifiable group of data; and instead, an overall traffic jam is the problem, then consider a hybrid storage system utilizing SSD as a high-performance adaptive cache in front of all your disk storage.

About CorData

CorData is a storage systems integrator based in the Washington, DC area. We work with US intelligence, military, and civilian agencies, commercial clients, and research and development organizations world-wide.

For more than thirteen years, CorData has provided data storage insights to help our clients achieve their key business objectives.

Headquarters
620 Herndon Parkway
Suite 340
Herndon, VA 20171
703-986-3808
info@CorDataSys.com
www.CorDataSys.com