

STORAGE SWITZERLAND

THE NEED FOR SERVER GRADE SSDs



George Crump, Senior Analyst

Servers are an ideal place to use solid state drives (SSDs). They can shorten boot time, improve virtual memory swap file performance and improve log file operations. Not only would standard servers benefit but also applications that are run on local servers, like read heavy cloud compute and virtual desktops applications.

There is confusion though over what type of SSD should be used in these servers and the applications they run. Unfortunately for IT managers it's a choice of two extremes. On one side is consumer grade solid state storage that is cost effective but has some reliability concerns. On the other are enterprise grade SSDs that are very reliable but create cost concerns. Server-grade SSDs provide the right balance between these two extremes, delivering enterprise class reliability and performance, at the right price point.

Server-grade SSD use cases have two requirements that cannot be separated. First, anything that is installed in a server must be very reliable and in order to be considered on a wide scale, it must also be affordable. There are several aspects of an SSD that directly impact its affordability and reliability: the type of flash that is used, the sophistication of the controller, and the hardware design itself.

SLC NAND flash is considered to be the most reliable and most commonly used in enterprise-grade SSDs. While it has about 10X the life span of MLC and 3X the life span of eMLC it is also significantly more expensive than either. Clearly SLC has its uses, especially in high write environments, like network storage caches and as storage for OLTP database applications. For servers however, it may be overkill.

The primary function of an SSD in a server environment is going to be for booting and for processing local virtual memory swap files. Cloud computing applications will also benefit since many of these environments use locally attached drives and are extremely high in read I/O. The server functions that have write I/O are relatively small as it's only being generated by a single server, where most network based SSDs have a much higher write profile and are being accessed by multiple servers at the same time.

eMLC NAND flash is less expensive than SLC and more reliable than MLC, but still holds a premium over MLC. MLC is at the ideal price point for server based SSDs, but there are legitimate concerns about its reliability. A potential solution to this impasse is to leverage other components inside the SSD that directly impact reliability, the flash controller and hardware design.

Most consumer-grade SSDs couple the use of MLC NAND flash with a consumer-grade flash controller. This means that it cannot correct writes to the MLC flash as effectively as an enterprise-grade controller can, nor can it perform garbage collection and wear leveling routines that keep the drives performance at an acceptable level as quickly. The use of consumer grade controllers and consumer quality MLC has led to significant reliability and write performance issues when they are used in server environments.

Data integrity can be further assured by the use of advanced error correction as well as a data fail recovery capability. Advanced detection actually fixes errors on the fly and prolongs each cell's usefulness. Data fail recoverability allows for a flash page or block to fail completely without data loss. When this happens the drive can rebuild the data elsewhere on the drive, seamlessly; once again, prolonging life.

There are significant advantages to breaking the one-size-fits-all model of SSD deployment in the enterprise. First, since the typical server/cloud compute workload is read-heavy, using MLC with an enterprise controller that runs enterprise level firmware is a very cost effective but still reliable way to move all of a data center's servers to SSDs. At a small price premium over consumer-grade SSDs and a significant cost savings over enterprise grade SSDs, server-grade SSDs provide the right balance between cost, reliability and performance.

Another feature that is important for server-grade SSDs is a reliable power backup circuitry. Most SSDs use a small amount of RAM as a write cache since write operations are disproportionately slower on NAND flash. Using RAM helps hide that differential from the user. Of course, the problem, especially in a server instance, is that any data in the RAM that has not been committed to flash will be lost in the event of a power failure. Since this RAM area is small, even a small data loss can lead to a corrupt database.

Additionally, you can also corrupt lower flash pages in MLC

(as discussed in [SMART Storage Systems' Power Failure Protection White Paper](#)) without a backup power circuitry.

Most consumer-grade SSDs do not include protection from a power failure and subsequent data loss, whereas most enterprise-grade SSDs do. A server-grade SSD needs to have similar protection to ensure data integrity on the drive. Leveraging the enterprise design with MLC NAND economics provides the solution here as well. Companies like [SMART Storage Systems](#) use a high reliability backup circuit that consists of an array of discrete capacitors, instead of super capacitors to protect data in RAM long enough for it to be written to flash.

The problem with super capacitors is that they can significantly degrade over time and when exposed to high temperatures. This can be especially problematic in small server (1U-2U) designs where space and airflow are limited and temperatures tend to rise quickly. Backup power circuitry designs based on discrete capacitors do not have issues when exposed to these extremes and are more ideally suited for the tight confines of a server.

Enterprise server and cloud computing applications can benefit greatly from the use of SSDs for faster boot and more efficient performance when dealing with virtual memory swap file, transaction logs or cloud compute applications. The challenge has been to find an SSD drive that can provide enterprise class reliability without the enterprise class price tag.

The solution may be to combine the use of MLC NAND flash with an enterprise class controller and hardware design that includes a high reliability backup power circuitry, such as Smart Storage Systems has done with it's XceedStor 500 SSD. This SSD is designed specifically for read-intensive workloads, commonly seen in enterprise servers and cloud computing applications. Leveraging the same controller code found in SMART's enterprise-grade SSDs and combining it with MLC based NAND flash and a special hardware design results in a server-grade SSD.

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>

Copyright © 2011 Storage Switzerland, Inc. - All rights reserved