



## **Mixed All-Flash Array Delivers Safer High Performance**

*Prepared by: George Crump, Lead Analyst  
Prepared on: 8/8/2013*

Almost two years ago, Storage Switzerland predicted that automated tiering technologies would be able to do more than just move data between hard disk and flash. It would begin to move data between different types of flash based storage as well. This kind of flash tiering has finally been introduced and not by a startup company trying to capture the attention of a crowded market. It's coming from a name brand storage vendor, Dell Compellent. They were one of the first companies to bring automated hard disk tiering to market but one that isn't banking on past history, instead continuing to innovate.

Storage Switzerland was brought in to validate the technology behind Dell's newest entry to the Compellent Storage System family - The Compellent All-Flash Array - as well as Dell's claim of "Flash for the price of disk".

In our testing we learned that Dell did more than just put flash into their storage system, they've actually had that flash support for a while. In this release Dell specifically changed the way they write data to the flash tier which allows them to take better advantage of a low-latency architecture. As a result Dell Compellent, unlike other large storage vendors in the market, didn't need to bring out a dedicated all-flash storage appliance with a new architecture and software. Instead they 'doubled down' on their existing architecture, tuned it for flash and delivered what their mainstream enterprise customers need: safe, reliable, full-featured performance.

The mainstream data center, made up of mixed workloads (databases, virtual server infrastructures, VDI) needs performance, typically in the 50k to 100k IOPS range per workload, not extreme performance as in millions of IOPS. In these mainstream environments storage system reliability and consistency are as important as extreme performance.

Safety is an important requirement because Flash technology wears out. Quality enterprise class SSDs should be used in conjunction with advanced data protection like RAID 10. Reliability of the storage system is also important. Many of the new all-flash vendors don't offer or are in their first release of high availability functionality. This is risky for a mainstream data center, whose entire data set maybe on a single storage system. Finally all-flash arrays are purchased for performance but that performance can't come at the sacrifice of storage management features. A balance is needed to appeal to the mainstream data center.

As this lab report will show Dell achieved this balance nicely.

## The Tech

The Dell All-Flash Array is based on their top-of-the-line Compellent Storage Controller and Storage Center 6.4 which provides the specific optimizations for the way flash is written to. To this controller Dell attaches a 24-drive, 2U flash optimized enclosure. The initial configuration must come with 6 SLC SSD drives and 6 MLC SSD drives. The remaining 12 drive slots can be of any combination, as expansion dictates. As we will explain, based on Storage Switzerland's study, we believe that most customers will be best served by only expanding with MLC.

Dell intelligently leverages the use of both SLC and MLC flash. All flash based SSD drives have a common enemy, writes, which cause the storage devices to wear out. But SLC is much better suited to handling write traffic than MLC both from a performance and endurance stand point. In fact SLC flash can last 5X to 10X longer than MLC and ingest writes 2-3X faster. Both flash types have similar read characteristics but MLC is significantly less expensive.

In a read-heavy environment MLC hits the right combination of price, performance and reliability, but in the mainstream data center isolating these read heavy workloads can be challenging. The Dell Compellent All-Flash array does this automatically for you using its well vetted Data Progression technology. In the all-flash array configuration, all inbound writes are sent to the SLC tier. The performance optimized SLC tier is flushed during normal snapshot activity but the system can automatically flush the SLC tier if it begins to get to full. Flushed data is automatically moved to the capacity optimized MLC tier for reads, freeing up the SLC tier for new inbound writes.

A concern would be the performance impact of moving data from one flash tier to another. Would the process itself impact flash performance? In our testing, other than a slight rise in controller CPU utilization, we saw no noticeable or measurable drop in performance when the Data Progression process was executed. In multiple runs through the progression schema performance remained consistent. In other words, automated use of the SLC/MLC combination could be achieved without negatively impacting performance.

The advantages of an automatically managed, mixed flash storage system are numerous. Overall flash reliability is increased since the more vulnerable MLC flash tier is used mostly for reads. And, capacity of the more expensive SLC tier can be kept to a minimum, just large enough to handle inbound write traffic. Storage Switzerland believes that most customers will never need to expand beyond the initial six SLC drives, meaning the cost of the all-flash array could become more affordable on a per-GB basis as it expands.

That expansion system is done by adding more MLC flash capacity and more storage enclosures as needed. Again, if these enclosures are filled with MLC SSD, the cost per GB of the system becomes that much more attractive. It also seems like the Dell Flash Array will continue to deliver on its 300-400K IOPS performance all the way through full capacity.

The attraction of all-flash array technology is the predictable nature of performance. Users never have to worry about a tier or cache miss causing data to be served from hard disk. Storage Switzerland has dealt with countless IT planners trying to understand the impact of that event and have found it's extremely hard to test for. All-flash arrays eliminate that problem, but at a price.

#### The Claim

When the Dell All-Flash Array was announced the big claim was "Flash for the Price of Disk". This essentially promised an enterprise flash system that delivered an 'effective' cost per GB comparable to an enterprise disk system using 15K RPM hard drives. It did this by providing far more MLC capacity than other enterprise storage vendors have been comfortable providing. We've seen a few other all-flash vendors begin to claim HDD-like pricing for their flash arrays. While the comparison may be fair, flash vs. 15K RPM HDD, we've found that the caveats or conditions that accompany those claims are often not fair.

The most common of these is the claim of price parity after some level of deduplication has been applied. The problem is that deduplication, especially on primary storage, will have unpredictable results. Some environments may see a 5X efficiency rating, others may see 2X or less. Vendors will try to strengthen their positions by aligning the claim with a specific use case, like VDI, but the mainstream data center needs a single storage system that can support mixed workloads. This means the deduplication rate will rarely be optimum.

Also, when making these comparisons the all-flash vendors typically don't factor in the storage efficiency features that hard disk based systems are good at like thin provisioning and writeable snapshots (clones). Instead, most count only on deduplication which also isn't fair. A native capacity-per-GB price comparison is the only fair way to make the claim.

Another caveat we find less than fair is that all-flash array vendors, while making their comparisons to enterprise class systems, often don't provide complete enterprise class feature sets. For example, many don't even have a highly available (HA) configuration or have just recently added it. Beyond that they don't have the above mentioned features like snapshotting, thin provisioning or replication. Since this is a Compellent system and not something that was started from scratch the Dell All-Flash Array gains all the features of Compellent's Storage Center software, including thin provisioning, snapshots and replication. Data Progression, Dell automated tiering software which provides the intelligent data placement and movement, has been validated across many customers.

Dell claims a \$5 to \$10-per-GB price point for this solution which certainly puts it in line with other hard drive equipped enterprise class arrays, so we find this claim to be valid. More importantly, it's based on native capacity not a theoretical statistic like deduplication efficiency. Finally, the comparison to an enterprise class system is fair since the Dell Compellent All-Flash Array comes with a very complete and well tested software feature set.

### The Twist

The Dell All-Flash Array has a unique twist to it that should be of interest to many potential customers when trying to curtail costs. The "twist" comes from its ability to add hard drive storage enclosures instead of just flash enclosures.

This ability to start with all-flash allows the IT planner to develop metrics based on months or even years of real world use. Then they can add hard disk enclosures later further reducing the overall cost per GB. This allows data that can be equally served from the HDD enclosure to migrate there. Examples would be data that's redundant because of a snapshot and data that simply hasn't been accessed in a while. Compare this to an all-flash array without an HDD tier where snapshots and copies, as well as old data, must be stored on the flash tier because there is no other option. Premium capacity is wasted on data that may never be accessed.

### The Reliability Factor

While they have made some inroads in niche use cases, the single biggest challenge facing all-flash array vendors is convincing mainstream data centers that the time has come to begin using this technology. Price is of course a big issue but so are reliability and feature set. In fact repeated Storage Switzerland studies show that flash safety is the biggest concern about flash based systems after price. In our lab analysis we found that Dell covers the reliability issue as well as any vendor.

First, the SLC tier is designed specifically to handle a more write intensive workload, which is important since all writes will be sent to it first. Second, the SLC tier leverages RAID 10 for protection, delivering ideal write performance characteristics while providing the highest level of redundancy. But since this tier just handles writes and is flushed to the MLC tier regularly, it can be much smaller in capacity than the MLC tier. In other words, the way Dell implemented it, you can actually afford to RAID10 the SLC tier.

The MLC tier is also fully protected from device failure by leveraging RAID 5. Since the read tier needs to be much larger the more efficient use of capacity makes RAID 5 the right choice. It's also much better suited for read I/O than it is for write I/O.

These drive enclosures are then fully protected by redundant Compellent controllers eliminating the single point of failure that plagues many all-flash arrays. Using the right flash type in each tier and the right level of protection on those tiers, combined with a redundant storage controller, makes this array one of the safest on the market. More importantly this performance comes from the same code base, other than specific flash optimization, that has been in production for over a decade. Most other vendors are forced to start from scratch and go through iterative releases to eliminate unknown bugs.

### Implementation Recommendations

The combination of SLC and MLC plus the potential addition of HDD in the future makes for an interesting recommended configuration. Assuming that they need more than the initial capacity, we think most environments will be well served populating the rest of the first flash enclosure with MLC drives. This will provide plenty of capacity plus the performance they need. We don't think that most environments will be so write intensive that anything beyond the initial 6 SLC drives will be needed. If more initial capacity than the first flash enclosure is required we suggest that most environments would expand with additional flash enclosures that are all MLC but the fact that the system has complete flexibility, SLC expansion could be used for very performance sensitive situations.. While less SLC may make Data Progression work harder,

based on our testing we don't feel that it will impact storage performance in any noticeable way. This also keeps the system all flash (at an HDD price), so the unpredictable nature of hard disk performance won't be a concern

Once initial capacity demand has been met, for future expansion we suggest seriously considering enclosures equipped with high capacity 4TB hard disks. Based on our testing this will be used only for data that is no longer primary in nature (an old snapshot block) or is so old that the chances of an HDD access are limited. Also, the occasional access to an HDD tier should be unnoticeable to most users even in performance demanding environments.

### Impact in the Real World

What is the impact of an all-flash array with enterprise features and reliability on the mainstream data center? The Dell All-Flash Array is not a niche product, it is designed squarely for the data center and able to handle the mixture of workloads that it will encounter in such an environment. With a system like this that can deliver 300K+ IOPS consistently and reliably, the mainstream data center should be able to create a storage infrastructure that allows the workloads it supports to achieve greater ROI.

This means databases, email servers and virtual desktop infrastructures should be able to support more transactions and users. The virtualized server environment should be able to support denser virtual machine counts per host, which means more ROI per host server investment. The environment should be able to support more hosts of the above types on fewer storage systems, potentially one.

Simply put, with an all-flash array like what Dell is delivering to the market, you should be able to do more with the same hardware while spending less time fine tuning the storage for performance. And you will be able to afford it, since it is the same cost as enterprise disk.