

# Flash Storage: A Necessary Disruption

It's been a long time coming, but there is finally a revolution in storage that is breaking through performance barriers and allowing agencies to realize greater benefit from their growing data stores. The revolution is the all-flash storage array—solid state systems that use multiple flash memory drives instead of spinning hard disk drives. They are growing in popularity so fast that 451 Research calls 2014 the year of the all-flash array.

The primary reasons why organizations are moving toward flash storage include:

- **Performance:** Flash can provide 100 times the performance of traditional hard disk drives and run multiple gigabytes per second of random data. It can produce about 10,000 input/output operations per second, compared to about 200 IOPS for hard disk drives. Latency is also greatly reduced; the latency delta between disk-based and flash storage can be measured in terms of 2.5 years of computation wait time per quarter per application. In its latest buyers' guide, consultancy DCIG said: "The most recent generation of flash memory storage arrays generally deliver twice the IOPS of the prior generation and a more complete set of features that enables them to address a broader set of use cases."
- **Efficiency:** Despite generally higher costs, efficiency gains make all-flash storage a bargain, with lower dollars per IOPs. All-flash storage also reduces the overall storage footprint by eliminating overprovisioning.
- **Productivity:** With all-flash storage, organizations can get actionable intelligence faster; many can now perform analysis and run critical reports in minutes instead of days.

While not all processes are good candidates for flash storage—for example, applications that aren't read-intensive and when data access is highly random—there are several key areas that are very good fits.

**VDI:** As virtual desktops have taken hold in government, agencies are looking for ways to improve performance, manageability and scalability. Flash memory arrays are ideal for supporting the demands of VDI by addressing I/O challenges, bandwidth requirements and bottlenecks that occur when many people are accessing virtual desktops at the same time.

**Applications:** Most mission-critical applications, based on underlying databases, require high availability and high performance. All-flash storage arrays often can meet those needs better than traditional storage because everything operates at the same speed, eliminating bottlenecks and greatly improving performance. It also solves the challenge of high availability, helping maintain predictable performance.

**Virtualized Infrastructures:** Virtualized environments can experience performance bottlenecks and unacceptable degrees of latency with traditional storage. All-flash arrays solve that problem by using a flash hypervisor to scale storage quickly, resulting in faster read and write performance and much lower latency.

**The Data Center:** While the all-flash data center is a few years away, it will happen, experts say, as speed increases, costs continue to decrease, performance becomes more important, and organizations remain pressured to reduce data center footprint.

## A Flash Array Buyers Guide

No matter what brand or model of all-flash array you buy, you'll experience significant benefits. But each has different features. Here's a rundown of what's important:

**Storage type:** All-flash arrays can be based on block, file or object storage, or any combination of the three. The newest kid on the block is object storage, which is ideal because of its ability to deal with granular metadata and its capability to scale. Block and file storage are still good options, providing excellent performance. Ideally, a flash product will use all three.

**Scale-Out or Scale-Up:** Scale-out technology allow arrays to be connected to add capacity, while scale-up technology allow a flash array to add capacity without adding more hardware. Most flash vendors use one technology or the other, but if you can get both, even better.

**Flash type:** MCL, for Multi-Level Cell, is the standard, and works well for smaller enterprises. For organizations with significant scalability and performance needs, go with eMLC, an upgraded version of flash memory designed for enterprises.

**Maximum capacity:** Depending on the vendor and model, maximum capacity per drive or module ranges from 512GB to 4TB, and maximum capacity per rack from about 6TB to 57TB. Deciding how much you need is a balance between cost and expected growth.

**Thin provisioning:** This function provisions storage capacity on an on-demand basis as data is being written. By removing unused reserve space, thin provisioning maximizes physical storage capacity.

**In-line deduplication and compression:** These methods of reducing data save space and increase efficiency and improve the performance of all-flash arrays.

**Snapshots, Cloning and Replication:** Snapshots are point-in-time copies of data, restorable as needed for backups and disaster recovery. Clones are the ability to create exact copies of databases and virtual machines. All are very important flash array features. Replication copies only compressed changed data, saving space.

# Is All Flash the Same?

**W**hile all-flash arrays continue to grow in popularity, there are other storage solutions that employ flash as one of several technologies.

	Hybrid Storage Array	Flash Appliance	PCIe Flash Cards
<b>Description</b>	Adds a small percentage of flash storage to a spinning disk array	Adds flash capabilities to an existing data center. Generally built with two controller cards in the same chassis.	These cards go directly into a server's PCIe slot, eliminating the need for a SAS or SATA controller in the server.
<b>Performance</b>	3-10 milliseconds read latency	<ul style="list-style-type: none"> <li>• Less than 1 ms average latency</li> <li>• 5-20 TB usable storage</li> </ul>	<ul style="list-style-type: none"> <li>• Up to 500,000 4K IOPS</li> <li>• 100s of microseconds latency</li> </ul>
<b>Best for</b>	<ul style="list-style-type: none"> <li>• Unpredictable workloads</li> <li>• Cost is an issue</li> <li>• Evolving storage infrastructure</li> </ul>	<ul style="list-style-type: none"> <li>• Data-intensive workloads</li> <li>• Data centers in transition</li> </ul>	<ul style="list-style-type: none"> <li>• Non-persistent database files</li> <li>• Small database applications</li> </ul>

## FLASH STORAGE BY THE NUMBERS

### 1.6

The amount, in billions, that the worldwide all-flash array market will reach by 2016.

### 22

The percentage of large and mid-sized organizations that are evaluating or piloting all-flash arrays in the first half of 2014, versus 9% in the first half of 2013

Source: <https://45iresearch.com/report-long?cid=3122>

### 38

The percentage of applications being moved to flash associated with databases. The next two most popular are analytics and VDI

Source: <https://45iresearch.com/report-long?cid=3122>

### 40

The percentage by which federal agency data increases each year

### 60

The compound annual growth rate (CAGR) globally for the enterprise flash storage market through 2016

Source: <http://www.businesswire.com/news/home/20140526005123/en/Research-Markets-Global-Enterprise-Flash-Storage-Report#.VAYFkmSwKkY>

### 100

Flash can provide 100 times the performance of traditional hard drives

Source: [http://info.purestorage.com/WPImpactofusingAllFlashArrays\\_Request.html](http://info.purestorage.com/WPImpactofusingAllFlashArrays_Request.html)

### 40,000

The amount of data the digital universe will contain by 2020, and double every two years thereafter

Source: <http://www.govloop.com/profiles/blogs/infographic-big-data-converges-with-government-it>

## Paying Off: The ROI of Flash Storage

Flash arrays have generally been considered out of reach for cash-strapped agencies, but new capabilities have driven the price down to the point where it is comparable with high-speed mechanical hard drive systems. As DCIG analysts put it, flash storage now provides "twice the capacity for half the price."

While flash storage still costs roughly five times more than disk-based storage on a per gigabyte basis, these facts bring the ROI in line with other alternatives:

- Unlike traditional storage volumes, where only 30-40% of storage allocated to applications is actually being used at any given time, 100% of allocated storage is always in use in flash-based systems.
- Thin provisioning means that less storage has to be purchased upfront. According to Storage Switzerland, this reduces the upfront purchase by up to 60%.
- Inline deduplication and compression mean that data is being compressed as it is received by the storage system. Storage Switzerland estimates that this reduces actual capacity by five to 19 times.
- Flash storage is all on one tier, saving space and capacity.
- MLC and eMLC memory, used to store data in a flash environment, is more economical.

# Flash Storage: No Flash in the Pan

As agencies work to manage, analyze and store more and more data, traditional tier 1 disk-based storage has become an inhibitor. In many cases, it isn't fast enough, reliable enough or scalable enough to get the job done. For many, the answer is flash storage—a purely solid state technology unencumbered by rotating and seeking mechanical parts and thus able to overcome the inherent limitations of disk.

Moving important workloads such as databases, consolidated virtual server infrastructures (VSI) and virtual desktop infrastructure (VDI) to all-flash arrays result in many benefits, from increased productivity and end user satisfaction to accelerated processing times and more agile delivery of IT services. For example, an organization with a mechanical disk-based storage area network (SAN) had trouble keeping up with growing performance demands. With a move to an all-flash array, the organization experienced a 20-fold reduction in latency, and as a result database backups were reduced from hours to minutes, thus keeping within defined backup windows and ensuring compliance. In another case, a city government's VDI implementation was experiencing performance challenges, and storage space was running out. The solution: flash storage.

"Flash storage is the single largest advancement in the storage industry in the past 30 years," said Vaughn Stewart, Chief Technical Evangelist at Pure Storage. "It increases agility, productively and efficiency with cloud computing, big data and mobile workforces."

Moving to flash storage clearly offers significant benefits, but all flash storage arrays aren't the same. Some are difficult to operate, others don't provide consistent always-on service levels, or lack the level of security that federal agencies require. And most importantly for some, most still aren't available at an acceptable cost point. Pure Storage, recently named a leader in Gartner's 2014 Magic Quadrant for Solid-State Arrays, addresses all of these issues with its FlashArray family of shared storage appliances.

#### **The Security US Agencies**

**Require:** Hardened with limited points of entry and sophisticated intrusion detection and prevention capabilities, the FlashArray is HIPAA-compliant, stores all data with AES-256 encryption and rapid data locking for forward deployments.

**Shared and Scalable:** Pure Storage delivers multi-tenant support with role based access that includes support for active directory and LDAP directories, system auditing and user logging.

**Always-on Architecture:** Pure Storage provides non-disruptive scaling of performance and capacity while delivering 100% storage performance during normal operations, system upgrades and unexpected failures. This is critical as every agency supports globally distributed workforces and no longer has time for downtime.

**Automated and Service Centric:** Pure Storage eliminates most, if not all, of the day-to-day tasks required with storage systems. This simplicity extends into ecosystem automation and orchestration via REST APIs

and workflow integrations with a number of key infrastructure partners including VMware, Microsoft and more.

**Disruptive Economics:** With FlashReduce, Pure Storage provides the most fine-grained and largest set of data reduction technologies in the industry. Together these technologies allow Pure Storage to offer flash at the price a tier 1 disk. No other vendor can deliver the same degree of data and cost reductions for databases, virtual infrastructures and virtual desktops. One county government's VDI deployment achieved an 8:1 data reduction and a 55 percent cost savings for hundreds of desktops. In another case, an organization reduced its cost per gigabyte from \$210 to \$30 with a 17:1 data reduction.

"Today when an agency adopts flash, the experiences are so beneficial; disk is quickly deemed unacceptable for future workloads," Stewart said. "This is just the beginning of how flash will revolutionize the data center. The economics in the flash space are outpacing the rest of the market. Pure Storage customers should expect to see flash expand into new areas like capacity based storage in as little as two years. It's an exciting time to be in storage."



For more information, go to  
[www.purestorage.com](http://www.purestorage.com)