Recipe For The Flash Revolution:

Focus On The Fundamentals In Your All-Flash Array Purchasing Decision

Introduction

John Colgrove
Founder and CTO

In the era of doing more with less, IT departments have rapidly evolved to become dynamic service providers that support nearly every aspect of the digital business. From accelerating productivity, streamlining operations and speeding time to market; to delivering better end user experiences, improving operational efficiency, and reducing waste. Keeping pace in the digital economy requires the right infrastructure foundation to streamline business and IT processes, and achieve competitive advantage.

To realize new levels of agility and efficiency, many IT organizations have invested in next generation technologies like virtualization and cloud that, when coupled with new business applications, command ever-growing levels of storage performance. The increased demand has pushed legacy storage infrastructures – which are slow, complex, inefficient, and ultimately detrimental to performance intensive applications and Tier-1 workloads – to their architectural limits.

Until recently, organizations have had few viable alternatives to mechanical disk-based storage. While servers and networks have continued to evolve, following Moore's Law

Featuring research from
Gartner

to become relentlessly faster and cheaper, mechanical disk ran into a fundamental problem: capacity continued to increase each year, but it couldn't get any faster. Limited by fundamental rotational physics, as disk capacities grow storage continues to become slower each year, on a performance-per-capacity basis. As a result, servers and networks are now thousands of times faster than the mechanical disk storage to which they connect.

Complexity is inherent to hardware-confined disk-based storage architectures, often requiring storage administrators to perform daily acts of heroism to layout and re-layout disk and RAID groups, manage pools, optimize tiering and caching configurations, troubleshoot performance hotspots, or maintain application availability during performance loss, maintenance events or failures. Storage provisioning can take several hours or days, causing administrators to work extended hours over nights and weekends.

Similarly, inefficiency is also inherent to disk-based architectures. They require short stroking to compensate for rotational head latency, which wastes capacity and increases cost. Inline deduplication isn't practical, due to the crippling random workloads that apply to disk media, and inline compression likewise adds randomness and complexity to disk storage architectures. As a result, waste (in the form of capacity bloat) continues to propagate, resulting in mounting storage acquisition costs, and consuming increasingly more data center space, power, cooling, and storage administration support to manage the sprawl.

We founded Pure Storage to address these challenges. Our unique recipe for all-flash storage combines reliable performance for real-world workloads, with enterprise-grade resiliency, robust software features, baked in simplicity, and comprehensive data reduction that shrinks the data footprint to maximize efficiency. Accomplishing this required a holistic reimagining of enterprise storage architectures and business models. Today, we're delivering products that offer an order of magnitude improvement in every dimension: 10x faster, 10x more power & space efficient, 10x simpler and more reliable. It's a unique recipe that's hard to imitate, aimed at bringing the flash storage revolution to the mainstream.

Anatomy of An Enterprise-Class All-Flash Array

Any flash-based storage can offer performance improvement, but delivering a robust, production-ready feature set that exceeds the Tier-1 storage requirements of the enterprise and service provider landscape is no easy feat. We believe the following capabilities are essential for an all-flash array (AFA) to be truly enterprise-grade:

- **Always-On**

1. **High availability with zero performance loss.** AFAs are high performance solutions, designed for multi-application consolidation with some of an organization's most mission-critical workloads. Data destructive upgrades, where data must be moved off an array to upgrade software or expand capacity, are unacceptable. Similarly, consistent performance is critical given the nature of consolidated workloads. Performance disruptive upgrades where the array suffers a 25%, or 50+% performance degradation while controllers are updated are likewise unacceptable, as they can result in significant application response issues or downtime. Definition of "failure" should include availability loss AND performance loss. Specifically, an AFA should maintain full performance under load during controller and/or SSD failure.

2. **Non-disruptive software and hardware upgrades with no performance loss.** Beyond unplanned downtime, users experience much larger amount of planned downtime: software upgrades, capacity upgrades, controller upgrades, etc. An AFA must be able to support all of these upgrades and maintenance events to happen without downtime or performance loss.

3. **Simultaneous Dual-drive failure protection/recovery.** Data protection is essential for deployment in mission-critical environments. An AFA must be able to protect data integrity and maintain data availability with no performance loss during simultaneous, dual SSD failures per RAID group or recover from latent bit errors discovered during rebuild from a single SSD failure

4. **Native, space-efficient snapshots and clones.** Snapshots and clones enable sharing of production data with development/test sites and rapid recovery from online, point-in-time copies, in the rare event that a recovery situation arises. AFAs should support native snapshots and clones to avoid dependency on bolt-on appliances, preserve space-efficiency, and provide the necessary flexibility to recover instantly.

5. **Native, space-efficient replication.** Disaster Recovery is a fundamental need for mission-critical workload deployments. Adding external appliances for replication is not a viable option – external appliances add cost (redundant controllers for replication), complexity (multiple layers and management interfaces for configuring and managing replication, and balloon your data (appliances are unaware of data reduction). AFAs must be able to support native, space-efficient replication.

6. **Encryption at rest with no performance loss.** Data security is a new reality for any modern enterprise all flash array. As and when failed SSDs are replaced, customers need the assurance that the data on the failed SSD can't be retrieved. Data encryption at rest within the AFA prevents unauthorized access to data on the failed SSDs. Of course, encryption support must be available with no performance loss.

- **Built to Scale & Share**

7. **Scale to 100s of usable TBs with online capacity expansion.** Data continues to grow annually. Organizations should have the flexibility to start small and grow capacity online and non-disruptively as storage needs grow.

8. **Flash-optimized data layout and metadata engine that enables use of commodity flash.** Commodity flash coupled with data reduction enables disruptive economics, enabling the transformational effects of flash to be deployed broadly across applications. Yet, the use of commodity flash requires flash-optimized data layout and metadata engine to ensure lifespan and resiliency.

- **Disruptive Economics**

9. **Built on modular, commodity hardware.** CPUs get faster every year and Flash geometry and density advances every year. AFAs built on modular commodity hardware can harness CPU and flash advances quickly and independently.

10. **Deduplication AND compression are necessary to enable economical multi-workload consolidation.** Comprehensive data reduction technologies are essential to lower the cost per usable GB of flash, and ensure cost-efficiency. While deduplication is beneficial for highly duplicative environments such as VDI and cloning, compression is the primary form of data reduction for databases and other applications. Multi-workload consolidation requires both data reduction technologies – one without the other is half the savings, at best. An AFA must support native deduplication and compression.

- **Automated & Service-Centric**

11. **Simple & Automated. AFAs should eliminate disk complexity.** Managing RAID, disk layout, tiers/pools/caching, performance tuning, etc., should not be required. In addition, AFAs should be built for automation with the API hooks in place for OpenStack or automation orchestration software of choice.

12. **No tuning at any layer.** Requiring tuning at the application, virtual machine, or database layer to achieve optimal data reduction and performance on the all flash array propagates complexity and elongates application deployment timeline. An AFA needs to support optimal data reduction and performance without requiring any tuning at any layer.

AFAs that deliver this critical feature set are well positioned to replace mechanical disk in enterprises, service providers, and web companies' Tier-1 storage environments, and can deliver the performance, simplicity, resiliency and efficiency to revolutionize their businesses.

Top Use Cases for All-Flash Arrays

The primary uses cases where customers are most frequently displacing legacy disk storage with flash are:

- **OLTP Databases and Data Analytics:** Two of the most obvious fits for flash. Storage IO is typically the culprit behind slow, inefficient database transactions, processing, and analytic queries. AFAs can increase the volume and velocity of transaction, store and analyze larger amounts of data at a lower cost, and often enable massive consolidation at the application tier.

- **Server Virtualization:** VSI is a key driver of AFA adoption. As virtualized workloads consolidate, the IO stream down to storage becomes highly random. Mechanical disk storage gets stressed beyond the random IO it can deliver, making it a primary bottleneck in expanding server virtualization and delivering virtualization benefits enterprise-wide, including to Tier-1 performance-intensive applications. AFAs deliver abundant random performance, and pave the way for enterprise-wide virtualization.

- **Virtual Desktop Infrastructure:** A perfect use case for flash. Elevated consolidation drives exceedingly random IO; boot/login storms apply bursts that exacerbate random IO demands on storage. Cost of storage per desktop is another significant factor and a key determinant of how broadly virtual desktops can be deployed. AFAs deliver a vastly better end-user experience, and since VDI is one of the most reducible workloads in the data center, flash can often drive the cost of deployments well below those supported by mechanical disk.

Flash is the New Tier-1

Maturation of the flash storage market beyond all-flash cards and appliances, coupled with price parity with 10K/15K disk storage on usable capacity and enterprise-grade resiliency, has lit the fuse on a massive disruption in the legacy storage industry for Tier-1 deployments. Core business software applications, where speed is essential to performance and function, are now moving to all-flash storage at an accelerated pace. The "tipping point" of the flash revolution is here.

At Pure Storage, we endeavor to set the bar for storage innovation:

- We pioneered a new approach for harnessing inline deduplication and compression, together with the abundant performance of commodity MLC flash, to deliver disruptive storage economics that are cost competitive with mechanical disk.

- We have delivered built-in resiliency for Tier-1 storage, with >99.999% availability, non-disruptive operations without performance loss, dual-SSD failure protection, and disaster recovery.

- Our software architecture was designed to allow customers to start small, and easily scale as their capacity needs grow – online and cost-effectively.

- Our FlashArrays were purpose-built to deliver ultimate storage simplicity, virtually eliminating storage management complexity, and ensuring they are invisible to application administrators and fully automated and service-centric for cloud environments.

Join the Flash Revolution

Customers that have joined the flash revolution are now making a significant leap forward, speeding their existing revenue-driving applications, and unlocking business value that was previously unimaginable with disk to deliver differentiated, compelling end-user experiences and achieve new levels of efficiency.

The following Gartner Critical Capabilities for Solid-State Arrays, 2014 report offers a comparative overview of available all-flash array products – or solid-state arrays (SSAs), as defined by Gartner. We are exceedingly proud that our Pure Storage FlashArrays have received top scores in 3 out of 5 SSA use cases: Online Transaction Processing, Server Virtualization, and Virtual Desktop Infrastructure – the most common use cases for all-flash storage – and consistently achieved high scores in the other two categories.

We hope this report will be a valuable resource in supporting your deliberations, as you consider joining the flash revolution.

Source: Pure Storage

From the Gartner Files:

Critical Capabilities for Solid-State Arrays

Solid-state arrays are capable of delivering significant improvements in performance, although high cost perceptions persist. This report analyzes 13 SSAs across six high-impact use cases and quantifies product attractiveness against seven critical capabilities that are important to IT leaders.

Key Findings

- The most common use cases for solid-state arrays (SSAs) are online transaction processing (OLTP), analytics and virtual desktop infrastructure (VDI), with performance being an inordinately important factor in the selection.
- SSAs are replacing high-end enterprise arrays configured for performance and are increasingly being used in business and mission-critical environments.
- Although most organizations today have deployed SSAs in a silo for specific workloads, Gartner inquiries reveal a keen interest to harness them for multiple workloads, given the maturing data services.
- The price gap between general-purpose storage arrays and SSAs is narrowing, particularly with products that exploit consumer-grade NAND flash/solid-state drives (SSDs) with in-line data reduction features.

Recommendations

- Mitigate product immaturity concerns by choosing vendors that offer guarantees and unconditional warranties around availability, durability, performance and usable capacity.
- Choose products that can deliver consistent performance across varying workloads, which are important in your current and future environment.
- Use data reduction simulation tools to verify data reduction suitability for your data and workload.
- Implement established SSAs in business and mission-critical environments because reliability has exceeded expectations.

Strategic Planning Assumption

By 2019, 50% of traditional general-purpose storage arrays used for low-latency, high-performance workloads will be replaced by SSAs.

What You Need to Know

Solid-state arrays are rapidly gaining adoption due to significant performance advantages that customers can gain. The products from late entrants are rapidly catching up, with features on par with general-purpose arrays and established SSAs. The SSA market is divided between several pure-play emerging vendors that have built up hardware and software capabilities optimized for SSD, while larger incumbent vendors are moving aggressively to stay relevant in this important market segment, through acquisitions and/or organic product development. Many vendors have chosen to take existing proven general-purpose disk array software operating systems and array hardware designs, and adapt these to fully dedicated SSAs, which are then marketed and sold as dedicated SSA products. While this is a quick and economical method of getting an SSA to market, many existing general-purpose storage arrays were not designed for or do not lend themselves to be used as SSA because they were tuned for HDDs. In some cases, a bifurcating product line, where some array models are tuned, maintained and patched for SSDs and others for HDD, can become a software development and patch consistency nightmare, leading to restricted product problem determination and development for customers.

SSAs are used to consolidate performance, with most customers preferring to use block protocols with these storage systems. The total cost of ownership (TCO) and storage utilization of an SSA is becoming cost-competitive with general-purpose storage arrays, especially when the workloads are suitable for data reduction and when the data reduction ratio is approximately 5-to-1. While performance benchmarks are important, many customers are moving beyond that to place a high degree of emphasis on features that can enhance SSD endurance and manageability, and deliver high availability on par with general-purpose systems and data reduction features that can reduce TCO. The performance gap between the leading SSA products is narrowing, which means customers can more closely consider data services, ecosystem, services and support as important factors during evaluation. Though the cost of SSDs is falling, only through in-line data reduction features can customers fully maximize the value of their SSD tier. In addition, data reduction features can extend the longevity of the SSD tier by reducing the volume of writes and erasures. Media reliability has not been an issue due to features such as wear leveling and better error correction methods, which are also making it possible to use consumer-grade NAND flash and PC SSDs in solid-state arrays, to lower acquisition costs.

Customers should recognize that this is a highly dynamic market with a great number of product features and upgrades announced in 2014. You should choose solutions that do not require extensive storage infrastructure changes and redesign, and that are backed by strong services and support with an ability to deliver product enhancements and new features.

Within five years, the expectation of consistent sub-500-µs storage I/O response times will become commonplace and will become a performance differentiator. Today, however, any SSA vendor that can improve general-purpose HDD disk array performance to

submillisecond levels or by an order of magnitude has a valuable product differentiator. When customer and service-level expectations of 150 µs become the norm, a 50 µs to 100 µs performance difference will be significant criteria in purchasing decisions, but today, this level of difference is not important. Software, price, support and data reduction are more important than 0.1 millisecond (ms) performance differences. Conversely, most SSA vendors still emphasize and sell "speeds and feeds," whereas features such as data reduction are more important.

Product rating evaluation criteria considerations include:

- Product features considered for inclusion must have been in general availability by 30 July 2014 to be considered in the vendors' product scores.
- Ratings in this Critical Capabilities report should not be compared with other research documents, because the ratings are relative to the products analyzed in this report, not ratings in other documents.
- Scoring for the seven critical capabilities and six use cases was derived from analyst research throughout the year and recent independent Gartner research on the SSA market. Each vendor responded in detail to a comprehensive, primary research questionnaire administered by the authors. Extensive follow-up interviews were conducted with all participating vendors, and reference checks were conducted with end users. This provided the objective process for considering the vendors' suitability for the use cases.

Analysis

Critical Capabilities Use-Case Graphics

FIGURE 1
Vendors' Product Scores for Overall Use Case

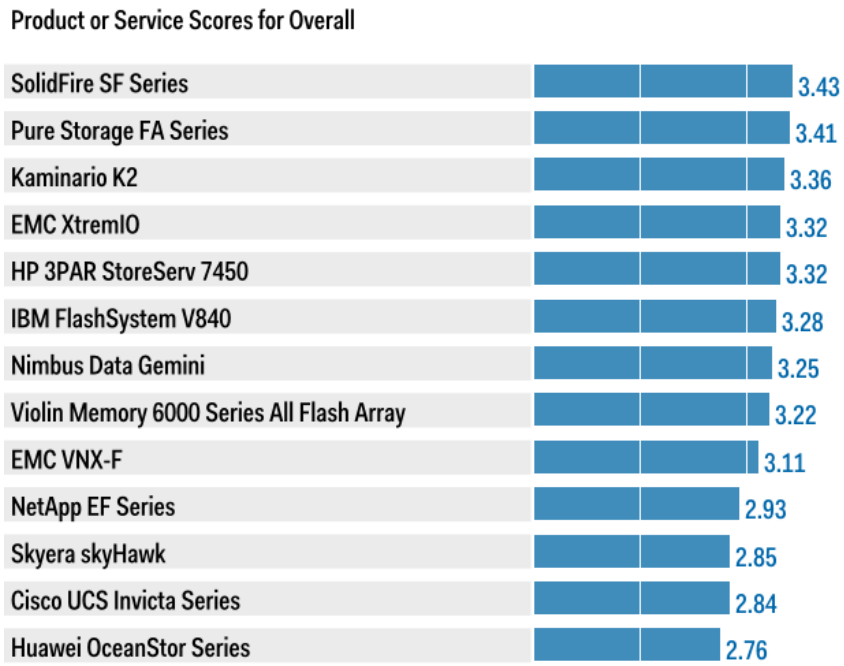

Source: Gartner (August 2014)

FIGURE 2
Vendors' Product Scores for Online Transaction Processing Use Case

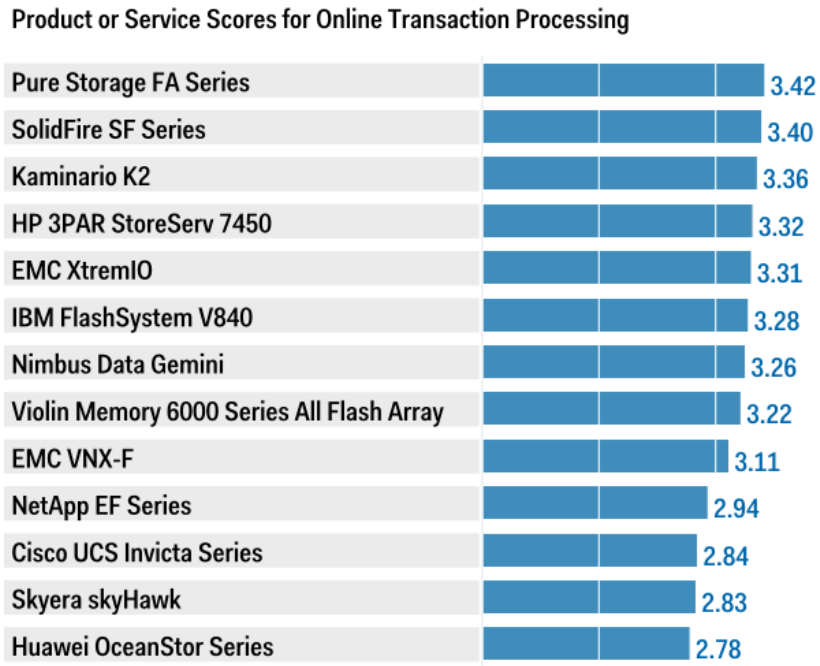

Source: Gartner (August 2014)

Vendors

Cisco UCS Invicta Series

Cisco entered the solid-state array market through the acquisition of Whiptail in 2013. Cisco has completed the process of porting the Whiptail OS onto the Unified Computing System (UCS) hardware, with plans to integrate the administration of the array with UCS Manager. The product is also being rebranded as the Cisco UCS Invicta Series, replacing the previous Accela and Invicta product names. The product uses enterprise multilevel cell (eMLC) NAND SSD with in-line deduplication and thin provisioning. The product has support for FC and iSCSI with asynchronous replication capabilities, and recently announced snapshot support.

Hypervisor support is limited to VMware, and integration with other enterprise independent software vendors (ISVs) remains limited at this point. Public performance benchmarks aren't widely available, leaving customers to use reference checks to verify claims of consistent performance. Microcode updates are disruptive, and the product currently lacks native encryption support. Customers purchasing the UCS and VCE integrated systems have the option of EMC or Cisco SSA and, therefore, will be able to leverage competing options to obtain the best purchase price.

EMC XtremIO

The XtremIO product was designed from inception to efficiently use external SSDs and currently uses robust, but more costly, enterprise SAS eMLC SSDs to deliver sustained and consistent performance. It has a purpose-built, performance-optimized scale-out architecture that leverages content-based addressing to achieve inherent balance, always-on and in-line data reduction, optimal resource utilization in its storage layout, a flash-optimized data protection scheme called XDP, and a very modern, simple-to-use graphical user interface (GUI). XtremIO arrays presently scale out to six X-Bricks, with each X-Brick having dual controllers providing a total of 120TB of physical flash, measured before the space-saving benefits of thin provisioning, data reduction and space-efficient writable snapshots. The addition of nodes currently requires a system outage, and upgrades to some ver sion 3 features such as compression will also require a disruptive upgrade, which EMC will mitigate with professional services to avoid interruptions to hosts and applications.

FIGURE 3
Vendors' Product Scores for Server Virtualization Use Case

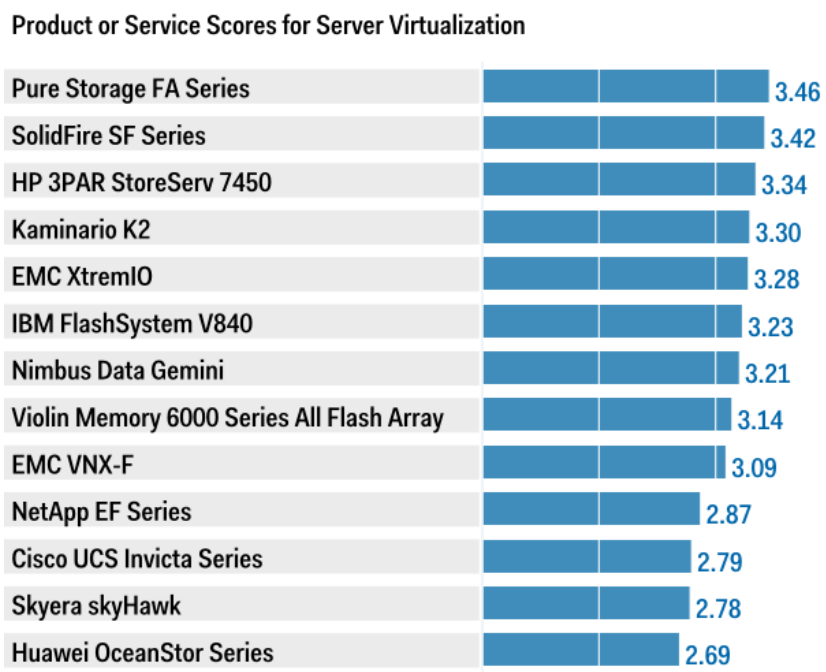

Source: Gartner (August 2014)

Compared with similar EMC scale-out product architectures, such as the Isilon scale-out array, which stores data across nodes and therefore can sustain a node outage, in an XtremIO cluster, blocks of data cannot be accessed if a single X-Brick has a complete outage, such as a simultaneous loss of both controllers, because data is stored only once on a single X-Brick.

EMC VNX-F

The lower-capacity 46TB VNX-F is based on the existing VNX unified general-purpose disk array. It has postprocess deduplication and a relatively more complex management interface due to the requirement to support the legacy (or its inherited) VNX architecture. We do not expect the VNX-F SSA and general-purpose VNX software and hardware architectures to diverge. As a result of the requirement for the software and hardware to support two different models/forks, new software features may take longer to become available due to the increased complexity of supporting two separate product lines and storage formats that use the same software stack and may require different fixes and firmware upgrades.

Different operational and administration GUIs between XtremIO and VNX-F and other products may require extra products to be purchased to provide a single management interface. Satisfaction guarantees and data service pricing for XtremIO data services are inclusive and simple, but customers need to buy data protection suite as separate package with the VNX-F.

HP 3PAR StoreServ 7450

The 7450 is based on the HP StoreServ general-purpose array architecture, which leverages HP's proprietary application-specific integrated circuit (ASIC) and additional DRAM capacity. The design uses a memory-mapping look-up implementation similar to an operating system's virtual to physical RAM translation, which is media-independent and lends itself well to virtual memory-mapping media such as SSDs. This is a particularly compelling attribute because of its efficient usage of the external SSDs by reducing the amount of overprovisioning required, as well as enabling a lean cost structure by leveraging consumer-grade MLC SSD. Another added benefit is maximizing SSD endurance as the granular, systemwide wear leveling extends the durability of the less reliable consumer MLC (cMLC) SSD media. Due to the media-independent memory-mapping

FIGURE 4
Vendor Product Scores for Virtual Desktop Infrastructure Use Case

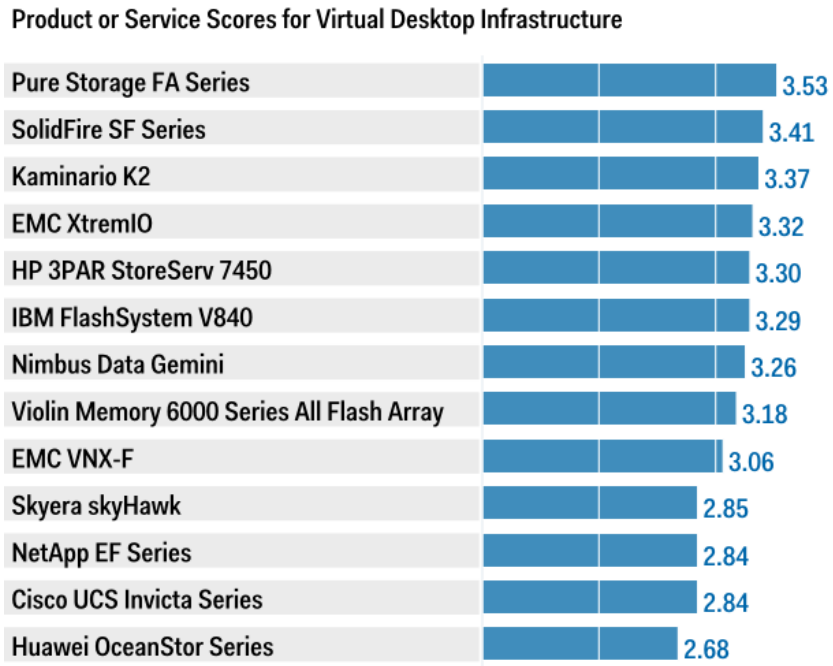

Source: Gartner (August 2014)

3PAR storage software architecture, which is implemented on SSD and general-purpose array models, we do not expect a software bifurcation. However, with more model variations, there will be longer testing and qualification periods.

The system scales to larger capacities than most competitors, with a maximum raw capacity of 460TB when configured with 1.9TB SSDs. The array does not currently have full in-line deduplication and compression, but does exploit existing 3PAR zero block bit pattern matching and thin provisioning to improve storage efficiency. The array performs well in shared environments due to its mature multitenancy and quality of service (QoS) features. However, no file protocols are supported. Pricing of all data services is tied to the general-purpose array 3PAR model and is based on host and capacity, making it complex compared to new entrants. The 3PAR 7450 platform has an extensive and proven compatibility matrix and reliability track record that is supported with a six 9s (99.9999%) high-availability guarantee during the first 12 months.

Huawei OceanStor Series

Huawei entered the solid-state array market in 2011 with the launch of the entry-level array, OceanStor Dorado2100. Since then, Huawei launched the OceanStor Dorado5100, the second generation of the 2100, and most recently, the OceanStor 18800F. Huawei currently uses self-developed SSDs that use single-level cell (SLC) and eMLC modules. The product supports thin provisioning, copy-on-write snapshots, and asynchronous and synchronous replication services. Customers can obtain competitively priced SSA from Huawei because it has an aggressive sales approach, offering steep discounts off the list price for qualified enterprise customers. Its maintenance and support pricing (as a percent of capital expenditure [capex]) also tends to be lower than many others, which is backed by a large postsales support team that is highly concentrated in Asia.

Huawei's R&D efforts in the past have been focused more on the hardware layer and only recently have software features started getting the attention that they deserve. Huawei's products currently lack data reduction features such as deduplication and compression. Firmware upgrades are disruptive, and native encryption support is lacking in the array. However, Huawei provides performance transparency via SPC benchmarks and it is one of the few unified block and file SSAs.

FIGURE 5
Vendors' Product Scores for High-Performance Computing Use Case

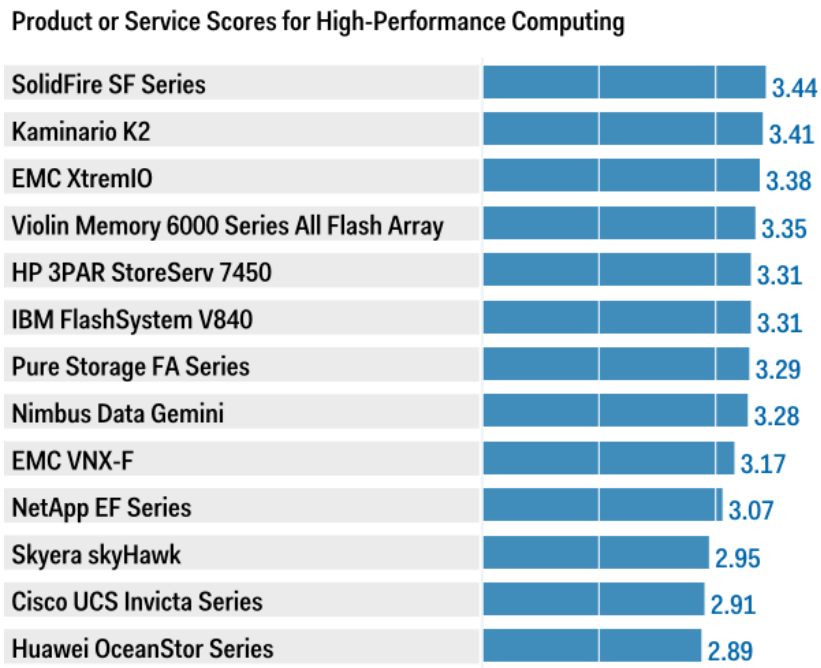

Source: Gartner (August 2014)

IBM FlashSystem V840

The FlashSystem family consists of the older 700 series and the newer 800 series SSA, and all models only support block protocols. The FlashSystem 840 has more connection options with QDR InfiniBand in addition to FC, FCoE and iSCSI connections. Alternatively, the FlashSystem V840, which adds in IBM's SVC, can scale to 320TB due to the internal FlashSystem Control Enclosure, and it inherits the broad SVC compatibility matrix but only supports FC protocols. Similarly, the V840 has richer data services in terms of QoS, compression, thin provisioning, snapshots and replication features, whereas the 840 lacks these. The 840 is designed to be a simple performance-oriented point product, whereas the V840 is for more general-purpose deployments. Both, however, lack deduplication. Additional features are provided by IBM's SVC product, FlashSystem Control Enclosure, which has a simple-to-learn-and-operate administrative GUI.

The addition of control enclosures with the V840 increases the number of separate products and components that come with the SVC layer, which reduces performance when using real-time compression compared to the 840. The SVC control enclosure layer increases product complexity, as it has separate software levels that need to be maintained and tested across IBM storage product families. In V840-based configurations, customers need to administer and operate two devices: (1) the control enclosure; and (2) the storage enclosure, which also increases system complexity, product upgrades and problem determination.

Kaminario K2

The K2 SSA is in its fifth generation and is a testament to the product's resiliency and flexible architectural approach, belying its original heritage as a DRAM appliance. Considering its genesis, K2 prides itself on its strong performance, which has been publicly scrutinized and has been verified via the Storage Performance Council, specifically the SPC-1 and SPC-2 benchmarks. In its latest generation, Kaminario also has added scale-up to its scale-out architecture, along with a more comprehensive suite of data management services. The array only supports FC and iSCSI block protocols and cMLC SATA SSDs, along with storage efficiency features that enable 30TB per rack unit that can extend both up and then out.

FIGURE 6
Vendors' Product Scores for Analytics Use Case

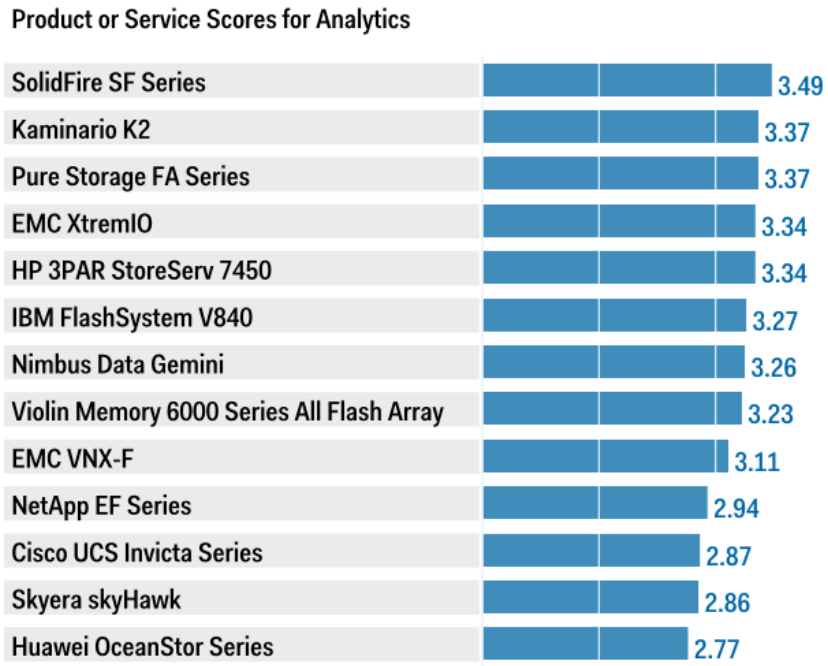

Source: Gartner (August 2014)

Its latest generation now features in-line compression, and selectable in-line deduplication and thin provisioning features. Pricing follows the customer-oriented approach of all options being inclusive in the base array price, and boasts a guaranteed effective capacity average price of $2 per GB, lending credence to storage efficiency claims. The product is reinforced with a seven-year unconditional warranty on SSD endurance, which negates customers' concerns with SSD reliability. However, replication is not yet available, and QoS performance features are limited, albeit the system does have sufficient reporting capabilities. Resiliency features and a scale-out design with very good nondisruptive software and system firmware upgrade features indicate flexibility and scalability as a fundamental requirement for its storage efficiency.

NetApp EF Series

NetApp's EF Series is an all-SSD version of the E-Series, a product line that NetApp inherited as part of the Engenio acquisition. There are two models in the EF product line — the EF540, which was launched in early 2013, and the EF550, which was launched in late 2013 with an SSD hardware refresh. The EF Series runs the SANtricity operating system and has its own management GUI. The product supports FC, iSCSI and InfiniBand. NetApp has made changes to the software to monitor SSD wear life and recently expanded the scalable raw capacity to 192TB. The EF Series product doesn't support any data reduction features. Existing NetApp OnCommand suite customers cite the need for improvement in the SANtricity management console. Given the focus of the EF Series on high-bandwidth workloads, InfiniBand has been a prominent interface, but now FC implementations have become the predominant interface as end-user acceptance of the product broadened. The long-term viability of the EF Series as a product line will remain in question, with NetApp's all-new FlashRay set for launch toward the end of the year with potentially better data services and manageability.

Nimbus Data Gemini

Nimbus has developed a purpose-built unified array from the ground up, and it features the broadest protocol support in the industry. The Gemini array is versatile due to its Halo OS, which is the epicenter for its data services and multiprotocol support, most notably including InfiniBand. The Halo operating system also offers a wide suite of data services. However, client feedback on the full depth of capability has been mixed and requires further diligence, especially the flexible and selectable data reduction features that we recommend should be verified in a proof of concept.

These arrays are cost-effective, given the use of advanced cMLC NAND designed directly into hot-swappable modules for the array. The usage of cMLC NAND and parallel memory architecture delivers resiliency across a wide spectrum of capacities ranging from 3TB to 48TB raw capacity. This approach not only provides considerable density in the 2U enclosures, but also has efficient power consumption, making it attractive from an opex perspective. The recent availability of the scale-out Gemini X-series has redundant directors that enable 10 nodes, reaching 960TB raw capacity. While the user interface and manageability is adequate, QoS features will need to evolve for this scale-out architecture.

Pure Storage FA Series

Pure Storage has focused on creating a purpose-built SSD-optimized storage array and controller software, which uses low-cost/low-capacity PC SSD cMLC media. Pure Storage is on its third-generation product, built on a foundation of granular block in-line deduplication and compression at a 512-byte level that allows compelling data reduction in workloads with various block sizes. Pure Storage has a good reputation for reliability, ease of use and extensive storage data services that now feature asynchronous replication. Overall, the arrays have relatively low raw capacities. The second-generation arrays, such as the 35TB FA-420, and the current, third-generation, more competitive 70TB FA-450 are based on Dell hardware. Only FC and iSCSI block protocols are supported with 16 Gbps FC on the FA-450 and the slower 8 Gbps FC on the FA-420. Traditional QoS features are not available, but consistent performance is provided by internal timing that skews I/Os toward the better-performing SSD. All data services are included in the base price of the array, plus product satisfaction guarantees are provided and controller investment protection is offered via the Forever Flash program if support and maintenance contracts are maintained.

Skyera skyHawk

Skyera designs its own controllers and software, and has developed its own wear-leveling algorithms to enhance and improve cMLC NAND reliability. The skyHawk is a relatively new entry-level iSCSI and NFS SSA with the extensive data services, but only a single power supply and controller. The dense packaging and exploitation of the most advanced consumer MLC NAND technology enables the product to be the most aggressively priced on a usable capacity basis, with prices starting at under $3 per GB formatted capacity before data reduction. Combined with in-line compression software, which is done only in hardware and not at the system level, this increases usable capacity. Array reporting is oriented toward internal metrics such as logical unit number (LUN) input/output operations per second (IOPS), bandwidth and latency. skyHawk also offers sophisticated array partitioning QoS features.

Firmware upgrades require a reboot and therefore an outage. Overall, from a hardware and single point of failure, this is not an enterprise data center array unless two or more Skyera arrays are mirrored or striped using a higher abstraction layer. Skyera is working with partners such as DataCore Software to provide integration certification to mitigate some of the platform compatibility and high-availability challenges. A dual-controller skyEagle array is in development, featuring 326TB of raw capacity in 1U, but given product delays with skyHawk, the already-announced skyEagle will also be subject to delays as well.

SolidFire SF Series

SolidFire sells scale-out solid-state arrays with a primary focus on service providers and large-enterprise customers. By leveraging external cMLC-based PC SSDs with in-line data reduction features, SolidFire is able to deliver competitive price/performance in a scale-out architecture. SolidFire's product is differentiated in the marketplace through QoS, where applications are delivered with guaranteed IOPS. SolidFire's QoS feature provides the ability to set minimum, maximum and burst performance settings, which enables enterprises and service providers to offer differentiated services. The product has close integration with common hypervisors, and the REST-based API support is commendable for a young company. The vendor offers broad support for cloud management platforms such as OpenStack, CloudStack and VMware, as well as support for public cloud APIs such as S3 and Swift. The product relies on a distributed replication algorithm rather than redundant array of independent disks (RAID) for data protection, which reduces rebuild times and creates a self-healing infrastructure.

Focus on enterprise private clouds and integration with traditional applications is fairly nascent and needs further development. Most current deployments are iSCSI-based, with FC support being only recently introduced.

Violin Memory 6000 Series All Flash Array

Violin has its own unique architecture based on its NAND chip-level expertise that it uses in its own Peripheral Component Interconnect Express (PCIe)-based memory module configurations that are organized and aggregated to enable a relatively dense array with strong performance and guaranteed sub-500-µs latency. Violin is one of the most cost-effective vendors on a raw $ per GB basis due to its usage of cMLC technology at advanced process geometries, allowing raw capacity up to 70TB in 3U and scale up to 280TB. Violin has strong block and file support and good ecosystem interoperability.

Violin has only recently introduced (in its June Concerto 7000 announcement) a more cohesive suite of data management features that can be upgraded from an existing 6000 series with a service disruption. The Concerto enhancements provide greater business continuity via remote asynchronous and synchronous replication along with mirroring and clones. Although Violin did introduce thin provisioning and snapshots, data reduction was only recently introduced on August 19 (and is not considered in the ratings). In June 2014, Violin also announced its Windows Flash Array, featuring tight integration with Microsoft Windows protocols and services that include data reduction, but this was not included in the rating. Pricing for data services is not fully inclusive and will require additional charges for certain features such as mirroring and replication.

Context

HDD-based general-purpose storage arrays have stagnated in performance compared to the order of magnitude of performance improvement of CPUs within servers. SSAs, which use SSDs instead of HDD, have addressed this performance imbalance by improving storage IOPS and latency performance by an order of magnitude or sometimes two orders of magnitude. While SSDs themselves are not new and have been available for decades, SSAs are new external storage offerings, which have been specifically designed or marketed to exploit the reduced cost and improved performance of NAND SSD. Latency or response time is what customers are mainly concerned with, but bandwidth or throughput is also improved by SSA. The reduced latency has also enabled new technologies such as in-line primary data reduction, deduplication, compression or both. These features were restricted by the mechanical constraints of HDDs. The reduced environmental requirements of SSAs such as power and cooling also have incidental and important advantages over general-purpose arrays and other HDD-based storage systems.

Product/Service Class Definition

The following description and criteria classify solid-state array architectures by their externally visible characteristics rather than vendor claims or other nonproduct criteria that may be influenced by fads in the solid-state array storage market.

Solid-State Array

The SSA category is a new subcategory of the broader external controller-based (ECB) storage market. SSAs are scalable, dedicated, solutions based solely on solid-state semiconductor technology for data storage that cannot be configured with HDD technology at any time. The SSA category is distinct from SSD-only racks within ECB storage arrays. An SSA must be a stand-alone product denoted with a specific name and model number, which typically (but not always) includes an operating system and data management software optimized for solid-state technology. To be considered a solid-state array then, the storage software management layer should enable most, if not all, of the following benefits: high availability, enhanced-capacity efficiency (perhaps through thin provisioning, compression or data deduplication), data management, automated tiering within SSD technologies and, perhaps, other advanced software capabilities, such as application and OS-specific acceleration based on the unique workload requirements of the data type being processed.

Scale-Up Architectures

- Front-end connectivity, internal bandwidth and back-end capacity scale independently of each other.
- Logical volumes, files or objects are fragmented and spread across user-defined collections such as solid-state pools, groups or RAID sets.
- Capacity, performance and throughput are limited by physical packaging constraints, such as the number of slots in a backplane and/or interconnected constraints.

Scale-Out Architectures

- Capacity, performance, throughput and connectivity scale with the number of nodes in the system.
- Logical volumes, files or objects are fragmented and spread across multiple storage nodes to protect against hardware failures and improve performance.
- Scalability is limited by software and networking architectural constraints, not physical packaging or interconnect limitations.

Unified Architectures

- These can simultaneously support one or more block, file and/or object protocols, such as FC, iSCSI, NFS, SMB (aka CIFS), FCoE and InfiniBand.
- Both gateway and integrated data flow implementations are included.
- These can be implemented as scale-up or scale-out arrays.

Gateway implementations provision block storage to gateways implementing NAS and object storage protocols. Gateway style implementations run separate NAS and SAN microcode loads on either virtualized or physical servers, and consequently, have different thin provisioning, auto-tiering, snapshot and remote copy features that are not interoperable. By contrast, integrated or unified storage implementations use the same primitives independent of protocol that enables them to create snapshots that span both SAN and NAS storage and dynamically allocate server cycles, bandwidth and cache based on QoS algorithms and/or policies.

Mapping the strengths and weaknesses of these different storage architectures to various use cases should begin with an overview of each architecture's strengths and weakness and an understanding of workload requirements (see Table 1).

Critical Capabilities Definition

Ecosystem

This refers to the ability of the platform to support multiple protocols, operating systems, third-party ISV applications, APIs and multivendor hypervisors.

Manageability

This refers to the automation, management, monitoring, and reporting tools and programs supported by the platform.

These tools and programs can include single-pane management consoles, monitoring and reporting tools designed to help support personnel to seamlessly manage systems, and monitor system usage and efficiencies. They can also be used to anticipate and correct system alarms and fault conditions before or soon after they occur.

Table 1. Solid-State Array Architecture

	Strengths	Weaknesses

Scale-up	• Mature architectures:
	  • Reliable
	  • Cost-competitive
	• Large ecosystems
	• Independently upgrade:
	  • Host connections
	  • Back-end capacity
	• May offer shorter recovery point objectives RPOs over asynchronous distances
		• Performance and bandwidth do not scale with capacity
		• Limited compute power can make a high impact
		• Electronics failures and microcode updates may be high-impact events

Scale-out	• IOPS and GB/sec scale with capacity • Nondisruptive load balancing • Greater fault tolerance than scale-up architectures	• There are high electronics costs relative to back-end storage costs.
Unified	• Maximal deployment flexibility • Comprehensive storage efficiency features	• Performance may vary by protocol (block versus file).
Source: Gartner (August 2014)		

Multitenancy and Security

This refers to the ability of a storage system to support a diverse variety of workloads, isolate workloads from each other, and provide user access controls and auditing capabilities that log changes to the system configuration.

Performance

This is the collective term that is often used to describe IOPS, bandwidth (MB/second) and response times (milliseconds per I/O) that are visible to attached servers.

RAS

Reliability, availability and serviceability (RAS) refers to a design philosophy that consistently delivers high availability by building systems with reliable components and "de-rating" components to increase their mean times between failures (MTBFs).

Systems are designed to tolerate marginal components, hardware and microcode designs that minimize the number of critical failure modes in the system, serviceability features that enable nondisruptive microcode updates, diagnostics that minimize human errors when troubleshooting the system, and nondisruptive repair activities. User-visible features can include tolerance of multiple disk and/or node failures, fault isolation techniques, built-in protection against data corruption, and other techniques (such as snapshots and replication; see Note 1) to meet customers' recovery point objectives (RPO) and recovery time objectives (RTO).

Scalability

This refers to the ability of the storage system to grow not just capacity, but performance and host connectivity. The concept of usable scalability links capacity growth and system performance to SLAs and application needs (see Note 2).

Storage Efficiency

This refers to the ability of the platform to support storage efficiency technologies, such as compression, deduplication and thin provisioning, to improve utilization rates while reducing storage acquisition and ownership costs.

Use Cases

Overall

This is an average of the following use cases. Please refer to Table 2 for the weightings of the use cases.

Online Transaction Processing

This use case is closely affiliated with business-critical applications, such as database management systems (DBMSs).

DBMSs require 24/7 availability and subsecond transaction response times — hence, the greatest emphasis is on performance and RAS features. Manageability and storage efficiency are important because they enable the storage system to scale with data growth while staying within budget constraints.

Server Virtualization

This use case encompasses business-critical applications, back-office and batch workloads, and development.

The need to deliver low I/O response times to large numbers of virtual machines or desktops that generate cache-unfriendly workloads, while providing 24/7 availability, heavily weights performance and storage efficiency, followed closely by RAS.

High-Performance Computing

High-performance computing (HPC) clusters can be made of large numbers of servers and storage arrays, which together deliver high compute densities and aggregated throughput.

Commercial HPC environments are characterized by the need for high throughput and parallel read-and-write access to large volumes of data. Performance, scalability and RAS are important considerations for this use case.

Analytics

This use case applies to all analytic applications that are packaged or provide business intelligence (BI) capabilities for a particular domain or business problem.

It does not apply to only storage consumed by big data applications using map/reduce technologies (see definition in "Hype Cycle for Advanced Analytics and Data Science, 2014").

Virtual Desktop Infrastructure

Virtual desktop infrastructure (VDI) is the practice of hosting a desktop operating system within a virtual machine (VM) running on a centralized server.

VDI is a variation on the client/server computing model, sometimes referred to as server-based computing. Performance and storage efficiency (in-line data reduction) features are heavily weighed for this use case for which solid-state arrays are emerging as a popular alternative.

Inclusion Criteria

- It must be a self-contained, SSD-only system that has a dedicated model name and model number.
- The SSD-only system must be exactly that. It must be initially sold with 100% SSD and cannot be reconfigured, expanded or upgraded at any future point in time with any form of HDDs within expansion trays via any vendor special upgrade or specific customer customization or vendor product exclusion process into a hybrid or general-purpose SSD and HDD storage array.
- The vendor must sell its product as stand-alone product, without the requirement to bundle it with other vendors' storage products in order for the product to be implemented in production.
- Vendors must be able to provide at least five references to Gartner that can be successfully interviewed by Gartner. At least one reference must be provided from each geographic market (Asia/Pacific, EMEA and North American) or the two within which the vendor has a presence.
- The vendor must provide an enterprise-class support and maintenance service, offering 24/7 customer support (including phone support). This can be provided via other service organizations or channel partners.
- The company must have established notable market presence, as demonstrated by the amount of terabytes sold, the number of clients or significant revenue.
- The product and a service capability must be available in at least two of the following three markets — Asia/Pacific, EMEA and North American — by either direct or channel sales.

The solid-state arrays evaluated in this research include scale-up, scale-out and unified storage architectures. Because these arrays have different availability characteristics, performance profiles, scalability, ecosystem support, pricing and warranties, they enable users to tailor solutions against operational needs, planned new application deployments, forecast growth rates and asset management strategies.

Although this SSA critical capabilities research represents vendors whose dedicated systems meet our inclusion criteria, ultimately, it is the application workload that governs which solutions should be considered, regardless of any criteria. The following vendors and products were considered for this research but did not meet the inclusion criteria, despite offering SSD-only configuration options to existing products. The following vendors may still warrant investigation based on application workload needs for their SSD-only offerings: American Megatrends, Dell Compellent, EMC VMAX,

Table 2. Weighting for Critical Capabilities in Use Cases

Critical Capabilities	Overall	Online Transaction Processing	Server Virtualization	High-Performance Computing	Analytics	Virtual Desktop Infrastructure
Performance	29.0%	30.0%	20.0%	42.0%	25.0%	30.0%
Storage Efficiency	16.0%	15.0%	20.0%	5.0%	15.0%	25.0%
RAS	17.0%	20.0%	15.0%	15.0%	20.0%	15.0%
Scalability	11.0%	8.0%	10.0%	15.0%	18.0%	4.0%
Ecosystem	7.0%	7.0%	10.0%	3.0%	5.0%	8.0%
Multitenancy and Security	6.0%	5.0%	5.0%	10.0%	6.0%	5.0%
Manageability	14.0%	15.0%	20.0%	10.0%	11.0%	13.0%
Total	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**	**100.0%**
						As of August 2014

Source: Gartner (August 2014)

Fusion-io ION (recently acquired by SanDisk), Fujitsu Eternus DX200F, Hitachi Unified Storage VM, IBM DS8000, NetApp FAS, Oracle ZFS and Tegile T-series.

This methodology requires analysts to identify the critical capabilities for a class of products/services. Each capability is then weighed in terms of its relative importance for specific product/service use cases.

Editor: please add the following boilerplate text below the table, "This methodology requires analysts to identify the critical capabilities for a class of products/services. Each capability is then weighed in terms of its relative importance for specific product/service use cases."

Critical Capabilities Rating

Each product or service that meets our inclusion criteria has been evaluated on several critical capabilities on a scale from 1.0 (lowest ranking) to 5.0 (highest ranking). Ratings are listed in Table 3, below.

Editor please add the following text above the next table: "Table 3 shows the product/service scores for each use case. The scores, which are generated by multiplying the use case weightings by the product/service ratings, summarize how well the critical capabilities are met for each use case."

Table 4 shows the product/service scores for each use case. The scores, which are generated by multiplying the use-case weightings by the product/service ratings, summarize how well the critical capabilities are met for each use case.

To determine an overall score for each product/service in the use cases, multiply the ratings in Table 3 by the weightings shown in Table 2.

Editor: please add the following boilerplate text below the table, "To determine an overall score for each product/service in the use cases, multiply the ratings in Table 2 by the weightings shown in Table 1."

Solid-state usage in solid-state arrays Gartner uses the commercial enterprise term "solid-state array" (SSA) to differentiate from electromechanical disk drives and avoid media dependence on a particular memory technology. The term "flash" is a consumer market term, and the products analyzed are not targeted for the consumer market. In the future, the present NAND circuits may not be used, but another type or derivative of semiconductor memory technology such as 3DNAND, Memristors, Phase Change Memory or any other solid-state technology could be used in SSAs. This gives the term "SSA" more inclusivity and longevity, and it is more flexible and not tied to a specific solid-state storage media or format.

Table 3. Product/Service Rating on Critical Capabilities

Product or Service Ratings	HP 3PAR StoreServ 7450	Violin Memory 6000 Series All Flash Array	Huawei OceanStor Series	NetApp EF Series	EMC VNX-F	Pure Storage FA Series	Nimbus Data Gemini	IBM FlashSystem V840	Cisco UCS Invicta Series	Kaminario K2	SolidFire SF Series	Skyera skyHawk	EMC XtremIO
Performance	3.1	3.7	3.1	3.2	3.2	3.3	3.4	3.5	3.1	3.7	3.3	3.3	3.6
Storage Efficiency	3.2	2.8	1.9	2.1	2.5	4.2	3.3	3.2	2.8	3.4	3.7	2.9	3.4
RAS	3.5	3.2	3.1	3.1	3.2	3.4	3.4	3.4	2.9	3.4	3.4	2.4	3.1
Scalability	3.5	3.4	2.8	3.1	3.2	2.8	3.2	3.1	3.0	3.4	4.0	3.2	3.6
Ecosystem	3.8	3.1	2.7	2.6	3.9	3.2	3.0	3.3	2.6	3.0	2.9	2.5	3.0
Multitenancy and Security	3.6	2.9	2.6	3.0	3.2	3.3	3.2	3.2	2.6	2.9	3.5	2.6	3.0
Manageability	3.2	2.8	2.7	3.1	3.0	3.4	2.9	2.9	2.4	2.9	3.2	2.4	3.0
As of August 2014													
Source: Gartner (August 2014)													

Table 4. Product Score in Use Cases

Use Cases	HP 3PAR StoreServ 7450	Violin Memory 6000 Series All Flash Array	Huawei OceanStor Series	NetApp EF Series	EMC VNX-F	Pure Storage FA Series	Nimbus Data Gemini	IBM FlashSystem V840	Cisco UCS Invicta Series	Kaminario K2	SolidFire SF Series	Skyera skyHawk	EMC XtremIO
Overall	3.32	3.22	2.76	2.93	3.11	3.41	3.25	3.28	2.84	3.36	3.43	2.85	3.32
Online Transaction Processing	3.32	3.22	2.78	2.94	3.11	3.42	3.26	3.28	2.84	3.36	3.40	2.83	3.31
Server Virtualization	3.34	3.14	2.69	2.87	3.09	3.46	3.21	3.23	2.79	3.30	3.42	2.78	3.28
High-Performance Computing	3.31	3.35	2.89	3.07	3.17	3.29	3.28	3.31	2.91	3.41	3.44	2.95	3.38
Analytics	3.34	3.23	2.77	2.94	3.11	3.37	3.26	3.27	2.87	3.37	3.49	2.86	3.34
Virtual Desktop Infrastructure	3.30	3.18	2.68	2.84	3.06	3.53	3.26	3.29	2.84	3.37	3.41	2.85	3.32
As of August 2014													
Source: Gartner (August 2014)													

Acronym Key and Glossary Terms

Solid-state usage in solid-state arrays	Gartner uses the commercial enterprise term "solid-state array" (SSA) to differentiate from electromechanical disk drives and avoid media dependence on a particular memory technology. The term "flash" is a consumer market term, and the products analyzed are not targeted for the consumer market. In the future, the present NAND circuits may not be used, but another type or derivative of semiconductor memory technology such as 3DNAND, Memristors, Phase Change Memory or any other solid-state technology could be used in SSAs. This gives the term "SSA" more inclusivity and longevity, and it is more flexible and not tied to a specific solid-state storage media or format.

Evidence

- Client interactions over the past year, vendor briefings and vendor references
- Detailed questionnaire responses and review calls with all profiled vendors
- Detailed reference checks with more than 50 customers

Critical Capabilities Methodology

This methodology requires analysts to identify the critical capabilities for a class of products or services. Each capability is then weighted in terms of its relative importance for specific product or service use cases. Next, products/services are rated in terms of how well they achieve each of the critical capabilities. A score that summarizes how well they meet the critical capabilities for each use case is then calculated for each product/service.

"Critical capabilities" are attributes that differentiate products/services in a class in terms of their quality and performance. Gartner recommends that users consider the set of critical capabilities as some of the most important criteria for acquisition decisions.

In defining the product/service category for evaluation, the analyst first identifies the leading uses for the products/services in this market. What needs are end-users looking to fulfill, when considering products/services in this market? Use cases should match common client deployment scenarios. These distinct client scenarios define the Use Cases.

The analyst then identifies the critical capabilities. These capabilities are generalized groups of features commonly required by this class of products/services. Each capability is assigned a level of importance in fulfilling that particular need; some sets of features are more important than others, depending on the use case being evaluated.

Each vendor's product or service is evaluated in terms of how well it delivers each capability, on a five-point scale. These ratings are displayed side-by-side for all vendors, allowing easy comparisons between the different sets of features.

Ratings and summary scores range from 1.0 to 5.0:

1 = Poor: most or all defined requirements not achieved

2 = Fair: some requirements not achieved

3 = Good: meets requirements

4 = Excellent: meets or exceeds some requirements

5 = Outstanding: significantly exceeds requirements

To determine an overall score for each product in the use cases, the product ratings are multiplied by the weightings to come up with the product score in use cases.

The critical capabilities Gartner has selected do not represent all capabilities for any product; therefore, may not represent those most important for a specific use situation or business objective. Clients should use a critical capabilities analysis as one of several sources of input about a product before making a product/service decision.

Source: Gartner Research, G00260421, Valdis Filks, Joseph Unsworth, Arun Chandrasekaran, 29 August 2014

Note 1
Replication Explanation

Replication distance is a function of two variables: storage array latency and network latency. SSAs reduce the storage array latency component so that the overall distance can be increased. However, because synchronous replication adds milliseconds of delay to the microsecond response time of SSAs, synchronous replication slows down the performance of all SSAs and SSDs due to the additional network latency induced by the requirement to acknowledge successful data replication at the remote destination. At a distance of 100 km, this reduces SSD performance by an order of magnitude (from 100 µs to 2.1 ms), and therefore, synchronous replication negates the performance improvement and is not an important factor during SSA acquisition. High-availability applications requiring high-performance and synchronous replication are limited, and the choice is mutually exclusive. Asynchronous replication with consistency groups is therefore the recommended technique if performance is more important than replication. From a positive perspective, existing synchronous replication distance limitations for HDD-based storage arrays can be relaxed. When an HDD-based array is replaced with an SSA, the new SSA's submillisecond performance will reduce the previous HDD array latency, thus allowing an increase in network latency and a longer distance — for example, if SSA performance improves by 4 ms (good HDD latency is 5 ms) to 1 µs or less. A customer can then increase synchronous replication distances by 200 km, which adds 4 ms, while maintaining the same performance, to the application when replacing a HDD storage array with an SSA.

Note 2
Capacity Explanation

All SSA capacities are in total or raw capacity and are not usable unless otherwise stated.

About Pure Storage

Pure Storage, Inc.
Twitter: @purestorage

650 Castro Street, Suite #400
Mountain View, CA 94041

T: 650-290-6088
F: 650-625-9667

Sales: sales@purestorage.com
Support: support@purestorage.com
Media: pr@purestorage.com
General: info@purestorage.com

Pure Storage, the all-flash enterprise storage company, enables the broad deployment of flash in the data center. When compared to traditional disk-centric arrays, Pure Storage all-flash enterprise arrays are 10x faster and 10x more space and power efficient at a price point that is less than performance disk per gigabyte stored. The Pure Storage FlashArray is ideal for high performance workloads, including server virtualization, desktop virtualization (VDI), database (OLTP, real-time analytics) and cloud computing. For more information, visit www.purestorage.com.