SIGMA SOLUTIONS

SigmaUptime

volume 13 number 2

NEWS FLASH

As flash storage dominates the headlines, EMC's XtremIO all-flash arrays make news thanks to extreme performance, scalability and data protection

UPTIME

PRSRT STD
U.S. POSTAGE
PAID
Tulsa, OK
Permit No. 2146

EMC²

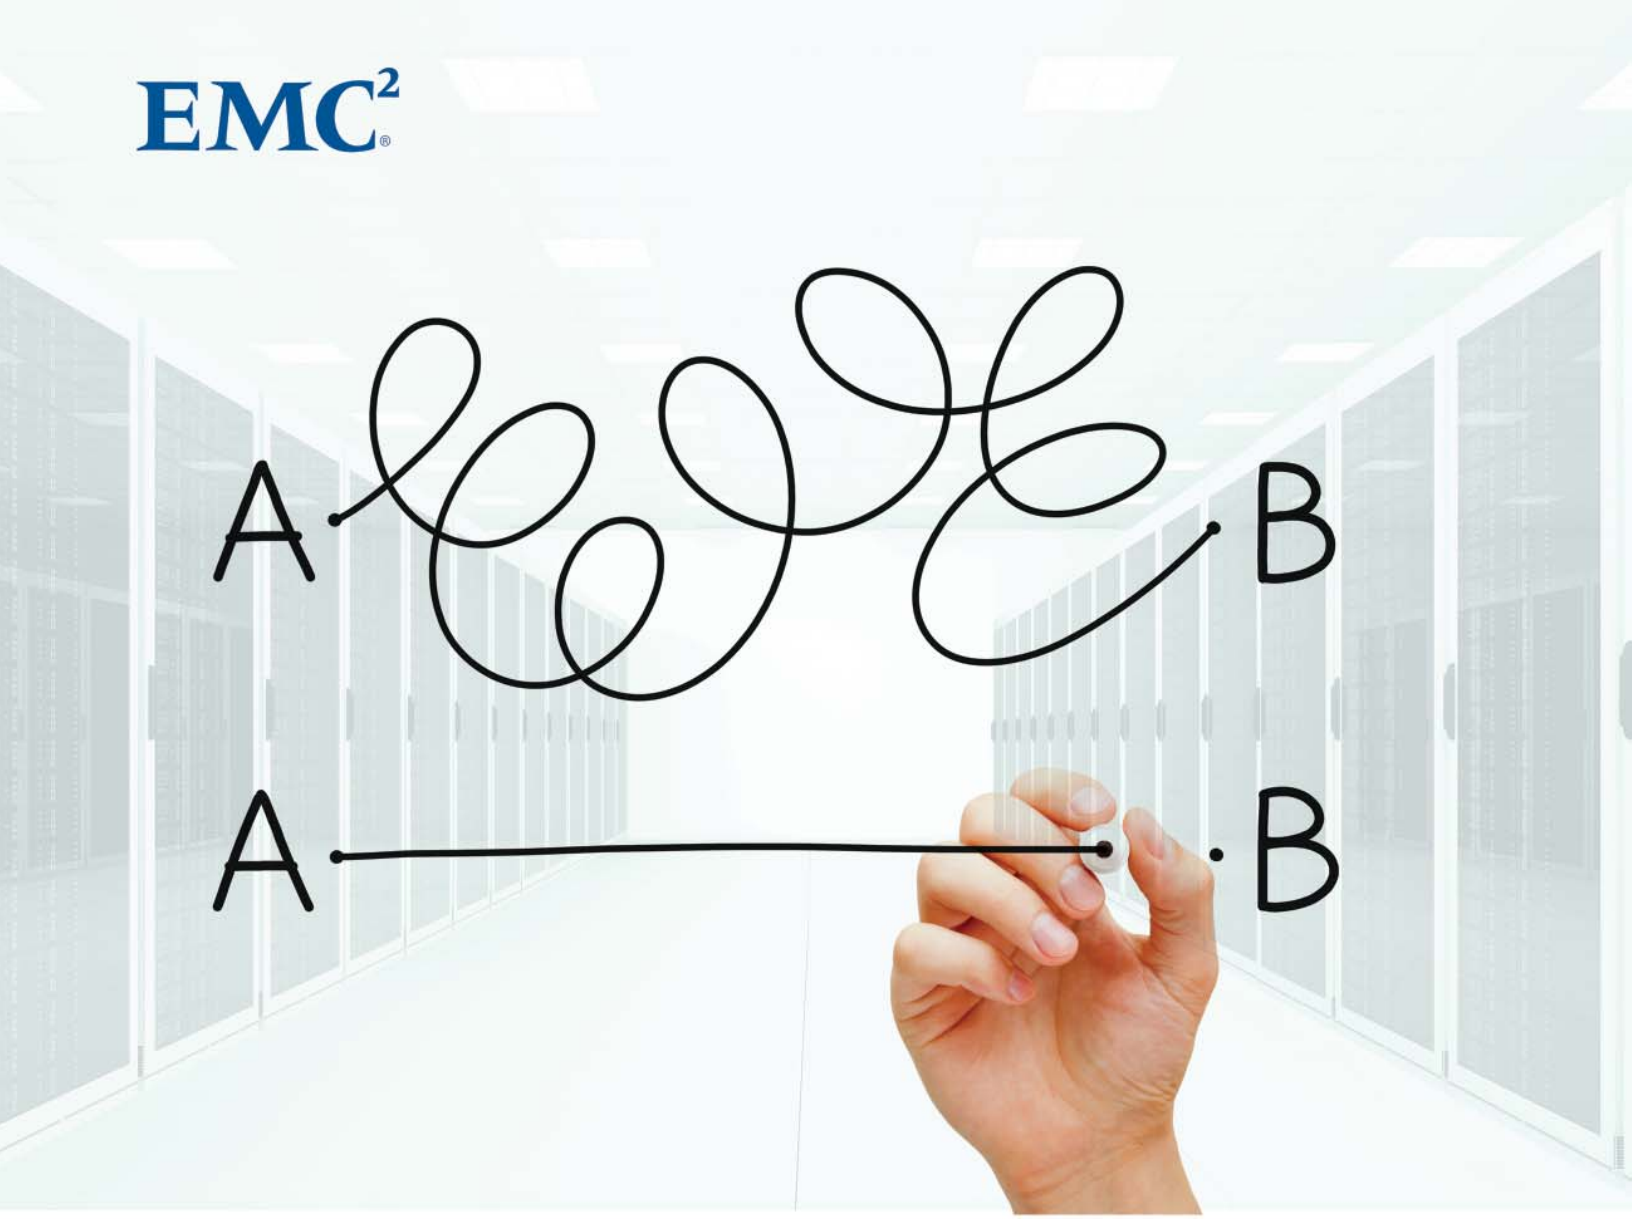

STORAGE. **SIMPLIFIED.**

Introducting EMC ViPR Software-Defined Storage.

EMC ViPR is a lightweight software platform that abstracts existing storage arrays into a single pool of virtual storage, similar to how PCs use universal device drivers to connect to peripherals. This approach creates an extensible "plug and play" storage environment that not only works with EMC arrays of all kinds but will also work for third-party storage solutions. With EMC ViPR, you are no longer bound by physical storage limitations, but can transform your existing infrastructure into a simple, extensible and open platform to extract more value from your data storage assets.

Contact your Sigma Solutions representative for more information about using EMC ViPR to solve your data storage challenges.

SIGMA SOLUTIONS

888.895.0495
www.sigmasol.com

Contents

4 **News Flash**

The emergence of all-flash storage arrays is a big story in the data storage world. They are exponentially faster than the best hard-disk arrays, but cost and other drawbacks have limited their usage in the data center. With its new XtremIO solution, EMC has overcome these drawbacks to deliver extreme performance, scalability and data protection.

8 **Redefining Storage**

Software-defined storage separates the storage controller from storage hardware, allowing storage to be managed and automated with intelligent software. With ViPR, EMC has taken an early lead in this nascent market.

10 **Ready for Prime Time?**

Data de-duplication's sweet spot has been in backup applications dealing with large data sets and many redundancies. Now, however, there is an industry-wide push to extend the benefits of de-duplication to primary, or Tier 1, storage.

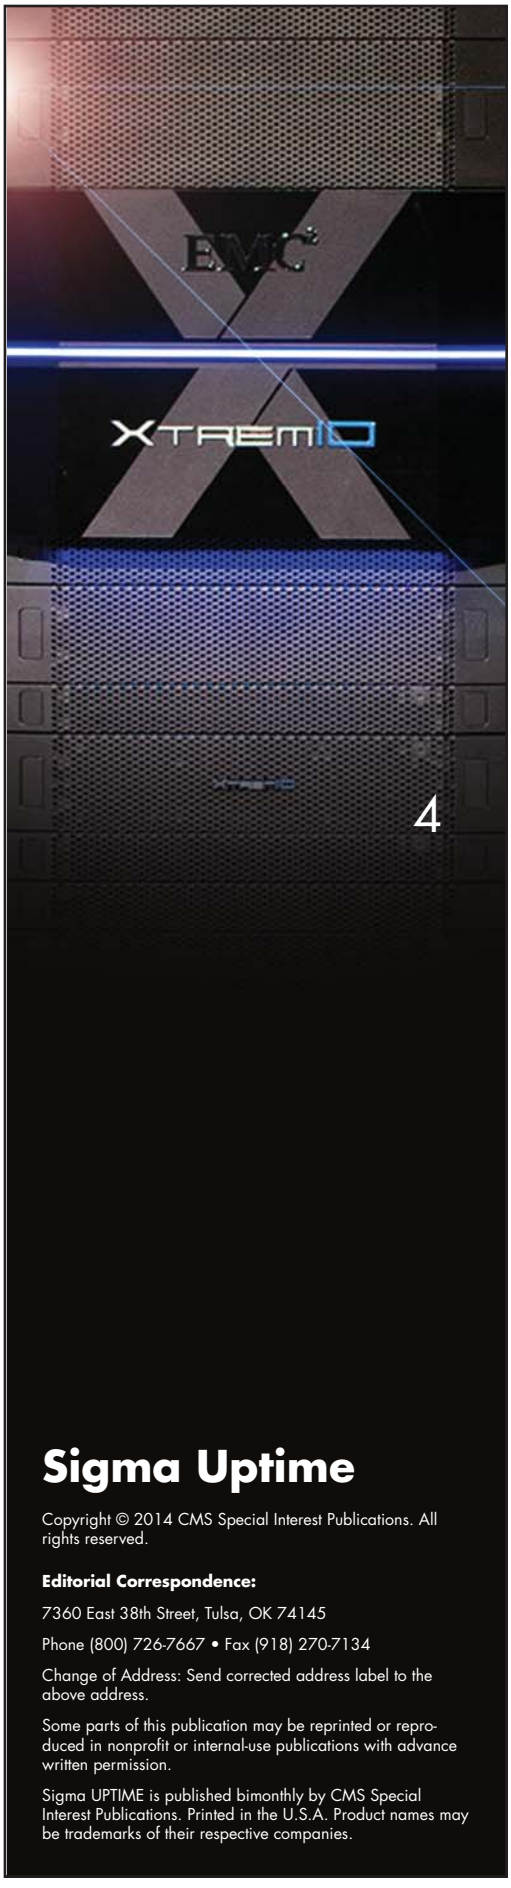

Sigma Uptime

Editorial Correspondence:

7360 East 38th Street, Tulsa, OK 74145

Phone (800) 726-7667 • Fax (918) 270-7134

Change of Address: Send corrected address label to the above address.

Sigma UPTIME is published bimonthly by CMS Special Interest Publications. Printed in the U.S.A. Product names may be trademarks of their respective companies.

NEWS FLASH

As flash storage dominates the headlines, EMC's XtremIO all-flash arrays make news thanks to extreme performance, scalability and data protection.

Flash was far and away the top storage story of 2013, and continues to dominate headlines in 2014. Several major storage vendors brought all-flash arrays to market, and the price of solid-state disks (SSDs) dropped to the point where they compete with hard-disk drives (HDDs). IT managers have stopped wondering if they should use flash in the data center and are now considering how they should best take advantage of it.

That's good news for organizations facing both massive data storage growth and ever-increasing performance demands. While CPU clock speeds have been getting faster all the time, traditional HDD storage creates performance bottlenecks due to the inherent inefficiencies of disk. According to most studies, the fastest HDDs have access times of about 5 milliseconds, causing unacceptable latency in mission-critical applications.

By using NAND flash memory instead of spinning magnetic disks to store and retrieve data, SSDs have read/write response times of as little as 20 microseconds — exponentially faster than the best HDDs. This resolves a vexing problem that has bedeviled the storage industry for years.

EMC became the first enterprise storage vendor to integrate SSDs into its core product portfolio in early 2008, when it began shipping Symmetrix DMX-4 arrays with optional enterprise flash drives. In November 2013, the company introduced another first, an all-flash array that provides consistent and predictable performance to

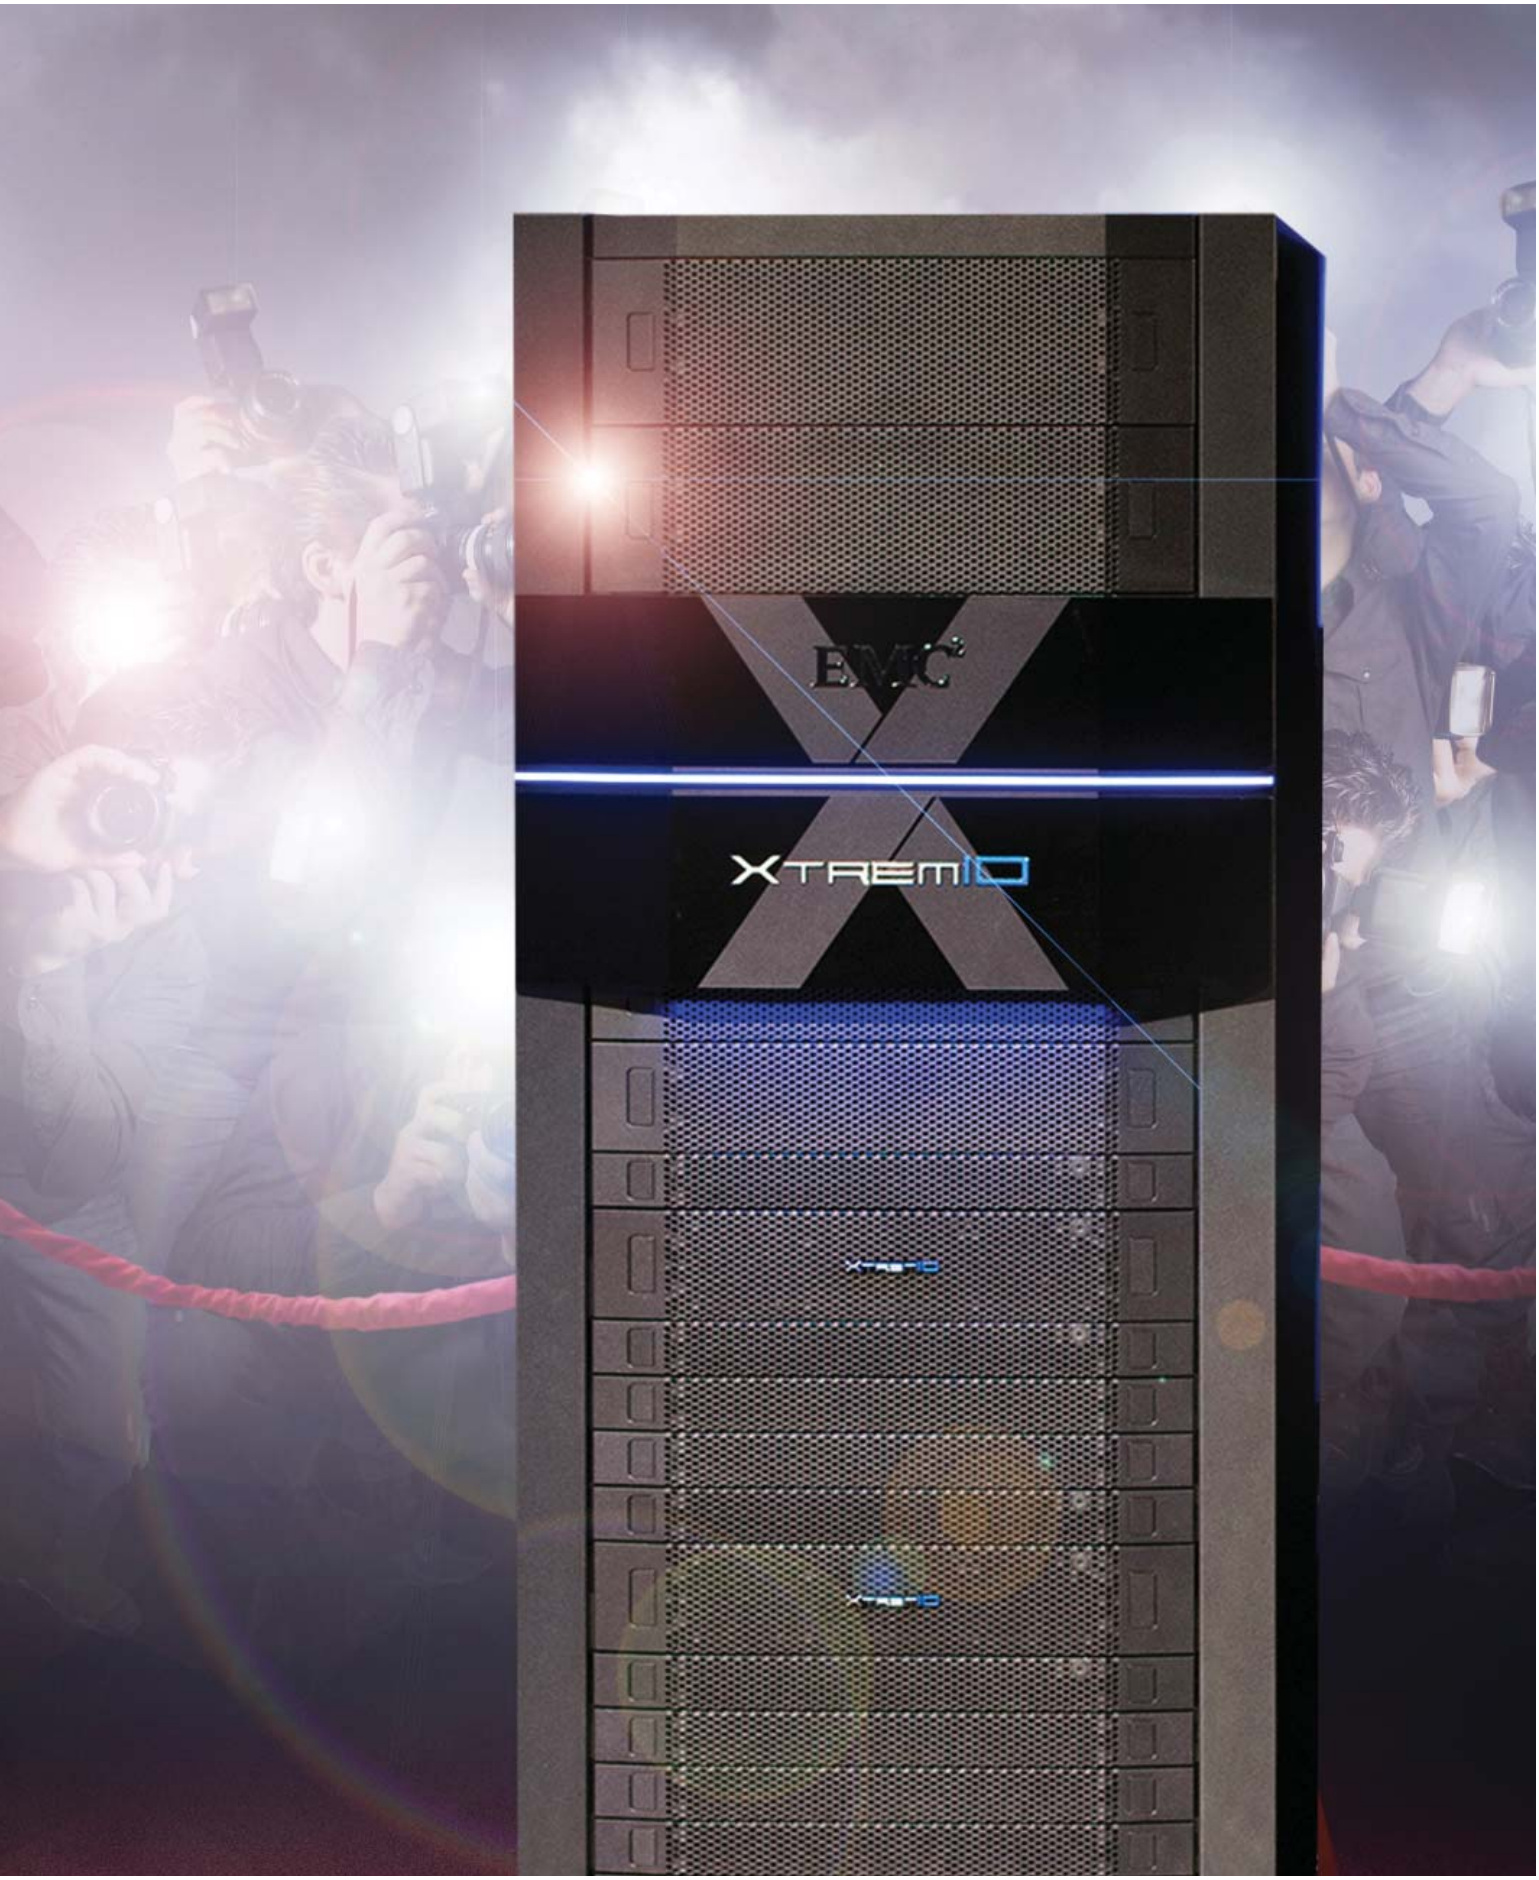
EMC²
XTREMIO

In addition to delivering higher performance than HDDs, flash storage can dramatically reduce power and cooling requirements in the data center and streamline storage administration.

any application workload over any period of time, regardless of whether the array is idle or busy, empty or full.

"Flash storage is seeing growing use in enterprise data centers thanks to lower price points and the need for higher-performance storage for virtualization, cloud computing and 'big data' applications," said Jon Chappell, EMC Competency Center Manager, Sigma Solutions. "With its new XtremIO solution, EMC provides the extreme performance these workloads demand, along with a scale-out architecture, enhanced de-duplication and data protection, and simplified management."

Compelling Story

Dr. Fujio Masuoka invented flash memory in the early 1980s, and the first commercial NAND chip was introduced in 1989. Since that time, continued developments in NAND flash storage have provided the industry with a relatively low-cost nonvolatile storage medium that is compact and efficient.

Today, there are a number of flash storage solutions available for the data center, including direct-attached storage (DAS) for servers, storage subsystems, standalone appliances and hybrid arrays that mix HDDs and flash storage. If one considers the entire realm of flash storage — not just SSDs but the more expensive PCIe flash cards — there is a wide array of options to meet various business requirements and application demands.

In addition to delivering higher performance than HDDs, flash storage can dramatically reduce power and cooling requirements in the data center and streamline storage administration. It does have drawbacks, however. Flash memory can endure only so many write-erase cycles before it wears out — although the longevity of today's SSDs exceeds 100,000 cycles, limited flash endurance still represents a significant downside. And while SSDs are generally more reliable than HDDs, when they do fail they tend to do so suddenly and catastrophically.

The EMC XtremIO solution uses a number of techniques to overcome these drawbacks and provide sustained performance. Industry analysts have taken notice.

"What I find most compelling about XtremIO is the unique architecture," said Laura Dubois, Research VP of Storage at IDC. "Core functions such as granular metadata processing, shared in-memory metadata handling and content-based data placement are enablers to XtremIO's impressive sustained IOPS metrics while offering core services — including deduplication and copy data services. The other standout capabilities with this system are native inline de-duplication, in-memory metadata-only copy and a scale-out architecture. These are attributes not all of the all-flash array solutions on the market offer."

In-Depth Coverage

Four key architectural differences set XtremIO apart from other all-flash arrays. The first is content-based data placement, which determines where to store each "chunk" of data based upon the unique fingerprint it is assigned. Content-based data placement keeps the array inherently load balanced and performance optimized, and automatically performs inline data de-duplication in the data path.

The dual-stage metadata engine allows XtremIO to place data anywhere in the array without the performance overhead associated with system-level garbage collection. Shared in-memory metadata enables XtremIO to deliver the widest range of performance and to rapidly clone information already in the array, accelerating common tasks such as virtual machine deployment.

XtremIO Data Protection (XDP) is a flash-specific algorithm that guards against SSD failures while delivering up to six times more usable capacity than traditional RAID. While other all-flash arrays will begin to show degraded performance when they reach 60 percent to 80 percent capacity, XDP maintains maximum performance at 100 percent capacity. In addition, the efficient XDP algorithms deliver up to four times better performance and flash endurance than RAID in long-term production data center conditions.

XtremIO's scale-out architecture is based upon building blocks called X-Bricks that can be clustered together to increase capacity and performance as needed. Each X-Brick is a high-performance, high-availability SAN appliance with 10TB or 20TB of storage and even greater logical usable capacity. An XtremIO array scales out linearly as X-Bricks are added, while the remote direct memory access (RDMA) Infiniband fabric maintains consistently low latency across the system.

"One XtremIO cluster can scale from two to eight controllers and up to 128 cores, and can handle large virtual desktop environments, OLTP databases and other demanding workloads," Chappell said. "Organizations that are considering a move to an all-flash array should take a look at EMC's extreme performance, highly scalable solution."

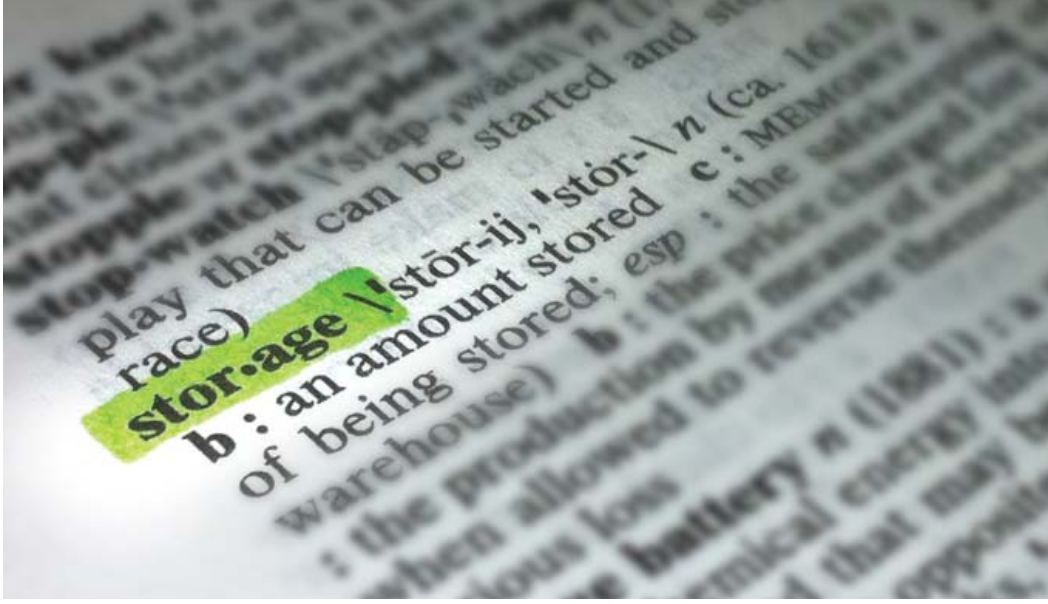

Redefining Storage

EMC's ViPR software-defined storage platform leverages existing IT investments and lays foundation for the future.

The term "software-defined storage" has emerged as the latest buzzword in the larger "software-defined" buzzword mill. The concept derives in part from software-defined networking, which essentially moves the "control plane" of the network from each network device to a controller that works with all the devices.

In a similar fashion, software-defined storage separates the storage controller from storage hardware. The storage controller functionality is performed by software that runs on an appliance or as a virtual machine.

EMC has taken an early lead in the nascent software-defined storage market. In October 2013, the company rolled out EMC ViPR, a software-defined storage platform that enables organizations to manage the control plane through the ViPR Controller and the data plane through ViPR Data Services. ViPR helps organizations drive improvements in automation and lay down a modern storage architecture for future application deployments.

"By decoupling the control plane from the data plane, EMC ViPR enables customers to manage both the storage infrastructure and the data stored within that infrastructure," said James Smith, Director of Solutions Engineering, Sigma Solutions. "Customers can use both together, as in a traditional storage environment, or use only the control plane to manage the underlying intelligence of the storage arrays. This is a radical departure from prior attempts to virtualize storage."

These capabilities are implemented entirely in software that will support EMC, non-EMC and commodity hardware. EMC ViPR also offers the ability to view objects as files and provides file access performance without the latency inherent to object storage.

Separating Control from Data

The EMC ViPR Controller can deliver dramatic improvements in automation because it virtualizes the underlying storage infrastructure. Common storage management functions such as provisioning and migration are abstracted so that different storage arrays can be managed as a single, pooled resource. Once created, these pools of storage are carved up for consumption by applications. For this task, ViPR provides a self-service portal so application owners can browse the storage service catalog and provision service resources best suited for their needs.

For most traditional storage infrastructures utilizing file and block storage, EMC ViPR will only provide the control plane. It discovers storage, creates virtual storage pools, pro-

visions those pools to the application and then gets out of the way.

"With ViPR, file and block storage run as virtual services that take advantage of the capabilities of the underlying storage array," Smith said. "It doesn't introduce any performance overhead — it provides centralized management while staying out of the data path. It leverages the intelligence of the underlying array and offloads processing wherever possible."

Traditional file and block storage represents the majority of application workloads within the data center, and EMC estimates those workloads will grow approximately 70 percent by 2016. But new application workloads are emerging — often operating on vast quantities of data and serving tens of thousands or millions of users. EMC estimates these workloads will grow approximately 700 percent by 2016.

Supporting New Workloads

These new applications require massive scale and a simpler approach to storage infrastructure — specifically object storage. Access methods are also changing from traditional protocols such as NFS and iSCSI to new protocols such as Hadoop Distributed File System (HDFS), the underlying foundation of the Hadoop database.

These new workloads benefit from ViPR's unique ability to view objects as files, which provides file access performance without the latency inherent to object storage. In addition, the ViPR HDFS Data Service enables customers to perform in-place analytics across the entire heterogeneous storage environment, eliminating the need to extract the data to a separate HDFS cluster. The entire storage infrastructure becomes a repository of big data.

EMC ViPR's Object Data Services further support these new application architectures by providing the flexibility and scale of cloud-based data services within the enterprise storage architecture. The ViPR Object Data Services supports Amazon S3- and OpenStack Swift-compatible REST APIs and HDFS access methods — existing applications written to these APIs should run seamlessly. ViPR Object Data Services also support existing EMC Atmos, VNX and Isilon arrays as a persistence layer in addition to third-party arrays and commodity hardware.

"Cloud-based storage platforms such as Amazon S3 are increasingly used to support web-based applications, but these solutions add costs and lack the redundancy, resiliency and data protection inherent in the enterprise storage environment," said Smith. "By supporting REST APIs, ViPR enables customers to determine what the performance, quality of service and cost characteristics should be based upon the application workload."

Part of the Software-Defined Data Center

With ViPR, EMC is doing for storage what VMware has already done for servers — abstracting, pooling and automating the infrastructure. Pools of storage created with EMC ViPR simply appear as an array within VMware vSphere. In addition, the ViPR Controller integrates with the vCloud Automation Center management and orchestration tool as well as vCenter Operations Manager. In this way, storage can be managed by ViPR as an entity in its own right or as an integral part of the VMware software-defined data center.

EMC ViPR is designed for service provider cloud environments but is also suitable for enterprise IT departments that are transforming themselves to offer IT-as-a-service and building out internal web-scale clouds. With this in mind, ViPR has a globally distributed architecture, enabling IT departments to avoid moving massive amounts of data across the network and perform functions like in-place analytics. As devices and data grow, the ViPR platform scales out with no single point of failure and provides an entirely self-managing and self-provisioning environment.

"EMC is sending a clear message that the combination of arrays with a powerful software layer is unbeatable in terms of speed and simplicity," said Vernon Turner, senior vice president, IDC Infrastructure Research Group. "Customers want to extract more value from their storage investments while scaling back on management, and ViPR meets these needs while embracing open architecture and catering to all arrays."

Traditional file and block storage represents the majority of application workloads within the data center, and EMC estimates those workloads will grow approximately 70 percent by 2016. But new application workloads are emerging — often operating on vast quantities of data and serving tens of thousands or millions of users. EMC estimates these workloads will grow approximately 700 percent by 2016.

Ready for Prime Time?

As de-duplication becomes mainstream technology, focus is shifting from backup to primary storage.

The relentless growth of data remains one of IT's biggest challenges. Annual data growth rates approach 70 percent in some organizations, fueled by "big data" initiatives, mobile data requirements, cloud storage and the rapid increase of connected devices composing the so-called "Internet of Things."

Data de-duplication is recognized as an effective tool in combatting this runaway growth. De-duplication eliminates redundant data segments during the backup and recovery process, capturing and storing only unique elements that have been added or changed since the last backup. This can reduce the amount of data to be backed up by 90 percent or more, slash bandwidth requirements by up to 99 percent and make it economical to store backups on disk rather than tape.

To date, de-duplication's sweet spot has been in backup applications dealing with large data sets and many redundancies. Now, however, there is an industry-wide push to extend the benefits of de-duplication to primary, or Tier 1, storage.

Eye on Performance

Primary storage is for active data that is frequently accessed by both end-users and key applications. With mission-critical software such as databases, email and transaction processing applications, it is the source of the vast amounts of data that organizations generate and use to run their businesses. It is also the most expensive storage tier, requiring high performance, low latency and high availability.

Previously, de-duplication methods have simply required too much performance overhead for Tier 1. Typically, a hashing algorithm identifies "chunks" of data, which are then inspected for duplicate patterns in the data. Once found, the de-duplication software replaces copies of the pattern with pointers in its file system to the initial instance. That is an I/O-intensive process that is acceptable for backup data, which can be de-duped during off-peak hours. Active data, however, can't spare that type of real-time computational resources without significant performance degradation.

Another reason for the focus on backup is that there simply isn't as much redundancy in primary storage. While de-duplication may result in up to 20:1 reductions in backup and archival data, it yields just 2:1 reductions in primary storage at best.

Game is Changing

However, refined implementation methods and the economics of virtualized environments now make primary storage de-duplication a much more practical and desirable proposition. While the data reduction ratio won't be as spectacular as it is in backup, that 2:1 reduction in primary storage can still deliver some impressive benefits. If organizations can reduce their primary storage footprint, they will not only save capacity in Tier 1 but create a waterfall of capacity savings

throughout the storage infrastructure and cost savings across the data lifecycle.

When combined with other improvements, the capacity savings make it feasible to leverage faster but more expensive solid-state drives (SSDs). The superior I/O performance of SSDs, which can be employed in al-flash arrays or more economical hybrid systems, dramatically reduces the performance overhead in primary storage de-duplication. But even in disk-based arrays, new techniques such as content-based caching, dynamic replica retrieval and selective de-duplication can deliver significant I/O performance improvements.

Increased dependence on virtualization is also creating a more compelling use case for primary storage de-duplication. Virtual machines (VMs) are typically cloned from a small number of master images, but these VM files tend to be large and take up a lot of primary disk space. In such virtualized environments, where hundreds of VMs may be hosted on a handful of servers, primary storage de-duplication can bring data reduction ratios of up to 100:1.

Manufacturers on Board

A stronger business case for primary storage de-duplication has led more vendors to offer such solutions. For example, data de-duplication is built into Microsoft Windows Server 2012 R2 for primary storage. Administrators can minimize performance issues by scheduling data de-duplication jobs at specific times and configuring policies to control which files should be processed. Microsoft's data de-duplication feature promises to provide data integrity, bandwidth efficiency and faster download times.

Data de-duplication was introduced in Windows Server 2012, but significant enhancements were made for R2. Where de-duplication initially focused on files at rest and would skip any file that was in active use, Windows Server 2012 R2 now has the ability to de-duplicate live virtual desktop infrastructure (VDI) workloads, with negligible impact on read/write performance.

Other manufacturers have also redoubled their efforts to bring effective primary storage de-duplication to market. Hitachi Data Systems released its Hitachi NAS (HNAS) platform in 2013, which uses an object-based file system offload engine (FOE) powered by field-programmable gate arrays (FPGAs) to accelerate hashing and chunking — the most processor-intensive part of de-duplication. EMC, Dell, HP and others have intensified their efforts in this area as well.

As it enters the mainstream, primary storage de-duplication is now likely to be built into enterprise-class storage arrays, all-flash platforms and cloud-integrated storage systems. While it is not ideal for all data types, primary storage data-de-duplication attacks the problem of rampant data growth at the source, helping organizations to reduce storage costs and improve efficiency throughout the storage infrastructure.

De-Duplication Improves Flash Efficiency

De-duplication has become a mainstream technology for backup, but its application in primary storage could hold the key to accelerating adoption of flash-based storage technologies in the data center.

NAND flash solid-state drives (SSDs) are smaller, faster and more durable than the highest-quality hard-disk drives (HDDs), and they have already become standard in laptops and notebooks. By using memory instead of spinning magnetic disks to store and retrieve data, SSDs enable read/write response times that are exponentially faster than the best HDDs. What's more, they use less power and deliver "instant-on" capabilities for lightning-fast access to cold data.

Price considerations have limited the practicality of SSDs in the data center to date. Enterprise-grade SSDs generally cost more than $1 per GB, and often more than $2 per GB. By comparison, high-quality SATA hard drives can be had for about $0.10 per GB. Although flash memory prices have fallen over the past few years, the all-flash data center won't be happening anytime soon.

Data de-duplication closes the cost gap significantly, however, and opens the door for opportunistic use of SSD — particularly in hybrid storage arrays in which flash acts as a performance booster for large pools of disk. De-duplication combined with other inline data reduction techniques such as compression and thin provisioning can dramatically reduce the data footprint, making the logical size of the drive appear to be up to 10 times larger than its actual raw physical size. That creates a much more tolerable effective cost for SSD.

De-duplication also reduces writes, greatly increasing the reliability and lifespan of an SSD. Current SSD technology suffers from a phenomenon called "write amplification," where write performance degrades throughout the life of the drive. "Wear leveling" algorithms mitigate the effect, but all flash-based drives eventually hit their limit. By examining data segments for redundancy prior to writing to disk, de-duplication reduces both the number of writes and the amount of data written to the drive.

Your solution to the cloud puzzle

CloudSource from Sigma Solutions

The benefits of cloud computing have been well established, but implementation challenges continue to perplex many organizations. Sigma Solutions' new CloudSource offering is designed to put all the pieces in place for a robust cloud solution. Leveraging partnerships with key cloud vendors, Sigma Solutions provides access to best-of-breed cloud services and architectures, which are combined with our own professional services in a comprehensive cloud delivery package. We'll help you determine the proper cloud strategy, migrate to the most appropriate environment and manage the infrastructure on an ongoing basis. Contact us today to learn more!

www.sigmasol.com | 888.895.0495

SIGMA SOLUTIONS