

# TrueNAS Plugins Converge Services for Simple Hybrid Cloud Enablement

iXsystems is taking simplified service delivery to a new level by enabling a curated set of third-party services to run directly on its TrueNAS arrays. TrueNAS already provided multi-protocol unified storage to include file, block and S3-compatible object storage. Now preconfigured plugins converge additional services onto TrueNAS for simple hybrid cloud enablement.

## TrueNAS Technology Provides a Robust Foundation for Hybrid Cloud Functionality

[iXsystems](#) is known for enterprise-class storage software and rock-solid storage hardware. This foundation lets iXsystems customers run select third-party applications as plugins directly on the storage arrays—whether TrueNAS, FreeNAS Mini or FreeNAS Certified. Several of these plugins dramatically simplify the deployment of hybrid public and private clouds.

## How it Works

iXsystems works with select technology partners to preconfigure their solutions to run on TrueNAS using FreeBSD jails, iocage plugins, and bhyve virtual machines. By collaborating with these technology partners, iXsystems enables rapid IT service delivery and drives down the total cost of technology infrastructure. The flexibility to extend TrueNAS functionality via these plugins transforms the appliances into complete solutions that streamline common workflows.

## Benefits of Curated Third-party Service Plugins

There are many advantages to this pre-integrated plugin approach:

- Plugins are preconfigured for optimal operation on TrueNAS
- Services can be added any time through the web interface
- Simply turn it on, download the plugin and enter the associated login credentials
- Plugins eliminate network latency by moving processing to the storage array
- Third party applications can be run in a virtual machine without purchasing separate server hardware

## Hybrid Cloud Data Protection

The integrated [Asigra](#) Cloud Backup software protects cloud, physical, and virtual environments. It is an enterprise-class backup solution that uniquely helps prevent malware from compromising backups. Asigra embeds cybersecurity software in its Cloud Backup software. It goes the extra mile to protect backup repositories, ensuring businesses can recover from malware attacks in their production environments.

Asigra is also one of the only enterprise backup solutions that offers agentless backup support across all types of environments: cloud, physical, and virtual. This flexibility makes adopting and deploying Asigra Cloud Backup easy with zero disruption to clients and servers. The integration of Asigra with TrueNAS is [Storage Magazine's Backup Product of the year for 2018](#).

## Hybrid Cloud Media Management

TrueNAS arrays from iXsystems are heavily used in the media and entertainment industry, including several major film and television studios. iXsystems storage accelerates workflows with any device file sharing, multi-tier caching technology,

and the latest interconnect technologies on the marketplace.

iXsystems recently announced a partnership with [Cantemo](#) to integrate its iconik software.

iconik is a hybrid cloud-based video and content management hub. Its main purpose is managing processes including ingestion, annotation, cataloging, collaboration, storage, retrieval, and distribution of digital assets. The main strength of the product is the support for managing metadata and transcoding of audio, video, and image files, but can store essentially all file formats. Users can choose to keep large original files on-premise yet still view and access the entire library in the cloud using proxy versions where required.

The Cantemo solutions are used to manage media across the entire asset lifecycle, from ingest to archive. iconik is used across a variety of industries including Fortune 500 IT companies, advertising agencies, broadcasters, houses of worship, and media production houses. Cantemo's clients include BBC Worldwide, Nike, Madison Square Garden, The Daily Telegraph, The Guardian and many other leading media companies.

Enabling iconik on TrueNAS streamlines multimedia workflows and increases productivity for iXsystems customers who choose to enable the Cantemo service.

## **Cloud Sync**

Both Asigra and Cantemo include hybrid cloud data management capabilities within their feature sets. iXsystems also supports file synchronization with many business-oriented and personal public cloud storage services. These enable staff to be productive anywhere—whether working with files locally or in the cloud.

Supported public cloud providers include Amazon Cloud Drive, Amazon S3, Backblaze B2, Box, Dropbox, Google Cloud Storage,

Google Drive, Hubic, Mega, Microsoft Azure Blob Storage, Microsoft OneDrive, pCloud and Yandex. The Cloud Sync tool also supports file sync via SFTP and WebDAV.

## **More Technology Partnerships Planned**

According to iXsystems, they will extend TrueNAS pre-integration to more technology partners where such partnerships provide win-win benefits for all involved. This intelligent strategy allows iXsystems to focus on enhancing core TrueNAS storage services, and it enables TrueNAS customers to quickly and confidently implement best-of-breed applications directly on their TrueNAS arrays.

## **All TrueNAS Owners Benefit**

TrueNAS plugins provide a simple and flexible way for all iXsystems customers to add sophisticated hybrid-cloud media management and data protection services to their IT environments. Existing TrueNAS customers can gain the benefits of this plugin capability by updating to the most recent version of the TrueNAS software.

---

# **Latest Enhancements to Dell DL1300 Provide the Out-of-the-box Backup, Recovery, and DR Experience that Mid-market Companies Demand**

More data to backup, less time to recover it, heightened recovery expectations and limited time to dedicate to manage

these tasks. These are the dilemmas that every mid-market business faces when backing up and recovering its data. The good news is that the DL1300 Backup and Recovery Appliance offers the specific features that mid-market companies need to address these issues. Delivered as a turn-key, easy-to-deploy solution, the DL1300 offers the comprehensive set of features that mid-market companies need to reduce their time spent on backups, replication and/or archiving data to low cost 3rd party cloud locations.

### **Faster, Easier Recoveries Top the List of Mid-market Company Needs**

Today, perhaps more so than ever, organizations want a backup appliance that makes all of the tasks associated with managing backup and recovery faster and easier. This includes making the appliance easier to deploy and manage that minimizes the amount of time and IT manpower needed to maintain and scale the appliance.

Foremost, organizations want a backup appliance that possesses the features to quickly and easily recover their applications and data. Most mid-market organizations already rely on disk-based backup in some form as part of their existing backup process to improve backup success rates as well as shorten backup windows. As such, one might assume that faster, easier recoveries go hand-in-glove with faster, easier backups.

Unfortunately, this rarely holds true as recovery has largely lagged behind backup in its ability to deliver on either “faster” or “easier”. Many organizations grapple with the aggressive recovery point objectives (RPOs) and recovery time objectives (RTOs) for their applications and data putting added pressure on them and any solution they evaluate to:

- Recovering their mission critical applications in minutes
- Quickly restore any of their applications and data

(under an hour)

- Scale to keep pace with their growing storage capacity requirements for their applications and data
- Update or upgrade the appliance with minimal to no impact to business operations
- Utilize 3rd party cloud services providers for archive and disaster recovery

In short, they want a solution that maximizes their investment in a backup appliance which offers them the best possible backup, recovery and archive experience for their ever changing environment.

### **The Fast, Easy Backup Experience that Mid-market Companies Expect**

The DL1300 Backup and Recovery Appliances positions midmarket organizations to address these key challenges that they face. The [DL1300](#) is built on the Dell 13G PowerEdge Server platform that hosts the DL1300's backup software and provides the resources that organizations need to consistently and quickly backup and recover large amounts of data. Marrying disk and application software in a turn-key solution, the [DL1300](#):

- Offers configuration wizards so organizations may go from box-to-backup in approximately 20 minutes
- Takes the guesswork out of deployment and ongoing management and maintenance
- Delivers data deduplication and compression to reduce data footprints and shorten backup windows
- Scales to quickly and easily add capacity with in-box capacity upgrades as well as supports external storage arrays for greater capacity



An initial deployment of a DL1300 Backup and Recovery Appliance may be future-proofed to meet their ever-growing and difficult-to-predict capacity growth without replacing the entire appliance.

Leveraging this option, mid-market organizations may size the [DL1300](#) to solve their immediate backup challenges and then easily scale it up if and when the amount of data they backup increases. The DL1300 ships in 3 capacities: 2TB, 3TB and 4TB with 2 VMs included with the 3TB and 4TB models. All three models can be expanded to include 8TB of capacity inside the appliance. However, the 4TB DL1300 model can be increased to a total capacity of 18TBs using one MD1400 external storage array. Additionally, the RAM on all models may be upgraded to 64GB.

These extra levels of availability, capacity and horsepower particularly come into play for mid-market organizations looking to consolidate and centralize backup, recovery, and cloud archive across their environment.

As many mid-market organizations have a mix of Linux, Windows and VMware in their environment, the DL1300 provides them the solution they need to backup and recover all of these operating systems as well as protect both physical and virtual machines. Equally important, it integrates tightly with leading Windows applications such as Active Directory (AD), SharePoint and SQL Server to provide the application consistent backups that each of these applications need.

Once setup and configured, organizations may concurrently backup up to 60 backup streams. The DL1300 then optimizes its

available disk capacity using RAID 5. This RAID configuration uses the DL1300's available disk capacity more efficiently than what a RAID 1 (mirrored disk drive) configuration natively offers while still protecting the backup data against data loss should a drive failure occur.

### **More Recovery Options for All Applications and Data**

The DL1300's increased amount of disk storage capacity coupled with its RAID 5 data protection scheme facilitates the ability of organizations to centralize the backup of all of their applications and data onto a single platform in one location. This serves an important purpose: organizations may then leverage the DL1300 to quickly, easily, and centrally recover all of their applications and data.

The DL1300 possibly does a better job of doing recovery than any other backup appliance in its class as organizations may recover across both physical and virtual machines, to include doing P2V, V2P, P2P and V2V conversions. Its Live Recovery feature can recover any protected application—physical or virtual—back to the production machine, usually in under an hour and often within minutes.

To perform the restore, Live Recovery restores data from the application's most recent backup checkpoint. As backups typically occur each hour, when an organization initiates a restore the organization can reasonably expect to recover application data that is less than an hour old.

Organizations that need even higher levels of application availability (no more than minutes of downtime) for their mission critical applications such as Microsoft Exchange or SQL Server, may leverage the optional Virtual Standby feature found on either the DL1300 3TB and 4TB models.

Using this feature, an organization may create a couple of virtual machines (VMs) on the [DL1300](#). Then, for up to two selected protected servers, it will create a standby VM on the



DL1300. Once set up, if the production physical or virtual machine goes offline, a virtual copy may be started almost immediately on the DL1300 appliance.

Finally, all mid-market organizations increasingly need a means to get their data offsite for archive and prepare an offsite disaster recovery (DR) plan. This is where the DL1300 shines. It supports multiple cloud storage providers to include Amazon, eFolder, Microsoft Azure, Rackspace and Zerolag. In so doing, it positions mid-market organizations to store data with their provider of choice providing them with cost-effective, non-proprietary options to move data and prepare their DR readiness plan.

### **DL1300 Provides the Out-of-the-Box Backup, Recovery and DR Solution that Mid-market Organizations Demand**

Many mid-market organizations find themselves at a crossroads. They need a solution that offers the scalability and technical features they need to protect their ever-changing environment in a turn-key and easy-to-manage package. Organizations need to improve application recovery, shorten backup windows, archive to non-proprietary 3rd party cloud locations and prepare a cohesive DR readiness plan. Further, they need to do so without breaking their budget, or stretching their technical limits.

The DL1300 meets these objectives by providing the out-of-the-box backup, recovery, and DR solution that organizations demand. Starting at an attractive price point under \$5,000 and available from Dell and its partners, the DL1300 puts mid-market organizations on a path toward consolidating and simplifying their backup operations even as they get new flexibility to recover and scale going forward.

---

# HP 3PAR StoreServ 8000 Series Lays Foundation for Flash Lift-off

Almost any hybrid or all-flash storage array will accelerate performance for the applications it hosts. Yet many organizations need a storage array that scales beyond just accelerating the performance of a few hosts. They want a solution that both solves their immediate performance challenges and serves as a launch pad to using flash more broadly in their environment.

Yet putting flash in legacy storage arrays is not the right approach to accomplish this objective. Enterprise-wide flash deployments require purpose-built hardware backed by Tier-1 data services. The HP 3PAR [StoreServ 8000 series](#) provides a fundamentally different hardware architecture and complements this architecture with mature software services. Together these features provide organizations the foundation they need to realize flash's performance benefits while positioning them to expand their use of flash going forward.

## A Hardware Foundation for Flash Success

Organizations almost always want to immediately realize the performance benefits of flash and the HP 3PAR [StoreServ 8000 series](#) delivers on this expectation. While flash-based storage arrays use various hardware options for flash acceleration, the 8000 series complements the enterprise-class flash HP 3PAR StoreServ 20000 series while separating itself from competitive flash arrays in the following key ways:

- **Scalable, Mesh-Active architecture.** An Active-Active controller configuration and a scale-out architecture are considered the best of traditional and next-generation array architectures. The HP 3PAR StoreServ

8000 series brings these options together with its Mesh-Active architecture which provides high-speed, synchronized communication between the up-to-four controllers within the 8000 series.

- **No internal performance bottlenecks.** One of the secrets to the 8000's ability to successfully transition from managing HDDs to SSDs and still deliver on flash's performance benefits is its programmable ASIC. The HP 3PAR ASIC, now it's 5<sup>th</sup> generation, is programmed to manage flash and optimize its performance, enabling the 8000 series to achieve **over 1 million IOPs**.
- **Lower costs without compromise.** Organizations may use lower-cost commercial MLC SSDs (cMLC SSDs) in any 8000 series array. Then leveraging its Adaptive Sparing technology and Gen5 ASIC, it optimizes capacity utilization within cMLC SSDs to achieve high levels of performance, extends media lifespan which are backed by a 5-year warranty, and increases usable drive capacity by up to 20 percent.
- **Designed for enterprise consolidation.** The 8000 series offers both 16Gb FC and 10Gb Ethernet host-facing ports. These give organizations the flexibility to connect performance-intensive applications using Fibre Channel or cost-sensitive applications via either iSCSI or NAS using the 8000 series' File Persona feature. Using the 8000 Series, organizations can start with configurations as small as 3TB of usable flash capacity and scale to 7.3TB of usable flash capacity.

## A Flash Launch Pad

As important as hardware is to experiencing success with flash on the 8000 series, HP made a strategic decision to ensure its converged flash and all-flash 8000 series models deliver the same mature set of data services that it has offered on its all-HDD HP 3PAR StoreServ systems. This frees organizations to move forward in their consolidation initiatives knowing that

they can meet enterprise resiliency, performance, and high availability expectations even as the 8000 series scales over time to meet future requirements.

For instance, as organizations consolidate applications and their data on the 8000 series, they will typically consume less storage capacity using the 8000 series' native thin provisioning and deduplication features. While storage savings vary, HP finds these features usually result in about 4:1 data reduction ratio which helps to drive down the effective price of flash on an 8000 series array to as low as \$1.50/GB.

Maybe more importantly, organizations will see minimal to no slowdown in application performance even as they implement these features, as they may be turned on even when running mixed production workloads. The 8000 series compacts data and accelerates application performance by again leveraging its Gen5 ASICs to do system-wide striping and optimize flash media for performance.

Having addressed these initial business concerns around cost and performance, the 8000 series also brings along the HP 3PAR StoreServ's existing data management services that enable organizations to effectively manage and protect mission-critical applications and data. Some of these options include:

- ***Accelerated data protection and recovery.*** Using HP's Recovery Manager Central ([RMC](#)), organizations may accelerate and centralize application data protection and recovery. RMC can schedule and manage snapshots on the 8000 series and then directly copy those snapshots to and from HP StoreOnce without the use of a third-party backup application.
- ***Continuous application availability.*** The HP 3PAR [Remote Copy](#) software either asynchronously or synchronously replicates data to another location. This provides recovery point objectives (RPOs) of minutes, seconds, or even non-disruptive application failover.

- **Delivering on service level agreements (SLAs).** The 8000 series' Quality of Service ([QoS](#)) feature ensures high priority applications get access to the resources they need over lower priority ones to include setting sub-millisecond response times for these applications. However QoS also ensures lower priority applications are serviced and not crowded out by higher priority applications.
- **Data mobility.** HP 3PAR StoreServ creates a federated storage pool to facilitate non-disruptive, bi-directional data movement between any of up to four (4) midrange or high end HP 3PAR arrays.

### **Onboarding Made Fast and Easy**

Despite the benefits that flash technology offers and the various hardware and software features that the 8000 series provides to deliver on flash's promise, migrating data to the 8000 series is sometimes viewed as the biggest obstacle to its adoption. As organizations may already have a storage array in their environment, moving its data to the 8000 series can be both complicated and time-consuming. To deal with these concerns, HP provides a relatively fast and easy process for organizations to migrate data to the 8000 series.

In as few as five steps, existing hosts may discover the 8000 series and then access their existing data on their old array through the 8000 series without requiring the use of any external appliance. As hosts switch to using the 8000 series as their primary array, Online Import non-disruptively copies data from the old array to the 8000 series in the background. As it migrates the data, the 8000 series also reduces the storage footprint by as much as 75 percent using its thin-aware functionality which only copies blocks which contain data as opposed to copying all blocks in a particular volume.

Maybe most importantly, data migrations from EMC, HDS or HP EVA arrays (*and others to come*) to the 8000 series may occur

in real time Hosts read data from volumes on either the old array or the new 8000 series with hosts only writing to the 8000 series. Once all data is migrated, access to volumes on the old array is discontinued.

### **Achieve Flash Lift-off Using the HP 3PAR StoreServ 8000 Series**

Organizations want to introduce flash into their environment but they want to do so in a manner that lays a foundation for their broader use of flash going forward without creating a new storage silo that they need to manage in the near term.

The HP 3PAR StoreServ 8000 series delivers on these competing requirements. Its robust hardware and mature data services work hand-in-hand to provide both the high levels of performance and Tier-1 resiliency that organizations need to reliably and confidently use flash now and then expand its use in the future. Further, they can achieve lift-off with flash as they can proceed without worrying about how they will either keep their mission-critical apps online or cost-effectively migrate, protect or manage their data once it is hosted on flash.

---

## **The Dell DL4300 Puts the Type of Thrills into Backup and Recovery that Organizations Really Want**

Organizations have long wanted to experience the thrills of non-disruptive backups and instant application recoveries. Yet the solutions delivered to date have largely been the exact

opposite offering only unwanted backup pain with very few of the types of recovery thrills that organizations truly desire. The new Dell DL4300 Backup and Recovery Appliance successfully takes the pain out of daily backup and puts the right types of thrills into the backup and recovery experience.

Everyone enjoys a thrill now and then. However individuals should want to get their thrills at an amusement park, not when they backup or recover applications or manage the appliance that hosts their software. In cases like these, boring is the goal when it comes to performing backups and/or managing the appliance that hosts the software with the excitement and thrills appropriately reserved for fast, successful application recoveries. This is where the latest [Dell DL4300](#) Backup and Recovery Appliance introduces the right mix of boring and excitement into today's organizations.

## Show Off

Being a show off is rarely if ever perceived as a “*good thing.*” However IT staff can now in good conscience show off a bit by demonstrating the DL4300's value to the business as it quickly backs up and recovers applications without putting business operations at risk. The Dell DL4300 Backup and Recovery Appliance's AppAssure software provides the following five (5) key features to give them this ability:

- ***Near-continuous backups.*** The Dell DL4300 may perform application backups as frequently as every five (5) minutes for both physical and virtual machines. During the short period of time it takes to complete a backup, it only consumes a minimal amount of system resources – no more than 2 percent. Since the backups occur so quickly, organizations have the flexibility to schedule as many as 288 backups in a 24 hour period which helps to minimize the possibility of data loss so organizations can achieve near-real time recovery point objectives (RPOs).

- ***Near-instantaneous recoveries.*** The Dell DL4300 complements its near-continuous backup functionality by also offering near-instantaneous application recoveries. Its [Live Recovery](#) feature works across both physical and virtual machines and is intended for use in situations where application data is corrupted or becomes unavailable. In those circumstances, Live Recovery can within minutes present data residing on non-system volumes to a physical or virtual machine. The application may then access that data and resume operations until the data is restored and/or available locally.
- ***Virtual Standby.*** The Dell DL4300's [Virtual Standby](#) feature complements its Live Recovery feature by providing an even higher level of availability and recovery for those physical or virtual machines that need this level of recovery. To take advantage of this feature, organizations identify production applications that need instant recovery. Once identified, these applications are associated with the up to four (4) virtual machines (VMs) that may be hosted by a Dell DL4300 appliance and which are kept in a "standby" state. While in this state, the Standby VM on the DL4300 is kept updated with changes on the production physical or virtual VM. Then should the production server ever go offline, the standby VM on the Dell DL4300 will promptly come online and take over application operations.
- ***Helps to insure application consistent recoveries.*** Simply being able to bring up a Standby VM on a moment's notice for some production applications may be insufficient. Some applications such as Microsoft Exchange create check points to ensure it is brought up in an application consistent state. In cases such as these, the DL4300 integrates with applications such as Exchange by regularly performing [mount checks](#) for specific Exchange server recovery points. These mount checks help to guarantee the recoverability of Microsoft



Exchange.

- **Open Cloud [support](#)**. As more organizations keep their backup data on disk in their data center, many still need to retain copies of data offsite without either moving it to tape or needing to set up a secondary site to which to replicate the data. This makes integration with public cloud storage providers to archive retention backup copies an imperative. The Dell DL4300 meets this requirement by providing one of the broadest levels of public cloud storage integration available as it natively integrates with Amazon S3, Microsoft Azure, OpenStack and Rackspace Cloud Block storage.

## **The Thrill of Having Peace of Mind**

The latest Dell DL4300 series goes a long way towards introducing the type of excitement that organizations really want to experience when they use an integrated backup appliance. It also goes an equally long way toward providing the type of peace of mind that organizations want when implementing a backup appliance or managing it long term.

For instance, the Dell DL4300 gives organization the flexibility to start small and scale as needed in both its Standard and High Capacity models with their capacity on demand license features. The Dell DL4300 Standard comes equipped with 5TB of licensed capacity and a total of 13TB of usable capacity. Similarly, the Dell DL4300 High Capacity ships with 40TB of licensed capacity and 78TB of usable capacity.

Configured in this fashion, DL4300 series minimizes or even eliminates the need for organizations to install additional storage capacity at a later date should its existing, available licensed capacity ever run out of room. If the 5TB threshold is reached on the DL4300 Standard or the 40TB limit is reached on the DL4300 High Capacity, organizations only need to acquire an upgrade license to access and use the pre-

installed and existing additional capacity. This takes away the unwanted worry about later upgrades as organizations may easily and non-disruptively add 5TB of additional capacity to the DL4300 Standard or 20TB of additional capacity to the DL4300 High Capacity.

Similarly the DL4300's Rapid Appliance Software Recovery (RASR) removes the shock of being unable to recover the appliance should it fail. RASR improves the reliability and recoverability of the appliance by taking regularly scheduled backups of the appliance. Then should the appliance itself ever experience data corruption or fail, organizations may first do a default restore to the original backup appliance configuration from an internal SD card and then restore from a recent backup to bring the appliance back up-to-date.

### **The Dell DL4300 Provides the Types of Thrills that Organizations Want**

Organizations want the sizzle that today's latest technologies have to offer without the unexpected worries that can too often accompany them. The Dell DL4300 provides this experience. It makes its ongoing management largely a non-issue so organizations may experience the thrills of near-continuous backup and near-instantaneous recovery of data and applications across their physical, virtual and/or cloud infrastructures.

It also delivers the new type of functionality that organizations want to meet their needs now and into the future. Through its native integration with multiple public cloud storage providers and giving organizations the flexibility to use its virtual standby feature for enhanced testing to insure consistent and timely recovery of their data, organizations get the type of thrills that they want and should rightfully expect from a solution such as the Dell DL4300 Backup and Recovery appliance that offers industry-leading self-recovery features and enhanced appliance

management.

---

# **HP 3PAR StoreServ's VVols Integration Brings Long Awaited Storage Automation, Optimization and Simplification to Virtualized Environments**

VMware Virtual Volumes ([VVols](#)) stands poised to fundamentally and positively change storage management in highly virtualized environments that use VMware vSphere. However enterprises will only realize the full benefits that VVols have to offer by implementing a backend storage array that stands ready to take advantage of the VVols architecture. The HP 3PAR StoreServ family of arrays provide the virtualization-first architecture along with the simplicity of implementation and ongoing management that organizations need to realize the benefits that the VVols architecture provide short and long term.

## **VVols Changes the Storage Management Conversation**

VVols eliminate many of the undesirable aspects associated with managing external storage array volumes in networked virtualized infrastructures today. Using storage arrays that are externally attached to ESXi servers over either Ethernet or Fibre Channel (FC) storage networks, organizations currently struggle with issues such as:

- Deciding on the optimal block-based protocol to achieve the best mix of cost and performance
- Provisioning storage to ESXi servers
- Lack of visibility into the data placed on LUNs assigned to specific VMs on ESXi servers
- Identifying and reclaiming stranded storage capacity
- Optimizing application performance on these storage arrays

The VVols architecture changes the storage management conversation in virtualized environments that use VMware in the following ways:

- **Protocol agnostic.** VVols minimize or even eliminate deciding on which protocol is “best” as VVols work the same way whether block or file-based protocols are used.
- **Uses pools of storage.** Storage arrays make raw capacity available in a unit known as a VVol Storage Container to one or more ESXi servers. As each VM is created, the VMware ESXi server allocates the proper amount of array capacity that is part of the VVol Storage Container to the VM.
- **Heightened visibility.** Using the latest VMware APIs for Storage Awareness (VASA 2.0), the ESXi server lets the storage array know exactly which array capacity is assigned to and used by each VM.
- **Automated storage management.** Knowing where each VM resides on the array facilitates the implementation of automated storage reclamation routines as well as performance management software. Organizations may also offload functions such as snapshots, thin provisioning and the overhead associated with these tasks onto the storage array.

VVols’ availability make it possible for organizations to move much closer to achieving the automated, non-disruptive, hassle-free storage array management experience in virtualized environments that they want and have been waiting for years to

implement.

## **Robust, VMware ESXi-aligned Storage Platform a Prerequisite to Realizing VVols Potential**

Yet the availability of VVols from VMware does not automatically translate into organizations being able to implement them by simply purchasing and installing any storage array. To realize the potential storage management benefits that VVols offer requires deploying a properly architected storage platform that is aligned with and integrated with VMware ESXi. These requirements make it a prerequisite for organizations to select a storage array that:

- ***Is highly virtualized.*** Each time array capacity is allocated to a VM, a virtual volume must be created on the storage array. Allocating a virtual volume that performs well and uses the most appropriate tier of storage for each VM requires a highly virtualized array.
- ***Supports VVols.*** VVols represent a significant departure from how storage capacity has been managed to date in VMware environments. As such, the storage array must support VVols.
- ***Tightly integrates with VMware VASA.*** Simplifying storage management only occurs if a storage array tightly integrates with VMware VASA. This integration automates tasks such as allocating virtual volumes to specific VMs, monitoring and managing performance on individual virtual volumes and reclaiming freed and stranded capacity on those volumes.

## **HP 3PAR StoreServ: Locked and Loaded with VVols Support**

The [HP 3PAR StoreServ](#) family of arrays come locked and loaded with VVols support. This enables any virtualized environment running VMware vSphere 6.0 on its ESXi hosts to use a VVol protocol endpoint to directly communicate with HP 3PAR StoreServ storage arrays running the HP 3PAR OS 3.2.1 MU2 P12

or later software.

Using FC protocols, the ESXi server(s) integrates with the HP 3PAR StoreServ array using the various APIs natively found in VMware vSphere. A VASA Provider is directly built into HP 3PAR StoreServ arrays which recognizes vSphere commands. It then automatically performs the appropriate storage management operations such as carving up and allocating a portion of the HP 3PAR StoreServ storage array capacity to a specific VM or reclaiming the capacity associated with a VM that has been deleted and is no longer needed.

Yet perhaps what makes HP 3PAR StoreServ's support of VVols most compelling is that the pre-existing HP 3PAR OS software carries forward. This gives the VMs created on a VVols Storage Container on the HP 3PAR StoreServ array access to all of the same, powerful data management services that were previously only available at the VMFS level on HP 3PAR StoreServ LUNs. These services include:

- **Adaptive Flash Cache** that dedicates a portion of the HP 3PAR StoreServ's available SSD capacity to augment its available primary cache and then accelerates response times for applications with read-intensive I/O workloads.
- **Adaptive Optimization** that optimizes service levels by matching data with the most cost-efficient resource on the HP 3PAR StoreServ system to meet that application's service level agreement (SLA).
- **Priority Optimization** that identifies exactly what storage capacity is being utilized by each VM and then places that data on the most appropriate storage tier according to each application's SLA so a minimum performance goal for each VM is assured and maintained.
- **Thin Deduplication** that first assigns a unique hash to each incoming write I/O. It then leverages HP 3PAR's Thin Provisioning metadata lookup table to quickly do hash comparisons, identify duplicate data and, when

matches are found, to deduplicate like data.

- **Thin Provisioning** that only allocates very small chunks of capacity (16 KB) when writes actually occur.
- **Thin Persistence** that reclaims allocated but unused capacity on virtual volumes without manual intervention or VM timeouts.
- **Virtual Copy** that can create up to 2,048 point-in-time snapshots of each virtual volume with up to 256 of them being available for read-write access.
- **Virtual Domains**, also known as virtual private arrays, offer secure multi-tenancy for different applications and/or user groups. Each Virtual Domain may then be assigned its own service level.
- **Zero Detect** that is used when migrating volumes from other storage arrays to HP 3PAR arrays. The Zero Detect technology identifies “zeroes” on existing volumes which represent allocated but unused space on those volumes. As HP 3PAR migrates these external volumes to HP 3PAR volumes, the zeroes are identified but not migrated so the space may be reclaimed on the new HP 3PAR volume.

## **HP 3PAR StoreServ and VVols Bring Together Storage Automation, Optimization and Simplification**

HP 3PAR StoreServ arrays are architected and built from the ground up to meet the specific storage requirements of virtualized environments. However VMware’s introduction of VVols further affirms this virtualization-first design of the HP 3PAR StoreServ storage arrays as together they put storage automation, optimization and simplification within an organization’s reach.

HP 3PAR StoreServ frees organizations to immediately implement the new VVols storage architecture and take advantage of the granularity of storage management that they offer. By HP 3PAR StoreServ immediately integrating and supporting VVols and bringing forward its existing, mature set of data management services, organizations can take a long awaited step forward

to automate and simplify the deployment and ongoing storage management of VMs in their VMware environment.

---

# HP StoreOnce Deduplicating Backup Appliances Put Organizations on Path to Ending Big Data Backup Headaches

During the recent HP Deep Dive Analyst Event in its Fremont, CA, offices, HP shared some notable insights into the percentage of backup jobs that complete successfully (*and unsuccessfully*) within end-user organizations. Among its observations using the anonymized data gathered from hundreds of backup assessments at end-user organizations of all sizes, HP found that over 60% of them had backup job success rates of 98% or lower, with 12% of organizations showing backup success rates of lower than 90%. Yet what is more noteworthy is through HP's use of Big Data analytics, it has identified large backups (*those that take more than 12 hours to complete*) as being the primary contributor to the backup headaches that organizations still experience.

About once every nine (9) months (*give-or-take*) HP invites storage analysts to either its Andover, MA, or Fremont, CA, offices to have a series of in-depth discussion about its portfolio of products in its Storage division. During these 2-day events, the product managers from the various groups (3PAR StoreServ, [StoreOnce Backup](#), StoreAll Archive, StoreVirtual,



etc.) are given time to present to the analysts in attendance. It is during these times that candid and frank discussions ensue where each HP product is examined in-depth with the HP product managers providing context as to why they made the product design decisions that they have.

One of the more enlightening pieces of information to come out of these sessions was the amount of data that HP has collected from organizations into which its StoreOnce appliances are being considered for deployment. To date, HP has assessed environments with more than **half an exabyte** of backup data with the vast majority of backup data analyzed comprised of file system backups, either performed directly or thru NDMP.

This amount of data gives HP a rather unique perspective on backup successes and failures. For instance, HP shared that of the approximately 4.5 million backup jobs for which it has collected data, 94.7% of them have completed successfully.

HP also revealed that organizations in particular struggle with long-running backups. Over 50% of the assessed environments had backup windows of 24 hours or more. Of these, 30% of the organizations that it had assessed had at least one backup that ran in excess of 192 or more hours – or 8 days or more. Further, the data indicates a correlation between file system backups and long backup windows.

Granted, these statistics from HP are by no means “official” and subject to some interpretation. However they possibly provide some of the first, large scale empirical evidence that for the vast majority of organizations that data growth goes hand-in-hand with elongated backup windows and is a major contributor if not the primary source of why backups still fail today.

Organizations moving to StoreOnce appliances, which provide high levels of performance in conjunction with source-side deduplication, are addressing this common organizational pain

point as they both shorten backup windows and increase the probability that backups complete successfully. Further, using HP's StoreOnce [Recovery Manager Central](#) solution, organizations may perform virtual machine and file system backups based on block level changes as backup data flows from HP 3PAR StoreServ to StoreOnce. This combination of solutions provides the keys that organizations need to solve backup in their environments as many organizations using the HP StoreOnce deduplicating backup appliances have already discovered.

---

## **Advanced Encryption and VTL Features Give Organizations New Impetus to Use the Dell DR Series as their “One Stop Shop” Backup Target**

To simplify their backup environments, organizations desire backup solutions that essentially function as “one-stop shops” to satisfy their multiple backup requirements. To succeed in this role, they should provide needed software, offer NAS and virtual tape library (VTL) interfaces, scale to high capacities and deliver advanced encryption capabilities to secure backup data. By Dell introducing advanced encryption and VTL options into its latest DR 3.2 OS software release for its [DR Series](#), it delivers this “one-stop shop” experience that organizations want to implement in their backup infrastructure.

## **The More Backup Changes, the More It Stays the Same**

Deduplicating backup appliances have replaced tape as a backup target in many organizations. By accelerating backups and restores, increasing backup success rates and making disk-based backup economical, these appliances have fundamentally transformed backup.

Yet their introduction does not always change the underlying backup processes. Backup jobs may still occur daily; are configured as differential, incremental or full; and, are managed centrally. The only real change is using disk in lieu of tape as a target.

Even once in place, many organizations still move backup data to tape for long term data retention and/or offsite disaster recovery. Further, organizations in the finance, government and healthcare sectors typically encrypt data such as SEC [Rule 17a-4](#) specifies or the 2003 [HIPAA Security Rules](#) and more recent [2009 HITECH Act](#) strongly encourage.

## **Continued Relevance of Encryption and VTLs in Enterprises**

This continued widespread use of tape as a final resting place for backup data leads organizations to keep current backup processes in place. While they want to use deduplicating backup appliances, they simply want to swap out existing tape libraries for these solutions. This has given rise to the need for deduplicating backup appliances to emulate physical tape libraries as virtual tape libraries (VTLs).

A VTL requires minimal to no changes to existing backup-to-tape processes nor does it require many changes to how the backup data is managed after backup. The backup software now backs up data to the VTL's virtual tape drives where the data is stored on virtual tape cartridges. Storing data this way facilitates its movement from virtual to real or physical tape cartridges and enables the backup software to track its location regardless of where it resides.

VTLs also accelerate backups. They give organizations more flexibility to keep data on existing SANs which negates the need to send data over corporate LANs where it has to contend with other network traffic. SAN protocols also better support the movement of larger block sizes of data which are used during backup.

Finally, VTLs free backup from the constraints of physical tape libraries. Creating new tape drives and tape cartridges on a VTL may be done with the click of a button. In this way organizations may quickly create multiple new backup targets to facilitate scheduling multiple, concurrent backup jobs.

Encrypting backup data is also of greater concern to organizations as data breaches occur both inside and outside of corporate firewalls. This behooves organizations to encrypt backup data in the most secure manner regardless if the data resides on disk or tape.

### **Advanced Encryption and VTL Functionality Central to Dell DR Series 3.2 OS Release**

Advanced encryption capabilities and VTL functionality are two new features central to Dell's 3.2 operating system (OS) release for its [DR Series](#) of deduplicating backup appliances. The 3.2 OS release provides organizations a key advantage over competitive solutions as Dell makes all of its software features available without requiring additional licensing fees. This applies to both new DR Series appliances as well as existing Dell DR Series appliances which may be upgraded to this release to gain full access to these features at no extra cost.

The 3.2 OS release's advanced encryption capabilities use the FIPS 140-2 compliant 256-bit Advanced Encryption Standard (AES) standard to encrypt data. By encrypting data that conforms to this standard ensures that it is acceptable to federal agencies in both Canada and the United States. This

also means that organizations who are in these countries and need to comply with their regulations are typically, by extension, in compliance when they use the DR Series to encrypt their backup data.

The 3.2 OS release implements this advanced encryption capability by encrypting data after its inline deduplication of the backup data is complete. In this way, each DR Series appliance running the 3.2 OS release deduplicates backup data as it is ingested to achieve the highest possible deduplication ratio as encrypting data prior to deduplication negatively impacts deduplication's effectiveness. Encrypting the data after it is deduplicated also reduces the amount of overhead associated with encryption since there is less data to encrypt while keeping the overhead associated with the encryption on the DR Series appliance. In cases where existing DR4100s are upgraded to the 3.2 OS release, encryption may be done post-process on those data volumes that have previously been stored unencrypted in the DR4100's storage repository.

The VTL functionality that is part of the 3.2 OS release includes options to present a VTL interface on either corporate LANs or SANs. If connected to a corporate LAN, the NDMP protocol is used to send data to the DR Series while, if it is connected to a corporate SAN, the iSCSI protocol is used.

Every DR Series appliance running the 3.2 OS release may be configured to present up to four (4) containers that each operate as separate VTLs. Each of these individual VTL containers may emulate one (1) StorageTek STK L700 tape library or an OEM version of the STK L700; up to ten (10) IBM ULT3580-TD4 tape drives; and, up to 10,000 tape cartridges that may each range in size from 10GB to 800GB.

As each individual VTL container on the DR Series appears as an STK L700 library to backup software, the backup software manages the VTL in the same way it does a physical tape

library: it copies the data residing on virtual tape cartridges to physical tape cartridges and back again, if necessary. With this functionality available on leading enterprise backup software products such as Dell NetVault, CommVault Simpana, EMC Networker, IBM TSM, Microsoft Data Protection Manager (iSCSI only), Symantec Backup Exec and Symantec NetBackup, each of these can recognize and manage the Dell DR Series VTL as a physical STK L700 tape library, carry forward existing tape copy processes, implement new ones if required, and manage where copies of tape cartridges—physical or virtual—reside.

### **Dell's 3.2 OS Release Gives Organizations New Impetus to Make Dell DR Series Their "One Stop Shop" Backup Target**

All size organizations want to consolidate and simplify their backup environments and using a common deduplicating backup appliance platform is one excellent way to do so. Dell's 3.2 OS release for its DR Series gives organizations new impetus to start down that path. The introduction of advanced encryption and VTL features along with the introduction of 6TB HDDs on expansion shelves for the DR6000 and the availability of Rapid NFS/Rapid CIFS protocol accelerators for the DR4100 provide the additional motivation that organizations need to non-disruptively introduce and use the DR Series in this broader role to improve their backup environments even as they keep existing backup processes in place.

---

## **HP 3PAR StoreServ Management Console Answers the Call for**

# Centralized, Simplified Storage Operations Management without Technical Compromise

Scalable. Reliable. Robust. Well performing. Tightly integrated with hypervisors such as Microsoft Windows and VMware ESXi. These attributes are what every enterprise expects production storage arrays to possess and deliver. But as enterprises grow their infrastructure, they need to manage more storage arrays with the same or fewer number of IT staff. This requirement moves storage array manageability center stage which plays directly into the strengths of HP 3PAR StoreServ storage arrays and HP 3PAR StoreServ Management Console (SSMC).

## HP 3PAR's Legacy of Autonomic Storage Management

Since their inception [HP 3PAR StoreServ](#) systems have always delivered a robust, sophisticated set of features that are easy to implementation as a result of its autonomic storage management. The beauty of an HP 3PAR implementation is that its features do **NOT** require IT staff members to spend numerous hours learning and mastering each one to master them. Rather enterprises may reap the benefits of these features as they are seamlessly managed as part of an HP 3PAR StoreServ deployment.

This autonomic approach to storage management grants enterprises access to features such as:

- Adaptive Optimization
- Autonomic Groups
- Consolidated management of block and file
- Dynamic Optimization
- Priority Optimization
- Rapid Provisioning

These and other features have led enterprises to deploy multiple HP 3PAR StoreServ systems to address their numerous challenges. But as enterprises deploy more HP 3PAR systems, a new, separate challenge emerges: centrally managing these multiple HP 3PAR StoreServ systems.

## **HP SSMC Answers the Call for Centralized Storage Management**

All of the capacity and performance management features used to manage a single HP 3PAR StoreServ array are now available through the HP StoreServ Management Console ([SSMC](#)) which centralizes and consolidates the management of up to sixteen (16) HP 3PAR StoreServ systems. Further, HP plans to extend the SSMC's capabilities to manage even more HP 3PAR StoreServ systems.

The SSMC creates a common storage management experience for any HP 3PAR StoreServ system. Whether it is a high end HP 3PAR StoreServ 10000, the all-flash HP 3PAR StoreServ 7450 or a member of the midrange HP 3PAR StoreServ 7000 family, all of these systems may be managed through the HP 3PAR SSMC.

### **Top Level and System Views**

The HP 3PAR SSMC provides both top level and system views. The top level view displays the health of each managed HP 3PAR array. Administrators may view real time capacity and performance metrics as well as historical data for both of these items to monitor and identify longer term trends. Administrators also have the flexibility to put individual arrays into groups so they may collectively visualize and manage each array group's capacity and performance by application, department or company.

In the system view, administrators may select individual HP 3PAR StoreServ systems and view information specific to it. For instance, they may view: the available capacity of each storage tier type to include block and file storage management; the features licensed on that system; and, the



system's resource utilization. By understanding how many or how few resources are available on each system, administrators may better determine where to place new applications and their data to align each application's needs with the StoreServ's available resources and features.

Centralizing management of all HP 3PAR StoreServ arrays under the SSMC also makes it easier to move an application and its data from one array to another. As the anticipated capacity and performance characteristics of a new application rarely align with how it actually performs in production, the SSMC helps administrators first understand how the application uses resources on the array and then, if a change in array is needed, helps them identify another array where the application might be better placed to give it access to needed storage capacity or improve its performance.

### **End-to-End Mapping**

Degraded application performance, hardware failures, system upgrades and storage system firmware patches are realities with which every modern data center contends and must manage in order to ensure continuous application availability and deliver optimal application performance. Yet delivering on these objectives in today's highly virtualized infrastructures without a view into the end-to-end mapping may become almost impossible to achieve.

Doing so requires visibility into: how file shares and/or virtual volumes map to a storage array's underlying disk drives; on which storage array ports they are presented; and, which applications access these file shares and/or virtual volumes. Only by having this visibility into how virtualized objects use the underlying physical infrastructure can they verify that each application is appropriately configured for continuous availability or begin to understand how a failed component in the infrastructure might impact the performance of a specific application.

The HP 3PAR SSMC provides this end-to-end mapping of the underlying infrastructure that is critical to maintaining application availability and ensuring optimal application performance. By identifying and visualizing the exact physical components used by each physical or virtual machine, enterprises can better understand the impact of system component upgrades or outages as well as identify, isolate and troubleshoot performance issues before they have influence an application.

### **Capacity and Performance Reporting**

The System Reporter component of SSMC automatically and in the background collects data on a number of different object data points on all managed HP 3PAR StoreServ systems without needing any additional setup. Using this collected data, the System Reporter can generate hundreds of different, customizable reports that contain detailed capacity and performance information on any of these managed systems.

The System Reporter contains predefined reports, settings, templates and values that further help enterprises accelerate their SSMC deployment. They frees them to quickly gather data information about their environment and then analyze it using its analytical engine that helps enterprises interpret collected performance data. Once analyzed, they may configure any of the default settings to meet their specific needs.

### **Simplified Ongoing Management**

The frequency and quality of management of storage for client-attached systems can vary as widely as the types of applications hosted on the client-attached systems. In some cases, administrators may only need to administer the storage array on a quarterly or annual basis. While this simplifies storage management, in large environments infrequent array administration has some unintended consequences such as simply remembering a client server's name or which applications or

data reside on a specific array.

The SSMC resolves these issues. Using its search functionality, administrators may search for specific clients that are attached to HP 3PAR StoreServ arrays and can quickly identify the storage array(s) in the environment that the clients are accessing.

### **HP 3PAR SSMC Answers Call for Centralized Storage Operations Management without Technical Compromise**

HP 3PAR StoreServ systems host critical application data that are the heart and soul of many enterprise data centers as they are optimized for hosting mixed physical and virtual machine workloads. But as more enterprises implement greater numbers of HP 3PAR systems they need a better way to manage them.

The HP 3PAR SSMC answers this call for a centralized storage operations management console as it ensures all HP 3PAR systems under management remain simple to manage even as organizations add more of them. The SSMC globally manages multiple HP 3PAR StoreServ systems from a single console while preserving the automation and simplicity associated with managing a single HP 3PAR StoreServ. This serves as testament to HP's commitment to delivering technology that accelerates business and technical operations while remaining easy to implement, use and manage.

---

## **The Performance of a \$500K Hybrid Storage Array Goes**

# Toe-to-Toe with Million Dollar All-Flash and High End Storage Arrays

On March 17, 2015, the Storage Performance Council (SPC) [updated](#) its “Top Ten” list of SPC-2 results that includes performance metrics going back almost three (3) years to May 2012. Noteworthy in these updated results is that the three storage arrays ranked at the top are, in order, a high end mainframe-centric, monolithic storage array (the [HP XP7](#), OEMed from Hitachi), an all-flash storage array (from startup Kaminario, the [K2 box](#)) and a hybrid storage array (Oracle ZFS Storage [ZS4-4 Appliance](#)). Making these performance results particularly interesting is that the hybrid storage array, the Oracle ZFS Storage ZS4-4 Appliance, can essentially go toe-to-toe from a performance perspective with both the million dollar HP XP7 and Kaminario K2 arrays and do so at approximately half of their cost.

Right now there is a great deal of debate in the storage industry about which of these three types of arrays – all-flash, high end or hybrid – can provide the highest levels of performance. In recent years, all-flash and high end storage arrays have generally gone neck-and-neck though all-flash arrays are generally now seen as taking the lead and pulling away.

However, when price becomes a factor (*and when isn't price a factor?*) such that enterprises have to look at price and performance, suddenly hybrid storage arrays surface as very attractive alternatives for many enterprises. Granted, hybrid storage arrays may not provide all of the performance of either all-flash or high end arrays, but they can certainly deliver superior performance at a much lower cost.

This is what makes the recently updated Top Ten results on the SPC website so interesting. While the breadth of arrays covered in the published SPC results by no means cover every storage array on the market, they do provide enterprises with some valuable insight into:

- How well hybrid storage arrays can potentially perform
- How comparable their storage capacity is to high-end and all-flash arrays
- How much more economical hybrid storage arrays are

In looking at these three arrays that currently sit atop the SPC-2 [Top Ten](#) list and how they were configured for this test, they were comparable in one of the ways that enterprises examine when making a buying decision. For instance, all three had comparable amounts of raw capacity.

### Raw Capacity

High-End XP7	230TB	HP
All-Flash K2	179TB	Kaminario
Hybrid Oracle ZFS Storage ZS4-4 Appliance	175TB	

Despite using comparable amounts of raw capacity for testing purposes, they got to these raw capacity totals using decidedly different media. The high end, mainframe-centric HP XP7 used 768 300GB 15K SAS HDDs to get to its 230TB total while the all-flash Kaminario K2 used 224 solid state drives (SSDs) to get to its 179TB total. The Oracle ZS4-4 stood out from these other two storage arrays in two ways. First, it used 576 300GB **10K** SAS HDDs. Second, its storage media costs were a fraction of the other two. Comparing strictly list prices, its media costs were only about 16% of the cost of the

HP XP7 and 27% of the cost of the Kaminario K2.

These arrays also differed in terms of how many and what types of storage networking ports they each used. Both the HP XP7 and the Kaminario K2 used a total of 64 and 56 8Gb FC ports respectively for connectivity between the servers and their storage arrays. The Oracle ZS4-4 only needed 16 ports for connectivity though it used Infiniband for server-storage connectivity as opposed to 8Gb FC. The HP XP7 and Oracle ZS4-4 also used cache (512GB and ~3TB respectively) while the Kaminario K2 used no cache at all. It instead used a total of 224 solid state drives (SSDs) packaged in 28 flash nodes (8-800GB SSDs in each flash node.)

This is not meant to disparage the configuration or architecture of any of these three different storage arrays as each one uses proven technologies in the design of their arrays. Yet what is notable is the end results when these three arrays in these configurations are subjected to the same SPC2 performance benchmarking tests.

While the HP XP7 and Kaminario K2 came out on top from an overall performance perspective, it is interesting to note how well the Oracle ZS4-4 performs and what its price/performance ratio is when compared to the high end HP XP7 and the all-flash Kaminario K2. It provides 75% to over 90% of the performance of these other arrays at a cost per MB that is up to 46% less.

SPC-2 Top 10 Performance Results					
Vendor	Date	GB/s	\$/MB/s	TBs	Price
HP XP7	2014	43.01	\$28.30	192.76	\$1,217,462
Kaminario K2	2013	33.47	\$29.79	60.12	\$997,348
<b>Oracle ZS4-4</b>	<b>2015</b>	<b>31.48</b>	<b>\$17.09</b>	<b>30.96</b>	<b>\$538,050</b>
Oracle ZS3-4	2013	17.24	\$22.53	31.61	\$388,472
Oracle ZS3-2	2014	16.20	\$12.08	24.20	\$195,916
Fujitsu Eternus DX8700	2013	16.04	\$79.40	71.40	\$1,275,163
IBM DS8870	2012	15.42	\$131.21	30.92	\$2,023,742
IBM Storwize V7000	2012	14.58	\$129.14	74.49	\$1,883,036
NEC M700	2013	14.41	\$25.10	53.55	\$361,612
Hitachi VSP	2012	13.15	\$95.38	129.11	\$1,254,093
HP P9500 XP	2012	13.15	\$88.34	129.11	\$1,161,504
Hitachi HUS VM	2014	11.27	\$32.64	193.67	\$368,065

Source: "Top Ten"

SPC-2 Results,

[https://www.storageperformance.org/results/benchmark\\_results\\_spc2\\_top-ten](https://www.storageperformance.org/results/benchmark_results_spc2_top-ten)

It is easy for enterprises to become enamored with all-flash arrays or remain transfixed on high-end arrays because of their proven and perceived performance characteristics and benefits. But these recent SPC-2 performance benchmarks illustrate that hybrid storage arrays such as the Oracle ZFS Storage ZS4-4 Appliance can deliver levels of performance that are comparable to million-dollar all-flash and high-end arrays at half of their cost which are numbers that any enterprise can take to the bank.

---

## Give the I/Os of Your Mission

# and Business Critical Apps the VIP Treatment They Deserve

Features such as automated storage tiering and storage domains on today's enterprise storage arrays go a long way toward making it feasible for organizations to successfully host multiple applications with different performance and priority requirements on a single array. However prioritizing the order in which data and I/Os are tiered is an entirely differently matter as organizations typically want the data and I/Os associated with their mission and business critical I/Os serviced ahead of lower priority applications. This is where the Quality of Service (QoS) Plus feature found on the Oracle FS1 comes into play as it does more than provide the "brains" behind its auto-tiering feature. It also re-prioritizes and re-orders application I/O according to each application's business value to the enterprise.

## **Historical Treatment of Application I/O by Storage Arrays**

Storage arrays by default treat incoming read and write I/Os from all applications the same (*Figures 1 and 2 below.*) Whether I/Os are issued by an Oracle Database or an application retrieving archival data, storage arrays ingest these I/Os in the order in which they arrive and then process and reply to them in the same order.





**Figure 1**

**Conventional Storage Arrays' "First-In-First-Out" I/O Input**

*Source: Oracle*

The issue that this "First-In-First-Out" process potentially creates in consolidated environments is that I/O from an application doing archiving is handled in the same manner as I/O from an OLTP application trying to access an Oracle Database. This puts I/O processing out of alignment with business priorities.



## Figure 2

### Conventional Storage Arrays' "First-In-First-Out" I/O Output

*Source: Oracle*

#### Oracle FS1 Realigns Application I/O with Business Priorities

The Oracle FS1's QoS Plus addresses this misalignment between available technical resources and business priorities with Application Profiles that have Archive, Low, Medium, High and Premium priorities. Enterprises may use pre-tuned and tested Application Profiles for Oracle Database and key Oracle and third-party applications that come with the [Oracle FS1](#) to expedite storage provisioning with just one click or create their own. These profiles are associated with each application accessing the FS1 and serve two purposes:

1. Data associated with each application is placed on the tier or tiers of storage associated with its FS1 application profile.
2. Read and write I/Os from each application are then ingested and prioritized according to its application profile to create a "Priority-In-Priority-Out" means of handling I/O.

Using Priority-In-Priority-Out, the FS1 services I/Os associated with the highest priority applications first and then services I/Os from other, lower priority apps according to how they are categorized on the FS1. As Figure 3 illustrates, I/Os still come into the FS1 the same way as before – the order in which they were sent.



Figure 3

FS1's "Priority-In-Priority-Out" I/O Input

Source: Oracle

The difference is that the FS1 recognizes each application's respective I/O, correlates it to its profile and then responds to each application I/O based upon how it is prioritized as Figure 4 below illustrates.



Figure 4

FS1's "Priority-In-Priority-Out" I/O Output

*Source: Oracle*

The Oracle FS1's "Priority-In-Priority-Out" component of its QoS Plus feature gives enterprises more flexibility to mix and match applications of different priorities on the same FS1 array or within FS1 Storage Domains without worrying about how I/O from low priority applications will affect their Premium or High Priority applications. QoS Plus ensures that I/Os from Premium and High priority applications are serviced before I/Os from lower priority applications. However it also ensures that I/Os from lower priority applications are queued up so that their I/Os are serviced in a timely manner to meet their respective service level agreements (SLAs).

---

## **A Glimpse into the Next Decade of Backup and Recovery; Interview with Dell Software's General Manager, Data Protection, Brett Roscoe, Part IX**

Today backup and recovery looks almost nothing like it did 10 years ago. But as one looks at all of the changes still going on in backup and recovery, one can only guess what backup and recovery might look like in another 5-10 years. In this ninth and final installment of my interview series with Brett Roscoe, General Manager, Data Protection for Dell Software, Brett provides some insight into where he sees backup and recovery going over the next decade.

**Jerome:** *There is a lot excitement out there right now around data protection and how much backup and recovery has changed in the last 5 – 10 years. To a certain degree, it does not even look like it did 10 years ago. It makes me wonder what it is going to look like in 5 or 10 more years in terms of what new technologies are going to come to market or how they are going take advantage of new technologies. Do you have any thoughts about what the future of backup and recovery looks like and is it even going to be called backup and recovery?*

**Brett:** That's a great question, and boy, if I could accurately predict what's going to happen ten years from now, that would be something!. But you're absolutely right in saying that this is a market that's quickly evolving and changing. That is one nice thing about the software side of the business: you can quickly change with the market and meet these changing customer needs.

But if I had to predict what I think is going to happen, it is clear to me that we are going to continue to move to real-time or near real-time data protection. That world of 10 years ago, where you scheduled backup jobs at night, is just going to fade away. The real-time backup means that as soon as something changes in your environment, it's immediately protected. I do not think we are going to get away from that. I think it's going to be driven by the technology, as well as the demands of our customers.

Data is becoming more and more critical. More and more of my company depends on this data being available and being recoverable. If a business has a disaster and loses their data, a large percentage of them never recover and never stay in business. That is just becoming a reality.

In addition, I think you're going to more and more cloud and backup and recovery as a service models, especially in the smaller SMB side of the market. Customers are looking to offload and maybe simplify their IT infrastructure. When you

can start using very efficient technologies such as WAN optimization, deduplication and compression to minimize the bandwidth required to move data into the cloud, that reduces costs while making bandwidth more efficient. This makes the cloud much more usable as a backup and recovery option.

Further, virtual standby technology is coming into its own. Using this technology, you can actually run your applications in the cloud for some period of time. Using this approach, you may lease some additional compute and storage resources in the cloud to deliver on this capability. However, it is a temporal thing and allows you to meet an SLA which would normally require much more cost if you did it locally. So, in that vein, I expect to see expanded use of the cloud and a continuation of today's hybrid cloud environment IT initiatives.

Another trend is using backup for additional IT functions like analytics, testing, and data migration. We have all this great data that we've captured through the backup process. Now there is a new push to create more value for this idle data and exploring what else we can do with it.

There are all kinds of great tools coming out in terms of analytics and data mining capabilities that can provide additional value to any company that can mine through that data and use it in some other manner.

I also expect to see tighter integration of backup to traditional management infrastructure. This idea that backup becomes a little bit more of the mainstream tool set that you see in IT through your larger IT frameworks or management infrastructures is going to continue. Many big companies are starting to build backup tools into their applications or into their hypervisors. We will continue to see that.

Lastly, there's the trend we hit on previously when we talked about the Dell Backup & Disaster Recovery Suite, which is the

desire for more flexibility to leverage Dell or whichever vendor's IP across the portfolio. That will continue to grow. How do you get customers to a place where they feel like they can leverage more and more of your IP no matter what product they buy? Dell is leading the charge on that. We really have a vision for making sure all of our IP is available to our customers.

*In [Part I](#) of this interview series, Brett and I discussed the biggest backup and recovery challenges that organizations face today.*

*In [Part II](#) of this interview series, Brett and I discussed the imperative to move ahead with next gen backup and recovery tools.*

*In [Part III](#) of this interview series, Brett and I discussed four (4) best practices that companies should be implementing now to align the new capabilities in next gen backup and recovery tools with internal business processes.*

*In [Part IV](#) of this interview series, Brett and I discussed the main technologies in which customers are currently expressing the most interest.*

*In [Part V](#) of this interview series, Brett and I examine whether or not one backup software product can "do it all" from a backup and recovery perspective.*

*In [Part VI](#) of this interview series, Brett and I discuss Dell's growing role as a software provider.*

*In [Part VII](#) of this interview series, Brett provides an in-depth look at Dell's new Backup and Disaster Recovery Suite.*

*In [Part VIII](#) of this interview series, Brett explains how Dell now provides a single backup solution for multiple backup and recovery challenges.*