

Nanoseconds, Stubborn SAS, and Other Takeaways from the Flash Memory Summit 2019

Every year at the Flash Memory Summit held in Santa Clara, CA, attendees get a firsthand look at the technologies that will impact the next generation of storage. This year many of the innovations centered on forthcoming interconnects that will better deliver on the performance that flash offers today. Here are DCIG's main takeaways from this year's event.

Takeaway #1 – Nanosecond Response Times Demonstrated



PCI Express (PCIe) fabrics can deliver nanosecond response times using resources (CPU, memory, storage) situated on different physical enclosures. In meeting with PCIe provider, [Dolphin Interconnect Solutions](#), it demonstrated how an application could access resources (CPU, flash storage & memory) on different devices across a PCIe fabric in nanoseconds. Separately, [GigaIO announced](#) 500 nanosecond end-to-end latency using its [PCIe FabreX](#) switches. While everyone else at the show was boasting about microsecond response times, Dolphin and GigaIO introduced nanoseconds into the conversation. Both these companies ship their solutions now.

Takeaway #2 – Impact of NVMe/TCP Standard Confirmed

Ever since we heard the industry planned to port NVMe-oF to TCP, DCIG [thought](#) this would accelerate the overall adoption of NVMe-oF. Toshiba confirmed our suspicions. In discussing

its [Kumoscale](#) product with DCIG, it shared that it has seen a 10x jump in sales since the industry ratified the NVMe/TCP standard. This stems from all the reasons DCIG stated in a previous [blog entry](#) such as TCP being well understood, Ethernet being widely deployed, its low cost, and its use of existing infrastructure in organizations.

Takeaway #3 – Fibre Channel Market Healthy, Driven by Enterprise All-flash Array

According to FCIA leaders, the Fibre Channel (FC) market is healthy. FC vendors are selling 8 million ports per year. The enterprise all-flash array market is driving FC infrastructure sales, and 32 Gb FC is shipping in volume. Indeed, DCIG's research revealed 37 all-flash arrays that support 32 Gb FC connectivity.

Front-end connectivity is often the bottleneck in all-flash array performance, so doubling the speed of those connections can double the performance of the array. Beyond 32 Gb FC, the FCIA has already ratified the 64 Gb standard and is working on the 128 Gb FC. Consequently, FC has a long future in enterprise data centers.

FC-NVMe brings the benefits of NVMe-oF to Fibre Channel networks. FC-NVMe reduces protocol overhead, enabling GEN 5 (16 Gb FC) infrastructure to accomplish the same amount of work while consuming about half the CPU of standard FC.

Takeaway #4 – PCIe Will Not be Denied

All resources (CPU, memory and flash storage) can connect with one another and communicate over PCIe. Further, using PCIe eliminates the need for introducing the overhead associated with storage protocols (FC, InfiniBand, iSCSI, SCSI). All these resources talk the PCIe protocol. With the PCIe 5.0 standard formally ratified in May 2019 and discussions about PCIe 6.0 occurring, the future seems bright for the growing

adoption of this protocol. Further, AMD and Intel having both thrown their support behind it.

Takeaway #5 – SAS Will Stubbornly Hang On

DCIG's research finds that over 75% of AFAs support 12Gb/second SAS now. This predominance makes the introduction of 24G a logical next step for these arrays. A proven, mature, and economical interconnect, few applications can yet drive the performance limits of 12Gb, much less the forthcoming 24G standard. Adding to the likelihood that 24G moves forward, the SCSI Trade Association (STA) reported that the recent 24G plug fest went well.

Editor's Note: This blog entry was updated on August 9, 2019, to correct grammatical mistakes and add some links.

iXsystems FreeNAS Mini XL+ and Mini E Expand the Reach of Open Source Storage to Small Offices

On July 25, 2019, iXsystems® [announced](#) two new storage systems. The **FreeNAS® Mini XL+** provides a new top-end model in the FreeNAS Mini product line, and the **FreeNAS Mini E** provides a new entry-level model. These servers are mini-sized yet provide professional-grade network-attached storage.

FreeNAS Minis are Professional-grade and

Whisper Quiet



Source: iXsystems

The FreeNAS Mini XL+ and Mini E incorporate technologies normally associated with enterprise servers, such as ECC memory, out-of-band management, and NAS-grade hard drives. Both are engineered for and powered by the widely adopted ZFS-based FreeNAS Open Source storage OS. Thus, the Mini XL+ and Mini E provide file, block, and S3 object storage to meet nearly any SOHO/SMB storage requirement.

Early in my IT career, I purchased a tower server that was marketed to small businesses as a convenient under-desk solution. The noise and heat generated by this server quickly helped me understand why so many small business servers were running in closets. The FreeNAS Mini is not this kind of server.

All FreeNAS Mini models are designed to share space with people. They are compact and “whisper quiet” for use in offices and homes. They are also power-efficient, drawing a maximum of 56 to 106 Watts for the Mini E and Mini XL+, respectively.

Next-Generation Technology Powers Up the FreeNAS Mini XL+ and Mini E

The Mini XL+ and E bring multiple technology upgrades to the FreeNAS Mini platform. These include:

- Intel Atom C3000 Series CPUs
- DDR4 ECC DRAM
- +1 2.5" Hot-swappable Bay (Mini XL+)
- PCI Express 3.0 (Mini XL+)
- IPMI iKVM (HTML5-based)
- USB 3.0
- Standard Dual 10 Gb Ethernet Ports (Mini XL+)
- Quad 1 Gb Ethernet Ports (Mini E)

FreeNAS Mini is a Multifunction Solution

FreeNAS Mini products are well-equipped to compete against other small form factor NAS appliances; and perhaps even tower servers because of their ability to run network applications directly on the storage appliance.

Indeed, the combination of more powerful hardware, application plug-ins, and the ability to run hypervisor or containerized applications directly on the storage appliance makes the FreeNAS Mini a multi-function SOHO/ROBO solution.

[FreeNAS plugins](#) are based on pre-configured FreeBSD containers called jails that are simple to install. iXsystems refers to these plugins as "Network Application Services". The plugins are available across all TrueNAS® and FreeNAS products, including the new FreeNAS Mini E and XL+.

The available plugins include quality commercial and open source applications covering a range of use cases, including:

- Backup (Asigra)
- Collaboration (NextCloud)

- DevOps (GitLab)
- Entertainment (Plex)
- Hybrid cloud media management (Iconik)
- Security (ClamAV)
- Surveillance video (ZoneMinder)

FreeNAS Mini Addresses Many Use Cases

The FreeNAS Mini XL+ and Mini E expand the range of use cases for the FreeNAS product line.

Remote, branch or home office. The FreeNAS Mini creates value for any business that needs professional-grade storage. It will be especially appealing to organizations that need to provide reliable storage across multiple locations. The Mini's combination of a dedicated management port, IPMI, and [TrueCommand](#) management software enables comprehensive remote monitoring and management of multiple Minis.

FreeNAS Mini support for S3 object storage includes bidirectional file sync with popular cloud storage services and private S3 storage. This enables low-latency local file access with off-site data protection for home and branch offices.

Organizations can also deploy and manage FreeNAS systems at the edge and use TrueNAS systems where enterprise-class support and HA are required. Indeed, iXsystems has many clients that deploy both TrueNAS and FreeNAS. In doing so, they gain the benefit of a single storage operating environment across all their locations, all of which can be managed centrally via TrueCommand.

Managed Service Provider. TrueCommand and IPMI also enable managed service providers (MSPs) to cost-effectively manage a whole fleet of FreeNAS or TrueNAS systems across their entire client base. TrueCommand enables role-based access controls, allowing MSPs to assign systems into teams broken down by

separate clients and admins..

Bulk data transfer. FreeNAS provides robust replication options, but sometimes the fastest way to move large amounts of data is to physically ship it from site to site. Customers can use the Mini XL+ to rapidly ingest, store, and transfer over 70 TB of data.

Convenient Purchase of Preconfigured or Custom Configurations

iXsystems has increased the appeal of the FreeNAS Mini by offering multiple self-service purchasing options. It offers a straightforward [online ordering tool](#) that allows the purchaser to configure and purchase any of the FreeNAS Mini products directly from iXsystems. iXsystems also makes preconfigured systems available for rapid ordering and delivery via Amazon Prime. Either method enables purchase with a minimal amount of fuss and a maximum amount of confidence.

Thoughtfully Committed to Expanding the Reach of Open Source Storage

Individuals and businesses that purchase the new FreeNAS Mini XL+ or Mini E are doing more than simply acquiring high-quality storage systems for themselves. They are also supporting the ongoing development of Open Source projects such as FreeBSD and OpenZFS.

iXsystems has decades of expertise in system design and development of Open Source software including FreeNAS, FreeBSD, OpenZFS, and TrueOS®. Its recent advances in GUI-based management for simplified operations are making sophisticated Open Source technology more comfortable to non-technical users.

iXsystems has thoughtfully engineered the FreeNAS Mini E and XL+ for FreeNAS, the world's most widely deployed Open Source

storage software. In doing so, they have created high-quality storage systems that offer much more than just NAS storage. Quietly. Affordably.

For a thorough hands-on technical review of the FreeNAS Mini XL+, see this [article](#) on ServetheHome.

Additional product information, including detailed specifications and documentation, is available on the iXsystems [FreeNAS Mini product page](#).

DCIG Quick Look: iXsystems TrueNAS X10 Offers an Affordable Offramp from Public Cloud Storage

For many of us, commuting in rush hour with its traffic jams is an unpleasant fact of life. But I once had a job on the outer edge of a metropolitan area. I was westbound when most were eastbound. I often felt a little sorry for the mass of people stuck in traffic as I zoomed—with a smile on my face—in the opposite direction. Today there is a massive flow of workloads and their associated storage to the public cloud. But there are also a lot of companies moving workloads off the public cloud, and their reason is cloud economics.

Cloud Economics Are Not Always Economical

In a recent conversation with [iXsystems](#), it indicated that many of its new customers are coming to it in search of lower-than-public-cloud costs. Gary Archer, Director of Storage

Marketing at iXsystems met with DCIG earlier this month to brief us on a forthcoming product. It turns out the product was not the rumored hyperconverged infrastructure appliance. Instead, he told us iXsystems was about to reach a new low as in a new low starting price and cost per gigabyte for enterprise-grade storage.

A lot of companies look at iXsystems because they want to reduce costs by migrating workloads off the public cloud. These customers find the Z-Series enterprise-grade open source storage attractive, but asked for a lower entry price and lower cost per GB.

iXsystems TrueNAS X10 is Economical by Design

To meet this demand, iXsystems chose current enterprise-grade, but not the highest-end, hardware for its new TrueNAS X10. For example, each controller features a single 6-core Intel Broadwell Xeon CPU. In an era of ever-larger DRAM caches, each X10 controller has just 32GB of ECC DRAM. Dual one-gigabit Ethernet is built in. 10 GbE is optional. Storage capacity is provided exclusively by SAS-attached hard drives. Flash memory is used, but only as cache.

The TrueNAS X10 retains all the redundancy and reliability features of the Z-Series, but at a starting price of just \$5,500. A 20 TB system costs less than \$10,000, and a 120 TB system costs less than \$18,000 street. So, the X10 starts at \$0.50/GB and ranges down to \$0.15/GB. Expansion via disk shelves should drive the \$/GB even lower.

iXsystems positions the TrueNAS X10 as entry-level enterprise-grade unified storage. As such, the TrueNAS X10 will make a cost-effective storage target for backups, video surveillance and file sharing workloads; but not for workloads characterized by random writes. Although iXsystems lists in-line deduplication and compression on its spec sheet, the relatively limited DRAM cache and CPU performance mean you

should probably only implement deduplication with caution. By way of example, the default setting for deduplication is off.

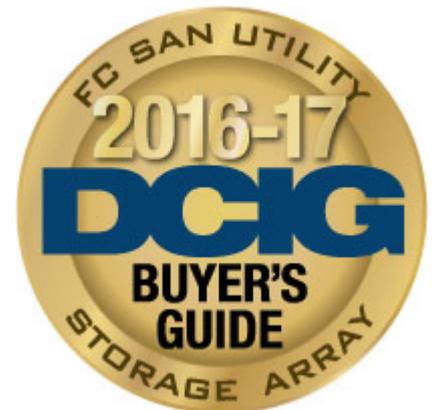
In the TrueNAS X10, iXsystems delivers enterprise-grade storage for companies that want to save money by moving off the public cloud. The X10 will also be attractive to companies that have outgrown the performance, capacity or limited data services offered by SMB-focused NAS boxes.

The TrueNAS X10 is not for every workload. But companies with monthly public cloud bills that have climbed into the tens of thousands may find that “cloud economics” are driving them to seek out affordable on premise alternatives. Seek and ye shall find.

DCIG 2016-17 FC SAN Utility Storage Array and Utility SAN Storage Array Buyer's Guides Now Available

DCIG is pleased to announce the availability of its **2016-17 FC SAN Utility Storage Array Buyer's Guide** and **2016-17 Utility SAN Storage Array Buyer's Guide** that each weight more than 100 features and rank 62 arrays from thirteen (13) different storage providers. These Buyer's Guide Editions are products of DCIG's updated research methodology where DCIG creates specific Buyer's Guide Editions based upon a larger, general body of research on a topic. As past Buyer's Guides have done, it continues to rank products as Recommended, Excellent, Good and Basic as well as offer the product information that organizations need to make informed buying decisions on FC SAN

Utility and multiprotocol Utility SAN storage arrays.



Over the years organizations have taken a number of steps to better manage the data that they already possess as well as prepare themselves for the growth they expect to experience in the future. These steps usually involve either deleting data that they have determined they do not need or should not keep while archiving the rest of it on a low cost media such as optical, tape or even with public cloud storage providers.

Fibre Channel (FC) and multiprotocol SAN storage arrays configured as utility storage arrays represent a maturation of the storage array market. Storage arrays using hard disk drives (HDDs) are still the predominant media used to host and service high performance applications. But with the advent of flash and solid state drives (SSDs), this reality is rapidly changing. Flash-based arrays are rapidly supplanting all-HDD storage arrays to host business-critical, performance sensitive applications as flash-based arrays can typically provide sub-two millisecond read and write response times.

However, the high levels of performance these flash-based arrays offer comes with a price – up to 10x more than all HDD-based utility storage arrays. This is where HDD-based arrays in general, and SAN utility storage arrays in particular, find a new home. These array may host and service applications with infrequently accessed or inactive data such as archived, backup and file data.

Many if not most organizations still adhere to a “*keep it all forever*” mentality when it comes to managing data for various reasons. These factors have led organizations to adopt a “*delete nothing*” approach to managing their data as this is often their most affordable and prudent option. The challenge with this technique is that as data volumes continue to grow and retention periods remain non-existent, organizations need to identify solutions on which they can affordably store all of this data.

Thanks to the continuing drop per GB in disk’s cost that day has essentially arrived. The emergence of highly available and reliable utility storage arrays that scale into the petabytes at a cost of well below \$1/GB opens the doors for organizations to confidently and cost-effectively keep almost any amount of data online and accessible for their business needs.

Utility storage arrays also offer low millisecond response times (8 – 10 ms) for application reads and writes. This is more than adequate performance for most archival or infrequently accessed data. These arrays deliver millisecond response times while supporting hundreds of terabytes if not petabytes of storage capacity at under a dollar per gigabyte.

The 2016-17 FC SAN Utility Storage Array Buyer’s Guide specifically covers those storage arrays that support the Fibre Channel storage networking protocol. The 2016-17 Utility SAN Storage Array Buyer’s Guide scores the arrays for their support for both FC and iSCSI storage networking protocols. All of the included utility storage arrays are available in highly available, reliable configurations and list for \$1/GB or less. While the arrays in this Guide may support other storage networking protocols, other specific protocols were not weighted in arriving in the conclusions in these Buyer’s Guide Editions.

DCIG’s succinct analysis provides insight into the state of

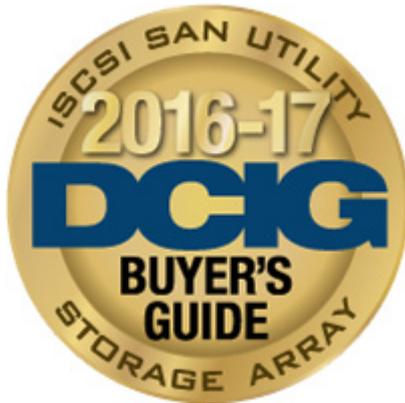
the SAN utility storage array marketplace. It identifies the significant benefits organizations can expect to realize by implementing a utility storage array, key features that organizations should evaluate on these arrays and includes brief observations about the distinctive features of each array. The storage array rankings provide organizations with an “at-a-glance” overview of this marketplace. DCIG complements these rankings with standardized, one-page data sheets that facilitate side-by-side product comparisons so organizations may quickly get to a short list of products that may meet their requirements.

Registration to access these Buyer’s Guides may be done via the DCIG Analysis Portal which includes access to DCIG Buyer’s Guides in PDF format as well as the DCIG Interactive Buyer’s Guide (IBG). Using the IBG, organizations may dynamically drill down and compare and contrast FC SAN and Utility SAN arrays by generating custom reports, including comprehensive strengths and weaknesses reports that evaluate a much broader base of features than what is found in the published Guide. Both the IBG and this Buyer’s Guide may be accessed after [registering](#) for the [DCIG Analysis Portal](#).

DCIG 2016-17 iSCSI SAN Utility Storage Array Buyer’s Guide Now Available

DCIG is pleased to announce the availability of its **2016-17 iSCSI SAN Utility Storage Array Buyer’s Guide** that weights more than 100 features and ranks 67 arrays from fourteen (14)

different storage providers. This Buyer's Guide Edition reflects the first use of DCIG's updated research methodology where DCIG creates specific Buyer's Guide Editions based upon a larger, general body of research on a topic. As past Buyer's Guides have done, it continues to rank products as Recommended, Excellent, Good and Basic as well as offer the product information that organizations need to make informed buying decisions on iSCSI SAN utility storage arrays.



Over the years organizations have taken a number of steps to better manage the data that they already possess as well as prepare themselves for the growth they expect to experience in the future. These steps usually involve either deleting data that they have determined they do not need or should not keep while archiving the rest of it on a low cost media such as optical, tape or even with public cloud storage providers.

iSCSI SAN storage arrays configured as utility storage arrays represent a maturation of the storage array market. Storage arrays using hard disk drives (HDDs) are still the predominant media used to host and service high performance applications. But with the advent of flash and solid state drives (SSDs), this reality is rapidly changing. Flash-based arrays are rapidly supplanting all-HDD storage arrays to host business-critical, performance sensitive applications as flash-based arrays can typically provide sub-two millisecond read and write response times.

However the high levels of performance these flash-based

arrays offer comes with a price – up to 10x more than all HDD-based utility storage arrays. This is where HDD-based arrays in general, and iSCSI SAN utility storage arrays in particular, find a new home. These array may host and service applications with infrequently accessed or inactive data such as archived, backup and file data.

Many if not most organizations still adhere to a “*keep it all forever*” mentality when it comes to managing data for various reasons. These factors have led organizations to adopt a “*delete nothing*” approach to managing their data as this is often their most affordable and prudent option. The challenge with this technique is that as data volumes continue to grow and retention periods remain non-existent, organizations need to identify solutions on which they can affordably store all of this data.

Thanks to the continuing drop per GB in disk’s cost that day has essentially arrived. The emergence of highly available and reliable iSCSI SAN utility storage arrays that scale into the petabytes at a cost of well below \$1/GB opens the doors for organizations to confidently and cost-effectively keep almost any amount of data online and accessible for their business needs.

iSCSI SAN utility storage arrays also offer low millisecond response times (8 – 10 ms) for application reads and writes. This is more than adequate performance for most archival or infrequently accessed data. These arrays deliver millisecond response times while supporting hundreds of terabytes if not petabytes of storage capacity at under a dollar per gigabyte.

This Buyer’s Guide edition specifically covers those storage arrays that support the iSCSI storage networking protocol, are available in highly available, reliable configurations and list for \$1/GB or less. While the arrays in this Guide may and often do support other storage networking protocols, those specific protocols were not weighted in arriving in the

conclusions in this Buyer's Guide Edition

DCIG's succinct analysis provides insight into the state of the iSCSI SAN utility storage array marketplace. It identifies the significant benefits organizations can expect to realize by implementing an iSCSI SAN utility storage array, key features that organizations should evaluate on these arrays and includes brief observations about the distinctive features of each array. The iSCSI SAN utility storage array ranking provide organizations with an "at-a-glance" overview of this marketplace. DCIG complements these rankings with standardized, one-page data sheets that facilitate side-by-side product comparisons so organizations may quickly get to a short list of products that may meet their requirements.

Registration to access this Buyer's Guide may be done via the DCIG Analysis Portal which includes access to DCIG Buyer's Guides in PDF format as well as the DCIG Interactive Buyer's Guide (IBG). Using the IBG, organizations may dynamically drill down and compare and contrast iSCSI SAN arrays by generating custom reports, including comprehensive strengths and weaknesses reports that evaluate a much broader base of features that what is found in the published Guide. Both the IBG and this Buyer's Guide may be accessed after [registering](#) for the [DCIG Analysis Portal](#).

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 5th 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.

Four Early Insights from the Forthcoming DCIG 2015-16 Enterprise Midrange Array Buyer's Guide

DCIG is preparing to release the DCIG 2015-16 Enterprise Midrange Array Buyer's Guide. The Buyer's Guide will include data on 33 arrays or array series from 16 storage providers. The term "Enterprise" in the name Enterprise Midrange Array, reflects a class of storage system that has emerged offering key enterprise-class features at prices suitable for mid-sized budgets.

In many businesses, there is an expectation that applications and their rapidly growing data will be available 24x7x365. Consequently, their storage systems must go beyond traditional expectations for scalable capacity, performance, reliability

and availability. For example, not only must the storage system scale, it must scale without application downtime.

These expectations are not new to large enterprises and the high end storage systems that serve them. What is new is that these expectations are now held by many mid-sized organizations—the kind of organizations for which the products in this guide are intended.

While doing our research for the upcoming Buyer's Guide, DCIG has made the following observations regarding the fit between the expectations of mid-sized organizations and the features of the enterprise midrange arrays that will be included in the Buyer's Guide:

Non-disruptive upgrades. In order to meet enterprises' expectations, storage systems must go beyond the old standard availability features like hot swap drives and redundant controllers to provide for uninterrupted operations even during storage system software and hardware upgrades. Consequently, this year's guide evaluates multiple NDU features and puts them literally at the top of the list on our data sheets. Over one third of the Enterprise Midrange Arrays support non-disruptive upgrade features.

Self-healing technologies. While self-healing features are relatively new to midrange storage arrays, these technologies help an array achieve higher levels of availability by enabling the array to detect and resolve certain problems quickly, and with no or minimal human intervention.

Self-healing technologies have been implemented by some storage vendors, but these are seldom mentioned on product specification sheets. DCIG attempted to discover which arrays have implemented self-healing technologies such as bad block repair, failed disk isolation, low-level formatting and power cycling of individual drives; but we suspect (and hope) that more arrays have implemented self-healing capabilities than we

were able to confirm through our research.

Automation. Data center automation is an area of growing emphasis for many organizations because it promises to reduce the cost of data center management and enable IT to be more agile in responding to changing business requirements. Ultimately, automation means more staff time can be spent addressing business requirements rather than performing routine storage management tasks.

Organizations can implement automation in their environment through management interfaces that are scriptable or through APIs and SDKs provided by storage vendors. Last year's Enterprise Midrange Array Buyer's Guide prediction that 'support for automated provisioning would improve in the near future' was correct. While less than 20% of midrange arrays in last year's Buyer's Guide exposed an API for third-party automation tools, the percentage has more than doubled to 50% in this year's guide. Provision of an SDK for integration with management platforms saw a similar increase, rising from 11% to 25%.

Multi-vendor virtualization. A growing number of organizations are embracing a multi-vendor approach to virtualization. Reflecting this trend, support for Microsoft virtualization technologies is gaining ground on VMware among enterprise midrange arrays.

The percentage of arrays that can be managed from within Microsoft's System Center Virtual Machine Manager (SCVMM) now matches vSphere/vCenter support at 33%. Support for Microsoft Windows Offloaded Data Transfer (ODX), a Windows Server 2012 technology that enhances array throughput, is now at 19%.

Although the gap between Microsoft and VMware support is narrowing, support for VMware storage integrations also continues to grow. VAAI 4.1 is supported by 90% of the arrays, while SIOC, VASA and VASRM are now supported by over 50% of

the arrays.

The *DCIG 2015-16 Enterprise Midrange Array Buyer's Guide* will provide organizations with a valuable tool to cut time and cost from the product research and purchase process. DCIG looks forward to providing prospective storage purchasers and others with an interest in the storage marketplace with this tool in the very near future.

A Primer on Today's Storage Array Types

Anyone who managed IT infrastructures in the late 1990's or early 2000's probably still remembers how external storage arrays were largely a novelty reserved for high end enterprises with big data centers and deep pockets. Fast forward to today and a plethora of storage arrays exist in a variety of shapes and sizes at increasingly low price points. As such it can be difficult to distinguish between them. To help organizations sort them out, my blog entry today provides a primer on the types of storage arrays currently available on the market.

The large number of different storage arrays on the market today would almost seem to suggest that there are too many on the market and that a culling of the herd is inevitable. While there may be some truth to that statement, storage providers have been forced to evolve, transform and develop new storage arrays to meet the distinctive needs of today's organizations. This has resulted in the emergence of multiple storage arrays that have the following classifications.

- **Enterprise midrange arrays.** These are the original

arrays that spawned many if not all of the array types that follow. The primary attributes of these arrays are high availability, high levels of reliability and stability, moderate to high amounts of storage capacity and mature and proven code. Features that typify these arrays include dual, redundant controllers, optimized for block level traffic (FC & iSCSI), and hard disk drives (HDDs). These are generally used as general purpose arrays to host a wide variety of applications with varying capacity and performance requirements. (The most recent DCIG Buyer's Guide on midrange arrays may be accessed via this [link](#).)

- **Flash memory storage arrays.** These are the new speed demons of storage arrays. Populated entirely with flash memory, many of these arrays can achieve performance of 500,000 to 1 million IOPS with latency at under a millisecond.

The two potential "gotchas" here are their high costs and relative immaturity of their code. To offset these drawbacks, many providers include compression and deduplication on their arrays to increase their effective capacity. Some also use open source versions of ZFS as a means to mature their code and overcome this potential client objection. Making these distinctively different from the other array types in this list of array types is their ability to manage flash's idiosyncrasies (garbage collection, wear leveling, etc.) as well as architecting their controllers to facilitate the faster throughputs that flash provides so they do not become a bottleneck. (The most recent DCIG Buyer's Guide on flash memory storage arrays may be accessed via this [link](#).)

- **Hybrid storage arrays.** These arrays combine the best of what both flash memory and midrange arrays have to offer. Hybrid storage arrays offer both flash memory and HDDs though what distinguishes them from a midrange

array is their ability to place data on the most appropriate tier of storage at the best time. To accomplish this feat they use sophisticated caching algorithms. A number also use compression and deduplication to improve storage efficiencies and lower the effective price per GB of the array. (*The most recent DCIG Buyer's Guide on hybrid storage arrays may be accessed via this [link](#).*)

- **Private cloud storage arrays.** Private cloud storage arrays (*sometimes referred to as scale-out storage arrays*) are defined by their ability to dynamically add (or remove) more capacity, performance or both to an existing array configuration by simply adding (or removing) nodes to the array.

The appeals of these arrays are three-fold. 1.) They give organizations the flexibility to start small with only as much capacity and performance as they need and then scale out as needed. 2.) They simplify management since administrators only need to manage one logical array instead of multiple smaller physical arrays. 3.) Organizations can mitigate and often eliminate the need to migrate data to new arrays as the array automatically and seamlessly redistributes the data across the physical nodes in the logical array.

*While these arrays possess many of the same attributes as public storage clouds in terms of their data mobility and scalability, they differentiate themselves by being intended for use behind corporate firewalls. (*The most recent DCIG Buyer's Guide on private cloud storage arrays may be accessed via this [link](#).*)*

- **Public cloud storage gateway arrays.** The defining characteristic of these storage arrays is their ability to connect to public storage clouds on their back end. Data is then stored on their local disk cache before it

is moved out to the cloud on some schedule based upon either default or user-defined policies.

The big attraction of these arrays to organizations is that it eliminates their need to continually scale and manage their internal storage arrays. By simply connecting these arrays to a public storage cloud, they essentially get the capacity they want (potentially unlimited but for a price) and they eliminate the painful and often time-consuming need to migrate data every few years. (A DCIG Buyer's Guide on this topic is scheduled to be released sometime next year.)

- **Unified storage arrays.** Sometimes called converged storage arrays, the defining characteristic of these storage arrays is their ability to deliver both block (FC, iSCSI, FCoE) and file (NFS, CIFS) protocols from a single array. In almost every other respect they are similar to midrange arrays in terms of the capabilities they offer.

The main difference between products in this space is that some use a single OS to deliver both block and file services while others use two operating systems running on separate controllers (this alternate architecture gave rise to the term "converged.") The "unified" name has stuck in large part because both block and file services are managed through a single (i.e. "unified") interface though the "converged" and "unified" terms are now used almost interchangeably.. (The most recent DCIG Buyer's Guide on midrange unified storage arrays may be accessed via this [link](#).)

Organizations should take note that even though multiple storage array types exist, many storage arrays exist that satisfy multiple classifications. While no one array model yet ships that fits neatly into all of them, DCIG expects that by the end of 2014 there will be a number of storage array models that will. This becomes important to those organizations that

want the flexibility to configure a storage array in a way that best meets their specific business and/or technical requirements while eliminating the need for them to buy another storage array to do so.

Early Insights from the Upcoming DCIG 2014 Enterprise Midrange Array Buyer's Guide

The time for the release of the refreshed DCIG 2014 Enterprise Midrange Array Buyer's Guide is rapidly approaching. As that date approached, we have been evaluating and reviewing the data on the current crop of midrange arrays that will be included in the published Buyer's Guide (*information on over 50 models*) as well as the models that will be included in DCIG's online, cloud-based Interactive Buyer's Guide (*over 100 models.*) Here is a peak into some of what we are finding out about these models in regards to their ability to deliver on data center automation, VMware integration and flash memory support.

Data Center Automation

Data center automation is an area of emphasis for many organizations because it promises to facilitate efficient management of their data center infrastructure and enable a more agile response from IT to changing business requirements. Ultimately, automation means more staff time can be spent addressing business requirements rather than managing the routine tasks of a data center.

Organizations can implement automation in their environment

through management interfaces that are scriptable and offer additional enhancements with API and SDK support.

Support for automated provisioning is an area where improvement in the near future is expected. Currently, **less than 20%** of midrange arrays featured in this upcoming Buyer's Guide expose an API for third-party automation tools, while **11% provide an SDK** for integration with management platforms. As more organizations place a premium on automating their storage environment, these numbers should go up.

A higher percentage of these arrays support **automated storage tiering, which is offered by 45% of arrays**. This automated tiering capability can be important for achieving maximum benefits from flash memory when using flash for more than just a larger cache.

Similarly, **40% natively support the reclamation of freed blocks** of thinly provisioned storage. These freed blocks are then available for reuse. Native support for this capability eliminates the cost and additional infrastructure complexity associated with licensing a third party product or the inefficiency associated with manual reclamation processes.

Along the same lines, 21% of arrays are recognized by third party software, such as Symantec Storage Foundation, that can simplify storage management by reclaiming freed blocks of thinly provisioned storage automatically.

VMware Integration

In general, DCIG emphasizes advanced software features in the DCIG 2014 Enterprise Midrange Array Buyer's Guide. This is especially true of integration with VMware vStorage APIs such as VAAI (vStorage APIs for Array Integration) and VASA (vSphere Storage APIs for Storage Awareness). The VAAI and VASA APIs can dramatically improve overall data center performance.

Given the wide adoption of VMware by enterprises, it follows that they are seeking hardware that can take advantage of the “force multiplication” these APIs provide for existing and future VMware deployments.

The good news is that **62% of the midrange arrays** included in this upcoming Buyer’s Guide support all of the VAAI 4.1 APIs. However, **only 5 arrays out of 53 support the full set of VAAI v5 features**. Of the VAAI v5 features, **Dead Space Reclamation (SCSI UNMAP) fares best** with 26% of arrays supporting this feature.

Similar to the currently low support for VAAI v5.0, less than a fourth of the arrays support VASA. These integrations are key to the software defined data center and to minimizing ongoing management overhead for the large number of data centers that utilize VMware.

Robust VMware support is a product differentiator that matters to many potential array purchasers, and is an area where we expect to see further improvement in the coming year. Those organizations embracing VMware as their primary hypervisor will want to pay particular attention to how an array’s VMware support maps to their requirements

Flash Memory Support

Flash memory is clearly of growing importance in data center storage. Within the enterprise midrange array segment of the market, the importance of flash memory is demonstrated by the fact that **77% of the arrays** in this upcoming Buyer’s Guide now support the use of flash memory in addition to traditional disk drives.

Nevertheless, just 45% support automated storage tiering, a technology that helps get the most benefit from the available flash memory. Also, **only 15% of arrays implement any of the flash memory optimization techniques**—such as write coalescing—that enhance both performance and reliability. So

while support for flash memory in midrange arrays has grown dramatically, the depth of integration still varies widely.

Fusion-io Acquires NexGen to Accelerate the Process of Data Center Transformation

Last week's acquisition of NexGen Storage by Fusion-io was greeted with quite a bit of fanfare by the storage industry. But as an individual who has covered Fusion-io for many years and talked one-on-one with their top executives on multiple occasions, its acquisition of NexGen signaled that [Fusion-io](#) wanted to do more than deliver an external storage array that had its technology built-in. Rather Fusion-io felt it was incumbent for it to take action and accelerate the coming data center transformation that it has talked and written about for years.

When I first learned of Fusion-io's decision to [acquire NexGen Storage](#), it caught me a bit by surprise. I had talked with its CEO David Flynn as well as with Rick White (its CTO, CMO and general person extraordinaire) on a number of occasions and their overriding opinion (paraphrased) was that disk storage, as it is currently used in production data center applications, was on the way out and would need to be replaced by a new faster tier of storage.

While many in the industry hold to this opinion, Fusion-io's position differed (at least early on) in that it contended this new tier of storage would emerge not as a new tier of storage (such as Tier 0 or flash memory/SSD-based storage) but as a new memory tier. It contended (and still does to the best

of my knowledge) that the processing overhead and time involved to convert data in memory to data storage and back again so that it is usable for processing would become unacceptable over time.

They contend (*and rightfully I might add,*) *“Why not just store data in memory all of the time and eliminate this unnecessary transformation of data that needs to occur during processing?”* This thought process and positioning explains why Fusion-io puts its [ioDrive](#) technology inside of servers, refers to it as ioMemory and ships with its own driver that eliminates all of this back and forth data transformation that is currently needed to store data on disk. It also explains why others such as EMC and Dell are actively working on solutions that compete with Fusion-io because of the huge acceleration in performance that this approach offers over even using flash memory based storage arrays.

However if Fusion-io's approach is so much better than any flash memory array shipping today, then why buy NexGen? This raises the following questions that I wish to examine in this blog entry:

- One, while Fusion-io's technology has been widely accepted, it probably has not experienced the broad market adoption and displaced traditional storage solutions as quickly as many thought it might despite the vast improvement in performance it offers over these arrays, to include flash-based arrays?. The question is, ***“Why has this not occurred?”***
- Second, ***if Fusion-io's technology is so much better than using storage arrays, why buy NexGen Storage?*** Granted, NexGen used Fusion-io cards in its storage systems. But the ugly truth is that NexGen does not eliminate the data transformation process that needs to occur from memory to storage back to memory again since NexGen Storage is connected to servers via traditional 1Gb and

10 Gb Ethernet connections. So while NexGen may be faster than other traditional storage arrays that use SSDs because it uses Fusion-io cards, it does not offer the same level of performance as a native Fusion-io card internal to our server.

I'll answer these two questions in the order that I posed them.

As to why Fusion-io's technology has not taken off the way some thought it might, there are at least four reasons for it.

1. **Too expensive.** While Fusion-io ioDrives are high performing, that performance comes at a cost and the vast majority of application cannot justify paying that much for the boost in performance it offers.
2. **Cannot be shared.** When a Fusion-io ioDrive is put in a server, it can only be used by the application(s) or virtual machine(s) on that physical server. This adds to the difficulty in cost-justifying the purchase of a Fusion-io card and explains why companies like NexGen that used Fusion-io in their systems emerged in the first place: to better share and optimize the available performance and capacity on Fusion-io's ioDrives.
3. **Flash memory and SSD solutions are "good enough" for now.** Going from 10K or 15K RPM spinning FC or SAS disk in a storage array to flash memory or SSD has resulted in 3-10x or greater improvements in application performance. While the performance is not as good as running Fusion-io inside the server, it is still a substantial improvement in performance at a cost that more organizations can absorb.
4. **Difficult to upgrade and/or maintain.** Adding more storage capacity and performance inside a server or replacing a faulty ioDrive can be problematic since it may require taking the server offline to do so. Adding more capacity or making a repair on an external storage array is much easier since they are typically designed

for 24x7 uptime.

It is for these reasons and others that Fusion-io has arguably not achieved the adoption rate that many thought it might. This also explains why hybrid and flash memory storage providers like NexGen have emerged since they do a better job of making flash available to a greater number of applications at a lower cost.

But if that is the case and where all storage will eventually end up, Fusion-io should simply look to partner with as many of these providers as possible and put its ioDrive technology inside of them. If anything, ***its acquisition of NexGen Storage could almost be seen as a detriment to its plans for growth.*** After all, why would other storage providers want to use Fusion-io's technology if Fusion-io is building its own storage arrays? In that sense, Fusion-io would have been better off staying independent so it could become the Intel (*so to speak*) of flash memory based arrays.

This gets to the real heart of why I believe Fusion-io bought NexGen Storage. Yes, Fusion-io wants to capitalize on the current craze in flash memory based arrays and have its own product offering in this space. This it accomplishes with its acquisition of NexGen Storage.

By acquiring NexGen Storage, Fusion-io in the near term provides some answers to the valid criticisms mentioned above about its current suite of product offerings. It drives down costs, it makes its technology shareable by storage attached servers, it provides the "good enough" storage solution that SMEs need now and it delivers the solution in a manner that is easier for them to maintain and upgrade.

Bringing NexGen into its fold also addresses another key concern that SMEs have: Support. While anyone can take a server, install a couple of Fusion-io ioDrives in it, put a Windows or Linux OS on it and call it a "storage system," by

Fusion-io bringing NexGen in house, SMEs get the interoperability matrices and application support guarantees that they need to host their applications on such a storage solution.

Fusion-io having solved these near term tactical issues, it can leverage NexGen to achieve its more strategic goal: creating a new memory tier in the data center that will accelerate the data center transformation.

As Fusion-io has previously [told](#) DCIG, being a hardware company is really not core to its DNA: deep down Fusion-io wants to be a software company but needs hardware to deliver on the promise of its software.

So here is what I think is going on. NexGen Storage already uses Fusion-io internally and offers [ioControl](#) software on its arrays. Therefore it is not a stretch to believe and which, in fact, may be the most logical conclusion to draw, is that Fusion-io will look to extend its existing server-based ioMemory capabilities into arrays residing in the the storage network starting with its newly acquired NexGen arrays. While I would not look for such capabilities in the next few months or even the next year, I see it as almost inevitable.

While companies are currently enthralled with the 3-10x improvement that flash memory arrays provide, that thrill will soon be gone and they will again be looking for another 3-10x or greater jump in storage performance without a huge price increase, By Fusion-io extending its ioMemory capabilities into its NexGen arrays, there is no reason to believe Fusion-io cannot meet these forthcoming SME expectations . More importantly, Fusion-io is arguably better equipped to make this leap than competitors since Fusion-io's DNA is to be a software company that operates at the memory as opposed to the storage level in the data center stack.

Everyone knows that a data center transformation is underway

and it is occurring in ways which few of us fully grasp. However from my conversations with Fusion-io's executives, I do get the sense that they do more than grasp the level of change that is about to occur. They see the role that they could potentially play in facilitating it. Fusion-io's acquisition of NexGen signals that it is tired of playing a passive role in bringing about this data center transformation and is going to, in the very near future, start shipping products that hastens this occurrence..

The Two Forces Driving the Evolution (or is it Revolution?) in Enterprise Midrange Array Architectures

In May 2010 DCIG released its first-ever Midrange Array Buyer's Guide in which we covered 70+ models from over 20 vendors. Fast forward just three (3) short years later and DCIG is on track to release not one, not two, not three no, not even four Buyer's Guides on enterprise midrange arrays but five distinct Buyer's Guides on this topic! So what has changed in just three (3) short years that DCIG feels the need to produce so many? To understand this requires a closer look at the forces that are driving the evolution and revolution in enterprise midrange arrays.

In 2010 when DCIG released its first Midrange Array Buyer's Guide, the midrange array market was already very mature. There were multiple providers of storage arrays (over 30,) multiple models from these providers (nearly 150 models) and

an increasingly sophisticated set of software available on these arrays.

The storage management software (*or firmware as it is commonly called*) was generally not as sophisticated as found on larger enterprise arrays (*the EMC VMAX or the HDS VSP.*) However it certainly offered many advanced features. Even three (3) years ago, automated storage tiering, snapshots, replication, thin provisioning and many others were commonly found on these arrays.

Despite the maturity of midrange arrays, a lot has changed in the last three years that DCIG now sees it necessary and can justify producing five Buyer's Guides in a single year on enterprise midrange arrays. In short, there are two specific forces driving midrange array segmentation. These are:

1. Unstructured Data Growth/Big Data. *As an analyst I regularly run across statistics like 30%, 50%, 80%, and, in some extreme cases, even 400% data growth in some environments. However organizations are feeling the impact of this data growth in real time and, they assure me, their storage budgets are growing nowhere near as fast as their data is.*

If they get single digit increases in their budgets year-over-year, they are thrilled. So their annual challenge is to make single digit increases in budget stretch to cover double and triple digit percentages in data growth.

One way in which they are doing so – especially small and midsized organizations – is by turning to Unified Storage Arrays (access free download of DCIG Buyer's Guide on this topic [here](#).) These can be tuned to achieve high capacity, high performance or some combination of both. This is done by deploying a mix of high performance storage capacity (flash memory/SSDs) and higher storage capacity, lower performing and more economical 3 & 4 TB SATA drives in a single array.

Then so any application can access this various types of storage capacity, these arrays make the storage accessible over any available storage networking protocol. These could be high performance SAN protocols (8 Gb FC or 10 Gb Ethernet) or 1 Gb NAS protocols (CIFS or NFS). In this way, organizations can buy a single storage array, configure it with the type of storage and networking interfaces they need to accommodate their mixed needs of unstructured data growth and performance hungry applications, and do so economically.

Enterprises are also turning to unified storage arrays but in these environments, they are often architected as scale-out storage arrays. In these configurations, organizations can add or even remove performance, capacity or both on an ad hoc basis with minimal effort and without increasing their ongoing management workload. More notably, these tend to scale to much higher capacities (into the petabytes) whereas other midrange arrays only scale into the hundreds of terabytes.

2. Performance Hungry Apps. Even as recently as a few years ago, if an array – any array – did read or write I/O in as little as a few milliseconds (around 5 ms) it was considered blazing fast. Today it seems 5 ms response times will barely get you in the performance conversation when discussing databases.

Further, as organizations virtualize more of their applications and put more VMs on fewer physical machines, this puts a lot of pressure on storage arrays to keep up. Aggravating the situation, server and networking technologies have literally experienced ten-fold or greater increases in performance over the last few years while storage arrays have only seen incremental increases in performance.

This has led to the emergence of two different types of midrange storage arrays – flash memory and hybrid – that have contributed to giving these arrays the 2 – 10x increases in

performance that they have needed to keep up with application demands and improvements in other parts of the technology stack.

Both of these arrays use flash memory and/or solid state drives (SSDs) to accelerate performance. The main difference between the two is that flash memory storage arrays only offer flash memory as a storage option while hybrid storage arrays use both flash memory and spinning disk to store data. As a result, flash memory arrays are generally faster though more expensive than hybrid storage arrays.

The primary use cases for both of these arrays due to their cost and more limited capacities are primarily for specific high performance applications workloads. However as their capacities increase, flash memory prices drop and other technologies such as compression, deduplication and thin provisioning are implemented on these arrays, expect them to be used more widely for other applications.

The combination of these two forces has led to dramatic changes in the architecture of enterprise midrange arrays. While one can still get big boxes full of spinning disks connected via FC to servers, there are now many more options than what was available in the past. They can be capacity focused. They can be performance focused. Storage can be delivered over a number of storage networking protocols. These combined are leading to an evolution – and some would even say a revolution – in how midrange arrays are architected and what they will look like in the years to come.

DCIG 2013 Midrange Unified Storage Array Buyer's Guide Now Available

DCIG is pleased to announce the availability of its inaugural DCIG 2013 Midrange Unified Storage Array Buyer's Guide that weights, scores and ranks over 100 features on 30 different storage arrays from eight (8) different storage providers. This Buyer's Guide provides the critical information that small and midsize enterprises particularly need in regards to storage arrays that will need to serve a variety of purposes within their organization. These purposes may include storing large amounts of unstructured data such as files and emails, hosting virtualized and high performance applications and even serving as a target for archival and backup data stores.



Many small and midsize enterprises (SMEs) deal with their own version of Big Data. Maybe best described as "**Big Data Lite**," their data stores are unlikely to cross the "*magic*" 1 PB number that has come to be associated with Big Data. Rather they are more likely to have a couple of hundred of TBs of data under management.

A November 2012 IDG Enterprise study provides some insight

into what is going on in SMEs. This IDG Enterprise study defined “*Big Data*” as “*large volumes of a wide variety of data collected from various sources across the enterprise.*” It is when “*Big Data*” is defined in this context that it quickly gets interesting in SMEs.

The study [found](#) that the average organization already manages 194.4TB of data and expects its data to grow by over 50% to 296.7TB in the next 12 to 18 months. Even a couple hundred terabytes at a 50% growth rate qualifies as high velocity in these size shops.

It is as more SMEs find themselves in this Big Data Lite category that they recognize it is time to move from direct attached storage located in hot, dusty closets into the “*midrange array*” class of storage solutions. These midrange arrays bridge the gap between the mid-teens of terabytes up to the low petabytes in a standalone appliance that can be accessed and shared by a number of devices.

The midrange array category is quite large and is usually broken down into several additional categories. The most basic breakdown is by how the storage is accessed: Storage Area Network (SAN), Network Attached Storage (NAS), or both. Solutions that can support both SAN and NAS are referred to as “***Unified Storage.***”

This Buyer’s Guide represents DCIG’s first foray into midrange unified storage arrays as DCIG believes this is the new sweet spot for storage arrays that most clearly align with SMEs and their “***Big Data Lite***” needs.

This DCIG 2013 Midrange Unified Storage Array Buyer’s Guide should help organizations quickly ascertain what midrange unified storage arrays are on the market, what features they possess and then help expedite their decision making and buying process.

DCIG sees midrange unified storage arrays as being well-suited

for SMEs as they:

- Support both NAS and SAN protocols thereby reducing duplication of resources, simplifying the IT infrastructure, and easing the transition of legacy systems from DAS to NAS or SAN
- Leverage standard NAS and SAN protocols so most devices will be “plug and play” when connecting to the midrange unified storage array
- Reduce cost by eliminating redundant processing power and wasted storage capacity
- Ease storage management by centralizing storage into a single namespace and user interface
- Facilitate centralized security using existing authentication schemes such as Active Directory and/or Kerberos/LDAP
- Scale up storage capacity through the addition of new disks and/or nodes

The DCIG 2013 Midrange Unified Storage Array Buyer’s Guide Top Ten products include (*in alphabetical order*): the EMC VNX 5500, 5700 and 7500 models, the HDS Unified Storage 110, 130 and 150 models and the NetApp FAS 3220, 3240, 3250 and 3270 models.

Of note is that the NetApp FAS3200 series models took the top four spots in this Buyer’s Guide. This is the first time a storage provider has ever done so in any DCIG Buyer’s Guide. Factors that particularly contributed to the NetApp FAS3200 models scoring so well were its full integration with VMware vSphere, the same management software across its entire line of midrange unified storage models and its read and write flash-based caching.

In doing its research for this Buyer’s Guide, DCIG uncovered some interesting statistics about midrange unified storage arrays in general:

- 100% support both user and group quotas
- 97% support some form of thin provisioning
- 84% support sub-volume tiering
- 78% support automated storage reclamation
- 65% have a starting list price of under \$50,000
- 30% support block-level deduplication
- 23% support file-level deduplication

As with prior DCIG Buyer's Guides, it accomplishes the following objectives for end users:

- Lists each midrange unified storage array model by vendor
- Lists out features of each midrange unified storage array showing key features supported or not supported by each product
- Scores the features most relevant to end users
- Provides "at a glance" reference for companies evaluating specific midrange unified storage arrays or midrange unified storage array features
- Provides a midrange unified storage array ranking showing how vendor products compare against similar products on the market
- Offers recommendations as to which midrange unified storage array rankings and products best align with their specific data storage objectives
- Provides 30 midrange unified storage array e data sheets from 8 different vendors so organizations may compare systems for one or many technology providers.
- Facilitates and accelerates the process of organizations obtaining bids on competitive products

The DCIG 2013 Midrange Unified Storage Array Buyer's Guide is immediately available in both a condensed and a full version. These may be downloaded for no charge with registration by following the appropriate link listed below.

- DCIG 2013 Midrange Unified Storage Arra – [Condensed](#)

Why FCoE and iSCSI Trump Infiniband in Today's SSD Deployments; WhipTail CTO Interview Series Part V

In this final installment of our blog series on WhipTail Technologies, a Solid State Drive ([SSD](#)) array provider with some impressive features and capabilities, I am continuing my discussion with WhipTail Technologies Chief Technology Officer, James Candelaria. [Last time](#), we looked at how WhipTail implements software RAID on its devices. Today, we will be discussing the different transport protocols supported by the WhipTail array and why the FCoE and iSCSI protocols trump Infiniband in today's SSD deployments.

Ben: *We have talked a bit about the wide array of hardware you both use and support. Can you tell me a little bit about the storage protocols supported by WhipTail?*

James: [WhipTail](#) is fully multi-protocol. We support fibre channel ([FC](#)), [iSCSI](#), [CIFS](#) and [NFS](#), as well as the SCSI [RDMA](#) protocol. SSDs are all about latency. It is all about the microsecond response time. The minute you get onto a fiber channel or Ethernet fabric, you are talking about 120 microseconds. You are talking about doing a Remote Directory Memory Access ([RDMA](#)) over IP and you are talking about maybe three to ten micros – massive latency advantage there. So we are fully multi-protocol serviceable.

Ben: And that is because of where you are in the stack naturally, right?

James: Exactly.

Ben: I did a DCIG Buyer's Guide looking at [Enterprise Scale-out Storage](#). I was surprised how little Infiniband there is out there in the market.

James: So am I. It is a great transport. But you know what, I understand. **People just do not want another Layer 2 typology in their environment.** It costs money to maintain it. Even though the Infiniband switches and Host Channel Adapters today are dirt cheap, there is another cable I have to run. I have to find somebody who understands what a subnet manager is. It is not without its intangible costs.

Ben: I think that is why a lot of places are going with Fiber Channel Over Ethernet and iSCSI – just because they understand Ethernet.

James: Yes, Ethernet is a lowest common denominator in the data center. To be honest with you, when you can make iSCSI go end-to-end in the 120 – 150 microseconds realm, do you really need to go much faster? You have to ask yourself that question. I am a speed freak at heart. I race cars on the weekend, so I like to push the envelope.

But in the data center, it is not about that. It is about meeting your business objectives. I honestly think that Whiptail combined with a high speed Layer 2 fabric you already own, like 10 gigabit Ethernet, meets 90 percent of all business objectives. So I try to be less emphatic about my Layer 2 choices. We all have our preferences.

Ben: Do you see WhipTail as being part of the VDI revolution right now?

James: VDI's been a huge success for us because some of our

corporate DNA. If you look at some of our senior executives, myself, our CMO Brian Feller, Andy Flesch, our VP in sales, all of us come from a virtualization background. So we get it. And we are effectively able to leverage our knowledge and contacts in that space to solve a glaring hole in the VDI deployment model. Which was “we got all the brokers, we could fix all the display problems, but no one ever told us we need to account for storage performance!”

One day somebody woke up and said “we actually have to do this.” We were standing right there with the solution and it has been very successful for us. That is for a couple reasons. Obviously first – right place, right time. But the other guiding factor is that there is a new type of buyer that has been empowered to buy storage.

The VDI buyer has zero allegiance to the incumbents. He has never bought storage before. That buyer has been told “you have to make VDI work, you have to get 5,000 seats up tomorrow.” He has never bought a [NetApp](#). He has never bought a [Symmetrix](#). He has never bought a Hitachi [USP V](#).

He is looking for something to solve his problem. He is not looking for a storage platform. So that has been one of the reasons we have been really successful in that market.

But as the product progresses and we mature as a company, other markets are obviously naturals for us to get in to – transaction processing, transaction. We have been very, very successful there, too.

We have a lot of customers that are Oracle shops that found us just too attractive – it was not really their first line of thinking, but somebody approached them and they said “*Hey, how about an extremely low latency piece of storage to speed up that report?*”

They gave it a shot and they took report run time down from 24

hours to one hour. All of a sudden their eyes lit up.

I had one customer in particular bought one for a specific use case, and then bought four more because he just decided everything in my data center needs to be this fast -just started buying them like they were candy.

So we have been successful in marketing towards the application owners with problems – tactical solution right now for that report that does not run very well, that PDI performance crunch, that one online transaction database thing.

And just again, I cannot really talk about futures, but watch the space – we are going to go from a tactical company to a strategic company in the near future. So there is some pretty cool stuff coming.

Ben: *Another thing I find interesting about WhipTail is that, by extending the life of these flash RAMs, you are becoming a force multiplier.*

James: That is actually a phrase we have used around the office quite a bit is “force multiplier.” Alone, SSDs and NAND can only do so much. But when you couple them with an error mitigation strategy, they can become so much more to the data center, more than just a point solution that you think is going to be disposable at the end of the year because of wear, which will stop you from deploying them widely.

Once you get those constraints under control, you can start thinking about using them much more widely across every workload that you have that needs to go faster. Let’s face it, IT, the governing factor on business intelligence has always been performance.

Ben: *I think what is unique here is that you are catching something that is not just improving performance; it is drastically improving the lifespan of the product that you are*

using.

James: Exactly. WhipTail is one of those rare things where you get additional performance and you actually get additional endurance simultaneously, which is fairly rare. Intel is able to do it with processors and what not. But on everything else they will usually go together. It is like the old adage: p; “Good, fast, cheap – pick two.”

Ben: Thank you, James. This has been a lot of good information.

James: Thanks, Ben. I appreciate your time.

In [Part I](#) of this interview series, James explained the SSD garbage collection problem and how WhipTail handles it.

In [Part II](#) of the series, James discussed how WhipTail is optimizing SSD performance while minimizing the deficiencies of MLC flash.

In [Part III](#), Candelaria and I discussed how WhipTail deals with manufacturer variations in SSD drives.

In part [IV](#) in this series, Candelaria explains how and why WhipTail uses software RAID in its SSD appliance.

SSD Optimization Techniques that Minimize MLC Flash Deficiencies – Interview

Series Part II with WhipTail CTO Candalaria

Today is part 2 of an interview I recently did with WhipTail Technologies Chief Technology Officer, James Candalaria, an emerging provider of SSD storage solutions. In my [last entry](#), he and I discussed one major roadblock to widespread enterprise SSD adoption: the performance penalty incurred by garbage collection. This time, we'll look at how WhipTail optimizes SSD performance while minimizing the deficiencies of MLC flash.

Ben: *We've talked a bit about the challenges of SSD and how [WhipTail](#) addresses them as well as how you're solving performance-related issues with SSDs used in RAID arrays. Are you doing any other optimizations in your stack or are you just trying to quiet down the interaction between the flash cells?*

James: WhipTail's claim to fame right now is the mitigation of the deficiencies in MLC flash. In most of our intellectual property, we sit directly underneath the transport stack and above the RAID stack. So we optimize as data comes in on the fiber channel or iSCSI, or even out of an ext3 or xfs file system.

As data comes into our block device, it manipulates the data and submits it to the RAID layer using a fairly standard RAID stack. So it's an ideal place for things like enhanced ECC data reduction strategies possibly later on, things like that. So I can't speak to some features, but you can imagine by our architecture that we're strategically placed.

Ben: *So what WhipTail is doing is optimizing your block sizes to fit exactly within the size of a RAID stripe?*

James: Yes, exactly, right to the borders of the write

stripes and the borders of an array slot on the flash media. One of the things I always like to point out here is that we've been shipping products for over two and a half years to customers.

So we have a lot of empirical data about how these flash controllers work over time and the behaviors of different controllers and firmware over a large install base, not necessarily just how they work on a spec sheet.

We have over 100 customers deployed in the field and probably the largest install base of VDI (Virtual Desktop Infrastructure) flash storage out there. So there are thousands and thousands of VDIs sitting on top of us, and that gives us a real good feel for how accurate our projections have been.

One of the things WhipTail did most recently was to look at the hardware and software layers of one of our oldest customers who have been in production for over two and a half years. We looked at them and did the math on figuring out how long their array would last on current endurance rates. The answer I got back was 22 more years of service before we get flash endurance levels that were unacceptable.

Ben: *What do you think the service life would have been without your software stack?*

James: Two and a half years ago we were shipping 54 nanometer flash, which had endurance levels of 10,000 cycles, double the industry's standard endurance. If we had not done something about write amplification, I think we probably would have seen a wear out in about a year, maybe less, primarily due to the fact that the workload was a fairly heavy 8K random write I/O.

Ben: *With a 48 MB buffer for 24 two MB erasures, are there additional buffering requirements to prevent blocking while you're flushing?*

James: We do keep a set of overlapping buffers. Currently I think it's a total of four overlapping outbound buffers. So while one set of buffers is flushing, the other one is filling, so we don't end up with too many stalls. Even if you're pounding at a full line rate all day long, you'll see a fairly linear performance experience.

Ben: *Any time that things are sitting in a RAM buffer like this and have not been written, you risk data loss, right?*

James: Absolutely.

Ben: *So how does WhipTail handle that situation?*

James: Very good question. We handle it in two different ways: First, we do have a maximum data age timer. WhipTail will flush buffers if they are not full within half a second. That's tunable, though, depending on the customer's tolerance for a hole in their data set. Second, customers are required to run a solid UPS. Your storage has got to be behind a good UPS. So those are the two mitigation strategies right now.

Ben: *And then what does WhipTail do for redundancy? Is this a single point of failure here? Or can I set up multiple controllers?*

James: We do sell HA devices. Currently our HA relies on synchronous replication between two units. So we have two units side by side over a 10 GB umbilical. And that allows you to ensure that your data exists in both places.

In case of a fail over, one will assume the identity of the other. It's not a traditional HA, dual controller architecture. Instead, you have two storage processors in one single shelf of storage. In my experience, if you have a failure, nine times out of ten it's going to be the underlying media, not the underlying controller that's going to be the problem.

Ben: Believe it or not, I have had a controller fail on me once. And then the HA failed. And chaos ensued. So I'm sensitive to it.

James: I've been down that road too. Prior to founding Whiptail, I spent 10 years in high end Fortune 500 consulting. I've found that it's not "if" it will fail, it's "when." We're not dealing with perfect machines here. With that in mind, we always provide for redundancy and think about contingency situations.

In the next [blog entry](#) in this series, I will continue my discussion with Candelaria looking at how WhipTail deals with variances between each SSD manufacturer's hardware and firmware.

In [Part I](#) of this interview series, James explained the SSD garbage collection problem and how WhipTail handles it.

In the fourth [installment](#) in this series, Candelaria explains how and why WhipTail uses software RAID in its SSD appliance.

In [Part V](#) of this series, James and I discuss the hardware and software supported by WhipTail and why FCoE and iSCSI trump Infiniband in today's SSD deployments.

DCIG 2012 Midrange Array Buyer's Guide Now Available for Free Download

Today DCIG is very excited to announce the availability of its updated [DCIG 2012 Midrange Array Buyer's Guide](#) that weights,

scores and ranks **over 90 features** on **more than 50 midrange arrays** from **18 different storage providers**. However the reason that DCIG believes users will find this guide even more helpful and insightful than the prior **DCIG 2010 Midrange Array Buyer's Guide** is that it does more than just look at the newest midrange array models and evaluate the latest and greatest features on each one. It takes an in-depth look into how well each midrange array integrates with VMware and supports its vStorage APIs.



From DCIG's perspective, as 2012 approaches **there has been a fundamental shift** since 2010 in **how most organizations approach purchasing a midrange storage array**. While the emphasis on hardware characteristics is still an important consideration, features such as thin provisioning, replication, integration with new vStorage APIs and the ability to scale-out are increasing in significance.

The statistics surrounding the adoption rate and growth of server virtualization in general and VMware specifically over the last few years are staggering. Consider:

- Over 19,000 customers, partners, press and analysts [attended](#) the 2011 VMworld conference.
- A new VM is [created](#) somewhere in the world every 6 seconds.
- 10 million VMs [existed](#) at the end of 2010.
- VMware customers will have [crossed](#) the 50 percent virtualization mark by the end of 2011.

But as server virtualization in general and VMware vSphere specifically continue on this explosive growth rate it puts tremendous pressure on the backend storage infrastructure. As more VMs are created organizations must grow and manage their backend storage infrastructure in a dynamic, flexible manner to keep up with this growth.

To do so each VMware deployment requires a backend storage infrastructure that is manageable and scalable in both storage capacity and performance. In response to these needs, new VMware vStorage APIs included in vSphere 4.0 and vSphere 5.0 dramatically improve both the storage management experience for organizations and the performance that midrange arrays should be able to deliver.

Yet realizing these improvements in how midrange arrays are managed and perform in virtualized environments only occurs if the storage provider takes advantage of these vStorage APIs and implements them on their respective midrange arrays.

It is for this reason that the [**DCIG 2012 Midrange Array Buyer's Guide**](#) put a heightened focus on both the midrange array's vSphere integration and its array management software. The level of vSphere integration as covered in this Buyer's Guide should provide organizations a high level of insight into how well positioned each midrange array is to optimally host and manage the growing number of VMware virtual machines (VMs) that exist in their environment.

This is not to imply that a midrange array's hardware or storage networking connections should be ignored. If anything, the importance of deploying a reliable, robust midrange array has increased, not decreased, with the advent of storage virtualization.

Midrange arrays are likely hosting the data for many more applications than before they were virtualized so any type of hardware failure has a more dramatic impact on the

environment. So in this respect, the type of hardware a midrange arrays offers and how well a midrange array accounts for “routine” events like HDD rebuilds should still be carefully considered.

It is in this context that DCIG is pleased to present its [**DCIG 2012 Midrange Array Buyer's Guide**](#). As prior Buyer's Guides have done, it puts at the fingertips of organizations a Buyer's Guide that provides them with a comprehensive list of midrange arrays that can assist them in this all-important buying decision while removing much of the mystery around how midrange arrays are configured and which ones are suitable for which purposes.

This [**DCIG 2012 Midrange Array Buyer's Guide**](#) accomplishes the following objectives:

- Provides an objective, third party evaluation of midrange arrays that evaluates and scores their features from an end user's viewpoint
- Includes recommendations on how to best use this Buyer's Guide
- Scores and ranks the features on each midrange array based upon the criteria that matter most to end users so they can quickly know which midrange arrays are the most appropriate for them to use and under what conditions
- Provides data sheets for 53 midrange arrays from 18 different storage providers so end users can do quick comparisons of the features that are supported and not supported on each midrange array
- Provides insight into which features on a midrange array will result in improved performance
- Provide insights into what features midrange arrays offer to optimize their VMware integration
- Gives any organization the ability to request competitive bids from different providers of midrange arrays that are “apples-to-apples” comparisons

In terms of how users intend to leverage this Buyer's Guide, Gareth Segree, a systems engineer with The Gleaner Company Limited., says, *"The Gleaner Company was looking to upgrade its SAN and the DCIG 2012 Midrange Array Buyer's Guide was the primary resource that it used to evaluate available solutions. We even used the Buyer's Guide to evaluate midrange arrays that were unfamiliar to us but appeared to meet our requirements."*

The [DCIG 2012 Midrange Array Buyer's Guide](#) is available immediately and may be downloaded for no charge with registration by following this [link](#).

Initial Observations from the Forthcoming DCIG 2012 Midrange Array Buyer's Guide

If you are a regular follower of the DCIG blog site you may have noticed that there has been a noticeable lack of blogging activity on DCIG's site this week. Unfortunately it is not because I have been taking a vacation, fishing or merely lounging by the lake. Rather I have been locked away in my office completing the background research associated with the upcoming release of the DCIG 2012 Midrange Array Buyer's Guide due out in the 4th quarter of 2011. Out of that some interesting early observations have emerged.

In preparing this updated version of the Midrange Array [Buyer's Guide](#), DCIG had a lot of feedback to draw upon to improve this Buyer's Guide from the prior one as it looks to make it even more pertinent to help enterprise users make enterprise storage buying decisions. Here are some of the

changes people will notice when the DCIG 2012 Midrange Array Buyer's Guide comes out in Q4 2011:

- **Fewer products.** The DCIG 2010 Midrange Array Buyer's Guide released in Q2 2010 evaluated over 70 products from 20+ providers. This one will still cover about the same number of providers but fewer products – right now it looks to be around 50. While I initially expected the number of products that would be covered to increase, the number actually decreased. This was due to stricter criteria for inclusion in the Buyer's Guide and the inability to find reliable and comprehensive publicly available information on the Internet for those companies that did not complete the DCIG survey about their midrange arrays.
- **More details on VMware vSphere integration.** VMware was important in 2010 but it was almost all on the server side. In the last 18 months VMware has extended its influence such that enterprise users are now more aggressively looking for and buying midrange arrays that support the VMware vStorage APIs such as VAAI and the newly announced VASA.
- **Storage capacity has been de-emphasized.** In the last Buyer's Guide the storage capacity on each midrange array was given a great deal of weighting in the final results. That will be less so the case in this release of the Buyer's Guide for the next reason.
- **Storage efficiency and intelligence have been emphasized.** Storage vendors have been adding efficiency and intelligence features into storage arrays over the last decade. But it has been difficult for most enterprises to capitalize on that extra intelligence unless they were using a specific application or had storage administrators skilled in the software of their particular storage array (hence the heavy emphasis on capacity in the last Buyer's Guide.)

The wide adoption of VMware by small, midsize and large enterprises has changed all of that. VMware IS the application that everyone owns that can take advantage of the efficiency and intelligence features in these midrange arrays without these organizations having to hiring costly storage administrators to do so since the management of these features is integrated with vSphere and vCenter. VMWare integration will contribute to a shake-up as to which storage arrays achieve the top ranking in the upcoming Buyer's Guide.

- **Better accounting for scale-out architectures.** The full capabilities of midrange arrays from Dell (EqualLogic,) HP (P4000 LeftHand) and Infotrend (ESVA) were not fully reflected in the last Buyer's Guide. So even though individual models that are part of these product portfolios were covered in the Buyer's Guide, there was insufficient consideration and weighting given to how the models can function together and support a scale-out SAN architecture. That deficiency is addressed in this Buyer's Guide.

In addition to these changes in focus and how specific features will be emphasized (or de-emphasized) in the upcoming Midrange Array Buyer's Guide there are some other general trends that I have noticed as I have reviewed and tabulated the results.

- **High capacity 5.4K RPM SATA drives are supported by very few providers.** 18 months ago the adoption rate of these HDDs drives appeared to be on the uptick as they offered high capacity at an even lower price than comparably sized 7.2K RPM SATA drives. 18 months later, the availability of these drives as an option on storage arrays has almost disappeared as the interest by users in them seems to have waned with maybe 10% of storage providers offering these 5.4K RPM drives as an option on

their arrays.

- **SSDs have gone main stream.** SSDs were just starting to come into their own in the first half of 2010 as option. Now it can be safely said that the availability of SSDs as an option has gone main stream with over 75% of midrange array models now supporting them.
 - **Midrange array support of VMware vStorage APIs is NOT as common as one might think.** As mentioned earlier, the adoption rate of VMware is putting a new emphasis on midrange arrays to integrate with vSphere and vCenter to do everything from array management to managing snapshots on the array. However a surprising 41% of midrange array models offered no support for VAAI at all while only 25% supported all or a majority of the features that VAAI has to offer. (Please note these percentages may change by the time the Buyer's Guide is released in Q4 2011 as DCIG expects storage vendors to announce additional support for VMware vStorage APIs at the upcoming VMworld in late August 2011.)
-

Expanded Edition of the DCIG 2011 Small Enterprise Storage Array Buyer's Guide Now Available

Today DCIG, LLC, and [Foskett Services, LLC](#), are pleased to jointly announce the availability of an Expanded Edition of the DCIG 2011 Small Enterprise Storage Array Buyer's Guide that weights, scores and ranks over 35 small enterprise storage array models priced from \$5,000 – 30,000 from 19

different vendors.



According to IDC Storage Systems senior research analyst Amita Potnis, the [trend](#) toward purchasing lower-priced systems continued in 2010 as revenues for external storage systems in the lower-price bands (\$0-\$24,999) **increased by 21.1% year over year** as end users looked to get the most for their money. This IDC number reflects the growing impact that small enterprise storage arrays are having on organizations of all sizes and the increase in money that they are spending on this particular category of external storage.

Declining storage capacity cost and increased network bandwidth coupled with constantly improving performance are contributing to make these external storage systems both affordable and practical. Yet it is the introduction and rapid adoption of server virtualization that is now making external storage a necessity in even the smallest organizations.

Among the types of external storage systems available, small enterprise storage arrays are riding this wave of server virtualization adoption. While large enterprises led the initial surge of server virtualization adoption, analysts at Gartner [expect](#) that by 2012, **50% of all application workloads will run on virtual machines**, with **most of that growth coming from businesses with 100 to 999 employees**.

This is leading more of these size organizations to look for storage arrays that are affordable, easy to deploy and reliable, and perform well. The only problem is that neither the storage providers nor the storage arrays in this class are well known as the majority of the focus to date has been on

midrange and large enterprise storage arrays.

This **Expanded Edition of the DCIG 2011 Small Enterprise Storage Array Buyer's Guide** changes that. Organizations for the first time now have at their fingertips a Buyer's Guide that provides them with a comprehensive list of small enterprise storage arrays that are **priced between \$5,000 and \$30,000**. This Buyer's Guide can assist in making this all-important buying decision while removing much of the mystery around how small enterprise storage arrays are configured.

This Buyer's Guide accomplishes the following objectives

- **Provides an objective, third-party evaluation of small enterprise storage arrays** that evaluates, scores and ranks their features from an end user's viewpoint
- **Includes recommendations** on how to best use this Buyer's Guide
- **Scores and ranks the features on each small enterprise storage array** based upon the criteria that matter most to users so they can quickly know which small enterprise storage arrays are the most appropriate for them to use and under what conditions
- **Provides data sheets for 37 small enterprise storage arrays** from 19 different storage providers so users can do quick comparisons of the features that are supported and not supported on each small enterprise storage array
- **Provides insight into the features that a small enterprise storage array offers** that will result in improved performance without needing to run timely and costly internal performance benchmarks on multiple products
- **Provides insight** into which features on a small enterprise storage array will result in **increased power and space efficiency** so your costs stay down over time
- Gives any organization the ability to **request competitive bids from different providers of small enterprise storage arrays** while still having confidence

that they are getting a small enterprise storage array that meets their application needs

The criteria for a storage array's inclusion in this Buyer's Guide were as follows:

- It had to allow networked storage connectivity (FC, iSCSI, and/or NAS) to multiple servers
- It had to support at least eight (8) internal hard disk drives (HDDs)
- There had to be sufficient public available information to make a meaningful decision
- The model had to be shipping by March 1, 2011
- It had to be approximately priced between \$5,000 and \$30,000

Some of the key findings that came out of the research that went into preparing this Expanded Edition of the DCIG 2011 Small Enterprise Storage Array Buyer's Guide include:

- 86% of the arrays support both 1 Gb Ethernet and the iSCSI protocol
- 52% of the arrays support FC (4 Gb or higher)
- Only 38% of the arrays support 10 Gb Ethernet
- Only 24% of the arrays support NAS
- All arrays that support NAS also support iSCSI
- 67% offer support for a single storage controller
- 47% offer support for dual storage controllers
- 19% offer support for either single or dual storage controller configurations

The Expanded Edition of the DCIG 2011 Small Enterprise Storage Array Buyer's Guide Top Five solutions in the \$5,000 to \$15,000 model category include (in alphabetical order):

- [Aberdeen](#) LLC [AberNAS](#) 260/360 Series (#1 overall in this class)
- [D-Link](#) Corp [DSN-5110-10](#)

- [Dell](#) PowerVault [MD3200i/MD3220i](#)
- EMC [VNXe3100](#)
- [PROMISE](#) Technology, Inc. VessRAID [1000i Series](#)

The Expanded Edition of the DCIG 2011 Small Enterprise Storage Array Buyer's Guide Top Ten solutions in the \$15,000 to \$30,000 model category include (in alphabetical order):

- [Aberdeen](#) LLC [AberSAN](#) Z20/Z30/Z50 (#1 overall in this class)
- [Dot Hill](#) Systems Corp. [AssuredSAN 2700](#) series
- EMC [VNXe3300](#)
- HP StorageWorks [P2000 G3 10GbE iSCSI/FC](#)
- [StoneFly](#), Inc. [Voyager TSC16](#)
- [Winchester Systems](#) Inc. [VX-2300/3400](#) and SX-2300/3400 series

The Expanded Edition of the DCIG 2011 Small Enterprise Storage Array Buyer's Guide is **available immediately** for [free download](#) with registration via DCIG's website.

Time to Refresh the DCIG 2010 Midrange Array Buyer's Guide

About a year ago DCIG decided to do something completely different in the analyst space: a side-by-side independent comparison of products in a particular market segment in the form of a Buyer's Guide. The end result of that was the DCIG 2010 Midrange Array Buyer's Guide. But believe it or not, a year has already passed since that was produced and it is now time to update and refresh that Buyer's Guide for a number of reasons.

Despite the tremendous success of the first DCIG Midrange Array Buyer's Guide, the need to update that Buyer's Guide has become apparent over the past few months. Aside from the release of a number of new models from midrange array storage providers, *midrange arrays are rapidly changing to adapt to the influx of virtualization at all layers of the computing stack.*

Server virtualization is clearly getting the most attention with VMware becoming almost the default in terms of which technology organizations are choosing to virtualize their servers. But that is creating new requirements in storage arrays in terms of the features they need to offer in order to optimize VMware deployments in user environments.

However server virtualization is not the only form of virtualization gaining momentum. Virtualization technology on storage arrays is enabling storage arrays to more efficiently provision and manage storage while requiring less storage capacity in the process. This is leading to new technologies appearing in storage arrays such as automated storage tiering, deduplication, thin provisioning and virtual storage pools.

It is these reasons and others that have made it necessary to refresh the DCIG 2010 Midrange Array Buyer's Guide. Some of the key ways that the DCIG 2011 Midrange Array Buyer's Guide will change for 2010 include:

- More attention will be paid to storage array vendor certification with virtual operating systems such as Citrix, Microsoft and VMware as certification with these OSes is viewed as a prerequisite in many user environments
- Midrange arrays that accommodate scale-out SAN designs will be examined in more detail.
- The following midrange array features will be examined in more detail to include:

- *Automated storage tiering (AST)*
 - *Deduplication*
 - *Replication and snapshots*
 - *Thin provisioning*
 - *Storage networking protocols*
 - *Virtual storage pools*
 - *VMware vSphere integration*
-
- The approximate starting list price of each midrange array will be requested.

Surveys for the 2011 Midrange Array Buyer's Guide have already been sent out so if you are a midrange array provider and did not receive this survey, please let DCIG know and we will get one sent out to you.

If you are an end user and looking forward to this updated Buyer's Guide, look for it to be published and released this summer.

DCIG 2011 Small Enterprise Storage Array Buyer's Guide Now Available

Today DCIG and Foskett Services are pleased to jointly announce the availability of DCIG's latest Buyer's Guide, the 2011 Small Enterprise Storage Array Buyer's Guide. This Buyer's Guide is focused on storage arrays that are priced from approximately \$10,000 – \$30,000 and is specifically targeted at small enterprises and business units within large enterprises who are looking for the most value for their IT

dollar when spending on storage.



Small enterprise storage arrays have become a critical and vibrant part of the storage market though it is one that has been largely overlooked because they are perceived as having limited feature sets, storage capacity and storage network connectivity. As such, small and midsize organizations have for the most part been left to fend for themselves in terms of figuring out which of these storage arrays are the most appropriate fit for their environments.

However the advent of server virtualization and virtual server operating systems, such as VMware ESX and vSphere, has driven a massive spike in demand for networked storage with even the smallest IT shops possessing applications that require advanced features and high performance.

This demand is driving innovation in this class of storage arrays with advanced features being introduced on small enterprise storage arrays at a dizzying pace. It may even be safe to say that the features found on small enterprise storage arrays will eventually trickle up to high-end arrays!

In preparing this Buyer's Guide, DCIG and [Foskett Services](#) looked at over 50 features on 21 models from thirteen (13) different vendors who offered arrays that met our definition of "small enterprise storage array." The criteria for a storage array's inclusion in this Buyer's Guide were as follows:

- It had to allow networked storage connectivity (FC, iSCSI, and/or NAS) to multiple servers
- It had to support at least twelve (12) internal hard disk drives (HDDs)
- There had to be sufficient public available information to make a meaningful decision
- The model had to be shipping by December 1, 2010
- It had to be approximately priced between \$10,000 and \$30,000

Some of the key findings that came out of the research that went into preparing this 2011 Small Enterprise Storage Array Buyer's Guide include:

- 86% of the arrays support both 1 Gb Ethernet and the iSCSI protocol
- 52% of the arrays support FC (4 Gb or higher)
- Only 38% of the arrays support 10 Gb Ethernet
- Only 24% of the arrays support NAS
- All arrays that support NAS also support iSCSI
- 67% offer support for a single storage controller
- 47% offer support for dual storage controllers
- 19% offer support for either single or dual storage controller configurations

The following vendors have storage arrays that are included in this Buyer's Guide. Those vendors having an asterisk next to their name possess storage arrays that achieved a "Top 10" overall score in this Buyer's Guide:

- [Aberdeen](#), LLC (*)

- Arena-[Maxtronic](#), Inc
- [Dot Hill](#) Systems Corp. (*)
- [D-Link](#) Corp.
- [EMC](#) Corporation
- [Hewlett Packard](#)
- [iStoragePro](#) (*)
- [NetApp](#) (*)
- [Nexsan](#) Corp. (*)
- [Overland Storage](#) (*)
- [PROMISE](#) Technology, Inc.
- [StoneFly](#), Inc. (*)
- [Winchester Systems](#), Inc (*)

The raw data featured in this Buyer's Guide primarily came from two sources: an online survey completed by manufacturers, and publicly available information on the web. DCIG and Foskett Services attempted to validate all of the information contained in this Buyer's Guide with all manufacturers contacted after the research was done, in order to verify our findings.

This Buyer's Guide is immediately available from DCIG for \$995 and may be obtained by contacting DCIG's VP of Business Development, Jim Nash, at jim.nash@www.d cig.com or at [1.844.324.4552](tel:1.844.324.4552) (844.DCIGLLC).

Please note that DCIG hopes to have a licensing agreement worked out in the very near future where interested end-users can download this Buyer's Guide at no charge. Once that licensing agreement is in place, DCIG will provide the link to where it may be downloaded.

Update!

DCIG is now making the Expanded Edition of the DCIG 2011 Small Enterprise Storage Array Buyer's Guide available for free with registration. This Buyer's Guide is accessible at this [link](#).

HP Shares Details about the Future of 3PAR within its StorageWorks Division

Now that the acquisition of 3PAR by HP is a [done deal](#), there are three big questions on the minds of many. How will 3PAR's InServ Storage Servers fit into HP's overall storage portfolio? Is HP's relationship with HDS over? Does HP keep its EVA line of storage? These are some of the questions I was able to get answered this week when I met with Craig Nunes, the new HP Director of StorageWorks Marketing at Storage Networking World ([SNW](#)) 2010.

First, in regards to where does 3PAR fit into the HP StorageWorks portfolio, Nunes referenced an presentation made by HP's Executive Vice President and General Manager, Enterprise Storage, Servers and Networks, Dave Donatelli, provided. Donatelli views 3PAR as **THE** architecture for the next decade and delivering advanced features today that customers desire.

He positions 3PAR as addressing a variety of markets from mid to enterprise to the cloud. On that point, Nunes indicated ***the 3PAR family is largely complementary within the HP storage line-up in the context of addressing application workloads which break down into two broad classifications, predictable and unpredictable.***

Predictable workloads are those one might associate with traditional enterprise deployments like **SAP** or a **Microsoft Exchange** workgroup, whereby capacity growth and workload type are largely understood and resistant to large,

unforecast changes.

Unpredictable workloads are those most often found in virtualized environments where capacity demand and workload type can vary widely since content and applications are often coming from outside the enterprise (as is the case with social networking or cloud hosting) or within the enterprise in a more dynamic form (as is the case with large scale server virtualization deployments). In this case, **workload demands can peak and change at almost any time.**

Prior to HP's acquisition of 3PAR, the XP and EVA storage solutions that HP offered could best be described as a fit for predictable workloads across a broad range of applications and deployments. However with the acquisition of 3PAR and its high end T-Class and midrange F-Class models, HP now has an offering for the growing number of virtual environments with unpredictable workloads.

So in terms of where 3PAR fits in HP's current storage stack, it can probably best be summed up as follows:

[HP XP](#) (HDS)

3PAR [T-Class](#)

3PAR [F-Class](#)

HP [EVA](#)

HP [P4000](#) (formerly Lefthand Networks)

Seeing that lineup somewhat answers the next two questions as to what HP plans to do with its current XP and EVA lines of storage. The short answer is that for now all of HP's current storage offerings are still on the table.

While there is arguably some overlap between the HP XP and 3PAR T-Class, one area where the **HP XP still has a distinct advantage** over all of HP's other storage offerings, 3PAR or otherwise, is **its mainframe connectivity**. Further, to the best

of my knowledge, HP has no plans in the near term to invest in providing mainframe connectivity for the 3PAR T-Class though in the long term, who knows?

So in all likelihood, **HP's relationship with Hitachi Ltd will not go away** and the HP XP will continue to be a part of the HP StorageWorks portfolio for the foreseeable future. But my gut feeling is that HP will more aggressively push 3PAR storage in all of its enterprise accounts and only bring up the HP XP in accounts that need mainframe connectivity or where installed base preference exists for the XP.

A good measure of this will be next week on Wednesday when I attend the Q4 2010 VMware User Group ([VMUG](#)) in Omaha where HP's Technical Team will be in attendance and presenting at the event. I will be curious to see what storage HP will be pushing but my bet is already on 3PAR as I was just contacted yesterday by a former 3PAR sales rep (now HP) who will coincidentally be in town the same time as the event.

As to the future of HP EVA, long term I see the EVA continuing to provide fibre channel connectivity below the 3PAR F-Series and alongside the iSCSI-based HP P4000. The EVA installed base is massive and for those who appreciate the ease of use of the EVA and are deploying in more traditional, predictable workload environments, the EVA will continue to rule.

One other minor question that people have also been wondering is, *"Will the 3PAR name stick around?"* While no one really knows for sure, my guess is probably not. HP wants its brand on everything and while 3PAR was fairly well known in the enterprise storage space, outside of that space, not so much.

So my sense is that based upon what HP did with Lefthand Networks and re-branding it the "P-Series", I would expect sometime in the near future that the 3PAR brand will suffer the same fate as Lefthand Networks. However I can see the 3PAR "T-Class" and "F-Class" model designations for its InServ

Storage Servers persevering as those seem to fit within the HP storage branding philosophy.