

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.

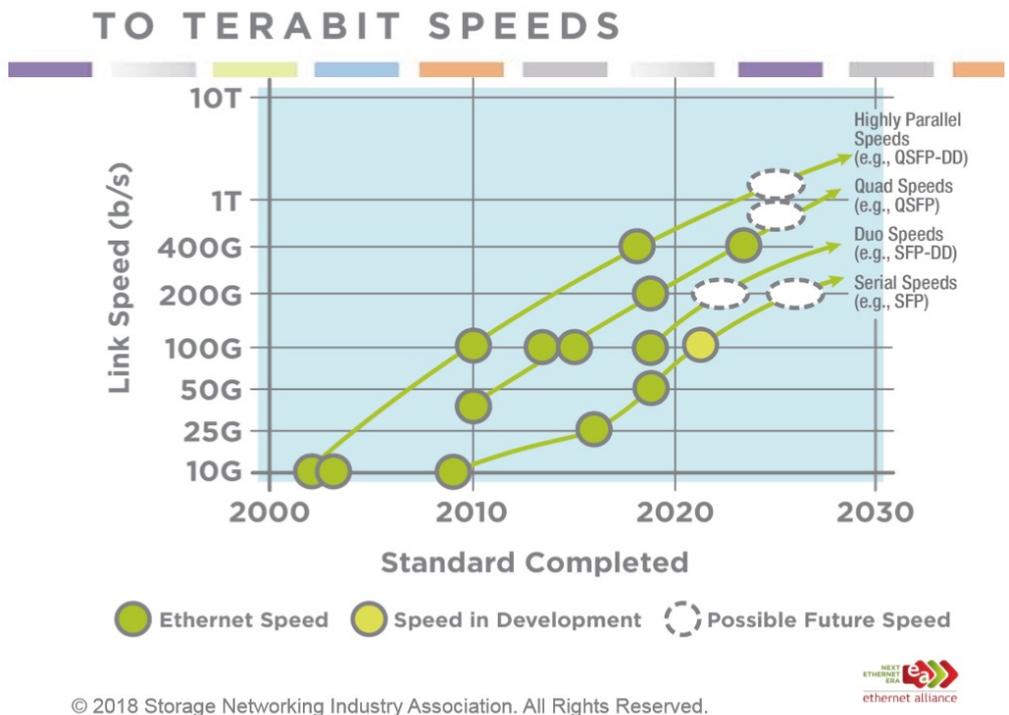
Editor's Note: This blog entry was updated on August 9, 2019, to correct grammatical mistakes and add some links.

Fast Network Connectivity Key to Unlocking All-flash Array Performance

The current generation of all-flash arrays offers enough performance to saturate the network connections between the arrays and application servers in the data center. In many scenarios, the key limiter to all-flash array performance is storage network bandwidth. Therefore, all-flash array vendors have been quick to adopt the latest advances in storage network connectivity.

Fast Networks are Here, and Faster

Networks are Coming



Ethernet is now available with connection speeds up to 400 Gb per second. Fibre Channel now reaches speeds up to 128 Gb per second. As discussed during a recent SNIA [presentation](#), the roadmaps for both technologies forecast another 2x to 4x increase in performance.

While the fastest connections are generally used to create a storage network fabric among data center switches, many all-flash arrays support fast storage network connectivity.

All-flash Arrays Embrace Fast Network Connectivity

[DCIG](#)'s research into all-flash arrays identified thirty-seven (37) models that support 32 Gb FC, seventeen (17) that support 100 Gb Ethernet, and ten (10) that support 100 Gb InfiniBand connectivity. These include products from Dell EMC, FUJITSU Storage, Hitachi Vantara, Huawei, Kaminario, NEC Storage, NetApp, Nimbus Data, Pure Storage and Storbyte.

Connectivity	AFA Models
32 Gb FC	37
100 Gb Ethernet	17
100 Gb InfiniBand	10

Source: DCIG

Other Drivers of Fast Network Connectivity

Although all-flash storage is a key driver behind fast network connectivity, there are also several other significant drivers. Each of these has implications for the optimal balance between compute, storage, network bandwidth, and the cost of creating and managing the infrastructure.

These other drivers of fast networking include:

- Faster servers that offer more capacity and performance density per rack unit
- Increasing volumes of data require increasing bandwidth
- Increasing east-west traffic between servers in the data center due to scale-out infrastructure and distributed cloud-native applications
- The growth of GPU-enabled AI and data mining
- Larger data centers, especially cloud and co-location facilities that may house tens of thousands of servers
- Fatter pipes yield more efficient fabrics with fewer switches and cables

Predominant All-Flash Array Connectivity Use Cases

How an all-flash array connects to the network is frequently based on the type of organization deploying the array. While there are certainly exceptions to the rule, the predominant

connection methods and use cases can be summarized as follows:

- Ethernet = Cloud and Service Provider data centers
- Fibre Channel = Enterprise data centers
- InfiniBand = HPC environments

Recent advances in network connectivity—and the adoption of these advances by all-flash array providers—creates new opportunities to increase the amount of work that can be accomplished by an all-flash array. Therefore, organizations intending to acquire all-flash storage should consider each product's embrace of fast network connectivity as an important part of the evaluation process.

Five Ways to Measure Simplicity on All-flash Arrays

Simplicity is one of those terms that I love to hate. On one hand, people generally want the products that they buy to be “*simple*” to deploy and manage so they can “*set them and forget them.*” The problem that emerges when doing product evaluations, especially when evaluating all-flash arrays (AFAs), is determining what features contribute to making AFAs simple to deploy and manage. The good news is that over the last few years five key features have emerged that organizations can use to measure the simplicity of an AFA to select the right one for their environment.

Simplicity is one of those attributes that everyone generally knows what it is when they see it. However, it can be challenging to quantify exactly what features contribute to making a product simple to deploy and manage. This difficulty stems from the fact that the definition as to what constitutes simplicity is subjective and varies from organization to organization and even from individual to individual.



Individuals and organizations may look at multiple features to ascertain the simplicity of a product. As they then apply their definitions and interpretation of simplicity to AFAs, arriving at a conclusion of what simplicity means and that everyone agrees upon can be problematic.

The good news is that as the adoption of AFAs has increased, the list of features that deliver on simplicity and which one should look for has coalesced to a list of features that you can get your arms around. These five features that contribute toward delivering on this ideal of simplicity on all-flash arrays are:

1. **All-inclusive software licensing.** Nothing is worse than trying to figure out how many or what type of software licenses you need on your AFA. Many AFAs now solve this dilemma by including software licenses for all the features on their array. While some still do tie licensing to storage capacity, number of hosts, processing power, or some mix thereof, the overhead and time associated with managing software licenses on each array should be much less than in the past.
2. **Evergreen.** The capital costs associated with hardware refreshes that occur every 3-5 years put a large hole in corporate budgets in the year that they hit. More AFAs now include “evergreen” options that, when purchased as part of their support contracts, refresh the existing hardware at its end of life, usually three years.
3. **Pre-built integration with automation frameworks.** As organizations look to automate the management of their IT infrastructure, AFAs are falling right in line. While using web-based GUIs to manage an AFA is handy, AFAs that can be discovered and managed as part of the organization’s broader automation framework make it more seamless for organizations to quickly roll new AFAs into their environment, discover them, and put them into productions.
4. **Proactive maintenance.** The last thing any IT manager wants to get is a notification in the middle of the night, while on vacation, on a weekend is that there is an application performance problem or a hardware failure. Many AFAs now proactively maintain their products using software that constantly optimizes performance or identifies and remediates hardware problems **before** they impact production applications. While IT managers still may be notified of these proactive activities performed by the AFA, the unpredictable, reactive nature of managing them is greatly reduced.
5. **Scale-out architectures.** Hardware upgrades and refreshes

as well as the data migrations that are often associated with performing those routine system admin activities have been a bugaboo for years in enterprise data centers. New scale-out architectures, sometimes referred to as web scale, now found on many AFAs mitigate if not put an end to the long hours and application disruptions that performing these activities have historically caused.

This list may not represent a comprehensive list of all the features that make an AFA simple to deploy and manage. However, this list does certainly represent the primary features that individuals and organizations should review to verify an AFA delivers on this attribute of simplicity to ensure an AFA's easy deployment and management in your environment.

These features and many others are what DCIG take into consideration as it prepares each of its [Buyer's Guides](#). Further, licensing the [DCIG Competitive Intelligence Portal](#), a SaaS offering from DCIG, includes DCIG research. You can then use this research as starting point to initiate and/or augment your own research. This Portal serves to centralize your internal competitive intelligence that can then be easily shared throughout your organization to whoever needs it wherever they need it. To learn more, click [here](#) to have someone from DCIG contact you.

**DCIG 2016-17 High End Storage
Array Buyers Guide Now**

Available

DCIG is pleased to announce the availability of the *DCIG 2016-17 High End Storage Array Buyer's Guide* developed from the enterprise storage array body of research. Other Buyer's Guide Editions based on this body of research will be published in the coming weeks and months, including the *2016-17 Midrange Unified Storage Array Buyer's Guide*.

The *DCIG 2016-17 High End Storage Array Buyer's Guide* weights, scores and ranks more than 100 features of fifteen (15) products from seven (7) different storage vendors. Using ranking categories of Best-in-Class, Recommended and Excellent, this Buyer's Guide offers much of the information an organization should need to make a highly informed decision as to which high end storage array will suit their needs.

Each array included in the *DCIG 2016-17 High End Storage Array Buyer's Guide* had to meet the following criteria:

- Be identified by the vendor as a high end storage array
- Support multiple controllers in an Active-Active configuration
- Be intended for the storage of production data (as opposed to archive or backup data)
- Provide synchronous replication for non-disruptive operations across two or more physical locations
- Have the ability to either scale-out or scale-up to at least 3 PB of raw capacity
- Provide sufficient information for DCIG to draw a meaningful conclusion
- Must be formally announced and/or generally available for purchase as of September 1, 2016

DCIG's succinct analysis provides insight into the state of the high end storage array marketplace. The Buyer's Guide identifies the specific benefits organizations can expect to achieve using a high end storage array and key features

organizations should be aware of as they evaluate products. It also provides brief observations about the distinctive features of each product. Ranking tables enable organizations to get an “at-a-glance” overview of the products; while DCIG’s standardized one-page data sheets facilitate side-by-side comparisons assisting organizations to quickly create a short list of products that may meet their requirements.

End users [registering](#) to access this report via the [DCIG Analysis Portal](#) also gain access to the DCIG Interactive Buyer’s Guide (IBG). The IBG enables organizations take the next step in the product selection process by generating custom reports, including comprehensive side-by-side feature comparisons of the products in which the organization is most interested.

Fibre Channel (FC) HBAs Will Not Be Embedded on Server Motherboards Anytime Soon; Interview with QLogic’s Vikram Karvat, Part 2

Ethernet adapters began migrating to LAN on motherboard solutions in the late 1990s. Yet this practice never took hold for other technologies like Fibre Channel. The Fibre Channel (FC) market even today, as Gen 6 (32Gb) is being introduced, is dominated by host bus adapters (HBAs). In this second installment in my interview with QLogic’s Vice President of Products, Marketing and Planning, Vikram Karvat, he explains

why 32Gb FC HBAs are still installed separately in servers, as well provides insight into what new features may be released in the Gen 7 FC protocol.



[QLogic 2700](#) HBA; Source: QLogic

Jerome: *Are the new QLogic 32Gb FC HBAs embedded in the server and/or storage array mother boards? If not, are there any plans to do so?*

Vikram: The HBAs being discussed here are pretty much entirely add in cards on the server side. There are no embedded FC HBAs on servers. As a result, the FC HBA port counts analysts report represent not only the ports that are shipped from the vendors, but ports that are actually being deployed for use on an annual basis. This represents as close to a natural demand in any market as you could hope to measure.

Jerome: *Why haven't FC HBAs gone to being embedded?*

Vikram: A network card typically goes embedded when it hits north of 50 percent connectivity. To get to north of 50 percent for FC, you would probably have to quintuple its volume. It's a different set of economics.

We previously talked a little bit about the increased use of FC on the all-flash array (AFA) side, but [QLogic](#) is also actually seeing an increase in use and deployment of FC SANs in emerging markets like China. FC SAN deployments in China grew by 15 percent last year. That is huge and the growth rate has been like that for probably the last two to three years. Obviously, in the early years, growing even faster than that, but from a relatively modest base.

But it's no longer just a modest base anymore. It's significant at a global scale in terms of how many SANs are being deployed. Again, not to the same scale as in North America. But nonetheless, it's measurable and is making an impact towards keeping the market relatively stable.

From a use case perspective, it's interesting because it's a market that tends not to want to spend money on something unless it's absolutely necessary. It's an indicator of the stability of the FC market. And FC remains the predominant storage interconnect for storage arrays, as well as, servers. There are areas of growth like AFAs and emerging markets. All in all, FC is not a bad story. FC offers the availability, reliability, security, and lossless fabric that enterprises want.

Further, there is a lot of discussion about Remote Direct Memory Access (RDMA) and storage options with very low CPU utilization, etc. But FC has always been a fully offloaded architecture with ultra-low CPU utilization (*in the single digits*) which is why it is used for Online Transaction Processing (OLTP) types of infrastructures and it has always been [zero copy](#) (*i.e. – does not require the CPU to perform the task of copying data from one memory area to another.*)

The notion that there are new storage networking implementations out there that are more efficient is potentially a bit of a fallacy there. FC, as an industry, has not made a big deal out of its strengths because the industry just assumed everybody understood these concepts. We are having to remind people now.

Jerome: *As FC is so mature and stable, what innovation is occurring?*

Vikram: There a number of areas of innovation where the industry is investing. Obviously [Gen 6](#) FC is good. Moving forward, the FC industry is actually in the process of

defining Gen 7 FC as a next step up. Layering on to that, we are innovating in the flash space with Fibre Channel over Non-Volatile Memory Express (FC-NVMe).

FC-NVMe is an industry initiative to directly map the NVMe drive over a fabric. Why you would ask? The normal reasons why you map something over a fabric is for the ability to share, create pools, provision, and manage storage more effectively when it is connected, as opposed to having islands of flash floating around in servers.

The unique thing about FC-NVMe is that instead of using the standard SCSI stack, it actually bypasses the SCSI stack and uses native NVMe semantics to reduce both the latency of the access and the CPU associated with the SCSI infrastructure on both the storage and server sides.

You are effectively taking a technology that was initially very focused on driving latency and performance within a server and extending it out of the box to get some of these additional benefits. We were recently demonstrated the ability to run FC-NVMe, as well as traditional FCP traffic, simultaneously, on existing fabrics.

When we talk about developing a new technology, it's like, *"Hey, here's my new thing. Oh by the way to get this, you have to go buy a whole bunch of new stuff."*

What QLogic is doing is layering this functionality onto the infrastructure that's already in place. It effectively comes for free.

We are pretty excited about that. We got a lot of interest from our OEM customers. I suspect that over the course of the next year, as this technology starts getting in front of end users via our OEM customers that our OEM customers will find it even more attractive. Again there's everything to gain and nothing to lose.

In [Part I](#) of this series, we took a look at why all-flash arrays are driving the need for 32Gb fibre channel.

In the third and final of this interview series, Vikram reveals what new FC HBA features service providers are most eager to see and use.

All-Flash Arrays Driving Need for 32Gb Fibre Channel; Interview with QLogic's Vikram Karvat, Part I

All-flash arrays, cloud computing, cloud storage, and converged and hyper-converged infrastructures may grab many of today's headlines. But the decades old Fibre Channel protocol is still a foundational technology present in many data centers with it holding steady in the U.S. and even gaining increased traction in countries such as China. In this first installment, QLogic's Vice President of Products, Marketing and Planning, Vikram Karvat, provides some background as to why fibre channel (FC) remains relevant and how all-flash arrays are one of the forces driving the need for 32Gb FC.



Jerome: Vik, thanks for taking time out of your schedule to share a bit about 32Gb fibre channel. Before we begin, for the benefit of DCIG's readers, can you share a bit about QLogic and what has been going on over there for the past few years?

Vikram: Thanks, Jerome. Many of your readers are probably familiar with [QLogic](#) from the fibre channel side as it has continued to be a preeminent player in that space. However, QLogic has had a few changes in the last few years.

Mostly notably, QLogic [acquired](#) Brocade's Fibre Channel HBA assets about two years ago. As a result of concluding that transaction in early 2014, QLogic was able to move that relationship to a new level in terms of technical cooperation, alignment on road maps and technologies, etc.

The other significant change was that QLogic [acquired](#) Broadcom's Ethernet controller assets. QLogic already had its own portfolio of Ethernet controllers with which it had been relatively successful on the host side, and very, very successful on the storage side; but the Broadcom assets brought a different level of scale to our overall Ethernet portfolio and immediately put QLogic in a very, very strong

number two position in Ethernet.

The net net is that today QLogic has the number one position in Fibre Channel and the number two position in the 10Gb Ethernet on the host/server side of the business. This is important because it allows QLogic to look at certain types of technology that would benefit from end-to-end integration. It also has some interesting benefits as QLogic moves forward.

***Jerome:** Tell me about 32Gb Fibre Channel. What is happening on that front?*

***Vikram:** The next instantiation of the Fibre Channel roadmap is Gen 6 (32Gb) Fibre Channel (FC) which QLogic is releasing today. A lot of people ask me, "Why do you need Gen 6 FC? Do we need more performance?"*

There is always some of that. You do need more performance to support today's latest technologies, such as multi-core processors and multichannel memory on servers, but then you also have the move towards non-volatile storage in servers, as well as in the storage arrays. Further, databases just keep getting bigger and bigger and the response time requirements for accessing these content repositories keeps getting shorter. Gen 6 FC performance advantages play directly into all of these demands from both a bandwidth and an IOPS perspective.

But there's more than just performance advantages with the shift to Gen 6 FC. IT organizations are under tremendous OPEX pressure. They need to maintain service-level agreements (SLAs) but with fewer people so they have to find ways to work more efficiently. Further, they are under pressure to increase scalability and deliver faster provisioning of new storage on demand.

This is where some of the features and functions that QLogic offers with its new Gen 6 FC adapters deliver as much value and, in some cases, maybe even more value than the performance

benefits of Gen 6 FC.

Jerome: Isn't QLogic introducing new technology and innovating in a market that is in decline?

Vikram: There has been a general sense in the industry that Fibre Channel is on a steep decline. I would propose to you today that that may not be entirely true. It's certainly not the growing market that it was a decade ago, but it's not ending any time soon.

The data points here just serve to underscore that. On the external block-based storage side, a lot of Fibre Channel connectivity has actually gone up in terms of a mix of total ports. Some of that is driven by the still significant need for Fibre Channel in traditional arrays.

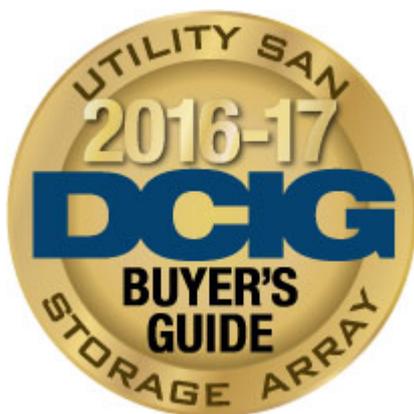
Some of this demand is also being driven by all-flash arrays. Almost 80 percent of these are connected via Fibre Channel. Then, if you look at Fibre Channel just in raw terms of how many Fibre Channel ports there are per unit of storage capacity, it's actually higher on all-flash arrays than it is on traditional storage arrays, just because of the performance levels associated with flash.

The result is that we've actually seen a slight uptick over the last three years in overall mix of Fibre Channel connectivity on external storage controllers. The actual number of port shipments has been holding steady for the last couple of years. We expect the same to hold true for 2016, with just slightly north of two million ports of server side HBA connectivity. Again, this might take some people by surprise because there's been the general sense that the market has been in decline, but the numbers actually show that from a standard HBA perspective, it's pretty stable.

In Part II of this interview series, Vikram shares his thoughts about industry initiatives to directly map the NVMe drive over Fibre Channel fabric.

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](#).

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

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

VVols Changes the Storage Management Conversation

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

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

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

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

and the overhead associated with these tasks onto the storage array.

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

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

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

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

HP 3PAR StoreServ: Locked and Loaded with VVols Support

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

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

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

- **Adaptive Flash Cache** that dedicates a portion of the HP 3PAR StoreServ's available SSD capacity to augment its available primary cache and then accelerates response times for applications with read-intensive I/O workloads.
- **Adaptive Optimization** that optimizes service levels by matching data with the most cost-efficient resource on the HP 3PAR StoreServ system to meet that application's service level agreement (SLA).
- **Priority Optimization** that identifies exactly what storage capacity is being utilized by each VM and then

places that data on the most appropriate storage tier according to each application's SLA so a minimum performance goal for each VM is assured and maintained.

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

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

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

organization's reach.

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

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.

DCIG 2014-15 High End Storage Array Buyers Guide Now Available

DCIG is pleased to announce the release of its **2014-15 High End Storage Array Buyer's Guide** that weights, scores and ranks more than 100 features of thirteen (13) different storage arrays from five different storage providers.



Due to the scalability and high availability criteria that were used to evaluate these high end storage arrays, the number of reviewed products is relatively small compared to other DCIG Buyer's Guides. However, this inaugural guide provides enterprises with a comprehensive list of high end storage arrays' supported features and functionality to assist them in this all-important buying decision.

High end storage arrays are especially well-suited for large enterprises because the arrays:

- Scale up storage capacity and performance through the addition of disks and/or nodes
- Provide high availability by implementing active-active controllers
- Support multiple ports and storage networking interfaces such as Ethernet, Fibre Channel (FC) and Fibre Connection (FICON)
- Provide a mature and feature-rich suite of data management services
- Support virtualization integration leveraging VMware's API's to offload some processing to the storage array
- Provide performance monitoring for the entire array and various components
- Provide storage efficiencies via automated storage tiering
- Support non-disruptive upgrades

The DCIG 2014-15 High End Storage Array Buyer's Guide Top 5 solutions include (in alphabetical order):

- HDS [VSP G1000](#)
- HP [XP7](#)
- HP 3PAR [StoreServ 10400](#)
- HP 3PAR [StoreServ 10800](#)
- NetApp [FAS8080 EX](#)

The **HP 3PAR StoreServ 10800** earned the "Best-in-Class" ranking among the high end storage arrays evaluated this year. In comparison to its counterparts, this array stood out in the following ways:

- Achieved the highest overall score
- Achieved the highest score in Software and VMware Integration categories
- Supports all VAAI 4.x and 5.x features earning a "Best-in-Class" ranking in VMware integration
- It represented the best balance of strengths across all the scoring categories

About the DCIG 2014-15 High End Storage Array Buyer's Guide

Selecting and comparing vendors and researching their products can be a daunting task. DCIG creates Buyer's Guides in order to help end users accelerate the product research and selection process—driving cost out of the research process while simultaneously increasing confidence in the results.

Organizations should therefore use this Buyer's Guide as a handbook to understand who the high end storage players are, what products they offer, what features and functions are available on each, how these solutions scale, what networking and storage protocols they offer, and how organizations might manage any solution they purchase.

The DCIG 2014-15 High End Storage Array Buyer's Guide achieves the following objectives:

- Provides an objective, third party evaluation of products that evaluates and scores their features *from an end user's perspective*
- Ranks each array in each scoring category and then presenting these results in an *easy to understand* table
- Provides a *standardized data sheet* for each of the arrays so users may do quick side-by-side comparisons of products
- Provides insights into the high availability and scalability of the arrays as well as what features the arrays offer
- Gives any organization a solid foundation for getting competitive bids from different providers that are based on “apples-to-apples” comparisons

The DCIG 2014-15 High End Storage Array Buyer's Guide is available immediately through the DCIG Analyst Portal for subscribing users by following this [link](#).

Six Observations about Today's High End Storage Arrays

Enterprises investing in today's high end storage arrays understand the value that these arrays offer in regard to their availability and performance as it can [cost](#) upwards of \$5,000 for every minute that an application is offline. Applications and data must be available all of the time as any interruption in service can seriously impact a corporation's revenue and reputation.

Selecting the appropriate high end storage array that offers the correct combination of features and functionality for an enterprise can mitigate the possibility of outages and the costs associated with them. This explains why high end storage arrays, even many years after their introduction, remain more than a popular choice to host today's centralized, virtualized applications. They are, in essence, experiencing a rebirth of sorts.

However, choosing any high end storage array requires a substantial investment in both time and money to research and implement. Further, there are notable differences between each array that DCIG classifies as "*high end*." This is why DCIG is producing a Buyer's Guide on High End Storage Arrays that it anticipates releasing in the very near future.

In that vein, as DCIG has done research on these arrays, it has made the following six observations about them and the environments into which they are going into.

- **High application availability.** Possibly the most

desirable feature on these arrays that prompts so many enterprises to deploy them is their high availability (HA). Yet what differentiates them is that vendors employ various methodologies to deliver HA with options to scale up, scale out or both to deliver HA. At the most fundamental level, these arrays support multiple pairs of Active-Active controllers (also implemented as “blade pairs” or “processor pairs” on some of these arrays) on the same physical array that are all part of the same logical array configuration.

- **Large hardware capacities.** Each high end array also has high capacities in regard to its cache, raw storage and processing. Over half of the arrays scale to support upwards of 3,000 GB (3 TBs) of cache, 67 percent of the arrays support at least 4,500TB (4.5 petabytes) of raw storage capacity and 60 percent of the arrays scale out to support at least 64 processor cores.
- **Multiple storage networking interfaces.** All of the storage arrays covered in the forthcoming Buyer’s Guide have a minimum of 20 storage networking ports available while 80% had up to 64 networking ports. The interface types vary by vendor and product, but all of the arrays support 8Gb Fibre Channel (FC), 75 percent support 16Gb FC, 87 percent offer 10Gb Ethernet and one-third of the arrays support 8Gb FICON (used in mainframe environments.)
- **Robust VMware integration.** Given VMware’s predominance in enterprise data centers today, it follows that they want storage arrays that can integrate with VMware vStorage APIs such as VAAI (vStorage APIs for Array Integration) and VASA (vSphere Storage APIs for Storage Awareness) to take advantage of the “force multiplication” that these APIs provide. The good news is that all of the high end storage arrays in this upcoming Buyer’s Guide support all of the VAAI v4 APIs and the majority of the options in VAAI v5. Similarly, all of the arrays support VASA.

- **OS and application performance monitoring.** Managing any enterprise data center is challenging, but managing one without visibility into each application's performance so that one can diagnose and even anticipate proactively anticipate problems can be foolhardy.

Using these performance monitoring and management tools, administrators can quickly pinpoint performance bottlenecks or what piece of hardware inside of the array is malfunctioning. All of the arrays provide some level of performance monitoring, with over 80 percent of the arrays providing physical drive monitoring while over 50 percent provide monitoring on a per application level. However, there is still some disparity in the ability of these arrays to monitor performance at the OS (Operating System) and VM (Virtual Machine) level so enterprises need to exercise some caution in which one they select as not all of these arrays may offer the full suite of software that they need to fully monitor and manage performance for all OSes and applications.

- **Automation.** Data center automation is another growing area of emphasis for many enterprises as it facilitates efficient management of their data center infrastructure and more agile responses 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.

Currently, 71% of the high end arrays support policy-based storage selection. However, only 28% expose their APIs for third-party automation tools, while 42% provide an SDK for integration with management platforms. As more enterprises place a premium on automating their storage environment, look for these numbers to increase.

10 Characteristics That Help to Define Today's High End Storage Arrays

It has been said that everyone knows what "*normal*" is but that it is often easier to define "*abnormal*" than it is to define "*normal*." To a certain degree that axiom also applies to defining "high end storage arrays." Everyone just seems to automatically assume that a certain set of storage arrays are in the "*high end*" category but when push comes to shove, people can be hard-pressed to provide a working definition as to what constitutes a high end storage array in today's crowded storage space.

Over the last few weeks the analysts at DCIG have certainly wrestled with some of those same issues regarding the definition of a high end storage array. Whereas the highest levels of availability, capacity and performance were once the defining attributes of these arrays, the providers of these arrays can no longer claim that they exclusively deliver these features. Many storage arrays classified as "*enterprise midrange*" or "*midrange*" offer similar or even higher levels of availability, capacity and performance than the storage arrays typically classified as "*high end*."

This is not to imply that a high end class of arrays does not exist. Such arrays do exist and it is important that organizations and enterprises recognize these arrays for what they are. However the features or characteristics that make them "*high end*" may, in some cases, differ from even a few years ago. To shed some light on what makes these storage arrays "*high end*," DCIG has come up with 10 characteristics

that organizations should look for to distinguish between an array that is “*high end*” and one that is “*midrange*.”

1. ***FICON connectivity to an IBM mainframe.*** In talking to a number of end users, VARs and vendors, FICON connectivity to IBM mainframes running z/OS is often where the difference between mainframe and midrange begins and ends. In short, if it does not offer FICION connectivity to a mainframe, it is not a high end storage array.
2. ***Fibre Channel (FC) block-based storage connectivity.*** Absent FICON connectivity, the storage array must minimally offer block-based FC connectivity to even have a shot at being considered a high end storage array. While a number of storage arrays considered high end may support Ethernet block-based protocols such as iSCSI or FCoE (Fibre Channel over Ethernet,) support for these protocols alone are not enough to bridge the midrange to high end gulf.
3. ***Multiple Active-Active controller/blade/processor pairs.*** A number of midrange arrays offer an “Active-Active” controller configuration where a pair of controllers permits concurrent access to data on the same backend disk. What differentiates a high end array from a midrange array is the availability of multiple pairs of these Active-Active controllers (also called “*blade pairs*” or “*processor pairs*” on some arrays) on the same physical array that are all part of the same logical array configuration.
4. ***High levels of cache and capacity.*** Despite the encroachment on this territory by multiple midrange arrays, high end storage arrays as a group still generally support far higher levels of cache and storage capacity than most midrange arrays. One should generally expect the amount of cache available on a high end storage array to scale into the hundreds if not thousands of GBs and provide support for PBs of storage

capacity.

5. ***Large number of multi-core processors.*** The multiple blade/controller/processor pairs in a high end storage array deliver much more than high availability. They also provide access to much higher levels of performance. This becomes critically important in environments that are handling mixed workloads that may include sequential reads, sequential writes and random access, small block transactions.
6. ***Scale-out and scale-up configurations.*** Midrange array providers often tout the scale-out or scale-up capabilities of their arrays like they are the best thing since sliced bread. High end storage providers tend to yawn, stretch and say, "*It is about time you offer those features on your array.*" In other words, scale-out and scale-up are part and parcel to the configuration of every high end storage array.
7. ***Detailed system analysis, performance monitoring and troubleshooting.*** High end storage arrays give organizations unparalleled flexibility to gather and analyst system data. This may then be used to quickly, accurately and confidently pinpoint where a performance bottleneck is occurring or what piece of hardware inside of the storage array is malfunctioning. Most midrange storage arrays do not offer this level of diagnostics or capabilities to troubleshoot a performance or system issue.
8. ***Tested, certified configurations.*** While midrange array also "certify" their arrays with certain OSes and applications, the certification process in my mind for midrange arrays has always been a little suspect. This concern stems from the large number of applications and operating systems for which midrange arrays must be certified and the diverse environments into which they are deployed. Due to the smaller number of application- and OS-specific environments into which high end storage arrays are deployed, the level of confidence that

enterprises may have about the quality and thoroughness of the interoperability testing and the quality of the features available can be higher.

9. ***Starting list price of \$250,000 or higher.*** All of these features, high levels of capacity and performance and certifications come at a price. While these high end storage arrays may actually be price competitive on a per GB basis with some midrange arrays, you first need an environment that justifies the scale that these high end arrays bring to the table.
10. ***Non-disruptive operations across two or more data centers.*** Many storage arrays offer one or more forms of replications. But what is arguably becoming a defining feature on high end arrays is their ability to deliver synchronous replication to at least two storage arrays and then sync the applications (think VMs) with the underlying replication activities so as to guarantee non-disruptive operation of applications. While this feature was initially designed to deliver disaster recovery, more enterprises are looking to leverage this capability for load balancing, non-disruptive failovers and failbacks and even to lower their data center operating costs.

The HP XP7 Storage Virtual Array Capability Marks the Beginning of the End of the

Pain of Data Consolidations and Migrations

Delivering always-on application availability accompanied by the highest levels of capacity, management and performance are the features that historically distinguish high end storage arrays from other storage arrays available on the market. But even these arrays struggle to *easily* deliver on a fundamental data center task: migrating data from one physical array to another. The [introduction](#) of the storage virtual array feature into the new HP XP7 dramatically eases this typically complex task as it facilitates data consolidations and migrations by migrating entire storage virtual arrays from one physical array frame to another while simplifying array management in the process.

Data Consolidations and Migrations Create High End Pain

Organizations with business and mission critical applications find high end storage arrays highly desirable for multiple reasons. They are highly available. They scale to hold up to petabytes of storage capacity. They deliver performance in the millions of IOs per second (IOPS.) They can handle mixed application workloads. Their operating systems are mature, stable and well documented. These represent the standards against which all other storage arrays are measured.

Despite these advantages, the pain of non-disruptively and seamlessly migrating data from one high end physical array frame to another persists. Like any other array, high end arrays still have capacity and performance limitations. Further, as their technology ages or warranties expire, their application data must be migrated to a new storage array. Here is where the challenges surface.

While all high end storage arrays provide software to

facilitate the migration of data from one array to another or the consolidation of data on a single array, these tasks are both complex and laborious. Planning and then executing upon them to avoid applications downtime and/or disruptions in performance may take weeks, months or even years to complete.

Organizations typically first document the placement of the application data on their **existing** high end storage array(s) before beginning any type of data consolidation or migration. Once documented, organizations must then determine where they want to place that data on the new array. At this point zoning and LUN masking on the new storage array is done so application servers may concurrently access capacity on both the old and new storage arrays. Only once those activities are complete may data on a LUN-by-LUN basis be migrated from an existing to a new array so the cutover to the new array may occur.

Even assuming all of these manual processes are accomplished flawlessly, there is still no guarantee the data consolidation or migration will go exactly as planned. Administrators over different applications need to learn to share array resources as well as schedule and resolve the change control requirements of their respective applications. Firmware on the servers' host bus adapters (HBAs) or converged network adapters (CNAs) may be out-of-date and not recognize the LUNs presented by the new storage array. The volume manager and/or operating system on the physical or virtual machines may experience similar issues. Should any of these challenges arise, organizations may need to fail back to the old array.

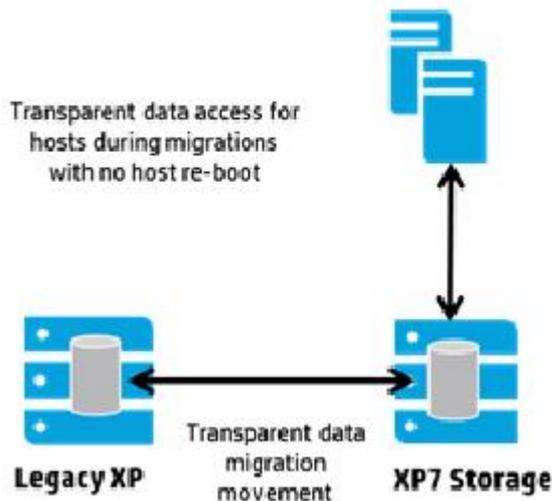
In a worst case scenario, a data consolidation migration only partially succeeds. Should this occur, both the old and new storage arrays must remain in use as some applications run on the new array while the rest remain on the older storage array. In this situation an organization may need to keep using the older storage array for an indeterminate amount of time until the data migration is complete.

The Storage Virtual Array Impact

The introduction of the storage virtual array capability into the Next Gen [HP XP7](#) removes these persisting complexities associated with data consolidations and migrations. To create a storage virtual array, organizations must first identify storage capacity resources such as hard disk drives (HDDs) and solid state drives (SSDs) within the frame of a physical HP XP arrays and then mark them for inclusion in a specific storage virtual array.

This feature reduces the current complexities and risks of migrating data as well as improves the manageability of the storage infrastructure in the following ways:

- ***Granular management through the creation of multiple storage virtual arrays.*** Organizations often consolidate the data of multiple applications and departments onto a single high end storage array to reduce costs and improve availability. The downside is that multiple individuals may need to access and manage the array. By creating up to eight (8) storage virtual arrays and placing each application's and/or department's data in its own one, administrators may then securely access and migrate only the data for which they are responsible.
- ***Simplified migrations by moving entire storage virtual arrays.*** Migrating LUNs from one physical XP array to another on a LUN-by-LUN basis is, at best, complex to setup and time-consuming to execute upon. Using the storage virtual array capability, organizations may migrate an entire virtual storage array from one physical XP array to another. Each storage virtual array has its own "*personality*" – array model, administrative privileges, LUN masking, etc. – so all of these characteristics are included with the storage virtual array as it is migrated. This reduces the setup time and simplifies the task of migration.



Source: HP

- **Reduced data migration risk through transparent data mobility.** Leveraging the HP XP7's existing data management and replication software, the storage virtual array may non-disruptively and transparently migrate a storage virtual array from one physical XP array to another. The physical and/or virtual hosts may then access the storage virtual array on the new XP array in the same way that they did on the old physical XP array once they are zoned to access the new XP array. Further, since the storage virtual array can continue to present to the hosts the same model number as the prior host, it reduces the chances of incompatibilities between the hosts' CNA, HBA and/or volume manager software and the storage virtual array residing on the new physical XP .
- **Access to additional resources.** Organizations invariably find themselves in a position where application servers need more storage capability, performance or both over time. The XP7 addresses both of these ongoing organizations requirements by offering up performance improvements of **up to 300 percent or more** versus the HP XP P9500. It also gives organizations the flexibility to put more HDDs and SSDs into an XP7 as well as a wider range of each media type.

- **Lays the groundwork for a seamless disaster recovery solution.** Most organizations envision a day where their applications and data are always available regardless of the circumstances. Storage virtual arrays that may be non-disruptively migrated across physical XP arrays bring that vision closer to a reality.

HP XP7 Storage Virtual Array Marks the Beginning of a New Reality without the Pain of Data Consolidations and Migrations

Organizations want the pain associated with data consolidations and migrations to end. The introduction of the storage virtual array capability into the Next Gen HP XP7 serves as a point of demarcation as to when companies can start to expect the pain associated with these tasks to stop. While organizations will need to utilize professional services to initially adopt and implement this technology on the HP XP7, once that investment is made, they can look forward to the storage virtual array feature facilitating the easy and secure sharing of XP resources while making data consolidations and migrations a much simpler task to plan and execute upon going forward.

**DCIG 2014-15 Flash Memory
Storage Array Buyer's Guide
Now Available**



DCIG is pleased to announce the March 30 release of the **DCIG 2014-15 Flash Memory Storage Array Buyer's Guide** that weights, scores and ranks more than 130 features of thirty-nine (39) different storage arrays from twenty (20) different storage providers.

Flash Memory Array Performance and Sales Rising Rapidly

Flash Memory Storage Arrays promise to deliver the dramatic performance benefits of flash memory including hundreds of thousands to millions of IOPS with sub-millisecond latencies while using as little as 1/10th the rack space, power and cooling of traditional enterprise storage arrays. The most recent generation of flash memory storage arrays generally deliver twice the IOPS of the prior generation and deliver a more complete set of features that enables them to address a broader set of use cases. Although enterprise storage professionals are traditionally cautious about adopting new technologies, many all-flash array vendors report that sales are growing in excess of 100% per year.

Flash Memory Storage Arrays Now Replacing Traditional Enterprise Arrays

Multiple vendors we spoke with indicated that prospective customers are now looking to do a complete tier 1 storage refresh; transitioning to an all-flash environment for their critical business applications. Reflecting this trend, International Data Corporation (IDC) [forecasts](#) "capacity

shipped in 2016 will increase to 611PB [petabyte] with a 2012–2016 CAGR [compound annual growth rate] of 110.8%”.

Driver #1: IT Budget Savings

For those with a responsibility for the technology budget, a *flash-storage-enabled rethinking of the data center can generally achieve hard cost savings of over 30% in data center hardware and software, and realize an ROI of less than 11 months.*² In some cases, the Flash Memory Storage Array may prove less expensive than just the maintenance cost of the former SAN.

Driver #2: Accelerating the Enterprise

Flash-based storage systems typically create a seven-fold improvement in application performance. This accelerated performance is enabling progress on initiatives that were hampered by storage that could not keep up with business requirements. Savvy business people are finding many ways to generate business returns that make the flash storage investment easy to justify. In one [case study](#) the installation of a flash memory storage system was directly attributed with avoiding the need to hire between 10 and 40 additional employees. Flash storage enabled the company to grow their business without growing their head count.

About the DCIG 2014-15 Flash Memory Storage Array Buyer's Guide

The plethora of vendors and products in the all-flash array marketplace—combined with a lack of readily available comparative data—can make product research and selection a daunting task. DCIG creates Buyer's Guides in order to help end users accelerate the product research and selection process—driving cost out of the research process while simultaneously increasing confidence in the results.

The DCIG 2014-15 Flash Memory Storage Array Buyer's Guide

achieves the following objectives:

- Provides an objective, third party evaluation of hybrid storage arrays that evaluates and scores their features from an end user's perspective
- Ranks each array in each scoring category and then presenting these results in an easy to understand table
- Provides a standardized data sheet for each of the arrays so users may do quick side-by-side comparisons of products
- Provides insights into what features the arrays offer to optimize integration into VMware environments, as well as support for other hypervisors and operating systems
- Provides insight into which features will result in improved performance
- Gives any organization a solid foundation for getting competitive bids from different providers that are based on "apples-to-apples" comparisons

The DCIG 2014-15 Flash Memory Storage Array Buyer's Guide Top 10 solutions include (in alphabetical order):

Hitachi Data Systems HUS VM
HP 3PAR StoreServ 7450 Storage
NetApp FAS3250 Series AFA
NetApp FAS6290 Series AFA
Nimbus Data Gemini F400
Nimbus Data Gemini F600
PureStorage FA-400 Series Controller
SolidFire SF6010
SolidFire SF9010
Tegile Zebi HA2800

The **Nimbus Data Gemini F600** earned the *Best-in-Class* ranking among all Flash Memory Storage Arrays in this buyer's guide. The [Gemini F600](#) multiprotocol unified all-flash storage system stood out in the following ways:

- Captured the highest score in the Management & Software as well as Hardware categories, and scored near the top in VMware Integration. In addition to excellent VMware integration, the Gemini F600 is one of just a handful of arrays supporting Microsoft System Center (SCVMM SMAPI), Microsoft Offloaded Data Transfer (ODX), and SMB 3.0.
- Raw flash storage density of 24 TB per rack unit, scalable to 1PB in a 42U cluster. This raw flash density was exceeded by only two other arrays in this guide.
- The already high flash storage density is further enhanced through comprehensive flash optimization support, including lossless in-line deduplication and compression.
- One of the few arrays supporting 56 Gb FDR InfiniBand and 40 Gigabit Ethernet , the fastest and lowest latency interfaces for host connectivity.

[Nimbus Data](#) claims that the Gemini F600 can deliver up to 1 million 4K write IOPS and 2 million 4K read IOPS with latencies of approximately 50 microseconds; at a 35% lower cost compared to the prior generation. The Gemini F600 starts under US\$80,000, and typically sells for between \$150,000 and \$250,000 as configured by customers.

Flash Memory Storage Arrays Are a Systemic Business Opportunity

The purchase of a Flash Memory Storage Array will be most easily justified and have the greatest benefit in businesses that think this through as a systemic opportunity. Many who do so will discover that “Flash is free.” That is, the return on investment within the IT budget will be rapid, and the business benefits of accelerating all enterprise applications could truly present an opportunity to accelerate the enterprise.

The DCIG 2014-15 Flash Memory Storage Array Buyer’s Guide is immediately available through the DCIG analyst portal for

subscribing users by following this [link](#).

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.

Insights and Observations about the Forthcoming DCIG 2014 Enterprise Midrange Array Buyer's Guide

Anytime DCIG prepares a Buyer's Guide – whether a net new Buyer's Guide or a refresh of an existing Buyer's Guide – it always uncovers a number of interesting trends and developments about that technology. Therefore it is no surprise (*at least to us anyway*) that as DCIG prepares to release its DCIG 2014 Enterprise Midrange Array Buyer's Guide that it observed a number of interesting data points about enterprise midrange arrays. As DCIG looks forward to releasing this Buyer's Guide, we wanted to share some of these observations and insights that we gained as we prepared this Guide as well as why we reached some of the conclusions that we did.

Value of Included Software

Vendors that sell their midrange arrays with all software features fully licensed as part of the standard array package create extra value for purchasers by reducing the number of decision points in the purchasing process and by smoothing the path to full utilization of the array's capabilities.

Separate license fees for features can reduce the agility of the IT department in responding to changing business requirements because the purchase approval and ordering process may take several weeks. If implementation services are required, that may add additional weeks to the process.

Separate licensing fees may be minor, or they can have a noticeable impact on the overall cost of ownership for an enterprise midrange array. Therefore, the annual cost of software licenses and associated support contracts should be incorporated into TCO (total cost of ownership) and ROI (return on investment) calculations.

This forthcoming Buyer's Guide acknowledges the value of included licenses by awarding a significant number of points to those arrays that ship with features already licensed. In particular, DCIG give attention to licensing for snapshots, replication and thin provisioning features.

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 forthcoming 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 vSphere 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 vSphere 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 in this forthcoming Buyer's Guide support all of the VAAI 4.1 APIs. However, only 10% of the arrays 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 DCIG expects 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 forthcoming 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.

**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.

Fibre Channel and Ethernet

“Waist Deep in Politics”

As convergence and SDE (*software-defined-everything*) make their way into the mainstream and add real value, organizations both large and small battle with the question of, “*What should we do about our storage networks?*” Stick with Fibre-Channel based approach or, as depreciation cycles end and/or new data-center locations come online, refresh to an Ethernet only solution?

This is not as simple as deciding what technology to go with. There are many factors involved in making this decision which include:

- Acquisition cost
- On-going cost
- Pool of talented technicians
- Complexity
- The internal politics between the LAN/WAN and SAN teams

As vendors and consultants continue to battle it out (“*Fibre Channel rocks Ethernet*” and “*Ethernet Rules the World,*”) end-users are stuck in the middle trying to make a very difficult and complex decision with the end result being one that will have a significant impact to the performance and availability of there applications (*which at the end of the day is all that really matters.*)

There are many commonalities between the requirements for Fibre Channel and Ethernet based storage networks. An enormous one that is often overlooked is that the cable plant required for each one is basically identical.

Generally speaking, OM4 -50um fibre is used for both 8/16 Bb Fibre Channel and 10/40/100 Gb Ethernet in short haul and/or campus based environments. The expense of these cable plants can be quite extensive but when considering Fibre Channel or Ethernet, the cable plant cost for each will be basically a

wash.

With this harmony come a few disparities which need to be taken into account when making this decision. Some of the glaring ones include:

Speeds and Feeds

When looking at speeds and feeds Fibre Channel has 4/8/16 Gb (32 Gb standard no gear yet – Targeted 2015 Deployment) and Ethernet has 10/40/100 Gb although 40/100 Gb are usually dedicated to internal switch-to-switch interconnects. With the pervasiveness of Flash/DRAM based storage and alike attaining these speeds in the real world are not as far out as they may have once seemed.

Design Requirements

When designing a Fibre Channel or Ethernet storage network it is important to take the same methodical approach to both. The design concepts around physical switch separation, dual HBAs-TOEs-CNAs, zoning, and vlans should be applied in either scenario. This will ensure some of the unfounded reliability claims of Ethernet from years past are just that past. This also helps when constructing the financial model debunking the myth that Fibre Channel is so much more expensive. (Like for Like)

With this in mind, making the determination which technology to choose should be boiled down to an apples-to-apples comparison. As I have traveled down this path in my day job, I made every attempt to take emotions and politics out of the equation (A very difficult task indeed) and focus on the cost, resources and technological benefits of both solutions.

The simplest way to begin is to construct a matrix of features and a TCO (Total Cost of Ownership) / ROI (Return on

Investment). This helps greatly reduce the amount of politics in the equation and allows you make the right decision for your specific environment based on your business and application requirements.

In summary, I'd say don't come down in either camp (*Why fight the battle?*) Use the matrix and financial models described in this post to take the emotional and political in-fighting out of the equation and rely on the metrics and like for like comparison.

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](#)

- DCIG 2013 Midrange Unified Storage Array – [Full](#)