All-inclusive Software Licensing: Best Feature Ever … with Caveats

On the surface, all-inclusive software licensing sounds great. You get all the software features that the product offers at no additional charge. You can use them – or not use them – at your discretion. It simplifies product purchases and ongoing licensing.

But what if you opt not to use all the product's features or only need a small subset of them? In those circumstances, you need to take a hard look at any product that offers all-inclusive software licensing to determine if it will deliver the value that you expect.

Why We Like All-Inclusive Software Licensing

All-inclusive software licensing has taken off in recent years with more enterprise data storage and data protection products than ever delivering their software licensing in this manner. Further, this trend shows no signs of abating for the following reasons:

- It makes lives easier for the procurement since they do not have manage and negotiate software licensing separately.
- It makes lives easier for the IT staff who want to use its features only to find out they cannot use them because they do not have a license to use them.
- It helps the vendors because their customers use their features. The more they use and like the features, the more apt they are to keep using the product long term.
- It provides insurance for the companies involved that if

they do unexpectedly need a feature, they do not have to go back to the proverbial well and ask for more money to license it.

- It helps IT be more responsive to changes in business requirements. Business need can change unexpectedly. It happens where IT is assured that a certain feature will never be of interest to the end user. Suddenly, this "never gonna need it" becomes a "gotta have it" requirement.

All-inclusive software licensing solves these dilemmas and others.

The Best Feature Ever … Has Some Caveats

The reasons as to why companies may consider all-inclusive software licensing the best feature ever are largely self-evident. But there are some caveats as to why companies should minimally examine all-inclusive software licensing before they select any product that supports it.

1. ***Verify you will use the features offered by the platform.*** It is great that a storage platform offers deduplication, compression, thin provisioning, snapshots, replication, metro clusters, etc., etc. at no

extra charge. But if you do not use these features now and have no plans to use them, guess what? You are still going to indirectly pay for them if you buy the product.

2. ***Verify the provider measures and knows which of its features are used.*** When you buy all-inclusive software licensing, you generally expect the vendor to support it and continue to develop it. But how does the vendor know which of its features are being used, when they are being used, and for what purposes? It makes no sense for the provider to staff its support lines with experts in replication or continue developing its replication features if no one uses it. Be sure you select a product that regularly monitors and reports back to the providers which of its features are used, how they are used and actively supports and develops them.
3. ***Match your requirements to the features available on the product.*** It still pays to do your homework. Know your requirements and then evaluate products with all-inclusive software licensing based upon them.
4. ***Verify the software works well in your environment.*** I have run across a few providers who led the way in providing all-inclusive software licensing. Yet the ones who selected the product based on this offering found out the features were not as robust as they anticipated or were so difficult to use that they had to abandon using them. In short, having a license to use software that does not work in your environment does not help anyone.
5. ***Try to quantify if other companies use the specific software features.*** Ideally, you want to know that others like you use the feature in production. This can help you avoid become an unsuspecting beta-tester for that feature.

Be Grateful but Wary

I, for one, am grateful that providers have come around with

more of them making all-inclusive software licensing available as a licensing option for their products. But the software features that vendors include with their all-inclusive software licensing vary from product to product. They also differ in their maturity, robustness, and fullness of support.

It behooves everyone to hop on the all-inclusive software licensing bandwagon. But as you do, verify to which train you hitched your wagon and that it will take you to where you want to go.

Predictive Analytics in Enterprise Storage: More Than Just Highfalutin Mumbo Jumbo

Enterprise storage startups are pushing the storage industry forward faster and in directions it may never have gone without them. It is because of these startups that flash memory is now the preferred place to store critical enterprise data. Startups also advanced the customer-friendly all-inclusive approach to software licensing, evergreen hardware refreshes, and pay-as-you-grow utility pricing. These startup-inspired changes delight customers, who are rewarding the startups with large follow-on purchases and Net Promoter Scores (NPS) previously unseen in this industry. Yet the greatest contribution startups may make to the enterprise storage industry is applying predictive analytics to storage.

The Benefits of Predictive Analytics for

Gilbert advises Anne to stop using “highfalutin mumbo jumbo” in her writing. (Note 1)

The end goal of predictive analytics for the more visionary startups goes beyond eliminating downtime. Their goal is to enable data center infrastructures to autonomously optimize themselves for application availability, performance and total cost of ownership based on the customer’s priorities.

The vendors that commit to this path and execute better than their competitors are creating value for their customers. They are also enabling their own organizations to scale up revenues without scaling out staff. Vendors that succeed in applying predictive analytics to storage today also position themselves to win tomorrow in the era of software-defined data centers (SDDC) built on top of composable infrastructures.

To some people this may sound like a bunch of “highfalutin mumbo jumbo”, but vendors are making real progress in applying

predictive analytics to enterprise storage and other elements of the technical infrastructure. These vendors and their customers are achieving meaningful benefits including:

- Measurably reducing downtime
- Avoiding preventable downtime
- Optimizing application performance
- Significantly reducing operational expenses
- Improving NPS

HPE Quantifies the Benefits of InfoSight Predictive Analytics

Incumbent technology vendors are responding to this pressure from startups in a variety of ways. HPE purchased Nimble Storage, the prime mover in this space, and plans to extend the benefits of Nimble's InfoSight predictive analytics to its other enterprise infrastructure products. HPE claims its Nimble Storage array customers are seeing the following benefits from InfoSight:

- 99.9999% of measured availability across its installed base
- 86% of problems are predicted and automatically resolved before customers even realize there is an issue
- 85% less time spent managing and resolving storage-related problems
- 79% savings in operational expense (OpEx)
- 54% of issues pinpointed are not storage, identified through InfoSight cross-stack analytics
- 42 minutes: the average level three engineer time required to resolve an issue
- 100% of issues go directly to level three support engineers, no time wasted working through level one and level two engineers

The Current State of Affairs in Predictive Analytics

HPE is certainly not alone on this journey. In fact, vendors are claiming some use of predictive analytics for more than half of the all-flash arrays DCIG researched.

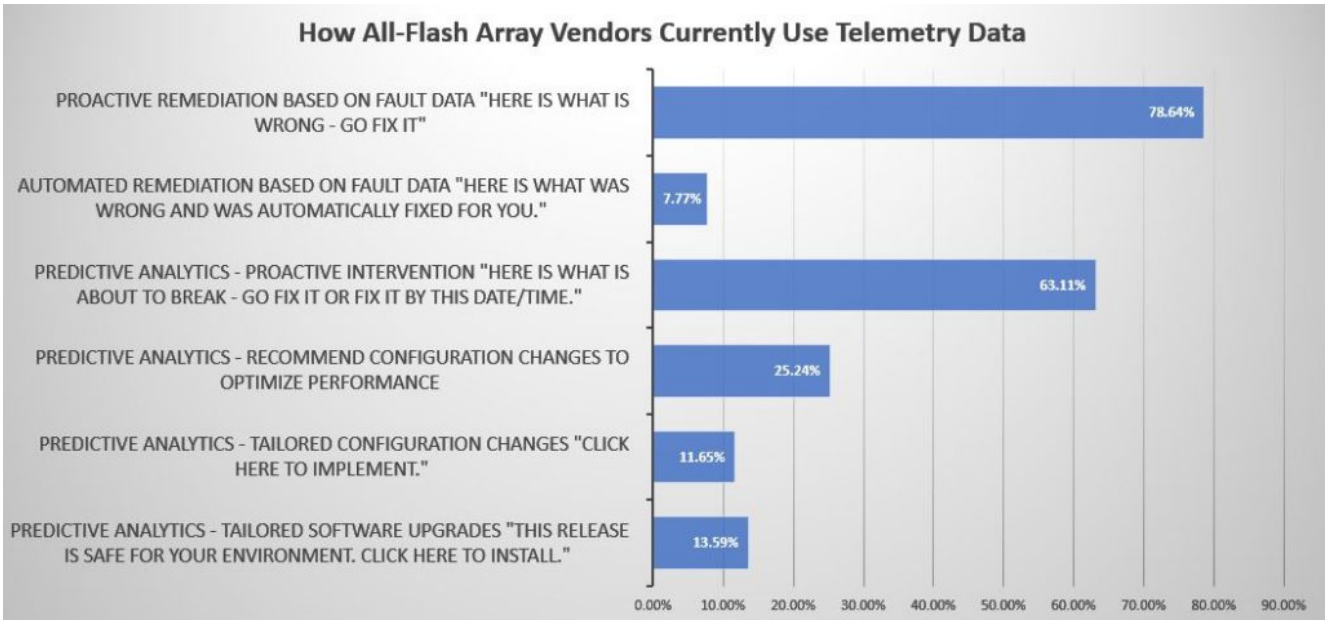

Source: DCIG; N = 103

Telemetry Data is the Foundation for Predictive Analytics

Storage array vendors use telemetry data collected from the installed product base in a variety of ways. Most vendors evaluate fault data and advise customers how to resolve problems, or they remotely log in and resolve problems for their customers.

Many all-flash arrays transmit not just fault data, but extensive additional telemetry data about workloads back to the vendors. This data includes IOPS, bandwidth, and latency associated with workloads, front end ports, storage pools and more. Some vendors apply predictive analytics and machine learning algorithms to data collected across the entire

installed base to identify potential problems and optimization opportunities for each array in the installed base.

Predictive Analytics Features that Matter

Proactive interventions identify something that is going to create a problem and then notify clients about the issue. Interventions may consist of providing guidance in how to avoid the problem or implementing the solution for the client. A wide range of interventions are possible including identifying the date when an array will reach full capacity or identifying a network configuration that could create a loop condition.

Recommending configuration changes enhances application performance at a site by comparing the performance of the same application at similar sites, discovering optimal configurations, and recommending configuration changes at each site.

Tailored configuration changes prevent outages or application performance issues based on the vendor seeing and fixing problems caused by misconfigurations. The vendor deploys the fix to other sites that run the same applications, eliminating potential problems. The vendor goes beyond recommending changes by packaging the changes into an installation script that the customer can run, or by implementing the recommended changes on the customer's behalf.

Tailored software upgrades eliminate outages based on the vendor seeing and fixing incompatibilities they discover between a software update and specific data center environments. These vendors use analytics to identify similar sites and avoid making the software update available to those other sites until they have resolved the incompatibilities. Consequently, site administrators are only presented with software updates that are believed to be safe for their environment.

Predictive Analytics is a Significant Yet Largely Untapped Opportunity

Vendors are already creating much value by applying predictive analytics to enterprise storage. Yet no vendor or product comes close to delivering all the value that is possible. A huge opportunity remains, especially considering the trends toward software-defined data centers and composable infrastructures. Reflecting for even a few minutes on the substantial benefits that predictive analytics is already delivering should prompt every prospective all-flash array purchaser to incorporate predictive analytics capabilities into their evaluation of these products and the vendors that provide them.

Note 1: Image source: https://jamesmacmillan.wordpress.com/2012/04/02/highfalutin-mumbo-jumbo/

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 0S 3.2.1 MU2 P12 or later software.

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

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

- ***Adaptive Flash Cache*** that dedicates a portion of the HP 3PAR StoreServ's available SSD capacity to augment its available primary cache and then accelerates response times for applications with read-intensive I/O workloads.
- ***Adaptive Optimization*** that optimizes service levels by matching data with the most cost-efficient resource on the HP 3PAR StoreServ system to meet that application's service level agreement (SLA).
- ***Priority Optimization*** that identifies exactly what storage capacity is being utilized by each VM and then places that data on the most appropriate storage tier according to each application's SLA so a minimum performance goal for each VM is assured and maintained.
- ***Thin Deduplication*** that first assigns a unique hash to each incoming write I/O. It then leverages HP 3PAR's Thin Provisioning metadata lookup table to quickly do hash comparisons, identify duplicate data and, when matches are found, to deduplicate like data.
- ***Thin Provisioning*** that only allocates very small chunks of capacity (16 KB) when writes actually occur.
- ***Thin Persistence*** that reclaims allocated but unused capacity on virtual volumes without manual intervention or VM timeouts.
- ***Virtual Copy*** that can create up to 2,048 point-in-time snapshots of each virtual volume with up to 256 of them being available for read-write access.
- ***Virtual Domains,*** also known as virtual private arrays, offer secure multi-tenancy for different applications and/or user groups. Each Virtual Domain may then be assigned its own service level.
- ***Zero Detect*** that is used when migrating volumes from other storage arrays to HP 3PAR arrays. The Zero Detect technology identifies "zeroes" on existing volumes which represent allocated but unused space on those volumes. As HP 3PAR migrates these external volumes to HP 3PAR volumes, the zeroes are identified but not migrated so

the space may be reclaimed on the new HP 3PAR volume.

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

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

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

HP 3PAR StoreServ Management Console Answers the Call for Centralized, Simplified Storage Operations Management without Technical Compromise

Scalable. Reliable. Robust. Well performing. Tightly integrated with hypervisors such as Microsoft Windows and

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

HP 3PAR's Legacy of Autonomic Storage Management

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

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

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

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

HP SSMC Answers the Call for Centralized Storage Management

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

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

Top Level and System Views

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

In the system view, administrators may select individual HP 3PAR StoreServ systems and view information specific to it. For instance, they may view: the available capacity of each storage tier type to include block and file storage management; the features licensed on that system; and, the system's resource utilization. By understanding how many or how few resources are available on each system, administrators may better determine where to place new applications and their data to align each application's needs with the StoreServ's available resources and features.

Centralizing management of all HP 3PAR StoreServ arrays under

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

End-to-End Mapping

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

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

The HP 3PAR SSMC provides this end-to-end mapping of the underlying infrastructure that is critical to maintaining application availability and ensuring optimal application performance. By identifying and visualizing the exact physical components used by each physical or virtual machine, enterprises can better understand the impact of system

component upgrades or outages as well as identify, isolate and troubleshoot performance issues before they have influence an application.

Capacity and Performance Reporting

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

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

Simplified Ongoing Management

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

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

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

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

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

Three Specific Use Cases for the Successful Implementation of Software-defined Storage

The introduction of first generation software-defined storage solutions (*often implemented as appliance and storage controller-based virtualization*) went terribly awry when they were originally introduced years ago for reasons that the industry probably only now fully understands and can articulate well. While the value of software-defined storage has never been disputed, best practices associated with its implementation, management and support short and long term

took time to develop. We are now seeing the fruits of these efforts as evidenced by some of the successful ways in which software-defined storage solutions are packaged and shipped.

The impact that software-defined storage solutions are poised to have on the traditional storage market is significant. Recent IDC research suggests that traditional stand-alone hybrid systems (*mix of disk and flash*) are expected to decline at a 13 percent compound annual rate while new system (*all-flash, hyperconverged and software-defined*) adoption will grow at 22 percent clip from 2014 to 2018.

The exact percentage that software-defined storage solutions will contribute to this overall 22% growth rate is unclear. However it is clear that doubts about their short and long term viability have largely evaporated.

Contributing to this increased confidence in using software-defined storage is the growing number of successful implementations of this technology on appliances and storage controllers. While software-defined storage has had a presence on these devices for well over a decade, the increased availability of software-defined storage solutions from vendors and growing adoption by end-users stems from the ability to better mitigate the issues associated with the use of software-defined storage and improved best practices for its initial implementation and ongoing management that optimize its inherent strengths.

Specific use cases where DCIG is aware of software-defined storage (SDS) solutions being successfully implemented and used on appliance and storage controller-based devices include:

- ***Non-disruptive (or near non-disruptive) data migrations.*** This is historically where appliance and storage controller-based SDS solutions have been used successfully for years. By inserting the appliance or

storage controller SDS solution into an existing storage network between the server and back end storage, the SDS solution is then used to virtualize the storage volumes on both existing and new storage arrays and then migrate the data from the existing array to a new storage array.

The appeal of using this approach was that the appliance or storage controller could be inserted non-disruptively or nearly non-disruptively (*application downtime of only seconds or minutes*) into the environment. Data may then be migrated from one storage array to another while the application continues to operate unaware that a data migration is occurring.

The HP 3PAR StoreServ storage arrays with their SDS solution now provide such an option. When migrating from an existing HP 3PAR, EMC VNX or EMC VMAX array to a new HP 3PAR StoreServ array, organizations may deploy the new HP 3PAR StoreServ, virtualize the volumes on the existing storage arrays, non-disruptively migrate the data to the storage on the new HP 3PAR StoreServ array and then cut the application(s) over to the new HP 3PAR StoreServ array with minimal to no application downtime.

- ***Better managing deployments of utility storage.*** Many if not most organizations have a growing need for deployments of large amounts of utility storage in their environments. Organizations increasingly have vast amounts of data for which they cannot quantify its value but know that it is sufficiently valuable that they cannot easily or justifiably delete it. In these cases they often want to use storage arrays that are reliable, stable, economical (*e. – provide storage capacity at well under $1/GB*,) perform moderately well and remain easy to manage and scale.

The storage upon which this data resides needs relatively few bells and whistles. In other words, it typically does

not need integration with any VMware APIs, will not host any Oracle databases, does not need any flash nor will it need any special automated storage tiering features. In short, the storage array deployed needs to be cheap and deep.

SDS solutions play nicely in these environments. Whether the SDS software resides on a storage controller (such as on a Dell EqualLogic, EMC Isilon, ExaBlox OneBlox or HP P4000 array) or on an appliance (DataCore SANSymphony, FalconStor FreeStor or IBM SVC), more storage capacity can be quickly and easily added to these environments and then just as easily managed and scaled since many of the interoperability and performance issues that have hindered SDS deployments in the past do not really come into play in these situations.

- ***Heterogeneous vendor multi-tiered storage environments.*** One of the big issues with appliance and storage controller-based SDS solutions is that they attempted to do it all by virtualizing every vendors' storage arrays. But by attempting to do it all, they often failed to deliver on one of the biggest benefit that SDS has to offer – creating a single pane of glass to manage all of the storage capacity and provide a common, standardized set of storage management features. Virtualizing all storage from all vendors made it too complicated to implement all of the features associated with each of the underlying arrays that were virtualized.

IBM with its SAN Volume Controller (SVC) has smartly avoided this pitfall. Rather than trying to virtualize every vendor's storage arrays and deliver all of their respective capabilities, its primary focus is to virtualize the various IBM storage arrays and deliver their respective capabilities. While organizations arguably sacrifice some choice and flexibility to buy from any storage vendor, many would rather have less choice with a more predictable environment than more choices with more risk. Further, IBM provides organizations with a sufficient number of storage

array options (flash, hybrid, disk, etc.) that they get most if not all of the tiers of disk that they will need, the flexibility to manage all of this storage capacity centrally and the ability to present a common set of storage array features to all attached applications.

Software-defined storage may not yet be fully mature but neither is it a half-baked or poorly thought out solution anymore. Vendors have largely figured out how to best implement it so they can take advantage of its strengths while mitigating its risks and have developed best practices to do so. Ultimately, this developing and maturing set of best practices will probably contribute more to SDS's long term success than any other new features that SDS solutions may offer now or in the future.

Three Specific Use Cases for the Successful Implementation of Software-defined Storage

The introduction of first generation software-defined storage solutions (*often implemented as appliance and storage controller-based virtualization*) went terribly awry when they were originally introduced years ago for reasons that the industry probably only now fully understands and can articulate well. While the value of software-defined storage has never been disputed, best practices associated with its implementation, management and support short and long term took time to develop. We are now seeing the fruits of these efforts as evidenced by some of the successful ways in which software-defined storage solutions are packaged and shipped.

The impact that software-defined storage solutions are poised to have on the traditional storage market is significant. Recent IDC research suggests that traditional stand-alone hybrid systems (*mix of disk and flash*) are expected to decline at a 13 percent compound annual rate while new system (*all-flash, hyperconverged and software-defined*) adoption will grow at 22 percent clip from 2014 to 2018.

The exact percentage that software-defined storage solutions will contribute to this overall 22% growth rate is unclear. However it is clear that doubts about their short and long term viability have largely evaporated.

Contributing to this increased confidence in using software-defined storage is the growing number of successful implementations of this technology on appliances and storage controllers. While software-defined storage has had a presence on these devices for well over a decade, the increased availability of software-defined storage solutions from vendors and growing adoption by end-users stems from the ability to better mitigate the issues associated with the use of software-defined storage and improved best practices for its initial implementation and ongoing management that optimize its inherent strengths.

Specific use cases where DCIG is aware of software-defined storage (SDS) solutions being successfully implemented and used on appliance and storage controller-based devices include:

- ***Non-disruptive (or near non-disruptive) data migrations.*** This is historically where appliance and storage controller-based SDS solutions have been used successfully for years. By inserting the appliance or storage controller SDS solution into an existing storage network between the server and back end storage, the SDS solution is then used to virtualize the storage volumes on both existing and new storage arrays and then migrate

the data from the existing array to a new storage array.

The appeal of using this approach was that the appliance or storage controller could be inserted non-disruptively or nearly non-disruptively (*application downtime of only seconds or minutes*) into the environment. Data may then be migrated from one storage array to another while the application continues to operate unaware that a data migration is occurring.

The HP 3PAR StoreServ storage arrays with their SDS solution now provide such an option. When migrating from an existing HP 3PAR, EMC VNX or EMC VMAX array to a new HP 3PAR StoreServ array, organizations may deploy the new HP 3PAR StoreServ, virtualize the volumes on the existing storage arrays, non-disruptively migrate the data to the storage on the new HP 3PAR StoreServ array and then cut the application(s) over to the new HP 3PAR StoreServ array with minimal to no application downtime.

- ***Better managing deployments of utility storage.*** Many if not most organizations have a growing need for deployments of large amounts of utility storage in their environments. Organizations increasingly have vast amounts of data for which they cannot quantify its value but know that it is sufficiently valuable that they cannot easily or justifiably delete it. In these cases they often want to use storage arrays that are reliable, stable, economical (*e. – provide storage capacity at well under $1/GB*,) perform moderately well and remain easy to manage and scale.

The storage upon which this data resides needs relatively few bells and whistles. In other words, it typically does not need integration with any VMware APIs, will not host any Oracle databases, does not need any flash nor will it need any special automated storage tiering features. In short, the storage array deployed needs to be cheap and deep.

SDS solutions play nicely in these environments. Whether the SDS software resides on a storage controller (such as on a Dell EqualLogic, EMC Isilon, ExaBlox OneBlox or HP P4000 array) or on an appliance (DataCore SANSymphony, FalconStor FreeStor or IBM SVC), more storage capacity can be quickly and easily added to these environments and then just as easily managed and scaled since many of the interoperability and performance issues that have hindered SDS deployments in the past do not really come into play in these situations.

- ***Heterogeneous vendor multi-tiered storage environments.*** One of the big issues with appliance and storage controller-based SDS solutions is that they attempted to do it all by virtualizing every vendors' storage arrays. But by attempting to do it all, they often failed to deliver on one of the biggest benefit that SDS has to offer – creating a single pane of glass to manage all of the storage capacity and provide a common, standardized set of storage management features. Virtualizing all storage from all vendors made it too complicated to implement all of the features associated with each of the underlying arrays that were virtualized.

IBM with its SAN Volume Controller (SVC) has smartly avoided this pitfall. Rather than trying to virtualize every vendor's storage arrays and deliver all of their respective capabilities, its primary focus is to virtualize the various IBM storage arrays and deliver their respective capabilities. While organizations arguably sacrifice some choice and flexibility to buy from any storage vendor, many would rather have less choice with a more predictable environment than more choices with more risk. Further, IBM provides organizations with a sufficient number of storage array options (flash, hybrid, disk, etc.) that they get most if not all of the tiers of disk that they will need, the flexibility to manage all of this storage capacity centrally and the ability to present a common set of storage array

features to all attached applications.

Software-defined storage may not yet be fully mature but neither is it a half-baked or poorly thought out solution anymore. Vendors have largely figured out how to best implement it so they can take advantage of its strengths while mitigating its risks and have developed best practices to do so. Ultimately, this developing and maturing set of best practices will probably contribute more to SDS's long term success than any other new features that SDS solutions may offer now or in the future.

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.

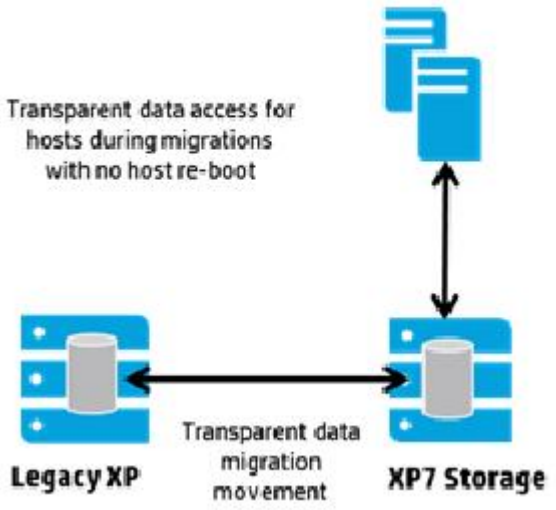

Source: HP

- ***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.

Heterogeneous Environments Still Alive and Kicking Despite Decline of Independent Third Party Storage Conferences

As I was planning my 2014 calendar over the past two weeks, I noticed that two storage conferences that focused on heterogeneous computing environments and were popular during from 2000-2010 have either gone the way of the dodo bird or are only a shell of what they formerly were. Yet during that same period of time, I met with some storage engineers and architects in the Omaha area who were telling me their environments are more heterogeneous than possibly ever before. While these trends on the surface may seem contradictory, they underscore the growing frustration that management in companies have with IT in general and how they are desperately looking for IT solutions that just work.

In the decade ranging from about 2000-2010, the two “can’t miss” conferences in the data storage world were Storage Decisions and Storage Networking World. Data storage was

undergoing a huge transformation from being direct-attached to network-attached and these two conferences were at the center of the vortex. Anyone who was anyone in the storage industry – analyst, vendor or end-user – was at these events as they showcased the best of what traditional players had to offer as well as many of the emerging technologies that were promising to re-shape the storage market.

Having attended many of these conferences, these are where I first saw many technologies such as backup appliances, deduplication, public storage clouds, scale-out storage, storage virtualization, storage resource management, thin provisioning and virtual server backup just to name a few. Each of these promised-and largely delivered-on solving key pain points that users were experiencing.

Yet these conferences fell short over time in an important aspect leading to their demise. They brought together competing vendors and put them in one place so users could view their wares, evaluate their products and bring them in-house to test and/or implement them. These conferences ultimately failed to transform themselves from solving specific customer pain paints to hosting vendors that offered holistic, macro-management solutions that could manage all of the product-specific solutions they had acquired over the years.

This was brought clearly into focus for me over lunch a couple of weeks ago that I had with a storage architect and a storage engineer. These two individuals are part of a global storage team that is responsible for managing all of the point solutions from various vendors brought in over the past decade. While it is a truly a heterogeneous environment, they find it very complex to manage, skill sets acquired in managing one technology do not easily transfer to managing other similar technologies, and vendor support for managing this heterogeneous environment is sketchy at best. Adding to their frustration, they have to support this environment while

trying to support the latest management initiative that is going to fix all of these issues (aka – the cloud.)

This leads us to why organizations have largely shifted away from attending conferences sponsored by independent third parties such as TechTarget and IDC to vendor-sponsored events. Vendors like Dell, EMC, HP, IBM, Microsoft, Symantec and VMware now host their conferences that attract thousands if not tens of thousands of users in large part because they are feeding on this end-user belief that if they adopt their cloud solution, they can easily and effectively manage this heterogeneous cludge created by buying decisions from 2000-2010.

While the approach varies slightly by provider, the general theme is this. Buy all new stuff from us that now all magically works together. Pay us a bunch of money for services to migrate data off of your old IT gear onto this new gear. Stand back and get ready to enjoy all of the benefits of our cloud once all of your data is hosted on our gear. This may sound a bit simplistic but this seems to be the common theme in every pitch I hear. It is also why organizations, hoping against hope that what these vendors are saying is true, are attending these vendor-sponsored conferences in growing numbers.

My thoughts are these. First, the cloud solutions these providers are promising will fix all of your existing problems probably will not – at least not all of them. They certainly may solve a subset of the problems, but they will likely only contribute to making your existing heterogeneous environment even more heterogeneous – if that is possible – simply because there are always too many legacy products with their own proprietary protocols or requirements to be stand-alone that organizations will never be able to virtualize away by putting them into a cloud.

Second, this infatuation with vendor-sponsored conferences is

likely just a near-term trend that has not fully run its course. At some point attendees are going to realize that no set of solutions from one provider is ever going to fully solve of their problems despite what vendors may promise. As this realization sinks in (and it may take a few years,) users will again start to seek out conferences that offer holistic solutions that have matured to the point they can manage products from multiple other vendors.

Third, such vendors still do exist and are even thriving despite the homogeneous, converged infrastructure mindset in which many organizations find themselves. Even as I write this, I am sitting in Colorado Springs, CO, attending a STORServer conference which recently inked a deal with CommVault so it could deliver a CommVault-powered backup appliance that is better suited to protecting today's mobile, enterprise IT environments.

Organizations are understandably frustrated by the lack of interoperability and inability to manage the heterogeneous assortment of solutions they purchased over the years. However they need to be wary about falling into the trap of today's vendor-sponsored conferences and the idea that homogeneous solutions are going to solve all of the problems. While they might, I would not bet the farm on it. Rather, I am more inclined to believe that heterogeneous IT environments are going to be alive and kicking for many years yet to come and that the sooner organizations recognize that and find (or even build) a solution to manage them, the sooner they will be happier with how their IT environment operates.

The Era of Application-Storage Convergence Has Arrived

Converged infrastructures are emerging as the next “Big Thing” in enterprise datacenters with servers, storage and networking delivered as a single SKU. Yet what providers are beginning to recognize – and what organizations should begin to expect – is that unprecedented jumps in application performance and resource optimization are now possible. The first examples of these jumps are seen in today’s ZS3 Storage Systems announcement from Oracle as it raises the bar in terms of how Oracle Database performance and resource utilization can be delivered while ushering in a new era of application-storage convergence.

Precipitated in large part by server virtualization, a tidal wave of change is sweeping through data centers. Data centers large and small recognize that endless racks of server, storage and networking are yesterday’s news as these configurations leave large amounts of capacity and performance sitting idle even as they waste power and generate heat.

In their place, smaller, more efficient and more powerful hardware solutions have emerged and are being deployed that today’s server virtualization platforms can better utilize. This trend has already led to a number of providers such as Nutanix, Oracle and Scale Computing to offer converged hardware solutions. These solutions bundle server, storage and networking hardware together as a single SKU to accelerate and simplify their purchase and subsequent deployment and management in virtualized environments.

As this occurs, providers are recognizing that delivering all of the hardware components (*servers, storage and networking*)

also gives them new levels of control over and insight into the hardware stack. This opens up unprecedented opportunities for them to accelerate application performance and optimize hardware resource utilization.

Application providers have historically had little or no visibility into the hardware stack in modern data centers as there was little they could do to accelerate performance or optimize resource utilization even if they did have this insight. However the introduction of technologies such as flash memory, large amounts of DRAM and multiple tiers of storage inside of storage arrays makes the ability of applications to communicate with storage arrays a much more interesting proposition as placing the right data on the right tier of storage at the right time can accelerate application performance even as it improves storage utilization.

This proposition becomes even more compelling when one considers the sizes of today's database that can easily reach into the tens if not hundreds of terabytes (TBs). Optimizing the placement of data within these databases can both accelerate their performance and reduce storage costs. This moves the concept of application-storage convergence from being "*highly desirable*" to a "*must-have*" for those organizations that expect to take full advantage of the performance and storage efficiency benefits that current storage arrays now deliver.

Yet what has stood in the way of such application-storage convergence to date is that both the application and the storage are still largely oblivious of one another. The application usually has some knowledge of the tier of disk (flash, FC or SATA) on which its data resides and the network connection (CFS, FC, iSCSI, NFS) it uses to access it and the storage array generally knows what operating systems and applications are accessing it.

Unfortunately that information does little to nothing to

improve application performance. Data optimization techniques employed by the storage array are "*best guesses*" and largely reactive while the application is left without any means to instruct the storage array where to place the data and when to do so.

The Oracle ZS3 Storage Systems change the paradigm of database-blind storage . With ZS3 and Oracle Database 12c, organizations can achieve unprecedented levels of Oracle Database efficiency, performance and storage utilization on hybrid ZS3 Storage System arrays. Here's how Oracle does it.

First, Oracle databases are now storage-aware. They can send IO metadata directly to the ZS3 Storage Systems using the new Oracle Intelligent Storage Protocol (OISP) that tells these arrays how the data is structured, what data is coming, what data it needs next and even what is optimal block size to use when sending data to and from the storage array.

On the storage side, it recognizes and responds to these instructions from Oracle applications and database to pre-stage frequently accessed data in cache or flash or, conversely, allows large sequential writes to pass directly through to disk. OISP is only available and supported by Oracle ZFS which gives Oracle ZS3 Storage Systems a significant advantage over all other storage arrays in delivering the highest levels of Oracle Database performance.

Further adding to the appeal of this technology particularly notable is that only Oracle storage enables organizations to utilize the full potential of Oracle Database 12c as it automates database-to-storage tuning and data compression. This eliminates the needs for organizations to manually tune database as the storage systems and database work in conjunction with one another to automatically execute these tasks.

Organizations can also benefit from the combination of

Automatic Data Optimization (ADO) and Hybrid Columnar Compression (HCC). As DCIG highlighted in a previous blog entry, this automates data compression at various levels based on data heat maps and usage patterns in such a way that it automatically recognizes which data is infrequently accessed content for deep archives to recent data being left uncompressed.

Second, to deliver the high levels of performance that Oracle databases frequently demand, the Oracle ZS3 Storage Systems use a highly threaded symmetric multiprocessing (SMP) architecture. The benefit this provides over many other storage system architectures is that it can handle multiple concurrent streams of I/O without overwhelming the storage system's processor.

Historically storage system processors have largely been underutilized as organizations were unable to send enough I/O to the storage system to seriously tax its CPUs. However, virtualized workloads have increased both the number of applications and the amount of throughput that applications send to storage systems. Then toss into the mix the high performance requirements of Oracle Databases and you suddenly have storage systems that are running out of CPU to process incoming data.

The SMP architecture used on the ZS3 Storage Systems eliminates this bottleneck in the storage stack allowing it to achieve ***450,000 IOPs in 700 microseconds*** according to the latest SPECsfs results. According to Oracle, this is lower latency than EMC's recent VNX-2 announcement at less than a quarter of its price.

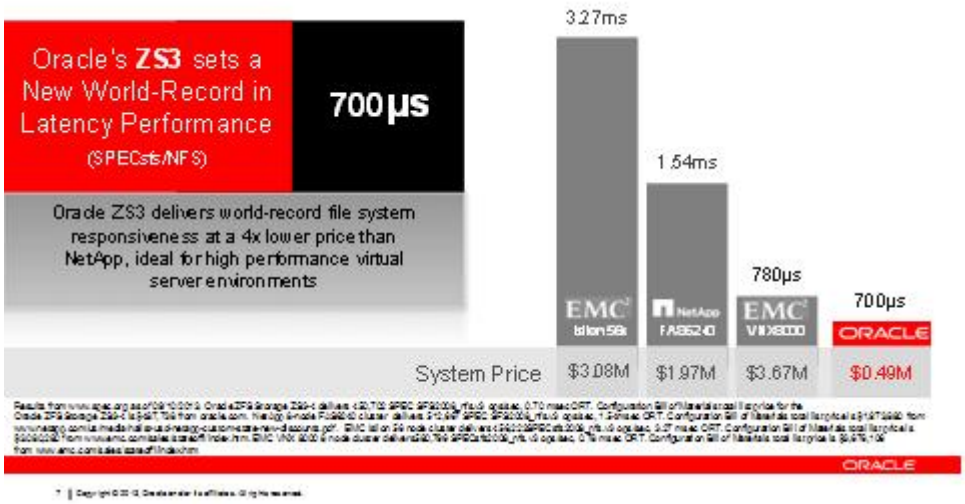

Source: Oracle

In other internal benchmarked results, it can handle the application processing on over 2000 VMs with its CPU running at only 33% utilization while the CPUs on other NAS filers max out at 100% supporting the workload of only 250 VMs. This is where the SMP operating system of ZS3 excels as it enables the system to scale linearly in VMware environm
ents, which are also a heavy SMP workload.

Server virtualization has already led to one massive wave of changes: organizations of all sizes consolidating and converging their infrastructures. Now another wave of change is about to be unleashed. As organizations converge their infrastructures it opens up new ways for applications to capitalize on the varying sets of features found within these converged infrastructures to both improve application performance and resource utilization.

The only way organizations will ever be able to fully utilize all of the features found in these newly converged and highly virtualized infrastructures is if the management of them from the application level down to the spindle is fully automated. The Oracle ZS3 Storage Systems in conjunction with Oracle Database 12c reflects the first of this next generation of converged application-to-storage solutions and, in so doing, foreshadows the transformation in data center management that is coming to enterprises in the not too distant future.

“Doc Lock-in” May not be the Villain at EMC World

The main theme at this year’s EMC World is “*Lead the Transformation*” that EMC is illustrating through the use of superhero characters. The superheroes are represented as end users who come up with solutions to manage today’s complex storage environment while the villain is pictured as “*Doc Lock-in*” who requires our superheroes to “*lock-in*” on a single vendor to mitigate this complexity. Yet for those users who think strategically about their storage acquisitions, Doc Lock-in may not be the full-fledged villain that EMC World portrays him to be.

EMC World 2013 has already provided some interesting insight into the underlying psyche that must currently prevail within the EMC culture. As I mentioned in yesterday’s blog entry, one of EMC’s objectives at EMC World was to once again make the topics of storage virtualization and, to a certain degree, storage resource management (SRM), palatable to end-users. To accomplish that, it coined the term “*Software-defined Storage*” to promote these two concepts without bringing along the emotional baggage that users often associate with them.

But by creating the Doc Lock-in character, EMC also inadvertently tipped its hand that it is sometimes (maybe often?) viewed by end-users as fulfilling the Doc Lock-in role. This likely motivated it to introduce ViPR, its software-defined storage solution. Through it, users can theoretically manage any vendor’s storage and EMC may cast itself on the side of the hero as opposed to playing the role of the villain.

But is "*lock-in*" really bad and should EMC be classified as "*evil*"? A group of us analysts were discussing this "*lock-in is evil*" mindset over dinner last night and we quickly came to the following conclusion: Vendor lock-in is a choice, not a decision any vendor requires any company to make.

Granted, some companies fail to fully comprehend the gravity of choices made and their consequences so they end up in a predicament where they feel "*locked-in*" to a particular vendor's solution. Yet I do not blame that on the vendor – that is on the end-user.

Further, even when they are in state where they feel "*locked-in*" to a particular vendor's solution, ***one is never truly "locked-in" as they are not out of choices***. ***Rather they are simply out of pleasant choices*** with "*pleasant choices*" being defined as *"the freedom to pick any solution at a discount* without having to spend weeks or months implementing it ."

Even then one should not be so quick to dismiss the benefits of being "*locked-in*." It takes a lot of end-user time – and I mean a lot of time – to investigate different vendor solutions, examine their benefits, choose one, test it and then bring it in-house and implement it. The amount of time it takes to accomplish this is what prompted DCIG to develop and release its Buyer's Guides to help buyers quickly assess what products are available in a particular space and identify the most appropriate one for them.

Conversely, consider the benefits, if you will, of being "*locked-in*" to a particular vendor. Yes, you may pay more for the technology but all of the time spent evaluating different solutions goes away. Instead, you can more quickly and simply pick that vendor's technology that matches the requirements of the job, implement it and then focus on leveraging that technology to address the issue.

Further, this mindset is certainly in line with what I hear

from end-users. They want simpler environments. They want to better leverage and optimize the technology they already have. They want one throat to choke. You probably only get that by having "*Doc Lock-in*" on your side of the field. Realize that if you do, you are probably going to pay up to have him there with the upside being that you will have more time to focus on doing what your business does best.

Even take EMC which is promoting no vendor lock-in with ViPR. C'mon. Really? If EMC was really promoting no vendor lock-in, then its keynote speakers should be onstage encouraging users to go out and buy as much HDS, IBM, NetApp, Dell, HP, Imation/Nexsan or whatever other brand of storage as they want and then use ViPR to manage it. I did not hear that message at all. EMC wants you to buy more VMAX, VNX, Isilon, Atmos and Data Domain and then use ViPR to manage it – pure and simple.

Locking in on a specific vendor's technology is a choice and it may be the only choice that some users feel they can make. But lock-in does have its upsides and users should not be so quick to dismiss them. Granted, they are times when Doc Lock-in may feel like and even be the villain. But having too many choices can have just as many downsides as feeling like Doc Lock-in is your only choice.

2013: ARM, SSD and Common Slot Servers

Bad news is only bad until you hear it, and then it's just information followed by opportunity. Information may arrive in political, personal, technological and economic forms. It creates opportunity which brings people, vision, ideas and

investment together. When thinking about a future history of 2013 opportunities, three (3) come to mind:

- Solid state storage
- 64bit ARM servers
- Common slot architecture for servers

While two of these are not new by themselves, an amalgamated version of them is a recipe for necessity. The most novel of the three is common slot architecture for servers. Common slot architecture allows an Intel or AMD x86, and Samsung or Calxeda CPU to be plugged into the same board and system. But, let's start by looking at solid state storage impact on storage architecture. It can eliminate or mitigate at least four (4) storage architecture constraints:

1. **IOPS** – Inputs/Outputs per Second
2. **Latency** – The time between when the workload generator makes an IO request and when it receives notification of the request's completion.
3. **Cooling** – The design and implementation of racks and HVAC for data centers
4. **Ampre** – The design and implementation of electrical systems for data centers and cities

While some may disagree with the assertion, a majority will agree that solid state storage modules or disks (SSD) are fast, much faster than their hard disk drive (HDD) brethren. In less than two years the measurement for IOPS has increased from a few hundred thousand to over one (1) million as expressed in "Performance of Off-the-Shelf Storage Systems is Leaving Their Enterprise Counterparts in the Dust." Thus it can be assumed the median IOPS requirement is some where between a few hundred thousand to one (1) million.

In that regard, it's fair to say that most applications and systems would perform quite well with the median output of a solid state storage system. Thus, when implementing an all

solid state storage system the median IOPS requirement can be met – CHECK.

Secondary to IOPS is Latency. Latency is a commonly overlooked requirement when gauging the adaptability of a storage system to an application. While defined above, Latency is referred to as "overall response time (ORT)" as commented by Don Capps Chair at SpecSFS. In 2012 Mr. Capps wrote to DCIG suggesting this format when sharing SpecSFS results ""XXX SPECsfs2008_cifs ops per second with an **overall response time** of YYY ms."

ORT and IOPS are not on par with each other. In that regard, a high IOPS number doesn't result in a lower latency. For example, *Alacritech, Inc. ANX 1500-20 has 120,954 SpecSFS 2008_nfs ops per second with an overall response time of 0.92ms*, whereas the *Avere Systems, Inc. FXT 3500 (44 Node Cluster) has 1,564,404 SPECsfs2008_nfs ops per second with an overall response time of 0.99 ms*. In both cases the ORT is under 1 ms and meets the latency requirements for the broadest application cases, but the IOPS are nearly 10x different.

The examples above are designed to illustrate a point *architecting a system to meet a balance of IOPS and Latency can go on for hours discussing controllers, memory, disk and networking (as do all performance baseline and bottleneck detection)*. Conversely, SSD has the ability to meet performance requirements of IOPS, while delivering low with little modification or discussion. Consequently latency and IOPS are easily balanced when using SSD – CHECK.

The final two constraints mitigated when using SSD compound each other – cooling and power. Let's take cooling first. For a system to be properly cooled it must be properly powered or geographically located. For simplicity, let's assume you can't build your data center in Prineville, OR. In that regard, it must be properly powered.

Since power must be adequate, the first thing a storage architect must consider is whether or not they can cool and power storage devices. Larger capacity systems offering higher IOPS and balanced latency require more power to cool and run them, thus compounding requirements. An architect must work with data center operations to balance cooling power with storage device power.

Here is where borrowing from Jerome Wendt, Lead Analyst and President of DCIG, is prudent:

> *Quantifying performance on storage systems has always been a bit like trying to understand Russia. Winston Churchill once famously said in October 1939, "I cannot forecast to you the action of Russia. It is a riddle, wrapped in a mystery, inside an enigma; but perhaps there is a key. That key is Russian national interest."*

Power is limited to the amperage available from an public utility. Limitations on available amperage creates a fixed constraint. Choosing storage with reduced power and cooling needs would mitigate the consequences of the fixed constraint. In that regard, SSD reduces complexities introduced by conflicting power and cooling architectures. While some may disagree, we know SSD requires less power and less cooling, and with less cooling, power needs are further reduced. SSD can or will eliminate the complexity related to power and cooling requirements – CHECK (Reference: The real costs to power and cool, IDC 06/2008).

Articulating storage needs isn't based solely on capacity. Storage architects must consider IOPS, latency (ORT), capacity required to meet IOPS & latency needs, data center rack space in the form of "U space", square footage for cooling, and physical (cooling/power) operations. While some will disagree that these are required with SSD, there complexity is significantly reduced if not eliminated in a broad deployment

of SSD.

While SSD can meet capacity, IOPS and ORT while reducing power and cooling costs, many of today's all flash memory storage arrays are based on x86 software and hardware. It is x86 that creates a barrier to entry for data center deployment of SSD. It is true that some may argue for x86 processing despite a high power-to-heat requirement, we know that ARM can deliver processing with substantially lower power-to-heat requirements.

To the point of power-to-heat, Calxeda published benchmarks indicating a 10x difference comparing x86-to-ARM power-consumption-versus-data-production with a +5 ms storage response time. From a marketing standpoint 10x is a great number, but a 5x difference is "good enough." 5x enables one to start thinking about replacing individual network attached storage (NAS) systems with private cloud scale-out storage systems using ARM processors and solid state storage modules or disk.

In that regard, it is my opinion that the market will desire ARM servers based on common slot architectures. Common slot architecture allows an Intel or AMD x86, and Samsung or Calxeda CPU to be plugged into the same board and system. Slot homogenization will reduce dependency on specific manufacturer motherboard designs (e.g. Intel) and allow for better elasticity in data center deployments.

As a result of homogenization, market pressure will pressure ARM processor vendors to enter scale-out NAS space in 2013. To that end, Calxeda silently noted their desire in late 2012 to enter the enterprise storage market in this piece by Stacy Higginbotham of GigaOM. Ms Higginbotham writes:

> *Its t*
> *ests show roughly a 4X improvement in IOPs for a rack of Calxeda SoCs versus x86-based systems. Adding Calexeda's*

> *SoCs also cuts complexity because the entire system of processing and networking components are integrated on the SoC, and the terabit-plus fabric between cores also offers more network capacity between cores in a system -the so-called east-west networking traffic.*

Calxeda's commentary muted the value of SSD, because Calxeda believes that the power hungry storage systems aren't concerned about power consumption. Instead they believe storage systems are looking for more IOPS by adding process and memory capability to a backlog of disk operations. While that assertion has some flaws, real value for ARM is power consumption and reduced heat signature. ARM combined with SSD delivers an investment annuity in operational and capital expenditure savings.

Complementing Calxeda's commitment to ARM is Apache's port of popular software that meets big data processing and storage requirements. Some will argue that SSD doesn't make sense in big data. But, common sense indicates storing 10 PB of data on spinning disk (HDD) over a period of a few years requires you to start migrating when the last TB is added. Controller aging alone require data to find a new home in immediately, or have a common slot server upgrade.

Factoring in an all SSD and ARM based scale-out storage system using open compute common slot architecture reduces or eliminates the top four (4) storage architect requirements ***AND*** delivers a storage ecosystem with the flexibility to exist in excess of 10 years.

Further complementing the marriage, ARM and SSD should have similar data center architecture requirement as tape. For example, let's track a company like Quantum with StorNext. They may port StorNext to ARM and take advantage of the $1/GB SSD prices as a way to transition tape customers from tape to new storage systems. Using ARM and SSD, very little would need

to change with the data center power and cooling.

Finally, look for companies like Samsung to be a powerful force in componentry as they continue to produce SSD and start the development of their ARM server processors. DCIG believes that as 2013 progresses, we'll experience a pull from the market for these storage systems long before the manufacturers are geared up to push them.

Fusion-io's Take on EMC's VFCache (Formerly Known as 'Project Lightning'); Interview with Fusion-io CMO Rick White Part V

EMC's VFCache announcement caused a lot of the buzz in the storage industry a few months ago as it was seen by some to be done in direct response to Fusion-io's very disruptive ioMemory architecture. Today in the conclusion of my interview series with Fusion-io's CMO Rick White, he provides his take on EMC's recent VFcache announcement and how he sees this impacting both Fusion-io and EMC. (*Editor's Note: This interview with Rick was conducted when EMC's VFCache was still known as "Project Lightning."*)

Jerome: *There have been a lot of rumblings coming out of EMC about its "Project Lightning" and how it is a Fusion-io killer. How does Fusion-io view the potential threat that Project Lightning presents to Fusion-io?*

Rick: Today EMC is talking about Project Lightning (*formally announced as VFCache on February 6, 2012*) which is such a huge shift for them. EMC has traditionally never had a footprint in the server. It is not what it does.

From what I am hearing the overall cost of ownership is not changing much. Hypothetically speaking, say you have a SAN that costs a million dollars. Now you have a new server caching solution, the SAN costs are cut in half.

That sounds pretty cool until you discover the cost for your 10 servers to have this new caching system installed is $50,000 per server so $50,000 times 10 for the caching solution is $500,000. So that cost plus the half a million dollars and suddenly you are back at a million dollars. Nothing has changed. The only changes are how EMC is going to invoice you for it.

I would be surprised if EMC builds a solution that decouples performance and capacity by deploying a scalable solution that lowers overall cost and improves efficiency. This is a fundamental difference we see between our two companies. For us – it is performance plus capacity. EMC is performance times capacity.

It happened to Digital Equipment Corporate (DEC). It happened to all mainframe manufacturers. The client-server environment was tough for them. They were selling a quarter of million dollar proprietary systems and suddenly a competitive solution emerged that had essentially the same performance for $10-15,000 per box based on commodity, off-the-shelf hardware and software components. It was a huge shift then. It will be a huge shift now for anyone in the storage business including EMC. It will be interesting to watch what happens.

Jerome: *Can't Fusion-io just go to server manufacturers and beat EMC at its own game?*

Rick: Manufacturers are getting closer. Fusion-io has been

establishing relationships with worlld's largest server manufacturers for the last couple of years. EMC is the newcomer to this space and is being forced to play catch-up with us.

That is probably why they need to OEM key technology and are looking at acquiring other pieces of technology. The market is moving fast and they just do not have time to do it themselves.

But it is frustrating. Others think flash on PCI-Express is all we do. Somehow they think they are entering 'SSD nirvana' because they have put flash on a PCI-Express card. You cannot put the flash on a PCI-Express card and call it 'good.' All they have done is taken the metal coverings off of flash drives and stuck them on a RAID controller. To be like Fusion-io, you also have to eliminate the large storage protocols and have the CPU interact with flash natively over the PCI-Express bus.

Flash drives have been speaking to the CPU through PCI-Express since they have first launched. They are only two ways to talk to the CPU – system bus (PCI-Express) or memory bus. That's it. There is no other way. Everything, host bus adaptor, RAID controller, graphics card, all communicate with the CPU through PCI-Express.

So just because you put the flash drives on a RAID controller and put them in the PCI-Express slot, it is no different than a RAID controller with eight (8) drives hanging off of it. You are still going to have a ton of context switching, which can cause dramatic and unpredictable swings in latency.

Jerome: *So is that all it takes to be like Fusion-io? Lose the storage protocols and interact natively with the flash?*

Rick: That is only the first step. Once you do that, you have to onload to the host CPU. This is similar to RAM. I have not seen a memory DIMM with an embedded CPU. I have not seen a

memory DIMM with SRAM as cache either. Most of us expect that, with more RAM, we can get our CPUs to do more work which means our CPU utilization goes up.

This idea that server flash has to use CPU offload and RAM as cache are both concepts inherited from hard disk drives. Hopefully the industry catches on to the fact that to unleash flash's true potential they need to treat it like memory rather than a hard disk drive.

The disk infrastructure was designed for a very slow medium – magnetic
disk. Lose the storage protocols. Use the DMA straight to the NAND
flash. Let the CPU make calls and access the NAND flash directly like as
if it were accessing a disk.

If they do then maybe we will stop hearing competitors say things like, "*You use host CPU cycles*." Although the simple answer to this is, "*Yes, yes we do. Just like your server's RAM does. We are persistent memory, not disk.*"

Another misnomer is that processing is the bottleneck. It is not. The reason many of today's biggest data centers do scale-out is not to get more processors because customers do not need more CPUs. They need memory. Often it is the easiest way to get the memory they need to keep data hot. It is not like they can go to disk for everything because of latency and what they are doing is not CPU-intensive.

I cannot tell you how many Fusion-io customers use less than 20 percent of their CPUs before adding Fusion-io. We help them improve the efficiency of their servers by allowing each CPU to do more work and increase utilization which ultimately increases the overall work output and productivity.

In