



ESG Lab Validation Report™



Hitachi USP V Enterprise-Class Services Oriented Storage Solutions

A validation study
by
ESG Lab
March 2008

Authors:
Brian Garrett, Tony Palmer

Table of Contents

Introduction	3
ESG Lab Validation	5
Dynamic Provisioning and Logical Pools.....	5
Performance	16
<i>OLTP Scalability</i>	<i>18</i>
<i>Wide Striping.....</i>	<i>20</i>
<i>Microsoft Exchange</i>	<i>22</i>
<i>SPC-1 Benchmark</i>	<i>23</i>
ESG Lab Validation Highlights	25
Issues to Consider.....	25
ESG Lab's View	26
Appendix	27

ESG Validation Reports

The goal of ESG Lab Validation reports is to educate customers about various storage and storage-related products, including storage systems, backup-to-disk solutions, storage management applications, backup and recovery software, storage virtualization platforms, etc. ESG Lab reports are not meant to replace the necessary evaluation process that end-users should conduct before making purchasing decisions, but rather to provide insight into these technologies. Our objective is to go over some of the highlighted features/functions of such products, show how they can be used to solve real customer problems, and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by HDS.

Introduction

ESG Lab has performed multiple hands-on tests of the Hitachi Universal Storage Platform (USP), starting in August, 2004 just prior to when it was first released. At that time, we focused on its performance, scalability and various features, paying special attention to its external storage virtualization capability. At the time, the notion of an Enterprise-class storage system that supported external storage virtualization was unique, a notion which persists to this day, and yet the value of such a system is obvious and compelling.

ESG defines virtualization as the logical view and control of physical infrastructure with the goal of greater optimization, utilization and simplification. Therefore, external storage virtualization provides a logical view and control of heterogeneous storage systems.

The industry has done a great job creating a partial network in which host systems share storage system resources. However, the storage systems themselves share no communication with one another. Aside from being on the same network, these storage systems do not, in any way, interact with each other. They are discrete and isolated devices that are networked to host systems, but not to each other. In reality, network storage systems are not fully networked. They do not work in concert, but rather separately, individually, discretely and therefore, inefficiently.

External storage virtualization networks your networked storage. It provides value from host system to storage system *and* storage system to storage system. The power of external storage virtualization is that it reinvents storage networking while minimizing disruption. For years, we have discussed the notion of a fluid data center, and virtualization is the liquid within this construct¹.

ESG was impressed with the USP external storage virtualization capability and observed it grow from a handful of early adopters to hundreds and now thousands of users. It is a brilliant strategy, leveraging a leading, field-proven Enterprise-class storage system that supports mission-critical environments, with all of its intelligence, software and reliability being leveraged to extend these capabilities to heterogeneous systems. It enables end-users to acquire the HDS storage system (which is budgeted for) and also implement storage virtualization (which may not have any allocated budget) within the same system.

Internal Storage Virtualization

ESG is a major proponent of storage virtualization. In fact, we believe that the more highly virtualized a storage system is, the more overall value it brings. Since its inception, the main focus of the USP has been on external storage virtualization. The new Hitachi Universal Storage Platform (USP V) goes the next step by adding compelling and valuable internal storage virtualization.

Internal Storage Virtualization comes in different forms and capabilities within a storage system. For example, the ability to abstract all of the disks in large logical pools is a core storage virtualization technology that provides significant value. Logical storage pools make it easier to manage and provision your disk capacity. Additionally, it can provide a huge performance boost by striping data across dozens, even hundreds or more, disk drives, making all of the disk actuators work in concert to perform read/write operations. Another valuable internal storage virtualization technology is thin provisioning. This allows for greater capacity utilization, brings more applications online without having to buy capacity in advance and simplifies the storage provisioning management process.

The USP V supports large logical pools and thin provisioning. ESG believes that these two capabilities are powerful forms of internal storage virtualization that raise the value of a storage system. Storage administrators have told ESG time-and-again that the management of individual storage systems is one of their biggest challenges. Cumbersome tasks such as storage provisioning and performance tuning consume too much time and too many resources. Storage systems that provide internal virtualization technologies significantly reduce complexity and enable more fully optimized systems without constant tuning and manual management.

¹ ESG Report: *Service Oriented Storage Solutions*, March, 2007.

Services Oriented Storage Solutions

Hitachi is extremely focused on providing a storage platform that can be used to enable IT to provide storage services more rapidly to their businesses. The HDS USP V and USP VM provide a number of capabilities to enable Services Oriented Storage Solutions including:

- External storage virtualization
- Thin provisioning of internal and external storage
- Large logical pools of capacity
- Virtual partitions
- Quality-of-service
- High performance, and massively parallel performance/mixed workloads
- Enterprise-class reliability and interoperability
- High levels of internal and external capacity support
- Full volume local and remote copies and mirroring
- Logical snapshots

This ESG Lab Validation analysis focused on verifying thin provisioning of internal and external storage, creating large logical pools of capacity and the improved performance of the USP V, which Hitachi delivers within the context of their Services Oriented Storage Solutions architecture.

ESG Lab Validation

ESG Lab validated the Hitachi USP V storage system during a day of hands-on testing at HDS corporate headquarters in Santa Clara, California. A fully configured, five-bay USP V system was used to test the enhanced performance and new capabilities of the USP V, including its Dynamic Provisioning software.


Dynamic Provisioning and Logical Pools

Dynamic Provisioning is the name Hitachi uses for a new and powerful volume creation capability that is generally known in the industry as thin provisioning. Dynamic Provisioning can be used to present a virtual pool of shared capacity that is larger than the actual amount of physical storage available. Dynamic Provisioning enables system administrators to deliver capacity on demand from a common pool of storage.

What differentiates Hitachi Dynamic Provisioning from the rest of the industry is the ability to dynamically provision from both internal and external storage, thus allowing advanced virtualization capabilities on older or lower tier arrays.

The difference between traditional provisioning and Dynamic Provisioning is shown in Figure One. In the case of a server that requires 300 GB today but is projected to grow to 2 TB, traditional methods would require the full 2 TB to be allocated today to avoid lengthy downtime for volume migration or expansion as the data set grows (as shown in red). Dynamic Provisioning allows a 2TB volume to be presented to the server with only 300GB of physical capacity (as shown in green). If and when the amount of capacity approaches 300 GB and exceeds defined thresholds, additional physical capacity can be added online with no impact to users and applications.

Figure One: Traditional vs. Thin Provisioning




Hitachi Dynamic Provisioning (HDP) is implemented using a concept that Hitachi refers to as HDP Pools.

A large number of drives can be part of an HDP pool and the USP V automatically stripes data across all available disk drives in the pool. Implementing a pool over a large number of disk drives is generally known in the industry as "wide striping."

HDP pools and wide striping are implemented in the HDS USP V as an extension to the traditional provisioning method shown in Figure Two. Traditionally, groups of drives with a defined RAID level (Array Groups) are carved into logical devices (LDEVs) which are presented to servers as Fibre Channel logical units (LUNS).

Figure Two: Traditional Provisioning



For applications that require striped volumes composed of many disk drives for capacity and/or performance, system administrators typically use host-based logical volume managers. The volume manager is used to stripe data over a number of LUNS. While volume managers can be used to make many physical hard drives appear as a single logical drive to applications, ideally this function would be performed in the disk array to avoid the CPU overhead, licensing costs and complexity commonly associated with volume management software running on multiple servers throughout the data center.


Historically, Hitachi has addressed this requirement with two capabilities referred to as Logical Unit-size Extension (LUSE) and Striped Logical Devices (Striped LDEV)s. LUSE volumes are a concatenation of multiple array groups into a single logical device. Because LUSE volumes lack the performance benefits of striping, Hitachi introduced the Striped LDEV feature in recent years. While powerful, a Striped LDEV is limited to using only two or four RAID5 7+1 array groups. Like other enterprise-class storage arrays which support these capabilities, Striped LDEV adoption has been limited to a few customers with extreme performance requirements due to complexity and configuration limitations.

With HDP Pools, USP V administrators now have the option of configuring host volumes from centrally managed logical pools. As shown in Figure Three, groups of drives with a defined RAID level (Array Groups) are assigned to a Pool. Pool capacity is used to define logical devices which are presented to servers as Fibre Channel logical units (LUNS). The array groups can be carved from both internal and external storage.

HDP Pools are automatically striped by the USP V. This allows large pools to be created with different performance and availability characteristics tailored to meet different application requirements. Dynamically provisioned LUNS carved from a pool are completely virtualized, utilizing capacity only as

data is written to disk. All volumes are able to utilize all of the disks in the pool with no volume management required on the host.

Figure Three: USP V Logical Storage Pools



Dynamic Provisioning and HDP Pools are a powerful combination which provides:


- Simplified capacity management from a centrally managed storage pool.
- Reduced capital costs, as over-provisioning is reduced or eliminated.
- Improved performance, as data sets are striped over a large number of drives.
- Advanced virtualization for external storage.

ESG Lab Testing

ESG Lab validation of Hitachi Dynamic Provisioning (HDP) was performed by presenting a large Dynamically Provisioned volume to a Windows server, filling the available space in the volume with enough data to exceed pre-set thresholds and expanding the volume using available disk in the array. ESG Lab also tested a Dynamically Provisioned volume created using external storage virtualized by the USP V. In these tests, ESG Lab focused on impact to the server and ease of administration.

The first step in building an HDP volume is by creating a large logical pool. This was accomplished from the Hitachi Storage Navigator management interface via the new "Pool" tab as shown in Figure Four. A Pool ID of 2 was chosen and a notification threshold was set. The notification threshold was used to generate a warning that the host volume was approaching the physical capacity allocated and required more physical capacity. ESG Lab chose a notification threshold of 70 percent for this test. ESG Lab noted that besides the user-configurable threshold of 70 percent, a system-level threshold fixed at 80 percent is supported to make sure that administrators are alerted when it is time to add more storage capacity.

Figure Four: Creating an HDP Pool




Two array groups were then added to the pool, yielding a total physical capacity of 531 GB. An HDP volume was configured and assigned to the recently created pool. The HDP volume was defined at 2,000 GB in size (2 TB). Note that the 2,000 GB of capacity presented to the Windows server was a completely virtual value and that the physical space allocated to the pool was only 531 GB.

From this point, the configuration of the HDP volume felt exactly the same as the configuration of a legacy LDEV volume. A logical LDEV was configured and the LUN manager user interface was used to define the USP V port and address presented to the Windows server.

Moving from the USP V console to the Windows server, the "Rescan Disks" option on the Windows Disk Administrator utility was used to recognize the new LUN. The 2,000 GB volume was detected normally and formatted using the Quick Format option as drive H:. Figure Five shows the Dynamically Provisioned volume as seen from both the HDS Storage Navigator screen and the Windows properties page.


Figure Five: A Dynamically Provisioned volume



The next stage of testing sought to demonstrate what would happen when the data stored on the volume exceeded the user defined threshold of 70 percent and the system level threshold of 80 percent. ESG Lab used an interesting method to simulate a drive being filled with data. A random write workload was created using the industry standard IOMETER tool with the maximum block defined to just cross the 70 percent threshold. This method simulated a drive that was 70 percent full in a matter of seconds.


ESG Lab noted that the Storage Navigator icon representing the status of the Logical Pool turned from blue to yellow when the 70 percent threshold was exceeded and then to red when the 80 percent threshold was exceeded. The screen shot shown in Figure Six, taken after a series of tests, indicates that Pool One is normal (blue), Pool Two has exceeded the 70 percent threshold (yellow) and Pool Three has exceeded the 80 percent threshold (red). System log entries were reviewed as each threshold was exceeded. It was noted that SNMP traps can be configured for automated e-mail notification and integration with industry-standard network management tools.

Figure Six: A Logical Pool Threshold Warning



After thresholds had been exceeded, additional array groups were added to the Logical Pool to restore it to normal status as shown in Figure Seven. In this case, a third array group is being added to Logical Pool Two to increase physical capacity from 531 GB to 804 GB. Once complete, the volume icon reverted to blue (normal) and the additional capacity was immediately available to the thin provisioned volume. The Windows server experienced no interruption in access to the volume and no action was required to utilize the additional space.

Figure Seven: Expanding a Logical Pool




When a new Array group is added to a logical pool, the USP V automatically begins using the wider stripes for new data to ensure that I/O distribution uses all disks in the pool.

Hitachi Dynamic Provisioning with External Storage

ESG Lab also tested Dynamic Provisioning using external storage virtualized by the USP V. ESG Lab chose a VMWare ESX server with a Windows client OS for this portion of the testing. Server virtualization affords organizations the ability to leverage largely unutilized processing power in mostly idle physical machines, while addressing growing power and cooling issues in the data center, by allowing organizations to deploy less physical hardware. The USP V's internal and external storage virtualization synergizes with these benefits, allowing organizations to purchase and deploy less physical storage with thin, just-in-time provisioning of multiple tiers of storage


The USP V under test had several external storage arrays that were attached and actively being virtualized. Figure Eight shows a simplified diagram of the test environment. A single VMWare ESX server was attached to a USP V via Fibre Channel SAN connections. Arrays from the major enterprise storage vendors were then attached to the back end of the USP V.

Figure Eight: The USP V with external storage



The procedures and tasks to dynamically provision volumes are exactly the same for external storage as for internal storage. In Figure Nine, three of the logical volumes ESG Lab selected for Dynamic Provisioning testing can be seen.

Figure Nine: Selecting External Volumes



External storage attached to the USP V is available for use by Dynamic Provisioning for Pool Storage, by Hitachi ShadowImage software for local copies, and by Hitachi TrueCopy software for remote replication. ESG Lab walked through the same steps as with internal storage to complete Dynamic Provisioning of external storage.

- Created new Storage Pool Four
- Assigned DMX LUNs to Storage Pool Four
- Configured HDP Volumes and assigned to Storage Pool Four
- Created logical LDEVs and assigned to the attached server.

Figure Ten: Thin Provisioned external storage





Figure Ten shows three 1TB thin provisioned volumes created from legacy external storage. ESG Lab confirmed that the USP V under test had a total of 886.43 GB of physical external storage allocated in Storage Pool Four and over 2.7PB had been Dynamically Provisioned and presented to servers from this pool.


ESG Lab confirmed that a previously created thin provisioned volume (LDEV 87) had been assigned to the ESX Server and used as the datastore to house Virtual Machine files. As seen in Figure Eleven, a Microsoft Windows 2003 guest OS system was installed on a virtual machine file allocated by ESX from the thin provisioned volume.

Figure Eleven: HDP Volumes in use by Windows under ESX



Finally, ESG Lab created a ShadowImage copy of LDEV 87 to validate business continuance functionality against thin provisioned, external storage volumes. The procedure for this was no different than when creating ShadowImage copies using internal storage. ESG Lab simply created a ShadowImage pair using the volume in use by the ESX server (LDEV 87) as the primary volume and an unassigned volume (LDEV 88) from HDP Pool Four as the ShadowImage copy. Figure Twelve shows the volumes after the pairing operation.

Figure Twelve: ShadowImage using external HDP storage



Why This Matters

There are a number of reasons why thin provisioning matters. Less storage is required initially when purchasing a new storage system with thin provisioning. Customers can buy what they need with a minimal amount of spare capacity for short-term growth.

Since there is no stranded storage capacity with thin provisioning, less storage is required over the life of the storage system.

Additional storage systems will not be required based on having stranded storage. ESG has found that end-users actually acquire and implement new storage systems because of allocated but unused storage capacity. Thin provisioning can eliminate this potentially major cost.

Thin provisioning can enable more applications and servers per storage system, providing greater levels of consolidation. Since capacity doesn't have to be dedicated and fixed on a per volume basis with thin provisioning, customers can create more volumes. As a result, more applications and servers can be attached to a single storage system.

The time and resource required to perform storage provisioning tasks is reduced. Thin provisioning can essentially eliminate the need for storage provisioning. Customers can create large logical volumes without any cost to them, allowing the application to keep growing and growing as needed. ESG has spoken to customers who have implemented thin provisioning and they report that they spend literally no time provisioning storage except when creating new volumes.

"Thin provisioning-aware" local and remote volume copy services consume less capacity. With traditional provisioning, if you have a 2 TB volume but only 100 GB is actual data, all 2 TB will still be copied. In some cases, customers will create two or more local copies of their primary volumes. With traditional provisioning this would consume lots of capacity for copies of empty blocks—thin provisioning eliminates this.

Thin provisioning of external storage provides all of these benefits for third party and lower tier storage. This extends users investment in existing business continuity software and management tools to external storage devices and enhances the value of the storage, software and tools.


Since thin provisioning is capacity-efficient, it requires fewer disks, which results in lower power and cooling costs. Often this is essential, for data centers that are running out of power, for instance—and it always makes sense to cut costs. Additionally, many organizations are mandating a green initiative, which raises the discussion to an even higher level.

Performance

The Hitachi Universal Storage Platform is an enterprise-class storage system that has been deployed by thousands of customers since it was first introduced in 2004. The USP was architected to meet the high-performance requirements of midsized and large organizations while providing a foundation for matching application requirements with different classes of storage. In previous reports, ESG Lab has confirmed that the massively parallel crossbar switch architecture of the USP can be used to support internal storage, heterogeneous external storage virtualization and a number of advanced capabilities including logical partitioning, quality-of-service and tiered storage migration².

The USP V provides a number of significant performance enhancements compared to the USP, as shown in Figure Thirteen. The architecture of the USP V is similar to the original USP, but has been enhanced in each of the areas indicated in green. Servers access the USP V through front end directors. Shared memory and cache are used to move control information and data through back-end directors to internal drives and externally virtualized heterogeneous storage systems.

Figure Thirteen: The Performance Enhanced USP V Architecture



² ESG Lab Validation: *Hitachi TagmaStore Universal Storage Platform*, September, 2004.
<http://www.enterprisestrategygroup.com/ESGPublications/download.asp?ReportAttachment=Attachment1&ReportID=234>
 And ESG Lab Validation: *Hitachi HiCommand Tiered Storage Manager*, November, 2005,.
<http://www.enterprisestrategygroup.com/ESGPublications/download.asp?ReportAttachment=Attachment1&ReportID=513>

The hardware and software performance enhancements in the USP V architecture are summarized in Table One.

Table One: The Performance Enhanced USP V Architecture

	What was Improved	Why this Matters
1	More bandwidth to shared memory	Commands start earlier and finish faster
2	More shared memory	Accomplish more in the same amount of time
3	More paths to shared memory	Faster access to more internal and external drives
4	2x faster processors on half-sized boards	Faster task completion with right-sized flexibility
5	4 Gbps FC switched back-end	Faster disk access, better serviceability
6	Faster path to external storage	Quicker access to virtualized storage
7	Wide-striped HDP pools	Improved response times for more users

In general, the enhancements shown in Table One can be used to support more users, applications and capacity on a centrally managed platform that leverages years of mature and reliable microcode. Specifically, ESG Lab has verified that the performance enhancements in the USP V yield extremely impressive low level characterization capabilities including:

- Significant improvements in internal bandwidth and processing capabilities.
- Major performance improvements at all levels.
- Up to 4 million I/O's per second.
- Up to 247 petabytes (247,000 terabytes) of heterogeneous external capacity support.
- Large wide-striped logical storage pool support.

The engineering level analysis of the low level performance capabilities of a storage system presented in this section is not an exact science. It does however provide a bit of insight into how the system will perform with real-world applications. In a way, it's like comparing the specifications of the engines of two cars. Those specifications provide a good starting place, but there are a variety of other factors, including the condition of the road and the skill of the driver, that matter as well. Much like buying a car, a test drive with real-world application traffic is the best way to determine how the system will perform in production environments.

While ESG Lab is impressed with the horsepower and theoretical scalability of the USP V architecture, readers are cautioned not to compare enterprise-class architectures based on internal bandwidth alone. What matters is the number of real-world applications and users that can take advantage of the internal bandwidth in parallel. Much like the electrical system in your home, figuring out how many appliances you can run in parallel before blowing a fuse is not a function of the number of wires behind the walls. What matters is the design of the circuits used to distribute the right amount of power to appliances when needed.


ESG Lab Testing

ESG Lab performed a series of tests to verify how the architectural improvements of the USP V translate into real-world application performance improvements. Performance scalability was verified using a simulated online transaction processing application and a simulated Microsoft Exchange Email deployment. The performance benefits of wide striping over an HDP pool were also examined.

OLTP Scalability

The test bed used for the simulated on-line transaction processing environment is shown in Figure Fourteen³. A top of the line IBM p595 server with 24 PowerPC processor cores running the AIX operating system was connected to a five-bay USP1100V using 64 four Gbps Fibre Channel connections.

Figure Fourteen: The OLTP Test Bed




A workload generator running on the IBM p-Series server was used to simulate a massively scalable large-scale on-line transaction processing system. The I/O workload was composed of 60 percent reads (40 percent writes) which were 90 percent random (10 percent sequential) and a 20 percent cache hit rate. Standard LDEVs built over 1,024 drives on the USP V and 512 drives on the USP were tested with 128 GB of cache. The USP V was tested with 15K RPM drives with 146 GB of capacity and the USP was tested with 15K RPM drives with 73 GB of capacity. Tests were run over 85 percent of total usable capacity.

The results shown in Figure Fifteen plot how long each transaction took on the Y-axis as the amount of traffic measured in database I/O operations per second (IOPS) was increased along the X-axis. This graph, known in the industry as a "response time throughput curve," provides an excellent visual picture of how well a storage subsystem performs under duress.

Note that as traffic increases (on the X axis); response time tends to increase (on the Y-axis). When a storage system becomes overloaded, it has a hard time keeping up, causing queues to reach their maximum, response times to elongate and the system to feel slow to end-users. This point-of-no-return, known as the "knee-of-the-curve," represents the breaking point after which more traffic cannot be handled gracefully. The best performing system is therefore represented by the lowest curve extending the furthest to the right.

³ Configuration details for this, and for all performance results presented in this report, can be found in the Appendix.

Figure Fifteen: The OLTP Results



What the Numbers Mean

- Comparing the response time throughput curves shows that the USP V not only performs significantly more work (furthest curve to the right), but also does so while responding quickly (a low riding curve).
- The USP V was tested with twice the number of drives and twice the usable capacity of the previous generation USP, yet was able to sustain significantly more transactions.
- At a response time of 4 milliseconds, the internal architectural improvements of the USP V delivered 64 percent more concurrent transactions.

Wide Striping


Mechanically spinning hard drives are the slowest component in the storage performance chain. The implementation of additional drives behind each application volume enables more work to be done in parallel. This is particularly true for massively interactive applications including online transaction processing and e-mail systems. This phenomenon is well-known by database and storage administrators who have been using host-based volume management software for years to make multiple Fibre Channel LUNS appear to an application as a single volume. Hitachi Dynamic Provisioned (HDP) volumes defined over a wide-striped pool of drives can be used to achieve the same performance benefits.

ESG Lab Tested

The configurations shown in Figure Sixteen were used for this round of testing to evaluate random I/O performance for three configurations:


- A single Windows Logical Disk Manager (LDM) volume defined over eight array groups
- A single Hitachi LDEV volume defined over a single array group
- A single Hitachi HDP volume defined over eight array groups

Figure Sixteen: Wide-Stripe Performance Testing



A Windows server was connected over a single 4 Gbps Fibre Channel interface to a USP V. The workload used for this test was 100 percent random, 75 percent reads and an 8 KB block size. Thirty-two 146 GB 15 RPM Fibre Channel disk drives were configured as eight array groups (2+2)⁴. Multiple tests were performed on each configuration, using a range of active I/O threads between eight and 256.

Figure Seventeen: The Wide Striping Performance Advantage



What the Numbers Mean

- Wide striped volumes over eight array groups can do significantly more in-parallel work than an LDEV configured over one array group. For example, the HDP volume with eight times as many drives as the LDEV performed 716 percent more transactions at a response time of 15 milliseconds or less.
- Wide striped volumes not only process more work in parallel, but they also do it faster. For example, the response time of the HDP volume was 118 percent faster than the LDEV volume at the first data point on the curve.

⁴ The diagram depicts 16 instead of 32 drives for simplicity.

Microsoft Exchange


Microsoft Exchange is an industry leading e-mail and office productivity tool that is very I/O intensive. The workload is mostly random in nature and requires good sequential write performance for logging. The I/O-intensive nature of e-mail messaging traffic benefits from a large number of drives working in parallel. Predictable performance scalability is a requirement for a growing number of organizations consolidating e-mail servers and users onto a consolidated pool of shared storage.

ESG Lab Tested

The Microsoft Jetstress tool was used to simulate the e-mail traffic for a number of “heavy” Exchange users⁵. Jetstress is a standard utility that is used by Exchange administrators for system sizing and burn-in. Jetstress is also the tool that vendors use to publish best practices and sizing recommendations as part of the Microsoft Exchange Solution Reviewed Program (ESRP).

ESG Lab performed a series of tests to evaluate the performance capabilities of thin-provisioned, wide-striped Hitachi Dynamic Provisioned (HDP) volumes on a USP V storage system and concatenated LUSE volumes on a USP. The same number of 15 K RPM Fibre Channel hard drives was used for each test. The maximum number of users that could be supported on each configuration with an average response time of 16 milliseconds or less is shown in Figure Eighteen⁶.

Figure Eighteen: Microsoft Exchange Results



What the Numbers Mean

- HDP volumes on the USP V, which are easier to configure and manage than LUSE volumes on the USP, can be used to support 38 percent more Exchange users.
- An ESG Lab audit of internal metrics revealed the source of the performance gains in the USP V. Faster processors on front and back-end directors, architectural improvements and the new switched 4 Gbps backplane provided dramatic improvements in CPU utilization and write latency.

⁵ Microsoft guidelines indicate that a heavy Exchange user generates one I/O per second on average.


⁶ The Microsoft ESRP guidelines indicate that these latencies should average under 20 milliseconds.

- Random I/O accesses are clustered around hot spots in real-world environments as opposed to this test, which executed random I/Os over the entire range of available capacity. ESG Lab is therefore confident that Exchange stores deployed over large pools of wide striped HDP USP V volumes can be used to deliver even better performance improvements than the 38 percent shown in Figure Thirteen.

SPC-1 Benchmark

In addition to running physical tests against the USP V, ESG Lab audited HDS' published results of the SPC-1 application-level industry standard benchmark suite maintained by the Storage Performance Council. SPC-1 testing generates a single workload designed to emulate the typical functions of transaction-oriented, real-world database applications. Transaction-oriented applications are generally characterized by largely random I/O and generate both queries (reads) and updates (writes). Examples of those types of applications include OLTP, database operations, and mail server implementations. SPC results can be roughly mapped by users into easily understood metrics. For a credit card database system, for instance, it might be the number of credit card authorizations that can be executed per second.

Figure Nineteen: SPC-1 Results



HDS has published an outstanding result of 200,245 SPC-1 IO requests per second at 100 percent load with an average response time of only 4.99 milliseconds; In fact, this is the best published result for a disk based enterprise class storage controller⁷.

An extremely important component of SPC results, response time is critical as this is the delay that an application will experience (and pass on to users) when a storage system is stressed to its limits. The generally accepted threshold for online transactional applications is 10 ms, beyond which application delay becomes apparent to users. The SPC will invalidate any results higher than 30 ms. In this context, the USP V's performance is even more impressive. It should be noted that the numbers published in these reports are for disk subsystem IO requests. Since real world applications typically issue multiple IO's per transaction, there is no easy way to say how much real world performance these numbers represent. Suffice it to say that a system that can process 200,245, random, small I/O requests per second with an extremely low response time can process huge volumes of credit card authorizations per second.

⁷ <http://www.ideasinternational.com/benchmark/ben020.aspx?b=99084824-4650-44ce-996b-93b06f10bef1>

SPC-1 results are audited by the Storage Performance Council and peer reviewed to ensure consistency. Executive Summary and Full Disclosure Reports (FDRs) for each SPC benchmark result are publicly available for download and review⁸. While this can be useful for comparison between vendors, it is important to note that not all vendors participate and publish results. ESG Lab hopes that participation by HDS will encourage other vendors of enterprise class storage systems to participate.

Why This Matters

ESG Lab evaluated the USP V performance and found that on multiple levels, Hitachi has raised the bar to a significant degree. The USP V performance improvements are due to key architectural improvements and the use of wide striping—a combination of both hardware and software core competencies.

The USP V takes performance to another level. This is also innovation at its best—improving the architecture of a storage system to raise its performance at nearly all levels, in leaps. Performance is not discussed as often as it should be. However, if your applications don't perform well then your business suffers.

A major benefit of large logical pools of capacity is the ability to stripe data across a large number of disk drives. This improves performance significantly since you can have dozens or hundreds or even thousands of disk drives operating on I/O operations simultaneously.

End-users do create large virtual volumes using host-based technologies but this is not required with the USP V, which provides this functionality internally. Centralizing this functionality enables easier management.

Additionally, wide striping significantly reduces performance management tasks. There is often a great deal of time needed to analyze, tune and re-tune performance with traditional provisioning methods. Wide striping minimizes and potentially eliminates the need to tune and analyze since it is perpetually optimized for performance.

Finally, ESG Lab was particularly impressed with HDS' peer-reviewed, industry standard SPC-1 benchmark results which demonstrate that the USP V is delivering industry leading performance while providing exceptional storage economics to end-users.

⁸ http://www.storageperformance.org/results/benchmark_results_spc1

ESG Lab Validation Highlights

- ☑ Hitachi Dynamic Provisioned volumes were configured using a centrally managed pool of thin-provisioned capacity drawn from both internal and external third-party storage.
- ☑ The Storage Navigator interface was used to monitor pool capacity as a 70 percent user defined threshold and an 80 percent system threshold were exceeded.
- ☑ Capacity was added to a pool using the Storage Navigator interface while the USP V was online and actively processing application I/O requests.
- ☑ Hitachi Dynamic Provisioned Volumes using external storage were used by VMware's ESX server to house OS images of Virtual Machines.
- ☑ A simulation of a massively parallel online transaction processing system on an IBM p-Series server proved that the hardware and architectural improvements in the USP V can be used to serve more than 100,000 concurrently active users on a singly managed system.
- ☑ A simulated Microsoft Exchange test was used to show that wide-striped HDP volume support on the USP V provides an excellent platform for storage consolidation and savings.
- ☑ A performance increase of 716 percent was measured for a random workload running over a wide-striped HDP volume with eight times as many drives as a traditional LDEV volume. Similar performance benefits were measured using a host-based volume manager from Microsoft.
- ☑ Audited SPC-1 results verified the USP V's performance capabilities in an open, peer-reviewed forum.

Issues to Consider

- ☑ Dynamic Provisioning is not supported, or needed, on mainframes.
Due to the historically small size of devices recognized by the IBM Z-OS operating system, and a plethora of mature capacity management and reporting tools available to mainframe storage administrators, the capacity utilization problem that Dynamic Provisioning was designed to solve on open systems platforms is better solved in mainframe environments using existing tools and processes.
- ☑ Application-specific best practices are recommended.
Thin provisioning is a relatively new concept and there is a lot of confusion about how well it works with real-world applications. ESG Lab has spoken with dozens of early adopters who have realized significant savings in production environments using a wide variety of applications. In addition, Hitachi has published a number of best practice articles advising users on various approaches of deploying Dynamic Provisioning. With that said, ESG Lab suggests that customers work with Hitachi to ensure that application-specific best practices are being met. In particular, applications and utilities which pre-allocate space and write towards the end of a drive can negate the benefits of thin provisioning. As an example, a full format of an NTFS file system on a Windows server pre-allocates space and writes data over the entire extent of the file system. As a result, ESG Lab adhered to the best practice of using the Microsoft quick format option when formatting NTFS file systems during the ESG Lab Validation.

ESG Lab's View

It is important to note that the USP V is not just about technology and features. It is also a vision that HDS is executing on: a highly virtualized storage system, with a rich set of system and data management features on a high performance reliable platform that enables end-users to provide best-in-class storage services to the business.

HDS is delivering innovation—the USP V is not a “me too” solution, and it’s not innovation for innovation’s sake. The level of capability and scalability is compelling and extremely valuable to end-users.

The USP V is a best-in-class enterprise storage system and storage virtualization platform. The USP V provides a series of capabilities that individually provide great value. When these capabilities are conjoined, it raises the stakes, offering value as a single storage system, as a storage virtualization platform and as a Services Oriented Storage Solution—provided by a leading, world-class storage vendor.

Storage services need to be delivered rapidly. Are you willing to wait weeks or even months to get a new service like cable or telephone to your home? Are you tolerant of a service that is slow or intermittently unavailable? Are you okay with someone losing your personal data or, worse, having someone steal it? And would you personally be willing to pay top dollar for this service? The answer is a resounding and self-evident No. However, this is often the case with storage services.

The ultimate driver for Services Oriented Storage Solutions is to enable the IT organization to provide world-class storage services to the business—to make the limitations of IT invisible, but make the triumphs apparent. The goal is for IT to transcend the nuts and bolts, becoming strategic and a core contributor to the business. It is necessary for the Services Oriented Storage platform to address the challenges of IT: to network their storage systems, enable centralized management, automate cumbersome manual processes and provide innovation without disruption.

ESG believes that infrastructure virtualization has to live within multiple layers of the data center and server virtualization will have a profound impact on storage. ESG has observed wide and massive adoption of both virtual server and storage virtualization technologies. Server virtualization: affords organizations the ability to leverage under-utilized server hardware while addressing growing power and cooling issues in the data center by allowing organizations to deploy less physical hardware to address data processing requirements. The USP V’s internal and external storage virtualization synergizes with these benefits, allowing organizations to purchase and deploy far less physical storage with thin, just-in-time provisioning of multiple tiers of storage providing advanced performance, data protection, replication and virtualization.

The USP V raises the level of play by supporting external and internal virtualization. Both are essential to getting the most out of your storage assets. External virtualization creates a network of otherwise isolated storage systems in order to move and replicate data between them. Internal virtualization enables IT to more fully optimize individual storage systems and make them far easier to manage. Storage virtualization leverages storage assets to be greater than the sum of their parts. But virtualization isn’t a panacea. It is important to look at solutions holistically and demand best-in-class systems. In ESG’s view, Hitachi truly understands this and the USP V embodies these goals.

Appendix

The configuration used during ESG Lab testing is listed in the following table.

Table Two: ESG Lab Validation Configuration Details

The OLTP Test Bed	
Server	IBM p595 Model 9119, 24 1.65 GHz Power5 CPU, 96 GB RAM, 64 4Gbps FC adapters, AIX version 5.3 ML3
Storage	USP V 1100, 256 GB cache, 12 GB shared memory, 8 front-end directors, 8 back-end directors, 1024 146 15 RPM FC drives, Microcode 60-00-31
The Microsoft Exchange Test Bed	
Server	Dual Xeon 2.8 GHz, 4 GB RAM, Windows 2003
Storage	USP V 1100, 256 GB cache, 12 GB shared memory, 8 front end directors, 8 back-end directors, 1024 146 15 RPM FC drives, Microcode 60-00-31
The HDP Pool vs. Host Based LVM Test Bed	
Server	Dual Xeon 2.8 GHz, 4 GB RAM, Windows 2003
Storage	USP V 1100, 128 GB cache, 12 GB shared memory, 8 front-end directors, 8 back-end directors, 1024 146 15 RPM FC drives, Microcode 60-00-31