

Hitachi “Innovation Machine” Strikes Again with USP V and USP VM Enhancements, November 2007

Josh Krischer is an expert IT advisor with 37 years of experience in high-end computing, storage, disaster recovery, and data center consolidation. Currently working as an independent analyst at Krischer & Associates GmbH, he was formerly a Research Vice President at Gartner, covering enterprise servers and storage from 1998 until 2007. During his career at Gartner he was responsible for high-end storage-subsystems and spoke on this topic at a multitude of worldwide IT events, including Gartner conferences and symposia, industry and educational conferences, and major vendor events.

Before the ink has dried on the Hitachi Data Systems USP VM announcement from September 10th 2007, the Hitachi “innovation machine” strikes again, raising the bar on externally virtualized storage with new enhancements for the USP V and USP VM (technically equivalent to HP’s XP24000, XP 12000 and Sun StorageTek 9990V and 9985V). These enhancements, which include thin provisioning of external storage from any third-party vendor, help extend the value of an enterprises’ storage investment.

This is the third HDS (and partners) high-end enterprise storage enhancement in the last six months, starting with the introduction of the USP V and followed by the smaller, more affordable, USP VM.

November, 5th, 2007 announcements at glance

- Hitachi Dynamic Provisioning™ enhancements to support *externally attached arrays* delivering thin provisioning with wide disk striping across multiple logical storage pools
- Support for internally attached, lower cost 750GB SATA II HDDs
- More than 4 million IOPS from cache, which is a 60 percent improvement over the previous-generation USP

Reminder: May 14th, October 1st, 2007 USP V main announcement points

- Introduction of the Hitachi Universal Storage Platform™ V with over 40% overall system performance improvement
- Internal virtualization, which includes Hitachi Dynamic Provisioning™ (thin provisioning) software with wide disk striping across large logical storage
- Extension of external virtualization capability to manage up to 247PB
- Five-fold (x5) performance improvement for virtualized storage ports
- Doubling of remote replication port performance
- Numerous availability, scalability, configurability, and serviceability enhancements

- Hitachi Data Systems published official SPC-1¹ benchmark figures for the USP V (1st October 2007)

November 2007 enhancements analysis

Hitachi engineers have achieved an incredible technological advancement by enabling the USP V to support **thin provisioning (Hitachi Dynamic Provisioning) of external storage**. This also enables all storage connected to a USP V to enjoy the “green” effects of lower power and cooling requirements, along with simplified capacity management.

Thin provisioning enables allocation of virtual storage as needed without the requirement to dedicate physical disk storage up-front. Additional capacity can be allocated on demand from existing storage assets, as well as newly-installed capacity, without disrupting applications. This essentially delivers new capabilities to storage systems that were never designed to encompass these advanced features.

Thin provisioning, in addition to saving investment and running costs (less energy, smaller floor space requirements), also improves performance by striping the data across all the disks in the array. Striping the data among a large number of physical devices practically eliminates “hot spots,” resulting in almost uniform performance.

The USP V can now support internal disks, as well as all storage connected to the USP V or VM from EMC, HDS, IBM and Sun. An additional utilization benefit of thin provisioning external storage is that enterprises will realize additional capacity for free, from their currently installed assets. This capacity can also be instantaneously provisioned, further reducing administration costs.

As a fail-safe mechanism to ensure continuous application availability, Hitachi Dynamic Provisioning software has built-in threshold monitoring to alert when the capacity pool is close to being exhausted. There is a “soft” warning alert when the used capacity is 60% of the installed capacity and a user selectable “hard” warning when the pool is almost empty. The USP supports non-disruptive upgrades if additional capacity is required.

Additionally, Hitachi announced support for internally attached, lower cost 750GB **SATA II HDDs** which increase the available options for building tiered storage. SATA disks have a lower availability than enterprise FC HDDs, however, this handicap is compensated by a RAID-6 configuration. RAID-6 in 6D+2P configurations

¹ “SPC-1 consists of a single workload designed to demonstrate the performance of a storage subsystem while performing the typical functions of business critical applications. These applications are characterized by predominantly random I/O operations and require both queries as well as update operations. Examples of those types of applications include OLTP, database operations, and mail server implementations.” (Source: SPC)

dramatically decreases the dangers of data loss with SATA disks. This technique consumes 12.5 percent more storage than RAID-5 in 7D+1P configurations but ensures an almost indefinite mean time between data loss (MTBL) and reduces the rebuild time by 60-percent in comparison to RAID-5 groups on the same system. In random write operations, the RAID-6 technique may impact performance by increasing the “write penalties,” but no such impact should be registered in large blocks of sequential writes.

Up until this announcement, tiered storage configurations with SATA disks could be built only by using external subsystems connected to the back-end of the USP.

Hitachi announced over 4 million IOPS of **maximum performance** on the USP V—more than 60 percent faster performance of the original Universal Storage Platform first introduced in September 2004.

Transactional performance is measured as response time in milliseconds and as maximum throughput (which is measured in number of I/O operations per second - IOPS) before the machine enters saturation causing response time to grow exponentially. Performance of sequential operation is measured in maximum data transfer rates of MB/s or GB/s.

For the first time, on October 1st, Hitachi Data Systems has published **Storage Performance Council benchmark** results for its USP. The USP V achieved the highest SPC-1 benchmark result in enterprise storage system history with 200,245.73 SPC-1 IOPS in the SPC-1 benchmark with a single storage controller. In comparison², according to the SPC-1 figures from December 5 2006, IBM's System Storage DS8300 Turbo achieved 123,033 SPC-1 IOPS. SPC benchmarks are modeled after real-world applications and therefore help in providing customers with meaningful performance results. Many vendors participate in SPC, and all vendors are encouraged to publish SPC benchmark results; however, HDS and IBM are the only two of the top three enterprise disk vendors who have posted results with the Storage Performance Council. EMC has never officially published its performance figures for the Symmetrix DMX.

Future enhancements and Statements of Directions

Hitachi plans enhancements to the USP V that will be available in late December and early into 2008, including:

- Doubling of cache capacity from 256GB to 512GB on USP V and from 64GB 128GB on USP VM

2 The figures above do not include the performance improvements October 2007 enhancements of the IBM DS8000 series.

- FICON connected external storage to simplify mainframe migration – currently, external storage is connected by FC only
- Additional Improvements of application performance and availability
- Expanded virtual partitioning with larger cache resources
- Quick Format – Improved storage management by shortening the time required to provision new disk space
- Snapshot for Thin Provisioned storage
- Audit Log in-band to improve security options
- Doubling the supported number of “pairs” in remote copy
- HyperPAV – Improving mainframe application access to data while lowering addressing requirements (planned for end November 2007)

Quick Format – Usually, formatting a volume may take several hours in which the volume is not usable by applications. The quick format feature allows use of the volume in a few minutes after initiation of the formatting process without the need to wait until it is finished. Users benefit from much quicker availability, in particular in situations requiring urgent changes such as recovery from problems or disasters.

Snapshot – currently, customers using Hitachi Dynamic Provisioning have the ShadowImage feature at their disposal in order to create point-in-time copies. This enhancement will add a Snapshot option, which only consumes capacity when preserving pre-change source volume tracks, and therefore, requires far less storage capacity compared to the “full volume” of the ShadowImage technique. Users will benefit from better storage utilization and lower TCO.

Many large data centers (in particular external service providers) use Hitachi Universal Replicator (HUR) to facilitate remote mirroring for external, heterogeneous storage systems. **Consistency groups** support will extend for up to 512,000 “pairs” across multiple systems to improve disaster recovery capabilities for such data centers.

HyperPAV - relieves System z logical volume size constraints and performance limitations of static and dynamic PAV (Parallel Access Volumes introduced in 1999 along with the IBM 2105 ESS) and Workload Manager (WLM)-managed aliases. HyperPAV was announced in October of 2006 for operating systems higher than z/OS 1.6.

Alias addresses managed by HyperPAV are assigned by the I/O Supervisor (IOS) according to the request priority of the WLM. These aliases are kept only for the duration of an I/O operation. After completion of the I/O operation, the aliases are returned to a “free pool” to be used by another high priority I/O request.

The dynamic operation of HyperPAV may significantly reduce the number of aliases required to meet the WLM's performance objectives compared to dynamic PAVs. Large System z sites which are constrained by the lack of free PAV aliases should consider exploiting the HyperPAV feature.

Summary and Conclusions

It seems that the Hitachi “innovation machine” in Odawara, Japan never sleeps, and with every enhancement, Hitachi and its partners (HP and Sun Microsystems) increase their competitive edge. This announcement significantly increases USP V systems' functionality, performance and flexibility.

By introducing thin provisioning for externally-attached arrays, Hitachi continues to advance the state-of-the-art for managing large, multi-vendor enterprise storage environments. Today's announcement is no exception and should be evaluated by a vast majority of the world's storage administrators.

Hitachi is the first company to deliver thin provisioning for external storage, which allows heterogeneous storage to be combined into a single pool that can be more easily managed from a central console.

The USP V is the most “virtualizable” storage system available today supporting internal and external volume virtualization, thin provisioning, partitioning, and virtual ports. It leads the market in every aspect: performance, functionality, availability, ease of use, and flexibility and supports of large scalabilities while delivering superior performance.

These features are a welcome complement to Hitachi's recent accomplishment and are a noteworthy follow-on to HDS's recent announcement as the first vendor to obtain external storage virtualization certification from VMware and support for VMware ESX Server 3.0. The VMware certification facilitates the integration of server and storage virtualization infrastructures to achieve even greater levels of reduced IT costs, simplified management, improved operational efficiencies and quality-of-service for applications.

The USP V high-end enterprise storage system is well positioned to cover current and future user requirements and should be considered by any large or small organization searching for excellence and improved application performance and availability in their storage infrastructure.