Reference Code: OI00070-042

Publication Date: March 2011

Author: Tim Stammers

TECHNOLOGY AUDIT

Virtual Storage Platform

Hitachi Data Systems

SUMMARY

CATALYST

The VSP is Hitachi's flagship block-level disk array, and the latest evolution of a series of devices that originated in the 1990s as the Hitachi Lightning. The VSP inherits the reputation of its predecessors as a tier-one disk array that offers very high levels of performance, scalability, availability, and data protection, alongside support for mainframe and open systems hosts.

KEY FINDINGS

Strengths:	✓ Highly modular design ✓ Ability to virtualize third-party disk arrays ✓ Use of high-density 2.5-inch SAS disk drives
Weaknesses:	✗ Only suitable for enterprise use
Key Facts:	i The smallest VSP configuration costs around $200,000

OVUM VIEW

The VSP has a number of rivals in the high-end storage market, including Hewlett-Packard's 3PAR disk arrays, and NetApp's FAS6000 series of products. But at the top of its market sector, the Hitachi device has only two major competitors – EMC's Symmetrix and IBM's DS8000 disk arrays. Against these two products the VSP's predecessors have been extremely competitive. The VSP itself was launched in late 2010, which currently gives it the advantage of being the most recent of the trio to have undergone a major hardware and software update.

The VSP introduced a more modular and higher performing hardware architecture, alongside software features such as fine-grained automated movement of data across storage tiers, and an entirely GUI-based management interface.

© Ovum. This Technology Audit is a licensed product and is not to be photocopied

Among the qualities of the VSP that Hitachi is pitching hardest are its low power and physical space consumption. As well as helping to reduce the overall ownership costs, this will appeal to many businesses whose data centers have reached the limits of power supply, cooling capacity, and floor space. One major reason for Hitachi's power and space claims is the VSP's use of 2.5-inch SAS disk drives, which offer greater physical density and lower power consumption than the 3.5-inch Fibre Channel drives that they have replaced. Although this move is part of an overall industry trend, EMC's Symmetrix array has yet to make the same switch to the smaller disk drives.

The increased modularity of the VSP has brought its entry-level price down to approximately $200,000. This allows a relatively small VSP to act as a storage platform with the potential to scale to very large sizes, while providing the high-end performance that is now needed in even mid-range IT shops, because of the loads placed on storage by the proliferation of virtual servers. In addition, the VSP is the only high-end disk array that can virtualize other disk arrays that are attached to it. Like other block-level disk arrays the VSP can provide file-level data access using a NAS gateway. It can also manage part of its capacity as object-oriented storage, which is well suited for archival or compliant data. Overall, the VSP can be considered as both a modular, scalable storage system, and the central pillar of a tiered storage infrastructure embracing block, file, and object storage.

FUNCTIONALITY

SOLUTION OVERVIEW

The VSP includes software functions that are rapidly becoming standard for mid-range to high-end devices, such as thin or dynamic provisioning of disk space, and automatic sub-LUN movement of data across storage tiers. However, the VSP sports other features that are restricted to high-end devices, such as its scalability and support for mainframe FICON connections, application-aware data replication between up to four data centers, and very high levels of availability.

The salient features of the VSP are:

- maximum capacity 2.5PB across 1,280 3.5-inch SATA drives
- maximum capacity 1.2PB across 2,048 2.5-inch SAS drives
- maximum 256 flash drives, 200GB or 400GB each
- maximum 1TB global cache
- maximum 284 Fibre Channel or FICON ports, and 96 10Gbit FCoE ports
- maximum 65,536 Open Systems or IBM Z/OS LUNs
- high-end reliability and availability features including non-disruptive update to processors, front-end ports, and cache, and flash protection of write-cache during power outages
- scaling from one to four pairs of storage directors, and from one to six 19-inch cabinets.

© Ovum. This Technology Audit is a licensed product and is not to be photocopied

SOLUTION ANALYSIS

Architecture

The predecessors to the VSP featured a back-end switch which linked cache memory to front-end and back-end processors. The VSP is an evolution from those designs. It features a version of its predecessors' grid switch architecture that has been upgraded to provide 50% more throughput, but also introduces a new, third type of blade-mounted processor. Although I/O and RAID operations are still handled by front-end directors (FEDs) and back-end directors (BEDs), all other data manipulation is now handled by blades that HDS calls virtual storage directors (VSDs). HDS describes this as a unique and tightly coupled architecture, which allows all three types of director to share access to a single image global cache. Unlike its rivals, this allows the VSP to spread processing workload across all its directors.

As well as claiming that this improves efficiency and performance, HDS says that this has advantages in terms of flexible scaling, cost, and energy consumption. For scaling, the number of redundant pairs of FEDs, BEDs, and VSDs can be increased independently of each other – up to six, two, and four pairs respectively. This provides customers with a wide choice of configurations to suit their individual needs. Rival high-end disk arrays do not provide this level of flexibility.

HDS' biggest reason for introducing VSDs was the increasing workload created by thin provisioning and automated movement of data across tiers. The VSDs are powered by quad-core, Intel x64 processors, while the FEDs and BEDs are powered by ASICs designed specifically to perform IO operations. The introduction of x64 processors into the VSP reflects an industry trend to build storage arrays using as many low-cost volume components as possible to reduce costs. However, HDS believes that for high-performance disk arrays, a mix of ASIC and general-purpose processors provides the best trade-off between cost and performance. The use of ASICs also reduces power consumption.

Figure 1: VSP architecture

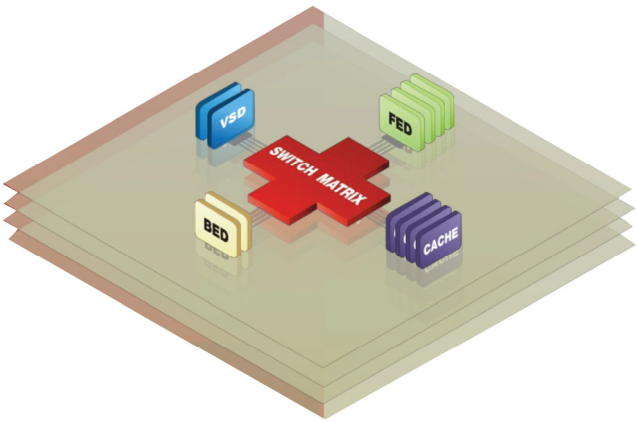

Source: HDS

OVUM

© Ovum. This Technology Audit is a licensed product and is not to be photocopied

Performance

HDS says that the VSP has twice the power of its predecessor the USP V. That would make the VSP very powerful, as the USP V itself recorded the highest performance of any disk array tested to the Storage Performance Council's SPC-1 benchmark. As this report went to press, the SPC was auditing test results for the VSP

The SPC-1 benchmark measures random I/O performance, and has been criticized by some vendors as being far from representative of real-world I/O loads. However, the majority of large storage suppliers have enthusiastically supported it. Although vendors sometimes test unrepresentatively expensive racing configurations of their arrays, the SPC requires them to clearly state the list price of whatever they tested, and involves a rigorous physical audit. As such, the benchmark does indicate the potential performance of an array. Customers can use this information to help decide whether the device is worth considering for more specific tests using their own applications and intended configurations.

VSP as third-party virtualizer

The VSP can absorb or virtualize the capacity of other disk arrays that are attached to it, and treats the capacity of those arrays as if it were internal to the VSP. This ability was introduced by HDS several years ago, in the HDS USP. As such it is a mature, well tested technology. Arrays from both HDS and third-party suppliers can be attached to the VSP. HDS has qualified approximately 100 devices for this purpose, including devices from EMC, IBM, HP, Sun, and other suppliers. Almost 255PB of extra capacity can be added to a VSP, which is obviously a theoretical limit far beyond any practical usage.

Extra capacity is subject to all of the functions of the VSP, such as thin provisioning, automated tiering, and replication. Although the VSP cannot improve the performance of third-party arrays, HDS says that putting them under the control of a VSP may boost overall, end-to-end application performance – presumably because of the effect of the VSP's tiering, wide striping, and larger cache.

Using the VSP to virtualize other disk arrays allows older storage systems to continue active service, under the control of the VSP, and with the benefit of its advanced features. Disparate disk arrays each with different management interfaces are pooled into one logical device, with one management console. Older disk arrays that might have been acquired through a company merger can be used as a lower storage tier.

Fine-grained automatic tiering

Automatic movement of data across storage tiers significantly improves overall performance by increasing the overall storage efficiency. By moving infrequently accessed data to slower tiers of SATA drives, it frees up space to store more frequently accessed data in faster tiers comprising flash and SAS disk drives. The VSP offers automated tiering as an optional function.

Automated tiering is rapidly becoming a common feature in mid-range to high-end disk arrays, and competition between suppliers has moved on to the effectiveness of their tiering systems. One of the major factors of this is the size or granularity of individual data movements.

© Ovum. This Technology Audit is a licensed product and is not to be photocopied

The smaller the size of each movement, the greater the potential improvement to overall storage efficiency. Although the USP V featured automatic tiering, it moved data only as entire volumes. The VSP now moves data in pages of just 42MB. Choosing a tiering page size is a trade-off, because if they are too small, they result in too many data movements and a processing overhead that defeats the purpose of the tiering. HDS says it chose 42MB after considering disk sector sizes across different RAID groups.

Like other systems, the VSP decides when to move data between tiers by considering the frequency of data access, and time since last access. Customers can define up to three storage tiers, and for each tier they can choose both RAID levels and physical drive type, including flash. Also like other systems, VSP tiering can be scheduled to be inactive during periods such as weekends, or overnight backups. If it were left running during those times, it would re-locate data to suit those conditions, so that when normal working resumed on a Monday morning, data would no longer be distributed appropriately. Storage pools can be exempted from the tiering system, so that data volumes in those pools are locked into a specific tier.

Management

The VSP is managed using a suite of software modules collectively labeled the Hitachi Command Suite. HCS was updated at the same time that the VSP was launched, and Hitachi says that it is now entirely GUI-driven, and shares interfaces across modules in order to simplify training and operations. Among a number of major changes were the introduction of task management with scheduling for multi-thread operations, and greater access to configuration and tier information.

Alongside functions such as data replication, thin provisioning, multi-pathing and end-to-end performance monitoring, the suite includes options for higher-level functions such as storage capacity planning and application-level service monitoring. The suite can be used to manage not just the VSP and any third-party storage attached to it, but also the full range of HDS storage products, including the mid-range AMS disk array, and HDS' Content Platform object-oriented storage system.

HDS' overall direction with HCS matches a wider trend in the storage industry, which has seen a major reduction in the time taken to manage storage systems, via automation of common tasks. Like other vendors, HDS says that, initially, customers were distrustful of automated decision making about storage systems, and often opted for manual confirmation of those decisions. Acceptance of automation has however been increasing steadily, in part because of limited resources, and the relentless growth of data volumes that has driven up administrative overheads.

Multi-tenancy, or the ability to share a VSP securely among multiple applications, is enabled by virtualized front-end ports, cache partitioning across up to 32 virtual storage machines, and controls on administrators' rights that are based on resource groups. The VSP also promotes multi-tenancy by allowing service levels to be defined for specific data volumes, and issuing alerts if those levels are not met.

© Ovum. This Technology Audit is a licensed product and is not to be photocopied

Integration with server virtualization

The VSP ships with a plug-in for VMware's vCenter management software, which allows VMware administrators to see which storage volumes are allocated to which virtual servers, and to allocate new volumes or resize existing ones. The device is also integrated with VMware's Site Recovery Manager for automated failover from one data center to another. Finally, it is integrated with VMware's recently released VAAI interface, which allows virtual servers to offload processing work to a storage array and so significantly improves overall application performance.

Although the VSP is certified by Citrix for use with XenServer, HDS currently has no plans to step up to a higher level of XenServer integration using Citrix' equivalent of VAAI. Microsoft's Hyper-V virtualization software does not yet offer an equivalent to VAAI, but HDS is currently certifying the VSP for use with Hyper-V.

Unified with both NAS and object storage

File-level access to data held in the VSP can be achieved by attaching a NAS head or gateway to the VSP. Since 2006 HDS has been re-branding BlueArc gateways to do this for HDS' USP arrays, and the company says that just over a quarter of its installed base uses the devices. Now, HDS is offering the re-branded gateways for use with the VSP. BlueArc is a successful supplier of very high performance NAS devices, which are powered by unusual specialist silicon, and can be scaled from two to eight controllers. The BlueArc gateways can handle up to 16PB in a single namespace, hugely simplifying the management of large volumes of file-level data.

Taking the unification theme one step further, the VSP can provide object-based data storage for long-term archiving of file-level data, or data that must comply with data retention regulation. This is done by putting a part of the VSP's capacity under the control of a blade that runs HDS' Content Platform object-based management system. As well as protecting data so securely that backup is no longer required, HCP meets regulatory requirements with features such as the ability to define data retention time periods, digital data shredding, deletion logs, the use of digital signatures to prove data authenticity, and a content and metadata search engine that supports 370 file formats and 77 languages.

Footprint and power consumption

Unlike its predecessor, the VSP is built around standard 19-inch racks. The base configuration is two 19-inch × 42U modules, which can be expanded to a total of six modules – two controller modules, and four storage modules – to reach the full internal capacity of the VSP.

The VSP has several features that allow it to deliver greater performance and denser capacity than its USP V predecessor. The most important of these is its use of 2.5-inch SAS disk drives, which for a given performance consume less space and power than the 3.5-inch Fibre Channel drives used in the USP V. HDS says that the 2.5-inch drives allow most customers to achieve around 150TB to 170TB capacity per rack.

© Ovum. This Technology Audit is a licensed product and is not to be photocopied

The overall redesign of the device also contributes to lower power consumption, as does a change from bottom-to-top to front-to-back cooling airflow. The latter allows the VSP to be used in datacenters with alternating hot and cold aisles, for greater cooling efficiency and lower load on data center air conditioning systems.

Overall, HDS says that compared to the USP, the VSP gives up to 105% more IOPs per square foot, and 159% more IOPS per watt of consumed electricity. For capacity, the VSP gives up to 40% more TB per square foot and 89% more TB per watt. Overall, the device requires 24% less air conditioning load per square foot. These numbers are all obviously dependent on configuration,

PRODUCT STRATEGY

The VSP is developed and manufactured by HDS' parent company, Hitachi. Founded in 1989 and based in California, HDS is responsible for the sale and support of all of Hitachi's storage products outside of Japan. Since the VSP was launched late in 2010, HDS has shipped the device to around 300 customers worldwide. That adds to the company's existing customer base of over 7,500 customers, which include 44 of the top 50 Fortune Global 500 companies.

The VSP was the first major update to Hitachi's flag-ship storage device since the USP V appeared in 2007. That is about typical for the length of time between significant updates to the device. Ovum fully expects that because of the size of the installed customer base, HDS will continue to develop the VSP to this pattern for many years yet.

For several years before Sun Microsystems was acquired by Oracle, Sun sold Hitachi's high-end storage re-branded as the Sun StorageTek 9000 series of disk arrays. Following Sun’s merger with Oracle, this arrangement was cancelled, which closed off a route to market for Hitachi. However, HDS says that because Oracle-Sun now offers no high-end storage products, former Sun customers are turning to Hitachi when they replace their existing 9000-series devices. Recent industry estimates of storage shipments support this claim, showing Hitachi's market share increasing slightly to 9% of total for external disk arrays at the end of 2010.

Hewlett-Packard has also enjoyed a reselling arrangement with Hitachi for several years, and currently sells versions of both the USP V and the VSP, renamed as the HP P9000 and P9500 respectively. HP’s recent purchase of 3PAR, a high-profile and successful mid-range to high-end storage supplier may eventually change this situation. 3PAR’s products have the potential to be developed into full competitors to the VSP, which would allow HP to present a portfolio based entirely on its own products, and end its reselling deal with Hitachi. Ovum believes that this is very likely to happen, but not for some while yet because of the development work required to close the functional gaps between 3PAR’s devices and the VSP.

© Ovum. This Technology Audit is a licensed product and is not to be photocopied

IMPLEMENTATION

HDS sells the VSP directly, and via a network of channel partners. Reselling partners are supported by a program with bronze, silver, gold, and platinum certifications, and offer development, consulting and implementation services. HDS itself provides a full range of assessment, design, implementation, and remote and local management services. Assessment services include the option to calculate costs and reconcile them with business needs, while a data classification service inventories and classifies data according to its value and compliance requirements. Other services include data migration, remote copy design, implementation of tiered storage and thin provisioning, and compliant data destruction. Services typically account for 25% of overall implementation costs.

The smallest VSP implementation involves a single 19-inch rack, costing about $200,000. Payment methods include outright purchase and leasing. Training is either self-paced via the Web, or via on-site instructors. HDS' Certified Professional program includes a Storage Manager track that is intended for customers, and covers planning and architecture as well as day-to-day operations.

Implementation times vary according to the number of applications and servers, and storage network connections and complexity. Basic implementation is typically completed within a few days. Usually this involves a single consultant from HDS or one of its partners, and requires no involvement from the customer.

Software licenses typically account for 25% of overall implementation costs. The HCS management suite consists of the HDS Basic Operating System, which comprise management tools and partitioning and virtualization functions, and a set of optional modules. All software can be bought on a permanent, subscription or fixed-term basis, with discounts for purchases of multiple modules. Optional areas include the virtualization external disk arrays, and automatic tiering. Prices are based on the number of VSD director pairs running in the device and the total usable capacity of the array. Replication software is priced according to the amount of data being replicated. A three-year standard warranty provides 24×7 support with a two-hour maximum for initial response, and four hours for an onsite visit within covered geographical regions. An elevated Warranty Plus guarantees initial response within 30 minutes, customized onsite response time, with additional training and consulting, and access to capacity planning and performance tools.

The VSP supports the following host operating systems:

- IBM z/OS and z/OS.e, OS/390, z/VM and VM/ESA, zVSE, VSE/ESA, MVS/XA, MVS/ESA,
- IBM AIX and Linux for S/390, and zSeries
- HP Tru64 Unix, HP-UX, and Open VMS
- Microsoft Windows Server 2000, 2003, 2008
- Novell NetWare, SUSE Linux
- Red Hat Enterprise Linux
- Oracle Solaris.

© Ovum. This Technology Audit is a licensed product and is not to be photocopied

Table 1: Contact Details	
Hitachi Data Systems 750 Central Expressway Santa Clara California 95050 USA Tel: +1 408 970 1000 Fax: +1 408 727 8036 www.hds.com	**Hitachi Data Systems** Sefton Park Stoke Poges Buckinghamshire SL2 4HD UK Tel: +44 (0)1753 618 000 Email: info.uk@hds.com www.hds.com/uk
Source: Ovum	OVUM

Ovum's Knowledge Centers are new premium services offering the entire suite of Ovum information in fully interactive formats.
To find out more about Knowledge Centers and our research, contact us:

Ovum Europe
119 Farringdon Road
London, EC1R 3DA
United Kingdom
t: +44 (0)20 7551 9000
f: +44 (0)20 7551 9090/1
e: info@ovum.com

Ovum Australia
Level 5, 459 Little Collins Street
Melbourne 3000
Australia
t: +61 (0)3 9601 6700
f: +61 (0)3 9670 8300
e: info@ovum.com

Ovum New York
245 Fifth Avenue, 4th Floor
New York, NY 10016
United States
t: +1 212 652 5302
f: +1 212 202 4684
e: info@ovum.com

All Rights Reserved

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior permission of the publisher, Ovum Europe Limited. Whilst every care is taken to ensure the accuracy of the information contained in this material, the facts, estimates and opinions stated are based on information and sources which, while we believe them to be reliable, are not guaranteed. In particular, it should not be relied upon as the sole source of reference in relation to the subject matter. No liability can be accepted by Ovum Europe Limited, its directors or employees for any loss occasioned to any person or entity acting or failing to act as a result of anything contained in or omitted from the content of this material, or our conclusions as stated. The findings are Ovum's current opinions; they are subject to change without notice. Ovum has no obligation to update or amend the research or to let anyone know if our opinions change materially.

© Ovum. Unauthorised reproduction prohibited

This report is a licensed product and is not to be reproduced without prior permission.