



Hitachi enters the high-performance NAS market

Carl Greiner: CXG@ovum.com

January 2007





Hitachi enters the high-performance NAS market

Hitachi Data Systems (HDS), a wholly owned subsidiary of Hitachi, announced a five-year, worldwide OEM agreement – including a private investment – with BlueArc Corporation. This will allow for the immediate availability of the Hitachi High-performance NAS Platform, powered by BlueArc. This agreement provides HDS with highly scalable, reliable and proven file-intensive NAS gateways, with advanced file virtualization capabilities that, we believe, will be integrated with HDS's existing block virtualization and management capabilities over time. Carl Greiner says that this undertaking expands HDS's capabilities and will enable enterprises to fully consolidate their ever-expanding NAS solutions (file servers/appliances) and simplify the supporting infrastructure, increasing productivity while reducing management complexity and expenses.

When performance and scale matter

The Hitachi High-performance NAS Platform, powered by BlueArc, resets the NAS gateway capability bar in much the same way as the HDS Universal Storage Platform, Adaptable Modular Storage and Network Storage Controller solutions reset SAN-based scalability, performance and virtualization. As HDS evolves, the integration of its file- and block-based solutions, including common and comprehensive management with the HiCommand Suite, the convergence of SAN and NAS will become a reality for the most demanding and diverse storage environments.

The challenge

File serving over Ethernet networks using network attached storage (NAS) appliances or file servers is ubiquitous in today's data centers. However, as usage has grown and performance characteristics have dynamically changed, IT organizations have faced multiple challenges.

- **Manageability** – ever-changing user/business needs and ballooning data storage requirements with increasing compliance requirements have driven most enterprises to deploy multiple file servers or NAS appliances, introducing management complexity and overheads in maintaining redundant copies, including protection/compliance data.
- **Scalability** – the continued proliferation of file servers/appliances is creating complexity within the storage infrastructure that requires a consistently scalable solution (both in terabytes and performance) with a unified directory structure, enabling a dynamic environment that is not just for storage of files, but a true



consolidation solution that can optimize and converge NAS and SAN environments.

- **Availability** – as the requirements for files become more critical, with today's businesses demanding near-instant access to data, high availability and data protection capability have become necessities for most NAS deployments, as have non-disruptive maintenance and upgrades. No longer are these requirements only for SAN solutions, as common requirements are evolving for all data, be it block or file data.

A testament to flexibility and ease of deployment of NAS solutions is the continued robust growth rate of NAS solutions within data centers, often exceeding the terabyte growth of SANs. However, as file storage capacity requirements increase, the problem of managing multiple NAS and other file storage devices becomes pronounced. Moreover, multiple NAS devices can be implemented on a single network to meet increasing capacity requirements, as each device most often exists as an island of storage, making it nearly impossible to consistently manage all the devices as a single consolidated resource. The result is that it becomes difficult to optimize storage investments by systematically and intelligently placing files on specific devices based on the files' varying needs for performance, security and availability.

Fortunately, file storage virtualization technologies (similar to block virtualization) have been developed to resolve these issues. The goal is to allow all NAS devices and file storage systems to be effectively consolidated and viewed by the end user and applications as a single integrated pool, with immediate access to files wherever they are located. In addition, the virtualization solution needs to cost-effectively and non-disruptively migrate files across multiple storage resources, transparently mapping end-user or application I/O requests to those new locations. Thus, virtualization can ensure consistent and continuous file access, even as files are migrated to the various storage devices required to meet business and management needs.

HDS expands its NAS portfolio

Realizing that there is a growing requirement for high-performance NAS solutions with robust virtualization and scaling capability, HDS augmented its NAS blade solutions with the Hitachi High-performance NAS Platform (gateway), which addresses performance-driven applications as well as high-performance computing storage requirements. The platform utilizes both the modular, highly parallel, hardware-accelerated BlueArc silicon pipeline architecture and its function-rich software. The platform is targeted at large enterprise NAS and file server consolidations, and the performance and scalability requirements attributed to file-intensive applications within such industries as life sciences, Internet services, entertainment, electronic engineering, government and education.



The hardware

The Hitachi high-performance NAS hardware is currently capable of managing up to 512 terabytes (TB) of data in a single storage pool and supports a cluster name space for a unified directory structure, with global access to data for common Internet file system (CIFS) and network file system (NFS) clients through any node within the cluster.

Hardware features include:

- hardware accelerator supporting up to 800MB throughput and over 200,000 SpecSFS IOPS in a dual node configuration; 600,000 IOPS in real world testing, according to BlueArc
- scalable file system with up to 512TB in a single storage pool
- up to 4 million files per directory.

The modular chassis design currently supports 40Gbit/s total throughput in a passive back-plane design with no active components, and creates a high-availability platform. The chassis employs four physical modules – the Network Interface Module (NIM), two file system modules (FSA and FSB), and the Storage Interface Module (SIM). Each module is capable of delivering 5Gbit/s throughput. Each module has clear responsibilities and typically operates independently from the others, although FSA and FSB modules do have a co-operative relationship.

Network Interface Module (NIM)

This module is responsible for:

- high-performance Gigabit Ethernet connectivity
- hardware processing of protocols
- OSI layers 1–4
- out-of-band management access.

The NIM is responsible for all Ethernet-facing I/O functions corresponding to OSI layer 1–4. Functions include handling Ethernet and Jumbo Ethernet frames up to 9000 bytes, ARP, IP and routing, TCP and user datagram protocol (UDP).

File system modules (FSA and FSB)

These modules are responsible for:

- advanced features
- OSI layers 5, 6, and 7; NFS, CIFS, iSCSI and NDMP protocols
- security and authentication
- object store layer
- file system attribute caching
- metadata cache management



- NVRAM logging.

The two file system modules work collaboratively, with the FSB board handling data movement and the FSA handling data management. Examples of processes managed by the FSA module include snapshots, quotas, file and directory locking. The FSA module sends instructions to the FSB module, which handles the data control and movement required for the specified tasks.

Storage Interface Module (SIM)

This module is responsible for:

- fibre channel processing
- SCSI command processing
- sector cache management
- parallel RAID striping
- cluster interconnect
- NVRAM mirroring.

The SIM module has two distinct responsibilities. The first is the handling and management of raw data on the SAN storage back-end. The second is for the high availability features of the SAN and cluster interconnect (when configured in a cluster). The SIM provides the redundant back-end SAN connection to the storage pool using 4 Gigabit channel ports.

The software

The supporting software for the Hitachi High-performance NAS Platform enables virtualization, which administrators can employ to partition and allocate storage, protects data with local and remote replication and anti-virus capabilities, and manages data via snapshot, query and migration features. Key features include:

- **cluster name space** – creates a unified directory structure across storage pools and servers. Multiple file systems can be linked and appear under a single common root, and both CIFS and NFS clients obtain global access through any node in the cluster. This enables simpler provisioning, better performance and consistent data protection
- **virtual file system** – supporting CIFS and NFS protocols deliver the flexibility and capability to handle very large files and millions of small files, as well as several concurrent file system operations, without affecting performance
- **virtual volumes** – enables storage administrators to quickly allocate storage required for different applications or users. Virtual volumes mask the complexity of the underlying physical disks while automatically reflecting capacity changes from clients. More than 4 million files are supported per directory and up to 1,024 snapshots per file system
- **virtual servers** – enables partitioning of storage resources, allows server consolidation, and provides multi-protocol support. When use patterns change or



spikes in I/O demand occur, administrators can rapidly balance workloads. They can also create up to eight logical servers within the same management framework, easily co-ordinating throughput by dedicating ports and separate IP addresses to virtual servers

- **multi-protocol support** – for both block and file-level application data within a single system. Included are CIFS, NFS, file transfer protocol (FTP), UDP, TCP, network data management protocol (NDMP) and iSCSI (supports up to 8,192 logical unit numbers – LUNs – on each system)
- **local and remote replication** – asynchronous replication functionality for both block data and files
- **support for anti-virus scanning** – in network data-sharing environments offering full integrated support for Symantec, McAfee, Sophos and Trend Micro
- **snapshots and quotas** – enables high-speed snaps or point-in-time copies (not full copies on writes) for rapid data back-up, minimizing multiple duplicates or copies of data. Hard and soft quotas, limiting available capacity or the number of files stored, can be set for individual users, virtual volumes or directories
- **NAS data migration feature** – allows automated NAS data migration from one storage tier to another based on predefined policies, improving storage utilization and simplifying data management.

Client re-use disclaimer

- This is a verbatim reproduction of independent material that has previously been published by Ovum within the last 6 months
- Ovum operates under an Independence Charter. For full details please see www.ovum.com/about/charter.asp
- Ovum may have been paid by the client for the right to re-use the material
- Ovum may have a deal with the client to supply research or consultancy. However, no other relationship exists between the 2 companies (e.g. shareholdings, loans, non-executive directorships etc)
- Ovum does not endorse companies or their products
- While we take every care 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 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
- This material is the copyright of Ovum Europe Ltd.



Ovum does not endorse companies or their products. Ovum operates under an Independence Charter. For full details please see www.ovum.com/about/charter.asp.

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 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.