Hitachi NAS Storage for Unstructured Data

Silverton Consulting, Inc. StorInt™ Briefing

Silverton Consulting
Strategy, Storage & Systems

Introduction

Hitachi Vantara™ has a long and rich tradition of investment in and enhancement of NAS systems and solutions dating back to an acquisition completed ~10 years ago. Recently, Hitachi released a new series of Hitachi™ network-attached storage (HNAS) hardware and software.

The HNAS system, known for its high performance and scalability, benefits from its state-of-the-art object file system and its unique field-programmable gate array (FPGA) offload engine hardware architecture, which is used to accelerate compute-intensive tasks. Hitachi Vantara has updated this FPGA hardware and NAS software in its latest HNAS systems.

HNAS systems come in NAS gateway configurations with attached Hitachi Virtual Storage Platform (VSP). Here, we'll discuss the significant scalability, performance and enterprise-class functionality offered by HNAS systems, along with the benefits that HNAS as a NAS gateway solution provides coming from a leading systems vendor today.

HNAS 5000 file gateway

HNAS Platform 5000 systems are unstructured file, object and block storage solutions that consist of the following:

- **HNAS 5200**, which can scale out a system with up to 40 NAS nodes, supporting a maximum of 2 trillion objects per namespace with up to 192 GB/s of (read) aggregate throughput.

- **HNAS 5300**, which can scale out a system with up to 80 NAS nodes, supporting a maximum of 4 trillion objects namespace with up to 384 GB/s of (read) aggregate throughput.

Customers can use a single, cluster-wide namespace across 8 HNAS systems. Moreover, HNAS can take advantage of multiple VSP storage systems for tiered data storage pools. Multiple HNAS systems can even share the same VSP storage pool. With up to 10 HNAS cluster-wide namespaces under management, older and newer versions of HNAS gateway hardware can co-exist in a virtual cluster and access back-end storage pool(s).

HNAS 5000 systems support up to 16 million/130 billion objects or files per directory/filesystem. This capability is especially useful for customers with large numbers of small files. While filesystem physical capacity is limited to 1PB on-premises, in practice with its tiering to the cloud feature, filesystem capacity is virtually unlimited.

Each HNAS node can support a maximum of 64K concurrent TCP client active connections. These file share and client connection capabilities exceed by 20x those offered by most competitive solutions.

In addition, all HNAS 5000 nodes come with multiple FC ports to support Network Data Management Protocol (NDMP) FC VSP storage access. NDMP FC provides faster backups/restores of HNAS data without adding to LAN traffic.

HNAS is the only NAS gateway left on the market today from major system vendors. As such, it's worth discussing the differences between integrated NAS and gateway NAS solutions:

- **Integrated NAS solutions** combine the storage back end with the NAS front end into a single hardware appliance. Most integrated systems support both file and block storage. With integrated systems, it's almost impossible to scale the front end independently from the back end. This limitation poses challenges when the front end or back end starts to run out of performance while its counterpart is performing adequately and can often lead to hardware overprovisioning.
- **Gateway NAS solutions** have back-end storage and a separate front-end NAS gateway to that storage. NAS gateway systems support both block and file data protocols and enable independent scaling and/or system refresh of the front end, the back end or both, depending on maintenance requirement or performance bottleneck. As such, system performance can be tailored and optimized to support customer requirements. Both NAS systems and backend storage technology refreshes can be performed non-disruptively, a huge benefit to organization with limited maintenance windows.

For example, an HNAS 5200 system can be upgraded to HNAS 5300 or next-generation HNAS gateway technology to speed up front-end processing without changing or migrating back-end VSP data storage (as discussed below). An HNAS 5200 or 5300 system attached to VSP hybrid flash storage can also be upgraded to use VSP all-flash storage without replacing the HNAS gateway node hardware.

Hitachi HNAS 5000 hardware

Each HNAS 5000 node uses dual Intel scalable Xeon quad-core CPUs with 64GB of system memory. HNAS 5000 series nodes also include NVDIMMs (non-volatile memory) used to hold data for power interruption offload. The 64GB of system memory doesn't include NVDIMM memory or the HNAS system's metadata cache, data cache or other external data structure information. As a result, an HNAS 5000 node has a total of more than 176GB of memory. Each HNAS 5000 node comes with four FPGAs for offload processing.

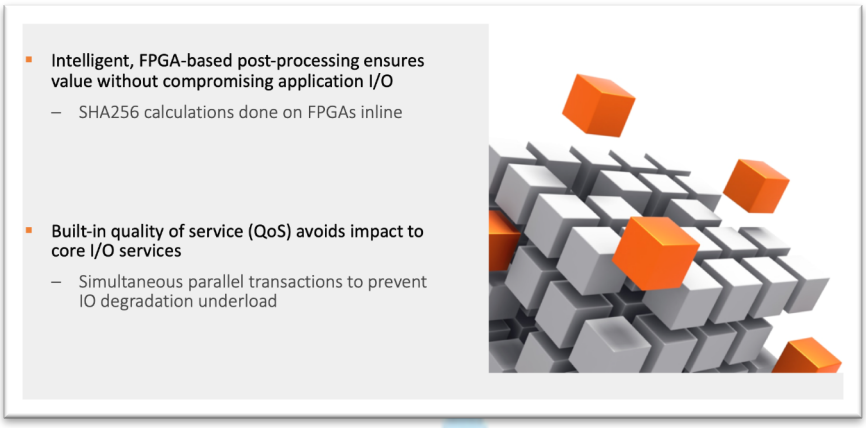

As noted above, all HNAS systems use FPGA hardware to offload compute-intensive, time-consuming work for NAS IO. Such compute offloading is pervasive in the system and occurs almost anywhere it can speed up IO processing.

The FPGAs manage system data paths so data transfers can take place in an accelerated and parallel fashion. They also provide deduplication/cryptographic hashing, support caching and free space management, help the system perform snapshots/cloning and support other time-critical file system operations.

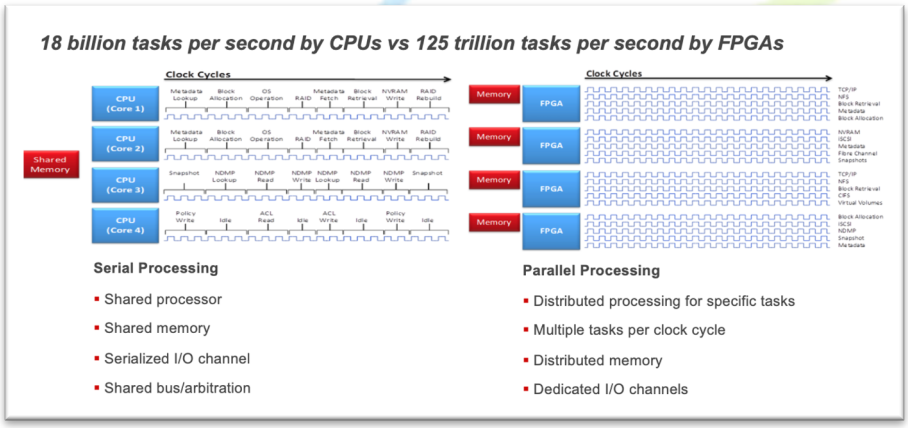

For instance, HNAS series deduplication occurs post-process to reduce impacts on IO performance. An FPGA performs most of the extensive cryptographic hashing calculations required for deduplication of customer data, which further speeds the process.

Cyclic redundancy check (CRC) codes are calculated by an FPGA to protect system metadata and back-end data. FPGAs are also used periodically to help validate back-end data using CRC codes.

Silverton Consulting
Strategy, Storage & Systems

FPGA unique and proprietary offload processing allows HNAS systems to support larger file systems, trillions of objects per system and more file systems in a configuration.

A two-node HNAS 5200 system has 1.5X more streaming throughput than previous-generation HNAS systems, and a two-node HNAS 5300 system has 2X more streaming throughput than previous-generation HNAS systems.

Hitachi HNAS 5000 software

The HNAS 5000 series has increased the efficiency and performance of small file IO over prior-generation systems[1]. The challenges with small file IO are that they often leave a lot of empty space in storage blocks and often require a lot of metadata operations. The latest generation of HNAS software has been redesigned to be able to pack multiple metadata structures into one metadata block and, similarly, pack more than one small file into one data block. This redesign helps with the following:

- **Small file storage efficiency/density** – Small file metadata and data are now able to use fewer physical blocks, resulting in up to 28% capacity improvement for small file repositories.
- **Small file IO performance** – More file data and metadata can fit into a single data block. Similarly, writing small file data with file packing can reduce the number of physical writes to back-end storage by up to 20%. These and other optimizations for small files can help reduce small file IO latency by 5X and increase the number of concurrent small file data users by up to 25%.

Next-gen applications like data analytics and AI/ML/DL applications predominantly process small files. Any improvements in small file space efficiency and performance can make these next-gen applications work much faster and more efficiently on HNAS systems.

Nevertheless, there's often a tradeoff between small file and large file IO performance. It is important to note that large file streaming also performs very well with HNAS 5000 systems.

The original HNAS system was designed for scalable streaming performance, enabling large file size repositories to be read/processed extremely quickly. This kind of streaming IO performance is especially important for high-performance computing (HPC) IO activities, such as oil and gas seismic analysis and media and entertainment streaming and transcoding. HNAS streaming performance also scales linearly as HNAS nodes are added to the cluster. Server Message Block (SMB) performance can be up to 4.8GB/s for a single HNAS 5300 node and up to

[1] Available post GA feature support/enhancement; contact your Hitachi representative for more details

38.4GB/sec for an eight-node HNAS 5300 cluster-wide namespace. Similar speeds are available with Network File System (NFS) protocols.

HNAS software provides **inter-protocol locking**. This feature allows files to be accessed by both NFS and SMB protocols while ensuring that one protocol does not corrupt files while the other protocol is accessing those files.

As discussed above, HNAS offers deduplication while also supporting storage pool and file system **thin provisioning**. Systems that use deduplication and thin provisioning allow customers to store more data in less physical storage, significantly reducing their data footprint. HNAS also offers a **high-performance deduplication option** that can further speed up data deduplication.

Furthermore, HNAS provides **multi-tenancy** through the use of **Enterprise Virtual Servers (EVSs)**. A base HNAS system offers up to 64 EVS, but additional licenses can be purchased in order to use up to 128 EVS on HNAS 5200 and 256 EVS on HNAS 5300 systems[2]. Each EVS offers separate security/administration context and individual routing tables with potentially overlapping IP addresses with parallel IO access across HNAS 5000 nodes.

An EVS and its services can be migrated across HNAS nodes within an HNAS system as well as across HNAS systems within a cluster that uses the same VSP back-end storage. EVS migration could be used to balance performance and resource consumption or to upgrade HNAS front-end hardware without migrating data or disrupting host IO activity.

HNAS storage also offers a **single namespace** that can aggregate file systems within an EVS and across EVSs within a system or a cluster. This single namespace enables a single mount point for all file systems on a cluster. The single namespace is available for both NFS and SMB protocols.

Transparent data multi-cloud offload is another feature available in HNAS systems. As customer files become less active, data can be archived from HNAS storage to AWS, Azure and IBM clouds, as well as to Hitachi Content Platform (HCP). While the file data payload is migrated, file metadata always remains on HNAS storage. Transparent, multi-cloud offload frees up physical capacity on local HNAS-VSP storage systems by using cloud/object storage while always retaining access to the data.

HNAS also offers local, on-storage, read-only **snapshots** at file system level using **redirect-on-write (RoW)** technology. HNAS provides **read-writeable clones** at the directory level that use RoW technology to create production-level copies of file data.

[2] Available post GA feature support/enhancement; contact your Hitachi representative for more details

All HNAS RoW snapshots are fast and space efficient, enabling customers to take up to 1,024 snapshots/clones per file system as frequently as once per second. With HNAS snapshots, admins can roll back file systems to a point-in-time snapshot with only a few clicks. Read-only snapshots are often used to provide copies of production data for backups. Further, development, testing and QA often use read-writeable clones to gain quick access to full copies of production data.

HNAS snapshots and clones integrate with **Microsoft Volume Shadow Copy Service (VSS)**. VSS integration enables the use of application-consistent backups/copies for file recovery from the Previous Versions tab within Windows Explorer.

In addition, HNAS storage offers extensive support for VMware® environments. This support includes **VAAI** for storage offload; **vCenter plugins** for provisioning and managing NFS datastores; **vRealize Orchestrator connectors** for automated workflows; **vRealize Operations management packs** for health, capacity and performance visibility; **vRealize Log Insight content packs** for log analysis; and Site Recovery Manager/Site Recovery Adapter (**SRM/SRA**) for VMware disaster recovery (DR) automation.

HNAS also provides **Kubernetes (K8s) persistent volume support** using the Hitachi CSI (container storage interface) plugin. The CSI plugin enables K8s containerized workloads to access HNAS storage for persistent volume requirements.

HNAS systems also support the **Red Hat Ansible Automation Platform** to automate HNAS provisioning, application deployment, configuration management and services orchestration for the data center[3]. Ansible can be used to speed up software release cycles and create a more reliable, scalable release process.

HNAS resiliency services

HNAS systems supply high-speed **file/object replication** for DR. As noted, HNAS systems use an object storage back end. With object replication, customers can replicate file objects using a scheduled activity that replicates files changed since last replication.

In addition, HNAS systems offer world-class resiliency with **global-active device (GAD) metro clustering**. With GAD, customers can configure a stretched-cluster, dual-site environment where both sites have active (read-write) access to the same data at the same time. It mirrors writes from one system to the other in a bidirectional manner, and its data is coherently updated across both sites. Failover

[3] Available post GA feature support/enhancement; contact your Hitachi representative for more details

of storage from one site to the other is completely automatic, meaning client systems access replicated data from whichever site has the shortest path. Since GAD is built on Hitachi-Vantara's VSP active-active architecture (vs active-standby) failover time takes a few seconds to execute. Similarly, all failback activities are completely automated once initiated by an administrator.

GAD requires two HNAS and VSP storage systems plus a third quorum (non-storage) system with access to the two sites. It can support two sites that are up to 500km apart using continuous, coherent synchronous replication or mirroring between the sites, and it supplies a 0 RPO (recovery point objective) for storage outages.

GAD services depend on stretched clusters where storage networking (Ethernet and FC) allows concurrent access from hosts/clients to both storage systems across the two sites.

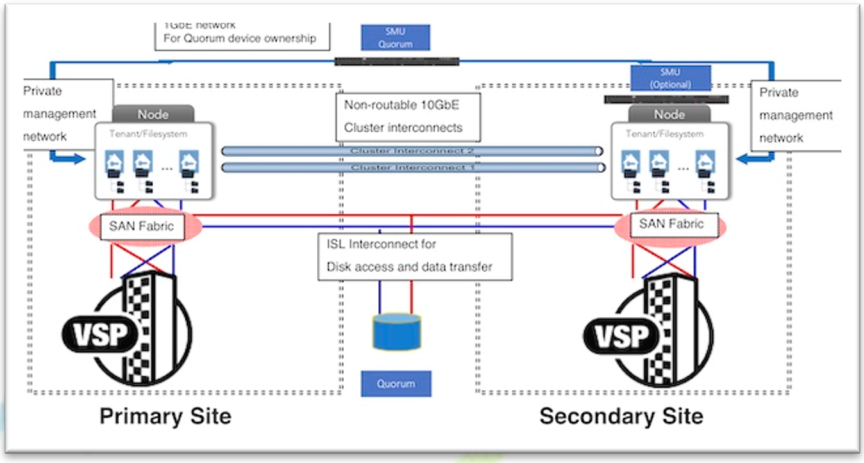

HNAS systems also support a **3DC (data center) replication** solution. This solution uses synchronous replication between the primary and secondary site and then asynchronous replication between the secondary and a tertiary (out-of-region) site. With 3DC replication, customer data centers can undergo a region-wide disaster, and customers can still access data from a third, out-of-region site.

HNAS series operations

HNAS integrates with **Hitachi Ops Center**, which can help improve operational efficiency with AI operations to optimize, orchestrate, plan and protect HNAS infrastructure and data. Ops Center can reduce the time spent analyzing, deploying and managing HNAS and other Hitachi Vantara infrastructure.

HNAS systems offer a **plug-in adapter for Splunk** software that automates the collection of HNAS performance and monitoring data to enable better management and visualization of a data center's storage environment. With Splunk and the adapter, customers can generate graphs, reports, alerts, dashboards and other views of HNAS performance and operations to better manage their storage environment.

HNAS systems also supports NAS management internal to the system. HNAS internal NAS management supports the use of administration panels, CLIs and APIs. With internal NAS management software, admins can configure any of the advanced HNAS services functionality.

Summary

Hitachi Vantara HNAS solutions have always provided state-of-the-art performance, scalability and functionality for unstructured data largely due to the unique functionality of their hardware-based architecture and object storage back end.

The latest HNAS platforms can scale to more NAS capacity and IO performance than previous-generation systems, and the new software improves small file performance and capacity utilization. This front-end scalability combined with a VSP F series all-flash array storage back end can greatly improve performance.

In the end, HNAS unstructured data storage solutions offer the best of all worlds: scalability, performance and enterprise-class functionality.

Silverton Consulting, Inc., is a U.S.-based Storage, Strategy & Systems consulting firm offering products and services to the data storage community.

Silverton Consulting
Strategy, Storage & Systems