

# Hitachi NAS Storage for Unstructured Data

Silverton Consulting, Inc. StorInt™ Briefing



## Introduction

Hitachi Vantara has a long and rich tradition of investment in and enhancement of network attached storage (NAS) systems and solutions dating back to acquisitions completed earlier this century. Recently, Hitachi released a new series of integrated NAS systems: Hitachi Virtual Storage Platform (VSP) N series systems. VSP N series systems are based on Hitachi NAS Platform (HNAS) hardware and software combined with VSP G series storage.

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

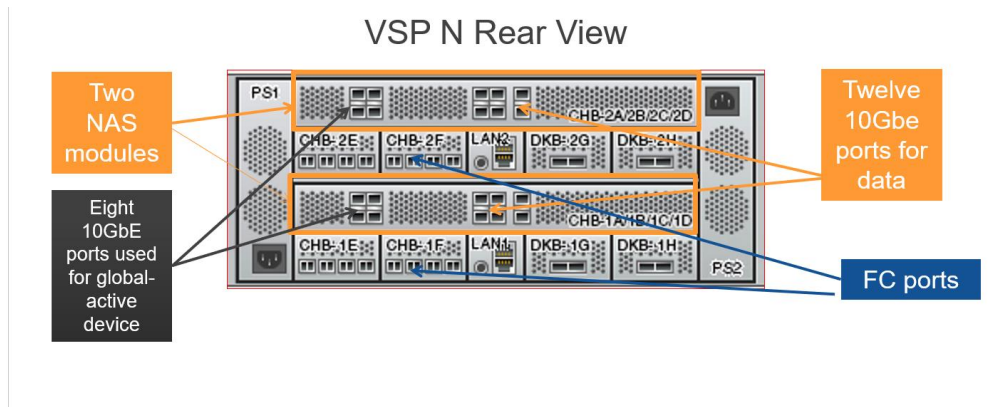
HNAS platforms come as NAS gateways with attached Hitachi VSP G series or VSP F series storage. In contrast, the new Hitachi VSP N series is a completely integrated NAS storage offering that supports unstructured and file-block storage access in a single system package.

This paper will review the new Hitachi VSP N series storage and current HNAS platform functionality and capabilities.

## Hitachi VSP N series storage

The VSP N series systems are dual-NAS module/storage controller systems that consist of the following:

- **VSP N400**, which contains 320GB of DRAM and up to 480 drives for a maximum of 2PB of physical capacity supporting 3GB/s of throughput and 700K concurrent open files;
- **VSP N600**, which contains 448GB of DRAM and up to 720 drives for a maximum of 6.5PB of physical capacity supporting 4GB/s of throughput and 1,178K concurrent open files; and
- **VSP N800**, which contains 704GB of DRAM and up to 1440 drives for a maximum of 6.5PB of physical capacity supporting 8GB/s of throughput and 1,178K concurrent open files;



Each VSP N series system has 12 10GbE file/iSCSI access ports and offers up to 24 FC ports for block storage access. Eight more 10GbE ports are used for NAS metro clustering with Hitachi's global-active device software (see below).

All VSP N series storage can support a maximum of 20K file shares. Each file share maps to a nearly unlimited virtual file system capacity size when combined with multi-cloud offloading. To access file system storage, customers can use up to 128K connections per VSP N series spread across all clients in a data center.

Some unique capabilities are available only with the VSP N series' new generation of FPGA. Specifically, the VSP N series FPGA supports SMB protocol offload functions including SMB Signing and SMB3 Encryption. The VSP N series FPGA also includes more logic gates, which improves performance, and enhances data integrity with features such as T10-DIF (Data Integrity Field) support and SEU (Single Event Upsets – FPGA soft failure) mitigation. With the new FPGA and its object storage backend, all N series storage file systems can support up to 130 billion objects (files), which far exceeds the capabilities offered by competitive solutions. The internal connectivity between the block storage controllers and the integrated NAS modules is facilitated via direct PCIe links. On the HNAS gateway platforms, these interconnects utilize standard FC ports.

Another new feature is Quality of Service (QoS), which allows the VSP N series to prevent highly parallel clients from monopolizing system resources. This feature can be useful when consolidating multiple applications on the same NAS cluster, avoiding impact to mission critical applications sharing the same infrastructure.

VSP N series systems also use hybrid (disk-flash) storage media, which means customers can trade off the lower cost (\$/GB) of disk against the high performance of flash. The optimal use of disk or flash storage is fully automated and orchestrated by **VSP's active-flash tiering**. This functionality moves active data into flash and inactive data down to disk. Alternatively, customers can dedicate flash or disk storage to any specific data, allowing them to choose between automated optimization or hand tuning of their backend storage.

Finally, the VSP N series storage offers a **100% system data availability guarantee** as part of Hitachi Vantara's flash availability program. Hitachi Vantara's commitment to excellence is one of the reasons many Global 100 companies depend on Hitachi Vantara for their mission-critical systems and applications.

## HNAS 4000 Series

Hitachi NAS Platform 4000 series systems for unstructured and file-block storage solutions include:

- **HNAS 4060**, which contains up to 2 NAS nodes, 8-10GbE file/iSCSI ports and up to 8 FC links, supporting a maximum of 8PB of physical capacity with up to 2GB/s of throughput and 148K IOPS;
- **HNAS 4080**, which can scale out to 4 NAS nodes, 16-10GbE file/iSCSI ports and up to 16 FC links, supporting a maximum of 16PB of physical capacity with up to 6GB/s of throughput and 419K IOPS; and
- **HNAS 4100**, which can scale out to 8 NAS nodes, 32-10GbE file/iSCSI ports and up to 32 FC links, supporting a maximum of 32PB of physical capacity with up to 16GB/s of throughput and 1,173K IOPS.

As noted earlier, HNAS systems can connect as a NAS gateway to Hitachi VSP G series hybrid and VSP F series all-flash array storage. Customers can upgrade backend storage separately from front-end HNAS hardware. That is, the HNAS 4060 can be upgraded to HNAS 4100 or next-generation HNAS gateway without the need to change backend VSP storage or scheduling a maintenance window. An HNAS 4080 system attached to VSP G series storage can also be upgraded to use VSP F series storage without changing the HNAS gateway. Customers have great flexibility in upgrading HNAS systems as a result.

All HNAS systems support up to 130 billion objects or files per file system, which is especially useful for customers with large numbers of small files.

Customers can have a maximum of 64K concurrent TCP client connections to the HNAS 4100 and a maximum of 45K connections to HNAS 4080 and HNAS 4060 per NAS node. This number of file shares and client connections exceeds that offered by most competitive solutions.

In addition, HNAS NAS modules have their own FC ports, so they can support both NDMP FC and NDMP IP. VSP N series storage supports only NDMP IP. NDMP offloads data transfer activity from media servers and client networks by using direct connections between storage and backup media. NDMP FC should provide more configuration flexibility and slightly faster backups/restores of HNAS data.

## Hitachi VSP N series and HNAS shared hardware

As noted above, the VSP N series and HNAS series have FPGA differences, including some packaging enhancements in the VSP N series systems. However, VSP N and HNAS hardware share many characteristics and functionality, including superior performance and scalability, which are largely the result of their unique FPGA hardware architecture.

The original rationale for the FPGA was to offload compute-intensive, time-consuming work for the NAS system. This offloading is pervasive in the system and occurs almost anywhere it can speed up IO activity. For example, the FPGA manages data paths so data transfers can take place in an accelerated and parallel fashion.

HNAS and VSP N series deduplication occurs post-process to reduce impacts on IO performance. The FPGA performs the extensive calculations required for deduplication of customer data, which further speeds the process. Cyclic redundancy check (CRC) codes calculated by the FPGA protect system metadata and backend data. The FPGA also periodically helps validate backend data.

- Intelligent, FPGA-based post-processing ensures value without compromising application I/O
  - SHA256 calculations done on FPGAs inline
  
- Built-in quality of service (QoS) avoids impact to core I/O services
  - Simultaneous parallel transactions to prevent IO degradation underload



As noted earlier, the FPGA on the VSP N series performs SMB protocol offload. Similar FPGA functionality is used for Network File System (NFS) protocols. The FPGA also offloads a number of time-critical metadata operations that help speed up NAS data access and support larger file systems, more objects per system and more file systems.

Hitachi Vantara is one of the few major storage vendors that offer a custom designed **Flash Module (FMD)**, which allows for even faster data access and a denser storage footprint for flash data. FMDs offer onboard, hardware-based, inline data compression. Similar to FPGA, hardware compression reduces the data footprint beyond deduplication while minimizing performance impacts. Customers can store more data on FMDs with no impact on IO performance.

## Hitachi VSP N series and HNAS 4000 software

The software for VSP N series and HNAS also shares many similarities.

Both the HNAS and VSP N series storage provide **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 it.

HNAS and VSP N series storage offers storage pool and file system **thin provisioning**. Systems that use deduplication, FMD hardware compression and thin provisioning can significantly reduce customer data footprint, allowing customers to store more data in less physical storage.

Both product families support multi-tenancy for secure user access through the use of **Enterprise Virtual Servers (EVS)** and supply up to 64 EVS per system. Each EVS offers a separate security/administration context and individual routing tables with potentially overlapping IP addresses and parallel IO access across NAS modules. An EVS and its services can be migrated across NAS modules to balance performance and resource consumption. Multiple EVS can even share the same storage pool.

The VSP N series and HNAS storage also offer a **cluster-wide namespace** that can aggregate file systems within an EVS or across EVS. This can provide a single mount point for all file systems on a VSP N series system. This cluster-wide namespace is available on both NFS and SMB protocols.

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

Both systems supply local, on-the-storage **snapshots** at file system level using **redirect on write (RoW)** technology. HNAS and VSP N series RoW snapshots are fast and space efficient, enabling customers to take up to 1024 snapshots per file system as frequently as once per second. With HNAS and VSP N series snapshots, admins can roll back file systems to a point-in-time snapshot with only a few clicks. Snapshots are often used to provide copies of production data for backups.

HNAS and VSP N series snapshots integrate with **Microsoft VSS** services, to allow end users the ability to perform file recovery from NAS snapshots via the previous version tab within Windows Explorer.

HNAS and VSP N series provide **read-writeable directory clones** that also use RoW technology for fast, space-efficient, production-level copies of file data. In this way, development, test and QA can quickly access read-writeable copies of production data.

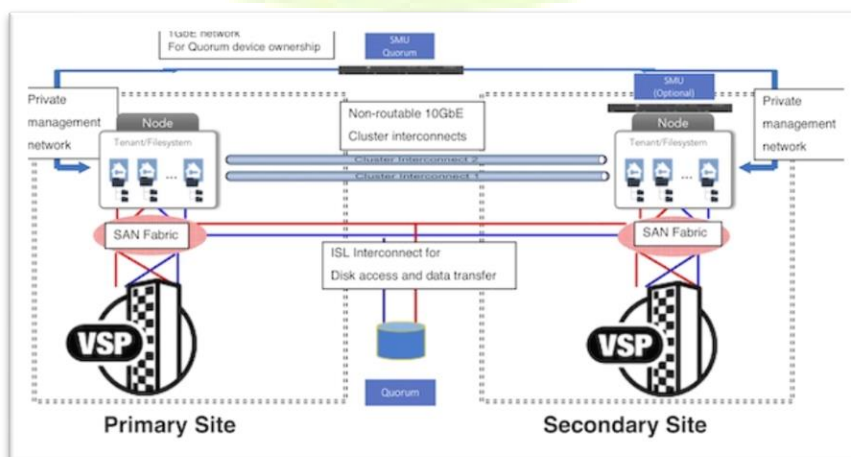
In addition to all of the above, both the VSP N series and HNAS storage offer extensive support for VMware environments. This support includes **VAAI** for storage offload, **vCenter plugin** for provisioning and managing NFS datastores, **vRealize Orchestrator connector** for automated workflows, **vRealize Operations management pack** for health, capacity and performance visibility, **vRealize Log Insight content pack** for log analysis and Site Recovery Manager adapter (**SRM/SRA**) for DR automation.

### HNAS and VSP N series resiliency services

HNAS and VSP N series systems supply high-speed **file/object replication** for disaster recovery (DR). As noted, HNAS and VSP N series systems use an object storage backend. With object replication, customers can replicate file objects using a scheduled activity that replicates files changed since last replication.

Both the HNAS and VSP N series systems offer best-in-class resiliency with **global active-device metro clustering**. With global-active device 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 from one site to the other is completely automatic, meaning client systems access replicated data from whichever site has the shortest path. Similarly, all failback activities are also completely automated once requested by an administrator.

Global-active device requires two VSP N series or HNAS 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 supplies a 0 RPO (recovery point objective) for site/storage outages.



concurrent access from hosts/clients to both storage systems across the two sites.

There are slight differences in how HNAS and VSP N series systems are configured for global-active device, but in all cases customer data is available in both locations. Global-active device services depend on stretched clusters where storage networking (Ethernet and FC) allows for

As noted earlier, it also requires a quorum system located at a third site that has access to the other two sites.

In addition VSP N series and HNAS storage support a **3DC (data center) replication** solution. This uses synchronous replication between the primary and secondary site and then asynchronous replication between the secondary and a tertiary (out of region) site. With the 3DC solution, customers data centers can undergo a region wide disaster and still have their data available at a 3<sup>rd</sup> site, outside the region.

Object replication, global-active device and 3DC replication are only some of the ways to replicate file and block data with HNAS and VSP N series storage. To learn more about data resiliency services for these systems, contact your Hitachi Vantara representative.

### HNAS and VSP N series operations

The two systems also offer a **plug-in adapter for Splunk** software that automates collection of HNAS and VSP N series performance and monitoring data to allow for 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 and VSP N series performance and operations to better manage their storage environment.

Hitachi Vantara's VSP N series and HNAS systems have NAS management software installed. With NAS management software, admins can configure any of the advanced HNAS or VSP N series file services functionality. Both HNAS and VSP N series support the same NAS administration panels, CLIs and APIs.

### Summary

Hitachi's NAS solutions have always provided state-of-the-art performance, scalability and functionality for unstructured data. Their high performance and broad scalability are largely due to the unique functionality of their hardware-based architecture and its object storage backend.

The new Hitachi VSP N series systems follow in this long tradition, introducing a single, integrated, file-block storage solution specifically packaged to meet enterprise unstructured data requirements. The VSP N series merges the best of HNAS and VSP G series storage to provide a unique hybrid unstructured storage solution.

The HNAS systems continue to offer the flexibility of a NAS gateway and a file-block storage solution, allowing customers to upgrade NAS modules and storage separately. HNAS systems can scale to more NAS modules, capacity and IO



performance than the VSP N series alone. The combination of scalability with a VSP F series all-flash array storage backend can improve performance considerably.

In the end, both HNAS and VSP N series storage continue to lead the way in enterprise unstructured data storage solutions that 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.*

