IDC PERSPECTIVE

Hitachi Vantara Brings NVMe to Its Flagship Enterprise Storage Arrays

Eric Burgener

EXECUTIVE SNAPSHOT

FIGURE 1

Executive Snapshot: The New Hitachi VSP 5000 Series Delivers Strong NVMe Performance with a Growth Path to Other NVMe Technologies

In October 2019, Hitachi Vantara announced the availability of its new high-end enterprise Hitachi Virtual Storage Platform (VSP) 5000 Series, as well as a number of other cloud integration and unified management capabilities that will clearly be of interest to enterprises undergoing digital transformation (DX). These enhancements will create strong motivation for existing VSP customers to upgrade and provide many features that will be enticing to noncustomers as well.

Key Takeaways

- The VSP 5000's hardware architecture and operating system software have been optimized for NVMe to drive industry-leading throughput and scalability numbers: up to 21 million IOPS, 69PB of raw storage capacity, and latencies as low as 70 microseconds (even without NVMe over Fabrics [NVMe-oF] support!).
- To enable increased scalability, up to 12 controllers, Hitachi has moved away from the crossbar switch design on its previous enterprise arrays to a new, switched PCIe node interconnect that supports much higher performance and resiliency.
- The VSP 5000 offers an excellent upgrade path from older VSP systems as well as to newer NVMe technologies such as persistent memory (PM), storage class memory (SCM), and NVMe-oF.

Recommended Actions

- Existing Hitachi VSP F/G 1500 customers should already be starting to think about how they will deploy and use NVMe technologies over the next several years.
- Given the expanded scalability of Hitachi Vantara's new high-end enterprise storage platforms, customers may consider extending the depreciation life cycle of newly purchased VSP 5000 Series systems.
- Customers should look to better understand and leverage the new orchestration, automation, and artificial intelligence/machine learning (AI/ML) capabilities in the new Hitachi Ops Center and other management products to help ease management challenges in high data growth environments.

Source: IDC, 2019

November 2019, IDC #US45596019

SITUATION OVERVIEW

As enterprises undergo digital transformation (DX), the transition to a more real-time, data-centric business model often drives the need to modernize information technology (IT) infrastructure. This is generally done to meet increasingly stringent agility, higher performance and availability, and ease of management and scalability requirements while staying within budget constraints. In addition, as part of this transition many enterprises are also deploying next-generation applications (NGAs) that feature new technologies such as artificial intelligence/machine learning (AI/ML) and big data analytics more prominently than ever before in customer-facing and business analytics operations. IT is continuously tasked to do more (and better) with less. Taken together, these new demands present challenges that are very difficult for traditional storage infrastructure to meet. Enterprise storage vendors have been responding to these evolving demands with the introduction of new technologies into their storage platform offerings such as NVMe, software-defined infrastructure, AI/ML, and new, more fault-resilient hardware architectures. These technologies all bolster customers' abilities to keep up with the evolving demands of their own businesses.

NVMe

NVMe is the next-generation storage protocol that will be replacing SCSI in many enterprise storage platforms over the next several years. Relative to SCSI, NVMe offers lower latencies, higher throughput and bandwidth, three orders of magnitude greater parallelism, and access to emerging memory technologies such as persistent memory (PM) and storage class memory (SCM). The infusion of NVMe technology into bespoke storage platforms results in what IDC refers to as NVMe-based all-flash arrays (NAFAs).

While NAFA start-ups were most often pursuing dedicated application environments that required NVMe performance for one or a few workloads, established enterprise storage providers integrated NVMe into their mature, proven flagship enterprise storage platforms. These systems could provide sub-100 microsecond latencies, millions of IOPS of throughput, "six nines" plus availability, and petabytes (PBs) of storage capacity — against which all of their enterprise-class data services (RAID, thin provisioning, inline data reduction, snapshots, encryption, quality of service, replication, etc.) could be selectively applied. While they offered sufficient performance for many of the more real-time-oriented NGAs, these systems offered multitenant management capabilities not available on the start-up NAFA products and were generally purchased for mixed enterprise workload consolidation (which might include one or more workloads that individually required the low NVMe latencies). While performance requirements drove NVMe purchases, these customers were actually most interested in NVMe's ability to maintain consistently predictable storage latencies at scale as additional workloads were added to the platforms (rather than just the lowest possible latency). Based on these two types of customer targets, IDC's NAFA taxonomy refers to "performance-oriented" NAFAs and "data services-oriented" NAFAs.

Several considerations played into customers' "data services-oriented" NAFA purchases. Even if customers did not absolutely require NVMe performance today, many felt they would require that within a year or two and so it made sense to purchase a system today that could deliver that when they were ready for it. Some established vendors continued to offer both SCSI and NVMe versions of their flagship storage platforms, pricing the NVMe version either at par or below where the older SCSI arrays were priced. For customers comfortable with the maturity level of NVMe, buying a NAFA from one of these vendors was an easy decision. In 2018, NAFAs generated $1 billion in revenue (most of which flowed to established enterprise storage players rather than the start-ups), and IDC predicts that by 2021, NAFAs will generate over 50% of external, primary (i.e., performance sensitive, mission critical) storage revenues.

PM and SCM are two of the emerging memory technologies that require NVMe (they cannot be accessed using SCSI). Both PM and SCM are based on the same, persistent solid-state media type (Intel's datacenter Optane is probably the best-known brand of this class of media today) but each uses a different interface type. PM uses a DDR4/DDR5 interface, while SCM uses a block-based NVMe interface. Because of the interface access characteristics, PM delivers higher performance than SCM, albeit at a higher cost. Both of these product types, however, offer near-DRAM performance at dollar-per-gigabyte costs that are closer to NAND flash-based products. There will clearly be more applications developed that will require this type of performance, and support for NVMe in a NAFA positions the customer to easily take advantage of PM and/or SCM as necessary when it is required.

There is one other consideration around NVMe, and that is host connection. SCSI host connections using fiber channel (FC) or Ethernet add latency to the storage access (relative to PCIe SSDs in a server), whereas NVMe over Fabrics (NVMe-oF) host connections can cut the storage network latency to as low as 5 microseconds (pretty much on par with internal PCIe storage access latencies). Those workloads requiring the lowest latencies will require an NVMe-oF host connection as well as an array that features NVMe controllers, devices, and backplanes. For enterprise customers that will be upgrading an existing SCSI-based all-flash array (SAFA) to a NAFA, many of them prefer to connect the NAFA to their existing SCSI-based FC or Ethernet networks, and then upgrade to NVMe-oF in the future once the higher-performance host connection is required. Given their focus on only the most performance-sensitive workloads that require NVMe performance today, the start-up NAFA vendors only offer NVMe-oF host connections, but for established enterprise storage vendors, offering both SCSI and NVMe-oF options, with an easy upgrade to NVMe-oF, is very important.

Software-Defined Infrastructure

As IT infrastructure becomes more software defined, it provides more agility and technology refresh options than more hardware-defined designs. This is clearly the general direction of the industry. While enterprise storage providers still offer some hardware-based functionality on their platforms where it provides performance, availability, efficiency, reliability, and/or cost benefits, they are also moving storage management functionality into software where it provides comparable capabilities at a lower cost. Software-based approaches also make it easier to integrate new, higher-performance hardware as that becomes available. IT infrastructure solutions that leverage standards-based hardware, whether those are CPUs, SSDs, or network adapters, can offer certain advantages to both customers and vendors alike.

Software-defined infrastructure can also make it easier to automate and orchestrate workflows that span different types of IT equipment (servers, storage, networks, etc.) and locations (tiering data between platforms and/or clouds). These two topics (automation and orchestration) are high on IT managers' lists as they look to improve the speed and reliability of operations while at the same time making administrative resources more versatile and increasing the span of administrative control.

AI/ML

While customers are deploying AI/ML-infused big data analytics to improve customer service, uncover new markets and innovative business opportunities, and streamline the efficiency of their operations, vendors are also leveraging the technologies to make their own products and services better. Hybrid cloud is becoming the way IT infrastructure will be built — 40% of customers are already using hybrid clouds in production, while 45% plan to implement it within the next two years — and management in a hybrid cloud environment that combines traditional on-premise infrastructure, private cloud, and public multicloud will make management more complex. Vendors are already starting to turn to AI/ML in two separate locations to help make systems management easier.

First, vendors are adding AI/ML algorithms into arrays themselves that can dynamically optimize a system in real time for better performance, higher availability, or increased efficiency. Generally, an administrator sets a policy, like a performance or availability objective, and the system makes adjustments in real time as the I/O load varies, capacity is expanded, workloads are added, and/or failures occur. This actually improves a system's ability to meet service-level agreements while actually reducing the time administrators need to spend worrying about it. It is also noticeably different from older, more statically defined approaches to systems management that were less flexible and less granular in their responses.

Second, we have seen the introduction of cloud-based predictive analytics platforms that represent the next generation of telemetrics for the enterprise. These systems leverage AI/ML against large, cloud-based information stores that include all data collected from installed base systems over time. Vendors mine this data to drive performance and availability management, risk identification, troubleshooting, upgrade verification, performance and capacity trend analysis, best practice dissemination, data placement optimization (tiering between arrays and/or various clouds), and resource utilization, as well as other benefits. Importantly, vendors are also extending this type of coverage to far more than just storage, which generates only 9.8% of downtime (according to primary research performed by IDC in 1H19). Vendors are starting to cover not only more of their own storage platforms but also servers, networks, applications, and other logical constructs (virtual machines, containers, etc.). The bottom line for customers on these platforms is that they drive meaningful value in a variety of areas that improve the customer experience, drive higher overall infrastructure availability, and help lower costs while creating differentiation between vendors.

Fault Resilience

High availability requirements have been one of the most common reasons that enterprises continue to use SANs even as the use of other options, such as hyperconverged infrastructure (HCI) and public cloud-based storage offerings, grows. HCI is still maturing and proving itself for mission-critical usage in the enterprise, and public cloud offerings generally will not contractually guarantee availability higher than "four nines" (99.99%). As enterprises undergo DX and IT infrastructure becomes ever more critical in driving more data-centric business models, availability requirements are on the increase. IDC's primary research from 2019 indicates that 57.8% of enterprises deem at least 26.0% of their workloads as "mission critical," and 34.6% of enterprises manage their mission-critical workloads to "five nines" or higher levels of availability.

Many high-end enterprise storage arrays have delivered "five nines" of availability for the past several decades, and this is why they are still most often chosen to host a company's most mission-critical workloads. Despite this, there are few vendors that have enough confidence in their solutions to guarantee 100% data availability on their enterprise storage platforms. Customers looking for the most highly available environments may want to consider buying from vendors that do offer that 100% data availability guarantee.

The New Hitachi VSP 5000 Series

In October 2019, Hitachi Vantara made a number of announcements at its Hitachi NEXT 2019 user conference around edge, core, and cloud infrastructure solutions. The introduction of the Hitachi VSP 5000 Series let the company tell a very strong story around next-generation core infrastructure. The VSP 5000 represents a complete architectural refresh of Hitachi's enterprise-class VSP F/G 1500 models and is important for enterprises undergoing DX as it infuses NVMe technology into the vendor's time-proven storage platform. Leveraging NVMe technology in the platform allows Hitachi to

announce some industry-leading metrics that are clearly of interest to enterprises looking for higher performance and more scalable storage: the VSP 5000 can deliver up to 21 million IOPS and supports up to 69PB of storage capacity, both of which represent high watermarks for enterprise-class storage systems shipping today. The system is also capable of delivering storage latencies as low as 70 microseconds – and that's with 32Gb FC and NAND flash–based NVMe SSDs, not NVMe-oF or SCM (which will become available for the platform in 2020).

There have also been some changes to the VSP architecture to enable higher scalability. First, the vendor has expanded the scale-out capability of the array, allowing the dual controller VSP 5100 to be nondisruptively upgraded to larger configurations in the VSP 5500 that can include up to 12 controllers (six node pairs). A new switched PCIe fabric connects all the controllers, replacing the older crossbar switch and supporting significantly greater throughput and parallelism. Each controller also includes field programmable gate arrays (FPGAs) that offload I/O processing from the storage controllers to enable even higher scalability.

Interestingly, the vendor will allow customers to mix media types in the system – NVMe node pairs support NVMe SSDs while SAS node pairs support SAS SSDs and HDDs. The Hitachi custom-designed flash module drives (FMDs) are also supported on SAS controller blocks, but the vendor will be slowly moving away from PCIe-based FMDs as the NVMe standard matures. Host connections on the new systems can be either 32Gb FC or 10GbE, making it easy for current customers to integrate the new system into their existing SANs. The system does, however, provide a future-ready growth path to additional NVMe-based technologies including SCM (e.g., Intel datacenter Optane SSDs), PM, and NVMe-oF options. For NVMe-oF, the vendor intends to support both FC and Ethernet-based transport options.

The new VSP 5000 offers a significantly more streamlined footprint than the prior VSP F/G 1500 systems. In a single 42U datacenter rack, the VSP 5000 yields over six times the capacity density of a VSP F/G 1500 configured for equal performance. This is with an equal number of FC ports (although these are higher-performance 32Gb FC ports rather than the 16Gb FC ports supported by the VSP F/G 1500). An expansion-only rack for the VSP 5000 using SAS (not NVMe) has almost four and a half times the capacity of a VSP F/G 1500 expansion-only rack using FMDs.

Hitachi Vantara has a long history of providing extremely reliable enterprise storage solutions for enterprises' most mission-critical workloads. It was in fact the first enterprise storage vendor to provide a 100% data availability guarantee on its systems, doing so several years before other vendors began to reach this level. With the VSP 5000, Hitachi has retained all the reliability of the prior designs but built additional redundancy into the hardware architecture, which better enables nondisruptive scale-out expansion and technology refresh. These changes make it even more reliable than the previous generation. All models of the VSP 5000 continue to be covered by the vendor's 100% data availability guarantee.

A new version of the Hitachi Storage Virtualization Operating System RF (SVOS RF) that is optimized for NVMe ships with the new systems. Write coalescing and other aspects of how the VSP system writes to persistent storage help maximize solid-state media endurance and make operations more efficient. With SVOS RF 9, the vendor is moving some functions that were performed in hardware before into software, making the system more portable and better able to harness the features of new hardware as it becomes available. The vendor has enhanced its data reduction capabilities as well, supplementing its inline compression and deduplication with post-process algorithms that allow it to guarantee up to 7:1 total efficiency (taking all storage efficiency technologies into account) and 4:1

effective capacity (for just data reduction) ratios. The new 4:1 effective capacity guarantee for compression and deduplication requires no preassessment of the customers' data. All of the other features of SVOS RF, from RAID, thin provisioning, and encryption to snapshots, replication, Global Active Device (stretch clusters) and, of course, mainframe support have been moved forward to SVOS RF 9 as well.

Putting the Hitachi VSP 5000 Announcement into Perspective

The VSP 5000 announcement was important for Hitachi storage customers as it gives them access to NVMe performance today with an NVMe-ready growth path to future NVMe technologies that will not require a forklift upgrade. While Hitachi tends to sell to more conservative customers that highly value reliability, most of its customers are undergoing DX and will require NVMe technologies at some point in the next several years (if they do not already). Key enhancements in both hardware and software move Hitachi and its customers in the right direction with a system architecture supporting much higher parallelism (very important to get the most out of NVMe technology) and a more software-defined offering. The move away from FMDs to off-the-shelf NVMe SSDs also rides an industry trend — we have seen other vendors that used custom PCIe technology early on to deliver NVMe performance well before NVMe was ready for enterprise use, such as IBM, follow a similar strategy as NVMe matures.

The performance numbers of the VSP 5000 are impressive, and this type of industry leadership could give more prospective customers a reason to look at Hitachi as they consider how they will migrate to NVMe over time. The ability to hit storage latencies as low as 70 microseconds without NVMe-oF or SCM highlights the efficiency of Hitachi's architecture, and its long-standing ability to deliver 100% data availability will also turn the heads of prospective customers even as it encourages the VSP installed base to stay within the fold. Although the Hitachi Ops Center and DataOps announcements are not covered in this IDC Perspective, Hitachi has a newly invigorated, AI/ML-driven management strategy to cover edge, core (including the new VSP 5000), and cloud-based infrastructure, and customers looking for a single pane of glass to manage their hybrid cloud environments will want to take a closer look at these Hitachi management platforms as well.

ADVICE FOR THE TECHNOLOGY BUYER

- Enterprise storage technology refresh is the time to consider the transition from SCSI to NVMe for primary storage workloads, and it is a rare enterprise undergoing DX that will not need NVMe performance at some point in the future as it brings on newer, more real-time workloads that leverage technologies such as AI/ML (even if they do not need it today).
- Existing Hitachi VSP enterprise storage customers will find much to like in the new VSP 5000, as well as easy upgrade paths from older systems and to newer NVMe technologies such as PM, SCM, and NVMe-oF. Customers modernizing their IT infrastructure for the hybrid cloud future should take note that Hitachi continues to deliver extremely reliable arrays, has extended its Flash Assurance Program (which includes its 100% data availability guarantee) to these new systems, and has introduced new data reduction technologies (among other enhancements) in SVOS RF 9 that offer data efficiency guarantees as high as 7:1 — all enhancements that deliver an improved customer experience.
- Customers should be considering how their storage vendors implement automation and orchestration solutions, leverage self-optimizing technologies such as AI/ML, and offer unified management for hybrid cloud environments as they seek to optimize IT infrastructures that span both on-premise (traditional, private cloud) and off-premise (public multicloud) locations.

LEARN MORE

Related Research

- *The Evolving Market for Enterprise Storage: Hybrid Cloud, NVMe, and Software Defined Are the Next Wave* (forthcoming)
- *Market Analysis Perspective: Worldwide Enterprise Storage Systems, 2019* (IDC #US45556019, September 2019)
- *Transition to NVMe Should Include a Review of System-Level Power Consumption* (IDC #US45151919, June 2019)

Synopsis

This IDC Perspective reviews new technology developments in enterprise storage and discusses Hitachi's NVMe strategy, as well as highlights the evolution and capabilities of their new NVMe-based system. Hitachi NEXT 2019 represented the largest set of infrastructure announcements from Hitachi Vantara this year, and possibly ever for the vendor. One of the announcements was about the new Hitachi VSP 5000 Series, the vendor's flagship enterprise storage system that now for the first time includes NVMe technology. Readers should note that IDC is publishing a separate IDC Market Note that provides a more general summary of the Hitachi NEXT 2019 announcements as well.

"NVMe-optimized storage platforms provide significant opportunities for customers to cost-effectively consolidate mixed enterprise workloads onto fewer systems," said Eric Burgener, research vice president in the Infrastructure Systems, Platforms and Technologies Group at IDC. "For IT leaders considering this type of consolidation, issues such as performance, availability, and scalability are critical. With the new VSP 5000 Series storage arrays, Hitachi is delivering across all of these metrics, turning in industry-leading high-end throughput on systems that deliver sub-100 microsecond latencies backed by its venerable 100% data availability guarantee."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com