Technical Validation

Hitachi Vantara VSP 5000 Series with Hitachi Accelerated Fabric

Leveraging NVMe-based Storage for Demanding Data-intensive Applications

By Alex Arcilla, Validation Analyst; and Tony Palmer, Senior Validation Analyst

June 2020

This ESG Technical Validation was commissioned by Hitachi Vantara and is distributed under license from ESG.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Contents

ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Introduction

This ESG Technical Validation documents our auditing of performance test results conducted on the Hitachi Vantara Virtual Storage Platform (VSP) 5000. We evaluated how the non-volatile memory express (NVMe)-based platform with Hitachi Accelerated Fabric (HAF) enables the VSP 5000 to achieve high performance and low latency.

Background

ESG research uncovered that the rise of NVMe-based storage has captured organizations' interest, with 56% of respondents seeing the potential to reap benefits such as increased performance for new, more demanding workloads and 48% seeing the potential for storage infrastructure consolidation in order to optimize capital and operational costs (see Figure 1).[1]

Figure 1. What is Driving Use/Interest in NVMe-based Flash Storage

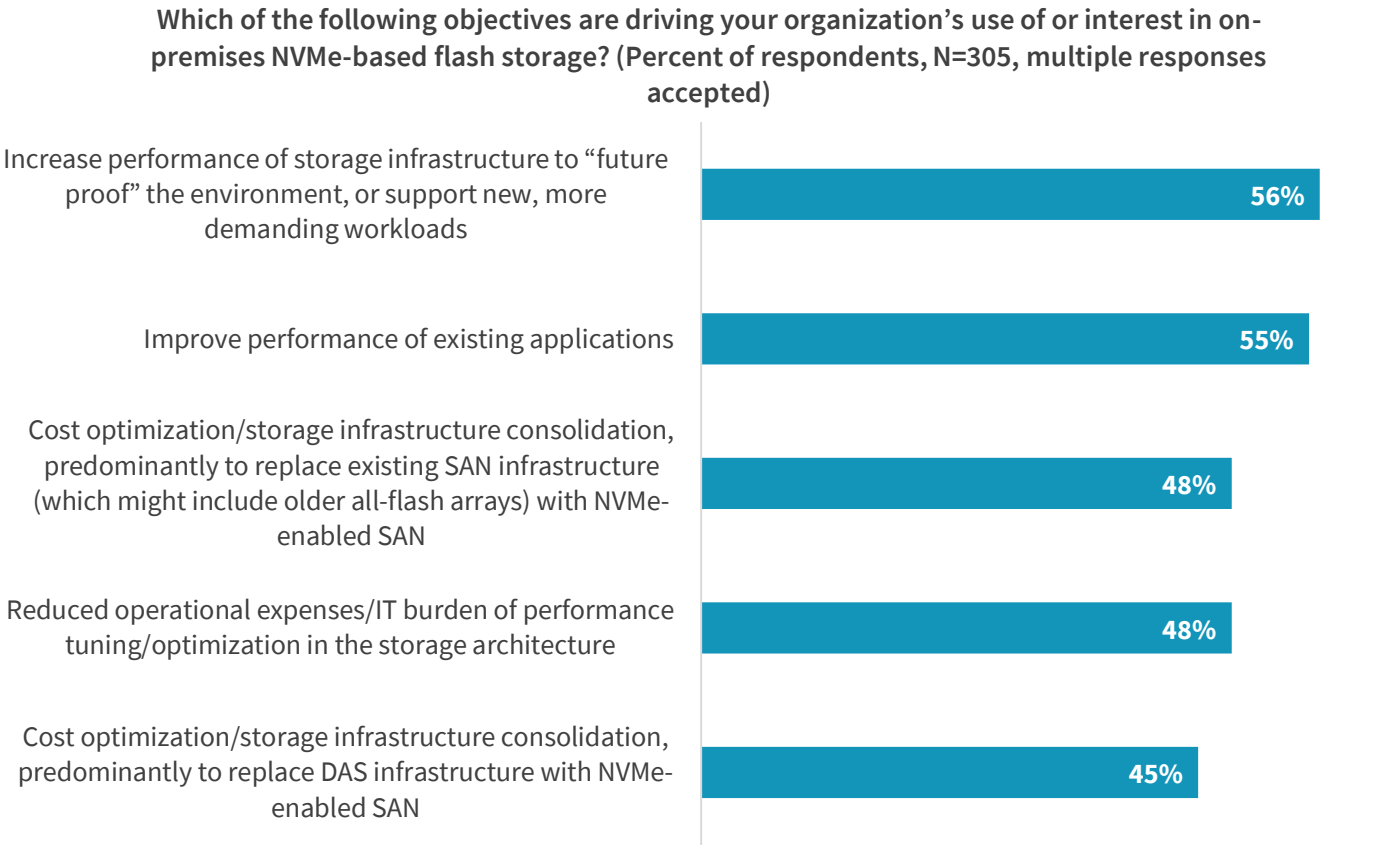

Source: Enterprise Strategy Group

It is no wonder that organizations have been looking more closely at NVMe-based storage. Figure 2 illustrates how storage performance has improved over the past 20 years between traditional hard disk drives (HDDs), solid-state drives (SSDs), and NVMe media, with a 10x performance improvement achieved between HDDs and SSDs, as well as between SSDs and NVMe. The timing of the performance improvement could not be more opportune, as organizations are exploring the use of or developing applications requiring highly performant storage, such as big data analytics, the internet of things (IoT), machine learning, and artificial intelligence.

[1] Source: ESG Master Survey Results, *2019 Data Storage Trends*, November 2019. All ESG research references and charts in this technical review were taken from this master survey results set, unless otherwise noted.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Figure 2. Historical Storage Media Performance

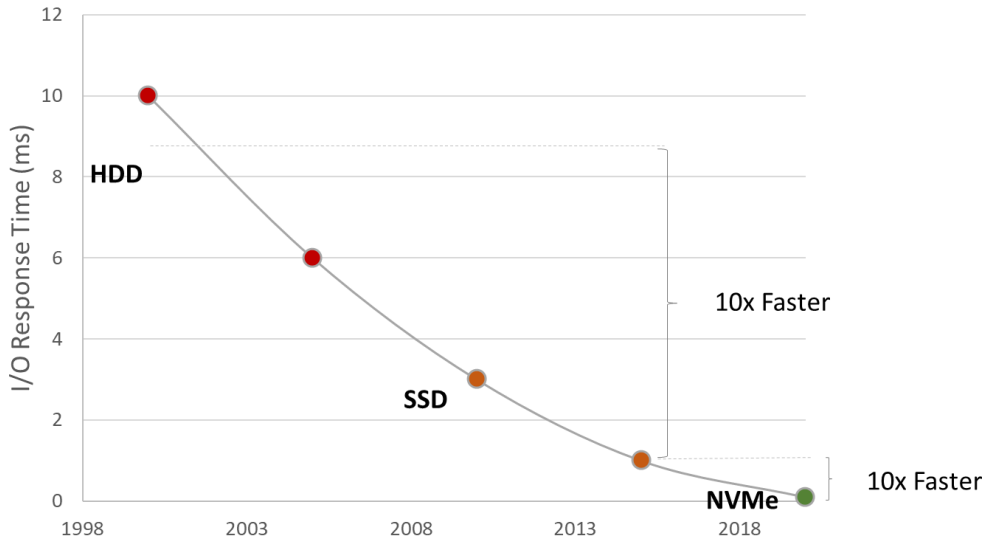

Source: Enterprise Strategy Group

While organizations see the benefits of NVMe, they also recognize that their path and timeline to adopting NVMe may vary, depending on current and planned workloads to support. 49% of organizations ESG surveyed have 30% or less of existing workloads running on flash storage, let alone NVMe. To that end, enterprises would benefit from a storage solution that can support multiple media types so that they can migrate from HDDs and SSDs to NVMe as business needs dictate, without sacrificing data access performance as required by the workloads to be supported.

Hitachi Vantara VSP 5000 Series with HAF

The Hitachi Vantara VSP 5000 series with HAF has been designed to help organizations support and consolidate a variety of workloads with different performance and latency requirements, while providing a path to migrate from spinning media to flash drives and, eventually, NMVe without the burden and expense of upgrades and migration. Leveraging the Hitachi Storage Virtualization Operating System (SVOS) RF, the VSP 5000 series with HAF has been designed to deliver a maximum of 21 million IOPS with as little as 70-microseconds latency.

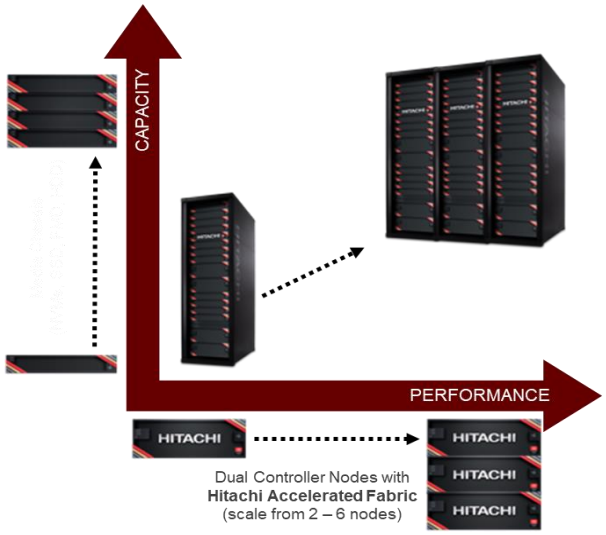

Why NVMe?

While the storage industry has touted the performance of flash drives, it is limited by existing serial attached SCSI (SAS) interfaces, as they have been designed for use with HDDs, not for high-speed flash media. The NVMe protocol was developed specifically for non-volatile, high-speed flash storage. Key features that enable maximum performance and lowest latency are NVMe's ability to run greater numbers of parallel I/O operations, its streamlined connection to the host CPU, and the simpler software stack that reduces I/O processing time and CPU overhead. These features accelerate existing applications and enable new applications that demand real-time processing. NVMe storage can also reduce overall TCO by handling hefty workloads in a smaller footprint.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

VSP 5000 Series Hardware Architecture

The Hitachi Vantara VSP 5000 series architecture is designed to meet varying performance, capacity, and storage technology requirements as dictated by current and future business needs. The key components of the architecture include the controller block and HAF.

To scale performance and capacity, the VSP leverages the controller block. The controller block acts as the base for the VSP 5000 series architecture. Each controller block contains two nodes, while each node contains dual controllers. As the number of controllers grow, the resiliency of the VSP platform increases. Failover of storage resources can occur across node controllers, nodes within the controller block, and controller blocks, minimizing the VSP 5000 series platform's downtime, thus increasing availability and resiliency. The VSP 5000 can scale to two, four, or six nodes (or one to three controller blocks, respectively). The media chassis, optionally available in any of the controller blocks, can enable organizations to further scale out capacity with any mix of storage media types simultaneously. Each media chassis is connected to the two nodes in the same controller block to achieve high system availability and growth of both performance and capacity. Because the VSP 5000 can support multiple storage media types, organizations can consolidate workloads and simplify storage management, thus decreasing both capital and operational expenses.

HAF enables the VSP 5000 to achieve high performance, regardless of the mix of storage protocols and media supported—HDD, SAS SSD, PCIe NVMe SSD, storage class memory (SCM) media ready, and Hitachi's flash modules (FMD). Considered the backbone of the node fabric, HAF consists of a pair of node interconnection switches (contained within the controller block) that act as the paths for data between all the controllers in a system, regardless of their location in a node or controller block. Not only can HAF enable performance and capacity to scale up and out, respectively, for efficient use and sharing of controller resources across the system, but also it enables organizations to tier data across controller blocks to maximize overall system performance.

Hitachi Vantara optimizes the interconnect path between the interconnect switches on the controllers via a PCI-Express Gen3 4Lane link (4GB/s). Each controller block has two fabric acceleration modules, and each module has two ports, for a total of four ports. Four interconnect paths link the controller to four separate infrastructure switch ports, providing quadruple redundancy for controller connections. The quadruple redundancy contributes further to the high performance, availability, and resiliency of the VSP 5000.

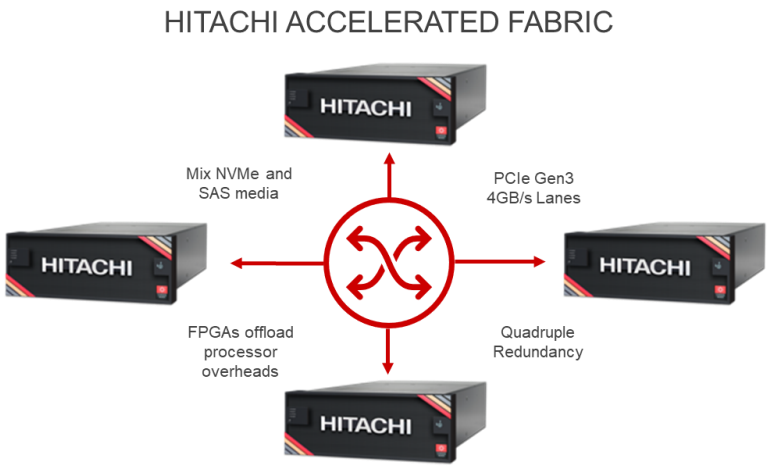

Hitachi Vantara bases the design of the fabric acceleration modules on field-programmable gate arrays (FPGAs) that are embedded in the interconnect on the controllers. The FPGAs offload processing from the controller to the interconnects, enabling SVOS RF 9 to use the flash-optimized code paths. Because fewer CPU cycles are consumed, the Hitachi Vantara VSP 5000 can achieve a maximum of 21 million IOPS. The design can also minimize latency to as low as 70 microseconds to reduce application response times.

The VSP 5000 series consists of four models, the 5100 and 5100H, and the 5500 and 5500H. The VSP 5100 and 5500 support all flash media, and the VSP 5100H and 5500H support both flash and spinning media for use cases such as data tiering—hot data stored within flash while reserving HDDs for cold data.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

ESG Technical Validation

ESG audited performance test results of the Hitachi Vantara VSP 5000 platform. Testing was designed to demonstrate the scalable performance of the VSP 5000 with HAF and validate the tangible business value offered by the platform.

NVMe Performance in the VSP 5000

ESG reviewed test results generated using industry-standard tools and methodologies. We first compared the performance of the VSP 5500 populated with SAS SSDs with another VSP 5500 with NVMe drives.

ESG Testing

The test bed consisted of two VSP 5500 systems, each with a single controller block (two nodes). One VSP 5500 was populated with 96 SAS SSDs totaling 3.8 TB, while the other was populated with 48 NVMe drives totaling 1.9 TB. Each VSP 5500 contained a 392GB cache. The test results recorded both IOPS and response time for an OLTP simulated reference workload generated with VDBench with a 70/30 read-write ratio, while accessing an 8TB data set. Storage media were arranged in a RAID 6 configuration (six primary drives with two protection drives). Figure 3 shows the results.

Figure 3. Comparing Performance of SAS SSD and NVMe-based VSP 5000

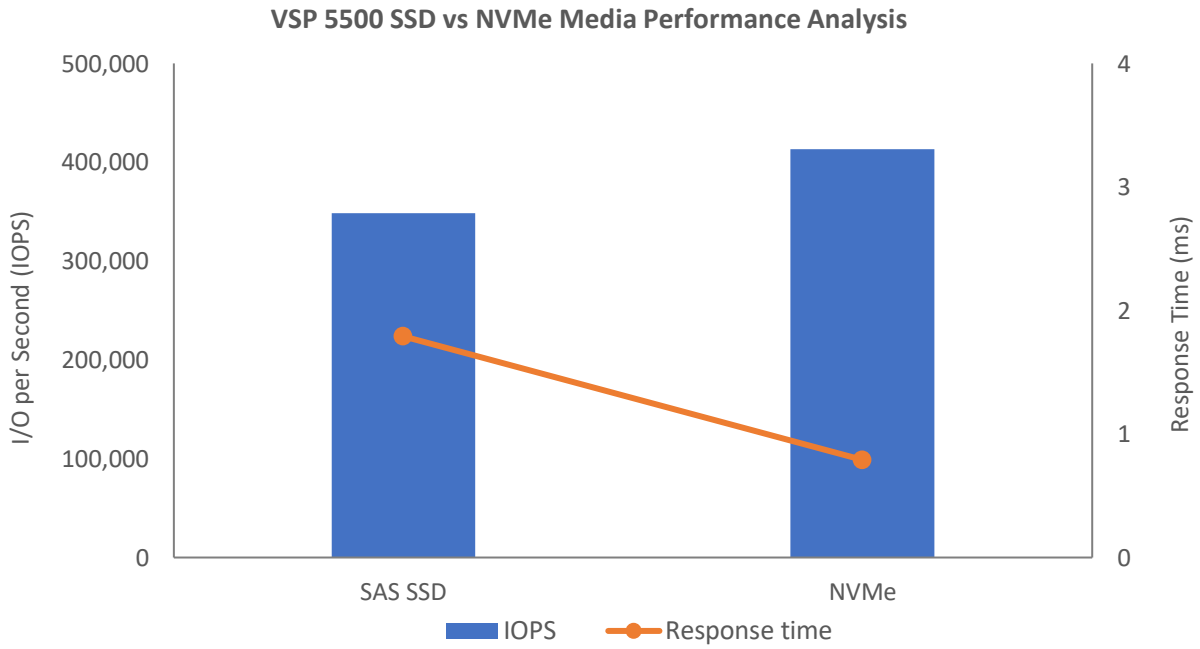

Source: Enterprise Strategy Group

Based on the results, ESG found that the performance of the VSP 5000 populated with NVMe drives exceeded that of the VSP 5000 populated with SAS SSDs, as it achieved 19% more IOPS with 56% faster response times. Based on these results, we see how the NVMe drives help with the overall performance of the VSP 5500 when compared with using all SAS SSDs. We can also attribute the increased IOPS to the dual controllers within each of the two nodes in the controller block. Since all four controllers process the OLTP workload, we observed the decreased response time.

VSP 5000 with NVMe exhibited **19% more IOPS and 56% faster response time** over VSP 5000 with SAS SSDs

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Organizations that are considering the use of data-intensive applications can benefit from the use of NVMe media. Given the large data sets that are typically accessed by big data analytics, artificial intelligence, and machine learning, the increased performance delivered using NVMe can ultimately reduce the time to process data and generate results. The less time spent on processing data translates into increased business agility, which is especially critical when responding to a constantly changing business environment.

VSP Performance with HAF

ESG then reviewed results of tests designed to illustrate the performance impact on the VSP 5000 when using HAF. We compared both the switch fabric I/O utilization and throughput of the VSP 5000 platform with those exhibited by a previous generation of the VSP with no HAF.

ESG Testing

The test bed consisted of two VSP platforms. We populated a two-module VSP G/F1500, a four-node VSP 5500, and a six-node VSP 5500 each with 96 3.8TB SAS SSDs.[2] Using VDBench to generate a random 60% read/40% write workload, the results showed the impact of HAF on VSP 5500 performance compared with the VSP G/F1500 performance without HAF, as illustrated in Figure 4.

Figure 4. HAF's Impact on VSP 5000 Performance

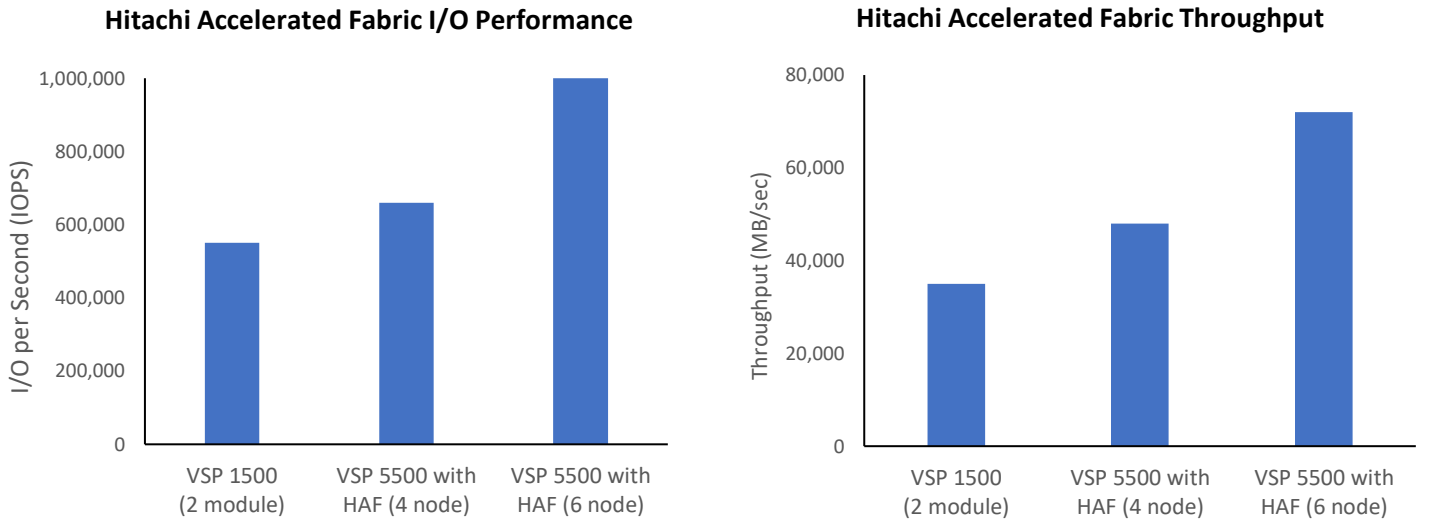

Source: Enterprise Strategy Group

Upon examining the results, ESG found that the VSP with HAF exhibited increased fabric I/O performance and throughout compared with the VSP G/F1500 without HAF. Specifically, we observed that the four-node VSP 5500 achieved 20% more IOPS and 37% more throughput than the previous generation VSP. Similarly, the six-node VSP 5500 achieved 82% more IOPS and 106% more throughput.

With the observed increases in both IOPS and throughput, ESG can infer that HAF improves overall VSP 5500 performance. We can see that the HAF architecture can offload processing from the controllers, enabling the VSP 5500 to increase IOPS

[2] A four-node VSP 5000 fully populated with SAS SSDs approximates a fully populated VSP G/F1500. A six-node VSP 5000 fully populated with SAS SSDs approximates 1.5x fully populated VSP G/F1500.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

as less CPU cycles are consumed. This is especially important to consider with data-intensive applications with high read/write ratios.

VSP 5000 Peak Performance

ESG proceeded to examine the results that illustrate the peak performance and minimal latency that can be achieved.

ESG Testing

To test peak performance, a three-node VSP 5500, populated with 96 3.8TB SAS SSD, was connected to eight servers via 64 32Gbps Fibre Channel (FC) ports. Using VDBench, an 8K Random Read Cache hit workload was generated. As the number of threads increased, response time was recorded. Results are graphed in Figure 5. The chart shows that the VSP 5000 achieved close to 21 million IOPS with a response time of 1.14 milliseconds.

Figure 5. Maximum IOPS Achieved by VSP 5000

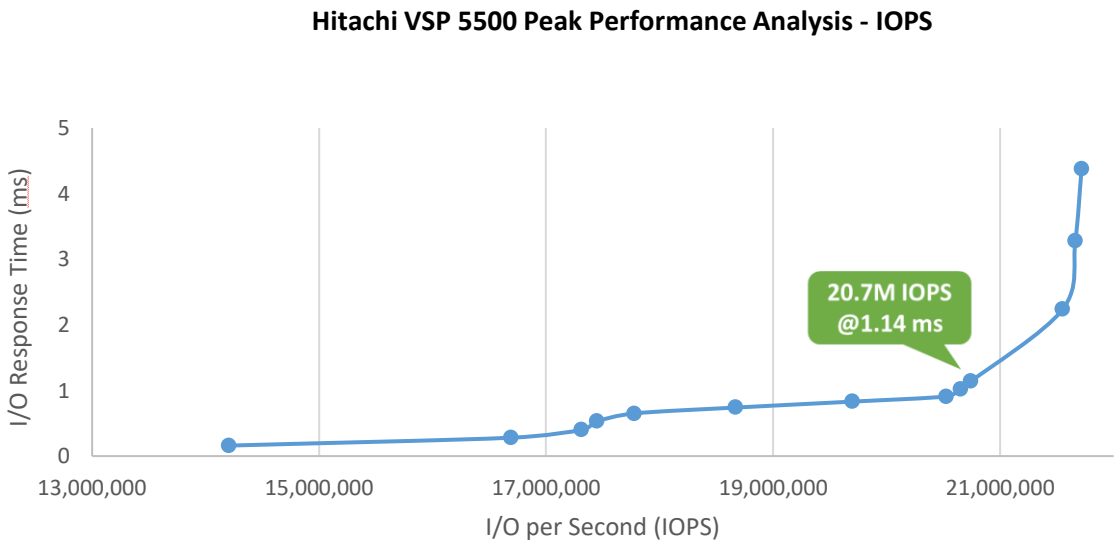

Source: Enterprise Strategy Group

We then looked at how well the VSP 5500 can reduce latency. The test leveraged a six-node VSP 5500 with 1702GB cache, populated with eight 1.9TB NVMe SSDs in a RAID6 configuration. Twenty-four host clients connected to the VSP 5500 via 192 32Gbps FC ports. Using VDBench, we generated an OLTP workload. The test measured response time as the number of threads increased. Results are graphed in Figure 6. The chart reveals that the VSP can achieve latency as low as 70 microseconds.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Figure 6. Lowest Latency Achieved by VSP 5000

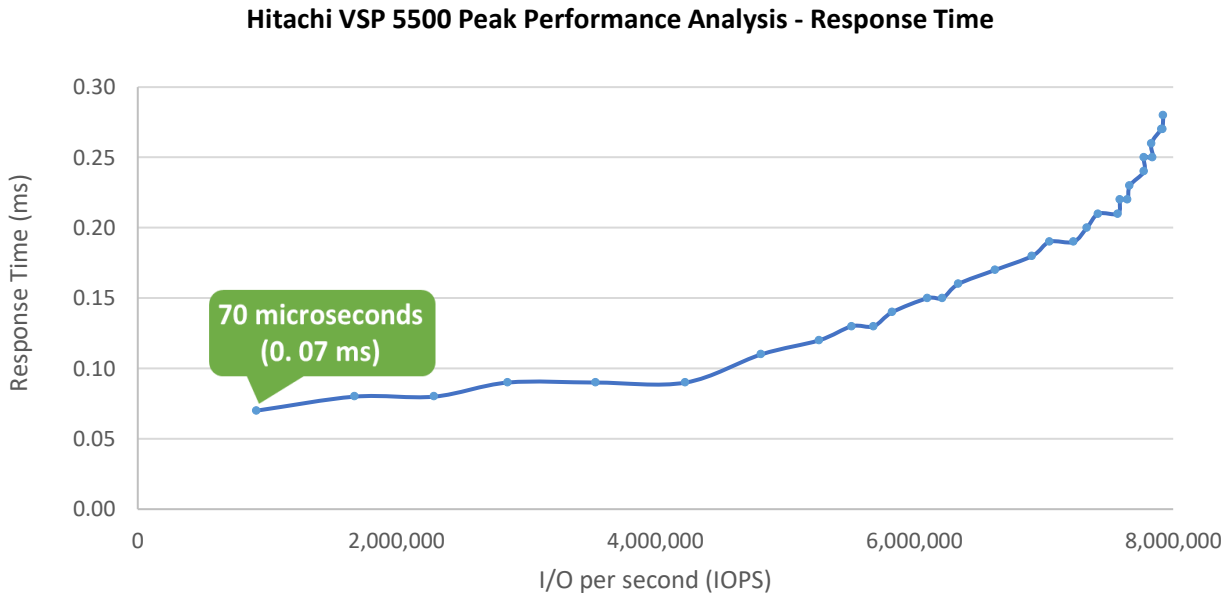

Source: Enterprise Strategy Group

Based on our audit of the test results, ESG validated that the VSP 5500 populated with NVMe drives can achieve close to 21 million IOPS. We also validated that a similarly configured VSP 5500 can achieve response times as low as 70-microseconds latency. The combination of NVMe SSD media, the controller block architecture, and HAF enables the VSP 5000 series to provide the high performance and low latency for applications requiring real-time processing of large data sets.

Why This Matters

While organizations are adopting applications such as big data analytics, machine learning, and artificial intelligence to increase their business advantage, they have been considering the use of NVMe-based platforms. The real-time processing and storage requirements for data-intensive applications must be addressed with a solution that delivers high performance at the lowest levels of latency.

ESG validated that the Hitachi VSP 5000 series with HAF has been architected to enhance the performance and reduced latency enabled by NVMe-based media. We saw that an NVMe-populated VSP 5500 outperforms the same system populated with SAS SSDs, showing a 19% improvement in IOPS and 56% improvement in response time. When considering the contribution that HAF makes to VSP performance, we found that storage fabric I/O and throughput increased as the number of VSP nodes increased when compared with a previous VSP generation. Finally, considering both the use of NVMe and HAF, ESG verified that the VSP 5500 can achieve close to 21 million IOPS. We also validated test results that revealed that the VSP 5000 series can exhibit response times as low as 70-microseconds latency.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

The Bigger Truth

As organizations continue to adopt and use data-intensive applications—big data analytics, the internet of things (IoT), artificial intelligence, and machine learning—they require a solution that can efficiently store and process data at high performance levels while reducing response times. While NVMe is being seriously considered to support such applications, organizations must also consider how to integrate it into existing SAS SSD and HDD-based storage environments supporting their traditional enterprise applications. How can these organizations future-proof their storage environments to allow a graceful migration to NVMe, while maintaining support for existing applications without incurring architectural complexity and, subsequently, additional capital and operational expenses?

The Hitachi Vantara VSP 5000 series with HAF has been designed to meet varying performance, capacity, and storage technology requirements as dictated by current and future business needs. The storage platform can support multiple storage media and protocols simultaneously while maximizing performance and minimizing latency for traditional enterprise and big data-intensive applications. Scaling performance and capacity is accomplished via the controller block, which helps to distribute the processing workload across multiple controllers. HAF is also key for the VSP 5000 series to achieve high performance levels with low latency as it offloads processing to the links connecting fabric acceleration modules between controller blocks. The VSP 5000 series can scale up to 6 nodes (12 controllers) with a maximum raw capacity of 69 PB.

ESG's examination of performance test results showed that the VSP 5000 series achieves high performance and low response times. By observing how the use of NVMe and HAF affect VSP 5000 series performance and latency, ESG saw how the combination of NVMe drives, the controller block, and HAF enable the VSP 5000 series to achieve a maximum of 21 million IOPS and exhibit response time as low as 70 microseconds of latency.

If your organization is looking to gain competitive edge with data-intensive applications, ESG recommends that you take a closer look at the performance advantages that the Hitachi Vantara VSP 5000 series with HAF can offer your business to take better advantage of your data.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

 www.esg-global.com

 contact@esg-global.com

 P. 508.482.0188