



CONTENTS

- 2 SUPERMICRO HYPER-CONVERGED INFRASTRUCTURE FOR VMWARE VSAN, ENABLING EFFICIENT AND AGILE DEPLOYMENTS
- 2 SUPERMICRO OFFERS THE BROADEST AND MOST OPTIMIZED PORTFOLIO OF NVME SERVERS
- 3 THE ALL-FLASH NVME TESTING CLUSTER WITH EIGHT SUPERMICRO 2U ULTRA SUPERSERVERS
- 6 SUPERMICRO ALL NVME HCI SOLUTION ACHIEVES 50% GREATER PERFORMANCE AND 7X LOWER CAPEX IN A VIRTUAL SAN ENVIRONMENT
- 7 CONCLUSIONS
- 7 ACKNOWLEDGMENT

WHITE PAPER

OPTIMIZING VMWARE VSAN WITH SUPERMICRO ALL-FLASH NVME SYSTEMS

Supermicro® Hyper-converged Solution with Intel® NVMe SSDs Achieves 50% Greater Performance and 7x CAPEX Improvement in VMware® vSAN™

June 2017

Super Micro Computer, Inc.
980 Rock Avenue
San Jose, CA 95131 USA
www.supermicro.com

Intel Inside®. Powerful Data Center Outside.



SUPERMICRO HYPER-CONVERGED INFRASTRUCTURE FOR VMWARE VSAN, ENABLING EFFICIENT AND AGILE DEPLOYMENTS

Hyper-Converged Infrastructure (HCI), based on VMware® vSAN, with All-Flash NVMe based servers and storage from Supermicro transforms industry-standard x86 servers and direct attached storage into a radically simplified high performance infrastructure that reduces IT costs, management complexity, and performance limitations. vSAN 6.2 and later releases add support for major data efficiency software features such as deduplication, compression, and erasure coding which require NVMe all-flash optimized hardware to deliver on the promise of dramatic improvements in the affordability and availability of VMware’s vSAN. Testing conducted on Supermicro NVMe all-flash SuperServer demonstrated a 50% improvement in performance and an up to 7x reduction in effective cost (\$/GB) compared to previous hybrid SAS and NVMe SSD configuration.




Figure 1. Traditional Infrastructure and Hyper-Converged Infrastructure of IT

SUPERMICRO OFFERS THE BROADEST AND MOST OPTIMIZED PORTFOLIO OF NVMe SERVERS

As the undisputed leader in NVMe server technology, Supermicro offers over 120 server models that support NVMe SSDs and the latest Intel® Xeon® processor E5-2600 v4 product family. Supermicro is the first to market with hot-plug capability for NVMe drives and first to market with dual-port NVMe drives for improved serviceability and availability.

Thanks to more streamlined I/O stack, the All-Flash NVMe servers from Supermicro provide order of magnitude performance improvements (up to 10x) over system based on traditional hard disk drives (HDD) and SAS3 solid state drives (SSD). These servers are also more power efficient than the traditional systems due to the fact



that less hardware is required to achieve the same performance goals resulting in less power consumption and lower TCO.

The all-flash NVMe products in the Supermicro portfolio provide three fundamental design alternatives for customers:

- Throughput/IOPS optimized (e.g., 1U 10 NVMe Drives)
- Balanced Throughput and Capacity (e.g., 2U 24 NVMe Drives)
- Capacity Optimized (e.g., 2U 48 NVMe Drives)

Depending on the requirements for IOPS, latency, capacity, endurance, and workloads (read/write), Supermicro provides a variety of choices that customers can choose from to build NVMe all flash vSAN clusters. All these NVMe server products from Supermicro can be found at www.supermicro.com/NVMe. As an illustration, two of the common NVMe products are shown in Figure 2 and Figure 3.



Figure 2. Supermicro 1U Ultra SuperServer 1028U-TN10RT+: 1U 10 NVMe Drives; AF-8 vSAN Ready Nodes™

Next, we are going to put one of the many Supermicro NVMe servers to test and see how it performs against previous hybrid SAS and NVMe flash configuration in a vSAN environment.



THE ALL-FLASH NVMe TESTING CLUSTER WITH EIGHT SUPERMICRO 2U ULTRA SUPERSERVERS

The Supermicro 2U Ultra SuperServer as shown in Figure 3 supporting 24 NVMe SSDs is used for this test, and runs vSAN 6.2. Each Supermicro SuperServer 2028U-TN24R4T server is configured with the following components

- Dual Intel® Xeon® E5-2687W v4 (12 Core @ 3.0 Ghz)
- 256 GB DDR4 RAM
- Boot Drive - 1x Intel® SSD DC S3710 Series (200 GB, 2.5")



- vSAN SSDs - 4 Disk Groups comprised of
 - 4x Intel® SSD DC P3700 Series (800 GB, 2.5" SFF) for caching
 - 20x Intel® SSD DC P3520 Series (2 TB, 2.5" SFF) for capacity
- Intel® Ethernet Server Adapter X540-DA2

The cluster also contains a 48 port 10Gbase-CX4 switch with 6x 40 Gigabit QSFP uplinks.

As a total, the eight server cluster provides

- 2,000 Virtual Machines With Windows* Server
- 8 Hyper-Converged VMware* ESXi 6.2 Hosts
- 192 Intel® Xeon® Cores
- 2 TB DDR4 Memory
- ~345 TB Raw Flash Memory
 - 25 TB of vSAN Cache
 - 320 TB of Raw Datastore Storage

Each VM profile is comprised of 1 vCPU, 4GB DRAM, and 2 virtual drives spread across 2 virtual controllers. One of the virtual drives is LSI logic SAS with a 40GB thin provisioned OS disk and another is a Paravirtualized 10GB thick provisioned data disk.



Figure 3. Supermicro 2U Ultra SuperServer 2028U-TN24R4T+: 24 NVMe drives is used to set a new performance record.



The vSAN storage policy is configured as follows

- Number of Disk Stripes is set to 1
- Number of Failures to Tolerance is set to 1
- Checksum Disabled is set to True

The vSAN host network is configured as follows

- 1 Virtual switch with 2x 10G uplinks
- Each vmnic is configured as a trunk, 4 port groups with 3 reserved for VMKs containing load balancing across both 10G ports in the VM network where vmk0 is for management and vMotion
- Each vSAN VMK is dedicated to different 10G ports where vmk1 is vSAN "A" network, denoted "vSAN.A" in Figure 4 and vmk2 is vSAN "B" network, denoted "vSAN.B" in Figure 4 respectively.

In addition, each single node's profile under testing is set as follows

- CPU power and performance policy is set to Performance
- Workload configuration is set to Balanced
- Memory RAS configuration is set to Maximum Performance
- Fan profile is set to Performance

The following features on the CPUs are enabled in the test: hyper-threading, NUMA optimized, enhanced Intel SpeedStep® Technology, Intel® Turbo Boost Technology, Uncore Frequency Scaling, and Performance P-Limit, while the other features listed below are disabled: Cluster on Die, Early Snoop, and Energy Efficient Turbo.




Figure 4. The 8 Node Cluster and the Networking Configuration of VM in the Test.

BUILD NAME	REPLICA PREFERENCE	DEDUPLICATION & COMPRESSION	FAILURES TO TOLERATE
Performance	Mirror	Disabled	1
Balance	Mirror	Enabled	1
Capacity	Erasure Coding	Enabled	1
Availability	Erasure Coding	Enabled	2

Table 1. The Configuration of the 4 Build Groups in The All NVMe vSAN Test

In the test, a total of 4 different build groups are formed to showcase the trade-offs between performance, capacity, and availability. Their configuration details can be found in Table 1.

SUPERMICRO ALL NVME HCI SOLUTION ACHIEVES 50% GREATER PERFORMANCE AND 7X LOWER CAPEX IN A VIRTUAL SAN ENVIRONMENT

The new Supermicro all NVMe HCI cluster achieves record breaking performance

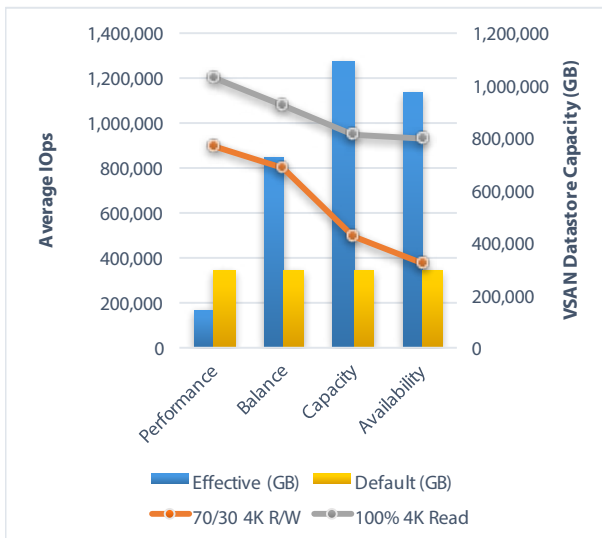


Figure 5. The Empirical Test Result of IOPS among 4 Build Groups

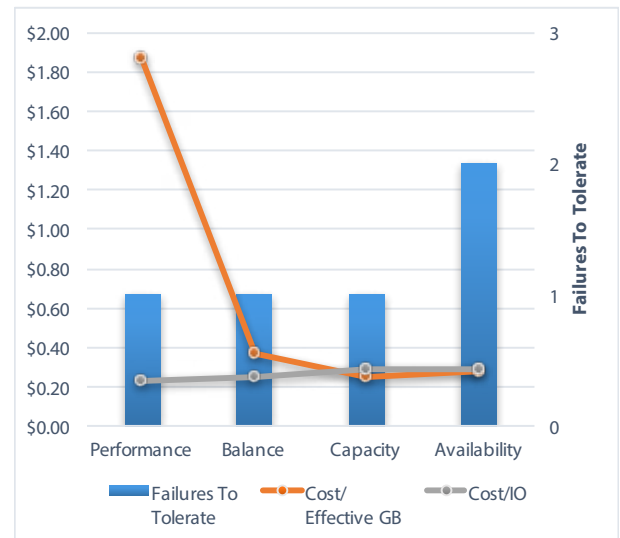


Figure 6. The Empirical Test Result of TCO among 4 Build Groups





FOR MORE INFORMATION

Supermicro® SuperServer®
SYS-1028U-TN10RT+ Datasheet

www.supermicro.com/products/system/1u/1028/sys-1028u-tn10rt.cfm

Supermicro® SuperServer®
SYS-2028U-TN24R4T+ Datasheet

www.supermicro.com/products/system/2u/2028/SYS-2028U-TN24R4T.cfm

VMware vSAN

www.vmware.com/products/virtual-san.html

Intel® NVMe SSDs

www.intel.com/content/www/us/en/solid-state-drives/solid-state-drives-ssd.html

and CAPEX reduction. As shown in Figure 5, the Performance Build yields 1.2M 4K read IOPS from the eight nodes, or 150K IOPS per node, which is 50% higher than the previous all-flash configuration we have tested.

The IOPS yielded by the Balanced Build is 89% of Performance Build but the Balanced Build offered 250% additional effective capacity. The Capacity Build increased effective capacity by 376%, the largest effective capacity among all Builds. However its 70/30 (70% read and 30% write ratio workload) and read IOPS dropped to 55% and 78% of Performance Build, respectively. The Availability Build had the second largest effective capacity with an increase of 333% but lower 70/30 and read only performance at 41% and 76% of

Performance Build, respectively. The Performance Build achieved the highest performance while the Capacity Build gained the largest effective capacity with the all-NVMe vSAN cluster.

As shown in Figure 6, the system acquisition cost consisting of server hardware, software and the network switches is compared across the four builds. The Supermicro all-flash NVMe cluster achieved build yields ranging from \$1.86/GB to an impressive \$0.25/GB. Deduplication and Compression alone reduce cost per Effective GB by over 7x.

CONCLUSIONS

In this paper, we demonstrated a 50% improvement in performance and an up to 7x reduction in effective cost (\$/GB) in a vSAN environment using Supermicro's NVMe all-flash SuperServer 2028U-TN24R4T compared to previous hybrid SAS and NVMe SSD configuration. Although the actual performance number varies, each Supermicro All Flash NVMe SuperServer™ will greatly outperform its non-NVMe counterparts currently available in the marketplace. With the broadest portfolio of NVMe optimized servers in the industry, Supermicro customers have the flexibility to select the right NVMe product that meets their performance as well as budgetary needs. The Supermicro all-flash NVMe product portfolio makes the deployments of the high performance and cost-effective HCI solutions feasible today.

ACKNOWLEDGMENT

This white paper is based on the previous works of **C. Brown, J. Hubbard, B. Billo**, "Re-architect storage: Intel® SSDs and 3D XPoint™ technology for Hyper-Converged Real World Guidance for Implementing VMware® vSAN Ready Nodes™ to Build Your Own," STO9396-SPO, VMworld Conference, Las Vegas, NV. August 29-September 1, 2016.



About Super Micro Computer, Inc.

Supermicro® (NASDAQ: SMCI), the leading innovator in high-performance, high-efficiency server technology is a premier provider of advanced server Building Block Solutions® for Data Center, Cloud Computing, Enterprise IT, Hadoop/Big Data, HPC and Embedded Systems worldwide. Supermicro is committed to protecting the environment through its “We Keep IT Green®” initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.

www.supermicro.com

The information contained in this document is subject to change without notice.

Results are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. Performance tests are measured using specific computer systems, components, software, operations, functions, and workloads. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

No part of this document covered by copyright may be reproduced in any form or by any means — graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system — without prior written permission of the copyright owner.

Supermicro, the Supermicro logo, Building Block Solutions, We Keep IT Green, SuperServer, TwinPro™, TwinPro²™, SuperDoctor are trademarks and/or registered trademarks of Super Micro Computer, Inc.

Ultrabook, Celeron, Celeron Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, vPro Inside, Xeon, Xeon Phi, and Xeon Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

© Copyright 2016 Super Micro Computer, Inc. All rights reserved.

Printed in USA

MKT-0006-09/2016-10-R13



Please Recycle

