

Broadcom NetXtreme® 25GbE Adapter

Ethernet Adapter Impact on vSAN Performance: Competitive Comparison

EXECUTIVE SUMMARY

Hyperconverged infrastructure (HCI) combines virtualization, computer servers, the storage network and storage to provide a simplified architecture that offers both agility and lower total cost of ownership (TCO). HCI implementations such as VMware vSphere/vSAN can be optimized and deliver higher performance using 25GbE networking.

Broadcom commissioned Tolly to benchmark the vSAN storage networking performance of the Broadcom NetXtreme 25Gb Ethernet Adapter in a VMware hyperconverged environment and compare those results to similar Ethernet adapters from Intel, Marvell and Mellanox.

The Broadcom NetXtreme adapter outperformed the competing solutions in both IOPS and data throughput in all scenarios of 16KB blocks and higher.


THE BOTTOM LINE

The Broadcom NetXtreme 25GbE Adapter delivered vSAN HCIBench throughput that was:

- 1 Up to 96% greater than when using a 10GbE Ethernet adapter
- 2 Up to 25% greater than a Mellanox 25GbE Ethernet adapter
- 3 Up to 29% greater than an Intel 25GbE Ethernet adapter
- 4 Up to 26% greater than a Marvell 25GbE Ethernet adapter

25GbE vs 10GbE vSAN Performance: HCIBench 70% Read Random

Broadcom NetXtreme P225P



Note: Oracle Vdbench 5.04.07. Four-node VMware vSAN.

Source: Tolly, April 2019

Figure 1



Overview

Hyperconverged infrastructure simplifies architecture and deployment by concentrating and combining elements of compute and storage into a single physical infrastructure using a software layer on x86 servers to replace purpose-built storage hardware.

Given the concentration of essential elements, the HCI virtual storage area networks (vSAN) environment needs to be high performance to support mission-critical applications.

VMware developed HCIbench to help users evaluate VMware vSAN performance by automating the low-level Oracle Vdbench storage benchmark. HCIbench/Vdbench provided the basis for the current test.

Test Results

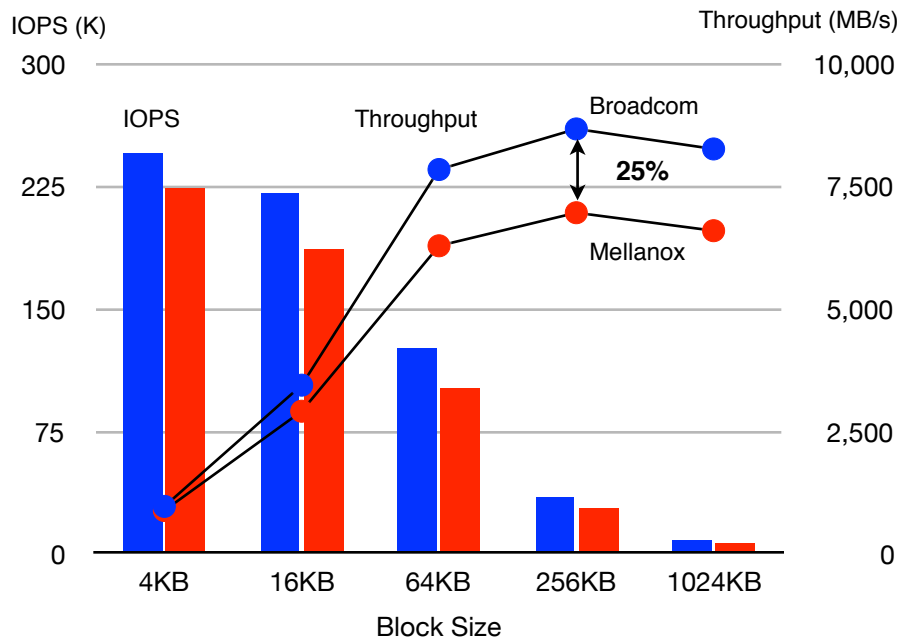
Benefits of 25GbE

For years, 10GbE was the fastest server connection available and became the standard for data center networking fabrics. Today, 25GbE network interface adapters and switches are readily available and the additional bandwidth can provide dramatic improvements to storage networking applications.

To illustrate these benefits, Tolly first ran the HCIbench suite with the vSAN environment connected using the Broadcom NetXtreme adapter configured for 10GbE and then again configured for 25GbE.

As the storage block size increased, the benefits of 25GbE become increasingly apparent. In fact, with 256KB blocks, 10GbE delivered 4,440 MB/s throughput where 25GbE was 96% higher at 8,689 MB/s. See Figure 1 and Table 1.

25Gb Ethernet vSAN Performance: HCIbench 70% Read Random Broadcom NetXtreme P225P vs. Mellanox CX512A

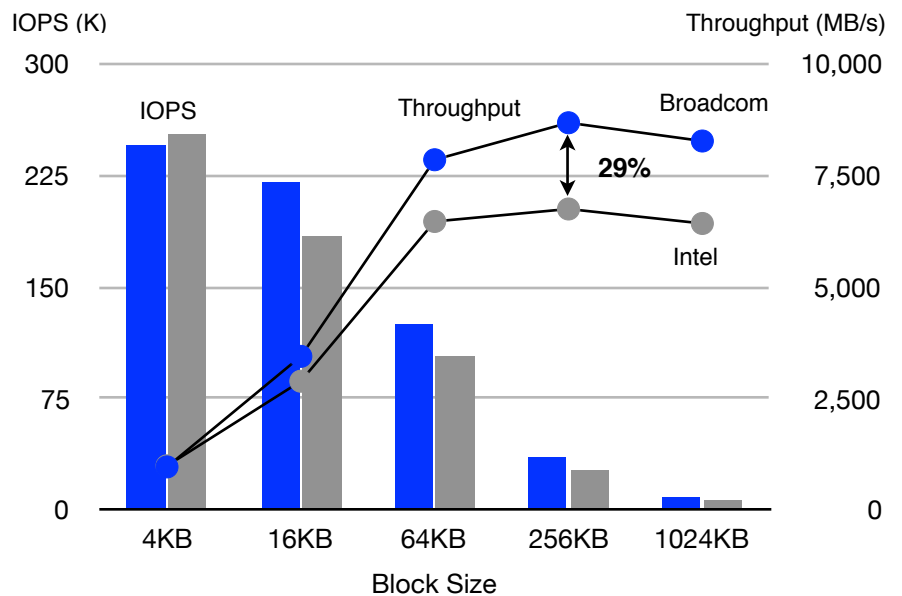


Note: Oracle Vdbench 5.04.07. Four-node VMware vSAN.

Source: Tolly, April 2019

Figure 2

25Gb Ethernet vSAN Performance: HCIbench 70% Read Random Broadcom NetXtreme P225P vs. Intel XXV710



Note: Oracle Vdbench 5.04.07. Four-node VMware vSAN.

Source: Tolly, April 2019

Figure 3



Broadcom vs. Competing Ethernet Adapters

The tests were run using four different vendors' 25GbE Ethernet adapters with all other elements of the test environment remaining the same.

In every test case, save one, the Broadcom Ethernet adapter outperformed the competing solutions.

Mellanox

In this test, Broadcom NetXtreme delivered more IOPS and greater data throughput at every block size. The performance delta was the greatest with 256KB blocks.

Broadcom's data throughput of 8,689 MB/s was 25% greater than the 6,973 MB/s of throughput measured with the Mellanox Ethernet adapter. See Figure 2 and Table 1.

Intel

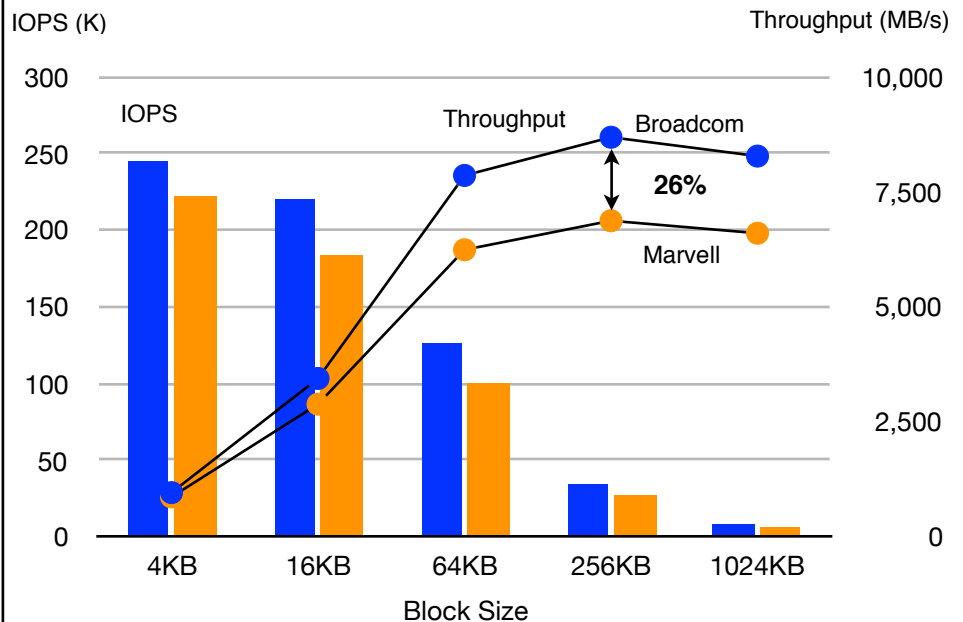
In this test, Broadcom NetXtreme delivered more IOPS and greater data throughput at every block size with the exception of 4KB blocks. At this smallest block size tested, Intel's throughput of 984 MB/s was 2.5% higher than Broadcom's 959 MB/s. The performance delta, again, was the greatest with 256KB blocks.

Broadcom's data throughput of 8,689 MB/s was 29% greater than the 6,752 MB/s of throughput measured with the Intel Ethernet adapter. See Figure 3 and Table 1.

Marvell

In this test, Broadcom NetXtreme delivered more IOPS and greater data throughput at every block size. The performance delta was the greatest with 256KB blocks.

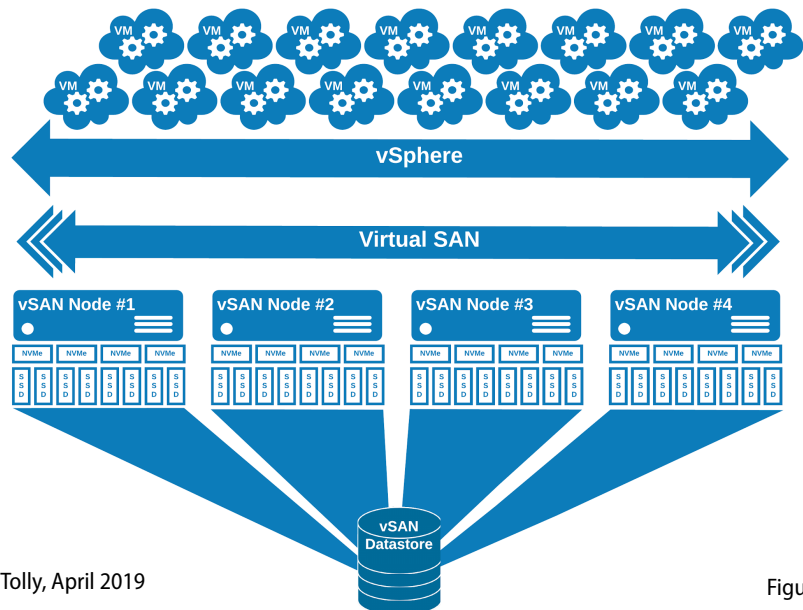
25Gb Ethernet vSAN Performance: HCI Bench 70% Read Random Broadcom NetXtreme P225P vs. Marvell FastLinQ 41000



Note: Oracle Vdbench 5.04.07. Four-node VMware vSAN.
Source: Tolly, April 2019

Figure 4

Logical Test Bed Topology Four-Node VMware vSAN



Source: Tolly, April 2019

Figure 5

Broadcom's data throughput of 8,689 MB/s was 25% greater than the 6,873 MB/s of

throughput measured with the Marvell Ethernet adapter. See Figure 4 and Table 1.



Test Setup & Methodology

All Ethernet adapters under test used the latest and most current firmware and driver sets available on the respective vendor's website. The Ethernet adapters were two-port adapters but testing involved one port only. For model numbers, driver and firmware information, see Table 2.

The test environment was a four-node VMware vSAN environment. Each server was a Dell R740XD outfitted with 2x Intel Xeon Gold 3.3GHz processors, 12x 32GB RAM and high performance SSD SAS drives along with NVMe cache. 25GbE network connectivity was provided by an Edgecore Networks AS7712-32X switch. Additional details can be found in Table 3.

Solution-Level Testing

Tolly/Broadcom measured VMware vSAN in an industry standard way using HClBench/Vdbench as a benchmark of total solution performance. In fact, VMware developed HClBench for this very purpose.

Both Marvell and Mellanox were given their results. Mellanox did not comment.

Marvell noted to Tolly, in part: *"The current vSAN perf data does not indicate how much of the vSAN traffic is local to the server and how much is going over the network/NIC. [VMware] vSAN has complex algorithms that decide what data blocks are served and where and iterations of such tests without accurately characterizing network I/O performance are not comparable."*

The intention of this method and white paper, again, is to represent solutions-level performance. It is true that not all storage is "going over the network/NIC", however, this does not change the fact that the performance of the overall solution (including local and remote traffic) is improved with certain Ethernet adapters.

Broadcom, Inc.

NetXtreme
25GbE Adapter

Hyperconverged
Infrastructure
vSAN
Performance



Tested
April
2019

The vSAN cluster generally used VMware defaults. Relevant configuration parameters can be found in Table 5.

VMware HClBench and Oracle Vdbench were used for the benchmarking. Vdbench is the low-level test tool that drives the storage access. HClBench is an automation-wrapper that provides an easy way to run the benchmark in a clustered environment. Relevant test information and parameters for these tools can be found in Table 6.

HClBench Detailed Storage Test Results: Four-Node vSAN

Vendor	Solution Under Test	Port Speed	Metric	vSAN Performance - Vdbench 70% Read/100% Random				
				Block Size				
				4KB	16KB	64KB	256KB	1024KB
Broadcom	NetXtreme P225P	10GbE	IOPS (K)	239	229	72	18	4
			Throughput (MB/s)	933	3,439	4,492	4,440	4,388
		25GbE	IOPS	245	220	126	35	8
			Throughput (MB/s)	959	3,442	7,858	8,689	8,279
Intel	XXV710DA2	25GbE	IOPS	252	185	104	27	6
			Throughput (MB/s)	984	2,884	6,472	6,752	6,427
Marvell	FastLinQ QL41262HLCU	25GbE	IOPS	221	184	100	27	7
			Throughput (MB/s)	865	2,883	6,241	6,873	6,601
Mellanox	ConnectX-5 CX512A	25GbE	IOPS	224	186	101	28	7
			Throughput (MB/s)	875	2,908	6,296	6,973	6,605

Note: Dual-port adapters. Single port tested.
Source: Tolly, April 2019

Table 1



Ethernet Adapters: Solutions Under Test

Vendor	Product Name	Part No.	Driver	Driver Version	Firmware Version
Broadcom	NetXtreme P225P	BCM957414A4142CC	bnxtnet	214.0.195.0	214.0.191.0
Intel	Ethernet Network Adapter XXV710-DA2	XXV710DA2	i40en	1.7.17	6.80.0x80003d05 1.2007.0
Marvell	FastLinQ QL41262HLCU	QL41262HLCU	qedentv	3.10.23.1	mfw: 8.40.30.0, storm: 8.37.9.0
Mellanox	ConnectX-5 CX512A	MCX512A-ACAT	nmlx5_core	4.17.15.16	16.24.1000

Table 2

Test Infrastructure

Server Hardware Configuration

Number of Nodes	4
Make & Model	Dell R740XD
CPU	2 x Intel Xeon Gold 6134 8C/16T 3.2GHz
Memory (RAM)	12 x 32GB RDIMM 2666MT/s Dual Rank
Disk Controller	HBA330 12Gbps
OS Drive	BOSS controller card + with 2 M.2 Sticks 240G + RAID 1
Cache Drives	4 x Dell 800GB NVMe PM1725a
Capacity Drives	8 x 1.92TB SSD SAS Read Intensive 12Gb/s
Switches	Edgecore Networks AS7712-32X-O-AC-F (25GbE), and Dell PowerConnect 8132F (10GbE)

Table 3

Software/Firmware Configuration

VMware OS	6.7.0-u1-10302 608
VMware vCenter	6.7.0-102 44745
HCIBench	1.6.8.7
Vdbench	5.04.07
Disk Controller	16.17.00.03
Edgecore Networks Switch	5.2.19.47
Dell Switch	5.1.0.1

Table 4

vSAN Cluster Configuration

Default Storage Policy	Failures to tolerate: 1 failure - RAID-1 (Mirroring)
vSphere DRS	Off
vSphere HA	Off
VMware EVC	Disabled
Dedupe/Compression	Disabled
Encryption	Disabled
iSCSI Target Service	Disabled
Fault Domain	1 host failure maximum
DSwitch	VLAN ID: 100, V: 6.6.0

Table 5

HCIBench/Vdbench Configuration

# VMs	16 total
# Data Disks	8 per VM
Disk Size	10GB per disk
Working Set %	100%
# Thread per Disk	4
Block Sizes (KB)	4/16/64/256/1024
Read Pct.	70%
Random pct.	100%
Test Time (s)	3600
Warmup Time (s)	1800
Clear Cache Before Testing	Yes
Prep. VDisk Before Testing	Zero
I/O Rate	Default (unlimited)
Reporting Interval (s)	60
Worker Threads Total	512 (16 VMs x 8 Disks x 4 threads)

Table 6

Source: Tolly, April 2019



About Tolly

The Tolly Group companies have been delivering world-class IT services for 30 years. Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

You can reach the company by E-mail at sales@tolly.com, or by telephone at +1 561.391.5610.

Visit Tolly on the Internet at:
<http://www.tolly.com>

Interaction with Competitors

In accordance with Tolly's Fair Testing Charter, Tolly personnel invited representatives from the competing vendors to view the test details and their results. Intel did not respond. Marvell and Mellanox responded and were provided with their results prior to publication. See report for relevant comments.

For more information on the
Tolly Fair Testing Charter, visit:
<http://reports.tolly.com/FTC.aspx>



Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is," and Tolly Enterprises, LLC (Tolly) gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com. No part of any document may be reproduced, in whole or in part, without the specific written permission of Tolly. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.