



Do budget network switches negatively impact VSAN performance and should this be a key consideration in VSAN design?

www.synchro.net

1-877-952-0104

Technical Whitepaper

AUTHORS	Brandon Wardlaw Ritesh Kapadia John Nicholson
TARGET	Technical Managers, Systems Architects, and Engineers who are evaluating VMware VSAN.
HARDWARE	Various Manufacturers
SOFTWARE	VMware VSAN

SynchroNet specializes in VMware virtualization solutions for a wide variety of national and international clients.

Introduction

As a VMware Elite VSAN partner, SynchroNet is committed to fully understanding VSAN best practices, from a practical, hands-on perspective.

A major component of any VMware VSAN deployment is the network switching, but it is curiously absent from the VMware Hardware Compatibility List (HCL). In fact, several VMware VSAN technical sessions have glossed over the topic stating that any switch that is gigabit and supports multicast is adequate. 10GbE is recommended, especially for maintenance and data migration. Our customers have many different switching vendors within their environments and we felt it important to discover what level of impact switch selection has on VMware VSAN performance, especially since other protocols (e.g., iSCSI) are severely impacted by small-buffer, low-PPS network switches.

We obtained a wide gamut of switches to test for VSAN performance to validate whether or not VMware is justified in not including network switches on their VMware VSAN HCL. We ran performance tests on a variety of switches from leading vendors to identify equipment that supports reliable and efficient VSAN performance, while keeping equipment costs within reasonable boundaries for small to medium deployments.

About VSAN

Virtual SAN (VSAN) is VMware's hyper-converged and software-defined storage solution. This product is integrated into vSphere 5.5 and later versions, and works "out-of-the-box" with any hardware on VMware's hardware compatibility list.

VSAN Capabilities

- **Enhanced performance:** server-side caching on flash storage media backed by spindle disks provides for performance gains associated with converging storage and compute resources within a host
- **Improved manageability:** VSAN is a policy-driven storage solution that simplifies provisioning and management of storage resources
- **Tiered redundancy:** VM Storage Profiles provide flexibility when defining failure tolerance for virtual machine resources
- **Flexible scalability:** Changing and adding storage to a VSAN cluster requires no special hardware or reconfiguration

Testing Approach and Methodology

The test was conducted in SynchroNet's lab, using the following hardware. A series of tests were run using 4 instances of the VMware Flings I/O Analyzer (3 workers, 1 manager/data collector).

Server Hardware (x3)

- Server: ASUS RS720-X7/RS8
- Processor: Intel Xeon E5-2620 @ 2.0GHz
- Memory: 48GB
- SSD: 240 GB Intel DC S3500 Series
- HDD: 6x WD 2 TB SATA
- Controller: ASUS Pike 2008
- Intel 82574L Gigabit Ethernet (management)
- Intel X540-AT2 10 Gigabit Ethernet (VSAN)

Networking Hardware

- Brocade: ICX6610-48-E
- Cisco: Catalyst 3500 XL; Catalyst 3750X; Meraki MS220-8P
- DLink: DGS-2208
- Extreme: X450A-24T
- HP: V1910-24G
- Juniper: EX2200-48T-4G
- Netgear: XS712T

The switches selected were all entry-level or mid-range business-class switches, except for the D-Link DGS-2208. The D-Link was selected to help provide a valuable control – to determine whether a home-office switch could keep up with the more costly options. Every switch tested was reset to its factory default configuration. No performance tuning was performed on the switches, and each switch was dedicated entirely to VSAN traffic for the duration of the test.

With each switch, we ran a battery of test workloads with I/O Analyzer. Each worker was configured to perform operations on a 100GB virtual disk located on the VSAN datastore. Each test workload ran for 120 seconds, and was repeated 3 times. The following test workloads were selected from the I/O Analyzer toolkit:

- 4k_100read_0rand.icf (Workload specs: 4k_100%Read_0%Random)
- Max_Write_IOPS.icf (Workload specs: 0.5k_0%Read_0%Random)

Results

On the whole, our results indicated fairly minor variation in performance between network switches. In fact, variance between workload runs on the same switch were often times greater than between completely different switch vendors. The charts below (A,B) represent an average of three (3) result sets for each switch. Please note that the charts use non-zero vectors to amplify the separation between data points. *Note: Raw test results are available by sending e-mail to vsan@synchronet.com.*

Chart A.

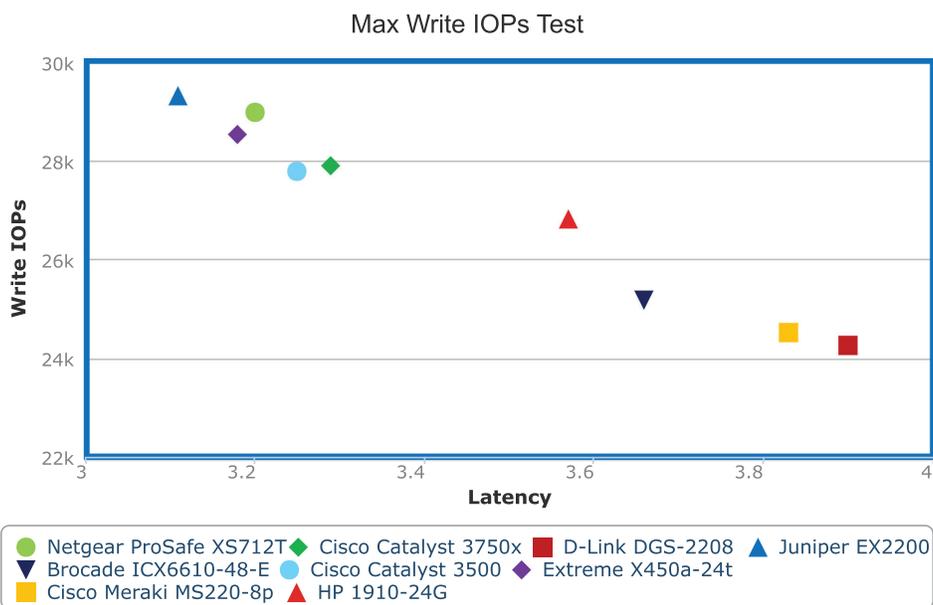
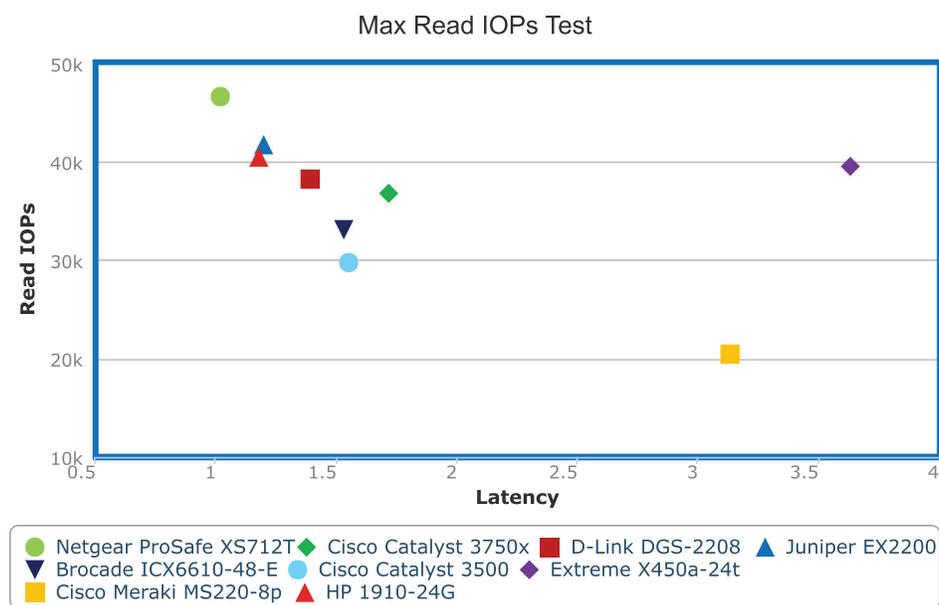


Chart B.



Results (cont.)

All samples indicated average latencies under 4 milliseconds, which is within acceptable tolerances for storage performance, therefore latency is excluded as a metric in the following ranking.

Rank	Switch
1 (Tie)	Netgear ProSafe XS712T
1 (Tie)	Juniper EX2200
3	Extreme x450a-24t
4	HP 1910-24G
5	Cisco Catalyst 3750X
6	Cisco Catalyst 3500
7 (Tie)	D-Link – DGS-2208
7 (Tie)	Brocade ICX6610-48-E
9	Cisco Meraki MS220-8p

Conclusions

The results illustrate that relatively good VSAN performance metrics can be achieved with a combination of mid-tier server hardware and any appropriate business-class networking hardware. This is a significant testament to VSAN's ability to lower overall cost of IT infrastructure, and VMware is validated in their lack of a switch HCL for VSAN.

As demonstrated in the results, 10 gigabit Ethernet (Netgear) provided only a slight advantage over the top 1 gigabit contenders. This pattern will likely hold in the majority of environments because the largest advantage conveyed by utilizing 10 gigabit switching for the VSAN backbone is for maintenance tasks, such as situations where the cluster initiates a rebuild operation on one or more hosts.

SynchroNet Best Practices

When configuring small to medium VSAN solutions, it is best to focus on host resources such as flash, memory, and CPU to improve performance rather than investing in high PPS and large-buffer network switches.

Low cost, unmanaged switches can be used for Proof of Concept or Pilots, but the lack of management and reliable support precludes their use in production.