



WHITE PAPER

SanDisk® SSD Boot Storm Testing for Virtual Desktop Infrastructure (VDI)

August 2014

Table of Contents

- Introduction 3
- Deployment Strategy 3
- Testing Process..... 4
- Test Execution and Results..... 5
- Test Result Observations..... 8
- Conclusion..... 9

Introduction

For virtualized desktops, Virtual Desktop Infrastructure (VDI) administrators rely on underlying storage performance to meet the necessary storage I/O (input/output) operations that become centralized when desktop operations move from the client-side flash or hard disk drives (HDDs) in laptops or desktops, to a centralized storage repository .

Fast response times from underlying storage for VDI are important to deliver productive end-user experiences . Prolonged or short spans of insufficient storage I/O during peak workloads can result in poor end-user satisfaction and the inability to perform end-user tasks .

Because VDI storage read/write operations are centralized, selection of underlying storage is critical to satisfy heavy I/O performance, such as boot storms, for end-user productivity and satisfaction.

To help VDI administrators understand the performance characteristic of solid state drives (SSDs) and the benefits, SanDisk® provides this lab validation report to highlight how VDI administrators may leverage SanDisk flash storage products to satisfy the unique read-intensive storage I/O requirements of VDI boot storms relative to traditional HDDs .

Deployment Strategy

VDI boot storm storage I/O requirements are much higher than steady state . Proper VDI architectural planning is required for successful deployment of the VDI environment.

From a VDI deployment perspective, a desktop pool could be persistent or non-persistent; both have their own use cases and are beyond the scope of this paper. In either case, if a boot storm storage peak demand is not properly sized, the overall performance and experience will not satisfy end-user productivity.

From the SanDisk testing perspective, we created a stateless (non-persistent) desktop pool where a single master image is used to deploy the end-user desktops . In this case, there is one golden copy, and delta drives (linked clones) are created for each user desktop from this master image .

When boot storms occur, all users read from a single master drive to boot-up, thus creating the aforementioned “boot storm” that is common in VDI deployment. It is the recommendation of SanDisk that the master drive image be placed on the SSD in order to address the boot storm issue.

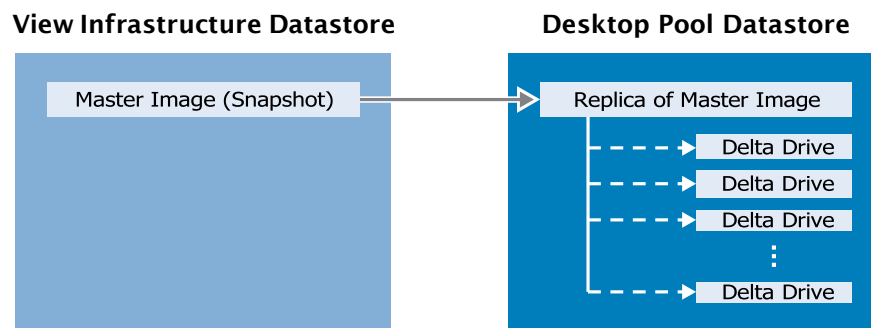


Figure 1: Typical Linked Clone Pool Deployment Strategy

Testing Process

VMware Horizon View™ environment is used to test the boot storm . Two VMware® Virtual Machine File System (VMFS) datastores were created virtually, one on HDDs and one on SSDs, to deploy the Horizon View user desktops . A desktop pool of 50 virtual machines (VMs) was created in each datastore to measure the boot storm impact.

To measure the impact of boot storms, this datastore was recreated each time with one, two, four and eight drives in the backend, with a RAID 0 configuration for both the HDD- and SSD-backed datastore . Caching is disabled at the storage controller so that each time these 50 desktops are powered on simultaneously, there is no cache data available at the storage controller layer and the master image (desktop operating system (OS) drive) is read from the drives configured in the backend . For convenience, the VMware Horizon View infrastructure was created on a separate datastore so that each time we redeploy the 50-VM desktop pool for a different number of drives, the whole environment does not need to be rebuilt.

Each time the datastore is recreated, the Horizon View pool is redeployed on the HDDs or SSDs datastore where the boot storm test was run.

VMware View™ Planner, VMware's proprietary VDI workload generator tool, is used to power-on these desktops simultaneously . The storage I/O and latency measurements were done by capturing the "esxtop" performance data in batch mode while the desktops were booted .

Complete boot time for each pool of 50 desktops is measured by using the VMware View Planner total boot time measurement script.

For the SSD used in this test, we selected our Optimus Ascend™ SAS SSD which is designed for typical mixed-use application workloads . Through a combination of unique and innovative IP, and the use of enterprise Multi-Level Cell (eMLC) flash, SanDisk's newest generation of 19nm SSDs feature a native SAS 6Gb/s interface, outstanding performance metrics, and a comprehensive set of high-end features making them ideal to integrate with existing infrastructure for a wide variety of enterprise environments such as servers, external storage arrays, and storage area networks (SAN) where performance and high-availability are required .



Test Configuration A




Figure 2 (B): Eight HDDs and SSDs are configured in storage controllers 1 and 2

Test Configuration B

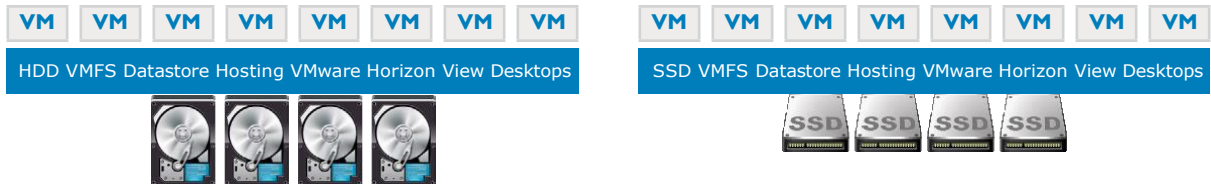


Figure 2 (B): Four HDDs and SSDs are configured in storage controllers 1 and 2

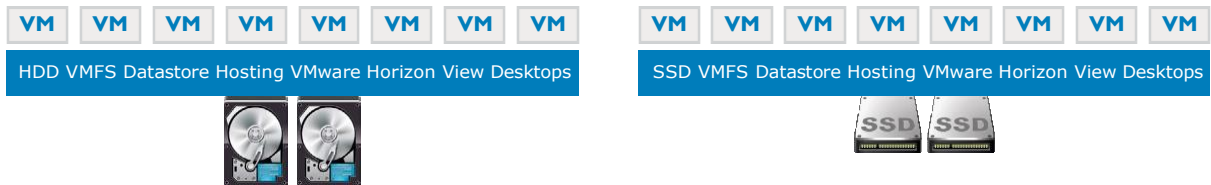


Figure 2 (C): Two HDDs and SSDs are configured in storage controllers 1 and 2

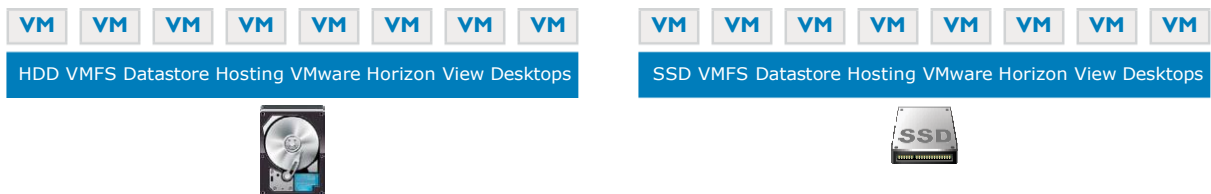


Figure 2 (D): One HDD and SSD are configured in storage controllers 1 and 2

As shown in the above figures, the server is powered down each time the drives' configuration needed to be changed. Once reconfiguration is done, the server is powered on and then the desktops' pool is recreated on top of that particular storage profile and a boot storm test for 50 VMs is executed.

Test Execution and Results

Following are the test results of 50-VM simultaneous boot time. The total boot time for the boot storm is measured using the View Planner time measurement script.

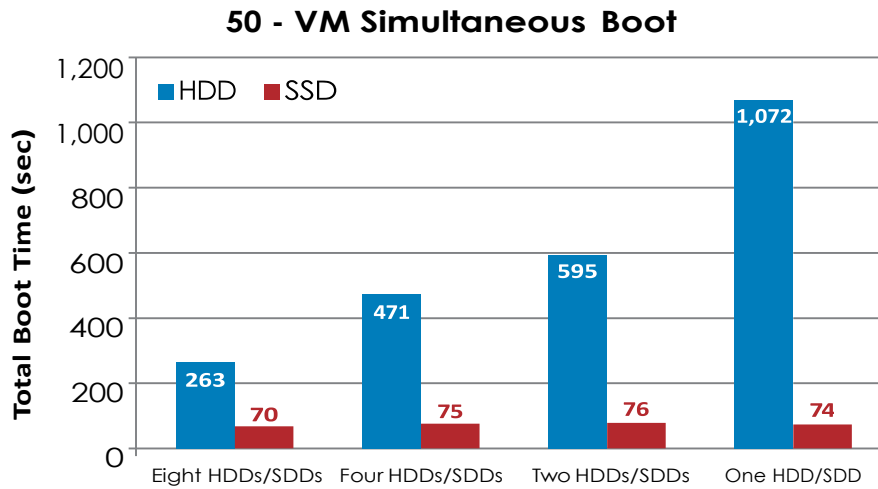


Figure 3. Total boot time of 50 VMs

As we can observe from the above total boot time bar graph, the total time taken for 50 VMs to boot-up is almost the same for each of the four SSD test combinations. Below the graph shows comparison of one SSD vs eight HDDs storage results while the boot storm is occurring in the server. The below measurements were made at the adapter level to accommodate all 50 VMs concurrent operations.

Drive IOPS

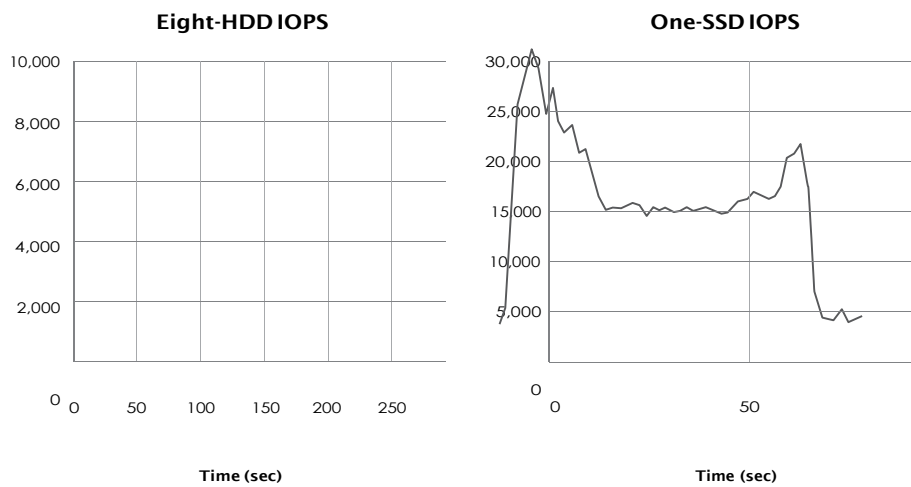


Figure 4. 50-VM drive IOPS during boot storm operation

Drive Latency

Drive Throughput—Read Throughput (MB/sec)

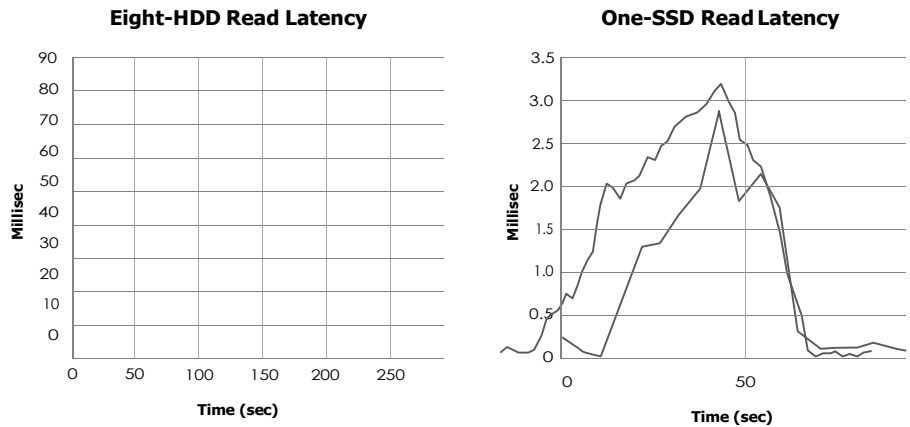


Figure 5 (A): 50-VM read latency during boot storm operation

Drive Throughput—Write Throughput (MB/sec)

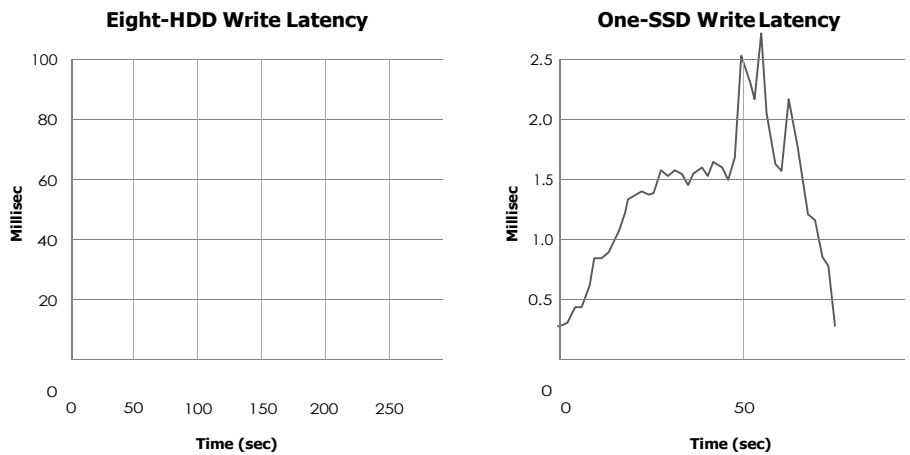


Figure 5 (B): 50-VM write latency during boot storm operation

The above charts show that the read latency is 30x lower and write latency is 50x lower during the boot storm for the SSD compared to the HDDs. Low latency is an important SSD benefit and is what provides the faster boot time of the compared user desktops.

High latency limits the IOPS of each HDD when all the VMs tried to read the OS image at the same point in time, whereas for the SSD; very low latency could drive high IOPS without any issue.

Drive Throughput—Read Throughput (MB/sec)

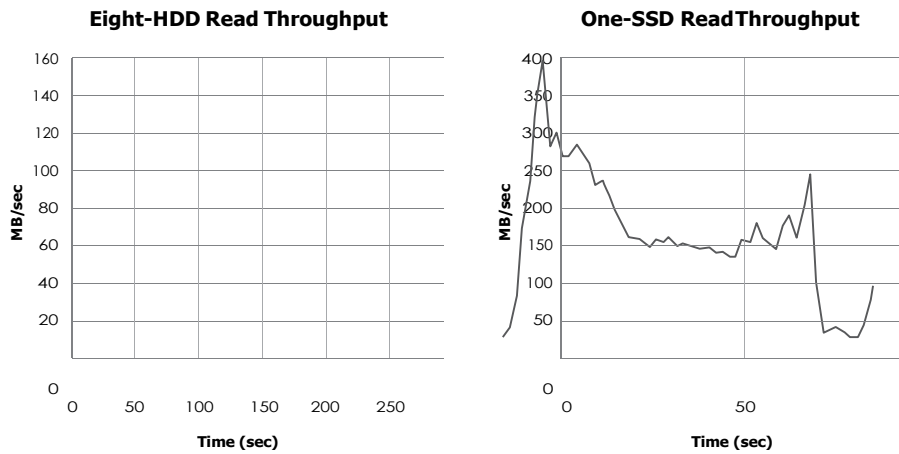


Figure 6 (A): 50-VM read throughput during 50-VM boot storm

Drive Throughput—Write Throughput (MB/sec)

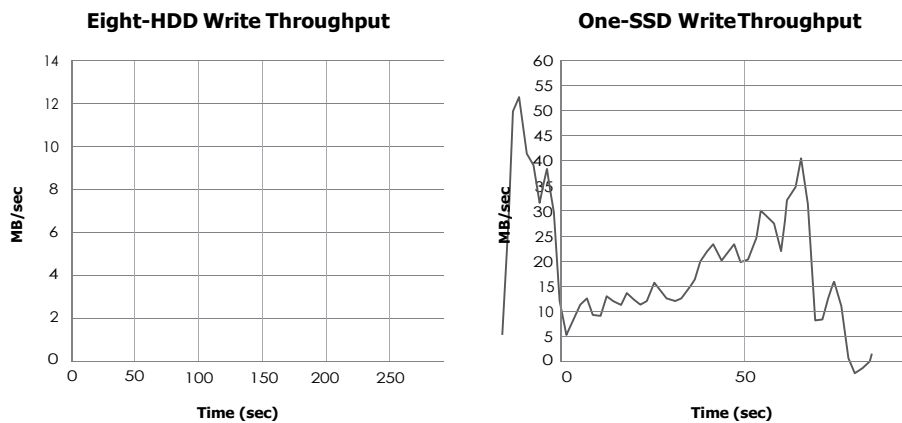


Figure 6 (B): 50-VM write throughput during 50-VM boot storm

From the throughput charts, we have a similar observation. As the latency in the SSD is much lower than the HDDs, it can drive higher throughput thus enabling quicker availability of end-user desktops for usage.

Test Result Observations

From the above results, we can observe that the SSD boot time for 50 VMs is 4X faster than the best possible deployment of eight HDDs. Though we started off with eight SSDs for the same configuration, it has been found that after testing of the one SSD, the performance is good enough to address the boot storm scenario. Also it can be seen that the latency for the single SSD is much lower than the HDDs and hence the SSD IOPS and throughput are higher.

Test Bed

The following tables describe the test bed used for the testing:

ESXi Host and Storage Configuration

Hardware	Specification
Servers	<ul style="list-style-type: none"> • 2-socket 8-core Intel® Xeon CPU E5-2667 v2 @ 3.30GHz • 256GB of RAM
Storage	<ul style="list-style-type: none"> • 15KRPMHDD • Optimus Ascend™ SAS SSDs from SanDisk®

Table 1. ESXi host hardware

Software Configuration

Software Installed
<ul style="list-style-type: none"> • VMware vSphere® 5.5 • Win 7—64bit OS for Desktop • VMware Horizon View™ 5.3 • Windows® 2008 R2 for View Infrastructure • VMware View™ Planner 3.x

Table 2. Software installed for swap-to-host cache testing

Conclusion

VDI is a type of workload that needs different drive sizing strategy at different times. As we saw, the boot storm is a read-intensive operation, whereas when the user desktops are in a steady state, it becomes a more write-intensive than read-intensive operation.

It is important that each of these unique requirements be addressed in its own way to have a successful VDI deployment. Other VDI administrative operation such as pool creation, recompose of pool, anti-virus patch updates, etc., need similar attention for the right drive sizing and deployment.

It's very clear from the boot storm study that the SSD solution is the key for a successful deployment of VDI desktops. If we delve into the details of the study, we can see that for a pool of 50 user desktops, eight HDDs are not good enough to have an acceptable boot time service level agreement (SLA). As we scale more and desktop density increases in a given server, the individual and total boot time of all desktops will be significantly higher, leading to failure of VDI success.

Since boot storm management is one of the critical administrative operations in the VDI environment, a detailed cost analysis will not make a significant impact unless all required sizing parameters are considered.

Specifications are subject to change. ©2014 - 2016 Western Digital Corporation or its affiliates. All rights reserved. SanDisk and the SanDisk logo are trademarks of Western Digital Corporation or its affiliates, registered in the U.S. and other countries. Optimus is a registered trademark of Western Digital Corporation or its affiliates. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s).

Western Digital Technologies, Inc. is the seller of record and licensee in the Americas of SanDisk® products.

