



WHITE PAPER

# SanDisk® Lab Validation: VMware vSphere® Swap-to-Host Cache on SanDisk SSDs

August 2014

**SanDisk®**  
a Western Digital brand

Western Digital Technologies, Inc.  
951 SanDisk Drive, Milpitas, CA 95035

[www.SanDisk.com](http://www.SanDisk.com)



## Table of Contents

Introduction .....	4
Overcommitting Memory .....	4
Deployment Strategies on Memory Overcommitment.....	5
Testing Process and Workload Used .....	6
Test Execution and Results .....	7
Test Result One.....	7
Test Result Two.....	9
Test Result Observations .....	11
Application SLA.....	11
VM Density Increments .....	11
Test Bed .....	13
ESXi Host and Storage Configuration.....	13
Software Configuration .....	13
Testing Conclusion .....	14
Acknowledgments .....	14
Resources .....	14
SanDisk Resources .....	14
VMware Resources .....	14

## Introduction

VMware® virtualization solutions have proven to be invaluable for enterprises running business and mission-critical applications. Today, VMware solutions are key enablers for business agility and flexibility to help enterprises quickly deploy applications to take advantage of new business opportunities and improve efficiencies.

Prior to server virtualization, data centers deployed a single application per server, which led to under-utilized servers and the inability to optimize IT resources. Adoption of VMware server virtualization has enabled better utilization of IT resources by pooling storage, memory, and networking to enable applications to leverage the appropriate computing resources when required.

With continued business requirements to maintain application peak performance and the addition of more business and mission-critical applications on virtualized infrastructure, further technological advances are needed to meet IT service-level agreements (SLAs) of application performance, the processing of larger data sets, and real-time access to analytics and file sharing.

SanDisk® documents its advances in flash storage performance when used in a virtualized environment with VMware “Swap-to-Host Cache” enablement.

SanDisk solid state drive (SSD) flash storage increases overall system performance while enabling much higher virtual machine (VM) host-to-server ratio because of the inherent I/O (input/output) performance characteristics of SanDisk SSDs.

VMware customers using SanDisk SSDs are able to achieve the following benefits:

- Consolidation of existing VMs onto fewer physical hosts
- Drive \$/VM costs down
- Better memory pool utilization through memory overcommitment
- Accommodate application memory demand spikes

## Overcommitting Memory

Overcommitting memory—i.e., when the total memory utilized by VMs running on a vSphere host exceeds the physical memory on that host—is a common practice in VMware environments. VMware provides several advanced memory management technologies such as transparent page sharing (TPS), ballooning, compression and memory swapping to manage memory overcommitment. When memory swapping occurs, the impact of swapping is manyfold, unlike physical environments, because many VMs are running on the host. Applications running in each VM will degrade drastically versus in a physical host case, where only one application will suffer.

Overcommitting memory provides opportunity for the VMware administrator to increase VM density and reduce the cost per VM. However, if application SLAs are not met, such a feature will not add much value to the deployment strategy. The silver bullet would be to overcommit memory by adding more VMs as much as possible, yet control application performance degradation such that applications’ SLAs are still met. This dual benefit of increased VM density and meeting application SLAs at the same time will improve the total cost of ownership (TCO) and return on investment (ROI) in a virtualized environment.

VMware vSphere's swap-to-host cache feature can help address this need. Using SanDisk SSDs as a memory swapping area for all these VMs will make sure that memory swapping is fast enough not to severely impact the application performance, and help increase VM density.

## Deployment Strategies on Memory Overcommitment

When memory utilization exceeds the available capacity in a given VMware host, it extends the memory to the drive (spinning media, hard-disk drive (HDD) or flash media, solid-state drive (SSD)) and helps manage the memory pressure. This memory pressure is typically managed by advanced memory management technique such as TPS, ballooning, compression and swap-to-host cache. The first three techniques try to manage when the pressure is building up, but once the available limit is exceeded, swap-to-host cache is the only way to manage it.

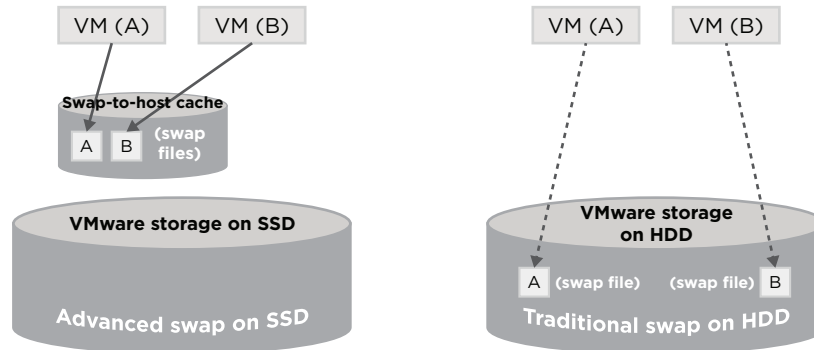
Swap-to-host cache can be configured only to a SSD. It can be configured either as a dedicated drive or as part of it. It is important to understand that partial assignment of a SSD has a great benefit. As overcommitment occurs, VMs not only suffer from memory pressure but there could be occasions when it might observe drive delay (latency) and running the VMs on SSDs will further improve VM performance. Partial assignment of drives not only helps address the overcommitment but allows the VM to run on the same SSD, enabling both purposes and helping reduce cost.

Enterprises can greatly benefit by introducing SSDs and adopting the right deployment scenario. Based on the application workload (mixed-use, read-intensive, or write-intensive) the correct workload-optimized SSD can be chosen. The same SSD can act as swap-to-host cache as mentioned above.

SanDisk offers a complete portfolio of SSD solutions for various workloads, endurance levels and storage interfaces such as SAS, SATA, PCIe and DIMM. Enterprises can choose from a wide range of available SSD products to address a particular data center need.

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.





**Figure 1: VMware swap-to-host cache vs. traditional HDD swap**

In the above figure we can see how swap-to-host cache can be carved out from a VMware storage volume and reap the benefit of the SSD. Though in either case, VM and its swap file can reside together in the same volume, application performance is significantly reduced when it is configured in the traditional way. The following section will show the results of the tests carried out.

## Testing Process and Workload Used

The testing was carried out at the SanDisk Technical Marketing lab. There are many workloads, and multiple possibilities exist to create a memory overcommitment scenario, with the scenario being completely dependent on an individual workload. From a deployment perspective, the user needs to test the application, and based on its behavior, application performance under memory overcommitment will vary.

For our tests, we used the DVD-store SQL v2.1 workload, online E-commerce load generator tool. We installed an SQL server and configured the default instance of the database in each VM. Each VM was also running an individual DVD-store workload driver.

A 20GB SQL database data file was created in each VM and the load driver ran the workload on it for a 20 minute duration to gather the operation-per-minute (OPM) value.

We ran Memtest ISO VM to saturate ESXi host memory so that the DVD store VM starves for memory and can therefore generate the need for swapping. We reserved the Memtest ISO VM memory so that swapping (memory overcommitment) occurred only in the DVD-store VMs. Memtest ISO VM is the key for saturating the memory of the host. The host memory was 128GB and creating a lot of VMs to saturate the host memory is a longer testing process. So we created a large size of Memtest ISO VM and 2-3 VMs where the actual workload was running. We created each DVD Store VM of memory size as 8GB. We observed that when the Memtest ISO VM is configured to run at 113GB, swapping starts in the host with two DVD-store VMs. We changed the Memtest ISO VM memory size to increase the swapping pressure on the DVD-store VMs and measured the impact. Later we added a third DVD-store VM of 8GB size and measured the incremental impact on increased VM density and swapping.

The overcommitment range was chosen based on a VMware study (please refer to the Resources section for details *VMware Customer Study on Memory Overcommitment*). As pointed out in the study, the performance impact is roughly a 1x to 3x memory overcommitment level which is typically used by VMware customers. We used this information and designed our experiments around these parameters.

Memory overcommitment was calculated by the following formula:

$$\text{Memory overcommitment} = \frac{\sum \text{VM Memory}}{\text{ESXi Host Memory} - \text{vSphere Overhead} - \sum \text{VM Overhead}}$$


VMware’s Memstat utility was used to capture the swapping of data and the esxtop utility was used for monitoring the CPU, memory, drive and network utilization while the workload was running in the VM.

## Test Execution and Results

Following section discusses the test results in detail for different configurations.

### Test Result One

In this instance we configured the Memtest ISO VM as 113GB, 118GB and 121GB. Two of the DVD store VMs were configured with 8GB RAM.



**Figure 2: Operation per minute (OPM) with minimal (1.24x) memory overcommitment**

Figure 2 shows how the OPM value was impacted during the test run. In this test the Memtest ISO VM was configured to run at 113GB and two of the DVD-store VMs were configured at 8GB each and were running the DVD store workload. As we can see, OPM started at the same level when the test began, and as memory overcommitment kicked in, the OPM value started falling from its value when there was no memory pressure. It is worth noting here that with a small amount of overcommitment, the SSD and HDDs swapping results in a similar amount of impact.


Let’s discuss how we have calculated the memory overcommitment ratio. In the above test:

Size of each VM = 8GB (Unreserved)

Size of Memtest ISO VM = 113GB (Reserved)


$$\begin{aligned} \text{Memory overcommitment} &= 16 / (128 - 113 - 2.1) \\ &= 1.4 \end{aligned}$$

vSphere overhead and VM overhead were obtained from the VMware utility esxtop during the test run. This variable differed in each run. For the remaining test run, we used the same methodology to calculate the overcommitment value.




**Figure 3: Operation per minute (OPM) with considerable (2.0x) overcommitment**

In Figure 3, we increased the memory of Memtest ISO VM to 118GB to further increase the memory pressure. Here we can see that the OPM value started falling much earlier and the swapping configured with a HDD started falling more than a SSD configured with swap-to-host cache.



**Figure 4: Operation per minute (OPM) with significant (3.4x) overcommitment**


In Figure 4, Memtest ISO VM was pushed up to 121GB. This can be considered as significant overcommitment. As we can see, swapping on HDDs failed to keep the pressure at all. Though SSD swap-to-host cache fell considerably, it still managed to have a much higher OPM value than swapping HDDs.



**Figure 5: Operation per minute (OPM) at different memory overcommitment**



In Figure 5, we have shown the bar chart comparison of OPM at different levels of memory overcommitment. With the increase of memory pressure, the OPM value falls. The OPM dramatically decreases when swapping occurs on the HDD in comparison to the higher OPM that occurs on the SSD in addressing the memory pressure.

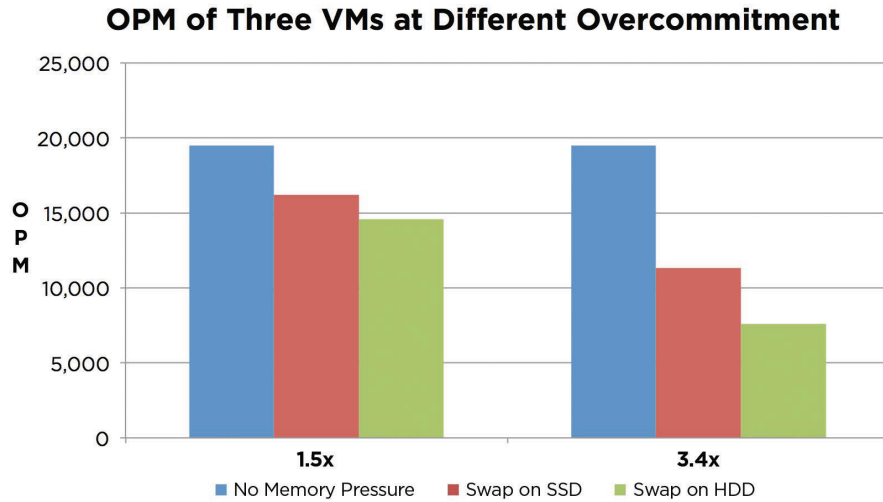


**Figure 6: Total amount of memory swap during test run**

Figure 6 shows the amount of memory swapped from an all-DVD-store VM during the test run. Please note that this swapping was only from the DVD-store VM as the Memtest-ISO VM was reserved with the memory and no swapping occurred from it. We captured this using the Memstat command and saw that swapping was always zero for the Memtest ISO VM.

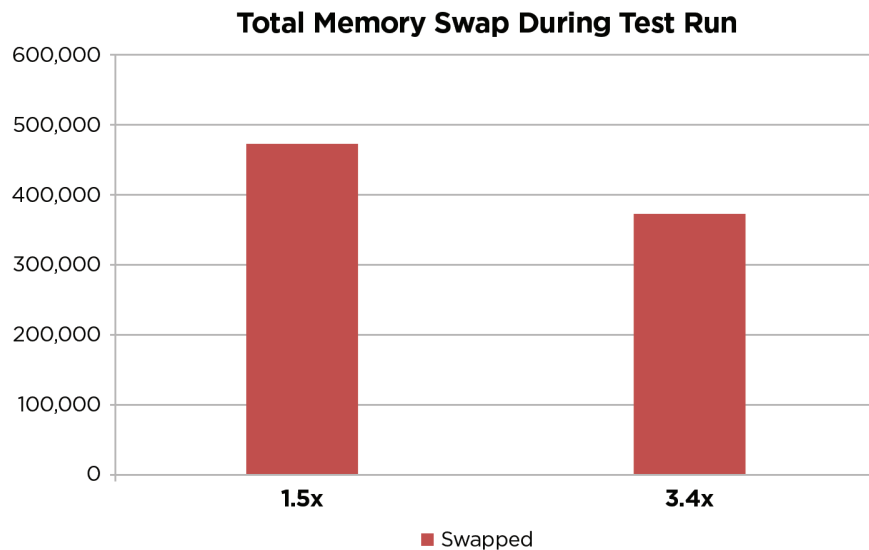
### Test Result Two

In these tests, we configured the Memtest ISO VM as 113GB and 118GB. Three of the DVD-store VMs were configured with 8GB RAM. The objective of this test was to see the impact when an additional VM was running in the host. This was to test whether an additional VM causes increased degradation or if it could maintain the OPM value when memory pressure increased. This way we could get a feeling of VM density impact on memory pressure. In this test, we ran only two combinations with Memtest ISO VM as 113GB and the other test at 118GB. The reason for not running the 121GB was that we saw the overcommitment level going way beyond the 3x range, which we wanted to restrict based on the VMware customer study of memory overcommitment.



**Figure 7: Operation per minute (OPM) at different memory overcommitment**

In Figure 7, we saw similar behavior as Test Results One and we observed that increasing the VM density did not impact the OPM value significantly. We have described this observation in detail in the next section, Test Result Observations.



**Figure 8: Total amount of memory swap during test run**

In Figure 8, we captured the swapping using Memstat, as we had for Test Result One. We again made sure—by analyzing the Memtest capture—that no swapping occurred in the Memtest-ISO VM and that all swapping was from the DVD-store VM.

## Test Result Observations

As shown above, we ran different combinations of tests to validate the swapping impact. One thing is clear: the swap-to-host cache using SSDs can manage the memory pressure extremely well as the memory pressure increases. Though application performance reduction in HDDs vs. SSDs during the start of overcommitment was minimal, it increased dramatically once overcommitment increased, and at some point the HDDs completely failed to maintain the pressure.

Many benefits were seen from the swap-to-host cache, but we will focus on two main areas.

### Application SLA

If we look at Figures 5 and 7, we can see OPM values at different overcommitment levels. The results show that with the increase in VM density, the OPM was not impacted when swapping occurred.

For example, if we look at Figures 5 and 7, the following table is derived from them:

Figure	No. of VMs	Total OPM – Swap on SSD	Memory Overcommitment	Per VM OPM (approx.)
5	2	11500	2x	4500
7	3	15500	2x	4500

We observed that at 2x memory overcommitment, the OPM value was the same. Though it was difficult to create the same memory pressure during testing for Figure 7, we conditioned the data between two data points of 1.5x and 3.4x. There could be little deviation from the actual but it definitely could be seen that allowing memory overcommitment increased VM density and the overall application SLA could be maintained.

Also if we look at the VMware study, we can see that average overcommitment is in the range of 1.8. It signifies that customers are packing more VMs in a host and maintaining the overall application SLA. The study proves this point very well, that memory overcommitment provides significant benefits towards meeting application SLAs.

### VM Density Increments

From an application SLA observation, it's been seen that VM density can be increased without impacting overall SLA. This shows that for clustered environments more VMs can be accommodated in a single host, thus reducing server hardware.

Typically, a server to VM ratio in a deployment is 1:8 to 1:10. This is a relative number and may vary. From our study, if the overall application SLA can be maintained in a memory overcommitment scenario—by increasing the number of VMs—we increased the consolidation ratio and thus reduced the overall number of servers needed as well as reducing the cost.

To illustrate this further we expanded the study mathematically from an application SLA where we saw VM density did not impact the overall OPM (application SLA) much. As mentioned in our testing process, it is tedious to execute real tests with many VMs running and we were restricted to a few VMs. The graph below shows the theoretical extrapolation from the given results in 2x overcommitment level.

In Figure 5, if we look at the 2x overcommitment numbers, we can see the approximate OPM for each VM.

Overcommitment Level	OPM for SSD Swap (each VM)	OPM for HDD Swap (each VM)
2x	4500	3400

For simplicity, let us ignore the memory overhead and use the following formula to calculate the total number of VMs which can be accommodated in each host during overcommitment.

$$\text{Memory overcommitment} = \frac{\sum \text{VM Memory}}{\text{ESXi Host Memory}}$$

In our case, host memory = 128GB

Memory of each VM = 8GB

Total number of VMs in each host when 2x overcommitment allowed = 32 VMs

	Total OPM - Swap in HDD	Total OPM - Swap in SSD
One Host	112000	144000
Three Hosts		432000
Four Hosts	435200	

As you can see in the above table, with the addition of one SSD as swap space, we could achieve the similar amount of OPM with a reduced number of hosts.

Figure 9 below shows the extrapolated result of our tests where we show that with 2x memory overcommitment, we could get equal OPM with three hosts when SSDs are used as swap space.

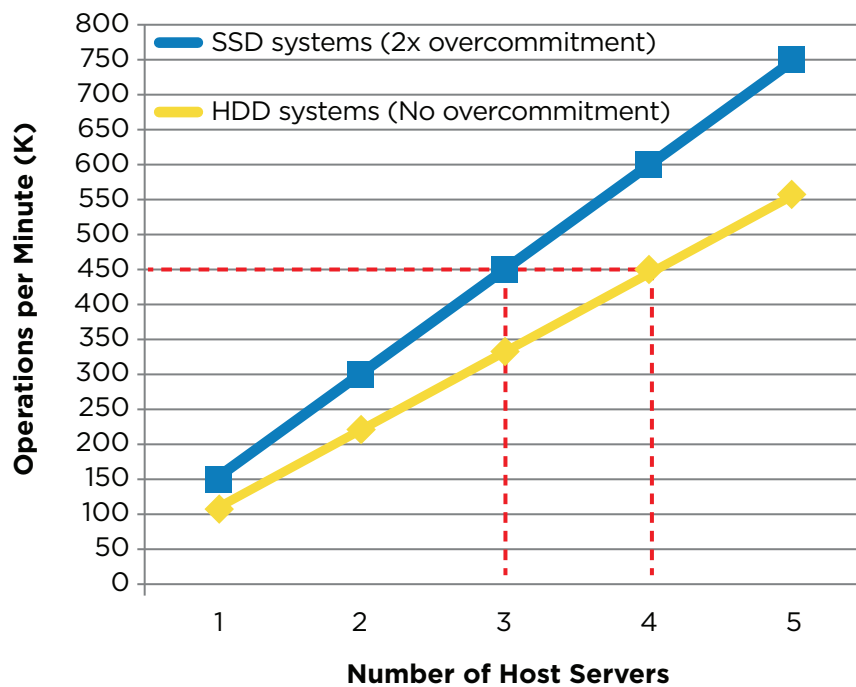


Figure 9: Total amount of memory swap during test run

It can be observed that if there is a clustered environment and many ESXi hosts are running the same workload (what we have tested) in each and every VM, with 2x overcommitment we could achieve the same overall OPM in three hosts instead of four.

In actual deployment, this will vary a bit. Individual VMs will be running different applications and overcommitment will vary based on many factors, such as the number of VMs, memory size and utilization in each VM, VM memory reservation, size of VMware host memory, host memory overhead, type of workload, VM memory overhead, etc. But one thing is clear; there is a definite benefit when SSDs are used for swapping in the VMware host.

In the above diagram, we have considered uniform workloads across each VM. It can be seen that a benefit of 25% can be achieved by reducing from four to three hosts. By allowing overcommitment to a considerable level (2x) practiced by enterprises as per the study, but considering the factors mentioned above, we can definitely expect a benefit of 10-15% with this feature when a varied workload is in use with different parameters.

### Test Bed

The tables below describe the test bed used for the testing.

#### ESXi Host and Storage Configuration

Hardware	Specification
<b>Servers</b>	<ul style="list-style-type: none"> <li>• Supermicro® 2-socket 8-core E5-2637 v2 @ 3.50GHz</li> <li>• 128GB RAM</li> <li>• Optimus Ascend™ SAS SSD from SanDisk® (370GB) configured for swap-to-host cache</li> </ul>
<b>Storage</b>	Dell® 300GB 10K RPM HDD for VM provisioning

**Table 1. ESXi Host Hardware**

#### Software Configuration

Software Installed	
• VMware vSphere® 5.5	• DVD Store v2.1
• Windows® Server 2008 R2 OS	• Memtest ISO VM
• SQL Server 2012 Enterprise Edition	

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

## Testing Conclusion

The testing suggests that enterprises are able to adopt an SSD solution for swapping and achieve their overall application performance requirements. VMware swap-to-host cache is an important feature which helps applications run in a better state during swapping. This benefit can only be achieved because of SSDs. Without them, the additional value is not realized.

Though we chose to run a particular workload and demonstrate the benefits of SSDs, this can be extended to other workloads (similar or different nature) as long as overcommitment is allowed based on VMware study.

We saw that VM density can be increased and overall application SLA can be achieved by allowing swapping in the SSD.

Ideally, an all-flash storage solution can address HDD issues but it is considered to be expensive. Swap-to-host cache is an opportunity to understand and realize the true benefits of SSDs. Implementing such features can help enterprises adopt more flash-based solutions and pave the way for embracing an all-flash storage infrastructure in a phased manner.

## Acknowledgments

The following groups helped in executing the tests and preparing the document:

- SanDisk Technical Marketing
- SanDisk Product Marketing
- SanDisk Outbound Marketing

## Resources

Customers can find more information about SanDisk and VMware via the links listed below:

### SanDisk Resources

- SanDisk official website: [www.sandisk.com](http://www.sandisk.com)
- SanDisk ESS Products: [www.sandisk.com/enterprise](http://www.sandisk.com/enterprise)

### VMware Resources

- VMware official website: [www.vmware.com](http://www.vmware.com)
- VMware memory management technique:  
[www.vmware.com/files/pdf/perf-vsphere-memory\\_management.pdf](http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf)
- VMware customer study on memory overcommitment:  
<http://blogs.vmware.com/virtualreality/2008/10/memory-overcomm.html>

For more information, please visit:

[www.sandisk.com/enterprise](http://www.sandisk.com/enterprise)

---

**SanDisk<sup>®</sup>**  
a Western Digital brand

At SanDisk, we're expanding the possibilities of data storage. For more than 25 years, SanDisk's ideas have helped transform the industry, delivering next generation storage solutions for consumers and businesses around the globe.

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 Ascend is a 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 holder(s). 06.22.16