



WHITE PAPER

Improving Microsoft Exchange Performance Using SanDisk® Solid State Drives (SSDs)

Hemant Gaidhani, SanDisk Enterprise Storage Solutions

Hemant.Gaidhani@SanDisk.com

Executive Summary

Enterprise datacenters have long relied on storage configurations based on hard disk drives (HDDs) to deploy mission-critical tier-1 applications. Due to the inherent performance limitations in HDDs, storage administrators have often adopted costly approaches such as overprovisioning storage or short-stroking HDDs to meet the service level agreements (SLAs) in place with these application owners.

SSDs can dramatically improve performance for these applications by delivering very high I/O performance at a very low latency. They do so without relying on mechanical devices, but rather by leveraging semiconductor memory. That memory is nonvolatile, so that power outages will not cause data to be lost; on the contrary, the data will be retained for later use.

Microsoft Exchange Server is one such tier-1 application – and it is pervasive in enterprise environments as an e-mail server and collaboration platform. Optimal storage configuration is a key to maximize the end user experience of Microsoft Exchange Server by reducing access times that would affect the customer experience. Both transactional input/output operations per second (IOPS) and latency must perform at the highest levels in order to successfully service large numbers of users accessing the mail server, without adversely affecting their productivity, especially during times of peak usage. These two parameters directly impact the usability of the e-mail system, as well as the number of mailboxes that can be hosted in an Exchange Server environment.

This white paper examines how SanDisk Enterprise SSDs can dramatically improve Microsoft Exchange Server performance and response times, resulting in a higher number of users per system, and in better user experience for all end-users. The paper also discusses how SanDisk SSD-based storage configurations also offer savings in terms of power, cooling and data center floor space requirements, leading to a reduction in overall Total Cost of Ownership (TCO) related to operational costs.

This document is intended for Microsoft Exchange administrators, and for storage administrators who are involved in the deployment of Exchange architectures.

Test Configuration

The section describes the detailed hardware and software configuration used for Microsoft Exchange server testing.

Hardware Configuration

A single Dell R720 server with two 6-core Intel Xeon E5-2620 2.00GHz CPUs and 64GB 1600MHz RAM was used for these tests.

Storage Configuration

The storage was configured as follows:

- C: drive 1x 500GB 7.2K SATA boot drive (RAID-0 for OS)
- D: drive 3x 300GB 15K SAS drives (RAID-5 for Exchange log files),
- E: drive 4x 400GB Optimus SSD (RAID-5 for Exchange database for SSD configuration) or
4x 300GB 15K SAS drive (RAID-5 for Exchange database for HDD configuration)

The SSDs were pre-conditioned to deliver sustained performance over time.

Software Configuration

The operating system used for this testing was Windows Server 2012 R2 Enterprise x64 (6.2.9200). The Jetstress 2013 Benchmark was used to simulate the Microsoft Exchange server operations.

A typical Microsoft Exchange server enterprise deployment in production includes a high-availability architecture with several mailbox servers in database availability groups that leverage data-replication for site resiliency. In addition, it also

employs anti-spam and anti-malware functionality, as well as third-party services such as mobile devices, search and indexing applications. These tests do not implement any of these features, because the objective of these tests was to simply find the maximum number of I/O operations per second (IOPS) that the storage configuration can support.

Jetstress configuration

The Jetstress 2013 benchmark is widely used to size Microsoft Exchange storage. Jetstress is a Microsoft software tool that simulates the Exchange database and log file loads that a specific number of users produce. This allows administrators to verify the performance and stability of the disk subsystem before deploying their Exchange server in a production environment. The primary storage performance test runs for two hours, and identifies the maximum sustainable Exchange I/O that the storage system can handle while delivering acceptable response times. This test should be regarded as mandatory for each Exchange server released into production.

Jetstress tests collect the results in terms of latency, IOs, and bandwidth for database and log operations. It then provides a Pass/Fail rating based on these results. The tests are marked as failed when the latency exceeds acceptable user experience. The values observed from these tests can be used both to qualify the solution ready for production and also to calculate available system I/O headroom once the service is in production.

Jetstress simulates Exchange database I/O load without requiring Exchange to be installed. To simulate the complex Exchange database I/O pattern effectively, Jetstress makes use of the same ESE.DLL that Exchange uses in production.

The following list provides the detailed Jetstress settings used during these tests:

Jetstress Version: 15.00.0775.000
ESE Version: 15.00.0516.026
Background Database Maintenance Tasks Enabled
Verify Checksum Enabled

Capacity and Throughput Percentage	100%
Minimum Database Cache	32.0 MB
Maximum Database Cache	256.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Number of Copies per Database	1

JetStress Results - Success Criteria

Microsoft Exchange transactional throughput and latencies are key performance parameters designated by Microsoft to determine the number of users that can be supported in a given environment, and in turn, how the user experience will be perceived. Jetstress, therefore, uses the following thresholds to determine if the test is successful. For the Performance Test in a “Strict” mode (i.e. test running ≤ 6 hour), the read and write latencies should be as follows:

Average Database Read Latency ≤ 20 ms

Average Log File Write Latency ≤ 10 ms

Max Database Read Latency ≤ 100 ms

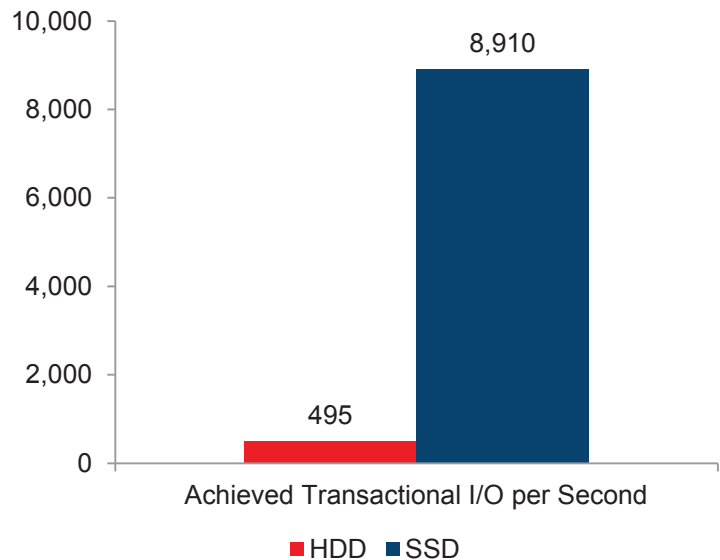
Max Log File Write Latency ≤ 100 ms

Test Results

The JetStress results demonstrate that using SanDisk SAS SSDs provides a dramatic increase in transactional IOPS for MS Exchange Server. Figure 1 shows that the transactional IOPS increased from 495 IOPS for HDDs to 8910 IOPS for SSDs – a factor of more than 18 times. This increase in IOPS performance was accompanied by an even better 25-fold decrease in the database read latency (from 18.3ms to 0.67ms), significantly lower than the Microsoft acceptable requirement of 20ms.

Figure 1: Jetstress Transactional IOPS comparison

Figure 1 illustrates this 18x increase in transactional IOPS performance. For Exchange administrators, this translates into support for many more users in their environment as well as the capability to support heavier user profiles and usage peaks without affecting end-user response times.



Figures 2 and 3 compare the read and write throughput between the HDD and SSD based configurations. As expected, the number of read and write I/Os increased very significantly for both Exchange database and log files.

Figure 2: Jetstress Exchange Database Reads and Writes Throughput comparison


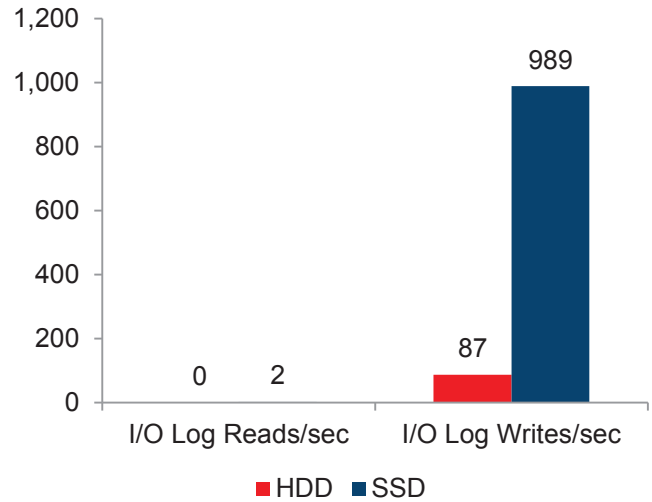


Figure 3: Jetstress Exchange Log Reads and Writes Throughput comparison



The usage of SSDs dramatically reduces the latency of the read I/Os. In order for the JetStress test to be successful, the database read latency should be below 20ms. Figures 4 and 5 illustrate the improvements in database and log latencies with SSDs compared to HDDs. The database read latencies improved by 25 times when using the SanDisk SAS SSDs. The HDD configuration did not scale beyond the reported results because it failed the Microsoft recommended threshold requiring database read latencies of less than 20ms. With SSD configuration, latency was less than 1ms and hence easily passed Microsoft’s requirements.

Figure 4: Jetstress Exchange Database Reads and Writes Latency comparison

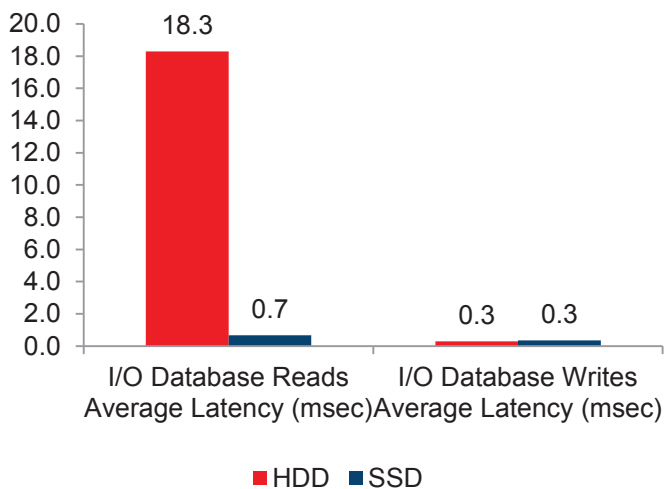


Figure 5: Jetstress Exchange Log Reads and Writes Latency comparison

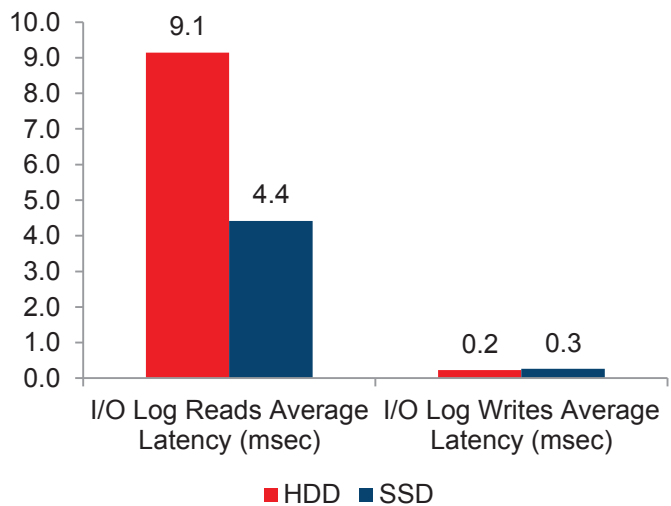


Table 1 below summarizes the improvements in transactional IOPS and I/O-related key performance indicators for HDDs and SSDs.

Metrics	4x HDD	4x SSD	Improvement
Achieved Transactional I/O per Second	495.15	8909.77	18x
I/O Database Reads/sec	322.59	6120.98	19x
I/O Database Writes/sec	172.56	2788.78	16x
I/O Log Reads/sec	0.185	2.032	11x
I/O Log Writes/sec	86.74	988.774	11.4x
I/O Database Reads Average Latency	18.30 msec	0.67 msec	25x
I/O Database Writes Average Latency	0.28 msec	0.34 msec	Same
I/O Log Reads Average Latency	9.14 msec	4.42 msec	2x
I/O Log Writes Average Latency	0.23 msec	0.26 msec	Same

TCO Analysis

The test results demonstrate that four SanDisk Optimus SAS SSDs delivered more than 18 times higher transactions IOPS, compared to four standard 15K RPM SAS hard disk drives. In other words, to deliver comparable number of transactions, 72 HDDs would be needed to do the same tasks. This equivalent HDD solution cannot fit into the drive slots that are available in a typical rackmount server, and would instead require up to 3 drive enclosures with 24 HDDs each. While the SSD configuration can leverage the existing drive slots, the HDD-based solution would require an additional 6Us of rack space to do the same amount of work. Based on SanDisk SSD pricing from the CDW website, the Solid State Drive (SSD) configuration using four (4) SanDisk Optimus 400GB SSDs will cost approximately \$5,600.

As of June, 2014, the Dell MD1220 enclosure’s starting price was \$4,000 each. The Dell (Toshiba) 15K RPM HDD cost was \$307.99. Using these numbers, the Hard Drive (HDD) configuration consisting of 72 Dell 300GB 15K RPM HDDs in three (3) Dell MD1220 enclosures would cost about \$34,200.

This cost analysis heavily favors using SSDs, showing the HDD configuration required to support same transactional IOPS as the SSD configuration is almost seven (7) times more expensive. Note that the larger number of HDDs is consistent with the industry standard practice of short-stroking HDDs to 10% of their available capacity, to achieve maximum performance for Exchange databases running on spinning hard-disk drives (HDDs).

The replacement of larger number of short-stroked HDDs with fewer SSDs can also deliver significant savings in power and cooling requirements. For example, many HDDs require 15 watts of power; however, the SanDisk SSDs require about half that much power – in the 5-7 watts range.

Furthermore, unlike HDDs, SSDs do not have any mechanical or moving parts. This reduces the likelihood of component failure, compared with HDDs. SSDs offer a more reliable and resilient solution compared to HDDs, with a Mean Time Between Failure (MTBF) of 2 million hours, SSDs offer a more reliable and resilient solution compared to HDDs. .

The potential savings in hardware, power and cooling costs clearly make SSDs a wise investment.

Summary

For Microsoft Exchange deployments, SanDisk SSDs can deliver unprecedented performance and Quality of Service (QoS) and better reliability than traditional HDD spinning disks for tier-1 mission critical applications. Enterprise IT departments can deliver on their Service Level Agreements (SLAs) with confidence, even for I/O-intensive applications like MS Exchange Server. The technical whitepaper demonstrates that using SSDs can increase MS Exchange performance by a factor of 18 times, at one-seventh the total cost of an equivalent HDD-based configuration.

This increased performance translates into a higher ROI because the same server infrastructure can now support a much larger number of users. Because latency is significantly lower with SSDs, end-user productivity is no longer adversely affected during peak usage. Furthermore, reduced MTBF and power and cooling costs realized when using SSDs can lead to efficiencies in capital and operating expenses, resulting in lower total cost of ownership (TCO) for customers.

References

Enclosure pricing from Dell:

<http://www.dell.com/us/business/p/powervault-md1220/pd>

SanDisk SSD pricing from CDW:

<http://www.cdw.com/shop/products/SanDisk-Optimus-solid-state-drive-400-GB-SAS-2/3266305.aspx>

Dell (Toshiba) MK3001GRRB HDD Pricing from CDW:

<http://www.cdw.com/shop/products/Toshiba-MK3001GRRB-hard-drive-300-GB-SAS-2/3010513.aspx?enkwr=ALLPROD%3a%7cMK3001GRRB%7cAll%20Product%20Catalog>

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. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their holder(s). 06.27.2016

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