

Evaluation Report: Microsoft Exchange for the Small-to-Medium Business on Dot Hill AssuredSAN 4824

Evaluation report prepared under contract with Dot Hill

Executive Summary

Email is one of the critical applications many businesses rely on. In addition to facilitating intra-company communications, many small to medium businesses rely on email as a method of communicating directly with customers, be it primary contact, confirmation of orders, support requests, or for other customer service functions. Adding the inevitable personal use and an irritatingly massive spam load, email systems can see a lot of use. It's vital for business continuity to know that corporate email remains operational and efficient 7x24.

The Dot Hill AssuredSAN 4824 is positioned in the market as a high performance hybrid storage system. Upgradeable with a customer configurable number of solid state drives (SSD), a base-model all hard disk drive (HDD) array can be enhanced to provide SSD read caching or an SSD performance tier to accelerate array performance. Dot Hill commissioned Demartek to evaluate the AssuredSAN 4824 storage array as backend storage for Microsoft Exchange Server, configured as it might be deployed by a typical small or medium sized business.

We ran Microsoft Exchange Jetstress 2013 on an all-HDD array as a baseline and again with SSD read caching and performance tier upgrades to simulate the Microsoft Exchange I/O load and database I/O patterns. We considered the Jetstress output reports along with Windows Perfmon data to evaluate the suitability of the AssuredSAN 4824 in these three configurations for this purpose.

We found that with the I/O load configured for a heavy usage simulation, the AssuredSAN array supported 500 mailboxes in an all-HDD configuration with acceptable maximum database latencies of 17.5 milliseconds. However, SSD tiering produced an impressive improvement on this baseline. The AssuredSAN array accelerated performance with the addition of SSD real-time tiering by reducing latency 75% while increasing write throughput by a factor of 4.4X.

The Dot Hill AssuredSAN 4824

The Dot Hill AssuredSAN 4824 Storage Array (Figure 1) is a dual controller, 2 rack unit, 24 small form factor drive storage array designed and manufactured by Dot Hill as part of the AssuredSAN 4004 product line. The 4824 model supports 8Gb and 16Gb Fibre Channel along with 1Gb and 10Gb iSCSI. Each array can support 7 expansion units for up to 384 TB of total storage capacity. The base configuration of 24 HDDs is upgradable by replacing some HDDs with flash storage.



Figure 1 - Dot Hill AssuredSAN 4824

RealStor™ is Dot Hill's next generation real-time storage operating system that supports SSDs in either a read cache or tiering operation. In the AssuredSAN 4824 array hybrid configuration, RealStor's patented autonomic tiering technology, RealTier™, dynamically stages high priority data on SSDs, with lower priority data on HDDs. This capability enables end-users to gain the benefits of flash performance without engaging in an expensive all-flash array.

For this evaluation, Dot Hill provided Demartek with an AssuredSAN 4824 array with 20 900 GB 10K RPM HDDs and 4 400 GB SSDs. With this hardware, we were able to set up a 20 HDD baseline test case, with 10 drives provisioned to each controller in a RAID 6 configuration. Then to evaluate read caching, we added a single SSD to each controller and assigned those SSDs as read cache disk groups. For measuring the effect of RealTier™, 2 SSDs in a RAID 1 disk group were added to each controller.

Microsoft Exchange Jetstress 2013

Microsoft Exchange is one of the dominant email server applications in use today. Many small and medium sized businesses that do not offload their email to third parties run their own Exchange servers. Email is a latency sensitive application (just ask anyone waiting for an important email), and therefore properly designing and scaling storage is critical to an acceptable user experience.

Microsoft Exchange Jetstress 2013 models the I/O loading and patterns that would be generated by an Exchange 2013 server, including email database reads and writes as well as logging. Its configuration includes the number and size of mailboxes, number of

mailbox databases and logs, and intensity of use, among other things. Jetstress uses the same Extensible Storage Engine files used by Microsoft Exchange to ensure that the Jetstress modelled performance is representative of a bona fide Exchange Server of the same version. The key metric used by Jetstress to pass or fail an email system is latency. Jetstress marks a test a failure if any Exchange database experiences an average transactional latency of 20 milliseconds or higher.

Microsoft recommends Exchange Jetstress be used to validate email systems before putting those systems into production.

Real vs. Synthetic Workloads

Demartek prefers to test storage systems with real workloads that users are likely to run on their systems as opposed to synthetic workloads that tightly control I/O. While Jetstress does not deploy a full Microsoft Exchange package, it is our opinion that it does a satisfactory job of simulating an operational Exchange environment. It achieves this by requiring the same physical hardware and configuration that would be used for the production email system as well as the same underlying software engine used by a full installation of Microsoft Exchange.

Workload Definition and Evaluation Objectives

Microsoft Exchange Jetstress 2013 exercises the storage system with read and write I/O to simulate email database I/O. We envisioned a hypothetical small-to-medium sized business with 500 mailboxes and sized the storage to support each mailbox up to 2 GB in size. To support this scale of mail server, we configured Jetstress to build 6 mailbox databases on 200 GB volumes, with corresponding log volumes. We spread these volumes equally across the two storage controllers of the array.

Jetstress can be configured to exercise all the storage space given it, and we did so. This has the potential to be somewhat cache unfriendly as it reduces hot spots on the drives, but we wanted to make this a difficult workload for the storage. We also configured Jetstress to perform .5 IOPS per mailbox, which is considered a heavy load for an Exchange 2013 email server.

We performed an initial 30 minute tuning run to allow Jetstress to build the database and identify a thread count, then a 2 hour execution at the proposed thread count followed by a 24 hour test to make sure that the system reached a steady I/O state and to allow the

read caching and RealTier™ scenarios time to fully migrate data to SSD (this is also the test method that Microsoft requires for publishing results through the ESRP program, though it was not our intention to publish outside of this report). The data and analysis included in this report is from the 24 hour test run. The databases were restored prior to beginning each new test.

The purpose of this test was to show performance comparisons in a typical SMB data center, not the maximum load. The environment with 500 mailboxes represents a common configuration to demonstrate the benefit of adding SSD as read cache or SSD as a tier in the array.

Performance Metrics

Key metrics for storage system performance analysis are I/Os per second (IOPS), bandwidth, and latency or response time. These metrics are defined as follows:

- ◆ **IOPS** – I/Os per second – a measure of the total I/O operations (reads and writes) issued by the application servers.
- ◆ **Bandwidth** – a measure of the data transfer rate, or I/O throughput, measured in Megabytes per second (MBPS).
- ◆ **Latency** – a measure of the time taken to complete an I/O request, also known as response time. This is frequently measured in milliseconds (one thousandth of a second). Latency is introduced into the SAN at many points, including the server and HBA, SAN switching, and at the storage target(s) and media.

It is important to consider all three metrics when evaluating the performance of storage systems because all three contribute to how the storage will support an application. IOPS drive bandwidth. The number of IOPS times the I/O request size determines the amount of bandwidth delivered.

Latency is important even though it doesn't necessarily have a direct effect on IOPS and bandwidth. It can have a very significant effect on application performance and user experience. Jetstress sets an upper limit of 20 milliseconds for database reads and writes. Unlike IOPS and bandwidth, where more is better, the goal with latency is to keep it as low as possible.

Jetstress collects its own performance metrics, which are viewable through Windows Perfmon. We converted these to CSV files for our analysis. Jetstress also performs its own

analysis to score test runs which we used to confirm that the tests were within Jetstress's parameters for a successful run and validation of the storage system.

Results and Analysis

Unlike some other applications, where a business can set its own standards for acceptable performance, Microsoft Exchange Jetstress expects to see database transaction response times below 20 milliseconds or it will fail a configuration. Our test of the Dot Hill AssuredSAN 4824, simulating 500 mailboxes on a heavily loaded email system, met the Jetstress response time requirements for all configurations: HDD-only, HDD plus an SSD read cache, and HDD plus a Dot Hill RealTier™ SSD performance tier.

Figure 2 records the highest average database transaction latencies measured by the Exchange Jetstress 2013 application. The highest latencies were consistently generated by read transactions. These values are the total of all system contributions to latency, including the storage device and SAN infrastructure plus server and application latencies. Scores above 20 milliseconds would prevent Jetstress from validating the system for Exchange 2013 with the number of

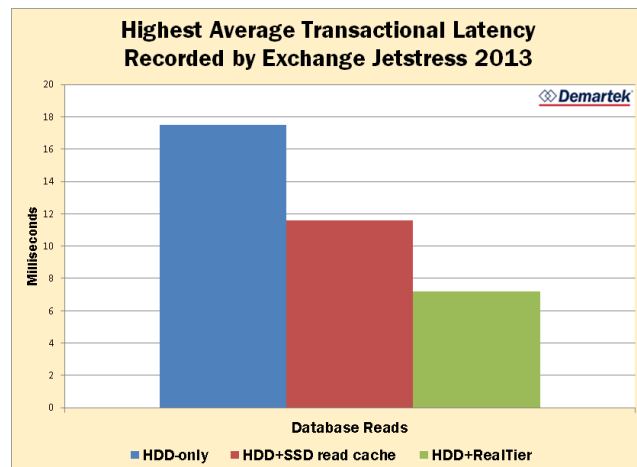


Figure 2 - Jetstress latencies

mailboxes and IOPS we configured. As can be seen here, the highest response time, at 17.5 milliseconds for the HDD-only configuration, came close to the maximum threshold, but still fell within acceptable limits. The flash accelerated configurations were even better. This tells us that the AssuredSAN 4824 is a suitable Exchange Server 2013 storage device for 500 rather busy mailboxes, even without any flash upgrades.

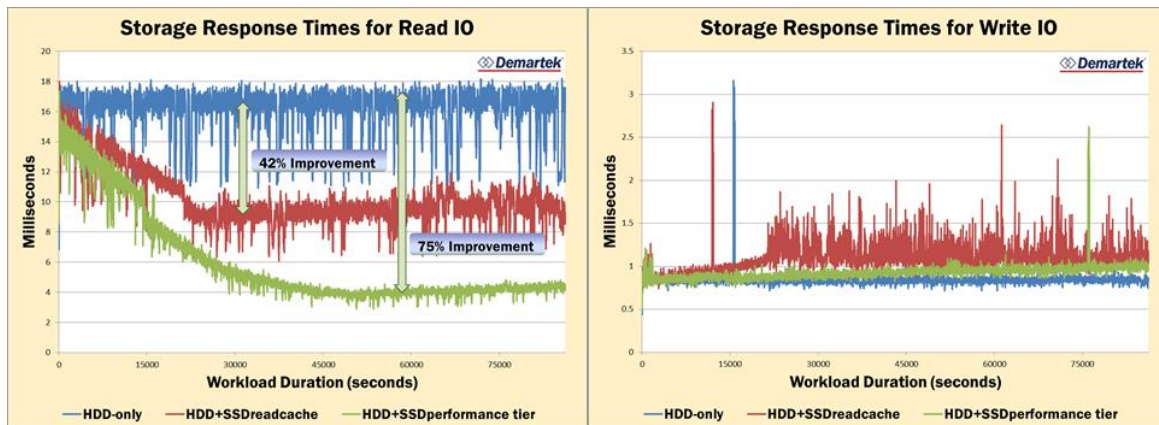


Figure 3 – Database volume response times

The data displayed in Figure 3 represents the average response times of I/O on the data and log volumes provisioned from the Dot Hill AssuredSAN 4824, not including application-introduced latency. Again, we can see that read I/O latency clearly received the largest benefit from flash upgrades in the storage. The baseline HDD-only configuration performed I/O transactions with a maximum average read I/O response time of 17.5 milliseconds, just about the same level as recorded for the full transactional latency in Figure 2. With the addition of SSDs for read caching and RealTier™ performance tiering, latency drops by 42% and 75%, respectively, to roughly 10 milliseconds and 4 milliseconds.

We can also see that database writes were never an issue, with every configuration maintaining a very low latency; the average being around 1 millisecond.

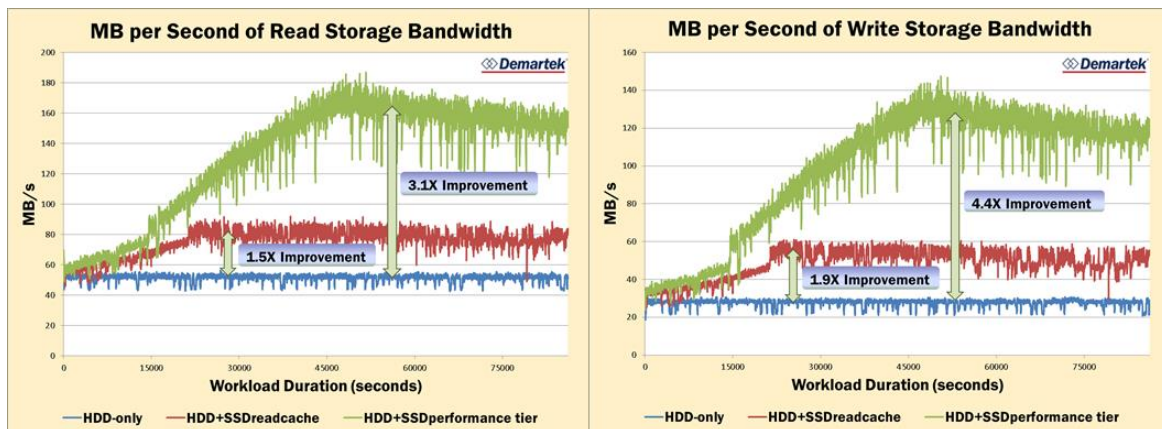


Figure 4 - Database volume bandwidths

Figure 4 demonstrates the corresponding increase in read/write bandwidths delivered by the AssuredSAN 4824 as we employed the flash upgrades. The workload is cache-unfriendly by definition. The Jetstress parameters were configured to minimize hot spots on the drives by spreading I/O across all provisioned drive space, and we can see that in the data. The abruptly halted cache warming curve and moderate performance gains of the HDD plus read cache configuration point to a strong degree of randomness in the I/O pattern, which can be hard to cache when the drive space in use is greater than the amount of cache provisioned.

In spite of this, read caching still provided a 1.5 times read bandwidth gain over a baseline of 51 MB/s and 1.9 times the baseline write bandwidth of 28 MB/s. If the Exchange server modelled had been larger, we would expect to see a smaller change over baseline unless more cache was deployed. This is not an indictment against read caching, but instead an acknowledgement that understanding a workload's I/O profile is important when selecting the appropriate storage solutions. What the data clearly indicate is that the real improvements for this workload come from the RealTier™ performance tiering solution.

The RealTier™ tiering algorithm and five second data migration interval seem to be quite effective in getting active data onto the flash tier for this workload. We recorded an average 3.1 times improvement in read bandwidth, up to 158MB/s and a 4.4 times improvement in write bandwidth, to 123MB/s. If we consider this along with the impressive improvement in latency, there seems to be a great deal of headroom in which we could add additional mailboxes or potentially another application.

Summary and Conclusion

AssuredSAN hybrid arrays with real-time tiering clearly improved performance in terms of reduced latency and increased read and write throughput for this test case. Based on Jetstress modelling, the AssuredSAN 4824 easily supports an email application with 500 very busy Microsoft Exchange 2013 mailbox users. While the addition of SSD's to improve performance is not required, we would advise businesses not to ignore the advantages that a hybrid array with real-time tiering offers over all-HDD configurations. Deploying RealTier™ SSD tiering will significantly reduce latency for a 500 mailbox environment and provide “future-proofing” in the event that the Exchange server usage goes up. A growing business would be wise to consider where its Exchange requirements might be in two to five years.

Another consideration for a flash enhanced AssuredSAN 4824 would be to support additional workloads. The environment we deployed in this evaluation did not demand anywhere near the entire drive space provided by the array. To keep latency below 20 milliseconds for the HDD-only array configuration, we are essentially short stroking the drives. This is wasteful in drive space and in return on investment of the array. RealTier™ SSD tiering improved latency by 75%, giving us the option to deploy more data in the form of additional mailboxes or an entirely different application altogether onto the array while still keeping latency within the recommended threshold for Microsoft Exchange. We would recommend that a business seeking to deploy storage in a cost effective manner give this option serious consideration.

Appendix A – Test Description and Environment




Figure 5 – Test Infrastructure

Server

- ◆ HP ProLiant DL380 Gen 8
- ◆ 2 Intel E5-2630 2.3GHz CPUs
- ◆ 16 GB RAM
- ◆ 1 HP SN1000E (16Gb FC) dual port HBA
- ◆ Microsoft Windows Server 2012 R2
- ◆ Microsoft Exchange Jetstress 2013, Microsoft Exchange 2013 ESE files

Fibre Channel Switch

- ◆ Brocade 6510 16Gb Fibre Channel Switch

Storage Array

- ◆ Dot Hill AssuredSAN 4824 array
- ◆ RealStor™ Storage Operating System

- ◆ 24 900GB 10k RPM 6Gb SAS HDD
 - ◇ 10 drive RAID 6 per storage controller
 - ◇ 6 Data volumes – 3 per controller
 - ◇ 6 Log volumes – 3 per controller
- ◆ 4 400GB SSD
 - ◇ 1 drive per storage controller for read caching
 - ◇ 2 drives RAID 1 per storage controller for RealTier
- ◆ 4 16Gb FC target ports per controller (1 port active per controller)

The original version of this document is available at

http://www.demartek.com/Demartek_Dot_Hill_AssuredSAN_4824_Improving_Exchange_Server_Performance_2015-08.html on the Demartek website.

Dot Hill, AssuredSAN, RealStor and RealTier are trademarks or registered trademarks of Dot Hill Systems Corp.

Intel and Xeon are registered trademarks of Intel Corporation.

Microsoft, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Demartek is a trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.