Exchange 2013 reference architecture using Dell SC4020 and Windows 2012 R2 Hyper-V

Dell Storage Engineering Solutions
October 2014

This document has been archived and will no longer be maintained or updated. For more information go to the Storage Solutions Technical Documents page on Dell TechCenter or contact support.

A Dell Reference Architecture

Revisions

Date	Description
October 2014	Initial release

Table of contents

Executive summary

Microsoft Exchange Server 2013 storage can be provided in many different ways. This paper presents instructions for obtaining higher mailbox quotas along with first class SAN features such as storage virtualization, Data Instant Replay snapshots, and automatic data tiering while keeping per-mailbox costs down by utilizing the smaller-scale Dell SC4020 storage array.

SC4000 Series arrays are based on the Dell flagship SC8000 platform. Offering similar benefits at a smaller scale, the multiprotocol-capable SC4020 all-in-one array and its optional expansion enclosures may be populated with any combination of hard disk drives (HDD) or Solid State Disks (SSD). The SC4020 dual-controller array features a 2U all-in-one chassis with easy expansion via standard Dell SC Series drive enclosures (up to 120 drives). Virtualized multi-tier, multi-RAID-level storage policies are applied quickly and automatically, taking full advantage of the unique characteristics of the disk, and allowing for target application-specific price and performance requirements with minimum of planning or effort.

Dell SC4020 provides an easily scaled-out storage platform on a versatile virtualization infrastructure supported by Microsoft Windows Hyper-V technologies. This solution is able to address such challenges through a modular or building block approach.

The ability of an SC storage solution to seamlessly scale from one to multiple drive enclosures or additional arrays provides an easy-to-deploy resolution for a sudden or planned change in the storage demand of both capacity and performance. Virtualization technologies assist with growth by providing automation mechanisms and flexibility in provisioning new resources and in distributing the workloads across the virtual infrastructure.

The solution presented in this paper focuses on presenting a moderately heavy-messaging workload for 4500 mailboxes with large storage quotas (6GB) on Microsoft Exchange Server 2013 in a virtualized infrastructure based on Hyper-V technologies and with a back-end storage on Dell SC4020 with 7K drives. Tested scenarios in this document include:

- A Microsoft Jetstress validated 7K only disk solution for 4500 6GB mailboxes using RAID 6
- The storage response to single disk failure in a RAID6 solution during degraded mode and RAID rebuild

1 Introduction

The reference architecture, validated for the purpose of this paper, included some of the most common planning variables (mailbox size and user profile) for a Microsoft Exchange solution. The deployment demonstrated a moderate number of users and portrayed an example of scaling the capacity and performance up and out on the back-end SAN. Besides these exercises, an assessment was made of the response to the simulated failure of a single drive in the RAID 6 disk page pool.

1.1 Purpose and scope

This paper is primarily intended for IT professionals (IT managers, solution architects, Exchange and storage administrators, and system and virtualization engineers) who are involved in defining, planning, or implementing Exchange Server infrastructures and who would like to investigate the benefits of using Compellent storage. This document assumes the reader is familiar with Exchange Server functions, Compellent SAN operation, and Microsoft Hyper-V architecture and system administration. The scope of this paper is restricted to a local datacenter topology and does not include specific or detailed server sizing information.

1.2 Terminology

The following terms will be used throughout this document:

Disk Folder: Disks are grouped into disk folders to provide a virtualized pool of storage that can be used by volumes as needed. The assigned disk folder is where new disks can be placed to become part of the usable storage. Custom disk folders can be used to split disks into different resource pools.

Hypervisor: Denotes the software layer that manages the access to the physical host hardware resources, residing above the hardware, and in between the guest VM operating systems.

Virtual Machine (VM): An operating system implemented on a software representation of hardware resources (processor, memory, storage, network, and others). VMs are usually identified as guests in relation with the host operating system that executes the processes to allow them to run over an abstraction layer of the hardware.

Balanced tree (B-Tree): A tree data structure where a node can have a variable number of child nodes. This structure is commonly used in databases to maintain data sorted in a hierarchical arrangement. It allows efficient data access to the pages for insertion, deletion, and searches.

Key performance indicators (KPI): A set of quantifiable measures or criteria used to define the level of success of a particular activity.

2 Solution architecture overview

The solution presented and evaluated in this paper is built on a virtual infrastructure supported by Microsoft Windows 2012 R2 with Hyper-V and a back-end, fiber-channel SAN comprised of Dell Compellent Storage. The operating system of the VMs simulating the Microsoft Exchange workload and of the host hypervisor machines running the infrastructure and monitoring is Windows Server 2012 R2.

2.1 Conceptual system design

The elements of the infrastructure supporting the simulated environment, their main relationships and connectivity links are represented in the conceptual diagram in Figure 1.

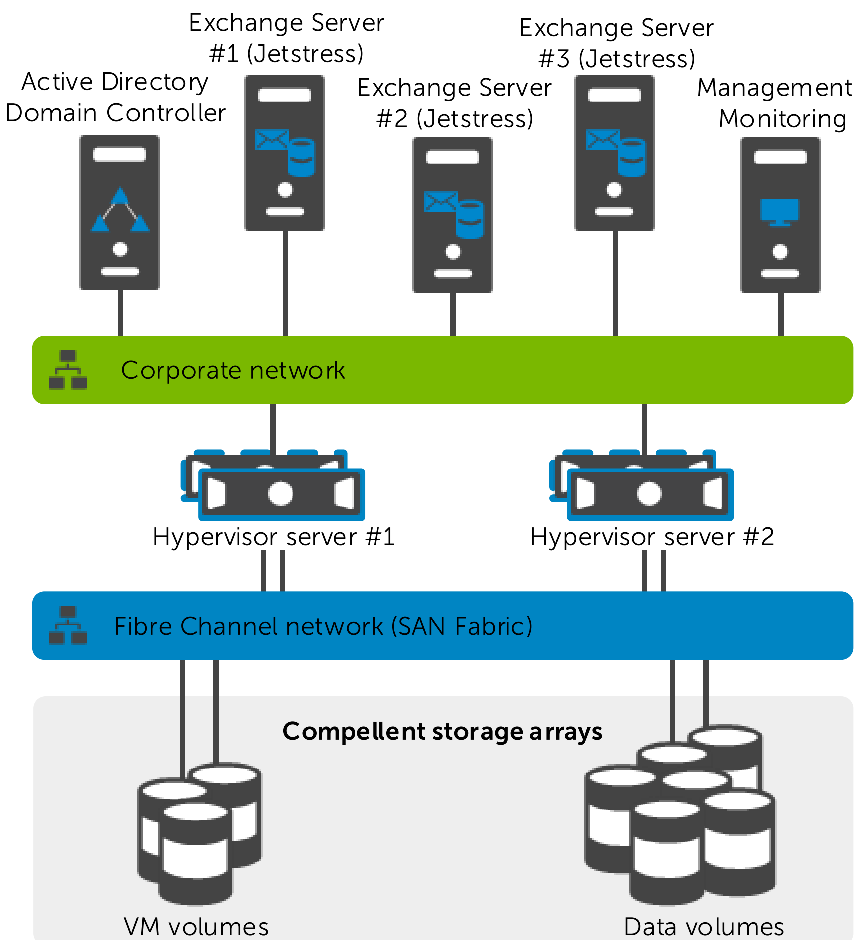

Figure 1 Conceptual system design for the solution components

The key elements of this design are:

- Single Active Directory forest, single domain, single site
- Centralized management and monitoring with dedicated resources (both physical and virtual)
- Building block design approach for mailbox server with Jetstress 2013

2.2 Physical system configuration tested

The physical components and the connections beneath the virtual infrastructure are shown in Figure 2.

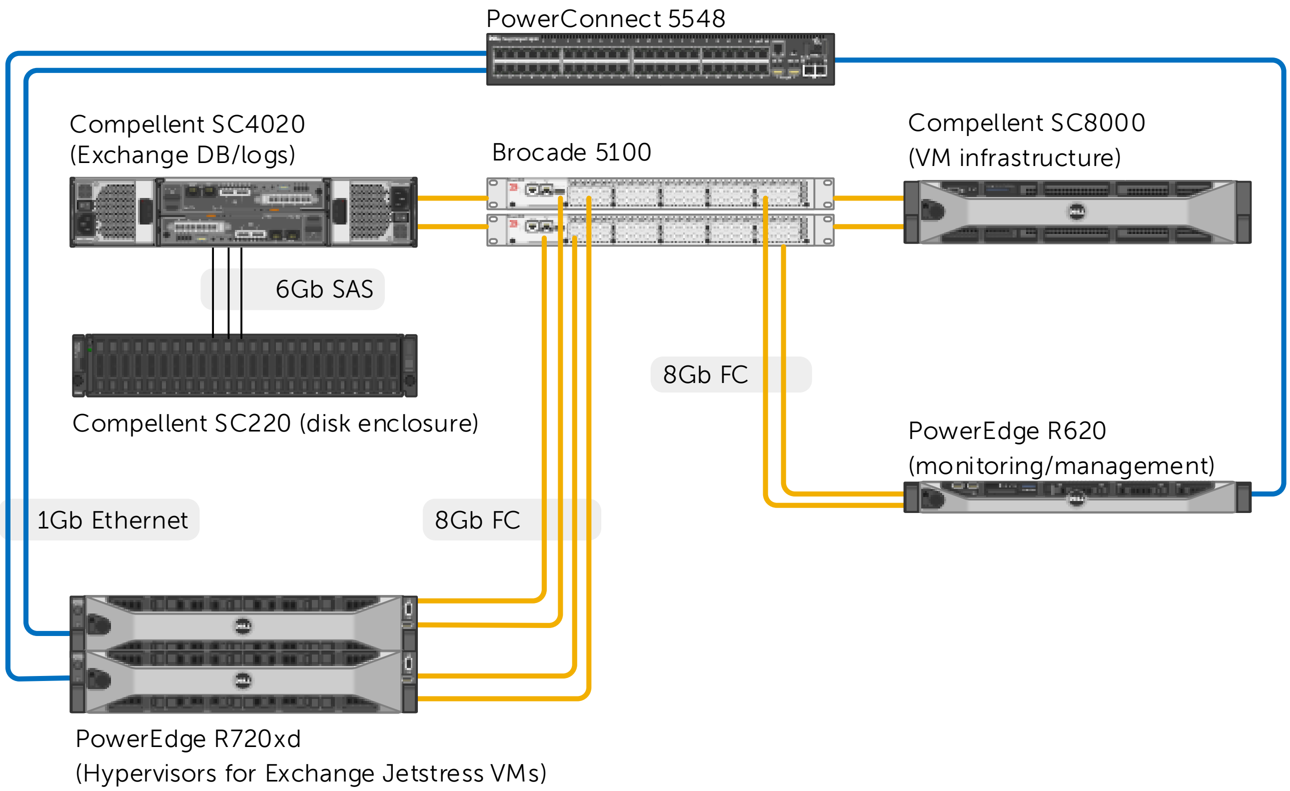

Figure 2 Physical system design for the components of the solution

The solution architecture was deployed on Dell rack mount servers with top–of-rack (ToR) Ethernet network switches dedicated to IP traffic (traditional client/server, management, and hypervisor communications). The hardware elements included in the architecture were:

- Two PowerEdge R720xd rack mount servers that powered the hypervisors beneath the simulated VMs
- One PowerEdge R620 rack mount server that powered the the Dell Storage Manager centralized management and monitoring infrastructure
- Dell SC4020 SAN provisioned with fiber-channel front-end and SAS back-end with one SC220 drive enclosure connected.
- Dell SC8000 SAN provisioned with fiber-channel front-end and SAS back-end for VM infrastructure storage. This SAN can scale to larger infrastructures and was used in this lab test to offload VM infrastructure from the Exchange workload. This could easily be another SC4020 in smaller environments.
- One PowerConnect 5548 Ethernet switch to support LAN IP traffic configured in stack
- Two Brocade 5100 Fiber-channel switches to support the SAN traffic

> **Note**: The configuration details used in the solution infrastructure, including a hardware and software list, SAN array characteristics, hypervisor and VM relationships, and physical and virtual network connections, are listed in Appendix A.

2.3 Storage and volume layout

The configuration of Compellent SAN arrays and the volumes underlying the Exchange databases included:

- One Storage Center disk folder, configured with 10 active 900GB 2.5 inch 10K drives for VM boot volumes on a separate SC8000 array designated for virtual machine infrastructure
- One SC4020, configured with one SC220 drive enclosure with one disk folder of 46 active 7K 1TB drives for Exchange databases and logs
- RAID 6-10 storage profile applied as a reference configuration to Exchange database/log volumes
- One volume allocated in the SC8000 disk folder for each host to store the file images, configuration and temporary files for the corresponding hosted VMs
- A set of data volumes allocated within the SC4020 disk folder, dedicated to the Exchange databases and respectively mapped to the R720 hosts containing the virtual disks in Exchange VMs hosting the simulation
- Exchange mailbox databases and their private set of log files are deployed in the same volumes

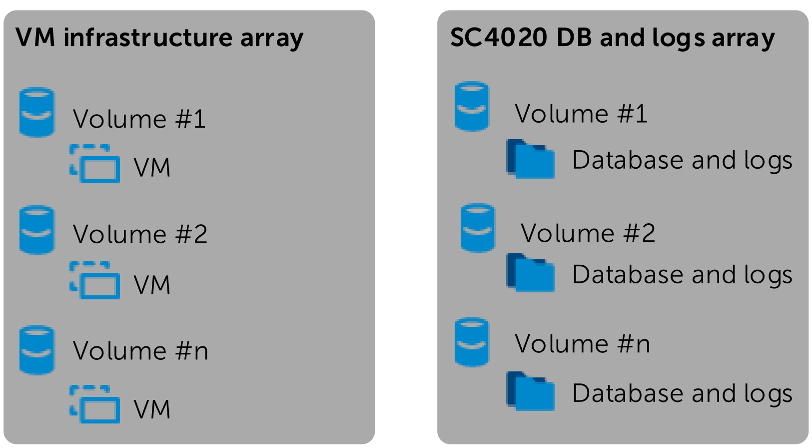

Figure 3 Volumes and database/logs layout

3 Virtualization of Microsoft Exchange storage workload

Exchange Server is a product built on a wide set of components and services that cooperate to support the various requirements needed to design and deploy a messaging infrastructure with advanced capabilities in every organization. The latest Exchange Server version offers a simplified topology that reduces the number of server roles required while consolidating and optimizing the services in two layers. The front-end services deal with the clients (Client Access Server role), and the back-end services are in charge of data management and message transport (Mailbox Server role).

Because more capabilities can be now implemented in fewer layers, the first outcome is an increase in processing demand for the servers deployed to support Exchange Server 2013. This is a challenge that should be planned for carefully from both a capacity and performance standpoint and tailored to each organization in order to be successful with virtualized infrastructures.

The appropriate sizing of the storage subsystem is a key factor in the mailbox role performance and can contribute heavily to easing bottlenecks or administrative overhead for that layer. To set the stage for the analysis provided in the later sections of this paper, the components underneath the Managed Store of Microsoft Exchange Server 2013 are examined first in section 3.1.

3.1 Exchange store elements

Accessing mailbox databases is the primary activity that generates I/O on an Exchange Server storage subsystem. A database is a logical representation of a collection of user or system mailboxes, and it is also an aggregation of files on the disk that are accessed and manipulated by a set of Exchange services following distinct rules (for example, the Information Store, Search feature or Replication Service).

Database file (*.edb): This is the container for user mailbox data. Its content, broken into database pages of 32 KB, is primarily read and written to randomly as required by the Exchange services running on the mailbox server role. A database has a 1:1 ratio with its own *.edb database file. The maximum supported database size in Exchange Server 2013 is still 16 TB, where the Microsoft guidelines recommend a maximum 200 GB database file in a standalone configuration and 2 TB if the database participates in a replicated DAG environment.

Transaction logs (*.log): This is the container where all the transactions that occur on the database (such as create, modify and delete messages) are recorded. Each database owns a set of logs and keeps a one-to-many ratio with them. The logs are written to the disk sequentially, appending the content to the file. The logs are read only when they are in a replicated database configuration within a DAG or in the event of a recovery. The log truncation method is the process to remove old log files preventing them from accumulating. The truncation happens when a consistent backup is performed or when the log configuration is set to circular logging.

Checkpoint file (*.chk): This is a container for metadata tracking when the last flush of data from the memory cache to the database occurred. Its size is limited to 8 KB and, although repeatedly accessed, its overall amount of I/O is so minimal that it can be ignored. The database keeps a 1:1 ratio with its own checkpoint file and positions it in the same folder location as the log files.

Search Catalog: This is a collection of flat files (content index files) built by the Microsoft Search Service, having several different file extensions that all reside in the same folder. The client applications connected to Exchange Server benefit from this catalog as it enables them to perform faster searches based on indexes instead of full scans.

Exchange Server uses a proprietary format called Extensible Storage Engine (ESE) to access, manipulate, and save data to its own mailbox databases. This same format is employed on the Exchange HUB server role for the queue databases. ESE technology, previously known as Jet Database Engine (JDE), has evolved through several versions of Exchange Server and has been a part of several Microsoft products since its inception (for example, Microsoft Access, Active Directory, File Replication Service, WINS server, and Certificate Services).

The ESE is an Indexed Sequential Access Method (ISAM) technology that organizes database data in B-Tree structures. Ideally, these databases are populated by data kept together or adjacent. When this does not occur, external reorganization or defragmentation processes should be used to restore the optimal data contiguity in these structured databases.

To summarize, an Exchange mailbox database is subject to a subset of tasks performing storage access.

- The regular read and write access required to retrieve and store user mailbox data (according to the Exchange cache policy)
- The online defragmentation and compacting activities due to the B-Tree optimization
- The background database maintenance including recoverable items cleanup, deleted mailboxes purge, and other activities addressing logical object support
- The checksum database scan to verify data block integrity (sequential read activity), which can be set to run constantly in the background or at a scheduled time

Furthermore, Exchange Server offers a specialized offline manual defragmentation task that runs while the database is dismounted by taking advantage of the ESEUTIL.EXE command line tool. The principal goal of this task is to reclaim the empty space left in a database by online defragmentation and to shrink the size of the *.edb file itself. This returns the free space to the operating system volume.

Note: It is not recommended to include offline defragmentation in a regular maintenance plan due to the disruption in the availability of the database, the rupture of the logs chain, and the need for database reseeding in case of DAG configuration.

Exchange DAG is a pool of up to 16 networked servers that hosts multiple copies of the same Exchange database or databases where only one of the copies is active at a specific point-in-time within the group. The other copies are passive and contain data sourced from replicated and replayed transaction logs.

Log Checkpoint depth refers to the amount of logs written to the disk and that contain transactions not yet flushed to the database file. In Exchange Server 2013, during a DAG failover the database cache is no longer flushed, since it is treated as a persistent object. Therefore, the log checkpoint for the passive databases is increased to 100 to reduce the write I/O and to reduce the failover time since the passive database has to pre-read less data.

4 Validating the Exchange workload

Microsoft Exchange is a flexible deployment platform offering many different configurations and set of thresholds depending on the needs of the organization and the end-users. The selection of some variables influences the performance, administration and cost effectiveness of the messaging infrastructure.

The focus of this paper is the behavior of larger mailbox sizes and achievable IOPS with acceptable latencies with a fixed number of 7K RPM drives. Storage page sizes from 2MB to 4MB are explored. The drive physics with slower-RPM higher-capacity drives translates to higher I/O latencies. Higher drive failure rates with larger-capacity slower-RPM drives necessitate dual redundancy by leveraging RAID 6, which can handle up to two simultaneous disk failures prior to needing to fail over to additional DAG copies. This paper demonstrates that this cost-effective SAN storage solution provides the needed performance and capacity required for an Exchange Server 2013 workload.

Note: the results presented in this white paper are from simulations executed in a lab built to the listed specifications. Other generally available tools or sizing calculators might have results slightly different based on the assumptions made from each particular tool.

The reference Exchange server deployment used is detailed in Table 1. For each scenario evaluated, an explicit description of the relevant differences is reported in a corresponding section of this paper.

Table 1 Reference configuration for Microsoft Jetstress 2013 tests

Reference configuration: factors under study	
Number of simulated users / mailboxes	4,500 concurrent users
Number of databases	16 databases (active)
RAID 6 striping	RAID 6-10 wide striped across entire disk folder/pool
Array model, SAN configuration	One SC4020, one SC220 2.5" enclosure, 46 active 7K 1TB drives, 2 spares
Reference configuration: consistent factors across each scenario	
Messages per day per mailbox / IOPS per mailbox	150 messages / 0.102 IOPS (with DAG)
Mailbox size	6 GB each
Mailbox allocation per database	281 mailboxes per each mailbox database
Database size	2.2 TB each max (1.67 TB @ 75% for Jetstress)
Number of database replica copies	2 (two node DAG)
Background database maintenance	Enabled
Windows Disk/Partition File System	Basic disk, GPT partition, default alignment NTFS, 64KB allocation unit size
RAID policy	RAID 6
Test duration	2 hours + time required to complete DBs checksum

Note: For a comprehensive list of storage configuration options and supported scenarios for Exchange Server, refer to the article *Exchange 2013 Storage Configuration Options* available at: http://technet.microsoft.com/en-us/library/ee832792%28v=exchg.150%29.aspx

The following list provides the metrics and pass/fail criteria recorded while completing the tests. Most of this information is outlined by the Jetstress tool report or is verified through the recording of Windows Performance Monitor and Dell Storage Manager counters. Microsoft thresholds for Exchange Server storage validation are reported as well.

Database Reads Average Latency (ms) is the average length of time to wait for a database read operation (random reads). It should be less than 20 ms according to Microsoft threshold criteria.

Database Writes Average Latency (ms) is the average length of time to wait for a database write operation (random writes). It should be less than 20 ms according to Microsoft threshold criteria.

Logs Writes Average Latency (ms) is the average length of wait time for a log file write operation (sequential writes). It should be less than 10 ms according to Microsoft threshold criteria.

Planned Transactional IOPS are the target amount of IOPS for the test (calculated by multiplying the number of users by the IOPS per mailbox).

Achieved Transactional IOPS are the amount of IOPS actually performed by the storage subsystem to address the transactional requests. The result should not diverge more than 5% from the planned IOPS to be considered a successful test iteration according to Microsoft Jetstress.

LOGs IOPS are the IOPS performed against the log files during the transactional test. They are not directly taken into account as part of the transactional IOPS, but are tracked separately.

Additional IOPS are the IOPS generated for the database (DB) maintenance, log files replication and all the remaining activities on the storage subsystem, calculated as the difference between the IOPS provisioned by the SAN and the previously reported transactional and logs IOPS.

Total IOPS of the SAN is the sum of achieved transactional IOPS, Logs IOPS and additional IOPS. It represents the entire IOPS footprint performed against the back-end SAN during a test. It is recorded at the SAN level and verified with the Exchange host.

Note: For details about the simulation tool, Microsoft Jetstress 2013, refer to Appendix B.

5 Performance characteristics

As the need for high capacity mailbox storage for Exchange increases, there is no need to sacrifice storage functionality and flexibility that is needed in a dynamic virtual server infrastructure. The SC4020 SAN solution below offers all the functionality of higher-end SAN solutions with a lower price point and acceptable performance. The RAID protection capabilities along with virtualized storage performance make it attractive as compared to a DAS/JBOD configuration for Exchange.

5.1 Performance at full workload

The full workload described and tested for this paper was a medium-sized organization with 4500 mailboxes that required a 6GB mailbox quota. This workload was well satisfied by the solution configuration tested. Two Hyper-V virtual machines were utilized with their database/log storage mapped to volumes provisioned on the SC4020 array. In an Exchange DAG environment, an additional storage array and hypervisor server/s would be required to maintain high availability according to best practices.

Table 2 Test parameters: 4500 concurrent Mailboxes at full workload

Reference configuration: factors under study	
Performance characteristics required of RAID-6 7K architecture at maximum workload	Read IOPS achieved =< 450
	Read Latency < 20ms
Reference configuration: consistent factors within this scenario	
Messages per day per mailbox / IOPS per mailbox	150 messages/0.102 IOPS (with DAG)
Number of simulated users / mailboxes	4,500 concurrent users
Mailbox size	6 GB (6144 MB) each
Number of databases	16 databases (active)
Mailbox allocation per database	281 mailboxes per each mailbox database
Database size	2.2 TB each (max)
Number of database replica copies	2 (two node DAG)
Array model, SAN configuration	One SC4020, one SC220 2.5" SAS enclosure, 46 active 7K 1TB drives, 2 spares
RAID policy	RAID 6

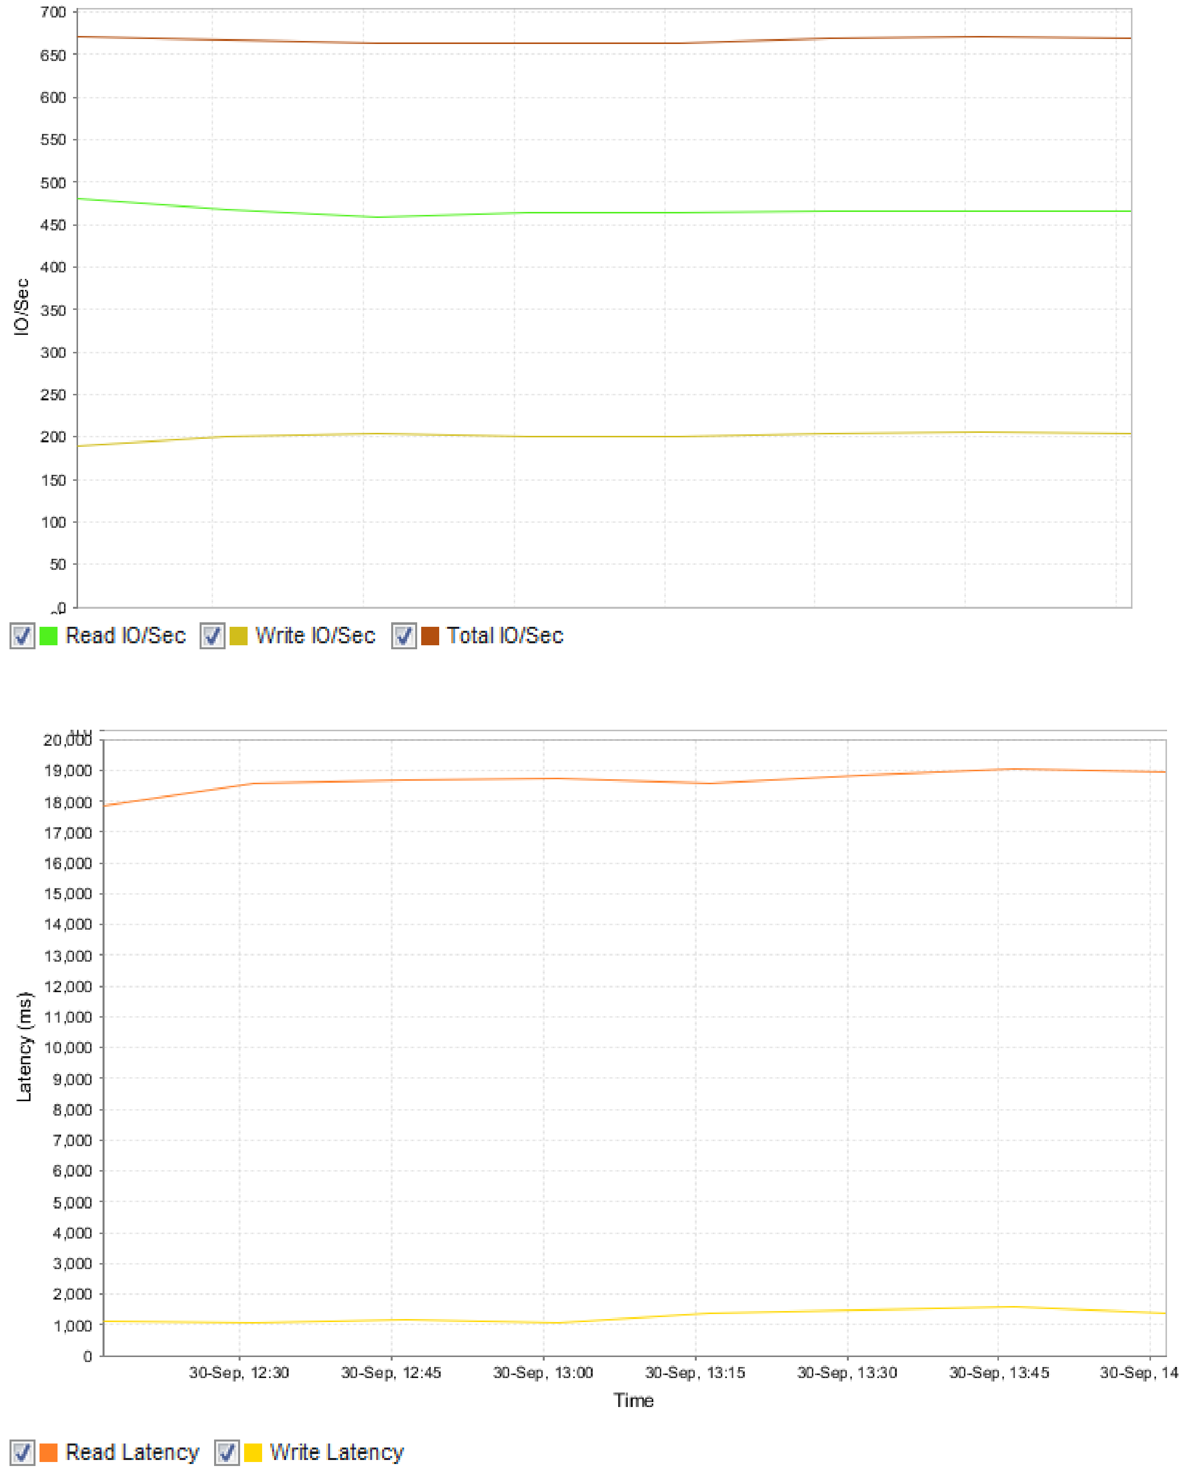

Figure 4 Full workload IOPS/Latency historical graphs from Dell Storage Manager

5.1.1 Storage Center Manager results

Table 3 Test results from Storage Manager - Full workload I/O usage

IO Usage for SC 31
Volume Reporting I/O Usage

	Total IOPS				**Read IOPS**				**Write IOPS**		
Volume Name	Min	Max	**Avg**	StDev	Min	Max	**Avg**	StDev	Min	Max	**Avg**
HV-JS1-DB1	42	42	42	0	29	30	29.12	0.35	12	13	12.88
HV-JS1-DB2	42	43	42.12	0.35	29	31	29.25	0.71	12	13	12.88
HV-JS1-DB3	41	43	42	0.53	29	30	29.38	0.52	12	13	12.62
HV-JS1-DB4	42	43	42.38	0.52	29	30	29.5	0.53	12	13	12.88
HV-JS1-DB5	40	42	41.62	0.74	28	30	29	0.53	12	13	12.62
HV-JS1-DB6	40	43	41.62	0.92	28	31	29	0.93	12	13	12.62
HV-JS1-DB7	42	43	42.12	0.35	29	30	29.25	0.46	12	13	12.88
HV-JS1-DB8	41	42	41.62	0.52	29	30	29.12	0.35	12	13	12.5
HV-JS3-DB1	39	42	41.5	1.07	27	30	28.88	0.83	12	13	12.62
HV-JS3-DB2	39	42	41.12	0.99	27	30	28.88	0.83	11	13	12.25
HV-JS3-DB3	39	43	41.62	1.19	27	30	29.12	0.99	12	13	12.5
HV-JS3-DB4	39	43	41.88	1.25	27	31	29.25	1.16	12	13	12.62
HV-JS3-DB5	40	43	42.12	1.13	27	31	29.38	1.19	12	13	12.75
HV-JS3-DB6	40	42	41.38	0.74	28	30	29	0.53	12	13	12.38
HV-JS3-DB7	40	43	41.62	0.92	27	30	29	0.93	12	13	12.62
HV-JS3-DB8	40	42	41.38	0.74	27	30	28.88	0.83	11	13	12.5
Total			**668.1**				**466.01**				**202.12**

Table 4 Test results from Storage Manager - Full workload I/O latencies

I/O Usage for SC 31
Volume Reporting I/O Latencies

	Read Latency				**Write Latency**			
Volume Name	Min	Max	**Avg**	StDev	Min	Max	**Avg**	StDev
HV-JS1-DB1	18	19	18.79	0.34	0	1	0.76	0.15
HV-JS1-DB2	15	16	15.95	0.21	0	1	1.03	0.26
HV-JS1-DB3	15	16	16.05	0.24	0	1	0.96	0.33
HV-JS1-DB4	15	16	16.08	0.33	0	1	0.89	0.27
HV-JS1-DB5	15	16	16.29	0.51	0	0	0.75	0.1
HV-JS1-DB6	15	16	15.95	0.23	0	1	1	0.31
HV-JS1-DB7	15	16	16.09	0.32	0	0	0.72	0.06
HV-JS1-DB8	15	16	15.92	0.25	0	1	1.03	0.32
HV-JS3-DB1	17	18	18.23	0.51	0	1	0.93	0.25
HV-JS3-DB2	15	16	15.97	0.34	0	1	0.85	0.24
HV-JS3-DB3	15	16	15.91	0.39	0	1	0.95	0.23
HV-JS3-DB4	15	16	15.99	0.4	0	1	0.86	0.23
HV-JS3-DB5	15	16	15.94	0.35	0	1	0.88	0.21
HV-JS3-DB6	15	16	15.88	0.19	0	1	0.86	0.21
HV-JS3-DB7	15	16	15.99	0.39	0	1	0.83	0.2
HV-JS3-DB8	15	16	15.94	0.43	0	1	0.84	0.19
Total			**16.3106**				**0.88375**	

5.1.2 Analysis of workload key performance indicators (KPI)

5.1.2.1 Summary

- HV-JS1
  - IO Size average = 82K
  - DB Reads/sec = 229.82
  - Total IOPS = 335

- HV-JS3
  - IO Size average = 82K
  - DB Reads/sec = 230.94
  - Total IOPS = 336

- TOTAL
  - 671 IOPS F/E = 14.58 IOPS F/E per disk
  - 2,300 IOPS B/E = 50 IOPS B/E per disk

5.1.2.2 Analysis

- On average, this configuration performed up to the necessary IOPS for the workload of 450 read IOPS.
- Read Latency on average was 16.31ms, which is comfortably under the 20 millisecond threshold that Exchange requires for excellent user experience.
- The higher latency instances on each server were discovered to have a smaller percentage of pages on the fast tracks of the drive. Therefore the drive physics of seek time at slower RPMs caused some slightly higher latency, but still under the 20ms seek time.
- Using the same type of drives in the external SC200 enclosure as the internal enclosure (1TB 7K in this case), allowed adding the disks to the same disk folder and show consistent results of performance and latency across all database volumes. This same drive preparation is best practice for a disk folder. If different size/class drives are used, these can be used to form a separate disk folder or tier of storage.
- Tests with 4MB page vs. the default 2MB page size did not appear to change the key performance characteristics for these tests. It was postulated that since the average I/O sizes were quite large, that increasing the page size to 4MB would decrease read latency due to fewer I/Os needed to read the same amount of data. There was no appreciable change seen in these specific tests.

5.2 Performance with single disk failure

The full workload was tested again with the same medium-sized organization and the 4500 mailboxes that require a 6GB mailbox quota. In this test, a simulated drive failure was introduced to demonstrate the continuation of the workload at acceptable performance in RAID degraded mode and RAID rebuild.

In an Exchange DAG environment, the additional storage array and hypervisor servers could be utilized to manually switch over the workload to the additional DAG database copies in the unlikely event of performance becoming unacceptable during rebuild. The tests showed that at full workload capacity, on average the read latency was acceptable.

Table 5 Test parameters: 4500 concurrent Mailboxes at full workload with single drive failure

Reference configuration: factors under study	
Performance characteristics required of RAID-6 7K architecture at maximum workload	Read IOPS achieved ~ 450
	Read Latency < 20ms
Reference configuration: consistent factors within this scenario	
Messages per day per mailbox / IOPS per mailbox	150 messages/0.102 IOPS (with DAG)
Number of simulated users / mailboxes	4,500 concurrent users
Mailbox size	6 GB (6144 MB) each
Number of databases	16 databases (active)
Mailbox allocation per database	281 mailboxes per each mailbox database
Database size	2.2 TB each (max)
Number of database replica copies	2 (two node DAG)
Array model, SAN configuration	One SC4020, one SC220 2.5" SAS enclosure, 45 active 7K 1TB drives, 2 spares – **RAID degraded to single redundancy, then RAID rebuild on spare disk**
RAID policy	RAID 6

Notes: The Storage Manager graphs below show that although there was a perceptible decrease in I/O during the RAID rebuild mode, there was not any I/O interruption or large latency spike during or after the drive failure.
The latency graph shown below is not the average I/O over the volumes, but the highest among all volumes. As shown in later tables, the majority of the volumes were under 20ms read latency even after drive failure.

Notice that as the RAID rebuild started, the IOPS dropped slightly due to the back-end I/O being generated by the rebuild. Read latency bumped up by a maximum of 2ms during RAID rebuild.

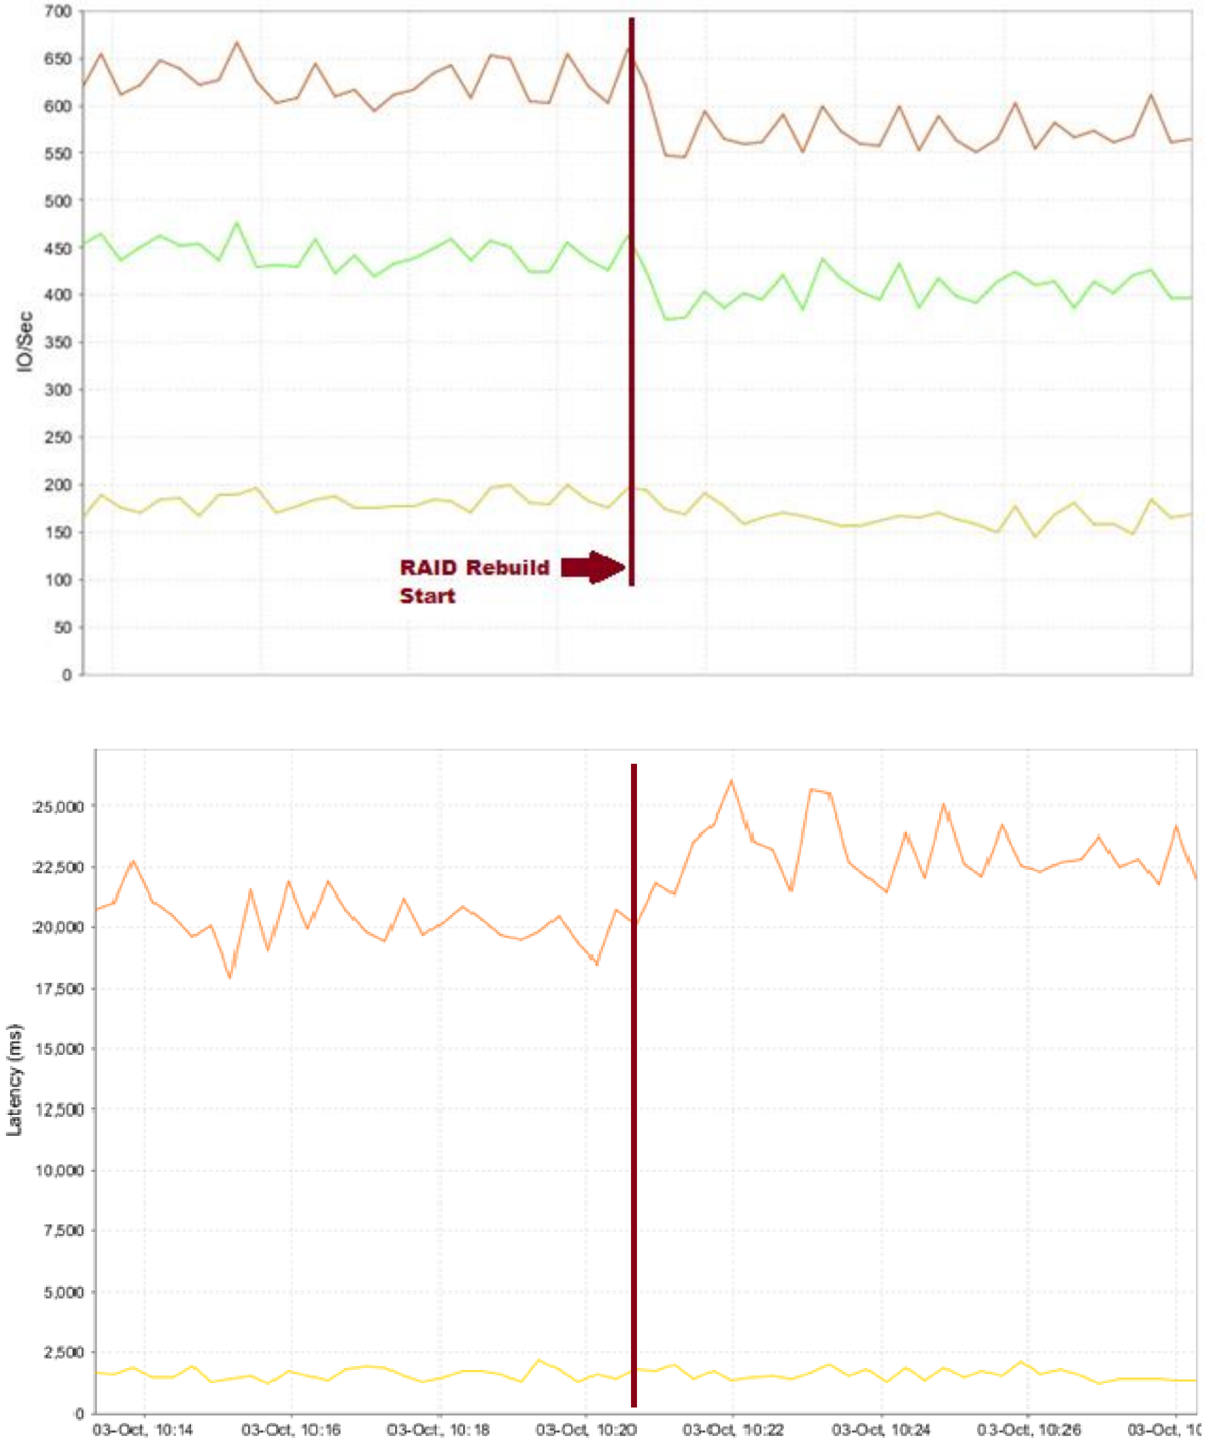

Figure 5 Full workload IOPS/Latency historical graphs during drive failure from Dell Storage Manager

5.2.1 Storage Center Manager results

Table 6 Storage Center Manager I/O usage report during disk failure test

IO Usage for SC 31
Volume Reporting I/O Usage

	Total IOPS				**Read IOPS**				**Write IOPS**		
Volume Name	Min	Max	**Avg**	StDev	Min	Max	**Avg**	StDev	Min	Max	**Avg**
HV-JS1-DB1	35	41	37.14	2.41	25	29	26.43	1.51	10	12	10.71
HV-JS1-DB2	35	41	37.43	1.99	25	29	26.43	1.51	10	12	11
HV-JS1-DB3	35	41	37.29	2.36	25	29	26.29	1.6	10	12	11
HV-JS1-DB4	35	42	37.14	2.73	25	29	26.29	1.6	10	13	10.86
HV-JS1-DB5	34	42	37.43	2.51	25	30	26.57	1.81	9	12	10.86
HV-JS1-DB6	36	41	38	2.08	26	29	26.71	1.25	10	13	11.29
HV-JS1-DB7	35	40	37	1.91	25	28	26.29	1.25	10	12	10.71
HV-JS1-DB8	35	41	37.43	2.23	25	29	26.43	1.51	10	12	11
HV-JS3-DB1	36	41	37.71	1.7	25	29	26.57	1.4	11	12	11.14
HV-JS3-DB2	35	40	37.57	1.62	25	28	26.43	1.13	10	12	11.14
HV-JS3-DB3	36	41	37.29	1.98	25	29	26.57	1.4	10	12	10.71
HV-JS3-DB4	35	41	37.14	1.95	25	29	26.43	1.27	10	12	10.71
HV-JS3-DB5	36	41	37.43	2.15	25	29	26.57	1.4	10	12	10.86
HV-JS3-DB6	35	41	37.29	2.29	25	29	26.57	1.4	10	12	10.71
HV-JS3-DB7	35	41	37.57	2.15	25	29	26.57	1.4	10	12	11
HV-JS3-DB8	35	41	37.29	2.06	25	29	26.57	1.4	10	12	10.71
Total			**598.15**				**423.72**				**174.41**

Table 7 Storage Center Manager I/O latency report during disk failure test

I/O Usage for SC 31
Volume Reporting I/O Latency

	Read Latency				**Write Latency**			
Volume Name	Min	Max	**Avg**	StDev	Min	Max	**Avg**	StDev
HV-JS1-DB1	16	18	17.95	0.98	0	1	1.06	0.09
HV-JS1-DB2	18	22	20.62	1.13	0	1	1	0.08
HV-JS1-DB3	16	18	17.8	0.87	0	1	1.04	0.1
HV-JS1-DB4	16	18	17.7	0.85	0	1	0.98	0.06
HV-JS1-DB5	16	19	17.89	0.98	0	1	1.02	0.06
HV-JS1-DB6	16	18	17.86	0.79	0	1	0.97	0.06
HV-JS1-DB7	16	18	17.9	1.03	0	1	0.96	0.08
HV-JS1-DB8	16	18	17.68	0.85	0	1	0.98	0.06
HV-JS3-DB1	18	22	20.25	1.21	0	1	0.99	0.05
HV-JS3-DB2	16	18	17.77	0.8	0	1	1.03	0.03
HV-JS3-DB3	16	18	17.88	0.83	0	1	0.99	0.05
HV-JS3-DB4	16	18	17.61	0.88	0	1	1.01	0.06
HV-JS3-DB5	16	18	17.96	0.99	0	1	0.99	0.05
HV-JS3-DB6	16	18	17.75	0.85	0	1	1.01	0.03
HV-JS3-DB7	16	19	18.06	1	0	1	0.98	0.03
HV-JS3-DB8	16	19	18.28	0.95	0	1	1	0.04
Total			**18.185**				**1.00063**	

5.2.2 Analysis of workload key performance indicators (KPI)

- HV-JS1
  - IO Size average = 88K
  - DB Reads/sec = 211.44
  - Total IOPS = 298.86

- HV-JS3
  - IO Size average = 88K
  - DB Reads/sec = 212.28
  - Total IOPS = 299.29

- TOTAL
  - 598 IOPS F/E = 13.00 IOPS F/E per disk
  - 2,407 IOPS B/E = 52 IOPS B/E per disk

- On average, this configuration in RAID rebuild mode performed adequately to the necessary IOPS for the workload: 423 of 450 read IOPS.
- The key point is that I/O operations continued without any perceptible interruption by Exchange ESE engine. All IO's were satisfied.
- Read Latency on average was 18.18ms, which is under the 20ms threshold that Exchange requires for an optimal user experience.
- The higher latency instances on each server (1 database each) were previously discovered to have a smaller percentage of pages on the fast tracks of the drives, again the drive physics of seek time at slower RPMs and higher back-end IOPS due to RAID rebuild caused some higher latency, slightly higher than 20ms seek time just on those instances.

6 Best Practice recommendations

Refer to these best practices to plan and configure Dell Storage arrays, Exchange Server 2013 and VMs.

6.1 Storage best practices

- Use Multipath I/O (MPIO) on Dell Compellent to improve storage performance and reliability. This support is built into Windows Server 2008 and newer.
- Choose the most appropriate RAID policy when designing the environment according to the performance, capacity, and tolerance to failure requirements of the environment. In this reference architecture, this was RAID 6 due to the larger mailbox capacity needs and single disk tier used.

6.2 Hypervisor and VM best practices

- Install a Windows Server Core version in the root partition of the Hyper-V role server to reduce the software maintenance, the attack surface, the memory, and disk space footprint. Otherwise, when installing a traditional Windows Server with Hyper-V technology with the GUI, minimize the use of additional software, components and/or roles in the root partition.
- The use of Non-Uniform Memory Access (NUMA) is advised to address the management of VMs with large and very large memory settings. Verify the number of NUMA nodes available in the system, based on the number of processors, and then design and size the VMs to have their memory resources entirely contained in a single NUMA node when possible. Spanning a VM memory across multiple NUMA nodes can result in less efficient usage of the memory and can decrease performance. See Appendix A for additional information on NUMA.
- Exchange Server 2013 is a memory-intensive workload. Configure static memory in the settings of each VM to avoid possible resource contention created by the dynamic memory management and to comply with the current support directives from Microsoft on this matter.
- Utilize Dell Replay Manager for Microsoft Hyper-V and Exchange for VSS consistent Replay snapshots.
  - Use the Hyper-V extension to protect all of the virtual machines along with their Exchange databases with a consistent VSS-aware Replay.
  - Use the Exchange extension when the operational needs require a more granular recovery of Exchange databases. Use this for complete database restores or Replay Manger Expose in conjunction with Exchange recovery databases or other offline database tool such as Dell Software Recovery Manager for Exchange.
  - See Replay Manager documentation and demo videos on these extensions at
    Dell Compellent Replay Manager 7.0 and Exchange 2013 Demo
    Dell Compellent Storage Center Replay Manager 7 and Microsoft Hyper-V Demo Video

- Investigate the opportunity to reserve the resources assigned to the Exchange VMs by using the resource control settings.
- Plan the reserve size for the volumes hosting the hard disk files of the VMs (*.vhdx, *.vhd) that takes into account the extra space required for memory image files (*.bin), saved states (*.vsv), or snapshot files (*.avhdx, *.avhd).

- Consider the supportability constraints of Exchange Server 2013 when deployed in a virtualized environment around suspended state, VM snapshots and replicas, and differencing disk image files.

  See Microsoft Technet - Exchange 2013 virtualization at: http://technet.microsoft.com/en-us/library/jj619301(v=exchg.150).aspx

- While performance of dynamically expanding disks has improved, fixed disks are preferred to deploy production environment VMs due to the risk of elevated fragmentation or high latency while disk expansion occurs. See the Technet article link above on Exchange 2013 virtualization for more information on this recommendation.
- Isolate the host management traffic from the VM traffic by using virtual switches not enabled for management.
- Design isolated LAN and iSCSI networks.

6.3 Exchange Server installation best practices

- Use Basic disk type for all volumes.
- Use GUID partition table (GPT) for Exchange volumes.
- Use default disk alignment provided by Windows 2008 or greater.
- Use NTFS file system with 64 KB allocation unit for Exchange database and log partitions.
- Evaluate the use of mount points for all the SAN volumes or the attached virtual disks to increase management flexibility and database portability. Mount points become necessary when the number of volumes exceeds the number of available drive letters on a server.
- When using mount points, prevent Windows Server from assigning drive letters automatically to new volumes by disabling the auto-mount option.
- Deploy Windows operating system and Exchange data on separate physical disk drives, or in separate disk folders in a Compellent storage environment.
- Database and log file isolation is not required when deployed in a DAG environment.
- Leave background database maintenance (BDM) enabled (24x7) and account for the additional load. The BDM is activated by default on every replica copy of a DAG configuration.
- Do not share the disk drives for active and replicated copies of an Exchange mailbox database in a DAG environment. If there is a failure of a set of drives with multiple copies of the same data, the resilience or the perceived availability of the applications would be affected. Dedicate separate storage arrays for each replicated instance of the data instead.

6.4 Know the workload

- Do not begin a deployment without having a solid understanding of the current messaging workload. In the case of a new deployment, collect estimates based on business cases matching your organization size and drive conservative figures for the average user profiles. The use of the Microsoft Exchange 2013 Server role calculator is advised to learn basic estimates for new or changing messaging environments.

- Forecast the workload gap between the current version of Exchange Server and Exchange Server 2013 when planning to design a storage solution jointly with a migration.

6.5 Distribution of databases and user mailbox count per database

- Carefully plan the number of databases and the mailbox count per database in the entire organization. The right balance of the number of mailbox databases to support a defined number of users is mostly based on administrative policies.
- Larger databases fitted with a high number of users have a smaller impact on the storage subsystem, but provide less flexibility because of large database files maintenance and replication.

6.6 Growth and overhead

- Understand the wider business requirements and company strategy to predict the organization growth instead of following it.
- Design an environment based on building blocks that are easily replicable to address sudden changes in business requirements.
- Always account for at least 20% overhead when calculating maximum sizes to address unforeseen growth.

A Configuration details

A.1 Hardware components

Table 8 lists the details of the hardware components used for the configuration.

Table 8 Hardware components

Component	Description
Servers	Dell PowerEdge R620 server, Firmware 2.2.2 • 2x Eight Core Intel Xeon E5-2665 Processors, 2.4 Ghz, 20M Cache • RAM 32 GB (4x 8GB) • iDRAC7 Enterprise, Firmware 1.56.55 • PERC H710 Mini RAID controller, Firmware 21.2.0-0007 • 4x 146 GB 15K SAS (2x RAID-1, stripe 1MB) • 4x Broadcom NetXtreme 5720 Quad Port 1GbE Base-T onboard, Firmware 7.8.53 • 2x Broadcom NetXtreme II 57810 Dual Port 10GbE Base-T, Firmware 7.8.53 • 2x QLogic QLE2532 2-port 8Gb Fibre Channel adapter 2x Dell PowerEdge R720xd servers, Firmware 1.57.57 (Build 04) • 2x Eight Core Intel(R) Xeon(R) CPU E5-2660 0 @ 2.20GHz • RAM 256 GB (16 x 16GB) • iDRAC-7 Enterprise, Firmware 1.57.57 (Build 04) • Intel(R) 1Gb Ethernet 4-Port I350-t rNDC, Firmware 13.1.10 • Intel(R) 10Gb Ethernet 2-Port X520 Adapter, Firmware 13.1.10 • 2x QLogic QLE2532 2-port 8Gb Fibre Channel adapter
Network	Dell PowerConnect 5548 Ethernet switch, Firmware 5.1.1.7 • 48x 1GbE interfaces & 2x 10GbE interfaces • Installed top of the rack • Connected by 2x16GbE redundant uplinks (STACK)

Component	Description
Storage	2x Brocade 5100 Fiber-channel switches (SAN Fabric) • 40x 8Gb FC interfaces • Fabric OS version 7.0.0b 1 to 2x Dell SC4020 (2 nodes per array) • Storage Array OS version 6.05.10 • 1x Qlogic QLE2564L 4-port 8Gb Fibre Channel adapter • 1x (up to 7) SC220 2.5" SAS drive enclosure with 24 1TB disks • Total 48x 1TB 7.2K 2.5" SAS disk drives, raw capacity 48 TB 1x Dell SC8000 (for virtual infrastructure) • Storage Array OS version 6.5.02 • 1x QLogic QLE2562 4-port 8Gb Fibre Channel adapter • 10x 900GB 10K 2.5" SAS disk drives, raw capacity 9 TB

A.2 Software components

The environment required to perform the simulations described in this paper included the following software components:

- Hypervisor: Windows Server 2012 R2 with Hyper-V on every physical host
- Dell OpenManage Server Administrator on every physical host
- Operating System: Windows Server 2012 R2 on every VM
- Operating System: Windows Server 2008 R2 on management VM
- Dell Compellent Multi-pathing Extension Module to provide Dell MPIO access to the back-end SAN on the hypervisor directly accessing the SAN (host initiator scenarios)
- Dell Storage Manager (formerly Enterprise Manager) to monitor the health and performance of the SAN
- Microsoft Exchange Jetstress to simulate the access to the storage subsystem from the mailboxes store simulated VM

The following software components were installed and configured to simplify the management of the environment and to support the failover cluster configuration, while they were not strictly required to accomplish the tests.

- Active Directory Domain Services and DNS Server roles for the domain controllers
- Microsoft SCVMM for the management VM (not strictly required to accomplish the tests)

Table 9 lists the details of the software components used for the configuration.

Table 9 Software components

Component	**Description**
Operating Systems	Host servers: • Microsoft Windows Server 2012 R2 Datacenter Edition (build 9600) with Hyper-V • MPIO enabled using built-in Compellent DSM for Windows • Guest VMs: o Microsoft Windows Server 2012 R2 Datacenter Edition (build 9600)
Applications	Microsoft System Center 2012 Virtual Machine Manager Service Pack 1 (version 3.1.6011.0)
Component	Description
Monitoring tools	Dell Storage Manager (Enterprise Manager) version 15.1.1.21 Microsoft Performance Monitor from the Windows operating system
Simulation tools	Microsoft Exchange Jetstress 2013 (build 15.00.0775.000) • Exchange 2013 Server Database Storage Engine and Library SP1 CU6 (build 15.00.0995.021)

A.3 Host hypervisor and VMs configuration

A virtual infrastructure built on Windows Server with Hyper-V hosted all the components of the test infrastructure. The primary elements of the virtual infrastructure configuration were:

- Windows Server 2012 R2 with Hyper-V deployed on all hosts, managed by the Hyper-V Role Administration tools or centrally by the SCVMM server
- Two identical hypervisor hosts (R720xd) configured as member servers on the domain
- One Enterprise Manager Data Collector Server (R620) configured as member server on the domain
- All guests deployed from one image template of Windows Server 2012 R2 operating system

Table 10 Configuration: guest to host placement - Lists the relation between each hypervisor host and its respective set of VMs, with a brief summary of the virtual resources allocated for each VM.

Host	VM	Purpose	vCPU	Memory	Storage	Network adapters
R620	DC01	Active Directory Domain Controller	2	4GB	250GB VHDX	2x VMBus Net Adapters (NIC teaming with LBFO)
	SCVMM01	System Center Virtual Machine Manager	4	8GB	250GB VHDX	2x VMBus Net Adapters (NIC teaming with LBFO)
R720s	HV-JS#n	Exchange Server mailbox server simulation (Jetstress)	4	98GB	250GB VHDX + Exchange data volumes	2x VMBus Net Adapters (NIC teaming with LBFO)

A.3.1 Guest VMs memory

The memory assigned to every VM in the infrastructure is configured as static, to avoid any possible occurrence of VMs exhausting the available host server memory pool.

A.3.2 Hyper-V configuration of NUMA

NUMA capabilities were enabled on the PowerEdge R720 physical hosts (with Node Interleaving disabled in the server BIOS) to allow memory access across CPUs. These R720s had two NUMA nodes each managing 192GB of memory. The Hyper-V NUMA Spanning setting was left enabled (the default).

A.3.3 Guest VMs disks

The virtual disks used to host the operating system of each VM were VHDX fixed type disks. The VHDX files for the VMs hosted from the R620 hypervisor were deployed on the second pair of local RAID-1 disks, all the VMs hosted by the R720XD hypervisors were deployed on the SAN as described below:

- VHDXs for the system disks deployed on the SC8000 array
- VHDXs for the Exchange data deployed on the SC4020 arrays (host initiator scenarios)

A.3.4 Host network adapters and virtual network configuration

The host and VM network adapters were configured as follows:

- Two physical network adapters, sourced from the onboard Intel(R) 1Gb Ethernet ports, provided connectivity for host domain access and management
- Two physical network adapters, sourced from the Intel(R) 10Gb Ethernet ports, connected independently with two virtual network switches, and provided connectivity for the VMs to both domain and intra-VM traffic
- Two (on the R620 or the R720s) physical fibre channel host adapters, sourced from the QLogic QLE2532 ports, connected independently with two Brocade switches, and provided access to the SAN from the host hypervisor environments for these use-case configurations.
- MPIO (Multi-path I/O) enabled and provided by the default Compellent DSM module.

The non-default settings implemented on all physical Ethernet network adapters included:

- Jumbo frames enabled
- Flow Control enabled
- Large send offload enabled
- Receive and transmit buffers maximized
- Virtual Machine Queues (for each network adapter used by a virtual network switch)

A.3.5 Virtual network adapter configuration

The assignment of the virtual network adapters of the VMs was configured as listed below:

- Two virtual network adapters to access the LAN traffic with switch independent mode NIC teaming by Load Balance with Fail Over (LBFO) (for all active adapters)

The hardware acceleration settings implemented on the virtual network adapters were:

- Virtual machine queue enabled on every adapter

B Microsoft Jetstress reports

B.1 Jetstress considerations

Microsoft Exchange Server Jetstress 2013 is a simulation tool that is able to reproduce the database and logs I/O workload of an Exchange mailbox role server. It is usually used to verify and validate the conformity of a storage subsystem solution before the full Exchange software stack is deployed. Some elements worth considering include:

- Does not require and should not be hosted on a server where Exchange Server is running
- Performs only Exchange storage access and not host processes simulations. It does not contribute in assessing or sizing the Exchange memory and processes footprints
- Is an ESE application requiring access to the ESE dynamic link libraries to perform database access. It takes advantage of the same API used by the Exchange Server application software stack and as such it is a reliable simulation application
- Requires, and provides, an initialization step to create and populate the database(s) that will be used for the subsequent test phases. The database(s) should be the same capacity as the one(s) planned for the Exchange Server future deployment
- Its topology layout includes number and size of simulated mailboxes, number and placement of databases and log files, and number of database replica copies (it simulates only active databases)
- While carrying out a mailbox profile test, it executes a pre-defined mix of insert, delete, replace and commit operations against the database objects during the transactional step, then it performs a full database checksum
- Collects application and system event logs, performance counter values for the criteria metrics of both operating system resources and ESE instances during transactional and DB checksum phases. It then generates a detailed HTML-based report
- Throttles the disk I/O generation using the assigned IOPS per mailbox, thread count (global per all databases) and SluggishSessions threads property (fine tuning for threads execution pace)

B.2 Jetstress report for full workload – HV-JS1

Figure 6 HV-JS1 Jetstress Performance Test Result Report

Microsoft Exchange Jetstress 2013

Performance Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	HV-JS1
Test Description	2250 Users x 2 Servers = 4500 mailboxes 8 databases .10 IOPS per mailbox 6144 MB (6GB) mailbox quota 8 threads
Test Start Time	9/30/2014 9:55:26 AM
Test End Time	9/30/2014 12:08:13 PM
Collection Start Time	9/30/2014 10:07:11 AM
Collection End Time	9/30/2014 12:07:10 PM
Jetstress Version	15.00.0775.000
ESE Version	15.00.0995.021
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Performance_2014_9_30_9_55_43.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	229.828
Target Transactional I/O per Second	225
Initial Database Size (bytes)	14495783059456
Final Database Size (bytes)	14496412205056
Database Files (Count)	8

Jetstress System Parameters

Thread Count	8
Minimum Database Cache	256.0 MB
Maximum Database Cache	2048.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Run Background Database Maintenance	True
Number of Copies per Database	2

Database Configuration

Instance1980.1 Log path: C:\DB\JS1-DB1
Database: C:\DB\JS1-DB1\Jetstress001001.edb

Instance1980.2 Log path: C:\DB\JS1-DB2
Database: C:\DB\JS1-DB2\Jetstress002001.edb

Instance1980.3 Log path: C:\DB\JS1-DB3
Database: C:\DB\JS1-DB3\Jetstress003001.edb

Instance1980.4 Log path: C:\DB\JS1-DB4
Database: C:\DB\JS1-DB4\Jetstress004001.edb

Instance1980.5 Log path: C:\DB\JS1-DB5
Database: C:\DB\JS1-DB5\Jetstress005001.edb

Instance1980.6 Log path: C:\DB\JS1-DB6
Database: C:\DB\JS1-DB6\Jetstress006001.edb

Instance1980.7 Log path: C:\DB\JS1-DB7
Database: C:\DB\JS1-DB7\Jetstress007001.edb

Instance1980.8 Log path: C:\DB\JS1-DB8
Database: C:\DB\JS1-DB8\Jetstress008001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance1980.1	19.377	1.937	20.210	8.598	33686.343	35474.248	0.000	1.457	0.000	2.291	0.000	20917.939
Instance1980.2	16.950	1.851	20.130	8.503	33791.356	35507.605	0.000	1.427	0.000	2.238	0.000	20865.362
Instance1980.3	16.691	1.761	20.190	8.536	33821.289	35481.637	0.000	1.497	0.000	2.254	0.000	21133.103
Instance1980.4	16.786	1.947	20.191	8.544	33778.859	35545.607	0.000	1.398	0.000	2.288	0.000	20833.049
Instance1980.5	16.796	2.201	20.113	8.522	33672.760	35594.566	0.000	1.477	0.000	2.258	0.000	20931.568
Instance1980.6	17.069	2.330	19.996	8.519	33746.178	35546.725	0.000	1.436	0.000	2.242	0.000	21583.087
Instance1980.7	16.842	2.061	20.366	8.663	33784.438	35300.126	0.000	1.433	0.000	2.284	0.000	20265.440
Instance1980.8	16.749	1.933	20.225	8.521	33745.497	35461.944	0.000	1.483	0.000	2.264	0.000	20912.303

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1980.1	9.109	261975.126
Instance1980.2	9.072	261885.808
Instance1980.3	9.117	261861.422
Instance1980.4	9.110	261889.818
Instance1980.5	9.109	261948.659
Instance1980.6	9.052	261944.200
Instance1980.7	9.111	261909.267
Instance1980.8	9.112	261845.277

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance1980.1	0.201	79275.309
Instance1980.2	0.198	77343.801
Instance1980.3	0.199	77785.915
Instance1980.4	0.199	77830.332
Instance1980.5	0.200	79227.622
Instance1980.6	0.203	79155.941
Instance1980.7	0.196	76857.270
Instance1980.8	0.199	77358.416

Total I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance1980.1	19.377	1.937	29.319	8.598	104615.130	35474.248	3.978	1.457	0.201	2.291	79275.309	20917.939
Instance1980.2	16.950	1.851	29.202	8.503	104652.078	35507.605	3.970	1.427	0.198	2.238	77343.801	20865.362
Instance1980.3	16.691	1.761	29.307	8.536	104758.244	35481.637	4.219	1.497	0.199	2.254	77785.915	21133.103
Instance1980.4	16.786	1.947	29.301	8.544	104700.032	35545.607	3.943	1.398	0.199	2.288	77830.332	20833.049
Instance1980.5	16.796	2.201	29.222	8.522	104828.699	35594.566	3.835	1.477	0.200	2.258	79227.622	20931.568
Instance1980.6	17.069	2.330	29.048	8.519	104855.520	35546.725	3.719	1.436	0.203	2.242	79155.941	21583.087
Instance1980.7	16.842	2.061	29.477	8.663	104294.672	35300.126	3.713	1.433	0.196	2.284	76857.270	20265.440
Instance1980.8	16.749	1.933	29.337	8.521	104594.881	35461.944	3.642	1.483	0.199	2.264	77358.416	20912.303

B.3 Jetstress report for full workload – HV-JS3

Figure 7 HV-JS3 Jetstress Performance Test Result Report

Microsoft Exchange Jetstress 2013

Performance Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	HV-JS3
Test Description	2250 Users x 2 Servers = 4500 mailboxes 8 databases .10 IOPS per mailbox 6144 MB (6GB) mailbox quota 8 threads
Test Start Time	9/30/2014 9:55:24 AM
Test End Time	9/30/2014 12:08:31 PM
Collection Start Time	9/30/2014 10:06:47 AM
Collection End Time	9/30/2014 12:06:38 PM
Jetstress Version	15.00.0775.000
ESE Version	15.00.0995.021
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Performance_2014_9_30_9_55_41.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	230.945
Target Transactional I/O per Second	225
Initial Database Size (bytes)	14495783059456
Final Database Size (bytes)	14496454148096
Database Files (Count)	8

Jetstress System Parameters

Thread Count	8
Minimum Database Cache	256.0 MB
Maximum Database Cache	2048.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Run Background Database Maintenance	True
Number of Copies per Database	2

Database Configuration

Instance3272.1 Log path: C:\DB\JS3-DB1
Database: C:\DB\JS3-DB1\Jetstress001001.edb

Instance3272.2 Log path: C:\DB\JS3-DB2
Database: C:\DB\JS3-DB2\Jetstress002001.edb

Instance3272.3 Log path: C:\DB\JS3-DB3
Database: C:\DB\JS3-DB3\Jetstress003001.edb

Instance3272.4 Log path: C:\DB\JS3-DB4
Database: C:\DB\JS3-DB4\Jetstress004001.edb

Instance3272.5 Log path: C:\DB\JS3-DB5
Database: C:\DB\JS3-DB5\Jetstress005001.edb

Instance3272.6 Log path: C:\DB\JS3-DB6
Database: C:\DB\JS3-DB6\Jetstress006001.edb

Instance3272.7 Log path: C:\DB\JS3-DB7
Database: C:\DB\JS3-DB7\Jetstress007001.edb

Instance3272.8 Log path: C:\DB\JS3-DB8
Database: C:\DB\JS3-DB8\Jetstress008001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance3272.1	19.327	1.993	20.359	8.576	33716.986	35581.914	0.000	1.525	0.000	2.271	0.000	20828.158
Instance3272.2	16.828	1.755	20.379	8.549	33648.245	35406.578	0.000	1.409	0.000	2.229	0.000	20573.994
Instance3272.3	16.845	1.790	20.342	8.731	33711.945	35544.559	0.000	1.499	0.000	2.313	0.000	20782.127
Instance3272.4	16.912	2.119	20.360	8.795	33727.245	35383.395	0.000	1.491	0.000	2.325	0.000	20261.544
Instance3272.5	16.782	2.000	20.272	8.498	33648.209	35515.371	0.000	1.406	0.000	2.246	0.000	21085.758
Instance3272.6	16.708	1.827	20.013	8.330	33607.714	35599.670	0.000	1.439	0.000	2.243	0.000	21263.810
Instance3272.7	16.812	1.944	20.291	8.663	33760.628	35491.976	0.000	1.469	0.000	2.260	0.000	21200.836
Instance3272.8	16.726	1.895	20.287	8.501	33708.566	35595.460	0.000	1.472	0.000	2.257	0.000	20868.694

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance3272.1	9.105	261907.417
Instance3272.2	9.115	261864.865
Instance3272.3	9.108	261843.032
Instance3272.4	9.111	261922.414
Instance3272.5	9.109	261855.315
Instance3272.6	9.110	261899.007
Instance3272.7	9.106	261963.342
Instance3272.8	9.111	261924.853

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance3272.1	0.198	77934.343
Instance3272.2	0.193	75617.140
Instance3272.3	0.202	78495.694
Instance3272.4	0.201	78346.785
Instance3272.5	0.198	77505.947
Instance3272.6	0.202	79012.823
Instance3272.7	0.202	78495.694
Instance3272.8	0.197	77018.396

Total I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance3272.1	19.327	1.993	29.463	8.576	104230.819	35581.914	3.426	1.525	0.198	2.271	77934.343	20828.158
Instance3272.2	16.828	1.755	29.494	8.549	104179.224	35406.578	4.063	1.409	0.193	2.229	75617.140	20573.994
Instance3272.3	16.845	1.790	29.451	8.731	104267.266	35544.559	4.131	1.499	0.202	2.313	78495.694	20782.127
Instance3272.4	16.912	2.119	29.471	8.795	104273.417	35383.395	3.885	1.491	0.201	2.325	78346.785	20261.544
Instance3272.5	16.782	2.000	29.381	8.498	104398.714	35515.371	3.791	1.406	0.198	2.246	77505.947	21085.758
Instance3272.6	16.708	1.827	29.124	8.330	105022.049	35599.670	3.671	1.439	0.202	2.243	79012.823	21263.810
Instance3272.7	16.812	1.944	29.397	8.663	104447.826	35491.976	3.421	1.469	0.202	2.260	78495.694	21200.836
Instance3272.8	16.726	1.895	29.397	8.501	104436.743	35595.460	3.217	1.472	0.197	2.257	77018.396	20868.694

B.4 Jetstress report disk loss test – HV-JS1

Figure 8 Jetstress report for RAID 6 disk loss test for HV-JS1

Microsoft Exchange Jetstress 2013

Performance Test Result Report

Test Summary

Overall Test Result	**Fail**
Machine Name	HV-JS1
Test Description	2250 Users x 2 Servers = 4500 mailboxes 8 databases .10 IOPS per mailbox 6144 MB (6GB) mailbox quota 8 threads Trial 4 - disk console downing
Test Start Time	10/3/2014 9:35:42 AM
Test End Time	10/3/2014 12:23:35 PM
Collection Start Time	10/3/2014 9:47:10 AM
Collection End Time	10/3/2014 11:46:56 AM
Jetstress Version	15.00.0775.000
ESE Version	15.00.0995.021
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Performance_2014_10_3_9_36_0.blg

Test Issues

Fail The process has average database read latencies higher than 20.000 msec.

Fail The process failed to meet target IOPS (base: 225.0 and margin: 5.0 %).

Database Sizing and Throughput

Achieved Transactional I/O per Second	199.332
Target Transactional I/O per Second	225
Initial Database Size (bytes)	14501831245824
Final Database Size (bytes)	14502502334464
Database Files (Count)	8

Jetstress System Parameters

Thread Count	8
Minimum Database Cache	256.0 MB
Maximum Database Cache	2048.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Run Background Database Maintenance	True
Number of Copies per Database	2

Database Configuration

Instance1548.1 Log path: C:\DB\JS1-DB1
Database: C:\DB\JS1-DB1\Jetstress001001.edb

Instance1548.2 Log path: C:\DB\JS1-DB2
Database: C:\DB\JS1-DB2\Jetstress002001.edb

Instance1548.3 Log path: C:\DB\JS1-DB3
Database: C:\DB\JS1-DB3\Jetstress003001.edb

Instance1548.4 Log path: C:\DB\JS1-DB4
Database: C:\DB\JS1-DB4\Jetstress004001.edb

Instance1548.5 Log path: C:\DB\JS1-DB5
Database: C:\DB\JS1-DB5\Jetstress005001.edb

Instance1548.6 Log path: C:\DB\JS1-DB6
Database: C:\DB\JS1-DB6\Jetstress006001.edb

Instance1548.7 Log path: C:\DB\JS1-DB7
Database: C:\DB\JS1-DB7\Jetstress007001.edb

Instance1548.8 Log path: C:\DB\JS1-DB8
Database: C:\DB\JS1-DB8\Jetstress008001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance1548.1	**21.428**	2.029	17.435	7.422	34032.219	36434.385	0.000	1.586	0.000	1.830	0.000	20999.560
Instance1548.2	18.463	1.899	17.363	7.329	34326.260	36440.772	0.000	1.613	0.000	1.824	0.000	21098.408
Instance1548.3	18.568	1.883	17.618	7.768	34178.129	36328.098	0.000	1.532	0.000	1.873	0.000	21011.175
Instance1548.4	18.467	1.947	17.507	7.470	34150.379	36182.825	0.000	1.553	0.000	1.862	0.000	20918.545
Instance1548.5	18.709	2.128	17.375	7.452	34064.503	36309.906	0.000	1.561	0.000	1.864	0.000	21067.293
Instance1548.6	18.643	2.090	17.424	7.498	34149.492	36298.774	0.000	1.548	0.000	1.820	0.000	21483.264
Instance1548.7	18.648	2.053	17.514	7.412	34377.247	36636.074	0.000	1.565	0.000	1.826	0.000	21225.332
Instance1548.8	18.715	1.972	17.391	7.352	34209.585	36480.856	0.000	1.557	0.000	1.827	0.000	20929.957

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1548.1	9.105	261761.408
Instance1548.2	9.115	261862.141
Instance1548.3	9.109	261854.268
Instance1548.4	9.111	261882.184
Instance1548.5	9.113	261921.980
Instance1548.6	9.110	261906.646
Instance1548.7	9.111	261874.274
Instance1548.8	9.113	261743.087

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance1548.1	0.161	62834.911
Instance1548.2	0.159	62453.266
Instance1548.3	0.166	64829.624
Instance1548.4	0.162	64268.272
Instance1548.5	0.164	63869.167
Instance1548.6	0.164	63869.167
Instance1548.7	0.161	62406.515
Instance1548.8	0.158	62433.100

Total I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance1548.1	21.428	2.029	26.540	7.422	112156.333	36434.385	4.084	1.586	0.161	1.830	62834.911	20999.560
Instance1548.2	18.463	1.899	26.478	7.329	112653.224	36440.772	3.342	1.613	0.159	1.824	62453.266	21098.408
Instance1548.3	18.568	1.883	26.728	7.768	111776.559	36328.098	4.348	1.532	0.166	1.873	64829.624	21011.175
Instance1548.4	18.467	1.947	26.618	7.470	112101.416	36182.825	5.068	1.553	0.162	1.862	64268.272	20918.545
Instance1548.5	18.709	2.128	26.488	7.452	112456.539	36309.906	3.487	1.561	0.164	1.864	63869.167	21067.293
Instance1548.6	18.643	2.090	26.534	7.498	112344.942	36298.774	2.838	1.548	0.164	1.820	63869.167	21483.264
Instance1548.7	18.648	2.053	26.625	7.412	112224.743	36636.074	3.631	1.565	0.161	1.826	62406.515	21225.332
Instance1548.8	18.715	1.972	26.504	7.352	112442.407	36480.856	3.947	1.557	0.158	1.827	62433.100	20929.957

B.5 Jetstress report disk loss test – HV-JS3

Figure 9 Jetstress report for RAID 6 disk loss test for HV-JS3

Microsoft Exchange Jetstress 2013

Performance Test Result Report

Test Summary

Overall Test Result	**Fail**
Machine Name	HV-JS3
Test Description	2250 Users x 2 Servers = 4500 mailboxes 8 databases .10 IOPS per mailbox 6144 MB (6GB) mailbox quota 8 threads Trial 4 - disk console downing
Test Start Time	10/3/2014 9:35:40 AM
Test End Time	10/3/2014 12:24:34 PM
Collection Start Time	10/3/2014 9:47:05 AM
Collection End Time	10/3/2014 11:46:53 AM
Jetstress Version	15.00.0775.000
ESE Version	15.00.0995.021
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Performance_2014_10_3_9_35_58.blg

Test Issues

Fail The process has average database read latencies higher than 20.000 msec.

Fail The process failed to meet target IOPS (base: 225.0 and margin: 5.0 %).

Database Sizing and Throughput

Achieved Transactional I/O per Second	201.188
Target Transactional I/O per Second	225
Initial Database Size (bytes)	14501697028096
Final Database Size (bytes)	14502376505344
Database Files (Count)	8

Jetstress System Parameters

Thread Count	8
Minimum Database Cache	256.0 MB
Maximum Database Cache	2048.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%
Run Background Database Maintenance	True
Number of Copies per Database	2

Database Configuration

Instance432.1 Log path: C:\DB\JS3-DB1
Database: C:\DB\JS3-DB1\Jetstress001001.edb

Instance432.2 Log path: C:\DB\JS3-DB2
Database: C:\DB\JS3-DB2\Jetstress002001.edb

Instance432.3 Log path: C:\DB\JS3-DB3
Database: C:\DB\JS3-DB3\Jetstress003001.edb

Instance432.4 Log path: C:\DB\JS3-DB4
Database: C:\DB\JS3-DB4\Jetstress004001.edb

Instance432.5 Log path: C:\DB\JS3-DB5
Database: C:\DB\JS3-DB5\Jetstress005001.edb

Instance432.6 Log path: C:\DB\JS3-DB6
Database: C:\DB\JS3-DB6\Jetstress006001.edb

Instance432.7 Log path: C:\DB\JS3-DB7
Database: C:\DB\JS3-DB7\Jetstress007001.edb

Instance432.8 Log path: C:\DB\JS3-DB8
Database: C:\DB\JS3-DB8\Jetstress008001.edb

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance432.1	**21.112**	2.008	17.619	7.483	34044.094	36345.076	0.000	1.667	0.000	1.864	0.000	20727.972
Instance432.2	18.682	1.999	17.588	7.466	34144.404	36582.373	0.000	1.501	0.000	1.826	0.000	21106.132
Instance432.3	18.812	1.972	17.619	7.510	34092.070	36340.863	0.000	1.553	0.000	1.866	0.000	20988.918
Instance432.4	18.872	2.067	17.682	7.560	34085.262	36429.355	0.000	1.630	0.000	1.878	0.000	21005.638
Instance432.5	18.485	2.037	17.566	7.606	34017.657	36432.660	0.000	1.649	0.000	1.892	0.000	20958.096
Instance432.6	18.440	2.016	17.674	7.493	33957.573	36385.703	0.000	1.633	0.000	1.851	0.000	20485.741
Instance432.7	18.538	1.991	17.609	7.485	34055.616	36429.374	0.000	1.638	0.000	1.866	0.000	21033.480
Instance432.8	19.090	1.986	17.760	7.468	33928.917	36405.830	0.000	1.655	0.000	1.866	0.000	20330.414

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance432.1	9.104	261858.966
Instance432.2	9.115	261845.652
Instance432.3	9.110	261857.133
Instance432.4	9.110	261875.530
Instance432.5	9.113	261857.203
Instance432.6	9.113	261891.525
Instance432.7	9.111	261855.744
Instance432.8	9.022	261767.116

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance432.1	0.163	63381.617
Instance432.2	0.162	63232.707
Instance432.3	0.164	63869.167
Instance432.4	0.166	64873.847
Instance432.5	0.166	65302.529
Instance432.6	0.161	62406.515
Instance432.7	0.163	63915.919
Instance432.8	0.159	61987.421

Total I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance432.1	21.112	2.008	26.723	7.483	111658.322	36345.076	3.114	1.667	0.163	1.864	63381.617	20727.972
Instance432.2	18.682	1.999	26.703	7.466	111870.684	36582.373	3.277	1.501	0.162	1.826	63232.707	21106.132
Instance432.3	18.812	1.972	26.729	7.510	111723.784	36340.863	3.433	1.553	0.164	1.866	63869.167	20988.918
Instance432.4	18.872	2.067	26.792	7.560	111538.446	36429.355	4.104	1.630	0.166	1.878	64873.847	21005.638
Instance432.5	18.485	2.037	26.680	7.606	111842.836	36432.660	3.420	1.649	0.166	1.892	65302.529	20958.096
Instance432.6	18.440	2.016	26.787	7.493	111503.707	36385.703	2.975	1.633	0.161	1.851	62406.515	20485.741
Instance432.7	18.538	1.991	26.720	7.485	111731.353	36429.374	4.738	1.638	0.163	1.866	63915.919	21033.480
Instance432.8	19.090	1.986	26.782	7.468	110679.408	36405.830	5.119	1.655	0.159	1.866	61987.421	20330.414

C Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Referenced or recommended Dell publications:

Title	Published/Updated	Type
Sizing and Best Practices for Deploying Microsoft Exchange Server 2013 with Dell Compellent Storage Arrays	August 2014	Deployment Sizing Guide
Microsoft ESRP - Dell Compellent SC4020 v6.5 10,000 Mailbox Exchange 2013 Mailbox Resiliency Solution	June 2014	Technical Report
Microsoft ESRP - Dell Compellent SC4020 v6.5 4,500 Mailbox Exchange 2013 Mailbox Resiliency Solution	June 2014	Technical Report
Microsoft ESRP - Dell Compellent Storage Center 6.4 / 25,000 mailboxes	January 2014	Technical Report
Microsoft ESRP - Dell Compellent Storage Center 6.4 / 50,000 mailboxes	January 2014	Technical Report
Dell Compellent Jetstress Overview - Testing Best Practice Requirements	November 2013	Solutions Guide
Dell Compellent Microsoft Exchange Server 2013 Best Practices	January 2014	Best Practices
Dell Compellent SC8000 Controller Details — Data Center SAN	2014	Product Specs
Dell Storage SC4020 all-in-one array	June 2014	Product Specs

Referenced or recommended Microsoft publications:

- Microsoft Technet article, Exchange 2013 storage configuration options: http://technet.microsoft.com/en-us/library/ee832792(v=exchg.150).aspx