



Dell PowerVault MD3820f 500 user Mailbox Exchange 2013 Resiliency Storage Solution — FC SAN using dual QLogic QLE2662 16Gb FC adapters and Brocade 6505 16Gb FC switch

Microsoft ESRP 4.0

Dell MD3 Series storage solutions
September 2015



Revisions

Date	Description
Sept. 2015	Initial release

Disclaimer

This technical paper has been produced independently of Microsoft Corporation. Microsoft Corporation expressly disclaims responsibility for, and makes no warranty, express or implied, with respect to the accuracy of the contents of this document.

The information in this document represents the current view of Dell on the issues discussed as of the date of publication. Due to changing market conditions, it should not be interpreted to be a commitment on the part of Dell and cannot guarantee the accuracy of any information presented after the date of publication.

About Microsoft ESRP-Storage program

The Microsoft ESRP-Storage program focuses on storage solution testing to address performance and reliability issues with storage design. However, storage is not the only factor to take into consideration when designing a scale up Exchange solution. Other factors which affect the server scalability are: server processor utilization, server physical and virtual memory limitations, resource requirements for other applications, directory and network service latencies, network infrastructure limitations, replication and recovery requirements, and client usage profiles. All these factors are beyond the scope of this paper. Therefore, the number of mailboxes hosted per server as part of the tested configuration may not necessarily be viable for some customer deployments.

For more information on identifying and addressing performance bottlenecks in an Exchange system, please refer to Microsoft's Troubleshooting Microsoft Exchange Server Performance, available at <http://go.microsoft.com/fwlink/?LinkId=23454>.

© 2015 Dell Inc. All Rights Reserved. Dell, the Dell logo, and other Dell names and marks are trademarks of Dell Inc. in the US and worldwide. All other trademarks mentioned herein are the property of their respective owners.



Table of Contents

1 Executive summary	4
1.1 Overview	4
1.2 Simulated environment	4
1.3 Solution description	5
2 The Dell MD3820f solution for Microsoft ESRP	6
2.1 A modular hardware design	6
2.2 Dell PowerEdge R720 features	9
2.3 QLogic QLE2662 FC adapter	9
2.4 Brocade 6505 FC switch	9
2.5 Storage sizing	10
2.6 Targeted customer profile	10
2.7 Volume sizing	11
3 Tested deployment	11
3.1 Simulated Exchange configuration	11
3.2 Primary storage hardware	12
3.3 Primary storage software	12
3.4 Primary storage disk configuration (Mailbox store/Log disks)	12
4 Best practices	13
5 Test results summary	16
5.1 Reliability	16
5.2 Storage performance results	16
5.3 Database backup/recovery performance	17
5.3.1 Database read-only performance	17
5.3.2 Transaction log recovery/Replay performance	17
6 Conclusion	17
7 Additional resources	18
8 Appendix - Performance testing details	19
A. Performance testing details	19
B Stress testing	21
C Backup testing	23
D Recovery testing	24



1 Executive summary

1.1 Overview

This technical paper describes a tested and validated resilient storage solution for a 500 user mailbox Microsoft Exchange 2013 site, with Data Availability Group (DAG). A DAG is a high availability mechanism in Microsoft Exchange 2013.

The “Low Maintenance” concept of this configuration is based on the self-healing data protection capability of the Dell PowerVault MD3820f storage array using Dynamic Disk Pooling (DDP) technology. DDP enables the solution to withstand multiple drive failures over time without requiring drive maintenance actions by the customer. In addition to up to 8x faster rebuilds during a drive failure, DDP also provides higher levels of system performance during drive failures, delivering improved service to the infrastructure end-users. This capability can be used to design system solutions that require no drive maintenance for multiple years, significantly lowering the operational and therefore total cost of system ownership. Dynamic Disk Pooling is a standard (no-cost) feature of the PowerVault MD3 Series storage. DDP requires a minimum of 11 drives in the pool, so to see the benefits of “low maintenance” it is recommended to add two additional drives to the pool. This will provide at least two years of predicted “no drive maintenance” based on standard drive failure rates.

This mailbox resiliency model supports multiple copies (up to 16) of Exchange database in a DAG. There can be only one active copy of a given Exchange 2013 database at any given time. Secondary copies, including the copies located at remote sites, are periodically synched with the primary copy. Mail clients access the primary (active) copy, and database changes to the primary copy are copied to the secondary (passive) copies in the form of transaction logs. The copied log records are played on the secondary copy to keep the secondary database copies consistent with the primary copy. All hosts within a DAG are configured to be identical in terms of storage resources for Exchange 2013 databases and logs. The primary and secondary copies do not share any storage resources and reside on their own dedicated storage resources, as discussed later in this document.

This document provides information on a specific Dell MD3820f solution for Microsoft Exchange Server, based on the Microsoft Exchange Solution Reviewed Program (ESRP) Storage program.

The ESRP–Storage program was developed by Microsoft Corporation to provide a common storage testing framework for vendors for information on its storage solutions with Microsoft Exchange Server software. Details about the Microsoft ESRP – Storage program are available at <http://technet.microsoft.com/en-us/exchange/ff182054.aspx>.

1.2 Simulated environment

This Mailbox Resiliency solution utilizes one Database Availability Group (DAG) and two copies of every database with (DDP) Dynamic Disk Pool technology. The tested environment simulates all users in this DAG running on a single MD3820f array. The tested environment simulates up to 500 users with a 2GB Mailbox size and 200 messages a day, or 0.12 IOPS for every user, including 20% headroom.



1.3 Solution description

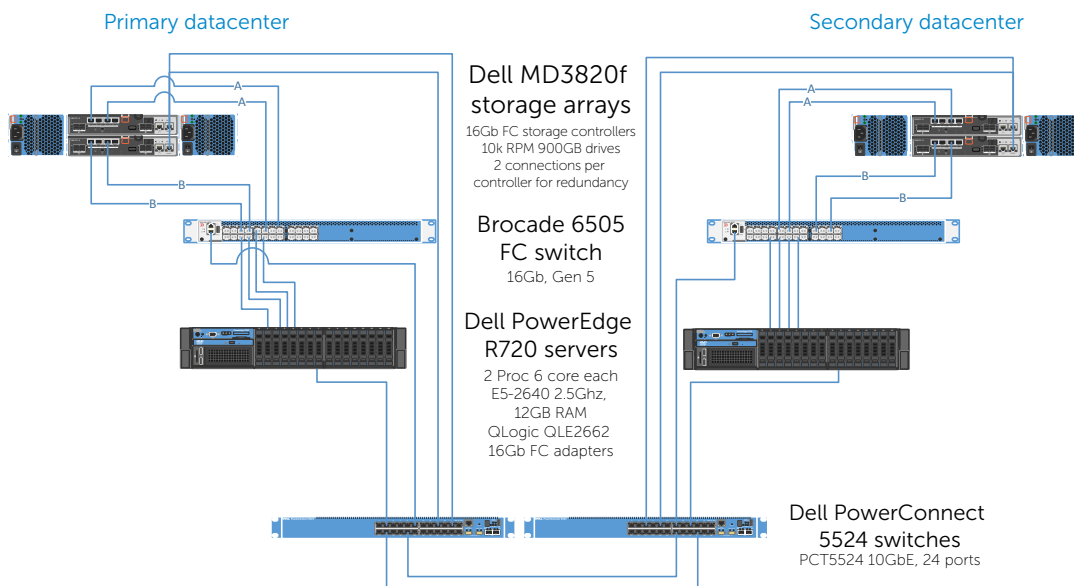
Testing was performed on a Dell R720 server, dual QLogic QLE2662 16Gb FC HBAs, Brocade 6505 16Gb FC switch and a Dell MD3820f storage array with redundant controller pair; front-end connections and back-end connections. Exchange is a critical application in most businesses today and the design of the system supporting Exchange should have redundant components and a design to support continued operation in case a single component fails. This solution was designed with the ability to support continuous operation during component failure.

The MD3820f is a 2U drive enclosure with 24 2.5" drive slots offering four 16Gb Fibre Channel and two 12Gb SAS host connections per controller. Eleven 10k 900GB 6Gbps SAS drives were used in the dedicated dynamic disk pool (DDP). As a redundant solution, databases and logs were stored together on the same volumes using Microsoft best practices. Given the self healing benefits of DDP consideration should be given to add additional HDDs to provide for a long term "no drive replacement" scenario. Adding 5% drive overhead to the drive pool provides for a predicted two years, or more, of no drive maintenance, based on typical drive failure rates. The cost of two additional drives is very low when compared to a skilled professionals time to have to order a new drive and travel to a remote site to replace a single drive.

Information about compatibility is available at <http://www.windowsservercatalog.com/item.aspx?idItem=467135f9-8f78-bfed-b511-f62d42b2d1cb&bCatID=1338>.

This figure illustrates the architectural design of the solution showing both primary site and secondary site configurations. This solution was tested on the primary site. The secondary site illustrates what a typical configuration would look like if a redundant Exchange environment were implemented.

Fibre Channel SAN storage diagram




2 The Dell MD3820f solution for Microsoft ESRP

2.1 A modular hardware design

The PowerVault MD3820f enclosure is designed to scale the needs of applications requiring large amounts of data storage. The MD3820f is a 24-drive, 2U standard rack enclosure and can scale up to 192 drives using MD1220 expansion enclosures. The MD3 Series is available in 16Gb Fibre Channel and 12Gb SAS host interfaces, 10Gb iSCSI and 12Gb SAS host interfaces or 12Gb SAS host interfaces. The MD3 Series also comes in a 2U 12-drive 3.5 inch drive module, 2U 24-drive 2.5 inch drive module or 4U 60-drive module supporting either 2.5 or 3.5 inch drives. The PowerVault MD3 Series supports simultaneous use of multiple host protocols making it highly adaptable to customer infrastructure environments. The solution described in this paper utilizes the 16Gb FC interface.

Figure 1 Dell PowerVault MD3820f front and back view



The MD3820f supports SAS, SED SAS, near-line SAS (NL-SAS), SED NL-SAS and SSD drives. The ability to mix SAS, near-line SAS and SSD drives within the same enclosure enables the user to blend drives to best suit their application storage needs across three tiers of performance offerings. Non-disruptive and on-line firmware upgrades are designed to enable high availability.

The storage management software, PowerVault Modular Disk Storage Manager (MDSM), was used to configure the storage for this solution. The MD storage management software has three major components:

- Client management software
- Host-agent management software
- Multi-path driver software

The client management software contains the graphical user interface for managing the storage array. It also contains an optional monitor service that sends alerts when an event occurs in the storage array.

The host-agent management software is installed on one or more hosts that are connected to the storage arrays to enable in-band management. The host-agent management software, along with the Ethernet connection on the host, provides another network management connection to the storage array, rather than using the individual Ethernet connections on each RAID controller module in the storage array.

The multi-path driver is also referred to as the I/O path failover driver. With the redundant pair of active RAID controller modules in a storage array, when a virtual disk is created, one of the RAID controller modules is automatically or manually chosen to “own” the virtual disk. The I/O between the virtual disk and the application host along the I/O path is controlled by the RAID controller “owning” virtual disk. When a component along the I/O path to a RAID controller module or the RAID controller module itself fails, ownership of the virtual disks that had been assigned to that RAID controller module automatically transfer to the other RAID controller module. The multi-path driver manages this failover process.

Figure 2 shows the view of disk groups, virtual disks, and the physical disks as displayed in PowerVault Modular Disk Storage Manager. Figure 3 provides an overall summary view of the PowerVault MD3820f. The features of Dell PowerVault MD3820f are detailed in Table 1.

Figure 2 MDSM view of disk groups, virtual disks and physical disks

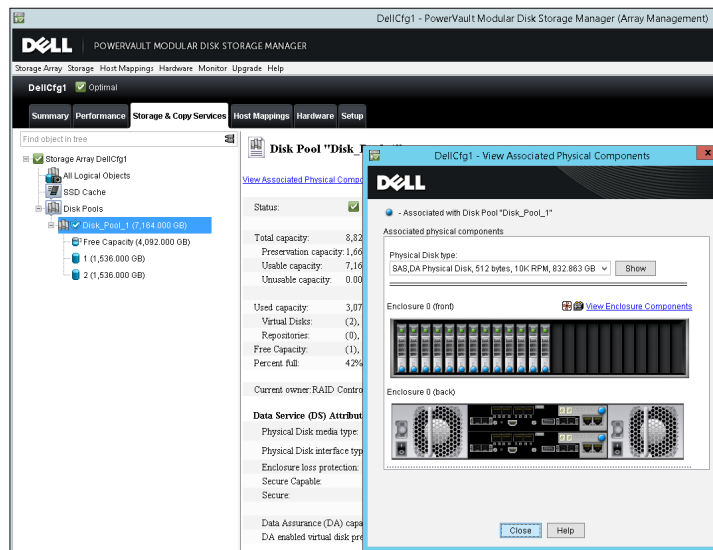


Figure 3 MDSM summary view

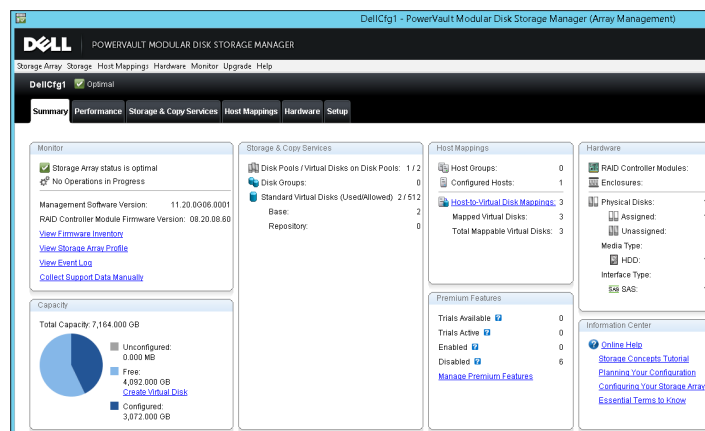


Table 1 Dell PowerVault MD3820f Features

Feature	Details
2U, 24 drive FC enclosure	Designed to fit standard 1000mm cabinets (32" max depth).
6Gb/s SAS-based storage system	Provides a high availability and high capacity storage offering when using 6GB near-line SAS drives.
Ports	Eight 16Gb/s FC ports (4 per each controller) and four 12Gb/s SAS ports (2 per controller)
Scales to support up to 192 2.5-in SAS drives	<ul style="list-style-type: none"> Up to 120 drive slots are supported as part of the base; moving from 121-180 drive slots requires purchase of the Premium Feature Key (PFK) for firmware High performance SAS, NL-SAS, SEDs and SSDs drives Configuration supports up to seven additional MD1220 expansion modules.
Support for SAS, near-line SAS and Solid State Disk drives	The ability to mix SAS, near-line SAS and SSD drives within the same enclosure supports a user's ability to blend drives to best suit their applications' storage needs across three tiers of performance offerings.
Non-disruptive, on-line firmware upgrades	Improved data availability
High Performance Tiering (HPT)	Increases system performance
SSD Cache (included as part of HPT)	Increases execution speed of applications by caching previously read data.
Thin Provisioning	<ul style="list-style-type: none"> Allocate and consume physical storage capacity as needed Thin virtual disk can only be created from a disk pool Reduces the likelihood of having excess, unused capacity in the disk pool
Support for self-encrypting drives (SED)	Secures data at rest.
VMware VAAI support	<ul style="list-style-type: none"> The ability to integrate array commands with VMware, allowing for an increased number of VM's. Reduces SAN traffic as functions are executed in the array.
Dynamic Disk Pools	<ul style="list-style-type: none"> Dynamically rebalances data in the event of a drive failure Allows for the creation of pools without the complexity of RAID Enables Thin Provisioning
Asymmetric Logical Unit Access (ALUA)	Enables the array to service I/O requests through either RAID controller module

2.2 Dell PowerEdge R720 Features

Dell PowerEdge™ R720 is a 2-socket CPU, 1U, multi-purpose server, offering an excellent balance of redundancy and value in a compact form factor. It is a most suitable hardware building block for any mid-size or large business. It delivers enormous performance in a dense 1U form-factor, enabling larger and more efficient databases and mail servers.

Major features of the server/storage system include:

- Intel® Xeon® processor E5-2600 or E5-2600 v2 product family
- Dual processor sockets
- Up to 768 GB of Memory with 24 DIMMs
- Integrated RAID support through PERC H310, PERC H710, PERC H710P
- Up to three PCIe 3.0 expansion slots
- Choice of NIC technologies
- Dell OpenManage™ Essentials and Dell Management Console, Dell OpenManage Power Center and Dell OpenManage Connections

For more information, see Dell PowerEdge R720 Server product page.

2.3 QLogic QLE2662 FC adapter

The QLogic Fibre Channel adapters have the following design characteristics:

- 16Gbps per port maximum throughput for high bandwidth (SAN) traffic
- Over 1.2 million IOPS reduces latency in high transaction intensive applications and virtualized environments
- Optimization for virtualized environments: with increasing numbers of VMs on virtualized servers it is essential that the I/O performance scales as the VM count grows and doesn't become a bottleneck
- Decreased power and cooling costs by using the fewest PCI Express® lanes in PCIe® Gen 3 environments
- Overlapping protection domains (OPDs) to ensure a high level of reliability as data moves to and from the PCI bus and Fibre Channel network
- Complete investment protection for legacy 8Gb Fibre Channel infrastructure

2.4 Brocade 6505 FC switch

The Brocade 6505 SAN Switch has the following design features:

- Provides exceptional price/performance value, combining flexibility, simplicity, and enterprise-class functionality in a 24-port, 1U entry-level switch
- Enables fast, easy, and cost-effective scaling from 12 to 24 ports using Ports on Demand (PoD) capabilities
- Simplifies management through Brocade Fabric Vision technology, reducing operational costs and optimizing application performance



- Simplifies deployment and supports high-performance fabrics by using Brocade ClearLink Diagnostic Ports (D_Ports) to identify optic and cable issues
- Maximizes resiliency with non-disruptive software upgrades and an optional redundant power supply
- Simplifies deployment with the Brocade EZSwitch Setup wizard
- Simplifies server connectivity by deploying as a full-fabric switch or a Brocade Access Gateway

2.5 Storage Sizing

Storage sizing typically involves the type of data protection chosen, type of disks, and the number of disks, both from a capacity and IOPS perspective. Selecting the right storage is crucial to achieve the balance between cost and performance. Jetstress tools provide a way of capturing the storage subsystem IOPS. Storage design also depends on the actual size of the mailbox on the disk, content indexing space and log space required. Microsoft Exchange 2013 Server Role Requirements Calculator can be used to derive the required IOPS for a particular user profile. Figure 4 shows the Mailbox Calculator output for 500 users with 200 messages/day profile. The recommended IOPS per server is 60. This will be the target IOPs that will be verified and tested as part of ESRP Jetstress verification. More details on this are provided in Section 6.

Figure 4 Recommended IOPS from the Microsoft Exchange 2013 Server Role Requirements Calculator

Role Requirements Results Pane - Log, Disk Space, and IO Requirements				
Transaction Log Requirements	/ Database	/ Server	/ DAG	/ Environment
User Transaction Logs Generated / Day	5000	5000	10000	20000
Average Move Mailbox Transaction Logs Generated / Day	1945	1945	3889	7779
Average Transaction Logs Generated / Day	6945	6945	13889	27779
Disk Space Requirements	/ Database	/ Server	/ DAG	/ Environment
Transport Database Space Required	--	64 GB	257 GB	515 GB
Database Space Required	1329 GB	1329 GB	10635 GB	21270 GB
Log Space Required	47 GB	47 GB	380 GB	760 GB
Database+Log Volume Space Required	2009 GB	2009 GB	16072 GB	32144 GB
Log Volume Space Required	0 GB	0 GB	0 GB	0 GB
Restore Volume Space Required	--	1449 GB	5797 GB	11594 GB
Host IO and Throughput Requirements	/ Database	/ Server	/ DAG	/ Environment
Total Database Required IOPS	20	20	80	161
Total Log Required IOPS	4	4	18	35
Database Read I/O Percentage	60%	--	--	--
Background Database Maintenance Throughput Requirements	1.0 MB/s	1 MB/s	4 MB/s	8 MB/s

2.6 Targeted customer profile

This solution is targeted for a medium-sized organization. Capacity can be dynamically scaled from 600GB to over a petabyte.

1. A Dell MD3 Series storage solution can be sized for any organization
2. Up to four servers can be directly connected to the storage array via Fibre Channel or iSCSI, two via SAS
3. User I/O profile (0.09 IOPs per user, 0.12 tested, giving 20% headroom).
4. User mailbox size (2GB quota)
5. Backup strategy — VSS backup using SAN based snapshots, use Mailbox Resiliency as primary data protection mechanism.



6. Using SAN-based snapshots and boot from SAN, a complete server can be restored in minutes.
7. Dynamic Disk Pooling was chosen for data protection of the database volumes and log volumes.

2.7 Volume sizing

The volume size tested was just large enough to support the database size. Volumes on Dell MD3 storage can be grown dynamically, without affecting service. As database sizes approach volume sizes, any volume can be automatically increased on demand. This simplifies sizing, as capacity can be added as needed.

Using Dell Dynamic Volume Expansion and hot upgrades, additional disk capacity can be added as needed. If more spindles are required to accommodate growth, they can simply be added to the disk pool to grow volume space. Since volumes are not tied to spindle boundaries, adding spindles will increase performance and capacity as the system grows.

The testing environment was configured for 88% storage utilization. If the storage requirement grows beyond the design specified, additional spindles will provide additional capacity for any volume to be expanded.

3 Tested Deployment

The following tables summarize the testing environment.

3.1 Simulated Exchange configuration

Configuration Item	Detail
Number of Exchange mailboxes simulated	500
Number of DAG	1
Number of servers/DAG	2
Number of active mailboxes/server	500
Number of databases/host	2
Number of copies/database	2
Number of mailboxes/database	250
Simulated profile: I/O per second per mailbox (IOPS, include 20% headroom)	0.12
Database/Log LUN size	592.39GB
Total database size for performance testing	1.18TB
% storage capacity used by Exchange database*	18.08%

* Note: Database size and capacity utilized may not match on a thin-provisioned system, as only used pages will consume space. Pages that are allocated, but contain no data, will consume no disk space.



3.2 Primary storage hardware

Configuration Item	Detail
Storage Connectivity (Fibre Channel, SAS, SATA, iSCSI)	FC
Storage Model and OS/firmware revision	Dell MD3820f: 08.20.08.60
Storage Cache	16GB
Number of storage controllers	2
Number of storage ports	4 active FC port per controller
Maximum bandwidth of storage connectivity to host	64Gb/s (2x16Gb HBA)
Switch type/model/firmware revision	FC: Brocade 6505: 7.3.0a
HBA model and firmware	QLogic QLE2662 16Gb FC HBA: 9.1.11.3
Number of HBA's/host	2
Host server type	Dell PowerEdge R720
Total number of disks tested in solution	11
Maximum number of spindles that can be hosted in the storage	24 drive bay + dual controllers in a 2U chassis Scalable to 192 drives via modular expansion enclosures

3.3 Primary storage software

Configuration Item	Detail
HBA driver	9.1.11.3
HBA Queue Depth Setting	254
Multi-Pathing (MPI/O)	Microsoft Windows Server 2012 R2 MPI/O Round-Robin (InBox DSM)
Host OS	Windows Server 2012 R2 Datacenter (6.3.9600)
ESE.dll file version	15.00.0847.030
Replication solution name/version	Microsoft Exchange Server 2013 DAG replication

3.4 Primary storage disk configuration (Mailbox store/Log disks)

Configuration Item	Detail
Disk Type, speed and firmware revision	SAS 10k 900GB, B556
Raw capacity per disk (GB)	838.363GB
Number of physical disks in test	11
Total raw storage capacity (TB)	6.398TB
Data protection	DDP
Total formatted capacity	837.863GB
Storage capacity utilization	18.52%
Database capacity utilization	87.35%



4 Best practices

- Ensure Multipath I/O is installed and configured on the server before installing MS Exchange. This feature provides alternate paths between storage devices and hosts in case the primary path fails. This feature also provides load balancing between paths.
- Configure the page file size to be 10MB larger than the physical RAM installed in the server.
- Assign an allocation unit size of 64KB when creating volumes in Windows Server 2012. This option increases the block size of the volume being created. This setting can result in increased performance because it uses the most efficient block size for data transfer on the system bus.
- Set the start demand cache flushing value to 80% in the Dell Modular Disk Storage Manager.
- When creating volumes in the Modular Disk Storage Manager, make sure read and write cache are both enabled. Also confirm that dynamic cache read pre-fetch is enabled. These three settings increase the performance of the storage system.
- Adjust IOPs per user to 0.12 to allow for 20% headroom.
- From a controller resource allocation perspective, there are two user-modifiable reconstruction priorities within DDP. It is recommended to set these as Low or Medium priority settings for NL-SAS drives, this will increase the drive reconstruction time but will also lessen the impact of I/O performance during rebuild.
 - Degraded reconstruction priority is assigned for instances where only a single D-Piece needs to be rebuilt for affected D-Stripes. The default is 'high' setting 1.
 - Critical reconstruction priority is assigned for instances where a D-Stripe has two missing D-Pieces which need to be rebuilt. The default is 'highest'.
- Given the self healing benefits of DDP consideration should be given to add additional HDDs to provide for a long term "no drive replacement" scenario. Adding 5% drive overhead to the drive pool provides for a predicted two years, or more, of no drive maintenance, based on typical drive failure rates.
- It is best to use SAS drives with Exchange 2013 when a moderate amount of storage capacity is needed with high performance and balanced power consumption. It is also important to disable physical disk-write caching when the drives are used without an un-interruptible power supply (UPS). The 900GB 10k RPM SAS drives used in the testing were chosen for their average storage capacity, excellent random I/O performance, and great sequential I/O performance and power utilization.

Best Practice Exchange 2013 storage configuration options

[https://technet.microsoft.com/en-us/library/ee832792\(v=exchg.150\).aspx](https://technet.microsoft.com/en-us/library/ee832792(v=exchg.150).aspx)

Planning for high availability and site resilience, see [https://technet.microsoft.com/library/dd638104\(EXCHG.150\)#StoreReq](https://technet.microsoft.com/library/dd638104(EXCHG.150)#StoreReq)

Exchange Server 2013 has changed dramatically from previous versions, see [http://technet.microsoft.com/en-us/library/jj150540\(v=exchg.150\).aspx](http://technet.microsoft.com/en-us/library/jj150540(v=exchg.150).aspx)

Exchange 2013 requirements that you need to know before you install Exchange 2013, see <https://technet.microsoft.com/en-us/library/aa996719.aspx>

Exchange 2013 Sizing and Configuration Recommendations, see <https://technet.microsoft.com/en-us/library/dn879075.aspx>



Drive Best Practices

When initializing disks in Windows Server 2012, the disks should be initialized as Basic Disks. Initializing a disk as dynamic increases processor overhead as the server also becomes responsible for managing volumes. This is the recommended disk configuration by Microsoft. When formatting drives, use the GUID partition table (GPT) scheme as opposed to MBR. GPT allows volumes to reach 256TB in size.

It is also important to disable automatic disk optimization and defragmentation on Windows Server 2012. When this feature is enabled, additional processor overhead will be incurred because the system will monitor and move data around to prevent fragmentation. Confirm that NTFS compression is not enabled. Do not use NTFS encrypting file system (EFS) or resilient file system (ReFS) as these will also increase processor overhead.

Dynamic Disk Pools

Dell MD3 Series Dynamic Disk Pools (DDP) is a data protection technology designed to deliver consistent storage system performance, data protection, and efficiency throughout the lifecycle of the system. DDP simplifies the setup process and reduces the ongoing maintenance requirements of data protection. With DDP, customers do not have to define RAID array sizes, hot spares, and drive maintenance schedules. DDP distributes data, parity information, and spare capacity across a pool of drives. Its intelligent algorithm defines which drives are used for segment placement, making sure data is fully protected.

DDP is able to utilize every drive in the pool for the intensive process of rebuilding a failed drive. This dynamic rebuild technology is the key to its exceptional performance under failure and returns the system to optimal conditions up to eight times more quickly than traditional RAID technology. With shorter rebuild times and patented prioritization reconstruction technology, DDP significantly reduces exposure to numerous cascading disk failures. Flexible disk pool sizing provides optimal utilization of any configuration for maximum performance, protection, and efficiency. DDP can easily be grown by adding up to 12 additional disk drives at a time.

In addition to superior data protection, Dynamic Disk Pools enable customers to structure their storage infrastructure in a way that can greatly reduce drive maintenance schedules. Designing a disk pool with additional drive capacity for growth at system installation leverages the technology's automatic self-healing capability and can extend drive maintenance schedules by years, driving operational costs down.

Configuration flexibility enables DDP to address wide-ranging requirements. Drives can be configured into one large disk pool to maximize simplicity and protection or into numerous smaller pools to maximize sequential performance. Different drive types can be used to create storage tiers, such as performance pools and capacity pools, and disk pools can reside in the same system with traditional RAID groups.

The following are the four key benefits of DDP technology:

- Reduce performance degradation following a drive (or multiple-drive) failure
- Eliminate complex RAID management without sacrificing data protection
- Eliminate deployment and management of idle hot spare drives
- Expand or contract the disk pool without reconfiguring RAID



Backup strategy

Other features of the MD3 Series that protect data include mirroring and backing up controller cache. If power is lost to the system during operation, onboard batteries are used to de-stage the data from cache memory to internal controller flash so that it will be available when power is restored. The DDP algorithms allow the system to recreate any lost data in the rare case of drive failure. Users also have the option of confirming data with RAID parity at all times and even continuing a rebuild when hitting an un-readable sector.

Behind the scenes, the system performs other tasks that protect data at all times. The optional media scan feature looks for inconsistencies even on sectors not currently being accessed by any host. All types of diagnostic data are constantly collected for later use by support if necessary.


Not only does the MD3 Series offer the detailed reliability and availability features already described, but using the MDSM software features allows the possibility to maximize availability.

High-speed, high-efficiency Snapshot copies

- Robust disaster recovery protection
 - Synchronous mirroring for no-data-loss protection of content
 - Asynchronous mirroring for long-distance protection and compliance
- Flexible protection to maximize ROI
 - Recovery target can be flash, NL-SAS, or mixed based on cost/performance needs
 - Delivers speed without breaking budgets

Figure 5 illustrates a graphical representation of the HA possibilities using Snapshot copies and mirroring. For more information, refer to the Dell Support site Documentation library and the MDSM online help.

Figure 5: Snapshot copies



Additional information

For more information Dell MD3 Series storage solutions, visit our website at <http://www.dell.com/storage>.

5 Test results summary

This section provides a high level summary of the test data from ESRP. The detailed html reports which are generated by ESRP testing framework are shown in the Appendix later in this white-paper.

5.1 Reliability

Tests in this framework to check storage reliability are run over a 24 hour period. The goal of these "Stress tests" is to verify that the storage can handle high I/O load for a long period of time. Both log and database files were analyzed for integrity after the stress test to ensure no database/log corruption.

The following list provides an overview of reliability results:

- No errors were reported in either the application or system log
- No errors were reported during the database and log checksum process
- No errors were reported during either the backup or restore process

5.2 Storage performance results

The Primary Storage performance testing is designed to exercise the storage with maximum sustainable Exchange type I/O for 2 hours. The test illustrates how long it takes for the storage to respond to a specific mailbox profile I/O load. The data below is the sum of all the logical disk I/O and average of all the logical disks I/O latency in the 2 hour test duration. Each server is listed separately and the aggregate numbers across all servers is listed as well.

Multiple Server Metrics:

The sum of all transactional I/O performance across all mailbox databases and the average latency across all databases on a per server basis.

Database I/O	Value
Disks Reads/sec sum	992.516
Disks Writes/sec sum	643.242
Disk Read Latency (ms) average	10.170
Disk Write Latency (ms) average	2.199
Transaction Log I/O	
Log Disks Writes/sec sum	111.695
Log Disk Write Latency (ms) average	0.665

5.3 Database backup/recovery performance

There are two tests reports in this section. The first measures the sequential read rate of the database files, and the second measures the recovery/replay performance (playing transaction logs in to the database).



5.3.1 Database read-only performance

The test measures the maximum rate at which databases could be backed up via Volume Shadow Copy Service (VSS). The following table shows the average rate for a single database file.

Performance item	Detail
MB read/sec per database	371.36
MB read/sec total per server	742.72

5.3.2 Transaction log recovery/Replay performance

The purpose of this test is to measure the maximum rate at which the log files can be played against the databases. The following table shows the average rate for 500 log files played in a single database. Each log file is 1MB in size.

Performance item	Detail
Average time to play one Log file (sec)	1.61

6 Conclusion

This ESRP document presents a tested and validated Exchange solution for 500 mailboxes with 2GB mailbox size supporting up to 200 messages/day in a single DAG. The solution uses one Dell PowerEdge R720 servers attached to a PowerVault MD3820f storage array for Exchange mailbox databases and transactional logs.

Testing was carried out as part of the ESRP test framework using Microsoft Exchange Server 2013 Jetstress. The test results show that the proposed solution is more than capable of delivering the IOPs and meeting the capacity requirements to support 500 mailboxes with the set mailbox profile.

This document is developed by storage solution providers, and reviewed by the Microsoft Exchange Product team. The test results/data presented in this document are based on the tests introduced in the ESRP test framework. Customers should not quote the data directly for his/her pre-deployment verification. It is still necessary to go through the exercises to validate the storage design for a specific customer environment.

The ESRP program is not designed to be a benchmarking program, and the tests are not designed to deliver the maximum throughput for a given solution. Rather, the tests are focused on producing recommendations from vendors for Exchange application. The data presented in this document should not be used for direct comparisons among solutions.



7 Additional resources

Microsoft ESRP Program Website: <http://technet.microsoft.com/en-us/exchange/ff182054.aspx>

Dell Storage Website: <http://www.dell.com/storage/>

Dell TechCenter storage page: <http://en.community.dell.com/techcenter/storage/>



8 Appendix

Test results for each particular mailbox size, users and connection

A Performance testing

Test Summary

Overall Test Result **Pass**

Machine Name Server 1

Test Description Machine Name: Dell Poweredge R720 (non-virtual)

500 users Microsoft Exchange 2013

1 Dell Poweredge R720 server with Microsoft Server 2012 r2 installed

2GB Mailboxes, 500 users per server, 0.12 IOPs

2 DB and LOG on 2 LUNs (combined)

Dell MD3820f using Dynamic Disk Pool (11 drives) technology for data protection

Dual QLogic 16Gb FC HBAs

Brocade 16Gb Fabric Switch

Test Start Time 7/15/2015 3:01:57 PM

Test End Time 7/15/2015 7:04:25 PM

Collection Start Time 7/15/2015 5:04:23 PM

Collection End Time 7/15/2015 7:04:16 PM

Jetstress Version 15.00.0995.000

ESE Version 15.00.0847.030

Operating System Windows Server 2012 R2 Datacenter (6.2.9200.0)

Performance Log C:\Program Files\Exchange Jetstress\Performance_2015_7_15_17_3_46.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second 1635.757

Target Transactional I/O per Second 60

Initial Database Size (bytes) 1073758601216

Final Database Size (bytes) 1076929495040

Database Files (Count) 2

Jetstress System Parameters

Thread Count 20

Minimum Database Cache 64.0 MB

Maximum Database Cache 512.0 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

Instance5028.1 Log path: C:\Users\Administrator\Desktop\Volume2\Log1
Database: C:\Users\Administrator\Desktop\Volume1\DB1\Jetstress001001.edb

Instance5028.2 Log path: C:\Users\Administrator\Desktop\Volume1\Log2
Database: C:\Users\Administrator\Desktop\Volume2\DB2\Jetstress002001.edb

Transactional I/O Performance

MSEExchange 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
Instance5028.1	10.202	2.508	495.912	322.270	32907.666	34704.623	0.000	0.641	0.000	55.836	0.000	17656.108
Instance5028.2	10.138	1.890	496.604	320.972	32916.053	34699.311	0.000	0.689	0.000	55.859	0.000	17660.914



Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance5028.1	9.348	261383.342
Instance5028.2	9.348	261338.116

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance5028.1	4.141	232561.946
Instance5028.2	4.142	232560.994

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/sec	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
Instance5028.1	10.202	2.508	505.259	322.270	37134.609	34704.623	0.856	0.641	4.141	55.836	232561.946	17656.108
Instance5028.2	10.138	1.890	505.952	320.972	37136.590	34699.311	1.144	0.689	4.142	55.859	232560.994	17660.914

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.058	0.792	1.358
Available MBytes	30033.827	29980.000	30041.000
Free System Page Table Entries	16610187.497	16609914.000	16610455.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	112851067.992	111587328.000	113135616.000
Pool Paged Bytes	100450261.244	100380672.000	100618240.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

7/15/2015 3:01:57 PM -- Preparing for testing ...
7/15/2015 3:01:57 PM -- Creating C:\Users\Administrator\Desktop\Volume1\DB1\Jetstress001001.edb.
7/15/2015 3:01:57 PM -- Database cache settings: (minimum: 32.0 MB, maximum: 256.0 MB)
7/15/2015 3:01:57 PM -- Database flush thresholds: (start: 2.5 MB, stop: 5.1 MB)
7/15/2015 4:39:34 PM -- 100.0% of 500.0 GB complete (170878753 records inserted).
7/15/2015 4:39:34 PM -- 100.0% of 500.0 GB complete (170880394 records inserted).
7/15/2015 4:39:35 PM -- Duplicating 1 databases:
7/15/2015 5:03:41 PM -- 100.0% of 500.0 GB complete (500.0 GB duplicated).
7/15/2015 5:03:44 PM -- Attaching databases ...
7/15/2015 5:03:44 PM -- Preparations for testing are complete.
7/15/2015 5:03:44 PM -- Starting transaction dispatch ..
7/15/2015 5:03:44 PM -- Database cache settings: (minimum: 64.0 MB, maximum: 512.0 MB)
7/15/2015 5:03:44 PM -- Database flush thresholds: (start: 5.1 MB, stop: 10.2 MB)
7/15/2015 5:03:46 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
7/15/2015 5:03:46 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
7/15/2015 5:03:48 PM -- Operation mix: Sessions 20, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
7/15/2015 5:03:48 PM -- Performance logging started (interval: 15000 ms).
7/15/2015 5:03:48 PM -- Attaining prerequisites:
7/15/2015 5:04:23 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 487387100.0 (lower bound: 483183800.0, upper bound: none)
7/15/2015 7:04:24 PM -- Performance logging has ended.
7/15/2015 7:04:24 PM -- JetInterop batch transaction stats: 111070 and 111070.
7/15/2015 7:04:24 PM -- Dispatching transactions ends.
7/15/2015 7:04:24 PM -- Shutting down databases ...
7/15/2015 7:04:25 PM -- Instance5028.1 (complete) and Instance5028.2 (complete)
7/15/2015 7:04:25 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_7_15_17_3_46.blg has 481 samples.
7/15/2015 7:04:25 PM -- Creating test report ...
7/15/2015 7:04:27 PM -- Instance5028.1 has 10.2 for I/O Database Reads Average Latency.
7/15/2015 7:04:27 PM -- Instance5028.1 has 0.6 for I/O Log Writes Average Latency.
7/15/2015 7:04:27 PM -- Instance5028.1 has 0.6 for I/O Log Reads Average Latency.
7/15/2015 7:04:27 PM -- Instance5028.2 has 10.1 for I/O Database Reads Average Latency.
7/15/2015 7:04:27 PM -- Instance5028.2 has 0.7 for I/O Log Writes Average Latency.
7/15/2015 7:04:27 PM -- Instance5028.2 has 0.7 for I/O Log Reads Average Latency.
7/15/2015 7:04:27 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
7/15/2015 7:04:27 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
7/15/2015 7:04:27 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_7_15_17_3_46.xml has 478 samples queried.



B Stress testing

Test Summary

Overall Test Result **Pass**

Machine Name Server 1

Test Description Machine Name: Dell Poweredge R720 (non-virtual)

500 users Microsoft Exchange 2013

1 Dell Poweredge R720 server with Microsoft Server 2012 r2 installed

2GB Mailboxes, 500 users per server, 0.12 IOPs

2 DB and LOG on 2 LUNs (combined)

Dell MD3820f using Dynamic Disk Pool (11 drives) technology for data protection

Dual QLogic 16Gb FC HBAs

Brocade 16Gb Fabric Switch

Test Start Time 7/16/2015 9:37:49 AM

Test End Time 7/17/2015 9:38:33 AM

Collection Start Time 7/16/2015 9:38:32 AM

Collection End Time 7/17/2015 9:38:24 AM

Jetstress Version 15.00.0995.000

ESE Version 15.00.0847.030

Operating System Windows Server 2012 R2 Datacenter (6.2.9200.0)

Performance Log C:\Program Files\Exchange Jetstress\Stress_2015_7_16_9_37_54.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second 1575.242

Target Transactional I/O per Second 60

Initial Database Size (bytes) 1077407645696

Final Database Size (bytes) 1111188570112

Database Files (Count) 2

Jetstress System Parameters

Thread Count 20

Minimum Database Cache 64.0 MB

Maximum Database Cache 512.0 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

Instance4932.1 Log path: C:\Users\Administrator\Desktop\Volume2\Log1
Database: C:\Users\Administrator\Desktop\Volume1\DB1\Jetstress001001.edb

Instance4932.2 Log path: C:\Users\Administrator\Desktop\Volume1\Log2
Database: C:\Users\Administrator\Desktop\Volume2\DB2\Jetstress002001.edb

Transactional I/O Performance

MSEExchange 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
Instance4932.1	10.753	2.358	481.150	305.522	32882.685	33525.593	0.000	0.601	0.000	50.679	0.000	17202.111
Instance4932.2	10.615	1.810	482.319	306.250	32886.480	33524.287	0.000	0.649	0.000	50.730	0.000	17173.398



Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4932.1	9.846	261218.446
Instance4932.2	9.845	261283.934

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance4932.1	3.659	232561.897
Instance4932.2	3.656	232561.454

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
Instance4932.1	10.753	2.358	490.997	305.522	37461.613	33525.593	0.755	0.601	3.659	50.679	232561.897	17202.111
Instance4932.2	10.615	1.810	492.164	306.250	37455.185	33524.287	1.435	0.649	3.656	50.730	232561.454	17173.398

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.139	0.767	12.898
Available MBytes	29768.152	29661.000	29782.000
Free System Page Table Entries	16609689.694	16609173.000	16609979.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	116864338.011	115654656.000	117612544.000
Pool Paged Bytes	104332095.944	104112128.000	109547520.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

7/16/2015 9:37:49 AM -- Preparing for testing ...
7/16/2015 9:37:51 AM -- Attaching databases ...
7/16/2015 9:37:51 AM -- Preparations for testing are complete.
7/16/2015 9:37:51 AM -- Starting transaction dispatch ..
7/16/2015 9:37:51 AM -- Database cache settings: (minimum: 64.0 MB, maximum: 512.0 MB)
7/16/2015 9:37:51 AM -- Database flush thresholds: (start: 5.1 MB, stop: 10.2 MB)
7/16/2015 9:37:54 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
7/16/2015 9:37:54 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
7/16/2015 9:37:55 AM -- Operation mix: Sessions 20, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
7/16/2015 9:37:55 AM -- Performance logging started (interval: 15000 ms).
7/16/2015 9:37:55 AM -- Attaining prerequisites:
7/16/2015 9:38:32 AM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 493514800.0 (lower bound: 483183800.0, upper bound: none)
7/17/2015 9:38:32 AM -- Performance logging has ended.
7/17/2015 9:38:32 AM -- JetInterop batch transaction stats: 1186538 and 1186537.
7/17/2015 9:38:32 AM -- Dispatching transactions ends.
7/17/2015 9:38:33 AM -- Shutting down databases ...
7/17/2015 9:38:33 AM -- Instance4932.1 (complete) and Instance4932.2 (complete)
7/17/2015 9:38:33 AM -- C:\Program Files\Exchange Jetstress\Stress_2015_7_16_9_37_54.big has 5755 samples.
7/17/2015 9:38:33 AM -- Creating test report ...
7/17/2015 9:38:41 AM -- Instance4932.1 has 10.8 for I/O Database Reads Average Latency.
7/17/2015 9:38:41 AM -- Instance4932.1 has 0.6 for I/O Log Writes Average Latency.
7/17/2015 9:38:41 AM -- Instance4932.1 has 0.6 for I/O Log Reads Average Latency.
7/17/2015 9:38:41 AM -- Instance4932.2 has 10.6 for I/O Database Reads Average Latency.
7/17/2015 9:38:41 AM -- Instance4932.2 has 0.6 for I/O Log Writes Average Latency.
7/17/2015 9:38:41 AM -- Instance4932.2 has 0.6 for I/O Log Reads Average Latency.
7/17/2015 9:38:41 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
7/17/2015 9:38:41 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
7/17/2015 9:38:41 AM -- C:\Program Files\Exchange Jetstress\Stress_2015_7_16_9_37_54.xml has 5752 samples queried.



C Backup testing

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance3164.1	513736.03	00:23:13	368.76
Instance3164.2	513744.03	00:22:53	373.96
Avg			371.36
Sum			742.72

Jetstress System Parameters

Thread Count	20
Minimum Database Cache	64.0 MB
Maximum Database Cache	512.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%

Database Configuration

Instance3164.1	Log path: C:\Users\Administrator\Desktop\Volume2\Log1 Database: C:\Users\Administrator\Desktop\Volume1\DB1\Jetstress001001.edb
Instance3164.2	Log path: C:\Users\Administrator\Desktop\Volume1\Log2 Database: C:\Users\Administrator\Desktop\Volume2\DB2\Jetstress002001.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
Instance3164.1	1.042	0.000	1468.539	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance3164.2	1.078	0.000	1494.891	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.729	1.261	3.086
Available MBytes	30458.522	30313.000	30489.000
Free System Page Table Entries	16610059.065	16609557.000	16610324.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	115899347.478	115810304.000	116359168.000
Pool Paged Bytes	104235453.217	104198144.000	104382464.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

```

7/16/2015 9:00:06 AM -- Preparing for testing ...
7/16/2015 9:00:08 AM -- Attaching databases ...
7/16/2015 9:00:08 AM -- Preparations for testing are complete.
7/16/2015 9:00:11 AM -- Performance logging started (interval: 30000 ms).
7/16/2015 9:00:11 AM -- Backing up databases ...
7/16/2015 9:23:25 AM -- Performance logging has ended.
7/16/2015 9:23:25 AM -- Instance3164.1 (100% processed) and Instance3164.2 (100% processed)
7/16/2015 9:23:25 AM -- C:\Program Files\Exchange Jetstress\DatabaseBackup_2015_7_16_9_0_8.blg has 46 samples.
7/16/2015 9:23:25 AM -- Creating test report ...

```



D Recovery testing

Test Summary

Overall Test Result **Pass**

Machine Name Server 1

Test Description Machine Name: Dell Poweredge R720 (non-virtual)

500 users Microsoft Exchange 2013

1 Dell Poweredge R720 server with Microsoft Server 2012 r2 installed

2GB Mailboxes, 500 users per server, 0.12 IOPs

2 DB and LOG on 2 LUNs (combined)

Dell MD3820f using Dynamic Disk Pool (11 drives) technology for data protection

Dual QLogic 16Gb FC HBAs

Brocade 16Gb Fabric Switch

Test Start Time 7/16/2015 8:31:15 AM

Test End Time 7/16/2015 8:49:36 AM

Collection Start Time 7/16/2015 8:31:35 AM

Collection End Time 7/16/2015 8:49:22 AM

Jetstress Version 15.00.0995.000

ESE Version 15.00.0847.030

Operating System Windows Server 2012 R2 Datacenter (6.2.9200.0)

Performance Log C:\Program Files\Exchange Jetstress\Performance_2015_7_16_8_31_20.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second 1663.3

Target Transactional I/O per Second 60

Initial Database Size (bytes) 1076929495040

Final Database Size (bytes) 1077407645696

Database Files (Count) 2

Jetstress System Parameters

Thread Count 20

Minimum Database Cache 64.0 MB

Maximum Database Cache 512.0 MB

Insert Operations 40%

Delete Operations 20%

Replace Operations 5%

Read Operations 35%

Lazy Commits 70%

Database Configuration

Instance4888.1 Log path: C:\Users\Administrator\Desktop\Volume2\Log1
Database: C:\Users\Administrator\Desktop\Volume1\DB1\Jetstress001001.edb

Instance4888.2 Log path: C:\Users\Administrator\Desktop\Volume1\Log2
Database: C:\Users\Administrator\Desktop\Volume2\DB2\Jetstress002001.edb

Transactional I/O Performance

MSEExchange Database ==> Instances	I/O Database Reads	I/O Database Writes	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
	Average Latency (msec)	Average Latency (msec)										
Instance4888.1	10.207	2.467	506.415	325.679	32870.525	33890.854	0.000	0.644	0.000	54.401	0.000	17876.880
Instance4888.2	10.068	1.891	506.896	324.310	32880.024	33868.483	0.000	0.694	0.000	54.345	0.000	18041.794



Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.054	0.867	1.227
Available MBytes	29962.111	29951.000	30233.000
Free System Page Table Entries	16609929.708	16609673.000	16610068.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	115196928.000	114413568.000	115757056.000
Pool Paged Bytes	103961713.778	103927808.000	104116224.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

7/16/2015 8:31:15 AM -- Preparing for testing ...
 7/16/2015 8:31:17 AM -- Attaching databases ...
 7/16/2015 8:31:17 AM -- Preparations for testing are complete.
 7/16/2015 8:31:17 AM -- Starting transaction dispatch ..
 7/16/2015 8:31:17 AM -- Database cache settings: (minimum: 64.0 MB, maximum: 512.0 MB)
 7/16/2015 8:31:17 AM -- Database flush thresholds: (start: 5.1 MB, stop: 10.2 MB)
 7/16/2015 8:31:20 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
 7/16/2015 8:31:20 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
 7/16/2015 8:31:20 AM -- Operation mix: Sessions 20, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
 7/16/2015 8:31:20 AM -- Performance logging started (interval: 15000 ms).
 7/16/2015 8:31:20 AM -- Generating log files ...
 7/16/2015 8:49:35 AM -- C:\Users\Administrator\Desktop\Volume2\Log1 (100.2% generated) and C:\Users\Administrator\Desktop\Volume1\Log2 (100.8% generated)
 7/16/2015 8:49:35 AM -- Performance logging has ended.
 7/16/2015 8:49:35 AM -- JetInterop batch transaction stats: 16690 and 16689.
 7/16/2015 8:49:35 AM -- Dispatching transactions ends.
 7/16/2015 8:49:36 AM -- Shutting down databases ...
 7/16/2015 8:49:36 AM -- Instance4888.1 (complete) and Instance4888.2 (complete)
 7/16/2015 8:49:36 AM -- C:\Program Files\Exchange Jetstress\Performance_2015_7_16_8_31_20.blg has 72 samples.
 7/16/2015 8:49:36 AM -- Creating test report ...
 7/16/2015 8:49:37 AM -- Instance4888.1 has 10.2 for I/O Database Reads Average Latency.
 7/16/2015 8:49:37 AM -- Instance4888.1 has 0.6 for I/O Log Writes Average Latency.
 7/16/2015 8:49:37 AM -- Instance4888.1 has 0.6 for I/O Log Reads Average Latency.
 7/16/2015 8:49:37 AM -- Instance4888.2 has 10.1 for I/O Database Reads Average Latency.
 7/16/2015 8:49:37 AM -- Instance4888.2 has 0.7 for I/O Log Writes Average Latency.
 7/16/2015 8:49:37 AM -- Instance4888.2 has 0.7 for I/O Log Reads Average Latency.
 7/16/2015 8:49:37 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
 7/16/2015 8:49:37 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
 7/16/2015 8:49:37 AM -- C:\Program Files\Exchange Jetstress\Performance_2015_7_16_8_31_20.xml has 71 samples queried.

