Dell Storage Center 6.6 SCv2000 iSCSI Arrays and 2,500 Mailbox Exchange 2013 Resiliency Storage Solution

Microsoft ESRP 4.0

Daniel J. Curran: Principal Engineer
Dell Storage Engineering
July 2015

A Dell Technical White Paper

Revisions

Date	Description
July 2015	Initial release

Table of contents

1 Executive summary

This document provides information on a Dell™ SCv2000 Storage Center solution for Microsoft® Exchange Server, based 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 to provide information on its storage solutions for Microsoft Exchange Server software. For more details on the Microsoft ESRP – Storage program, please click http://technet.microsoft.com/en-us/exchange/ff182054.aspx

1.1 Simulated environment

The solution presented in this document is designed to simulate a small to medium-sized number of mailboxes hosted on highly redundant hardware. Application level redundancy is augmented with redundant storage to create a highly available and fault tolerant solution.

The Mailbox Resiliency features of Exchange 2013 have greatly enhanced the availability of Exchange Server, while also improving I/O performance. The solution presented here is a Mailbox Resiliency solution utilizing one Database Availability Group (DAG) and two copies of every database. The tested environment simulates all users in this DAG running on a single Storage Center, or half of the solution. The number of users simulated was 2,500 across two servers, with 1,250 users per server. The mailbox size was 2GB per user. Each server has four databases, with one copy local and the second copy replicated to the second server. This provides redundancy through hardware and software.

The replication mechanism is the native Exchange 2013 DAG database replication engine. This is a very efficient and reliable replication mechanism and is the recommended method for providing highly available and redundant Exchange solutions.

1.2 Solution description

Testing was performed on a Dell Storage Center (SCv2000) v6.6, a redundant controller pair with redundant front-end and back-end connections. The front-end connections are iSCSI-based over redundant 10GbE fabrics with two 10GbE ports per server and two 10GbE ports per storage controller. One 12 bay 3.5" built-in drive enclosure was utilized with each Storage Center.

The disk connectivity uses SAS 6Gbps. Disk drives used are SAS 7.2K 4TB. The spindle count is 11 disks/1 spare for database and logs, on a dedicated disk pool on each Storage Center. As this is a redundant solution, databases and logs are stored together on the same volumes. All volumes are RAID 6.

For information about compatibility is available at:
http://www.windowsservercatalog.com/item.aspx?idItem=f2b1c4fb-ee97-7d0d-b671-bf5a75f63c2d&bCatID=1338

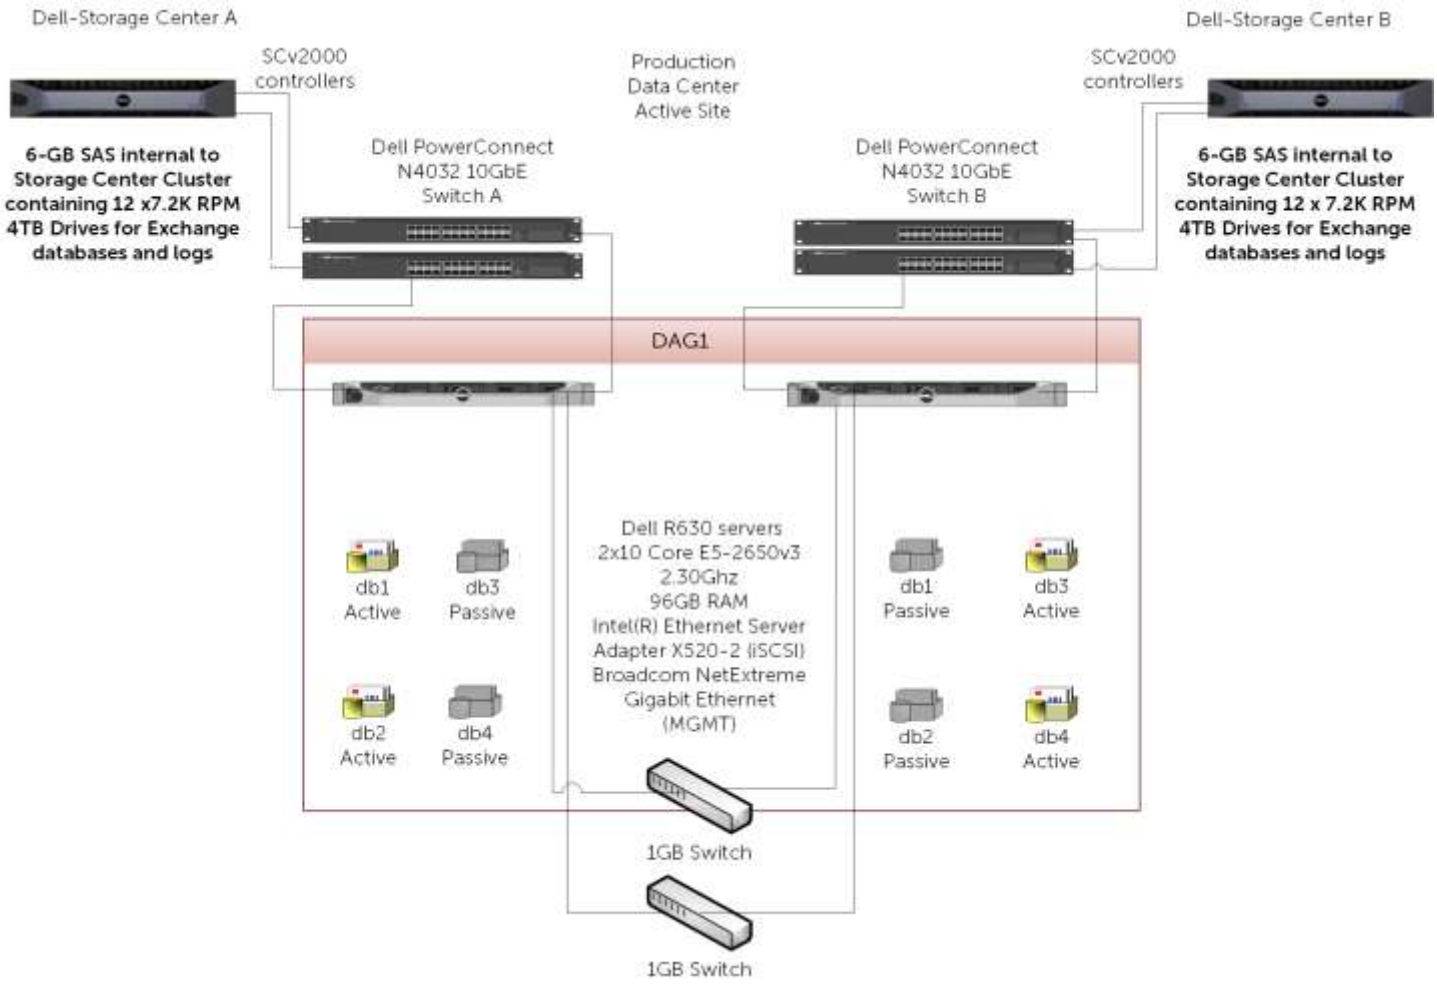

Figure 1 Highly available data center design

The solution is designed around a highly available data center model (Figure 1). There are two disk arrays, for complete redundancy. The Exchange configuration is one DAG. The LAN ports are in a dedicated replication VLAN for traffic isolation. There are two networks for redundancy.

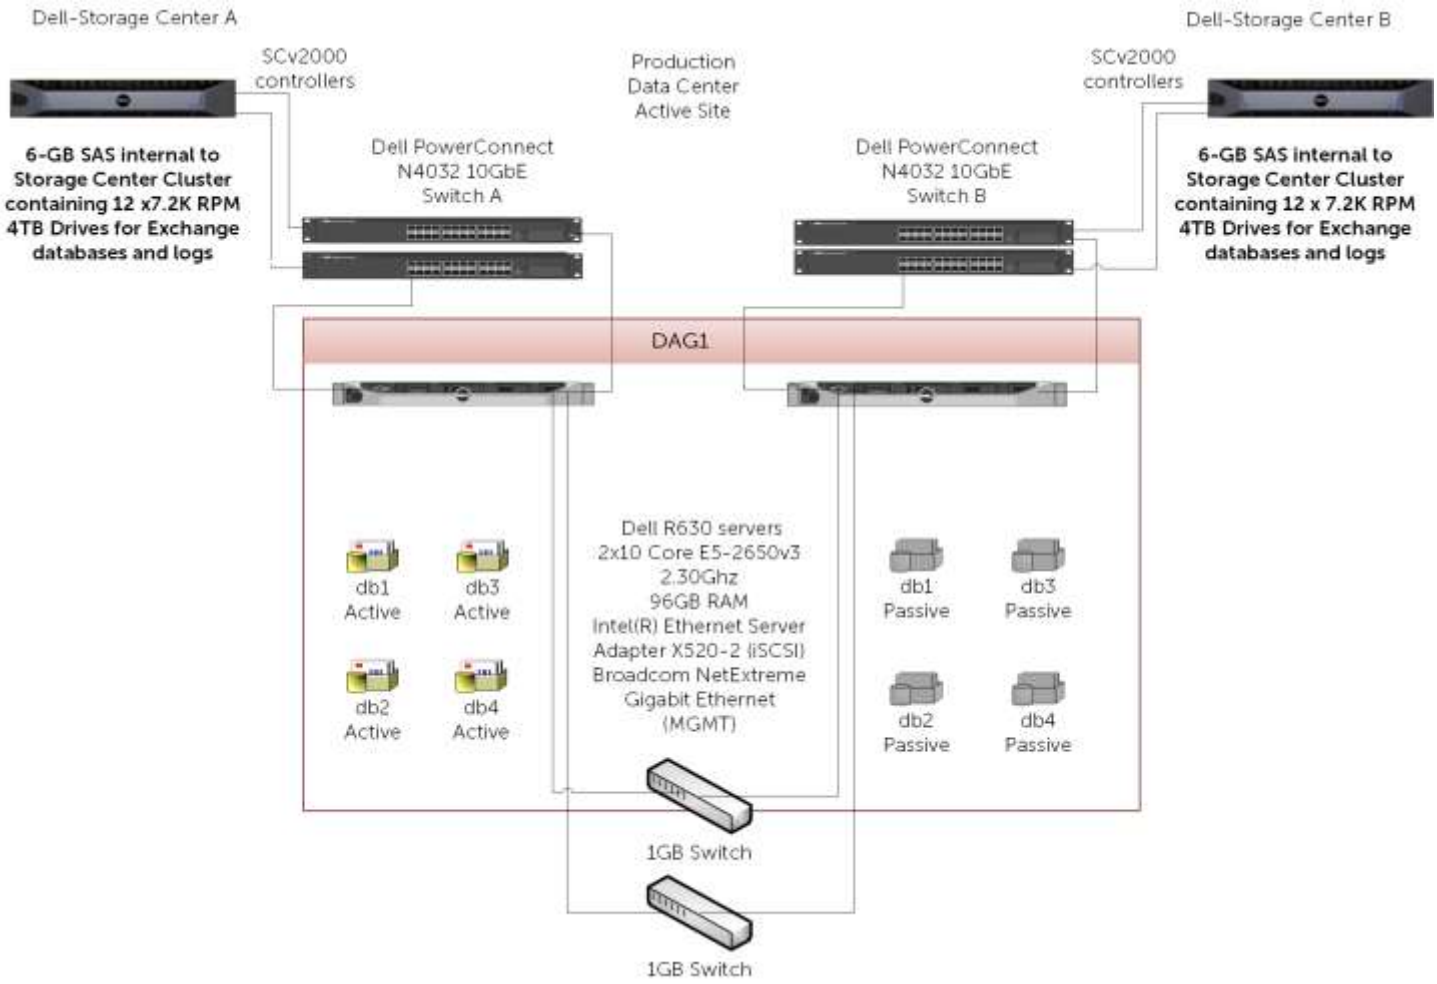

Figure 2 Tested configuration showing Storage Center A with full user load and Storage Center B offline

The tested configuration is a single Storage Center array (Figure 2) running with the full user load. This is to clearly show a single array can handle the user load in an array failure scenario. Under normal operating conditions, the preferred activation scenario would be to run half of the mailbox databases active on each Storage Center array, while either array could handle the entire workload at any given time.

The ability to handle the entire workload on a single Storage Center array means no I/O performance degradation will occur if an array or any volume(s) were to fail. All mailbox servers would have volumes mapped to both arrays, with one copy of each database on each array.

2 The Dell Storage Center SCv2000 solution

2.1 A modular hardware design

The SCv2000 hardware design supports up to 12 internal 3.5-inch hot-swappable SAS drives. It also includes dual controllers that provide automatic failover in a single chassis. Dell SCv2000 arrays have a flex-port option that provides multi-protocol connectivity to any open-systems server without the need for server side agents; organizations can utilize iSCSI, Fiber Channel or SAS connectivity. Disk enclosures support any external interface and disks based on Solid State, Fiber Channel, and/or Serial ATA.

The new SCv2000 arrays combine the benefits of proven Dell Fluid Data™ architecture with a resilient Dell hardware design to provide efficiency, quality and durability. The SCv2000 array offers future flexibility as an entry-level storage array within the existing SC family of products. This series joins a market space where growth and flexibility can collide, making purchase decisions and product investments a challenge. The SCv2000 array leverages existing investments through data migration to enterprise-level products within the SC product line, serving current needs while focusing on future growth strategy.

As a part of the Storage Center product line, the management of all SC products is streamlined into a single, consistent interface. Optimize IT resources with Enterprise Manager by managing through a single pane of glass — from the entry-level SCv2000 array to the enterprise-level SC4000 and SC8000 arrays. For an enhanced, simplified out-of-the box experience, you can deploy the SCv2000 storage arrays with a simple wizard-based tool, minimizing the need for IT expertise for small businesses or improving data center administration by getting projects up and running quickly.

2.2 Powerful suite of software

Storage Center offers a powerful suite of enterprise capabilities to manage data differently. Building on the Dell Storage Center Dynamic Block Architecture, Storage Center software intelligently optimizes data movement and access at the block-level to maximize utilization, automate tiered storage, simplify replication and speed data recovery.

2.3 Intuitive, unified interface

A centralized management interface streamlines administration and speeds common storage management tasks. The interface features a point-and-click wizard-based setup and management, comprehensive Phone Home capabilities, automatic notifications for user-defined capacity thresholds, and advanced storage consumption and chargeback reporting.

Dell Storage Enterprise Manager further simplifies storage management by providing comprehensive monitoring of all local and remote Storage Center environments.

Enterprise Manager allows you to gain better insight into your Storage Center deployments and reduces planning and configuration time for remote replications.

The 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 that 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 of these factors are beyond the scope for ESRP-Storage. Therefore, the number of mailboxes hosted per server, as part of the tested configuration, may not be viable for some customer deployments.

For more information on identifying and addressing performance bottlenecks in an Exchange system, please refer to the Microsoft Exchange Server 2013 Performance Recommendations, available at https://technet.microsoft.com/en-us/library/dn879084%28v=exchg.150%29.aspx.

2.4 Targeted customer profile

This solution is targeted for a medium-sized organization. Capacity can be dynamically scaled from 1TB to over a Petabyte. This provides excellent growth potential with no downtime required for upgrades.

- A Storage Center solution can be sized for any size organization
- An Unlimited number of hosts can be attached using an iSCSI connection
- User IO profile (.07 IOPS per user, .084 tested, giving 20% headroom)
- User mailbox size (2 GB quota)
- The Backup strategy: VSS backup using SAN based snapshots; use Mailbox Resiliency as the primary data protection mechanism
- Using SAN based snapshots and boot from SAN, a complete server can be restored in minutes.
- The tested RAID type was RAID 6 for database volumes and log volumes, while a mix of RAID 10, RAID 5, and RAID 6 can be blended, with fully automated tiered storage providing the most efficient and best performing storage where needed.

2.5 Volume sizing

The volume size tested was just large enough to support the database size. Volumes on Dell SCv2000 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 Storage Dynamic Capacity and hot upgrades additional disk capacity can be added as needed. If more spindles are required to accommodate growth, they can simply be cabled and added to the disk pool to grow volume space. Since volumes are not tied to spindle boundaries, adding spindles increases performance and capacity as the system grows.

3 Tested deployment

The following tables summarize the testing environment.

3.1 Simulated Exchange configuration

Table 1 Simulated Exchange configuration

Configuration Item	Detail
Number of Exchange mailboxes simulated	2,500
Number of Database Availability Groups (DAGs)	1
Number of servers/DAG	2
Number of active mailboxes/server	1250
Number of databases/host	4
Number of copies/database	2
Number of mailboxes/database	625
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	.07 (.084 tested)
Database/Log LUN size	2 TB
Total database size for performance testing per Storage Center	8 TB
% storage capacity used by Exchange database**	25%

** **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 blank data, will consume no disk space.

3.2 Primary storage hardware

Table 2 Primary storage hardware

Configured hardware	Detail
Storage Connectivity (Fiber Channel, SAS, SATA, iSCSI)	iSCSI
Storage model and OS/firmware revision	Dell Storage Center (SCv2000) v6.6 http://www.windowsservercatalog.com/item.aspx?idItem=f2b1c4fb-ee97-7d0d-b671-bf5a75f63c2d&bCatID=1338
Storage cache	16 GB

Configured hardware	Detail
Number of storage controllers	2
Number of storage ports	2 active iSCSI ports per controller
Maximum bandwidth of storage connectivity to host	20 Gb/sec (2x10Gb GB HBA)
Switch type/model/firmware revision	Dell PowerConnect N4032 10GbE Switches Firmware version 6.1.2.4
HBA model and firmware	Intel(R) Ethernet 10G 4P X520/I350 Jumbo frames enabled
Number of HBA's/host	1 Quad-port Intel(R) Ethernet 10G 4P X520/I350
Host server type	2x10 Core E5-2650 v3 2.30Ghz 96GB RAM
Total number of disks tested in solution	11 Active for DB and logs + 1 hot spare = 12 total spindles
Maximum number of spindles can be hosted in the storage	12 drive bay + dual controllers in a 2U chassis Scalable to 168 drives (504TB) via modular expansion enclosures

3.3 Primary storage software

Table 3 Primary storage software

Configuration	Detail
HBA driver	Intel(R) Ethernet 10G 4P X520/I350 Driver 3.4.47.2
HBA Queue Depth Setting	65535
Multi-Pathing	Microsoft Windows Server 2012 R2 MPI/O Round-Robin(In-Box DSM)
Host OS	Windows Server 2012 R2 Datacenter
ESE.dll file version	15.00.0995.021
Replication solution name/version	Microsoft Exchange Server 2013 DAG replication

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

Table 4 Primary storage disk configuration

Configuration	Detail
Disk type, speed and firmware revision	SAS 7k 4TB, XRC0
Raw capacity per disk (GB)	3.64 TB
Number of physical disks in test	11
Total raw storage capacity (GB)	40.0 TB
Raid level	RAID-6
Total formatted capacity	32.0 TB
Storage capacity utilization	80%
Database capacity utilization	20%

4 Best practices

Exchange Server 2013 has changed dramatically from previous versions. For a list of what has changed see the following: http://technet.microsoft.com/en-us/library/jj150540(v=exchg.150).aspx

The best practices have also changed based on the changes in behavior in Exchange 2013. Significant I/O reduction in Exchange 2013 has made it preferable to utilize RAID 6 volumes for both database and logs. This provides overall storage savings due to the smaller capacity overhead versus RAID 10.

Because processor performance has increased dramatically, and servers support much larger memory models, sizing requirements for servers have changed to reflect this. For server sizing, refer to the Microsoft Exchange Server Role Calculator.

For general sizing and requirements, visit: http://technet.microsoft.com/en-us/library/aa996719.aspx

One Microsoft best practice calls for transaction logs and databases to be separated from each other and dedicated to their own set of spindles. Dell SCv2000 arrays address this by virtualizing at the disk level within Storage Center and accelerating data access by spreading read/write operations across all disk drives in the SAN so that multiple requests are processed in parallel. Dell SCv2000 virtualization allows the creation of high performance, highly efficient virtual volumes in seconds without allocating drives to specific servers, without complicated capacity planning and without manual performance tuning. By managing disk drives as a single resource, Dell SCv2000 arrays provide increased storage performance, availability and utilization.

Dell SCv2000 storage virtualization is optimized to take advantage of all available spindles as part of a single disk folder, but is flexible enough to be configured to allow storage configurations where specific spindles are dedicated to a particular volume.

Another best practice in past versions of Exchange Server has been to align Exchange I/O with disk page boundaries. With Windows Server 2012, this is no longer required, because Windows 2012 automatically aligns to a 1024K page boundary.

The volume where transaction logs are stored is critical to a well performing Exchange environment. Since all transactions are first written to a transaction log before being committed to the information store database, it is important that this volume have the lowest possible write latency. Exchange 2013 no longer requires log files to be stored on a volume separate from the database volumes when configured in a Database Availability Group (DAG); The Dell Storage Center can be flexibly designed for separate disk folders or as a single disk folder configuration.

For issues related to performance and server health, see the Microsoft TechNet article: http://technet.microsoft.com/en-us/library/jj150551(v=exchg.150).aspx

For more information on Exchange best practices when implemented with Dell Storage Center, visit the Dell TechCenter Exchange technical content collateral page: http://en.community.dell.com/techcenter/storage/w/wiki/5018.compellent-technical-content#Exchange.

4.1 Core storage

Characteristically, Storage Center volumes do not need disk sector alignment to perform properly. Dell Storage Center virtualizes all disk reads and writes, and applies them across system managed data pages, so that the disk I/O is isolated from sector boundaries. The page to sector alignment for all volumes and data pages is handled automatically by the system.

- The Dell Storage Center method of I/O and disk capacity aggregation provides maximum I/O to all hosted applications.

  IOPS for the assigned drives can be applied to all applications hosted on a Storage Center. If IOPS need to be dedicated to an application, such as Exchange, a dedicated disk pool can be created for each I/O type, such as database or log files. As Exchange 2013 I/O is mostly sequential, using a smaller number of database files will greatly improve the performance. This is because a higher number of sequential streams make it appear more random. Minimizing the number of file streams while meeting business requirements will provide a more responsive solution. Isolating the log files can also provide a performance benefit in an I/O constrained system. Dell SC Series Dynamic Storage makes it possible to start a small system with all volumes sharing spindles. The volumes are dynamic and can be moved to dedicated spindles as load increases.

- Dell Storage Center is a true thin provisioned system.

  Volumes only consume space when and where data is written. Create volume sizes to reflect the maximum size they will achieve. The volumes will only consume the space used by data, so the storage can be sized to host the actual storage requirement, rather than the volume sizes allocated. This allows the volumes to be sized properly to meet growth while requiring the minimum number of disks to meet the storage and IOPS requirement.

- Fluid Data architecture uses an IOPS and storage aggregation model.

  The IOPS and storage capacity of all available disks are accessible to the entire disk pool providing a huge performance boost to all applications and all LUNs. The combined I/O performance of all the spindles applies to all of the configured storage. If dedicated spindles are desired, a disk pool can be created that will dedicate those spindles to the LUNs created in that pool. All disks in a disk pool have multiple RAID types applied to them because the RAID pools are virtualized on the disks. For example, a write action delivered in RAID 10 would be mirrored at the block level across a pair of disks. In essence, each write could hit a different pair of disks, dramatically improving performance. Another write on a RAID 5 block would have the blocks striped across all the disks available to the pool. In this method, a disk pool balances the I/O across all the available spindles.

- Latency and I/O load can be measured in real time, or logged historically for reporting purposes.

  This means if a volume is performing poorly, its I/O can be reported over time, and compared to I/O load on the server, for any length of time needing to be stored. If reporting on the last month of I/O history, a report can be generated showing the I/O graphically or as a summary chart. This provides the ability to trend and determine when I/O performance changed. Volumes can also be

summarized as a group, to determine if I/O load is shifting, increasing, or disk performance is changing. Reporting can be done at any level, including at the disk device level. This allows reporting on the latency at the Server, LUN, or disk level to provide more accurate performance monitoring and diagnostics.

- Fluid Data also allows disks to be added to a pool to increase performance dynamically.

  This allows for accurate sizing on day one and disks to be added as performance requirements increase. If after one year I/O requirements double, additional disks could simply be added (without any downtime), and RAID stripes rebalanced.

- The most common cause of performance issues is low spindle count. To achieve a given I/O level requires a spindle count equal to or greater than the IOP target. If the I/O load exceeds the capabilities of the spindles, poor performance will result. Dell along with a business partner, will work with customers to determine the correct spindle count. As I/O load grows, the spindle count must increase to maintain performance. Using Dell Storage Enterprise Manager, current I/O loads can be tracked, and thresholds can be set for alerting, to warn of I/O usage approaching or exceeding acceptable performance levels. Because I/O patterns can be very diverse, creating a baseline and using historical reporting is a key strategy in planning for and managing growth. With an accurate growth plan, a disk can be added before it is needed, and performance as well as capacity can be increased with no down time.

4.2 Backup Strategy

The Dell Storage Center has an integrated snapshot facility that provides basic volume based snapshots. In order to provide VSS integration with a graphical management interface, the Dell Replay Manager option should be implemented. This provides a full interface for scheduling database backups. Using Replay Manger, Exchange Servers can be restored in minutes to any available restore point. It also provides detailed reporting on snapshots. Because Dell Storage Center has the ability to manage thousands of snapshots, a fine-grained backup strategy can be defined to reduce the need to rely on tapes for historical data recovery. Combined with a lagged database copy, data can be recovered very quickly with minimal administrative effort.

Since Dell Replays do not require page pre-allocation or disk allocation, disk space requirements are much smaller for snapshots. Backup verification can also be passed to a secondary server to isolate the impact of backups on the production Exchange environment. By automating the creation and verification process using a secondary server, more frequent database backups and more frequent database scans can be implemented to reduce exposure.

Replay restore points can also be replicated and tested in a remote environment without breaking replication. This allows Disaster Recovery testing of a production restore point without pausing replication, reducing exposure even further.

4.3 Additional Information

For more information on Dell Storage Center and other Dell 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 that are generated by ESRP testing framework are shown in the Appendices later in this whitepaper.

5.1 Reliability

A number of tests in the framework were run to check reliability for 24 hours. The goal is to verify the storage can handle high IO load for a long period of time. Both log and database files will be analyzed for integrity after the stress test to ensure no database/log corruption.

The following list provides an overview:

- No relevant errors were reported in the event log for the storage reliability test.
- No errors were reported by 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 I/O for two hours. The test is to show how long it takes for the storage to respond to an I/O under load. The data below is the sum of all of the logical disk I/O's and average of all the logical disks I/O latency in the two hours test duration. Each server is listed separately and the aggregate numbers across all servers is listed as well.

5.2.1 Individual Server Metrics:

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

Table 5 Sum of I/O's and average latency

Database I/O	
Database Disks Transfers/sec	254.33
Database Disks Reads/sec	175.713
Database Disks Writes/sec	78.62
Average Database Disk Read Latency (ms)	17.023
Average Database Disk Write Latency (ms)	2.354
Transaction Log I/O	
Log Disks Writes/sec	19.011
Average Log Disk Write Latency (ms)	1.614

5.3 Database Backup/Recovery performance

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

5.3.1 Database read-only performance

The test is to measure the maximum rate at which databases could be backed up via VSS. The following table shows the average rate for a single database file.

Table 6 Maximum backup rate

Performance item	Detail
MB read/sec per database	118.96
MB read/sec total per server	475.84

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 1 MB in size.

Table 7 Average rate for 500 log files played in a single database

Performance item	Detail
Average time to play one Log file (sec)	4.368

6 Conclusion

The testing shows the scalability and performance of the Dell Storage Center SCv2000 arrays.

This document is developed by storage solution providers, and reviewed by the Microsoft Exchange Product team. The test results and data presented in this document are based on the tests introduced in the ESRP v4.0 test framework. Customers should not quote the data directly for a specific 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; the tests are not designed to obtain the maximum throughput for a given solution. Rather, it is focused on producing recommendations from vendors for the Exchange application. So the data presented in this document should not be used for direct comparisons among the solutions.

7 Additional resources

Web sites that provide additional papers and information for Microsoft and Dell products and solutions are:

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

- Dell Storage Website:
  http://www.dellstorage.com/

- Dell TechCenter:
  http://en.community.dell.com/techcenter/storage/

A Stress testing

Stress Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	R7U36R630DC
Test Description	Raid 6 7 threads
Test Start Time	6/13/2015 12:15:17 PM
Test End Time	6/14/2015 12:20:57 PM
Collection Start Time	6/13/2015 12:20:42 PM
Collection End Time	6/14/2015 12:20:32 PM
Jetstress Version	15.00.0995.000
ESE Version	15.00.0995.021
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Exchange Jetstress\Stress_2015_6_13_12_15_26.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	252.932
Target Transactional I/O per Second	210
Initial Database Size (bytes)	5382382551040
Final Database Size (bytes)	5389680640000
Database Files (Count)	4

Jetstress System Parameters

Thread Count	7
Minimum Database Cache	128.0 MB
Maximum Database Cache	1024.0 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

Instance3228.1 Log path: E:\log1
Database: E:\db1\Jetstress001001.edb

Instance3228.2 Log path: F:\log2
Database: F:\db2\Jetstress002001.edb

Instance3228.3 Log path: G:\log3
Database: G:\db3\Jetstress003001.edb

Instance3228.4 Log path: H:\log4
Database: H:\db4\Jetstress004001.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
Instance3228.1	17.921	2.340	43.582	19.679	33005.436	34740.906	0.000	1.586	0.000	4.614	0.000	20455.645
Instance3228.2	16.454	2.349	43.581	19.678	33016.484	34741.435	0.000	1.594	0.000	4.611	0.000	20462.277
Instance3228.3	17.898	2.337	43.525	19.589	33005.418	34745.236	0.000	1.598	0.000	4.601	0.000	20455.101
Instance3228.4	16.854	2.260	43.588	19.712	33015.961	34749.102	0.000	1.586	0.000	4.624	0.000	20503.661

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance3228.1	8.893	261822.117
Instance3228.2	8.958	261818.740
Instance3228.3	8.880	261817.890
Instance3228.4	8.874	261791.533

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance3228.1	0.402	157064.644
Instance3228.2	0.401	157063.002
Instance3228.3	0.400	156352.505
Instance3228.4	0.403	157447.444

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
Instance3228.1	17.921	2.340	52.475	19.679	71784.912	34740.906	10.517	1.586	0.402	4.614	157064.644	20455.645
Instance3228.2	16.454	2.349	52.538	19.678	72027.629	34741.435	10.762	1.594	0.401	4.611	157063.002	20462.277
Instance3228.3	17.898	2.337	52.405	19.589	71776.917	34745.236	11.461	1.598	0.400	4.601	156352.505	20455.101
Instance3228.4	16.854	2.260	52.461	19.712	71712.945	34749.102	10.561	1.586	0.403	4.624	157447.444	20503.661

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.222	0.000	11.833
Available MBytes	94722.265	94676.000	94839.000
Free System Page Table Entries	16366686.482	16366264.000	16366965.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	583734770.899	583184384.000	584073216.000
Pool Paged Bytes	112032377.472	111566848.000	117424128.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

A.1 Test log

Test Log

```
6/13/2015 12:15:17 PM -- Preparing for testing ...
6/13/2015 12:15:21 PM -- Attaching databases ...
6/13/2015 12:15:21 PM -- Preparations for testing are complete.
6/13/2015 12:15:22 PM -- Starting transaction dispatch ..
6/13/2015 12:15:22 PM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)
6/13/2015 12:15:22 PM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)
6/13/2015 12:15:26 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
6/13/2015 12:15:26 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
6/13/2015 12:15:27 PM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
6/13/2015 12:15:27 PM -- Performance logging started (interval: 15000 ms).
6/13/2015 12:15:27 PM -- Attaining prerequisites:
6/13/2015 12:20:42 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 967364600.0 (lower bound: 966367600.0, upper bound: none)
6/14/2015 12:20:42 PM -- Performance logging has ended.
6/14/2015 12:20:42 PM -- JetInterop batch transaction stats: 130303, 130302, 130302 and 130302.
6/14/2015 12:20:42 PM -- Dispatching transactions ends.
6/14/2015 12:20:43 PM -- Shutting down databases ...
6/14/2015 12:20:57 PM -- Instance3228.1 (complete), Instance3228.2 (complete), Instance3228.3 (complete) and Instance3228.4 (complete)
6/14/2015 12:20:57 PM -- C:\Exchange Jetstress\Stress_2015_6_13_12_15_26.blg has 5765 samples.
6/14/2015 12:20:57 PM -- Creating test report ...
6/14/2015 12:21:13 PM -- Instance3228.1 has 17.9 for I/O Database Reads Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.1 has 1.6 for I/O Log Writes Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.1 has 1.6 for I/O Log Reads Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.2 has 16.5 for I/O Database Reads Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.2 has 1.6 for I/O Log Writes Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.2 has 1.6 for I/O Log Reads Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.3 has 17.9 for I/O Database Reads Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.3 has 1.6 for I/O Log Writes Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.3 has 1.6 for I/O Log Reads Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.4 has 16.9 for I/O Database Reads Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.4 has 1.6 for I/O Log Writes Average Latency.
6/14/2015 12:21:13 PM -- Instance3228.4 has 1.6 for I/O Log Reads Average Latency.
6/14/2015 12:21:13 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
6/14/2015 12:21:13 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
6/14/2015 12:21:13 PM -- C:\Exchange Jetstress\Stress_2015_6_13_12_15_26.xml has 5744 samples queried.
```

B Performance testing

Performance Test Result Report

Test Summary

Overall Test Result	**Pass**
Machine Name	R7U36R630DC
Test Description	Raid 6 7 threads
Test Start Time	6/6/2015 6:22:04 PM
Test End Time	6/6/2015 8:27:45 PM
Collection Start Time	6/6/2015 6:27:29 PM
Collection End Time	6/6/2015 8:27:15 PM
Jetstress Version	15.00.0995.000
ESE Version	15.00.0995.021
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Exchange Jetstress\Performance_2015_6_6_18_22_13.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	254.332
Target Transactional I/O per Second	210
Initial Database Size (bytes)	5374044274688
Final Database Size (bytes)	5374690197504
Database Files (Count)	4

Jetstress System Parameters

Thread Count	7
Minimum Database Cache	128.0 MB
Maximum Database Cache	1024.0 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

Instance1468.1 Log path: E:\log1
Database: E:\db1\Jetstress001001.edb

Instance1468.2 Log path: F:\log2
Database: F:\db2\Jetstress002001.edb

Instance1468.3 Log path: G:\log3
Database: G:\db3\Jetstress003001.edb

Instance1468.4 Log path: H:\log4
Database: H:\db4\Jetstress004001.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
Instance1468.1	17.697	2.367	43.605	19.283	32992.570	35634.006	0.000	1.672	0.000	4.710	0.000	20521.422
Instance1468.2	15.979	2.428	44.074	19.766	33007.493	35597.787	0.000	1.557	0.000	4.743	0.000	20505.445
Instance1468.3	17.700	2.317	44.104	19.876	33003.627	35659.209	0.000	1.610	0.000	4.808	0.000	20409.809
Instance1468.4	16.717	2.304	43.930	19.695	33026.032	35600.048	0.000	1.616	0.000	4.750	0.000	20543.262

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1468.1	8.842	261851.070
Instance1468.2	9.107	261859.501
Instance1468.3	8.854	261733.585
Instance1468.4	8.915	261814.454

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance1468.1	0.411	159650.455
Instance1468.2	0.415	160982.702
Instance1468.3	0.417	162676.046
Instance1468.4	0.415	160691.831

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
Instance1468.1	17.697	2.367	52.447	19.283	71576.245	35634.006	10.203	1.672	0.411	4.710	159650.455	20521.422
Instance1468.2	15.979	2.428	53.181	19.766	72197.238	35597.787	11.315	1.557	0.415	4.743	160982.702	20505.445
Instance1468.3	17.700	2.317	52.958	19.876	71243.094	35659.209	13.299	1.610	0.417	4.808	162676.046	20409.809
Instance1468.4	16.717	2.304	52.845	19.695	71622.192	35600.048	11.040	1.616	0.415	4.750	160691.831	20543.262

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.163	0.018	0.307
Available MBytes	94625.958	94616.000	94719.000
Free System Page Table Entries	16365823.061	16365508.000	16366064.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	596388254.731	596295680.000	596488192.000
Pool Paged Bytes	121510165.912	121339904.000	121602048.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

B.1 Test log

Test Log

```
6/6/2015 6:22:04 PM -- Preparing for testing ...
6/6/2015 6:22:08 PM -- Attaching databases ...
6/6/2015 6:22:08 PM -- Preparations for testing are complete.
6/6/2015 6:22:08 PM -- Starting transaction dispatch ..
6/6/2015 6:22:08 PM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)
6/6/2015 6:22:08 PM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)
6/6/2015 6:22:13 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
6/6/2015 6:22:13 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
6/6/2015 6:22:14 PM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
6/6/2015 6:22:14 PM -- Performance logging started (interval: 15000 ms).
6/6/2015 6:22:14 PM -- Attaining prerequisites:
6/6/2015 6:27:29 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 969769000.0 (lower bound: 966367600.0, upper bound: none)
6/6/2015 8:27:30 PM -- Performance logging has ended.
6/6/2015 8:27:30 PM -- JetInterop batch transaction stats: 11669, 11669, 11668 and 11668.
6/6/2015 8:27:30 PM -- Dispatching transactions ends.
6/6/2015 8:27:30 PM -- Shutting down databases ...
6/6/2015 8:27:45 PM -- Instance1468.1 (complete), Instance1468.2 (complete), Instance1468.3 (complete) and Instance1468.4 (complete)
6/6/2015 8:27:45 PM -- C:\Exchange Jetstress\Performance_2015_6_6_18_22_13.blg has 499 samples.
6/6/2015 8:27:45 PM -- Creating test report ...
6/6/2015 8:27:47 PM -- Instance1468.1 has 17.7 for I/O Database Reads Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.1 has 1.7 for I/O Log Writes Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.1 has 1.7 for I/O Log Reads Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.2 has 16.0 for I/O Database Reads Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.2 has 1.6 for I/O Log Writes Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.2 has 1.6 for I/O Log Reads Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.3 has 17.7 for I/O Database Reads Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.3 has 1.6 for I/O Log Writes Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.3 has 1.6 for I/O Log Reads Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.4 has 16.7 for I/O Database Reads Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.4 has 1.6 for I/O Log Writes Average Latency.
6/6/2015 8:27:47 PM -- Instance1468.4 has 1.6 for I/O Log Reads Average Latency.
6/6/2015 8:27:48 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
6/6/2015 8:27:48 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
6/6/2015 8:27:48 PM -- C:\Exchange Jetstress\Performance_2015_6_6_18_22_13.xml has 478 samples queried.
```

C Backup testing

Database backup Test Result Report

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance3700.1	1281656.03	02:57:36	120.26
Instance3700.2	1281640.03	03:05:18	115.27
Instance3700.3	1281640.03	02:50:33	125.24
Instance3700.4	1281640.03	03:05:38	115.06
Avg			118.96
Sum			475.84

Jetstress System Parameters

Thread Count	7
Minimum Database Cache	128.0 MB
Maximum Database Cache	1024.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%

Database Configuration

Instance3700.1 Log path: E:\log1
Database: E:\db1\Jetstress001001.edb

Instance3700.2 Log path: F:\log2
Database: F:\db2\Jetstress002001.edb

Instance3700.3 Log path: G:\log3
Database: G:\db3\Jetstress003001.edb

Instance3700.4 Log path: H:\log4
Database: H:\db4\Jetstress004001.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
Instance3700.1	3.188	0.000	481.459	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance3700.2	3.386	0.000	461.233	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance3700.3	3.038	0.000	501.485	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance3700.4	3.399	0.000	460.559	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	2.260	1.624	2.732
Available MBytes	95706.782	95661.000	95725.000
Free System Page Table Entries	16365900.887	16365460.000	16366133.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	598994204.291	598974464.000	599068672.000
Pool Paged Bytes	124955678.361	124805120.000	125104128.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

C.1 Test log

Test Log

6/10/2015 4:02:11 PM -- Preparing for testing ...
6/10/2015 4:02:15 PM -- Attaching databases ...
6/10/2015 4:02:15 PM -- Preparations for testing are complete.
6/10/2015 4:02:20 PM -- Performance logging started (interval: 30000 ms).
6/10/2015 4:02:20 PM -- Backing up databases ...
6/10/2015 7:07:59 PM -- Performance logging has ended.
6/10/2015 7:07:59 PM -- Instance3700.1 (100% processed), Instance3700.2 (100% processed), Instance3700.3 (100% processed) and Instance3700.4 (100% processed)
6/10/2015 7:07:59 PM -- C:\Exchange Jetstress\DatabaseBackup_2015_6_10_16_2_15.blg has 371 samples.
6/10/2015 7:07:59 PM -- Creating test report ...

D Recovery testing

SoftRecovery Test Result Report

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance4640.1	504	2204.7141927
Instance4640.2	501	2183.8900091
Instance4640.3	502	2191.0693017
Instance4640.4	502	2191.3349385
Avg	502	2192.752
Sum	2009	8771.008442

Database Configuration

Instance4640.1 Log path: E:\log1
Database: E:\db1\Jetstress001001.edb

Instance4640.2 Log path: F:\log2
Database: F:\db2\Jetstress002001.edb

Instance4640.3 Log path: G:\log3
Database: G:\db3\Jetstress003001.edb

Instance4640.4 Log path: H:\log4
Database: H:\db4\Jetstress004001.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
Instance4640.1	15.988	0.864	93.545	0.912	38471.887	15570.824	6.456	0.000	1.141	0.000	99297.569	0.000
Instance4640.2	15.605	0.834	92.807	0.915	38645.337	15654.471	6.327	0.000	1.148	0.000	98417.238	0.000
Instance4640.3	15.956	0.925	92.616	0.914	38553.006	15914.154	7.103	0.000	1.143	0.000	99983.512	0.000
Instance4640.4	16.253	0.996	93.036	0.914	38500.703	15701.965	6.830	0.000	1.143	0.000	99191.214	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4640.1	0.000	0.000
Instance4640.2	0.000	0.000
Instance4640.3	0.000	0.000
Instance4640.4	0.000	0.000

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
Instance4640.1	15.988	0.864	93.545	0.912	38471.887	15570.824	6.456	0.000	1.141	0.000	99297.569	0.000
Instance4640.2	15.605	0.834	92.807	0.915	38645.337	15654.471	6.327	0.000	1.148	0.000	98417.238	0.000
Instance4640.3	15.956	0.925	92.616	0.914	38553.006	15914.154	7.103	0.000	1.143	0.000	99983.512	0.000
Instance4640.4	16.253	0.996	93.036	0.914	38500.703	15701.965	6.830	0.000	1.143	0.000	99191.214	0.000

5

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.218	0.000	1.825
Available MBytes	94643.675	94507.000	95698.000
Free System Page Table Entries	16365908.545	16365299.000	16366166.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	599034177.587	598884352.000	599314432.000
Pool Paged Bytes	124724911.368	124567552.000	125329408.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

D.1 Test log

Test Log

```
6/10/2015 11:44:39 AM -- Preparing for testing ...
6/10/2015 11:44:43 AM -- Attaching databases ...
6/10/2015 11:44:43 AM -- Preparations for testing are complete.
6/10/2015 11:44:43 AM -- Starting transaction dispatch ..
6/10/2015 11:44:43 AM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)
6/10/2015 11:44:43 AM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)
6/10/2015 11:44:47 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
6/10/2015 11:44:47 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
6/10/2015 11:44:49 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
6/10/2015 11:44:49 AM -- Performance logging started (interval: 15000 ms).
6/10/2015 11:44:49 AM -- Generating log files ...
6/10/2015 2:41:06 PM -- E:\log1 (100.8% generated), F:\log2 (100.2% generated), G:\log3 (100.4% generated) and H:\log4 (100.4% generated)
6/10/2015 2:41:06 PM -- Performance logging has ended.
6/10/2015 2:41:06 PM -- JetInterop batch transaction stats: 16912, 16912, 16912 and 16912.
6/10/2015 2:41:06 PM -- Dispatching transactions ends.
6/10/2015 2:41:06 PM -- Shutting down databases ...
6/10/2015 2:41:17 PM -- Instance4640.1 (complete), Instance4640.2 (complete), Instance4640.3 (complete) and Instance4640.4 (complete)
6/10/2015 2:41:17 PM -- C:\Exchange Jetstress\Performance_2015_6_10_11_44_47.blg has 704 samples.
6/10/2015 2:41:17 PM -- Creating test report ...
6/10/2015 2:41:20 PM -- Instance4640.1 has 19.3 for I/O Database Reads Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.1 has 1.6 for I/O Log Writes Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.1 has 1.6 for I/O Log Reads Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.2 has 18.1 for I/O Database Reads Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.2 has 1.6 for I/O Log Writes Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.2 has 1.6 for I/O Log Reads Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.3 has 19.3 for I/O Database Reads Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.3 has 1.6 for I/O Log Writes Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.3 has 1.6 for I/O Log Reads Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.4 has 18.3 for I/O Database Reads Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.4 has 1.6 for I/O Log Writes Average Latency.
6/10/2015 2:41:20 PM -- Instance4640.4 has 1.6 for I/O Log Reads Average Latency.
6/10/2015 2:41:20 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
6/10/2015 2:41:20 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
6/10/2015 2:41:20 PM -- C:\Exchange Jetstress\Performance_2015_6_10_11_44_47.xml has 703 samples queried.
6/10/2015 2:41:20 PM -- C:\Exchange Jetstress\Performance_2015_6_10_11_44_47.html was saved.
6/10/2015 2:41:21 PM -- Performance logging started (interval: 2000 ms).
6/10/2015 2:41:21 PM -- Recovering databases ...
6/10/2015 3:18:06 PM -- Performance logging has ended.
6/10/2015 3:18:06 PM -- Instance4640.1 (2204.7141927), Instance4640.2 (2183.8900091), Instance4640.3 (2191.0693017) and Instance4640.4
(2191.3349385)
6/10/2015 3:18:06 PM -- C:\Exchange Jetstress\SoftRecovery_2015_6_10_14_41_20.blg has 1089 samples.
6/10/2015 3:18:06 PM -- Creating test report ...
```