Dell Storage Center 6.6 SCv2000 SAS Front-end Arrays and 2,500 Mailbox Exchange 2013 Resiliency Storage Solution

Microsoft ESRP 4.0

Dell Storage Engineering
October 2015

A Dell Technical White Paper

Revisions

Date	Description
October 2015	Initial release

Acknowledgements

Authored by: Mark Boeser, Senior Development Engineer

Table of contents

Revisions 2

Acknowledgements 2

1 Introduction 5

1.1 How to use this document 5

1.2 Simulated environment 5

1.3 Solution description 6

2 Dell Storage SCv2000 solution 8

2.1 A modular hardware design 8

2.2 Powerful suite of software 8

2.3 Intuitive, unified interface 8

2.4 Targeted customer profile 9

2.5 Volume sizing 9

3 Tested deployment 10

3.1 Simulated Exchange configuration 10

3.2 Primary storage hardware 10

3.3 Primary storage software 11

3.4 Primary storage disk configuration (mailbox store/log disks) 12

4 Best practices 13

4.1 Core storage 14

4.2 Backup strategy 15

5 Test results summary 16

5.1 Reliability 16

5.2 Storage performance results 16

5.2.1 Server metrics 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 18

A Stress testing 19

A.1 Stress test report 19

A.2 Test log 21

B Performance testing 22
B.1 Performance test report 22
B.2 Test log 24
C Backup testing 25
C.1 Database backup test report 25
C.2 Test log 26
D Recovery testing 27
D.1 SoftRecovery test report 27
D.2 Test log 29
E Additional resources 30
E.1 Technical support and resources 30
E.2 Related documentation 30

1 Introduction

This document provides information on a Dell™ Storage SCv2000 Series 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, refer to the Microsoft TechNet page, Exchange Solution Reviewed Program (ESRP) – Storage.

1.1 How to use this document

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. Therefore, the data presented in this document should not be used for direct comparisons among the solutions.

1.2 Simulated environment

The solution presented in this document is designed to simulate a small-to-medium 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 2 GB per user. Each server has four databases, with one local copy and a 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.

1.3 Solution description

Testing was performed on a Dell Storage SCv2000 Series array with Storage Center v6.6 and a redundant controller pair with redundant front-end and back-end connections. The front-end connections are SAS-based with two SAS ports per server and two SAS ports per storage controller. One 12-bay, 3.5-inch built-in drive enclosure was utilized with each array.

The disk connectivity is SAS 6Gbps and the disk drives are SAS 7.2K 2TB. The spindle count includes 11 disks and 1 spare for the database and logs on a dedicated disk pool on each Storage Center. Because this is a redundant solution, the database and logs are stored together on the same volumes, all of which are RAID 6. For information about compatibility, visit the page, Windows Server Catalog.

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.

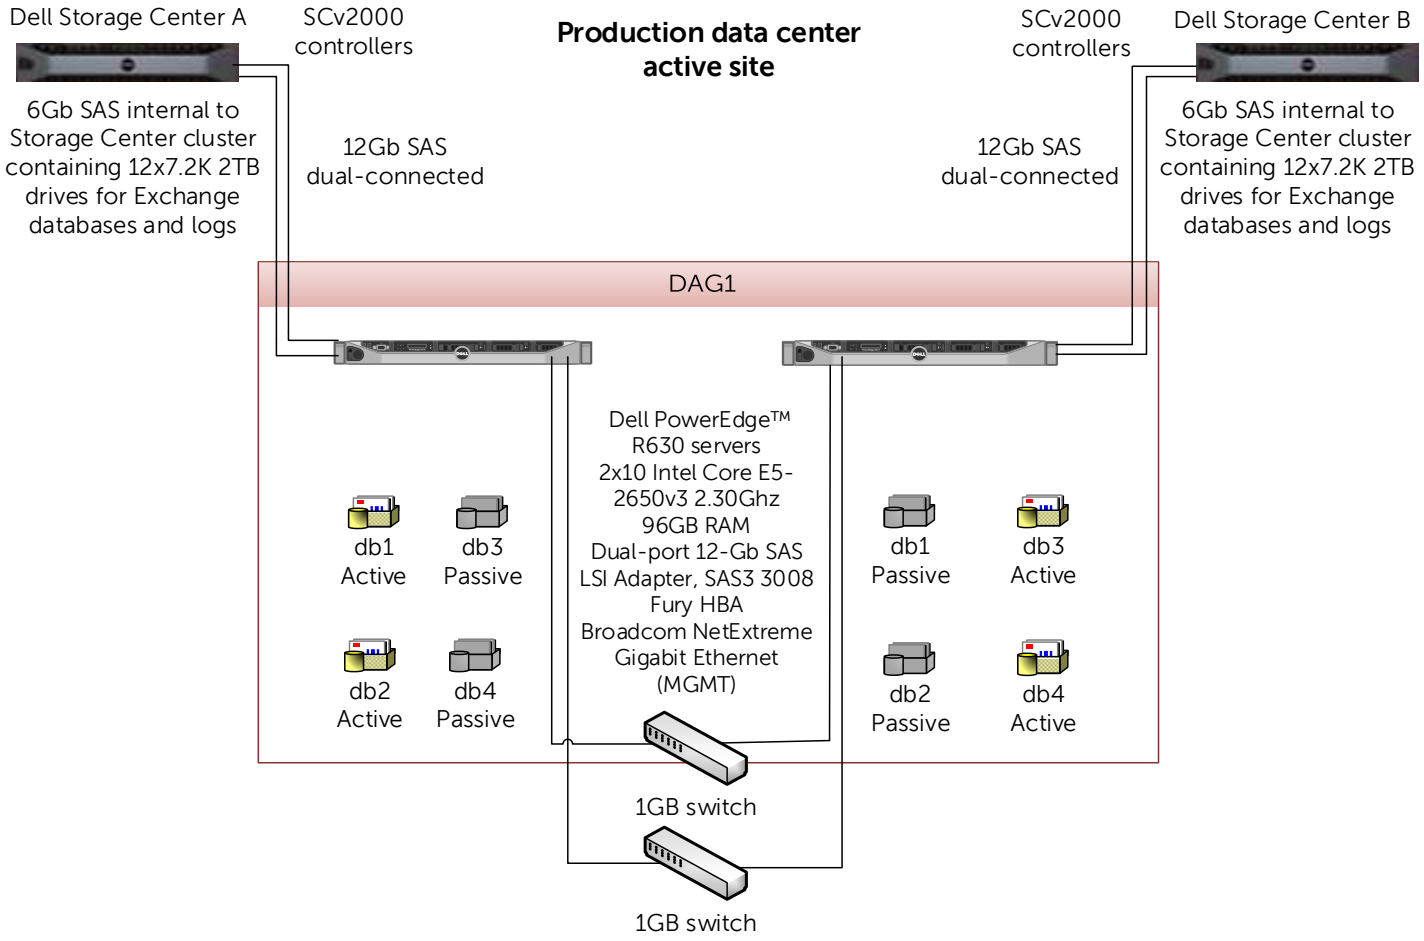

Figure 1 Highly available data center design

The tested configuration includes a single SC Series 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 SC Series array, while either array could handle the entire workload at any given time.

The ability to handle the entire workload on a single SC Series 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.

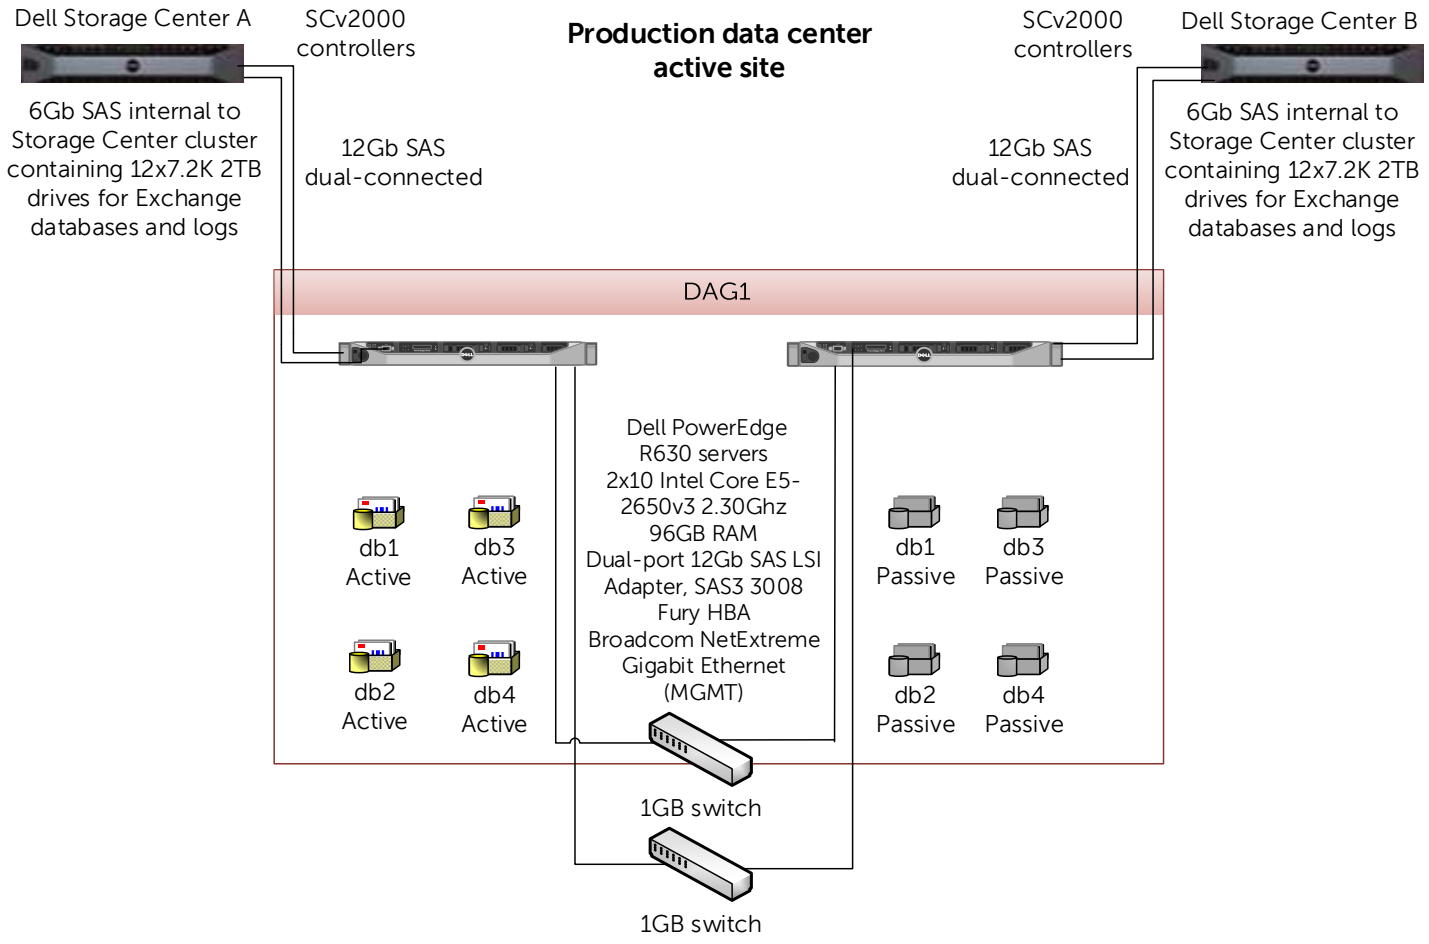

Figure 2 Tested configuration showing Storage Center A with full user load and Storage Center B offline

2 Dell Storage SCv2000 solution

2.1 A modular hardware design

The Dell Storage SCv2000 Series 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. 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, Fibre Channel, or SAS connectivity. Disk enclosures support any external interface and disks based on solid state, Fibre Channel, and/or Serial ATA.

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 Series 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 Series product line, serving current needs while focusing on future growth strategy.

As a part of the SC Series product line, the management of all products is streamlined into a single, consistent interface. Enterprise Manager can help optimize IT resources by offering management through a single pane of glass — from the entry-level SCv2000 array to the enterprise-level SC4000 and SC8000 arrays. For a simplified out-of-the box experience, you can deploy the SCv2000 storage arrays with a wizard-based tool, minimizing the need for IT expertise for small businesses and 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 Dell SupportAssist 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, refer to the page, Microsoft Exchange Server 2013 Performance Recommendations.

2.4 Targeted customer profile

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

- Organization: An SC Series solution can be sized to meet an organization's needs
- User IO profile: .09 IOPS per user, .11 tested, giving ~20% headroom
- User mailbox size: 2 GB quota
- Backup strategy: VSS backup using SAN based snapshots; use Mailbox Resiliency as the primary data protection mechanism
- Restore method: Using SAN-based snapshots and boot from SAN, a complete server can be restored in minutes
- RAID configuration: 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 and tiered RAID 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 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/Os per second, per mailbox (IOPS, include 20% headroom)	.09 (.11 tested)
Database/log LUN size	2 TB
Total database size for performance testing per Storage Center	8 TB
% storage capacity used by Exchange database[1]	31%

[1]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 (Fibre Channel, SAS, SATA, iSCSI)	SAS
Storage model, OS/firmware revision	Dell Storage SCv2000 Series, Storage Center v6.6, Windows Server® (view compatibility)
Storage cache	16 GB
Number of storage controllers	2
Number of storage ports	2 of 4 active SAS ports per controller

Configured hardware	Detail
Maximum bandwidth of storage connectivity to host	24 Gb/sec (2x12Gb SAS HBA)
Switch type, model, firmware revision	N/A
HBA model and firmware	Dell 12Gb SAS HBA (Dell P/N: 405-AAES) (LSI chipset)
Number of HBAs/host	1 Dell 12Gb SAS HBA (Dell P/N: 405-AAES) (LSI chipset)
Host server type	2x10 Intel® 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 (504 TB) using modular expansion enclosures

3.3 Primary storage software

Table 3 Primary storage software

Configuration	Detail
HBA driver	LSI® Adapter, SAS3 3008 Fury -StorPort
HBA queue depth setting	65535
Multipathing	Microsoft Windows Server 2012 R2 MPIO Round-Robin (in-box DSM)
Host OS	Windows Server 2012 R2 Datacenter
ESE.dll file version	15.00.1104.002
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, 2 TB, XRC0
Raw capacity per disk (GB)	1.82 TB
Number of physical disks in test	11
Total raw storage capacity (GB)	20.0 TB
Raid level	RAID 6
Total formatted capacity	16.0 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 Microsoft TechNet page, What's new in Exchange 2013.

The best practices have also evolved 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 the database and logs. This provides overall storage savings due to the smaller capacity overhead compared to 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 the TechNet page, Exchange 2016 system requirements.

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 Storage SCv2000 Series 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. 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, SCv2000 arrays provide increased storage performance, availability, and utilization. In addition, SCv2000 storage virtualization is optimized to take advantage of all available spindles as part of a single disk folder, but is flexible enough 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 TechNet page, Server health and performance.

For more information on Exchange best practices when implemented with Dell SC Series storage, visit the SC Series technical content page.

4.1 Core storage

Characteristically, SC Series 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 SC Series 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 an SC Series system. 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 can 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 SC Series storage is a true thin provisioned system.

  Volumes only consume space when and where data is written. Volume sizes should be created 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 or increasing, or if 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 for disks to be added as performance requirements increase. If I/O requirements double after one year, additional disks can be added (without any downtime), and RAID stripes can be rebalanced.

  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 downtime. 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.

4.2 Backup strategy

SC Series storage 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.

5 Test results summary

This section provides a high-level summary of the ESRP test data. Detailed HTML reports generated by the ESRP testing framework are listed in the appendices.

5.1 Reliability

A number of tests in the framework were run to check reliability for 24 hours. The goal was to verify the storage can handle a 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 or 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/Os 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 Server metrics

Table 5 lists the sum of I/Os across all mailbox databases and the average latency across all databases on a per server basis.

Table 5 Sum of I/Os and average latency

Database I/O	
Database disks transfers/sec	322.088
Database disks reads/sec	233.282
Database disks writes/sec	88.806
Average database disk read latency (ms)	18.723
Average database disk write latency (ms)	1.633
Transaction log i/o	
Log disks writes/sec	23.223
Average log disk write latency (ms)	1.107

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

This test measures the maximum rate at which databases could be backed up using VSS. Table 6 shows the average rate for a single database file.

Table 6 Maximum backup rate

Performance item	Detail
MB read/sec per database	125.51
MB read/sec total per server	502.03

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. Table 7 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)	4.468

6 Conclusion

The testing presented in this paper demonstrates the scalability and performance of Dell Storage SCv2000 Series arrays.

A Stress testing

A.1 Stress test report

Test Summary

Overall Test Result	**Pass**
Machine Name	JS16
Test Description	2500 Mailboxes .10 IOPS Profile 2GB Mailbox size
Test Start Time	8/10/2015 2:22:05 PM
Test End Time	8/11/2015 2:26:47 PM
Collection Start Time	8/10/2015 2:26:45 PM
Collection End Time	8/11/2015 2:26:38 PM
Jetstress Version	15.00.0775.000
ESE Version	15.00.1104.002
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Stress_2015_8_10_14_22_14.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	284.159
Target Transactional I/O per Second	250
Initial Database Size (bytes)	5369715752960
Final Database Size (bytes)	5378196635648
Database Files (Count)	4

Jetstress System Parameters

Thread Count	8
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

Instance4104.1 Log path: C:\DB\DB1
Database: C:\DB\DB1\Jetstress001001.edb

Instance4104.2 Log path: C:\DB\DB2
Database: C:\DB\DB2\Jetstress002001.edb

Instance4104.3 Log path: C:\DB\DB3
Database: C:\DB\DB3\Jetstress003001.edb

Instance4104.4 Log path: C:\DB\DB4
Database: C:\DB\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
Instance4104.1	19.287	1.611	49.116	21.893	32972.050	35885.317	0.000	1.100	0.000	5.359	0.000	20409.968
Instance4104.2	19.017	1.594	49.189	22.027	32977.974	35865.920	0.000	1.083	0.000	5.386	0.000	20364.387
Instance4104.3	18.453	1.591	49.020	21.920	32982.120	35854.210	0.000	1.099	0.000	5.381	0.000	20546.476
Instance4104.4	18.574	1.598	49.067	21.927	32984.161	35873.768	0.000	1.070	0.000	5.386	0.000	20445.067

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4104.1	8.468	261810.988
Instance4104.2	8.449	261839.069
Instance4104.3	8.643	261781.365
Instance4104.4	8.594	261795.560

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance4104.1	0.466	180277.812
Instance4104.2	0.467	181120.277
Instance4104.3	0.471	182687.441
Instance4104.4	0.469	181934.139

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
Instance4104.1	19.287	1.611	57.583	21.893	66622.501	35885.317	11.968	1.100	0.466	5.359	180277.812	20409.968
Instance4104.2	19.017	1.594	57.638	22.027	66524.959	35865.920	10.967	1.083	0.467	5.386	181120.277	20364.387
Instance4104.3	18.453	1.591	57.663	21.920	67277.454	35854.210	11.631	1.099	0.471	5.381	182687.441	20546.476
Instance4104.4	18.574	1.598	57.661	21.927	67085.537	35873.768	12.009	1.070	0.469	5.386	181934.139	20445.067

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.050	0.000	0.413
Available MBytes	126705.860	126652.000	126764.000
Free System Page Table Entries	16460789.699	16459984.000	16461252.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	218633896.175	218525696.000	223809536.000
Pool Paged Bytes	427951590.551	427802624.000	433029120.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

A.2 Test log

Test Log

```
8/10/2015 2:22:05 PM -- Preparing for testing ...
8/10/2015 2:22:10 PM -- Attaching databases ...
8/10/2015 2:22:10 PM -- Preparations for testing are complete.
8/10/2015 2:22:10 PM -- Starting transaction dispatch ..
8/10/2015 2:22:10 PM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)
8/10/2015 2:22:10 PM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)
8/10/2015 2:22:14 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
8/10/2015 2:22:14 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
8/10/2015 2:22:15 PM -- Operation mix: Sessions 8, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
8/10/2015 2:22:15 PM -- Performance logging started (interval: 15000 ms).
8/10/2015 2:22:15 PM -- Attaining prerequisites:
8/10/2015 2:26:45 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 971943900.0 (lower bound: 966367600.0, upper bound: none)
8/11/2015 2:26:46 PM -- Performance logging has ended.
8/11/2015 2:26:46 PM -- JetInterop batch transaction stats: 151941, 151941, 151941 and 151940.
8/11/2015 2:26:46 PM -- Dispatching transactions ends.
8/11/2015 2:26:46 PM -- Shutting down databases ...
8/11/2015 2:26:47 PM -- Instance4104.1 (complete), Instance4104.2 (complete), Instance4104.3 (complete) and Instance4104.4 (complete)
8/11/2015 2:26:47 PM -- C:\Program Files\Exchange Jetstress\Stress_2015_8_10_14_22_14.blg has 5771 samples.
8/11/2015 2:26:47 PM -- Creating test report ...
8/11/2015 2:27:09 PM -- Instance4104.1 has 19.3 for I/O Database Reads Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.1 has 1.1 for I/O Log Writes Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.1 has 1.1 for I/O Log Reads Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.2 has 19.0 for I/O Database Reads Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.2 has 1.1 for I/O Log Writes Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.2 has 1.1 for I/O Log Reads Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.3 has 18.5 for I/O Database Reads Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.3 has 1.1 for I/O Log Writes Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.3 has 1.1 for I/O Log Reads Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.4 has 18.6 for I/O Database Reads Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.4 has 1.1 for I/O Log Writes Average Latency.
8/11/2015 2:27:09 PM -- Instance4104.4 has 1.1 for I/O Log Reads Average Latency.
8/11/2015 2:27:09 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
8/11/2015 2:27:09 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
8/11/2015 2:27:09 PM -- C:\Program Files\Exchange Jetstress\Stress_2015_8_10_14_22_14.xml has 5753 samples queried.
```

B Performance testing

B.1 Performance test report

Test Summary

Overall Test Result	**Pass**
Machine Name	JS16
Test Description	2500 Mailboxes .09 IOPS Profile .11 IOPS Tested 2GB Mailbox size
Test Start Time	8/10/2015 9:01:18 AM
Test End Time	8/10/2015 11:06:04 AM
Collection Start Time	8/10/2015 9:06:02 AM
Collection End Time	8/10/2015 11:05:51 AM
Jetstress Version	15.00.0775.000
ESE Version	15.00.1104.002
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Performance_2015_8_10_9_1_27.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	288.044
Target Transactional I/O per Second	225
Initial Database Size (bytes)	5368944001024
Final Database Size (bytes)	5369715752960
Database Files (Count)	4

Jetstress System Parameters

Thread Count	8
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

Instance4332.1 Log path: C:\DB\DB1
Database: C:\DB\DB1\Jetstress001001.edb

Instance4332.2 Log path: C:\DB\DB2
Database: C:\DB\DB2\Jetstress002001.edb

Instance4332.3 Log path: C:\DB\DB3
Database: C:\DB\DB3\Jetstress003001.edb

Instance4332.4 Log path: C:\DB\DB4
Database: C:\DB\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
Instance4332.1	19.822	1.682	49.837	22.177	32958.001	36271.332	0.000	1.133	0.000	5.757	0.000	20543.545
Instance4332.2	18.480	1.694	49.927	22.452	32992.297	36277.034	0.000	1.078	0.000	5.869	0.000	20461.380
Instance4332.3	17.857	1.589	49.693	22.083	32983.621	36349.144	0.000	1.098	0.000	5.807	0.000	20468.128
Instance4332.4	18.734	1.568	49.781	22.094	32980.005	36307.145	0.000	1.122	0.000	5.790	0.000	20388.470

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4332.1	8.157	261783.750
Instance4332.2	8.589	261813.716
Instance4332.3	8.816	261756.966
Instance4332.4	8.483	261816.029

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance4332.1	0.503	192895.111
Instance4332.2	0.512	198385.531
Instance4332.3	0.507	194597.751
Instance4332.4	0.504	194772.471

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
Instance4332.1	19.822	1.682	57.994	22.177	65143.095	36271.332	13.134	1.133	0.503	5.757	192895.111	20543.545
Instance4332.2	18.480	1.694	58.515	22.452	66578.317	36277.034	12.153	1.078	0.512	5.869	198385.531	20461.380
Instance4332.3	17.857	1.589	58.509	22.083	67453.235	36349.144	12.562	1.098	0.507	5.807	194597.751	20468.128
Instance4332.4	18.734	1.568	58.264	22.094	66297.757	36307.145	13.229	1.122	0.504	5.790	194772.471	20388.470

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.049	0.000	0.299
Available MBytes	126682.470	126640.000	126741.000
Free System Page Table Entries	16460739.148	16460038.000	16461127.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	218655109.077	218607616.000	219013120.000
Pool Paged Bytes	427714290.639	427663360.000	428158976.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

B.2 Test log

Test Log

```
8/10/2015 9:01:18 AM -- Preparing for testing ...
8/10/2015 9:01:22 AM -- Attaching databases ...
8/10/2015 9:01:22 AM -- Preparations for testing are complete.
8/10/2015 9:01:22 AM -- Starting transaction dispatch ..
8/10/2015 9:01:22 AM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)
8/10/2015 9:01:22 AM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)
8/10/2015 9:01:27 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
8/10/2015 9:01:27 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
8/10/2015 9:01:28 AM -- Operation mix: Sessions 8, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
8/10/2015 9:01:28 AM -- Performance logging started (interval: 15000 ms).
8/10/2015 9:01:28 AM -- Attaining prerequisites:
8/10/2015 9:06:02 AM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 971292700.0 (lower bound: 966367600.0, upper bound: none)
8/10/2015 11:06:03 AM -- Performance logging has ended.
8/10/2015 11:06:03 AM -- JetInterop batch transaction stats: 14197, 14197, 14196 and 14196.
8/10/2015 11:06:03 AM -- Dispatching transactions ends.
8/10/2015 11:06:03 AM -- Shutting down databases ...
8/10/2015 11:06:04 AM -- Instance4332.1 (complete), Instance4332.2 (complete), Instance4332.3 (complete) and Instance4332.4 (complete)
8/10/2015 11:06:04 AM -- C:\Program Files\Exchange Jetstress\Performance_2015_8_10_9_1_27.blg has 497 samples.
8/10/2015 11:06:04 AM -- Creating test report ...
8/10/2015 11:06:08 AM -- Instance4332.1 has 19.8 for I/O Database Reads Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.1 has 1.1 for I/O Log Writes Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.1 has 1.1 for I/O Log Reads Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.2 has 18.5 for I/O Database Reads Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.2 has 1.1 for I/O Log Writes Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.2 has 1.1 for I/O Log Reads Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.3 has 17.9 for I/O Database Reads Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.3 has 1.1 for I/O Log Writes Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.3 has 1.1 for I/O Log Reads Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.4 has 18.7 for I/O Database Reads Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.4 has 1.1 for I/O Log Writes Average Latency.
8/10/2015 11:06:08 AM -- Instance4332.4 has 1.1 for I/O Log Reads Average Latency.
8/10/2015 11:06:08 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
8/10/2015 11:06:08 AM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
8/10/2015 11:06:08 AM -- C:\Program Files\Exchange Jetstress\Performance_2015_8_10_9_1_27.xml has 478 samples queried.
```

C Backup testing

C.1 Database backup test report

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance2952.1	1282232.03	03:05:05	115.46
Instance2952.2	1282248.03	02:11:00	163.12
Instance2952.3	1282272.03	03:07:57	113.71
Instance2952.4	1282264.03	03:14:44	109.74
Avg			125.51
Sum			502.03

Jetstress System Parameters

Thread Count	8
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

Instance2952.1 Log path: C:\DB\DB1
Database: C:\DB\DB1\Jetstress001001.edb

Instance2952.2 Log path: C:\DB\DB2
Database: C:\DB\DB2\Jetstress002001.edb

Instance2952.3 Log path: C:\DB\DB3
Database: C:\DB\DB3\Jetstress003001.edb

Instance2952.4 Log path: C:\DB\DB4
Database: C:\DB\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
Instance2952.1	2.451	0.000	461.116	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2952.2	3.032	0.000	652.403	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2952.3	2.476	0.000	454.677	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2952.4	2.786	0.000	438.339	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	0.171	0.071	0.312
Available MBytes	127526.432	127473.000	127544.000
Free System Page Table Entries	16460879.717	16460342.000	16461243.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	218832158.931	218804224.000	218968064.000
Pool Paged Bytes	431354772.072	431230976.000	431730688.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

C.2 Test log

Test Log

```
8/12/2015 7:52:38 AM -- Preparing for testing ...
8/12/2015 7:52:43 AM -- Attaching databases ...
8/12/2015 7:52:43 AM -- Preparations for testing are complete.
8/12/2015 7:52:48 AM -- Performance logging started (interval: 30000 ms).
8/12/2015 7:52:48 AM -- Backing up databases ...
8/12/2015 11:07:32 AM -- Performance logging has ended.
8/12/2015 11:07:32 AM -- Instance2952.1 (100% processed), Instance2952.2 (100% processed), Instance2952.3 (100% processed) and Instance2952.4 (100%
processed)
8/12/2015 11:07:32 AM -- C:\Program Files\Exchange Jetstress\DatabaseBackup_2015_8_12_7_52_43.blg has 389 samples.
8/12/2015 11:07:32 AM -- Creating test report ...
```

D Recovery testing

D.1 SoftRecovery test report

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance4568.1	507	2266.9289832
Instance4568.2	513	2287.3779462
Instance4568.3	501	2236.34039
Instance4568.4	506	2253.788597
Avg	506	2261.109
Sum	2027	9044.4359164

Database Configuration

Instance4568.1 Log path: C:\DB\DB1
Database: C:\DB\DB1\Jetstress001001.edb

Instance4568.2 Log path: C:\DB\DB2
Database: C:\DB\DB2\Jetstress002001.edb

Instance4568.3 Log path: C:\DB\DB3
Database: C:\DB\DB3\Jetstress003001.edb

Instance4568.4 Log path: C:\DB\DB4
Database: C:\DB\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
Instance4568.1	16.134	0.484	93.273	0.893	38388.010	15348.146	6.704	0.000	1.120	0.000	97210.821	0.000
Instance4568.2	17.670	0.468	93.190	0.895	38465.439	15212.681	6.856	0.000	1.122	0.000	96540.876	0.000
Instance4568.3	17.073	0.429	92.718	0.892	38399.164	15319.336	7.353	0.000	1.119	0.000	97214.682	0.000
Instance4568.4	16.941	0.449	92.577	0.896	38353.328	15195.903	6.851	0.000	1.122	0.000	96802.439	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4568.1	0.000	0.000
Instance4568.2	0.000	0.000
Instance4568.3	0.000	0.000
Instance4568.4	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
Instance4568.1	16.134	0.484	93.273	0.893	38388.010	15348.146	6.704	0.000	1.120	0.000	97210.821	0.000
Instance4568.2	17.670	0.468	93.190	0.895	38465.439	15212.681	6.856	0.000	1.122	0.000	96540.876	0.000
Instance4568.3	17.073	0.429	92.718	0.892	38399.164	15319.336	7.353	0.000	1.119	0.000	97214.682	0.000
Instance4568.4	16.941	0.449	92.577	0.896	38353.328	15195.903	6.851	0.000	1.122	0.000	96802.439	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.137	0.000	1.377
Available MBytes	126518.202	126439.000	127536.000
Free System Page Table Entries	16460762.123	16460426.000	16461053.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	218682077.235	218664960.000	218730496.000
Pool Paged Bytes	431638874.751	431603712.000	431681536.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

D.2 Test log

Test Log

```
8/12/2015 11:09:13 AM -- Preparing for testing ...
8/12/2015 11:09:18 AM -- Attaching databases ...
8/12/2015 11:09:18 AM -- Preparations for testing are complete.
8/12/2015 11:09:18 AM -- Starting transaction dispatch ..
8/12/2015 11:09:18 AM -- Database cache settings: (minimum: 128.0 MB, maximum: 1.0 GB)
8/12/2015 11:09:18 AM -- Database flush thresholds: (start: 10.2 MB, stop: 20.5 MB)
8/12/2015 11:09:22 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
8/12/2015 11:09:22 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
8/12/2015 11:09:23 AM -- Operation mix: Sessions 8, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
8/12/2015 11:09:23 AM -- Performance logging started (interval: 15000 ms).
8/12/2015 11:09:23 AM -- Generating log files ...
8/12/2015 1:50:19 PM -- C:\DB\DB1 (101.4% generated), C:\DB\DB2 (102.6% generated), C:\DB\DB3 (100.2% generated) and C:\DB\DB4 (101.2% generated)
8/12/2015 1:50:20 PM -- Performance logging has ended.
8/12/2015 1:50:20 PM -- JetInterop batch transaction stats: 17082, 17082, 17082 and 17082.
8/12/2015 1:50:20 PM -- Dispatching transactions ends.
8/12/2015 1:50:20 PM -- Shutting down databases ...
8/12/2015 1:50:22 PM -- Instance4568.1 (complete), Instance4568.2 (complete), Instance4568.3 (complete) and Instance4568.4 (complete)
8/12/2015 1:50:22 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_8_12_11_9_22.blg has 643 samples.
8/12/2015 1:50:22 PM -- Creating test report ...
8/12/2015 1:50:23 PM -- Instance4568.1 has 19.8 for I/O Database Reads Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.1 has 1.0 for I/O Log Writes Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.1 has 1.0 for I/O Log Reads Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.2 has 19.3 for I/O Database Reads Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.2 has 1.0 for I/O Log Writes Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.2 has 1.0 for I/O Log Reads Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.3 has 19.3 for I/O Database Reads Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.3 has 1.0 for I/O Log Writes Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.3 has 1.0 for I/O Log Reads Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.4 has 19.2 for I/O Database Reads Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.4 has 1.0 for I/O Log Writes Average Latency.
8/12/2015 1:50:23 PM -- Instance4568.4 has 1.0 for I/O Log Reads Average Latency.
8/12/2015 1:50:23 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
8/12/2015 1:50:23 PM -- The test has 0 Database Page Fault Stalls/sec samples higher than 0.
8/12/2015 1:50:23 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_8_12_11_9_22.xml has 642 samples queried.
8/12/2015 1:50:23 PM -- C:\Program Files\Exchange Jetstress\Performance_2015_8_12_11_9_22.html was saved.
8/12/2015 1:50:24 PM -- Performance logging started (interval: 2000 ms).
8/12/2015 1:50:24 PM -- Recovering databases ...
8/12/2015 2:28:32 PM -- Performance logging has ended.
8/12/2015 2:28:32 PM -- Instance4568.1 (2266.9289832), Instance4568.2 (2287.3779462), Instance4568.3 (2236.34039) and Instance4568.4 (2253.788597)
8/12/2015 2:28:32 PM -- C:\Program Files\Exchange Jetstress\SoftRecovery_2015_8_12_13_50_23.blg has 1134 samples.
8/12/2015 2:28:32 PM -- Creating test report ...
```

E Additional resources

E.1 Technical support and resources

For SCv2000 support, visit Dell.com/support.

Dell TechCenter is an online technical community for IT professionals and is a great resource to discover and learn about a wide range of technologies such as storage, servers, networking, software, and cloud management.

E.2 Related documentation

See the following referenced or recommended resources related to this document:

- Microsoft ESRP program
- SC Series storage and other Dell storage solutions
- SC Series technical content