

Hitachi Adaptable Modular Storage 2300 Dynamically Provisioned 58,400 User Exchange 2010 Mailbox Resiliency Storage Solution

Tested with: ESRP – Storage Version 3.0

Test Date: May 2010

Notices and Disclaimer

Copyright © 2010 Hitachi Data Systems Corporation. All rights reserved.

The performance data contained herein was obtained in a controlled isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While Hitachi Data Systems Corporation has reviewed each item for accuracy in a specific situation, there is no guarantee that the same results can be obtained elsewhere.

All designs, specifications, statements, information and recommendations (collectively, "designs") in this manual are presented "AS IS," with all faults. Hitachi Data Systems Corporation and its suppliers disclaim all warranties, including without limitation, the warranty of merchantability, fitness for a particular purpose and non-infringement or arising from a course of dealing, usage or trade practice. In no event shall Hitachi Data Systems Corporation or its suppliers be liable for any indirect, special, consequential or incidental damages, including without limitation, lost profit or loss or damage to data arising out of the use or inability to use the designs, even if Hitachi Data Systems Corporation or its suppliers have been advised of the possibility of such damages.

This document has been reviewed for accuracy as of the date of initial publication. Hitachi Data Systems Corporation may make improvements and/or changes in product and/or programs at any time without notice.

Table of Contents

Overview	1
Disclaimer	1
Features	1
Solution Description	2
Targeted Customer Profile	7
Tested Deployment	8
Replication Configuration	10
Best Practices	12
Core Storage.....	13
Storage-based Replication.....	14
Backup Strategy.....	14
Test Result Summary	14
Reliability.....	14
Storage Performance Results.....	14
Database Backup and Recovery Performance	18
Conclusion	18
Appendix A — RAID-5 Drive Failure and Rebuild	19
Appendix B — Test Reports	20
Performance Test Result: SUN159.....	20
Performance Test Database Checksums Result: SUN159.....	25
Stress Test Database Performance Result: SUN159.....	27
Stress Test Database Checksums Result: SUN159	32
Database Backup Test Result: SUN159	34
Soft Recovery Test Result: SUN159	36
Soft Recovery Test Performance Result: SUN159.....	40

Hitachi Adaptable Modular Storage 2300 Dynamically Provisioned 58,400 User Exchange 2010 Mailbox Resiliency Storage Solution

Tested with: ESRP – Storage Version 3.0

Test Date: May 2010

Overview

This document provides information on a Hitachi Adaptable Modular Storage 2300 storage solution for Microsoft Exchange Server 2010, based on the Microsoft Exchange Solution Reviewed Program (ESRP) – Storage program. For more information about the contents of this document or Hitachi Data Systems best practice recommendations for Microsoft Exchange Server 2010 storage design, see Hitachi Data Systems [Microsoft Exchange Solutions Web page](#).

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 information about the Microsoft ESRP – Storage program, see [TechNet's overview of the program](#).

Disclaimer

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

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

Features

The purpose of this testing was to measure the ESRP 3.0 results on a Microsoft Exchange 2010 environment with 58,400 users and eight servers. This testing used the Hitachi Adaptable Modular Storage 2300 storage system Hitachi Dynamic Provisioning software in a two-pool RAID-5 configuration (one for databases and one for logs) in a resiliency configuration. These results help answer questions about the kind of performance capabilities to expect with a large-scale Exchange deployment on the 2300.

This solution includes Exchange 2010 Mailbox Resiliency by using the new Database Availability Group (DAG) feature. This tested configuration uses eight DAGs, each containing two database copies and two servers. The test configuration was capable of supporting 58,400 users with a 0.12 IOPS per user profile and user mailbox size of 1024MB. A 2300 with 234 450GB 15K RPM SAS disks, 16GB of cache and eight 4Gbit/s paths was used for these tests. Testing used eight Sun Fire 4270 servers with 32GB of

RAM, two quad-core Intel E5540 2.53GHz CPUs each, 16 Emulex 4Gbit/s Fibre Channel adapters, and Windows Server 2008 R2 Enterprise.

The Hitachi Adaptable Modular Storage 2300 is a medium-sized, high-performance, highly reliable midrange storage system that can scale to 240 disks while maintaining 99.999% availability. It is highly suitable for a variety of applications and host platforms and is modular in scale. With the option of in-system and cross-system replication functionality, the 2300 is fully capable of being used as the core underlying storage platform for high-performance Exchange Server 2010 architectures.

Solution Description

Deploying Microsoft Exchange Server 2010 requires careful consideration of all aspects of the solution architecture. Host servers need to be configured so that they are robust enough to handle the required Exchange load. The storage solution must be designed to provide the necessary performance while also being reliable and easy to administer. Of course, an effective backup and recovery plan should be incorporated into the solution as well. The aim of this solution report is to provide a tested configuration that utilizes the 2300 to meet the needs of a large Exchange Server deployment.

This solution uses Hitachi Dynamic Provisioning software, which is enabled on the 2300 via a license key. In the most basic sense, Hitachi Dynamic Provisioning software is similar to the use of a host-based logical volume manager (LVM), but with several additional features available within the 2300 and without the need to install software on the host or incur host processing overhead. Hitachi Dynamic Provisioning software is a superior solution. Hitachi Dynamic Provisioning software provides for one or more pools of wide striping across many RAID groups within a 2300. One or more Dynamic Provisioning virtual volumes (DP-VOLs) of a user-specified logical size (with no initial physical space allocated) are created against each pool.

Hitachi Dynamic Provisioning software also carries the side benefit of thin provisioning, where physical space is only assigned from the pool to the DP-VOL as needed using 1GB chunks, up to the logical size specified for each DP-VOL. A pool can also be dynamically expanded by adding more RAID groups without disruption or requiring downtime. Upon expansion, a pool can easily be rebalanced so that the data and workload is wide striped evenly across the current and newly added RAID groups that make up the pool.

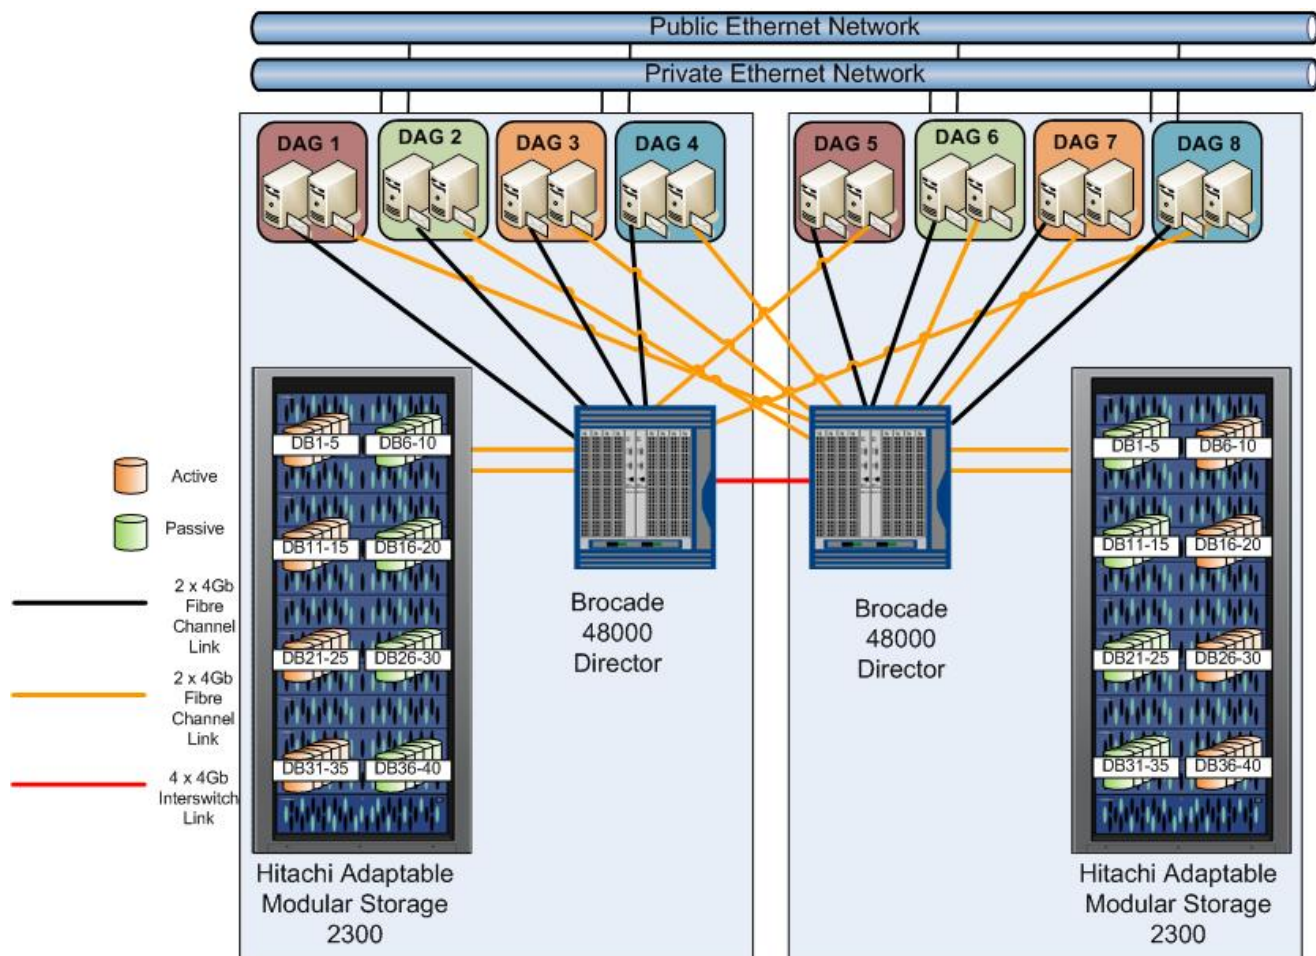
High availability is also a part of this solution with the use of the new DAG feature, which is the base component of the high availability and site resilience framework built into Microsoft Exchange Server 2010. A DAG is a group of up to 16 Mailbox servers that host a set of databases and provide automatic database-level recovery from failures that affect individual servers or databases.

Any server in a DAG can host a copy of a mailbox database from any other server in the DAG. When a server is added to a DAG, it monitors and works with the other servers in the DAG to provide automatic recovery, delivering a robust, highly-available Exchange solution without the administrative complexities of traditional failover clustering. For more information about the DAG feature in Exchange Server 2010, see <http://technet.microsoft.com/en-us/library/dd979799.aspx>.

This solution includes two copies of each Exchange database using eight DAGs, each configured with two servers and that host active mailboxes in five databases. To target the 58,400 user resiliency solution, a Hitachi Adaptable Modular Storage 2300 configured with 234 disks (240 is the maximum) and eight host servers, each configured with 7,300 mailboxes, were used to host the 40 active databases and the simulated database copies for the tests

Each DAG contained two copies of every database; a local, active copy on a server connected to the primary 2300 and the passive copy on another server connected to a second 2300. This recommended configuration can support both high-availability and disaster-recovery scenarios when the active and passive database copies are allocated among both DAG members and dispersed across both 2300s. Each simulated DAG server node in this solution maintains a mirrored configuration and possesses adequate capacity and performance capabilities to support the second set of replicated databases.

Figure 1 illustrates the two systems that make up the recommended DAG configuration that was simulated.




This solution enables organizations to consolidate Exchange Server 2010 DAG deployments on two 2300 storage systems. Using identical hardware and software configurations guarantees that an active database and its replicated copy do not share storage paths, disk spindles or storage controllers, making it a very reliable, high-performing, highly available Exchange Server 2010 solution that is cost effective and easy to manage. This helps ensure that performance and service levels related to storage are maintained regardless of which server is hosting the active database. If further protection is needed in a production environment, additional Exchange Server 2010 mailbox servers can be easily added to support these failover scenarios.

Table 1 illustrates how the 2300 storage system's disks were organized into RAID groups for use by either databases or logs. Each set of colored disks represents a RAID-5 (8D+1P) group. Except for RKA-0 (with 15 internal SAS disks), each RKA is an external disk enclosure with 15 SAS disks.

Table 1. Adaptable Modular Storage 2300 RAID Groups by RKA Tray Layout

Drive Slot	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
RKA 15	26	26	26	26	26	26	26	26	26						
RKA 14	24	24	24	24	24	24	25	25	25	25	25	25	25	25	25
RKA 13	22	22	22	23	23	23	23	23	23	23	23	23	24	24	24
RKA 12	21	21	21	21	21	21	21	21	21	22	22	22	22	22	22
RKA 11	19	19	19	19	19	19	20	20	20	20	20	20	20	20	20
RKA 10	17	17	17	18	18	18	18	18	18	18	18	18	19	19	19
RKA 9	16	16	16	16	16	16	16	16	16	17	17	17	17	17	17
RKA 8	14	14	14	14	14	14	15	15	15	15	15	15	15	15	15
RKA 7	12	12	12	13	13	13	13	13	13	13	13	13	14	14	14
RKA 6	11	11	11	11	11	11	11	11	11	12	12	12	12	12	12
RKA 5	9	9	9	9	9	9	10	10	10	10	10	10	10	10	10
RKA 4	7	7	7	8	8	8	8	8	8	8	8	8	9	9	9
RKA 3	6	6	6	6	6	6	6	6	6	7	7	7	7	7	7
RKA 2	4	4	4	4	4	4	5	5	5	5	5	5	5	5	5
RKA 1	2	2	2	3	3	3	3	3	3	3	3	3	4	4	4
RKA 0	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2



Two Dynamic Provisioning pools were created, one for the databases and other for the logs. The database pool was created from 23 RAID-5 groups and the log pool was created from three RAID-5 groups. From the database pool, 40 DP-VOLs (each specified to have a 1750GB size limit) were created for 40 databases (five per server). From the log pool, 40 DP-VOLs (each specified to have a size limit of 175GB) were created for 40 logs (five per server).

Table 2 outlines the port layout for the servers.

Table 2. Adaptable Modular Storage 2300 Port to Server Layout

<i>Server</i>	<i>Primary Path</i>	<i>Secondary Path</i>
SUN170	0A	1A
SUN169	0B	1B
SUN164	0C	1C
SUN163	0D	1D
SUN161	0E	1E
SUN160	0F	1F
SUN159	0G	1G
SUN158	0H	1H

Table 3 outlines the port layout with the database DP-VOL assignments for the primary storage and servers. An identical configuration is deployed on the replicated storage and servers for this solution.

Table 3. Adaptable Modular Storage 2300 Port to Log DP-VOL Layout

<i>Port</i>	<i>Database</i>	<i>DB DP-VOL</i>
0A	1-5	14-20
0B	6-10	25-29
0C	11-15	30-34
0D	16-20	35-39
1A	21-25	0-4
1B	26-30	5-9
1C	31-35	10-14
1D	36-40	15-19

Table 4 outlines the port layout with the log DP-VOL assignments for primary storage and servers. An identical configuration is deployed on the replicated storage and servers for this solution.

Table 4. Adaptable Modular Storage 2300 Port to Log DP-VOL Layout

<i>Port</i>	<i>Log</i>	<i>DP-VOL</i>
0A	1-5	60-64
0B	6-10	65-69
0C	11-15	70-74
0D	16-20	75-79
1A	21-25	40-44
1B	26-30	45-49
1C	31-35	50-54
1D	36-40	55-59

Table 5 provides the detailed specifications for the storage configuration which uses RAID-5 (8D+1P) groups and 450GB 15K disks. Dynamic Provisioning pool 0 is dedicated for the database and Dynamic Provisioning pool 1 is dedicated for the logs.

Table 5. Adaptable Modular Storage 2300 Configuration Details

<i>Host</i>	<i>Pool</i>	<i>Port</i>	<i>DP-VOL</i>	<i>Size (GB)</i>	<i>Description</i>
SUN170	0	0A/1A	0-4	1750	Databases 1-5
SUN169	0	0B/1B	5-9	1750	Databases 6-10
SUN164	0	0C/1C	10-14	1750	Databases 11-15
SUN163	0	0D/1D	15-19	1750	Databases 16-20
SUN161	0	1A/0A	20-24	1750	Databases 21-25
SUN160	0	1B/0B	24-29	1750	Databases 26-30
SUN159	0	1C/0C	30-34	1750	Databases 31-35
SUN158	0	1D/0D	35-39	1750	Databases 36-40
SUN170	1	0A/1A	40-44	175	Logs 1-5
SUN169	1	0B/1B	45-49	175	Logs 6-10
SUN164	1	0C/1C	50-54	175	Logs 11-15
SUN163	1	0D/1D	55-59	175	Logs 16-20
SUN161	1	1A/0A	60-64	175	Logs 21-25
SUN160	1	1B/0B	65-69	175	Logs 26-30
SUN159	1	1C/0C	70-74	175	Logs 31-35
SUN158	1	1D/0D	75-79	175	Logs 36-40

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. These factors also affect server scalability:

- 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
- Client usage profiles

These factors are all beyond the scope of the ESRP – Storage program. Therefore, the number of mailboxes hosted per server as part of the tested configuration might not necessarily be viable for some customer deployments.

For more information about identifying and addressing performance bottlenecks in an Exchange system, see Microsoft's [Troubleshooting Microsoft Exchange Server Performance](#).

Targeted Customer Profile

This solution is designed for medium to large organizations that plan to consolidate their Exchange Server 2010 storage on high-performance, high-reliability storage systems. This configuration is designed to support 58,400 Exchange users with the following specifications:

- 16 Exchange servers (eight tested, simulating 16 for the database copies)
- Eight Database Availability Groups each maintaining two servers and two copies per database
- Two Adaptable Modular Storage 2300 storage systems (one tested)
- 0.1 IOPS per user (0.12 tested for 20 percent growth)
- 1GB mailbox size
- Mailbox resiliency provides high-availability and used as primary data protection mechanism.
- Adaptable Modular Storage RAID protection against physical failure or loss.
- 24x7 background database maintenance enabled.

Tested Deployment

The following tables summarize the testing environment.

Table 6. Simulated Exchange Configuration

<i>Number of Exchange mailboxes simulated</i>	58,400
<i>Number of database availability groups (DAGs)</i>	8
<i>Number of servers per DAG</i>	2
<i>Number of active mailboxes per server</i>	7,300
<i>Number of databases per host</i>	5
<i>Number of copies per database</i>	2
<i>Number of mailboxes per database</i>	1,460
<i>Simulated profile: I/Os per second per mailbox (IOPS, include 20% headroom)</i>	0.12
<i>Database LU size</i>	1750GB
<i>Log LU size</i>	175GB
<i>Total database size for performance testing</i>	58,400GB
<i>% storage capacity used by Exchange database**</i>	79.5%

**Storage performance characteristics change based on the percentage utilization of the individual disks. Tests that use a small percentage of the storage (~25%) might exhibit reduced throughput if the storage capacity utilization is significantly increased beyond what was tested for this paper.

Table 7. Primary Storage Hardware

Storage connectivity (Fibre Channel, SAS, SATA, iSCSI)	Fibre Channel
Storage model and OS/firmware revision	1 Hitachi Adaptable Modular Storage 2300 Firmware: 0890/B-Y WHQL listing: Hitachi Adaptable Modular Storage 2300
Storage cache	16GB
Number of storage controllers	2
Number of storage ports	8
Maximum bandwidth of storage connectivity to host	32Gb/s (8 x 4Gbit/s ports)
Switch type/model/firmware revision	Brocade 5320, Fabric OS v6.3.0b
HBA model and firmware	Emulex LPe11002, FW:2.82A3
Number of HBAs per host	2 dual-ported HBA per host, 1 4Gbit/s port used per HBA
Host server type	Sun Fire 4270 2 2.54 GHz quad-core Intel Xeon CPUs, 32 GB memory
Total number of disks tested in solution	234
Maximum number of spindles can be hosted in the storage	240

Table 8. Primary Storage Software

HBA Driver	STOR Miniport 7.2.0.12
HBA QueueTarget Setting	0
HBA QueueDepth Setting	32
Multipathing	Hitachi Dynamic Link Manager v6.2.0
Host OS	Microsoft Windows Server 2008 R2 Enterprise
ESE.dll file version	14.00.0639.019
Replication solution name/version	N/A

Table 9. Primary Storage Disk Configuration (Mailbox Store Disks)

Disk type, speed and firmware revision	SAS Disk 4540GB 15K 4C57
Raw capacity per disk (GB)	450GB
Number of physical disks in test	207 (Dynamic Provisioning pool)
Total raw storage capacity (GB)	93,150GB
Disk slice size (GB)	N/A
Number of slices per LUN or number of disks per LUN	N/A
RAID level	RAID 5 (8+1) at storage level
Total formatted capacity	73,416GB (Dynamic Provisioning database pool)
Storage capacity utilization	78.8%
Database capacity utilization	75.1%

Table 10. Primary Storage Disk Configuration (Transaction Log Disks)

Disk type, speed and firmware revision	SAS Disk 4540GB 15K 4C57
Raw capacity per disk (GB)	450GB
Number of spindles in test	27 (Dynamic Provisioning pool)
Total raw storage capacity (GB)	12,150GB
Disk slice size (GB)	N/A
Number of slices per LU or number of disks per LU	N/A
RAID level	RAID 5 (8+1) at storage level
Total formatted capacity	9,576GB

Replication Configuration

The following tables summarize the replication environment.

Table 11. Replicated Configuration

Replication mechanism	Exchange Server 2010 Database Availability Group (DAG)
Number of links	2
Simulated link distance	N/A
Link type	IP
Link bandwidth	GigE (1Gbps)

Table 12. Replicated Storage Hardware

Storage connectivity (Fibre Channel, SAS, SATA, iSCSI)	Fibre Channel
Storage model and OS/firmware revision	1 Hitachi Adaptable Modular Storage 2300 Firmware: 0890/B-Y WHQL listing: Hitachi Adaptable Modular Storage 2300
Storage cache	16GB
Number of storage controllers	2
Number of storage ports	8
Maximum bandwidth of storage connectivity to host	32Gb/s (8 x 4Gbit/s ports)
Switch type/model/firmware revision	Brocade 5320, Fabric OS v6.3.0b
HBA model and firmware	Emulex LPe11002, FW:2.82A3
Number of HBAs per host	2 dual-ported HBA per host, 1 4Gbit/s port used per HBA
Host server type	Sun Fire 4270 2 2.54 GHz quad-core Intel Xeon CPUs, 32 GB memory
Total number of disks tested in solution	240
Maximum number of spindles can be hosted in the storage	240

Table 13. ReplicatedStorage Software

HBA Driver	STOR Miniport 7.2.0.12
HBA QueueTarget Setting	0
HBA QueueDepth Setting	32
Multipathing	Hitachi Dynamic Link Manager v6.2.0
Host OS	Microsoft Windows Server 2008 R2 Enterprise
ESE.dll file version	14.00.0639.019
Replication solution name/version	N/A

Table 14. ReplicatedStorage Disk Configuration (Mailbox Store Disks)

Disk type, speed and firmware revision	SAS Disk 4540GB 15K 4C57
Raw capacity per disk (GB)	450GB
Number of physical disks in test	207 (Dynamic Provisioning pool)
Total raw storage capacity (GB)	93,150GB
Disk slice size (GB)	N/A
Number of slices per LUN or number of disks per LUN	N/A
RAID level	RAID 5 (8+1) at storage level
Total formatted capacity	73,416GB (Dynamic Provisioning database pool)
Storage capacity utilization	78.8%
Database capacity utilization	75.1%

Table 15. ReplicatedStorage Disk Configuration (Transaction Log Disks)

Disk type, speed and firmware revision	SAS Disk 4540GB 15K 4C57
Raw capacity per disk (GB)	450GB
Number of spindles in test	27 (Dynamic Provisioning pool)
Total raw storage capacity (GB)	12,150GB
Disk slice size (GB)	N/A
Number of slices per LU or number of disks per LU	N/A
RAID level	RAID 5 (8+1) at storage level
Total formatted capacity	9,576GB

Best Practices

Microsoft Exchange Server 2010 is a very disk-intensive application. It presents two distinct workload patterns to the storage, with 32KB random read/write operations to the databases, and sequential write operations of varying size (between 512 bytes up to the log buffer size) to the transaction logs. For this reason, designing an optimal storage configuration can prove challenging in practice. Based on the testing run using the ESRP framework, Hitachi Data Systems recommends these best practices to improve the performance of the Hitachi Adaptable Modular Storage 2300 running Exchange 2010.

For more information about Exchange 2010 best practices for storage design, see the Microsoft TechNet article "[Mailbox Server Storage Design](#)."

Core Storage

1. When formatting a newly partitioned LUN, HDS recommends that for database and log files the ALU is set to 64K and 4K respectively.
2. Disk alignment is no longer required when using Microsoft Windows Server 2008.
3. Keep the Exchange workload isolated from other applications. Mixing another I/O intensive application whose workload differs from Exchange can cause the performance for both applications to degrade.
4. Due to the difference in I/O patterns, isolate the Exchange database from the log groups. Create a dedicated Dynamic Provisioning Pool for the databases and a separate Pool for the logs.
5. Use Hitachi Dynamic Link Manager multipathing software to provide fault tolerance and high availability for host connectivity.
6. Use Hitachi Dynamic Provisioning software to simplify storage management of the Exchange database and log volumes
7. Due to the difference in I/O patterns, isolate the Exchange database from the log groups. Create a dedicated Dynamic Provisioning pool for the databases and a separate pool for the logs.
8. Hitachi Data Systems recommends RAID-5 or RAID-1+0 groups for both the database pools and for the log pool. Use of RAID-1+0 allows more writes at a lower response time under heavier loads. RAID-1+0 also supports a shorter RAID group rebuild time on failure of a disk.
9. Log LUs should be at least 10 percent of the size of the database LUs.
10. Hitachi Data Systems does not recommend LU concatenation.
11. Hitachi Data Systems recommends implementing Mailbox Resiliency using the Exchange Server 2010 Database Availability Group feature.
12. Ensure that each DAG maintains at least two database copies to provide high availability.
13. Isolate active databases and their replicated copies in separate Dynamic Provisioning pools or ensure that they are located on a separate 2300.
14. Use fewer, larger LUs for Exchange 2010 databases (up to 2TB) with Background Database Maintenance (24x7) enabled.
15. Size storage solutions for Exchange based primarily on performance criteria. The number of disks, RAID level and percent utilization of each disk directly affect the level of achievable performance. Factor in capacity requirements only after performance is addressed.
16. Disk size is unrelated to performance with regards to IOPS or throughput rates. Disk size is related to the usable capacity of all of the LUs from a RAID group, which is a choice users make.
17. The number of spindles, coupled with the RAID level, determines the physical IOPS capacity of the RAID group and all of its LUs. If the disk has too few spindles, the response times grow to large values very quickly.
18. For more information about RAID-5 drive failure and rebuild, see Appendix A.

Storage-based Replication

N/A

Backup Strategy

N/A

Test Result Summary

This section provides a high-level summary of the test data from ESRP and the link to the detailed HTML reports that are generated by ESRP testing framework.

Reliability

A number of tests in the framework check reliability spanning a 24-hour window. The goal is to verify the storage can handle high I/O load for a long period of time. Following these stress tests, both log and database files are analyzed for integrity to ensure that no database or log corruption occurs.

- No errors were reported in the event log file for the storage reliability testing
- No errors were reported for the [database](#) and [log](#) checksum process
- Backup to disk test is N/A
- Database checksum on the remote storage database is N/A

Storage [Performance](#) Results

Primary storage performance testing exercises the storage with maximum sustainable Exchange type of I/O for two hours. The test shows how long it takes for the storage to respond to an I/O under load. The following data is the sum of all of the logical disk I/Os and average of all the logical disks I/O latency in the two-hour test duration.

Individual Server Metrics

Individual server metrics show the sum of I/Os across database and the average latency across all databases on a per-server basis.

Table 16. Individual Server Metrics for Exchange Server (SUN170)

Database I/O	
<i>Database disk transfers per second</i>	1025
<i>Database disk reads per second</i>	684
<i>Database disk writes per second</i>	341
<i>Average database disk read latency (ms)</i>	16.1
<i>Average database disk write latency (ms)</i>	8.2
Transaction Log I/O	
<i>Log disk writes per second</i>	258
<i>Average log disk write latency (ms)</i>	4.0

Table 17. Individual Server Metrics for Exchange Server (SUN169)

Database I/O	
<i>Database disk transfers per second</i>	875
<i>Database disk reads per second</i>	554
<i>Database disk writes per second</i>	321
<i>Average database disk read latency (ms)</i>	17.2
<i>Average database disk write latency (ms)</i>	10.1
Transaction Log I/O	
<i>Log disk writes per second</i>	231
<i>Average log disk write latency (ms)</i>	4.9

Table 18. Individual Server Metrics for Exchange Server (SUN164)

Database I/O	
<i>Database disk transfers per second</i>	970
<i>Database disk reads per second</i>	647
<i>Database disk writes per second</i>	323
<i>Average database disk read latency (ms)</i>	17.2
<i>Average database disk write latency (ms)</i>	9.8
Transaction Log I/O	
<i>Log disk writes per second</i>	234
<i>Average log disk write latency (ms)</i>	4.8

Table 19. Individual Server Metrics for Exchange Server (SUN163)

Database I/O	
<i>Database disk transfers per second</i>	1001
<i>Database disk reads per second</i>	668
<i>Database disk writes per second</i>	333
<i>Average database disk read latency (ms)</i>	16.2
<i>Average database disk write latency (ms)</i>	8.5
Transaction Log I/O	
<i>Log disk writes per second</i>	252
<i>Average log disk write latency (ms)</i>	4.1

Table 20. Individual Server Metrics for Exchange Server (SUN161)

Database I/O	
<i>Database disk transfers per second</i>	1025
<i>Database disk reads per second</i>	684
<i>Database disk writes per second</i>	341
<i>Average database disk read latency (ms)</i>	16.1
<i>Average database disk write latency (ms)</i>	8.6
Transaction Log I/O	
<i>Log disk writes per second</i>	254
<i>Average log disk write latency (ms)</i>	4.2

Table 21. Individual Server Metrics for Exchange Server (SUN160)

Database I/O	
<i>Database disk transfers per second</i>	1019
<i>Database disk reads per second</i>	679
<i>Database disk writes per second</i>	340
<i>Average database disk read latency (ms)</i>	16.4
<i>Average database disk write latency (ms)</i>	8.4
Transaction Log I/O	
<i>Log disk writes per second</i>	255
<i>Average log disk write latency (ms)</i>	4.2

Table 22. Individual Server Metrics for Exchange Server (SUN159)

Database I/O	
<i>Database disk transfers per second</i>	967
<i>Database disk reads per second</i>	647
<i>Database disk writes per second</i>	320
<i>Average database disk read latency (ms)</i>	16.9
<i>Average database disk write latency (ms)</i>	9.1
Transaction Log I/O	
<i>Log disk writes per second</i>	230
<i>Average log disk write latency (ms)</i>	5.0

Table 23. Individual Server Metrics for Exchange Server (SUN158)

Database I/O	
<i>Database disk transfers per second</i>	1016
<i>Database disk reads per second</i>	679
<i>Database disk writes per second</i>	337
<i>Average database disk read latency (ms)</i>	15.7
<i>Average database disk write latency (ms)</i>	7.4
Transaction Log I/O	
<i>Log disk writes per second</i>	251
<i>Average log disk write latency (ms)</i>	4.3

Aggregate Performance Across All Servers Metrics

The aggregate performance across all server metrics shows the sum of I/Os across all servers in the solution and the average latency across all servers in the solution.

Table 24. Aggregate Performance for Exchange Server 2010

Database I/O	
<i>Database disk transfers per second</i>	7898
<i>Database disk reads per second</i>	5242
<i>Database disk writes per second</i>	2656
<i>Average database disk read latency (ms)</i>	16.5
<i>Average database disk write latency (ms)</i>	8.7
Transaction Log I/O	
<i>Log disk writes per second</i>	1965
<i>Average log disk write latency (ms)</i>	4.5

Database Backup and Recovery Performance

This section has two tests: The first measures the sequential read rate of the database files and the second measures recovery/replay performance (playing transaction logs in to the database).

Database Read-only Performance

This test measures the maximum rate at which databases can be backed up via VSS. The following tables show the average rate for a single database file.

Table 25. Database Read-only Performance

<i>MB read per second per database</i>	34.6
<i>MB read per second total per server</i>	173.4

Transaction Log Recovery/Replay Performance

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

Table 26. Transaction Log Recovery/Replay Performance

<i>Average time to play one log file (sec)</i>	1.50
--	------

Conclusion

This document details a robust Exchange Server 2010 Resiliency solution capable of supporting 58,400 users with a 0.12 IOPS per user profile and user mailbox size of 1GB using eight DAGs, each configured with two server nodes. An Adaptable Modular Storage 2300, with 16GB of cache and eight 4Gbit/s Fibre Channel host paths, using Hitachi Dynamic Provisioning (with two pools) and 240 450GB 15K RPM SAS disks in a RAID-5 configuration was used for these tests. Testing confirmed that the 2300 is capable of delivering the IOPS and capacity requirements needed to support 58,400 Exchange mailboxes configured with the specified user profile, while maintaining additional headroom to support peak throughput.

The solution outlined in this document does not include data protection components such as local or remote replication. Adding these technologies may affect performance and capacity requirements and each need to be factored into the storage design accordingly.

For more information about planning Exchange Server 2010 storage architectures for the Hitachi Adaptable Modular Storage 2000 family, see <http://www.hds.com/assets/pdf/hitachi-ams-2000-family.pdf>.

This document is developed by Hitachi Data Systems and reviewed by Microsoft Exchange Product team. The test results and data presented in this document are based on the tests introduced in the ESRP test framework. Do not quote the data directly for pre-deployment verification. It is still necessary to validate the storage design for a specific customer environment.

The ESRP program is not designed to be a benchmarking program; tests do not generate the maximum throughput for a given solution. Rather, it is focused on producing recommendations from vendors for Exchange application. Thus, do not use the data presented in this document for direct comparisons among the solutions.

Appendix A — RAID-5 Drive Failure and Rebuild

These ESRP tests used RAID-5 (8D+1P) rather than RAID-6 (for example, 4D+2P) or RAID-1+0 (for example, 4D+4D). RAID-5 is a much more capacity-efficient RAID level than the others, losing only 12.5 percent of the usable space (using 8D+1P) instead of 33 percent (4D+2P) or 50 percent (4D+4D). One downside with the use of parity RAID instead of mirrored and striped (RAID-1+0) is that for *writes*, the internal disk write penalty is higher. For SAS or Fibre Channel disks, RAID-5 requires four physical disk I/Os on the backend for every host write, whereas RAID-1+0 consumes two physical I/Os. RAID-6 requires six physical I/Os for each host write.

The other downside is the RAID group rebuild time after a sudden disk failure. The Hitachi Adaptable Modular Storage 2000 family is always scanning the storage system looking for *soft fails*, because excessive soft fails often predict a hard failure. If the number of soft fails exceeds the failure threshold in a 24-hour period (user parameter driven), the 2000 family storage system first executes a disk-to-disk copy to a global hot spare (thus avoiding a RAID-5 or RAID-6 rebuild), and then marks the disk as *failed* and replaces it.

If hard fail does occur, for RAID-1+0, the contents of the good disk are mirrored onto a spare disk (these 'hot spares' are user defined to be in several disk enclosures on a storage system). For RAID-5 and RAID-6, all disks in the RAID group must be read to recreate the missing data and parity that was on the failed disk onto the spare disk. This rebuild mode is called *Corrective Copy*. An associated array setting called *[Drive] Restore Options* determines how aggressive the rebuild operation is in the face of ongoing host I/Os. This setting has three levels: aggressive, moderate and background.

Lab tests show that, on a RAID-6 group using Fibre Channel disks (the only sample available), and an *aggressive* Restore Option setting, a RAID-6 (8D+2P) group Corrective Copy operation requires about 30 minutes to complete in the absence of host workloads on LUs from that RAID group. In the presence of sustained 100 percent sequential write workloads to LUs from that RAID group, this rebuild time increased to 18 hours. The host performance on a LU from that RAID group was measured at 154MB/s (normal state) and 95MB/s (Corrective Copy state). Had this been RAID-5, the Corrective Copy times would have been reduced.

Appendix B — Test Reports

This appendix contains Jetstress test results for one of the servers used in testing this storage solution. These test results are representative of the results obtained for all of the servers tested.

Performance Test Result: SUN159

Test Summary

Overall Test Result	Pass
Machine Name	SUN159
Test Description	
Test Start Time	6/9/2010 8:09:50 PM
Test End Time	6/9/2010 10:35:52 PM
Collection Start Time	6/9/2010 8:11:36 PM
Collection End Time	6/9/2010 10:11:35 PM
Jetstress Version	14.01.0043.000
Ese Version	14.00.0639.019
Operating System	Windows Server 2008 R2 Enterprise (6.1.7600.0)
Performance Log	C:\Exchange 2010\SUN159\RAID5\7300 users\Performance_2010_6_9_20_10_2.blg C:\Exchange 2010\SUN159\RAID5\7300 users\DBCchecksum_2010_6_9_22_35_52.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	871.108
Target Transactional I/O per Second	876
Initial Database Size (bytes)	7838373314560
Final Database Size (bytes)	7841611317248
Database Files (Count)	5

Jetstress System Parameters

Thread Count	7 (per database)
Minimum Database Cache	160.0 MB
Maximum Database Cache	1280.0 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

Instance1448.1	Log Path: C:\alogluns\log1 Database: C:\asgluns\srg1\Jetstress001001.edb
Instance1448.2	Log Path: C:\alogluns\log2 Database: C:\asgluns\srg2\Jetstress002001.edb
Instance1448.3	Log Path: C:\alogluns\log3 Database: C:\asgluns\srg3\Jetstress003001.edb
Instance1448.4	Log Path: C:\alogluns\log4 Database: C:\asgluns\srg4\Jetstress004001.edb
Instance1448.5	Log Path: C:\alogluns\log5 Database: C:\asgluns\srg5\Jetstress005001.edb

Transactional I/O Performance

MSExchange Database => Instance	I/O Database Reads /sec	I/O Database Writes /sec	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
Instance 1448.1	17.524	11.128	110.867	64.521	34697.264	37218.807	0.000	5.028	0.000	46.197	0.000	5905.630
Instance 1448.2	16.990	9.080	109.333	63.351	34738.848	37310.178	0.000	5.109	0.000	45.916	0.000	5869.221
Instance 1448.3	16.665	9.062	110.222	63.798	34717.609	37190.940	0.000	4.825	0.000	45.999	0.000	5794.304
Instance 1448.4	16.921	7.996	110.035	63.742	34803.037	37282.963	0.000	5.082	0.000	46.055	0.000	5886.545
Instance 1448.5	16.651	8.272	110.938	64.301	34626.555	37185.366	0.000	5.111	0.000	46.359	0.000	5864.933

Background Database Maintenance I/O Performance

MSExchange Database => Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance1448.1	17.167	261898.571
Instance1448.2	19.219	261950.062
Instance1448.3	19.886	261829.789
Instance1448.4	19.270	261813.989
Instance1448.5	19.859	261932.414

Log Replication I/O Performance

MSExchange Database => Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance1448.1	1.126	231975.056
Instance1448.2	1.114	230560.509
Instance1448.3	1.101	232100.932
Instance1448.4	1.118	231467.295
Instance1448.5	1.123	232615.719

Total I/O Performance

<i>MSExchange => Instance</i>	<i>I/O Database Reads Average Latency (msec)</i>	<i>I/O Database Writes Average Latency (msec)</i>	<i>I/O Database Reads/s</i>	<i>I/O Database Writes/s</i>	<i>I/O Database Reads Average Bytes</i>	<i>I/O Database Writes Average Bytes</i>	<i>I/O Log Reads Average Latency (msec)</i>	<i>I/O Log Writes Average Latency (msec)</i>	<i>I/O Log Reads/s</i>	<i>I/O Log Writes/s</i>	<i>I/O Log Reads Average Bytes</i>	<i>I/O Log Writes Average Bytes</i>
Instance 1448.1	17.524	11.128	128.034	64.521	65160.699	37218.807	41.814	5.028	1.126	46.197	231975.056	5905.630
Instance 1448.2	16.990	9.080	128.552	63.351	68707.227	37310.178	48.000	5.109	1.114	45.916	230560.509	5869.221
Instance 1448.3	16.665	9.062	130.108	63.798	69430.521	37190.940	42.177	4.825	1.101	45.999	232100.932	5794.304
Instance 1448.4	16.921	7.996	129.305	63.742	68633.113	37282.963	42.811	5.082	1.118	46.055	231467.295	5886.545
Instance 1448.5	16.651	8.272	130.797	64.301	69138.807	37185.366	38.640	5.111	1.123	46.359	232615.719	5864.933

Host System Performance

<i>Counter</i>	<i>Average</i>	<i>Minimum</i>	<i>Maximum</i>
% Processor Time	0.819	0.000	3.251
Available MBytes	29161.635	29124.000	29224.000
Free System Page Table Entries	33555643.894	33555639.000	33555645.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	66536644.267	65175552.000	67416064.000
Pool Paged Bytes	94281881.600	94162944.000	94339072.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

6/7/2010 8:11:08 PM -- Prepare testing begins ...
6/7/2010 8:11:08 PM -- Creating C:\asgluns\sg1\Jetstress001001.edb.
6/7/2010 8:11:08 PM -- Database cache settings: (minimum: 32.0 MB, maximum: 256.0 MB)
6/7/2010 8:11:08 PM -- Database flush thresholds: (start: 2.5 MB, stop: 5.1 MB)
6/9/2010 1:21:41 AM -- 100.0% of 1.4 TB complete (498968822 records inserted).
6/9/2010 1:21:41 AM -- 100.0% of 1.4 TB complete (498968825 records inserted).
6/9/2010 1:21:43 AM -- Duplicating 4 database(s):
6/9/2010 10:00:15 AM -- 100.0% of 5.7 TB complete (5.7 TB duplicated).
6/9/2010 10:00:20 AM -- Attaching databases ...
6/9/2010 10:00:20 AM -- Prepare testing ends.

6/9/2010 8:09:50 PM -- Jetstress testing begins ...
6/9/2010 8:09:50 PM -- Prepare testing begins ...
6/9/2010 8:09:56 PM -- Attaching databases ...
6/9/2010 8:09:56 PM -- Prepare testing ends.
6/9/2010 8:09:56 PM -- Dispatching transactions begins ...
6/9/2010 8:09:56 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
6/9/2010 8:09:56 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
6/9/2010 8:10:02 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
6/9/2010 8:10:02 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
6/9/2010 8:10:08 PM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
6/9/2010 8:10:08 PM -- Performance logging begins (interval: 15000 ms).
6/9/2010 8:10:09 PM -- Attaining prerequisites:
6/9/2010 8:11:36 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1213895000.0 (lower bound: 1207960000.0, upper bound: none)
6/9/2010 10:11:36 PM -- Performance logging ends.
6/9/2010 10:35:29 PM -- JetInterop batch transaction stats: 46727, 46188, 46499, 46467 and 46614.
6/9/2010 10:35:30 PM -- Dispatching transactions ends.
6/9/2010 10:35:30 PM -- Shutting down databases ...
6/9/2010 10:35:52 PM -- Instance1448.1 (complete), Instance1448.2 (complete), Instance1448.3 (complete), Instance1448.4 (complete) and Instance1448.5 (complete)
6/9/2010 10:35:53 PM -- Performance logging begins (interval: 30000 ms).
6/9/2010 10:35:53 PM -- Verifying database checksums ...
6/10/2010 1:37:00 PM -- C:\asgluns\sg1 (100% processed), C:\asgluns\sg2 (100% processed), C:\asgluns\sg3 (100% processed), C:\asgluns\sg4 (100% processed) and C:\asgluns\sg5 (100% processed)
6/10/2010 1:37:00 PM -- Performance logging ends.
6/10/2010 1:37:00 PM -- C:\Exchange 2010\SUN159\RAID5\7300 users\DBChecksum_2010_6_9_22_35_52.blg has 1801 samples.
6/10/2010 1:37:08 PM -- C:\Exchange 2010\SUN159\RAID5\7300 users\DBChecksum_2010_6_9_22_35_52.html is saved.
6/10/2010 1:37:08 PM -- Verifying log checksums ...
6/10/2010 1:37:09 PM -- C:\alogluns\log1 (11 log(s) processed), C:\alogluns\log2 (11 log(s) processed), C:\alogluns\log3 (11 log(s) processed), C:\alogluns\log4 (11 log(s) processed) and C:\alogluns\log5 (11 log(s) processed)
6/10/2010 1:37:09 PM -- C:\Exchange 2010\SUN159\RAID5\7300 users\Performance_2010_6_9_20_10_2.blg has 485 samples.
6/10/2010 1:37:09 PM -- Creating test report ...
6/10/2010 1:37:12 PM -- Instance1448.1 has 17.5 for I/O Database Reads Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.1 has 5.0 for I/O Log Writes Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.1 has 5.0 for I/O Log Reads Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.2 has 17.0 for I/O Database Reads Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.2 has 5.1 for I/O Log Writes Average Latency.

6/10/2010 1:37:12 PM -- Instance1448.2 has 5.1 for I/O Log Reads Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.3 has 16.7 for I/O Database Reads Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.3 has 4.8 for I/O Log Writes Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.3 has 4.8 for I/O Log Reads Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.4 has 16.9 for I/O Database Reads Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.4 has 5.1 for I/O Log Writes Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.4 has 5.1 for I/O Log Reads Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.5 has 16.7 for I/O Database Reads Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.5 has 5.1 for I/O Log Writes Average Latency.
6/10/2010 1:37:12 PM -- Instance1448.5 has 5.1 for I/O Log Reads Average Latency.
6/10/2010 1:37:12 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
6/10/2010 1:37:12 PM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
6/10/2010 1:37:12 PM -- C:\Exchange_2010\SUN159\RAID5\7300 users\Performance_2010_6_9_20_10_2.xml has 479 samples queried.

Performance Test Database Checksums Result: SUN159

Checksum Statistics - All

<i>Database</i>	<i>Seen pages</i>	<i>Bad pages</i>	<i>Correctable pages</i>	<i>Wrong page-number pages</i>	<i>File length / seconds taken</i>
C:\asgluns\sg1\Jetstress001001.edb	47861602	0	0	0	1495675 MBytes / 54066 sec
C:\asgluns\sg2\Jetstress002001.edb	47861602	0	0	0	1495675 MBytes / 43102 sec
C:\asgluns\sg3\Jetstress003001.edb	47861090	0	0	0	1495659 MBytes / 49665 sec
C:\asgluns\sg4\Jetstress004001.edb	47861346	0	0	0	1495667 MBytes / 43235 sec
C:\asgluns\sg5\Jetstress005001.edb	47861346	0	0	0	1495667 MBytes / 49811 sec
(Sum)	239306986	0	0	0	7478343 MBytes / 54066 sec

Disk Subsystem Performance of Checksum

Logical Disk	Avg. Disk sec/Read	Avg. Disk sec/Write	Disk Reads/sec	Disk Writes/sec	Avg. Disk Bytes/Read
C:\asgluns\sg1	0.131	0.000	442.482	0.000	65535.996
C:\asgluns\sg2	0.095	0.000	555.172	0.000	65536.000
C:\asgluns\sg3	0.106	0.000	481.731	0.000	65536.000
C:\asgluns\sg4	0.094	0.000	553.161	0.000	65536.000
C:\asgluns\sg5	0.106	0.000	480.046	0.000	65536.000

Memory System Performance of Checksum

Counter	Average	Minimum	Maximum
% Processor Time	0.603	0.000	3.427
Available MBytes	30515.586	30501.000	30529.000
Free System Page Table Entries	33555643.186	33555643.000	33555645.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	67015563.443	66883584.000	67424256.000
Pool Paged Bytes	94005931.425	92856320.000	97460224.000

Test Log

6/7/2010 8:11:08 PM -- Prepare testing begins ...
6/7/2010 8:11:08 PM -- Creating C:\asgluns\sg1\Jetstress001001.edb.
6/7/2010 8:11:08 PM -- Database cache settings: (minimum: 32.0 MB, maximum: 256.0 MB)
6/7/2010 8:11:08 PM -- Database flush thresholds: (start: 2.5 MB, stop: 5.1 MB)
6/9/2010 1:21:41 AM -- 100.0% of 1.4 TB complete (498968822 records inserted).
6/9/2010 1:21:41 AM -- 100.0% of 1.4 TB complete (498968825 records inserted).
6/9/2010 1:21:43 AM -- Duplicating 4 database(s):
6/9/2010 10:00:15 AM -- 100.0% of 5.7 TB complete (5.7 TB duplicated).
6/9/2010 10:00:20 AM -- Attaching databases ...
6/9/2010 10:00:20 AM -- Prepare testing ends.
6/9/2010 8:09:50 PM -- Jetstress testing begins ...
6/9/2010 8:09:50 PM -- Prepare testing begins ...
6/9/2010 8:09:56 PM -- Attaching databases ...
6/9/2010 8:09:56 PM -- Prepare testing ends.
6/9/2010 8:09:56 PM -- Dispatching transactions begins ...
6/9/2010 8:09:56 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
6/9/2010 8:09:56 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
6/9/2010 8:10:02 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
6/9/2010 8:10:02 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

6/9/2010 8:10:08 PM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
6/9/2010 8:10:08 PM -- Performance Logging begins (interval: 15000 ms).
6/9/2010 8:10:09 PM -- Attaining prerequisites:
6/9/2010 8:11:36 PM -- \MSEExchange Database(JetstressWin)\Database Cache Size, Last: 1213895000.0 (Lower bound: 1207960000.0, upper bound: none)
6/9/2010 10:11:36 PM -- Performance Logging ends.
6/9/2010 10:35:29 PM -- JetInterop batch transaction stats: 46727, 46188, 46499, 46467 and 46614.
6/9/2010 10:35:30 PM -- Dispatching transactions ends.
6/9/2010 10:35:30 PM -- Shutting down databases ...
6/9/2010 10:35:52 PM -- Instance1448.1 (complete), Instance1448.2 (complete), Instance1448.3 (complete), Instance1448.4 (complete) and Instance1448.5 (complete)
6/9/2010 10:35:53 PM -- Performance Logging begins (interval: 30000 ms).
6/9/2010 10:35:53 PM -- Verifying database checksums ...
6/10/2010 1:37:00 PM -- C:\asgluns\sg1 (100% processed), C:\asgluns\sg2 (100% processed), C:\asgluns\sg3 (100% processed), C:\asgluns\sg4 (100% processed) and C:\asgluns\sg5 (100% processed)
6/10/2010 1:37:00 PM -- Performance Logging ends.
6/10/2010 1:37:00 PM -- C:\Exchange 2010\SUN159\RAID5\7300 users\DBChecksum_2010_6_9_22_35_52.blg has 1801 samples.

Stress Test Database Performance Result: SUN159

Test Summary

Overall Test Result	Pass
Machine Name	SUN159
Test Description	
Test Start Time	6/2/2010 1:01:45 AM
Test End Time	6/3/2010 1:04:19 AM
Collection Start Time	6/2/2010 1:03:47 AM
Collection End Time	6/3/2010 1:03:47 AM
Jetstress Version	14.01.0043.000
Ese Version	14.00.0639.019
Operating System	Windows Server 2008 R2 Enterprise (6.1.7600.0)
Performance Log	C:\Exchange 2010\SUN159\RAID5\7300 users\Stress_2010_6_2_1_1_58.blg C:\Exchange 2010\SUN159\RAID5\7300 users\DBChecksum_2010_6_3_1_4_19.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	947.531
Target Transactional I/O per Second	876
Initial Database Size (bytes)	7844178231296
Final Database Size (bytes)	7876885413888
Database Files (Count)	5

Jetstress System Parameters

Thread Count	7 (per database)
Minimum Database Cache	160.0 MB
Maximum Database Cache	1280.0 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

Instance3368.1	Log Path: C:\alogluns\log1 Database: C:\asgluns\sg1\Jetstress001001.edb
Instance3368.2	Log Path: C:\alogluns\log2 Database: C:\asgluns\sg2\Jetstress002001.edb
Instance3368.3	Log Path: C:\alogluns\log3 Database: C:\asgluns\sg3\Jetstress003001.edb
Instance3368.4	Log Path: C:\alogluns\log4 Database: C:\asgluns\sg4\Jetstress004001.edb
Instance3368.5	Log Path: C:\alogluns\log5 Database: C:\asgluns\sg5\Jetstress005001.edb

Transactional I/O Performance

MSExchange Database => Instance	I/O Database Reads /sec	I/O Database Writes /sec	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
Instance3368.1	18.588	12.324	118.925	70.630	34142.886	35386.425	0.000	3.842	0.000	47.065	0.000	5733.095
Instance3368.2	18.323	10.479	118.711	70.273	34248.154	35387.166	0.000	6.065	0.000	45.826	0.000	5900.408
Instance3368.3	15.571	9.211	118.641	70.441	34566.395	35383.720	0.000	3.834	0.000	47.474	0.000	5702.400
Instance3368.4	18.338	9.323	119.527	70.763	34193.139	35394.819	0.000	5.608	0.000	46.357	0.000	5853.283
Instance3368.5	15.565	8.022	119.010	70.611	34538.914	35382.561	0.000	3.891	0.000	47.741	0.000	5668.036

Background Database Maintenance I/O Performance

MSExchange Database => Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance3368.1	13.813	261876.408
Instance3368.2	16.828	261900.644
Instance3368.3	21.836	261867.221
Instance3368.4	16.829	261860.511
Instance3368.5	21.834	261870.606

Log Replication I/O Performance

MSExchange Database => Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance3368.1	1.115	231680.582
Instance3368.2	1.116	231743.457
Instance3368.3	1.118	231590.354
Instance3368.4	1.121	231678.547
Instance3368.5	1.117	231752.290

Total I/O Performance

MSExch Database e => Instances	I/O Data base Read s Aver age Laten cy (mse c)	I/O Data base Write s Aver age Laten cy (mse c)	I/O Datab ase Read s/sec	I/O Datab ase Write s/sec	I/O Datab ase Read s Avera ge Bytes	I/O Datab ase Write s Avera ge Bytes	I/O Log Rea ds Aver age Late ncy (mse c)	I/O Log Writ es Aver age Late ncy (mse c)	I/O Log Read s/sec	I/O Log Write s/sec	I/O Log Reads Avera ge Bytes	I/O Log Writes Avera ge Bytes
Instance 3368.1	18.58 8	12.32 4	132.7 37	70.63 0	57840 .609	35386 .425	28.4 61	3.84 2	1.115	47.06 5	23168 0.582	5733. 095
Instance 3368.2	18.32 3	10.47 9	135.5 39	70.27 3	62512 .171	35387 .166	65.2 68	6.06 5	1.116	45.82 6	23174 3.457	5900. 408
Instance 3368.3	15.57 1	9.211	140.4 77	70.44 1	69898 .093	35383 .720	28.5 54	3.83 4	1.118	47.47 4	23159 0.354	5702. 400
Instance 3368.4	18.33 8	9.323	136.3 56	70.76 3	62292 .341	35394 .819	65.2 54	5.60 8	1.121	46.35 7	23167 8.547	5853. 283
Instance 3368.5	15.56 5	8.022	140.8 44	70.61 1	69779 .936	35382 .561	28.1 33	3.89 1	1.117	47.74 1	23175 2.290	5668. 036

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.121	0.000	4.023
Available MBytes	29209.067	29174.000	29413.000
Free System Page Table Entries	33555643.446	33555637.000	33555645.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	70570909.730	69517312.000	71409664.000
Pool Paged Bytes	99661286.364	97742848.000	102895616.000
Database Page Fault Stalls/sec	0.587	0.000	191.051

Test Log

6/2/2010 1:01:45 AM -- Jetstress testing begins ...
6/2/2010 1:01:45 AM -- Prepare testing begins ...
6/2/2010 1:01:51 AM -- Attaching databases ...
6/2/2010 1:01:51 AM -- Prepare testing ends.
6/2/2010 1:01:51 AM -- Dispatching transactions begins ...
6/2/2010 1:01:51 AM -- Database cache settings: (minimum: 160.0 MB, maximum:
1.2 GB)
6/2/2010 1:01:51 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6
MB)
6/2/2010 1:01:58 AM -- Database read latency thresholds: (average: 20
msec/read, maximum: 200 msec/read).
6/2/2010 1:01:58 AM -- Log write latency thresholds: (average: 10 msec/write,

maximum: 200 msec/write).
 6/2/2010 1:02:05 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%,
 Replaces 5%, Reads 35%, Lazy Commits 70%.
 6/2/2010 1:02:05 AM -- Performance logging begins (interval: 15000 ms).
 6/2/2010 1:02:05 AM -- Attaining prerequisites:
 6/2/2010 1:03:47 AM -- \MSEExchange Database(JetstressWin)\Database Cache Size,
 Last: 1211335000.0 (lower bound: 1207960000.0, upper bound: none)
 6/3/2010 1:03:48 AM -- Performance logging ends.
 6/3/2010 1:03:57 AM -- JetInterop batch transaction stats: 466887, 465138,
 466087, 467590 and 466882.
 6/3/2010 1:03:58 AM -- Dispatching transactions ends.
 6/3/2010 1:03:58 AM -- Shutting down databases ...
 6/3/2010 1:04:19 AM -- Instance3368.1 (complete), Instance3368.2 (complete),
 Instance3368.3 (complete), Instance3368.4 (complete) and Instance3368.5
 (complete)
 6/3/2010 1:04:20 AM -- Performance logging begins (interval: 30000 ms).
 6/3/2010 1:04:20 AM -- Verifying database checksums ...
 6/3/2010 4:11:35 PM -- C:\asgluns\sg1 (100% processed), C:\asgluns\sg2 (100%
 processed), C:\asgluns\sg3 (100% processed), C:\asgluns\sg4 (100% processed)
 and C:\asgluns\sg5 (100% processed)
 6/3/2010 4:11:35 PM -- Performance logging ends.
 6/3/2010 4:11:35 PM -- C:\Exchange 2010\SUN159\RAID5\7300
 users\DBChecksum_2010_6_3_1_4_19.blg has 1813 samples.
 6/3/2010 4:11:43 PM -- C:\Exchange 2010\SUN159\RAID5\7300
 users\DBChecksum_2010_6_3_1_4_19.html is saved.
 6/3/2010 4:11:43 PM -- Verifying log checksums ...
 6/3/2010 4:11:45 PM -- C:\alogluns\log1 (11 log(s) processed),
 C:\alogluns\log2 (12 log(s) processed), C:\alogluns\log3 (12 log(s)
 processed), C:\alogluns\log4 (13 log(s) processed) and C:\alogluns\log5 (13
 log(s) processed)
 6/3/2010 4:11:45 PM -- C:\Exchange 2010\SUN159\RAID5\7300
 users\Stress_2010_6_2_1_1_58.blg has 5758 samples.
 6/3/2010 4:11:45 PM -- Creating test report ...
 6/3/2010 4:12:16 PM -- Instance3368.1 has 18.6 for I/O Database Reads Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.1 has 3.8 for I/O Log Writes Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.1 has 3.8 for I/O Log Reads Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.2 has 18.3 for I/O Database Reads Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.2 has 6.1 for I/O Log Writes Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.2 has 6.1 for I/O Log Reads Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.3 has 15.6 for I/O Database Reads Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.3 has 3.8 for I/O Log Writes Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.3 has 3.8 for I/O Log Reads Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.4 has 18.3 for I/O Database Reads Average
 Latency.
 6/3/2010 4:12:16 PM -- Instance3368.4 has 5.6 for I/O Log Writes Average
 Latency.

6/3/2010 4:12:16 PM -- Instance3368.4 has 5.6 for I/O Log Reads Average Latency.
 6/3/2010 4:12:16 PM -- Instance3368.5 has 15.6 for I/O Database Reads Average Latency.
 6/3/2010 4:12:16 PM -- Instance3368.5 has 3.9 for I/O Log Writes Average Latency.
 6/3/2010 4:12:16 PM -- Instance3368.5 has 3.9 for I/O Log Reads Average Latency.
 6/3/2010 4:12:16 PM -- Test has 191.050837207238 Maximum Database Page Fault Stalls/sec.
 6/3/2010 4:12:16 PM -- Test has 73 Database Page Fault Stalls/sec samples higher than 0.
 6/3/2010 4:12:16 PM -- C:\Exchange 2010\SUN159\RAID5\7300 users\Stress_2010_6_2_1_1_58.xml has 5751 samples queried.

Stress Test Database Checksums Result: SUN159

Checksum Statistics - All

Database	Seen pages	Bad pages	Correctable pages	Wrong page-number pages	File length / seconds taken
C:\asgluns\sg1\Jetstress001001.edb	48076130	0	0	0	1502379 MBytes / 54434 sec
C:\asgluns\sg2\Jetstress002001.edb	48076642	0	0	0	1502395 MBytes / 47996 sec
C:\asgluns\sg3\Jetstress003001.edb	48076898	0	0	0	1502403 MBytes / 46208 sec
C:\asgluns\sg4\Jetstress004001.edb	48077410	0	0	0	1502419 MBytes / 47965 sec
C:\asgluns\sg5\Jetstress005001.edb	48076386	0	0	0	1502387 MBytes / 46368 sec
(Sum)	240383466	0	0	0	7511983 MBytes / 54434 sec

Disk Subsystem Performance of Checksum

Logical Disk	Avg. Disk sec/Read	Avg. Disk sec/Write	Disk Reads/sec	Disk Writes/sec	Avg. Disk Bytes/Read
C:\asgluns\sg1	0.131	0.000	441.025	0.000	65535.987
C:\asgluns\sg2	0.101	0.000	500.886	0.000	65536.000
C:\asgluns\sg3	0.101	0.000	520.242	0.000	65536.000
C:\asgluns\sg4	0.099	0.000	501.260	0.000	65536.000
C:\asgluns\sg5	0.100	0.000	518.174	0.000	65536.000

Memory System Performance of Checksum

Counter	Average	Minimum	Maximum
% Processor Time	1.030	0.000	3.339
Available MBytes	30531.277	30518.000	30545.000
Free System Page Table Entries	33555643.179	33555641.000	33555645.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	70667872.300	70610944.000	70840320.000
Pool Paged Bytes	99926004.704	98525184.000	102813696.000

Test Log

6/2/2010 1:01:45 AM -- Jetstress testing begins ...
6/2/2010 1:01:45 AM -- Prepare testing begins ...
6/2/2010 1:01:51 AM -- Attaching databases ...
6/2/2010 1:01:51 AM -- Prepare testing ends.
6/2/2010 1:01:51 AM -- Dispatching transactions begins ...
6/2/2010 1:01:51 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
6/2/2010 1:01:51 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
6/2/2010 1:01:58 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
6/2/2010 1:01:58 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
6/2/2010 1:02:05 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
6/2/2010 1:02:05 AM -- Performance logging begins (interval: 15000 ms).
6/2/2010 1:02:05 AM -- Attaining prerequisites:
6/2/2010 1:03:47 AM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1211335000.0 (lower bound: 1207960000.0, upper bound: none)
6/3/2010 1:03:48 AM -- Performance logging ends.
6/3/2010 1:03:57 AM -- JetInterop batch transaction stats: 466887, 465138, 466087, 467590 and 466882.
6/3/2010 1:03:58 AM -- Dispatching transactions ends.
6/3/2010 1:03:58 AM -- Shutting down databases ...
6/3/2010 1:04:19 AM -- Instance3368.1 (complete), Instance3368.2 (complete), Instance3368.3 (complete), Instance3368.4 (complete) and Instance3368.5 (complete)
6/3/2010 1:04:20 AM -- Performance logging begins (interval: 30000 ms).
6/3/2010 1:04:20 AM -- Verifying database checksums ...
6/3/2010 4:11:35 PM -- C:\asgluns\sg1 (100% processed), C:\asgluns\sg2 (100% processed), C:\asgluns\sg3 (100% processed), C:\asgluns\sg4 (100% processed) and C:\asgluns\sg5 (100% processed)
6/3/2010 4:11:35 PM -- Performance logging ends.
6/3/2010 4:11:35 PM -- C:\Exchange 2010\SUN159\RAID5\7300 users\DBChecksum_2010_6_3_1_4_19.blg has 1813 samples.

Database Backup Test Result: SUN159

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance3808.1	1495667.09	12:38:12	32.88
Instance3808.2	1495667.09	11:39:17	35.65
Instance3808.3	1495651.09	12:38:45	32.85
Instance3808.4	1495659.09	11:37:44	35.73
Instance3808.5	1495659.09	12:38:23	32.87

Jetstress System Parameters

Thread Count	7 (per database)
Minimum Database Cache	160.0 MB
Maximum Database Cache	1280.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%

Database Configuration

Instance3808.1	Log Path: C:\alogluns\log1 Database: C:\asgluns\sg1\Jetstress001001.edb
Instance3808.2	Log Path: C:\alogluns\log2 Database: C:\asgluns\sg2\Jetstress002001.edb
Instance3808.3	Log Path: C:\alogluns\log3 Database: C:\asgluns\sg3\Jetstress003001.edb
Instance3808.4	Log Path: C:\alogluns\log4 Database: C:\asgluns\sg4\Jetstress004001.edb
Instance3808.5	Log Path: C:\alogluns\log5 Database: C:\asgluns\sg5\Jetstress005001.edb

Transactional I/O Performance

<i>MSExch Database => Instances</i>	<i>I/O Data base Reads Average Laten cy (mse c)</i>	<i>I/O Data base Writes Average Laten cy (mse c)</i>	<i>I/O Datab ase Reads /sec</i>	<i>I/O Datab ase Writes /sec</i>	<i>I/O Datab ase Reads Avera ge Bytes</i>	<i>I/O Data base Write s Aver age Bytes</i>	<i>I/O Log Rea ds Aver age Late ncy (mse c)</i>	<i>I/O Log Writ es Aver age Late ncy (mse c)</i>	<i>I/O Log Reads /sec</i>	<i>I/O Log Writes /sec</i>	<i>I/O Log Rea ds Aver age Byte s</i>	<i>I/O Log Writ es Aver age Byte s</i>
Instance 3808.1	11.41	0.000	131.399	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00
Instance 3808.2	10.53	0.000	142.179	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00
Instance 3808.3	11.11	0.000	131.334	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00
Instance 3808.4	10.38	0.000	142.633	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00
Instance 3808.5	11.10	0.000	131.441	0.000	262144.000	0.000	0.00	0.00	0.000	0.000	0.00	0.00

Host System Performance

<i>Counter</i>	<i>Average</i>	<i>Minimum</i>	<i>Maximum</i>
% Processor Time	0.476	0.000	3.696
Available MBytes	30603.351	30594.000	30607.000
Free System Page Table Entries	33555643.214	33555643.000	33555645.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	66729673.288	66727936.000	66760704.000
Pool Paged Bytes	94861990.164	94482432.000	98320384.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

6/10/2010 5:04:27 PM -- Jetstress testing begins ...
6/10/2010 5:04:27 PM -- Prepare testing begins ...
6/10/2010 5:04:32 PM -- Attaching databases ...
6/10/2010 5:04:32 PM -- Prepare testing ends.
6/10/2010 5:04:42 PM -- Performance Logging begins (interval: 30000 ms).
6/10/2010 5:04:42 PM -- Backing up databases ...
6/11/2010 5:43:28 AM -- Performance Logging ends.
6/11/2010 5:43:28 AM -- Instance3808.1 (100% processed), Instance3808.2 (100% processed), Instance3808.3 (100% processed), Instance3808.4 (100% processed) and Instance3808.5 (100% processed)
6/11/2010 5:43:28 AM -- C:\Exchange 2010\SUN159\RAID5\7300 users\DatabaseBackup_2010_6_10_17_4_32.blg has 1516 samples.
6/11/2010 5:43:28 AM -- Creating test report ...

Soft Recovery Test Result: SUN159

Soft Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance3168.1	505	830.6078589
Instance3168.2	505	722.9364698
Instance3168.3	500	814.4306305
Instance3168.4	511	753.9649243
Instance3168.5	501	813.9002295

Database Configuration

Instance3168.1	Log Path: C:\alogluns\log1 Database: C:\asgluns\sg1\Jetstress001001.edb
Instance3168.2	Log Path: C:\alogluns\log2 Database: C:\asgluns\sg2\Jetstress002001.edb
Instance3168.3	Log Path: C:\alogluns\log3 Database: C:\asgluns\sg3\Jetstress003001.edb
Instance3168.4	Log Path: C:\alogluns\log4 Database: C:\asgluns\sg4\Jetstress004001.edb
Instance3168.5	Log Path: C:\alogluns\log5 Database: C:\asgluns\sg5\Jetstress005001.edb

Transactional I/O Performance

MSExchange Database => Instance	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/s/sec	I/O Database Writes/s/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/s/sec	I/O Log Writes/s/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance3168.1	25.512	18.904	657.023	3.656	36921.681	32290.796	18.702	0.000	5.486	0.000	22905.4503	0.000
Instance3168.2	25.179	15.878	745.138	4.190	36332.820	32401.877	15.029	0.000	6.281	0.000	22991.3735	0.000
Instance3168.3	24.843	16.991	686.945	3.697	36448.530	32605.782	19.587	0.002	5.555	0.015	23129.3578	2.535
Instance3168.4	25.236	16.647	738.809	4.070	36018.823	32768.000	15.175	0.000	6.105	0.000	23236.5997	0.000
Instance3168.5	25.414	16.061	679.852	3.662	36036.322	32768.000	20.398	0.000	5.493	0.000	23237.5663	0.000

Background Database Maintenance I/O Performance

MSExchange Database => Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance3168.1	19.296	262014.169
Instance3168.2	23.006	261851.476
Instance3168.3	21.318	261877.403
Instance3168.4	23.131	261805.049
Instance3168.5	21.294	261791.820

Total I/O Performance

MSExch Database e=> Instances	I/O Data base Read s Aver age Laten cy (mse c)	I/O Data base Write s Aver age Laten cy (mse c)	I/O Datab ase Read s/sec	I/O Datab ase Write s/sec	I/O Datab ase Read s Avera ge Bytes	I/O Datab ase Write s Avera ge Bytes	I/O Log Rea ds Aver age Late ncy (mse c)	I/O Log Writ es Aver age Late ncy (mse c)	I/O Log Read s/sec	I/O Log Write s/sec	I/O Log Reads Avera ge Bytes	I/O Log Writes Avera ge Bytes
Instance 3168.1	25.51 2	18.90 4	676.3 19	3.656	43343 .740	32290 .796	18.7 02	0.00 0	5.486	0.000	22905 4.503	0.00 0
Instance 3168.2	25.17 9	15.87 8	768.1 44	4.190	43087 .037	32401 .877	15.0 29	0.00 0	6.281	0.000	22991 3.735	0.00 0
Instance 3168.3	24.84 3	16.99 1	708.2 63	3.697	43233 .689	32605 .782	19.5 87	0.00 2	5.555	0.015	23129 3.578	2.53 5
Instance 3168.4	25.23 6	16.64 7	761.9 40	4.070	42873 .299	32768 .000	15.1 75	0.00 0	6.105	0.000	23236 5.997	0.00 0
Instance 3168.5	25.41 4	16.06 1	701.1 45	3.662	42892 .439	32768 .000	20.3 98	0.00 0	5.493	0.000	23237 5.663	0.00 0

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	2.233	0.000	13.132
Available MBytes	29208.478	29175.000	30383.000
Free System Page Table Entries	33555645.000	33555645.000	33555645.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	72853137.932	68096000.000	73506816.000
Pool Paged Bytes	95203407.150	95182848.000	95248384.000
Database Page Fault Stalls/sec	0.001	0.000	0.249

Test Log

6/11/2010 6:19:46 AM -- Jetstress testing begins ...
6/11/2010 6:19:46 AM -- Prepare testing begins ...
6/11/2010 6:19:51 AM -- Attaching databases ...
6/11/2010 6:19:51 AM -- Prepare testing ends.
6/11/2010 6:19:51 AM -- Dispatching transactions begins ...
6/11/2010 6:19:51 AM -- Database cache settings: (minimum: 160.0 MB, maximum:
1.2 GB)
6/11/2010 6:19:51 AM -- Database flush thresholds: (start: 12.8 MB, stop:
25.6 MB)
6/11/2010 6:19:58 AM -- Database read latency thresholds: (average: 20
msec/read, maximum: 100 msec/read).
6/11/2010 6:19:58 AM -- Log write latency thresholds: (average: 10 msec/write,

maximum: 100 msec/write).

6/11/2010 6:20:02 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

6/11/2010 6:20:02 AM -- Performance logging begins (interval: 15000 ms).

6/11/2010 6:20:02 AM -- Generating log files ...

6/11/2010 7:20:42 AM -- C:\alogluns\log1 (101.2% generated), C:\alogluns\log2 (101.2% generated), C:\alogluns\log3 (100.2% generated), C:\alogluns\log4 (102.4% generated) and C:\alogluns\log5 (100.4% generated)

6/11/2010 7:20:42 AM -- Performance logging ends.

6/11/2010 7:20:42 AM -- JetInterop batch transaction stats: 21984, 21674, 21800, 21941 and 21735.

6/11/2010 7:20:42 AM -- Dispatching transactions ends.

6/11/2010 7:20:42 AM -- Shutting down databases ...

6/11/2010 7:20:45 AM -- Instance3168.1 (complete), Instance3168.2 (complete), Instance3168.3 (complete), Instance3168.4 (complete) and Instance3168.5 (complete)

6/11/2010 7:20:45 AM -- C:\Exchange 2010\SUN159\RAID5\7300 users\Performance_2010_6_11_6_19_58.blg has 242 samples.

6/11/2010 7:20:45 AM -- Creating test report ...

6/11/2010 7:20:46 AM -- Instance3168.1 has 16.0 for I/O Database Reads Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.1 has 5.5 for I/O Log Writes Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.1 has 5.5 for I/O Log Reads Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.2 has 14.7 for I/O Database Reads Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.2 has 3.6 for I/O Log Writes Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.2 has 3.6 for I/O Log Reads Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.3 has 16.3 for I/O Database Reads Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.3 has 5.0 for I/O Log Writes Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.3 has 5.0 for I/O Log Reads Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.4 has 14.7 for I/O Database Reads Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.4 has 3.6 for I/O Log Writes Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.4 has 3.6 for I/O Log Reads Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.5 has 16.3 for I/O Database Reads Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.5 has 5.5 for I/O Log Writes Average Latency.

6/11/2010 7:20:46 AM -- Instance3168.5 has 5.5 for I/O Log Reads Average Latency.

6/11/2010 7:20:46 AM -- Test has 11.0614456545642 Maximum Database Page Fault Stalls/sec.

6/11/2010 7:20:46 AM -- Test has 1 Database Page Fault Stalls/sec samples higher than 0.

6/11/2010 7:20:46 AM -- C:\Exchange 2010\SUN159\RAID5\7300 users\Performance_2010_6_11_6_19_58.xml has 241 samples queried.

6/11/2010 7:20:46 AM -- C:\Exchange_2010\SUN159\RAID5\7300
 users\Performance_2010_6_11_6_19_58.html is saved.
 6/11/2010 9:16:00 AM -- Performance logging begins (interval: 4000 ms).
 6/11/2010 9:16:00 AM -- Recovering databases ...
 6/11/2010 9:29:51 AM -- Performance logging ends.
 6/11/2010 9:29:51 AM -- Instance3168.1 (830.6078589), Instance3168.2
 (722.9364698), Instance3168.3 (814.4306305), Instance3168.4 (753.9649243) and
 Instance3168.5 (813.9002295)
 6/11/2010 9:29:52 AM -- C:\Exchange_2010\SUN159\RAID5\7300
 users\SoftRecovery_2010_6_11_9_15_56.blg has 207 samples.
 6/11/2010 9:29:52 AM -- Creating test report ...

Soft Recovery Test Performance Result: SUN159

Test Summary

Overall Test Result	Pass
Machine Name	SUN159
Test Description	
Test Start Time	6/11/2010 6:19:46 AM
Test End Time	6/11/2010 7:20:45 AM
Collection Start Time	6/11/2010 6:20:17 AM
Collection End Time	6/11/2010 7:20:35 AM
Jetstress Version	14.01.0043.000
Ese Version	14.00.0639.019
Operating System	Windows Server 2008 R2 Enterprise (6.1.7600.0)
Performance Log	C:\Exchange_2010\SUN159\RAID5\7300 users\Performance_2010_6_11_6_19_58.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	1048.152
Capacity Percentage	100%
Throughput Percentage	100%
Initial Database Size (bytes)	7841611317248
Final Database Size (bytes)	7843129655296
Database Files (Count)	5

Jetstress System Parameters

Thread Count	7 (per database)
Minimum Database Cache	160.0 MB
Maximum Database Cache	1280.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%

Database Configuration

Instance3168.1	Log Path: C:\alogluns\log1 Database: C:\asgluns\sg1\Jetstress001001.edb
Instance3168.2	Log Path: C:\alogluns\log2 Database: C:\asgluns\sg2\Jetstress002001.edb
Instance3168.3	Log Path: C:\alogluns\log3 Database: C:\asgluns\sg3\Jetstress003001.edb
Instance3168.4	Log Path: C:\alogluns\log4 Database: C:\asgluns\sg4\Jetstress004001.edb
Instance3168.5	Log Path: C:\alogluns\log5 Database: C:\asgluns\sg5\Jetstress005001.edb

Transactional I/O Performance

<i>MSExchange Database => Instance</i>	<i>I/O Database Reads Average Latency (msec)</i>	<i>I/O Database Writes Average Latency (msec)</i>	<i>I/O Database Reads /sec</i>	<i>I/O Database Writes /sec</i>	<i>I/O Database Reads Average Bytes</i>	<i>I/O Database Writes Average Bytes</i>	<i>I/O Log Reads Average Latency (msec)</i>	<i>I/O Log Writes Average Latency (msec)</i>	<i>I/O Log Reads /sec</i>	<i>I/O Log Writes /sec</i>	<i>I/O Log Reads Average Bytes</i>	<i>I/O Log Writes Average Bytes</i>
Instance 3168.1	16.009	11.994	134.061	78.411	32770.354	36553.580	0.000	5.498	0.000	51.426	0.000	5844.881
Instance 3168.2	14.705	9.115	131.100	76.738	32770.739	36507.788	0.000	3.649	0.000	52.635	0.000	5746.565
Instance 3168.3	16.339	10.470	132.053	77.151	32770.453	36574.694	0.000	4.959	0.000	51.285	0.000	5814.319
Instance 3168.4	14.683	7.639	133.011	78.100	32770.474	36593.215	0.000	3.631	0.000	53.128	0.000	5734.615
Instance 3168.5	16.348	8.131	131.004	76.522	32769.727	36567.664	0.000	5.483	0.000	50.936	0.000	5891.398

Host System Performance

<i>Counter</i>	<i>Average</i>	<i>Minimum</i>	<i>Maximum</i>
% Processor Time	0.986	0.000	3.527
Available MBytes	29217.818	29188.000	30136.000
Free System Page Table Entries	33555643.760	33555642.000	33555645.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	67013386.579	66760704.000	67715072.000
Pool Paged Bytes	94993602.645	94924800.000	95100928.000
Database Page Fault Stalls/sec	0.046	0.000	11.061

Test Log

6/11/2010 6:19:46 AM -- Jetstress testing begins ...
6/11/2010 6:19:46 AM -- Prepare testing begins ...
6/11/2010 6:19:51 AM -- Attaching databases ...
6/11/2010 6:19:51 AM -- Prepare testing ends.
6/11/2010 6:19:51 AM -- Dispatching transactions begins ...
6/11/2010 6:19:51 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
6/11/2010 6:19:51 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
6/11/2010 6:19:58 AM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
6/11/2010 6:19:58 AM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
6/11/2010 6:20:02 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
6/11/2010 6:20:02 AM -- Performance logging begins (interval: 15000 ms).
6/11/2010 6:20:02 AM -- Generating log files ...
6/11/2010 7:20:42 AM -- C:\al o g l u n s \ l o g 1 (101.2% generated), C:\al o g l u n s \ l o g 2 (101.2% generated), C:\al o g l u n s \ l o g 3 (100.2% generated), C:\al o g l u n s \ l o g 4 (102.4% generated) and C:\al o g l u n s \ l o g 5 (100.4% generated)
6/11/2010 7:20:42 AM -- Performance logging ends.
6/11/2010 7:20:42 AM -- JetInterop batch transaction stats: 21984, 21674, 21800, 21941 and 21735.
6/11/2010 7:20:42 AM -- Dispatching transactions ends.
6/11/2010 7:20:42 AM -- Shutting down databases ...
6/11/2010 7:20:45 AM -- Instance3168.1 (complete), Instance3168.2 (complete), Instance3168.3 (complete), Instance3168.4 (complete) and Instance3168.5 (complete)
6/11/2010 7:20:45 AM -- C:\Exchange 2010\SUN159\RAID5\7300 users\Performance_2010_6_11_6_19_58.blg has 242 samples.
6/11/2010 7:20:45 AM -- Creating test report ...
6/11/2010 7:20:46 AM -- Instance3168.1 has 16.0 for I/O Database Reads Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.1 has 5.5 for I/O Log Writes Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.1 has 5.5 for I/O Log Reads Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.2 has 14.7 for I/O Database Reads Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.2 has 3.6 for I/O Log Writes Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.2 has 3.6 for I/O Log Reads Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.3 has 16.3 for I/O Database Reads Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.3 has 5.0 for I/O Log Writes Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.3 has 5.0 for I/O Log Reads Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.4 has 14.7 for I/O Database Reads Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.4 has 3.6 for I/O Log Writes Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.4 has 3.6 for I/O Log Reads Average

Latency.
6/11/2010 7:20:46 AM -- Instance3168.5 has 16.3 for I/O Database Reads Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.5 has 5.5 for I/O Log Writes Average Latency.
6/11/2010 7:20:46 AM -- Instance3168.5 has 5.5 for I/O Log Reads Average Latency.
6/11/2010 7:20:46 AM -- Test has 11.0614456545642 Maximum Database Page Fault Stalls/sec.
6/11/2010 7:20:46 AM -- Test has 1 Database Page Fault Stalls/sec samples higher than 0.
6/11/2010 7:20:46 AM -- C:\Exchange 2010\SUN159\RAID5\7300 users\Performance_2010_6_11_6_19_58.xml has 241 samples queried.



Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.uk@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names mentioned in this document are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect and that may be configuration dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

© Hitachi Data Systems Corporation 2010. All Rights Reserved.
ESRP-056-00 May 2010