

Hitachi Adaptable Modular Storage 2300 Dynamically Provisioned 40,000 Mailbox Exchange 2010 Mailbox Resiliency Storage Solution

Tested with: ESRP – Storage Version 3.0

Test Date: January-February 2011

Notices and Disclaimer

Copyright © 2011 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	3
Disclaimer	3
Features	3
Solution Description	4
Targeted Customer Profile	10
Test Deployment	10
Replication Configuration	12
Best Practices	15
Core Storage	15
Storage-based Replication	16
Backup Strategy	16
Test Results Summary	16
Reliability	16
Storage Performance Results	16
Database Backup and Recovery Performance	20
Conclusion	20
Appendix A — RAID-5 Drive Failure and Rebuild	21
Appendix B — Test Reports	22
Performance Test Result Report: SUN170	22
Performance Test Database Checksums Result: SUN170	28
Stress Test Result Report: SUN170	30
Stress Test Database Checksums Result: SUN170	36
Backup Test Result Report: SUN170.....	38
Soft Recovery Test Result Report: SUN170	40
Soft Recovery Test Performance Result: SUN170	46

Hitachi Adaptable Modular Storage 2300 Dynamically Provisioned 40,000 Mailbox Exchange 2010 Mailbox Resiliency Storage Solution

Tested with: ESRP – Storage Version 3.0

Test Date: January-February 2011

Overview

This document provides information on a Hitachi Adaptable Modular Storage 2300 Mailbox Resiliency storage solution using Hitachi Dynamic Provisioning for Microsoft Exchange Server 2010. This solution is 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 40,000 users and eight servers. This testing used a Hitachi Adaptable Modular Storage 2300 Rev 2 (AMS 2300) storage system using Hitachi Dynamic Provisioning in a two-pool RAID-5 (8D+1P) 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 an AMS 2300.

Testing used eight Sun Fire 4270 servers, each with the following:

- 32GB of RAM
- Two quad-core Intel E5540 2.53GHz CPUs
- Two Emulex 4Gb/sec Fibre Channel adapters
- Microsoft Windows Server 2008 R2 Enterprise

This solution includes Exchange 2010 Mailbox Resiliency by using the Database Availability Group (DAG) feature. This tested configuration uses eight DAGs, each containing two database copies and two servers (one simulated). The test configuration was capable of supporting 40,000 users with a 0.12 IOPS per user profile and user mailbox size of 1GB.

An AMS 2300 with the following was used for these tests:

- 189 600GB 15k RPM SAS disks
- 16GB of cache
- Eight 8Gb/sec paths

Hitachi Adaptable Modular Storage 2300 is a 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, AMS 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 uses AMS 2300 to meet the needs of a large Exchange Server deployment.

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

Primarily, Hitachi Dynamic Provisioning is deployed to avoid the routine issue of hot spots that occur on logical units (LUs) from individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. By using many RAID groups as members of a striped dynamic provisioning pool underneath the virtual or logical volumes seen by the hosts, a host workload is distributed across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots and results in fewer mailbox moves for the Exchange administrator.

Hitachi Dynamic Provisioning 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 Database Availability Group (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 logs and uses continuous replication to 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, with each DAG configured with two servers (one simulated) to host active mailboxes in five databases.

To target the 40,000-user resiliency solution, a Hitachi Adaptable Modular Storage 2300 (AMS 2300) was configured with 189 disks (of a maximum of 240). Eight servers (one per DAG) were used, with each server configured with 5,000 mailboxes. There were 40 active databases and the simulated database copies for the tests.

Each DAG contained two copies of the five databases hosted by that DAG;

- A local, active copy on a server connected to the primary AMS 2300
- A passive copy (simulated) on another server connected to a second AMS 2300 (simulated)

This recommended configuration can support high-availability and disaster-recovery scenarios when the active and passive database copies are allocated among both DAG members and dispersed across both storage systems. 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. For more information, see the Hitachi Data Systems [Storage Systems web page](#).

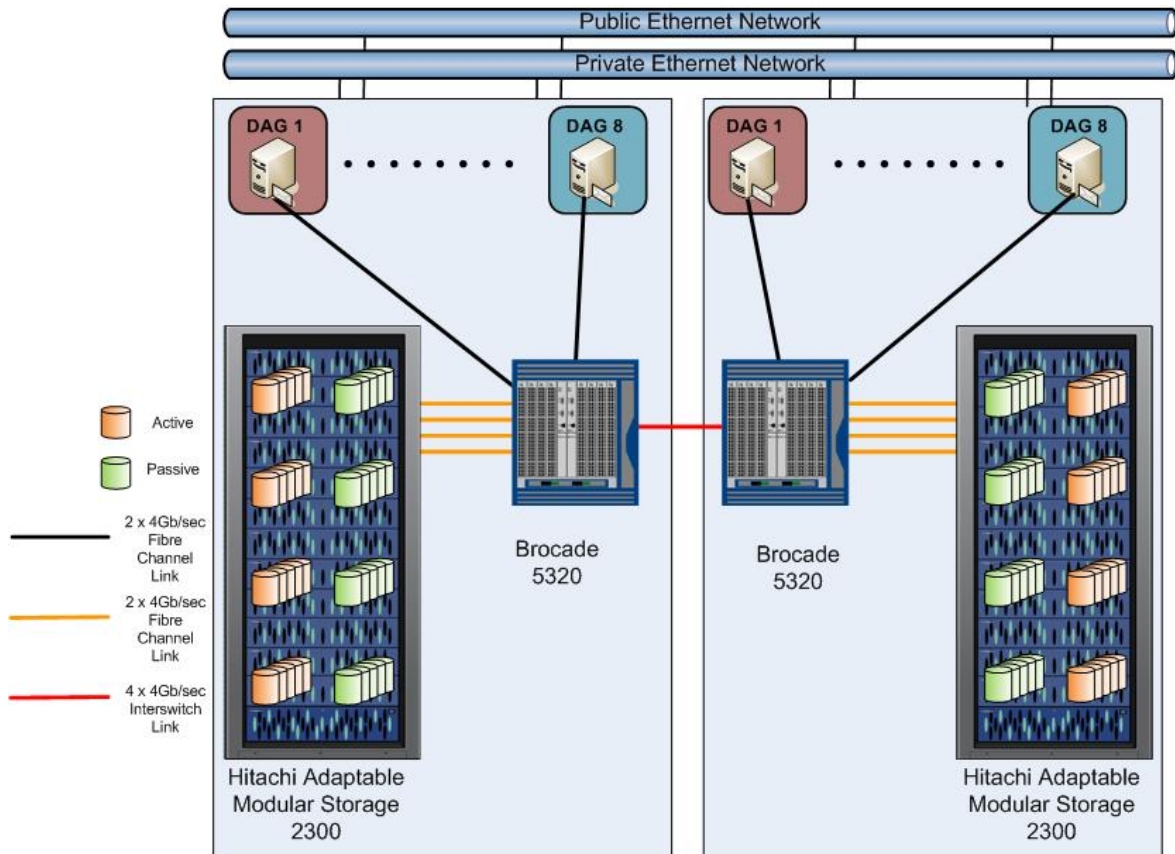



Figure 1

This solution enables organizations to consolidate Exchange Server 2010 DAG deployments on two AMS 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 disks in an AMS 2300 were organized into RAID groups for use by the databases or logs. Each set of colored disks represents a RAID group. There were 189 disks used in these tests configured as 21 RAID groups (using RAID-5 8D+1P) for the Exchange databases and logs.

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	15	16	17	18	
RKA 10	20	20	20	20	20	20	20	20	20	21	21	21	21	21	21	21	21	21	21	5
RKA 9	17	18	18	18	18	18	18	18	18	18	19	19	19	19	19	19	19	19	19	19
RKA 8	15	15	16	16	16	16	16	16	16	16	16	17	17	17	17	17	17	17	17	17
RKA 7	13	13	13	14	14	14	14	14	14	14	14	14	14	15	15	15	15	15	15	15
RKA 6	11	11	11	11	12	12	12	12	12	12	12	12	12	12	13	13	13	13	13	13
RKA 5	9	9	9	9	9	10	10	10	10	10	10	10	10	10	11	11	11	11	11	11
RKA 4	7	7	7	7	7	7	8	8	8	8	8	8	8	8	8	9	9	9	9	9
RKA 3	5	5	5	5	5	5	5	6	6	6	6	6	6	6	6	6	6	7	7	7
RKA 2	3	3	3	3	3	3	3	3	4	4	4	4	4	4	4	4	4	4	5	5
RKA 1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	2	2	3
RKA 0	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-



Disk tray RKA-0 was the internal 15-disk tray that wasn't used during these tests. Trays RKA-1 through tray RKA-10 each held 19 600GB 15K SAS disks. Tray RKA-10 had one disk used as a spare. There were actually five dense trays, but each is accessed as two separate trays, so trays 1 and 2 are dense enclosure 1, and so forth.

Two dynamic provisioning pools were created, one for the databases and the other for the logs. The database pool was created from 18 RAID-5 RAID groups and the log pool was created from three RAID-5 RAID groups. From the database pool, 40 DP-VOLs (each specified to have a 1,850GB size limit) were created for 40 databases (5 per server). From the log pool, 40 DP-VOLs (each specified to have a size limit of 185GB) were created for 40 logs (5 per server).

Table 2 outlines the port layout for the primary storage and servers. An identical configuration would be deployed on the replicated storage and servers for this solution.

Table 2. Adaptable Modular Storage 2300 Port to Server Layout

Server	Primary path	Secondary path
SUN170	0A	1A
SUN163	0B	1B
SUN164	0C	1C
SUN165	0D	1D
SUN166	1A	0A
SUN167	1B	0B
SUN171	1C	0C
SUN169	1D	0D

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

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

<i>Port</i>	<i>Database</i>	<i>DB DP-VOL</i>
0A	Database 1-5	0-4
0B	Database 6-10	5-9
0C	Database 11-15	10-14
0D	Database 16-20	15-19
1A	Database 21-25	20-24
1B	Database 26-30	25-29
1C	Database 31-35	30-34
1D	Database 36-40	35-39

Table 4 outlines the port layout with the log DP-VOL assignments for the primary storage and servers. An identical configuration would be 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>Log DP-VOL</i>
0A	Log 1-5	40-44
0B	Log 6-10	45-49
0C	Log 11-15	50-54
0D	Log 16-20	55-59
1A	Log 21-25	60-64
1B	Log 26-30	65-69
1C	Log 31-35	70-74
1D	Log 36-40	75-79

Table 5 provides the detailed specifications for the storage configuration which uses RAID-5 (8D+1P) groups and 600GB 15K disks. Dynamic provisioning pool 0 is dedicated for the databases 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-VOLs</i>	<i>Size (GB)</i>	<i>RAID Level</i>	<i>Description</i>
SUN170	0	0A/1A	0-4	1850	RAID-5	Database 1-5
SUN163	0	0B/1B	5-9	1850	RAID-5	Database 6-10
SUN164	0	0C/1C	10-14	1850	RAID-5	Database 11-15
SUN165	0	0D/1D	15-19	1850	RAID-5	Database 16-20
SUN166	0	1A/0A	20-24	1850	RAID-5	Database 21-25
SUN167	0	1B/0B	25-29	1850	RAID-5	Database 26-30
SUN171	0	1C/0C	30-34	1850	RAID-5	Database 31-35
SUN169	0	1D/0D	35-39	1850	RAID-5	Database 36-40

<i>Host</i>	<i>Pool</i>	<i>Port</i>	<i>DP-VOLs</i>	<i>Size (GB)</i>	<i>RAID Level</i>	<i>Description</i>
SUN170	1	0A/1A	40-44	185	RAID-5	Log 1-5
SUN163	1	0B/1B	45-49	185	RAID-5	Log 6-10
SUN164	1	0C/1C	50-54	185	RAID-5	Log 11-15
SUN165	1	0D/1D	55-59	185	RAID-5	Log 16-20
SUN166	1	1A/0A	60-64	185	RAID-5	Log 21-25
SUN167	1	1B/0B	65-69	185	RAID-5	Log 26-30
SUN171	1	1C/0C	70-74	185	RAID-5	Log 31-35
SUN169	1	1D/0D	75-79	185	RAID-5	Log 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 40,000 Exchange users with the following specifications:

- 16 Exchange servers (eight servers tested, eight servers simulated for the database copies)
- 8 Database Availability Groups each with two servers (one simulated) and two copies per database
- 2 Adaptable Modular Storage 2300s (one tested, using 8 8Gb/sec paths operated at 4Gbps)
- 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.

Test Deployment

The following tables summarize the testing environment.

Table 6. Simulated Exchange Configuration

<i>Number of Exchange mailboxes simulated</i>	40,000
<i>Number of database availability groups (DAGs)</i>	8
<i>Number of servers per DAG</i>	2 (1 simulated)
<i>Number of active mailboxes per server</i>	5,000
<i>Number of databases per host</i>	5
<i>Number of copies per database</i>	2
<i>Number of mailboxes per database</i>	1000
<i>Simulated profile: I/Os per second per mailbox (IOPS, include 20% headroom)</i>	0.12
<i>Database LU size</i>	1850GB
<i>Log LU size</i>	185GB
<i>Total database size for performance testing</i>	40,000GB
<i>% storage capacity used by Exchange database**</i>	52.4%

**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. Storage Hardware

<i>Storage connectivity (Fibre Channel, SAS, SATA, iSCSI)</i>	Fibre Channel
<i>Storage model and OS/firmware revision</i>	1 Hitachi Adaptable Modular Storage 2300 Firmware: 0897/A-Y WHQL listing: Hitachi Adaptable Modular Storage 2300
<i>Storage cache</i>	16GB
<i>Number of storage controllers</i>	2
<i>Number of storage ports</i>	8
<i>Maximum bandwidth of storage connectivity to host</i>	32Gb/sec (8x 4Gb/sec ports)
<i>Switch type/model/firmware revision</i>	Brocade 5320, Fabric OS v6.4.0b
<i>HBA model and firmware</i>	Emulex LightPulse LPe 11002-S FW : 2.82A3
<i>Number of HBAs per host</i>	2 dual-ported HBA per host, 1 4Gb/sec port used per HBA
<i>Host server type</i>	Sun Fire 4270 2 2.54GHz quad-core Intel Xeon CPUs, 32GB memory
<i>Total number of disks tested in solution</i>	189
<i>Maximum number of spindles that can be hosted in the storage</i>	240

Table 8. Storage Software

<i>HBA driver</i>	Storport Miniport 7.2.30.016
<i>HBA QueueTarget setting</i>	0
<i>HBA QueueDepth setting</i>	32
<i>Multipathing</i>	Hitachi Dynamic Link Manager v6.4.0-00
<i>Host OS</i>	Microsoft Windows Server 2008 R2 Enterprise
<i>ESE.dll file version</i>	14.00.0639.019
<i>Replication solution name/version</i>	N/A

Table 9. Storage Disk Configuration (Mailbox Store Disks)

<i>Disk type, speed and firmware revision</i>	SAS 600GB 15K 5C53
<i>Raw capacity per disk (GB)</i>	600GB
<i>Number of physical disks in test</i>	162 (dynamic provisioning pool)
<i>Total raw storage capacity (GB)</i>	97,200
<i>Disk slice size (GB)</i>	N/A
<i>Number of slices per LU or number of disks per LU</i>	N/A
<i>RAID level</i>	RAID-5(8D+1P) at storage level
<i>Total formatted capacity</i>	76,320GB
<i>Storage capacity utilization</i>	78.5%
<i>Database capacity utilization</i>	76.1%

Table 10. Storage Disk Configuration (Transaction Log Disks)

<i>Disk type, speed and firmware revision</i>	SAS 600GB 15K 5C53
<i>Raw capacity per disk (GB)</i>	600GB
<i>Number of spindles in test</i>	27 (dynamic provisioning pool)
<i>Total raw storage capacity (GB)</i>	16,200
<i>Disk slice size (GB)</i>	N/A
<i>Number of slices per LU or number of disks per LU</i>	N/A
<i>RAID level</i>	RAID-5 (8D+1P) at storage level
<i>Total formatted capacity</i>	12,720GB

Replication Configuration

Table 11. Replication Configuration

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

Table 12. Replicated Storage Hardware

<i>Storage connectivity (Fiber Channel, SAS, SATA, iSCSI)</i>	Fibre Channel
<i>Storage model and OS/firmware revision</i>	1 Hitachi Adaptable Modular Storage 2300 Firmware: 0897/A-Y WHQL listing: Hitachi Adaptable Modular Storage 2300
<i>Storage cache</i>	16GB
<i>Number of storage controllers</i>	2
<i>Number of storage ports</i>	8
<i>Maximum bandwidth of storage connectivity to host</i>	32Gb/sec (8 x 4Gb/sec ports)
<i>Switch type/model/firmware revision</i>	Brocade 5320, Fabric OS v6.4.0b
<i>HBA model and firmware</i>	Emulex LightPulse LPe 11002-S FW : 2.82A3
<i>Number of HBAs per host</i>	2 dual-ported HBA per host, 1 4Gb/sec port used per HBA
<i>Host server type</i>	Sun Fire 4270 2 2.54GHz quad-core Intel Xeon CPUs, 32GB memory
<i>Total number of disks tested in solution</i>	189
<i>Maximum number of spindles that can be hosted in the storage</i>	240

Table 13. Replicated Storage Software

<i>HBA driver</i>	Storport Miniport 7.2.30.016
<i>HBA QueueTarget setting</i>	0
<i>HBA QueueDepth setting</i>	32
<i>Multipathing</i>	Hitachi Dynamic Link Manager v6.4.0-00
<i>Host OS</i>	Microsoft Windows Server 2008 R2 Enterprise
<i>ESE.dll file version</i>	14.00.0639.019
<i>Replication solution name/version</i>	N/A

Table 14. Replicated Storage Disk Configuration (Mailbox Store Disks)

<i>Disk type, speed and firmware revision</i>	SAS 600GB 15K 5C53
<i>Raw capacity per disk (GB)</i>	600GB
<i>Number of physical disks in test</i>	162 (dynamic provisioning pool)
<i>Total raw storage capacity (GB)</i>	97,200
<i>Disk slice size (GB)</i>	N/A
<i>Number of slices per LU or number of disks per LU</i>	N/A
<i>RAID level</i>	RAID-5(8D+1P) at storage level
<i>Total formatted capacity</i>	76,320GB
<i>Storage capacity utilization</i>	78.5%
<i>Database capacity utilization</i>	76.1%

Table 15. Replicated Storage Disk Configuration (Transactional Log Disks)

<i>Disk type, speed and firmware revision</i>	SAS 600GB 15K 5C53
<i>Raw capacity per disk (GB)</i>	600GB
<i>Number of spindles in test</i>	27 (dynamic provisioning pool)
<i>Total raw storage capacity (GB)</i>	16,200
<i>Disk slice size (GB)</i>	N/A
<i>Number of slices per LU or number of disks per LU</i>	N/A
<i>Raid level</i>	RAID-5(8D+1P) at storage level
<i>Total formatted capacity</i>	12,720GB

Best Practices

Microsoft Exchange Server 2010 is a 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 LU, Hitachi Data Systems recommends setting the ALU to 64K for the database files and 4K for the log files.
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. Use Hitachi Dynamic Link Manager multipathing software to provide fault tolerance and high availability for host connectivity.
5. Use Hitachi Dynamic Provisioning software to simplify storage management of the Exchange database and log volumes.
6. 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.
7. Hitachi Data Systems recommends RAID-5 or RAID-10 groups for both the database pools and for the log pool. Use of RAID-10 allows more writes at a lower response time under heavier loads. RAID-10 also supports a shorter RAID group rebuild time on failure of a disk.
8. The log LUs should be at least 10 percent of the size of the database LUs.
9. Hitachi Data Systems does not recommend using LU concatenation.
10. Hitachi Data Systems recommends implementing Mailbox Resiliency using the Exchange Server 2010 Database Availability Group feature.
11. Ensure that each DAG maintains at least two database copies to provide high availability.
12. Isolate active databases and their replicated copies in separate dynamic provisioning pools or ensure that they are located on a separate AMS 2300.
13. Use fewer, larger LUs for Exchange 2010 databases (up to 2TB) with Background Database Maintenance (24x7) enabled.
14. 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.

15. 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.
16. 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.

Storage-based Replication

N/A

Backup Strategy

N/A

Test Results 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.
- If done, no errors were reported during the backup to disk test [process](#).
- No errors were reported for the database checksum on the remote storage database.

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

These individual server metrics show the sum of the input/outputs across the storage groups and the average latency across all storage groups on a per-server basis.

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

<i>Database I/O</i>	
<i>Database Disk Transfers Per Second</i>	772
<i>Database Disk Reads Per Second</i>	479
<i>Database Disk Writes Per Second</i>	293
<i>Average Database Disk Read Latency (ms)</i>	16.9
<i>Average Database Disk Write Latency (ms)</i>	7.4
<i>Transaction Log I/O</i>	
<i>Log Disk Writes Per Second</i>	216
<i>Average Log Disk Write Latency (ms)</i>	3.4

Table 17. Individual Server Metrics for Exchange Server (SUN163)

<i>Database I/O</i>	
<i>Database Disk Transfers Per Second</i>	801
<i>Database Disk Reads Per Second</i>	497
<i>Database Disk Writes Per Second</i>	304
<i>Average Database Disk Read Latency (ms)</i>	16.4
<i>Average Database Disk Write Latency (ms)</i>	7.5
<i>Transaction Log I/O</i>	
<i>Log Disk Writes Per Second</i>	221
<i>Average Log Disk Write Latency (ms)</i>	3.7

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

<i>Database I/O</i>	
<i>Database Disk Transfers Per Second</i>	793
<i>Database Disk Reads Per Second</i>	492
<i>Database Disk Writes Per Second</i>	301
<i>Average Database Disk Read Latency (ms)</i>	16.7
<i>Average Database Disk Write Latency (ms)</i>	7.3
<i>Transaction Log I/O</i>	
<i>Log Disk Writes Per Second</i>	221
<i>Average Log Disk Write Latency (ms)</i>	3.4

Table 19. Individual Server Metrics for Exchange Server (SUN165)

<i>Database I/O</i>	
<i>Database Disk Transfers Per Second</i>	801
<i>Database Disk Reads Per Second</i>	497
<i>Database Disk Writes Per Second</i>	304
<i>Average Database Disk Read Latency (ms)</i>	16.3
<i>Average Database Disk Write Latency (ms)</i>	7.4
<i>Transaction Log I/O</i>	
<i>Log Disk Writes Per Second</i>	223
<i>Average Log Disk Write Latency (ms)</i>	3.6

Table 20. Individual Server Metrics for Exchange Server (SUN166)

<i>Database I/O</i>	
<i>Database Disk Transfers Per Second</i>	783
<i>Database Disk Reads Per Second</i>	486
<i>Database Disk Writes Per Second</i>	297
<i>Average Database Disk Read Latency (ms)</i>	16.7
<i>Average Database Disk Write Latency (ms)</i>	7.3
<i>Transaction Log I/O</i>	
<i>Log Disk Writes Per Second</i>	218
<i>Average Log Disk Write Latency (ms)</i>	3.5

Table 21. Individual Server Metrics for Exchange Server (SUN167)

<i>Database I/O</i>	
<i>Database Disk Transfers Per Second</i>	800
<i>Database Disk Reads Per Second</i>	496
<i>Database Disk Writes Per Second</i>	304
<i>Average Database Disk Read Latency (ms)</i>	16.4
<i>Average Database Disk Write Latency (ms)</i>	7.5
<i>Transaction Log I/O</i>	
<i>Log Disk Writes Per Second</i>	221
<i>Average Log Disk Write Latency (ms)</i>	3.6

Table 22. Individual Server Metrics for Exchange Server (SUN171)

<i>Database I/O</i>	
<i>Database Disk Transfers Per Second</i>	804
<i>Database Disk Reads Per Second</i>	499
<i>Database Disk Writes Per Second</i>	305
<i>Average Database Disk Read Latency (ms)</i>	16.5
<i>Average Database Disk Write Latency (ms)</i>	7.2
<i>Transaction Log I/O</i>	
<i>Log Disk Writes Per Second</i>	224
<i>Average Log Disk Write Latency (ms)</i>	3.4

Table 23. Individual Server Metrics for Exchange Server (SUN169)

<i>Database I/O</i>	
<i>Database Disk Transfers Per Second</i>	793
<i>Database Disk Reads Per Second</i>	492
<i>Database Disk Writes Per Second</i>	301
<i>Average Database Disk Read Latency (ms)</i>	16.3
<i>Average Database Disk Write Latency (ms)</i>	7.6
<i>Transaction Log I/O</i>	
<i>Log Disk Writes Per Second</i>	221
<i>Average Log Disk Write Latency (ms)</i>	3.6

Aggregate Performance Across All Servers Metric

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

Table 24. Aggregate Performance for Exchange Server 2010

<i>Database I/O</i>	
<i>Database Disk Transfers Per Second</i>	6347.977
<i>Database Disk Reads Per Second</i>	3938.140
<i>Database Disk Writes Per Second</i>	2409.837
<i>Average Database Disk Read Latency (ms)</i>	16.523
<i>Average Database Disk Write Latency (ms)</i>	7.413
<i>Transaction Log I/O</i>	
<i>Log Disk Writes Per Second</i>	1764.137
<i>Average Log Disk Write Latency (ms)</i>	3.520

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 34. Database Read-only Performance

<i>MB Read Per Second Per Database</i>	34.95
<i>MB Read Per Second Total Per Server</i>	174.763

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 storage group. Each log file is 1MB in size.

Table 35. Transaction Log Recovery/Replay Performance

<i>Average Time to Play One Log File (sec)</i>	1.645591
--	----------

Conclusion

This document details a tested and robust Exchange Server 2010 Resiliency solution capable of supporting 40,000 users with a 0.12 IOPS per user profile and user mailbox size of 1GB using eight DAGs, each configured with 2 server nodes (one simulated).

A Hitachi Adaptable Modular Storage 2300 (AMS 2300), with 16GB of cache and eight 8Gb/sec Fibre Channel host paths, using Hitachi Dynamic Provisioning (with two pools) and 189 600GB 15K RPM SAS disks in a RAID-5 configuration was used for these tests.

Testing confirmed that an AMS 2300 is more than capable of delivering the IOPS and capacity requirements needed to support the active and replicated databases for 40,000 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 VSS snapshot or clone backups, and relies on the built-in Mailbox Resiliency features of Exchange Server 2010 coupled with Adaptable Modular Storage RAID technology to provide high-availability and protection from logical and physical failures. Adding additional protection requirements may affect performance and capacity requirements of the underlying storage configuration, and as such need to be factored into the storage design accordingly.

For more information to 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 the 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-10 (for example, 4D+4D).

RAID-5 is a more capacity-efficient RAID level than the others. It loses only 12.5 percent of the usable space when using 8D+1P. This compares to 33 percent for 4D+2P or 50 percent for 4D+4D.

One downside with the use of parity RAID instead of mirrored and striped (RAID-10) is that the internal disk write penalty for writes is higher. For SAS or Fibre Channel disks, RAID-5 requires four physical disk I/Os on the backend for every host write. In comparison, RAID-10 requires two physical I/Os and RAID-6 requires six physical I/Os for each host write.

The other downside is the rebuild time for the RAID group after a sudden disk failure. The Hitachi Adaptable Modular Storage 2000 family always scans the storage system looking for soft fails, because excessive soft fails is a predictor of a hard failure. If the number of soft fails exceeds the user-set failure threshold in a 24-hour period, an Adaptable Modular Storage 2000 family storage system does the following, in order:

1. Executes a disk-to-disk copy to a global hot spare to avoid a RAID-5 or RAID-6 rebuild.
2. Marks the disk as failed.
3. Replaces the disk.

If a hard fail does occur, the following happens:

- When using RAID-10, the contents of the good disk are mirrored onto a spare disk. These “hot spares” are user-defined in several disk enclosures on a storage system.
- When using RAID-5 or 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 while there are still ongoing host I/Os. This setting has three levels: aggressive, moderate, and background.

Lab tests were conducted on a RAID-6 group using Fibre Channel disks with an aggressive restore option setting. A RAID-6 (8D+2P) group corrective copy operation takes about 30 minutes to complete without any host workload on the LUs from that RAID group. When there was a sustained 100 percent sequential write workloads to the LUs from the same RAID group, the rebuild time increased to 18 hours. The host performance on a LU from that RAID group was measured at 154MB/sec (normal state) and 95MB/sec (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 Report: SUN170

Test Summary

Overall Test Result	Pass
Machine Name	SUN170
Test Description	
Test Start Time	1/26/2011 8:17:31 PM
Test End Time	1/26/2011 11:06:12 PM
Collection Start Time	1/26/2011 8:19:22 PM
Collection End Time	1/26/2011 10:19:17 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:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Performance Test\Performance_2011_1_26_20_17_42.blg C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Performance Test\DBCchecksum_2011_1_26_23_6_12.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	772.278
Target Transactional I/O per Second	600
Initial Database Size (bytes)	6489943900160
Final Database Size (bytes)	6493123182592
Database Files (Count)	5

Jetstress System Parameters

<i>Thread Count</i>	8 (per database)
<i>Minimum Database Cache</i>	160.0 MB
<i>Maximum Database Cache</i>	1280.0 MB
<i>Insert Operations</i>	40%
<i>Delete Operations</i>	20%
<i>Replace Operations</i>	5%
<i>Read Operations</i>	35%
<i>Lazy Commits</i>	70%
<i>Run Background Database Maintenance</i>	True
<i>Number of Copies per Database</i>	2

Database Configuration

<i>Instance2880.1</i>	Log Path: C:\logluns\log1 Database: C:\dbluns\db1\Jetstress001001.edb
<i>Instance2880.2</i>	Log Path: C:\logluns\log2 Database: C:\dbluns\db2\Jetstress002001.edb
<i>Instance2880.3</i>	Log Path: C:\logluns\log3 Database: C:\dbluns\db3\Jetstress003001.edb
<i>Instance2880.4</i>	Log Path: C:\logluns\log4 Database: C:\dbluns\db4\Jetstress004001.edb
<i>Instance2880.5</i>	Log Path: C:\logluns\log5 Database: C:\dbluns\db5\Jetstress005001.edb

Transactional I/O Performance

<i>MSExchange Database => Instances</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>
Instance2880.1	17.217	8.909	95.477	58.408	36316.110	35133.030	0.000	3.537	0.000	42.702	0.000	5211.060
Instance2880.2	16.815	7.682	95.765	58.670	36155.901	35101.811	0.000	3.460	0.000	43.359	0.000	5229.983
Instance2880.3	16.842	7.644	95.910	58.810	36187.875	35111.228	0.000	3.504	0.000	43.003	0.000	5251.407
Instance2880.4	16.779	6.503	96.281	58.907	36269.902	35072.858	0.000	3.317	0.000	43.529	0.000	5154.648
Instance2880.5	16.780	6.349	95.572	58.478	36114.530	35149.718	0.000	3.235	0.000	43.008	0.000	5184.181

Background Database Maintenance I/O Performance

<i>MSEExchange Database ==> Instances</i>	<i>Database Maintenance IO Reads/sec</i>	<i>Database Maintenance IO Reads Average Bytes</i>
Instance2880.1	22.474	261912.900
Instance2880.2	22.810	261809.931
Instance2880.3	22.803	261852.535
Instance2880.4	22.755	261851.832
Instance2880.5	22.813	261850.451

Log Replication I/O Performance

<i>MSEExchange Database ==> Instances</i>	<i>I/O Log Reads/sec</i>	<i>I/O Log Reads Average Bytes</i>
Instance2880.1	0.914	224744.665
Instance2880.2	0.932	227788.909
Instance2880.3	0.927	225234.508
Instance2880.4	0.918	225240.647
Instance2880.5	0.915	225753.304

Total I/O Performance

<i>MSExchange Database => Instances</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>
Instance2880.1	17.217	8.909	117.951	58.408	79299.740	35133.030	15.892	3.537	0.914	42.702	224744.665	5211.060
Instance2880.2	16.815	7.682	118.575	58.670	79564.499	35101.811	9.982	3.460	0.932	43.359	227788.909	5229.983
Instance2880.3	16.842	7.644	118.713	58.810	79535.278	35111.228	11.999	3.504	0.927	43.003	225234.508	5251.407
Instance2880.4	16.779	6.503	119.036	58.907	79392.875	35072.858	10.801	3.317	0.918	43.529	225240.647	5154.648
Instance2880.5	16.780	6.349	118.385	58.478	79614.396	35149.718	10.177	3.235	0.915	43.008	225753.304	5184.181

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.092	0.000	16.695
Available MBytes	29444.885	29436.000	29465.000
Free System Page Table Entries	33555644.706	33555637.000	33555645.000
Transition Pages Repurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	66940047.232	66854912.000	66994176.000
Pool Paged Bytes	101358187.958	101310464.000	101425152.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

1/26/2011 8:17:31 PM -- Jetstress testing begins ...
 1/26/2011 8:17:31 PM -- Prepare testing begins ...
 1/26/2011 8:17:36 PM -- Attaching databases ...
 1/26/2011 8:17:36 PM -- Prepare testing ends.
 1/26/2011 8:17:36 PM -- Dispatching transactions begins ...
 1/26/2011 8:17:36 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
 1/26/2011 8:17:36 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
 1/26/2011 8:17:42 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
 1/26/2011 8:17:42 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
 1/26/2011 8:17:48 PM -- Operation mix: Sessions 8, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
 1/26/2011 8:17:48 PM -- Performance logging begins (interval: 15000 ms).
 1/26/2011 8:17:48 PM -- Attaining prerequisites:
 1/26/2011 8:19:22 PM -- \MSEExchange Database(JetstressWin)\Database Cache Size, Last: 1210806000.0 (lower bound: 1207960000.0, upper bound: none)
 1/26/2011 10:19:23 PM -- Performance logging ends.
 1/26/2011 11:06:02 PM -- JetInterop batch transaction stats: 44933, 45244, 45176, 45276 and 45185.
 1/26/2011 11:06:03 PM -- Dispatching transactions ends.
 1/26/2011 11:06:03 PM -- Shutting down databases ...
 1/26/2011 11:06:12 PM -- Instance2880.1 (complete), Instance2880.2 (complete), Instance2880.3 (complete), Instance2880.4 (complete) and Instance2880.5 (complete)
 1/26/2011 11:06:13 PM -- Performance logging begins (interval: 30000 ms).
 1/26/2011 11:06:13 PM -- Verifying database checksums ...
 1/27/2011 11:25:36 AM -- C:\dbluns\db1 (100% processed), C:\dbluns\db2 (100% processed), C:\dbluns\db3 (100% processed), C:\dbluns\db4 (100% processed) and C:\dbluns\db5 (100% processed)
 1/27/2011 11:25:36 AM -- Performance logging ends.
 1/27/2011 11:25:36 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Performance Test\DBChecksum_2011_1_26_23_6_12.blg has 1478 samples.
 1/27/2011 11:25:43 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Performance Test\DBChecksum_2011_1_26_23_6_12.html is saved.
 1/27/2011 11:25:43 AM -- Verifying log checksums ...
 1/27/2011 11:25:44 AM -- C:\logluns\log1 (12 log(s) processed), C:\logluns\log2 (11 log(s) processed), C:\logluns\log3 (12 log(s) processed), C:\logluns\log4 (11 log(s) processed) and

C:\logluns\log5 (11 log(s) processed)
 1/27/2011 11:25:44 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB
 Mbox\Sun170\Performance Test\Performance_2011_1_26_20_17_42.blg has 485 samples.
 1/27/2011 11:25:44 AM -- Creating test report ...
 1/27/2011 11:25:47 AM -- Instance2880.1 has 17.2 for I/O Database Reads Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.1 has 3.5 for I/O Log Writes Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.1 has 3.5 for I/O Log Reads Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.2 has 16.8 for I/O Database Reads Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.2 has 3.5 for I/O Log Writes Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.2 has 3.5 for I/O Log Reads Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.3 has 16.8 for I/O Database Reads Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.3 has 3.5 for I/O Log Writes Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.3 has 3.5 for I/O Log Reads Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.4 has 16.8 for I/O Database Reads Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.4 has 3.3 for I/O Log Writes Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.4 has 3.3 for I/O Log Reads Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.5 has 16.8 for I/O Database Reads Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.5 has 3.2 for I/O Log Writes Average Latency.
 1/27/2011 11:25:47 AM -- Instance2880.5 has 3.2 for I/O Log Reads Average Latency.
 1/27/2011 11:25:47 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
 1/27/2011 11:25:47 AM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
 1/27/2011 11:25:47 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB
 Mbox\Sun170\Performance Test\Performance_2011_1_26_20_17_42.xml has 477 samples queried.

Performance Test Database Checksums Result: SUN170

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:\dbluns\db1\Jetstress001001.edb	39631186	0	0	0	1238474 MBytes / 44305 sec
C:\dbluns\db2\Jetstress002001.edb	39630930	0	0	0	1238466 MBytes / 35139 sec
C:\dbluns\db3\Jetstress003001.edb	39631186	0	0	0	1238474 MBytes / 44336 sec
C:\dbluns\db4\Jetstress004001.edb	39630162	0	0	0	1238442 MBytes / 35209 sec
C:\dbluns\db5\Jetstress005001.edb	39630930	0	0	0	1238466 MBytes / 44362 sec
(Sum)	198154394	0	0	0	6192324 MBytes / 44362 sec

Disk Subsystem Performance (of checksum)

<i>LogicalDisk</i>	<i>Avg. Disk sec/Read</i>	<i>Avg. Disk sec/Write</i>	<i>Disk Reads/sec</i>	<i>Disk Writes/sec</i>	<i>Avg. Disk Bytes/Read</i>
C:\dbluns\db1	0.111	0.000	447.471	0.000	65536.000
C:\dbluns\db2	0.093	0.000	564.083	0.000	65536.000
C:\dbluns\db3	0.110	0.000	447.128	0.000	65536.000
C:\dbluns\db4	0.093	0.000	563.072	0.000	65536.000
C:\dbluns\db5	0.109	0.000	446.957	0.000	65536.000

Memory System Performance (of checksum)

<i>Counter</i>	<i>Average</i>	<i>Minimum</i>	<i>Maximum</i>
% Processor Time	0.858	0.000	3.690
Available MBytes	30800.978	30756.000	30814.000
Free System Page Table Entries	33555643.138	33555643.000	33555645.000
Transition Pages Repurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	67253548.687	67092480.000	67694592.000
Pool Paged Bytes	101567095.859	100974592.000	129298432.000

Test Log

1/26/2011 8:17:31 PM -- Jetstress testing begins ...
 1/26/2011 8:17:31 PM -- Prepare testing begins ...
 1/26/2011 8:17:36 PM -- Attaching databases ...
 1/26/2011 8:17:36 PM -- Prepare testing ends.
 1/26/2011 8:17:36 PM -- Dispatching transactions begins ...
 1/26/2011 8:17:36 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
 1/26/2011 8:17:36 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
 1/26/2011 8:17:42 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
 1/26/2011 8:17:42 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
 1/26/2011 8:17:48 PM -- Operation mix: Sessions 8, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
 1/26/2011 8:17:48 PM -- Performance logging begins (interval: 15000 ms).
 1/26/2011 8:17:48 PM -- Attaining prerequisites:
 1/26/2011 8:19:22 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1210806000.0 (lower bound: 1207960000.0, upper bound: none)
 1/26/2011 10:19:23 PM -- Performance logging ends.
 1/26/2011 11:06:02 PM -- JetInterop batch transaction stats: 44933, 45244, 45176, 45276 and 45185.
 1/26/2011 11:06:03 PM -- Dispatching transactions ends.
 1/26/2011 11:06:03 PM -- Shutting down databases ...
 1/26/2011 11:06:12 PM -- Instance2880.1 (complete), Instance2880.2 (complete), Instance2880.3 (complete), Instance2880.4 (complete) and Instance2880.5 (complete)
 1/26/2011 11:06:13 PM -- Performance logging begins (interval: 30000 ms).
 1/26/2011 11:06:13 PM -- Verifying database checksums ...
 1/27/2011 11:25:36 AM -- C:\dbluns\db1 (100% processed), C:\dbluns\db2 (100% processed), C:\dbluns\db3 (100% processed), C:\dbluns\db4 (100% processed) and C:\dbluns\db5 (100%

processed)

1/27/2011 11:25:36 AM -- Performance logging ends.

1/27/2011 11:25:36 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB

Mbox\Sun170\Performance Test\DBChecksum_2011_1_26_23_6_12.blg has 1478 samples.

Stress Test Result Report: SUN170

Test Summary

Overall Test Result	Pass
Machine Name	SUN170
Test Description	
Test Start Time	2/6/2011 8:49:27 PM
Test End Time	2/7/2011 10:38:55 PM
Collection Start Time	2/6/2011 8:51:28 PM
Collection End Time	2/7/2011 8:51:27 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:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Stress Test\Stress_2011_2_6_20_49_39.blg C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Stress Test\DBChecksum_2011_2_7_22_38_55.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	770.193
Target Transactional I/O per Second	600
Initial Database Size (bytes)	6559527403520
Final Database Size (bytes)	6588610707456
Database Files (Count)	5

Jetstress System Parameters

<i>Thread Count</i>	8 (per database)
<i>Minimum Database Cache</i>	160.0 MB
<i>Maximum Database Cache</i>	1280.0 MB
<i>Insert Operations</i>	40%
<i>Delete Operations</i>	20%
<i>Replace Operations</i>	5%
<i>Read Operations</i>	35%
<i>Lazy Commits</i>	70%
<i>Run Background Database Maintenance</i>	True
<i>Number of Copies per Database</i>	2

Database Configuration

<i>Instance3000.1</i>	Log Path: C:\logluns\log1 Database: C:\dbluns\db1\Jetstress001001.edb
<i>Instance3000.2</i>	Log Path: C:\logluns\log2 Database: C:\dbluns\db2\Jetstress002001.edb
<i>Instance3000.3</i>	Log Path: C:\logluns\log3 Database: C:\dbluns\db3\Jetstress003001.edb
<i>Instance3000.4</i>	Log Path: C:\logluns\log4 Database: C:\dbluns\db4\Jetstress004001.edb
<i>Instance3000.5</i>	Log Path: C:\logluns\log5 Database: C:\dbluns\db5\Jetstress005001.edb

Transactional I/O Performance

<i>MSExchange Database => Instances</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>
Instance3000.1	17.212	8.836	95.920	58.533	36269.921	34988.846	0.000	3.765	0.000	41.958	0.000	5232.694
Instance3000.2	16.906	7.797	95.576	58.286	36342.308	34962.338	0.000	3.684	0.000	41.921	0.000	5224.271
Instance3000.3	16.898	7.726	95.374	58.196	36287.686	34963.623	0.000	3.819	0.000	41.781	0.000	5260.997
Instance3000.4	16.857	6.596	95.872	58.501	36294.680	34972.249	0.000	3.698	0.000	41.945	0.000	5237.240
Instance3000.5	16.919	6.479	95.612	58.323	36393.106	34977.391	0.000	3.848	0.000	41.813	0.000	5245.680

Background Database Maintenance I/O Performance

<i>MSExchange Database ==> Instances</i>	<i>Database Maintenance IO Reads/sec</i>	<i>Database Maintenance IO Reads Average Bytes</i>
Instance3000.1	22.337	261833.684
Instance3000.2	22.527	261733.385
Instance3000.3	22.532	261830.059
Instance3000.4	22.659	261750.142
Instance3000.5	22.410	261836.842

Log Replication I/O Performance

<i>MSExchange Database ==> Instances</i>	<i>I/O Log Reads/sec</i>	<i>I/O Log Reads Average Bytes</i>
Instance3000.1	0.902	225792.067
Instance3000.2	0.898	225483.361
Instance3000.3	0.900	225558.231
Instance3000.4	0.901	225260.977
Instance3000.5	0.899	225548.217

Total I/O Performance

<i>MSExchange Database => Instances</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>
Instance3000.1	17.212	8.836	118.257	58.533	78875.682	34988.846	13.794	3.765	0.902	41.958	225792.067	5232.694
Instance3000.2	16.906	7.797	118.104	58.286	79334.146	34962.338	13.311	3.684	0.898	41.921	225483.361	5224.271
Instance3000.3	16.898	7.726	117.906	58.196	79388.359	34963.623	13.760	3.819	0.900	41.781	225558.231	5260.997
Instance3000.4	16.857	6.596	118.530	58.501	79393.742	34972.249	11.933	3.698	0.901	41.945	225260.977	5237.240
Instance3000.5	16.919	6.479	118.022	58.323	79200.608	34977.391	12.791	3.848	0.899	41.813	225548.217	5245.680

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.922	0.000	4.484
Available MBytes	29517.978	29508.000	29619.000
Free System Page Table Entries	33555643.054	33555634.000	33555645.000
Transition Pages Repurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	66132816.111	66113536.000	66240512.000
Pool Paged Bytes	97274167.900	95473664.000	100487168.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

2/6/2011 8:49:27 PM -- Jetstress testing begins ...
 2/6/2011 8:49:27 PM -- Prepare testing begins ...
 2/6/2011 8:49:33 PM -- Attaching databases ...
 2/6/2011 8:49:33 PM -- Prepare testing ends.
 2/6/2011 8:49:33 PM -- Dispatching transactions begins ...
 2/6/2011 8:49:33 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
 2/6/2011 8:49:33 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
 2/6/2011 8:49:39 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
 2/6/2011 8:49:39 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
 2/6/2011 8:49:46 PM -- Operation mix: Sessions 8, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
 2/6/2011 8:49:46 PM -- Performance logging begins (interval: 15000 ms).
 2/6/2011 8:49:46 PM -- Attaining prerequisites:
 2/6/2011 8:51:28 PM -- \MSEExchange Database(JetstressWin)\Database Cache Size, Last: 1211146000.0 (lower bound: 1207960000.0, upper bound: none)
 2/7/2011 8:51:28 PM -- Performance logging ends.
 2/7/2011 10:38:51 PM -- JetInterop batch transaction stats: 408248, 406570, 407281, 408194 and 407186.
 2/7/2011 10:38:51 PM -- Dispatching transactions ends.
 2/7/2011 10:38:51 PM -- Shutting down databases ...
 2/7/2011 10:38:55 PM -- Instance3000.1 (complete), Instance3000.2 (complete), Instance3000.3 (complete), Instance3000.4 (complete) and Instance3000.5 (complete)
 2/7/2011 10:38:56 PM -- Performance logging begins (interval: 30000 ms).
 2/7/2011 10:38:56 PM -- Verifying database checksums ...
 2/8/2011 11:03:17 AM -- C:\dbluns\db1 (100% processed), C:\dbluns\db2 (100% processed), C:\dbluns\db3 (100% processed), C:\dbluns\db4 (100% processed) and C:\dbluns\db5 (100% processed)
 2/8/2011 11:03:17 AM -- Performance logging ends.
 2/8/2011 11:03:17 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Stress Test\DBChecksum_2011_2_7_22_38_55.blg has 1487 samples.
 2/8/2011 11:03:23 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Stress Test\DBChecksum_2011_2_7_22_38_55.html is saved.
 2/8/2011 11:03:23 AM -- Verifying log checksums ...
 2/8/2011 11:03:24 AM -- C:\logluns\log1 (10 log(s) processed), C:\logluns\log2 (12 log(s) processed), C:\logluns\log3 (12 log(s) processed), C:\logluns\log4 (11 log(s) processed) and

C:\logluns\log5 (13 log(s) processed)
 2/8/2011 11:03:24 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Stress
 Test\Stress_2011_2_6_20_49_39.blg has 5758 samples.
 2/8/2011 11:03:24 AM -- Creating test report ...
 2/8/2011 11:03:54 AM -- Instance3000.1 has 17.2 for I/O Database Reads Average Latency.
 2/8/2011 11:03:54 AM -- Instance3000.1 has 3.8 for I/O Log Writes Average Latency.
 2/8/2011 11:03:54 AM -- Instance3000.1 has 3.8 for I/O Log Reads Average Latency.
 2/8/2011 11:03:54 AM -- Instance3000.2 has 16.9 for I/O Database Reads Average Latency.
 2/8/2011 11:03:54 AM -- Instance3000.2 has 3.7 for I/O Log Writes Average Latency.
 2/8/2011 11:03:54 AM -- Instance3000.2 has 3.7 for I/O Log Reads Average Latency.
 2/8/2011 11:03:54 AM -- Instance3000.3 has 16.9 for I/O Database Reads Average Latency.
 2/8/2011 11:03:54 AM -- Instance3000.3 has 3.8 for I/O Log Writes Average Latency.
 2/8/2011 11:03:54 AM -- Instance3000.3 has 3.8 for I/O Log Reads Average Latency.
 2/8/2011 11:03:55 AM -- Instance3000.4 has 16.9 for I/O Database Reads Average Latency.
 2/8/2011 11:03:55 AM -- Instance3000.4 has 3.7 for I/O Log Writes Average Latency.
 2/8/2011 11:03:55 AM -- Instance3000.4 has 3.7 for I/O Log Reads Average Latency.
 2/8/2011 11:03:55 AM -- Instance3000.5 has 16.9 for I/O Database Reads Average Latency.
 2/8/2011 11:03:55 AM -- Instance3000.5 has 3.8 for I/O Log Writes Average Latency.
 2/8/2011 11:03:55 AM -- Instance3000.5 has 3.8 for I/O Log Reads Average Latency.
 2/8/2011 11:03:55 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
 2/8/2011 11:03:55 AM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
 2/8/2011 11:03:55 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Stress
 Test\Stress_2011_2_6_20_49_39.xml has 5751 samples queried.

Stress Test Database Checksums Result: SUN170

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:\dbluns\db1\Jetstress001001.edb	40214866	0	0	0	1256714 MBytes / 44546 sec
C:\dbluns\db2\Jetstress002001.edb	40214610	0	0	0	1256706 MBytes / 36197 sec
C:\dbluns\db3\Jetstress003001.edb	40213074	0	0	0	1256658 MBytes / 44649 sec
C:\dbluns\db4\Jetstress004001.edb	40213586	0	0	0	1256674 MBytes / 36468 sec
C:\dbluns\db5\Jetstress005001.edb	40212306	0	0	0	1256634 MBytes / 44660 sec
(Sum)	201068442	0	0	0	6283388 MBytes / 44660 sec

Disk Subsystem Performance (of checksum)

<i>LogicalDisk</i>	<i>Avg. Disk sec/Read</i>	<i>Avg. Disk sec/Write</i>	<i>Disk Reads/sec</i>	<i>Disk Writes/sec</i>	<i>Avg. Disk Bytes/Read</i>
C:\dbluns\db1	0.116	0.000	451.482	0.000	65536.000
C:\dbluns\db2	0.099	0.000	555.688	0.000	65536.000
C:\dbluns\db3	0.116	0.000	449.751	0.000	65536.000
C:\dbluns\db4	0.099	0.000	551.678	0.000	65536.000
C:\dbluns\db5	0.115	0.000	448.848	0.000	65536.000

Memory System Performance (of checksum)

<i>Counter</i>	<i>Average</i>	<i>Minimum</i>	<i>Maximum</i>
% Processor Time	0.903	0.000	3.966
Available MBytes	30870.138	30851.000	30880.000
Free System Page Table Entries	33555643.298	33555642.000	33555645.000
Transition Pages Repurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	66560272.699	66461696.000	66797568.000
Pool Paged Bytes	97978614.531	96620544.000	100737024.000

Test Log

2/6/2011 8:49:27 PM -- Jetstress testing begins ...
 2/6/2011 8:49:27 PM -- Prepare testing begins ...
 2/6/2011 8:49:33 PM -- Attaching databases ...
 2/6/2011 8:49:33 PM -- Prepare testing ends.
 2/6/2011 8:49:33 PM -- Dispatching transactions begins ...
 2/6/2011 8:49:33 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
 2/6/2011 8:49:33 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
 2/6/2011 8:49:39 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
 2/6/2011 8:49:39 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
 2/6/2011 8:49:46 PM -- Operation mix: Sessions 8, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
 2/6/2011 8:49:46 PM -- Performance logging begins (interval: 15000 ms).
 2/6/2011 8:49:46 PM -- Attaining prerequisites:
 2/6/2011 8:51:28 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1211146000.0 (lower bound: 1207960000.0, upper bound: none)
 2/7/2011 8:51:28 PM -- Performance logging ends.
 2/7/2011 10:38:51 PM -- JetInterop batch transaction stats: 408248, 406570, 407281, 408194 and 407186.
 2/7/2011 10:38:51 PM -- Dispatching transactions ends.
 2/7/2011 10:38:51 PM -- Shutting down databases ...
 2/7/2011 10:38:55 PM -- Instance3000.1 (complete), Instance3000.2 (complete), Instance3000.3 (complete), Instance3000.4 (complete) and Instance3000.5 (complete)
 2/7/2011 10:38:56 PM -- Performance logging begins (interval: 30000 ms).
 2/7/2011 10:38:56 PM -- Verifying database checksums ...
 2/8/2011 11:03:17 AM -- C:\dbluns\db1 (100% processed), C:\dbluns\db2 (100% processed),

C:\dbluns\db3 (100% processed), C:\dbluns\db4 (100% processed) and C:\dbluns\db5 (100% processed)

2/8/2011 11:03:17 AM -- Performance logging ends.

2/8/2011 11:03:17 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Stress Test\DBChecksum_2011_2_7_22_38_55.blg has 1487 samples.

Backup Test Result Report: SUN170

Database Backup Statistics - All

<i>Database Instance</i>	<i>Database Size (MBytes)</i>	<i>Elapsed Backup Time</i>	<i>MBytes Transferred/sec</i>
Instance3000.1	1256706.59	10:14:24	34.09
Instance3000.2	1256698.59	09:45:35	35.77
Instance3000.3	1256650.59	09:52:06	35.37
Instance3000.4	1256666.59	09:49:12	35.55
Instance3000.5	1256626.59	09:54:32	35.23

Jetstress System Parameters

<i>Thread Count</i>	8 (per database)
<i>Minimum Database Cache</i>	160.0 MB
<i>Maximum Database Cache</i>	1280.0 MB
<i>Insert Operations</i>	40%
<i>Delete Operations</i>	20%
<i>Replace Operations</i>	5%
<i>Read Operations</i>	35%
<i>Lazy Commits</i>	70%

Database Configuration

<i>Instance3000.1</i>	Log Path: C:\logluns\log1 Database: C:\dbluns\db1\Jetstress001001.edb
<i>Instance3000.2</i>	Log Path: C:\logluns\log2 Database: C:\dbluns\db2\Jetstress002001.edb
<i>Instance3000.3</i>	Log Path: C:\logluns\log3 Database: C:\dbluns\db3\Jetstress003001.edb
<i>Instance3000.4</i>	Log Path: C:\logluns\log4 Database: C:\dbluns\db4\Jetstress004001.edb
<i>Instance3000.5</i>	Log Path: C:\logluns\log5 Database: C:\dbluns\db5\Jetstress005001.edb

Transactional I/O Performance

<i>MSExchange Database => Instances</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>
Instance3000.1	11.804	0.000	136.278	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance3000.2	9.907	0.000	143.048	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance3000.3	9.946	0.000	141.411	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance3000.4	9.902	0.000	142.070	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance3000.5	9.972	0.000	140.808	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Host System Performance

<i>Counter</i>	<i>Average</i>	<i>Minimum</i>	<i>Maximum</i>
% Processor Time	0.732	0.000	2.322
Available MBytes	30878.024	30864.000	30881.000
Free System Page Table Entries	33555643.075	33555643.000	33555645.000
Transition Pages Repurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	66347084.717	66330624.000	66371584.000
Pool Paged Bytes	100541493.368	99651584.000	104235008.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log

2/8/2011 9:29:38 PM -- Jetstress testing begins ...
2/8/2011 9:29:38 PM -- Prepare testing begins ...
2/8/2011 9:29:43 PM -- Attaching databases ...
2/8/2011 9:29:43 PM -- Prepare testing ends.
2/8/2011 9:29:53 PM -- Performance logging begins (interval: 30000 ms).
2/8/2011 9:29:53 PM -- Backing up databases ...
2/9/2011 7:44:17 AM -- Performance logging ends.
2/9/2011 7:44:17 AM -- Instance3000.1 (100% processed), Instance3000.2 (100% processed),
Instance3000.3 (100% processed), Instance3000.4 (100% processed) and Instance3000.5 (100%
processed)
2/9/2011 7:44:17 AM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Backup
Test\DatabaseBackup_2011_2_8_21_29_43.blg has 1228 samples.
2/9/2011 7:44:17 AM -- Creating test report ...

Soft Recovery Test Result Report: SUN170

Soft-Recovery Statistics - All

<i>Database Instance</i>	<i>Log files replayed</i>	<i>Elapsed seconds</i>
Instance3000.1	511	867.2055232
Instance3000.2	500	832.7294626
Instance3000.3	504	855.0063017
Instance3000.4	502	857.1279055
Instance3000.5	503	867.4707236

Database Configuration

<i>Instance3000.1</i>	Log Path: C:\logluns\log1 Database: C:\dbluns\db1\Jetstress001001.edb
<i>Instance3000.2</i>	Log Path: C:\logluns\log2 Database: C:\dbluns\db2\Jetstress002001.edb
<i>Instance3000.3</i>	Log Path: C:\logluns\log3 Database: C:\dbluns\db3\Jetstress003001.edb
<i>Instance3000.4</i>	Log Path: C:\logluns\log4 Database: C:\dbluns\db4\Jetstress004001.edb
<i>Instance3000.5</i>	Log Path: C:\logluns\log5 Database: C:\dbluns\db5\Jetstress005001.edb

Transactional I/O Performance

<i>MSExchange Database => Instances</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>
Instance3000.1	26.085	15.730	740.274	3.540	35576.443	32768.000	14.663	0.000	5.313	0.000	232390.528	0.000
Instance3000.2	27.111	14.765	734.336	3.595	35468.690	32768.000	13.621	0.000	5.393	0.000	232403.263	0.000
Instance3000.3	25.990	16.813	732.143	3.540	35547.970	32768.000	12.059	0.000	5.309	0.000	232419.345	0.000
Instance3000.4	25.868	15.204	743.107	3.508	35546.370	32768.000	13.756	0.000	5.261	0.000	231492.996	0.000
Instance3000.5	25.353	15.810	722.217	3.459	35547.836	32768.000	15.829	0.000	5.188	0.000	232516.399	0.000

Background Database Maintenance I/O Performance

<i>MSExchange Database ==> Instances</i>	<i>Database Maintenance IO Reads/sec</i>	<i>Database Maintenance IO Reads Average Bytes</i>
Instance3000.1	21.715	261782.720
Instance3000.2	22.342	261804.865
Instance3000.3	21.748	261809.941
Instance3000.4	22.428	261762.543
Instance3000.5	21.996	261866.476

Total I/O Performance

<i>MSExchange Database => Instances</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>
Instance3000.1	26.085	15.730	761.990	3.540	42022.889	32768.000	14.663	0.000	5.313	0.000	232390.528	0.000
Instance3000.2	27.111	14.765	756.677	3.595	42151.459	32768.000	13.621	0.000	5.393	0.000	232403.263	0.000
Instance3000.3	25.990	16.813	753.892	3.540	42075.209	32768.000	12.059	0.000	5.309	0.000	232419.345	0.000
Instance3000.4	25.868	15.204	765.535	3.508	42173.884	32768.000	13.756	0.000	5.261	0.000	231492.996	0.000
Instance3000.5	25.353	15.810	744.213	3.459	42236.904	32768.000	15.829	0.000	5.188	0.000	232516.399	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	2.281	0.000	14.494
Available MBytes	29523.454	29504.000	30669.000
Free System Page Table Entries	33555643.625	33555642.000	33555645.000
Transition Pages Repurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	73315138.370	69050368.000	74067968.000
Pool Paged Bytes	100577431.704	100544512.000	100646912.000
Database Page Fault Stalls/sec	0.002	0.000	0.249

Test Log

2/9/2011 8:23:59 PM -- Jetstress testing begins ...
 2/9/2011 8:23:59 PM -- Prepare testing begins ...
 2/9/2011 8:24:04 PM -- Attaching databases ...
 2/9/2011 8:24:04 PM -- Prepare testing ends.
 2/9/2011 8:24:04 PM -- Dispatching transactions begins ...
 2/9/2011 8:24:04 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
 2/9/2011 8:24:04 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
 2/9/2011 8:24:10 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
 2/9/2011 8:24:10 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
 2/9/2011 8:24:14 PM -- Operation mix: Sessions 8, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
 2/9/2011 8:24:14 PM -- Performance logging begins (interval: 15000 ms).
 2/9/2011 8:24:14 PM -- Generating log files ...
 2/9/2011 9:38:13 PM -- C:\logluns\log1 (102.4% generated), C:\logluns\log2 (100.2% generated), C:\logluns\log3 (100.8% generated), C:\logluns\log4 (100.6% generated) and C:\logluns\log5 (100.8% generated)
 2/9/2011 9:38:13 PM -- Performance logging ends.
 2/9/2011 9:38:13 PM -- JetInterop batch transaction stats: 22286, 21931, 22119, 22167 and 21805.
 2/9/2011 9:38:13 PM -- Dispatching transactions ends.
 2/9/2011 9:38:13 PM -- Shutting down databases ...
 2/9/2011 9:38:19 PM -- Instance3000.1 (complete), Instance3000.2 (complete), Instance3000.3 (complete), Instance3000.4 (complete) and Instance3000.5 (complete)
 2/9/2011 9:38:19 PM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Soft Recovery\Performance_2011_2_9_20_24_10.blg has 295 samples.
 2/9/2011 9:38:19 PM -- Creating test report ...
 2/9/2011 9:38:20 PM -- Instance3000.1 has 15.1 for I/O Database Reads Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.1 has 2.6 for I/O Log Writes Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.1 has 2.6 for I/O Log Reads Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.2 has 15.2 for I/O Database Reads Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.2 has 4.8 for I/O Log Writes Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.2 has 4.8 for I/O Log Reads Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.3 has 14.7 for I/O Database Reads Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.3 has 2.5 for I/O Log Writes Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.3 has 2.5 for I/O Log Reads Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.4 has 15.2 for I/O Database Reads Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.4 has 4.2 for I/O Log Writes Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.4 has 4.2 for I/O Log Reads Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.5 has 14.7 for I/O Database Reads Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.5 has 2.5 for I/O Log Writes Average Latency.
 2/9/2011 9:38:20 PM -- Instance3000.5 has 2.5 for I/O Log Reads Average Latency.
 2/9/2011 9:38:20 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
 2/9/2011 9:38:20 PM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
 2/9/2011 9:38:20 PM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Soft
 Recovery\Performance_2011_2_9_20_24_10.xml has 294 samples queried.
 2/9/2011 9:38:20 PM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Soft
 Recovery\Performance_2011_2_9_20_24_10.html is saved.
 2/9/2011 9:56:53 PM -- Performance logging begins (interval: 4000 ms).
 2/9/2011 9:56:53 PM -- Recovering databases ...
 2/9/2011 10:11:21 PM -- Performance logging ends.
 2/9/2011 10:11:21 PM -- Instance3000.1 (867.2055232), Instance3000.2 (832.7294626),
 Instance3000.3 (855.0063017), Instance3000.4 (857.1279055) and Instance3000.5 (867.4707236)
 2/9/2011 10:11:21 PM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Soft
 Recovery\SoftRecovery_2011_2_9_21_56_50.blg has 216 samples.
 2/9/2011 10:11:21 PM -- Creating test report ...

Soft Recovery Test Performance Result: SUN170

Test Summary

Overall Test Result	Pass
Machine Name	SUN170
Test Description	
Test Start Time	2/9/2011 8:23:59 PM
Test End Time	2/9/2011 9:38:19 PM
Collection Start Time	2/9/2011 8:24:29 PM
Collection End Time	2/9/2011 9:38:02 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:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Soft Recovery\Performance_2011_2_9_20_24_10.blg

Database Sizing and Throughput

<i>Achieved Transactional I/O per Second</i>	869.986
<i>Capacity Percentage</i>	100%
<i>Throughput Percentage</i>	100%
<i>Initial Database Size (bytes)</i>	6588610707456
<i>Final Database Size (bytes)</i>	6590187765760
<i>Database Files (Count)</i>	5

Jetstress System Parameters

<i>Thread Count</i>	8 (per database)
<i>Minimum Database Cache</i>	160.0 MB
<i>Maximum Database Cache</i>	1280.0 MB
<i>Insert Operations</i>	40%
<i>Delete Operations</i>	20%
<i>Replace Operations</i>	5%
<i>Read Operations</i>	35%
<i>Lazy Commits</i>	70%

Database Configuration

<i>Instance3000.1</i>	Log Path: C:\logluns\log1 Database: C:\dbluns\db1\Jetstress001001.edb
<i>Instance3000.2</i>	Log Path: C:\logluns\log2 Database: C:\dbluns\db2\Jetstress002001.edb
<i>Instance3000.3</i>	Log Path: C:\logluns\log3 Database: C:\dbluns\db3\Jetstress003001.edb
<i>Instance3000.4</i>	Log Path: C:\logluns\log4 Database: C:\dbluns\db4\Jetstress004001.edb
<i>Instance3000.5</i>	Log Path: C:\logluns\log5 Database: C:\dbluns\db5\Jetstress005001.edb

Transactional I/O Performance

<i>MSExchange Database => Instances</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>
Instance3000.1	15.069	9.204	109.918	66.668	32768.000	34950.262	0.000	2.564	0.000	48.792	0.000	5147.263
Instance3000.2	15.213	8.397	107.426	65.081	32768.000	34969.570	0.000	4.799	0.000	46.773	0.000	5274.267
Instance3000.3	14.746	7.675	107.706	65.189	32768.057	34970.969	0.000	2.510	0.000	47.894	0.000	5169.537
Instance3000.4	15.232	6.556	109.196	65.990	32768.264	34888.349	0.000	4.183	0.000	47.726	0.000	5207.533
Instance3000.5	14.703	5.791	107.618	65.194	32768.000	34930.329	0.000	2.549	0.000	48.251	0.000	5147.016

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	1.018	0.000	3.969
Available MBytes	29528.654	29513.000	30449.000
Free System Page Table Entries	33555643.946	33555642.000	33555645.000
Transition Pages Repurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	66668259.363	66437120.000	66772992.000
Pool Paged Bytes	101410739.634	100524032.000	102477824.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

Test Log2/9/2011 8:23:59 PM -- Jetstress testing begins ...

2/9/2011 8:23:59 PM -- Prepare testing begins ...

2/9/2011 8:24:04 PM -- Attaching databases ...

2/9/2011 8:24:04 PM -- Prepare testing ends.

2/9/2011 8:24:04 PM -- Dispatching transactions begins ...

2/9/2011 8:24:04 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)

2/9/2011 8:24:04 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)

2/9/2011 8:24:10 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).

2/9/2011 8:24:10 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).

2/9/2011 8:24:14 PM -- Operation mix: Sessions 8, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.

2/9/2011 8:24:14 PM -- Performance logging begins (interval: 15000 ms).

2/9/2011 8:24:14 PM -- Generating log files ...

2/9/2011 9:38:13 PM -- C:\logluns\log1 (102.4% generated), C:\logluns\log2 (100.2% generated), C:\logluns\log3 (100.8% generated), C:\logluns\log4 (100.6% generated) and C:\logluns\log5 (100.8% generated)

2/9/2011 9:38:13 PM -- Performance logging ends.

2/9/2011 9:38:13 PM -- JetInterop batch transaction stats: 22286, 21931, 22119, 22167 and 21805.

2/9/2011 9:38:13 PM -- Dispatching transactions ends.

2/9/2011 9:38:13 PM -- Shutting down databases ...

2/9/2011 9:38:19 PM -- Instance3000.1 (complete), Instance3000.2 (complete), Instance3000.3 (complete), Instance3000.4 (complete) and Instance3000.5 (complete)

2/9/2011 9:38:19 PM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Soft Recovery\Performance_2011_2_9_20_24_10.blg has 295 samples.

2/9/2011 9:38:19 PM -- Creating test report ...

2/9/2011 9:38:20 PM -- Instance3000.1 has 15.1 for I/O Database Reads Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.1 has 2.6 for I/O Log Writes Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.1 has 2.6 for I/O Log Reads Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.2 has 15.2 for I/O Database Reads Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.2 has 4.8 for I/O Log Writes Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.2 has 4.8 for I/O Log Reads Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.3 has 14.7 for I/O Database Reads Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.3 has 2.5 for I/O Log Writes Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.3 has 2.5 for I/O Log Reads Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.4 has 15.2 for I/O Database Reads Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.4 has 4.2 for I/O Log Writes Average Latency.

2/9/2011 9:38:20 PM -- Instance3000.4 has 4.2 for I/O Log Reads Average Latency.
2/9/2011 9:38:20 PM -- Instance3000.5 has 14.7 for I/O Database Reads Average Latency.
2/9/2011 9:38:20 PM -- Instance3000.5 has 2.5 for I/O Log Writes Average Latency.
2/9/2011 9:38:20 PM -- Instance3000.5 has 2.5 for I/O Log Reads Average Latency.
2/9/2011 9:38:20 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
2/9/2011 9:38:20 PM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
2/9/2011 9:38:20 PM -- C:\ESRP 3 AMS2300 Res R5 HDP 600GB SAS 1GB Mbox\Sun170\Soft
Recovery\Performance_2011_2_9_20_24_10.xml has 294 samples queried.

 **Hitachi Data Systems Corporation**

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 Corporation

© Hitachi Data Systems Corporation 2011. All Rights Reserved. ESRP-067-00 January-February 2011

Corporate Headquarters

750 Central Expressway,
Santa Clara, California 95050-2627 USA
www.HDS.com

Regional Contact Information

Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com