HITACHI
Inspire the Next

# Hitachi Adaptable Modular Storage 2100 Dynamically Provisioned 28,000 User Exchange 2010 Mailbox Resiliency Storage Solution

Tested with: ESRP – Storage Version 3.0

Test Date: April 2010

## Notices and Disclaimer



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** ..... **6**

**Tested Deployment** ..... **6**

**Replication Configuration** ..... **8**

**Best Practices** ..... **9**

Storage – Mailbox Resiliency ..... 10

Storage-based Replication ..... 10

Backup Strategy ..... 10

**Test Result Summary** ..... **11**

Reliability ..... 11

Storage Performance Results ..... 11

Database Backup and Recovery Performance ..... 13

**Conclusion** ..... **14**

**Appendix A — RAID-5 Drive Failure and Rebuild** ..... **15**

**Appendix B — Test Reports** ..... **16**

Performance Test Result: SUN149 ..... 16

Performance Test Database Checksums Result: SUN149 ..... 21

Stress Test Database Performance Result: SUN149 ..... 22

Stress Test Database Checksums Result: SUN149 ..... 27

Database Backup Test Result: SUN149 ..... 28

Soft Recovery Test Result: SUN149 ..... 30

Soft Recovery Test Performance Result: SUN149 ..... 33

# Hitachi Adaptable Modular Storage 2100 Dynamically Provisioned 28,000 User Exchange 2010 Mailbox Resiliency Storage Solution

Tested with: ESRP – Storage Version 3.0

Test Date: April 2010

## Overview

This document provides information on a Hitachi Adaptable Modular Storage 2100 resiliency storage solution using Hitachi Dynamic Provisioning software 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 28,000 users and four servers. This testing used the Hitachi Adaptable Modular Storage 2100 storage system using 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 2100.

This solution includes Exchange 2010 Mailbox Resiliency by using the new Database Availability Group (DAG) feature. This tested configuration utilizes 4 DAG's, each containing two database copies and two servers. The test configuration was capable of supporting 28,000 users with a 0.12 IOPS per user profile and user mailbox size of 1GB. A 2100 with 117 450GB 15K RPM SAS disks, 8GB of cache and 4Gbit/s paths was used for these tests. Testing used four Sun Fire 4270 servers with 32GB of RAM, two quad-core Intel E5540 2.53GHz CPUs, eight Emulex 4Gbit/s Fibre Channel adapters, and Windows Server 2008 R2 Enterprise.

The Hitachi Adaptable Modular Storage 2100 is a medium-sized, high-performance, highly reliable midrange storage system that can scale to 120 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 2100 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 the 2100 to meet the needs of a large Exchange Server deployment.

This solution uses Hitachi Dynamic Provisioning software, which is enabled on the 2100 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 2100 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 2100. 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.

Primarily, Hitachi Dynamic Provisioning software 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 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 make up the pool.

High availability is also a part of this solution with the use of the new 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 4 DAG's, each configured with two servers that host active mailboxes in five databases. To target the 28,000-user resiliency solution, a Hitachi Adaptable Modular Storage 2100 configured with 117 disks (120 is the maximum) and four servers, each configured with 7,000 mailboxes, were used to host the 20 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 2100 and the passive copy on another server connected to a second 2100. 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 2100 storage systems. Each simulated DAG server node in this solution would maintain a mirrored configuration and possess 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.

**Figure 1. Recommended Database Availability Group Configuration**

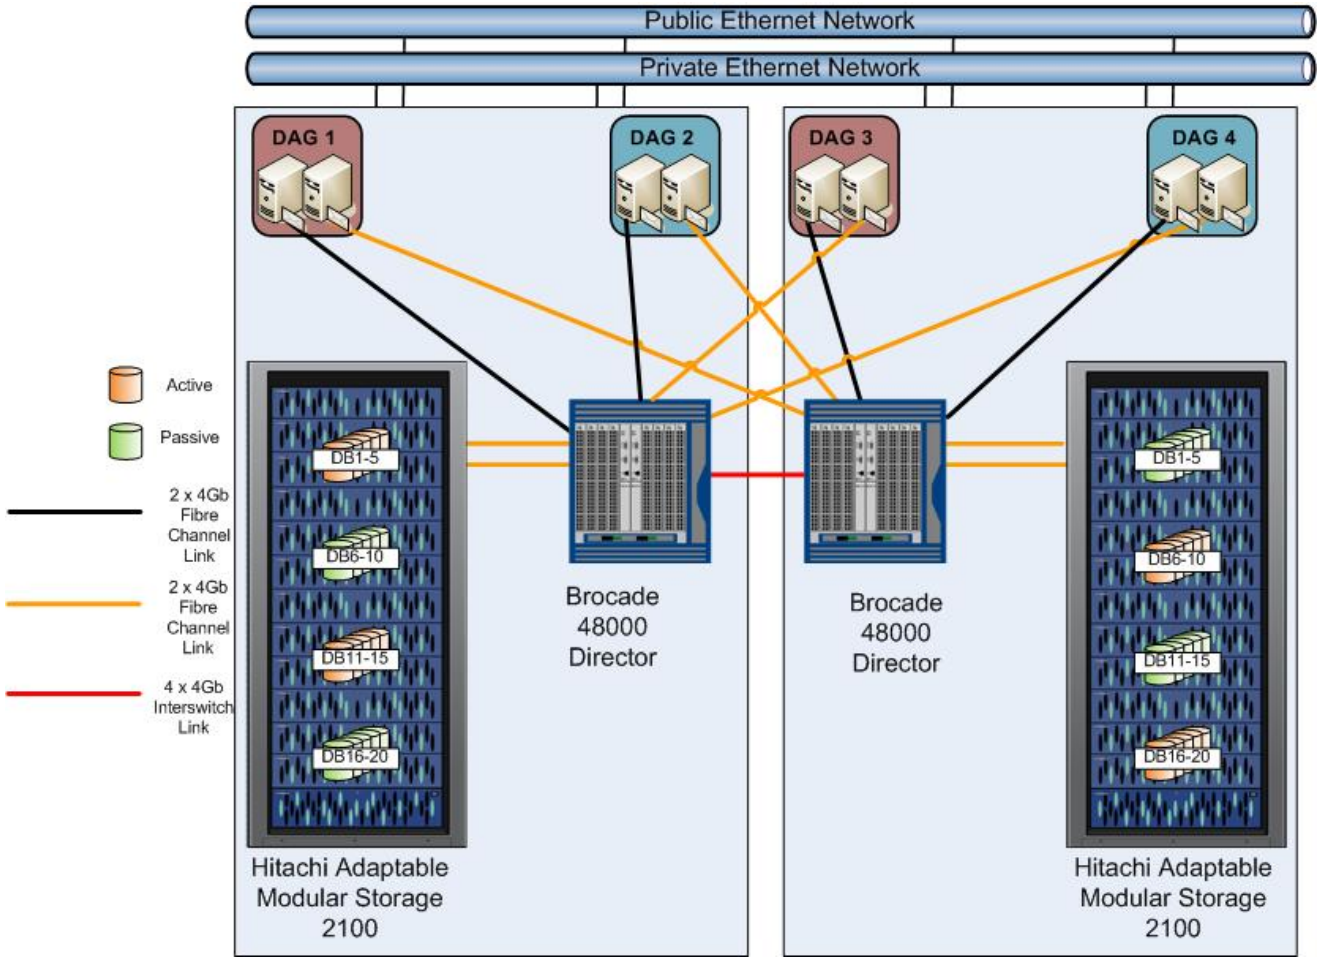


This solution enables organizations to consolidate Exchange Server 2010 DAG deployments on two 2100 storage systems. Using identical hardware and software configurations guarantees that an active database and its replicated copy do not share storage paths, disk spindles or storage controllers, making 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, Exchange Server 2010 mailbox servers can be easily added to support these failover scenarios.

Table 1 illustrates how the 2100'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 2100 RAID Groups by RKA Tray Layout**

| *Drive Slot* | *0* | *1* | *2* | *3* | *4* | *5* | *6* | *7* | *8* | *9* | *10* | *11* | *12* | *13* | *14* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| RKA 7 | 11 | 11 | 11 | 12 | 12 | 12 | 12 | 12 | 12 | 12 | 12 | 12 | | | |
| RKA 6 | 10 | 10 | 10 | 10 | 10 | 10 | 10 | 10 | 10 | 11 | 11 | 11 | 11 | 11 | 11 |
| RKA 5 | 8 | 8 | 8 | 8 | 8 | 8 | 9 | 9 | 9 | 9 | 9 | 9 | 9 | 9 | 9 |
| RKA 4 | 6 | 6 | 6 | 7 | 7 | 7 | 7 | 7 | 7 | 7 | 7 | 7 | 8 | 8 | 8 |
| RKA 3 | 5 | 5 | 5 | 5 | 5 | 5 | 5 | 5 | 5 | 6 | 6 | 6 | 6 | 6 | 6 |
| RKA 2 | 3 | 3 | 3 | 3 | 3 | 3 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4 |
| RKA 1 | 1 | 1 | 1 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 3 | 3 | 3 |
| RKA 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 1 | 1 |

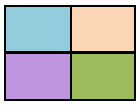

Database RAID Groups

Log RAID Groups

Two Dynamic Provisioning pools were created, one for the databases and the other for the logs. The database pool was created from 11 RAID-5 (8D+1P) RAID groups and the log pool was created from two RAID-5 (8D+1P) RAID groups. From the database pool, 20 DP-VOLs (each specified to have a 1,750GB size limit) were created for 20 databases (five per server). From the Log pool, two DP-VOLs (each specified to have a size limit of 175GB) were created for 20 logs (five per server).

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

**Table 2. Adaptable Modular Storage 2100 Port to Server Layout**

| *Server* | *Primary Path* | *Secondary Path* |
|---|---|---|
| SUN149 | 0A | 1A |
| SUN150 | 0B | 1B |
| SUN151 | 1A | 0A |
| SUN152 | 1B | 0B |

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 2100 Port to Database DP-VOL Layout**

| *Port* | *Database* | *DP-VOL* |
|---|---|---|
| 0A | 1-5 | 0-4 |
| 0B | 6-10 | 5-9 |
| 1A | 11-15 | 10-14 |
| 1B | 16-20 | 15-19 |

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

**Table 4. Adaptable Modular Storage 2100 Port to Log DP-VOL Layout**

| Port | Log | DP-VOL |
|---|---|---|
| 0A | 1-5 | 20-24 |
| 0B | 6-10 | 25-29 |
| 1A | 11-15 | 30-34 |
| 1B | 16-20 | 35-39 |

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 databases and Dynamic Provisioning pool 1 is dedicated for the logs.

**Table 5. Adaptable Modular Storage 2100 Configuration Details**

| Host | Pool | Port | DP-VOL | Size GB) | Description |
|---|---|---|---|---|---|
| Sun149 | 0 | 0A/1A | 0-4 | 1750 | Databases 1-5 |
| Sun150 | 0 | 0B/1B | 5-9 | 1750 | Databases 6-10 |
| Sun151 | 0 | 1A/0A | 10-14 | 1750 | Databases 11-15 |
| Sun152 | 0 | 1B/0B | 15-19 | 1750 | Databases 16-20 |
| Sun149 | 1 | 0A/1A | 20-24 | 175 | Logs 1-5 |
| Sun150 | 1 | 0B/1B | 25-29 | 175 | Logs 6-10 |
| Sun151 | 1 | 1A/0A | 30-34 | 175 | Logs 11-15 |
| Sun152 | 1 | 1B/0B | 35-39 | 175 | Logs 16-20 |

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 28,000 Exchange users with the following specifications:

- Eight Exchange servers (four tested, simulating eight for the database copies)
- Two Adaptable Modular Storage 2100 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**

| | |
|---|---|
| ***Number of Exchange mailboxes simulated*** | 28,000 |
| ***Number of database availability groups (DAGs)*** | 4 |
| ***Number of servers per DAG*** | 2 |
| ***Number of active mailboxes per server*** | 7,000 |
| ***Number of databases per host*** | 5 |
| ***Number of copies per database*** | 2 |
| ***Number of mailboxes per database*** | 1,400 |
| ***Simulated profile: I/Os per second per mailbox (IOPS, include 20% headroom)*** | 0.12 |
| ***Database LU size*** | 1750GB |
| ***Log LU size*** | 175GB |
| ***Total database size for performance testing*** | 28,000GB |
| ***% storage capacity used by Exchange database***** | 79.7% |

**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 2100<br>Firmware: 0872/B-S<br>WHQL listing: Hitachi Adaptable Modular Storage 2100 |
| ***Storage cache*** | 8GB |
| ***Number of storage controllers*** | 2 |
| ***Number of storage ports*** | 4 |
| ***Maximum bandwidth of storage connectivity to host*** | 16Gb/s (4 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<br>2 2.54GHz quad-core Intel Xeon CPUs, 32GB memory |
| ***Total number of disks tested in solution*** | 117 |
| ***Maximum number of spindles can be hosted in the storage*** | 120 |

**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 450GB 15K 4C57 |
| ***Raw capacity per disk (GB)*** | 450GB |
| ***Number of physical disks in test*** | 99 (Dynamic Provisioning pool) |
| ***Total raw storage capacity (GB)*** | 44,550GB |
| ***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*** | 35,112GB (Dynamic Provisioning database pool) |
| ***Storage capacity utilization*** | 78.8% |
| ***Database capacity utilization*** | 78.6% |

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

| | |
|---|---|
| ***Disk type, speed and firmware revision*** | SAS Disk 450GB 15K 4C57 |
| ***Raw capacity per disk (GB)*** | 450GB |
| ***Number of spindles in test*** | 18 (Dynamic Provisioning pool) |
| ***Total raw storage capacity (GB)*** | 8,100GB |
| ***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*** | 6,384GB |

# 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 2100<br>Firmware: 0872/B-S<br>WHQL listing: Hitachi Adaptable Modular Storage 2100 |
| ***Storage cache*** | 8GB |
| ***Number of storage controllers*** | 2 |
| ***Number of storage ports*** | 4 |
| ***Maximum bandwidth of storage connectivity to host*** | 16Gb/s (4 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<br>2 2.54GHz quad-core Intel Xeon CPUs, 32GB memory |
| ***Total number of disks tested in solution*** | 117 |
| ***Maximum number of spindles can be hosted in the storage*** | 120 |

**Table 13. Replicated 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*** | Exchange Server 2010 Database Availability Group (DAG) |

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

| | |
|---|---|
| ***Disk type, speed and firmware revision*** | SAS Disk 450GB 15K 4C57 |
| ***Raw capacity per disk (GB)*** | 450GB |
| ***Number of physical disks in test*** | 99 (Dynamic Provisioning pool) |
| ***Total raw storage capacity (GB)*** | 44,550GB |
| ***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*** | 35,112GB (Dynamic Provisioning database pool) |
| ***Storage capacity utilization*** | 78.8% |
| ***Database capacity utilization*** | 78.6% |

**Table 15. Replicated Storage Disk Configuration (Transaction Log Disks)**

| | |
|---|---|
| ***Disk type, speed and firmware revision*** | SAS Disk 450GB 15K 4C57 |
| ***Raw capacity per disk (GB)*** | 450GB |
| ***Number of spindles in test*** | 18 (Dynamic Provisioning pool) |
| ***Total raw storage capacity (GB)*** | 8,100GB |
| ***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*** | 6,384GB |

# 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 2100 running Exchange 2010.

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

## Storage – Mailbox Resiliency

1. When formatting a newly partitioned LU, Hitachi Data Systems recommends setting the ALU to 64K and 4K respectively for the database and 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-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.
8. Log LUs should be at least 10 percent of the size of the database LUs.
9. LU concatenation is not recommended.
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 2100.
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.
17. 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 (SUN149)**

| | |
|---|---|
| **Database I/O** | |
| ***Database disk transfers per second*** | 1005 |
| ***Database disk reads per second*** | 630 |
| ***Database disk writes per second*** | 375 |
| ***Average database disk read latency (ms)*** | 11.6 |
| ***Average database disk write latency (ms)*** | 3.2 |
| **Transaction Log I/O** | |
| ***Log disk writes per second*** | 302 |
| ***Average log disk write latency (ms)*** | 1.3 |

**Table 17. Individual Server Metrics for Exchange Server (SUN150)**

| Database I/O | |
|---|---|
| *Database disk transfers per second* | 869 |
| *Database disk reads per second* | 543 |
| *Database disk writes per second* | 326 |
| *Average database disk read latency (ms)* | 13.3 |
| *Average database disk write latency (ms)* | 4.1 |
| **Transaction Log I/O** | |
| *Log disk writes per second* | 264 |
| *Average log disk write latency (ms)* | 1.6 |

**Table 18. Individual Server Metrics for Exchange Server (SUN151)**

| Database I/O | |
|---|---|
| *Database disk transfers per second* | 879 |
| *Database disk reads per second* | 549 |
| *Database disk writes per second* | 330 |
| *Average database disk read latency (ms)* | 14.1 |
| *Average database disk write latency (ms)* | 5.4 |
| **Transaction Log I/O** | |
| *Log disk writes per second* | 269 |
| *Average log disk write latency (ms)* | 1.3 |

**Table 19. Individual Server Metrics for Exchange Server (SUN152)**

| Database I/O | |
|---|---|
| *Database disk transfers per second* | 868 |
| *Database disk reads per second* | 542 |
| *Database disk writes per second* | 326 |
| *Average database disk read latency (ms)* | 13.3 |
| *Average database disk write latency (ms)* | 4.1 |
| **Transaction Log I/O** | |
| *Log disk writes per second* | 262 |
| *Average log disk write latency (ms)* | 1.6 |

### *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 20. Aggregate Performance for Exchange Server 2010**

| | |
|---|---|
| **Database I/O** | |
| ***Database disk transfers per second*** | 3621 |
| ***Database disk reads per second*** | 2264 |
| ***Database disk writes per second*** | 1357 |
| ***Average database disk read latency (ms)*** | 13.1 |
| ***Average database disk write latency (ms)*** | 4.2 |
| **Transaction Log I/O** | |
| ***Log disk writes per second*** | 1097 |
| ***Average log disk write latency (ms)*** | 1.4 |

## 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 21. Database Read-only Performance**

| | |
|---|---|
| ***MB read per second per database*** | 43.2 |
| ***MB read per second total per server*** | 215.8 |

### *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 was 1MB in size.

**Table 22. Transaction Log Recovery/Replay Performance**

| | |
|---|---|
| ***Average time to play one log file (sec)*** | 1.10 |

# Conclusion

This document details a tested and robust Exchange Server 2010 Resiliency solution capable of supporting 28,000 users with a 0.12 IOPS per user profile and user mailbox size of 1GB using four DAGs each configured with 2 server nodes. A Hitachi Adaptable Modular Storage 2100, with 8GB of cache and four 4Gbit/s Fibre Channel host paths, using Hitachi Dynamic Provisioning (with two pools) and 117 450GB 15K RPM SAS disks in a RAID-5 configuration, was used for these tests. Testing confirmed that the 2100 is more than capable of delivering the IOPS and capacity requirements needed to support the active and replicated databases for 28,000 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 protection requirements might affect performance and capacity requirements of the underlying storage configuration, and as such needs 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 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: SUN149

### *Test Summary*

| | |
|---|---|
| ***Overall Test Result*** | Pass |
| ***Machine Name*** | SUN149 |
| ***Test Description*** | |
| ***Test Start Time*** | 4/6/2010 12:49:44 AM |
| ***Test End Time*** | 4/6/2010 2:55:40 AM |
| ***Collection Start Time*** | 4/6/2010 12:51:28 AM |
| ***Collection End Time*** | 4/6/2010 2:51:27 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:\ESRP 3.0_1GB MBoxx_RAID 5\Performance Test\7Threads\Performance_2010_4_6_0_49_55.blg<br>C:\ESRP 3.0_1GB MBoxx_RAID 5\Performance Test\7Threads\DBChecksum_2010_4_6_2_55_40.blg |

### *Database Sizing and Throughput*

| | |
|---|---|
| ***Achieved Transactional I/O per Second*** | 1005.579 |
| ***Target Transactional I/O per Second*** | 840 |
| ***Initial Database Size (bytes)*** | 7590146015232 |
| ***Final Database Size (bytes)*** | 7593216245760 |
| ***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

| | |
|---|---|
| ***Instance3676.1*** | Log Path: C:\alogluns\log1<br>Database: C:\asgluns\sg1\Jetstress001001.edb |
| ***Instance3676.2*** | Log Path: C:\alogluns\log2<br>Database: C:\asgluns\sg2\Jetstress002001.edb |
| ***Instance3676.3*** | Log Path: C:\alogluns\log3<br>Database: C:\asgluns\sg3\Jetstress003001.edb |
| ***Instance3676.4*** | Log Path: C:\alogluns\log4<br>Database: C:\asgluns\sg4\Jetstress004001.edb |
| ***Instance3676.5*** | Log Path: C:\alogluns\log5<br>Database: C:\asgluns\sg5\Jetstress005001.edb |

## Transactional I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads /sec | I/O Database Writes /sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads /sec | I/O Log Writes /sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance3676.1 | 10.567 | 3.370 | 125.952 | 75.177 | 34722.881 | 35036.139 | 0.000 | 0.677 | 0.000 | 61.600 | 0.000 | 4702.513 |
| Instance3676.2 | 9.940 | 2.879 | 125.578 | 74.830 | 34505.011 | 34983.949 | 0.000 | 2.152 | 0.000 | 58.609 | 0.000 | 4910.867 |
| Instance3676.3 | 13.878 | 4.349 | 126.582 | 75.370 | 34178.748 | 35071.069 | 0.000 | 0.658 | 0.000 | 61.914 | 0.000 | 4724.178 |
| Instance3676.4 | 9.914 | 1.987 | 126.071 | 75.186 | 34900.836 | 35008.337 | 0.000 | 2.217 | 0.000 | 58.492 | 0.000 | 4938.241 |
| Instance3676.5 | 13.809 | 3.235 | 125.937 | 74.894 | 34440.245 | 35006.992 | 0.000 | 0.659 | 0.000 | 61.125 | 0.000 | 4711.170 |

## Background Database Maintenance I/O Performance

| MSExchange Database ==> Instances | Database Maintenance IO Reads/sec | Database Maintenance IO Reads Average Bytes |
|---|---|---|
| Instance3676.1 | 30.709 | 261878.027 |
| Instance3676.2 | 30.216 | 261878.630 |
| Instance3676.3 | 26.019 | 261853.336 |
| Instance3676.4 | 30.235 | 261892.158 |
| Instance3676.5 | 26.118 | 261836.912 |

## Log Replication I/O Performance

| MSExchange Database ==> Instances | I/O Log Reads/sec | I/O Log Reads Average Bytes |
|---|---|---|
| Instance3676.1 | 1.177 | 232076.263 |
| Instance3676.2 | 1.173 | 232085.149 |
| Instance3676.3 | 1.189 | 232076.263 |
| Instance3676.4 | 1.176 | 232050.052 |
| Instance3676.5 | 1.169 | 230616.351 |

## Total I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads /sec | I/O Database Writes /sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads /sec | I/O Log Writes /sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance3676.1 | 10.567 | 3.370 | 156.662 | 75.177 | 79250.557 | 35036.139 | 1.422 | 0.677 | 1.177 | 61.600 | 232076.263 | 4702.513 |
| Instance3676.2 | 9.940 | 2.879 | 155.795 | 74.830 | 78603.819 | 34983.949 | 7.704 | 2.152 | 1.173 | 58.609 | 232085.149 | 4910.867 |
| Instance3676.3 | 13.878 | 4.349 | 152.602 | 75.370 | 72998.108 | 35071.069 | 1.430 | 0.658 | 1.189 | 61.914 | 232076.263 | 4724.178 |
| Instance3676.4 | 9.914 | 1.987 | 156.307 | 75.186 | 78809.139 | 35008.337 | 7.953 | 2.217 | 1.176 | 58.492 | 232050.052 | 4938.241 |
| Instance3676.5 | 13.809 | 3.235 | 152.056 | 74.894 | 73500.036 | 35006.992 | 1.480 | 0.659 | 1.169 | 61.125 | 230616.351 | 4711.170 |

## Host System Performance

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| % Processor Time | 1.162 | 0.000 | 4.199 |
| Available MBytes | 29927.750 | 29923.000 | 30017.000 |
| Free System Page Table Entries | 33555642.977 | 33555640.000 | 33555643.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72208204.800 | 72089600.000 | 72335360.000 |
| Pool Paged Bytes | 136054707.200 | 135987200.000 | 136122368.000 |
| Database Page Fault Stalls/sec | 0.000 | 0.000 | 0.000 |

## Test Log

```
4/6/2010 12:49:44 AM -- Jetstress testing begins ...
4/6/2010 12:49:44 AM -- Prepare testing begins ...
4/6/2010 12:49:49 AM -- Attaching databases ...
4/6/2010 12:49:49 AM -- Prepare testing ends.
4/6/2010 12:49:49 AM -- Dispatching transactions begins ...
4/6/2010 12:49:49 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2
GB)
4/6/2010 12:49:49 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6
MB)
4/6/2010 12:49:55 AM -- Database read latency thresholds: (average: 20 msec/read,
maximum: 100 msec/read).
4/6/2010 12:49:55 AM -- Log write latency thresholds: (average: 10 msec/write,
maximum: 100 msec/write).
4/6/2010 12:50:01 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%,
Replaces 5%, Reads 35%, Lazy Commits 70%.
4/6/2010 12:50:01 AM -- Performance logging begins (interval: 15000 ms).
4/6/2010 12:50:01 AM -- Attaining prerequisites:
4/6/2010 12:51:28 AM -- \MSExchange Database(JetstressWin)\Database Cache Size,
```

```
Last: 1215697000.0 (lower bound: 1207960000.0, upper bound: none)
4/6/2010 2:51:29 AM -- Performance logging ends.
4/6/2010 2:55:18 AM -- JetInterop batch transaction stats: 42845, 42452, 42733,
42219 and 42754.
4/6/2010 2:55:19 AM -- Dispatching transactions ends.
4/6/2010 2:55:19 AM -- Shutting down databases ...
4/6/2010 2:55:40 AM -- Instance3676.1 (complete), Instance3676.2 (complete),
Instance3676.3 (complete), Instance3676.4 (complete) and Instance3676.5
(complete)
4/6/2010 2:55:42 AM -- Performance logging begins (interval: 30000 ms).
4/6/2010 2:55:42 AM -- Verifying database checksums ...
4/6/2010 2:36:07 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)
4/6/2010 2:36:07 PM -- Performance logging ends.
4/6/2010 2:36:07 PM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Performance
Test\7Threads\DBChecksum_2010_4_6_2_55_40.blg has 1400 samples.
4/6/2010 2:36:13 PM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\ Performance Test
t\7Threads\DBChecksum_2010_4_6_2_55_40.html is saved.
4/6/2010 2:36:13 PM -- Verifying log checksums ...
4/6/2010 2:36:16 PM -- C:\alogluns\log1 (13 log(s) processed), C:\alogluns\log2
(12 log(s) processed), C:\alogluns\log3 (12 log(s) processed), C:\alogluns\log4
(12 log(s) processed) and C:\alogluns\log5 (13 log(s) processed)
4/6/2010 2:36:16 PM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\ Performance Test
\7Threads\Performance_2010_4_6_0_49_55.blg has 485 samples.
4/6/2010 2:36:16 PM -- Creating test report ...
4/6/2010 2:36:19 PM -- Instance3676.1 has 10.6 for I/O Database Reads Average
Latency.
4/6/2010 2:36:19 PM -- Instance3676.1 has 0.7 for I/O Log Writes Average Latency.
4/6/2010 2:36:19 PM -- Instance3676.1 has 0.7 for I/O Log Reads Average Latency.
4/6/2010 2:36:19 PM -- Instance3676.2 has 9.9 for I/O Database Reads Average
Latency.
4/6/2010 2:36:19 PM -- Instance3676.2 has 2.2 for I/O Log Writes Average Latency.
4/6/2010 2:36:19 PM -- Instance3676.2 has 2.2 for I/O Log Reads Average Latency.
4/6/2010 2:36:19 PM -- Instance3676.3 has 13.9 for I/O Database Reads Average
Latency.
4/6/2010 2:36:19 PM -- Instance3676.3 has 0.7 for I/O Log Writes Average Latency.
4/6/2010 2:36:19 PM -- Instance3676.3 has 0.7 for I/O Log Reads Average Latency.
4/6/2010 2:36:19 PM -- Instance3676.4 has 9.9 for I/O Database Reads Average
Latency.
4/6/2010 2:36:19 PM -- Instance3676.4 has 2.2 for I/O Log Writes Average Latency.
4/6/2010 2:36:19 PM -- Instance3676.4 has 2.2 for I/O Log Reads Average Latency.
4/6/2010 2:36:19 PM -- Instance3676.5 has 13.8 for I/O Database Reads Average
Latency.
4/6/2010 2:36:19 PM -- Instance3676.5 has 0.7 for I/O Log Writes Average Latency.
4/6/2010 2:36:19 PM -- Instance3676.5 has 0.7 for I/O Log Reads Average Latency.
4/6/2010 2:36:19 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
4/6/2010 2:36:19 PM -- Test has 0 Database Page Fault Stalls/sec samples higher
than 0.
4/6/2010 2:36:19 PM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\ Performance Test
\7Threads\Performance_2010_4_6_0_49_55.xml has 479 samples queried.
```

# Performance Test Database Checksums Result: SUN149

## Checksum Statistics - All

| Database | Seen pages | Bad pages | Correctable pages | Wrong page-number pages | File length / seconds taken |
|---|---|---|---|---|---|
| C:\asgluns\sg1\Jetstress001001.edb | 46345826 | 0 | 0 | 0 | 1448307 MBytes / 21545 sec |
| C:\asgluns\sg2\Jetstress002001.edb | 46344546 | 0 | 0 | 0 | 1448267 MBytes / 21622 sec |
| C:\asgluns\sg3\Jetstress003001.edb | 46345058 | 0 | 0 | 0 | 1448283 MBytes / 42025 sec |
| C:\asgluns\sg4\Jetstress004001.edb | 46346082 | 0 | 0 | 0 | 1448315 MBytes / 21570 sec |
| C:\asgluns\sg5\Jetstress005001.edb | 46345058 | 0 | 0 | 0 | 1448283 MBytes / 21599 sec |
| (Sum) | 231726570 | 0 | 0 | 0 | 7241455 MBytes / 42025 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 | 0.053 | 0.000 | 1075.157 | 0.000 | 65536.000 |
| C:\asgluns\sg2 | 0.053 | 0.000 | 1070.413 | 0.000 | 65536.000 |
| C:\asgluns\sg3 | 0.083 | 0.000 | 551.046 | 0.000 | 65536.000 |
| C:\asgluns\sg4 | 0.053 | 0.000 | 1073.310 | 0.000 | 65536.000 |
| C:\asgluns\sg5 | 0.052 | 0.000 | 1071.737 | 0.000 | 65536.000 |

## Memory System Performance of Checksum

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| % Processor Time | 1.099 | 0.000 | 4.508 |
| Available MBytes | 31290.329 | 31270.000 | 31302.000 |
| Free System Page Table Entries | 33555642.990 | 33555611.000 | 33555645.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72266892.434 | 72155136.000 | 74846208.000 |
| Pool Paged Bytes | 135013247.269 | 134090752.000 | 137969664.000 |

## Test Log

```
4/6/2010 12:49:44 AM -- Jetstress testing begins ...
4/6/2010 12:49:44 AM -- Prepare testing begins ...
4/6/2010 12:49:49 AM -- Attaching databases ...
4/6/2010 12:49:49 AM -- Prepare testing ends.
4/6/2010 12:49:49 AM -- Dispatching transactions begins ...
4/6/2010 12:49:49 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2
GB)
```

```
4/6/2010 12:49:49 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6
MB)
4/6/2010 12:49:55 AM -- Database read latency thresholds: (average: 20 msec/read,
maximum: 100 msec/read).
4/6/2010 12:49:55 AM -- Log write latency thresholds: (average: 10 msec/write,
maximum: 100 msec/write).
4/6/2010 12:50:01 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%,
Replaces 5%, Reads 35%, Lazy Commits 70%.
4/6/2010 12:50:01 AM -- Performance logging begins (interval: 15000 ms).
4/6/2010 12:50:01 AM -- Attaining prerequisites:
4/6/2010 12:51:28 AM -- \MSExchange Database(JetstressWin)\Database Cache Size,
Last: 1215697000.0 (lower bound: 1207960000.0, upper bound: none)
4/6/2010 2:51:29 AM -- Performance logging ends.
4/6/2010 2:55:18 AM -- JetInterop batch transaction stats: 42845, 42452, 42733,
42219 and 42754.
4/6/2010 2:55:19 AM -- Dispatching transactions ends.
4/6/2010 2:55:19 AM -- Shutting down databases ...
4/6/2010 2:55:40 AM -- Instance3676.1 (complete), Instance3676.2 (complete),
Instance3676.3 (complete), Instance3676.4 (complete) and Instance3676.5
(complete)
4/6/2010 2:55:42 AM -- Performance logging begins (interval: 30000 ms).
4/6/2010 2:55:42 AM -- Verifying database checksums ...
4/6/2010 2:36:07 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)
4/6/2010 2:36:07 PM -- Performance logging ends.
4/6/2010 2:36:07 PM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Performance
Test\7Threads\DBChecksum_2010_4_6_2_55_40.blg has 1400 samples.
```

## Stress Test Database Performance Result: SUN149

### *Test Summary*

| | |
|---|---|
| ***Overall Test Result*** | Pass |
| ***Machine Name*** | SUN149 |
| ***Test Description*** | |
| ***Test Start Time*** | 4/4/2010 12:02:14 PM |
| ***Test End Time*** | 4/5/2010 12:26:42 PM |
| ***Collection Start Time*** | 4/4/2010 12:04:01 PM |
| ***Collection End Time*** | 4/5/2010 12:03:58 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 3.0_1GB MBoxx_RAID 5\Stress Test\7Thread\Stress_2010_4_4_12_2_25.blg<br>C:\ESRP 3.0_1GB MBoxx_RAID 5\Stress Test\7Thread\DBChecksum_2010_4_5_12_26_42.blg |

## Database Sizing and Throughput

| | |
|---|---|
| ***Achieved Transactional I/O per Second*** | 973.426 |
| ***Target Transactional I/O per Second*** | 840 |
| ***Initial Database Size (bytes)*** | 7555626893312 |
| ***Final Database Size (bytes)*** | 7590146015232 |
| ***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

| | |
|---|---|
| ***Instance3676.1*** | Log Path: C:\alogluns\log1<br>Database: C:\asgluns\sg1\Jetstress001001.edb |
| ***Instance3676.2*** | Log Path: C:\alogluns\log2<br>Database: C:\asgluns\sg2\Jetstress002001.edb |
| ***Instance3676.3*** | Log Path: C:\alogluns\log3<br>Database: C:\asgluns\sg3\Jetstress003001.edb |
| ***Instance3676.4*** | Log Path: C:\alogluns\log4<br>Database: C:\asgluns\sg4\Jetstress004001.edb |
| ***Instance3676.5*** | Log Path: C:\alogluns\log5<br>Database: C:\asgluns\sg5\Jetstress005001.edb |

## Transactional I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads /sec | I/O Database Writes /sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads /sec | I/O Log Writes /sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance3676.1 | 13.478 | 5.767 | 122.164 | 72.805 | 33769.182 | 35078.791 | 0.000 | 1.033 | 0.000 | 58.763 | 0.000 | 4780.031 |
| Instance3676.2 | 12.848 | 5.150 | 121.486 | 72.394 | 33895.612 | 35072.000 | 0.000 | 1.249 | 0.000 | 58.775 | 0.000 | 4761.935 |
| Instance3676.3 | 12.811 | 4.481 | 122.175 | 72.811 | 33950.315 | 35068.381 | 0.000 | 1.024 | 0.000 | 58.545 | 0.000 | 4797.991 |
| Instance3676.4 | 12.883 | 3.724 | 122.072 | 72.767 | 33823.489 | 35085.040 | 0.000 | 1.262 | 0.000 | 59.154 | 0.000 | 4771.622 |
| Instance3676.5 | 12.804 | 2.795 | 122.002 | 72.749 | 33938.846 | 35072.339 | 0.000 | 1.025 | 0.000 | 58.590 | 0.000 | 4786.747 |

## Background Database Maintenance I/O Performance

| MSExchange Database ==> Instances | Database Maintenance IO Reads/sec | Database Maintenance IO Reads Average Bytes |
|---|---|---|
| Instance3676.1 | 28.561 | 261875.875 |
| Instance3676.2 | 28.014 | 261874.447 |
| Instance3676.3 | 27.999 | 261863.728 |
| Instance3676.4 | 27.903 | 261874.145 |
| Instance3676.5 | 28.025 | 261867.958 |

## Log Replication I/O Performance

| MSExchange Database ==> Instances | I/O Log Reads/sec | I/O Log Reads Average Bytes |
|---|---|---|
| Instance3676.1 | 1.143 | 232362.981 |
| Instance3676.2 | 1.137 | 232305.601 |
| Instance3676.3 | 1.143 | 232399.173 |
| Instance3676.4 | 1.147 | 232265.277 |
| Instance3676.5 | 1.141 | 232155.946 |

## Total I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads /sec | I/O Database Writes /sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads /sec | I/O Log Writes /sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance3676.1 | 13.478 | 5.767 | 150.724 | 72.805 | 76992.869 | 35078.791 | 2.291 | 1.033 | 1.143 | 58.763 | 232362.981 | 4780.031 |
| Instance3676.2 | 12.848 | 5.150 | 149.500 | 72.394 | 76614.886 | 35072.000 | 4.481 | 1.249 | 1.137 | 58.775 | 232305.601 | 4761.935 |
| Instance3676.3 | 12.811 | 4.481 | 150.175 | 72.811 | 76443.800 | 35068.381 | 2.515 | 1.024 | 1.143 | 58.545 | 232399.173 | 4797.991 |
| Instance3676.4 | 12.883 | 3.724 | 149.975 | 72.767 | 76252.705 | 35085.040 | 4.267 | 1.262 | 1.147 | 59.154 | 232265.277 | 4771.622 |
| Instance3676.5 | 12.804 | 2.795 | 150.027 | 72.749 | 76515.948 | 35072.339 | 2.682 | 1.025 | 1.141 | 58.590 | 232155.946 | 4786.747 |

## Host System Performance

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| % Processor Time | 1.168 | 0.000 | 4.666 |
| Available MBytes | 29957.462 | 29941.000 | 30084.000 |
| Free System Page Table Entries | 33555642.988 | 33555639.000 | 33555643.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72447051.483 | 71950336.000 | 72536064.000 |
| Pool Paged Bytes | 131839250.159 | 130576384.000 | 134791168.000 |
| Database Page Fault Stalls/sec | 0.000 | 0.000 | 0.000 |

## Test Log

```
4/4/2010 12:02:14 PM -- Jetstress testing begins ...
4/4/2010 12:02:14 PM -- Prepare testing begins ...
4/4/2010 12:02:19 PM -- Attaching databases ...
4/4/2010 12:02:19 PM -- Prepare testing ends.
4/4/2010 12:02:19 PM -- Dispatching transactions begins ...
4/4/2010 12:02:20 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2
GB)
4/4/2010 12:02:20 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6
MB)
4/4/2010 12:02:25 PM -- Database read latency thresholds: (average: 20 msec/read,
maximum: 200 msec/read).
4/4/2010 12:02:25 PM -- Log write latency thresholds: (average: 10 msec/write,
maximum: 200 msec/write).
4/4/2010 12:02:32 PM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%,
Replaces 5%, Reads 35%, Lazy Commits 70%.
4/4/2010 12:02:32 PM -- Performance logging begins (interval: 15000 ms).
4/4/2010 12:02:32 PM -- Attaining prerequisites:
4/4/2010 12:04:01 PM -- \MSExchange Database(JetstressWin)\Database Cache Size,
```

```
Last: 1214288000.0 (lower bound: 1207960000.0, upper bound: none)
4/5/2010 12:04:02 PM -- Performance logging ends.
4/5/2010 12:26:26 PM -- JetInterop batch transaction stats: 481617, 479099,
480859, 480979 and 481151.
4/5/2010 12:26:29 PM -- Dispatching transactions ends.
4/5/2010 12:26:29 PM -- Shutting down databases ...
4/5/2010 12:26:42 PM -- Instance3676.1 (complete), Instance3676.2 (complete),
Instance3676.3 (complete), Instance3676.4 (complete) and Instance3676.5
(complete)
4/5/2010 12:26:43 PM -- Performance logging begins (interval: 30000 ms).
4/5/2010 12:26:43 PM -- Verifying database checksums ...
4/6/2010 12:01:16 AM -- 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)
4/6/2010 12:01:16 AM -- Performance logging ends.
4/6/2010 12:01:16 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Stress
Test\7Thread\DBChecksum_2010_4_5_12_26_42.blg has 1388 samples.
4/6/2010 12:01:23 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Stress
Test\7Thread\DBChecksum_2010_4_5_12_26_42.html is saved.
4/6/2010 12:01:23 AM -- Verifying log checksums ...
4/6/2010 12:01:25 AM -- C:\alogluns\log1 (11 log(s) processed), C:\alogluns\log2
(11 log(s) processed), C:\alogluns\log3 (11 log(s) processed), C:\alogluns\log4
(13 log(s) processed) and C:\alogluns\log5 (12 log(s) processed)
4/6/2010 12:01:25 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Stress
Test\7Thread\Stress_2010_4_4_12_2_25.blg has 5757 samples.
4/6/2010 12:01:25 AM -- Creating test report ...
4/6/2010 12:01:54 AM -- Instance3676.1 has 13.5 for I/O Database Reads Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.1 has 1.0 for I/O Log Writes Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.1 has 1.0 for I/O Log Reads Average Latency.
4/6/2010 12:01:54 AM -- Instance3676.2 has 12.8 for I/O Database Reads Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.2 has 1.2 for I/O Log Writes Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.2 has 1.2 for I/O Log Reads Average Latency.
4/6/2010 12:01:54 AM -- Instance3676.3 has 12.8 for I/O Database Reads Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.3 has 1.0 for I/O Log Writes Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.3 has 1.0 for I/O Log Reads Average Latency.
4/6/2010 12:01:54 AM -- Instance3676.4 has 12.9 for I/O Database Reads Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.4 has 1.3 for I/O Log Writes Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.4 has 1.3 for I/O Log Reads Average Latency.
4/6/2010 12:01:54 AM -- Instance3676.5 has 12.8 for I/O Database Reads Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.5 has 1.0 for I/O Log Writes Average
Latency.
4/6/2010 12:01:54 AM -- Instance3676.5 has 1.0 for I/O Log Reads Average Latency.
4/6/2010 12:01:54 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
4/6/2010 12:01:54 AM -- Test has 0 Database Page Fault Stalls/sec samples higher
than 0.
4/6/2010 12:01:54 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Stress
Test\7Thread\Stress_2010_4_4_12_2_25.xml has 5751 samples queried.
```

# Stress Test Database Checksums Result: SUN149

## *Checksum Statistics - All*

| Database | Seen pages | Bad pages | Correctable pages | Wrong page-number pages | File length / seconds taken |
|---|---|---|---|---|---|
| C:\asgluns\sg1\Jetstress001001.edb | 46327138 | 0 | 0 | 0 | 1447723 MBytes / 41479 sec |
| C:\asgluns\sg2\Jetstress002001.edb | 46325858 | 0 | 0 | 0 | 1447683 MBytes / 41486 sec |
| C:\asgluns\sg3\Jetstress003001.edb | 46326114 | 0 | 0 | 0 | 1447691 MBytes / 41673 sec |
| C:\asgluns\sg4\Jetstress004001.edb | 46327394 | 0 | 0 | 0 | 1447731 MBytes / 41393 sec |
| C:\asgluns\sg5\Jetstress005001.edb | 46326370 | 0 | 0 | 0 | 1447699 MBytes / 41669 sec |
| (Sum) | 231632874 | 0 | 0 | 0 | 7238527 MBytes / 41673 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 | 0.095 | 0.000 | 558.484 | 0.000 | 65536.000 |
| C:\asgluns\sg2 | 0.094 | 0.000 | 558.313 | 0.000 | 65536.000 |
| C:\asgluns\sg3 | 0.094 | 0.000 | 555.949 | 0.000 | 65536.000 |
| C:\asgluns\sg4 | 0.093 | 0.000 | 559.567 | 0.000 | 65536.000 |
| C:\asgluns\sg5 | 0.092 | 0.000 | 555.946 | 0.000 | 65536.000 |

## *Memory System Performance of Checksum*

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| % Processor Time | 1.121 | 0.000 | 4.404 |
| Available MBytes | 31296.823 | 31290.000 | 31310.000 |
| Free System Page Table Entries | 33555643.003 | 33555642.000 | 33555645.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72527025.061 | 72376320.000 | 72966144.000 |
| Pool Paged Bytes | 132319606.778 | 131284992.000 | 134361088.000 |

*Test Log*

```
4/4/2010 12:02:14 PM -- Jetstress testing begins ...
4/4/2010 12:02:14 PM -- Prepare testing begins ...
4/4/2010 12:02:19 PM -- Attaching databases ...
4/4/2010 12:02:19 PM -- Prepare testing ends.
4/4/2010 12:02:19 PM -- Dispatching transactions begins ...
4/4/2010 12:02:20 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2
GB)
4/4/2010 12:02:20 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6
MB)
4/4/2010 12:02:25 PM -- Database read latency thresholds: (average: 20 msec/read,
maximum: 200 msec/read).
4/4/2010 12:02:25 PM -- Log write latency thresholds: (average: 10 msec/write,
maximum: 200 msec/write).
4/4/2010 12:02:32 PM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%,
Replaces 5%, Reads 35%, Lazy Commits 70%.
4/4/2010 12:02:32 PM -- Performance logging begins (interval: 15000 ms).
4/4/2010 12:02:32 PM -- Attaining prerequisites:
4/4/2010 12:04:01 PM -- \MSExchange Database(JetstressWin)\Database Cache Size,
Last: 1214288000.0 (lower bound: 1207960000.0, upper bound: none)
4/5/2010 12:04:02 PM -- Performance logging ends.
4/5/2010 12:26:26 PM -- JetInterop batch transaction stats: 481617, 479099,
480859, 480979 and 481151.
4/5/2010 12:26:29 PM -- Dispatching transactions ends.
4/5/2010 12:26:29 PM -- Shutting down databases ...
4/5/2010 12:26:42 PM -- Instance3676.1 (complete), Instance3676.2 (complete),
Instance3676.3 (complete), Instance3676.4 (complete) and Instance3676.5
(complete)
4/5/2010 12:26:43 PM -- Performance logging begins (interval: 30000 ms).
4/5/2010 12:26:43 PM -- Verifying database checksums ...
4/6/2010 12:01:16 AM -- 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)
4/6/2010 12:01:16 AM -- Performance logging ends.
4/6/2010 12:01:16 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Stress
Test\7Thread\DBChecksum_2010_4_5_12_26_42.blg has 1388 samples.
```

## Database Backup Test Result: SUN149

*Database Backup Statistics - All*

| Database Instance | Database Size (MBytes) | Elapsed Backup Time | MBytes Transferred/sec |
|---|---|---|---|
| Instance3676.1 | 1448299.09 | 12:41:21 | 31.70 |
| Instance3676.2 | 1448259.09 | 12:39:06 | 31.80 |
| Instance3676.3 | 1448275.09 | 12:38:39 | 31.82 |
| Instance3676.4 | 1448307.09 | 12:24:30 | 32.42 |
| Instance3676.5 | 1448275.09 | 12:42:56 | 31.64 |

## 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

| | |
|---|---|
| **Instance3676.1** | Log Path: C:\alogluns\log1<br>Database: C:\asgluns\sg1\Jetstress001001.edb |
| **Instance3676.2** | Log Path: C:\alogluns\log2<br>Database: C:\asgluns\sg2\Jetstress002001.edb |
| **Instance3676.3** | Log Path: C:\alogluns\log3<br>Database: C:\asgluns\sg3\Jetstress003001.edb |
| **Instance3676.4** | Log Path: C:\alogluns\log4<br>Database: C:\asgluns\sg4\Jetstress004001.edb |
| **Instance3676.5** | Log Path: C:\alogluns\log5<br>Database: C:\asgluns\sg5\Jetstress005001.edb |

## Transactional I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads /sec | I/O Database Writes /sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads /sec | I/O Log Writes /sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance3676.1 | 11.402 | 0.000 | 126.774 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Instance3676.2 | 10.877 | 0.000 | 127.178 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Instance3676.3 | 10.844 | 0.000 | 127.261 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Instance3676.4 | 10.598 | 0.000 | 129.687 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Instance3676.5 | 10.999 | 0.000 | 126.353 | 0.000 | 262144.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |

### Host System Performance

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| % Processor Time | 0.869 | 0.000 | 6.695 |
| Available MBytes | 31281.077 | 31245.000 | 31312.000 |
| Free System Page Table Entries | 33555642.827 | 33555637.000 | 33555645.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72166437.627 | 72085504.000 | 72384512.000 |
| Pool Paged Bytes | 135866295.433 | 134303744.000 | 138981376.000 |
| Database Page Fault Stalls/sec | 0.000 | 0.000 | 0.000 |

### Test Log

```
4/6/2010 4:48:29 PM -- Jetstress testing begins ...
4/6/2010 4:48:29 PM -- Prepare testing begins ...
4/6/2010 4:48:34 PM -- Attaching databases ...
4/6/2010 4:48:34 PM -- Prepare testing ends.
4/6/2010 4:48:43 PM -- Performance logging begins (interval: 30000 ms).
4/6/2010 4:48:43 PM -- Backing up databases ...
4/7/2010 5:31:40 AM -- Performance logging ends.
4/7/2010 5:31:40 AM -- Instance3676.1 (100% processed), Instance3676.2 (100%
processed), Instance3676.3 (100% processed), Instance3676.4 (100% processed) and
Instance3676.5 (100% processed)
4/7/2010 5:31:40 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Database
Backup\DatabaseBackup_2010_4_6_16_48_34.blg has 1524 samples.
4/7/2010 5:31:40 AM -- Creating test report ...
```

## Soft Recovery Test Result: SUN149

### Soft-Recovery Statistics - All

| Database Instance | Log files replayed | Elapsed seconds |
|---|---|---|
| Instance3676.1 | 504 | 644.1719315 |
| Instance3676.2 | 508 | 644.4371319 |
| Instance3676.3 | 500 | 616.3258826 |
| Instance3676.4 | 505 | 632.2379105 |
| Instance3676.5 | 500 | 638.0723208 |

### Database Configuration

| | |
|---|---|
| **Instance3676.1** | Log Path: C:\alogluns\log1<br>Database: C:\asgluns\sg1\Jetstress001001.edb |
| **Instance3676.2** | Log Path: C:\alogluns\log2<br>Database: C:\asgluns\sg2\Jetstress002001.edb |
| **Instance3676.3** | Log Path: C:\alogluns\log3<br>Database: C:\asgluns\sg3\Jetstress003001.edb |
| **Instance3676.4** | Log Path: C:\alogluns\log4<br>Database: C:\asgluns\sg4\Jetstress004001.edb |
| **Instance3676.5** | Log Path: C:\alogluns\log5<br>Database: C:\asgluns\sg5\Jetstress005001.edb |

## Transactional I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads /sec | I/O Database Writes /sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads /sec | I/O Log Writes /sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance3676.1 | 27.552 | 16.489 | 962.272 | 4.696 | 36090.642 | 32561.912 | 8.393 | 0.000 | 7.044 | 0.000 | 231049.682 | 0.000 |
| Instance3676.2 | 25.941 | 16.198 | 998.245 | 4.724 | 35991.235 | 32561.912 | 4.956 | 0.000 | 7.087 | 0.000 | 231091.397 | 0.000 |
| Instance3676.3 | 28.592 | 16.081 | 955.175 | 4.892 | 35709.813 | 32768.000 | 8.578 | 0.000 | 7.338 | 0.000 | 232569.372 | 0.000 |
| Instance3676.4 | 27.501 | 15.778 | 977.359 | 4.809 | 35874.115 | 32768.000 | 5.569 | 0.000 | 7.221 | 0.000 | 232528.916 | 0.000 |
| Instance3676.5 | 27.408 | 15.714 | 965.277 | 4.698 | 35566.151 | 32768.000 | 8.104 | 0.000 | 7.047 | 0.000 | 232473.857 | 0.000 |

## Background Database Maintenance I/O Performance

| MSExchange Database ==> Instances | Database Maintenance IO Reads/sec | Database Maintenance IO Reads Average Bytes |
|---|---|---|
| Instance3676.1 | 26.199 | 261959.382 |
| Instance3676.2 | 25.601 | 261940.770 |
| Instance3676.3 | 25.539 | 261944.490 |
| Instance3676.4 | 25.683 | 261936.650 |
| Instance3676.5 | 25.656 | 261872.405 |

## Total I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads /sec | I/O Database Writes /sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads /sec | I/O Log Writes /sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance3676.1 | 27.552 | 16.489 | 988.471 | 4.696 | 42077.211 | 32561.912 | 8.393 | 0.000 | 7.044 | 0.000 | 231049.682 | 0.000 |
| Instance3676.2 | 25.941 | 16.198 | 1023.846 | 4.724 | 41641.099 | 32561.912 | 4.956 | 0.000 | 7.087 | 0.000 | 231091.397 | 0.000 |
| Instance3 | 28.59 | 16.08 | 980.71 | 4.892 | 41601. | 32768. | 8.578 | 0.000 | 7.338 | 0.000 | 232569 | 0.000 |

| | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 676.3 | 2 | 1 | 4 | | 207 | 000 | | | | | .372 | |
| Instance3676.4 | 27.501 | 15.778 | 1003.042 | 4.809 | 41662.381 | 32768.000 | 5.569 | 0.000 | 7.221 | 0.000 | 232528.916 | 0.000 |
| Instance3676.5 | 27.408 | 15.714 | 990.932 | 4.698 | 41425.309 | 32768.000 | 8.104 | 0.000 | 7.047 | 0.000 | 232473.857 | 0.000 |

## *Host System Performance*

| *Counter* | *Average* | *Minimum* | *Maximum* |
|---|---|---|---|
| % Processor Time | 2.769 | 0.000 | 12.548 |
| Available MBytes | 29931.250 | 29905.000 | 31095.000 |
| Free System Page Table Entries | 33555640.994 | 33555640.000 | 33555641.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 79719193.600 | 73531392.000 | 80596992.000 |
| Pool Paged Bytes | 137068467.200 | 137027584.000 | 137125888.000 |
| Database Page Fault Stalls/sec | 0.002 | 0.000 | 0.249 |

## *Test Log*

```
4/7/2010 5:46:12 AM -- Jetstress testing begins ...
4/7/2010 5:46:12 AM -- Prepare testing begins ...
4/7/2010 5:46:17 AM -- Attaching databases ...
4/7/2010 5:46:17 AM -- Prepare testing ends.
4/7/2010 5:46:17 AM -- Dispatching transactions begins ...
4/7/2010 5:46:17 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2
GB)
4/7/2010 5:46:17 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
4/7/2010 5:46:23 AM -- Database read latency thresholds: (average: 20 msec/read,
maximum: 100 msec/read).
4/7/2010 5:46:23 AM -- Log write latency thresholds: (average: 10 msec/write,
maximum: 100 msec/write).
4/7/2010 5:46:27 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%,
Replaces 5%, Reads 35%, Lazy Commits 70%.
4/7/2010 5:46:27 AM -- Performance logging begins (interval: 15000 ms).
4/7/2010 5:46:27 AM -- Generating log files ...
4/7/2010 6:44:59 AM -- C:\alogluns\log1 (101.0% generated), C:\alogluns\log2
(101.8% generated), C:\alogluns\log3 (100.2% generated), C:\alogluns\log4 (101.2%
generated) and C:\alogluns\log5 (100.2% generated)
4/7/2010 6:44:59 AM -- Performance logging ends.
4/7/2010 6:44:59 AM -- JetInterop batch transaction stats: 21688, 21932, 21544,
21670 and 21557.
4/7/2010 6:44:59 AM -- Dispatching transactions ends.
4/7/2010 6:44:59 AM -- Shutting down databases ...
4/7/2010 6:45:04 AM -- Instance3676.1 (complete), Instance3676.2 (complete),
Instance3676.3 (complete), Instance3676.4 (complete) and Instance3676.5
(complete)
4/7/2010 6:45:04 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Soft
Recovery\Performance_2010_4_7_5_46_23.blg has 233 samples.
4/7/2010 6:45:04 AM -- Creating test report ...
4/7/2010 6:45:05 AM -- Instance3676.1 has 13.5 for I/O Database Reads Average
Latency.
4/7/2010 6:45:05 AM -- Instance3676.1 has 1.2 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.1 has 1.2 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.2 has 12.5 for I/O Database Reads Average
Latency.
4/7/2010 6:45:05 AM -- Instance3676.2 has 1.0 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.2 has 1.0 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.3 has 12.5 for I/O Database Reads Average
```

```
Latency.
4/7/2010 6:45:05 AM -- Instance3676.3 has 1.2 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.3 has 1.2 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.4 has 12.6 for I/O Database Reads Average
Latency.
4/7/2010 6:45:05 AM -- Instance3676.4 has 1.0 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.4 has 1.0 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.5 has 12.5 for I/O Database Reads Average
Latency.
4/7/2010 6:45:05 AM -- Instance3676.5 has 1.2 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.5 has 1.2 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
4/7/2010 6:45:05 AM -- Test has 0 Database Page Fault Stalls/sec samples higher
than 0.
4/7/2010 6:45:05 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Soft
Recovery\Performance_2010_4_7_5_46_23.xml has 232 samples queried.
4/7/2010 6:45:05 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Soft
Recovery\Performance_2010_4_7_5_46_23.html is saved.
4/7/2010 6:46:08 AM -- Performance logging begins (interval: 4000 ms).
4/7/2010 6:46:08 AM -- Recovering databases ...
4/7/2010 6:56:52 AM -- Performance logging ends.
4/7/2010 6:56:52 AM -- Instance3676.1 (644.1719315), Instance3676.2
(644.4371319), Instance3676.3 (616.3258826), Instance3676.4 (632.2379105) and
Instance3676.5 (638.0723208)
4/7/2010 6:56:52 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Soft
Recovery\SoftRecovery_2010_4_7_6_46_4.blg has 160 samples.
4/7/2010 6:56:52 AM -- Creating test report ...
```

## Soft Recovery Test Performance Result: SUN149

### *Test Summary*

| | |
|---|---|
| ***Overall Test Result*** | Pass |
| ***Machine Name*** | SUN149 |
| ***Test Description*** | |
| ***Test Start Time*** | 4/7/2010 5:46:12 AM |
| ***Test End Time*** | 4/7/2010 6:45:04 AM |
| ***Collection Start Time*** | 4/7/2010 5:46:42 AM |
| ***Collection End Time*** | 4/7/2010 6:44:45 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:\ESRP 3.0_1GB MBoxx_RAID 5\Soft Recovery\Performance_2010_4_7_5_46_23.blg |

### *Database Sizing and Throughput*

| | |
|---|---|
| ***Achieved Transactional I/O per Second*** | 1091.715 |
| ***Capacity Percentage*** | 100% |
| ***Throughput Percentage*** | 100% |
| ***Initial Database Size (bytes)*** | 7593216245760 |
| ***Final Database Size (bytes)*** | 7594784915456 |
| ***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

| | |
|---|---|
| **Instance3676.1** | Log Path: C:\alogluns\log1<br>Database: C:\asgluns\sg1\Jetstress001001.edb |
| **Instance3676.2** | Log Path: C:\alogluns\log2<br>Database: C:\asgluns\sg2\Jetstress002001.edb |
| **Instance3676.3** | Log Path: C:\alogluns\log3<br>Database: C:\asgluns\sg3\Jetstress003001.edb |
| **Instance3676.4** | Log Path: C:\alogluns\log4<br>Database: C:\asgluns\sg4\Jetstress004001.edb |
| **Instance3676.5** | Log Path: C:\alogluns\log5<br>Database: C:\asgluns\sg5\Jetstress005001.edb |

## Transactional I/O Performance

| MSExchange Database ==> Instances | I/O Database Reads Average Latency (msec) | I/O Database Writes Average Latency (msec) | I/O Database Reads /sec | I/O Database Writes /sec | I/O Database Reads Average Bytes | I/O Database Writes Average Bytes | I/O Log Reads Average Latency (msec) | I/O Log Writes Average Latency (msec) | I/O Log Reads /sec | I/O Log Writes /sec | I/O Log Reads Average Bytes | I/O Log Writes Average Bytes |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance3676.1 | 13.486 | 9.937 | 137.970 | 81.774 | 32768.000 | 35059.916 | 0.000 | 1.178 | 0.000 | 65.976 | 0.000 | 4788.609 |
| Instance3676.2 | 12.544 | 8.688 | 139.242 | 82.479 | 32769.085 | 35038.475 | 0.000 | 1.019 | 0.000 | 67.010 | 0.000 | 4768.954 |
| Instance3676.3 | 12.491 | 7.254 | 135.117 | 80.122 | 32768.000 | 35019.515 | 0.000 | 1.185 | 0.000 | 64.553 | 0.000 | 4849.346 |
| Instance3676.4 | 12.570 | 5.636 | 136.576 | 81.064 | 32768.000 | 35031.097 | 0.000 | 1.005 | 0.000 | 66.041 | 0.000 | 4793.826 |
| Instance3676.5 | 12.506 | 3.928 | 136.467 | 80.904 | 32768.264 | 34991.871 | 0.000 | 1.201 | 0.000 | 65.295 | 0.000 | 4806.090 |

## Host System Performance

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| % Processor Time | 1.273 | 0.000 | 3.342 |
| Available MBytes | 29930.755 | 29913.000 | 30875.000 |
| Free System Page Table Entries | 33555641.000 | 33555641.000 | 33555641.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 73219480.721 | 72335360.000 | 73428992.000 |
| Pool Paged Bytes | 136903877.768 | 136802304.000 | 136974336.000 |
| Database Page Fault Stalls/sec | 0.000 | 0.000 | 0.000 |

## Test Log

```
4/7/2010 5:46:12 AM -- Jetstress testing begins ...
4/7/2010 5:46:12 AM -- Prepare testing begins ...
4/7/2010 5:46:17 AM -- Attaching databases ...
4/7/2010 5:46:17 AM -- Prepare testing ends.
4/7/2010 5:46:17 AM -- Dispatching transactions begins ...
4/7/2010 5:46:17 AM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2
GB)
4/7/2010 5:46:17 AM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
4/7/2010 5:46:23 AM -- Database read latency thresholds: (average: 20 msec/read,
maximum: 100 msec/read).
4/7/2010 5:46:23 AM -- Log write latency thresholds: (average: 10 msec/write,
maximum: 100 msec/write).
4/7/2010 5:46:27 AM -- Operation mix: Sessions 7, Inserts 40%, Deletes 20%,
Replaces 5%, Reads 35%, Lazy Commits 70%.
4/7/2010 5:46:27 AM -- Performance logging begins (interval: 15000 ms).
4/7/2010 5:46:27 AM -- Generating log files ...
4/7/2010 6:44:59 AM -- C:\alogluns\log1 (101.0% generated), C:\alogluns\log2
(101.8% generated), C:\alogluns\log3 (100.2% generated), C:\alogluns\log4 (101.2%
generated) and C:\alogluns\log5 (100.2% generated)
4/7/2010 6:44:59 AM -- Performance logging ends.
4/7/2010 6:44:59 AM -- JetInterop batch transaction stats: 21688, 21932, 21544,
21670 and 21557.
4/7/2010 6:44:59 AM -- Dispatching transactions ends.
4/7/2010 6:44:59 AM -- Shutting down databases ...
4/7/2010 6:45:04 AM -- Instance3676.1 (complete), Instance3676.2 (complete),
Instance3676.3 (complete), Instance3676.4 (complete) and Instance3676.5
(complete)
4/7/2010 6:45:04 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Soft
Recovery\Performance_2010_4_7_5_46_23.blg has 233 samples.
4/7/2010 6:45:04 AM -- Creating test report ...
4/7/2010 6:45:05 AM -- Instance3676.1 has 13.5 for I/O Database Reads Average
Latency.
4/7/2010 6:45:05 AM -- Instance3676.1 has 1.2 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.1 has 1.2 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.2 has 12.5 for I/O Database Reads Average
Latency.
4/7/2010 6:45:05 AM -- Instance3676.2 has 1.0 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.2 has 1.0 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.3 has 12.5 for I/O Database Reads Average
Latency.
4/7/2010 6:45:05 AM -- Instance3676.3 has 1.2 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.3 has 1.2 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.4 has 12.6 for I/O Database Reads Average
Latency.
4/7/2010 6:45:05 AM -- Instance3676.4 has 1.0 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.4 has 1.0 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.5 has 12.5 for I/O Database Reads Average
```

```
Latency.
4/7/2010 6:45:05 AM -- Instance3676.5 has 1.2 for I/O Log Writes Average Latency.
4/7/2010 6:45:05 AM -- Instance3676.5 has 1.2 for I/O Log Reads Average Latency.
4/7/2010 6:45:05 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
4/7/2010 6:45:05 AM -- Test has 0 Database Page Fault Stalls/sec samples higher
than 0.
4/7/2010 6:45:05 AM -- C:\ESRP 3.0_1GB MBoxx_RAID 5\Soft
Recovery\Performance_2010_4_7_5_46_23.xml has 232 samples queried.
```

Hitachi Data Systems Corporation

**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


ESRP-051-00 April 2010