HITACHI
Inspire the Next

# Hitachi Adaptable Modular Storage 2100 Dynamically Provisioned 11,200 Mailbox Exchange 2010 Mailbox Resiliency Storage Solution

Tested with: ESRP – Storage Version 3.0

Test Date: March-April 2011

# 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

# Hitachi Adaptable Modular Storage 2100 Dynamically Provisioned 11,200 Mailbox Exchange 2010 Mailbox Resiliency Storage Solution

Tested with: ESRP – Storage Version 3.0

Test Date: March-April 2011

## Overview

This document provides information on a Hitachi Adaptable Modular Storage 2100 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 11,200 users and four servers. This testing used a Hitachi Adaptable Modular Storage 2100 (AMS 2100) 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 AMS 2100.

Testing used four 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 four DAGs, each containing two database copies and two servers (one simulated). The test configuration was capable of supporting 11,200 users with a 0.18 IOPS per user profile and user mailbox size of 3GB.

An AMS 2100 with the following was used for these tests:

- 108 600GB 15k RPM SAS disks
- 8GB of cache
- Eight 8Gb/sec paths (four used)

Hitachi Adaptable Modular Storage 2100 is a 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, AMS 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 an AMS 2100 to meet the needs of a large Exchange Server deployment.

This solution uses Hitachi Dynamic Provisioning, which is enabled on AMS 2100 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 2100 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 2100. 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 four DAGs, with each DAG configured with two servers (one simulated) that host active mailboxes in five databases.

To target the 11,200-user resiliency solution, a Hitachi Adaptable Modular Storage 2100 (AMS 2100) was configured with 108 disks (of a maximum of120). Four servers (one per DAG) were used, with each server configured with 2,800 mailboxes. There were 20 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 simulated DAG configuration. For more information, see the Hitachi Data Systems Storage Systems web page.

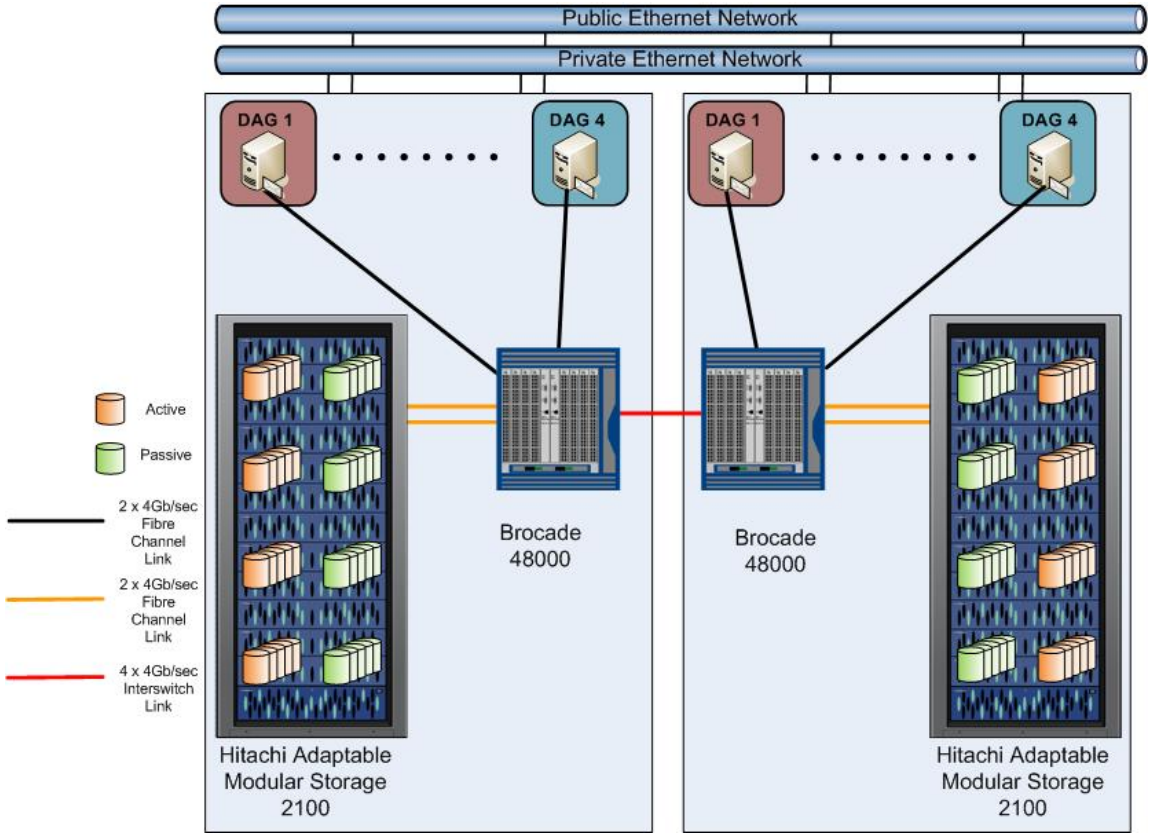


**Figure 1**

This solution enables organizations to consolidate Exchange Server 2010 DAG deployments on two AMS 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, additional Exchange Server 2010 mailbox servers can be easily added to support these failover scenarios.

Table 1 illustrates how the disks in an AMS 2100 were organized into RAID groups for use by the databases or logs. Each set of colored disks represents a RAID-5 (8D+1P) RAID group. There were 108 disks used in these tests configured as 12 RAID groups for the Exchange databases and logs.

**Table1. 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 | 15 | 16 | 17 | 18 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| RKA 6 | 11 | 11 | 11 | 11 | 12 | 12 | 12 | 12 | 12 | 12 | 12 | 12 | 12 | 0 | 0 | 0 | 0 | 0 | 0 |
| RKA 5 | 9 | 9 | 9 | 9 | 9 | 10 | 10 | 10 | 10 | 10 | 10 | 10 | 10 | 10 | 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 |
| RKA 3 | 5 | 5 | 5 | 5 | 5 | 5 | 5 | 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 | 5 | 5 |
| RKA 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 3 |
| RKA 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | | | | |

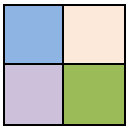

DB RAID Groups
Log RAID Groups

Disk tray RKA-0 was the internal 15-disk tray that wasn't used during these tests. Trays RKA-1 through tray RKA-6 each held 19 600GB 15K SAS disks. Tray RKA-6 had 6 disks used as spares. There were actually three dense trays, but each is accessed as two separate trays, so trays RKA-1 and RKA-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 10 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 2,000GB size limit) were created for 20 databases (five per server). From the log pool, 20 DP-VOLs (each specified to have a size limit of 200GB) were created for 20 logs (five 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.

**Table2. Adaptable Modular Storage 2100 Port to Server Layout**

| Server | Primary path | Secondary path |
|---|---|---|
| SUN158 | 0A | 1A |
| SUN159 | 0B | 1B |
| SUN160 | 1A | 0A |
| SUN161 | 1B | 0B |

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.

**Table3. Adaptable Modular Storage 2100 Port to Database DP-VOL Layout**

| Port | Database | DB DP-VOL |
|---|---|---|
| 0A | Databases 1-5 | 0-4 |
| 0B | Databases 6-10 | 5-9 |
| 1A | Databases 11-15 | 10-14 |
| 1B | Databases 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 would be deployed on the replicated storage and servers for this solution.

**Table4. Adaptable Modular Storage 2100 Port to Log DP-VOL Layout**

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

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.

**Table5. Adaptable Modular Storage 2100 Configuration Details**

| Host | Pool | Port | DP-VOLs | Size (GB) | Description |
|---|---|---|---|---|---|
| Sun158 | 0 | 0A/1A | 0-4 | 2000 | Databases 1-5 |
| Sun159 | 0 | 0B/1B | 5-9 | 2000 | Databases 6-10 |
| Sun160 | 0 | 1A/0A | 10-14 | 2000 | Databases 11-15 |
| Sun161 | 0 | 1B/0B | 15-19 | 2000 | Databases 16-20 |
| Sun158 | 1 | 0A/1A | 20-24 | 200 | Log 1-5 |
| Sun159 | 1 | 0B/1B | 25-29 | 200 | Log 6-10 |
| Sun160 | 1 | 1A/0A | 30-34 | 200 | Log 11-15 |
| Sun161 | 1 | 1B/0B | 35-39 | 200 | Log 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 11,200 Exchange users with the following specifications:

- 8 Exchange servers (four servers tested, four servers simulated for the database copies)
- 4 database availability groups each with two servers (one simulated) and two copies per database
- 2 Adaptable Modular Storage 2100s (one tested)
- 0.15 IOPS per user (0.18 tested for 20 percent growth)
- 3GB mailbox size
- Mailbox resiliency provides high-availability and used as primary data protection mechanism.
- Adaptable Modular Storage RAID protection against physical failure or loss.
- 24×7 background database maintenance enabled.

# Test Deployment

The following tables summarize the testing environment.

**Table 6. Simulated Exchange Configuration**

| | |
|---|---|
| ***Number of Exchange mailboxes simulated*** | 11,200 |
| ***Number of database availability groups (DAGs)*** | 4 |
| ***Number of servers per DAG*** | 2 (1 simulated) |
| ***Number of active mailboxes per server*** | 2,800 |
| ***Number of databases per host*** | 5 |
| ***Number of copies per database*** | 2 |
| ***Number of mailboxes per database*** | 560 |
| ***Simulated profile: I/Os per second per mailbox (IOPS, include 20% headroom)*** | 0.18 |
| ***Database LU size*** | 2000GB |
| ***Log LU siz*** | 200GB |
| ***Total database size for performance testing*** | 33,600GB |
| ***% storage capacity used by Exchange database**** | 79.2% |

**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**

| | |
|---|---|
| *Storage connectivity (Fibre Channel, SAS, SATA, iSCSI)* | Fibre Channel |
| *Storage model and OS/firmware revision* | 1 Hitachi Adaptable Modular Storage 2100<br>Firmware: 0897/A-X<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/sec (4 x 4Gb/sec ports) |
| *Switch type/model/firmware revision* | Brocade 5320, Fabric OS v6.4.0b |
| *HBA model and firmware* | Emulex LightPulse LPe 11002-S<br>FW : 2.82A3 |
| *Number of HBAs per host* | 2 dual-ported HBA per host, 1 4Gb/sec 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* | 108 |
| *Maximum number of spindles that can be hosted in the storage* | 120 |

**Table 8. Storage Software**

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

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

| | |
|---|---|
| *Disk type, speed and firmware revision* | SAS 600GB 15K 5C53 |
| *Raw capacity per disk (GB)* | 600GB |
| *Number of physical disks in test* | 90 (dynamic provisioning pool) |
| *Total raw storage capacity (GB)* | 54,000 |
| *Disk slice size (GB)* | N /A |
| *Number of slices per LU or number of disks per LU* | N/A |
| *RAID level* | RAID-5 (8D+1P) at storage level |
| *Total formatted capacity* | 42,400GB |
| *Storage capacity utilization* | 78.5% |
| *Database capacity utilization* | 74.1% |

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

| | |
|---|---|
| *Disk type, speed and firmware revision* | SAS 600GB 15K 5C53 |
| *Raw capacity per disk (GB)* | 600GB |
| *Number of spindles in test* | 18(dynamic provisioning pool) |
| *Total raw storage capacity (GB)* | 10,800 |
| *Disk slice size (GB)* | N/A |
| *Number of slices per LU or number of disks per LU* | N/A |
| *RAID level* | RAID-5 (8D+1P) at storage level |
| *Total formatted capacity* | 8,480GB |

# Replication Configuration

**Table 11. Replication 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 (1Gb/sec) |

**Table 12. Replicated Storage Hardware**

| | |
|---|---|
| *Storage connectivity (Fiber Channel, SAS, SATA, iSCSI)* | Fibre Channel |
| *Storage model and OS/firmware revision* | 1 Hitachi Adaptable Modular Storage 2100<br>Firmware: 0897/A-X<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/sec (4 x 4Gb/sec ports) |
| *Switch type/model/firmware revision* | Brocade 5320, Fabric OS v6.4.0b |
| *HBA model and firmware* | Emulex LightPulse LPe 11002-S<br>FW : 2.82A3 |
| *Number of HBAs per host* | 2 dual-ported HBA per host, 1 4Gb/sec 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* | 108 |
| *Maximum number of spindles that can be hosted in the storage* | 120 |

**Table 13. Replicated Storage Software**

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

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

| | |
|---|---|
| *Disk type, speed and firmware revision* | SAS 600GB 15K 5C53 |
| *Raw capacity per disk (GB)* | 600GB |
| *Number of physical disks in test* | 90 (dynamic provisioning pool) |
| *total raw storage capacity (GB)* | 54,000 |
| *Disk slice size (GB)* | N /A |
| *Number of slices per LU or number of disks per LU* | N/A |
| *Raid level* | RAID-5 (8D+1P) at storage level |
| *Total formatted capacity* | 42,400GB |
| *Storage capacity utilization* | 78.5% |
| *Database capacity utilization* | 74.1% |

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

| | |
|---|---|
| *Disk type, speed and firmware revision* | SAS 600GB 15K 5C53 |
| *Raw capacity per disk (GB)* | 600GB |
| *Number of spindles in test* | 18 (dynamic provisioning pool) |
| *Total raw storage capacity (GB)* | 10,800 |
| *Disk slice size (GB)* | N/A |
| *Number of slices per LU or number of disks per LU* | N/A |
| *Raid level* | RAID-5 (8D+1P) at storage level |
| *Total formatted capacity* | 8,480GB |

# 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."

## 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 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 2300.
13. Use fewer, larger LUs for Exchange 2010 databases (up to 2TB) with Background Database Maintenance (24×7) 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*

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 (SUN158)**

| | |
|---|---|
| ***Database I/O*** | |
| ***Database Disk Transfers Per Second*** | 899 |
| ***Database Disk Reads Per Second*** | 566 |
| ***Database Disk Writes Per Second*** | 333 |
| ***Average Database Disk Read Latency (ms)*** | 11.5 |
| ***Average Database Disk Write Latency (ms)*** | 3.1 |
| ***Transaction Log I/O*** | |
| ***Log Disk Writes Per Second*** | 263 |
| ***Average Log Disk Write Latency (ms)*** | 1.3 |

**Table 17. Individual Server Metrics for Exchange Server (SUN159)**

| *Database I/O* | |
|---|---|
| *Database Disk Transfers Per Second* | 943 |
| *Database Disk Reads Per Second* | 595 |
| *Database Disk Writes Per Second* | 347 |
| *Average Database Disk Read Latency (ms)* | 11.2 |
| *Average Database Disk Write Latency (ms)* | 3.0 |
| *Transaction Log I/O* | |
| *Log Disk Writes Per Second* | 280 |
| *Average Log Disk Write Latency (ms)* | 1.0 |

**Table 18. Individual Server Metrics for Exchange Server (SUN160)**

| *Database I/O* | |
|---|---|
| *Database Disk Transfers Per Second* | 920 |
| *Database Disk Reads Per Second* | 581 |
| *Database Disk Writes Per Second* | 339 |
| *Average Database Disk Read Latency (ms)* | 11.4 |
| *Average Database Disk Write Latency (ms)* | 2.9 |
| *Transaction Log I/O* | |
| *Log Disk Writes Per Second* | 274 |
| *Average Log Disk Write Latency (ms)* | 1.0 |

**Table 19. Individual Server Metrics for Exchange Server (SUN161)**

| *Database I/O* | |
|---|---|
| *Database Disk Transfers Per Second* | 933 |
| *Database Disk Reads Per Second* | 592 |
| *Database Disk Writes Per Second* | 341 |
| *Average Database Disk Read Latency (ms)* | 11.2 |
| *Average Database Disk Write Latency (ms)* | 2.9 |
| *Transaction Log I/O* | |
| *Log Disk Writes Per Second* | 298 |
| *Average Log Disk Write Latency (ms)* | 1.0 |

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

| | |
|---|---|
| ***Database I/O*** | |
| ***Database Disk Transfers Per Second*** | 3694.034 |
| ***Database Disk Reads Per Second*** | 2334.668 |
| ***Database Disk Writes Per Second*** | 1359.366 |
| ***Average Database Disk Read Latency (ms)*** | 11.312 |
| ***Average Database Disk Write Latency (ms)*** | 2.967 |
| ***Transaction Log I/O*** | |
| ***Log Disk Writes Per Second*** | 1115.149 |
| ***Average Log Disk Write Latency (ms)*** | 1.099 |

## 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*** | 45.60 |
| ***MB Read Per Second Total Per Server*** | 228.02 |

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

| | |
|---|---|
| ***Average Time to Play One Log File (sec)*** | 1.19776 |

# Conclusion

This document details a tested and robust Exchange Server 2010 Resiliency solution capable of supporting 11,200 users with a 0.18 IOPS per user profile and user mailbox size of 3GB using four DAG's each configured with 2 server nodes (one simulated).

A Hitachi Adaptable Modular Storage 2100 (AMS 2100), with 8GB of cache and four 8Gb/sec Fibre Channel host paths (limited to 4Gbps by the host ports), using Hitachi Dynamic Provisioning (with two pools) and 108 600GB 15K RPM SAS disks in a RAID-5 (8D+1P) configuration was used for these tests.

Testing confirmed that an AMS 2100 is more than capable of delivering the IOPS and capacity requirements needed to support the active and replicated databases for 11,200 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: SUN158

**Test Summary**

| | |
|---|---|
| ***Overall Test Result*** | Pass |
| ***Machine Name*** | SUN158 |
| ***Test Description*** | |
| ***Test Start Time*** | 2/7/2011 11:10:54 PM |
| ***Test End Time*** | 2/8/2011 1:44:41 AM |
| ***Collection Start Time*** | 2/7/2011 11:12:54 PM |
| ***Collection End Time*** | 2/8/2011 1:12:53 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 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Performance Test\Performance_2011_2_7_23_11_6.blg<br>C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Performance Test\DBChecksum_2011_2_8_1_44_41.blg |

**Database Sizing and Throughput**

| | |
|---|---|
| ***Achieved Transactional I/O per Second*** | 898.932 |
| ***Target Transactional I/O per Second*** | 504 |
| ***Initial Database Size (bytes)*** | 9276032352256 |
| ***Final Database Size (bytes)*** | 9279303909376 |
| ***Database Files (Count)*** | 5 |

### Jetstress System Parameters

| | |
|---|---|
| ***Thread Count*** | 6 (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

| | |
|---|---|
| ***Instance2116.1*** | Log Path: C:\logluns\log1<br>Database: C:\dbluns\db1\Jetstress001001.edb |
| ***Instance2116.2*** | Log Path: C:\logluns\log2<br>Database: C:\dbluns\db2\Jetstress002001.edb |
| ***Instance2116.3*** | Log Path: C:\logluns\log3<br>Database: C:\dbluns\db3\Jetstress003001.edb |
| ***Instance2116.4*** | Log Path: C:\logluns\log4<br>Database: C:\dbluns\db4\Jetstress004001.edb |
| ***Instance2116.5*** | Log Path: C:\logluns\log5<br>Database: C:\dbluns\db5\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* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance2116.1 | 12.306 | 3.891 | 112.188 | 66.020 | 33568.038 | 34950.162 | 0.000 | 1.383 | 0.000 | 52.450 | 0.000 | 4797.178 |
| Instance2116.2 | 11.031 | 3.148 | 113.920 | 66.899 | 33756.388 | 34925.368 | 0.000 | 1.164 | 0.000 | 52.828 | 0.000 | 4771.425 |
| Instance2116.3 | 11.480 | 3.361 | 111.530 | 65.461 | 33667.078 | 34948.820 | 0.000 | 1.388 | 0.000 | 51.785 | 0.000 | 4807.665 |
| Instance2116.4 | 11.249 | 2.388 | 114.565 | 67.306 | 33644.261 | 34919.261 | 0.000 | 1.154 | 0.000 | 53.357 | 0.000 | 4763.756 |
| Instance2116.5 | 11.464 | 2.485 | 114.040 | 67.003 | 33620.960 | 34890.827 | 0.000 | 1.399 | 0.000 | 52.572 | 0.000 | 4797.022 |

**Background Database Maintenance I/O Performance**

| *MSExchange Database ==> Instances* | *Database Maintenance IO Reads/sec* | *Database Maintenance IO Reads Average Bytes* |
|---|---|---|
| Instance2116.1 | 28.875 | 261938.699 |
| Instance2116.2 | 30.172 | 261873.800 |
| Instance2116.3 | 28.998 | 261878.697 |
| Instance2116.4 | 29.591 | 261917.912 |
| Instance2116.5 | 29.188 | 261928.300 |

**Log Replication I/O Performance**

| *MSExchange Database ==> Instances* | *I/O Log Reads/sec* | *I/O Log Reads Average Bytes* |
|---|---|---|
| Instance2116.1 | 1.024 | 232614.046 |
| Instance2116.2 | 1.024 | 232549.982 |
| Instance2116.3 | 1.013 | 232561.778 |
| Instance2116.4 | 1.033 | 232561.778 |
| Instance2116.5 | 1.027 | 232048.800 |

**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* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance2116.1 | 12.306 | 3.891 | 141.063 | 66.020 | 80314.904 | 34950.162 | 5.688 | 1.383 | 1.024 | 52.450 | 232614.046 | 4797.178 |
| Instance2116.2 | 11.031 | 3.148 | 144.093 | 66.899 | 81523.316 | 34925.368 | 1.889 | 1.164 | 1.024 | 52.828 | 232549.982 | 4771.425 |
| Instance2116.3 | 11.480 | 3.361 | 140.528 | 65.461 | 80757.819 | 34948.820 | 4.615 | 1.388 | 1.013 | 51.785 | 232561.778 | 4807.665 |
| Instance2116.4 | 11.249 | 2.388 | 144.157 | 67.306 | 80502.471 | 34919.261 | 2.115 | 1.154 | 1.033 | 53.357 | 232561.778 | 4763.756 |
| Instance2116.5 | 11.464 | 2.485 | 143.228 | 67.003 | 80146.923 | 34890.827 | 5.322 | 1.399 | 1.027 | 52.572 | 232048.800 | 4797.022 |

Host System Performance

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| % Processor Time | 0.986 | 0.000 | 2.992 |
| Available MBytes | 29405.206 | 29393.000 | 29506.000 |
| Free System Page Table Entries | 33555643.023 | 33555642.000 | 33555645.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72084462.933 | 72056832.000 | 72142848.000 |
| Pool Paged Bytes | 126206711.467 | 126169088.000 | 126259200.000 |
| Database Page Fault Stalls/sec | 0.000 | 0.000 | 0.000 |

Test Log2/7/2011 11:10:54 PM -- Jetstress testing begins ...
2/7/2011 11:10:54 PM -- Prepare testing begins ...
2/7/2011 11:11:00 PM -- Attaching databases ...
2/7/2011 11:11:00 PM -- Prepare testing ends.
2/7/2011 11:11:00 PM -- Dispatching transactions begins ...
2/7/2011 11:11:00 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
2/7/2011 11:11:00 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
2/7/2011 11:11:06 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
2/7/2011 11:11:06 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
2/7/2011 11:11:12 PM -- Operation mix: Sessions 6, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
2/7/2011 11:11:12 PM -- Performance logging begins (interval: 15000 ms).
2/7/2011 11:11:12 PM -- Attaining prerequisites:
2/7/2011 11:12:54 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1222320000.0 (lower bound: 1207960000.0, upper bound: none)
2/8/2011 1:12:55 AM -- Performance logging ends.
2/8/2011 1:44:37 AM -- JetInterop batch transaction stats: 45771, 46245, 45386, 46121 and 46190.
2/8/2011 1:44:37 AM -- Dispatching transactions ends.
2/8/2011 1:44:37 AM -- Shutting down databases ...
2/8/2011 1:44:41 AM -- Instance2116.1 (complete), Instance2116.2 (complete), Instance2116.3 (complete), Instance2116.4 (complete) and Instance2116.5 (complete)
2/8/2011 1:44:43 AM -- Performance logging begins (interval: 30000 ms).
2/8/2011 1:44:43 AM -- Verifying database checksums ...
2/8/2011 11:26:09 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 11:26:09 AM -- Performance logging ends.
2/8/2011 11:26:09 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Performance Test\DBChecksum_2011_2_8_1_44_41.blg has 1162 samples.
2/8/2011 11:26:14 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Performance Test\DBChecksum_2011_2_8_1_44_41.html is saved.
2/8/2011 11:26:14 AM -- Verifying log checksums ...
2/8/2011 11:26:15 AM -- C:\logluns\log1 (12 log(s) processed), C:\logluns\log2 (12 log(s) processed), C:\logluns\log3 (11 log(s) processed), C:\logluns\log4 (12 log(s) processed) and

C:\logluns\log5 (11 log(s) processed)
2/8/2011 11:26:15 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Performance Test\Performance_2011_2_7_23_11_6.blg has 486 samples.
2/8/2011 11:26:15 AM -- Creating test report ...
2/8/2011 11:26:18 AM -- Instance2116.1 has 12.3 for I/O Database Reads Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.1 has 1.4 for I/O Log Writes Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.1 has 1.4 for I/O Log Reads Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.2 has 11.0 for I/O Database Reads Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.2 has 1.2 for I/O Log Writes Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.2 has 1.2 for I/O Log Reads Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.3 has 11.5 for I/O Database Reads Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.3 has 1.4 for I/O Log Writes Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.3 has 1.4 for I/O Log Reads Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.4 has 11.2 for I/O Database Reads Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.4 has 1.2 for I/O Log Writes Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.4 has 1.2 for I/O Log Reads Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.5 has 11.5 for I/O Database Reads Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.5 has 1.4 for I/O Log Writes Average Latency.
2/8/2011 11:26:18 AM -- Instance2116.5 has 1.4 for I/O Log Reads Average Latency.
2/8/2011 11:26:18 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
2/8/2011 11:26:18 AM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
2/8/2011 11:26:18 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Performance Test\Performance_2011_2_7_23_11_6.xml has 479 samples queried.

## Performance Test Database Checksums Result: SUN158

**Checksum Statistics - All**

| *Database* | *Seen pages* | *Bad pages* | *Correctable pages* | *Wrong page-number pages* | *File length / seconds taken* |
|---|---|---|---|---|---|
| C:\dbluns\db1\Jetstress001001.edb | 56635250 | 0 | 0 | 0 | 1769851 MBytes / 34885 sec |
| C:\dbluns\db2\Jetstress002001.edb | 56637042 | 0 | 0 | 0 | 1769907 MBytes / 28722 sec |
| C:\dbluns\db3\Jetstress003001.edb | 56638066 | 0 | 0 | 0 | 1769939 MBytes / 34868 sec |
| C:\dbluns\db4\Jetstress004001.edb | 56635250 | 0 | 0 | 0 | 1769851 MBytes / 28704 sec |
| C:\dbluns\db5\Jetstress005001.edb | 56636274 | 0 | 0 | 0 | 1769883 MBytes / 34834 sec |
| (Sum) | 283181882 | 0 | 0 | 0 | 8849433 MBytes / 34886 sec |

**Disk Subsystem Performance (of checksum)**

| LogicalDisk | Avg. Disk sec/Read | Avg. Disk sec/Write | Disk Reads/sec | Disk Writes/sec | Avg. Disk Bytes/Read |
|---|---|---|---|---|---|
| C:\dbluns\db1 | 0.068 | 0.000 | 811.241 | 0.000 | 65536.000 |
| C:\dbluns\db2 | 0.057 | 0.000 | 985.931 | 0.000 | 65536.000 |
| C:\dbluns\db3 | 0.067 | 0.000 | 810.264 | 0.000 | 65536.000 |
| C:\dbluns\db4 | 0.057 | 0.000 | 986.135 | 0.000 | 65536.000 |
| C:\dbluns\db5 | 0.066 | 0.000 | 812.266 | 0.000 | 65536.000 |

**Memory System Performance (of checksum)**

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| % Processor Time | 1.453 | 0.000 | 4.476 |
| Available MBytes | 30777.249 | 30766.000 | 30791.000 |
| Free System Page Table Entries | 33555642.250 | 33555613.000 | 33555645.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72136778.024 | 72052736.000 | 74739712.000 |
| Pool Paged Bytes | 126660276.654 | 125571072.000 | 128311296.000 |

Test Log2/7/2011 11:10:54 PM -- Jetstress testing begins ...
2/7/2011 11:10:54 PM -- Prepare testing begins ...
2/7/2011 11:11:00 PM -- Attaching databases ...
2/7/2011 11:11:00 PM -- Prepare testing ends.
2/7/2011 11:11:00 PM -- Dispatching transactions begins ...
2/7/2011 11:11:00 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
2/7/2011 11:11:00 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
2/7/2011 11:11:06 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
2/7/2011 11:11:06 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
2/7/2011 11:11:12 PM -- Operation mix: Sessions 6, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
2/7/2011 11:11:12 PM -- Performance logging begins (interval: 15000 ms).
2/7/2011 11:11:12 PM -- Attaining prerequisites:
2/7/2011 11:12:54 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1222320000.0 (lower bound: 1207960000.0, upper bound: none)
2/8/2011 1:12:55 AM -- Performance logging ends.
2/8/2011 1:44:37 AM -- JetInterop batch transaction stats: 45771, 46245, 45386, 46121 and 46190.
2/8/2011 1:44:37 AM -- Dispatching transactions ends.
2/8/2011 1:44:37 AM -- Shutting down databases ...
2/8/2011 1:44:41 AM -- Instance2116.1 (complete), Instance2116.2 (complete), Instance2116.3 (complete), Instance2116.4 (complete) and Instance2116.5 (complete)
2/8/2011 1:44:43 AM -- Performance logging begins (interval: 30000 ms).
2/8/2011 1:44:43 AM -- Verifying database checksums ...
2/8/2011 11:26:09 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 11:26:09 AM -- Performance logging ends.
2/8/2011 11:26:09 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Performance Test\DBChecksum_2011_2_8_1_44_41.blg has 1162 samples.

# Stress Test Result Report: SUN158

**Test Summary**

| | |
|---|---|
| ***Overall Test Result*** | Pass |
| ***Machine Name*** | SUN158 |
| ***Test Description*** | |
| ***Test Start Time*** | 2/8/2011 7:50:40 PM |
| ***Test End Time*** | 2/9/2011 8:52:22 PM |
| ***Collection Start Time*** | 2/8/2011 7:52:44 PM |
| ***Collection End Time*** | 2/9/2011 7:52:40 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 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Stress Test\Stress_2011_2_8_19_50_52.blg<br>C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Stress Test\DBChecksum_2011_2_9_20_52_22.blg |

**Database Sizing and Throughput**

| | |
|---|---|
| ***Achieved Transactional I/O per Second*** | 897.05 |
| ***Target Transactional I/O per Second*** | 504 |
| ***Initial Database Size (bytes)*** | 9279303909376 |
| ***Final Database Size (bytes)*** | 9311256117248 |
| ***Database Files (Count)*** | 5 |

### Jetstress System Parameters

| | |
|---|---|
| ***Thread Count*** | 6 (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

| | |
|---|---|
| ***Instance2116.1*** | Log Path: C:\logluns\log1<br>Database: C:\dbluns\db1\Jetstress001001.edb |
| ***Instance2116.2*** | Log Path: C:\logluns\log2<br>Database: C:\dbluns\db2\Jetstress002001.edb |
| ***Instance2116.3*** | Log Path: C:\logluns\log3<br>Database: C:\dbluns\db3\Jetstress003001.edb |
| ***Instance2116.4*** | Log Path: C:\logluns\log4<br>Database: C:\dbluns\db4\Jetstress004001.edb |
| ***Instance2116.5*** | Log Path: C:\logluns\log5<br>Database: C:\dbluns\db5\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* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance2116.1 | 12.019 | 2.932 | 112.923 | 66.438 | 33586.004 | 34920.229 | 0.000 | 1.255 | 0.000 | 52.808 | 0.000 | 4759.652 |
| Instance2116.2 | 11.400 | 3.982 | 113.137 | 66.569 | 33651.691 | 34911.334 | 0.000 | 1.247 | 0.000 | 52.596 | 0.000 | 4777.295 |
| Instance2116.3 | 11.656 | 3.417 | 112.802 | 66.375 | 33637.638 | 34921.345 | 0.000 | 1.270 | 0.000 | 52.562 | 0.000 | 4779.056 |
| Instance2116.4 | 11.651 | 2.650 | 112.759 | 66.298 | 33675.805 | 34908.335 | 0.000 | 1.238 | 0.000 | 52.390 | 0.000 | 4769.146 |
| Instance2116.5 | 11.133 | 2.204 | 113.184 | 66.564 | 33701.943 | 34909.424 | 0.000 | 1.246 | 0.000 | 52.624 | 0.000 | 4764.085 |

### Background Database Maintenance I/O Performance

| *MSExchange Database ==> Instances* | *Database Maintenance IO Reads/sec* | *Database Maintenance IO Reads Average Bytes* |
|---|---|---|
| Instance2116.1 | 29.057 | 261921.744 |
| Instance2116.2 | 29.823 | 261929.097 |
| Instance2116.3 | 29.211 | 261919.824 |
| Instance2116.4 | 29.173 | 261913.525 |
| Instance2116.5 | 29.437 | 261921.831 |

### Log Replication I/O Performance

| *MSExchange Database ==> Instances* | *I/O Log Reads/sec* | *I/O Log Reads Average Bytes* |
|---|---|---|
| Instance2116.1 | 1.023 | 232470.736 |
| Instance2116.2 | 1.023 | 232395.220 |
| Instance2116.3 | 1.022 | 232424.041 |
| Instance2116.4 | 1.017 | 232484.606 |
| Instance2116.5 | 1.020 | 232437.630 |

**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* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance2116.1 | 12.019 | 2.932 | 141.980 | 66.438 | 80316.360 | 34920.229 | 4.396 | 1.255 | 1.023 | 52.808 | 232470.736 | 4759.652 |
| Instance2116.2 | 11.400 | 3.982 | 142.960 | 66.569 | 81272.587 | 34911.334 | 2.333 | 1.247 | 1.023 | 52.596 | 232395.220 | 4777.295 |
| Instance2116.3 | 11.656 | 3.417 | 142.013 | 66.375 | 80593.522 | 34921.345 | 4.387 | 1.270 | 1.022 | 52.562 | 232424.041 | 4779.056 |
| Instance2116.4 | 11.651 | 2.650 | 141.932 | 66.298 | 80588.134 | 34908.335 | 2.408 | 1.238 | 1.017 | 52.390 | 232484.606 | 4769.146 |
| Instance2116.5 | 11.133 | 2.204 | 142.622 | 66.564 | 80807.056 | 34909.424 | 4.489 | 1.246 | 1.020 | 52.624 | 232437.630 | 4764.085 |

**Host System Performance**

| *Counter* | *Average* | *Minimum* | *Maximum* |
|---|---|---|---|
| % Processor Time | 1.020 | 0.000 | 4.984 |
| Available MBytes | 29394.644 | 29377.000 | 29409.000 |
| Free System Page Table Entries | 33555642.991 | 33555635.000 | 33555645.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72093488.743 | 72089600.000 | 72163328.000 |
| Pool Paged Bytes | 126104569.768 | 125612032.000 | 129900544.000 |
| Database Page Fault Stalls/sec | 0.000 | 0.000 | 0.000 |

Test Log2/8/2011 7:50:40 PM -- Jetstress testing begins ...
2/8/2011 7:50:40 PM -- Prepare testing begins ...
2/8/2011 7:50:46 PM -- Attaching databases ...
2/8/2011 7:50:46 PM -- Prepare testing ends.
2/8/2011 7:50:46 PM -- Dispatching transactions begins ...
2/8/2011 7:50:46 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
2/8/2011 7:50:46 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
2/8/2011 7:50:52 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
2/8/2011 7:50:52 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
2/8/2011 7:50:58 PM -- Operation mix: Sessions 6, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
2/8/2011 7:50:58 PM -- Performance logging begins (interval: 15000 ms).
2/8/2011 7:50:58 PM -- Attaining prerequisites:
2/8/2011 7:52:44 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1214628000.0 (lower bound: 1207960000.0, upper bound: none)
2/9/2011 7:52:44 PM -- Performance logging ends.
2/9/2011 8:52:19 PM -- JetInterop batch transaction stats: 448482, 448075, 447867, 446457 and 448392.
2/9/2011 8:52:19 PM -- Dispatching transactions ends.
2/9/2011 8:52:19 PM -- Shutting down databases ...
2/9/2011 8:52:22 PM -- Instance2116.1 (complete), Instance2116.2 (complete), Instance2116.3 (complete), Instance2116.4 (complete) and Instance2116.5 (complete)
2/9/2011 8:52:23 PM -- Performance logging begins (interval: 30000 ms).
2/9/2011 8:52:23 PM -- Verifying database checksums ...
2/10/2011 7:00:35 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)
2/10/2011 7:00:35 AM -- Performance logging ends.
2/10/2011 7:00:35 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Stress Test\DBChecksum_2011_2_9_20_52_22.blg has 1215 samples.
2/10/2011 7:00:40 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Stress Test\DBChecksum_2011_2_9_20_52_22.html is saved.
2/10/2011 7:00:40 AM -- Verifying log checksums ...
2/10/2011 7:00:41 AM -- C:\logluns\log1 (11 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 (12 log(s) processed)
2/10/2011 7:00:41 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Stress Test\Stress_2011_2_8_19_50_52.blg has 5758 samples.
2/10/2011 7:00:41 AM -- Creating test report ...
2/10/2011 7:01:15 AM -- Instance2116.1 has 12.0 for I/O Database Reads Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.1 has 1.3 for I/O Log Writes Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.1 has 1.3 for I/O Log Reads Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.2 has 11.4 for I/O Database Reads Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.2 has 1.2 for I/O Log Writes Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.2 has 1.2 for I/O Log Reads Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.3 has 11.7 for I/O Database Reads Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.3 has 1.3 for I/O Log Writes Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.3 has 1.3 for I/O Log Reads Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.4 has 11.7 for I/O Database Reads Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.4 has 1.2 for I/O Log Writes Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.4 has 1.2 for I/O Log Reads Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.5 has 11.1 for I/O Database Reads Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.5 has 1.2 for I/O Log Writes Average Latency.
2/10/2011 7:01:15 AM -- Instance2116.5 has 1.2 for I/O Log Reads Average Latency.
2/10/2011 7:01:15 AM -- Test has 0 Maximum Database Page Fault Stalls/sec.
2/10/2011 7:01:15 AM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
2/10/2011 7:01:15 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Stress Test\Stress_2011_2_8_19_50_52.xml has 5750 samples queried.

## Stress Test Database Checksums Result: SUN158

**Checksum Statistics - All**

| *Database* | *Seen pages* | *Bad pages* | *Correctable pages* | *Wrong page-number pages* | *File length / seconds taken* |
|---|---|---|---|---|---|
| C:\dbluns\db1\Jetstress001001.edb | 56830578 | 0 | 0 | 0 | 1775955 MBytes / 28132 sec |
| C:\dbluns\db2\Jetstress002001.edb | 56832370 | 0 | 0 | 0 | 1776011 MBytes / 36339 sec |
| C:\dbluns\db3\Jetstress003001.edb | 56833138 | 0 | 0 | 0 | 1776035 MBytes / 36492 sec |
| C:\dbluns\db4\Jetstress004001.edb | 56829554 | 0 | 0 | 0 | 1775923 MBytes / 35393 sec |
| C:\dbluns\db5\Jetstress005001.edb | 56831346 | 0 | 0 | 0 | 1775979 MBytes / 28027 sec |
| (Sum) | 284156986 | 0 | 0 | 0 | 8879905 MBytes / 36492 sec |

**Disk Subsystem Performance (of checksum)**

| LogicalDisk | Avg. Disk sec/Read | Avg. Disk sec/Write | Disk Reads/sec | Disk Writes/sec | Avg. Disk Bytes/Read |
|---|---|---|---|---|---|
| C:\dbluns\db1 | 0.056 | 0.000 | 1006.924 | 0.000 | 65536.000 |
| C:\dbluns\db2 | 0.070 | 0.000 | 781.402 | 0.000 | 65536.000 |
| C:\dbluns\db3 | 0.070 | 0.000 | 778.299 | 0.000 | 65536.000 |
| C:\dbluns\db4 | 0.070 | 0.000 | 800.821 | 0.000 | 65536.000 |
| C:\dbluns\db5 | 0.056 | 0.000 | 1013.014 | 0.000 | 65536.000 |

**Memory System Performance (of checksum)**

| Counter | Average | Minimum | Maximum |
|---|---|---|---|
| % Processor Time | 1.414 | 0.000 | 4.594 |
| Available MBytes | 30761.496 | 30749.000 | 30782.000 |
| Free System Page Table Entries | 33555637.600 | 33555635.000 | 33555643.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72243926.492 | 72187904.000 | 72486912.000 |
| Pool Paged Bytes | 127007884.747 | 126537728.000 | 128634880.000 |

Test Log2/8/2011 7:50:40 PM -- Jetstress testing begins ...
2/8/2011 7:50:40 PM -- Prepare testing begins ...
2/8/2011 7:50:46 PM -- Attaching databases ...
2/8/2011 7:50:46 PM -- Prepare testing ends.
2/8/2011 7:50:46 PM -- Dispatching transactions begins ...
2/8/2011 7:50:46 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
2/8/2011 7:50:46 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
2/8/2011 7:50:52 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 200 msec/read).
2/8/2011 7:50:52 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 200 msec/write).
2/8/2011 7:50:58 PM -- Operation mix: Sessions 6, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
2/8/2011 7:50:58 PM -- Performance logging begins (interval: 15000 ms).
2/8/2011 7:50:58 PM -- Attaining prerequisites:
2/8/2011 7:52:44 PM -- \MSExchange Database(JetstressWin)\Database Cache Size, Last: 1214628000.0 (lower bound: 1207960000.0, upper bound: none)
2/9/2011 7:52:44 PM -- Performance logging ends.
2/9/2011 8:52:19 PM -- JetInterop batch transaction stats: 448482, 448075, 447867, 446457 and 448392.
2/9/2011 8:52:19 PM -- Dispatching transactions ends.
2/9/2011 8:52:19 PM -- Shutting down databases ...
2/9/2011 8:52:22 PM -- Instance2116.1 (complete), Instance2116.2 (complete), Instance2116.3 (complete), Instance2116.4 (complete) and Instance2116.5 (complete)
2/9/2011 8:52:23 PM -- Performance logging begins (interval: 30000 ms).
2/9/2011 8:52:23 PM -- Verifying database checksums ...
2/10/2011 7:00:35 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)

2/10/2011 7:00:35 AM -- Performance logging ends.
2/10/2011 7:00:35 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Stress Test\DBChecksum_2011_2_9_20_52_22.blg has 1215 samples.

# Backup Test Result Report: SUN 158

**Database Backup Statistics - All**

| *Database Instance* | *Database Size (MBytes)* | *Elapsed Backup Time* | *MBytes Transferred/sec* |
|---|---|---|---|
| Instance2116.1 | 1775947.59 | 13:41:10 | 36.04 |
| Instance2116.2 | 1776003.59 | 14:27:04 | 34.14 |
| Instance2116.3 | 1776027.59 | 03:27:58 | 142.33 |
| Instance2116.4 | 1775915.59 | 14:25:59 | 34.18 |
| Instance2116.5 | 1775971.59 | 14:25:19 | 34.21 |

**Jetstress System Parameters**

| | |
|---|---|
| *Thread Count* | 6 (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**

| | |
|---|---|
| *Instance2116.1* | Log Path: C:\logluns\log1<br>Database: C:\dbluns\db1\Jetstress001001.edb |
| *Instance2116.2* | Log Path: C:\logluns\log2<br>Database: C:\dbluns\db2\Jetstress002001.edb |
| *Instance2116.3* | Log Path: C:\logluns\log3<br>Database: C:\dbluns\db3\Jetstress003001.edb |
| *Instance2116.4* | Log Path: C:\logluns\log4<br>Database: C:\dbluns\db4\Jetstress004001.edb |
| *Instance2116.5* | Log Path: C:\logluns\log5<br>Database: C:\dbluns\db5\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* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance2116.1 | 9.728 | 0.000 | 144.139 | 0.000 | 262144.00 0 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Instance2116.2 | 10.790 | 0.000 | 136.416 | 0.000 | 262144.00 0 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Instance2116.3 | 2.789 | 0.000 | 568.892 | 0.000 | 262144.00 0 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Instance2116.4 | 10.636 | 0.000 | 136.703 | 0.000 | 262144.00 0 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |
| Instance2116.5 | 10.626 | 0.000 | 136.643 | 0.000 | 262144.00 0 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 | 0.000 |

**Host System Performance**

| *Counter* | *Average* | *Minimum* | *Maximum* |
|---|---|---|---|
| % Processor Time | 0.725 | 0.000 | 2.637 |
| Available MBytes | 30769.080 | 30753.000 | 30777.000 |
| Free System Page Table Entries | 33555642.965 | 33555635.000 | 33555645.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72218384.102 | 72212480.000 | 72232960.000 |
| Pool Paged Bytes | 130065901.387 | 129777664.000 | 133816320.000 |
| Database Page Fault Stalls/sec | 0.000 | 0.000 | 0.000 |

Test Log2/10/2011 8:00:59 PM -- Jetstress testing begins ...
2/10/2011 8:00:59 PM -- Prepare testing begins ...
2/10/2011 8:01:04 PM -- Attaching databases ...
2/10/2011 8:01:04 PM -- Prepare testing ends.
2/10/2011 8:01:13 PM -- Performance logging begins (interval: 30000 ms).
2/10/2011 8:01:13 PM -- Backing up databases ...
2/11/2011 10:28:18 AM -- Performance logging ends.
2/11/2011 10:28:18 AM -- Instance2116.1 (100% processed), Instance2116.2 (100% processed), Instance2116.3 (100% processed), Instance2116.4 (100% processed) and Instance2116.5 (100% processed)
2/11/2011 10:28:18 AM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Backup Test\DatabaseBackup_2011_2_10_20_1_4.blg has 1733 samples.
2/11/2011 10:28:18 AM -- Creating test report ...

## Soft Recovery Test Result Report: SUN158

**Soft-Recovery Statistics - All**

| *Database Instance* | *Log files replayed* | *Elapsed seconds* |
|---|---|---|
| Instance2116.1 | 519 | 643.37633 |
| Instance2116.2 | 500 | 636.4811179 |
| Instance2116.3 | 505 | 355.3686241 |
| Instance2116.4 | 504 | 620.038689 |
| Instance2116.5 | 501 | 632.5031109 |

**Database Configuration**

| | |
|---|---|
| ***Instance2116.1*** | Log Path: C:\logluns\log1<br>Database: C:\dbluns\db1\Jetstress001001.edb |
| ***Instance2116.2*** | Log Path: C:\logluns\log2<br>Database: C:\dbluns\db2\Jetstress002001.edb |
| ***Instance2116.3*** | Log Path: C:\logluns\log3<br>Database: C:\dbluns\db3\Jetstress003001.edb |
| ***Instance2116.4*** | Log Path: C:\logluns\log4<br>Database: C:\dbluns\db4\Jetstress004001.edb |
| ***Instance2116.5*** | Log Path: C:\logluns\log5<br>Database: C:\dbluns\db5\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* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance2116.1 | 24.666 | 15.470 | 949.249 | 4.847 | 35610.427 | 32768.000 | 8.612 | 0.000 | 7.270 | 0.000 | 232571.462 | 0.000 |
| Instance2116.2 | 25.609 | 15.552 | 962.976 | 4.711 | 35415.857 | 32768.000 | 3.318 | 0.000 | 7.070 | 0.000 | 232418.028 | 0.000 |
| Instance2116.3 | 20.060 | 11.978 | 1723.268 | 8.607 | 35419.491 | 32768.000 | 8.830 | 0.000 | 12.923 | 0.000 | 232549.281 | 0.000 |
| Instance2116.4 | 25.742 | 16.244 | 979.472 | 4.870 | 35489.964 | 32768.000 | 3.143 | 0.000 | 7.305 | 0.000 | 232598.427 | 0.000 |
| Instance2116.5 | 25.514 | 15.355 | 967.170 | 4.767 | 35482.919 | 32768.000 | 7.983 | 0.000 | 7.151 | 0.000 | 232463.744 | 0.000 |

**Background Database Maintenance I/O Performance**

| *MSExchange Database ==> Instances* | *Database Maintenance IO Reads/sec* | *Database Maintenance IO Reads Average Bytes* |
|---|---|---|
| Instance2116.1 | 23.663 | 261915.889 |
| Instance2116.2 | 23.544 | 261959.812 |
| Instance2116.3 | 26.413 | 261820.631 |
| Instance2116.4 | 23.162 | 261970.307 |
| Instance2116.5 | 23.108 | 261918.768 |

**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* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance2116.1 | 24.666 | 15.470 | 972.912 | 4.847 | 41114.653 | 32768.000 | 8.612 | 0.000 | 7.270 | 0.000 | 232571.462 | 0.000 |
| Instance2116.2 | 25.609 | 15.552 | 986.519 | 4.711 | 40822.409 | 32768.000 | 3.318 | 0.000 | 7.070 | 0.000 | 232418.028 | 0.000 |
| Instance2116.3 | 20.060 | 11.978 | 1749.681 | 8.607 | 38837.245 | 32768.000 | 8.830 | 0.000 | 12.923 | 0.000 | 232549.281 | 0.000 |
| Instance2116.4 | 25.742 | 16.244 | 1002.634 | 4.870 | 40721.839 | 32768.000 | 3.143 | 0.000 | 7.305 | 0.000 | 232598.427 | 0.000 |
| Instance2116.5 | 25.514 | 15.355 | 990.278 | 4.767 | 40766.730 | 32768.000 | 7.983 | 0.000 | 7.151 | 0.000 | 232463.744 | 0.000 |

**Host System Performance**

| *Counter* | *Average* | *Minimum* | *Maximum* |
|---|---|---|---|
| % Processor Time | 3.030 | 0.000 | 11.065 |
| Available MBytes | 29385.525 | 29344.000 | 30544.000 |
| Free System Page Table Entries | 33555611.044 | 33555610.000 | 33555613.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 75525913.600 | 72654848.000 | 76587008.000 |
| Pool Paged Bytes | 131322137.600 | 131284992.000 | 131354624.000 |
| Database Page Fault Stalls/sec | 0.002 | 0.000 | 0.249 |

Test Log2/13/2011 9:11:36 PM -- Jetstress testing begins ...
2/13/2011 9:11:36 PM -- Prepare testing begins ...
2/13/2011 9:11:41 PM -- Attaching databases ...
2/13/2011 9:11:41 PM -- Prepare testing ends.
2/13/2011 9:11:41 PM -- Dispatching transactions begins ...
2/13/2011 9:11:41 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
2/13/2011 9:11:41 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
2/13/2011 9:11:52 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
2/13/2011 9:11:52 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
2/13/2011 9:11:56 PM -- Operation mix: Sessions 6, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
2/13/2011 9:11:56 PM -- Performance logging begins (interval: 15000 ms).
2/13/2011 9:11:56 PM -- Generating log files ...
2/13/2011 10:26:51 PM -- C:\logluns\log1 (104.0% generated), C:\logluns\log2 (100.2% generated), C:\logluns\log3 (101.2% generated), C:\logluns\log4 (101.0% generated) and C:\logluns\log5 (100.4% generated)
2/13/2011 10:26:51 PM -- Performance logging ends.
2/13/2011 10:26:51 PM -- JetInterop batch transaction stats: 22487, 21950, 21966, 22107 and 22224.
2/13/2011 10:26:51 PM -- Dispatching transactions ends.
2/13/2011 10:26:51 PM -- Shutting down databases ...
2/13/2011 10:26:56 PM -- Instance2116.1 (complete), Instance2116.2 (complete), Instance2116.3 (complete), Instance2116.4 (complete) and Instance2116.5 (complete)
2/13/2011 10:26:56 PM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Soft Recovery\Performance_2011_2_13_21_11_52.blg has 299 samples.
2/13/2011 10:26:56 PM -- Creating test report ...
2/13/2011 10:26:57 PM -- Instance2116.1 has 13.7 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.1 has 2.1 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.1 has 2.1 for I/O Log Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.2 has 13.0 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.2 has 1.0 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.2 has 1.0 for I/O Log Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.3 has 10.7 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.3 has 2.2 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.3 has 2.2 for I/O Log Reads Average Latency.

2/13/2011 10:26:57 PM -- Instance2116.4 has 12.9 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.4 has 1.0 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.4 has 1.0 for I/O Log Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.5 has 12.8 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.5 has 2.0 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.5 has 2.0 for I/O Log Reads Average Latency.
2/13/2011 10:26:57 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
2/13/2011 10:26:57 PM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
2/13/2011 10:26:57 PM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Soft Recovery\Performance_2011_2_13_21_11_52.xml has 298 samples queried.
2/13/2011 10:26:57 PM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Soft Recovery\Performance_2011_2_13_21_11_52.html is saved.
2/13/2011 10:39:03 PM -- Performance logging begins (interval: 4000 ms).
2/13/2011 10:39:03 PM -- Recovering databases ...
2/13/2011 10:49:46 PM -- Performance logging ends.
2/13/2011 10:49:46 PM -- Instance2116.1 (643.37633), Instance2116.2 (636.4811179), Instance2116.3 (355.3686241), Instance2116.4 (620.038689) and Instance2116.5 (632.5031109)
2/13/2011 10:49:46 PM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Soft Recovery\SoftRecovery_2011_2_13_22_38_58.blg has 160 samples.
2/13/2011 10:49:46 PM -- Creating test report ...

## Soft Recovery Test Performance Result: SUN158

**Test Summary**

| | |
|---|---|
| ***Overall Test Result*** | Pass |
| ***Machine Name*** | SUN158 |
| ***Test Description*** | |
| ***Test Start Time*** | 2/13/2011 9:11:36 PM |
| ***Test End Time*** | 2/13/2011 10:26:56 PM |
| ***Collection Start Time*** | 2/13/2011 9:12:11 PM |
| ***Collection End Time*** | 2/13/2011 10:26:44 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 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Soft Recovery\Performance_2011_2_13_21_11_52.blg |

### Database Sizing and Throughput

| | |
|---|---|
| *Achieved Transactional I/O per Second* | 878.468 |
| *Capacity Percentage* | 100% |
| *Throughput Percentage* | 100% |
| *Initial Database Size (bytes)* | 9311256117248 |
| *Final Database Size (bytes)* | 9312841564160 |
| *Database Files (Count)* | 5 |

### Jetstress System Parameters

| | |
|---|---|
| *Thread Count* | 6 (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

| | |
|---|---|
| *Instance2116.1* | Log Path: C:\logluns\log1<br>Database: C:\dbluns\db1\Jetstress001001.edb |
| *Instance2116.2* | Log Path: C:\logluns\log2<br>Database: C:\dbluns\db2\Jetstress002001.edb |
| *Instance2116.3* | Log Path: C:\logluns\log3<br>Database: C:\dbluns\db3\Jetstress003001.edb |
| *Instance2116.4* | Log Path: C:\logluns\log4<br>Database: C:\dbluns\db4\Jetstress004001.edb |
| *Instance2116.5* | Log Path: C:\logluns\log5<br>Database: C:\dbluns\db5\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* |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Instance2116.1 | 13.656 | 9.538 | 113.338 | 66.439 | 32768.000 | 34885.833 | 0.000 | 2.056 | 0.000 | 51.832 | 0.000 | 4886.730 |
| Instance2116.2 | 12.954 | 8.335 | 110.221 | 64.320 | 32768.475 | 34934.350 | 0.000 | 1.047 | 0.000 | 51.447 | 0.000 | 4750.665 |
| Instance2116.3 | 10.688 | 2.679 | 110.175 | 64.625 | 32768.061 | 34941.421 | 0.000 | 2.223 | 0.000 | 49.652 | 0.000 | 4956.638 |
| Instance2116.4 | 12.921 | 6.725 | 110.510 | 64.675 | 32768.000 | 34891.096 | 0.000 | 1.042 | 0.000 | 51.745 | 0.000 | 4758.984 |
| Instance2116.5 | 12.819 | 4.927 | 109.856 | 64.309 | 32768.000 | 34939.439 | 0.000 | 2.028 | 0.000 | 50.106 | 0.000 | 4880.781 |

**Host System Performance**

| *Counter* | *Average* | *Minimum* | *Maximum* |
|---|---|---|---|
| % Processor Time | 0.997 | 0.000 | 3.800 |
| Available MBytes | 29387.686 | 29364.000 | 30330.000 |
| Free System Page Table Entries | 33555643.043 | 33555642.000 | 33555645.000 |
| Transition Pages RePurposed/sec | 0.000 | 0.000 | 0.000 |
| Pool Nonpaged Bytes | 72242645.191 | 72216576.000 | 72335360.000 |
| Pool Paged Bytes | 131449914.221 | 131108864.000 | 131645440.000 |
| Database Page Fault Stalls/sec | 0.000 | 0.000 | 0.000 |

Test Log2/13/2011 9:11:36 PM -- Jetstress testing begins ...
2/13/2011 9:11:36 PM -- Prepare testing begins ...
2/13/2011 9:11:41 PM -- Attaching databases ...
2/13/2011 9:11:41 PM -- Prepare testing ends.
2/13/2011 9:11:41 PM -- Dispatching transactions begins ...
2/13/2011 9:11:41 PM -- Database cache settings: (minimum: 160.0 MB, maximum: 1.2 GB)
2/13/2011 9:11:41 PM -- Database flush thresholds: (start: 12.8 MB, stop: 25.6 MB)
2/13/2011 9:11:52 PM -- Database read latency thresholds: (average: 20 msec/read, maximum: 100 msec/read).
2/13/2011 9:11:52 PM -- Log write latency thresholds: (average: 10 msec/write, maximum: 100 msec/write).
2/13/2011 9:11:56 PM -- Operation mix: Sessions 6, Inserts 40%, Deletes 20%, Replaces 5%, Reads 35%, Lazy Commits 70%.
2/13/2011 9:11:56 PM -- Performance logging begins (interval: 15000 ms).
2/13/2011 9:11:56 PM -- Generating log files ...
2/13/2011 10:26:51 PM -- C:\logluns\log1 (104.0% generated), C:\logluns\log2 (100.2% generated), C:\logluns\log3 (101.2% generated), C:\logluns\log4 (101.0% generated) and C:\logluns\log5 (100.4% generated)
2/13/2011 10:26:51 PM -- Performance logging ends.
2/13/2011 10:26:51 PM -- JetInterop batch transaction stats: 22487, 21950, 21966, 22107 and 22224.
2/13/2011 10:26:51 PM -- Dispatching transactions ends.
2/13/2011 10:26:51 PM -- Shutting down databases ...
2/13/2011 10:26:56 PM -- Instance2116.1 (complete), Instance2116.2 (complete), Instance2116.3 (complete), Instance2116.4 (complete) and Instance2116.5 (complete)
2/13/2011 10:26:56 PM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Soft Recovery\Performance_2011_2_13_21_11_52.blg has 299 samples.
2/13/2011 10:26:56 PM -- Creating test report ...
2/13/2011 10:26:57 PM -- Instance2116.1 has 13.7 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.1 has 2.1 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.1 has 2.1 for I/O Log Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.2 has 13.0 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.2 has 1.0 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.2 has 1.0 for I/O Log Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.3 has 10.7 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.3 has 2.2 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.3 has 2.2 for I/O Log Reads Average Latency.

2/13/2011 10:26:57 PM -- Instance2116.4 has 12.9 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.4 has 1.0 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.4 has 1.0 for I/O Log Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.5 has 12.8 for I/O Database Reads Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.5 has 2.0 for I/O Log Writes Average Latency.
2/13/2011 10:26:57 PM -- Instance2116.5 has 2.0 for I/O Log Reads Average Latency.
2/13/2011 10:26:57 PM -- Test has 0 Maximum Database Page Fault Stalls/sec.
2/13/2011 10:26:57 PM -- Test has 0 Database Page Fault Stalls/sec samples higher than 0.
2/13/2011 10:26:57 PM -- C:\ESRP 3 AMS2100 Res R5 HDP 600GB SAS 3GB Mbox 11,200 Users\Sun158\Soft Recovery\Performance_2011_2_13_21_11_52.xml has 298 samples queried.

Hitachi Data Systems Corporation



**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