

**Oracle Flash Storage FS1-2
10000 Mailbox Microsoft Exchange Server
2013 Mail Box Resiliency Solution**

Tested with: ESRP – Storage Version 4.0

Table of Contents

1.	Overview.....	3
2.	Disclaimer.....	3
3.	Features.....	3
4.	Solution Description.....	9
5.	Targeted Customer Profile.....	11
6.	Tested Deployment.....	12
	6.1 Exchange Configuration.....	12
	6.2 Primary Storage Hardware.....	12
	6.3 Primary Storage Disk Configuration (Mailbox Store Disks).....	13
	6.4 Storage Disk Configuration (Transactional Log Disks).....	13
7.	Best Practices.....	13
	7.1 Core Storage.....	13
	7.2 Backup strategy.....	14
8.	Test Results Summary.....	14
	8.1 Reliability.....	14
9.	Storage Performance Results.....	15
	9.1 Aggregate Performance Metrics across All Servers.....	15
	9.2 Database Backup/Recovery Performance.....	16
	9.3 Database Read-Only Performance.....	16
	9.4 Transaction Log Recovery/Replay Performance.....	16
10.	Conclusion.....	16
11.	Test Result Report.....	17
	11.1 Performance Test Result Report.....	17
	11.2 Stress Test Result Report.....	20
	11.3 Database backup Test Result Report.....	23
	11.4 Soft Recovery Test Result Report.....	25

1. Overview

This document provides information on Oracle Flash Storage solution performance with 10000 Microsoft Exchange server 2013 users, based on the *Microsoft Exchange Solution Reviewed Program (ESRP) – Storage program**.

*The *ESRP – Storage* program was developed by Microsoft Corporation to provide a common storage testing framework for vendors to provide information on its storage solutions for Microsoft Exchange Server software. For more details on the *Microsoft ESRP – Storage* program, please click <http://technet.microsoft.com/en-us/exchange/ff182054.aspx>

2. 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 Oracle Flash storage 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 Oracle Flash storage, and Oracle Flash storage cannot guarantee the accuracy of any information presented after the date of publication.

3. Features

The Oracle FS1-2 Flash Storage System, Oracle's premier preferred SAN storage solution, delivers enterprise-grade storage capabilities that are optimized for flash media and co-engineered with Oracle software. Using the Quality of Service Plus (QoS Plus) feature, the Oracle FS1-2 Flash Storage System places data across flash and disk storage to maximize performance, efficiency, and cost based on usage profiles and business priorities. The Oracle FS1-2 flash storage system takes application-engineered storage to a new level by providing out-of-the-box tuned storage provisioning profiles for Oracle Database and key applications, including Microsoft SharePoint and Exchange. With the Oracle FS1-2 Flash Storage System, you can consolidate storage while achieving predictable performance for multiple diverse workloads in enterprise computing or multitenant environments.

ORACLE FS1-2 CONTROLLER SPECIFICATION

Cache and I/O ports

CPU	4 Intel E5-2620 CPUs (24 cores, 2.0 GHz)
Cache	384 GB RAM cache/32 GB NV-DIMM
Maximum cache hold-up time (after power failure)	Infinite hold-up time (using Oracle FS1-2 energy storage modules with super capacitors and Oracle-designed flash-backed DIMM modules)
Host ports	4 to 12 ports (2 to 6 HBAs) –16 Gbit FC, 10 Gbit NAS, 10 Gbit iSCSI Any HBA pair combination—customer-selected Four 10 Gbit NAS copper ports (motherboard)
Storage ports	Standard 12 ports (six 6 Gbit 4-lane SAS-2 HBAs)
Maximum disk configurations	SSD capacity: 912 TB maximum Disk drive capacity: 2,880 TB maximum Maximum of 30 drive enclosures

Drive Enclosure Types

Oracle Storage Drive Enclosure DE2-24P	2U rack size with twenty-four 2.5" drive bays
Oracle Storage Drive Enclosure DE2-24C	4U rack size with twenty-four 3.5" drive bays

Drive Enclosures

SSD Types/Usage	Drive Enclosure Layout	Total Capacity
Performance SSD (2.5" SAS-2)	13 x 400 GB drives	5.2 TB
	7 x 400 GB drives	2.8TB
Capacity SSD (2.5" SAS-2)	19 x 1.6 TB drives	30.4 TB
	13 x 1.6 TB drives	20.8 TB
	7 x 1.6 TB drives	11.2 TB
HDD Types/Usage	Drive Enclosure Layout	Total Capacity
Performance disk drive (2.5" SAS-2)	24 x 300 GB 10 K RPM drives	7.2 TB
	24 x 900 GB 10 K RPM drives	21.6 TB
Capacity SSD (3.5" SAS-2)	24 x 4 TB 7,200 RPM drives	96 TB

ORACLE MAXREP FOR SAN

Oracle MaxRep Asynchronous Replication	For remote long-haul (LAN/WAN) replication (transcontinental capable).
Oracle MaxRep Asynchronous Replication with Application Protection	For cases where restoring replicated volumes to a consistent point-in-time image is desired.
Oracle MaxRep Synchronous Replication	For local or metro replication.
Oracle MaxRep Synchronous Replication with Application Protection	For cases where restoring replicated volumes to a consistent point-in-time image is desired.



Flash-Optimized Performance

Leveraging a foundation of flash innovation leadership, the Oracle FS1-2 Flash Storage System is designed from the ground up to exploit the unique characteristics of flash storage to provide high IOPS and throughput without compromising expandability. It scales to 912 TB of flash, and up to 2.9 PB combined flash and disk, to meet the most demanding performance requirements. By supporting performance-optimized flash and capacity-optimized flash, you can create flash tiers to optimize read-intensive and mixed-use I/O.

QoS Plus

QoS Plus is a policy-based virtualization feature incorporating business priority I/O queue management fused with sub LUN automatic tiering into one simple management framework. Built on Oracle's patented storage quality-of-service technology, QoS Plus collects detailed information on your storage usage profile, evaluates data granules for movement to different storage tiers, then automatically migrates data to the most cost effective media (flash or disk) from a \$/IOP and \$/GB standpoint based on the usage profile and the importance of that data to the business. QoS Plus performs data collection, evaluation, and movement based on the most efficient data granularity in the storage industry—up to 1,600 times more granular than competitive systems, making it the most efficient auto-tiering system in the market.

Application Profiles

The Oracle FS1-2 Flash Storage System comes with predefined application profiles that provide tuned and tested out-of-the-box storage optimization for Oracle Database and key enterprise applications, including non-Oracle applications such as Microsoft Exchange. With one-click provisioning you can optimize flash performance and manage Oracle Applications with a minimum of administration. The Oracle FS1-2 Flash Storage System database storage profiles can disaggregate database components such as index files, database tables, archive logs, redo logs, control files, and temp files so provisioning automatically optimizes Oracle Database performance without requiring detailed knowledge of the database components. New application profiles can be added to Oracle FS1-2, existing ones can be modified, and all profiles can be exported to other Oracle FS1-2 systems to standardize Storage provisioning across global data centers.

Oracle FS1-2 Controllers

The Oracle FS1-2 controllers provide high performance and redundancy with dual active/active control units, fans, power, and nonvolatile mirrored cache—all requirements for Enterprise-grade Storage. One Oracle FS1-2 Flash Storage System can support as many as 30 drive enclosures, enabling storage capacity scaling with no need to purchase another system and without introducing additional management points, extra software licenses, or costly downtime. Oracle FS1-2 controllers support multiple I/O ports for different infrastructure requirements—16 Gbit Fibre Channel, 10 Gbit iSCSI, and 10 Gbit Ethernet, providing a flexible multiprotocol platform for nearly any deployment

Coengineered with Oracle Database and Oracle Applications

The Oracle FS1-2 Flash Storage System is coengineered with Oracle Database and Oracle Applications, leveraging unique features such as the Hybrid Columnar Compression feature of Oracle Database, and one-click best practices provisioning for key Oracle Applications and Oracle Database. Hybrid Columnar Compression compresses data by up to 50x and speeds queries up to 5x when compared to competitive storage systems. Hybrid Columnar Compression is supported only on Oracle storage systems such as the Oracle FS1-2 Flash Storage System.

Application Profiles

The Oracle FS1-2 Flash Storage System comes with predefined application profiles that provide tuned and tested out-of-the-box storage optimization for Oracle Database and key enterprise applications, including non-Oracle applications such as Microsoft Exchange. With one-click provisioning you can optimize flash performance and manage Oracle Applications with a minimum of administration. The Oracle FS1-2 Flash Storage System database storage profiles can disaggregate database components such as index files, database tables, archive logs, redo logs, control files, and temp files so provisioning automatically optimizes Oracle Database performance without requiring detailed knowledge of the database components. New application profiles can be added to Oracle FS1-2, existing ones can be modified, and all profiles can be exported to other Oracle FS1-2 systems to standardize storage provisioning across global data centers.

Storage Domains

Storage domain software enables multiple, virtual storage systems within a single Oracle FS1-2 Flash Storage System. Each storage domain is a “data container” that isolates data from other storage domains, providing independence in multitenant environments for private or public cloud deployments, regulatory compliance requirements, or chargeback models. With Storage domains, you can custom-tailor QoS settings for multiple unique environments, with all Storage domains still residing on a single physical Oracle FS1-2 flash storage system for reduced power, cooling, and management administration expense.

Enterprise Grade

The Oracle FS1-2 Flash Storage System is designed with enterprise-grade hardware and software features to provide maximum uptime, data availability, and application access. These features deliver ~one second controller failover capability, warm start technology for upgrades and error containment, and no single point of failure (SPOF). Additional enterprise-grade data integrity capabilities include pre-emptive copying of data away from potentially failing media, monitoring of SSD usage patterns and wear levels, T10 Protection Information (T10-PI) for application-to-storage data integrity verification, and complete copy services—all included with the Oracle FS1-2 system price. Oracle FS1-2 Flash Storage System ship fully tested and racked from the factory to accelerate deployments and support remote real-time diagnostic telemetry, and they are available with optional Oracle Business Critical Service for Systems to maximize system and data availability (see below for details).

Integrated with Oracle's Red Stack Management for Improved IT Staff Productivity

The Oracle MaxMan feature of the Oracle FS1-2 Flash Storage System enables the management of multiple Oracle FS1-2 flash storage systems and/or Oracle's Pillar Axiom systems from a single console. Oracle FS1-2 provides management plug-ins for Oracle Enterprise Manager and the Storage Connect feature of Oracle VM for singlepane-of-glass monitoring and management.

T10 Protection Information

The Oracle FS1-2 Flash Storage System uses T10-PI for end-to-end data integrity checking. T10-PI protects against silent data corruption from the application to the storage devices.

Solution Description

Modular Enterprise Storage

The Oracle FS1-2 Flash Storage System supports enterprise-grade deployments for your mission-critical applications and data with a modular approach that enables you to scale throughput, media types, and capacity independently. This architecture enhances the resiliency of the Oracle FS1-2 Flash Storage System by simplifying management, upgrades, reconfigurations, disaster recovery, and business continuance while minimizing the impact on your enterprise.

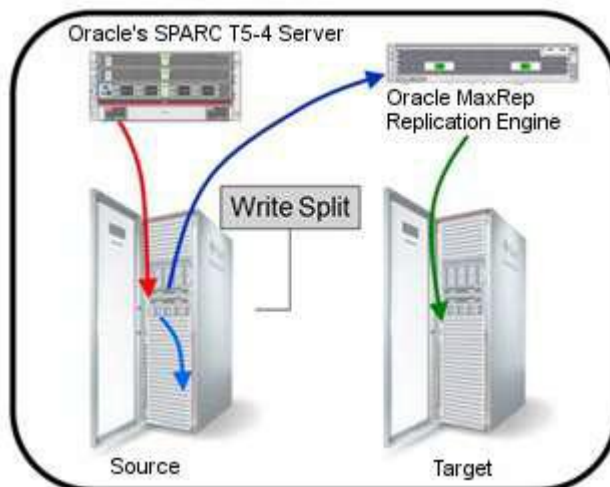
Oracle's Storage Drive Enclosures Oracle's

Storage drive enclosures can be configured with a range of SSD flash drives and disk drives to meet complete business needs. Drive enclosure media options include 400 GB performance SSDs, 1.6 TB capacity SSDs, 300 GB and 900 GB performance disk drives, and 4 TB capacity disk drives. A single Oracle FS1-2 flash storage system supports any combination of these drives and uses QoS Plus to optimize both \$/IOPS and \$/TB. By scaling out either SSD or disk drive enclosures, a single Oracle FS1-2 Flash Storage System can support up to 912 TB of flash or 2.88 PB of disk-based storage. Oracle FS Pilot Oracle FS Pilot is the primary management interface for Oracle FS Series. Oracle FS Pilot is fully redundant for high availability



Oracle MaxRep Replication Engine

For the Oracle FS1-2 Flash Storage System, the optional Oracle MaxRep Replication Engine supports both high-performance synchronous and asynchronous replication to local and remote locations, including many-to-one, one-to-many, and multihop replication. Multiple recovery point objectives (RPOs) and recovery time objectives (RTOs) are supported. Application-consistent recovery options restore applications to a consistent point. Oracle MaxRep Replication Engine supports both the Oracle FS1-2 Flash Storage System as well as Oracle's legacy Pillar Axiom 600 storage system to enable replication between the two storage systems.



Oracle FS Data Protection Manager Overview

The Oracle FS Data Protection Manager (DPM) runs in a physical or virtual environment and gives you the capability to schedule backups of your data on a regular basis. Data Protection Manager recognizes each Exchange database or storage group that is set up on a Oracle FS System as a consistency group. A consistency group contains all the data necessary to represent the Exchange application instance. DPM manages the process of making backups of these consistency groups on the Oracle FS System.

4. Solution Description

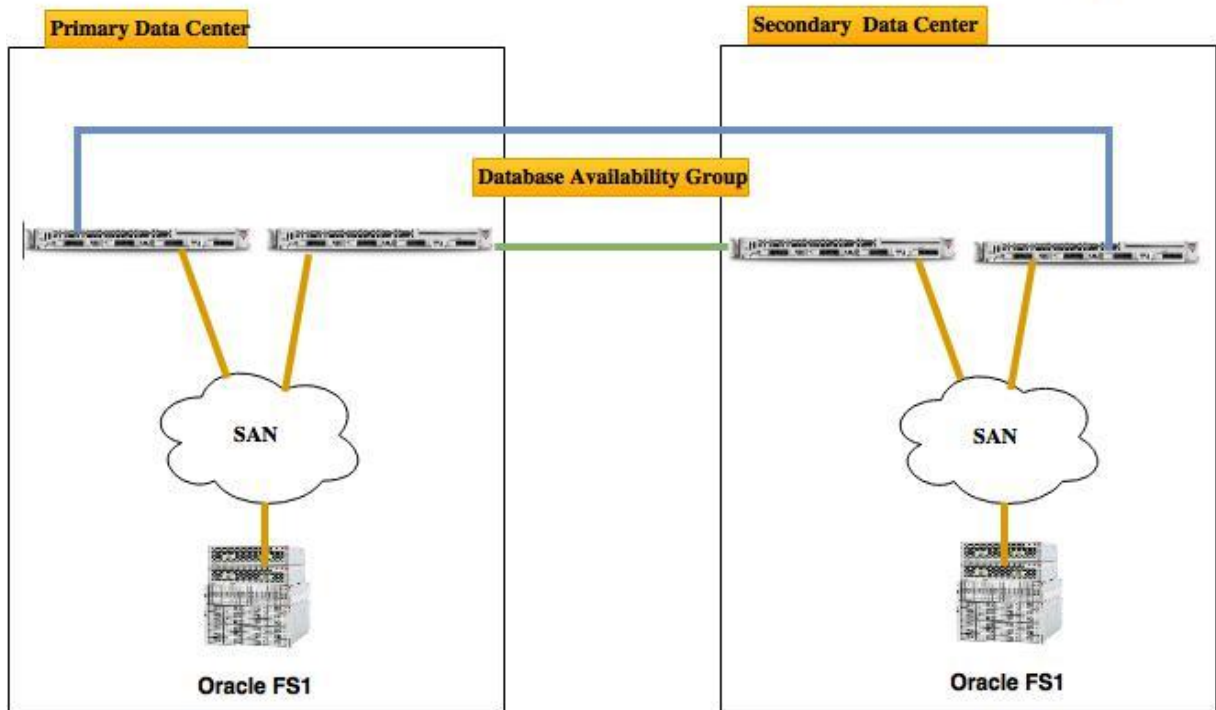
This solution is intended for medium to large enterprises that are planning to deploy Microsoft Exchange server 2013 on Microsoft and Oracle Flash Storage. It uses the Exchange Server 2013 Database Availability Group (DAG) feature to provide mailbox resiliency and high availability for exchange users. The solution design represents Exchange Server 2013 environment supporting 10,000 users in a mailbox resiliency configuration across two Oracle Flash storage arrays. This solution is designed to provide performance and flexibility for today's and tomorrow's Exchange users.

In this solution 10,000 users are distributed across two physical server hosts which host four Exchange Mailbox Servers in a DAG. This DAG has two RAID protected copies of every Exchange database a primary (active) copy and a secondary (passive) copy that are evenly split between physical hosts and storage pools on two Oracle Flash Storage arrays. Each Exchange database replicates to an alternate mailbox server that resides on different physical host through the use of the Exchange native DAG host based log shipping mechanism.

Each Mailbox Server physical machine is configured to support up to 5000 users (2,500 active and 2,500 passive) with a 1 GB mailbox capacity and 0.10 I/O operations per second (IOPS) per user (including an additional 20 percent I/O headroom). This user profile corresponds to about 75 messages per user per day.

This solution is designed to eliminate a single point and can handle the loss of an array, mailbox server, database volume, host bus adapter.

Exchange 2013 Mailbox Resiliency Solution for 10,000 Users on Oracle FS1 (2 DAG Copies)

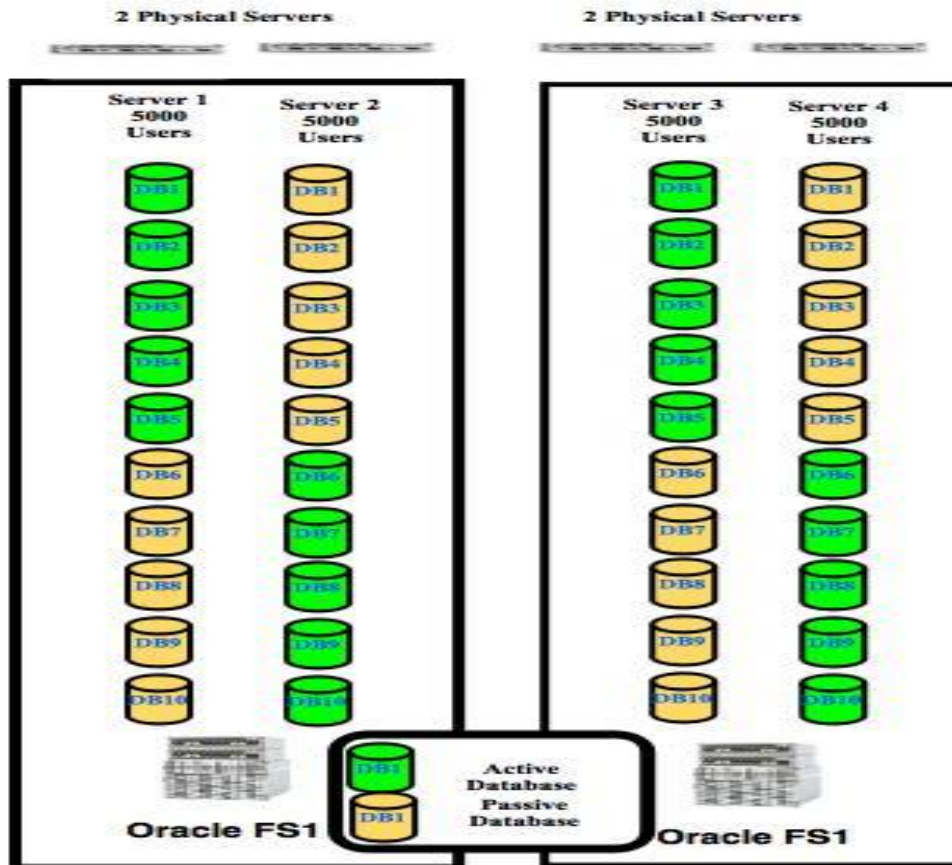


Storage Design is an important element for ensuring successful deployment of Microsoft Exchange Server 2013. A mailbox server building block represents the amount of disk, network and server resources required to support a specific number of Exchange Server 2013 users. The amount of resources is derived from specific user profile type, mailbox size and availability requirements.

Once an initial building block is designed, it can be easily reproduced to support the required number of users in your enterprise. By using this approach, Oracle FS1 customers can now create their own building blocks that are based on their company's specific Exchange environment requirements. This approach is helpful when future growth is expected because it makes Exchange environment expansion simple.

The mailbox server used in this solution supports up to 5,000 users per server with a 1 GB mailbox size and 0.10 IOPS per user. In this solution we have used RAID 6 for storage pool with 31*400 GB SSD disks thus supporting database I/O and mailbox capacity. For Log LUNs we have used 3 * 400GB SSD drive.

Storage Design for 10,000 Exchange 2013 users in a DAG with 2 copies on Oracle Flash Storage



5. Targeted Customer Profile

This solution design which supports 10,000 Exchange users has the following characteristics:

- 10,000 active mailboxes
- Unlimited number of hosts can be attached, via Fiber-Channel
- User IO profile (0.10 IOPS per user, 0.11 tested, giving 10% headroom).
- User mailbox size (1GB quota)
- Four Exchange servers per DAG, each supporting 5,000 users (2,500 active and 2,500 passive)
- Two Oracle Flash storage arrays (the solution was validated with two mailbox servers connected to a single storage array)
- Backup strategy VSS backup
- The tested RAID type was RAID 6 for database volumes and RAID 6 for log volumes, while a mix of RAID10, RAID5, and RAID6 can be blended, with fully automated tiered storage providing the most efficient and best performing storage where needed

6. Tested Deployment

The following tables summarize the testing environment:

6.1 Exchange Configuration

Number of Exchange mailboxes simulated	10000
Number of active mailboxes/server	5000
Number of databases/host	10
Number of copies/database	2
Number of Servers(2 Tested)	4
Number of DAGs	1
Number of mailboxes/database	500
Simulated profile: I/O's per second per mailbox (IOPS, include 20% headroom)	0.10 IOPS
Database LUN size	600GB
Log LUN size	100GB
Total database size for performance testing	10TB
% storage capacity used by Exchange database**	100% (10TB/10.0TB)

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

6.2 Primary Storage Hardware

Storage Connectivity (Fiber Channel, SAS, SATA, iSCSI)	Fiber Channel
Storage model and OS/firmware revision	http://www.windowsservercatalog.com/item.aspx?itemId=45be2cae-4c0a-432e-79e0-f987bf54d127&bCatID=1282
Storage cache	64 GB RAM cache/16 GB NV-DIMM
Number of storage controllers	2
Number of storage ports	4
Maximum bandwidth of storage connectivity to host	2GB
Switch type/model/firmware revision	Brocade 6510 Fiber Channel 4Gbit, v7.0.1a
HBA model and firmware	ISP2432-based 4Gb Fibre Channel HBA, qlnativefc version 1.0.12.0-1vmw
Number of HBA's/host	2
Host server type	SUN FIRE X4170 M3, 32 Processor,256 GB Memory
Total number of disks tested in solution	34
Maximum number of spindles can be hosted in the storage	

6.3 Primary Storage Disk Configuration (Mailbox Store Disks)

Disk type, speed and firmware revision	Performance SSD (2.5" SAS-2)
Raw capacity per disk (GB)	400
Number of physical disks in test	34
Total raw storage capacity (GB)	13600
Raid level	RAID 6
Total formatted capacity	10 TB

6.4 Storage Disk Configuration (Transactional Log Disks)

Disk type, speed and firmware revision	400 GB SAS 2.5
Raw capacity per disk (GB)	400
Total raw storage capacity (GB)	800
Number of disks per LUN	3
Raid level	RAID 6
Total formatted capacity	400GB

7. Best Practices

Exchange server is a disk-intensive application. Based on the testing run using the ESRP framework, we would recommend the following to improve the storage performance.

For Exchange 2013 best practices on storage design, please visit: <http://technet.microsoft.com> .

7.1 Core Storage

The Oracle FS1 family offers a range of configuration options for maximizing performance in any environment. For the Exchange environment we found that RAID 6 provides the best performance. The following is a list of the best practices for configuring an Exchange environment on the Oracle FS1 array. For a more detailed look at the Oracle FS1 array features Please visit our website at www.oracle.com .

- With Windows 2008R2 and above Diskpart is no longer needed to align sector boundaries.
- All database and log volumes should be formatted with NTFS using a cluster size of 64 KB.
- Best practice is to isolate I/O for Exchange from any other I/O on the array. Therefore, the Exchange database and logs should reside on isolated RAID Groups on the array.

- It is no longer necessary in Exchange 2013 to separate the database and log I/O, therefore a single volume can be used for both database and log capacity.
- Higher performance can be achieved using RAID 6 groups for the Exchange data. It is recommended that each RAID 6 group contains 1 active database (and log) and 1 passive database (and log).

7.2 Backup strategy

This solution does not cover the backup strategy. However Oracle FS1 does offer an Exchange backup and recovery feature that integrates the Oracle FS1 storage arrays Data Protection Manager (DPM) which functions with VSS (Volume Shadow Copy Services).

8. 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 the ESRP testing framework.

The ESRP – Storage Version 4.0 uses Jetstress 2013 to verify that the storage design meets the required disk performance and capacity requirements. Jetstress simulates an Exchange disk I/O load on a server to verify the performance and stability of a disk subsystem before it is placed in production. Jetstress is not designed to test server CPU and memory configuration. Microsoft Exchange Load Generator is a simulation tool to measure the impact of MAPI, OWA, ActiveSync, IMAP, POP and SMTP clients on Exchange servers.

8.1 Reliability

Reliability tests run for 24 hours. The goal is to verify the storage can handle high I/O load for a long period of time. Both log and database files will be analyzed for integrity after the stress test to ensure no database/log corruption.

Results:

1. No errors reported in the saved event log file
2. No errors reported in the checksum process for the database and log files

Please refer to the attachments to this document for the detailed results of the Database Checksum and Log Checksum Tests.

9. Storage Performance Results

The Primary Storage performance testing is designed to exercise the storage with maximum sustainable Exchange type of I/O for 2 hours. The test is to show how long it takes for the storage to respond to an I/O under load. The data below is the sum of all of the logical disk I/Os and average of all the logical disks I/O latency in the 2 hours test duration.

9.1 Aggregate Performance Metrics across All Servers

This is the sum of I/Os on 2 active server in the solution and the average latencies on 2 servers in the solution.

Database I/O	Server 1	Server 2
Target Transactional I/O	550	550
Achieved Transactional I/O	886	875
I/O Database reads/sec	66.85	66.14
I/O Database writes/sec	30.97	30.61
I/O database average read latency (ms)	1.04	1.05
I/O database average write latency (ms)	1.78	1.72
Transaction Log I/O		
Log Disks writes/sec	21.58	21.21
Log Disk Write Latency (ms)	0.67	0.67

Performance across Servers

Table below represents the sum of the I/O across servers and the average latency across the servers. This configuration is designed to achieve a target of 1,000 Exchange server 2013 IOPS for 10,000 users. The result shows I/O performance with 1761 Exchange Server 2013 user IOPs achieved. This provides an additional 761 IOPs. Note that these results already include 20 percent IOPs overhead recommended by Microsoft.

Database I/O	Value
Target Transactional I/O	1100
Achieved Transactional I/O	1761
I/O Database reads/sec	132.99
I/O Database writes/sec	61.58
I/O database average read latency (ms)	1.04
I/O database average write latency (ms)	1.75

Transaction Log I/O	
Log Disks writes/sec	42.79
Log Disk Write Latency (ms)	0.67

9.2 Database Backup/Recovery Performance

There are two test reports in this section. The first one measures the sequential read rate of the database files. The second measures the recovery/replay performance (playing transaction logs into the database).

9.3 Database Read-Only Performance

This test measures the maximum rate at which databases can be backed up via VSS. The following table shows the average rate for a single database file, a single mailbox server

MB read/sec per database	195.2
MB read/sec total per server	1952.08
Total read/sec for two servers	3881.12

9.4 Transaction Log Recovery/Replay Performance

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

Average time to play one log file (sec)	0.239
---	-------

10. Conclusion

This document is developed by Oracle America and reviewed by Microsoft Exchange Product team. The test results/data presented in this document is based on the tests introduced in the ESRP test framework. A customer should not quote the data directly for his/her pre-deployment verification. It is still necessary to go through the exercises to validate the storage design for a specific customer environment.

ESRP program is not designed to be a benchmarking program; tests are not designed to getting the maximum throughput for a giving solution. Rather, it is focused on producing

recommendations from vendors for Exchange application. So the data presented in this document should not be used for direct comparisons among the solutions

Oracle Flash Storage provides complete sizing information for Exchange Server deployments through its systems engineers. Contact your local Oracle Flash Storage sales office for additional information.

The testing of this solution demonstrates the following:

- The Oracle FS1 storage array has the performance and capacity requirements to easily handle a Microsoft Exchange environment with the ability to handle 10,000 mailboxes while maintaining Storage Resiliency through Exchange’s Data Resiliency Groups.
- The building block approach to designing an Exchange environment provides scalability options to your Exchange designs.

11. Test Result Report

11.1 Performance Test Result Report

Test Summary

Overall Test Result	Pass
Machine Name	CO-IOPSUN-09
Test Description	
Test Start Time	10/23/2015 11:39:20 PM
Test End Time	10/24/2015 2:10:01 AM
Collection Start Time	10/24/2015 1:10:12 AM
Collection End Time	10/24/2015 2:09:59 AM
Jetstress Version	15.00.0995.000
ESE Version	15.00.0516.026
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Results\Performance_23102015\Tuning_2015_10_23_23_39_55.blg C:\Program Files\Exchange Jetstress\Results\Performance_23102015\Performance_2015_10_24_09_56.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	886.731
Target Transactional I/O per Second	500
Initial Database Size (bytes)	5408775864320
Final Database Size (bytes)	5414714998784
Database Files (Count)	10

Jetstress System Parameters

Thread Count	2
Minimum Database Cache	320.0 MB
Maximum Database Cache	2560.0 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

Instance2296.1	Log path: F:\log1 Database: G:\master\Jetstress001001.edb
Instance2296.2	Log path: F:\log2 Database: H:\db1\Jetstress002001.edb
Instance2296.3	Log path: F:\log3 Database: M:\db2\Jetstress003001.edb
Instance2296.4	Log path: F:\log4 Database: J:\db3\Jetstress004001.edb
Instance2296.5	Log path: F:\log5 Database: K:\db4\Jetstress005001.edb
Instance2296.6	Log path: F:\log6 Database: L:\db5\Jetstress006001.edb
Instance2296.7	Log path: F:\log7 Database: N:\db6\Jetstress007001.edb
Instance2296.8	Log path: F:\log8 Database: O:\db7\Jetstress008001.edb
Instance2296.9	Log path: F:\log9 Database: P:\db8\Jetstress009001.edb
Instance2296.10	Log path: F:\log10 Database: Q:\db9\Jetstress010001.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
Instance2296.1	1.041	0.912	57.526	30.735	32969.083	35592.980	0.000	0.670	0.000	21.616	0.000	7989.422
Instance2296.2	1.045	1.046	57.749	31.091	32955.374	35590.148	0.000	0.671	0.000	21.717	0.000	7972.381
Instance2296.3	1.043	1.209	57.902	31.181	32951.604	35581.582	0.000	0.670	0.000	21.548	0.000	7959.483
Instance2296.4	1.045	1.435	58.059	31.373	32952.879	35529.119	0.000	0.671	0.000	21.741	0.000	7932.722
Instance2296.5	1.036	1.654	58.150	31.431	32984.399	35478.156	0.000	0.671	0.000	21.455	0.000	7958.004
Instance2296.6	1.045	1.896	57.599	30.634	33001.427	35592.955	0.000	0.671	0.000	21.458	0.000	7958.865
Instance2296.7	1.041	2.120	57.381	30.932	32985.319	35560.497	0.000	0.672	0.000	21.570	0.000	8006.910
Instance2296.8	1.039	2.358	57.192	30.566	32981.908	35622.790	0.000	0.672	0.000	21.660	0.000	8021.151

Instance2296.9	1.043	2.527	57.752	30.888	32956.241	35559.743	0.000	0.670	0.000	21.448	0.000	8010.154
Instance2296.10	1.046	2.731	57.627	30.963	32965.930	35619.935	0.000	0.671	0.000	21.630	0.000	7969.677

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance2296.1	9.167	261312.080
Instance2296.2	9.165	261401.444
Instance2296.3	9.164	261432.684
Instance2296.4	9.163	261445.116
Instance2296.5	9.169	261262.234
Instance2296.6	9.165	261388.250
Instance2296.7	9.165	261404.326
Instance2296.8	9.163	261455.154
Instance2296.9	9.166	261393.543
Instance2296.10	9.162	261488.992

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance2296.1	0.497	190544.314
Instance2296.2	0.497	189567.163
Instance2296.3	0.490	187612.863
Instance2296.4	0.495	190544.314
Instance2296.5	0.487	186635.712
Instance2296.6	0.490	188590.013
Instance2296.7	0.497	190544.314
Instance2296.8	0.500	190544.314
Instance2296.9	0.497	191521.464
Instance2296.10	0.497	190544.314

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
Instance2296.1	1.041	0.912	66.693	30.735	64355.194	35592.980	0.792	0.670	0.497	21.616	190544.314	7989.422
Instance2296.2	1.045	1.046	66.914	31.091	64244.756	35590.148	0.805	0.671	0.497	21.717	189567.163	7972.381
Instance2296.3	1.043	1.209	67.066	31.181	64171.996	35581.582	0.790	0.670	0.490	21.548	187612.863	7959.483
Instance2296.4	1.045	1.435	67.222	31.373	64099.280	35529.119	0.799	0.671	0.495	21.741	190544.314	7932.722
Instance2296.5	1.036	1.654	67.319	31.431	64075.734	35478.156	0.777	0.671	0.487	21.455	186635.712	7958.004
Instance2296.6	1.045	1.896	66.764	30.634	64353.831	35592.955	0.797	0.671	0.490	21.458	188590.013	7958.865
Instance2296.7	1.041	2.120	66.546	30.932	64442.891	35560.497	0.810	0.672	0.497	21.570	190544.314	8006.910
Instance2296.8	1.039	2.358	66.356	30.566	64532.332	35622.790	0.799	0.672	0.500	21.660	190544.314	8021.151
Instance2296.9	1.043	2.527	66.918	30.888	64245.341	35559.743	0.802	0.670	0.497	21.448	191521.464	8010.154
Instance2296.10	1.046	2.731	66.789	30.963	64313.802	35619.935	0.796	0.671	0.497	21.630	190544.314	7969.677

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.441	0.238	0.939
Available MBytes	255130.172	255115.000	255139.000
Free System Page Table Entries	16303739.167	16303445.000	16303984.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	242063659.916	241987584.000	242176000.000
Pool Paged Bytes	114786766.728	114778112.000	114814976.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

11.2 Stress Test Result Report

Test Summary

Overall Test Result	Pass
Machine Name	CO-IOPSUN-09
Test Description	
Test Start Time	10/22/2015 10:22:35 PM
Test End Time	10/23/2015 10:45:57 PM
Collection Start Time	10/22/2015 11:46:09 PM
Collection End Time	10/23/2015 10:45:55 PM
Jetstress Version	15.00.0995.000
ESE Version	15.00.0516.026
Operating System	Windows Server 2012 R2 Datacenter (6.2.9200.0)
Performance Log	C:\Program Files\Exchange Jetstress\Results\24 hours\Tuning_2015_10_22_22_23_30.blg C:\Program Files\Exchange Jetstress\Results\24 hours\Stress_2015_10_22_22_45_52.blg

Database Sizing and Throughput

Achieved Transactional I/O per Second	839.005
Target Transactional I/O per Second	500
Initial Database Size (bytes)	5377352138752
Final Database Size (bytes)	5408775864320
Database Files (Count)	10

Jetstress System Parameters

Thread Count	2
Minimum Database Cache	320.0 MB
Maximum Database Cache	2560.0 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

Instance2296.1	Log path: F:\log1 Database: G:\master\Jetstress001001.edb
Instance2296.2	Log path: F:\log2 Database: H:\db1\Jetstress002001.edb
Instance2296.3	Log path: F:\log3 Database: M:\db2\Jetstress003001.edb
Instance2296.4	Log path: F:\log4 Database: J:\db3\Jetstress004001.edb
Instance2296.5	Log path: F:\log5 Database: K:\db4\Jetstress005001.edb
Instance2296.6	Log path: F:\log6 Database: L:\db5\Jetstress006001.edb
Instance2296.7	Log path: F:\log7 Database: N:\db6\Jetstress007001.edb
Instance2296.8	Log path: F:\log8 Database: O:\db7\Jetstress008001.edb
Instance2296.9	Log path: F:\log9 Database: P:\db8\Jetstress009001.edb
Instance2296.10	Log path: F:\log10 Database: Q:\db9\Jetstress010001.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
Instance2296.1	1.123	0.906	54.638	29.217	33045.458	36121.639	0.000	0.682	0.000	20.638	0.000	8032.130
Instance2296.2	1.120	1.027	54.493	29.100	33036.329	36144.822	0.000	0.682	0.000	20.702	0.000	8017.561
Instance2296.3	1.114	1.203	54.683	29.287	33036.578	36104.278	0.000	0.682	0.000	20.714	0.000	8025.450
Instance2296.4	1.124	1.405	54.629	29.261	33042.482	36126.484	0.000	0.683	0.000	20.769	0.000	8017.884
Instance2296.5	1.114	1.634	54.663	29.271	33039.745	36107.130	0.000	0.682	0.000	20.725	0.000	8015.593
Instance2296.6	1.125	1.850	54.715	29.305	33032.480	36110.784	0.000	0.683	0.000	20.721	0.000	8001.948
Instance2296.7	1.119	2.083	54.612	29.217	33036.061	36139.902	0.000	0.682	0.000	20.688	0.000	8025.049
Instance2296.8	1.129	2.297	54.738	29.364	33038.075	36122.875	0.000	0.683	0.000	20.777	0.000	8009.787
Instance2296.9	1.115	2.512	54.632	29.263	33037.488	36149.432	0.000	0.682	0.000	20.782	0.000	8001.354

Instance2296.10	1.129	2.684	54.635	29.280	33043.540	36118.079	0.000	0.682	0.000	20.821	0.000	8002.927
-----------------	-------	-------	--------	--------	-----------	-----------	-------	-------	-------	--------	-------	----------

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance2296.1	9.164	261434.422
Instance2296.2	9.163	261455.091
Instance2296.3	9.163	261463.129
Instance2296.4	9.163	261447.919
Instance2296.5	9.162	261472.706
Instance2296.6	9.163	261462.251
Instance2296.7	9.163	261464.055
Instance2296.8	9.163	261467.791
Instance2296.9	9.163	261446.693
Instance2296.10	9.163	261458.091

Log Replication I/O Performance

MSExchange Database ==> Instances	I/O Log Reads/sec	I/O Log Reads Average Bytes
Instance2296.1	0.478	184782.144
Instance2296.2	0.479	184875.575
Instance2296.3	0.479	184921.623
Instance2296.4	0.479	185300.278
Instance2296.5	0.478	185039.963
Instance2296.6	0.477	184112.447
Instance2296.7	0.478	184677.196
Instance2296.8	0.479	185157.391
Instance2296.9	0.479	184994.605
Instance2296.10	0.480	185009.537

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
Instance2296.1	1.123	0.906	63.801	29.217	65848.168	36121.639	0.788	0.682	0.478	20.638	184782.144	8032.130
Instance2296.2	1.120	1.027	63.656	29.100	65915.997	36144.822	0.788	0.682	0.479	20.702	184875.575	8017.561
Instance2296.3	1.114	1.203	63.846	29.287	65818.818	36104.278	0.793	0.682	0.479	20.714	184921.623	8025.450
Instance2296.4	1.124	1.405	63.792	29.261	65850.494	36126.484	0.791	0.683	0.479	20.769	185300.278	8017.884
Instance2296.5	1.114	1.634	63.826	29.271	65832.062	36107.130	0.782	0.682	0.478	20.725	185039.963	8015.593
Instance2296.6	1.125	1.850	63.878	29.305	65798.622	36110.784	0.788	0.683	0.477	20.721	184112.447	8001.948
Instance2296.7	1.119	2.083	63.774	29.217	65855.224	36139.902	0.779	0.682	0.478	20.688	184677.196	8025.049
Instance2296.8	1.129	2.297	63.901	29.364	65792.020	36122.875	0.788	0.683	0.479	20.777	185157.391	8009.787
Instance2296.9	1.115	2.512	63.795	29.263	65844.830	36149.432	0.794	0.682	0.479	20.782	184994.605	8001.354
Instance2296.10	1.129	2.684	63.798	29.280	65848.824	36118.079	0.789	0.682	0.480	20.821	185009.537	8002.927

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	0.422	0.205	1.673
Available MBytes	255092.820	255024.000	255111.000
Free System Page Table Entries	16303499.914	16302654.000	16303859.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	241242914.509	241016832.000	241872896.000
Pool Paged Bytes	109701929.588	109527040.000	114798592.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

11.3 Database backup Test Result Report

Database Backup Statistics - All

Database Instance	Database Size (MBytes)	Elapsed Backup Time	MBytes Transferred/sec
Instance2296.1	516649.09	00:45:43	188.33
Instance2296.2	516681.09	00:46:17	185.99
Instance2296.3	516673.09	00:42:42	201.60
Instance2296.4	516665.09	00:46:15	186.18
Instance2296.5	516673.09	00:42:46	201.34
Instance2296.6	516657.09	00:45:25	189.55
Instance2296.7	516665.09	00:43:10	199.46
Instance2296.8	516681.09	00:45:38	188.70
Instance2296.9	516673.09	00:43:01	200.13
Instance2296.10	516697.09	00:45:51	187.77
Avg			192.90
Sum			1929.04

Jetstress System Parameters

Thread Count	80
Minimum Database Cache	320.0 MB
Maximum Database Cache	2560.0 MB
Insert Operations	40%
Delete Operations	20%
Replace Operations	5%
Read Operations	35%
Lazy Commits	70%

Database Configuration

Instance2296.1	Log path: F:\log1 Database: G:\master\Jetstress001001.edb
Instance2296.2	Log path: F:\log2 Database: H:\db1\Jetstress002001.edb
Instance2296.3	Log path: F:\log3 Database: M:\db2\Jetstress003001.edb

Instance2296.4	Log path: F:\log4 Database: J:\db3\Jetstress004001.edb
Instance2296.5	Log path: F:\log5 Database: K:\db4\Jetstress005001.edb
Instance2296.6	Log path: F:\log6 Database: L:\db5\Jetstress006001.edb
Instance2296.7	Log path: F:\log7 Database: N:\db6\Jetstress007001.edb
Instance2296.8	Log path: F:\log8 Database: O:\db7\Jetstress008001.edb
Instance2296.9	Log path: F:\log9 Database: P:\db8\Jetstress009001.edb
Instance2296.10	Log path: F:\log10 Database: Q:\db9\Jetstress010001.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
Instance2296.1	2.065	0.000	754.363	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2296.2	2.130	0.000	742.443	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2296.3	1.970	0.000	804.302	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2296.4	2.100	0.000	744.039	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2296.5	1.981	0.000	804.239	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2296.6	2.069	0.000	756.177	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2296.7	1.963	0.000	798.295	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2296.8	2.085	0.000	753.571	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2296.9	1.996	0.000	798.450	0.000	262144.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Instance2296.10	2.068	0.000	747.490	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	2.350	1.264	2.725
Available MBytes	257726.783	257704.000	257750.000
Free System Page Table Entries	16303985.250	16303605.000	16304285.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	246113146.435	246091776.000	246185984.000
Pool Paged Bytes	123211241.739	123183104.000	123297792.000
Database Page Fault Stalls/sec	0.000	0.000	0.000

11.4 Soft Recovery Test Result Report

Soft-Recovery Statistics - All

Database Instance	Log files replayed	Elapsed seconds
Instance2296.1	506	118.6254895
Instance2296.2	507	121.2445791
Instance2296.3	501	124.5970858
Instance2296.4	508	122.0236563
Instance2296.5	507	120.9789471
Instance2296.6	512	123.3164315
Instance2296.7	506	119.1556956
Instance2296.8	507	122.5399037
Instance2296.9	506	123.0639879
Instance2296.10	510	118.6254895
Avg	507	121.417
Sum	5070	1214.171266

Database Configuration

Instance2296.1	Log path: F:\log1 Database: G:\master\Jetstress001001.edb
Instance2296.2	Log path: F:\log2 Database: H:\db1\Jetstress002001.edb
Instance2296.3	Log path: F:\log3 Database: M:\db2\Jetstress003001.edb
Instance2296.4	Log path: F:\log4 Database: J:\db3\Jetstress004001.edb
Instance2296.5	Log path: F:\log5 Database: K:\db4\Jetstress005001.edb
Instance2296.6	Log path: F:\log6 Database: L:\db5\Jetstress006001.edb
Instance2296.7	Log path: F:\log7 Database: N:\db6\Jetstress007001.edb
Instance2296.8	Log path: F:\log8 Database: O:\db7\Jetstress008001.edb
Instance2296.9	Log path: F:\log9 Database: P:\db8\Jetstress009001.edb
Instance2296.10	Log path: F:\log10 Database: Q:\db9\Jetstress010001.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
Instance2296.1	2.367	0.872	3308.800	17.131	40484.183	32768.000	2.392	0.000	21.414	0.000	209716.222	0.000
Instance2296.2	2.387	0.872	3214.674	16.848	40469.434	32768.000	2.364	0.017	21.137	0.017	209786.819	141.241
Instance2296.3	2.317	0.851	3116.909	15.859	40600.600	32768.000	2.340	0.000	19.824	0.000	209715.633	0.000
Instance2296.4	2.302	0.874	3179.776	16.695	40559.653	32768.000	2.405	0.000	20.890	0.000	209718.516	0.000
Instance2296.5	2.334	0.855	3225.243	16.849	40385.181	32768.000	2.365	0.000	21.061	0.000	209703.171	0.000
Instance2296.6	2.317	0.873	3089.400	16.507	40480.745	32768.000	2.355	0.000	20.634	0.000	209712.323	0.000
Instance2296.7	2.330	0.858	3290.842	17.061	40417.699	32768.000	2.399	0.000	21.326	0.000	209724.914	0.000
Instance2296.8	2.340	0.877	3280.121	16.541	40497.285	32768.000	2.338	0.000	20.676	0.000	209713.436	0.000
Instance2296.9	2.326	0.864	3208.297	16.405	40325.825	32768.000	2.391	0.000	20.506	0.000	209715.200	0.000
Instance2296.10	2.351	0.866	3283.558	17.237	40542.494	32768.000	2.504	0.000	21.547	0.000	209715.200	0.000

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance2296.1	9.175	261017.893
Instance2296.2	9.157	261072.395
Instance2296.3	9.166	261579.034
Instance2296.4	9.149	262144.000
Instance2296.5	9.175	260997.020
Instance2296.6	9.191	260789.913
Instance2296.7	9.140	262144.000
Instance2296.8	9.149	261876.385
Instance2296.9	9.140	261876.385
Instance2296.10	9.148	262112.371

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
Instance2296.1	2.367	0.872	3317.975	17.131	41094.008	32768.000	2.392	0.000	21.414	0.000	209716.222	0.000
Instance2296.2	2.387	0.872	3223.831	16.848	41096.049	32768.000	2.364	0.017	21.137	0.017	209786.819	141.241
Instance2296.3	2.317	0.851	3126.074	15.859	41248.513	32768.000	2.340	0.000	19.824	0.000	209715.633	0.000
Instance2296.4	2.302	0.874	3188.924	16.695	41195.353	32768.000	2.405	0.000	20.890	0.000	209718.516	0.000
Instance2296.5	2.334	0.855	3234.418	16.849	41010.982	32768.000	2.365	0.000	21.061	0.000	209703.171	0.000
Instance2296.6	2.317	0.873	3098.592	16.507	41134.243	32768.000	2.355	0.000	20.634	0.000	209712.323	0.000
Instance2296.7	2.330	0.858	3299.982	17.061	41031.790	32768.000	2.399	0.000	21.326	0.000	209724.914	0.000
Instance2296.8	2.340	0.877	3289.270	16.541	41113.021	32768.000	2.338	0.000	20.676	0.000	209713.436	0.000
Instance2296.9	2.326	0.864	3217.437	16.405	40955.209	32768.000	2.391	0.000	20.506	0.000	209715.200	0.000
Instance2296.10	2.351	0.866	3292.707	17.237	41158.103	32768.000	2.504	0.000	21.547	0.000	209715.200	0.000

Host System Performance

Counter	Average	Minimum	Maximum
% Processor Time	8.484	7.195	9.389
Available MBytes	255110.300	255038.000	256438.000
Free System Page Table Entries	16303722.833	16303478.000	16303863.000
Transition Pages RePurposed/sec	0.000	0.000	0.000
Pool Nonpaged Bytes	243044352.000	242753536.000	243527680.000
Pool Paged Bytes	115437431.467	115404800.000	115494912.000
Database Page Fault Stalls/sec	0.000	0.000	0.000